



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

DOTTORATO DI RICERCA IN

FISICA

Ciclo 38

Settore Concorsuale: 02/A1 - FISICA SPERIMENTALE DELLE INTERAZIONI
FONDAMENTALI

Settore Scientifico Disciplinare: FIS/01 - FISICA SPERIMENTALE

DISCOVERY POTENTIAL FOR HEAVY NEUTRAL LEPTONS AT THE SAND
DETECTOR OF THE DUNE EXPERIMENT

Presentata da: Alessandro Ruggeri

Coordinatore Dottorato

Alessandro Gabrielli

Supervisore

Nicoletta Mauri

Co-supervisore

Alessandro Montanari

Esame finale anno 2026

Abstract

Beyond its primary goal of determining the remaining neutrino oscillation parameters, the DUNE next-generation long-baseline neutrino experiment will pursue a broad program of precision measurements of neutrino cross-sections and standard model observables, as well as searches for Physics Beyond the Standard Model (BSM). At the DUNE Near Detector complex, the SAND detector will exploit the intense broadband neutrino beam and a high-resolution gas-detector tracker in a magnetic field coupled to a novel 1-ton liquid Argon detector and a high-granularity calorimeter to perform searches of BSM-phenomena. In particular, the high neutrino flux and a low-density tracker are optimal conditions to search for Dark Neutrinos, a class of Heavy Neutral Lepton (HNL) models in which HNLs are produced in the scattering of standard model neutrinos and undergo displaced decays to di-lepton pairs through the exchange of a dark sector gauge boson Z' , or *dark photon*.

This dissertation presents a first estimate of the discovery potential in SAND for dark neutrino models featuring decays to e^+e^- pairs. To this end, a dedicated event generator was integrated in the standard SAND simulation chain, allowing the definition of the decay signatures in the tracker and the development of a dedicated track fitting algorithm for e^+/e^- . A multivariate classification was then performed to estimate the SAND sensitivity to the models under study. A significant improvement over the current limits after one year of exposure was found across the model range.

Contents

Introduction	vii
1 Neutrino Physics	1
1.1 The physics of neutrinos	1
1.1.1 Massive neutrinos in the Standard Model	1
1.1.2 Neutrino mixing and oscillations	4
1.1.3 Neutrino Oscillation experiments	11
1.2 Heavy Neutral Leptons	29
1.2.1 Heavy neutrinos in seesaw models	29
1.2.2 Minimal HNL models	31
1.2.3 HNLs in Dark Sector models	33
2 The Deep Underground Neutrino Experiment	41
2.1 The design of DUNE	41
2.1.1 The LBNF neutrino beam	42
2.1.2 The DUNE Far Detector	42
2.1.3 The Near Detector Complex	46
2.2 The scientific programme of DUNE	49
2.2.1 Neutrino oscillation program	51
2.2.2 Low-energy physics and Supernova neutrinos	53
2.2.3 Beyond the Standard Model searches	56
3 The SAND detector	59
3.1 The SAND detector design	60
3.1.1 The SAND magnet	60
3.1.2 The Electromagnetic Calorimeter (ECAL)	60
3.1.3 The GRAIN LAr active target	61
3.1.4 The SAND Straw Tube Tracker	62
3.2 The SAND physics program	65
3.2.1 Beam monitoring	65
3.2.2 Flux measurements	65
3.2.3 Constraints on nuclear smearing	67
3.2.4 Background rejection	67
3.2.5 Solid Hydrogen Measurement	67
3.2.6 Precision measurements and searches	68

4	The SAND simulation framework	71
4.1	The standard SAND simulation chain	71
4.1.1	ND and SAND detector geometries	71
4.1.2	The GENIE neutrino event generator	75
4.1.3	Particle propagation in edep-sim	77
4.1.4	The SAND reconstruction framework	77
4.2	Dark neutrino event generators	78
4.2.1	The DarkNews event generator	78
4.2.2	The DarkNeutrino module in GENIE	80
4.2.3	Correction to the DarkNeutrino integrator	83
5	Simulation of dark neutrino events in SAND	85
5.1	Event generation	85
5.1.1	Geometry and generation volumes	86
5.1.2	Signal productions	87
5.1.3	Background productions	88
5.2	Signal event features	89
5.2.1	Interaction volumes	89
5.2.2	Flight lengths	89
5.2.3	Signal event kinematics	91
5.3	Event signature and backgrounds	92
5.3.1	Topological selection cuts	93
5.3.2	Background kinematics	95
5.3.3	Sources of background	95
5.4	Reconstruction of e^+ and e^- tracks	98
5.4.1	Helicoidal fit	98
5.4.2	Reconstruction performance	101
6	Discovery potential for dark neutrinos in SAND	109
6.1	Simulated models	109
6.2	Multivariate classification with TMVA	110
6.2.1	Training and application pipelines in TMVA	110
6.2.2	Model implementation and training	111
6.3	Estimation of the discovery potential	114
6.3.1	Score optimisation	114
6.3.2	Model selection	117
6.3.3	Discovery potential	119
6.4	Future prospects	122
	Conclusions	125
A	Tracker response simulations with Garfield++	127
A.0.1	Fast digitisation	127
A.0.2	Response simulation with Garfield++	127
A.0.3	Realistic digitisation with lookup tables	129

B	TMVA workflow	135
B.0.1	Training and testing pipeline	135
B.0.2	Model output and application	136
Bibliography		137

Introduction

The experimental study of neutrino oscillations has made it possible to determine many of the parameters related to neutrino masses through the combined use of multiple experimental channels and sources. Several parameters are yet to be determined, most importantly the neutrino mass ordering and the CP-violating phase of the neutrino mixing matrix, the latter of which would have a profound impact on the whole lepton sector of particle physics.

The search for the remaining unknowns in neutrino physics is the aim of the next-generation Deep Underground Neutrino Experiment (DUNE), leveraging innovative detector technologies and the most powerful neutrino beam available at the time of completion. DUNE will feature two experimental sites, the Near Detector (ND) at Fermilab, a few hundred meters from the neutrino beam source, and the Far Detector (FD) located at SURF, 1300 km away from the source. In order to determine the oscillation parameters through the data of the multi-kiloton Liquid Argon Time Projection Chambers (LArTPC) at the FD, the neutrino beam must be characterized at the ND before the onset of oscillations: the three proposed detectors at the ND will be able to provide complementary information on the neutrino beam and to refine the interaction models that will be necessary to analyse the FD data.

The Near Detector will provide an optimal environment for searches for Physics Beyond the Standard Model (BSM), owing to its high neutrino flux at a short baseline and the presence of a complementary suite of subsystems. It comprises a LArTPC detector with a 70 t target mass (ND-LAr) as well as two detectors featuring low-density trackers and EM-calorimeters within a magnetised volume (SAND and ND-GAr). This setup will enable excellent momentum and vertex resolution, and a sensitivity to BSM signatures beyond those accessible to LAr detectors. The low-density tracker allows a more precise measurement of vertex activity and the detection of electromagnetic BSM signatures that are difficult to observe in liquid argon. The use of a magnetic field for the detectors will further enable charge identification.

Among the broad range of BSM physics possible at DUNE-ND, ranging from Light Dark Matter to neutrino tridents and sterile neutrino mixing, searches for weakly-interacting long-lived particles such as Heavy Neutral Leptons (HNLs) and hidden sector portals have attracted a significant interest over the past years. Such particles could be produced at the DUNE beam target and, due to the small expected mixing angles, decay inside the ND detectors into Standard Model particles.

A class of HNL models that has recently gained attention within the community is that of Dark Neutrinos, which postulates a minimal $U(1)'$ gauge group and one or more HNLs whose associated flavour states, the *dark neutrinos*, carry the new gauge charge. In such a scenario, HNLs are produced in the scattering of an active flavour

state on a nuclear target, and can decay to di-lepton pairs such as e^+e^- and $\mu^+\mu^-$.

This thesis work aims to estimate the discovery potential of the SAND detector at the DUNE-ND for dark neutrinos decaying to e^+e^- . An analysis pipeline for dark neutrino events was developed, starting from the integration of a dedicated event generator for dark neutrino production and decay within the standard SAND simulation framework. The analysis focused on e^+e^- pair signatures in the SAND inner tracker. This subsystem features low-density gas detectors with diffuse target layers within the inner magnetised volume, providing an optimal solution for accurate momentum reconstruction and charge identification. The potential signatures of a range of dark neutrino models in the SAND tracker were studied in order to identify the background topologies from pair production in standard neutrino interactions. A dedicated track reconstruction algorithm was developed to extract the kinematics of signal and background e^+e^- pairs. The reconstructed variables were used as inputs to a Boosted Decision Tree (BDT) classifier for the final signal-background separation step, enabling an estimation of the discovery potential for the dark neutrino models under study.

The dissertation is structured as follows:

- Chapter 1 outlines the state of the art in the physics of massive neutrino oscillations before introducing the theory and current experimental constraints for HNL models, with a particular focus on Dark Neutrinos.
- Chapter 2 presents the design and scientific program of the future DUNE experiment.
- Chapter 3 discusses the design and physics program of the SAND detector.
- Chapter 4 outlines the standard SAND simulation chain, and introduces the dark neutrino event generators used in the analysis and their integration within the standard simulation framework.
- Chapter 5 details the first steps in the analysis pipeline: from an assessment of the dark neutrino and background topologies, the event signature is defined. A track reconstruction algorithm for the signature e^+e^- pairs is developed and its performance evaluated.
- Chapter 6 describes the final stage in the analysis, with a pipeline for the training and testing of a Boosted Decision Tree classifier to optimise signal-background separation. The optimal classifier model and working point are determined, and the final estimation of the discovery potential from the selected signal and background is presented.

Chapter 1

Neutrino Physics

1.1 The physics of neutrinos

Neutrinos have been at the forefront of particle physics ever since they were proposed by Pauli in 1930, the study of their interactions having proven crucial for the description of phenomena ranging from β -decays and the quark structure of nucleons to stellar thermonuclear reactions and the gravitational collapse in supernovae.

Neutrinos are the only known colourless and neutral fundamental fermions, and, as such, they only interact through weak charged and neutral current processes, with much smaller cross sections at all available energies than those of strong and electromagnetic processes. The detection of neutrino interactions thus required the development of advanced experimental techniques which have provided unique handles on numerous facets of fundamental physics.

In the original formulation of the Standard Model, neutrinos were assumed to be massless, but the observation of *Neutrino flavour oscillations* in many separate experimental regimes have since proven that their mass is in fact non-zero. While the absolute scale of neutrino masses is still an open question, the available upper limits are so much smaller compared to the masses of other fundamental fermions that a *natural* production mechanism may require Beyond the Standard Model physics.

Neutrinos can thus be a probe for New Physics beyond the Standard Model, with flavour oscillations potentially being the first hint of new such physics, while searches for neutrinoless double β -decay may provide evidence for the Majorana nature of neutrinos and a non-standard origin of their masses.

This section is intended to provide a brief theoretical overview of the main aspects of neutrino physics, introducing the relevant observables and the current experimental status, presenting finally the questions that are still open for the next generation of experiments to investigate.

1.1.1 Massive neutrinos in the Standard Model

Three types, or *flavour states*, of neutrinos and antineutrinos have been observed at present. They are categorized into electron-type, muon-type and tau-type neutrinos (ν_e , ν_μ and ν_τ) based on the charged lepton that they couple to in weak charged current interactions. The cancellation of quantum anomalies in the Standard Model

constrains the number of left-handed interacting flavour neutrino states to be the same as that of charged leptons. Indeed, results from LEP have ruled out additional neutrino flavours with masses below $M_W/2$ [1, 2].

The weak charged and neutral current (CC and NC respectively) of flavour neutrino states are described, to a great accuracy, by the Standard Model lagrangian terms:

$$\mathcal{L}_\ell^{\text{CC}} = -\frac{g}{2\sqrt{2}}j_\alpha^{\text{CC}}W^\alpha + \text{h.c.}, \quad j_\alpha^{\text{CC}} = 2 \sum_{\ell=e,\mu,\tau} \bar{\nu}_{\ell L}\gamma_\alpha \ell_L \quad (1.1.1)$$

$$\mathcal{L}_\ell^{\text{NC}} = -\frac{g}{2\cos\theta_W}j_\alpha^{\text{NC}}Z^\alpha, \quad j_\alpha^{\text{NC}} = \sum_{\ell=e,\mu,\tau} \bar{\nu}_{\ell L}\gamma_\alpha \nu_{\ell L}$$

The phenomenon of neutrino flavour oscillations, which has been observed by now over a wide region of the parameter space [3], has furthermore established that neutrinos have a non-zero mass. As a matter of fact, neutrino flavour mixing requires mass terms that feature a combination of the flavour fields.

A Dirac neutrino mass term can be generated through the same Higgs mechanism as with the other Standard Model fermions, introducing right-handed (RH) components $\nu_{\alpha R}$ of the neutrino fields ($\alpha = e, \mu, \tau$). The presence of RH neutrino fields in this *Minimally Extended Standard Model* eliminates the asymmetry between the lepton and quark sector, allowing to write an extended SM Higgs-lepton Yukawa Lagrangian:

$$\mathcal{L}_{H,L} = - \sum_{\alpha,\beta=e,\mu,\tau} Y_{\alpha\beta}^{\ell\ell} \bar{L}_{\alpha L} \Phi \ell'_{\beta R} - \sum_{\alpha,\beta=e,\mu,\tau} Y_{\alpha\beta}^{\ell\nu} \bar{L}_{\alpha L} \tilde{\Phi} \nu'_{\beta R} + \text{H.c.}, \quad (1.1.2)$$

where $L_{\alpha L}$ are the left-handed lepton doublets and the conjugate Higgs doublet $\tilde{\Phi}$ is

$$\tilde{\Phi} = i\sigma_2 \Phi^* = \frac{1}{\sqrt{2}} \begin{pmatrix} v + h \\ 0 \end{pmatrix}. \quad (1.1.3)$$

It is thus possible to write the complete Higgs-lepton Lagrangian as:

$$\mathcal{L}_{H,L} = - \left(\frac{v + h}{\sqrt{2}} \right) \left[\bar{\ell}'_L Y^{\ell\ell} \ell'_R + \bar{\nu}'_L Y^{\ell\nu} \nu'_R \right] + \text{H.c.}, \quad (1.1.4)$$

with $\ell'_{L,R}$ and $\nu'_{L,R}$ being the chiral lepton arrays:

$$\ell'_i \equiv \begin{pmatrix} e'_i \\ \mu'_i \\ \tau'_i \end{pmatrix}, \quad \nu'_i \equiv \begin{pmatrix} \nu'_{ei} \\ \nu'_{\mu i} \\ \nu'_{\tau i} \end{pmatrix}, \quad i = L, R. \quad (1.1.5)$$

The matrices $Y^{\ell\ell}$ and $Y^{\ell\nu}$ can then be diagonalized as:

$$\begin{aligned} V_L^{\ell\dagger} Y^{\ell\ell} V_R^\ell &= Y^\ell, & \text{with } Y_{\alpha\beta}^\ell &= y_\alpha^\ell \delta_{\alpha\beta} & (\alpha, \beta = 1, 2, 3) \\ \text{and} & & & & \\ V_L^{\nu\dagger} Y^{\ell\nu} V_R^\nu &= Y^\nu, & \text{with } Y_{kj}^\nu &= y_k^\nu \delta_{kj} & (k, j = 1, 2, 3), \end{aligned} \quad (1.1.6)$$

y_α^ℓ and y_k^ν being real and positive and $V_L^\ell, V_R^\ell, V_L^\nu, V_R^\nu$ being appropriate 3×3 unitary matrices. The *massive lepton fields*, having, that is, a diagonalised Yukawa Lagrangian, can be defined as:

$$\ell_L = V_L^{\ell\dagger} \ell'_L \equiv \begin{pmatrix} e_L \\ \mu_L \\ \tau_L \end{pmatrix}, \quad \ell_R = V_R^{\ell\dagger} \ell'_R \equiv \begin{pmatrix} e_R \\ \mu_R \\ \tau_R \end{pmatrix} \quad (1.1.7)$$

$$\mathbf{n}_L = V_L^{\nu\dagger} \mathbf{\nu}'_L \equiv \begin{pmatrix} \nu_{1L} \\ \nu_{2L} \\ \nu_{3L} \end{pmatrix}, \quad \mathbf{n}_R = V_R^{\nu\dagger} \mathbf{\nu}'_R \equiv \begin{pmatrix} \nu_{1R} \\ \nu_{2R} \\ \nu_{3R} \end{pmatrix},$$

leading to the diagonalised Higgs-lepton Yukawa Lagrangian:

$$\begin{aligned} \mathcal{L}_{H,L} &= - \left(\frac{v+h}{\sqrt{2}} \right) [\overline{\ell}_L Y^\ell \ell_R + \overline{\mathbf{n}}_L Y^\nu \mathbf{n}_R] + \text{h.c.} \\ &= - \left(\frac{v+h}{\sqrt{2}} \right) \left[\sum_{\alpha=e,\mu,\tau} y_\alpha^\ell \overline{\ell}_{\alpha L} \ell_{\alpha R} + \sum_{k=1}^3 y_k^\nu \overline{\nu}_{kL} \nu_{kR} \right] + \text{h.c.}, \end{aligned} \quad (1.1.8)$$

which yields the Dirac mass terms m_α and m_k respectively for the charged lepton fields $\ell_\alpha = \ell_{\alpha L} + \ell_{\alpha R}$ and the neutrino fields $\nu_k = \nu_{kL} + \nu_{kR}$:

$$m_\alpha = \frac{y_\alpha^\ell v}{\sqrt{2}} \quad (\alpha = e, \mu, \tau) \quad \text{and} \quad m_k = \frac{y_k^\nu v}{\sqrt{2}} \quad (k = 1, 2, 3). \quad (1.1.9)$$

The Standard Model leptonic weak charged current interaction lagrangian in eqn. (1.1.1) can be written in terms of the chiral lepton arrays ℓ' and $\mathbf{\nu}'$ as:

$$\mathcal{L}_\ell^{CC} = -\frac{g}{2\sqrt{2}} j_\alpha^{CC} W^\alpha + \text{h.c.}, \quad j_\alpha^{CC} \equiv 2\overline{\mathbf{\nu}}'_L \gamma_\alpha \ell'_L, \quad (1.1.10)$$

and through eqn. (1.1.7) the lepton current j_α^{CC} can be written as :

$$j_\alpha^{CC} = 2\overline{\mathbf{\nu}}'_L \gamma_\alpha \ell'_L = 2\overline{\mathbf{n}}_L V_L^{\nu\dagger} \gamma_\alpha V_L^\ell \ell_L = 2\overline{\mathbf{n}}_L \gamma_\alpha V_L^{\nu\dagger} V_L^\ell \ell_L. \quad (1.1.11)$$

The charged lepton current is as such dependent on the product

$$U \equiv V_L^{\nu\dagger} V_L^\ell, \quad (1.1.12)$$

which describes the mixing in the lepton sector in an analogous way to the CKM matrix for quarks. The leptonic weak charged current specifically reads, in terms of the mixing matrix

$$j_\alpha^{CC} = 2\overline{\mathbf{n}}_L \gamma_\alpha U^\dagger \ell_L. \quad (1.1.13)$$

Defining, as is customary, the left-handed *flavour neutrino* fields as

$$\mathbf{\nu}_L \equiv U \mathbf{n}_L = V_L^{\ell\dagger} \mathbf{\nu}'_L, \quad \text{with} \quad \mathbf{\nu}_L \equiv \begin{pmatrix} \nu_{eL} \\ \nu_{\mu L} \\ \nu_{\tau L} \end{pmatrix} \quad (1.1.14)$$

makes it possible to express the leptonic weak charged current in the form of eqn. (1.1.1):

$$j_\alpha^{CC} = 2\overline{\nu}_L\gamma_\alpha\ell_L = 2\sum_{\ell=e,\mu,\tau}\overline{\nu}_{\ell L}\gamma_\alpha\ell_L. \quad (1.1.15)$$

The introduction of left-handed flavour neutrino fields makes it possible to define flavour lepton numbers as in the SM, the weak charged current in eqn. (1.1.15) being invariant under global $U(1)$ transformations

$$\ell_{\alpha L} \rightarrow e^{i\phi_\alpha}\ell_{\alpha L}, \quad \nu_{\alpha L} \rightarrow e^{i\phi_\alpha}\nu_{\alpha L} \quad (\alpha = e, \mu, \tau), \quad (1.1.16)$$

with each flavour having different phases ϕ_α . The Yukawa lagrangian terms for charged leptons are invariant for right-handed chiral fields transforming with the same phases:

$$\ell_{\alpha R} \rightarrow e^{i\phi_\alpha}\ell_{\alpha R} \quad (\alpha = e, \mu, \tau). \quad (1.1.17)$$

By contrast, it is not possible to find any transformation of the right-handed flavour neutrino fields $\nu_{\alpha R}$ that leaves simultaneously invariant both the Yukawa terms and the kinetic part of the neutrino lagrangian, meaning that flavour lepton numbers are not conserved in the neutrino sector. This feature of the neutrino lagrangian is at the origin of the phenomenon of neutrino oscillations, which will be discussed in the next section [1].

1.1.2 Neutrino mixing and oscillations

The phenomenology of neutrino flavour oscillations is determined by the mixing matrix U in eqn. (1.1.12): known as Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, its parameters are also crucial to the issue of CP-violation in the lepton sector and the Majorana nature of neutrinos. The measurement of the PMNS parameters has therefore been one of the main objectives of oscillation experiments.

Being a unitary 3×3 matrix, the PMNS is characterised by $3^2 = 9$ real parameters, which can be expressed as $n_\theta = 3(3-1)/2 = 3$ mixing angles and $n_\phi = 3(3+1)/2 = 6$ phases. In the case of Dirac neutrinos, though, five of the six physical phases can be eliminated by a rephasing of the neutrino and charged lepton fields since the current terms are invariant under the global phase transformations

$$\ell_\alpha \rightarrow \ell_\alpha e^{i\varphi_\alpha}, \quad \nu_k \rightarrow \nu_k e^{i\varphi_k} \quad \text{and} \quad U_{\alpha k} \rightarrow U_{\alpha k} e^{i(\varphi_\alpha - \varphi_k)}. \quad (1.1.18)$$

An overall phase $Ue^{i\varphi}$ in the PMNS matrix has thus no physical consequence, allowing five of the six phases to be absorbed in the lepton fields. The case of Majorana neutrinos, whose lagrangian mass term is not invariant under $U(1)$ transformations, will be discussed in more detail in sect. 1.1.2, assuming Dirac neutrinos in the following treatment of mixing and oscillations [1, 2].

The most widely used parametrization of the PMNS matrix in the field of neutrino mixing is:

$$U \equiv \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{-i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad (1.1.19)$$

where $c_{ij} \equiv \cos \theta_{ij}$ and $s_{ij} \equiv \sin \theta_{ij}$. As will be discussed in the following, this notation has proven particularly convenient since θ_{13} has been measured to be relatively small, leading to an approximately diagonal central matrix [4]. The individual matrix elements can thus be written as:

$$\begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{13} - c_{12}s_{23}s_{13}e^{-i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{-i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{-i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{-i\delta} & c_{23}c_{13} \end{pmatrix}. \quad (1.1.20)$$

Neutrino oscillations in vacuum

In the standard theory of neutrino oscillations, neutrinos with flavour α and momentum $\vec{p}$ are produced in weak charged-current (CC) interactions from a charged lepton ℓ_α^- or together with a charged antilepton ℓ_α^+ . In all experimental setups so far, particle momenta and energies cannot be measured with the accuracy needed to distinguish the neutrino mass states. Thus the flavour state can be expressed as a coherent superposition of the mass eigenstates, according to 1.1.14:

$$|\nu_\alpha\rangle = \sum_k U_{\alpha k}^* |\nu_k\rangle \quad (\alpha = e, \mu, \tau) \quad (1.1.21)$$

The massive neutrino states $|\nu_k\rangle$ are eigenstates of the Hamiltonian

$$\mathcal{H} |\nu_k\rangle = E_k |\nu_k\rangle, \quad (1.1.22)$$

with energy eigenvalues

$$E_k = \sqrt{\vec{p}^2 + m_k^2}. \quad (1.1.23)$$

The Shrödinger equation therefore implies that the massive eigenstates evolve in time as a plane wave:

$$|\nu_k(t)\rangle = e^{-iE_k t} |\nu_k\rangle, \quad (1.1.24)$$

so that the time evolution of a flavour state $|\nu_\alpha(t)\rangle$ describing a neutrino created with a definite flavour α at $t = 0$ is given by

$$|\nu_\alpha(t)\rangle = \sum_{k=1}^3 U_{\alpha k}^* e^{-iE_k t} |\nu_k\rangle, \quad (1.1.25)$$

such that $|\nu_\alpha(t=0)\rangle = |\nu_\alpha\rangle$. Through the unitarity relation $U^\dagger U = \mathbf{1}$, the massive states can be expressed in terms of the flavour states

$$|\nu_k\rangle = \sum_\alpha U_{\alpha k} |\nu_\alpha\rangle, \quad (1.1.26)$$

which, substituting into eqn. (1.1.25), gives the time evolution in terms of the flavour eigenstates:

$$|\nu_\alpha(t)\rangle = \sum_{\beta=e,\mu,\tau} \left(\sum_k U_{\alpha k}^* e^{-iE_k t} U_{\beta k} \right) |\nu_\beta\rangle. \quad (1.1.27)$$

Hence, a pure flavour state $|\nu_\alpha\rangle$ at $t = 0$ becomes a superposition of different flavour states at $t > 0$, with an amplitude for the transition $\nu_\alpha \rightarrow \nu_\beta$ at time t being:

$$A_{\nu_\alpha \rightarrow \nu_\beta}(t) \equiv \langle \nu_\beta | \nu_\alpha(t) \rangle = \sum_k U_{\alpha k}^* U_{\beta k} e^{-iE_k t}, \quad (1.1.28)$$

and a probability for the transition

$$P_{\nu_\alpha \rightarrow \nu_\beta}(t) = \left| A_{\nu_\alpha \rightarrow \nu_\beta}(t) \right|^2 = \sum_{k,j} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* e^{-i(E_k - E_j)t}. \quad (1.1.29)$$

For ultrarelativistic neutrinos, the energy-momentum relation in eqn. (1.1.23) can be approximated by

$$E_k \simeq E + \frac{m_k^2}{2E}, \quad E = |\vec{p}|, \quad (1.1.30)$$

so that

$$E_k - E_j \simeq \frac{\Delta m_{kj}^2}{2E}, \quad (1.1.31)$$

where $\Delta m_{kj}^2 = m_k^2 - m_j^2$ is the squared mass difference. The transition probability can therefore be approximated as:

$$P_{\nu_\alpha \rightarrow \nu_\beta}(t) = \sum_{k,j} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* \exp \left(-i \frac{\Delta m_{kj}^2 t}{2E} \right). \quad (1.1.32)$$

While the neutrino propagation time is not measured in oscillation experiments, the distance L between the source and the detector is instead known, allowing to approximate, for ultrarelativistic neutrinos $t \simeq L$, leading to:

$$P_{\nu_\alpha \rightarrow \nu_\beta}(t) = \sum_{k,j} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* \exp \left(-i \frac{\Delta m_{kj}^2 L}{2E} \right). \quad (1.1.33)$$

Thus, the phases of neutrino oscillations

$$\Phi_{kj} = \frac{\Delta m_{kj}^2 L}{2E} \quad (1.1.34)$$

are determined by two experimental parameters, the source-detector distance L and the neutrino energy as well as by the squared-mass differences Δm_{kj}^2 , which are physical constants to be measured. The amplitude of the oscillations, on the other hand, only depends on the elements of the PMNS matrix, which are physical constants. The main driver behind neutrino oscillation experiments has therefore been the determination of the PMNS matrix elements and of the squared-mass differences. While the phenomenon of neutrino oscillations implies massive neutrinos, it doesn't provide information on the absolute scale of their masses, except that either m_k^2 or m_j^2 must be larger than $|\Delta m_{kj}^2|$ [1].

***CP*-violation in the lepton sector**

The observation of *CP*-violation in the neutral K and B meson systems has been a crucial result in the field of electroweak interactions, with the observed amount of *CP*-violation being compatible with that generated by the remaining physical phase of the CKM [4].

