

**Neutrino physics
beyond the minimal Standard Model
with IceCube DeepCore and PINGU**

Dissertation zur Erlangung des Grades
„Doktor der Naturwissenschaften“

am Fachbereich Physik, Mathematik und Informatik
der Johannes Gutenberg-Universität Mainz

vorgelegt von

Thomas Ehrhardt
geboren am 05.05.1990 in Frankfurt am Main

Mainz, den 02.07.2024

-
1. Gutachter: [in der elektronischen Version aus Datenschutzgründen entfernt]
 2. Gutachter: [in der elektronischen Version aus Datenschutzgründen entfernt]
- Datum der mündlichen Prüfung: 03.11.2025



Publiziert unter der Creative Commons-Lizenz Namensnennung (CC BY) 4.0 International.

Published under a Creative Commons Attribution (CC BY) 4.0 International License.

<https://creativecommons.org/licenses/by/4.0/>

Abstract

The observation of neutrino flavour transitions including the periodic “oscillation” phenomenon has led to increasingly deep insights into the nature of leptonic mass-flavour mixing and the neutrino mass spectrum. Ever since the earliest experiments, new neutrino interactions have also been scrutinised as an additional mechanism that could give rise to flavour transitions in the terrestrial or solar medium. This thesis presents two studies with very-large-volume Cherenkov detectors in deep glacial ice at the Geographic South Pole which investigate the flavour-, energy-, and zenith-dependent signatures of the aforementioned physics beyond the minimal Standard Model on the flux spectra of neutrinos produced in the Earth’s atmosphere.

The first study makes use of a sample of atmospheric neutrino events created from three years of operation with the existing IceCube DeepCore detector, with energies from 5.6 GeV to 100 GeV, in order to search for and constrain neutral-current non-standard interactions which modify the potential for coherent forward scattering encountered by atmospheric neutrinos as they propagate through the Earth. The measurement consists of separate tests of several hypotheses about the flavour structure of the interactions in the presence of standard oscillations and reports limits on effective coupling strengths for Earth matter. While oscillation measurements in general are not sensitive to the scale of flavour-diagonal interactions, which corresponds to an unobservable phase common to all of the propagating neutrino states, the vanishing four-momentum transfer in forward scattering renders them independent of the energy scale of the physics beyond the Standard Model.

The second analysis solely draws from the Monte Carlo simulation of the detector response of the proposed IceCube low-energy upgrade PINGU, whose dense instrumentation is expected to lead to a lowered energy threshold of ~ 1 GeV and improved reconstruction capabilities for GeV-scale atmospheric neutrinos. Working within a frequentist statistical framework, different approaches to projecting PINGU’s sensitivity to the still unresolved ordering of the neutrino mass eigenstates are discussed, assuming standard neutrino interactions with Earth matter. In particular, the sensitivity dependence on the values of the atmospheric neutrino oscillation parameters, the leptonic mixing angle θ_{23} and the mass-squared difference Δm_{31}^2 , is investigated in detail.

Contributions

Undoubtedly, the combined efforts of the IceCube collaboration as a whole have laid the foundations of this work and enabled my research with the DeepCore experiment in Part III and with PINGU in Part IV in the first place. Below is a synopsis of the specific contributions made by me, beginning chronologically with the content of Part IV.

- Sensitivity studies on the neutrino mass ordering with PINGU:
 - porting of the Fisher information formalism from the PaPA software framework for preceding PINGU sensitivity projections to the generalised PISA codebase used to obtain the results presented in the revised Letter of Intent (LoI) [1]
 - identification of limitations of the Fisher information matrix and extension for application to semianalytic minimisation
 - implementation of the Asimov $\overline{\Delta\chi^2}$ approach in PISA and study and resolution of discrepancies with respect to predictions made by the Fisher information matrix and by fits to large ensembles of pseudoexperiments
 - analysis of the performance of various detector deployment candidates in preparation for the LoI, thereby informing the selection of the new baseline geometry with 26 strings (publication of sensitivities to the neutrino mass ordering and to atmospheric oscillation parameters in Refs. [1, 2, 3])
 - significant contributions to the verification of the various smoothing methods of the PISA staged approach to the PINGU event template generation (publication in Ref. [4])
- Search for non-standard neutrino interactions with DeepCore:
 - addition of a more general description of the Earth matter potential to the code which calculated atmospheric neutrino oscillation probabilities on a GPU for the PINGU sensitivity studies
 - development of the full analysis approach with the aim of extending the physics reach with respect to any similar prior measurements made by individual neutrino oscillation experiments:
 - * selection of a suitable pre-existing DeepCore event sample

-
- * selection of the hypotheses and parameters to be tested
 - * design of customised fitting techniques tailored to the encountered test-statistic landscapes
 - * nuisance-parameter and sensitivity studies
 - * conducting and interpreting the measurement (summary of results published in Ref. [5])

Contents

1	Introduction	1
I	The physics of Standard Model neutrinos and beyond	7
2	Massless neutrinos in the minimal Standard Model	9
2.1	Fundamental interactions from symmetry	9
2.2	Development of Standard Model building blocks	11
2.3	Free Dirac fermions	14
2.4	Elements of electroweak theory	17
2.4.1	Gauge fields and fermion fields	17
2.4.2	Interactions	19
2.4.3	Mass generation	22
2.4.4	Effective low-energy theory	24
2.5	Neutrino scattering	25
2.5.1	General concepts	26
2.5.2	Neutrino-electron scattering	27
2.5.3	Neutrino-nucleon scattering	29
2.5.3.1	Elastic and quasielastic interactions	33
2.5.3.2	Resonance production	35
2.5.3.3	Deep inelastic scattering	36
2.5.4	Nuclear targets	38
2.5.4.1	Coherence	39
3	Neutrino mass and mixing	41
3.1	Massive neutrinos	41
3.1.1	Right-chiral or sterile?	43
3.1.2	Majorana neutrinos	44
3.1.2.1	Active neutrinos only	45
3.1.2.2	Coexisting active and sterile neutrinos	46
3.1.3	PMNS mixing matrix	48
3.1.4	Weinberg operator	49
3.2	Neutrino propagation in vacuum	50
3.2.1	Foundations	51

3.2.2	Standard vacuum oscillation formula	54
3.2.3	Discrete symmetries	55
3.2.4	Physical parameter space and neutrino mass ordering	57
3.3	Neutrino propagation with standard matter effects	60
3.3.1	Effective matter potential	61
3.3.2	Standard oscillation formula in uniform matter	63
3.3.3	Discrete symmetries	64
3.3.4	Flavour transition characteristics in different types of matter .	65
3.3.4.1	Constant density: resonance enhancement	65
3.3.4.2	Slow density changes: MSW effect	67
3.3.4.3	Periodic density changes: parametric enhancement .	70
4	Non-standard interactions of propagating neutrinos	73
4.1	Model-independent effective formalism	73
4.2	Neutral-current NSI	75
4.3	Parameterisations	76
4.4	Generalised mass-ordering degeneracy	78
4.5	New-physics origin	79
5	Global knowledge of the PMNS paradigm and NSI	83
5.1	Neutrino sources	84
5.1.1	Solar neutrinos	84
5.1.2	Atmospheric neutrinos	86
5.1.2.1	Flux characteristics	86
5.1.2.2	Standard oscillations	90
5.1.3	Accelerator neutrinos	99
5.1.4	Reactor neutrinos	101
5.2	Status of global fits	102
5.2.1	Bounds on PMNS parameters with standard matter potential	103
5.2.2	Simultaneous bounds on PMNS and NSI parameters	105
II	Atmospheric neutrinos in IceCube: present and future	113
6	Dual role of glacial ice in IceCube	115
6.1	Ice as target medium	115

6.2	Ice as detection medium	116
6.2.1	Cherenkov emission	116
6.2.2	Energy loss processes	117
6.2.3	Optical ice properties	120
7	The present: IceCube and DeepCore	127
7.1	Detector design	127
7.2	Data acquisition and processing at the South Pole	130
7.2.1	Digital optical module	130
7.2.2	Local-coincidence logic	131
7.2.3	Triggering	132
7.2.4	Pulse extraction from hit records	133
7.2.5	Removal of noise pulses	133
7.2.6	Filtering	135
7.3	Detector calibration	136
8	The future: PINGU and the IceCube Upgrade	137
8.1	The proposed PINGU experiment	137
8.2	The imminent IceCube Upgrade	139
III	Search for non-standard neutrino interactions with DeepCore	143
9	Atmospheric neutrino oscillations in the presence of NSI	145
9.1	Rescaled standard matter potential	146
9.2	μ - τ non-universality and flavour violation	147
9.3	e - μ or e - τ flavour violation	150
9.4	Arbitrary NSI flavour structure	151
10	Event simulation and reconstruction	153
10.1	Detector response modelling: from neutrino interactions to filtered events	154
10.2	Undoing the detector response: event reconstruction	162
11	Event sample	167
11.1	Context	167

11.2	Selection criteria	168
11.2.1	Data-taking runs	168
11.2.2	Events	168
12	Simulation predictions	177
12.1	Event-by-event reweighting	177
12.2	Atmospheric μ^\pm background	179
12.3	Event-sample characterisation	180
12.3.1	Event template	181
12.3.2	Distributions of true neutrino properties	184
12.3.3	Reconstruction resolutions	185
12.4	NSI signatures	188
12.5	Systematic uncertainties	195
13	Analysis method	203
13.1	Frequentist inference	204
13.2	A weighted-least-squares statistic	204
13.2.1	Relation to the χ^2 distribution	206
13.3	Construction of confidence regions	207
13.3.1	Choice of test statistic and Wilks' theorem	207
14	Monte Carlo studies	209
14.1	Characterisation of test-statistic landscapes	210
14.2	Statistics-only sensitivity	212
14.3	Treatment and impact of nuisance parameters	217
14.3.1	Profiling techniques	218
14.3.2	Systematic uncertainty budgets	220
14.3.3	Systematic fit biases	222
14.3.4	Sensitivity with additional nuisance parameters	223
14.4	Ensemble studies	224
14.4.1	Distribution of the weighted-least-squares statistic	227
14.4.2	NSI sensitivity	228
14.4.3	Coverage of confidence regions	232

15 Measurement results	237
15.1 Observed event distribution and fit outcomes	237
15.1.1 One-by-one fits	240
15.1.2 Generalised matter potential	245
15.2 Summary and experiment comparison	247
15.2.1 Assessment of subsequent NSI studies	249
15.3 Rejecting SI: from Wilks to Feldman–Cousins	252
IV Monte Carlo studies on the neutrino mass ordering with PINGU	
	257
16 Event sample	259
16.1 Simulation and reconstruction	259
16.2 Selection criteria	260
16.3 Practicability	261
17 Simulation predictions	263
17.1 Staged approach	264
17.2 Event-sample characterisation	270
17.3 Atmospheric neutrino oscillation signatures	271
17.4 Systematic uncertainties	278
18 Statistical techniques	281
18.1 Frequentist hypothesis test: NO or IO?	281
18.1.1 Sensitivity definition	282
18.1.2 Challenges of testing discrete composite hypotheses	283
18.1.3 Sensitivity from Asimov templates	285
18.1.3.1 Fisher information	288
18.1.3.2 Semianalytic minimisation: pull approach	294
18.1.4 Sensitivity from ensemble studies	299
18.2 Statistical uncertainties only	300
18.2.1 Sampling extreme test-statistic outcomes	301
18.2.2 Sensitivity evolution with exposure time	304

19 Sensitivity projections under the Gaussian approximation	307
19.1 Sensitivity impact by uncertainty category	307
19.2 One-by-one sensitivity potential	309
19.3 Event-classification synergies	310
19.4 Coverage in dependence of truth assumptions	314
V Conclusion	323
20 Summary and outlook	325
Appendices	330
A Analytic calculation of oscillation probabilities	331
B Analytic breakdown of atmospheric (anti)neutrino event rates	334
C DeepCore NSI search: supplementary material	337
D PINGU NMO studies: supplementary material	345
D.1 Parameter pulls from Fisher information	346
Bibliography	348
List of Figures	380
List of Tables	386

1

Introduction

Nearly 40 years ago, the IMB (Irvine–Michigan–Brookhaven) experiment revealed a surprising result which would come to be known as the “atmospheric neutrino anomaly” [6]. In fact designed to search for proton decay—a key prediction made by “Grand Unified Theories” such as the SU(5) model [7] whose gauge symmetry is broken to that of the SM below some extremely high unification energy scale—the experiment consisted of a nearly cubical tank filled with around 8 kt of ultrapure water located 610 m underground in a salt mine near Lake Erie, US. 2 048 photomultiplier tubes that observed Cherenkov photons on the tank’s inner surface had sufficient directional reconstruction capabilities for distinguishing downgoing from upgoing events. Based on 417 days of data taking, 401 events with reconstructed interaction vertices within the detector’s fiducial volume were observed and assumed to originate from atmospheric neutrinos [8]. Compared to simulation, the fraction of events showing a delayed signal due to μ^\pm decay¹, indicative of a charged-current interaction of an atmospheric $\bar{\nu}_\mu$, exhibited a deficit of greater than three standard deviations [9, 10, 11]. The IMB collaboration cautiously provided several possible interpretations thereof [10]:

This discrepancy could be a statistical fluctuation or a systematic error due to (i) an incorrect assumption as to the ratio of muon neutrinos to electron neutrinos in the atmospheric fluxes, (ii) an incorrect estimate of the efficiency for our observing a muon decay, or (iii) some other as-yet-unaccounted-for physics.

A clear deficit soon after also showed up in μ^\pm -like events in the Kamiokande experiment [12], simultaneously giving credence to the third interpretation above. When

¹Either $\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$ or $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$, with an observed average time delay of the decay signal of $2.02 \mu\text{s}$ [8].

Super-Kamiokande reported the exceedingly high-significance evidence of the disappearance of upgoing atmospheric $\bar{\nu}_\mu$'s, but not $\bar{\nu}_e$'s, in 1998 [13], a breakthrough was finally achieved. This result represents a major cornerstone of our modern understanding of the anomaly—as a manifestation of $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ oscillations within the six-parameter Pontecorvo–Maki–Nakagawa–Sakata (PMNS) paradigm of three active massive (mixed) neutrinos [14], comprising the three leptonic mixing angles θ_{12} , θ_{13} , θ_{23} , the charge-parity-violating phase δ_{CP} , and the two differences of squared neutrino masses Δm_{21}^2 and Δm_{31}^2 . These are not part of the (minimal) Standard Model of particle physics, which requires neutrinos to be strictly massless [15].

The discovery of atmospheric neutrino oscillations followed a century of experimental efforts motivated by quantitative measurements in 1900 [16, 17] that had shown a spontaneous discharge of electroscopes in air despite intact insulation [18], suggesting the existence of highly penetrating ionising radiation. In 1911, systematic measurements of electroscope discharge rates off-shore, at sea level, and underwater led Italian meteorologist Pacini to conclude at high significance “that a sizable cause of ionisation exists in the atmosphere, originating from penetrating radiation, independent of the direct action of radioactive substances in the crust” [19]. The same year, Austrian physicist Hess embarked on his first two balloon flights [20], followed by a series of seven thoroughly planned ones during which he found strongly increased radiation intensities at altitudes of up to 5 350 m, thus asserting the existence of a cosmic radiation component [21]. With the demonstration of a latitude dependence of the cosmic-ray flux in 1933 [22] and of an “east-west effect” [23] it became evident that the primaries were predominantly positively charged particles. While the existence of atmospheric neutrinos was unambiguously inferred from observing the decays of cosmic-ray muons in 1940 [24], they were only directly detected for the first time in 1965, by deep-underground experiments performed in gold mines in India [25] and South Africa [26]. Relying on the exponential suppression of the atmospheric μ^\pm flux with depth, as well as on its primarily vertical direction of incidence, both experiments were optimised to detect near-horizontally travelling μ^\pm 's from charged-current interactions of atmospheric $\bar{\nu}_\mu$'s interacting in the surrounding rock. To this end, each detector was formed by two parallel walls of scintillators in the vertical plane whose light outputs were measured by photomultiplier tubes.

Today, atmospheric neutrino interactions are routinely recorded in large quantities, contributing to increasingly precise constraints on the parameters of the PMNS framework. Nevertheless, even after combining all available measurements from atmospheric with those from solar, reactor, or accelerator neutrino experiments, open

questions of fundamental importance remain [15], such as which neutrino is the heaviest (sign of Δm_{31}^2), whether the mixing angle responsible for the pronounced disappearance of atmospheric $\bar{\nu}_\mu$'s is maximal ($\theta_{23} = 45^\circ$), or whether leptons are subject to charge-parity violation ($\delta_{\text{CP}} \neq 0, \pi$). Furthermore, the possibility of neutrinos undergoing new interactions with regular matter during propagation [27] as another, subdominant, mechanism giving rise to flavour transformations has not been ruled out, which in turn impedes the determination of the PMNS parameters [28, 29, 30]. Obtaining accurate experimental answers to these questions is paramount in distinguishing between various theoretical models extending the Standard Model to incorporate neutrino masses [31, 32, 33].

The IceCube Neutrino Observatory is a highly versatile neutrino experiment embedded deep in the Antarctic ice sheet at the Geographic South Pole, where it encompasses one cubic kilometre of extremely clear glacial ice. Since the beginning of its data taking in 2005 [34], IceCube has made important contributions to the emerging modern understanding of neutrinos across about ten decades in neutrino energy, including the first detection of high-energy cosmic neutrinos [35], the identification of active galaxies as their possible sources [36, 37], or the measurement of the total neutrino-nucleon cross section up to centre-of-mass energies of ~ 1 TeV, corresponding to neutrino energies as high as ~ 10 PeV [38]. Near the low-energy threshold, the DeepCore infill array [39, 40] has extended IceCube's neutrino physics reach much beyond what was originally conceived [41], resulting in the highest-precision measurement using atmospheric neutrinos of θ_{23} and Δm_{31}^2 to date [42]. Still, at both the low- and high-energy detection limits promising efforts are being undertaken to expand IceCube's capabilities even further [43].

Drawing on atmospheric neutrino events at the GeV energy scale as observed in DeepCore or simulated with the proposed IceCube extension PINGU [1, 2], this thesis investigates the issues of non-standard neutrino interactions and the ordering of neutrino masses raised above. It is organised into five main parts. The first lays the theoretical foundations and consists of four chapters, of which Chapter 2 presents the embedding of massless neutrinos into the minimal Standard Model of particle physics, Chapter 3 investigates neutrino mass mechanisms and the associated effects of leptonic mixing and neutrino flavour transformations in vacuum and matter, Chapter 4 explores a well-motivated model-independent effective formalism which describes non-standard interactions of propagating neutrinos, and Chapter 5 summarises the production and propagation processes of neutrinos from a variety of natural and artificial sources and gives an overview of current experimental con-

straints on neutrino theory parameters relevant to this thesis.

The second part introduces the experimental setting at the South Pole with a focus on the detection of atmospheric neutrinos, beginning with a discussion of the interaction and optical properties of the glacial ice sheet in Chapter 6, followed by an overview of the working principles behind the existing IceCube detector and its low-energy extension DeepCore in Chapter 7 and a brief outlook to the planned future extension PINGU and the imminent IceCube Upgrade aiming at enhancing the detection capabilities for GeV-scale neutrinos in Chapter 8.

A measurement of non-standard neutrino interactions using tens of thousands of atmospheric neutrino events collected by DeepCore over a span of three years comprises the third part of this thesis. Since the nature of these effective interactions diverges between theoretical models, which encompass flavour-changing or flavour-non-universal couplings to the new interaction mediators, a number of distinct hypotheses are separately studied and constrained. Chapter 9 reveals the impacts of each hypothesis on the flavour transition probabilities of atmospheric neutrinos. Chapter 10 outlines the step-by-step process of transforming a simulated neutrino interaction into a detailed and accurate DeepCore detector response, which in turn is converted into a set of observable neutrino properties through a maximum-likelihood reconstruction process which is identically applied to the real detector response. The set of event selection criteria, which may only depend on observables and not true neutrino properties, is slightly adapted from preceding atmospheric neutrino studies with DeepCore and recapitulated in Chapter 11. Chapter 12 scrutinises collective properties of the resultant event sample as well as the characteristics of the expected event histograms from which statistical inferences will be made. In the case of atmospheric neutrino events, an event-by-event reweighting approach is adopted in order to efficiently modify any underlying physics assumptions, either about non-standard neutrino interactions or about nuisance parameters which represent various sources of systematic uncertainty. A histogram of background events is added in the form of an adequately rescaled sideband of the observed event sample. Chapter 13 introduces the statistical analysis methodology with which frequentist confidence regions are constructed in the parameters of interest. Here, a weighted-least-squares statistic is used to create a test statistic motivated by Wilks' theorem. Chapter 14 draws upon the thus established analysis framework to quantify the expected DeepCore sensitivity to the parameters of the various hypotheses under study, based on the "typical" event histogram for any given assumption about the true values of model parameters. This is done first in the absence of nuisance para-

meters to understand how statistical uncertainties limit the measurement precision. Complementary investigations of the simulated measurement outcomes of large ensembles of pseudoexperiments are made in order to deduce the expected distribution of the test statistic and validate the frequentist coverage probability of inferred confidence regions. The results of all measurements are presented in Chapter 15 and contextualised within the broader experimental landscape related to non-standard neutrino interactions.

The fourth part of this thesis employs a sample of simulated atmospheric neutrino events which have undergone a rudimentary selection procedure with the proposed PINGU detector in order to study the latter’s capabilities of distinguishing between the two allowed orderings of neutrino masses. Analysing and optimising the sensitivity of an experiment that is yet to be built comes with its own unique set of challenges. Crucially, one needs an analysis framework that allows for the quick adaptation of design assumptions, taking into account changing knowledge about the existing IceCube observatory, progress in research and development, as well as evolving budgetary constraints. Furthermore, in principle all of these considerations, together with their associated uncertainties, should be projected into the future, so that the external constraints applying to PINGU at the time of its operation are accurately captured. However, the expected experimental outcome for any detector configuration can only be found after a detailed simulation. Projecting the statistical distribution of PINGU’s sensitivity to the neutrino mass ordering given all conceivable designs is a task that is therefore computationally far out of reach. As a result, efforts have concentrated on finding sensitivity estimates for conservative experimental conditions, which are likely to be achieved with the technologies and data processing methods that have proven reliable for DeepCore. Still, this has required the simulation and analysis of a large number of different PINGU layouts. Exacerbating the above computational challenge, even for a single layout the number of simulated events required for a “brute-force” sensitivity calculation with a reasonable precision would be enormous [4]. An early example of PINGU sensitivity projections for different detector designs, event selection efficiencies, and statistical analysis methods can be found in Ref. [44].

During the PINGU design optimisation process, a statistical analysis software framework was developed in Ref. [45] for the fast computation of the sensitivity to the mass ordering, even when drawing on low event statistics. Given a set of reconstructed events generated for a particular PINGU design, parameterisations are created of the intermediate “stages” (e.g., detection, reconstruction, event classifica-

tion) into which the detector response can be divided. Then, assuming a multivariate normal distribution of the model-parameter estimates [45, 46] of the staged approach has been the collaboration's baseline method of determining PINGU's sensitivity to the neutrino mass ordering. Enabling to relax the latter assumption, a methodological study of the frequentist inference of PINGU's sensitivity using an optimised detector design that accounts for more up-to-date experimental assumptions [1, 2] is performed in this work. Chapter 16 outlines the construction of the underlying event sample. Chapter 17 reviews our methods of generating the expected event distribution, characterises typical imprints of the physics effects under study, and introduces the sources of systematic uncertainty included as nuisance parameters in the analysis. Chapter 18 sketches the challenges associated with a discrete hypothesis test as it applies to the mass ordering, derives and scrutinises computationally efficient sensitivity proxies, and compares these to the predictions from large ensembles of pseudoexperiments in the absence of systematic uncertainties. Finally, Chapter 19 reports a detailed study of the impact of nuisance parameters on the hypothesis test and of the sensitivity dependence on their true values, simultaneously pointing out regions of parameter space where a previously adopted pragmatic analysis approach succeeds and where it is overly optimistic.

The fifth part concludes the thesis with a summary and outlook in Chapter 20.

Part I

The physics of Standard Model neutrinos and beyond

2

Massless neutrinos in the minimal Standard Model

The Standard Model of particle physics (SM) as we know it today is the result of the collective endeavour of both theoretical and experimental physics toward finding a unified theory of the interactions of fundamental particles. As a relativistic quantum field theory¹, it succeeds in the joint description of electromagnetism, the weak interaction, and the strong interaction (but not the gravitational interaction) by combining the principles of special relativity, quantum mechanics, and classical field theory. This chapter provides an introduction into the Standard Model’s theoretical underpinning, in particular into the electroweak sector, in which neutrino interactions are embedded.

2.1 Fundamental interactions from symmetry

The “language” of the SM is group theory [48, 49]. A group can be considered a mathematical prescription for modifying a system without changing some property of the system, i.e., applying a symmetry transformation. According to Noether’s theorem [50], the invariance of the Lagrangian² of some system under such a group of transformations results in a conserved current, which in turn results in an observable conserved charge. In fact, all the dynamics of interactions between fundamental

¹The conceptual developments behind modern quantum field theories are discussed in detail in Ref. [47].

²Lagrangian *density* \mathcal{L} in field theory, which by integration over spacetime yields the “action” functional $S = \int d^4x \mathcal{L}$. By requiring the action to be stationary, the equations of motion associated with a given Lagrangian can be derived.

particles in the SM can be derived from “gauge” symmetries of its Lagrangian. This was already envisioned by Salam and Ward—two of the main authors of the SM—in 1961 [51]:

Our basic postulate is that it should be possible to generate strong, weak and electro-magnetic interaction terms (with all their correct symmetry properties and also with clues regarding their relative strengths), by making local gauge transformations on the kinetic-energy terms in the free Lagrangian for all particles.

In the gauge theory underlying the SM, gauge fields are brought about by requiring the invariance of a free (non-interacting) Lagrangian under local matter-field transformations that form a Lie group [49]. Gauge bosons are the gauge fields’ quanta and act as carriers of a force. The strength of the force is given by the gauge coupling constant associated with the symmetry group. Of particular relevance to the SM is the fact that when the symmetry group is the direct product of at least two subgroups, each of these is associated with an independent coupling constant.

An *unbroken* gauge symmetry is unable to account for masses of gauge bosons, in contradiction with the observed short range of the weak interaction. Contrariwise, the (explicit or spontaneous) breaking of a gauge symmetry allows its gauge fields to acquire non-zero masses. In the SM, the “Brout–Englert–Higgs mechanism” [52, 53, 54, 55, 56, 57] is responsible for the spontaneous breaking of the symmetry group of the unified electroweak interaction.

Gauge groups of the Standard Model

The SM gauge symmetry is described by the Lie groups $U(n)$ and $SU(n)$, the sets of all $n \times n$ unitary matrices u , with $\det u = \exp(i\phi)$ and $\det u = 1$, respectively. $U(n)$ is of dimension $d = n^2$, corresponding to the $2n^2$ independent real parameters of a complex $n \times n$ matrix of which n^2 are removed by the unitarity constraint. $SU(n)$ is of dimension $n^2 - 1$ owing to the additional constraint on the determinant.

In general, there is one “generator” for each continuous parameter of the groups above. The set of d linearly independent generators—which must satisfy certain commutation relations (“Lie algebra”)—then completely determines the group elements. An “Abelian” group has mutually commuting generators only, whereas a “non-Abelian” group exhibits a non-trivial Lie algebra. Any given “representation” of the Lie algebra consists of d $m \times m$ generator matrices which preserve the original commutation relations. It allows finding a group representation in an m -dimensional

vector space. The “fundamental” representation has $m = n$ and is denoted as \mathbf{n} . For example, the generators of the fundamental representation, $\mathbf{2}$, of $SU(2)$ are the matrices $\sigma_i/2$ ($i = 1, 2, 3$), where σ_i is the i th Pauli spin matrix. This representation is also referred to as the “spin-1/2” or “doublet” representation.

2.2 Development of Standard Model building blocks

Stated succinctly, the SM corresponds to a spontaneously broken quantised Yang–Mills theory with a (local) $G_{\text{SM}} \equiv SU(3)_C \times SU(2)_L \times U(1)_Y$ gauge symmetry: the elements of the $SU(3)_C$ symmetry group act non-trivially only on matter fields carrying “colour charge”, those of $SU(2)_L$ only on “left-chiral” field components, and those of $U(1)_Y$ only on fields carrying “weak hypercharge”. As sketched below (mainly based on reviews from Refs. [58, 59, 60, 47]), it was only the combination of experiment, theoretical reasoning, mathematical calculations, and aesthetic ideals that would eventually bring into existence the vision of Salam and Ward.

Quantum electrodynamics

The first building block of the SM is the quantum field theory of electromagnetism, or “quantum electrodynamics” (QED). The theoretical foundations of QED specifically, and of quantum field theory (QFT) more broadly, were laid by painstaking, often confused efforts in the 1920s to quantise the classical electromagnetic field [47]. Most notably, in a groundbreaking work [61] in 1927, Dirac devised a quantum theory of a non-relativistic atom interacting with quanta of light, or “photons” [62]. This work also laid the conceptual foundation of QFT in general, the inappropriately named [47] framework of “second quantisation”.

QED has a $U(1)_Q$ gauge symmetry, which implies the existence of the photon as the gauge boson responsible for the electromagnetic interaction between two electrically charged particles. It represents the residual gauge symmetry of the electroweak interaction at energies below the unification scale.

The QED prediction of the electron anomalous magnetic moment has a relative precision of the order of 0.1 parts per billion and has impressively withstood experimental tests to the tenth significant digit [63].³

³Confronting the predicted values of both the electron and the muon anomalous magnetic moment with experiment at ever-increasing precision provides a powerful test of the SM and leads to constraints on possible physics beyond the SM [64, 65].

Electroweak unification

By unifying the electromagnetic with the weak interaction, “electroweak theory” represents the second core building block of the SM.

The theoretical foundations of the weak interaction were laid by Fermi in 1934. He formulated his theory of β -decay [66, 67]—which adopted Pauli’s earlier proposal of the neutrino [68]—by analogy to photon emission from an excited atom in QED: a four-fermion contact interaction is expressed as the product of two vector currents (see Sec. 2.3 below) at a single spacetime point, i.e., the weak current of the proton-neutron pair and that of the electron-neutrino pair. The theory led to extensive research into the neutrino hypothesis (including the neutrino mass) or the shape of the electron energy spectrum.

Eventually, experimental studies culminated in the “ θ - τ ” puzzle, the intrinsic left-right (parity) mismatch between two weak decays whose parent mesons share the same mass and lifetime. It was resolved in 1956 when Lee and Yang suggested a number of experimental tests to look for a breakdown of parity in weak interactions [69], which was confirmed only a few months later in weak decays [70, 71, 72]. As a result, the theory of the massless “two-component” neutrino was invoked [73, 74, 75] as a natural explanation for the observed non-conservation of parity, at least for weak charged-current processes involving the neutrino.

The two-component theory was confirmed by Goldhaber et al. in 1958 [76], but a generalisation realised in “ $V - A$ theory” [77, 78, 79, 80] was still necessary. Here, the Hamiltonian of the weak charged-current interaction is the product of a pair of currents which are differences between a term transforming like a vector (V) and a term transforming like an axial vector (A) under parity.

A complete joint theory of the weak and electromagnetic interaction was developed through the seminal work of Glashow, Salam, and Weinberg [81, 82, 83, 84] during the 1960s. Its final formulation relied on two major theoretical advancements which had occurred a few years before: Yang and Mills generalised the local gauge principle behind QED [85], and the Brout–Englert–Higgs mechanism [52, 53, 54, 55, 56, 57] allowed for the existence of massive charged gauge bosons as the result of spontaneous symmetry breaking. However, it was only after the proof [86, 87] that spontaneously broken Yang–Mills theories were indeed “renormalisable” [47] that electroweak theory received significant attention [88]. Its underlying $SU(2)_L \times U(1)_Y$ symmetry structure—which exhibits a spontaneous breakdown to the $U(1)_Q$ gauge group of QED—gives rise to one neutral gauge boson, Z^0 , and two charged gauge bosons, W^\pm , responsible for the interactions of neutrinos, among others.

Experiments at CERN soon confirmed the existence of the weak neutral current [89, 90, 91]—involving a neutrino interaction with an electron or a nucleon—and of the postulated Z^0 [92, 93] and W^\pm bosons [94, 95]. In 2012, the ATLAS and CMS experiments, also at CERN, announced the observation of a particle compatible with the Higgs boson predicted by the Brout–Englert–Higgs mechanism [96, 97, 15].

Quantum chromodynamics

Strong interactions within the SM are described by “quantum chromodynamics” (QCD), characterised by $SU(3)_C$. Significant insight into its nature was gained through the discovery of an approximate, “accidental” $SU(3)$ symmetry obeyed by hadrons, referred to as the “eightfold way” [98, 99, 100], by Gell-Mann and Ne’eman in 1961 [101]. This is a classification scheme devised in response to the preceding discovery of a number of strongly interacting particles in cosmic-ray and collider experiments, the first being the charged pion in 1947 [102].

Another significant discovery was made around the same time when “two photographs containing forked tracks of a very striking character” [103]—taken with a cloud chamber exposed to cosmic rays—proved the existence of the neutral kaon. In order to explain the unexpectedly large decay times of the kaon, and of the heavier hyperons Σ , Λ^0 , and Ξ , all of which had been observed only shortly thereafter [104, 105], Nakano, Nishijima, and Gell-Mann introduced the concept of “strangeness” [106, 107, 108, 109, 110], a quantum number conserved by strong but violated by weak interactions.

According to the eightfold way, mesons and baryons are classified into different representations (“multiplets”) of $SU(3)$, with each multiplet member described by either its electric charge and strangeness, or, equivalently, by its “strong hypercharge” and by the third component of its “strong isospin”. The eightfold way famously predicted the Ω^- with strangeness -3 , which was discovered in 1964 [111, 112]. In addition, it led to the identification of the “up” (u), “down” (d), and “strange” (s) quarks with the fundamental ($\mathbf{3}$ or “triplet”) representation of the approximate $SU(3)$ flavour symmetry by Gell-Mann and Zweig [113, 114].

Shortly after, the notion of “colour charge” was introduced as an additional quantum number in order to solve quark-statistics problems [115, 116]. “Deep inelastic” electron-nucleon scattering experiments conducted at SLAC in 1968 exposed the proton’s quark structure [117, 118, 119, 120], thereby directly confirming the existence of the up and the down quark and proving that they were not purely “fictitious, mathematical devices” [60] responsible for the symmetry pattern of hadronic

bound states. The fourth, “charm” (c), quark was predicted in 1970 as a requirement of the “GIM” mechanism [121]. Neutrino scattering in the Gargamelle bubble chamber [122, 123] was employed to gain further insight into the nature of quarks as point-like constituents of the nucleon.

Soon, during the early 1970s, QCD was proposed as the gauge theory describing the strong interactions of quarks [124, 125, 126, 127]. The non-Abelian character of QCD implies that the exchange particles, massless neutral gauge bosons called “gluons”, interact with each other as a result of carrying colour charge themselves. “Asymptotic freedom” [126, 127] dictates that neither quarks nor gluons be observable as free particles.

2.3 Free Dirac fermions

The theory of free Dirac fermions, which is briefly summarised below based on Ref. [128], serves as the foundation for the description of neutrinos in the SM. As particles with spin $1/2$ —as are all other fundamental fermions of the SM—their kinematics are described by the Dirac equation [129]

$$(i\rlap{/}\partial - m)\Psi(x) = 0. \quad (2.1)$$

Here, $\Psi(x)$ is a non-interacting *four-component* fermion field, m is the associated fermion’s mass, and $\rlap{/}\partial \equiv \gamma^\mu \partial_\mu$, where γ^μ ($\mu = 0, 1, 2, 3$) are 4×4 matrices. These (“Dirac” or “ γ ”) matrices are defined by the anticommutation relations

$$\gamma^\mu \gamma^\nu + \gamma^\nu \gamma^\mu = 2g^{\mu\nu} \mathbf{1}, \quad (2.2)$$

where $g^{\mu\nu}$ is the metric tensor, and $\gamma^0 \gamma^{\mu\dagger} \gamma^0 = \gamma^\mu$. Despite the notation, the γ matrices are not four-vectors but constants that may occur in different representations.

The “adjoint” field $\bar{\Psi}(x) \equiv \Psi^\dagger(x) \gamma^0$ can be employed to construct the conserved current $j^\mu \equiv \bar{\Psi}(x) \gamma^\mu \Psi(x)$, which transforms as a four-vector (Lorentz vector). More generally, any fermion “bilinear” $\bar{\Psi}(x) \Gamma \Psi(x)$, where Γ is any constant 4×4 matrix, can be decomposed into 16 linearly independent bilinears with well-defined Lorentz transformation properties:

- the single component of $S \equiv \bar{\Psi} \Psi$ transforms as a Lorentz scalar and
- the single component of $\tilde{S} \equiv \bar{\Psi} \gamma^5 \Psi$ (where $\gamma^5 \equiv i\gamma^0 \gamma^1 \gamma^2 \gamma^3$) as a pseudoscalar,

- the four components of $V^\mu \equiv \bar{\Psi}\gamma^\mu\Psi$ transform as a vector and
- the four components of $\tilde{V}^\mu \equiv \bar{\Psi}\gamma^5\gamma^\mu\Psi$ as an axial vector,
- and the six components of $T^{\mu\nu} \equiv \bar{\Psi}\sigma^{\mu\nu}\Psi$ (where $\sigma^{\mu\nu} \equiv i\gamma^\mu\gamma^\nu$, $\mu \neq \nu$) transform as a tensor.

Such bilinear currents serve the construction of the necessarily Lorentz-scalar terms of the Hermitian Lagrangian of the theory, whether it contains only a single free fermion field or several different interacting fermion and boson fields (such as the SM). For example, the Lagrangian from which the Dirac equation (2.1) can be derived reads

$$\mathcal{L}_D = \frac{i}{2} (\bar{\Psi}\gamma^\mu\partial_\mu\Psi - \partial_\mu\bar{\Psi}\gamma^\mu\Psi) - m\bar{\Psi}\Psi, \quad (2.3)$$

where the first two are “kinetic terms” and the third is a “mass term”.

Each current also exhibits specific transformation properties under space inversion (parity) P , charge conjugation C , and time reversal T ; the combination CPT of these, no matter the order of the individual operations, leaves each bilinear invariant up to a possible and irrelevant phase. Hence, CPT corresponds to an exact symmetry of a Lorentz-invariant quantum field theory such as the SM [130, 131, 132, 133]. It implies the simultaneous conservation or violation of

$$CP \leftrightarrow T, \quad CT \leftrightarrow P, \quad PT \leftrightarrow C. \quad (2.4)$$

Chirality

A general fermion field $\Psi(x)$ which solves the Dirac equation is a “bispinor”, constructed from two two-component, “Weyl” spinors. A Weyl spinor transforms under either one of two distinct irreducible two-dimensional representations (“left-” vs. “right-chiral”) of the Lorentz group $SO(1,3)$, which transform in opposite manners under Lorentz boosts and which are related to each other by the parity transformation [49, 128]. Using the “chiral representation” of the Dirac matrices, a bispinor can be written as

$$\Psi \equiv \begin{pmatrix} \chi_R \\ \chi_L \end{pmatrix}, \quad (2.5)$$

where χ_R and χ_L are both two-component spinors. The Dirac equation (2.1) then splits into the two coupled equations

$$i\not{\partial}\Psi_R = m\Psi_L, \quad (2.6)$$

$$i\not{\partial}\Psi_L = m\Psi_R, \quad (2.7)$$

with right- and left-chiral fields

$$\Psi_R = \begin{pmatrix} \chi_R \\ 0 \end{pmatrix}, \quad \Psi_L = \begin{pmatrix} 0 \\ \chi_L \end{pmatrix}, \quad (2.8)$$

respectively.

In any representation of the Dirac matrices, the chiral components of a general Dirac spinor Ψ are projected out by the two projection matrices

$$P_L \equiv \frac{1 - \gamma^5}{2}, \quad P_R \equiv \frac{1 + \gamma^5}{2}, \quad (2.9)$$

with $\Psi_L = P_L\Psi$ and $\Psi_R = P_R\Psi$. The chirality of a field Ψ is well-defined only if Ψ is an eigenfunction of the ‘‘chirality matrix’’ $\gamma^5 \equiv \gamma_5$, which anticommutes with all Dirac matrices. On the one hand, this property implies that chirality is not conserved, since γ^5 anticommutes with the mass term in Eq. (2.3). On the other, chirality is Lorentz-invariant, since γ^5 commutes with the generators of the Lorentz group [49]. By convention, the right-chiral field Ψ_R is assigned the eigenvalue +1, whereas the left-chiral field Ψ_L is assigned the eigenvalue -1 . In any case, a bispinor Ψ can be written as the sum of its chiral components, $\Psi = \Psi_L + \Psi_R$.

The fact that the Dirac equation reduces to a simpler version in the case of a massless fermion—requiring fields with only two independent components—was demonstrated by Weyl in 1929 [134]. Due to the decoupling of the dynamics of the chiral fields Ψ_R and Ψ_L for $m = 0$, as seen from Eqs. (2.6) and (2.7), a massless fermion can in principle be described by either of these alone. Note that both the quantised left-chiral field (operator) as well as its adjoint, Ψ_L respectively $\overline{\Psi}_L$, describe only particles with left-handed helicity (antiparallel spin and momentum) *and* antiparticles with right-handed helicity (parallel spin and momentum). The right-chiral field as well as its adjoint, Ψ_R respectively $\overline{\Psi}_R$, in contrast, describe only particles with right-handed helicity and antiparticles with left-handed helicity.⁴ The

⁴More precisely, the left-chiral field, for instance, contains annihilation operators for left-handed particles and creation operators for right-handed antiparticles, whereas its adjoint contains cre-




Figure 2.1: Helicity states of left- and right-chiral fermion fields, for particles and antiparticles, with spin \mathbf{s} and momentum \mathbf{p} either parallel (right-handed) or antiparallel (left-handed). Figure adapted from Ref. [128].

possible helicity states are sketched in Fig. 2.1.

From a historical perspective, the parity-violating nature of Weyl’s two-component theory of massless fermions was pointed out by Pauli in 1933 [135]: since P transforms between the left- and right-chiral components Ψ_L and Ψ_R , a theory which only implements one of these necessarily violates invariance under parity. Once parity violation had been experimentally observed, integrating the neutrino as a massless left-handed Weyl fermion into the SM represented the most economical route—thereby mimicking the left-chiral ($V - A$) nature of the weak charged current within the electroweak Standard Model of Sec. 2.4 below. A two-component spinor also fully specifies a special solution of the Dirac equation referred to as “Majorana fermion” [136, 137]. As explained in Sec. 3.1, a massive neutrino could be such a type of field.

2.4 Elements of electroweak theory

Neutrino interactions are embedded in the Standard Model of electroweak interactions (EWSM), whose unified theory demands the joint consideration of the $SU(2)_L$ and $U(1)_Y$ gauge groups. No reference to strong interactions is required, owing to the lacking mixing between the electroweak, $SU(2)_L \times U(1)_Y$, and the QCD, $SU(3)_C$, sectors of the SM. This section elucidates salient features of the EWSM by strongly drawing from Ref. [128].

2.4.1 Gauge fields and fermion fields

One of the key elements of gauge theory is requiring the invariance of a non-interacting Lagrangian under a gauge transformation $U(x)$, such that a matter field

ation operators for left-handed particles and annihilation operators for right-handed antiparticles.

$\Psi(x)$ transforms as $\Psi(x) \rightarrow \Psi'(x) = U(x)\Psi(x)$, where $U(x)$ is an element of the considered gauge group, chosen *independently* at each point x in spacetime. In general, the symmetry of the Lagrangian is ensured by introducing a gauge field $A_\mu(x)$ and replacing the four-derivative ∂_μ with the derivative $D_\mu \equiv \partial_\mu + igA_\mu(x)$.

Given a particular gauge group representation R and its associated generators T_R^a ($a = 1, \dots, d$), the gauge field can be written as the linear combination $A_\mu(x) = A_\mu^a(x)T_R^a$, with a set of vector gauge fields A_μ^a . Accordingly, the EWSM introduces $d = 4$ of the latter:

$$\begin{aligned} \text{SU}(2)_L : & \quad W_\mu^1, W_\mu^2, W_\mu^3, \\ \text{U}(1)_Y : & \quad B_\mu. \end{aligned}$$

On the one hand, the gauge group choice uniquely fixes the interaction properties and number of vector gauge boson fields. On the other, experimental observations dictate the number and types of matter fields and guide the way of assigning them to representations of the group.

Leptons and quarks constitute the fundamental fermion fields, which for reasons unknown occur as members of three “families” or “generations” with identical properties except for their mass. The left-chiral fermion fields within each family are assigned to an $\text{SU}(2)_L$ doublet, which has “weak isospin” $T = 1/2$, whereas their right-chiral counterparts are singlets, with isospin $T = 0$. The lepton respectively quark doublets $L'_{\alpha L}$ and Q'_{iL} are

$$L'_{eL} = \begin{pmatrix} \nu'_{eL} \\ e'_L \end{pmatrix}, \quad L'_{\mu L} = \begin{pmatrix} \nu'_{\mu L} \\ \mu'_L \end{pmatrix}, \quad L'_{\tau L} = \begin{pmatrix} \nu'_{\tau L} \\ \tau'_L \end{pmatrix}, \quad (2.10)$$

$$Q'_{1L} = \begin{pmatrix} u'_L \\ d'_L \end{pmatrix}, \quad Q'_{2L} = \begin{pmatrix} c'_L \\ s'_L \end{pmatrix}, \quad Q'_{3L} = \begin{pmatrix} t'_L \\ b'_L \end{pmatrix}. \quad (2.11)$$

The right-chiral charged-lepton and quark singlets are given by $\ell'_{\alpha R} = e'_R, \mu'_R, \tau'_R$ respectively $q'_{iR}^U = u'_R, c'_R, t'_R$ and $q'_{iR}^D = d'_R, s'_R, b'_R$, where the superscript U (D) denotes up-type (down-type) quark fields.

The matter fields above couple to the gauge bosons via the covariant derivative

$$D_\mu = \begin{cases} \partial_\mu + igW_\mu^j \frac{\sigma^j}{2} + ig'B_\mu \frac{Y}{2} & (L) \\ \partial_\mu + ig'B_\mu \frac{Y}{2} & (R) \end{cases} \quad (2.12)$$

for doublets (L) respectively singlets (R). The coupling constant g is associated with

		T	T_3	Y	Q
lepton doublet	$L'_{\alpha L} = \begin{pmatrix} \nu'_{\alpha L} \\ \ell'_{\alpha L} \end{pmatrix}$	1/2	1/2 -1/2	-1	0 -1
lepton singlet	$\ell'_{\alpha R}$	0	0	-2	-1
quark doublet	$Q'_{iL} = \begin{pmatrix} q'_{iL}^U \\ q'_{iL}^D \end{pmatrix}$	1/2	1/2 -1/2	1/3	2/3 -1/3
quark singlets	q'_{iR}^U q'_{iR}^D	0	0	4/3 -2/3	2/3 -1/3

Table 2.1: Weak isospin T , its third component T_3 , weak hypercharge Y , and electric charge Q (in units of the elementary charge $|e|$) of the $SU(2)_L$ fundamental fermion doublets and singlets. Table adapted from Ref. [128].

$SU(2)_L$, whereas g' is associated with $U(1)_Y$. Together with the hypercharge generator Y of $U(1)_Y$, the third component of the weak isospin, T_3 , satisfies the ‘‘Gell-Mann–Nishijima relation’’

$$Q = T_3 + \frac{Y}{2}, \quad (2.13)$$

where Q is the fermion’s electric charge in units of the elementary charge. Table 2.1 specifies the different quantum numbers from Eq. (2.13) for all fermion fields in the SM.

As motivated in Sec. 2.3, right-chiral neutrino fields (ν'_{eR} , $\nu'_{\mu R}$, $\nu'_{\tau R}$)—which Eq. (2.13) would require to have $Y = 0$ —are not included in the SM.

2.4.2 Interactions

The EWSM is given by the most general renormalisable Lagrangian invariant under local $SU(2)_L \times U(1)_Y$ gauge transformations, constructed from the fermion and gauge fields as well as the Higgs field. Leaving the latter aside to begin with—for the case of massless fermions and gauge bosons—the gauge-invariant Lagrangian is given by the sum of a kinetic term \mathcal{L}_f for the fermion fields, a term $\mathcal{L}_{\text{gauge}}$ for the gauge fields’ kinetics and self-interactions, and an interaction term \mathcal{L}_{int} containing the couplings between the fermion and gauge fields,

$$\mathcal{L}_{\text{massless}} \equiv \mathcal{L}_f + \mathcal{L}_{\text{gauge}} + \mathcal{L}_{\text{int}}. \quad (2.14)$$

Here, \mathcal{L}_f is just the free Dirac Lagrangian for the fermion fields of Sec. 2.4.1. The interaction term \mathcal{L}_{int} follows naturally from the requirement of gauge invariance:

$$\mathcal{L}_f + \mathcal{L}_{\text{int}} = i\overline{L'_{\alpha L}}\not{D}L'_{\alpha L} + i\overline{Q'_{jL}}\not{D}Q'_{jL} + i\overline{\ell'_{\alpha R}}\not{D}\ell'_{\alpha R} + iq'_{\alpha R}\not{D}q'_{\alpha R} + iq'_{\alpha R}\not{D}q'_{\alpha R} + i\overline{q'_{\alpha R}{}^D}\not{D}q'_{\alpha R}{}^D. \quad (2.15)$$

Its leptonic part describes neutrino interactions, whose investigation can proceed without taking into account the mass generation mechanism, to which we briefly return afterwards.

Substituting the covariant derivative (2.12) together with the isospin and hypercharge eigenvalues of the fermion fields from Table 2.1 into Eq. (2.15) yields the following expression for the leptonic part of \mathcal{L}_{int} :

$$\mathcal{L}_{\text{int},\ell} = -\frac{1}{2} \begin{pmatrix} \overline{\nu'_{\alpha L}} & \overline{\ell'_{\alpha L}} \end{pmatrix} \begin{pmatrix} gW_3 - g'\not{B} & g(W_1 - iW_2) \\ g(W_1 + iW_2) & -gW_3 - g'\not{B} \end{pmatrix} \begin{pmatrix} \nu'_{\alpha L} \\ \ell'_{\alpha L} \end{pmatrix} + g'\overline{\ell'_{\alpha R}}\not{B}\ell'_{\alpha R}. \quad (2.16)$$

While its off-diagonal terms describe charged-current (CC) interactions, the diagonal terms give rise to neutral-current (NC) interactions.

By defining the linear combination $W^\mu \equiv (W_1^\mu - iW_2^\mu)/\sqrt{2}$, the leptonic CC Lagrangian can be written as

$$\begin{aligned} \mathcal{L}_{\text{CC},\ell} &= -\frac{g}{\sqrt{2}} \left(\overline{\nu'_{\alpha L}} W \ell'_{\alpha L} + \overline{\ell'_{\alpha L}} W^\dagger \nu'_{\alpha L} \right) \\ &= -\frac{g}{2\sqrt{2}} \overline{\nu'_{\alpha L}} \gamma^\mu (1 - \gamma^5) \ell'_{\alpha L} W_\mu + \text{h.c.} \\ &= -\frac{g}{2\sqrt{2}} j_{W,\ell}^\mu W_\mu + \text{h.c.}, \end{aligned} \quad (2.17)$$

with the leptonic charged current

$$j_{W,\ell}^\mu = 2\overline{\nu'_{\alpha L}} \gamma^\mu \ell'_{\alpha L}. \quad (2.18)$$

The description of NC interactions is more involved because electroweak unification requires these to contain the leptonic QED Lagrangian

$$\mathcal{L}_{\gamma,\ell} = -ej_{\gamma,\ell}^\mu A_\mu, \quad (2.19)$$

with the leptonic electromagnetic current $j_{\gamma,\ell}^\mu = -\overline{\ell'_{\alpha L}} \gamma^\mu \ell'_{\alpha L}$. This can be achieved via a

rotation in the W_3^μ - B^μ plane, through the “weak mixing angle” ϑ_w :

$$\begin{pmatrix} A^\mu \\ Z^\mu \end{pmatrix} = \begin{pmatrix} \cos \vartheta_w & \sin \vartheta_w \\ -\sin \vartheta_w & \cos \vartheta_w \end{pmatrix} \begin{pmatrix} B^\mu \\ W_3^\mu \end{pmatrix}. \quad (2.20)$$

Here, Z^μ is a vector boson field that is orthogonal to A^μ . Combining the mixing relation (2.20) with the NC part of the interaction Lagrangian (2.16) results in the leptonic NC Lagrangian

$$\begin{aligned} \mathcal{L}_{\text{NC},\ell} = & -\frac{g}{2 \cos \vartheta_w} \left[\overline{\nu'_{\alpha L}} \not{Z} \nu'_{\alpha L} - (1 - 2 \sin^2 \vartheta_w) \overline{\ell'_{\alpha L}} \not{Z} \ell'_{\alpha L} + 2 \sin^2 \vartheta_w \overline{\ell'_{\alpha R}} \not{Z} \ell'_{\alpha R} \right] \\ & + g \sin \vartheta_w \overline{\ell'_{\alpha}} \not{A} \ell'_{\alpha}. \end{aligned} \quad (2.21)$$

The relation $\tan \vartheta_w = g'/g$, which prevents any coupling between the neutrinos and the photon, has been employed in order to arrive at this expression. In addition, by identifying the NC Lagrangian as the sum $\mathcal{L}_{\text{NC},\ell} = \mathcal{L}_{\gamma,\ell} + \mathcal{L}_{Z,\ell}$ of the QED Lagrangian (2.19) and the weak NC Lagrangian $\mathcal{L}_{Z,\ell}$, one finds that g and g' are related to the elementary charge e as $g \sin \vartheta_w = e$ respectively $g' \cos \vartheta_w = e$. The weak NC Lagrangian follows as

$$\mathcal{L}_{Z,\ell} = -\frac{g}{2 \cos \vartheta_w} j_{Z,\ell}^\mu \not{Z}_\mu, \quad (2.22)$$

where the leptonic weak neutral current $j_{Z,\ell}^\mu$ has been introduced. The latter is typically written as

$$\begin{aligned} j_{Z,\ell}^\mu = & 2g_L^\nu \overline{\nu'_{\alpha L}} \gamma^\mu \nu'_{\alpha L} + 2g_L^\ell \overline{\ell'_{\alpha L}} \gamma^\mu \ell'_{\alpha L} + 2g_R^\ell \overline{\ell'_{\alpha R}} \gamma^\mu \ell'_{\alpha R} \\ = & \overline{\nu'_{\alpha}} \gamma^\mu (g_V^\nu - g_A^\nu \gamma^5) \nu'_{\alpha} + \overline{\ell'_{\alpha}} \gamma^\mu (g_V^\ell - g_A^\ell \gamma^5) \ell'_{\alpha}. \end{aligned} \quad (2.23)$$

The coefficients of the NC terms coupling only left-chiral respectively right-chiral fermions in the first line of Eq. (2.23) are $g_L^f = T_3 - Q \sin^2 \vartheta_w$ and $g_R^f = -Q \sin^2 \vartheta_w$. However, $j_{Z,\ell}^\mu$ can also be expressed in terms of its vector and axial components, as done in the second line. The vector coefficients then correspond to the sum of the left- and right-chiral coefficients for the fermion in question, $g_V^f = g_L^f + g_R^f$, whereas the axial coefficients correspond to their difference, $g_A^f = g_L^f - g_R^f$.

2.4.3 Mass generation

If gauge invariance is to hold, the addition of charged-fermion⁵ and gauge-boson mass terms to the Lagrangian (2.14) is forbidden. In the case of a charged SM fermion field $f = \ell_\alpha, q_i^{U(D)}$ (unprimed for reasons that will become clear below), a mass term couples left- and right-chiral components, whose transformation behaviours under $SU(2)_L$ differ though. In the presence of a complex $SU(2)_L$ doublet consisting of a charged (ϕ^+) and a neutral (ϕ^0) scalar field,

$$\Phi(x) \equiv \begin{pmatrix} \phi^+(x) \\ \phi^0(x) \end{pmatrix}, \quad (2.24)$$

however, the following combination (considering leptons only) is invariant under the $SU(2)_L \times U(1)_Y$ gauge transformation:

$$\mathcal{L}_{\text{Yukawa},\ell} = -Y'_{\alpha\beta}{}^\ell \overline{L'_{\alpha L}} \Phi \ell'_{\beta R} - Y'^{\ell*}_{\alpha\beta} \overline{\ell'_{\beta R}} \Phi^\dagger L'_{\alpha L}, \quad (2.25)$$

where $Y'_{\alpha\beta}{}^\ell$ is a complex 3×3 matrix of ‘‘Yukawa couplings’’. These gauge-invariant terms give rise to lepton masses via the Brout–Englert–Higgs mechanism [52, 53, 54, 55, 56, 57], which assigns the Lagrangian

$$\mathcal{L}_{\text{Higgs}} = (D_\mu \Phi)^\dagger (D^\mu \Phi) - V(\Phi) \quad (2.26)$$

to the scalar doublet. The gauge-invariant Higgs potential reads

$$V(\Phi) = \mu^2 \Phi^\dagger \Phi + \lambda (\Phi^\dagger \Phi)^2, \quad (2.27)$$

with a self-coupling $\lambda > 0$. When the parameter $\mu^2 < 0$, the potential $V(\Phi)$ exhibits an infinite number of states with minimal energy, all degenerate with each other. These states satisfy $\Phi^\dagger \Phi = -\frac{\mu^2}{2\lambda} \equiv \frac{v^2}{2}$, where $v \equiv \sqrt{-\frac{\mu^2}{\lambda}}$ is the vacuum expectation value (VEV).

After spontaneous symmetry breaking, which corresponds to arbitrarily selecting a ground state from the set of degenerate vacuum states,⁶ the Higgs doublet can be

⁵A charged fermion has to be a ‘‘Dirac’’ fermion, which has a distinct antiparticle with opposite electric charge, etc. [138].

⁶The original gauge symmetry becomes hidden [128].

expressed as

$$\Phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix}, \quad (2.28)$$

where $H(x)$ is the (real) scalar Higgs field.

Charged-lepton masses

Evaluating the charged-lepton Yukawa Lagrangian (2.25) after symmetry breaking allows identifying the charged-lepton mass matrix $M'_{\alpha\beta} \equiv Y'_{\alpha\beta} v / \sqrt{2}$, which need not be diagonal. As shown in Ref. [128], a “biunitary” transformation can be performed by introducing two suitable 3×3 unitary matrices V_L^ℓ and V_R^ℓ which yield a diagonal mass matrix $M^\ell \equiv V_L^{\ell\dagger} M'_{\alpha\beta} V_R^\ell$ with real and positive diagonal elements m_α^ℓ . The charged lepton fields with definite masses are then $\ell_\alpha = \ell_{\alpha L} + \ell_{\alpha R}$, with components that comprise the column vector

$$\boldsymbol{\ell}_{L(R)} \equiv \begin{pmatrix} e_{L(R)} \\ \mu_{L(R)} \\ \tau_{L(R)} \end{pmatrix} = V_{L(R)}^{\ell\dagger} \begin{pmatrix} e'_{L(R)} \\ \mu'_{L(R)} \\ \tau'_{L(R)} \end{pmatrix}. \quad (2.29)$$

The associated Yukawa coupling strengths to the Higgs field, $y_\alpha^\ell = m_\alpha^\ell \sqrt{2}/v$, are not constrained by the SM.

Charged-lepton flavour and mass basis

With the mass generation mechanism in place, the question arises what the change of basis (2.29)—from the flavour to the mass basis—implies for CC and NC interactions of leptons.

According to Eq. (2.29), the leptonic charged current (2.18) becomes $j_{W,\ell}^\rho = 2\overline{\boldsymbol{\nu}'_L} \gamma^\rho V_L^\ell \boldsymbol{\ell}_L$, where $\overline{\boldsymbol{\nu}'_L} \equiv (\overline{\nu'_{eL}}, \overline{\nu'_{\mu L}}, \overline{\nu'_{\tau L}})$. Since all neutrinos are massless, any choice for their basis leaves the Yukawa Lagrangian unaffected. We are thus free to apply the same rotation through $V_L^{\ell\dagger}$ as for the charged lepton fields,

$$\boldsymbol{\nu}_L \equiv V_L^{\ell\dagger} \boldsymbol{\nu}'_L. \quad (2.30)$$

The resulting charged current $j_{W,\ell}^\rho = 2\overline{\boldsymbol{\nu}_L} \gamma^\rho \boldsymbol{\ell}_L = 2\overline{\nu_{\alpha L}} \gamma^\rho \ell_{\alpha L}$ differs from Eq. (2.18) only via the replacement of primed fields (flavour basis) by unprimed fields (mass basis). The same applies to the leptonic weak neutral current (2.23).

In summary, the fields of definite masses created through the Yukawa interaction of the charged lepton fields with the Higgs field are identical to the flavour fields—a direct consequence of the assumption of massless neutrinos.

Quark masses and mixing

There need to exist mass terms for both members of each weak-isospin quark doublet. The Higgs–Yukawa Lagrangian for down-type quarks is constructed in the same way as the leptonic one in Eq. (2.25), and that for up-type quarks with the aid of the “conjugate” Higgs doublet $\tilde{\Phi} \equiv i\sigma_2\Phi^*$, which transforms just as Φ under $SU(2)_L$ but has opposite hypercharge. The fields of definite masses are then determined in analogy to the case of charged leptons, but with a total of four suitable 3×3 unitary matrices, V_L^U , V_R^U , V_L^D , and V_R^D , instead.

An important consequence of the Yukawa interaction of quarks with the Higgs field is revealed through the quark charged current: it is not diagonal in the mass basis but depends on the complex 3×3 product matrix

$$V_{\text{CKM}} \equiv V_L^{U\dagger} V_L^D, \quad (2.31)$$

called the “Cabibbo–Kobayashi–Maskawa” (CKM) matrix [139, 140]. Therefore, it is possible to choose the flavour fields of the up-type quarks to correspond to definite-mass fields, $\mathbf{q}_L^{\prime U} \rightarrow \mathbf{q}_L^U$, and the flavour fields of the down-type quarks to correspond to the rotated fields

$$\mathbf{q}_L^{\prime D} \rightarrow V_{\text{CKM}} \mathbf{q}_L^D. \quad (2.32)$$

Since the quark neutral current (both weak and electromagnetic) is unaffected by the change of basis, the SM does not accommodate flavour-changing neutral currents (FCNCs).

2.4.4 Effective low-energy theory

Low-energy processes for which the four-momentum transfer $|q^2| \ll m_{W,Z}^2$ appear as point-like, zero-range, interactions, in analogy to Fermi’s theory of β -decay. Accordingly, the *effective* low-energy weak CC and NC interactions are described by products of two currents,

$$\mathcal{L}_{W,\ell}^{(\text{eff})} = -\frac{G_F}{\sqrt{2}} j_{W,\ell\mu}^\dagger j_{W,\ell}^\mu, \quad (2.33)$$

and

$$\mathcal{L}_{Z,\ell}^{(\text{eff})} = -\frac{G_F}{\sqrt{2}} j_{Z,\ell}^\mu j_{Z,\ell\mu}, \quad (2.34)$$

where $j_{W,\ell}^\mu$ is given in Eq. (2.18), $j_{Z,\ell}^\mu$ in Eq. (2.23), and where G_F is Fermi's coupling constant, with

$$\frac{G_F}{\sqrt{2}} = \frac{g^2}{8m_W^2} = \frac{g^2}{8\cos^2\vartheta_w m_Z^2}. \quad (2.35)$$

Figure 2.2 shows how the W^\pm gauge boson propagator in the Feynman diagram (see Sec. 2.5.1) for the CC interaction $\nu_\alpha + \ell_\alpha^- \rightarrow \nu_\alpha + \ell_\alpha^-$ is contracted to a four-lepton vertex in this limit.




Figure 2.2: Lowest-order Feynman diagram of the CC process $\nu_\alpha + \ell_\alpha^- \rightarrow \nu_\alpha + \ell_\alpha^-$ (left), leading to an effective Fermi interaction in the case of sufficiently small four-momentum transfer (right). Figure adapted from Ref. [128].

2.5 Neutrino scattering

Ever since the experimental discovery of the weak neutral current via the reactions $\bar{\nu}_\mu + N \rightarrow \bar{\nu}_\mu + X$ and $\bar{\nu}_\mu + e^- \rightarrow \bar{\nu}_\mu + e^-$ and the first measurements of the weak mixing angle $\sin^2\vartheta_w$ (using the CC reaction $\bar{\nu}_\mu + N \rightarrow \mu^\pm + X$ in addition) [89, 90, 91], the use of neutrinos as scattering projectiles has not only furthered our understanding of electroweak interactions but also shed light on the internal nucleon structure. Nowadays, in particular experiments studying neutrino oscillations stand to benefit from improved neutrino interaction models for nuclear targets, as the associated systematic uncertainties limit the measurement precision and potentially introduce biases [141].

This thesis exploits the coherent (elastic) forward scattering on Earth matter of propagating neutrinos that are naturally produced in the atmosphere to search for neutrino properties beyond the minimal SM considered throughout this chapter. A minor fraction of these neutrinos can subsequently be detected as they undergo incoherent interactions—which give rise to charged leptons or lead to the breaking-up of a nucleus—with the constituents of large volumes of transparent media. In light of the magnitudes of the four-momentum transfers associated with both, coherent and incoherent, types of interactions in this work, it is possible to adopt the low-energy limit from Sec. 2.4.4.

To gain insight into the above scattering processes, an introduction of useful general concepts is followed by brief accounts of the specific processes’ theoretical descriptions as well as selected phenomenological aspects and experimental applications. All discussions are restricted to the lowest order, or “tree level”.⁷

2.5.1 General concepts

The probability for a neutrino interaction to occur is represented by the interaction cross section, which allows obtaining the total interaction rate from an impinging neutrino flux (per unit area and time) as

$$\text{interaction rate} = \text{flux} \times \text{number of target particles} \times \text{cross section} . \quad (2.36)$$

In principle, the cross section is fully specified by the SM Lagrangian and the kinematics of the process under consideration. Indeed, Lorentz-covariant scattering theory allows deriving a general formula for the (differential or total) cross section of a process with two incoming particles and an arbitrary number of outgoing particles (see, e.g., Refs. [143, 144]). Crucially, this expression contains the squared magnitude of the “scattering matrix” element, $|T_{fi}|^2$.

The scattering matrix T is related to the “ S -matrix” of probability amplitudes for the transitions (including no scattering, represented by the identity matrix) from initial to final states via [143]

$$S_{fi} - \delta_{fi} = i(2\pi)^4 \delta^{(4)}(P_i - P_f) T_{fi} , \quad (2.37)$$

where P_i and P_f are the total four-momenta of the initial and final states, respectively.

⁷See Ref. [142] for a recent and comprehensive overview of advanced theoretical approaches to the description of neutrino scattering.

The “reaction matrix” elements (2.37) can be obtained from “Feynman rules”, which are graphically represented by diagrams such as those in Fig. 2.2. Each diagram corresponds to an amplitude contributing to the elements’ perturbative expansion. In the case of electroweak interactions, lowest-order diagrams—containing the lowest number of vertices—contribute with the highest weight.

Specifically, at tree level, the scattering matrix element factorises into the weak leptonic current, a propagator term, and the second, e.g., weak hadronic, current [145]. Single-particle-current matrix elements, with which we are concerned, are typically evaluated at the (arbitrary) spacetime point $x = 0$, as each can be straightforwardly translated to any other point. When the current is unknown a priori—as a result of incalculable effects of virtual strong interactions—its matrix element is expressed in the most general way compatible with the required Lorentz transformation behaviour, namely as a sum of covariant (e.g., hadron field) bilinears constructed from the available particle momenta and, if applicable, spins. In (quasi)elastic scattering processes, each term now contributes not with a constant coupling strength but with an a-priori unknown Lorentz-invariant “form factor”, which only depends on Q^2 . In inelastic processes, the “hadronic tensor”—a sum over spin-averaged bilinear products of hadronic current matrix elements—is expressed in the most general way through terms proportional to Lorentz-invariant “structure functions”, which generically depend on two linearly independent kinematic invariants (one of which could be Q^2). Additional symmetry arguments often further restrict the structure of the hadronic current or tensor and thereby determine the number and properties of non-zero form factors respectively structure functions.

The reader is referred to Refs. [143, 146] for a more in-depth discussion of the scattering matrix element and its various ingredients above.

2.5.2 Neutrino-electron scattering

The elastic scattering process

$$\bar{\nu}_\alpha + e^- \rightarrow \bar{\nu}_\alpha + e^- \quad (2.38)$$

allows low-energy neutrinos and antineutrinos of any flavour to interact with electrons because it is without an energy threshold. Depending on the flavour, different interaction channels contribute to the scattering amplitude.

When $\alpha = \mu, \tau$ the scattering matrix element only receives a single contribution,

from the exchange of a Z^0 , [128, 143]

$$T^{(\text{eff})} = T_Z^{(\text{eff})} \propto \frac{G_F}{\sqrt{2}} [\bar{\nu}_\alpha \gamma^\rho (1 - \gamma^5) \nu_\alpha] [\bar{e} \gamma_\rho (g_V^e - g_A^e \gamma^5) e] \quad (\alpha = \mu, \tau), \quad (2.39)$$

where g_V^e is the vector and g_A^e the axial coupling constant for the electron defined in Sec. 2.4.2.

The electron (anti)neutrino in addition interacts with the electron via the CC process depicted in Fig. 2.2. Accordingly, the scattering amplitude takes the form

$$T^{(\text{eff})} = T_W^{(\text{eff})} + T_Z^{(\text{eff})} \propto \frac{G_F}{\sqrt{2}} \left\{ [\bar{e} \gamma^\mu (1 - \gamma^5) \nu_e] [\bar{\nu}_e \gamma_\mu (1 - \gamma^5) e] + [\bar{\nu}_e \gamma^\mu (1 - \gamma^5) \nu_e] [\bar{e} \gamma_\mu (g_V^e - g_A^e \gamma^5) e] \right\}. \quad (2.40)$$

By making a ‘‘Fierz transformation’’ (see Ref. [128], for example), Eq. (2.40) can be cast into the same form as the NC amplitude (2.39), resulting in

$$T_W^{(\text{eff})} + T_Z^{(\text{eff})} \propto \frac{G_F}{\sqrt{2}} [\bar{\nu}_e \gamma^\mu (1 - \gamma^5) \nu_e] [\bar{e} \gamma_\mu ((g_V^e + 1) - (g_A^e + 1) \gamma^5) e], \quad (2.41)$$

where the interference between the CC and NC amplitudes is reflected by the effective shifts $g_{V,A}^e \rightarrow g_{V,A}^e + 1$ of the couplings between the Z^0 and the electron.

For any neutrino type and in any frame, the total cross section is proportional to the invariant squared centre-of-mass energy $s = (p_A + p_B)^2$ [128],

$$\sigma_{\bar{\nu}_\alpha}^{(\text{eff})} \equiv \sigma^{(\text{eff})} (\bar{\nu}_\alpha e \rightarrow \bar{\nu}_\alpha e) \propto G_{FS}^2 s. \quad (2.42)$$

In the limit $\sqrt{s} \gg m_e$, the complete expressions (including antineutrino cross sections) [128, 147] correspond to the relative ratios $\sigma_{\nu_e}^{(\text{eff})} : \sigma_{\bar{\nu}_e}^{(\text{eff})} : \sigma_{\nu_{\mu,\tau}}^{(\text{eff})} : \sigma_{\bar{\nu}_{\mu,\tau}}^{(\text{eff})} \approx 1 : 0.42 : 0.16 : 0.14$. For reference, the numerical value $\sigma_{\nu_e}^{(\text{eff})} \approx 93 \times 10^{-46} \text{ cm}^2 \frac{s}{\text{MeV}^2}$ [128]. Evaluated in the laboratory frame, with the electron at rest, for a neutrino energy of $E_\nu = 1 \text{ GeV}$ this cross section amounts to approximately 10^{-41} cm^2 .

All processes’ cross sections in principle allow for the determination of the weak mixing angle $\sin^2 \vartheta_w$ [148]. For example, a precision of a few percent has been achieved by measuring the ratio of differential cross sections for ν_μ and $\bar{\nu}_\mu$, using accelerator (anti)neutrinos with energies of some tens of GeV [149]. Furthermore, the combination of existing elastic scattering measurements of lower-energy ν_e ’s and

$\bar{\nu}_e$'s, produced in accelerators respectively nuclear reactors, constrains $\sin^2 \vartheta_w$ to the level of 10 % [150].

The purely leptonic “quasielastic” scattering of neutrinos or antineutrinos on electrons is possible via the CC processes

$$\nu_\alpha + e^- \rightarrow \ell_\alpha + \nu_e \quad \text{and} \quad \bar{\nu}_e + e^- \rightarrow \ell_\alpha + \bar{\nu}_\alpha,$$

where $\alpha = \mu, \tau$. These are strongly kinematically suppressed with respect to the reaction (2.38): in order to produce a μ^- (τ^-) in the laboratory frame, the energy of the incoming neutrino needs to exceed approximately 11 GeV (3 TeV). The mass m_{ℓ_α} of the final-state lepton leads to a suppression factor $\propto 1 - m_{\ell_\alpha}^2/s$ [128] relative to the cross section (2.42) for the elastic scattering processes.

Finally, $\bar{\nu}_e$ and e^- can annihilate to produce negatively charged hadronic states with zero baryon number, such as the semileptonic process resulting in a pair of quarks,

$$\bar{\nu}_e + e^- \rightarrow d + \bar{u}, s + \bar{u}, d + \bar{c}, \dots \quad (2.43)$$

For sufficiently high centre-of-mass energies, these processes become more complicated to describe than their purely leptonic counterparts due to the occurrence of negatively charged intermediate vector-meson ($J^P = 1^-$) resonances; the two lowest corresponding resonant $\bar{\nu}_e$ energies are 580 GeV for the ρ^- and 780 GeV for the K^{*-} , with a cross section of $\mathcal{O}(10^{-38} \text{ cm}^2)$ [151]. As of today, these are unobserved, but are considered realistic near-term targets for the LHC forward physics facility or IceCube [152]. The “Glashow resonance” [153], $\bar{\nu}_e e^- \rightarrow W^-$, has a resonant energy of 6.3 PeV and a strongly enhanced cross section of $\mathcal{O}(10^{-32} \text{ cm}^2)$; a single candidate event, hinting at the existence of astrophysical $\bar{\nu}_e$'s, has been detected so far, by IceCube [154].

2.5.3 Neutrino-nucleon scattering

When a neutrino incident on a nuclear target is high enough in energy, the target's substructure becomes resolvable. A composite particle serving as a neutrino scattering target is probed by the exchanged virtual gauge boson, whose de Broglie wavelength λ determines the relevant distance scale of the process at hand. Since, for a given neutrino energy E_ν in the laboratory frame, the maximum squared four-momentum transfer $Q_{\text{max}}^2 = 2m_T E_\nu$ (m_T : target mass), it follows that smaller distance scales can be probed with increasing neutrino energy. For neutrino-nucleon scatter-




Figure 2.3: From left to right, example Feynman diagrams for CC quasielastic scattering, resonant single pion production, and deep inelastic scattering of (anti)neutrinos on protons, respectively.

ing, Q_{\max}^2 only reaches the order of $m_{W,Z}^2$ when $E_\nu \sim \mathcal{O}(\text{TeV})$, which means that a contact interaction can be assumed as long as the incident energy is well below this scale.

In the case of the charged current, the target’s structure causes a variety of processes to contribute to the total neutrino-nucleon cross section depending on neutrino energy, as illustrated by the Feynman diagrams in Fig. 2.3. From the μ^- production threshold to several hundred MeV, only quasielastic scattering (QES) is of relevance, which transforms between a proton and a neutron. As the neutrino energy increases (and the maximal Q^2 transfer grows), baryonic resonances of the nucleon can be created, followed by their subsequent decay into various multi-particle final states. These scattering processes are of the “resonance production” (RES) type. Clearly, RES is accompanied by a “continuum” of non-resonant inelastic processes that give rise to the same final states. Once the neutrino energy exceeds a few GeV, “deep inelastic” scattering (DIS) is by far the dominant interaction. Here, the neutrino scatters directly on the constituent quarks inside the nucleon, which in turn is reflected by the characteristic cross-section behaviour $\propto s$ also observed for neutrino scattering on an electron in Eq. (2.42).

One should keep in mind that there is ambiguity in the definition of the kinematic regions across which inelastic processes generally, as well as DIS specifically, occur. Usually, the onset of “shallow” inelastic scattering (SIS) is taken as the pion production threshold, corresponding to an invariant mass (centre-of-mass energy) of the final-state hadronic system of $W \approx 1.08 \text{ GeV}$, whereas DIS is primarily defined by $Q^2 > 1 \text{ GeV}^2$ with an additional kinematic criterion of typically $W > 2 \text{ GeV}$ [155] to remove resonant contributions. The accessible region in the W - Q^2 plane for the

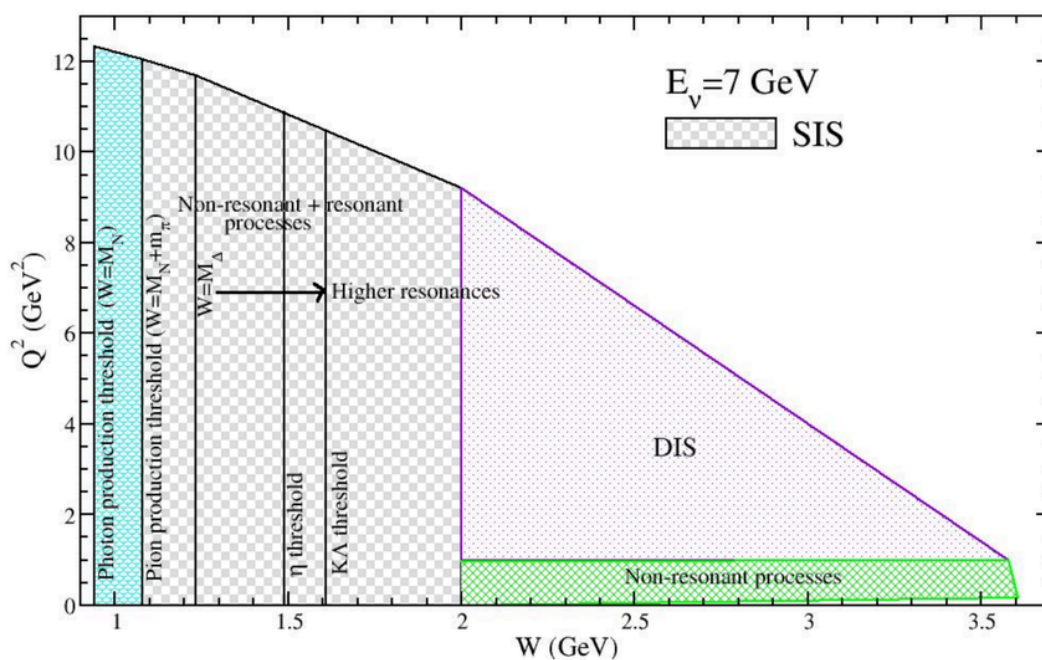


Figure 2.4: W - Q^2 plane for inelastic neutrino-nucleon scattering at $E_\nu = 7$ GeV, with an assumed onset of DIS at $Q^2 = 1$ GeV² and $W = 2$ GeV. Figure taken from Ref. [142]. See text for details.

inelastic neutrino-nucleon scattering at an example neutrino energy of $E_\nu = 7$ GeV in the laboratory frame is depicted in Fig. 2.4, subdivided into the aforementioned kinematic regimes (among others). Many theoretical and experimental efforts are dedicated to improving our understanding of the transition region from SIS to DIS—from scattering on hadrons to scattering on quarks—reflecting its strong contributions to event samples of current and upcoming long-baseline neutrino oscillation experiments [156, 142].

Figure 2.5 shows specific predictions as well as a compilation of measurements for the total ν_μ CC scattering cross section per “isoscalar” nucleon⁸ across neutrino energies from approximately 100 MeV to several hundreds of GeV [157]. The expected contributions from QES, RES, and DIS (here, different criterion of $W > 1.4$ GeV) are shown in addition, as is the predicted total NC cross section, which is flavour-independent. Note that the measurement of the ratio of NC to CC DIS event numbers induced by (predominantly) ν_μ ’s with energies of the order of 100 GeV has permitted the CCFR experiment [158] included in the figure to constrain $\sin^2 \vartheta_w$ to a somewhat higher $\sim 2\%$ precision than what has been achieved in elastic $\bar{\nu}_\mu e$ scattering.

⁸An isoscalar target has an equal number of protons and neutrons and is in a state of zero total strong isospin [145].




Figure 2.5: Predictions for CC $\nu_\mu + N \rightarrow \mu^- + X$ and all-flavour (elastic and inelastic) NC $\nu_\alpha + N \rightarrow \nu_\alpha + X$ cross section per unit neutrino energy, as a function of energy, together with measurements, for the scattering on an isoscalar nucleon, with individual calculations for QES, RES, and DIS. The central values of “axial mass” parameters, discussed in Sec. 2.5.3.1, are specified in addition, as is the invariant mass $W = 1.4 \text{ GeV}$ of the final-state hadron system at the assumed boundary between the phase spaces of RES and DIS. Modelling uncertainties are represented by coloured bands, and a separate prediction of the CC cross section by the GENIE event generator (revisited in Sec. 10.1) by the thick red dashed line. Figure taken from Ref. [157].

2.5.3.1 Elastic and quasielastic interactions

Lowest-energy neutrino interactions with nucleons either leave the latter intact, in which case the interaction is referred to as “elastic”, or lead to a change in type, $p \leftrightarrow n$, in which case it is referred to as “quasielastic”. Reactions involving a neutron in the initial state always occur in a nuclear target. However, we proceed by assuming the absence of nuclear effects in any case.

In analogy to elastic neutrino-electron scattering, the elastic (NC) process $\bar{\nu}_\alpha + N \rightarrow \bar{\nu}_\alpha + N$ with $N = n, p$ has no energy threshold. The quasielastic (CC) processes $\nu_\alpha + n \rightarrow \ell_\alpha + p$ and $\bar{\nu}_\alpha + p \rightarrow \ell_\alpha^+ + n$, in contrast, create massive (charged) leptons. Among the possible quasielastic interactions with a proton, “inverse β -decay” (IBD, also: “inverse neutron decay”),

$$\bar{\nu}_e + p \rightarrow e^+ + n, \quad (2.44)$$

has the lowest threshold. It can be induced by an electron antineutrino as low as 1.806 MeV in energy impinging on a free proton at rest. This renders IBD suitable for the detection of $\bar{\nu}_e$'s produced by nuclear fission. Notably, it was the reaction of choice in the Cowan–Reines experiment at the Savannah River Plant, which resulted in the first experimental observation of the electron antineutrino (or in fact any type of neutrino) in 1956 [159, 160].

When computing elastic and quasielastic neutrino-nucleon scattering cross sections, it is common to employ the formalism of Llewellyn Smith [161]. One of its core elements consists of constraining the structure of the weak nucleon current entering the scattering matrix element. Key steps and assumptions in deriving cross sections are summarised below based on Llewellyn Smith and Ref. [128].

Considering elastic scattering for definiteness and owing to its major impact on the propagation of atmospheric neutrinos through the Earth (see Sec. 3.3.1), a priori, since the electroweak SM only implements vector and axial-vector couplings, one typically splits the quark weak neutral current into the same two components, i.e.,

$$j_{Z,Q}^\mu(x) \equiv v_Z^\mu(x) - a_Z^\mu(x). \quad (2.45)$$

Both here and in the case of QES, for each component Lorentz covariance permits three independent complex form factors, one of which vanishes due to the absence of “second-class” currents [162].⁹ The vector and axial-vector current matrix elements

⁹The third vector, but not axial-vector, form factor is expected to vanish also as a consequence of the “conserved vector current hypothesis” (CVC).

between the two nucleon states follow as [128]

$$\langle N(p_f) | v_Z^\mu(0) | N(p_i) \rangle = \bar{u}_N(p_f) \left[\gamma^\mu F_1^{ZN}(Q^2) + \frac{i}{2m_N} \sigma^{\mu\nu} q_\nu F_2^{ZN}(Q^2) \right] u_N(p_i) , \quad (2.46)$$

$$\langle N(p_f) | a_Z^\mu(0) | N(p_i) \rangle = \bar{u}_N(p_f) \left[\gamma^\mu \gamma^5 G_A^{ZN}(Q^2) + \frac{q^\mu}{m_N} \gamma^5 G_P^{ZN}(Q^2) \right] u_N(p_i) , \quad (2.47)$$

where $u_N(p_{i,f})$ is the free nucleon's Dirac spinor and F_1^{ZN} , F_2^{ZN} , G_A^{ZN} , G_P^{ZN} are the "Dirac" (vector), "Pauli", axial, and pseudoscalar weak NC form factors, respectively. F_2^{ZN} and G_P^{ZN} are form factors of "induced" currents [163], whose Lorentz structures deviate from the intrinsic ("bare") vector (γ^μ) and axial-vector ($\gamma^\mu \gamma^5$) types and which do not contribute in the $Q^2 = 0$ limit. All form factors have to be real if strong interactions are not to violate T (CP) invariance.

Proceeding with the calculation of the cross section for elastic $\bar{\nu}_\alpha N$ scattering, based on the scattering matrix element

$$T_{\bar{\nu}_\alpha N \rightarrow \bar{\nu}_\alpha N}^{(\text{eff})} \propto \frac{G_F}{\sqrt{2}} j_{Z,\ell}^\mu \langle N(p_f) | j_{Z,Q}^\mu(0) | N(p_i) \rangle , \quad (2.48)$$

it turns out that no contribution from G_P^{ZN} remains, due to the vanishing neutrino mass. For $Q^2 \rightarrow 0$, the finite extent of the nucleon is unresolved and all form factors become coupling constants. Generally, in the low-energy regime $E_\nu \ll m_N$, any dependence on F_2^{ZN} is also removed, and the total cross section for NC elastic scattering follows as [128]

$$\sigma_{\bar{\nu}_\alpha N \rightarrow \bar{\nu}_\alpha N}^{(\text{eff})} = \frac{G_F^2 E_\nu^2}{\pi} \left[(F_1^{ZN}(0))^2 + 3 (G_A^{ZN}(0))^2 \right] . \quad (2.49)$$

Here, for the proton $F_1^{Zp}(0) = 1/2 - 2 \sin^2 \vartheta_w \approx 0.04$ and $G_A^{Zp}(0) \approx 1.27/2$, and for the neutron $F_1^{Zn}(0) = -1/2$ and $G_A^{Zn}(0) \approx -1.27/2$.

An expression closely resembling Eq. (2.49) emerges in the low-energy limit of QES. Its cross section depends on the squares of the CC couplings $F_1(0) = 1$ and $G_A(0) \approx 1.27$, as well as on the squared modulus $|V_{ud}|^2 \approx 0.95$ of the CKM matrix element for the conversion between up and down quarks, cf. Eq. (2.32).

The values substituted above for the NC and CC form factors at $Q^2 = 0$ satisfy several theoretically motivated relations (for arbitrary Q^2) expressing, first, weak CC form factors in terms of their (known) electromagnetic counterparts, and, second, weak NC form factors in terms of both, weak CC and electromagnetic form factors.


Their details can be found in Ref. [155] for example. As a consequence, the only unknown form factor entering the QES cross section is the axial one. It is typically fit by assuming a “dipole” shape

$$G_A(Q^2) = \frac{G_A(0)}{\left(1 + \frac{Q^2}{M_A^2}\right)^2} \quad (2.50)$$

with some “axial mass” parameter M_A . Beyond this, the only unknown form factors entering the general expression for the elastic scattering cross section are associated with the virtual heavier quarks’ bare vector and axial-vector neutral currents, referred to as “strange(ness) form factors” [128, 155].

2.5.3.2 Resonance production

As can be seen from Fig. 2.5, the excitation of nucleon resonances is an appreciable contribution to the total ν_μ CC cross section only in the energy range of a few GeV. Most likely, the decay of the resonance N^* results in the creation of a single pion, in which case the full reaction chain is [164]



where $N, N' = n, p$ and $\pi = \pi^+, \pi^0, \pi^-$. In total, there are 14 processes producing a single pion, whose cross sections are subject to interference between the resonant and non-resonant contributions to the matrix element of the hadronic current [155]. The popular yet outdated model by Rein and Sehgal [165] implements 18 resonances with central invariant masses of $W_{\pi N} < 2 \text{ GeV}$ [165, Table 1], including the $\Delta(1232)$ resonance, which is dominant at low W . While the structure of the current matrix elements responsible for the transitions from a ground-state nucleon to spin-1/2 resonances is precisely that introduced in the investigation of (Q)ES, a more complicated structure arises for transitions to spin-3/2 resonances such as $\Delta(1232)$ [155]. Processes producing final states containing, for example, photons, multiple pions, or kaons, are subdominant compared to single pion production, but contribute to the overall modelling complexity of SIS.

2.5.3.3 Deep inelastic scattering

The elastic scattering on an individual quark within the nucleon may result in the breakup of the latter, which in turn gives rise to a hadronic shower due to quark recombination. Since CC or NC DIS is possible, DIS generally is defined by the (inclusive) processes

$$\bar{\nu}_\alpha + N \rightarrow \bar{\nu}_\alpha(\ell_\alpha^\mp) + X, \quad (2.51)$$

where X represents the set of final-state hadrons (not restricted to a single nucleon). The quark-level counterparts of Eq. (2.51) are [128]

$$\nu_\alpha + N \text{ CC} : \quad W^+ + d, s, b (\bar{u}, \bar{c}, \bar{t}) \rightarrow u, c, t (\bar{d}, \bar{s}, \bar{b}), \quad (2.52)$$

$$\bar{\nu}_\alpha + N \text{ CC} : \quad W^- + u, c, t (\bar{d}, \bar{s}, \bar{b}) \rightarrow d, s, b (\bar{u}, \bar{c}, \bar{t}), \quad (2.53)$$

$$\bar{\nu}_\alpha + N \text{ NC} : \quad Z^0 + \bar{q}_i \rightarrow \bar{q}_i. \quad (2.54)$$

The description of DIS kinematics commonly invokes the “Bjorken scaling variable”

$$x \equiv \frac{Q^2}{2p_N \cdot q} \quad (2.55)$$

and the “inelasticity”

$$y = \frac{p_N \cdot q}{p_N \cdot p_\nu}, \quad (2.56)$$

where p_N is the nucleon and p_ν the neutrino four-momentum before scattering. The pair of invariants x and y is related to the pair Q^2 and s as $xy = \frac{Q^2}{s - m_N^2}$ [128].

The most general weak hadronic tensor relevant to any of the above DIS processes, $W_{\mu\nu}$, is typically initially expressed through six real dimensionful weak nucleon structure functions $W_i(x, Q^2)$ ($i = 1, 2, \dots, 6$), the last three of which can be neglected in the limit of vanishing lepton masses [146, 155].

Quark-parton model

While the DIS differential cross sections can be given directly in terms of the W_i by contracting the hadronic tensor with the leptonic tensor, the “quark-parton” model (QPM) [166, 167] provides an intuitive and physically well-motivated framework. Here, the nucleon is a composite system of three valence quarks and a “sea” of virtual gluons which give rise to quark-antiquark pairs. In the “Breit frame”—the frame of reference in which the nucleon carries infinite momentum along the axis of collision with the virtual gauge boson—quark transverse momenta and masses vanish. In

this system, x corresponds to the nucleon four-momentum fraction carried by the quark with which the gauge boson with four-momentum q ($q^2 = -|\mathbf{q}|^2$) interacts. Furthermore, the interacting quark is considered to be (asymptotically) free [128] (“impulse approximation”, see Sec. 2.5.4 below).

The QPM was developed in response to the “scaling” phenomenon predicted by Bjorken [168] and observed in the inelastic electron-proton scattering program of the late 1960s [169]: the dimensionful structure functions can be transformed into the dimensionless counterparts

$$F_1(x, Q^2) \equiv m_N W_1(x, Q^2), \quad F_{2,3}(x, Q^2) \equiv \frac{Q^2}{2m_N x} W_{2,3}(x, Q^2), \quad (2.57)$$

which depend only on x in the “Bjorken limit” of $Q^2, q_0 \rightarrow \infty$ with finite x [146, 155]. With these structure functions, the doubly differential weak CC DIS cross section is given by [128, 155]

$$\frac{\partial^2 \sigma_{\text{CC}}^{\bar{\nu}N}}{\partial x \partial y} = \frac{G_{\text{FS}}^2}{2\pi} \left(1 + \frac{Q^2}{m_W^2}\right)^{-2} \left(xy^2 F_1^{W^{\pm N}} + (1-y) F_2^{W^{\pm N}} \pm xy \left(1 - \frac{y}{2}\right) F_3^{W^{\pm N}} \right). \quad (2.58)$$

To obtain the corresponding NC cross sections, the substitutions $m_W^2 \rightarrow m_Z^2$ and $F_i^{W^{\pm N}} \rightarrow F_i^{ZN}$ are necessary.

The QPM considers neutrino-nucleon DIS as the incoherent scattering on free and point-like constituents, such that a given nucleon structure function follows from the sum over all “point” structure functions, i.e., quark and antiquark contributions, $F_i = \sum_q F_{iq} + \sum_{\bar{q}} F_{i\bar{q}}$. Hence, the nucleon structure functions depend on the “parton distribution functions” $f_{\bar{q}}^N(x)$, which encode the distributions of the fractional nucleon momentum x for the various quark flavours. The relations can be found in Refs. [128, 170], for example. A general expectation of the QPM is that the first two structure functions are related by $F_2 = 2xF_1$ [171].

In the experimentally relevant case of an isoscalar target, each structure function is the average of the corresponding proton and neutron ones. To halve the number of parton distribution functions required to describe the cross section for a given DIS process, simple relations [128] between the various distribution functions of the proton and neutron can be applied under the assumptions of strong-isospin symmetry and of a symmetric sea. Even though the QPM does not predict the individual parton distribution functions, in the absence of sea quarks it would imply a ratio of the integrated isoscalar neutrino and antineutrino CC cross sections of $\sigma_{\text{CC}}^{\nu}/\sigma_{\text{CC}}^{\bar{\nu}} = 3$.

Experiment however obtains a neutrino cross section that is only roughly twice as large as its antineutrino counterpart, due to the antiquark component within the nucleon [145]. As a further example, the two integrated CC cross sections also allow determining the fraction of the total nucleon momentum carried by quarks and antiquarks combined, measured to be $r_{q+\bar{q}} \approx 0.5$, with a relative contribution from (sea) antiquarks of $r_{\bar{q}} \approx 0.15r_{q+\bar{q}}$ [155]; the other half must thus originate from gluons.

2.5.4 Nuclear targets

When scattering takes place on a nucleus—for example in modern neutrino experiments that rely on heavy targets such as oxygen, carbon, or argon to maximise event rates—the description of neutrino-nucleon interactions needs to be supplemented with a characterisation of nuclear medium effects. It should account for, among others, the nucleon’s Fermi motion and nuclear binding energy, the Pauli exclusion principle (or “Pauli blocking”, which leads to a reduction of the scattering phase space for a nucleon remaining in the nucleus), nucleon-nucleon correlations beyond the Pauli principle, and the intranuclear hadron transport including the possibility of reinteraction [155]. Figure 2.6 sketches and contrasts this complex picture with a simpler interaction model, the “impulse approximation” (IA) [173], in which the nucleus is characterised by a “spectral function” describing the momentum and energy distributions of the constituent nucleons and in which scattering is assumed to proceed on a single nucleon, thereby producing some hadronic state that does not interact with the remaining $A - 1$ non-interacting “spectator” nucleons. This picture constitutes a rough nuclear equivalent of the QPM employed for DIS, one of the main differences being the ability of an ejected hadron to remain free. Electron scattering data suggests that the IA is only applicable in the regime $|\mathbf{q}| \gtrsim 400 \text{ MeV}$ [174, 175]. Processes with a lower three-momentum transfer probe more than a single nucleon.




Figure 2.6: Sketch of two different approaches to modelling neutrino-nucleus interactions, with interactions between individual nucleons modifying the scattering dynamics (top) or scattering on a single independent nucleon surrounded by “spectators” (bottom). Figure adapted from Ref. [172].

In fact, in the limit

$$|\mathbf{q}|R_A \ll 1, \quad (2.59)$$

where R_A is the nuclear radius of a few fm, the scattering takes place on the nucleus as a whole, which allows for the occurrence of coherent neutrino-nucleus scattering (CE ν NS). The associated experimental challenges—in particular detecting nuclear recoils with small kinetic energy $T_A \sim \mathcal{O}(\text{keV})$ —are reflected by the passing of more than 40 years between the first suggestions [176, 177] to exploit CE ν NS as a probe of the weak neutral current and its first observation [178] in 2017 by the COHERENT experiment [179].

To conclude this chapter, below we investigate coherence in the context of the inclusive NC neutrino-nucleus scattering process, making use of the arguments and cross-section derivations in Refs. [180, 181].

2.5.4.1 Coherence

NC scattering processes of the type

$$\bar{\nu}_\alpha + A_N^Z \rightarrow \bar{\nu}_\alpha + A_N^{(*)Z} \quad (2.60)$$

conserve the integrity of the nucleus A_N^Z but might leave it in some excited state A_N^{*Z} . Without excitation, the scattering is elastic, otherwise inelastic. With characteristic inverse nuclear sizes ranging from 25 MeV to 150 MeV, neutrinos from a variety of sources (reactor, supernovae, solar, accelerator) are likely to satisfy the coherence limit (2.59) [182], in which “all nucleons take part in the scattering at the level of scattering amplitudes (and not cross sections)” [181]—independent of whether elastic or inelastic scattering is considered.

In theory, the internal excitation of the nucleus by an amount $\Delta\epsilon_{fi}$ modifies its recoil energy with respect to elastic ($\Delta\epsilon_{fi} = 0$) scattering. When the neutrino is scattered forward, the recoil energy T_A takes its minimum value, and when it is scattered backward, the recoil energy becomes maximal. An experiment with $E_\nu \gg \Delta\epsilon_{fi}$ and without sensitivity to γ -rays from nuclear de-excitation is unable to distinguish elastic from inelastic processes. According to Refs. [180, 181], a constructive interference of the individual neutrino-nucleon scattering amplitudes, with a cross section proportional to the squared number of nucleons, on theoretical grounds is exclusively obtained in the case of elastic scattering. Crucially, the latter is assumed to preclude the possibility of an individual nucleon’s spin flip, as opposed to inelastic scattering. The nuclear recoils observed in the COHERENT experiment, for instance,

have contributions from elastic (or coherent) and inelastic (or incoherent) scattering.

The three-momentum transfer \mathbf{q} effectuates a smooth transition between the two scattering regimes. Illustratively, the elastic process contributes to the cross section a term $\propto A_N^2 |F_N(\mathbf{q}^2)|^2$ ($\mathbf{q}^2 \equiv |\mathbf{q}|^2$), while the inelastic process contributes a term $\propto A_N (1 - |F_N(\mathbf{q}^2)|^2)$, where $F_N(\mathbf{q}^2)$ is the nucleon form factor¹⁰ of the nucleus, whose limits are

$$\lim_{|\mathbf{q}| \rightarrow 0} F_N(\mathbf{q}^2) = 1 \quad (\text{full coherence}), \quad (2.61)$$

$$\lim_{|\mathbf{q}| \rightarrow \infty} F_N(\mathbf{q}^2) = 0 \quad (\text{full loss of coherence}). \quad (2.62)$$

This implies that as $|\mathbf{q}|$ grows beyond R_A^{-1} , and the “classical” coherence limit (2.59) is violated, elastic scattering becomes exceedingly improbable.

In the case of an unpolarised target, starting from Ref. [181, Eq. (131)], one finds that the coherent contribution to the observable (inclusive) differential cross section for a spin-zero nucleus can be expressed as

$$\frac{d\sigma_{\text{coh}}}{dT_A} \approx \frac{G_F^2 m_A}{\pi} \left(1 - \frac{T_A}{T_{A,\text{max}}}\right) |F|^2 \frac{Q_W^2}{4}, \quad (2.63)$$

where the form factors for neutrons and protons are assumed to be the same [179, 181]. This expression also introduces a quantity of which literature makes frequent use [155], namely the “weak charge” Q_W of the nucleus, which is related to the latter’s NC vector coupling G_V as

$$Q_W \equiv 2G_V = 2(Zg_V^p + Ng_V^n). \quad (2.64)$$

The quantity $g_V^N \equiv F_1^{ZN}(0)$ is the nucleon’s weak NC vector coupling previously encountered in Sec. 2.5.3.1. Due to the smallness of g_V^p , Q_W is dominated by the neutron number N of the target nucleus.

The observable cross section is nearly independent of whether a neutrino or an antineutrino is impinging on the target. Any differences are suppressed by at least one power in the small inelasticity $y = T_A/E_\nu$ or by products of the numbers of unpaired nucleon spins. Furthermore, in the forward direction, where the nuclear recoil energy $T_A \rightarrow 0$ ($\mathbf{q} \rightarrow 0$), the coherent cross section is independent of E_ν .

¹⁰Form factors are not explicitly constructed to be Lorentz scalars here. Their indicated sole dependence on $|\mathbf{q}|$ applies to spherically symmetric proton/neutron density distributions [143].

3

Neutrino mass and mixing

The minimal Standard Model laid out in Chapter 2 must be incomplete because its combination of particle content and gauge group leaves no room for massive neutrinos. Generically, neutrino mass is expected to lead to a wealth of intriguing effects that can be investigated experimentally, such as transitions between the known neutrino flavours in vacuum and matter, or new interactions with fermions included in the SM. After a necessarily condensed introduction into the vast topic of neutrino mass and the leptonic mixing it entails, the concept of an oscillation probability is examined for vacuum. The theory is then generalised to the case of ordinary matter, in which SM neutrino interactions (SI) are expected to manifest themselves in diverse ways.

3.1 Massive neutrinos

To date, only neutrino oscillations—more precisely, the oscillatory or non-oscillatory evolution of lepton flavour as a function of distance for propagating neutrinos—have provided undisputable evidence for the existence of new physics beyond the SM: at least two non-zero neutrino masses are required to adequately describe the large variety of experimental data.¹

The minimal SM implements the two-component theory of massless neutrinos because—besides only the left-chiral component taking part in weak interaction processes—at the time of its formulation there was no evidence for the existence of neutrino mass. Neutrinos must be strictly massless in the minimal SM even if higher orders in perturbation theory, beyond the tree level, and non-perturbative effects

¹As explained in Ref. [183], oscillations do not trivially imply neutrino masses.

are taken into account. This is due to the “accidental” and “non-anomalous” global symmetry of the SM Lagrangian under $U(1)_{B-L}$, where B is total baryon number and L total lepton number (see, e.g., Ref. [184] for details).² Consequently, without additional SM gauge singlet leptons among the SM’s field content, neutrino mass necessarily implies additional degrees of freedom. These could be new fields or a qualitatively different mass generation mechanism. Conceivable theories for massive neutrinos can be broadly assigned to three types [186]:

1. Yukawa interactions between neutrinos and the SM Higgs field could be enabled by the existence of N_s neutrino singlets ν_{sR} . These are called “sterile” [187] in contrast to the “active” neutrinos which participate in weak interactions. Electroweak symmetry breaking (EWSB) then allows for Dirac mass terms. If total lepton number is imposed as a fundamental symmetry of nature, massive neutrinos are Dirac fermions. If not, bare Majorana mass terms, constructed from the Weyl spinors ν_{sR} alone, are also possible, with the result that massive neutrinos are Majorana fermions. This possibility is scrutinised in more detail in Sec. 3.1.2.
2. Instead of adding neutrino fields to the SM, a scalar $SU(2)$ triplet (“Higgs”) field with a non-zero VEV may act as a new source of EWSB, thereby generating distinct Majorana neutrino masses.
3. The existence of some new source of mass, independent of the EWSB scale, could give rise to Majorana neutrino masses whose values are determined by the relative magnitudes of the two scales.

As Fig. 3.1 illustrates, the three known neutrinos have masses at the eV scale or below and are lighter than the charged fermions by at least six orders of magnitude. Neutrino Yukawa couplings to the SM Higgs field would have to be more than nine orders of magnitude smaller than those of the heaviest, third-generation quarks and leptons. Moreover, the mass range between the top quark and the electron is populated, whereas no particles fill the gap between the heaviest neutrino and the electron. These observations in turn are usually considered as indications that the SM Higgs mechanism is not the origin of neutrino mass, motivating the potential Ma-

²Given the SM field content, local gauge invariance and renormalisability requirements alone prohibit terms in the SM Lagrangian which violate the symmetry under global [185] $G_{SM}^{glob} = U(1)_B \times U(1)_{L_e} \times U(1)_{L_\mu} \times U(1)_{L_\tau}$ transformations, where $L_{e,\mu,\tau}$ are lepton flavour numbers. $U(1)_{L,B,B-L}$ are all subgroups of G_{SM}^{glob} .




Figure 3.1: Schematic overview of the masses of the known fundamental fermions, grouped into three generations, each comprising three charged fermions (charged leptons, down-type quarks, or up-type quarks) and one definite-mass neutrino. Here, the lower bound on the mass of the neutrino ν_3 is assumed to be the largest, whereas ν_1 is allowed to be massless (see Sec. 3.2.4). Figure taken from Ref. [188].

Majorana nature of neutrinos and the associated requirement of total lepton number violation.³

3.1.1 Right-chiral or sterile?

According to Sec. 2.3, chiral fermions are Weyl fermions and a general fermion field is described using the combination of a left- and a right-chiral Weyl fermion field. In a CPT-invariant theory, given a left-chiral fermion field, its CP conjugate, a right-chiral field, exists as well (and vice versa)—the two states are “Lorentz-covariant conjugates” (LCCs) of each other [138]. A sterile neutrino field is thus not per se right-chiral in nature, but a singlet of G_{SM} (or some extended gauge group). Mass mixing with active neutrinos and the exchange of a Higgs boson are possible. While sterile with respect to G_{SM} , a right-chiral neutrino might take part in new gauge interactions beyond the SM [189]. Oscillations between active and sterile neutrinos can only exist if both are superpositions of the same neutrino fields with definite mass, which in turn depends on whether the neutrino is a Dirac or Majorana fermion and on the number of active and sterile neutrino fields.

