Doctoral Thesis

Atmospheric neutrino oscillations in IceCube-DeepCore within and beyond the unitary framework

Tetiana Kozynets

This thesis has been submitted to the PhD School of The Faculty of Science, University of Copenhagen

To the brave, the honest, and the kind;
to those who light the way without knowing.

Abstract

The Standard Model of particle physics postulates the existence of three weakly interacting neutrino flavor states but fails to explain their observed oscillations, which require neutrinos to propagate as superpositions of distinct mass states. A key unresolved question is the origin of neutrino masses, which is assumed to be different from the Higgs mechanism of the Standard Model. A class of models explain the very small yet non-vanishing neutrino masses through the existence of one or multiple heavy and non-weakly interacting neutrino states, leading to a breach in unitarity of the three-flavor mixing framework.

This thesis focuses on atmospheric neutrinos as a tool to constrain the unitarity of the neutrino mixing matrix, which links the flavor and the mass states. For this purpose, we utilize three years of data from IceCube-DeepCore, the sub-array of IceCube Neutrino Observatory at the South Pole, which detects neutrinos with energies from a few GeV to a few hundreds of GeV. By combining this dataset with the reactor neutrino data from the Daya Bay and KamLAND experiments, we constrain the individual elements of the mixing matrix and derive the ensuing non-unitarity metrics. Our results are consistent with unitarity, with precision levels of $\mathcal{O}(1\%)$ the electron sector and $\mathcal{O}(10\%)$ in the muon and tau sectors. The projections for the upcoming IceCube-Upgrade and JUNO experiments suggest further improvements, halving uncertainties in the muon and tau sectors and reducing those in the electron sector by 25%.

Additionally, we present the contributions towards an improved event selection in IceCube-DeepCore, incorporating nine years of data to probe unitarity in the tau sector with greater precision. We address systematic issues in muon Cherenkov light yield simulation and energy reconstruction, demonstrating their impact on oscillation analyses. We derive updated sensitivities to the tau neutrino normalization, which is estimated to be 9% at the 1σ level with the final sample. This sensitivity is contingent upon the correlation of tau and muon neutrino cross section uncertainties – a relationship validated through a dedicated study in this thesis.

Abstrakt

Standardmodellen for partikelfysik postulerer eksistensen af tre svagt interagerende neutrinotyper, men kan ikke forklare deres observerede oscillationer, som kræver, at neutrinoer udbreder sig som superpositioner af forskellige massetilstande. Et centralt uafklaret spørgsmål er neutrinoerne massers oprindelse, som antages at være forskellig fra standardmodellens Higgs-mekanisme. En klasse af neutrinomasse-modeller forklarer neutrinomassernes lave masser gennem eksistensen af en eller flere tunge og ikke-svagt vekselvirkende neutrinotilstande, hvilket fører til et brud på unitariteten i tre-type-beskrivelsen (*three-flavor mixing framework*).

Denne afhandling fokuserer på atmosfæriske neutrinoer som et værktøj til at begrænse unitariteten af neutrinoblandingsmatricen (*neutrino mixing matrix*), som forbinder type- og massetilstandene. Til dette formål bruger vi tre års data fra IceCube-DeepCore, en del af IceCube Neutrino Observatory på Sydpolen, som detekterer neutrinoer med energier fra nogle få GeV til nogle få hundrede GeV. Ved at kombinere dette datasæt med reaktor-neutrinodata fra Daya Bay- og KamLAND-eksperimenterne indsnævrer vi de enkelte elementer i blandingsmatricen og udleder de deraf følgende ikke-unitaritetsmålinger. Vores resultater er i overensstemmelse med unitaritet med $\mathcal{O}(1\%)$ præcision i elektronsektoren og $\mathcal{O}(10\%)$ i myon- og tau-sektoren. Prognoserne for de kommende IceCube-Upgrade- og JUNO-eksperimenter tyder på yderligere forbedringer, der halverer usikkerhederne i myon- og tau-sektoren og reducerer dem i elektronsektoren med 25%.

Derudover præsenterer vi bidragene til en forbedret udvælgelse af hændelser i IceCube-DeepCore, hvor vi inddrager ni års data for at undersøge unitariteten i tau-sektoren med større præcision. Vi tager fat på systematiske problemer i simulering af myon-cherenkov-lysudbytte (*light yield*) og energirekonstruktion og demonstrerer deres indvirkning på oscillationsanalyser. Vi udleder opdaterede sensitiviteter for tau-neutrino-normaliseringen, som anslås til at være 9% på 1σ-niveau med den endelige prøve. Denne sensitivitet er betinget af korrelationen mellem tau- og myonneutrino-tværsnitsusikkerheder – et forhold, der er valideret gennem en dedikeret undersøgelse i denne afhandling.

Statement of contributions

This thesis presents predominantly my independent work as well as the work carried out collaboratively on specific projects. Importantly, Sections 2.3 and 2.4, Chapters 7, 8 and 11, and Appendices B and C of the thesis contain the following two publications:

1. **Tetiana Kozynets**, Philipp Eller, Alan Zander, Manuel Ettengruber, and David Jason Koskinen, *Constraints on Non-Unitary Neutrino Mixing in Light of Atmospheric and Reactor Neutrino Data*, submitted to J. High Energy Phys. (2024), arXiv:2407.20388 [hep-ph] [1].

2. **Tetiana Kozynets**, Thomas Stuttard, and David Jason Koskinen, *Deep Inelastic Scattering Cross Section Uncertainties in Tau Neutrino Appearance Measurements*, submitted to Phys. Rev. D. (2024), arXiv:2409.01258 [hep-ph] [2].

Both of these papers are under review by the respective journals at the time of thesis submission, with their full pre-prints available online in the arXiv database.

The first study utilizes the public three-year IceCube-DeepCore dataset [3], with the respective event selection developed in [4] by M. Larson. The basis of the software framework enabling global analyses of neutrino oscillation data was developed by P. Eller. I have implemented the non-unitary neutrino mixing model within this framework as well as a third-party software for neutrino oscillation probability calculations [5]. I have additionally contributed to the integration IceCube-DeepCore analysis into the global fitting framework (in particular with regards to the usage of the Sample A variant of the IceCube-DeepCore data and the treatment of matter effects in the non-unitary mixing scenario). Code development for reactor neutrino experiments was done primarily by M. Ettengruber and A. Zander, where I contributed with the implementation of systematic uncertainties. Finally, I independently performed and documented the entire analysis presented in the publication, including development, optimization, and execution of the code for the Bayesian fit, conducting all of the core measurements, performing post-fit checks, and interpreting the results.

The research question for the second publication was formulated by D. J. Koskinen and answered independently by me, including the identification all of the methods and computational/software tools necessary to perform the calculation in question, carrying out the calculation, and documenting the findings in a manuscript.

Chapter 10 presents the work on the oscNext event selection carried out in close collaboration with T. Stuttard, whose explicit contributions are specified in text. In particular, I highlight that all of the official oscNext Monte Carlo simulation was produced and processed by T. Stuttard, with important contributions from K. Leonard de Holton, É. Bourbeau, and R. Ørsøe. I have analyzed the individual stages and aspects of this event selection, produced smaller batches of Monte Carlo simulation for testing and interpretation purposes, developed physical explanations for the shortcomings of the simulation, and contributed to the interpretation of the data/Monte Carlo discrepancies.

Unless explicitly stated in the figure/table captions, all of the figures/tables in this thesis were produced by me.

Acknowledgements

This thesis, while a reflection of my work over the last few years, is also a reflection of an incredible community to which I have been lucky and privileged to be warmly welcomed, through which I have grown scientifically and personally, and which supported my journey throughout this PhD in more ways than I could have hoped for.

First and foremost, I would like to express my infinite gratitude to *Jason Koskinen*, my advisor, for his guidance and protection, for seeing me as the scientist and the human being that I am, for accepting me with the entire spectrum of strengths and weaknesses, for giving me the opportunities to utilize the former and work on the latter, and for so much more. Jason, the last five years would not have been the same without you, and the next five years will not be the same without you either. To think critically, to find inconsistencies in complex logical chains, to prioritize, to simplify, to be truthful, and to navigate the twists and turns of life are only some of the very many things that I learned from you, which will keep maturing in my mindset and inspiring me throughout my future. Perhaps even more importantly, thank you, *April*, and *Kai* for making me feel at home and helping me get back on my feet, sometimes literally, sometimes metaphorically.

Another part of my PhD life at the Niels Bohr Institute which has a very special place in my heart is the entire IceCube family, of which I am not ready to talk in the past tense. My dearest academic sister, *Kathrine*: what a privilege it has been to be by your side and to have you by my side in the last few years. Thank you for the inspiration you have always been to me, for your unique support and humor, for being the embodiment of the fully reciprocal "sharing and caring", for experiencing the beautiful randomness of our office with me, and for keeping the hope for academia alive in both of us. *Markus*: thank you for being a human example of an exceptionally insightful scientist and a kind and compassionate soul. I will deeply cherish the sense of community and belonging brought by you through all the cookies and all the DHL runs you organized. *Janni*: thank you for being the sunshine of our office, for raising its average positivity and contributing to the aforementioned beautiful randomness. You are a pure soul and a brilliant mind, and the gloomiest of my days became so much better with your hugs and our talks over tea. I wish the best of the PhD adventures to both you and *Marc*, who we are very happy to have back in the group after all the years and experiences. And finally, my fellow firefighter, *Tom*: what an incredible adventure we have had! While I am struggling to find the right words to describe the thrill and the horror of our shared quest for the truth, my gratitude for having been through it together will always express itself with salt and vinegar crisps, pistachios, and Belgian beers.

Next, I would like to thank *Philipp Eller* for welcoming me for my short PhD exchange at the Technical University of Munich and for all of our work together, which has been extremely fun and productive. I will have the greatest memories of the "seven" weeks of Tania, our whiteboard discussions that materialized in a paper, and your magic power to get us unstuck. Thank you for your positive and goal-oriented attitude and for being such a great and thoughtful person to be around. My days in the Garching bunker and its snowy surroundings will also be warmly remembered thanks to *Rasmus*, *Alan*, and

Manuel, whose company and humor lifted my mood and spirit.

Next up are my dearest Copenhagen friends, whom I met throughout my journey as a graduate student – Étienne, Arianna, Rita, Ada, Antonio, Marie, Kevin, Arnau, Greta, James, Alex, Karolis, Jessa, Tymon, and Kayla – as well as my friends from back home, Tania and Ania – you are all so very special to me, and each of you played such a big part in making this thesis happen. Thank you for the love and care, the trust, the inspiration, and the incredible times, good and bad, that we have shared. And of course, I cannot forget my lovely friends and colleagues from NBIA, current and former, who made my PhD experience so much more memorable. Dear Manuel, Bernie, Ersilia, Victor, Damiano, Enrico, Pablo – thank you for having been such a great company to read and discuss books with.

To the PhD students anywhere the world reading this in anticipation of their own writing period: this thesis has been written with the help of background jazz music and the Pomodoro timer, hugs of my friends and colleagues, as well as regular physical exercise. For the latter, I would like to thank my fitness coaches Mikka, Celina, and Anna, who contributed more to my sanity and progress than they probably realize.

Finally, I thank my family – my dear mom Iryna, my dad Oleksandr, my brother Andrii, and my uncle Ihor – for your love, for accepting this lifelong adventure for what it is, and our closeness and connection throughout the challenging years.

With that, let the Thesis commence.

Contents

Introduction

Among a multitude of hints pointing towards physics beyond the Standard Model [6–8], *neutrinos* play a particularly prominent role, with some of their most fundamental particle properties still lacking a solid theoretical explanation [9, 10][1]. Over the last few decades, the interest in these mysterious particles has further spread to astrophysics and cosmology, raising the questions of their astrophysical origin [11–13] and the impact on the large-scale structure of the Universe [14, 15].

Today, the neutrino is known as the second most abundant particle in the Universe [16], which is produced in a variety of natural environments such as the Sun [17, 18], active galactic nuclei [19–21], supernovae [22, 23] and their remnants [24, 25], as well as the Earth's atmosphere [26]. Additionally, neutrinos are actively "manufactured" for research purposes on Earth, e.g., in nuclear reactor and particle accelerator facilities [27, 28]. Despite their ubiquity, neutrinos are anything but ordinary; they are the only fermions in the Standard Model that interact exclusively through the weak force, a uniquely peculiar force of nature due to its maximal violation of parity [29]. The eigenstates of the weak interactions are called *flavor* states, with the three Standard Model flavors comprising electron neutrinos ν_e, muon neutrinos ν_μ, and tau neutrinos ν_τ – along with their corresponding antiparticles $\bar{\nu}_\alpha$ ($\alpha \in \{e, \mu, \tau\}$). Through observations of *neutrino oscillations* – in which flavor eigenstates transform into one another during propagation [9, 17, 30, 31] – it is now known that there are at least three neutrino non-degenerate mass states, with at least two of them having non-zero masses. In modern neutrino physics, the composition of the flavor eigenstates in terms of mass eigenstates is described through the *mixing matrix* [32], which is the central mathematical construct in this thesis. Under the Standard Model assumption of three neutrino flavors, and conservation of the total transition probability from a given flavor to any of the available three, this matrix is unitary.

A key motivation for exploring the mixing matrix structure is the unresolved origin of neutrino masses. Unlike other Standard Model fermions whose masses are generated through the Higgs mechanism [33], neutrino masses are exceedingly small, constrained by direct measurements and cosmological data to be below approximately $\mathcal{O}(1\,\text{eV})$ and $\mathcal{O}(0.1\,\text{eV})$, respectively [34]. Neutrino mass generation is therefore an active research topic (for a review, see [35, 36]), with many of the proposed models resulting in non-unitarity of the neutrino mixing matrix. In particular, the seesaw-type models explain

---

[1]The author defines a solid theoretical explanation as the one convincingly supported by experimental evidence.

the smallness of the three known neutrino masses through the existence of new, much heavier degrees of freedom – for example, neutrino states that do *not* couple to the weak force (the so-called *sterile* neutrinos). The latter have gained large attention not only in the neutrino mass context, but also as a potential dark matter candidate [37–39].

The above considerations motivate the main topic of this thesis – namely, constraining the unitarity of the neutrino mixing matrix using neutrino oscillation data. To achieve this, we utilize primarily atmospheric neutrinos as observed in the IceCube Neutrino Observatory at the South Pole [40, 41]. Atmospheric neutrinos originate from cosmic ray-induced atmospheric air showers and cover a wide range of energies from GeV to PeV [42]. While they are produced in all three flavors, at sub-TeV energies muon and electron neutrino fluxes dominate over the negligible tau neutrino flux [43, 44]. This motivates and enables an experimental probe of the *tau neutrino appearance* from atmospheric neutrino oscillations, particularly through the $\nu_\mu \rightarrow \nu_\tau$ channel at $\mathbb{O}(1\text{–}100\,\text{GeV})$ energies. The IceCube Collaboration leads an active program in this regard through the detection of atmospheric neutrinos of all flavors in the low-energy sub-array called IceCube-DeepCore [45].

One of the related physical quantities that can be constrained with the IceCube-DeepCore data is the so-called ν_τ normalization, N_{ν_τ}. A value of $N_{\nu_\tau} = 1$ would indicate that the observed number of tau neutrinos ($\nu_\tau + \bar{\nu}_\tau$) is consistent with the standard three-flavor picture of neutrino oscillations under a unitary mixing matrix. In 2019, the IceCube Collaboration reported the measurement of $N_{\nu_\tau} = 0.73^{+0.30}_{-0.24}$, using three years of IceCube-DeepCore data collected between 2012–2015 [46]. Concurrently, a measurement of $N_{\nu_\tau}^{\text{CC}} = 0.57^{+0.36}_{-0.30}$ was reported in the same study, representing the number of tau neutrinos participating in charged current ("CC") interactions relative to the nominal expectation. This result is comparable to the measurements of $N_{\nu_\tau}^{\text{CC}}$ by the OPERA [47] and the Super-Kamionande [48] Collaborations, yielding $1.2^{+0.6}_{-0.5}$ and 1.47 ± 0.32, respectively. All of the mentioned results are therefore consistent with unitarity of the neutrino mixing matrix in the tau sector.

The focus and the objectives of this thesis are twofold. First, we aim to expand the previous measurement of N_{ν_τ} to a measurement of all elements of the neutrino mixing matrix – the effort reported in Chapters 7 and 8, as well as the pre-print of the respective publication [1]. This formulation of the unitarity test is less ad-hoc, as instead of a simple scaling of the overall number of tau neutrinos, it enables a comprehensive treatment of non-unitary mixing throughout neutrino production, propagation, and detection. Furthermore, from the individual elements of the mixing matrix, the non-unitarity metrics can be derived not only in the tau sector, but also in the muon and electron neutrino sectors. To aid the latter, we combine three years of atmospheric neutrino data from IceCube-DeepCore with reactor $\bar{\nu}_e$ data from the Daya Bay [49] and KamLAND [50] experiments. This "minimal" global fit improves on previous similar studies in literature [51–53] by incorporating atmospheric neutrino systematic uncertainties, whose impact on the non-unitarity metrics is studied in detail. Secondly, we present a follow-up effort related to the tau neutrino appearance (N_{ν_τ} measurement) analysis with 9.3 years of IceCube-DeepCore data, which was initiated by É. Bourbeau in [54] and is summarized in Chapter 9. As this analysis had returned a poor goodness of fit, we thoroughly investigate symptoms and sources of disagreement between data and Monte Carlo simulation that could have caused this outcome. This includes a deep scrutiny of the IceCube Monte Carlo simulation at low level as well as reconstruction of the analysis variables, where we find and resolve multiple issues as described in Chapter 10. In parallel, we study the relevance of various systematic uncertainties that

can affect the N_{ν_τ} measurement – in particular, the relative scaling of the tau neutrino cross section with respect to the muon neutrino cross section. This investigation is reported in a pre-print of a publication [2], which is included in Chapter 11. Overall, while this thesis puts competitive constraints on neutrino mixing non-unitarity using atmospheric and reactor neutrino data, the investigation of systematic effects and the intricacies of Monte Carlo simulation in the context of new physics searches remains the overarching theme. This becomes apparent both through the impact of the systematic uncertainties on the three-year measurement and through the shift from 3 to 9.3 years of IceCube-DeepCore data, which marks the transition from a statistics-dominated to a systematics-dominated regime.

To put the aforementioned analyses into broader context, Chapter 1 opens this thesis with an overview of the main theoretical concepts in the Standard Model of particle physics. In Chapter 2, we focus specifically on neutrinos, including their known and unknown properties and the phenomenology of their oscillations, where we devote particular attention to the non-unitary mixing case. Further, we take a closer look at atmospheric neutrinos in Chapter 3, namely their production in atmospheric air showers and the specifics of atmospheric neutrino oscillations. We then review the physics of neutrino-matter interactions, as well as neutrino detection through Cherenkov radiation by neutrino interaction products, in Chapter 4. This leads us directly to the comprehensive description of the IceCube detector in Chapter 5, which covers its operational principles and data acquisition methodologies. Finally, in Chapter 6, we introduce the reader to the Monte Carlo simulation and the event selection pipelines used in the two aforementioned analyses. The remainder of the thesis is devoted to the analyses themselves, followed by a short summary and an outlook towards future research directions.

I

Theory

THE STANDARD MODEL

This chapter describes the foundations of the Standard Model of particle physics – namely, its particle (field) constituents and the generalized prescriptions for their dynamics. The material in the sections below draws inspiration from [9, 55–58], which were used as the main references and can be consulted for completeness.

1.1 FIELDS AND PARTICLES

The most fundamental objects of Nature are **fields** – continuous entities permeating all of spacetime, such that each spacetime point $x^\mu = \langle t, x, y, z \rangle$ has a specific value[1] $\phi(x^\mu)$ of the field assigned [55]. Particles, on the other hand, can be viewed as localizations of the respective field's energy in spacetime, and the framework which describes free propagation and interactions of all particles is that of quantum field theory (QFT).

Just as classical field theories, QFTs are subject to the fundamental principle of least action, i.e., the requirement that

$$S = \int \mathcal{L}(\phi, \partial_\mu \phi) \, \mathrm{d}^4 x, \tag{1.1}$$

where $\mathcal{L}(\phi, \partial_\mu \phi)$ is the field's Lagrangian[2], is minimized. The dynamics of multiple free or interacting fields can be encoded in the master Lagrangian of the theory, i.e., that combining all participating fields ϕ_i. The corresponding equations of motion of all particles are then derived as Euler-Lagrange equations, one for each field ϕ_i:

$$\frac{\partial \mathcal{L}}{\partial \phi_i} - \partial_\mu \left(\frac{\partial \mathcal{L}}{\partial (\partial_\mu \phi_i)} \right) = 0. \tag{1.2}$$

The theory of special relativity postulates that the obtained equations of motion must be independent of the observer – i.e., have the same form in all inertial (non-accelerating) reference frames. Mathematically, this translates to the requirement for the Lagrangian $\mathcal{L}$ to preserve its structure when the fields ϕ_i transform into $\phi_i' = \Lambda \phi_i$ under Lorentz transformations Λ. A Lagrangian that satisfies this property is called **Lorentz invariant**.

---

[1] As we will see later in this section, this value does not need to be a simple scalar and is generally a multi-valued object obeying specific transformation rules.

[2] To be precise, $\mathcal{L}$ is the Lagrangian density, and the full Lagrangian is its volume integral, $L = \int \mathcal{L} \, \mathrm{d}^3 x$. It is this integral that is defined as the difference between the kinetic and the potential energy of the system, $L = T - V$.

The Lorentz transformations include Lorentz boosts (shifts between reference frames moving with a non-zero velocity $v = \beta c$ with respect to one another) and rotations in three spatial dimensions[3]. The particle fields of the Standard Model are classified according to their behavior under such transformations, yielding the broad classes of *scalar*, *vector*, and *spinor* fields. While the scalar fields are not affected by Lorentz transformations, this does not apply to the vector fields a_μ. They transform as four-vectors of special relativity, for example:

- $a_\mu \to R^x_{\mu\nu}(\theta)a^\nu$ under rotation about the x-axis by an angle θ;
- $a_\mu \to \Lambda^z_{\mu\nu}(v)a^\nu$ under the boost along the z-axis by velocity $\mathbf{v} = \langle 0, 0, v \rangle$,

where $R^x_{\mu\nu}(\theta)$ and $\Lambda^z_{\mu\nu}(v)$ are the standard 4×4 rotation and Lorentz boost matrices. The spinors obey even more complex transformation rules, particularly distinguishing between two fundamental types: *left-chiral* spinors, ψ_L, and *right-chiral* spinors, ψ_R, each represented as a two-component object. The ψ_L and ψ_R spinors transform identically under rotations but differently under boosts – and, crucially, convert into one another under the action of the parity operator $\mathbb{P}(\mathbf{x} \to -\mathbf{x})$:

$$\begin{aligned}
\mathbb{P}[\psi_L(\langle t, \mathbf{x} \rangle)] &= \psi_L(\langle t, -\mathbf{x} \rangle) = \psi_R(\langle t, \mathbf{x} \rangle); \\
\mathbb{P}[\psi_R(\langle t, \mathbf{x} \rangle)] &= \psi_R(\langle t, -\mathbf{x} \rangle) = \psi_L(\langle t, \mathbf{x} \rangle).
\end{aligned} \tag{1.3}$$

The chiral behavior of ψ_L and ψ_R is therefore encoded in Equation (1.3), which shows that each of these spinors transforms into a different physical object in the mirrored coordinate system.

These transformation rules have physical implications for the particle content of the Standard Model, which is shown in Figure 1.1. For example, fermions – which encompass quarks and leptons, and compose all of the ordinary matter – are superpositions of left-chiral and right-chiral spinors:

$$\psi = \begin{pmatrix} \psi_L \\ \psi_R \end{pmatrix}, \tag{1.4}$$

where ψ is a Dirac fermion that obeys the Dirac equation when propagating freely. The bosons that mediate the interactions between the fermions are vectors; this includes the photon, the gluon, and the three weak bosons ($W^\pm$ and Z^0). Finally, the Higgs boson is a scalar, which couples to the fermions[4] to give them mass (this will be explained in more detail in Section 1.3). Before we dive into the fermion mass generation mechanism, let us discuss the internal symmetries of the *massless* Standard Model Lagrangian, which, as we will see in Section 1.2, give rise to the aforementioned vector boson fields and their coupling to the fermions.

1.2 THE MASSLESS LAGRANGIAN AND ITS SYMMETRIES

The symmetry group of the Standard Model ("SM") is $SU(3)_C \times SU(2)_L \times U(1)_Y$[5], which implies that the SM Lagrangian is invariant with respect to the transformations that compose the specified product of the three groups. In this notation, "U" stands for **u**nitary,

$$\mathcal{U}\mathcal{U}^\dagger = \mathcal{U}^\dagger\mathcal{U} = \mathbb{I}, \tag{1.5}$$

---

[3]A larger group of transformations called the Poincaré group also includes spacetime translations.

[4]Strictly speaking, the Higgs couples only to the fermions which have both a left-chiral and right-chiral component. In the minimal version of the Standard Model, this is not the case for neutrinos (antineutrinos), which are purely left-chiral (right-chiral).

[5]The subscripts will be explained later in this section.
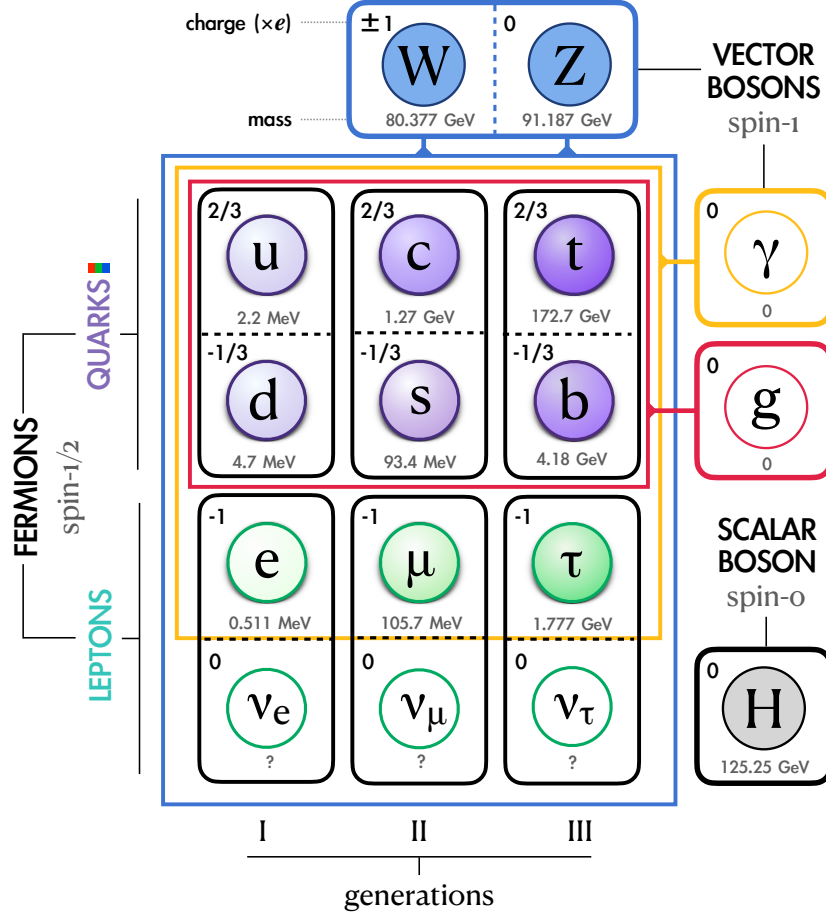
FIGURE 1.1 – The elementary particles of the Standard Model, including the matter particles (fermions), the force carriers (vector bosons), and the Higgs. The antifermions are not included in this figure but can be obtained from the respective fermions via charge conjugation. The listed particle masses correspond to the central fit values from the 2022 Particle Data Group Review [34]. *This figure was originally produced by the author for [55].*

and "S" – for **s**pecial, imposing an additional constraint on the transformations:

$$\det \mathcal{U} = 1. \tag{1.6}$$

The element U of a generic group $\mathrm{SU}(N)$ is a complex-valued matrix, which satisfies Equations (1.5) and (1.6) and is generated as

$$\mathcal{U} = \exp\left\{ \sum_{m=1}^{N^2-1} i\lambda_m T_m \right\} \equiv e^{i\lambda_m T_m} \tag{1.7}$$

from the $N^2 - 1$ group generators T_m[6]. This means that the SU(3) group has $3^2 - 1 = 8$ generators (resulting in 8 gluon fields), while SU(2) has 3 generators (yielding 3 weak boson fields). The U(1) group has only one generator – the 1×1 identity "matrix" itself (i.e., the number 1). Each of the three groups $\mathcal{G} \in \{\mathrm{SU}(3), \mathrm{SU}(2), \mathrm{U}(1)\}$ is a *local* symmetry of the SM Lagrangian, i.e., the latter remains invariant under the corresponding transformations $\mathcal{U} \in \mathcal{G}$ even if the multipliers λ_m in Equation (1.7) depend on the spacetime coordinates:

$$\lambda_m \to \lambda_m(x_\mu) \equiv \lambda_m(x). \tag{1.8}$$

---

[6]These generators are the basis elements of the corresponding algebra. Although the most general $N \times N$ complex matrix $\mathcal{U}$ has $2N^2$ degrees of freedom, N^2 of them are fixed by the unitarity constraints, and one more – by $\det\mathcal{U} = 1$ [58].

Below, we show how the different pieces of the SM Lagrangian emerge from the local symmetry requirement for each of the three groups. Note that we will use different notations for the generator coefficients λ_m as is conventional for each group.

U(1) symmetry and the QED Lagrangian

Let us start with a Lagrangian for a free Dirac fermion ψ, where, for the time being, we will assume that this fermion is massless:

$$\mathcal{L}_{\text{Dirac}} = -m\bar{\psi}\psi + i\bar{\psi}\gamma_\mu\partial^\mu\psi \xrightarrow{m=0} i\bar{\psi}\gamma_\mu\partial^\mu\psi, \tag{1.9}$$

where γ^μ are the usual "gamma matrices" and $\bar{\psi} = \psi^\dagger\gamma^0$. The Lagrangian in Equation (1.9) is Lorentz-invariant and has a global U(1) symmetry (is invariant under spacetime-independent transformations $\mathcal{U} = e^{i\alpha}$):

$$\psi \to \psi' = e^{i\alpha}\psi \implies \tag{1.10}$$

$$\mathcal{L}_{\text{Dirac}}(\psi, \partial_\mu\psi) \to \mathcal{L}'_{\text{Dirac}}(\psi', \partial_\mu\psi') = \mathcal{L}_{\text{Dirac}}. \tag{1.11}$$

Since fixing α to the same value at every point in spacetime is not favored by special relativity[7], our goal is promote this to a local symmetry, i.e., invariance under

$$\mathcal{U}(x) = e^{i\alpha(x)} = e^{iga(x)}. \tag{1.12}$$

In the above, we explicitly separated $\alpha(x)$ into a constant g and a spacetime-dependent part $a(x)$. The Lagrangian in Equation (1.9) is **not** invariant under such $\mathcal{U}(x)$:

$$\psi \to \psi' = e^{iga(x)}\psi \implies \tag{1.13}$$

$$\mathcal{L}_{\text{Dirac}}(\psi, \partial_\mu\psi) \to \mathcal{L}'_{\text{Dirac}}(\psi', \partial_\mu\psi') = \mathcal{L}_{\text{Dirac}} - g\left(\partial_\mu a(x)\right)\bar{\psi}\gamma^\mu\psi. \tag{1.14}$$

We see that the introduction of the spacetime dependence into $\mathcal{U}$ through $a(x)$ leads to the appearance of an extra term in the Lagrangian, which is proportional to $\partial_\mu a(x)$.

To restore the invariance of the Lagrangian under $\mathcal{U}$, we introduce a vector (spin-1) field A_μ, which has an internal local symmetry with respect to a shift by $\partial_\mu a(x)$. In other words, the following transformation leaves the Lagrangian of a free A_μ invariant:

$$A_\mu \to A'_\mu = A_\mu + \partial_\mu a(x). \tag{1.15}$$

Now, if we let the fermion ψ couple to such a field (for example, an electron coupling to a photon), we must introduce the a new term, $gA_\mu\bar{\psi}\gamma^\mu\psi$ into Equation (1.9). Under $\mathcal{U}$ acting on ψ and the internal symmetry of A_μ, this term transforms as

$$gA_\mu\bar{\psi}\gamma^\mu\psi \to gA'_\mu\bar{\psi}'\gamma^\mu\psi' = gA_\mu\bar{\psi}\gamma^\mu\psi + g\left(\partial_\mu a(x)\right)\bar{\psi}\gamma^\mu\psi. \tag{1.16}$$

We notice that the last term in Equation (1.16) cancels exactly that in Equation (1.14); therefore, if our fermions ψ are allowed to couple to the external spin-1 fields A_μ, the Lagrangian is invariant under local U(1) transformations. The constant g can be now understood as the coupling strength. The final ingredient in such a theory is the dynamics of a free A_μ, which is encoded in the tensor $F_{\mu\nu}$ (the *strength* of the A_μ field):

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu. \tag{1.17}$$

---

[7]This would imply that we are capable of making such a choice immediately for everywhere in the Universe, which contradicts the fundamental principle that no information may be propagated faster than the speed of light [56].

To add this term to the Lagrangian, we must make it Lorentz-invariant, which is achieved by contracting this tensor with itself. The final Lagrangian of our theory becomes

$$\mathcal{L}_{\text{QED}} = i\bar{\psi}\gamma_\mu\partial^\mu\psi + gA_\mu\bar{\psi}\gamma^\mu\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu}, \tag{1.18}$$

where **QED** stands for **Q**uantum **E**lectro**d**ynamics – the part of the Standard Model that describes *electromagnetic* interactions of charged fermions ψ, which are mediated by the photon A_μ. Equation (1.18) can be conveniently rewritten if we define the *covariant derivative*,

$$D^\mu = \partial^\mu - igA^\mu, \tag{1.19}$$

which leads to[8]

$$\mathcal{L}_{\text{QED}} = i\bar{\psi}\gamma_\mu D^\mu\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} \equiv i\bar{\psi}\slashed{D}\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu}. \tag{1.20}$$

At last, we note that the *global* U(1) symmetry leads to the conservation of the electric charge Q according to the Noether's theorem [59]. The corresponding conserved quantity for the global SU(2) symmetry (discussed in the next subsection) is the isospin I. These two quantities are related through the Gell-Mann–Nishijima formula [60, 61]:

$$Q = I_3 + \frac{1}{2}Y, \tag{1.21}$$

where I_3 is the third component of the isospin and Y is the hypercharge. In a combined electroweak theory with an unbroken SU(2) × U(1) symmetry, Y is therefore conserved, which is why U(1) is conventionally given the "Y" subscript.

SU(2) symmetry and the Electroweak Lagrangian

Next, let us consider the SU(2) transformations, which are fundamental to understanding the *weak* force mediating the interactions of neutrinos. The transformations of the SU(2) group rotate the members of the weak *doublets* into one another. The members of the same doublet share a property called **flavor** – for example, the electron e^- and the electron neutrino ν_e form a lepton doublet associated with the electron flavor. The full set of doublets includes (e^-, ν_e), (μ^-, ν_μ) and (τ^-, ν_τ), representing the three generations in the leptonic sector. The same doublet structure applies to quarks, which are also arranged in three generations of doublets: (u, d), (c, s), and (t, b)[9]. The SU(2) symmetry of the Lagrangian implies freedom in the choice of the orientation of the doublet axes – and, if the symmetry is local, there is such freedom at every point in spacetime.

Let $\Psi = (\psi_1, \psi_2)^\top$ be a generic doublet, where ψ_1 and ψ_2 are Dirac fermions. The Lagrangian for this system (where the fermions are again assumed to be massless) is

$$\mathcal{L}(\psi_1, \psi_2) = i\bar{\psi}_1\gamma_\mu\partial^\mu\psi_1 + i\bar{\psi}_2\gamma_\mu\partial^\mu\psi_2 \tag{1.22a}$$

$$= i\bar{\Psi}\gamma_\mu\partial^\mu\Psi \equiv \mathcal{L}(\Psi). \tag{1.22b}$$

This Lagrangian is invariant the global SU(2) transformations, i.e.,

$$\mathcal{U} = e^{\frac{ia_i\sigma_i}{2}}, \tag{1.23}$$

---

[8]Note that the factor of $1/4$ is included before $F_{\mu\nu}F^{\mu\nu}$ in order to have the conventional normalization of the kinetic energy of the A_μ field [58].

[9]The reason behind the existence of precisely three generations is presently unknown.

where index i now runs from 1 to 3, and σ_i are the Pauli matrices (generators of SU(2) in the two-dimensional representation). However, if we attempt to promote this to a local symmetry by introducing the spacetime dependence to the generator coefficients ($a_i \to a_i(x)$), we will find that $\mathscr{L}(\Psi)$ is no longer invariant under such a transformation, as extra terms appear in the transformed $\mathscr{L}'(\Psi')$. In case of the U(1) symmetry discussed earlier, we fixed this by introducing a spin-1 field A_μ that couples to the fermion ψ with a coupling constant g. A similar idea applies to SU(2), except now we have to introduce three spin-1 fields W_i^μ, one for each generator σ_i, which couple to the fermions with a different strength g'. The covariant derivative in this case becomes [56]

$$D^\mu = \partial^\mu - ig' \frac{\sigma_i W_i^\mu}{2} \equiv \partial^\mu - ig' \mathcal{W}^\mu, \tag{1.24}$$

and the final locally SU(2)-invariant and Lorentz-invariant Lagrangian (which also includes the dynamics of the W_i^μ fields themselves) reads[10]:

$$\mathscr{L}_{\text{loc. SU(2)}} = i\bar{\Psi}\slashed{D}\Psi - \frac{1}{4}\text{Tr}\left[\mathcal{W}_{\mu\nu}\mathcal{W}^{\mu\nu}\right]. \tag{1.25}$$

Combining this with a Lagrangian that is locally invariant with respect to both U(1) and SU(2), we get the *electroweak* ("EW") theory with massless fields:

$$\mathscr{L}_{\text{loc. U(1)+SU(2)}} \equiv \mathscr{L}_{\text{EW}} = i\bar{\Psi}\slashed{D}\Psi - \frac{1}{4}B_{\mu\nu}B^{\mu\nu} - \frac{1}{2}\text{Tr}\left[\mathcal{W}_{\mu\nu}\mathcal{W}^{\mu\nu}\right] \tag{1.26a}$$

$$= i\bar{\Psi}\gamma_\mu(\partial^\mu - igB^\mu - ig'\mathcal{W}^\mu)\Psi - \frac{1}{4}B_{\mu\nu}B^{\mu\nu} - \frac{1}{2}\text{Tr}\left[\mathcal{W}_{\mu\nu}\mathcal{W}^{\mu\nu}\right]. \tag{1.26b}$$

Note that in Equation (1.26b), we replaced $A^\mu \to B^\mu$, $F_{\mu\nu} \to B_{\mu\nu}$ (compared to Equation (1.20)), since in the combined EW theory, B^μ is not immediately identified as the photon. In the same way, W_i^μ, $i \in \{1, 2, 3\}$, are not immediately identified as the familiar $W^\pm$ and Z^0. In fact, both the photon and the weak bosons are linear combinations of the B^μ and the W_i^μ fields, which will be explained in Section 1.3.1.

To finish this section, we note that the weak force maximally violates parity, which was discovered in the famous ^{60}Co decay experiment by Wu [29]. This implies that the weak bosons couple only to the left-chiral fermions (and right-chiral antifermions) of the Standard Model. For this reason, SU(2) is usually given the subscript "L" in the definition of the Standard Model symmetry group.

SU(3) symmetry and the QCD Lagrangian

Although not of prime relevance for the remainder of this thesis, we complete our discussion of the SM symmetries with the massless Lagrangian of **Q**uantum **C**hromo**d**ynamics (QCD), which describes *strong* interactions of quarks:

$$\mathscr{L}_{\text{QCD}} = i\bar{\mathcal{Q}}\slashed{D}\mathcal{Q} - \frac{1}{2}\text{Tr}\left[\mathcal{G}_{\mu\nu}\mathcal{G}^{\mu\nu}\right]. \tag{1.27}$$

In Equation (1.27), $\mathcal{Q} \equiv (q_1, q_2, q_3)^\top$ are quark *color triplets*, where the indices 1, 2, 3 are typically given the color labels of red ("r"), green ("g"), and blue ("b"). The local SU(3)

---

[10]This is almost completely analogous to Equation (1.20), except $\mathcal{W}_{\mu\nu} = \partial^\mu \mathcal{W}^\nu - \partial^\nu \mathcal{W}^\mu - ig'[\mathcal{W}^\mu, \mathcal{W}^\nu]$, and the trace (Tr) is needed to ensure the SU(2) invariance of the $\mathcal{W}_{\mu\nu}\mathcal{W}^{\mu\nu}$ object [56].

symmetry implies that the orientation of the color axes at each point in spacetime might be arbitrary, and the eight gluon fields G_i^μ are introduced as connections between the different color coordinate systems[11]. By analogy with our discussion of U(1) and SU(2) groups, it is these connector fields that insure the local invariance of the Lagrangian, whose covariant derivative is

$$D^\mu = \partial^\mu - ig''\mathcal{G}^\mu = \partial^\mu - ig''\frac{\lambda_i G_i^\mu}{2}, \tag{1.28}$$

where we defined a new coupling constant g'' and new generators λ_i of the SU(3) transformations. These generators are referred to as the Gell-Mann matrices, and $\mathcal{G}_{\mu\nu}$ in Equation (1.27) is the gluon field strength tensor analogous to $\mathcal{W}_{\mu\nu}$ in Equation (1.26b).

1.3 HIGGS MECHANISM

So far, we have worked with massless fermion fields, whose interactions were mediated by the massless bosons. From experiments, however, we know that most of the fermions and some of the bosons are massive. In the sections below, we will explain how their corresponding masses are generated through the coupling to the scalar Higgs field – the process accordingly known as the **Higgs mechanism**.

1.3.1 *Gauge boson mass generation*

As the Higgs field is still missing from the electroweak Lagrangian in Section 1.2, we introduce a complex scalar doublet $\Phi = (\phi_1, \phi_2)^\top$. A locally SU(2) and U(1) invariant Lagrangian for this field is

$$\mathcal{L}(\Phi, D_\mu\Phi) = (D_\mu\Phi^\dagger)(D^\mu\Phi) + \rho^2\Phi^\dagger\Phi - \lambda(\Phi^\dagger\Phi)^2, \tag{1.29}$$

where the first term is the kinetic term for Φ, with the covariant derivative

$$D^\mu = \partial^\mu - igB^\mu - ig'\mathcal{W}^\mu \tag{1.30}$$

as in Equation (1.26b). As before, the coupling of the Higgs to the B^μ and the $\mathcal{W}^\mu$ fields is introduced to ensure the specified local symmetries. The remainder of the Lagrangian is the negative Higgs potential:

$$V(\Phi) = -\rho^2\Phi^\dagger\Phi + \lambda(\Phi^\dagger\Phi)^2. \tag{1.31a}$$

$$= -\rho^2(\phi_1^\dagger\phi_1) + \lambda(\phi_1^\dagger\phi_1)^2 - \rho^2(\phi_2^\dagger\phi_2) + \lambda(\phi_2^\dagger\phi_2)^2 \tag{1.31b}$$

$$\equiv V_1(\phi_1) + V_2(\phi_2). \tag{1.31c}$$

This potential is visualized in Figure 1.2 for one of the doublet fields ϕ_i (we will use $\phi = \phi_i$ for brevity). We see that this potential has an infinite number of minima,

$$\phi_{\min} = \sqrt{\frac{\rho^2}{\lambda}}e^{i\theta}, \; \theta \in [0, 2\pi], \tag{1.32}$$

which are also referred to as the *vacua* of the field ϕ and represent the states of lowest energy. One of these infinitely many minima was spontaneously chosen in the early Universe, thereby breaking the global U(1) symmetry of ϕ – the phenomenon known as **spontaneous symmetry breaking**[12]. By convention, the minimum of the doublet Φ is

---

[11] For this reason, SU(3) is also referred to as the color group with the subscript "C".

[12] The Lagrangian in Equation (1.29) is still globally U(1) invariant, but the U(1) transformations applied to ϕ result in a different physical state ϕ'. This would not be the case if the ground state was at 0, where the $U(1)$ rotations would result in the same unique state.
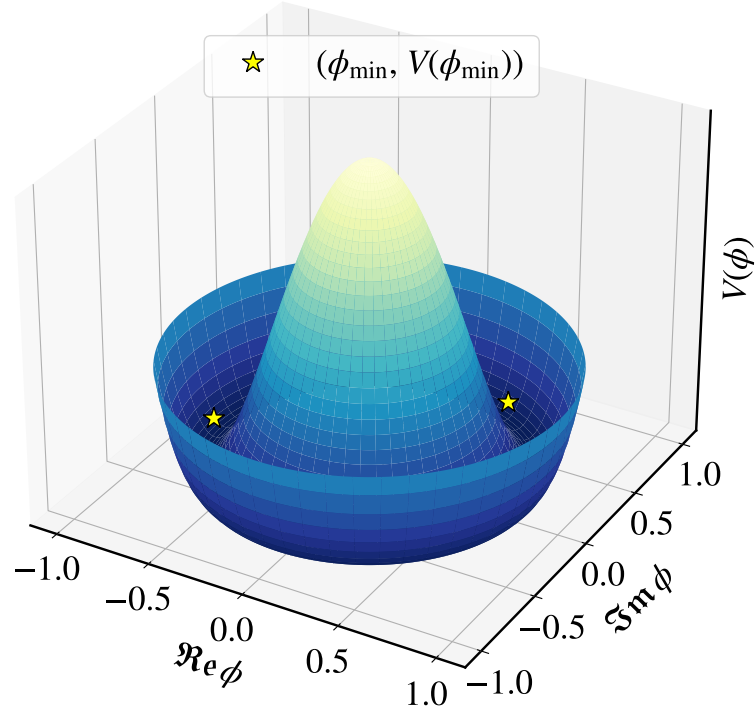
FIGURE 1.2 – "Mexican hat" potential $V(\phi)$ of a complex scalar field ϕ, shown as a function of the real and the imaginary component of ϕ. This plot was produced with $\rho = 1, \lambda = 0.8$ (see Equation (1.31)). The yellow stars indicate two of the infinite number of minima of the potential.

then economically constructed as [56]

$$\Phi_{\min} = \begin{pmatrix} 0 \\ \phi_{2,\min} \end{pmatrix} = \begin{pmatrix} 0 \\ \frac{v}{\sqrt{2}} \end{pmatrix}, \tag{1.33}$$

where we defined the **vacuum expectation value** v,

$$v = \frac{\rho^2}{\lambda}. \tag{1.34}$$

Expanding each of the doublet components about the chosen minimum, we end up with four real scalar fields (by adding a real and an imaginary component to each of the rows in Equation (1.33)). Three of these fields can be eliminated by exploiting the local SU(2) symmetry of the Lagrangian in Equation (1.29) (remember that there are 3 generators and therefore 3 degrees of freedom in SU(2)). By making a gauge choice called a *unitary gauge*, we can therefore write

$$\Phi = \Phi_{\min} + \begin{pmatrix} \phi_{\text{re},1} + i\phi_{\text{im},1} \\ \phi_{\text{re},2} + i\phi_{\text{im},2} \end{pmatrix} = \begin{pmatrix} \phi_{\text{re},1} + i\phi_{\text{im},1} \\ \frac{v}{\sqrt{2}} + \phi_{\text{re},2} + i\phi_{\text{im},2} \end{pmatrix} \tag{1.35a}$$

$$\xrightarrow{\text{unitary gauge}} \begin{pmatrix} 0 \\ \frac{v+h}{\sqrt{2}} \end{pmatrix}. \tag{1.35b}$$

Here, the only remaining physical field is the **Higgs field** h. This explicit gauge choice breaks the local SU(2) symmetry, and the remnant symmetry of the Standard Model is

$U(1)_Q$. On the bright side, if we define the following fields:

$$A^\mu = g W_3^\mu + g' B_\mu; \tag{1.36a}$$

$$Z^\mu = g' W_3^\mu - g B_\mu; \tag{1.36b}$$

$$(W^\pm)^\mu = W_1^\mu \pm i W_2^\mu, \tag{1.36c}$$

and reexpress the Lagrangian in terms of these fields and Φ from Equation (1.35b), we will find that Z^μ and $(W^\pm)^\mu$ acquire mass terms, with their masses being proportional to v. Explicitly,

$$M_W = \frac{v g'}{2} \simeq 80.4\,\text{GeV}; \tag{1.37a}$$

$$M_Z = \frac{v\sqrt{g^2 + g'^2}}{2} = M_W \frac{\sqrt{g^2 + g'^2}}{g'} = \frac{M_W}{\cos\theta_W} \simeq 91.2\,\text{GeV}, \tag{1.37b}$$

where θ_W is the Weinberg angle ($\sin^2\theta_W = \frac{g^2}{g^2+g'^2} \simeq 0.223$ [34]). Now, we are ready to identify the fields in Equation (1.36) as the gauge bosons of the Standard Model, of which three are massive and one (the photon) is massless. We reiterate that the masses of $W^\pm$ and Z were acquired at the expense of introducing a scalar field, letting it couple to the initially massless bosons, minimizing the potential energy of this field by choosing one of the infinitely many possible ground states, and finally breaking the local SU(2) invariance.

1.3.2 *Fermion mass generation*

Although the weak bosons have now acquired mass through their coupling to the Higgs, the fermion fields ψ_i entering the doublet $\Psi = (\psi_1, \psi_2)^\top$ so far are still massless in the locally SU(2) and U(1)invariant Lagrangian (see Equation (1.26b)). Specifically, the latter is missing terms of the form $-m\bar{\psi}_i\psi_i$, where m would be identified as the mass of ψ_i. These terms cannot be immediately inserted into the Lagrangian without breaking the SU(2) invariance. This is due to the fact that, as explained in Section 1.1, a physical fermion ψ satisfying the Dirac equation is always a superposition of left-chiral and right-chiral fields (see Equation (1.4)). Therefore,

$$\bar{\psi}\psi = \bar{\psi}_L \psi_R + \bar{\psi}_R \psi_L. \tag{1.38}$$

However, SU(2) transformations only affect the left-chiral fields, and therefore the doublets $\Psi_L = (\psi_{1L}, \psi_{2L})^\top$ transform under two-dimensional representations of SU(2). The right-chiral fields ψ_R, on the other hand, are singlets of SU(2), meaning that they do not transform under this group (physically, this is interpretable as the fact as only the left-chiral fields couple to the weak force carriers). Therefore, the terms $\bar{\Psi}_L \psi_R$ and $\bar{\psi}_R \Psi_L$, which naively would have been part of the Lagrangian, are not SU(2) invariant, and it is not possible to include them immediately.

To resolve this, we use the same trick as in Section 1.3.1 and let the fermion fields ψ_1 and ψ_2 couple to the scalar doublet Φ. Specifically, we introduce the following terms to the Lagrangian:

$$-y_2\left[\bar{\Psi}_L \Phi \psi_{2R} + \bar{\psi}_{2R} \bar{\Phi} \Psi_L\right] \tag{1.39a}$$

$$\text{and} \quad -y_1\left[\bar{\Psi}_L \widetilde{\Phi} \psi_{1R} + \bar{\psi}_{1R} \widetilde{\bar{\Phi}} \Psi_L\right], \tag{1.39b}$$

where y_i are the constants defining the ψ_i-Higgs coupling strength (also called *Yukawa coupling constants*), and $\widetilde{\Phi}$ is the charge conjugate of Φ. Once again, Φ randomly

chooses one of the lowest-energy configurations in the process of spontaneous symmetry breaking. By making the same choice of this configuration as in Equation (1.33) and choosing the unitary gauge as in Equation (1.35b), we arrive at the Lagrangian with the fermion mass terms:

$$\mathscr{L}_{\text{massive }\Psi} = \overbrace{\mathscr{L}_{\text{massless }\Psi}}^{\text{Equation (1.26)}} - \sum_{i=1,2} \overbrace{\frac{y_i v}{\sqrt{2}}}^{m_{\psi_i}} (\bar{\psi}_i \psi_i) - \sum_{i=1,2} \overbrace{\frac{y_i h}{\sqrt{2}}}^{\psi_i\text{-Higgs coupling}} (\bar{\psi}_i \psi_i) \quad . \tag{1.40}$$

Note that in Equation (1.40), we identified

$$m_{\psi_i} = \frac{y_i v}{\sqrt{2}} \tag{1.41}$$

as the mass of the fermion ψ_i[13], where v, as before, is the vacuum expectation value of the Higgs field ($\simeq 246\,\text{GeV}$).

In our discussion so far, the components ψ_1 and ψ_2 of the fermionic doublet Ψ were completely generic. Let us see what happens when we for concreteness choose the leptonic doublet with $\psi_1 = e^-$, $\psi_2 = \nu_e$. We know from experiments that neutrinos are *almost* massless, and when the Standard Model was originally formulated, they were postulated to have zero mass. The latter statement is equivalent to postulating the absence of right-chiral neutrinos, since the mass terms must combine left-chiral and right-chiral fields. In this "minimal" Standard Model framework, the only non-zero Yukawa constant is therefore $y_1 \equiv y_e$, yielding the mass of the electron.

This simplified treatment applies to the case of *one* leptonic generation. We know, however, that there are *three* generations of leptons, as per Figure 1.1. This means that we could in principle introduce the following mass and lepton-Higgs coupling terms for the charged leptons ℓ [9]:

$$-\sum_{\alpha,\beta \in \{e,\mu,\tau\}} Y'_{\alpha\beta} \bar{\ell}'_{\alpha\text{L}} \Phi \ell'_{\beta\text{R}} = \left(\frac{v+h}{\sqrt{2}}\right) \sum_{\alpha,\beta \in \{e,\mu,\tau\}} Y'_{\alpha\beta} \bar{\ell}'_{\alpha\text{L}} \ell'_{\beta\text{R}}, \tag{1.42}$$

where Y' is a generally non-diagonal matrix of Yukawa couplings. Thus, ℓ'_α are not immediately interpretable as the states of definite mass. To obtain the latter, one must diagonalize Y' with appropriate unitary matrices V_L and V_R,

$$Y = V_\text{L}^\dagger Y' V_\text{R}, \tag{1.43}$$

transforming Equation (1.42) as

$$\left(\frac{v+h}{\sqrt{2}}\right) \sum_{\alpha,\beta \in \{e,\mu,\tau\}} Y_{\alpha\beta} \bar{\ell}_{\alpha\text{L}} \ell_{\beta\text{R}}. \tag{1.44}$$

Now, Y is diagonal ($Y_{\alpha\beta} = y_\alpha \delta_{\alpha\beta}$), and the states ℓ_α such that

$$\ell_\alpha = \ell_{\alpha\text{L}} + \ell_{\alpha\text{R}} = V_\text{L}^\dagger \ell'_\text{L} + V_\text{R}^\dagger \ell'_\text{R} \tag{1.45}$$

have definite masses as per Equation (1.41). This notation allows us to define neutrinos of flavor α as the fields $\nu_{\alpha\text{L}}$ that weakly couple to the charged leptons ℓ_α with definite mass.

---

[13]After all, we see that it was possible to include the terms of the desired form $\bar{\psi}_i \psi_i$, but only thanks to the introduction of the Φ doublet, which also transforms after SU(2) and cancels the effect of this transformation acting on the left-chiral fermion fields.

Neutrinos

We now shift our attention to the main actors of this thesis – neutrinos, which we introduced very briefly in the last chapter. First, in Section 2.1, we will summarize very generally what we know and do not know about neutrinos. We will subsequently focus on the topic of neutrino masses in Section 2.2, covering, in particular, the current experimental limits and how – if not via the Higgs mechanism – neutrino masses could be generated. Finally, we will set up the framework of neutrino oscillations – experimentally observed transitions between different neutrino flavors. This will cover the standard three-flavor ("unitary") oscillation formalism, which we further expand to non-unitary oscillations.

N.B.: Sections 2.3 and 2.4 of this chapter contain partial reproduction of Sections 2.1, 2.2, and 2.3 of the following paper:

Tetiana Kozynets, Philipp Eller, Alan Zander, Manuel Ettengruber, and David Jason Koskinen, *Constraints on Non-Unitary Neutrino Mixing in Light of Atmospheric and Reactor Neutrino Data*, submitted to J. High Energy Phys. (2024).

At the time of writing, the paper is under review by the Journal of High Energy Physics. The full pre-print of the paper is available online in the arXiv database [1].

2.1 THE PROPERTIES OF NEUTRINOS

2.1.1 *The known properties*

Neutrinos are neutral leptons, which appear in the Standard Model as components of three flavor doublets together with the respective charged leptons. Accordingly, there are three known and experimentally observed neutrino flavors – electron, muon, and tau neutrinos (antineutrinos), denoted as $\nu_e(\bar{\nu}_e)$, $\nu_\mu(\bar{\nu}_\mu)$, and $\nu_\tau(\bar{\nu}_\tau)$. As mentioned at the very end of Section 1.3.2, each of these flavor states weakly couples to the lepton of the same flavor and a definite mass. This coupling occurs in the so-called *charged current interactions* (see more in Section 4.1), where a neutrino ν_α is rotated into the charged lepton ℓ_α via the SU(2) transformation at the $W^\pm$ vertex. In the reverse process, an unstable charged lepton ($\mu^\pm$ or $\tau^\pm$) decays, among other particles, into a neutrino of the same flavor. *Neutral current* neutrino interactions are also possible, which are mediated by Z^0 instead of $W^\pm$. The generic charged and neutral current vertices and

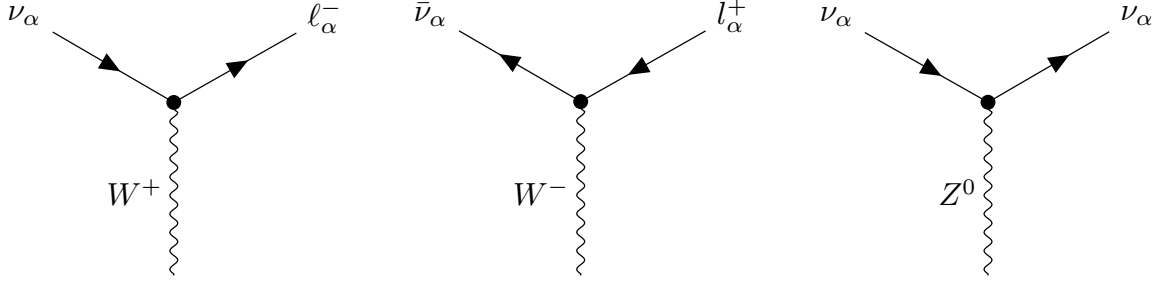the Feynman diagrams illustrating charged lepton decays are shown in Figures 2.1 and 2.2, respectively.

FIGURE 2.1 – *Left and middle*: charged current neutrino and antineutrino vertices. *Right*: neutral current neutrino vertex.
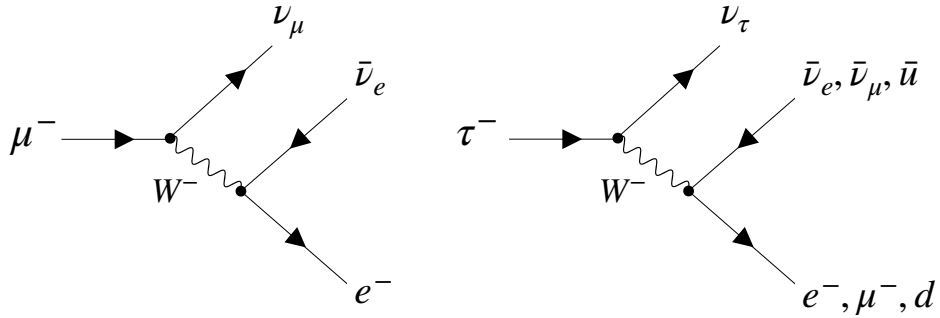
FIGURE 2.2 – Muon and tau lepton decays. Note that the tau lepton is relatively heavy ($m_\tau \simeq 1.7\,\text{GeV}$) and can decay both leptonically and hadronically, which is illustrated in the right panel.

While charged current weak decays are the dominant channel of neutrino production in natural environments[1], it is the neutral current decays that help us understand how many active (i.e. coupled in any way to the weak bosons) neutrino flavors are there. In particular, the measurements of the decay width of the Z^0 produced in the e^+e^- collisions at the LEP collider constrain the number of active neutrino species to $N_\nu = 2.984 \pm 0.008$ [64]. This number is within 2σ of (and thus, claimed to be in a good agreement with) $N_\nu^{\text{SM}} = 3$ – the three active neutrino flavors of the Standard Model.

Most importantly for this thesis, neutrinos have been experimentally observed to change their flavor from the point of production to the point of detection. This phenomenon, known as **neutrino oscillations**, is consistent with the picture of neutrinos propagating as a superposition of three states $|\nu_i\rangle$ with definite and non-degenerate masses m_i[2]. Explicitly, most of the neutrino oscillation data gathered to date can be explained if one lets

$$|\nu_\alpha\rangle = \sum_{i=1}^{3} U_{\alpha i}^* |\nu_i\rangle, \tag{2.1}$$

where $\alpha \in \{e, \mu, \tau\}$, $i = 1..3$, and U is a 3×3 unitary matrix which we describe in detail in Section 2.3. In this framework, the mass splittings

$$\Delta m_{ij}^2 = m_i^2 - m_j^2 \tag{2.2}$$

---

[1]Note that in exotic environments such as the supernovae, neutrinos can also be produced via the neutral current e^+e^- annihilation; see e.g. [62, 63].

[2]The wording "non-degenerate" implies that at least two of the three neutrino mass states must have non-zero masses, and that the latter must not be equal to each other.

have been experimentally measured, with the 2022 global fit values being $\Delta m_{31}^2 = 2.511 \cdot 10^{-3}\,\text{eV}^2$ and $\Delta m_{21}^2 = 7.41 \cdot 10^{-5}\,\text{eV}^2$ [65, 66]. The reported values assume the mass state $|\nu_3\rangle$ is the heaviest, that $m_3 > m_2 > m_1$. However, the latter has not been decisively confirmed through experimental evidence, which leads us directly to the topic of the unknown neutrino properties.

2.1.2 *The unknown properties*

Below, we provide a non-exhaustive list of fundamental yet so far undetermined neutrino properties, which are of relevance to the topic of this thesis.

- **Neutrino mass ordering:** The measurement of the solar (Δm_{21}^2) mass splitting indicates that the second mass eigenstate is heavier than the first. However, the long-baseline (atmospheric or accelerator) neutrino data can at present accommodate either $\Delta m_{31}^2 > 0$ or $\Delta m_{32}^2 < 0$ – the scenarios in which the third mass eigenstate is either the heaviest or the lightest. These scenarios are referred to as the normal mass ordering/hierarhy ("NO") and the inverted mass ordering ("IO"), respectively. Currently, the NO is preferred at a $2.7\,\sigma$ level (considering joint constraints from neutrino oscillation data and cosmological measurements [67]), which is considered a moderate level of evidence. The combination of the upcoming JUNO and IceCube-Upgrade experiments is expected to reject the wrong ordering with a 5σ significance within 3–7 years of joint operation [68].

- **Dirac *vs.* Majorana nature:** All charged leptons of the Standard Model are Dirac fermions, such that the particles ℓ^- are physically distinct from the corresponding antiparticles ℓ^+. Neutrino, however, is electrically neutral and in principle could be its own antiparticle (a so-called *Majorana fermion*). In this case, the neutrino field ν is equivalent to $\nu^C \equiv C\bar{\nu}^\top$, where $C = i\gamma^2\gamma^0$ in the chiral basis [69, 70]. One possible experimental signature of the Majorana nature of neutrinos would be neutrinoless double beta decay, i.e., a nuclear decay of the type [71]

$$(A, Z) \rightarrow (A, Z+2) + e^- + \cancel{2\bar{\nu}_e}, \tag{2.3}$$

  where A is the atomic mass and Z is the nuclear charge. In this "$0\nu\beta\beta$" process, no final state electron antineutrinos are observed, as they are emitted in one beta decay and absorbed in the other. Current and upcoming experiments searching for this signature include CUORE/CUPID [72, 73], LEGEND [74], NEXT [75], and nEXO [76]. The neutrino mass ordering would have implications for the effective Majorana neutrino mass measured from such experiments ($m_{\beta\beta}$; see Section 2.2.1), with the NO favoring lower effective masses [77].

- **CP violation:** While weak interactions maximally violate parity and charge conjugation symmetries (corresponding to the reversal of spatial coordinates and the flipping of all "charge-like" quantum numbers, respectively [69]), the combined CP symmetry is still *mostly* preserved. Notable exceptions appear in the quark sector, where differences between decay rates of mesons and antimesons have been observed in K^0 [78–80], B^0 [81, 82], B^+ [83–85], and D^0 [86] decays. The "strength" of the CP violation in the quark sector can be condensed to the phase δ_{13} of the quark mixing (CKM [87, 88]) matrix, where $\delta_{13} = 0$ or $\pm\pi$ would indicate no CP violation and $\delta_{13} = \pi/2$ – maximal CP violation. The current global measurement is $\delta_{13} = (1.144 \pm 0.027)\,\text{rad} = (65.5 \pm 1.5)\,\text{deg}$ [34], corresponding to percent-level precision. At the same time, this degree of CP violation is not sufficient to explain the observed matter-antimatter asymmetry in the Universe [89, 90]. This

raises the question of additional sources of CP violation, in particular in the leptonic sector, which could be probed by comparing neutrino and antineutrino oscillations. The measurements of the corresponding $\delta_{\rm CP}$ phase in the neutrino sector (see Section 2.3) have been performed independently by the T2K [91] and NOνA [92] Collaborations, originally resulting in a mild tension for the normal neutrino mass ordering scenario (with $\delta_{\rm CP}^{\rm T2K, NO} = -1.97^{+0.97}_{-0.70}$ rad $= -113^{+56}_{-40}$ deg, which was excluded with 2σ confidence by the NOνA data). The preliminary results of the joint T2K+NOνA analysis indicate that CP conservation is excluded at $> 3\sigma$ in the IO scenario, while it is favored in the NO case [93]. Given an overall poor precision of the CP phase measurements in the neutrino sector compared to the quark sector, the efforts in this domain will continue in the next years, particularly by the DUNE [94], Hyper-Kamiokande [95], and KM3NeT/ORCA [96] Collaborations, whose projected 1σ $\delta_{\rm CP}$ resolutions are ~10–20 deg[3] in 10 years of operation.

- **Existence of right-chiral/sterile neutrinos:** While the Z^0 decay width provides a constraint on the number of the active neutrino species, it does not constrain the number of possible *sterile* states, i.e., neutrinos which do not participate in weak interactions. Since we know that the weak bosons only couple to the left-chiral fermions and the right-chiral antifermions, sterile neutrinos are also associated with the right-chiral neutrino (left-chiral antineutrino) fields. The light (eV-scale) sterile neutrinos have been proposed to explain the short-baseline neutrino oscillation anomalies (particularly in LSND [97], SAGE [98], GALLEX [99], and MiniBooNE [100, 101] experiments), but have been largely excluded by recent observations [102–105]. Heavier (keV-scale) sterile neutrinos have further been proposed as dark matter candidates, particularly in connection with the 3.55 keV line observed in the X-ray data, which possibly indicates sterile neutrino decays [106–108]. Finally, *very heavy* right-chiral neutrinos, with masses much larger than the electroweak scale defined by the vacuum expectation value of the Higgs field (246 GeV), are motivated from the perspective of neutrino mass generation [35, 36], as we discuss in more detail in Section 2.2.2.

- **Unitarity of neutrino mixing:** The possibility of existence of sterile states leads directly to the main question raised in this thesis – namely, whether the standard 3×3 neutrino mixing matrix is unitary. If the extra states do exist, the three-flavor matrix becomes a part of a larger unitary matrix (that including the sterile states) and is itself non-unitary. While the extremely heavy states can be kinematically inaccessible (depending on the energy reach of a given experiment), the very fact of their existence affects the active part of the mixing matrix. Specifically, the normalizations of rows and columns of the latter become smaller than unity, which ultimately propagates into the three-flavor neutrino oscillation probabilities. Rather stringent (sub-percent level) constraints on the neutrino mixing unitarity have been placed through e.g. charged lepton decay measurements, searches for lepton flavor violating processes, and weak decay universality tests (see [109] for a full scope of tests). At the same time, "in-situ" constraints utilizing neutrino oscillation data alone yield much looser bounds, ranging from $\mathcal{O}(1\%)$ in the electron sector to tens of % in the tau sector [46, 48, 53], thus prompting further effort in this direction.

---

[3]The exact numbers depend on the value of $\delta_{\rm CP}$ and specifics of the detector configurations.

- **Origin of neutrino mass:** The mechanism of neutrino mass generation is at present considered unknown, as the standard Higgs mechanism described in Section 1.3.2 would yield masses for Dirac neutrinos only if the corresponding Yukawa couplings y_ν were extremely small ($\lesssim 10^{-7} y_e$ where $y_e \simeq 2.9 \cdot 10^{-6}$ is the Yukawa coupling of the next particle on the mass scale – the electron). One class of the actively researched neutrino mass generation mechanisms is the seesaw-type models (see Section 2.2.2), where the very heavy right-chiral Majorana neutrinos give rise to small masses of the mostly-active neutrinos. Thus, the question of the neutrino mass origin is closely connected to that of unitarity of the neutrino mixing matrix.

2.2 NEUTRINO MASSES

2.2.1 *Overview of the current experimental bounds*

In Table 2.1, we summarize the bounds on neutrino masses as derived from neutrino oscillation data, direct measurements of the neutrino mass from tritium (^{3}H) decay, neutrinoless double beta decay experiments, and cosmological observations.

Experiment type	Experiment name	Observable	Bound (confidence)
Neutrino oscillations	Multiple (global fit) [65, 66]	$\Delta m_{ij}^2 \implies \sum_i m_i$	NO: $> 0.058\,\mathrm{eV}$ (99.7%)
			IO: $> 0.097\,\mathrm{eV}$ (99.7%)
Direct (^{3}H decay)	KATRIN [110]	$m_{\nu_e} \equiv \sqrt{\sum_i \lvert U_{ei} \rvert^2 m_i^2}$	$< 0.8\,\mathrm{eV}$ (90%)
		$\implies \sum_i m_i$	$< 2.4\,\mathrm{eV}$ (90%)
$0\nu\beta\beta$	GERDA [111]	$\lvert m_{\beta\beta} \rvert \equiv \lvert \sum_i U_{ei}^2 m_i \rvert$	$< 0.079\text{–}0.180\,\mathrm{eV}$ (90%)
	KamLAND-Zen [112]		$< 0.061\text{–}0.165\,\mathrm{eV}$ (90%)
Cosmology	Planck [113]	$\sum_i m_i$	$< 0.12\,\mathrm{eV}$(95%)
	SDSS [114]		$< 0.115\,\mathrm{eV}$ (95%)
	DESI + Planck + ACT [115]		$< 0.072\,\mathrm{eV}$ (95%)[4]

TABLE 2.1 – Neutrino mass bounds as derived from different types of experiments. Only the two most stringent constraints for neutrinoless double beta decay ($0\nu\beta\beta$) experiments are listed here. For a more complete selection of bounds (as of 2021), see [116].

First, we note that neutrino oscillation experiments only measure the mass splittings Δm_{ij}^2 and are not in principle sensitive to the absolute mass scale. However, assuming that the lightest neutrino has the lowest possible mass (0), we can derive a lower bound on the sum of neutrino masses, $\sum_i m_i$ (also referred to as $\sum m_\nu$ in literature). As we show in the first row of Table 2.1, this estimate depends significantly on the assumed neutrino mass ordering, with the lower bound being nearly twice as high for IO compared to NO. These estimates can be directly contrasted with the cosmological constraint on $\sum_i m_i$, which are derived from cosmic microwave background (CMB) and matter power spectrum measurements[5]. The typical bound derived from such observations is $\sum_i m_i \lesssim 0.1\,\mathrm{eV}$, with the most recent measurement by DESI [115] being in tension with the inverted mass ordering scenario.

---

[4]Note that this measurement is prior-dependent [115]. The reported value corresponds to the prior $\sum_i m_i > 0$, although a tighter prior ($\sum_i m_i \gtrsim 0.06\,\mathrm{eV}$) could be defined based on neutrino oscillation constraints.

[5]These measurements are possible thanks to the impact of massive neutrinos on structure formation in the Universe, in particular the hindered growth of small-scale structures (large wavenumber k in the matter power spectrum $P(k)$) due to neutrino free-streaming [15, 117].

Another important channel through which the absolute neutrino mass scale can be determined is the so-called direct neutrino mass measurement. Specifically, the KATRIN experiment measures the endpoint of the electron kinetic energy spectrum in the decay of tritium,

$$^3\text{H} \rightarrow {}^3\text{He} + e^- + \bar{\nu}_e. \tag{2.4}$$

The lower the effective mass m_{ν_e} of the electron antineutrino in Equation (2.4), the closer is the endpoint of the e^- energy spectrum to its maximum allowed value of 18.57 keV [118]. The sum of neutrino masses can also be inferred from m_{ν_e}, with the resulting upper bound being an order of magnitude looser than that derived from cosmological measurements.

Finally, if neutrino is a Majorana fermion, the neutrinoless double beta decay rate $\Gamma_{0\nu\beta\beta}$ is a function of the effective Majorana mass $m_{\beta\beta}$ (specifically, $\Gamma_{0\nu\beta\beta} \propto |m_{\beta\beta}|^2$) [9]. In this case, an observation of the $0\nu\beta\beta$ process would yield a measurement of $|m_{\beta\beta}|$, while a non-observation results in an upper bound as shown in Table 2.1.

Figure 2.3 summarizes the relationships between the mass of the lightest neutrino state and the three neutrino mass metrics discussed above, and shows the uncertainties permitted by current (2022) neutrino oscillation data[6].
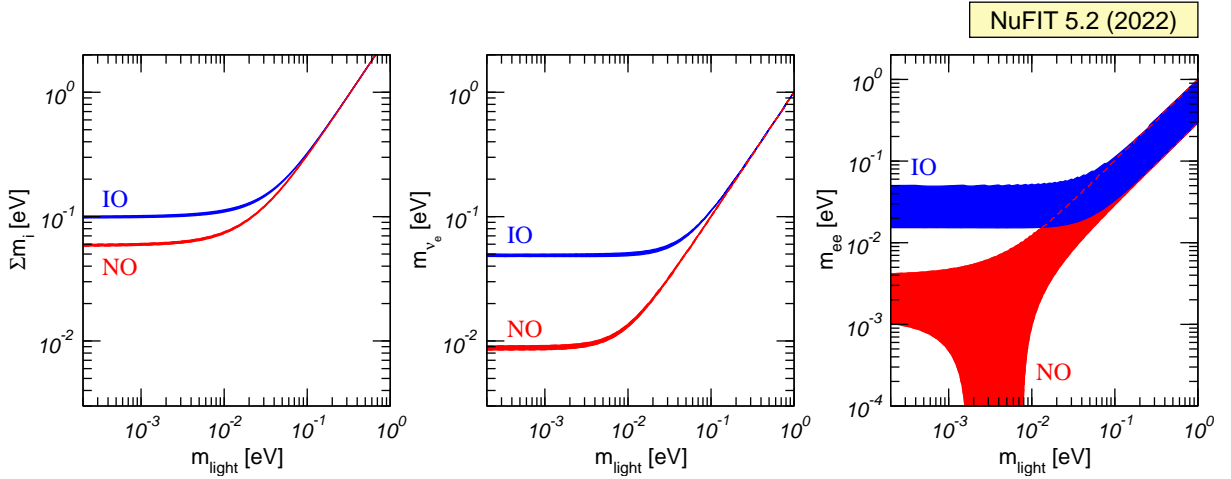
FIGURE 2.3 – Relations between the mass of the lightest neutrino state (m_{light}) and three neutrino mass metrics, as derived in the NuFit 5.2 [65, 66] global fit. The mass metrics include the sum of neutrino masses $\sum_i m_i$; effective electron neutrino mass m_{ν_e}; and effective Majorana neutrino mass $m_{ee} \equiv m_{\beta\beta}$ in Table 2.1. The bands correspond to the 2σ (95%) confidence limit, reflecting the uncertainty in the mass splittings and the mixing angles derived from neutrino oscillation data.

2.2.2 Neutrino mass models

The structure of the neutrino mass terms that could be added to the Standard Model Lagrangian depends on whether neutrino is of Dirac or Majorana type. In either case, it requires addition of new degrees of freedom to the Standard Model, as we explain below.

First, let us consider the possibility that neutrino is a pure Dirac fermion, i.e., $\nu = \nu_L + \nu_R$, where ν_R is a right-chiral field added to the SM. Then, the Dirac mass term can be written as [9]

$$\mathcal{L}^D = -m_D \bar{\nu}\nu = -m(\bar{\nu}_R \nu_L + \bar{\nu}_L \nu_R), \tag{2.5}$$

---

[6]Note that the cosmological limits are not shown in Figure 2.3.

combining left-chiral and right-chiral terms as in Section 1.3.2. Then, a coupling between ν and the Higgs field can be introduced by analogy to Equation (1.42), and neutrinos would acquire masses after electroweak symmetry breaking (EWSB) just as the charged leptons do. However, as mentioned in Section 2.1.2, neutrino-Higgs Yukawa couplings y_ν would have to be seven orders smaller than those of the next most massive particles, indicating a huge mass gap in the SM particle landscape. While this scenario is not impossible, it is typically considered unnatural and discarded for aesthetic reasons [35].

What happens if the neutrino was instead a Majorana fermion? In this case, we could define a Majorana neutrino field

$$\nu = \nu_L + \nu_L^C = \nu_L + C\bar{\nu}_L^\top, \tag{2.6}$$

which is equal to its antiparticle conjugate, $\nu = \nu^C$, as prescribed by the Majorana nature. The corresponding Majorana mass term for the left-chiral field ν_L would be

$$\mathscr{L}_L^M = -\frac{1}{2}m\bar{\nu}\nu = \frac{1}{2}m_L\nu_L^\top C^\dagger \nu_L + \text{h.c.}^7. \tag{2.7}$$

At the same time, there is no lepton-Higgs coupling term that could be written within the realms of the Standard Model such that is results in Equation (2.7) after the symmetry breaking. Indeed, all terms in the SM Lagrangian have the dimension of E^4, which can be understood from the requirement of the action $S = \int \mathscr{L}[\phi_i(x)]\mathrm{d}^4(x)$ being dimensionless. However, the lowest-dimensional addition to the SM Lagrangian that could give rise to Equation (2.7) after EWSB is the **dimension-5 operator** $\mathscr{L}_5$, which has the structure[8]

$$\mathscr{L}_5 = g\frac{LHLH}{\Lambda}, \tag{2.8}$$

where L (H) is the lepton (Higgs) doublet (equivalent to Ψ (Φ) in the earlier notation), g is a coupling coefficient, and Λ is a constant with the dimension of mass [9]. The $\mathscr{L}_5$ in Equation (2.8) is also referred to as the **Weinberg operator** [119], which, after H acquires the vacuum expectation value, results precisely into the Majorana mass term for ν_L:

$$m_L = \frac{gv^2}{\Lambda}. \tag{2.9}$$

Requiring $m_L \sim 0.1\,\text{eV}$, we obtain $\Lambda \sim g \times 10^{14}\,\text{GeV}$, which defines the scale of new physics responsible for neutrino mass generation. The effective formulation of this new physics as per Equation (2.8) does not, however, explicitly specify the new degrees of freedom or the mechanism through which they generate small m_L. All that Equation (2.8) gives is parametrically suppressed neutrino masses, at the expense of a non-renormalizable theory with a higher-dimensional operator. This theory is valid only up until $\sim\Lambda$ and must be "UV completed" by explicitly specifying the new degrees of freedom at higher energies [120].

One way to do this is to introduce heavy right-chiral neutrinos, whose Majorana mass term is

$$\mathscr{L}_R^M = \frac{1}{2}m_R\nu_R^\top C^\dagger \nu_R + \text{h.c.}, \tag{2.10}$$

---

[7] The shorthand notation "h.c." stands for "Hermitian conjugate."

[8] The expression below provides a rather schematic structure of this operator; see [9] for the full expansion of $LHLH$. The wording "dimension 5" comes from the fact that the $\dim L = E^{3/2}$ and $\dim H = E^1$.

by analogy with Equation (2.7). Now, if ν_R exists, a Dirac mass term (Equation (2.5)) can also be written, resulting in a combined Dirac-Majorana mass term:

$$\mathscr{L}_{\mathrm{mass}} = -m_D \bar{\nu}_R \nu_L + \frac{1}{2} m_R \nu_R^\top C^\dagger \nu_R + \frac{1}{2} m_L \nu_L^\top C^\dagger \nu_L + \mathrm{h.c.} \tag{2.11a}$$

$$= \frac{1}{2} N_L^\top C^\dagger M N_L + \mathrm{h.c..} \tag{2.11b}$$

In Equation (2.11), we defined $N_L = (\nu_L, \nu_R^C)^\top$ and the mass matrix M,

$$M = \begin{pmatrix} m_L & m_D \\ m_D & m_R \end{pmatrix}. \tag{2.12}$$

We see that neither ν_L nor ν_R have definite masses; to find the latter, M needs to be diagonalized. If we require the Majorana mass m_R and the Dirac mass m_D to be real and positive, the remaining Majorana mass m_L is in general complex [9]. We will consider the simplest case of $m_L = 0$, in which case M has the following eigenvalues:

$$m_1 \simeq \frac{m_D^2}{m_R}; \tag{2.13a}$$

$$m_2 \simeq m_R. \tag{2.13b}$$

Equation (2.13) is a classic illustration of the so-called **seesaw mechanism** (specifically, the "type-I" seesaw [35]), where the heavier is one of the mass states (m_2), the lighter is the other (m_1). The Dirac mass scale m_D in this case does not have to be tiny, as would have to be the case if neutrino masses were generated by the Higgs mechanism alone. In fact, m_D could be comparable to the mass scale of the charged fermions, thus solving the aesthetic problem of unexplainably small Yukawa couplings.

While the above arguments were applicable to the simple one-generation case, they straightforwardly generalize to three generations of active neutrinos and an arbitrary number n_s of right-chiral ("sterile") neutrino species. In this case, the combined Dirac-Majorana mass matrix becomes

$$M = \begin{pmatrix} M_L & M_D^\top \\ M_D & M_R \end{pmatrix}, \tag{2.14}$$

where the dimensions of matrices M_L, M_R, and M_D are 3×3, $n_s \times n_s$, and $n_s \times 3$, respectively. The diagonalization of this matrix yields a total of $n = 3 + n_s$ neutrino mass states. Furthermore, the complete neutrino mixing matrix in this case becomes

$$\mathcal{U} = \begin{pmatrix} U_{ll} & U_{lh} \\ U_{hl} & U_{hh} \end{pmatrix}, \tag{2.15}$$

where the index l (h) stands for light (heavy). The U_{ll} matrix is the usual 3×3 neutrino mixing matrix, which parameterizes the three active neutrino flavors in terms of the three light mass eigenstates ν_i, $i = 1..3$. We see that U_{ll} is no longer unitary, as it becomes a submatrix of a larger unitary matrix $\mathcal{U}$; for this reason, we refer to this matrix as N in Section 2.4. The first three rows of $\mathcal{U}$, i.e., a combination of U_{ll} and U_{lh}, give a $3 \times n$ mixing matrix, which allows one to express the active flavors in terms of *all* of the mass eigenstates, i.e., ν_i with $4 \leq i \leq n$. Whether such ν_i can practically participate in the mixing depends on their masses m_i and the energy reach of a given process/experiment.

We note that the aforementioned type-I seesaw mechanism is not only just one of the many neutrino mass models, but also one of the many neutrino mass models that

result in non-unitarity of the active neutrino mixing matrix. For example, instead of a SM singlet introduced in type-I seesaw, the type-III seesaw mechanism features the addition of a fermionic triplet (with two new charged and one new neutral fermion) to the Standard Model [35]. In this case, the 3×3 matrix $N \equiv U_{ll}$ again becomes non-unitary [121]. Finally, non-unitarity of the neutrino mixing matrix can also result from the so-called radiative mass models [36, 122, 123], where neutrino masses are generated in one or more loops.

2.3 THREE-FLAVOR NEUTRINO OSCILLATIONS

2.3.1 *Unitarity of the three-neutrino mixing*

In the standard three-flavor neutrino oscillation framework, flavor states $|\nu_\alpha\rangle$ produced together with the respective charged leptons are related to the propagating mass states $|\nu_i\rangle$ via a unitary matrix U, as per Equation (2.1). The latter is also known as the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix. It usually parameterized as a product of three rotation matrices [34]:

$$
U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{\mathrm{CP}}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{\mathrm{CP}}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{2.16}
$$

where δ_{CP} is the CP-violating phase, $c_{ij} \equiv \cos\theta_{ij}$, $s_{ij} \equiv \sin\theta_{ij}$, and θ_{ij} is the mixing angle between mass eigenstates i and j. The *unitarity* condition postulated for the mixing matrix U is

$$
UU^\dagger = U^\dagger U = \mathbb{I}, \tag{2.17}
$$

where $\mathbb{I}$ is the identity matrix. This translates to nine real constraints on the individual elements of U. First, all three of its rows are required to normalize to 1 (in the ℓ^2 sense):

$$
N_\alpha \equiv \sum_{i=1}^{3} |U_{\alpha i}|^2 = 1. \tag{2.18}
$$

Furthermore, the *closure* conditions must be satisfied by the row unitarity triangles:

$$
t_{\alpha\beta} \equiv \sum_{i=1}^{3} U_{\alpha i}^* U_{\beta i} = 0, \tag{2.19}
$$

contributing six additional real constraints. Alternatively, instead of Equation (2.18) and Equation (2.19), one may require that all three matrix columns normalize to 1 and that all of the column unitarity triangles close [53], i.e.,

$$
N_k \equiv \sum_{\alpha \in \{e,\mu,\tau\}} |U_{\alpha k}|^2 = 1 \tag{2.20}
$$

and

$$
t_{kl} \equiv \sum_{\alpha \in \{e,\mu,\tau\}} U_{\alpha k}^* U_{\alpha l} = 0. \tag{2.21}
$$

Physically, the unitarity conditions imply that the total probability for one of the $\{e, \mu, \tau\}$ flavor states to oscillate to either one of these three flavors is conserved and equal to 1.

2.3.2 *Three-flavor oscillation probabilities*

Once produced as flavor states in charged current weak decays, neutrinos propagate as a superposition of mass states. Let us consider the most general case of neutrino propagation in matter, with the understanding that neutrino propagation in vacuum can be derived by imposing zero matter density. The evolution of the three mass states $|\nu_i\rangle$ is governed by the following Hamiltonian (expressed in the mass basis[9]):

$$H = H_{\text{vacuum}} + A = \frac{1}{2E}\begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta m_{21}^2 & 0 \\ 0 & 0 & \Delta m_{31}^2 \end{pmatrix} + U^\dagger \begin{pmatrix} V_{\text{CC}} + V_{\text{NC}} & 0 & 0 \\ 0 & V_{\text{NC}} & 0 \\ 0 & 0 & V_{\text{NC}} \end{pmatrix} U. \quad (2.22)$$

The first term in Equation (2.22) corresponds to the free propagation in vacuum, where $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$ are the mass splittings between the physical states i and j. The second term is the contribution of the matter potential A due to the charged current (CC) and the neutral current (NC) interactions of neutrinos with electrons and nucleons in matter. As electrons are the only leptons that compose the ordinary stable matter, the CC interactions via $W^\pm$ exchange are accessible only to electron neutrinos and antineutrinos. The V_{CC} component of the matter potential thus only depends on electron number density n_e:

$$V_{\text{CC}} = \pm\sqrt{2}G_F n_e = \pm\sqrt{2}G_F E_\nu \rho_m Y_e N_A, \quad (2.23)$$

where G_F is the Fermi constant, ρ_m is the mass density of the medium in g cm^{-3}, Y_e is the electron fraction per unit molar mass (also known as the "Z/A" factor and expressed in mol g^{-1}), and N_A is the Avogadro number. In Equation (2.23), the "+" sign corresponds to the matter potential experienced by neutrinos (ν_α), and the "-" sign – that seen by antineutrinos ($\bar{\nu}_\alpha$). The neutral current (Z^0-exchange) interactions of neutrinos could occur with protons, neutrons, and electrons alike. However, a typical assumption is that the medium that neutrinos travel through is neutral and unpolarized, in which case the contributions due to proton and electron NC potentials cancel each other out [124–126]. The only remaining component is the flavor-independent neutrino-neutron scattering, with the respective NC potential:

$$V_{\text{NC}} = \mp\frac{\sqrt{2}}{2}G_F n_n = \mp\frac{\sqrt{2}}{2}G_F \rho_m Y_n N_A, \quad (2.24)$$

where $Y_n = 1 - Y_e$ is the neutron fraction per unit molar mass, and the rest of the notations are the same as in Equation (2.23). In the unitary case, the neural current potential is usually omitted in Equation (2.22), as it is identical for all active neutrino flavors and contributes only an unobservable phase to the neutrino oscillation amplitude [126]. This is not the case when the mixing matrix is non-unitary (as explored later in Section 2.4), which necessitates the explicit inclusion of V_{NC} in Equation (2.22).

With the Hamiltonian defined as in Equation (2.22), the Schrödinger equation for the propagating neutrino states is:

$$i\hbar \frac{\partial}{\partial t} \boldsymbol{\nu}_{\text{m}}(t) = H\boldsymbol{\nu}_m(t), \quad (2.25)$$

where $\boldsymbol{\nu}_{\text{m}}(t)$ describes the time dependence of the state vector $\boldsymbol{\nu}_m \equiv (\nu_1, \nu_2, \nu_3)^\top$. This dependence can be expressed through the evolution operator S^0, such that

$$\boldsymbol{\nu}_{\text{m}}(t) = S^0(t)\boldsymbol{\nu}_{\text{m}}(0). \quad (2.26)$$

---

[9]We could also express this Hamiltonian in the flavor basis, where the matter potential contribution would be diagonal in the unitary three-flavor case. However, computations are much simplified in the non-unitary case if H is expressed in the mass basis, which is why we make the same choice in this section.