The case of three-flavour mixing of Dirac neutrinos is analogous to that of quarks. For *CP*-violation in the lepton sector to occur, the following conditions must be met:

1. No two charged leptons and neutrinos must be degenerate in mass,
2. no mixing angle must be either 0 or $\pi/2$,
3. The PMNS matrix must be complex; i.e., the physical phase δ must be different from 0 or π .

These 14 conditions can be concisely expressed as

$$-2J (m_{\nu_2}^2 - m_{\nu_1}^2) (m_{\nu_3}^2 - m_{\nu_1}^2) (m_{\nu_3}^2 - m_{\nu_2}^2) \cdot (m_\mu^2 - m_e^2) (m_\tau^2 - m_e^2) (m_\tau^2 - m_\mu^2) \neq 0, \quad (1.1.35)$$

where $J \equiv [U_{e2}U_{e3}^*U_{\mu 2}^*U_{\mu 3}]$ is the Jarlskog Invariant quantifying the *CP*-violation in the lepton sector in a rephasing-invariant way.

A *CP*-transformation exchanges neutrinos with antineutrinos and reverses their helicities

$$\nu_\alpha \xrightarrow{\text{CP}} \bar{\nu}_\alpha, \quad (1.1.36)$$

therefore, the oscillation channel $\nu_\alpha \rightarrow \nu_\beta$ is transformed to $\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta$, so that experiments can be performed to the asymmetry

$$A_{\alpha\beta}^{CP} = P_{\nu_\alpha \rightarrow \nu_\beta} - P_{\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta}. \quad (1.1.37)$$

As *CPT*-conservation implies $A_{\alpha\beta}^{CP} = -A_{\beta\alpha}^{CP}$, it follows that the *CP*-asymmetry can only be measured in transitions between different flavours [1].

Majorana Neutrinos

While the Minimally Extended Standard Model introduces mass terms for Dirac neutrino chiral fields, the possibility of neutrinos being Majorana particles has not yet been ruled out by experiments. Indeed, a potential Majorana nature of neutrinos is crucial to the See-Saw mechanisms, some of the most promising explanations for *natural* light neutrino masses.

In Majorana neutrinos, the chiral components ν_L and ν_R are not independent, being instead related by the Majorana condition:

$$\nu_R = \nu_L^C = \mathcal{C}\bar{\nu}_L^T, \quad (1.1.38)$$

where $C \equiv i\gamma^2\gamma^0$ is the charge conjugation operator. The previous definition can be substituted in the chiral decomposition of the neutrino fields, yielding the relation

$$\nu = \nu_L + \nu_L^C = \nu^C. \quad (1.1.39)$$

Majorana particles are thus equal to their own antiparticles, hence why only neutral fermions like neutrinos can be described by a Majorana field. Since for massless neutrinos the chiral components obey the same Weyl equations in both the Dirac and Majorana cases, and only the left-handed component can interact, the two descriptions can only be distinguished by measuring some effect due to the neutrino mass. A Majorana mass term for a single neutrino type ν can be generated with a single chiral field ν_L thanks to the Majorana condition in eqn. (1.1.38):

$$\mathcal{L}_{\text{mass}}^{\text{M}} = -\frac{1}{2}m\overline{\nu_L^C}\nu_L + \text{h.c.}, \quad (1.1.40)$$

the full Majorana lagrangian consisting of the kinetic terms for ν_L and ν_L^C in addition to $\mathcal{L}_{\text{mass}}^{\text{M}}$. The generalization to three generations of massive Majorana neutrinos requires the diagonalization of the mass term, as with Dirac neutrinos, expressing the left-handed flavour fields ν'_L as linear combinations of massive neutrino fields $\mathbf{n}_L$:

$$\nu'_L = V_L^\nu \mathbf{n}_L \quad \text{with} \quad \mathbf{n}_L = \begin{pmatrix} \nu_{1L} \\ \nu_{2L} \\ \nu_{3L} \end{pmatrix}. \quad (1.1.41)$$

The Majorana mass term can thus be written as

$$\mathcal{L}_{\text{mass}}^{\text{M}} = -\frac{1}{2}\overline{\mathbf{n}_L^C} M \mathbf{n}_L + \text{h.c.} = \frac{1}{2} \sum_{k=1}^3 m_k \overline{\nu_{kL}^C} \nu_{kL} + \text{h.c.}, \quad (1.1.42)$$

M being the diagonalised mass matrix. Differently from the Dirac lagrangian mass term, the Majorana mass term is not invariant under global $U(1)$ gauge transformations

$$\nu_{kL} \rightarrow e^{i\varphi} \nu_{kL} \quad (k = 1, 2, 3), \quad (1.1.43)$$

with the same phase φ for all massive neutrinos. This corresponds to a violation of the total lepton number conservation, leading to distinctive phenomena such as *neutrinoless* double β -decay, which represent the best experimental handle to probe the Majorana nature of neutrinos.

The result above can be extended to global $U(1)$ gauge transformations $\nu_{kL} \rightarrow e^{i\varphi_k} \nu_{kL}$, implying, recalling sect. 1.1.2, that the left-handed massive neutrinos cannot be rephased to cancel two of the PMNS phases, leaving three physical CP -violating phases. It is thus possible to write the mixing matrix of Majorana neutrino as the product of a Dirac unitary mixing matrix U^D similar to the standard PMNS, and a diagonal unitary matrix D^M featuring the two extra phases:

$$U = U^D D^M, \quad D^M \equiv \text{diag}(e^{i\lambda_1}, e^{i\lambda_2}, e^{i\lambda_3}) \quad \text{with} \quad \lambda_1 = 0. \quad (1.1.44)$$

Using the previous expression, the oscillation formula in eqn. (1.1.33) can be extended to Majorana neutrinos, with the products $U_{\alpha k}^* U_{\beta k}$ reducing to:

$$U_{\alpha k}^* U_{\beta k} = U_{\alpha k}^{D*} D_{\alpha k}^{M*} U_{\beta k}^D D_{\beta k}^M = U_{\alpha k}^{D*} U_{\beta k}^D, \quad (1.1.45)$$

since $D_{\alpha k} = \delta_{\alpha k} e^{i\lambda_k}$. Hence, the extra Majorana phases do not enter into the transition probabilities for neutrino oscillations, which are the same as for Dirac neutrinos.

Lepton number violating processes such as the previously mentioned neutrinoless double beta decay remain the only handles to probe the Majorana nature of neutrinos [1, 2].

Determining the nature of neutrinos and quantifying the amount of CP -violation in the lepton sector through the measurement of δ_{CP} are essential steps for the understanding of the processes of *leptogenesis* and *baryogenesis* [5].

Neutrino oscillations in matter

As was first shown by Wolfenstein [6], the propagation of neutrinos through matter is not only influenced by the oscillation parameters (squared-mass differences and PMNS elements) as when through vacuum: the coherent scattering of neutrinos in matter must additionally be taken into account. The coherent scattering contribution to the neutrino Hamiltonian is determined by the electron number-density of the medium, giving rise to a varied phenomenology according to the dependence of the latter on the distance [2].

The effect of coherent interactions with the medium through coherent forward elastic weak CC and NC scattering can be modelled as an effective potential

$$V_\alpha = V_{CC}\delta_{\alpha e} + V_{NC} = \sqrt{2}G_F \left(N_e\delta_{\alpha e} - \frac{1}{2}N_n \right), \quad (1.1.46)$$

V_{CC} and V_{NC} being respectively the CC and NC scattering contributions. Only electron neutrinos can undergo CC coherent scattering in ordinary matter, while the contribution of NC is the same for all flavours and is proportional to the neutron density N_n . This effective potential is very small, since

$$\sqrt{2}G_F \simeq 7.63 \times 10^{-14} \frac{\text{eVcm}^3}{N_A}, \quad (1.1.47)$$

N_A being Avogadro's number.

The total neutrino Hamiltonian in matter is thus

$$\mathcal{H} = \mathcal{H}_0 + \mathcal{H}_I \quad \text{with} \quad \mathcal{H}_I |\nu_\alpha\rangle = V_\alpha |\nu_\alpha\rangle, \quad (1.1.48)$$

where $\mathcal{H}_0$ is the vacuum hamiltonian. A neutrino with initial flavour α obeys the evolution equation

$$i \frac{d}{dt} |\nu_\alpha(t)\rangle = \mathcal{H} |\nu_\alpha(t)\rangle \quad \text{with} \quad |\nu_\alpha(0)\rangle = |\nu_\alpha\rangle. \quad (1.1.49)$$

The amplitude of a $\nu_\alpha \rightarrow \nu_\beta$ transition after a time t

$$\psi_{\alpha\beta}(t) = \langle \nu_\beta | \nu_\alpha(t) \rangle \quad \text{with} \quad \psi_{\alpha\beta}(0) = \delta_{\alpha\beta}, \quad (1.1.50)$$

has therefore a time evolution equation given by, using eqn. (1.1.48):

$$i \frac{d}{dt} \psi_{\alpha\beta}(t) = \sum_\eta \left(\sum_k U_{\beta k} E_k U_{\eta k}^* + \delta_{\beta\eta} V_\beta \right) \psi_{\alpha\eta}(t). \quad (1.1.51)$$

It can be shown that for ultrarelativistic neutrinos (with $t = x$), the equation above becomes [1]:

$$i \frac{d}{dt} \psi_{\alpha\beta}(t) = \sum_{\eta} \left(\sum_k U_{\beta k} \frac{\Delta m_{k1}^2}{2E} U_{\eta k}^* + \delta_{\beta e} \delta_{\eta e} V_{CC} \right) \psi_{\alpha\eta}(x), \quad (1.1.52)$$

which implies that neutrino oscillations in matter, as is the case for oscillations in vacuum, depend on the squared-mass differences. Equation 1.1.52 can be written in a matrix form as

$$i \frac{d}{dx} \Psi_{\alpha} = \mathcal{H}_F \Psi_{\alpha}, \quad (1.1.53)$$

where the effective Hamiltonian matrix $\mathcal{H}_F$ in the flavour basis is given by

$$\mathcal{H}_F = \frac{1}{2E} (U \mathbb{M}^2 U^{\dagger} + \mathbb{A}). \quad (1.1.54)$$

In the three-neutrino mixing framework one has

$$\Psi_{\alpha} = \begin{pmatrix} \psi_{\alpha e} \\ \psi_{\alpha \mu} \\ \psi_{\alpha \tau} \end{pmatrix}, \quad \mathbb{M}^2 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \Delta m_{21}^2 & 0 \\ 0 & 0 & \Delta m_{31}^2 \end{pmatrix}, \quad A = \begin{pmatrix} A_{CC} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad (1.1.55)$$

where $A_{CC} \equiv 2EV_{CC} = 2\sqrt{2}EG_F N_e$. From eqn. (1.1.54) follows that the Majorana phases in the vacuum mixing matrix do not have any effect on neutrino oscillations in matter, since the matrix D^M of Majorana phases in eqn. (1.1.44) cancels in the product $U \mathbb{M}^2 U^{\dagger}$. In a two-flavour framework, eqn. (1.1.53) can be written, neglecting a common phase, as:

$$i \frac{d}{dx} \begin{pmatrix} \psi_{ee} \\ \psi_{eX} \end{pmatrix} = \frac{1}{4E} \begin{pmatrix} -\Delta m^2 \cos 2\theta + A_{CC} & \Delta m^2 \sin 2\theta \\ \Delta m^2 \sin 2\theta & \Delta m^2 \cos 2\theta - A_{CC} \end{pmatrix} \begin{pmatrix} \psi_{ee} \\ \psi_{eX} \end{pmatrix}. \quad (1.1.56)$$

The effective Hamiltonian matrix $\mathcal{H}_F$ can be diagonalised by the orthogonal transformation

$$U_M^T \mathcal{H}_F U_M = \mathcal{H}_M, \quad (1.1.57)$$

where

$$\mathcal{H}_M = \frac{1}{4E} \text{diag}(-\Delta m_M^2, \Delta m_M^2) \quad (1.1.58)$$

is the effective mixing matrix in matter, and

$$\Delta m_M^2 = \sqrt{(\Delta m^2 \cos 2\theta - A_{CC})^2 + (\Delta m^2 \sin 2\theta)^2} \quad (1.1.59)$$

is the effective squared-mass difference. The effective mixing angle is then given by

$$\tan 2\theta_M = \frac{\tan 2\theta}{1 - \frac{A_{CC}}{\Delta m^2 \cos 2\theta}}. \quad (1.1.60)$$

A special resonance condition occurs when A_{CC} becomes

$$A_{CC}^R = \Delta m^2 \cos 2\theta, \quad (1.1.61)$$

which corresponds to the number density

$$N_e^R = \frac{\Delta m^2 \cos 2\theta}{2\sqrt{2}EG_F}. \quad (1.1.62)$$

Indeed, at the resonance density N_e^R , according to eqn. (1.1.60), $\theta_M = \pi/4$ and the mixing is maximal, allowing total flavour transitions if the resonance region is wide enough. This mechanism is known as *Mikheyev-Smirnov-Wolfenstein (MSW) effect* [1].

1.1.3 Neutrino Oscillation experiments

The aim of neutrino oscillation experiments is to retrieve, through the relations in sect. 1.1.2, the PMNS parameters and the squared-mass differences. Measuring the neutrino fluxes at the production and detection points allows to determine the oscillation probability into different neutrino flavours than the initial one. Two broad categories of experiments can be identified [1]:

- *Appearance experiments*, which measure transitions between different neutrino flavours. Since the final flavour may be either absent or present as a small contamination in the initial beam, such experiments can have very small backgrounds and thus be sensitive to small values of the mixing angle.
- *Disappearance experiments*, which measure the survival probability of a neutrino flavour

$$P_{\nu_\alpha \rightarrow \nu_\alpha}(L, E) = 1 - 4 \sum_{k>j} |U_{\alpha k}|^2 |U_{\alpha j}|^2 \sin^2 \left(\frac{\Delta m_{kj}^2 L}{2E} \right), \quad (1.1.63)$$

by comparing the initial and final interaction rates. Since interaction rates are affected by statistical fluctuations on top of flavour oscillations, small disappearance probabilities are difficult to reveal, making these experiments less suitable for measuring small mixing angles.

Two-flavour effective mixing

A two-neutrino mixing approximation, in which the coupling of flavour neutrinos to the third massive eigenstate is neglected, has proven particularly useful for interpreting experimental data, since

- it leads to simpler oscillation formulas that depend on fewer parameters than in the case of three-neutrino mixing,
- many experiments are not sensitive to the influence of three-neutrino mixing, allowing to analyse the data in an effective two-neutrino model.

The two flavour states ν_α and ν_β in the two-neutrino mixing model can be either pure flavour neutrinos ($\alpha, \beta = e, \mu, \tau$) or linear combinations of such states. The two

flavour neutrinos can be written as a superposition of the two massive neutrinos ν_1 and ν_2 with coefficients given by the effective two-neutrino mixing matrix

$$U = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}, \quad (1.1.64)$$

where $0 \leq \theta \leq \pi/2$. In the case of two-neutrino mixing, there is only one squared-mass difference:

$$\Delta m^2 \equiv \Delta m_{12}^2 \equiv m_2^2 - m_1^2, \quad (1.1.65)$$

where ν_1 can be defined as the lightest state, so that $\Delta m^2 > 0$. From eqn. (1.1.33), the probability of the $\nu_\alpha \rightarrow \nu_\beta$, $\alpha \neq \beta$ transitions can be derived straightforwardly:

$$P_{\nu_\alpha \rightarrow \nu_\beta}(L, E) = \frac{1}{2} \sin^2 2\theta \left[1 - \cos \left(\frac{\Delta m^2 L}{2E} \right) \right] = \sin^2 2\theta \sin^2 \left(\frac{\Delta m^2 L}{4E} \right) \quad (\alpha \neq \beta). \quad (1.1.66)$$

From the expression above, the survival probability $P_{\nu_\alpha \rightarrow \nu_\alpha}(L, E)$ follows trivially:

$$P_{\nu_\alpha \rightarrow \nu_\alpha}(L, E) = 1 - P_{\nu_\alpha \rightarrow \nu_\beta}(L, E) = 1 - \sin^2 2\theta \sin^2 \left(\frac{\Delta m^2 L}{4E} \right). \quad (1.1.67)$$

Averaged oscillation probability

Due to the finite intrinsic spatial and energetic uncertainties in any experiment, it is impossible to measure the oscillation probabilities for exact values of the L/E ratio. Hence, in practice it is always necessary to average any measurement over the appropriate distributions of L and E . In the two-neutrino mixing model, this average probability is

$$\langle P_{\nu_\alpha \rightarrow \nu_\beta}(L, E) \rangle = \frac{1}{2} \sin^2 2\theta \left[1 - \left\langle \cos \left(\frac{\Delta m^2 L}{2E} \right) \right\rangle \right], \quad (\alpha \neq \beta) \quad (1.1.68)$$

with

$$\left\langle \cos \left(\frac{\Delta m^2 L}{2E} \right) \right\rangle = \int \cos \left(\frac{\Delta m^2 L}{2E} \right) \phi \left(\frac{L}{E} \right) d\frac{L}{E}.$$

Assuming a Gaussian distribution of the experimental L/E ratio with average $\langle L/E \rangle$ and standard deviation $\sigma_{L/E}$, the average oscillation probability follows the behaviour in fig. 1.1. In particular, for

$$\frac{\Delta m^2 L}{2E} \gg 1, \quad (1.1.69)$$

the transition probability averages out to

$$\langle P_{\nu_\alpha \rightarrow \nu_\beta} \rangle = \frac{1}{2} \sin^2 2\theta \quad (\alpha \neq \beta), \quad (1.1.70)$$

so that the experiment can only yield information on $\sin^2 2\theta$ [1].




Figure 1.1: $\nu_\alpha \rightarrow \nu_\beta$ oscillation probability for $\sin^2 2\theta = 1$ as a function of $\langle L/E \rangle \Delta m^2$. *Solid line:* probability averaged over a Gaussian distribution for L/E with $\sigma_{L/E} = 0.2 \langle L/E \rangle$. *Dashed line:* un-averaged transition probability with $L/E = \langle L/E \rangle$ [1].

Classes of oscillation experiments

Since the phase of oscillations depends only on the physical parameter Δm^2 , experiments can be designed to be sensitive to different values of the squared-mass difference by tuning the ratio L/E . Namely, transitions to different flavours cannot be observed if

$$\frac{\Delta m^2 L}{2E} \ll 1, \quad (1.1.71)$$

since they would have low probabilities according to eqn. (1.1.66), while they would be averaged out for $\Delta m^2 L/2E \gg 1$ and would not depend on Δm^2 , as stated in eqn. (1.1.70). The *sensitivity* of an experiment to Δm^2 is therefore the value of Δm^2 for which

$$\frac{\Delta m^2 L}{2E} \sim 1. \quad (1.1.72)$$

Neutrino oscillation experiments have been historically classified according to their average L/E , which corresponds to their sensitivity to Δm^2 : the following will present the main classes of experiments according to their sensitivities and discuss briefly their possible designs [1].

Short Baseline experiments (SBL). Reaching sensitivities down to $\Delta m^2 \sim 0.1 \text{ eV}^2$, with $L \lesssim 1 \text{ km}$. Depending on the neutrino source, they are divided as follows:

- *Reactor SBL:* utilizing the large isotropic fluxes of $\bar{\nu}_e$ produced in nuclear reactors. Since reactor antineutrinos have energies of the order of a few MeV, baselines of $\sim 10 \text{ m}$ correspond to sensitivities of

$$L/E \lesssim 10 \text{ m/MeV} \Rightarrow \Delta m^2 \gtrsim 0.1 \text{ eV}^2. \quad (1.1.73)$$

As reactor antineutrino energies are too low to produce μ s or τ s, reactor-SBL experiments are limited to the $\bar{\nu}_e$ -disappearance channel, which can be probed in liquid scintillator detectors through the *inverse- β decay* channel.

- *Accelerator SBL*: experiments operating with neutrino beams produced by decay of pions, kaons and muons created by a proton beam impinging on a fixed target. Two designs in this class, *Pion Decay in Flight (DIF)* and *Muon Decay at Rest (DAR)* experiments, according to the reaction which produces their muon neutrino beam, reach sensitivities $\Delta m^2 \gtrsim 1 \text{ eV}^2$. *Beam Dump* experiments, where neutrinos are produced in the prompt decays of charmed hadrons from high energy protons impinging on a target, can reach neutrino energies of $\sim 10^2 \text{ GeV}$, yielding sensitivities $\Delta m^2 \gtrsim 10^2 \text{ eV}^2$ for typical propagation distances of $L \sim 1 \text{ km}$.

Long Baseline experiments (LBL). Experiments generally featuring the same sources as SBL experiments, but source-detector distances larger than those by two to three orders of magnitude, allowing $\Delta m^2 \gtrsim 10^{-4} \text{ eV}^2$ sensitivities. LBL-experiments can be classified as follows:

- *Reactor LBL*: reactor neutrino experiments with $L \sim 1 \text{ km}$, covering a sensitivity range of $\Delta m^2 \gtrsim 10^{-3} \text{ eV}^2$.
- *Accelerator LBL*: accelerator experiments in which the source-detector distance is about 10^{-2} - 10^{-3} km , corresponding, for the $\sim 1 \text{ GeV}$ energies of accelerator neutrinos, to

$$L/E \lesssim 10^3 \text{ km/GeV} \Rightarrow \Delta m^2 \gtrsim 10^{-3} \text{ eV}^2. \quad (1.1.74)$$

- *Atmospheric neutrino experiments (ATM)*: experiments detecting neutrinos that are produced in cosmic-ray showers in the upper atmosphere. The energies of cosmic-ray neutrinos cover a wide range, $500 \text{ MeV} \lesssim E \lesssim 100 \text{ GeV}$, while the source-detector distance varies from $L \sim 20 \text{ km}$ of neutrinos coming from the zenith, to $L \sim 1.4 \times 10^4 \text{ km}$ for neutrinos produced at the opposite side of the planet. The typical values of L/E and the associated sensitivity in atmospheric neutrino experiments are thus:

$$L/E \lesssim 10^4 \text{ km/GeV} \Rightarrow \Delta m^2 \gtrsim 10^{-4} \text{ eV}^2. \quad (1.1.75)$$

Very Long Baseline Experiments (VLB). Experiments with a source-detector distance larger than LBL experiments by one or two orders of magnitude. They are:

- *Reactor VLB*: reactor experiments with $L \sim 100 \text{ km}$, corresponding to sensitivities

$$L/E \lesssim 10^5 \text{ m/MeV} \Rightarrow \Delta m^2 \gtrsim 10^{-5} \text{ eV}^2. \quad (1.1.76)$$

- *Accelerator VLBs*: experiments in which accelerator neutrinos propagate for distances $L \gtrsim 10^3 \text{ km}$, covering a sensitivity range of

$$L/E \lesssim 10^4 \text{ km/GeV} \Rightarrow \Delta m^2 \gtrsim 10^{-4} \text{ eV}^2. \quad (1.1.77)$$

- *Solar neutrino experiments*: experiments which detect neutrinos generated by the thermonuclear reaction processes in the core of the Sun. Since the energy of detectable solar neutrinos covers the 0.2 - 15 MeV and the Sun-Earth distance

is of about 1.5×10^{11} m, the theoretical sensitivity to Δm^2 of solar neutrino experiments is

$$L/E \lesssim 10^{12} \text{ m/MeV} \Rightarrow \Delta m^2 \gtrsim 10^{-12} \text{ eV}^2, \quad (1.1.78)$$

much smaller than that of any other experimental technique presented above.

Experimental results in the two-neutrino framework

By the end of the first period in the history of neutrino oscillation experiments, corresponding to the results of Super-Kamiokande and SNO, two features of neutrinos had been established: that (a) neutrinos have small non-zero masses, so that flavour neutrinos are combinations of mass eigenstates, and (b) that all oscillation data can be interpreted assuming a three-neutrino mixing model [2]. In fact, the data obtained in the first solar and atmospheric neutrino experiments could be modelled, in a two-neutrino approximation, according to two separate sets of oscillation parameters: a solar mixing angle $\theta_\odot$ and mass difference $\Delta m_\odot^2$ and an atmospheric mixing angle θ_{ATM} and mass difference Δm_{ATM}^2 . This clearly implied the existence of three massive neutrino states.

This section will present the original results on the solar and atmospheric oscillation parameters, followed by a discussion of the full experimental scenario within the three-neutrino mixing framework.

Solar oscillation parameters. Despite having an extremely large flux on Earth of $\sim 6 \times 10^{10} \text{ cm}^{-2}\text{s}^{-1}$, solar neutrinos are particularly difficult to detect due to their low energies of $\sim 1 \text{ MeV}$, requiring large underground detectors to compensate for the small interaction cross sections and to mitigate the cosmic ray background.

Solar neutrinos were first detected in 1970 by the Homestake Mine radiochemical experiment, which found a deficit in the observed ν_e flux compared to the prediction of the standard model of solar thermonuclear reactions (SSM) [7]. This deficit, known as *Solar neutrino problem*, was confirmed by later radiochemical experiments such as GALLEX [8], GNO [9] and SAGE [10], as well as by the Kamiokande [11] and the later Super-Kamiokande [12] Cherenkov detectors. The latest updated measurement of the ${}^8\text{Be}$ solar neutrino flux from Super-Kamiokande is [13]:

$$\Phi({}^8\text{Be}) = [2.314 \pm 0.014 \text{ (stat.)} \pm 0.040 \text{ (sys.)}] \times 10^6 \text{ cm}^{-2}\text{s}^{-1}. \quad (1.1.79)$$

Neither the radiochemical experiments nor the Cherenkov detectors could discriminate a deficit produced by neutrino oscillations from an incorrect prediction of the SSM. The decisive measurement for the interpretation of the Solar neutrino problem was performed by the Sudbury Neutrino Observatory (SNO) [14]. Consisting of a Cherenkov detector located 2 km underground and containing 1 kt of heavy water, SNO is sensitive to three separate channels

$$\begin{aligned} \text{CC:} \quad & \nu_e + d \rightarrow p + p + e^-, \\ \text{NC:} \quad & \nu_\alpha \rightarrow p + n + n + \nu_\alpha, \\ \text{ES:} \quad & \nu_\alpha + e^- \rightarrow \nu_\alpha + e^-, \end{aligned} \quad (1.1.80)$$

where the Neutral Current channel is equally sensitive to all three active flavours. SNO could therefore separately measure the total flux of ^8Be solar neutrinos [15]

$$\Phi_{^8\text{Be}}^{\text{NC}}(\nu_{\text{active}}) = [5.54_{-0.31}^{+0.33} (\text{stat.})_{-0.34}^{+0.36} (\text{sys.})] \times 10^6 \text{ cm}^{-2} \text{ s}^{-1} \quad (1.1.81)$$

in agreement with the SSM, and the flux of ^8Be ν_e in the CC channel:

$$\Phi_{^8\text{Be}}^{\text{CC}}(\nu_e) = [1.67_{-0.04}^{+0.05} (\text{stat.})_{-0.08}^{+0.07} (\text{sys.})] \times 10^6 \text{ cm}^{-2} \text{ s}^{-1}. \quad (1.1.82)$$

The results from SNO proved that the total solar neutrino flux was consistent with the solar models, so that the deficit in the flux of ν_e required neutrino oscillations to be explained.

As anticipated, solar neutrino oscillation data can be modelled in a first order approximation as resulting from two-neutrino mixing with the effective parameters $\Delta m_{\odot}^2$ and $\theta_{\odot}$, since the mixing of the third mass eigenstate into ν_e is small. A global fit that included data from SNO and several other solar neutrino experiments together with the KamLAND reactor experiment [16] yielded:

$$\Delta m_{\odot}^2 = 7.46_{-0.19}^{+0.20} \times 10^{-5} \text{ eV}^2, \quad \tan^2 \theta_{\odot} = 0.427_{-0.024}^{+0.027}, \quad (1.1.83)$$

thus corresponding to a large but non maximal mixing angle [17]. The corresponding two-flavour oscillation contours are shown in fig. 1.2.