³Total lepton number and baryon number are violated also within the minimal SM, but by non-perturbative effects [184].

3.1.2 Majorana neutrinos

In 1937, Majorana employed a purely imaginary representation of the Dirac matrices, $\tilde{\gamma}_\mu = -\tilde{\gamma}_\mu^*$, which renders the Dirac equation (2.1) real and allows for real solutions [138]

$$\tilde{\Psi} = \tilde{\Psi}^* \quad (3.1)$$

to find that “there is no reason now to infer [...] the existence of [...] antineutrinos” [136, 137]. These solutions represent Majorana fermions. Switching from the Majorana representation to any representation of the Dirac matrices, the reality condition (3.1) transforms into the Lorentz-covariant expression [128, 138]

$$\Psi = \widehat{\Psi} \equiv C\overline{\Psi}^T = -\gamma_0 C\Psi^*, \quad (3.2)$$

where C is the antisymmetric unitary “charge conjugation matrix”, defined by $C^{-1}\gamma_\mu C = -\gamma_\mu^T$ in any representation. Hence, any fermion field that satisfies Eq. (3.2) is of the Majorana type. $\widehat{\Psi}$ is an LCC of Ψ since it transforms in the same way as the latter under Lorentz transformations [138]. It is also equivalent to the charge conjugate of Ψ , $C\Psi C$ (with the unitary charge conjugation operator C [128]), up to a possible phase.

Weyl fields, as the fundamental building blocks of general fermion fields, can be combined to allow for mass in a way that is compatible with the Lorentz-covariant Majorana condition (3.2). If χ denotes a left-chiral Weyl field, it is straightforward to conclude that its LCC, $\hat{\chi}$, satisfies $P_L\hat{\chi} = 0$, i.e., that it is right-chiral. The field $\chi + \hat{\chi}$ therefore satisfies Eq. (3.2). It constitutes a Majorana fermion, whose Lagrangian can contain a mass term [138].

As pointed out in Refs. [190, 138, 191], charge conjugation applied to a chiral field does not change chirality. Accordingly, the right-chiral component Ψ_R of a Majorana field must not be identified as the charge conjugate of the field’s left-chiral component Ψ_L (and vice versa). Instead, the two chiral fields are LCCs, and (effectively, see Ref. [138] and references therein) CP conjugates, of each other.

In the following, we examine neutrino mixing for the example of active Majorana neutrinos only and illustrate how it generalises to the more complex case of coexisting active and sterile neutrinos. Both discussions are based on Ref. [128], to which the reader is referred for much more detail. The leptonic mixing matrix frequently relied on throughout the remainder of this work is also defined.

3.1.2.1 Active neutrinos only

In the case of a single neutrino species, assuming no sterile neutrino field ν_R exists,

$$\mathcal{L}_{\text{mass}}^L = -\frac{1}{2}m_L\overline{(\widehat{\nu}_L)}\nu_L - \frac{1}{2}m_L\overline{\nu}_L(\widehat{\nu}_L) \stackrel{(3.2)}{=} \frac{1}{2}m_L\nu_L^T C^\dagger \nu_L - \frac{1}{2}m_L\overline{\nu}_L C\overline{\nu}_L^T \quad (3.3)$$

is the only possible (Majorana) mass term [128]. While it is forbidden by G_{SM} , the term could arise as the low-energy manifestation of some high-energy theory which encompasses the symmetries of the SM above the EWSB scale (see Sec. 3.1.4). For three generations, Eq. (3.3) reads (dropping the explicit Hermitian conjugate of the first term)

$$\mathcal{L}_{\text{mass}}^L = \frac{1}{2}\nu_{\alpha L}^T C^\dagger M_{\alpha\beta}^L \nu'_{\beta L} + \text{h.c.}, \quad (3.4)$$

with a complex symmetric mass matrix M^L . It can be diagonalised by the transformation $(V_L^\nu)^T M^L V_L^\nu = M$ with $M_{ij} = m_i\delta_{ij}$ ($i, j = 1, 2, 3$). In the mass basis,

$$\mathcal{L}_{\text{mass}}^L = \frac{1}{2}\sum_{k=1}^3 m_k \nu_{kL}^T C^\dagger \nu_{kL} + \text{h.c.}, \quad (3.5)$$

with ν_{kL} and $(\widehat{\nu}_{kL})$ combining into the massive Majorana neutrino fields $\nu_k \equiv \nu_{kL} + (\widehat{\nu}_{kL})$.

The leptonic weak charged current expressed in terms of the three mass eigenstates of charged leptons, ℓ_L , and neutrinos, $\mathbf{n}_L^T \equiv (\nu_{1L}^T, \nu_{2L}^T, \nu_{3L}^T)$, depends on the complex symmetric 3×3 ‘‘leptonic mixing matrix’’

$$U_{3 \times 3} \equiv V_L^{\ell\dagger} V_L^\nu, \quad (3.6)$$

which can be viewed as the lepton analogue of the CKM matrix in the quark sector, compare Eq. (2.31). The mixing is assigned without ambiguity to the upper (neutrino) members of the lepton doublets (2.10),

$$\begin{pmatrix} \nu_{eL} \\ \nu_{\mu L} \\ \nu_{\tau L} \end{pmatrix} = U_{3 \times 3} \mathbf{n}_L = V_L^{\ell\dagger} \begin{pmatrix} \nu'_{eL} \\ \nu'_{\mu L} \\ \nu'_{\tau L} \end{pmatrix}, \quad (3.7)$$

such that the leptonic weak charged current can be expressed as

$$j_{W,\ell}^\rho = 2\overline{\nu}_{\alpha L} \gamma^\rho \ell_{\alpha L}, \quad (3.8)$$

with $\ell_{\alpha L}$ given in Eq. (2.29).

The choice of identifying the weak eigenstates of charged leptons with their mass eigenstates and the weak eigenstates of neutrinos to correspond to the rotated states (3.7) is motivated by the large mass differences between the three charged leptons. These ensure that they are either produced and detected incoherently, as mass eigenstates (with mass determining flavour), or lose coherence over microscopic distances in practice [192].

The leptonic mixing matrix of three Majorana neutrinos can be parameterised using three physical phases and three mixing angles, and written as the product of two unitary matrices,

$$U_{3 \times 3} = U^D D^M, \quad (3.9)$$

where

$$U^D = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{\text{CP}}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{\text{CP}}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{\text{CP}}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{\text{CP}}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{\text{CP}}} & c_{23}c_{13} \end{pmatrix}, \quad (3.10)$$

with $s_{ij} \equiv \sin \theta_{ij}$, $c_{ij} \equiv \cos \theta_{ij}$, the mixing angles θ_{12} , θ_{13} , and θ_{23} , the ‘‘Dirac phase’’ δ_{CP} , and

$$D^M = \text{diag}(e^{i\lambda_1}, e^{i\lambda_2}, e^{i\lambda_3}), \quad (3.11)$$

choosing, e.g., $\lambda_3 = 0$. The ranges to which all mixing parameters can be restricted are detailed in Sec. 3.2.4.

3.1.2.2 Coexisting active and sterile neutrinos

In the single-generation example (3.3), if the right-chiral field ν_R exists in addition, $\mathcal{L}_{\text{mass}}^R = -\frac{1}{2}m_R(\widehat{\nu_R})\nu_R + \text{h.c.}$, provides another Majorana mass term. It does not violate SM gauge invariance in contrast to $\mathcal{L}_{\text{mass}}^L$.

When the coexistence of (three) active and N_s sterile neutrino fields is assumed, the most general mass term is the sum of the two mentioned Majorana mass terms and of a familiar Dirac mass term from Sec. 2.4.3:

$$\mathcal{L}_{\text{mass}}^{\text{D+M}} \equiv \mathcal{L}_{\text{mass}}^L + \mathcal{L}_{\text{mass}}^R + \mathcal{L}_{\text{mass}}^D. \quad (3.12)$$

All $N = 3 + N_s$ available left-chiral neutrino fields can be placed in a column matrix,

$$\mathbf{N}'_L{}^T \equiv \left(v'_{eL}{}^T, v'_{\mu L}{}^T, v'_{\tau L}{}^T, \overline{(v_{s_1 R})}^T, \dots, \overline{(v_{s_{N_s} R})}^T \right), \quad (3.13)$$

in order to express Eq. (3.12) with the help of an $N \times N$ mass matrix,

$$M^{\text{D+M}} \equiv \begin{pmatrix} M^L & M^{\text{D}^T} \\ M^{\text{D}} & M^R \end{pmatrix}, \quad (3.14)$$

as

$$\mathcal{L}_{\text{mass}}^{\text{D+M}} = \frac{1}{2} \mathbf{N}'_L{}^T C^\dagger M^{\text{D+M}} \mathbf{N}'_L + \text{h.c.}, \quad (3.15)$$

with 3×3 , $N_s \times N_s$, and $N_s \times 3$ mass matrices M^L , M^R , and M^{D} , respectively. The mass eigenstates can be obtained in a manner analogous to Sec. 3.1.2.1, but now through a square N -dimensional unitary neutrino mixing matrix V_L^V . This approach results in the same structure for $\mathcal{L}^{\text{D+M}}$ as Eq. (3.5), with the difference that the sum now runs over N massive left-chiral fields. Again, v_{kL} and $\overline{(v_{kL})}$ combine into the massive Majorana neutrino fields $\nu_k \equiv v_{kL} + \overline{(v_{kL})}$.

The leptonic mixing matrix $U_{3 \times N}$ entering the leptonic weak charged current is of dimension $3 \times N$ and is non-unitary,

$$U_{3 \times N} \equiv V_L^{\ell^\dagger} V_L^V, \quad (3.16)$$

as it describes the mixing of N massive neutrino fields $\mathbf{n}_L^T \equiv (v_{1L}^T, v_{2L}^T, \dots, v_{NL}^T)$ into the three active ones,

$$\begin{pmatrix} v_{eL} \\ v_{\mu L} \\ v_{\tau L} \end{pmatrix} = U_{3 \times N} \mathbf{n}_L = V_L^{\ell^\dagger} \begin{pmatrix} v'_{eL} \\ v'_{\mu L} \\ v'_{\tau L} \end{pmatrix}, \quad (3.17)$$

in analogy to Eq. (3.7). This gives rise to the same form of the leptonic weak charged current as Eq. (3.8). $U_{3 \times N}$ is parameterised by an equal number ($3 + 3N_s$) of mixing angles and physical phases ($1 + 2N_s$ Dirac phases, $2 + N_s$ Majorana phases).

The non-unitarity of $U_{3 \times N}$ implies the possibility of oscillations between active and sterile neutrinos, as both are linear combinations of the ν_{kL} ,

$$\mathbf{N}_L = U_{N \times N} \mathbf{n}_L, \quad (3.18)$$

with the column matrix \mathbf{N}_L of N left-chiral fields and the *unitary* $N \times N$ mixing matrix

$U_{N \times N}$ respectively defined as

$$\mathbf{N}_L^T \equiv \left(v_{eL}^T, v_{\mu L}^T, v_{\tau L}^T, \overline{(v_{s_1 R})}^T, \dots, \overline{(v_{s_{N_s} R})}^T \right) \quad (3.19)$$

and

$$U_{N \times N} \equiv \begin{pmatrix} U_{3 \times N} \\ V_L^{\nu}|_{N_s \times N} \end{pmatrix} = \begin{pmatrix} V_L^{\ell\dagger} V_L^{\nu}|_{3 \times N} \\ V_L^{\nu}|_{N_s \times N} \end{pmatrix}. \quad (3.20)$$

Here, $V_L^{\nu}|_{3 \times N}$ and $V_L^{\nu}|_{N_s \times N}$ are the matrices corresponding to the first three respectively last N_s rows of V_L^{ν} .

The frequently considered pure Dirac scenario with three Dirac neutrinos is constructed by assuming vanishing Majorana mass matrices $M^L = M^R = 0$. In this case, there are no sterile neutrino fields with admixtures from the left-chiral components of the neutrino fields with definite masses because the diagonalisation of the Dirac mass matrix M^D does not mix active and sterile fields—in contrast to Eqs. (3.18) to (3.20). This precludes oscillations between active and sterile neutrinos.

3.1.3 PMNS mixing matrix

The various types of mixing matrices defined above either describe changes of basis or more generally the relations between two different sets of neutrino fields. However, the remainder of this work exclusively deals with the upper left 3×3 submatrix of the unitary mixing matrix $U_{N \times N}$, Eq. (3.20), describing the mixing of the three light neutrinos with definite masses, ν_{1L} , ν_{2L} , and ν_{3L} , into the three known active neutrinos. The matrix is henceforth denoted simply as U and referred to by convention [14] as the ‘‘Pontecorvo–Maki–Nakagawa–Sakata’’ (PMNS) matrix [193, 194, 195].

The prominent role of the 3×3 PMNS matrix is motivated by neutrino mass mechanisms of the ‘‘seesaw’’ type [196], which naturally bring about small neutrino masses (see, e.g., Ref. [128] and references therein). They assume that the eigenvalues of M^R are much greater than the elements of M^D and M^L [197], with $M^L = 0$ in ‘‘type-I’’ and ‘‘type-III’’ and $M^L \neq 0$ in ‘‘type-II’’ seesaw models [186]. This gives rise to one ‘‘light’’ 3×3 and one ‘‘heavy’’ $N_s \times N_s$ mass matrix,

$$(V_L^{\nu})^T M^{D+M} V_L^{\nu} = \begin{pmatrix} M_{\text{light}} & 0 \\ 0 & M_{\text{heavy}} \end{pmatrix}. \quad (3.21)$$

As a consequence, the three active neutrinos are composed predominantly of the three light, and the N_s sterile neutrinos mainly of the N_s heavy neutrinos. In particular, the entries of the upper-right $3 \times N_s$ sub-block of the neutrino mixing matrix V_L^ν are small compared to those in its upper-left 3×3 sub-block. From Eq. (3.17) one thus finds

$$\begin{pmatrix} \nu_{eL} \\ \nu_{\mu L} \\ \nu_{\tau L} \end{pmatrix} \simeq U \begin{pmatrix} \nu_{1L} \\ \nu_{2L} \\ \nu_{3L} \end{pmatrix}, \quad (3.22)$$

such that U is unitary up to small corrections [198] and can be parameterised according to Eqs. (3.9) to (3.11).

The approximate relation (3.22) becomes exact, and U unitary, in the “minimal” model with three massive Majorana neutrinos scrutinised in Sec. 3.1.2.1 ($N_s = 0$, with unphysical M^R and M^D). The same applies to the phenomenologically relevant cases of two or three massive Dirac neutrinos ($M^{L,R} = 0$, $D^M = \mathbb{1}$), in which—whether $N_s = 2$ or $N_s = 3$ —the parameterisation of U requires the same number of mixing angles (three) and physical phases (one) [199]. More generally, all models with exactly three massive Dirac neutrinos ($N_s > 3$) also have a unitary matrix U that is parameterisable in the same manner, as they contain $N_s - 3$ completely unobservable massless sterile neutrinos [200].

3.1.4 Weinberg operator

The most economic way of generating neutrino masses—that is, without extending the SM field content and without imposing additional symmetries—is by means of an effective Lagrangian. It is assumed to be brought about by new physics at an energy scale Λ well beyond the energies at which the predictions of the SM hold. Once the SM is viewed as such an effective low-energy theory, non-renormalisable contributions (terms with energy dimension greater than four) to its Lagrangian have to be taken into account, as explained for example in Ref. [201]:

$$\mathcal{L}_{4+n}^{\text{eff}} = \sum_{n=1,2,\dots} \frac{\mathcal{O}_{4+n}}{\Lambda^n} + \text{h.c.}, \quad (3.23)$$

with G_{SM} -invariant operators \mathcal{O}_{4+n} ($n = 1, 2, \dots$) of energy dimension $4 + n$. The presence of the factors Λ^{-n} implies a suppression of the low-energy effects due to the new physics beyond the SM, in analogy to the dimension-six Lagrangian of the effective low-energy weak interaction discussed in Sec. 2.4.4, which is proportional

to $G_F \propto m_{W(Z)}^{-2}$.

The lowest-dimensional non-renormalisable term which can be constructed from SM fields and which respects the SM gauge symmetry is the well-known ‘‘Weinberg dimension-5 operator’’ [202] (set of operators implied), which can be written as [128]

$$\mathcal{L}_5^{\text{eff}} = \frac{g_{\alpha\beta}}{\Lambda} (L'_{\alpha L} \sigma_2 \Phi) C^\dagger (\Phi^T \sigma_2 L'_{\beta L}) + \text{h.c.} . \quad (3.24)$$

After EWSB, Eq. (3.24) brings about the Majorana mass term

$$\mathcal{L}_{\text{mass}}^L = \frac{1}{2} \frac{v^2}{\Lambda} g_{\alpha\beta} \nu'_{\alpha L} C^\dagger \nu'_{\beta L} + \text{h.c.} . \quad (3.25)$$

If the new-physics scale significantly exceeds the Higgs VEV, $\Lambda \gg v$, the Weinberg operator’s observable effect at low energies is the non-zero Majorana neutrino mass matrix $M_{\alpha\beta}^L = g_{\alpha\beta} v^2 / \Lambda$ encountered in Eq. (3.4). Its entries are suppressed as v/Λ with respect to the charged-fermion masses from Sec. 2.4.3.

Different implementations of the seesaw mechanism are commonly invoked as means of generating the effective Lagrangian (3.24) at the tree level. Higher-dimensional effective operators which may give rise to new neutrino interactions are considered in Chapter 4.

3.2 Neutrino propagation in vacuum

The idea that neutrinos could oscillate, that is, undergo periodic flavour transformations, was pioneered by Pontecorvo in 1957 and 1958 [193, 194]. To this end, he assumed that there could be vacuum transitions between the electron neutrino and antineutrino—processes at variance with the prevailing two-component neutrino theory [203]. In 1962, Maki, Nakagawa, and Sakata discussed neutrino mixing in the case of two flavours and foresaw the ‘‘virtual transmutation $\nu_e \leftrightarrow \nu_\mu$ ’’ [195]. The muon neutrino was discovered in the first high-energy accelerator neutrino experiment at the Brookhaven National Laboratory [204] in the same year, after independent proposals had been made by Pontecorvo [205], Markov [206], and Schwartz [207]. Shortly after, Pontecorvo introduced the term ‘‘sterile neutrino’’ to refer to left-handed antiparticles of ν_e and ν_μ [187]. The same work also anticipated oscillations between ν_e and ν_μ , as well as an observable ν_e deficit from the Sun (see

also Ref. [208]).⁴

This section briefly outlines the nowadays widely employed “standard oscillation formula” along with its implied assumptions, as well as the parameters governing its behaviour. The review is based on Refs. [210, 211, 212, 213], which provide further details and references.

3.2.1 Foundations

We consider the evolution of neutrinos that are produced with definite flavour, by the leptonic charged current

$$j_{W,\ell}^\rho = 2 \sum_{\alpha=e,\mu,\tau} \sum_k U_{\alpha k}^* \bar{\nu}_{kL} \gamma^\rho \ell_{\alpha L}, \quad (3.26)$$

which either induces transitions $\ell_\alpha^- \rightarrow \nu_k$ or the creation of $\ell_\alpha^+ \nu_k$ pairs [128]. The neutrino of definite flavour is a coherent superposition of the mass eigenstates—the single-particle states created by the field operators with definite masses—only if the process of neutrino production does not allow identifying which massive neutrino is emitted. For this to apply, the intrinsic uncertainty on the square of the mass of the produced neutrino state, σ_{m^2} , must exceed the squared-mass differences separating the neutrino mass eigenstates, Δm^2 . More generally, this requirement can be recast into the condition that the neutrino production and detection regions both have much smaller spatial extents than the “oscillation length” [210]. If the emitted neutrino is a mass eigenstate, on the other hand, no flavour transformation is possible in vacuum.

Plane waves, quantum mechanics, or QFT

Due to its simplicity, the plane-wave approximation to the neutrino states of definite masses adopted in Ref. [209] still today serves as the de-facto standard approach when neutrino oscillations are derived, often together with either the “equal-momentum” or “equal-energy” assumption. Plane waves are not localised, though, and thus cannot describe propagating particles. Moreover, neutrino mass eigenstates composing the flavour eigenstates in general have different energies as well as different momenta [214].

⁴An important early review on neutrino oscillations by Bilenky and Pontecorvo, made prior to any dedicated experimental efforts, can be found in Ref. [209].

Even though the second issue can be overcome even in the plane-wave approach [128], the first means that one has to resort to a quantum-mechanical (QM) wave-packet ansatz or a QFT approach. The goal in each case is to obtain the probability $P_{\alpha\beta} \equiv P_{\nu_\alpha \rightarrow \nu_\beta}$ for producing a neutrino of flavour α and detecting it as a neutrino of flavour β after propagation over the distance (“baseline”) L , with universal applicability to experiment. The key points underlying both approaches can be contrasted in the following manner [211]:

- The QM approach describes neutrinos produced in weak-interaction processes as spatially extended propagating wave packets. The transition amplitude

$$\mathcal{A}_{\alpha\beta} = \langle \nu_{\beta D} | \nu_{\alpha S}(t_D) \rangle \quad (3.27)$$

is the projection of the evolved state $|\nu_{\alpha S}(t_D)\rangle$ onto the detected state $\langle \nu_{\beta D}|$. Both are flavour eigenstates, identified as linear combinations of mass eigenstates. The oscillation probability corresponds to the squared modulus of the amplitude, integrated over the unknown time $T \equiv t_D - t_S$ elapsed between production and detection,

$$P_{\alpha\beta}(L) \propto \int dT |\mathcal{A}_{\alpha\beta}(L, T)|^2. \quad (3.28)$$

The normalisation factor is determined from the requirement $\sum_\beta P_{\alpha\beta} = 1$. While the neutrino production and detection processes themselves are not treated explicitly, only a small number of assumptions about the wave packets of the involved states are required, as demonstrated in Ref. [210]. Other important assumptions such as the unitarity condition of the oscillation probability need to be preimposed by hand.

- QFT treats production, propagation, and detection as a single process described by a tree-level Feynman diagram with two vertices that distinguish the neutrino flavour, connected by an internal line that corresponds to a neutrino propagator. External particles, which take part in the production and detection processes, are described by wave packets in case they are detected and as plane waves otherwise. The overall amplitude $\mathcal{A}_{\alpha\beta}$ for the whole process follows from applying Feynman rules. The time-integrated squared modulus of $\mathcal{A}_{\alpha\beta}$ does not yield an oscillation probability but a time-integrated probability for the combined process to occur. Only because this quantity factorises into individual terms for almost all relevant applications can the concept of an




Figure 3.2: Spacetime diagram representing a simplified neutrino oscillation process. At the source, a muon neutrino is first produced in the decay of a charged pion at the time T_S and location X_S . The oscillated neutrino, taken as the superposition of the two mass eigenstates ν_1 and ν_2 , is then observed via the CC scattering on a target nucleus N at the distance $L = X_D - X_S$ and after the time $T = T_D - T_S$ at the detector. Coloured bands represent the localisation of the various wave packets. The interaction regions for production and detection are sketched with filled rectangles. Figure taken from Ref. [212].

“oscillation probability” be adopted in the first place. It is given by

$$P_{\alpha\beta}(L, E_\nu) = \frac{\frac{d\Gamma_{\alpha\beta}^{\text{tot}}(E_\nu)}{dE_\nu}}{\frac{\sigma_\beta(E_\nu)}{4\pi L^2} \frac{d\Gamma_\alpha^{\text{prod}}}{dE_\nu}}, \quad (3.29)$$

where $\Gamma_{\alpha\beta}^{\text{tot}}$ is the overall rate for the process under consideration, $\Gamma_\alpha^{\text{prod}}$ the neutrino production rate, and σ_β the detection cross section.

Figure 3.2 depicts an example spacetime diagram of the neutrino oscillation process [212], initiated by a charged-pion decay (approximately at rest) producing a muon neutrino. The latter undergoes oscillation and is detected as an electron neutrino via the CC interaction with a target nucleus (also at rest)—a process that is relevant in both accelerator and atmospheric neutrino experiments (discussed in Chapter 5). In the figure, each coloured band shows the localisation in spacetime of the wave packet corresponding to a given particle. The propagating neutrino is assumed to be the superposition of only two mass eigenstates, for example ν_1 and ν_2 .

Their different group velocities appear as slightly different slopes of the two corresponding bands. Owing to the small range of the weak-interaction processes involved, the two interaction regions are defined by the overlaps of the wave packets.

3.2.2 Standard vacuum oscillation formula

Given known neutrino mass-squared differences and leptonic mixing matrix elements, a universal neutrino oscillation probability only depends on the energy and flavour of the neutrino produced at the source, as well as on the distance to the detector. It is independent of the details of the production and detection processes.

Two conditions have to be met for such a universal expression to exist: the neutrinos are ultrarelativistic or have quasidegenerate masses, and all decoherence effects are negligible, such that it is impossible to observe the mass of the produced neutrino [210, 212]. Both apply to current neutrino oscillation experiments that make use of terrestrial baselines and that are compatible with transitions between three active neutrino flavours [215]. Then, for neutrinos in vacuum one has the oscillation probability

$$P_{\nu_\alpha \rightarrow \nu_\beta}(L, E_\nu) = \sum_{j,k} U_{\alpha j}^* U_{\beta j} U_{\alpha k} U_{\beta k}^* e^{-i \frac{\Delta m_{jk}^2}{2E_\nu} L}, \quad (3.30)$$

whereas for antineutrinos one has

$$P_{\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta}(L, E_\nu) = \sum_{j,k} U_{\alpha j} U_{\beta j}^* U_{\alpha k}^* U_{\beta k} e^{-i \frac{\Delta m_{jk}^2}{2E_\nu} L}. \quad (3.31)$$

The case $\alpha = \beta$ describes a “survival” or “disappearance” probability, the case $\alpha \neq \beta$ a “transition” or “appearance” probability.

Oscillation amplitudes

From the expressions (3.30) and (3.31) one can tell that the oscillation amplitudes are fully determined by the quartic products of elements of the leptonic mixing matrix, which are unaffected by any possible Majorana phases.⁵ All amplitudes depend only on the mixing angles θ_{12} , θ_{13} , and θ_{23} , as well as on the phase δ_{CP} .

⁵The Dirac or Majorana nature of neutrinos has no observable consequence on neutrino oscillations [216, 217].

Oscillation phases

The oscillation phases

$$\Phi_{jk} \equiv \omega_{\text{vac}} L \equiv \frac{\Delta m_{jk}^2}{2E_\nu} L, \quad (3.32)$$

are governed by the differences $\Delta m_{jk}^2 \equiv m_j^2 - m_k^2$ of squared neutrino masses, the oscillation baseline L , and the neutrino energy E_ν . ω_{vac} is the oscillation frequency. Each of the ratios $L_{jk}^{\text{osc}} \equiv \frac{2\pi}{\omega_{\text{vac}}} = \frac{4\pi E_\nu}{\Delta m_{jk}^2}$ can be referred to as an ‘‘oscillation length’’, because it corresponds to the distance after which the phase arising from Δm_{jk}^2 reaches 2π .

For atmospheric (and accelerator long-baseline) neutrino oscillation experiments, the choice of units L [km], Δm_{jk}^2 [eV^2], E_ν [GeV] is appropriate, in which case

$$\Phi_{jk} \approx 2.53 \frac{\Delta m_{jk}^2 [\text{eV}^2] L [\text{km}]}{E_\nu [\text{GeV}]} . \quad (3.33)$$

In order to be able to observe an oscillatory behaviour due to the mass-squared difference Δm_{jk}^2 , the corresponding phase should be of $\mathcal{O}(1)$. Since atmospheric neutrino experiments study neutrino oscillations with baselines L up to Earth’s diameter and at energies E_ν of several GeV, the minimal mass-squared difference they are sensitive to is $\Delta m_{jk}^2 \sim \mathcal{O}(10^{-4} \text{ eV}^2)$.

The oscillation phases are always averaged over some distribution of L/E_ν in practice [128]. One reason is that the baseline L is subject to spatial uncertainty that arises as a result of the finite extent of both the source and detector, and possibly also as a result of the detector’s directional resolution. The other is that neutrinos are usually produced with some energy spectrum, and their detection process is subject to a finite energy resolution.

3.2.3 Discrete symmetries

The study of neutrino oscillations in vacuum allows searching for fundamental violations of discrete symmetries such as T or CP, which are uniquely related if CPT symmetry holds, as motivated in Sec. 2.3. To gain more insight, it is useful to split

Eqs. (3.30) and (3.31) into their real and imaginary parts [128],

$$\begin{aligned}
 P_{\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta}^{\langle \rangle}(L, E_\nu) &= \delta_{\alpha\beta} - 4 \sum_{k>j} \text{Re} [U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*] \sin^2 \left(\frac{\Delta m_{kj}^2 L}{4E_\nu} \right) \\
 &\quad \pm 2 \sum_{k>j} \text{Im} [U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^*] \sin \left(\frac{\Delta m_{kj}^2 L}{2E_\nu} \right), \quad (3.34)
 \end{aligned}$$

where “ \pm ” evaluates to “+” for neutrinos and “-” for antineutrinos. Figure 3.3 depicts the theoretically expected relationships between all possible neutrino and anti-neutrino oscillation channels under the individual T and CP transformations, as well as under the combined CPT operation.




Figure 3.3: Relations between different (anti)neutrino flavour transition channels in terms of discrete transformations. Figure adapted from Ref. [128].

First of all, a direct reflection of CPT symmetry is the equality $P_{\nu_\alpha \rightarrow \nu_\beta} = P_{\bar{\nu}_\beta \rightarrow \bar{\nu}_\alpha}$, which can be easily verified by comparing Eq. (3.30) to Eq. (3.31) with exchanged flavour indices. Any experimentally observed difference between these oscillation probabilities (or between mass-squared splittings or mixing parameters from neutrino vs. antineutrino probabilities [218, 219]) would imply a violation of CPT symmetry and thereby challenge our fundamental approach to the description of nature.

CP symmetry holds if the neutrino and antineutrino oscillation probabilities are found to be identical, $P_{\nu_\alpha \rightarrow \nu_\beta} = P_{\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta}$. Since the last term on the right-hand side of Eq. (3.34) is zero when $\alpha = \beta$, a CP asymmetry can only be established through appearance channels.

Finally, one expects $P_{\nu_\alpha \rightarrow \nu_\beta}^{\langle \rangle} = P_{\bar{\nu}_\beta \rightarrow \bar{\nu}_\alpha}^{\langle \rangle}$ to hold under T invariance. The only known potential source of violation of CP (and, assuming the CPT theorem, T) symmetry in neutrino oscillations in vacuum is δ_{CP} . On the one hand, $\delta_{\text{CP}} = 0, \pi$ renders U real and no violation of the aforementioned symmetries should be observed. On the other hand, their maximal violation is expected for $\delta_{\text{CP}} = \pi/2, 3\pi/2$.

3.2.4 Physical parameter space and neutrino mass ordering

Consistency in choosing the space of fundamental mixing and mass parameters such that it allows “[...] to describe all physically distinguishable values of all observables” [220] is a rather subtle problem in the lepton sector. The physical ranges of the mixing parameters are intricately connected to the definition of the neutrino mass eigenstates: it is physically meaningful how the various mass eigenstates mix into each flavour eigenstate, but a given ordering or numbering scheme (see, e.g., Ref. [221]) must lead to the same physics as any other scheme. This definition is not as straightforward as in the quark sector. Not only are neutrinos for all intents and purposes never produced or detected as mass eigenstates in the laboratory, but leptonic mixing angles are also larger than quark mixing angles [15]. We are free to choose real and positive masses of charged leptons (m_e, m_μ, m_τ) and neutrinos (m_1, m_2, m_3) [220, 190].

In a hypothetical two-neutrino case [220] with the Majorana mixing matrix

$$U_{2\times 2} \equiv \begin{pmatrix} U_{e1} & U_{e2} \\ U_{\mu 1} & U_{\mu 2} \end{pmatrix} \equiv \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} 0 & 0 \\ 0 & e^{i\lambda} \end{pmatrix}, \quad (3.35)$$

whilst the neutrino flavour eigenstates are uniquely defined by the charged leptons to which they are coupled through W^\pm exchange, with $m_e < m_\mu$, freedom remains in defining ν_1 and ν_2 :

- (a) ν_1 is lighter than ν_2 , $m_1 < m_2$;
- (b) ν_1 has a larger admixture of ν_e than ν_2 , $|U_{e1}|^2 > |U_{e2}|^2$, and ν_2 has a larger admixture of ν_μ than ν_1 , $|U_{\mu 2}|^2 > |U_{\mu 1}|^2$.

Different physical mixing parameter ranges result simply from these two definitions of the neutrino mass eigenstates [220]. Choice (a) does not place any restrictions on the mixing matrix elements and therefore leads to continuous 2π ranges $\theta, \lambda \in [-\pi, \pi]$. Choice (b) leads to the same domain for λ , but to the discontinuous domain $\theta \in [-\pi, -3\pi/4] \cup [-\pi/4, \pi/4] \cup [3\pi/4, \pi]$. Their minimal ranges can then be determined from using freedom in redefining the charged-lepton and neutrino fields. These do not remove the discrepancy in the physical parameter space for the mixing angle θ though: in case (a) it can only be restricted to $\theta \in [0, \pi/2]$, but to $\theta \in [0, \pi/4]$ in case (b). This is the result of the two discrete orderings of neutrino mass eigenstates that are possible under choice (b), namely $m_1 < m_2$ and $m_1 > m_2$.




Figure 3.4: Graphical representation of the two possible orderings of the neutrino mass eigenstates, together with their flavour compositions. The normal ordering is shown on the left, and the inverted one on the right. The widths of the coloured regions encode vacuum admixture probabilities of the various flavour eigenstates in each mass eigenstate, based on global knowledge of the three mixing angles. The effect of varying δ_{CP} is also illustrated. Diagram taken from Ref. [222].

The (more complex) discussion for three neutrinos can also be found in Ref. [220], which distinguishes three unambiguous mass-eigenstate definitions. One of these has become the de-facto standard throughout the literature:

$$m_1 < m_2 \text{ with } |\Delta m_{21}^2| < |\Delta m_{31}^2|, |\Delta m_{32}^2|, \quad (3.36)$$

which allows for two distinct orderings of the pair of m_1 and m_2 with respect to m_3 . This is referred to as the “neutrino mass ordering” (NMO) problem. Given the “normal” ordering (NO), ν_1 is the lightest state, ν_3 the heaviest, $m_1 < m_2 < m_3$, such that $\Delta m_{31}^2 > \Delta m_{32}^2 > 0$. Given the “inverted” ordering (IO), ν_3 is the lightest state, ν_2 the heaviest, $m_3 < m_1 < m_2$, such that $0 > \Delta m_{31}^2 > \Delta m_{32}^2$. By choosing all mixing angles to be positive, δ_{CP} is allowed to vary over the full 2π range. This yields the physical parameter intervals

$$\theta_{12}, \theta_{13}, \theta_{23} \in [0, \pi/2], \quad \delta_{\text{CP}} \in [-\pi, \pi], \quad \lambda_2, \lambda_3 \in [0, \pi]. \quad (3.37)$$

The two possible mass orderings are depicted schematically in Fig. 3.4. Based on current knowledge of the three mixing angles (see Sec. 5.2), the diagram shows the probabilities for encountering each of the three active flavour eigenstates in a given

mass eigenstate, using the fact that their incoherent sums yield unity:

$$|U_{ei}|^2 + |U_{\mu i}|^2 + |U_{\tau i}|^2 = 1. \quad (3.38)$$

No knowledge is assumed about δ_{CP} . The standard parameterisation (3.10) of the leptonic mixing matrix reveals that only $|U_{\mu 1,2}|^2$ and $|U_{\tau 1,2}|^2$ depend on δ_{CP} .

Indeterminability of NMO in vacuum

The physical parameter space (3.37) suffers from an intrinsic degeneracy which prevents the determination of the NMO in vacuum. When the PMNS matrix is rephased as [223]

$$U' \equiv PUP^\dagger, \text{ with } P \equiv \text{diag}(e^{i\delta_{\text{CP}}}, 1, 1) \quad (3.39)$$

or parameterised slightly differently than in the standard convention (3.10), by associating δ_{CP} with the 1–2 instead of the 1–3 rotation [224], the following combined set of parameter transformations reproduces the CPT operation and leaves the neutrino evolution invariant [225, 226, 224, 223]:⁶

$$\begin{aligned} \Delta m_{31}^2 &\rightarrow -\Delta m_{31}^2 + \Delta m_{21}^2 = -\Delta m_{32}^2, \\ \theta_{12} &\rightarrow \pi/2 - \theta_{12}, \\ \delta_{\text{CP}} &\rightarrow \pi - \delta_{\text{CP}}. \end{aligned} \quad (3.40)$$

The first line corresponds to a change of the NMO because $\Delta m_{21}^2 < |\Delta m_{31}^2|$ (the signs of Δm_{31}^2 and Δm_{32}^2 coincide within both orderings).

“Light” and “dark” sides of physical parameter space

Data from solar neutrino experiments originally used to be analysed in an effective two-neutrino framework with the relevant (“solar”) oscillation parameters Δm_{21}^2 and θ_{12} . For historical reasons, only the “light” side $0 \leq \theta_{12} \leq \pi/4$ used to be considered, while Δm_{21}^2 was assumed to be positive definite, corresponding to case (a) above. Matter effects in the Sun, however, introduce an observable dependence on the octant of θ_{12} . Accordingly, one has to additionally consider the “dark” side $\pi/4 < \theta_{12} \leq \pi/2$ in order to cover the full physically observable parameter space [227, 228, 229]. Claims about NMO sensitivity of oscillations in vacuum usually implicitly rely on

⁶These transformations correspond to the mapping $H_\nu \rightarrow -H_\nu^*$ up to an overall phase, which does not affect the oscillation probabilities.

the assumption that the true octant of θ_{12} is known.⁷

3.3 Neutrino propagation with standard matter effects

The case of neutrino oscillations in (non-uniform) matter was pioneered by Wolfenstein in his seminal papers [27, 231]. The standard approach (reviewed, e.g., in Ref. [213]) proceeds via the Schrödinger-like equation

$$i \frac{d}{dx} |\nu\rangle = H_\nu(x) |\nu\rangle, \quad (3.41)$$

which is assumed to describe the flavour evolution of ultrarelativistic left-handed neutrinos $|\nu\rangle \equiv (|\nu_e\rangle, |\nu_\mu\rangle, |\nu_\tau\rangle)^T$ in one dimension ($x \equiv |\mathbf{x}|$) in matter. The Hamiltonian

$$H_\nu(x) = H_{\text{vac}} + H_{\text{mat}}(x) \quad (3.42)$$

is the sum of the vacuum (kinetic) term H_{vac} and the position-dependent interaction (matter) term $H_{\text{mat}}(x)$. Substituting

$$H_{\text{vac}} = U \frac{M_{\text{diag}}^2}{2E_\nu} U^\dagger, \quad (3.43)$$

$$H_{\text{mat}}(x) = V(x), \quad (3.44)$$

where $M_{\text{diag}}^2 \equiv \text{diag}(m_1^2, m_2^2, m_3^2)$ and $V(x)$ is the matrix of neutrino potentials in the flavour basis, one arrives at the standard evolution equation

$$i \frac{d}{dx} |\nu\rangle = \left(U \frac{M_{\text{diag}}^2}{2E_\nu} U^\dagger + V(x) \right) |\nu\rangle. \quad (3.45)$$

For antineutrinos the full Hamiltonian reads [225]

$$H_{\bar{\nu}} = (H_{\text{vac}} - H_{\text{mat}})^*. \quad (3.46)$$

The Schrödinger-like macroscopic evolution equation (3.45) has served as a near-exclusive starting point of studies of neutrino oscillations in literature, despite its heuristic nature. QFT treatments—both for the case of vacuum oscillations [211], for which $V(x) = 0$, as well as for non-uniform matter [232, 213, 233]—have only

⁷For instance, compare Ref. [230] to Ref. [226].



Figure 3.5: All possible tree-level Feynman diagrams for (anti)neutrino coherent forward scattering on electrons or nucleons in Earth matter (CC on the left, NC on the right).

recently confirmed its validity for typical experimental settings.

3.3.1 Effective matter potential

In Earth matter, electrons, protons, and neutrons constitute an unpolarised background of fermions at rest. In the evolution equation (3.45), $V(x)$ is the matrix of neutrino potentials for coherent forward scattering. Coherence requires each background fermion to remain in the same quantum state [180]. Since forward scattering is characterised by a lack of momentum transfer to the scattering centres, the low-energy effective formalism of Sec. 2.4.4 applies. The absence of FCNCs in the SM results in the diagonal nature of the matrix potential in the flavour basis,

$$V(x) \equiv \text{diag} \left(V_e(x), V_\mu(x), V_\tau(x) \right). \quad (3.47)$$

Figure 3.5 shows the tree-level Feynman diagrams that have to be taken into account. Neutrinos and antineutrinos of all flavours interact with electrons and nucleons through the exchange of a Z^0 , but only electron neutrinos and antineutrinos exchange a W^\pm with electrons. Any deviations between the potentials for the different neutrino flavour components yield additional phase differences for the propagating states.

The effective Hamiltonian density for the NC interaction of a neutrino of flavour α with background fermions of type f , averaged over the background, $\langle \mathcal{H}_Z^{(\text{eff})} \rangle(x) = -\langle \mathcal{L}_Z^{(\text{eff})} \rangle(x)$, depends solely on the true *vector* component of the background fermion current, and not on its axial component [234]. Computing the integrated matrix elements of the averaged Hamiltonian densities yields the following NC potentials

for ultrarelativistic left-handed neutrinos [128]:

$$V_{\text{NC}}^f(x) = \sqrt{2}G_F N_f(x) g_V^f. \quad (3.48)$$

The averaged Hamiltonian density for the CC process $\nu_e e^- \rightarrow \nu_e e^-$, $\langle \mathcal{H}_W^{\text{(eff)}} \rangle$, can be cast into the required form containing separate neutrino and background-fermion bilinears by means of the Fierz transformation mentioned in Sec. 2.5.2. The result deviates from Eq. (3.48) only through the absence of a factor g_V^e :

$$V_{\text{CC}}(x) \equiv V_{\text{CC}}^e(x) = \sqrt{2}G_F N_e(x). \quad (3.49)$$

The effective matter potentials can be equivalently calculated assuming either quarks or nucleons as the background scattering centres [235]. First, for $Q^2 = 0$ only the (CC or NC) Dirac form factor F_1 contributes to the respective weak nucleon vector current, cf. Eq. (2.46). Second, $F_1(Q^2 = 0)$ simply reduces to the appropriately weighted sum of the constituent quarks' vector couplings.⁸

Continuing in the nucleon picture, using the medium's electrical neutrality and $g_V^e = -g_V^p$ leads to $V_{\text{NC}}^e(x) = -V_{\text{NC}}^p(x)$, while $V_{\text{NC}}^n(x) = -\sqrt{2}G_F N_n(x)/2$. As a result,

$$\begin{aligned} V(x) &= \text{diag}(V_{\text{CC}}(x), 0, 0) + \sum_{f=e,n,p} V_{\text{NC}}^f(x) \mathbb{1}_{3 \times 3} \\ &= \sqrt{2}G_F \left(N_e(x) \text{diag}(1, 0, 0) - \frac{1}{2} N_n(x) \mathbb{1}_{3 \times 3} \right). \end{aligned} \quad (3.50)$$

Each diagonal entry $V_\alpha(x)$ gives the potential energy of an ultrarelativistic left-handed neutrino of flavour $\alpha = e, \mu, \tau$ in regular matter. For right-handed antineutrinos the sign of $V_\alpha(x)$ has to be flipped, in agreement with Eq. (3.46). The potential energies for ultrarelativistic right-handed neutrinos and left-handed antineutrinos are negligibly small [128].

Since any term proportional to the identity matrix in the Hamiltonian generates an overall phase, by convention one considers

$$\begin{aligned} H_\nu(x) &\rightarrow H_\nu(x) - \frac{m_1^2}{2E_\nu} \mathbb{1}_{3 \times 3} - \sum_f V_{\text{NC}}^f(x) \mathbb{1}_{3 \times 3} \\ &= U \frac{\text{diag}(0, \Delta m_{21}^2, \Delta m_{31}^2)}{2E_\nu} U^\dagger + \text{diag}(V_{\text{CC}}(x), 0, 0). \end{aligned} \quad (3.51)$$

⁸ $F_1^{Zp}(0) = 1/2 - 2 \sin^2 \vartheta_w \equiv g_V^p = 2g_V^u + g_V^d$ and $F_1^{Zn}(0) = -1/2 \equiv g_V^n = g_V^u + 2g_V^d$ (cf. Secs. 2.4.2 and 2.5.3.1).

A handy expression for $V_{CC}(x)$ follows from introducing the relative local electron number density

$$Y_e(x) = \frac{N_e(x)}{N_p(x) + N_n(x)} , \quad (3.52)$$

where $N_p(x) = N_e(x)$ and $N_n(x)$ are the local proton and neutron number densities, respectively. Given $N_p(x) + N_n(x) \approx \rho(x)/m_N$, with the average nucleon mass $m_N = (m_p + m_n)/2$ and the local matter density $\rho(x)$, we find

$$N_e(x) \approx Y_e(x) \frac{\rho(x)}{m_N} , \quad (3.53)$$

such that

$$V_{CC}(x) \approx 0.76 \times 10^{-13} \text{ eV} \cdot Y_e(x) \left(\frac{\rho(x)}{\text{g/cm}^3} \right) \quad (3.54)$$

in units that are useful when applied to Earth matter.

Y_e is determined by the chemical and isotopic composition of the medium and is equivalent to the local average of the ratio between the atomic number and the mass number, $\langle Z/A \rangle$, which is ~ 0.5 for heavy elements. The Earth's average matter density of $\langle \rho \rangle \sim 5.5 \text{ g cm}^{-3}$ results in a characteristic matter potential $V_{CC} \sim \mathcal{O}(10^{-13} \text{ eV})$.

3.3.2 Standard oscillation formula in uniform matter

In matter the neutrino oscillation probability factorises from the probability for the overall production-propagation-detection process only for ultrarelativistic neutrinos, when their energy is large compared to the matter potential, and when they are produced and detected coherently [213]. Under these conditions a local mixing matrix in matter, $\tilde{U}(x)$, transforms between the flavour eigenstates $|\nu_\alpha\rangle$ and the local eigenstates $|\nu_k(x)\rangle$ of $H_\nu(x)$, such that [213]

$$|\nu_\alpha\rangle = \sum_k \tilde{U}_{\alpha k}^*(x) |\nu_k(x)\rangle . \quad (3.55)$$

The presence of \tilde{U}^* is due to the definition (3.22) and the fact the adjoint field creates particles [14].

In uniform matter the Hamiltonian H_ν remains constant, as do the local matter eigenstates and the mixing matrix. As a consequence of Eq. (3.55), the expression for the oscillation probability has the same structure as Eq. (3.30), after replacing

$U \rightarrow \tilde{U}$ and the mass-squared difference with an effective one, $\Delta m^2 \rightarrow \widetilde{\Delta m^2}$:

$$P_{\nu_\alpha \rightarrow \nu_\beta}(L, E_\nu) \simeq \sum_{j,k} \tilde{U}_{\alpha j}^* \tilde{U}_{\beta j} \tilde{U}_{\alpha k} \tilde{U}_{\beta k}^* e^{-i \frac{\widetilde{\Delta m^2}_{jk}}{2E_\nu} L} \quad (V(x) = \text{const.}) . \quad (3.56)$$

In non-uniform matter in general no analytic expression for the oscillation probability exists, and one has to resort to numerical methods for solving the neutrino evolution equation.

3.3.3 Discrete symmetries

Since ordinary matter violates both CP and CPT symmetry [236], CP- and T-violation effects on neutrino oscillations are not necessarily related—in contrast to the vacuum case. Whilst δ_{CP} still gives rise to an intrinsic (“genuine”) violation of CP (T) symmetry, the mere presence of matter provides an additional, extrinsic source of CP violation. In matter, an overall CP violation therefore generally occurs not just in appearance channels, but also in survival channels. Also, only neutrinos traversing an asymmetric matter density profile are subject to extrinsic T violation. This implies, for example, that the oscillations of atmospheric neutrinos—which encounter the symmetric Earth density profile [237]—are symmetric under the exchange of the produced and the detected neutrino flavour in the absence of intrinsic CP violation. A thorough discussion of these and further aspects of symmetries of neutrino oscillation probabilities under discrete transformations in various media can be found in Ref. [236].

Remarkably, matter with SM interactions lifts the degeneracy (3.40) due to CPT: given $H_{\text{mat}} \neq 0$, no parameter transformation allows for the mapping $H_\nu \rightarrow -H_\nu^*$. As a consequence, regular, CP-asymmetric matter with SM interactions generally allows distinguishing between the NO and the IO, which is exploited in Part IV. This also allows distinguishing between the “light” and the “dark” side of the neutrino parameter space. In fact, the detection of the matter effect in the Sun excludes the second octant of θ_{12} at a high level of confidence.⁹ The assumption of non-standard neutrino interactions, however, can “restore” CPT symmetry in matter (see Chapter 4). This impacts the neutrino physics reach of the analysis conducted in Part III.

⁹This problem is explicitly addressed, e.g., in the global analysis [238] from 2002 and the more recent one in Ref. [223].

3.3.4 Flavour transition characteristics in different types of matter

Medium properties strongly impact the nature of the neutrino flavour transitions that take place as a consequence of neutrino mixing and interfering propagating states. Discussed below are the most relevant mechanisms affecting the propagation of neutrinos in the Sun and the Earth. A two-neutrino scenario suffices to demonstrate the salient features that arise in the two media. It also exhibits an inherent quantitative usefulness because one can often describe the three-neutrino propagation by means of an effective two-neutrino system.

The “evolution matrix” $S(t, t_0)$ [239] relates the neutrino system at some time t to that at the initial time t_0 and satisfies the Schrödinger equation, just like the neutrino state vector. In the flavour basis, the elements $S(t, t_0)_{\alpha\beta}$ represent the transition amplitudes from which the probabilities $P_{\alpha\beta} = \left| S(t, t_0)_{\alpha\beta} \right|^2$ can be obtained. The unitary time evolution due to probability conservation, $S^\dagger S = \mathbb{1}$, allows parameterising the evolution matrix of a two-neutrino system as [239]

$$S = \begin{pmatrix} a & b \\ -b^* & a^* \end{pmatrix}, \quad |a|^2 + |b|^2 = 1. \quad (3.57)$$

The dominant effects relevant to neutrino propagation in the Sun and the Earth can be classified roughly according to the three categories below. The first two summaries are largely based on Ref. [239] (with exceptions noted) and the third in addition draws upon Refs. [240, 241, 242].

3.3.4.1 Constant density: resonance enhancement

Neutrino oscillations in matter of constant density proceed in a straightforward manner. Since the Hamiltonian H_ν is constant, its eigenstates $|\tilde{\nu}\rangle \equiv (|v_{1m}\rangle, |v_{2m}\rangle)^T$ are the “propagation eigenstates” with matter eigenvalues $H_{im} \equiv \frac{m_i^2}{2E_\nu}$. The evolution of the eigenstates is thus described by the diagonal matrix

$$\tilde{S}(L, 0) = \begin{pmatrix} e^{i\phi_m(L)} & 0 \\ 0 & e^{-i\phi_m(L)} \end{pmatrix}, \quad \text{with } \phi_m(L) = \frac{1}{2}(H_{2m} - H_{1m})L. \quad (3.58)$$

The mixing matrix \tilde{U} in the medium, which relates the matter eigenstates to the flavour eigenstates according to Eq. (3.55), follows from diagonalising H_ν :

$$\tilde{U}^\dagger H_\nu \tilde{U} = \text{diag}(H_{1m}, H_{2m}). \quad (3.59)$$

For the two-flavour system of $|\nu_e\rangle$ and $|\nu_a\rangle$, the effective matter potential is $V = \text{diag}(V_{CC}, 0)$. When the mixing matrix is parameterised with one mixing angle $\theta_m \in [0, \pi/2]$ (for $\Delta m^2 > 0$), i.e., Eq. (3.35) with the substitution $\theta \rightarrow \theta_m$, the latter is found from Eq. (3.59) as

$$\sin^2 2\theta_m = \frac{1}{R} \sin^2 2\theta, \quad \text{with } R \equiv \left(\cos 2\theta - \frac{2V_{CC}E_\nu}{\Delta m^2} \right)^2 + \sin^2 2\theta. \quad (3.60)$$

The “resonance factor” R determines whether the two-neutrino mixing in the medium is enhanced or suppressed with respect to that in vacuum. R also enters the expression for the oscillation frequency, which corresponds to the difference between the eigenvalues of the Hamiltonian, the “level splitting”

$$\omega_m \equiv H_{2m} - H_{1m} = \frac{\Delta m^2}{2E_\nu} \sqrt{R}, \quad (3.61)$$

from which the oscillation length in the medium follows as $L_m = \frac{2\pi}{\omega_m}$.

The “Mikheyev–Smirnov–Wolfenstein” (MSW) resonance condition

$$\Delta m^2 \cos 2\theta = 2V_{CC}E_\nu \Leftrightarrow R = \sin^2 2\theta \quad (3.62)$$

can be met (for neutrinos) only if $\theta < \pi/4$, because otherwise $\cos 2\theta \leq 0$. At resonance, the mixing in matter is maximal, $\sin^2 2\theta_m^{\text{res}} = 1$, independent of the value of θ in the first octant. Simultaneously, the level splitting is minimal, $\omega_m^{\text{res}} = \frac{\Delta m^2}{2E_\nu} \sin 2\theta$. For a given electron number density N_e , Eq. (3.62) implies that there is always a neutrino energy for which oscillations are resonantly enhanced, the “resonance energy”

$$E_R \equiv \frac{\Delta m^2 \cos 2\theta}{2V_{CC}}. \quad (3.63)$$

In the low-energy limit $E_\nu \ll E_R$, one has $\theta_m \rightarrow \theta$, whereas the high-energy limit $E_\nu \gg E_R$ suppresses the mixing, $\theta_m \rightarrow \pi/2$. Furthermore, in the regime of small vacuum mixing, with $\cos 2\theta \approx 1$, the MSW condition corresponds to the coincidence between the vacuum oscillation frequency of the two-neutrino system, $\omega_{\text{vac}} = \frac{\Delta m^2}{2E_\nu}$, and the medium’s “characteristic” frequency $\omega_0 \equiv V_{CC}^{-1}$.

Figure 3.6 shows the transition probability

$$P_{ea}(L, E_\nu) = \left| \left(\tilde{U} \tilde{S}(L, 0) \tilde{U}^\dagger \right)_{ea} \right|^2 = \sin^2 2\theta_m(E_\nu) \sin^2 \phi_m(L, E_\nu) \quad (3.64)$$

for different baselines L in matter of constant density, as a function of neutrino energy E_ν scaled by the ratio $2V_{CC}/\Delta m^2$. This yields the dimensionless quantity $\omega_0/\omega_{\text{vac}} = 2V_{CC}E_\nu/\Delta m^2$, which can be evaluated for any combination of values for N_e , Δm^2 , and E_ν .¹⁰ The oscillation amplitude $\sin^2 \theta_m$ is the envelope: only at the resonance energy E_R can a full flavour transformation take place in principle. As can be seen in the figure, whether $P_{ea}(L, E_R) = 1$ is indeed realised, however, depends on the oscillation “half-phase” (on the baseline):¹¹

$$\phi_m = \frac{1}{2} \omega_m L \stackrel{!}{=} \frac{1}{2} n \pi, \quad \text{with } n = 1, 3, 5, \dots \quad (3.65)$$

Expressing $\phi_m = \pi L/L_m = \pi m L_m^{\text{res}}/L_m$ makes evident that a full transformation takes place for a baseline L that is a half-integer multiple of the oscillation length L_m^{res} at resonance, i.e., for $m = 1/2, 3/2, \dots$ (phase condition). Thus, $P_{ea} = 1$ when the amplitude (MSW resonance) and the phase condition are satisfied simultaneously. The full width at half maximum of $\sin^2 2\theta_m$ is

$$\Delta \left(\frac{\omega_0}{\omega_{\text{vac}}} \right) = \Delta \left(\frac{2V_{CC}E_\nu}{\Delta m^2} \right) = 2 \sin 2\theta. \quad (3.66)$$

The smaller the vacuum mixing, the more the parameter space shrinks throughout which a conversion probability of more than 50% is possible. The dimensionless resonance width (3.66) can be rescaled straightforwardly into its commonly encountered dimensionful counterparts in E_ν or N_e .

3.3.4.2 Slow density changes: MSW effect

When neutrinos traverse a non-uniform medium, the local (instantaneous) eigenstates $|\tilde{\nu}\rangle$ evolve in time according to [239]

$$i \frac{d}{dt} |\tilde{\nu}\rangle = \begin{pmatrix} H_{1m} & -i\dot{\theta}_m \\ i\dot{\theta}_m & H_{2m} \end{pmatrix} |\tilde{\nu}\rangle. \quad (3.67)$$

¹⁰The continuation of the plot to $\omega_0/\omega_{\text{vac}} < 0$ would yield the antineutrino scenario. For two-flavour antineutrino oscillations a resonance is only possible for $\theta > \pi/4$.

¹¹Note the factor of two distinguishing the definitions (3.65) and (3.32).




Figure 3.6: 2ν oscillation probability (3.64) as a function of neutrino energy E_ν (scaled by the constant $2V_{CC}/\Delta m^2$ to construct the dimensionless frequency ratio $\omega_0/\omega_{\text{vac}}$), for three different values of the baseline L . The vacuum mixing angle, which determines the matter resonance position and width, is chosen as $\theta = 0.16$ rad, close to the current global best-fit value of θ_{13} . Figure inspired by Ref. [239].

Since the off-diagonal Hamiltonian elements $\propto \dot{\theta}_m = d\theta_m/dt$ induce transitions $\nu_{1m} \leftrightarrow \nu_{2m}$, the local eigenstates are not propagation eigenstates. However, if the rate of change of the mixing angle in matter is small compared to the level splitting, such that the “adiabaticity parameter”

$$\left| \frac{\dot{\theta}_m}{H_{2m} - H_{1m}} \right| \ll 1, \quad (3.68)$$

the local eigenstates evolve independently. In this adiabatic approximation, the Hamiltonian in the mass basis is diagonal just as in the case of propagation in vacuum or uniform matter, but its eigenvalues vary along the neutrino trajectory. Equivalently, the flavour compositions of the local eigenstates vary, due to the dependence of θ_m on the matter potential, as governed by Eq. (3.60). One obtains the evolution matrix (3.58), but the half-phase is now the integral of the level splitting along the

trajectory,

$$\phi_m(L) = \frac{1}{2} \int_0^L (H_{2m} - H_{1m}) dx . \quad (3.69)$$

The adiabatic flavour evolution reflects several aspects:

- the flavour compositions of the local eigenstates, $\langle v_\alpha | v_{im}(x) \rangle = \tilde{U}_{\alpha i}(x)$, follow the varying matter density;
- the propagating states are constant admixtures of the local eigenstates (due to the lack of transitions $\nu_{1m} \leftrightarrow \nu_{2m}$), fixed by the value of the mixing angle $\theta_m(0)$ at the neutrino source;
- and there is a phase difference between the local eigenstates, $\phi_m(x)$, whose rate of increase is determined by the level splitting $H_{2m}(x) - H_{1m}(x)$.

A simple, albeit physically relevant, example [183] is that of an electron neutrino produced where the electron number density is high, $N_e(0) \gg N_e^{\text{res}}$, i.e., without mixing ($\theta_m = \pi/2$). This means that $|v_e\rangle \simeq |v_{2m}\rangle$ according to Eq. (3.35). If ν_{2m} propagates adiabatically, its flavour composition changes as dictated by the evolution of θ_m , and ν_{1m} is never produced in the first place. The electron neutrino survival probability is given by

$$P_{ee}(L) = \left| \left(\tilde{U}(L) \tilde{S}(L,0) \tilde{U}(0)^\dagger \right)_{ee} \right|^2 = |\langle v_e | v_{2m}(L) \rangle|^2 . \quad (3.70)$$

Assuming the density decreases from $N_e(0) \gg N_e^{\text{res}} \rightarrow N_e(L) = 0$, the local eigenstate after propagation over the distance L is given by that in vacuum, $|v_{2m}(L)\rangle = |v_2\rangle$. As a consequence, the survival probability is determined by the vacuum mixing alone, $P_{ee}(L) = \sin^2 \theta$. This kind of non-oscillatory flavour conversion is a realisation of the ‘‘MSW effect’’. It is sketched in Fig. 3.7, which depicts the propagating neutrino state as a wave packet, together with its two flavour admixtures as these evolve from high to low density (vacuum).

If the neutrino production instead happens at $N_e(0) > N_e^{\text{res}}$, $|v_{1m}\rangle$ is present in $|v_e\rangle$ with a non-negligible admixture. There is interference between $|v_{2m}\rangle$ and $|v_{1m}\rangle$ and a flavour transformation $\nu_e \rightarrow \nu_\alpha$ arises through the coexistence of adiabatic conversion and oscillation.

Finally, neutrino production with $N_e(0) < N_e^{\text{res}}$ results in matter effects which constitute minor modifications to the vacuum oscillation scenario.

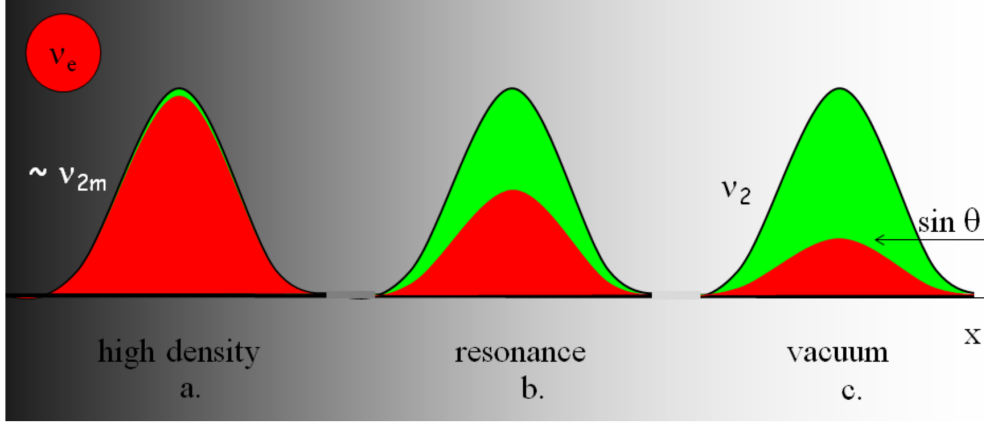


Figure 3.7: Two-neutrino example of adiabatic conversion in the case of an electron neutrino produced at $N_e \gg N_e^{\text{res}}$ (a), such that it nearly coincides with the effective mass eigenstate $|\nu_{2m}\rangle$. The flavour composition of the propagating state follows the flavour composition of $|\nu_{2m}\rangle$, which at N_e^{res} (b) is composed of equal parts $|\nu_e\rangle$ and $|\nu_a\rangle$. For $N_e = 0$ (c), the vacuum mixing angle θ fixes the flavour admixtures of the state that has now become $|\nu_2\rangle$. Figure taken from Ref. [183].

3.3.4.3 Periodic density changes: parametric enhancement

The propagation of neutrinos in a medium with a periodically changing matter potential provides the only possibility for strong, $P_{\alpha\beta} \sim 1$, flavour transitions in the absence of large mixing in both vacuum and the medium itself. Strong transitions occur if the period L_T of the variation of the matter potential equals an integer multiple of the matter oscillation length, $L_T/\langle L_m \rangle = k$, or [239]

$$\int_0^{L_T} \frac{dx}{L_m} = k, \quad \text{with } k = 1, 2, 3, \dots \quad (3.71)$$

A density profile with alternating electron number densities N_1 and N_2 and a period $L_T = L_1 + L_2$, where L_1 and L_2 are the widths of the two layers, serves as an example with applicability to the Earth. Here, the evolution matrix after n periods of density modulation is given by the n th power of the product of the pair of evolution matrices describing propagation in the two layers,

$$S_n \equiv (S_T)^n \equiv (S_2 S_1)^n. \quad (3.72)$$

In layer k , $\phi_k = \frac{1}{2}\omega_{mk}L_k = \frac{\Delta m^2}{4E_\nu}\sqrt{R_k}L_k$ is the half-phase acquired by the local eigenstates, where ω_{mk} , R_k , and L_k are the layer's level splitting, resonance factor, and thickness, respectively. A convenient expression for the evolution matrix in the layer

is [239]

$$S_k = \mathbb{1}_{2 \times 2} \cos \phi_k - \mathbb{1}_{2 \times 2} (\boldsymbol{\sigma} \cdot \mathbf{n}_k) i \sin \phi_k, \quad (3.73)$$

with the vector of Pauli matrices $\boldsymbol{\sigma} \equiv (\sigma_1, \sigma_2, \sigma_3)$ and $\mathbf{n}_k \equiv (\sin 2\theta_{mk}, 0, -\cos 2\theta_{mk})$. Substituting the representation (3.73) into Eq. (3.72) yields the evolution matrix for n periods,

$$S_n = \mathbb{1}_{2 \times 2} \cos n\Phi - \mathbb{1}_{2 \times 2} (\boldsymbol{\sigma} \cdot \mathbf{X}/|\mathbf{X}|) i \sin n\Phi, \quad (3.74)$$

with $\mathbf{X} \equiv \sin \phi_1 \cos \phi_2 \mathbf{n}_1 + \sin \phi_2 \cos \phi_1 \mathbf{n}_2 - \sin \phi_1 \sin \phi_2 \sin(2\theta_{m1} - 2\theta_{m2}) \mathbf{e}_y$ and $\Phi \equiv \arcsin |\mathbf{X}|$. The two-neutrino transition probability after n full periods of density modulation is

$$P_{ea}(nL_T) = \left| (S_n)_{ea} \right|^2 = \frac{X_1^2 + X_2^2}{|\mathbf{X}|^2} \sin^2 n\Phi. \quad (3.75)$$

In case the neutrino system traverses a non-integer number of periods, $L = nL_T + L'$ (with $L' < L_T$), the evolution matrix follows as $S(L) = S(L', 0) S_n$.

The parametric resonance conditions are satisfied when the factor in front of the sine-squared function in Eq. (3.75) is unity, i.e., when $X_3 = -\sin \phi_1 \cos \phi_2 \cos 2\theta_{m1} - \sin \phi_2 \cos \phi_1 \cos 2\theta_{m2} = 0$. This can be achieved, for example, when each term in X_3 vanishes separately (see Ref. [242] and references therein):

$$\phi_1 = k \frac{\pi}{2}, \quad \phi_2 = k' \frac{\pi}{2}, \quad \text{with } k, k' = 1, 3, 5, \dots \quad (3.76)$$

Evaluating Eq. (3.75) under these assumptions yields a two-neutrino transition probability with a maximal amplitude:

$$P_{ea}^{\text{res}}(nL_T) \equiv \left| S_n \left(\phi_i = k^{(i)} \frac{\pi}{2} \right) \right|_{ea}^2 = \sin^2 [2n(\theta_{m1} - \theta_{m2})]. \quad (3.77)$$

Whether this (resonant) probability is subject to an enhancement with respect to the maximally achievable transition probabilities in either of the two layers, $P_{ea}^{\text{res}}(nL_T) \stackrel{?}{>} \sin^2 2\theta_{m1, m2}$, depends on the relation between $N_{1,2}$ and the MSW resonance density N_e^{res} and on the number n of periods traversed.

If $N_1, N_2 < N_e^{\text{res}}$, the effective mixing angles are in the first octant ($\theta_{m1}, \theta_{m2} < \pi/4$), whereas $N_1, N_2 > N_e^{\text{res}}$ yields effective mixing angles in the second octant ($\theta_{m1}, \theta_{m2} > \pi/4$). Either way, after evolution over a single period ($n = 1$), the resulting transition probability (3.75) is smaller than one of the individual probability maxima, either $\sin^2 2\theta_{m1}$ or $\sin^2 2\theta_{m2}$. For sufficiently large values of n , however, it will be larger than both, $P_{ea}(nL_T) > \sin^2 2\theta_{m1, m2}$. A more rapid increase of the transition probability ensues if $N_1 < N_e^{\text{res}} < N_2$, as this relation between densities is equivalent to the




Figure 3.8: Two-neutrino example of parametric transition probability enhancement as a function of the number n of periods traversed by the neutrino system, with $\phi_1 = \phi_2 = \pi/2$ and with realistic assumptions for θ and $\omega_0/\omega_{\text{vac}}$ (cf. Fig. 3.6). In the upper panel, the solid line is the $\nu_e \rightarrow \nu_a$ transition probability for any (integer and non-integer) number of periods, while the red markers represent the resonant probability after n full periods, given by Eq. (3.77). The lower panel depicts the assumed “castle wall” density profile, with $L_1/L_2 \approx 0.5$.

mixing configuration $\theta_{m1} < \pi/4 < \theta_{m2}$. Here, the evolution over a small number of periods already leads to an enhancement of the transition probability.

Figure 3.8 shows how the two-flavour transition probability builds up and decreases again as a function of $n \leq 10$ for an example configuration of the “castle wall” density profile introduced above. In both layers, it is assumed that $V_{\text{CC}}/V_{\text{CC}}^{\text{res}} < 1$, or equivalently $\omega_0/\omega_{\text{vac}} = 2V_{\text{CC}}E_\nu/\Delta m^2 < \cos 2\theta$. For a given E_ν , Δm^2 , and vacuum mixing angle θ , the density configuration is $N_1 < N_2 < N_e^{\text{res}}$. The distance travelled through each layer is taken to be half the oscillation length in the layer, $L_{1(2)} = L_{m1(m2)}/2$, so that the phase conditions (3.76) hold. Hence, the transition probability after each full period follows from Eq. (3.77), whose maximum is encountered at $n = 4$. To the left of this maximum there is an oscillatory increase of the transition probability over the course of any given period. To the right, a non-monotonic decrease is observed up to $n = 8$. This behaviour is a consequence of the density configuration underlying the depicted example.

4

Non-standard interactions of propagating neutrinos

The discovery of the weak neutral current prompted investigations into the possibility of FCNCs such as those by Wolfenstein in 1975 [243]. He realised that neutrino mass was not a requirement for the occurrence of neutrino oscillations in matter, because the necessary phase differences between the propagating states could also be induced solely by flavour-changing (or “flavour-violating”) effective potentials experienced by the neutrino flavour states. Assuming a system of two massless neutrinos—both with and without including the flavour-diagonal standard potential (3.50)—Wolfenstein estimated the refraction length in Earth matter (here identical to the oscillation length) to approximately correspond to the Earth’s radius [27]. This would render the effects of FCNCs observable in long-baseline oscillation experiments.

FCNCs only constitute a subset of all the interactions nowadays collectively referred to as “non-standard neutrino interactions” (NSI, instead of the less frequently used acronym “NSIs”). These are usually assumed to include both new flavour-violating and flavour-diagonal NC neutrino interactions, in addition to new CC interactions.

4.1 Model-independent effective formalism

In the effective theory approach of Eq. (3.23), neutrino interactions beyond the SM may arise at energy dimension $d \geq 6$ [244]. Below the scale of EWSB, the effective Lagrangian has to respect the residual $SU(3)_C \times U(1)_Q$ symmetry. Above, it has to

be invariant under G_{SM} , as the high-energy theory needs to encompass the latter. In agreement with most literature, the acronym “NSI” is reserved for $d = 6$ four-fermion Lorentz-vector operators resulting in new neutrino interactions, while the parent set of all possible new Lorentz-invariant neutrino interactions are referred to as “general neutrino interactions” (GNI) [245]. Generically, one expects the largest low-energy effects from operators of dimension $d = 6$, because higher-dimension operators are suppressed by higher powers of the new-physics scale Λ .

Considering $d = 6$, a complete overview of the CC and NC effective GNI Lagrangians below the EWSB scale can be found in Ref. [245]. The possible operators are obtained from the five bilinear currents with definite Lorentz transformation behaviour in Sec. 2.3. Allowing for chiral fermions increases the number of possible operators to ten. NSI correspond to the two configurations with Lorentz structures $(V - A)(V \pm A)$. These only involve left-chiral neutrinos—together with either left- or right-chiral charged leptons or quarks—and are therefore realisable using the SM field content. All other operators involve right-chiral neutrino fields [245].

This work investigates the effects of NC NSI on neutrino propagation. Incoherent scattering processes can be safely neglected for neutrino energies of up to $\mathcal{O}(1 \text{ TeV})$ [128].¹ Coherence requires the fermion on which the neutrino scatters to retain its flavour. Correspondingly, the relevant NC NSI Lagrangian to which we restrict ourselves reads [27, 246, 247, 248, 249]

$$\mathcal{L}_{\text{NSI}}^{\text{NC}} = -2\sqrt{2}G_F\epsilon_{\alpha\beta}^{fC} (\bar{\nu}_\alpha\gamma_\mu P_L\nu_\beta) (\bar{f}\gamma^\mu P_C f) , \quad (4.1)$$

with $f = e, u, d$.

The CC NSI Lagrangian reads (see Ref. [250] and references therein)

$$\mathcal{L}_{\text{NSI}}^{\text{CC}} = -2\sqrt{2}G_F\epsilon_{\alpha\beta}^{ff'C} (\bar{\ell}_\alpha\gamma_\mu P_L\nu_\beta) (\bar{f}'\gamma^\mu P_C f) , \quad (4.2)$$

with $f \neq f' = u, d$. CC NSI are relevant for processes which lead to the production or detection of neutrinos. In this work, however, their effects will not be considered: new-physics models leading to CC NSI of observable strengths are hard to reconcile with existing data, with upper limits on the effective coupling strengths that are roughly one order of magnitude stronger than those for NC NSI [251] (see also Sec. 4.5).

¹Assuming the NSI cross sections are not larger than their weak CC and NC counterparts.

4.2 Neutral-current NSI

In addition to the couplings of neutrinos to chiral fermions, $\epsilon_{\alpha\beta}^{fL}$, $\epsilon_{\alpha\beta}^{fR}$, which appear in Eq. (4.1), it is useful to introduce the vector coupling $\epsilon_{\alpha\beta}^{fV}$ and the axial-vector coupling $\epsilon_{\alpha\beta}^{fA}$ as well. These arise when Eq. (4.1) is split into a term containing a background-fermion vector current and another term containing a background-fermion axial-vector current [252],

$$\mathcal{L}_{\text{NSI}}^{\text{NC}} = -\sqrt{2}G_F\epsilon_{\alpha\beta}^{fV}(\bar{\nu}_\alpha\gamma^\mu P_L\nu_\beta)(\bar{f}\gamma_\mu f) - \sqrt{2}G_F\epsilon_{\alpha\beta}^{fA}(\bar{\nu}_\alpha\gamma^\mu P_L\nu_\beta)(\bar{f}\gamma_\mu\gamma^5 f), \quad (4.3)$$

where we have defined

$$\epsilon_{\alpha\beta}^{fV} \equiv \epsilon_{\alpha\beta}^{fL} + \epsilon_{\alpha\beta}^{fR}, \quad (4.4)$$

$$\epsilon_{\alpha\beta}^{fA} \equiv \epsilon_{\alpha\beta}^{fL} - \epsilon_{\alpha\beta}^{fR}. \quad (4.5)$$

NSI of neutrinos propagating in terrestrial or solar matter can be described by a matrix of effective potentials $V_{\text{NSI}}^{\oplus(\ominus)}$, in analogy to the SM scenario of Sec. 3.3.1. In particular, as demonstrated there, it is only the true vector component of the fermion background current in the Lagrangian (4.3) that is relevant, i.e., the Lorentz structure $(V - A)(V)$. The matrix of effective potentials is the linear combination of the potentials for coherent forward scattering on electrons, up quarks, and down quarks [250]:

$$V_{\text{NSI},\alpha\beta}^{\oplus(\ominus)}(x) = \sqrt{2}G_F(\epsilon_{\alpha\beta}^{eV}N_e(x) + \epsilon_{\alpha\beta}^{uV}N_u(x) + \epsilon_{\alpha\beta}^{dV}N_d(x)). \quad (4.6)$$

In contrast to the standard potential, $V_{\text{NSI}}^{\oplus(\ominus)}$ need not be diagonal: its off-diagonal components are potentials for flavour-violating interactions that do not modify the fermion background.

Normalising $V_{\text{NSI}}^{\oplus(\ominus)}$ to the SM potential given in Eq. (3.49) yields the following compact expression for the neutrino potential in the presence of NSI:

$$H_{\text{mat}}(x) = V_{\text{CC}}(x) \begin{pmatrix} 1 + \epsilon_{ee}^V(x) & \epsilon_{e\mu}^V(x) & \epsilon_{e\tau}^V(x) \\ \epsilon_{\mu e}^V(x) & \epsilon_{\mu\mu}^V(x) & \epsilon_{\mu\tau}^V(x) \\ \epsilon_{\tau e}^V(x) & \epsilon_{\tau\mu}^V(x) & \epsilon_{\tau\tau}^V(x) \end{pmatrix}, \quad (4.7)$$

where the position-dependent effective NSI couplings [223]

$$\epsilon_{\alpha\beta}^V(x) \equiv \frac{V_{\text{NSI},\alpha\beta}^{\oplus(\odot)}(x)}{V_{\text{CC}}(x)} = \epsilon_{\alpha\beta}^{eV} + \frac{N_u(x)}{N_e(x)} \epsilon_{\alpha\beta}^{uV} + \frac{N_d(x)}{N_e(x)} \epsilon_{\alpha\beta}^{dV} \quad (4.8)$$

$$= \epsilon_{\alpha\beta}^{eV} + \epsilon_{\alpha\beta}^{pV} + \frac{N_n(x)}{N_e(x)} \epsilon_{\alpha\beta}^{nV} \quad (4.9)$$

have been introduced. The second equality follows from the equivalence of the quark and nucleon descriptions already exploited in Sec. 3.3.1:

$$\epsilon_{\alpha\beta}^{pV} \equiv 2\epsilon_{\alpha\beta}^{uV} + \epsilon_{\alpha\beta}^{dV}, \quad (4.10)$$

$$\epsilon_{\alpha\beta}^{nV} \equiv \epsilon_{\alpha\beta}^{uV} + 2\epsilon_{\alpha\beta}^{dV}. \quad (4.11)$$

The number densities of up and down quarks expressed in terms of the proton and neutron number densities are given by $N_u(x) = 2N_p(x) + N_n(x)$ respectively $N_d(x) = N_p(x) + 2N_n(x)$.

While all of the above without distinction applies to neutrinos propagating in the Earth or the Sun, the effective couplings (4.8) of the former can however be considered constant. This is due to the fact that the ratio $Y_n(x) \equiv N_n(x)/N_e(x)$ within the Earth is constant to approximately $\pm 5\%$, with a mean value of $Y_n^\oplus \equiv \langle Y_n(x) \rangle \approx 1.051$ [223]. Thus, to a good approximation one can view a measurement of neutrino oscillations in the Earth as a means of constraining the constant effective couplings

$$\epsilon_{\alpha\beta}^\oplus \approx \epsilon_{\alpha\beta}^{eV} + \epsilon_{\alpha\beta}^{pV} + Y_n^\oplus \epsilon_{\alpha\beta}^{nV}. \quad (4.12)$$

Modifications of the neutrino detection cross sections arising from NC NSI are not accounted for in this work. Their model-independent inclusion would be computationally intractable because detection cross sections are not determined uniquely by the vector couplings considered above. This becomes evident from the presence of the axial-vector fermion currents in the NC NSI Lagrangian (4.3). A more extensive discussion of these and related aspects can be found in Ref. [253].

4.3 Parameterisations

Standard parameterisation

Since neutrino oscillations are not sensitive to an overall phase shift, it is possible to eliminate one of the diagonal entries of H_{mat} without affecting the oscillation physics.

It is common to redefine

$$H_{\text{mat}}(x) \rightarrow H_{\text{mat}}(x) - V_{\text{CC}}(x) \epsilon_{\mu\mu}^{\oplus} \mathbb{1}_{3 \times 3}, \quad (4.13)$$

though any other choice is just as valid. Hermiticity ($H_{\text{mat}}^{\dagger} = H_{\text{mat}}$) gives rise to the parameter constraints

$$\epsilon_{ee,\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} \in \mathbb{R}, \quad (4.14)$$

$$\epsilon_{\beta\alpha}^{\oplus} = (\epsilon_{\alpha\beta}^{\oplus})^* \quad (\alpha \neq \beta). \quad (4.15)$$

These imply that the observable neutrino potential is fully described by eight real parameters, one real part for each of the two non-universality (NU) strengths $\epsilon_{ee,\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$, and one absolute value and complex phase for each of the three flavour-violating (FV) couplings $\epsilon_{e\mu}^{\oplus}$, $\epsilon_{e\tau}^{\oplus}$, and $\epsilon_{\mu\tau}^{\oplus}$:

$$\epsilon_{ee,\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} = \text{Re} [\epsilon_{ee,\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}], \quad (4.16)$$

$$\epsilon_{\alpha\beta}^{\oplus} \equiv |\epsilon_{\alpha\beta}^{\oplus}| e^{i\delta_{\alpha\beta}} \quad (\alpha \neq \beta). \quad (4.17)$$

Hence,

$$H_{\text{mat}}(x) = V_{\text{CC}}(x) \begin{pmatrix} 1 + \epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} & \epsilon_{e\mu}^{\oplus} & \epsilon_{e\tau}^{\oplus} \\ \epsilon_{e\mu}^{\oplus*} & 0 & \epsilon_{\mu\tau}^{\oplus} \\ \epsilon_{e\tau}^{\oplus*} & \epsilon_{\mu\tau}^{\oplus*} & \epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} \end{pmatrix}. \quad (4.18)$$

H_{vac} -like parameterisation

Not all eight parameters in the standard parameterisation above are intrinsic properties of NSI: if the kinetic Hamiltonian was not present (or not observable, as in the case of massless or degenerate-mass neutrinos), only one combination of the three off-diagonal couplings' complex phases would be physical [254]. This is demonstrated by an alternative parameterisation whose structure resembles that of the vacuum Hamiltonian (3.43) [255, 225, 254, 223]:

$$H_{\text{mat}}(x) = Q_{\text{rel}} U_{\text{mat}} D_{\text{mat}}(x) U_{\text{mat}}^{\dagger} Q_{\text{rel}}^{\dagger}. \quad (4.19)$$

This parameterisation, referred to as ‘‘generalised matter potential’’ (GMP) henceforth, introduces three new matrices with a total of eight real parameters as follows.

Matter-potential eigenvalues $D_{\text{mat}}(x) = V_{\text{CC}}(x) \text{diag}(\epsilon_{\oplus}, \epsilon'_{\oplus}, 0)$ gives the local eigenvalues of the GMP, $V_{\text{CC}}(x) \epsilon_{\oplus}$ and $V_{\text{CC}}(x) \epsilon'_{\oplus}$, in analogy to the vacuum Hamiltonian, which has the eigenvalues $\frac{\Delta m_{21}^2}{2E_\nu}$ and $\frac{\Delta m_{31}^2}{2E_\nu}$.

Matter-potential rotation $U_{\text{mat}} = R_{12}(\varphi_{12}) R_{13}(\varphi_{13}) \tilde{R}_{23}(\varphi_{23}, \delta_{\text{NS}})$ rotates the matter potential into the off-diagonal entries of H_{mat} via the two real Euler rotations R_{12} and R_{13} , by φ_{12} in the 1–2 plane and by φ_{13} in the 1–3 plane, and one complex Euler rotation \tilde{R}_{23} , by φ_{23} and δ_{NS} . The association of the CP-violating NSI phase δ_{NS} with the 2–3 rotation is not a unique choice. U_{mat} plays a similar role in matter as the PMNS matrix U does in vacuum.

Relative matter-vacuum rephasing $Q_{\text{rel}} = \text{diag}(e^{i\alpha_1}, e^{i\alpha_2}, e^{-i(\alpha_1+\alpha_2)})$ rephases the matter part of the Hamiltonian with respect to its vacuum part by means of the two phases α_1, α_2 . It is a matter of convention whether these are assigned to H_{mat} or H_{vac} .

Out of the three phases α_1, α_2 , and δ_{NS} , only δ_{NS} is an intrinsic feature of NSI. With respect to the standard parameterisation, this implies that the three effective phases $\delta_{e\mu}, \delta_{e\tau}$, and $\delta_{\mu\tau}$ associated with the FV coupling strengths combine into one “genuine” complex phase.

Section 5.2.2 presents a convenient reduction of the GMP parameter space when the parameterisation is applied to data from atmospheric neutrino oscillation experiments, as done in Part III of this work. It also provides the resulting relations between the GMP parameterisation and the standard parameterisation and discusses the former’s minimal physical parameter space.

4.4 Generalised mass-ordering degeneracy

In the presence of NSI in the standard parameterisation, CPT is realised by the set of transformations (3.40) of the vacuum Hamiltonian in conjunction with

$$\begin{aligned} \left[\epsilon_{ee}^V(x) - \epsilon_{\mu\mu}^V(x) \right] &\rightarrow - \left[\epsilon_{ee}^V(x) - \epsilon_{\mu\mu}^V(x) \right] - 2, \\ \left[\epsilon_{\tau\tau}^V(x) - \epsilon_{\mu\mu}^V(x) \right] &\rightarrow - \left[\epsilon_{\tau\tau}^V(x) - \epsilon_{\mu\mu}^V(x) \right], \\ \epsilon_{\alpha\beta}^V(x) &\rightarrow -\epsilon_{\alpha\beta}^{V*}(x) \quad (\alpha \neq \beta). \end{aligned} \tag{4.20}$$

The equivalence between these parameter transformations and CPT is referred to as the “generalised mass-ordering degeneracy” (GMOD) [224] because it prevents

the determination of the NMO if the non-standard couplings to neutrons vanish ($\epsilon_{\alpha\beta}^{nV} = 0$) or if the relative neutron number density $Y_n(x) = \text{const.}$ along the neutrino trajectory. Crucially, the latter applies to the Hamiltonian (4.18) for Earth matter.

4.5 New-physics origin

Many well-motivated new-physics scenarios—often invoked to explain the small neutrino masses—naturally give rise to NSI in neutrino propagation. However, building models which lead to potentially observable NSI without violating constraints from low-energy processes—such as neutrino oscillations, charged-lepton decay, or β -decay—is challenging. The main characteristics of different types of NSI models frequently encountered in the literature are briefly sketched below. A comprehensive overview with many relevant references can be found in Ref. [31].

As an illustrative example which gives some insight into rather generic constraints on NSI coupling strengths, consider the existence of the tree-level process $\nu_\alpha + f \rightarrow \nu_\beta + f$, in which the initial-state neutrino ν_α interacts with a charged fundamental fermion f by exchanging a (scalar or vector) boson X of mass m_X . On the one hand, in the limit of an effective contact interaction as described in Sec. 2.4.4 ($|q^2| \ll m_X^2$) the amplitude of the process is proportional to $g_{\alpha f} g_{\beta f} / m_X^2$, where $g_{\alpha f}$ and $g_{\beta f}$ represent the couplings associated with the two vertices. The effective coupling strength relative to Fermi's coupling constant then reads [256]

$$\epsilon_{\alpha\beta}^f \propto \frac{g_{\alpha f} g_{\beta f}}{m_X^2 G_F}. \quad (4.21)$$

As a consequence, the new couplings are comparable in strength to the weak interaction if $g^2/m_X^2 \sim \mathcal{O}(G_F)$. If $|q^2| \gtrsim m_X^2$, on the other hand, the NSI scattering amplitude is suppressed by m_X^2/q^2 [257] with respect to the contact-interaction limit. This generic behaviour of the effective coupling strengths has to be kept in mind when comparing the NSI reach of different experiments. A fundamental assumption underlying the expression (4.6) for the potential due to NSI is that the mediator is much heavier than the inverse extent of the matter distribution traversed by the neutrinos, i.e., $m_X \gg R_\oplus^{-1} \sim 3 \times 10^{-14} \text{ eV}$ in the case of the Earth (see for example Refs. [258, 259] and references therein). For $m_X \lesssim R_\oplus^{-1}$, the local value of the potential would depend not simply on the local background-fermion number densities but on their spatial distributions (long-range potential).

Tree-level NSI

Among the models considered as viable candidates to bring about NSI, there are approaches to extend the SM with new scalar bosons which generate new tree-level four-fermion operators of energy dimension $d = 6$ that transform as Lorentz vectors after a Fierz transformation, for example in scalar- $SU(2)_L$ -singlet models, two-Higgs-doublet models, or the type-II ($SU(2)_L$ -triplet) seesaw model (see Ref. [260] and references therein, as well as Ref. [261] for a comprehensive review in the context of radiative neutrino mass generation). All these share the need for the new effective scalar operators to be of the form $\mathcal{O}_{\text{scalar}} \propto \bar{\Psi} \nu_L \bar{\nu}_L \Psi$ (without family indices), where Ψ is a (not necessarily chiral) charged fermion field. As a result, the scalar boson must be electrically charged and in addition carry colour charge (leptoquark) in case NSI with quarks are to exist. Only under these conditions may vector NSI couplings as given in Eq. (4.4) arise [260]. However, its electric charge demands the scalar boson to be heavy—preventing production in existing experiments—owing to its fixed gauge coupling to the photon [262, 260]. This model-building problem similarly applies to CC NSI, for which new physics above the TeV scale is therefore usually required [262].

Notably, model-independent studies of the maximally allowed strengths of NC NSI produced at tree level by ($d = 6$ or $d = 8$) gauge-invariant operators above the EWSB scale have been conducted in Refs. [263, 244]. Since the high-energy theory must encompass G_{SM} , processes involving four charged leptons are generically expected to occur also. For NSI to be large, these interactions should however be strongly suppressed. At $d = 6$, a systematic operator analysis reveals [244] that there is just one² possibility of generating such NSI with a *single* tree-level mediator—a heavy scalar $SU(2)_L$ singlet [264, 265].³ For regular matter, this simple model only permits non-zero couplings in the μ - τ sector, namely $\epsilon_{\mu\mu}^{eL}$, $\epsilon_{\mu\tau}^{eL}$, and $\epsilon_{\tau\tau}^{eL}$ [263]. More generally, as summarised in Ref. [270], the magnitudes of all NSI coupling strengths $|\epsilon_{\alpha\beta}^{\oplus}|$ (from gauge-invariant operators and tree-level mediators) are limited to $\mathcal{O}(10^{-2})$ or below.

Loop-induced NSI

In contrast to the above “Fierz-transformed” NSI, loop-induced NSI [260] are the result of neutrino Yukawa interactions with new—electrically neutral or charged—

²An additional possibility arises when the existence of right-chiral neutrinos is allowed [263].

³The scalar singlet occurs in a class of models referred to as “Zee models” [266, 267, 268], which provide economical mechanisms for the radiative generation of neutrino mass [269].




Figure 4.1: “Triangle” (left) and “box” (right) diagrams which induce NC vector NSI via loop corrections to the tree-level Z^0 exchange and via pure Yukawa interactions, respectively. On the left, the two occurrences of ψ_{int} represent (SM or new, possibly different) fermions (either both electrically charged or neutral). In each diagram, ϕ is a new scalar boson. Figure taken from Ref. [260].

scalar (or vector) bosons. Two possible diagrams generating NSI in this way are shown in Fig. 4.1. In the “triangle” diagram on the left, external neutrinos and internal fermions with couplings to the Z^0 boson are converted into each other through Yukawa interactions with a new scalar boson (ϕ). In the “box” diagram on the right, only pure Yukawa interactions are involved. Under the assumption that the mass of the scalar boson exceeds the fermion masses, both diagrams generate effective four-fermion operators $\mathcal{O}_{\text{scalar}}^{\text{tri,box}} \propto \bar{\psi}\gamma^\mu\psi\bar{\nu}_\alpha\gamma_\mu P_L\nu_\beta$. These can be combined into effective Lagrangian terms for loop-induced NSI with (chiral) charged fermions [260],

$$\mathcal{L}_{\text{scalar}}^{\text{tri+box},fC} = \frac{G_F}{\sqrt{2}} \left(\epsilon_{\alpha\beta}^{\text{tri},fC} + \epsilon_{\alpha\beta}^{\text{box},fC} \right) \bar{\psi}\gamma^\mu\psi\bar{\nu}_\alpha\gamma_\mu P_L\nu_\beta, \quad (4.22)$$

where $\psi = f_C = e_L, e_R, u_L, u_R, d_L, d_R$. Vector NSI can be constructed from these via the summation (4.4).

Loop-induced NSI may occur with both charged leptons and quarks serving as scattering partners, without the need for the new scalar boson to be electrically charged or colour-charged. This particular feature of loop-induced NSI in turn is conducive to their reconcilability with existing experimental limits. The detailed arguments and calculations, together with general and model-dependent formulae and constraints for $\epsilon_{\alpha\beta}^{\text{tri,box}}$, can be found in Ref. [260].

New gauge symmetries

The third approach to generating NSI considered here consists of extending the SM gauge sector with additional gauge symmetries accompanied by new (vector) gauge bosons, such as new Abelian $U(1)'$ symmetries with their associated Z' bosons [271]. The large wealth of realisations deviate for example in the masses and couplings of the Z' bosons, the scale of $U(1)'$ symmetry breaking, or the $U(1)'$ charges of the SM fermions and the Higgs doublet. Vector NSI couplings could then emerge by integrating out a “flavour-sensitive” Z' with direct couplings to SM fermions, though other possibilities without direct couplings such as Z - Z' mixing are also conceivable [262, 272]. More specifically, in the light of the GMOD exhibited by neutrino oscillation experiments, models with a light Z' with a mass of the order of 10 MeV which can viably generate the large e - μ NU NSI strength required have received much attention [273].

5

Global knowledge of the PMNS paradigm and NSI

By combining multiple types of experiments (solar, atmospheric, accelerator, and reactor neutrino, as well as neutrino scattering experiments), global fits exploit complementarities between different observables that are sensitive to the same fundamental neutrino parameters. They thereby often reduce degeneracies that impede the determination of the parameters of interest when the experiments are analysed individually. While standard three-neutrino oscillation physics has transitioned from the phase of discovery to an era of precision measurements of the PMNS paradigm, the existence of NSI is yet to be proven.

Throughout this chapter, we restrict ourselves to the three-neutrino scenario and begin with an overview of the defining characteristics of the relevant types of neutrino fluxes—before and after propagation including flavour transitions—in Sec. 5.1. Solar neutrinos are discussed first, owing to the striking impact these have had on our understanding of neutrino flavour transitions [274]. The flux of neutrinos produced in the Earth’s atmosphere, which is in the centre of this work, is examined next. In both cases, the key phenomenological aspects of flavour transitions are illustrated in some detail, with reference to the matter effects with SM interactions introduced in Sec. 3.3.4. It turns out that the Sun and the Earth represent ideal natural “laboratories” in this regard [275, 240, 253]: first and foremost, because the energies of neutrinos produced within the Sun or the Earth’s atmosphere have similar relative magnitudes with respect to the relevant mass-squared differences (the “solar” Δm_{21}^2 and the “atmospheric” Δm_{31}^2) and this scale is similar to that of the standard matter potential and to the inverse of the Earth’s radius, R_{\oplus}^{-1} . Furthermore, the allowed values of the mixing angles θ_{12} and θ_{23} allow for reasonably large or even maximal

effects in terrestrial oscillation experiments, while the relative smallness of θ_{13} permits adopting a two-neutrino system in the description of the evolution of both solar and atmospheric neutrino fluxes (in the absence of NSI). Accelerator and reactor neutrino experiments are also touched upon.

Section 5.2 concludes this chapter with a summary of existing global constraints on the parameters of the PMNS paradigm and considers NSI in neutrino propagation as additional degrees of freedom in global fits. At the same time, statistical inference concepts which are re-encountered in the analyses in Parts III and IV are mentioned.

5.1 Neutrino sources

5.1.1 Solar neutrinos

The Sun is an intense source of electron neutrinos, emitted predominantly by two groups of thermonuclear fusion reactions, the pp chain and the catalytic CNO cycle (see, e.g., Ref. [276] and references therein). In effect, both of these convert four protons into a ${}^4\text{He}$ nucleus under the emission of two electron neutrinos, $4p \rightarrow {}^4\text{He} + 2e^+ + 2\nu_e$.

“Standard Solar Models”¹ predict the neutrino energy spectrum that is produced by any particular reaction, as exemplified on the left of Fig. 5.1 for the Earth. “ pp neutrinos” are emitted by the initial proton-proton fusion reaction of the pp chain, $p+p \rightarrow {}^2\text{H}+e^++\nu_e$. Their mean energy is 0.27 MeV and they contribute approximately 90% of the total solar neutrino flux on Earth, $\Phi_{\odot} \sim 6.5 \times 10^{10} \text{ cm}^{-2} \text{ s}^{-1}$ [128]. The maximal solar neutrino energy of 18.8 MeV is the Q value of the reaction ${}^3\text{He} + p \rightarrow {}^4\text{He} + e^+ + \nu_e$ [128]. These neutrinos are known as “ hep neutrinos” and are also a product of the pp chain. However, the low fusion rate renders their flux almost seven orders of magnitude smaller than that of pp neutrinos [128]. As evidenced by Fig. 5.1, in terms of both overall flux and mean energy, ν_e ’s from other reactions that are part of the pp chain (${}^7\text{Be}$, pep , ${}^8\text{B}$) or the CNO cycle (${}^{13}\text{N}$, ${}^{15}\text{O}$, ${}^{13}\text{N} + e^-$, ${}^{15}\text{O} + e^-$) lie in between. Fluxes of solar $\bar{\nu}_e$ ’s on Earth are exceeded by Earth-originating and reactor antineutrino backgrounds by several orders of magnitude, and thermal, keV-range, neutrinos and antineutrinos of all flavours produced in the Sun are similarly undetectable as of today [277].

Based on up-to-date global knowledge of the PMNS parameters, the panel on the right of Fig. 5.1 shows the expected probability as a function of neutrino energy for

¹Or “Simplified Solar Models” according to Ref. [277].

5.1 Neutrino sources




Figure 5.1: Calculations and measurements of the properties of solar neutrino flux components on Earth. *Left panel:* solar neutrino flux spectra on Earth. Fluxes from pp -chain reactions are shown in blue, with measured normalisations, whereas fluxes generated by the CNO cycle are shown in orange (calculations and upper limits). *Right panel:* calculations of the solar ν_e survival probability for major flux components. Figures taken from Ref. [277]. See text for details.

ν_e 's produced in the Sun—by the various nuclear reactions that comprise the pp chain and the CNO cycle—to retain their flavour during propagation from their production region to the Earth, $P_{ee}(E_\nu)$. This survival probability depends on the process of origin because of the different radial ν_e production distributions (see, for example, Ref. [277, Fig. 11]).

As detailed in Refs. [278, 183], to high accuracy P_{ee} is the result of the independent and incoherent adiabatic propagation of the neutrino matter eigenstates ν_{im} inside the Sun (cf. Sec. 3.3.4), which transform as $\nu_{im} \rightarrow \nu_i$ at the Sun's surface, propagate without modification to the Earth, at the surface of which each is decomposed into the eigenstates in Earth matter, which then undergo oscillations as they travel to the detector. Depending on energy, at night these matter oscillations lead to a small regeneration or suppression of the ν_e flux arriving at a detector on Earth. The fact that matter effects are negligible during day time at the detector location, when detected neutrinos only have to cross the Earth's crust, results in a few-percent day-night asymmetry [278] (not included in Fig. 5.1).

The horizontal lines in the right panel mark the low- and high-energy limits of the survival probability. Under the currently accepted “large-mixing angle (LMA) MSW solution” [128] to the solar neutrino problem, matter effects in the Sun are

not relevant at low energy, because the encountered electron number density is always far below the resonance value (cf. Sec. 3.3.4.2). In the corresponding limit, one simply observes the averaged survival probability for propagation through vacuum, $\lim_{E_\nu \rightarrow 0} P_{ee} \approx \langle P_{ee}^{\text{vacuum}} \rangle = 1 - \frac{1}{2} \sin^2 2\theta_{12}$ [128, 278] at leading order. In the high-energy limit, $\lim_{E_\nu \rightarrow \infty} P_{ee} \approx \sin^2 \theta_{12}$ [128, 278], as the resonance density is crossed adiabatically.

The range of intermediate energies, referred to as the “transition region”, contains the value $E_\nu \approx 2$ MeV for which the MSW resonance appears at the maximal density within the Sun [278]. The size of the 1–2 mixing angle θ_{12} governs the resonance width (cf. Fig. 3.6) and thus the width of the transition region.

5.1.2 Atmospheric neutrinos

The fact that Earth’s atmosphere is subject to a constant influx of primary cosmic rays means that it continuously produces copious amounts of neutrinos², whose flux is sufficiently large to be studied with terrestrial detectors across the energy range from \odot (100 MeV) to \odot (10 TeV). The early phenomenological flux discussions in Refs. [279, 280] stand in contrast with today’s sophisticated computational modelling tools, such as Refs. [281, 282, 283], which differ in the underlying computational methods or hadronic-interaction and atmospheric-density models. Analytic methods based on simplifying assumptions about the primary cosmic-ray spectrum and hadronic interactions are able to capture the essential flux characteristics (see Ref. [284] and references therein).

5.1.2.1 Flux characteristics

The spectrum of primary cosmic rays spans more than ten orders of magnitude in energy. Approximately three out of four cosmic-ray nucleons are free protons; nucleons bound in helium nuclei contribute about 70 % of the remaining fraction [15]. At energies in the range from 1 GeV to 100 TeV, this nucleon flux is reasonably well described by a power law $\propto E^{-\alpha}$ ($\alpha \equiv \gamma + 1 \approx 2.7$), where E and α respectively γ are the energy per nucleon and the differential respectively integral spectral index.³

²The term “neutrinos” here is used to jointly refer to neutrinos and antineutrinos for brevity. Similarly, “muon” refers to the particle and the antiparticle, unless an explicit distinction is made.

³For protons, deviations from such a simple power-law form have begun to emerge from data taken by various cosmic-ray experiments, and there are similar indications as to the existence of more complex spectral features for nuclei [285].

Cosmic-ray reactions typically occur in the stratosphere—at altitudes from approximately 10 km to 40 km [128]. They produce particle showers comprising large numbers of unstable secondary hadrons, which give rise to atmospheric neutrino fluxes that depend on energy and direction. Here, we are only concerned with the “conventional” contribution to these fluxes, namely decays of charged pions and of kaons. Above TeV-scale neutrino energies, weak decays of hadrons containing a charm (anti)quark produce the “prompt” flux. Particles containing the charm flavour have a lifetime of $\mathcal{O}(10^{-12} \text{ s})$ [15] and almost certainly decay rather than interact in the atmosphere up to the highest energies, but their production cross section is kinematically strongly suppressed compared to that of pions and kaons [286, 287]. In the absence of oscillations, charm decays would be the dominant source of atmospheric tau neutrinos by a large margin [288]. The intrinsic conventional flux of tau neutrinos, however, is negligible [287].

As the main contribution to the flux of atmospheric muon neutrinos with GeV energies, the decay of a charged pion results in a $\bar{\nu}_\mu \mu^\pm$ pair with a branching ratio of $\sim 99.99\%$ [15]:

$$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu, \quad \pi^+ \rightarrow \mu^+ + \nu_\mu. \quad (5.1)$$

The corresponding electron decay mode only has a branching ratio of $\sim 10^{-4}$ [15]. Heavier charged kaons directly generate neutrinos via the leptonic and semileptonic decays

$$K^- \rightarrow \mu^- + \bar{\nu}_\mu, \quad K^+ \rightarrow \mu^+ + \nu_\mu, \quad (5.2)$$

$$K^- \rightarrow \pi^0 + e^- + \bar{\nu}_e, \quad K^+ \rightarrow \pi^0 + e^+ + \nu_e, \quad (5.3)$$

and indirectly through hadronic decays that produce one or two charged pions. Similar semileptonic and hadronic neutral-kaon (K_L^0, K_S^0) decays are also relevant [15, 289, 290]. μ^\pm 's generated in any of the above reactions may decay before reaching the ground, thereby producing both electron and muon neutrinos via the two reactions

$$\mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu, \quad \mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu. \quad (5.4)$$

The resulting neutrino flux spectra (without oscillations) from 100 MeV to 10 TeV expected at the South Pole—at an observation height of 2.84 km above sea level—are shown on the left of Fig. 5.2, calculated in Ref. [281]. On the right, the corresponding flux ratios, namely flavour ratios and neutrino-to-antineutrino ratios, are displayed. Both fluxes and their ratios are averaged over all directions and over the span of one year.




Figure 5.2: Predicted properties of the atmospheric neutrino flux spectra at the South Pole. *Left panel:* direction- and year-averaged atmospheric electron and muon neutrino flux spectra (scaled by neutrino energy to the third power). *Right panel:* various direction- and year-averaged atmospheric neutrino flux ratios as a function of neutrino energy. Figures taken from Ref. [281]. See text for details.

Since the μ^\pm decay probability above ground decreases the higher the energy, the neutrino flux at higher energy is dominated by direct charged-pion and kaon decays (up to the cross-over energy above which the prompt component dominates). These mesons' relative contributions depend on their production ratios, as well as on their decay and reinteraction probabilities and kinematics. Whether decay or reinteraction is the dominant effect for a given parent meson is governed by the relative size of its decay length compared to its interaction length: above the critical energy E_c —which depends on the atmospheric density profile and on zenith—the meson is more likely to interact than to decay. For vertical trajectories, the critical energies are $E_c^{\pi^\pm} \approx 115 \text{ GeV}$, $E_c^{K_L^0} \approx 210 \text{ GeV}$, $E_c^{K^\pm} \approx 850 \text{ GeV}$, and $E_c^{K_S^0} \approx 1.2 \times 10^5 \text{ GeV}$ [286, 290].