Plugging Equation (2.26) into Equation (2.25), we find

$$i\hbar\frac{\partial S^0}{\partial t} = HS^0,\tag{2.27}$$

where $\nu_m(0)$ is eliminated since Equation (2.25) holds for any choice of initial condition. When the densities of electrons and neutrons are constant throughout the distance L propagated by the neutrinos in time t, we can easily solve for the time dependence of the evolution operator [7, 127, 128]:

$$S^0(t) \to S^0(L) = \exp(-iHL).\tag{2.28}$$

The probability of oscillations from flavor α to flavor β can then be found as follows:

$$P_{\alpha\beta}(E,L) = |(US^0(E,L)U^\dagger)_{\beta\alpha}|^2 = |(Ue^{-iHL}U^\dagger)_{\beta\alpha}|^2.\tag{2.29}$$

In Figure 2.4, we use Equation (2.29) to compute the unitary three-flavor oscillation probabilities in vacuum.
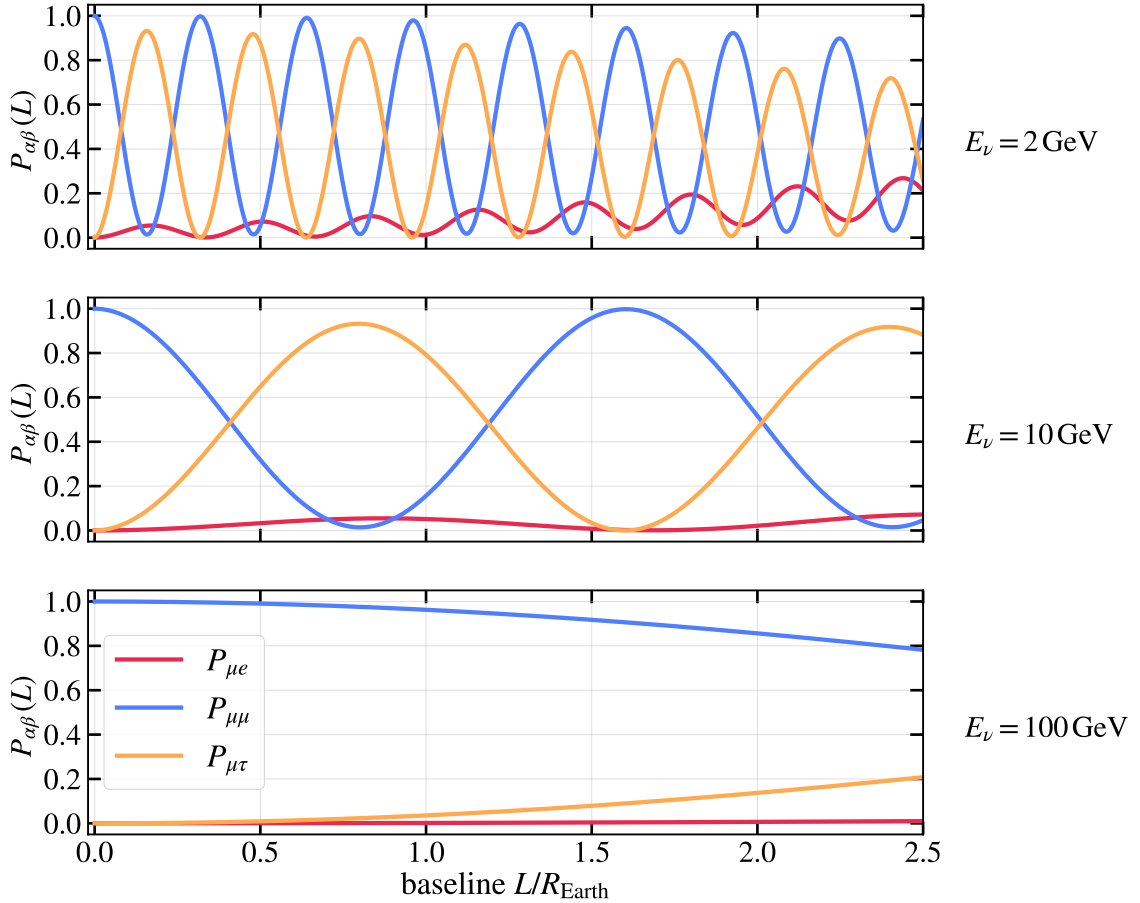
FIGURE 2.4 – The probability of ν_μ oscillation into ν_e (red), ν_τ (orange), as well as its survival as ν_μ (blue), shown at three characteristic atmospheric neutrino energies. The oscillation probabilities were computed numerically using Equation (2.29) with the NuFit 5.2 (2022) oscillation parameters for the normal neutrino mass ordering [65, 66]. *A similar version of this plot has been previously produced by the author for [55].*

2.4 NEUTRINO MIXING NON-UNITARITY

As discussed in Section 2.2.2, the existence of new physics motivated by non-zero neutrino masses (such as additional heavy neutrino states) may render the 3×3 neutrino mixing matrix non-unitary. In the most general form, such a matrix may be

parameterized as

$$N = \begin{pmatrix} |N_{e1}| & |N_{e2}|e^{i\phi_{e2}} & |N_{e3}|e^{i\phi_{e3}} \\ |N_{\mu1}| & |N_{\mu2}| & |N_{\mu3}| \\ |N_{\tau1}| & |N_{\tau2}|e^{i\phi_{\tau2}} & |N_{\tau3}|e^{i\phi_{\tau3}} \end{pmatrix}, \qquad (2.30)$$

where both $NN^\dagger$ and $N^\dagger N$ may deviate from identity. The parameterization (2.30) includes 9 real non-negative matrix element magnitudes $|N_{\alpha i}|$ and 4 complex phases $\phi_{\alpha i}$, which can be assigned to any 2×2 submatrix of N [53].

In Chapters 7 and 8 of this thesis, we focus on the regime of "minimal unitarity violation" (MUV) [53, 129], where non-unitarity is induced by physics that exists at energy scales much larger than the electroweak scale $v \simeq 246\,\mathrm{GeV}$. In this regime, only the usual low-mass neutrino mass eigenstates $|v_i\rangle$ with masses m_i ($i \in 1, 2, 3$) are kinematically accessible in an experiment. The effective neutrino flavor state $|v_\alpha^{\mathrm{eff.}}\rangle$ at the time of production or detection can only be a superposition of these accessible mass states, i.e.,

$$|v_\alpha^{\mathrm{eff.}}\rangle = \frac{1}{\sqrt{(NN^\dagger)_{\alpha\alpha}}} \sum_{i=1}^{3} N_{\alpha i}^* |v_i\rangle \equiv \frac{1}{\sqrt{N_\alpha}} \sum_{i=1}^{3} N_{\alpha i}^* |v_i\rangle, \qquad (2.31)$$

where the sum must be truncated at $i = 3$. The normalization factor $(\sqrt{(NN^\dagger)_{\alpha\alpha}})^{-1}$ ensures that $\langle v_\alpha^{\mathrm{eff.}} | v_\alpha^{\mathrm{eff.}}\rangle = 1$ [53, 129, 130]. However, under a non-unitary mixing matrix N, the set of effective flavor states is not orthonormal, i.e., $\langle v_\beta^{\mathrm{eff.}} | v_\alpha^{\mathrm{eff.}}\rangle \neq \delta_{\alpha\beta}$. An orthonormal basis could be defined in a complete high-energy theory covering the new physics energy scale [129].

Once U is replaced with the more general N in Equation (2.22), the non-unitary oscillation probabilities are found analogously to Equation (2.29):

$$P_{\alpha\beta}(E, L) = \frac{|(NS^0(E,L)N^\dagger)_{\beta\alpha}|^2}{(NN^\dagger)_{\alpha\alpha}(NN^\dagger)_{\beta\beta}} = \frac{|(Ne^{-iHL}N^\dagger)_{\beta\alpha}|^2}{(NN^\dagger)_{\alpha\alpha}(NN^\dagger)_{\beta\beta}} \equiv \frac{|(Ne^{-iHL}N^\dagger)_{\beta\alpha}|^2}{N_\alpha N_\beta}, \qquad (2.32)$$

where the N_α, N_β factors in the denominator originate from the normalization of the effective flavor states as per Equation (2.31).

The non-unitarity of the N matrix can be quantified as any deviation of $(NN^\dagger)_{ij}$ or $(N^\dagger N)_{ij}$ from δ_{ij}. Specific cases include off-nominal normalizations of any of the matrix rows ($N_\alpha = \sum_i |N_{\alpha i}|^2 \neq 1$) or columns ($N_i = \sum_\alpha |N_{\alpha i}|^2 \neq 1$), as well as non-zero closures of rows ($t_{\alpha\beta} = \sum_i N_{\alpha i}^* N_{\beta i} \neq 0$) or columns ($t_{kl} = \sum_\alpha N_{\alpha k}^* N_{\alpha l} \neq 0$). Such deviations directly affect not only neutrino oscillations as per Equation (2.32), but also production and detection of neutrinos [52, 53, 128–130]. This requires treating carefully the flux and the cross section inputs to the projected number of events in an experiment. In particular, if the unoscillated flux Φ_α of the initial neutrino flavor α or the interaction cross section σ_β of the final flavor β are based on Standard Model ("SM") calculations assuming unitarity, then they need to be corrected by appropriate combinations of N_α or N_β in the non-unitary ("NU") case. In particular,

$$\Phi_\alpha^{\mathrm{NU}} = N_\alpha \Phi_\alpha^{\mathrm{SM}}; \qquad (2.33a)$$

$$\sigma_\beta^{\mathrm{NU,CC}} = N_\beta \sigma_\beta^{\mathrm{SM}}; \qquad (2.33b)$$

$$\sigma_\beta^{\mathrm{NU,NC}} = N_\beta^2 \sigma_\beta^{\mathrm{SM}}, \qquad (2.33c)$$

where "CC" stands for charged current, and "NC" – for neutral current interactions. We refer the reader to [52, 53, 128–130] for details and derivations of the prefactors in Equation (2.33). We note that these prefactors can be absorbed in the definition of the oscillation probability itself, such that the "effective" oscillation probability $\hat{P}_{\alpha\beta}$ is the product of $P_{\alpha\beta}$ from Equation (2.32) and the appropriate factors from Equation (2.33).

To demonstrate the impact of non-unitarity on the oscillation probabilities compared to the unitary (Standard Model[10]) expectations, let us focus on the case of the off-nominal row normalizations and the following three oscillation channels: electron antineutrino disappearance (relevant for reactor neutrino experiments such as Daya Bay, KamLAND, and JUNO), muon neutrino disappearance, and tau neutrino appearance (relevant for atmospheric neutrino experiments such as IceCube-DeepCore, IceCube-Upgrade, and KM3NeT/ORCA). The respective oscillation probabilities are denoted as $P_{\bar{e}\bar{e}}$, $P_{\mu\mu}$, and $P_{\mu\tau}$. For the purpose of this example, we rescale the row elements of the unitary matrix U (Equation (2.16)) such that $N_e = 0.95$ when probing $P_{\bar{e}\bar{e}}$, $N_\mu = 0.9$ when probing $P_{\mu\mu}$, and $N_\tau = 0.9$ when probing $P_{\mu\tau}$. All other row normalizations, except for the specific one modified in each case, are fixed at 1. We further assume that neutrinos are propagating in a medium with density $\rho = 2.7\,\mathrm{g\,cm}^{-3}$. Our results are shown in Figure 2.5, both with and without rescaling the probabilities by the factors of N_α, N_β as discussed above.

From the top left panel of Figure 2.5, we see that the "raw" $P_{\bar{e}\bar{e}}$ probabilities computed from Equation (2.32) are almost completely unaffected by the off-nominal N_e in the case of a short-baseline experiment such as Daya Bay. This is due to the cancellation of the factors of N_e in the numerator and denominator of Equation (2.32) and is in agreement with the result obtained by [53]. However, the differences with the SM case are clearly visible when rescaling the effective oscillation probability by N_e or N_e^2, which is necessary if either the cross section or both the flux and the cross section need to be corrected for non-unitarity. In the remaining three panels, which are representative of the scenarios probed by KamLAND and IceCube-DeepCore, the non-unitary oscillation probability visibly differs from the SM one even if no extra corrections for the cross section (N_β) or the flux (N_α) are introduced to $P_{\alpha\beta}$. This occurs due to the matter effects, which add an extra N-dependent term in the Hamiltonian (see Figure B.1 for comparison with the vacuum case).

In the global fit for neutrino mixing matrix non-unitarity, which is described in Chapters 7 and 8, we probe the entire structure of the non-unitary mixing matrix N (as defined in Equation (2.30)) using combined atmospheric and reactor neutrino data. This implies incorporating both the non-unitary oscillation probabilities from Equation (2.32) and the experiment-specific flux and cross section corrections due to non-unitarity (see Section 7.2). On the other hand, the tau neutrino appearance analysis introduced in Section 9.1 only probes the overall scale N_{ν_τ} of tau neutrinos relative to the unitary expectation. We stress that N_{ν_τ} measured in such an analysis is not equal to the row normalization N_τ of the matrix in Equation (2.30), as the scaling factor N_{ν_τ} is applied to tau neutrinos after they have been produced through the $\nu_\mu \to \nu_\tau$ channel with the unitary oscillation probabilities as in Equation (2.29). This means that the tau neutrino appearance analysis does not take into account the impact of non-unitarity on neutrino matter potential, as well as other potential manifestations of non-unitarity (e.g., off-unitary normalizations of other rows/columns or non-zero closures of the unitarity triangles).

---

[10]Even though neutrino oscillations cannot be explained in the Standard Model without neutrino masses, we use terms "unitary" and "Standard Model" interchangeably to denote the setup with only 3 neutrino mass states and 3 neutrino flavor states.
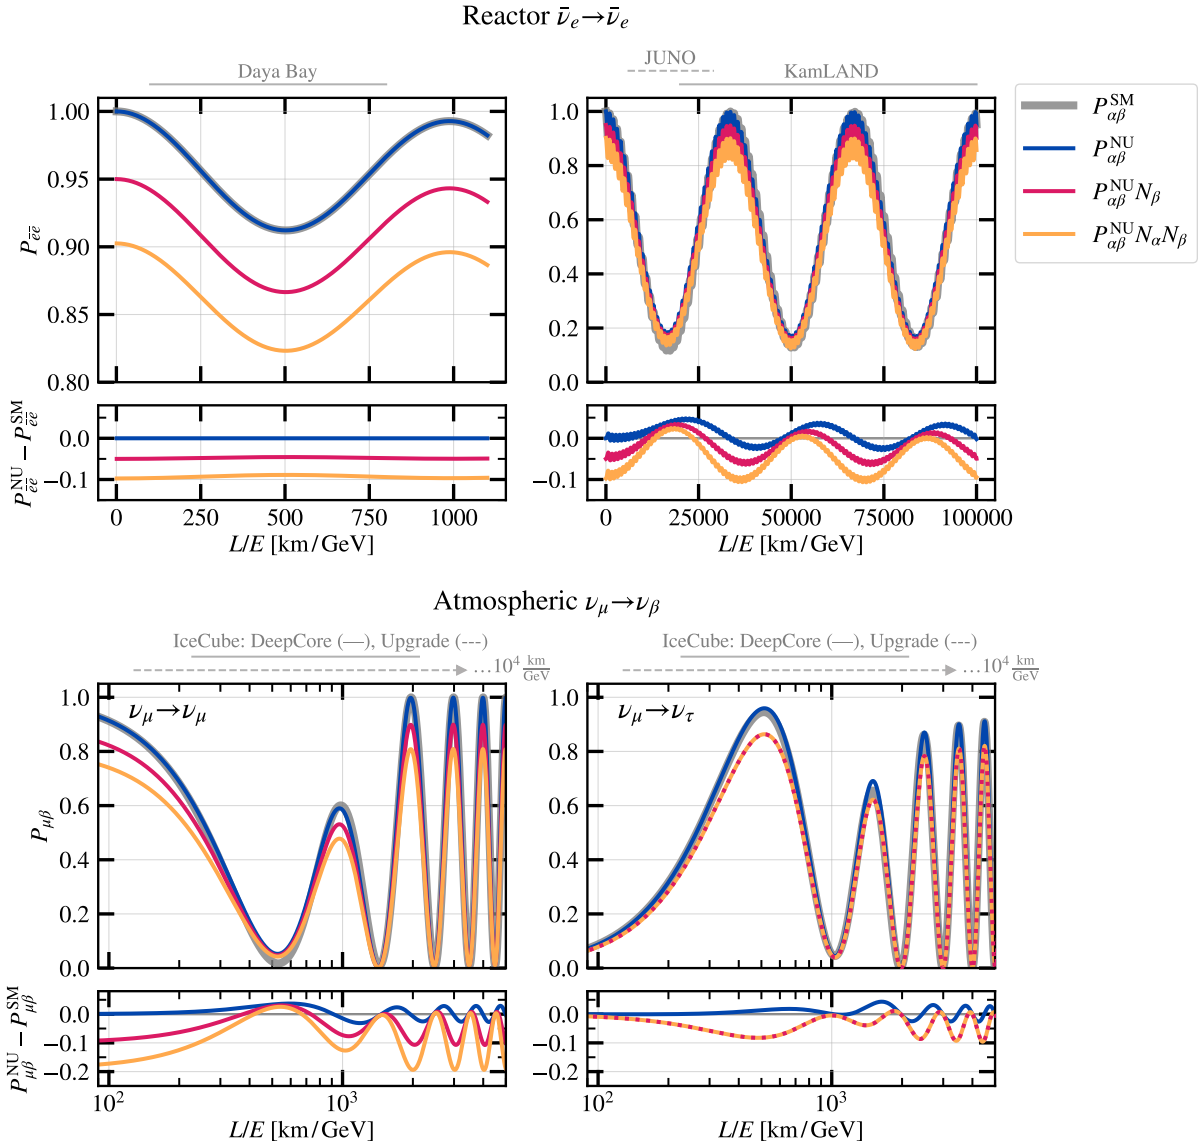
FIGURE 2.5 – Comparison of the unitary (Standard Model, "SM") and the non-unitary ("NU") oscillation probabilities for the case of $\bar{\nu}_e$ disappearance (top), ν_μ disappearance (bottom left), and ν_τ appearance (bottom right), including matter effects. $P_{\alpha\beta}^{\mathrm{NU}}$ corresponds to Equation (2.32), while other non-unitary probabilities are scaled by the normalization factors of N_α, N_β for illustration purposes (see text). The approximate ranges of the baseline-to-neutrino-energy ratio (L/E) probed by the reactor and atmospheric neutrino oscillation experiments considered in Chapters 7 and 8 are shown in gray above each panel. For the atmospheric oscillation probabilities, a fixed baseline of $12\,742\,\mathrm{km}$ is assumed. Note that the $P_{\alpha\beta}^{\mathrm{NU}} N_\beta$ and $P_{\alpha\beta}^{\mathrm{NU}} N_\alpha N_\beta$ lines are overlapping in the bottom right panel, since $N_\alpha = N_\mu = 1$ in this case.

II

Atmospheric neutrinos

Production and oscillations

Atmospheric neutrinos are produced in *atmospheric air showers* – cascades of particles in the Earth's atmosphere induced by high-energy charged particles called cosmic rays. The latter are composed of predominantly protons but also of heavier nuclei such as helium, oxygen, and iron [131]. When the cosmic ray hits one of the nuclei in the atmosphere, e.g., ^{14}N or ^{16}O, we speak of *hadronic interactions*, which, given a sufficiently energetic primary particle, will break apart the target nucleus and lead to production of several new hadrons. This is the beginning of the hadronic cascade component of an air shower. The mesons produced in such a cascade, such as $\pi^{\pm}$, $K^{\pm}$, and K^0_L, are unstable and, after traveling a certain path length and possibly re-interacting, will decay into atmospheric leptons – neutrinos and muons. This is illustrated in the left panel of Figure 3.1.
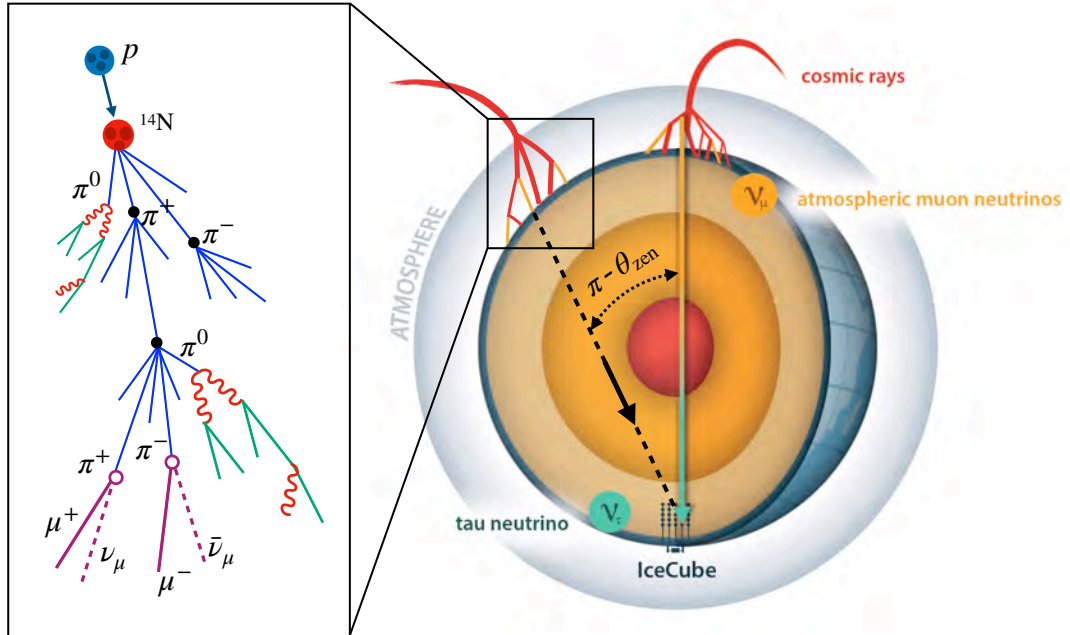

Figure 3.1 – Schematic illustration of atmospheric neutrino production in cosmic ray-induced air showers (left), followed by neutrino propagation through the Earth and flavor oscillations on the way to the IceCube detector at the South Pole. *The right portion of the figure is adapted from the IceCube press release* [132] *(credit: IceCube Collaboration).*

Neutrinos produced in such hadronic cascades are primarily of the muon flavor, ν_μ and $\bar\nu_\mu$, obtained directly from the $\pi^\pm$ decays (e.g., $\pi^+ \to \mu^+ \nu_\mu$) and from the subsequent $\mu^\pm$ decays (e.g., $\mu^+ \to e^+ \bar\nu_\mu \nu_e$). Since the majority of electron neutrinos are produced in the muon decays, the approximate $\nu_\mu : \nu_e$ flavor ratio for atmospheric neutrinos is $2:1$. The ν_τ flux obtained from the air showers themselves, at the energies relevant for neutrino oscillation studies ($\sim$1–100 GeV), is practically negligible. The only source of tau neutrinos of $\mathcal{O}(10\,\text{GeV})$ energies[1] is therefore $\nu_\mu \to \nu_\tau$ oscillations, as illustrated in the right panel of Figure 3.1.

The muon and electron neutrinos are produced continuously as a function of altitude in the Earth's atmosphere (with a characteristic production height being 15–20 km above sea level) and emitted at different angles. Therefore, those of them that reach a terrestrial detector, such as the IceCube Neutrino Observatory at the South Pole, arrive from different directions and travel different distances towards the detector. For example, the neutrinos that are produced in the atmosphere above the North Pole need to cross the entire Earth (12 742 km in diameter) to reach IceCube. This corresponds to the zenith angle $\theta_{\text{zen}} = \pi$ in Figure 3.1 – or, equivalently, $\cos\theta_{\text{zen}} = -1$. These Earth-crossing neutrinos are the most interesting from the perspective of neutrino oscillation measurements, as they provide an exceptionally long baseline for $\nu_\mu \to \nu_\tau$ oscillations, which lead to observable signatures of atmospheric muon neutrino disappearance and tau neutrino appearance.

In the remainder this chapter, we will cover the calculations of atmospheric neutrino fluxes prior to oscillations (Section 3.1), as well as the atmospheric neutrino oscillations themselves (Section 3.2), in more detail.

3.1 ATMOSPHERIC NEUTRINO FLUXES

Neutrino flux Φ is defined as the energy spectrum $\mathrm{d}N/\mathrm{d}E$ of neutrinos produced per unit area A, unit time t, and solid angle Ω as viewed from the detector [55]:

$$\Phi_\nu(E, \Omega) = \frac{\mathrm{d}N}{\mathrm{d}E\,\mathrm{d}A\,\mathrm{d}t\,\mathrm{d}\Omega}. \tag{3.1}$$

In spherical coordinates, the unit solid angle is $\mathrm{d}\Omega = \mathrm{d}\cos\theta_{\text{zen}}\mathrm{d}\varphi$, where θ_{zen} is the zenith angle (as defined by the orientation of the neutrino arrival direction relative to the detector axis; see Figure 3.1), and φ is the respective azimuthal angle.

One of the most natural approaches to atmospheric neutrino flux calculations lies through Monte Carlo simulations of atmospheric air showers, which start with an input cosmic ray flux spectrum and proceed with event-by-event treatment of cosmic ray-atmospheric nucleus collisions and subsequent secondary particle production. The interactions of primary and secondary particles, as well as the decays of unstable secondaries, are handled by event generators, which yield a distribution of interaction or decay products with a complete per-particle kinematic information according to the underlying model of the generator. The Honda-Kajita-Kasahara-Midorikawa-Sajjad Athar (shortly "HKKMS," or "Honda") atmospheric neutrino flux model is a widely used example of such a calculation, which relies on a combination of the JAM+DPMJET [133–136] hadronic interaction models to simulate the development of hadronic cascades. The HKKMS fluxes predicted for ν_e, $\bar\nu_e$, ν_μ, and $\bar\nu_\mu$ are shown in Figure 3.2. This is the nominal atmospheric neutrino flux model used in all analyses presented in this thesis.

---

[1] At energies above a few TeV, the so-called *prompt* component of the ν_τ flux becomes significant, which originates from decays of heavy charmed mesons [43, 44].

It takes into account the essential physics processes as well as the full three-dimensional geometry of the air shower development and the deflection of charged particles in the geomagnetic field.
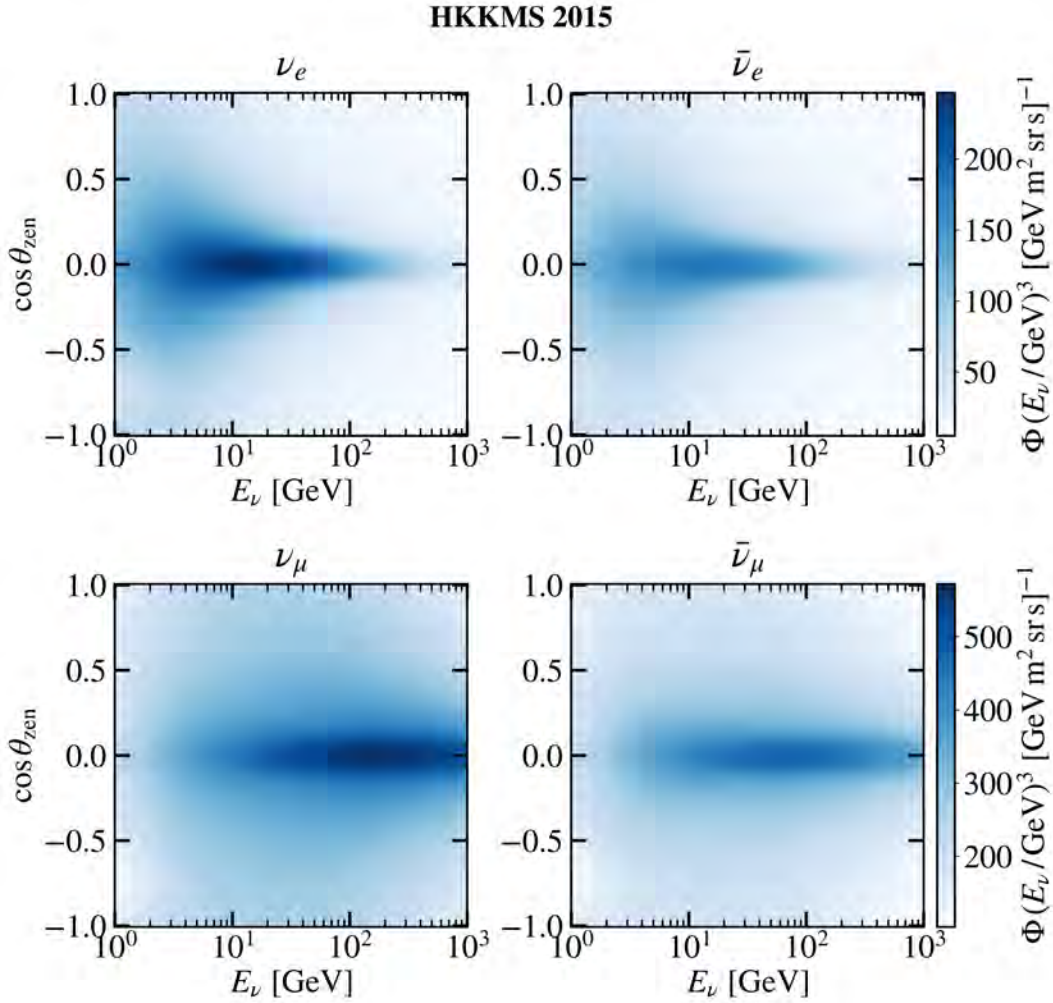

FIGURE 3.2 – Atmospheric neutrino flux as predicted by the HKKMS 2015 ("Honda") model [137] through the full-sky Monte Carlo simulations of air showers. The flux is given for the South Pole location at the solar minimum, averaged over the azimuthal angle, and scaled by E_ν^3.

An alternative way to calculate atmospheric neutrino fluxes is via *cascade equations* – the multi-species coupled transport equations describing particle production, interaction, and decay in the atmosphere [42]. A fast numerical framework for solving these equations was developed in [43, 44, 138] and is available as an open-source software called MCEQ. Through the MCEQ framework, the contributions of individual hadronic cascade species (e.g., $\pi^\pm$, $K^\pm$, $\mu^\pm$) to the atmospheric neutrino flux can be evaluated separately, as shown in Figure 3.3. This allows for flexible propagation of systematic uncertainties associated with meson production in hadronic cascades; in particular, the Bartol uncertainty scheme [140] (also referred to as the "Barr scheme") is implemented in IceCube analyses via MCEQ. This scheme assigns 1σ uncertainties on meson production in hadronic interactions across the (E_{proton}, $x_{\mathrm{lab}} \equiv \frac{E_{\mathrm{meson}}}{E_{\mathrm{proton}}}$) phasespace split into discrete blocks, as shown in Figure 3.4. The magnitude of the quoted uncertainties is based on availability of accelerator data in the respective regions of phasespace as of 2006 and is therefore considered conservative at the present time. A modern successor of the Barr scheme is the data-driven hadronic interaction model ("DDM") utilizing the fixed-target data from the NA49 and the NA61/SHINE experiments [141]. The
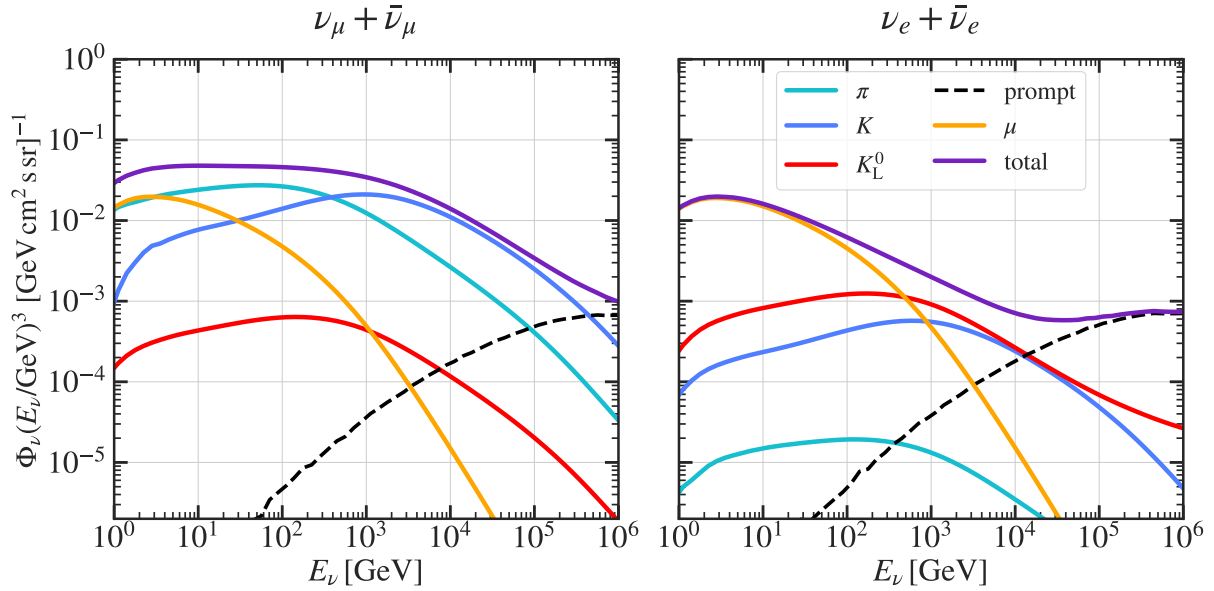
FIGURE 3.3 – Energy dependence of the $\cos\theta_{\text{zen}}$-averaged atmospheric neutrino fluxes, scaled by E_ν^3. The fluxes were computed via the MCEq software [43, 44], using the DPMJet-III 19.1 hadronic interaction model [138], Global Spline Fit cosmic ray flux [131], and the NRLMSISE-00 atmosphere for the South Pole in January [139]. The different curves represent the contributions of the different parent particles to the total neutrino flux (solid purple). *The figure was originally produced by the author for [55].*
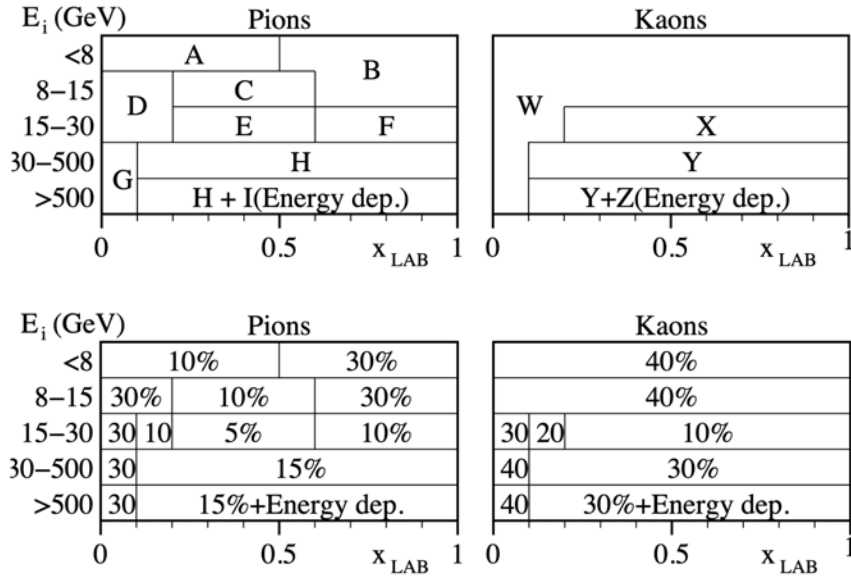

FIGURE 3.4 – Bartol ("Barr") uncertainties assigned in discrete (E_i, x_{lab}) blocks, denoted by letters A-I, where E_i is the incident proton energy and x_{lab} is the ratio of the secondary meson to the primary proton energy [140]. *Figure extracted from [140].*

combination of DDM with the data-driven cosmic ray flux model (Global Spline Fit, or GSF [131]) was calibrated using the muon spectrometer data in [142], yielding the state-of-the-art DAEMONFLUX model. The predictions of this model for atmospheric electron and muon neutrino fluxes are presented in Figure 3.5 along with their estimated data-driven uncertainties (which are based on those of GSF and DDM propagated through MCEq).

It is apparent that the DAEMONFLUX uncertainties are significantly lower than those obtained by propagation of the Barr scheme prescriptions through MCEq (~1–2%
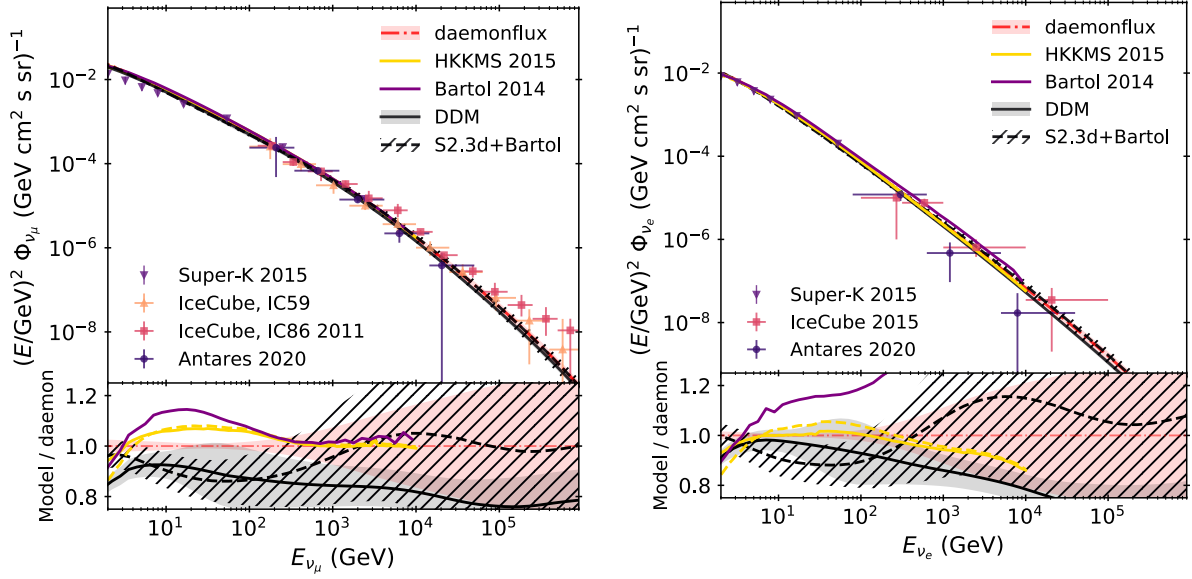
FIGURE 3.5 – Atmospheric neutrino flux as predicted by the DAEMONFLUX model [142] calibrated with muon spectrometer data. The DAEMONFLUX prediction is compared to a few other calculation schemes, including the HKKMS 2015 model [137] shown in Figure 3.2 and the MCEq [43, 44] calculations with the SIBYLL 2.3D [143] hadronic interaction model. *Figure extracted from* [142].

compared to ∼5–10%, respectively, in the 1–100 GeV energy range). This is a favorable and an expected outcome, as much stricter constraints on the hadronic interaction model and the cosmic ray flux parameters possible with the modern data naturally lead to higher precision of the ensuing atmospheric neutrino flux estimates. Many analyses in the IceCube Collaboration are therefore transitioning to DAEMONFLUX as the nominal atmospheric neutrino flux model; in particular, it has been used in the sterile neutrino search with the high-energy (>500 GeV) atmospheric neutrino data [144]. The analyses presented in this thesis, however, rely on the older prescription of the HKKMS flux model combined with the conservative Barr uncertainties, as the DAEMONFLUX model became available either after or during the development of the respective event selections.

3.2 ATMOSPHERIC NEUTRINO OSCILLATIONS

In Section 2.4, we have shown examples of neutrino oscillation probabilities in vacuum and in constant matter profiles as a function of the neutrino oscillation phase, L/E_ν (the ratio of the neutrino path length to its energy). Let us now consider a realistic scenario of atmospheric neutrino flavor conversions happening inside the Earth, whose approximate matter density profile is shown in Figure 3.6 as a function of distance from the Earth's center. This includes the Earth's inner and outer core with the average densities $\langle\rho\rangle$ of approximately $13\,\mathrm{g\,cm^{-3}}$ and $11\,\mathrm{g\,cm^{-3}}$; the mantle with $\langle\rho\rangle \simeq 5\,\mathrm{g\,cm^{-3}}$; and the Earth's crust with $\langle\rho\rangle \simeq 3\,\mathrm{g\,cm^{-3}}$.

The $\cos\theta_{\mathrm{zen}}$ of the neutrino directly translates to the number, the length, and the density of the layers of the Earth traversed. In practical calculations, the propagation of neutrinos through the individual layers of approximately constant density is therefore treated sequentially, i.e., on a layer-by-layer basis (see Section 7.2.1). In Figure 3.7, we compare the $\nu_\mu \to \nu_\tau$ oscillation probability in vacuum (i.e., assuming zero matter density) to that computed with the four-layer Earth density profile [145]. We observe
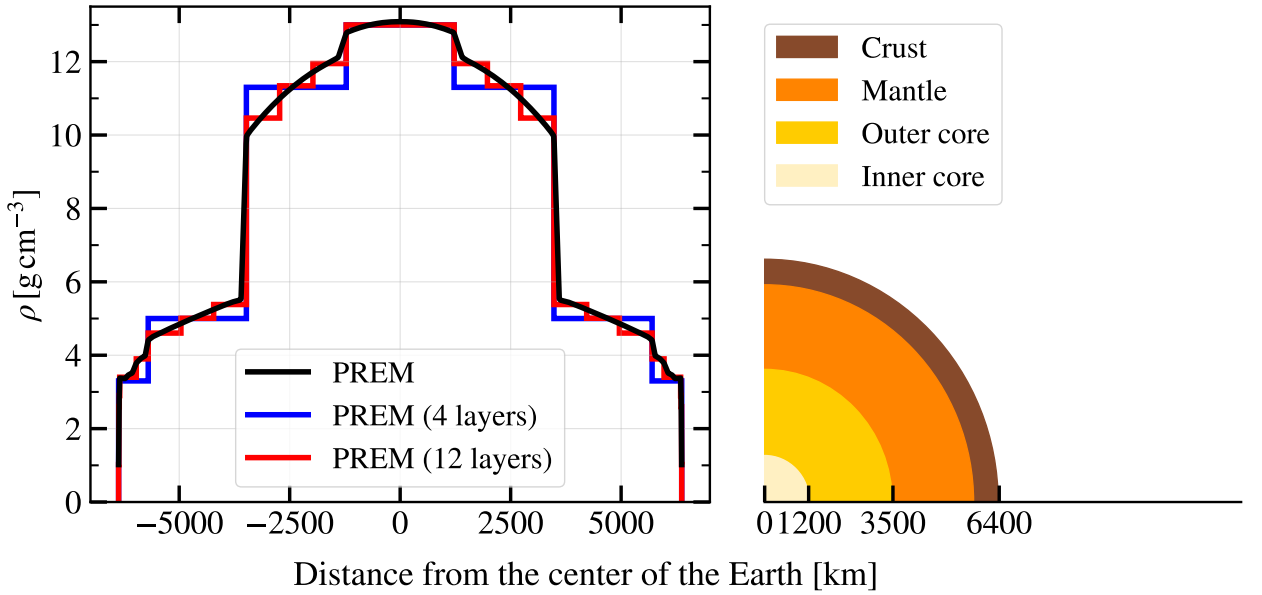
FIGURE 3.6 – The density of the Earth according to the Preliminary Earth Model (PREM) and its approximations with 4 or 12 layers of constant density [145], shown as a function of distance from the Earth's center. Here, ρ is the total matter density, and the electron (neutron) fractions Y_e (Y_n) must be applied to calculate the CC (NC) potential contributions to the Hamiltonian in Equation (2.22). The figure is reproduced from a similar variant in [146].

that the modification of the oscillation probability due to the presence of matter is the strongest at $E_\nu < 10\,\mathrm{GeV}$ and forms an irregular banded pattern in the (E_ν, $\cos\theta_{\mathrm{zen}}$) space. The irregularities in the modified oscillation probability are particularly strong for the most upgoing neutrinos, i.e., those entering and leaving the Earth's core. This is a manifestation of the phenomenon called *parametric resonance* [149], i.e., enhancement of the $\nu_\mu \to \nu_e$ oscillation probability due to periodic changes in the electron matter potential[2]. The neutrinos which pass through the Earth's mantle, the core, and the mantle again on their way to the detector experience one and a half periods of the "castle wall-like" density profile, with abrupt density changes at the core-mantle boundary and the subsequent modifications to the oscillation phase. The specific conditions (in terms of neutrino energy, zenith angle, and the layer densities) for the parametric resonances to occur have been studied analytically in [149] and numerically in [150].

In addition to the parametric resonances, atmospheric neutrinos undergo the so-called **M**ikheev-**S**mirnov-**W**olfenstein (**MSW**) resonance [147, 148] – once in the core and once in the mantle. While the parametric resonances occur due to the shifts in the phase of the oscillation probability in periodic density profiles, the MSW resonance is driven by the strongly enhanced effective mixing angle in matter – and the subsequent increase in the oscillation amplitude. In particular, neglecting the mass splitting Δm_{21}^2 – and, thus, the mixing in the 1-2 sector – the effective mixing angle in matter in the 1-3 sector becomes [9, 150]:

$$\sin 2\tilde{\theta}_{13} = \frac{\Delta m_{31}^2 \sin 2\theta_{13}}{\Delta \tilde{m}_{31}^2}, \tag{3.2}$$

where the effective mass splitting $\Delta \tilde{m}_{31}^2$ is

$$\Delta \tilde{m}_{31}^2 = \sqrt{(\Delta m_{31}^2 \cos 2\theta_{13} - 2E_\nu V_{\mathrm{CC}})^2 + (\Delta m_{31}^2 \sin 2\theta_{13})^2}. \tag{3.3}$$

---

[2]As the result, the $\nu_\mu \to \nu_\mu$ and the $\nu_\mu \to \nu_\tau$ oscillation probabilities are also affected, which is what we observe in Figure 3.7.
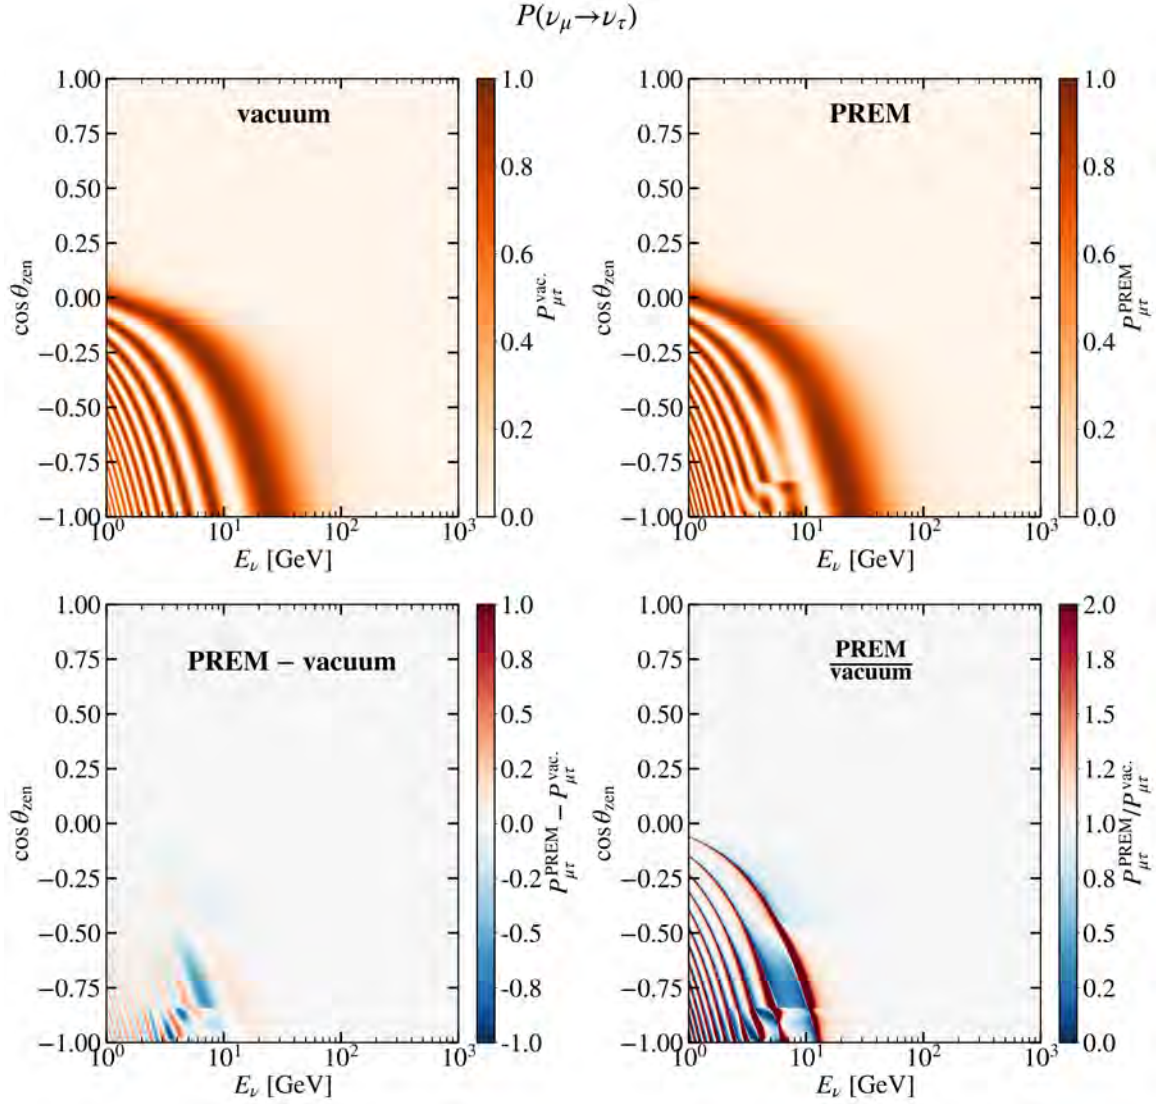
$$P(\nu_\mu \rightarrow \nu_\tau)$$

FIGURE 3.7 – The probability of $\nu_\mu \rightarrow \nu_\tau$ oscillations for atmospheric neutrinos, as computed using the NEURTHINO software [5] with the NuFit 5.2 (2022) oscillation parameters [65, 66] for normal mass ordering. The PREM Earth model [145] is assumed for the computation of matter effects [147, 148] in the top right panel, while the top left panel assumes zero matter density. The bottom left (right) panel shows the difference (the ratio) between the vacuum and the PREM-based oscillation probabilities.

The mixing in the 1-3 sector can therefore become maximal ($\sin 2\tilde{\theta}_{13} = 1$) when

$$\Delta m_{31}^2 \cos 2\theta_{13} - 2E_\nu V_{CC} = 0 \implies E_\nu = \frac{\Delta m_{31}^2 \cos 2\theta_{13}}{2V_{CC}} = \frac{\Delta m_{31}^2 \cos 2\theta_{13}}{2\sqrt{2}G_F N_A Y_e \rho}. \tag{3.4}$$

The condition in Equation (3.4) is the MSW resonance condition for the neutrino energy E_ν. As an example, let us consider neutrinos with $\cos\theta_{zen} \simeq -0.75$, i.e., $\theta_{zen} \simeq 138°$. These neutrinos only pass through the Earth's crust and mantle, and the density in Equation (3.4) can be well approximated by the density of the mantle alone. The $\nu_\mu \rightarrow \nu_e$ oscillation probability for this scenario is shown in Figure 3.8, which we overlay with the vacuum oscillation probability and the energy that satisfies the MSW resonance condition from Equation (3.4). We find a good agreement between the predicted and the observed extrema of $P(\nu_\mu \rightarrow \nu_e)$.

Finally, we note that the $\nu_\mu \rightarrow \nu_e$ oscillation probability is enhanced only in case of the normal neutrino mass ordering. In case of the inverted ordering, both parametric and MSW resonances appear in the $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ channel instead, as illustrated in Figure 3.9.
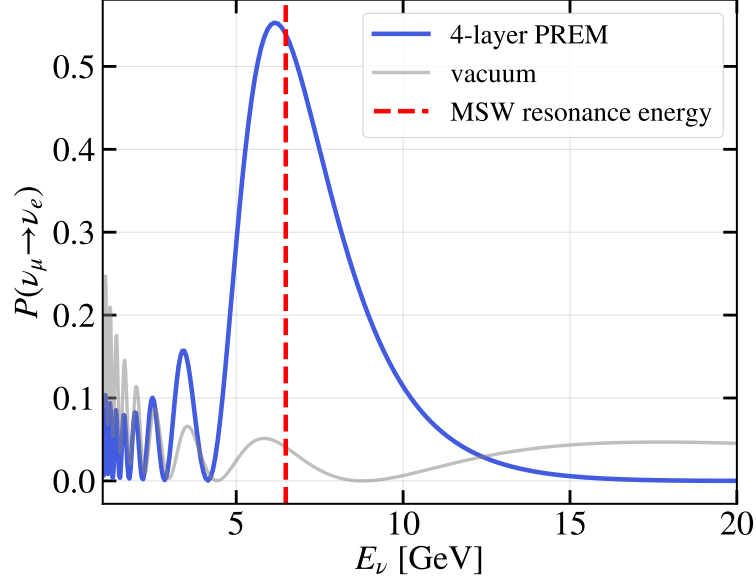
FIGURE 3.8 – The probability of $\nu_\mu \to \nu_e$ oscillations for atmospheric neutrinos with $\cos\theta_{\text{zen}} \simeq -0.75$ in the 4-layer PREM profile and in vacuum (see Figure 3.7 for the technical details of the calculation). Additionally, the MSW resonance energy predicted from Equation (3.4) is shown as the dashed vertical line, where we assumed $\rho = 5\,\text{g\,cm}^3$.
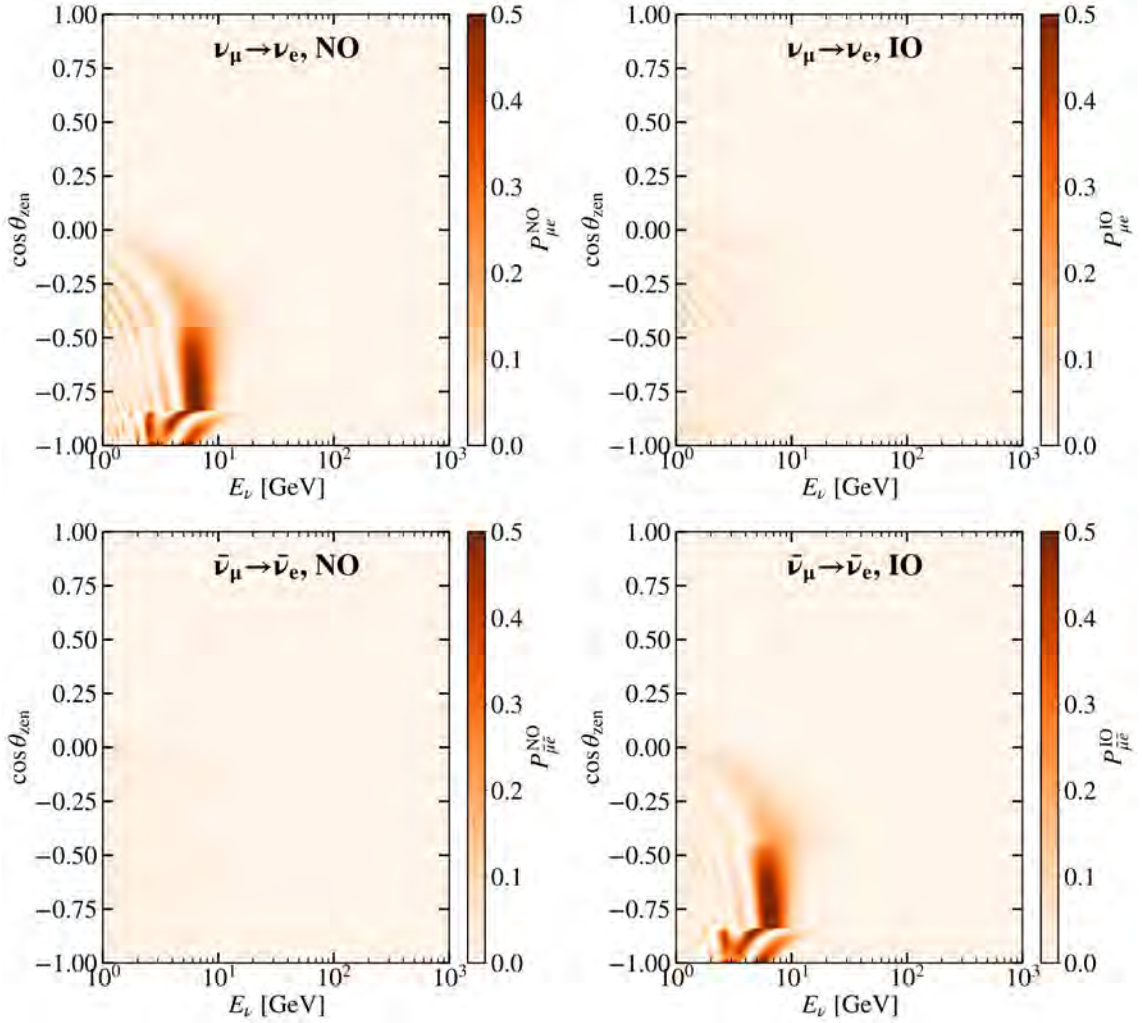
FIGURE 3.9 – The probability of $\nu_\mu \to \nu_e$ (top) and $\bar\nu_\mu \to \bar\nu_e$ oscillations for atmospheric neutrinos, as computed using the NEURTHINO software [5] with the NuFit 5.2 (2022) oscillation parameters [65, 66]. The left (right) column assumes the normal (inverted) neutrino mass ordering.

Detection in terrestrial media

4.1 neutrino interactions with matter

At the energies of relevance to the atmospheric neutrino oscillations resolvable with IceCube-DeepCore – i.e., few GeV to few hundreds of GeV – the three main channels of neutrino interactions in matter are **charged current quasielastic scattering** (shortly **CCQE** or QEL), **resonant scattering (RES)**, and **deep inelastic scattering (DIS)**. In this section, we devote our attention to neutrino-nucleon interactions, i.e., disregarding any nuclear medium effects.

Let us consider a generic neutrino-nucleon interaction, $\nu_\alpha N \rightarrow \ell_\alpha X$, where N is the target nucleon, ℓ_α is either a charged or a neutral final-state lepton (i.e., $\ell_\alpha^\pm$ in case of a charged current (CC) or ν_α in case of a neutral current (NC) interaction), and X is the final-state hadronic system. The regime into which this neutrino-nucleon interaction falls depends on the neutrino energy E_ν, the momentum transfer Q (or, more conventionally, its square Q^2), and the energy $\mathcal{V}$ transferred to the nucleon. The latter is equal to

$$\mathcal{V} = E_\nu - E_\ell = E_X - M_N, \tag{4.1}$$

where E_ℓ is the outgoing lepton energy, E_X is the total energy of the final-state hadronic system, and M_N is the nucleon mass. From here, we can derive the *inelasticity* (y) and the *invariant mass of the hadronic final state* (W):

$$y = \frac{E_\nu - E_\ell}{E_\nu} = \frac{\mathcal{V}}{E_\nu}; \tag{4.2}$$

$$W^2 = M_N^2 - Q^2 + 2 M_N \mathcal{V}. \tag{4.3}$$

The main distinction between the CCQE, the RES, and the DIS interactions, from an experimental perspective, is in the types and the masses of the final-state particles X. For example, when X consists of only one nucleon, the mass W of the final state is simply the mass of that nucleon, i.e., $M_N \approx 0.938\,\text{GeV}$[1]. For Equation (4.3), this means that $Q^2 = 2 M_N \mathcal{V}$; this relationship defines the (quasi)elastic limit, which is satisfied only by two types of interactions covered in this thesis: $\nu_\alpha n \rightarrow \ell_\alpha^- p$ and $\bar{\nu}_\alpha p \rightarrow \ell_\alpha^+ n$[2]. Thus, the CCQE interactions can be visualized as the 1:1 line in the

---

[1] Although we used M_N earlier to denote the mass of the *target* rather than final-state nucleon, the difference between the proton and the neutron masses can be considered negligible for the purposes of all calculations in this thesis.

[2] Note that production of hyperons Λ^0, Σ^0, and Σ^- is also possible in $\bar{\nu}_\alpha$-CCQE interactions – however, it is Cabibbo-suppressed due to the presence of a strange quark in these baryons [151].

Q^2-$2M_N\mathcal{V}$ plane (see Figure 4.1). In the resonant scattering case, where short-lived resonances are produced and decay into multiple hadrons, W is larger than the mass of a single nucleon. For example, in the single pion production case through the Δ^+ resonance, $\nu_\alpha n \to \ell_\alpha^- \Delta^+ \to n\pi^+$, the mass of the intermediate hadronic state is $W = m_\Delta \approx 1.232\,\mathrm{GeV}$, and the mass of the final hadronic state is $W = M_N + m_\pi \approx 1.077\,\mathrm{GeV}$. The exact upper bound of the resonant scattering region in terms of W (which corresponds to the onset of the deep inelastic scattering region) is not strictly defined. The typical choices are $W = 1.4\,\mathrm{GeV}$, $W = 1.7\,\mathrm{GeV}$, or $W = 2.0\,\mathrm{GeV}$, depending on which resonances are allowed to be included in the DIS region. It is common [152, 153] to define the region $m_\Delta < W < 2.0\,\mathrm{GeV}$ as the **shallow inelastic scattering (SIS)** region, which denotes a transition between the RES and the DIS regimes.
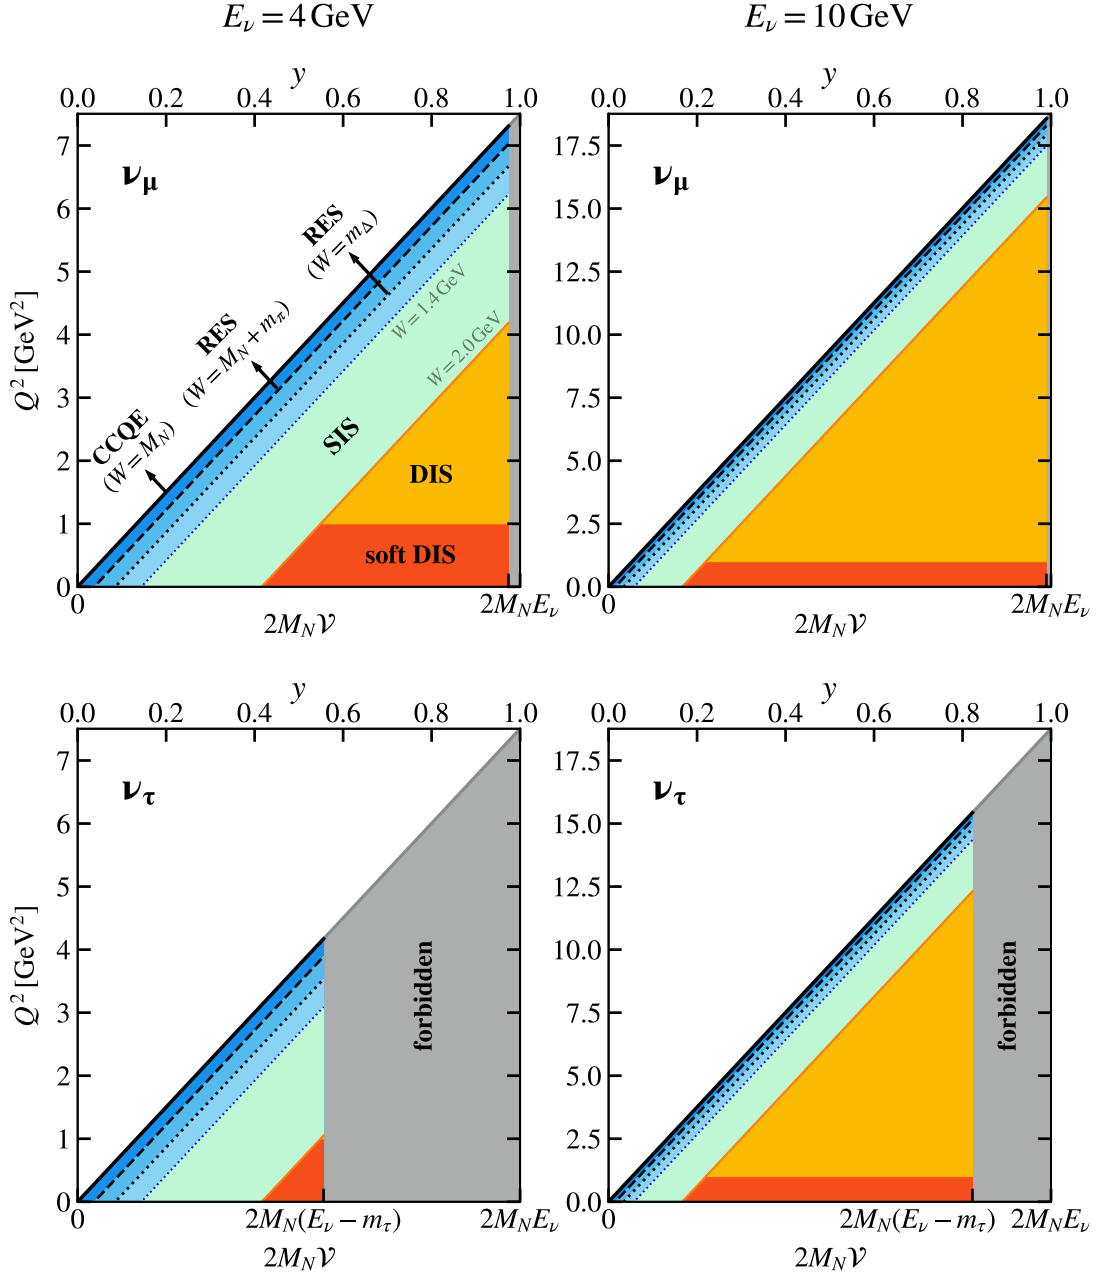
FIGURE 4.1 – The relationships between the kinematic variables in the CCQE, RES, SIS, and DIS regimes of ν-nucleon scattering, shown for ν_μ (top) and ν_τ (bottom) at two different energies. The soft DIS region is defined as DIS at $Q^2 < 1\,\mathrm{GeV}^2$. An important feature to compare between the different subpanels is the amount of DIS phasespace available for ν_μ *vs.* ν_τ, and at $4\,\mathrm{GeV}$ *vs.* $10\,\mathrm{GeV}$. The figure is inspired by [152].

In the subsections that follow, we describe each of the mentioned neutrino-nucleon interaction channels individually.

4.1.1 *Quasielastic scattering*

In Figure 4.2, we show the Feynman diagrams of the CCQE neutrino-nucleon scattering. As already mentioned in the preface to this subsection, this interaction predominantly yields a single nucleon in the final state and is therefore accessible to neutrinos with energies as low as a few hundreds of MeV.

FIGURE 4.2 – Feynman diagrams illustrating charged current neutrino-neutron (left) and antineutrino-proton (right) quasielastic scattering.

The differential cross section of the CCQE interaction with respect to the momentum transfer Q^2 is commonly calculated with the Llewelyn Smith model [154, 155]:
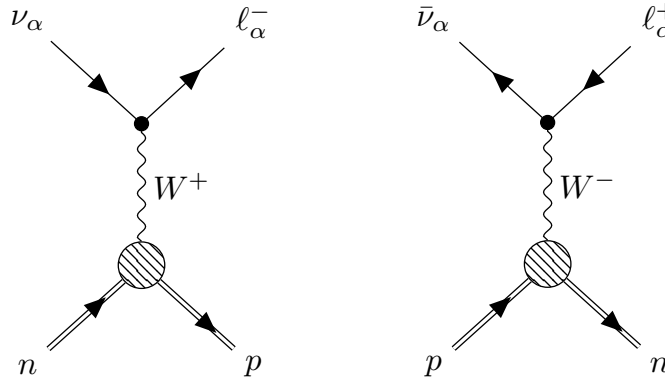
$$\frac{d\sigma_{\text{CCQE}}}{dQ^2} = \frac{M_N^2 G_F^2 \cos^2\theta_{\text{Cab}}}{8\pi E_\nu^2} \left(\frac{M_W^2}{Q^2 + M_W^2}\right)^2 \left[A(Q^2) \pm B(Q^2)\frac{s-u}{M_N^2} + C(Q^2)\left(\frac{s-u}{M_N^2}\right)^2\right], \quad (4.4)$$

where the "+"("-") sign corresponds to neutrinos (antineutrinos), $s-u = 4M_N E_\nu - Q^2 - m_\ell^2$, G_F is the Fermi constant, θ_{Cab} is the Cabibbo angle[3], and $M_W \simeq 80.4\,\text{GeV}$ is the W boson mass. The coefficients A, B, and C are functions of three Q^2-dependent *form factors*, whose full expressions can be found in [155]. Under the conserved vector current (CVC) hypothesis, two of these form factors can be related to the electromagnetic form factors and therefore determined from the eN scattering data. This leaves freedom only in the axial-vector form factor $F_A(Q^2)$,

$$F_A(Q^2) = \frac{F_A(Q^2 = 0)}{1 + \frac{Q^2}{M_{A,\text{QE}}^2}} = \frac{-g_A}{1 + \frac{Q^2}{M_{A,\text{QE}}^2}} = \frac{-1.27}{1 + \frac{Q^2}{M_{A,\text{QE}}^2}}, \quad (4.5)$$

where the g_A constant is determined from neutron β^- decay, and the axial mass $M_{A,\text{QE}}$ is a free parameter to be constrained by neutrino-nucleon scattering experiments. At the time of writing, the latest global axial mass constraint is $M_{A,\text{QE}} = 1.014 \pm 0.014\,\text{GeV}$ [156, 157]. It is a free systematic parameter in the tau neutrino appearance analysis with the oscNext event selection covered in Chapter 9, with a nominal value of $0.99\,\text{GeV}$ and -15%/+25% relative 1σ uncertainty [158]. In the non-unitarity analysis described in Chapters 7 and 8, $M_{A,\text{QE}}$ is fixed at $0.99\,\text{GeV}$.

To obtain the total CCQE cross section from Equation (4.4), the differential CCQE cross section must be integrated over Q^2 over the kinematically allowed range, $Q^2 \in$

---

[3]Note that the usual notation for the Cabibbo angle is θ_C; however, we use this notation for the angle of the Cherenkov cone in Section 4.2.

$[Q^2_{\min}, Q^2_{\max}]$. The bounds of this integration range are [159]:

$$Q^2_{\min(\max)} = \frac{2E_\nu^2 M_N - M_N m_\ell^2 \mp E_\nu m_\ell^2 \mp E_\nu \sqrt{(s - m_\ell^2)^2 - 2(s + m_\ell^2)M_N^2 + M_N^4}}{2E_\nu + M_N}, \qquad (4.6)$$

where $s = M_N^2 + 2M_N E_\nu$. In Figure 4.3, we plot the total CCQE cross section thereby predicted by the Llewelyn Smith model for neutrino and antineutrino scattering on free nucleons and compare it with the historical experimental data. We use $M_{A,\mathrm{QE}} = 1.014 \pm 0.014\,\mathrm{GeV}$ for the Llewelyn Smith model prediction and the Budd-Bodek-Arrington form factor parameterizations [160].
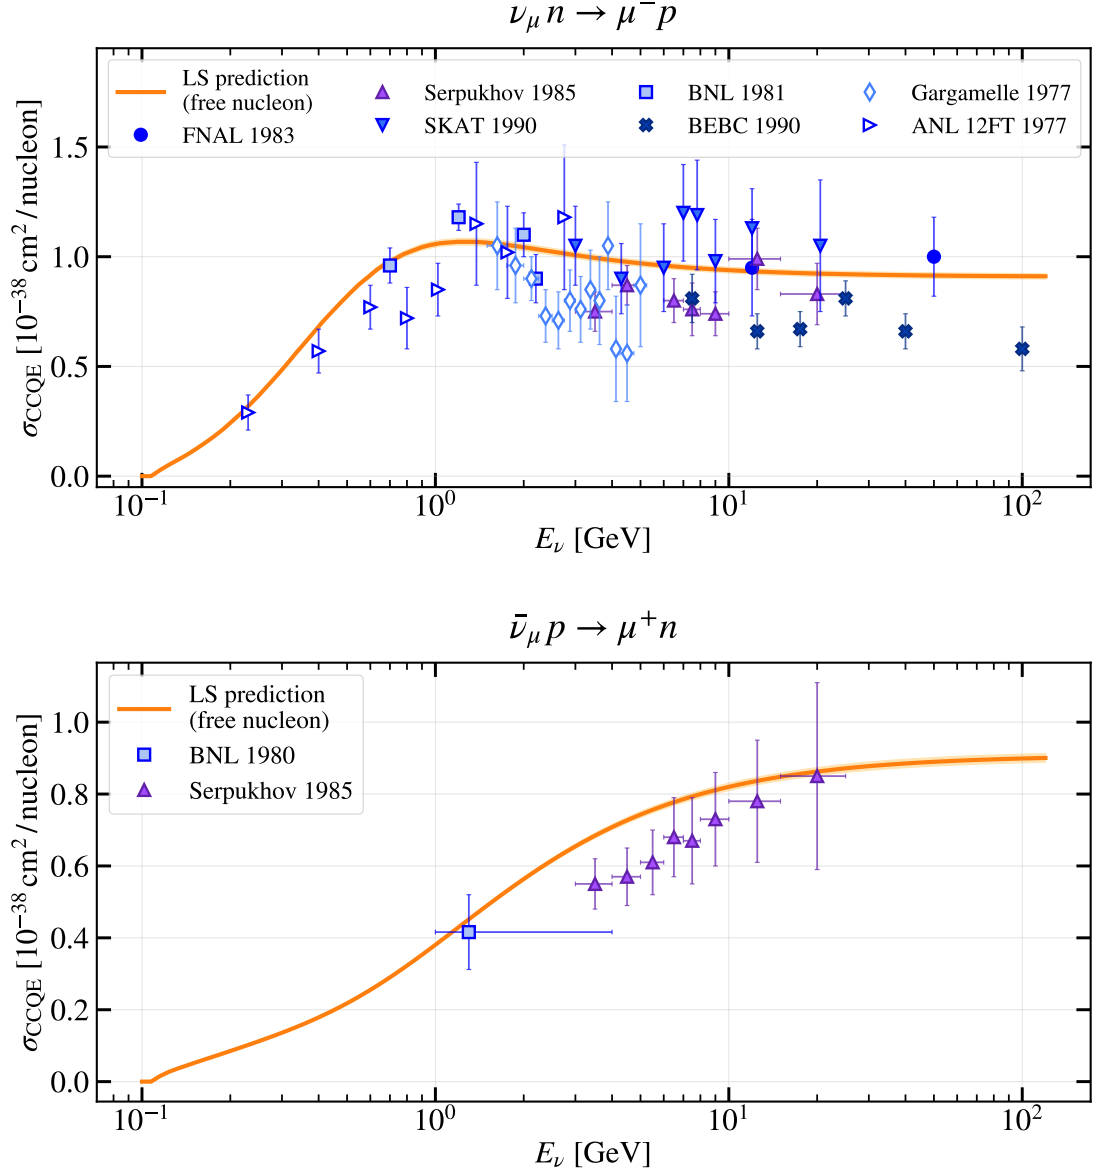
Figure 4.3 – Charged current quasielastic scattering (CCQE) cross section for neutrinos (top) and antineutrinos (bottom), as predicted by the Llewelyn Smith ("LS") model [154, 155, 161] and measured in deuterium target (ANL [162], BNL [163], FNAL [164], BEBC [165]) and heavier-target (Gargamelle:$C_3H_8CF_3Br$ [166], Serpukhov:Al [167], and SKAT:CF_3Br [168]) experiments. Note that only the experiments whose measurements were available in the HEPDATA online database at the time of writing are included in this figure. Figures including more datasets can be found in e.g. [34, 151].

4.1.2 *Resonant scattering*

While CCQE is the dominant neutrino interaction channel in the 0.1 GeV–1 GeV energy range [151], at ~1 GeV the resonant scattering takes over – a process in which excited baryon states are produced and quickly decay, typically yielding one or more pions in the final state. Two example charged current resonant scattering diagrams are shown in Figure 4.4.

FIGURE 4.4 – Feynman diagrams illustrating charged current neutrino-proton (left) and neutrino-neutron (right) resonant scattering with single-pion production through the Δ resonance with mass $m_\Delta \simeq 1.232\,\mathrm{GeV}$ and mean lifetime $\tau_\Delta \simeq 5.63 \cdot 10^{-24}\,\mathrm{s}$ [34].

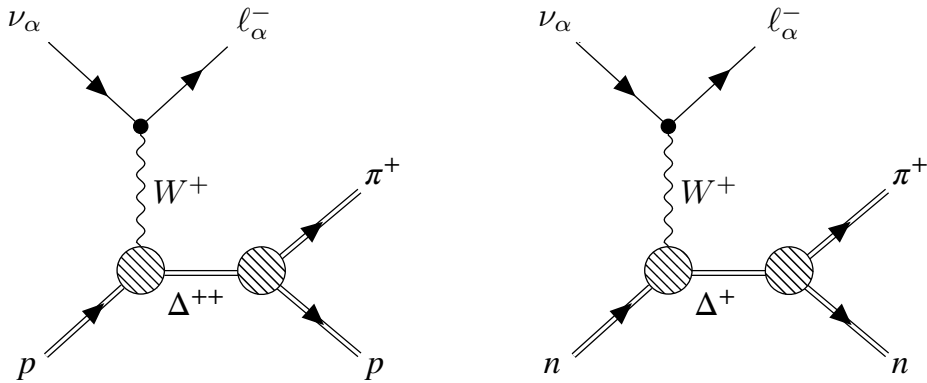Resonant scattering can occur on both protons and neutrons for both neutrinos and antineutrinos, and through both W and Z exchange. In addition, $\mathcal{O}(20)$ different resonances can be excited in the intermediate state, with complex interferences between the neighboring resonant states. A model taking into account the different interaction channels, resonances, and interferences of the latter was developed by Rein and Sehgal [169] and is most commonly used in neutrino event generation codes. Similarly to the quasielastic scattering case, the Q^2 dependence of the resonant scattering cross section is captured through form factors, where the axial mass $M_{A,\mathrm{RES}}$ is a free parameter determined from neutrino-nucleon scattering experiments. In Figure 4.5, we plot the Rein-Sehgal model prediction for the $\nu_\mu p \rightarrow \mu^- p \pi^+$ process, as evaluated in GENIE [158] with $M_{A,\mathrm{RES}} = 1.12\,\mathrm{GeV}$, and overlay this prediction with historical experimental data.

This choice of $M_{A,\mathrm{RES}}$ corresponds to the latest global average value, $M_{A,\mathrm{RES}} = 1.12 \pm 0.03\,\mathrm{GeV}$ [157, 172], and is the nominal value of the corresponding systematic parameter in the tau neutrino appearance analysis with the oscNext event selection. Despite $\lesssim 3\%$ relative uncertainty on the aforementioned global estimate, the oscNext ν_τ analysis assigns $\pm 20\%$ 1σ relative uncertainty, as prescribed by the GENIE Collaboration [158]. As is the case with the quasielastic axial mass, $M_{A,\mathrm{RES}}$ is fixed at the nominal value for the non-unitarity analysis in Chapters 7 and 8.

4.1.3 *Deep inelastic scattering*

Deep inelastic scattering is the dominant neutrino interaction channel at $E_\nu \gtrsim 10\,\mathrm{GeV}$. These high energies allow for higher momentum transfers Q^2 to the target system compared to the CCQE and the RES processes, meaning that the parton structure of the nucleon can be resolved. As illustrated in Figure 4.6, DIS begins with a neutrino scattering off an individual valence or sea quark in the nucleon, q, which carries away a fraction of the nucleon momentum and causes the breakup of the nucleon with a
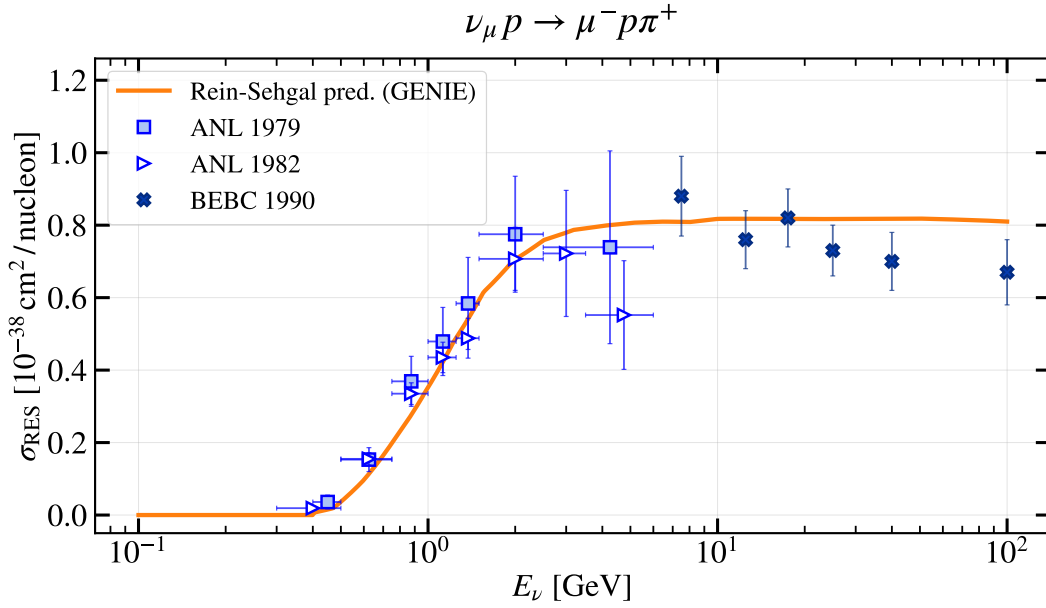
FIGURE 4.5 – Exclusive charged current resonant scattering (RES) cross section for the $\nu_\mu p \rightarrow \mu^- p \pi^+$ channel, as predicted by the Rein-Sehgal model [169] in GENIE [158] and measured in deuterium target experiments (ANL [170, 171] and BEBC [165]). Note that this figure includes only the experiments from [151] whose measurements were recorded in a tabular form in the corresponding publications. Figures including more datasets can be found in e.g. [34, 151].

production of a multi-hadron system X. This can proceed as both charged current and neutral current scattering.

Chapter 11 of this thesis is dedicated to CC DIS calculations at leading order, which is why we omit any analytical descriptions of the DIS cross section in the present section. Instead, in Figure 4.7, we show its prediction obtained directly from the GENIE event generator, as compared to select inclusive cross section data.
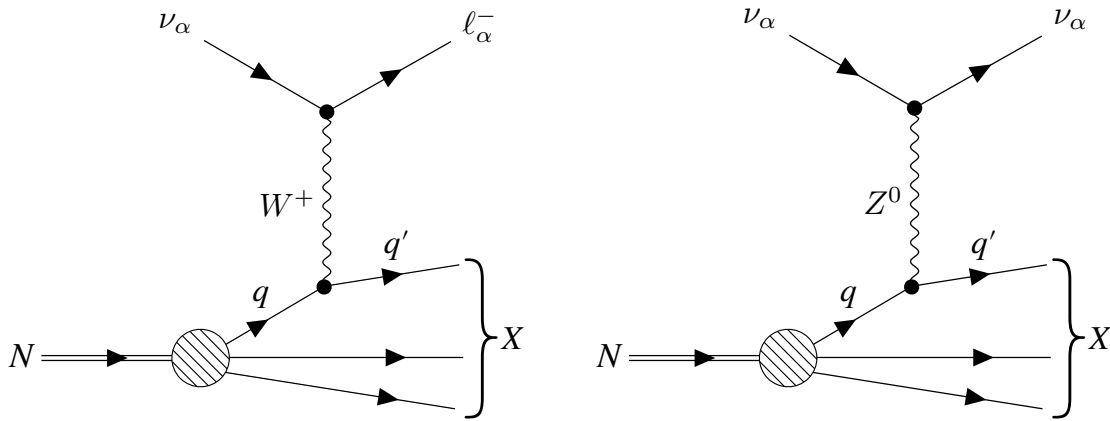
FIGURE 4.6 – Feynman diagrams illustrating charged current (left) and neutral current (right) neutrino-nucleon scattering.

The shown GENIE cross sections are evaluated at effective leading order (LO) with the GRV98 parton distribution functions (PDFs), which include special Bodek-Yang corrections for the region of low Q^2 ($\lesssim 1\,\mathrm{GeV}^2$) [177, 178]. This prescription represents the nominal cross section model in all of the analyses covered in this thesis, for both CC and NC interactions. Many more PDF sets have been derived by multiple experimental
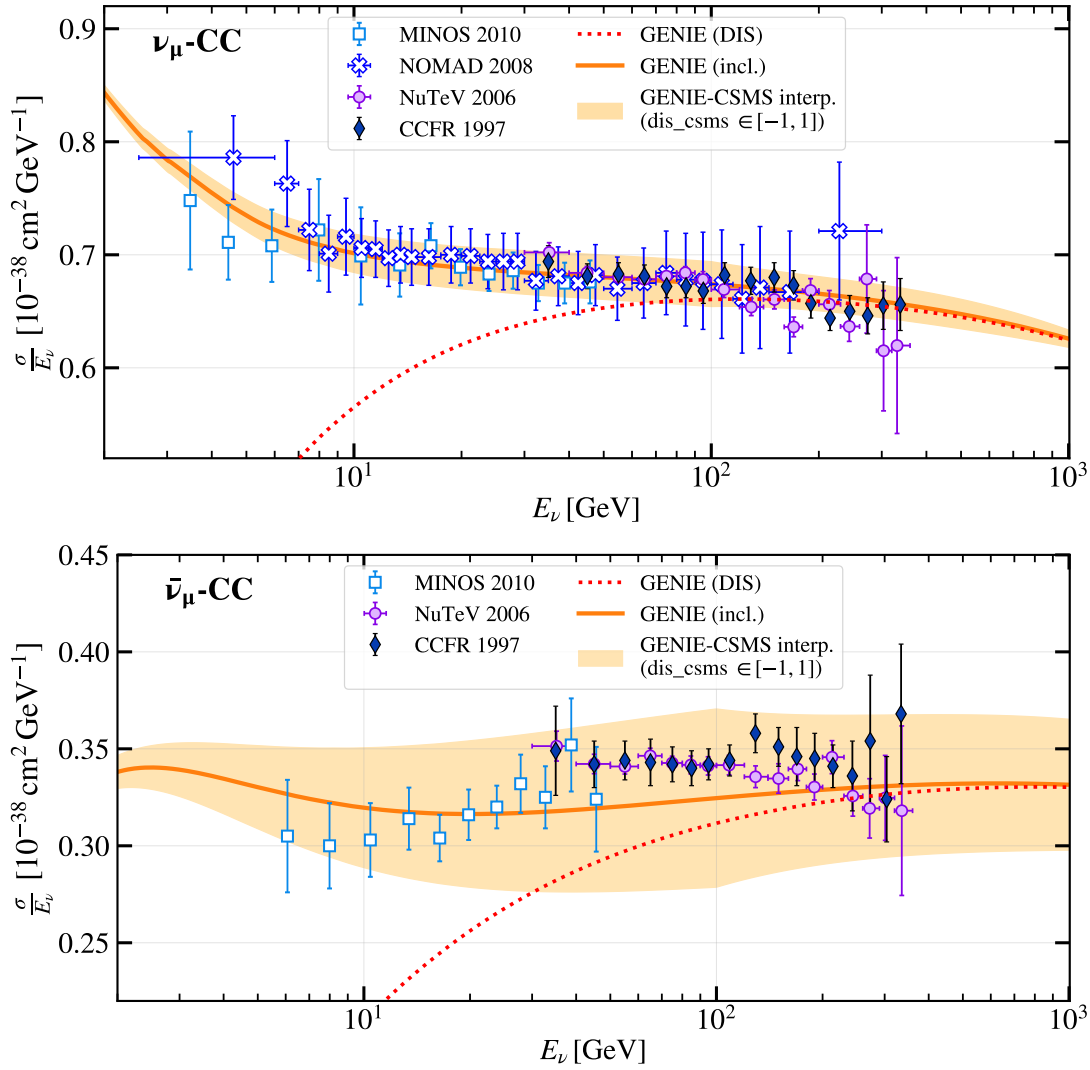
FIGURE 4.7 – Inclusive charged current ν_μ (top) and $\bar{\nu}_\mu$ (bottom) cross section, as predicted by GENIE [158] for an isoscalar target and measured in CCFR [173], NuTeV [174], NOMAD [175], and MINOS [176] experiments. The GENIE DIS cross section is evaluated at effective LO with the GRV98 PDFs [177, 178]. The shaded band around the GENIE prediction corresponds to the 1σ range of the "dis_csms" systematic parameter interpolating between the GENIE and the CSMS [179] DIS cross sections (see [180] and text for details). *The data for this figure was provided by G. P. Zeller.*

collaborations since 1998, resulting in updated cross section predictions. In particular, the HERA PDFs measured from the combined H1 and ZEUS experimental data [181] form the basis of the Cooper-Sarkar-Mertsch-Sarkar (**CSMS**) cross section model [179]. At present, this is a recommended DIS cross section model at $E_\nu \geq 100\,\text{GeV}$ in all Ice-Cube analyses; however, it is not applicable below 100 GeV due to the underlying HERA PDFs being valid only down to $Q^2 = 1\,\text{GeV}^2$. In order to incorporate the predictions of both the CSMS model and the GENIE-based LO calculation, the "dis_csms" systematic parameter has been devised for neutrino oscillation analyses with IceCube-DeepCore (see e.g. [54, 180] for a detailed description). This parameter enables interpolation of both total and single-differential (with respect to inelasticity) cross sections between the two models, such that dis_csms = 0 corresponds to the GENIE-like cross section, and dis_csms = 1 – to the CSMS-like cross section. This parameter has been further artificially extended to negative values in order to approximately cover the experimental data in Figure 4.7. In the ν_τ appearance analysis with the oscNext event selection (Chapter 9), it covers a range of [-3.0, 3.0], with a Gaussian prior of width 1. In the global fit for non-unitarity (Chapters 7 and 8), dis_csms is fixed at 0.

For atmospheric neutrinos, the main detection technique at GeV energies and above is the detection of Cherenkov radiation emitted due to the passage of charged neutrino interaction products (such as charged leptons or hadrons) through dielectric media. To record the Cherenkov signal with light sensors situated at a sizeable distance from the interaction point, the medium has to be transparent to the Cherenkov radiation, which is satisfied by e.g. water or ice. The two presently operating detectors exploiting such naturally occurring media are KM3NeT [182] and IceCube [41, 45]; additionally, the Super-Kamiokande detector uses a water-filled tank for the same purpose [183]. While future detectors such as JUNO [184, 185] and DUNE [150, 186, 187] aim to exploit liquid scintillator technology for $\mathcal{O}$ (0.1–10 GeV) atmospheric neutrino detection, we focus our attention only on Cherenkov radiation in this chapter. We first describe the origin of such radiation qualitatively in Section 4.2.1 and then proceed with a thorough derivation of the Cherenkov frequency spectrum in Section 4.2.2. The material in these sections partly follows the original references on Cherenkov radiation such as [188, 189] as well as the classic textbooks [190, 191], and these sources can be advised for further reading.

4.2.1 *Qualitative description*

For concreteness and for the purpose of relevance to this thesis, we will consider that the charged particle in question is a muon (μ^-) traveling through ice in the direction of the positive z-axis. This is illustrated in the left panel of Figure 4.8, where the muon travels from point A to point B. The electric field of the negatively charged muon forces
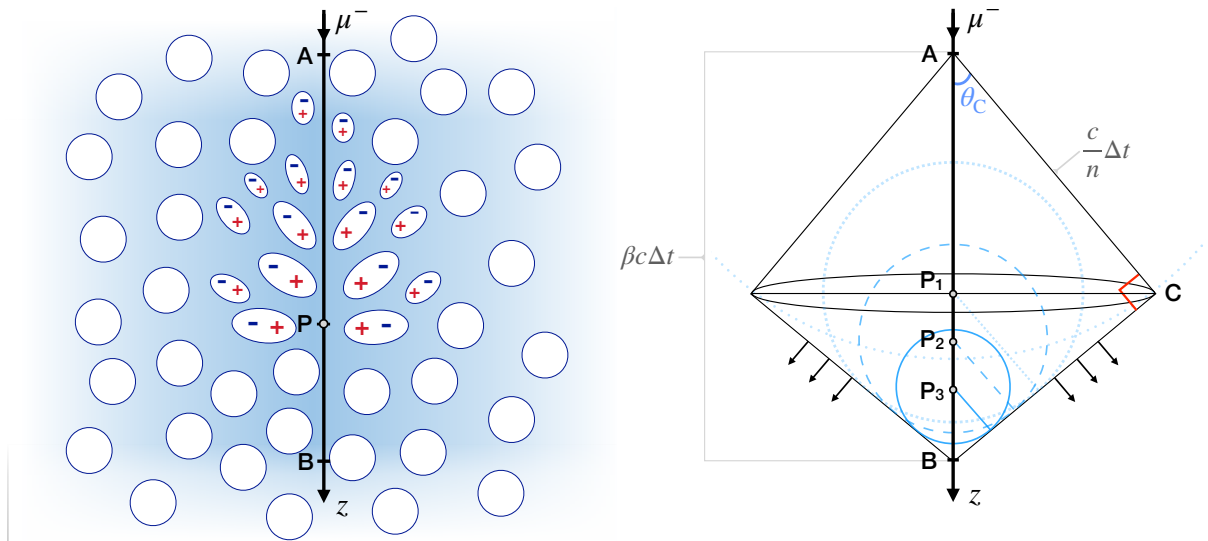
FIGURE 4.8 – *Left*: Illustration of a negatively charged muon traveling from point A to point B in a dielectric medium faster than the speed of light in this medium. The instantaneous electric dipoles induced in proximity of the muon track (drawn at the time the muon passes point P) eventually relax to their original charge distributions, and Cherenkov photons are emitted. *Right*: The geometry of the Cherenkov radiation emission (see text for details). The figure is inspired by [188].

the electrons in the water molecules to redistribute such that there is an excess of positive charge closer to the muon track, and an excess of negative charge on the other side of the molecule. This creates a population of *induced electric dipoles* around the muon track, which, as the muon travels further along, relax back to the ordinary charge distributions with the emission of *Cherenkov photons*. The required condition for the Cherenkov radiation to occur at large distances from the muon track is the non-zero

net polarization field, arising if and only if the muon travels ***faster than the speed of light in the medium***, thereby breaking the polarization symmetry along the z-axis as in Figure 4.8. The speed of light in ice is modified compared to its value in vacuum by the refractive index n,

$$c \to c' = \frac{c}{n},$$ (4.7)

where $n \simeq 1.31$ for visible light [4] [192]. Thus, it is required that the muon velocity $v = \beta c$, with β being the Lorentz velocity factor, is greater than $\frac{c}{n}$, or, alternatively,

$$\beta > \frac{1}{n}.$$ (4.8)

For a muon, this translates to the condition that its kinetic energy is greater than $\sim 58\,\mathrm{MeV}$. In Table 4.1, we list the equivalent requirements for Cherenkov radiation emission by several other particle types relevant for neutrino detection in ice.

particle	mass [MeV]	$E_{\min}^{\mathrm{kin}}$ [MeV]
$e^\pm$	0.511	0.28
$\mu^\pm$	105.7	57.9
$\pi^\pm$	139.6	76.5
$K^\pm$	493.7	270.6
$p, \bar{p}$	938.3	514.2

TABLE 4.1 – Minimum kinetic energy, $E_{\min}^{\mathrm{kin}}$, required for Cherenkov radiation emission in ice by various particle types, assuming the refractive index $n = 1.31$.