Figure 1.2: Two-flavour oscillation analysis contours for (a) SNO data up to Phase-III, (b) a global fit including also measurements from Homestake, Gallex, GNO, SAGE, Borexino [18], SuperKamiokande up to Phase-III, (c) KamLAND data [16], (d) a global fit of both the solar neutrino experiments and KamLAND. The best fit value for the solar-only global fit is $\Delta m_{\odot}^2 = 5.13_{-0.96}^{+1.29} \times 10^{-5} \text{ eV}^2$, $\tan^2 \theta_{\odot} = 0.427_{-0.028}^{+0.028}$ [17].

The sign of the solar squared mass difference could be determined as a result of the *Mikheyev-Smirnov-Wolfenstein* (MSW) effect, which introduces a dependence on Δm^2 in the oscillation probabilities due to the interaction of neutrinos inside the solar

volume. Since the Sun produces ν_e in its core, the propagation can be modelled as that of a two-neutrino ν_e - ν_X system, where X is an arbitrary superposition of μ and τ , as ν_μ and ν_τ have only equal neutral interactions in ordinary matter.

The density of solar matter is a monotonically decreasing function of the distance from the centre of the Sun R , and it can be approximated by an exponential for $0.1 \lesssim R/R_\odot \lesssim 0.9$:

$$n_e(R) = n_e(0) \exp(-R/r_0), \quad (1.1.84)$$

where $r_0 = R_\odot/10.54 = 6.6 \times 10^7$ m and $n_e(0) = 245 N_A/\text{cm}^3$ [3]. The composition of the neutrino state exiting the Sun will thus depend on the relative size of the interaction potential at the production point $A_0 \simeq 2\sqrt{2}EG_F n_e(R=0)$ compared to the resonance potential $A_{CC}^R = \Delta m^2 \cos 2\theta$, from eqn. (1.1.61). If, in particular, $0 < A_{CC}^R < A_0$, neutrinos may cross the resonance on their way out, since $n_e(R)$ decreases farther from the centre.

Indeed, the combined measurements of the solar neutrino flux and reactor-neutrino data from KamLAND [19] are consistent with such a crossing for $E_\nu \gtrsim 10$ MeV, constraining, through eqn. (1.1.61), $\Delta m^2 \cos 2\theta$ to be positive. The non-maximal value of $\theta_\odot$, thus implies $\Delta m_{12} > 0$ [3, 20, 21].

Atmospheric oscillation parameters. Atmospheric neutrinos are a product of the interactions of primary cosmic rays (mainly protons with a small fraction of heavier nuclei) with nuclei in the atmosphere. These interactions generate secondary cosmic rays, consisting of hadrons and their decay products with a spectrum peaking in the GeV range. The large quantity of secondary pions, in particular, decays mainly into muons and ν_μ :

$$\pi^+ \rightarrow \mu^+ + \nu_\mu, \quad \pi^- \rightarrow \mu^- + \bar{\nu}_\mu, \quad (1.1.85)$$

with kaons also contributing to the muon and neutrino production at higher energies. Muons will in turn decay to electrons, ν_e and ν_μ before reaching the surface:

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu, \quad \mu^- \rightarrow e^- + \bar{\nu}_e + \nu_\mu. \quad (1.1.86)$$

The *atmospheric neutrinos* produced in these reactions, with energies between about 100 MeV and 100 GeV, can be detected via scattering reactions on nuclei.

The first detections of atmospheric neutrinos were performed by in the second half of the 1980s by the first large underground water-Cherenkov experiments, Kamiokande [22] and IMB [23]. Being capable of discriminating muons from electrons using the information from Cherenkov ring patterns, these experiments could measure the flux ratio $\nu_\mu/\nu_e \equiv (\nu_\mu + \bar{\nu}_\mu)/(\nu_e + \bar{\nu}_e)$, observing, however, a deficit in the $(\nu_\mu + \bar{\nu}_\mu)$ compared to the expectation from the cosmic ray spectrum [22, 23].

The follow-up experiment to Kamiokande, Super-Kamiokande (SK), brought about the solution of this so-called *atmospheric neutrino anomaly*, observing a dependence of the distribution of muon neutrino events on the zenith-angle, a compelling evidence of atmospheric neutrino oscillations [24]. Over the years, Super-Kamiokande has provided high-statistics data on the flux of atmospheric neutrinos, allowing to determine that atmospheric oscillations are predominantly between ν_μ and ν_τ , and to precisely measure the oscillation parameters. The SK results were corroborated by

numerous experiments, such as Soudan-2 [25] and MACRO [26] as the K2K Long-Baseline accelerator experiment [27]. A definitive proof of $\nu_\mu \rightarrow \nu_\tau$ oscillations was provided by the observation of ν_τ appearance by the OPERA long-baseline experiment [28]. More recently IceCube has also reported ν_τ appearance in atmospheric neutrinos [29].

As was the case for solar oscillations, atmospheric neutrino oscillations could be effectively described in a two-neutrino approximation. The combined analysis of data collected over the first phase of SK operations [30] yielded the following 90% CL ranges for the oscillation parameters, as shown in fig. 1.3:

$$\sin^2 2\theta_{ATM} > 0.92, \quad 1.5 \times 10^{-3} < \Delta m_{ATM}^2 < 3.4 \times 10^{-3} \text{ eV}^2, \quad (1.1.87)$$

with the best fit values being

$$\sin^2 2\theta_{ATM} = 1.00, \quad \Delta m_{ATM}^2 = 2.1 \times 10^{-3} \text{ eV}^2. \quad (1.1.88)$$




Figure 1.3: SK 90% CL allowed parameter regions for atmospheric $\nu_\mu \rightarrow \nu_\tau$ oscillations from six sub-samples with different topologies. The allowed region from the combined fit is shown as the bold contour [30].

Experimental results in the three-neutrino framework

The minimal scenario capable of accommodating the two independent squared mass differences $\Delta m_\odot^2$ and Δm_{ATM}^2 is the three-neutrino mixing framework, introduced in sect. 1.1.2. The neutrino mass eigenstates can be labelled arbitrarily in order to have

$$\Delta m_\odot^2 \equiv \Delta m_{21}^2, \quad \Delta m_{ATM}^2 \equiv |\Delta m_{31}^2|, \quad (1.1.89)$$

so that, according to equations (1.1.83) and (1.1.88)

$$\Delta m_{21}^2 \ll \Delta m_{31}^2 \simeq \Delta m_{32}^2. \quad (1.1.90)$$

It thus follows that either $|\Delta m_{31}^2|$ or $|\Delta m_{32}^2|$ can be taken as the atmospheric squared-mass difference, since atmospheric and LBL experiments are not sensitive to $\Delta m_{\odot}^2$. Conversely, solar and VLBL experiments are sensitive to $\Delta m_{\odot}^2$, but observe only the averaged probabilities of the transitions generated by $|\Delta m_{31}^2|$ and $|\Delta m_{32}^2|$. A distinction between $|\Delta m_{31}^2|$ and $|\Delta m_{32}^2|$ is only necessary for experiments sensitive to subleading oscillations due to $\Delta m_{\odot}^2$, as is the case of some LBL experiments.

While the observed squared-mass differences are independent of the mixing framework used in the analysis, the inferred mixing parameters could change significantly when replacing separate two-neutrino analyses with a full three-flavour mixing framework, since the PMNS matrix introduces more degrees of freedom than a mere doubling of two-neutrino mixing. In practice, however, this is not the case, as $|U_{e3}|$, the only PMNS element which enters in both solar and atmospheric oscillations, is small.

Indeed, solar neutrino experiments are sensitive to the ν_e -disappearance channel, so that their oscillation probabilities only depend on the first row $|U_{ej}|$ of the PMNS. At the same time, the probabilities for atmospheric oscillations depend on the θ_{13} and θ_{23} angles which depend on the elements of the third column of the PMNS, as

$$\sin \theta_{23} = \frac{|U_{\mu 3}|}{\sqrt{1 - |U_{e3}|^2}}, \quad \sin \theta_{13} = |U_{e3}|. \quad (1.1.91)$$

Solar and atmospheric oscillations are therefore only correlated by $|U_{e3}|$ [1]. The measurement of θ_{13} has been, historically, the next critical step in the understanding of the PMNS, as a non-zero value of θ_{13} is required both for CP-violation to be observable and for determining the neutrino mass ordering. The following paragraphs outline the experimental constraints on θ_{13} and examine the channels through which the remaining parameters of the PMNS matrix can be measured.

Measurement of θ_{13} . An almost pure measurement of $\sin^2 2\theta_{13}$ is possible in Medium Baseline reactor neutrino experiments, with $L \sim 1$ km, from the disappearance of $\bar{\nu}_e$. For such a baseline, the ν_e survival probability can be approximated as [3]

$$P_{\nu_e \rightarrow \nu_e} \simeq 1 - \sin^2 2\theta_{13} \sin^2 \left(\frac{\Delta m_{ee}^2 L}{4E_\nu} \right) \quad (1.1.92)$$

where

$$\Delta m_{ee}^2 \equiv \cos^2 \theta_{12} \Delta m_{31}^2 + \sin^2 \theta_{12} \Delta m_{32}^2. \quad (1.1.93)$$

The first experiments to measure the θ_{13} angle were Double-Chooz, Daya-Bay and RENO, all of which were based around liquid scintillator detectors. The best fit values for $\sin^2 2\theta_{13}$ from the first phases of data-taking are shown in Table 1.1.

The small value of the θ_{13} angle thus justifies the independence, at the leading order, of solar and atmospheric oscillations, with the remaining PMNS angles being approximately

$$\theta_{12} \simeq \theta_{\odot}, \quad \theta_{23} \simeq \theta_{ATM}. \quad (1.1.94)$$

Measurement of the θ_{23} octant. The current global fits of 3ν oscillation parameters are not capable of resolving whether θ_{23} is below 45° (*first octant*) or above it (*second octant*) [3]. Indeed, the ν_μ -disappearance and ν_e -appearance probabilities at

θ_{13} measurements	
Daya Bay [31]	$\sin^2 2\theta_{13} = 0.092 \pm 0.016(\text{stat}) \pm 0.005(\text{syst})$
RENO [32]	$\sin^2 2\theta_{13} = 0.113 \pm 0.013(\text{stat}) \pm 0.019(\text{syst})$
Double-Chooz [33]	$\sin^2 2\theta_{13} = 0.086 \pm 0.041(\text{stat}) \pm 0.030(\text{syst})$

Table 1.1: $\sin^2 2\theta_{13}$ PMNS parameters obtained from the Daya Bay, RENO and Double-Chooz experiments. The most stringent constraints on the mixing angle were those from Daya Bay, with a 5.2σ significance on a non-zero θ_{13} [3, 31].

long-baseline atmospheric and accelerator neutrino experiments, which are sensitive to θ_{23} , can be approximated, using $\Delta m_{21}^2/\Delta m_{31}^2 \ll 1/30 \ll 1$, as:

$$\begin{aligned} P(\nu_\mu \rightarrow \nu_\mu) &\simeq 1 - \sin^2 2\theta_{23} \sin^2 \left(\frac{\Delta m_{31}^2 L}{4E} \right) \\ P(\nu_\mu \rightarrow \nu_e) &\simeq \sin^2 2\theta_{23} \sin^2 2\theta_{13} \sin^2 \left(\frac{\Delta m_{31}^2 L}{4E} \right), \end{aligned} \quad (1.1.95)$$

neglecting matter and δ_{CP} -related terms. While both Δm_{31}^2 and $\sin^2 2\theta_{23}$ can be accurately determined through the disappearance channel, it can be noticed that for each $\theta_{23} \neq \pi/4$ there is a twofold solution of θ_{23} for any value of $\sin^2 2\theta_{23}$:

$$\sin^2 \theta_{23} = \frac{1}{2} \left[1 \pm \sqrt{1 - \sin^2 2\theta_{23}} \right], \quad (1.1.96)$$

leading to the aforementioned *octant degeneracy*.

Lifting the θ_{23} octant degeneracy will require the combination of ν_μ -disappearance and ν_e -appearance data with independent measurements of θ_{13} from reactor experiments [34]. As of 2024 [35], global analyses do not find a statistically significant preference for the octant of θ_{23} , with 3σ confidence intervals ranging between $41.3^\circ \rightarrow 49.9^\circ$ and $41.5^\circ \rightarrow 49.8^\circ$ under the assumption of Normal and Inverted ordering respectively. The current status of the maximality/non-maximality of θ_{23} is therefore not conclusive.

Determination of the Mass Ordering The precise measurement of several PMNS parameters, such as θ_{23} and δ_{CP} , hinges on resolving the *Mass Ordering* of the neutrino mass spectrum, with *Normal Ordering* (NO)

$$m_1 < m_2 < m_3 \Rightarrow \Delta m_{31}^2 \simeq \Delta m_{32}^2 > 0, \quad \Delta m_{21}^2 > 0 \quad (1.1.97)$$

and *Inverted Ordering* (IO)

$$m_3 < m_1 < m_2 \Rightarrow \Delta m_{31}^2 \simeq \Delta m_{32}^2 < 0, \quad \Delta m_{21}^2 > 0, \quad (1.1.98)$$

being the two non-equivalent possibilities, as shown in fig. 1.4.

Current experiments are not capable of sufficient individual sensitivities to enable discrimination [3], and global fits are not yet statistically significant [35], as sect. 1.1.3 will discuss. In order to determine the mass ordering, the next generation of experiments will therefore leverage three different experimental configurations [37]:




Figure 1.4: Neutrino mass eigenstate flavour composition and mass pattern in the two cases of Normal Ordering (left) and Inverted Ordering (right) [36].

- Medium-baseline ($L \sim 50 \text{ km}$) reactor $\bar{\nu}_e \rightarrow \bar{\nu}_e$ experiments, such as JUNO [38].
- Long-baseline accelerator muon (anti-)neutrino $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ experiments including the current NO ν A [39] and the future DUNE [40] and Hyper-K [41].
- Atmospheric neutrino experiments with the same oscillation channel as that above, including PINGU [42], ORCA [43] as well as Hyper-K [41] and DUNE [40].

The discrimination of Mass Ordering in medium baseline experiments relies on the interference in the disappearance probability caused by the difference between Δm_{31}^2 and Δm_{32}^2 :

$$\begin{aligned}
 P_{\bar{\nu}_e \rightarrow \bar{\nu}_e} \simeq & 1 - \cos^4 \theta_{13} \sin^2 2\theta_{12} \sin^2 \left(\frac{\Delta m_{21}^2 L}{4E} \right) \\
 & - \sin^2 2\theta_{13} \sin^2 \left(\frac{\Delta m_{31}^2 L}{4E} \right) \\
 & - \sin^2 \theta_{12} \sin^2 2\theta_{13} \sin^2 \left(\frac{\Delta m_{21}^2 L}{4E} \right) \cos \left(\frac{2|\Delta m_{31}^2| L}{4E} \right) \\
 & \pm \frac{\sin^2 \theta_{12}}{2} \sin^2 2\theta_{13} \sin \left(\frac{2\Delta m_{21}^2 L}{4E} \right) \sin \left(\frac{2|\Delta m_{31}^2| L}{4E} \right),
 \end{aligned} \tag{1.1.99}$$

where the $\pm$ at the fourth term distinguished the effect for normal and inverted ordering. As discussed in sect. 1.1.3, a nonzero value of θ_{13} is essential for the effect to be observable. The probability in eqn. (1.1.99) does not depend on the δ_{CP} phase, and the Matter Effect can be neglected at the $L \sim 50 \text{ km}$ distances characteristic of medium-baseline experiments. The modulation expected at JUNO on the oscillated $\bar{\nu}_e$ energy spectra resulting from the two possible orderings is illustrated in fig. 1.5 as a function of L/E , compared to the unoscillated flux.

Determining the mass ordering through the medium-baseline channel is the primary aim of the JUNO experiment in China, whose construction was completed in August 2025 [44]. Consisting of a 20 kt liquid scintillator detector located at $\sim 53 \text{ km}$ from two nuclear power plants so as to maximise the flux modulation, JUNO will leverage an excellent energy resolution and a large fiducial volume [38].




Figure 1.5: Modulation in the JUNO reactor $\bar{\nu}_e$ oscillated fluxes in the two cases of Normal and Inverted Ordering (here referred to as hierarchies). The unoscillated flux is shown as a reference [37].

The alternative experimental handle for the determination of mass ordering is based on the *matter effect* on the propagation of neutrinos through the Earth's crust over Long Baselines ($L \sim 10^2 - 10^3$ km) in the $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ channels. Assuming a constant matter density, the ν_e -appearance probability at LBL experiments can be expanded, at the second order, in the small parameters $\sin \theta_{13}$ and $\alpha \equiv \Delta m_{21}^2 / \Delta m_{31}^2$:

$$\begin{aligned}
 P_{\nu_\mu \rightarrow \nu_e, (\bar{\nu}_\mu \rightarrow \bar{\nu}_e)} &\simeq 4 \sin^2 \theta_{13} \sin^2 \theta_{23} \frac{\sin^2 \Delta (1 - A)}{(1 - A)^2} \\
 &+ \alpha^2 \sin^2 2\theta_{12} \cos^2 \theta_{23} \frac{\sin^2 A \Delta}{A^2} \\
 &+ 8\alpha J_{CP}^{max} \cos(\Delta \pm \delta_{CP}) \frac{\sin \Delta A \sin \Delta (1 - A)}{A (1 - A)},
 \end{aligned} \tag{1.1.100}$$

with

$$J_{CP}^{max} = \cos \theta_{12} \sin \theta_{12} \cos \theta_{23} \sin \theta_{23} \cos^2 \theta_{13} \sin \theta_{13}, \tag{1.1.101}$$

and

$$\Delta \equiv \frac{\Delta m_{31}^2 L}{4E_\nu}, \quad A \equiv \frac{2E_\nu V}{\Delta m_{31}^2}, \tag{1.1.102}$$

V being the effective matter potential in the Earth's crust and the $(-)+$ in the third term refers to the (anti)neutrino channel. The ν_e -appearance channel at LBL is therefore sensitive to the sign of Δm_{31}^2 through the α , Δ and A terms, and thus to Mass Ordering. The dependence of LBL oscillations on δ_{CP} , from eqn. (1.1.100), furthermore provides the main handle on leptonic CP-violation. Determining the mass ordering is thus crucial also for measuring δ_{CP} , as both factors appear in the same expression, leading to degenerate solutions for $P_{\nu_\mu \rightarrow \nu_e, (\bar{\nu}_\mu \rightarrow \bar{\nu}_e)}$. Independent measurements such as

those from medium baseline experiments are thus important [38]. Figure 1.6 shows the appearance probabilities for ν_e and $\bar{\nu}_e$ at the $L = 1300$ km baseline of the future DUNE experiment, for both mass orderings and for several values of δ_{CP} .




Figure 1.6: Electron neutrino and antineutrino appearance probabilities as a function of E_ν at the DUNE baseline of $L = 1300$ km for four selected values of δ_{CP} . The top (bottom) two panels correspond to the NO (IO), while the left (right) panels refer to ν_e ($\bar{\nu}_e$). In NO (IO), the appearance of ν_e ($\bar{\nu}_e$) is enhanced, while that of $\bar{\nu}_e$ (ν_e) is suppressed [45].

Measurement of δ_{CP} . As discussed in sect. 1.1.2, neutrino flavour oscillations provide an handle for observing CP-violation in the lepton sector. Since all mixing angles were determined to have non-zero values, the CP-violating phase δ_{CP} may be experimentally accessible.

Referring to eqn. (1.1.37), CP-invariance can be tested by searching for an asymmetry between the neutrino and antineutrino channels $P_{\nu_\alpha \rightarrow \nu_\beta}$ and $P_{\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta}$. LBL experiments are particularly suited to such measurements through their $\bar{\nu}_e$ -appearance and $\bar{\nu}_\mu$ -disappearance channels. The dependence of the oscillation probabilities on δ_{CP} in LBL experiments is then given by eqn. (1.1.100), which, neglecting the matter effect, can be written as:

$$\begin{aligned}
 P_{\nu_\mu \rightarrow \nu_e}(\bar{\nu}_\mu \rightarrow \bar{\nu}_e) &\simeq \sin^2 \theta_{23} \sin^2 2\theta_{13} \sin^2 \Delta \\
 &\pm \frac{\sin 2\theta_{12} \sin 2\theta_{23}}{2 \sin \theta_{13}} \sin \frac{\Delta m_{21}^2}{4E} \sin^2 2\theta_{13} \sin^2 \Delta \sin \delta_{CP} \\
 &+ (\text{CP-conserving term, solar term, matter term})
 \end{aligned} \tag{1.1.103}$$

where Δ has been defined in eqn. (1.1.102) and the $+$ ($-$) at the second term refers to the (anti)neutrino channel [46]. The CP-violation asymmetry

$$A_{CP} \equiv \frac{P(\nu_\mu \rightarrow \nu_e) - P(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)}{P(\nu_\mu \rightarrow \nu_e) + P(\bar{\nu}_\mu \rightarrow \bar{\nu}_e)} \tag{1.1.104}$$

thus becomes, to the leading order in Δm_{12}^2 :

$$A_{CP} \simeq \frac{\cos \theta_{21} \sin 2\theta_{12} \sin \delta_{CP}}{\sin \theta_{23} \sin \theta_{13}} \left(\frac{\Delta m_{12}^2 L}{4E} \right) + \text{matter terms.} \quad (1.1.105)$$

For a fixed E_ν and θ_{13} , therefore, the asymmetry grows linearly with L [47]. As discussed in the following, the current global fits of LBL oscillation data do not allow a statistically significant determination of δ_{CP} [35].

Global fits

The global analysis of neutrino oscillation data provides the comprehensive description of the structure of three-neutrino flavour mixing. Over the last decade, global analyses have provided consistent results with good agreement on the dominant effects in oscillations. Conversely, subdominant or quantitatively small effects in oscillations, such as the mass ordering, the θ_{23} octant and the leptonic CP-violation phase δ_{CP} , remain open questions in the current experimental landscape, and represent the main goal of the upcoming generation of experiments.

The most recent global analysis at the time of writing is provided by the NuFIT program [35], which incorporates updated solar neutrino models [48] and data from the latest runs of multiple experiments, spanning the solar neutrino sector (SK and Borexino), reactor-based LBL and MBL experiments (Daya-Bay and SNO+), as well as accelerator-based LBL and atmospheric neutrino experiments (SK, NOvA and IceCube). Updated χ^2 -maps provided by the collaborations have been included in the analysis.

The global fit results for selected oscillation parameters are shown in fig. 1.7 as one-dimensional $\Delta\chi^2$ curves, for two versions of the analysis, respectively including and excluding χ^2 -maps from atmospheric neutrino data in Super-Kamiokande and IceCube/DeepCore. Figure 1.8 furthermore displays the two-dimensional projections of the allowed parameter regions. The best fit values of the oscillation parameters, with the corresponding 1σ and 3σ intervals, are given in Tab. 1.9 in for both mass orderings [35].

The main results that can be gathered from the latest global fit are the following [35]:

- The determination of the θ_{12} , θ_{13} , Δm_{21}^2 and $|\Delta m_{3\ell}^2|$ parameters is very stable, with relative precisions at 3σ of about 13%, 8%, 16%, 5-6% respectively.
- The 3σ precision for θ_{23} is still about 20% with the determination still suffering from the octant ambiguity. Despite a slight preference for $\theta_{23} > 45^\circ$ (except for NO and the analysis including SK-atmospheric data), the local minimum for the first octant has always $\Delta\chi^2 < 4$ for all combinations of datasets and mass ordering.
- The determination of δ_{CP} has a strong dependence on the mass ordering. For NO the best-fit point is close to the CP-conserving value of 180° with $\Delta\chi^2 < 1$. For IO, conversely, the best fit point for both dataset variants are close to

the maximal CP-violation $\delta_{CP} = 270^\circ$ within 1σ , with CP-conservation being disfavoured at 3.6σ (4σ) for the analysis excluding (including) SK atmospheric neutrino data.

- For what concerns the mass ordering, $\Delta\chi^2_{\text{IO,NO}} = -0.6$ (6.1) is found respectively in the analysis excluding (including) SK atmospheric neutrino data. The indecisive result for the former case is related to opposite trends in the LBL-accelerator appearance data from T2K and NOvA and the combination of disappearance data from accelerator and reactor experiments. This corresponds to a 1.7σ (1.4σ) median sensitivity to reject NO (IO). The addition of SK atmospheric data provides instead additional preference for NO, leading to the $\Delta\chi^2_{\text{IO,NO}} = 6.1$ quoted above.

The amount of CP-violation in neutrino oscillations can be expressed, in a rephasing-invariant way, in terms of the Jarlskog Invariant (see sect. 1.1.2):

$$\begin{aligned} J_{CP} &\equiv \text{Im} [U_{\alpha i} U_{\alpha j}^* U_{\beta i}^* U_{\beta j}] \\ &\equiv J_{CP}^{max} \sin \delta_{CP} = \cos \theta_{12} \sin \theta_{12} \cos \theta_{23} \sin \theta_{23} \cos^2 \theta_{13} \sin \theta_{13} \sin \delta_{CP}. \end{aligned} \quad (1.1.106)$$

The current global fit finds, for its maximum value J_{CP}^{max}

$$J_{CP}^{max} = 0.0333 \pm 0.0007 (\pm 0.0017), \quad (1.1.107)$$

at 1σ (3σ) for both orderings. For NO, the best fit value excluding (including) SK atmospheric data $J_{CP}^{best} = 0.0017$ (-0.018) is only favoured over CP-conservation with $\Delta\chi^2 = 0.02$ (0.55). By contrast, in IO CP-conservation is disfavoured with respect to $J_{CP}^{best} = -0.032$ with $\Delta\chi^2 = 13(16)$, corresponding to $3.6\sigma(4\sigma)$ at 1 d.o.f. [35].




Figure 1.7: $\Delta\chi^2$ profiles for the global 3ν oscillation analysis. The profiles are minimised with respect to all undisplayed parameters. The red (blue) curves correspond to Normal (Inverted) Ordering. The solid (dashed) lines correspond to the analysis excluding (including) atmospheric neutrino data from Super-Kamiokande [35].




Figure 1.8: Two-dimensional projections of the allowed region for the global 3ν oscillation analysis. The profiles are minimised with respect to all undisplayed parameters. The four lower panels correspond to a $\Delta\chi^2$ minimised with respect to Mass Ordering. The different contours refer to the 1σ, 90%, 2σ, 99%, 3σ CL, with 2 d.o.f.. Coloured regions (black contours) correspond to the analysis variant that excludes (includes) atmospheric neutrino data from Super-Kamiokande [35].