Depending on whether one considers meson decay below or above the critical energy, the resulting neutrino flux properties differ. Far below the critical energy of the parent meson, the spectral index coincides with that of the primary nucleon flux, α , whereas the slope at high energies is steeper by one unit, $\alpha + 1$. On the left of Fig. 5.2, this transition is seen to occur in the few-GeV range for $\bar{\nu}_e$'s and to extend over several

hundreds of GeV for $\bar{\nu}_\mu$'s.

While neutrinos with $E_\nu \ll E_c^\pi$ predominantly originate from the decays of the lighter charged pions, the relation $E_c^K \gg E_c^\pi$ renders kaons increasingly important for atmospheric neutrino production at higher energy. Moreover, the rather close masses of the pion and the muon result in the latter carrying away much of the energy of the decaying pion, whereas the dominant, leptonic kaon decay splits the kaon energy almost equally between the muon and the neutrino on average [280]. Despite a comparably small K/π production ratio [291], kaons are the dominant neutrino source for $E_\nu \gtrsim 100$ GeV, for both electron and muon neutrinos [292].

Also, the critical energies increase for more inclined trajectories, $E_c \propto F(\cos \vartheta)$, where $F \approx \cos^{-1} \vartheta$ for zenith angles $\vartheta \lesssim 60^\circ$, rising monotonically up to factor of around 10 at the horizontal [293]. Thus, the larger the zenith angle, the higher the energy above which the neutrino flux spectra steepen. Not taking into account neutrino oscillations or the small interaction probability with Earth matter, this anisotropy exhibits an up-down symmetry [128],

$$\bar{\Phi}_{\text{unosc}}^\alpha(\cos \vartheta) = \bar{\Phi}_{\text{unosc}}^\alpha(-\cos \vartheta), \quad (5.5)$$

where $\bar{\Phi}_{\text{unosc}}^\alpha$ is the atmospheric flux of (anti)neutrinos of flavour α .

In the few-GeV energy range, the flux is neither isotropic nor up-down symmetric, as a result of effects of the geomagnetic field on the incident charged primary nucleons and on charged secondaries within the atmosphere. Furthermore, solar activity leads to a temporal modulation of the neutrino fluxes at similar energies, though the impact on neutrinos is dampened with respect to that on primary nucleons due to the minimal ("cutoff") rigidity imposed on these by the geomagnetic field [289].

If all muons decayed and the kaon component was negligible, one would expect a flavour ratio

$$\frac{\Phi_{\text{unosc}}^\mu + \bar{\Phi}_{\text{unosc}}^\mu}{\Phi_{\text{unosc}}^e + \bar{\Phi}_{\text{unosc}}^e} = 2 \quad (5.6)$$

of the atmospheric fluxes on the ground. As the right panel of Fig. 5.2 underscores, this is a good approximation at neutrino energies around and below 1 GeV. At these energies, one half of the muon neutrino flux originates from the charged-pion decay (5.1) directly, and the other from the muon decay (5.4) [281]. Correspondingly, the flux ratio $\Phi_{\text{unosc}}^\mu/\bar{\Phi}_{\text{unosc}}^\mu \approx 1$. It almost reaches twice this value at 10 TeV. The flux ratio $\Phi_{\text{unosc}}^e/\bar{\Phi}_{\text{unosc}}^e$ at low energy reflects the parent pion and secondary muon charge ratio, $\Phi_{\pi^+}/\Phi_{\pi^-} = \Phi_{\mu^+}/\Phi_{\mu^-}$, which is somewhat larger than unity because of the excess of positive charge in cosmic rays [281].

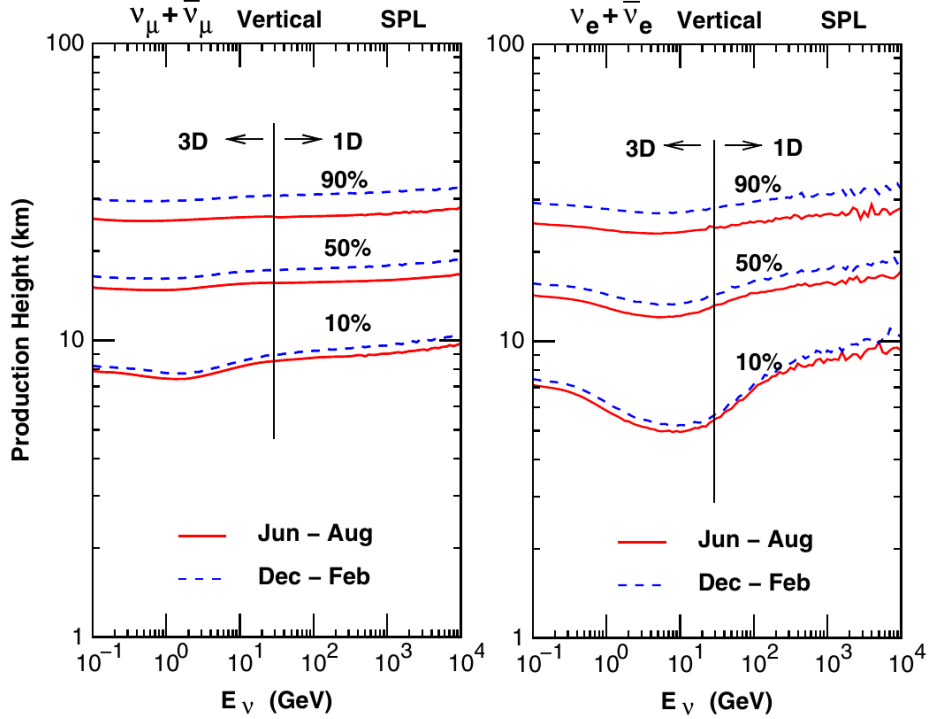


Figure 5.3: Predicted atmospheric neutrino production height quantiles for vertically downward trajectories at the South Pole, averaged over the two annual periods shown, for the combined flux of $\bar{\nu}_\mu$'s on the left, and for that of $\bar{\nu}_e$'s on the right. Figure adapted from Ref. [281].

5.1.2.2 Standard oscillations

The possible oscillation channels for conventional atmospheric neutrinos are $\bar{\nu}_\alpha \rightarrow \bar{\nu}_e, \bar{\nu}_\mu, \bar{\nu}_\tau$ with $\alpha = e, \mu$. Isocontours of constant probability for each channel as a function of energy E_ν and zenith angle $\cos \vartheta$ are visualised with the help of “oscillograms” [294]. Assuming a universal oscillation probability exists for the problem at hand (cf. Sec. 3.2.2), inputs to the calculation of a given oscillogram are the production height (distribution), the Earth’s electron density distribution, the leptonic mixing matrix, the neutrino mass-squared differences, and the detector depth.

Production height and baseline

For a neutrino produced at the altitude h impinging on a detector at a depth d below sea level under the zenith angle ϑ , the baseline follows as [128]

$$L(\vartheta; h, d) = \sqrt{(R_\oplus + h)^2 - (R_\oplus - d)^2 \sin^2 \vartheta} - (R_\oplus - d) \cos \vartheta, \quad (5.7)$$

with $R_{\oplus} \approx 6371$ km. Usually, the detector is located underground—shielded by several kilometres of water equivalent—in order to suppress the background of atmospheric μ^{\pm} 's, which are penetrating enough for their flux at the detector to nonetheless significantly exceed that of $\bar{\nu}$'s.

Figure 5.3 gives an idea of the neutrino production height distributions, showing their simulated 10th, 50th, and 90th percentiles for ν_{μ} 's and $\bar{\nu}_{\mu}$'s or ν_e 's and $\bar{\nu}_e$'s combined, restricted to zenith angles $\cos \vartheta > 0.9$ at the South Pole as a function of neutrino energy, averaged over the months June to August or December to February. As a result of the transition to kaon decay as the dominant neutrino production mechanism, the high-energy production height distribution closely follows that for kaons. Large seasonal variations of the atmospheric density at the South Pole are responsible for the $\mathcal{O}(10\%)$ greater production heights encountered during the austral summer months December to February [281].

As a consequence of the rather stable average production height $h \sim 15$ km to 20 km—also at sites other than the South Pole [281]—in analysis it is common to assume h to be fixed. The relative variations are particularly small compared to the baselines of neutrinos reaching the detector from below the horizon ($\cos \vartheta < 0$)—the only directions for which significant flavour transitions are expected to occur—with $L_{\max} \approx D_{\oplus} \equiv 2R_{\oplus}$. This is illustrated in Fig. 5.4, which represents Eq. (5.7) as a function of $\cos \vartheta$ for two fixed production heights of $h = 15$ km and $h = 20$ km. The baselines for the latter are larger by between 0.04% ($\cos \vartheta = -1$) and 29% ($\cos \vartheta = 1$). Neutrinos impinging horizontally on the underground detector have a baseline of either 530 km or 466 km, almost half of which crosses the Earth's crust (see below).

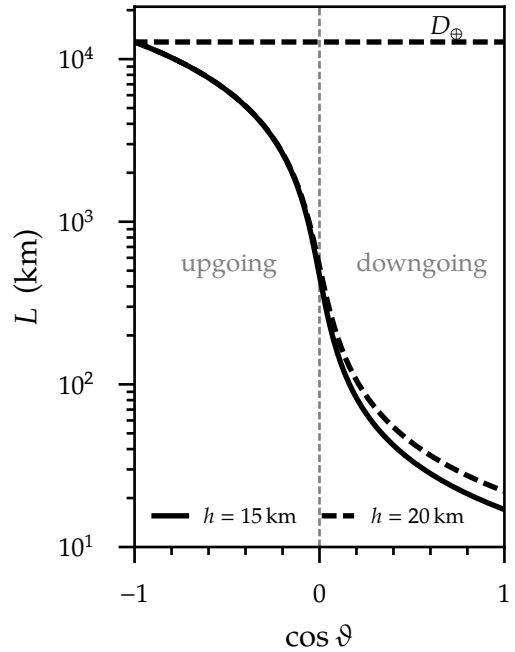


Figure 5.4: Atmospheric neutrino baseline versus zenith, Eq. (5.7), for a detector at a depth of $d = 2$ km, for two representative production heights.

Density profile

The conventionally employed density profile is provided by the “preliminary reference Earth model” (PREM) [237] in terms of a radial mass density distribution $\rho_{\oplus}(r)$. PREM has to be combined with the relative local number density of electrons, $Y_e(r)$, via the relation (3.53), in order to obtain the potential $V_{CC}(x; \cos \vartheta)$ encountered by an atmospheric neutrino on a trajectory under the zenith angle ϑ . Broadly speaking, PREM defines an “average Earth model, [...], a mathematical abstraction” [237] with four principal entities in the form of concentric shells as follows.

- The *crust* has a thickness of some tens of kilometres and an average density of $\langle \rho_{\text{crust}} \rangle \sim 2.8 \text{ g cm}^{-3}$.
- The *mantle* is the thickest layer, extending over 2 900 km in radius. With an average density of $\langle \rho_{\text{mantle}} \rangle \sim 4.5 \text{ g cm}^{-3}$, it contributes approximately two thirds of the Earth’s total mass.
- The liquid *outer core* has a radius of nearly 2 200 km and an average density of $\langle \rho_{\text{outer}} \rangle \sim 11.5 \text{ g cm}^{-3}$. A sharp jump in density characterises the transition from the (lower) mantle to the outer core.
- The *inner core* is a rather spherical ellipsoid extending about 1 220 km in radius. It is separated from the outer core by a steep density transition. Its density varies between approximately 12.8 g cm^{-3} close to its boundary and 13.1 g cm^{-3} at its centre.

This work assumes an electron fraction of $Y_e = 0.496$ for the crust and mantle and $Y_e = 0.466$ for the (inner and outer) core [1, 295].

Figure 5.5 provides a sketch of the zenith ranges subtended by the mantle, the outer core, and the inner core. Earth-crossing neutrinos on shallow trajectories with $\cos \vartheta \gtrsim -0.84$ only traverse the mantle (“mantle domain”), whereas paths with $\cos \vartheta \lesssim -0.84$ traverse the core (“core domain”), with the transition between the outer and the inner core occurring at $\cos \vartheta \approx -0.98$.

Figure 5.6 depicts different matter-potential profiles, Eq. (3.54), encountered depending on zenith angle. These determine how matter effects shape the respective transition probabilities. Among the displayed choices, the neutrino on a trajectory with $\cos \vartheta = -0.1$ is the only one with a notable propagation distance through the atmosphere, where the value of the potential V_{CC} is at least three orders of magnitude smaller than inside the Earth. Multiple layer transitions are visible for every trajectory.

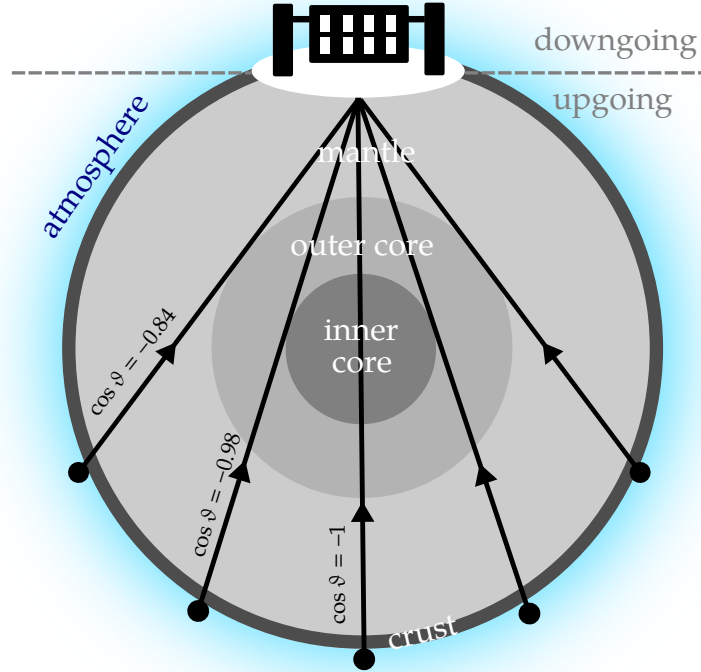


Figure 5.5: Sketch of various Earth-crossing trajectories of atmospheric neutrinos that are detected at the South Pole, and hence are referred to as “upgoing”. Sharp density changes exist at the boundary between the mantle and the core and at the transition between the outer and the inner core. Only neutrinos with $\cos \vartheta \lesssim -0.84$ ($L \gtrsim 10\,726$ km) traverse the core.

Three-flavour oscillation probabilities from two-neutrino effects

All salient features of the neutrino evolution relevant to current and next-generation very-large-volume atmospheric neutrino experiments (as well as very-long-baseline accelerator neutrino experiments) can be understood by neglecting the 1–2 mass-squared difference Δm_{21}^2 (based on its relative magnitude of only $\sim \Delta m_{31}^2/30$, see Sec. 5.2). It is then possible to describe neutrino evolution in terms of a 2ν system. This is shown in Ref. [294], on which the following discussion of standard three-flavour atmospheric neutrino oscillations at the GeV scale is based, together with Ref. [239]. The flavour oscillation probabilities derived in this framework are valid to a good approximation for neutrino energies exceeding 1 GeV to 2 GeV.

Employing a “propagation basis”⁴ $\nu' \equiv (\nu_e, \nu'_2, \nu'_3)^T$ defined by the transformation

⁴There is no unique definition of the propagation basis in the literature, as the optimal choice depends on the considered oscillation regime. For example, in Refs. [296, 297], which study sub-GeV neutrino evolution in the Earth, the basis states are additionally rotated through the 1–3 mixing.




Figure 5.6: Expected matter-potential profiles, Eq. (3.54), encountered by atmospheric neutrinos with different zenith angles, based on PREM.

(following the field, not state, notation of Refs. [294, 239])

$$\nu_f = U' \nu' \equiv U_{23} I_{\delta_{\text{CP}}} \nu' \quad (5.8)$$

with $I_{\delta_{\text{CP}}} \equiv \text{diag}(1, 1, e^{i\delta_{\text{CP}}})$ has the advantage that neither θ_{23} nor δ_{CP} enter the dynamical problem. Instead, both parameters' impact only arises through the projection from the flavour basis onto the propagation basis and vice versa. Moreover, the oscillation probabilities in the 3ν system can be expressed in terms of 2ν amplitudes for transitions between the various propagation basis states. A further simplification arises from the fact that—at least to a good approximation for upgoing atmospheric neutrinos with deep trajectories—the symmetric Earth density profile implies a nearly symmetric potential (cf. Sec. 3.3.3)

$$V_{\text{CC}}(x) \approx V_{\text{CC}}(L - x) . \quad (5.9)$$

The vector of the propagation basis states evolves as governed by the Hamiltonian

An earlier discussion of the GeV oscillation regime employing the basis (5.8) under the assumption of vanishing δ_{CP} can be found in Ref. [298]. This approach gives the same results that are obtained in the limit $\Delta m_{21}^2 \ll \Delta m_{31}^2$ [294].

$H'(x) = U_{13}U_{12}\frac{M_{\text{diag}}^2}{2E_\nu}U_{12}^\dagger U_{13}^\dagger + V(x)$, whose effective potential V is unmodified with respect to the Hamiltonian in the flavour basis. Given $\Delta m_{21}^2 \ll \Delta m_{31}^2$, H' suggests that the state ν'_2 evolves independently—it “decouples”. The 3ν evolution matrix in the propagation basis reflects this,

$$S'(x) \approx \begin{pmatrix} A_{ee}(x) & 0 & A_{e3}(x) \\ 0 & A_{22}(x) & 0 \\ A_{3e}(x) & 0 & A_{33}(x) \end{pmatrix}, \quad (5.10)$$

with $A_{22}(x) \equiv \exp\{-i\phi_2(x)\}$, $\phi_2(x) \equiv \frac{\cos^2\theta_{12}\Delta m_{21}^2 x}{2E_\nu}$.

The evolution of the 2ν subsystem (ν_e, ν'_3) is determined by θ_{13} , Δm_{31}^2 , and V_{CC} , with a $\nu_e \leftrightarrow \nu'_3$ transition probability $P_{e3} \equiv |A_{e3}|^2 = |A_{3e}|^2 = 1 - |A_{ee}|^2 = 1 - |A_{33}|^2$ (dependence on x and E_ν suppressed).⁵

In the flavour basis, $S(x) = U'S'(x)U'^\dagger$, which implies that the observable oscillation probabilities are given by the matrix elements

$$P_{\alpha\beta} = \left| \left(U_{23} I_{\delta_{\text{CP}}} S' I_{\delta_{\text{CP}}}^\dagger U_{23}^\dagger \right)_{\alpha\beta} \right|^2. \quad (5.11)$$

Explicitly, in terms of P_{e3} and A_{33} and in the limit $\Delta m_{21}^2 \rightarrow 0$ the oscillation probabilities follow as [298, 294]:

$$P_{ee} = 1 - P_{e3}, \quad (5.12)$$

$$P_{e\mu} = P_{\mu e} = s_{23}^2 P_{e3}, \quad (5.13)$$

$$P_{e\tau} = P_{\tau e} = c_{23}^2 P_{e3}, \quad (5.14)$$

$$P_{\mu\mu} = 1 - s_{23}^4 P_{e3} - 2s_{23}^2 c_{23}^2 (1 - \text{Re } A_{33}), \quad (5.15)$$

$$P_{\mu\tau} = P_{\tau\mu} = -s_{23}^2 c_{23}^2 P_{e3} + 2s_{23}^2 c_{23}^2 (1 - \text{Re } A_{33}), \quad (5.16)$$

$$P_{\tau\tau} = 1 - c_{23}^4 P_{e3} - 2s_{23}^2 c_{23}^2 (1 - \text{Re } A_{33}), \quad (5.17)$$

where the shorthand notation $s_{ij} \equiv \sin \theta_{ij}$ and $c_{ij} \equiv \cos \theta_{ij}$ is used. The probabilities involving ν_e either equal P_{e3} up to a constant factor (s_{23}^2 for $\nu_e \leftrightarrow \nu_\mu$, c_{23}^2 for $\nu_e \leftrightarrow \nu_\tau$) or its complementary probability $1 - P_{e3}$ ($\nu_e \rightarrow \nu_e$). Also, the replacement $\nu_\mu \leftrightarrow \nu_\tau$ is equivalent to switching $s_{23} \leftrightarrow c_{23}$.

⁵ $A_{e3} = A_{3e}$ is the consequence of assuming a symmetric density profile. Similarly, $A_{e2} = A_{2e}$ and $A_{23} = A_{32}$, which would be relevant if terms proportional to small parameters such as $\Delta m_{21}^2/\Delta m_{31}^2$ were taken into account [239].




Figure 5.7: Neutrino transition scheme in the propagation basis ν' (squares), related to the flavour basis ν_f (circles) according to Eq. (5.8). On the left, the initial flavour state is projected onto the propagation basis states. In the centre, transitions between the propagating states are shown, with ν'_2 evolving independently, cf. Eq. (5.10). On the right, the evolved propagation basis states are projected back onto the flavour basis. Figure adapted from Ref. [239].

Figure 5.7 illustrates how these expressions arise through the various projection and evolution amplitudes [239]. All transitions involving ν_e can only proceed through a single path and therefore exhibit no interference of amplitudes. For example, the $\nu_e \rightarrow \nu_\tau$ transition results from the ν_e transformation into ν'_3 and the subsequent projection of ν'_3 onto ν_τ . The $\nu_\mu \rightarrow \nu_\tau$ transition, in contrast, can proceed either through the projection of ν_μ onto ν'_2 , which evolves independently and then projects onto ν_τ , or through the projection of ν_μ onto ν'_3 , whose amplitude of survival after oscillation then projects onto ν_τ .

The Hamiltonian $H^{(2\nu)}$ for the evolution of the (ν_e, ν'_3) system is the result of removing the second row and column from H' . It can be expressed in a symmetric form by the subtraction [294]

$$H^{(2\nu)} \rightarrow H^{(2\nu)} - \left(\frac{\Delta m_{31}^2}{4E_\nu} + \frac{V_{CC}}{2} \right) \mathbb{1}_{2 \times 2} \quad (5.18)$$

$$= \frac{\Delta m_{31}^2}{4E_\nu} \begin{pmatrix} -\cos 2\theta_{13} & \sin 2\theta_{13} \\ \sin 2\theta_{13} & \cos 2\theta_{13} \end{pmatrix} + \begin{pmatrix} \frac{V_{CC}}{2} & 0 \\ 0 & -\frac{V_{CC}}{2} \end{pmatrix}. \quad (5.19)$$

The neutrino system evolving according to this Hamiltonian is subject to the matter effects investigated in Sec. 3.3.4.

In the case of uniform matter, the effective mixing angle of the 2ν system is given by Eq. (3.60) with the substitutions $\theta \rightarrow \theta_{13}$ and $\Delta m^2 \rightarrow \Delta m_{31}^2$. The same substitutions have to be applied to the level splitting (3.61). It is then straightforward to find the

full $\nu_e \leftrightarrow \nu_3'$ transition probability, which reads

$$P_{e3}(L, E_\nu) = \frac{\sin^2 2\theta_{13}}{R_{13}(E_\nu)} \sin^2 \left(\frac{L}{4E_\nu} \Delta m_{31}^2 \sqrt{R_{13}(E_\nu)} \right), \quad (5.20)$$

where we have defined the resonance factor

$$R_{13}(E_\nu) \equiv \left(\cos 2\theta_{13} - \frac{2V_{CC}E_\nu}{\Delta m_{31}^2} \right)^2 + \sin^2 2\theta_{13}. \quad (5.21)$$

Equation (5.20) underscores that each oscillation channel involving ν_e , described by the probabilities (5.12) to (5.14), exhibits an effective 2ν behaviour, with the following effective mixing angles governing the oscillation amplitudes [128]:

$$\sin^2 2\theta_{ee}^m \equiv \sin^2 2\theta_{13}^m = \frac{\sin^2 2\theta_{13}}{R_{13}}, \quad (5.22)$$

$$\sin^2 2\theta_{e\mu}^m \equiv \sin^2 \theta_{23} \sin^2 2\theta_{13}^m, \quad (5.23)$$

$$\sin^2 2\theta_{e\tau}^m \equiv \cos^2 \theta_{23} \sin^2 2\theta_{13}^m. \quad (5.24)$$

Earth oscillograms

As seen above, the transition probabilities of atmospheric neutrinos involving ν_e at the GeV energy scale are governed by two-flavour matter effects. Their features can be understood through the adoption of a zenith-dependent constant-density approximation for the layers traversed. In this framework, introduced in Ref. [294], either one or three constant-density layers are encountered along any given trajectory, with at most two distinct densities. The two possible layer configurations consist of only the mantle or the mantle, the core, and again the mantle. The average potential faced by a neutrino travelling the distance L_i in the i th layer (mantle or core),

$$\langle V_i(\vartheta) \rangle = \frac{1}{L_i(\vartheta)} \int_0^{L_i} V_i(x) dx, \quad (5.25)$$

is substituted wherever V_{CC} occurs in Sec. 3.3.4. Points in the E_ν - $\cos \vartheta$ plane for which the initial ν_e flux is completely depleted ($P_{e3} = 1 \Leftrightarrow P_{ee} = 0$) are determined by the amplitude and phase conditions introduced in Sec. 3.3.4. The manner in which the conditions are realised depends on the considered trajectory.

For propagation through the *mantle only*, the resonance energy can be estimated by assuming some typical (average) matter potential in Eq. (3.63), which yields $E_R \approx$

6 GeV. The value of the zenith angle required to satisfy the phase condition then follows as (neglecting both h and d)

$$\cos \vartheta \approx \frac{-\pi E_R}{R_\oplus \sin 2\theta_{13} \Delta m_{31}^2} \approx -0.77. \quad (5.26)$$

These conditions constitute the MSW resonance peak of P_{e3} (and therefore P_{ee} , $P_{e\mu}$, and $P_{e\tau}$), Eq. (5.20), in the mantle.

For core-crossing trajectories, transitions are parametrically enhanced if the conditions are adequate. For the Earth, for any given zenith angle the “castle wall” density profile from Sec. 3.3.4.3 is a good approximation. The first layer corresponds to the mantle, and the second one to the core. Since the density profile encountered upon the second crossing of the mantle is T-reversed with respect to the first, the evolution matrix follows as $S = S_1^T S_2 S_1$. Applying the parametric resonance conditions (3.76) results in three “ridges” of parametric enhancement of P_{e3} appearing in the core domain for $E_\nu \gtrsim 3$ GeV. These represent different oscillation phases acquired during the evolution in the core. In addition, at a slightly lower energy $2.5 \text{ GeV} \lesssim E_\nu \lesssim 2.8 \text{ GeV}$, there is a ridge which can be attributed to the MSW resonance enhancement in the core.

Panel (a) of Fig. 5.8 shows the oscillogram of the appearance channel $\nu_e \rightarrow \nu_\mu$ subject to the matter effects above. This channel has a maximal transition probability of $\sin^2 \theta_{e\mu, \max}^m = \sin^2 \theta_{23}$, whereas the appearance channel $\nu_e \rightarrow \nu_\tau$ has a maximum of $\sin^2 \theta_{e\tau, \max}^m = \cos^2 \theta_{23}$. The oscillogram of the ν_e disappearance channel follows from appropriately rescaling, inverting, and superimposing all features of both appearance channels.

Panel (c) depicts the oscillogram of the disappearance channel $\nu_\mu \rightarrow \nu_\mu$. Its main features, determined by the superposition of multiple terms, can be interpreted as deviations from the 2ν vacuum oscillation probability

$$P_{\mu\mu}^{2\nu} \equiv 1 - \sin^2 2\theta_{23} \sin^2 \left(\frac{\Delta m_{31}^2 L}{4E_\nu} \right), \quad (5.27)$$

and are typically of the same order as the latter. In general, these deviations are significant for regions in E_ν and $\cos \vartheta$ where P_{e3} exhibits a large enhancement due to matter effects. A similar statement applies to the transition channel $\nu_\mu \rightarrow \nu_\tau$.

The corresponding antineutrino channels are shown in panels (b) and (d) of Fig. 5.8. For the normal ordering, matter generally suppresses the oscillations of antineutrinos; the strongest transitions occur where matter effects are negligible.




Figure 5.8: Earth oscillograms for atmospheric (anti)neutrino transitions, taken from Ref. [239]. Panel (a) shows $P_{\nu_e \rightarrow \nu_\mu}$, (b) $P_{\nu_e \rightarrow \bar{\nu}_\mu}$, (c) $P_{\nu_\mu \rightarrow \nu_\mu}$, and (d) $P_{\bar{\nu}_\mu \rightarrow \bar{\nu}_\mu}$. See text for details.

A more extensive discussion of all oscillograms is found in Refs. [294, 239].

5.1.3 Accelerator neutrinos

A neutrino beam can be created artificially by aiming a beam of accelerated protons at a thick nuclear target, thereby leading to the production of hadronic secondaries such as pions and kaons. These are able to produce neutrinos via the reactions (5.1) to (5.3).

If the target is thick enough, the positively charged antimesons, π^+ and K^+ , will come to rest and decay (DAR), whereas the negatively charged mesons, π^- and K^- , will mostly be absorbed, as illustrated schematically in Fig. 5.9. In the case of a stopped-pion source, the resulting neutrino flux will be isotropic and made up of a prompt monoenergetic ν_μ component at $E_\nu \approx 29.8$ MeV and a continuous spec-

trum of delayed ν_e 's and $\bar{\nu}_\mu$'s with an endpoint energy of $E_\nu \approx 53$ MeV from the μ^+ decay. Stopped kaons produce a higher-energy monochromatic ν_μ component at $E_\nu \approx 236$ MeV in approximately 64 % of their decays, but their production rate is typically one order of magnitude below that of pions [15].

Given their isotropy upon emission, DAR neutrinos are typically detected at a distance L of several tens of metres from their source. Therefore, these experiments are characterised by $L/E_\nu \sim 1$ m/MeV, which results in sensitivity to oscillations involving a sterile neutrino state with a characteristic mass-squared difference $\Delta m^2 \sim 1$ eV² [128]. To this end, one can consider the disappearance channel $\nu_e \rightarrow \nu_e$ or the appearance channel $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$. In addition, the pion and muon DAR flux is employed to constrain the cross section for the CE ν NS process (2.60) and to search for any deviations from the SM [300]. The ν_μ flux from kaon DAR is considered as

a ‘‘standard candle’’ that can benefit the determination of the CC scattering cross section in the energy range of a few hundred MeV [301]. In the future, measuring the oscillated DAR neutrino flux at long baselines of $L \sim \mathcal{O}(10^2 \text{ km})$ could provide useful consistency checks of the standard three-neutrino oscillation paradigm [302].

When the hadronic secondaries are allowed to decay in flight (DIF) instead (or in addition), such as in a decay tunnel of several hundreds of metres in length, depending on the primary proton energy the decay neutrinos may reach energies of several GeV. In this case, if a detector is placed at a distance $L \sim \mathcal{O}(1 \text{ km})$, the experiment also has sensitivity to $\Delta m^2 \sim 1$ eV² [128]. However, the fact that the DIF neutrino flux is boosted in the forward direction generally allows adopting longer baselines of up to $L \sim \mathcal{O}(10^3 \text{ km})$, such that mass-squared differences of $\Delta m^2 \gtrsim 10^{-3}$ eV² become accessible [128]. In the context of the aforementioned Earth oscillograms, these long-baseline (LBL) accelerator neutrino experiments essentially study 3ν oscillations at a fixed zenith angle and over a narrow range of neutrino energies, depending on the underlying flux spectrum. As a consequence of the rather shallow neutrino trajectories with $\cos \vartheta \gtrsim -0.3$, matter effects in LBL experiments are not expected to be as pronounced as in the case of atmospheric neutrino experiments. This makes the




Figure 5.9: Illustration of pion decay at rest, after an accelerated proton hits a thick enough nuclear target. Figure taken from Ref. [299].




Figure 5.10: Calculations of the properties of reactor antineutrino fluxes. *Left panel:* expected number flux spectra (per fission and MeV energy output) of reactor $\bar{\nu}_e$'s from ^{235}U , ^{238}U , ^{239}Pu , and ^{241}Pu fissions and from neutron capture on ^{238}U , with the latter scaled down by a factor of 20 for display purposes. The dashed vertical line marks the IBD energy threshold. Figure adapted from Ref. [277]. *Right panel:* 3ν oscillation probabilities for a 3 MeV $\bar{\nu}_e$ as a function of distance between reactor and detector. The vertical axis represents the flavour composition of the initially pure $\bar{\nu}_e$ flux. Shown as an inset is the δ_{CP} -sensitive appearance probability $P_{\bar{\nu}_e \rightarrow \bar{\nu}_\mu}$ for three different values of δ_{CP} . Figure adapted from Ref. [304].

former more susceptible to parameter degeneracies brought about by the unknown neutrino mass ordering, θ_{13} , δ_{CP} , and the octant of θ_{23} [239].

5.1.4 Reactor neutrinos

Nuclear reactors are the second common source of artificial neutrino fluxes: electron antineutrinos are emitted in β^- -decays of neutron-rich fission fragments and represent a few percent of the total reactor energy output, with a number flux of the order of 10^{20} s^{-1} . Its energy spectrum is shown in the left panel of Fig. 5.10. The isotopes ^{235}U , ^{238}U , ^{239}Pu , and ^{241}Pu are responsible for more than 99% of all fission processes taking place in the reactor core [15]. At the low-energy end of the spectrum, the flux is dominated by β^- -decays of nuclides produced by neutron-capture processes [277]. However, above the threshold of the IBD process (2.44)—which serves as the primary detection channel for reactor $\bar{\nu}_e$'s—their impact on the flux drops below a few percent [303].

Akin to neutrino fluxes from pion, muon, or kaon DAR at accelerators, the reactor antineutrino flux is emitted isotropically, limiting the maximal distance at which a detector can be placed from the reactor while still retaining a sufficiently high event rate—even more so since the IBD threshold renders around 75% of the integrated

reactor flux output undetectable. With short baselines of a few tens of metres, it is possible to probe characteristic mass-squared differences $\Delta m^2 \gtrsim 0.1 \text{ eV}^2$. When the detector is placed hundreds of kilometres away from the reactor, values down to $\Delta m^2 \gtrsim 10^{-5} \text{ eV}^2$ can be probed [128].

Oscillations of reactor antineutrinos are analysed by observing the disappearance channel $\bar{\nu}_e \rightarrow \bar{\nu}_e$ through the well understood IBD process in liquid scintillator detectors. The plot on the right of Fig. 5.10 gives the expected reactor antineutrino flux composition by flavour as a function of distance from the reactor core, assuming an initially pure monoenergetic 3 MeV $\bar{\nu}_e$ flux. The evolution of its $\bar{\nu}_e$ component can be understood by expressing the corresponding survival probability in vacuum, Eq. (3.31) with $\alpha = \beta = e$, as [304]

$$\begin{aligned}
 P_{\bar{\nu}_e \rightarrow \bar{\nu}_e} = & 1 - \underbrace{\cos^4 \theta_{13} \sin^2 2\theta_{12} \sin^2 \left(\frac{\Delta m_{21}^2 L}{4E_\nu} \right)}_{\text{large amplitude, low frequency ("solar")}} \\
 & - \underbrace{\sin^2 2\theta_{13} \left(\cos^2 \theta_{12} \sin^2 \left(\frac{\Delta m_{31}^2 L}{4E_\nu} \right) + \sin^2 \theta_{12} \sin^2 \left(\frac{\Delta m_{32}^2 L}{4E_\nu} \right) \right)}_{\text{small amplitude, high frequency ("atmospheric")}}. \quad (5.28)
 \end{aligned}$$

For a fixed energy E_ν , the large-amplitude but slowly oscillating term with a frequency $\propto \Delta m_{21}^2/E_\nu$ is superimposed with two small-amplitude but rapidly oscillating terms with frequencies $\propto \Delta m_{31}^2/E_\nu$ and $\propto \Delta m_{32}^2/E_\nu$. One expects the first $\bar{\nu}_e$ disappearance maximum due to Δm_{21}^2 at a distance of $L \approx 50 \text{ km}$ when $\Delta m_{21}^2 \approx 7.5 \times 10^{-5} \text{ eV}^2$, cf. Eq. (3.33). The first disappearance maximum due to the rapid oscillation is found at a value of $L \sim \mathcal{O}(1 \text{ km})$. As can be seen in Fig. 5.10, its amplitude (which is sensitive to θ_{13}) is expected to be at the 10% level only.

5.2 Status of global fits

Several groups of researchers regularly provide analyses of up-to-date experimental neutrino data in order to constrain the neutrino mixing parameters and mass-squared differences (henceforth, simply “PMNS” or “oscillation” parameters) of the three known active neutrinos, assuming that they interact as in the SM. Among these, there is the “Bari group”, whose most recent analysis at the time of this writing is found in Ref. [305], the “NuFIT group” [306, 307], with their latest analysis in Refs. [308, 309], and the “Valencia group” [310], whose latest analysis is found

in Ref. [311]. Experiments on neutrino flavour conversion from all four types of sources enumerated in the previous section enter into each global fit. In order to infer the absolute neutrino mass scale, data from β -decay and neutrinoless double β -decay is also included sometimes, as are cosmological observations (for example in Refs. [305, 311]).

Recently, some of the authors from the NuFIT group performed a global analysis of oscillation data with NSI couplings to up and down quarks serving as nuisance parameters [28]. A minor update was later published in Ref. [29] in order to also include δ_{CP} -sensitive $\bar{\nu}_e$ and $\bar{\nu}_\mu$ appearance data from the LBL accelerator experiments T2K and NO ν A. Omitting the latter, the reverse configuration—with PMNS parameters serving as nuisance parameters in the NSI measurement—was also studied by assuming vector NSI with up and down quarks to be present in oscillation experiments and in CE ν NS at COHERENT [223, 312]. Finally, the most recent and significantly expanded global analysis in Ref. [30] accounts for vector and axial-vector NSI couplings to both quarks and electrons. With respect to the included oscillation experiments, while vector NSI with electrons affect the matrix of effective neutrino potentials due to coherent forward scattering, vector or axial-vector NSI with electrons and axial-vector NSI with quarks are assumed to give rise to modified detection cross sections for solar neutrinos, provided the new mediator is sufficiently massive compared to the typical four-momentum transfers of the scattering processes.

For definiteness, the scenario of neutrino interactions as in the SM in Sec. 5.2.1 is examined based on the global fit “NuFIT 5.2 with Super-Kamiokande (SK) atmospheric data” in Refs. [308, 313]. Here, the matrix of effective neutrino potentials due to coherent forward scattering for neutrino oscillation experiments is hence assumed to be given by Eq. (3.50).

The possible presence of NSI is taken into account in Sec. 5.2.2. First, resulting global constraints on PMNS parameters are discussed based on Refs. [28, 29]. Then, the global NSI constraints from Ref. [223] are reviewed, which is the most up-to-date global fit to explicitly construct confidence regions applying to the parameterisation of Eq. (4.19).

5.2.1 Bounds on PMNS parameters with standard matter potential

Table 5.1 shows the current global best-fit values (“bfp”) and 1σ and 3σ ranges—equivalent to 68.3% respectively 99.7% confidence intervals—for the oscillation parameters θ_{12} , θ_{13} , θ_{23} , δ_{CP} , Δm_{21}^2 , and $\Delta m_{32(1)}^2$ as obtained with experimental data available before November 2022 [308, 313], when neutrino interactions in matter are

5 Global knowledge of the PMNS paradigm and NSI

	Normal Ordering (best fit)		Inverted Ordering ($\Delta\chi^2 = 6.4$)		
	bfp $\pm 1\sigma$	3σ range	bfp $\pm 1\sigma$	3σ range	
with SK atmospheric data	$\sin^2 \theta_{12}$	$0.303^{+0.012}_{-0.012}$	$0.270 \rightarrow 0.341$	$0.303^{+0.012}_{-0.011}$	$0.270 \rightarrow 0.341$
	$\theta_{12}/^\circ$	$33.41^{+0.75}_{-0.72}$	$31.31 \rightarrow 35.74$	$33.41^{+0.75}_{-0.72}$	$31.31 \rightarrow 35.74$
	$\sin^2 \theta_{23}$	$0.451^{+0.019}_{-0.016}$	$0.408 \rightarrow 0.603$	$0.569^{+0.016}_{-0.021}$	$0.412 \rightarrow 0.613$
	$\theta_{23}/^\circ$	$42.2^{+1.1}_{-0.9}$	$39.7 \rightarrow 51.0$	$49.0^{+1.0}_{-1.2}$	$39.9 \rightarrow 51.5$
	$\sin^2 \theta_{13}$	$0.02225^{+0.00056}_{-0.00059}$	$0.02052 \rightarrow 0.02398$	$0.02223^{+0.00058}_{-0.00058}$	$0.02048 \rightarrow 0.02416$
	$\theta_{13}/^\circ$	$8.58^{+0.11}_{-0.11}$	$8.23 \rightarrow 8.91$	$8.57^{+0.11}_{-0.11}$	$8.23 \rightarrow 8.94$
	$\delta_{\text{CP}}/^\circ$	232^{+36}_{-26}	$144 \rightarrow 350$	276^{+22}_{-29}	$194 \rightarrow 344$
	$\frac{\Delta m_{21}^2}{10^{-5} \text{ eV}^2}$	$7.41^{+0.21}_{-0.20}$	$6.82 \rightarrow 8.03$	$7.41^{+0.21}_{-0.20}$	$6.82 \rightarrow 8.03$
	$\frac{\Delta m_{3\ell}^2}{10^{-3} \text{ eV}^2}$	$+2.507^{+0.026}_{-0.027}$	$+2.427 \rightarrow +2.590$	$-2.486^{+0.025}_{-0.028}$	$-2.570 \rightarrow -2.406$

Table 5.1: Summary of global constraints on three-neutrino mixing parameters and mass-squared differences from NuFIT 5.2 [308, 313]. Table taken from Ref. [313].

assumed to be governed by the SM. The left column shows the results under the NO fit assumption, while the right column shows those under the IO. In each case, the best fit corresponds to the minimum of a weighted-least-squares test statistic.⁶ The various parameter ranges are derived relative to the local minimum under the respective NMO hypothesis, after the projection of the multidimensional test-statistic surface onto the parameter of interest.⁷ The IO is disfavoured by 6.4 test-statistic units with respect to the NO, which can be considered a modest indication for the latter. In fact, the preference for the NO has remained rather close to the 3σ (when SK is excluded: 2σ) level over the preceding analysis iterations starting with NuFIT 4.1 [314, 315] in 2019.

The 3σ ranges for all parameters within the NO from Table 5.1 are strictly equivalent to the ones encountered when the NMO itself is treated as an unknown. Based on the definition of the relative precision as $2(x_+ - x_-)/(x_+ + x_-)$, where x_+ is a given parameter's upper bound at the desired confidence level (CL) and x_- its corresponding lower bound, the atmospheric mass-squared difference $\Delta m_{3\ell}^2$ ($\ell = 1$ for NO and $\ell = 2$ for IO) is the most precisely measured parameter. At the 3σ CL, it is known

⁶The test statistic from which confidence intervals are calculated is suggestively denoted as “ $\Delta\chi^2$ ” by the authors of Refs. [308, 313], owing to the assumption of it being approximately χ_m^2 -distributed. More details can be found in Sec. 13.2.

⁷A commonly employed method for the construction of frequentist confidence regions is examined in Sec. 13.3.

to a relative precision of approximately 6%. For θ_{13} one finds a slightly worse precision of 8%, for θ_{12} one of 13%, and for θ_{23} one of 25%. The solar mass-squared difference Δm_{21}^2 is known to a precision of 16%, and δ_{CP} is only weakly constrained to a precision of 83%.

When the allowed ranges for the neutrino mixing parameters in Table 5.1 are combined following the procedure of Ref. [316], the magnitudes of the elements of the PMNS matrix (unitarity preimposed) are restricted to the following ranges at the 3σ CL [308, 313]:

$$|U|_{3\sigma} = \begin{pmatrix} 0.803 \rightarrow 0.845 & 0.514 \rightarrow 0.578 & 0.143 \rightarrow 0.155 \\ 0.244 \rightarrow 0.498 & 0.502 \rightarrow 0.693 & 0.632 \rightarrow 0.768 \\ 0.272 \rightarrow 0.517 & 0.473 \rightarrow 0.672 & 0.623 \rightarrow 0.761 \end{pmatrix}. \quad (5.29)$$

Note that the inclusion of oscillation data up to March 2024 in NuFIT 5.3 [309] has very little impact on the preferred parameter values and confidence intervals quoted above.

5.2.2 Simultaneous bounds on PMNS and NSI parameters

The past four decades have been characterised by important phenomenological progress in the interpretation of the available experimental data (reviewed, e.g., in Ref. [262]) in terms of NSI in neutrino propagation. Above all, it was found that both flavour-changing and flavour-diagonal interaction types could play a significant role in the MSW solution of the solar neutrino problem, both in the massive- and in the massless-neutrino scenario [317, 318, 246, 248, 247, 319, 320, 321].

Immediately following the observation of a decisive deficit of upgoing atmospheric muon neutrinos by Super-Kamiokande in 1998 [13], NSI began to be investigated as a possible explanation thereof [322, 256, 323, 324, 255], becoming more strongly ruled out as the exclusive cause of the deficit with the accumulation of several years of atmospheric neutrino data. The authors of Refs. [325, 253] subsequently performed a novel three-flavour analysis of NSI in the e - τ sector based on atmospheric (“Super-Kamiokande phase I”) and accelerator (“KEK to Kamioka”) neutrino data, and found much relaxed bounds as compared to the earlier studies, which had only considered a two-flavour scenario in the μ - τ sector.

Today, with increasing experimental sensitivity to fundamental unknowns such as the NMO, δ_{CP} , or the octant and non-maximality of θ_{23} , it is of essence to ensure that the emerging indications are not actually brought about by new physics that

has so far been unaccounted for. With regards to these three major unknowns, the most comprehensive global analysis to date in Ref. [28] (based on the same datasets as NuFIT 4.0 [314, 326]) respectively its minor update in Ref. [29] (same datasets as NuFIT 4.1 [315]) fits for the three mixing angles, the two mass-squared differences, and the CP-violating phase within the PMNS framework, as well as for the most general matter potential due to NSI. It reports constraints derived exclusively from oscillation experiments for δ_{CP} , $\sin^2 \theta_{23}$, and the NMO, which thus apply for NSI mediator masses exceeding R_{\oplus}^{-1} (cf. Sec. 4.5). In the case of $|\Delta m_{32(1)}^2|$ and all NSI couplings, however, only results with the fit restricted to select regions of the parameter space are provided, which are moreover derived using data from COHERENT (cf. Sec. 2.5.4). Their validity is thus limited to NSI models with mediator masses above $\mathcal{O}(10 \text{ MeV})$. Not least to avoid such a model dependence in this discussion, we additionally review the global fit results from Ref. [223] (similar datasets as NuFIT 3.2 [327, 328]). The latter only uses datasets with limited or without sensitivity to CP-violating effects and therefore only draws inferences about CP-conserving parameters in the full neutrino Hamiltonian, both from oscillation experiments only and from their combination with COHERENT.

Both studies above draw upon the same general analysis methodology [223], which makes the assumption that the NSI flavour structure is independent of the type of charged background fermion $f = u, d$, so that the following factorisation applies to the NC NSI coupling strengths:

$$\epsilon_{\alpha\beta}^{fV} \equiv \epsilon_{\alpha\beta} \xi^f, \quad (5.30)$$

where $\epsilon_{\alpha\beta}$ is a matrix of universal couplings, which are independent of f and rescaled by ξ^f . The factorisation encompasses the special case of equal NSI couplings to up and down quarks, i.e., $\xi^u = \xi^d$, which is a prediction made for example by many SM extensions with a light Z' associated with a new $U(1)'$ gauge symmetry (cf. Sec. 4.5), or the similarly well-motivated case of vanishing NSI with either up or down quarks, i.e., $\xi^u = 0$ or $\xi^d = 0$ (see Ref. [30] and references therein). Couplings to electrons are not explicitly considered, without any loss of generality for neutrino propagation according to Eq. (4.9).

An important distinction between the analysis of solar experiments (and KamLAND) on the one hand and that of atmospheric and LBL (and trivially MBL reactor) experiments on the other is that the latter can be performed in terms of effective NSI parameters for Earth matter: the Earth's neutrino matter potential behaves as if it were “composed of a unique effective fermion” [254], which is expressed

formally through Eq. (4.18). The Sun’s chemical composition, however, exhibits substantial variations along the neutrino trajectory, with the consequence that a universal effective coupling strength which could serve as the phenomenologically relevant parameter cannot be defined. Instead, in addition to the eight independent real parameters related to $\epsilon_{\alpha\beta}$ (cf. Sec. 4.3), the authors define the direction in the ξ^p – ξ^n plane via an angle η that uniquely fixes the relation between $\epsilon_{\alpha\beta}^{pV}$ and $\epsilon_{\alpha\beta}^{nV}$:

$$\xi^p = \sqrt{5} \cos \eta, \quad \xi^n = \sqrt{5} \sin \eta, \quad (5.31)$$

with $-\pi/2 \leq \eta \leq \pi/2$. Only in the special case $\eta = 0$, when all $\epsilon_{\alpha\beta}^{nV} = 0$, does the overall matter potential have the same structure as Eq. (4.18), and, e.g., an exact cancellation of the SM matter potential could occur at each point along the trajectory.

Neutrino mass ordering, 2–3 mixing, and leptonic CP-violating phase

Given the aforementioned methodology, a least-squares function fit to the observables from LBL experiments (MINOS, T2K, and NO ν A) is evaluated in the full 15-dimensional parameter space, whereas the observables from other neutrino oscillation experiments (solar, KamLAND, MBL reactor, and atmospheric) are fit assuming a reduced 11-dimensional, CP-conserving, parameter space. In addition, matter effects are neglected in MBL reactor experiments due to the short oscillation baselines, the oscillation parameters θ_{12} , θ_{13} , and Δm_{21}^2 are fixed in the fit to LBL data (with both octants of θ_{12} explicitly tested), and the constraints on θ_{23} and Δm_{31}^2 from atmospheric neutrino experiments are neglected. In contrast to NuFIT in Sec. 5.2.1, the data analysis including NSI does not incorporate more recent Super-Kamiokande datasets optimised for the NMO determination.

The best fit of Ref. [28] (and of its minor update [29])—resulting in the minimal test-statistic value when all oscillation experiments are taken into account—is located in the parameter space of the NO with θ_{12} in the “light” octant (“light-NO”).⁸ The three other combinations of the octant of θ_{12} and the NMO are disfavoured by different degrees of confidence, namely “dark-IO” by approximately 2 units of $\Delta\chi^2$, “light-IO” by 3.5, and “dark-NO” by 5.5 units with respect to the global minimum. The “dark” octant is hence rendered a statistically acceptable solution in the presence of NSI. By lifting the GMOD in Eq. (4.20), solar neutrino data is solely responsible

⁸The same main conclusions, only with stated exclusion levels increased by $\Delta(\Delta\chi^2) \sim 1$ to 2, are reached when LBL antineutrino appearance data is additionally included (whether COHERENT is considered or not only has a minor impact) [29].




Figure 5.11: $\Delta\chi^2$ as a function of $\sin^2 \theta_{23}$ (left) and δ_{CP} (right) from fits to global neutrino oscillation data assuming the presence of NSI [28]. The four curves labelled “light-NO”, “dark-NO”, “light-IO”, “dark-IO” have been obtained by optimising the respective undisplayed standard oscillation and NSI parameters while fixing the NMO and octant of θ_{12} to the specified domains. Shown in addition are the $\Delta\chi^2$ projections under the assumption of the SM matter potential, obtained using the same global data (NuFIT 4.0) [314, 326].

for the weak indications in favour of the NO and, equivalently, the “light” octant.

The $\Delta\chi^2$ offsets separating the four parameter domains can be read off from Fig. 5.11, which shows the $\Delta\chi^2$ function, relative to the global minimum, in the presence of NSI projected onto $\sin^2 \theta_{23}$ or δ_{CP} . Also included are the projections of the fit to the same data in the absence of NSI (NuFIT 4.0), analogously after subtracting the fit’s global test-statistic minimum. Comparing the two global analyses reveals a somewhat reduced preference for the NO when NSI are accounted for. The statistical exclusion level of maximal 2–3 mixing is almost unaffected, but the first octant of θ_{23} is rendered more acceptable, both overall and under each individual NMO hypothesis. The most pronounced impact of NSI on the determination of δ_{CP} is the reduction of the exclusion level of maximal CP violation with $\delta_{\text{CP}} = 3\pi/2$ in the favoured (“light-NO”) parameter domain. Vanishing intrinsic CP violation remains excluded at a similar confidence level as compared to the scenario without NSI.

NSI couplings and generalised matter potential

The global fit in Ref. [223] reports results for the GMP parameterisation (4.19), which has a variety of features that are conducive to performing a combined ana-

lysis of atmospheric and LBL accelerator neutrino oscillation data in the framework of NSI [225]. First, one can impose the equality of two eigenvalues of the matter potential—realised by fixing $\epsilon'_{\oplus} = 0$ —because it causes the disappearance of atmospheric $\bar{\nu}_{\mu}$'s with energies $E_{\nu} \gtrsim 10$ GeV to proceed with the same L/E_{ν} dependence as in the case without NSI, i.e., to mimic the oscillation pattern in vacuum [325]. When appropriate datasets are analysed, the NSI constraints obtained within this restricted subspace are therefore expected to be a good proxy of those in the most general parameter space [225, 254, 223, 30]. Adopting this subspace renders the angle φ_{23} and the CP-violating NSI phase δ_{NS} unphysical. Together with the approximation $\Delta m_{21}^2 = 0$, which additionally renders θ_{12} and δ_{CP} unphysical, the assumption of a single non-degenerate eigenvalue further implies an effective 2ν flavour evolution that only depends on the parameters Δm_{31}^2 , θ_{13} , θ_{23} , ϵ_{\oplus} , φ_{12} , φ_{13} , α_1 , and α_2 [329, 225]. Moreover, a CP-conserving formalism is implemented by setting $\alpha_1 = \alpha_2 = 0$, which reduces the number of parameters required to describe the oscillations of atmospheric and LBL accelerator neutrinos to six, all of which are real. A given point in the space of ϵ_{\oplus} , φ_{12} , and φ_{13} then yields the following NU strengths and FV couplings in the standard parameterisation:

$$\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} = \epsilon_{\oplus} (\cos^2 \varphi_{12} - \sin^2 \varphi_{12}) \cos^2 \varphi_{13} - 1, \quad (5.32)$$

$$\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} = \epsilon_{\oplus} (\sin^2 \varphi_{13} - \sin^2 \varphi_{12} \cos^2 \varphi_{13}), \quad (5.33)$$

$$\epsilon_{e\mu}^{\oplus} = -\epsilon_{\oplus} \cos \varphi_{12} \sin \varphi_{12} \cos^2 \varphi_{13}, \quad (5.34)$$

$$\epsilon_{e\tau}^{\oplus} = -\epsilon_{\oplus} \cos \varphi_{12} \cos \varphi_{13} \sin \varphi_{13}, \quad (5.35)$$

$$\epsilon_{\mu\tau}^{\oplus} = \epsilon_{\oplus} \sin \varphi_{12} \cos \varphi_{13} \sin \varphi_{13}. \quad (5.36)$$

To allow for the most general real NSI hypotheses within this reduced parameter space, the ranges of the Euler rotation angles have to be chosen as $-\pi/2 \leq \varphi_{ij} \leq \pi/2$ [225]. Also, only the relative sign of ϵ_{\oplus} and Δm_{31}^2 is physical [225], implying that one can restrict the atmospheric mass-squared difference to positive values while allowing both positive and negative values of ϵ_{\oplus} . The choice $\varphi_{12} = 0 = \varphi_{13}$ with $\epsilon_{\oplus} = +1$ then corresponds to standard oscillations given NO, and that with $\epsilon_{\oplus} = -1$ to standard oscillations given IO.

Figure 5.12 shows the regions in the Earth matter parameters ϵ_{\oplus} , φ_{12} , and φ_{13} that are compatible with three different combinations of global data at different confidence levels. For each parameter pair, the test statistic is optimised over all undisplayed parameters. No matter whether one considers only atmospheric, LBL, and MBL reactor experiments, or removes IceCube and reactor experiments, or considers

the full set of experiments, SM interactions with $\epsilon_{\oplus} = 1$ (and $\varphi_{12} = \varphi_{13} = 0$) are contained within all confidence regions. In other words, no indications for (CP-conserving) NSI can currently be inferred from the globally available neutrino data. At the same time, the full global analysis clearly allows placing limits on all three matter parameters, even at the 3σ CL.

A prominent feature of the two analyses not including solar and KamLAND data (green and yellow hatched regions) is that the flavour structure of the GMP is unconstrained at the 3σ CL: the full φ_{12} – φ_{13} plane is allowed. This is the result of the optimisation over small values of $|\epsilon_{\oplus}| \lesssim 0.1$, as gauged from the two rightmost panels. In the vacuum limit $|\epsilon_{\oplus}| = 0$ the two rotation angles φ_{12} and φ_{13} are no longer observable, rendering the fit independent of their assumed values. The atmospheric, LBL, and MBL reactor data is compatible with the vacuum scenario at the 3σ CL for two degrees of freedom. As a consequence, no NSI flavour structure is decisively rejected. Contrariwise, when large values $|\epsilon_{\oplus}| \rightarrow 10^2$ are assumed, the 3σ bounds $\varphi_{12} \lesssim 15^\circ$ and $\varphi_{13} \lesssim 20^\circ$ emerge. Furthermore, the overall strength ϵ_{\oplus} of the Earth matter potential can neither be bounded from below or above, because “the considered data sample is mainly sensitive to NSI through $[\vec{\nu}_{\mu}]$ disappearance, and lacks robust constraints on matter effects in the $[\vec{\nu}_e]$ sector” [223]. The mentioned limitations disappear once solar and KamLAND data is included.

Evaluating the same test statistic with all datasets included in dependence of the NU and FV NSI parameters in Earth matter that can be constrained by oscillation experiments yields one-dimensional $\Delta\chi^2$ projections onto $\epsilon_{ee(\tau\tau)}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$, $\epsilon_{e\mu}^{\oplus}$, $\epsilon_{e\tau}^{\oplus}$, and $\epsilon_{\mu\tau}^{\oplus}$, from which 90 % confidence intervals can be derived [223, Fig. 9]. The extracted intervals are summarised in Table 5.2. The most precisely determined parameter is the μ – τ FV coupling $\epsilon_{\mu\tau}^{\oplus}$, followed by $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ and $\epsilon_{e\mu}^{\oplus}$, whose magnitudes may amount to at most 2 %, 7 %, and 10 %, respectively, of Fermi’s coupling strength. The confidence intervals on $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ and $\epsilon_{e\tau}^{\oplus}$ both exhibit a notable asymmetry, with a higher precision for negative values in the former case, and for positive values (corresponding to the phase $\delta_{e\tau} = 0$) in the latter. The (global) test-statistic minimum is located in the “light” θ_{12} octant and at $\eta \approx -44^\circ$ [223].⁹ The intervals in Table 5.2 do not extend into the “dark” octant, because it is disfavoured at greater than 90 % confidence.

⁹Coincidentally, this best fit is located close to $\eta \approx -43.6^\circ$, where the NSI contribution to the Earth’s matter potential vanishes: at this point, the nucleon couplings $\epsilon_{\alpha\beta}^{pV}$ and $\epsilon_{\alpha\beta}^{nV}$ cannot be separately constrained, which becomes evident by combining Eqs. (5.30) and (5.31) with the definition (4.12). In the Sun, the interval $-70^\circ \lesssim \eta \lesssim -60^\circ$ gives rise to the weakest (but not vanishing) NSI constraints.

Parameter	90 % CL interval
$\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$	$[-0.24, 0.25]$
$\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$	$[-0.014, 0.07]$
$\epsilon_{e\mu}^{\oplus}$	$[-0.10, 0.10]$
$\epsilon_{e\tau}^{\oplus}$	$[-0.26, 0.06]$
$\epsilon_{\mu\tau}^{\oplus}$	$[-0.015, 0.021]$

Table 5.2: 90 % confidence intervals on the real NSI parameters for Earth matter from a global fit to CP-conserving neutrino oscillation data [223].

Reference [223] also underlines the robustness of the global knowledge of the PMNS parameters Δm_{21}^2 , $|\Delta m_{31}^2|$, θ_{13} , and θ_{23} even in the presence of (CP-conserving) NSI. A reduction of the precision with which these are determined is only encountered when synergies between the different types of oscillation experiments are not fully exploited. For instance, omitting solar and KamLAND as well as MBL reactor data from the global NSI fit renders large values $|\Delta m_{31}^2| \sim 3 \times 10^{-3} \text{ eV}^2$ acceptable and severely impedes its joint determination with θ_{23} [223, Fig. 6].




Figure 5.12: Two-dimensional confidence regions in pairs of the GMP parameters ϵ_{\oplus} , φ_{12} ($^{\circ}$), φ_{13} ($^{\circ}$) from a global fit to neutrino oscillation data without sensitivity to CP-violating effects [223]. The light and dark hatched green patches cover the regions allowed by atmospheric, LBL, and MBL reactor data at 90% and 3σ CL, respectively. The yellow hatched regions (same confidence levels) follow from omitting IceCube and reactor data. A fit to all datasets (including solar and KamLAND) combined yields the coloured 1σ , 90%, 2σ , 99%, and 3σ confidence regions. The star is located at the best-fit point of the full analysis. Figure taken from Ref. [223]. See text for details.

Part II

Atmospheric neutrinos in IceCube: present and future

6

Dual role of glacial ice in IceCube

The glacial ice at the South Pole in which IceCube is embedded serves two major physics purposes: it provides a target medium for the interactions of neutrinos with energies spanning many orders of magnitude and simultaneously radiates optical Cherenkov photons due to the passage of charged secondaries produced by the neutrino interaction. The bulk properties of the ice not only determine the probability for an individual neutrino to interact but also affect the propagation of secondary particles. Owing to the natural origin of the ice, its optical properties deviate from those of an ideal medium and have to be carefully calibrated in situ. Since the deployment of the “AMANDA-A” prototype detector in 1993 and 1994 [330], efforts to this end have been pursued by first the AMANDA and later the IceCube collaboration. Today, the microscopic structure of the ice and its impact on photon propagation over macroscopic scales remain active subjects of investigation [331].

6.1 Ice as target medium

For a given impinging neutrino rate, Eq. (2.36) shows that the interaction rate per unit volume is proportional to the product of the cross section for incoherent scattering on ice nuclei and their number density. The surveyable ice volume places an upper limit on the rate of detectable neutrino interactions. For a first estimate, consider a cube of ice with a volume of $V_{\text{ice}} = 1 \text{ km}^3$ subject to the total incident flux of atmospheric neutrinos in the energy range 0.1 GeV to 10^4 GeV at the South Pole, $\Phi_{\text{tot}} \sim \mathcal{O}(10^5 \text{ cm}^{-2} \text{ s}^{-1})$ [281]. Using the nucleon density $n_{\text{ice}} \approx 6 \times 10^{23} \text{ cm}^{-3}$, the mass density $\rho_{\text{ice}} \approx 0.92 \text{ g cm}^{-3}$, and a typical neutrino-nucleon cross section of $\sigma_{\text{tot}} \sim \mathcal{O}(10^{-38} \text{ cm}^2)$ then yields a neutrino interaction rate of the order of 10^6 s^{-1} .

6.2 Ice as detection medium

Only an ideal detector would be able to observe the above interaction rate. In practice, considerable efficiency losses can occur, even if the instrumented (detector) volume V_{det} and the target volume V_{ice} fully coincide. Such losses are encapsulated in the detector’s “effective” volume V_{eff} , which quantifies the fraction of V_{ice} that is surveyed effectively. It usually has to be determined from simulation. V_{eff} is bounded from above by the target volume V_{ice} , not the instrumented volume V_{det} . A detector is referred to as “fully efficient” (e.g., above some sufficiently high neutrino energy) when $V_{\text{eff}} \geq V_{\text{det}}$, or when effectively every neutrino interaction occurring within the instrumented volume is detected. V_{eff} is converted into a corresponding effective area A_{eff} via the relation (any dependence on neutrino energy or arrival direction suppressed)

$$A_{\text{eff}} \equiv \sigma_{\text{tot}} n_{\text{ice}} V_{\text{eff}}. \quad (6.1)$$

Instrumenting the ice with sensors for Cherenkov radiation allows not only to detect a given neutrino scattering process (with some efficiency $V_{\text{eff}}/V_{\text{ice}} \leq 1$) but also to infer properties of the neutrino itself, by measuring the “visible” energy deposited in the ice and tracking charged daughter particles. Hence, the detector serves as an absorption calorimeter with tracking capabilities.

6.2.1 Cherenkov emission

Particle	E_{th} (GeV)
e^{\pm}	7.83×10^{-4}
μ^{\pm}	1.62×10^{-1}
τ^{\pm}	2.72
\bar{p}	1.44
π^{\pm}	2.14×10^{-1}
K^{\pm}	7.56×10^{-1}

Table 6.1: Total energy E_{th} corresponding to the Cherenkov threshold for a refractive index of $n = 1.33$ for leptons and various hadrons.

Any charged particle with a velocity β exceeding the local phase velocity of light, c/n , emits a cone of Cherenkov radiation with a half-angle of opening (Cherenkov angle) of $\theta_C = \arccos\left(\frac{1}{n\beta}\right)$, which is azimuthally symmetric around the direction of particle propagation [15]. Since the refractive index n is a function of frequency ω in a dispersive medium, θ_C also depends on ω —given the Cherenkov condition $\beta > c/n(\omega)$. In the wavelength range $\lambda \sim 0.2 \mu\text{m}$ to $1 \mu\text{m}$, in ice n slowly decreases from 1.39 to 1.30 [332]. For $n = 1.33$ (corresponding to $\lambda \approx 300 \text{ nm}$), Table 6.1 specifies the total energy E_{th} above which various charged particles emit Cherenkov

photons.

The number of Cherenkov photons radiated per unit path length and wavelength interval by a particle with the charge ze is given by [15]

$$\frac{\partial^2 N}{\partial x \partial \lambda} = \frac{2\pi\alpha z^2}{\lambda^2} \left(1 - \frac{1}{\beta^2 n^2(\lambda)} \right) = \frac{2\pi\alpha z^2}{\lambda^2} \sin^2 [\theta_C(\lambda)]. \quad (6.2)$$

A unit charge with $\beta = 1$ radiates approximately 250 optical Cherenkov photons for each cm of path length in ice [333].

Energy loss processes of massive secondaries in the ice determine the spatial extent over which Cherenkov photons are emitted and impact the photon spectrum and overall yield as compared to that of a $\beta = 1$ charge, according to Eq. (6.2). Losses due to the Cherenkov emission itself are at the subpercent level and small compared to those due to ionisation even for a minimum-ionising particle [334]. The optical properties of the ice dominate the shapes of the photon arrival time distributions at the photosensor locations.

6.2.2 Energy loss processes

As detailed in Sec. 2.5.3, at energies above a few GeV neutrino-nucleon interactions proceed predominantly through DIS. According to Eq. (2.51), both CC and NC interactions result in a system of final-state hadrons, which in turn gives rise to a hadronic cascade. A charged lepton emerges in addition in the case of a CC interaction. Sketches of the four interaction categories are shown in Fig. 6.1. The outgoing lepton at the vertex of the original neutrino interaction is highlighted in each case, but individual cascade constituents are not resolved. This reflects the fact that the cascade extents are small compared to the typical photosensor spacing.

μ^\pm energy losses Among the charged leptons, only μ^\pm 's are able to propagate and deposit energy over distances exceeding a few metres. Their mean energy loss rate in a given material can be viewed as a function of β alone and exhibits a broad minimum dominated by continuous ionisation energy losses for $\beta\gamma \sim 1$ to 10^3 (or a kinetic energy $T_\mu \sim 0.1$ GeV to 100 GeV) [15]. For a minimum-ionising μ^\pm in ice, the mean energy loss rate is given by [335]

$$\left\langle -\frac{dE_\mu}{dx} \right\rangle_{\text{ice}} \approx 0.238 \text{ GeV m}^{-1}. \quad (6.3)$$

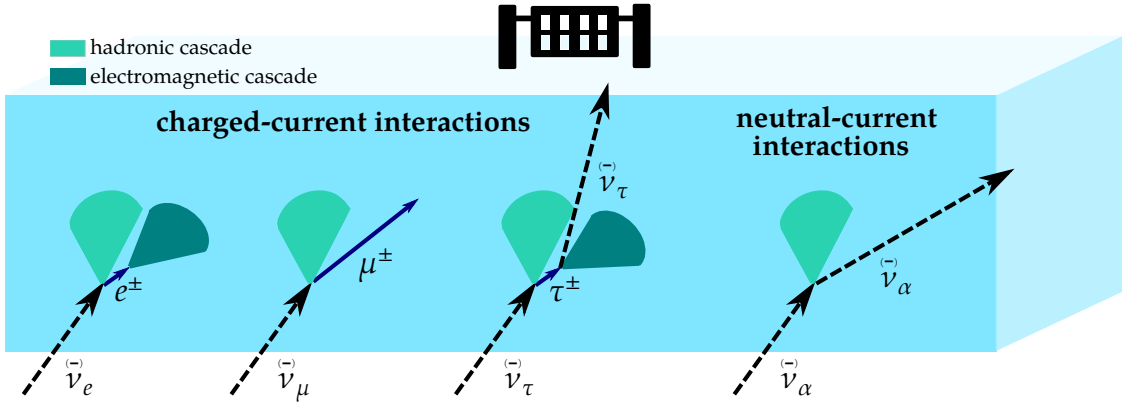


Figure 6.1: Sketch of some typical atmospheric neutrino interaction signatures in ice (not to scale), cf. Eq. (2.51). The three leftmost processes represent CC interactions, which produce a hadronic cascade and a charged lepton. The rightmost process represents NC interactions, which produce a hadronic cascade and an outgoing neutrino.

In the above kinematic regime, radiative processes—bremsstrahlung, pair production, or photonuclear interactions—contribute negligibly to the overall energy loss rate.

e^\pm energy losses An e^\pm with the energy of a minimum-ionising μ^\pm initiates an electromagnetic cascade [15] because the (“critical”) energy above which the loss rate due to bremsstrahlung exceeds that due to ionisation is given by $E_{c,\text{ice}}^{e^-} = 78.6$ MeV respectively $E_{c,\text{ice}}^{e^+} = 76.5$ MeV [336].

The characteristic distance scale that governs the longitudinal development of the cascade in ice is the “radiation length” $X_0 = 39.3$ cm [336], which applies to the energy losses of high-energy e^\pm ’s (emitting bremsstrahlung) and photons (producing e^+e^- pairs). To a first approximation [337, 338], the cascade contains a maximum of $N_{\text{max}} = E_0/E_{c,\text{ice}}$ particles, where E_0 is the energy of the primary e^\pm . The corresponding depth of the cascade along its longitudinal axis depends on E_0 logarithmically, $\propto X_0 \ln(E_0/E_{c,\text{ice}})$. This simple dependence of the depth of the shower maximum on E_0 is confirmed by realistic simulations [333], according to which for example a primary e^\pm of energy $E_0 = 10$ GeV results in a maximum energy deposition at a depth of $z_{\text{max}} \approx 1.55$ m.

A large fraction ϵ_e^{vis} of the energy deposited in the electromagnetic cascade is “visible” via Cherenkov emission. The total Cherenkov-radiating (“effective”) track length of an electromagnetic cascade linearly depends on E_0 , with an average constant of proportionality of ~ 5 m GeV $^{-1}$ and with subpercent statistical fluctuations at the GeV energy scale [333].

While an electromagnetic cascade may also be initiated by a decaying μ^\pm [339], its lifetime $\tau_\mu = 2.2 \mu\text{s}$ causes a minimum-ionising μ^\pm to lose a large fraction of its energy prior to the decay. The energy of the resulting cascade is therefore typically negligibly small compared to the overall energy deposited in the ice. This is why no μ^\pm decay is included in Fig. 6.1.

τ^\pm energy losses As the heaviest of the charged leptons, the mean range of a τ^\pm in the considered energy range is just given by its mean decay length, which reaches the cm scale at most. The τ^\pm decays leptonically with a branching ratio of approximately 35%, to approximately equal parts into e^\pm and μ^\pm , under the emission of a neutrino and an antineutrino [15]. In nearly two thirds of cases, the τ^\pm decays hadronically, most likely via the process $\tau^\pm \rightarrow \pi^\pm + \pi^0 + \bar{\nu}_\tau$ (branching ratio $\sim 25\%$). As a result, a $\bar{\nu}_\tau$ CC interaction gives rise to a hadronic cascade and either another hadronic cascade, an electromagnetic cascade, or a low-energy μ^\pm . Independent of the decay channel, some fraction of the energy of the primary $\bar{\nu}_\tau$ is undetectable because it is carried away by at least one outgoing (anti)neutrino. Figure 6.1 shows a hadronic τ^\pm decay.

Hadronic energy losses Hadronic cascades accompany all of the above interactions, at the primary interaction vertex itself and possibly at the τ^\pm decay vertex in the case of $\bar{\nu}_\tau$ CC scattering. The hadronic system produced at such a vertex causes a cascade of hadronic interactions, with a rapidly multiplying number of secondaries. The characteristic distance scale between subsequent hadronic interactions is given by ~ 90 cm in ice, which is only weakly dependent on energy [336].

In general, a hadronic cascade consists of a hadronic and an electromagnetic sector, summarised in Fig. 6.2: during the development of the cascade, π^0 production removes energy from the hadronic sector through the fast electromagnetic decay $\pi^0 \rightarrow \gamma\gamma$ [15]. This results in the overall energy transfer of $E_{\pi^0} = f_{\pi^0} E_0$ throughout all particle generations within the cascade. In Ref. [340], the phenomenological model

$$f_{\pi^0}(E_0) = 1 - \left(\frac{E_0}{E'}\right)^{m-1} \quad (6.4)$$

describes the mean energy fraction in the electromagnetic sector, where E' and m are model parameters that depend on the calorimeter material and the incident hadron type. The dedicated simulation study in Ref. [341]—performed for ice and based on realistic energy spectra and species distributions for the hadronic particle system induced by a neutrino NC interaction—finds that $m = 0.87$ and $E' = 0.399$ GeV yield

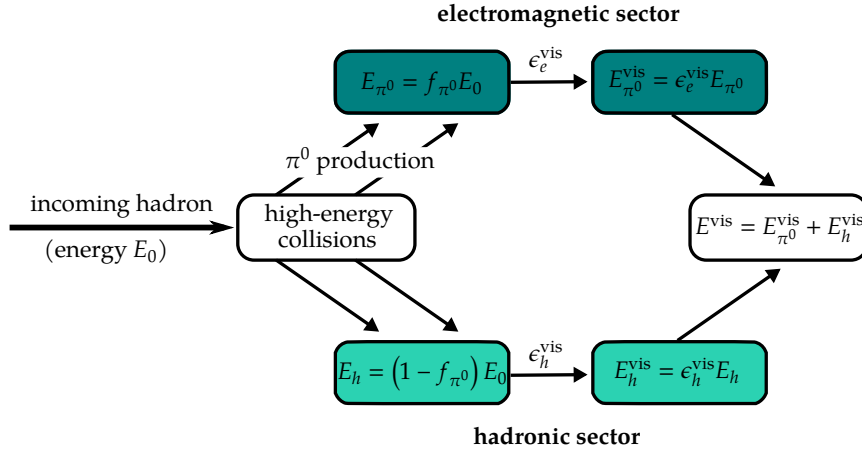


Figure 6.2: Energy flow within a hadronic cascade initiated by a high-energy hadron with the incident energy E_0 . On average, the fraction $f_{\pi^0}E_0$ is transferred to the electromagnetic part of the hadronic cascade through π^0 production. Only the fractions ϵ_e^{vis} respectively ϵ_h^{vis} of the energies in the electromagnetic and the hadronic sector are visible. Figure adapted from Ref. [342].

a reasonable fit across many orders of magnitude of the incident energy $E_0 > E'$. If E_0 is below the pion production threshold, all energy remains in the hadronic sector ($E_{\pi^0} = 0$), whereas in the high-energy limit E_0 is fully transferred to the electromagnetic sector ($E_{\pi^0} = E_0$). The particle multiplication in the hadronic sector subsides once energy loss processes such as ionisation, nuclear excitation, spallation, and fission become important.

The relative suppression of the Cherenkov light yield of a hadronic cascade with respect to that of an electromagnetic cascade with the same incident energy, $\epsilon_h^{\text{vis}}(E_0)/\epsilon_e^{\text{vis}}(E_0)$, is governed by the ratio of the two corresponding total effective track lengths. For $E_0 = 100$ GeV, a hadronic cascade generates only approximately 75 % of the amount of Cherenkov light produced by an electromagnetic cascade of the same incident energy [341]. Due to the large variety of possible hadronic physics processes, below the 100 GeV energy scale statistical fluctuations in the Cherenkov light yield of a hadronic cascade amount to some tens of percent [341].

6.2.3 Optical ice properties

Until recently, light propagation in the glacial ice at the South Pole was thought to be fully characterised by the local distributions of the absorption and scattering coefficients, together with their wavelength dependencies, $k_a(\mathbf{x}, \lambda)$ and $k_s(\mathbf{x}, \lambda)$, respectively. While the inverse of the absorption coefficient gives the absorption length λ_a —the distance over which the photon survival probability due to absorption drops

by $1/e$ —the inverse of the scattering coefficient gives the scattering length λ_s —the mean distance between consecutive scattering processes. However, recent findings suggests a more complex photon propagation picture which includes optical anisotropy.

Absorption

For pure ice, there is no known absorption mechanism that is centred in the near-ultraviolet (UV) to visible wavelength range $\lambda \sim 200$ nm to 700 nm: near the lower end of this range lies the electronic UV absorption edge, whereas the tails of the near-infrared (IR) molecular vibration absorption bands extend into this range from above [344]. Both regimes are reflected in the measurements of the absorption coefficient of laboratory-grown ice and of deep glacial ice at the South Pole in Fig. 6.3. The measurements at the South Pole were performed with various pulsed and continuous optical light sources and receivers deployed with precursors of IceCube operated by the AMANDA collaboration.¹ The measurements allowed disentangling light extinction due to absorption from extinction due to scattering by means of analytic or numerical fits to the observed photon delay time distributions [345, 343]. The smallest of the measured absorption coefficients of the ice at the South Pole correspond to absorption lengths of $\lambda_a \sim 200$ m; no other known solid exhibits such a high transparency (see Ref. [343] and references therein). Light-absorbing impurities, such as natural fallout of dust, limit the transparency of the deep glacial ice in the near-UV regime and result in the depth dependence of k_a observed in Fig. 6.3.² The concentration and nature of the impurities as a function of

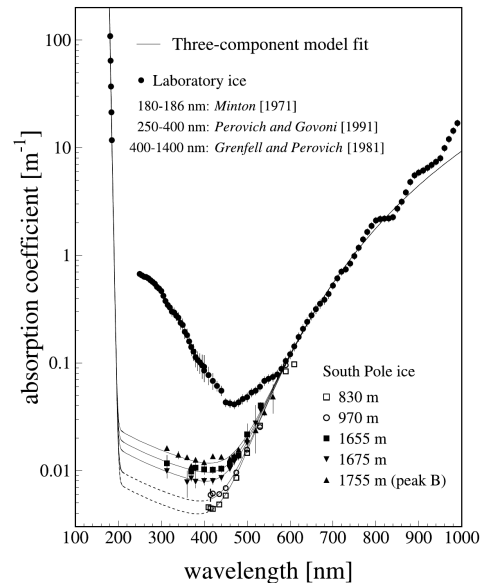


Figure 6.3: Absorption coefficient k_a from UV to IR measured in laboratory ice and deep South Pole ice. Figure taken from Ref. [343].

¹The AMANDA-A prototype detector was installed at depths between 0.8 km and 1.0 km. It was superseded by AMANDA-B, which was installed at depths between 1.55 km and 1.95 km, with some photosensors placed up to 350 m below and above this range. AMANDA-B was decommissioned in 2009.

²As discussed in Ref. [343], the absorption coefficient of ice is estimated to increase with temperature by $1\% \text{ K}^{-1}$, which cannot account for its observed depth dependence in the near-UV range, where

depth reflect the prevailing climate conditions over the time span the glacier formed. By correlating the depth-dependent optical properties with an ice core extracted in East Antarctica, it was inferred that the glacial ice at the South Pole provides a climate record which dates back at least 10^5 years [346].

Scattering

In contrast to the case of absorption, for near-UV and visible light even pure ice has a finite scattering coefficient. It stems from small thermally induced optical inhomogeneities (thermal Rayleigh scattering) or acoustic waves (Brillouin scattering) [347]. The only measurement of the intrinsic Rayleigh scattering coefficient of ice published as of today is that by Raman in 1923, who found $k_s(\lambda = 436 \text{ nm}) = 8 \times 10^{-6} \text{ cm}^{-1}$ at a temperature of $T \approx 270 \text{ K}$ [348, 347].³ Another intrinsic contribution to the scattering coefficient originates from line defects in its crystal structure, or “dislocations” [351]. The associated scattering coefficient for visible light in the deep glacial ice is expected to be subdominant with respect to that of thermal scattering and to have a different wavelength dependence [352, 347]. However, the scattering coefficient is estimated to be approximately four orders of magnitude larger when dislocations are “decorated” with impurity precipitations, $k_s(\lambda = 436 \text{ nm}) \sim \mathcal{O}(10^{-5} \text{ cm}^{-1})$ [347].

More generally, impurities by far dominate the overall scattering coefficient of visible light in the deep ice. In fact, a four-component model consisting of mineral grains, sea salt crystals, acid droplets, and soot was introduced in Ref. [353] to account for the AMANDA measurements of the local scattering and absorption coefficients as well as their wavelength dependencies in the depth interval from 1.6 km to 1.83 km. The model assumes distinct refractive indices, size distributions, and depth-dependent mass concentrations for the constituent types. Since the mean sizes of the impurities range from approximately 12 nm (for soot) to 400 nm (for salt), the scattering properties for visible light lie in the transition regime between Rayleigh scattering and geometric scattering. Hence, Mie scattering theory has served as the method of choice for modelling the optical properties of the ice analytically [353].

Existing scattering measurements in South Pole ice constrain the *effective* scattering

ice has no intrinsic absorptivity.

³Raman attributed the blue colour of glaciers to Rayleigh scattering on ice molecules [348]—in clear contradiction with observation [349, 350]. Instead, ice is intrinsically blue, as is the impure South Pole ice as evidenced by Fig. 6.3.

coefficient k_e . The latter is related to k_s as

$$k_e(\lambda) \equiv k_s(\lambda) \left(1 - \langle \cos \vartheta_s \rangle(\lambda)\right), \quad (6.5)$$

where $\langle \cos \vartheta_s \rangle(\lambda)$ is the mean cosine of the photon scattering angle ϑ_s as a function of wavelength. It is also referred to as the “asymmetry parameter” and denoted by $g(\lambda)$. The value of the asymmetry parameter depends on the underlying angular distribution $f(\vartheta_s)$ of scattering on any given scattering centre. The asymmetry is positive when photons are predominantly forward-scattered, negative when they are predominantly scattered in the backward direction, and zero when the scattering proceeds with equal probability in the forward and backward directions, $f(\pi - \vartheta_s) = f(\vartheta_s)$ [354]. In the latter case, the effective scattering length $\lambda_e \equiv 1/k_e$ is equal to the geometric scattering length λ_s . For an average dust composition of the deep glacial ice at the South Pole, Mie theory predicts $g \approx 0.94$ when averaged over the visible-wavelength interval [353]. Such strongly forward-scattered visible light has an effective scattering length λ_e that is approximately one order of magnitude greater than its geometric counterpart λ_s .

Figure 6.4 shows the (depth-averaged) wavelength dependence $k_e(\lambda)$ as measured in deep ice by AMANDA-B [343], normalised to its value at $\lambda = 532$ nm. k_e approximately scales as a power law $\propto \lambda^\alpha$, with an exponent of $\alpha = -(0.90 \pm 0.03)$. This suggests, for example, that the effective scattering length for UV light at $\lambda = 337.1$ nm as emitted by a nitrogen laser is 2/3 as large as that for green light at $\lambda = 532$ nm. The exponent α is highly sensitive to the size distribution of the dust constituents. Mie theory predicts a value $\alpha \sim -1$ when the typical modal radius of the impurities is of the order of $0.1 \mu\text{m}$ [353]. For deeper ice with lower dust concentrations than those underlying Fig. 6.4, and for wavelengths further into the UV region—where the Cherenkov spectrum (6.2) peaks—it has been posited that Rayleigh-like scattering with $\alpha = -4$ due to thermal fluctuations and decorated dislocations should become important [353]. Since the optical acceptance of the commonly deployed photosensors in the UV region

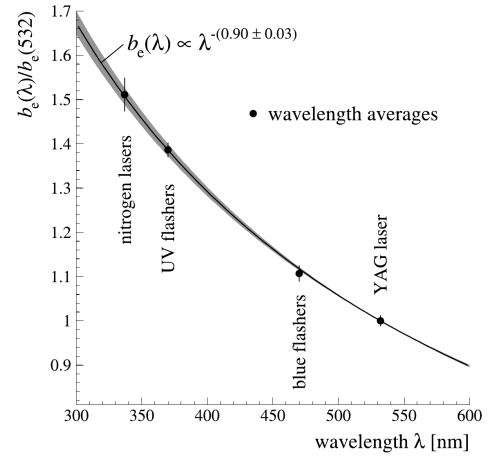


Figure 6.4: Depth-averaged effective scattering coefficient k_e (here denoted as b_e) from near-UV to visible wavelengths measured in deep South Pole ice. Figure taken from Ref. [343].

$\lambda \lesssim 300$ nm is extremely small though, the corresponding effective scattering coefficient does not have to be modelled in practice.

Depth dependence of absorption and scattering




Figure 6.5: Depth dependence of absorption coefficient/length (left) and effective scattering coefficient/length (right) for light with a wavelength of $\lambda = 400$ nm as obtained by measurements with IceCube (“SPICE MIE”) and AMANDA-B (“AHA”). Figure taken from Ref. [355].

An overall reduction in both absorption and scattering with increasing depth in South Pole ice is observed in Fig. 6.5, which compares two joint measurements of the absorption coefficient k_a ($\lambda = 400$ nm) (left panel) and the effective scattering coefficient k_e ($\lambda = 400$ nm) (right panel) as a function of depth from 1 400 m to 2 500 m. “AHA” represent an ice model that is based on the AMANDA-B measurements from Ref. [343] mentioned in the preceding paragraphs. For the most densely instrumented depths between 1.55 km and 1.95 km, the quoted uncertainty on k_a is approximately 14 %, and that on k_e approximately 5 %. The parameter values at depths below those instrumented by AMANDA have been obtained by comparing the depths of the various observed peaks in the optical measurements at the South Pole to dust concentrations in an ice core and extrapolating [356]. A measurement of the same two parameters with the still incomplete IceCube detector in 2008 is overlaid (labelled “SPICE MIE”) [355]. The associated uncertainties are at the level of 10 %. In contrast to previous ice-modelling attempts, the IceCube model has an additional degree of freedom affecting the shape of the scattering function $f(\vartheta_s)$ in order to improve the quality of the fit.

Tilted layers and anisotropy

Figure 6.5 hides at least two significant findings about the nature of the South Pole ice, namely tilted ice layers and an optical anisotropy.

First, the depicted depths are measured along a straight vertical line connecting one particular surface point with the bedrock. High-resolution, sub-cm dust profiles as a function of depth taken at different surface coordinates reveal shifts in the depths corresponding to the various peaks (“dust layers”). The magnitudes of the shifts are of the order of 10 %, corresponding to depth changes of 0.1 km over a distance of 1 km along a direction that is approximately perpendicular to the direction of the flow of the glacier ice [346].

Second, an asymmetry in the amount of light detected by photosensors located at different azimuthal angles with respect to the light source has been revealed, providing evidence of an optical anisotropy of the South Pole ice. When averaged over all depths, this anisotropy manifests itself as a light-yield excess (with respect to the azimuthal average) at the level of 16 % per 100 m of propagation distance along a particular azimuthal axis, accompanied by a similar deficit for light propagating along the perpendicular azimuthal axis [357]. The fact that the axis along which the excess is observed closely aligns (to within 5°) with the direction of the ice flow has resulted in explanation attempts calling upon a direction-dependent scattering coefficient caused by elongated dust constituents, or a direction-dependent absorption coefficient caused by impurities that have accumulated along the boundaries of elongated ice grains. However, as both hypotheses exhibit clear deficiencies when compared to the observed light yields and photon delay time distributions, neither seems to be the main origin of the anisotropy [358]. Instead, the birefringent nature of the ice has been identified as a more plausible underlying mechanism [359, 331].

7

The present: IceCube and DeepCore

7.1 Detector design

The IceCube detector is embedded in the ice cap in the direct vicinity of the South Pole. It consists of an array of 5160 “digital optical modules” (DOMs) attached to 86 vertical cables (“strings”) containing twisted copper-wire pairs. Serving as a nearly autonomous fundamental detection unit, each DOM comprises a downward-facing 25 cm-diameter photomultiplier tube (PMT) [360] together with several circuit boards enclosed in a 35.6 cm-diameter pressure-resistant spherical glass housing. IceCube is located in a radio-quiet area known as the “dark sector” [361] and has been taking data since February 2005, shortly after the water in the borehole surrounding the first string had refrozen [34]. In total, it took seven field seasons of hot water drilling and string deployment until the detector was completed in December 2010. The refrozen columns of ice, called “hole ice” (as opposed to the unaltered bulk ice in between strings discussed in Sec. 6.2.3), are characterised by a much reduced transparency caused by air bubbles released from the meltwater during the drilling procedure.

78 strings make up the primary “IceCube” detector array as seen in Fig. 7.1. These strings are arranged on a triangular grid, within a hexagonal footprint roughly 1 km^2 in area, with a horizontal spacing of 125 m between strings. Each string holds 60 DOMs deployed at depths between 1450 m and 2450 m below the surface. The vertical distance between neighbouring DOMs is 17 m.

A denser infill array is located at large depths $> 1750 \text{ m}$ in the central region of IceCube, also shown in Fig. 7.1. Its eight strings have inter-string spacings between 41 m and 105 m. A “veto cap”, which is almost 100 m high, is formed by the 10 uppermost DOMs on these strings. The remaining 50 DOMs on each string are deployed

at depths between 2100 m to 2450 m, below a layer with an elevated dust concentration. These DOMs have a vertical spacing of only 7 m and are surrounded by the clearest ice.

By convention, the combined instrumentation at depths $\gtrsim 2100$ m of the seven centralmost IceCube strings and the eight specialised strings is referred to as the “DeepCore” subarray. It is predominantly made up of DOMs with a higher quantum efficiency (HQE). Six of the specialised DeepCore strings exclusively hold DOMs with PMTs whose quantum efficiency exceeds that of regular IceCube DOMs by $\sim 35\%$, while the other two carry a mixture. As a result of all of the above, the subarray’s energy threshold is as low as a few GeV.

Neutrinos selected for physics analyses with DeepCore are commonly required to have interacted within a minimal fiducial volume whose boundaries are also depicted in Fig. 7.1 (in the horizontal: dashed cyan hexagon enclosing the centralmost IceCube and the specialised DeepCore strings in the top panel; in the vertical: green rectangle below 2100 m in the bottom panel). In addition, the footprint of an extended fiducial volume is shown in the upper panel. The deep part of the 20 strings it comprises are used for triggering the detector on dim light signals as outlined in Sec. 7.2.3.

7.1 Detector design




Figure 7.1: Top- and sideviews of the IceCube detector including the DeepCore infill array, together with a depth profile of the absorption coefficient of the bulk ice (cf. Sec. 6.2.3). In the top panel, the filled green circles represent locations of strings of the primary IceCube array, whereas the eight red circles show the locations of the more densely instrumented specialised DeepCore strings. The horizontal footprint of the DeepCore fiducial volume also includes the seven centralmost IceCube strings, as depicted by the dashed cyan hexagon. An extended footprint containing five additional IceCube strings is indicated by the purple dotted shape. The bottom panel shows a selection of IceCube and DeepCore strings in a sideview. Figure taken from Ref. [362].

7.2 Data acquisition and processing at the South Pole

Data acquisition (DAQ) in IceCube relies on a decentral process: a large ensemble of DOMs detect optical Cherenkov photons and digitise and timestamp the resulting signals. Accordingly, IceCube can be considered as a gigantic digital network of DOMs. Each pair of DOMs communicates with the surface via a dedicated twisted-wire pair that is routed through the main in-ice cable, which is connected to a surface junction box. The twisted-wire pair also provides each connected DOM with 96 V of total DC power. Starting from the surface junction boxes, trenched surface cables between 300 m and 800 m in length lead to the “IceCube Laboratory” (ICL). The ICL houses (among others) dedicated computers, called “DOMHubs”, which are responsible for the DOM readout, power supply, timing, and communication. A GPS-driven master clock synchronises surface clocks running in the DOMHubs to UTC. Various crucial elements in the acquisition and basic processing of data are examined below. A more comprehensive and technical overview of the working principles behind IceCube is given in Refs. [363, 364].

7.2.1 Digital optical module

Figure 7.2 contains an overview of the main components of a DOM. The main board (MB) placed around the PMT’s neck serves as the DOM’s “central processor”: it is in charge of capturing and digitising the analog PMT anode signal and controlling all devices inside the DOM. The MB also controls the high-voltage (HV) subsystem responsible for the PMT power supply, which generates voltages of around 1.3 kV (at a target PMT gain of 10^7). A crystal oscillator provides local time stamps for any operations occurring inside the DOM. In addition, the MB drives a set of devices for self-calibration purposes, controls sensors for pressure, temperature, and power supply voltage, and communicates with adjacent DOMs and the surface DAQ system. It is also interfaced with the “flasher board”, which holds 12 LEDs. These enable inter-module calibration, ice property studies, or the simulation of physics events, by means of adjustable high-intensity UV light




Figure 7.2: Illustration of the main components of a DOM without its glass housing. Figure taken from Ref. [364].

pulses.

The analog PMT signal is split into three paths at the input to the MB. The “discriminator” serves triggering purposes: amplified PMT waveforms are compared to a discriminator threshold set to a typical voltage-equivalent of one fourth of the signal induced by a single photoelectron (SPE). The detection of a threshold crossing results in the high-speed waveform capture and digitisation, a “launch”. In addition, two types of digitisers provide sampled PMT waveforms. The “Analog Transient Waveform Digitiser” (ATWD) produces a set of analog PMT voltage samples at an interval of 3.3 ns over a time span of 427 ns, which is digitally converted when certain logical conditions are met (see below). The “fast analog-to-digital converter” (fADC) continuously samples the PMT signal at a lower rate, with a sampling interval of 25 ns, over a total duration of 6.4 μ s. For small long-lasting signals the fADC can provide essential information.

All of the above components are enclosed in a 1.3 cm-thick borosilicate glass housing which is able to withstand the extreme long-term pressure of 250 bar in the deep ice (even higher during freeze-in). The PMT, which is shielded from the ambient South Pole magnetic field by a mu-metal grid, is optically coupled to the glass sphere via silicone gel. In conjunction with the glass transmissivity, the gel leads to the short-wavelength cutoff of the DOM at around 350 nm.

7.2.2 Local-coincidence logic

DOMs are able to receive, send, and relay “local coincidence” (LC) signals to neighbouring DOMs on the same string. The amount of information included in a DOM’s digital output record (“hit”) depends on the state of neighbouring DOMs.

The case of no coincidence (“soft local coincidence”, SLC) is characterised by a lack of discriminator threshold crossings within a time window of $\pm 1 \mu$ s in any of the nearest and next-to-nearest neighbour DOMs on the same string above and below a given DOM that has been launched. It will lead to an early abortion of the digitisation and only three samples of the fADC record centred around the peak value being sent to the DAQ system at the surface. These samples are chosen such that a coarse charge stamp can be assigned and the time of the underlying SPE inferred.

If a neighbouring DOM has also signalled a launch, however, the entire digitised ATWD and fADC waveform information is also transmitted (“hard local coincidence”, HLC). The reasoning behind this logic is that isolated hits are more likely to originate from PMT dark noise than from actual physics events of interest. As a result, dead times and noise-induced data flows are significantly limited. On the other

hand, as physics events sometimes do cause isolated hits, the coarse charge information recorded for SLC hits can in principle still aid in the formation of global triggers, as discussed in the following section.

7.2.3 Triggering

At the surface, software running on the DOMHubs converts DOM hit time stamps into a calibrated “DAQ time” and performs a time ordering. HLC hit information is then employed by a set of software algorithms which isolate time windows around configurable spatial or temporal clusters of hits (or both). Each trigger window is padded with longer readout windows before and after to ensure the capture of early and late hits accompanying a physics event. Finally, a global trigger merges and unifies overlapping trigger readout windows, from which an “event” is created.

Trigger conditions in IceCube are typically simple, and a significant event rate reduction happens in online and offline filtering at a later stage. The fundamental and most commonly used “simple multiplicity” trigger (SMT) searches for time windows of duration T_0 containing at least N HLC hits throughout the detector, without any spatial requirements. For the primary IceCube array, the corresponding parameters are $T_0 = 5 \mu\text{s}$ and $N = 8$ (SMT8).

In the case of DeepCore, we require $N = 3$ HLC hits (SMT3) within a time interval of $T_0 = 2.5 \mu\text{s}$ throughout the 50 deepest DOMs on each of the eight specialised DeepCore strings and the 22 deepest DOMs on each of the twelve IceCube strings which are part of the extended DeepCore footprint in Fig. 7.1 (“SMT3 DOMs” in the following). This configuration reflects the lower energy threshold of DeepCore as compared to the surrounding primary IceCube array. The trigger condition is evaluated by sliding a time window of duration T_0 across the time series consisting of the times of all relevant DOM launches. The trigger time is determined by the earliest launch within the first time window containing three HLC DOM launches. The time window keeps sliding until the number of HLC DOM launches within it drops below the threshold. All launches (HLC and SLC) within a time window for which the threshold condition is satisfied are included in the trigger time window, whose duration is $T' \geq T_0$. Finally, this window is padded by $4 \mu\text{s}$ before its beginning and $6 \mu\text{s}$ after its end. All hits within the resulting readout window constitute what will be referred to as an “SMT3 event” in the following.

The SMT3 trigger is initiated at a typical rate of 250 Hz, roughly one order of magnitude less frequently than IceCube’s SMT8. For a broad overview of triggers used in the IceCube Neutrino Observatory, see Ref. [364, Sec. 6.4.2].

7.2.4 Pulse extraction from hit records




Figure 7.3: Result of the simultaneous unfolding of the calibrated ATWD and fADC waveforms (in green and purple, respectively) recorded by a DOM in response to four simulated photons (blue dotted vertical lines) hitting the photocathode. Figure taken from Ref. [365].

As a part of the standard online processing chain, the digitised PMT waveforms included in a given hit are “deconvolved” into a single time series of photoelectrons ejected from the photocathode of each DOM, in SPE units. The unfolding algorithm [366] is applied to calibrated voltage waveforms [367], which are generated by correcting the raw ADC counts recorded by the ATWDs or the fADC for their baselines, their amplification factors, for certain distortions [368] induced by a transformer on the PMT base board, and for time delays due to a delay line on the MB and the PMT traversal of a photomultiplied pulse. Remaining temporal shaping effects applied to the SPE waveform produced by the PMT anode are due to the transformer and the amplifier stages preceding the ATWD channels and the fADC. Assuming the electronics response to act in a linear manner—i.e., that each observed waveform is a linear superposition of several SPE response functions shifted in time with respect to each other—allows deconvolving the original series of photoelectron “pulses”, as described in Refs. [366, 369].

An example is shown in Fig. 7.3, where four simulated photons hit the photocathode of a DOM within a time span of 200 ns, resulting in pulses with charges between 0.8 PE and 1.5 PE. Here, the unfolding succeeds in extracting the photon arrival times and the charges of the associated pulses at high accuracy from the two calibrated waveforms recorded by the ATWD and the fADC.

7.2.5 Removal of noise pulses

The inclusion of SLC hits in SMT3 events is conducive both to vetoing events due to atmospheric μ^\pm 's and to the reconstruction of dim neutrino events, depositing only a few PE in the detector. However, the presence of SLC hits brought about by dark noise in general complicates both. The “seeded R - T cleaning” method devised in

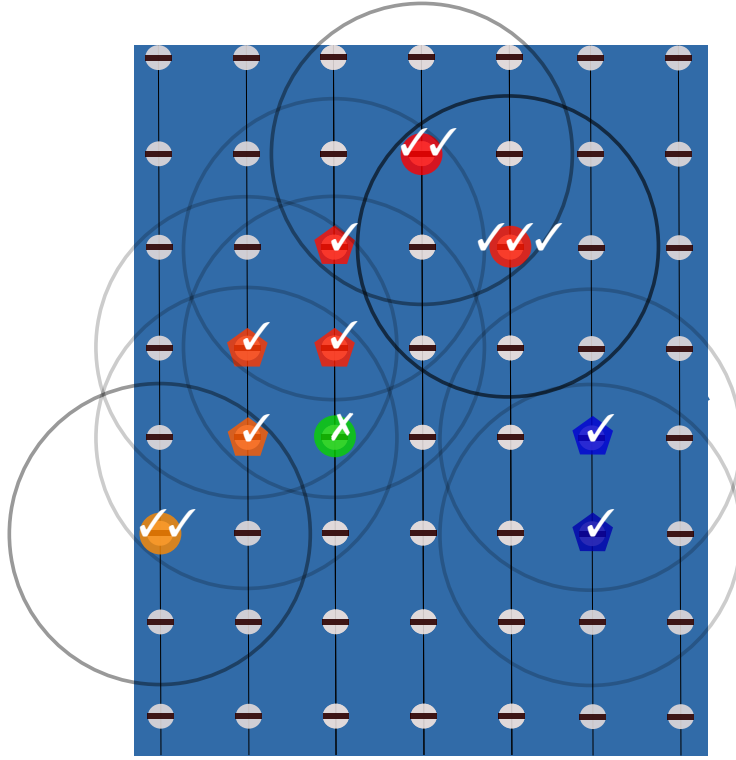


Figure 7.4: Simplified sketch of the “seeded R - T ” noise-removal procedure, using HLC signals as seeds. DOMs detecting signals are highlighted by filled pentagons (for HLC hits) or circles (SLC hits). The large empty circles represent the spatial search radius R_{\max} . The temporal search criterion is not shown explicitly; instead, the signal times are colour-coded (red: early, blue: late, green: in between). In this sketch, signals detected by DOMs designated with at least one tick would be retained, while signals with a cross would be discarded. The number of ticks corresponds to the iteration number in which the signal would be accepted. The discarded SLC signal is assumed to fail the temporal search criterion, e.g., because it is not causally connected to the HLC signals in its spatial vicinity.

Ref. [39] represents an effective general-purpose noise-removal technique, which is able to simultaneously reject a large fraction of noise signals (either launches/hits or individual pulses) and retain a large fraction of Cherenkov signals. Its working principle is straightforward: starting from a set of “seeds”—some suitably chosen initial ensemble of spacetime points $\{(\mathbf{x}_{\text{seed},i}, t_{\text{seed},i})\}_i$ —it searches for further signals $(\mathbf{x}_{\text{test}}, t_{\text{test}})$ within a fixed predefined radial (R_{\max}) and temporal (T_{\max}) vicinity of each seed. All signals with

$$|\mathbf{x}_{\text{test}} - \mathbf{x}_{\text{seed},i}| \leq R_{\max} \quad \text{and} \quad |t_{\text{test}} - t_{\text{seed},i}| \leq T_{\max} \quad (7.1)$$

are accepted and added to the original set of signals. In each of the following iterations, the procedure begins from the signals accepted during the previous iteration. Once an iteration does not find any further signals, the procedure is terminated. What remains is an ensemble of signals clustered in space and time. Figure 7.4 provides a sketch of the algorithm for the case of HLC signals serving as the initial seeds.

7.2.6 Filtering




Figure 7.5: Sketch of the DeepCore online filtering scheme, taken from Ref. [40].

In order to reduce the experimental data to an amount transmittable to the IceCube data centre in Madison, WI (US) via satellite, a processing and filtering software in place at the ICL filters the triggered events it receives from the DAQ with a latency of only a few tens of seconds. A rather small set of core filter selections has been employed with only small modifications since the start of detector operation [364, Sec. 6.5.4].

The “DeepCore online filter” [40]¹ aims at removing events due to atmospheric μ^\pm energy losses according to the scheme sketched in Fig. 7.5. It is the only filter relevant to this work. The implementation in place at the time of this

writing has been unchanged since the beginning of the 2012 data-taking season [370, 371]. The filter decision is based on those pulses in an SMT3 event that are retained by the R - T method above when using the times and DOM coordinates of all HLC pulses as initial seeds, with $R_{\max} = 150$ m and $T_{\max} = 1$ μ s. For historical reasons, only the earliest pulse on each DOM is considered. Among these, the algorithm then counts the number of pulses detected by DOMs that serve as a “veto” for DeepCore, i.e., all but the 664 SMT3 DOMs. In case at least one such veto pulse is present, a simple initial guess of the event’s vertex position and time is constructed based on a refined [39] average time and position (“centre of gravity”, COG) of the cleaned pulses recorded by the SMT3 DOMs. The filter then determines the apparent velocities v_i of hypothetical particles

¹Introduced under the name “causally related hit veto” in Ref. [39].

propagating toward the COG from all spacetime points corresponding to the pulses in the veto region. If two or more velocities are found to be “close” to the speed of light c ($0.25 \text{ m ns}^{-1} \leq v_i \leq 0.4 \text{ m ns}^{-1}$), the event is discarded, and otherwise retained. Events devoid of pulses in the veto region are also kept. This scheme results in an overall rejection of $\sim 95\%$ of atmospheric μ^\pm 's, at a neutrino signal efficiency of $\sim 99\%$, leading to a filtered data rate of $\mathcal{O}(20 \text{ Hz})$ [40].

Events passing the filter selection are transmitted to the north on a daily basis, while those that are rejected by the filter criterion are written to disk. Disks containing archival data are shipped to the IceCube data centre for indefinite retainment annually during the austral summer.