Now, let us discuss the geometry of the Cherenkov radiation emission as illustrated in the right panel of Figure 4.8. Any point along the muon track AB is a source of spherical electromagnetic waves, as depicted e.g. for points P_1, P_2, and P_3. In time Δt it takes for the muon to reach point B, the wave originally emitted from point A will propagate a shorter distance AC in all directions, as it is delayed compared to the muon. The waves emitted from P_1, P_2, and P_3 will cover even shorter distances within this time. The interference between the waves emitted from point A and all possible P_i between A and B is destructive everywhere except the surface of the cone obtained by rotating the line element BC about the z-axis. The geometry of this cone is defined by the condition that

$$\mathrm{AC} = \mathrm{AB}\cos\theta_{\mathrm{C}} \implies \beta c \Delta t = \frac{c}{n}\Delta t,$$ (4.9)

or, equivalently,

$$\cos\theta_{\mathrm{C}} = \frac{1}{\beta n},$$ (4.10)

where $\theta_{\mathrm{C}} \equiv \angle\mathrm{BAC}$ is the angle that the Cherenkov photons appear to travel at relative to the muon track. Assuming $\beta = 1$ and $n = 1.31$, we find $\theta_{\mathrm{C}} \simeq 40.2°$ for an ultrarelativistic charged particle in ice. However, we note that in general n depends on the light frequency, making the medium such as ice mildly dispersive and resulting in a frequency-dependent opening angle of the Cherenkov cone.

4.2.2 Derivation of the Cherenkov spectrum

The passage of a charged particle through a medium generates a current collinear with the direction of the particle motion, which we choose to be in the direction of the positive z-axis as in Section 4.2.1. This current creates an electric field $\mathbf{E}$ with

---

[4]This refers to the real part of the complex refractive index. In the visible part of the electromagnetic spectrum, the imaginary part is of $\mathcal{O}(10^{-8})$ [192].

longitudinal and radial components E_z and E_ρ, as well as a magnetic field **H** with an azimuthal component H_ϕ, where ρ, z, and φ constitute a standard set of cylindrical coordinates.

In the most general medium, some amounts of internal electrical polarization **P** and magnetization **M** can be present either as permanent material properties or those induced by external polarizing/magnetizing forces. This necessitates the definition of two more fields, **D** and **B**, which represent the total electric and magnetic flux densities, respectively:

$$\mathbf{D} = \epsilon_0 \mathbf{E} + \mathbf{P} = \epsilon \mathbf{E} \equiv \epsilon_r \epsilon_0 \mathbf{E}; \tag{4.11a}$$

$$\mathbf{B} = \mu_0 (\mathbf{H} + \mathbf{M}) = \mu \mathbf{H} \equiv \mu_r \mu_0 \mathbf{H}. \tag{4.11b}$$

In Equation (4.11), ϵ_0 (ϵ) and μ_0 (μ) are the absolute electric permittivity and magnetic permeability of vacuum (medium in question), and the dimensionless ϵ_r and μ_r are the permittivity and permeability of the medium relative to vacuum. The speed of light connects ϵ_0 and μ_0 through

$$c = \frac{1}{\sqrt{\epsilon_0 \mu_0}}. \tag{4.12}$$

For a non-magnetic material such as ice, there is no external or induced magnetization, i.e., $\mathbf{M} = 0$ and $\mu_r = 1$. However, as discussed in Section 4.2.1, ice is a dielectric, whose bound electrons get redistributed in response to the electric field of the passing charged particle and result in a net polarization in the vicinity of the particle track. This implies that $\mathbf{P} \neq 0$ and $\epsilon_r > 1$. In fact, ϵ_r is related to the refractive index n of the medium:

$$n = \sqrt{\epsilon_r}. \tag{4.13}$$

These observations transform Equation (4.11) as follows:

$$\mathbf{D} = n^2 \epsilon_0 \mathbf{E}; \tag{4.14a}$$

$$\mathbf{B} = \mu_0 \mathbf{H}. \tag{4.14b}$$

With these definitions, we write down Maxwell's equations for spatial and temporal evolution of the electromagnetic field in a medium:

$$\nabla \cdot \mathbf{D} = \rho_e \quad \Longleftrightarrow \quad \nabla \cdot \mathbf{E} = \frac{\rho_e}{n^2 \epsilon_0}; \tag{4.15a}$$

$$\nabla \times \mathbf{H} = \mathbf{J} + \frac{\partial \mathbf{D}}{\partial t} \quad \Longleftrightarrow \quad \nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \frac{n^2}{c^2} \frac{\partial \mathbf{E}}{\partial t}, \tag{4.15b}$$

where ρ_e is the electric charge density and **J** is the current density. To simplify Equations (4.15a) and (4.15b), we introduce two auxiliary fields: the vector magnetic potential **A** and the scalar electric potential ϕ, such that

$$\mathbf{E} = -\frac{\partial \mathbf{A}}{\partial t} - \nabla \phi; \tag{4.16a}$$

$$\mathbf{B} = \nabla \times \mathbf{A}. \tag{4.16b}$$

The substitution of Equation (4.16) into Equation (4.15) yields

$$\nabla^2 \phi = -\frac{\partial (\nabla \cdot \mathbf{A})}{\partial t} - \frac{\rho_e}{n^2 \epsilon_0}; \tag{4.17a}$$

$$\nabla^2 \mathbf{A} - \frac{n^2}{c^2} \frac{\partial^2 \mathbf{A}}{\partial t^2} = -\mu_0 \mathbf{J} + \nabla \left[\nabla \cdot \mathbf{A} + \frac{n^2}{c^2} \frac{\partial \phi}{\partial t} \right], \tag{4.17b}$$

where we used the identity $\nabla \times (\nabla \times \mathbf{A}) = \nabla(\nabla \cdot \mathbf{A}) - \nabla^2 \mathbf{A}$. By adding $-\frac{n^2}{c^2}\frac{\partial^2 \phi}{\partial t^2}$ to both sides of Equation (4.17a), we obtain

$$\nabla^2 \phi - \frac{n^2}{c^2}\frac{\partial^2 \phi}{\partial t^2} = -\frac{\rho_e}{n^2 \epsilon_0} - \frac{\partial}{\partial t}\left[\nabla \cdot \mathbf{A} + \frac{n^2}{c^2}\frac{\partial \phi}{\partial t}\right]. \tag{4.18}$$

Importantly, since $\mathbf{A}$ and ϕ are only the auxiliary gauge fields, they can be chosen freely as long as the ultimate physical fields $\mathbf{E}$ and $\mathbf{B}$ from Equation (4.16) remain unchanged. The similarity in the right-hand sides of Equation (4.17b) and Equation (4.18) points us to the choice of the *Lorentz gauge* [190], with the condition

$$\nabla \cdot \mathbf{A} + \frac{n^2}{c^2}\frac{\partial \phi}{\partial t} = 0. \tag{4.19}$$

This significantly simplifies Equations (4.17b) and (4.18):

$$\nabla^2 \phi - \frac{n^2}{c^2}\frac{\partial^2 \phi}{\partial t^2} = -\frac{\rho_e}{n^2 \epsilon_0}; \tag{4.20a}$$

$$\nabla^2 \mathbf{A} - \frac{n^2}{c^2}\frac{\partial^2 \mathbf{A}}{\partial t^2} = -\mu_0 \mathbf{J}. \tag{4.20b}$$

Since we are interested in the frequency (ω) spectrum of the emitted radiation, we introduce the Fourier transform $\tilde{F}(\omega)$ of a generic time-dependent function $F(t)$ (either scalar or vector) as

$$\tilde{F}(\omega) = \frac{1}{2\pi}\int_{-\infty}^{\infty} F(t)e^{-i\omega t}\mathrm{d}t; \quad F(t) = \int_{-\infty}^{\infty} \tilde{F}(\omega)e^{i\omega t}\mathrm{d}\omega. \tag{4.21}$$

The time derivatives of the frequency-space function are simply

$$\frac{\partial \tilde{F}(\omega)}{\partial t} = -i\omega\tilde{F}(\omega); \quad \frac{\partial^2 \tilde{F}(\omega)}{\partial t^2} = -\omega^2 \tilde{F}(\omega). \tag{4.22}$$

where we applied the Leibniz integral rule to perform the differentiation with respect to t under the integral sign. This allows us to write down the Fourier-transformed versions of Equation (4.20) as follows:

$$\nabla^2 \tilde{\phi}(\omega) + \frac{\omega^2 n^2}{c^2}\tilde{\phi}(\omega) = -\frac{\rho_e}{n^2 \epsilon_0}; \tag{4.23a}$$

$$\nabla^2 \tilde{\mathbf{A}}(\omega) + \frac{\omega^2 n^2}{c^2}\tilde{\mathbf{A}}(\omega) = -\mu_0 \tilde{\mathbf{J}}(\omega). \tag{4.23b}$$

Importantly, Equation (4.23b) represents three individual equations – one for each component of the vector potential $\mathbf{A}$ in the cylindrical coordinates. However, the current density $\mathbf{J}$ on the right-hand side of Equation (4.23) only has the z component collinear with the direction of the charged particle motion. As in Section 4.2.1, we assume that the travelling particle is a muon μ^- with the electric charge $-e$. Then, the current density in the time domain is

$$J_z(t) = \rho_e v = -\frac{ev}{2\pi\rho}\delta(\rho)\delta(z - vt), \tag{4.24}$$

where v is the muon velocity. In the frequency domain,

$$\tilde{J}_z(\omega) = -\frac{ev}{4\pi^2}\frac{\delta(\rho)}{\rho}\int_{-\infty}^{\infty} e^{-i\omega t}\delta(z - vt)\mathrm{d}t = -\frac{e}{4\pi^2}\frac{\delta(\rho)}{\rho}e^{-\frac{i\omega z}{v}}. \tag{4.25}$$

Since $\tilde{\mathbf{J}}(\omega) = \tilde{J}_z(\omega)\hat{z}$, where $\hat{z}$ is the unit vector in the direction of the z axis, the same must apply to the vector potential $\tilde{\mathbf{A}}$, i.e., $\tilde{\mathbf{A}}(\omega) = \tilde{A}_z(\omega)\hat{z}$. This can be seen from Equation (4.23b). Then, factorizing $\tilde{A}_z(\omega)$ into the z- and the ρ-dependent functions, we can write

$$\tilde{A}_z(\omega) = u(\rho)e^{-\frac{i\omega z}{v}}, \tag{4.26}$$

where the exponent $e^{-\frac{i\omega z}{v}}$ comes from the Fourier transform analogous to that of Equation (4.25). Plugging this into Equation (4.23b) and cancelling out the factors of $e^{-\frac{i\omega z}{v}}$, we obtain an equation for the radially dependent function only:

$$\frac{\partial^2 u}{\partial \rho^2} + \frac{1}{\rho}\frac{\partial u}{\partial \rho} + s^2 u(\rho) = \frac{e\mu_0}{4\pi^2}\frac{\delta(\rho)}{\rho}, \tag{4.27}$$

where we defined

$$s^2 \equiv \frac{\omega^2 n^2}{c^2} - \frac{\omega^2}{v^2} \overset{\beta = v/c}{=} \frac{\omega^2}{v^2}(\beta^2 n^2 - 1). \tag{4.28}$$

Let us consider Equation (4.27) in the region $\rho > 0$, i.e., far away from the pole at $\rho = 0$ which represents the radial coordinate of the muon. Here, the right-hand side vanishes and we obtain the Bessel equation,

$$\frac{\partial^2 u}{\partial \rho^2} + \frac{1}{\rho}\frac{\partial u}{\partial \rho} + s^2 u(\rho) = 0. \tag{4.29}$$

The general solution to Equation (4.29) is the linear superposition of the zeroth-order Hankel functions of the first and the second kind, i.e.,

$$u(\rho) = C_1 H_0^{(1)}(s\rho) + C_2 H_0^{(2)}(s\rho). \tag{4.30}$$

We are interested in the solution which represents an outgoing cylindrical wave at infinity and satisfies the inhomogeneous equation (4.27) at $\rho \to 0$. Such a solution is found when s is real, which translates to the requirement that $\beta^2 n^2 \geq 1$ (equivalently, $v \geq \frac{c}{n}$), and is equal to

$$u(\rho) = \frac{ie\mu_0}{8\pi}H_0^{(2)}(s\rho) \overset{s\rho \to \infty}{=} \frac{e\mu_0}{4\pi\sqrt{2\pi s\rho}}e^{-is\rho + \frac{3\pi i}{4}} + \mathcal{O}(\rho^{-3/2}) \tag{4.31}$$

for positive frequencies ω. Then, the magnetic potential A_z in the time domain becomes

$$A_z(t) = \frac{e\mu_0}{4\pi\sqrt{2\pi\rho}}\int_0^\infty \frac{e^{i\omega t - is\rho + \frac{3\pi i}{4}}}{\sqrt{s}}\mathrm{d}\omega. \tag{4.32}$$

Applying the same procedure to the scalar potential ϕ and ultimately taking the derivatives in Equation (4.16), we can find the electric and magnetic field components E_z, E_ρ, and H_φ. We reproduce the results from [188, 189] below:

$$E_z \propto \frac{1}{\sqrt{\rho}}\int_0^\infty \left(1 - \frac{1}{\beta^2 n^2(\omega)}\right)\frac{1}{\sqrt{s}}\cos\chi\,\omega\mathrm{d}\omega; \tag{4.33a}$$

$$E_\rho \propto \frac{1}{\sqrt{\rho}}\int_0^\infty \frac{\sqrt{\beta^2 n^2(\omega) - 1}}{\beta^2 n^2(\omega)\sqrt{s}}\cos\chi\,\omega\,\mathrm{d}\omega; \tag{4.33b}$$

$$H_\phi \propto \frac{1}{\sqrt{\rho}}\int_0^\infty \sqrt{s}\cos\chi\,\mathrm{d}\omega, \tag{4.33c}$$

where $\chi = \omega\left[t - \frac{z\cos\theta_C + \rho\sin\theta_C}{c/n}\right] + \frac{\pi}{4}$ and θ_C is defined in Equation (4.10). To find the total electromagnetic power P radiated due to the passage of the muon, we first compute the Poynting vector,

$$\mathbf{S} = \frac{1}{\mu_0}\mathbf{E}\times\mathbf{H}, \tag{4.34}$$

which represents the electromagnetic energy flow per unit area per unit time. Then, we integrate the Poynting vector over the surface $\mathcal{A}$ of a cylinder of radius ρ and length ℓ collinear with the direction of the muon:

$$P = \left|\int_{\mathcal{A}}\mathbf{S}\cdot\mathrm{d}\mathcal{A}\right| = \frac{2\pi\rho\ell}{\mu_0}\left|[\mathbf{E}\times\mathbf{H}]_\rho\right|, \tag{4.35}$$

where $[\mathbf{E}\times\mathbf{H}]_\rho$ is the radial component of the cross product of the $\mathbf{E}$ and $\mathbf{H}$ fields. In cylindrical coordinates and taking into account only the non-zero field components, it is found as

$$[\mathbf{E}\times\mathbf{H}]_\rho = -E_z H_\phi. \tag{4.36}$$

Taking this result and introducing dummy variables ω' and ω'' for integrals in Equations (4.33a) and (4.33c), respectively, we expand the expression for the radiated power:

$$P \propto \ell\int_0^\infty\int_0^\infty\left(1 - \frac{1}{\beta^2 n^2(\omega')}\right)\cos\left(\omega't^* + \frac{\pi}{4}\right)\cos\left(\omega''t^* + \frac{\pi}{4}\right)\omega'\mathrm{d}\omega'\mathrm{d}\omega'', \tag{4.37}$$

with $t^* = t - \frac{z\cos\theta_C + \rho\sin\theta_C}{c/n}$. The total radiated energy W is obtained by integrating Equation (4.37) over time. Since $\int_{-\infty}^\infty\cos\left(\omega't^* + \frac{\pi}{4}\right)\cos\left(\omega''t^* + \frac{\pi}{4}\right)\mathrm{d}t = \pi\delta(\omega' - \omega'')$, we get

$$W \propto \ell\int_0^\infty\left(1 - \frac{1}{\beta^2 n^2(\omega)}\right)\omega\mathrm{d}\omega. \tag{4.38}$$

The energy radiated per unit length per unit frequency is then simply

$$\frac{\partial^2 W}{\partial\omega\partial\ell} \propto \omega\left(1 - \frac{1}{\beta^2 n^2(\omega)}\right). \tag{4.39}$$

Reinserting the omitted constants and rewriting Equation (4.39) in terms of the number of photons N_γ radiated per unit photon wavelength λ, we finally arrive at the **Frank-Tamm formula for the Cherenkov radiation spectrum** [189]:

$$\frac{\partial^2 N_\gamma}{\partial\lambda\partial\ell} = \frac{2\pi\alpha}{\lambda^2}\left(1 - \frac{1}{\beta^2 n^2(\lambda)}\right), \tag{4.40}$$

where $\alpha \approx \frac{1}{137}$ is the fine-structure constant. Importantly, the Frank-Tamm formula is valid only for photon wavelengths and charged particle velocities where the condition $\beta n(\lambda) > 1$ is satisfied. This implies that as soon as the $n(\lambda)$ approaches 1, the number of radiated Cherenkov photons must drop to 0, as $\beta > 1$ can never be satisfied (no particle can travel faster than the speed of light in vacuum). This is illustrated in Figure 4.9, where we overlay the Cherenkov spectrum computed from Equation (4.40) with the wavelength dependence of the refractive index of ice.

We can see that in the 400–700 nm range, where n is approximately constant, the spectrum follows the characteristic $\frac{1}{\lambda^2}$ dependence. However, when n begins to rapidly decrease (at $\lambda \lesssim 100$ nm), so does the Cherenkov spectrum, and there is virtually no Cherenkov radiation below $\sim$70 nm. However, despite the amplitude of the spectrum
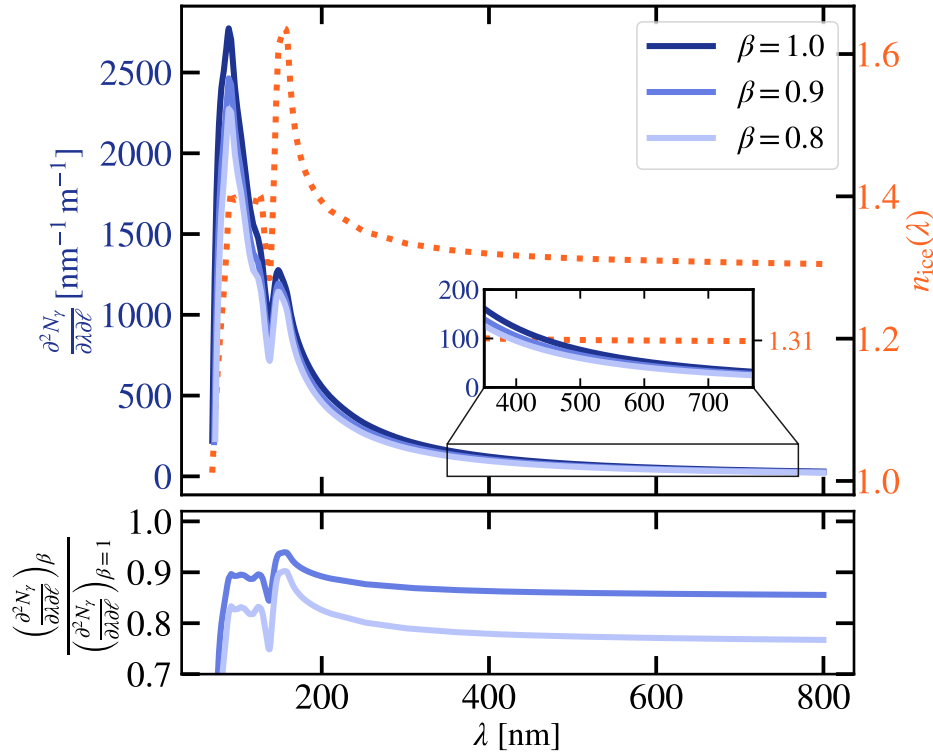
FIGURE 4.9 – *Top panel*, *right axis* (solid blue lines): The wavelength spectrum of the Cherenkov radiation in ice, shown for three different values of the charged particle Lorentz velocity factor $\beta = v/c$. *Top panel*, *left axis* (dotted orange line): The refractive index of ice as a function of wavelength, as extracted from [192]. *Bottom panel*: the ratio of the Cherenkov spectrum at $\beta \in \{0.8, 0.9\}$ to that at $\beta = 1$.

increasing almost monotonically all the way down to ~ 80 nm wavelength, Cherenkov light appears blue in ice (400 nm $\leq \lambda \leq 450$ nm). This originates from the fact that the wavelength dependence of the ice absorption coefficient has a minimum at 400 nm, and light in the ultraviolet range is subject to ~ 5–10 times stronger absorption [193–195].

Finally, we comment on the dependence of the Cherenkov light yield on the charged particle velocity. As intuitively expected, we see from Figure 4.9 that the particles which travel at the $\beta = 0.9$ ($\beta = 0.8$) make about 15% (25%) less light than those propagating at the speed of light ($\beta = 1$). This highlights the importance of accounting for the β dependence of the Cherenkov spectrum when modelling the light emission by neutrino interaction products, which we will come back to in Section 10.1 in the context of IceCube Monte Carlo simulation.

III

IceCube detector

The basics

5.1 operational principle and science goals

The IceCube Neutrino Observatory is a neutrino detector situated at the Amundsen-Scott South Pole Station in Antarctica. The thick ice sheet of the Antarctic Plateau, which elevates the station nearly 3 km above sea level, provides a suitably large volume and an excellent optical medium for detection of Cherenkov photons radiated by the products of neutrino interactions in the ice. The IceCube detector takes up nearly one cubic kilometer of this ice volume, which is instrumented with 5160 digital optical modules (DOMs) sensitive to Cherenkov radiation.

The full detector, with 86 strings housing 60 DOMs each (see Figure 5.1), was commissioned by 2011 and has been taking data continuously since. The layout of the detector was optimized for the detection of high-energy (TeV and above) astrophysical neutrinos, with the goal of discovering and understanding the behavior of neutrino sources within and beyond our galaxy. However, the inner sub-array of IceCube (called IceCube-DeepCore and described in detail in Section 5.3) is instrumented on average twice as densely as the rest of the detector volume, lowering the detector threshold to ~5–10 GeV. This enables the studies of standard [180, 196, 197] and non-standard [46, 103, 198, 199] atmospheric neutrino oscillations, detection of low-energy astrophysical transients [200, 201], and searches for GeV-scale dark matter annihilation or decay signals [202–204].

The spatiotemporal patterns of light deposition in the IceCube detector provide information necessary to reconstruct the basic neutrino properties such as its energy, arrival direction, and flavor. At 1–100 GeV energies relevant for atmospheric neutrino oscillation measurements, the reconstruction of the latter is generally constrained to two classes. The first one is ν_μ-CC events, which produce a muon with a characteristic *track* pattern of Cherenkov light deposition, as shown in the top left panel of Figure 5.2. The second class encompasses all other events (ν_e-CC, ν_τ-CC, and all-flavor NC), which deposit light in spherical topologies known as *cascades* (shown in the top right panel of Figure 5.2). At low (< 10 GeV) energies, the track and cascade topologies become practically indistinguishable, as seen from the bottom row of the same figure. This example clearly shows that the average accuracy of low-energy event classification in IceCube is significantly hindered by the sparsity of the detector instrumentation. This presents an obvious experimental limitation for the flavor-sensitive measurements,
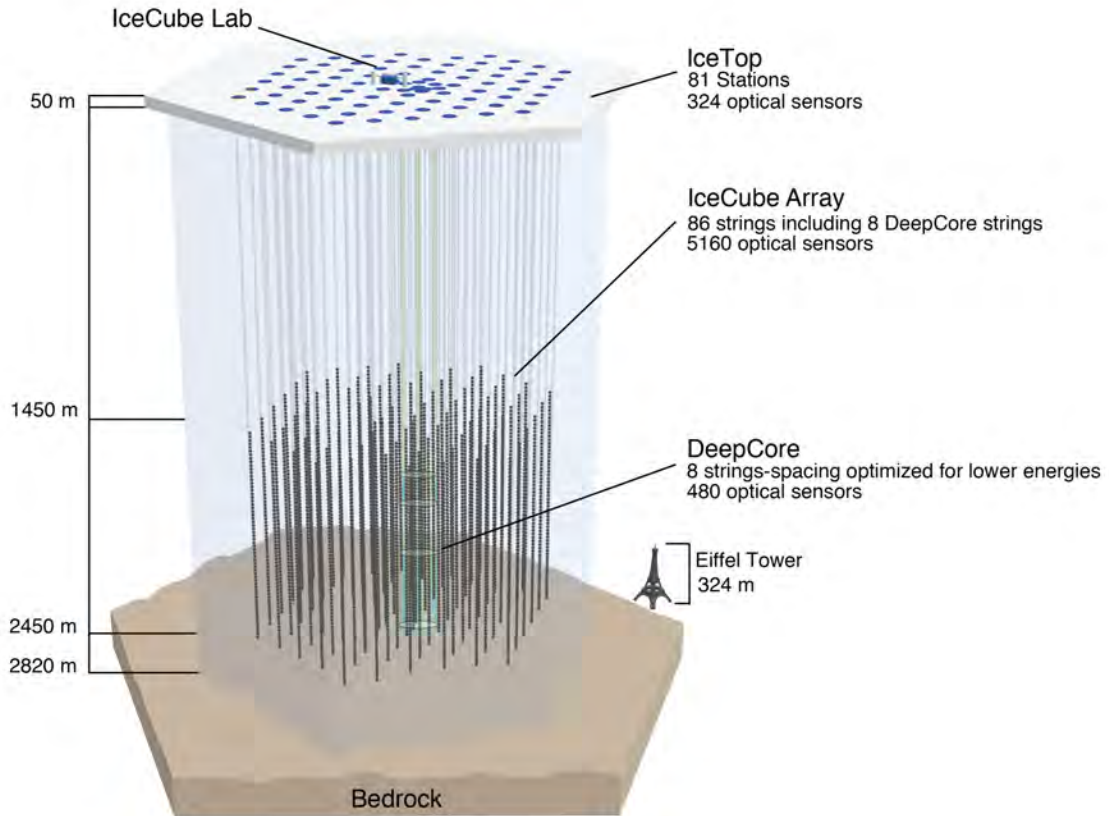
FIGURE 5.1 – The side view of the IceCube detector, including the main in-ice array (1450 m..2450 m below the
IceTop de·

True tracks (ν_μ-CC) **True cascades (ν_e-CC)**

$E_\nu \simeq 70\,\mathrm{GeV}$

$E_\nu \simeq 9\,\mathrm{GeV}$
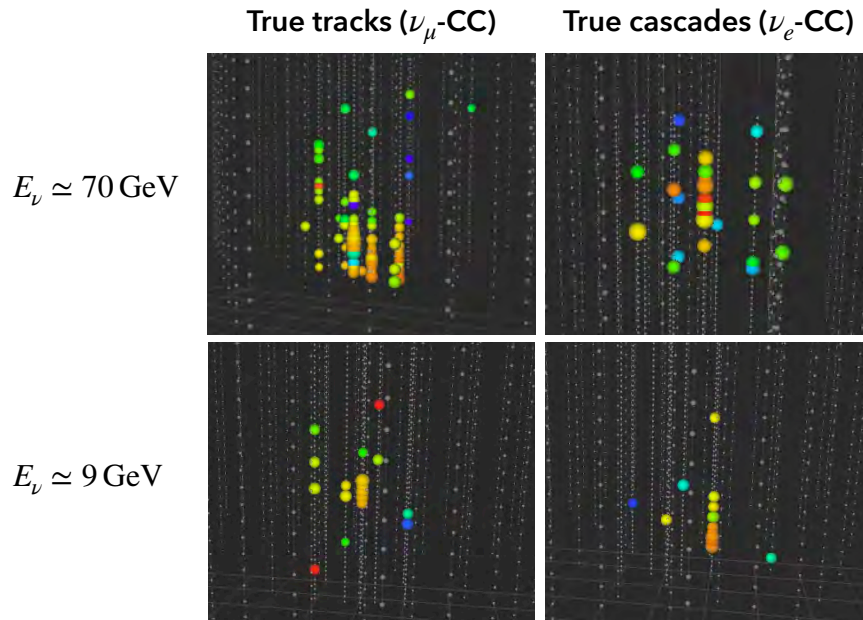
FIGURE 5.2 – Two types of event topologies in IceCube at "high" ($\sim 70\,\mathrm{GeV}$) *vs.* "low" ($\sim 9\,\mathrm{GeV}$) neutrino energies. The events are extracted from the IceCube Monte Carlo simulation and visualized with the STEAMSHOVEL IceCube software. The color-filled spheres represent the DOMs that detected either Cherenkov light signal or background noise in a given event, and the exact color represents the timestamp of the detection (red color corresponds to earlier times, and blue – to later times).

such as those of neutrino oscillation parameters[1]. The upcoming IceCube-Upgrade detector, with 7 additional strings scheduled for deployment within the DeepCore fiducial volume during the polar season 2025–2026 [205], will address this issue by increasing the density of instrumentation and introducing new types of optical modules with better directional sensitivity. This will both lower the detection energy threshold down to ~1 GeV and, as shown in Section 5.4, significantly improve the reconstruction of the lowest energy neutrino events currently detectable with DeepCore.

5.2 FUNCTIONAL COMPONENTS OF THE ICECUBE DETECTOR

5.2.1 *Digital optical module (DOM)*

A DOM, shown schematically in Figure 5.3, is the fundamental light detection unit of IceCube. Inside a glass sphere of a ~30 cm (12”) diameter designed to protect the inner electronics from extreme pressures in deep ice, it contains a photomultiplier tube (PMT) and several circuit boards. The inner workings of a PMT, which converts the incident light to an electric signal via the photoelectric effect [206], are described in Section 5.2.2. The circuit boards include, first and foremost, the Main Board, which
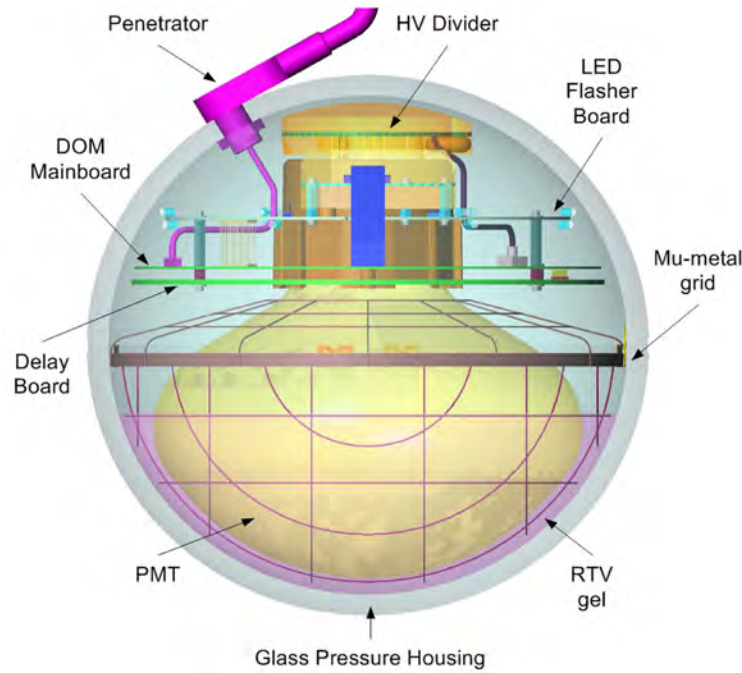

FIGURE 5.3 – Schematic visualization of the IceCube DOM. The main components are the downward-facing photomultiplier tube (PMT), the main board housing the electronics for PMT pulse digitization, the LED flasher board used for calibration studies, and the glass pressure sphere encapsulating all of the electronics. *Figure courtesy of the IceCube Collaboration.*

is responsible for managing all internal DOM devices such as the high-voltage (HV) PMT power supply and a range of sensors used for monitoring purposes. Apart from overseeing the function of these components, the Main Board digitizes and processes the PMT waveform data and handles communication with the surface data acquisition (DAQ) system. Additionally, it exchanges information with the neighboring DOMs with the goal of identifying signals in a *local coincidence* with each other. The LED Flasher Board, which is responsible for generating light signals for in-situ calibration studies, is also controlled by the Main Board.

---

[1] Note that neutrino energy and direction reconstruction resolutions are likewise affected by the detector sparsity and decrease rapidly at $E_\nu < 10\,\text{GeV}$; see, for example, Figure 5.8 for typical resolutions.

5.2.2 *Photomultiplier tube (PMT)*

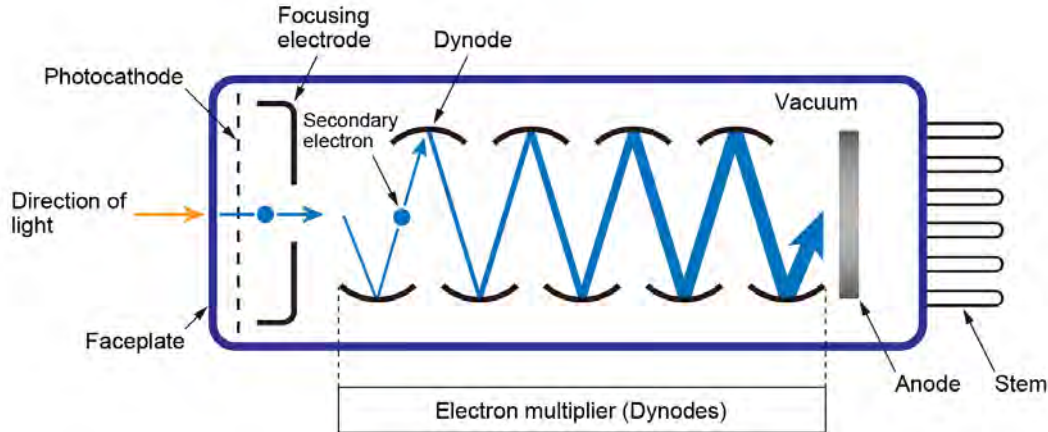A typical layout of a photomultiplier tube is shown in Figure 5.4. The photons incident

FIGURE 5.4 – Schematic diagram of a photomultipler tube. *Image courtesy of Matsusada Precision Inc.* [207]

onto the PMT hit a metal plate called a photocathode, from which they eject photo-electrons (PEs). These electrons move in the electric fields established by the sequence of dynodes on their way, such that each subsequent dynode is held at a more positive potential than the previous one [208]. As the primary electrons gain energy due to acceleration in the electric fields, more and more secondary electrons are progressively ejected from each dynode, resulting in an amplification of the primary electric signal. The standard gain of the IceCube PMTs is 10^7 [41], which represents the ratio of the final output current at the anode to the initial current at the photocathode [208]. This current is then read out and digitized as described in Section 6.1.1.

The PMTs of IceCube are approximately 25 cm (10") in diameter and are oriented downwards with respect to the ice surface, which increases the detector's sensitivity to upward-going (Earth-crossing) neutrinos. The other two important characteristics of a PMT are its angular acceptance and quantum efficiency. The angular acceptance refers to the relative optical efficiency of the photon conversion to the electrical signal as a function of angle η between the incident photon direction and the PMT rotation axis. The angular response of an entire IceCube optical module, as measured in a laboratory setting, is shown in Figure 5.5. One a module is submerged in the melted ice during installation and the ice is refrozen, the angular response changes due to the formation of bubble columns in the drill holes. The PMT quantum efficiency (QE), on the other hand, is defined as the absolute efficiency of a single photon-to-photoelectron conversion as a function of wavelength. The QE of bare PMTs is approximately 25% for 390 nm light incident on IceCube's "normal quantum efficiency" (NQE) DOMs and 34% for "high quantum efficiency" (HQE) DOMs [41]. The wavelength dependence of a typical NQE DOM photon acceptance is given in Figure 5.6, which combines the effects of the PMT QE and the photon transmission through the OM glass and gel.

5.3 ICECUBE-DEEPCORE: LAYOUT AND PERFORMANCE

The fiducial volume of the IceCube-DeepCore detector encompasses the central string 36, the six surrounding strings of the main IceCube array (strings 26–27, 35, 37, and 45–46), and the eight special DeepCore strings (79–86), as illustrated in Figure 5.7. The first aspect in which the eight DeepCore ("DC") strings differ from the rest of IceCube ("IC") is their geometrical layout, with both a denser interstring spacing ($\Delta\rho_{DC} = 41\,\text{m}..105\,\text{m}$,
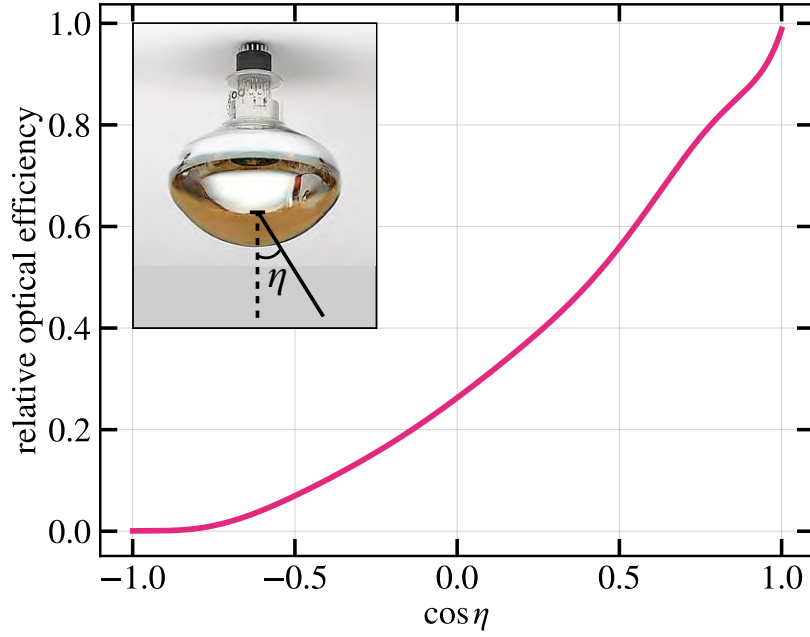
FIGURE 5.5 – Angular response of an IceCube OM as measured in a laboratory setting and extracted from the IceCube Collaboration internal software. The shown PMT is the R7081 Hamamatsu model (*credit: Hamamatsu Photonics*).
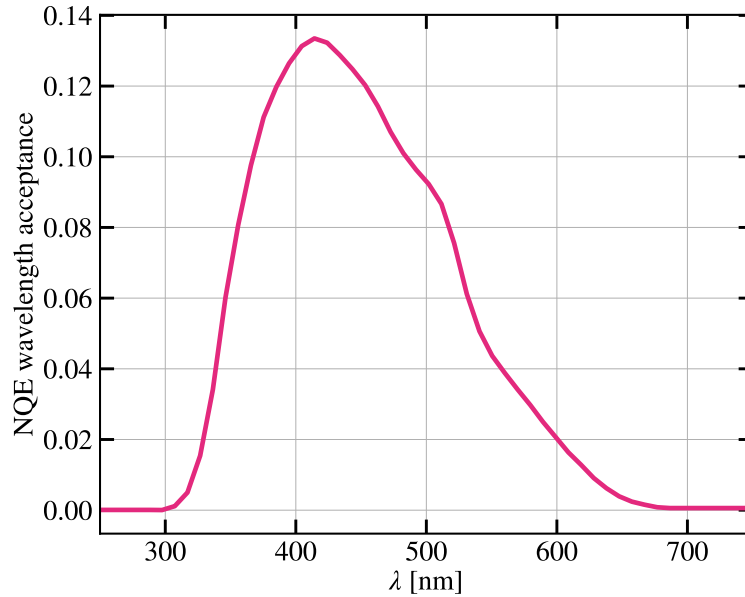
FIGURE 5.6 – Wavelength acceptance of an NQE IceCube DOM [209].

compared to $\Delta\rho_{\mathrm{IC}} = 125\,\mathrm{m}$) and a denser vertical spacing Δz of the optical modules on each string [41]. While $\Delta z_{\mathrm{IC}} = 17\,\mathrm{m}$ for the IceCube strings, the vertical spacing of the DeepCore DOMs is either $\Delta z_{\mathrm{DC}} = 7\,\mathrm{m}$ (for the 50 DOMs per string within the main DC volume located at depths between $2100\,\mathrm{m}..2450\,\mathrm{m}$) or $\Delta z_{\mathrm{DC}} = 10\,\mathrm{m}$ (for the 10 DOMs per string within the veto cap region[2] located above the $2000\,\mathrm{m}$ depth). The DeepCore strings have no optical modules at $\sim 2000\,\mathrm{m}..2100\,\mathrm{m}$ depths, i.e., the region known as the *dust layer*. This ice layer has a high concentration of dust particles, which increases absorption and scattering as shown in Figure 5.7 and thereby hinders the detection of light deposited in this region.

---

[2]The hits in this veto region are used for atmospheric muon background rejection.
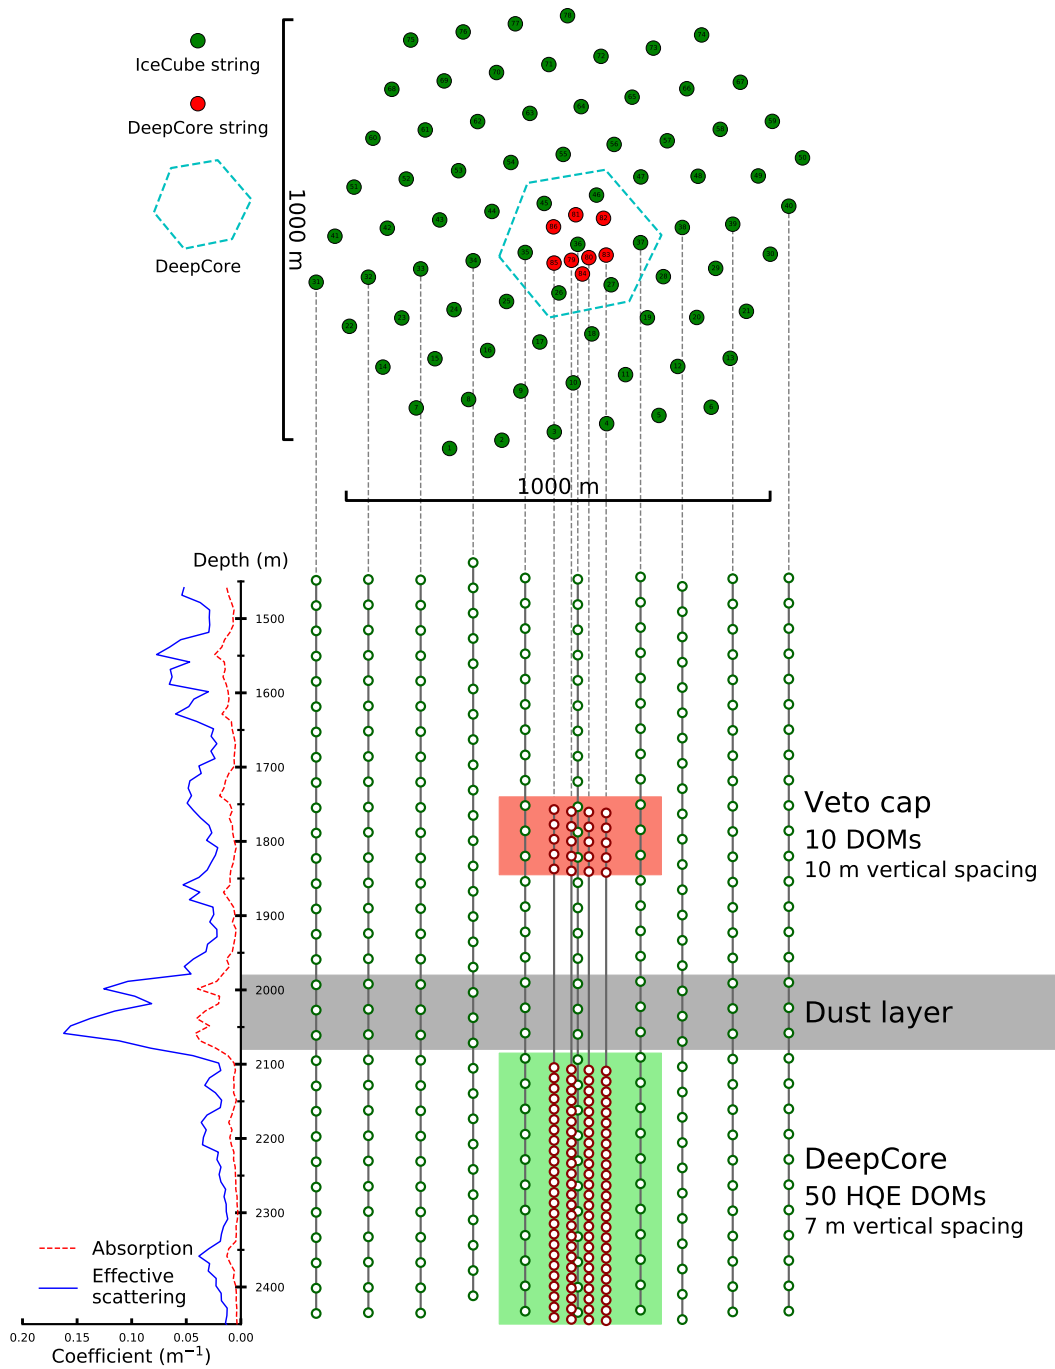
FIGURE 5.7 – Top and side view of the geometrical layout of the main IceCube array and the DeepCore subarray. *Figure extracted from* [210].

The second main distinction between the DeepCore strings and those of the rest of IceCube is the higher quantum efficiency of most of the optical modules. While most strings of the main IceCube array have only NQE DOMs[3], six of the DeepCore strings (81–86) are fully populated with HQE DOMs, and the other two strings (79–80) are populated with a mix of HQE and NQE modules.

---

[3]This applies to strings 1–78, excluding string 36 with 14 HQE DOMs and string 43 with one HQE DOM.

With the help of state-of-the-art machine learning algorithms such as graph neural networks (GNNs), 10 GeV neutrinos in IceCube-DeepCore can be reconstructed with a ~50% energy resolution and a 20°-25° zenith angle resolution. As more light is deposited by the products of higher-energy neutrino interactions, the resolutions improve monotonically with increasing energy. This is shown in Figure 5.8, where the performance of the DYNEDGE GNN [210] on the DeepCore neutrino events is compared with that of the likelihood-based RETRO reconstruction algorithm [211]. Both reconstruction methods are described in detail in Section 6.4.
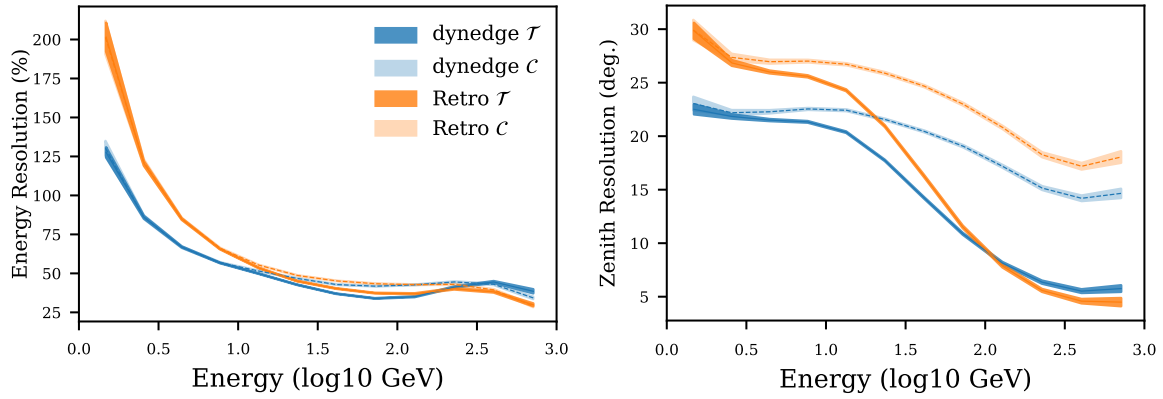
FIGURE 5.8 – Reconstruction resolutions of the DeepCore track ($\mathcal{T}$) and cascade ($\mathcal{C}$) events, compared between the DYNEDGE GNN reconstruction and the RETRO likelihood-based reconstruction *Figure courtesy of R. Ørsøe* (adapted from [210]).

5.4 ICECUBE-UPGRADE

The deployment of the IceCube-Upgrade ("ICU") will introduce 7 strings in the middle of DeepCore with new types optical modules. As shown in Figure 5.9, these will include mDOMs (multi-PMT modules with 24 PMTs each, pointing in all directions around the sphere) and D-Eggs (two-PMT modules with one PMT facing down, and one facing up). With ~700 optical modules planned for the IceCube-Upgrade [205], the total number of

FIGURE 5.9 – New types of optical modules introduced in the IceCube-Upgrade: D-Eggs with 2 PMTs (left) and mDOMs with 24 PMTs (right). *Figure courtesy of the IceCube Collaboration.*

PMT channels with respect to the current IceCube configuration will be increased nearly threefold [212]. These modules will be densely spaced, with the vertical intermodule distance $\Delta z_{ICU} = 3\,m$ and the radial interstring distance $\Delta\rho_{ICU} = 7\,m$, and therefore detect a lot more light than would be possible with the IceCube-DeepCore alone from

identical events. This is illustrated in Figure 5.10, where the photons from the same simulated ν_τ-CC event are propagated to and detected by the modules of the current IceCube-DeepCore configuration as well as that augmented with the Upgrade strings.

DeepCore only   **DeepCore + Upgrade**
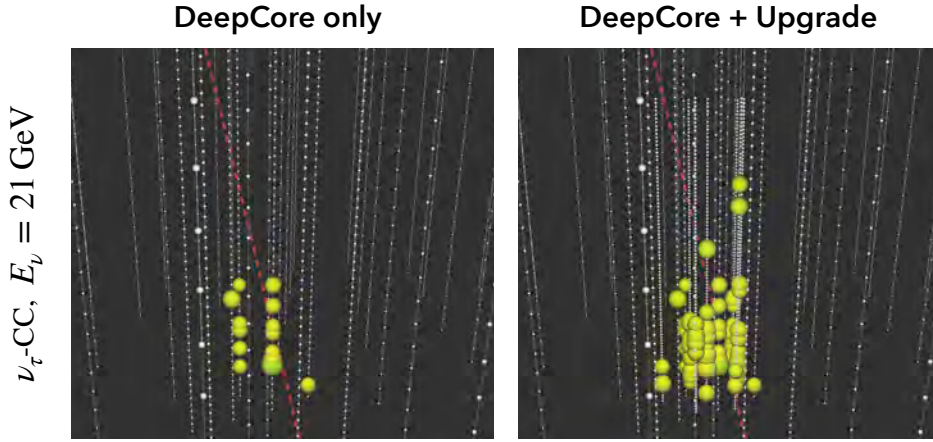
FIGURE 5.10 – *Left panel*: Simulated 21 GeV tau neutrino interaction in the current IceCube-DeepCore detector. *Right panel*: the same event simulated with 7 additional IceCube-Upgrade strings (2019 geometry). *Image courtesy of T. Stuttard.*

The projections of the IceCube-Upgrade zenith resolutions as available at the time of the public IceCube-Upgrade Monte Carlo simulation release [213] (dated 2020) are shown in Figure 5.11. These projections are conservative, as they do not utilize
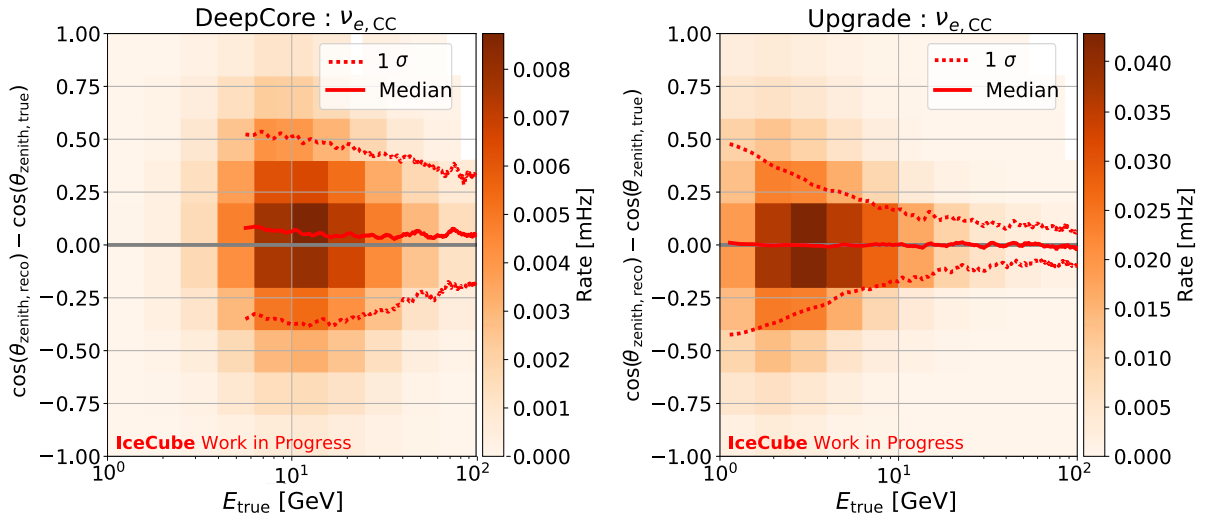
FIGURE 5.11 – Projected IceCube-Upgrade zenith angle resolutions (right) as compared to those of IceCube-DeepCore (left) for ν_e-CC events. *Figure extracted from* [205].

the most recent progress in the IceCube-Upgrade Monte Carlo simulation [212] and rely on a parameterization of an outdated likelihood-based cascade reconstruction[4] [214]. However, they preliminarily indicate at least a factor 3 improvement in the reconstruction of $\cos\theta_{\mathrm{zen}}$, which highlights the importance of the IceCube-Upgrade for physics analyses with GeV-scale neutrinos. Furthermore, these estimated resolutions provide a baseline for the machine learning-based reconstructions such as DYNEDGE currently in development within the IceCube Collaboration for use in the IceCube-Upgrade.

---

[4]Within IceCube, this is known as the Monopod reconstruction.

The making of an IceCube atmospheric neutrino event

We open this chapter by mapping out the path from the Cherenkov photon detection by individual PMTs to the event *triggering* in IceCube-DeepCore and the first stage of event selection known as *filtering* (covered in Section 6.1.1). Further, in Section 6.1.2 we describe how events that triggered DeepCore and passed the DeepCore filter are prepared for reconstruction through *noise cleaning*. These are the initial processing stages that are equally applicable to data and Monte Carlo events. While we already explained how events come to life in data by describing atmospheric neutrino fluxes and interactions (in Chapter 3 and Section 4.1, respectively), in Section 6.2 we elaborate on the Monte Carlo simulation chain that aims to mimic atmospheric neutrino events and sources of background as realistically as possible. Thereafter, Section 6.3 explains the general procedure and provide the historical examples of event selection in IceCube-DeepCore, followed by an overview of event reconstruction methods relevant for this thesis in Section 6.4.

6.1 DATA READOUT AND INITIAL STAGES OF PROCESSING

6.1.1 *Triggering and filtering*

When a single photon hits the photocathode of a PMT, the ultimate current output at the anode is not deterministic but instead is subject to statistical fluctuations of the PMT response. These fluctuations are driven by the variation in the number of electrons ejected from each PMT dynode as well as the uncertainty in the photoelectron trajectories [215]. Therefore, if a single-photon experiment is repeated multiple times, the charge collected at the anode represents a broad distribution rather than a single value. This is called a **Single Photoelectron (SPE) distribution**, whose charge values are scaled such that the distribution peak is centered at 1 PE. In IceCube, an effort is made to measure this distribution for each DOM, as illustrated in Figure 6.1 for DOM #1 on string #1. This distribution is fitted with a template whose functional form combines the Gaussian distribution centered around 1 PE[1] and two exponential distributions. The latter are introduced to fit the low-charge peak of the SPE distribution (i.e., the rapid increase in the probability density of charge $q_{\text{SPE}} \to 0$), which originates

---

[1] As seen from Figure 6.1, the mean μ of the Gaussian distribution is left free in the fit and is generally expected to fit close to 1. The rescaling of the entire distribution by the fitted value of the mean is applied in both data and MC.
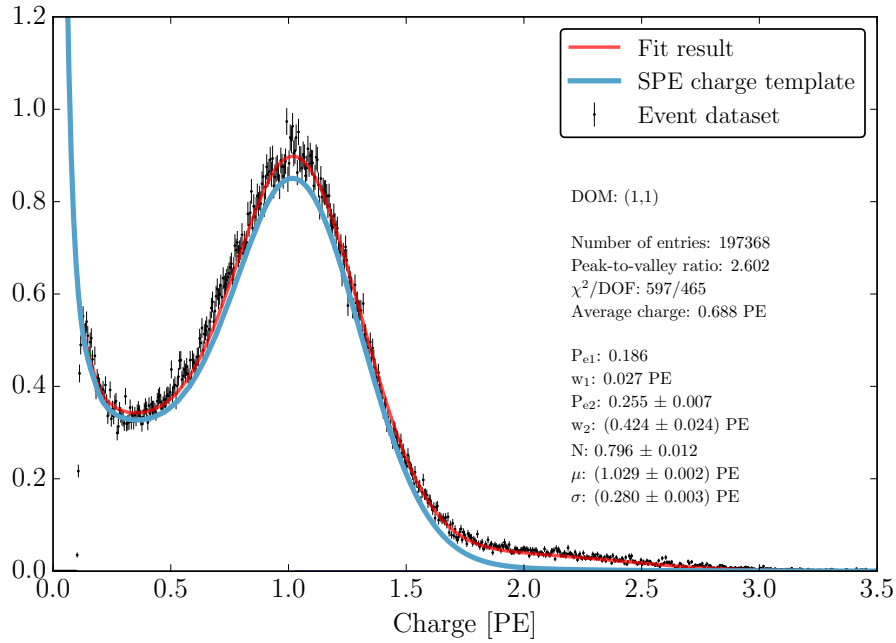
FIGURE 6.1 – Charge distribution expected from a single photon detected by an IceCube DOM, here shown specifically for DOM #1 on string #1. The datapoints represent all the charge measurements for this DOM in the 2011–2016 IceCube seasons. *Figure extracted from* [215].

from dark noise (e.g., spontaneous PE emission by the photocathode and the dynodes) and electronic noise. One of these exponentials is fitted by combining data from all IceCube DOMs, while the other one is specific to each DOM. Finally, the SPE template is convolved with itself to derive a high-charge tail of the SPE distribution expected from multi-photoelectron contamination. Only the 2 PE contribution is considered in practice, which represents the scenario when two photoelectrons are separated only by a few nanoseconds in their arrival time at the first dynode. The original SPE template, plus its convolution with itself, thereby form the basis of a *convolution fit* [215]. The fit parameters for each DOM are stored in the *Geometry, Calibration, and Detector (GCD) file*, which is used for IceCube Monte Carlo simulation production. Specifically, the charge of a single reconstructed photoelectron in Monte Carlo (shortly MCPE) is sampled from the SPE distribution with the parameters stored in the GCD file for a given DOM.

Once the charge is read out in data or sampled from the SPE template in Monte Carlo, the decision is made about whether or not further information recording by the DOM will be initiated. The quantitative threshold for this decision is called the **discriminator threshold** and is equal to 0.25 PE (after charge recalibration such that the peak of the SPE distribution is located at 1). Should the charge q on a given DOM be above the discriminator threshold, the DOM will record the PMT waveform with the time resolution dependent on whether this signal is observed in local coincidence (LC) with other DOMs. Specifically, if the $q > 0.25$ PE condition is met by a nearest or a next-to-nearest neighbor of a given DOM on the same string within $\Delta t = 1\,\mu s$, the **hard local coincidence (HLC)** criterion is satisfied. In this case, the PMT waveform is digitized with both an **Analog Transient Waveform Digitizer (ATWD)**, which records the charge information over 128 time samples of 3.336 ns duration each (covering the timespan of 427 ns) and **Fast Analog-to-Digital Converter (FADC)** (recording 256 time samples with 25 ns resolution over a longer time window of 6400 ns). If the HLC criteria are not satisfied, the DOM is said to have a **soft local coincidence (SLC)** hit. The SLC waveforms are digitized only through the FADC channel, and only three

time-charge stamps are stored (the bin with the highest charge amplitude and its two nearest temporal neighbors).

Whether an event is kept for further processing through a given event selection pipeline (i.e., passed onto the Level 2 of event selection) is decided by one or more event triggers, which form the conditions for a specific *filter*. For all IceCube-DeepCore event selections, passing the DeepCore filter is necessary for an event to arrive at Level 2. The DC filter consists of only one trigger condition, namely that of the Single Multiplicity Trigger (SMT) requiring 3 HLC hits recorded within $\Delta t = 2.5\,\mu s$ (thus referred to as the SMT3 trigger). Additionally, a condition on the effective particle speed derived from the center of gravity of the event hits is applied to remove the majority of the muon background events at this stage [4, 45].

6.1.2 *Pulse reconstruction and cleaning*

The PMT waveform digitized through the FADC channel (in case of an SLC hit) or the combination of FADC and ATWD channels (in case of an HLC hit) is further passed through a deconvolution algorithm known as Wavedeform. This algorithm uses the non-negative least squares method [216] to reconstruct the individual pulses (including their times and number of underlying PEs) whose convolutions with the SPE templates best reproduce the recorded waveform. The time sequence of these pulses forms the pulse series of an event, which may include both signal pulses (originating from Cherenkov photons radiated by neutrino interaction products) and noise pulses (originating from e.g. decays of radioactive isotopes in the OM glass, PMT effects such as late, pre- or after-pulsing (for details, see [215]), or dark noise mentioned in Section 6.1.1). To maximize the sensitivity of a given event selection to the physics process of interest, the goal of an effective cleaning algorithm is to remove as much noise pulses from the events as possible while keeping most of the interesting signal pulses.

The standard algorithm used in IceCube for noise cleaning is the **Seeded Radius-Time (SRT)** algorithm, which identifies DOM hits in spatial and temporal coincidence with one another to arrive at a population of pulses which could have originated from a single physics event. The procedure followed by the SRT cleaning is schematically visualized in Figure 6.2.
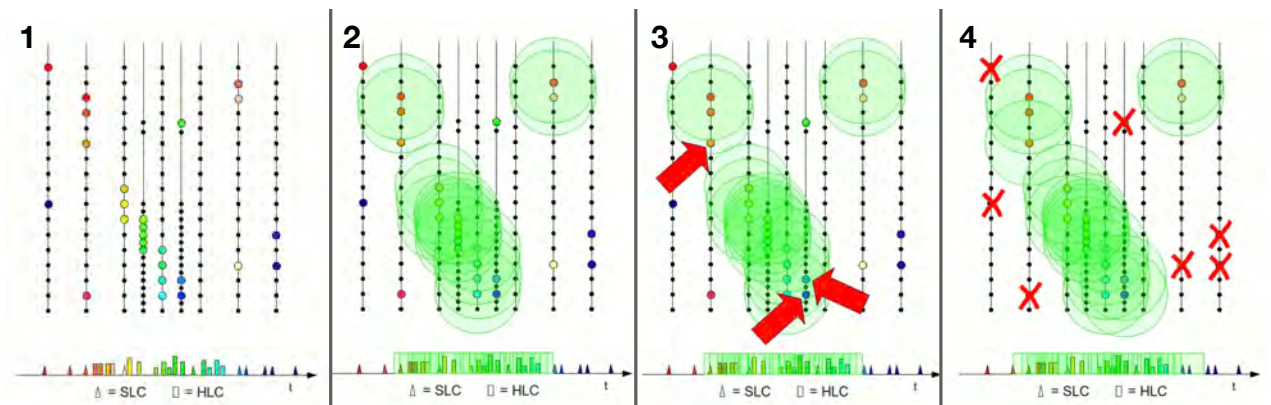
FIGURE 6.2 – Seeded Radius-Time (SRT) noise cleaning algorithm. *Figure courtesy of the IceCube Collaboration (contributed by O. Schulz).*

First, all HLC hits (i.e., DOMs found in hard local coincidence with one another) are identified in a given event, such as that in panel #1 of Figure 6.2. These HLC hits are used as the *seeds* to construct the cleaned pulse series; they are not removed by the SRT algorithm. For each of the HLC-triggered DOMs, SRT then searches for

any SLC-triggered DOMs within the radius r which were hit within the time t, as illustrated in panel #2. Any SLC-hit DOMs falling within this spatiotemporal window are added to the cleaned pulse series (as shown in panel #3), and the algorithm moves to the next HLC seed. The whole procedure is repeated ≤ 3 times until no more SLC DOMs can be added. All remaining SLC hits that did not make it into any of the RT windows are removed, which is illustrated in panel #4. The specific cleaning settings used in the IceCube-DeepCore event selections designed for oscillation analyses are $r = 150\,\text{m}$, $t = 1\,\mu\text{s}$. Finally, a static **time window (TW)** filter is applied to remove any pulses recorded outside of the [8800 ns, 14 000 ns] window. The resulting cleaning algorithm is therefore referred to as `SRTTW`, and the ultimate collection of cleaned pulses is called the `SRTTWOfflinePulsesDC` pulse series. This pulse series, with the exact settings specified above, was used as the input to event reconstructions at the time this work began. In Sections 10.3.1 and 10.3.2, we highlight significant issues in the event reconstruction driven by the remnant noise in this pulse series and describe a solution that effectively mitigates them.

In recent years, machine learning-based algorithms have likewise gained attention in IceCube for noise cleaning applications. In particular, the DYNEDGE GNN [210] has been successfully applied to noise rejection in the IceCube-Upgrade simulation [212] and is set to become the baseline noise cleaning algorithm for ICU. The GNN-based noise tagging is particularly advantageous for multi-PMT modules such as mDOMs and D-Eggs, where the SRT cleaning algorithm would automatically identify all PMTs on a module as being in spatiotemporal coincidence with one another.

6.2 MONTE CARLO SIMULATION

The Monte Carlo simulation of the expected neutrino events and the experimental backgrounds (such as atmospheric muons and noise) is the foundation of any physics analyses in IceCube. It produces a viable model of the neutrino signal under an assumed set of parameters and allows one to measure those parameters by reweighting the MC until a good match with data is achieved. In Section 6.2.1, we describe neutrino and muon simulation production in IceCube-DeepCore. In Section 6.2.2, we further summarize the Monte Carlo simulation of noise, which is injected into all physics MC events.

6.2.1 *General simulation chain: Neutrinos and muons in DeepCore*

Prior to any event selections, the MC simulation of neutrinos and muons is conducted in three steps, which we detail below.

- **Step 1: Primary particle generation**

    At this step, neutrino events are simulated using the GENIE event generator[2] [158] according to a requested geometrical distribution and energy spectrum. In the simulation used in this thesis, the energy spectrum at generation is chosen to be a power law ($\propto E_\nu^{-2}$), according to which the primary neutrino energies are sampled within 4-5 discrete energy bands covering the energy range from 1 GeV for 10 TeV. The geometrical distributions of neutrino interaction vertices are represented by cylindrical volumes whose radii are tuned per energy band. This is done for simulation efficiency purposes, such that the lowest-energy neutrinos ($E_\nu \sim \mathcal{O}(1\text{–}10\,\text{GeV})$) are generated only

---

[2]GENIE is used for almost all neutrino MC in this thesis. In the three-year GRECO event selection, the NuGen generator was used for neutrinos with true energies above 100 GeV [4, 217]. The version of GENIE used in the GRECO (oscNext) event selection was 2.8.6 (2.12.8).

within DeepCore, and the highest-energy neutrinos ($E_\nu \gtrsim 100\,\text{GeV}$) are generated both inside and outside of the DeepCore fiducial volume. GENIE then simulates neutrino interactions with the nuclei in the ice (i.e., hydrogen and oxygen) and handles the kinematics of the final-state particle production according to a specified cross section model. In particular, the deep inelastic scattering interactions are simulated according to the leading-order cross section calculated with the Bodek-Yang corrected GRV98 parton distribution functions [3][177, 178]. Any final-state muons (e.g. the secondaries of the ν_μ-CC interactions) are propagated until stopping in the PROPOSAL software [218], which takes into account ionization and radiative muon energy losses. At this stage, the stopping range ("track length") of such secondary muons is recorded, which is used for calculating their Cherenkov light yield at Step 2.

The atmospheric background muons are simulated in the native IceCube software called MUONGUN [4, 219], where muons are directed from the surface of a cylinder volume enveloping the entire IceCube towards the surface of the DeepCore fiducial volume. Their energy and angular distributions (including both azimuth and zenith) at generation are sampled such that these muons are likely to reach the final levels of the event selection, thereby increasing the simulation efficiency[4]. For the GRECO sample ("Sample A" of the public three-year IceCube-DeepCore data release [3, 46], which is used in Chapters 7 and 8 of this thesis), the target muon distribution was that of the final level (Level 7) of the GRECO event selection. For the oscNext sample described in Chapters 9 and 10, the target distribution was that of the muons reaching Level 5. Prior to the described biased sampling, the true muon energies cover the 160–500 GeV (150 GeV–5 TeV) in GRECO (oscNext). The generated muon events can then be weighted according to any preferred cosmic ray flux model, which was chosen to be H4A [220] for both of the event selections. The MUONGUN muons are likewise propagated with PROPOSAL at Step 1.

No other particles but muons are propagated at this stage. All other secondaries (such as electrons, positrons, photons, and hadrons) are stored in the particle interaction tree[5] and passed onto Step 2 of the simulation.

- **Step 2: Propagation of remaining secondaries; Photon generation and propagation**

Photon generation and propagation in the IceCube Monte Carlo simulation are handled via the GPU-accelerated CLSIM software [222]. The treatment of secondary particles and their Cherenkov light yield in CLSIM depends on the particle type and energy, which we elaborate on below.

  ▷ **Case 1: The secondary particle is a muon**

  In this case, the muon is already propagated in PROPOSAL at Step 1, and the only remaining step is to generate Cherenkov photons and propagate them towards the DOMs. This is done in the PPC submodule of CLSIM [223], which takes the muon track segments returned by PROPOSAL and generates a number of Cherenkov photons emitted by each segment according the Frank-Tamm formula. Specifically, the number of photons expected per unit length as per Equation (4.40) is multiplied by the length of each segment extracted from PROPOSAL. Poisson or Gaussian fluctuctuations applied on top of this expected number of photons N_γ, depending on its magnitude (if $N_\gamma > 10^7$, the Gaussian approximation is used). An important

---

[3]In the oscNext analysis, there is an option to reweight these events to the CSMS cross section [179], which is more accurate for neutrinos with energies above 100 GeV (see Section 4.1.3).

[4]This method was developed in [4] and is referred to as the kernel density estimator (KDE) prescale.

[5]Except for the muon secondaries with energies below 500 MeV, which are not stored in the tree. Their light yield is directly parameterized via a contribution to the effective Cherenkov-radiating track length of the muon as described in [221].

simplification made by the ppc code is that the muon Lorentz velocity factor β is assumed to be equal to 1 in the Frank-Tamm formula throughout the entire muon trajectory, i.e., the muon energy losses are disregarded in the context of Cherenkov light yield calculation. In Section 10.1, we show that this approximation results in a major issue for simulations of light yield at muon energies below $\sim$10 GeV.

▷ **Case 2: The secondary particle is <u>not</u> a muon**

This applies to all other secondaries of neutrino interaction products in the ice, such as hadrons (e.g., $\pi^{\pm}, K^{\pm}, K^0, p, \bar{p}, n, \bar{n}\ldots$) and electromagnetic (EM) shower primaries. The latter include $e^{\pm}$, γ, and π^0, which almost immediately decays into two photons. In Figure 6.3, we summarize how the propagation of the secondary particles other than muons and their Cherenkov photon generation is treated in CLSim.
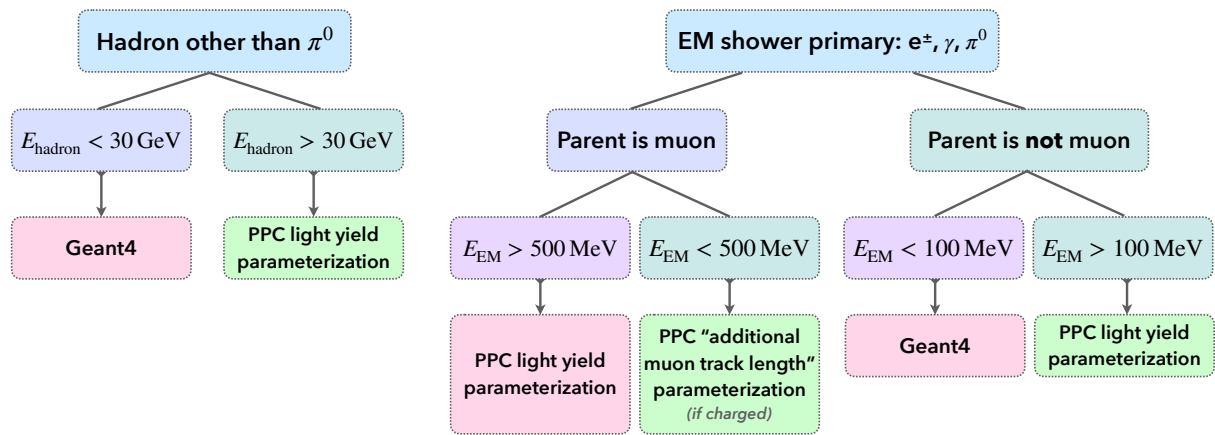

Figure 6.3 – Decision tree followed at Step 2 of the MC chain to decide whether a given secondary particle should be propagated in Geant4 [224] (followed by Cherenkov photon emission in Geant4), or whether its Cherenkov light yield should be parameterized in the ppc code via the energy-dependent functions developed in [221, 225, 226]. This applies to all secondaries other than muons; for the description of how muons are handled, see Case 1 in text.

As seen from the above diagram, there are two broad ways of handling hadrons and EM shower primaries: either via the widely used and well-supported Geant4 software [224], which both propagates these particles and generates their Cherenkov photons, or via ppc, which bypasses particle propagation and immediately generates photons based on parameterizations of the Cherenkov light yield of these particles. These parameterizations are derived based on Geant4 simulations and are meant to provide a significant speed-up compared to the Geant4 propagation initiated from within CLSim for every secondary. We find no issues with hadronic light yield parameterizations derived in [225]; however, in Section 10.2, we identify discrepancies between the ppc electromagnetic light yield parameterization (a linear dependence of the expected number of photons on the EM particle energy) and the Geant4 simulations.

A special case is that of the charged low-energy (< 500 MeV) muon secondaries, namely $e^{\pm}$ produced as the result of muon energy losses due to ionization or pair production. These particles are not propagated in Geant4, nor is their light yield parameterized in ppc on a per-particle basis. Instead, a parameterization of the "effective Cherenkov-radiating track length" that these particles collectively add to the muon is derived based on the Geant4 simulations [221]. Specifically, the path lengths of the low-energy $e^{\pm}$ travelled until stopping are (1) extracted from

the G\textsc{eant}4 muon simulations; (2) corrected for the $\beta_{e^{\pm}}$ velocity factors; and (3) summed together. This total Cherenkov-emitting length is ultimately used to parameterize the additional light yield to be added on top of that of their parent muon, depending on the muon energy E_μ. Importantly, this parameterization was derived in [221] for $E_\mu > 100\,\text{GeV}$ and has not been tested at lower energies. In Appendix A.1, we comment on discrepancies between our standalone G\textsc{eant}4 simulations, extended down to $E_\mu = 1\,\text{GeV}$, and the PPC parameterization of the light yield of the muon low-energy secondaries.

Once the generation of Cherenkov photons is complete, they are propagated towards the IceCube DOMs within CLS\textsc{im}. Then, these photons can be either accepted or rejected, depending on their arrival direction and wavelength (see Figures 5.5 and 5.6). As mentioned in Section 5.2.2, the angular acceptance of photons in ice is not the same as for the DOMs tested in laboratory conditions and depends on the local ice properties in the drill hole (known as *hole ice*). The dependence of the optical efficiency on the angle of photon incidence is described by two parameters, referred to as head-on efficiency and lateral efficiency in the GRECO (three-year Sample A) analyses, or p_0 and p_1 in the oscNext sample analyses[6]. The example variations in the angular acceptance possible through changing p_0 and p_1 are shown in Figure 6.4.
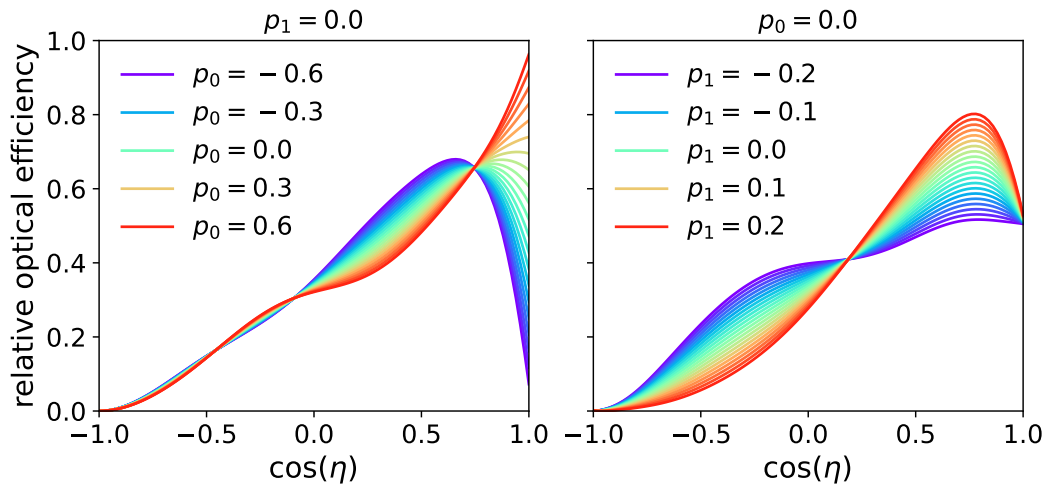

F\textsc{igure} 6.4 – Modifications in the photon angular acceptance curve (see Figure 5.5) possible by changing the head-on (p_0) and the lateral (p_1) optical efficiency parameters. *Figure extracted from* [227].

The laboratory-based curve from Figure 5.5 is reproduced with $p_0 \simeq 0.56$, $p_1 \simeq 0.16$. In the IceCube oscillation analyses, these parameters are typically left free in order to account for the systematic uncertainty in the hole ice angular acceptance. In [180] ([197]), the best-fitting values were $p_0 = -0.27\,(-0.14)$, $p_1 = -0.04\,(-0.07)$. The accepted photons are converted into Monte Carlo photoelectrons (MCPEs); from here, the PMT response simulation begins.

- **Step 3: PMT response simulation**

The simulation of the PMT response encompasses the production of PMT waveforms with and an amplitude and a time delay (relative to the MCPE production time) expected after the MCPE passage through the PMT tube and the ensuing charge amplification. This is handled in the IceCube software called PMTR\textsc{esponse}S\textsc{imulator} [228]. The latter takes care of the PMT response not only to the neutrino (or background muon)

---

[6]These two parameters correspond to the first two components in the principal component analysis (PCA) decomposition of the measurement-based variations in the angular acceptance (see [180, 227]).

signal MCPEs, but also to the noise MCPEs, which are added to the MCPE series as described in Section 6.2.2. The pre-, late and afterpulses (see [215]) are also simulated in PMTRᴇsᴘᴏɴsᴇSɪᴍᴜʟᴀᴛᴏʀ.

Thereafter, the simulation chain proceeds with DOMLᴀᴜɴᴄʜᴇʀ, where the local coincidence criteria are checked in the same way as described in Section 6.1.1 for real data events. Depending on which LC criteria are met, the PMT waveform digitization is handled within DOMLᴀᴜɴᴄʜᴇʀ with the software implementations of FADC or FADC+ATWD digitizers. Finally, Wᴀᴠᴇᴅᴇғᴏʀᴍ is run to reconstruct the pulses, their times, and their charges with the help of the SPE templates stored in the MC GCD file.

6.2.2 *Noise simulation*

The noise simulation in IceCube is performed with the ᴠᴜᴠᴜᴢᴇʟᴀ module. Both the uncorrelated thermal noise and the correlated noise due to radioactive decays in the OM glass are taken into account. The characteristic frequencies associated with these noise components (∼20 Hz and ∼200 Hz) are derived from fits to the untriggered IceCube data on a per-DOM basis [4, 229]. These fits are designed to describe the distribution of time differences between temporally adjacent hits (the so-called Δt distribution). The noise MCPEs are added to the event MCPE series over a time window that begins 10 µs before the first signal MCPE and ends 10 µs after the last signal MCPE [230].

In the IceCube-Upgrade mDOM modules, the expected dark noise rate per PMT was ∼450 Hz by original design. However, due to the usage of a new furnace in the PMT production process by the Hamamatsu company, the noise rate increased to ∼750 Hz. This is taken into account in the IceCube-Upgrade Monte Carlo simulation produced after 2022. Noise tagging and removal will be handled in the IceCube-Upgrade with the ᴅʏɴᴇᴅɢᴇ GNN, which was shown to be extremely efficient based on ICU Monte Carlo studies (specifically, reducing the amount of noise by a factor ∼10, with only ∼7% of noise pulses and as much as 90% of the signal remaining on average [212]).

6.3 ᴇᴠᴇɴᴛ sᴇʟᴇᴄᴛɪᴏɴ

Without any filter conditions or other forms of cuts applied, the event rate in IceCube is dominated by the atmospheric muon background, with a median rate of ∼2.7 kHz. The rate of muons passing the SMT3 trigger is approximately 250 Hz and is further reduced down to $\mathbb{O}(10\,\text{Hz})$ with the full DeepCore filter [4, 41, 45].

This reduction in the muon background, however, is still not sufficient for a neutrino-based physics analysis, with a characteristic atmospheric neutrino rate of 8 mHz after the DC filter. For this reason, multi-stage procedures called *event selections* are developed in IceCube, with the goal of arriving at a neutrino-dominated sample with an optimal purity and selection efficiency for a given analysis. The two event selections of relevance for this thesis are GRECO and oscNext, which are used in the non-unitarity analysis in Chapters 7 and 8 and for development of the high-statistics tau neutrino appearance analysis in Chapters 9 and 10. We provide a brief overview of these event selections below.

- **GRECO event selection**

    This event selection was originally developed for the ν_τ appearance analysis with the three-year IceCube-DeepCore dataset, covering the data-taking period from April 2012 to May 2015. It remains the only analysis of its kind published by the IceCube Collaboration to date [46], and the corresponding data and Monte

Carlo simulation have been publicly released [3]. The sequential reduction of the background rates throughout the GRECO event selection levels are summarized in Figure 6.5 and in text below[7] (for all levels following the filtering procedure described in Section 6.1.1).
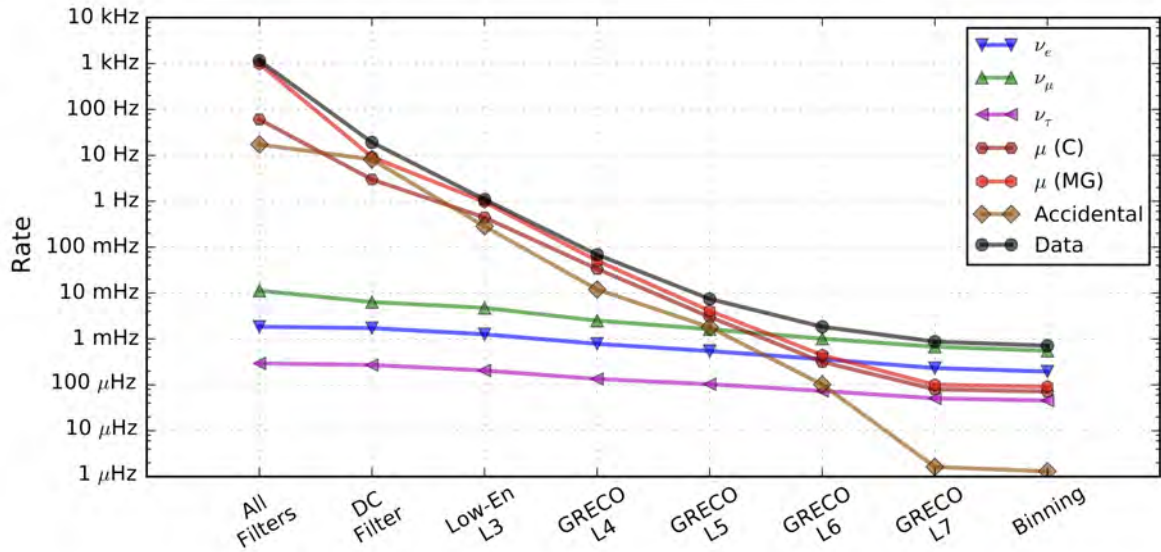

FIGURE 6.5 – The evolution of the neutrino signal and background event rates throughout the GRECO event selection. The pure noise events that trigger DeepCore appear under the label "accidental" in this figure. For muons, the two curves correspond to the full-sky CORSIKA ("C") [231] and the DeepCore-targeted MuonGun ("MG") [4, 219] simulations. *Figure courtesy of M. Larson (extracted from [4]).*

Level 3: The rejection of atmospheric muon events at this level is based on their expected light emission patterns and hit positions. The Level 3 cuts utilize the z position of the first hit, the charge in the top part of the detector (NAbove200), and employ the RTVeto algorithm [232] to detect clusters of hits in the veto region. Additionally, the effective speed of light deposition, evaluated as the fraction of the total charge collected within the first 600 ns, allows to distinguish muons (depositing light over few-microsecond timescales) from low-energy neutrino events (depositing light much quicker). Separately, requirements are imposed on the number of pulses, the total charge, and the relative directionality of hits in the cleaned pulse series, which removes over 96% of accidental noise triggers [4].

Level 4: A stricter hit cleaning compared to Level 3 is applied here. Then, the atmospheric muons are identified using the same variables as at Level 3, plus the Tensor of Inertia and the LineFit speed. The former quantifies the spatial patterns of hits as the three (x, y, and z) charge moments of inertia, which are together used to differentiate between the topologies expected for muon-like and neutrino-like events. The latter fits a plane wave to the event hits and derives its effective speed – a variable motivated by the fact that muons would move through ice at the speed of light. These variables are fed into a Boosted Decision Tree (BDT) [233], which is trained as a classifier to distinguish between the background muons and the neutrino signal. All events with a BDT score (nominally covering the range from -1 to 1, where 1 corresponds to signal-like events) smaller than 0.04 are removed, which reduces the muon background by a factor of ~20.

---

[7]A more complete description of each level can be found in the original source, [4].

Level 5: At this level, GRECO uses yet another BDT to remove atmospheric muon events based on the variables such as t_{75} (time required to collect 75% of all event charge), the total charge in the veto region passing the causality criteria, and radial distance of the first HLC hit to string 36. Additionally, the distance between the centers of gravity of the first and the last temporal quartiles, as well as the z-distance between the charge-weighted average position of all hits and the event vertex estimate, are included into the BDT. As at Level 4, all events with the BDT score below 0.04 are classified as muons are removed from the sample. Finally, the likelihood-based SPE reconstruction [234] is run at this stage, providing an estimate of the particle direction and interaction vertex. This reconstruction requires at least 6 hit DOMs in an event, which provides an implicit reduction of the pure noise events.

Level 6: Here, special attention is devoted to the so-called *muon corridors*, i.e., long gaps between strings in the IceCube layout which can serve as pathways for atmospheric muons to enter DeepCore unnoticed. Specifically, if more than 1 hit is observed at a string at the beginning of a corridor, the event is rejected at this stage. Furthermore, a cut is applied on the starting vertex position, which is required to be within the DeepCore fiducial volume both radially and vertically. Furthermore, the *fill ratio* of an event, i.e., the ratio number of hit DOMs contained within a certain radius to the total number of DOMs within that radius, is required to be greater than 0.05. This is done to prioritize the selection of low-energy neutrino events with compactly distributed hits and the rejection of noise events with sparser hit distributions. To prepare for the Level 7 reconstruction, a cut on the number of hit DOMs is once again applied, this time requiring at least 8 hit DOMs in the cleaned pulse series.

Level 7: This is the final level of the GRECO event selection, where the likelihood-based PEGLEG reconstruction [235] is run. PEGLEG yields the refined estimates of the particle direction, interaction vertex position, and interaction time, supplementing them with the energy reconstruction as described in Section 6.4. A two-dimensional geometric cut is applied using the updated z and ρ vertex coordinates for the purpose of further muon background reduction. Finally, a set of requirements is imposed to remove events with flaring DOMs (reflected in large values of reconstructed energy per hit DOM) as well as remaining noise events with a large amount of hit time scatter.

Table 6.1 reports how the rates of the different event types in IceCube-DeepCore change between the DeepCore filter level and the Level 7 of the GRECO event selection. This shows that the goal of a neutrino-dominated sample was achieved exceptionally well with GRECO, with the neutrino rate exceeding the muon (noise) rate by a factor of ~10 (1000).

Sel. level	$\nu_e + \bar{\nu}_e$ [mHz]	$\nu_\mu + \bar{\nu}_\mu$ [mHz]	$\nu_\tau + \bar{\nu}_\tau$ [mHz]	$\mu_{\rm atm}$ [mHz]	noise [mHz]
DC filter	1.721	6.360	0.270	9178	8117
Level 7, w/o cuts	0.325	0.676	0.051	0.08	0.002
Level 7, w/ cuts	0.194	0.552	0.045	0.07	0.001

TABLE 6.1 – Signal and background rates in the GRECO event selection at two selection levels: DeepCore filter (Level 2) and final level (Level 7, before and after analysis cuts). The rates are extracted from [4].

- **oscNext event selection**

  The oscNext event selection was developed as a successor to the GRECO event selection, with the goal of combining higher-statistics IceCube data with improved Monte Carlo simulation, detector calibration, and data processing. The original effort aimed at incorporating 8.5 years of IceCube data (2011–2019), which was ultimately used for the "golden tracks" oscNext subsample and the corresponding ν_μ disappearance analysis published in [180]. In this subsample, also known as the *oscNext verification sample*, only events with strict cuts on direct light (i.e. extremely low amounts of photon scattering) were selected. This resulted in a total of ~22 000 events, dominated by $\nu_\mu + \bar{\nu}_\mu$ CC interactions (80% of the sample). The development of the full sample without the aforementioned direct light cuts, known as the *oscNext high-statistics sample*, proceeded in parallel, with ~200 000 events in total expected in ~9 years[8]. The goal of this full sample was to enable both a higher-precision ν_μ disappearance measurement and the first ν_τ appearance measurement with an unprecedented statistics of ~10 000 $(\nu_\tau + \bar{\nu}_\tau)$-CC events[9]. The original high-statistics sample, which formed the basis for the first iteration of the two mentioned analyses [54, 236], relied on the likelihood-based RETRO reconstruction [211] (see Section 10.3). Both of these analyses resulted in failed blind fits (i.e., not meeting the goodness-of-fit condition of a p-value being greater than 5%), thus having launched a long sequence of investigations partially covered in Chapter 10. A parallel ν_μ disappearance analysis branched out after the first 5 levels of the oscNext event selection, replacing the likelihood-based reconstruction with one based on convolutional neural networks (CNNs) and introducing stricter cuts on the analysis variables. This analysis was successfully unblinded and published in [197]. Since a significant fraction of the work presented in Chapter 10 was set out to explain why the RETRO-based analyses did not succeed, we will describe the levels of event selection as designed prior to the introduction of the CNN reconstruction. A visual representation of what we will therefore refer to as the "oscNext-high statistics" event selection is given in Figure 6.6, with the per-level description of all levels after the DeepCore filter following in text below.