		Normal Ordering ($\Delta\chi^2 = 0.6$)		Inverted Ordering (best fit)	
		bfp $\pm 1\sigma$	3σ range	bfp $\pm 1\sigma$	3σ range
IC19 without SK atmospheric data	$\sin^2 \theta_{12}$	$0.307^{+0.012}_{-0.011}$	$0.275 \rightarrow 0.345$	$0.308^{+0.012}_{-0.011}$	$0.275 \rightarrow 0.345$
	$\theta_{12}/^\circ$	$33.68^{+0.73}_{-0.70}$	$31.63 \rightarrow 35.95$	$33.68^{+0.73}_{-0.70}$	$31.63 \rightarrow 35.95$
	$\sin^2 \theta_{23}$	$0.561^{+0.012}_{-0.015}$	$0.430 \rightarrow 0.596$	$0.562^{+0.012}_{-0.015}$	$0.437 \rightarrow 0.597$
	$\theta_{23}/^\circ$	$48.5^{+0.7}_{-0.9}$	$41.0 \rightarrow 50.5$	$48.6^{+0.7}_{-0.9}$	$41.4 \rightarrow 50.6$
	$\sin^2 \theta_{13}$	$0.02195^{+0.00054}_{-0.00058}$	$0.02023 \rightarrow 0.02376$	$0.02224^{+0.00056}_{-0.00057}$	$0.02053 \rightarrow 0.02397$
	$\theta_{13}/^\circ$	$8.52^{+0.11}_{-0.11}$	$8.18 \rightarrow 8.87$	$8.58^{+0.11}_{-0.11}$	$8.24 \rightarrow 8.91$
	$\delta_{\text{CP}}/^\circ$	177^{+19}_{-20}	$96 \rightarrow 422$	285^{+25}_{-28}	$201 \rightarrow 348$
	$\frac{\Delta m_{21}^2}{10^{-5} \text{ eV}^2}$	$7.49^{+0.19}_{-0.19}$	$6.92 \rightarrow 8.05$	$7.49^{+0.19}_{-0.19}$	$6.92 \rightarrow 8.05$
	$\frac{\Delta m_{3\ell}^2}{10^{-3} \text{ eV}^2}$	$+2.534^{+0.025}_{-0.023}$	$+2.463 \rightarrow +2.606$	$-2.510^{+0.024}_{-0.025}$	$-2.584 \rightarrow -2.438$
IC24 with SK atmospheric data		Normal Ordering (best fit)		Inverted Ordering ($\Delta\chi^2 = 6.1$)	
		bfp $\pm 1\sigma$	3σ range	bfp $\pm 1\sigma$	3σ range
	$\sin^2 \theta_{12}$	$0.308^{+0.012}_{-0.011}$	$0.275 \rightarrow 0.345$	$0.308^{+0.012}_{-0.011}$	$0.275 \rightarrow 0.345$
	$\theta_{12}/^\circ$	$33.68^{+0.73}_{-0.70}$	$31.63 \rightarrow 35.95$	$33.68^{+0.73}_{-0.70}$	$31.63 \rightarrow 35.95$
	$\sin^2 \theta_{23}$	$0.470^{+0.017}_{-0.013}$	$0.435 \rightarrow 0.585$	$0.550^{+0.012}_{-0.015}$	$0.440 \rightarrow 0.584$
	$\theta_{23}/^\circ$	$43.3^{+1.0}_{-0.8}$	$41.3 \rightarrow 49.9$	$47.9^{+0.7}_{-0.9}$	$41.5 \rightarrow 49.8$
	$\sin^2 \theta_{13}$	$0.02215^{+0.00056}_{-0.00058}$	$0.02030 \rightarrow 0.02388$	$0.02231^{+0.00056}_{-0.00056}$	$0.02060 \rightarrow 0.02409$
	$\theta_{13}/^\circ$	$8.56^{+0.11}_{-0.11}$	$8.19 \rightarrow 8.89$	$8.59^{+0.11}_{-0.11}$	$8.25 \rightarrow 8.93$
	$\delta_{\text{CP}}/^\circ$	212^{+26}_{-41}	$124 \rightarrow 364$	274^{+22}_{-25}	$201 \rightarrow 335$
	$\frac{\Delta m_{21}^2}{10^{-5} \text{ eV}^2}$	$7.49^{+0.19}_{-0.19}$	$6.92 \rightarrow 8.05$	$7.49^{+0.19}_{-0.19}$	$6.92 \rightarrow 8.05$
	$\frac{\Delta m_{3\ell}^2}{10^{-3} \text{ eV}^2}$	$+2.513^{+0.021}_{-0.019}$	$+2.451 \rightarrow +2.578$	$-2.484^{+0.020}_{-0.020}$	$-2.547 \rightarrow -2.421$

Figure 1.9: Three-neutrino flavour oscillation parameters from the two global fit variants described in the text. The values in the first and second column are obtained assuming respectively NO and IO, with $\Delta m_{3\ell}^2 \equiv \Delta m_{31}^2 > 0$ for NO and $\Delta m_{3\ell}^2 \equiv \Delta m_{32}^2 < 0$ for IO [35].

1.2 Heavy Neutral Leptons

As shown in the previous section, neutrinos have proven particularly valuable in elucidating the structure of the Standard Model of particle physics. The observation of three neutrino-flavour oscillations in a wide variety of experimental settings has demonstrated the existence of non-zero neutrino masses and lepton flavour mixing, marking the first confirmed evidence of physics beyond the original formulation of the Standard Model.

While the framework of the SM describes how neutrinos interact with charged leptons and quarks through weak interactions, it leaves several fundamental questions about the nature of neutrinos unanswered. Chief among these is perhaps the question of why neutrino masses, while non-zero, are so small compared to those of the other fermions. A simple potential explanation for this is the Seesaw-Mechanism (see sect. 1.2.1), which introduces multiple electroweak-singlet (or *sterile*) neutrino interaction eigenstates. The Seesaw-Mechanism does not specify the number of singlet neutrinos and, importantly, is viable at both high-scales and low-scales. The low-scale scenario is particularly attractive as the additional neutrino mass eigenstates, commonly referred to as *Heavy Neutral Leptons* (HNLs), might be detectable in current or near future experiments [49]. Together with the origin of neutrino masses, searches for HNLs in the keV to GeV scale can probe a wide range of phenomenological behaviours [50], from Dark Matter [51] to leptogenesis [52].

This section provides an overview of the theoretical motivation and phenomenology of HNL models and discusses their present experimental status and the prospects for future searches at upcoming experiments.

1.2.1 Heavy neutrinos in seesaw models

The smallness of neutrino masses relative to other fermions and the markedly different structure of leptonic mixing compared to that of quarks both suggest a different origin for neutrino masses other than the Higgs mechanism. A key ingredient of many models are right handed neutrinos, referred to as HNLs for masses $\gg$ eV.

The simplest extension of the SM to include a neutrino mass term can be constructed, as detailed in sect. 1.1.1, by adding gauge singlets ν_R , typically called *sterile neutrinos*, and constructing the Yukawa lagrangian in eqn. (1.1.2), which, after electroweak symmetry breaking, yields the neutrino Dirac mass terms of eqn. (1.1.9). The Yukawa lagrangian and the resulting mass term conserve the lepton number. Indeed, while the ν_L Majorana mass term in eqn. (1.1.40) is not allowed in the SM due to ν_L having $I = 1/2$, the corresponding term for ν_R is allowed and its absence must be imposed by requiring lepton number conservation.

Seesaw mechanisms

While a pure Yukawa Lagrangian provides a viable explanation for neutrino masses, it presents several shortcomings: it requires the imposition of lepton number conservation discussed above, demands Yukawa couplings many orders of magnitude smaller than those of the charged leptons, and finally it does not account for the significant

differences in the hierarchy of masses and mixing angles between the lepton and quark sector [49].

Alternative production mechanisms for the Majorana mass terms can be found going beyond dimension-four operators, with the SM admitting the dimension-five gauge invariant Weinberg effective operator

$$\mathcal{L}_{M,BSM} = \frac{c^{[5]}}{\Lambda} L^T \cdot \tilde{H}^* C^\dagger \tilde{H}^\dagger \cdot L + \text{h.c.}, \quad (1.2.1)$$

which leads to Majorana masses after electroweak symmetry breaking:

$$\mathcal{L}_{M,BSM} \xrightarrow{\langle \tilde{H} \rangle \neq 0} \frac{c^{[5]} v_H^2}{2\Lambda} \nu_L^T C^\dagger \nu_L + \text{h.c.} \quad (1.2.2)$$

The Weinberg operator is a low-energy effective term corresponding to a full theory which is completed above the energy scale Λ and includes new particles with $M \approx \Lambda$. There are three possible options for the origin of this operator, at the tree level, that conserve Lorentz and gauge symmetries:

- Type-I seesaw for a singlet fermion [53];
- Type-II seesaw using heavy triplet scalars [54];
- Type-III seesaw for triplet fermions [55],

as illustrated in the diagrams in fig. 1.10 [49].




Figure 1.10: Diagrams contributing to neutrino masses for the three tree-level seesaw scenarios. $\langle H \rangle$ indicates the v.e.v. of the neutral component of the Higgs field. Δ is a scalar triplet and Σ is the neutral component of a fermion triplet of mass M_Σ [49].

As no indication of the scale of new physics is provided by neutrino experiments and theoretical considerations, both high energy, $\Lambda \gg \text{TeV}$, and low energy $\Lambda < \text{TeV}$ models can be considered. Low scale Seesaw models, which feature HNLs in the MeV to TeV range, in particular, are directly testable at experiments and, in certain mass ranges, offer a framework to explain the baryon asymmetry and are connected to dark sector models [49].

Type-I seesaw mechanism

The mechanism by which light neutrino masses are produced in seesaw models is clearly illustrated by the Type-I seesaw mechanism. In the case of three neutrino

families, the most general Dirac and Majorana mass term can be written in matrix form as [56]:

$$\mathcal{L}^{D+M} = \frac{1}{2} \bar{S}_L M S_L^C + \text{h.c.}, \quad (1.2.3)$$

where $S_L \equiv (\nu_L, (\nu_R)^C)^T$ with

$$\nu_L \equiv (\nu_{eL}, \nu_{\mu L}, \nu_{\tau L})^T \quad \text{and} \quad (\nu_R)^C \equiv ((\nu_{eR})^C, (\nu_{\mu R})^C, (\nu_{\tau R})^C)^T, \quad (1.2.4)$$

while the Seesaw mass matrix M has the form

$$M = \begin{pmatrix} m_L & m_D \\ m_D^T & M_R \end{pmatrix} \quad (1.2.5)$$

m_D , m_L and M_R being 3×3 matrices and $M_R = M_R^T$. In the Type-I seesaw, the left-handed Majorana mass term is absent ($m_L = 0$), and the Dirac mass arises from the Standard Model Higgs mechanism. In contrast, the right-handed Majorana mass of the heavy singlet fermion is not constrained by the electroweak scale, with $M_R \gg m_D$. A matrix m can be thus introduced by the relation $m \equiv U^T M U$, where U is a unitary matrix of the form

$$U = \begin{pmatrix} 1 & (m_D^T)^\dagger (M_R^{-1})^\dagger \\ -M_R^{-1} m_D^T & 1 \end{pmatrix}, \quad (1.2.6)$$

for which $U U^\dagger \simeq 1$ in the linear approximation over $m_D M_R^{-1}$. It follows that m takes a block-diagonal form up to $\mathcal{O}(m_D M_R^{-1})$ [2]:

$$m \simeq \begin{pmatrix} -m_D M_R^{-1} m_D^T & 0 \\ 0 & M_R \end{pmatrix} \quad (1.2.7)$$

Thus the Majorana masses for the light neutrinos arise:

$$m_\nu = -m_D M_R^{-1} m_D^T + \mathcal{O}([m_D M_R^{-1}]^2), \quad (1.2.8)$$

while the masses of the heavy neutrinos N_i are instead given by the diagonal entries of M_R . Low scale seesaw models are realised by taking $m_D \sim \text{keV}$, with the corresponding nearly-sterile states having masses in the GeV range. The heavy eigenstates are mainly composed of sterile neutrinos and mixing with active neutrinos is suppressed. A constraint to such mixing is imposed by the contribution to light neutrino masses, so that one can expect

$$|U_{\alpha N}|^2 \lesssim \frac{m_\nu}{m_N} \lesssim 10^{-10} \frac{\text{GeV}}{m_N}, \quad (1.2.9)$$

where $m_\nu \lesssim 0.1 \text{ eV}$ is taken [50].

1.2.2 Minimal HNL models

The Heavy Neutral Leptons that result from low scale seesaw models may be probed at laboratory energies: as their mixing with active neutrinos is suppressed to small values, as shown in eqn. (1.2.9), they will undergo *weaker-than-weak* interactions featuring long lifetimes and suppressed production.

Phenomenology and current constraints

Experimental searches have not found, so far, any significant evidence for their existence, but have provided strong limits on the mixing elements $U_{\alpha N}$. It is customary, when discussing experimental constraints, to assume a simplified flavour structure where mixing with one flavour dominates over the other two elements [57]. The following paragraphs present the main experimental channels and constraints for each of the three possible cases of flavour dominance.

Electron flavour dominance ($|U_{eN}|$). In the lower mass region ($M_N \lesssim 100$ MeV) the best limits are obtained by searches for an invisible resonance in $\pi^+ \rightarrow e^+ N$ decays such as those performed by TRIUMF [58] and PIENU [59]. Borexino has provided competitive limits below 16 MeV with a search for $N \rightarrow \nu e^+ e^-$ decays from HNLs produced in ${}^8\text{B}$ in the Sun [60]. Above the pion mass, the dominant constraints come from peak searches in leptonic kaon decays $K \rightarrow e^+ N$, particularly, those performed by NA62 [61]. A complementary technique in a similar mass region with competitive constraints is that of HNL decay-in-flight signatures at neutrino experiments. A search for HNL charged-current (CC) decay channels was performed by the T2K experiment at its near detector, ND280, using HNLs produced by kaon decays within the neutrino beam [62]. Above the kaon mass, the strongest constraints are currently set by CHARM [63] and BEBC [63]. For HNLs above ~ 22 GeV, finally, colliders are the best environment for detection: one of the most competitive limits was provided by the DELPHI collaboration with a search for HNL produced from Z decays with prompt or displaced decays in the detector [64]. This channel proceeds through neutral currents and as such the bound is largely insensitive to the flavour structure of HNL mixing. The ATLAS and CMS collaborations at LHC have furthermore set some of the strongest limits by using both W and Z decays with searches for HNL decays to dilepton final states [65, 66]. Beyond the energy range of colliders, HNL mixing can be probed indirectly by testing the unitarity of the standard PMNS matrix, through flavour and electroweak precision data [57].

Muon flavour dominance ($|U_{\mu N}|$). The pattern for this case is similar to the one above. Since the muon mass is similar to that of the pion, only HNLs lighter than ~ 30 MeV can be produced in pion decays. The best limits, as before, are provided by peak searches at PSI [67] and PIENU [68]. Above the pion mass, the decay-in-flight searches at PS191 and T2K [62, 69] have yielded the strongest constraints. Phenomenological studies on older data from MicroBooNE, T2K and KEK [70–72] have provided bounds on the region between $m_\pi - m_\mu \sim 34$ MeV and m_π . Above $M_N \sim 180$ MeV, the E949 and NA62 kaon factories provide the strongest constraints through peak searches in leptonic kaon decays [73, 74]. Low-energy neutrino experiments and kaon factories become insensitive above the kaon mass, and high-energy beam dumps are the most competitive through searches for decays of HNLs produced in D meson decays and ν_μ upscattering in the detector. The leading constraints are provided by BEBC, NuTeV and CHARM [75–77]. Above approximately ~ 2 GeV, the strongest constraints come from collider experiments. ATLAS and CMS searches [66, 78] focused on displaced decays of HNLs, while DELPHI [64] set competitive

limits from both prompt and displaced HNL decays produced in Z boson events. At even higher HNL masses, constraints can be derived from tests of the unitarity of the leptonic mixing matrix [57].

Tau flavour dominance ($|U_{\tau N}|$). Mixing with ν_τ has been the least probed so far. Light HNLs ($M_N < 16 \text{ MeV}$) can be produced by solar neutrinos inside the Earth and the search for their subsequent decays in the large-volume detector of Borexino has provided the best limits [79]. A similar strategy was pursued by SK using HNLs produced by atmospheric neutrinos scattering in the Earth, obtaining a limit of $|U_{\tau N}|^2 < 0.13$ (90% C.L.) for HNL masses in the $10 \text{ eV} \lesssim M_N \lesssim 10 \text{ MeV}$ [80]. Only high-energy beam dumps can produce more massive HNLs through D and τ decays: the best limits are those obtained in phenomenological re-analyses of CHARM [81] and BEBC [63] data together with a dedicated search by ArgoNeuT [82]. A peak search in $\tau \rightarrow N 3\pi$ decays was furthermore performed by the BaBar collider [83]. As discussed above, the most competitive bounds for HNLs above the GeV scale come from colliders, in particular from the search of displaced decays of HNLs produced in Z decays. Similarly to the other flavours, the limits for the highest masses stem from bounds on unitarity [57].

The updated bounds discussed in the paragraphs above are displayed in fig. 1.11 as a function of the HNL mass for each lepton flavour dominance case. Besides experimental constraints, the cosmological effects of HNLs, specifically their contribution to the energy density of the universe and the CMB temperature polarization and power spectra, allow to set stringent bounds, which, as fig. 1.11 displays, are stronger than laboratory constraints in some parts of the parameter space.

1.2.3 HNLs in Dark Sector models

In addition to being considered in isolation, that is in minimal extensions of the Standard Model as in the previous section, HNLs can be included in extended hidden sectors with multiple new particles and interactions [49].

A set of models that has gathered much interest in recent years is that of *Dark Sectors*, extensions of the standard model that advocate for new particles and interactions below the electroweak scale. Such particles would not partake in SM-interactions (or have extremely weak couplings) while displaying strong *dark interactions* [85].

Dark sectors can include different types of extension, with a new gauge sector, new scalars and/or fermions, or a combination of them, each option having a connection to the SM, known as *vector*, *scalar* and *neutrino portals* respectively. The popularity of such models stems from their ability to introduce new properties for neutrinos, to embed dark matter beyond the WIMP paradigm and to explain several low energy anomalies. Having relatively low energy scales, the dark sectors may furthermore be within reach of present and future non-collider experiments.

HNLs have been considered in the context of dark sectors with fermionic extensions, as they can connect to the SM via mixing between sterile neutrinos, *dark fermions* charged under the new interactions and neutrinos. While in minimal models HNL production and decays are controlled by SM interactions and the mixing with active neutrinos typically result in relatively long lifetimes in the MeV-GeV mass

range, additional interactions introduce new production and decay channels. In the dark sector, such interactions can be strong, leading to enhanced scattering cross section and fast decays [49].

Dark neutrino models

A class of models originally proposed to explain the excess of ν_e -like events observed in the MiniBooNE short-baseline neutrino experiment [86], via decays of dark sector particles into e^+e^- pairs with energies of $\mathcal{O}(100)\text{MeV}$, has gained traction in the community and has recently been targeted by experimental searches [87].

The model assumes a minimal $U(1)'$ gauge group with one or more HNLs whose associated flavour states, known as *dark neutrinos*, are the only particles carrying the group charge. After electroweak symmetry breaking, the dark neutrinos ν_D are a superposition of mass eigenstates. In the case of a single HNL [88]:


$$\nu_\alpha = \sum_{i=1}^4 U_{\alpha i} \nu_i, \quad (\alpha = e, \mu, \tau, D). \quad (1.2.10)$$

Where U is a 4×4 unitary matrix. It is expected that $|U_{\alpha 4}|$ is small for $\alpha = e, \mu, \tau$, while $|U_{D4}|$ is of $\mathcal{O}(1)$.

In this context, the dark sector can be connected to the SM by the kinetic mixing between the new $U(1)'$ gauge boson Z' (referred to as *dark photon*), hypercharge, and by neutrino mass mixing. The interaction lagrangian coupling Z' to both the neutral leptons and the standard EM charge can be parametrised as:

$$\mathcal{L}_{\text{int}} \supset \sum_{i,j}^{3+n} g_D Z'_\mu V_{ij} \bar{\nu}_i \gamma^\mu \nu_j - e \epsilon Z'_\mu J_{\text{EM}}^\mu, \quad (1.2.11)$$

where n is the number of dark neutrinos, J_{EM}^μ is the SM EM current with kinetic mixing ϵ . Assuming dark neutrino states to have a +1 charge under $U(1)'$, $V_{ij} = \sum_{k=0}^n U_{ki}^* U_{kj}$ is the interaction vertex for the dark current.



In such scenarios, HNLs are produced from an active flavour state by *upscattering* on a nuclear target A and decay into di-lepton pairs, e.g.:

$$\nu_\mu A \rightarrow (N \rightarrow \nu Z'^{(*)} \rightarrow \nu e^+ e^-) A, \quad (1.2.12)$$

where N is the HNL and the dark photon Z' may be produced on or off shell [89].

The upscattering cross section (here for a single dark neutrino) is proportional to $\alpha_D \alpha_{\text{QED}} \epsilon^2 |U_{\alpha 4}|^2$, where $|U_{\mu 4}|$ dominates since current accelerator neutrino beams are mostly composed of ν_μ . Production can occur off the whole nucleus, as in coherent scattering, or off individual nucleons. For $m_{Z'} \lesssim 100$ MeV, the scattering will be mainly coherent, but for heavier masses $\gtrsim 1$ GeV [90] incoherent scattering dominates [88]. A comparison between the standard quasi-elastic (CCQE) cross section and those of upscattering on Carbon for the benchmark parameters in [91] is shown in fig. 1.12.




Figure 1.12: Upscattering cross section compared to the quasi-elastic. The ν_e CCQE cross section on Carbon ($6p^+$) is shown as a function of E_ν (solid black line). The coherent (solid blue) and incoherent (dashed blue) upscattering cross sections are superimposed, for the benchmark point in [91]: $m_N = 420$ MeV, $m_{Z'} = 30$ MeV. The ν_μ flux of BNB at MiniBooNE and NuMI at MINERνA, the latter for both low (LE) and medium energy (ME) runs are shown in the background [88].

If the dark photon is produced on-shell, the decay chain is prompt, with forward going final states. Conversely, if Z' is off-shell, the HNL will decay via a three-body process and will be longer lived. In this latter case, HNL production in the rock or dirt upstream of the detector gives the largest contribution [89].

Dual dark neutrino models include an additional ν_5 HNL with $M_{\nu_5} > M_{\nu_4}$, and the decay occurs via $\nu_5 \rightarrow \nu_4 (Z' \rightarrow e^+ e^-)$, proportional to the $\mathcal{O}(1)$ coupling between ν_5 and ν_4 instead of the small coupling between active and dark neutrinos. It is thus possible to explore regions of the parameters space with $m_{Z'} \gtrsim 500$ MeV through prompt decay topologies [87].

Dark neutrino searches

The best opportunities for dark neutrino searches are offered by accelerator neutrino facilities and in particular by short-baseline experiments and by the near detectors of long-baseline ones. The host of liquid argon (LAr) experiments in the Short-baseline Neutrino Program (SBN) at Fermilab [92], including MicroBooNE [93], SBND [94] and ICARUS [95], are well-suited for such searches, as LAr detectors feature high calorimetric and spatial resolutions, potentially distinguishing the EM final states and the final state hadron multiplicity of interactions. The typical signatures of neutrino upscattering include a primary neutrino interaction with a displaced HNL decay vertex and EM showers. The showers may be reconstructed as coming from separate electrons and do not point back to the original vertex due to missing energy. A proton track at the initial vertex is also present for incoherent scattering. The distinct topologies in LAr experiments are shown in fig. 1.13, taking MicroBooNE as a reference.




Figure 1.13: The Four main topologies of e^+e^- dark neutrino decays at MicroBooNE: coherent and incoherent scattering having well-separated or overlapping e^+e^- pairs [49].

As of 2025, a first search for dark neutrino decays to e^+e^- pairs and missing energy has been conducted at MicroBooNE [87], focusing on coherent upscattering of neutrinos in single and dual dark neutrino models. The analysis used an exposure of 6.87×10^{20} protons-on-target (POT) from the BNB beam. The world leading limits exclude at the 95% C.L. the majority of the parameter space that was viable as a solution to the MiniBooNE anomaly [91]. The exclusion contours are shown in fig. 1.14.

The next generation of long-baseline accelerator neutrino experiments, such as T2HK [41] and DUNE [96], will offer new opportunities for dark neutrino searches at their near detector complexes, thanks to advances in detector technology and increased beam power.

The Near Detector (ND) complex of DUNE will take advantage of the high-intensity LBNF beam, which is expected to deliver 1.1×10^{21} protons on target (POT) per year at an initial beam power of 1.2 MW, with a planned upgrade to 2.4 MW after six years. Two components of the ND will use magnetized, gas-based detectors surrounded by electromagnetic calorimeters: ND-GAr, a large-volume high-

pressure gaseous argon TPC to be installed during Phase II of operations, and SAND, which features a straw-tube tracker. Gas-based detectors help mitigate the challenging backgrounds from hadronic activity in Standard Model neutrino scattering, since the signal rate for displaced HNL decays scales with detector volume, while the background scales with detector mass. This makes ND-GAr and SAND particularly well-suited for searching for off-shell decays of dark neutrinos.

The study of the sensitivity of SAND to dark neutrino signatures represents the main topic of this work.

Beyond accelerator neutrino facilities, fixed target experiments are another testbed for dark neutrino models: HNLs can be produced in kaon decays at NA62 [97], leading to cascade decays with displaced vertices. Searches for inelastic dark matter produced in electron bremsstrahlung are also sensitive to models of this kind [98].



(a) Single dark neutrino scenario, $\epsilon = 8 \times 10^{-4}$.(b) Dual dark neutrino scenario, $\Delta = 1$, $\epsilon = 8 \times 10^{-4}$.

Figure 1.14: Resulting MicroBooNE 95% C.L. exclusion contours for (a) a single dark neutrino model and (b) a representative dual dark neutrino scenario with $\Delta \equiv (m_5 - m_4)/m_4 = 1$. The coloured solid lines show the MicroBooNE limits for six $m_{Z'}$ values between 30 MeV and 1.25 GeV, with the region above each line excluded. The solid coloured regions indicate to the MiniBooNE 95% C.L. preferred regions for a viable explanation of the MiniBooNE anomaly [89] at each $m_{Z'}$ value. Stars denote the corresponding best-fit points. The gray contours correspond to prior independent limits [87].

Chapter 2

The Deep Underground Neutrino Experiment

The Deep Underground Neutrino Experiment (DUNE) is a future long-baseline neutrino oscillation experiment, currently under construction in the United States. DUNE will consist of a Near Detector complex at Fermilab, Illinois, at a distance of 574 m from the neutrino source and of a Far Detector site located at the Sanford Underground Research Facility (SURF) in South Dakota, 1.5 km underground at a baseline of 1300 km. DUNE will conduct a broad scientific programme addressing open questions in standard neutrino physics, the search for proton decay, the study of neutrinos from core-collapse supernovas and the Sun, and other beyond-the-Standard-Model (BSM) searches. The design and features of the DUNE components are discussed in sect. 2.1, while sect. 2.2 outlines the scientific programme of the experiment.

2.1 The design of DUNE

In order to reach the primary goals of its scientific program, the determination of the neutrino mass ordering and of the lepton-sector CP-violating phase, DUNE will leverage:

- The LBNF wide-band (anti)neutrino beam, the most intense in the world.
- A multi-kiloton Far Detector (FD) based on Liquid Argon Time Projection Chambers (LArTPCs).
- A set of Near Detector (ND) components capable of constraining the systematic uncertainties to an unprecedented level.

The early physics goals of DUNE are expected to be achieved by the end of Phase-I of operations, using a 1.2 MW proton beam, two far detector modules providing a total of 20 kt of liquid argon (LAr) fiducial mass, and a temporary near detector setup. Realizing the full scientific potential of DUNE will require the Phase-II upgrades: an increased beam power of 2.4 MW, four far detector modules with at least 40 kt of fiducial LAr mass, and a near detector complex in its final configuration.

2.1.1 The LBNF neutrino beam

The neutrino beam for DUNE will be provided by the Long Baseline Neutrino Facility (LBNF) at Fermilab. LBNF will be based on the Proton Improvement Plan II (PIP II) upgrade [99] of the Main Injector accelerator at Fermilab, which will allow to deliver a proton beam of between 1 and 1.2 MW in the 60 to 120 GeV energy range from the start of DUNE operations. Such beam power corresponds to 1.1×10^{21} protons-on-target (POT) per year. The PIP II upgrade will also provide a platform for extending the beam power to 2.4 MW for Phase-II.

As is conventional for neutrino beams, a proton beam will be extracted from the Fermilab Main Injector (MI) and made to impinge on a fixed graphite target. The secondary mesons, primarily $\pi^\pm$ and $K^\pm$, that are thus produced are focused by the magnetic horn system and left to decay in a 194 m long decay pipe, producing a beam of mostly $\mu^\pm \nu_\mu/\bar{\nu}_\mu$. An absorber pile at the end of the decay pipe removes the residual hadrons and most $\mu^\pm$, leaving a small contamination of $\nu_e/\bar{\nu}_e$. The wide-band neutrino beam will be aimed at the Far Detector complex at SURF, 1.5 km underground at a distance of 1300 km from the production point [100]. Mostly positive (or negative) particles are focused by selecting the horn magnetic field polarity. The horn operational modes are referred to as Forward Horn Current (FHC) and Reverse Horn Current (RHC) and correspond to a neutrino and antineutrino beam, respectively. The main proton beam parameters for Phase-I are listed in Table 2.1.

The focusing system will provide a wide-band neutrino beam in a 0.5-10 GeV range with a peak at ~ 3 GeV so as to cover the first and second oscillation maxima (approximately 2.4 and 0.8 GeV at a 1300 km baseline). In fig. 2.1 the unoscillated fluxes for both horn modes during Phase-I are shown.

Parameter	Value
Energy	120 GeV
Protons per cycle	7.5×10^{13}
Spill duration	1.0×10^{-5} s
POT per year	1.1×10^{21}
Cycle time	1.2 s
Beam power	1.2 MW

Table 2.1: Principal parameters of the LBNF primary proton beam for DUNE Phase-I [100].