7.3 Detector calibration

The natural stability of the glacial medium in which IceCube is embedded makes it affordable to limit certain detector calibration campaigns to an annual basis, while others proceed more continuously. Constantly ongoing calibration efforts include (see Ref. [364] for details):

- the waveform calibration (once per year), which determines constants required for translating the waveforms recorded by each IceCube DOM into Cherenkov photon arrival times by making use of the reference signal sources (oscillator, DC bias voltage circuit, programmable electrical pulser, low-intensity UV LED) integrated into each DOM;
- the detector time base calibration (in parallel to regular data taking), which ensures a global time base for the whole detector at an accuracy of $\mathcal{O}(\text{ns})$ via the reciprocal emission and timestamping of electronic pulses between each DOMHub and its connected DOMs;
- the overall optical efficiency calibration, for which laboratory measurements are supplemented by analyses of the PMT charges observed in dedicated offline event samples containing single low-energy, minimum-ionising, μ^\pm 's, which stand out as calibration tools due to their well understood and approximately constant Cherenkov light yield; and
- the bulk ice calibration (annual flasher campaigns), during which optical ice properties are fit to the hits produced by well-controlled light pulses emitted by the 12 LEDs on each DOM.

8

The future: PINGU and the IceCube Upgrade

In the case of DeepCore, it barely took two years from its first proposal in 2007 to the beginning of its deployment [39]. PINGU (“Precision IceCube Next Generation Upgrade”), on the other hand, was first proposed already in 2010 [372], when six of the eight DeepCore strings had just begun taking data, and remains at the proposal stage as of today. Instead, in 2019 the IceCube “Upgrade” was approved [373, 374] as the “Phase 1” extension of the envisaged “IceCube-Gen2 Observatory”.¹ The idea behind the IceCube Upgrade project is based on the necessarily phased deployment procedure of in-ice detectors at the South Pole together with funding constraints. Upon the Upgrade’s first proposal in 2016 [375], the foreseen detector layout therefore essentially corresponded to a seven-string subset of PINGU, which could be completed during later deployment seasons. There is a high degree of overlap between the physics goals of the Upgrade and PINGU.

8.1 The proposed PINGU experiment

PINGU provides an even more densely instrumented multi-megaton volume within the DeepCore volume. The proven principle of operation behind IceCube—and in particular its DeepCore extension—serves as the foundation of PINGU’s design,

¹This is the major next-generation extension of IceCube, which will integrate the existing IceCube detector with a much larger, $\mathcal{O}(10 \text{ km}^3)$, in-ice Cherenkov telescope and a dense low-energy core [43]. The in-ice arrays will be complemented by an array of radio antennas at different depths (to probe the astrophysical neutrino flux at energies above $\sim 10 \text{ PeV}$) and by a diverse array of surface detectors.

which has undergone extensive optimisation since the proposal.

The PINGU physics analyses in this work assume the detector layout reported as the “new baseline” geometry in the second version of the PINGU Letter of Intent (LoI) [1, 2]. This layout comprises 26 strings instrumented with 192 HQE DOMs each, at depths between 2150 m and 2450 m—slightly inset with respect to the boundaries of the DeepCore fiducial volume in order to enable an efficient vetoing of the atmospheric μ^\pm background. Neighbouring strings are spaced apart 24 m on average, neighbouring DOMs on the same string 1.5 m. Whereas earlier PINGU geometries comprising a larger number of strings suffered from considerably higher cost, driven by drilling, cable, and deployment expenses, the suggested number of strings for the adopted layout allows for a full detector construction within two seasons, based on experience gathered during the completion of IceCube and DeepCore: here, the drilling crew achieved a maximum installation rate of 20 strings per season (in 2009/2010) [376]. In contrast to past deployment procedures, PINGU could profit from dedicated degassing and filtering of the meltwater in the boreholes, which would prevent the refrozen ice in the vicinity of PINGU’s strings and DOMs from exhibiting similarly enhanced scattering compared to the surrounding bulk ice.

Figure 8.1 shows the positions of the strings in the horizontal x - y plane as well as the z -coordinates (origin at depth of 1948 m) of the PINGU DOMs in relation to both the centralmost IceCube as well as the dedicated DeepCore strings and DOMs. PINGU’s fiducial volume corresponds to a cylinder of 320 m in height and 85 m in radius around the detector centre, yielding a fiducial mass of ~ 6 Mt. Together with the existing IceCube and DeepCore strings, this results in a detector with 178 strings and 10152 DOMs.

Data acquisition with the PINGU detector is foreseen to be fully compatible with the system in place for IceCube [1, 2]. The reference optical module (PDOM) closely follows the design of the DeepCore DOM. While the external components of the PDOM are essentially unchanged with respect to those of the DOM, the internal readout electronics are simplified and updated, reducing module cost, power consumption, and MB space. In addition, multiple alternative optical module designs could benefit PINGU’s physics goals if deployed alongside or instead of the PDOM. These alternate designs aim at minimising module noise, maximising photon sensitive area, or improving directional or temporal photoelectron resolutions. Some of these technologies will soon be deployed with the Upgrade and are briefly mentioned in Sec. 8.2.

The PINGU sensitivity studies in this work are conservatively based on the un-

modified DeepCore DOM, which allowed employing the well-proven standard IceCube and DeepCore software tools for simulation and reconstruction. Similarly, while one can expect optimised trigger algorithms to be developed in the future [1], for the purpose of this work only those events are considered which satisfy an SMT3 trigger criterion just like the one in use for DeepCore, but which is applied to DOMs on the 41 strings depicted in Fig. 8.1: all 192 DOMs on each of the 26 PINGU strings, the 50 deepest DOMs on each of the eight specialised DeepCore strings, and the 23 deepest DOMs on each of the seven central IceCube strings. The basic filtering procedure is outlined together with the full event selection in Part IV.

8.2 The imminent IceCube Upgrade

The seven strings forming the Upgrade array are expected to be deployed during the Antarctic summer season 2025/2026 [377]. They will be equipped with a wide variety of novel optical sensor types and calibration devices, for a total of around 700 [378].

One major type of sensor is the “multi-PMT optical module” (mDOM) [379, 380], shown on the left of Fig. 8.2. By equipping each module with 24 small PMTs—each approximately 80 mm in diameter—a nearly uniform angular photon detection efficiency is achieved. What further distinguishes mDOMs from regular IceCube DOMs is their increased effective area for photon detection by more than a factor of two (when averaged over the Cherenkov spectrum and the full solid angle), and the potential benefit on event reconstruction from knowing the direction of the incident photon that gives rise to a PMT output charge.

The second major module type is the “D-Egg” [377], shown on the right of Fig. 8.2. It features a pair of larger PMTs in an ellipsoidal glass housing, one facing vertically upward and the other vertically downward. It improves on the PDOM in several aspects, such as the UV photon detection efficiency, the uniformity of the angular acceptance, the intrinsic photon directionality—albeit limited to one hemisphere versus the other—the effective area for photon detection, or the smaller diameter of its pressure housing.

Studies of the Upgrade’s atmospheric neutrino oscillation physics potential based on a dedicated simulation of the novel module types are ongoing [381].




Figure 8.1: Layout of the proposed PINGU detector considered in this work. The 26 PINGU strings are surrounded by the seven centralmost IceCube strings and DeepCore. The upper panel shows a topview (x - y plane), whereas the lower panel shows a sideview (x - z plane). The coordinate value $z = 0$ corresponds to a depth of 1 948 m below the surface of the South Pole ice sheet.



Figure 8.2: Prototype of the mDOM (left) and the D-Egg (right). Both types of modules will be used to instrument the seven Upgrade strings. Figure taken from Ref. [43].

Part III

Search for non-standard neutrino interactions with DeepCore

9

Atmospheric neutrino oscillations in the presence of NSI

Depending on the assumed flavour structure of the effective NC NSI Hamiltonian (4.18) for Earth matter, the oscillations of atmospheric neutrinos at the GeV energy scale will proceed very differently than in the case of standard Earth matter effects investigated in Sec. 5.1.2.2. This chapter provides a detailed account of the expected impacts of the standard NSI couplings on the relevant atmospheric neutrino oscillation probabilities and points out the phenomenological implications for the event sample utilised in the analysis. The underlying calculations are based on a dedicated adaptation of the Prob3++ software [382], which is briefly summarised in Appendix A.

In order to scrutinise the atmospheric neutrino oscillation probabilities $P_{\alpha\beta}$ under the continuous variation of a given NSI parameter at a time, an inclined trajectory with $\cos \vartheta = -0.75$, which only crosses the Earth's mantle, is chosen. This choice of zenith angle corresponds to the baseline $L \approx 9.6 \times 10^3$ km (cf. Fig. 5.4). We assume an atmospheric neutrino production height of $h = 20$ km above the surface [281] and a detection depth of $d = 2$ km below. The values for the PMNS parameters are taken from the NuFIT 3.2 NO global best fit [327, 328], except for δ_{CP} , which is set to zero. These settings also represent the nominal assumptions for the NSI analysis in this work.

The oscillation probabilities are plotted for each of the six neutrino channels $\nu_e \rightarrow \nu_e$, $\nu_e \leftrightarrow \nu_\mu$, $\nu_e \rightarrow \nu_\tau$, $\nu_\mu \rightarrow \nu_\mu$ and $\nu_\mu \rightarrow \nu_\tau$, and as a function of neutrino energy $2 \text{ GeV} \leq E_\nu \leq 1000 \text{ GeV}$ (the antineutrino channels can be found in Figs. C.1 to C.5 in the appendix). While the focus is on real NSI couplings for simplicity, the effects of complex phases are touched upon where applicable. Interference effects between

the transition amplitudes in the mantle and the core are not discussed (the fraction of events from core-crossing neutrinos included in the analysed DeepCore dataset is small, see Sec. 12.3.2).

9.1 Rescaled standard matter potential




Figure 9.1: Oscillation probabilities of atmospheric neutrinos crossing the Earth under the zenith angle $\cos \vartheta = -0.75$ vs. neutrino energy E_ν , for an initial flux of electron neutrinos in the top, and for an initial flux of muon neutrinos in the bottom row. Shown are different realisations of the effective matter-potential strength ϵ_\oplus , with $-5 \leq \epsilon_\oplus \leq 5$. The blue dashed lines (hardly visible) show the probabilities obtained for $\epsilon_\oplus < 0$, while the red dashed lines show those obtained for $\epsilon_\oplus > 0$. Darker shades represent larger $|\epsilon_\oplus|$. The two cases of SI ($\epsilon_\oplus = 1$) and no interactions (vacuum, $\epsilon_\oplus = 0$) are highlighted. See text for details.

The six panels in Fig. 9.1 illustrate the result of varying ϵ_\oplus (or equivalently $1 + \epsilon_{ee}^\oplus - \epsilon_{\mu\mu}^\oplus$) when the matter potential remains restricted to the ee entry. All scenarios represent the simple rescaling of the SM matter potential, $V_{CC}(x) \rightarrow V'(x) \equiv \epsilon_\oplus V_{CC}(x) = (1 + \epsilon_{ee}^\oplus - \epsilon_{\mu\mu}^\oplus) V_{CC}(x)$. The formalism of standard matter effects from Sec. 5.1.2.2 thus applies. Two cases that are nested within the considered parameter interval are highlighted: no interactions or vacuum ($\epsilon_\oplus = 0$) and SI ($\epsilon_\oplus = 1$).

At the energies considered here, all transitions involving ν_e are suppressed in vacuum compared to those not involving ν_e . Once a positive non-zero matter potential is introduced for ν_e , its transitions can be enhanced, while a negative matter potential leads to their suppression. On the one hand, the ν_e disappearance probability $1 - P_{ee}$ in vacuum remains small in the limit $\Delta m_{21}^2 \rightarrow 0$: $1 - P_{ee} \leq \sin^2 2\theta_{13}$. When $\epsilon_\oplus > 0$, on the other hand, the MSW resonance condition can be satisfied (given $\Delta m_{31}^2 > 0$ and

$\theta_{13} < \pi/4$), and a complete disappearance of the initial ν_e flux can be observed when the amplitude and phase conditions are satisfied. The example trajectory in Fig. 9.1 ensures that the oscillation half-phase at resonance is given by $\phi_R(L) \approx \pi/2$, such that there is a nearly complete disappearance of ν_e 's at the energy $E_R^{\text{SI}} \approx 6$ GeV for SI. The appearance probabilities of ν_μ 's and ν_τ 's are simply given by $P_{e\mu} = \sin^2 \theta_{23} (1 - P_{ee})$ and $P_{e\tau} = \cos^2 \theta_{23} (1 - P_{ee})$, respectively. Since the vacuum 2–3 mixing is close to maximal, it follows that ν_e 's convert to ν_μ 's and ν_τ 's with approximately the same probability, $P_{e\mu}(E_R^{\text{SI}}) \approx P_{e\tau}(E_R^{\text{SI}}) \approx 0.5$.

Depending on the assumed value $\epsilon_\oplus > 0$, both the resonance energy E_R and the resonance half-phase $\phi_R(L)$ shift. When $0 < \epsilon_\oplus < 1$, the half-phase $\phi_R(L) < \pi/2$, and no full ν_e disappearance can occur for the chosen trajectory. In contrast, the half-phase becomes an odd multiple of $\pi/2$ when $\epsilon_\oplus \approx 3$. Accordingly, one observes full ν_e disappearance also at $E_R \approx 2$ GeV in Fig. 9.1. Negative values of V' together with $\Delta m_{31}^2 > 0$ do not give rise to a similar enhancement, because $\cos 2\theta_{13} > 0$. Consequently, there are no significant transitions $\nu_e \rightarrow \nu_{\mu,\tau}$ for $\epsilon_\oplus < 0$. Instead, the antineutrino transitions $\bar{\nu}_e \rightarrow \bar{\nu}_{e,\mu,\tau}$ are then subject to the matter effects detailed above.

Figure 9.1 further demonstrates that the transitions $\nu_\mu \rightarrow \nu_{\mu,\tau}$ proceed as in vacuum for sufficiently high energy, $E_\nu \gtrsim 20$ GeV for the considered trajectory—irrespective of the value of ϵ_\oplus . As ν_e 's decouple from the evolution at high energy, the $\nu_\mu \rightarrow \nu_\tau$ transition probability is given by the complementary ν_μ survival probability, namely $P_{\mu\tau} \approx 1 - P_{\mu\mu}^{2\nu}$, with $P_{\mu\mu}^{2\nu}$ given in Eq. (5.27).

9.2 μ - τ non-universality and flavour violation

In case $\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus$ is the only source of NSI as in Fig. 9.2, the flavour non-universality gives rise to an effective potential in the decoupled μ - τ system mentioned in the previous section. The 2–3 mixing in matter is modified according to the standard MSW mechanism, but with a potential $V'(x) = (\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus) V_{\text{CC}}(x)$, leading to a resonance energy of the two-neutrino system that follows from Eq. (3.63) by substituting θ_{23} for θ , Δm_{31}^2 for Δm^2 , and V' for V_{CC} [383]. If $\Delta m_{31}^2 \cos 2\theta_{23} > 0$ (< 0),¹ the resonance occurs in the neutrino channel for $\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus > 0$ (< 0), and in the antineutrino channel for $\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus < 0$ (> 0). The resonant enhancement of the 2–3 mixing in matter is much weaker than that of the 1–3 mixing investigated in Sec. 9.1, or even

¹Only the relative sign of Δm_{31}^2 and $\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus$ is relevant for the evolution of the two-neutrino system, compare relations (4.20).




Figure 9.2: Same as Fig. 9.1, but for different realisations of the NSI NU strength $\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus$, with $-0.10 \leq \epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus \leq 0.10$ (negative values in blue, positive values in red, larger magnitudes darker).

fully absent. In fact, the introduction of $\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus$ in general reduces 2–3 mixing. As discussed in Ref. [383], the main observable consequence of $\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus$ is the increased survival probability and therefore flux of both atmospheric ν_μ 's and $\bar{\nu}_\mu$'s across the broad range of energies at which the μ - τ system is decoupled. This can be understood by realising that both ν_μ and $\bar{\nu}_\mu$ become eigenstates of the Hamiltonian as $|\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus| \rightarrow \infty$.

On the other hand, if $\epsilon_{\mu\tau}^\oplus$ is the only non-zero NSI coupling strength as in Fig. 9.3, the flavour-violating, off-diagonal, elements $V_{CC}(x)\epsilon_{\mu\tau}^{\oplus(*)}$ of the two-neutrino interaction Hamiltonian result in qualitatively different effects on the neutrino evolution [383]. For $\Delta m_{31}^2 > 0$ and real NSI, a resonance occurs for neutrinos when $\epsilon_{\mu\tau}^\oplus < 0$ and for antineutrinos when $\epsilon_{\mu\tau}^\oplus > 0$ (independent of the θ_{23} octant), with a resonance energy of

$$E_R \approx \pm \frac{\Delta m_{31}^2 \sin 2\theta_{23}}{4\epsilon_{\mu\tau}^\oplus \langle V_{CC} \rangle}, \quad (9.1)$$

and an oscillation half-phase at resonance of

$$\phi_R(L) \approx \frac{\Delta m_{31}^2 L \cos 2\theta_{23}}{4E_R} = \pm \frac{\epsilon_{\mu\tau}^\oplus \langle V_{CC} \rangle L}{\tan 2\theta_{23}}. \quad (9.2)$$

Clearly, when $\epsilon_{\mu\tau}^\oplus$ is varied as in Fig. 9.3, resonances at $E_R \gtrsim 60$ GeV are observed. Since the corresponding phases are small, $\phi_R(L) \ll \pi/2$, $P_{\mu\mu}$ becomes nearly maximally enhanced at high energies when $\epsilon_{\mu\tau}^\oplus < 0$. Sufficiently far away from the reson-




Figure 9.3: Same as Fig. 9.1, but for different realisations of the NSI FV coupling $\epsilon_{\mu\tau}^\oplus$, with $-0.05 \leq \epsilon_{\mu\tau}^\oplus \leq 0.05$ (negative values in blue, positive values in red, larger magnitudes darker).

ance, the overall oscillation half-phase is simply given by the sum of that in vacuum, $\phi_{\text{vac}} = \frac{\Delta m_{31}^2 L}{4E_\nu}$, and that induced by NSI, $\phi_{\text{NSI}} = \epsilon_{\mu\tau}^\oplus \langle V_{CC} \rangle L$ [383]. For fixed L , at energies at which ϕ_{vac} is significant, $\epsilon_{\mu\tau}^\oplus$ results in a shift in energy of the oscillation pattern in the μ - τ system. When $\epsilon_{\mu\tau}^\oplus > 0$, a shift to higher energies appears for neutrinos, and a shift to lower energies for antineutrinos (reversed effects for $\epsilon_{\mu\tau}^\oplus < 0$). Asymptotically, the vacuum phase becomes negligible and the oscillation half-phase in matter approaches the constant ϕ_{NSI} . Since the mixing angle in matter simultaneously becomes maximal, the two-neutrino survival probability of both ν_μ and $\bar{\nu}_\mu$ at high energies tends toward [383]

$$P_{\mu\mu}^{2\nu} \approx 1 - \sin^2 \phi_{\text{NSI}}. \quad (9.3)$$

A comparison with Eq. (5.27) reveals that $\epsilon_{\mu\tau}^\oplus$ reduces the survival probability and therefore the fluxes of both atmospheric ν_μ 's and $\bar{\nu}_\mu$'s at high energies, for which ϕ_{vac} can be neglected. At energies below this regime, the inverse effects of a given value of $\epsilon_{\mu\tau}^\oplus$ on the neutrino and antineutrino oscillation probabilities leads to cancellations in the observed event spectra.

For the more general case $\epsilon_{\mu\tau}^\oplus \in \mathbb{C}$, the value of the complex phase $\delta_{\mu\tau}$ affects the impact of the magnitude $|\epsilon_{\mu\tau}^\oplus|$ on the oscillation probabilities in the μ - τ sector. For example, their leading-order perturbative expansions [384, 385] reveal that a purely imaginary coupling (corresponding to $\delta_{\mu\tau} = 90^\circ, 270^\circ$) results in a sensitivity loss at the probability level. However, this does not necessarily imply minimal NSI effects

on the combined oscillated neutrino and antineutrino fluxes across the full spectrum. The reasons for this are that no MSW-like enhancement is possible, and the high-energy survival probabilities of both ν_μ and $\bar{\nu}_\mu$ are subject to the same reduction compared to the SI expectation, reaching an asymptotic limit at a lower energy than in the case of a real coupling. This implies a smaller reduction of $P_{\mu\mu}^{2\nu}$ at the energies observed in this NSI analysis, but the oscillated ν_μ and $\bar{\nu}_\mu$ flux spectra do not suffer from the previously mentioned cancellations in the energy range from approximately 30 GeV to 100 GeV.

9.3 e - μ or e - τ flavour violation

Similar to $\epsilon_{ee}^\oplus - \epsilon_{\mu\mu}^\oplus$, the FV couplings involving the electron flavour, $\epsilon_{e\mu}^\oplus$ and $\epsilon_{e\tau}^\oplus$, are not typically the focus of atmospheric neutrino studies, partly due to their weaker impacts on the $\bar{\nu}_\mu$ disappearance probabilities. It has been shown perturbatively that they contribute only at second order when one considers oscillation probabilities far away from the 1–3 MSW resonance regime [384, 385]. However, $\epsilon_{e\mu}^\oplus$ and $\epsilon_{e\tau}^\oplus$ also enter the oscillation probabilities involving the electron flavour at the second order, lower by one order compared to the four remaining couplings [385]. In the light of this, it is not possible to neglect any of the oscillation channels when considering DeepCore's sensitivity to either of these two FV couplings. Below we restrict ourselves to the basic phenomenological implications.

Fundamentally, transforming the flavour indices $\mu \leftrightarrow \tau$ of a given oscillation channel in the presence of $\epsilon_{e\mu}^\oplus$ and $\epsilon_{e\tau}^\oplus$ is equivalent to substituting $\epsilon_{e\mu}^\oplus \rightarrow \epsilon_{e\tau}^\oplus$ and $\epsilon_{e\tau}^\oplus \rightarrow -\epsilon_{e\mu}^\oplus$ together with $\sin \theta_{23} \rightarrow \cos \theta_{23}$ and $\cos \theta_{23} \rightarrow -\sin \theta_{23}$ in the calculation of the corresponding oscillation probabilities [385]. Only $\epsilon_{e\mu}^\oplus$ results in modifications of the important atmospheric oscillation channels involving the muon flavour across the full range of energies, as seen by comparing the appropriate panels in Fig. 9.4. Characteristically, at high energies e - μ FV is expected to manifest itself in the disappearance of ν_μ 's and $\bar{\nu}_\mu$'s and the simultaneous appearance of ν_e 's and $\bar{\nu}_e$'s, whereas e - τ FV induces the conversion $\bar{\nu}_e \leftrightarrow \bar{\nu}_\tau$. Due to the lacking intrinsic atmospheric tau component, the comparably low intrinsic atmospheric $\bar{\nu}_e$ flux, as well as the reduced $\bar{\nu}_\tau$ CC detection cross section, a priori one expects DeepCore to have a lower sensitivity to $\epsilon_{e\tau}^\oplus$ than to $\epsilon_{e\mu}^\oplus$.

For detailed phenomenological and numerical discussions of $\epsilon_{e\mu}^\oplus$ and $\epsilon_{e\tau}^\oplus$ up to energies of the order of 20 GeV, relevant to future long-baseline and atmospheric neutrino experiments, see for example Refs. [386, 387, 388].




Figure 9.4: Same as Fig. 9.1, but for different realisations of the NSI FV couplings $\epsilon_{e\mu}^\oplus$ (upper six panels) and $\epsilon_{e\tau}^\oplus$ (lower six panels), with $-0.30 \leq \epsilon_{e\mu(\tau)}^\oplus \leq 0.30$ (negative values in blue, positive values in red, larger magnitudes darker).

9.4 Arbitrary NSI flavour structure

Once H_{mat} is allowed to take an arbitrary flavour structure, atmospheric neutrino oscillation probabilities do not in general lend themselves to a transparent analytic discussion any more. Nevertheless, Refs. [325, 253, 329] discuss regimes of neutrino propagation in which the three-neutrino evolution in the presence of NSI can be reduced to an analytically treatable effective two-neutrino system, which is rotated with respect to the flavour basis. They specifically investigate the case in which $\epsilon_{ee}^\oplus - \epsilon_{\mu\mu}^\oplus$, $\epsilon_{e\tau}^\oplus$, and $\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus$ are the only non-zero NSI parameters. Here, two identical

eigenvalues result in the “atmospheric parabola” relation [325, 253, 329]

$$\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} = \frac{|\epsilon_{e\tau}^{\oplus}|^2}{1 + \epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}}, \quad (9.4)$$

which is able to accommodate two-flavour vacuum-like $\bar{\nu}_{\mu}$ disappearance at high energy, independent of the magnitudes of the involved NSI parameters. The assumption of the existence of two degenerate H_{mat} eigenvalues—which leads to the parameterisation described by Eqs. (5.32) to (5.36)—constitutes the key prerequisite for such a two-flavour reduction. It can be considered to be a “natural” generalisation [329] of the two-flavour approximation applied to neutrino oscillations with SI in regimes with negligible impact from Δm_{21}^2 .

10

Event simulation and reconstruction

The IceCube collaboration mostly employs its own set of software projects [389] in the simulation of its detector and the processing and reconstruction of events. Simulation is often produced centrally, with dedicated Monte Carlo (MC) simulation (see, e.g., Ref. [390]) campaigns adapted to and targeting the various types of physics studies performed by the analysers. Essential common processing steps are applied in the process. For example, trigger and filter conditions are evaluated on all simulated events. Different analyses—starting off from these preprocessed events—then diverge as the respective analysers apply their own targeted processing steps (though typically based on common tools and methods), thereby creating dedicated MC event samples that mirror the event samples in data. The software constantly undergoes development to be able to account for the ever-improving understanding of the detector’s response.

This chapter first describes the techniques used to simulate the neutrino events underlying this NSI search in Sec. 10.1, which provides further insight into the existing knowledge of DOM properties, such as their optical efficiencies and angular acceptance functions, both of which constitute important sources of systematic uncertainty for DeepCore neutrino oscillation analyses. Reconstruction—the process of inferring neutrino properties such as the energy or direction from any given event (detector response)—is briefly reviewed in the context of this analysis in Sec. 10.2. When applied identically to observed and simulated events, it allows us to compare the predicted distributions of reconstructed neutrino properties under different physics hypotheses to the observed distribution. Both the MC and observed events required for this work were already prepared with slight variations by the earlier studies in Refs. [391, 392, 393, 394, 395].

10.1 Detector response modelling: from neutrino interactions to filtered events




Figure 10.1: Typical DeepCore neutrino simulation scheme, taken from Ref. [393]: neutrino interactions are generated by GENIE (top left), the energy loss processes of their secondaries are simulated by PROPOSAL (in the case of μ^\pm 's, bottom left) or GEANT4 (in the case of hadrons, τ^\pm 's, and low-energy e^\pm 's and γ 's, centre), Cherenkov photons are propagated through a given ice model by CLSim (bottom centre), and custom software produces the detector response through successive stages, including triggering and filtering (right). Intermediate results are written to disk and may be discarded after completion. See text for details.

In a typical simulation chain used for low-energy analyses with DeepCore, all the relevant processes requiring simulation are contained in three consecutive steps, as illustrated by the flow chart in Fig. 10.1 [393]. The first comprises the neutrino interaction and the μ^\pm propagation, the second the propagation of any other secondary particles, their Cherenkov light production, as well as the photon propagation to the DOMs, while the third models the detector readout and applies the standard IceCube and DeepCore triggering and filtering algorithms. The following paragraphs summarise the essential assumptions made at each step.

Neutrino interaction

The GENIE MC event generator [396, 397] is used to model the interactions of atmospheric neutrinos with energies from several MeV to 1000 GeV. For each desired neutrino flavour, it generates an isotropic flux with energies distributed according to a user-defined power law. A cylindrical generation volume is defined as illustrated in Fig. 10.2—centred approximately on DeepCore, aligned with the propagation direction of the generated neutrino—with a flavour- and energy-dependent length L_{gen} and radius ρ_{gen} . Each neutrino is drawn randomly from a uniform distribution over the full generation surface upstream of DeepCore and




Figure 10.2: Illustration of low-energy neutrino generation scheme for DeepCore events with GENIE (not to scale).

is forced to interact. The point of interaction along the longitudinal axis is sampled from a uniform distribution over L_{gen} , and the interaction types implemented by GENIE are drawn with equal probability. The generation volume is chosen small enough to prevent a significant fraction of interactions to occur too far outside of the fiducial volume to result in an SMT3 trigger, but also large enough to preclude a significant fraction of unsimulated events when compared to measurement. The cylinder dimensions adopted by the simulation used in this work can be found in Ref. [393].

To each neutrino interaction, GENIE assigns a weight w_i that allows removing the above generation bias at a later step. This weight is the inverse of the probability density function (PDF) from which the neutrino energy E_i , zenith angle ϑ_i , and azimuthal angle φ_i are drawn, multiplied with the generation area ($\pi\rho_{\text{gen}}^2$) and the “true” probability for the generated interaction to occur (based on some nominal cross-section model). After normalising w_i to the total number of interactions generated, it can be multiplied with any neutrino flux model $\Phi(E_\nu, \vartheta_\nu, \varphi_\nu)$ to obtain one sample of the true neutrino interaction rate, evaluated at $(E_i, \vartheta_i, \varphi_i)$.

GENIE outputs a tree-like list containing the primary neutrino together with all of its secondaries, including their energies and directions at the interaction vertex. This

information in turn can be exploited by software which models the particle propagation in the target medium. Here we employ the GENIE release version 2.8.6. Newer versions as from 2.10.0 provide an improved accuracy of the high-energy cross sections for $E_\nu \gtrsim 100$ GeV due to the inclusion of the vector-boson propagator terms for DIS [398].

Propagation of secondaries and light production

Different tools and techniques simulate the propagation of all daughter particles of the initial neutrino interaction. The goal is to obtain an accurate prediction of the event's overall Cherenkov light emission profile. Where possible, analytic parameterisations are exploited to speed up the process.

Both ionisation and stochastic energy losses of any μ^\pm 's are modelled by the PROPOSAL software [399]. In order to prevent simulating a diverging number of stochastic energy losses, PROPOSAL implements a scheme by which processes in which the muon loses less than 0.5 GeV are treated as continuous and are parameterised analytically. Stochastic processes with energies above this threshold are considered individually, as is their Cherenkov emission.

e^\pm 's and γ 's with energies below 0.1 GeV and all hadrons and τ^\pm 's are propagated by the GEANT4 toolkit [400, 401, 402]. Here, the trajectory of any secondary or daughter Cherenkov emitter is split up into a series of light-emitting segments, each of which is assumed to have a constant velocity β and which radiates Cherenkov photons as described by Eq. (6.2). Due to the regular nature of energetic electromagnetic cascades, the Cherenkov light outputs of cascades initiated by e^\pm 's and γ 's with energies exceeding 0.1 GeV are calculated using parameterisations from Ref. [333].

Photon propagation

Individual photons created by the various Cherenkov emitters are subsequently propagated through the ice by the CLsim software package [403]. It exploits the publicly available OpenCL library [404] in order to allow efficiently tracking photons in multiple parallel threads on GPUs or multi-core CPUs.

The Cherenkov angle under which a photon is radiated by a track segment with constant velocity β is calculated based on the refractive index $n(z, \lambda)$ of ice given the photon's wavelength λ and its emission depth z in the ice, whereas the azimuthal emission angle is uniformly sampled. The number of absorption lengths determining the photon's total track length is sampled from an exponential distribution, as is

the spatial distance between two scattering instances. The propagation terminates once a photon, which carries information about its direction, position, wavelength, and time of incidence, intersects the surface of a DOM. The photon's properties are used in the last simulation step of Fig. 10.1 to generate the responses of all DOMs that are hit and the readout of the whole detector. A photon not incident on a DOM prior to covering its total track length is considered absorbed by the medium.

The properties of the pristine glacial ice in simulation here are described by the first model to incorporate a parameterisation of the optical anisotropy [357] (cf. Sec. 6.2.3).¹ Discrete deviations of the absorption and effective scattering coefficients from their nominal values are also simulated in order to include the optical properties as a source of systematic uncertainty in the analysis. More recent ice models improve on the above in particular by calibrating the ice with the full IceCube detector, that is, by flashing LEDs on more or all IceCube strings.

Photon detection

For each photon stored after the previous step of the simulation chain, one has to evaluate the probability for the creation of a photoelectron reaching the first dynode, giving rise to charge multiplication. This would be achieved most straightforwardly by discarding any photon that impinges on the upper DOM hemisphere, and otherwise scaling the overall optical DOM efficiency $\eta_0(\lambda)$ (cf. Sec. 7.3) by the value of the collection efficiency $\langle\eta_{\text{CE}}(\mathbf{x})\rangle$ for the given point of incidence \mathbf{x} on the photocathode. However, a different strategy is pursued in practice [360] in order to account for the presence of the refrozen borehole ice: the DOM is considered as a point-like receiver with a universal effective angular acceptance function, $g_{0,i}(\cos\Theta)$, which only depends on the cosine of the polar angle between the impinging photon's path and the string (z) axis, $\cos\Theta$. This function serves as a wavelength-independent rescaling factor to the position-averaged probability for a photon to produce a signal at the PMT anode:

$$p_{i,\text{hit}}(\lambda, \Theta) = \eta_{0,i}(\lambda) \langle\eta_{\text{CE}}(\mathbf{x})\rangle g_{0,i}(\cos\Theta) . \quad (10.1)$$

We assume that $g_{0,i}(\cos\Theta) \equiv g_0(\cos\Theta)$ is identical for all DOMs, whereas $\eta_{0,i}(\lambda)$ only varies by a wavelength-independent relative efficiency factor, $\eta_{0,i}(\lambda) \equiv k_i\eta_0(\lambda)$. Figure 10.3 shows the product $\eta_0(\lambda) \langle\eta_{\text{CE}}(\mathbf{x})\rangle$ as a function of wavelength for a regular-QE DOM, together with a 10% uncertainty band [364]. Dedicated simulation runs for overall optical efficiencies scaled by between 88% and 112% are available.

¹The model is referred to as “SPICE LEA” by the IceCube collaboration [405].




Figure 10.3: Position-averaged probability for a photon of wavelength λ incident on the lower hemisphere of a regular-QE DOM to produce a charge at the PMT anode, together with a 10% uncertainty band, whose impact is probed by means of multiple dedicated simulation runs. For a DeepCore high-QE DOM, the function is scaled up by 35%.




Figure 10.4: Relative angular DOM acceptance as a function of the polar photon angle $\cos \Theta$ as measured in the laboratory (dashed line) and modified by absorption and scattering properties of the borehole ice surrounding the strings (solid). The range of acceptances resulting from all simulated hole-ice variations is indicated by the uncertainty band.

The angular acceptance function of a DOM is an effective description that accounts for both the intrinsic PMT angular acceptance measured in the laboratory as well as scattering and absorption modifications of the hole ice surrounding the strings [362]. We parameterise the angular acceptance as a function of $x \equiv \cos \Theta$ as [406]

$$g_0(x) = C \left[0.34 (1 + 1.5x - 0.5x^3) + p_1 x (x^2 - 1)^3 + p_2 e^{10(x-1.2)} \right], \quad (10.2)$$

where C is a normalisation constant. The dimensionless parameter p_1 controls the lateral acceptance and is relatively well constrained by LED flasher calibration studies, whereas p_2 , also dimensionless, allows for large acceptance variations in the forward region ($\cos \Theta \approx 1$). The latter cannot be easily probed via existing in-situ measurements due to the lack of bright upward-pointing LEDs [362].

Several simulation runs allow modelling the systematic impact of hole-ice uncertainties, each performed for a specific pair (p_1, p_2) and with the normalisation factor C chosen such that the integrated acceptance remains unchanged.² The nominal model corresponds to $p_1 = 0.25, p_2 = 0$. It is depicted in Fig. 10.4 together with the

²We require the integral $\int_{-1}^1 g_0(x) dx = 0.68$ for all values of p_1 and p_2 .

envelope of the generated systematic variations, covering the ranges $p_1 \in [0.15, 0.35]$ and $p_2 \in [-5, 2]$.

DOM noise

Not all photoelectrons emitted from the photocathode are caused by photons originating from physics events: “dark noise” occurs at a typical rate of 560 Hz and 780 Hz for regular- and high-QE DOMs, respectively [364]. This rate has a variety of contributions, including thermal electronic noise³ and scintillation and fluorescence emission from radioactive decays in the PMT glass and in the pressure housing. Two types of noise are distinguished—the “uncorrelated” (Poissonian) and “correlated” (non-Poissonian) components—with similar overall contributions to the total rate. Compared to the pure Poissonian expectation, the correlated component leads to a significant excess of hits separated by short time intervals below some 100 μ s, for which it is the dominant source. A large fraction of uncorrelated noise hits can be attributed to Cherenkov emission from β -decays of ^{40}K in the glass pressure sphere, at an estimated rate of 100 Bq per DOM. The origin of the uncorrelated noise has not yet been unambiguously identified. The leading explanation is scintillation luminescence of the glass, with trace elements such as cerium acting as active scintillators [408]. Resulting noise hits typically occur in “bursts”—temporal clusters of hits separated by short time intervals up to the ms range—with a rate whose temperature dependence is empirically well described by an exponential [409, 410].

All the known noise sources above are simulated using the Vuvuzela software [411]. In Vuvuzela, the noise model for each DOM is characterised by five parameters: (i) a thermal noise rate, (ii) a radioactive decay rate, (iii) a mean number of scintillation hits for each decay, and (iv) the mean and (v) standard deviation of a log-normal distribution of the time intervals between successive scintillation hits within a burst. In order to obtain the parameters for each individual DOM, the noise model was fit to ten minutes of raw, untriggered, IceCube data taken in early 2015, as detailed in Ref. [412].⁴ Crucially, hits with temporal separations below 2 μ s are included.

³This type of noise is known as “thermionic emission”, first discovered and characterised by Richardson [407].

⁴The underlying “HitSpool” data stream was developed to improve IceCube’s data acquisition for Supernova detection [410].




Figure 10.5: PMT jitter distribution, described by a right-skewed Gumbel distribution shifted to the right by 0.15 ns and with a scale parameter of 1.92 ns.




Figure 10.6: SPE response distribution, described by the superposition of a Gaussian centred at 1 PE with a standard deviation of 0.29 PE and an exponential with a rate parameter of 0.51 PE.

PMT response

Once—both physics- and noise-induced—photoelectrons ejected from the photocathode have been generated, these are converted into amplified pulses at the PMT anode [413]. Each photoelectron undergoes the effects of the amplification process initiated by the PMT dynode chain. This process can result in four different pulse types: “regular” pulses, “prepulses”, “late” pulses, and “afterpulses”.

The variation in the amount of time an (amplified) photoelectron that initiates a regular pulse needs to reach the anode is sampled from the distribution depicted in Fig. 10.5, which has a standard deviation of ~ 2 ns. The observed charge of the resulting pulse is distributed according to the superposition of a Gaussian peaked at the charge-equivalent of one photoelectron (1 PE) and a subdominant exponential component, leading to a mean charge of $\bar{q}_{\text{SPE}} \equiv 0.85 \text{ PE}$,⁵ as shown in Fig. 10.6. This SPE charge response, $p_c(q|1)$, is taken as a representative template across all DOMs, based on dedicated laboratory measurements of a large number of SPE charge distributions taken at different gains which exhibit only small deviations [360]. Recent calibration efforts have allowed replacing the universal SPE charge template in Fig. 10.6 with individual charge response distributions for each DOM [414, 358, 415]. Here, significant improvements in the modelling are expected especially for the low-charge re-

⁵Note that the amplified charge Q follows as the product of photoelectron charge and PMT gain. For example, $Q = 1.6 \text{ pC}$ for a typical gain of 10^7 .

gion, where a second, steeply-falling exponential provides a better fit to each DOM's response distribution.

The other three types of pulses in general have different arrival-time and charge distributions than a regular pulse and are produced with a combined probability of only $\sim 10\%$. As a result, their number is extremely low for typical DeepCore event energies. Details on these types of pulses can be found in Ref. [360].

PMT saturation is applied to each pulse using a universal parameterisation of the PMT current response measured in the laboratory [360]. It describes the (gain-dependent) reduction of the "ideal" instantaneous anode current, leading to a deviation from the linear relationship between the PMT output and the incident number of photoelectrons. Since this deviation only sets in above an instantaneous current of about 50 mA—corresponding to a rate of 30 PE/ns—saturation effects are negligible at typical DeepCore event energies.

DOM launching and signal digitisation

Once PMT pulses with associated times and charges have been output by the PMT response simulation step, the signal processing and waveform capture circuitry implemented on the DOM MB (cf. Sec. 7.2.1) are modelled [416]. After the charges of the pulses on all DOMs have been compared to each DOM's discriminator threshold level, the local-coincidence logic is applied. In particular, if at least one of the nearest or next-to-nearest neighbouring DOMs signals such a threshold crossing, an HLC launch is created (unless the considered DOM is still busy processing an earlier signal). This signals the digitisation of the PMT waveform, sampled by the ATWDs and the fADC, using waveform constants retrieved during the detector calibration campaign (cf. Sec. 7.3).

At this point, the simulated detector readout is qualitatively indistinguishable from real (untriggered) detector data. Hence, one can proceed to simulate the online triggers and filters in place at the South Pole. Other standard processing steps applied to observed detector readouts are also reproduced, such as the pulse extraction and noise removal described in Secs. 7.2.4 and 7.2.5, which are both prerequisites for the DeepCore online filter.

10.2 Undoing the detector response: event reconstruction

The event reconstruction procedure outlined below is applied identically to the detector readouts of observed and simulated events in order to arrive at estimators of neutrino properties expected to carry imprints of the physics hypotheses under study. In this analysis, these estimators are the neutrino energy, zenith angle, and an event classification category, whose joint distribution is sought. Sometimes, faster “first-guess” and less precise reconstruction algorithms⁶ are employed in order to create the final analysis sample that is then subjected to the more computationally expensive reconstruction. How such an approach applies to the event selection for this analysis is summarised in Chapter 11.

The estimation of neutrino properties at GeV-scale energies presents distinct challenges from that at higher energies of TeV and above [369], to which the primary IceCube array is adapted. While the spatial extents of low-energy events allow these to be fully contained (within DeepCore or IceCube), the energy deposited in Cherenkov emitters is not in general a suitable proxy of the true neutrino energy, owing to the significant relative average suppression and comparably large statistical fluctuations of the overall Cherenkov light yield for hadronic cascades compared to electromagnetic cascades (cf. Sec. 6.2). In addition, outgoing neutrinos carry away energy in the case of NC interactions or $\bar{\nu}_\tau$ CC interactions. A further complication stems from the optical scattering properties of ice, which smear out the differences in the longitudinal, transverse, and angular Cherenkov light emission profiles that exist between electromagnetic and hadronic cascades, even over small distance scales of only a few tens of metres [418]. Accordingly, established DeepCore event reconstruction techniques do not attempt to distinguish between the two cascade types. Instead, a single cascade template is assumed, which also does not account for the opening angle between the electromagnetic and the hadronic cascade in the case of $\bar{\nu}_e$ CC or $\bar{\nu}_\tau$ CC interactions. This opening angle decreases with increasing energy of the incoming neutrino, and is larger when the inelasticity of the interaction approaches zero or unity [418]. When considering $\bar{\nu}_\mu$ CC interactions, the scattering of the Cherenkov photons emitted by the hadronic cascade in practice justifies reducing the number of degrees of freedom by assuming the cascade axis to be aligned with the direction of the outgoing μ^\pm .

⁶Essential examples with extensive usage throughout the IceCube collaboration can be found, for example, in Ref. [417].

Hybrid template likelihood fit

Event hypothesis Based on the considerations above, DeepCore events are tested for their compatibility with the hypothesis of a single cascade accompanied by a minimum-ionising μ^\pm track [391, 419]. The point of origin at the neutrino interaction vertex is described by the associated time t_0 and coordinate vector \mathbf{x}_0 . The direction of the collinear cascade and track is described by the azimuthal angle ϕ and the zenith angle ϑ . The angles give the direction of the radial vector pointing back toward the source of the incident particle, i.e., a vertically downgoing particle has a zenith angle of $\vartheta = 0$, while one that is vertically upgoing has a zenith angle of $\vartheta = \pi$. The cascade is additionally characterised by its visible energy $E_{\text{casc}}^{\text{vis}}$, and the track by its length L_{track} (which uniquely fixes the expected total energy of a minimum-ionising μ^\pm).

Poisson likelihood The event hypothesis $\boldsymbol{\theta} \equiv (t_0, \mathbf{x}_0, \phi, \vartheta, E_{\text{casc}}^{\text{vis}}, L_{\text{track}})$ is adapted to an observed ensemble of extracted pulse charges \mathbf{q}^{obs} by means of the Poisson likelihood function

$$L(\mathbf{q}^{\text{obs}}|\boldsymbol{\theta}) = \prod_{i=1}^{N_{\text{DOM}}} \prod_{j=1}^{N_t^i} p_{\text{Pois}}[q_{i,j}^{\text{obs}}|q_{i,j}^{\text{exp}}(\boldsymbol{\theta})], \text{ with} \quad (10.3)$$

$$p_{\text{Pois}}(q_{i,j}^{\text{obs}}|q_{i,j}^{\text{exp}}) = \frac{(q_{i,j}^{\text{exp}})^{q_{i,j}^{\text{obs}}}}{\Gamma(q_{i,j}^{\text{obs}} + 1)} e^{-q_{i,j}^{\text{exp}}}, \quad (10.4)$$

where the first product in the likelihood function at the top is taken over all considered DOMs, and the second over a variable number of time intervals into which the pulse charges are binned. The likelihood function depends on the hypothesis parameter values through the ensemble of expected charges $\mathbf{q}^{\text{exp}}(\boldsymbol{\theta})$ (dependence on $\boldsymbol{\theta}$ omitted for clarity in the Poisson probability mass function (PMF) p_{Pois}). The Γ function in the denominator of Eq. (10.4) allows for a continuous extension of the Poisson PMF, thereby accounting for the non-integer nature of pulse charges.

The observed pulse ensemble \mathbf{q}^{obs} serving as the data in Eq. (10.3) is chosen to be that generated by a standard processing step applied to each event before filtering (cf. Fig. 10.1 for simulation). This pulse ensemble follows from applying at most three iterations of the “seeded R - T ” noise removal explained in Sec. 7.2.5 with $R_{\text{max}} = 150$ m and $T_{\text{max}} = 1$ μs to an initial set of HLC pulses, with the additional requirement that for each such seed there are two or more additional HLC pulses that satisfy the




Figure 10.7: The expected PMT charge distribution of an IceCube DOM (solid) for some expectation value (denoted as λ in the figure, but μ_γ in the text) of the number of photons incident on the photocathode is the convolution of the Poisson PMF $p_{\text{Pois}}(n|\lambda)$ for the ejection of n photoelectrons (dotted) with the PMT charge response distribution $p_c(q|n)$ to n photoelectrons (not shown). The dashed line represents the approximation that the charge expectation value and the photon number expectation value λ are related by \bar{q}_{SPE} . Figure taken from Ref. [369].

R-T criterion. If no such seed exists, the algorithm falls back to starting from all HLC pulses. In addition, a given pulse returned by the algorithm is only kept if it falls into a time window $[t_{\text{trigger}} - 4 \mu\text{s}, t_{\text{trigger}} + 5 \mu\text{s}]$ around the SMT3 trigger time, in order to further suppress pulses due to dark noise.

Cherenkov light-yield templates The lack of an accurate analytic functional form of the photon arrival time distribution for arbitrary configurations of Cherenkov emitters and DOMs—which necessitates the tracking of individual photons in simulation—poses an even greater challenge in reconstruction. Since reconstruction has to be applied to large numbers of (in particular MC) events, the direct simulation of the detector response to many different realisations of θ is computationally prohibitive.

An alternative approach that is both computationally feasible and reasonably accurate is the use of precomputed light-yield templates [369]. In the case of a cascade it is justified to employ a single universal template because of (i) the similarity of the observed light yield from an electromagnetic and from a hadronic cascade with the same visible energy, (ii) the nearly point-like nature,⁷ (iii) the nearly universal angular emission profile, and (iv) the linear relationship between the energy in the electromagnetic sector and the emitted number of Cherenkov photons. The method

⁷Both (i) and (ii) apply at the scale of DeepCore’s and, in particular, IceCube’s instrumentation density.

of choice is therefore to precalculate the mean normalised photon arrival time distribution $p_\gamma(t)$ —with t corresponding to the time delay with respect to the time of arrival of an unscattered photon—together with the mean total number of photons, μ_γ , expected at any observer position relative to an electromagnetic cascade with some reference energy E_0 occurring at a given depth and under a given zenith angle (assuming azimuthal symmetry and lateral translation invariance). By repeating the procedure for many different cascade depths and zenith angles and creating a multidimensional spline fit, one obtains a smooth representation of $p_\gamma(t)$ and μ_γ , which can be evaluated for any desired set of cascade properties (position, energy, direction) and DOM positions [369]. In fact, since a regular-QE DOM is assumed to be located at each observer position in the template production process, μ_γ simultaneously accounts for the probability of a photon to produce a signal at the PMT anode, Eq. (10.1), with the wavelength-independent relative efficiency factor $k_i = 1$. Figure 10.7 demonstrates that the assumption of a purely Poissonian process with the extension to real numbers adopted in Eq. (10.3) is a reasonable approximation for any number of photons incident on the photocathode, given that the charge expectation is calculated as

$$q_{i,j}^{\text{exp,casc}} = \bar{q}_{\text{SPE}} k_i \frac{E_{\text{casc}}}{E_0} \mu_\gamma \int_{t_j}^{t_{j+1}} dt p_\gamma(t) , \quad (10.5)$$

where E_{casc} is the energy in the electromagnetic cascade and t_j and t_{j+1} are the bounds of the considered time interval.

Similarly, templates are available for a minimum-ionising μ^\pm with some reference energy $E_{0,\mu}$ (3.3 GeV in our case), corresponding to an average track length of

$$\langle L_{\text{track}} \rangle = \frac{E_{0,\mu}}{\left\langle -dE_\mu/dx \right\rangle_{\text{ice}}} \stackrel{\text{Eq. (6.3)}}{\approx} 4.2 \text{ m} \left(\frac{E_{0,\mu}}{\text{GeV}} \right) . \quad (10.6)$$

Any integer number $N_\mu = 0, 1, 2, \dots$ of these templates can be combined to obtain the charge expectation due to a μ^\pm of energy $E_\mu = N_\mu E_{0,\mu}$,

$$q_{i,j}^{\text{exp},\mu^\pm} = \bar{q}_{\text{SPE}} k_i \sum_{k=1}^{N_\mu} \mu_{\gamma,k} \int_{t_j}^{t_{j+1}} dt p_{\gamma,k}(t) . \quad (10.7)$$

In likelihood fits to observed and simulated events for this analysis, the cascade predictions $q_{i,j}^{\text{exp,casc}}$ have been generated under the “SPICE MIE” ice model intro-

duced in Sec. 6.2.3, while the μ^\pm predictions $q_{i,j}^{\text{exp},\mu^\pm}$ follow from a slightly earlier model [420]. While both models are based on the same calibration study, the latter differs in the employed fitting technique and lacks the degree of freedom in the shape of the scattering function. In the most general case, including a contribution due to dark noise that depends on the i th DOM's overall noise rate $\mu_{\text{noise},i}$, the charge expectation for any combination of DOM and time bin is then the sum of three source terms,

$$q_{i,j}^{\text{exp}}(\boldsymbol{\theta}) = q_{i,j}^{\text{exp,casc}}(\boldsymbol{\theta}) + q_{i,j}^{\text{exp},\mu^\pm}(\boldsymbol{\theta}) + q_{i,j}^{\text{noise}}(\mu_{\text{noise},i}). \quad (10.8)$$

Solving for maximum-likelihood estimators The Pegleg [419] event reconstruction method yields the neutrino estimators entering our statistical inferences. It builds on the various components of the Poisson likelihood approach discussed in the preceding paragraphs. These are applied widely throughout the IceCube collaboration, albeit with different event hypotheses [369]. Pegleg employs the MULTINEST [421, 422, 423] nested sampling algorithm in order to minimise the negative logarithm of the likelihood (“log-likelihood”) in Eq. (10.3), $-\ln L(\boldsymbol{q}^{\text{obs}}|\boldsymbol{\theta})$. MULTINEST is a Bayesian inference tool, capable of efficiently sampling model evidences and posterior parameter probability distributions. As such, it is not adapted to finding the global minimum of a negative log-likelihood, as is a common goal in frequentist inference. Nevertheless, it has proven to be an efficient numerical minimisation algorithm when applied to the multimodal and irregular log-likelihood hypersurfaces encountered here [391, 424]. Pegleg makes use of MULTINEST only to sample the six-dimensional parameter subspace $\boldsymbol{\theta}_{\text{MN}} \equiv (t_0, x_0, \phi, \vartheta)$. For each sample of $\boldsymbol{\theta}_{\text{MN}}$, the number of minimum-ionising μ^\pm track segments N_μ is scanned manually, and for each hypothesised value the maximum-likelihood estimator \hat{E}_{casc} is determined with the Newton method as laid out in Ref. [419]. The efficiency and robustness of this three-layered minimisation compared to letting MULTINEST explicitly sample the full eight-dimensional parameter space [391] led to the adoption of Pegleg for the reconstruction of events in the data sample used in this work [425]. Its performance is shown in Chapter 12.

11

Event sample

The NSI measurement presented in this work examines events detected by DeepCore over the period of time between April 2012 and May 2015, corresponding to the second, third, and fourth seasons of data taking with the full, 86-string IceCube detector. The time range was chosen in this way due to effects which would have required significant changes to extant simulation routines: prior to this period, the DOMs’ dark noise rates were still stabilising [364], whereas after it a recalibration was applied to the observed PMT charges [419].

For the measurement to be feasible, one first needs to identify atmospheric neutrino candidates, which only contribute negligibly to the total DeepCore trigger and filter rates. Due to the large effort involved in developing an event selection scheme with a good sensitivity to atmospheric neutrino oscillations, the preceding DeepCore oscillation studies were performed on variants of two (significantly overlapping) event samples.

11.1 Context

The data sample used here is an extension of that labelled “ \mathcal{B} ” in Ref. [362], which introduces a second, higher-statistics sample “ \mathcal{A} ” that deviates from \mathcal{B} by accepting a larger fraction of non-neutrino events. Sample \mathcal{B} was discussed in the context of a measurement of standard atmospheric neutrino oscillations [391] and allows placing competitive constraints on the atmospheric oscillation parameters $|\Delta m_{31}^2|$ and $\sin^2 \theta_{23}$ [426] (as does sample \mathcal{A} [362]). The results of a $\bar{\nu}_\tau$ appearance search conducted using the same set of events are discussed in Refs. [394, 362]. Both samples have also been employed to probe the NMO [419, 392, 427], albeit with modified acceptance regions in the reconstructed neutrino energy and cosine zenith. Finally, the

first measurement of the atmospheric neutrino flux spectrum using DeepCore [395] was performed by accepting reconstructed neutrino energies between 5.6 GeV and 180 GeV and by relaxing the event containment criteria of sample \mathcal{B} to increase event statistics above reconstructed neutrino energies of around 50 GeV. A similar approach is adopted by this NSI search in order to enhance the sensitivity of sample \mathcal{B} to NSI couplings responsible for deviations from standard three-flavour oscillations at energies of the order of 100 GeV (in particular $\epsilon_{e\mu}^{\oplus}$ and $\epsilon_{\mu\tau}^{\oplus}$, cf. Chapter 9).

11.2 Selection criteria

Data taking by the IceCube detector proceeds in the form of a sequence of time intervals called “runs” of typically and at most 8 h duration, with short breaks in between. Some runs serve calibration purposes, while others suffer from parts of the detector (individual DOMs or whole strings) dropping out from the data acquisition. Instead of simulating runs with such irregular configurations, it is more practical to exclude these from the analysis altogether. This approach is feasible because IceCube’s high reliability combined with the low frequency of calibration runs results in a tiny fraction of detector exposure time (“livetime”) lost. Section 11.2.1 discusses the criteria that runs have to satisfy in order to enter this NSI search. The criteria for selecting events from a given run are outlined in Sec. 11.2.2.

11.2.1 Data-taking runs

A run must meet three criteria to be considered: (i) a duration of at least 1 h, (ii) 86 active strings, and (iii) at least 5 380 active DOMs (including the 324 IceTop DOMs). For each of the three data-taking seasons considered, Fig. 11.1 shows a histogram of the run durations and event numbers contributed to the final-level event sample (after undergoing the full event selection below). The vast majority of runs have a duration of close to 8 h. No more than 32 events in any given run end up in the final sample. The typical number of events per run is ~ 15 . The livetimes of all considered runs add up to 2.8 years.

11.2.2 Events

The iterative application of event selection criteria (“cuts”) ensures that the final event sample is dominated by atmospheric neutrinos and prevents the wasteful application of computational resources to events that are of little to no physics value.

The selection levels 1 to 3 are common to most DeepCore analyses and consist of triggering and filtering criteria. Three additional levels originally tailored to the mentioned oscillation measurement yield additional background suppression power and improve the overall quality of the selected events. As the elements of this selection process have been laid out in great detail in Refs. [391, 362], we merely provide a condensed text summary below. However, Table 11.1 specifies all of the criteria that are applied at each selection level and designates their main aims: the removal of noise events, the removal of atmospheric μ^\pm events, or the enhancement of event quality¹. Furthermore, Fig. 11.2 illustrates how these criteria impact the expected rates of the various MC event types, starting from the combined event rates surviving at least one filter.

First (level 1), only events passing the DeepCore SMT3 trigger are considered, as explained in Sec. 7.2.3.

At level 2, the DeepCore online filter detailed in Sec. 7.2.6 removes the most obvious atmospheric μ^\pm candidates among the triggered events, leading to a loss of only roughly 1 % of atmospheric neutrinos.

At level 3, a set of general cuts developed by the collaboration for low-energy DeepCore analyses only keep events with spatial and temporal hit distributions that are unlikely to be the result of random noise or atmospheric μ^\pm 's. These cuts reduce the filtered data rate by one additional order of magnitude, down to $\mathcal{O}(1 \text{ Hz})$.² Here, restrictions are placed for example on the charge detected near the top of DeepCore and in the veto region, both of which naturally have power to discriminate between neutrino interactions and penetrating atmospheric μ^\pm 's. Also, we exploit the fact that the fraction of charge deposited early in an event is in general higher when an atmospheric neutrino interaction causes the launch of the SMT3 trigger, as sketched in Fig. 11.3 [428].

Similar to level 3, variables used at level 4 do not rely on computationally intensive algorithms. Instead, they are directly derived from the observed pattern of signals in the fiducial and veto regions. Among others, they demand that at least three DeepCore DOMs record signals due to minimally scattered Cherenkov photons ("direct" hits), that at least eight DeepCore DOMs record signals that are not obviously due to random noise, and a charge-equivalent of at least 7 PE associated with clusters of pulses (in space and time) within DeepCore. Two refined versions of the DeepCore online filter are also applied. The first employs the same range

¹"Quality" refers to desirable event properties, which benefit the precision of the reconstruction.

²The event selection used in this work employs a slightly modified version of these common criteria.

of apparent velocities connecting pulses on a more central subset of SMT3 DOMs to those observed throughout an extended veto region, and requires the total charge of pulses corresponding to such velocities to amount to no more than 5 PE. The second represents a complementary strategy of limiting the total charge of isolated signals due to an incoming μ^\pm in a region of time differences and spatial distances that suggest a causal connection with the observed signals within DeepCore, but here relative to the HLC pulse closest in time to the SMT3 time, as introduced in Ref. [429]. The integrated charge within this region is limited to 7 PE. The idea is sketched in Fig. 11.4.

The following level 5 applies a cut on the score calculated for each event by a “boosted decision tree” (BDT). The BDT was trained on a set of eleven event variables, some of which are closely related or identical to those employed in previous selection levels, and some of which are based on the results of fast event reconstruction algorithms. The resulting multivariate cut³ further reduces the atmospheric μ^\pm contamination of the sample.⁴

At level 6, the overall event rate is reduced sufficiently for the computationally expensive reconstruction algorithm of Sec. 10.2 to be applied to all remaining events.⁵ Cuts are placed on the “starting” position—the reconstructed position of the interaction vertex—but not on the “stopping” position (which in turn is specific to the event selection used for the NSI search at hand), as sketched in Fig. 11.5. We require the reconstructed interaction vertex to be located within a volume that is inset with respect to DeepCore’s fiducial volume, namely within a cylinder of 100 m radius and 200 m height, with a lower base at a depth of -475 m, and which is extended by a spherical cap with a radius of 125 m and a height of 50 m at the top. As a final means of suppressing background, “corridors” of sparse detector instrumentation through which an atmospheric μ^\pm could have penetrated, depicted in Fig. 11.6, are identified based on the event’s signal pattern. As explained in Ref. [428], if there is any such corridor with two or more nearby DOMs detecting direct signals (with respect to the respective hypothetical μ^\pm trajectory), the event is discarded.

At the final level of the event selection, the values of two of the three analysis variables are restricted. Since the range of MC neutrino energies is limited to

³While the original iteration of the event selection kept all events with a score greater than 0.1 (cf. Ref. [391]), this threshold was later set to a value of 0.2 instead [425, 362].

⁴Atmospheric μ^\pm ’s evading straight veto cuts and making it to late selection levels of DeepCore event samples are typically referred to as “sneaky”.

⁵The layered approach of the reconstruction ansatz presented in Sec. 10.2 differs from that used in the original iteration of the event selection. It was adopted for computational efficiency reasons [425] prior to the publication of Ref. [426].

$E_\nu \in [1 \text{ GeV}, 1 \text{ TeV}]$, we have to limit reconstructed energies to a range that is significantly smaller, here chosen as $E_{\text{reco}} \in [5.62 \text{ GeV}, 100 \text{ GeV}]$. Furthermore, an event is only kept if the ratio between the likelihood $L_{\text{cascade+track}}$ at the best-fit parameter values for the combined (cascade and track) hypothesis and that for the single-cascade hypothesis, L_{cascade} , satisfies $\ln(L_{\text{cascade+track}}/L_{\text{cascade}}) \in [-3, 1000]$. As the single-cascade hypothesis is nested within the more general combined hypothesis, a smaller likelihood of the event under the combined hypothesis is the result of the numerical minimisation routine failing to locate the global minimum in $-L$. Conversely, an exceedingly high relative compatibility of a given event with the combined hypothesis could indicate an inaccurate minimisation under the single-cascade hypothesis. Such a likelihood preference is bounded from above.

The three histograms in the right column of Fig. 11.1 show the total number of events resulting from applying all of the above criteria to the selected runs of a given data-taking campaign. Summed over all three campaigns, 47 855 events are observed.




Figure 11.1: Livetimes and numbers of selected events for all runs with at least one selected event, split up according to data-taking seasons with the full 86-string IceCube detector (top row: 2012/2013, middle row: 2013/2014, bottom row: 2014/2015).

11.2 Selection criteria

Level	Criterion	Requirement	Target
1 + 2	DeepCore SMT3 and online filter		noise, μ^\pm
3	max. clustered charge in 300 ns	> 2 PE	noise
	max. no. of clustered pulses in 300 ns	> 2	noise
	max. no. of causal pulses in 48 solid-angle bins	> 2	noise
	charge fraction within 600 ns after 2nd hit	> 0.4	μ^\pm
	total charge above -200 m within $2 \mu\text{s}$ before trigger	< 12 PE	μ^\pm
	total causal veto charge	< 7 PE	μ^\pm
	ratio of veto charge to fiducial charge	< 1.5	μ^\pm
	depth of first pulse	< -120 m	μ^\pm
4	no. of DOMs in fiducial seeded R - T search	≥ 8	noise
	max. clustered charge in fiducial seeded R - T search	≥ 7	noise
	std. dev. of depth-COG of cleaned pulses	$\in [7 \text{ m}, 100 \text{ m}]$	noise
	std. dev. of time-COG of cleaned pulses	$\leq 1 \mu\text{s}$	noise
	spacetime separation of 1st and 4th charge quartiles	$\in [-(400 \text{ m})^2, 0 \text{ m}^2]$	noise
	no. of DOMs detecting unscattered light	≥ 3	noise
	total causal charge in extended veto region	≤ 5 PE	μ^\pm
	complementary causal track veto charge	≤ 7 PE	μ^\pm
	depth of earliest HLC DOM	$\in [-475 \text{ m}, -200 \text{ m}]$	μ^\pm
	radial distance of earliest HLC DOM	≤ 150 m	μ^\pm
	depth of COG of pulses of 1st charge quartile	$\in [-475 \text{ m}, -150 \text{ m}]$	μ^\pm
	radial distance of COG of pulses of 1st charge quartile	≤ 150 m	μ^\pm
5	11-variable BDT score	≥ 0.2	μ^\pm
6	reconstructed starting depth	≥ -475 m	quality
	reconstructed starting radial distance	≤ 100 m	quality
	reconstructed spherical radial distance from centre above -350 m	≤ 125 m	quality
	max. no. of hit DOMs in "blind" corridors	≤ 1	μ^\pm

Table 11.1: Summary of the event selection criteria employed in this search for NSI, except for requirements on analysis variables. Exhaustive explanations of the various variables on which cuts are placed can be found in Refs. [391, 362].




Figure 11.2: Event rates after each level of event selection criteria for the various MC event types. The (accidental) noise background is negligible after level 4. The atmospheric μ^\pm background contribution is reduced to below that from all atmospheric neutrinos and anti-neutrinos by the criteria imposed at level 5. See text for details.




Figure 11.3: Sketch of the light depositions from a $\bar{\nu}_\mu$ interacting in DeepCore and a μ^\pm passing through DeepCore, with a significant fraction of charge detected near the interaction point of the neutrino. Figure taken from Ref. [428].

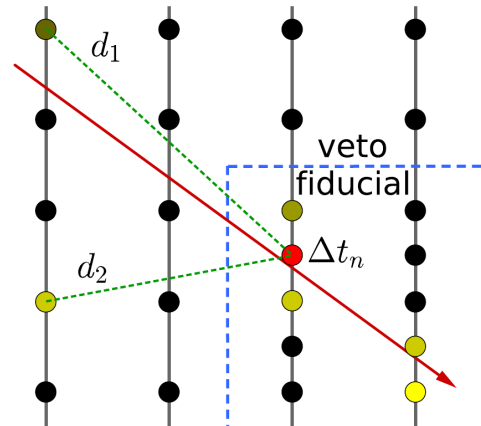


Figure 11.4: Sketch of isolated signals in the DeepCore veto region from random noise and an atmospheric μ^\pm , with distances and time differences d_i and Δt_i to the signal causing the SMT3 trigger (red circle). Figure taken from Ref. [428].




Figure 11.5: Illustration of final-level event containment requirements, highlighting deviations from the original set of criteria: only the interaction vertex (filled circles) is required to be contained.




Figure 11.6: "Blind" corridors, represented by dashed arrows, of a hypothetical DeepCore event with a first strong signal detected by the DeepCore string marked in blue. Figure taken from Ref. [428].

12

Simulation predictions

Correctly weighted, the reconstructed MC events passing the selection criteria of the preceding chapter allow predicting the expected event distribution for any set of NSI parameter values. In order to avoid a measurement that wrongly interprets systematic deviations from the nominal MC model as a signature of NSI, the combined effects on the event distribution from systematic uncertainties on the atmospheric neutrino fluxes, standard oscillation parameters, neutrino interaction cross sections, or the detector response have to be modelled. Similar considerations apply to the subdominant atmospheric μ^\pm background distribution. Below, we detail the computational methods underlying the calculation of the expected event distribution in the presence of several known sources of systematic uncertainty. A high-level discussion of the event sample's properties is followed by an investigation of the possible imprints on it left by NSI. An overview of the implementation details for all considered sources of systematic uncertainty concludes the chapter.

12.1 Event-by-event reweighting

Reweighting an ensemble of weighted MC events simulated according to a nominal set of model assumptions represents an efficient path toward the generation of expected event distributions for many different realisations of the physics parameters under study. The approach is conceptually feasible in our case since the details of the neutrino interaction in the ice, the detector response to the interaction, and the subsequent event reconstruction process are independent of assumptions about NSI in neutrino propagation. As a result, the range of NSI hypotheses that can be probed is not limited by having to constantly repeat these time-consuming steps.

In more detail, each reconstructed event R_i that satisfies the selection criteria is

the result of exactly one generator-level neutrino of flavour α , G_i^α . It has the nominal (GENIE) weight w_i^α attached to it, which does not account for neutrino oscillations. The event is reweighted as¹

$$w_i^\alpha \rightarrow w_i^\alpha \Phi_{\text{osc}}^\alpha(\boldsymbol{\theta}_{\text{osc}}; E_i, \vartheta_i, \phi_i) t_{\text{live}}, \quad (12.1)$$

where $\Phi_{\text{osc}}^\alpha(\boldsymbol{\theta}_{\text{osc}}; E_i, \vartheta_i, \phi_i)$ is the flux of atmospheric neutrinos of flavour α (after oscillations) at the location of the detector, evaluated at the energy E_i , zenith ϑ_i , and azimuth ϕ_i of the generator-level neutrino G_i^α , and t_{live} is the effective livetime of the event sample. This flux depends on a set of parameters $\boldsymbol{\theta}_{\text{osc}}$ that affect neutrino oscillation probabilities, of which NSI parameters are a subset. Φ_{osc}^α is represented as the oscillation probability-weighted sum of unoscillated fluxes of atmospheric electron and muon neutrinos at the detector location:

$$\Phi_{\text{osc}}^\alpha(\boldsymbol{\theta}_{\text{osc}}; E_i, \vartheta_i, \phi_i) = \sum_{\beta=e,\mu} P_{\beta\alpha}(\boldsymbol{\theta}_{\text{osc}}; E_i, \vartheta_i) \Phi_{\text{unosc}}^\beta(E_i, \vartheta_i, \phi_i). \quad (12.2)$$

The oscillation probabilities $P_{\beta\alpha}$ are taken to be independent of azimuth ϕ (cf. Sec. 5.1.2.2).

Nominal atmospheric neutrino fluxes

Nominal fluxes of atmospheric electron and muon neutrinos at the Geographic South Pole are obtained from MC simulations performed by Honda *et al.* [281]. Their results are tabulated in three dimensions, characterising the fluxes as a function of E_ν , $\cos \vartheta$, and ϕ . The simulations span the energy range from $E_{\text{min}} = 0.1$ GeV to $E_{\text{max}} = 10$ TeV and the full 4π solid angle, employing 101 logarithmic bins in E_ν , 20 bins in $\cos \vartheta$, and 12 bins in ϕ . Seasonal variations are included via the ‘‘NRLMSISE-00’’ model of the atmosphere [430]. These calculations have been made public via tables providing the flux average across each bin. In addition, two-dimensional tables are available with the azimuthal dimension averaged, as well as one-dimensional tables of fluxes averaged over the full solid angle. This analysis makes use of the year- and azimuth-averaged flux calculations (at the neutrino energies of interest, the azimuthal neutrino flux variations at the South Pole are sufficiently small, as are seasonal variations [392]).

¹Keep in mind that the original weight w_i^α has the inverse dimensions of a particle number per unit area, energy, and solid angle, i.e., we could express it as $w_i^\alpha = (\Phi_{\text{orig}}^\alpha(E_i, \vartheta_i, \phi_i) t_{\text{live}})^{-1}$. Equation (12.1) would then look like a ‘‘common’’ reweighting relation.

In order to not be confined to the granularity offered by these flux tables, an integral-preserving spline interpolation is performed to obtain $\Phi_{\text{unosc}}^\beta(E_i, \vartheta_i)$, according to the method introduced in Ref. [392]. Here, for each of the 20 bins in $\cos \vartheta$, the 101 integrals of the tabulated flux over the E_ν intervals $[E_{\min}, E_j]$ ($j = 1, 2, \dots, 101$) are interpolated with a cubic spline. Subsequently, the first derivative of each spline with respect to E_ν is evaluated at the energy E_i of the MC event G_i^β . The resulting binwise fluxes are integrated over the $\cos \vartheta$ intervals $[-1, \cos \vartheta_j]$ ($j = 1, 2, \dots, 20$), and interpolated with another cubic spline. The interpolated value of $\Phi_{\text{unosc}}^\beta(E_i, \vartheta_i)$ follows from evaluating the first derivative of this spline function with respect to $\cos \vartheta$ at the value $\cos \vartheta_i$. This approach has been shown to yield an accurate continuous representation of the flux tables [392]. While it is associated with a considerable computational cost, a flux value only needs to be calculated once per MC event. In the reweighting framework, systematic flux uncertainties are implemented as functions that result in externally computed modifications of the shapes or normalisations of the nominal fluxes. They are introduced in Sec. 12.5, alongside other sources of systematic uncertainty.

12.2 Atmospheric μ^\pm background

In principle, an MC simulation method (using the CORSIKA [291] or MuonGun [431] software) can be employed to predict the distribution of the background of atmospheric μ^\pm events. In practice, an event selection procedure that suppresses the contribution of these events as strongly as the one used by this work poses an extreme challenge in terms of generator-level MC statistics. We therefore follow the approach taken by previous analyses using subsets of the event sample here [426, 362, 427]: two ensembles of observed events (“data sidebands”) are employed to estimate the expected three-dimensional shape of the distribution of μ^\pm events at the analysis level, including a shape uncertainty. A set of MC events from MuonGun was nevertheless subjected to the full event selection procedure in order to estimate the overall normalisation of the distribution. It also permitted selecting suitable data sidebands.

The shape of the background distribution is found from the data itself by modifying two event selection criteria: the maximum number of hit DOMs in uninstrumented corridors imposed at level 6 in Table 11.1 is required to be at least two, and the cut on the integrated charge in the spacetime region defined by the complementary causal track veto applied at level 4 is removed. All other criteria are kept in place. Using the mentioned MC events, the μ^\pm purity of the resulting sample has been shown

to amount to $\sim 97\%$, and its overall size to exceed the number of μ^\pm events at the analysis level by a factor of almost three. A second background-dominated sample with an even higher purity of $\sim 99.9\%$ is obtained by requiring a maximum number of hit DOMs in corridors of at least three. This second sample is only 10% greater in size than the sample of μ^\pm events at analysis level. Its deviation from the former, higher-statistics, data-sideband estimate of the background distribution is assumed to correspond to an uncorrelated shape uncertainty. That this uncertainty is indeed large enough to capture possible deviations from the shape of the distribution expected after applying the actual selection criteria has also been verified using the MuonGun simulation.

12.3 Event-sample characterisation

Table 12.1 specifies the predicted composition of the event sample assuming SI and standard three-flavour neutrino oscillations, given NuFIT 3.2 [327, 328] global best-fit values for θ_{12} , θ_{13} , θ_{23} , Δm_{21}^2 , and Δm_{31}^2 under the NO hypothesis, while δ_{CP} is set to zero. The neutrino prediction follows from reweighting neutrino MC events to the Honda *et al.* atmospheric neutrino flux model as discussed in Sec. 12.1, while the μ^\pm prediction is determined as described in Sec. 12.2. In all cases, an effective livetime of 2.5 years underlies the event-count calculation. This $\sim 10\%$ reduction compared to the effective livetime of the selected IceCube data-taking runs reflects the expected fraction of vetoed neutrino events due to (unsimulated) coincidences between atmospheric neutrinos and muons [362]. In contrast, the indicated event rates are normalised to the effective livetime of all data-taking runs. All investigations in this section are based on the above assumptions.

As expected, ν_μ events by far dominate the sample, selected at a rate of nearly 250 μHz , or 0.9 per hour of data taking. In general, the fact that CC events for a given neutrino or antineutrino type strongly outnumber their NC counterparts is due to the larger CC interaction cross sections and higher selection efficiencies. The two CC-to-NC ratios are smaller in the case of $\bar{\nu}_\tau$ events—whose presence is only expected as the result of oscillations—because their CC cross sections are kinematically suppressed and because the outgoing $\bar{\nu}_\tau$ carries away energy. Antineutrino CC interactions of a given flavour are observed at a relative frequency of between 40% to 50% compared to CC interactions of a neutrino of the same flavour. This is the combined effect of smaller intrinsic atmospheric antineutrino fluxes, smaller cross sections, but on average higher inelasticities. This latter antineutrino selection bias does not apply to

Event type	Count	Rate (μHz)
ν_e CC	7234 ± 22	81.87 ± 0.25
$\bar{\nu}_e$ CC	2904 ± 13	32.87 ± 0.15
ν_μ CC	19415 ± 38	219.73 ± 0.43
$\bar{\nu}_\mu$ CC	8950 ± 26	101.29 ± 0.29
ν_τ CC	1260 ± 9	14.26 ± 0.10
$\bar{\nu}_\tau$ CC	516 ± 6	5.84 ± 0.07
ν_e NC	613 ± 7	6.94 ± 0.08
$\bar{\nu}_e$ NC	163 ± 3	1.84 ± 0.04
ν_μ NC	2435 ± 14	27.56 ± 0.16
$\bar{\nu}_\mu$ NC	746 ± 8	8.44 ± 0.09
ν_τ NC	499 ± 6	5.65 ± 0.06
$\bar{\nu}_\tau$ NC	123 ± 3	1.39 ± 0.03
μ^\pm	2451 ± 53	27.74 ± 0.60
Sum	47309 ± 78	535.42 ± 0.88

Table 12.1: Nominal predicted composition of the event sample at analysis level according to neutrino MC, a data sideband, and μ^\pm MC, together with statistical uncertainties. Neutrino events are categorised according to flavour and interaction type. The rightmost column indicates the rate of selected events per event type, assuming an effective exposure time of 2.8 years.

NC interactions. Here, the relative antineutrino selection frequency is only at a level of 25 % to 30 %.

12.3.1 Event template

The observed events and their predicted distributions are binned in three reconstructed parameters: reconstructed neutrino energy E_{reco} , the cosine of the reconstructed zenith angle, $\cos \vartheta_{\text{reco}}$, and the reconstructed classification metric $\Delta L_{\text{reco}} \equiv \ln L_{\text{cascade+track}} - \ln L_{\text{cascade}}$. Eight bins spaced uniformly in $\log_{10}(E_{\text{reco}}/\text{GeV})$ cover the range from $E_{\text{reco,min}} = 10^{0.75} \text{ GeV} \approx 5.62 \text{ GeV}$ to $E_{\text{reco}} = 10^{1.75} \text{ GeV} \approx 56.2 \text{ GeV}$. A single high-energy bin extends this range to $E_{\text{reco,max}} = 100 \text{ GeV}$. While neutrinos originating from above the horizon are expected to yield negligible NSI sensitivity, the inclusion of events with $\cos \vartheta_{\text{reco}} > 0$ provides a simultaneous control measurement of the atmospheric neutrino fluxes. However, the contamination from atmospheric μ^\pm 's is more significant in this region of parameter space. Eight uniformly spaced bins are chosen to cover $\cos \vartheta_{\text{reco}} \in [-1, 1]$. In addition, an event is classified as ‘‘cascade-like’’ if $\Delta L_{\text{reco}} \in [-3, 2)$, and as ‘‘track-like’’ if it has $\Delta L_{\text{reco}} \in [2, 10^3]$.

Except for the addition of one high-energy bin, this choice of binning is identical to that employed by preceding neutrino oscillation measurements based on the more restrictive event selection in Refs. [426, 362]. In principle, a finer binning is desirable as it reduces cancellations and smearing out of the NSI signatures. However, bin-to-bin migration owing to statistical uncertainty on the best-fit values returned by the event reconstruction, as well as the finite amount of neutrino MC statistics and event statistics in the μ^\pm -dominated data sideband are limiting factors.

The neutrino MC prediction for this binning, referred to as the “neutrino template” in the following, is presented in Fig. 12.1. This three-dimensional event distribution is depicted by means of two-dimensional slices in the event-classification dimension, with cascade-like events on the left and track-like events on the right. The overlaid bin labels give the number of unweighted MC events available per bin. The unweighted event count in each bin typically exceeds the reweighted count (colour scale) by at least one order of magnitude. The relative statistical precision of the prediction, $\sigma_{MC,i}/N_{events,i}$ is better than 5% in all bins, and typically $\sim 1\%$ to 2% .

Figure 12.2 shows the analogous information for the atmospheric μ^\pm distribution (“ μ^\pm template”). Here, the universal scaling factor w_{μ^\pm} relates the unweighted and weighted event counts by construction, as motivated in Sec. 12.2. The resulting statistical precision of the prediction for a given bin can therefore be straightforwardly read off from Fig. 12.2, as

$$\Delta_{\mu^\pm,i} = \frac{\sqrt{N_{\text{sideband},i} w_{\mu^\pm}^2}}{N_{\text{sideband},i} w_{\mu^\pm}} = \frac{1}{\sqrt{N_{\text{sideband},i}}}, \quad (12.3)$$

where $N_{\text{sideband},i}$ is the event count observed in the corresponding bin of the data sideband. Δ_{μ^\pm} ranges from $\sim 7\%$ to 58% .

The combined distribution, or “event template”, follows by summing each bin’s weighted neutrino and μ^\pm counts. In the region $\cos \vartheta_{\text{reco}} < 0$, no bin exhibits a relative μ^\pm contamination that exceeds 6% of the total expectation in that bin. In contrast, the event-count expectations in bins with reconstructed directions $\cos \vartheta_{\text{reco}} > 0.25$ and energies $E_{\text{reco}} \gtrsim 25 \text{ GeV}$ typically have relative μ^\pm contributions ranging from $\sim 20\%$ to 40% .

12.3 Event-sample characterisation




Figure 12.1: Predicted distribution of neutrino events (colour map) for standard oscillations using the analysis binning in E_{reco} and $\cos \vartheta_{\text{reco}}$, with cascade-like events on the left and track-like events on the right. In addition, bin labels show the unweighted counts of MC neutrinos.




Figure 12.2: Predicted distribution of atmospheric μ^\pm events (colour map) using the analysis binning, in analogy to Fig. 12.1. In addition, bin labels show the unweighted counts of events observed in the data sideband employed to find the distribution's shape. A universal scaling factor of 2.74, determined by MC simulation, relates unweighted and weighted counts.




Figure 12.3: Predicted distributions of true neutrino energies (left) and zenith angles (right) according to interaction type (CC or NC) for standard oscillations at analysis level. See text for details.

12.3.2 Distributions of true neutrino properties

Figure 12.3 shows the expected distributions of E_ν and $\cos \vartheta$ of neutrinos that pass the selection criteria, depending on neutrino and interaction type. One can see that the average neutrino energy is higher for NC events than for CC events. The true neutrino energy lies outside of the considered E_{reco} range for $\sim 10\%$ of the events. The fraction of neutrinos with $E_\nu > 100 \text{ GeV}$ is largest for $\bar{\nu}_\mu$ NC events, at $\sim 22\%$.

The $\cos \vartheta$ distributions suggest an overall $\sim 10\%$ excess of upgoing neutrinos, $\cos \vartheta < 0$, despite the up-down symmetry of the intrinsic atmospheric neutrino flux in the absence of oscillations, Eq. (5.5), and the appearance $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$ through oscillations. This observation suggests a preference for selecting upgoing neutrinos. In fact, an up-down asymmetry in the distributions of selected events exists for all combinations of neutrino and interaction type except for those involving $\bar{\nu}_\mu$. Here, whether we consider neutrinos or antineutrinos, and CC or NC interactions, the fraction of downgoing events is nearly identical to the fraction of upgoing events. Approximately 6% of the detected neutrinos have trajectories crossing any part of the Earth's core, $\cos \vartheta < -0.84$, for which parametrically enhanced transitions (cf. Sec. 3.3.4.3) could occur. Fewer than 1% traverse both the Earth's outer and inner

core, $\cos \vartheta < -0.98$. Approximately one out of four selected $\bar{\nu}_\tau$'s appears from atmospheric $\bar{\nu}_e$'s or $\bar{\nu}_\mu$'s with $\cos \vartheta < -0.84$.

12.3.3 Reconstruction resolutions

Neutrino energy and zenith resolutions play a crucial role in DeepCore's sensitivity to the oscillation-probability distortions induced by NSI. Finite resolutions imply an averaging of the oscillation probabilities over the detector resolution functions for L and E_ν . This reduces the oscillation amplitudes and potentially results in a complete loss of oscillatory signatures [128]. In the case of a Cherenkov telescope such as DeepCore, the fact that neutrino events cannot in general be distinguished from antineutrino events on an event-by-event basis presents an additional complication. As a consequence of the comparably sparse detector instrumentation with respect to distance scales over which electromagnetic- and hadronic-cascade constituents interact, only a rudimentary discrimination of $\bar{\nu}_\mu$ CC events is possible. The resolution functions enter the event-by-event reweighting scheme naturally and thereby account for all resolution dependencies on the underlying event properties and for all correlations (e.g., between the energy and zenith reconstruction residuals). For illustration purposes, the only dependence investigated below is that on E_ν (undisplayed dimensions on which the resolution functions depend are integrated over). For example, in a given range of E_ν , the distribution of $\cos \vartheta$ will differ from that in another range of E_ν , which in turn impacts the reconstruction resolutions beyond the effect due to the change in E_ν alone. We employ the weighted neutrino event distributions, whose resolutions as a function of E_ν are not independent of the assumed model parameter values.

To begin with, the predicted fraction of events of a given interaction type in the sample to be classified as track-like, r_{trck} , is depicted in Fig. 12.4, for the range $4 \text{ GeV} \leq E_\nu \leq 150 \text{ GeV}$. As the analysis employs exactly two classification bins, the cascade-like fraction is just the complementary probability $r_{\text{cscd}}(E_\nu) \equiv 1 - r_{\text{trck}}(E_\nu)$. The fraction of track-like events among all ν_μ CC and $\bar{\nu}_\mu$ CC events grows more rapidly with E_ν than that among all other event types. An increased selection bias affects events with E_ν outside of the E_{reco} range considered by the analysis, since only the tails of the energy resolution functions are sampled (see below). This implies a limited reconstruction quality for these types of events, which is also expected to affect their classification accuracy.

Figure 12.5 summarises the predicted distributions of the fractional neutrino energy reconstruction residual $\Delta E_\nu/E_\nu \equiv (E_{\text{reco}} - E_\nu)/E_\nu$ and of the cosine-zenith resid-




Figure 12.4: Predicted track-like classification fraction as a function of neutrino energy according to interaction type for standard oscillations at analysis level. See text for details.

ual $\Delta \cos \vartheta \equiv \cos \vartheta_{\text{reco}} - \cos \vartheta$ as a function of E_ν . In each panel, the coloured band extends from the distribution's 15.84th to its 84.16th percentile, indicating the central 1σ (Gaussian-equivalent) interval. Neutrino and antineutrino energy resolutions are nearly indistinguishable by eye and are shown separately only in the case of the muon flavour. In general, the precision (here: width of the central 1σ interval on the residual) with which $\cos \vartheta$ is reconstructed is best for $\bar{\nu}_\mu$ CC events, because of the comparably large spatial extent over which Cherenkov light is deposited by the μ^\pm track on average. $\cos \vartheta$ of the incident $\bar{\nu}_\mu$ is measured with a precision of approximately 0.7 at $E_\nu \approx 5$ GeV, and with a precision of better than 0.2 at $E_\nu \approx 100$ GeV. Weaker improvements with increasing neutrino energy are observed for other event types. A slight but consistent bias of $\Delta \cos \vartheta \approx 0.1$ exists for $\bar{\nu}_\tau$ CC events. The energy reconstruction precision depends on the underlying event type only weakly. The reconstructed energies of the central 1σ population of $\bar{\nu}_\mu$ CC events with $E_\nu \approx 50$ GeV lie within $0.6E_\nu$. Energy resolutions are worse by a few tens of percent for other event types. The bias depends more strongly on event type. As expected, it is particularly skewed toward small energies for NC and $\bar{\nu}_\tau$ CC interactions.




Figure 12.5: Predicted fractional neutrino energy reconstruction residuals (left side of each panel pair) and zenith reconstruction residuals (right side of each pair) as a function of neutrino energy according to interaction type for standard oscillations at analysis level. In each panel, the solid line represents the median, whereas the coloured band corresponds to the central 1σ interval. See text for details.

12.4 NSI signatures

When propagated to the event template, the imprints left by NSI at the oscillation-probability level are diluted by several experimental effects, such as the precision of the event reconstruction, the summation over different event types, and the presence of backgrounds without NSI sensitivity (combined NC and atmospheric μ^\pm events). In order to understand the nature of the residual distortions of the event template induced solely by NSI, here we investigate the binwise statistical compatibility between the SI (or “null”) hypothesis and the assumption of a given NSI coupling strength set to the lower or upper bound of the corresponding 90 % confidence interval allowed by the fit to global neutrino oscillation data discussed in Sec. 5.2.2. The statistical compatibility, or “pull”, is expressed through the quantity

$$n_i \equiv \frac{\mu_{\text{NSI},i} - \mu_{\text{SI},i}}{\sqrt{\mu_{\text{SI},i}}}, \quad (12.4)$$

where $\mu_{\text{SI},i}$ is the expected event count in the i th bin under the SI hypothesis (cf. Sec. 12.3.1), and $\mu_{\text{NSI},i}$ is that under the considered NSI hypothesis. Dividing the difference between these by the expected statistical (Poissonian) uncertainty $\sqrt{\mu_{\text{SI},i}}$ of the SI event count yields a number of standard deviations n_i at which the NSI hypothesis would be excluded on average² if the measurement was subject to statistical uncertainties only. A negative sign of n_i implies that the assumed NSI strength results in a bin-count deficit with respect to SI, whereas a positive sign implies a bin-count excess. The combined significance from all bins is the square root of the individual significances added in quadrature: $n_{\text{tot}} \equiv \sqrt{\sum_i n_i^2}$.

Since the nature of the mentioned dilution effects subtly depends on the specific NSI example under investigation, quantitative and qualitative conclusions drawn from studying the statistical pulls cannot be easily transferred to other NSI hypotheses.

NU NSI

e - μ NU Figure 12.6 shows the expected signature (12.4) arising from the assumption $\epsilon_{ee}^\oplus - \epsilon_{\mu\mu}^\oplus = 0.25$, which can be interpreted as a 25 % rise of the strength of the standard matter potential V_{CC} . In a given bin, the magnitude of the change in the event count is at most ~ 15 % of the statistical uncertainty. Qualitatively, since we are considering the scenario $\epsilon_{ee}^\oplus - \epsilon_{\mu\mu}^\oplus > 0$ and $\Delta m_{31}^2 > 0$, the observed patterns can

²The underlying “Asimov” approach is explained in Chapter 13.




Figure 12.6: Predicted binwise statistical significance of $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} = 0.25$ with respect to SI. Cascade-like events are displayed on the left, track-like events on the right. See text for details.

be understood as the net impact on ν_e CC and ν_μ CC event counts in the presence of backgrounds from all NC events, atmospheric μ^\pm events, as well as from ν_τ CC events and all antineutrino CC events.

Within the cascade-like event class, the strongest signature occurs in a cluster of bins with deficits of similar significance in the upgoing region and below reconstructed energies of ~ 20 GeV. A more complex pattern occurs for track-like events, where there are several sign flips between neighbouring bins. Also here, the overall statistical significance is largest at low energies and for reconstructed zenith angles close to vertically upgoing.

Within the simplified analytic framework developed in Appendix B, the above phenomenology is approximated by the following expression for the event-rate asymmetry in either cascade- or track-like events:

$$\Delta N(\Delta(\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}) > 0) \propto \left(\frac{\Delta P_{ee}}{R_{\mu/e}^0} + \Delta P_{\mu e} \right) r_e + \left(\Delta P_{\mu\mu} + \frac{\Delta P_{\mu e}}{R_{\mu/e}^0} \right) r_\mu. \quad (12.5)$$

On the one hand, for cascade-like events only the first term $\propto \Delta P_{\mu e}$ contributes without a suppression, because $\Delta P_{\mu\mu}$ is suppressed by the small misidentification probability r_μ compared to $r_e \sim 80\%$ (cf. $r_{\text{trck}} = 1 - r_{\text{cscd}}$ in Fig. 12.4). On the other, for track-like events the first two terms are suppressed by the small misidentification probability r_e . In general, unless near the detection threshold, $\Delta P_{\mu\mu}$ is negligible.




Figure 12.7: Same as Fig. 12.6, but for $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} = 0.07$.

Any deficit with respect to the SI scenario is due to the ν_e CC contribution (first two terms), while any excess is due to the ν_μ CC contribution (last two terms). The higher E_{reco} one considers, the higher is the $\bar{\nu}_\mu$ CC purity of track-like events in general ($r_\mu > r_e$). Accordingly, the (weak) high-energy track-like NSI signature is driven by ν_μ CC events.

μ - τ NU The pattern of statistical significances for the μ - τ NU strength $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} = 0.07$ shown in Fig. 12.7 is straightforwardly motivated by the corresponding probability-level discussion in Sec. 9.2. Most prominently, there is an increased $\bar{\nu}_\mu$ CC event count across a broad E_{reco} range in the upgoing region for both cascade- and track-like events, but around four times more significant for the latter. Simultaneously, there is a (less significant) reduction of the $\bar{\nu}_\tau$ CC event count, though shifted to somewhat lower E_{reco} due to the energy reconstruction bias (cf. Fig. 12.5). At sufficiently high E_{reco} , the net signature in both cascade- and track-like events is thus to a first approximation the result of the rate asymmetry

$$\Delta N(\Delta(\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}) > 0) \propto \Delta P_{\mu\mu} r_\mu + \Delta P_{\mu\tau} \frac{r_\tau}{a_{e/\tau}^0} + \frac{1}{R_{\nu_\mu/\bar{\nu}_\mu}^0 a_{\nu_e/\bar{\nu}_e}^0} \left(\Delta \bar{P}_{\mu\mu} \bar{r}_\mu + \Delta \bar{P}_{\mu\tau} \frac{\bar{r}_\tau}{\bar{a}_{e/\tau}^0} \right).$$

Minor reductions of the $\nu_\mu \leftrightarrow \nu_e$ transition probabilities for $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} > 0$ at the matter resonances give rise to a low-significance, $n_i \sim \mathcal{O}(-0.1)$, net decrease of the ν_e CC event count at low energies. In Fig. 12.7, the weak deficits at low energies have




Figure 12.8: Same as Fig. 12.6, but for $\epsilon_{e\mu}^{\oplus} = 0.10$.

their origin in precisely these $\bar{\nu}_{\tau}$ CC and ν_e CC event types.

FV NSI

e - μ FV In contrast to our previous examples, the complex pattern of statistical pulls that results from assuming $\epsilon_{e\mu}^{\oplus} = 0.10$ in Fig. 12.8 can only be understood by considering the interplay between several neutrino and antineutrino CC event types.

A rather uniform increase of the ν_e CC event count across the upgoing region at the level of one standard deviation is encountered for cascade-like events as the result of the enhancement of the $\nu_{\mu} \rightarrow \nu_e$ transition probabilities, $\Delta P_{\mu e(e\mu)} > 0$, as visible in Fig. 9.4. In terms of the contributions to Eq. (B.2), this means that any reduced survival probabilities $\Delta P_{ee} < 0$ are not sufficient to overcome increased appearance probabilities due to the higher atmospheric ν_{μ} flux, or $\Delta P_{\mu e} > -\Delta P_{ee}/R_{\mu/e}^0$. Hence, ν_e CC events are largely responsible for the net $0 < n_i \lesssim 0.5$ signature in cascade-like events, after partial cancellations due to reduced ν_{μ} , ν_{τ} , $\bar{\nu}_{\tau}$, and (for $E_{\text{reco}} \lesssim 25$ GeV) $\bar{\nu}_e$ CC event counts.

The most significant negative pulls in the track-like template arise through ν_{μ} CC events, as the result of $\Delta P_{\mu\mu} < 0$ at the level of several tens of percent (not overcome by $\Delta P_{e\mu} > 0$). Track-like $\bar{\nu}_e$ and $\bar{\nu}_{\mu}$ CC events are responsible for the cluster of positive significances at values of E_{reco} in the range from approximately 25 GeV to 55 GeV in the two most upgoing $\cos \vartheta_{\text{reco}}$ bins.




Figure 12.9: Same as Fig. 12.6, but for $\epsilon_{e\tau}^{\oplus} = -0.26$.

e - τ FV Since $\epsilon_{e\tau}^{\oplus}$ at high energies mainly induces the conversion between $\vec{\nu}_e$ and $\vec{\nu}_\tau$ ($\Delta P_{e\tau}^{(\vec{\nu})} > 0$ and $\Delta P_{ee}^{(\vec{\nu})} < 0$, cf. Fig. 9.4), whose intrinsic atmospheric fluxes are small (or negligible) compared to the $\vec{\nu}_\mu$ fluxes, we only encounter comparably low-significance high-energy signatures in Fig. 12.9, given the large magnitude of the coupling strength $\epsilon_{e\tau}^{\oplus} = -0.26$ assumed.

Still, we find count increases exceeding the statistical uncertainty by more than a factor of two for $15 \text{ GeV} \lesssim E_{\text{reco}} \lesssim 30 \text{ GeV}$ for vertically upgoing track-like events. These large pulls are brought about by a strong attenuation of the first ν_μ disappearance maximum (here, $\Delta P_{\mu\mu} \sim 0.15$ at $\cos \vartheta = -0.75$ in Fig. 9.4), the more significant the more vertical the zenith angle, combined with an enhanced appearance probability $\Delta P_{e\mu} > 0$. This ν_μ CC event signature is also encountered in the cascade-like template on the left, though at a lower overall level.

Even though this e - τ FV hypothesis has a similar impact on ν_μ CC events as the μ - τ NU strength considered above, the pulls of cascade-like events are more significant and their patterns more complex (compare pull patterns in Fig. 12.7). The mechanism behind the high-energy deficit of cascade-like—and, to a reduced extent, track-like—events near the horizontal is analogous to that responsible for the deficit of high-energy track-like events near the horizontal for non-zero $\epsilon_{e\mu}^{\oplus}$ in Fig. 12.8. In the case of e - τ FV, however, it is a deficit of ν_e CC events, which is only partially cancelled by the simultaneous excess of ν_τ CC events, because in these bins (assuming




Figure 12.10: Same as Fig. 12.6, but for $\epsilon_{\mu\tau}^{\oplus} = 0.021$.

$r_e \approx r_\tau$)

$$-\left(\frac{\Delta P_{ee}}{R_{\mu/e}^0} + \Delta P_{\mu e}\right) \gtrsim \left(\frac{\Delta P_{e\tau}}{R_{\mu/e}^0} + \Delta P_{\mu\tau}\right) \frac{1}{a_{e/\tau}^0}, \quad (12.6)$$

with $\Delta P_{ee}/R_{\mu/e}^0 < -\Delta P_{\mu e} < 0$ (ν_e count deficit) and $\Delta P_{e\tau} > -\Delta P_{\mu\tau} > 0$ (ν_τ count excess).

Almost without exception, the binwise pulls of antineutrino CC events have the same signs as their neutrino counterparts above, but are less significant.

μ - τ FV Given $\epsilon_{\mu\tau}^{\oplus} = 0.021$, the origin of nearly all signs of the pulls of track-like events is the ν_μ CC template. When the latter is considered alone, the maximal strength of the NSI signature is about twice as strong than that in Fig. 12.10. For example, the track-like ν_μ CC event counts in the two highest-energy bins with $\cos \vartheta_{\text{reco}} < -0.75$ exhibit pulls of $n_i \sim -4$.

Only few track-like bins represent an exception, in that the sign of the pull is determined by that encountered in $\bar{\nu}_\mu$ CC events. For this purpose, starting from Eq. (B.2), it is sufficient to consider the rate asymmetry

$$\Delta N(\Delta\epsilon_{\mu\tau}^{\oplus} > 0) \propto \Delta P_{\mu\mu} r_\mu + \frac{1}{R_{\nu_\mu/\bar{\nu}_\mu}^0 a_{\nu_e/\bar{\nu}_e}^0} \Delta \bar{P}_{\mu\mu} \bar{r}_\mu, \quad (12.7)$$

where the neutrino term $\propto \Delta P_{\mu\mu}$ carries more weight than its antineutrino counterpart because $\bar{r}_\mu < R_{\nu_\mu/\bar{\nu}_\mu}^0 a_{\nu_e/\bar{\nu}_e}^0 r_\mu$. $\bar{\nu}_\mu$ CC events only dominate if $|\Delta \bar{P}_{\mu\mu}| >$

$|\Delta P_{\mu\mu}| R_{\nu_\mu/\bar{\nu}_\mu}^0 a_{\nu_e/\bar{\nu}_e}^0 r_\mu/\bar{r}_\mu$, with $r_\mu/\bar{r}_\mu < 1$ typically (cf. Fig. 12.4). For the vast majority of track-like bins, the summation over ν_μ CC and $\bar{\nu}_\mu$ CC events merely results in a partial cancellation of the NSI signature of the former.

The net excess of cascade-like events seen in most bins in the upgoing region is driven by a broad enhancement of the ν_τ CC event count, whose excess reaches nearly four times the magnitude of its statistical uncertainty, due to $\Delta P_{\mu\tau} > 0$ (cf. Fig. 9.3):

$$\Delta N(\Delta\epsilon_{\mu\tau}^\oplus > 0) \propto \Delta P_{\mu\tau} \frac{r_\tau}{a_{e/\tau}^0}. \quad (12.8)$$

This excess is partially cancelled by adding cascade-like $\bar{\nu}_\tau$ CC ($\Delta\bar{P}_{\mu\tau} < 0$, antineutrino rate suppression) and ν_μ CC ($\Delta P_{\mu\mu} < 0$, suppressed by the small probability for the cascade-like classification) events, whereas cascade-like $\bar{\nu}_\mu$ CC events also exhibit a moderate enhancement for $E_{\text{reco}} \gtrsim 20$ GeV.

12.5 Systematic uncertainties

Many of the sources of systematic uncertainty considered in this study generalise the reweighting prescription (12.1) for the event weight w_i^α : their effects are encapsulated by (multiplicative) reweighting functions that only depend on the properties of the generator-level neutrino G_i^α . While some apply to all events universally (such as an overall event-rate normalisation), others only affect neutrino events of a certain flavour or of a certain interaction type. Detector response-related uncertainties require a dedicated treatment because they produce a new mapping from the generator-level to the reconstruction-level parameter space. This mapping can only be determined by repeating the whole simulation and reconstruction chain. With several such datasets at hand, these sources of uncertainty are ultimately implemented as rescaling factors of binned event counts, not of individual event weights.

In this NSI search, each source of systematic uncertainty is represented by an associated nuisance parameter with the potential of mitigating the impact of one or more NSI parameters on the event template, either by varying freely or subject to some prior (external) constraint. The discussion below introduces all such uncertainty sources (“systematics”) studied for this analysis and describes how their effects are modelled. A summary is provided in Table 12.2, which also specifies those parameters whose impacts on the NSI sensitivity are found to be negligible (see also later Sec. 14.3). The systematics can be classified according to four broad categories.

Atmospheric neutrino fluxes

One of the most detailed treatments of the uncertainties on the atmospheric neutrino production process to date is provided in Ref. [432], making use of flux predictions from Ref. [282]. This treatment serves as the basis for the four systematics described in the following paragraphs, representing flavour-, energy-, and zenith-dependent deviations from the nominal flux model. The first two systematics make use of empirical parameterisations [433, 434], whose parameter values and flux impacts are detailed in Ref. [393].

In some cases below, flux dependencies on energy and direction are omitted for brevity.

$\nu/\bar{\nu}$ flux ratio The uncertainty $\Delta_{\nu_\alpha/\bar{\nu}_\alpha}$ on the relative production yields of atmospheric neutrinos and antineutrinos is parameterised separately for $\alpha = e, \mu$, as a

function of energy and cosine zenith:

$$\Delta_{\nu_\alpha/\bar{\nu}_\alpha}(E_\nu, \cos \vartheta) = \Delta_{\nu_\alpha/\bar{\nu}_\alpha}^{\text{avg}}(E_\nu) - \Delta_{\nu_\alpha/\bar{\nu}_\alpha}^{\text{shape}}(E_\nu) \left(\frac{k_{\alpha,1}}{\sqrt{2\pi\sigma_\alpha^2}} e^{-\frac{\cos^2 \vartheta}{2\sigma_\alpha^2}} - k_{\alpha,2} \right). \quad (12.9)$$

$\Delta_{\nu_\alpha/\bar{\nu}_\alpha}^{\text{avg}}(E_\nu)$ characterises the uncertainty average over all directions and is implemented as a power function of E_ν ,

$$\log_{10} \left(\Delta_{\nu_\alpha/\bar{\nu}_\alpha}^{\text{avg}}(E_\nu) \right) = \frac{\log_{10}(\Delta_{\alpha,2}/\Delta_{\alpha,1})}{\log_{10}(E_{\alpha,2}/E_{\alpha,1})} \log_{10}(E_\nu/E_{\alpha,1}) + \log_{10}(\Delta_{\alpha,1}) - 2, \quad (12.10)$$

where $\Delta_{\alpha,1}$ and $\Delta_{\alpha,2}$ are the overall uncertainties at some appropriately chosen energies $E_{\alpha,1}$ and $E_{\alpha,2}$. The ‘‘shape’’ term on the right of Eq. (12.9) describes the directional dependence of the uncertainty as a deviation from the average and is the product between the functional form of Eq. (12.10) (with different parameter choices) and an exponential energy suppression $\propto \exp(-E_\nu/E_{\alpha,\text{cutoff}})$. Similarly, $k_{\alpha,1}$, $k_{\alpha,2}$, and σ_α are flavour-dependent fit parameters.