  Level 3: The Level 3 of the oscNext event selection is very similar to that of the GRECO event selection described earlier. One of the differences is that the requirement of at least 6 hit DOMs, which was introduced in GRECO only at Level 5, is applied already at this level in oscNext. Furthermore, the variables that previously depended on charge are replaced with more robust variables dependent on the number of hit DOMs, thereby avoiding any potential charge mismodeling issues [54]. Finally, a new cut is introduced on the overall duration of both cleaned and uncleaned pulse series, which are required to be shorter than 5 µs and 13 µs in duration, respectively. This eliminates coincident events in data (i.e., the cases where two events occur within a single trigger time window), which are not simulated in the oscNext MC [236].

  Level 4: Analogously to the Level 4 of GRECO, a BDT is trained here to distinguish between neutrino signal and muon background. This BDT has largely the same input variables as in GRECO, but uses the atmospheric muon data rather

---

[8]At this point, season 2011 was dropped from the sample due to unusually high event rates and failure to pass seasonal compatibility checks with the rest of the seasons [54]; however, subsequent data seasons (2020 and later) were continuously added to the sample alongside its development.

[9]In comparison, the three-year GRECO event selection contained $\lesssim$2 000 $(\nu_\tau + \bar{\nu}_\tau)$-CC events. Note that these events were not identified on an individual basis; instead, this number provides their estimated contribution to the overall sample.

FIGURE 6.6 – The evolution of the neutrino signal and background event rates throughout the oscNext-high statistics event selection. *Figure courtesy of K. Leonard DeHolton (extracted from [236]).*

than CORSIKA simulation to provide representative muon events for the BDT training. Unlike in GRECO, where simple cuts were applied at this level to remove noise events, the Level 4 noise rejection in oscNext is also handled via a BDT. One of the new input variables for this BDT is the ratio of the event duration in the SRTTW-cleaned pulse series to that in the uncleaned pulse series, with pure noise events shown to result in larger values of this ratio. With both of the BDT classifier scores ranging from 0 (background-like) to 1 (signal-like), only events with L4_muon_classifier > 0.7 and L4_noise_classifier > 0.65 are kept in the sample.

Level 6: The main purpose of Level 6 is the high-level reconstruction of the physics variables, namely neutrino energy, direction, interaction vertex position, and interaction time. In the variant of the oscNext sample discussed here, this was handled with the likelihood-based RETRO algorithm, which we describe in detail in Section 6.4. As a preparation for the RETRO reconstruction, only events with at least 8 pulses are selected at this stage. Furthermore, the events where the RETRO reconstruction fails (either by not converging or by terminating after very few iterations) are removed from the sample. Finally, the SANTA reconstruction [237], which was used for the oscNext verification sample described above, is run on events that have at least five direct pulses (i.e., those likely coming from unscattered photons). This reconstruction provides an additional handle on the neutrino arrival direction, whose deviation from the RETRO-reconstructed direction is used as one of the inputs to the Level 7 muon classifier.

Level 7: This is the final stage of the oscNext event selection, where the high-level muon and PID classifiers are run, and additional cuts[10] on the high-level variables are applied to prepare the sample for analyses. In preparation for the muon classifier, a corridor hit identification algorithm from Level 5 is re-run here

---

[10]These include cuts such as the radial and the vertical containment of the interaction vertex reconstructed at Level 6, the requirement of at least 3 direct hits in an event, and an upper limit on the number of hit DOMs in the outermost IceCube strings. For more details, see [54, 236].

with tighter spatiotemporal settings. Then, a BDT is trained to classify MuonGun-simulated muons against neutrinos, based on the Level 4 muon classifier output, the retro-reconstructed vertex coordinates, the difference between the santa- and the retro-reconstructed directions, and the geometry of the identified Level 7 corridor hits[11]. Only events with the L7_muon_classifier score larger than 0.8 are kept for analysis[12], thereby making the muon contribution to the overall sample less than 1%. Separately, another BDT is trained as a PID classifier in order to distinguish between events containing a muon track – i.e., $(\nu_\mu + \bar{\nu}_\mu)$-CC interactions – and all other events[13]. The input variables to this PID classifier consist of the retro-reconstructed track length, the log-likelihood difference between the track-only and the track+cascade reconstruction hypothesis, the retro-reconstructed cascade energy, the reconstructed neutrino zenith angle, and its estimated uncertainty. The output PID score, ranging from 0 (cascade-like) to 1 (track-like), is used to bin the Level 7 events into three analysis bins: *cascades* (PID $\in$ [0.0, 0.5]), *mixed* (PID $\in$ [0.5, 0.85]), and *tracks* (PID $\in$ [0.85, 1.0]).

In Table 6.2, we summarize the rates of the different event types in IceCube-DeepCore at the DeepCore filter level and the Level 7 of the oscNext event selection. With the latter, we arrive at a sample where the neutrino rate exceeds the muon rate by a factor of ~100 – a significant improvement compared to the GRECO selection. Furthermore, the final noise rate in oscNext is effectively negligible; it is estimated to account for less than 0.03% of the final-level sample [236].

Sel. level	$(\nu_e + \bar{\nu}_e)$-CC [mHz]	$(\nu_\mu + \bar{\nu}_\mu)$-CC [mHz]	$(\nu_\tau + \bar{\nu}_\tau)$-CC [mHz]	all NC [mHz]	μ_{atm} [mHz]	noise [mHz]
DC filter	1.610	6.160	0.193	0.860	7273	6621
Level 7, w/o cuts	0.433	1.340	0.067	0.130	0.970	0.057
Level 7, w/ cuts	0.168	0.455	0.033	0.050	0.005	–

Table 6.2 – Signal and background rates in the oscNext event selection at two selection levels: DeepCore filter (Level 2) and final level (Level 7). The DC filter rates are extracted from [54]. The Level 7 rates correspond to the oscNext Monte Carlo simulation produced with the birefringence ice model [240] (not available at the time of the analyses in [54, 236]), but without the muon light yield correction discussed in Section 10.1.

6.4 EVENT RECONSTRUCTION

6.4.1 *Pegleg reconstruction*

Pegleg [235] is a reconstruction developed for low-energy ($\mathcal{O}(10\,\text{GeV})$) events, which is based on the track + cascade hypothesis fitted to the observed event hits. Specifically, the Pegleg hypothesis consists of 8 parameters: cascade energy E_{cascade}, muon track length ℓ_{track} (equivalently, ℓ_μ), interaction vertex position $\langle x, y, z \rangle$ and time t, as well as the original neutrino arrival direction represented by the zenith angle θ_{zen} and the azimuthal angle ϕ[14]. The track length is converted into the estimate of the track energy

---

[11]This LightGBM-based [238] classifier was developed in [54]. The input variables were selected based on the SHAP importance metrics [239].

[12]Note that in the original ν_τ appearance analysis attempted in [54], the threshold for the muon classifier score was 0.4.

[13]This BDT was based on the XGBoost algorithm and developed in [236].

[14]Note that these angles represent the direction from which the neutrino came from, such that $\theta_{\text{zen}} = \pi$ corresponds to Earth-crossing neutrinos and $\phi = 0$ indicates the direction from the positive x-axis towards 0 (in the IceCube detector coordinates).

through the assumption of a minimally ionizing muon with the average energy loss of $\langle \frac{dE}{dx} \rangle = 0.22 \text{GeV m}^{-1}$. Given this eight-dimensional particle hypothesis, a likelihood to observe a hit on a DOM[15] at a given timestep and at a given distance from the light source is evaluated. This is done with the help of splines derived from large-scale CLSim [222] simulations of photon emission, propagation, and detection, which are run in advance of the reconstruction and stored in a table format. Due to the complex likelihood space and a large number of local minima expected for the low-energy events, the likelihood optimization is performed with the nested sampling algorithm via the
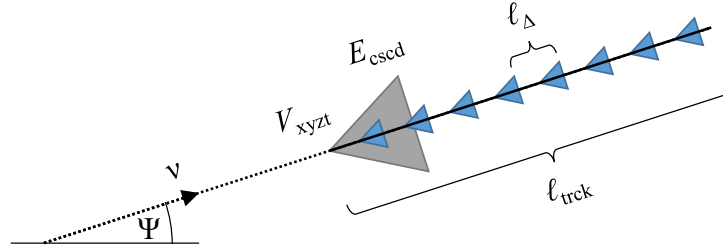

FIGURE 6.7 – The eight-dimensional RETRO hypothesis: a neutrino arriving from the direction $\Psi = (\theta_{\text{zen}}, \phi)$ and interacting at the vertex V with the spatial coordinates $\langle x, y, z \rangle$ at the time t, with the subsequent production of a cascade (hadronic and/or electromagnetic shower) with total energy E_{cascade} and a muon with track length ℓ_{track}. Note that in the absence of a muon, the true ℓ_{track} is equal to 0. *Figure extracted from* [211].

The muon track is discretized into steps of length $\Delta \ell = 3\,\text{ns} \cdot c \approx 0.9\,\text{m}$, each of which is assumed to emit $2450\,\text{photons m}^{-1} \cdot \Delta \ell \approx 2200$ photons (according to the Frank-Tamm formula convolved with the DOM wavelength acceptance and integrated over the 265–675 nm wavelength range). Thus, the actual fitted variable in the RETRO reconstruction is the integer number n of Cherenkov emitters along the muon track. The cascade part of the event is modeled as a point emitter yielding 12 800 Cherenkov photons per GeV cascade energy[16] [211, 225]. The likelihood is optimized with the global controlled random search (CRS) minimizer [243] with spherical awareness[17] (i.e., adequate treatment of the angular distributions on a sphere). As mentioned in Section 6.3, a few of the RETRO-reconstructed quantities (E_{cascade}, ℓ_{track}, log-likelihood difference between the optimized cascade-only and track+cascade hypotheses, θ_{zen}, and its 1σ uncertainty estimate) are used for PID classification in the oscNext sample. The rest of the reconstructed variables are used either for the final-level cuts or directly in the oscNext oscillation analyses.

From a technical standpoint, the main novelty of the RETRO reconstruction is the process through which the lookup tables for the likelihood computation are generated.

---

[15] Originally, the full charge expectation was calculated, but the charge information was eventually dropped due to concerns regarding charge mismodeling [4].

[16] Note that this number was derived from simulations of electromagnetic (EM) cascades. Therefore, the cascade energy is reconstructed under the assumption that all Cherenkov photons come from an EM rather than hadronic shower. To account for this, an energy-dependent correction factor derived from GEANT4 simulations of hadronic and EM cascades is applied to the reconstructed cascade energy. For more details, see [54, 211].

[17] https://pypi.org/project/spherical-opt

Compared to the PEGLEG table generation, where the light sources are distributed
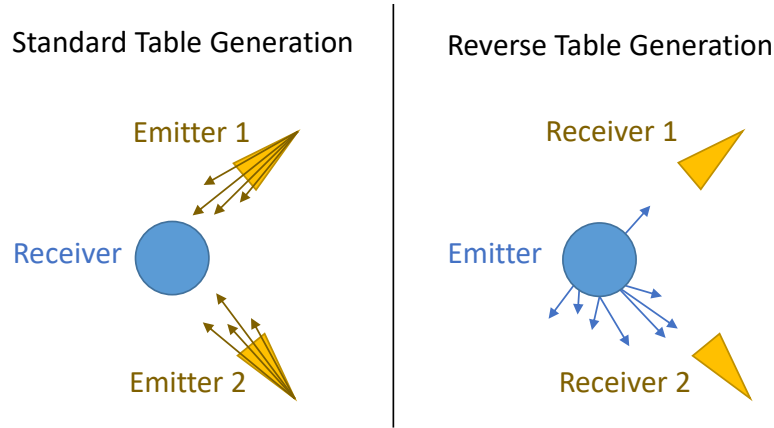isotro
they
Figur

FIGURE 6.8 – Illustration of the photon table generation for the RETRO reconstruction, where the roles of
the receiver and the emitter are reversed. *Figure extracted from* [211].

are propagated away from the DOMs according to the DOM angular acceptance and
the ice properties. This provides a significant speedup to the lookup table generation,
as the computational resources are not wasted on generating and propagating photons
that may never reach a DOM.

At the time of the original oscNext high-statistics sample development, the RETRO
reconstruction relied on individual *pulses* (of which a minimum of 8 was required for
the reconstruction to commence at Level 6 of the event selection) and their charges.
This is in contrast with the the hit/no hit information used in the final version of PEGLEG
for the reconstruction of the GRECO events. As we discover later in Section 10.3, the
explicit inclusion of the pulse+charge information results in RETRO reconstruction being
highly sensitive to noise pulses and ultimately leads to data/Monte Carlo disagreement
due to noise charge mismodeling in simulation.

6.4.3 *DynEdge reconstruction*

DYNEDGE is a specific architecture of a graph neural network (GNN) proposed in [210]
as a method of low-energy event reconstruction in IceCube-DeepCore. It was developed
and is continuously maintained as part of the open-source GRAPHNET project [244]
aiming to provide flexible reconstructions for a variety of neutrino telescopes. The GNN
architecture is particularly well-suited to the IceCube geometry, which represents an
irregular grid in all three spatial dimensions[18] (see Figure 5.7 and text of Section 5.3).

The idea of a GNN in application to IceCube is to represent the hit DOMs or pulses
on the hit DOMs as *graph nodes*. The usual task is then to map the features of the
individual nodes (such as the pulse spatial coordinates, times, and charges) and their
local neighborhoods onto the necessary *global* information about the events (such as the
primary neutrino energy, arrival direction, and vertex position)[19]. This mapping is done

---

[18]The conventional convolutional neural networks (CNNs), which are typically used in learning from image data and require
regular grids (equidistant pixels), can in principle still be applied to IceCube's irregular geometry (as was done in, e.g., [197]).
However, GNNs provide a much more natural architecture for this task.

[19]These are examples of *graph-level* tasks, i.e., predicting the properties of an entire event. As mentioned in Section 6.2.2,
DYNEDGE has also been successfully used for *node-level* tasks such as signal/noise classification of individual pulses.

via a sequence of *graph convolutional layers*, which find the neighbors of each node and combine the features of the neighbors with those of the central node. This collection of information from and exchange of information between the local neighborhoods is called *message passing*, which is schematically illustrated in Figure 6.9 for a generic graph convolutional neural network.
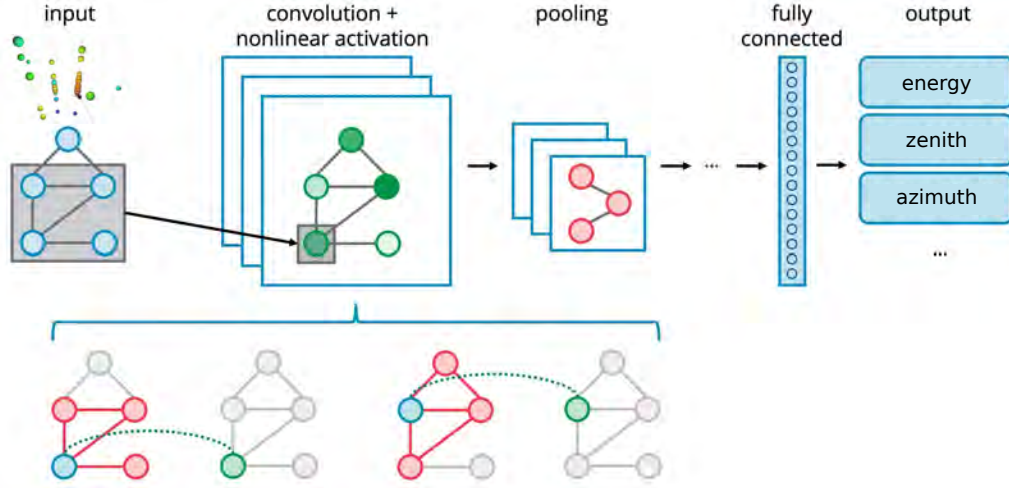

FIGURE 6.9 – A simplified visualization of a graph neural network (GNN), which takes the unordered IceCube pulse series from an event as an input (left) and is trained to predict the global event properties (right). Note that this figure (*adapted from* [245]) serves as an illustration only. The real DYNEDGE architecture can be found in [210].

As seen from Figure 6.9, the aggregation of information from the node neighborhoods is generally followed by nonlinear activation (i.e., passing the aggregated features through a nonlinear function) and pooling (reducing the aggregated information by summing, averaging, or taking a minimum/maximum value across nodes). After pooling, all information is flattened into a single vector, whose entries are combined via a linear superposition (followed by another nonlinear transformation, if necessary) into one or many output variables. The network is trained by optimizing the weights, which multiply the units of information from the preceding layers to form the input to the subsequent ones, such that the target loss function is minimized. The exact DYNEDGE loss functions for the regression of particle properties and the classification of event topologies (track/cascade) are given in [210].

The DYNEDGE GNN is the fastest among the reconstruction methods presented in this section, given a typical per-event reconstruction speed of 10 min with PEGLEG [4], 40 s with RETRO, and 10 ms with DYNEDGE (on a single CPU and after the network has been trained) [210]. Furthermore, as previously shown in Figure 5.8, DYNEDGE provides highly competitive resolutions for all the essential oscillation analysis variables, unmatched by any other algorithm at the time of writing.

IV

Global fit for neutrino mixing non-unitarity

Analysis description

N.B.: Chapters 7 and 8 contain the reproduction of the following paper:

Tetiana Kozynets, Philipp Eller, Alan Zander, Manuel Ettengruber, and David Jason Koskinen, *Constraints on Non-Unitary Neutrino Mixing in Light of Atmospheric and Reactor Neutrino Data*, submitted to J. High Energy Phys. (2024).

The full pre-print of the paper is available online in the arXiv database [1].

7.1 MOTIVATION AND GOALS

Percent and sub-percent constraints on leptonic non-unitarity in the minimal unitarity violation regime (see Section 2.4) have been placed via e.g. searches for flavor-violating decays of charged leptons, electroweak universality tests, measurements of the invisible Z-boson decay width, and other searches detailed in [109, 246, 247]. At the same time, neutrino oscillation data provides a way to study the unitarity of neutrino mixing directly, without invoking complementary channels or constraints from the electroweak sector. This approach has been taken in e.g. [51–53, 248], where neutrino oscillation amplitudes were inferred from published experimental results and reinterpreted as constraints on the individual mixing matrix elements and combinations thereof. Such constraints were typically derived as the result of global fit studies, in which neutrino experiments with different baselines, energy ranges, and oscillation channels were combined to probe the entire neutrino mixing matrix. All of short- and long-baseline, solar, reactor, and atmospheric neutrino experiments were fully or partially incorporated in the literature dedicated to global fits [51, 53, 248, 249], and projections for selected next-generation experiments have similarly been made [125, 250, 251].

The main goal of the present study is to draw attention to the treatment of atmospheric neutrino data in the non-unitarity analyses, which, to the best of our knowledge, has so far been incomplete. Although several preceding studies [51, 52] reinterpreted the Super-Kamiokande results [48, 252, 253] and the IceCube-DeepCore oscillation results [46, 196] to constrain unitarity in the muon and the tau rows of the mixing matrix,

these analyses have not considered atmospheric neutrino systematic uncertainties as prescribed by the respective collaborations. As we show in this work, the latter are crucial for placing accurate non-unitarity constraints, since the nuisance parameters may introduce energy- and direction-dependent effects correlated with those of the non-unitarity physics. In particular, the relevant uncertainties include those related to the unoscillated atmospheric neutrino flux, such as the overall normalization, spectral index, $\nu/\bar{\nu}$ and ν_e/ν_μ ratios, and angular distributions [140], as well as detector-specific systematics. We seek to address these uncertainties in the context of non-unitarity constraints for the first time and study their degeneracies with the non-unitarity metrics (such as the normalizations of the mixing matrix rows and columns). We stress that fitting for the non-unitarity physics parameters and the experimental systematic parameters at the same time is a significantly more robust approach than the reinterpretation of the best-fit constraints from the three-flavor analyses, since the best-fit systematic parameters might already be absorbing some non-unitarity effects.

The approach we propose above is possible with the public atmospheric neutrino data and Monte Carlo simulation from the IceCube-DeepCore experiment [3], which has not been considered in literature as part of any global fit studies focused on non-unitarity[1]. This dataset covers atmospheric muon neutrino disappearance and tau neutrino appearance channels, thereby providing access to the elements of the muon and the tau rows of the neutrino mixing matrix. To form a "minimal" selection of datasets that would let us probe all three matrix rows, we supplement the IceCube-DeepCore data with reactor neutrino data from Daya Bay [49] and KamLAND [50] experiments. The reactor experiments provide a handle on the elements of the electron row of the mixing matrix, whose measurements are similarly subject to the reactor systematic uncertainties and are implemented in this study to the extent possible with the publicly available information. To assess how the non-unitarity constraints possible with this minimal selection of experiments will evolve in the future, we further develop an equivalent analysis including the next-generation IceCube-Upgrade [205] and JUNO [255] experiments (both in deployment at the time of writing). For these future projections, we once again perform the global fit for the individual mixing matrix elements not constrained by unitarity alongside the systematic parameters known at this stage. While the future constraints will evolve as the experiments become operational and settle on the event selection and data analysis pipelines, our study provides the first attempt to utilize the preliminary IceCube-Upgrade simulation in this context [213] and serves as a proof of concept for the upcoming studies with an improved detector simulation and the upcoming data.

7.2 DATASETS

7.2.1 *Atmospheric neutrino data*

- **IceCube-DeepCore**

In this study, we use the public release of the DeepCore data collected between 2012–2015 [3], which is referred to as "Sample A" in [46] and the GRECO event selection in Section 6.3. The reconstructed neutrino energies (E_ν^{reco}) span the 6–56 GeV range, while the reconstructed zenith angle ($\theta_{\mathrm{zen}}^{\mathrm{reco}}$) covers the full atmospheric neutrino sky ($\cos\theta_{\mathrm{zen}}^{\mathrm{reco}} \in [-1, 1]$). The release is supplemented by the simulated Monte Carlo neutrino

---

[1]This dataset corresponds to the previous generation of the IceCube-DeepCore analyses [46, 196]. Although it contains only ~3 years of data, it remains the only publicly available atmospheric neutrino dataset that contains all of the information necessary for performing a non-unitarity analysis. The recent public data release from the Super-Kamiokande experiment [254] does not provide any prescriptions for implementing the systematic uncertainties, which is why we are not considering it in this study.

events as well as the estimated muon background, which together pass through the same event selection stages as the data. The PID classification of all events is performed on the basis of the length of the reconstructed track (ℓ_{track}) fitted to the recorded light patterns in the detector. All events with $\ell_{\text{track}} > 50\,\text{m}$ are classified as track-like, while the rest fall into the cascade-like PID bin.

While the event-by-event information is provided for both data and neutrino+muon Monte Carlo samples, the statistical analysis is ultimately performed in the binned 3D (reconstructed energy, reconstructed $\cos\theta_{\text{zen}}$, PID) space. The prescriptions for incorporating the impact of the detector systematic uncertainties on the event count in individual bins are also supplied with the data release. This includes the bin count gradients with respect to the variations in the OM efficiency, the ice absorption/scattering coefficients, and the angular dependence of the OM acceptance. The systematic uncertainties related to the atmospheric neutrino flux and the interaction cross section are implemented separately in the public PISA software [256]. The full list of the systematic uncertainties ($\bar{\lambda}_{\text{syst}}$) used for the present analysis follows that of [46], with the exception of the quasielatic (QE) and the resonance (RES) cross section parameters[2].

Our analysis of the DeepCore data closely follows that of [46] and is based on the procedure devised in [256]. It consists of a staged multiplication of the event weights due to the unoscillated flux, oscillation probability, and the interaction cross section, ultimately followed by binning into the 3D analysis histograms and applying the systematic uncertainty gradients. The expected event rate in each analysis bin therefore depends on the oscillation hypothesis $\bar{\lambda}_{\text{osc}}$ and the systematic parameters $\bar{\lambda}_{\text{syst}}$. For DeepCore, the relevant parameters of $\bar{\lambda}_{\text{osc}}$ include $N_{\mu\{2,3\}}$, $N_{\tau\{2,3\}}$, and the mass splitting Δm^2_{32}. These parameters directly enter the oscillation probabilities $P_{\mu\mu}(E_\nu, L[\cos\theta_{\text{zen}}])$ and $P_{\mu\tau}(E_\nu, L[\cos\theta_{\text{zen}}])$ as per Equation (2.32).

Following the discussion of Section 2.4, we also need to consider the non-unitary corrections for the unoscillated fluxes and the interaction cross sections. The nominal atmospheric flux model for the DeepCore analysis is HKKMS-2015 [137], which relies on a hadronic interaction model calibrated by the muon spectrometer measurements [134, 257]. We therefore assume that the flux prediction is already contaminated by non-unitarity effects, in particular in the muon sector, and do not apply any additional normalization corrections related to flux. The nominal cross section model of the analysis, however, relies on Standard Model-based calculations within the GENIE framework [158] and therefore requires a normalization correction N_β for the final flavor β according to Equations (2.33b) and (2.33c)[3]. We absorb this correction factor in the effective oscillation probabilities $\hat{P}_{\alpha\beta}$, such that

$$\hat{P}^{\text{DC}}_{\mu\mu}(E_\nu, L_l \,|\, \bar{\lambda}_{\text{osc}}) = P_{\mu\mu}(E_\nu, L_l \,|\, \bar{\lambda}_{\text{osc}}) \cdot N_\mu = \frac{|(N e^{-iH(\rho_l)L_l} N^\dagger)_{\mu\mu}|^2}{N_\mu}; \tag{7.1a}$$

$$\hat{P}^{\text{DC}}_{\mu\tau}(E_\nu, L_l \,|\, \bar{\lambda}_{\text{osc}}) = P_{\mu\tau}(E_\nu, L_l \,|\, \bar{\lambda}_{\text{osc}}) \cdot N_\tau = \frac{|(N e^{-iH(\rho_l)L_l} N^\dagger)_{\tau\mu}|^2}{N_\mu} \tag{7.1b}$$

---

[2]The public DeepCore data release does not include the GENIE [158] coefficients necessary to implement the quasielastic and resonance cross section systematic uncertainties. Since the CC-QE (CC-RES) CC events are subdominant at energies $\gtrsim 6\,\text{GeV}$ and constitute only $\sim 9\%$ (14%) of the sample, we find it acceptable to proceed without these two systematics.

[3]An argument was made in [130] that only the deep inelastic scattering (DIS) cross sections need to be corrected for non-unitarity. Since the DIS events constitute the majority ($\gtrsim 70\%$) of the public DeepCore data release, we choose to apply the final-flavor correction to all events regardless of the interaction type. Furthermore, we apply a single correction factor of N_β to both CC and NC cross sections, since the relative normalization of the NC events relative to CC is already included as a systematic parameter [46].

for the Earth layer l of thickness L_l and constant density ρ_l. In this study, we consider a four-layer PREM Earth profile with constant-density layers [145] (see Figure 3.6) and generalize the phases in the exponents of Equations (7.1a) and (7.1b) to $-i \sum_l H(\rho_l) L_l$ to account for the propagation through the entire baseline of length $L = \sum_l L_l$.

In Figure 7.1, we show the impact of the off-nominal muon and tau row normalizations on the DeepCore analysis histograms, where we consider the cases $(N_\mu = 0.9, N_\tau = 1)$ and $(N_\mu = 1, N_\tau = 0.9)$ separately.
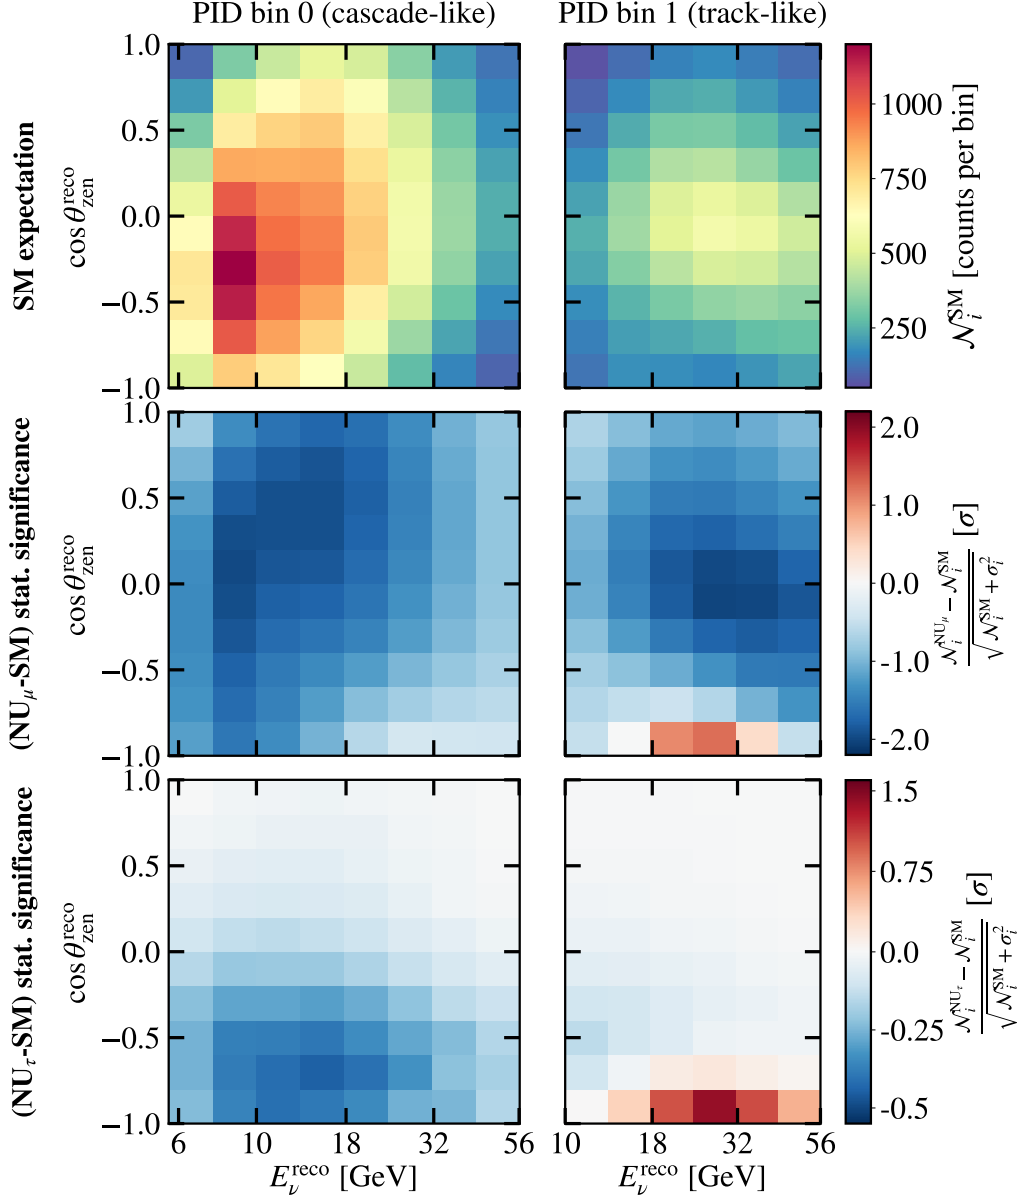

FIGURE 7.1 – *Top:* the unitary (Standard Model, "SM") expectation of the DeepCore event rates in the reconstructed $(E_\nu, \cos\theta_{\mathrm{zen}})$ space, assuming the livetime of 2.5 years, NuFit 5.2 oscillation parameters [65, 66], and the nominal values of the DeepCore systematic parameters [3, 46]. *Middle*: Statistical significance of the non-unitary ("NU") expectation with $(N_e, N_\mu, N_\tau) = (1, 0.9, 1)$ compared to the unitary case. *Bottom*: same as the middle panel, but for $(N_e, N_\mu, N_\tau) = (1, 1, 0.9)$.

As seen from the middle panel of Figure 7.1, $N_\mu < 1$ lowers the expected number of events in both track-like and cascade-like bins, which are dominated by ν_μ-CC events[4]. Similarly, the bottom panel shows the deficit of events predominantly occurring in the

---

[4]Even though ν_μ-CC events contain a true muon track, at low energies it might be classified as a cascade due to the small muon length ($\sim 4.5\,\mathrm{m/GeV}$ [258]) compared to the large DeepCore (75 m) interstring spacing [45].

cascade-like bin for the case of $N_\tau < 1$. The excesses of events at $\cos\theta_{\mathrm{zen}}^{\mathrm{reco}} \simeq -1$ in both cases are due to strong matter effects in this region, as it includes neutrinos crossing through the Earth's core (see Figure B.2 for the equivalent figure assuming neutrino propagation in vacuum). The statistical significances of the event deficits/excesses in the non-unitary mixing case compared to the unitary one are defined as

$$\Sigma_i = \frac{\mathcal{N}_i^{\mathrm{NU}} - \mathcal{N}_i^{\mathrm{SM}}}{\sqrt{\mathcal{N}_i^{\mathrm{SM}} + \sigma_i^2}} \tag{7.2}$$

for each analysis bin i, where σ_i is the Monte Carlo uncertainty reflecting the limited simulation statistics. When comparing the observed event rates ($\mathcal{N}^{\mathrm{obs}}$) to the expected ones ($\mathcal{N}^{\mathrm{exp}}$) under the hypothesis $\bar{\lambda} = \{\bar{\lambda}_{\mathrm{osc}}, \bar{\lambda}_{\mathrm{syst}}\}$, we employ a similar expression for the χ^2 test statistic,

$$\chi^2_{\mathrm{DC}}(\bar{\lambda}) = \sum_{i=1}^{n_{\mathrm{bins}}} \frac{(\mathcal{N}_i^{\mathrm{obs}} - \mathcal{N}_i^{\mathrm{exp}}(\bar{\lambda}))^2}{\mathcal{N}_i^{\mathrm{exp}}(\bar{\lambda}) + \sigma_i^2} + \sum_{j=1}^{\dim \bar{\lambda}_{\mathrm{syst}}} \left(\frac{s_j - \hat{s}_j}{\sigma_{s_j}}\right)^2. \tag{7.3}$$

The last term of Equation (7.3) is the penalty due to the systematic parameters $s_j \in \bar{\lambda}_{\mathrm{syst}}$, with $\hat{s}_j$ being their nominal values and σ_{s_j} – the widths of the respective Gaussian priors[5] [46]. We use the χ^2 from Equation (7.3) directly to reproduce the standard oscillation contours for ν_μ disappearance (see Section 7.3). For our main Bayesian analysis of the non-unitary mixing, we convert it to the log-likelihood:

$$\ln\mathcal{L}_{\mathrm{DC}}(\bar{\lambda}_{\mathrm{osc}}) = -\frac{1}{2}\chi^2_{\mathrm{DC}}, \tag{7.4}$$

assuming that Wilks' theorem [259] holds and dropping the $\bar{\lambda}_{\mathrm{osc}}$-independent log-likelihood term for the null hypothesis.

• **IceCube-Upgrade**

To incorporate the sensitivity of the upcoming IceCube-Upgrade (see Section 5.4) into the future projections for neutrino mixing matrix non-unitarity, we utilize the preliminary public release of the IceCube-Upgrade neutrino Monte Carlo simulation [213]. We note that this release corresponds to the sensitivities published in [205] and does not reflect the more recent progress in simulation, event selection, reconstruction, and rejection of random noise and muon background [210, 212]. At the time of the IceCube-Upgrade simulation release, the projected minimum improvement in the θ_{zen} reconstruction compared to the DeepCore-only configuration was a factor of ~ 3. This allows us to bin the $\cos\theta_{\mathrm{zen}}$ dimension of the IceCube-Upgrade expectation templates more finely compared to the corresponding DeepCore templates in Figure 7.1. We further assume a uniform logarithmic energy binning with 15 bins in the 1–100 GeV range. This choice is conservative in the energy range overlapping with that of the three-year DeepCore analysis but includes additional bins at both lower (1–6 GeV) and higher (56–100 GeV) reconstructed energies. The dimensions of the grid are tuned such that the least populated bin has $\gtrsim 20$ counts, which justifies the use of the same test statistic as in Equations (7.3) and (7.4).

In Figure 7.2, we show the impact of the off-nominal muon and tau row normalizations ($N_\mu = 0.9$ and $N_\tau = 0.9$, considered separately as before). When comparing these non-unitary expectations to the SM case, we find the same features as in Figure 7.1,

---

[5]If the priors are uniform, no penalty term is added.

which now provide higher statistical power for placing unitarity constraints due to the finer angular binning and the extended energy range.
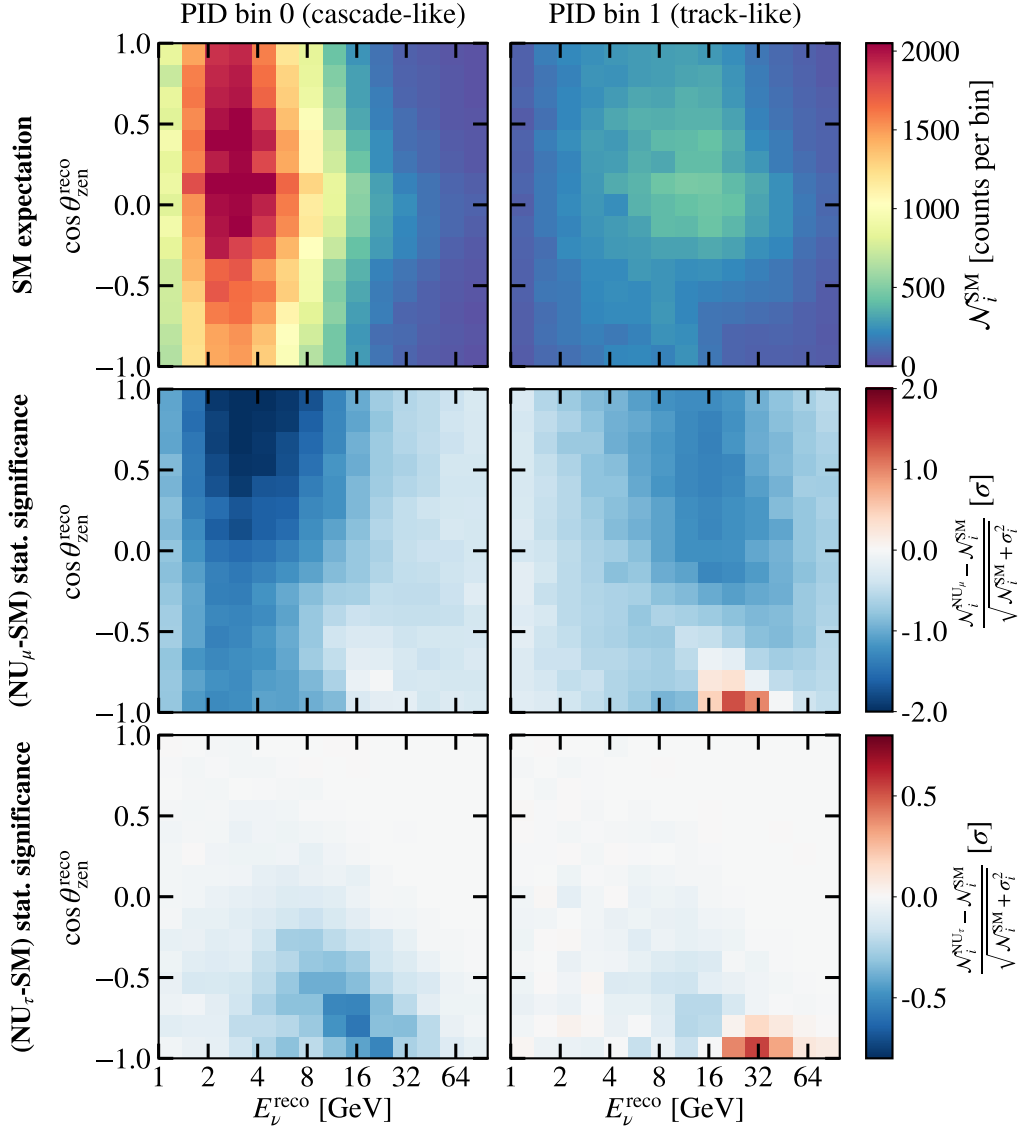

FIGURE 7.2 – Same as Figure 7.1, applied to the ICU dataset and binning configuration. The assumed livetime is 3 years.

For the IceCube-Upgrade analysis, we implement the same atmospheric neutrino flux uncertainties as for the DeepCore analysis, as well as the systematic uncertainty for the relative scaling of the NC cross section compared to CC. The detector systematic uncertainties are not possible to implement for the public IceCube-Upgrade Monte Carlo, as no prescriptions for treating the off-nominal detector systematics (e.g. in the form of the bin count gradients with respect to the systematic parameter variation, as in [3]) are supplied with this release. The per-event GENIE coefficients necessary to implement the QE and the RES cross section systematic uncertainties are similarly not available. We note that the latter will become particularly relevant for the IceCube-Upgrade, as the relative contributions of the CC-QE and CC-RES events to the overall 1–100 GeV sample will constitute ~21% and ~25%, respectively. Given these limitations, we choose a 3-year livetime for our projections involving the IceCube-Upgrade, which we expect to be statistics-limited. However, we stress that more accurate projections will be possible and necessary in the future once the remaining theoretical and experimental uncertainties are taken into account, along with the recent progress in the technical

development of the IceCube-Upgrade simulation [212].

7.2.2 *Reactor neutrino data*

- **Daya Bay**

The Daya Bay Reactor Neutrino Experiment (shortly: "Daya Bay", later abbreviated as "DB" in equations and figures) is probing disappearance of electron antineutrinos from the Daya Bay and the Ling Ao nuclear power plants. The $\bar{\nu}_e$ flux originates from the fission of the ^{235}U, ^{238}U, ^{239}Pu, and ^{241}Pu isotopes and is detected through the inverse beta decay (IBD) reaction, $\bar{\nu}_e + p \rightarrow e^+ + n$, in eight antineutrino detectors (ADs). Four of these detectors are located at the near experimental halls, EH1 and EH2, and the remaining four are placed at the far experimental hall, EH3. The distances between the reactor cores and the far detectors range between 1.5 km and 1.9 km, making Daya Bay a *short-baseline* detector. The IBD reaction products ionize the liquid scintillator filling the ADs, which leads to the emission of photons detected by the PMTs in each detector. This enables the reconstruction of the deposited "prompt energy":

$$E_{\text{prompt}} \approx E_{\bar{\nu}} + m_p - (m_n + m_e + T_n) + 2E_\gamma, \tag{7.5}$$

where $E_{\bar{\nu}}$ is the energy of the incoming antineutrino; m_p, m_n, and m_e are the masses of the proton, neutron, and electron, respectively; T_n is the kinetic energy of the neutron; and $2E_\gamma = 2 \times 0.511\,\text{MeV}$ is the energy of the two photons emitted as the result of the e^+e^- annihilation in the scintillator medium. The typical range of T_n is $\mathcal{O}(10\,\text{keV})$ [260] and is therefore neglected. This leads to the following approximation:

$$E_{\text{prompt}} \approx E_{\bar{\nu}} - 0.782\,\text{MeV}, \tag{7.6}$$

which we adopt in this study[6]. To probe the electron row of the mixing matrix N entering Equation (2.32), we employ a method similar to the Analysis Method B from [263], as was done in the most recent analysis by the Daya Bay Collaboration at the time of writing [49].

To predict the event rates at the far hall EH3, Method B relies on convolution of the reactor $\bar{\nu}_e$ flux model, the standard IBD cross section [264], the oscillation probabilities $P_{\bar{e}\bar{e}}(\bar{\lambda}_{\text{osc}})$, and the detector response. Here, $\bar{\lambda}_{\text{osc}}$ denotes the parameters driving the $\bar{\nu}_e$ survival probability, which for Daya Bay are the electron row elements N_{ei} and the mass splitting Δm^2_{32}. The reactor flux model is effectively constrained by the near hall measurements [263], and therefore is already contaminated by any non-unitarity effects. However, the IBD cross section used in this analysis is based on a calculation assuming the Standard Model and thus requires an extra correction factor of $N_e \equiv (NN^\dagger)_{ee}$. We include this factor in the effective oscillation probabilities $\hat{P}_{\bar{e}\bar{e}}(\bar{\lambda}_{\text{osc}})$, which become

$$\hat{P}^{\text{DB}}_{\bar{e}\bar{e}}(E_{\bar{\nu}_e}, L_r \,|\, \bar{\lambda}_{\text{osc}}) = P_{\bar{e}\bar{e}}(E_{\bar{\nu}_e}, L_r \,|\, \bar{\lambda}_{\text{osc}}) \cdot N_e = \frac{|(Ne^{-iHL_r}N^\dagger)_{ee}|^2}{N_e} \tag{7.7}$$

for a fixed baseline L_r connecting the far hall with one of the six reactors r. We further perform averaging over reactor baselines as follows:

$$\langle \hat{P}^{\text{DB}}_{\bar{e}\bar{e}}(E_{\bar{\nu}_e} \,|\, \bar{\lambda}_{\text{osc}}) \rangle_L = \left[\sum_r \frac{1}{L_r^2} \hat{P}^{\text{DB}}_{\bar{e}\bar{e}}(E_{\bar{\nu}_e}, L_r \,|\, \bar{\lambda}_{\text{osc}}) \right] \cdot \left[\sum_r \frac{1}{L_r^2} \right]^{-1}, \tag{7.8}$$

---

[6]We subsequently assume that the 26 prompt energy bins used in the official analyses by the Daya Bay Collaboration [49, 261] have a 1:1 correspondence with the 26 energy bins of the IBD neutrino spectrum from [260], which is an approximation to the complete energy response matrix of the detector (see e.g. [260, 262]).

such that the average oscillation probability $\langle \hat{P}_{\bar{e}\bar{e}}^{\mathrm{DB}}(E_{\bar{\nu}_e} | \bar{\lambda}_{\mathrm{osc}}) \rangle_L$ only depends on neutrino energy and the oscillation parameters.

In this study, we analyze the IBD candidate selection from 3158 days of Daya Bay data [49]. Together with the release of this dataset, the Daya Bay Collaboration provides the event rates predicted for EH3 with the best-fit oscillation parameters: $\Delta m_{32}^2 = (2.466 \pm 0.060) \cdot 10^{-3}\,\mathrm{eV}^2$ (under normal mass ordering) and $\sin^2 2\theta_{13} = 0.0851 \pm 0.0024$. We unfold this best-fit oscillation prediction to compute the expected unoscillated event rates at EH3:

$$\mathcal{N}_{\mathrm{EH3}}^{\mathrm{exp,\,no\,osc}}(E_{\mathrm{prompt}}) = \frac{\mathcal{N}_{\mathrm{EH3}}^{\mathrm{exp}}(E_{\mathrm{prompt}}, \bar{\lambda}_{\mathrm{osc}}^{\mathrm{best\,fit}})}{\langle \hat{P}_{\bar{e}\bar{e}}^{\mathrm{DB}}(E_{\bar{\nu}_e} | \bar{\lambda}_{\mathrm{osc}}^{\mathrm{best\,fit}}) \rangle_L}, \tag{7.9}$$

where $\bar{\lambda}_{\mathrm{osc}}^{\mathrm{best\,fit}}$ assumes unitary mixing ($NN^\dagger = N^\dagger N = \mathbb{I}$). We then use this unoscillated prediction to compute the expected spectrum $\mathcal{N}_{\mathrm{EH3}}^{\mathrm{exp}}$ under the tested non-unitary mixing hypothesis $\bar{\lambda}_{\mathrm{osc}}$:

$$\mathcal{N}_{\mathrm{EH3}}^{\mathrm{exp}}(E_{\mathrm{prompt}}, \bar{\lambda}_{\mathrm{osc}}) = \mathcal{N}_{\mathrm{EH3}}^{\mathrm{exp,\,no\,osc}}(E_{\mathrm{prompt}}) \cdot \langle \hat{P}_{\bar{e}\bar{e}}^{\mathrm{DB}}(E_{\bar{\nu}_e} | \bar{\lambda}_{\mathrm{osc}}) \rangle_L. \tag{7.10}$$

Finally, the expected spectrum $\mathcal{N}_{\mathrm{EH3}}^{\mathrm{exp}}$ is compared to the observed spectrum $\mathcal{N}_{\mathrm{EH3}}^{\mathrm{obs}}$ at EH3. To optimize for the parameters of $\bar{\lambda}_{\mathrm{osc}}$, the following χ^2 test statistic is constructed[7]:

$$\chi_{\mathrm{DB}}^2(\bar{\lambda}_{\mathrm{osc}}) = (\mathcal{N}_{\mathrm{EH3}}^{\mathrm{obs}} - \mathcal{N}_{\mathrm{EH3}}^{\mathrm{exp}})^\top V_{\mathrm{tot}}^{-1} (\mathcal{N}_{\mathrm{EH3}}^{\mathrm{obs}} - \mathcal{N}_{\mathrm{EH3}}^{\mathrm{exp}}), \tag{7.11}$$

where V_{tot} is the total covariance matrix consisting of the statistical and the systematic components, $V_{\mathrm{tot}} = V_{\mathrm{stat}} + V_{\mathrm{syst}}$. The statistical component of V is calculated as a diagonal matrix with $V_{\mathrm{stat},ii} = \mathcal{N}_{\mathrm{EH3},i}^{\mathrm{exp}}$, where index i runs over the prompt energy bins $E_{\mathrm{prompt},i}$. The estimation of the systematic component follows the prescription from [260] and is detailed in Appendix B.2.1[8]. As in Section 7.2.1, the χ^2 from Equation (7.11) is used to reproduce the three-flavor Daya Bay contours in Section 7.3, and converted to the log-likelihood $\ln \mathcal{L}_{\mathrm{DB}}$ analogously to Equation (7.4) for the Bayesian analysis of the non-unitary mixing.

In Figure 7.3, we show the predicted prompt energy spectrum at EH3 assuming no oscillations, $\mathcal{N}_{\mathrm{EH3}}^{\mathrm{exp,\,no\,osc.}}$; the actual spectrum measured by Daya Bay in 3158 days of livetime, $\mathcal{N}_{\mathrm{EH3}}^{\mathrm{obs}}$; and the corresponding best-fit spectrum, $\mathcal{N}_{\mathrm{EH3}}^{\mathrm{exp}}(\bar{\lambda}_{\mathrm{osc}}^{\mathrm{best\,fit}})$. All of the spectra are background-subtracted and summed over the data-taking periods. To illustrate the effect of non-unitarity on the predicted event rate, we scale the electron row of the mixing matrix N such that the normalization $N_e = 0.95$, i.e., 5% lower than the unitary expectation. We see from Figure 7.3 that the deviation from the measured Daya Bay spectrum in the non-unitary case is highly non-statistical, which suggests that the constraint on the electron row normalization that could be put with Daya Bay will be much stronger than 5%. We note that this is the result of the IBD cross section correction by one factor of N_e. This factor would cancel out if one were to use an analysis method based on computing the far/near detector event ratios, such as the Method A of Ref. [263], resulting in the loss of sensitivity to the electron row normalization with Daya Bay alone.

---

[7]We note that the original version of the Analysis Method B from [263] relies on profiling of the systematic parameters with respective penalty terms added to the χ^2. However, since we do not have access to the Daya Bay simulation chain and implementaton of the individual systematics, we resort to the covariance matrix method instead.

[8]The ultimate impact of the non-zero V_{syst} on the measured electron row elements is shown in Figure B.3.
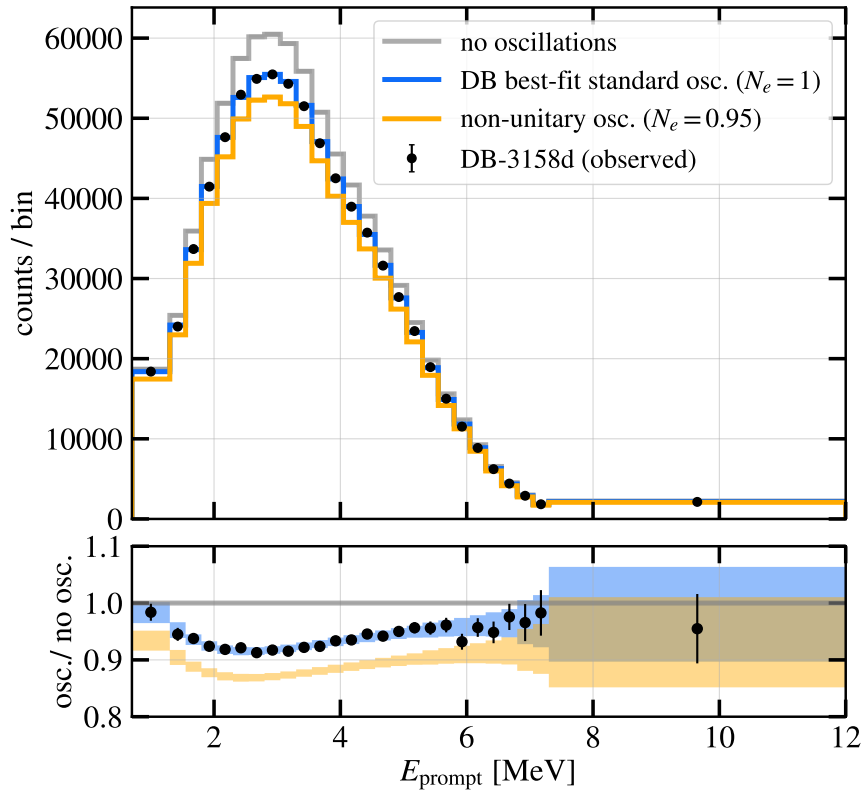
FIGURE 7.3 – *Top*: Daya Bay prompt energy spectrum predicted at the far hall EH3 for the case of no oscillations (gray), Daya Bay best-fit three-flavor oscillations (blue) [49], and non-unitary oscillations with normalization of the electron row $N_e = 0.95$ (yellow). The black dots represent the spectrum observed by Daya Bay at EH3 in 3158 days. *Bottom*: ratios of the expected or observed spectra including oscillations to the no-oscillation prediction. The color scheme follows that of the top panel. The errorbars represent the statistical uncertainty for data and the square-roots of the diagonal elements of the covariance matrix ($\sqrt{V_{ii}}$) for expectations.

● **KamLAND**

The concept of the Kamioka Liquid Scintillator Antineutrino Detector ("KamLAND", later abbreviated as "KL" in equations and figures) is similar to that of Daya Bay. KamLAND measures electron antineutrinos from nuclear reactors via the IBD process in a single tank filled with liquid scintillator, and the light emitted post-interaction due to the scintillation and Cherenkov radiation processes is collected by PMTs. The main difference compared to Daya Bay is that there are more than 50 (57 in this study) contributing nuclear reactors located at much larger distances to the detector, with the effective flux-averaged distance $L_0 \simeq 180$ km. This means that KamLAND can probe ~ 100 times smaller mass splittings compared to Daya Bay, which makes it sensitive to Δm_{21}^2 and $N_{e\{1,2\}}$ (alternatively, θ_{12} in the mixing angle parameterization). Furthermore, due to both long baselines and the depth at which the KamLAND detector is located (2700 m w.e.), matter effects start playing a role in neutrino oscillations and have to be taken into account explicitly in the Hamiltonian of Equation (2.22).

In this study, we analyze the selection of IBD candidates collected by KamLAND over 7 years (2002–2009), as reported in [50, 265]. From the referenced publication, we extract the no-oscillation prediction for the binned prompt energy spectrum at the detector and subtract the best-fit background distributions. We assume that this energy spectrum directly corresponds to the energy deposited by the IBD interaction products, i.e., we do not apply any deposited-to-visible energy conversions due to non-linearity

of the detector response[9]. However, we take into account the 1.9% uncertainty on the energy scale determination (ϵ_E) and the 4.1% bin-to-bin correlated uncertainty on the event rate ($\epsilon_\mathcal{N}$), which are the dominant systematic uncertainties for the determination of Δm_{21}^2 and θ_{12} [50, 267][10].

We compute the expected energy spectrum under the oscillation hypothesis $\bar{\lambda}_{\mathrm{osc}}$ as

$$\mathcal{N}_{\mathrm{exp}}(E_{\mathrm{prompt}}, \bar{\lambda}_{\mathrm{osc}}, \bar{\lambda}_{\mathrm{syst}}) = \mathcal{N}_{\mathrm{exp,noosc}}(E_{\mathrm{prompt}}, \bar{\lambda}_{\mathrm{syst}}) \langle \hat{P}_{\bar{e}\bar{e}}^{\mathrm{KL}}(E_{\bar{\nu}_e} | \bar{\lambda}_{\mathrm{osc}}) \rangle_L, \tag{7.12}$$

where $\langle \rangle_L$ denotes averaging over the baselines as in Equation (7.8) and $\bar{\lambda}_{\mathrm{syst}} = (\epsilon_E, \epsilon_\mathcal{N})$ are the systematic parameters. To arrive at the effective survival probability $\hat{P}_{\bar{e}\bar{e}}^{\mathrm{KL}}$ for KamLAND, the standard oscillation probability is multiplied by two factors of N_e. This is needed to correct for non-unitarity in both the IBD cross section and the reactor flux, which are based on calculations assuming the Standard Model rather than calibration at the near detector [264, 268]. This removes the factors of N_e from the denominator of Equation (2.32) and yields

$$\hat{P}_{\bar{e}\bar{e}}^{\mathrm{KL}}(E_{\bar{\nu}_e}, L_r | \bar{\lambda}_{\mathrm{osc}}) = |(N e^{-iHL_r} N^\dagger)_{ee}|^2 \tag{7.13}$$

for a fixed baseline L_r.

Due to the relatively low numbers of counts in the KamLAND spectra, we are using the binned Poisson likelihood[11] to quantify the agreement of the observed data with the expected spectrum under a given oscillation hypothesis:

$$\ln \mathcal{L}_{\mathrm{KL}}(\bar{\lambda}_{\mathrm{osc}}, \bar{\lambda}_{\mathrm{syst}}) = \mathcal{N}_{\mathrm{obs}} \cdot \ln\left[\mathcal{N}_{\mathrm{exp}}(\bar{\lambda}_{\mathrm{osc}}, \bar{\lambda}_{\mathrm{syst}})\right] - \mathcal{N}_{\mathrm{exp}}(\bar{\lambda}_{\mathrm{osc}}, \bar{\lambda}_{\mathrm{syst}}) - \ln[\mathcal{N}_{\mathrm{obs}}!]. \tag{7.14}$$

The penalty terms due to $\bar{\lambda}_{\mathrm{syst}}$ are added in the same way as in Equation (7.3).

In Figure 7.4, we show the predicted prompt energy spectrum assuming no oscillations, the actual spectrum measured by KamLAND over 7 years, and the spectrum expected at the KamLAND best-fit point assuming $\theta_{13} = 0°$ ($\Delta m_{21}^2 = 7.5 \cdot 10^{-5}\,\mathrm{eV}^2$, $\tan^2 \theta_{12} = 0.492$) [50]. As was done in Figure 7.3, we also provide the corresponding prediction under the assumption of non-unitary mixing, where we rescale all of the elements of the electron row such that $N_e = 0.95$. We see that due to limited statistics, KamLAND is not nearly as sensitive to the electron row normalization as Daya Bay, and therefore only a weak improvement to the Daya Bay constraint on N_e is expected from inclusion of KamLAND. However, KamLAND provides an additional handle on the N_{e1} and the N_{e2} elements of the mixing matrix, which enter the KamLAND oscillation probability in a different combination ($|N_{e1}|^2 |N_{e2}|^2$) than that of Daya Bay ($|N_{e1}|^2 + |N_{e2}|^2$) [53]. Thus, we include KamLAND into the analysis for the purpose of constraining $|N_{e1}|$ and $|N_{e2}|$ individually. At the same time, we note that this cannot be done completely unambibuously without the solar neutrino oscillation data (such as that from SNO [269]), which helps break the degeneracy in the octant of θ_{12}. We therefore use prior information from the solar neutrino measurements that θ_{12} lies in the lower octant, such that $|N_{e1}| > |N_{e2}|$, which helps avoid the bimodality in our ultimate matrix element posteriors.

---

[9] We investigated this option in our analysis, applying the $E_{\mathrm{vis}}(E_{\mathrm{dep}})$ correction as derived in [266], but found nearly no change in our reproduction of the standard oscillation results.

[10] The impact of these uncertainties on the measured electron row elements is shown in Figure B.4.

[11] We note that this treatment differs from the unbinned and time-dependent likelihood approach officially used by the KamLAND collaboration [50, 267]; however, since we do not have access to the event-by-event information from KamLAND, we find this approach to be optimal given the public information available.
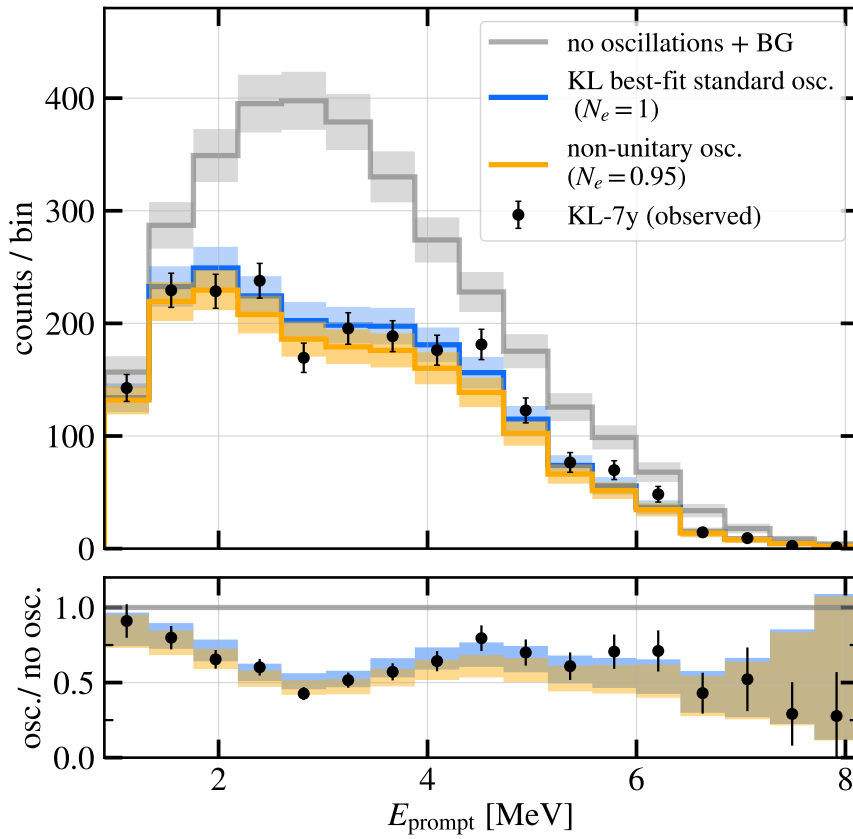
FIGURE 7.4 – *Top*: KamLAND prompt energy spectrum predicted for the case of no oscillations (gray), KamLAND best-fit two-flavor oscillations (blue) [50], and non-unitary oscillations with normalization of the electron row $N_e = 0.95$ (yellow). The black dots represent the spectrum observed by KamLAND between 2002-2009. *Bottom panel*: same as Figure 7.3. The errorbars represent the statistical uncertainty for data and the total (statistical + systematic) uncertainty on the event rate for expectations.

- ## JUNO

The Jiangmen Underground Neutrino Observatory ("JUNO") is an upcoming multipurpose neutrino experiment, which is undergoing installation in Southern China at the time of writing [255, 270]. It will measure, in particular, the disappearance of reactor $\bar\nu_e$ from two nuclear power plants located at 52.5 km distance each from a liquid scintillator detector. The combination of this baseline length with the typical energy range of the IBD spectrum (~1.8–9 MeV) results in a sufficiently large L/E reach suitable for measuring the slow oscillations due to $\sin^2\theta_{12}$ and Δm_{21}^2. At the same time, the energy deposited by the IBD interaction products will be measured in JUNO with very high resolution (~3% at 1 MeV using the main JUNO PMT system, LPMT [255]). This will allow for simultaneous measurements of Δm_{31}^2 and $\sin^2\theta_{13}$, which induce fast oscillations.

The JUNO Collaboration provides the prediction for the unoscillated event rates $\mathcal{N}_{\mathrm{exp,noosc}}$ as a function of true neutrino energy, which results from a convolution of the reactor fluxes with the IBD cross section [255]. Since the assumed cross section is following the Standard Model-based calculation by [264], we correct it by the normalization N_e as per Equation (2.33b). This correction factor is absorbed into the effective oscillation probability $\hat{P}_{ee}$ as in Equation (7.7). The reactor flux, on the other hand, does not require a non-unitarity correction, since its calibration at the satellite TAO detector is envisioned [255, 271].

To represent the expected oscillated event rates $\mathcal{N}_{\mathrm{exp,noosc}}(E_\nu)\cdot\hat{P}_{ee}(L,E_\nu)$ as a function

of the visible energy E_{vis} in the detector, we apply the shift from E_ν to the prompt (positron) energy as per Equation (7.6), which we further correct according to the nonlinear response of the LPMT system and smear according to its expected resolution. Following [255], we use a 20 keV bin width for the resulting visible energy range. For the final oscillated spectrum expectation, we assume a 2.2% flux uncertainty ($\epsilon_\mathcal{N}$) and a 1% detection uncertainty (ϵ_{eff}), both bin-to-bin correlated. We further include the geoneutrino background, which is the dominant source of background in JUNO [255], as well as an associated 30% uncertainty (ϵ_{geo}). We do not take into account other sources of background or spectrum shape uncertainties in this analysis. The $\epsilon_\mathcal{N}$, ϵ_{eff}, and ϵ_{geo} systematic parameters acquire respective penalty terms in the χ^2_{JUNO} test statistic. The latter is equivalent to Equation (7.3), modulo the Monte Carlo simulation uncertainty σ_i, which only applies to DeepCore. As in all other experiments for which the χ^2 metric is defined, we convert it to the log-likelihood $\ln \mathcal{L}_{JUNO}$ assuming Wilks' theorem.

In Figure 7.5, we show our predictions for the visible energy spectra in JUNO after 6 years of exposure for the same physics cases as in Figures 7.3 and 7.4. By comparing the spectrum generated under the standard three-flavor scenario to the non-unitary mixing case with $N_e = 0.95$, we can expect that JUNO will similarly have the capability to constrain the electron row normalization but not improve significantly on the Daya Bay-only constraints.
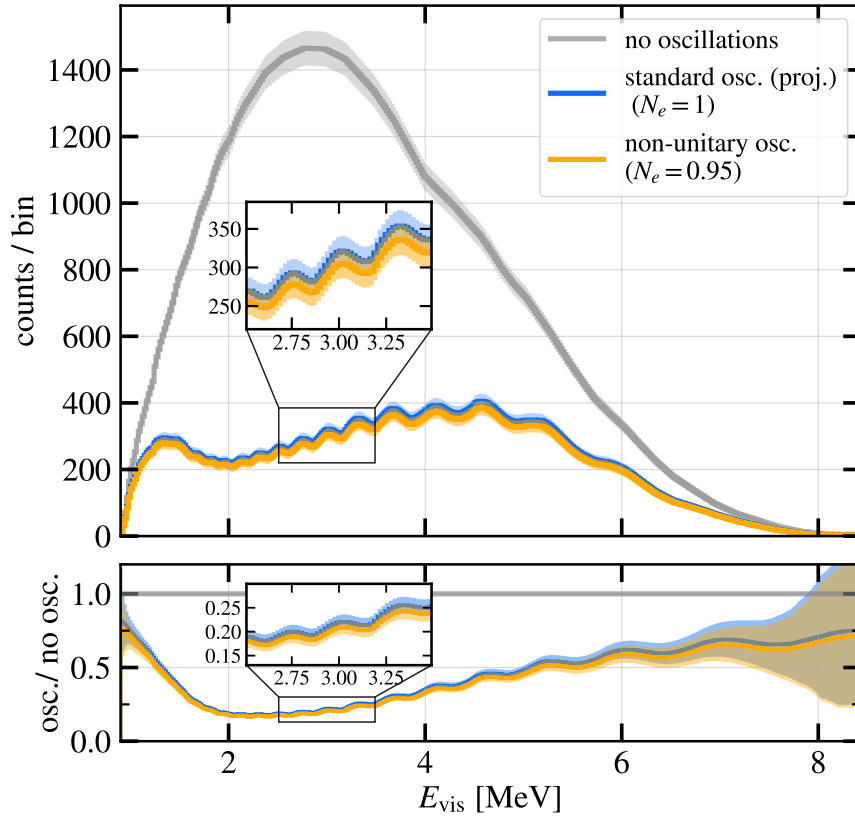

FIGURE 7.5 – *Top*: JUNO visible energy spectrum predicted for the case of no oscillations (gray), best-fit oscillations with parameters from [255], and non-unitary oscillations with normalization of the electron row $N_e = 0.95$ (yellow). *Bottom*: same as Figure 7.3. The errorbars represent the total (statistical + systematic) uncertainty on the expected event rate. The backgrounds are not included in this figure. The assumed livetime is 6 years.

7.3 REPRODUCTION OF THE THREE-FLAVOR OSCILLATION RESULTS

Prior to fitting for the elements of the non-unitary mixing matrix N, we validate our setup of the experiments, systematic uncertainties, and test statistics by reproducing the standard three-flavor oscillation results. In particular, we construct the following confidence level contours: ($\sin^2 2\theta_{13}$, Δm^2_{32}) for Daya Bay; ($\tan^2 \theta_{12}$, Δm^2_{21}) for Kam-LAND; ($\sin^2 \theta_{12}$, Δm^2_{21}) for JUNO; and ($\sin^2 \theta_{23}$, Δm^2_{32}) for IceCube-DeepCore and IceCube-Upgrade. The contours are derived from the frequentist χ^2 scans across the relevant 2D parameter spaces and compared to the official results from the respective collaborations [46, 49, 50, 205, 255]. The solar mass splitting and mixing angle are fixed at $\Delta m^2_{21} = 7.53 \cdot 10^{-5} \, \text{eV}^2$ ($7.5 \cdot 10^{-5} \, \text{eV}^2$) and $\theta_{12} = 33.48° \, (33.64°)$ for the Daya Bay (DeepCore) analysis. For future projections with JUNO and IceCube-Upgrade, the injected true parameters are ($\sin^2 \theta_{12}$, Δm^2_{21}) = (0.307, $7.53 \cdot 10^{-5} \, \text{eV}^2$) and ($\sin^2 \theta_{23}$, Δm^2_{32}) = (0.51, $2.31 \cdot 10^{-3} \, \text{eV}^2$), respectively, following [46, 255]. The DeepCore systematic parameters $\bar{\lambda}_{\text{syst}}$ [46] are profiled over, i.e., fitted by minimizing $\chi^2_{\text{DC}}(\bar{\lambda}_{\text{osc}}, \bar{\lambda}_{\text{syst}})$ at a given scan point. The same treatment applies to the simplified KamLAND and JUNO systematic parameters. For reproduction of the KamLAND results, we additionally fix $\theta_{13} = 0°$[12].

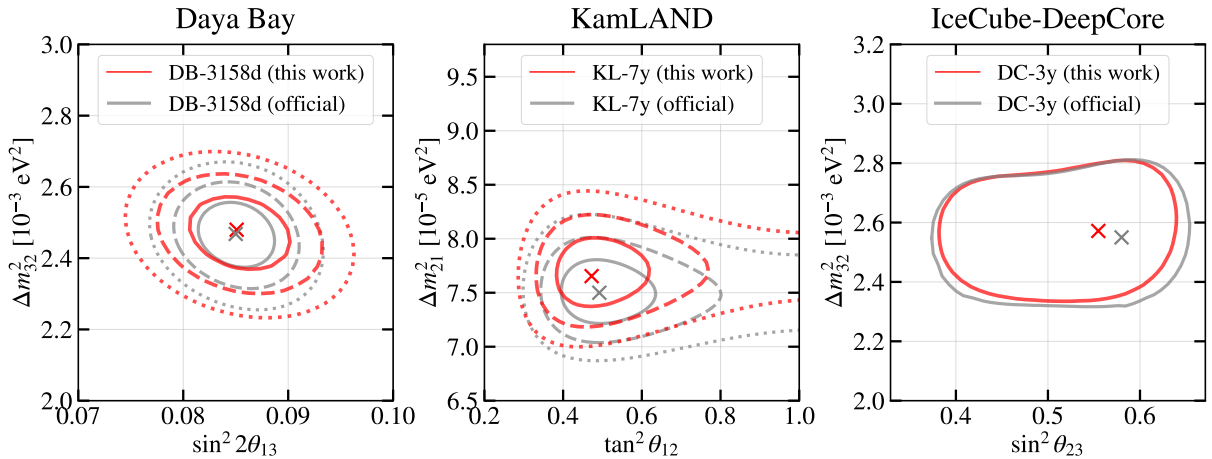Our results are summarized in Figures 7.6 and 7.7.

FIGURE 7.6 – Reproduction of the standard oscillation results from Daya Bay [49], KamLAND [50], and IceCube-DeepCore [46] experiments. For Daya Bay and KamLAND, the solid, dashed, and dotted contours indicate 1σ, 2σ, and 3σ confidence levels. For IceCube-DeepCore, the contours correspond to the 90% confidence level. The crosses represent the best fit points from this work (red) or the published experimental result (gray). The official IceCube contours were extracted from [46] (2019 result).

Generally, we find a good agreement between the official collaboration contours and our reproductions, and note the following:

- For Daya Bay, our recovered 2σ and 3σ contours are 30–35% wider than the official result of the Daya Bay collaboration. This level of agreement is expected, as we implemented an approximate covariance matrix with the simplifications detailed in Appendix B.2.1, rather than fully profiling the systematic parameters as done in [49].

- For KamLAND, the best-fit mass splitting Δm^2_{21} obtained in our reproduction is 2% (0.5σ) higher than the collaboration reported-value, and the best-fit $\tan^2 \theta_{12}$ is 4% (0.2σ) lower than the official KamLAND result. All three confidence levels

---

[12]Note that we fix θ_{13} only for the one-time comparison with the KamLAND-provided χ^2 map at $\theta_{13} = 0°$ [265], and do not rely on this simplification anywhere else in this analysis.
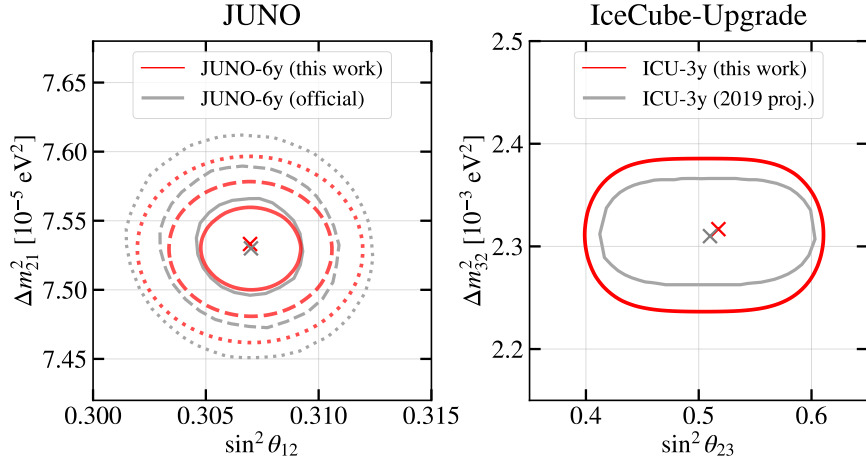
FIGURE 7.7 – Reproduction of the projected standard oscillation contours for the JUNO [255] and the IceCube-Upgrade [205] experiments. For JUNO, the solid, dashed, and dotted contours indicate 1σ, 2σ, and 3σ confidence levels. For IceCube-Upgrade, the contours correspond to the 90% confidence level. The crosses represent the best fit points from this work (red) or the injected truth (gray).

contours are similarly shifted to higher Δm_{21}^2 and lower $\tan^2\theta_{12}$. We attribute this to the difference in the likelihood implementations, namely the binned and time-independent Poisson likelihood used in this study as opposed to the unbinned likelihood incorporating the time-dependent reactor power, which was used in the official KamLAND fit [50].

- The $(\sin^2\theta_{23}, \Delta m_{32}^2)$ contours derived in this work are in a good agreement with the official IceCube result [46]. We achieve a nearly perfect match of the contour widths in the Δm_{32}^2 dimension and $\lesssim 20\%$ discrepancy in the $\sin^2\theta_{23}$ dimension, with our contours being narrower. This could be connected with the cross section systematic uncertainties being omitted in our analysis and the resulting shift of the best-fit point by $\Delta\sin^2\theta_{23} = 0.026$.

- The 2σ and 3σ contours obtained in the $(\sin^2\theta_{12}, \Delta m_{21}^2)$ space for JUNO are underestimated only at 0.1%-0.2% level in our reproduction compared the official JUNO projection [255]. This is expected since we considered only geoneutrinos as a source of background in JUNO, and wider contours would be obtained with more backgrounds and associated systematic uncertainties incorporated into the analysis.

- The 90% confidence level contours derived in this work for the IceCube-Upgrade are 13–18% wider in the $\sin^2\theta_{23}$ dimension and 40–50% wider in the Δm_{32}^2 dimension, as compared to the official projections from [205]. This is caused by the different form of the test statistic, such that the uncertainty due to limited Monte Carlo statistics is taken into account in this work but not in [205]. We confirmed that the contours match exactly if the Monte Carlo statistics uncertainty term is excluded from our test statistic (see Equation (7.2)).

7.4 GLOBAL FIT PROCEDURE

For the primary analysis of this work, we take a Bayesian approach to constrain the individual elements of the leptonic mixing matrix, including their magnitudes and phases. We combine the likelihoods defined in Section 7.2 for each experiment with the uniform $[0, 2\pi]$ prior on each of the fitted phases $\phi_{\alpha i}$ and the uniform $[0, 1]$ prior on most of the fitted magnitudes $|N_{\alpha i}|$. The exceptions are $|N_{e1}|$, $|N_{e2}|$, which we constrain to be in the $[0.7, 0.9]$ and $[0.4, 0.7]$ ranges, respectively, using the prior

information from the solar oscillation experiments as discussed in Section 7.2.2. The constraints for the mass splittings are $\Delta m_{21}^2 \in [6.5 \cdot 10^{-5}\,\mathrm{eV}^2, 9.0 \cdot 10^{-5}\,\mathrm{eV}^2]$ and $\Delta m_{31}^2 \in [1.9 \cdot 10^{-3}\,\mathrm{eV}^2, 2.9 \cdot 10^{-3}\,\mathrm{eV}^2]$. Normal neutrino mass ordering is assumed throughout this study[13]. The priors on the IceCube-DeepCore and the IceCube-Upgrade systematic parameters are defined in Appendix B.2.3.

Our fit is performed using the ULTRANEST [272, 273] package, which utilizes the MLFriends Monte Carlo nested sampling algorithm [274, 275]. We couple the likelihood evaluations to the oscillation probability calculations within the NEURTHINO package [5], which we modified to include the physics specific to non-unitarity (see Section 2.4).

When making sensitivity projections for both the current and the next-generation experiments, we use the unitary mixing matrix U generated with the NuFit 5.2 oscillation parameters [65, 66] as the truth. We then constrain the individual matrix elements of N by fitting the full non-unitary model to the fake data templates produced with U.

---

[13]We tested the possibility of fitting the normal mass ordering model to pseudodata with injected inverted ordering and found virtually identical sensitivities to the mixing matrix elements, normalizations, and closures with the current set of experiments (IceCube-DeepCore, Daya Bay, KamLAND).

Results

N.B.: During the preparation of this thesis, which coincides with the journal review period, a misinterpretation of the KamLAND data used in the analysis was found. The results are being recreated with a correct data implementation, and updated results will be available for the defense, next draft of this thesis, and for the journal article. The updated result is not expected to significantly change the scientific conclusions regarding the impact of the KamLAND data on the unitarity global fits.

8.1 GLOBAL FIT WITH THE CURRENT DATA

8.1.1 *Constraints on matrix elements, normalizations, and closures*

In Figure 8.1, we show the posterior densities of each of the matrix element magnitudes $|N_{\alpha i}|$ obtained with the available public data from IceCube-DeepCore, Daya Bay, and KamLAND experiments. We see that the electron row elements are significantly better constrained than the muon and the tau row elements. In particular, the $\sim$6% precision of the measured $|N_{e3}|$ is driven by the extremely high statistics of the $\bar{\nu}_e$ neutrino sample detected by Daya Bay, while the constraints on $|N_{e1}|$ and $|N_{e2}|$ are the result of combining the Daya Bay data with that of KamLAND and the prior knowledge from the solar neutrino experiments about the octant of θ_{12}. The primary advantage of including the IceCube-DeepCore data into the global fit lies in the measurement $|N_{\mu3}|$ and $|N_{\tau3}|$ elements, which we find to be constrained with $\sim$30–40 % precision given the three-year dataset. Thanks to the constraints on the overall row normalizations (discussed further in this section), the $|N_{\mu\{1,2\}}|$ and the $|N_{\tau\{1,2\}}|$ elements could also be measured, albeit with large (up to 100%) uncertainties and strong degeneracies between the elements of the first and the second column in each row.

From the posterior densities of the matrix element magnitudes, we construct the posteriors of the row normalizations, N_α ($\alpha \in \{e, \mu, \tau\}$), and the column normalizations, N_i ($i \in \{1, 2, 3\}$). The results are shown in Figures 8.2 and 8.3, respectively. The 68%, 95%, and 99.7% credible intervals derived from the 1D normalization posteriors are additionally reported in Table 8.1.
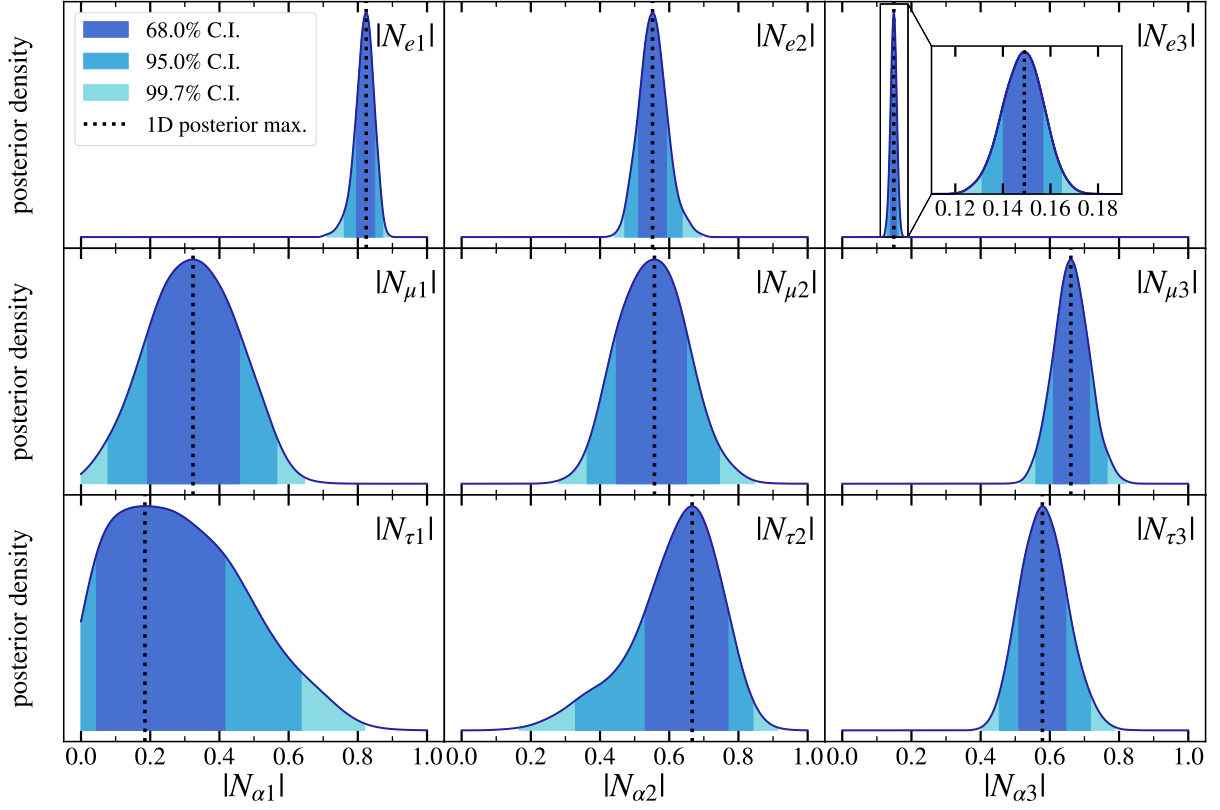
FIGURE 8.1 – Posterior densities of the non-unitary matrix element magnitudes, $|N_{\alpha i}|$, as obtained with the current data from IceCube-DeepCore, Daya Bay, and KamLAND. All other fitted physics and nuisance parameters are marginalized over. The three shaded regions correspond to 68%, 95%, and 99.7% credible intervals (C.I.), from darker to lighter.

	68% C.I.	95% C.I.	99.7% C.I.
N_e	[0.993, 1.009]	[0.985, 1.018]	[0.977, 1.026]
N_μ	[0.76, 0.97]	[0.67, 1.10]	[0.57, 1.24]
N_τ	[0.73, 0.96]	[0.62, 1.09]	[0.52, 1.24]
N_1	[0.75, 0.98]	[0.67, 1.20]	[0.59, 1.50]
N_2	[0.84, 1.19]	[0.66, 1.36]	[0.53, 1.53]
N_3	[0.70, 0.90]	[0.61, 1.02]	[0.52, 1.14]

TABLE 8.1 – Credible intervals (C.I.) for the 1D posteriors of the row normalizations (first three rows) and the column normalizations (last three rows) of the unitarity-agnostic neutrino mixing matrix N, as measured with the current public data from IceCube-DeepCore, Daya Bay, and KamLAND.

We find that the electron row normalization is well constrained by the current data, such that the 99.7% credible interval covers only a 5% range around the unitary expectation of $N_e = 1$. We remind the reader that this constraint comes predominantly from the Daya Bay data and is the result of our specific choice of the analysis where the input IBD cross section is based on the Standard Model calculation from [264]. This necessitates the correction of the Daya Bay oscillation probabilities by one factor of N_e, as per Equation (7.7), which would have cancelled out in an analysis relying on near-to-far detector event ratios but does not cancel out in a direct calculation of the expected event rates with a SM cross section. Our constraint on N_e therefore represents the statistical power of Daya Bay to constrain the $\bar{\nu}_e$ spectrum normalization affected by the non-unitarity effects on the cross section. We point out that if no extra factors of N_e were modifying the effective oscillation probability, the non-unitarity in the electron
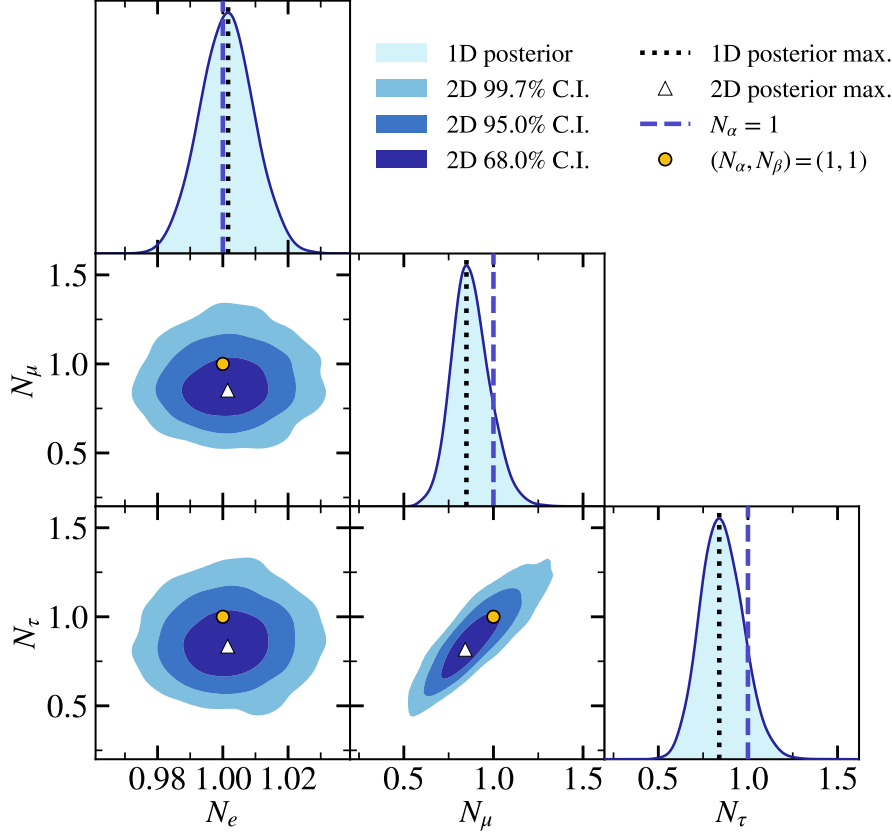
FIGURE 8.2 – Diagonal elements: 1D posterior densities of the row normalizations N_α, where the dotted black line corresponds to the measured posterior maximum, and the dashed blue line – to the normalization expected under the assumption of unitarity. Off-diagonal elements: correlations between the normalizations of different rows (N_α, N_β), where the triangular marker is the maximum of the 2D posterior in the (N_α, N_β) space, and the circular marker is the unitary expectation.

row would only be observed by KamLAND due to the presence of non-standard matter effects (see top left panels of Figures 2.5 and B.1), but would not be observed in Daya Bay.

Further, we observe that the normalization of the muon and the tau rows have much wider contours for the same choice of the credible intervals, covering ∼70% ranges about the unitary expectation. These ranges are asymmetric, with ∼88% of the one-dimensional N_μ and N_τ posteriors contained below 1, and the respective posterior maxima situated at 0.85[1]. We note a strong positive correlation between the N_μ and N_τ posterior densities, as well as between N_2 and N_3. In Section 8.2, we reveal that these correlations are driven by the atmospheric neutrino systematic uncertainties and would not have existed if the atmospheric neutrino flux was perfectly constrained. On the contrary, the negative correlation between the N_1 and N_2 column normalizations is the result of the degeneracies between $|N_{\alpha 1}|$ and $|N_{\alpha 2}|$ ($\alpha \in [\mu, \tau]$), which arise due to the fact that the atmospheric oscillation probabilities only depend on the sum of squares of these elements [53].