2.1.2 The DUNE Far Detector

The Far Detector (FD) of DUNE, located at a depth of 1.5 km at SURF, will consist, by the end of the staging plan for Phase-II, of four detector modules. Only two modules, FD1 and FD2, are planned to be operational in Phase-I. Both will be LArTPCs holding 17.5 kt of LAr.

The single-phase LArTPC technology selected for FD1 and FD2 offers high performance in tracking, calorimetry, and particle identification. Since no signal amplification occurs in the liquid argon, low-noise electronics are essential to achieve a




Figure 2.1: Expected unoscillated fluxes at the Far Detector for the FCH (left) and RHC (right) modes. Uncertainties on the ratio between FN and ND unoscillated fluxes are at $\sim 1\text{-}2\%$ [101].

good signal-to-noise ratio [101]. FD1 and FD2 will employ different TPC designs: FD1 will use a horizontal-drift chamber, where ionization charges drift horizontally toward anode wire planes, while FD2 will adopt a vertical-drift TPC, where charges drift vertically toward anode PCB planes located at the top and bottom of the detector.

The two TPC designs are being tested at the CERN Neutrino Platform on a hadron test-beam using the ProtoDUNE-HD (PDHD) and ProtoDUNE-VD (PDVD) prototypes, which implement the horizontal-drift and vertical-drift configurations of FD1 and FD2, respectively. Each prototype is approximately 1/20 the size of a full-scale FD module and uses the same components as the final detectors. ProtoDUNE-HD has completed two data-taking runs, the first between 2018 and 2022 [101], and the second between May and November 2024. Data-taking for ProtoDUNE-VD began in May 2025 [102]. Results obtained with the ProtoDUNE detectors have already provided valuable information on the design, calibration and simulation of the FD modules.

At the time of writing, the third module is planned to use a vertical-drift TPC design, while the design of the fourth module has yet to be finalized.

The FD1 Horizontal Drift LArTPC

The FD1 single-phase horizontal-drift LArTPC will contain a total of 17.5 kt of LAr, for a ~ 10 kt fiducial mass. The TPC will be housed inside a $65.8\text{ m} \times 17.8\text{ m} \times 18.9\text{ m}$ cryostat and will feature four 3.5 m wide drift volumes between five alternating $58\text{ m} \times 12\text{ m}$ anode and cathode walls. The internal structure of the module is shown in fig. 2.2.

Each of the cathode walls consists of an array of 150 Cathode Plane Assemblies, $1.2 \times 4\text{ m}$ panels held at -180 kV . A uniform 500 V/cm field is produced across the drift volume as the anode walls are kept close to ground. A field cage surrounds the




Figure 2.2: Structure of the inner fiducial volume of the FD1 module, showing the alternating anode (A) and cathode (C) planes, and the field cage (FC) that surrounds the drift region between anode and cathode planes [103].

remaining open sides of the TPC, ensuring a $< 1\%$ field uniformity throughout the active volume.

Anode walls consist of 50 $6\text{ m} \times 2.3\text{ m}$ Anode Plane Assemblies (APAs) arranged in a 2×25 grid. APAs, which measure $6 \times 2.3\text{ m}$, are two-sided, with three active wire layers in a XUV configuration and ten photon detectors. The wire spacing on the layers is 5 mm. An additional shielding wire layer is wrapped around the assembly.

Drifting ionization electrons will produce bipolar induction signals on the outer U and V induction layers and will be absorbed by the X collection layer, yielding a monopolar signal. The induction and collection signals are shaped, amplified and digitized by the front-end cold electronics, which are connected to both the top and bottom ends of an APA wall unit shown in fig. 2.3. The low temperature of LAr allows the reduction of thermal noise.

In addition to ionization, charged particles passing through LAr will produce Vacuum Ultraviolet (127 nm) scintillation light at approximately 2.4×10^3 photons per MeV, providing the starting time for the ionization [103]. In FD1, the scintillation photons are collected by the X-Arapuca devices, set in bars running the 2.3 m width of an APA between the wire layers. Each X-Arapuca cell consists of dichroic filters transparent to LAr scintillation light alternated to wavelength-shifter (WLS) plates that convert the UV photons in the visible spectrum at 430 nm. The dichroic filters reflect visible photons escaping the WSL back to SiPMs at the edge of the plates, as illustrated in fig. 2.4.




Figure 2.3: Left: single unit of an APA wall formed by two linked APAs. Photodetection system bars are set across the width of the APAs. Right: top and bottom ends of an APA stack showing the readout electronics [103].




Figure 2.4: Operating principle of an X-Arapuca device. A dichroic filter cutoff at 400 nm is assumed [103].

The FD2 Vertical Drift LArTPC

A vertical drift configuration has been selected for the FD2 module LArTPC. In this setup, the cathode plane is suspended at mid-height within the module and placed at -300 kV, while the anode planes are positioned at the top and bottom as fig. 2.5 illustrates. This arrangement creates two vertically stacked drift volumes, each 6.5 m long, with a uniform electric field of 500 V/cm [104].

In a departure from the design of FD1, the TPC anodes are constructed of perforated PCB boards with etched electrodes, forming a three-view charge readout while




Figure 2.5: Scheme of the FD2 vertical drift LArTPC, with a close-up of the cathode plane with mounted photo-sensors [105].

avoiding deformations when hung horizontally. Each section of the anode plane is composed of two double-sided PCB boards, the innermost one hosting a guard electrode and the first induction plane, and the second one featuring the second induction plane and the collection plane, as shown in fig. 2.6. The induction and collection planes are readout by strip electrodes in a $+30^\circ$, -30° , 90° configuration relative to the beam direction, providing three separate views. The anodes are mounted within Charge Readout Plane structures on the top and bottom of the module.

The LArTPC photodetection system for FD2 will once again be based on X-Arapuca devices, which will be installed on the cryostat walls, designed to be semi-transparent, as well as on the cathode plane. As fig. 2.7 illustrates, the X-Arapuca modules will be $60 \times 60 \text{ cm}^2$ tiles equipped with 160 SiPMs. Since the modules mounted on the cathode are referenced to the cathode voltage, they will employ non-conductive optoelectronic systems for both power and signal transmission to eliminate any direct path to ground. Power will be delivered via a power-over-fibre system, and an analog optical transmitter will be used to relay output signals to warm digitizers [104].

2.1.3 The Near Detector Complex

The Near Detector Complex (ND) of DUNE will be located at approximately 574 m from the neutrino source for the LBNF beam. For Phase-I, the ND will consist of three detectors, shown in fig. 2.8: ND-LAr, TMS (Temporary Muon Spectrometer) and SAND (System for On-Axis Neutrino Detection). During Phase-II TMS will be replaced with the ND-GAr gaseous argon TPC. The ND-LAr and TMS/ND-GAr detectors will be capable of moving to take data in positions off the beam axis, a programme referred to as Precision Reaction-Independent Spectrum Measurement (DUNE-PRISM) [106].

The role of the ND will be crucial for the oscillation programme of DUNE [106]:

- ND will characterise the beam close to the source with high-statistics. This will provide the initial state of the beam that is compared to the FD measurements to extract the oscillation parameters. Functionally similar LArTPCs in both ND and FD will be instrumental in reducing the systematics.




Figure 2.6: Structure of the anode PCB (left) and simulation of the electric field between two holes in the planes (right) [105].




Figure 2.7: Structure of an X-Arapuca module for FD2 [105].

- SAND will be a powerful spectral beam monitor, capable of detecting changes in the beam flux in a timely fashion.
- The high statistics collected and the optimal particle-ID at ND will allow to tune the neutrino interaction model used to move between beam model and experimental data.
- ND-LAr and TMS/ND-GAr will take data at different off-axis positions, thus providing datasets with different beam spectra, thereby making it possible to

deconvolve flux and cross-section models and constrain separately each one.

The following section will briefly discuss the technology of the ND components.




Figure 2.8: View of the ND-complex detectors in the on-axis configuration during Phase-I.

The ND-LAr LArTPC

Based on the same fundamental principles as the FD modules, the ND-LAr LArTPC will adopt novel solutions to cope with the high intensity of the beam at the ND, which would lead to pile-up in a traditional LArTPC. The detector will be composed of multiple smaller optically-isolated TPC modules with individual pixellated readouts. The timing information will be sufficient to associate tracks and events across the matrix. The subdivision of the volume and optical insulation, will lead to smaller drift distances and times, hence mitigating the issue of pile-up. Currently, a 5×7 matrix of modules, pictured in fig. 2.9, is planned. The pixellated readout will enable a full 3D reconstruction of tracks with enhanced robustness for a high-multiplicity environment. ND-LAr will have a total active volume of $5 \times 7 \times 3 \text{ m}^3$ for a fiducial mass of 67 t of LAr, leading to an expected $10^8 \nu_\mu\text{-CC}$ events/year on-axis [106].

The ND-GAr Gaseous Argon TPC

ND-GAr will feature a high-pressure gaseous argon TPC with an outer electromagnetic calorimeter (ECAL) within a 0.5 T magnetic field, as pictured in fig. 2.10. An outer muon system will complete the setup. As ND-LAr begins to lose acceptance for muons above $\sim 0.7 \text{ GeV}$, ND-GAr will provide coverage for the momentum and charge reconstruction of events that are not contained in the upstream detector. The low momentum threshold for charged particle tracking, excellent resolution and nearly uniform angular coverage guaranteed by the use of a gaseous argon TPC will make it possible to obtain a large sample of $\nu\text{-Ar}$ interactions independent from the ND-LAr




Figure 2.9: External view of the ND-LAr detector, one of the seven columns of five modules [106].

systematics, with approximately 1.6×10^6 ν_μ -CC events per year of on-axis running with 1.0 t argon fiducial volume. Together with ND-LAr, ND-Gar will characterise the off-axis beam spectra within the PRISM programme [106].

As mentioned above, ND-Gar is planned for operation during DUNE Phase-II. In its place, the Temporary Muon Spectrometer (TMS) will run throughout Phase-I. TMS, pictured in fig. 2.11, will consist of magnetised steel layers interleaved with planes of scintillator strips, providing a muon momentum resolution of up to 5%.

The System For On-Axis Neutrino Detection (SAND)

SAND will be the only detector at the ND to be fixed in an on-axis position. Leveraging, in this way, the large neutrino flux from the LBNF beam, SAND will monitor the flux of neutrinos going towards FD on a weekly basis, on top of a rich programme of neutrino physics. SAND will reuse the magnet and electromagnetic calorimeter (ECAL) of the KLOE experiment [107, 108]. The magnetised inner volume of the ECAL will be instrumented with a low-density tracking system with distributed target for cross-section studies and a small liquid argon active target volume. The next Chapter provides a detailed description of SAND and of its physics programme.

2.2 The scientific programme of DUNE

The primary objective of DUNE will be to perform a comprehensive programme of neutrino oscillation measurements, aimed at determining the δ_{CP} CP-violating




Figure 2.10: External scheme of ND-GAr, showing the High-Pressure Gas TPC and its external pressure vessel, the ECAL, and the magnet with its return yoke [106].




Figure 2.11: External schematics of the TMS detector, showing the interleaved scintillator planes and magnetised steel layers.

phase of the lepton sector, resolving the neutrino mass ordering and the octant of θ_{23} , and testing unitarity in the three-neutrino mixing framework. Together with the oscillation program, one of the main objectives of DUNE is the detection and measurement of ν_e from a core-collapse supernova, should one event occur during the

lifetime of the experiment.

A wide range of BSM-physics measurements will be possible at the ND and FD: from proton decay searches to dark-matter candidates, non-standard interactions, CPT-violation, sterile neutrinos and HNLs and tau-neutrino appearance. Neutrino cross section measurements and studies of nuclear effects will furthermore be possible at the ND [101].

2.2.1 Neutrino oscillation program

As a long baseline accelerator experiment, DUNE will observe the $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ channel. The probability for such oscillations, given by eqn. (1.1.100) in a constant matter density approximation, is such that both δ_{CP} and the mass ordering introduce an asymmetry between the $\nu_\mu \rightarrow \nu_e$ and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ channels. The asymmetry increases linearly with L in the few-GeV range, making DUNE particularly sensitive owing to its 1300 km baseline. The asymmetry, the sign of whom depends on the mass ordering, amounts to about $\pm 40\%$ in the peak flux region, thus being larger than the maximal CP-violation asymmetry related to δ_{CP} . Hence, DUNE will be capable of removing the degeneracy between the two asymmetry contributions and determine both mass ordering and δ_{CP} unambiguously and to a high confidence level. The $\bar{\nu}_e$ -appearance probabilities at the FD baseline are shown in fig. 2.12 as a function of neutrino energy for multiple values of δ_{CP} . The difference in probability between different values of δ_{CP} is larger at the higher oscillation nodes below 1.5 GeV. Being a broadband experiment, DUNE will be capable of mapping out the spectrum of oscillations down to 500 MeV, proving therefore optimal for studying δ_{CP} [101].




Figure 2.12: Probability of $\bar{\nu}_e$ -appearance at the DUNE FD 1300 km baseline as a function of neutrino energy for $\delta_{CP} = -\pi/2$ (blue), 0 (red), $\pi/2$ (green). The oscillation behaviour for $\theta_{13} = 0$ is shown as a black line. Normal ordering is assumed in both plots [101].

By observing the oscillation structure over a wide energy range, DUNE will be able to perform precision measurements of the individual PMNS parameters with improved sensitivities and without the need for external constraints. The current

projections for seven, ten and fifteen years are presented in the following paragraphs for the standard staging plan of DUNE [101]:

- Beam run start: two FD modules (FD1 and FD2) for a 20 kt fiducial mass, 1.2 MW beam.
- After one year: FD3 added for a total fiducial mass of 30 kt.
- After three years: FD4 added, for 40 kt of fiducial mass.
- After six years: beam upgraded to 2.4 MW.

Sensitivity to CP-violation The significance for the observation of CP-violation as a function of the true value of $\sin^2 2\theta_{13}$ is shown in fig. 2.13a for exposures equivalent to seven and ten years of data, assuming equal run times in FHC and RDH modes. The significance of the CPV measurement drops necessarily to zero at the CP-conserving values of $\delta_{CP} \neq 0, \pm\pi$, producing the double peak structure. The variation in the fit projections, shown by the transparent bands indicating 68% of fits, is significantly affected by the true value of $\sin^2 2\theta_{23}$. The significance for CP-violation discovery for 75% and 50% of the δ_{CP} values and for $\delta_{CP} = -\pi/2$ is plotted in fig. 2.13b as a function of the exposure, accounting for the staging plan: CP-violation can be observed at a 5σ significance after 7 years for $\delta_{CP} = -\pi/2$ and after about 10 years for 50% of δ_{CP} values. 3σ is reached at 13 years of running for 75% of δ_{CP} values [101].

Sensitivity to mass ordering Fig. 2.14a shows the significance for neutrino mass ordering determination as a function of the true δ_{CP} value. The near degeneracy between matter and CP-violating effects occurring near $\delta_{CP} = \pi/2$ for true mass ordering results in the characteristic shape. The staging plan and exposure assumptions discussed above are assumed. The significance for the measurement of mass ordering for 100% of δ_{CP} values and for $\delta_{CP} = -\pi/2$ are plotted in fig. 2.14b as a function of the exposure in years. The plot extends to seven years of exposure, corresponding to 500kt-MW-years, as DUNE will be able to establish the mass ordering at a $5\text{-}\sigma$ level for 100% of δ_{CP} values after between two and three years [101].

Sensitivity to the oscillation parameters DUNE will improve the precision of measurements for several key neutrino oscillation parameters, including δ_{CP} , $\sin^2 2\theta_{13}$, Δm_{31}^2 , as well as $\sin^2 \theta_{23}$ and its octant. The resolution, in degrees, of the DUNE measurement of δ_{CP} as a function of the true δ_{CP} value is shown in fig. 2.15. The measurement resolution is significantly better close to CP-conserving values δ_{CP} , compared to maximally CP-violating values. Depending on the true δ_{CP} values, resolutions between 5° and 15° are possible for fifteen years of exposure.

Figures fig. 2.16a to fig. 2.16d plot the resolution of DUNE measurements of δ_{CP} , $\sin^2 2\theta_{13}$, $\sin^2 2\theta_{23}$ and Δm_{32}^2 as a function of the exposure in kt-MW-years. At high exposures, the measurement of $\sin^2 2\theta_{13}$ will reach a precision comparable to that of reactor experiments, enabling a direct comparison between the two. Such a comparison would serve as an important test of the unitarity of the PMNS.




Figure 2.13: (a) Significance of the DUNE determination of CP-violation as a function of the true δ_{CP} for seven (blue) and ten (orange) years of exposure. Normal ordering is assumed. The width of the bands corresponds to 68% of the fits. The median sensitivities are shown by the solid lines. (b) Significance of the DUNE determination of CP-violation for $\delta_{CP} = -\pi/2$, 50% and 75% of the true values as a function of exposure time. Normal ordering is assumed. The width of the bands corresponds to the external constraint on $\sin^2 2\theta_{13}$ [101].

2.2.2 Low-energy physics and Supernova neutrinos

DUNE will be sensitive to low-energy neutrinos between 5 MeV and a few tens of MeV. This energy range is particularly interesting for the detection of neutrino-bursts produced by a galactic core-collapse supernova. Neutrinos from other astrophysical sources, such as the Sun, and the Diffuse Supernova Neutrino Background may also be detectable.

Detection of Supernova neutrino bursts

The LArTPCs at the DUNE FD will have a particular sensitivity to the ν_e component of the neutrino burst through the dominant charged-current absorption of ν_e on ^{40}Ar

$$\nu_e + {}^{40}\text{Ar} \rightarrow e^- + {}^{40}\text{K}^*, \quad (2.2.1)$$

which will be detected via the e^- plus de-excitation products of the K^* final state. Other potential detection channels are currently under investigation.

The expected number of signal events from a supernova at 10 kpc in the 40 kt FD full fiducial volume according to the updated GKVM model [109] is ~ 3700 , over a timespan of ~ 10 s. With such a statistics, DUNE is expected to characterise the time-dependent spectrum of ν_e , which is the dominant component of the flux. Preliminary simulations indicate that radiological backgrounds in Ar will have a minor effect on the triggering of a supernova signal. DUNE will be capable of reconstructing, to an extent, the direction of the incoming neutrino flux, providing a valuable early




Figure 2.14: (a) Significance for the DUNE determination of mass ordering as a function of the true δ_{CP} for seven (blue) and ten (orange) years of exposure. Normal ordering is assumed. The width of the bands corresponds to 68% of the fits. The median sensitivities are shown by the solid lines. (b) Significance of the DUNE determination of mass ordering for $\delta_{CP} = -\pi/2$ and for 100% of the δ_{CP} values as a function of exposure time. Normal ordering is assumed. The width of the bands corresponds to the external constraint on $\sin^2 2\theta_{13}$ [101].




Figure 2.15: Resolution, in degrees, for the DUNE measurement of δ_{CP} as a function of its true value, for seven (blue), ten (orange) and (fifteen) years of exposure. Normal ordering is assumed. The width of the bands is mainly affected by external constraints on $\sin^2 2\theta_{13}$ [101].

warning for astronomers. An overall angular resolution of $\sim 4.5^\circ$ can be reached using a sample of ν_e elastic scattering [101].




Figure 2.16: (a) Resolution of the DUNE measurement of δ_{CP} as a function of the exposure, assuming normal ordering. Curves for $\delta_{CP} = -\pi/2$ and $\delta_{CP} = 0$ are shown. The width of the bands is due to constraints on $\sin^2 2\theta_{13}$. (b) Resolution of DUNE’s measurement of $\sin^2 2\theta_{13}$ as a function of the exposure. External constraints are not significant so that the unconstrained curve is shown. (c) Resolution of DUNE’s measurement of $\sin^2 2\theta_{23}$ as a function of the exposure. The width of the band is due to external constraints on $\sin^2 2\theta_{13}$. (d) Resolution of DUNE’s measurement of Δm_{32}^2 as a function of the exposure. The unconstrained curve is shown as external constraints are irrelevant [101].

On top of their significance for astrophysics, core-collapse supernovae represent a particularly interesting environment for fundamental neutrino physics. Flavour oscillations occurring during the event will be non-trivial, being highly dependent on the dynamics of the explosion and involving both the neutrino and antineutrino fluxes as well as the solar and atmospheric mass splittings. Crucial information

on the core-collapse process, can thus be extracted from the neutrino signal [101]. Several features of the neutrino burst signal will in particular be sensitive to the effect of neutrino mass ordering during the different eras of neutrino emission. The most robust prospect for mass ordering determination will be the signal from the neutronization burst, which will be highly suppressed in the case of Normal Ordering due to the complete oscillation of ν_e to ν_x [110].

Sensitivity to Solar Neutrinos

Detecting solar and other astrophysical low-energy neutrinos will be a challenge for the DUNE LArTPCs due to the relatively high detection threshold (> 5 MeV) for CC-interactions in LAr. Compared to other technologies, though, LArTPCs offer a large cross section and unique channel-tagging signatures for de-excitation photons. The overall solar neutrino event rate for the final FD fiducial mass of 40 kt is ~ 100 per day [101]. The most challenging backgrounds for solar neutrino detection at DUNE are radiological, particularly neutron and alpha-capture reactions on argon—the latter being especially poorly understood. Sophisticated event selection techniques, together with additional shielding and improved radiopurity may allow to obtain a high-statistics neutrino sample, enabling competitive measurements of the ^8B and ^7Be ¹ solar neutrino fluxes and of the Δm_{21}^2 solar mass difference [111].

Study of the Diffuse Supernova Neutrino Background

While individual galactic supernovae are relatively rare (~ 1 -4 per century), the neutrinos produced by supernova explosions all over the universe form a diffuse background (DSNB) with energies in the few-to-30 MeV range. The DSNB is as of yet unobserved, and large LArTPC detector such as those at the DUNE FD, being sensitive to its ν_e component, complement water Cherenkov and scintillator detectors, which are sensitive to the antineutrino component. The search for the DSNB will be limited by background: the lower bound for the detection energy range is the 18.8 MeV endpoint of ^7Be neutrinos, while the upper bound is determined by the atmospheric ν_e flux at ~ 40 MeV. The expected number of DSNB in the 10 kt fiducial volume of an FD module is of 1-2 per MeV over 20 years within the 19-31 MeV window, making even rare radiological cosmogenic backgrounds challenging [101, 112].

2.2.3 Beyond the Standard Model searches

DUNE will be able to leverage the high-intensity LBNF neutrino beam together with its ND-Complex detectors and the massive FD LArTPCs at a 1300 km baseline to carry out a variety of searches for Physics Beyond the Standard Model with unprecedented sensitivity.

Searches for sterile-active neutrino mixing. DUNE will be sensitive to a broad range of potential mass splittings in sterile neutrino oscillations by searching for the disappearance of CC and NC interactions over both the long baseline between the

¹Solar neutrinos produced in the rare $^3\text{He} + p \rightarrow ^4\text{He} + e^+ + \nu_e$ reaction [1].

ND and FD and the short baseline of the ND. The longer baseline, intense beam and high-resolution large-mass FD detectors will make it possible to improve significantly the sensitivities to existing probes.

Non-unitarity of the PMNS matrix. As mentioned in sect. 1.2, most models proposed to explain the neutrino mass pattern feature heavy neutrino states in addition to the light SM ones. This would imply a deviation of the 3×3 PMNS matrix from unitarity. A deviation of $\mathcal{O}(10^{-2})$ would decrease the reach of DUNE to the standard parameters.

Searches for Non-standard Interactions (NSI). NSI can affect neutrino propagation through the earth, altering significantly the data collected by DUNE if the parameters associated to new physics are large enough. DUNE will be uniquely sensitive to these probes thanks to its very long baseline and wide-band beam.

Searches for Lorentz or CPT symmetries violation. Potential violations of the Lorentz and CPT symmetries would significantly upend the Standard Model of particle physics. DUNE can improve on the current constraints by several orders of magnitude, providing valuable tests of these fundamental assumptions.

Searches for Light-mass Dark Matter (LDM). Light Dark Matter models, featuring candidates below the electroweak GeV scale, have gained traction in recent years, often including a new *vector portal* through which the dark matter candidates interact with ordinary matter. High-flux neutrino beam experiments such as DUNE can cover a region of the DM and mediator parameter space that cannot be accessed by direct detection or collider experiments [113]. Dark matter particles could be detected at the ND via neutral-current-like interactions with electrons or nucleons in the detector material. The kinematics of the scattered electron and timing information can be used for background suppression [101]. The ability to perform off-axis measurements with DUNE-PRISM will allow for sampling with lower neutrino-to-DM flux ratios, since DM production is due to the decay of un-focused neutral mesons [114].

Boosted Dark Matter (BDM) searches. DUNE will be able to test Boosted Dark Matter, models, where a lighter relativistic DM component is produced in the annihilation of a heavier one in an astrophysical source, leveraging its large FD fiducial mass [115]. The light DM component can have energies above the DUNE detection thresholds in a wide region of parameter space [101].

Heavy Neutral Lepton searches As detailed in sect. 1.2, the high beam power of the LBNF beam and the use of magnetised gas-based detectors at the ND will provide new opportunities for HNL searches, as this work sets out to evaluate.

Chapter 3

The SAND detector

SAND (System for On-Axis Neutrino Detection) will monitor the neutrino flux at the DUNE Near Detector complex from a fixed on-axis position and perform a broad range of additional measurements and searches. An overview of the detector design and performance is provided in sect. 3.1, while sect. 3.2 outlines the proposed SAND physics program.




Figure 3.1: Exploded view of the SAND detector, illustrating (from left to right) the inner volume detectors (GRAIN and the tracker), the ECAL barrel and endcap modules, and the magnet.

3.1 The SAND detector design

The design of SAND is based on the refurbished magnet and electromagnetic calorimeter from the KLOE experiment [107], and incorporating novel detectors in the inner magnetised volume. An exploded view of the detector is shown in fig. 3.1.

3.1.1 The SAND magnet

The SAND superconducting magnet has been successfully operated over the course of the KLOE experimental runs. The magnet is designed in conjunction with its iron return yoke to provide a 0.6 T magnetic field over a 4.3 m long, 4.8 m diameter inner volume. The solenoidal longitudinal field component measured during the KLOE installation phase is shown in fig. 3.2. The magnet coil is enclosed within a cryostat measuring 4.40 m in length and 5.76 m in diameter, which is positioned inside the return yoke. The coil operates at a nominal current of 2902 A, storing an energy of 14.32 MJ. It is composed of a single conductor layer of Nb-Ti Rutherford cable co-extruded with high-purity aluminium. The cooling system of the coil uses 5.2 K He gas at a 3 bar absolute pressure, which is then liquefied into a reservoir in thermal contact with the coil [116].




Figure 3.2: Longitudinal component, along the magnetic axis, of the KLOE solenoid field (in Gauss) [116].

3.1.2 The Electromagnetic Calorimeter (ECAL)

The KLOE Electromagnetic Calorimeter (ECAL) [108] will be refurbished for its use in SAND. The ECAL is a lead/scintillating-fibre sampling calorimeter, offering a good light transmission over distances of several meters, a sub-ns timing accuracy and a very high hermeticity. The barrel calorimeter is located inside the magnetised volume, close to the cryostat. It has a cylindrical shape subdivided into 24 modules,

each 4.3 m long and 23 cm thick, having a trapezoidal cross-section with 52 and 59 cm bases as pictured in fig. 3.3a. The ECAL endcap modules will be mounted on the endcap return yoke section. Each endcap comprises 32 modules, with lengths ranging from 0.7 to 3.9 meters, and a rectangular cross-section. Both ends of each module are bent in a C-like shape to allow their insertion into the ECAL barrel.




Figure 3.3: (a) Scheme of the ECAL barrel module segmentation into PMT cells. (b) Placement of the light guides at one end of an ECAL barrel module [116].