A single systematic parameter $R_{\nu/\bar{\nu}}$ is put in place to simultaneously modify the electron and muon neutrino-to-antineutrino flux ratios according to the parametrisation of Eq. (12.9). The flux ratios are treated as correlated because the same hadronic processes are responsible for their generation during the development of the air shower [432]. For an MC *neutrino* with the unoscillated flux weight $\Phi_{\text{unosc}}^\alpha(E_i, \vartheta_i, \phi_i)$, $R_{\nu/\bar{\nu}}$ implements the transformation

$$\Phi_{\text{unosc}}^\alpha \rightarrow \Phi_{\text{unosc}}^{\alpha'} = \Phi_{\text{unosc}}^\alpha(E_i, \vartheta_i, \phi_i) \left(1 + \frac{R_{\nu/\bar{\nu}}}{2} \Delta_{\nu_\alpha/\bar{\nu}_\alpha}(E_i, \cos \vartheta_i) \right), \quad (12.11)$$

whereas for an MC *antineutrino*, $R_{\nu/\bar{\nu}}$ transforms the input flux weight as

$$\bar{\Phi}_{\text{unosc}}^\alpha \rightarrow \bar{\Phi}_{\text{unosc}}^{\alpha'} = \frac{\bar{\Phi}_{\text{unosc}}^\alpha(E_i, \vartheta_i, \phi_i)}{1 + \frac{R_{\nu/\bar{\nu}}}{2} \Delta_{\nu_\alpha/\bar{\nu}_\alpha}(E_i, \cos \vartheta_i)}. \quad (12.12)$$

Thus, half of the uncertainty is assigned to the neutrino flux, and the other half to the antineutrino flux. The modified ratio is then given as

$$\frac{\Phi_{\text{unosc}}^{\alpha'}}{\bar{\Phi}_{\text{unosc}}^{\alpha'}} = \frac{\Phi_{\text{unosc}}^\alpha}{\bar{\Phi}_{\text{unosc}}^\alpha} \left(1 + \frac{R_{\nu/\bar{\nu}}}{2} \Delta_{\nu_\alpha/\bar{\nu}_\alpha} \right)^2, \quad (12.13)$$

where $R_{\nu/\bar{\nu}}\Delta_{\nu,\alpha/\bar{\nu},\alpha}(E_{\nu}, \cos \vartheta) \ll 1$. The value of $R_{\nu/\bar{\nu}}$ corresponds approximately to the number of standard deviations by which the neutrino-to-antineutrino ratio differs from its nominal value. Setting $R_{\nu/\bar{\nu}} = 0$ recovers the nominal ratio as predicted by the Honda *et al.* model.

Upward-going-to-horizontal $\bar{\nu}$ flux ratio An additional nuisance parameter, $R_{\text{up/hor}}$, is applied only to the angular distribution of the combined electron neutrino and antineutrino flux. Here, the previously introduced parameter $R_{\nu/\bar{\nu}}$ does not cover the full range of uncertainty on the ratio between the flux in the upward-going and that in the horizontal direction.

While the uncertainty $\Delta_{\text{up/hor}}$ is parameterised in an analogous manner to Eqs. (12.9) and (12.10), the nuisance parameter $R_{\text{up/hor}}$ is implemented as in Eq. (12.13), but is now understood to modify the flux ratio $(\Phi_{\text{up,unosc}}^e + \bar{\Phi}_{\text{up,unosc}}^e) / (\Phi_{\text{hor,unosc}}^e + \bar{\Phi}_{\text{hor,unosc}}^e)$.

$\bar{\nu}_e/\bar{\nu}_\mu$ flux ratio The ratio between atmospheric electron and muon neutrino fluxes is known to a precision of about 5%. The corresponding uncertainty is realised through a single rescaling parameter $R_{e/\mu}$ which conserves the total atmospheric flux of neutrinos, $\Phi_{\text{tot,unosc}} \equiv \Phi_{\text{unosc}}^e + \Phi_{\text{unosc}}^\mu$, and of antineutrinos, $\bar{\Phi}_{\text{tot,unosc}} \equiv \bar{\Phi}_{\text{unosc}}^e + \bar{\Phi}_{\text{unosc}}^\mu$. With the nominal flux ratios $\bar{R}_{e/\mu}^0 \equiv \bar{\Phi}_{\text{unosc}}^e / \bar{\Phi}_{\text{unosc}}^\mu$, the transformations are given by

$$\bar{\Phi}_{\text{unosc}}^e \rightarrow \bar{\Phi}_{\text{unosc}}^{e'} = R_{e/\mu} \bar{R}_{e/\mu}^0 \frac{\bar{\Phi}_{\text{tot,unosc}}}{1 + R_{e/\mu} \bar{R}_{e/\mu}^0}, \quad (12.14)$$

$$\bar{\Phi}_{\text{unosc}}^\mu \rightarrow \bar{\Phi}_{\text{unosc}}^{\mu'} = \frac{\bar{\Phi}_{\text{tot,unosc}}}{1 + R_{e/\mu} \bar{R}_{e/\mu}^0}, \quad (12.15)$$

$$\Rightarrow \frac{\bar{\Phi}_{\text{unosc}}^{e'}}{\bar{\Phi}_{\text{unosc}}^{\mu'}} = R_{e/\mu} \bar{R}_{e/\mu}^0 \quad \text{and} \quad \bar{\Phi}'_{\text{tot,unosc}} = \bar{\Phi}_{\text{tot,unosc}}. \quad (12.16)$$

$R_{e/\mu}$ is subjected to a Gaussian constraint with a standard deviation of 5%, centred around $R_{e/\mu} = 1$.

$\bar{\nu}$ flux spectral index The energy dependencies of the atmospheric neutrino and antineutrino fluxes are allowed to vary via a shift $\Delta\gamma$ of the spectral index γ , such that $-\gamma \rightarrow -\gamma + \Delta\gamma$. $\Delta\gamma$ transforms a flux weight $\bar{\Phi}_{\text{unosc}}^\alpha$ as an energy-dependent scaling

factor $s(\Delta\gamma, E_\nu)$,

$$\bar{\Phi}_{\text{unosc}}^\alpha \rightarrow \bar{\Phi}_{\text{unosc}}^{\alpha'} = s(\Delta\gamma, E_i) \bar{\Phi}_{\text{unosc}}^\alpha(E_i, \vartheta_i, \phi_i), \quad (12.17)$$

$$\text{with } s(\Delta\gamma, E_\nu) = \left(\frac{E_\nu}{E_{\text{piv}}}\right)^{\Delta\gamma}, \quad (12.18)$$

where the energy pivot point $E_{\text{piv}} = 24.1$ GeV. A Gaussian constraint with a standard deviation of 0.10 is assumed, centred around $\Delta\gamma = 0$.

Neutrino oscillations

The six independent vacuum-Hamiltonian parameters $\Delta m_{21}^2, \Delta m_{31}^2, \theta_{12}, \theta_{13}, \theta_{23}, \delta_{\text{CP}}$ are considered as unconstrained nuisance parameters.

Neutrino interactions

Neutrino interaction uncertainties are obtained from GENIE (cf. Sec. 10.1). In GENIE, a dimensionless systematic x_p changes a given physics quantity p by the amount $x_p \delta p$, where δp is the quantity's prior standard deviation [397]:

$$p \rightarrow p' = p \left(1 + x_p \frac{\delta p}{p}\right). \quad (12.19)$$

For example, a value of $x_p = -1$ gives rise to a -1σ deviation of p from its nominal value. The fractional uncertainty $\delta p/p$ on each systematic is provided by GENIE itself. Technically, each neutrino MC event is assigned interaction-probability reweighting factors for the discrete variations $x_p = -2, -1, 0, 1, 2$. The reweighting factors are fit with a second-order polynomial in x_p in order to create a continuous nuisance parameter.

CC QES cross section The axial mass for CC quasielastic neutrino scattering (cf. Sec. 2.5.3.1) has a nominal value of $M_A^{\text{CC, QE}} = 0.990$ GeV in GENIE, with 1σ deviations of -15% and 25% .

CC RES cross section The axial mass for neutrino CC resonance production (cf. Sec. 2.5.3.2) has a nominal value of $M_A^{\text{CC, RES}} = 1.120$ GeV in GENIE, with a 1σ range of $\pm 20\%$.

CC/NC DIS cross section Four parameters of the “Bodek–Yang” interaction model [435] for deep inelastic neutrino scattering (cf. Sec. 2.5.3.3) are subject to uncertainty: $A_{\text{HT}}^{\text{BY}} = 0.538 \pm 25\%$, $B_{\text{HT}}^{\text{BY}} = 0.305 \pm 25\%$, $C_{V1}^u = 0.291 \pm 30\%$, and $C_{V2}^u = 0.189 \pm 40\%$.

Event-count normalisation

Overall $\bar{\nu}$ event-count normalisation No constraint is applied to the overall normalisation of the number of neutrino and antineutrino events. We employ a fully correlated scaling factor $N_{\bar{\nu}}$, which acts as an effective parameter representing various physics effects such as interaction cross sections, the effective exposure time, or coincidences between atmospheric $\bar{\nu}$'s and μ^\pm 's.

Overall $\bar{\nu}$ NC event-count normalisation The NC $\bar{\nu}$ event count is allowed to vary independently via a fully correlated scaling factor N_{NC} between ν 's and $\bar{\nu}$'s, which is subject to a Gaussian constraint with a standard deviation of 20% centred around unity.

Overall atmospheric μ^\pm event-count normalisation In addition to statistical and shape uncertainties of the atmospheric μ^\pm template included in the test statistic (see later Sec. 13.2), the overall normalisation of the template serves as an unconstrained nuisance parameter N_{μ^\pm} .

Neutrino detector response

The impact of systematic uncertainties in IceCube DeepCore's response to neutrino interactions is found by means of dedicated MC datasets which have undergone the whole selection and reconstruction chain. Five nuisance parameters are considered, which are associated with three uncertainty categories.

Overall optical DOM efficiency Seven MC simulation runs in total have been performed for different scaling factors s_{eff} applied to all DOMs' nominal overall optical efficiencies, such that $\eta_{0,i}(\lambda) \rightarrow s_{\text{eff}}\eta_{0,i}(\lambda)$ in Eq. (10.1). The simulated values are $s_{\text{eff}} = 0.88, 0.94, 0.97, 1.00, 1.03, 1.06, 1.12$, with the central value of unity corresponding to the nominal model. We impose a Gaussian constraint with a standard deviation of 0.10.

Optical properties of hole ice The two angular-acceptance parameters p_1 (lateral) and p_2 (forward) in Eq. (10.2) are used to account for uncertainty on the optical properties of the refrozen column of ice that used to be the drill hole. Simulation runs employed by this analysis have been performed for the points $(p_1, p_2) = (15, 0), (20, 0), (25, 0), (30, 0), (35, 0), (20, -3), (30, 2), (30, -5), (30, -3), (30, 1), (30, -1)$. Only p_1 is subjected to a prior Gaussian constraint, which is centred around 25 and has a standard deviation of 10.

Optical properties of bulk ice Four different assumptions about the bulk ice absorption and effective scattering coefficients are made by dedicated MC runs. In addition to the nominal ice model, a joint 7% reduction, a 10% increase of the absorption coefficient alone, and a 10% increase of the effective scattering coefficient alone have been simulated. A prior Gaussian constraint of 10% is imposed on the deviation of each coefficient with respect to its value in the nominal ice model.

The bin counts of the event template in dependence of the $n = 5$ detector-response nuisance parameters are parameterised by multiple linear regression. A dedicated regression model is created for $\nu_e + \bar{\nu}_e$ CC, $\nu_\mu + \bar{\nu}_\mu$ CC, and $\nu_\tau + \bar{\nu}_\tau$ CC events each, as well as for the flavour-agnostic sum of all $\bar{\nu}$ NC events. For a given event group and the i th template bin, the ratios between all possible bin counts (given all simulated systematic variations) and the nominal count are subjected to a least-squares fit assuming a hyperplane. Accordingly, the ratio is parameterised as

$$f_i(\theta_1, \dots, \theta_n) = f_{0,i} + \sum_{k=1}^n m_{i,k} \theta_k, \quad (12.20)$$

where θ_k is the nuisance parameter representing the k th systematic uncertainty, $m_{i,k}$ the differential change of the ratio due to θ_k , and $f_{0,i} \equiv f_i(0, \dots, 0)$ is the ratio with respect to the nominal detector model when all $\theta_i = 0$ (not necessarily physical). Finally, the expected $\bar{\nu}$ bin count $\mu_{\bar{\nu},i}^{\bar{\nu}}$ in the presence of n detector systematics $\theta_1, \dots, \theta_n$ follows as

$$\begin{aligned} \mu_{\bar{\nu},i}^{\bar{\nu}}(\theta_1, \dots, \theta_n) &= \left[\mu_{0,i}^{\nu_e \text{CC}} + \mu_{0,i}^{\bar{\nu}_e \text{CC}} \right] f_i^{\nu_e + \bar{\nu}_e \text{CC}}(\theta_1, \dots, \theta_n) \\ &+ \left[\mu_{0,i}^{\nu_\mu \text{CC}} + \mu_{0,i}^{\bar{\nu}_\mu \text{CC}} \right] f_i^{\nu_\mu + \bar{\nu}_\mu \text{CC}}(\theta_1, \dots, \theta_n) \\ &+ \left[\mu_{0,i}^{\nu_\tau \text{CC}} + \mu_{0,i}^{\bar{\nu}_\tau \text{CC}} \right] f_i^{\nu_\tau + \bar{\nu}_\tau \text{CC}}(\theta_1, \dots, \theta_n) \\ &+ \left[\mu_{0,i}^{\nu \text{NC}} + \mu_{0,i}^{\bar{\nu} \text{NC}} \right] f_i^{\nu + \bar{\nu} \text{NC}}(\theta_1, \dots, \theta_n). \end{aligned} \quad (12.21)$$

The MC simulation runs enumerated above have been performed by varying the nuisance parameter(s) of only one of the three uncertainty categories at a time, resulting in a total of 21 systematically perturbed MC datasets. Five additional datasets have been created for which both the DOM efficiency and the hole-ice properties deviate from their nominal assumptions, namely $(s_{\text{eff}}, p_1, p_2) = (0.93, 20, 0), (0.95, 30, -1), (0.98, 35, 1), (1.03, 25, -2), (1.05, 30, 0)$. Despite the underlying joint variations, the corresponding event templates are found to be reasonably well described by the multiple linear regression model without interaction terms in Eq. (12.20) and are therefore included in the least-squares fit. The sum over

$$\underbrace{26}_{\text{no. of MC datasets}} \cdot \underbrace{4}_{\text{no. of event categories}} \cdot \underbrace{9 \cdot 8 \cdot 2}_{\text{no. of analysis bins}} = \underbrace{14\,976}_{\text{no. of residuals}} \quad (12.22)$$

weighted squared residuals between the simulated count ratios and those predicted by the optimised hyperplane model yields $\chi_{\text{min}}^2 \approx 11\,224$. 4 032 parameters are estimated, resulting in an expected number of degrees of freedom of 10 944.³ This implies a goodness of fit (see Sec. 14.4) with a p -value of $\sim 3\%$ for the linear regression model.

³Technically, a sixth parameter for characterising the detector response, a measure of the relative variation of the optical DOM efficiency, is taken into account during the hyperplane fit but not considered as a nuisance parameter.

Nuisance parameter	Nominal value \pm uncertainty	Fit range
<i>Free parameters:</i>		
$R_{e/\mu}$	1.00 ± 0.05	$\pm 3\sigma$
$R_{\nu/\bar{\nu}} (\sigma)$	0.0 ± 1.0	$\pm 3\sigma$
$\Delta\gamma$	0.0 ± 0.1	$\pm 3\sigma$
$N_{\bar{\nu}}$	1.0	[0.0, 1.5]
$x_{M_A^{\text{CC,QE}}} (\sigma)$	0.0 ± 1.0	$\pm 3\sigma$
$x_{M_A^{\text{CC,RES}}} (\sigma)$	0.0 ± 1.0	$\pm 3\sigma$
N_{NC}	1.0 ± 0.2	$\pm 3\sigma$
$\theta_{23} (\circ)$	47.2	[30, 60]
$\Delta m_{31}^2 (10^{-3} \text{ eV}^2)$	2.494	[1.0, 4.0]
$s_{\text{eff}}^{\text{DOM}}$	1.0 ± 0.1	$\pm 2\sigma$
k_{lat}	25 ± 10	$^{+2.5\sigma}$ $^{-2\sigma}$
k_{fwd}	0.0	[-5.0, 2.0]
$s_{\text{scatt}}^{\text{ice}}$	1.0 ± 0.1	$\pm 1\sigma$
$s_{\text{abs}}^{\text{ice}}$	1.0 ± 0.1	$\pm 1\sigma$
N_{μ^\pm}	0.146	[0.0, 1.0]
<i>Fixed parameters (also tested):</i>		
$R_{\text{up/hor}}$	0.0 ± 1.0	$\pm 3\sigma$
$x_{A_{\text{HT}}}^{\text{BY}}$	0.0 ± 1.0	$\pm 3\sigma$
$x_{B_{\text{HT}}}^{\text{BY}}$	0.0 ± 1.0	$\pm 3\sigma$
$x_{C_{\text{V1u}}}^{\text{BY}}$	0.0 ± 1.0	$\pm 3\sigma$
$x_{C_{\text{V2u}}}^{\text{BY}}$	0.0 ± 1.0	$\pm 3\sigma$
$\theta_{12} (\circ)$	33.62	[30, 60]
$\theta_{13} (\circ)$	8.54 ± 0.15	$\pm 3\sigma$
$\Delta m_{21}^2 (10^{-5} \text{ eV}^2)$	7.40	[6.99, 8.02]
$\delta_{\text{CP}} (\circ)$	0	[0, 360]

Table 12.2: Nuisance parameters investigated for inclusion in the NSI analysis, together with their nominal values and (if applicable) prior uncertainties and fit ranges. The upper 15 parameters (“free parameters”) are included in the analysis, while the nine parameters at the bottom (“fixed parameters”) are kept fixed, based on their negligible impacts.

13

Analysis method

The DeepCore event sample introduced in Chapter 11 is interpreted assuming the evolution of three active neutrinos according to Eq. (3.45), using six different interaction Hamiltonians. Each represents a distinct NSI hypothesis, as summarised in Table 13.1.

Hypothesis	Parameters	Sampling grid
e - μ NU	$\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$	$[-5, 5]$
μ - τ NU	$\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$	$[-0.10, 0.10]$
e - μ FV	$ \epsilon_{e\mu}^{\oplus} , \delta_{e\mu}$	$[0, 0.30] \times [0^{\circ}, 360^{\circ}]$
e - τ FV	$ \epsilon_{e\tau}^{\oplus} , \delta_{e\tau}$	$[0, 0.35] \times [0^{\circ}, 360^{\circ}]$
μ - τ FV	$ \epsilon_{\mu\tau}^{\oplus} , \delta_{\mu\tau}$	$[0, 0.07] \times [0^{\circ}, 360^{\circ}]$
GMP	$\epsilon_{\oplus}, \varphi_{12}, \varphi_{13}$	$[-10, 10] \times [-90^{\circ}, 90^{\circ}]^2$

Table 13.1: Overview of the NSI hypotheses studied in this analysis. The middle and right columns specify the parameters of the hypothesis as well as the grid used to sample them. The first two hypotheses allow only for flavour non-universality, and the following three only for flavour violation. The last one is based on the generalised matter potential in Eq. (4.19) and does not place as many restrictions on the NSI flavour structure.

The five phenomenological NSI parameters from the standard parameterisation given by Eq. (4.18) are assumed to be non-zero “one-by-one”: the four remaining parameters are fixed to zero in each case. In the absence of evidence of NSI, this necessarily model-dependent approach is widespread throughout the field (see review in Ref. [262] for example). The most generally applicable constraints result from accounting for the correlations between all couplings. These correlations can lead to cancellations—as in the case of Eq. (9.4) and the study of high-energy $\vec{\nu}_{\mu}$ disappearance—and thereby to weakened constraints compared to those resulting

from assuming one coupling at a time. Nevertheless, the simplified approach is adopted in the first part of this NSI search, not least because there are several theoretical NSI models that accommodate the possibility of the existence of only a single or a small number of sizeable coupling strengths relevant to neutrino propagation.¹ So far, no dedicated analyses of IceCube (DeepCore) event samples [383, 438, 439, 440] have explored e - μ non-universality or new sources of CP violation through complex couplings. Testing the generalised matter potential (4.19) with the three non-zero parameters ϵ_{\oplus} , φ_{12} , and φ_{13} imposes fewer model constraints than the one-by-one fits. Simultaneously, this parameter space was still found to present a computationally tractable problem within the frequentist statistical framework adopted below.

13.1 Frequentist inference

We employ standard frequentist inference methods in order to search for NSI and constrain the associated parameters in Table 13.1. As the analysis was developed in a “blind” manner [441, 442] in order to avoid experimenter bias, extensive MC studies were performed to fully define the analysis procedure before the real data could be revealed and analysed itself. Crucially, this involved the choice of a statistic [443] able to (i) discriminate the tested NSI hypotheses from SI and (ii) constrain the parameters constituting a given NSI hypothesis. In principle, (i) represents a “hypothesis testing” problem, whereas (ii) represents a “parameter determination” problem, where the latter is assumed to imply both parameter estimation and confidence interval estimation [444]. However, the SI hypothesis is nested within all NSI hypotheses, which are composite hypotheses that reduce to the SI hypothesis in different limits of the matter-potential parameters. As a result, the two types of problems (i) and (ii) coincide [445]: if the SI hypothesis is not contained within a confidence region [446] of CL 100% ($1 - \alpha$) in the parameters of interest, the hypothesis is excluded with the corresponding confidence at least.

13.2 A weighted-least-squares statistic

Independent of the hypothesis under consideration, its free matter-potential parameters p_{NSI} and nuisance parameters p_{nuis} are jointly adjusted to best match the ob-

¹Some examples of this kind of NSI can be found in Refs. [436, 437, 272].

served event distribution through the minimisation of [426, 362]

$$\chi_{\text{mod}}^2(\mathbf{p}_{\text{NSI}}, \mathbf{p}_{\text{nuis}}) = \sum_{i=1}^{N_{\text{bins}}} \frac{[n_{\text{obs},i} - n_{\text{exp},i}(\mathbf{p}_{\text{NSI}}, \mathbf{p}_{\text{nuis}})]^2}{n_{\text{exp},i}(\mathbf{p}_{\text{NSI}}, \mathbf{p}_{\text{nuis}}) + \sigma_{\text{exp},i}^2(\mathbf{p}_{\text{NSI}}, \mathbf{p}_{\text{nuis}})} + \sum_{j=1}^{N_{\text{prior}}} \frac{(\Delta p_{\text{nuis},j})^2}{\sigma_{p_{\text{nuis},j}}^2}, \quad (13.1)$$

where $n_{\text{obs},i}$ is the observed number of events in bin i and $n_{\text{exp},i}$ is the combined expectation due to $\bar{\nu}$ atmospheric and μ^\pm background events in the same bin.

The expectation's squared residual with respect to the observed event count in bin i is divided by the sum of two uncertainties added in quadrature: the Poisson variance $n_{\text{exp},i}$ and the variance of the expectation ("MC variance"), $\sigma_{\text{exp},i}^2$. The latter itself is given by the sum

$$\sigma_{\text{exp},i}^2 = \sigma_{\nu,i}^2 + \sigma_{\mu^\pm,i}^2 = \sigma_{\nu,i}^2 + \left[(\sigma_{\mu^\pm,i}^{\text{stat}})^2 + (\sigma_{\mu^\pm,i}^{\text{shape}})^2 \right] \quad (13.2)$$

of the variance $\sigma_{\nu,i}^2$ of the expected number of $\bar{\nu}$ events and the variance $\sigma_{\mu^\pm,i}^2$ of the expected number of atmospheric μ^\pm events in the bin. The variances $\sigma_{\nu,i}^2$ depend on \mathbf{p}_{NSI} and on the subset of nuisance parameters that affect the $\bar{\nu}$ expectation. The variances $\sigma_{\mu^\pm,i}^2$ are taken as constant, set to their nominal values.

The second sum contributing to the statistic (13.1) is taken over all N_{prior} nuisance parameters subject to external Gaussian constraints: a deviation $\Delta p_{\text{nuis},j}$ of the j th such parameter from its nominal value is penalised depending on the parameter's prior standard deviation $\sigma_{p_{\text{nuis},j}}$.

Figure 13.1 shows the nominal "uncertainty template", $\sigma_{\text{total},i} \equiv (n_{\text{exp},i} + \sigma_{\text{exp},i}^2)^{1/2}$, broken down into its contributions. The top (bottom) row corresponds to the uncertainty template of cascade-like (track-like) events. Each panel corresponds to one particular E_{reco} bin along the $\cos \vartheta_{\text{reco}}$ dimension, with E_{reco} increasing from left to right. The Poissonian standard deviation $\sqrt{n_{\text{exp}}}$ is always the largest fraction of the total uncertainty. The MC uncertainty on the prediction itself is typically dominated by the statistical uncertainty of the combined reweighted neutrino and antineutrino MC expectation. Bins for which the μ^\pm background uncertainty dominates instead are almost exclusively encountered in the region $\cos \vartheta_{\text{reco}} > 0$. This is expected from the larger overall μ^\pm count expectations in these bins (cf. Fig. 12.2). The relative MC uncertainty on each bin's predicted total, $\sigma_{\text{exp}}/n_{\text{exp}}$, is below 6.5% (lowest-energy bin) for cascade-like events, and below 12% (second-highest-energy bin) for track-like events. Only two bins have $\sigma_{\text{exp}} > \sqrt{n_{\text{exp}}}$. For these, the relative size of the total uncertainty with respect to the expectation, $\sigma_{\text{total}}/n_{\text{exp}}$, therefore appreciably deviates




Figure 13.1: Nominal binwise event-count uncertainties (σ_{total}) used in the statistic (13.1), broken down into separate contributions that add in quadrature, namely the Poisson standard deviation $\sqrt{n_{\text{exp}}}$ and the MC standard deviation σ_{exp} of the expected event count. Displayed in addition are the neutrino MC (σ_{ν}) and data-sideband ($\sigma_{\mu^{\pm}}$) contributions to the latter in Eq. (13.2). See text for details.

from the simple Poissonian $1/\sqrt{n_{\text{exp}}}$ relation. The maximal fractional total uncertainty in Fig. 13.1 is 19%, its minimum 6%.

13.2.1 Relation to the χ^2 distribution

As a test statistic, χ_{mod}^2 is a stochastic variable of the observed data \mathbf{n}_{obs} whose distribution follows some a-priori unknown PDF for each point in the model-parameter space. χ_{mod}^2 can be considered a variant of the classical Pearson's χ^2 statistic χ_p^2 [447]. The observed event count in a given histogram bin is an independent sample from a Poisson distribution with an unknown mean. The suggestive notation χ_{mod}^2 is motivated by the fact that the minimum of χ_p^2 follows the χ^2 PDF when applied to mutually independent Poisson samples (or multinomial event distributions [447]) with sufficiently large $n_{\text{exp},i}$ [444, 15]. For the nominal event distribution in this analysis, $n_{\text{exp},i} \gtrsim 25 \forall i$, as can be seen in Fig. 13.1. Substituting the increased variance $n_{\text{exp},i} \rightarrow n_{\text{exp},i} + \sigma_{\text{exp},i}^2$ represents the attempt to conserve the asymptotic distribution of the statistic. Its naïvely expected number of degrees of freedom is $N_{\text{bins}} - N_{\text{free}}$

where N_{free} corresponds to the number of estimated parameters [15].

13.3 Construction of confidence regions

As reviewed for example in Ref. [445], in the frequentist framework the coverage probability of a confidence region on the parameter(s) of interest is given by the fraction of boundaries in an ensemble of outcomes of statistically independent experiments that enclose the true parameter value(s). Constructing the confidence region requires a test statistic, specifying the CL of the confidence region (corresponding to the coverage probability), as well as possibly an “ordering rule” which uniquely determines the range of test-statistic values summed or integrated over. All of the above can be equivalently achieved by implementing a hypothesis test for all physical points in the parameters under study, where a given point is included in the confidence region if the hypothesis is not rejected at the desired CL.

In practice, a rigorous frequentist construction of joint confidence regions in all considered parameters (i.e., the parameters of interest and nuisance parameters) quickly becomes intractable [448, 445]. Furthermore, the construction of frequentist confidence regions in the subspace of the parameters of interest (or any other subspace) is subject to conceptual issues [449] and not well studied [445]. As a result, a variety of “pragmatic” approaches are encountered whose coverage probabilities should be evaluated by means of pseudoexperiments when feasible.

13.3.1 Choice of test statistic and Wilks’ theorem

A pragmatic approach is also adopted in this work: confidence regions on NSI parameters are constructed using differences [448, 442] between locally minimised values of χ_{mod}^2 and the global minimum of χ_{mod}^2 . Formally, given some observed data \mathbf{n}_{obs} , the corresponding test statistic in dependence of some fit parameter values \mathbf{p}_{NSI} is given by

$$\Delta\chi_{\text{mod}}^2(\mathbf{p}_{\text{NSI}}) \equiv \chi_{\text{mod}}^2(\mathbf{p}_{\text{NSI}}, \hat{\mathbf{p}}_{\text{nuis}}) - \chi_{\text{mod},\text{min}}^2, \quad (13.3)$$

where $\hat{\mathbf{p}}_{\text{nuis}}$ denote the “conditional” nuisance-parameter values—which minimise χ_{mod}^2 at the point \mathbf{p}_{NSI} —and where $\chi_{\text{mod},\text{min}}^2 \equiv \chi_{\text{mod}}^2(\hat{\mathbf{p}}_{\text{NSI}}, \hat{\mathbf{p}}_{\text{nuis}})$, with $\hat{\mathbf{p}}_{\text{NSI}}$ and $\hat{\mathbf{p}}_{\text{nuis}}$ denoting the “unconditional” matter-potential and nuisance-parameter values at the global minimum. Equation (13.3) is an example of “profiling” (over) nuisance parameters in order to eliminate them from the inference problem (see, e.g., Ref. [445]).

The use of $\Delta\chi_{\text{mod}}^2$ is motivated by a result derived by Wilks in 1938 [450] for the profile likelihood ratio λ in the large-sample limit, or when the estimators of the m parameters of interest (here: the dimensionality of \mathbf{p}_{NSI}) are Gaussian-distributed around their respective true values. Wilks' theorem states that under these circumstances $-2 \ln \lambda \sim \chi_m^2$, where m corresponds to the number of degrees of freedom.² For Wilks' theorem to hold independently of the true parameter values, certain regularity conditions must be met, which are concisely reviewed and illustrated in Ref. [442]. In this work, owing to the large number of parameters of interest probed separately according to Table 13.1, the baseline method of constructing confidence regions is the application of Wilks' theorem to $\Delta\chi_{\text{mod}}^2$ profiles in up to three parameters of interest (henceforth also called " $\Delta\chi_{\text{mod}}^2$ profile scans" due to the discrete grids of points on which $\Delta\chi_{\text{mod}}^2$ is computed). Thus, all confidence levels or significances employ Wilks' theorem unless stated otherwise.

²A detailed investigation of the relationships between Poissonian or Gaussian likelihoods and weighted-least-squares statistics and their performance in counting experiments can be found in Ref. [451] and references therein.

14

Monte Carlo studies

An MC event template that yields the values of all injected parameters when used for parameter estimation is referred to as “Asimov” template, which can always be obtained in the case of a binned analysis [452]. When parameters are fit (estimated) via the minimisation of Eq. (13.1) with $n_{\text{obs}} = n_{\text{exp}}(\mathbf{p}_{\text{inj}})$ under some “truth model”¹ \mathbf{p}_{inj} , one finds $\chi_{\text{mod},\text{min}}^2 = 0$ by construction.²

The Asimov template can be employed to approximate an experiment’s median sensitivity to the parameters of interest without performing pseudoexperiments, by assuming the test-statistic distribution found in the large-sample limit [452]. Here, “sensitivity” refers to the statistical significance of excluding any point \mathbf{p}_{NSI} given \mathbf{p}_{inj} , where the latter will usually correspond to SI. Conversely, the statistical significance of excluding SI when a non-standard matter potential is injected constitutes a “discovery potential”.

Throughout Secs. 14.1, 14.2, and 14.3, the basic properties of χ_{mod}^2 and $\Delta\chi_{\text{mod}}^2$ are studied with the Asimov approach, as is the impact of nuisance parameters. In all of these cases, to indicate that the pseudodata is given by an Asimov template, we identify $\overline{\chi_{\text{mod}}^2} \equiv \chi_{\text{mod}}^2$ and $\overline{\Delta\chi_{\text{mod}}^2} \equiv \Delta\chi_{\text{mod}}^2$. Then, Eq. (13.3) dictates the equality $\overline{\Delta\chi_{\text{mod}}^2}(\mathbf{p}_{\text{NSI}}) = \overline{\chi_{\text{mod}}^2}(\mathbf{p}_{\text{NSI}}, \hat{\mathbf{p}}_{\text{nuis}})$ of the two Asimov metrics. Studies of the actual distributions of both statistics based on pseudoexperiment ensembles are conducted in Sec. 14.4.

For definiteness, all MC studies assume the NO. In particular, whenever the PMNS parameters are not varied (during a fit or for the pseudoexperiment generation) they are again fixed to the NuFIT 3.2 NO global best-fit values (except $\delta_{\text{CP}} = 0$).

¹In the context of these MC studies the terms “injected” and “true” are used interchangeably.

²This statement holds under the condition that the preferred values of any prior parameter constraints match the injected parameter values.




Figure 14.1: Distributions of Asimov fit outcomes in the indicated NSI strengths in dependence of the true parameter value. 100 minimisation trials are performed for each true value. See text for details.

14.1 Characterisation of test-statistic landscapes

Equation (13.3) implies that the frequentist profiling method assumes not only that the global test-statistic minimum in the combined parameter space of \mathbf{p}_{NSI} and \mathbf{p}_{nuis} is reliably found, but also that in \mathbf{p}_{nuis} at a given fixed \mathbf{p}_{NSI} . This section ignores the subspace of the nuisance parameters entirely and focusses only on the parameters of interest.

The two-dimensional histograms in Figs. 14.1 and 14.2 demonstrate that $\overline{\chi^2_{\text{mod}}}$ in general exhibits more than one minimum in each standard NSI parameter as a function of its true value.³ Each histogram is produced by fitting the Asimov template corresponding to a given true value 100 times through the numerical minimisation of $\overline{\chi^2_{\text{mod}}}$, which is initialised with statistically independent starting points (“seeds”) drawn from a uniform distribution across the shown range of values. In this analysis, minimisation always refers to the “sequential least-squares quadratic programming” algorithm SLSQP [453]. Dark blue bins are empty, i.e., no minimisation converged on the associated range of fit values, whereas yellow-coloured bins indicate convergence in a large fraction of trials.

A highly irregular structure characterises the fit values of $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$. Only true values close to zero have a high probability of being accurately recovered by the fit. For the four remaining NSI parameters we never encounter more than one local minimum

³Here we only allow for real coupling strengths for simplicity.




Figure 14.2: Same as Fig. 14.1, but for real FV couplings.

in addition to the global minimum. All such local minima are approximately found by switching the sign of the true parameter value.

The significance levels at which the various local minima in $\overline{\chi_{\text{mod}}^2}$ are excluded vary widely (not shown). The local minima encountered in $\epsilon_{\mu\tau}^{\oplus}$ are disfavoured by $\mathcal{O}(10^2)$, and those in $\epsilon_{e\mu}^{\oplus}$ and $\epsilon_{e\tau}^{\oplus}$ by $\mathcal{O}(10)$ units. Stronger degeneracies occur for the two NU strengths, whose local minima are excluded by between $\mathcal{O}(10^{-4})$ and $\mathcal{O}(1)$ $\overline{\chi_{\text{mod}}^2}$ units. Taking into account that we expect of the order of 10^2 degrees of freedom for χ_{mod}^2 , it is unlikely that the convergence on a local minimum in an NU strength would be revealed by a low goodness of fit. For the same reason, in particular in the case of $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$, such a convergence would have little impact on the accuracy of confidence intervals (e.g., at $> 1\sigma$ CL).

Note that the introduction of nuisance parameters will in general change the sig-

nificances and locations of local minima compared to the above scenario of statistical uncertainties only.

14.2 Statistics-only sensitivity

In the absence of systematic uncertainties, statistical uncertainties govern $\chi_{\text{mod}}^2(\mathbf{p}_{\text{NSI}})$ and limit the precision with which the parameters of interest can be determined. Figures 14.3, 14.4, and 14.5 depict the corresponding statistics-only Asimov sensitivities of the event sample to the NU, FV, and GMP parameters, respectively.

Method

In the case of the NU hypotheses, the median sensitivity is simply obtained by evaluating the one-dimensional $\overline{\Delta\chi_{\text{mod}}^2}$ function for a sequence of hypothesised values of $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ and $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$. Horizontal lines are located at certain critical values $\chi_{1,\alpha}^2$ of a χ^2 distribution with one degree of freedom, with upper-tail probabilities of $\alpha \approx 0.317$ (two-sided tail probability of a central Gaussian 1σ interval), $\alpha = 0.1$ (90% interval), and $\alpha \approx 2.7 \times 10^{-3}$ (3σ interval). In accordance with Sec. 13.3, parameter values with $\overline{\Delta\chi_{\text{mod}}^2} \leq \chi_{1,\alpha}^2$ are included in the 100% $(1 - \alpha)$ confidence interval. Since the pseudodata is the Asimov template, a given bound of a confidence interval should be interpreted as a median: half of the experiments with true SI would yield larger and half would yield smaller bounds.

In the case of the FV hypotheses, the value of the two-dimensional $\overline{\Delta\chi_{\text{mod}}^2}$ function is computed for each point on a predefined grid of hypothesised points in the $|\epsilon_{\alpha\beta}^{\oplus}| - \delta_{\alpha\beta}$ plane. The resulting samples are colour-coded in the three central panels of the figure, overlaid with contours of critical values $\chi_{2,\alpha}^2$ with the same choices of α as before. Points to the left of a given contour are accepted at the 100% $(1 - \alpha)$ CL, whereas those to the right are excluded. The panels on top and on the right show the one-dimensional profiles in $|\epsilon_{\alpha\beta}^{\oplus}|$ respectively $\delta_{\alpha\beta}$. Since $\delta_{\alpha\beta}$ is unphysical when $|\epsilon_{\alpha\beta}^{\oplus}| = 0$, in the Asimov approach the sensitivity to the complex NSI phase always vanishes exactly. By construction there can thus be no closed two-dimensional confidence region and $\overline{\Delta\chi_{\text{mod}}^2}(\delta_{\alpha\beta})$ cannot give the median sensitivity to the phase. The thin red lines in the three central panels of Fig. 14.4 represent the conditional estimators $\hat{\delta}_{\alpha\beta}$ for the sequence of probed $|\epsilon_{\alpha\beta}^{\oplus}|$ values. The discrete nature of the scans is the cause for the observed steps. However, the resulting systematic overestimates of the three profiles in $|\epsilon_{\alpha\beta}^{\oplus}|$ are negligible.




Figure 14.3: Statistics-only sensitivity to NU strengths with the Asimov approach given pseudodata generated assuming SI. The horizontal lines can be used to read off the parameter intervals corresponding to the indicated confidence levels.

When considering the parameters ϵ_{\oplus} , φ_{12} , and φ_{13} of the GMP in Fig. 14.5, the fact that $\Delta\chi_{\text{mod}}^2$ is a three-dimensional function means that only its one- and two-dimensional profiles are shown, with one or two of the matter parameters eliminated at a time via profiling. The rotation angles φ_{12} and φ_{13} are rendered unphysical not for the hypothesis of a standard overall strength of the matter potential, $\epsilon_{\oplus} = 1$, but for a vanishing matter potential, $\epsilon_{\oplus} = 0$. Hence, in the Asimov approach the sensitivity to both is not expected to vanish, but to be bounded from above by the value of $\Delta\chi_{\text{mod}}^2$ with which the hypothesis of vacuum oscillations is disfavoured. Consequently, in principle there may be closed two-dimensional confidence regions in all three planes, as well as finite confidence intervals on φ_{12} and φ_{13} . In each of the three central panels of Fig. 14.5, the points connected by thin lines illustrate the succession of conditional estimators along which the two-dimensional profiles $\Delta\chi_{\text{mod}}^2$ are projected onto their one-dimensional counterparts. There are two non-trivial paths for each two-dimensional profile (in contrast to the three preceding FV hypotheses, where the three vertical lines through the origin are omitted), and each path is associated with the corresponding one-dimensional profile via its colour. As the order in which parameters are profiled is irrelevant, each one-dimensional $\Delta\chi_{\text{mod}}^2$ profile occurs twice.

Both the one- and two-dimensional GMP profiles are subject to profiling over discrete samples of the undisplayed parameters of interest. However, any discreteness effects on the displayed confidence regions and intervals are small because of the high density of the three-dimensional scan, which uses $\sim 1.4 \times 10^5$ grid points.




Figure 14.4: Same as Fig. 14.3, but for FV couplings. The isocontours in the central panels trace the critical values $\chi^2_{2,\alpha}$ for the same choices of α as in Fig. 14.3 and yield the corresponding two-dimensional confidence regions. The one-dimensional profiles on top are the result of evaluating $\overline{\Delta\chi^2_{\text{mod}}}$ along the thin red paths ($\hat{\delta}_{\alpha\beta}$ vs. fit $|\epsilon_{\alpha\beta}^{\oplus}|$).

NU sensitivity

The 1σ confidence interval on $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ in the left panel of Fig. 14.3 consists of two disjoint pieces, separated by a range of values close to $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} = -1$, where the SM matter potential is exactly cancelled. The central values within this range have $\overline{\Delta\chi^2_{\text{mod}}} > \chi^2_{1,0.1}$, implying that they are excluded at 90% CL. $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} = -1$ is not necessarily expected to be the most disfavoured hypothesis, because the neutrino event distributions (of any flavour) disfavour other parameter intervals than their antineutrino counterparts. The prominent local minimum at $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} \approx -2$ reflects that this NSI hypothesis is physically equivalent to switching from the NO to the IO




Figure 14.5: Same as Fig. 14.4, but for GMP parameters. The pair of differently coloured paths in each central panel illustrates how $\Delta\chi^2_{\text{mod}}$ has to be evaluated to find the associated one-dimensional profiles after eliminating the remaining parameter (e.g., upper left, red path: $\hat{\phi}_{12}$ vs. fit ϵ_{\oplus} , yellow path: $\hat{\epsilon}_{\oplus}$ vs. fit ϕ_{12}).

in the SI scenario, giving rise to standard matter-enhanced oscillations of antineutrinos. Clearly, DeepCore has little sensitivity to the NMO even in the absence of NSI. A large fraction of the considered e - μ NU range is expected to be barely excluded at the 1σ CL. Strikingly, the sensitivity flattens for large absolute values. This can be explained through a combination of several probability-level and detector effects, which would hold—at a higher overall level of $\Delta\chi^2_{\text{mod}}$ —even if neutrinos could be distinguished from antineutrinos and if interactions of different neutrino flavours could be told apart. For large positive values of $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ in the case of neutrinos, the matter resonance in transitions involving ν_e shifts below the detection threshold, leading to the suppression (compared to SI) of oscillations at energies just above the

detection threshold. A similar suppression of oscillations occurs for negative values with large moduli. The summation over neutrinos and antineutrinos, as well as over appearance and disappearance channels, results in a further weakening of the NSI signature. Moreover, near the detection threshold the discrimination power between $\bar{\nu}_\mu$ CC events and events of other types is impeded due to the small propagation distance of the μ^\pm emerging at the interaction vertex.

As expected from the probability and template discussions in Chapter 9 respectively Sec. 12.4, the analysis is more sensitive to μ - τ than to e - μ NU. The test statistic keeps growing for the largest values of $|\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus|$ probed in the right panel of Fig. 14.3, resulting in absolute values of ~ 0.1 being excluded at more than 3σ . In order to understand the weak asymmetry of $\Delta\chi_{\text{mod}}^2$ around zero, it is sufficient to consider how the μ - τ NU strength impacts the survival probabilities of ν_μ 's and $\bar{\nu}_\mu$'s depending on its sign. $\Delta m_{31}^2 \cos 2\theta_{23} > 0$ in our nominal scenario, such that the main impact of $\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus < 0$ is the increase of the ν_μ survival probability. $\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus > 0$, on the other hand, mainly increases the $\bar{\nu}_\mu$ survival probability. As a result, the sensitivity to negative values stems mostly from ν_μ CC events, while the sensitivity to positive values stems mostly from $\bar{\nu}_\mu$ CC events.

FV sensitivity

Continuing with the FV couplings in Fig. 14.4, we find qualitatively and quantitatively similar Asimov sensitivities to $\epsilon_{e\mu}^\oplus$ and $\epsilon_{e\tau}^\oplus$, with upper limits at the 90% CL on the couplings' magnitudes of 0.13 and 0.19, respectively. The sensitivity to $|\epsilon_{e\tau}^\oplus|$ has a slightly stronger dependence on the complex phase than the sensitivity to $|\epsilon_{e\mu}^\oplus|$. Both NSI hypotheses exhibit two (shallow) minima of $\Delta\chi_{\text{mod}}^2$ in $\delta_{\alpha\beta}$ for many values of $|\epsilon_{\alpha\beta}^\oplus|$, evidenced by the shapes of the two-dimensional isocontours. The conditional estimators $\hat{\delta}_{e\mu}$ lie between approximately 90° and 135° —implying that the best fits to the Asimov template for non-zero $\epsilon_{e\mu}^\oplus$ are achieved with $\text{Re } \epsilon_{e\mu}^\oplus < 0$ and $\text{Im } \epsilon_{e\mu}^\oplus > 0$. The estimators $\hat{\delta}_{e\tau}$ lie between approximately 270° and 225° , implying best fits with $\text{Re } \epsilon_{e\tau}^\oplus < 0$ and $\text{Im } \epsilon_{e\tau}^\oplus < 0$.

At the 90% CL we expect to be able to exclude $|\epsilon_{\mu\tau}^\oplus| \gtrsim 0.033$ after profiling $\delta_{\mu\tau}$. $\Delta\chi_{\text{mod}}^2$ has two prominent and nearly fully degenerate minima in $\delta_{\mu\tau}$ for a given fixed value of $|\epsilon_{\mu\tau}^\oplus|$. These minima are so deep that while $|\epsilon_{\mu\tau}^\oplus| \gtrsim 0.02$ is incompatible with the Asimov pseudodata at significances exceeding 3σ (assuming χ_2^2) for $\delta_{\mu\tau} = 0^\circ, 180^\circ$, even $|\epsilon_{\mu\tau}^\oplus| \approx 0.03$ is still allowed at 1σ for $\delta_{\mu\tau} \approx 120^\circ, 240^\circ$. The conditional estimators $\hat{\delta}_{\mu\tau}$ grow from approximately 90° to 130° as the assumed value

of the coupling's magnitude increases over the considered range. A priori, there is no reason to expect DeepCore's $\epsilon_{\mu\tau}^\oplus$ sensitivity to be worst for phases of precisely $\delta_{\mu\tau} = 90^\circ$ or $\delta_{\mu\tau} = 270^\circ$ (cf. Sec. 9.2), though this seems to be the case as $|\epsilon_{\mu\tau}^\oplus| \rightarrow 0$.

GMP sensitivity

If φ_{12} and φ_{13} were both fixed to zero, we would have $\overline{\Delta\chi_{\text{mod}}^2}(\epsilon_\oplus) = \Delta\chi_{\text{mod}}^2(\epsilon_{ee}^\oplus - \epsilon_{\mu\mu}^\oplus + 1)$ by construction. The weakly though visibly reduced sensitivity to $\epsilon_\oplus \neq 1$ in Fig. 14.5 compared to that to $\epsilon_{ee}^\oplus - \epsilon_{\mu\mu}^\oplus \neq 0$ in Fig. 14.3 is the result of profiling over φ_{12} and φ_{13} . The cross-like shapes of the two-dimensional confidence regions corresponding to the 90% and 3σ confidence levels on ϵ_\oplus and φ_{12} or φ_{13} are the result of the lacking constraints on the latter two parameters at the corresponding levels of $\overline{\Delta\chi_{\text{mod}}^2}$ when $\epsilon_\oplus \rightarrow 0$ and of the low sensitivity even to large values of $|\epsilon_\oplus|$ (whether the flavour structure of the matter potential is constrained or not). For large $|\epsilon_\oplus|$ the Asimov template is best fit with $\hat{\varphi}_{12}$ and $\hat{\varphi}_{13}$ close to zero, that is, with the standard flavour structure of the matter potential. Conversely, for large $|\varphi_{12}|$ or $|\varphi_{13}|$, $\hat{\epsilon}_\oplus \approx 0$ is preferred. Like this, large oscillation-probability changes at high energy and in channels involving the μ flavour are avoided. For a given hypothesised sign of φ_{12} , the sign of the corresponding conditional estimator $\hat{\epsilon}_\oplus$ is the same, and the same is true when ϵ_\oplus is profiled for the various hypothesised values of φ_{13} . While two-dimensional confidence regions exist up to large confidence levels in the overall strength of the matter potential and either of the rotation angles, the sensitivity to any point in the φ_{12} - φ_{13} plane only reaches the low level of $\overline{\Delta\chi_{\text{mod}}^2} \approx 3$ at which vacuum oscillations are disfavoured—corresponding to a CL of approximately 80% (χ_2^2). The closed 1σ contour present in the lower central panel of Fig. 14.5 extends furthest into the quadrant with $\varphi_{12} < 0$ and $\varphi_{13} < 0$. All points outside of it are excluded at confidence levels between 1σ and 80%. The expected constraints on φ_{13} are weaker than those on φ_{12} (cf. Fig. 5.12).

14.3 Treatment and impact of nuisance parameters

All MC studies so far have been performed by minimising the test statistics only over the parameters of interest. This gives an idea of the sensitivity that would be obtained if all sources of systematic uncertainty were known with perfect accuracy and precision. Minimising over relevant nuisance parameters in addition is of paramount importance for a robust determination of the confidence regions in the parameters

of interest. More generally, systematics with noticeable effects on the extents of confidence regions, or which are expected to bias the measurement or reduce its overall goodness of fit if our nominal assumptions deviate from the (unknown) truth in nature, need to be incorporated into the analysis as nuisance parameters. Conversely, any source of systematic uncertainty that can be removed from the analysis by fixing the associated nuisance parameter to its nominal value in general reduces computational cost and improves numerical stability. In the following, these questions are investigated using the Asimov method, because a prohibitively large number of fits to ensembles of pseudoexperiments would otherwise be required to produce samples of the test-statistic distributions.

14.3.1 Profiling techniques

Continuous profiling of nuisance parameters

Evaluating χ_{mod}^2 on a dense grid ensures that the parameter estimators at the global minimum are found. However, scanning the full parameter space including nuisance parameters constitutes an intractable problem due to the exponential scaling of the computational expense with the number of parameters. For example, assuming even just a small number of $n = 5$ samples in each of $d = 15$ (cf. Table 12.2) dimensions, $n^d \sim \mathcal{O}(10^{10})$ event templates would have to be generated. For a template generation time of ~ 1 s—typical of the analysis at hand—the minimisation would take $\sim 10^3$ years if performed sequentially, and still a single year even if performed on a realistic number of 10^3 processing units in parallel. Consequently, nuisance parameters have to be profiled not by scanning but by numerical optimisation.

SLSQP is initially seeded with the nominal values of all nuisance parameters. Subsequently, it starts from the opposite octant of θ_{23} compared to that found in the initial execution, again seeded with the nominal values of all remaining nuisance parameters. This profiling technique is feasible because the minimisation is a convex problem in every nuisance parameter except θ_{23} . χ_{mod}^2 may have one minimum in the first octant and another in the second—depending on the true value of θ_{23} and the hypothesised values of the other parameters. For Asimov pseudodata generated assuming SI and different examples of fixed NSI fit hypotheses, the dependence of χ_{mod}^2 on θ_{23} is illustrated in Fig. 14.6. Two shallow and nearly fully degenerate minima, one in the first and one in the second octant, are encountered when all NSI parameters are fixed to zero in the fit (SI hypothesis), and when the e - μ NU strength $\epsilon_{ee}^\oplus - \epsilon_{\mu\mu}^\oplus = 1$ is tested. Only a single minimum, in the first (wrong) octant of θ_{23} is




Figure 14.6: $\chi_{\text{mod}}^2(\theta_{23})$ under different fit hypotheses (labelled) for Asimov pseudodata generated assuming true SI. See text for details.

encountered for the two remaining examples, though there is a strong degeneracy between the two octants when $|\epsilon_{e\mu}^{\oplus}| = 0.02$.

In order to verify that the profiling technique above succeeds in identifying the global minimum in the space of nuisance parameters, it was compared to a dedicated, more computationally expensive global optimisation routine also made use of in this analysis.⁴ 140 Asimov templates were generated for SI and with all values of nuisance parameters randomised within their allowed ranges. These templates were fit by fixing one NSI parameter to a random non-zero value at a time and minimising χ_{mod}^2 . The SLSQP minimisation was performed in the default manner as explained above, and the resulting $\chi_{\text{mod,min}}^2$ recorded. For comparison, ten separate SLSQP minimisation runs were performed, each starting from a different seed in the nuisance-parameter space, and the lowest among all ten minima was recorded as $\chi'_{\text{mod,min}}^2$. Here, each seed was chosen as the position of the lowest χ_{mod}^2 value among an initial population of 200 points randomly sampled from the space of nuisance parameters. The relative deviation between $\chi_{\text{mod,min}}^2$ and $\chi'_{\text{mod,min}}^2$ was found to be smaller than 10^{-5} in $\sim 95\%$ of all Asimov pseudoexperiments. Only in a single pseudoexperiment did $\chi_{\text{mod,min}}^2$ exceed $\chi'_{\text{mod,min}}^2$ by more than one unit. Since DeepCore has a weak sensitivity to the octant of θ_{23} (cf. Fig. 14.6), not finding the global minimum in θ_{23} in general only has little impact on the $\Delta\chi_{\text{mod}}^2$ profiles in NSI parameters. Among the ten fits to each Asimov pseudoexperiment seeded with a random population's best estimator, the largest absolute spread between the test-statistic outcomes of any two minimisation runs amounts to approximately six units in χ_{mod}^2 (largest relative spread of around 30%).

⁴Whenever even a scan of just the subspace of NSI parameters is too time-consuming.

Profiling over NSI grid parameters

When the use of a similarly high grid density as in Figs. 14.4 and 14.5 is not viable, profiling over discrete samples of one (FV or GMP) or two (GMP) parameters of interest in order to obtain some lower-dimensional projection of the $\Delta\chi_{\text{mod}}^2$ profile would yield too small confidence regions or intervals, suggesting too high of a measurement precision.⁵

In order to prevent a lower-dimensional projection to suffer from a bias of this kind, for each point in the NSI parameters onto which the high-dimensional $\Delta\chi_{\text{mod}}^2$ profile is to be projected, we search for local minima on the ($d = 1$ - or $d = 2$ -dimensional) grid spanning the space of parameters of interest that have to be optimised. Each detected local minimum is employed as a seed for an additional local SLSQP minimisation process. Any nuisance parameters present—profiled in the aforementioned manner—are simultaneously optimised starting from their conditional best-fit values at this local minimum. The best fit among all the outcomes is recorded and employed in the lower-dimensional projection. Projecting from $d = 2 \rightarrow d = 1$ and from $d = 3 \rightarrow d = 2$ proceeds in an identical manner. Projecting from $d = 3 \rightarrow d = 1$ represents a straightforward generalisation. In the following, this approach is employed wherever a lower-dimensional projection is required, independent of the used grid.

14.3.2 Systematic uncertainty budgets: one nuisance parameter at a time and totals

For the five standard NSI strengths constrainable with the Asimov dataset, $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$, $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$, and $|\epsilon_{e\mu}^{\oplus}|$, $|\epsilon_{e\tau}^{\oplus}|$, $|\epsilon_{\mu\tau}^{\oplus}|$, and the overall scale ϵ_{\oplus} of the GMP, Figs. 14.7 and 14.8 illustrate the individual, “one-at-a-time”, sensitivity impacts of all nuisance parameters. In each figure panel, the statistics-only $\Delta\chi_{\text{mod}}^2$ profile already investigated earlier (e.g., with the complex phase also profiled in the case of the FV hypotheses) is compared to the envelope of all possible profiles from the optimisation of a single nuisance parameter. A small number of such profiles, for the three or four most impactful nuisance parameters in each case, are shown explicitly. The selection criterion is the precision at 90 % CL for all NSI strengths except $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ and ϵ_{\oplus} , where the 30 % respectively 20 % CL is assumed instead.⁶ Comparisons to the full profiles—with the 15 free nuisance parameters from Table 12.2 jointly profiled—allow us to gauge how

⁵Here, the profiled parameters of interest technically act as (temporary) nuisance parameters.

⁶At these levels, the prominent local $\Delta\chi_{\text{mod}}^2$ minimum at the location of the IO hypothesis given SI remains excluded no matter which nuisance parameter is profiled.




Figure 14.7: One-at-a-time impacts of selected nuisance parameters on the sensitivity to NU strengths with the Asimov approach, compared to the statistics-only sensitivity (“stats.”). Also shown is the envelope of all sensitivities in the presence of a single nuisance parameter (“0+1 range”), as well as the nominal sensitivity (“full”). See text for details.

important the individual impacts are with respect to the overall systematic uncertainty budgets.

A summary of the Asimov 90 % confidence intervals from statistical uncertainties only and after the inclusion of systematic uncertainties is provided in Table 14.1. For all parameters of interest with upper and lower bounds at the 90 % CL, the precision (here: width of the confidence interval) is expected to be limited by statistical uncertainties. The relative overall systematic uncertainty budget amounts to 15 % for $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$, 19 % for $|\epsilon_{e\mu}^{\oplus}|$, 14 % for $|\epsilon_{e\tau}^{\oplus}|$, and only 6 % for $|\epsilon_{\mu\tau}^{\oplus}|$. θ_{23} alone is responsible for more than half the systematic uncertainty budget on both $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ and $|\epsilon_{e\tau}^{\oplus}|$. The budget on $|\epsilon_{e\mu}^{\oplus}|$ is dominated by the contribution from Δm_{31}^2 , whereas no single nuisance parameter alone is responsible for the loss in precision on $|\epsilon_{\mu\tau}^{\oplus}|$. Here, the overall normalisation of the NC event template and θ_{23} have the largest impact. Discrepant nuisance-parameter impacts on $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ and ϵ_{\oplus} stem from the simultaneous profiling over φ_{12} and φ_{13} in the GMP case. The two associated profiles are not presented in detail due to the computational cost of ensuring the accurate profiling of ϵ_{\oplus} . Nonetheless, we find that Δm_{31}^2 and θ_{23} are the most impactful nuisance parameters in the outer regions of the φ_{12} profile, $|\varphi_{12}| \gtrsim 10^\circ$, and together with δ_{CP} also in the outer regions of the φ_{13} profile. The associated test-statistic reduction is typically of the order of some 0.5 units.

As a general remark, the one-at-a-time sensitivity impact of any given nuisance parameter is not necessarily independent of the considered CL (or hypothesised NSI strength). For example, δ_{CP} results in a notable reduction of the sensitivity to e - μ NU only at the $\Delta\chi_{\text{mod}}^2$ peak around the vacuum hypothesis, and Δm_{31}^2 loses its dominance




Figure 14.8: Same as Fig. 14.7, but for the magnitudes of the FV couplings (top and lower left) and the overall scale of the GMP (lower right).

over $R_{e/\mu}$ and N_{NC} the higher the CL on $|\epsilon_{e\mu}^{\oplus}|$.

14.3.3 Systematic fit biases

Any known systematics that are not included as nuisance parameters in the analysis should be unable to lead to a significant bias in the measurement of any parameter of interest. To this end, Asimov pseudoexperiments were generated by injecting a non-nominal value for one nuisance parameter at a time. These were then fit with each NSI hypothesis while keeping the nuisance parameter under study fixed to its nominal value, by minimising over the NSI parameters of interest and the other 14 (free) nuisance parameters from Table 12.2. This models the scenario in which the unknown true value of a nuisance parameter deviates from the fit assumption about the same parameter.

Here, only the studies of the nuisance parameters ultimately kept fixed are relevant: the neutrino flux parameter $R_{\text{up/hor}}$, the cross-section parameters $x_{A_{\text{HT}}}^{\text{BY}}$, $x_{B_{\text{HT}}}^{\text{BY}}$, $x_{C_{\text{V1u}}}^{\text{BY}}$, $x_{C_{\text{V2u}}}^{\text{BY}}$, the mixing angles θ_{12} and θ_{13} , the mass-squared difference Δm_{21}^2 , and

Parameter	Statistics	Full	Most impactful
$\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$	$[-\infty, -1.02] \cup [-0.76, +\infty]$	-	$\delta_{\text{CP}}, x_{M_A}^{\text{CC,QE}}, x_{M_A}^{\text{CC,RES}}, \Delta\gamma$
$\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$	$[-0.043, 0.051]$	$[-0.055, 0.055]$	$\theta_{23}, R_{\nu/\bar{\nu}}, k_{\text{fwd}}$
$ \epsilon_{e\mu}^{\oplus} $	$[0.0, 0.13]$	$[0.0, 0.16]$	$\Delta m_{31}^2, R_{e/\mu}, \delta_{\text{CP}}$
$ \epsilon_{e\tau}^{\oplus} $	$[0.0, 0.19]$	$[0.0, 0.22]$	$\theta_{23}, k_{\text{fwd}}, k_{\text{lat}}$
$ \epsilon_{\mu\tau}^{\oplus} $	$[0.0, 0.033]$	$[0.0, 0.035]$	$N_{\text{NC}}, \theta_{23}, s_{\text{abs}}^{\text{ice}}$
ϵ_{\oplus}	-	-	$\Delta\gamma, \Delta m_{31}^2, \theta_{12}, \delta_{\text{CP}}$

Table 14.1: Asimov 90% confidence intervals on the standard NSI strengths and the GMP scale (from statistics only and with the final selection of nuisance parameters) and the most impactful nuisance parameters, determined at a fixed CL as explained in the text.

the CP-violating phase δ_{CP} . For the former five, the injected values are ± 1 (corresponding to $\pm 1\sigma$ variations). For θ_{12} , θ_{13} , and Δm_{21}^2 , the injected values are the bounds of the respective 1σ confidence intervals of NuFIT 3.2 [327, 328]. Lastly, the injected values of δ_{CP} are $45^\circ, 90^\circ, 135^\circ, 180^\circ, 225^\circ, 270^\circ$, and 315° .

As an example, Fig. 14.9 gives the fit outcomes for the μ - τ NU hypothesis, with the best-fit values of $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ compared to the expected 90% confidence intervals from statistical uncertainties only and with the final selection of nuisance parameters from Table 14.1 in the left panel, and the corresponding values of $\chi_{\text{mod,min}}^2$ in the right panel. It is possible to fix all of the enumerated nuisance parameters, because the fit magnitudes of $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ are extremely small compared to the expected measurement precision and the values of $\chi_{\text{mod,min}}^2$ deviate from zero by $\sim 10^{-2}$ units at most. Outcomes of analogous studies for the four remaining standard NSI hypotheses can be found in Figs. C.6 to C.9 in the appendix. The same conclusions hold in all cases and by extension also apply to the GMP analysis.

14.3.4 Sensitivity with additional nuisance parameters

Figure C.10 in the appendix evidences that the previously derived nominal NSI sensitivities are unaffected by the inclusion of the additional systematics investigated for their bias potentials in Sec. 14.3.3. While the $\Delta\chi_{\text{mod}}^2$ profiles for the scenario in which p_{nuis} only encompasses the nominal set of 15 nuisance parameters were computed with the standard profiling techniques from Sec. 14.3.1, a different approach had to be applied to the extended set of 24, namely a global optimisation routine consisting




Figure 14.9: Expected μ - τ NU fit bias from nuisance-parameter offsets. *Left panel:* best-fit values of $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ compared to the expected measurement precision at 90 % CL with and without nuisance parameters (cf. Table 14.1). *Right panel:* corresponding best-fit values of $\overline{\chi^2_{\text{mod}}}$. A triangle pointing upward (downward) represents the outcome of a fit to an Asimov pseudo-experiment with an accordingly shifted injected value of the indicated nuisance parameter. Diamonds represent the fit outcomes for the different injected values of δ_{CP} enumerated in the text.

of several tens of SLSQP minimisation trials on a coarse grid of hypothesised NSI parameter values. At values of $\Delta\chi^2_{\text{mod}}$ corresponding to confidence levels of 1σ , 90 %, and 3σ , barely any degradation of the corresponding confidence intervals is encountered. Thus, the nominal set of nuisance parameters represents a suitable choice for all six fit hypotheses.

14.4 Ensemble studies

For a given truth model \mathbf{p}_{inj} we generate an ensemble of statistically independent pseudoexperiments as

$$n_{\text{obs},i}(\mathbf{p}_{\text{inj}}) \sim \text{Pois} \left\{ n \left[n_{\text{exp},i}(\mathbf{p}_{\text{inj}}), \sigma_{\text{exp},i}^2(\mathbf{p}_{\text{inj}}) \right] \right\}, \quad (14.1)$$

where $n_{\text{obs},i}$ is a given pseudoexperiment's i th bin count, $n(\mu, \sigma^2)$ the Gaussian distribution with mean μ and variance σ^2 , and $\text{Pois}(\lambda)$ the Poissonian distribution with mean λ . A comparably small impact is expected from the Gaussian fluctuations that reflect the finite amount of neutrino MC and data-sideband statistics, owing to the dominance of the Poisson uncertainty demonstrated in Fig. 13.1.

Identifying the global χ^2_{mod} minimum and calculating the $\Delta\chi^2_{\text{mod}}$ profile in the

parameters of interest for each pseudoexperiment allows addressing the following questions:

- What is the overall goodness of fit of the observed measurement outcome $\chi_{\text{mod,min}}^{2,\text{obs}}$?
- Given the truth model \mathbf{p}_{inj} and a hypothesised point \mathbf{p}_{NSI} , which measurement outcomes $\Delta\chi_{\text{mod}}^{2,\text{obs}}(\mathbf{p}_{\text{NSI}})$ are compatible with a statistical fluctuation in the experiment?
- Do confidence regions constructed with Wilks' theorem have the stated coverage or are corrections—which may depend on \mathbf{p}_{NSI} —expected?

The compatibility of the observed minimum $\chi_{\text{mod,min}}^{2,\text{obs}}$ under a given NSI hypothesis with the prediction, the goodness of fit, can be quantified by calculating the probability

$$p_{\text{obs}} = \int_{\chi_{\text{mod,min}}^{2,\text{obs}}}^{\infty} f(\chi_{\text{mod,min}}^2) d\chi_{\text{mod,min}}^2 \quad (14.2)$$

of obtaining an outcome that is equal to or greater than the one observed, $\chi_{\text{mod,min}}^2 \geq \chi_{\text{mod,min}}^{2,\text{obs}}$. An outcome exceedingly far into either tail of the PDF $f(\chi_{\text{mod,min}}^2)$, resulting in an exceedingly large or small value of p_{obs} , could indicate a bad modelling of the data. The PDF could show a weak dependence on the underlying truth model \mathbf{p}_{inj} , which is suppressed in Eq. (14.2).

Independent of the overall goodness of fit, the measurement will yield a profile $\Delta\chi_{\text{mod}}^{2,\text{obs}}(\mathbf{p}_{\text{NSI}})$ constraining the parameters of interest of the hypothesis under study. Certain parameter values might be disfavoured at extreme values of $\Delta\chi_{\text{mod}}^2$ compared to the expectation for true SI, because of statistical fluctuations or because the unknown truth deviates from our nominal MC assumptions.⁷ Knowledge of the distributions of $\Delta\chi_{\text{mod}}^2(\mathbf{p}_{\text{NSI}})$ in dependence of the unknown truth \mathbf{p}_{inj} , here denoted as $f[\Delta\chi_{\text{mod}}^2(\mathbf{p}_{\text{NSI}}|\mathbf{p}_{\text{inj}})]$, permits quantifying the statistical significances of the observed fluctuations throughout parameter space. In the following, the particular ensemble of distributions, for various discrete fit assumptions for \mathbf{p}_{NSI} , found by injecting the nominal SI model is referred to as a “sensitivity distribution” for brevity.

The construction of confidence regions with a coverage of 100% $(1 - \alpha)$ on \mathbf{p}_{NSI}

⁷Modelling inaccuracies could also result in a deviation between the observed profile and the MC expectation.

would require determining the critical values $\Delta\chi_{\text{mod},\alpha}^2(\mathbf{p}_{\text{NSI}})$ such that

$$\alpha = \int_{\Delta\chi_{\text{mod},\alpha}^2(\mathbf{p}_{\text{NSI}})}^{\infty} f[\Delta\chi_{\text{mod}}^2(\mathbf{p}_{\text{NSI}}|\mathbf{p}_{\text{NSI}})] d\Delta\chi_{\text{mod}}^2. \quad (14.3)$$

In the approach popularised by Feldman and Cousins [454], given an observed profile $\Delta\chi_{\text{mod}}^{2,\text{obs}}(\mathbf{p}_{\text{NSI}})$, only those hypothesised points would be included in the sought confidence region that have

$$\Delta\chi_{\text{mod}}^{2,\text{obs}}(\mathbf{p}_{\text{NSI}}) \leq \Delta\chi_{\text{mod},\alpha}^2(\mathbf{p}_{\text{NSI}}). \quad (14.4)$$

Together with the sensitivity distributions, the critical values $\Delta\chi_{\text{mod},\alpha}^2(\mathbf{p}_{\text{NSI}})$ enable one to explicitly obtain ensembles of confidence regions with the correct coverage for any \mathbf{p}_{inj} . Using these, one could for example determine the median upper limit at the desired CL on any parameter of interest for true SI.

With SI (denoted *symbolically* as $\mathbf{p}_{\text{inj}} = 0$) serving as the null hypothesis for all NSI hypotheses, pseudoexperiments $\{\mathbf{n}_{\text{obs},k}(0)\}_{k=1,\dots,n_{\text{pseudo},0}}$ are produced according to Eq. (14.1). $\Delta\chi_{\text{mod}}^2$ profiles are obtained by scanning \mathbf{p}_{NSI} within the parameter space of each analysis hypothesis from Table 13.1 and profiling over the nuisance parameters. Contrary to the Asimov approach from the previous sections, both the global best-fit point and the associated value $\chi_{\text{mod},\text{min}}^2$ at the global minimum are unknown a priori, but scanning and profiling yields an approximation of both. In order to avoid constraining the fit to the predefined grid points in \mathbf{p}_{NSI} , all parameter values at local minima are refined by means of one further SLSQP minimisation per local minimum, seeded with the position of the latter. The lowest χ_{mod}^2 value among the outcomes of the refined fits is taken as $\chi_{\text{mod},\text{min}}^2$.⁸ The ensemble of minima obtained in this manner for each NSI analysis hypothesis is a sample of their expected distribution. For brevity, unless there is ambiguity, hereafter we will denote the distribution of $\chi_{\text{mod},\text{min}}^2$ as $f(\chi_{\text{mod}}^2)$ (even after minimisation over free parameters).

Even if the Asimov approach reasonably approximates the median sensitivity, Wilks' theorem is not guaranteed to give the correct coverage. In general, both problems thus need to be addressed separately using pseudoexperiments. To reduce the associated computational cost, in principle existing pseudoexperiments could be repeatedly reweighted to sample different truth models (see, e.g., Ref. [455]), but an investigation of this technique is beyond the scope of this work.

⁸This refinement routine for global minimisation closely resembles the method of profiling over NSI parameters introduced in Sec. 14.3.1, for the projection of two- or three-dimensional profiles onto some subspace.

14.4.1 Distribution of the weighted-least-squares statistic

The statistics-only distribution of χ_{mod}^2 obtained by generating pseudoexperiments for SI is shown in Fig. 14.10 (labelled “w/ MC fluct.”). It is superimposed with the χ_{144}^2 distribution function, which—because the event templates have 144 bins—is the expected asymptotic distribution of Eq. (13.1) under the truth hypothesis. For comparison, a second distribution of χ_{mod}^2 is shown (labelled “w/o MC fluct.”) which results from replacing the Gaussian random variates in Eq. (14.1) with their means. It resembles a χ_{127}^2 distribution, implying an average reduction per bin of the number of degrees of freedom by $17/144 \approx 0.12$. This stems from the underestimated weight of each term contributing to χ_{mod}^2 . As long as the asymptotic conditions hold, the χ_{mod}^2 distribution without free fit parameters is expected to be independent of the truth model.




Figure 14.10: Statistics-only distributions of χ_{mod}^2 overlaid with different χ_m^2 PDFs.

In the case of free fit parameters (ranging from 15 to 18, ten of which are subject to an external constraint), the latter assumption is only guaranteed to be accurate sufficiently far away from the boundary of the physical parameter space. In any case, in the asymptotic limit one would expect the χ_{mod}^2 distribution to be well approximated by χ_m^2 with $m \sim 130$. Figure 14.11 confirms this expectation: it provides histograms of $\chi_{\text{mod},\text{min}}^2$ for $\mathcal{O}(10^3)$ pseudoexperiments fit with each NSI hypothesis. Also included is the $\chi_{\text{mod},\text{min}}^2$ histogram after minimisation over \mathbf{p}_{nuis} only (SI hypothesis). The result of a χ^2 fit for m to each distribution is overlaid, with m largest for the SI and smallest for the GMP fit hypothesis. The observation that the distributions associated with the two-parameter (FV) hypotheses do not seem to exhibit the naïvely expected reduction in m is reasonable due to the unphysicality of $\delta_{\alpha\beta}$ under SI. In contrast, the SI model is not located at the boundary of the parameter spaces spanned by the NU or the GMP hypotheses. This effect is more explicitly demonstrated during the coverage discussion in Sec. 14.4.3.




Figure 14.11: Expected distributions of χ_{mod}^2 after minimisation over \mathbf{p}_{nuis} and, in all panels but the upper-left one, \mathbf{p}_{NSI} (see labels) for $\mathcal{O}(10^3)$ pseudoexperiments generated according to Eq. (14.1) for true SI. The outcome of a χ_m^2 fit to each distribution is superimposed.

14.4.2 NSI sensitivity

For a fixed truth model \mathbf{p}_{inj} , a given ensemble of $\Delta\chi_{\text{mod}}^2$ profiles—one for each pseudoexperiment and NSI hypothesis—yields a sample of the expected distribution of $\Delta\chi_{\text{mod}}^2$ at each of the \mathbf{p}_{NSI} grid points, $f[\Delta\chi_{\text{mod}}^2(\mathbf{p}_{\text{NSI}}|\mathbf{p}_{\text{inj}})]$. The median sensitivity to \mathbf{p}_{NSI} is defined as the median of the test-statistic distribution under that hypothesis, denoted as $\Delta\chi_{\text{mod},0.5}^2(\mathbf{p}_{\text{NSI}}|0)$ when true SI are considered. Approximations of the distributions for the various hypothesised grid points can be produced with the Asimov method as done in Sec. 14.3.2, where $\overline{\Delta\chi_{\text{mod}}^2}$ was assumed to be a proxy of the median sensitivity. When there are two parameters of interest, instead of investigating the distribution of the test statistic, we determine the probability of excluding each hypothesised point at a value of $\Delta\chi_{\text{mod}}^2$ which corresponds to the 90 %

CL according to Wilks' theorem. For the GMP parameters these exclusion probabilities are not computed, owing to the order-of-magnitude increase in the number of fits required to produce three-dimensional profiles.

NU sensitivity distribution

Figure 14.12 compares the Asimov profile $\overline{\Delta\chi_{\text{mod}}^2}(\epsilon_{ee(\tau\tau)}^\oplus - \epsilon_{\mu\mu}^\oplus | 0)$ to the distribution $f[\Delta\chi_{\text{mod}}^2(\epsilon_{ee(\tau\tau)}^\oplus - \epsilon_{\mu\mu}^\oplus | 0)]$ represented by its median and central 1σ and 90% ranges of statistical uncertainty. By construction, the Asimov test-statistic value must underestimate the median at the injected point, $\overline{\Delta\chi_{\text{mod}}^2}(0|0) = 0 \leq \Delta\chi_{\text{mod},0.5}^2(0|0)$, where $\Delta\chi_{\text{mod},0.5}^2(0|0) = \chi_{1,0.5}^2 \approx 0.45$ if Wilks' theorem holds.

Considering e - μ NU to begin with, the Asimov profile underestimates the median $\Delta\chi_{\text{mod},0.5}^2$ consistently, typically by almost one unit. Just as the Asimov profile, the median, which is around the 1σ exclusion level for most NU values, is relatively flat. For $\epsilon_{ee}^\oplus - \epsilon_{\mu\mu}^\oplus \lesssim -3$ and $\epsilon_{ee}^\oplus - \epsilon_{\mu\mu}^\oplus \gtrsim 2$, pseudoexperiments suggest that there is a probability of around 16% to find $\Delta\chi_{\text{mod}}^2 > 2.71$ —the 90% Wilks exclusion level. The probability to exclude flavour-universal NSI at the same level or higher is $\sim 5\%$ (this probability would be 10% if Wilks' theorem was exact). Vacuum oscillations are excluded at 90% CL in nearly 50% of experiments.

A different situation presents itself in the case of the μ - τ NU hypothesis, in which the Asimov profile is more compatible with the median throughout the studied parameter range. The largest deviations, at the level of $\mathcal{O}(0.1) \Delta\chi_{\text{mod}}^2$ units, occur for small $|\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus| \lesssim 0.05$. For the value $\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus = 0.076$ —close to the global 90% CL upper limit in Table 5.2—there is only a $\sim 5\%$ probability to encounter a rejection at less than 90% CL. The fact that there seems to be an extremely small probability of $\sim 5\%$ to exclude flavour-universal NSI even at moderate levels $\Delta\chi_{\text{mod}}^2 > 1 \approx \chi_{1,0.32}^2$ is revisited in Sec. 14.4.3.

FV sensitivity distribution

The three central panels of Fig. 14.13 contain selected two-dimensional isocontours of the probability of excluding various FV couplings at the level $\chi_{2,0.1}^2$. In analogy to the one-dimensional case, the region of parameter space excluded with a probability between 15.87% (5%) and 84.13% (95%) is defined as the central 1σ (90%) region of statistical uncertainty. The percentiles of $f(\Delta\chi_{\text{mod}}^2)$ after projection onto $|\epsilon_{\alpha\beta}^\oplus|$ respectively $\delta_{\alpha\beta}$ are displayed in the panels at the top and on the right of each plot. In each central panel, the Asimov isocontour $\overline{\Delta\chi_{\text{mod}}^2} = \chi_{2,0.1}^2$ agrees well with




Figure 14.12: Distribution of the sensitivity to NU strengths from 1.5×10^3 pseudoexperiments, expressed as percentiles of $f[\Delta\chi_{\text{mod}}^2(\epsilon_{ee(\tau\tau)}^\oplus - \epsilon_{\mu\mu}^\oplus | 0)]$ as detailed in the text, compared to the Asimov profiles $\overline{\Delta\chi_{\text{mod}}^2}(\epsilon_{ee(\tau\tau)}^\oplus - \epsilon_{\mu\mu}^\oplus | 0)$. The thin dashed lines represent the individual pseudoexperiments' outcomes.

the isocontour of 50% exclusion probability (median) at the same exclusion level.

A distinct feature is the relatively large probability of rejecting even small moduli $|\epsilon_{\mu\tau}^\oplus|$ when $130^\circ \lesssim \delta_{\mu\tau} \lesssim 230^\circ$, exceeding 5% for $|\epsilon_{\mu\tau}^\oplus|$ as small as 5×10^{-3} . The large range of statistical uncertainty in the one-dimensional $\Delta\chi_{\text{mod}}^2$ projection onto small moduli reflects this, as does the “bump” in the centre of the projection onto the phase. In both projections, within the corresponding parameter region the 95th percentile extends beyond the level $\chi_{1,0.1}^2$. The probability of a rejection at 90% CL exceeds 95% for $|\epsilon_{\mu\tau}^\oplus| \gtrsim 0.048$.

For $\epsilon_{e\mu}^\oplus$ and $\epsilon_{e\tau}^\oplus$, we are somewhat less likely to reject a vanishing coupling strength at the level of $\chi_{1,0.1}^2$ or higher. Similarly, the probability to exclude any value of the phase $\delta_{e\mu}$ or $\delta_{e\tau}$ at the same $\Delta\chi_{\text{mod}}^2$ level is well below 5%. In general, the Asimov profile of each FV coupling strength's magnitude is an accurate proxy of the median sensitivity, except near the null hypothesis. At 90% CL, the probability of rejecting $|\epsilon_{e\mu}^\oplus| \gtrsim 0.22$ and $|\epsilon_{e\tau}^\oplus| \gtrsim 0.27$ exceeds 95%. The median sensitivity to each of the phases $\delta_{\alpha\beta}$ is nearly independent of the hypothesised value and of the order of some 0.1 test-statistic units.




Figure 14.13: Distribution of the sensitivity to FV couplings from 7.5×10^2 pseudoexperiments, expressed in the same manner as in Fig. 14.12 for the one-dimensional projections onto $|\epsilon_{\alpha\beta}^{\oplus}|$ and $\delta_{\alpha\beta}$, and as two-dimensional regions bounded by isocontours representing the same set of percentiles of the 90 % CL exclusion probability in each central panel as detailed in the text. The one-dimensional Asimov profiles, the Asimov isocontours at 90 % CL, and the individual pseudoexperiments' one-dimensional projections, two-dimensional isocontours at 90 % CL, and best-fit points in the $|\epsilon_{\alpha\beta}^{\oplus}|$ - $\delta_{\alpha\beta}$ plane are also shown.

14.4.3 Coverage of confidence regions

Due to the number of NSI hypotheses studied and the fact that up to three dimensions of interest are involved, an extensive and high-precision coverage study throughout all parameter spaces is beyond the scope of this thesis. Wilks' theorem assumes that the surface of critical values is the constant $\chi_{m,\alpha}^2$, such that the corresponding isocontour of $\Delta\chi_{\text{mod}}^2$ is identified as the boundary of the 100% $(1 - \alpha)$ confidence region. However, at least one of the regularity conditions underlying Wilks' theorem is violated for true SI when either of the NSI FV hypotheses is considered, as the complex phase $\delta_{\alpha\beta}$ is unphysical at this point. Consequently, one degree of freedom is effectively removed from the problem (cf. Sec. 14.4.1). A significant correlation between a parameter of interest and a nuisance parameter would manifest itself in a similar manner.

The samples of $\Delta\chi_{\text{mod}}^2(0|0)$ are histogrammed in the left column of Fig. 14.14, while the cumulative histograms are shown in the right column. The latter correspond to the actual coverage probabilities as a function of the assumed critical value of $\Delta\chi_{\text{mod}}^2$. For comparison, the PDFs respectively cumulative distribution functions (CDFs) of $\chi_{1,2,3}^2$ are displayed in addition. In all of the rightmost panels, the horizontal line is located at a coverage of 90%, while the vertical lines mark the corresponding critical values of $\chi_{1,2,3}^2$.

The NU and FV hypotheses result in $\Delta\chi_{\text{mod}}^2$ distributions that typically fall more rapidly than χ_1^2 . The coverage is therefore generally larger than with Wilks' theorem. This discrepancy is most striking for μ - τ NU, for which $\Delta\chi_{\text{mod},0.1}^2 \approx 0.5$ —a consequence of its strong correlation with θ_{23} explained in Sec. 9.2. Indeed, we could have already estimated the discrepancy between the critical values for the two NU hypotheses from Fig. 14.12: in the left panel (e - μ NU), $\Delta\chi_{\text{mod},0.05}^2(0|0) \approx \chi_{1,0.1}^2 \approx 2.71$, whereas in the right (μ - τ NU), $\Delta\chi_{\text{mod},0.05}^2(0|0) \approx \chi_{1,0.32}^2 \approx 1$. Thus, if we simply assumed Wilks' theorem to hold, the probability of a rejection of μ - τ flavour universality at 90% CL given true SI would in fact be vanishingly small, prohibiting us from ever claiming an NSI discovery stemming from a statistical fluctuation given SI in nature. In the case of e - μ flavour universality, the probability of such a rejection would still be half as large as claimed, or 5%. Considering the GMP, the distribution of $\Delta\chi_{\text{mod}}^2$ resembles that of χ_m^2 with $2 < m < 3$, smaller than the naïve expectation of $m = 3$.

Coverage for a non-trivial NSI model

One non-trivial NSI scenario is explicitly scrutinised using a statistically independent set of $\sim 10^3$ dedicated pseudoexperiments generated given a non-zero μ - τ FV strength of $\epsilon_{\mu\tau}^{\oplus} = 0.035$, which equals the upper limit of the Asimov 90% confidence interval on the coupling's absolute value in Table 14.1. Figure 14.15 shows the pseudoexperiments' outcomes and can be compared to the corresponding distributions in the middle row of Fig. 14.14 in order to gauge how $\Delta\chi_{\text{mod}}^2$ is affected by the change in the assumed underlying truth model as it is shifted away from the boundary of the physical parameter space. The histograms now follow the χ_2^2 distribution more closely. Therefore, if the measured test-statistic outcome at the injected point ($|\epsilon_{\mu\tau}^{\oplus}| = 0.035$, $\delta_{\mu\tau} = 0$) was $\Delta\chi_{\text{mod}}^2 = \chi_{2,0.1}^2 \approx 4.61$, this particular NSI assumption would indeed be excluded at around 90% CL, in agreement with Wilks' theorem. The indicated actual critical value $\Delta\chi_{\text{mod},0.1}^2 (\epsilon_{\mu\tau}^{\oplus} = 0.035 | \epsilon_{\mu\tau}^{\oplus} = 0.035) \approx 5$ is only slightly larger.




Figure 14.14: Expected probability density and distribution functions of $\Delta\chi^2_{\text{mod}}$ ($p_{\text{NSI}} = 0$) for true SI compared to $\chi^2_{1,2,3}$ for the six considered NSI hypotheses. The left column shows the PDF $f(\Delta\chi^2_{\text{mod}})$, whereas the right column shows the CDF $F(\Delta\chi^2_{\text{mod}})$, which corresponds to the coverage probability as defined in Sec. 13.1. The dashed vertical lines indicate the critical values $\chi^2_{1,0.1}$, $\chi^2_{2,0.1}$, and $\chi^2_{3,0.1}$.




Figure 14.15: Expected PDF (left) and CDF (right) of $\Delta\chi^2_{\text{mod}} (\epsilon_{\mu\tau}^\oplus = 0.035 | \epsilon_{\mu\tau}^\oplus = 0.035)$ fit with the μ - τ FV hypothesis, which can be contrasted with the corresponding histograms in the middle row of Fig. 14.14.

15

Measurement results

15.1 Observed event distribution and fit outcomes




Figure 15.1: Observed distribution (colour maps and bin labels) of cascade-like events on the left and track-like events on the right.

Figure 15.1 presents the observed distribution of events at the final selection level of this analysis. Fitting the six separate NSI hypotheses to this distribution yields the outcomes in Table 15.1, none of which allow making an NSI discovery claim. The null hypothesis is identical across all six scenarios and is comprised of the nested best fit under the SI hypothesis, whose outcome is also specified. All fits are performed within the parameter space of the NO.

Hypothesis	Best-fit values	$\chi_{\text{mod,min}}^{2,\text{obs}}$	$\Delta\chi_{\text{mod,SI}}^{2,\text{obs}}$
SI	$\Delta m_{32}^2 = 0.002\,37\text{ eV}^2, \theta_{23} = 46.4^\circ$	146.4	0
$e\text{-}\mu$ NU	$\epsilon_{ee}^\oplus - \epsilon_{\mu\mu}^\oplus = -0.59$	145.1	1.3
$\mu\text{-}\tau$ NU	$\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus = 0.001\,6$	146.4	0.0
$e\text{-}\mu$ FV	$ \epsilon_{e\mu}^\oplus = 0.072, \delta_{e\mu} = 343.7^\circ$	145.3	1.1
$e\text{-}\tau$ FV	$ \epsilon_{e\tau}^\oplus = 0.060, \delta_{e\tau} = 35.5^\circ$	145.9	0.5
$\mu\text{-}\tau$ FV	$ \epsilon_{\mu\tau}^\oplus = 0.003\,0, \delta_{\mu\tau} = 175.0^\circ$	146.3	0.1
GMP	$\epsilon_\oplus = 0.40, \varphi_{12} = 2.3^\circ, \varphi_{13} = -4.7^\circ$	144.3	2.1

Table 15.1: Summary of the observed fit outcomes for the NSI hypotheses considered in Table 13.1, together with the best-fit values of all NSI parameters, the global minima $\chi_{\text{mod,min}}^{2,\text{obs}}$ and the $\Delta\chi_{\text{mod}}^{2,\text{obs}}$ values at which the SI hypothesis is excluded. In the SI case, owing to the lack of free matter-potential parameters, only the best-fit values of the two PMNS parameters Δm_{32}^2 and θ_{23} are displayed.

SI rejection

The compatibility of the SI hypothesis within each hypothesis' parameter space can be gauged from the indicated test-statistic value

$$\Delta\chi_{\text{mod,SI}}^{2,\text{obs}} \equiv \chi_{\text{mod}}^{2,\text{obs}}(0, \hat{\mathbf{p}}_{\text{nuis}}) - \chi_{\text{mod,min}}^{2,\text{obs}}. \quad (15.1)$$

If Wilks' theorem is assumed to hold and a universal 100% $(1 - \alpha) \approx 75\%$ CL is preimposed, in all cases the best-fit SI hypothesis is accepted: the strongest statistical exclusion of SM interactions is observed for the $e\text{-}\mu$ NU hypothesis, at a Gaussian-equivalent significance of $\approx \sqrt{1.3}\sigma$ (upper tail probability of 25%, $\chi_{1,\alpha=0.25}^2 \approx 1.3$). Even lower confidence levels are required to accept the null hypothesis in the five other hypothesis tests, namely 46% (GMP), 44% ($e\text{-}\mu$ FV), 21% ($e\text{-}\tau$ FV), 4% ($\mu\text{-}\tau$ FV), and 4% ($\mu\text{-}\tau$ NU).

Goodness of fit

Each of the fit outcomes is characterised by a reasonable goodness of fit with a p -value in the range $p_{\text{obs}} \sim 19\%$ to 22% , sufficiently high to pass the prerequisite of $p_{\text{obs}} > 10\%$. The goodness of fit a given fit hypothesis can only be derived from Table 15.1 in conjunction with Fig. 14.11: it is calculated according to Eq. (14.2) using the observed value of $\chi_{\text{mod,min}}^{2,\text{obs}}$ and the appropriate χ_{mod}^2 distribution from pseudoexperiments generated for SI. Any possible dependence of the latter on the underlying




Figure 15.2: Observed distribution of the statistical pulls (colour maps and bin labels) at the μ - τ NU best fit, for cascade-like events on the left and track-like events on the right. The statistical contribution of each event class to the total $\chi_{\text{mod,min}}^{2,\text{obs}}$ is indicated on top. The residual difference is the sum over the prior penalty terms in Eq. (13.1).

truth model—as alluded to in Sec. 14.4.1—is neglected in the goodness-of-fit calculations. This can be justified by the weak observed rejections of the SI hypothesis and the only minor expected changes of $\mathcal{O}(1)$ in the number of degrees of freedom of the various χ_m^2 distributions with $m \sim 130$.

Pull distributions

The pattern of statistical pulls (here: for the best fit under the μ - τ NU hypothesis) in Fig. 15.2 reveals the binwise contributions of the event distribution to the global minimum $\chi_{\text{mod,min}}^{2,\text{obs}}$. The pull associated with a given bin is the square root of the appropriate term in the first sum in Eq. (13.1) and is approximated as the number of standard deviations by which the observed event count differs from the best-fit expectation, with a negative (positive) sign implying a deficit (excess) in the observation. The pull patterns of the other fit outcomes are omitted because of their strong likeness.

The strongest single pull is encountered for cascade-like events with $-0.75 \leq \cos \vartheta_{\text{reco}} < -0.50$ in the highest-energy bin, where the best-fit expectation exceeds the observation by 3.0σ . For track-like events, the strongest pull is also encountered in the highest-energy bin, but in the downgoing region, with a 2.4σ excess in the observed count. The overall contribution to the weighted-least-squares sum from

cascade-like events to the total is lower by 20 units than that from track-like events, however, here $\chi_{\text{mod,cscd.}}^{2,\text{obs}} = 61.85$ and $\chi_{\text{mod,trck.}}^{2,\text{obs}} = 80.91$. Histograms of the observed binwise weighted-least-squares values for the μ - τ NU fit are available in Fig. C.11.

The sum over the ten nuisance-parameter penalty terms, $\chi_{\text{mod,prior}}^{2,\text{obs}} = 3.64$, is almost negligible compared to the aforementioned statistical contributions from cascade- and track-like event counts. This sum alone implies that no significant nuisance-parameter pull (beyond the level of $\sqrt{3.64}\sigma$) could have been observed. The detailed breakdown in Fig. C.12 shows that the only nuisance parameters with pulls of the order of $\pm 1\sigma$ are $x_{M_A^{\text{CC,RES}}}$, $R_{\nu/\bar{\nu}}$, and $s_{\text{eff}}^{\text{DOM}}$. The pulls are nearly constant across fit hypotheses.

Atmospheric oscillation parameters

The best-fit values of Δm_{32}^2 and θ_{23} under the six NSI hypotheses are within 2.5% and 4%, respectively, of their SI counterparts. The latter values of $\Delta m_{32}^2 = 0.00237_{-0.00019}^{+0.00021} \text{ eV}^2$ and $\theta_{23} = 46.4_{-5.8}^{+5.0} \circ$ (Wilks 90% CL) are compatible with the prior dedicated DeepCore measurements of the two parameters in Refs. [426, 362]. The two-dimensional observed $\Delta\chi_{\text{mod}}^{2,\text{obs}}$ profile and its projections can be examined in Fig. C.13.

15.1.1 One-by-one fits

NU NSI

Figure 15.3 shows the observed¹ $\Delta\chi_{\text{mod}}^2$ profiles as a function of the two differences of the flavour-diagonal NSI coupling strengths, $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ and $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$. The shaded bands are the central 68.3% and 90% ranges of statistical uncertainty also displayed in Fig. 14.12. The vertical lines are located at the parameter values that leave flavour transitions unchanged with respect to SI—either because there are no NSI or because these are flavour universal.

e - μ NU The panel on the left of Fig. 15.3 reveals that the data does not allow constraining $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ beyond the level $\Delta\chi_{\text{mod,max}}^2 \approx 7.2$, corresponding to a CL of around 99%. Values outside of the union of intervals $[-2.26, -1.27] \cup [-0.74, 0.32]$ are excluded at 90% CL. Vacuum oscillations, at $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} = -1$, where standard matter effects are exactly cancelled by NSI, are disfavoured by the data nearly at the

¹From this point onwards, when there is no ambiguity we drop the “obs” superscript to avoid further notational clutter.




Figure 15.3: Observed $\Delta\chi_{\text{mod}}^2$ profiles as a function of the effective NSI flavour-non-universality strengths $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ (left) and $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ (right), together with the central 68.3 % and 90 % statistical uncertainty regions of the sensitivity shown as shaded bands.

above maximal test-statistic value of $\Delta\chi_{\text{mod,max}}^2$ with respect to the best fit at $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} \approx -0.59$, somewhat more strongly than expected from the 90 % sensitivity range. Compared to the hypothesis of e - μ flavour universality or SI, vacuum oscillations are disfavoured by $\Delta\chi_{\text{mod}}^2 \approx 5.9$.

The GMOD (4.20) implies that the sign of $1 + \epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ is fully degenerate with the NMO at the level of the (anti)neutrino evolution equation. Only within the NO do we observe a minor preference for a positive sign. This also implies that flavour universality, at which neutrinos and antineutrinos are subject to the standard matter potentials, is favoured over $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} = -2$, at which the signs of the potentials are effectively interchanged. When the IO is assumed in the measurement instead, a degenerate best fit ($\chi_{\text{mod,min}}^2 \approx 145.1$) is encountered at $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} \approx -1.4$, and a slight preference for $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} = -2$ over flavour universality. These observations in turn hint at a small preference for the NO over the IO under the assumption of SI. Such a dedicated NMO measurement in the absence of NSI with two related DeepCore event samples is touched upon in Part IV.

μ - τ NU From the right panel of Fig. 15.3, one is able to tell that the observed event sample is highly compatible with flavour-universal NSI $\epsilon_{\tau\tau}^{\oplus} = \epsilon_{\mu\mu}^{\oplus}$. As a result, stringent constraints emerge, with values of $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ outside of the interval $[-0.041, 0.042]$ excluded at 90 % CL. The measurement precision exceeds the median sensitivity (cf. Fig. 14.12) for any CL, but it is still compatible with the expected central 68.3 % range of statistical uncertainty (even if close to or at its upper bound).

Taking into account the GMOD and the fact that the sensitivity to μ - τ NU originates from (decoupled) 2ν oscillations in the μ - τ sector, it is reasonable to expect

the measured profile under the IO to follow approximately by mirroring that from Fig. 15.3 about the origin, which has been explicitly verified. The best fit is not degenerate with that under the NO because the ee entry of the effective matter potential is not permitted to change sign.

FV NSI

The central panel of the each of the three plots in Fig. 15.4 shows the observed two-dimensional 90% confidence region (isocontour at $\Delta\chi_{\text{mod}}^2 = \chi_{2,0.1}^2 \approx 4.61$) in the NSI magnitude and complex phase from the fit of a given FV NSI coupling strength. Points to the right of the observed isocontour are excluded. Each isocontour is accompanied by the central 68.3% and 90% regions of statistical uncertainty displayed in the corresponding part of Fig. 14.13. The profile's projection onto the modulus of the coupling strength is depicted on top, and that onto the phase on the right. Lines and shaded bands have the same interpretation as in Fig. 15.3. Due to the fact that a vanishing coupling renders its phase unphysical, the appropriate entry for $\Delta\chi_{\text{mod,SI}}^2$ in Table 15.1 provides the maximal projected test-statistic value at which any given value of the phase can be excluded in principle.

Generically, only without restrictions on the NSI flavour structure, δ_{CP} , θ_{12} , and the NMO should $\epsilon_{\alpha\beta}^{\oplus}$ and $-\epsilon_{\alpha\beta}^{\oplus*}$ (or $\delta_{\alpha\beta}$ and $180^\circ - \delta_{\alpha\beta}$) yield physically equivalent solutions, cf. Eq. (4.20). Given that e - μ NU as a prerequisite for this exact degeneracy is not accommodated by our FV NSI hypotheses, we have explicitly verified that all three $|\epsilon_{\alpha\beta}^{\oplus}|$ measurements are only negligibly affected by switching from the NO to the IO hypothesis, and that optimising over both NMO realisations does not modify the confidence intervals compared to those specified below for the NO hypothesis.

e - μ FV From the top left part of Fig. 15.4, an upper bound of $|\epsilon_{e\mu}^{\oplus}| \leq 0.146$ at 90% CL is obtained on e - μ FV NSI when the full range of phases, $0^\circ \leq \delta_{e\mu} \leq 360^\circ$, is tested. The constraint becomes stronger when $\delta_{e\mu}$ is allowed to vary only over the range $(160 \pm 90)^\circ$, for which the conditional value of the magnitude that best fits the data is zero, $|\hat{\epsilon}_{e\mu}^{\oplus}| = 0$. This explains the plateau in the projection onto $\delta_{e\mu}$ with $\Delta\chi_{\text{mod}}^2 = \Delta\chi_{\text{mod,SI}}^2 \approx 1.1$. A somewhat stronger exclusion of real negative values of $\epsilon_{e\mu}^{\oplus}$ (corresponding to $\delta_{e\mu} = 180^\circ$) with respect to the expectation from pseudoexperiments is observed.

e - τ FV Compared to e - μ FV NSI, we find both qualitatively and quantitatively similar bounds on e - τ FV NSI from the upper right panels of Fig. 15.4. The mag-

nitude's upper bound at 90 % CL after optimising over $0^\circ \leq \delta_{e\tau} \leq 360^\circ$ is almost 20 % larger, $|\epsilon_{e\tau}^\oplus| \leq 0.173$, and the statistical compatibility between the best fit and the SI hypothesis is even higher. For values of $\delta_{e\tau}$ contained within the range $(200 \pm 90)^\circ$, the magnitude is somewhat more strongly bounded from above than expected from pseudoexperiments. As before, the conditional NSI hypotheses in this limited $\delta_{e\tau}$ interval are characterised by vanishing conditional absolute values $|\epsilon_{e\tau}^{\hat{\oplus}}| = 0$, giving rise to a plateau in the projection onto $\delta_{e\tau}$ with $\Delta\chi_{\text{mod}}^2 = \Delta\chi_{\text{mod,SI}}^2 \approx 0.5$.

μ - τ FV The lower part of Fig. 15.4 confirms that the selected event sample is able to more strongly constrain flavour violation in the μ - τ sector than in the electron sector. Under the full $\delta_{\mu\tau}$ range, $|\epsilon_{\mu\tau}^\oplus| \leq 0.0232$ at 90 % CL. We observe the strongest bounds for real and positive NSI couplings ($\delta_{\mu\tau} = 0^\circ$). Hypotheses with $\delta_{\mu\tau} \approx 125^\circ$ and $\delta_{\mu\tau} \approx 235^\circ$ result in the weakest upper limits on the magnitude. In turn, the measurement is almost unable to place any constraints on $\delta_{\mu\tau}$, as $\Delta\chi_{\text{mod,SI}}^2 \approx 0.1$ is the maximum of the parameter's test-statistic projection.

While the isocontour is well contained within the 90 % region of statistical uncertainty, the profile's projection onto $|\epsilon_{\mu\tau}^\oplus|$ is only barely compatible with the 90 % sensitivity range. Just from the shape of the latter, one can tell that the smallness of the observed best-fit value $|\epsilon_{\mu\tau}^\oplus| = 0.0030$ is rather untypical. This simultaneously brings about a stronger-than-observed exclusion of the SI hypothesis by the median pseudoexperiment and also a larger-than-observed upper limit.




Figure 15.4: Observed 90 % confidence regions in the moduli $|\epsilon_{\alpha\beta}^{\oplus}|$ and phases $\delta_{\alpha\beta}$ of the effective flavour-violating NSI coupling strengths $\epsilon_{e\mu}^{\oplus}$ (upper left), $\epsilon_{e\tau}^{\oplus}$ (upper right), and $\epsilon_{\mu\tau}^{\oplus}$ (bottom), together with each parameter's projected one-dimensional $\Delta\chi_{\text{mod}}^2$ profile. The best-fit point for each pair of parameters is indicated by a cross. The central 68.3 % and 90 % statistical uncertainty regions and intervals of the sensitivity are shown as shaded bands.

15.1.2 Generalised matter potential




Figure 15.5: Observed 1σ , 90% , and 3σ confidence regions in pairs of the GMP parameters ϵ_{\oplus} , φ_{12} , and φ_{13} , together with each parameter's projected one-dimensional $\Delta\chi_{\text{mod}}^2$ profile. In addition, the colour in each of the three central panels encodes the local value of the projected two-dimensional $\Delta\chi_{\text{mod}}^2$ profile. A cross indicates the best-fit point for each pair of parameters.

Our final fit to data employs the GMP that is characterised by the three intrinsic matter parameters ϵ_{\oplus} , φ_{12} , and φ_{13} . Figure 15.5 shows the resulting constraints, by means of the two- and one-dimensional $\Delta\chi_{\text{mod}}^2$ profiles, all of which are projections of the complete three-dimensional profile. In terms of the five standard NSI parameters, the global best fit, also provided in Table 15.1, corresponds to

$$\begin{aligned} \epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} &= -0.60, & \epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} &= 0.0020, \\ \epsilon_{e\mu}^{\oplus} &= -0.016, & \epsilon_{e\tau}^{\oplus} &= 0.033, & \epsilon_{\mu\tau}^{\oplus} &= -0.0013. \end{aligned} \quad (15.2)$$

This point is weakly favoured over the hypothesis of SI (or flavour-universal NSI), at the level of $\Delta\chi_{\text{mod,SI}}^2 = 2.2$ (upper tail probability of 55.2%, $\chi_{3,0.552}^2 \approx 2.2$)—a difference that cannot be derived from any of the projections in Fig. 15.5 because none of these contain the two points also contained within the SI parameter space ($\epsilon_{\oplus} = \pm 1, \varphi_{12} = 0^\circ, \varphi_{13} = 0^\circ$). The one-dimensional projections yield the 90% confidence intervals $-9^\circ \lesssim \varphi_{12} \lesssim 8^\circ$, $-14^\circ \lesssim \varphi_{13} \lesssim 9^\circ$, and the union $[-1.2, -0.3] \cup [0.2, 1.4]$ for ϵ_{\oplus} .

Compared to keeping φ_{12} and φ_{13} fixed at zero, their joint optimisation almost leaves the bounds on ϵ_{\oplus} at 90% CL unaffected, and it does not change the overall shape of its one-dimensional $\Delta\chi_{\text{mod}}^2$ profile, which follows from contrasting it with the profile for $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} = \epsilon_{\oplus} - 1$ in Fig. 15.3.

The two-dimensional $\Delta\chi_{\text{mod}}^2$ projection onto ϵ_{\oplus} and φ_{12} demonstrates that all hypotheses with $|\epsilon_{\oplus}| > 0.05$ and $|\varphi_{12}| > 10^\circ$ are excluded with a significance above 3σ (for any value of φ_{13}). Conversely, hypotheses with $|\epsilon_{\oplus}| > 0.1$ and $|\varphi_{13}| > 20^\circ$ are excluded. Hypotheses with any value of ϵ_{\oplus} and smaller values of $|\varphi_{12}|$ and $|\varphi_{13}|$ are allowed at the 3σ level.

In the projection onto φ_{12} and φ_{13} an isocontour corresponding to the 90% CL is observed, but not the 99.7% (3σ) CL. The maximal significance of excluding any particular pair of values of the rotation angles cannot exceed the $\Delta\chi_{\text{mod}}^2$ value of vacuum oscillations ($\epsilon_{\oplus} = 0$), which renders both parameters unphysical. In fact, no value of ϵ_{\oplus} is excluded at a higher CL than $\epsilon_{\oplus} = 0$, by $\Delta\chi_{\text{mod}}^2 \approx 8.0 < \chi_{1,0.003}^2 < \chi_{2,0.003}^2 \approx 11.62$. As a consequence, we also encounter the cross-like shapes of the 3σ isocontours in the two upper central panels in Fig. 15.5.