Finally, we make use of the fitted complex phases $\phi_{\alpha i}$ in combination with the matrix element magnitudes $|N_{\alpha i}|$ to derive the constraints on the row closures $t_{\alpha\beta}$ and the column closures t_{ij}, as defined in Section 2.4. The 2D posteriors in the complex $(\mathfrak{Re}, \mathfrak{Im})$ planes for each of the closures are given in Figure 8.4. Most of the closure posteriors

---

[1]For reference, the value of the ν_τ normalization reported in the official IceCube-DeepCore tau neutrino appearance result was $0.73^{+0.30}_{-0.24}$ ($0.57^{+0.36}_{-0.30}$) and obtained through a simple scaling of the expected number of the ν_τ-CC (CC+NC) events [46].
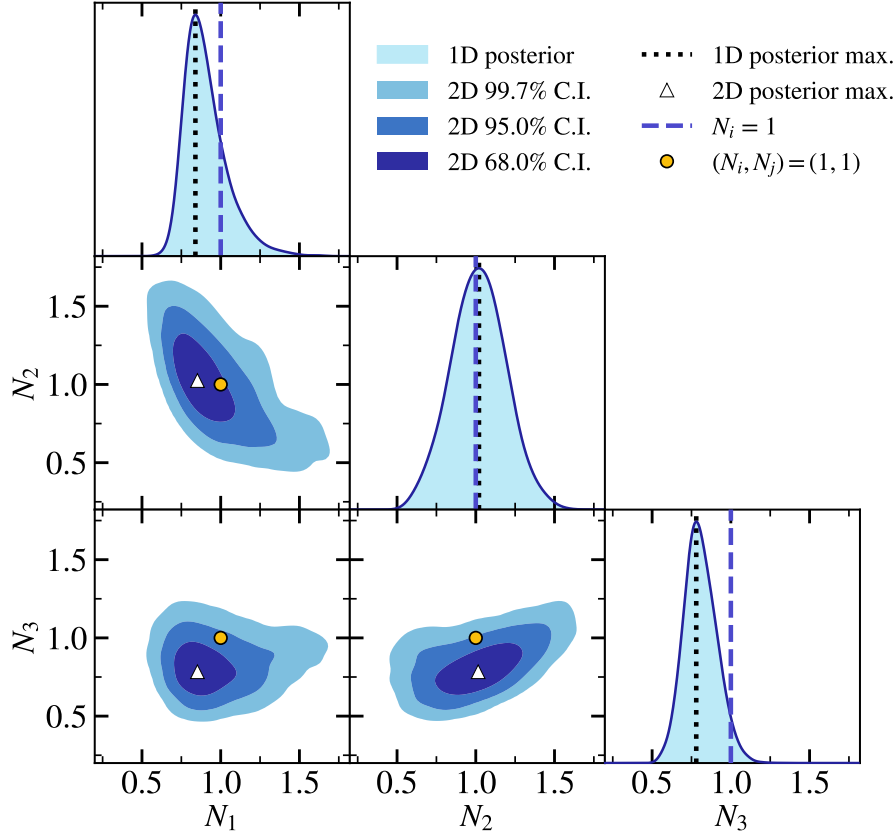
FIGURE 8.3 – Same as Figure 8.2, but applied to the matrix columns (i, j) instead of the matrix rows (α, β).

enclose the unitary expectation of $\mathfrak{Re}(t_{...}) = \mathfrak{Im}(t_{...}) = 0$ within the 68% credible interval, except for t_{12}, which contains the unitary expectation within the 95% credible interval. Thus, our measurements of the row and column closures demonstrate consistency with unitarity. The best constraints are obtained for $t_{\mu\tau}$, such that the 99.7% credible interval for $|t_{\mu\tau}|$ is contained within the radius of 0.2 from the unitary expectation.

8.1.2 Comparison with the model assuming unitarity

In Section 8.1, we showed that the data appears to prefer smaller-than-unity normalizations N_μ, N_τ, N_1, and N_3, when fitted with a unitarity-agnostic model relying on a generic mixing matrix N. We now seek to contrast this model with the one that assumes unitarity, i.e., the standard three-flavor neutrino mixing through the PMNS matrix as defined in Equation (2.16). For that purpose, we run an additional Bayesian fit for the unitary scenario and use ULTRANEST to extract the Bayesian evidence $\mathcal{Z}$ for both models. The evidence is computed as

$$\mathcal{Z} = \int \mathcal{L}(\mathcal{N}_{\text{obs}} | \bar{\lambda}_{\text{osc}}, \bar{\lambda}_{\text{syst}}) \pi(\bar{\lambda}_{\text{osc}}, \bar{\lambda}_{\text{syst}}) \mathrm{d}\bar{\lambda}_{\text{osc}} \mathrm{d}\bar{\lambda}_{\text{syst}}, \tag{8.1}$$

where $\mathcal{L}(\mathcal{N}_{\text{obs}} | \bar{\lambda}_{\text{osc}}, \bar{\lambda}_{\text{syst}})$ is the product of the per-experiment likelihoods combined in the global fit, $\pi(\bar{\lambda}_{\text{osc}}, \bar{\lambda}_{\text{syst}})$ are the priors for the oscillation (systematic) parameters $\bar{\lambda}_{\text{osc}}$ ($\bar{\lambda}_{\text{syst}}$), and $\mathcal{N}_{\text{obs}}$ are the observed data. Having obtained the evidence values $\mathcal{Z}_{\text{NU}}$ and $\mathcal{Z}_{\text{SM}}$ for the non-unitary and the unitary ("Standard Model") scenarios, we compute the logarithm of the Bayes factor as

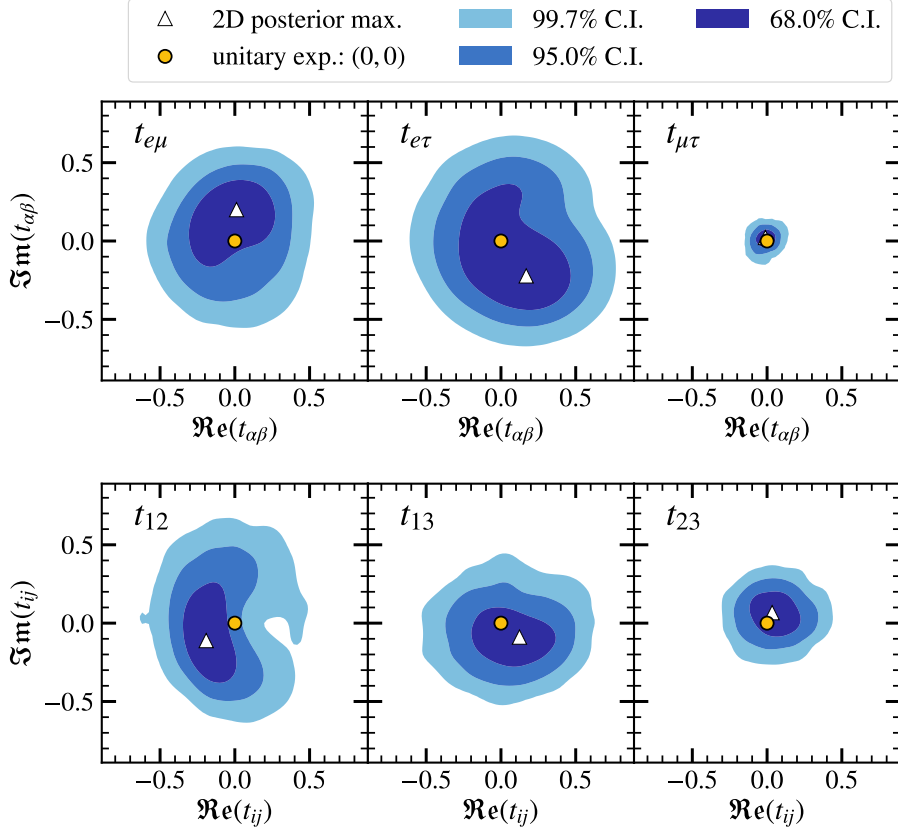$$\ln \mathcal{B} = \ln \mathcal{Z}_{\text{SM}} - \ln \mathcal{Z}_{\text{NU}}, \tag{8.2}$$

FIGURE 8.4 – *Top*: posteriors of the row closures $t_{\alpha\beta}$ obtained with the current data from IceCube-DeepCore, Daya Bay, and KamLAND, and displayed in the 2D complex planes. *Bottom*: same as top, but applied to the column closures t_{ij}.

such that a negative (positive) value of $\ln\mathcal{B}$ indicates a preference for the non-unitary (unitary) model. We obtain $\ln\mathcal{Z}_{\mathrm{NU}} = -120.04$ and $\ln\mathcal{Z}_{\mathrm{SM}} = -105.18$, resulting in $\ln\mathcal{B} = 14.86$. According to the Jeffreys' scale [276, 277], this implies a strong preference for the unitary model over the non-unitary one. This finding is in contrast with the fact that the fit with the generic mixing matrix N prefers off-unitary normalizations for certain rows and columns, as indicated above. Given that the non-unitary model introduces 9 extra oscillation parameters compared to the unitary case, the resulting Bayes factor can simply reflect the relative complexity of the models – i.e., the non-unitary model being unnecessarily complex. In addition, the non-unitary model can be penalized due to the rather agnostic priors on the matrix element magnitudes and phases (see Section 7.4), and the dependence of the evidence values on the choice of the prior ranges was left outside the scope of this work. To get further insight into whether the best-fit unitary and the best-fit non-unitary models are individually compatible with the data, we run posterior predictive checks for each model as described in Section 8.1.3.

8.1.3 *Posterior predictive checks*

To quantify the goodness of fit for both the unitary and the non-unitary models tested in this study, we use the method of posterior predictive checks via realized discrepancies, as devised in [278]. The purpose of the method is to assess whether the obtained posterior distributions of the fitted parameters can yield observations that "cover" the original data. A successful outcome of such a test would be that the data does not deviate from the expectation under a given set of parameters more than the majority of the statistical realizations of the same expectation. This test is repeated for every posterior sample of $\bar{\lambda} = \{\bar{\lambda}_{\mathrm{osc}}, \bar{\lambda}_{\mathrm{syst}}\}$, which can be used to generate an expectation

template $\mathcal{N}_{\text{exp}}$ for each experiment and further statistically fluctuated (e.g. by applying Poisson fluctuations) to obtain a mock "replica" of the data $\mathcal{N}_{\text{repl}}$. Then, given the current posterior sample of $\bar{\lambda}$, the log-likelihoods $\ln \mathcal{L}(\mathcal{N}_{\text{data}} | \bar{\lambda})$ and $\ln \mathcal{L}(\mathcal{N}_{\text{repl}} | \bar{\lambda})$ can be computed for the original data and the mock replica. One can eventually claim that the posteriors result in replicas that cover the data well if $-\ln \mathcal{L}(\mathcal{N}_{\text{repl}}) > -\ln \mathcal{L}(\mathcal{N}_{\text{data}})$ for a sufficiently large fraction of $\bar{\lambda}$ sampled from the posterior [278]. This fraction is then defined as the p-value describing the compatibility between the tested model with its posteriors and the data. In Figure 8.5, we plot the $-\ln \mathcal{L}(\mathcal{N}_{\text{repl}})$ values against the $-\ln \mathcal{L}(\mathcal{N}_{\text{data}})$ values obtained for 2000 equally weighted posterior samples pertaining to each of the two models. The p-values are found as the fractions of points above the 1:1 lines, which results in $p_{\text{NU}} \simeq 0.658$ and $p_{\text{SM}} \simeq 0.544$.
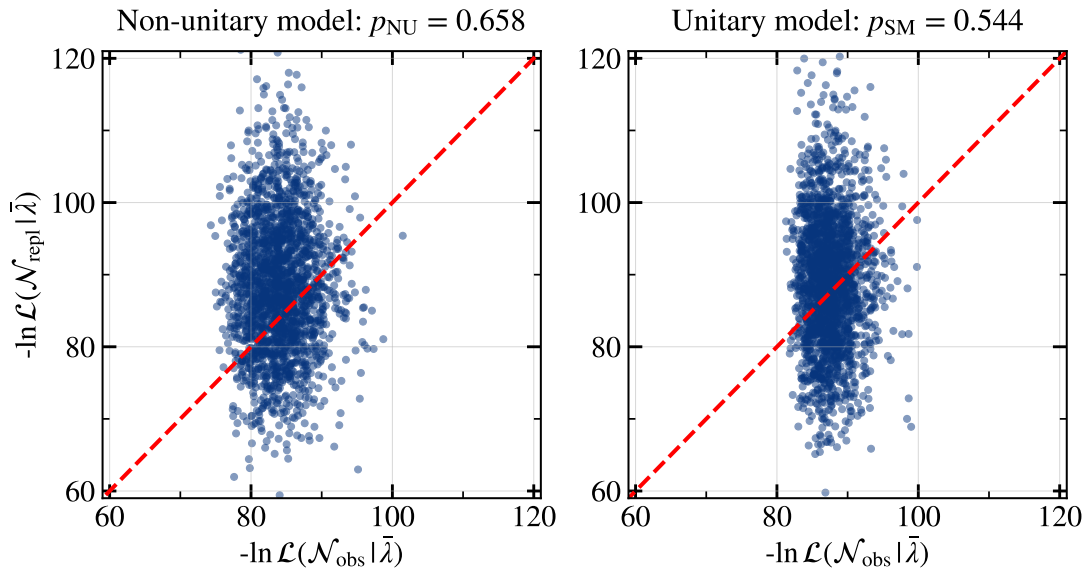

FIGURE 8.5 – The results of the posterior predictive checks via realized discrepancies for the non-unitary model (left) and the unitary model (right). In each of the scatter plots, the fraction of the points lying above the 1:1 line corresponds to the p-value for a given model. More details on the method are provided in text and Ref. [278].

We deem both of these p-values to be acceptable and conclude that the data is well described by both non-unitary and unitary models. This further strengthens our earlier argument in Section 8.1.2 that the strong preference for the unitary model according to the Bayes factor is likely a mere consequence of the prior definitions and/or the non-unitary model penalized for its complexity.

8.2 IMPACT OF THE ATMOSPHERIC NEUTRINO SYSTEMATIC UNCERTAINTIES

To test the impact of the atmospheric neutrino systematic uncertainties on the non-unitarity metrics evaluated in this study, we perform two additional Bayesian fits. In both of the fits, we replace the real data from the considered experiments with the pseudodata templates generated under the assumption of unitarity[2] and with the injected oscillation parameters fixed at the NuFit 5.2 values [65, 66]. In one of the fits, we further exclude the IceCube-DeepCore systematic uncertainties from consideration by fixing their values at the nominal expectations [46]. The results of these two fits are shown in Figure 8.6.

---

[2] Fitting pseudodata rather than real data is necessary for this test since the data would not be well described without the systematic uncertainty parameters. This would result in a bad goodness of fit and complicate the interpretation of the test outcomes.

Sensitivity projections with and w/o atmospheric systematics
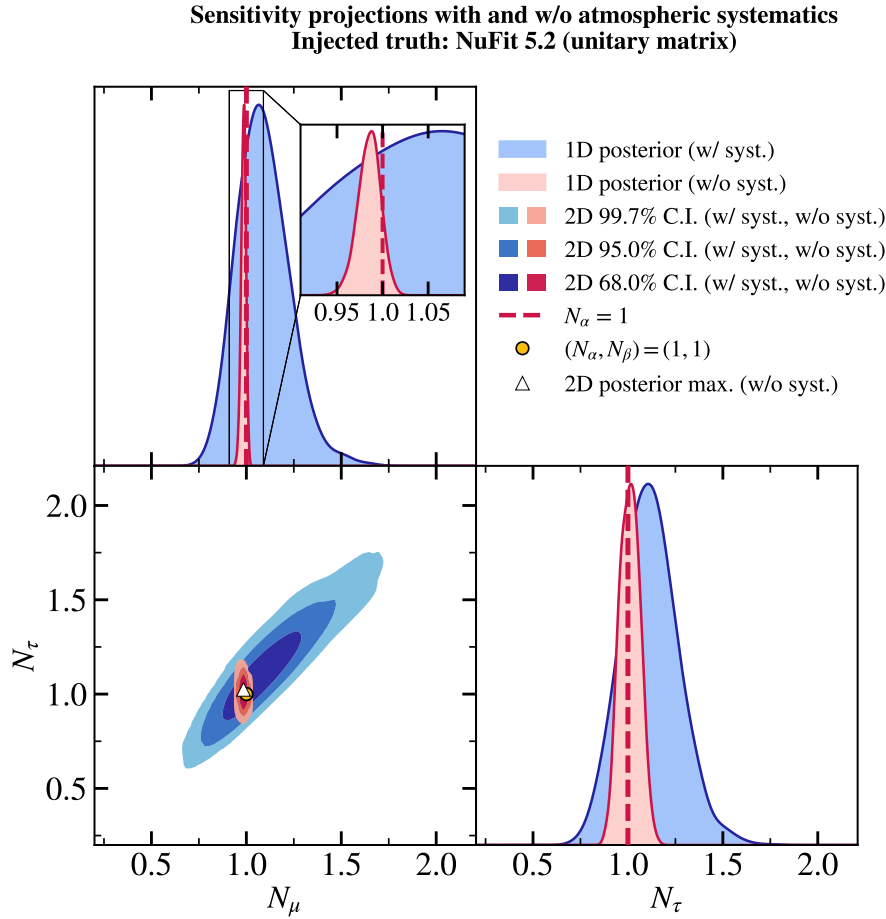**Injected truth: NuFit 5.2 (unitary matrix)**

FIGURE 8.6 – The projected posterior densities of the row normalizations N_μ and N_τ with and without atmospheric systematic uncertainties (shown as blue and red contours, respectively).

We find that if all of the IceCube-DeepCore systematic parameters were known perfectly, the constraints on N_μ and N_τ normalizations individually would be much stronger. In particular, the width of the 99.7% credible interval of N_μ would shrink by a factor of 12.4 (from 0.87 to 0.07), and that of N_τ – by a factor of 3.2 (from 0.89 to 0.28). Furthermore, the correlation between N_μ and N_τ would be broken, resulting in constraints similar to those obtained in [53].

To better understand which of the nuisance parameters induce the (N_μ, N_τ) correlations and lead to the wide one-dimensional N_μ and N_τ posteriors, we compare the nuisance parameter posterior densities between the non-unitary and the unitary model fits to the current data (see Figure B.5). The difference in the posterior distributions obtained with the two models for each parameter allows us deduce which of the IceCube-DeepCore parameters are degenerate with non-unitarity. In particular, the shift in the means of the distributions between the unitary and the non-unitary model fits suggests that the combination of the parameters in question is degenerate with the shift of the N_μ and N_τ posterior maxima to the off-unitary values. On the contrary, the change in the distribution widths of the individual systematic parameters implies that these parameters are responsible for the spread of the N_μ and N_τ distributions (in the unitary case, the spread is equal to 0, as $N_\mu = N_\tau = 1$). The two parameters with the largest change in the distribution width are the overall normalization, whose 99.7% credible interval shrinks by a factor 5 when unitarity is enforced, and the spectral index of the neutrino flux, which shrinks by 36%. These are followed by the head-on efficiency of the optical modules (25% change in the distribution width) and the rela-

tive normalization of the NC events (21% change). By directly plotting the posteriors of N_μ, N_τ against these four systematic parameters in Figure B.7 and calculating the respective Pearson correlation coefficients, we confirm that the strongest correlations are indeed observed between N_μ, N_τ and the two leading atmospheric neutrino flux parameters (normalization and spectral index). We therefore conclude that tightening the priors on these parameters would help place stronger constraints on the N_μ and N_τ normalizations.

8.3 FUTURE PROJECTIONS

With the upcoming IceCube-Upgrade and JUNO detectors, the constraints on both the individual mixing matrix elements and the non-unitarity metrics will be improved, as shown in Figures 8.7 and 8.8. A particular enhancement is seen in the sensitivity
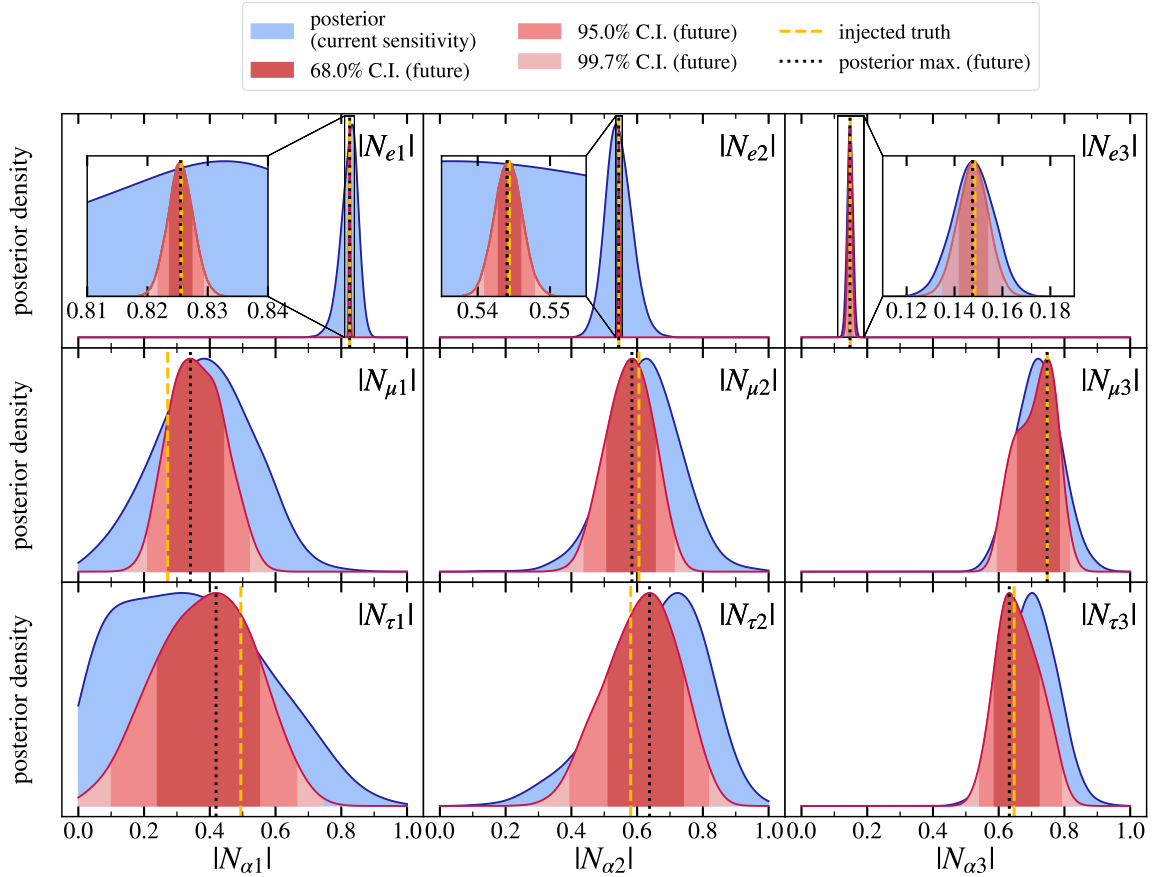
FIGURE 8.7 – Same as Figure 8.1, but applied to the future projections of the matrix element magnitudes using the combined IceCube-Upgrade, JUNO, and Daya Bay data. The injected truth is the unitary matrix generated with NuFit 5.2 oscillation parameters [65, 66].

to the $|N_{e1}|$ and $|N_{e2}|$ elements, with more than an order of magnitude decrease in the width of the respective 99.7% credible intervals. The width of the projected $|N_{e3}|$ element posterior does not change significantly with the inclusion of JUNO, as it is constrained predominantly by the high-statistics IBD selection of Daya Bay. Similarly, the N_e normalization constraint is mainly driven by the Daya Bay data as described in Section 8.1, and the large quoted improvement in the individual sensitivities to $|N_{e1}|$ and $|N_{e2}|$ does not lead to a comparably stronger N_e constraint due to the anticorrelations between these elements. On the other hand, the constraints on the muon and the tau row normalizations are expected to improve by factors ~2.5 and ~2, respectively, when replacing the three years of the current IceCube-DeepCore data with three years of the

Sensitivity projections
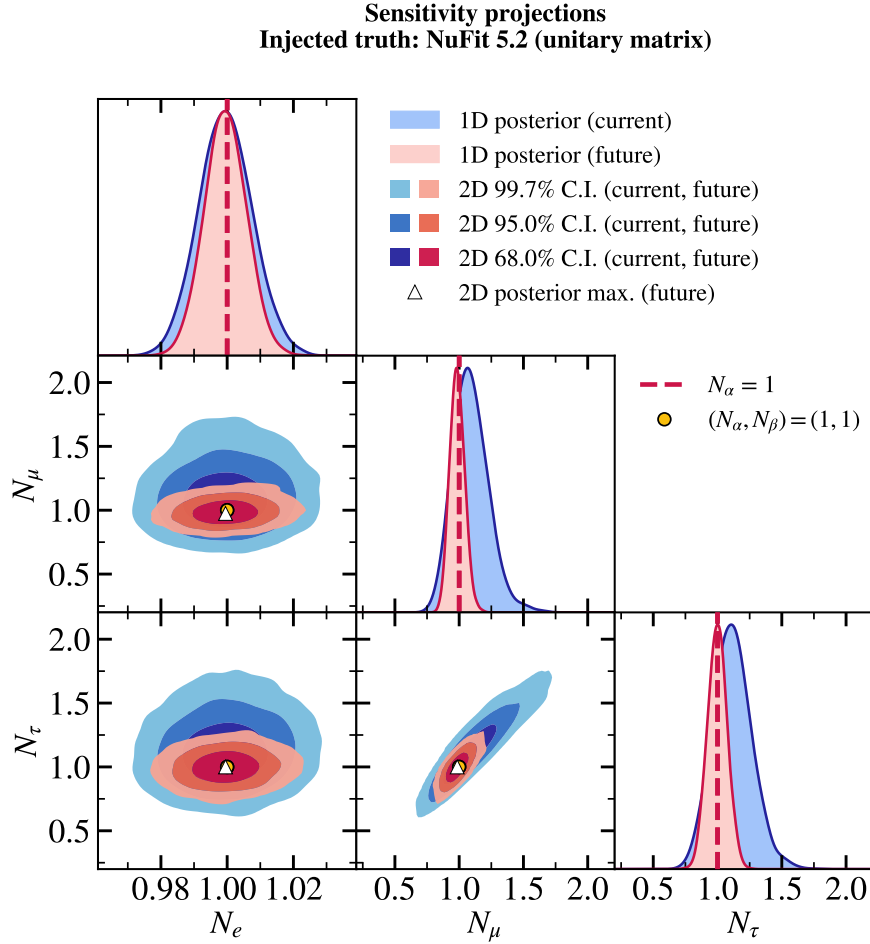**Injected truth: NuFit 5.2 (unitary matrix)**

FIGURE 8.8 – Same as Figure 8.2, but applied to the future projections of the matrix row normalizations using the combined IceCube-Upgrade, JUNO, and Daya Bay data. The injected truth is the unitary matrix generated with NuFit 5.2 oscillation parameters [65, 66].

future IceCube-Upgrade data. A similar degree of improvement is projected for the column normalizations and the closures, as seen from Figures B.8 and B.9.

8.4 CONCLUDING REMARKS

The purpose of this study was to test the unitarity of the neutrino mixing matrix by measuring its individual elements through neutrino oscillation data. We continued similar recent efforts in the literature [51–53], which combined multiple experiments into global fit analyses to access all of the elements of the mixing matrix. However, rather than building on these previous studies "horizontally," i.e., by expanding the scope of the considered experimental datasets, this work advanced the approach to the non-unitarity analysis "vertically" for a selected subset of experiments (reactor and atmospheric). This involved performing full fits of the non-unitary mixing model to the available public datasets and including the experimental nuisance parameters into the fits.

Our analysis provides constraints on the electron row normalization at the level of a few percent and the tau row normalization at the level of approximately 25-50%, results that are comparable with the state-of-the-art global fit studies [51–53]. All of the derived non-unitarity metrics in this study were found consistent with unitarity within their 95% Bayesian credible intervals. At the same time, the data is well-described by both unitary and non-unitary mixing models, with the respective posterior predictive

p-values of 66% and 54%. The computed Bayes factor shows a preference for the unitary model by ~ 15 units in log-evidence, indicating that the non-unitary model may be too complex for the current data.

This work is the first to comprehensively analyze atmospheric neutrino data with systematic uncertainties in a global fit for non-unitarity. An important outcome of this effort is the revealed degeneracy between the non-unitarity metrics (namely the muon and the tau row normalizations) and atmospheric neutrino flux systematic uncertainties. The overall flux normalization and spectral index uncertainty couple these normalizations, making their constraints positively correlated. This correlation, not seen in previous studies, highlights the importance of treating nuisance parameters carefully in new physics searches to avoid overly optimistic or otherwise misleading results. It also motivates the need to better constrain systematic uncertainties, potentially by incorporating the data-driven atmospheric neutrino flux model [141, 142].

As the row and column normalizations that were found to be in a mild tension with unitarity were constrained predominantly by the three-year IceCube-DeepCore dataset, it remains to be seen whether these results are simply a statistical feature of the latter through future analyses with more atmospheric neutrino data. This could include, for example, the oscNext event selection in IceCube-DeepCore (see Section 6.3) or the data from the upcoming IceCube-Upgrade, as well as other atmospheric neutrino experiments such as KM3NeT or Super-Kamiokande. In the context of future projections, we have shown that the addition of atmospheric neutrino data with a lower ($\mathcal{O}(1\,\text{GeV})$) energy reach and a ~ 2 times higher energy resolution – both possible with the IceCube-Upgrade compared to the IceCube-DeepCore – will help tighten the constraints on N_μ, N_τ, and all of the column normalizations by at least a factor of 2. Although these projections could be further improved by incorporating recent advancements in the IceCube-Upgrade detector simulation and the low-energy cross section systematic uncertainties, they already provide an estimate of the IceCube-Upgrade capabilities to constrain non-unitarity and set the foundation for future studies in this direction. In the electron sector, we found that the main advantage of including the data from the upcoming JUNO experiment will be in tightening the constraints on $|N_{e1}|$ and $|N_{e2}|$ by more than an order of magnitude. The constraint on the electron row normalization will be improved by nearly 25% but remain to be dominated by Daya Bay data even in the JUNO era. Similarly to the IceCube-Upgrade, the exact quantification of the sensitivity enhancement expected from JUNO is subject to a more careful treatment of the detector response and the sources of background, which will be possible to understand and model more accurately when the experiment becomes operational.

Going forward, the presented analysis could be expanded both horizontally and vertically, i.e., by adding the solar and the long-baseline experiments into the global fit while implementing the relevant systematic uncertainties for these experiments. This undertaking will be the more successful the more public information is provided by the experimental collaborations, and any other searches for new physics through global fits will similarly benefit from the availability of open data and analysis prescriptions.

V

oscNext event selection and ν_τ appearance analysis

Historical and conceptual overview

This chapter reviews the follow-up efforts to the tau neutrino appearance analysis carried out with three years of IceCube-DeepCore data [46]. A conceptually identical measurement was attempted in [54] with 8 years of data and the oscNext event selection, which replaced GRECO as described in Section 6.3. As we show in Section 9.1, this attempt reported a poor goodness of fit (frequentist p-value of 0%), preventing the target N_{ν_τ} quantity from being unblinded. This thesis work began with a close variant of this original sample, with a few modifications summarized in Section 9.1, which likewise returned a poor goodness of fit. To put this into context, we introduce the reader to the pre-fit data/Monte Carlo agreement in the analysis variables in Section 9.2, which reveal several unexplained trends and motivate our subsequent work in Chapter 10.

9.1 STATUS PRIOR TO THIS WORK

Analysis-level cuts

After the main oscNext event selection pipeline described in Section 6.3, a set of final-level cuts was applied in [54] to prepare the sample for the tau appearance analysis. These cuts included:

- Geometric containment cuts: $-500\,\mathrm{m} < z^{\mathrm{reco}} < -200\,\mathrm{m}$, $\rho^{\mathrm{reco}}_{36,\mathrm{starting}} < 300\,\mathrm{m}$, aimed at containing the events within the densely instrumented DeepCore fiducial volume both radially and vertically. Here, z^{reco} and $\rho^{\mathrm{reco}}_{36,\mathrm{starting}}$ are the reconstructed z coordinate of the interaction vertex and the radial distance from the reconstructed vertex to string 36, respectively;

- Number of "direct" hit DOMs > 2.5, aimed at selecting events with low amounts of scattering (such that at least 3 DOMs receive unscattered light);

- $t^{\mathrm{reco}} < 14\,000\,\mathrm{ns}$, with t^{reco} being the reconstructed interaction time, designed to remove a small yet unexplained excess of late events in data;

- Coincident event cuts, aimed at removing events in data where a neutrino and a muon appear in the same trigger window. These events are not simulated in

the IceCube Monte Carlo and were found to appear in data in the top 15 layers of DOMs and the outermost IceCube strings;

- Muon classifier probability > 0.4, which removed events classified as likely atmospheric muon candidates according the score of a BDT desribed in Section 6.3.

Furthermore, the total reconstructed energy was limited to the $5\,\mathrm{GeV} \leq E_{\mathrm{total}}^{\mathrm{reco}} \leq 300\,\mathrm{GeV}$ range, and the reconstructed cosine of the zenith angle – to the $-1 \leq \cos\theta_{\mathrm{zen}}^{\mathrm{reco}} \leq 0.3$ range, thereby removing the most downgoing events with high muon contamination. The final sample was histogrammed with 12 bins in energy, 10 bins in $\cos\theta_{\mathrm{zen}}$, and 3 PID bins as described in Section 6.3.

Systematic parameters

The following groups of systematic parameters were included in the original attempt at a tau appearance analysis with the oscNext sample [54]:

- *Atmospheric neutrino flux*, including the flux spectral index uncertainty $\Delta\gamma_\nu$ and the Barr/Bartol parameters described in Section 3.1, here limited to G_π, H_π, I_π, W_{K^+}, $Y_{K^\pm}$, and Z_{K^+};

- *Neutrino interaction cross sections*: "dis_csms" parameter for DIS cross section interpolation between the GENIE GRV98 and the CSMS model predictions (see Section 4.1.3), as well as the axial mass parameters $M_{A,\mathrm{RES}}$ and $M_{A,\mathrm{QE}}$ capturing the uncertainty on the resonant and the quasielastic cross sections (see Sections 4.1.1 and 4.1.2);

- *Detector*: absolute DOM efficiency relative to the nominal expectation, upscaled/-downscaled by the same amount for all DOMs;

- *Ice model*: uncertainties on the absorption and scattering coefficients in bulk ice, relative to the nominal values of the SPICE 3.2.1 ice model [279], as well as the angular acceptance parameters p_0 and p_1 (see Section 6.2.1) describing the uncertainties on the properties of the refrozen ice in the drill holes;

- *Neutrino (muon) rate normalization*: overall scale of the neutrino (muon) rate relative to the nominal expectation and livetime.

In practice, the events in the sample are pre-weighted according to the nominal value each of the above parameters as well as the physics parameters $\Delta m_{3\ell}^2$, θ_{23}, and N_{ν_τ}. The weighted events are then histogrammed according to the aforementioned binning scheme, such that the number of counts in each bin represents the sum of the final weights, $\mathcal{N}_i^{\mathrm{exp}} = \sum_i w_i$. This procedure yields the *nominal expectation templates* of the analysis, analogously to the procedure used in Chapter 7. The effects of the off-nominal parameter values are applied by modifying the weights of the sample events by the appropriate multiplicative factors, which is done in a staged manner in the PISA framework [256]. The impact of the listed systematic parameters parameters on the histogrammed event templates and the definitions of their priors can be found in [54][1]

Test statistic and sensitivity

The standard Poisson likelihood was chosen as the optimization metric in [54]:

---

[1] Specifically, see Table 7 in Chapter 8 of [54].

$$\mathrm{LH} = \sum_i^{n_{\mathrm{bins}}} \frac{\lambda_i^{k_i} e^{-\lambda_i}}{k_i!}, \tag{9.1}$$

where $\lambda_i \equiv \mathcal{N}_i^{\mathrm{exp}}$ (the expected number of events in bin i) and $k_i \equiv \mathcal{N}_i^{\mathrm{obs}}$ (the observed number of events in bin i). Therefore, no Monte Carlo uncertainty terms were taken into account for the analysis in [54]; their impact was studied in detail, with the conclusion that the oscNext sample has enough simulation statistics for the MC uncertainties to be neglected. The test statistic (TS) for the tau appearance analysis was subsequently defined as follows:

$$\mathrm{TS} = -2\ln\left[\frac{\mathrm{LH_{best\ fit}}}{\mathrm{LH}_{N_{\nu_\tau}=1}}\right] \equiv -2\Delta\mathrm{LLH}, \tag{9.2}$$

where $\mathrm{LH_{best\ fit}}$ is the maximum likelihood found by fitting the Monte Carlo simulation template to data, and $\mathrm{LH}_{N_{\nu_\tau}=1}$ is the maximum likelihood in a null fit (with the physics parameter N_{ν_τ} fixed to 1). The corresponding N_{ν_τ} sensitivity projected for the 8-year oscNext sample using statistically fluctuated Monte Carlo templates with different injected values of N_{ν_τ} is shown in Figure 9.1.
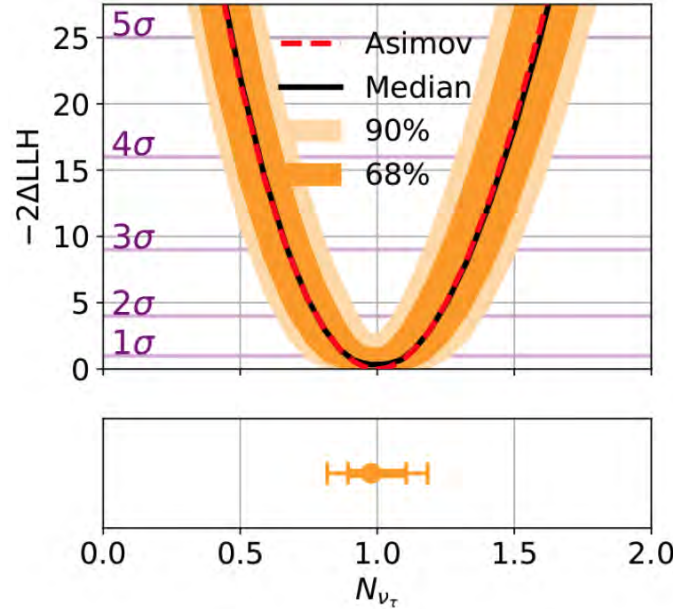
FIGURE 9.1 – Tau neutrino appearance sensitivity as projected for 8 years of IceCube-DeepCore data in [54], using the oscNext event selection with the RETRO event reconstruction. *Figure extracted from [54].*

Figure 9.1 indicates 10–13% (16–20%) sensitivity to N_{ν_τ} at a 68% (90%) confidence level – a world-leaving precision given the previous measurements by IceCube [46], Super-Kamiokande [48], and OPERA [47], with the reported 1σ bounds weaker by at least a factor of 2.

Goodness of fit issues

Despite originally promising sensitivities, the N_{ν_τ} value could not be measured from the fit to real data, as the blind analysis hit a stopping criterion: an extremely poor goodness of fit, as shown in Figure 9.2. This suggests that there is not enough freedom in the model to capture all of the data features; in other words, some of the physics effects present in data are either mismodelled or completely unsimulated in Monte Carlo. While this could in principle be a hint of new physics, the first natural suspect
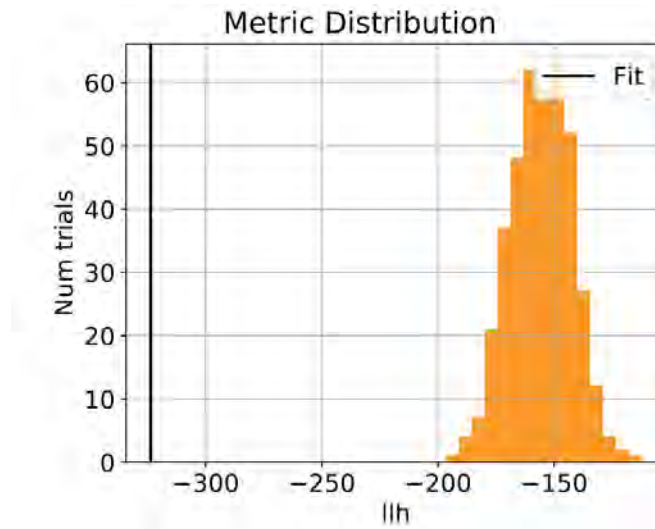
FIGURE 9.2 – Results of the tau neutrino appearance blind fit as performed by É. Bourbeau [54] using 8 years of IceCube-DeepCore data. The figure shows the LLH metric distribution (orange histogram) expected from statistically fluctuated pseudotrials at the best-fit point (BFP) and the LLH of the data-BFP agreement (solid line). This result indicates a poor goodness of fit (0% *p*-value). *Figure extracted from* [54].

is the accuracy of the Monte Carlo simulation itself. This prompts additional checks on signal and background modeling, as well as the search for any biases that could be introduced by event selection and reconstruction.

Initial sample modifications

Before the work on this thesis began, the oscNext sample underwent a few adjustments compared to the version described above. First, a new ice model was introduced to all of the IceCube simulation, internally referred to as the "BFR" ice model, which takes into account the birefringence property of ice [240]. Notably, the birefringence of ice crystals introduces an anisotropy in the light propagation direction, such that the preferred propagation direction is aligned with that of the glacial ice flow, and the diffusion in the orthogonal direction is attenuated. New Monte Carlo simulation was subsequently produced for the oscNext sample by T. Stuttard, followed by re-processing of the entire sample through the event selection chain. As expected, this transition was shown to have a direct affect on the reconstructed neutrino direction variables, namely the zenith and the azimuth angles. Additionally, a tighter cut on the Level 7 muon classifier was set, now selecting only events with the respective BDT score above 0.8 (instead of 0.4 used in [54]). This was done as a precaution in order to further reduce any potential contamination by the background muons, even though the latter were estimated to constitute less than 1% of the overall sample. While these changes were being introduced, the livetime of the sample increased from 8 to 9.3 years[2]. Finally, a few of the Barr parameters representing the uncertainty on the atmospheric neutrino flux were fixed to their nominal values due to high correlations with the rest of the flux parameters. As the result, only G_π, H_π, I_π, W_{K^+}, and Y_{K^+} were left free in the analysis.

Starting point of the present work

Figure 9.3 shows the binned event templates corresponding to the 9.3-year oscNext sample as defined at the start of this work. The per-bin magnitude of the corresponding

---

[2]At the time of the thesis submission, the livetime has further increased to 11 years, but we are presenting the work done on the 9.3 year sample only.

tau neutrino signal relative to "background" (all non-tau neutrino events) is subsequently presented in Figure 9.4. The Asimov (median expectation) sensitivities to the tau neutrino normalization obtained with this sample are shown in Figure 9.5.
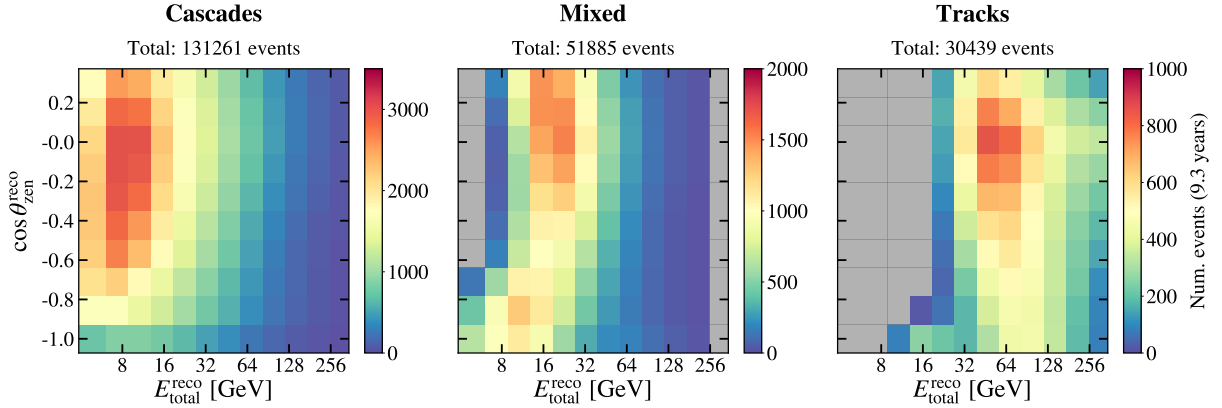

FIGURE 9.3 – Distribution of events in the oscNext sample (as defined in the beginning of this work) across the reconstructed energy (E_{total}^{reco}), reconstructed cosine of the zenith angle ($\cos\theta_{zen}^{reco}$), and PID space. The figure is based on the nominal Monte Carlo simulation weighted to 9.3-year livetime.
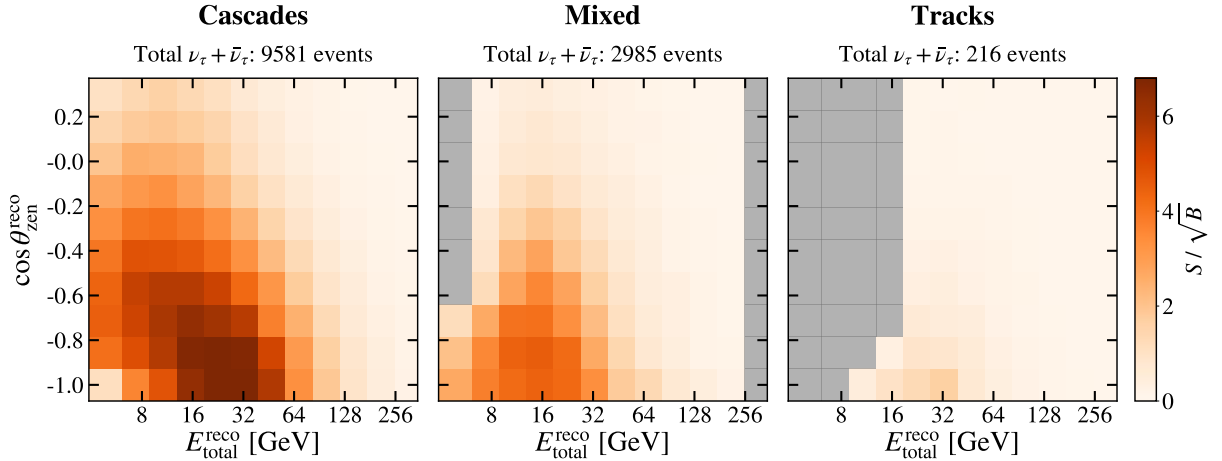
FIGURE 9.4 – Expected tau neutrino appearance signal (S) in the 9.3-year oscNext sample relative to the background (B) consisting of all non-ν_τ events. Here, the signal events include both ν_τ and $\bar{\nu}_\tau$, and both CC and NC interactions. The tau neutrino normalization (N_{ν_τ}) is set to its nominal value of 1.

The updated sample configuration yields a 1σ N_{ν_τ} sensitivity of 11.7%, closely consistent with the previous projection for the 8-year oscNext sample in Figure 9.1. We note that while the livetime of the sample increased by 16%, it was compensated by the reduction in the neutrino rate due to the tighter cut on the muon classifier score, as well as the shift to the BFR ice model. The latter significantly improved the data-Monte Carlo agreement in the reconstructed zenith and azimuth variables, which had a strongly positive impact on the goodness of fit as shown in Figure 9.6. Specifically, the data test statistic has moved to the very tail of the test statistic distribution expected from the statistically fluctuated templates at the best-fit point, which is a drastic (~100 units of LLH) improvement compared to Figure 9.2. However, this level of improvement is far from sufficient, as the p-value is still measured to be 0% (in a test with 200 pseudotrials), while the collaboration-accepted benchmark is $p \geq 5\%$.
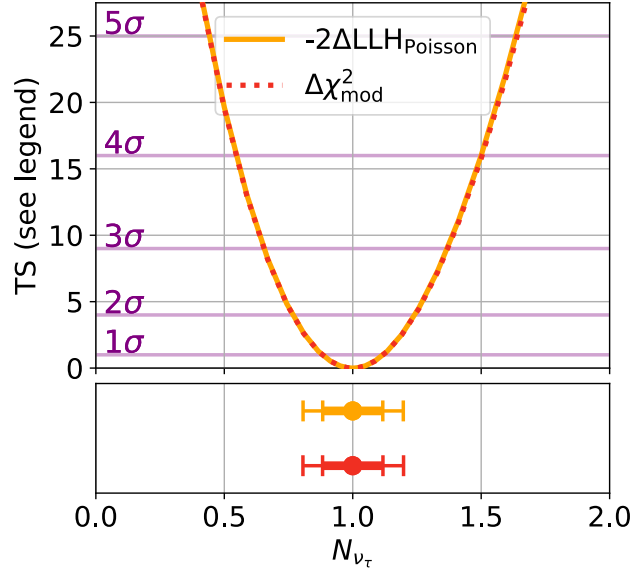
FIGURE 9.5 – Asimov sensitivities to the N_{ν_τ} parameter projected for the 9.3-year oscNext sample as defined at the start of this work, shown for two choices of the test statistic (Poisson likelihood *vs.* modified χ^2 (see Equation (7.3))). The lower panel shows the 68% and the 90% confidence intervals.
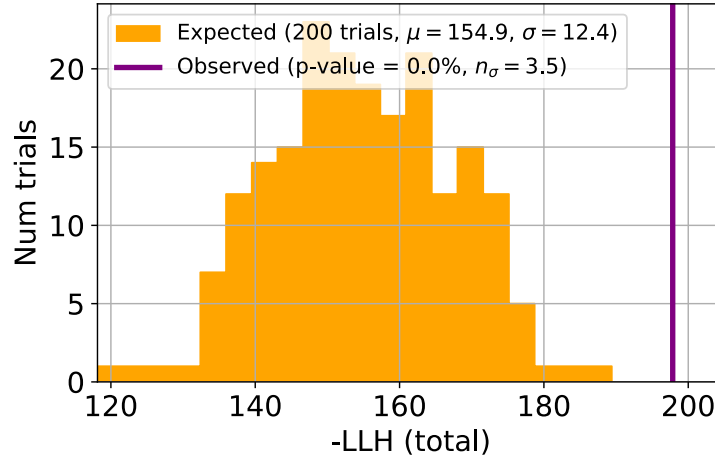

FIGURE 9.6 – Results of the tau neutrino appearance blind fit as performed with 9.3 years of IceCube-DeepCore data, with sample modifications relative to Figure 9.2 described in text. Note that the axis of this figure is reversed compared to Figure 9.2.

9.2 HIGH-LEVEL INDICATORS OF DATA/MONTE CARLO DISAGREEMENT

We now turn our attention to the potential "early warning" symptoms of the data/Monte Carlo disagreement that could be causing a bad goodness of fit in Figure 9.6. In particular, in Figures 9.7 and 9.8, we show the data and the pre-fit Monte Carlo distributions for the total reconstructed energy (sum of track and cascade energies returned by the RETRO reconstruction algorithm). We find an apparent excess of high-energy events and a deficit of low-energy events in data across all PID bins, with the resulting "tilt" in the data/MC ratio having a strong PID dependence (as seen from Figure 9.8). Specifically, the events that are classified as more track-like appear to exhibit a stronger tilt, or bias, of the data distribution with respect to the Monte Carlo. Thus, it is possible that the issue is associated with true ν_μ-CC events, which are strongly dominant in the track PID bin. Further, we note a peculiar shape of the data distribution at reconstructed energies above ∼50 GeV in the mixed bin and ∼70 GeV in the track bin, which might be indicative of a presence of an extra population of events in data compared to MC.
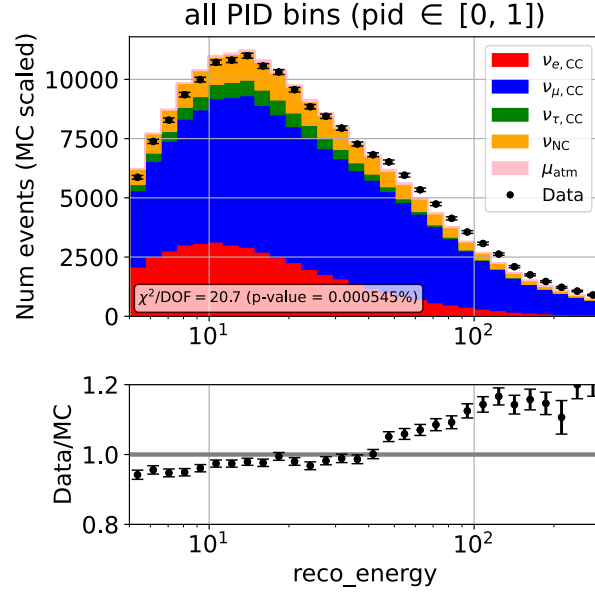
FIGURE 9.7 – Pre-fit data/Monte Carlo agreement in the total reconstructed neutrino energy variable ("reco_energy", shown in GeV), here combined into a single histogram for all PID bins.
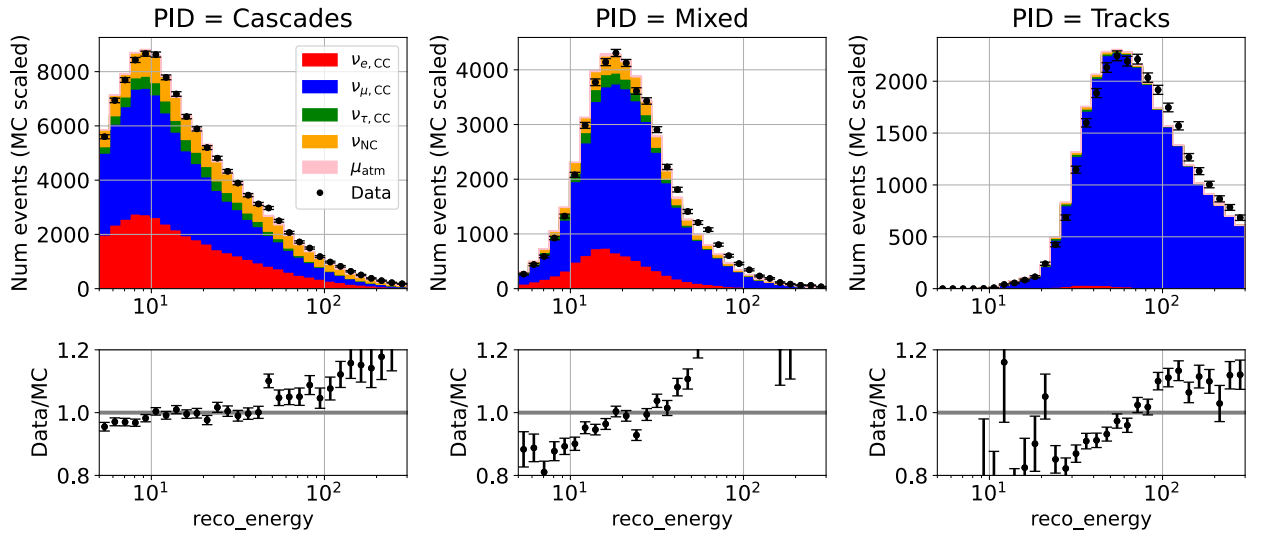

FIGURE 9.8 – Same as Figure 9.7, but split between cascades (PID $\in [0, 0.5)$), mixed (PID $\in [0.5, 0.85)$), and tracks (PID $\in [0.85, 1]$) PID bins.

By splitting the total energy into its cascade and track components, we find that the cascade energy data/MC ratio exhibits a simple tilt prior to the fit (not shown here), which is ultimately inherited by the total reconstructed energy distributions. At the same time, the track energy data/MC agreement has a more complicated shape, as seen from Figures 9.9 and 9.10. In particular, the aforementioned "feature" (an irregularly shaped excess) in data appears particularly prominent in the mixed bin above reconstructed track energy of 30–40 GeV, and is similarly present in the track bin at higher energies. While in principle the pre-fit disagreement in the individual 1D data/Monte Carlo distributions may not necessarily lead to a poor goodness of fit, any remaining disagreement in the post-fit distributions is likely to be correlated with the latter. With that in mind, we note that the complex shape of the reconstructed track energy data/MC ratio is not corrected post-fit (i.e., after finding the best-fitting physics and systematic parameters via likelihood optimization), which is shown in the right
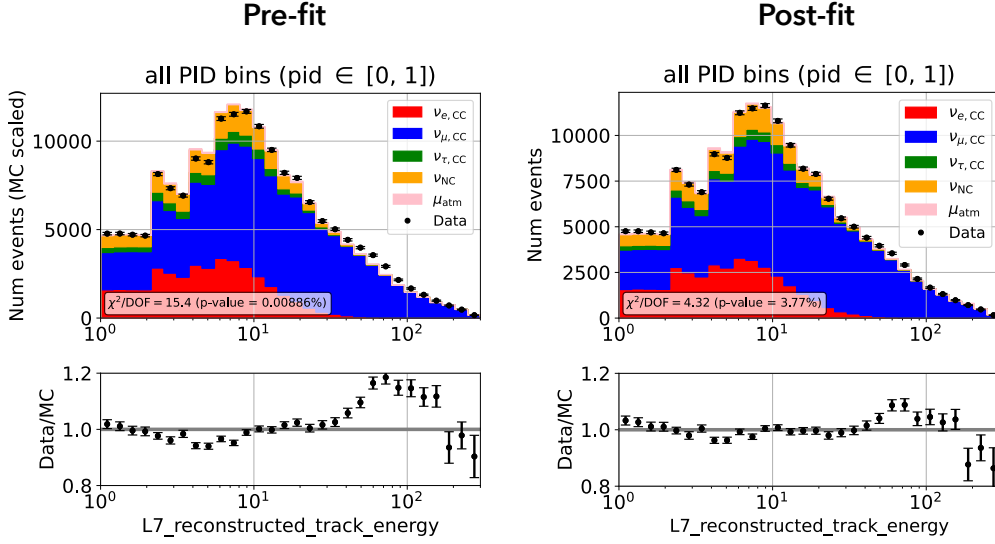
FIGURE 9.9 – *Left (right) panel*: Pre-fit (post-fit) data/Monte Carlo agreement in the reconstructed track energy variable ("L7_reconstructed _track_energy", shown in GeV), here combined into a single histogram for all PID bins.
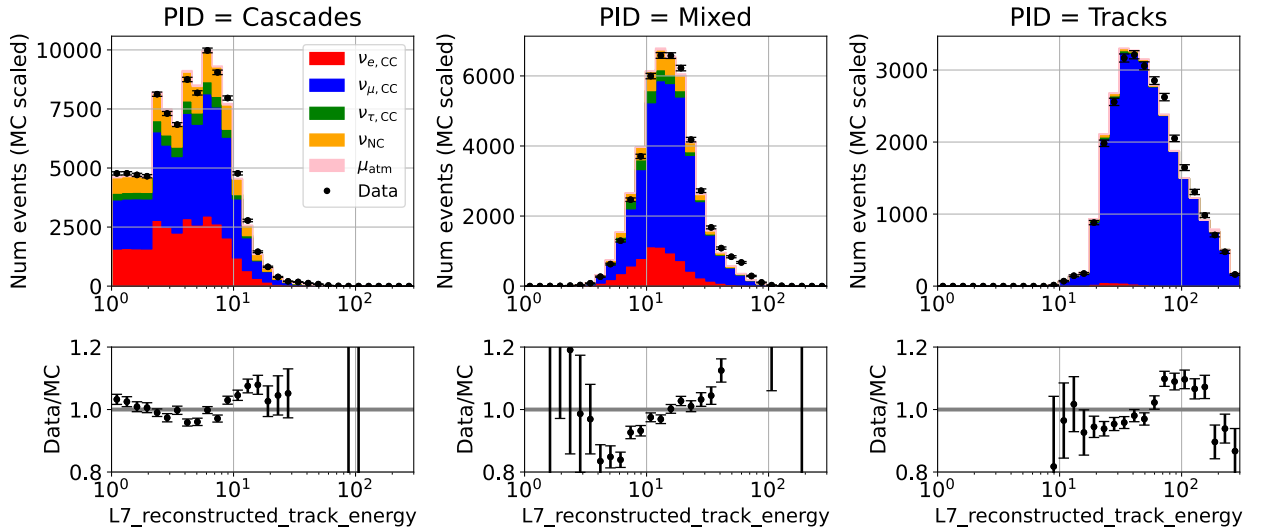

FIGURE 9.10 – Same as the left panel of Figure 9.9 (pre-fit agreement), but split between cascades (PID ∈ [0, 0.5)), mixed (PID ∈ [0.5, 0.85)), and tracks (PID ∈ [0.85, 1]) PID bins.

panel of Figure 9.9. This implies that there is no systematic parameter, or combination thereof, that would be capable of flattening out the complex shape of the data/MC ratio in the reconstructed track energy. This observation is a starting point for a series of detailed investigations related to muon simulation and track energy reconstruction, to which we dedicate the subsequent chapter.

Identified issues in the oscNext Monte Carlo simulation

The energy- and PID-dependent data-MC disagreement discussed in Section 9.2 prompts further investigation in two complementary directions – the accuracy of the Cherenkov light yield simulation in IceCube and the robustness of the track energy reconstruction. The first item concerns the low levels (Step 2) of the simulation chain described in Section 6.2. At this stage, a certain amount of Cherenkov photons is generated depending on the types and the energies of the charged neutrino interaction products, followed by photon propagation towards the IceCube DOMs. The number of detected photons ultimately affects the reconstructed energy, meaning that mismodeled light yield in simulation could be a source of energy-dependent data/MC disagreement. Thus, the first goal of this chapter is to validate the implementation of the Cherenkov light yield from first principles, with a particular focus on any energy-dependent components of the simulation. Given that the data/MC disagreement appears to worsen continuously as the events become more track-like, this study begins with a closer look on the *muon* light yield, which we examine in Section 10.1. In Section 10.2, we proceed with an investigation of Cherenkov light yield by electrons, which can be produced either directly in ν_e-CC interactions or as secondary particles due to muon stochastic energy losses. Finally, in Section 10.3, we address the second possible culprit of the data/MC issues, namely energy reconstruction with the retro algorithm and its behavior in the region of data excess. We devise several variables which reflect the track energy reconstruction quality and use them to identify populations of badly reconstructed events. We further present evidence of noise being the origin of such misreconstructions and the retro algorithm's susceptibility to the different levels of noise between data and Monte Carlo.

Both of these investigations lead to important findings and the corresponding adjustments in the oscNext simulation and processing chain. However, we warn the reader that they do not fully mitigate the main data/MC trend as a function of energy and PID, and that more work is needed to reveal the true origin of this trend.

10.1 MUON LIGHT YIELD

10.1.1 *Issues with the original simulation*

To begin the investigation into the track energy-dependent mismodelling in the oscNext Monte Carlo simulation, we ask the following questions:

1. Is the muon propagation simulated properly, i.e, does the muon track have the expected length?
2. Is the Cherenkov light emission per muon track length in agreement with the Frank-Tamm formula, Equation (4.40)?

To answer these questions, we design a toy study where small Monte Carlo samples of fixed-energy muons are generated at Step 1, followed by Cherenkov photon emission due to these muons at Step 2. For simplicity of interpretation, we fix the starting position of all muons at $\langle x_0, y_0, z_0 \rangle = \langle 0\,\text{m}, 0\text{m}, -450\,\text{m} \rangle$ and direct them vertically towards the center of IceCube ($\cos\theta_{\text{zen}} = -1$)[1]. We assign the following initial total energies to the muons: $[1\,\text{GeV}, 2\,\text{GeV}, 5\,\text{GeV}, 10\,\text{GeV}, 20\,\text{GeV}, 50\,\text{GeV}, 100\,\text{GeV}, 200\,\text{GeV}]$, and generate 10 000 events per energy.

As in the original oscNext simulation, we use the PROPOSAL software to propagate muons at Step 1, which takes into account continuous and stochastic muon energy losses. PROPOSAL requires setting either an absolute (E_{cut}) or a relative (v_{cut}) energy threshold below which energy losses are treated continously, and above which – stochastically with emission of individual secondary particles[2]. For the initial version of this toy study, we use the same setting as in the full oscNext MC, namely $E_{\text{cut}} = 0.5\,\text{GeV}$. When the muon stops and decays, its stopping position is used to calculate the total propagated distance, which we extract to answer Question 1. Upon inspecting the resulting track length distributions, we identify the following issue.

- **Issue 1:** Unphysical artefacts in the muon stopping range

At low (1–20 GeV) muon energies, the track length distributions obtained with the $E_{\text{cut}} = 0.5\,\text{GeV}$ setting show an unphysical clumping at discrete track length values, as shown in the orange histogram of Figure 10.1. This is a known artefact of PROPOSAL when the energy cut setting is too high, and has been illustrated, for example, in Figures 17-18 of [218]. We further attempt to lower the absolute energy cut to $E_{\text{cut}} = 0.1\,\text{GeV}$ as well as to set the relative energy cut to $v_{\text{cut}} = 10^{-4}$, adding these results to Figure 10.1. Our findings confirm that for $\mathcal{O}(1\,\text{GeV})$ energy muons, energy losses of $\mathcal{O}(0.1\,\text{GeV})$ should be treated stochastically rather than via the continuous loss approximation, since they occur at most a few times before the muon stops, and the variance in the magnitude of these individual energy losses is large. The best-case scenario among the tested options is $v_{\text{cut}} = 10^{-4}$, where only the energy losses below 0.1 MeV at $E_\mu = 1\,\text{GeV}$ and below 20 MeV at $E_\mu = 200\,\text{GeV}$ are treated continuously. At 200 GeV, the track length distribution obtained with this v_{cut} setting can be well approximated with the tested absolute energy cut options.

---

[1]Although such muons move in the z direction similarly to our setup in Section 4.2.2, we denote the distance travelled by the muon along this dimension as "x", such that we can be consistent with the convention of the energy losses being represented by $\frac{\text{d}E}{\text{d}x}$.

[2]Note that if both E_{cut} and v_{cut} are specified, PROPOSAL chooses $\min\{E_{\text{cut}}, v_{\text{cut}}E_\mu\}$ for the energy threshold.
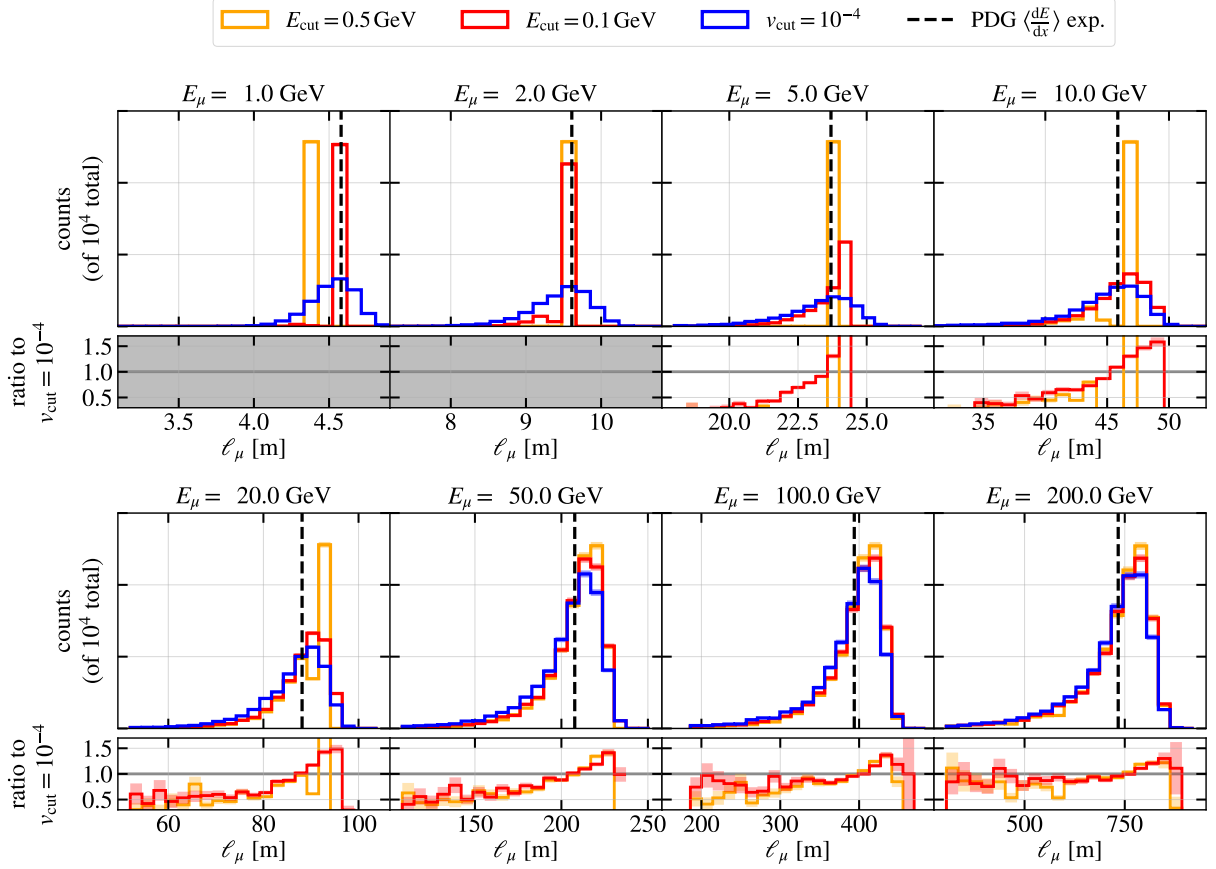
FIGURE 10.1 – Distributions of distances ℓ_μ travelled by muons with fixed total energies E_μ, as obtained from the PROPOSAL software [218]. The histograms of different colors represent the different settings of the energy threshold defining the transition between the stochastic and the continuous energy losses. The $v_{\text{cut}} = 10^{-4}$ case (dark blue histogram) provides the most accurate muon track length distribution among the tested options. The dashed vertical line represents the track length estimate obtained assuming average energy losses in ice as extracted from the PDG Review [34] (see Figure 10.2 and text for details). The ratios of the track lengths distributions to the reference case of $v_{\text{cut}} = 10^{-4}$ are not shown for the lowest two muon energies due to significant differences in the distribution shapes.

Despite resulting in unphysical shapes of the track length distributions, the original PROPOSAL settings return reasonably accurate values for the median track length. To prove this, we set up a simple propagator of muons with the selected discrete energies and the average energy losses per unit length, $\langle \frac{dE}{dx} \rangle$, following the Bethe-Bloch prescription [280, 281]. The exact $\langle \frac{dE}{dx} \rangle$ values are extracted from the Particle Data Group (PDG) 2022 Review [34] and plotted in Figure 10.2. We propagate the muons in small steps of length $\Delta x = 1\,\text{cm}$ and calculate the kinetic energy loss $\Delta E_{\mu,i}^{\text{kin}} = \langle \frac{dE}{dx} \rangle (E_{\mu,i}^{\text{kin}}) \cdot \Delta x$ at step i. At the subsequent propagation step, we update the muon kinetic energy to $E_{\mu,i+1}^{\text{kin}} = E_{\mu,i}^{\text{kin}} - \Delta E_{\mu,i}^{\text{kin}}$, and continue the propagation until the muon reaches the minimum kinetic energy available in the PDG tables (1 MeV). We record the muon track lengths as $\ell_\mu = n\Delta x$, where n is the number of propagation steps reached until the muon stopping point, and display these results as the dashed vertical lines in Figure 10.1. Except for the lowest-energy case ($E_\mu = 1\,\text{GeV}$), the PROPOSAL distribution medians agree well with the toy propagator estimates even with the $E_{\text{cut}} = 0.5\,\text{GeV}$ setting used in the oscNext simulation. This implies that, on average, the oscNext muons propagated in PROPOSAL travel the right distance, but the shape of the track length distribution for monoenergetic muon beams is not accurate.

To understand whether the track length artefacts appearing in monoenergetic
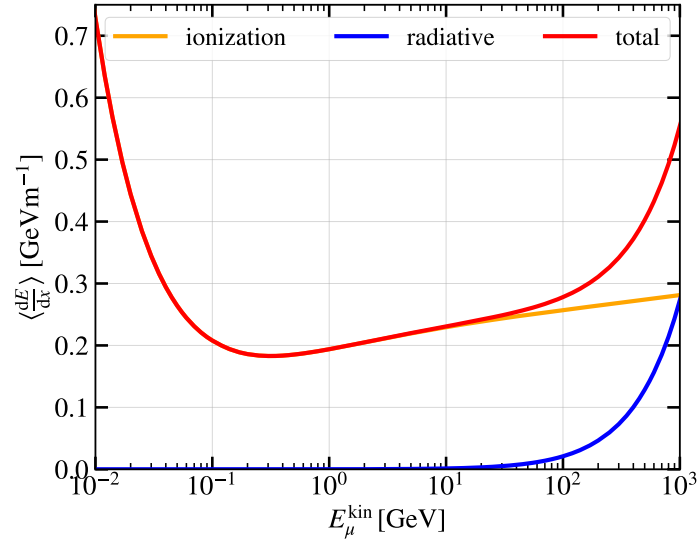
FIGURE 10.2 – Average muon energy losses per unit length in ice ($\rho_{ice} = 0.9216\,\mathrm{g\,cm^{-3}}$) as a function of muon kinetic energy E_μ^{kin}. The values are extracted from the PDG Review [34].

muon simulations still have an impact on the full-scale simulations of ν_μ-CC interactions with a wide range of neutrino energies, we produce another benchmark Monte Carlo sample of 50 000 muons with their total energies uniformly distributed in the $[1\,\mathrm{GeV}, 200\,\mathrm{GeV}]$ range. We simulate only the two limiting cases from Figure 10.1 here, namely $E_{\mathrm{cut}} = 0.5\,\mathrm{GeV}$ and $v_{\mathrm{cut}} = 10^{-4}$. The results are given in Figure 10.3. We can see
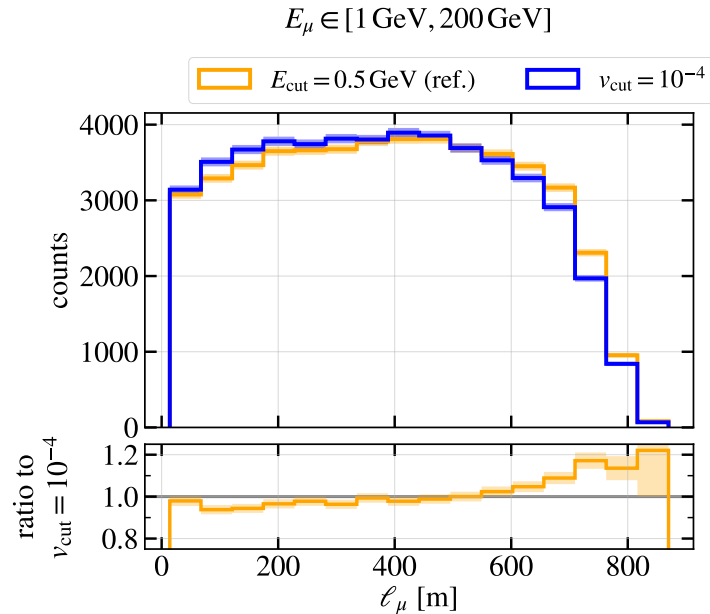

FIGURE 10.3 – Muon track lengths as obtained from PROPOSAL for a muon sample with a uniform energy distribution, $E_\mu \in [1\,\mathrm{GeV}, 200\,\mathrm{GeV}]$.

that even though the overall distribution of track lengths obtained for this muon sample with $E_{\mathrm{cut}} = 0.5\,\mathrm{GeV}$ does not display any obvious artefacts, it is still not statistically compatible with the $v_{\mathrm{cut}} = 10^{-4}$ case. This can be seen from the ratio of the respective distributions in Figure 10.3, which shows statistically significant deviations from 1, as well as an extremely low two-sample Kolmogorov-Smirnov test p-value ($2.7 \cdot 10^{-9}$). While we investigated the uniform energy distribution case, the atmospheric neutrino spectrum follows a steeper ($\sim E_\nu^{-3.7}$) power law, where the low-energy track length artefacts will be exacerbated. We therefore conclude that **the true muon track length**

distribution in the original oscNext simulation is <u>not</u> accurate even for a large ensemble of muon neutrinos with different energies. This implies that the daughter muons in the simulated ν_μ-CC interactions have to be repropagated either with a lower energy cut setting in proposal or with another software package such as Geant4, where the default continuous-to-stochastic energy loss transition threshold is very low[3].

We now turn our attention to the second question with which we began this section, namely the accuracy of the Cherenkov light emission modelling. Following the muon propagation in proposal, Cherenkov photon generation in the oscNext MC was handled by the ppc module within CLSim. A close examination of the inner workings of ppc reveals that its implementation of the Frank-Tamm formula (Equation (4.40)) assumes that muon travels at the speed of light, i.e., with $\beta = 1$ throughout the muon propagation regardless of the actual muon velocity at a given propagation step. As β decreases as the muon loses its energy, this simplification leads to the second major issue in the oscNext simulation, which we discuss below.

- **Issue 2:** <u>Overestimated light yield per unit track length</u>

To arrive at the expected number of photons as a function of initial muon energy, we can use the Frank-Tamm formula in conjunction with the toy muon propagator under average energy losses $\langle \frac{dE}{dx} \rangle$ set up in this section. The first step of such a combined calculation is tracking the $\beta(x)$ histories for each considered muon energy, which we show in Figure 10.4 for $E_\mu \leq 10\,\text{GeV}$. We find that even though the $\beta = 1$ approximation
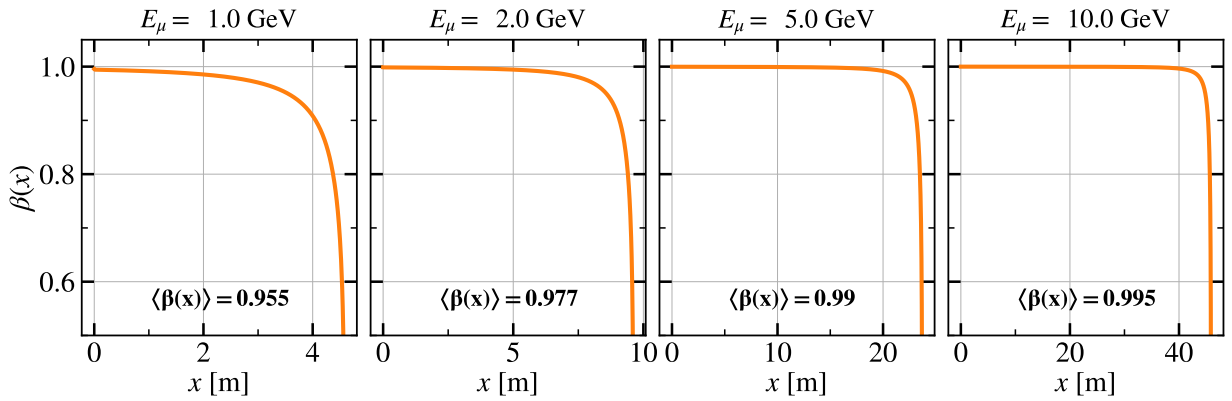

Figure 10.4 – Evolution of the muon Lorentz velocity factor β as a function of distance x travelled by the muon, shown for four starting total muon energies E_μ. Each panel reports the average β along the muon trajectory, computed as $\langle \beta(x) \rangle = \frac{1}{\ell_\mu} \int_0^{\ell_\mu} \beta(x) dx$, where ℓ_μ is the total muon track length.

is reasonable at the very beginning of the muon path, it is no longer valid within the last few meters before the muon stopping point. For muons with $\mathbb{O}(\text{few GeV})$ energies, these last few meters represent a large portion of the track, such that the average β drops noticeably below 1. The error that the $\beta = 1$ approximation brings to the muon light yield is further exacerbated by β appearing in the Frank-Tamm formula quadratically rather than linearly.

Having the $\beta(x)$ histories for each muon energy, we can integrate the Frank-Tamm

---

[3]By default, Geant4 treats stochastically the production of all secondaries that can travel at least 1 mm in the considered medium. For electrons in ice, this corresponds to ~0.3 MeV [282]. Only below this threshold, the energy losses due to ionization are treated continuously.

prediction for the double-differential number of photons $\frac{\partial^2 N_{\gamma,\mu}}{\partial\lambda\partial x}$ over the muon trajectory and over the photon wavelength, arriving at the total expected number of photons:

$$N_{\gamma,\mu} = \int_0^{\ell_\mu} \int_{\lambda_{\min}}^{\lambda_{\max}} \frac{\partial^2 N_{\gamma,\mu}}{\partial\lambda\partial x} d\lambda dx = \int_0^{\ell_\mu} \int_{\lambda_{\min}}^{\lambda_{\max}} \frac{2\pi\alpha}{\lambda^2} \left(1 - \frac{1}{\beta^2(x)n^2(\lambda)}\right) d\lambda dx. \qquad (10.1)$$

We set the wavelength integration limits at $\lambda_{\min} = 265\,\text{nm}$ and $\lambda_{\max} = 675\,\text{nm}$, as only the photons within these wavelength bounds are generated in CLSɪᴍ due the wavelength acceptance of the IceCube DOMs. The wavelength dependence of the refractive index, $n(\lambda)$, is extracted from Figure 4.9. Finally, we divide $N_{\gamma,\mu}$ by the total length of the muon track, ℓ_μ, to obtain the average number of photons emitted per unit length. The results are shown in Figure 10.5, where we compare the realistic evolution of β along the muon path with the simplified case of $\beta = 1$ enforced for the muon light yield calculation.
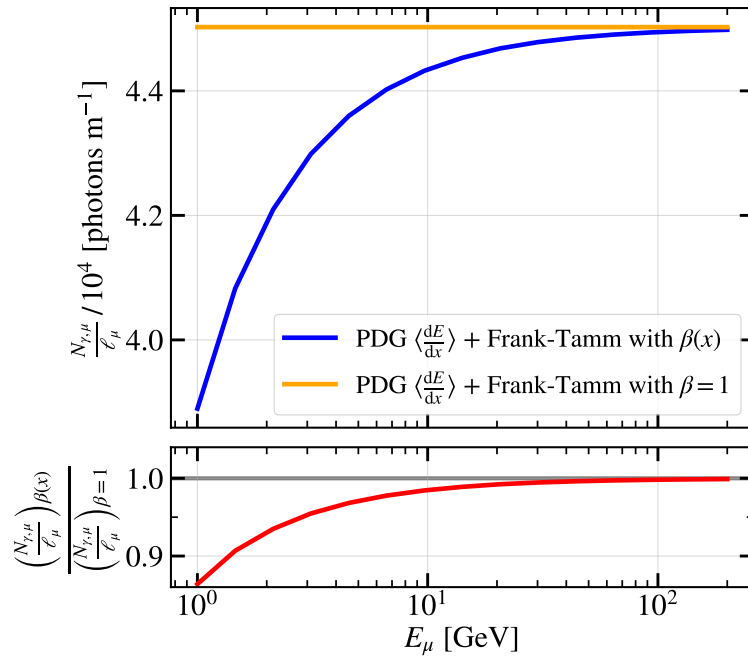

Figure 10.5 – Average muon light yield per unit length as computed from the Frank-Tamm formula with the full β dependence (blue curve) and the simplifying assumption of $\beta = 1$ along the entire muon trajectory (orange curve). Note that the muon length ℓ_μ is identical between the $\beta(x)$ and the $\beta = 1$ cases, which only impact the light yield.

As expected from Equation (10.1), the muon that is assumed to travel at $\beta = 1$ throughout its entire path produces more light than the muon with the realistic $\beta(x)$ evolution as in Figure 10.4. The discrepancy in the average number of photons per unit length reaches 14% at 1 GeV and approximately 2% at 10 GeV[4]. Since the Level 7 oscNext Monte Carlo simulation discussed in Chapter 9 contains $\gtrsim$40% of ν_μ-CC events with outgoing muon energies below 10 GeV (see Figure 10.6), ***the significant mismodelling of Cherenkov light yield at these low energies must be corrected in the simulation***.

Due to technical difficulties with incorporating the β evolution along the muon trajectory into the light yield calculation within ᴘᴘᴄ, *we switch to the Gᴇᴀɴᴛ4 software package*, where the β dependence of the Frank-Tamm formula is already integrated into the Cherenkov photon generation routines. Furthermore, as discussed earlier in the

---

[4]These numbers can be simply verified by substituting $\langle\beta(x)\rangle$ from Figure 10.4 into the following ratio: $\left(1 - \frac{1}{\langle\beta(x)\rangle^2 n^2}\right)\left(1 - \frac{1}{n^2}\right)^{-1}$, and assuming $n = \text{const} = 1.31$ for simplicity.
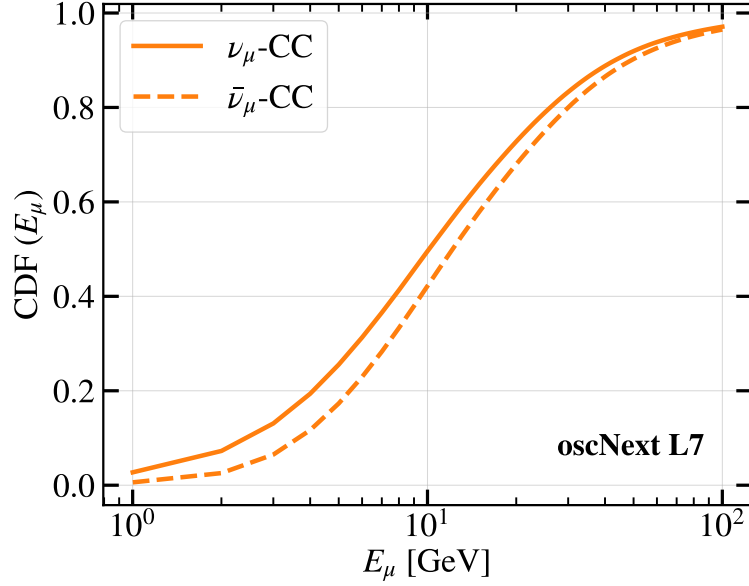
FIGURE 10.6 – Fraction of the charged current muon neutrino (antineutrino) events in the oscNext Level 7 Monte Carlo simulation with outgoing muon energies smaller than the value on the *x*-axis; equivalently, the cumulative distribution function (CDF) of the muon energy E_μ.

context of PROPOSAL propagation artefacts, GEANT4 has a sufficiently low continuous-to-stochastic energy loss transition threshold to not introduce any unphysical artefacts in the muon stopping range. In the following subsection, we validate the light yield output of the GEANT4 simulation.

10.1.2 *Geant4 simulation: Low-level impact*

Here, we replicate the toy muon Monte Carlo study from Section 10.1.1, but now using GEANT4 for both muon propagation and Cherenkov photon generation. Just like PPC, GEANT4 is called from within CLSIM, and the generated photons are propagated to and detected at the IceCube DOMs. Figures 10.7 and 10.8 show the distributions of the generated and the detected photons, respectively, comparing the GEANT4 output to that from the baseline PROPOSAL+PPC simulation (i.e., with $E_{\text{cut}} = 0.5\,\text{GeV}$). In both cases, we only keep track the photons generated by the muon alone and set the light yield output from any secondary particles to 0[5].

As expected from Figure 10.5, the muon light yield output from GEANT4 is lower than that from PROPOSAL+PPC, with the largest differences observed at $E_\mu \leq 10\,\text{GeV}$. The number of generated photons ($N_{\gamma,\mu}^{\text{generated}}$) in the PROPOSAL+PPC case inherits the muon track length artefacts from Figure 10.1, while the GEANT4 simulation always returns a smooth distribution of $N_{\gamma,\mu}^{\text{generated}}$. Furthermore, as shown in Figure 10.9, the mean of the GEANT4-based distribution of $N_{\gamma,\mu}^{\text{generated}}$ matches the Frank-Tamm prediction for the Cherenkov light yield with the full $\beta(x)$ evolution at a sub-percent level. Finally, the ratio of the GEANT4 and PROPOSAL+PPC distribution means is in excellent agreement with the ratio of the predicted number of photons under the $\beta(x)$ and the $\beta = 1$ assumptions, as seen from Figure 10.9 for both generated and detected photons. These checks validate

---

[5]This is done only for the purpose of simplifying the interpretation of tests in this section. The full oscNext Monte Carlo includes the light yield output from the secondary particles.
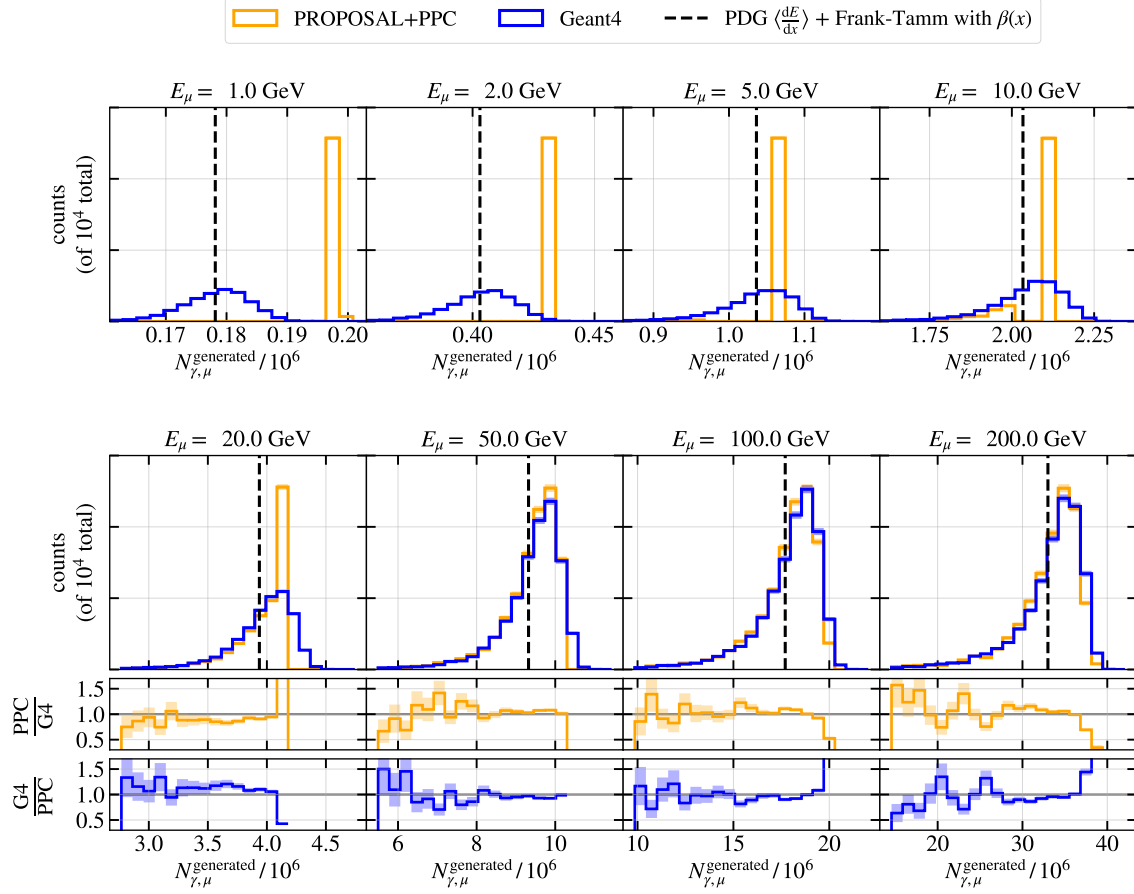
FIGURE 10.7 – Number of generated Cherenkov photons in ice due to the passage of a muon with a fixed total energy E_μ, as compared between the approach originally used in the oscNext simulation (PROPOSAL+PPC, shortly "PPC") and the approach advocated for in this study (GEANT4 simulation, shortly "G4"). The y-axis represents the probability density of a given number of photons as obtained from a Monte Carlo simulation of 10 000 muon events per panel. The PPC/G4 and G4/PPC ratios are shown only for $E_\mu \geq 20\,\mathrm{GeV}$ due to large differences in the distribution shapes below this energy.

the use of GEANT4 to simulate Cherenkov light yield at muon energies as low as 1 GeV[6].

10.1.3 *Geant4 simulation at analysis level*

Having shown that switching to the GEANT4-based Cherenkov photon generation and detection mitigates the identified issues of the PROPOSAL+PPC simulation, we entirely replace PROPOSAL+PPC with GEANT4 in the oscNext simulation chain. As with the rest of the official oscNext simulation sets, the production and processing of the GEANT4-based neutrino Monte Carlo was handled by T. Stuttard. In Figure 10.10, we show the impact of switching from PROPOSAL+PPC to GEANT4 on ν_μ-CC events at the final level (L7) of the oscNext event selection[7]. We find that significantly fewer muon neutrinos with true energies below 10 GeV make it into the final-level sample when GEANT4 is used to simulate the muon light yield instead of PROPOSAL+PPC. Specifically, at $E_\nu = 2\,\mathrm{GeV}$ the drop-off in the event rate is as large as 20%, and is approximately 10% at $E_\nu = 6\,\mathrm{GeV}$. This is explained by the fact that when the muon light yield is corrected to be dimmer than simulated previously, fewer low-energy events pass the trigger threshold, and the

---

[6]The origin of mild discrepancies between the GEANT4 and the PROPOSAL+PPC number of generated photons at $E_\mu \geq 100\,\mathrm{GeV}$, as seen in Figure 10.7, as well as the difference in the widths of the detected photon distributions in the same energy range in Figure 10.8, has not been identified. However, the means of all distributions mentioned here are in excellent agreement between GEANT4 and PROPOSAL+PPC, and these slight discrepancies are not a source of concern.