All ECAL modules are constructed as a stack of 200 grooved 0.5 mm thick lead foils alternated with 200 layers of 1 mm cladded scintillating fibres, glued together with compatible epoxy. The modules are read out by phototubes at both ends, with each end-face being subdivided into five planes: the first four are 4.4 cm deep and the fifth is 5.2 cm deep, as pictured in fig. 3.3a. Each plane is segmented into cells 4.4 cm wide in the transverse direction, for a total of 4880 phototubes. Light guides, pictured in fig. 3.3b, match the square readout cells to the circular photocathodes. The energy and time resolution of the ECAL was determined during the commissioning and running phases of KLOE, amounting to [108]:

$$\begin{aligned} \text{Energy resolution} \quad \frac{\sigma}{E} &= \frac{5.7\%}{\sqrt{E \text{ (GeV)}}}, \\ \text{Time resolution} \quad \sigma &= \frac{54 \text{ ps}}{\sqrt{E \text{ (GeV)}}} \oplus 140 \text{ ps}. \end{aligned} \tag{3.1.1}$$

3.1.3 The GRAIN LAr active target

The upstream section of the inner magnetised volume of SAND will host the GRAIN (GRanular Argon for Interaction of Neutrinos) ~ 1 t liquid argon detector. GRAIN will provide a sample of on-axis neutrino-argon interactions, thus supplementing the ND-LAr detector. GRAIN will follow a novel approach for track reconstruction in LAr, based exclusively on the detection of the UV scintillation emission. Liquid Argon scintillation produces Vacuum Ultraviolet (VUV) light with a large yield (~ 40 k photons/MeV) and two emission components, with decay times of 7 ns (fast) and 1.6 μ s (slow). The fast component can be leveraged to limit collection time to few nanoseconds, thus enabling high spatial resolutions. Equipped with an optical system coupled to fast, segmented SiPMs, GRAIN will provide calorimetric information and spatial reconstruction of the interactions on LAr [117].

The GRAIN cryostat will be meniscus-shaped, with a 147 cm high 150 cm wide and up to 47 cm deep inner vessel. A view of the inner and outer GRAIN vessels is shown in fig. 3.4. Two options for the VUV-sensitive optical system are currently being developed, one based on lenses [118] and one based on Coded Aperture Masks. While lens-based cameras are easier to reconstruct, they are difficult to build for a cryogenic environment and have a limited Field of View due to the high LAr refractive index. Mask-based cameras will conversely have a simpler and more compact design, but a more complex GPU-based high-performance algorithm will be needed [119]. Both camera designs will make use of SiPM-matrix sensors with a TPB wavelength shifter coating to match the 127 nm VUV emission to their optimal photon detection efficiency in the visible spectrum [117]. Schematics of the cameras for both optical systems are shown in fig. 3.5.




Figure 3.4: Renderings of the GRAIN outer and inner vessels in the YX and YZ views (left and right respectively).

3.1.4 The SAND Straw Tube Tracker

The remaining section of the inner magnetized volume of SAND will host the Straw Tube Tracker (STT), a gas-detector tracking system with diffuse C and CH₂ target layers. STT will be capable of high resolution measurements of neutrino interaction on multiple nuclear targets.

The two-fold need for a large enough target mass (~ 5 t fiducial mass) and for high space and momentum resolutions will be satisfied by way of a diffuse target design. The neutrino target mass is arranged in thin solid layers spread out uniformly throughout the tracking volume, accounting for more than 97% of the total detector mass. Layers of low-mass straw-tube detectors are alternated to the targets, providing high space ($< 200 \mu\text{m}$) and momentum (*e.g.* $\sigma(1/p)/(1/p) \simeq 4\%$ for 1 GeV μ) resolutions thanks to their low density and high granularity. Using dedicated target




Figure 3.5: (a) 3D model of a GRAIN mask-based camera, showing the Coded Aperture Mask and the inner SiPM matrix. (b) Cutaway view of a lens-based camera: the optical system consists of two plano-convex lenses with a nitrogen-containing interface between them [119].

layers will make it possible to change their configuration, in order to tune the average density of the detector ($5 \times 10^{-3} < \rho < 0.18 \text{ g/cm}^3$) and perform measurements with several target materials.

STT will be subdivided into independently operated and configured modules, for a flexible overall design. The default module configuration (in fig. 3.6a) aimed at ν -hydrogen interaction measurements, consists of: (a) a 5 mm thick polypropylene (CH_2) target slab, (b) a ~ 14 mm thick CH_2 radiator for e/π separation via transition radiation, (c) four straw tube layers arranged in a XXYY configuration [116].




Figure 3.6: (a) Cutaway view of a CH_2 STT module showing: (i) the tunable CH_2 target layer, (ii) radiator foils for e/π discrimination, (iii) four straw layers in a XXYY configuration. (b) 3D view of the SAND inner volume, with the upstream GRAIN LAr active target followed by STT. The barrel ECAL surrounds the inner volume.

The straw tubes have 5 mm diameter, 12 μm mylar walls with Al coating and a 20 μm diameter gold-plated tungsten sense wire. The default gas mixture for the straws is Xe/CO₂ (70%/30%) gas mixture at 1.9 atm, ensuring sensitivity to the transition radiation as well as a relief to the mechanical tension applied to the supporting frames thanks to their overpressure.

The main planned target material is graphite (pure C), for the direct measurement of the C-background in the selection of hydrogen interactions (see sect. 3.2.5). The baseline configuration, outlined in fig. 3.6b, foresees 8 C-target modules, 70 CH₂ modules and 6 target-less tracking modules (at the upstream and downstream ends), for a total CH₂ fiducial mass of ~ 5 t. The ratio of C to CH₂ modules is tuned to optimize the purity of the H samples within technical constraints [120], corresponding to ~ 600 kg of graphite per ~ 700 kg of H for the aforementioned CH₂ 5 t fiducial mass. In this configuration, the average density of STT is $\rho \sim 0.18$ g/cm³ with a radiation length $X_0 \sim 2.6$ m. Each of the CH₂ and graphite modules corresponds to 1.5% and 2% of X_0 respectively [116]. While the default STT configuration is aimed at the measurement of $(\bar{\nu} - H)$ interactions, alternative nuclear targets are foreseen, including Ca, Fe and Pb.

The SAND STT is currently in the advanced development phase, with ongoing work on prototypes and on the design of the final modules. A medium scale (120×80 cm²) STT module prototype was assembled at CERN over 2023 and has undergone successful testing with a ⁵⁵Fe source, and a further iteration has recently been exposed at the SPS muon test beam [121].

Drift chamber backup design for STT

A backup design for STT based on large-area planar drift chambers is being developed, aiming at a reduced mechanical complexity and at a lower number of readout channels, while maintaining a comparable physics performance to the base solution. The drift-chamber tracker will feature the same modular layout of straw-tube tracker: each detector module, pictured in fig. 3.7a, will consist of a layer of target material (CH₂ or C) and three stereo-wire planes in a staggered $-5^\circ, 0^\circ + 5^\circ$ configuration with respect to the SAND magnetic field axis. The use of stereo wires will reduce the left-right ambiguity while maintaining an optimal spatial resolution in the bending direction and minimising the number of channels read from the top and bottom sides of the module assemblies, thus simplifying the mechanical design. Multiple detector modules will be placed inside a gas-tight frame, constituting a *super-module*. A preliminary layout of the *super-modules* in the SAND inner volume is shown in fig. 3.7b.

The wire-planes consist of alternating anode sense wires (gold-plated tungsten, $d = 20$ μm) and cathode field wires (copper-beryllium alloy, $d = 100$ μm). Grounded mylar foils (20 μm thick) are used to separate the wire planes. Gaps (2 mm wide) between the foils ensure the flow of gas between the planes. An Ar/CO₂ (85%/15%) mixture at an overpressure of ~ 5 mbar is currently foreseen.

A medium-scale prototype of a drift-chamber module, with an active area of 120×80 cm² has been built as a testbed for technical solutions that will be applied to the full-scale module and has recently completed a run with a muon test beam at

SPS [121]. A detailed discussion of the development and test of the drift-chamber prototype can be found in [122].




Figure 3.7: (a) Layout of a unit module of the drift chamber design. (b) Preliminary layout of the *super-modules* in the inner volume of SAND. Each *super-module* contains 9 CH_2 and a C target module.

3.2 The SAND physics program

Located in a fixed on-axis position, SAND will enable continuous and high-rate beam monitoring, carry out neutrino interaction measurements essential for the Near Detector (ND), and support a comprehensive program of precision studies and BSM physics searches.

3.2.1 Beam monitoring

A key aspect in the operation of DUNE ND will be the continuous monitoring of the $\nu_\mu/\bar{\nu}_\mu$ -beam to detect potential variations which may affect the oscillation analysis. To this purpose, SAND will provide: (a) continuous measurements from a fixed on-axis position, (b) a large target mass (22.8 tons fiducial from the upstream ECAL), (c) an accurate reconstruction of the neutrino energy, (6% and 7% from ECAL and STT respectively), (d) measurements of the beam profile over a wide cross-sectional area. Such features will enable a sensitivity of $\Delta\chi^2 > 9$ in one week of data taking with 3.78×10^{19} POT (Protons On Target) for variations in most of the beam parameters [116].

3.2.2 Flux measurements

The precise knowledge of neutrino fluxes is crucial for both the primary Far Detector (FD) oscillation analysis and all other ND measurements. Specifically, it is necessary

for unfolding the separate terms in the number of events for any exclusive process X either at ND or FD:

$$N_X(E_{\text{rec}}) = \int_{E_\nu} dE_\nu \Phi(E_\nu) P_{\text{osc}}(E_\nu) \sigma_X(E_\nu) R_{\text{phys}}(E_\nu, E_{\text{vis}}) R_{\text{det}}(E_{\text{vis}}, E_{\text{rec}}), \quad (3.2.1)$$

where $\Phi(E_\nu)$ is the incoming (anti)neutrino flux, P_{osc} is the survival probability ($P_{\text{osc}} \sim 1$ at ND), σ_X is the X -process cross-section, R_{phys} is the response function introduced by nuclear smearing and R_{det} is the detector response function, which accounts for resolution and acceptance. E_ν and E_{rec} are respectively the true and reconstructed neutrino energies.

The folded terms in eqn. (3.2.1) cannot be decoupled using a single detector or nuclear target. In addition to this, in the $0.5 < E_\nu < 10 \text{ GeV}$ energy range relevant for oscillation analyses at DUNE, (anti)neutrino interactions receive relevant contributions from multiple processes: $\sim 25\%$ quasi-elastic (QE), 42% resonances (RES) and 33% deep inelastic scattering (DIS). The current uncertainties in neutrino cross-sections are not consistent with the precision required by the DUNE oscillation analysis [123], as shown in fig. 3.8 for the discovery potential of the CP-violation search at the FD [116]. SAND will be capable of an accurate determination of the absolute $\bar{\nu}_\mu^{(-)}$ and relative $\bar{\nu}_\mu^{(-)}, \bar{\nu}_e^{(-)}$ fluxes through several complementary processes. In particular, the energy dependence of the relative $\nu_\mu/\bar{\nu}_\mu$ flux can be determined at a sub-percent level with a 5 years exposure through the exclusive $\nu_\mu p \rightarrow \mu^- p \pi^+$, $\bar{\nu}_\mu p \rightarrow \mu^+ p \pi^-$ channels on STT (see sect. 3.2.5) [116].




Figure 3.8: Impact of systematics uncertainties on the sensitivity for the CP-violation search at the FD. Small changes in the relative uncertainties (from 1% to 3%) can increase the exposure needed to reach a 5σ significance substantially [124].

3.2.3 Constraints on nuclear smearing

Owing to its integration of the GRAIN LAr detector with the STT modules, SAND will be able to constrain the R_{phys} factor in eqn. (3.2.1), as both systems have a similar acceptance R_{det} with respect to the lighter targets. Once the fluxes are precisely measured, see sect. 3.2.2, the remaining $\sigma_X R_{\text{phys}}$ can be constrained by comparison between the LAr target and the lighter targets integrated in STT, which have the same R_{det} .

To this end, the critical role is played by the hydrogen interaction sample in STT, for which $R_{\text{phys}} \equiv 1$ (see sect. 3.2.5). An accurate measurement of σ_X on hydrogen is possible given the large expected statistics, so that the unfolding of E_ν only depends on the detector response R_{det} , which is determined essentially by the momentum resolution $\delta p/p$. Constraining directly the product $\sigma_X R_{\text{phys}}$ on LAr is therefore possible by comparison between LAr and H interactions [116].

3.2.4 Background rejection

Interactions occurring in the magnet and ECAL materials that surround the STT fiducial volume can produce particles (mostly neutral) that reach the inner tracker thus representing a source of background for genuine STT interactions. This effect is exacerbated by the ~ 110 mass ratio between the magnet+ECAL and the STT fiducial volume. Through a multivariate analysis of timing and topological information from both ECAL and STT, SAND will efficiently reject the background from beam-related neutrino interactions in the material surrounding the detector, reaching a rejection factor of 3×10^{-5} against CC+NC external interactions, with 92.7% efficiency and 99.6% purity [116].

3.2.5 Solid Hydrogen Measurement

The measurement of (anti)neutrino scattering cross-sections on hydrogen (*i.e.* free proton) which will be carried out by STT is crucial for the physics potential of the DUNE ND, as it is one of the only two types of interactions to be free from nuclear effects (the other being ν -e elastic scattering) [116]. The combination of C and CH₂ solid targets in STT overcomes the technical issues related to solid or gaseous H-targets in underground experiments, hence the name *solid hydrogen* for this kind of analysis. The selection of H-events makes use of the excellent angular, momentum, vertex and timing resolutions of STT, to precisely locate the interaction vertex in the specific target layers.

For the exclusive CC topology $\nu_\mu p \rightarrow \mu^- p \pi^+$, coming mainly from resonance production (resonances make up $\sim 68\%$ of ν -interactions in STT), the selection of H-interactions is performed through likelihood functions incorporating the complete event kinematics and the information from internal hadron system structure, respectively $\mathcal{L}^{\text{H}}$ and $\mathcal{L}_{\text{IN}}^{\text{H}}$. Such kinematic selections can achieve a high rejection power [120]. Table 3.1 shows the expected CC statistics in CH₂ and H in STT with the $\nu/\bar{\nu}$ beam modes and for the two planned beam power options (1.2 and 2.4 MW), assuming a 5 years data taking with each beam mode [125].

1.2 MW beam statistics	CH ₂	H
ν_μ -C (FHC, 5 yrs.)	35×10^6	3.6×10^6
ν_μ -C (RHC, 5 yrs.)	13×10^6	2.9×10^6
2.4 MW beam statistics	CH ₂	H
ν_μ -C (FHC, 2 yrs.)	66×10^6	6.5×10^6
ν_μ -C (RHC, 2 yrs.)	24×10^6	4.3×10^6

Table 3.1: Expected CC statistics in STT (5 tons of CH₂) for interactions in CH₂ and H, with the different options considered for the DUNE (anti)neutrino beam [125].

3.2.6 Precision measurements and searches

Making use of its high available statistics and the accurate determination of the (anti)neutrino fluxes through $(\bar{\nu})$ -H interactions, SAND will be capable of exploiting the unique properties of the (anti)neutrino probe for the study of fundamental physics. A broad program of measurements and searches is planned, spanning a range complementary to collider, fixed-target and nuclear physics experiments.

Electroweak precision measurements

The precise determination of the weak mixing angle $\sin^2 \theta_W$ in (anti)neutrino scattering at the DUNE energy range is of interest as it: (a) provides a direct measurement of neutrino couplings to the Z boson and (b) it probes a different momentum transfer scale compared to LEP as it does not operate on the Z mass peak. SAND will be capable of determining the Weinberg angle from the ratio $\mathcal{R}^\nu \equiv \sigma_{NC}^\nu / \sigma_{CC}^\nu$ of NC to CC-DIS neutrino interactions by leveraging its efficient identification of ν_e -CC interactions and an event-by-event kinematic analysis separating NC and CC interactions. The measurement of $\mathcal{R}^\nu$ will be dominated by theoretical systematic uncertainties on the structure functions of the target nuclei [126–128]. Overall, a total relative uncertainty on the value of $\sin^2 \theta_W$ of 0.35% can be feasibly extracted from ν N-DIS.

An independent measurement of $\sin^2 \theta_W$ can be obtained via the neutral current (NC) $\nu_\mu e$ elastic scattering channel, which is free from hadronic uncertainties but suffers from a very small cross-section. The value of the weak mixing angle can be extracted from the ratio $\mathcal{R}_{\nu e}(Q^2) \equiv \sigma(\bar{\nu}_\mu e \rightarrow \bar{\nu}_\mu e) / \sigma(\nu_\mu e \rightarrow \nu_\mu e)$ [129], in which the systematics due to selection and electron identification cancel out. An overall precision of 1% or better on $\sin^2 \theta_W$ is feasible from a combined analysis of SAND and ND-LAr. The DIS and νe -elastic channels have substantially different Q^2 -scales, allowing a test of the running of $\sin^2 \theta_W$ in the same experiment [116].

Tests of Isospin physics

Isospin physics involving both neutrino and antineutrino interactions is a particularly compelling area of study in SAND. The experiment’s ability to collect large samples of $(\bar{\nu})$ -H interactions [120] enables a high-precision test of the Adler sum rule $S_A = 0.5 \int_0^1 dx/x (F_2^{\bar{\nu}p} - F_2^{\nu p}) = I_p$ [130], which so-far has only been tested with limited statistics [131]. In SAND it will be possible to measure S_A as a function of Q^2 from the

structure functions $F_2^{\bar{\nu}p}$ and $F_2^{\nu p}$ determined from the corresponding differential cross-sections on hydrogen. Such a measurement would be sensitive to potential violations of the isospin (charge) symmetry, to the production of heavy (charm) quarks and asymmetries in the strange sea ($s - \bar{s}$). The measurement on H made in STT can be compared to the value obtained from the C-target, for which $S_A = 0$ [116].

Nucleon structure physics and QCD

Owing to the precise calibration of energy scale uncertainties and its highly precise determination of the $\nu/\bar{\nu}$ fluxes, SAND will be capable of measuring structure functions and cross-sections up to accuracies comparable to electron scattering experiments. The structure functions F_2 , xF_3 , F_L and F_T can be determined through both neutrino and antineutrino DIS [132]. Moreover, both the Parton Distribution Functions (PDFs) and non-perturbative corrections can be studied through global QCD analyses of (anti)neutrino data in a broad range of Q^2 and Bjorken x given the energies and statistics available in DUNE. The potential use of multiple nuclear targets, in addition to H and C, at STT will make it possible to separate valence and sea quark distributions, d and u distributions and the strange quark s and $\bar{s}$ distributions. Such nucleon structure studies in SAND can provide complementary information to the ongoing programs at colliders and fixed-target DIS experiments [116].

Searches for New Physics

In addition to the precision tests discussed in the previous sections, each sensitive to various proposed New Physics models through potential deviations from Standard Model predictions, SAND will also enable a broad range of direct searches for BSM physics.

A prime example will be the test of the MiniBooNE low-energy anomaly at a similar L/E to the original experiment, for which the optimal resolution and electron-identification capabilities of STT can be leveraged to search for anomalies in the $\bar{\nu}_\mu^-$ -CC (disappearance) and $\bar{\nu}_e^-$ -CC (appearance) spectra. Oscillations to sterile neutrinos can be thus detected using both the CC-ratios $\mathcal{R}_{e\mu}(L/E) \equiv (\nu_e N \rightarrow e^- X)/(\nu_\mu N \rightarrow \mu^- X)$ and $\bar{\mathcal{R}}_{e\mu}(L/E) \equiv (\bar{\nu}_e N \rightarrow e^+ X)/(\bar{\nu}_\mu N \rightarrow \mu^+ X)$ and the NC/CC ratios $\mathcal{R}_{\nu p}(Q^2) \equiv \sigma(\nu_\mu p \rightarrow \nu_\mu p)/\sigma(\nu_\mu p \rightarrow \mu^- n)$ and $\bar{\mathcal{R}}_{\bar{\nu} p}(Q^2) \equiv \sigma(\bar{\nu}_\mu p \rightarrow \bar{\nu}_\mu p)/\sigma(\bar{\nu}_\mu p \rightarrow \mu^+ n)$. In addition to anomalies in the beam spectra, STT will provide a high sensitivity to potential ν_τ appearance as a result of sterile neutrino oscillations or Non Standard Interactions.

Finally, SAND will further enhance the sensitivity of DUNE-ND to Dark Sector physics searches, including HNLs (see sect. 1.2), axion-like particles, dark photons and sub-GeV dark matter among other models. The present work specifically aims at estimating the sensitivity of SAND to dark neutrino models, in which Heavy Neutral Lepton states are connected to the Standard Model via a new $U(1)'$ gauge boson Z' , the dark photon (see sect. 1.2.3).

Chapter 4

The SAND simulation framework

Detailed Monte Carlo simulations for the DUNE ND components have been developed in order to study reconstruction and analysis algorithms and evaluate the performance of the detectors. Such frameworks include multiple steps, from the generation of neutrino interactions to the simulation of detector effects and, lastly, reconstruction algorithms.

This chapter will detail the standard simulation and reconstruction framework developed for SAND in sect. 4.1, describing the main steps in the production and analysis of any event dataset. The steps taken to integrate available generators of dark neutrino events in the main SAND simulation chain will be presented in sect. 4.2, discussing the current limitations and potential developments.

4.1 The standard SAND simulation chain

The simulation of standard neutrino interactions in SAND involves several stages, beginning with event generation within the Near Detector geometry. This is followed by the propagation of particles and energy deposits, and concludes with the simulation of the detector readout. Reconstruction algorithms, based on the simulated detector response, are currently being developed for the ECAL, GRAIN, and STT sub-detectors. A schematic overview of these simulation steps is provided in fig. 4.1, illustrating the overall simulation workflow. As the present analysis will not make use of GRAIN, its simulation pipeline will not be discussed in detail (see [117, 119] as reference).

4.1.1 ND and SAND detector geometries

The geometries on the DUNE ND components are implemented using the Geometry Description Markup Language (GDML), a format based on the XML language that is widely used by HEP frameworks including GENIE [133], Geant4 [134] and ROOT [135].

The geometry GDML files are built using DUNE-ND-GGD [136] the official DUNE tool for the implementation of complex geometries at the ND. This software system allows to define constructive solid geometries of any complexity based on basic shapes through boolean operations. Each component solid can be assigned a material in order




Figure 4.1: Scheme of the simulation framework for SAND. The main flow of the simulation is shown by the solid red arrows, while the dashed indicate inputs to non-adjacent stages. The stages of the Sandreco software and GRAIN simulation pipeline are highlighted.

to generate the GDML file. DUNE-ND-GGD is highly modular, making it possible to produce geometries with different configurations (e.g. the two options for the SAND tracker). Both single detectors and the full ND-geometry, including the cavern space and surrounding rock can be modelled. A view of the ND detectors in the Phase-II configuration is pictured in fig. 4.2.




Figure 4.2: View of the geometry for the Phase-II ND detectors, showing, from left to right, ND-LAr, ND-GAr and SAND. The beam is at a 5.3° angle with respect to the z-axis.

A detailed geometry of SAND has been implemented, modelling the detector elements discussed in sect. 3.1.

The magnet system. The components of the superconducting magnet and coil have been modelled together with the iron yoke, the latter divided into separate sections according to the magnitude of the return solenoidal field in the barrel and endcap components.

ECAL geometry The ECAL barrel consists of 24 trapezoidal modules, each 4.3 m long and 23 cm thick, with major and minor base widths of 59.6 cm and 52.5 cm respectively. The modules are arranged along a cylindrical profile of 2 m radius. A detailed endcap model has also been implemented, comprising 32 C-shaped modules with lengths ranging from 75 cm to 386 cm and widths from 8.88 cm to 26.64 cm, as illustrated in fig. 4.3.

Both barrel and endcap modules share the same internal structure: 209 alternating layers of scintillator fibres and lead foils, with thicknesses of 700 μm and 400 μm respectively. The 2.5 cm aluminium back-plate has also been included in the module geometry.




Figure 4.3: View of the ECAL geometry showing (a) the full set of barrel and endcap modules and (b) the structure of a single endcap module. The 2.5 cm thick module back-plates are shown in blue.

GRAIN geometry. Both the inner and outer vessels of GRAIN have been modelled as elliptical cylinders, the outer vessel having 192.4 cm $\times$ 85.4 cm axes and a length of 193.2 cm and the inner vessel having semi-axes of 72.8 cm $\times$ 23.75 cm with a length of 65 cm. The multi-layered structure of the outer vessel, consisting of carbon-fiber and honeycomb-composite layers with a thickness of 62 mm has been modelled. The inner vessel consists of a 6 mm layer of stainless steel. Stainless steel plates, 20 mm thick, are also used to model the inner vessel endcaps. A cutout scheme of the GRAIN geometry is shown in fig. 4.4.

A separate model of the GRAIN inner optical system has been developed to provide a flexible framework for studying camera and lens configurations independently of the event simulations in the argon volume. The same neutrino event sample can thus be used to study several readout configurations.




Figure 4.4: View of the GRAIN model implemented for the full SAND geometry. The dimensions of the axes are given for the outer vessel, while the semi-axes dimensions are shown for the inner vessel. All dimensions are expressed in mm.

Tracker geometry. Geometries for both the straw-tube and drift-chamber based solutions discussed in sect. 3.1.4 have been implemented with the aim of performing comparative performance analyses on the principal SAND physics items. In the straw-tube STT geometry, a basic module consists of the target layer, the separate radiator foils (TRD) and the four XXYY tube layers. The STT modules fill the inner volume, each having a different length according to its position in SAND, as shown in fig. 4.5a. As mentioned in sect. 3.1, a graphite-target module is followed by 9 modules with C_3H_6 targets. Four XXYY straw-tube tracking layers, two upstream and two downstream, complete the setup, pictured in fig. 4.6a.

The unit volume of the drift-chamber tracker geometry, coincides with the drift-chamber module introduced in sect. 3.1.4: it includes a target layer and three wire planes, or *tracking stations*, as shown in fig. 4.6b. In the current geometry wires are simulated with an XXY orientation perpendicular to the beam axis. The material budget of the Cu-Be and tungsten wires is overall negligible, hence the plan to remove them from future iterations of the geometry and define the wire configuration at the digitization stage of the simulation (see sect. 4.1.4). In the present model, the wire planes have a 1 cm spacing, and are separated by 200 μm thick Mylar planes. The graphite-target module and the following 9 C_3H_6 ones form a *supermodule*, packed within a single graphite frame. In the current design, the tracker volume comprises eight supermodules, whose heights are adjusted to fit the inner volume, as pictured in fig. 4.5b. An additional targetless tracking station is modelled upstream of the supermodules. The drift-chamber geometry accounts for all clearances within the

inner-volume components, according to the current design.




Figure 4.5: View of the SAND tracker layout in (a) the standard straw tube configuration and (b) the backup drift chamber solution. In the drift chamber geometry 1 + 9 modules are grouped into supermodules, for which the clearances have been modelled.




Figure 4.6: View of the base tracker modules in (a) the Straw Tube geometry and (b) the drift chamber backup solution. The base straw tube module features TRD layers. The drift chamber wire planes are shown here in a staggered XXY configuration.

4.1.2 The GENIE neutrino event generator

The standard DUNE-ND simulation pipeline relies on the GENIE Monte Carlo generator [133] to produce neutrino events. GENIE is a ROOT-based [135] framework largely used by modern neutrino experiments. GENIE aims at supporting the simulation of all processes for all neutrino species and nuclear targets in the MeV to PeV

energy range. Its modular framework includes nuclear physics, hadronization and cross section models, as well as advanced physics and BSM channels.

GENIE integrates both complex GDML detector geometries and external flux models in the generation.

Generation for DUNE-ND is done using the `gevgen_fnal` event generation application, which integrates GENIE with the standard beam-line simulations and geometry descriptions of Fermilab experiments. Events can be generated in the geometry discussed in sect. 4.1.1 using the official Near Detector LBNF flux models [137], referred to a distance of 574 m from the horns, and including the $\bar{\nu}_\mu^{(-)}$ and $\bar{\nu}_e^{(-)}$ components for both FHC and RHC modes, as illustrated in fig. 4.7.

The GENIE generation chain. Beyond detector inputs such as geometry and flux models, GENIE also requires the neutrino interaction model to be specified. A model configuration defines the set of processes to be simulated (e.g. QE, RES, DIS), the specific model assigned to each interaction, the nuclear structure model, and the hadronization scheme. It also specifies the treatment of cascades and final-state interactions.