Hypothesis	Parameter	90 % confidence interval
e - μ NU	$\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$	$[-2.26, -1.27] \cup [-0.74, 0.32]$
μ - τ NU	$\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$	$[-0.041, 0.042]$
e - μ FV	$ \epsilon_{e\mu}^{\oplus} $	≤ 0.146
	$\delta_{e\mu}$	$[0^{\circ}, 360^{\circ}]$
e - τ FV	$ \epsilon_{e\tau}^{\oplus} $	≤ 0.173
	$\delta_{e\tau}$	$[0^{\circ}, 360^{\circ}]$
μ - τ FV	$ \epsilon_{\mu\tau}^{\oplus} $	≤ 0.0232
	$\delta_{\mu\tau}$	$[0^{\circ}, 360^{\circ}]$
GMP	ϵ_{\oplus}	$[-1.2, -0.3] \cup [0.2, 1.4]$
	φ_{12}	$[-9^{\circ}, 8^{\circ}]$
	φ_{13}	$[-14^{\circ}, 9^{\circ}]$

Table 15.2: Summary of the observed 90 % confidence intervals according to Wilks' theorem on NU and FV NSI parameters within the one-by-one fits and on the GMP parameters. All measurements make the assumption $\Delta m_{31}^2 > 0$, which does not necessarily result in a loss of generality (see text).

15.2 Summary and experiment comparison

Table 15.2 compiles a summary of the intervals at 90 % CL placed on the parameters of interest to this work. Only the constraints on the strength of e - μ NU NSI are at the level of the strength of SM weak interactions, all others are stronger by almost one (e - μ and e - τ FV) or two (μ - τ NU or FV) orders of magnitude. None of the FV phases are restricted.

For comparison with existing measurements, in Fig. 15.6 we restrict the FV coupling strengths to the real plane, defined by $\delta_{\alpha\beta} = 0^{\circ}, 180^{\circ}$, and show the allowed intervals at 90 % CL for the real-valued signed quantities $\epsilon_{\alpha\beta}^{\oplus}$. In the cases of $\epsilon_{e\mu}^{\oplus}$ and $\epsilon_{e\tau}^{\oplus}$, the lower limits ($\delta_{\alpha\beta} = 180^{\circ}$) are stronger than the upper limits ($\delta_{\alpha\beta} = 0^{\circ}$). The latter reproduce the observed constraints on $|\epsilon_{e\mu}^{\oplus}|$ and $|\epsilon_{e\tau}^{\oplus}|$ under the hypotheses of complex coupling strengths in Table 15.2. In the case of $\epsilon_{\mu\tau}^{\oplus}$, the upper limit is $\sim 20\%$ stronger than the lower limit, $-0.0165 \leq \epsilon_{\mu\tau}^{\oplus} \leq 0.0130$. This range improves upon the 90 % confidence interval reported by the preceding DeepCore measurement, $-0.020 \lesssim \epsilon_{\mu\tau}^{\oplus} \lesssim 0.024$ [439].² Neither limit's modulus reproduces the limit on

²After translating from the assumed NSI with down quarks, $\epsilon_{\mu\tau}^{dV}$, to the effective coupling strength in Eq. (4.12) employed in this work.




Figure 15.6: Summary of the one-by-one 90 % confidence intervals on real NSI parameters obtained in this work compared to previous limits. Constraints on the complex NSI coupling strengths are given for the associated phase restricted to $\delta_{\alpha\beta} = 0^\circ, 180^\circ$. See text for details.

$|\epsilon_{\mu\tau}^\oplus|$ in Table 15.2 because the precision on $|\epsilon_{\mu\tau}^\oplus|$ is weakest for a complex coupling strength.

Data from a number of other neutrino experiments has been used to set limits on the NSI quark coupling strengths $\epsilon_{\alpha\beta}^{uV}$ and $\epsilon_{\alpha\beta}^{dV}$, which are also presented in Figure 15.6 after rescaling them to the convention (4.12) adopted by this work.

Limits on $\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus$ and $\epsilon_{\mu\tau}^\oplus$ (two-dimensional analysis including correlations) have been obtained from atmospheric neutrino data collected by Super-Kamiokande [456], as well as one-by-one limits on $\epsilon_{\mu\tau}^\oplus$ from LBL accelerator $\bar{\nu}_\mu$ disappearance data from MINOS [457] and high-energy, $\mathcal{O}(\text{TeV})$, atmospheric $\bar{\nu}_\mu$ disappearance data from IceCube (“IC 2017 (public)”) [438].

The figure also allows gauging the combined impact of the increased event statistics and the inclusion of higher-energy events in the event sample prepared for this work compared to an external analysis [440] on the published DeepCore sample “B” mentioned in Sec. 11.1 (“IC DC 2020 (public)”). The widths of the 90 % confidence intervals in this thesis are smaller by between $\sim 25\%$ (for $\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus, \epsilon_{e\tau}^\oplus, \epsilon_{\mu\tau}^\oplus$) and $\sim 50\%$

($\epsilon_{e\mu}^\oplus$) with respect to their counterparts from the external measurement. The latter does not study e - μ NU, complex NSI phases, or a generalised matter potential with a less restrictive flavour structure.

Furthermore, we show the limits on FV couplings reported by a measurement using the published timing (or flavour) data from CE ν NS at COHERENT [273]. Here, the assumed underlying NSI model based on the exchange of a Z' mediator with $M_{Z'} \sim \mathcal{O}(10 \text{ MeV})$ dictates $\epsilon_{\alpha\beta}^u = \epsilon_{\alpha\beta}^d$, such that no cancellations between NSI with different quark flavours occur (see Ref. [458] for a comprehensive analysis). While CE ν NS only yields constraints that are valid for a new-physics energy scale above $\mathcal{O}(10 \text{ MeV})$, in contrast to oscillation experiments it is sensitive to the individual flavour-diagonal coupling strengths $\epsilon_{ee}^{uV,dV}$ and $\epsilon_{\mu\mu}^{uV,dV}$. The results of this work are also not directly comparable to NSI measurements from collider experiments, as these commonly depend strongly on the underlying new-physics model and energy scale [459].

Limits derived by the combined analysis [223] of global neutrino oscillation data with negligible sensitivity to CP-violating effects scrutinised in Sec. 5.2.2 are depicted in addition (“global 2018 (w/ correl.)”). These constraints are no more stringent than ours because all NSI couplings are simultaneously fit to the global set of experimental data, such that correlations between couplings with different flavour indices reduce the precision with which any single parameter can be measured.

15.2.1 Assessment of subsequent NSI studies

Results from a small number of NSI studies with atmospheric neutrinos have appeared after the conclusion of the analysis presented in this thesis and are briefly reviewed and contextualised below.

IceCube

A new high-energy IceCube measurement [460, 461] of μ - τ FV based on more than 3×10^5 upgoing track-like events with reconstructed energies between 500 GeV and 10 TeV has constrained the associated coupling strength with an unprecedented precision to the range $-0.0041 \lesssim \epsilon_{\mu\tau}^{dV} \lesssim 0.0031$ (or $-0.0123 \lesssim \epsilon_{\mu\tau}^\oplus \lesssim 0.0093$) at 90 % CL according to Wilks under the assumption of real NSI and the NO.

For the case of complex NSI, Fig. 15.7 (reproduced from Ref. [461, Fig. 4]) also provides a comparison to the 90 % confidence region observed in this work, after conversion to its representation in the $\text{Re} \epsilon_{\mu\tau}^{dV}$ - $\text{Im} \epsilon_{\mu\tau}^{dV}$ plane. For this purpose, the

confidence region of the high-energy measurement has been constructed by applying Wilks' theorem with a single degree of freedom to the observed two-dimensional profile likelihood ratio (cf. Sec. 13.3). Consequently, the critical test-statistic value for a given confidence level is reduced with respect to the assumption of two degrees of freedom. The existence of only one degree of freedom is expected, however, from the degeneracies of the high-energy expression for the $\bar{\nu}_\mu$ survival probability (9.3) and has been validated via the Feldman–Cousins procedure (cf. Sec. 14.4). In contrast, the confidence region on the modulus and phase of $\epsilon_{\mu\tau}^\oplus$ in the bottom of Fig. 15.4 assumes Wilks' theorem with two degrees of freedom. The latter might not be exceedingly conservative in light of the coverage study conducted for the non-trivial NSI model $\epsilon_{\mu\tau}^\oplus = 0.035$ in Fig. 14.15.

Contrary to this work, the high-energy measurement omits Δm_{31}^2 and θ_{23} as nuisance parameters, an approach adopted from preceding IceCube searches for sterile neutrinos for which the event selection was originally developed [462, 463]. Completely vanishing effects from both parameters are only expected in the asymptotic limit represented by Eq. (9.3) though, which is accompanied by a full degeneracy between the two signs of $\epsilon_{\mu\tau}^\oplus$ (for real NSI). This limit is not quite satisfied by the high-energy measurement, as evidenced by the weak preference for a particular sign of $\epsilon_{\mu\tau}^\oplus$ under a given NMO [461, Fig. 2]. When vacuum oscillation terms are considered, at a neutrino energy of, e.g., $E_\nu = 1$ TeV, their contribution to the $\bar{\nu}_\mu$ survival probability half-phase [383] is $\frac{\Delta m_{31}^2}{4E_\nu} \sim \pm 10^{-24}$ GeV. Since $\langle V_{CC} \rangle \sim 10^{-22}$ GeV (as illustrated in Fig. 5.6), the NSI contribution $\epsilon_{\mu\tau}^\oplus \langle V_{CC} \rangle$ is of the same order as the standard one when $|\epsilon_{\mu\tau}^\oplus| \sim 10^{-2}$. This coincides with the precision to which the high-energy measurement constrains the coupling.




Figure 15.7: Observed 90 % confidence regions on complex $\epsilon_{\mu\tau}^{dV}$ from a high-energy IceCube measurement (solid black) and this work (blue). Figure taken from Ref. [461].

ANTARES

Also focussing solely on NSI in the μ - τ sector, the underwater neutrino telescope “ANTARES” [464] off the coast of France in the Mediterranean Sea has performed a joint measurement [465] of $\epsilon_{\tau\tau}^{dV} - \epsilon_{\mu\mu}^{dV}$ and real $\epsilon_{\mu\tau}^{dV}$ based on 7710 track-like events with reconstructed energies from 16 GeV to 100 GeV collected between 2007 and 2016 [466]. This NSI hypothesis is constructed in the same manner as in the Super-Kamiokande study [456] whose results are depicted in Fig. 15.6. In the case of $|\epsilon_{\tau\tau}^{dV} - \epsilon_{\mu\mu}^{dV}|$, the ANTARES upper bound is roughly 60 % weaker than that from Super-Kamiokande. In the case of $|\epsilon_{\mu\tau}^{dV}|$, it is stronger by a factor of approximately two (and, after rescaling, slightly stronger than the limits on real $\epsilon_{\mu\tau}^{\oplus}$ obtained in this work). In fact, ANTARES observes a preference for μ - τ NU over SI at 90 % CL according to Wilks when $\epsilon_{\mu\tau}^{dV}$ is treated as a nuisance parameter. This preference in turn evidently contributes to a weaker upper bound on $|\epsilon_{\tau\tau}^{dV} - \epsilon_{\mu\mu}^{dV}|$ with respect to Super-Kamiokande. The ANTARES fit considers the atmospheric oscillation parameters as unconstrained sources of systematic uncertainty, together with δ_{CP} (unconstrained but without impact) and θ_{13} (subject to a Gaussian prior).

KM3NeT/ORCA




Figure 15.8: Various observed 90 % confidence intervals on real couplings $\epsilon_{\alpha\beta}^{dV}$, including this work (“IC DeepCore 2021”). Figure taken from Ref. [467].

Serving as a low-energy successor to ANTARES and situated close to the latter’s location, the densely instrumented Gt-scale “KM3NeT/ORCA” neutrino telescope

in its partial, six-string configuration (“ORCA6”) has already allowed conducting proof-of-concept NSI measurements [468, 467]. These closely follow the analysis approach of this work (with fits of the hypotheses in Table 13.1 except for e - μ NU and the GMP). A study based on an exposure of 433 kt-years [467], corresponding to 5 828 events distributed over two track-like bins (with different $\bar{\nu}_\mu$ CC purities) and one cascade-like bin, reports intervals allowed at 90% CL (Wilks) on the real NSI couplings $\epsilon_{\tau\tau}^{dV}$, $\epsilon_{\mu\mu}^{dV}$, $\epsilon_{e\mu}^{dV}$, $\epsilon_{e\tau}^{dV}$, and $\epsilon_{\mu\tau}^{dV}$. Figure 15.8 shows that the widths of these intervals deviate by at most some tens of percent from the rescaled results of this work. Most prominently, ORCA6 yields a weaker bound on the absolute value of $\epsilon_{e\mu(e\tau)}^{dV}$ when the phase $\delta_{e\mu(e\tau)} = 180^\circ$, but a somewhat stronger bound for $\delta_{e\mu(e\tau)} = 0^\circ$. Also depicted are the aforementioned measurement results for real $\epsilon_{\mu\tau}^{dV}$ from IceCube and ANTARES, representing the highest-precision constraints on $\epsilon_{\mu\tau}^{dV}$ to date, as well as the MINOS measurement of the same parameter [457]. As is the case in this work, the atmospheric oscillation parameters serve as unconstrained nuisance parameters in the ORCA6 measurement; the NMO is treated as a nuisance parameter throughout.

Ongoing DeepCore studies

Today, as the number of recorded DeepCore trigger events exceeds that underlying this work by a factor greater than three, analysis efforts aiming at constantly expanding the experiment’s standard and non-standard oscillation physics reach are ongoing. Reference [469] has made use of several modified data-processing and simulation procedures to investigate the NSI sensitivity of an all-flavour event sample of $\sim 1.5 \times 10^5$ exclusively upgoing events with reconstructed energies of up to 100 GeV and a low (μ^\pm and noise) background contamination of less than 1%. One-by-one fits of the μ - τ NU and the three FV NSI scenarios to the Asimov pseudodata reveal expected improvements at the level of several tens of percent with respect to this work’s observed 90% confidence intervals on real NSI couplings. For example, the median interval to which real μ - τ FV is expected to be constrained given true NO is $-0.009 \lesssim \epsilon_{\mu\tau}^\oplus \lesssim 0.011$, whose width is almost identical to that of the aforementioned high-energy IceCube measurement.

15.3 Rejecting SI: from Wilks to Feldman–Cousins

The available pseudoexperiments generated under the SI hypothesis and examined in Sec. 14.4 enable us in principle to refine our prior statements about the significance




Figure 15.9: Feldman–Cousins approach for determining the significance of rejecting the SI hypothesis in the two one-by-one NU NSI measurements. In each case, the left panel shows a small interval of the observed $\Delta\chi^2_{\text{mod}}$ profile around the SI hypothesis, and the right panel the CDF $F(\Delta\chi^2_{\text{mod}}|0)$ from which the confidence level of the SI rejection can be read off (solid and dashed red lines). The CDF of χ^2_1 , used in conjunction with Wilks’ theorem, is also provided for comparison as a dotted line.

of rejecting the SI hypothesis (or flavour-universal NSI in the case of either NSI NU measurement) by locally adopting the Feldman–Cousins approach. An enhanced rejection will be encountered if the PDF of $\Delta\chi^2_{\text{mod}}(0|0)$ displayed in Fig. 14.14 associated with a given NSI scenario drops off much more rapidly than expected from Wilks’ theorem. In the cases of e - μ , e - τ , and μ - τ FV, in which the complex phases become unphysical under the SI hypothesis, the observed confidence level of rejecting the latter will coincide with that of rejecting $|\epsilon_{\alpha\beta}^{\oplus}| = 0$. In the case of the GMP, the rejection level for the SI hypothesis ($\epsilon_{\oplus} = 1, \varphi_{12} = 0, \varphi_{13} = 0$) need not coincide with that for a standard strength of the matter potential ($\epsilon_{\oplus} = 1 \forall \varphi_{12}, \varphi_{13}$).

Nominal nuisance-parameter values underlie the generation of the available pseudoexperiments, not their observed best-fit values. It is conceivable for these offsets to affect the shapes of the $\Delta\chi^2_{\text{mod}}(0|0)$ PDFs—even after profiling over all nuisance parameters. Here, however, we address the above questions by neglecting any dependence on the true nuisance-parameter values. This is justified because the deviations between the nominal and best-fit points are not significant or neither are close to the bounds of their manually preimposed or physical ranges. We thus use the nominal distributions displayed in Fig. 14.14 directly in order to compute the tail probabilities with respect to the observed outcomes $\Delta\chi^2_{\text{mod,SI}}$ in Table 15.1. Then, as illustrated in Figs. 15.9 to 15.11, the SI hypothesis is rejected by the data at the following confidence levels and significances (compared to using Wilks’ theorem):

- e - μ NU: 77% (vs. 75% Wilks and χ^2_1), or 1.20σ (vs. 1.15σ);




Figure 15.10: Same as Fig. 15.9, but for the FV NSI measurements. The CDF of χ^2_2 represents Wilks' theorem applied to a hypothesis with two parameters of interest in spite of the unphysical complex phases for SI (see text).

- μ - τ NU: 34 % (vs. 4 % Wilks and χ^2_1), or 0.44σ (vs. 0.05σ);
- e - μ FV: 78 % (vs. 44 % Wilks and χ^2_2), or 1.23σ (vs. 0.58σ);
- e - τ FV: 56 % (vs. 21 % Wilks and χ^2_2), or 0.77σ (vs. 0.27σ);
- μ - τ FV: 18 % (vs. 4 % Wilks and χ^2_2), or 0.23σ (vs. 0.05σ);
- GMP: 56 % (vs. 46 % Wilks and χ^2_3), or 0.77σ (vs. 0.61σ).

While the confidence level of excluding SI rises considerably in all measurements but those of e - μ NU and the GMP, it remains low overall and far below any conventional threshold required for a discovery claim. With the Feldman–Cousins method, the strongest statistical exclusion of SI is encountered in the e - μ FV and e - μ NU measurements, at a Gaussian-equivalent significance of $\sim 1.2\sigma$.




Figure 15.11: Same as Fig. 15.9, but for the GMP fit. The displayed profile is a slice along ϵ_{\oplus} through the full three-dimensional profile for fixed $\varphi_{12} = \varphi_{13} = 0^\circ$, allowing to read off the SI rejection level at the value $\epsilon_{\oplus} = 1$. By definition, the profile slice is shifted upward by a constant number of units relative to the $\Delta\chi^2_{\text{mod}}$ profile in $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ (after taking into account the translation $\epsilon_{\oplus} = \epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus} + 1$).

Part IV

Monte Carlo studies on the neutrino mass ordering with PINGU

16

Event sample

The MC event sample employed in this analysis of the PINGU experiment has been generated with the same IceCube software tools used in Part III. Therefore, Sec. 16.1 only gives a condensed summary of the assumptions and procedures underlying the PINGU event simulation and reconstruction. The criteria according to which events are selected for the NMO sensitivity study are presented in Sec. 16.2, which is followed by a brief assessment of the criteria’s feasibility once PINGU is deployed in Sec. 16.3.

16.1 Simulation and reconstruction

Given the PINGU geometry detailed in Sec. 8.1, for each atmospheric neutrino event generated by GENIE the propagation of Cherenkov photons is subjected to the “SPICE MIE” model of the bulk ice mentioned in Sec. 6.2.3. The nominal hole-ice assumptions represented by the effective angular acceptance functions of IceCube and DeepCore DOMs in Fig. 10.4 are conservatively equally applied to PINGU DOMs. The latter are assumed to have universal response functions and dark-noise distributions, given by the nominal model of a high-QE DOM.

The event reconstruction algorithm producing the neutrino properties’ analysis-level estimators is in line with the general considerations of Sec. 10.2. It consists of the hybrid template likelihood ansatz for the hypothesis of a single cascade accompanied by a collinear minimum-ionising μ^\pm track. The charge expectations for both types of sources are generated with the bulk- and hole-ice models of the simulation. The reference energy of each minimum-ionising μ^\pm template is assumed to be ~ 1.7 GeV, which corresponds to an average track length of 7 m according to Eq. (10.6). MULTINEST explicitly samples the full eight-dimensional parameter space.

16.2 Selection criteria

The PINGU event sample at the analysis level is extracted from all simulated atmospheric neutrino events by applying three sets of event criteria applied in succession and summarised in Table 16.1.

The first set (level 1) corresponds to the dedicated PINGU SMT3 defined in Sec. 8.1. It results in an expected event rate of the order of 100 Hz, with the atmospheric μ^\pm contribution still exceeding that of atmospheric neutrinos by almost four orders of magnitude [470].

As alluded to in Sec. 8.1, the second level of the event selection is intended to serve as a proxy of the online filter, in terms of both background-removal efficiency and computational cost. Here, we encounter a number of simple cuts based on the temporal and spatial distribution of preselected ensembles of pulses throughout the event (“seeded $R-T$ ” and time-window cleaning procedures with parameters adapted to the PINGU geometry). At least eight DOMs are required to record signals, 50 % respectively 90 % of the total charge has to be detected within 600 ns after the second hit and within 2 μ s after the first hit of the event (cf. Fig. 11.3), and the earliest quartile of hits must have a mean depth which is on average at least 30 m closer to the bottom of the detector than the full set of hits. The latter criterion is a spatial complement to the preceding pair of temporal μ^\pm -rejection criteria. A fast estimate of the interaction position, taken as the charge-weighted spatial average of hit DOMs, must yield a depth $z < -200$ m for the event to be considered. Any event that passes the previous cuts then undergoes a first likelihood-based cascade reconstruction. The event is only retained if the radial distance of the reconstructed vertex from the centremost IceCube string is below 95 m. This ensures that only interactions occurring at most a few tens of metres outside of the densely instrumented PINGU volume in the horizontal are considered. Together, the above criteria reduce the atmospheric μ^\pm event rate by approximately two orders of magnitude, whereas the atmospheric neutrino rate is nearly cut in half [470].

Prior to the application of the third level of selection criteria, events surviving all preceding ones undergo the computationally expensive MULTINEST-based hybrid event reconstruction algorithm. The outcome of the third level depends solely on the resulting estimators: an event is only retained if the radial distance of its reconstructed vertex to the centremost IceCube string is at most 85 m, and if in addition its reconstructed interaction depth lies in the range between -500 m and -180 m. This reduces the expected atmospheric μ^\pm event rate to the order of 10 mHz, which is comparable to the rate of atmospheric neutrino events [470].

Level	Criterion	Requirement	Target
1	PINGU SMT3		noise
2	number of DOMs hit	≥ 8	quality
	charge fraction within 600 ns after 2nd hit	> 0.5	μ^\pm
	charge fraction within 2 μs after 1st hit	> 0.9	μ^\pm
	mean hit depth difference to mean depth of 1st hit quartile	$> 30 \text{ m}$	μ^\pm
	fast estimate of interaction depth	$< -200 \text{ m}$	quality, μ^\pm
	first estimate of radial cascade distance	$< 95 \text{ m}$	quality, μ^\pm
3	reconstructed starting radial distance	$\leq 85 \text{ m}$	quality, μ^\pm
	reconstructed starting depth	$\in [-500 \text{ m}, -180 \text{ m}]$	quality, μ^\pm

Table 16.1: Summary of the event selection criteria employed by the PINGU NMO sensitivity studies, except for requirements on analysis variables. The criteria of both level 2 and level 3 are applied to variables derived from pulse ensembles that have undergone noise removal. The formal definitions of the various cut criteria can be found in Refs. [45, 470].

The variables considered at the analysis level are the reconstructed neutrino energy, zenith angle, and an event classification score. Only upgoing trajectories are included, $\cos \vartheta_{\text{reco}} \leq 0$, the energy is restricted to the range $E_{\text{reco}} \in [1 \text{ GeV}, 80 \text{ GeV}]$, and events are separated into cascade- and track-like categories. The event classification score is obtained by a neural network whose input variables are taken from the event's observed hit pattern and the hybrid event reconstruction algorithm applied prior to level 3. The reconstructed variables include the length of the minimum-ionising μ^\pm track, the fraction of the total neutrino energy carried away by the latter, and the likelihood ratio $L_{\text{cascade+track}}/L_{\text{cascade}}$ (cf. Sec. 11.2.2). More details about this procedure can be found in Ref. [1].

16.3 Practicability

The event selection procedure considered in this part consists of a relatively simple set of criteria which nonetheless result in a large relative suppression of the background compared to the atmospheric neutrino event rate. The criteria's preliminary and ad-hoc nature was necessitated by (the need for) repeated changes in

PINGU design assumptions and the continuous adaptation of the event simulation and reconstruction software to these. Their practicability is investigated in detail in Ref. [470], whose key findings provide support for the assumed neutrino selection efficiency and assumed lack of background contamination of the PINGU event sample on which this work relies.

It is found that sequentially estimating the radial cascade distance of the event—which underlies the last criterion of level 2 in Table 16.1—proceeds at a rate two orders of magnitude below the expected (online) event rate of the aggregate of all six prior requirements. Furthermore, in order to be able to perform the third level’s cuts at the same rate at which events pass the second level, on the order of 10^3 concurrently running reconstruction processes would be necessary. While the latter is achievable in practice because the reconstruction would be performed offline, after the transmission of the filtered event data to the north, the resource capacity of the ICL renders the online suitability of the requirements of level 2 questionable. However, it turns out that it is possible to find an extended set of computationally lightweight background veto criteria with proven performance in DeepCore instead, whose background rejection and atmospheric neutrino efficiencies exceed those of the level 2 at hand.

At the same time, Ref. [470] also suggests that by further filtering the resulting event data stream arriving in the north the number of events to reconstruct could be reduced by roughly one order of magnitude. As a desired side effect of all of the above, the background contamination of the final event sample also decreases by nearly one order of magnitude without a loss in the overall neutrino efficiency of the selection.

17

Simulation predictions

Finding the expected event distribution (template) from the PINGU event sample described in the previous chapter for any set of oscillation parameter values is the core prerequisite for an accurate projection of the experiment's sensitivity to any of these parameters. To this aim, in Sec. 17.1 we discuss a staged approach consisting of consecutive transformations of histograms representing the intrinsic atmospheric neutrino fluxes. This allows us to efficiently and accurately calculate the sought template and quantify the detector response characteristics even in the case of comparably low MC event statistics [4]. Many of the issues pertaining to the prediction of the neutrino template via the event-by-event reweighting scheme from Sec. 12.1 are also relevant here. For example, the concept of the staged approach also lacks the need to repeatedly simulate and reconstruct events each time the injected value of a model parameter is changed. While the nominal template is characterised in Sec. 17.2, the imprints left by the atmospheric oscillation parameters are scrutinised in Sec. 17.3. An overview of various sources of systematic uncertainty that will likely be of importance in the operating experiment follows in Sec. 17.4.

17.1 Staged approach




Figure 17.1: Schematic illustration of the staged approach adopted in the generation of the neutrino template for PINGU, where separate stages are numbered as follows: 1) atmospheric flux, 2) oscillation, 3) detection, 4) reconstruction, 5) event classification. All fluxes, transformations, and event counts are shown as one-dimensional functions of energy for simplicity only. Transformation functions relying on PINGU MC events are shown in cyan, while transformations (and the atmospheric neutrino flux) relying on external MC or numerical calculations are shown in blue. The multiplication and convolution operators “ \times ” respectively “ \otimes ” should only be understood symbolically, as the actual operations mapping the output of one stage to that of the next are more complex (e.g., additional summation over different oscillation channels in stage 2). Figure adapted from Ref. [4]. See text for details.

A simplified high-level overview of the staged approach is shown in Fig. 17.1. We factorise the description of the physics into neutrino production in the atmosphere, propagation including flavour transitions, detection, reconstruction, and event classification. For illustration, a particular neutrino type is tracked from its intrinsic

(unoscillated) atmospheric flux to the associated distributions of cascade-like and track-like (CC or NC) events for a given detector exposure time.

The starting point is a histogram representation of the flux. Each subsequent stage (2 to 5) outputs and applies a transformation which is evaluated on a discrete set of grid points in the—generator-level or reconstructed—space of neutrino energy and zenith (only energy displayed for simplicity). The number of grid points considered by each stage is smaller than the number of MC events based on which the corresponding transformation function is calculated. At each stage and for each transformation, a suitable parameterisation or smoothing method can be applied that is adapted to the nature of the neutrino physics processes governing that particular component of the experiment. In this way, statistical fluctuations in the underlying MC event sample are more easily mitigated than by applying generic smoothing methods to the final event template (either before or after superimposing the different event types). The latter is the convolution of many physics processes with highly disparate impacts on the shapes of the original unoscillated atmospheric neutrino fluxes.

In order to complete this high-level overview, Fig. 17.2 illustrates the various paths across neutrino and interaction type connecting the outputs of successive stages. Each stage computes transformations either for each individual neutrino (and possibly interaction) type or for several suitably chosen groupings (such as all NC events).

Initially, unoscillated flux histograms are obtained separately for ν_e , $\bar{\nu}_e$, ν_μ , and $\bar{\nu}_\mu$. These are used as inputs to the oscillation stage, which computes the respective probabilities for neutrino or antineutrino oscillation into any of $\bar{\nu}_{e,\mu,\tau}$ and returns the six resulting oscillated flux histograms according to Eq. (12.2).

Drawing on PINGU generator-level events and their associated GENIE weights (cf. Sec. 10.1), the detection stage calculates a dedicated effective area (scaled with the assumed exposure time) for each of the six CC event types, for neutrino NC events, and for antineutrino NC events, where (anti)neutrino NC events of all flavours are grouped together to enhance statistics. Multiplying each effective area with the appropriate oscillated flux yields the expected event counts in the space of the generator-level variables for all twelve types.

The reconstruction stage uses PINGU MC events which have already undergone the reconstruction algorithm of Sec. 16.1 to generate its transformations. These represent mappings from the event-count histograms in the space of generator-level to those in the space of reconstructed variables. In this process, CC neutrino and an-

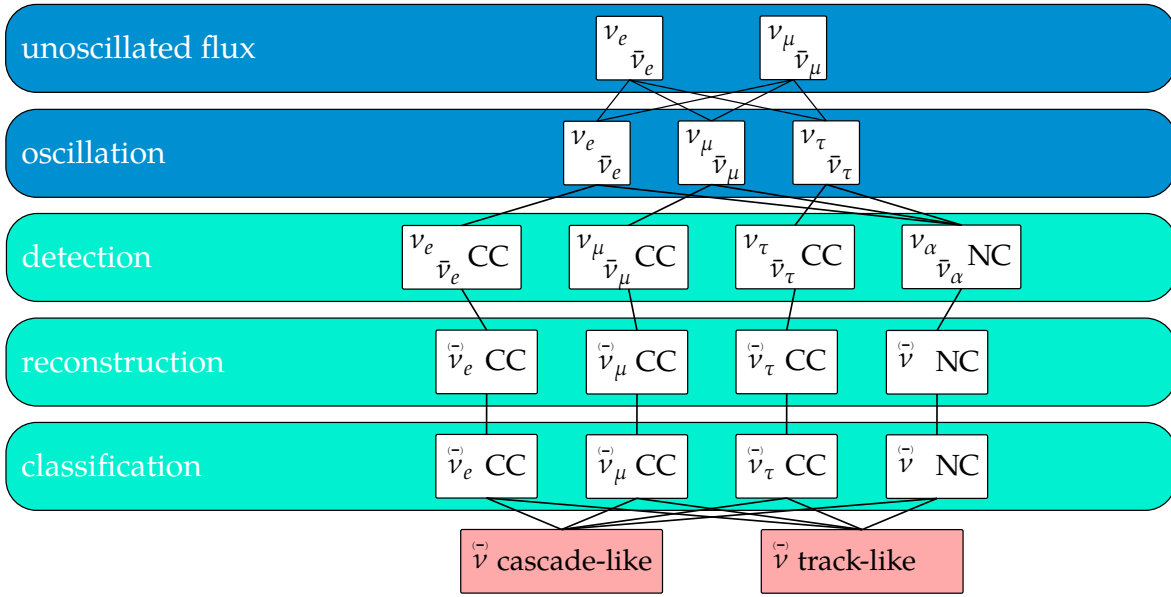


Figure 17.2: Neutrino- and interaction-type flow chart of the staged approach for PINGU, where the contents of the white boxes specify the event types output by the various stages. $\bar{\nu}_\alpha$ indicates summation over neutrinos and antineutrinos of flavour α , and the absence of a subscript summation over all three flavours in addition. Figure adapted from Ref. [4]. See text for details.

tineutrino events of a particular flavour are grouped together, as are all NC events. For convenience, the stage also only outputs the four corresponding histograms.

Finally, the last stage translates the distribution of the classification score (precomputed for each MC event) within a given input event group into track- and cascade-like classification probabilities. These split each input histogram into a track-like and a cascade-like output histogram. All four output histograms of either class are then summed to produce the analysis-level event template. For all practical intents and purposes, we can consider the event class as a third binning dimension of this template, just as in Part III, where the event template is *directly* populated with DeepCore MC events, however.

Stage implementations

Below we describe the assumptions and computational techniques employed in the calculation of each stage's transformations (or, in the case of the flux stage, outputs) in greater depth. Minor deviations exist with respect to Ref. [4] because this thesis is based on a slightly earlier iteration of the staged approach.

Flux The atmospheric neutrino flux predictions are obtained from the tables by Honda *et al.* [281] discussed in Sec. 12.1, which are interpolated with cubic splines in order to obtain smooth representations [392] of the unoscillated fluxes $\Phi_{\text{unosc}}^\alpha(E_\nu, \vartheta)$. These are histogrammed on a two-dimensional grid consisting of $n_E = 39$ bins spaced uniformly in $\log_{10}(E_\nu/\text{GeV})$ that cover the range [1 GeV, 80 GeV] and of $n_\vartheta = 20$ bins spaced uniformly in $\cos \vartheta$ that cover the range $[-1, 0]$. Contributions from downgoing neutrinos and those with energies outside the specified range, from which no NMO sensitivity is expected, are not accounted for by this analysis.

Oscillation The neutrino oscillation probabilities are calculated via a GPU-accelerated adaptation of Prob3++ according to the approach detailed in Appendix A, assuming the standard Earth matter potential. For a given oscillation channel $\nu_\alpha \rightarrow \nu_\beta$ and each histogram bin of the unoscillated flux $\Phi_{\text{unosc}}^\alpha$, a single oscillation probability is obtained as the average over an enclosed 10×10 grid in E_ν and $\cos \vartheta$. This degree of sampling granularity—necessitated by the histogram nature of the staged approach—has been found to yield accurate representations of the oscillated flux in any given bin while keeping the overall computational cost of the template generation process manageable.

Detection In order to convert the oscillated fluxes into event counts still in the space of the generator-level variables E_ν and $\cos \vartheta$, the detection stage is the first of three consecutive stages to employ dedicated PINGU MC events. The detector's effective areas are smoothed two-dimensional functions of E_ν and $\cos \vartheta$, extracted from MC events satisfying the selection criteria of Table 16.1.

The stage first performs MC integration on a two-dimensional grid with a user-defined binning in E_ν and $\cos \vartheta$ (here: same as flux stage). For a given neutrino- and interaction-type grouping, the energy dependence of the effective area is then characterised by cubic splines through the individual $\cos \vartheta$ slices, $\{A_{\text{eff}}(E_\nu)\}_{i=1, \dots, n'_\vartheta}$, which are subsequently evaluated at the n'_E energy bins' centres. One such slice (vertically upgoing) is shown on the left of Fig. 17.3 for ν_e CC events. Afterwards, the zenith dependencies are characterised in an analogous manner, $\{A_{\text{eff}}(\cos \vartheta)\}_{i=1, \dots, n'_E}$, using the aforementioned smoothed energy dependencies as the values to be interpolated.

Finally, the above samples $\{A_{\text{eff}}(E_\nu, \cos \vartheta)\}_{i=1, \dots, n'_E n'_\vartheta}$ are interpolated with a two-dimensional cubic spline that can be evaluated on any grid specified by the user (here: also that of the flux stage). The panel on the right of Fig. 17.3 gives the full effective-area transformation for ν_e CC events as computed by the stage.




Figure 17.3: Raw MC (left) and smoothed representations (left and right) of PINGU’s effective area for the detection of (vertically upgoing) ν_e CC events selected for analysis. See text for details.

Reconstruction The reconstruction stage maps the various event distributions in the space of the generator-level variables E_ν and $\cos \vartheta$ onto their counterparts in the space of the reconstructed observables E_{reco} and $\cos \vartheta_{\text{reco}}$ by means of smoothed detector resolution functions. The integration of a given event group’s resolution functions over the stage’s $n_E n_\vartheta$ output bins in E_{reco} and $\cos \vartheta_{\text{reco}}$ then yields the transformation to be applied to the histograms of its constituent event types, which is a fourth-order “migration” tensor (here: with 780 entries).

For a given grouping, an ensemble of resolution functions is built up across the n_E input bins in E_ν . The dependence of the resolution functions on $\cos \vartheta$ is weak and thus not taken into account. Within a given energy bin, the energy and zenith reconstruction residuals of all n_{evts} contained events are considered as uncorrelated samples from the same one-dimensional resolution functions in $\Delta E \equiv E_{\text{reco}} - E_\nu$ and $\Delta \cos \vartheta \equiv \cos \vartheta_{\text{reco}} - \cos \vartheta$, respectively. The width of the bin is allowed to expand for two reasons: first, in order to reach a predefined minimum threshold of $n_{\text{min}} = 100$ events, and, second, in order to reach a predefined target threshold of $n_{\text{tgt}} = 300$ events. The expansion is performed by increasing the maximally allowed absolute deviation between an event’s true energy and the bin centre E_{mid} , $|E_\nu - E_{\text{mid}}|$. For the first purpose, no maximum is imposed on the deviation, whereas for the second a maximum of the quadruple width of the original bin is imposed. Only the condition $n_{\text{evts}} \geq n_{\text{min}}$ is therefore guaranteed.

For each bin’s MC events, we generate one variable-bandwidth kernel density estimate (VBWKDE) of the distribution of ΔE , and one of the distribution of $\Delta \cos \vartheta$. Each VBWKDE construction is initiated via a fixed-bandwidth KDE [471] based




Figure 17.4: VBWKDEs of PINGU’s energy (left) and cosine-zenith (right) reconstruction residuals for the four event groups of the reconstruction stage, with true energies restricted to the narrow interval indicated at the top. For comparison, histograms of the raw MC residuals of $\nu_e + \bar{\nu}_e$ CC events are also shown.

on an improved version of the “Sheather-Jones” bandwidth selection criteria [472] which does not assume a Gaussian distribution of the reconstruction residuals. Each individual kernel’s bandwidth is subsequently adaptively rescaled by the inverse square root [473] of the initial probability density estimate of the residuals at its location; a second, constant rescaling factor ensures that the narrowest bandwidth corresponds to the initial estimate and occurs at the density maximum [4]. The shapes of the individual kernels to be summed over are Gaussian. As an example, Fig. 17.4 contrasts the resulting VBWKDEs of the four event groups for the energy interval $10.59 \text{ GeV} \leq E_\nu \leq 11.85 \text{ GeV}$. The original residual distributions themselves are provided for $\nu_e + \bar{\nu}_e$ CC events in addition.

Classification The final, classification, stage carries over the four two-dimensional event distributions in the space of E_{reco} and $\cos \vartheta_{\text{reco}}$ to a single three-dimensional event distribution, with an added dual-bin event-class dimension. This is achieved through smoothed classification probabilities, assumed to be one-dimensional functions of E_{reco} .

After predefining a cascade-like and a track-like range for the event classification score, the probability of classifying a given grouping’s events as either of the two is described in dependence of E_{reco} . The $\cos \vartheta_{\text{reco}}$ dimension is disregarded because the classification-score distributions show a negligible dependence on it. The final transformation functions are computed by applying “Blackman-window” smooth-




Figure 17.5: Smoothed probabilities of classifying PINGU events as track-like for the four event groups of the last stage, together with a histogram of the raw MC values of $\nu_\mu + \bar{\nu}_\mu$ CC events.

ing [474] to finely binned (300 intervals spaced uniformly in $\log_{10}(E_{\text{reco}}/\text{GeV})$) raw classification probabilities for the cascade-like event class and downsampling the result (respectively its complement) to the $n_E = 39$ energy bins of the input and output event-count histograms of the stage. The transformation functions of the four event groups are displayed in Fig. 17.5. The original probabilities for the classification of $\nu_\mu + \bar{\nu}_\mu$ CC events as track-like are provided in addition.

17.2 Event-sample characterisation

For the remainder of this part, the nominal neutrino oscillation parameter values are chosen to correspond to those obtained by NuFIT 2.0 [475, 476], both for the NO and the IO hypothesis. Where predictions about the sample are made without reference to the NMO choice, it is understood that we assume the best-fit parameters within the global fit’s NO parameter space. A single exception to this rule is again made for δ_{CP} , which is set to zero. Under these assumptions, the staged approach predicts an overall event rate of 2.2 mHz at the analysis level, to which cascade-like and track-like events contribute 1.6 mHz and 0.6 mHz, respectively.

The nominal PINGU event template, based on which statistical inference is made, is depicted in Fig. 17.6 (event count normalised to one year, or 365 days, of exposure). It employs 39 bins spaced uniformly in $\log_{10}(E_{\text{reco}}/\text{GeV})$, covering the range from $E_{\text{reco,min}} = 1 \text{ GeV}$ to $E_{\text{reco,max}} = 80 \text{ GeV}$, and 20 bins spaced uniformly across $-1 \leq \cos \vartheta_{\text{reco}} \leq 0$. The highest rate is expected for rather weakly inclined cascade-like events (left panel) with reconstructed energies of around 2 GeV. Less frequently




Figure 17.6: Predicted distribution of neutrino events for standard oscillations after one year of PINGU operation using the analysis binning in E_{reco} and $\cos \vartheta_{\text{reco}}$, with cascade-like events on the left and track-like events on the right.

detected track-like events (right panel) are most likely to occur at slightly higher reconstructed energies. The visible “valley” extending almost diagonally across to the histogram’s upper left corner corresponds to the first survival minimum of $\bar{\nu}_{\mu}$ ’s (and the associated appearance maximum of $\bar{\nu}_{\tau}$ ’s). As demonstrated in the following section, the net effect of this disappearance is both a localised and an overall event rate reduction compared to the no-oscillation scenario. Its amplitude and location are governed by θ_{23} and Δm_{31}^2 .

17.3 Atmospheric neutrino oscillation signatures

Using the statistical pull from Sec. 12.4 allows us to quantify the impact on the PINGU event template due to oscillation in the effective 2ν subsystem (ν_e, ν'_3) defined in Sec. 5.1.2.2. A few discrete combinations of injected and hypothesised values of the atmospheric oscillation parameters θ_{23} and Δm_{31}^2 (including its sign) are considered in order to gain a basic understanding of the origins and limitations of PINGU’s sensitivity to these.

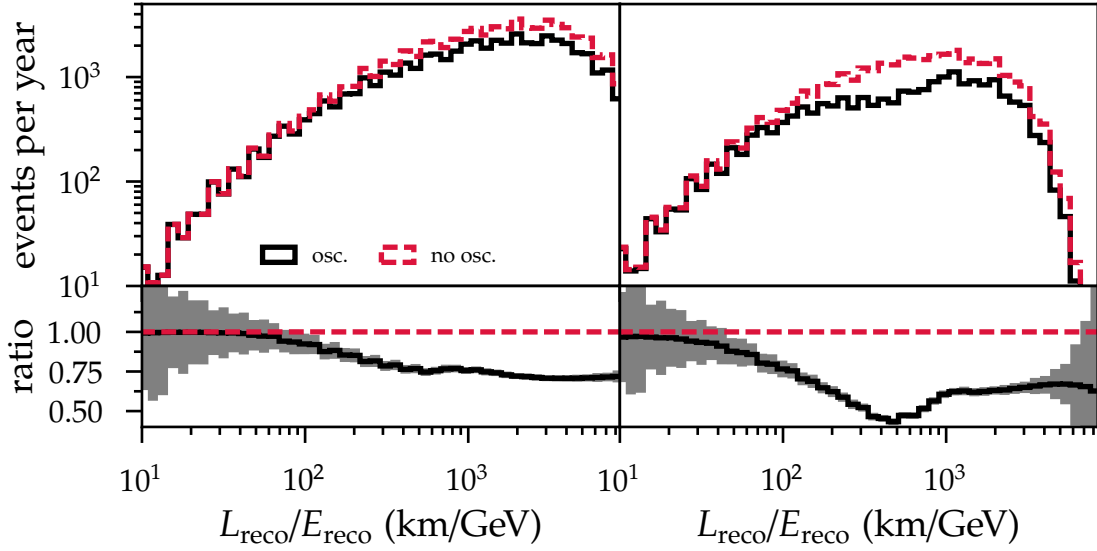


Figure 17.7: Expected annual event rate distribution in PINGU as a function of the ratio of reconstructed neutrino propagation distance to reconstructed energy, $L_{\text{reco}}/E_{\text{reco}}$, assuming standard oscillations (black) or no oscillations (red dashed), for cascade-like events on the left and track-like events on the right. The lower panels show the ratios of the two pairs of distributions. The gray bands therein correspond to the sizes of the expected statistical (Poissonian) uncertainties. See text for details.

General oscillation signature

Figure 17.7 contrasts the nominal event template as a function of the reconstructed L/E_ν ratio to the event template under the assumption of no oscillations (equivalent to setting $\Delta m_{21}^2 = \Delta m_{31}^2 = 0$). All histograms peak at $L_{\text{reco}}/E_{\text{reco}} \sim 2 \times 10^3$ km/GeV, where oscillations are expected to suppress the event rate by some 30% with respect to the no-oscillation scenario. This suppression remains flat for even higher values of $L_{\text{reco}}/E_{\text{reco}}$, but a pronounced disappearance maximum of the order of 50% is encountered for track-like events as $L_{\text{reco}}/E_{\text{reco}}$ approaches 5×10^2 km/GeV. Much less structure exists in the ratio of the cascade-like event counts.

According to Eq. (3.33), requiring the Δm_{31}^2 -induced oscillation phase $\Phi_{31} = \pi$ with $\Delta m_{31}^2 = 2.5 \times 10^{-3}$ eV² yields $L/E_\nu \approx 5 \times 10^2$ km/GeV. The fact that this is precisely where the most pronounced disappearance occurs in the track-like event sample hints at the applicability of the 2ν vacuum survival probability $P_{\mu\mu}^{2\nu}$ in Eq. (5.27).

The second minimum of $P_{\mu\mu}^{2\nu}$ at $(L/E_\nu)_{2\text{nd}} \approx 1.5 \times 10^3$ km/GeV, is largely indiscernible even in the distribution of track-like events. Due to the upper bound $L \lesssim d_\oplus$, the corresponding true neutrino energy is also bounded from above, $E_{\nu,2\text{nd}} \lesssim 8.5$ GeV. The largest contribution stems from events with energies well below this bound, be-




Figure 17.8: PINGU’s predicted binwise annual statistical significance of $\theta_{23} = \pi/4$ with respect to the nominal, first-octant model. Cascade-like events are displayed on the left, track-like events on the right. See text for details.

cause the intrinsic atmospheric neutrino flux does not peak near the vertical direction. Hence, the second minimum is substantially smeared out in the reconstructed ratio $L_{\text{reco}}/E_{\text{reco}}$. In addition, at energies of a few GeV the discriminating power of the event classification is weak, with a large majority of events being classified as cascade-like. For illustration, Fig. D.1 in the appendix shows the oscillation signature in the ν_{μ} event template without dilution or smearing effects. Here, the higher-order survival minima are easily discernible.¹

Atmospheric mixing angle θ_{23}

Setting $\theta_{23} = \pi/4$, corresponding to a relative increase (decrease) of approximately 10% in $\sin^2 \theta_{23}$ ($\cos^2 \theta_{23}$) and a 1% increase in $\sin^2 2\theta_{23}$ with respect to the nominal, first-octant model leads to the annual event-count pull distribution of Fig. 17.8, with magnitudes reaching $\mathcal{O}(0.1)$. For cascade-like events, a weak deficit is brought about in the energy regime of a few GeV, whereas a slightly more pronounced excess appears above, in particular for E_{reco} between 5 GeV and 12 GeV. Track-like events on the other hand exhibit a deficit across all bins. The strongest occurs along the first

¹The assumption that the intrinsic ν_{μ} flux is reduced by $P_{\mu\mu}^{2\nu}$ is a good approximation except for small regions in the E_{ν} - $\cos \vartheta$ plane where matter effects dominate the survival probability. Accordingly, matter effects result in a small decrease of the amplitudes of the oscillations observed in Fig. D.1 with respect to propagation in vacuum.

survival minimum of $P_{\mu\mu}^{2\nu}$ and along a horizontal band extending across the energy range of 6 GeV to 8 GeV.

It is straightforward to understand the signature of θ_{23} in the regime of high energies or weak matter effects on 1–3 mixing. Here, PINGU essentially observes oscillations only through the 2ν vacuum-like disappearance of the intrinsic atmospheric ν_μ and $\bar{\nu}_\mu$ fluxes and the accompanying appearance of ν_τ and $\bar{\nu}_\tau$ fluxes. The small increase in $\sin^2 2\theta_{23}$ from assuming maximal 2–3 mixing instead of the first octant either reduces the $\bar{\nu}_\mu$ flux by the same factor or causes it to vanish precisely at the survival minimum. It also gives rise to a correspondingly higher $\bar{\nu}_\tau$ flux, whose events are identified predominantly as cascade-like. Since the first $\bar{\nu}_\mu$ survival minimum has $\sin^2 \phi_{\text{vac}} = 1$, the largest probability-level sensitivity to θ_{23} is expected along the corresponding valley in the nominal event template on the right of Fig. 17.6. This is directly reflected in the valley of statistical pulls at the same location in the right panel of Fig. 17.8. Along this valley itself, the significance is lowest for more inclined and higher-energy events, owing to their reduced intrinsic $\bar{\nu}_\mu$ fluxes. The fact that no pronounced ridge of positive pulls exists along the first $\bar{\nu}_\tau$ appearance maximum in the cascade-like event category is due to their smaller detection cross sections and worse reconstruction resolutions, combined with the non-trivial superposition of contributions from all flavours in the cascade-like event category.

In the region near the mantle resonance, the positive pull pattern observed in the cascade-like event distribution is dominated by the transition channel $\nu_\mu \rightarrow \nu_e$, with subleading contributions from the transitions $\nu_e \rightarrow \nu_{\mu,\tau}$ and the disappearance $\nu_\mu \rightarrow \nu_\mu$. Making use of Eq. (B.2), neglecting antineutrinos as well as ΔP_{ee} and $\Delta P_{\mu\tau}$, it follows that

$$\Delta N(\Delta\theta_{23} > 0) \propto \Delta P_{\mu e} r_e + \left(\Delta P_{\mu\mu} + \frac{\Delta P_{\mu e}}{R_{\mu/e}^0} \right) r_\mu + \frac{\Delta P_{e\tau}}{R_{\mu/e}^0 a_{e/\tau}^0} r_\tau. \quad (17.1)$$

Here, only the first term $\propto \Delta P_{\mu e} > 0$ is unsuppressed and is therefore largely responsible for the net positive θ_{23} signature. $\Delta P_{\mu\mu} < 0$ only contributes in proportion to the comparably small ν_μ misidentification probability $r_\mu < r_e$. The second term $\propto \Delta P_{\mu e} > 0$ is doubly suppressed, as is the contribution $\propto \Delta P_{e\tau} < 0$. A similar reasoning can be adopted for the same region in the track-like event category, where $r_\mu \gg r_{e,\tau}$ though.

For cascade-like events with energies $E_{\text{reco}} \lesssim 5$ GeV near the threshold, $P_{\mu e}$ and $P_{e\tau}$ become exceedingly small, and therefore so do $\Delta P_{\mu e}$ and $\Delta P_{e\tau}$. The factors r_e, r_τ , and $R_{\mu/e}^0 \sim 2$ remain nearly constant, but r_μ rapidly grows to become comparable in




Figure 17.9: Same as Fig. 17.8, but for increasing the value of Δm_{31}^2 to $2.5 \times 10^{-3} \text{ eV}^2$ with respect to the nominal model.

magnitude to r_e . The negative pulls are thus the residual signature of $\Delta P_{\mu\mu} < 0$.

Modulus of atmospheric mass-squared difference Δm_{31}^2

Increasing Δm_{31}^2 by almost 2% to $2.5 \times 10^{-3} \text{ eV}^2$ leaves a qualitatively similar imprint on cascade-like events as it does on track-like events, as evidenced by the pull distributions in Fig. 17.9. The overall scale of the significance is at the same level as the significance of observing maximal 2–3 mixing, but more sign changes are observed here because the parameter shift is equivalent to an increase of the oscillation frequency.

The strongest impact of Δm_{31}^2 , at the level of 10% of the expected statistical uncertainty in a given bin, appears at $E_{\text{reco}} \approx 15 \text{ GeV}$ and close to the vertically up-going region of the track-like event distribution. This primarily reflects the shift of the first survival maximum of $\bar{\nu}_\mu$'s to a smaller value of L/E_ν . The signature does not strengthen when moving away from $\cos \vartheta_{\text{reco}} = -1$ along fixed L/E_ν because the amplitude of the first survival maximum of ν_μ 's (but not $\bar{\nu}_\mu$'s) that only traverse the mantle is reduced with respect to those that cross the core, which weakens the sensitivity to Δm_{31}^2 at the $(P_{\mu\mu})$ probability level. The surrounding regions of negative pulls reflect the corresponding displacements of the first and second survival minima of both ν_μ 's and $\bar{\nu}_\mu$'s. The most significant deficit exists in the core region and for $E_{\text{reco}} \sim 7 \text{ GeV}$ to 9 GeV .

The cascade-like event distribution exhibits weaker pulls overall and lacks the sign flip for events with nearly vertical trajectories at energies $E_{\text{reco}} \gtrsim 25$ GeV. Instead, the band containing the first $\bar{\nu}_\mu$ survival maximum appears to be smeared out to higher energies. It is helpful once more to investigate the relative weights of the various contributions to Eq. (B.2), where we can neglect all antineutrino oscillation channels except for $\bar{\nu}_\mu \rightarrow \bar{\nu}_{\mu,\tau}$ a priori:

$$\begin{aligned} \Delta N(\Delta(\Delta m_{31}^2) > 0) \propto & \left(\frac{\Delta P_{ee}}{R_{\mu/e}^0} + \Delta P_{\mu e} \right) r_e + \left(\Delta P_{\mu\mu} + \frac{\Delta P_{\mu e}}{R_{\mu/e}^0} \right) r_\mu + \left(\frac{\Delta P_{e\tau}}{R_{\mu/e}^0} + \Delta P_{\mu\tau} \right) \frac{r_\tau}{a_{e/\tau}^0} \\ & + \frac{1}{R_{\nu_\mu/\bar{\nu}_\mu}^0 a_{\nu_e/\bar{\nu}_e}^0} \left[\Delta \bar{P}_{\mu\mu} \bar{r}_\mu + \frac{\Delta \bar{P}_{\mu\tau}}{\bar{a}_{e/\tau}^0} \bar{r}_\tau \right]. \end{aligned} \quad (17.2)$$

In the regime of L/E_ν between the position of the first $\bar{\nu}_\mu$ survival maximum and the first survival minimum ($E_\nu \sim 12$ GeV to 25 GeV for $\cos \vartheta = -1$), the excess of misidentified $\bar{\nu}_\mu$ CC events drives the net significance, since the $\bar{\nu}_e$ survival probabilities are close to unity and cancellations owing to $\Delta \bar{P}_{\mu\tau} < 0$ are weak due to the small $\bar{\nu}_\tau$ CC effective areas.

At even smaller L/E_ν for $\cos \vartheta_{\text{reco}} \lesssim -0.5$, below the first $\bar{\nu}_\mu$ survival minimum, the terms $\propto \Delta \bar{P}_{\mu\tau} > 0$ gain more weight as $\bar{A}_{\text{eff}}^\tau \rightarrow \bar{A}_{\text{eff}}^{e,\mu}$. Hence, the enhanced appearance of high-energy $\bar{\nu}_\tau$'s due to the increase in Δm_{31}^2 is the origin of the extended region of positive pulls in the cascade-like event distribution.

Neutrino mass ordering

The signature in Fig. 17.10 of simply switching the sign of Δm_{31}^2 in the nominal NO model is characterised by a pronounced deficit of cascade-like events—at the level of up to 25% of the expected statistical uncertainty—across an extended region roughly covering the mantle's and the core's MSW resonance peaks and the parametric ridges of the $\nu_e \rightarrow \nu'_3$ transition. A more significant excess of track-like events is observed along a band similar to the one seen in the right panel of Fig. 17.9, but with the largest significances occurring outside of the core region. Overall, the considered NMO flip affects track-like events in a similar manner as does raising Δm_{31}^2 within the NO.

Why does the NMO give rise to the above signature in PINGU? First of all, under the IO hypothesis matter effects only enhance the antineutrino equivalent of the transition $\nu_e \rightarrow \nu'_3$ in exactly the same manner as described in Sec. 5.1.2.2. Secondly, even though PINGU is unable to distinguish between neutrinos and antineutrinos




Figure 17.10: Same as Fig. 17.8, but for the simple flip of sign $\Delta m_{31}^2 > 0 \rightarrow -\Delta m_{31}^2$.

on an event-by-event basis, it has statistical discrimination power, governed by the deviations from unity of the neutrino-to-antineutrino flux and cross-section ratios. Neglecting Δm_{21}^2 -driven transitions implies that $\Delta P_{ee} = -\Delta \bar{P}_{ee}$, $\Delta P_{\mu e} = -\Delta \bar{P}_{\mu e}$, and $\Delta P_{e\tau} = -\Delta \bar{P}_{e\tau}$ in Eq. (B.2), but one has to in general take $\Delta P_{\mu\mu} \neq -\Delta \bar{P}_{\mu\mu}$ and $\Delta P_{\mu\tau} \neq -\Delta \bar{P}_{\mu\tau}$. This leads to the expression

$$\begin{aligned} & \Delta N(\Delta m_{31}^2 \rightarrow -\Delta m_{31}^2) \\ & \propto \left(\frac{\Delta P_{ee}}{R_{\mu/e}^0} + \Delta P_{\mu e} \right) r_e + \left(\Delta P_{\mu\mu} + \frac{\Delta P_{\mu e}}{R_{\mu/e}^0} \right) r_\mu + \left(\Delta P_{\mu\tau} + \frac{\Delta P_{e\tau}}{R_{\mu/e}^0} \right) r_\tau \\ & - \frac{1}{R_{\nu_\mu/\bar{\nu}_\mu}^0 a_{\nu_e/\bar{\nu}_e}^0} \left[\left(\frac{\Delta P_{ee}}{\bar{R}_{\mu/e}^0} + \Delta P_{\mu e} \right) \bar{r}_e + \left(-\Delta \bar{P}_{\mu\mu} + \frac{\Delta P_{\mu e}}{\bar{R}_{\mu/e}^0} \right) \bar{r}_\mu + \left(-\Delta \bar{P}_{\mu\tau} + \frac{\Delta P_{e\tau}}{\bar{R}_{\mu/e}^0} \right) \bar{r}_\tau \right]. \end{aligned} \quad (17.3)$$

Here, in general each neutrino term carries more weight than its (opposite-sign) antineutrino counterpart, because $\bar{r}_\alpha/r_\alpha < R_{\nu_\mu/\bar{\nu}_\mu}^0 a_{\nu_e/\bar{\nu}_e}^0$.

In the case of cascade-like events, the deficit in regions where matter effects induce strong transitions $\nu_e \leftrightarrow \nu_\mu$ given NO is driven primarily by the net negative contribution from the four terms containing $\Delta P_{\mu e} < 0$. The net-positive terms containing $\Delta P_{ee} > 0$ only cause a partial cancellation, while the net-negative ones containing $\Delta P_{e\tau} < 0$ are strongly suppressed at these energies. $\Delta P_{\mu\mu}$ and $\Delta P_{\mu\tau}$ are subject to rapid sign changes of similarly substantial sizes, which are severely averaged out

due to PINGU's limited resolutions. As an example, at the MSW resonance peak in the mantle our parameter assumptions result in $\Delta P_{ee} \sim 1$ and $\Delta P_{\mu e} \sim -0.5$. Given $R_{\mu/e}^0 \sim 2.5$, in order to determine the sign of ΔN even after event reconstruction it is sufficient to notice that the full survival of the intrinsic ν_e flux is more than compensated by the lacking appearance of ν_e 's from the intrinsic ν_μ flux. The excess of nearly vertically upgoing high-energy events is due to the enhanced transitions $\bar{\nu}_\mu \rightarrow \bar{\nu}_\tau$, i.e., $\Delta P_{\mu\tau} \approx \Delta \bar{P}_{\mu\tau} > 0$ (see below). That these transitions govern the overall signature in this particular region of the cascade-like event distribution is due to the combination of the following high-energy conditions: $\bar{a}_{e/\tau}^0 \rightarrow 1$, $\Delta P_{ee, e\mu} \rightarrow 0$, $R_{\mu/e}^0 \gg 1$, and $\bar{\nu}_\mu$ survival is irrelevant because $\bar{r}_\mu \ll \bar{r}_{e,\tau}$.

The more significant maximal pull of up to ~ 0.45 on track-like event counts stems from ν_μ survival proceeding as in vacuum under the IO, unsuppressed by matter effects. The strong enhancement occurs between the location of the first survival minimum and that of the first survival maximum and is also visible in cascade-like events. At higher levels of L/E_ν , the deficit of track-like events is the result of a minor shift of the $\bar{\nu}_\mu$ survival probability patterns to smaller values of L/E_ν .

17.4 Systematic uncertainties

This analysis takes into account the eight sources of systematic uncertainty listed in Table 17.1, which are implemented as nuisance parameters. Five of these are assigned a Gaussian prior constraint.

Atmospheric neutrino fluxes

$\nu/\bar{\nu}$ flux ratio The uncertainty on the relative production yields of atmospheric neutrinos and antineutrinos is realised through a universal rescaling parameter $R_{\nu/\bar{\nu}}$ that conserves the total atmospheric flux of neutrinos and antineutrinos, $\Phi_{\text{unosc}}^\alpha + \bar{\Phi}_{\text{unosc}}^\alpha$, for each flavour $\alpha = e, \mu$. With the nominal flux ratios $R_{\nu_\alpha/\bar{\nu}_\alpha}^0 \equiv \Phi_{\text{unosc}}^\alpha / \bar{\Phi}_{\text{unosc}}^\alpha$

the transformations are given by²

$$\Phi_{\text{unosoc}}^{\alpha} \rightarrow \Phi_{\text{unosoc}}^{\alpha'} = R_{\nu/\bar{\nu}} R_{\nu_{\alpha}/\bar{\nu}_{\alpha}}^0 \frac{\Phi_{\text{unosoc}}^{\alpha} + \bar{\Phi}_{\text{unosoc}}^{\alpha}}{1 + R_{\nu/\bar{\nu}} R_{\nu_{\alpha}/\bar{\nu}_{\alpha}}^0}, \quad (17.4)$$

$$\bar{\Phi}_{\text{unosoc}}^{\alpha} \rightarrow \bar{\Phi}_{\text{unosoc}}^{\alpha'} = \frac{\Phi_{\text{unosoc}}^{\alpha} + \bar{\Phi}_{\text{unosoc}}^{\alpha}}{1 + R_{\nu/\bar{\nu}} R_{\nu_{\alpha}/\bar{\nu}_{\alpha}}^0}, \quad (17.5)$$

$$\Rightarrow \frac{\Phi_{\text{unosoc}}^{\alpha'}}{\bar{\Phi}_{\text{unosoc}}^{\alpha'}} = R_{\nu/\bar{\nu}} R_{\nu_{\alpha}/\bar{\nu}_{\alpha}}^0. \quad (17.6)$$

$R_{\nu/\bar{\nu}}$ is subjected to a Gaussian constraint with a standard deviation of 10 %, centred around unity.

$\bar{\nu}_e/\bar{\nu}_{\mu}$ flux ratio The uncertainty on the ratio between atmospheric electron and muon (anti)neutrino fluxes is realised through a rescaling parameter $R_{e/\mu}$ as defined in Eqs. (12.14) to (12.16). $R_{e/\mu}$ is subjected to a Gaussian constraint with a standard deviation of 3 %, centred around unity.

$\bar{\nu}$ flux spectral index The energy dependencies of the atmospheric neutrino and antineutrino fluxes are allowed to vary via a shift $\Delta\gamma$ of the spectral index γ , as defined in Eqs. (12.17) to (12.18). A Gaussian constraint with a standard deviation of 0.05 is assumed, centred around zero.

Neutrino oscillations

Both Δm_{31}^2 and θ_{23} are considered as unconstrained nuisance parameters. θ_{13} is subjected to a Gaussian constraint with a standard deviation of 0.2° , centred around $\theta_{13} = 8.5^\circ$. The solar oscillation parameters θ_{12} and Δm_{21}^2 as well as δ_{CP} remain fixed at their nominal values specified in Table 17.1.

Detector response

Overall $\bar{\nu}$ event-count normalisation No constraint is applied to the overall normalisation of the number of neutrino and antineutrino events, which is scaled by the nuisance parameter $N_{\bar{\nu}}$ (cf. Sec. 12.5).

$\bar{\nu}$ energy scale A universal energy scaling factor s_E is introduced, which accounts for the possibility of a systematic under- or overestimation of the true (anti)neutrino

²Compare Eqs. (12.14) to (12.16).

Nuisance parameter	Nominal value \pm uncertainty	Fit range
<i>Free parameters:</i>		
$R_{e/\mu}$	1.00 ± 0.03	[0.70, 1.30]
$R_{\nu/\bar{\nu}}$	1.00 ± 0.10	[0.70, 1.30]
$\Delta\gamma$	0.00 ± 0.05	[-0.20, 0.20]
$N_{\bar{\nu}}$	1.0	[0.70, 1.30]
θ_{23} ($^\circ$)	42.3 (49.5)	[38.3, 53.3]
Δm_{31}^2 (10^{-3} eV 2)	2.46 (-2.37)	[2.10 (-2.70), 3.00 (-1.90)]
θ_{13} ($^\circ$)	8.50 ± 0.20	[7.85, 9.10]
s_E	1.00 ± 0.10	[0.70, 1.30]
<i>Fixed parameters (also tested):</i>		
θ_{12} ($^\circ$)	33.50	
Δm_{21}^2 (10^{-5} eV 2)	7.50	
δ_{CP} ($^\circ$)	0	

Table 17.1: Nuisance parameters investigated for inclusion in the PINGU sensitivity studies, together with their nominal values and (if applicable) prior uncertainties and fit ranges. The upper eight parameters are allowed to vary, whereas the three parameters at the bottom are kept fixed at the indicated values.

energy. s_E is subjected to a Gaussian constraint with a standard deviation of 10 %, centred around unity, which intends to capture the uncertainty on the overall DOM efficiency.

18

Statistical techniques

While not the core physics goal at PINGU’s conception, the NMO came under scrutiny [477, 478, 44, 46] when the first indications for a non-zero value of θ_{13} appeared, in the T2K and MINOS neutrino accelerator experiments (see Ref. [479] and references therein). The study of PINGU’s NMO sensitivity lends itself to a variety of frequentist analysis approaches which exhibit distinct challenges and tradeoffs examined in this chapter.

18.1 Frequentist hypothesis test: NO or IO?

For all practical intents and purposes, the NMO measurement in PINGU represents a binary hypothesis-testing problem, consisting of the two mutually exclusive hypotheses of the NO and the IO, as defined in Sec. 3.2.4. Following common conventions, our test-statistic choices are guided by the Neyman–Pearson lemma [480]. It states that for two *simple* hypotheses, the likelihood ratio λ yields the “best” critical region for the null hypothesis (\mathcal{H}_0) with regard to its alternative (\mathcal{H}_1). That is, for each choice of critical value, corresponding to the probability α of *rejecting* \mathcal{H}_0 when it is in fact true (“type-I error”), the use of the likelihood ratio results in the smallest probability β of *accepting* \mathcal{H}_0 when in fact the alternative \mathcal{H}_1 is true (“type-II error”). The complementary probability $1 - \beta$ is referred to as the “power” of the hypothesis test.

In this work we produce pseudoexperiments and explicitly generate distributions of the test statistic $-2 \ln \lambda$ for the two NMO hypotheses (cf. Sec. 13.3). In addition, merely by generating Asimov templates, the distributions of a closely related “ $\Delta\chi^2$ ” test statistic, a difference between two weighted-least-squares statistics, are derived.

18.1.1 Sensitivity definition

It is crucial to evaluate how a chosen test statistic can be translated into an NMO sensitivity proxy, with reasonable choices for the sizes of the type-I and type-II errors. This problem is scrutinised in detail in Ref. [481]. Its findings are at the core of the the statistical methodology underlying the frequentist NMO inference approaches adopted in this work. In particular, we seek the confidence level $1 - \alpha$ at which the “wrong ordering” (WO, the null hypothesis) has a probability of 50% of being rejected (or accepted). The outcome is interpreted as the median sensitivity.

We take the test statistic, \mathcal{T} , to be governed by the PDF $f(\mathcal{T}|\text{IO})$ under the IO and by $f(\mathcal{T}|\text{NO})$ under the NO. Choosing the IO as our null hypothesis for definiteness and assuming we reject it if \mathcal{T} is exceedingly small, the type-I error probability is defined as

$$\alpha = \int_{-\infty}^{\mathcal{T}_\alpha} f(\mathcal{T}|\text{IO}) . \quad (18.1)$$

It is determined by equating the critical value \mathcal{T}_α with the median of the distribution of \mathcal{T} given the true ordering (TO, in this case the NO), or

$$\beta = \int_{-\infty}^{\mathcal{T}_\alpha} f(\mathcal{T}|\text{NO}) = 0.5 = \int_{\mathcal{T}_\alpha}^{\infty} f(\mathcal{T}|\text{NO}) , \quad (18.2)$$

where β is the type-II error. Hence, we define PINGU’s median NMO sensitivity as $\alpha(\beta = 0.5)$. It is converted into a corresponding number of standard deviations, the significance n_σ , by assuming a two-sided Gaussian test [481],

$$n_\sigma(\alpha) = \sqrt{2} \operatorname{erfc}^{-1}(\alpha) , \quad (18.3)$$

where erfc^{-1} is the inverse of the complementary error function. This choice implies that $n_\sigma = 1$ corresponds to a CL of $1 - \alpha = 68.27\%$ for example, and $n_\sigma = 1.65$ to the 90% CL.¹ Below, the median test-statistic outcome ($\beta = 0.5$) is assumed except where otherwise stated. Accordingly, when given without an argument, n_σ refers to the median significance $n_\sigma(\beta = 0.5)$.

¹The two-sided relationship between the confidence level and the number of standard deviations is a common convention in the field and agrees with that in Part III, but deviates from the one-sided choice in Ref. [427], for example. The two-sided choice results in a non-zero median significance of $n_\sigma \approx 0.48$ when the distributions under the two orderings coincide.

18.1.2 Challenges of testing discrete composite hypotheses

The discussion in the preceding paragraphs does not explicitly account for the fact that both hypotheses under consideration are *composite*. Since they depend on free parameters, the definitions (18.2) and (18.3) do not uniquely define the NMO inference approach.

When the sensitivity associated with given sizes of the type-II error β is sought, the critical value \mathcal{T}_α —which decides whether \mathcal{H}_0 (WO) is rejected—is determined uniquely by the assumed true values of the (nuisance) parameters within \mathcal{H}_1 (TO). Consequently, without reference to a specific point p_{WO} in the nuisance-parameter space of the WO, it is understood that all possible realisations of the latter should be excluded at the stated confidence level [481]. Theoretically, this would require finding the maximum type-I error $\alpha_{\text{WO,max}} \equiv \max_{p_{\text{WO}}} \alpha(p_{\text{WO}})$ for the selected value of β by generating all test-statistic distributions allowed under the WO via pseudoexperiments. Only this Feldman–Cousins-like approach ensures the correct coverage of the stated confidence level.

For common test-statistic choices (see Ref. [481] and sections below), the actual measurement would yield a single observed value \mathcal{T}_{obs} . The knowledge of the test-statistic distributions over the parameter spaces of both NMO hypotheses would then allow looking up the maximum type-I error $\alpha_{\mathcal{H},\text{max}}$ for each. In turn, confidence levels $1 - \alpha_{\text{NO,max}}$ and $1 - \alpha_{\text{IO,max}}$ for the exclusion of either would be obtained. Alternatively, given any predefined CL $1 - \alpha$, the measurement has four distinct outcomes: 1) both NMO hypotheses are rejected, 2) the IO is rejected but the NO is accepted, 3) the NO is rejected but the IO is accepted, and 4) both are accepted.

Existing neutrino oscillation measurements of the NMO typically make use of approximate coverage studies, due to the computational cost prohibiting the full procedure laid out above. For example, in the case of a recent proof-of-concept measurement by DeepCore [427], only two sets of test-statistic outcomes of pseudoexperiments (“trials”) are compared to the observation to derive associated type-I errors. These trials are generated by injecting the observed best-fit point within each NMO hypothesis. The result is shown in Fig. 18.1 for one of the two DeepCore analyses that have been performed.² Here, the hypothesis test is defined such that the NO is rejected for exceedingly large test-statistic values, whereas the IO is rejected for exceedingly small test-statistic values, resulting in the observed type-I error probability of 11.4 % for true NO and 84.5 % for true IO. This pair is converted into a pair of ra-

²Similar to the NSI measurement in Part III, “Analysis B” makes use of a slightly modified version of the event sample of the same name in Ref. [362].




Figure 18.1: DeepCore NMO measurement outcome for “Analysis B”, taken from Ref. [427]. Shown are the trial histograms of the test statistic $\Delta\chi^2$ (see Sec. 18.1.3) given the observed best fits, together with its measured value (“data”). See text for details.

tios, referred to as “ CL_S ” values, that can be constructed from the observed exclusion confidence level for one NMO hypothesis and the observed type-I error probability for its alternative, thereby removing the directionality from the hypothesis test. If CL_S equals unity, no NMO preference is observed; if CL_S vanishes, there is a maximal preference for one NMO over the alternative. However, unless the distributions are independent of the true values of the nuisance parameters, this procedure is not guaranteed to minimise the confidence level for excluding either NMO.

A study of the above issues is performed in this analysis using PINGU MC events and pseudoexperiments alone. For computational feasibility and to adhere to convention (see, e.g., Refs. [46, 1, 482, 427]), our baseline approach adopts the best WO fit to the Asimov experiment under a given TO model p_{TO} as the null hypothesis. In searching for this hypothesis, only the two impactful oscillation parameters θ_{23} and Δm_{31}^2 are allowed to vary—within the WO half-space, but otherwise in an unconstrained manner—while the other nuisance parameters remain at their nominal WO values. The resulting model, denoted as p_{test} , serves as a proxy of the model $p_{\text{WO},\alpha_{\text{max}}}$ which gives rise to the maximum type-I error probability $\alpha_{\text{WO},\text{max}}$ in the rigorous but computationally prohibitive analysis.

18.1.3 Sensitivity from Asimov templates

The goal of the Asimov ansatz is to quantify the NMO sensitivity for arbitrary type-II error probabilities without having to resort to fits to pseudoexperiments, in close analogy to Chapter 14. To this end, we assume that the measurement will make use of the weighted sum of squared residuals

$$\chi^2(\mathbf{p}_{\mathcal{H}}) \equiv \sum_{i=1}^{N_{\text{bins}}} \frac{[n_{\text{obs},i} - n_{\text{exp},i}(\mathbf{p}_{\mathcal{H}})]^2}{n_{\text{exp},i}(\mathbf{p}_{\mathcal{H}})} + \sum_{j=1}^{N_{\text{prior}}} \frac{(\Delta p_{\mathcal{H},j})^2}{\sigma_{p_{\mathcal{H},j}}^2}, \quad (18.4)$$

where $\mathbf{p}_{\mathcal{H}}$ are free parameters within $\mathcal{H} = \text{NO}, \text{IO}$, $n_{\text{obs},i}$ the observed event count in bin i , and $n_{\text{exp},i}(\mathbf{p}_{\mathcal{H}})$ the expectation due to $\bar{\nu}$ events evaluated at the point $\mathbf{p}_{\mathcal{H}}$, retrieved according to Chapter 17. The expectation's squared residual with respect to the observed event count in bin i is divided by the Poisson variance. The second sum contributing to Eq. (18.4) is taken over all N_{prior} parameters subject to external Gaussian constraints, in analogy to Eq. (13.1).

Optimising $\mathbf{p}_{\mathcal{H}}$ to best match the observed event distribution yields the two minima

$$\chi_{\text{NO}}^2 \equiv \chi^2(\hat{\mathbf{p}}_{\text{NO}}) \text{ and } \chi_{\text{IO}}^2 \equiv \chi^2(\hat{\mathbf{p}}_{\text{IO}}), \quad (18.5)$$

where $\hat{\mathbf{p}}_{\mathcal{H}}$ are the nuisance-parameter values at the conditional χ^2 minimum (given \mathcal{H}). The two hypotheses are distinguished by the half-space to which $\Delta m_{31}^2 \subset \mathbf{p}_{\mathcal{H}}$ is restricted during the minimisation of Eq. (18.4).

Finally, the two statistics of Eq. (18.5) are combined into a single suitable statistic for the NMO hypothesis test:

$$\Delta\chi^2 \equiv \chi_{\text{NO}}^2 - \chi_{\text{IO}}^2. \quad (18.6)$$

The sign of $\Delta\chi^2$ indicates which hypothesis constitutes a better match to the observed event distribution, with $\Delta\chi^2 < 0$ ($\Delta\chi^2 > 0$) in the case of the NO (IO). We require the distributions of $\Delta\chi^2(\mathbf{p}_{\text{TO}})$ and $\Delta\chi^2(\mathbf{p}_{\text{test}})$ in order to quantify PINGU's NMO sensitivity given \mathbf{p}_{TO} in nature.

In the Asimov approach $n_{\text{obs},i}$ is set to its expectation value $n_{\text{exp},i}(\mathbf{p}_{\text{inj}})$ given the injected model $\mathbf{p}_{\text{inj}} = \mathbf{p}_{\text{TO}}, \mathbf{p}_{\text{test}}$. When the latter is centred with respect to all prior constraints, the conditional minimum given \mathcal{H}_{inj} must be located at $\hat{\mathbf{p}}_{\mathcal{H}_{\text{inj}}} = \mathbf{p}_{\text{inj}}$, where the Asimov-equivalent of Eq. (18.4) identically vanishes, $\overline{\chi_{\mathcal{H}_{\text{inj}}}^2} \equiv \overline{\chi^2(\hat{\mathbf{p}}_{\mathcal{H}_{\text{inj}}})} = 0$. Moving forward, we only add the *injected* hypothesis \mathcal{H}_{inj} as a subscript to the Asimov-

equivalent of the $\Delta\chi^2$ statistic (18.6):

$$\overline{\Delta\chi^2_{\mathcal{H}_{\text{inj}}}} \equiv \begin{cases} -\overline{\chi^2_{\text{IO}}}, & \mathcal{H}_{\text{inj}} = \text{NO} \\ \chi^2_{\text{NO}}, & \mathcal{H}_{\text{inj}} = \text{IO} \end{cases}. \quad (18.7)$$

Whether the injected hypothesis is assumed to correspond to some truth realised in nature or represents the WO hypothesis will always be apparent from the context.

Since we are dealing with non-nested hypotheses, there is no a-priori justification for applying Wilks' theorem to Eq. (18.6). However, under certain regularity conditions $\Delta\chi^2$ is distributed as [481]

$$f(\Delta\chi^2|\mathcal{H}_{\text{inj}}) \sim n\left(\overline{\Delta\chi^2_{\mathcal{H}_{\text{inj}}}}, 2\sqrt{|\overline{\Delta\chi^2_{\mathcal{H}_{\text{inj}}}}|}\right), \quad (18.8)$$

where $n(\mu, \sigma)$ is the Gaussian distribution. Under such conditions the type-I respectively type-II error is given by

$$\alpha = \frac{1}{2} \operatorname{erfc}\left(\frac{\pm\overline{\Delta\chi^2_{\text{WO}}} \mp \overline{\Delta\chi^2_{\alpha}}}{\sqrt{8|\overline{\Delta\chi^2_{\text{WO}}}|}}\right), \quad (18.9)$$

$$\beta = \frac{1}{2} \operatorname{erfc}\left(\frac{\pm\overline{\Delta\chi^2_{\alpha}} \mp \overline{\Delta\chi^2_{\text{TO}}}}{\sqrt{8|\overline{\Delta\chi^2_{\text{TO}}}|}}\right), \quad (18.10)$$

where the upper (lower) signs apply when the NO (IO) is assumed to be true.

The case $\beta = 0.5$ requires the critical value to coincide with the TO distribution mean, $\overline{\Delta\chi^2_{\alpha}} = \overline{\Delta\chi^2_{\text{TO}}}$. The corresponding median significance, expressed through the equivalent number of two-sided Gaussian standard deviations (18.3), becomes

$$n_{\sigma}(\beta = 0.5) = \sqrt{2} \operatorname{erfc}^{-1}\left[\frac{1}{2} \operatorname{erfc}\left(\frac{\pm\overline{\Delta\chi^2_{\text{WO}}} \mp \overline{\Delta\chi^2_{\text{TO}}}}{\sqrt{8|\overline{\Delta\chi^2_{\text{WO}}}|}}\right)\right]. \quad (18.11)$$

The values it can take for different combinations of $\overline{\Delta\chi^2_{\text{WO}}}$ and $\overline{\Delta\chi^2_{\text{TO}}}$ are provided in Fig. 18.2, which shows different fixed choices of $|\overline{\Delta\chi^2_{\text{WO}}}|$ (opposite signs of the




Figure 18.2: Possible values (shaded region) and minimum (thick solid line) of the Asimov median NMO significance (18.11) in dependence of the TO test-statistic mean’s absolute value, compared to the latter’s square root (dashed line)—a commonly used NMO significance proxy. Thin solid lines represent different choices of the WO test-statistic mean, with larger moduli corresponding to smaller slopes. The WO test-statistic mean that minimises the median significance (maximises the median type-I error probability) can be read off from the secondary vertical axis (dotted line).

two means are implied). Without any constraint on the latter, for a given value of $\overline{\Delta\chi_{\text{TO}}^2}$ the median significance has a lower but no upper bound. The figure also demonstrates that the lower bound is realised by the symmetric configuration $\overline{\Delta\chi_{\text{WO}}^2} = -\overline{\Delta\chi_{\text{TO}}^2}$. In an actual sensitivity study, without sampling the allowed range of $\Delta\chi_{\text{WO}}^2$ across the (reasonably constrained) WO parameter space—from which the null hypothesis emerges—quoting the lower bound could amount to underestimating the confidence level at which the WO can be excluded. Similar considerations apply to the quantity $\sqrt{|\overline{\Delta\chi_{\text{TO}}^2}|}$, which represents an NMO significance proxy that is frequently encountered in the literature and that approaches the lower bound of n_σ as $|\overline{\Delta\chi_{\text{TO}}^2}|$ grows. Its interpretation is elucidated in Sec. 18.1.3.1 below.

Note that critical values $\Delta\chi_\alpha^2 \neq \overline{\Delta\chi_{\text{TO}}^2}$ correspond to different probabilities of being able to exclude the WO at higher or lower confidence levels $1 - \alpha$ than in the median

case. The central $m\sigma$ band of confidence levels due to statistical fluctuations is obtained by varying $\Delta\chi_\alpha^2$ over the range $\overline{\Delta\chi_{\text{TO}}^2} \pm 2m\sqrt{|\Delta\chi_{\text{TO}}^2|}$. In general, increasing the probability $1 - \beta$ of rejecting the WO by suitably choosing $\Delta\chi_\alpha^2$ comes at the cost of a reduced confidence level of the NMO measurement [481, Fig. 2].

Similar to Sec. 14.3.1, a numerical minimisation routine serves as the reference method for obtaining the Asimov metrics (18.7). Here, the L-BFGS-B algorithm [453] is seeded with points in both octants of θ_{23} and the best outcome is recorded. The combination of this minimisation technique with the statistical assumptions above is referred to as the “(Asimov) $\overline{\Delta\chi^2}$ ” approach in the remainder of this chapter. Two closely related analysis methods are investigated in addition: by imposing more restrictive regularity conditions, the need for numerical minimisation is obviated. As a result, computational costs are further reduced because fewer templates have to be produced in order to arrive at an NMO significance projection—albeit at the cost of its accuracy.

18.1.3.1 Fisher information

Given the likelihood function $L(\mathbf{p})$ for a sample of random variables, the observed “Fisher information” is a measure of the amount of information about \mathbf{p} contained in the sample. In the univariate case, the Fisher information is identified with the negative of the second derivative of the log-likelihood with respect to the parameter p , and in practice is typically evaluated at the maximum-likelihood estimate \hat{p} [443, 483]. The metric’s multivariate equivalent is the observed Fisher information matrix (“Fisher matrix” for short), whose elements are the second partial derivatives (see, e.g., Ref. [484] and references therein)

$$\mathcal{F}_{ij}(\mathbf{p}) = -\frac{\partial^2}{\partial p_i \partial p_j} \ln L(\mathbf{p}) . \quad (18.12)$$

The Fisher matrix evaluated at $\hat{\mathbf{p}}$ is commonly interpreted as the inverse of the covariance matrix of the maximum-likelihood estimators of all parameters:

$$\mathcal{F}^{-1}(\hat{\mathbf{p}}) \equiv \mathcal{C}(\hat{\mathbf{p}}) . \quad (18.13)$$

The validity of this relation is discussed at the bottom of this section.

In the case of the Asimov experiment there is no distinction between the observed

Fisher matrix and its expectation [452]. Substituting the Poisson likelihood³ into Eq. (18.12) yields [452, 45]

$$\mathcal{F}_{ij}(\mathbf{p}) = \sum_{k=1}^{N_{\text{bins}}} \frac{1}{n_{\text{exp},k}(\mathbf{p})} \frac{\partial n_{\text{exp},k}(\mathbf{p})}{\partial p_i} \frac{\partial n_{\text{exp},k}(\mathbf{p})}{\partial p_j}, \quad (18.14)$$

where the sum is taken over all template bins. The covariance matrix of parameter estimators at some injected point \mathbf{p}_{inj} , $C(\mathbf{p}_{\text{inj}})$, allows inferring confidence regions for these parameters under the assumption of \mathbf{p}_{inj} being true ($\hat{\mathbf{p}} = \mathbf{p}_{\text{inj}}$ by definition). A prior Gaussian uncertainty σ_{p_j} centred around $p_{\text{inj},j}$ is accounted for by adding the inverse variance $\sigma_{p_j}^{-2}$ to the diagonal element \mathcal{F}_{jj} [45].

By employing finite differences for the partial derivatives at \mathbf{p}_{inj} , in principle only two Asimov templates need to be generated per free parameter p_i in order to obtain the Fisher matrix and its inverse. However, the question arises how this method can be reconciled with the inference approach to the discrete NMO hypothesis test. To this end, Ref. [45] introduces a continuous parameter $0 \leq h \leq 1$ which interpolates between the two orderings as

$$n_{\text{exp},k}(h; \mathbf{p}_{\text{NO}}, \mathbf{p}_{\text{IO}}) \equiv h n_{\text{exp},k}(\mathbf{p}_{\text{NO}}) + (1-h) n_{\text{exp},k}(\mathbf{p}_{\text{IO}}), \quad (18.15)$$

assuming $h_{\text{IO}} = 0$ and $h_{\text{NO}} = 1$ for definiteness and without loss of generality. The parameter implies the following prescription for computing PINGU's NMO sensitivity:

1. Let $\mathbf{p}_{\text{inj}} = \mathbf{p}_{\text{WO}}$ represent the desired WO hypothesis (referred to as the "fiducial model" in Ref. [45]). This is \mathbf{p}_{test} in this work's baseline method.
2. Evaluate the partial derivatives with respect to h of the expected bin counts (18.15) as

$$\frac{\partial n_{\text{exp},k}}{\partial h} = \pm n_{\text{exp},k}(\mathbf{p}_{\text{WO}}) \mp n_{\text{exp},k}(\mathbf{p}_{\text{TO}}), \quad (18.16)$$

where the upper (lower) signs apply for true IO (NO).

3. Evaluate the partial derivatives with respect to $p_i \neq h$, keeping the remaining

³Discussed later in the context of ensemble studies, see Eq. (18.33).

parameters $p_{j \neq i}$ as well as h fixed at their injected values,

$$\left. \frac{\partial n_{\text{exp},k}}{\partial p_i} \right|_{p_{\text{WO}}} \approx \left. \frac{n_{\text{exp},k}(p_{\text{WO},i} + \epsilon) - n_{\text{exp},k}(p_{\text{WO},i})}{\epsilon} \right|_{p_{j \neq i} = p_{\text{WO},j}}. \quad (18.17)$$

4. Calculate the covariance matrix using Eqs. (18.14) and (18.13) to find $\sigma_h \equiv \sqrt{C_{hh}}$.

5. Calculate the type-I error probability α by assuming that h is Gaussian-distributed under the WO hypothesis,

$$f(h|\text{WO}) = \mathcal{N}(h_{\text{WO}}, \sigma_h), \quad (18.18)$$

and that h_{TO} is the median of the distribution of h given the TO hypothesis,

$$\int_{-\infty}^{h_{\text{TO}}} f(h|\text{TO}) = 0.5. \quad (18.19)$$

If one converts this type-I error probability into the equivalent number of *one-sided* Gaussian standard deviations, Eq. (18.19) gives rise to the median NMO significance [45]

$$n_{\sigma, \text{one-sided}} = \frac{1}{\sigma_h} = \frac{1}{\sqrt{(\mathcal{F}^{-1})_{hh}}}, \quad (18.20)$$

which was chosen as the baseline sensitivity proxy in the original PINGU LoI [46]. However, for later reference, σ_h can be also converted into the two-sided median significance

$$n_{\sigma} = \sqrt{2} \operatorname{erfc}^{-1} \left[\frac{1}{2} \operatorname{erfc} \left(\frac{1}{\sqrt{2} \sigma_h} \right) \right]. \quad (18.21)$$

Note that any dependence of the significance on the truth p_{TO} originates from Eq. (18.16) in the second step above.

An exact equivalence to the significance proxies based on the Gaussian $\Delta\chi^2$ test statistic (18.8) is encountered in the absence of nuisance parameters and when the means of the two $\Delta\chi^2$ distributions satisfy $\Delta\chi_{\text{TO}}^2 = -\Delta\chi_{\text{WO}}^2$, that is, when the symmetric Gaussian configuration is realised, which bounds the NMO sensitivity from below. On the one hand, the former assumption in the Fisher matrix approach gives

rise to the one-sided median significance [45]

$$n_{\sigma, \text{one-sided}} = \sqrt{\mathcal{F}_{hh}} = \sqrt{\sum_{k=1}^{N_{\text{bins}}} \frac{[n_{\text{exp},k}(\mathbf{p}_{\text{TO}}) - n_{\text{exp},k}(\mathbf{p}_{\text{WO}})]^2}{n_{\text{exp},k}(\mathbf{p}_{\text{WO}})}}. \quad (18.22)$$

and the two-sided median significance

$$n_{\sigma} = \sqrt{2} \operatorname{erfc}^{-1} \left[\frac{1}{2} \operatorname{erfc} \left(\sqrt{\frac{\mathcal{F}_{hh}}{2}} \right) \right]. \quad (18.23)$$

On the other hand, with or without nuisance parameters in the $\overline{\Delta\chi^2}$ approach under the symmetric configuration one has

$$n_{\sigma, \text{one-sided}} = \sqrt{|\overline{\Delta\chi_{\text{TO}}^2}|} \quad (18.24)$$

and

$$n_{\sigma} = \sqrt{2} \operatorname{erfc}^{-1} \left[\frac{1}{2} \operatorname{erfc} \left(\sqrt{\frac{|\overline{\Delta\chi_{\text{TO}}^2}|}{2}} \right) \right]. \quad (18.25)$$

The equivalence of the two approaches then follows from the equality $\mathcal{F}_{hh} = |\overline{\Delta\chi_{\text{TO}}^2}|$.

Validity of Fisher matrix approach

The strict equivalence (18.13) between the covariance matrix of all parameter estimators and the inverse of the Fisher matrix is an asymptotic formula which is contingent upon a variety of regularity conditions. Key pillars of its derivation are the multivariate central limit theorem and the law of large numbers, whether the independent data samples are identically [485] distributed or not [484]. By expanding the gradient of the log-likelihood function into a Taylor series around the injected point \mathbf{p}_{inj} and neglecting terms of higher than second order, their application allows demonstrating the convergence in distribution of the maximum-likelihood estimator to

$$\hat{\mathbf{p}} \sim \mathcal{N}[\mathbf{p}_{\text{inj}}, \mathcal{F}^{-1}(\hat{\mathbf{p}})], \quad (18.26)$$

where no distinction is made between the observed and the expected Fisher information due to their asymptotic convergence in probability. In addition, the likelihood (log-likelihood) function itself becomes Gaussian (hyperparabolic) [15]. Equ-




Figure 18.3: 1σ , 90%, and 3σ error ellipses from two variants of the Fisher matrix approach (thick dashed and solid lines) compared to the corresponding isocontours of $\overline{\chi^2}$ surfaces in PINGU toy NMO sensitivity studies with a single nuisance parameter (overall normalisation on the left, θ_{23} on the right). See text for details.

tion (18.13) also always holds in the special case of Gaussian samples whose means are linear in all model parameters [15].

In the NMO sensitivity analysis, the choice of the injected model \mathbf{p}_{inj} for the Fisher matrix entails the assumption that the distribution of the maximum-likelihood estimator for pseudoexperiments generated under the WO is approximately given by Eq. (18.26), where \mathbf{p}_{inj} is the WO hypothesis which maximises the type-I error for a fixed TO model, cf. Eqs. (18.18) and (18.19). In our baseline approach, it is the same numerically determined function $\mathbf{p}_{\text{inj}} \equiv \mathbf{p}_{\text{test}}(\mathbf{p}_{\text{TO}})$ as in the case of the $\overline{\Delta\chi^2}$ method.

The suitability of the Fisher matrix formalism for determining the NMO sensitivity of an earlier PINGU baseline geometry has been scrutinised with empirical methods in Ref. [45] for one fixed pair of injected models with identical values of $\theta_{23} \approx 39^\circ$ far in the first octant and assuming true IO. However, depending on the choice of \mathbf{p}_{inj} , the octant degeneracy in PINGU's θ_{23} sensitivity, partially discernible for example in Ref. [45, Fig. B.3b], will necessarily invalidate constraints derived from the Fisher matrix at least at sufficiently high confidence levels.

Substantiating the above assertion, Fig. 18.3 presents outcomes of two bivariate toy PINGU NMO studies with the Fisher matrix approach (for true NO), whose error ellipses of 1σ , 90%, and 3σ CL are contrasted with the corresponding isocontours $\overline{\chi^2} \approx \chi_{2,0.32}^2, \chi_{2,0.1}^2, \chi_{2,0.003}^2$. The sole nuisance parameter in the panel on the left is

the overall normalisation, $N_{\bar{\nu}}$, and in the one on the right θ_{23} . Two variants of the Fisher matrix approach are shown: in the first (“finite diffs.”), the partial derivatives (18.17) are determined with the finite-difference method, yielding the average rates of change of the expected event counts across the full fit range of the nuisance parameter (cf. Table 17.1); in the second (“polyn.”), the count expectations are calculated at 20 equally spaced values across the parameters’ full range and fit with second-order polynomials whose first derivatives are evaluated analytically at the injected value. Thus, the number of templates generated by the second variant to obtain count derivatives with respect to nuisance parameters is higher by a factor of ten. Disagreeing confidence regions between the two variants are evidence of count non-linearities in the nuisance parameter under study (at the scale of the fit range). In contrast to an actual NMO sensitivity analysis, in order to be able to verify the accuracy of the Fisher matrix approach by means of a comparison to $\overline{\chi^2}$ isocontours, it is necessary to inject the same nuisance-parameter value in the TO and the WO (here: $N_{\bar{\nu}} = 1$ respectively $\theta_{23} = 49.5^\circ$), or, if we considered $\overline{\Delta m_{31}^2}$ as nuisance parameter, the same absolute value. The pseudoexperiment of the $\overline{\chi^2}$ scan is the Asimov experiment under the injected WO model of the Fisher matrix approach. By definition, the location of the scan’s global minimum (here: $\overline{\chi_{\text{IO}}^2} = 0$) coincides with that model, whereas its minimum within the other NMO (here: $\overline{\chi_{\text{NO}}^2} > 0$)—the TO of the Fisher matrix approach—may be displaced with respect to the truth. Such a displacement is encountered in both toy studies, with a nearly 1% larger $N_{\bar{\nu}}$ respectively a first-octant θ_{23} best fit within the NO. The accuracy of the constraints from the Fisher matrix breaks down at the confidence level (here: 3σ) at which the octant degeneracy appears. This is the result of $\overline{\chi^2}$ deviating from a parabolic behaviour across θ_{23} octants. Contrariwise, the explicitly linear model comprising \bar{h} and $N_{\bar{\nu}}$ does not exhibit any internal inconsistencies, either between the $\overline{\chi^2}$ scan and the Fisher matrix approach or between the two implementation variants adopted for the latter.

The applicability to the NMO problem of the Fisher matrix approach in combination with the continuous parameter \bar{h} is not re-evaluated in detail in this work. Instead, certain sensitivities derived from it are mentioned in passing in the context of the results from the other, less approximate, methods discussed throughout this part. In particular, Sec. 18.1.3.2 below presents a modified statistical inference approach which makes use of the Fisher information matrix but mitigates some of its mentioned deficiencies.

18.1.3.2 Semianalytic minimisation: pull approach

The “pull approach” [486] explicitly describes the effects of model parameters on observables by linear error propagation. A change ξ_k in each of the n parameters p_k is propagated to a corresponding change $\propto \xi_k$ in the count expectation $n_{\text{exp},i}$ through the latter’s first-order Taylor expansion around the injected parameter value $p_{\text{inj},k}$ (see, e.g., Ref. [487]):

$$\tilde{p}_k \equiv p_{\text{inj},k} + \xi_k \implies n_{\text{exp},i}(\tilde{\mathbf{p}}) = n_{\text{exp},i}(\mathbf{p}_{\text{inj}}) + \underbrace{\sum_{k=1}^n \xi_k \left. \frac{\partial n_{\text{exp},i}(\mathbf{p})}{\partial p_k} \right|_{p_k=p_{\text{inj},k}}}_{\equiv \partial_k n_{\text{exp},i}}. \quad (18.27)$$

Substituting this expansion into Eq. (18.4) results in the weighted sum of squared residuals of the pull approach,

$$\chi_{\text{pull}}^2(\tilde{\mathbf{p}}) \equiv \sum_{i=1}^{N_{\text{bins}}} \frac{[n_{\text{obs},i} - n_{\text{exp},i}(\mathbf{p}_{\text{inj}}) - \sum_k \xi_k \partial_k n_{\text{exp},i}]^2}{n_{\text{exp},i}(\mathbf{p}_{\text{inj}}) + \sum_k \xi_k \partial_k n_{\text{exp},i}} + \sum_{j=1}^{N_{\text{prior}}} \frac{\xi_j^2}{\sigma_{p_j}^2}, \quad (18.28)$$

assuming that the injected value of each parameter with a prior Gaussian constraint coincides with the prior’s preferred value. Crucially, the assumption of linear parameter dependencies of the event counts allows us to minimise χ_{pull}^2 semianalytically (the partial derivatives in Eq. (18.27) cannot in general be found analytically), by simultaneously solving for all the parameter variations $\hat{\xi}_k$ that minimise χ_{pull}^2 :

$$\chi_{\text{pull},\text{min}}^2 \equiv \min_{\{\xi_k\}} \chi_{\text{pull}}^2 = \chi_{\text{pull}}^2(\hat{\xi}). \quad (18.29)$$

The best-fit parameter values, encountered at the minimum $\chi_{\text{pull},\text{min}}^2$, are given by

$$\hat{p}_k = p_{\text{inj},k} + \hat{\xi}_k. \quad (18.30)$$

For the purpose of extending the Fisher matrix approach, an expression for $\hat{\xi}_k$ has been derived which makes use of the partial derivatives required by the computation of the Fisher matrix and of its elements themselves. The derivation, whose steps are provided in Appendix D.1, yields

$$\hat{\xi}_k = \sum_l (\mathcal{F}^{-1})_{kl} \sum_{i=1}^{N_{\text{bins}}} \frac{[n_{\text{obs},i} - n_{\text{exp},i}(\mathbf{p}_{\text{inj}})] \partial_l n_{\text{exp},i}}{n_{\text{exp},i}(\mathbf{p}_{\text{inj}})}. \quad (18.31)$$

We refer to the optimal parameter variations $\hat{\xi}_k$ of the linearised model as the parameters' "pulls" even though these cannot be assumed to represent normally distributed random variables (not all of our parameters are subject to external constraints, in contrast to Ref. [486]). Any count non-linearities (and numerical errors) will appear and can be quantified through an overestimate of the unknown true minimum $\chi^2_{\min} \equiv \min_{\{p_k\}} \chi^2 \equiv \chi^2(\mathbf{p}_{\text{opt}})$, with a contribution $\Delta_{\text{pull,stat}}$ from counting statistics and a contribution $\Delta_{\text{pull,prior}}$ from priors, whose sum exceeds zero:

$$\begin{aligned} \Delta_{\text{pull}} &\equiv \chi^2(\mathbf{p}_{\text{inj}} + \hat{\xi}) - \chi^2_{\min} && (18.32) \\ &= \underbrace{\sum_{i=1}^{N_{\text{bins}}} \frac{[n_{\text{obs},i} - n_{\text{exp},i}(\mathbf{p}_{\text{inj}} + \hat{\xi})]^2}{n_{\text{exp},i}(\mathbf{p}_{\text{inj}} + \hat{\xi})} - \frac{[n_{\text{obs},i} - n_{\text{exp},i}(\mathbf{p}_{\text{opt}})]^2}{n_{\text{exp},i}(\mathbf{p}_{\text{opt}})}}_{\equiv \Delta_{\text{pull,stat}}} + \underbrace{\sum_{j=1}^{N_{\text{prior}}} \frac{\hat{\xi}_j^2 - (\Delta p_{\text{opt},j})^2}{\sigma_{p_j}^2}}_{\equiv \Delta_{\text{pull,prior}}} \\ &> 0. \end{aligned}$$

The least-squares statistic (18.29) of the linearised model itself is never explicitly evaluated in this work, because it effectively assumes the existence of event templates which are in fact unphysical (accordingly, for a given pseudoexperiment, $\chi^2_{\text{pull,min}}$ may turn out to be smaller than χ^2_{\min}), and because of the small computational overhead of generating the template $n_{\text{exp},i}(\mathbf{p}_{\text{inj}} + \hat{\xi})$ numerically with the staged approach.

Note that the pull approach is strictly equivalent to the "covariance approach", where the least-squares statistic accounts for observables' covariances [486]. When $n < N_{\text{bins}}$, performing the minimisation within the pull approach is less computationally expensive though: a system of n , not N_{bins} , linear equations has to be solved.

Pull approach in practice

In particular when template generation and numerical minimisation have an exceedingly high computational cost and the assumption of linear parameter dependencies made by Eq. (18.27) is accurate, the pull approach can serve as a drop-in replacement and thus for example be used to approximate the median NMO sensitivity (without any reference to the parameter h).

For the analysis at hand, the approach was primarily investigated during the initial development of the statistical inference framework to reduce the encountered limitations of the Fisher information matrix—specifically, in combination with a grid scan of the fit value of θ_{23} (within the WO, but also the TO). Any $\overline{\chi^2}$ profile yields con-

fidence intervals for θ_{23} assuming Wilks' theorem (cf. Sec. 13.3), whether the NMO is assumed to be correctly identified in the fit or not. Each hypothesised value of θ_{23} is also accompanied by the Fisher matrix entering the calculation of the conditional pulls of the remaining model parameters in Eq. (18.31). When the injected model coincides with that preferred by our priors within the hypothesised NMO, the matrix inverse is the covariance matrix. When it does not, the matrix inverse still permits us to determine the parameters' pulls through a generalisation of Eq. (18.31) which is also given in Appendix D.1.

Figure 18.4 demonstrates some of the above possibilities offered by the pull approach. Underlying are seven separate $\overline{\chi^2}$ profiles in dependence of θ_{23} within the NO, given the nominal IO Asimov pseudoexperiment, and performed by either numerically minimising $\overline{\chi^2}$ over a single nuisance parameter or by determining the parameter's pull within the linearised model of Eq. (18.27) starting from the nominal parameter value. The top part shows the nuisance-parameter estimates conditional on the fit value of θ_{23} , where each shaded band represents the parameter's central 68.3% range of (here: either statistical or, if applicable, combined statistical and prior) uncertainty, corresponding to twice the square root of the inverse of the (here: scalar) Fisher information at the nominal value. For comparison, the vertical extents of the hatched regions correspond to the various parameters' prior standard deviations. In all cases, the conditional parameter estimates from the pull approach deviate by at most a few percent from their numerically optimised counterparts, resulting in barely visible overestimates of the associated $\overline{\chi^2}$ (here: equivalent to $\overline{\Delta\chi_{\text{IO}}^2}$) profiles depicted in the lower part of Fig. 18.4. Since deviations between the depths and locations of each pair of profiles' minima are negligible, the approximate median NMO significances (here: of rejecting the NO) as given by Eq. (18.25), allowing any fit value of θ_{23} , also agree well.

18.1 Frequentist hypothesis test: NO or IO?




Figure 18.4: Results of seven $\overline{\Delta\chi^2_{\text{IO}}}$ profile scans in dependence of the fit value of θ_{23} within the NO, in the presence of one nuisance parameter at a time, using numerical minimisation and the pull approach. The upper part shows the conditional nuisance-parameter estimates and the lower part the seven profiles themselves, contrasted with the scenario of statistical uncertainties only. See text for details.