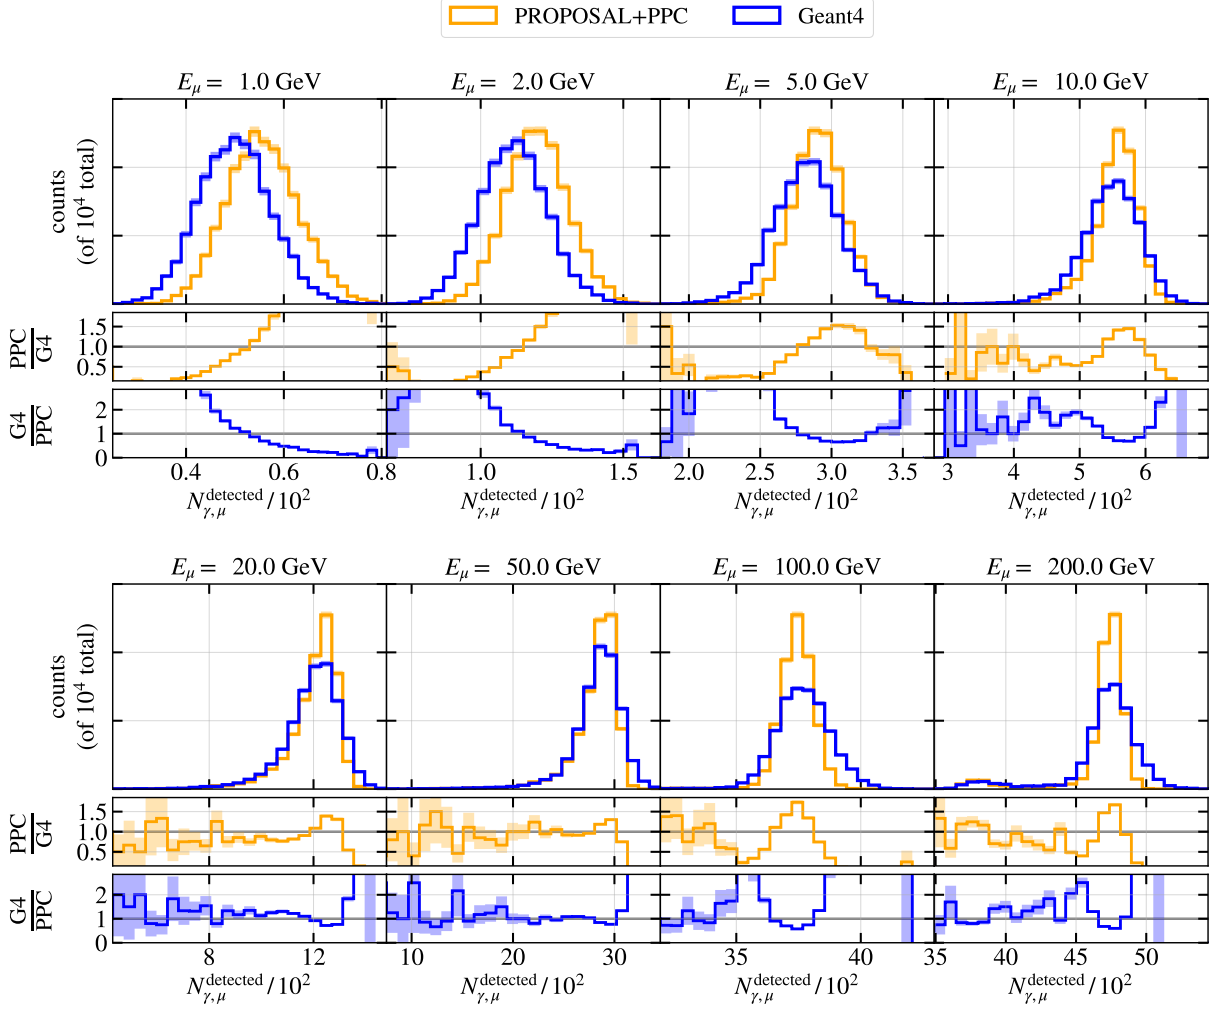[7]We note that electron and tau neutrinos were not resimulated at this stage.

FIGURE 10.8 – Same as Figure 10.7, but for photons detected at the IceCube DOMs.

dimmest events simply drop out from the sample already after Level 2. The effect is still visible in the reconstructed energy space, where the maximum drop-off in the event rate reaches ~6–7%.

Given that both true and reconstructed ν_μ energy distributions are significantly affected by correcting the muon light yield, it is natural to study what impact the light yield mismodeling makes on the measurement of the muon neutrino disappearance parameters, θ_{23} and Δm_{31}^2. First, we perform the ***inject-recovery test***, where the Geant4-based Monte Carlo template[8] is used as pseudodata and fitted with the proposal+ppc-based Monte Carlo template. Combinations of $(\theta_{23}, \Delta m_{31}^2)$ are sampled from a regular grid and injected into the pseudodata as the true physics parameters driving the oscillations. These injected configurations are expected to be recovered perfectly by the fit with the proposal+ppc template if the switch to Geant4 makes no impact on the ν_μ disappearance analysis. This is not what we find from the inject-recovery test results in Figure 10.11. None of the scan points in the $(\theta_{23}, \Delta m_{31}^2)$ grid are recovered perfectly, exhibiting two systematic offset features not desirable for an oscillations measurement. The first one is the bias in the fitted mass splitting to lower values across the entire range of θ_{23}. The magnitude of the bias is $\gtrsim 5 \times 10^{-5}\,\mathrm{eV}^2$, while, for reference, the 1σ uncertainty on the atmospheric mass splitting estimated from the golden IceCube-DeepCore event sample [180] is $7 \times 10^{-5}\mathrm{eV}^2$. Another systematic feature is the pull of

---

[8]The overall template includes neutrinos of all flavors as well as the background muons, and only the muon light yield output for ν_μ-CC daughter muons is simulated in Geant4.
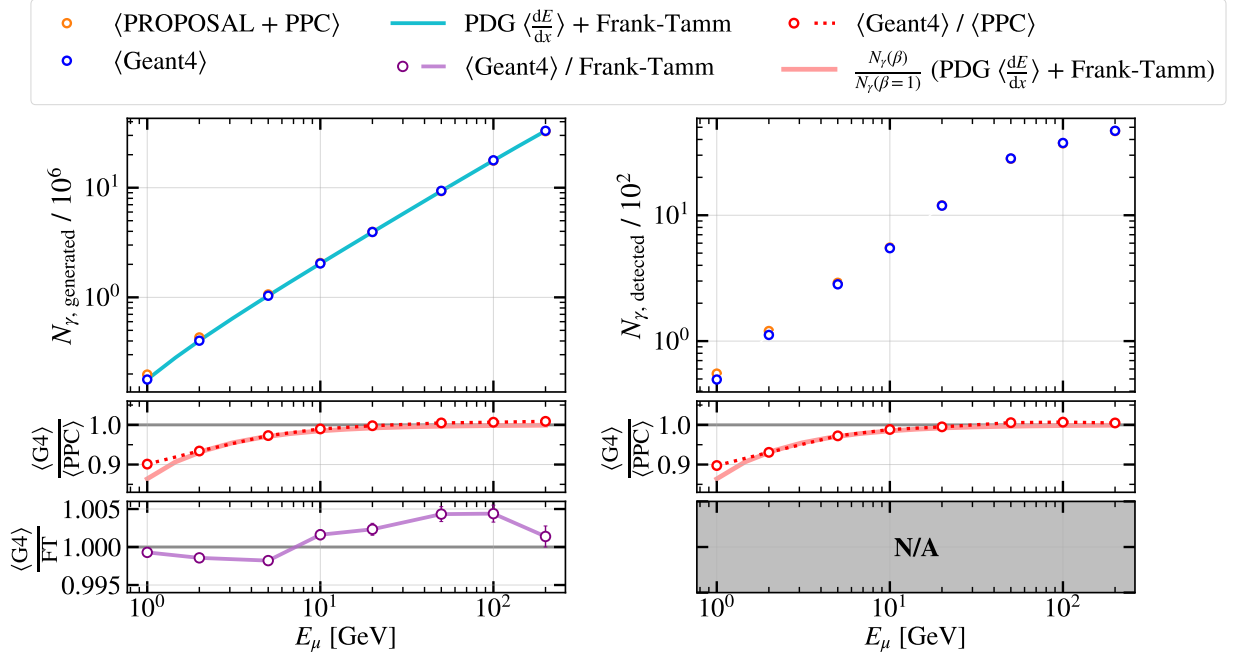
FIGURE 10.9 – *Top left*: average numbers of Cherenkov photons generated by muons with total energies E_μ, as extracted from the GEANT4 ("G4") and PROPOSAL+PPC ("PPC") simulations and compared to the Frank-Tamm prediction ("FT") with the full evolution of the muon Lorentz velocity factor β. *Top right*: same as top left, but for photons detected at the IceCube DOMs (no analytical prediction is shown for this case). *Middle*: ratio of the G4 and the PPC $N_{\gamma,\mu}$ averages for generated (left) and detected (right) photons, as compared to the ratio of $\beta(x)$ and the $\beta = 1$ cases expected from the Frank-Tamm formula. *Bottom left*: ratio of the G4-based $N_{\gamma,\mu}^{\text{generated}}$ distribution average to the Frank-Tamm prediction with $\beta(x)$.
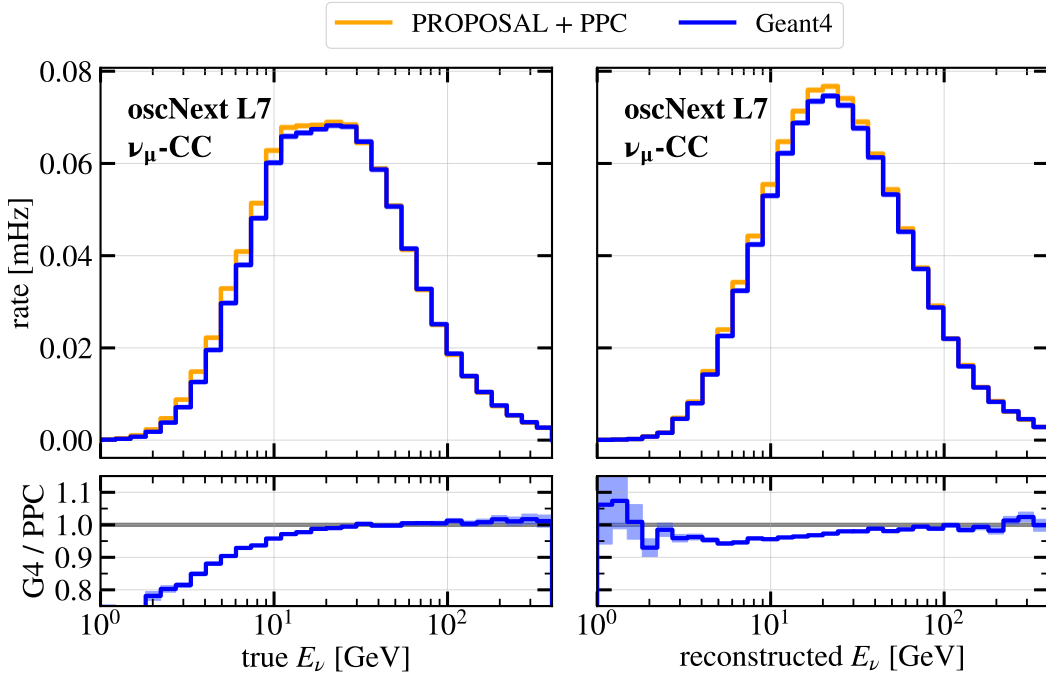

FIGURE 10.10 – Impact of switching from PROPOSAL+PPC ("PPC") to GEANT4 ("G4") simulation of the muon light yield on the distributions of true and reconstructed neutrino energies in the ν_μ-CC subsample of the Level 7 oscNext Monte Carlo.

the injected θ_{23} values towards maximal mixing (45°), particularly for the true values of θ_{23} in the upper octant currently preferred by the global fits for neutrino oscillations [65, 66].
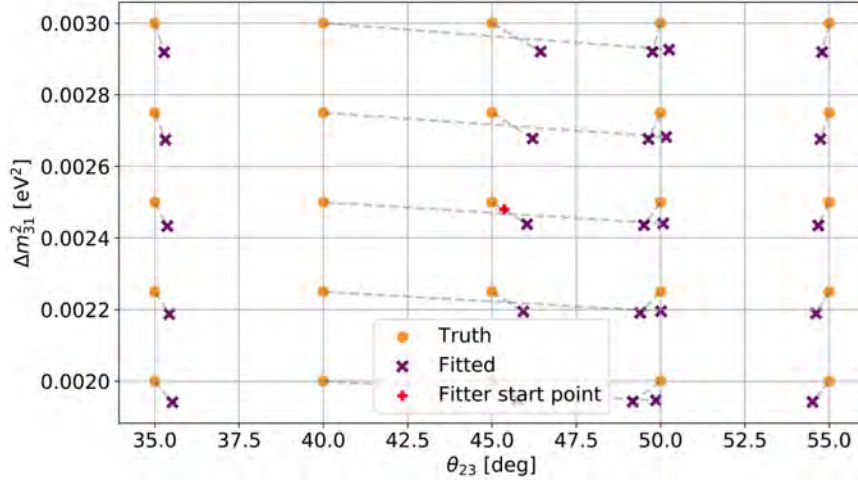
FIGURE 10.11 – Results of the inject-recovery test for the oscillation parameters $(\theta_{23}, \Delta m_{31}^2)$, which are injected into a pseudodata template with the GEANT4 muon light yield in the $(\nu_\mu + \bar{\nu}_\mu)$-CC simulation and fitted with the PROPOSAL+PPC-based Monte Carlo template. The dashed gray lines show the shift from the true set of parameters to the fitted one. *Figure courtesy of T. Stuttard.*

To quantify the amount of bias in the oscillation parameters measurement introduced by fitting the template with an overestimated muon light yield to the pseudodata with the corrected muon light yield, we perform an additional fit per each grid scan point in Figure 10.11, where the physics parameters $(\theta_{23}, \Delta m_{31}^2)$ are fixed at their true values. This forces the fitter to compensate for the muon light yield mismodeling using the systematic parameters only. The test statistic (log-likelihood, or LLH) is recorded both when $(\theta_{23}, \Delta m_{31}^2)$ are free and when they are fixed, and the resulting ΔLLH = LLH$_{\text{fixed}}$ − LLH$_{\text{free}}$ test statistic is calculated from each pair of fits. The fit with the fixed $(\theta_{23}, \Delta m_{31}^2)$ returns a lower likelihood, since it is allowed fewer "handles" to adjust the mismodeled muon light yield template. Thus, -2ΔLLH is positive-valued, with larger values indicating a greater impact of the muon light yield mismodeling on the oscillation parameters. The -2ΔLLH distribution constructed from the 25 grid scan points in Figure 10.11 is shown in Figure 10.12, where we also overlay the χ^2 critical values for two degrees of freedom[9].

We find that the entire -2ΔLLH distribution lies above the χ^2 critical value corresponding to 0.5σ, which implies that *at least* a 0.5σ bias is introduced to the measurement of $(\theta_{23}, \Delta m_{31}^2)$ if their true values lie anywhere on the conservatively wide $([35°, 55°], [2 \cdot 10^{-3}\,\text{eV}^2, 3 \cdot 10^{-3}\,\text{eV}^2])$ grid. Furthermore, nearly 50% of the trials exhibit biases at a level greater than 1σ, which is unacceptable for a precision measurement of oscillation parameters. We therefore conclude that ***the wrongly simulated muon light yield in the original oscNext simulation has a large impact on the ν_μ disappearance analysis***[10] and that switching to the GEANT4-based simulation[11] is absolutely necessary to avoid the demonstrated bias.

10.1.4 *Impact of the corrected muon light yield on the data/MC agreement*

Having produced analysis-ready neutrino Monte Carlo simulation with the GEANT4-based muon light yield, we can see whether it helps resolve the energy-dependent

---

[9]These χ^2 critical values are compared to the -2ΔLLH distribution assuming Wilks' theorem.

[10]The impact on the ν_τ appearance analysis is expectedly smaller ($\sim 0.5\sigma$ bias in the measured N_{ν_τ} normalization).

[11]Alternatively, any other simulation that would take into account the muon β evolution in the Cherenkov light yield calculations.
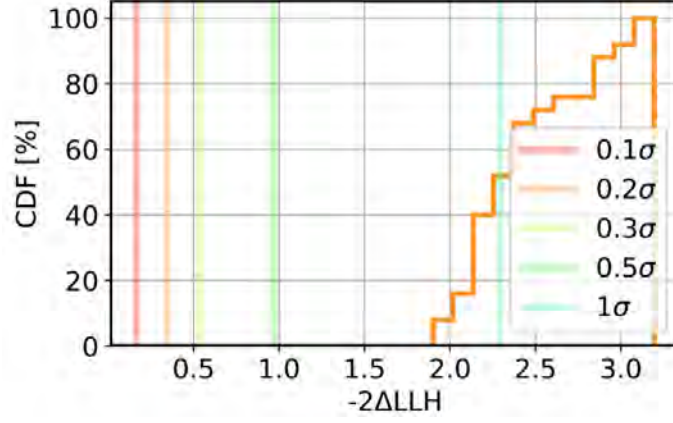
FIGURE 10.12 – The cumulative distribution function (CDF) of the $-2\Delta\text{LLH} \equiv -2(\text{LLH}_{\text{fixed}} - \text{LLH}_{\text{free}})$ test statistic distribution constructed from the $(\theta_{23}, \Delta m_{31}^2)$ grid scan points in Figure 10.11. To arrive at LLH_{free} ($\text{LLH}_{\text{fixed}}$), the two oscillations parameters are kept free (fixed) in a fit of the PROPOSAL+PPC-based template to the GEANT4-based pseudodata. All of the standard ν_μ disappearance analysis systematic parameters are kept free, and no statistical fluctuations are applied to the pseudodata template. *Figure courtesy of T. Stuttard.*

data/Monte Carlo agreement issues highlighted in Section 9.2. The updated data/MC ratios are shown in Figures 10.13 and 10.14 for the total reconstructed neutrino energy and the reconstructed track energy, respectively. Since the corrected muon light yield only affects true neutrino energies $\lesssim 20$ GeV and reconstructed neutrino energies $\lesssim 40$ GeV (see Figure 10.10), it is a natural outcome that the main "energy tilt" trend with a characteristic data excess at high energies has not been affected by switching to GEANT4. The data/MC ratio for the total reconstructed energy does, however, move closer to unity, and the overall pre-fit test statistic improves by nearly 60 units of log-likelihood (increasing from LLH =-626 in the PROPOSAL+PPC case to LLH =-568 in the GEANT4 case). This is driven by the decrease in the Monte Carlo event rate in the low-energy part of the reconstructed energy spectrum and the subsequent increase in the data/Monte Carlo ratio.

In conclusion, we have found significant issues in the low-energy muon propagation and the modeling of the muon light yield in the IceCube Monte Carlo simulation. This issue affects not only the oscNext event selection and the ν_μ disappearance analysis it set out to perform, but also all other IceCube analyses relying on charged current ν_μ events in the 1–20 GeV region. For the event selection described in this thesis, a decision has been made to switch to purely GEANT4-based Monte Carlo for muon propagation and Cherenkov light yield generation, which resolves the identified issues in simulation but does not fully mitigate the energy-dependent data/Monte Carlo disagreement. At the same time, a similar analysis closely resembling oscNext up to Level 5 has been recently unblinded without the muon light yield correction [197]. As an alternative to GEANT4, future analyses may consider implementing the muon β dependence into the PPC code in CLSim and adjusting the PROPOSAL energy loss settings to avoid the artefacts in the muon stopping range.

10.2 LIGHT YIELD FROM ELECTROMAGNETIC SHOWERS

Given that our investigation into the muon light yield simulation in IceCube revealed several major issues, we now proceed with the validation of the Cherenkov light yield from *electromagnetic (EM) showers*, i.e., cascades of particles initiated by an electron, a
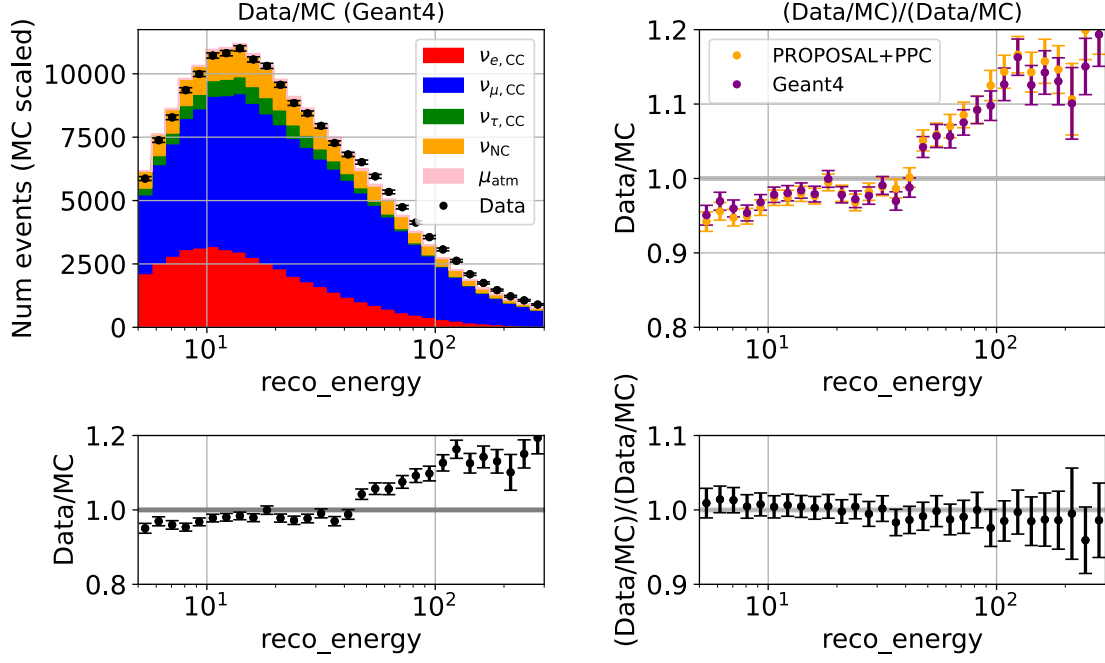
FIGURE 10.13 – Impact of switching from PROPOSAL+PPC to GEANT4 in the simulation of muon light yield in the ($\nu_\mu + \bar{\nu}_\mu$)-CC events on the data/Monte Carlo agreement in the total reconstructed neutrino energy variable ("reco_energy", shown in GeV).
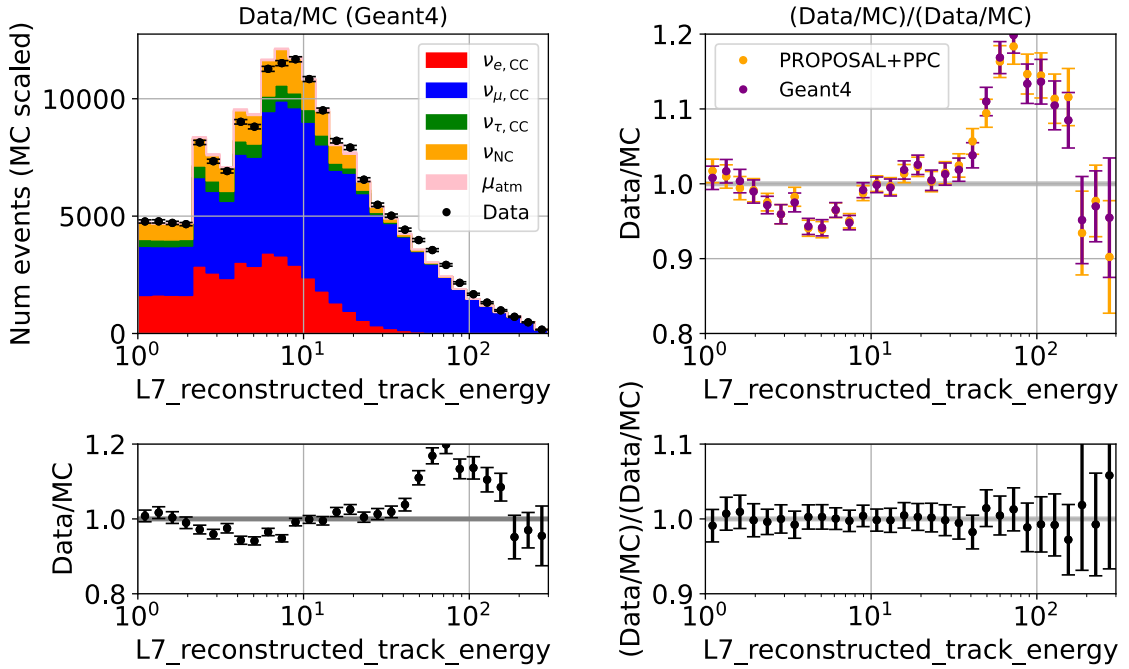

FIGURE 10.14 – Same as Figure 10.13, but for the reconstructed track energy variable ("L7_reconstructed _track_energy", shown in GeV).

positron, or a photon. All three of these particles can be produced as secondaries due to stochastic energy losses of high-energy muons. Additionally, electrons and positrons can appear directly from ν_e-CC and $\bar{\nu}_e$-CC interactions, while photon production can also result from the $\pi^0 \rightarrow 2\gamma$ decay. In [226], $e^\pm$ and γ primaries were shown to give identical light yield outputs as a function of energy from 1 GeV to 10 TeV. Therefore, in this study, we focus our attention only on *electrons* as representative electromagnetic shower primaries.

In the original oscNext simulation, the light yield from EM cascades was modeled following an outdated GEANT4-based parameterization from [225]. In this initial 2002 study by M. Kowalski, which we will refer to as the "reference (ref.) study A," the EM cascades were simulated in GEANT4, and all of the charged cascade particles i produced in a shower were tracked until stopping. The total distance ℓ_i traveled by each particle was recorded and corrected for the starting particle velocity β_i to calculate the effective Cherenkov-radiating track length, $\hat{\ell}_i$:

$$\hat{\ell}_i = \frac{1 - \frac{1}{n^2 \beta_i^2}}{1 - \frac{1}{n^2}} \cdot \ell_i, \tag{10.2}$$

with n being the refractive index of ice. These effective $\hat{\ell}_i$ were finally summed across all charged particles in the shower to arrive at the total Cherenkov-radiating track length, $\hat{\ell}$. By fitting the resulting $\hat{\ell}$ as a function of primary energy E with a simple power law,

$$\hat{\ell}(E) = \alpha_{\text{ref. A}} E^{\beta_{\text{ref. A}}}, \tag{10.3}$$

the following coefficients were derived:

$$\alpha_{\text{ref. A}} = 5.21 \, \text{m} \, \text{GeV}^{-1}; \; \beta_{\text{ref. A}} \approx 1, \tag{10.4}$$

implying a linear dependence of the Cherenkov-radiating track length on the primary energy. The reference study A was performed with the ice density $\rho_{\text{ref. A}} = 0.924 \, \text{g} \, \text{cm}^{-3}$ in the GEANT4 simulation. The ice density in the original oscNext simulation was slightly lower, $\rho_{\text{oscNext}} = 0.9216 \, \text{g} \, \text{cm}^{-3}$, and therefore the effective Cherenkov-radiating track length was parameterized in the PPC code within CLSIM as[12]

$$\hat{\ell}(E) = \alpha_{\text{ref. A}} \frac{\rho_{\text{ref. A}}}{\rho_{\text{oscNext}}} E. \tag{10.5}$$

Equation (10.5) was finally combined with the average number of photons expected per unit length from the Frank-Tamm formula (Equation (4.40)) with $\beta = 1$ – here being a valid assumption given that $\hat{l}_i$ has already been corrected for the β_i factor of each particle.

To validate the PPC parameterization of the electron light yield, we design a study similar to that described in Section 10.1. Specifically, we produce small Monte Carlo samples of fixed-energy electrons, $E_e \in [0.5 \, \text{GeV}, 1 \, \text{GeV}, 5 \, \text{GeV}, 10 \, \text{GeV}, 50 \, \text{GeV}, 100 \, \text{GeV}]$, with 1000 electrons per energy. We then either use PPC at Step 2 of the simulation to generate photons according to Equations (4.40) and (10.5), or resimulate the entire shower with Cherenkov photon emission directly in GEANT4. The geometry of the simulation (i.e., the electron direction and starting position) remains the same as in Section 10.1. The comparison of the detected number of photons ($N_{\gamma,e}^{\text{detected}}$) between the two simulation approaches is shown in Figure 10.15. We find that GEANT4 agrees well with PPC for $E_e \leq 10 \, \text{GeV}$, while at 50 GeV and 100 GeV, we observe a discrepancy in the form of $N_{\gamma,e}^{\text{detected}}$ being shifted to lower values in GEANT4 relative to PPC. This result is surprising, given that the PPC parameterization originates from the reference study A similarly based on GEANT4. However, one component of the reference study A was not explicitly recorded in [225], namely which index of refraction was used in the Frank-Tamm formula. The direct successor study of [225] performed by L. Radel [226], which

---

[12]This linear density rescaling stems from the linear dependence of the average Bethe-Bloch energy losses on the density of the medium [34, 281]. If the oscNext simulation had a higher density than the reference study A, the particle would travel a shorter distance, which is why ρ_{oscNext} appears in the denominator.
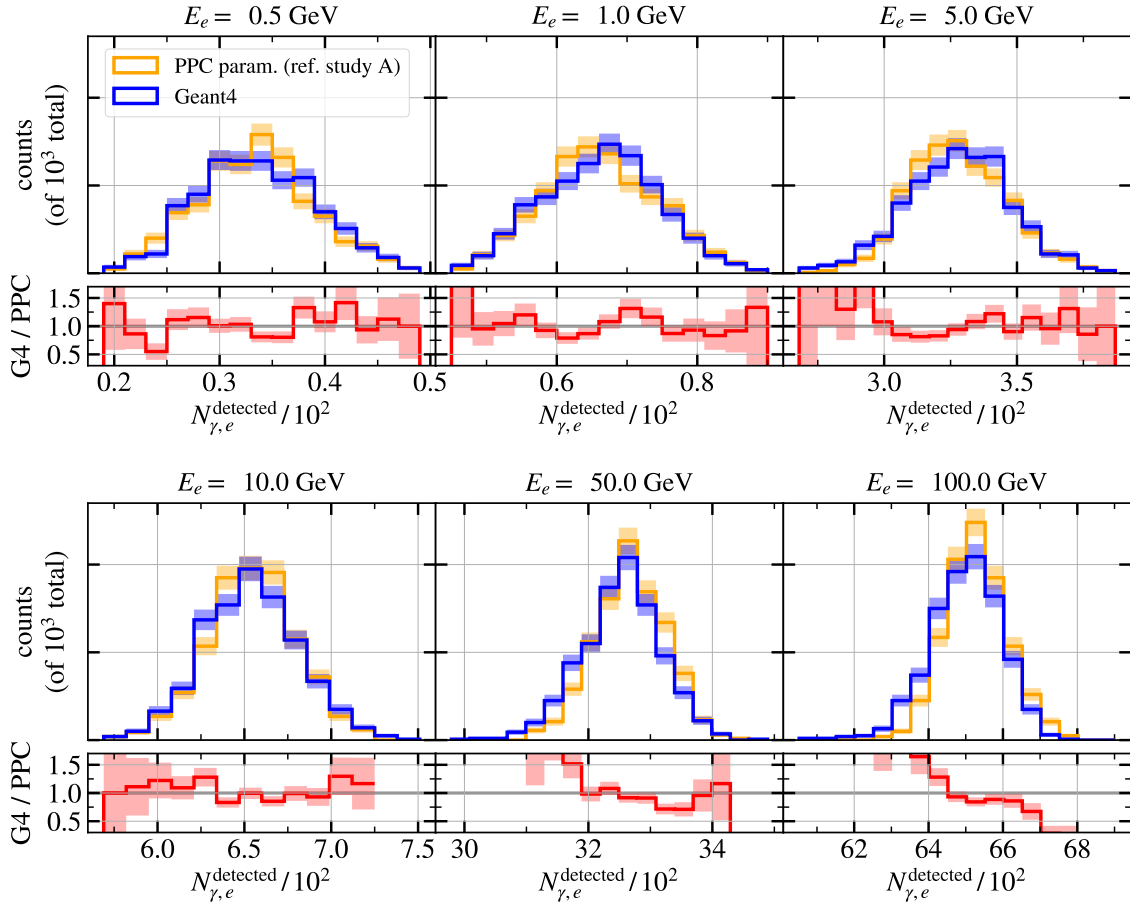
FIGURE 10.15 – Number of Cherenkov photons detected by the IceCube DOMs due to the passage of an electron in ice with a fixed total energy E_e, as compared between the approach originally used in the oscNext simulation (PPC parameterization, shortly "PPC" [225]) and the GEANT4 simulation (shortly "G4"). The y-axis represents the probability density of a given number of photons as obtained from a Monte Carlo simulation of 1000 electron events per panel.

we will call "reference study B,"[13] reports having used a constant value of $n = 1.33$ in the GEANT4 simulation, i.e., disregarding the wavelength dependence of the refractive index. We therefore hypothesize that the reference study A had also implemented a constant n, while our GEANT4 simulation takes into account its wavelength dependence as per Figure 4.9. To see whether changing from a constant to a wavelength-dependent index of refraction can result in an effect similar to that observed in Figure 10.15, we run an additional set of GEANT4 simulations with $n = 1.33$. The results are presented in Figure 10.16. We indeed observe a qualitatively similar effect to that in Figure 10.15, albeit with an even stronger shift of the $N_{\gamma,e}^{\text{detected}}$ distributions relative to one another, which appears already at 10 GeV. We conclude that the constant index of refraction is most likely the culprit of the high-energy discrepancy between the GEANT4 and the PPC simulations of the electron light yield, with the disclaimer that the passage of time has made the exact reproduction of the reference study A impossible.

Given the potential of this discrepancy to affect the high-energy part of the reconstructed neutrino energy spectrum, we attempt switching to GEANT4 for Cherenkov photon production due to EM cascades in the main oscNext simulation chain. This large-scale simulation production was executed by T. Stuttard for all neutrino flavors. The impact on the true and reconstructed energies of ν_e-CC events is shown in Fig-

---

[13]Note that reference study B derived a different set of coefficients for Equation (10.5), namely $\alpha_{\text{ref. B}} = 5.321\,\text{GeV m}^{-1}$ given $\rho_{\text{ref. B}} = 0.91\,\text{g cm}^{-3}$. The difference between $\alpha_{\text{ref. A}}\rho_{\text{ref. A}}$ and $\alpha_{\text{ref. B}}\rho_{\text{ref. B}}$ is ~0.8%.
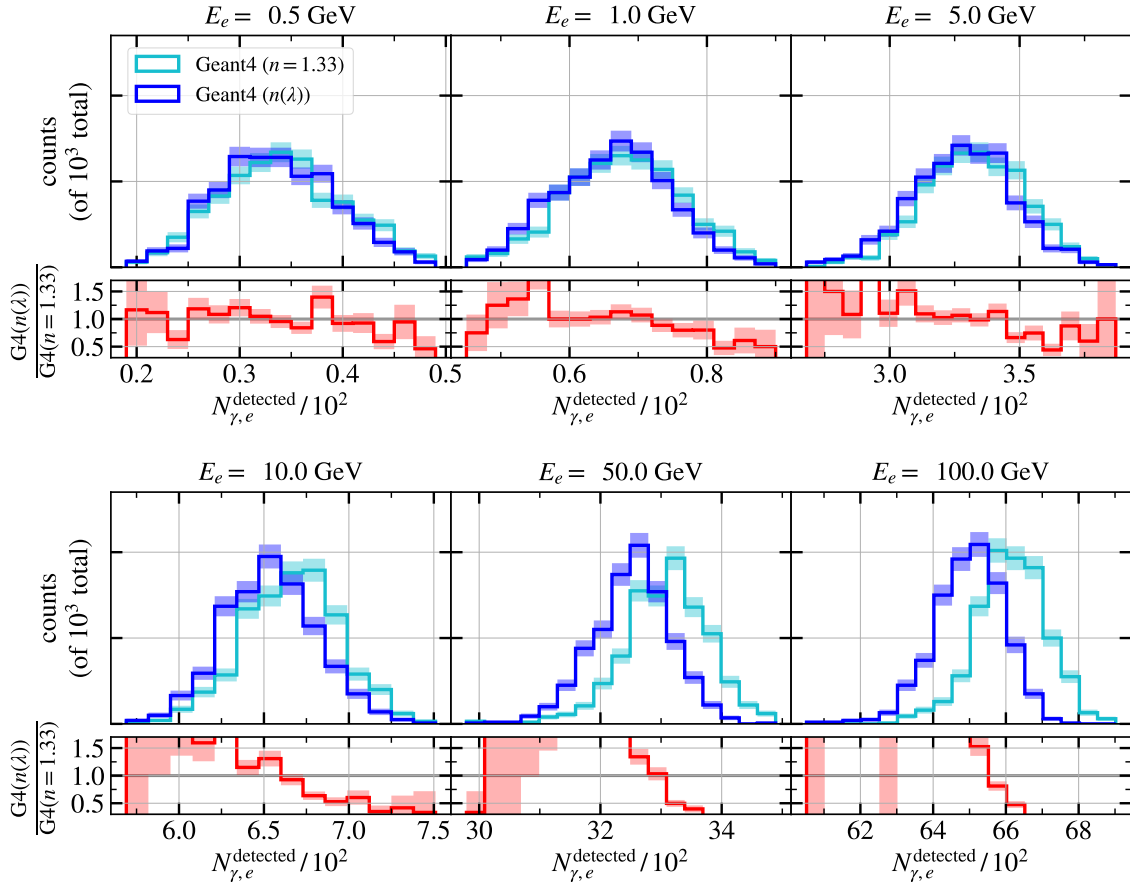
FIGURE 10.16 – Same as Figure 10.15, but compared between the GEANT4 simulation with the full wavelength dependence of the refractive index $n(\lambda)$ (see Figure 4.9) and the simplified assumption of a constant $n = 1.33$.

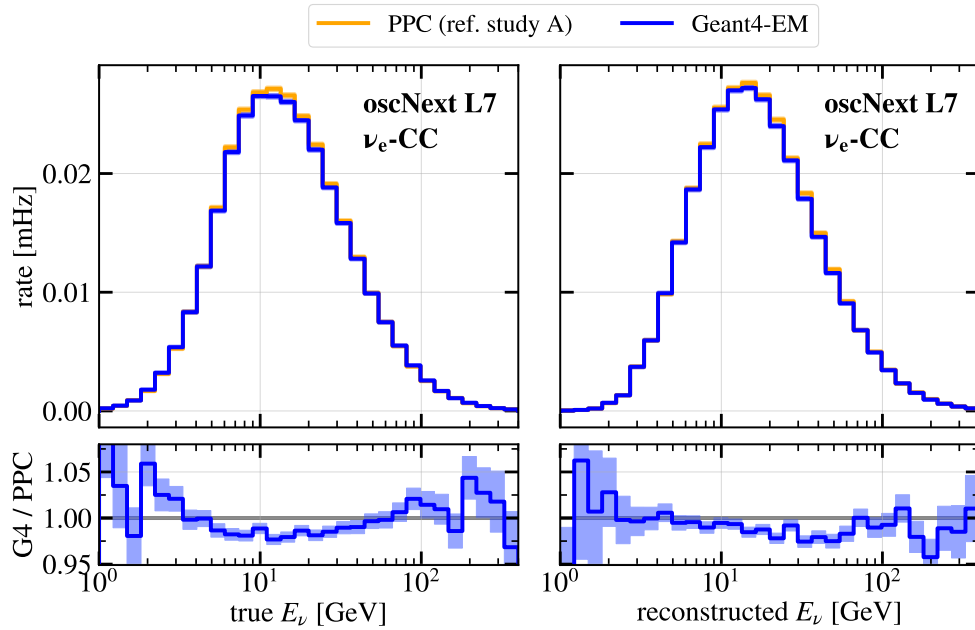ure 10.17. While switching to GEANT4 has modified both true and reconstructed energy

FIGURE 10.17 – Impact of switching from the PPC parameterization [225] ("PPC") to GEANT4 ("G4") simulation of the Cherenkov light yield in electromagnetic cascades on the distributions of true and reconstructed neutrino energies in the ν_e-CC subsample of the Level 7 oscNext Monte Carlo.

distributions in a non-trivial way complicated by the selection cuts, the magnitude of

the EM light yield correction impact is at most a few %. Since this is much smaller than what we previously found in the context of muon light yield (see Figure 10.10), we expect that the EM light yield correction will affect the data/MC ratios for reconstructed energy to a proportionally smaller degree. This is what we ultimately find in Figure 10.18, noticing negligible impact of the GEANT4-based EM light yield Monte Carlo on the data/MC ratios for reconstructed energy. For completeness, we also test the case where both the muon light yield and the EM light yield are simulated in GEANT4, but find no improvement compared to the muon light yield-only case.
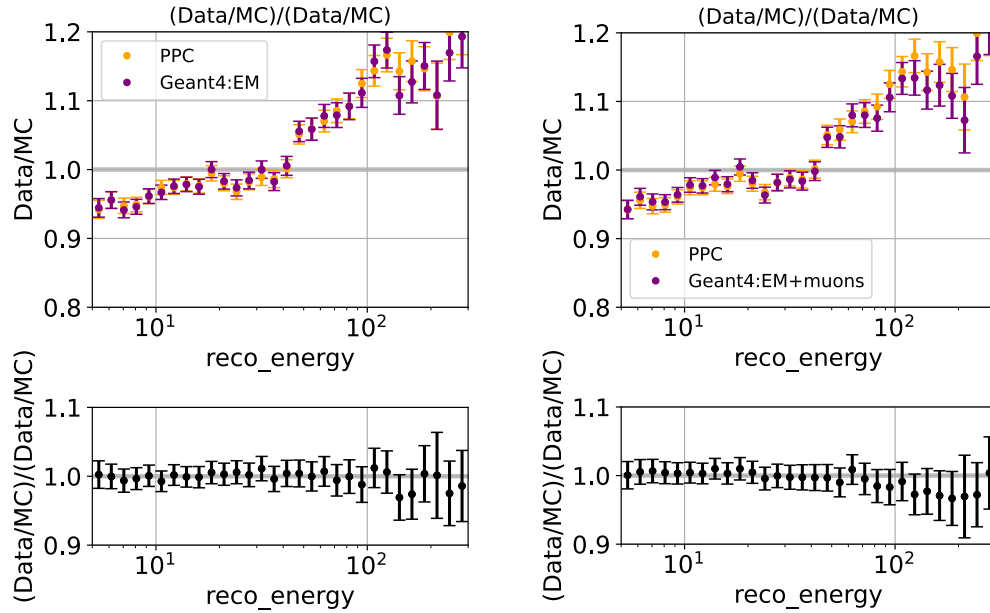

FIGURE 10.18 – Impact of switching from the PPC parameterization [225] ("PPC") to GEANT4 ("G4") simulation of the Cherenkov light yield in electromagnetic cascades on the data/Monte Carlo ratios of the 1D reconstructed energy distributions ("reco_energy", shown in GeV).

Given the near-negligible impact of the EM light yield correction on the agreement between data and Monte Carlo, *we decide not to incorporate the GEANT4-based simulation of the EM cascades into the oscNext simulation sets at the present moment*. However, we warn future analyzers that the error in the Cherenkov light yield due to the simplified assumption of the constant index of refraction grows with energy. Therefore, we strongly recommend that any studies relying on cascade events with energies above 100 GeV either update the PPC parameterization based on GEANT4 simulations with the full wavelength dependence of the refractive index, or run the latter directly during simulation production.

10.3 TRACK ENERGY RECONSTRUCTION WITH THE RETRO ALGORITHM

10.3.1 *Misreconstructed events and their characteristics*

We now consider another possible source of data/Monte Carlo disagreement in the reconstructed energy variables, namely the reconstruction technique itself. As of 2022, the likelihood-based RETRO algorithm [211] was the nominal reconstruction method for the oscNext high-statistics sample analysis variables, including, among others, track and cascade energy (see Section 6.4). Despite this algorithm being based on lookup tables

produced with an outdated ice model[14] and assuming muon energy losses in water rather than ice, the reconstruction bias caused by such limitations should be equally applicable to data and simulation without introducing a relative discrepancy between the two. In addition, the bias in the reconstructed *vs.* true energy was addressed explicitly in [211], where correction functions were devised for reconstructed track and cascade energies to make their sum representative, on average, of the true total deposited energy. The mechanism through which the data/MC disagreement can be introduced or exacerbated at the reconstruction stage is therefore highly non-trivial.

To shed some light on whether the energy reconstruction behaves as expected in the region of strong data/MC disagreement, we undertake the tedious task of visually inspecting such events and contrasting their reconstructed properties with the recorded pulse maps. We observe from Figures 9.9 and 9.10 that events with the reconstructed track energy ($E_{\text{track}}^{\text{reco}}$, or "L7_reconstructed_track_energy") greater than ~30–40 GeV exhibit a particularly bad mismatch between data and Monte Carlo rates. Although this boundary is not clearly defined, we focus on the region $E_{\text{track}}^{\text{reco}} \geq 30\,\text{GeV}$ in the mixed and track PID bins to alleviate the search for any extreme misreconstruction cases. This search turns out to be fruitful, as we are able to ***identify a class of misreconstructed events with a significantly overestimated track energy***. In Figure 10.19, we show a representative example of such an event – a ν_μ-CC interaction in which the true energy of a daughter muon (2.0 GeV) was reconstructed to as much as ~56 GeV.
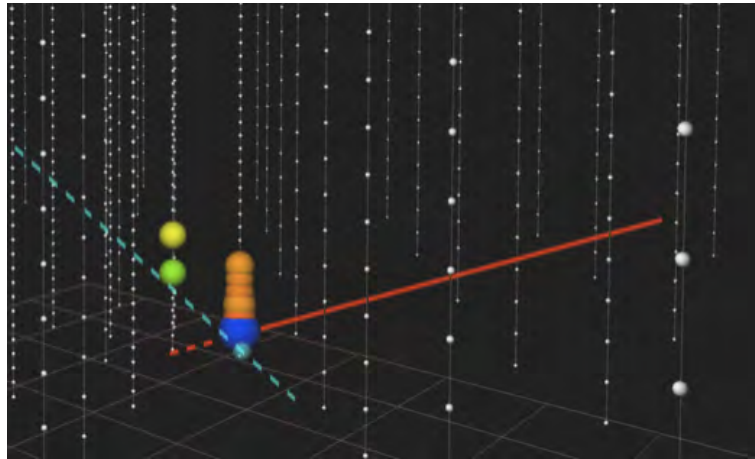

FIGURE 10.19 – Example event from the ν_μ-CC oscNext Level 7 Monte Carlo with a bad track energy reconstruction returned by the RETRO algorithm [211]. The dashed blue line shows the true particle direction, while the solid red line represents the reconstructed track direction and length. The true track (total) energy in this event is 2.0 GeV (9.0 GeV), while the reconstructed track (total) energy is 55.6 GeV (72.1 GeV). The image contains the SRTTWOfflinePulsesDC pulse series, which were used as input to the RETRO reconstruction, visualized here with the STEAMSHOVEL IceCube software.

A characteristic feature of such events is an unjustifiably long "ghost track" extended into a region of the detector where no light has been deposited. The likelihood optimization algorithm arrives at such a reconstructed event topology by directing the tracks into sparsely instrumented regions, such that the absence of light deposition can be explained by a large distance to the closest strings. Since this requires a very specific directional configuration of an event, the misreconstructed track energy is often accompanied by a significant mismatch between the true and the reconstructed directions. More examples that support this finding are presented in Table 10.1.

---

[14]The RETRO tables were produced with the Spice Lea model [283], which explains the ice optical anisotropy (i.e., anisotropy in the photon propagation direction) through modified scattering coefficients. It precedes the most recent ice model reported in publication by IceCube [240], where this anisotropy is modeled through intrinsic birefringence of ice.
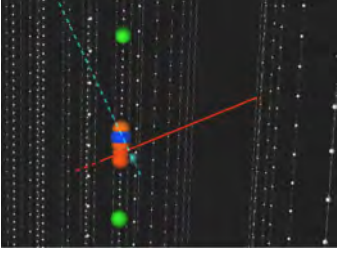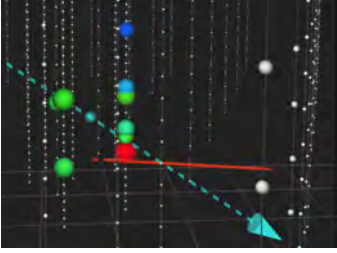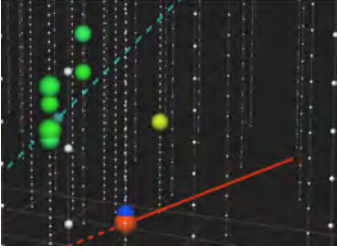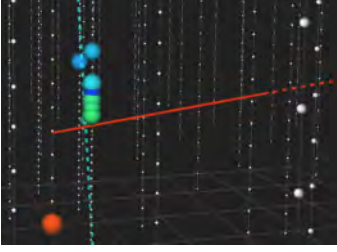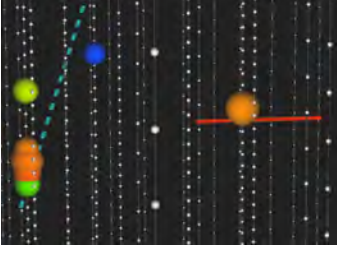
Event display	Event type	$E_{\text{track}}^{\text{true}}$ [GeV]	$E_{\text{track}}^{\text{reco}}$ [GeV]	$E_{\text{total}}^{\text{true}}$ [GeV]	$E_{\text{total}}^{\text{reco}}$ [GeV]
	ν_μ-CC	1.2	32.7	2.6	33.5
	ν_μ-CC	4.7	31.1	7.1	31.1
	ν_μ-CC	3.0	34.0	6.5	34.4
	ν_μ-NC	0.0	44.7	9.3	44.7
	ν_e-CC	0.0	42.0	4.3	42.0
	ν_τ-CC	0.0	34.3	4.5	36.0

TABLE 10.1 – Representative examples of Monte Carlo events with bad track energy reconstructions returned by the RETRO algorithm [211] (see Figure 10.19 for details). Note that non-ν_μ-CC events do not contain a true track by definition, yet their reconstructed track energy is several tens of GeV.

In Monte Carlo simulation, such misreconstructed events are by definition easy to spot by searching for significant discrepancies between the true and the reconstructed energies or directions. However, in data the truth is unknown, and these events can only be identified using the reconstructed metrics. One such metric could be the reconstructed radial stopping coordinate of an event ("L7_reconstructed_stopping_rho", or $\rho^{\text{reco}}_{36,\text{stopping}}$) – i.e., the radial distance of the stopping point to the center of the detector (string 36). The radius of the DeepCore fiducial volume is 150 m, which, as shown in Figure 10.20, covers all HQE DeepCore strings and one outer ring of NQE strings. Any events starting and stopping within with volume are considered fully contained.
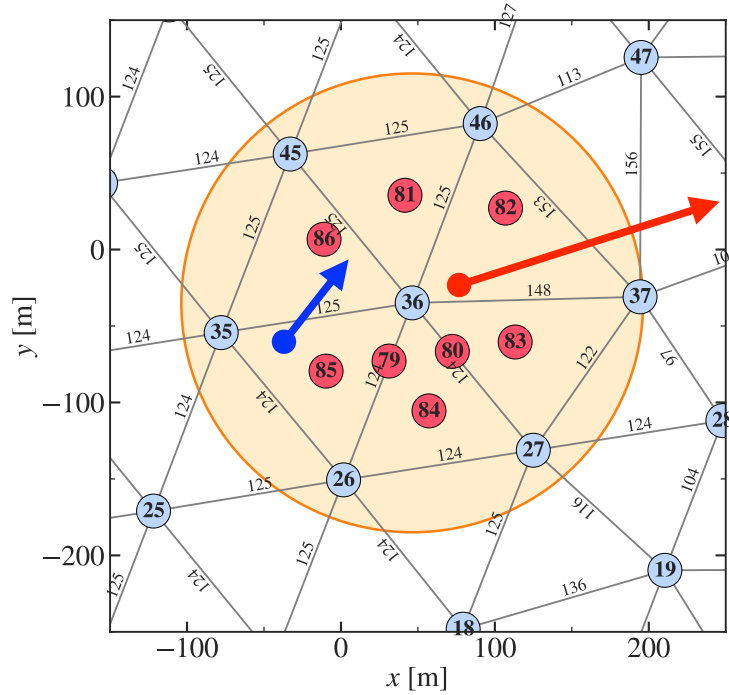

FIGURE 10.20 – A zoom-in into the xy projection of the DeepCore fiducial volume, which is bounded by the orange circle of a 150 m radius relative to string 36. The high quantum efficiency (HQE) strings are shown in red, and the normal quantum efficiency (NQE) strings – in light blue. An example event whose stopping point is contained (uncontained) within the DeepCore fiducial volume is represented schematically with the blue (red) arrow.

In Figure 10.21, we show that events with reconstructed track energies ≥30 GeV are more often uncontained within the DeepCore fiducial volume ($\rho^{\text{reco}}_{36,\text{stopping}} > 200$ m) in data compared to Monte Carlo. This is consistent with our observation that "ghost tracks" appear to escape to more sparsely instrumented regions to explain the non-detection of any light that they could have produced.

Another characteristic feature of the "ghost track" events is that they are much too dim to justify their reconstructed track and total energy. For example, the event in Figure 10.19 has only 8 hit DOMs in the cleaned pulse series used for the RETRO reconstruction, which is very low compared to the reconstructed track energy of 56 GeV. For reference, we show a well-reconstructed track in Figure 10.22 ($E^{\text{reco}}_{\text{track}} = 48.5$ GeV, $E^{\text{true}}_{\text{track}} = 47.6$ GeV), whose cleaned pulse series contains 36 hit DOMs. This suggests that we can use the number of hit DOMs per unit track energy ($n_{\text{hit DOMs}}/E^{\text{reco}}_{\text{track}}$, or "num_hit_doms_per_GeV_track") to partially identify the misreconstructed events. The distributions of this variable are compared between data and Monte Carlo in Fig-
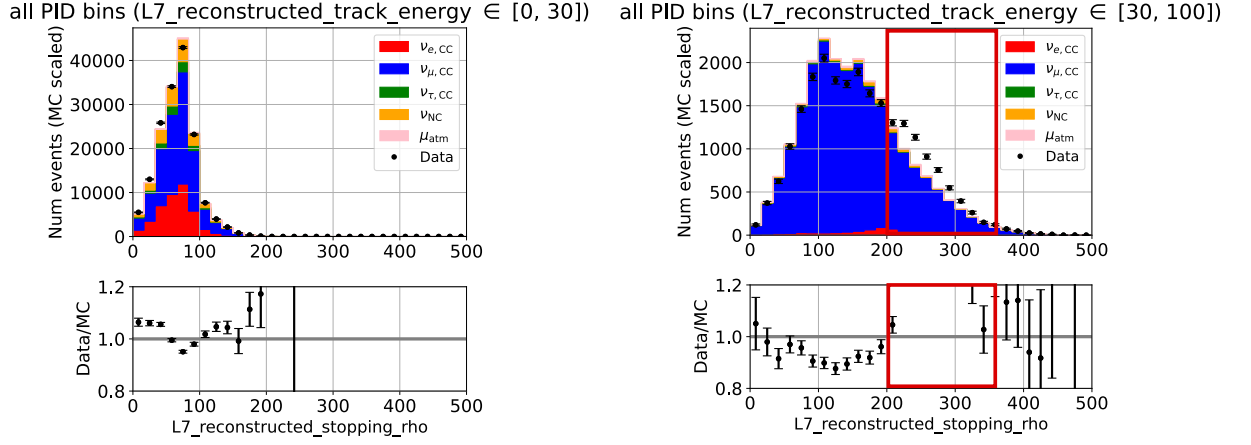
FIGURE 10.21 – Data/Monte Carlo agreement in the reconstructed radial stopping coordinate variable ("L7_reconstructed_stopping_rho", shown in meters) for events with $E_{\text{track}}^{\text{reco}} < 30\,\text{GeV}$ (left panel) and $30\,\text{GeV} \le E_{\text{track}}^{\text{reco}} \le 100\,\text{GeV}$. The region of a particularly strong data/MC disagreement, which is consistent with our discovery of the population of "ghost tracks" as in Figure 10.19 and Table 10.1, is highlighted with the red box.
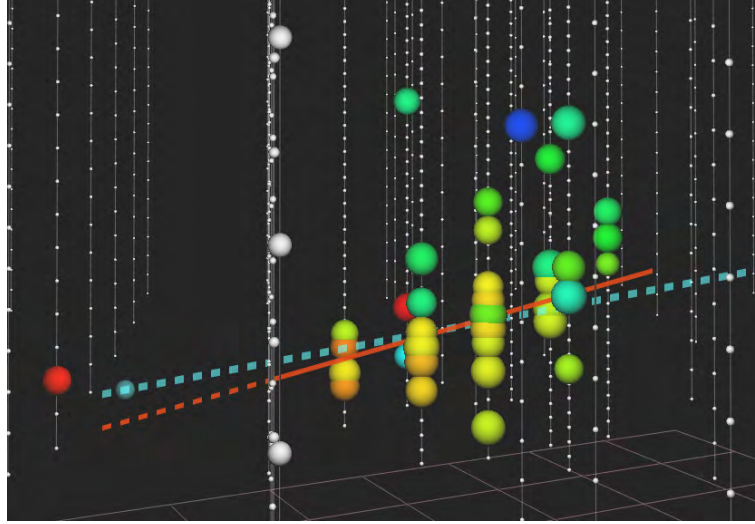
FIGURE 10.22 – Example event from the ν_μ-CC oscNext Level 7 Monte Carlo with a good track energy reconstruction returned by the RETRO algorithm [211]. The true track (total) energy in this event is 47.6 GeV (47.8 GeV), while the reconstructed track (total) energy is 48.5 GeV (56.4 GeV). For legend details, see Figure 10.19.

ure 10.23. We find a particular excess of events in data with $n_{\text{hit\,DOMs}}/E_{\text{track}}^{\text{reco}} \le 0.3$, which supports our discovery of very dim events with high reconstructed track energies.

Finally, we can characterize not only the overall the brightness of the event relative to the total reconstructed track energy, but also the spatial distribution of pulses compared to the reconstructed track geometry. To accomplish this, we project the pulses in the SRTTWOfflinePulsesDC pulse series onto the reconstructed track, which we slice into the segments of 10 m length. We then count the total charge[15] of the projected pulses within each track segment and build one-dimensional track light profiles, as shown in Figure 10.24. We find that in a typical misreconstructed event, the majority of charge accumulates at the very beginning of the reconstructed track, while the remaining

---

[15]For this calculation, the charge of each pulse is rounded to the nearest integer.

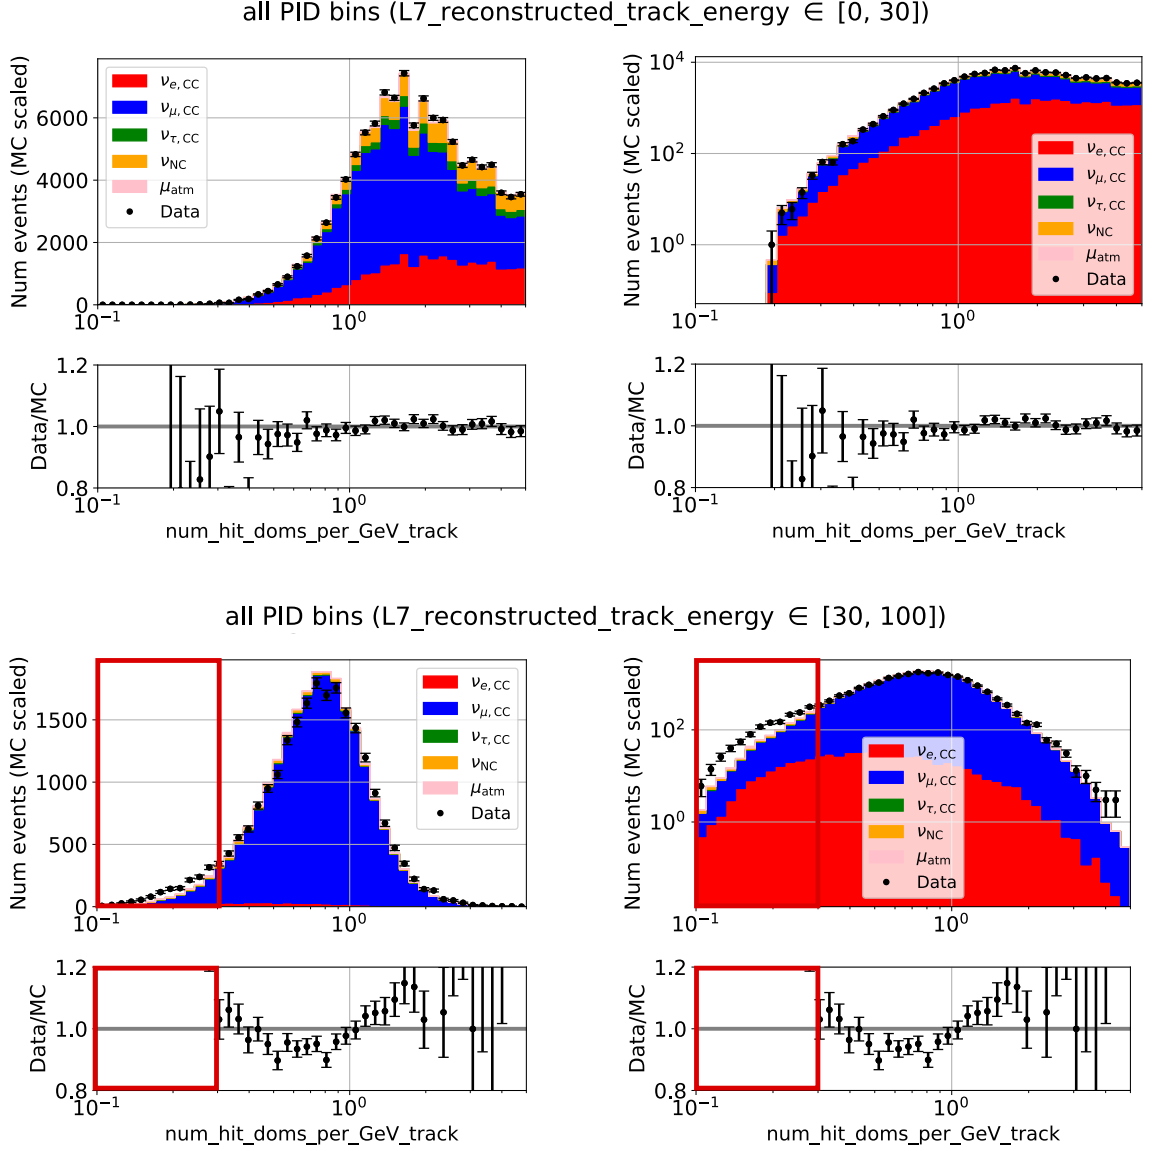FIGURE 10.23 – Number of hit DOMs in the SRTTWOfflinePulsesDC pulse series used in the RETRO reconstruction [211] per GeV of the RETRO-reconstructed track energy, as compared between data and oscNext Monte Carlo simulation. The top (bottom) panel shows events with $E_{\text{track}}^{\text{reco}} < 30\,\text{GeV}$ ($30\,\text{GeV} \leq E_{\text{track}}^{\text{reco}} \leq 100\,\text{GeV}$). The difference between the left and the right columns is in the scale of the y-axis (linear *vs.* logarithmic). The region of a particularly strong data/MC disagreement is highlighted with the red box.

fraction of the reconstructed track length is virtually empty of light deposition and therefore recorded charge. This is illustrated in the left column of Figure 10.24, where 90% of the total charge is collected within ~13% of the reconstructed track. On the other hand, a typical well-reconstructed event has a nearly uniform light deposition throughout the track. This is shown in the right column of Figure 10.24, where 90% of the total charge is collected within 115% of the reconstructed track length[16]. Thefore, the track length fraction within which 90% of the charge[17] is accumulated can be used as a complementary track energy/length reconstruction quality metric. In Figures 10.25

[16]Note that the light/charge deposition can continue well beyond the track stopping point, as the photons emitted immediately prior to the muon stopping continue propagating until they are absorbed or detected.

[17]Alternatively, any other preferred percentage representative of the *majority* of the charge.
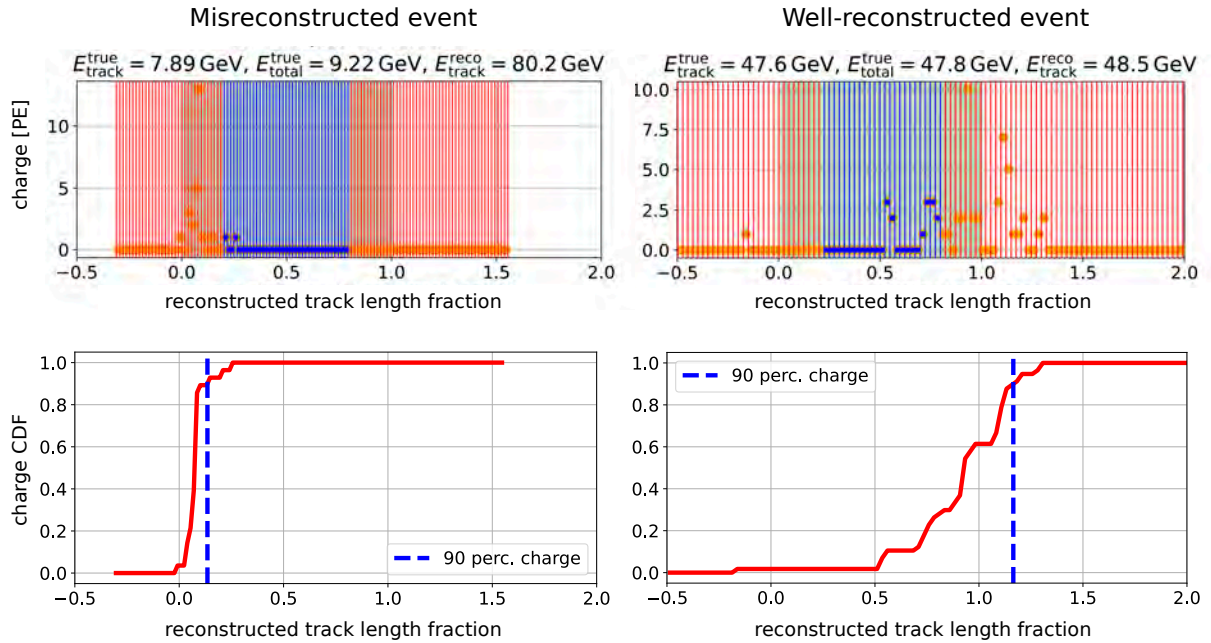
FIGURE 10.24 – *Left column*: one-dimensional projection of the track light deposition profile of a typical event with an overestimated track energy and length. *Right column*: 1D track light deposition profile of a typical well-reconstructed event. The *x*-axis in all plots shows the fraction of the reconstructed track length, where 0 represents the interaction vertex, and 1 – the reconstructed stopping point. Values below 0 (above 1) correspond to the regions of the reconstructed track that lie before the interaction vertex (after the stopping point). The vertical lines in the top panel divide the track into 10 m segments. Inside (outside) of the inner 60% of the track, the charge deposition is plotted in blue (orange). The bottom panel shows the cumulative distribution function (CDF) of the collected charge, which we use to find the track length fraction within which 90% of all charge is accumulated.

and 10.26, we show the agreement of this variable between data and Monte Carlo.
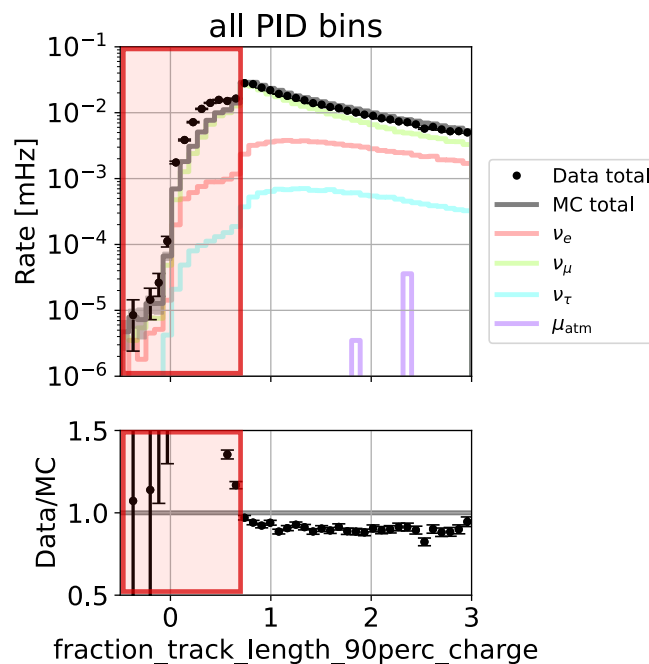
FIGURE 10.25 – Fraction of the reconstructed track length within which 90% of the total event charge is accumulated (see text and Figure 10.24 for details), shown here for all PID bins combined. The red box highlights the region of a particularly strong data/Monte Carlo disagreement in this variable.

FIGURE 10.26 – Same as Figure 10.25, but split between cascades (PID ∈ [0, 0.5)), mixed (PID ∈ [0.5, 0.85)), and tracks (PID ∈ [0.85, 1]) PID bins.

As expected, we find a significant excess of the data rates at low (≲ 0.7) values of the reconstructed track length fraction that contains the majority of the event charge.
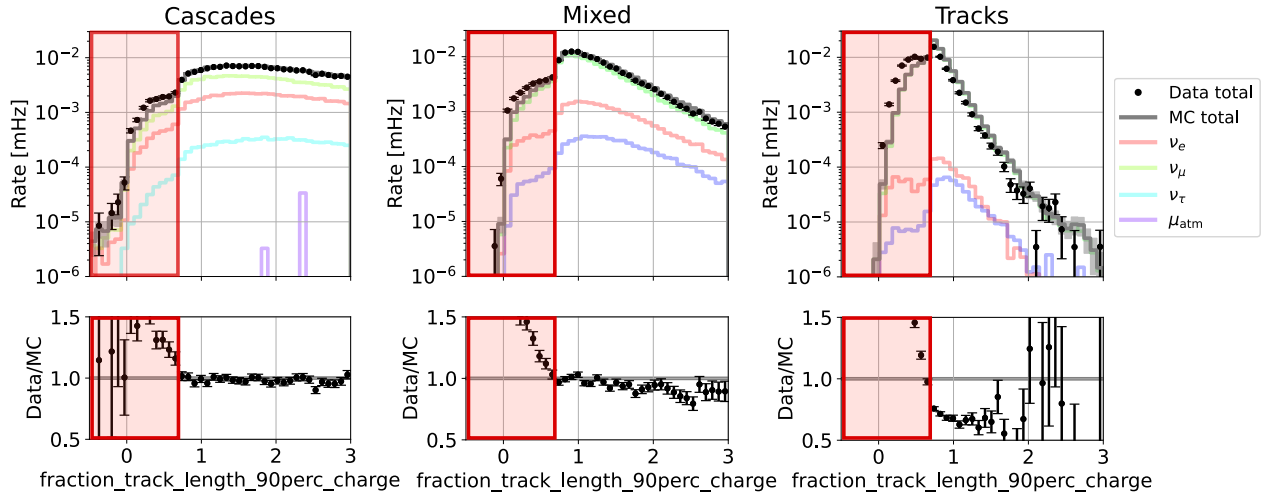
To summarize, we have identified a population of misreconstructed events in both data and Monte Carlo, where the track energy appears to be significantly overestimated. The radial stopping coordinate, the number of hit DOMs per unit track energy, and the fraction of the reconstructed track length within which the majority of all charge is collected can all be used as track energy reconstruction quality metrics. The regions of bad data/MC agreement in these variables correlate with an excess of events with high (≳30–40 GeV) reconstructed track energies in data. However, these derived variables are merely *symptoms* of the track energy misreconstructions and do not point to the physical origin of a significantly higher rate of misreconstructions in data. ***It is essential to find the root cause of the reconstruction failures***, since its effects may reach far beyond the variables discussed in this section and affect the data/Monte Carlo agreement in a continuous fashion throughout the entire phasespace. In the subsequent section, we present evidence that noise is the origin of track energy misreconstructions with the RETRO algorithm. We further summarize the follow-up work of T. Stuttard that definitively proves that PMT noise is not simulated accurately in the IceCube Monte Carlo, causing the data/MC disagreement in the reconstructed track energy.

10.3.2 *Noise as the origin of track energy reconstruction failures*

A careful visual inspection of events with failed track energy reconstructions, such as those in Table 10.1, points to a correlation between the reconstruction failures and the presence of very late pulses in the pulse series. Since these individual pulses appear temporally disconnected from the rest in a given event, we hypothesize that they are not a part of the *signal pulses*, which result directly from Cherenkov photons hitting the PMT cathode. Instead, we suggest that these late pulses are a form of noise, to which the likelihood reconstruction in a true low-energy event with only a few signal pulses is particularly sensitive. To test this hypothesis, we compare the reconstructions of identical events where the input for the reconstruction is either (a) the nominal (SRTTWOfflinePulsesDC) pulse series or (b) the same pulse series with the true noise

pulses removed[18]. The results of this comparison are shown in Figure 10.27. We find
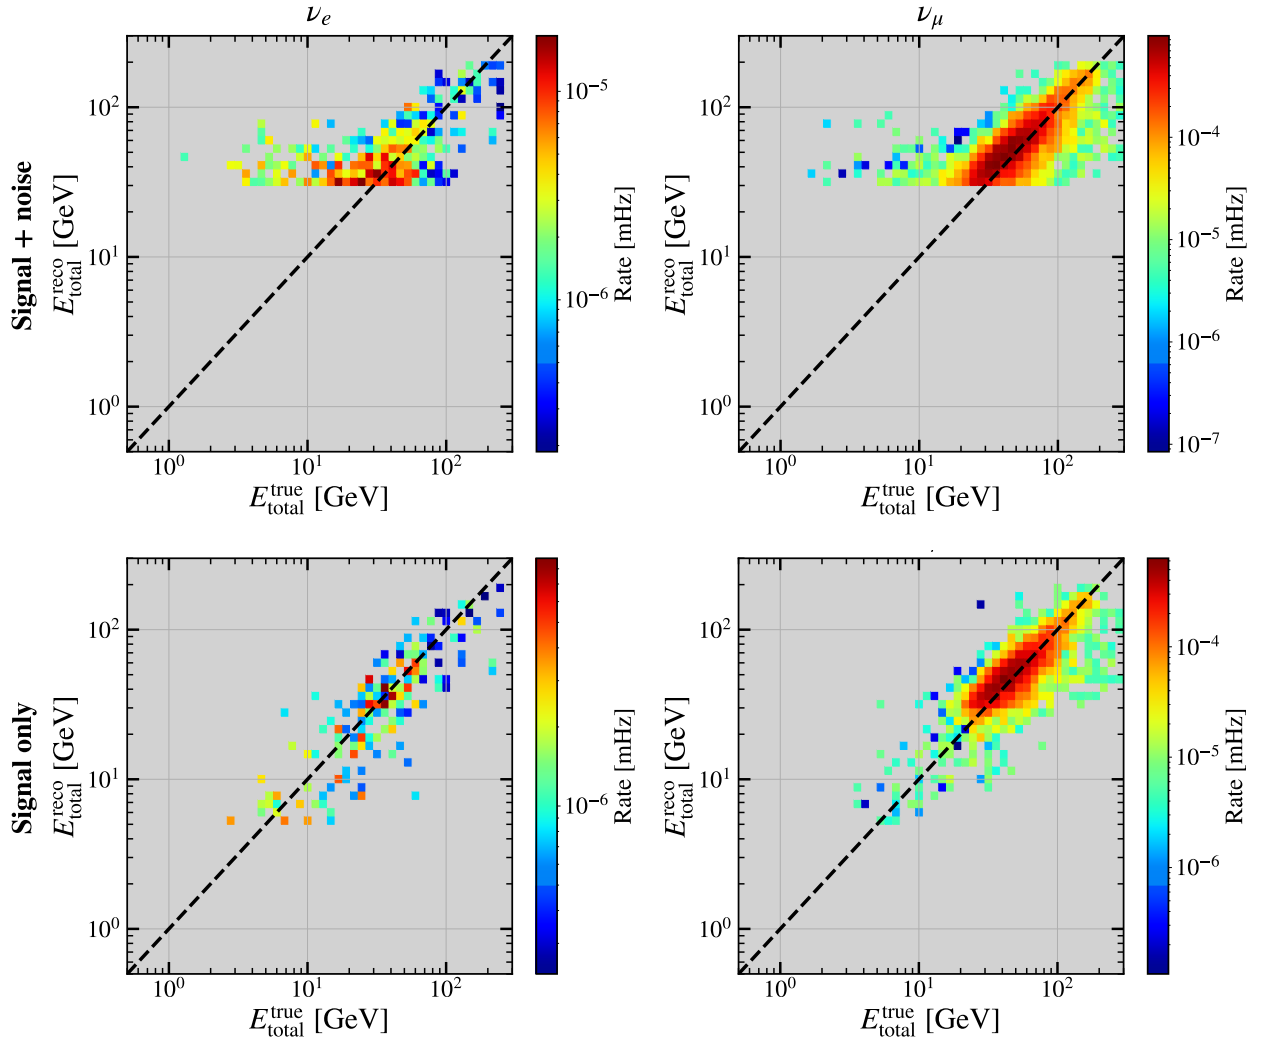

FIGURE 10.27 – Reconstructed ($E_{\text{total}}^{\text{reco}}$) *vs.* true ($E_{\text{total}}^{\text{true}}$) total energy in the oscNext L7 Monte Carlo events, as obtained when using the nominal SRTTWOfflinePulsesDC pulse series (top row) or the signal-only subset of the SRTTWOfflinePulsesDC pulse series (bottom row) for reconstruction. Only events with $E_{\text{track}}^{\text{reco}} > 30\,\text{GeV}$ as obtained from the nominal pulse series reconstruction are shown. Note that the SRTTWOfflinePulsesDC pulse series is already pre-cleaned with the SRTTW algorithm – however, it still contains a fraction of true noise pulses that lead to energy misreconstructions.

that, when the full SRTTWOfflinePulsesDC pulse series is used for reconstruction, the events with true total energies as low as 3 GeV get reconstructed to track and total energies as high as 30 GeV. This is seen from the top row of Figure 10.27 and is particularly peculiar for ν_e events, which never contain a true track. On the other hand, when the noise is removed from the same events and the RETRO reconstruction is re-run, the $E_{\text{total}}^{\text{reco}}$ ($E_{\text{total}}^{\text{true}}$) dependence follows a 1:1 line, and the "band" of events with low true energies and high reconstructed energies disappears. This presents a strong evidence that ***noise is the culprit of energy misreconstructions***.

While the use of the signal-only pulse maps mitigates the energy reconstruction bias, this cannot immediately be applied to resolve the reconstruction issues in data, where the "signalness" of a given pulse is unknown. Therefore, we must first understand the type and the properties of the noise pulses that could lead to a significantly higher rate

---

[18]The "signalness" truth flags were added to the SRTTWOfflinePulsesDC pulse map by J. Weldert, K. Leonard de Holton, and R. Ørsøe. These events were then processed to the final level, including the RETRO reconstruction at Level 6, by T. Stuttard.

of misreconstructions in data, and eventually remove those pulses without relying on truth information. A study in this direction by T. Stuttard reveals a strong mismodeling of *PMT afterpulses* in simulation, which we briefly summarize below.

Afterpulses are a form of PMT noise that results from ionization of the gas molecules in the PMT tube by the signal photoelectron. The ions drift back to the cathode and release several more photoelectrons, followed by a new wave of charge amplification with a ~500–1000 ns delay relative to the primary signal. As much as 45% of afterpulses in the DeepCore fiducial volume survive the SRTTW cleaning, compared to the ~30% survival rate of the radioactive decay noise in the PMT pressure vessel targeted by the algorithm. Given that afterpulses are correlated with the primary signal but are delayed by longer than the ATWD readout window (427 ns), they are typically a subset of HLC pulses digitized through the FADC circuitry. Indeed, ~ 1/3 of all HLC+FADC pulses that survive the SRTTW cleaning and are ultimately used for reconstructions are afterpulses. The charge and time distributions of these pulses, which are selected by imposing the pulse charge cut $q_{\text{pulse}} > 2.5\,\text{PE}$[19], are compared between data and Monte Carlo in Figure 10.28. A significant data/MC disagreement in both distributions indicates poor modeling of afterpulses in Monte Carlo and necessitates the development of a new pulse cleaning scheme targeting this particular noise component.
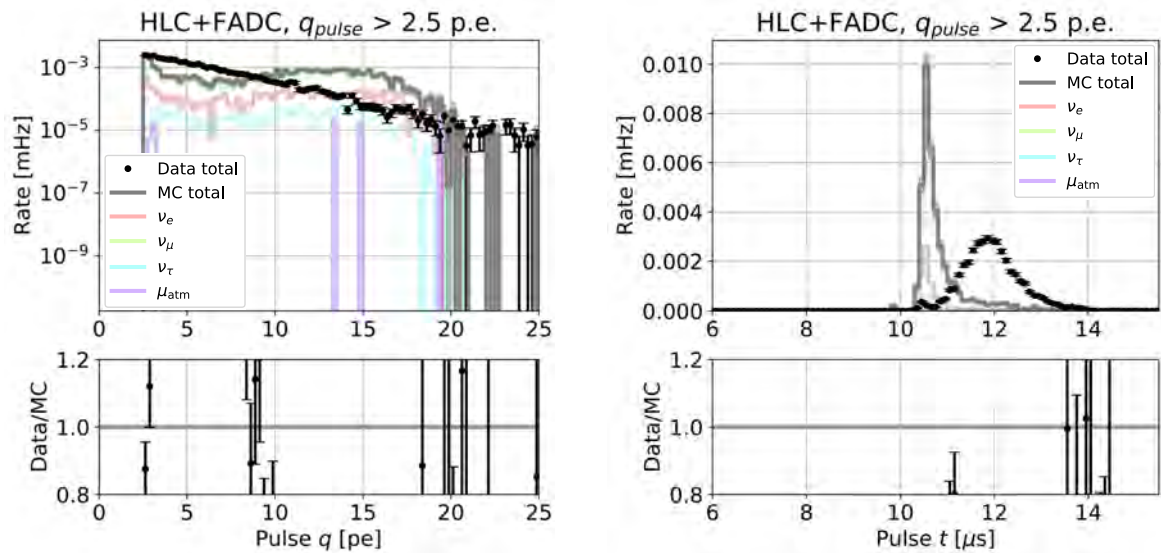
FIGURE 10.28 – Data/Monte Carlo agreement in the charge (left) and time (right) distributions of high-charge HLC+FADC pulses, which are dominated by afterpulses [284]. *Figure courtesy of T. Stuttard.*

To address the afterpulse mismodeling issue, the following noise cleaning scheme is proposed by T. Stuttard:

- Only pulses within the [9600 ns, 12 500 ns] time window are kept prior to the SRT cleaning (compared to the old time window of [8800 ns, 14 000 ns]);
- Within this time window, only the first pulse on each DOM is kept, thereby rejecting afterpulses and late pulses[20];
- The pulse time is rounded to the nearest nanosecond[21], and the pulse charge information is discarded to bypass any charge modeling issues;

---

[19]This requirement differentiates the afterpulses, which have high-charge tails as per [284], from the SPE charge distributions.
[20]Late pulses occur due to photoelectron backscattering from the first dynode and result in a secondary signal delayed by ~50–100 ns. These are commonly found in the HLC+ATWD channel and do not exhibit such a strong data/MC disagreement as in Figure 10.28.
[21]This avoids issues related to the presence of non-integer pulse times in data but never in Monte Carlo, which occurs due to re-extraction of compressed waveforms in individual events in data.

- The SRT cleaning settings are adjusted from $\{r = 150\,\text{m}, t = 1\,\text{μs}\}$ to $\{r = 150\,\text{m}, t = 500\,\text{ns}\}$.

This prescription removes all pulses in the HLC+FADC channel and nearly 60% of afterpulses in the HLC+ATWD channel. As the result, ~80% of all afterpulses are removed. This has a hugely positive impact on data/Monte Carlo agreement in both the track energy reconstruction quality metrics and the reconstructed track energy itself, as shown in Figures 10.29 and 10.30, respectively. Specifically, ***the new noise cleaning eliminates the previously identified data excesses*** in the regions of low numbers of hit DOMs per unit track energy, small fractions of track length within which 90% of all charge is accumulated, and high reconstructed track energy.
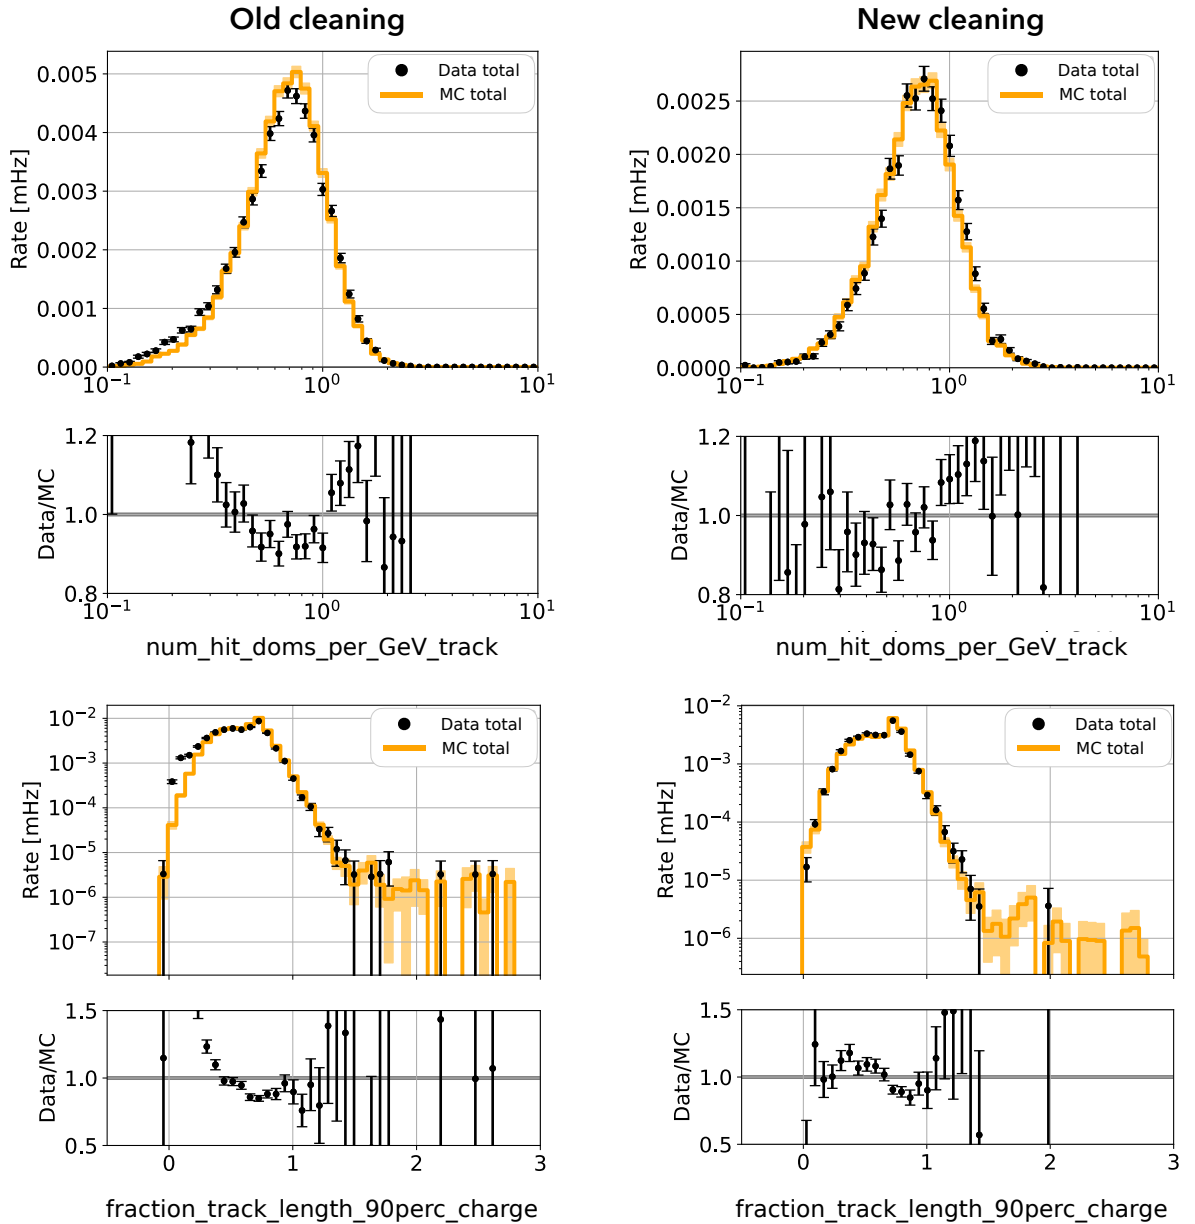
FIGURE 10.29 – Data/Monte Carlo agreement in the track energy reconstruction quality metrics derived in Section 10.3.1, compared between two choices of pulse maps used as input to the RETRO reconstruction [211]. In the left column, the SRTTWOfflinePulsesDC pulse map with the old noise cleaning settings is used. The right column shows the impact of the new pulse map with the noise cleaning settings devised by T. Stuttard (described in text). Only events with $40\,\text{GeV} < E_{\text{track}}^{\text{reco}} < 100\,\text{GeV}$ (using the old reconstruction) are selected here. *Figure courtesy of T. Stuttard (adapted).*
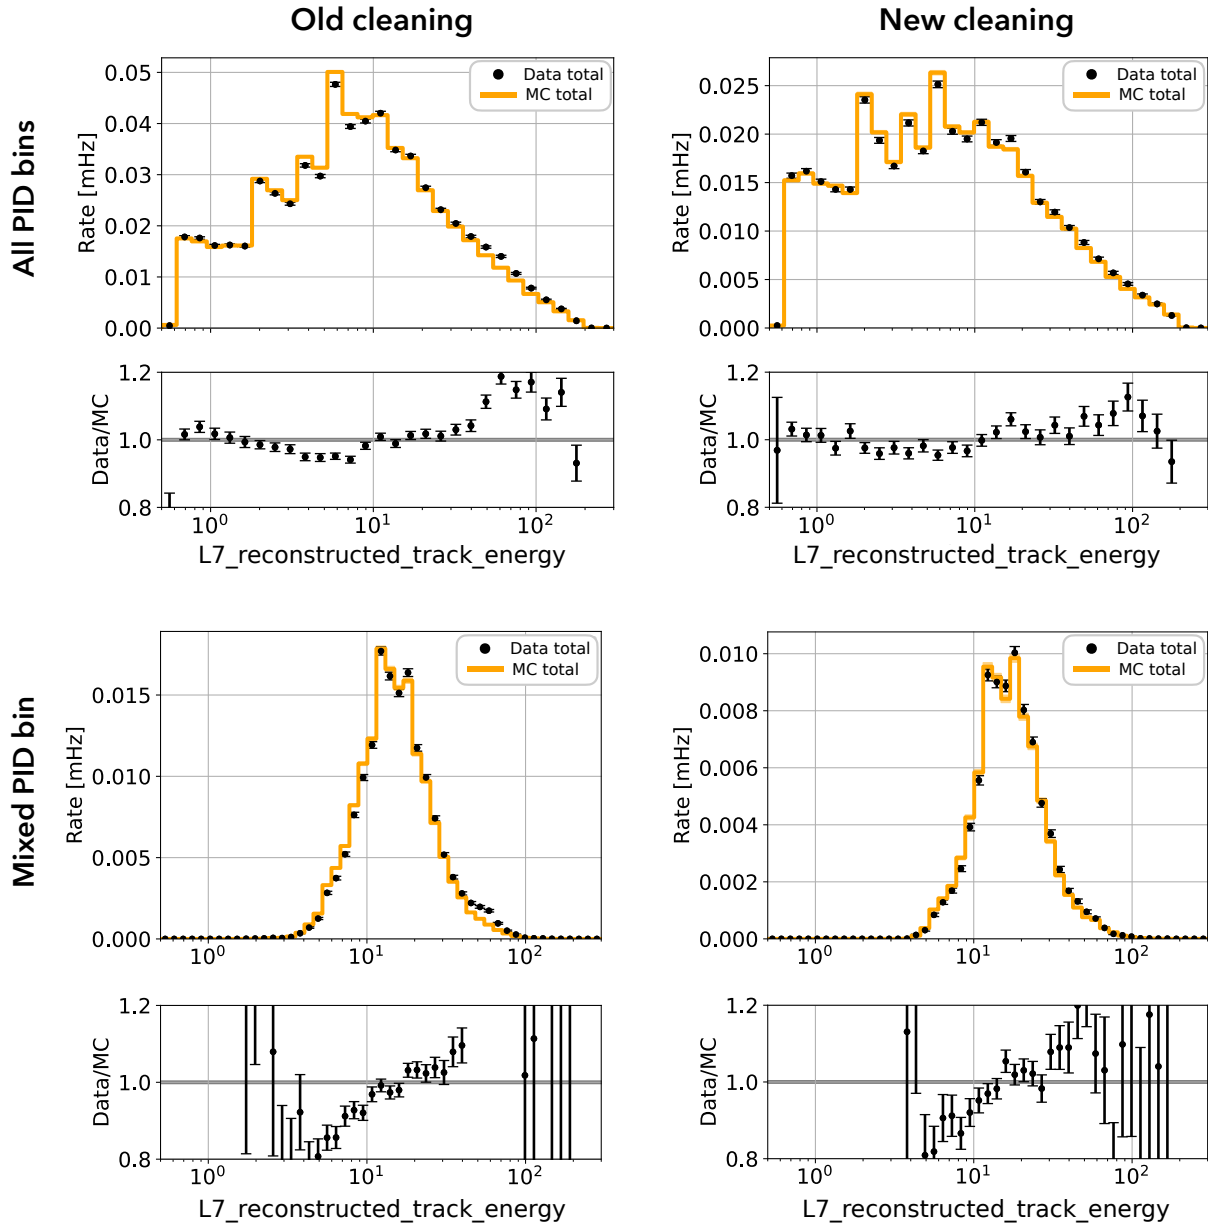
Figure 10.30 – Data/Monte Carlo agreement in the reconstructed track energy variable (E_{track}^{reco}, or "L7_reconstructed_track_energy", shown in GeV), compared between two choices of pulse maps used as input to the RETRO reconstruction [211]. For details, see the caption of Figure 10.29. *Figure courtesy of T. Stuttard (adapted).*

Although track energy reconstruction with the RETRO algorithm becomes much more robust when the new pulse map is used as input, re-running the RETRO reconstructions on all events in the oscNext Monte Carlo simulation sets is computationally infeasible on the timescale allocated for neutrino oscillation analyses with the oscNext-high statistics sample. Therefore, a decision is made to *replace the RETRO reconstruction with the DYNEDGE graph neural network* (GNN) [210], which was designed and developed by R. Ørsøe and trained by K. Leonard de Holton and T. Stuttard. The pulse map with the new noise cleaning is used for training the GNN to predict total neutrino energy, interaction vertex position, track length, and PID. The inference speed of a trained DYNEDGE network is $\mathcal{O}(10\,\text{ms})$ per event on a CPU, compared to ~40 s per event spent by the RETRO algorithm. This enables fast processing of all Monte Carlo sets,

including ~30 systematic sets, and 11 years of data (2012–2022). The preliminary pre-fit data/Monte Carlo agreement in the total reconstructed neutrino energy variable is shown in Figure 10.31.
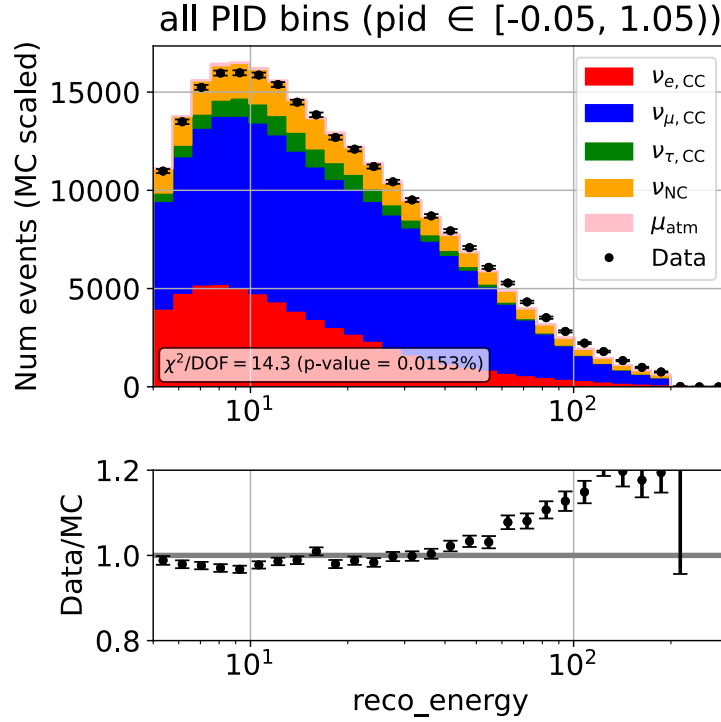
FIGURE 10.31 – Pre-fit data/Monte Carlo agreement in the total neutrino energy variable reconstructed with the DYNEDGE graph neural network [210] ("reco_energy", shown in GeV). The same L7 cuts as described in Section 9.1, complemented with a tighter starting containment cut ($\rho^{\text{reco}}_{36,\text{starting}} < 150\,\text{m}$) and an additional stopping containment cut ($\rho^{\text{reco}}_{36,\text{stopping}} < 150\,\text{m}$, $-500\,\text{m} < z^{\text{reco}}_{\text{stopping}} < -200\,\text{m}$), are applied in this figure. The muon light yield in the ν_μ-CC Monte Carlo shown here is simulated in GEANT4.