The actual parameter values for the models in the configuration list are subsequently specified in *tunes*, which are typically based on global fits to external data. Multiple tunes can therefore be defined for the same model configuration. The G18_02a_00_00 [138] and AR23_20i_00_000 tunes are widely employed in DUNE-ND standard productions, with the latter having been recently designated as the official tune.




Figure 4.7: Unoscillated fluxes at the Near Detector for the FHC (left) and RHC (right) modes from the available beam data release [139].

Once the interaction tune has been specified, together with the detector geometry and flux model, neutrino events can be generated. The most computationally demanding step in this process is the selection of the interaction channel, which requires

integrating the differential cross section of each possible channel over the available phase space at the given energy. To mitigate this, GENIE generally relies on pre-computed total cross sections stored as XML *spline* files. The spline-generation tools allow the user to specify the list of neutrino flavours, target nuclei, interaction tune, and energy range. Pre-computed splines for the most commonly used tunes are also available online [140].

The generator output is produced in the native *GHEP* format, a ROOT TTree file that stores *EventRecord* instances corresponding to individual interactions. Each *EventRecord* contains detailed particle-level information as well as event-level quantities, including total and differential cross sections, event weights, and a summary of the interaction. GHEP files may be converted into several alternative formats for interoperability with subsequent simulation stages. In DUNE, the `numi_roottracker` format is employed: this is a simplified ROOT TTree structure in which particle data are represented as arrays of fundamental types.

4.1.3 Particle propagation in edep-sim

The event generation output, in a `numi_roottracker` format, is taken as an input for the following simulation stage: the propagation of particles in the detector geometry and the simulation of energy deposition in the SAND active volumes. Such a step is carried out by the `edep-sim` software [141], a wrapper around the Geant4 [134] framework. `edep-sim` includes detailed models for energy deposition through both ionizing and non-ionizing processes as well as the NEST (Noble Element Simulation Technique) for LAr.

`Edep-sim` processes the GENIE event output by propagating the primary particles generated at the interaction vertex. It simulates their subsequent interactions and the resulting secondary particles, storing the interaction points information in *particle trajectory* objects. The energy deposits associated to each particle trajectory within the sensitive volumes of the detector, referred to as *hit-segments*, are also simulated. For each hit segment, `edep-sim` records the initial and final positions, the deposited energy, and the identifiers of both the depositing particle and its parent.

A SAND-specific library, `edep-reader` [142], has been developed as a wrapper to the `edep-sim` output in order to facilitate the interfacing with and querying of the Monte Carlo truth information during analyses. The library creates a tree structure with a parent-child relation between the trajectories of an `edep-sim` event. Each trajectory stores information on its own trajectory points and hit-segments.

4.1.4 The SAND reconstruction framework

The output of the propagation stage in `edep-sim` is the basis for the simulation of detector effects and the reconstruction in the separate SAND subdetectors. Currently, the simulation and reconstruction of ECAL and tracker is integrated in the `sandrec` framework [143], whereas, due to the high-performance requirements of its reconstruction (see sect. 3.1.3), a separate pipeline has been developed for GRAIN [117, 119].

The sandreco framework. The `sandreco` software is used to simulate the detector response and reconstruction in the tracker and ECAL. It features an initial *Digitization* module, which provides a simulation of digitised ADC and TDC values for the signals of the tracker and ECAL from the edep-sim hit-segments. The digitization algorithm for the tracker, which is adapted to both geometry options, still makes use of Monte Carlo information for its digitised values, while a response simulation based on the Garfield++ toolkit [144] is in an initial stage of development (see Appendix A). The digitisation for the ECAL, conversely, already includes PMT response and propagation effects. Reconstruction algorithms for the tracker and ECAL are implemented in the *Reconstruction* module. Multiple reconstruction algorithms are available for the tracker, from a Monte Carlo-based fast reconstruction, to a Kalman Filter implementation in active development. A clustering algorithm for ECAL has been developed and validated.

Simulation and reconstruction in GRAIN A detailed simulation of the scintillation light emission in GRAIN, starting from edep-sim hits and modelling the scintillation spectrum and emission profile of VUV light, together with scattering and absorption. Starting from the scintillation photons, a waveform simulation for the optical readout SiPMs has been developed. Starting from the optical system response, a 3D track reconstruction algorithm based on a combinatorial approach with an iterative maximisation of a likelihood for photon emission was developed. A complete discussion of the technique is found in [119].

4.2 Dark neutrino event generators

For the modelling of new physics in neutrino–nucleus interactions, the community currently relies on different approaches, including dedicated custom generators, additional modules within established neutrino generators such as GENIE, NuWro [145] and GiBUU [146], and, more recently, end-to-end Monte Carlo frameworks that go from a user-defined Lagrangian all the way to event generation in complex geometries.

The present analysis on the SAND sensitivity to Dark Neutrino models made use of two single-purpose generators: DarkNews [147], a fast generator with detailed cross-section models, and the DarkNeutrino module for the previously discussed GENIE generator.

4.2.1 The DarkNews event generator

DarkNews is a fast Python-based generator for the simulation of heavy-neutrino up-scattering and decay in neutrino detectors. It is built on the `vegas` numerical integrator [148], allowing it to sample from pre-computed analytical differential cross sections and decay rates.

Featured HNL models. The principal scope of DarkNews is the modelling of coherent neutrino-nucleus scattering to HNLs (Dark Neutrinos) followed by decays mediated by additional interactions, e.g. Dark Photons (see sect. 1.2.3). Cross

sections for incoherent scattering are provided, but no nuclear effect is currently modelled.

DarkNews features two types of HNL model implementation:

- as a generic HNL model, defined through a parametrization of the interaction vertices by specifying a set of simplified interaction terms and couplings,
- as a Three-portal model [85], which, as introduced in sect. 1.2.3, features dark neutrino states ν_D in addition to the standard sterile neutrino ν_N and both Z' dark photon and dark scalar h' mediators. In this case, the interaction vertices may be computed directly by specifying the relevant model parameters.

In its current state, DarkNews supports models with up to three dark neutrino states, with both Dirac and Majorana neutrino types being available. In addition to the contributions from the dark mediators, dark neutrino processes include terms from the Standard Model Z and W bosons as well as potential transition magnetic moment terms [147].

Production cross sections. The full analytical expressions for the upscattering cross sections are computed assuming a low-energy neutrino flavour state scattering on an hadronic target (either a nucleus or a free nucleon) and producing a heavy neutral lepton state. A diagonal mass matrix is considered for the three light neutrino states, while heavy neutrinos ($i \geq 4$) are assumed to undergo decoherence from the light states. The hadronic tensors for the interaction are determined for each scattering regime (coherent and incoherent) assuming spin-1/2 nuclei and nucleons.

In coherent scattering, the electric form factors are used for the γ , Z and Z' exchanges, while a separate formalism [149] is implemented for the h' scalar exchange. Elastic scattering on free protons and nucleons is modelled via the relevant Dirac, Pauli and Axial form factors. The nuclear form factors for the interaction are either given in a Fourier-Bessel series parametrization [150], if nuclear data is available, or by a Fermi-symmetrised function [151].

HNL decay simulation. DarkNews models both the prompt and off-shell Z' HNL decays introduced in sect. 1.2.3, together with the ν_5 and ν_6 decays in dual dark neutrino models. At the time of writing, only the e^+e^- and single photon HNL decay channels are implemented, with the di-muon decay being a future addition.

The DarkNews workflow

The event generator is initialised by passing as input the model parameters discussed above, either directly or via a custom model file, together with a set of experimental parameters:

- a text file containing the neutrino fluxes of each flavour for the chosen beam,
- a list of target nuclei with the corresponding fiducial mass,
- the required exposure in POTs.

Detector geometries are not, in general, available within DarkNews, with the scattering position being set at the geometric origin and at $t = 0$. So far, fiducial volume geometries have only been implemented for MiniBooNE [86], MicroBooNE [93] and SBND [94]. In all other cases, the events must be placed within the experiment’s own geometry after generation is complete.

Once the input parameters are defined, vegas is used to sample based on the event rate

$$\frac{dN}{dQ^2 dPS} = \frac{d\sigma}{dQ^2} \left(\frac{d\Gamma_{N \rightarrow \nu X}}{dPS} \frac{1}{\Gamma_{\text{tot}}} \right), \quad (4.2.1)$$

where $d\Gamma_{N \rightarrow \nu X}/dPS$ is the differential decay rate of HNL into neutral lepton and $X \equiv e^+e^-, \mu^+\mu^-, \gamma$ while Γ_{tot} is the total HNL decay rate. The samples thus generated are weighted according to the event rate and the generation is repeated for each combination of parameters in the process chain, e.g. for each nuclear target, scattering regime, ν -projectile flavour etc. The 4-momenta for the particles are subsequently built and their positions are assigned.

The event generation output features all event information, weights and 4-momenta built from the vegas samples for each process chain. The default output format is a `pandas` dataframe [152, 153] that stores the model and experiment information as metadata, but other widely used formats in HEP are available, including `numpy` arrays [154], and the standard `HEPevt` and `HepMC` [155] data formats. A chart of the DarkNews generation workflow is shown in fig. 4.8.

4.2.2 The DarkNeutrino module in GENIE

DarkNeutrino is a single purpose GENIE module [156] for the generation of dark neutrinos. The module has been available since the GENIE 3.2 production release and can be enabled upon installation.

At the time of writing, only the Z' -mediated coherent scattering channel is modelled, with the exclusion of free nucleons, as ^1H targets are, for which the form factors are not implemented. The decay channels are furthermore limited to the prompt HNL decay to e^+e^- through on-shell Z' , thus requiring $m_{HNL} > m_{Z'}$. Only models with single dark neutrino state are furthermore available. Other interaction and decay channels may be added to the framework according to the standard GENIE workflow.

Despite offering a more limited set of simulated channels compared to DarkNews, DarkNeutrino was selected for signal event generation in this analysis due to its ease of integration, as a GENIE module, into the official DUNE-ND simulation pipeline. The production scripts could be quickly adapted from the standard ones, and the resulting output, aside from minor corrections discussed below, remained fully compatible with the existing framework. Importantly, GENIE allows event generation directly within the DUNE-ND geometry (see sect. 4.1.2), thus avoiding the ad-hoc event positioning and re-weighting procedures that a DarkNews-based pipeline would require.

Overall, while currently limited in its range of available models, DarkNeutrino was deemed to be the more robust option for the analysis, allowing to potentially extend the pipeline to the other DUNE-ND components.




Figure 4.8: Schematic chart of the main steps and parameters in the DarkNews event generation pipeline (from [147]).

Model implementation. The differential cross section for the currently available channel of coherent scattering with a Z' mediator is implemented following [91], as, in the case of ν_μ -scattering:

$$\begin{aligned} \frac{d\sigma}{dT_T} &= 2\pi\alpha_{\text{EM}}\alpha_D F F^2 \epsilon^2 |U_{\mu 4}|^2 |U_{D4}|^2 Z_T^2 \\ &\times \frac{\left[M_N^2 (T_T - 2E_\nu - M_T) + 2M_T (2E_\nu^2 - 2T_T E_\nu + T_T (T_T - M_T)) \right]}{E_\nu^2 (M_{Z'}^2 + 2T_T M_T)^2}, \end{aligned} \quad (4.2.2)$$

where T_T is the target recoil kinetic energy, E_ν is the energy of the probe neutrino, M_N and $M_{Z'}$ are respectively the dark neutrino and Z' masses. The expression also depends on the kinetic mixing term ϵ , the dark gauge coupling α_D , the ν_μ -dark

neutrino mixing angle as well as on the coefficient

$$|U_{D4}|^2 \equiv 1 - |U_{e4}|^2 - |U_{\mu4}|^2 - |U_{\tau4}|^2, \quad (4.2.3)$$

with $|U_{D4}|^2 = 1$ if unitarity is imposed in the dark mixing.

The form factor term FF in eqn. (4.2.2) is estimated, from [157], as:

$$FF(q) \equiv \frac{3j_1(qR_0)}{qR_0} \exp \left[-\frac{1}{2}(qs)^2 \right], \quad (4.2.4)$$

with $R_0^2 \equiv R^2 - 5s^2$, where $R = 1.2A^{1/3}$ fm and $s \sim 1$ fm is a scale factor. Despite the simplified expression for the form factor, the DarkNeutrino and DarkNews differential cross sections are consistent for $Q^2 < 1$ GeV, which make-up the large majority of the statistics, as fig. 4.9 shows.




Figure 4.9: Comparison between the differential cross sections (as $d\sigma/dQ^2$) on ^{12}C , for a 1 GeV probe neutrino, computed in DarkNews and in the DarkNeutrino module of GENIE. A model with $M_N = 106$ MeV, $M_{Z'} = 30$ MeV, $|U_{\mu4}|^2 = 2 \times 10^{-10}$, $\epsilon = 8 \times 10^4$ is assumed [87]. The two functions are consistent for $Q^2 \lesssim 2$ GeV, at which point the cross section is significantly reduced.

The $d\sigma/dT_T$ differential cross sections are integrated over the dark neutrino energy $E_N = E_\nu - T_T$ operating a corresponding change of variable. The integration limits are found requiring $E_N > M_N$ and $\cos\theta_N(E_N) \in [-1; 1]$. A correction to the integration range implemented in DarkNeutrino, required to recover the correct physical behaviour of total cross sections in coherent scattering, is described in sect. 4.2.3.

At the time of writing, DarkNeutrino only models decays to an on-shell Z' and SM neutrino. As per [91], the decay rate is implemented as:

$$\Gamma_{N \rightarrow Z' + \nu} = \frac{\alpha_D}{2} |U_{\mu4}|^2 (1 - |U_{\mu4}|^2) \frac{M_N^3}{m_{Z'}^2} \left(1 - \frac{m_{Z'}^2}{M_N^2} \right) \left(1 + \frac{m_{Z'}^2}{M_N^2} - 2 \frac{m_{Z'}^4}{M_N^4} \right). \quad (4.2.5)$$

The subsequent decay of Z' to e^+e^- is modelled, with a decay rate

$$\Gamma_{Z' \rightarrow e^+e^-} = \frac{\alpha_{EM}\epsilon^2}{3} m_{Z'} \sqrt{1 - 4 \frac{m_e^2}{m_{Z'}^2}}, \quad (4.2.6)$$

which would be applicable for decays to $\mu^+\mu^-$ and $\tau^+\tau^-$ as well.

Module workflow. The generation workflow with DarkNeutrino closely follows the steps presented in sect. 4.1.2. The module interactions are available in a specific tune, GDNu20_01a_00_000, so that signal-only productions can be generated. A GENIE configuration file allows to set the model parameters before generating the splines and launching the simulation.

All particles involved in the scattering and subsequent decays are recorded in the generator output, with the recoil nucleus and e^+e^- pair, in particular, treated as *stable final state* particles and thus propagated by edep-sim. Unlike standard interactions, dark-neutrino decays can occur with a macroscopic displacement, so that a correction to edep-sim was required to allow the placement of primary particles away from the initial interaction point [158].

4.2.3 Correction to the DarkNeutrino integrator

Upon first examination of the DarkNeutrino simulation output, the total coherent scattering cross sections were found to quickly decrease with E_ν , after a peak, with the unphysical behaviour shown in fig. 4.10 against the correct behaviour from DarkNews.

As the differential cross sections were consistent with DarkNews, the issue could be traced back to the *relative tolerance* setting of the GENIE numerical integrator, based on the QUADPACK package [159]. The falling behaviour of the cross section was caused by integration over the higher values in the $Q^2 = 2M_T T_T = 2M_T(E_\nu - E_N)$ range, for which $d\sigma/dQ^2$ falls to small values. The behaviour could be tested independently with the QUADPACK-based integrator in `scipy` [153].

A general fix was implemented in the form of an upper cut of $Q^2 < 1 \text{ GeV}^2$, as done in DarkNews: contributions from $d\sigma/dQ^2$ at higher Q^2 were deemed to be negligible. A comparison between the corrected GENIE and DarkNews cross sections is shown in fig. 4.10: the residual differences are determined to stem from the difference in the implementation of the form factors and may be accounted for by reweighting the event statistics.

The correction, implemented on a fork [160], has been included in the recent (at the time of writing) GENIE v3.6.2 production release.





Figure 4.10: Comparison of the total ν_μ upscattering cross sections on (a) ^{12}C and (b) ^{40}Ar obtained with DarkNews (blue), GENIE v3.2 (orange), and the GENIE fork implementing the cut at $Q^2 < 1 \text{ GeV}^2$ (green). The results refer to a model with $M_N = 420 \text{ MeV}$, $M_{Z'} = 30 \text{ MeV}$, $|U_{\mu 4}|^2 = 9 \times 10^{-7}$, and $\epsilon = 1.66 \times 10^{-4}$.

Chapter 5

Simulation of dark neutrino events in SAND

This analysis aims at evaluating the discovery potential of dark neutrino models through a pipeline consistent, as closely as possible, with the current implementation of the SAND simulation and reconstruction framework. Thus, the DarkNeutrino generator is employed, as, while limited to coherent scattering events with on-shell $Z' \rightarrow e^+e^-$, it allows event generation and propagation with the standard SAND geometry and pipeline. The event signature is limited to e^+e^- pairs in the SAND inner tracker, allowing an optimal charge separation and reconstruction of the pair kinematics. An ad-hoc track reconstruction algorithm is developed to account for the propagation of electrons and positrons. The selection of signal and background events and the estimation of the discovery potential Z for each dark neutrino mass model is articulated in four main steps:

1. The definition of signal and background for the search by applying topological cuts, described in sect. 5.3. Since event classification algorithms are not available at present in SAND, these are based on the edep-sim Monte Carlo information.
2. The reconstruction of candidate e^+e^- pair tracks, discussed in sect. 5.4, and the extraction of characteristic pair kinematical features.
3. A multivariate classification based on the reconstructed pair features using the ROOT TMVA library, presented in sect. 6.2.
4. The estimation of the discovery potential for the model, detailed in sect. 6.3.

A reference set of dark neutrino models, with $M_{Z'} = 30$ MeV and dark neutrino masses varying between 35 MeV and 1 GeV was chosen for the analysis, as it currently features the most stringent experimental constraints [87]. The analysis pipeline will, nonetheless, be applicable to additional sets of models.

5.1 Event generation

As introduced in Chap. 4, the generation of dark neutrino events in SAND required several additions and modifications to the standard simulation chain. It is therefore

best to detail the setup adopted for producing the datasets used in the analysis, and the additions made to the standard simulation chain.

5.1.1 Geometry and generation volumes

The analysis used the standard SAND geometry in its drift-chamber (DC) tracker configuration, as described in sect. 4.1.1. The track-reconstruction algorithm in sect. 5.4 is compatible with both the main straw-tube tracker (STT) design and the DC option. As such, the present framework may be applied to the STT configuration with only minor changes.

Because the on-shell Z' models currently implemented in DarkNeutrino, as well as the relevant background topologies (sections 5.2 and 5.3), yield short flight lengths for the e^+e^- pair production, event generation was initially restricted to the inner SAND volume (GRAIN plus the tracker) to maximise statistics for a fixed production size. Within this volume, the tracker features a total of 81 tracking stations, corresponding to 72 polypropylene (C_3H_6) and 8 graphite target layers. The corresponding 1H and ^{12}C target masses, together with those of the GRAIN ^{40}Ar active target and the passive material from the support structures are listed in Table 5.1.

Tracker target masses	
^{12}C (C/C_3H_6 target layers)	3.70×10^3 kg
^{12}C (supermodule frames)	6.47×10^3 kg
1H (C_3H_6 target layers)	488 kg
GRAIN target masses	
^{40}Ar (active target)	1.19×10^3 kg
^{12}C (outer vessel)	210 kg
^{56}Fe (inner vessel)	302 kg

Table 5.1: Masses of the main targets in the GRAIN and drift-chamber tracker geometry used for the analysis.

Rate of scattering on hydrogen

The lack of an implementation in DarkNeutrino for the scattering cross section on hydrogen will necessarily lead to an underestimation of the dark neutrino event rate in SAND. However, since the dark neutrino cross section on free protons is lower than that on other target materials, as illustrated in fig. 5.1, this contribution is not expected to be significant, as discussed below.

The scattering rate on 1H from the tracker C_3H_6 layers was estimated quantitatively using the DarkNews generator, and compared to the corresponding rates on ^{12}C in the target layers and ^{40}Ar in GRAIN. For each target mass (1H , ^{12}C , and ^{40}Ar , as listed in Table 5.1) and for each neutrino flavour, datasets of $N_{prod} \sim 5 \times 10^3$ weighted events were produced in DarkNews, assuming a 10^{21} POT exposure of the DUNE-ND flux in FHC mode (see sect. 2.1.1). The datasets were generated for 19 benchmark models with dark neutrino masses M_N in the range 10^{-2} GeV–1 GeV, and with fixed parameters $M_{Z'} = 30$ MeV, $|U_{\mu 4}|^2 = 2 \times 10^{-10}$, and $\epsilon = 8 \times 10^4$.




Figure 5.1: Comparison between the cross sections for the main target nuclei in the SAND inner volume as a function of the probe neutrino energies, computed in DarkNews. A model with $M_N = 106$ MeV, $M_{Z'} = 30$ MeV, $|U_{\mu 4}|^2 = 2 \times 10^{-10}$, $\epsilon = 8 \times 10^4$ is assumed [87].

In DarkNews, the expected number of events in each target and its statistical uncertainty at the given exposure is obtained by summing the event weights (see sect. 4.2.1) in the corresponding datasets over all beam neutrino flavours:

$$N_{evt}^{tgt} = \sum_{\alpha=\nu_e, \bar{\nu}_e, \nu_\mu, \bar{\nu}_\mu} \sum_{i=0}^{N_{prod}} w_{\alpha i}^{tgt}, \quad \Delta N_{evt}^{tgt} = \left[\sum_{\alpha=\nu_e, \bar{\nu}_e, \nu_\mu, \bar{\nu}_\mu} \sum_{i=0}^{N_{prod}} (w_{\alpha i}^{tgt})^2 \right]^{1/2}. \quad (5.1.1)$$

As shown in fig. 5.2, the event rate on the ^{12}C and ^{40}Ar targets remains at about 10^3 per 10^{21} POT for $M_N < 100$ MeV, before decreasing at higher masses. By contrast, the number of events expected in the 488 kg ^1H tracker target mass is suppressed by $\mathcal{O}(10)$ relative to ^{12}C and ^{40}Ar across the considered range of M_N . Consequently, the ^1H contribution to the statistics is neglected in this preliminary sensitivity study. Nevertheless, the implementation of upscattering cross sections on ^1H in DarkNeutrino remains a key priority for a potential next stage of the analysis.

5.1.2 Signal productions

The GENIE version used for the signal productions is a dedicated fork [160] that implements the cross-section integration correction discussed in sect. 4.2.3. As described in sect. 4.2.2, dark neutrino event samples are generated in GENIE with the DarkNeutrino-specific tune GDNu20_01a_00_000. For each set of model parameters, a corresponding cross-section spline file [133] must be prepared in advance. For the present analysis, a set of 23 models with masses in the $35 \text{ MeV} \leq M_N \leq 106 \text{ MeV}$




Figure 5.2: Estimated number of events at 10^{21} POT in FHC mode for dark neutrino masses M_N between 10^{-2} and 1 GeV and for the main SAND target masses: the GRAIN ^{40}Ar volume and the ^1H and ^{12}C tracker target masses. Models with $M_{Z'} = 30$ MeV, $|U_{\mu 4}|^2 = 2 \times 10^{-10}$, $\epsilon = 8 \times 10^4$ are assumed [87]. Error bars are not visible due to the logarithmic scale.

range was considered, fixing the other parameters as $M_{Z'} = 30$ MeV, $|U_{\mu 4}|^2 = 2 \times 10^{-10}$, and $\epsilon = 8 \times 10^4$.

A fixed number of events ($\sim 5 \times 10^5$) is generated for each configuration to ensure sufficient statistics regardless of the interaction cross section. The effective beam exposure associated with each generation is then extracted from the GENIE output and applied as an event weight in the sensitivity estimation. Productions include both $\bar{\nu}_\mu^{(-)}$ and $\bar{\nu}_e^{(-)}$ with DUNE-ND fluxes in the standard `gsimple` format taken from [137] corresponding to the November 2017 data release. A Forward Horn Current (FHC) beam mode is used in the productions.

Particle propagation is performed, starting from the `numi_rootracker` GENIE output, using the public fork of `edep-sim` introduced in sect. 4.2.2 [158], which includes a fix allowing primary particles to be placed away from the interaction vertex. In parallel with the propagation stage, the GENIE `numi_rootracker` output is converted into a `MC_truth` ROOT file that stores the relevant primary particle information, such as event cross sections, lab-frame positions, and four-momenta, in an easily accessible format (i.e. using ROOT four-vector types).

5.1.3 Background productions

Background datasets are generated with a GENIE v3.6.0 installation using the standard `AR23_20i_00_000` tune with the corresponding cross section spline file. For the sake of consistency, particle propagation is performed with the same `edep-sim` in-

stallation used for the signal productions and discussed above. As with signal events, background generation is limited within the SAND inner volume and performed using the standard `gsimple` flux files with a FHC beam configuration.

5.2 Signal event features

A preliminary step in the analysis, following the implementation of the necessary corrections to the simulation chain, was to examine the signal event topologies and kinematic distributions in order to identify potential backgrounds and the corresponding optimal classification features.

A full description of the signal events at the Monte Carlo level is obtained by combining the `edep-sim` events with the GENIE output, since the probe neutrino, the dark neutrino N , and the on-shell Z' are not propagated in `edep-sim`. `HNLUpScattering` objects were defined to include both the GENIE event-particle information and pointers to the `edep-reader` trajectories of the e^+ and e^- . Through the `edep-reader` library, trajectory points and hit segments for the lepton pair can be accessed.

5.2.1 Interaction volumes

Estimating the relative interaction rate in each component of the SAND detector is crucial for determining the optimal event reconstruction and selection strategy. The distribution of scattering vertices among the main SAND volumes, obtained from a GENIE sample of $\sim 5 \times 10^5$ events generated within the GRAIN and tracker volumes for a benchmark model with $M_N = 106$ MeV, are shown in fig. 5.3. Consistently with the fast simulation in sect. 5.1, the majority of interactions occur within the graphite drift-chamber frames, followed by the overall tracker target ^{12}C mass and the GRAIN ^{40}Ar target.

Because a majority of upscattering events occurs in the tracker and the decays to e^+e^- pairs are prompt, as detailed in the following, an optimal signature may be identified in an e^+e^- pair reconstructed as two oppositely charged tracks in the drift-chamber volume.

5.2.2 Flight lengths

As discussed in sect. 1.2.3, the GENIE implementation of dark-neutrino decays via an on-shell Z' is effectively prompt, with proper decay lengths $c\tau \ll 10$ cm, and produces forward-going final states. Since the present analysis involves a broad range of dark neutrino masses M_N , the flight lengths in both steps of the decay chain were explicitly quantified in order to validate, across the model grid, the initial choice of the SAND inner detectors as the interaction region and the use of an oppositely charged e^+e^- pair reconstructed as two tracks in the drift-chamber volume as the primary signature.

The rest frame decay length distributions for the $N \rightarrow \nu Z'$ and $Z' \rightarrow e^+e^-$ steps of the decay chain were obtained from $\sim 5 \times 10^5$ events distributions with M_N in the 35 MeV to 1 GeV range. The distributions were fitted to an exponential, with the resulting proper decay lengths $c\tau$ being shown in fig. 5.4.




Figure 5.3: Number of interactions in the SAND component volumes out of $\sim 5 \times 10^5$ events in FHC mode. Vertices occurring in components other than those listed are categorised as *Other*. A model with $M_N = 106$ MeV, $M_{Z'} = 30$ MeV, $|U_{\mu 4}|^2 = 2 \times 10^{-10}$, $\epsilon = 8 \times 10^4$ is assumed.




Figure 5.4: Proper flight lengths of the dark neutrino N and Z' . The models feature $35 \text{ MeV} \leq M_N \leq 1 \text{ GeV}$, $M_{Z'} = 30 \text{ MeV}$, $|U_{\mu 4}|^2 = 2 \times 10^{-10}$ and $\epsilon = 8 \times 10^4$. All models show a *prompt* topology, with $c\tau \ll 10 \text{ cm}$

As all considered models feature $m_{Z'} = 30 \text{ MeV}$, the Z' proper length shows a constant behaviour across the range, in accordance to eqn. (4.2.6). Conversely, the dark neutrino proper lengths show the reduction with M_N expected from eqn. (4.2.5).