Figure 18.5: Same as the lower part of Fig. 18.4, but for the simultaneous optimisation over all nuisance parameters. In addition, an implementation of the pull approach which makes use of a second iteration is employed. The panel at the bottom shows the same results after subtracting the profile obtained via numerical minimisation. See text for details.

Figure 18.5 reveals limitations of the pull approach, in particular of the simple variant based on finite differences, when all seven nuisance parameters are jointly optimised. The $\overline{\Delta\chi_{\text{IO}}^2}$ profile is significantly overestimated (differences with respect to numerical minimisation at the bottom) across a large portion of the considered θ_{23} interval. Also, as Fig. D.2 in the appendix evidences, the conditional best-fit values of several nuisance parameters do not give accurate representations of their (numerically optimised) actual dependencies on the fit value of θ_{23} any more. When the initial pull calculation with finite differences is followed by a second iteration—here with polynomial fits of the count expectations, whose derivatives are evaluated at the best-fit parameter values from the first iteration—the overestimate of $\overline{\Delta\chi_{\text{IO}}^2}$ is strongly reduced for most θ_{23} values. While the NMO significance (18.25) for any fit value of θ_{23} from numerical minimisation over the nuisance parameters is $n_\sigma \approx 2.65$, those of the two implementations of the pull approach are $n_\sigma \approx 2.78$ and $n_\sigma \approx 3.36$. For reference, the significance found in the absence of other nuisance parameters besides θ_{23} is $n_\sigma \approx 3.67$.

In case one did not have recourse to a reliable or fast numerical minimisation routine, one could conceive of further optimisations of the pull approach, such as adaptively imposing more restrictive parameter intervals during the calculations of the partial derivatives (18.17) in order to improve their accuracy in the vicinity of the true minimum of the least-squares statistic.

In general, any (conditional or global) best fit from the pull approach can be used as a cross-check for a numerical least-squares minimisation routine. In fact, such a

cross-check uncovered an initial failure in the joint optimisation of several nuisance parameters. This was subsequently fixed by adapting the minimiser configuration until fit outcomes were achieved which consistently improved on those of the pull approach as seen above.

18.1.4 Sensitivity from ensemble studies

Based on the coverage discussion in Sec. 13.3, only the explicit sampling of test-statistic distributions under the two NMO hypotheses, by means of ensembles of fits to pseudoexperiments, guarantees accurate sensitivity projections across the full model-parameter space. Using the Poissonian log-likelihood function of observing the data \mathbf{n}_{obs} given the hypothesis $\mathbf{p}_{\mathcal{H}}$,

$$\begin{aligned} \mathcal{L}(\mathbf{p}_{\mathcal{H}}) &\equiv \ln L(\mathbf{n}_{\text{obs}}|\mathbf{p}_{\mathcal{H}}) \\ &= \sum_{i=1}^{N_{\text{bins}}} \left[n_{\text{obs},i} \ln n_{\text{exp},i}(\mathbf{p}_{\mathcal{H}}) - n_{\text{exp},i}(\mathbf{p}_{\mathcal{H}}) - \ln(n_{\text{obs},i}!) \right] - \sum_{j=1}^{N_{\text{prior}}} \frac{\Delta p_{\mathcal{H},j}}{2\sigma_{p_{\mathcal{H},j}}^2}, \end{aligned} \quad (18.33)$$

our test statistic is defined as double the difference between the log-likelihood maxima within the two NMO hypotheses,

$$\mathcal{S} \equiv -2 \left[\mathcal{L}(\hat{\mathbf{p}}_{\text{NO}}) - \mathcal{L}(\hat{\mathbf{p}}_{\text{IO}}) \right] = -2 \ln \frac{L(\mathbf{n}_{\text{obs}}|\hat{\mathbf{p}}_{\text{NO}})}{L(\mathbf{n}_{\text{obs}}|\hat{\mathbf{p}}_{\text{IO}})}. \quad (18.34)$$

The log-likelihood function is maximised numerically by minimising $-\mathcal{L}$, using the same routine relied upon by the $\overline{\Delta\chi^2}$ approach.

For a given truth model \mathbf{p}_{TO} , the two distributions of \mathcal{S} follow from generating a set of N_{pseudo} statistically independent pseudoexperiments for each NMO hypothesis $\mathbf{p}_{\text{inj}} = \mathbf{p}_{\text{TO}}, \mathbf{p}_{\text{test}}$ as

$$n_{\text{obs},i}(\mathbf{p}_{\text{inj}}) \sim \text{Pois} \left[n_{\text{exp},i}(\mathbf{p}_{\text{inj}}) \right]. \quad (18.35)$$

Since each Asimov template is a smoothed event distribution, no fluctuations owing to finite MC event statistics are considered.

The NMO sensitivity derived from the above procedure is subject to statistical fluctuations. As an example, if \mathcal{S} were Gaussian distributed and our experiment able to exclude the WO at $\sim 3\sigma$ in the median case, at least $N_{\text{pseudo}} \sim \mathcal{O}(10^4)$ pseudoexperiments would be required to achieve a relative statistical uncertainty of less than 1% [4]. Here, the most extreme outcomes of the distribution of \mathcal{S} under \mathbf{p}_{test} —beyond the median of the distribution under \mathbf{p}_{TO} —would be sampled by ~ 30




Figure 18.6: Test-statistic probability densities of a test for the IO without systematics, using pseudoexperiments and the Asimov approach. Here, the median significance of excluding the IO is estimated from an unbinned Gaussian likelihood fit (solid blue curve) to the sample of \mathcal{S} under the nominal IO model (blue histogram). The standard deviation $\sigma = 10.8$ of this fit is indicated. The vertical blue line marks the mean of the fit, whereas the vertical red line marks the median of the sample of \mathcal{S} given the nominal NO model. For comparison, the dashed black curve represents Eq. (18.8) under the IO in the Asimov approach, and the vertical dashed black lines are the means of the two Gaussians (Asimov Gaussian distribution under the NO not displayed).

pseudoexperiments on average.

The statistic \mathcal{S} coincides with the $\Delta\chi^2$ statistic defined in Eq. (18.6) in the asymptotic (Gaussian) limit, which is also a sufficient condition for the Gaussian distribution (18.8) of $\Delta\chi^2$ [481]. While the same sensitivity discussion would therefore apply to \mathcal{S} (see, e.g., Ref. [419]), the aim of the ensemble studies is to determine PINGU’s NMO sensitivity without a-priori assumptions about the distribution of the test statistic.

18.2 Statistical uncertainties only

We use the idealised scenario of only statistical uncertainties to provide an initial overview of PINGU’s NMO potential, compare the sensitivity predictions made by the different statistical approaches, and illustrate general computational challenges encountered with the ensemble method. For these purposes, only the two nominal NMO hypotheses are tested against each other.

For illustration, Fig. 18.6 displays two histogrammed samples of \mathcal{S} obtained from

Method	True NO	True IO
ensemble (Gauss.)	5.49 ± 0.06	5.54 ± 0.06
ensemble (reweight.)	5.50 ± 0.02	5.52 ± 0.02
$\overline{\Delta\chi^2}$	5.49	5.54
Fisher matrix	5.48	5.54

Table 18.1: Median NMO significance (18.3) without systematics after four years of exposure time using different analysis methods. See text for details.

5×10^3 pseudoexperiments per injected model, assuming four (calendar) years of PINGU exposure time. Owing to the lack of free parameters in this problem, no minimisation is required: when calculating \mathcal{S} , \hat{p}_{NO} and \hat{p}_{IO} are simply set to the injected hypotheses. For definiteness, the plot shows how the projected median significance of excluding the IO given true NO is found. The available number of pseudoexperiments prohibits *directly* sampling test-statistic outcomes that are sufficiently “extreme” for the determination of α ($\beta = 0.5$). Hence, the latter has to be obtained through indirect means, such as an unbinned Gaussian maximum-likelihood fit to the sample of \mathcal{S} under the WO. The result is superimposed on the histogram. The two values of the Asimov metric $\overline{\Delta\chi^2_{\mathcal{H}_{\text{inj}}}}$ are displayed in addition, as is the Gaussian approximation (18.8) of the distribution of $\Delta\chi^2$ under the WO. This approximation is completely indistinguishable from the aforementioned likelihood fit.

Table 18.1 summarises the median sensitivities from ensemble studies—with and without (see below) Gaussian assumptions—and from the Asimov $\overline{\Delta\chi^2}$ and the Fisher matrix approach. In the ensemble case, each entry corresponds to the mean of the median significances from 10^3 statistically independent analyses, with the bounds of the resulting (approximately symmetric) central 1σ interval indicated. The NMO choice has a negligible impact on the sensitivity given our pair of nominal truth models. The various estimates at the level of 5.5 standard deviations are in excellent agreement with each other.

18.2.1 Sampling extreme test-statistic outcomes

A high-significance scenario as the one encountered above would require producing—and, in the presence of nuisance parameters, fitting—on average at least $\mathcal{O}(10^7)$ pseudoexperiments under the WO hypothesis in order to sample correspond-

ingly extreme test-statistic outcomes.⁴ This is generally unfeasible in our case and represents a problem which is exacerbated during the detector design stage.

As an alternative to imposing a Gaussian test-statistic distribution, when the median sensitivity is sought the sample of test-statistic outcomes under the TO hypothesis by construction is well suited to be reweighted to the WO hypothesis [488]. A given unweighted (unit weight) pseudoexperiment produced under the model \mathbf{p}_{TO} and characterised by a set of observed bin counts has the statistical weight

$$w_{\text{obs}}^{\text{TO}}(\mathbf{p}_{\text{test}}) = \prod_{i=1}^{N_{\text{bins}}} \frac{p_{\text{Pois}}[n_{\text{obs},i}|n_{\text{exp},i}(\mathbf{p}_{\text{test}})]}{p_{\text{Pois}}[n_{\text{obs},i}|n_{\text{exp},i}(\mathbf{p}_{\text{TO}})]} \quad (18.36)$$

to be observed under the null hypothesis \mathbf{p}_{test} . The superscript TO specifies the originally underlying hypothesis. Equation (18.36) is nothing other than the ratio between the Poisson likelihoods of observing the pseudoexperiment under \mathbf{p}_{test} and \mathbf{p}_{TO} . In the *absence* of nuisance parameters, the weight follows directly from the test-statistic outcome \mathcal{S}_{obs} as

$$\ln w_{\text{obs}}^{\text{TO}}(\mathbf{p}_{\text{test}}) = \begin{cases} -\mathcal{S}_{\text{obs}}/2, & \text{TO} = \text{IO} \\ +\mathcal{S}_{\text{obs}}/2, & \text{TO} = \text{NO} \end{cases} . \quad (18.37)$$

If *all* unweighted pseudoexperiments produced under \mathbf{p}_{TO} were considered, the logarithm of the weight would thus be approximately Gaussian distributed. As an example, the weight under the nominal IO model of an originally unweighted pseudoexperiment with $\mathcal{S}_{\text{obs}} = \mathcal{S}_{\text{med}}^{\text{NO}} \sim -30$ generated under the assumption $\mathcal{H}_{\text{inj}} = \text{NO}$ (cf. Fig. 18.6) is merely of the order of 10^{-7} .

The type-I error probability from the set of weights $\{w_{\text{obs},1}^{\text{TO}}(\mathbf{p}_{\text{test}}), w_{\text{obs},2}^{\text{TO}}(\mathbf{p}_{\text{test}}), \dots, w_{\text{obs},N_{<(>)}}^{\text{TO}}(\mathbf{p}_{\text{test}})\}$ of all $N_{<(>)} \approx N_{\text{pseudo}}/2$ test-statistic outcomes under the TO with $\mathcal{S}_{\text{obs}} < (>) \mathcal{S}_{\text{med}}^{\text{TO}}$ corresponding to a rejection power of 50 % reads

$$\alpha(\beta = 0.5) = \frac{\sum_{i=1}^{N_{<(>)}} w_{\text{obs},i}^{\text{TO}}(\mathbf{p}_{\text{test}})}{N_{\text{pseudo}}} . \quad (18.38)$$

The result of reweighting the corresponding ensemble of trials produced under the NO in Fig. 18.6 to the nominal IO model is shown in Fig. 18.7, together with the original unweighted ensemble of trials produced under the IO. The weighted sample has a visibly Gaussian shape and results in a median significance of $n_{\sigma} = 5.51$. The

⁴Assuming that the tail of the distribution is Gaussian.




Figure 18.7: Use of the reweighted pseudoexperiments generated under the NO from Fig. 18.6 in order to sample the left tail of test-statistic outcomes under the IO. This allows determining the median significance without employing a Gaussian approximation of the test-statistic distribution shape. The vertical red line, the dashed lines, and the histogrammed samples to the right of $\mathcal{T} \approx -12$ are the same as in Fig. 18.6.

distribution of median significances from repeated reweighting procedures is in statistical agreement with the other results in Table 18.1.

The generalisation of Eq. (18.36) with appropriate substitutions for \mathbf{p}_{TO} and \mathbf{p}_{test} in principle enables sampling the test-statistic distribution across the space of true parameter values without repeatedly producing new trials. Simply reweighting each existing pseudoexperiment to the desired injected parameter values ($\mathbf{p}_{\mathcal{H}_{\text{inj}}} \rightarrow \mathbf{p}'_{\mathcal{H}_{\text{inj}}}$) is then not sufficient. In addition, the trial's test-statistic value \mathcal{S}_{obs} must be recalculated using the two new hypotheses to be tested against each other. Otherwise, the sensitivity after reweighting would reflect a situation in which the assumed values of fixed model parameters systematically deviated from the true values realised in nature. When at least one nuisance parameter is considered, the same reasoning implies that only the injected values of parameters which are among the set of parameters optimised in the search for the log-likelihood maxima $\mathcal{L}_{\text{obs}}(\hat{\mathbf{p}}_{\text{NO}})$ and $\mathcal{L}_{\text{obs}}(\hat{\mathbf{p}}_{\text{IO}})$ may be modified.

Whether the above generalised reweighting approach is a viable alternative throughout the parameter space of interest depends on the accuracy and precision of the resulting NMO significance estimate and is not further scrutinised in this thesis.

18.2.2 Sensitivity evolution with exposure time

Figure 18.8 demonstrates how the median significance without systematic uncertainties is projected to scale with PINGU exposure time t_{live} for true NO, according to the ensemble reweighting, the Asimov $\overline{\Delta\chi^2}$, and the Fisher matrix approach. The scenario of true IO is nearly indistinguishable from the one shown and therefore omitted. With the considered exposure time spanning a range from below one up to ten years, each template bin's count expectation $n_{\text{exp},i} \propto t_{\text{live}}$ also varies by more than one order of magnitude. Still, there is no indication of a disagreement between the predictions made by the three methods even for the lowest exposure times studied. The typical standard deviation of the ensemble method's significance projections is similar to that in Table 18.1, much smaller than the marker height.

Given the median experimental outcome under the nominal NO model, one can tell that the IO is expected to be excluded at the 3σ CL within slightly more than one year of measurement, and at the 5σ CL within four years. The impact of a more or less extreme statistical fluctuation of the observed test-statistic value is illustrated based on the Asimov $\overline{\Delta\chi^2}$ approach: deviations from the median by no more than one respectively two standard deviations result in the NMO significance ranges represented by the green and yellow bands. These correspond to rejection powers $0.159 \lesssim 1 - \beta \lesssim 0.841$ respectively $0.046 \lesssim 1 - \beta \lesssim 0.954$. After one year of measurement, a statistical fluctuation at the level of two standard deviations is sufficient to yield an NMO significance of less than 1σ . Around four (seven) years of exposure time are required to reduce the probability of a statistical fluctuation resulting in a significance below 3σ (5σ) to less than 5%.

Finally, the figure also explores the scaling with time of the median significance. While a scaling with the square root $\sqrt{t_{\text{live}}}$ would apply to the "naïve" sensitivity proxy (18.22), a somewhat slower scaling is encountered for the two-sided convention (18.3). In any case, the growth would become weaker in the presence of nuisance parameters, once statistical uncertainties become sufficiently small [45, 489].




Figure 18.8: Exposure time evolution of the median significance for true NO, using the pseudoexperiment ensemble reweighting (diamonds) and the Asimov $\overline{\Delta\chi^2}$ (solid line) and Fisher matrix (thick dashed line) approaches. The function $\propto \sqrt{t_{\text{live}}}$ (thin dashed line) is scaled to agree with the significance from reweighted pseudoexperiments at $t_{\text{live}} = 1$ year. The green (yellow) band corresponds to the range of significances obtained when allowing for statistical fluctuations within the central 1σ (2σ) interval of the Gaussian distribution of $\Delta\chi^2$ under the NO.

19

Sensitivity projections under the Gaussian approximation

The final NMO sensitivity projections for the PINGU MC event sample at hand include the impacts of the nuisance parameters introduced in Sec. 17.4 which represent known sources of systematic uncertainty. No possible improvements in our knowledge or description of these uncertainties over time are taken into account. During the eventual data-taking period, however, it is conceivable that the smallness of the encountered statistical uncertainties will require a high degree of scrutiny with respect to known and unknown sources of uncertainty [489].

All results presented in the following, pertaining to the roles of the various nuisance parameters in the NMO determination, are derived from the Asimov $\overline{\Delta\chi^2}$ approach for computational feasibility reasons. When not stated explicitly, an exposure time of four years underlies any given sensitivity projection.

19.1 Sensitivity impact by uncertainty category

Table 19.1 presents the median NMO significance resulting from including either one of three categories of nuisance parameters: those related to the intrinsic atmospheric $\bar{\nu}$ fluxes ($R_{e/\mu}$, $R_{\nu/\bar{\nu}}$, $\Delta\gamma$), to the detector (s_E , $N_{\bar{\nu}}$), and to oscillations (Δm_{31}^2 , θ_{13} , θ_{23}). Only when Δm_{31}^2 and θ_{23} are among the set of free nuisance parameters is the WO hypothesis p_{test} determined by a fit to the Asimov template of the TO, in accordance with the baseline sensitivity definition in Sec. 18.1.2. The significances without any systematic uncertainties (“none”, from Table 18.1), and with all eight nuisance parameters (“all”) are also given. We can interpret the difference between

the significance in the presence of statistical uncertainties only and the significance in the presence of any given group of nuisance parameters as the associated systematic error on the NMO.

Category	True NO	True IO
none	5.49	5.54
flux	4.09	4.36
detector	4.35	4.37
oscillations	3.27	3.04
all	2.78	2.64

Table 19.1: Median $\overline{\Delta\chi^2}$ NMO significance with select categories of nuisance parameters included, as well as without any and with all together.

included alone the systematic error amounts to $\sim 75\%$ of that of the pair. The fact that uncertainty on θ_{23} should in principle be able to reduce PINGU's NMO sensitivity to a considerable degree is expected just from a comparison by eye of the pair of template signatures in Fig. 17.8 and Fig. 17.10, which suggest that changes in θ_{23} and a flip of the NMO impact similar regions in E_{reco} and $\cos \vartheta_{\text{reco}}$, possibly with opposite signs.

Flux and detector nuisance parameters only induce approximately half as large systematic errors, ranging from 1.1σ to 1.4σ . All four corresponding analyses test the two nominal NMO hypotheses against each other. Among the flux parameters, $R_{\nu/\bar{\nu}}$ is by far the most impactful when optimised alone. Both nuisance parameters s_E and $N_{\bar{\nu}}$ in the detector category have almost the same individual impact as $R_{\nu/\bar{\nu}}$ (compare Fig. 18.4).

Note that the solar oscillation parameters θ_{12} and Δm_{21}^2 are omitted because their additional inclusion in the analysis is not found to reduce the sensitivity further. δ_{CP} has a similarly minor impact and would cause an exceedingly large computational burden on the optimisation process as a result of introducing additional degeneracies into the $\overline{\chi^2}$ function.

When the Fisher matrix method with polynomial fits is employed for a fast cross-check of the systematics impacts determined above, its significance projections agree to within $\sim 1.5\%$ for all categories.

When accounting for all nuisance parameters, the sensitivity is approximately halved, below the level of 3σ , compared to the case of purely statistical uncertainties.

Oscillation parameters are by far the dominant source of systematic uncertainty and give rise to a systematic error exceeding 2σ . This is the combined effect of a modified null hypothesis p_{test} and a reduction in $\overline{\Delta\chi^2_{\mathcal{H}_{\text{inj}}}}$ with respect to the statistics-only scenario when θ_{23} and Δm_{31}^2 are jointly optimised. When θ_{23} is




Figure 19.1: Median systematic errors on the NMO (as number of Gaussian standard deviations), as defined in Secs. 19.1 and 19.2, for true NO on the left and true IO on the right, using the Asimov $\overline{\Delta\chi^2}$ approach. See text for details.

19.2 One-by-one sensitivity potential

By determining the expected sensitivity gain from an improved knowledge of each individual source of systematic uncertainty, one at a time, a different type of systematic error on the NMO is uncovered. We consider the most optimistic scenario for each nuisance parameter, in which the latter is kept fixed at its true value, whereas the remaining seven are optimised together. The difference between the median significance obtained in this manner and the baseline significance, in the presence of all eight nuisance parameters, can also be considered an “incremental” systematic error on the NMO.

By and large, the individual sensitivity gains (or incremental systematic errors) are much smaller than the overall sensitivity loss from optimising all nuisance parameters simultaneously (full systematic error). This is demonstrated in Fig. 19.1, which also includes the systematic errors induced by the various uncertainty categories defined previously. For example, for true IO the overall systematic error amounts to 2.9σ , cf. Table 19.1, which can be contrasted with the incremental systematic error of close to 0.8σ for θ_{23} —the nuisance parameter with the strongest impact in this case. Fixing the energy scale or the overall rate normalisation leads to moderate gains corresponding to incremental systematic errors of a few tenths of standard deviations. Small or almost vanishing gains are encountered in the case of the flux parameters and θ_{13} .

In general, the sizes of the various types of systematic error on the NMO are rather

stable under the exchange of the TO, and the atmospheric oscillation parameters Δm_{31}^2 and θ_{23} are the two dominant sources of systematic uncertainty.

19.3 Event-classification synergies

So far, we have not focussed on the mechanism through which PINGU's sensitivity to the NMO is brought about at the highest and most straightforward level, namely the joint fit of the cascade- and track-like event templates to a given pseudoexperiment.

Consider the Asimov median significance (18.11) for illustration. Each of the two minima $\overline{\Delta\chi_{\text{WO}}^2}$ and $\overline{\Delta\chi_{\text{TO}}^2}$ can be written as a sum of three terms,

$$\overline{\Delta\chi_{\mathcal{H}_{\text{inj}}}^2} = \overline{\Delta\chi_{\mathcal{H}_{\text{inj}},\text{cscd}}^2} + \overline{\Delta\chi_{\mathcal{H}_{\text{inj}},\text{trck}}^2} + \overline{\Delta\chi_{\mathcal{H}_{\text{inj}},\text{prior}}^2}, \quad (19.1)$$

with a contribution from the cascade-like event template, the track-like event template, and the nuisance parameters' external constraints, respectively, at the minimum within the other ordering with respect to \mathcal{H}_{inj} . Considering only one contribution at a time and ignoring the remaining two yields the median significance originating from this contribution within the joint fit. In the case of either template this quantity is referred to as the "statistical" contribution [3]. Synergistic effects occur when there is tension between the parameter values preferred by the various contributions. The magnitude of the synergy corresponds to the enhancement of the median significance compared to the result of the simple addition in the sense of Eq. (19.1) but after performing *separate* fits.

The assumption of statistical uncertainties only as done in Sec. 18.2 removes any distinction between the joint analysis and the simple addition of individual contributions. It also renders $\overline{\Delta\chi_{\mathcal{H}_{\text{inj}},\text{prior}}^2} = 0$. Accordingly, the statistical contribution from cascade- or track-like events coincides with their standalone sensitivity to the NMO. For both, it is $\sim 4\sigma$ after four years of exposure time, independent of whether we assume true NO or true IO. Hence, both event classes contribute evenly to the overall sensitivity.

This picture is modified in the presence of nuisance parameters, which results in the cascade-like event template contributing most of the overall sensitivity. Figure 19.2 demonstrates this by means of the time evolution of the baseline PINGU median significance, the statistical contributions from cascade-like and track-like events, their standalone significances, and the significance resulting from the simple addition of the standalone analysis results, for the two nominal NMO models. First of




Figure 19.2: Exposure time evolution of the synergy effect on the median NMO significance from jointly fitting the cascade-like and the track-like event template, with respect to the simple addition of their standalone best-fit metric values, using the Asimov $\overline{\Delta\chi^2}$ approach. True NO is assumed on the left, and true IO on the right. The three thick lines in each panel give the baseline median significance, the statistical contribution to it from cascade-like events, and the statistical contribution from track-like events. The thin lines give the median significance from the simple addition of the standalone $\overline{\Delta\chi^2}$ values from separately fitting the two types of templates, and the significances from the standalone analyses themselves. See text for details.

all, note that the prior constraints on nuisance parameters are included in the standalone analyses, i.e., as penalty terms in every minimisation and as a contribution to the standalone significance. Hence, the observable synergy effect, given by the difference between the baseline significance and the significance derived from the simple sum, is a “conservative” estimate. In contrast, by definition there is no explicit contribution from prior constraints to the statistical contributions. The latter exceed the standalone significances visibly in the case of true IO, for which the joint analysis is projected to be subject to a non-negligible synergy, of 0.3σ after four years, and 0.5σ after ten years of measurement. The joint analysis benefits the significance originating from the cascade-like event template in particular.

The origin of the stronger synergy for true IO lies in a stronger tension between the best-fit points preferred by cascade-like and track-like events within each opposite half-space in relation to p_{TO} respectively p_{test} . To illustrate this point, we consider the

one-dimensional Asimov profile $\overline{\chi^2}(p_{\text{WO,scan}}, \hat{\mathbf{p}}_{\text{WO,nuis}})$, where $p_{\text{WO,scan}}$ is a parameter on a predefined grid of hypothesised points within the WO, and $\hat{\mathbf{p}}_{\text{WO,nuis}}$ represent the conditional nuisance-parameter best-fit values. The square root of the profile's minimum is equivalent to the “naïve” sensitivity proxy $\sqrt{|\Delta\chi_{\text{TO}}^2|}$ for excluding the WO. This can be interpreted as a one-sided median significance (under the assumptions stated in the derivation of Eq. (18.24)), which will deviate slightly from the corresponding significance (18.11) in Fig. 19.2.

True IO

Figure 19.3 shows profile scans of $\Delta m_{31}^2 > 0$ for true IO for the same types of analysis as in Fig. 19.2. As the NMO sensitivity is the primary interest, and not the constraints on Δm_{31}^2 within the WO, the synergistic effects that manifest themselves through comparisons of the $\overline{\chi^2}$ minima are studied, and not the various profiles' shapes.

The synergy due to tension in Δm_{31}^2 is visible in particular in the simple sum of the two standalone profiles. The sum profile exhibits a “global” synergy: its minimum is moderately enhanced with respect to the sum of the minima of the two standalone profiles, namely by $\sim 6\%$, or 0.27 units of $\overline{\chi^2}$. This enhancement is brought about by effectively forcing the fit of the cascade-like and the track-like template to assume the same hypothesised value of Δm_{31}^2 —while being allowed to find standalone minima in the space of the remaining, undisplayed, nuisance parameters.

The fact that this freedom is removed by the full joint fit is reflected in the strictly larger conditional $\overline{\chi^2}$ values of the corresponding profile (“baseline” vs. “simple sum”). The minimal additional increase in $\overline{\chi^2}$ from jointly optimising also the un-




Figure 19.3: Asimov $\overline{\chi^2}$ profile scans of Δm_{31}^2 within the NO, for true IO. The lines refer to the same analysis types as in Fig. 19.2.

displayed nuisance parameters is $\sim 10\%$, or 0.61 units of $\overline{\chi^2}$, and occurs at $\Delta m_{31}^2 \approx 2.62 \times 10^{-3} \text{ eV}^2$. The global synergy—irrespective of Δm_{31}^2 and relevant for the NMO sensitivity—stemming from these parameters amounts to 1.28 units of $\overline{\chi^2}$.

In total, the synergy of the combined fit with respect to the sum of the minima of the two standalone fits results in an enhancement of the $\overline{\chi^2}$ metric by 1.55 units. The combined fit generates an increase of 0.80 units of the $\overline{\chi^2}$ minimum for cascade-like events with respect to their standalone fit, but only a barely noticeable one of 0.07 units for track-like events.

True NO

The scenario of hypothesising $\Delta m_{31}^2 < 0$ when the NO is assumed to be true is displayed in Fig. 19.4. No tension occurs between the best-fit values of Δm_{31}^2 from cascade-like and track-like events. Hence, the parameter’s joint optimisation (“simple sum”) does not increase the χ^2 minimum beyond the simple addition of the two standalone minima, which is dominated by the contribution from cascade-like events.

The minima of the three profiles of the baseline analysis and its statistical contributions all appear at $\Delta m_{31}^2 \approx -2.24 \times 10^{-3} \text{ eV}^2$. At its minimum, the statistical contribution of the cascade-like (track-like) event template exceeds the minimum of its standalone counterpart by 0.26 (0.07) χ^2 units. Thus, there is synergy from the joint optimisation of the undisplayed nuisance parameters in addition to Δm_{31}^2 . Locally, there is a small interval around $\Delta m_{31}^2 \approx -2.2 \times 10^{-3} \text{ eV}^2$ where the joint fit of the remaining nuisance parameters has a nearly vanishing impact on the χ^2 values.

The global synergy of the baseline analysis with respect to the sum of the minima of the two standalone fits is given by 0.64 units of $\overline{\chi^2}$, less than half as much as in the case of true IO.




Figure 19.4: Same as Fig. 19.3, but for true NO and Δm_{31}^2 within the IO.

19.4 Coverage in dependence of truth assumptions

The most thorough understanding of PINGU’s NMO potential and the frequentist sensitivity approach is gained by studying a broader range of truth assumptions made about nuisance parameters. Our focus is on the atmospheric oscillation parameters Δm_{31}^2 and θ_{23} . These have already been shown to give rise to considerable systematic errors in Secs. 19.1 and 19.2. The question remains to which degree the test-statistic distribution under each NMO hypothesis, and therefore the median sensitivity, depends on the values of Δm_{31}^2 and θ_{23} . In this context, the issue of obtaining proper coverage raised in Sec. 18.1.2 is revisited with the $\overline{\Delta\chi^2}$ approach. Both problems are addressed simultaneously by determining the function $\overline{\Delta\chi_{\text{TO}}^2}(\boldsymbol{p}_{\text{TO}})$, which fully specifies the distribution of $\Delta\chi^2$ according to Eq. (18.8) for any nuisance-parameter values within the TO hypothesis. In practice, we keep all nuisance parameters except Δm_{31}^2 and θ_{23} fixed at their nominal values.¹

The results are depicted in Fig. 19.5 for the rectangular parameter regions defined by the 3σ confidence intervals on Δm_{31}^2 and θ_{23} as obtained by NuFIT 2.0 [475, 476]. On the one hand, large variations are observed in $\left|\overline{\Delta\chi_{\text{NO}}^2}\right|$, which ranges from around 5.2 to 34.9. $\left|\overline{\Delta\chi_{\text{IO}}^2}\right|$, on the other hand, only varies from about 5.3 to 12.7. Both metrics increase monotonically with the absolute value of Δm_{31}^2 . Non-monotonic behaviour is observed in dependence of θ_{23} : $\left|\overline{\Delta\chi_{\text{NO}}^2}\right|$ exhibits a valley in the first octant, at $\sim 41^\circ$, whereas $\left|\overline{\Delta\chi_{\text{IO}}^2}\right|$ exhibits one in the second octant, at $\sim 49^\circ$. The maximum is located at the largest sampled values of θ_{23} and $|\Delta m_{31}^2|$ in either case. The disparate upper bounds on the two Asimov metrics imply that the scenario of symmetric test-statistic distributions is not always realised, namely when the NO is true and $\theta_{23} \gtrsim 45^\circ$, almost regardless of Δm_{31}^2 . Under these conditions $\left|\overline{\Delta\chi_{\text{NO}}^2}\right| > \left|\overline{\Delta\chi_{\text{IO}}^2}\right|$, no matter which hypothesis from the allowed IO parameter region is tested.

For any grid point within the TO, the two $\overline{\Delta\chi_{\text{TO}}^2}$ maps enable us to compute the associated maximum type-I error probability $\alpha_{\text{WO,max}}^{3\sigma}$ by performing a dedicated hypothesis test with each point in the WO as the null hypothesis. This search is identically performed across the displayed 3σ parameter ranges. The resulting median NMO significance (18.11) is given in Fig. 19.6. It varies between 2.5σ and 6.8σ in the case of true NO, and between 2.5σ and 3.7σ in the case of true IO. The behaviour of

¹The distinct treatment of the atmospheric oscillation parameters applies only to the sampling of the TO model space—and, by extension, to the construction of the WO model within the NMO hypothesis test—not to the definition of the $\overline{\Delta\chi_{\text{TO}}^2}$ metric. The value of the latter is always found through the joint optimisation of all nuisance parameters within the WO.




Figure 19.5: Gaussian test-statistic mean $|\overline{\Delta\chi_{\text{TO}}^2}|$ in dependence of the true values of Δm_{31}^2 and θ_{23} , for TO = NO on the left and TO = IO on the right. The displayed parameter ranges correspond to the respective 3σ confidence intervals obtained by NuFIT 2.0 [475, 476].

the $|\overline{\Delta\chi_{\text{TO}}^2}|$ surfaces is reproduced qualitatively, because, regardless of the assumed point in the TO, the maximum type-I error probability for an outcome as or more extreme than $\pm\Delta\chi_{\text{TO}}^2$ is always selected from the same set of $\Delta\chi^2$ distributions under the WO. The ordering of the sampled TO models thus remains the same, irrespective of whether they are sorted according to $|\overline{\Delta\chi_{\text{TO}}^2}|$ or n_σ . Quantitatively, however, the impact of the maximisation over WO hypotheses depends strongly on the true values of Δm_{31}^2 and θ_{23} .

This dependence is evidenced as a function of θ_{23} for the two nominal values $\Delta m_{31}^2 = 2.46 \times 10^{-3} \text{ eV}^2$ and $\Delta m_{31}^2 = -2.37 \times 10^{-3} \text{ eV}^2$ in Fig. 19.7. In the case of true NO and the first octant, the width of the range of median significances obtained by testing all WO models is only at the level of a few tenths of standard deviations. For the second octant or true IO, the width is generally larger, and reaches up to one (true IO) or two (true NO, second octant) standard deviations. The lower bound of the range—determined by the WO model with the largest type-I error $\alpha_{\text{WO,max}}^{3\sigma}$ —can also be read off from the vertical slice at the corresponding value of Δm_{31}^2 in Fig. 19.6. In the considered θ_{23} interval this sensitivity minimum reaches up to 6.5σ , namely when the NO is true and θ_{23} is at its upper bound. The sensitivity is lowest, around 2.6σ , for true NO and $\theta_{23} \approx 41.5^\circ$, as well as for true IO and $\theta_{23} \approx 49^\circ$. The two latter values of θ_{23} roughly coincide with those of our nominal NMO models. For comparison, the median significance calculated with the baseline approach of excluding the




Figure 19.6: Dependence of the median NMO significance n_σ on the true values of Δm_{31}^2 and θ_{23} , for TO = NO on the left and TO = IO on the right, using the test-statistic means of Fig. 19.5. In order to ensure correct coverage, for each pair of true values the null hypothesis is chosen such that the type-I error is maximised over Δm_{31}^2 and θ_{23} within the WO.

WO model that yields the best fit to the TO Asimov experiment is also shown. The coverage of the approach is found to be accurate for true NO and θ_{23} in the first octant, as well as at the upper bound of its sampled range in the second octant, and for all but the largest values $\theta_{23} \gtrsim 52^\circ$ and true IO. Only for θ_{23} in the second octant and true NO does it consistently and significantly overestimate the minimal confidence level of excluding any reasonably constrained IO model. Furthermore, as expected from the investigation of Fig. 19.5, the theoretical significance minimum for a given TO model yields a lower confidence level than the other two methods for true NO and $\theta_{23} \gtrsim 46^\circ$.

The observed shape of the median significance is elucidated by considering the various hypothesised or best-fit values encountered for a given truth θ_{23} . We denote the best-fit value at the minimum of $\overline{\chi_{\text{WO}}^2}$ as $\hat{\theta}_{23,\text{WO}}$, the value within the null hypothesis (WO model) of the baseline analysis as $\hat{\theta}_{23,\text{WO}}^{\text{bl}}$, the value within the null hypothesis of the analysis that maximises the type-I error as $\hat{\theta}_{23,\text{WO}}^{\alpha_{\text{max}}}$, and the best-fit value within the TO model which constitutes the best Asimov fit to the null hypothesis as $\hat{\theta}_{23,\text{TO}}^{\text{bl}}$ or $\hat{\theta}_{23,\text{TO}}^{\alpha_{\text{max}}}$. Hence, depending on the type of analysis conducted, one




Figure 19.7: Dependence of the median NMO significance on the true value of θ_{23} , for different choices of the WO model, for TO = NO on the left and TO = IO on the right. Δm_{31}^2 is set to its nominal truth assumption in each case. The curve resulting from maximising the type-I error probability is a vertical slice of the two-dimensional surface given the corresponding TO in Fig. 19.6. See text for details.

of the following sets of parameter values is encountered:

$$\boldsymbol{\theta}_{23}^{\text{bl}} \equiv (\theta_{23}, \hat{\theta}_{23,\text{WO}}, \hat{\theta}_{23,\text{WO}}^{\text{bl}}, \hat{\theta}_{23,\text{TO}}^{\text{bl}}) \text{ or} \quad (19.2)$$

$$\boldsymbol{\theta}_{23}^{\alpha_{\text{max}}} \equiv (\theta_{23}, \hat{\theta}_{23,\text{WO}}, \hat{\theta}_{23,\text{WO}}^{\alpha_{\text{max}}}, \hat{\theta}_{23,\text{TO}}^{\alpha_{\text{max}}}) . \quad (19.3)$$

The six distinct quantities for each TO model are compared in Fig. 19.8. Once more the true parameter value is the independently varied quantity. $\hat{\theta}_{23,\text{WO}}^{\alpha_{\text{max}}}$ is the only other variable which is not the direct result of a fit to an Asimov template. The existence of rapid changes in its value reflect the fact that $\Delta\chi_{\text{WO}}^2$ is not an injective function of the two oscillation parameters (combined with the discreteness of the sampling grids). As a result, there can be a complete degeneracy between two or more WO models with the same type-I error probability $\alpha_{\text{WO,max}}^{3\sigma}$. Since $\hat{\theta}_{23,\text{TO}}^{\alpha_{\text{max}}}$ is the result of a full (eight-dimensional) fit to the Asimov template given $\hat{\theta}_{23,\text{WO}}^{\alpha_{\text{max}}}$ (and $\widehat{\Delta m_{31,\text{WO}}^{2,\alpha_{\text{max}}}}$), its variations are correlated with those of the latter.

$\hat{\theta}_{23,\text{WO}}$ and $\hat{\theta}_{23,\text{WO}}^{\text{bl}}$, which differ only in the set of jointly fit nuisance parameters, closely follow each other. They lie in the wrong octant in the case of true NO and




Figure 19.8: Dependence of various best-fit or hypothesised values of θ_{23} on the latter's true value in Asimov NMO hypothesis tests, for TO = NO on the left and TO = IO on the right. See text for details.

$39.5^{\circ} \lesssim \theta_{23} < 45^{\circ}$, and in the case of true IO and $\theta_{23} > 45^{\circ}$. This implies a degeneracy between the NMO and the octant of θ_{23} : the WO model that minimises the weighted-least-squares sum misidentifies the true octant of θ_{23} . For true NO, the jump to the second octant around 39.5° coincides with a local maximum of the baseline significance at which the IO is excluded in the left panel of Fig. 19.7. The local minimum at $\theta_{23} \approx 41.5^{\circ}$ follows from the fact that the two test-statistic distributions are both centred around similarly low absolute values at the level of five units, cf. Fig. 19.5. This can be contrasted with the situation at $\theta_{23} \approx 50^{\circ}$, where the significance obtained within the baseline approach is maximal: whereas the median absolute test-statistic value under the NO shows an increase by a factor of ~ 5 , the distribution under the null hypothesis (for which again $\hat{\theta}_{23,WO=IO}^{bl} \approx 50^{\circ}$) remains nearly unchanged. The same considerations in reverse explain why the significance of excluding the NO in the case of true IO exhibits a local minimum at the same value of θ_{23} . It starts to grow again for even higher values of θ_{23} due to the increase in the median $\Delta\chi_{IO}^2$, roughly where $\hat{\theta}_{23,TO=IO}^{bl}$ jumps from the second into the first octant. Both observed jumps in $\hat{\theta}_{23,TO=IO}^{bl}$ are in agreement with the coincidence of the octant of $\hat{\theta}_{23,WO=IO}$ with that of $\theta_{23} \lesssim 39.5^{\circ}$ within the NO in the panel on the left. For true NO, the baseline significance of excluding the IO drops for $\theta_{23} \gtrsim 50^{\circ}$ even though $|\overline{\Delta\chi_{NO}^2}|$




Figure 19.9: Same as Fig. 19.7, but in dependence of the true value of Δm_{31}^2 , with maximal 2–3 mixing in nature.

keeps growing. This growth is compensated by the increasing standard deviation of the test-statistic distribution under the baseline IO model for a similarly large value of $\hat{\theta}_{23, \text{WO}=\text{IO}}^{\text{bl}}$. The IO model that maximises the type-I error probability remains at the largest allowed value for all $\theta_{23} \gtrsim 44^\circ$, which follows from the aforementioned guaranteed regime of asymmetry with $|\overline{\Delta\chi_{\text{NO}}^2}| > |\overline{\Delta\chi_{\text{IO}}^2}|$ which takes effect here.² The associated monotonic increase of the significance above this θ_{23} threshold value is then just a reflection of the monotonic increase of $|\overline{\Delta\chi_{\text{NO}}^2}|$. The significance appears to become linear at $\theta_{23} \approx 48^\circ$, above which we also find an approximate proportionality between θ_{23} and $\hat{\theta}_{23, \text{WO}=\text{IO}}^{\text{bl}}$ or $\hat{\theta}_{23, \text{WO}=\text{IO}}^{\text{bl}}$. Given true IO, $\hat{\theta}_{23, \text{WO}=\text{NO}}^{\alpha_{\text{max}}}$ always remains in the first octant, as this is the only way the symmetric configuration can be realised. Correspondingly, for values of θ_{23} in the two aforementioned valleys of $\Delta\chi_{\text{IO}}^2$, $\hat{\theta}_{23, \text{WO}=\text{NO}}^{\alpha_{\text{max}}}$ jumps back and forth between the two degenerate parameter regions where $\Delta\chi_{\text{NO}}^2$ reaches equally low levels. For all other true parameter values, $40^\circ \lesssim \theta_{23} \lesssim 47^\circ$ and $\theta_{23} \gtrsim 51^\circ$, the null hypothesis remains comparably stable.

For the two slices at true maximal mixing, Fig. 19.9 exemplifies that the median NMO significance grows rather weakly with $|\Delta m_{31}^2|$. The significance of excluding the IO lies between 3.5σ and 4σ , whereas that of excluding the NO only slightly ex-

²In this regime the maximum of $\Delta\chi_{\text{IO}}^2$ always maximises $\alpha_{\text{IO}}^{3\sigma}$.




Figure 19.10: Same as Fig. 19.8, but for Δm_{31}^2 .

ceeds 3σ . Generally, the impact of sampling different WO hypotheses is larger for true IO than it is for true NO, evidenced by the consistently wider range of significances associated with the allowed values of $\alpha_{\text{WO}=\text{NO}}^{3\sigma}$. The baseline approach yields a consistent minor overestimate of the actual confidence level when the NO is true. Its prediction agrees with the lower bound of the allowed range of significances for true IO. The theoretical significance minimum given true IO also coincides with the lower bound because the symmetric test-statistic configuration is always achievable. In the case of the NO and maximal mixing, any asymmetric test-statistic configurations (which occur for $\Delta m_{31}^2 \gtrsim 2.44 \times 10^{-3} \text{ eV}^2$) are too insignificant to be discernible here.

The various hypothesised or best-fit values of Δm_{31}^2 —the counterparts of Eqs. (19.2) and (19.3)—encountered in dependence of the true value of Δm_{31}^2 when conducting sensitivity studies for maximal 2–3 mixing are presented in Fig. 19.10. All three best-fit quantities obtained within the baseline approach, $\widehat{\Delta m_{31, \text{WO}}^2}$ and $\widehat{\Delta m_{31, \text{WO}, \text{TO}}^{2, \text{bl}}}$, are proportional to the true value Δm_{31}^2 . In particular, the greater the true magnitude $|\Delta m_{31}^2|$, the greater is the magnitude of each of the three best-fit values. When the NO is true, a fit of the IO hypothesis (lower left plot) yields a magnitude of the mass-squared difference that is smaller than in nature. The opposite holds when the IO is true and a fit with the NO hypothesis (upper right plot) is performed. For

true NO and $\Delta m_{31}^2 \gtrsim 2.44 \times 10^{-3} \text{ eV}^2$, an outcome $\widehat{\Delta m}_{31, \text{WO}=\text{IO}}^{2, \alpha_{\text{max}}}$ at the lower bound of the respective 3σ range maximises $\left| \overline{\Delta \chi_{\text{IO}}^2} \right|$ and therefore the type-I error probability. For true IO, the observed behaviour of $\widehat{\Delta m}_{31, \text{WO}=\text{NO}}^{2, \alpha_{\text{max}}}$ reflects the degeneracy between different null hypotheses with respect to conserving the symmetry $\left| \overline{\Delta \chi_{\text{NO}}^2} \right| = \left| \overline{\Delta \chi_{\text{IO}}^2} \right|$: a small change in Δm_{31}^2 can be compensated by an appropriate small change in either $\widehat{\Delta m}_{31, \text{WO}=\text{NO}}^{2, \alpha_{\text{max}}}$ or $\hat{\theta}_{23, \text{WO}=\text{NO}}^{\alpha_{\text{max}}}$ or simultaneous changes in both. The specific locations of the grid points along which $\left| \overline{\Delta \chi_{\text{NO}}^2} \right|$ is evaluated result in null hypotheses with a total of three different discrete values of $\hat{\theta}_{23, \text{WO}=\text{NO}}^{\alpha_{\text{max}}}$ (not displayed). As hinted at in the figure, if the analysis were forced to fix the null hypothesis to either one of these, there would be a characteristic linear relationship between the associated value of $\widehat{\Delta m}_{31, \text{WO}=\text{NO}}^{2, \alpha_{\text{max}}}$ and Δm_{31}^2 .

Part V

Conclusion

20

Summary and outlook

Experimental advances achieved during the past few decades suggest that the Standard Model's description of the neutrino sector is lacking crucial elements, first and foremost a mechanism which explains the small overall neutrino mass scale and the tiny mass differences underlying the flavour transitions unambiguously observed in vacuum and in matter.

If the Standard Model was the effective low-energy manifestation of some complete physics theory at high energies, neutrino mass could be the dominant new-physics effect to appear at the experimentally accessible energy scales. As of yet, both the absolute masses and their ordering—given the most common definition of the three known mass eigenstates—remain unresolved. Interestingly, new, non-standard types of neutrino interactions may accompany neutrino masses. The experimentally allowed strengths of certain non-standard neutral-current couplings which affect neutrino propagation in matter are of the order of the strength of the weak interaction. More precise measurements of non-standard interactions and the determination of the elusive mass ordering will serve as crucial theory inputs and reduce uncertainties of measurements targeting other important parameters within the leptonic sector of the Standard Model. However, intrinsic degeneracies of the neutrino evolution for example in the Earth hamper the simultaneous resolution of both problems. The latter requires complementary observables from measurements of flavour transitions in the Sun or neutrino scattering cross-section measurements.

This thesis contributes to the investigation of both of these open issues. On the one hand, it makes use of a sample of around 45×10^3 atmospheric neutrinos collected over three years with the existing low-energy core (DeepCore) of the IceCube Observatory at the South Pole to present a comprehensive measurement of neutral-current non-standard interactions (NC NSI). On the other hand, it draws upon simu-

lated atmospheric neutrino interactions in a proposed low-energy IceCube upgrade (PINGU)—which is even more densely instrumented and expected to result in an event-rate enhancement by a factor of around four—to study the potential of resolving the neutrino mass ordering (NMO), by employing statistical techniques differing by many orders of magnitude in computational expense.

In the absence of any experimental evidence for NC NSI so far and due to the large number of unknown coupling strengths in the most general effective NC interaction Lagrangian, the observed DeepCore data sample was separately fit with hypotheses with manageable numbers—either one, two, or three—effective parameters of interest. This approach was also motivated by the existence of theoretical models predicting one (dominant) non-zero coupling and ensured the comparability with other measurements. The underlying event selection could be adopted with few modifications from a preceding DeepCore measurement of the atmospheric oscillation parameters Δm_{31}^2 and θ_{23} , the constraints on which mainly originate in the observation of $\bar{\nu}_\mu$ disappearance. While NSI couplings to muon and tau flavours were shown to modify this disappearance pattern in distinctive ways, couplings involving the electron flavour were found to give rise to other, harder-to-resolve distortions of the expected event distribution. In any case, only moderate impacts from several sources of systematic uncertainty were encountered on the expected 90% confidence intervals on the various parameters of interest.

We found no evidence of NSI in either of our six fits to the observed DeepCore event distribution. In turn, assuming the normal ordering for definiteness, limits at the 90% confidence level could be placed on all NSI strengths in their standard parameterisation. The most restrictive is that on μ - τ flavour violation, $|\epsilon_{\mu\tau}^\oplus| \leq 0.0232$, followed by that on μ - τ non-universality, $-0.041 \leq \epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus \leq 0.042$, on e - μ flavour violation, $|\epsilon_{e\mu}^\oplus| \leq 0.146$, and on e - τ flavour violation, $|\epsilon_{e\tau}^\oplus| \leq 0.173$. The only type of NSI whose strength could not be constrained beyond that of the Standard Model's weak interaction is e - μ non-universality, which effectively rescales the Earth's standard neutrino matter potential and thus allows for example for the complete cancellation of standard matter effects. Here, the disjoint intervals $-2.26 \leq \epsilon_{ee}^\oplus - \epsilon_{\mu\mu}^\oplus \leq -1.27$ and $-0.74 \leq \epsilon_{ee}^\oplus - \epsilon_{\mu\mu}^\oplus \leq 0.32$ are allowed at the 90% confidence level. In addition, limits were placed within an alternative parameterisation of the Earth's matter potential, after a physically motivated reduction from the most general, eight-dimensional, parameter space to one comprising three (intrinsic) matter parameters, namely an overall scale parameter ϵ_\oplus and two rotation angles φ_{12} and φ_{13} . This parameterisation has been employed by various global analyses of neutrino oscillation experi-

ments and accommodates a less restrictive NSI flavour structure than the aforementioned one-by-one fits. By construction, ϵ_{\oplus} could not be measured more precisely than $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$, while similar widths of $\sim 20^\circ$ were obtained for the 90 % confidence intervals on φ_{12} and φ_{13} .

The NSI measurements reported in this thesis are compatible with other results from the broader landscape of neutrino oscillation and scattering experiments. The high-energy reach of our DeepCore atmospheric neutrino event sample limits the sensitivity to μ - τ and e - μ , and to a more limited extent e - τ , flavour-violating NSI. A lower energy threshold would benefit the sensitivity to standard matter effects and equivalently e - μ non-universality, which in turn would give a better handle on the overall strength of the generalised matter potential with its arbitrary flavour structure. In this regard, near-term improvements can be expected from the anticipated IceCube Upgrade and KM3NeT/ORCA, whose partial configuration has already yielded competitive limits (albeit not yet on e - μ non-universality). An event-by-event discrimination between neutrinos and antineutrinos would increase the sensitivity to many NSI scenarios which currently suffer from cancellations owing to opposite-sign effects at the level of the oscillation probabilities. Similarly, cancellations occur between events with different flavours which currently can only be separated statistically, such as $\vec{\nu}_e$ and $\vec{\nu}_\tau$ CC events. The measurement of any given NSI parameter would benefit from better energy and zenith reconstruction resolutions, in particular when different parameter values result in rapid variations of the oscillation probabilities in both dimensions. While the sensitivity analyses suggest an only subdued role of systematic uncertainties, it stands to reason that those on $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ and $\epsilon_{e\tau}^{\oplus}$ could be somewhat reduced if the true value of the leptonic mixing angle θ_{23} —which each of our measurements simultaneously constrained to a certain extent—was known with a higher precision.

The NMO sensitivity studies in this thesis were built and expanded upon preceding efforts of projecting PINGU's performance. The focus was placed on methods which are particularly useful during the detector optimisation stage. First, different approaches which derive their sensitivity estimates from our prediction of the typical event distribution, the "Asimov" template, were discussed and compared in detail. The one with the least restrictive regularity conditions requires evaluating the difference between two minimised weighted-least-squares statistics to infer the shapes of Gaussian test-statistic distributions under the pair of discrete ordering hypotheses. Constraints from the expected Fisher information matrix were scrutinised by way of comparison to the same weighted-least-squares statistic and demonstrated

to be ill-suited to describe the octant degeneracy of the mixing angle θ_{23} : when the NMO is assumed to be correctly identified, in addition to the global best fit at the true value of θ_{23} a local solution within the wrong octant is often also present, but when the NMO is assumed to be misidentified, the wrong octant may even provide a better fit than the correct one. This motivated adopting the Fisher information matrix for parameter estimation via the semianalytic minimisation of a modified weighted-least-squares statistic that explicitly assumes linear dependencies of the event-count expectations on the model parameters. The proposed application of this “pull approach” permits constructing a confidence interval for θ_{23} while approximating the NMO sensitivity, without being fundamentally limited by the octant degeneracy.

Given the assumed PINGU detector and event-sample characteristics, it was shown for example that statistical uncertainties are small enough to allow for a probability of around 50 % to achieve a 5σ determination of the NMO within four years of measurement. Explicitly sampling the test-statistic distributions by reweighting ensembles of statistically fluctuated pseudoexperiments confirmed the suitability of the Asimov method. In the presence of sources of systematic uncertainty related to neutrino fluxes, oscillations, and the detector response, the median significance after four years was still found to exceed 2.6σ . For the lowest-significance scenarios, improving our knowledge of any one of the considered flux or detector nuisance parameters is unlikely to strongly enhance PINGU’s NMO sensitivity.

More generally, the importance of a systematic evaluation of the possible shapes of the test-statistic distribution, over the allowed space of true nuisance-parameter values within both orderings, in order to create a comprehensive evaluation of the experiment’s potential measurement outcomes and achieving proper frequentist coverage was highlighted. For example, depending on the value of θ_{23} in nature, the median significance after four years of PINGU operation could exceed 6σ . Also, the potential inaccuracies inherent to pragmatic methods of approximating the maximum probability to mistakenly accept the wrong ordering were quantified in dependence of θ_{23} and Δm_{31}^2 . Adopting the Gaussian approximation of the Asimov approach, a theoretical lower bound on the PINGU sensitivity follows directly from the true test-statistic distribution and is too conservative only when the normal ordering is true and θ_{23} lies in the second octant. Instead, explicitly selecting the wrong-ordering model whose expected PINGU event distribution resembles the true one the most as the null hypothesis yields accurate results when the inverted ordering is true.

The NMO significance projections presented in this thesis rely on expected event distributions which most likely deviate in both their shape and overall normalisation

from those that will be available in the presence of updated experimental designs and accordingly adapted detector-simulation and event-reconstruction tools. Hence, all of the quoted sensitivities should only be taken as indicative of the eventual performance of a very-large-volume neutrino telescope with a neutrino energy threshold of ~ 1 GeV and an effective mass of a few Mt of ice. Further enhancing the discrimination power between the various neutrino interaction types—including between neutrinos and antineutrinos—and improving reconstruction resolutions in both energy and zenith, especially in the few-GeV range for ν_e and $\bar{\nu}_e$ CC events and below ~ 20 GeV for ν_μ and $\bar{\nu}_\mu$ CC events, represent promising routes that would reduce cancellations and limit the smearing out of the NMO signature. In this case, the sensitivity to the NMO could indirectly benefit from an improved precision on nuisance parameters such as (the absolute value of) Δm_{31}^2 , which is known to be responsible for strong synergy effects that emerge when data from long-baseline atmospheric and medium-baseline reactor neutrino experiments is jointly analysed. Combining the staged template generation scheme with, e.g., the Asimov $\overline{\Delta\chi^2}$ hypothesis testing approach permits us not least to quickly modify the mentioned detector response characteristics in order to quantify their impact accurately.

While the construction of the PINGU experiment is not foreseen at the time of writing, the imminent deployment of the smaller IceCube Upgrade will allow addressing many of the same physics goals and applying the knowledge gained during the extensive PINGU design and sensitivity analysis phase.

Appendices

A

Analytic calculation of oscillation probabilities

There are several codes available for solving the Wolfenstein equation (3.45), which also governs the $\nu_\alpha \rightarrow \nu_\beta$ transition amplitudes $S_{\alpha\beta}$ (cf. Sec. 3.3.4) in typical experimental settings [213]. The Wolfenstein equation consists of a set of homogeneous linear ordinary differential equations (ODEs) of first order with respect to the flavour transition amplitudes. The ODEs have constant coefficients in uniform and variable coefficients in non-uniform matter. However, when the continuous matter density profile along a neutrino trajectory traversing non-uniform matter is approximated by an adequate number of layers of uniform matter, n , solving one set of variable-coefficient ODEs becomes equivalent to solving n sets of constant-coefficient ODEs. This is a generalisation of the case of periodic density changes discussed for two neutrino flavours in Sec. 3.3.4.3. In both scenarios, analytic expressions for the transition amplitudes exist if no more than three neutrino flavours are assumed.

For this work, a GPU-accelerated Python adaptation [490] of Prob3++¹ was generalised to the case of an arbitrary Hermitian matter potential in either of the parameterisations (4.18) or (4.19). Prob3++ solves the three-neutrino evolution in uniform matter following an approach introduced in 1980 by Barger *et al.* [492]. Here, the evolution matrix in the basis of the neutrino mass eigenstates in vacuum, $\mathcal{S} = USU^\dagger$, with initial conditions $\mathcal{S}(L=0) = \mathbb{1}$, is found as

$$\mathcal{S}(L) = \sum_k \left[\prod_{j \neq k} \frac{2E_\nu H_{\nu, \text{mass}} - \tilde{m}_j^2 \mathbb{1}}{\widetilde{\Delta m_{kj}^2}} \right] \exp\left(-i \frac{\tilde{m}_k^2 L}{2E_\nu}\right). \quad (\text{A.1})$$

¹Originally ported to CUDA C code for the case of a single constant-density layer in Ref. [491].

$H_{\nu,\text{mass}}$ is the constant-density neutrino Hamiltonian in the mass eigenstate basis, \tilde{m}_j^2 are the neutrinos' *effective* mass-squared eigenvalues, and $\widetilde{\Delta m^2}_{kj} \equiv \tilde{m}_k^2 - \tilde{m}_j^2$ the effective mass-squared differences in matter. Equation (A.1) holds for an arbitrary number of neutrinos, but the analytic diagonalisation of H_ν is only possible for up to three of them. In the original Prob3++ implementation, the mass eigenvalues in matter are explicitly evaluated for SI [492, Eqs. (21) and (22)]. It is also possible, however, to find the eigenvalues of the neutrino Hamiltonian H_ν in matter with NSI, i.e., with the matter part given by Eq. (4.7), by solving for the cubic roots of the characteristic polynomial of H_ν [493, 494, 495, 496]. This procedure yields the effective mass eigenvalues [496]

$$\frac{\tilde{m}_i^2}{2E_\nu} = \frac{2}{3}\sqrt{p} \cos \left[\frac{1}{3} \arctan \left(\frac{\sqrt{p^3 - q^2}}{q} \right) + k_i \right] - \frac{c_2}{3} \quad (i = 1, 2, 3), \quad (\text{A.2})$$

with $k_1 = \frac{2}{3}\pi$, $k_2 = -\frac{2}{3}\pi$, and $k_3 = 0$, $p = c_2^2 - 3c_1$ and $q = -\frac{27}{2}c_0 - c_2^3 + \frac{9}{2}c_1c_2$, and the coefficients

$$c_0 = H_{\nu,ee} |H_{\nu,\mu\tau}|^2 + H_{\nu,\mu\mu} |H_{\nu,e\tau}|^2 + H_{\nu,\tau\tau} |H_{\nu,e\mu}|^2 - 2 \text{Re} (H_{\nu,e\mu} H_{\nu,\mu\tau} H_{\nu,e\tau}) - H_{\nu,ee} H_{\nu,\mu\mu} H_{\nu,\tau\tau}, \quad (\text{A.3})$$

$$c_1 = H_{\nu,ee} H_{\nu,\mu\mu} + H_{\nu,ee} H_{\nu,\tau\tau} + H_{\nu,\mu\mu} H_{\nu,\tau\tau} - |H_{\nu,e\mu}|^2 - |H_{\nu,\mu\tau}|^2 - |H_{\nu,e\tau}|^2, \quad (\text{A.4})$$

$$c_2 = -H_{\nu,ee} - H_{\nu,\mu\mu} - H_{\nu,\tau\tau}. \quad (\text{A.5})$$

By substituting the \tilde{m}_i^2 into Eq. (A.1) and transforming to the flavour basis one obtains the flavour evolution matrix in the uniform medium. The oscillation probabilities correspond to the absolute squares of the various elements of the product of all n evolution matrices,

$$P_{\alpha\beta}(L) = \left| \left[S_n(L_n) S_{n-1}(L_{n-1}) \dots S_1(L_1) \right]_{\alpha\beta} \right|^2, \quad (\text{A.6})$$

where S_i is the evolution matrix and L_i the neutrino path length in the i th layer and the total baseline is given by $L = L_n + L_{n-1} + \dots + L_1$.

Computational efficiency is the main advantage of the preceding analytic calculation with respect to numerical approaches, which is why Prob3++ or the oscillation code included in the GLOBES software package [497, 498] are favoured in practice. Numerical approaches such as nuSQuIDS [499, 500] provide quasicontinuous density profiles and are typically able to handle more than three neutrino species. A per-

formance comparison between Prob3++ and nuSQuIDS for three-neutrino oscillations in Earth matter for both CPU and GPU implementations performed in Ref. [501] shows a runtime reduction by approximately two orders of magnitude.

In preparation for Part III of this work, the analytically computed oscillation probabilities have been found to be in excellent agreement with numerical approaches for typical settings of atmospheric and long-baseline neutrino oscillation experiments, and for several SI and NSI hypotheses.

B

Analytic breakdown of atmospheric (anti)neutrino event rates

Before attempting to analytically break down the rate of atmospheric neutrino and antineutrino events in DeepCore or PINGU, it is useful to introduce

1. the ratio of the unoscillated fluxes of $\bar{\nu}_\mu$'s and $\bar{\nu}_e$'s, $R_{\mu/e}^0 \equiv \bar{\Phi}_{\text{unosc}}^\mu / \bar{\Phi}_{\text{unosc}}^e$,
2. the ratio of the unoscillated fluxes of ν_μ 's and $\bar{\nu}_\mu$'s, $R_{\nu_\mu/\bar{\nu}_\mu}^0 \equiv \Phi_{\text{unosc}}^\mu / \bar{\Phi}_{\text{unosc}}^\mu$,
3. the differences $\Delta \bar{P}_{\alpha\beta} \equiv \bar{P}'_{\alpha\beta} - \bar{P}_{\alpha\beta}$ between the (anti)neutrino oscillation probabilities for different oscillation or NSI parameter values, for example $\theta'_{23} > \pi/4 > \theta_{23}$ or $\epsilon'_\oplus = -1$ vs. $\epsilon_\oplus = 0$,
4. the effective areas \bar{A}_{eff}^e , \bar{A}_{eff}^μ , and $\bar{A}_{\text{eff}}^\tau$ for detecting $\bar{\nu}_e$ CC, $\bar{\nu}_\mu$ CC, and $\bar{\nu}_\tau$ CC events, respectively,
5. the ratio of the effective areas for $\bar{\nu}_e$ CC and $\bar{\nu}_\tau$ CC events, $\bar{a}_{e/\tau}^0 \equiv \bar{A}_{\text{eff}}^e / \bar{A}_{\text{eff}}^\tau$,
6. the ratio of the effective areas for ν_e CC and $\bar{\nu}_e$ CC events, $a_{\nu_e/\bar{\nu}_e}^0 \equiv A_{\text{eff}}^e / \bar{A}_{\text{eff}}^e$
7. and the probabilities \bar{r}_e , \bar{r}_μ , and \bar{r}_τ to classify (anti)neutrino CC events of the three flavours as cascade-like, for example.

Then we can write the nominal total rate of CC events of a particular event class—

including both neutrino and antineutrino events—as (antineutrino terms implicit)

$$\begin{aligned}
N &= \sum_{\nu, \bar{\nu}} N_e + N_\mu + N_\tau \\
&= \sum_{\nu, \bar{\nu}} \Phi_{\text{unosc}}^\mu \left[\left(\frac{P_{ee}}{R_{\mu/e}^0} + P_{\mu e} \right) A_{\text{eff}}^e r_e + \left(P_{\mu\mu} + \frac{P_{e\mu}}{R_{\mu/e}^0} \right) A_{\text{eff}}^\mu r_\mu + \left(\frac{P_{e\tau}}{R_{\mu/e}^0} + P_{\mu\tau} \right) A_{\text{eff}}^\tau r_\tau \right]. \quad (\text{B.1})
\end{aligned}$$

To a first approximation, for DeepCore and PINGU we can assume $\bar{P}_{e\mu} \approx \bar{P}_{\mu e}$, as well as $\bar{A}_{\text{eff}}^e \approx \bar{A}_{\text{eff}}^\mu$. Now making the antineutrino contributions explicit and using all of the variables enumerated at the top, the rate difference $\Delta N \equiv N' - N$ induced by oscillation-parameter variations and normalised to the product of the ν_μ flux and the ν_e (or $\bar{\nu}_\mu$) CC effective area becomes

$$\begin{aligned}
\frac{\Delta N}{\Phi_{\text{unosc}}^\mu A_{\text{eff}}^e} &= \frac{\sum_{\nu, \bar{\nu}} \Delta N_e + \Delta N_\mu + \Delta N_\tau}{\Phi_{\text{unosc}}^\mu A_{\text{eff}}^e} \\
&\approx \left(\frac{\Delta P_{ee}}{R_{\mu/e}^0} + \Delta P_{\mu e} \right) r_e + \left(\Delta P_{\mu\mu} + \frac{\Delta P_{\mu e}}{R_{\mu/e}^0} \right) r_\mu + \left(\frac{\Delta P_{e\tau}}{R_{\mu/e}^0} + \Delta P_{\mu\tau} \right) \frac{r_\tau}{a_{e/\tau}^0} \\
&+ \frac{1}{R_{\nu_\mu/\bar{\nu}_\mu}^0 a_{\nu_e/\bar{\nu}_e}^0} \left[\left(\frac{\Delta \bar{P}_{ee}}{\bar{R}_{\mu/e}^0} + \Delta \bar{P}_{\mu e} \right) \bar{r}_e + \left(\Delta \bar{P}_{\mu\mu} + \frac{\Delta \bar{P}_{\mu e}}{\bar{R}_{\mu/e}^0} \right) \bar{r}_\mu + \left(\frac{\Delta \bar{P}_{e\tau}}{\bar{R}_{\mu/e}^0} + \Delta \bar{P}_{\mu\tau} \right) \frac{\bar{r}_\tau}{\bar{a}_{e/\tau}^0} \right]. \quad (\text{B.2})
\end{aligned}$$

Equation (B.2) holds for a given event class at any point in true neutrino energy and zenith (dependence on both suppressed). It enables a qualitative derivation of the origins and limitations of the DeepCore or PINGU sensitivity to NSI coupling strengths or oscillation parameters given the characteristics of Earth oscillograms and a coarse understanding of the relations between various detector response functions (relative effective areas and classification efficiencies)—in the limit of perfect (δ -like) detector resolution functions. A priori, it is straightforward to tell that

- the weights of all antineutrino terms are reduced by the product of the neutrino-to-antineutrino flux and effective-area ratios;
- the terms related to $\bar{\nu}_\tau$ appearance are suppressed by the smallness of the corresponding CC cross section;
- and that the induced differences of oscillation probabilities $\bar{\nu}_e \rightarrow \bar{\nu}_\beta$, that is, $\Delta \bar{P}_{ee}$, $\Delta \bar{P}_{e\mu}$, and $\Delta \bar{P}_{e\tau}$, are suppressed due to the reduced intrinsic $\bar{\nu}_e$ fluxes.

The sum of all NC events within a given event class does not need to be taken into account when considering oscillation-induced event rate differences, because the rate of NC events is a constant (antineutrino terms implicit):

$$N_{\text{NC}} = \sum_{\nu, \bar{\nu}} r_{\text{NC}} A_{\text{eff,NC}} (\Phi_{\text{unosc}}^e + \Phi_{\text{unosc}}^\mu) . \quad (\text{B.3})$$

This phenomenon is the combined result of the flavour-universal effective area $A_{\text{eff,NC}}$ for NC events (cf. Fig. 2.5), the indistinguishability of NC events of different flavours (cf. Fig. 6.1), which gives rise to a flavour-universal classification probability r_{NC} , and the unitarity of the leptonic mixing matrix (cf. Sec. 3.1.3). Accordingly, NC events constitute a background in the measurement of neutrino oscillations with DeepCore or PINGU.

C

DeepCore NSI search: supplementary material




Figure C.1: Same as Fig. 9.1, but for antineutrinos.




Figure C.2: Same as Fig. 9.2, but for antineutrinos.




Figure C.3: Same as Fig. 9.3, but for antineutrinos.




Figure C.4: Same as the upper six panels of Fig. 9.4, but for antineutrinos.




Figure C.5: Same as the lower six panels of Fig. 9.4, but for antineutrinos.




Figure C.6: Same as Fig. 14.9, but for $\epsilon_{ee}^\oplus - \epsilon_{\mu\mu}^\oplus$.




Figure C.7: Same as Fig. C.6, but for $|\epsilon_{e\mu}^\oplus|$.




Figure C.8: Same as Fig. C.6, but for $|\epsilon_{e\tau}^{\oplus}|$.




Figure C.9: Same as Fig. C.6, but for $|\epsilon_{\mu\tau}^{\oplus}|$.




Figure C.10: Asimov sensitivity to standard NSI parameters given true SM interactions. Solid lines (“full”) result from profiling the final selection of nuisance parameters (“free parameters” in Table 12.2), while the red markers (“full+”, connected with dashed lines to guide the eye) result from profiling all nuisance parameters in Table 12.2. Each $\overline{\Delta\chi^2_{\text{mod}}}$ profile in the magnitude of a flavour-violating coupling parameters is also profiled over the undisplayed corresponding complex phase.

Event type	Cascade-like	Track-like
ν_e CC	5756 ± 20	1799 ± 11
$\bar{\nu}_e$ CC	2481 ± 13	765 ± 7
ν_μ CC	9811 ± 27	9429 ± 27
$\bar{\nu}_\mu$ CC	4328 ± 18	4935 ± 20
ν_τ CC	835 ± 7	317 ± 4
$\bar{\nu}_\tau$ CC	374 ± 5	144 ± 3
ν_e NC	465 ± 6	141 ± 3
$\bar{\nu}_e$ NC	135 ± 3	43 ± 2
ν_μ NC	1731 ± 11	569 ± 7
$\bar{\nu}_\mu$ NC	584 ± 7	193 ± 4
ν_τ NC	342 ± 4	104 ± 3
$\bar{\nu}_\tau$ NC	93 ± 2	31 ± 1
μ^\pm	1187 ± 37	1353 ± 38
Sum	28123 ± 57	19823 ± 53

Table C.1: Predicted best-fit composition of the DeepCore event sample under the GMP hypothesis (cf. Table 15.1), split up into the two event classes of cascade-like and track-like events. The sums of the event counts in each row can be compared to our nominal expectations in Table 12.1.




Figure C.11: Histogram of the observed binwise weighted-least-squares values for the μ - τ NU measurement on the left, and cumulative histogram on the right. The samples are stacked across both event classes in the left panel—such that the sum over all histogram entries corresponds to the total number of template bins—but not in the one on the right. The (rescaled) PDF respectively CDF of χ^2_1 is only shown to guide the eye, not because it represents the exact expectation (see goodness of fit).




Figure C.12: Relative deviations between the best-fit and nominal values of nuisance parameters subject to external constraints observed in the various NSI measurements, given in numbers of prior standard deviations (“pulls”).




Figure C.13: Observed joint 1σ , 90% , and 3σ confidence regions in θ_{23} and Δm_{31}^2 within the SI measurement, together with each parameter’s projected one-dimensional $\Delta\chi_{\text{mod}}^2$ profile. The colour in the central panel encodes the local value of the two-dimensional $\Delta\chi_{\text{mod}}^2$ profile and the cross indicates the best-fit point. The latter serves as the truth model of the pseudo-experiments underlying the sensitivity distribution provided for θ_{23} on the right.

D

PINGU NMO studies: supplementary material




Figure D.1: Same as Fig. 17.7, but for ν_μ CC and NC events and as a function of the true ratio of neutrino propagation distance to energy, L/E_ν .

D.1 Parameter pulls from Fisher information

The definition of the χ_{pull}^2 function (18.28) within the pull approach transformed into index notation reads (summing over repeated indices within each product)

$$\chi_{\text{pull}}^2 = \left(\Delta_i^\alpha - \xi_k \partial_k n_{\text{exp},i}^\alpha \right) \frac{\delta_{ij} \delta_{\alpha\beta}}{\sigma_i^\alpha \sigma_j^\beta} \left(\Delta_j^\beta - \xi_l \partial_l n_{\text{exp},j}^\beta \right) + \frac{\xi_k}{\sigma_k} \delta_{kl} \frac{\xi_l}{\sigma_l}, \quad (\text{D.1})$$

where $n_{\text{exp},i}^\alpha \equiv n_{\text{exp},i}^\alpha(\mathbf{p}_{\text{inj}})$, $\Delta_i^\alpha \equiv n_{\text{obs},i}^\alpha - n_{\text{exp},i}^\alpha$, and $\sigma_i^\alpha \equiv \sqrt{n_{\text{exp},i}^\alpha}$. Note the minor notational variation here with respect to Eq. (18.28) through the use of greek indices, which in this work represent the different PINGU analysis ‘‘classes’’ (cascade-like and track-like events), but which generalise to independent experiments (e.g., PINGU and JUNO in Ref. [45]).

After expanding Eq. (D.1), such that

$$\chi_{\text{pull}}^2 = \frac{\delta_{ij} \delta_{\alpha\beta}}{\sigma_i^\alpha \sigma_j^\beta} \left(\Delta_i^\alpha \Delta_j^\beta - 2 \Delta_i^\alpha \xi_l \partial_l n_{\text{exp},j}^\beta + \xi_k \partial_k n_{\text{exp},i}^\alpha \xi_l \partial_l n_{\text{exp},j}^\beta \right) + \frac{\xi_k}{\sigma_k} \delta_{kl} \frac{\xi_l}{\sigma_l},$$

and extremising with respect to the r th parameter’s pull, by imposing a vanishing derivative $\left. \frac{\partial \chi_{\text{pull}}^2}{\partial \xi_r} \right|_{\hat{\xi}_r} = 0$, one finds

$$0 = \frac{\delta_{ij} \delta_{\alpha\beta}}{\sigma_i^\alpha \sigma_j^\beta} \left(-2 \Delta_i^\alpha \partial_r n_{\text{exp},j}^\beta + \partial_r n_{\text{exp},i}^\alpha \hat{\xi}_l \partial_l n_{\text{exp},j}^\beta + \hat{\xi}_k \partial_k n_{\text{exp},i}^\alpha \partial_r n_{\text{exp},j}^\beta \right) + \frac{2}{\sigma_r} \delta_{rl} \frac{\xi_l}{\sigma_l}.$$

‘‘Executing’’ $\delta_{\alpha\beta}$ and introducing $d_l \equiv \Delta_i^\alpha \frac{\delta_{ij}}{\sigma_i^\alpha \sigma_j^\alpha} \partial_l n_{\text{exp},j}^\alpha$ then leads to

$$\sum_{r=1}^n \sum_{i,j=1}^{N_{\text{bins}}} \sum_{\alpha \in \{\text{cscd}, \text{trck}\}} \left(\partial_r n_{\text{exp},i}^\alpha \frac{\delta_{ij}}{\sigma_i^\alpha \sigma_j^\alpha} \partial_l n_{\text{exp},j}^\alpha + \frac{\delta_{rl}}{\sigma_r \sigma_l} \right)^{-1} d_r = \hat{\xi}_l, \quad (\text{D.2})$$

which represents the multiplication of the inverse of the overall Fisher matrix (sum of the Fisher matrices of cascade- and track-like events and the diagonal matrix of inverse prior variances) with the n -dimensional column vector d_r . This expression is equivalent to Eq. (18.31).

Relaxing the assumption of the injected parameter values $p_{\text{inj},k}$ coinciding with all prior preferred values $p_{\text{nom},k}$, $\xi_k/\sigma_k \rightarrow (p_{\text{inj},k} - p_{\text{nom},k} + \xi_k)/\sigma_k$, necessitates the replacement $d_r \rightarrow d_r - (p_{\text{inj},r} - p_{\text{nom},r})/\sigma_r^2$ (for every parameter with a defined prior

uncertainty) in Eq. (D.2). This modification enables iterative pull calculations, for example.

A literature solution that is of the same form as Eq. (D.2) can be found in Ref. [486, Eq. (A6)]. This reference also provided the original motivation to search for a semi-analytic expression for the best-fit parameter values that depends only on known quantities employed in the construction of PINGU's Fisher information matrix in this thesis.




Figure D.2: Conditional nuisance-parameter estimates associated with the $\overline{\Delta\chi^2_{10}}$ profiles of Fig. 18.5. See also Fig. 18.4.

Bibliography

- [1] M. G. Aartsen et al. *Letter of Intent: the Precision IceCube Next Generation Upgrade (PINGU)*. 2017. arXiv: 1401.2046v2 [physics.ins-det].
- [2] M. G. Aartsen et al. 'PINGU: a vision for neutrino and particle physics at the South Pole'. In: *J. Phys. G* 44.5 (2017), p. 054006. doi: 10.1088/1361-6471/44/5/054006. arXiv: 1607.02671 [hep-ex].
- [3] M. G. Aartsen et al. 'Combined sensitivity to the neutrino mass ordering with JUNO, the IceCube Upgrade, and PINGU'. In: *Phys. Rev. D* 101.3 (2020), p. 032006. doi: 10.1103/PhysRevD.101.032006. arXiv: 1911.06745 [hep-ex].
- [4] M. G. Aartsen et al. 'Computational techniques for the analysis of small signals in high-statistics neutrino oscillation experiments'. In: *Nucl. Instrum. Meth. A* 977 (2020), p. 164332. doi: 10.1016/j.nima.2020.164332. arXiv: 1803.05390 [physics.data-an].
- [5] R. Abbasi et al. 'All-flavour constraints on nonstandard neutrino interactions and generalised matter potential with three years of IceCube DeepCore data'. In: *Phys. Rev. D* 104.7 (2021), p. 072006. doi: 10.1103/PhysRevD.104.072006. arXiv: 2106.07755 [hep-ex].
- [6] P. Lipari. 'Atmospheric neutrinos: from the pioneering experiments to IMB and Kamiokande'. In: *International Conference on History of the Neutrino: 1930-2018 Paris, France, September 5-7, 2018*. 2019. URL: <http://neutrinohistory2018.in2p3.fr/proceedings/lipari.pdf>.
- [7] H. Georgi and S. L. Glashow. 'Unity of all elementary particle forces'. In: *Phys. Rev. Lett.* 32 (1974), pp. 438–441. doi: 10.1103/PhysRevLett.32.438.
- [8] R. M. Bionta et al. 'Contained neutrino interactions in an underground water detector'. In: *Phys. Rev. D* 38 (1988), pp. 768–775. doi: 10.1103/PhysRevD.38.768.
- [9] J. M. LoSecco. 'Physics results from underground experiments'. In: *New Frontiers in Particle Physics. Proceedings, Winter Institute, Lake Louise, Canada, February 16-22, 1986*. 1986.
- [10] T. J. Haines et al. 'Calculation of atmospheric neutrino-induced backgrounds in a nucleon-decay search'. In: *Phys. Rev. Lett.* 57 (1986), pp. 1986–1989. doi: 10.1103/PhysRevLett.57.1986.
- [11] J. M. LoSecco. 'The history of "anomalous" atmospheric neutrino events: a first-person account'. In: *Phys. Perspect.* 18.2 (2016), pp. 209–241. doi: 10.1007/s00016-016-0185-0. arXiv: 1606.00665 [physics.hist-ph].
- [12] K. S. Hirata et al. 'Experimental study of the atmospheric neutrino flux'. In: *Phys. Lett.* B205 (1988). [447(1988)], p. 416. doi: 10.1016/0370-2693(88)91690-5.
- [13] Y. Fukuda et al. 'Evidence for oscillation of atmospheric neutrinos'. In: *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. doi: 10.1103/PhysRevLett.81.1562. arXiv: hep-ex/9807003 [hep-ex].
- [14] C. Giganti, S. Lavignac and M. Zito. 'Neutrino oscillations: the rise of the PMNS paradigm'. In: *Prog. Part. Nucl. Phys.* 98 (2018), pp. 1–54. doi: 10.1016/j.pnpnp.2017.10.001. arXiv: 1710.00715 [hep-ex].

Bibliography

- [15] R. L. Workman et al. ‘Review of Particle Physics’. In: *PTEP* 2022 (2022), p. 083C01. doi: 10.1093/ptep/ptac097.
- [16] J. Elster and H. Geitel. ‘Über Elektrizitätszerstreuung in der Luft’. In: *Annalen der Physik* 307 (1900), pp. 425–446. doi: 10.1002/andp.19003070702.
- [17] C. T. R. Wilson. ‘On the leakage of electricity through dust-free air’. In: *Proc. Cambridge Phil. Soc.* 11 (1900), p. 32.
- [18] A. de Angelis. ‘Atmospheric ionisation and cosmic rays: studies and measurements before 1912’. In: *Astropart. Phys.* 53 (2014), pp. 19–26. doi: 10.1016/j.astropartphys.2013.05.010. arXiv: 1208.6527 [physics.hist-ph].
- [19] D. Pacini. ‘La radiazione penetrante alla superficie ed in seno alle acque’. In: *Il Nuovo Cimento* 3 (1912). English translation, pp. 93–100. doi: 10.1007/BF02957440. arXiv: 1002.1810 [physics.hist-ph].
- [20] V. F. Hess. *Messungen der durchdringenden Strahlung bei zwei Freiballonfahrten*. Mitteilungen aus dem Institut für Radiumforschung. In German. 1911.
- [21] V. F. Hess. ‘Über Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten’. In: *Phys. Z.* 13 (1912). In German, pp. 1084–1091.
- [22] A. H. Compton. ‘A geographic study of cosmic rays’. In: *Phys. Rev.* 43 (6 1933), pp. 387–403. doi: 10.1103/PhysRev.43.387.
- [23] L. Alvarez and A. H. Compton. ‘A positively-charged component of cosmic rays’. In: *Phys. Rev.* 43 (10 1933), pp. 835–836. doi: 10.1103/PhysRev.43.835.
- [24] E. J. Williams and G. E. Roberts. ‘Evidence for transformation of mesotrons into electrons’. In: *Nature* 145.3664 (1940), pp. 102–103. doi: 10.1038/145102a0.
- [25] C. V. Achar et al. ‘Detection of muons produced by cosmic-ray neutrinos deep underground’. In: *Phys. Lett.* 18 (1965), pp. 196–199. doi: 10.1016/0031-9163(65)90712-2.
- [26] F. Reines et al. ‘Evidence for high-energy cosmic-ray neutrino interactions’. In: *Phys. Rev. Lett.* 15 (1965), pp. 429–433. doi: 10.1103/PhysRevLett.15.429.
- [27] L. Wolfenstein. ‘Neutrino oscillations in matter’. In: *Phys. Rev. D* 17 (1978), pp. 2369–2374. doi: 10.1103/PhysRevD.17.2369.
- [28] I. Esteban, M. C. Gonzalez-Garcia and M. Maltoni. ‘On the determination of leptonic CP violation and neutrino mass ordering in presence of non-standard interactions: present status’. In: *JHEP* 06 (2019), p. 055. doi: 10.1007/JHEP06(2019)055. arXiv: 1905.05203 [hep-ph].
- [29] I. Esteban, M. C. Gonzalez-Garcia and M. Maltoni. *On the effect of NSI in the present determination of the mass ordering*. 2020. arXiv: 2004.04745 [hep-ph].
- [30] P. Coloma et al. ‘Global constraints on non-standard neutrino interactions with quarks and electrons’. In: *JHEP* 08 (2023), p. 032. doi: 10.1007/JHEP08(2023)032. arXiv: 2305.07698 [hep-ph].
- [31] P. S. Bhupal Dev et al. ‘Neutrino non-standard interactions: a status report’. In: *SciPost Physics Proceedings* 2 (2019). doi: 10.21468/scipostphysproc.2.001. arXiv: 1907.00991 [hep-ph].
- [32] J. Gehrlein et al. ‘Testing neutrino flavour models’. In: *Snowmass 2021*. Mar. 2022. arXiv: 2203.06219 [hep-ph].

Bibliography

- [33] Y. Almumin et al. ‘Neutrino flavour model building and the origins of flavour and violation’. In: *Universe* 9.12 (2023), p. 512. DOI: 10.3390/universe9120512. arXiv: 2204.08668 [hep-ph].
- [34] A. Achterberg et al. ‘First-year performance of the IceCube Neutrino Telescope’. In: *Astropart. Phys.* 26 (2006), pp. 155–173. DOI: 10.1016/j.astropartphys.2006.06.007. arXiv: astro-ph/0604450 [astro-ph].
- [35] M. G. Aartsen et al. ‘Evidence for high-energy extraterrestrial neutrinos at the IceCube detector’. In: *Science* 342 (2013), p. 1242856. DOI: 10.1126/science.1242856. arXiv: 1311.5238 [astro-ph.HE].
- [36] M. G. Aartsen et al. ‘Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert’. In: *Science* 361.6398 (2018), pp. 147–151. DOI: 10.1126/science.aat2890. arXiv: 1807.08794 [astro-ph.HE].
- [37] R. Abbasi et al. ‘Evidence for neutrino emission from the nearby active galaxy NGC 1068’. In: *Science* 378.6619 (2022), pp. 538–543. DOI: 10.1126/science.abg3395. arXiv: 2211.09972 [astro-ph.HE].
- [38] R. Abbasi et al. ‘Measurement of the high-energy all-flavour neutrino-nucleon cross section with IceCube’. In: *Phys. Rev. D* 104 (2021), p. 022001. DOI: 10.1103/PhysRevD.104.022001. arXiv: 2011.03560 [hep-ex].
- [39] O. Schulz. ‘The design study of IceCube DeepCore: characterisation and veto studies’. PhD thesis. Ruprecht-Karls-Universität Heidelberg, 2010.
- [40] R. Abbasi et al. ‘The design and performance of IceCube DeepCore’. In: *Astropart. Phys.* 35 (2012), pp. 615–624. DOI: 10.1016/j.astropartphys.2012.01.004. arXiv: 1109.6096 [astro-ph.IM].
- [41] L. Anchordoqui and F. Halzen. ‘IceHEP high-energy physics at the South Pole’. In: *Annals Phys.* 321 (2006), pp. 2660–2716. DOI: 10.1016/j.aop.2005.11.015. arXiv: hep-ph/0510389.
- [42] R. Abbasi et al. *Measurement of atmospheric neutrino oscillation parameters using convolutional neural networks with 9.3 years of data in IceCube DeepCore*. May 2024. arXiv: 2405.02163 [hep-ex].
- [43] M. G. Aartsen et al. ‘IceCube-Gen2: the window to the extreme Universe’. In: *J. Phys. G* 48.6 (2021), p. 060501. DOI: 10.1088/1361-6471/abbd48. arXiv: 2008.04323 [astro-ph.HE].
- [44] M. G. Aartsen et al. ‘PINGU sensitivity to the neutrino mass hierarchy’. In: *Community Summer Study 2013: Snowmass on the Mississippi*. 2013. arXiv: 1306.5846 [astro-ph.IM].
- [45] L. Schulte. ‘Determining the neutrino mass hierarchy with the Precision IceCube Next Generation Upgrade (PINGU)’. PhD thesis. Rheinische Friedrich-Wilhelms-Universität Bonn, 2015.
- [46] M. G. Aartsen et al. *Letter of Intent: the Precision IceCube Next Generation Upgrade (PINGU)*. 2014. arXiv: 1401.2046v1 [physics.ins-det].
- [47] T. Y. Cao. *Conceptual Developments of 20th Century Field Theories*. 2nd ed. Cambridge University Press, 2019. DOI: 10.1017/9781108566926.

- [48] G. 't Hooft, M. J. G. Veltman and B. Q. P. J. de Wit. 'Lie groups in physics'. Lecture notes. 2007. URL: <https://webpace.science.uu.nl/~hooft101/lectures/lieg07.pdf> (visited on 27/07/2023).
- [49] M. B. Robinson. *Symmetry and the Standard Model: mathematics and particle physics*. New York, USA: Springer, 2011. doi: 10.1007/978-1-4419-8267-4.
- [50] E. Noether. 'Invariant variation problems'. In: *Gott. Nachr.* 1918 (1918). [Transp. Theory Statist. Phys.1,186(1971)], pp. 235–257. doi: 10.1080/00411457108231446. arXiv: physics/0503066 [physics].
- [51] A. Salam and J. C. Ward. 'On a gauge theory of elementary interactions'. In: *Nuovo Cim.* 19 (1961), pp. 165–170. doi: 10.1007/BF02812723.
- [52] F. Englert and R. Brout. 'Broken symmetry and the mass of gauge vector mesons'. In: *Phys. Rev. Lett.* 13 (1964). [157(1964)], pp. 321–323. doi: 10.1103/PhysRevLett.13.321.
- [53] P. W. Higgs. 'Broken symmetries, massless particles, and gauge fields'. In: *Phys. Lett.* 12 (1964), pp. 132–133. doi: 10.1016/0031-9163(64)91136-9.
- [54] P. W. Higgs. 'Broken symmetries and the masses of gauge bosons'. In: *Phys. Rev. Lett.* 13 (1964). [160(1964)], pp. 508–509. doi: 10.1103/PhysRevLett.13.508.
- [55] G. S. Guralnik, C. R. Hagen and T. W. B. Kibble. 'Global conservation laws and massless particles'. In: *Phys. Rev. Lett.* 13 (1964). [162(1964)], pp. 585–587. doi: 10.1103/PhysRevLett.13.585.
- [56] P. W. Higgs. 'Spontaneous symmetry breakdown without massless bosons'. In: *Phys. Rev.* 145 (1966), pp. 1156–1163. doi: 10.1103/PhysRev.145.1156.
- [57] T. W. B. Kibble. 'Symmetry breaking in non-abelian gauge theories'. In: *Phys. Rev.* 155 (1967). [165(1967)], pp. 1554–1561. doi: 10.1103/PhysRev.155.1554.
- [58] E. Amaldi. 'Beta decay opens the way to weak interactions'. In: *Journal de Physique Colloques* 43.C8 (1982), pp. C8-261-C8-300. doi: 10.1051/jphyscol:1982816.
- [59] S. F. Novaes. 'Standard Model: an introduction'. In: *Particles and fields. Proceedings, 10th Jorge Andre Swieca Summer School, Sao Paulo, Brazil, February 6-12, 1999*. 1999, pp. 5–102. arXiv: hep-ph/0001283 [hep-ph].
- [60] D. J. Gross. 'The discovery of asymptotic freedom and the emergence of QCD'. In: *Proc. Nat. Acad. Sci.* 102 (2005). [Rev. Mod. Phys.77,837(2005)], pp. 9099–9108. doi: 10.1142/S0217751X05029009, 10.1073/pnas.0503831102, 10.1103/RevModPhys.77.837.
- [61] P. A. M. Dirac and N. H. D. Bohr. 'The quantum theory of the emission and absorption of radiation'. In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 114.767 (1927), pp. 243–265. doi: 10.1098/rspa.1927.0039.
- [62] G. N. Lewis. 'The conservation of photons'. In: *Nature* 118.2981 (1926), pp. 874–875. doi: 10.1038/118874a0.
- [63] X. Fan et al. 'Measurement of the electron magnetic moment'. In: *Phys. Rev. Lett.* 130.7 (2023), p. 071801. doi: 10.1103/PhysRevLett.130.071801. arXiv: 2209.13084 [physics.atom-ph].
- [64] H. Davoudiasl and W. J. Marciano. 'Tale of two anomalies'. In: *Phys. Rev.* D98.7 (2018), p. 075011. doi: 10.1103/PhysRevD.98.075011. arXiv: 1806.10252 [hep-ph].

Bibliography

- [65] A. Crivellin, M. Hoferichter and P. Schmidt-Wellenburg. ‘Combined explanations of $(g - 2)_{\mu,e}$ and implications for a large muon EDM’. In: *Phys. Rev. D* 98.11 (2018), p. 113002. doi: 10.1103/PhysRevD.98.113002. arXiv: 1807.11484 [hep-ph].
- [66] E. Fermi. ‘An attempt of a theory of beta radiation. 1.’ In: *Z. Phys.* 88 (1934), pp. 161–177. doi: 10.1007/BF01351864.
- [67] F. L. Wilson. ‘Fermi’s theory of beta decay’. In: *Am. J. Phys.* 36.12 (1968), pp. 1150–1160. doi: 10.1119/1.1974382.
- [68] W. Pauli. ‘Dear radioactive ladies and gentlemen’. In: *Phys. Today* 31N9 (1978), p. 27.
- [69] T. D. Lee and C. N. Yang. ‘Question of parity conservation in weak interactions’. In: *Phys. Rev.* 104 (1956), pp. 254–258. doi: 10.1103/PhysRev.104.254.
- [70] C. S. Wu et al. ‘Experimental test of parity conservation in beta decay’. In: *Phys. Rev.* 105 (1957), pp. 1413–1414. doi: 10.1103/PhysRev.105.1413.
- [71] R. L. Garwin, L. M. Lederman and M. Weinrich. ‘Observations of the failure of conservation of parity and charge conjugation in meson decays: the magnetic moment of the free muon’. In: *Phys. Rev.* 105 (1957), pp. 1415–1417. doi: 10.1103/PhysRev.105.1415.
- [72] J. I. Friedman and V. L. Telegdi. ‘Nuclear emulsion evidence for parity non-conservation in the decay chain $\pi^+ \rightarrow \mu^+ \rightarrow e^+$ ’. In: *Phys. Rev.* 106 (1957), pp. 1290–1293. doi: 10.1103/PhysRev.106.1290.
- [73] A. Salam. ‘On parity conservation and neutrino mass’. In: *Nuovo Cim.* 5 (1957), pp. 299–301. doi: 10.1007/BF02812841.
- [74] L. D. Landau. ‘On the conservation laws for weak interactions’. In: *Nucl. Phys.* 3 (1957), pp. 127–131. doi: 10.1016/0029-5582(57)90061-5.
- [75] T. D. Lee and C. N. Yang. ‘Parity non-conservation and a two-component theory of the neutrino’. In: *Phys. Rev.* 105 (1957). [245(1957)], pp. 1671–1675. doi: 10.1103/PhysRev.105.1671.
- [76] M. Goldhaber, L. Grodzins and A. W. Sunyar. ‘Helicity of neutrinos’. In: *Phys. Rev.* 109 (1958), pp. 1015–1017. doi: 10.1103/PhysRev.109.1015.
- [77] E. C. G. Sudarshan and R. E. Marshak. ‘The nature of the four-fermion interaction’. In: *Proceedings of the Padua-Venice Conference on Mesons and Newly Discovered Particles, Sept. 1957*. Also in “A Gift of Prophecy”, E. C. G. Sudarshan (ed.), World Scientific, Singapore (1994), pp. 508-515.
- [78] R. P. Feynman and M. Gell-Mann. ‘Theory of the Fermi interaction’. In: *Phys. Rev.* 109 (1958), pp. 193–198. doi: 10.1103/PhysRev.109.193.
- [79] E. C. G. Sudarshan and R. E. Marshak. ‘Chirality invariance and the universal Fermi interaction’. In: *Phys. Rev.* 109 (5 1958), pp. 1860–1862. doi: 10.1103/PhysRev.109.1860.2.
- [80] J. J. Sakurai. ‘Mass reversal and weak interactions’. In: *Il Nuovo Cimento* 7 (1958), pp. 649–660. doi: 10.1007/BF02781569.
- [81] S. L. Glashow. ‘Partial symmetries of weak interactions’. In: *Nucl. Phys.* 22 (1961), pp. 579–588. doi: 10.1016/0029-5582(61)90469-2.
- [82] A. Salam and J. C. Ward. ‘Electromagnetic and weak interactions’. In: *Phys. Lett.* 13 (1964), pp. 168–171. doi: 10.1016/0031-9163(64)90711-5.

Bibliography

- [83] S. Weinberg. 'A model of leptons'. In: *Phys. Rev. Lett.* 19 (1967), pp. 1264–1266. doi: 10.1103/PhysRevLett.19.1264.
- [84] A. Salam. 'Weak and electromagnetic interactions'. In: *Conf. Proc.* C680519 (1968), pp. 367–377.
- [85] C. N. Yang and R. L. Mills. 'Conservation of isotopic spin and isotopic gauge invariance'. In: *Phys. Rev.* 96 (1 1954), pp. 191–195. doi: 10.1103/PhysRev.96.191.
- [86] G. 't Hooft. 'Renormalisation of massless Yang-Mills fields'. In: *Nucl. Phys.* B33 (1971), pp. 173–199. doi: 10.1016/0550-3213(71)90395-6.
- [87] G. 't Hooft and M. J. G. Veltman. 'Regularisation and renormalisation of gauge fields'. In: *Nucl. Phys.* B44 (1972), pp. 189–213. doi: 10.1016/0550-3213(72)90279-9.
- [88] S. Coleman. 'The 1979 Nobel Prize in Physics'. In: *Science* 206.4424 (1979), pp. 1290–1292. doi: 10.1126/science.206.4424.1290.
- [89] F. J. Hasert et al. 'Search for elastic muon-neutrino electron scattering'. In: *Physics Letters B* 46.1 (1973), pp. 121–124. doi: 10.1016/0370-2693(73)90494-2.
- [90] F. J. Hasert et al. 'Observation of neutrino-like interactions without muon or electron in the Gargamelle neutrino experiment'. In: *Phys. Lett.* B46 (1973). [,5.15(1973)], pp. 138–140. doi: 10.1016/0370-2693(73)90499-1.
- [91] F. J. Hasert et al. 'Observation of neutrino-like interactions without muon or electron in the Gargamelle neutrino experiment'. In: *Nucl. Phys.* B73 (1974), pp. 1–22. doi: 10.1016/0550-3213(74)90038-8.
- [92] G. Arnison et al. 'Experimental observation of lepton pairs of invariant mass around $95\text{ GeV}/c^2$ at the CERN SPS collider'. In: *Phys. Lett.* B126 (1983). [,7.55(1983)], pp. 398–410. doi: 10.1016/0370-2693(83)90188-0.
- [93] P. Bagnaia et al. 'Evidence for $Z^0 \rightarrow e^+e^-$ at the CERN $\bar{p}p$ collider'. In: *Phys. Lett.* B129 (1983). [,7.69(1983)], pp. 130–140. doi: 10.1016/0370-2693(83)90744-X.
- [94] G. Arnison et al. 'Experimental observation of isolated large-transverse-energy electrons with associated missing energy at $\sqrt{s} = 540\text{ GeV}$ '. In: *Phys. Lett.* B122 (1983). [,611(1983)], pp. 103–116. doi: 10.1016/0370-2693(83)91177-2.
- [95] M. Banner et al. 'Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the CERN $\bar{p}p$ collider'. In: *Phys. Lett.* B122 (1983). [,7.45(1983)], pp. 476–485. doi: 10.1016/0370-2693(83)91605-2.
- [96] G. Aad et al. 'Observation of a new particle in the search for the Standard-Model Higgs boson with the ATLAS detector at the LHC'. In: *Phys. Lett.* B716 (2012), pp. 1–29. doi: 10.1016/j.physletb.2012.08.020. arXiv: 1207.7214 [hep-ex].
- [97] S. Chatrchyan et al. 'Observation of a new boson at a mass of 125 GeV with the CMS Experiment at the LHC'. In: *Phys. Lett.* B716 (2012), pp. 30–61. doi: 10.1016/j.physletb.2012.08.021. arXiv: 1207.7235 [hep-ex].
- [98] M. Gell-Mann. *The Eightfold Way: a theory of strong interaction symmetry*. California Institute of Technology Report CTSL-20. 1961. doi: 10.2172/4008239.

Bibliography

- [99] Y. Ne'eman. 'Derivation of strong interactions from a gauge invariance'. In: *Nucl. Phys.* 26 (1961). [,34(1961)], pp. 222–229. DOI: 10.1016/0029-5582(61)90134-1.
- [100] M. Gell-Mann. 'Symmetries of baryons and mesons'. In: *Phys. Rev.* 125 (1962), pp. 1067–1084. DOI: 10.1103/PhysRev.125.1067.
- [101] L. B. Okun. 'The theory of weak interaction'. In: *1962 International Conference on High-Energy Physics at CERN*. Ed. by J. Prentki. 1962, p. 845.
- [102] C. M. G. Lattes et al. 'Processes involving charged mesons'. In: *Nature* 159 (1947). [,42(1947)], pp. 694–697. DOI: 10.1038/159694a0.
- [103] G. D. Rochester and C. C. Butler. 'Evidence for the existence of new unstable elementary particles'. In: *Nature* 160 (1947), pp. 855–857. DOI: 10.1038/160855a0.
- [104] R. Armenteros et al. 'LVI. The properties of charged V-particles'. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 43.341 (1952), pp. 597–611. DOI: 10.1080/14786440608520216.
- [105] W. B. Fowler et al. 'Production of heavy unstable particles by negative pions'. In: *Phys. Rev.* 93 (1954), pp. 861–867. DOI: 10.1103/PhysRev.93.861.
- [106] A. Pais. 'Some remarks on the V-particles'. In: *Phys. Rev.* 86 (1952), pp. 663–672. DOI: 10.1103/PhysRev.86.663.
- [107] M. Gell-Mann. 'Isotopic spin and new unstable particles'. In: *Phys. Rev.* 92 (1953), pp. 833–834. DOI: 10.1103/PhysRev.92.833.
- [108] T. Nakano and K. Nishijima. 'Charge independence for V-particles'. In: *Prog. Theor. Phys.* 10 (1953), pp. 581–582. DOI: 10.1143/PTP.10.581.
- [109] K. Nishijima. 'Charge-independence theory of V-particles'. In: *Prog. Theor. Phys.* 13.3 (1955), pp. 285–304. DOI: 10.1143/PTP.13.285.
- [110] M. Gell-Mann. 'The interpretation of the new particles as displaced charge multiplets'. In: *Nuovo Cim.* 4.S2 (1956), pp. 848–866. DOI: 10.1007/BF02748000.
- [111] M. Gell-Mann and Y. Ne'eman. *The Eightfold way: a review with a collection of reprints*. 1964.
- [112] V. E. Barnes et al. 'Observation of a hyperon with strangeness -3 '. In: *Phys. Rev. Lett.* 12 (1964), pp. 204–206. DOI: 10.1103/PhysRevLett.12.204.
- [113] M. Gell-Mann. 'A schematic model of baryons and mesons'. In: *Phys. Lett.* 8 (1964), pp. 214–215. DOI: 10.1016/S0031-9163(64)92001-3.
- [114] G. Zweig. 'An SU(3) model for strong interaction symmetry and its breaking. Version 2'. In: *Developments in the quark theory of hadrons. Vol. 1. 1964 - 1978*. Ed. by D.B. Lichtenberg and Simon Peter Rosen. 1964, pp. 22–101.
- [115] O. W. Greenberg. 'Spin and unitary-spin independence in a paraquark model of baryons and mesons'. In: *Phys. Rev. Lett.* 13 (1964), pp. 598–602. DOI: 10.1103/PhysRevLett.13.598.
- [116] M. Y. Han and Y. Nambu. 'Three-triplet model with double SU(3) symmetry'. In: *Phys. Rev.* 139 (1965). [,187(1965)], B1006–B1010. DOI: 10.1103/PhysRev.139.B1006.