Although the development of the DYNEDGE-based oscNext sample is not finalized at the time of writing, in particular with regards to the exact choice of the event quality cuts, we can already see from Figure 10.31 that the main data/MC trend as a function of reconstructed energy appears largely similar to Figure 9.7. Specifically, the apparent excess of high-energy events in data, which grows as a function of reconstructed energy, remains present even after switching to the accurate muon light yield simulation, the refined noise cleaning, and the state-of-the-art machine learning-based event reconstruction. While this trend does not necessarily forecast data/Monte Carlo disagreement in the post-fit distributions or a bad goodness of fit, its origin remains unclear. At the present moment, we can conclude that the detailed investigations into the energy data/Monte Carlo disagreement led to multiple crucial fixes and improvements in the IceCube Monte Carlo simulation and processing, yielding a significantly more robust version of the oscNext sample.

10.4 UPDATED ν_τ NORMALIZATION SENSITIVITIES

10.4.1 *The impact of new reconstruction and muon simulation*

In Figure 10.32, we show how the sensitivities to N_{ν_τ} are affected by replacing the RETRO reconstruction of the analysis variables with that using DYNEDGE, as well as the shift from the PROPOSAL+PPC to the GEANT4-based simulation of muon light yield.

FIGURE 10.32 – Comparison between the tau neutrino normalization sensitivities projected for the old configuration of the oscNext sample (RETRO+PPC) and the new configuration (DYNEDGE+GEANT4; see text for details). The lower panel shows the 68% and the 90% confidence levels.
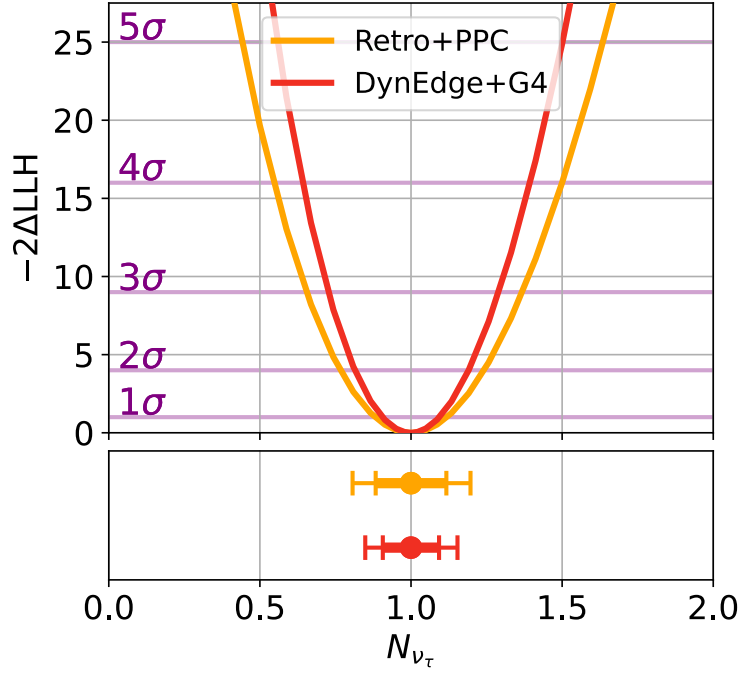
We find that the 68% (90%) sensitivity changes from 11.7% (19.5%) in the old configuration to 9.3% (15.3%) in the new configuration, indicating a small improvement. This outcome is a natural consequence of the superior DYNEDGE reconstruction resolutions [210] and the reduced muon neutrino rate at energies below $10\,\mathrm{GeV}$ (see Figure 10.10).

10.4.2   *The impact of the ν_τ cross section uncertainty*

Finally, we note that all of the N_{ν_τ} sensitivities presented so far relied on the assumption that the charged current tau neutrino cross section cross section is fully correlated with that of the lighter flavors. The validity of this assumption is justified in Chapter 11, where we calculate the uncertainty on the ν_τ / ν_μ CC DIS cross section ratio, and find it to be negligibly small for the purpose of the present analysis. A simple implementation of a separate absolute ν_τ cross section uncertainty was, however, considered in the IceCube Collaboration – namely, by devising a systematic parameter which interpolates between the different absolute ν_τ cross section models in [285]. In Figure 10.33, we show the impact of incorporating this parameter into the analysis for both the old configuration and the new configuration of the oscNext sample. We find that, had there been a theoretically motivated need to incorporate an additional ν_τ uncertainty (on top of the freedom provided by varying the cross sections of all three flavors in a correlated manner), the sensitivities to N_{ν_τ} would worsen in both variants of the sample. In particular, the sensitivity of the RETRO+PPC-based sample would drop from 11.7% (19.5%) to 18.2% (26.6%) at the 68% (90%) confidence level. For the DYNEDGE+GEANT4 configuration, the sensitivity would accordingly change from 9.3% (15.3%) to 15.2% (21.9%) at the same confidence levels[22].

Given that we show in Chapter 11 that any change to the tau neutrino cross section is

---

[22]Note that when the confidence intervals are mildly asymmetric, we are reporting the arithmetic average between the one-sided bounds.
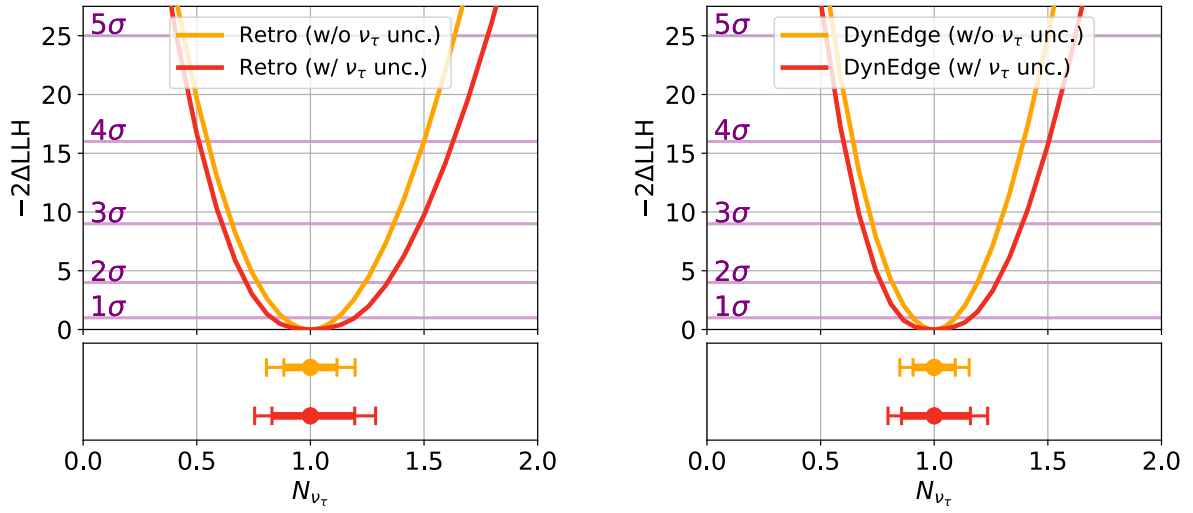
FIGURE 10.33 – Impact of the absolute tau neutrino cross section uncertainty parameter (explained in text) on the N_{ν_τ} sensitivities in the old (*left*) and the new (*right*) configurations of the oscNext sample. The lower panels in both subfigures show the 68% and the 90% confidence levels.

strongly correlated with that of the muon neutrino cross section[23], the inclusion of the additional ν_τ cross section uncertainty parameter is unmotivated, and the sensitivities reported in Figure 10.32 should be interpreted as accurate projections.

---

[23]Our projected ν_τ / ν_μ CC DIS cross section ratio uncertainty is $\sim 2\%$ (see Figure 11.5), while the parameter described in this section suggests up to 25% additional ν_τ cross section uncertainty at this energy.

VI

DIS CROSS SECTION SYSTEMATIC UNCERTAINTIES

RELATIVE SCALING OF ν_τ AND ν_μ DIS CROSS SECTIONS

N.B.: Sections 11.1 to 11.6 of this chapter contain the reproduction of the following paper:

Tetiana Kozynets, Thomas Stuttard, and David Jason Koskinen, *Deep Inelastic Scattering Cross Section Uncertainties in Tau Neutrino Appearance Measurements*, submitted to Phys. Rev. D. (2024).

At the time of writing, the paper is under review by the Physical Review D editorial board. The full pre-print of the paper is available online in the arXiv database [2].

11.1 INTRODUCTION

The uncertainty on the charged current deep inelastic scattering (CC DIS) neutrino cross section is an important systematic parameter in experimental analyses involving neutrino detection at energies of a few GeV and above. The scenario when a detector is sensitive to all three neutrino flavors, which applies to e.g. atmospheric neutrino oscillation experiments, is of particular interest, as it raises the question of correlations between the uncertainties on the cross sections of the different flavors. The flavor dependence of the DIS cross section appears in terms proportional to m_ℓ^2/E_ν^2 or $m_\ell^2/(E_\nu M_N)$, where E_ν is the incident neutrino energy, M_N is the mass of the target nucleon, and m_ℓ is the mass of the produced charged lepton [286–289]. These terms are typically neglected for ν_e and ν_μ cross sections due to the smallness of the electron and the muon masses relative to the neutrino energies in the DIS regime. However, the mass of the tau lepton ($m_\tau \simeq 1.7\,\mathrm{GeV}$) is comparable to the $\mathcal{O}(1\text{–}10\,\mathrm{GeV})$ energies relevant for atmospheric and accelerator neutrino oscillation measurements. Naively, this could

result in the conclusion that the ν_τ CC DIS cross section uncertainty should be decoupled from the ν_μ and ν_e uncertainties in an experimental analysis, i.e., by letting the energy-dependent ν_τ uncertainty vary independently from those of other flavors. The present study challenges this hypothesis by computing the ν_τ/ν_μ CC DIS cross section ratio at the leading order in perturbation theory and estimating its uncertainty. We achieve the latter by propagating the uncertainties on the parton distribution functions (PDFs) into the structure functions entering the DIS cross section, which we evaluate numerically. We additionally study the dependence of the cross section ratio on the threshold mass of the hadronic final state defining the transition between the resonant (RES) and the deep inelastic scattering. The outcomes of these tests are meant to guide neutrino experiments in defining the scope of the cross section systematic uncertainties necessary for neutrino oscillation analyses in the DIS regime. This concerns, in particular, the experiments such as IceCube-DeepCore [46, 180, 196, 197] and KM3NeT [290], whose ν_μ and ν_τ event selections are dominated by DIS interactions, as well as the upcoming IceCube-Upgrade [205, 212] and DUNE [291, 292], which cover the transition region between RES and DIS regimes (also referred to as "shallow inelastic scattering").

11.2 THE CHARGED CURRENT DIS CROSS SECTION

Neutrino-nucleon (νN) charged current deep inelastic scattering can be characterized through the kinematic variables x and y, representing the fraction of the nucleon momentum carried away by the struck quark (also called the Bjorken scaling variable) and the inelasticity of the interaction, respectively. For a neutrino ν_ℓ (antineutrino $\bar{\nu}_\ell$), the double-differential cross section with respect to these two variables reads [286–289]:

$$\frac{\mathrm{d}^2\sigma^{\nu_\ell N(\bar{\nu}_\ell N)}}{\mathrm{d}x\mathrm{d}y} = \frac{G_F^2 M_N E_\nu}{\pi(1+Q^2/M_W^2)^2}\left\{\left(y^2 x + \frac{m_\ell^2 y}{2E_\nu M_N}\right)F_1 + \left[\left(1-\frac{m_\ell^2}{4E_\nu^2}\right)-\left(1+\frac{M_N x}{2E_\nu}\right)y\right]F_2\right.$$
$$\left.\pm\left[xy\left(1-\frac{y}{2}\right)-\frac{m_\ell^2 y}{4E_\nu M_N}\right]F_3 + \frac{m_\ell^2(m_\ell^2+Q^2)}{4E_\nu^2 M_N^2 x}F_4 - \frac{m_\ell^2}{E_\nu M_N}F_5\right\}. \qquad (11.1)$$

Here, $Q^2 \equiv 2M_N E_\nu xy$ is the momentum transfer squared, G_F is the Fermi constant, and $F_i \equiv F_i(x, Q^2)$, $i = 1..5$, are the structure functions of the nucleon. The "+" sign in the coefficient multiplying F_3 corresponds to neutrinos, and the "−" sign – to antineutrinos. While the terms proportional to either m_ℓ^2/E_ν^2 or $m_\ell^2/(E_\nu M_N)$ enter the coefficients multiplying all of the 5 structure functions, the contributions due to F_4 and F_5 are effectively negligible when $m_\ell^2 \ll (E_\nu M_N)$. For this reason, they are typically not included in the calculation of the ν_e or ν_μ DIS cross sections (see e.g. [293]) but become relevant in the case of ν_τ [286, 288, 289]. We include these structure functions explicitly when computing both ν_μ and ν_τ cross sections, as described in Section 11.3.

All of the structure functions F_i have an underlying dependence on the experimentally determined PDFs, whose uncertainties ultimately propagate to the DIS cross section in Equation (11.1). At large Q^2 ($Q^2 \gtrsim 1\,\mathrm{GeV}^2$, where the strong coupling constant α_s is smaller than unity [34, 294]), the path from the PDFs to the structure functions lies though perturbation theory. In this context, the leading-order (LO), next-to-leading order (NLO), and next-to-next-to-leading order (NNLO) perturbative QCD calculations of the absolute ν_τ cross sections have been performed in literature [153, 286, 288, 289, 295]. The contribution of each scattering diagram of order n is weighted by $\alpha_s^n(Q^2)$,

which is monotonically decreasing with growing Q^2. This implies that the difference between the results of the leading-order and the higher-order calculations is the largest at the lowest Q^2 and, consequently, E_ν [1]. Furthermore, the low-Q^2 region is susceptible to non-perturbative effects such as the target mass corrections (concerning the masses of the target nucleon and the struck quark) and the dynamical higher twist effects (concerning the interactions the of struck quark with the surrounding quarks) [289, 295]. In this study, we evaluate the structure functions and the inclusive DIS cross sections at the leading order to capture the dominant impact of the PDF uncertainty propagation. Additionally, we take into account the corrections for the nucleon mass ($M_N \approx 0.938\,\text{GeV}$) and the mass of the final-state charm quark ($m_c \approx 1.27\,\text{GeV}$), as these effects introduce a small deviation of the F_4 structure function from 0 at leading order [286]. Finally, we perform all of our calculations for an isoscalar target, assuming that the nuclear medium effects apply equally to ν_μ and ν_τ scattering and do not influence the cross section ratio [2].

11.3 STRUCTURE FUNCTIONS AT LEADING ORDER

Throughout this study, we work with the conventional mass basis for quarks, such that the u, c, t mass eigenstates coincide with the weak interaction eigenstates, while the d, s, b mass eigenstates are related to the weak eigenstates d', s', and b' through the CKM matrix [34]:

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix}. \tag{11.2}$$

Additionally, we operate in the fixed-flavor number scheme with $n_f = 4$ quark flavors, i.e., consider only the first two generations [3]. From the charge conservation constraints, a neutrino ν_l corresponding to the lepton flavor l can engage in CC interactions with the d', s', $\bar{u}$, or $\bar{c}$ quarks inside a nucleon. Such interactions are mediated by the W^+ boson and result in the production of the u, c, $\bar{d}'$, and $\bar{s}'$ final-state quarks, accordingly. Analogously, a CC scattering of an antineutrino $\bar{\nu}_l$ can result in the $\bar{d}' \to \bar{u}$, $\bar{s}' \to \bar{c}$, $u \to d$, and $c \to s'$ conversions at the W^- vertex.

As the PDFs measured by the experimental collaborations probe the proton structure, we will use $\boldsymbol{q}(x, Q^2)$ to denote the density of the quark q inside a proton. These PDFs can be used to compute the structure functions F_i ($i = 1..5$) for both neutrino-proton and neutrino-neutron scattering by applying the isospin symmetry argument. At leading order and ignoring the target nucleon and the charm mass corrections (i.e., assuming $M_N = 0$, $m_c = 0$), the F_2 and F_3 structure functions for the νp scattering are [289, 293]

$$F_2^{\nu p}(x, Q^2)\bigg|_{\substack{M_N=0 \\ m_c=0}} = 2x[\boldsymbol{d}' + \boldsymbol{s}' + \bar{\boldsymbol{u}} + \bar{\boldsymbol{c}}](x, Q^2); \tag{11.3a}$$

$$F_3^{\nu p}(x, Q^2)\bigg|_{\substack{M_N=0 \\ m_c=0}} = 2[\boldsymbol{d}' + \boldsymbol{s}' - \bar{\boldsymbol{u}} - \bar{\boldsymbol{c}}](x, Q^2), \tag{11.3b}$$

---

[1] At $E_\nu = 10\,\text{GeV}$, the difference between the ν_τ CC DIS cross sections evaluated at LO vs NLO is ~10% [286].

[2] The effects such the W-boson interactions with the mesonic cloud and the nuclear shadowing were shown to introduce $\mathcal{O}$(tens of %) shifts to the ^{56}Fe structure functions in regions of intermediate x (0.1–0.5) and Q^2 (2–20 GeV^2) [295]. This is comparable to the difference in the structure functions due to the calculation order, e.g. LO vs. NLO (NNLO) [286, 293]. The impact of nuclear effects increases with the mass number and is therefore expected to be smaller for nuclei such as ^{16}O (part of the IceCube and KM3NeT water target) or ^{40}Ar (DUNE target).

[3] The density of the sea charm quark relative to that of the sea strange quark is given in Figure C.1. We find that the ultimate impact of including the sea charm into the ν_τ / ν_μ cross section ratio calculations is near-negligible (see Figure C.2).

where $\boldsymbol{d}' = |V_{ud}|^2 \boldsymbol{d} + |V_{us}|^2 \mathfrak{s}$ and $\mathfrak{s}' = |V_{cd}|^2 \boldsymbol{d} + |V_{cs}|^2 \mathfrak{s}$. For the $\bar{\nu}p$ scattering, the PDFs for all quarks in Equation (11.3) are replaced by those of their weak doublet partners, e.g., $\boldsymbol{d}' \to \boldsymbol{u}$, while for the νn scattering, the replacements are made only for the u and d quarks (see [289] for complete expressions). Furthermore, in the considered massless approximation at LO, $F_4(x, Q^2) = 0$, while F_1 and F_5 follow from the Callan-Gross [296] and the Albright-Jarlskog [287] relations:

$$\left[F_1(x, Q^2) - \frac{F_2(x, Q^2)}{2x} \right]\Bigg|_{\substack{M_N = 0 \\ m_c = 0}} = 0; \tag{11.4a}$$

$$\left[F_5(x, Q^2) - \frac{F_2(x, Q^2)}{2x} \right]\Bigg|_{\substack{M_N = 0 \\ m_c = 0}} = 0. \tag{11.4b}$$

To account for the non-zero nucleon mass M_N, the Bjorken x variable in the quark PDFs has to be replaced by the Nachtmann variable η, defined as [286, 297]

$$\frac{1}{\eta} = \frac{1}{2x} + \sqrt{\frac{1}{4x^2} + \frac{M_N^2}{Q^2}}. \tag{11.5}$$

The Nachtmann variable is further corrected by the final-state quark mass m_q:

$$\eta \to \bar{\eta}_q \equiv \frac{\eta}{\lambda_q} = \eta \left(\frac{Q^2}{Q^2 + m_q} \right)^{-1}, \tag{11.6}$$

where we assume only the charm quark to be massive ($m_{q, q \neq c} = 0$). Further, we define

$$\rho^2 = 1 + \left(\frac{2M_N x}{Q} \right)^2, \tag{11.7}$$

following [286, 289]. Then, the final leading-order expressions for the structure functions corrected by the masses of the nucleon and the final-state charm quark become[45] [286]:

$$F_1^{\nu p}(x, Q^2) = \boldsymbol{d}'(\eta, Q^2) + \mathfrak{s}'(\bar{\eta}_c, Q^2) + \bar{\boldsymbol{u}}(\eta, Q^2) + \bar{\boldsymbol{c}}(\eta, Q^2); \tag{11.8a}$$

$$F_2^{\nu p}(x, Q^2) = \frac{2x}{\rho^2} \left[\boldsymbol{d}'(\eta, Q^2) + \frac{1}{\lambda_c} \mathfrak{s}'(\bar{\eta}_c, Q^2) + \bar{\boldsymbol{u}}(\eta, Q^2) + \bar{\boldsymbol{c}}(\eta, Q^2) \right]; \tag{11.8b}$$

$$F_3^{\nu p}(x, Q^2) = \frac{2}{\rho} \left[\boldsymbol{d}'(\eta, Q^2) + \mathfrak{s}'(\bar{\eta}_c, Q^2) - \bar{\boldsymbol{u}}(\eta, Q^2) - \bar{\boldsymbol{c}}(\eta, Q^2) \right]; \tag{11.8c}$$

$$F_4^{\nu p}(x, Q^2) = \left[\frac{1 - \rho^2}{2\rho^2} \right] \cdot \left[\boldsymbol{d}'(\eta, Q^2) + \bar{\boldsymbol{u}}(\eta, Q^2) + \bar{\boldsymbol{c}}(\eta, Q^2) \right] + \left[\frac{(1-\rho)^2}{2\lambda_c \rho^2} + \frac{1-\rho}{\rho} \right] \mathfrak{s}'(\bar{\eta}_c, Q^2); \tag{11.8d}$$

$$F_5^{\nu p}(x, Q^2) = \frac{1}{\rho^2} \left[\boldsymbol{d}'(\eta, Q^2) + \bar{\boldsymbol{u}}(\eta, Q^2) + \bar{\boldsymbol{c}}(\eta, Q^2) \right] + \left[\frac{1}{\rho} - \frac{\rho - 1}{\lambda_c \rho^2} \right] \mathfrak{s}'(\bar{\eta}_c, Q^2). \tag{11.8e}$$

---

[4] The expressions in Equation (11.8) are given in a collinear approximation [298], i.e., where the momenta of the struck parton and the parent nucleon are aligned in the infinite momentum frame. This approximation can be relaxed to include the parton transverse momentum [299], which leads to the full target mass corrections as in [300, 301]. While these expressions would lead to more accurate absolute cross sections, our goal is to estimate the magnitude of the PDF-driven uncertainty on the ν_τ/ν_μ cross section ratio, and we deem the collinear approximation plausible for this purpose.

[5] The ultimate impact of the M_N and m_c corrections on the ν_τ/ν_μ cross section ratio is shown in Figure C.2.

We note that with $M_N \to 0$, Equation (11.7) gives $\rho = 1$, which recovers $F_4 = 0$ and the relations from Equation (11.4). The isospin symmetry arguments can again be applied to Equation (11.8) to derive the neutron structure functions $F_i^{\nu n}$. From here, the structure functions for an isoscalar target A are obtained as

$$F_i^{\nu A} = \frac{F_i^{\nu p} + F_i^{\nu n}}{2}. \tag{11.9}$$

In Figure 11.1, we present the $F_i^{\nu p}(x)$ structure functions evaluated at fixed Q^2, while Figure 11.2 shows $F_i^{\nu p}(Q^2)$ at fixed x. We are using the NNPDF4.0 PDF set [302] fitted at NNLO and applied at LO according to Equation (11.8), as done in the benchmark tests of [293]. The PDFs are accessed through the LHAPDF library [303] interfaced with the PDFFlow Python module [304]. As the NNPDF4.0 grids are provided and valid down to the minimum momentum transfer value of $Q_{\min} = 1.65\,\text{GeV}$, we use the Martin-Stirling-Thorne-Watt (MSTW) scheme[6] built into PDFFlow for extrapolation to the low-Q^2 region [305]. In addition to the central values of the PDFs, the NNPDF set includes 1000 PDF "replicas" obtained by fitting 1000 individual neural networks to the Monte Carlo replicas of the original data [306, 307]. We use these replicas to construct the 68% (1σ) uncertainty contours for the structure functions, which are shown in Figures 11.1 and 11.2 alongside their central values. The distribution of the structure function replicas ultimately yields the corresponding distribution of the total DIS cross sections for ν_μ and ν_τ flavors and allows us to derive the PDF-driven uncertainty on their ratio, as described in Section 11.4.
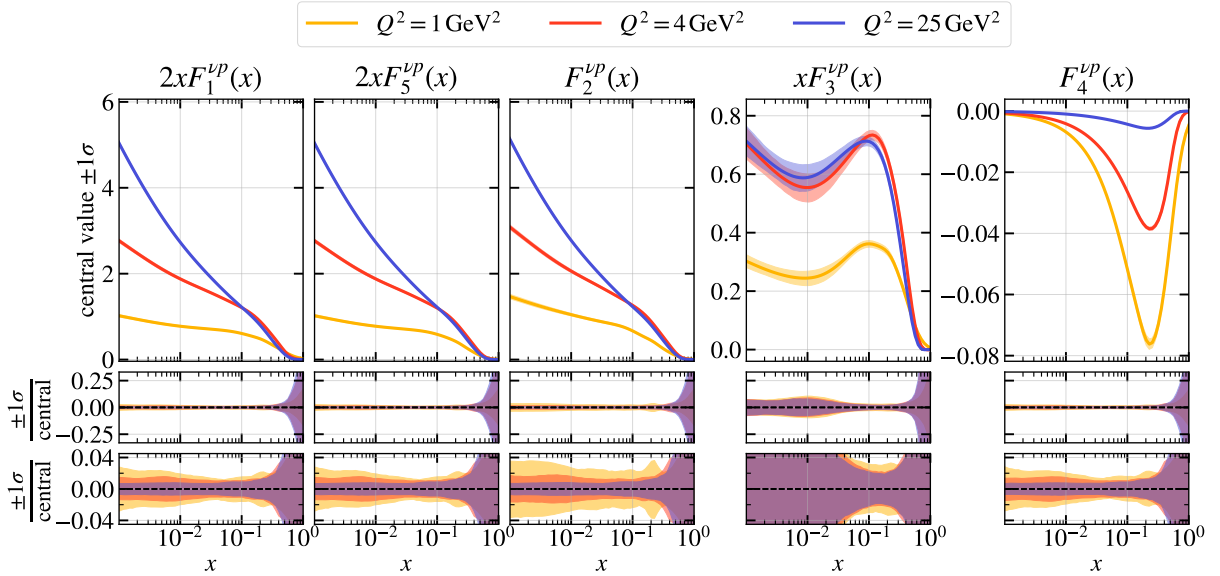
FIGURE 11.1 – Structure functions $F_i(x)$ evaluated at fixed Q^2 at LO and using the NNPDF4.0 PDF grids [302]. In the upper panels, the shaded bands represent the 68% (1σ) confidence level, while the middle and the bottom panels show the ratio of the 1σ band widths to the central value. At $Q^2 < Q_{\min}^2 = (1.65\,\text{GeV})^2$, the MSTW extrapolation scheme [303, 305] is used.

---

[6]This is the default extrapolation scheme for all PDF grids in the recent versions of the LHAPDF library (see Sec. 3.4.2 in [303]).
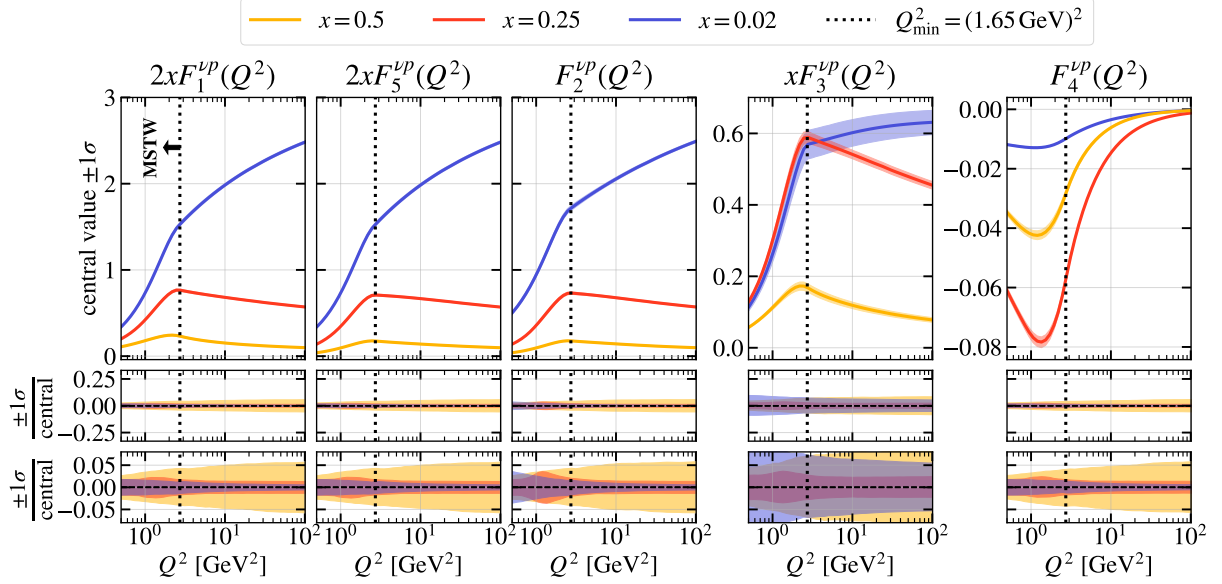
FIGURE 11.2 – Structure functions $F_i(Q^2)$ evaluated at fixed x at LO (for details, see Figure 11.1). The vertical dotted black line denotes the minimum Q^2 value of the NNPDF4.0 PDF grids, below which the MSTW extrapolation scheme is used [303, 305].

11.4 INCLUSIVE DIS CROSS SECTIONS AND ν_τ/ν_μ CROSS SECTION RATIO AT LO

The double-differential cross section $\frac{d^2\sigma}{dxdy}$ from Equation (11.1) can alternatively be expressed as a function of y and Q^2 through the change of variables:

$$\frac{d^2\sigma}{dydQ^2} = \frac{d^2\sigma}{dxdy}\frac{\partial x}{\partial Q^2} = \frac{1}{2M_N E_\nu y}\frac{d^2\sigma}{dxdy}. \tag{11.10}$$

Then, the single-differential cross section with respect to inelasticity can be found by integrating over Q^2:

$$\frac{d\sigma}{dy} = \int_{Q^2_{\min}}^{Q^2_{\max}} \frac{d^2\sigma}{dydQ^2} dQ^2. \tag{11.11}$$

We reproduce the Q^2 integration limits from [288] for completeness:

$$Q^2_{\min} = 2E_\nu^2(1-\epsilon)(1-y) - m_\ell^2; \tag{11.12a}$$
$$Q^2_{\max} = 2M_N E y + M_N^2 - W^2_{\min}, \tag{11.12b}$$

where $\epsilon = \sqrt{1 - \frac{m_\ell^2}{((1-y)E_\nu)^2}}$ and $W_{\min}$ is the minimum invariant mass of the hadronic final state required to classify an interaction as deep inelastic as opposed to resonant scattering. We use $W_{\min} = 1.4\,\text{GeV}$ throughout this study and further test the impact of the $W_{\min}$ choice in Section 11.5. In Figure 11.3, we show the result of the $\frac{d\sigma}{dy}$ calculation through Equation (11.11), including the 68% confidence limit (1σ) uncertainties obtained from the 1000 NNPDF replicas. Figure 11.4 further gives the total cross sections obtained by integrating $\frac{d\sigma}{dy}$ over the allowed inelasticity range[7].

We observe that the relative cross section uncertainty is larger at higher inelasticities and lower energies, with both of these effects being more prominent for ν_τ than ν_μ.

---

[7]These results can be compared to those from [286, 288] obtained at NLO.
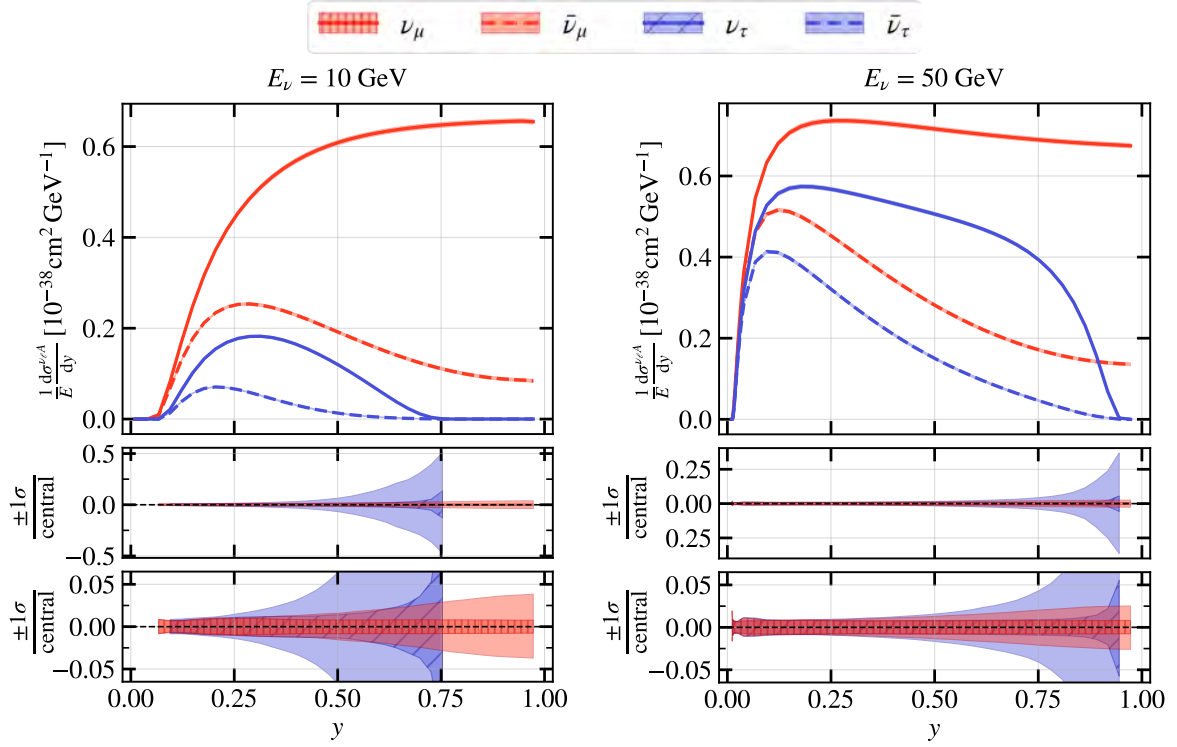
FIGURE 11.3 – Differential CC DIS cross sections of $\nu_\mu, \bar{\nu}_\mu$ (red curves) and $\nu_\tau, \bar{\nu}_\tau$ (blue curves) with an isoscalar target A, obtained at leading order using the structure functions from Equation (11.8). The details on the PDF set and the 1σ uncertainty range (shaded bands) derivation are given in Section 11.3.
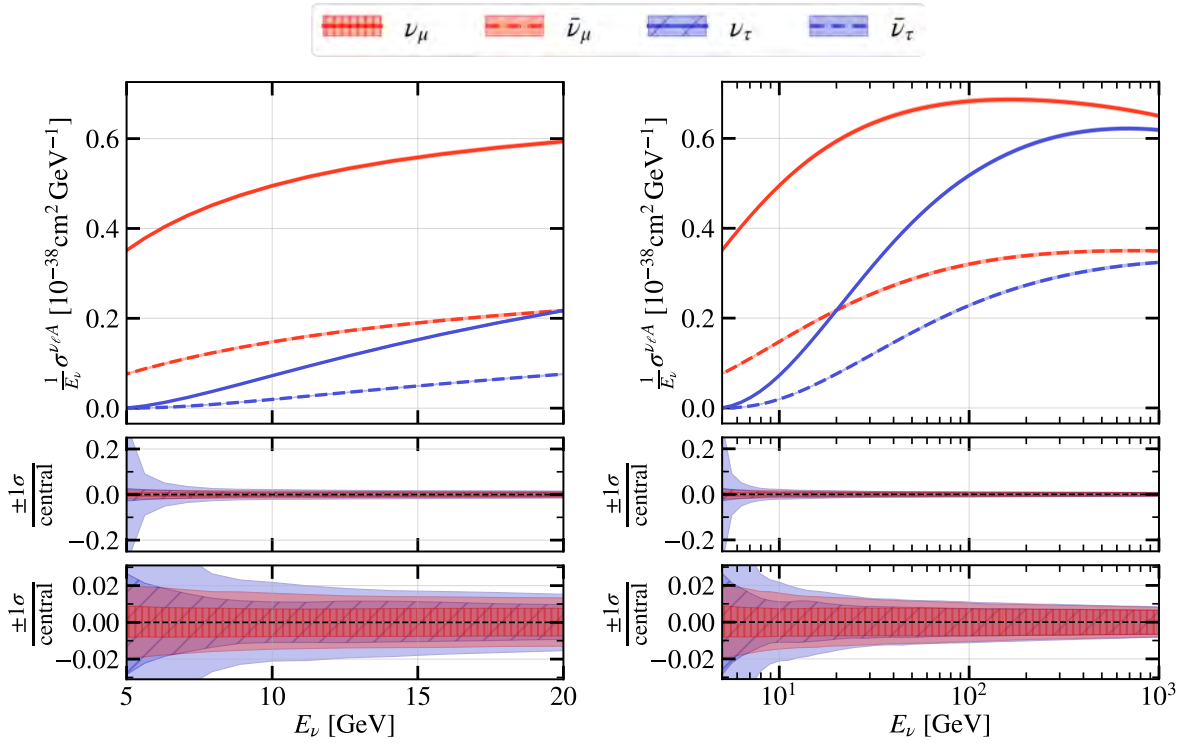
FIGURE 11.4 – Total CC DIS cross sections of $\nu_\mu, \bar{\nu}_\mu$ (red curves) and $\nu_\tau, \bar{\nu}_\tau$ (blue curves) with an isoscalar target A, obtained at leading order by integrating Equation (11.11) over the inelasticity range and using the structure functions from Equation (11.8). The details on the PDF set and the 1σ uncertainty range (shaded bands) derivation are given in Section 11.3. The left and the right panel respectively show the low energy range (5–20 GeV) and the full energy range (5–1000 GeV).

This could be explained by the larger relative contributions of the F_4 and F_5 structure functions to the DIS cross section in case of the ν_τ scattering, with their coefficients being proportional to y and inversely proportional to E_ν. We also find that the uncertainty on both single-differential and total cross sections is larger for antineutrinos than for neutrinos, which can be attributed to the PDF uncertainty cancellation being dependent on the sign with which F_3 structure function enters Equation (11.1).

Having the individual replicas of the total $\nu_\tau A$ and $\nu_\mu A$ cross sections at each probed E_ν, we build the distribution of the cross section ratios $r_{\tau\mu} \equiv \sigma(\nu_\tau A)/\sigma(\nu_\mu A)$ as a function of E_ν and construct the corresponding 68% confidence levels contours. The results are given in Figure 11.5.

FIGURE 11.5 – Ratio of the ν_τ CC DIS cross section with an isoscalar target A to that of ν_μ, shown as a function of neutrino energy E_ν. The purple curves correspond to neutrinos,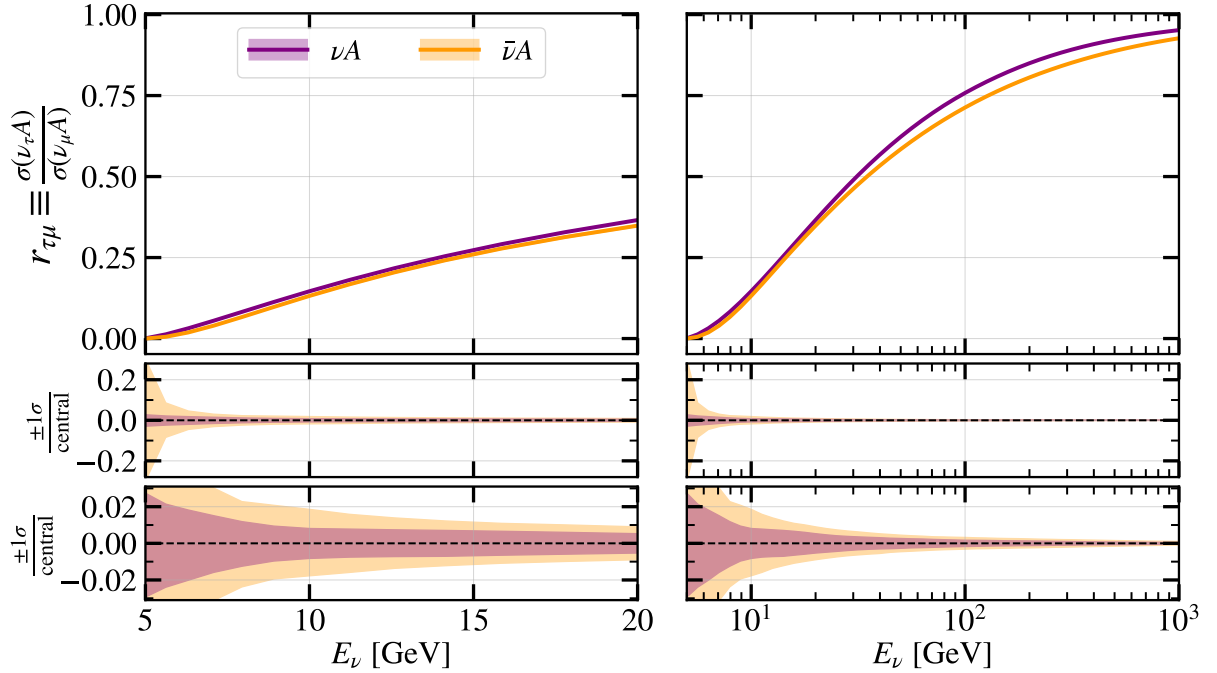 and the orange ones – to antineutrinos. The middle and the bottom panels show the relative 68% (1σ) uncertainty on the cross section ratios as derived from the NNPDF4.0 Monte Carlo replicas.

We find that the PDF-driven uncertainty on the $\sigma(\nu_\tau A)/\sigma(\nu_\mu A)$ cross section ratio is $\lesssim 3\%$ ($\lesssim 30\%$) at 5 GeV and $\lesssim 1\%$ ($\lesssim 2\%$) at 10 GeV for neutrinos (antineutrinos). The uncertainty monotonically decreases as a function of neutrino energy and reaches sub-percent level for both neutrinos and antineutrinos above 20 GeV. Whether the PDF-driven uncertainty on the ν_τ/ν_μ cross section ratio is significant enough to be included in an experimental analysis is therefore dependent on the energy reach of the experiment and the sensitivity to the $\bar\nu_\tau$ events. In the specific case of atmospheric neutrinos, the $\bar\nu_\tau$ flux (resulting mainly from $\bar\nu_\mu \to \bar\nu_\tau$ oscillations in the energy range from a few GeV to a few hundreds of GeV) is subdominant compared to the ν_τ flux. Additionally, the $\bar\nu_\tau$ CC DIS cross section is $\sim$2–10 times smaller at these energies than that of $\bar\nu_\tau$. As the result, one could expect $\sim$4 times fewer tau antineutrinos than neutrinos detected in an experiment such as IceCube in the 5–1000 GeV energy range[8]. The expected $\nu_\tau/\bar\nu_\tau$ event ratio further increases with decreasing energy. Therefore, given that only the antineutrinos with $E_\nu \lesssim 7$ GeV have a non-negligible tau-to-muon CC DIS cross section

---

[8]This estimate is derived from the ratio of the expected ν_τ and $\bar\nu_\tau$ {flux × cross section} products integrated over the specified energy range. The atmospheric neutrino fluxes for ν_μ and $\bar\nu_\mu$ are calculated via MCEq [43, 44] with the DPMJet-III 19.1 hadronic interaction model and oscillated into ν_τ and $\bar\nu_\tau$ using the Neurthino oscillation probability code [5]. The energy-dependent detector acceptance is not taken into account.

ratio uncertainty (>5%), and that their contribution to the overall $(\nu_\tau + \bar{\nu}_\tau)$-CC event rate is small, our recommendation is for the corresponding systematic uncertainty to *not* be included in the atmospheric neutrino oscillation analyses. However, the explicit consideration of the PDF-driven uncertainty on the $\bar{\nu}_\tau/\bar{\nu}_\mu$ CC DIS cross section ratio is generally advisable for beam experiments capable of operating in an antineutrino mode and detecting large samples of $\bar{\nu}_\tau$ CC events at energies below 10 GeV. While this could in principle apply to DUNE, only ~40 $\bar{\nu}_\tau$ CC events are projected for 3.5 years of exposure in the antineutrino mode, including DIS, RES, and QE interactions [292]. This implies that the initial $\bar{\nu}_\tau$ DIS samples from DUNE will be statistics-limited and likely insensitive to the PDF-driven ~2–30% uncertainties on the $\bar{\nu}_\tau/\bar{\nu}_\mu$ DIS cross section ratio.

11.5 IMPACT OF THE INVARIANT HADRONIC MASS THRESHOLD

When evaluating the differential cross section $\frac{\mathrm{d}\sigma}{\mathrm{d}y}$ in Equation (11.11), we imposed a cut on the invariant mass of the hadronic final state, requiring it to be larger than $W_{\min} = 1.4$ GeV for the interaction to be classified as deep inelastic scattering. However, this setting of $W_{\min}$ is not universal and depends on the analyzer's choice of the resonances to be included in the deep inelastic scattering region [153]. In the GENIE event generator, the default setting is $W_{\min} = 1.7$ GeV[9][158]. It is therefore natural to test what impact the choice of $W_{\min}$ makes on the $r_{\tau\mu}$ cross section ratio and compare it to the PDF-driven uncertainty.

In Figure 11.6, we show the relative impact of the PDF variation within its 1–3σ uncertainty limits as well as the impact of changing $W_{\min}$ from 1.4 GeV to 1.7 GeV. We find that at $E_\nu \lesssim 30$ GeV, the difference between the tau-to-muon CC DIS cross section ratios obtained with the two different choices of $W_{\min}$ is not contained in the 3σ PDF uncertainty bounds, and grows to nearly 100% at 5 GeV for both ν and $\bar{\nu}$. While these differences appear to be significant at first, one must remember that drop in the DIS cross section due to a higher $W_{\min}$ threshold is accompanied by the increase in the RES cross section. In the GENIE framework, this is handled by the RES-DIS joining scheme, the details of which can be found in [158]. In Figure 11.7, we zoom into the 5–30 GeV energy region and show the impact on the DIS, and DIS + RES ν_τ/ν_μ cross section ratios due to the shift in $W_{\min}$, as evaluated within GENIE. As expected, the impact on the ratios of the DIS+RES ν_τ/ν_μ cross sections is much smaller (reaching at most 6%) than the impact on the DIS cross section ratio alone (reaching 10–30% at $E_\nu \lesssim 10$ GeV). This implies that if a systematic uncertainty on the $r_{\tau\mu}$ ratio due to the arbitrariness in the choice of $W_{\min}$ is to be included in an analysis, it must be implemented for both DIS and RES channels to ensure a physical transition between the two.

Furthermore, the cross section in the RES-DIS transition region is constrained by inclusive and exclusive cross section data from the bubble chamber experiments such as ANL 12FT [171, 308], BNL 7FT [163, 309–311], BEBC [165, 312–318], and FNAL 15FT [319–321]. The $W_{\min}$ setting therefore cannot be varied on its own without the corresponding changes in the other parameters of the RES and the DIS models. In particular, $W_{\min}$ is anti-correlated with the $R^{CC2\pi}$ coefficients modulating the contribution of the two-pion final states, as well as the axial mass $M_{A,\mathrm{RES}}$ entering the form factors of the resonant scattering cross sections [157, 322]. The most accurate approach would therefore be to keep $W_{\min}$ fixed, tune the rest of the RES and the DIS parameters to the cross section data, and determine their covariance matrix as done e.g. in the tuning

---

[9]Note that in GENIE, the equivalent variable is called W_{cut}, while $W_{\min} \equiv M_N + m_\pi$, with m_π being the pion mass, corresponds to the lower end of the RES energy range.
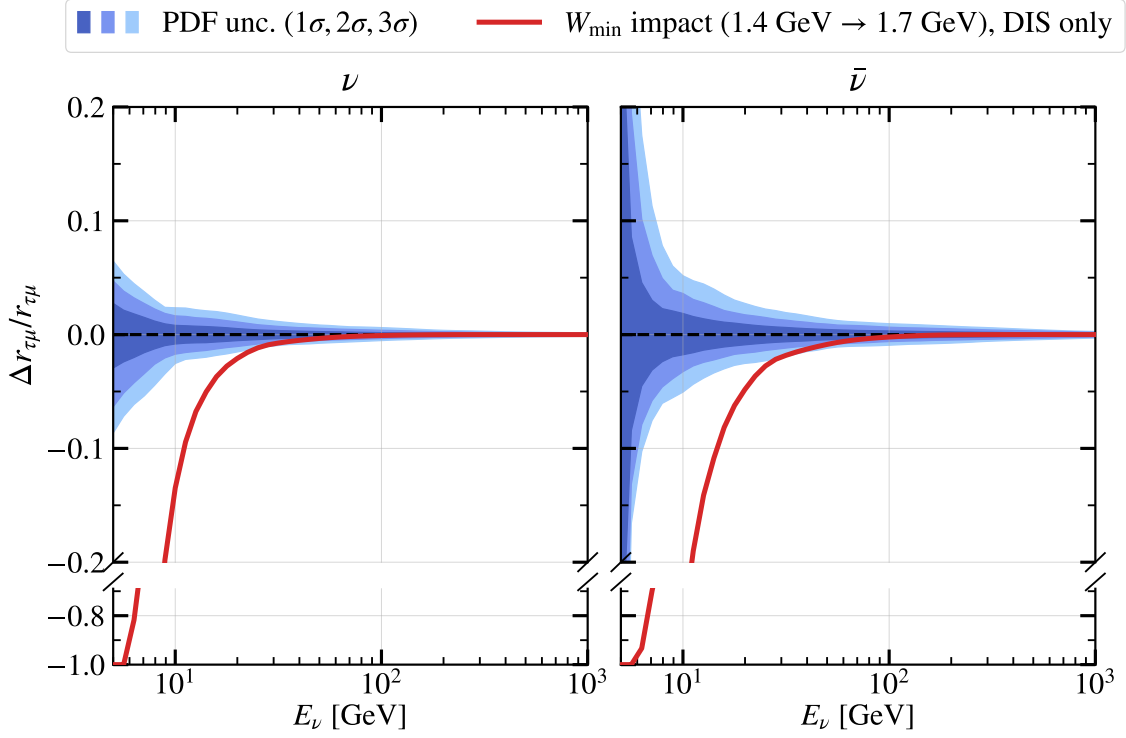
FIGURE 11.6 – Red curve: relative impact of a naive shift in the invariant hadronic mass threshold (W_{min}) from 1.4 GeV to 1.7 GeV on the ν_τ/ν_μ CC DIS cross section ratio ($r_{\tau\mu}$), disregarding the corresponding impact on the resonance cross section. Blue bands: relative impact of the NNPDF uncertainties on $r_{\tau\mu}$, as derived from the NNPDF Monte Carlo replicas. The y-axis is broken to show the full impact of W_{min} at the lowest considered neutrino energies. This figure serves only as an illustration of the W_{min} impact if not compensated by other parameters of the DIS and RES cross sections (see text for details).
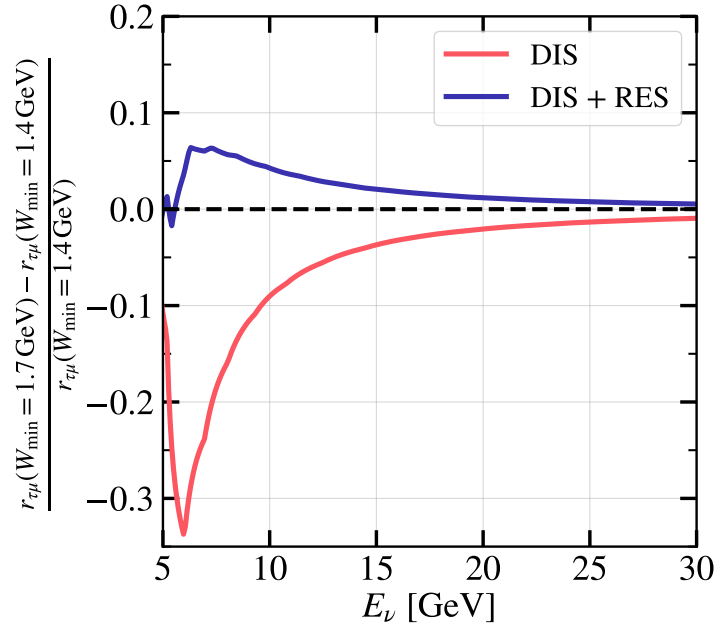
FIGURE 11.7 – Relative difference between the ν_τ/ν_μ charged current cross section ratios $r_{\tau\mu}$ evaluated with the invariant hadronic mass thresholds $W_{min} = 1.4$ GeV and $W_{min} = 1.7$ GeV. The red curve shows the difference in the ratios of the CC DIS cross sections, and the dark blue curve – CC DIS + RES cross sections. All cross sections are evaluated in GENIE 2.12.8 [158] with the Bodek-Yang-corrected GRV98 PDF at LO [177, 178].

efforts of the GENIE Collaboration [157]. One would then keep the tune parameters free in the experimental analysis in question and introduce a penalty term through

the covariance matrix. With this approach, the uncertainty on the $r_{\tau\mu}$ ratio would be automatically handled across all interaction channels, while the introduction of the $r_{\tau\mu}$ uncertainty due to $W_{\min}$ alone would not be complete and might lead to unphysical fit outcomes.

11.6 SUMMARY

In this study, we have considered the charged current deep inelastic scattering cross sections of muon and tau neutrinos and evaluated the uncertainty on their ratio at leading order in perturbation theory. This result was obtained by propagating the uncertainty on the parton distribution functions, for which we utilized the NNPDF Monte Carlo replicas, through the standard DIS calculation at LO. We found this uncertainty to be smaller than 3% for neutrinos with energies above 5 GeV and antineutrinos with energies above 10 GeV. Between 5 GeV and 10 GeV, the PDF-driven uncertainty on the antineutrino cross section ratio was found to vary between 2–30%. For the purpose of atmospheric $\nu_\mu \to \nu_\tau$ oscillation analyses with experiments such as IceCube-DeepCore and KM3NeT, these uncertainties can be considered negligible, as the atmospheric tau neutrino samples are dominated by ν_τ rather than $\bar{\nu}_\tau$. We therefore conclude that in atmospheric neutrino analyses, the ν_τ and the ν_μ DIS cross section uncertainties should be coupled through a single parameter instead of being represented by two independent parameters. For the beam experiments with an antineutrino mode, accounting for the PDF-driven uncertainty on the $\bar{\nu}_\tau/\bar{\nu}_\mu$ cross section ratio may be worthwhile. However, in the case of DUNE, we expect that the DIS event rate uncertainty in the $\gtrsim 5$ GeV region will be predominantly statistical in the first few years of exposure [292].

The smallness of the derived cross section ratio uncertainty stems from the correlations among the F_i structure functions at the PDF level. In particular, although the contribution of the F_5 structure function to the DIS cross section is modulated by the lepton mass, F_5 itself is strongly correlated with F_2 and F_1. This is reflected in the Albright-Jarlskog and the Callan-Gross relations at LO. This implies that one cannot allow for additional freedom in the tau neutrino CC DIS cross section through F_5, as varying F_5 must lead to the corresponding variations in F_1 and F_2. Any physical variations in F_5 therefore modify both tau and muon neutrino cross sections and leave their ratio largely unaffected.

We expect these conclusions to be stable with respect to the order in perturbation theory at which the calculation of the structure functions is performed, given that the correlations between F_1, F_2, and F_5 are well preserved at NLO [286]. However, this could be more rigorously confirmed by utilizing modern cross section calculation tools such as YADISM [323], which can be used up to N^{3}LO. Although the F_4 and F_5 structure functions – and therefore, the tau neutrino cross sections – are not implemented in the YADISM framework at the time of writing, making such functionality available would be a natural next step for this study.

VII

Summary and outlook

Summary

This thesis presented two complementary efforts to constrain the unitarity of neutrino mixing using neutrino oscillation data. First, we utilized a three-year atmospheric neutrino dataset from IceCube-DeepCore in combination with the reactor neutrino data from the Daya Bay and KamLAND experiments to provide a direct measurement of the individual elements of the neutrino mixing matrix. Subsequently, we set constraints on the non-unitarity metrics, namely the normalizations and closures of the rows and columns of the mixing matrix, and found a general agreement with unitarity across all channels. To the best of our knowledge, this was the first attempt at a comprehensive treatment of atmospheric neutrino data with the associated systematic uncertainties in a global fit for non-unitarity. The precision on the obtained row normalizations was found to be of order 10% in the muon and tau sector, and order 1% in the electron sector at a 1σ level – an indication of a much stronger constraining power of the reactor compared to the atmospheric neutrino data. We show this to be a consequence of large systematic uncertainties in the atmospheric neutrino fluxes, which induce a high degree of degeneracy between the constraints in the muon and the tau neutrino sectors.

Secondly, this work contributed to the development of the tau neutrino appearance analysis formerly initiated within the IceCube Collaboration. This analysis limits the manifestation of non-unitarity to the tau neutrino sector and is set to measure the tau neutrino normalization, N_{ν_τ}, with nine years of IceCube data. Motivated by the goodness-of-fit issues found in this analysis in the past, we conducted a deep investigation into the origins of disagreement between data and Monte Carlo simulation, identifying several major issues. The latter include the overestimation of the Cherenkov light yield by muons, as well as the misreconstruction of the muon energy due to high sensitivity of the likelihood-based reconstruction algorithm to noise pulses. These issues are now mitigated in the analysis pipeline by changing the software where the muon light yield simulation is handled, improving the robustness of noise cleaning, and shifting to a machine learning-based reconstruction algorithm. We showed that the sensitivity to N_{ν_τ} in this updated configuration is 9% at a 1σ level – a world-leading precision compared to the previous measurements by IceCube, OPERA, and Super-Kamiokande experiments. At the time of writing, the sample is undergoing the final checks prior to blind fits, which remain beyond the scope of this thesis.

Outlook

The overarching theme of this thesis is the role of systematic uncertainties and simulation mismodeling effects in searches for Beyond the Standard Model physics. A particularly important outcome of this work is the observation that the metrics of the mixing matrix non-unitarity are strongly correlated with the atmospheric neutrino systematic uncertainties. On one hand, this implies that the future iterations of this analysis would benefit from tighter priors on the nuisance parameters degenerate with the target physics quantities. This concerns, for example, the atmospheric neutrino flux parameters, whose scope and uncertainties can be re-defined through a modern data-driven framework in the future. On the other hand, this means that such parameters must necessarily be included in any future global fit studies and treated with extreme care, as their omission might lead to overly optimistic constraints or biased measurements. Our work on the oscNext event selection provides a prominent example of how the mistreatment of systematic effects can not just lead to a biased measurement, but also prevent such measurement from being made in the first place due to inability to describe the data with an insufficiently robust model.

Moving forward, the immediate goals are to finalize the measurement of the tau neutrino normalization with nine years of IceCube data, as well as to use this improved event selection to conduct a full non-unitarity analysis akin to the one presented in this thesis with three years of data. Moreover, the promising projections for the IceCube-Upgrade and JUNO experiments suggest that revisiting the non-unitary mixing analysis with the upcoming data from these experiments will be worthwhile.

The global fit for non-unitarity would further benefit from inclusion of more experiments than just reactor and atmospheric – by incorporating, for example, the solar and the accelerator neutrino data. Inspired by the findings of this thesis, we strongly recommend that the inclusion of these additional experiments in a global fit is accompanied by the identification and the full treatment of the relevant systematic uncertainties in the analysis.

VIII

Appendices

Additional Monte Carlo simulation checks

A.1 Light yield of the low-energy muon secondaries

As mentioned in Section 6.2, the parameterization of the Cherenkov light yield of the low-energy muon secondaries – which is currently employed in the IceCube Monte Carlo simulation chain at Step 2 – was based on the GEANT4 simulations of high-energy muons ($E_\mu > 100\,\text{GeV}$). The exact expression used in this parameterization reads

$$\frac{\ell_{\text{add}}}{\ell_\mu} = 0.1880 + 0.0206 \cdot \ln\left(\frac{E_\mu}{\text{GeV}}\right), \tag{A.1}$$

where ℓ_μ is the muon track length, and ℓ_{add} is the effective Cherenkov-radiated track length of all the charged low-energy secondaries ($e^\pm$) produced as the result of muon energy losses. The ratio $\frac{\ell_{\text{add}}}{\ell_\mu}$ is meant to approximate the ratio $\frac{N_{\gamma,\text{add}}}{N_{\gamma,\mu}}$, where $N_{\gamma,\mu}$ is the number of Cherenkov photons radiated by the muon, and $N_{\gamma,\text{add}}$ is the number of Cherenkov photons radiated collectively by all the charged low-energy secondaries.

As muons with much lower (down to sub-GeV) energies are part of the IceCube Monte Carlo simulation used in neutrino oscillation analyses, we attempt to test whether Equation (A.1) holds in this low-energy regime. For that purpose, we run standalone GEANT4 simulations for discrete muon energies in the $E_\mu \in [1\,\text{GeV}, 1\,\text{TeV}]$ range and track their secondary $e^\pm$ until stopping. The path length ℓ_i is saved for each secondary i along with its initial Lorentz velocity factor β_i. We then compute ℓ_{add} following the prescription of [221]:

$$\ell_{\text{add}} = \sum_i \hat{\ell}_i(E_i < E_{\text{max}}) = \frac{1 - \frac{1}{n^2\beta_i^2}}{1 - \frac{1}{n^2}} \cdot \ell_i, \tag{A.2}$$

where $E_{\text{max}} = 500\,\text{MeV}$ and n is set to 1.33 (the same value as used in [221], although a better representation of the average refractive index in the wavelength range of interest would be 1.31; see Figure 4.9).

In Figure A.1, we compare our results to the parameterization from Equation (A.1). We see that below $100\,\text{GeV}$, the dependence of the $\frac{\ell_{\text{add}}}{\ell_\mu}$ ratio on the muon energy is no
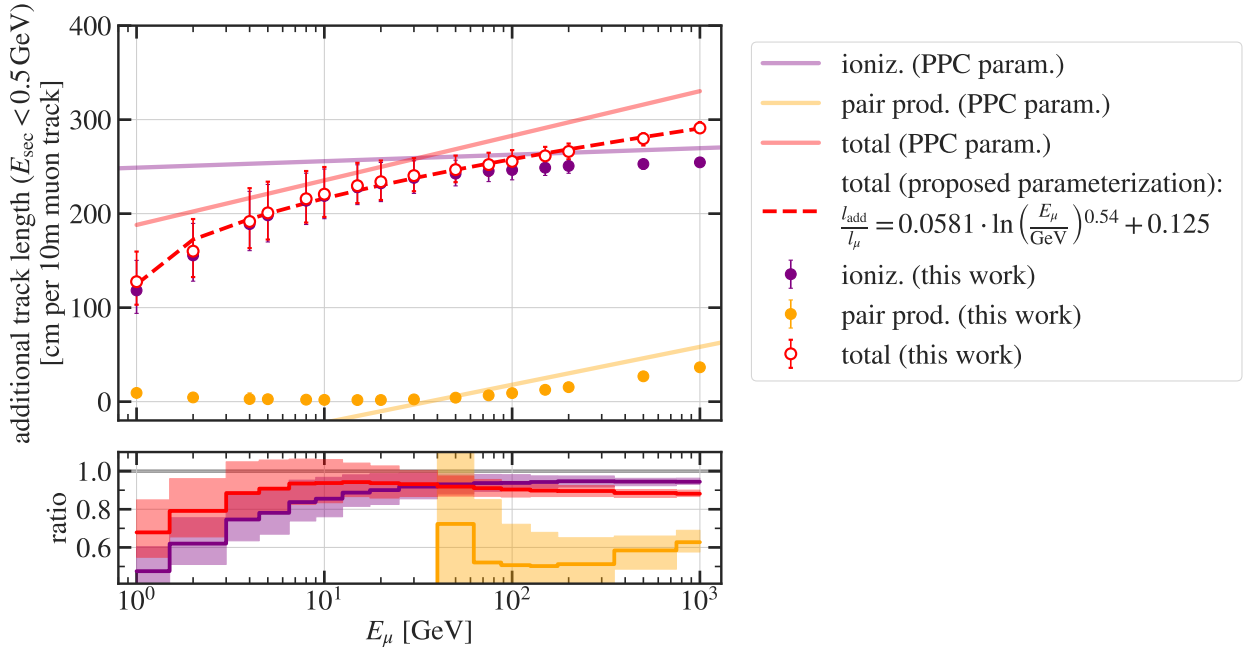
FIGURE A.1 – Comparison of the PPC parameterization [221] of the Cherenkov light yield of the low-energy muon secondaries ($E_{\text{sec}} < 500\,\text{MeV}$) with our standalone GEANT4 simulations, shown as a function of the parent muon energy E_μ. Note that instead of the number of photons, we report the fraction of the Cherenkov-radiating track length that these low-energy secondaries effectively contribute to the actual muon track length. The errorbars in our measured $\frac{\ell_{\text{add}}}{\ell_\mu}$ correspond to the 16%–84% quantile range in the $\frac{\ell_{\text{add}}}{\ell_\mu}$ distribution extracted from GEANT4.

longer linear with respect to $\ln E_\mu$, and there is significant (up to 30%) drop-off in our measured $\frac{\ell_{\text{add}}}{\ell_\mu}$ at low muon energies ($\sim1\,\text{GeV}$) compared to the PPC parameterization. We find that the following function better describes the shape of the $\frac{\ell_{\text{add}}}{\ell_\mu}$ dependence on E_μ:

$$\frac{\ell_{\text{add}}}{\ell_\mu} = 0.0581 \ln\left(\frac{E_\mu}{\text{GeV}}\right)^{0.54} + 0.125, \tag{A.3}$$

which we overlay with our measured $\frac{\ell_{\text{add}}}{\ell_\mu}$ in Figure A.1. We note that this function disagrees with the PPC parameterization not only in the low-energy range, which Equation (A.1) was not designed for, but also above $100\,\text{GeV}$, which this parameterization was tuned to. At $1\,\text{TeV}$, a relative discrepancy as large as 10% is found between our derived parameterization and Equation (A.1). The origin of this discrepancy was not identified at the time of writing, and could be speculatively attributed to the difference in the versions of the GEANT4 software[1] or the difference in the GEANT4 configuration settings such as physics lists[2]. For this reason, we do not explicitly use the newly derived parameterization in any of the studies described in this thesis.

---

[1] We used GEANT4 v11.1.0, which was first released in late 2022. The study [221] was completed in 2012, meaning that a much older GEANT4 version was used.

[2] Our work uses the QGSP_BERT physics list in GEANT4, which includes all standard EM processes and should in principle be equivalent to the physics list used in [221].

Supplemental material on neutrino mixing non-unitarity

b.1 non-unitary neutrino oscillations in vacuum
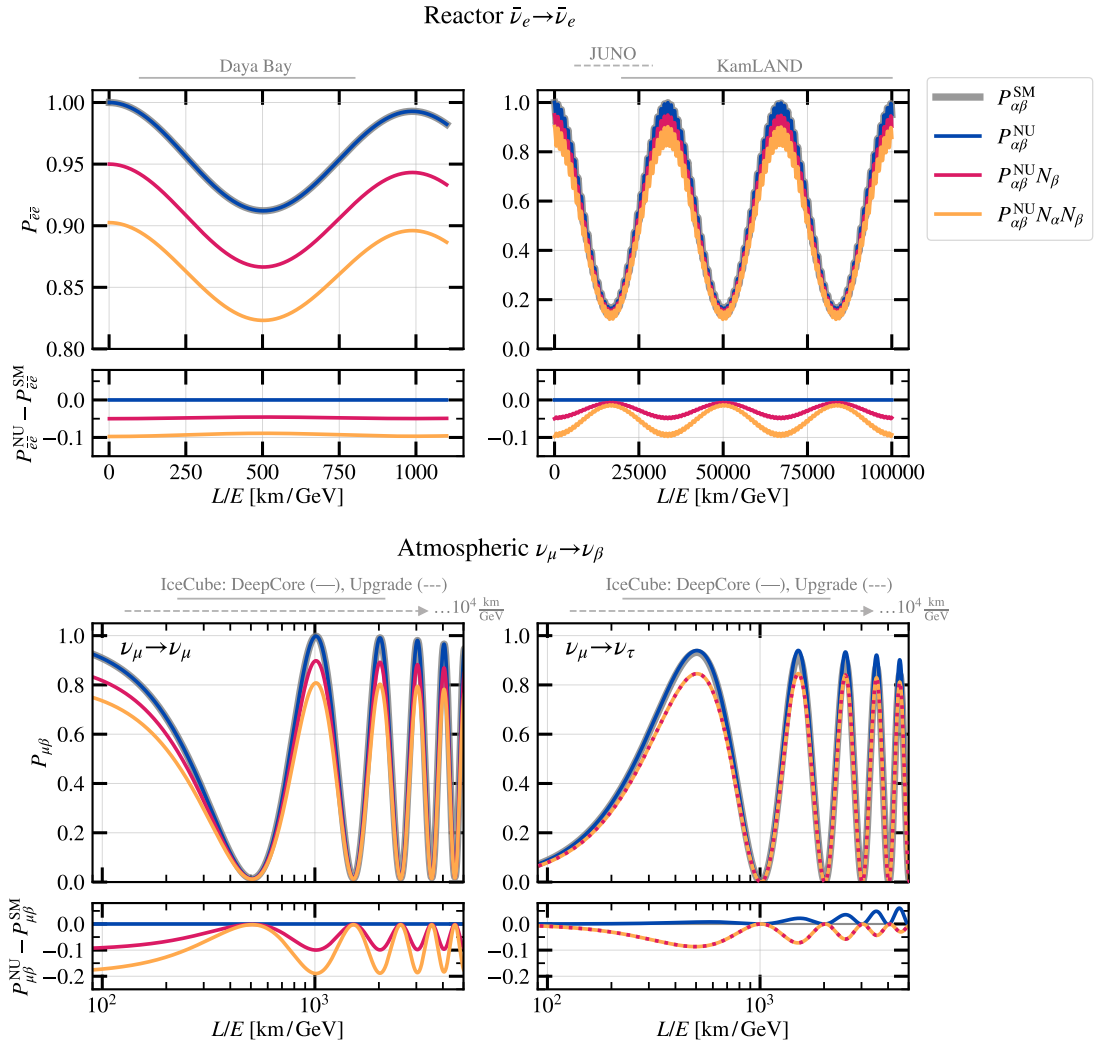
FIGURE B.1 – Comparison of the unitary (Standard Model, "SM") and the non-unitary ("NU") oscillation probabilities in vacuum. The legend follows that of Figure 2.5.
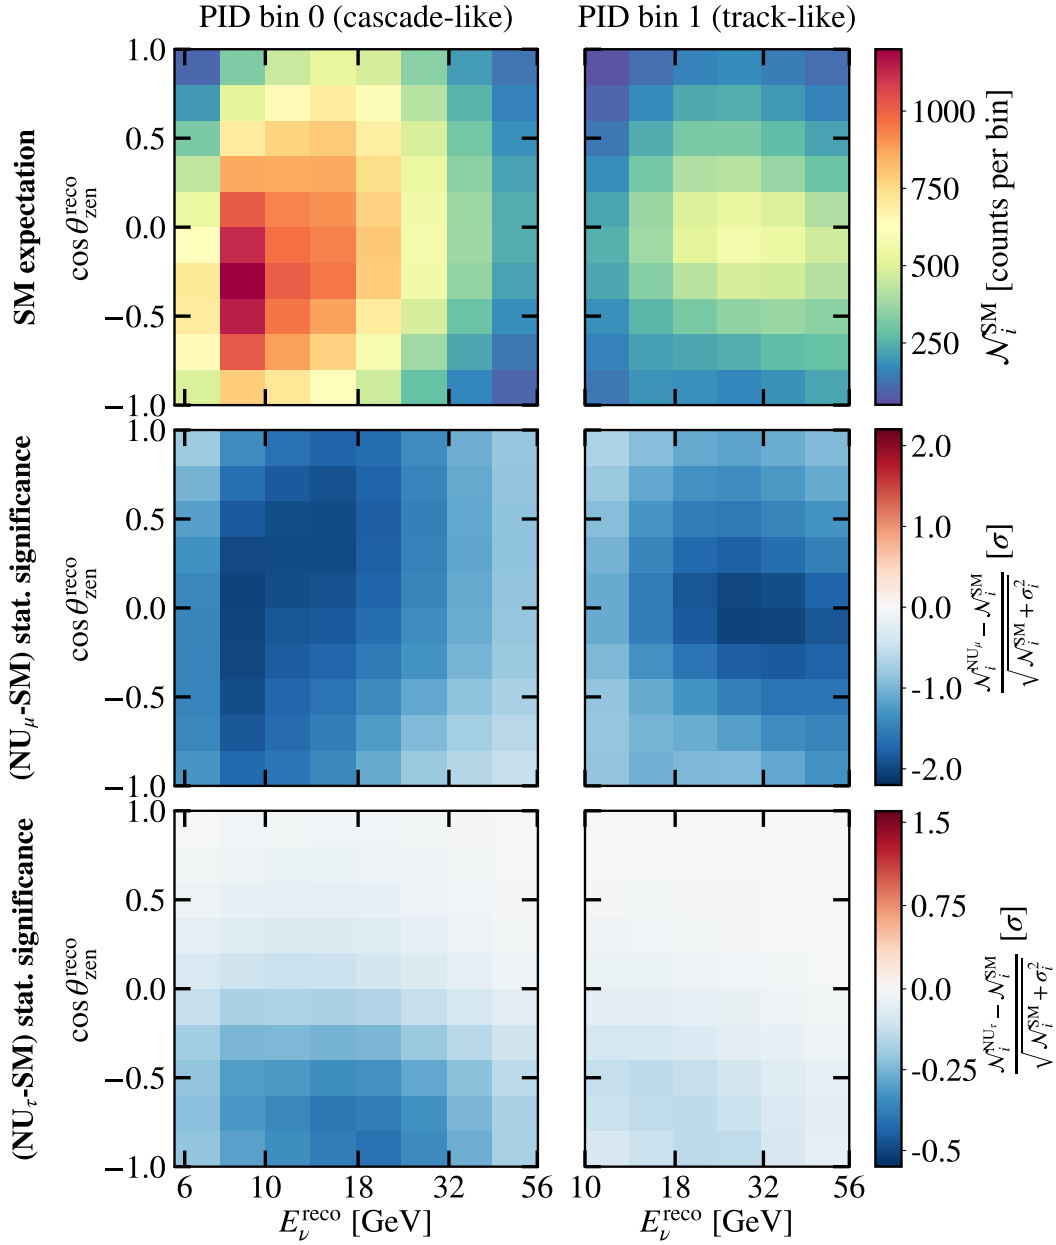
FIGURE B.2 – *Top:* the unitary (Standard Model, "SM") expectation of the DeepCore event rates, with the same parameter setup as in Figure 7.1. *Middle and bottom*: Statistical significances of the non-unitary ("NU") expectations. All three panels assume neutrino propagation in vacuum.

B.2 SYSTEMATIC UNCERTAINTIES IN THE GLOBAL FIT FOR NON-UNITARITY

B.2.1 *Daya Bay covariance matrix*

For the global non-unitarity analysis in Chapters 7 and 8, we use the covariance matrix of the unfolded $\bar{\nu}_e$ spectrum from Ref. [260], which we rescale to the expected spectrum in the 3158-day analysis. As the first step, we reproduce the correlation matrix ρ, whose elements ρ_{ij} represent the correlations between the different $\bar{\nu}_e$ energy bins $E_{\bar{\nu}_e, i}$ and $E_{\bar{\nu}_e, j}$:

$$\rho_{ij} = \frac{V_{ij}^{\text{DB, 2016}}}{\sqrt{V_{ii}^{\text{DB, 2016}} V_{jj}^{\text{DB, 2016}}}}, \tag{B.1}$$

where the "DB, 2016" superscript refers to the covariance matrix extracted from Table 13 of [260]. The diagonal elements of ρ represent self-correlations and are equal to 1.

The general relationship between the covariance matrix V and the correlation matrix ρ is

$$V_{ij} = \rho_{ij}\sigma_i\sigma_j, \tag{B.2}$$

where $\sigma_{i(j)}$ is the absolute uncertainty on the predicted event count $\mathcal{N}_i$ in the energy bin $i(j)$. We derive this absolute uncertainty as follows:

$$\sigma_i = \epsilon_i\mathcal{N}_i = \frac{\sqrt{V_{ii}^{\mathrm{DB,2016}}}}{\mathcal{N}_i^{\mathrm{DB,2016}}}\mathcal{N}_i, \tag{B.3}$$

where ϵ_i is the relative rate uncertainty computed from the covariance matrix element $V_{ii}^{\mathrm{DB,2016}}$ (Table 13 of [260]) and the corresponding IBD rate $\mathcal{N}_i^{\mathrm{DB,2016}}$ (Table 12 of [260]). Then, in the matrix form,

$$V = \sum_k \mathrm{diag}(\epsilon\mathcal{N})_k \cdot \rho \cdot \mathrm{diag}(\epsilon\mathcal{N})_k, \tag{B.4}$$

where "$\cdot$" denotes matrix multiplication and ϵ multiplies $\mathcal{N}$ elementwise as per Equation (B.3). The index k runs over the contributions of the six different reactors to the four different antineutrino detectors in EH3 and over the three data taking periods (6 AD, 7AD, and 8AD). We assume these contributions to be uncorrelated and do not add any cross-covariance terms in Equation (B.4). Each of the k subspectra $\mathcal{N}_k$ constitutes a fraction α_k of the total expected spectrum at EH3, such that $\mathcal{N}_k = \alpha_k\mathcal{N}_{\mathrm{EH3}}^{\mathrm{exp}}$. Then, the covariance matrix becomes

$$V = \left(\sum_k \alpha_k^2\right)\left[\mathrm{diag}(\epsilon\mathcal{N}_{\mathrm{EH3}}^{\mathrm{exp}}) \cdot \rho \cdot \mathrm{diag}(\epsilon\mathcal{N}_{\mathrm{EH3}}^{\mathrm{exp}})\right]. \tag{B.5}$$

We find the factors α_k by requiring that they are weighted according to the distance of each reactor core to each AD, the target mass and the efficiency of the ADs, as well as imposing the normalization condition $\sum_k \alpha_k = 1$. We obtain $\sum_k \alpha_k^2 \simeq 0.015$, and finally use V from Equation (B.5) with this prefactor to denote the systematic component V_{syst} of the covariance matrix in Equation (7.11).

We note that this is not the official approach used by the Daya Bay Collaboration, and that the above procedure reflects the steps taken in our specific study to approximate the systematic uncertainty on the expected event spectrum in the far hall of Daya Bay. We find that this treatment of the systematic uncertainties reproduces well the standard three-flavor oscillation contours from the 3158-day Daya Bay analysis, and the best-fit points in the $(\sin^2 2\theta_{13}, \Delta m_{32}^2)$ space are in excellent agreement (see Figure 7.3).

B.2.2 *Impact of the Daya Bay and KamLAND systematic uncertainties*
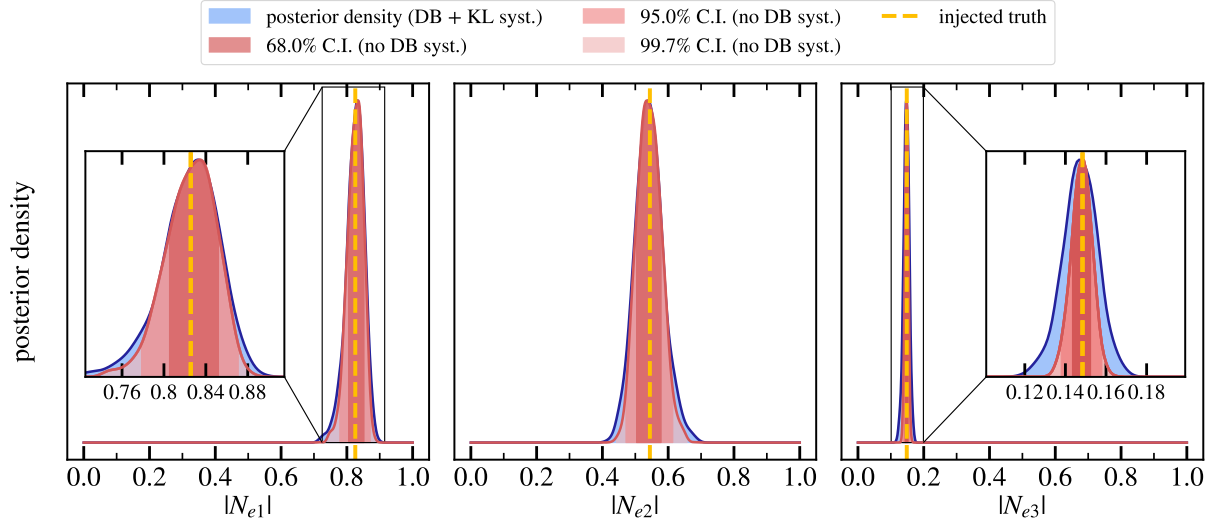
FIGURE B.3 – Projected impact of the Daya Bay flux systematic uncertainty implemented as described in Appendix B.2.1. The blue distributions show the posteriors obtained when including both statistical and systematic components of the covariance matrix. The red distributions correspond to the case of statistical-only uncertainty. The true values of $|N_{ei}|$ (dashed yellow lines) correspond to the elements of the unitary PMNS matrix computed with NuFit 5.2 oscillation parameters [65, 66].
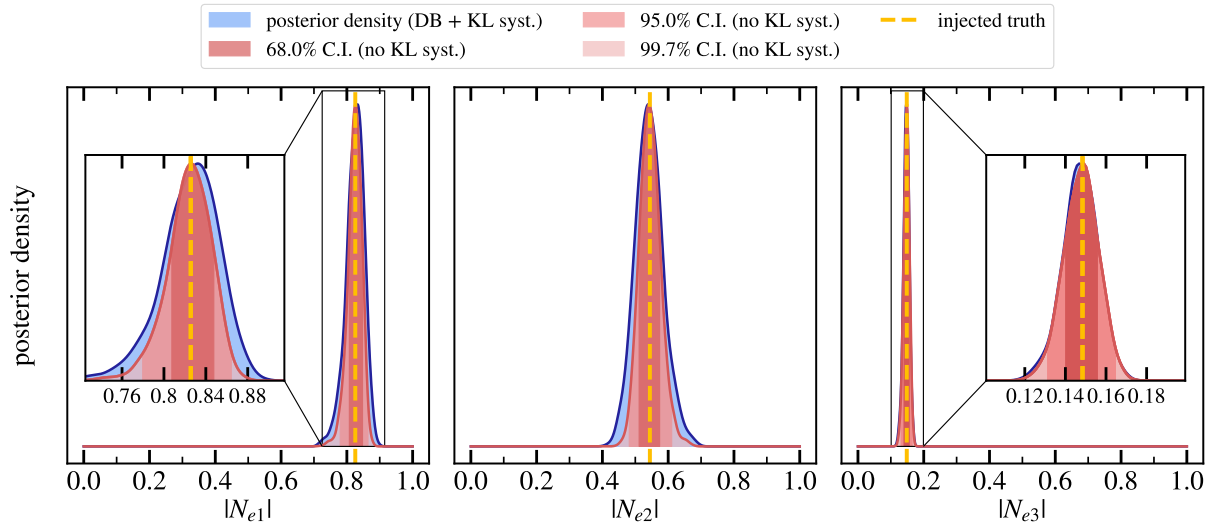
FIGURE B.4 – Same as Figure B.3, but applied to the KamLAND flux normalization ($\epsilon_{\mathcal{N}}$) and energy scale (ϵ_E) systematic uncertainties.