Bibliography

- [117] W. K. H. Panofsky. ‘Low- q electrodynamic-ics, elastic, and inelastic electron (and muon) scattering’. In: *Proceedings, 14th International Conference on High-Energy Physics (ICHEP 68): Vienna, Austria, 28 Aug-5 Sep, 1968*. 1968, pp. 23–42.
- [118] R. E. Taylor. ‘Inelastic electron-proton scattering in the deep continuum region’. In: *Conf. Proc. C690914* (1969), pp. 251–260.
- [119] E. D. Bloom et al. ‘High-energy inelastic ep scattering at 6-degrees and 10-degrees’. In: *Phys. Rev. Lett.* 23 (1969), pp. 930–934. doi: 10.1103/PhysRevLett.23.930.
- [120] M. Breidenbach et al. ‘Observed behaviour of highly-inelastic electron-proton scattering’. In: *Phys. Rev. Lett.* 23 (1969), pp. 935–939. doi: 10.1103/PhysRevLett.23.935.
- [121] S. L. Glashow, J. Iliopoulos and L. Maiani. ‘Weak interactions with lepton-hadron symmetry’. In: *Phys. Rev. D* 2 (1970), pp. 1285–1292. doi: 10.1103/PhysRevD.2.1285.
- [122] G. Myatt and D. H. Perkins. ‘Further observations on scaling in neutrino interactions’. In: *Phys. Lett.* 34B (1971), pp. 542–546. doi: 10.1016/0370-2693(71)90676-9.
- [123] D. H. Perkins. ‘Neutrino interactions’. In: *eConf C720906V4* (1972), pp. 189–247.
- [124] H. Fritzsch and M. Gell-Mann. ‘Current algebra: quarks and what else?’ In: *eConf C720906V2* (1972), pp. 135–165. arXiv: hep-ph/0208010 [hep-ph].
- [125] H. Fritzsch, M. Gell-Mann and H. Leutwyler. ‘Advantages of the colour-octet gluon picture’. In: *Phys. Lett.* 47B (1973), pp. 365–368. doi: 10.1016/0370-2693(73)90625-4.
- [126] H. D. Politzer. ‘Reliable perturbative results for strong interactions?’ In: *Phys. Rev. Lett.* 30 (1973). [274(1973)], pp. 1346–1349. doi: 10.1103/PhysRevLett.30.1346.
- [127] D. J. Gross and F. Wilczek. ‘Ultraviolet behaviour of non-abelian gauge theories’. In: *Phys. Rev. Lett.* 30 (1973). [271(1973)], pp. 1343–1346. doi: 10.1103/PhysRevLett.30.1343.
- [128] C. Giunti and C. W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford: Oxford Univ., 2007. doi: 10.1093/acprof:oso/9780198508717.001.0001.
- [129] P. A. M. Dirac. ‘The quantum theory of the electron’. In: *Proceedings of the Royal Society of London Series A* 117 (1928), pp. 610–624. doi: 10.1098/rspa.1928.0023.
- [130] G. Luders. ‘On the equivalence of invariance under time reversal and under particle-antiparticle conjugation for relativistic field theories’. In: *Kong. Dan. Vid. Sel. Mat. Fys. Med.* 28N5.5 (1954), pp. 1–17.
- [131] W. Pauli, L. Rosenfeld and V. Weisskopf. *Niels Bohr and the development of physics*. 1955.
- [132] J. S. Bell. ‘Time reversal in field theory’. In: *Proc. Roy. Soc. Lond. A* A231 (1955), pp. 479–495. doi: 10.1098/rspa.1955.0189.
- [133] R. Jost. ‘A remark on the C.T.P. theorem’. In: *Helv. Phys. Acta* 30 (1957), pp. 409–416.
- [134] H. Weyl. ‘Electron and gravitation. 1. (In German)’. In: *Z. Phys.* 56 (1929). [Surveys High Energ. Phys.5,261(1986)], pp. 330–352. doi: 10.1007/BF01339504.

- [135] W. Pauli. *General principles of quantum mechanics*. The original German edition appeared as Band 5, Teil 1 of: *Handbuch der Physik*, 1933. Berlin, Heidelberg: Springer, 1980. doi: 10.1007/978-3-642-61840-6.
- [136] E. Majorana. ‘Teoria simmetrica dell’elettrone e del positrone’. In: *Nuovo Cim.* 14 (1937), pp. 171–184. doi: 10.1007/BF02961314.
- [137] E. Majorana. ‘A symmetric theory of electrons and positrons’. In: *Scientific papers of Ettore Majorana: a new expanded edition*. Ed. by L. Cifarelli. Springer International Publishing, 2020, pp. 113–128. doi: 10.1007/978-3-030-23509-3_11.
- [138] P. B. Pal. ‘Dirac, Majorana, and Weyl fermions’. In: *Am. J. Phys.* 79 (2011), pp. 485–498. doi: 10.1119/1.3549729. arXiv: 1006.1718 [hep-ph].
- [139] N. Cabibbo. ‘Unitary symmetry and leptonic decays’. In: *Phys. Rev. Lett.* 10 (1963). [648(1963)], pp. 531–533. doi: 10.1103/PhysRevLett.10.531.
- [140] M. Kobayashi and T. Maskawa. ‘CP violation in the renormalisable theory of weak interaction’. In: *Prog. Theor. Phys.* 49 (1973), pp. 652–657. doi: 10.1143/PTP.49.652.
- [141] L. Alvarez-Ruso et al. ‘NuSTEC White Paper: status and challenges of neutrino–nucleus scattering’. In: *Prog. Part. Nucl. Phys.* 100 (2018), pp. 1–68. doi: 10.1016/j.ppnp.2018.01.006. arXiv: 1706.03621 [hep-ph].
- [142] L. Alvarez-Ruso et al. *Theoretical tools for neutrino scattering: interplay between lattice QCD, EFTs, nuclear physics, phenomenology, and neutrino event generators*. 2022. arXiv: 2203.09030 [hep-ph].
- [143] F. Scheck. *Quantum physics*. Springer Berlin Heidelberg, 2007.
- [144] M. Srednicki. *Quantum field theory*. Cambridge University Press, 2007.
- [145] G. Barr et al. *Particle physics in the LHC era*. Oxford master series in particle physics, astrophysics and cosmology. Oxford: Oxford University Press, 2016. URL: <https://cds.cern.ch/record/2034442>.
- [146] E. A. Paschos. *Electroweak Theory*. Cambridge University Press, 2007. doi: 10.1017/9781009402378.
- [147] F. Halzen and A. D. Martin. *Quarks and leptons: an introductory course in modern particle physics*. Wiley, 1984.
- [148] G. Rädcl and R. Beyer. ‘Neutrino-electron scattering’. In: *Mod. Phys. Lett. A* 8 (1993), pp. 1067–1088. doi: 10.1142/S0217732393002567.
- [149] P. Vilain et al. ‘Precision measurement of electroweak parameters from the scattering of muon-neutrinos on electrons’. In: *Phys. Lett. B* 335 (1994), pp. 246–252. doi: 10.1016/0370-2693(94)91421-4.
- [150] B. C. Cañas et al. ‘The weak mixing angle from low-energy neutrino measurements: a global update’. In: *Phys. Lett. B* 761 (2016), pp. 450–455. doi: 10.1016/j.physletb.2016.08.047. arXiv: 1608.02671 [hep-ph].
- [151] E. A. Paschos and O. Lalakulich. *Antineutrino-neutrino and antineutrino-electron resonant annihilation through rho and other vector mesons*. 2002. arXiv: hep-ph/0206273 [hep-ph].
- [152] V. Brdar et al. ‘Resonances in $\bar{\nu}_e e^-$ scattering below a TeV’. In: *Phys. Rev. D* 105.9 (2022), p. 093004. doi: 10.1103/PhysRevD.105.093004. arXiv: 2112.03283 [hep-ph].

Bibliography

- [153] S. L. Glashow. 'Resonant scattering of antineutrinos'. In: *Phys. Rev.* 118 (1960), pp. 316–317. doi: 10.1103/PhysRev.118.316.
- [154] M. G. Aartsen et al. 'Detection of a particle shower at the Glashow resonance with IceCube'. In: *Nature* 591.7849 (2021). [Erratum: *Nature* 592, E11 (2021)], pp. 220–224. doi: 10.1038/s41586-021-03256-1. arXiv: 2110.15051 [hep-ex].
- [155] M. Sajjad Athar, A. Fatima and S. K. Singh. 'Neutrinos and their interactions with matter'. In: *Prog. Part. Nucl. Phys.* 129 (2023), p. 104019. doi: 10.1016/j.pnpnp.2022.104019. arXiv: 2206.13792 [hep-ph].
- [156] M. Sajjad Athar and J. G. Morfin. 'Neutrino(antineutrino)–nucleus interactions in the shallow- and deep-inelastic scattering regions'. In: *J. Phys. G* 48.3 (2021), p. 034001. doi: 10.1088/1361-6471/abbb11. arXiv: 2006.08603 [hep-ph].
- [157] A. Gazizov et al. 'Neutrino-nucleon cross sections at energies of Megaton-scale detectors'. In: *EPJ Web Conf.* 116 (2016), p. 08003. doi: 10.1051/epjconf/201611608003. arXiv: 1604.02092 [hep-ph].
- [158] K. S. McFarland et al. 'A precision measurement of electroweak parameters in neutrino-nucleon scattering'. In: *Eur. Phys. J. C* 1 (1998), pp. 509–513. doi: 10.1007/s100520050099. arXiv: hep-ex/9701010.
- [159] C. L. Cowan et al. 'Detection of the free neutrino: a confirmation'. In: *Science* 124 (1956), pp. 103–104. doi: 10.1126/science.124.3212.103.
- [160] F. Reines et al. 'Detection of the free antineutrino'. In: *Phys. Rev.* 117 (1960), pp. 159–173. doi: 10.1103/PhysRev.117.159.
- [161] C. H. Llewellyn Smith. 'Neutrino reactions at accelerator energies'. In: *Phys. Rept.* 3 (1972), pp. 261–379. doi: 10.1016/0370-1573(72)90010-5.
- [162] S. Weinberg. 'Charge symmetry of weak interactions'. In: *Phys. Rev.* 112 (1958), pp. 1375–1379. doi: 10.1103/PhysRev.112.1375.
- [163] L. Grenacs. 'Induced weak currents in nuclei'. In: *Ann. Rev. Nucl. Part. Sci.* 35 (1985), pp. 455–499. doi: 10.1146/annurev.ns.35.120185.002323.
- [164] J. A. Formaggio and G. P. Zeller. 'From eV to EeV: neutrino cross sections across energy scales'. In: *Rev. Mod. Phys.* 84 (2012), pp. 1307–1341. doi: 10.1103/RevModPhys.84.1307. arXiv: 1305.7513 [hep-ex].
- [165] D. Rein and L. 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.
- [166] R. P. Feynman. 'Very high-energy collisions of hadrons'. In: *Phys. Rev. Lett.* 23 (1969). Ed. by L. M. Brown, pp. 1415–1417. doi: 10.1103/PhysRevLett.23.1415.
- [167] R. P. Feynman. 'What neutrinos can tell us about partons'. In: *Proceedings, 1st Europhysics Conference on Neutrinos (Neutrino '72): Balatonfured, Hungary, 11-17 June, 1972*. Ed. by A. Frenkel and G. Marx. 1972, pp. 75–100. doi: 10.5281/zenodo.5110200.
- [168] J. D. Bjorken. 'Asymptotic sum rules at infinite momentum'. In: *Phys. Rev.* 179 (1969), pp. 1547–1553. doi: 10.1103/PhysRev.179.1547.

- [169] H. W. Kendall. 'Deep inelastic scattering: experiments on the proton and the observation of scaling'. In: *Rev. Mod. Phys.* 63 (1991), pp. 597–614. doi: 10.1103/RevModPhys.63.597.
- [170] K. Zuber. *Neutrino Physics*. Boca Raton: Taylor & Francis, 2020. doi: 10.1201/9781315195612.
- [171] C. G. Callan Jr. and D. J. Gross. 'High-energy electroproduction and the constitution of the electric current'. In: *Phys. Rev. Lett.* 22 (1969), pp. 156–159. doi: 10.1103/PhysRevLett.22.156.
- [172] H. Gallagher, G. Garvey and G. P. Zeller. 'Neutrino-nucleus interactions'. In: *Ann. Rev. Nucl. Part. Sci.* 61 (2011), pp. 355–378. doi: 10.1146/annurev-nucl-102010-130255.
- [173] G. F. Chew. 'The inelastic scattering of high-energy neutrons by deuterons according to the impulse approximation'. In: *Phys. Rev.* 80 (1950), pp. 196–202. doi: 10.1103/PhysRev.80.196.
- [174] A. M. Ankowski. 'Breakdown of the impulse approximation and its consequences: the low- Q^2 problem'. In: *PoS NFACT08* (2008). Ed. by A. Breskin et al., p. 118. doi: 10.22323/1.074.0118. arXiv: 0810.1167 [nucl-th].
- [175] J. T. Sobczyk. 'Recent developments in modelling neutrino interactions in 1 GeV energy region'. In: *Acta Phys. Polon. B* 41 (2010). Ed. by A. Zalewska, pp. 1491–1507. arXiv: 1005.3401 [hep-ex].
- [176] D. Z. Freedman. 'Coherent neutrino-nucleus scattering as a probe of the weak neutral current'. In: *Phys. Rev. D* 9 (1974), pp. 1389–1392. doi: 10.1103/PhysRevD.9.1389.
- [177] V. B. Kopeliovich and L. L. Frankfurt. 'Isotopic and chiral structure of neutral current'. In: *JETP Lett.* 19 (1974), pp. 145–147.
- [178] D. Akimov et al. 'Observation of coherent elastic neutrino-nucleus scattering'. In: *Science* 357.6356 (2017), pp. 1123–1126. doi: 10.1126/science.1220990. arXiv: 1708.01294 [nucl-ex].
- [179] P. S. Barbeau, Yu. Efremenko and K. Scholberg. 'COHERENT at the Spallation Neutron Source'. In: *Annual Review of Nuclear and Particle Science* 73.1 (2021), pp. 41–68. doi: 10.1146/annurev-nucl-101918-023518. arXiv: 2111.07033 [hep-ex].
- [180] V. A. Bednyakov and D. V. Naumov. 'Coherency and incoherency in neutrino-nucleus elastic and inelastic scattering'. In: *Phys. Rev. D* 98.5 (2018), p. 053004. doi: 10.1103/PhysRevD.98.053004. arXiv: 1806.08768 [hep-ph].
- [181] V. A. Bednyakov and D. V. Naumov. 'Concept of coherence in neutrino and antineutrino scattering off nuclei'. In: *Phys. Part. Nucl.* 52.1 (2021), pp. 39–154. doi: 10.1134/S1063779620060039.
- [182] J. Barranco, O. G. Miranda and T. I. Rashba. 'Probing new physics with coherent neutrino scattering off nuclei'. In: *JHEP* 12 (2005), p. 021. doi: 10.1088/1126-6708/2005/12/021. arXiv: hep-ph/0508299.
- [183] A. Yu. Smirnov. *Solar neutrinos: oscillations or no oscillations?* 2016. arXiv: 1609.02386 [hep-ph].
- [184] Y. Nir. 'The Standard Model'. Lectures given at the 5th Chilean School of High Energy Physics, Universidad Técnica Federico Santa María, Chile, 15–18 January 2018. URL: https://indico.cern.ch/event/628497/attachments/1582420/2494937/Standard_Model-Yosef_Nir.pdf (visited on 27/07/2023).

Bibliography

- [185] M. C. Gonzalez-Garcia and M. Maltoni. ‘Phenomenology with massive neutrinos’. In: *Phys. Rept.* 460 (2008), pp. 1–129. doi: 10.1016/j.physrep.2007.12.004. arXiv: 0704.1800 [hep-ph].
- [186] A. de Gouvêa. ‘Neutrino mass models’. In: *Ann. Rev. Nucl. Part. Sci.* 66 (2016), pp. 197–217. doi: 10.1146/annurev-nucl-102115-044600.
- [187] B. Pontecorvo. ‘Neutrino experiments and the problem of conservation of leptonic charge’. In: *Sov. Phys. JETP* 26 (1968). [*Zh. Eksp. Teor. Fiz.* 53,1717(1967)], pp. 984–988.
- [188] A. de Gouvêa. ‘TASI lectures on neutrino physics’. In: *Physics in D \geq 4. Proceedings, Theoretical Advanced Study Institute in elementary particle physics, TASI 2004, Boulder, USA, June 6-July 2, 2004*. 2004, pp. 197–258. arXiv: hep-ph/0411274 [hep-ph].
- [189] R. R. Volkas. ‘Introduction to sterile neutrinos’. In: *Prog. Part. Nucl. Phys.* 48 (2002). [161(2001)], pp. 161–174. doi: 10.1016/S0146-6410(02)00122-9. arXiv: hep-ph/0111326 [hep-ph].
- [190] S. M. Bilenky and S. T. Petcov. ‘Massive neutrinos and neutrino oscillations’. In: *Rev. Mod. Phys.* 59 (1987). [Erratum: *Rev. Mod. Phys.* 61, 169 (1989), Erratum: *Rev. Mod. Phys.* 60, 575–575 (1988)], p. 671. doi: 10.1103/RevModPhys.59.671.
- [191] K. Fujikawa. ‘Operatorial characterisation of Majorana neutrinos’. In: *Eur. Phys. J. C* 80.3 (2020), p. 285. doi: 10.1140/epjc/s10052-020-7855-4. arXiv: 1910.03189 [hep-ph].
- [192] E. Kh. Akhmedov. ‘Do charged leptons oscillate?’ In: *JHEP* 09 (2007), p. 116. doi: 10.1088/1126-6708/2007/09/116. arXiv: 0706.1216 [hep-ph].
- [193] B. Pontecorvo. ‘Mesonium and anti-mesonium’. In: *Sov. Phys. JETP* 6 (1957). [*Zh. Eksp. Teor. Fiz.* 33,549(1957)], p. 429.
- [194] B. Pontecorvo. ‘Inverse beta processes and non-conservation of lepton charge’. In: *Sov. Phys. JETP* 7 (1958). [*Zh. Eksp. Teor. Fiz.* 34,247(1957)], pp. 172–173.
- [195] Z. Maki, M. Nakagawa and S. Sakata. ‘Remarks on the unified model of elementary particles’. In: *Prog. Theor. Phys.* 28 (1962). [34(1962)], pp. 870–880. doi: 10.1143/PTP.28.870.
- [196] F. Wilczek and A. Zee. ‘Operator analysis of nucleon decay’. In: *Phys. Rev. Lett.* 43 (1979), pp. 1571–1573. doi: 10.1103/PhysRevLett.43.1571.
- [197] W. Grimus and L. Lavoura. ‘The seesaw mechanism at arbitrary order: disentangling the small scale from the large scale’. In: *JHEP* 11 (2000), p. 042. doi: 10.1088/1126-6708/2000/11/042. arXiv: hep-ph/0008179.
- [198] H. Hettmansperger, M. Lindner and W. Rodejohann. ‘Phenomenological consequences of subleading terms in seesaw formulas’. In: *JHEP* 04 (2011), p. 123. doi: 10.1007/JHEP04(2011)123. arXiv: 1102.3432 [hep-ph].
- [199] J. Schechter and J. W. F. Valle. ‘Comment on the lepton mixing matrix’. In: *Phys. Rev. D* 21 (1980), p. 309. doi: 10.1103/PhysRevD.21.309.
- [200] K. N. Abazajian et al. *Light sterile neutrinos: a white paper*. 2012. arXiv: 1204.5379 [hep-ph].
- [201] S. M. Bilenky. ‘Neutrino in Standard Model and beyond’. In: *Phys. Part. Nucl.* 46.4 (2015), pp. 475–496. doi: 10.1134/S1063779615040024. arXiv: 1501.00232 [hep-ph].

Bibliography

- [202] S. Weinberg. 'Baryon and lepton non-conserving processes'. In: *Phys. Rev. Lett.* 43 (1979), pp. 1566–1570. doi: 10.1103/PhysRevLett.43.1566.
- [203] S. M. Bilenky. 'Neutrino. History of a unique particle'. In: *Eur. Phys. J. H* 38 (2013), pp. 345–404. doi: 10.1140/epjh/e2012-20068-9. arXiv: 1210.3065 [hep-ph].
- [204] G. Danby et al. 'Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos'. In: *Phys. Rev. Lett.* 9 (1962), pp. 36–44. doi: 10.1103/PhysRevLett.9.36.
- [205] B. Pontecorvo. 'Electron and muon neutrinos'. In: *Sov. Phys. JETP* 10 (1960). [*Zh. Eksp. Teor. Fiz.* 37,1751(1959)], pp. 1236–1240.
- [206] M. A. Markov. *Hyperons and K-mesons*. Trans. from Russian. Washington, DC: ERDA Div. Phys. Res., 1959. URL: <http://cds.cern.ch/record/104118>.
- [207] M. Schwartz. 'Feasibility of using high-energy neutrinos to study the weak interactions'. In: *Phys. Rev. Lett.* 4 (1960), pp. 306–307. doi: 10.1103/PhysRevLett.4.306.
- [208] V. N. Gribov and B. Pontecorvo. 'Neutrino astronomy and lepton charge'. In: *Phys. Lett.* 28B (1969), p. 493. doi: 10.1016/0370-2693(69)90525-5.
- [209] S. M. Bilenky and B. Pontecorvo. 'Lepton mixing and neutrino oscillations'. In: *Phys. Rept.* 41 (1978), pp. 225–261. doi: 10.1016/0370-1573(78)90095-9.
- [210] E. Kh. Akhmedov and A. Yu. Smirnov. 'Paradoxes of neutrino oscillations'. In: *Phys. Atom. Nucl.* 72 (2009), pp. 1363–1381. doi: 10.1134/S1063778809080122. arXiv: 0905.1903 [hep-ph].
- [211] E. Kh. Akhmedov and J. Kopp. 'Neutrino oscillations: quantum mechanics vs. quantum field theory'. In: *JHEP* 04 (2010). [Erratum: *JHEP*10,052(2013)], p. 008. doi: 10.1007/JHEP04(2010)008. arXiv: 1001.4815 [hep-ph].
- [212] E. Kh. Akhmedov and A. Yu. Smirnov. 'Neutrino oscillations: entanglement, energy-momentum conservation, and QFT'. In: *Found. Phys.* 41 (2011), pp. 1279–1306. doi: 10.1007/s10701-011-9545-4. arXiv: 1008.2077 [hep-ph].
- [213] E. Kh. Akhmedov and A. Wilhelm. 'Quantum field theoretic approach to neutrino oscillations in matter'. In: *JHEP* 01 (2013), p. 165. doi: 10.1007/JHEP01(2013)165. arXiv: 1205.6231 [hep-ph].
- [214] C. Giunti. 'Energy and momentum of oscillating neutrinos'. In: *Mod. Phys. Lett.* A16 (2001), p. 2363. doi: 10.1142/S0217732301005801. arXiv: hep-ph/0104148 [hep-ph].
- [215] E. Kh. Akhmedov. 'Quantum-mechanics aspects and subtleties of neutrino oscillations'. In: *International Conference on History of the Neutrino: 1930-2018 Paris, France, September 5-7, 2018*. 2019. arXiv: 1901.05232 [hep-ph].
- [216] S. M. Bilenky, J. Hosek and S. T. Petcov. 'On oscillations of neutrinos with Dirac and Majorana masses'. In: *Phys. Lett.* 94B (1980), pp. 495–498. doi: 10.1016/0370-2693(80)90927-2.
- [217] M. Doi et al. 'CP violation in Majorana neutrinos'. In: *Phys. Lett.* 102B (1981), pp. 323–326. doi: 10.1016/0370-2693(81)90627-4.
- [218] G. Barenboim, C. A. Ternes and M. Tórtola. 'Neutrinos, DUNE, and the world best bound on CPT invariance'. In: *Phys. Lett.* B780 (2018), pp. 631–637. doi: 10.

Bibliography

- 1016/j.physletb.2018.03.060. arXiv: 1712.01714 [hep-ph].
- [219] G. Barenboim. ‘CPT violation in neutrinos’. In: *Proceedings, 1st Workshop on Phenomenology for Particle and Anti-Particle 2018 (PPAP2018): Hiroshima, Japan, March 6-8, 2018*. Kyoto: YITP, Kyoto University, 2018, pp. 49–51.
- [220] A. de Gouvêa and J. Jenkins. ‘The physical range of Majorana neutrino mixing parameters’. In: *Phys. Rev. D* 78 (2008), p. 053003. doi: 10.1103/PhysRevD.78.053003. arXiv: 0804.3627 [hep-ph].
- [221] J. Gluza and M. Zralek. ‘Parameters’ domain in three-flavour neutrino oscillations’. In: *Phys. Lett. B* 517 (2001), pp. 158–166. doi: 10.1016/S0370-2693(01)00962-5. arXiv: hep-ph/0106283 [hep-ph].
- [222] P. F. De Salas et al. ‘Neutrino mass ordering from oscillations and beyond: 2018 status and future prospects’. In: *Front. Astron. Space Sci.* 5 (2018), p. 36. doi: 10.3389/fspas.2018.00036. arXiv: 1806.11051 [hep-ph].
- [223] I. Esteban et al. ‘Updated constraints on non-standard interactions from global analysis of oscillation data’. In: *JHEP* 08 (2018). [Addendum: *JHEP* 12, 152 (2020)], p. 180. doi: 10.1007/JHEP08(2018)180. arXiv: 1805.04530 [hep-ph].
- [224] P. Coloma and T. Schwetz. ‘Generalised mass ordering degeneracy in neutrino oscillation experiments’. In: *Phys. Rev. D* 94,5 (2016). [Erratum: *Phys. Rev. D* 95,no.7,079903(2017)], p. 055005. doi: 10.1103/PhysRevD.94.055005. arXiv: 1604.05772 [hep-ph].
- [225] M. C. Gonzalez-Garcia, M. Maltoni and J. Salvado. ‘Testing matter effects in propagation of atmospheric and long-baseline neutrinos’. In: *JHEP* 05 (2011), p. 075. doi: 10.1007/JHEP05(2011)075. arXiv: 1103.4365 [hep-ph].
- [226] P. Bakhti and Y. Farzan. ‘Shedding light on LMA-Dark solar neutrino solution by medium-baseline reactor experiments: JUNO and RENO-50’. In: *JHEP* 07 (2014), p. 064. doi: 10.1007/JHEP07(2014)064. arXiv: 1403.0744 [hep-ph].
- [227] G. L. Fogli, E. Lisi and D. Montanino. ‘Matter-enhanced three-flavour oscillations and the solar neutrino problem’. In: *Phys. Rev. D* 54 (1996), pp. 2048–2062. doi: 10.1103/PhysRevD.54.2048. arXiv: hep-ph/9605273 [hep-ph].
- [228] A. de Gouvêa, A. Friedland and H. Murayama. ‘The dark side of the solar neutrino parameter space’. In: *Phys. Lett. B* 490 (2000), pp. 125–130. doi: 10.1016/S0370-2693(00)00989-8. arXiv: hep-ph/0002064 [hep-ph].
- [229] M. C. Gonzalez-Garcia and C. Peña-Garay. ‘On the size of the dark side of the solar neutrino parameter space’. In: *Phys. Rev. D* 62 (2000), p. 031301. doi: 10.1103/PhysRevD.62.031301. arXiv: hep-ph/0002186 [hep-ph].
- [230] A. de Gouvêa, J. Jenkins and B. Kayser. ‘Neutrino mass hierarchy, vacuum oscillations, and vanishing $|U_{e3}|$ ’. In: *Phys. Rev. D* 71 (2005), p. 113009. doi: 10.1103/PhysRevD.71.113009. arXiv: hep-ph/0503079 [hep-ph].
- [231] L. Wolfenstein. ‘Effects of matter on neutrino oscillations’. In: *American Institute of Physics Conference Series*. Vol. 52. American Institute of Physics Conference Series. 1979, pp. 108–112. doi: 10.1063/1.31797.

- [232] C. Y. Cardall and D. J. H. Chung. ‘The MSW effect in quantum field theory’. In: *Phys. Rev. D* 60 (1999), p. 073012. doi: 10.1103/PhysRevD.60.073012. arXiv: hep-ph/9904291 [hep-ph].
- [233] E. Kh. Akhmedov. ‘Neutrino oscillations in matter: from microscopic to macroscopic description’. In: *JHEP* 02 (2021), p. 107. doi: 10.1007/JHEP02(2021)107. arXiv: 2010.07847 [hep-ph].
- [234] R. R. Lewis. ‘Coherent detector for low-energy neutrinos’. In: *Phys. Rev. D* 21 (1980), p. 663. doi: 10.1103/PhysRevD.21.663.
- [235] J. Linder. *Derivation of neutrino matter potentials induced by Earth*. 2005. arXiv: hep-ph/0504264 [hep-ph].
- [236] E. Kh. Akhmedov et al. ‘T violation in neutrino oscillations in matter’. In: *Nucl. Phys.* B608 (2001), pp. 394–422. doi: 10.1016/S0550-3213(01)00261-9. arXiv: hep-ph/0105029 [hep-ph].
- [237] A. M. Dziewonski and D. L. Anderson. ‘Preliminary reference Earth model’. In: *Physics of the Earth and planetary interiors* 25.4 (1981), pp. 297–356. doi: 10.1016/0031-9201(81)90046-7.
- [238] J. N. Bahcall, M. C. Gonzalez-Garcia and C. Peña-Garay. ‘Before and after: How has the SNO neutral-current measurement changed things?’ In: *JHEP* 07 (2002), p. 054. doi: 10.1088/1126-6708/2002/07/054. arXiv: hep-ph/0204314.
- [239] M. Blennow and A. Yu. Smirnov. ‘Neutrino propagation in matter’. In: *Adv. High Energy Phys.* 2013 (2013), p. 972485. doi: 10.1155/2013/972485. arXiv: 1306.2903 [hep-ph].
- [240] E. Kh. Akhmedov. ‘Parametric resonance of neutrino oscillations and passage of solar and atmospheric neutrinos through the Earth’. In: *Nucl. Phys.* B538 (1999), pp. 25–51. doi: 10.1016/S0550-3213(98)00723-8. arXiv: hep-ph/9805272 [hep-ph].
- [241] E. Kh. Akhmedov. ‘Parametric resonance in neutrino oscillations in matter’. In: *Pramana* 54 (2000). Ed. by A. S. Joshipura and P. B. Pal, pp. 47–63. doi: 10.1007/s12043-000-0006-4. arXiv: hep-ph/9907435.
- [242] E. Kh. Akhmedov, M. Maltoni and A. Yu. Smirnov. ‘Describing oscillations of high-energy neutrinos in matter precisely’. In: *Phys. Rev. Lett.* 95 (2005), p. 211801. doi: 10.1103/PhysRevLett.95.211801. arXiv: hep-ph/0506064 [hep-ph].
- [243] L. Wolfenstein. ‘Possible identification of the outgoing lepton in neutral-current reactions’. In: *Nucl. Phys.* B91 (1975), p. 95. doi: 10.1016/0550-3213(75)90169-8.
- [244] M. B. Gavela et al. ‘Large gauge-invariant non-standard neutrino interactions’. In: *Phys. Rev. D* 79 (2009), p. 013007. doi: 10.1103/PhysRevD.79.013007. arXiv: 0809.3451 [hep-ph].
- [245] I. Bischer and W. Rodejohann. ‘General neutrino interactions from an effective field theory perspective’. In: *Nucl. Phys.* B947 (2019), p. 114746. doi: 10.1016/j.nuclphysb.2019.114746. arXiv: 1905.08699 [hep-ph].
- [246] E. Roulet. ‘MSW effect with flavour-changing neutrino interactions’. In: *Phys. Rev. D* 44 (1991), R935–R938. doi: 10.1103/PhysRevD.44.R935.
- [247] M. M. Guzzo, A. Masiero and S. T. Petcov. ‘On the MSW effect with massless neutrinos and no mixing in the vacuum’. In:

- Phys. Lett.* B260 (1991), pp. 154–160. doi: 10.1016/0370-2693(91)90984-X.
- [248] V. D. Barger, R. J. N. Phillips and K. Whisnant. ‘Solar neutrino solutions with matter-enhanced flavour-changing neutral-current scattering’. In: *Phys. Rev. D* 44 (1991), pp. 1629–1643. doi: 10.1103/PhysRevD.44.1629.
- [249] Y. Grossman. ‘Non-standard neutrino interactions and neutrino oscillation experiments’. In: *Phys. Lett. B* 359 (1995), pp. 141–147. doi: 10.1016/0370-2693(95)01069-3. arXiv: hep-ph/9507344.
- [250] T. Ohlsson. ‘Status of non-standard neutrino interactions’. In: *Rept. Prog. Phys.* 76 (2013), p. 044201. doi: 10.1088/0034-4885/76/4/044201. arXiv: 1209.2710 [hep-ph].
- [251] C. Biggio, M. Blennow and E. Fernández-Martínez. ‘General bounds on non-standard neutrino interactions’. In: *JHEP* 08 (2009), p. 090. doi: 10.1088/1126-6708/2009/08/090. arXiv: 0907.0097 [hep-ph].
- [252] T. Han et al. ‘Non-standard neutrino interactions at COHERENT, DUNE, T2HK and LHC’. In: *JHEP* 11 (2019), p. 028. doi: 10.1007/JHEP11(2019)028. arXiv: 1910.03272 [hep-ph].
- [253] A. Friedland and C. Lunardini. ‘Test of tau neutrino interactions with atmospheric neutrinos and K2K’. In: *Phys. Rev. D* 72 (2005), p. 053009. doi: 10.1103/PhysRevD.72.053009. arXiv: hep-ph/0506143 [hep-ph].
- [254] M. C. Gonzalez-Garcia and M. Maltoni. ‘Determination of matter potential from global analysis of neutrino oscillation data’. In: *JHEP* 09 (2013), p. 152. doi: 10.1007/JHEP09(2013)152. arXiv: 1307.3092 [hep-ph].
- [255] N. Fornengo et al. ‘Probing neutrino non-standard interactions with atmospheric neutrino data’. In: *Phys. Rev. D* 65 (2002), p. 013010. doi: 10.1103/PhysRevD.65.013010. arXiv: hep-ph/0108043 [hep-ph].
- [256] M. C. Gonzalez-Garcia et al. ‘Atmospheric neutrino observations and flavour-changing interactions’. In: *Phys. Rev. Lett.* 82 (1999), pp. 3202–3205. doi: 10.1103/PhysRevLett.82.3202. arXiv: hep-ph/9809531 [hep-ph].
- [257] Y. Farzan and I. M. Shoemaker. ‘Lepton-flavour-violating non-standard interactions via light mediators’. In: *JHEP* 07 (2016), p. 033. doi: 10.1007/JHEP07(2016)033. arXiv: 1512.09147 [hep-ph].
- [258] A. Yu. Smirnov and X.-J. Xu. ‘Wolfenstein potentials for neutrinos induced by ultra-light mediators’. In: *JHEP* 12 (2019), p. 046. doi: 10.1007/JHEP12(2019)046. arXiv: 1909.07505 [hep-ph].
- [259] P. Coloma, M. C. Gonzalez-Garcia and M. Maltoni. ‘Neutrino oscillation constraints on U(1) models: from non-standard interactions to long-range forces’. In: *JHEP* 01 (2021). [Erratum: *JHEP* 11, 115 (2022)], p. 114. doi: 10.1007/JHEP01(2021)114. arXiv: 2009.14220 [hep-ph].
- [260] I. Bischer, W. Rodejohann and X.-J. Xu. ‘Loop-induced neutrino non-standard interactions’. In: *JHEP* 10 (2018), p. 096. doi: 10.1007/JHEP10(2018)096. arXiv: 1807.08102 [hep-ph].
- [261] K. S. Babu et al. ‘Non-standard interactions in radiative neutrino mass models’. In: *JHEP* 03 (2020), p. 006. doi: 10.1007/JHEP03(2020)006. arXiv: 1907.09498 [hep-ph].

- [262] Y. Farzan and M. Tórtola. ‘Neutrino oscillations and non-standard interactions’. In: *Front.in Phys.* 6 (2018), p. 10. doi: 10.3389/fphy.2018.00010. arXiv: 1710.09360 [hep-ph].
- [263] S. Antusch, J. P. Baumann and E. Fernández-Martínez. ‘Non-standard neutrino interactions with matter from physics beyond the Standard Model’. In: *Nucl. Phys.* B810 (2009), pp. 369–388. doi: 10.1016/j.nuclphysb.2008.11.018. arXiv: 0807.1003 [hep-ph].
- [264] M. S. Bilenky and A. Santamaria. ‘One-loop effective Lagrangian for a Standard Model with a heavy charged scalar singlet’. In: *Nucl. Phys.* B420 (1994), pp. 47–93. doi: 10.1016/0550-3213(94)90375-1. arXiv: hep-ph/9310302 [hep-ph].
- [265] S. Bergmann, Y. Grossman and D. M. Pierce. ‘Can lepton flavour-violating interactions explain the atmospheric neutrino problem?’ In: *Phys. Rev.* D61 (2000), p. 053005. doi: 10.1103/PhysRevD.61.053005. arXiv: hep-ph/9909390 [hep-ph].
- [266] A. Zee. ‘A theory of lepton-number violation, neutrino Majorana mass, and oscillation’. In: *Phys. Lett.* 93B (1980). [Erratum: *Phys. Lett.* 95B,461(1980)], p. 389. doi: 10.1016/0370-2693(80)90349-4.
- [267] A. Zee. ‘Charged scalar field and quantum number violations’. In: *Phys. Lett.* 161B (1985), pp. 141–145. doi: 10.1016/0370-2693(85)90625-2.
- [268] K. R. S. Balaji, W. Grimus and T. Schwetz. ‘The solar LMA neutrino oscillation solution in the Zee model’. In: *Phys. Lett.* B508 (2001), pp. 301–310. doi: 10.1016/S0370-2693(01)00532-9. arXiv: hep-ph/0104035 [hep-ph].
- [269] Y. Cai et al. ‘From the trees to the forest: a review of radiative neutrino mass models’. In: *Front.in Phys.* 5 (2017), p. 63. doi: 10.3389/fphy.2017.00063. arXiv: 1706.08524 [hep-ph].
- [270] E. Fernández-Martínez. ‘Are there consistent models giving observable NSI?’ In: *J. Phys. Conf. Ser.* 408 (2013). Ed. by A. Blondel, I. Efthymiopoulos and G. Prior, p. 012031. doi: 10.1088/1742-6596/408/1/012031.
- [271] P. Langacker. ‘The physics of heavy Z' gauge bosons’. In: *Rev. Mod. Phys.* 81 (2009), pp. 1199–1228. doi: 10.1103/RevModPhys.81.1199. arXiv: 0801.1345 [hep-ph].
- [272] J. Heeck et al. ‘Non-standard neutrino interactions and neutral gauge bosons’. In: *SciPost Phys.* 6.3 (2019), p. 038. doi: 10.21468/SciPostPhys.6.3.038. arXiv: 1812.04067 [hep-ph].
- [273] P. B. Denton, Y. Farzan and I. M. Shoemaker. ‘Testing large non-standard neutrino interactions with arbitrary mediator mass after COHERENT data’. In: *JHEP* 07 (2018), p. 037. doi: 10.1007/JHEP07(2018)037. arXiv: 1804.03660 [hep-ph].
- [274] M. Cribier, J. Dumarchez and D. Vignaud, eds. *Proceedings of the International Conference on History of the Neutrino: 1930-2018*. AstroParticle and Cosmology Laboratory (APC), 2019.
- [275] M. C. Goodman. ‘Neutrino mistakes: wrong tracks and hints, hopes, and failures’. In: *International Conference on History of the Neutrino: 1930-2018 Paris, France, September 5-7, 2018*. 2019. arXiv: 1901.07068 [hep-ex].

Bibliography

- [276] M. Wurm. ‘Solar neutrino spectroscopy’. In: *Phys. Rept.* 685 (2017), pp. 1–52. doi: 10.1016/j.physrep.2017.04.002. arXiv: 1704.06331 [hep-ex].
- [277] E. Vitagliano, I. Tamborra and G. Raffelt. *Grand-unified neutrino spectrum at Earth*. 2019. arXiv: 1910.11878 [astro-ph.HE].
- [278] M. Maltoni and A. Yu. Smirnov. ‘Solar neutrinos and neutrino physics’. In: *Eur. Phys. J. A* 52.4 (2016), p. 87. doi: 10.1140/epja/i2016-16087-0. arXiv: 1507.05287 [hep-ph].
- [279] M. A. Markov and I. M. Zheleznykh. ‘On high-energy neutrino physics in cosmic rays’. In: *Nucl. Phys.* 27 (1961), pp. 385–394. doi: 10.1016/0029-5582(61)90331-5.
- [280] G. T. Zatsepin and V. A. Kuz’min. ‘Neutrino production in the atmosphere’. In: *Sov. Phys. JETP* 14 (6 1962), pp. 1294–1300.
- [281] M. Honda et al. ‘Atmospheric neutrino flux calculation using the NRLMSISE-00 atmospheric model’. In: *Phys. Rev. D* 92.2 (2015), p. 023004. doi: 10.1103/PhysRevD.92.023004. arXiv: 1502.03916 [astro-ph.HE].
- [282] G. D. Barr et al. ‘A three-dimensional calculation of atmospheric neutrinos’. In: *Phys. Rev. D* 70 (2004), p. 023006. doi: 10.1103/PhysRevD.70.023006. arXiv: astro-ph/0403630 [astro-ph].
- [283] A. Fedynitch et al. ‘A state-of-the-art calculation of atmospheric lepton fluxes’. In: *PoS ICRC2017* (2018), p. 1019. doi: 10.22323/1.301.1019.
- [284] P. Lipari. ‘Lepton spectra in the Earth’s atmosphere’. In: *Astropart. Phys.* 1 (1993), pp. 195–227. doi: 10.1016/0927-6505(93)90022-6.
- [285] P. Lipari and S. Vernetto. ‘The shape of the cosmic ray proton spectrum’. In: *Astropart. Phys.* 120 (2020), p. 102441. doi: 10.1016/j.astropartphys.2020.102441. arXiv: 1911.01311 [astro-ph.HE].
- [286] P. Gondolo, G. Ingelman and M. Thunman. ‘Charm production and high-energy atmospheric muon and neutrino fluxes’. In: *Astropart. Phys.* 5 (1996), pp. 309–332. doi: 10.1016/0927-6505(96)00033-3. arXiv: hep-ph/9505417 [hep-ph].
- [287] A. Fedynitch et al. ‘Hadronic interaction model Sibyll 2.3c and inclusive lepton fluxes’. In: *Phys. Rev. D* 100.10 (2019), p. 103018. doi: 10.1103/PhysRevD.100.103018. arXiv: 1806.04140 [hep-ph].
- [288] A. Bulmahn and M. H. Reno. ‘Secondary atmospheric tau-neutrino production’. In: *Phys. Rev. D* 82 (2010), p. 057302. doi: 10.1103/PhysRevD.82.057302. arXiv: 1007.4989 [hep-ph].
- [289] M. Honda et al. ‘Calculation of the flux of atmospheric neutrinos’. In: *Phys. Rev. D* 52 (1995), pp. 4985–5005. doi: 10.1103/PhysRevD.52.4985. arXiv: hep-ph/9503439 [hep-ph].
- [290] T. K. Gaisser and S. R. Klein. ‘A new contribution to the conventional atmospheric neutrino flux’. In: *Astropart. Phys.* 64 (2015), pp. 13–17. doi: 10.1016/j.astropartphys.2014.10.006. arXiv: 1409.4924 [astro-ph.HE].
- [291] D. Heck et al. *CORSIKA: a Monte Carlo code to simulate extensive air showers*. 1998. URL: https://web.iap.kit.edu/corsika/physics_description/corsika_phys.pdf (visited on 31/01/2024).

- [292] T. K. Gaisser and M. Honda. ‘Flux of atmospheric neutrinos’. In: *Ann. Rev. Nucl. Part. Sci.* 52 (2002), pp. 153–199. doi: 10.1146/annurev.nucl.52.050102.090645. arXiv: hep-ph/0203272 [hep-ph].
- [293] J. I. Illana et al. ‘Atmospheric lepton fluxes at very high energy’. In: *Astropart. Phys.* 34 (2011), pp. 663–673. doi: 10.1016/j.astropartphys.2011.01.001. arXiv: 1010.5084 [astro-ph.HE].
- [294] E. Kh. Akhmedov, M. Maltoni and A. Yu. Smirnov. ‘1–3 leptonic mixing and the neutrino oscillograms of the Earth’. In: *JHEP* 05 (2007), p. 077. doi: 10.1088/1126-6708/2007/05/077. arXiv: hep-ph/0612285 [hep-ph].
- [295] C. Rott, A. Taketa and D. Bose. ‘Spectrometry of the Earth using neutrino oscillations’. In: *Sci. Rep.* 5 (2015), p. 15225. doi: 10.1038/srep15225. arXiv: 1502.04930 [physics.geo-ph].
- [296] O. L. G. Peres and A. Yu. Smirnov. ‘Atmospheric neutrinos: LMA oscillations, U_{e3} -induced interference and CP violation’. In: *Nucl. Phys.* B680 (2004), pp. 479–509. doi: 10.1016/j.nuclphysb.2003.12.017. arXiv: hep-ph/0309312 [hep-ph].
- [297] O. L. G. Peres and A. Yu. Smirnov. ‘Oscillations of very low energy atmospheric neutrinos’. In: *Phys. Rev.* D79 (2009), p. 113002. doi: 10.1103/PhysRevD.79.113002. arXiv: 0903.5323 [hep-ph].
- [298] E. Kh. Akhmedov et al. ‘Atmospheric neutrinos at Super-Kamiokande and parametric resonance in neutrino oscillations’. In: *Nucl. Phys.* B542 (1999), pp. 3–30. doi: 10.1016/S0550-3213(98)00825-6. arXiv: hep-ph/9808270 [hep-ph].
- [299] M. Shaevitz. ‘New compact low-energy neutrino source using isotope beta decay’. Presentation given at Coherent Neutrino Scattering 2012 Workshop. Livermore, California, USA. URL: <https://neutrinos.llnl.gov/workshops/coherent-neutrino-scattering-2012> (visited on 11/02/2024).
- [300] O. G. Miranda, G. Sanchez Garcia and O. Sanders. ‘Coherent elastic neutrino-nucleus scattering as a precision test for the Standard Model and beyond: the COHERENT proposal case’. In: *Adv. High Energy Phys.* 2019 (2019), p. 3902819. doi: 10.1155/2019/3902819. arXiv: 1902.09036 [hep-ph].
- [301] J. Spitz. ‘Cross-section measurements with monoenergetic muon neutrinos’. In: *Phys. Rev.* D89.7 (2014), p. 073007. doi: 10.1103/PhysRevD.89.073007. arXiv: 1402.2284 [physics.ins-det].
- [302] R. Harnik, K. J. Kelly and P. A. N. Machado. ‘Prospects of measuring oscillated decay-at-rest neutrinos at long baselines’. In: *Phys. Rev. D* 101.3 (2020), p. 033008. doi: 10.1103/PhysRevD.101.033008. arXiv: 1911.05088 [hep-ph].
- [303] P. Huber and P. Jaffke. ‘Neutron capture and the antineutrino yield from nuclear reactors’. In: *Phys. Rev. Lett.* 116.12 (2016), p. 122503. doi: 10.1103/PhysRevLett.116.122503. arXiv: 1510.08948 [hep-ph].
- [304] X. Qian and J.-C. Peng. ‘Physics with reactor neutrinos’. In: *Rept. Prog. Phys.* 82.3 (2019), p. 036201. doi: 10.1088/1361-6633/aae881. arXiv: 1801.05386 [hep-ex].
- [305] F. Capozzi et al. ‘Unfinished fabric of the three-neutrino paradigm’. In: *Phys. Rev. D* 104.8 (2021), p. 083031. doi: 10.1103/PhysRevD.104.083031. arXiv: 2107.00532 [hep-ph].
- [306] *NuFIT group neutrino fits*. URL: <http://www.nu-fit.org> (visited on 17/06/2024).

Bibliography

- [307] M. C. Gonzalez-Garcia, M. Maltoni and T. Schwetz. ‘NuFIT: three-flavour global analyses of neutrino oscillation experiments’. In: *Universe* 7.12 (2021), p. 459. doi: 10 . 3390 / universe7120459. arXiv: 2111.03086 [hep-ph].
- [308] I. Esteban et al. ‘The fate of hints: updated global analysis of three-flavour neutrino oscillations’. In: *JHEP* 09 (2020), p. 178. doi: 10 . 1007 / JHEP09(2020) 178. arXiv: 2007.14792 [hep-ph].
- [309] *NuFIT 5.3 (2024)*. URL: <http://www.nu-fit.org> (visited on 17/06/2024).
- [310] *Valencia group neutrino fits*. URL: <https://globalfit.astroparticles.es> (visited on 17/06/2024).
- [311] P. F. de Salas et al. ‘2020 global reassessment of the neutrino oscillation picture’. In: *JHEP* 02 (2021), p. 071. doi: 10.1007/JHEP02(2021) 071. arXiv: 2006 . 11237 [hep-ph].
- [312] P. Coloma et al. ‘Improved global fit to non-standard neutrino interactions using COHERENT energy and timing data’. In: *JHEP* 02 (2020), p. 023. doi: 10 . 1007 / JHEP02(2020) 023. arXiv: 1911 . 09109 [hep-ph].
- [313] *NuFIT 5.2 (2022)*. URL: <http://www.nu-fit.org> (visited on 21/09/2023).
- [314] I. 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: *JHEP* 01 (2019), p. 106. doi: 10 . 1007 / JHEP01(2019) 106. arXiv: 1811 . 05487 [hep-ph].
- [315] *NuFIT 4.1 (2019)*. URL: <http://www.nu-fit.org> (visited on 21/09/2023).
- [316] M. C. Gonzalez-Garcia and C. Peña-Garay. ‘Three-neutrino mixing after the first results from K2K and KamLAND’. In: *Phys. Rev. D* 68 (2003), p. 093003. doi: 10 . 1103 / PhysRevD . 68 . 093003. arXiv: hep-ph/0306001 [hep-ph].
- [317] J. W. F. Valle. ‘Resonant oscillations of massless neutrinos in matter’. In: *Phys. Lett. B* 199 (1987), pp. 432–436. doi: 10 . 1016/0370-2693(87)90947-6.
- [318] M. Fukugita and T. Yanagida. ‘Neutrino oscillation in matter induced by charged scalar particles’. In: *Phys. Lett. B* 206 (1988), pp. 93–96. doi: 10 . 1016 / 0370-2693(88)91268-3.
- [319] P. I. Krastev and J. N. Bahcall. ‘FCNC solutions to the solar neutrino problem’. In: *Flavour-changing neutral currents: present and future studies. Proceedings, Symposium, FCNC’97, Santa Monica, USA, February 19-21, 1997*. 1997, pp. 259–263. arXiv: hep-ph/9703267 [hep-ph].
- [320] S. Bergmann. ‘The solar neutrino problem in the presence of flavour-changing neutrino interactions’. In: *Nucl. Phys. B* 515 (1998), pp. 363–383. doi: 10 . 1016 / S0550-3213(97)00794-3. arXiv: hep-ph/9707398 [hep-ph].
- [321] S. Bergmann et al. ‘Status of the solution to the solar neutrino problem based on non-standard neutrino interactions’. In: *Phys. Rev. D* 62 (2000), p. 073001. doi: 10 . 1103 / PhysRevD . 62 . 073001. arXiv: hep-ph/0004049 [hep-ph].
- [322] E. Ma and P. Roy. ‘Neutrino oscillation data versus three light neutrino flavours’. In: *Phys. Rev. Lett.* 80 (1998), pp. 4637–4640. doi: 10 . 1103 / PhysRevLett . 80 . 4637. arXiv: hep-ph/9706309 [hep-ph].

- [323] P. Lipari and M. Lusignoli. ‘On exotic solutions of the atmospheric neutrino problem’. In: *Phys. Rev. D* 60 (1999), p. 013003. doi: 10.1103/PhysRevD.60.013003. arXiv: hep-ph/9901350 [hep-ph].
- [324] N. Fornengo, M. C. Gonzalez-Garcia and J. W. F. Valle. ‘On the interpretation of the atmospheric neutrino data in terms of flavour-changing neutrino interactions’. In: *JHEP* 07 (2000), p. 006. doi: 10.1088/1126-6708/2000/07/006. arXiv: hep-ph/9906539 [hep-ph].
- [325] A. Friedland, C. Lunardini and M. Maltoni. ‘Atmospheric neutrinos as probes of neutrino-matter interactions’. In: *Phys. Rev. D* 70 (2004), p. 111301. doi: 10.1103/PhysRevD.70.111301. arXiv: hep-ph/0408264 [hep-ph].
- [326] *NuFIT 4.0 (2018)*. URL: <http://www.nu-fit.org> (visited on 21/09/2023).
- [327] I. Esteban et al. ‘Updated fit to three-neutrino mixing: exploring the accelerator-reactor complementarity’. In: *J. High Energy Phys.* 2017.1, 87 (2017), p. 87. doi: 10.1007/JHEP01(2017)087.
- [328] *NuFIT 3.2 (2018)*. URL: <http://www.nu-fit.org> (visited on 21/09/2023).
- [329] M. Blennow and T. Ohlsson. ‘Approximative two-flavour framework for neutrino oscillations with non-standard interactions’. In: *Phys. Rev. D* 78 (2008), p. 093002. doi: 10.1103/PhysRevD.78.093002. arXiv: 0805.2301 [hep-ph].
- [330] P. Askebjerg et al. ‘Optical properties of the South Pole ice at depths between 0.8 km and 1 km’. In: *Science* 267 (1995), pp. 1147–1150. doi: 10.1126/science.267.5201.1147. arXiv: astro-ph/9412028.
- [331] R. Abbasi et al. ‘In situ estimation of ice crystal properties at the South Pole using LED calibration data from the IceCube Neutrino Observatory’. In: *The Cryosphere* 18.1 (2024), pp. 75–102. doi: 10.5194/tc-18-75-2024.
- [332] S. G. Warren and R. E. Brandt. ‘Optical constants of ice from the ultraviolet to the microwave: a revised compilation’. In: *Journal of Geophysical Research: Atmospheres* 113.D14 (2008). doi: 10.1029/2007JD009744.
- [333] L. Rädcl and C. Wiebusch. ‘Calculation of the Cherenkov light yield from electromagnetic cascades in ice with Geant4’. In: *Astropart. Phys.* 44 (2013), pp. 102–113. doi: 10.1016/j.astropartphys.2013.01.015. arXiv: 1210.5140 [astro-ph.IM].
- [334] J. V. Jelley. ‘Cherenkov radiation and its applications’. In: *British Journal of Applied Physics* 6.7 (1955), pp. 227–232. doi: 10.1088/0508-3443/6/7/301.
- [335] D. Chirkin and W. Rhode. *Muon Monte Carlo: a high-precision tool for muon propagation through matter*. 2004. arXiv: hep-ph/0407075.
- [336] D. E. Groom. *Atomic and nuclear properties of materials for more than 350 materials*. URL: http://pdg.lbl.gov/2019/AtomicNuclearProperties/HTML/water_ice.html (visited on 18/01/2024).
- [337] W. Heitler. *The quantum theory of radiation*. Vol. 5. International Series of Monographs on Physics. Oxford: Oxford University Press, 1936.
- [338] J. F. Carlson and J. R. Oppenheimer. ‘On multiplicative showers’. In: *Phys. Rev.* 51 (1937), pp. 220–231. doi: 10.1103/PhysRev.51.220.

Bibliography

- [339] L. Michel. 'Interaction between four half-spin particles and the decay of the μ meson'. In: *Proc. Phys. Soc. A* 63 (1950). Ed. by T. Damour, I. Todorov and B. Zhilinskii, pp. 514–531. doi: 10.1088/0370-1298/63/5/311.
- [340] T. A. Gabriel et al. 'Energy dependence of hadronic activity'. In: *Nucl. Instrum. Meth. A* 338 (1994), pp. 336–347. doi: 10.1016/0168-9002(94)91317-X.
- [341] M. Kowalski. *On the Cherenkov light emission of hadronic and electromagnetic cascades*. Tech. rep. AMANDA-IR/20020803. 2002. URL: <http://internal.icecube.wisc.edu/reports/amanda/data/20020803-track.pdf>.
- [342] D. E. Groom. 'Energy flow in a hadronic cascade: application to hadron calorimetry'. In: *Nucl. Instrum. Meth. A* 572 (2007). [Erratum: *Nucl. Instrum. Meth. A* 593, 628 (2008)], pp. 633–653. doi: 10.1016/j.nima.2006.11.070. arXiv: physics/0605164.
- [343] 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.
- [344] S. G. Warren. 'Optical properties of ice and snow'. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 377.2146 (2019), p. 20180161. doi: 10.1098/rsta.2018.0161.
- [345] P. Askebjør et al. 'Optical properties of deep ice at the South Pole: absorption'. In: *Appl. Opt.* 36 (1997), pp. 4168–4180. doi: 10.1364/AO.36.004168. arXiv: physics/9701025.
- [346] M. G. Aartsen et al. 'South Pole glacial climate reconstruction from multi-borehole laser particulate stratigraphy'. In: *J. Glaciol.* 59.218 (2013), pp. 1117–1128. doi: 10.3189/2013JGlG13J068.
- [347] P. B. Price and L. Bergström. 'Optical properties of deep ice at the South Pole: scattering'. In: *Appl. Opt.* 36.18 (1997), pp. 4181–4194. doi: 10.1364/AO.36.004181.
- [348] C. V. Raman. 'Thermal opalescence in crystals and the colour of ice in glaciers'. In: *Nature* 111.2775 (1923), pp. 13–14.
- [349] W. D. Bancroft. 'The colour of water'. In: *Journal of the Franklin Institute* 187.4 (1919), pp. 459–485. doi: 10.1016/S0016-0032(19)91151-3.
- [350] C. F. Bohren. 'Understanding colours in nature'. In: *Pigment Cell Research* 1.4 (1988), pp. 214–222. doi: 10.1111/j.1600-0749.1988.tb00419.x.
- [351] V. F. Petrenko and R. W. Whitworth. *Physics of ice*. Oxford: Oxford University Press, 1999.
- [352] D. Sahoo, A. K. Arora and R. Kesava-moorthy. 'Elasto-optic contribution to light scattering from dislocations in insulating crystals'. In: *Journal of Physics C Solid State Physics* 16.9 (1983), pp. 1687–1692. doi: 10.1088/0022-3719/16/9/013.
- [353] Y. D. He and P. B. Price. 'Remote sensing of dust in deep ice at the South Pole'. In: *Journal of Geophysical Research: Atmospheres* 103.D14 (1998), pp. 17041–17056. doi: 10.1029/98JD01643.
- [354] M. I. Mishchenko, J. W. Hovenier and L. D. Travis, eds. *Light scattering by non-spherical particles: theory, measurements, and applications*. San Diego, CA: Academic Press, 1999.

- [355] M. G. Aartsen et al. ‘Measurement of South Pole ice transparency with the IceCube LED calibration system’. In: *Nucl. Instrum. Meth.* A711 (2013), pp. 73–89. doi: 10.1016/j.nima.2013.01.054. arXiv: 1301.5361 [astro-ph.IM].
- [356] IceCube collaboration. *AHA ice model*. URL: https://wiki.icecube.wisc.edu/index.php/AHA_ice_model (visited on 18/01/2024).
- [357] D. Chirkin. ‘Evidence of optical anisotropy of the South Pole ice’. In: *International Cosmic Ray Conference*. Vol. 33. International Cosmic Ray Conference. 2013, p. 3338.
- [358] M. Rongen. ‘Calibration of the IceCube Neutrino Observatory’. PhD thesis. 2019. doi: 10.18154/RWTH-2019-09941. arXiv: 1911.02016 [astro-ph.IM].
- [359] D. Chirkin and M. Rongen. ‘Light diffusion in birefringent polycrystals and the IceCube ice anisotropy’. In: *PoS ICRC2019* (2020), p. 854. doi: 10.22323/1.358.0854. arXiv: 1908.07608 [astro-ph.HE].
- [360] R. Abbasi et al. ‘Calibration and characterisation of the IceCube photomultiplier tube’. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 618.1 (2010), pp. 139–152. doi: 10.1016/j.nima.2010.03.102.
- [361] US National Science Foundation (NSF). *Map of the South Pole Station and attached science instruments*. URL: https://www.nsf.gov/news/special_reports/livingsouthpole/pdf_files/instrument_map.pdf (visited on 18/01/2024).
- [362] M. G. Aartsen et al. ‘Measurement of atmospheric tau neutrino appearance with IceCube DeepCore’. In: *Phys. Rev.* D99.3 (2019), p. 032007. doi: 10.1103/PhysRevD.99.032007. arXiv: 1901.05366 [hep-ex].
- [363] R. Abbasi et al. ‘The IceCube data acquisition system: signal capture, digitization, and timestamping’. In: *Nucl. Instrum. Meth.* A601 (2009), pp. 294–316. doi: 10.1016/j.nima.2009.01.001. arXiv: 0810.4930 [physics.ins-det].
- [364] M. G. Aartsen et al. ‘The IceCube Neutrino Observatory: instrumentation and online systems’. In: *JINST* 12.03 (2017), P03012. doi: 10.1088/1748-0221/12/03/P03012. arXiv: 1612.05093 [astro-ph.IM].
- [365] T. Ehrhardt. ‘Studies of the reconstruction of cascade-like events in PINGU’. Master thesis. Rheinische Friedrich-Wilhelms-Universität Bonn, 2014.
- [366] N. Whitehorn. ‘A search for high-energy neutrino emission from gamma bursts’. PhD thesis. University of Wisconsin-Madison, 2012.
- [367] IceCube collaboration. *WaveCalibrator*. URL: <https://software.icecube.wisc.edu/icetray/main/projects/WaveCalibrator> (visited on 29/04/2024).
- [368] IceCube collaboration. *Droop correction*. URL: https://wiki.icecube.wisc.edu/index.php/Droop_correction (visited on 29/04/2024).
- [369] M. G. Aartsen et al. ‘Energy reconstruction methods in the IceCube Neutrino Telescope’. In: *JINST* 9 (2014), P03009. doi: 10.1088/1748-0221/9/03/P03009. arXiv: 1311.4767 [physics.ins-det].
- [370] IceCube collaboration. *DeepCore 2012 filter proposal*. URL: <https://docushare.icecube.wisc.edu/dsweb/Get/Document-59397/DeepCoreFilterProposal.pdf> (visited on 01/05/2020).

Bibliography

- [371] IceCube collaboration. *DeepCore 2013 filter proposal*. URL: https://docushare.icecube.wisc.edu/dsweb/Get/Document-62751/filter_proposal3.pdf (visited on 01/05/2020).
- [372] S. Bravo and L. Norris. *IceCube looks to the future with PINGU*. URL: <https://icecube.wisc.edu/news/view/186> (visited on 18/01/2024).
- [373] US National Science Foundation (NSF). *IceCube Gen2 Phase 1: an IceCube extension for precision neutrino physics and astrophysics*. URL: https://www.nsf.gov/awardsearch/showAward?AWD_ID=1719277 (visited on 18/01/2024).
- [374] IceCube collaboration. *NSF mid-scale award sets off the first extension of IceCube*. URL: <https://icecube.wisc.edu/news/view/661> (visited on 18/01/2024).
- [375] J. van Santen. ‘IceCube-Gen2: science case and strategy’. Presentation given at Mediterranean Antarctic Neutrino Telescope Symposium (MANTS) 2016. Mainz, Germany.
- [376] IceCube collaboration. *Deployment order*. URL: https://wiki.icecube.wisc.edu/index.php/Deployment_order (visited on 29/04/2024).
- [377] R. Abbasi et al. ‘D-Egg: a dual PMT optical module for IceCube’. In: *JINST* 18.04 (2023), P04014. DOI: 10.1088/1748-0221/18/04/P04014. arXiv: 2212.14526 [astro-ph.IM].
- [378] A. Ishihara. ‘The IceCube Upgrade – design and science goals’. In: *PoS ICRC2019* (2020), p. 1031. DOI: 10.22323/1.358.1031. arXiv: 1908.09441 [astro-ph.HE].
- [379] R. Abbasi et al. ‘Design and performance of the multi-PMT optical module for IceCube Upgrade’. In: *PoS ICRC2021* (2021), p. 1070. DOI: 10.22323/1.395.1070. arXiv: 2107.11383 [astro-ph.IM].
- [380] R. Abbasi et al. *Acceptance tests of more than 10 000 photomultiplier tubes for the multi-PMT digital optical modules of the IceCube Upgrade*. Apr. 2024. arXiv: 2404.19589 [astro-ph.IM].
- [381] P. Eller et al. ‘Sensitivity of the IceCube Upgrade to atmospheric neutrino oscillations’. In: *PoS ICRC2023* (2023), p. 1036. DOI: 10.22323/1.444.1036. arXiv: 2307.15295 [astro-ph.HE].
- [382] R. Wendell. *Prob3++ software for computing three-flavour neutrino oscillation probabilities*. URL: <http://webhome.phy.duke.edu/~raw22/public/Prob3++> (visited on 18/01/2024).
- [383] A. Esmaili and A. Yu. Smirnov. ‘Probing non-standard interaction of neutrinos with IceCube and DeepCore’. In: *JHEP* 06 (2013), p. 026. DOI: 10.1007/JHEP06(2013)026. arXiv: 1304.1042 [hep-ph].
- [384] J. Kopp et al. ‘Non-standard neutrino interactions in reactor and superbeam experiments’. In: *Phys. Rev. D* 77 (2008), p. 013007. DOI: 10.1103/PhysRevD.77.013007. arXiv: 0708.0152 [hep-ph].
- [385] T. Kikuchi, H. Minakata and S. Uchinami. ‘Perturbation theory of neutrino oscillation with non-standard neutrino interactions’. In: *JHEP* 03 (2009), p. 114. DOI: 10.1088/1126-6708/2009/03/114. arXiv: 0809.3312 [hep-ph].
- [386] T. Ohlsson, H. Zhang and S. Zhou. ‘Effects of non-standard neutrino interactions at PINGU’. In: *Phys. Rev. D* 88.1 (2013), p. 013001. DOI: 10.1103/

- PhysRevD. 88. 013001. arXiv: 1303. 6130 [hep-ph].
- [387] A. Chatterjee et al. ‘Testing non-standard neutrino-matter interactions in atmospheric neutrino propagation’. In: *Phys. Rev. D* 93.9 (2016), p. 093017. DOI: 10.1103/PhysRevD.93.093017. arXiv: 1409.8472 [hep-ph].
- [388] W.-J. Feng et al. ‘Non-standard interactions versus planet-scale neutrino oscillations’. In: *Phys. Rev. D* 100.11 (2019), p. 115034. DOI: 10.1103/PhysRevD.100.115034. arXiv: 1909.12674 [hep-ph].
- [389] IceCube collaboration. *IceCube software documentation*. URL: <https://software.icecube.wisc.edu> (visited on 18/01/2024).
- [390] F. James. ‘Monte Carlo theory and practice’. In: *Rept. Prog. Phys.* 43 (1980), p. 1145. DOI: 10.1088/0034-4885/43/9/002.
- [391] M. G. Dunkman. ‘Measurement of atmospheric muon neutrino disappearance with IceCube DeepCore’. PhD thesis. Pennsylvania State University, 2015.
- [392] S. Wren. ‘Neutrino mass ordering studies with IceCube-DeepCore’. PhD thesis. The University of Manchester, 2018.
- [393] A. Terliuk. ‘Measurement of atmospheric neutrino oscillations and search for sterile neutrino mixing with IceCube DeepCore’. PhD thesis. Humboldt-Universität zu Berlin, 2018.
- [394] F. Huang. ‘Measurement of atmospheric tau neutrino appearance with IceCube/Deepcore’. PhD thesis. Pennsylvania State University, 2018.
- [395] T. R. Wood. ‘Characterising the atmospheric neutrino spectrum with the IceCube Neutrino Observatory’. PhD thesis. University of Alberta, 2018.
- [396] C. Andreopoulos et al. ‘The GENIE neutrino Monte Carlo generator’. In: *Nucl. Instrum. Meth. A* 614 (2010), pp. 87–104. DOI: 10.1016/j.nima.2009.12.009. arXiv: 0905.2517 [hep-ph].
- [397] C. Andreopoulos et al. *The GENIE neutrino Monte Carlo generator: physics and user manual*. 2015. arXiv: 1510.05494 [hep-ph].
- [398] M. Alam et al. *GENIE production release 2.10.0*. 2015. arXiv: 1512.06882 [hep-ph].
- [399] J.-H. Koehne et al. ‘PROPOSAL: a tool for propagation of charged leptons’. In: *Computer Physics Communications* 184 (2013), pp. 2070–2090. DOI: 10.1016/j.cpc.2013.04.001.
- [400] S. Agostinelli et al. ‘GEANT4: a simulation toolkit’. In: *Nucl. Instrum. Meth. A* 506 (2003), pp. 250–303. DOI: 10.1016/S0168-9002(03)01368-8.
- [401] J. Allison et al. ‘Geant4 developments and applications’. In: *IEEE Trans. Nucl. Sci.* 53 (2006), p. 270. DOI: 10.1109/TNS.2006.869826.
- [402] J. Allison et al. ‘Recent developments in Geant4’. In: *Nucl. Instrum. Meth. A* 835 (2016), pp. 186–225. DOI: 10.1016/j.nima.2016.06.125.
- [403] IceCube collaboration. *CLsim*. URL: <https://software.icecube.wisc.edu/icetray/main/projects/clsim> (visited on 29/04/2024).
- [404] J. E. Stone, D. Gohara and G. Shi. ‘OpenCL: a parallel programming standard for heterogeneous computing systems’. In: *Computing in Science Engineering* 12.3 (2010), pp. 66–73. DOI: 10.1109/MCSE.2010.69.

Bibliography

- [405] IceCube collaboration. *Ice models*. URL: https://wiki.icecube.wisc.edu/index.php/Ice_models (visited on 18/01/2024).
- [406] IceCube collaboration. *MSU forward hole ice*. URL: https://wiki.icecube.wisc.edu/index.php/MSU_Forward_Hole_Ice (visited on 27/12/2023).
- [407] O. W. Richardson. *On the negative radiation from hot platinum*. Vol. 11. 1901, pp. 286–295.
- [408] K. Helbing et al. *Light emission in Amanda pressure spheres*. Amanda internal report AMANDA-IR/20030701. 2003. URL: https://internal-apps.icecube.wisc.edu/reports/amanda/data/20030701-Scinti_Intern_Docu_submit.pdf.
- [409] H. O. Meyer. ‘Spontaneous electron emission from a cold surface’. In: *EPL (Europhysics Letters)* 89.5 (2010), p. 58001. DOI: 10.1209/0295-5075/89/58001.
- [410] D. Heeremann. ‘HitSpooling: an improvement for the Supernova neutrino detection system in IceCube’. PhD thesis. Université Libre de Bruxelles, 2015.
- [411] IceCube collaboration. *Vuvuzela*. URL: <https://software.icecube.wisc.edu/icetray/main/projects/vuvuzela> (visited on 29/04/2024).
- [412] M. J. Larson. ‘A search for tau neutrino appearance with IceCube-DeepCore’. PhD thesis. Københavns Universitet, 2018.
- [413] IceCube collaboration. *PMTResponseSimulator*. URL: <https://software.icecube.wisc.edu/icetray/main/projects/DOMLauncher/PMTRes.html> (visited on 29/04/2024).
- [414] S. N. G. Axani. ‘Sterile-neutrino searches at the IceCube Neutrino Observatory’. PhD thesis. MIT, 2019. arXiv: 2003.02796 [hep-ex].
- [415] M. G. Aartsen et al. ‘In-situ calibration of the single-photoelectron charge response of the IceCube photomultiplier tubes’. In: *JINST* 15.06 (2020), P06032. DOI: 10.1088/1748-0221/15/06/P06032. arXiv: 2002.00997 [physics.ins-det].
- [416] IceCube collaboration. *DOMLauncher*. URL: <https://software.icecube.wisc.edu/icetray/main/projects/DOMLauncher/DOML.html> (visited on 29/04/2024).
- [417] M. Usner. ‘Search for neutrino-induced cascades in IceCube’. Master thesis. Rheinische Friedrich-Wilhelms-Universität Bonn, 2012.
- [418] J. Hofestädt. ‘Measuring the neutrino mass hierarchy with the future KM3NeT/ORCA detector’. PhD thesis. Friedrich-Alexander-Universität Erlangen-Nürnberg, 2017.
- [419] M. Leuermann. ‘Testing the neutrino mass ordering with IceCube DeepCore’. PhD thesis. Rheinisch-Westfälische Technische Hochschule Aachen, 2018.
- [420] D. Chirkin. *Study of ice transparency with IceCube flashers*. IceCube internal report. 2009. URL: <https://icecube.wisc.edu/~dima/work/WISC/ppc/fit/paper/a.pdf>.
- [421] F. Feroz and M. P. Hobson. ‘Multimodal nested sampling: an efficient and robust alternative to MCMC methods for astronomical data analysis’. In: *Mon. Not. Roy. Astron. Soc.* 384 (2008), p. 449. DOI: 10.1111/j.1365-2966.2007.12353.x. arXiv: 0704.3704 [astro-ph].

- [422] F. Feroz, M. P. Hobson and M. Bridges. ‘MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics’. In: *Mon. Not. Roy. Astron. Soc.* 398 (2009), pp. 1601–1614. DOI: 10.1111/j.1365-2966.2009.14548.x. arXiv: 0809.3437 [astro-ph].
- [423] F. Feroz et al. ‘Importance nested sampling and the MultiNest algorithm’. In: *Open J. Astrophys.* 2.1 (2019), p. 10. DOI: 10.21105/astro.1306.2144. arXiv: 1306.2144 [astro-ph.IM].
- [424] M. Hieronymus. ‘Reconstruction of low-energy neutrino events with GPUs at IceCube’. Master thesis. Johannes Gutenberg-Universität Mainz, 2019.
- [425] IceCube collaboration. *IC86 muon disappearance analysis*. URL: https://wiki.icecube.wisc.edu/index.php/IC86_Muon_Disappearance_Analysis (visited on 18/01/2024).
- [426] M. G. Aartsen et al. ‘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].
- [427] M. G. Aartsen et al. ‘Development of an analysis to probe the neutrino mass ordering with atmospheric neutrinos using three years of IceCube DeepCore data’. In: *Eur. Phys. J. C* 80.1 (2020), p. 9. DOI: 10.1140/epjc/s10052-019-7555-0. arXiv: 1902.07771 [hep-ex].
- [428] J. P. Yáñez Garza. ‘Measurement of neutrino oscillations in atmospheric neutrinos with the IceCube DeepCore detector’. PhD thesis. Humboldt-Universität zu Berlin, 2014.
- [429] S. Euler. ‘Observations of oscillations of atmospheric neutrinos with the IceCube Neutrino Observatory’. PhD thesis. Rheinisch-Westfälische Technische Hochschule Aachen, 2013.
- [430] J. M. Picone et al. ‘NRLMSISE-00 empirical model of the atmosphere: statistical comparison and scientific issues’. In: *Journal of Geophysical Research* 107 (2002). DOI: 10.1029/2002JA009430.
- [431] Y. Becherini et al. ‘A parameterisation of single and multiple muons in the deep water or ice’. In: *Astropart. Phys.* 25 (2006), pp. 1–13. DOI: 10.1016/j.astropartphys.2005.10.005. arXiv: hep-ph/0507228 [hep-ph].
- [432] G. D. Barr et al. ‘Uncertainties in atmospheric neutrino fluxes’. In: *Phys. Rev. D* 74 (2006), p. 094009. DOI: 10.1103/PhysRevD.74.094009. arXiv: astro-ph/0611266 [astro-ph].
- [433] J. P. Yáñez Garza. Presentation given at IceCube oscillations call, December 14, 2015. URL: https://wiki.icecube.wisc.edu/index.php/Low_Energy_and_Oscillations_Phone_Meetings (visited on 16/02/2024).
- [434] J. P. Yáñez Garza. IceCube internal communication.
- [435] A. Bodek, I. Park and U.-k. Yang. ‘Improved low- Q^2 model for neutrino- and electron-nucleon cross sections in few-GeV region’. In: *Nucl. Phys. B Proc. Suppl.* 139 (2005). Ed. by F. Cavanna et al., pp. 113–118. DOI: 10.1016/j.nuclphysbps.2004.11.208. arXiv: hep-ph/0411202.
- [436] M. Malinsky, T. Ohlsson and H. Zhang. ‘Non-standard neutrino interactions from a triplet seesaw model’. In: *Phys. Rev. D* 79 (2009), p. 011301. DOI: 10.1103/

Bibliography

- PhysRevD . 79 . 011301. arXiv: 0811 . 3346 [444] S. Baker and R. D. Cousins. ‘Clarification of the use of chi-square and likelihood functions in fits to histograms’. In: *Nuclear Instruments and Methods in Physics Research* 221.2 (1984), pp. 437–442. doi: 10 . 1016/0167-5087(84)90016-4.
- [437] T. Ohlsson, T. Schwetz and H. Zhang. ‘Non-standard neutrino interactions in the Zee–Babu model’. In: *Phys. Lett. B* 681 (2009), pp. 269–275. doi: 10 . 1016 / j . physletb . 2009 . 10 . 025. arXiv: 0909 . 0455 [hep-ph]. [445] R. D. Cousins. *Lectures on statistics in theory: prelude to statistics in practice*. 2018. arXiv: 1807.05996 [physics.data-an].
- [438] J. Salvado et al. ‘Non-standard interactions with high-energy atmospheric neutrinos at IceCube’. In: *JHEP* 01 (2017), p. 141. doi: 10 . 1007 / JHEP01(2017) 141. arXiv: 1609.03450 [hep-ph]. [446] J. Neyman. ‘Outline of theory of statistical estimation based on the classical theory of probability’. In: *Philosophical Transactions of the Royal Society of London Series A* 236.767 (1937), pp. 333–380. doi: 10 . 1098/rsta.1937.0005.
- [439] M. G. Aartsen et al. ‘Search for non-standard neutrino interactions with IceCube DeepCore’. In: *Phys. Rev. D* 97.7 (2018), p. 072009. doi: 10 . 1103 / PhysRevD . 97 . 072009. arXiv: 1709.07079 [hep-ex]. [447] K. Pearson. ‘X. On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling’. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 50.302 (1900), pp. 157–175. doi: 10.1080/14786440009463897.
- [440] S. V. Demidov. ‘Bounds on non-standard interactions of neutrinos from IceCube DeepCore data’. In: *JHEP* 03 (2020), p. 105. doi: 10 . 1007 / JHEP03(2020) 105. arXiv: 1912.04149 [hep-ph]. [448] L. Lyons. ‘Open statistical issues in particle physics’. In: *Ann. Appl. Stat.* 2.3 (2008), pp. 887–915. doi: 10 . 1214 / 08 - A0AS163.
- [441] J. R. Klein and A. Roodman. ‘Blind analysis in nuclear and particle physics’. In: *Ann. Rev. Nucl. Part. Sci.* 55 (2005), pp. 141–163. doi: 10 . 1146 / annurev . nucl . 55 . 090704 . 151521. [449] S. D. Biller and S. M. Oser. ‘Another look at confidence intervals: proposal for a more relevant and transparent approach’. In: *Nucl. Instrum. Meth. A* 774 (2015), pp. 103–119. doi: 10 . 1016 / j . nima . 2014 . 11 . 081. arXiv: 1405 . 5010 [physics.data-an].
- [442] L. Lyons and N. Wardle. ‘Statistical issues in searches for new phenomena in high-energy physics’. In: *J. Phys. G* 45.3 (2018), p. 033001. doi: 10 . 1088 / 1361 - 6471 / aa9408. [450] S. S. Wilks. ‘The large-sample distribution of the likelihood ratio for testing composite hypotheses’. In: *Ann. Math. Statist.* 9.1 (1938), pp. 60–62. doi: 10 . 1214/aoms/1177732360.
- [443] R. A. Fisher and E. J. Russell. ‘On the mathematical foundations of theoretical statistics’. In: *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* 222.594-604 (1922), pp. 309–368. doi: 10.1098/rsta.1922.0009.

- [451] T. Hauschild and M. Jentschel. ‘Comparison of maximum-likelihood estimation and chi-square statistics applied to counting experiments’. In: *Nuclear Instruments and Methods in Physics Research A* 457.1-2 (2001), pp. 384–401. DOI: 10.1016/S0168-9002(00)00756-7.
- [452] G. Cowan et al. ‘Asymptotic formulae for likelihood-based tests of new physics’. In: *Eur. Phys. J. C* 71 (2011). [Erratum: *Eur.Phys.J.C* 73, 2501 (2013)], p. 1554. DOI: 10.1140/epjc/s10052-011-1554-0. arXiv: 1007.1727 [physics.data-an].
- [453] P. Virtanen et al. ‘SciPy 1.0: fundamental algorithms for scientific computing in Python’. In: *Nature Methods* 17 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2. arXiv: 1907.10121 [cs.MS].
- [454] G. J. Feldman and R. D. Cousins. ‘A unified approach to the classical statistical analysis of small signals’. In: *Phys. Rev. D* 57 (1998), pp. 3873–3889. DOI: 10.1103/PhysRevD.57.3873. arXiv: physics / 9711021.
- [455] J. S. Gainer et al. ‘Exploring theory space with Monte Carlo reweighting’. In: *JHEP* 10 (2014), p. 078. DOI: 10.1007/JHEP10(2014)078. arXiv: 1404.7129 [hep-ph].
- [456] G. Mitsuka et al. ‘Study of nonstandard neutrino interactions with atmospheric neutrino data in Super-Kamiokande I and II’. In: *Phys. Rev. D* 84 (11 2011), p. 113008. DOI: 10.1103/PhysRevD.84.113008.
- [457] P. Adamson et al. ‘Search for flavour-changing nonstandard neutrino interactions by MINOS’. In: *Phys. Rev. D* 88 (7 2013), p. 072011. DOI: 10.1103/PhysRevD.88.072011.
- [458] C. Giunti. ‘General COHERENT constraints on neutrino non-standard interactions’. In: *Phys. Rev. D* 101.3 (2020), p. 035039. DOI: 10.1103/PhysRevD.101.035039. arXiv: 1909.00466 [hep-ph].
- [459] K. S. Babu et al. ‘Neutrino non-standard interactions: complementarity between LHC and oscillation experiments’. In: *Phys. Lett. B* 815 (2021), p. 136131. DOI: 10.1016/j.physletb.2021.136131. arXiv: 2003.03383 [hep-ph].
- [460] G. Parker. ‘Strong constraints on new physics from the IceCube South Pole Neutrino Observatory’. PhD thesis. Texas University, 2022.
- [461] R. Abbasi et al. ‘Strong constraints on neutrino non-standard interactions from TeV-scale ν_μ disappearance at IceCube’. In: *Phys. Rev. Lett.* 129.1 (2022), p. 011804. DOI: 10.1103/PhysRevLett.129.011804. arXiv: 2201.03566 [hep-ex].
- [462] M. G. Aartsen et al. ‘Searches for sterile neutrinos with the IceCube detector’. In: *Phys. Rev. Lett.* 117.7 (2016), p. 071801. DOI: 10.1103/PhysRevLett.117.071801. arXiv: 1605.01990 [hep-ex].
- [463] M. G. Aartsen et al. ‘Searching for eV-scale sterile neutrinos with eight years of atmospheric neutrinos at the IceCube Neutrino Telescope’. In: *Phys. Rev. D* 102.5 (2020), p. 052009. DOI: 10.1103/PhysRevD.102.052009. arXiv: 2005.12943 [hep-ex].
- [464] M. Ageron et al. ‘ANTARES: the first undersea neutrino telescope’. In: *Nucl. Instrum. Meth.* A656 (2011), pp. 11–38. DOI: 10.1016/j.nima.2011.06.103. arXiv: 1104.1607 [astro-ph.IM].
- [465] A. Albert et al. ‘Search for non-standard neutrino interactions with 10 years of ANTARES data’. In: *JHEP* 07 (2022),

- p. 048. DOI: 10.1007/JHEP07(2022)048. arXiv: 2112.14517 [hep-ex].
- [466] A. Albert et al. ‘Measuring the atmospheric neutrino oscillation parameters and constraining the 3+1 neutrino model with ten years of ANTARES data’. In: *JHEP* 06 (2019), p. 113. DOI: 10.1007/JHEP06(2019)113. arXiv: 1812.08650 [hep-ex].
- [467] A. Lazo. ‘Updated results on neutrino non-standard interactions with KM3NeT/ORCA6’. In: *PoS ICRC2023* (2023), p. 998. DOI: 10.22323/1.444.0998.
- [468] J. M. Mańczak. ‘Search for neutrino non-standard interactions with the KM3NET/ORCA 6 lines detector’. PhD thesis. Universitat de València, 2022.
- [469] E. Lohfink. ‘Testing non-standard neutrino interaction parameters with IceCube DeepCore’. PhD thesis. Johannes Gutenberg-Universität Mainz, 2023. DOI: 10.25358/openscience-9288.
- [470] E. Hansen. ‘Early atmospheric muon rejection with PINGU’. Master thesis. Københavns Universitet, 2016.
- [471] Z. I. Botev, J. F. Grotowski and D. P. Kroese. ‘Kernel density estimation via diffusion’. In: *The Annals of Statistics* 38.5 (2010). DOI: 10.1214/10-aos799.
- [472] S. J. Sheather and M. C. Jones. ‘A reliable data-based bandwidth selection method for kernel density estimation’. In: *Journal of the Royal Statistical Society: Series B (Methodological)* 53.3 (1991), pp. 683–690. DOI: 10.1111/j.2517-6161.1991.tb01857.x.
- [473] I. S. Abramson. ‘On bandwidth variation in kernel estimates—a square root law’. In: *The Annals of Statistics* 10.4 (1982), pp. 1217–1223. DOI: 10.1214/aos/1176345986.
- [474] C. R. Harris et al. ‘Array programming with NumPy’. In: *Nature* 585.7825 (2020), pp. 357–362. DOI: 10.1038/s41586-020-2649-2.
- [475] M. C. Gonzalez-Garcia, M. Maltoni and T. Schwetz. ‘Updated fit to three neutrino mixing: status of leptonic CP violation’. In: *JHEP* 11 (2014), p. 052. DOI: 10.1007/JHEP11(2014)052. arXiv: 1409.5439 [hep-ph].
- [476] *NuFIT 2.0* (2014). URL: <http://www.nu-fit.org> (visited on 21/09/2023).
- [477] S. Nowicki and D. Grant. ‘Status of neutrino hierarchy measurements for DeepCore and PINGU’. Presentation given at Fall 2011 IceCube Collaboration meeting. Uppsala, Sweden. URL: <https://events.icecube.wisc.edu/event/36/overview>.
- [478] K. Clark and D. F. Cowen. ‘IceCube/DeepCore and IceCube/PINGU: Prospects for few-GeV scale ν physics in the ice’. In: *Nucl. Phys. B Proc. Suppl.* 233 (2012). Ed. by Marco Bozzo et al., pp. 223–228. DOI: 10.1016/j.nuclphysbps.2012.12.081.
- [479] G. L. Fogli et al. ‘Evidence of $\theta_{13}>0$ from global neutrino data analysis’. In: *Phys. Rev. D* 84 (2011), p. 053007. DOI: 10.1103/PhysRevD.84.053007. arXiv: 1106.6028 [hep-ph].
- [480] J. Neyman and E. S. Pearson. ‘On the problem of the most efficient tests of statistical hypotheses’. In: *Phil. Trans. Roy. Soc. Lond. A* 231.694-706 (1933), pp. 289–337. DOI: 10.1098/rsta.1933.0009.
- [481] M. Blennow et al. ‘Quantifying the sensitivity of oscillation experiments to the neutrino mass ordering’. In: *Journal of High Energy Physics* 2014.3 (2014), p. 28. DOI: 10.1007/JHEP03(2014)028.

- [482] S. Adrián-Martínez et al. ‘Letter of Intent for KM3NeT 2.0’. In: *J. Phys.* G43.8 (2016), p. 084001. doi: 10.1088/0954-3899/43/8/084001. arXiv: 1601.07459 [astro-ph.IM].
- [483] R. A. Fisher. ‘Theory of statistical estimation’. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 22.5 (1925), pp. 700–725. doi: 10.1017/S0305004100009580.
- [484] X. Cao. ‘Relative performance of expected and observed Fisher information in covariance estimation for maximum likelihood estimates’. PhD thesis. Johns Hopkins University, 2013. arXiv: 1305.1056 [math.ST].
- [485] C. J. Geyer. ‘Stat 5102 lecture notes’. Lecture notes. 2001. URL: <https://www.stat.umn.edu/geyer/5102/notes/n2.pdf> (visited on 28/03/2024).
- [486] G. L. Fogli et al. ‘Getting the most from the statistical analysis of solar neutrino oscillations’. In: *Phys. Rev. D* 66 (2002), p. 053010. doi: 10.1103/PhysRevD.66.053010. arXiv: hep-ph/0206162.
- [487] F. Capozzi. ‘Phenomenology of neutrino oscillations and the mass hierarchy’. PhD thesis. Università degli studi di Bari Aldo Moro, 2016.
- [488] J. Weldert. ‘Combining JUNO and PINGU to determine the neutrino mass ordering’. Master thesis. Johannes Gutenberg-Universität Mainz, June 2018.
- [489] F. Capozzi, E. Lisi and A. Marrone. ‘PINGU and the neutrino mass hierarchy: statistical and systematic aspects’. In: *Phys. Rev. D* 91 (2015), p. 073011. doi: 10.1103/PhysRevD.91.073011. arXiv: 1503.01999 [hep-ph].
- [490] The IceCube collaboration. *PINGU Simulation and Analysis (PISA) software repository*. URL: <https://github.com/icecubeopensource/pisa> (visited on 18/01/2024).
- [491] R. G. Calland, A. C. Kaboth and D. Payne. ‘Accelerated event-by-event neutrino oscillation reweighting with matter effects on a GPU’. In: *JINST* 9 (2014), P04016. doi: 10.1088/1748-0221/9/04/P04016. arXiv: 1311.7579 [physics.data-an].
- [492] V. Barger et al. ‘Matter effects on three-neutrino oscillations’. In: *Phys. Rev. D* 22 (11 1980), pp. 2718–2726. doi: 10.1103/PhysRevD.22.2718.
- [493] T. Ohlsson and H. Snellman. ‘Three-flavour neutrino oscillations in matter’. In: *J. Math. Phys.* 41 (2000). [Erratum: *J.Math.Phys.* 42, 2345 (2001)], pp. 2768–2788. doi: 10.1063/1.533270. arXiv: hep-ph/9910546.
- [494] T. Ohlsson and H. Snellman. ‘Neutrino oscillations with three flavours in matter: applications to neutrinos traversing the Earth’. In: *Phys. Lett. B* 474 (2000). [Erratum: *Phys.Lett.B* 480, 419–419 (2000)], pp. 153–162. doi: 10.1016/S0370-2693(00)00008-3. arXiv: hep-ph/9912295.
- [495] J. Kopp. ‘Efficient numerical diagonalisation of Hermitian 3×3 matrices’. In: *Int. J. Mod. Phys. C* 19 (2008). [Erratum: *Int. J. Mod. Phys. C* 19,845(2008)], pp. 523–548. doi: 10.1142/S0129183108012303. arXiv: physics/0610206.
- [496] D. Meloni, T. Ohlsson and H. Zhang. ‘Exact and approximate formulas for neutrino mixing and oscillations with non-standard interactions’. In: *JHEP* 04 (2009), p. 033. doi: 10.1088/1126-6708/2009/04/033. arXiv: 0901.1784 [hep-ph].

Bibliography

- [497] P. Huber, M. Lindner and W. Winter. ‘Simulation of long-baseline neutrino oscillation experiments with GLoBES (General Long-Baseline Experiment Simulator)’. In: *Comput. Phys. Commun.* 167 (2005), p. 195. doi: 10.1016/j.cpc.2005.01.003. arXiv: hep-ph/0407333.
- [498] P. Huber et al. ‘New features in the simulation of neutrino oscillation experiments with GLoBES 3.0: General Long-Baseline Experiment Simulator’. In: *Comput. Phys. Commun.* 177 (2007), pp. 432–438. doi: 10.1016/j.cpc.2007.05.004. arXiv: hep-ph/0701187.
- [499] C. A. Argüelles Delgado, J. Salvado and C. N. Weaver. ‘A simple quantum integro-differential solver (SQuIDS)’. In: *Comput. Phys. Commun.* 196 (2015), pp. 569–591. doi: 10.1016/j.cpc.2015.06.022. arXiv: 1412.3832 [hep-ph].
- [500] C. A. Argüelles Delgado, J. Salvado and C. N. Weaver. ‘nuSQuIDS: A toolbox for neutrino propagation’. In: *Comput. Phys. Commun.* 277 (2022), p. 108346. doi: 10.1016/j.cpc.2022.108346. arXiv: 2112.13804 [hep-ph].
- [501] F. Kallenborn et al. ‘Massively-parallel computation of atmospheric neutrino oscillations on CUDA-enabled accelerators’. In: *Comput. Phys. Commun.* 234 (2019), pp. 235–244. doi: 10.1016/j.cpc.2018.07.022.

List of Figures

2.1	Helicity states of chiral fermion fields	17
2.2	Contraction of W^\pm propagator in effective low-energy CC process . .	25
2.3	Feynman diagrams for CC QES, RES, and DIS	30
2.4	W - Q^2 plane for inelastic νN scattering at $E_\nu = 7$ GeV [142]	31
2.5	Total $\nu_\mu N$ CC and $\nu_\alpha N$ NC cross section at GeV energies [157]	32
2.6	Approaches to modelling neutrino-nucleus scattering [172]	38
3.1	Masses of fundamental fermions [188]	43
3.2	Spacetime diagram of neutrino oscillation process [212]	53
3.3	Discrete transformations and (anti)neutrino flavour transition channels	56
3.4	Neutrino mass ordering and flavour-mass mixing [222]	58
3.5	Feynman diagrams for neutrino coherent forward scattering	61
3.6	2ν oscillation probability in uniform matter vs. $2E_\nu V_{CC}/\Delta m^2$	68
3.7	2ν example of adiabatic conversion [183]	70
3.8	2ν example of parametric transition probability enhancement	72
4.1	Feynman diagrams generating NSI via loop corrections [260]	81
5.1	Solar ν_e flux spectra and predicted survival probability on Earth [277]	85
5.2	Predicted atmospheric neutrino flux spectra at South Pole [281]	88
5.3	Predicted atmospheric neutrino production height distributions at South Pole [281]	90
5.4	Atmospheric neutrino baseline vs. zenith angle	91
5.5	Sketch of “upgoing” atmospheric neutrino trajectories	93
5.6	Standard matter-potential profiles encountered by atmospheric neutrinos with different zenith angles	94
5.7	Neutrino transition scheme in propagation basis	96
5.8	Earth oscillograms for transitions of atmospheric (anti)neutrinos [239]	99
5.9	Pion decay at rest [299]	100
5.10	Expected reactor antineutrino spectra [277] and 3ν oscillation probabilities vs. baseline [304]	101
5.11	One-dimensional $\Delta\chi^2$ profiles from global oscillation data for $\sin^2 \theta_{23}$ and δ_{CP} in presence of NSI [28]	108
5.12	Confidence regions from global oscillation data for $\epsilon_\oplus, \varphi_{12}, \varphi_{13}$ [223]	112
6.1	Signatures of atmospheric neutrino interactions in ice	118

List of Figures

6.2	Energy flow within hadronic cascade	120
6.3	Measurements of absorption coefficient for laboratory ice and deep South Pole ice from UV to IR [343]	121
6.4	Measurements of effective scattering coefficient for deep South Pole ice from near-UV to visible wavelengths [343]	123
6.5	Measurements of absorption and effective scattering coefficients for deep South Pole ice vs. depth for $\lambda = 400$ nm [355]	124
7.1	Top- and sideview of IceCube and DeepCore [362]	129
7.2	Main components of DOM [364]	130
7.3	Unfolding example of DOM waveforms [365]	133
7.4	“Seeded $R-T$ ” noise-removal scheme	134
7.5	DeepCore online filtering scheme [40]	135
8.1	Layout of proposed PINGU detector	140
8.2	Multi-PMT and D-Egg optical module prototypes [43]	141
9.1	Atmospheric neutrino oscillation probabilities for $\cos \vartheta = -0.75$ for varying ϵ_{\oplus}	146
9.2	Atmospheric neutrino oscillation probabilities for $\cos \vartheta = -0.75$ for varying $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$	148
9.3	Atmospheric neutrino oscillation probabilities for $\cos \vartheta = -0.75$ for varying $\epsilon_{\mu\tau}^{\oplus}$	149
9.4	Atmospheric neutrino oscillation probabilities for $\cos \vartheta = -0.75$ for varying $\epsilon_{e\mu}^{\oplus}$ and $\epsilon_{e\tau}^{\oplus}$	151
10.1	Typical DeepCore neutrino simulation scheme [393]	154
10.2	Low-energy neutrino generation scheme with GENIE	155
10.3	Photoelectron amplification probability vs. wavelength for regular-QE DOM	158
10.4	Intrinsic and effective angular DOM acceptance	158
10.5	Simulated PMT time jitter distribution	160
10.6	Simulated PMT charge response to single photoelectron	160
10.7	Poisson-like approximation of PMT charge distribution in template reconstruction approach [369]	164
11.1	Run livetimes and event numbers of DeepCore NSI measurement	172
11.2	Predicted event rates after each selection level	174

11.3	Sketch of light depositions from interacting muon (anti)neutrino and through-going μ^\pm [428]	175
11.4	Sketch of isolated veto signals from noise and μ^\pm [428]	175
11.5	Final-level event containment requirements	175
11.6	“Blind” corridors within DeepCore [428]	175
12.1	Predicted DeepCore neutrino event distribution and unweighted event counts	183
12.2	Predicted μ^\pm event distribution and unweighted sideband event counts	183
12.3	Predicted E_ν and $\cos \vartheta$ distributions	184
12.4	Predicted track-like classification fraction vs. E_ν	186
12.5	Predicted reconstruction residual distributions vs. E_ν	187
12.6	Predicted binwise statistical significance of $\epsilon_{ee}^\oplus - \epsilon_{\mu\mu}^\oplus = 0.25$ w.r.t. SI .	189
12.7	Predicted binwise statistical significance of $\epsilon_{\tau\tau}^\oplus - \epsilon_{\mu\mu}^\oplus = 0.07$ w.r.t. SI .	190
12.8	Predicted binwise statistical significance of $\epsilon_{e\mu}^\oplus = 0.10$ w.r.t. SI	191
12.9	Predicted binwise statistical significance of $\epsilon_{e\tau}^\oplus = -0.26$ w.r.t. SI	192
12.10	Predicted binwise statistical significance of $\epsilon_{\mu\tau}^\oplus = 0.021$ w.r.t. SI	193
13.1	Nominal binwise event-count uncertainties used in weighted-least-squares statistic of DeepCore NSI measurement	206
14.1	DeepCore Asimov test-statistic landscapes in NU strengths	210
14.2	Asimov test-statistic landscapes in real FV couplings	211
14.3	Statistics-only Asimov sensitivities to NU strengths	213
14.4	Statistics-only Asimov sensitivities to FV couplings	214
14.5	Statistics-only Asimov sensitivity to GMP	215
14.6	$\overline{\chi^2_{\text{mod}}}(\theta_{23})$ under different NSI hypotheses for SI pseudodata	219
14.7	Nuisance-parameter impacts on Asimov sensitivities to NU strengths	221
14.8	Nuisance-parameter impacts on Asimov sensitivities to FV couplings and GMP strength	222
14.9	Expected μ - τ NU fit bias from nuisance-parameter offsets	224
14.10	Statistics-only distributions of χ^2_{mod} for pseudoexperiments with and without MC fluctuations	227
14.11	Predicted χ^2_{mod} distributions for different NSI hypotheses overlaid with χ^2_m fits	228
14.12	Sensitivity distributions for NU strengths	230
14.13	Sensitivity distributions for FV couplings	231

14.14	Predicted $\Delta\chi_{\text{mod}}^2$ ($p_{\text{NSI}} = 0$) distributions for true SI fit with six NSI hypotheses	234
14.15	Predicted $\Delta\chi_{\text{mod}}^2$ ($\epsilon_{\mu\tau}^{\oplus} = 0.035 \epsilon_{\mu\tau}^{\oplus} = 0.035$) distribution fit with μ - τ FV hypothesis	235
15.1	Observed event distribution in DeepCore NSI measurement	237
15.2	Observed distribution of event-count pulls at μ - τ NU best fit	239
15.3	Observed $\Delta\chi_{\text{mod}}^2$ profiles for $\epsilon_{ee}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ and $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$ with statistical sensitivity bands	241
15.4	Observed 90 % confidence regions for $\epsilon_{e\mu}^{\oplus}$, $\epsilon_{e\tau}^{\oplus}$, and $\epsilon_{\mu\tau}^{\oplus}$ with statistical sensitivity bands	244
15.5	Observed 1σ , 90 %, and 3σ confidence regions for ϵ_{\oplus} , φ_{12} , and φ_{13}	245
15.6	Summary of observed one-by-one 90 % confidence intervals on standard NSI parameters compared to previous limits	248
15.7	Comparison of observed 90 % confidence regions on complex $\epsilon_{\mu\tau}^{dV}$ from high-energy IceCube measurement and this work [461]	250
15.8	Comparison of various observed 90 % confidence intervals on real NSI with down quarks including this work [467]	251
15.9	Feldman–Cousins approach for SI rejection in NU NSI measurements	253
15.10	Feldman–Cousins approach for SI rejection in FV NSI measurements	254
15.11	Feldman–Cousins approach for SI rejection in GMP measurement	255
17.1	Staged template generation approach in PINGU	264
17.2	Neutrino- and interaction-type flow chart of staged approach	266
17.3	Raw MC and smoothed ν_e CC effective area	268
17.4	VBWKDEs of reconstruction residuals	269
17.5	Smoothed track-like classification probabilities	270
17.6	Predicted annual neutrino event distribution in analysis binning	271
17.7	Predicted event rate vs. $L_{\text{reco}}/E_{\text{reco}}$ with and without oscillations	272
17.8	Predicted binwise annual statistical significance of maximal 2–3 mixing w.r.t. first-octant model	273
17.9	Predicted binwise annual statistical significance of $\Delta m_{31}^2 = 2.5 \times 10^{-3} \text{ eV}^2$ w.r.t. nominal model	275
17.10	Predicted binwise annual statistical significance of switching the sign $\Delta m_{31}^2 > 0 \rightarrow -\Delta m_{31}^2$	277
18.1	DeepCore NMO measurement outcome [427]	284

List of Figures

18.2	Median Asimov NMO significance vs. TO test-statistic mean and WO test-statistic mean at significance minimum	287
18.3	Error ellipses from Fisher matrix approach compared to isocontours of $\overline{\chi^2}$ surfaces in PINGU toy NMO sensitivity studies	292
18.4	Results of single-nuisance-parameter $\overline{\Delta\chi_{\text{IO}}^2}$ profile scans of θ_{23} using numerical minimisation and pull approach	297
18.5	$\overline{\Delta\chi_{\text{IO}}^2}$ profile scan of θ_{23} using numerical minimisation and pull approach including all nuisance parameters	298
18.6	Statistics-only test for IO using pseudoexperiment ensembles and $\overline{\Delta\chi^2}$ approach	300
18.7	Statistics-only test for IO using reweighted pseudoexperiments . . .	303
18.8	Time evolution of statistics-only NMO sensitivity using pseudoexperiment ensembles and Asimov approaches	305
19.1	PINGU median systematic errors on NMO	309
19.2	Time evolution of NMO synergy from combined cascade- and track-like template fit	311
19.3	Asimov $\overline{\chi^2}$ profile scans of Δm_{31}^2 within NO for true IO	312
19.4	Asimov $\overline{\chi^2}$ profile scans of Δm_{31}^2 within IO for true NO	313
19.5	Dependence of $\overline{\Delta\chi_{\text{IO}}^2}$ on true θ_{23} and Δm_{31}^2	315
19.6	Dependence of median NMO significance on true θ_{23} and Δm_{31}^2 . . .	316
19.7	Dependence of median NMO significance on true θ_{23} for different WO models	317
19.8	Various best-fit or hypothesised values of θ_{23} vs. truth in NMO hypothesis tests	318
19.9	Dependence of median NMO significance on true Δm_{31}^2 for different WO models	319
19.10	Various best-fit or hypothesised values of Δm_{31}^2 vs. truth in NMO hypothesis tests	320
C.1	Atmospheric antineutrino oscillation probabilities for $\cos \vartheta = -0.75$ for varying ϵ_{\oplus}	337
C.2	Atmospheric antineutrino oscillation probabilities for $\cos \vartheta = -0.75$ for varying $\epsilon_{\tau\tau}^{\oplus} - \epsilon_{\mu\mu}^{\oplus}$	338
C.3	Atmospheric antineutrino oscillation probabilities for $\cos \vartheta = -0.75$ for varying $\epsilon_{\mu\tau}^{\oplus}$	338

List of Figures

C.4	Atmospheric antineutrino oscillation probabilities for $\cos \vartheta = -0.75$ for varying $\epsilon_{e\mu}^{\oplus}$	339
C.5	Atmospheric antineutrino oscillation probabilities for $\cos \vartheta = -0.75$ for varying $\epsilon_{e\tau}^{\oplus}$	339
C.6	Expected DeepCore e - μ NU fit bias from nuisance-parameter offsets	340
C.7	Expected DeepCore e - μ FV fit bias from nuisance-parameter offsets .	340
C.8	Expected DeepCore e - τ FV fit bias from nuisance-parameter offsets .	341
C.9	Expected DeepCore μ - τ FV fit bias from nuisance-parameter offsets .	341
C.10	DeepCore Asimov sensitivities to NSI NU and FV strengths with extended set of nuisance parameters	342
C.11	(Cumulative) histogram of observed binwise weighted-least-squares values in DeepCore μ - τ NU measurement	343
C.12	Observed nuisance-parameter pulls in DeepCore NSI measurements	344
C.13	Observed joint and individual constraints on θ_{23} and Δm_{31}^2 from DeepCore SI fit	344
D.1	Predicted PINGU ν_{μ} CC and NC event rates vs. L/E_{ν} with and without oscillations	345
D.2	Joint nuisance-parameter estimates conditional on θ_{23} within NO for true IO using numerical minimisation and pull approach	347

List of Tables

2.1	Electroweak gauge quantum numbers of fundamental fermion fields	19
5.1	Global constraints on three-neutrino mixing parameters and mass-squared differences from NuFIT 5.2 [308, 313]	104
5.2	90 % confidence intervals from global oscillation data on real NSI parameters for Earth matter [223]	111
6.1	Cherenkov energy threshold for various leptons and hadrons in ice	116
11.1	Summary of event selection criteria of DeepCore NSI measurement	173
12.1	Predicted final-level event sample composition of DeepCore NSI measurement	181
12.2	Nuisance parameters tested for inclusion	202
13.1	Overview of hypotheses studied by DeepCore NSI measurement	203
14.1	Predicted Asimov 90 % confidence intervals and most impactful nuisance parameters of DeepCore NSI measurement	223
15.1	Summary of observed DeepCore NSI fit outcomes	238
15.2	Summary of observed 90 % confidence intervals	247
16.1	Summary of PINGU event selection criteria	261
17.1	Nuisance parameters tested for inclusion in PINGU NMO sensitivity study	280
18.1	PINGU median four-year statistics-only NMO significance using different analysis methods	301
19.1	PINGU median four-year Asimov NMO significance for different systematics categories	308
C.1	Predicted DeepCore event sample composition at GMP best fit	343