B.2.3   *IceCube-DeepCore and IceCube-Upgrade systematic parameters*
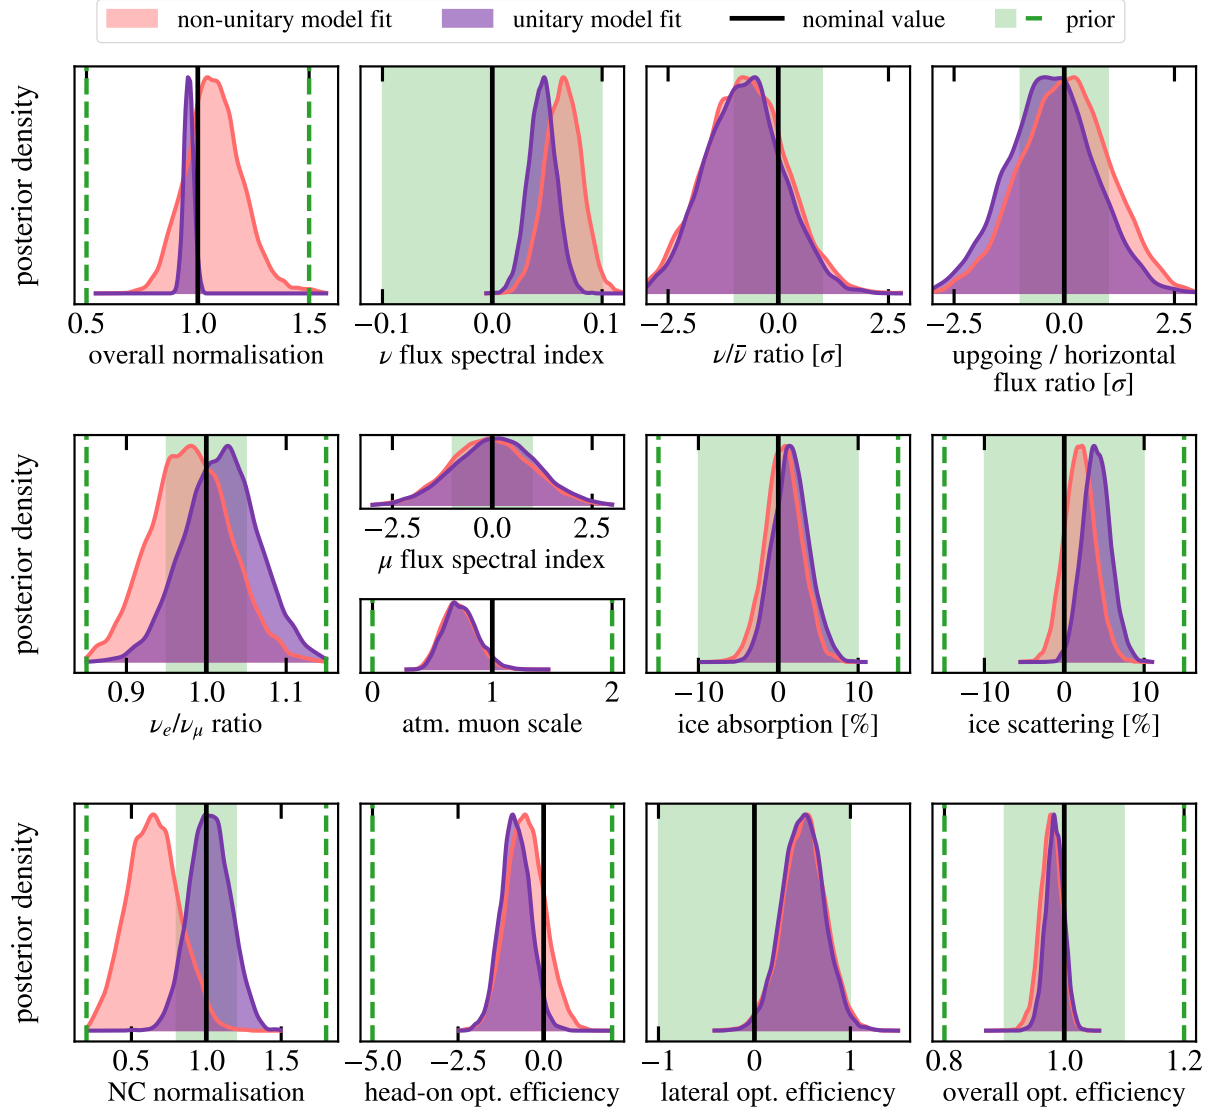
FIGURE B.5 – Posterior densities of the IceCube-DeepCore systematic parameters in the non-unitary (red) and the unitary (purple) model fits to the current atmospheric + reactor neutrino data. When applicable, the rectangular shaded green area corresponds to the 1σ range of a parameter with a Gaussian prior, and the dashed green lines – to the entire allowed range of this parameter. When only the dashed lines appear in a given panel, they represent the allowed range of a parameter with a uniform prior. The definitions and the units of the parameters follow [46].
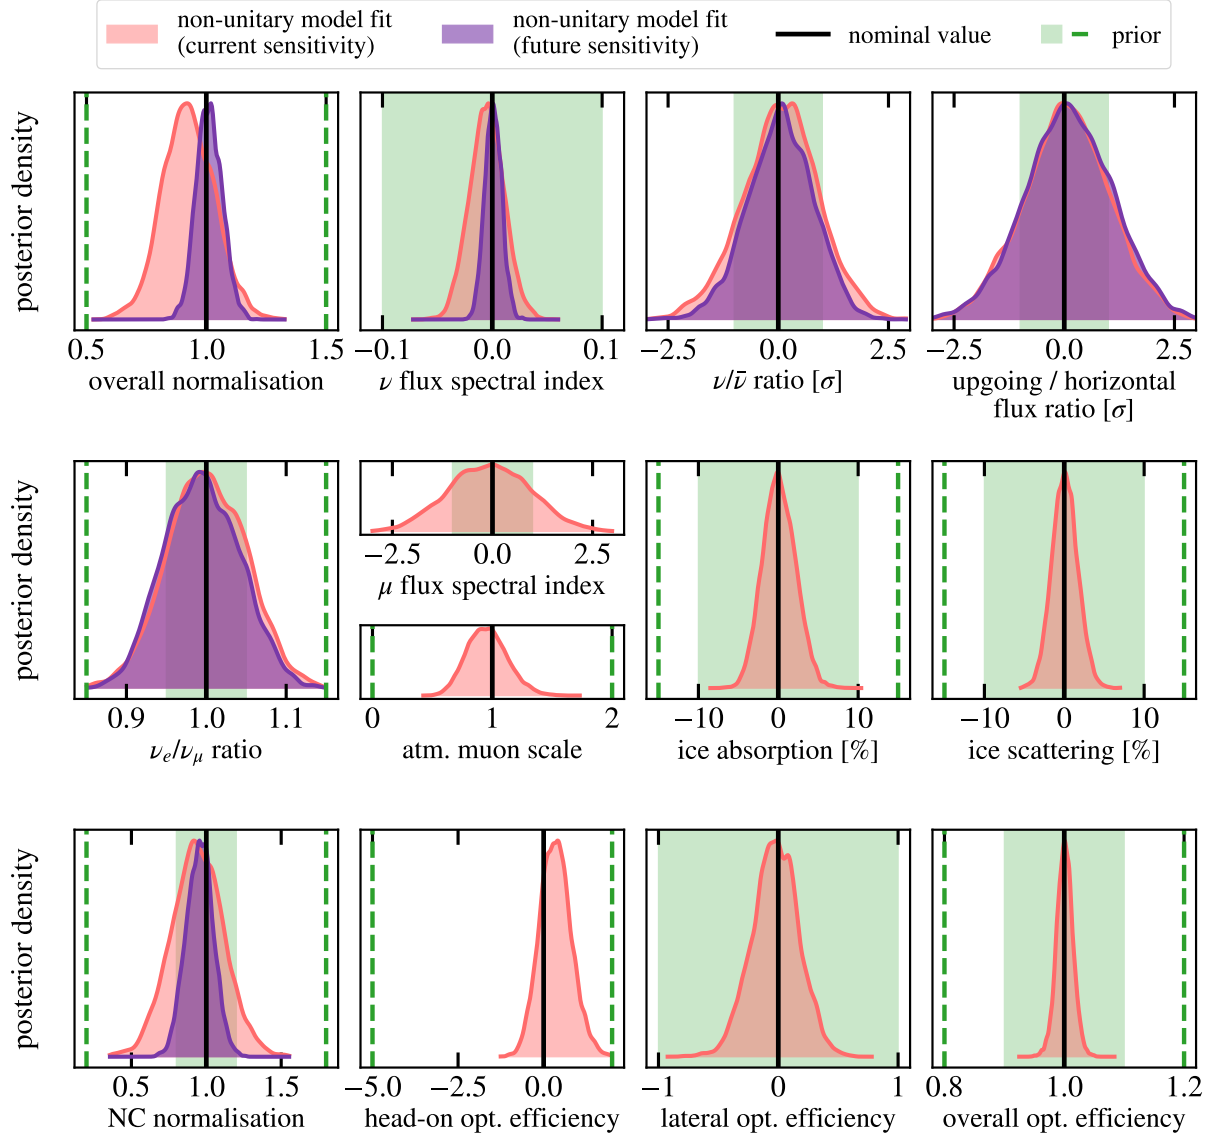
FIGURE B.6 – Posterior densities of the IceCube-DeepCore (red) and the IceCube-Upgrade (purple) systematic parameters in the non-unitary model fit to the pseudodata generated assuming unitarity and NuFit 5.2 oscillation parameters [65, 66]. The priors of the parameters are shown as the rectangular shaded area (1σ range) and the dashed lines (full range), analogously to Figure B.5. The definitions and the units of the parameters follow [46]. If no parameter posterior is shown for the IceCube-Upgrade in a given panel, this parameter has not been included in the IceCube-Upgrade projections.
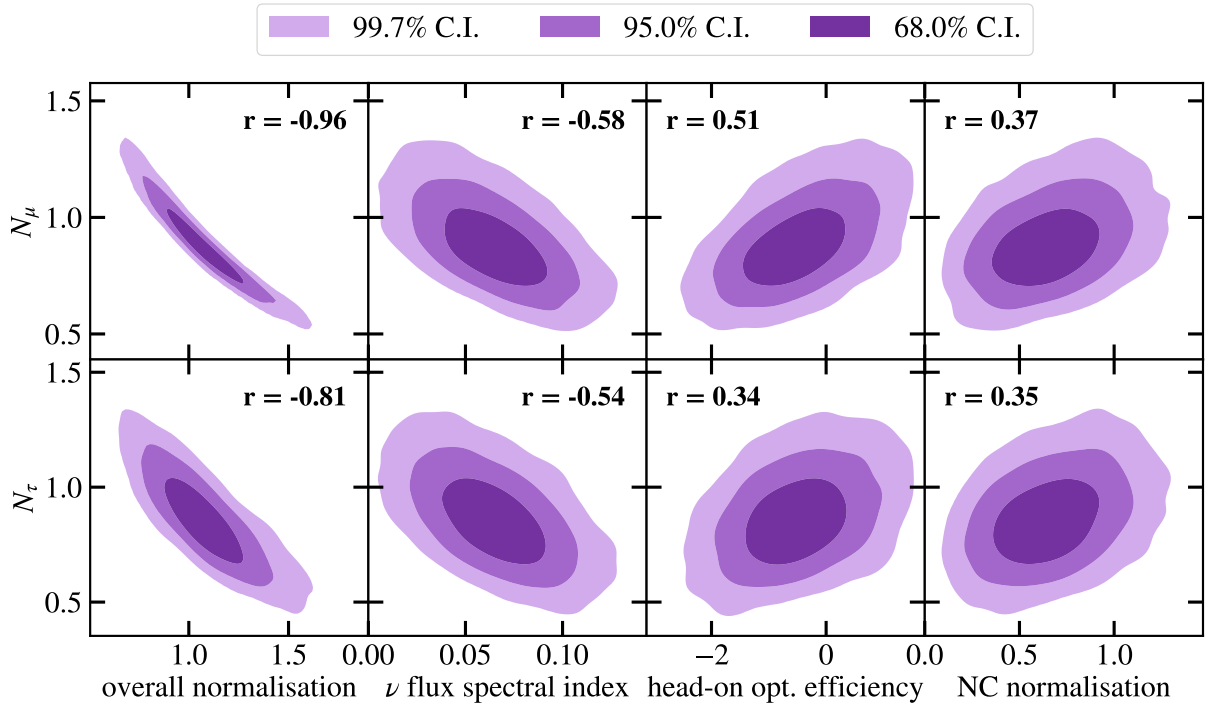
FIGURE B.7 – Correlations between the posteriors of the muon and tau row normalizations (N_μ and N_τ) and a subset of the IceCube-DeepCore systematic parameters, as found in the Bayesian fit of the non-unitary mixing model to the current atmospheric + reactor neutrino data. Each panel additionally reports the Pearson correlation coefficient value (r) between N_μ or N_τ and a given systematic parameter, and larger absolute values of r indicate stronger correlations.
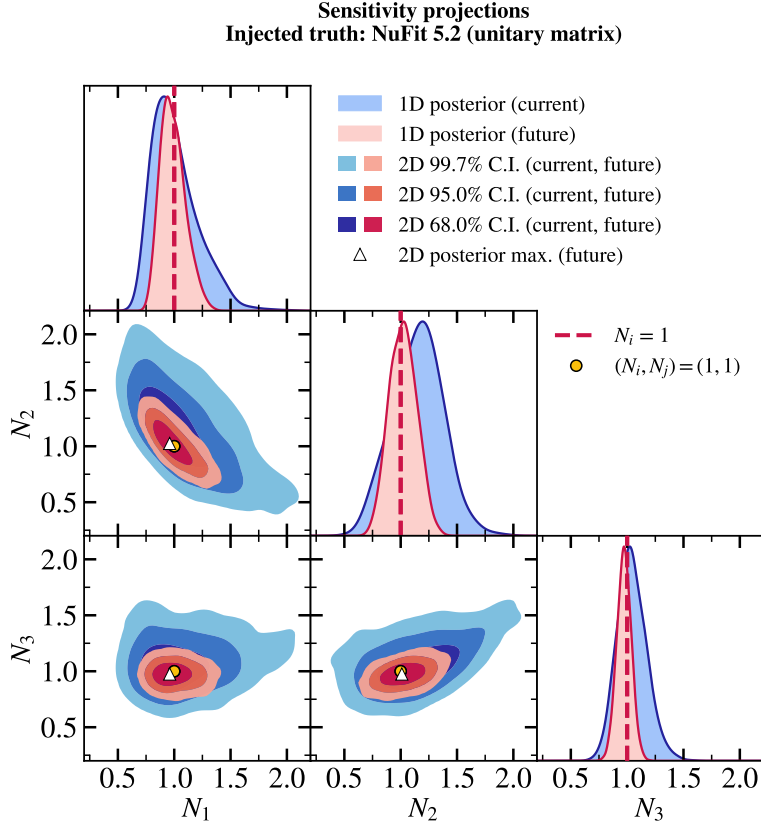
FIGURE B.8 – Same as Figure 8.2, but applied to the future projections of the matrix column normalizations using the combined IceCube-Upgrade, JUNO, and Daya Bay data.
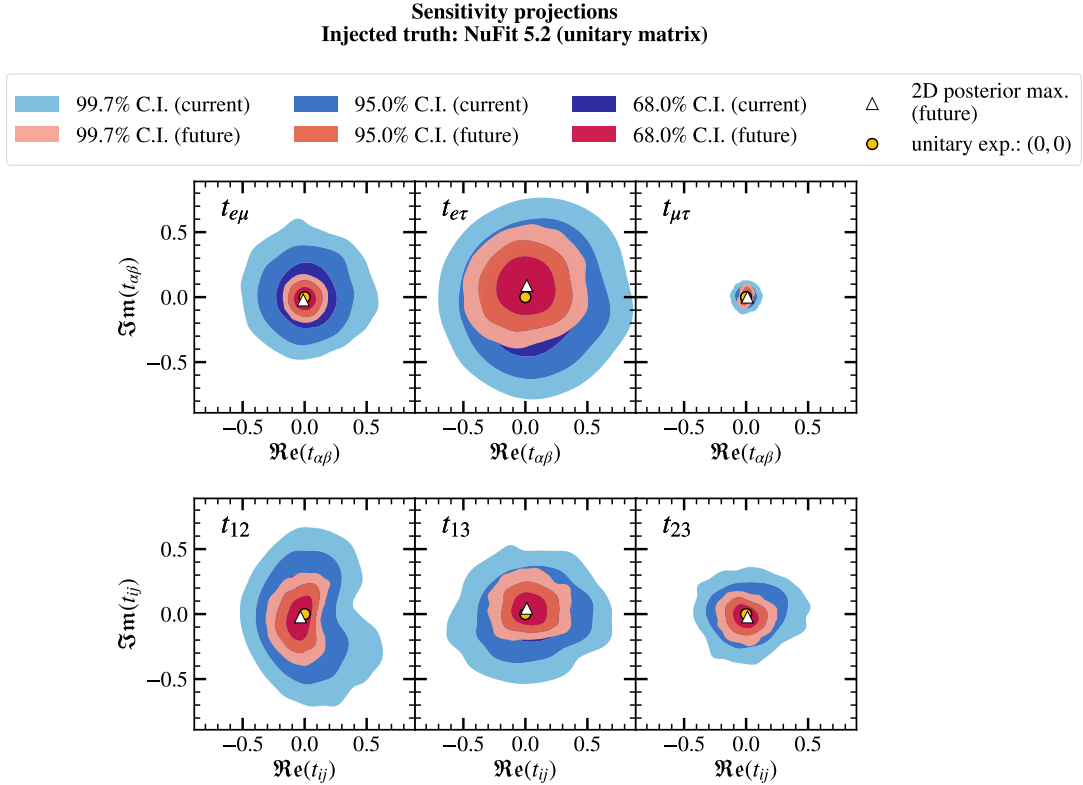
FIGURE B.9 – Same as Figure 8.4, but applied to the future projections of the matrix row and column closures using the combined IceCube-Upgrade, JUNO, and Daya Bay data.

SUPPLEMENTAL MATERIAL ON DEEP INELASTIC SCATTERING CROSS SECTIONS

C.1 IMPACT OF THE CHARM QUARK AND THE TARGET MASS CORRECTIONS

FIGURE C.1 – *First row*: densities of the sea strange and charm quarks inside inside a proton, shown at several fixed values of Q^2 as a function of Bjorken x. *Second to fourth rows*: ratio of the strange quark density to that of the charm quark (shown for three different sets of the y-axis limits for clarity). The displayed PDFs correspond to average of the 1000 NNPDF4.0 Monte Carlo replicas (0th members of the PDF set) [302].

FIGURE C.2 – Relative impact on the ν_τ/ν_μ CC DIS cross section ratio of the final-state charm mass correction as per Equation (11.5) (solid yellow curve); the target nucleon mass corrections as per Equations (11.7) and (11.8) (solid orange curve); the inclusion of the sea charm quark densities into the structure function calculations (dotted red curve); and the PDF uncertainties derived from the NNPDF Monte Carlo replicas [302] (blue bands).

Bibliography

[1] T. Kozynets et al. "Constraints on non-unitary neutrino mixing in light of atmospheric and reactor neutrino data". In: *J. High Energy Phys. (under review)* (July 2024). DOI: 10.48550/arXiv. 2407.20388. arXiv: 2407.20388 [hep-ph].

[2] T. Kozynets, T. Stuttard, and D. J. Koskinen. "Deep Inelastic Scattering Cross Section Uncertainties in Tau Neutrino Appearance Measurements". In: *Phys. Rev. D (under review)* (Sept. 2024). DOI: 10.48550/arXiv.2409.01258. arXiv: 2409.01258 [hep-ph].

[3] *Three-Year High-Statistics Neutrino Oscillation Samples*. DOI:10.21234/ac23-ra43. 2019. URL: https://doi.org/10.21234/ac23-ra43.

[4] M. J. Larson. "A Search for Tau Neutrino Appearance with IceCube-DeepCore". PhD thesis. Niels Bohr Institute, University of Copenhagen, 2018.

[5] *Neurthino.jl*. https://github.com/KM3NeT/Neurthino.jl (2024).

[6] T. Bose et al. "Report of the Topical Group on Physics Beyond the Standard Model at Energy Frontier for Snowmass 2021". In: (Sept. 2022). arXiv: 2209.13128 [hep-ph].

[7] C. A. Argüelles et al. "Snowmass White Paper: Beyond the Standard Model Effects on Neutrino Flavor". In: *Eur. Phys. J. C* 83.1 (2023), p. 15. DOI: 10.1140/epjc/s10052-022-11049-7. arXiv: 2203.10811 [hep-ph].

[8] 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].

[9] C. Giunti and C. W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. 2007. ISBN: 978-0-19-850871-7. DOI: 10.1093/acprof:oso/9780198508717.001.0001.

[10] 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].

[11] 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].

[12] M. G. Aartsen et al. "Observation of High-Energy Astrophysical Neutrinos in Three Years of IceCube Data". In: *Phys. Rev. Lett.* 113 (2014), p. 101101. DOI: 10.1103/PhysRevLett.113.101101. arXiv: 1405.5303 [astro-ph.HE].

[13] J. K. Becker. "High-energy neutrinos in the context of multimessenger physics". In: *Phys. Rept.* 458 (2008), pp. 173–246. DOI: 10.1016/j.physrep.2007.10.006. arXiv: 0710.1557 [astro-ph].

[14] J. Lesgourgues and S. Pastor. "Massive neutrinos and cosmology". In: *Phys. Rept.* 429 (2006), pp. 307–379. DOI: 10.1016/j.physrep.2006.04.001. arXiv: astro-ph/0603494.

[15] A. D. Dolgov. "Cosmological implications of neutrinos". In: *Surveys High Energ. Phys.* 17 (2002). Ed. by A. B. Kaidalov and M. I. Vysotsky, pp. 91–114. DOI: 10.1080/0124421021000054202. arXiv: hep-ph/0208222.

[16] S. Hannestad. "Neutrinos in cosmology". In: *New J. Phys.* 6 (2004), p. 108. DOI: 10.1088/1367-2630/6/1/108. arXiv: hep-ph/0404239.

[17] Q. R. Ahmad et al. "Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^{8}B solar neutrinos at the Sudbury Neutrino Observatory". In: *Phys. Rev. Lett.* 87 (2001), p. 071301. DOI: 10.1103/PhysRevLett.87.071301. arXiv: nucl-ex/0106015.

[18] R. Davis Jr., D. S. Harmer, and K. C. Hoffman. "Search for neutrinos from the sun". In: *Phys. Rev. Lett.* 20 (1968), pp. 1205–1209. DOI: 10.1103/PhysRevLett.20.1205.

[19] M. G. Aartsen et al. "Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A". In: *Science* 361.6398 (2018), eaat1378. DOI: 10.1126/science.aat1378. arXiv: 1807.08816 [astro-ph.HE].

[20] 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].

[21] 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].

[22] K. Hirata et al. "Observation of a Neutrino Burst from the Supernova SN 1987a". In: *Phys. Rev. Lett.* 58 (1987). Ed. by K. C. Wali, pp. 1490–1493. DOI: 10.1103/PhysRevLett.58.1490.

[23] A. Mirizzi et al. "Supernova Neutrinos: Production, Oscillations and Detection". In: *Riv. Nuovo Cim.* 39.1-2 (2016), pp. 1–112. DOI: 10.1393/ncr/i2016-10120-8. arXiv: 1508.00785 [astro-ph.HE].

[24] Q. Yuan, P.-F. Yin, and X.-J. Bi. "Neutrino emission of Fermi supernova remnants". In: *Astropart. Phys.* 35 (2011), pp. 33–38. DOI: 10.1016/j.astropartphys.2011.03.004. arXiv: 1010.1901 [astro-ph.HE].

[25] S. Chakraborty and I. Izaguirre. "Diffuse neutrinos from extragalactic supernova remnants: Dominating the 100 TeV IceCube flux". In: *Phys. Lett. B* 745 (2015), pp. 35–39. DOI: 10.1016/j.physletb.2015.04.032. arXiv: 1501.02615 [hep-ph].

[26] 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.

[27] C. Bemporad, G. Gratta, and P. Vogel. "Reactor Based Neutrino Oscillation Experiments". In: *Rev. Mod. Phys.* 74 (2002), p. 297. DOI: 10.1103/RevModPhys.74.297. arXiv: hep-ph/0107277.

[28] G. J. Feldman, J. Hartnell, and T. Kobayashi. "Long-baseline neutrino oscillation experiments". In: *Adv. High Energy Phys.* 2013 (2013), p. 475749. DOI: 10.1155/2013/475749. arXiv: 1210.1778 [hep-ex].

[29] C. S. Wu et al. "Experimental Test of Parity Conservation in β Decay". In: *Phys. Rev.* 105 (1957), pp. 1413–1414. DOI: 10.1103/PhysRev.105.1413.

[30] 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.

[31] B. Pontecorvo. "Inverse beta processes and nonconservation of lepton charge". In: *Zh. Eksp. Teor. Fiz.* 34 (1957), p. 247.

[32] Z. Maki, M. Nakagawa, and S. Sakata. "Remarks on the unified model of elementary particles". In: *Prog. Theor. Phys.* 28 (1962), pp. 870–880. DOI: 10.1143/PTP.28.870.

[33] P. W. Higgs. "Broken Symmetries and the Masses of Gauge Bosons". In: *Phys. Rev. Lett.* 13 (1964). Ed. by J. C. Taylor, pp. 508–509. DOI: 10.1103/PhysRevLett.13.508.

[34] R. L. Workman et al. "Review of Particle Physics". In: *PTEP* 2022 (2022), p. 083C01. DOI: 10.1093/ptep/ptac097.

[35] A. de Gouvêa. "Neutrino Mass Models". In: *Ann. Rev. Nucl. Part. Sci.* 66 (2016), pp. 197–217. DOI: 10.1146/annurev-nucl-102115-044600.

[36] E. Ma. "Neutrino Mass: Mechanisms and Models". In: (May 2009). arXiv: 0905.0221 [hep-ph].

[37] P. S. B. Dev et al. "Heavy Right-Handed Neutrino Dark Matter and PeV Neutrinos at IceCube". In: *JCAP* 08 (2016), p. 034. DOI: 10.1088/1475-7516/2016/08/034. arXiv: 1606.04517 [hep-ph].

[38] P. S. B. Dev, R. N. Mohapatra, and Y. Zhang. "Heavy Right-Handed Neutrino Dark matter in Left–Right Models". In: *Mod. Phys. Lett. A* 32.5 (2017), p. 1740007. DOI: 10.1142/S0217732317400077. arXiv: 1610.05738 [hep-ph].

[39] A. Boyarsky et al. "Sterile neutrino Dark Matter". In: *Prog. Part. Nucl. Phys.* 104 (2019), pp. 1–45. DOI: 10.1016/j.ppnp.2018.07.004. arXiv: 1807.07938 [hep-ph].

[40] F. Halzen and S. R. Klein. "IceCube: An Instrument for Neutrino Astronomy". In: *Rev. Sci. Instrum.* 81 (2010), p. 081101. DOI: 10.1063/1.3480478. arXiv: 1007.1247 [astro-ph.HE].

[41] M. G. Aartsen et al. "The IceCube Neutrino Observatory: Instrumentation and Online Systems". In: *JINST* 12.03 (2017). [Erratum: JINST 19, E05001 (2024)], P03012. DOI: 10.1088/1748-0221/12/03/P03012. arXiv: 1612.05093 [astro-ph.IM].

[42]   T. K. Gaisser, R. Engel, and E. Resconi. *Cosmic Rays and Particle Physics*. Cambridge University Press, 2016.

[43]   A. Fedynitch et al. "Calculation of conventional and prompt lepton fluxes at very high energy". In: *EPJ Web Conf.* 99 (2015). Ed. by D. Berge et al., p. 08001. DOI: 10.1051/epjconf/20159908001. arXiv: 1503.00544 [hep-ph].

[44]   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].

[45]   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].

[46]   M. G. Aartsen et al. "Measurement of Atmospheric Tau Neutrino Appearance with IceCube DeepCore". In: *Phys. Rev. D* 99.3 (2019), p. 032007. DOI: 10.1103/PhysRevD.99.032007. arXiv: 1901.05366 [hep-ex].

[47]   N. Agafonova et al. "Final Results of the OPERA Experiment on ν_τ Appearance in the CNGS Neutrino Beam". In: *Phys. Rev. Lett.* 120.21 (2018). [Erratum: Phys.Rev.Lett. 121, 139901 (2018)], p. 211801. DOI: 10.1103/PhysRevLett.120.211801. arXiv: 1804.04912 [hep-ex].

[48]   Z. Li et al. "Measurement of the Tau Neutrino Cross Section in Atmospheric Neutrino Oscillations with Super-Kamiokande". In: *Phys. Rev. D* 98.5 (2018), p. 052006. DOI: 10.1103/PhysRevD.98.052006. arXiv: 1711.09436 [hep-ex].

[49]   F. P. An et al. "Precision Measurement of Reactor Antineutrino Oscillation at Kilometer-Scale Baselines by Daya Bay". In: *Phys. Rev. Lett.* 130.16 (2023), p. 161802. DOI: 10.1103/PhysRevLett.130.161802. arXiv: 2211.14988 [hep-ex].

[50]   A. Gando et al. "Constraints on θ_{13} from a Three-Flavor Oscillation Analysis of Reactor Antineutrinos at KamLAND". In: *Phys. Rev. D* 83 (2011), p. 052002. DOI: 10.1103/PhysRevD.83.052002. arXiv: 1009.4771 [hep-ex].

[51]   S. Parke and M. Ross-Lonergan. "Unitarity and the Three Flavor Neutrino Mixing Matrix". In: *Phys. Rev. D* 93.11 (2016), p. 113009. DOI: 10.1103/PhysRevD.93.113009. arXiv: 1508.05095 [hep-ph].

[52]   P. B. Denton and J. Gehrlein. "New Oscillation and Scattering Constraints on the Tau Row Matrix Elements Without Assuming Unitarity". In: *JHEP* 06 (2022), p. 135. DOI: 10.1007/JHEP06(2022)135. arXiv: 2109.14575 [hep-ph].

[53]   S. A. R. Ellis, K. J. Kelly, and S. W. Li. "Current and Future Neutrino Oscillation Constraints on Leptonic Unitarity". In: *JHEP* 12 (2020), p. 068. DOI: 10.1007/JHEP12(2020)068. arXiv: 2008.01088 [hep-ph].

[54]   É. Bourbeau. "Measurement of Tau Neutrino Appearance in 8 Years of IceCube Data and a Search for Astrophysical Neutrinos from the Local Universe". PhD thesis. Niels Bohr Institute, University of Copenhagen, 2021.

[55]   T. Kozynets. "Angular distributions of atmospheric leptons via two-dimensional matrix cascade equations". Available at https://nbi.ku.dk/english/research/experimental-particle-physics/icecube/group_theses/Tetiana-Kozynets-MSc-Thesis_copy.pdf. Master's thesis. Copenhagen: Niels Bohr Institute, University of Copenhagen, 2022.

[56]   J. Schwichtenberg. *Physics from Symmetry*. Undergraduate Lecture Notes in Physics. Springer, 2018. DOI: 10.1007/978-3-319-66631-0.

[57]   J. Schwichtenberg. *No-Nonsense Quantum Field Theory: A Student-Friendly Introduction*. en. No-Nonsense Books, 2020.

[58]   A. J. Larkoski. *Elementary Particle Physics: An Intuitive Introduction*. Cambridge University Press, June 2019.

[59]   E. Noether. "Invariante Variationsprobleme". ger. In: *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* 1918 (1918), pp. 235–257. URL: http://eudml.org/doc/59024.

[60]   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.

[61]   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.

[62] T. A. Thompson, A. Burrows, and J. E. Horvath. "Mu and tau-neutrino thermalization and production in supernovae: Processes and time scales". In: *Phys. Rev. C* 62 (2000), p. 035802. DOI: 10.1103/PhysRevC.62.035802. arXiv: astro-ph/0003054.

[63] A. Betranhandy and E. O'Connor. "Impact of neutrino pair-production rates in Core-Collapse Supernovae". In: *Phys. Rev. D* 102 (2020), p. 123015. DOI: 10.1103/PhysRevD.102.123015. arXiv: 2010.02261 [astro-ph.HE].

[64] S. Schael et al. "Precision electroweak measurements on the *Z* resonance". In: *Phys. Rept.* 427 (2006), pp. 257–454. DOI: 10.1016/j.physrep.2005.12.006. arXiv: hep-ex/0509008.

[65] I. Esteban et al. "The Fate of Hints: Updated Global Analysis of Three-Flavor Neutrino Oscillations". In: *JHEP* 09 (2020), p. 178. DOI: 10.1007/JHEP09(2020)178. arXiv: 2007.14792 [hep-ph].

[66] *NuFit 5.2*. http://www.nu-fit.org/. 2022. URL: http://www.nu-fit.org/.

[67] S. Gariazzo et al. "Neutrino mass and mass ordering: no conclusive evidence for normal ordering". In: *JCAP* 10 (2022), p. 010. DOI: 10.1088/1475-7516/2022/10/010. arXiv: 2205.02195 [hep-ph].

[68] 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].

[69] E. Akhmedov. "Majorana neutrinos and other Majorana particles:Theory and experiment". In: Dec. 2014. arXiv: 1412.3320 [hep-ph].

[70] S. Esposito. *The Physics of Ettore Majorana*. Cambridge University Press, Dec. 2014.

[71] W. H. Furry. "On transition probabilities in double beta-disintegration". In: *Phys. Rev.* 56 (1939), pp. 1184–1193. DOI: 10.1103/PhysRev.56.1184.

[72] D. Q. Adams et al. "New Direct Limit on Neutrinoless Double Beta Decay Half-Life of Te128 with CUORE". In: *Phys. Rev. Lett.* 129.22 (2022), p. 222501. DOI: 10.1103/PhysRevLett.129.222501. arXiv: 2205.03132 [nucl-ex].

[73] G. Wang et al. "CUPID: CUORE (Cryogenic Underground Observatory for Rare Events) Upgrade with Particle IDentification". In: (Apr. 2015). arXiv: 1504.03599 [physics.ins-det].

[74] N. Abgrall et al. "The Large Enriched Germanium Experiment for Neutrinoless *ββ* Decay: LEGEND-1000 Preconceptual Design Report". In: (July 2021). arXiv: 2107.11462 [physics.ins-det].

[75] J. Martín-Albo et al. "Sensitivity of NEXT-100 to Neutrinoless Double Beta Decay". In: *JHEP* 05 (2016), p. 159. DOI: 10.1007/JHEP05(2016)159. arXiv: 1511.09246 [physics.ins-det].

[76] G. Adhikari et al. "nEXO: neutrinoless double beta decay search beyond 10^{28} year half-life sensitivity". In: *J. Phys. G* 49.1 (2022), p. 015104. DOI: 10.1088/1361-6471/ac3631. arXiv: 2106.16243 [nucl-ex].

[77] J. J. Gómez-Cadenas et al. "The search for neutrinoless double-beta decay". In: *Riv. Nuovo Cim.* 46.10 (2023), pp. 619–692. DOI: 10.1007/s40766-023-00049-2.

[78] H. Burkhardt et al. "First Evidence for Direct CP Violation". In: *Phys. Lett. B* 206 (1988), pp. 169–176. DOI: 10.1016/0370-2693(88)91282-8.

[79] V. Fanti et al. "A New measurement of direct CP violation in two pion decays of the neutral kaon". In: *Phys. Lett. B* 465 (1999), pp. 335–348. DOI: 10.1016/S0370-2693(99)01030-8. arXiv: hep-ex/9909022.

[80] A. Alavi-Harati et al. "Observation of direct CP violation in $K_{S,L} \to \pi\pi$ decays". In: *Phys. Rev. Lett.* 83 (1999), pp. 22–27. DOI: 10.1103/PhysRevLett.83.22. arXiv: hep-ex/9905060.

[81] B. Aubert et al. "Observation of direct CP violation in $B^0 \to K^+\pi^-$ decays". In: *Phys. Rev. Lett.* 93 (2004), p. 131801. DOI: 10.1103/PhysRevLett.93.131801. arXiv: hep-ex/0407057.

[82] Y. Chao et al. "Evidence for direct CP violation in B0 —> K+ pi- decays". In: *Phys. Rev. Lett.* 93 (2004), p. 191802. DOI: 10.1103/PhysRevLett.93.191802. arXiv: hep-ex/0408100.

[83] P. del Amo Sanchez et al. "Measurement of CP observables in $B^{+-} -> D_{CP}K^{+-}$ decays and constraints on the CKM angle γ". In: *Phys. Rev. D* 82 (2010), p. 072004. DOI: 10.1103/PhysRevD.82.072004. arXiv: 1007.0504 [hep-ex].

[84] A. Poluektov et al. "Evidence for direct CP violation in the decay B->D(*)K, D->KsPi+Pi- and measurement of the CKM phase phi3". In: *Phys. Rev. D* 81 (2010), p. 112002. DOI: 10.1103/PhysRevD.81.112002. arXiv: 1003.3360 [hep-ex].

[85] R. Aaij et al. "Observation of CP violation in $B^{\pm} \to DK^{\pm}$ decays". In: *Phys. Lett. B* 712 (2012). [Erratum: Phys.Lett.B 713, 351 (2012)], pp. 203–212. DOI: 10.1016/j.physletb.2012.04.060. arXiv: 1203.3662 [hep-ex].

[86] R. Aaij et al. "Observation of CP Violation in Charm Decays". In: *Phys. Rev. Lett.* 122.21 (2019), p. 211803. DOI: 10.1103/PhysRevLett.122.211803. arXiv: 1903.08726 [hep-ex].

[87] N. Cabibbo. "Unitary Symmetry and Leptonic Decays". In: *Phys. Rev. Lett.* 10 (1963), pp. 531–533. DOI: 10.1103/PhysRevLett.10.531.

[88] M. Kobayashi and T. Maskawa. "CP Violation in the Renormalizable Theory of Weak Interaction". In: *Prog. Theor. Phys.* 49 (1973), pp. 652–657. DOI: 10.1143/PTP.49.652.

[89] G. R. Farrar and M. E. Shaposhnikov. "Baryon asymmetry of the universe in the standard electroweak theory". In: *Phys. Rev. D* 50 (1994), p. 774. DOI: 10.1103/PhysRevD.50.774. arXiv: hep-ph/9305275.

[90] A. Riotto. "Theories of baryogenesis". In: *ICTP Summer School in High-Energy Physics and Cosmology*. July 1998, pp. 326–436. arXiv: hep-ph/9807454.

[91] K. Abe et al. "Measurements of neutrino oscillation parameters from the T2K experiment using 3.6×10^{21} protons on target". In: *Eur. Phys. J. C* 83.9 (2023), p. 782. DOI: 10.1140/epjc/s10052-023-11819-x. arXiv: 2303.03222 [hep-ex].

[92] M. A. Acero et al. "Improved measurement of neutrino oscillation parameters by the NOvA experiment". In: *Phys. Rev. D* 106.3 (2022), p. 032004. DOI: 10.1103/PhysRevD.106.032004. arXiv: 2108.08219 [hep-ex].

[93] K. Sutton. *Results from the Joint NOvA-T2K Analysis*. XXXI International Conference on Neutrino Physics and Astrophysics (NEUTRINO 2024). 2024. URL: https://agenda.infn.it/event/37867/contributions/227857/.

[94] B. Abi et al. "Long-baseline neutrino oscillation physics potential of the DUNE experiment". In: *Eur. Phys. J. C* 80.10 (2020), p. 978. DOI: 10.1140/epjc/s10052-020-08456-z. arXiv: 2006.16043 [hep-ex].

[95] J. Bian et al. "Hyper-Kamiokande Experiment: A Snowmass White Paper". In: *Snowmass 2021*. Mar. 2022. arXiv: 2203.02029 [hep-ex].

[96] J. Hofestädt, M. Bruchner, and T. Eberl. "Super-ORCA: Measuring the leptonic CP-phase with Atmospheric Neutrinos and Beam Neutrinos". In: *PoS* ICRC2019 (2020), p. 911. DOI: 10.22323/1.358.0911. arXiv: 1907.12983 [hep-ex].

[97] A. Aguilar et al. "Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam". In: *Phys. Rev. D* 64 (2001), p. 112007. DOI: 10.1103/PhysRevD.64.112007. arXiv: hep-ex/0104049.

[98] J. N. Abdurashitov et al. "Measurement of the solar neutrino capture rate with gallium metal". In: *Phys. Rev. C* 60 (1999), p. 055801. DOI: 10.1103/PhysRevC.60.055801. arXiv: astro-ph/9907113.

[99] W. Hampel et al. "GALLEX solar neutrino observations: Results for GALLEX IV". In: *Phys. Lett. B* 447 (1999), pp. 127–133. DOI: 10.1016/S0370-2693(98)01579-2.

[100] A. A. Aguilar-Arevalo et al. "Improved Search for $\bar{\nu}_\mu \to \bar{\nu}_e$ Oscillations in the MiniBooNE Experiment". In: *Phys. Rev. Lett.* 110 (2013), p. 161801. DOI: 10.1103/PhysRevLett.110.161801. arXiv: 1303.2588 [hep-ex].

[101] A. A. Aguilar-Arevalo et al. "Significant Excess of ElectronLike Events in the MiniBooNE Short-Baseline Neutrino Experiment". In: *Phys. Rev. Lett.* 121.22 (2018), p. 221801. DOI: 10.1103/PhysRevLett.121.221801. arXiv: 1805.12028 [hep-ex].

[102] P. Abratenko et al. "First Constraints on Light Sterile Neutrino Oscillations from Combined Appearance and Disappearance Searches with the MicroBooNE Detector". In: *Phys. Rev. Lett.* 130.1 (2023), p. 011801. DOI: 10.1103/PhysRevLett.130.011801. arXiv: 2210.10216 [hep-ex].

[103] R. Abbasi et al. "Search for a light sterile neutrino with 7.5 years of IceCube DeepCore data". In: (July 2024). arXiv: 2407.01314 [hep-ex].

[104] P. Adamson et al. "Improved Constraints on Sterile Neutrino Mixing from Disappearance Searches in the MINOS, MINOS+, Daya Bay, and Bugey-3 Experiments". In: *Phys. Rev. Lett.* 125.7 (2020), p. 071801. DOI: 10.1103/PhysRevLett.125.071801. arXiv: 2002.00301 [hep-ex].

[105] M. Aker et al. "Improved eV-scale sterile-neutrino constraints from the second KATRIN measurement campaign". In: *Phys. Rev. D* 105.7 (2022), p. 072004. DOI: 10.1103/PhysRevD.105.072004. arXiv: 2201.11593 [hep-ex].

[106] M. Drewes et al. "A White Paper on keV Sterile Neutrino Dark Matter". In: *JCAP* 01 (2017), p. 025. DOI: 10.1088/1475-7516/2017/01/025. arXiv: 1602.04816 [hep-ph].

[107] A. Boyarsky et al. "Unidentified Line in X-Ray Spectra of the Andromeda Galaxy and Perseus Galaxy Cluster". In: *Phys. Rev. Lett.* 113 (2014), p. 251301. DOI: 10.1103/PhysRevLett.113.251301. arXiv: 1402.4119 [astro-ph.CO].

[108] E. Bulbul et al. "Detection of An Unidentified Emission Line in the Stacked X-ray spectrum of Galaxy Clusters". In: *Astrophys. J.* 789 (2014), p. 13. DOI: 10.1088/0004-637X/789/1/13. arXiv: 1402.2301 [astro-ph.CO].

[109] E. Fernandez-Martinez, J. Hernandez-Garcia, and J. Lopez-Pavon. "Global Constraints on Heavy Neutrino Mixing". In: *JHEP* 08 (2016), p. 033. DOI: 10.1007/JHEP08(2016)033. arXiv: 1605.08774 [hep-ph].

[110] M. Aker et al. "Direct neutrino-mass measurement with sub-electronvolt sensitivity". In: *Nature Phys.* 18.2 (2022), pp. 160–166. DOI: 10.1038/s41567-021-01463-1. arXiv: 2105.08533 [hep-ex].

[111] M. Agostini et al. "Final Results of GERDA on the Search for Neutrinoless Double-β Decay". In: *Phys. Rev. Lett.* 125.25 (2020), p. 252502. DOI: 10.1103/PhysRevLett.125.252502. arXiv: 2009.06079 [nucl-ex].

[112] A. Gando et al. "Search for Majorana Neutrinos near the Inverted Mass Hierarchy Region with KamLAND-Zen". In: *Phys. Rev. Lett.* 117.8 (2016). [Addendum: Phys.Rev.Lett. 117, 109903 (2016)], p. 082503. DOI: 10.1103/PhysRevLett.117.082503. arXiv: 1605.02889 [hep-ex].

[113] N. Aghanim et al. "Planck 2018 results. VI. Cosmological parameters". In: *Astron. Astrophys.* 641 (2020). [Erratum: Astron.Astrophys. 652, C4 (2021)], A6. DOI: 10.1051/0004-6361/201833910. arXiv: 1807.06209 [astro-ph.CO].

[114] S. Alam et al. "Completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey: Cosmological implications from two decades of spectroscopic surveys at the Apache Point Observatory". In: *Phys. Rev. D* 103.8 (2021), p. 083533. DOI: 10.1103/PhysRevD.103.083533. arXiv: 2007.08991 [astro-ph.CO].

[115] A. G. Adame et al. "DESI 2024 VI: Cosmological Constraints from the Measurements of Baryon Acoustic Oscillations". In: (Apr. 2024). arXiv: 2404.03002 [astro-ph.CO].

[116] J. A. Formaggio, A. L. C. de Gouvêa, and R. G. H. Robertson. "Direct Measurements of Neutrino Mass". In: *Phys. Rept.* 914 (2021), pp. 1–54. DOI: 10.1016/j.physrep.2021.02.002. arXiv: 2102.00594 [nucl-ex].

[117] M. Gerbino et al. "Synergy between cosmological and laboratory searches in neutrino physics". In: *Phys. Dark Univ.* 42 (2023), p. 101333. DOI: 10.1016/j.dark.2023.101333. arXiv: 2203.07377 [hep-ph].

[118] M. Aker et al. "Improved Upper Limit on the Neutrino Mass from a Direct Kinematic Method by KATRIN". In: *Phys. Rev. Lett.* 123.22 (2019), p. 221802. DOI: 10.1103/PhysRevLett.123.221802. arXiv: 1909.06048 [hep-ex].

[119] S. Weinberg. "Baryon and Lepton Nonconserving Processes". In: *Phys. Rev. Lett.* 43 (1979), pp. 1566–1570. DOI: 10.1103/PhysRevLett.43.1566.

[120] A. V. Manohar. "Introduction to Effective Field Theories". In: (Apr. 2018). Ed. by S. Davidson et al. DOI: 10.1093/oso/9780198855743.003.0002. arXiv: 1804.05863 [hep-ph].

[121] J. F. Kamenik and M. Nemevsek. "Lepton flavor violation in type I + III seesaw". In: *JHEP* 11 (2009). [Erratum: JHEP 03, 033 (2014)], p. 023. DOI: 10.1088/1126-6708/2009/11/023. arXiv: 0908.3451 [hep-ph].

[122] J. Herrero-Garcia et al. "The Zee–Babu model revisited in the light of new data". In: *Nucl. Phys. B* 885 (2014), pp. 542–570. DOI: 10.1016/j.nuclphysb.2014.06.001. arXiv: 1402.4491 [hep-ph].

[123] S. M. Boucenna, S. Morisi, and J. W. F. Valle. "The low-scale approach to neutrino masses". In: *Adv. High Energy Phys.* 2014 (2014), p. 831598. DOI: 10.1155/2014/831598. arXiv: 1404.3751 [hep-ph].

[124] M. Wallraff and C. Wiebusch. "Calculation of oscillation probabilities of atmospheric neutrinos using nuCraft". In: *Comput. Phys. Commun.* 197 (2015), pp. 185–189. DOI: 10.1016/j.cpc.2015.07.010. arXiv: 1409.1387 [astro-ph.IM].

[125] F. J. Escrihuela et al. "Probing CP Violation with Non-Unitary Mixing in Long-Baseline Neutrino Oscillation Experiments: DUNE as a Case Study". In: *New J. Phys.* 19.9 (2017), p. 093005. DOI: 10.1088/1367-2630/aa79ec. arXiv: 1612.07377 [hep-ph].

[126] M. Blennow and A. Y. Smirnov. "Neutrino Propagation in Matter". In: *Adv. High Energy Phys.* 2013 (2013), p. 972485. DOI: 10.1155/2013/972485. arXiv: 1306.2903 [hep-ph].

[127] V. D. Barger et al. "Matter Effects on Three-Neutrino Oscillations". In: *Phys. Rev. D* 22 (1980), p. 2718. DOI: 10.1103/PhysRevD.22.2718.

[128] M. Blennow et al. "Non-Unitarity, Sterile Neutrinos, and Non-Standard Neutrino Interactions". In: *JHEP* 04 (2017), p. 153. DOI: 10.1007/JHEP04(2017)153. arXiv: 1609.08637 [hep-ph].

[129] S. Antusch et al. "Unitarity of the Leptonic Mixing Matrix". In: *JHEP* 10 (2006), p. 084. DOI: 10.1088/1126-6708/2006/10/084. arXiv: hep-ph/0607020.

[130] D. Aloni and A. Dery. "Revisiting Leptonic Non-Unitarity in Light of FASERν". In: (Nov. 2022). arXiv: 2211.09638 [hep-ph].

[131] H. P. Dembinski et al. "Data-driven model of the cosmic-ray flux and mass composition from 10 GeV to 10^{11} GeV". In: *PoS* ICRC2017 (2018), p. 533. DOI: 10.22323/1.301.0533. arXiv: 1711.11432 [astro-ph.HE].

[132] S. Bravo. *IceCube unveils the world's most precise search for mysterious new neutrino interactions.* en-US. Jan. 14, 2022. URL: https://icecube.wisc.edu/news/research/2022/01/icecube-unveils-the-worlds-most-precise-search-for-mysterious-new-neutrino-interactions/ (visited on 10/21/2022).

[133] M. Honda et al. "Calculation of atmospheric neutrino flux using the interaction model calibrated with atmospheric muon data". In: *Phys. Rev. D* 75 (2007), p. 043006. DOI: 10.1103/PhysRevD.75.043006. arXiv: astro-ph/0611418.

[134] M. Honda et al. "Improvement of low energy atmospheric neutrino flux calculation using the JAM nuclear interaction model". In: *Phys. Rev. D* 83 (2011), p. 123001. DOI: 10.1103/PhysRevD.83.123001. arXiv: 1102.2688 [astro-ph.HE].

[135] Y. Nara. "JAM: an event generator for high energy nuclear collisions". In: *European Physical Journal Web of Conferences*. Vol. 208. European Physical Journal Web of Conferences. May 2019, 11004, p. 11004. DOI: 10.1051/epjconf/201920811004.

[136] S. Roesler, R. Engel, and J. Ranft. "The Monte Carlo event generator DPMJET-III". In: *International Conference on Advanced Monte Carlo for Radiation Physics, Particle Transport Simulation and Applications (MC 2000)*. Dec. 2000, pp. 1033–1038. DOI: 10.1007/978-3-642-18211-2_166. arXiv: hep-ph/0012252.

[137] 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].

[138] A. Fedynitch. "Cascade equations and hadronic interactions at very high energies". PhD thesis. KIT, Karlsruhe, Dept. Phys., Nov. 2015. DOI: 10.5445/IR/1000055433.

[139] J. M. Picone et al. "NRLMSISE-00 empirical model of the atmosphere: Statistical comparisons and scientific issues". In: *Journal of Geophysical Research (Space Physics)* 107.A12, 1468 (2002), p. 1468. DOI: 10.1029/2002JA009430.

[140] 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.

[141] A. Fedynitch and M. Huber. "Data-driven hadronic interaction model for atmospheric lepton flux calculations". In: *Phys. Rev. D* 106.8 (2022), p. 083018. DOI: 10.1103/PhysRevD.106.083018. arXiv: 2205.14766 [astro-ph.HE].

[142] J. P. Yañez and A. Fedynitch. "Data-driven muon-calibrated neutrino flux". In: *Phys. Rev. D* 107.12 (2023), p. 123037. DOI: 10.1103/PhysRevD.107.123037. arXiv: 2303.00022 [hep-ph].

[143] F. Riehn et al. "Hadronic interaction model Sibyll 2.3d and extensive air showers". In: *Phys. Rev. D* 102.6 (2020), p. 063002. DOI: 10.1103/PhysRevD.102.063002. arXiv: 1912.03300 [hep-ph].

[144] R. Abbasi et al. "A search for an eV-scale sterile neutrino using improved high-energy ν_μ event reconstruction in IceCube". In: (May 2024). arXiv: 2405.08070 [hep-ex].

[145] A. M. Dziewonski and D. L. Anderson. "Preliminary reference earth model". In: *Phys. Earth Planet. Interiors* 25 (1981), pp. 297–356. DOI: 10.1016/0031-9201(81)90046-7.

[146] S. Wren. "Neutrino Mass Ordering Studies With IceCube-Deepcore". PhD thesis. University of Manchester, 2018.

[147] S. P. Mikheyev and A. Y. Smirnov. "Resonance Amplification of Oscillations in Matter and Spectroscopy of Solar Neutrinos". In: *Sov. J. Nucl. Phys.* 42 (1985), pp. 913–917.

[148] L. Wolfenstein. "Neutrino Oscillations in Matter". In: *Phys. Rev. D* 17 (1978), pp. 2369–2374. DOI: 10.1103/PhysRevD.17.2369.

[149] E. K. Akhmedov. "Parametric resonance of neutrino oscillations and passage of solar and atmospheric neutrinos through the earth". In: *Nucl. Phys. B* 538 (1999), pp. 25–51. DOI: 10.1016/S0550-3213(98)00723-8. arXiv: hep-ph/9805272.

[150] K. J. Kelly et al. "DUNE atmospheric neutrinos: Earth tomography". In: *JHEP* 05 (2022), p. 187. DOI: 10.1007/JHEP05(2022)187. arXiv: 2110.00003 [hep-ph].

[151] 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].

[152] M. Sajjad Athar and S. K. Singh. "Neutrinos and their interactions". In: *Eur. Phys. J. ST* 230.24 (2021), pp. 4215–4220. DOI: 10.1140/epjs/s11734-021-00302-x. arXiv: 2111.12328 [hep-ph].

[153] Y. S. Jeong and M. H. Reno. "Neutrino cross sections: Interface of shallow- and deep-inelastic scattering for collider neutrinos". In: *Phys. Rev. D* 108.11 (2023), p. 113010. DOI: 10.1103/PhysRevD.108.113010. arXiv: 2307.09241 [hep-ph].

[154] 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.

[155] A. Rubbia. *Phenomenology of Particle Physics*. Cambridge, UK: Cambridge University Press, June 2022.

[156] A. Bodek et al. "Vector and Axial Nucleon Form Factors:A Duality Constrained Parameterization". In: *Eur. Phys. J. C* 53 (2008), pp. 349–354. DOI: 10.1140/epjc/s10052-007-0491-4. arXiv: 0708.1946 [hep-ex].

[157] J. Tena-Vidal et al. "Neutrino-nucleon cross-section model tuning in GENIE v3". In: *Phys. Rev. D* 104.7 (2021), p. 072009. DOI: 10.1103/PhysRevD.104.072009. arXiv: 2104.09179 [hep-ph].

[158] 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].

[159] T. J. Leitner. "Neutrino Interactions with Nucleons and Nuclei". Available at https://gibuu.hepforge.org/trac/chrome/site/files/dipl/leitner.pdf. Master's thesis. Giessen: Institute for Theoretical Physics, 2005.

[160] H. S. Budd, A. Bodek, and J. Arrington. "Modeling quasielastic form-factors for electron and neutrino scattering". In: *2nd International Workshop on Neutrino-Nucleus Interactions in the Few GeV Region*. Aug. 2003. arXiv: hep-ex/0308005.

[161] D. I. Scully. "Neutrino induced coherent pion production." PhD thesis. Warwick U., 2013.

[162] S. J. Barish et al. "Study of Neutrino Interactions in Hydrogen and Deuterium. 1. Description of the Experiment and Study of the Reaction Neutrino d –> mu- p p(s)". In: *Phys. Rev. D* 16 (1977), p. 3103. DOI: 10.1103/PhysRevD.16.3103.

[163] N. J. Baker et al. "Quasielastic Neutrino Scattering: A Measurement of the Weak Nucleon Axial Vector Form-Factor". In: *Phys. Rev. D* 23 (1981), pp. 2499–2505. DOI: 10.1103/PhysRevD.23.2499.

[164] T. Kitagaki et al. "High-Energy Quasielastic Muon-neutrino n —> mu- p Scattering in Deuterium". In: *Phys. Rev. D* 28 (1983), pp. 436–442. DOI: 10.1103/PhysRevD.28.436.

[165] D. Allasia et al. "Investigation of exclusive channels in neutrino / anti-neutrino deuteron charged current interactions". In: *Nucl. Phys. B* 343 (1990), pp. 285–309. DOI: 10.1016/0550-3213(90)90472-P.

[166] S. Bonetti et al. "Study of Quasielastic Reactions of Neutrino and anti-neutrino in Gargamelle". In: *Nuovo Cim. A* 38 (1977), pp. 260–270. DOI: 10.1007/BF02730023.

[167] S. V. Belikov et al. "Quasielastic Neutrino and Anti-neutrinos Scattering: Total Cross-sections, Axial Vector Form-factor". In: *Z. Phys. A* 320 (1985), p. 625. DOI: 10.1007/BF01411863.

[168] J. Brunner et al. "Quasielastic Nucleon and Hyperon Production by Neutrinos and Anti-neutrinos With Energies Below 30-GeV". In: *Z. Phys. C* 45 (1990), p. 551. DOI: 10.1007/BF01556267.

[169] 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.

[170] S. J. Barish et al. "Study of Neutrino Interactions in Hydrogen and Deuterium: Inelastic Charged Current Reactions". In: *Phys. Rev. D* 19 (1979), p. 2521. DOI: 10.1103/PhysRevD.19.2521.

[171] G. M. Radecky et al. "Study of Single Pion Production by Weak Charged Currents in Low-energy Neutrino *d* Interactions". In: *Phys. Rev. D* 25 (1982). [Erratum: Phys.Rev.D 26, 3297 (1982)], pp. 1161–1173. DOI: 10.1103/PhysRevD.25.1161.

[172] K. S. Kuzmin, V. V. Lyubushkin, and V. A. Naumov. "Axial masses in quasielastic neutrino scattering and single-pion neutrinoproduction on nucleons and nuclei". In: *Acta Phys. Polon. B* 37 (2006). Ed. by K. M. Graczyk and J. T. Sobczyk, pp. 2337–2348. arXiv: hep-ph/0606184.

[173] W. G. Seligman. "A Next-to-Leading Order QCD Analysis of Neutrino - Iron Structure Functions at the Tevatron". PhD thesis. Nevis Labs, Columbia U., 1997. DOI: 10.2172/1421736.

[174] M. Tzanov et al. "Precise measurement of neutrino and anti-neutrino differential cross sections". In: *Phys. Rev. D* 74 (2006), p. 012008. DOI: 10.1103/PhysRevD.74.012008. arXiv: hep-ex/0509010.

[175] Q. Wu et al. "A Precise measurement of the muon neutrino-nucleon inclusive charged current cross-section off an isoscalar target in the energy range 2.5 < E(nu) < 40-GeV by NOMAD". In: *Phys. Lett. B* 660 (2008), pp. 19–25. DOI: 10.1016/j.physletb.2007.12.027. arXiv: 0711.1183 [hep-ex].

[176] P. Adamson et al. "Neutrino and Antineutrino Inclusive Charged-current Cross Section Measurements with the MINOS Near Detector". In: *Phys. Rev. D* 81 (2010), p. 072002. DOI: 10.1103/PhysRevD.81.072002. arXiv: 0910.2201 [hep-ex].

[177] A. Bodek and U. K. Yang. "Higher twist, xi(omega) scaling, and effective LO PDFs for lepton scattering in the few GeV region". In: *J. Phys. G* 29 (2003). Ed. by K. Long and R. Edgecock, pp. 1899–1906. DOI: 10.1088/0954-3899/29/8/369. arXiv: hep-ex/0210024.

[178] M. Glück, E. Reya, and A. Vogt. "Dynamical parton distributions revisited". In: *Eur. Phys. J. C* 5 (1998), pp. 461–470. DOI: 10.1007/s100520050289. arXiv: hep-ph/9806404.

[179] A. Cooper-Sarkar, P. Mertsch, and S. Sarkar. "The high energy neutrino cross-section in the Standard Model and its uncertainty". In: *JHEP* 08 (2011), p. 042. DOI: 10.1007/JHEP08(2011)042. arXiv: 1106.3723 [hep-ph].

[180] R. Abbasi et al. "Measurement of atmospheric neutrino mixing with improved IceCube Deep-Core calibration and data processing". In: *Phys. Rev. D* 108.1 (2023), p. 012014. DOI: 10.1103/PhysRevD.108.012014. arXiv: 2304.12236 [hep-ex].

[181] A. M. Cooper-Sarkar. "PDF Fits at HERA". In: *PoS* EPS-HEP2011 (2011), p. 320. DOI: 10.22323/1.134.0320. arXiv: 1112.2107 [hep-ph].

[182] S. Adrian-Martinez et al. "Letter of intent for KM3NeT 2.0". In: *J. Phys. G* 43.8 (2016), p. 084001. DOI: 10.1088/0954-3899/43/8/084001. arXiv: 1601.07459 [astro-ph.IM].

[183] Y. Fukuda et al. "The Super-Kamiokande detector". In: *Nucl. Instrum. Meth. A* 501 (2003). Ed. by V. A. Ilyin, V. V. Korenkov, and D. Perret-Gallix, pp. 418–462. DOI: 10.1016/S0168-9002(03)00425-X.

[184] A. Abusleme et al. "JUNO sensitivity to low energy atmospheric neutrino spectra". In: *Eur. Phys. J. C* 81 (2021), p. 10. DOI: 10.1140/epjc/s10052-021-09565-z. arXiv: 2103.09908 [hep-ex].

[185] A. Abusleme et al. "JUNO physics and detector". In: *Prog. Part. Nucl. Phys.* 123 (2022), p. 103927. DOI: 10.1016/j.ppnp.2021.103927. arXiv: 2104.02565 [hep-ex].

[186] B. Abi et al. "Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics". In: (Feb. 2020). arXiv: 2002.03005 [hep-ex].

[187] K. J. Kelly et al. "Sub-GeV Atmospheric Neutrinos and CP-Violation in DUNE". In: *Phys. Rev. Lett.* 123.8 (2019), p. 081801. DOI: 10.1103/PhysRevLett.123.081801. arXiv: 1904.02751 [hep-ph].

[188] J. V. Jelley. "Cerenkov radiation and its applications". In: *British Journal of Applied Physics* 6.7 (1955), p. 227. DOI: 10.1088/0508-3443/6/7/301.

[189] I. M. Frank and I. E. Tamm. "Coherent visible radiation of fast electrons passing through matter". In: *Compt. Rend. Acad. Sci. URSS* 14.3 (1937). Ed. by V. L. Ginzburg, B. M. Bolotovsky, and I. M. Dremin, pp. 109–114. DOI: 10.3367/UFNr.0093.196710o.0388.

[190] M. S. Longair. *High Energy Astrophysics*. 3rd ed. Cambridge University Press, 2011.

[191] J. D. Jackson. *Classical Electrodynamics*. Wiley, 1998. ISBN: 978-0-471-30932-1.

[192] S. G. Warren. "Optical constants of ice from the ultraviolet to the microwave". In: *Appl. Opt.* 23.8 (1984), pp. 1206–1225. DOI: 10.1364/AO.23.001206.

[193] S. G. Warren, R. E. Brandt, and T. C. Grenfell. "Visible and near-ultraviolet absorption spectrum of ice from transmission of solar radiation into snow". In: *Appl. Opt.* 45.21 (2006), pp. 5320–5334. DOI: 10.1364/AO.45.005320.

[194] M. G. Cooper et al. "Spectral attenuation coefficients from measurements of light transmission in bare ice on the Greenland Ice Sheet". In: *The Cryosphere* 15.4 (2021), pp. 1931–1953. DOI: 10.5194/tc-15-1931-2021.

[195] D. K. Perovich and J. W. Govoni. "Absorption coefficients of ice from 250 to 400 nm". In: *Geophysical Research Letters* 18.7 (1991), pp. 1233–1235. DOI: 10.1029/91GL01642.

[196] 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].

[197] R. Abbasi et al. "Measurement of atmospheric neutrino oscillation parameters using convolutional neural networks with 9.3 years of data in IceCube DeepCore". In: (May 2024). arXiv: 2405.02163 [hep-ex].

[198] R. Abbasi et al. "All-flavor constraints on nonstandard neutrino interactions and generalized 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].

[199] T. Stuttard and M. Jensen. "Neutrino decoherence from quantum gravitational stochastic perturbations". In: *Phys. Rev. D* 102.11 (2020), p. 115003. DOI: 10.1103/PhysRevD.102.115003. arXiv: 2007.00068 [hep-ph].

[200] R. Abbasi et al. "Search for sub-TeV Neutrino Emission from Novae with IceCube-DeepCore". In: *Astrophys. J.* 953.2 (2023). [Erratum: Astrophys.J. 971, (2024)], p. 160. DOI: 10.3847/1538-4357/acdc1b. arXiv: 2212.06810 [astro-ph.HE].

[201] R. Abbasi et al. "First all-flavor search for transient neutrino emission using 3-years of IceCube DeepCore data". In: *JCAP* 01 (2022), p. 027. DOI: 10.1088/1475-7516/2022/01/027. arXiv: 2011.05096 [astro-ph.HE].

[202] M. G. Aartsen et al. "All-flavour Search for Neutrinos from Dark Matter Annihilations in the Milky Way with IceCube/DeepCore". In: *Eur. Phys. J. C* 76.10 (2016), p. 531. DOI: 10.1140/epjc/s10052-016-4375-3. arXiv: 1606.00209 [astro-ph.HE].

[203] R. Abbasi et al. "Search for GeV-scale dark matter annihilation in the Sun with IceCube DeepCore". In: *Phys. Rev. D* 105.6 (2022), p. 062004. DOI: 10.1103/PhysRevD.105.062004. arXiv: 2111.09970 [astro-ph.HE].

[204] R. Abbasi et al. "Search for neutrino lines from dark matter annihilation and decay with IceCube". In: *Phys. Rev. D* 108.10 (2023), p. 102004. DOI: 10.1103/PhysRevD.108.102004. arXiv: 2303.13663 [astro-ph.HE].

[205] A. Ishihara. "The IceCube Upgrade - Design and Science Goals". In: *PoS* ICRC2019 (2021), p. 1031. DOI: 10.22323/1.358.1031. arXiv: 1908.09441 [astro-ph.HE].

[206] A. Einstein. "Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt (On a Heuristic Point of View about the Creation and Conversion of Light")". In: *Annalen der Physik* 322.6 (1905), pp. 132–148. DOI: https://doi.org/10.1002/andp.19053220607.

[207] *Photomultiplier Tube (PMT): Matsusada Precision Inc.* 2023. URL: https://www.matsusada.com/application/ps/photomultiplier_tubes/ (visited on 04/10/2024).

[208] *Photomultiplier Tubes: Basics and Applications (technical manual)*. Hamamatsu Photonics, 2017. URL: https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/etd/PMT_handbook_v4E.pdf (visited on 04/10/2024).

[209] M. G. Aartsen et al. "Measurement of South Pole ice transparency with the IceCube LED calibration system". In: *Nucl. Instrum. Meth. A* 711 (2013), pp. 73–89. DOI: 10.1016/j.nima.2013.01.054. arXiv: 1301.5361 [astro-ph.IM].

[210] R. Abbasi et al. "Graph Neural Networks for low-energy event classification & reconstruction in IceCube". In: *JINST* 17.11 (2022), P11003. DOI: 10.1088/1748-0221/17/11/P11003. arXiv: 2209.03042 [hep-ex].

[211] R. Abbasi et al. "Low energy event reconstruction in IceCube DeepCore". In: *Eur. Phys. J. C* 82.9 (2022), p. 807. DOI: 10.1140/epjc/s10052-022-10721-2. arXiv: 2203.02303 [hep-ex].

[212] 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].

[213] *IceCube Upgrade Neutrino Monte Carlo Simulation.* DOI:10.21234/qfz1-yh02. 2020. URL: https://doi.org/10.21234/qfz1-yh02.

[214] 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].

[215] 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].

[216] C. L. Lawson and R. J. Hanson. "23. Linear Least Squares with Linear Inequality Constraints". In: *Solving Least Squares Problems.* 1995, pp. 158–173. DOI: 10.1137/1.9781611971217.ch23. eprint: https://epubs.siam.org/doi/pdf/10.1137/1.9781611971217.ch23.

[217] A. Gazizov and M. P. Kowalski. "ANIS: High energy neutrino generator for neutrino telescopes". In: *Comput. Phys. Commun.* 172 (2005), pp. 203–213. DOI: 10.1016/j.cpc.2005.03.113. arXiv: astro-ph/0406439.

[218] J.-H. Koehne et al. "PROPOSAL: A tool for propagation of charged leptons". In: *Computer Physics Communications* 184.9 (2013), pp. 2070–2090. ISSN: 0010-4655. DOI: 10.1016/j.cpc.2013.04.001.

[219] J. van Santen. "Neutrino Interactions in IceCube above 1 TeV: Constraints on Atmospheric Charmed-Meson Production and Investigation of the Astrophysical Neutrino Flux with 2 Years of IceCube Data taken 2010–2012". PhD thesis. Wisconsin U., Madison, Nov. 2014.

[220] T. K. Gaisser, T. Stanev, and S. Tilav. "Cosmic Ray Energy Spectrum from Measurements of Air Showers". In: *Front. Phys. (Beijing)* 8 (2013), pp. 748–758. DOI: 10.1007/s11467-013-0319-7. arXiv: 1303.3565 [astro-ph.HE].

[221] L. Radel and C. Wiebusch. "Calculation of the Cherenkov light yield from low energetic secondary particles accompanying high-energy muons in ice and water with Geant 4 simulations". In: *Astropart. Phys.* 38 (2012), pp. 53–67. DOI: 10.1016/j.astropartphys.2012.09.008. arXiv: 1206.5530 [astro-ph.IM].

[222] D. Chirkin. "Photon tracking with GPUs in IceCube". In: *Nucl. Instrum. Meth. A* 725 (2013). Ed. by G. Anton et al., pp. 141–143. DOI: 10.1016/j.nima.2012.11.170.

[223] D. A. Chirkin. *icecube/ppc: Version 1.0.* Version V1.0. Dec. 2023. DOI: 10.5281/zenodo.10410726. URL: https://doi.org/10.5281/zenodo.10410726.

[224] 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.

[225] M. Kowalski. *On the Cherenkov light emission of hadronic and electromagnetic cascades.* Tech. rep. DESY-Zeuthen, 2002.

[226] L. Radel 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].

[227] P. Eller et al. "A model independent parametrization of the optical properties of the refrozen IceCube drill holes". In: *PoS* ICRC2023 (2023), p. 1034. DOI: 10.22323/1.444.1034. arXiv: 2307.15298 [astro-ph.HE].

[228] C. Weaver. "Evidence for Astrophysical Muon Neutrinos from the Northern Sky". PhD thesis. University of Wisconsin–Madison, 2015.

[229] M. Larson. "Simulation and Identification of Non-Poissonian Noise Triggers in the IceCube Neutrino Detector". Available at `https://docushare.icecube.wisc.edu/dsweb/Get/Document-68303/LarsonThesis_final.pdf`. Master's thesis. Tuscaloosa, AL: University of Alabama, 2013.

[230] M. Medici. "Search for Dark Matter Annihilation in the Galactic Halo using IceCube". PhD thesis. Niels Bohr Institute, University of Copenhagen, 2016.

[231] D. Heck et al. "CORSIKA: A Monte Carlo code to simulate extensive air showers". In: (Feb. 1998).

[232] S. Flis. "Extending IceCube Low Energy Neutrino Searches for Dark Matter with DeepCore". In: *33rd International Cosmic Ray Conference*. 2013, p. 0565.

[233] J. H. Friedman. "Greedy function approximation: A gradient boosting machine." In: *The Annals of Statistics* 29.5 (2001), pp. 1189–1232. DOI: `10.1214/aos/1013203451`.

[234] J. Ahrens et al. "Muon track reconstruction and data selection techniques in AMANDA". In: *Nucl. Instrum. Meth. A* 524 (2004), pp. 169–194. DOI: `10.1016/j.nima.2004.01.065`. arXiv: `astro-ph/0407044`.

[235] M. Leuermann. "Testing the Neutrino Mass Ordering with IceCube DeepCore". PhD thesis. RWTH Aachen U., 2018. DOI: `10.18154/RWTH-2018-231554`.

[236] K. Leonard DeHolton. "Measuring the Atmospheric Neutrino Oscillation Parameters with Icecube-Deepcore". PhD thesis. U. Wisconsin, Madison (main), University of Wisconsin Madison, 2022.

[237] J. P. Yáñez Garza. "Measurement of neutrino oscillations in atmospheric neutrinos with the IceCube DeepCore detector". PhD thesis. Humboldt U., Berlin, 2014.

[238] G. Ke et al. "LightGBM: A Highly Efficient Gradient Boosting Decision Tree". In: *Advances in Neural Information Processing Systems*. Ed. by I. Guyon et al. Vol. 30. Curran Associates, Inc., 2017.

[239] S. M. Lundberg and S.-I. Lee. "A unified approach to interpreting model predictions". In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. NIPS'17. Long Beach, California, USA: Curran Associates Inc., 2017, pp. 4768–4777. ISBN: 9781510860964. DOI: `10.5555/3295222.3295230`.

[240] 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`.

[241] F. Feroz, M. P. Hobson, and M. Bridges. "MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics". In: *Monthly Notices of the Royal Astronomical Society* 398.4 (Sept. 2009), pp. 1601–1614. ISSN: 0035-8711. DOI: `10.1111/j.1365-2966.2009.14548.x`.

[242] 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]`.

[243] P. Kaelo and M. M. Ali. "Some Variants of the Controlled Random Search Algorithm for Global Optimization". In: *Journal of Optimization Theory and Applications* 130.2 (2006), pp. 253–264. DOI: `10.1007/s10957-006-9101-0`.

[244] A. Søgaard et al. *GraphNeT: Graph neural networks for neutrino telescope event reconstruction*. Version v1.0.0. May 2023. DOI: `10.5281/zenodo.7928487`.

[245] M. Cheung et al. "Graph Signal Processing and Deep Learning: Convolution, Pooling, and Topology". In: *IEEE Signal Processing Magazine* 37.6 (2020), pp. 139–149. DOI: `10.1109/MSP.2020.3014594`.

[246] A. de Gouvêa and A. Kobach. "Global Constraints on a Heavy Neutrino". In: *Phys. Rev. D* 93.3 (2016), p. 033005. DOI: `10.1103/PhysRevD.93.033005`. arXiv: `1511.00683 [hep-ph]`.

[247] S. Antusch and O. Fischer. "Non-Unitarity of the Leptonic Mixing Matrix: Present Bounds and Future Sensitivities". In: *JHEP* 10 (2014), p. 094. DOI: `10.1007/JHEP10(2014)094`. arXiv: `1407.6607 [hep-ph]`.

[248] D. V. Forero et al. "Nonunitary Neutrino Mixing in Short and Long-Baseline Experiments". In: *Phys. Rev. D* 104.7 (2021), p. 075030. DOI: `10.1103/PhysRevD.104.075030`. arXiv: `2103.01998 [hep-ph]`.

[249] S. A. R. Ellis, K. J. Kelly, and S. W. Li. "Leptonic Unitarity Triangles". In: *Phys. Rev. D* 102.11 (2020), p. 115027. DOI: `10.1103/PhysRevD.102.115027`. arXiv: `2004.13719 [hep-ph]`.

[250] S. K. Agarwalla et al. "Model-Independent Constraints on Non-Unitary Neutrino Mixing from High-Precision Long-Baseline Experiments". In: *JHEP* 07 (2022), p. 121. DOI: 10.1007/JHEP07(2022)121. arXiv: 2111.00329 [hep-ph].

[251] S. Sahoo et al. "Constraining Non-unitary Neutrino Mixing Using Matter Effects in Atmospheric Neutrinos at INO-ICAL". In: (Sept. 2023). arXiv: 2309.16942 [hep-ph].

[252] R. Wendell et al. "Atmospheric Neutrino Oscillation Analysis with Sub-leading Effects in Super-Kamiokande I, II, and III". In: *Phys. Rev. D* 81 (2010), p. 092004. DOI: 10.1103/PhysRevD.81.092004. arXiv: 1002.3471 [hep-ex].

[253] K. Abe et al. "Limits on sterile neutrino mixing using atmospheric neutrinos in Super-Kamiokande". In: *Phys. Rev. D* 91 (2015), p. 052019. DOI: 10.1103/PhysRevD.91.052019. arXiv: 1410.2008 [hep-ex].

[254] *Data Release: Atmospheric Neutrino Oscillation Analysis with Neutron Tagging and an Expanded Fiducial Volume in Super-Kamiokande I-V*. DOI:10.5281/zenodo.8401262. 2023. URL: https://doi.org/10.21234/ac23-ra43.

[255] A. Abusleme et al. "Sub-percent Precision Measurement of Neutrino Oscillation Parameters with JUNO". In: *Chin. Phys. C* 46.12 (2022), p. 123001. DOI: 10.1088/1674-1137/ac8bc9. arXiv: 2204.13249 [hep-ex].

[256] 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].

[257] T. Sanuki et al. "Study of Cosmic Ray Interaction Model Based on Atmospheric Muons for the Neutrino Flux Calculation". In: *Phys. Rev. D* 75 (2007), p. 043005. DOI: 10.1103/PhysRevD.75.043005. arXiv: astro-ph/0611201.

[258] D. E. Groom, N. V. Mokhov, and S. I. Striganov. "Muon Stopping Power and Range Tables 10-MeV to 100-TeV". In: *Atom. Data Nucl. Data Tabl.* 78 (2001), pp. 183–356. DOI: 10.1006/adnd.2001.0861.

[259] S. S. Wilks. "The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses". In: *Annals Math. Statist.* 9.1 (1938), pp. 60–62. DOI: 10.1214/aoms/1177732360.

[260] F. P. An et al. "Improved Measurement of the Reactor Antineutrino Flux and Spectrum at Daya Bay". In: *Chin. Phys. C* 41.1 (2017), p. 013002. DOI: 10.1088/1674-1137/41/1/013002. arXiv: 1607.05378 [hep-ex].

[261] D. Adey et al. "Measurement of the Electron Antineutrino Oscillation with 1958 Days of Operation at Daya Bay". In: *Phys. Rev. Lett.* 121.24 (2018), p. 241805. DOI: 10.1103/PhysRevLett.121.241805. arXiv: 1809.02261 [hep-ex].

[262] F. P. An et al. "Antineutrino Energy Spectrum Unfolding Based on the Daya Bay Measurement and Its Applications". In: *Chin. Phys. C* 45.7 (2021), p. 073001. DOI: 10.1088/1674-1137/abfc38. arXiv: 2102.04614 [hep-ex].

[263] F. P. An et al. "Measurement of Electron Antineutrino Oscillation Based on 1230 Days of Operation of the Daya Bay Experiment". In: *Phys. Rev. D* 95.7 (2017), p. 072006. DOI: 10.1103/PhysRevD.95.072006. arXiv: 1610.04802 [hep-ex].

[264] P. Vogel and J. F. Beacom. "Angular Distribution of Neutron Inverse Beta Decay, Anti-Neutrino(e) + p —> e+ + n". In: *Phys. Rev. D* 60 (1999), p. 053003. DOI: 10.1103/PhysRevD.60.053003. arXiv: hep-ph/9903554.

[265] *Data Release Accompanying the 4th KamLAND Reactor Result*. 2012. URL: https://www.awa.tohoku.ac.jp/KamLAND/4th_result_data_release/4th_result_data_release.html.

[266] S. Enomoto. "Neutrino Geophysics and Observation of Geo-Neutrinos at KamLAND". PhD thesis. Tohoku University, 2005. URL: http://kamland.stanford.edu/GeoNeutrinos/GeoNuResult/SanshirosDoctoralDissertation.pdf.

[267] S. Abe et al. "Precision Measurement of Neutrino Oscillation Parameters with KamLAND". In: *Phys. Rev. Lett.* 100 (2008), p. 221803. DOI: 10.1103/PhysRevLett.100.221803. arXiv: 0801.4589 [hep-ex].

[268] J. A. Detwiler. "Measurement of Neutrino Oscillation with KamLAND". PhD thesis. Stanford University, 2005. URL: http://faculty.washington.edu/jasondet/DetwilerThesis.pdf.

[269]   B. Aharmim et al. "Combined Analysis of all Three Phases of Solar Neutrino Data from the Sudbury Neutrino Observatory". In: *Phys. Rev. C* 88 (2013), p. 025501. DOI: 10.1103/PhysRevC.88.025501. arXiv: 1109.0763 [nucl-ex].

[270]   C. Guo. "JUNO Detector Design and Status". In: *PoS* TAUP2023 (2024), p. 178. DOI: 10.22323/1.441.0178.

[271]   A. Abusleme et al. "TAO Conceptual Design Report: A Precision Measurement of the Reactor Antineutrino Spectrum with Sub-percent Energy Resolution". In: (May 2020). arXiv: 2005.08745 [physics.ins-det].

[272]   J. Buchner. "UltraNest – A Robust, General Purpose Bayesian Inference Engine". In: (Jan. 2021). arXiv: 2101.09604 [stat.CO].

[273]   J. Buchner. *UltraNest Documentation*. https://johannesbuchner.github.io/UltraNest/. 2024. URL: https://johannesbuchner.github.io/UltraNest/.

[274]   J. Buchner. "A Statistical Test for Nested Sampling Algorithms". In: *Statistics and Computing* 26.1–2 (Sept. 2014), pp. 383–392. ISSN: 1573-1375. DOI: 10.1007/s11222-014-9512-y. URL: http://dx.doi.org/10.1007/s11222-014-9512-y.

[275]   J. Buchner. "Collaborative Nested Sampling: Big Data versus Complex Physical Models". In: *Publications of the Astronomical Society of the Pacific* 131.1004 (Aug. 2019), p. 108005. ISSN: 1538-3873. DOI: 10.1088/1538-3873/aae7fc. URL: http://dx.doi.org/10.1088/1538-3873/aae7fc.

[276]   H. Jeffreys. *Theory of Probability*. Oxford, England: Clarendon Press, 1939.

[277]   R. Trotta. "Bayes in the Sky: Bayesian Inference and Model Selection in Cosmology". In: *Contemp. Phys.* 49 (2008), pp. 71–104. DOI: 10.1080/00107510802066753. arXiv: 0803.4089 [astro-ph].

[278]   A. Gelman, X.-L. Meng, and H. S. Stern. "Posterior Predictive Assessment of Model Fitness via Realized Discrepancies". In: 1996. URL: https://api.semanticscholar.org/CorpusID:17309196.

[279]   M. Rongen. "Calibration of the IceCube Neutrino Observatory". PhD thesis. RWTH Aachen U., 2019. DOI: 10.18154/RWTH-2019-09941. arXiv: 1911.02016 [astro-ph.IM].

[280]   D. E. Groom and S. R. Klein. "Passage of particles through matter". en. In: *The European Physical Journal C* 15.1-4 (Mar. 2000), pp. 163–173. ISSN: 1434-6044, 1434-6052. DOI: 10.1007/BF02683419. (Visited on 09/26/2022).

[281]   E. Segrè et al. *Experimental nuclear physics. volume I*. eng. OCLC: 494335476. New York, London: John Wiley & sons ; Chapman & Hall, 1953.

[282]   M. Berger, J. Coursey, and M. Zucker. *ESTAR, PSTAR, and ASTAR: Computer Programs for Calculating Stopping-Power and Range Tables for Electrons, Protons, and Helium Ions (version 1.21)*. 1999. DOI: 10.18434/T4NC7P.

[283]   D. Chirkin. "Evidence of optical anisotropy of the South Pole ice". In: *33rd International Cosmic Ray Conference*. 2013, p. 0580.

[284]   K. J. Ma et al. "Time and Amplitude of Afterpulse Measured with a Large Size Photomultiplier Tube". In: *Nucl. Instrum. Meth. A* 629 (2011), pp. 93–100. DOI: 10.1016/j.nima.2010.11.095. arXiv: 0911.5336 [physics.ins-det].

[285]   J. Conrad et al. "Atmospheric Tau Neutrinos in a Multi-kiloton Liquid Argon Detector". In: *Phys. Rev. D* 82 (2010), p. 093012. DOI: 10.1103/PhysRevD.82.093012. arXiv: 1008.2984 [hep-ph].

[286]   S. Kretzer and M. H. Reno. "Tau neutrino deep inelastic charged current interactions". In: *Phys. Rev. D* 66 (2002), p. 113007. DOI: 10.1103/PhysRevD.66.113007. arXiv: hep-ph/0208187.

[287]   C. H. Albright and C. Jarlskog. "Neutrino Production of m+ and e+ Heavy Leptons. 1." In: *Nucl. Phys. B* 84 (1975), pp. 467–492. DOI: 10.1016/0550-3213(75)90318-1.

[288]   Y. S. Jeong and M. H. Reno. "Tau neutrino and antineutrino cross sections". In: *Phys. Rev. D* 82 (2010), p. 033010. DOI: 10.1103/PhysRevD.82.033010. arXiv: 1007.1966 [hep-ph].

[289]   V. Ansari et al. "Nonperturbative and higher order perturbative effects in deep inelastic $\nu_\tau/\bar{\nu}_\tau$−nucleon scattering". In: *Phys. Rev. D* 102.11 (2020), p. 113007. DOI: 10.1103/PhysRevD.102.113007. arXiv: 2010.05538 [hep-ph].

[290]   S. Aiello et al. "Determining the neutrino mass ordering and oscillation parameters with KM3NeT/ORCA". In: *Eur. Phys. J. C* 82.1 (2022), p. 26. DOI: 10.1140/epjc/s10052-021-09893-0. arXiv: 2103.09885 [hep-ex].

[291] B. Abi et al. "Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume I Introduction to DUNE". In: *JINST* 15.08 (2020), T08008. DOI: 10.1088/1748-0221/15/08/T08008. arXiv: 2002.02967 [physics.ins-det].

[292] A. De Gouvêa et al. "Physics with Beam Tau-Neutrino Appearance at DUNE". In: *Phys. Rev. D* 100.1 (2019), p. 016004. DOI: 10.1103/PhysRevD.100.016004. arXiv: 1904.07265 [hep-ph].

[293] A. Candido et al. "Neutrino Structure Functions from GeV to EeV Energies". In: *JHEP* 05 (2023), p. 149. DOI: 10.1007/JHEP05(2023)149. arXiv: 2302.08527 [hep-ph].

[294] A. Deur, S. J. Brodsky, and G. F. de Teramond. "The QCD Running Coupling". In: *Nucl. Phys.* 90 (2016), p. 1. DOI: 10.1016/j.ppnp.2016.04.003. arXiv: 1604.08082 [hep-ph].

[295] V. Ansari et al. "Deep inelastic (anti)neutrino–nucleus scattering". In: *Eur. Phys. J. ST* 230.24 (2021), pp. 4433–4448. DOI: 10.1140/epjs/s11734-021-00277-9. arXiv: 2106.14670 [hep-ph].

[296] 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.

[297] O. Nachtmann. "Positivity constraints for anomalous dimensions". In: *Nucl. Phys. B* 63 (1973), pp. 237–247. DOI: 10.1016/0550-3213(73)90144-2.

[298] M. A. G. Aivazis, F. I. Olness, and W.-K. Tung. "Leptoproduction of heavy quarks. 1. General formalism and kinematics of charged current and neutral current production processes". In: *Phys. Rev. D* 50 (1994), pp. 3085–3101. DOI: 10.1103/PhysRevD.50.3085. arXiv: hep-ph/9312318.

[299] R. K. Ellis, W. Furmanski, and R. Petronzio. "Unraveling Higher Twists". In: *Nucl. Phys. B* 212 (1983), p. 29. DOI: 10.1016/0550-3213(83)90597-7.

[300] S. Kretzer and M. H. Reno. "Target mass corrections to electroweak structure functions and perturbative neutrino cross-sections". In: *Phys. Rev. D* 69 (2004), p. 034002. DOI: 10.1103/PhysRevD.69.034002. arXiv: hep-ph/0307023.

[301] I. Schienbein et al. "A Review of Target Mass Corrections". In: *J. Phys. G* 35 (2008), p. 053101. DOI: 10.1088/0954-3899/35/5/053101. arXiv: 0709.1775 [hep-ph].

[302] R. D. Ball et al. "The path to proton structure at 1% accuracy". In: *Eur. Phys. J. C* 82.5 (2022), p. 428. DOI: 10.1140/epjc/s10052-022-10328-7. arXiv: 2109.02653 [hep-ph].

[303] A. Buckley et al. "LHAPDF6: parton density access in the LHC precision era". In: *Eur. Phys. J. C* 75 (2015), p. 132. DOI: 10.1140/epjc/s10052-015-3318-8. arXiv: 1412.7420 [hep-ph].

[304] S. Carrazza, J. M. Cruz-Martinez, and M. Rossi. "PDFFlow: Parton distribution functions on GPU". In: *Comput. Phys. Commun.* 264 (2021), p. 107995. DOI: 10.1016/j.cpc.2021.107995. arXiv: 2009.06635 [hep-ph].

[305] A. D. Martin et al. "Parton distributions for the LHC". In: *Eur. Phys. J. C* 63 (2009), pp. 189–285. DOI: 10.1140/epjc/s10052-009-1072-5. arXiv: 0901.0002 [hep-ph].

[306] S. Forte et al. "Neural network parametrization of deep inelastic structure functions". In: *JHEP* 05 (2002), p. 062. DOI: 10.1088/1126-6708/2002/05/062. arXiv: hep-ph/0204232.

[307] R. D. Ball et al. "Parton distributions from high-precision collider data". In: *Eur. Phys. J. C* 77.10 (2017), p. 663. DOI: 10.1140/epjc/s10052-017-5199-5. arXiv: 1706.00428 [hep-ph].

[308] D. Day et al. "STUDY OF neutrino D CHARGED CURRENT TWO PION PRODUCTION IN THE THRESHOLD REGION". In: *Phys. Rev. D* 28 (1983), pp. 2714–2720. DOI: 10.1103/PhysRevD.28.2714.

[309] G. Fanourakis et al. "Study of Low-energy Anti-neutrino Interactions on Protons". In: *Phys. Rev. D* 21 (1980), p. 562. DOI: 10.1103/PhysRevD.21.562.

[310] T. Kitagaki et al. "Charged Current Exclusive Pion Production in Neutrino Deuterium Interactions". In: *Phys. Rev. D* 34 (1986), pp. 2554–2565. DOI: 10.1103/PhysRevD.34.2554.

[311] N. J. Baker et al. "Total Cross-sections for Muon-neutrino N and Muon-neutrino P Charged Current Interactions in the 7-ft Bubble Chamber". In: *Phys. Rev. D* 25 (1982), pp. 617–623. DOI: 10.1103/PhysRevD.25.617.

[312] M. A. Parker et al. "A Comparison of Charged Current Cross-sections and Structure Functions for Neutrino and Anti-neutrinos Beams on Hydrogen and Neon". In: *Nucl. Phys. B* 232 (1984), pp. 1–20. DOI: 10.1016/0550-3213(84)90358-4.

[313] D. C. Colley et al. "Cross-sections for Charged Current Neutrino and Anti-neutrino Interactions in the Energy Range 10-GeV to 50-GeV". In: *Z. Phys. C* 2 (1979), p. 187. DOI: 10.1007/BF01474659.

[314] P. C. Bosetti et al. "Total Cross-Sections for Charged Current Neutrino and anti-neutrino Interactions in BEBC in the Energy Range 20-GeV to 200-GeV". In: *Phys. Lett. B* 70 (1977), p. 273. DOI: 10.1016/0370-2693(77)90537-8.

[315] P. Bosetti et al. "Total Cross-sections for ν_μ and $\bar{\nu}_\mu$ Charged Current Interactions Between 20-GeV and 200-GeV". In: *Phys. Lett. B* 110 (1982), pp. 167–172. DOI: 10.1016/0370-2693(82)91028-0.

[316] P. Allen et al. "Single π^+ Production in Charged Current Neutrino - Hydrogen Interactions". In: *Nucl. Phys. B* 176 (1980), p. 269. DOI: 10.1016/0550-3213(80)90450-2.

[317] P. Allen et al. "A Study of Single Meson Production in Neutrino and Anti-neutrinos Charged Current Interactions on Protons". In: *Nucl. Phys. B* 264 (1986), pp. 221–242. DOI: 10.1016/0550-3213(86)90480-3.

[318] J. Bell et al. "Cross-section Measurements for the Reactions $\nu p \to \mu^- \pi^+ p$ and $\nu p \to \mu^- K^+ p$ at High-Energies". In: *Phys. Rev. Lett.* 41 (1978), p. 1008. DOI: 10.1103/PhysRevLett.41.1008.

[319] N. J. Baker et al. "Measurement of the muon-neutrino charged current cross-section". In: *Phys. Rev. Lett.* 51 (1983), pp. 735–738. DOI: 10.1103/PhysRevLett.51.735.

[320] A. E. Asratian et al. "Total antineutrino-nucleon charged current cross section in the energy range 10 GeV-50 GeV". In: *Phys. Lett. B* 137 (1984), pp. 122–124. DOI: 10.1016/0370-2693(84)91118-3.

[321] G. N. Taylor et al. "Anti-muon-neutrino Nucleon Charged Current Total Cross-Section for 5 GeV to 250 GeV". In: *Phys. Rev. Lett.* 51 (1983), pp. 739–742. DOI: 10.1103/PhysRevLett.51.739.

[322] C. Andreopoulos et al. "Summary of the NuSTEC Workshop on Shallow- and Deep-Inelastic Scattering". In: *NuSTEC Workshop on Shallow- and Deep-Inelastic Scattering*. July 2019. arXiv: 1907.13252 [hep-ph].

[323] A. Candido et al. "Yadism: yet another deep-inelastic scattering module". In: *Eur. Phys. J. C* 84.7 (2024), p. 698. DOI: 10.1140/epjc/s10052-024-12972-7. arXiv: 2401.15187 [hep-ph].