A key observable for experimental signatures is the vertex displacement of the pair in the laboratory frame. Figure 5.5 shows the median laboratory-frame flight lengths of N and Z' , together with the resulting displacement of the e^+e^- pair from

the primary interaction vertex. The largest pair displacement occurs near the off-shell Z' threshold, $M_N \simeq M_{Z'} = 30 \text{ MeV}$, where it reaches $\mathcal{O}(1 \text{ cm})$, driven by the dark neutrino flight length.




Figure 5.5: Laboratory frame median displacements in the decay chain. The total e^+e^- pair-vertex displacement is superimposed. The models feature $35 \text{ MeV} \leq M_N \leq 1 \text{ GeV}$, $M_{Z'} = 30 \text{ MeV}$, $|U_{\mu 4}|^2 = 2 \times 10^{-10}$ and $\epsilon = 8 \times 10^4$.

As the clearance between the most downstream drift-chamber supermodule and ECAL is 12 cm [161], the fraction of scattering events in the tracker volume that produce e^+e^- within ECAL can be neglected, to a first approximation, even at the lower end of the M_N range. Therefore, event generation within the GRAIN and tracker volumes can be regarded as providing a representative sample of e^+e^- tracks in the drift chamber.

5.2.3 Signal event kinematics

In addition to the topological features that will be discussed in sect. 5.3, the selection of signal events will exploit the kinematical features of the visible products in the dark neutrino decay chain, specifically, the nuclear target and the e^+e^- pair in the processes that GENIE currently implements.

At the time of writing, only coherent scattering interactions are implemented in DarkNeutrino. Consequently, the kinetic energies of the nuclear recoil will be limited to small values, as fig. 5.6 shows. Across the $35 \text{ MeV} \leq M_N \leq 106 \text{ MeV}$ mass range of the considered models, the kinetic energy remains limited to values $T_{rec} \lesssim 10^{-1} \text{ MeV}$. The small kinetic energy of the recoils, together with the majority of interactions occurring in the passive volumes of the tracker, imply that any potential signature from the vertex interactions is not realistically observable. Conversely, the presence of vertex activity may be used to reject background events, as sect. 5.3 will discuss.

The dominant experimental signature of dark neutrino events will therefore be the e^+e^- pair produced in the Z' decay in the tracker volume. In addition to the absence of vertex activity, background rejection will primarily rely on the reconstructed pair




Figure 5.6: Median kinetic energy of the nuclear recoil T_{rec} for the models under consideration. The ranges between the 16% and 84% quantiles of the corresponding distributions are superimposed. All models feature $35 \text{ MeV} \leq M_N \leq 1 \text{ GeV}$, $M_{Z'} = 30 \text{ MeV}$, $|U_{\mu 4}|^2 = 2 \times 10^{-10}$ and $\epsilon = 8 \times 10^4$.

features. Since the tracker was selected as the fiducial volume for the search, a set of observables derived from the reconstructed e^+ and e^- momenta was chosen:

- the invariant mass M_{ee} of the pair, which, for on-shell Z' corresponds to $M_{Z'}$,
- the total pair energy E_{pair} ,
- the pair transverse momentum P_{ee}^T with respect to the beam direction,
- the opening angle θ_{ee} between the the electron and positron trajectories,
- the angle θ_{beam} between the pair three-momentum $\vec{P}_{pair}$ and the beam direction.

The median Monte Carlo values of said feature distributions, extracted from $\sim 5 \times 10^5$ datasets, are plotted in fig. 5.7 for the M_N range of interest. A significant increase in the pair energy E_{pair} is expected, in particular, for $M_N \gtrsim 600 \text{ MeV}$.

5.3 Event signature and backgrounds

Providing an estimate of the discovery potential for dark neutrino models in SAND requires the definition of a signature for the signal events, and the estimation of the corresponding search background.

In the absence, at present, of validated algorithms for vertexing, tracker-to-ECAL matching and event classification, the selection of such signature must utilise Monte Carlo in a significant amount, and should therefore be based on more comprehensive criteria than ideally possible.




Figure 5.7: Medians of the selected e^+e^- pair kinematics features for the models under consideration. The ranges between the 16% and 84% quantiles of the corresponding distributions are superimposed. All models feature $35 \text{ MeV} \leq M_N \leq 1 \text{ GeV}$, $M_{Z'} = 30 \text{ MeV}$, $|U_{\mu 4}|^2 = 2 \times 10^{-10}$ and $\epsilon = 8 \times 10^4$.

5.3.1 Topological selection cuts

As shown in sect. 5.2, upscattering interactions in the SAND inner volume will occur predominantly within the tracker components and produce an e^+e^- pair within the same detector. The choice was therefore made to limit the signature to an e^+e^- pair produced within the drift chamber volume.

In sect. 5.2.3, it was shown that the recoil kinetic energy is negligible for signal

events across the considered range of models, justifying the use of vertex activity to reject background events. As vertexing algorithms for SAND are still in the development stage, a proxy for vertex activity was found in the edep-sim primary trajectories, which are produced in the main neutrino interaction vertex. A selection cut can thus be defined from the number of such trajectories that produce hits in ≥ 4 tracker modules, a standard track quality criterion used for analysis in SAND [116]. The number of primary trajectories crossing more than 4 tracking stations for a sample of signal and background events is shown in fig. 5.8: as expected, a maximum of two trajectories, corresponding to the signature e^+e^- pair is found for signal events, while the distribution for standard events features more than two trajectories in $\sim 34\%$ of events.




Figure 5.8: Number of edep-sim primary trajectories crossing more than the 4 drift-chamber modules threshold in a benchmark production with $M_N = 100$ MeV. A subset of 2×10^4 events from both signal and standard productions is used. The sample includes trajectories originating outside the tracker volume. A fraction of 34% of standard events features ≥ 2 trajectories over threshold.

It is therefore possible to identify two main background topologies from the standard interaction sample:

1. events with no primary trajectories above the module threshold and with ≥ 1 e^+e^- pair above threshold produced in the tracker volume,
2. events with a single e^+e^- pair in the tracker volume among the primaries produced above threshold.

For both background topologies, an e^+e^- pair *candidate* for reconstruction in the tracker can be identified: the most upstream pair above threshold for topology (1) and the primary pair for topology (2).

A corresponding cut, selecting events with the primary e^+e^- pair being above threshold and produced in the tracker is applied to the signal dataset. The resulting

efficiency of the topological cut was estimated for the signal models and for the standard interaction sample, using $\sim 5 \times 10^5$ events datasets, and is shown in fig. 5.9. The signal efficiencies are found to increase with the model M_N parameter, and are above the background efficiency of $\sim 3.5\%$ for any considered model.




Figure 5.9: Efficiencies of the topological selection cut for the selection of signal and background events from the signal model and standard interaction datasets respectively. Both samples for signal and background contained $\sim 5 \times 10^5$ events.

5.3.2 Background kinematics

The Monte Carlo distributions of the pair kinematic features introduced in sect. 5.2.3 for the signal and background datasets resulting from the topological cuts are compared in fig. 5.10: in addition to the invariant mass, which for signal events will be $M_{ee} = M_{Z'} = 30$ MeV, a good separation is also found for the transverse momentum and pair-beam angle θ_{beam} angle distributions. The further rejection of background events therefore hinges on the performance of the track reconstruction algorithm presented in sect. 5.4 on the candidate e^+e^- pairs.

5.3.3 Sources of background

A basic assessment of the sources of background was performed on the events resulting from the topological cuts. Compared to the initial sample of standard interactions, the background sample is enriched in Neutral Current interactions, as fig. 5.11 shows: this is a consequence of the topological cut on primary trajectories, which enables the rejection of primary charged lepton tracks.

The parent particles of the candidate e^+e^- pairs were also extracted and are shown in fig. 5.12, where the parents of pair-producing photons are listed. The majority of background pairs are thus related to π^0 decays, in agreement with [87], with the next-to-largest fraction due to pair production in EM showers.




Figure 5.10: Monte Carlo pair kinematical feature distributions of the $M_N = 106$ MeV benchmark signal events and of the background.




Figure 5.11: Neutrino interaction type for the background events passing the topological selection cut. An enriched fraction of NC interactions compared to CC is observed, consistently with the cut on primary trajectories.




Figure 5.12: Parent particles of background pair candidates as recorded by edep-sim. For photon-initiated pairs, the parent of the photon is reported. Rarer parents are indicated as *Other*. A fraction of photon parents is mislabelled owing to the edep-sim propagation mechanism and is indicated as *EdepSim art.*

5.4 Reconstruction of e^+ and e^- tracks

To ensure that the classification of e^+e^- pairs based on kinematical features is realistic, the effects of detector resolution and track reconstruction performance must be taken into account. An *ad-hoc* track-fitting algorithm for reconstructing particle momenta from edep-sim hits was developed for this purpose. As a full reconstruction that includes corrections for energy-loss processes is still in development, at the time of writing, the algorithm assumes a helicoidal track model. The effects of this approximation are discussed below.

5.4.1 Helicoidal fit

A charged particle propagating through the SAND tracking volume will be subject to a 0.6 T magnetic field which will bend its trajectory in the zy plane, the z axis being at 5.3° from the beam direction and x being the direction of the magnetic field $\vec{B}$. In an homogeneous magnetic field, if the cumulative effect of Multiple Coulomb scattering and particle energy loss is negligible, as can be assumed to hold for minimum ionising particles like muons, the track can be parametrised as a helix with the path length s as the parameter:

$$\begin{cases} x_h(s) = x_0 - s \cdot \sin \lambda \\ y_h(s) = y_0 + R \cdot \left[\sin \left(\Phi_0 + \frac{h \cdot s \cdot \cos \lambda}{R} \right) - \sin(\Phi_0) \right] \\ z_h(s) = z_0 + R \cdot \left[\cos \left(\Phi_0 + \frac{h \cdot s \cdot \cos \lambda}{R} \right) - \cos(\Phi_0) \right] \end{cases}, \quad (5.4.1)$$

where:

- (x_0, y_0, z_0) is the starting point of the helix,
- λ is the *dip-angle*,
- Φ_0 is the angle between the point (x_0, y_0, z_0) and the z axis,
- R is the curvature radius of the helix,
- h is the sense of rotation of the particle trajectory.

The particle trajectory can thus be described as the superposition of and independent circular motion in the perpendicular plane to $\vec{B}$ (the *bending plane*)

$$(z_h - z_0 + R \cos \Phi_0)^2 + (y_h - y_0 + R \sin \Phi_0)^2 = R^2, \quad (5.4.2)$$

and a linear motion $x_h(s) = x_0 - s \cdot \sin \lambda$ along the $\vec{B}$ direction. Correspondingly, the track fit can be factorised into two components: in the bending plane and along the magnetic field direction. Using the bending-plane path length $s_\perp \equiv -s \cdot \cos \lambda$ as an independent parameter and being $z_c \equiv z_0 - R \cos \Phi_0$, $y_c \equiv y_0 - R \sin \Phi_0$ the coordinates of the centre of the circumference in the bending plane, the parametrisation reads as:

$$\begin{cases} x_h(s_\perp) = x_0 + s_\perp \cdot \tan \lambda \\ (z_h - z_c)^2 + (y_h - y_c)^2 = R^2 \end{cases} \quad (5.4.3)$$

The total number of parameters to be fitted is therefore five: $(z_c, y_c, R, \tan \lambda, x_0)$.

Preliminary processing and selection cuts

Once the candidate e^+e^- pair for each signal and background event had been selected according to the criteria in sect. 5.3, the track fitting algorithm was applied to the edep-sim hits of the corresponding trajectories. As trajectory hits are purely Monte Carlo objects, the effects of tracker digitization were simulated by introducing a gaussian smearing with a $\sigma_{res} = 0.2$ mm on all coordinates, corresponding to the expected final resolution in the SAND tracker [116].

Electron and positron propagation is subject to significant energy losses from Bremsstrahlung (the critical energy in ^{12}C is $E_c \simeq 46$ MeV), compared to minimum ionising particles [162]. Consequently, the helicoidal track model only approximates their trajectories. To mitigate the cumulative effects of radiative energy loss in the reconstruction, the fit was restricted to the initial segment of each track: specifically, the first 25% of available hits ($N_{hits}/4$) were used, with a lower bound of 10 hits, corresponding to approximately three drift-chamber tracking stations. Trajectories with fewer than 10 original hits were discarded, while for short tracks with $N_{hits}/4 < 10$, all hits were retained in order to ensure an adequate fit quality.

Track fitting algorithm

After the smearing and selection steps discussed above, the track hits undergo a multi-step track reconstruction algorithm. The fitting procedure, outlined in fig. 5.13, consists of:

1. *Preliminary pass*: two initial iterations of the bending-plane fit to initialise the fit parameters and estimate the path lengths $s_\perp$ per-hit.
2. *Full fit*: three subsequent iterations combining the bending-plane and x -axis fits, after which the solution was found to stabilise.

If any between the preliminary pass and the full-fit iteration fails to converge, the trajectory is flagged and discarded in the subsequent analysis stages. The implementation uses the ROOT `ROOT::Fit::Fitter` interface for generic model functions, with χ^2 minimisation performed by the ROOT `Minuit2` implementation of the MINUIT algorithm [163, 164]. Seeding and minimisation strategies specific to each step are detailed below.

Bending-plane fit. As the first step of the procedure, a seed for the bending plane circle centre (z_c, y_c) is obtained through a Least Squares fit method based on [165]. The mean distance of hits from such guess coordinates provides the seed for the radius R . A first iteration of the circular fit is then performed, and the resulting parameters provide an estimate of the bending-plane path length $s_\perp$ for each hit as:

$$s_\perp^{hit} = h \cdot R \cdot (\phi^{hit} - \Phi_0), \quad \text{with} \quad \phi^{hit} = \arctan \left(\frac{y^{hit} - y_c}{z^{hit} - z_c} \right), \quad (5.4.4)$$




Figure 5.13: Scheme of the track fitting algorithm. After a first pass, the bending plane fit is repeated introducing a dependence of χ^2 on $s_\perp$ in order to compensate the reduction in the curvature radius.

where Φ_0 is a reference angle defined by the intersection of the segment through (z_c, y_c) with the fitted circumference. The sense of rotation $h \in \{-1, +1\}$ is determined from the sign of the summed cross products of consecutive hit positions $\mathbf{r}_i^{\text{hit}} \equiv (0, y_i, z_i)$ relative to $\mathbf{r}_c \equiv (0, y_c, z_c)$:

$$h = \text{sgn} \left[\hat{\mathbf{x}} \cdot \sum_{i=0}^{n-1} (\mathbf{r}_{i+1}^{\text{hit}} - \mathbf{r}_c) \times (\mathbf{r}_i^{\text{hit}} - \mathbf{r}_c) \right]. \quad (5.4.5)$$

A further iteration of the circular fit is then performed, introducing an $s_\perp$ dependent penalty term in the per-hit variance:

$$\sigma_i^2 \rightarrow \sigma_i^2 + \alpha s_\perp + \beta s_\perp^2, \quad (5.4.6)$$

where $\alpha \equiv \sigma_{res}$ and β is a numerical parameter optimised to $\beta \equiv 0.2$. An additional weight $w_i \equiv \sigma_{res}/\Delta R_i$ inversely proportional to the residue between hit i and the fitted circumference is applied to reduce the impact of outliers. Such corrections aim at downplaying the contribution of later trajectory hits to the χ^2 , which would lead to an underestimation of the fit R due to energy loss. Reconstruction is stopped for tracks whose second fit pass did not converge.

Fit on the x-axis. Once the path lengths $s_\perp$ for each hit have been computed, the linear fit on the x-axis is performed using a standard χ^2 function for fitting, with the seeding obtained by linear regression on the hit x coordinates. Three iterations of

the full bending-plane and x-axis fit are performed, at which point the fit parameters were found to converge.

Momentum reconstruction. The modules of the particle momentum components on the bending-plane p_{yz} and on the magnetic field axis p_x are obtained from the fit R and $\tan \lambda$ parameters:

$$\begin{cases} p_{yz}[\text{GeV}] = 0.3 \cdot R[\text{m}] \cdot B[\text{T}] \\ p_x = p_{yz} \cdot \tan \lambda \end{cases} \quad (5.4.7)$$

where the magnetic field in SAND is $B = 0.6 \text{ T}$. In order to compute the separate p_y and p_z components, the helix equation must be propagated back to the track starting point (x_0, y_0, z_0) . In the absence, at the time of writing, of a vertexing algorithm, such point is taken as the first point of the Monte Carlo trajectory, so that:

$$\begin{cases} p_y = h \cdot p_{yz} \cdot \cos \theta_{MC} \\ p_z = h \cdot p_{yz} \cdot \sin \theta_{MC} \end{cases}, \quad (5.4.8)$$

with $\theta_{MC} \equiv \arctan((y_{MC} - y_c)/(z_{MC} - z_c))$.

5.4.2 Reconstruction performance

The performance of the reconstruction algorithm was first assessed using a sample of electrons from signal pairs with valid fits for both e^+ and e^- . A model with $M_N = 106 \text{ MeV}$ was used as a benchmark.

A comparison between the true and reconstructed values of the R and λ fit parameters is presented in figures 5.14 and 5.15 respectively, showing the corresponding resolutions obtained from the residuals. A significant bias towards $R_{reco} < R_{MC}$ is observed, consistently with the significant particle energy loss affecting electrons.

The bias towards lower $R_{reco} < R_{MC}$ is reflected in the reconstructed momentum components, compared to the corresponding true values in fig. 5.16: both bending plane components p_y and p_z are underestimated for a significant fraction of events. The percentage residuals for the bending plane momentum p_{yz} , p_x and for the electron energy are shown in figures 5.17 to 5.19.

A list of the resolutions on the relevant reconstruction parameters is presented in Table 5.2.

Reconstruction efficiency on signal and background events

As discussed in sect. 5.2.3, the kinematics of the e^+e^- pair is expected to vary across the $35 \text{ MeV} \leq M_N \leq 106 \text{ MeV}$ mass range under consideration. The performance of the track reconstruction algorithm can similarly be expected to vary with M_N . As a preliminary check, the reconstruction efficiency for the pairs was determined for both signal and background events that passed the topological cuts, requiring both e^+ and e^- fits to be valid and the reconstructed charges to have the correct sign. The resulting values for all models are shown in fig. 5.20: both signal and background efficiencies are found to be $\geq 90\%$, with a slight decrease in signal models with




Figure 5.14: Left: comparison between the true and reconstructed curvature radius R of signal electrons. Right: curvature radius residuals relative to R_{MC} expressed as a percentage. A bias for $R_{reco} < R_{true}$ is observed. A fit to a Crystal ball distribution yields a resolution $\sigma = (9.4 \pm 0.1) \%$.




Figure 5.15: Left: comparison between the true and reconstructed dip angle λ of signal electrons. Right: distribution of the λ residuals. The resolution is estimated with the HWHM as $\sigma = 3 \times 10^{-3}$ rad.

$M_N \gtrsim 500$ MeV. The reduction in signal efficiency for increasing M_N , corresponding to higher pair energies, see sect. 5.2.3, may be related to the increasing curvature radius, which may hamper convergence in a higher fraction of the track fits.




Figure 5.16: Comparison between the true and reconstructed momentum components for signal electrons. A bias for $p_i^{reco} < p_i^{MC}$ is visible for p_y and p_z .




Figure 5.17: Left: comparison between the true and reconstructed bending plane momentum component p_{yz} . Right: p_{yz} residuals relative to p_{yz}^{MC} expressed as a percentage. A Crystal Ball fit yields a resolution $\sigma = (9.4 \pm 0.1) \%$.

Reconstructed kinematics distributions

The residuals between the true and reconstructed pair kinematic features introduced in sect. 5.2.3 were obtained for the $M_N = 106 \text{ MeV}$ benchmark model and used to estimate the corresponding resolutions, listed in Table 5.3. As the corresponding distributions in fig. 5.21 show, a significant dispersion is found about the 30 MeV pair invariant mass. A Crystal Ball fit to the peak of the energy residuals relative to E_{pair}^{MC} shows a bias consistent with that on the single electron energy, with a corresponding




Figure 5.18: Left: comparison between the true and reconstructed p_x momentum component for signal electrons. Right: p_x residuals relative to p_x^{MC} expressed as a percentage. A Crystal Ball fit yields a resolution $\sigma = (18.6 \pm 0.3) \%$.




Figure 5.19: Left: comparison between the true and reconstructed signal electron energies E . Right: E residuals relative to E^{MC} expressed as a percentage. A Crystal Ball fit yields a resolution $\sigma = (9.4 \pm 0.1) \%$.

$(7.91 \pm 0.07)\%$ resolution. Together with the pair energy, the best resolution, of $(8.2 \pm 0.1)\%$, is found for the pair-beam angle θ_{beam} , which, being well-separated from the background distribution (see fig. 5.10), may become the optimal classification feature.

The limited resolution on the pair kinematic features corresponds to an increased superimposition between the signal and background distributions, as fig. 5.22 shows

Fit parameter resolutions	
Curvature radius R	$\sigma = (9.4 \pm 0.1) \%$
Dip angle λ (HWHM)	$\sigma = 3 \times 10^{-3}$ rad
Kinematics resolutions	
Bending plane momentum p_{yz}	$\sigma = (9.4 \pm 0.1) \%$
B-field axis momentum p_x	$\sigma = (18.6 \pm 0.3) \%$
Electron energy E	$\sigma = (9.4 \pm 0.1) \%$

Table 5.2: Reconstruction algorithm resolutions on the main fit parameters and momentum components from a sample of signal e^- in a $M_N = 106$ MeV benchmark model. The resolution on λ is estimated as the HWHM of the distribution.




Figure 5.20: Reconstruction efficiencies on signal and background events from the signal model and standard interaction datasets respectively. Events with successfully reconstructed pair e^+ and e^- tracks are selected.

Resolutions on the pair kinematic variables	
Invariant mass M_{ee} (HWHM)	6.95 MeV
Pair energy E_{ee}	$(7.91 \pm 0.07) \%$
Pair transverse momentum P_T	$(13.3 \pm 0.2) \%$
Pair opening angle θ_{ee}	$(16.6 \pm 0.2) \%$
Pair-beam angle θ_{beam}	$(8.2 \pm 0.1) \%$

Table 5.3: Resolutions of the reconstruction algorithm on the pair kinematic variables for $M_N = 106$ MeV signal events. The invariant mass resolution is given as the HWHM of the residuals, while for the other variables, the percentage resolutions resulting from a Crystal Ball fit are listed.

for the benchmark $M_N = 106$ MeV model. In the absence of a sufficiently selective minimal criterion for separation from one of the kinematics features, the decision was

made to perform the final selection step in the analysis via a multivariate classifier, taking all reconstructed kinematic variables as input. The implementation and performance of the classification algorithm and the resulting estimates for the discovery potential are presented in the next chapter.




Figure 5.21: Top: reconstruction residuals on the signal pair invariant mass. Bottom: residuals of the signal pair kinematic features relative to their true value, expressed as a percentage. The resolution on the features is estimated through fits to a Crystal Ball. The $M_N = 106 \text{ MeV}$ benchmark model event sample is used.




Figure 5.22: Reconstructed pair kinematic feature distributions of the $M_N = 106$ MeV benchmark signal events and of the background. Smearing on the reconstructed momenta leads to a significant superimposition between the invariant mass distributions.

Chapter 6

Discovery potential for dark neutrinos in SAND

This chapter presents the final steps of the analysis procedure used to evaluate the discovery potential of dark neutrinos in SAND. Once the signal and background events have been selected through the topological cuts in sect. 5.3 and the e^+e^- pair kinematic features have been reconstructed using the track fitting algorithm discussed in sect. 5.4, a final selection is performed using a multivariate classifier, the implementation and optimisation of which is discussed in sect. 6.2. The procedure used to determine the discovery potential from the full analysis pipeline is discussed in sect. 6.3 and the preliminary estimates are presented. A short summary of the dark neutrino models considered for the analysis simulation is given in sect. 6.1.

6.1 Simulated models

This first assessment in SAND focused on the dark neutrino models with $M_{Z'} = 30$ MeV used as a reference throughout Chap. 5. These are currently the best constrained [87] and thus provide a useful benchmark for the analysis. Moreover, since they feature a lighter dark photon compared to other tested models, a wider region of the parameter space can be probed within the present limitation to on-shell Z' decays in GENIE (see sect. 4.2.2).

As introduced in sect. 5.1.2, a set of 23 models with dark neutrino masses spanning from 35 MeV to 1 GeV was considered in this analysis in order to cover the parameter space as much as possible. The dark neutrino coupling $|U_{\mu 4}|^2$ was set to 2×10^{-10} for all productions: since the cross section in eqn. (4.2.2) scales with this parameter, the coupling corresponding to any desired significance is found by re-scaling the number of signal events resulting from the analysis. The parameters of the models considered for the analysis are summarised in Table 6.1.

Datasets of $\sim 5 \times 10^5$ events are generated for each model in the SAND inner volume detectors using the setup detailed in sect. 5.1.2.

Dark Neutrino model parameters	
Dark neutrino mass M_N	35 MeV – 1 GeV
Dark photon mass $M_{Z'}$	30 MeV
$ U_{\mu 4} ^2$ coupling	2×10^{-10}
$ U_{e 4} ^2, U_{\tau 4} ^2$ couplings	10^{-14}
Dark gauge coupling α_D	0.25
Kinetic mixing term ϵ	8×10^4

Table 6.1: Parameters of the dark neutrino models included in the analysis. As introduced in sect. 4.2.2, model parameters for each generation are set through a GENIE configuration file. Values of $|U_{e 4}|^2, |U_{\tau 4}|^2 \ll |U_{\mu 4}|^2$ are used, in accordance to [91].

6.2 Multivariate classification with TMVA

As discussed in sect. 5.4, no single reconstructed kinematic observable was found to be sufficiently discriminating between signal and background events. To exploit correlations among these variables, the final stage of the selection was based on a multivariate classification technique. A training and testing pipeline based on the TMVA ROOT library [166] was implemented and designed so that classifier model parameters and input datasets could be easily configured and extended. This section describes the TMVA implementation within the analysis pipeline, the preprocessing of the event datasets, the choice of model parameters, and the performance of the classifier on the test sample.

6.2.1 Training and application pipelines in TMVA

As usual for Machine Learning models, a classification analysis with the TMVA library consists of three separate phases:

1. the *training* phase, where the algorithm is trained on a representative dataset,
2. the *testing* phase, where the performance of the algorithm is benchmarked on a *test* dataset in order to validate the training,
3. the *application* phase, where the trained model is applied to a final target sample, assigning each event a classifier score corresponding to its likelihood of being signal or background.

The datasets used in each of the three steps must be independent, and sufficiently large samples are required in the *training* and *testing* phases to avoid overtraining the model. To this end, *auxiliary productions* of $\sim 5 \times 10^5$ events were generated for all signal models and the background using the same setups as the target productions described in sect. 6.1, to serve as input for training and testing the classifier.

TMVA model training and application

The ROOT files, in an `RDataFrame` format [135], containing the reconstructed pair kinematic features for signal and background events that result from the selection

and reconstruction steps in sections 5.3 and 5.4 were used as input for the TMVA training and application pipelines. Events in which either trajectory of the candidate pair has a failed fit were excluded as an initial filtering step.

The input datasets were loaded in TMVA and split into training and test sets. Multiple classifier methods could be trained and benchmarked on the same signal and background input samples. In the present analysis, the list of training variables and model configuration parameters was defined in an external configuration file, enabling a quick reconfiguration and consistent tracking of the different training experiments.

The final output of the TMVA model training is a *weight file* containing trained model parameters, which is applied to the target datasets, thus evaluating the score corresponding to each sample event. Signal events were thus selected by applying a selection cut on the score values. The procedure through which the cut score value was optimised will be discussed in sect. 6.3.1. A detailed description of the standard TMVA workflow is given in Appendix B.

6.2.2 Model implementation and training

Once the training and testing pipeline for the TMVA models had been implemented, it was possible to determine an optimal algorithm for background rejection over the entire range of signal models. As introduced at the beginning of the section, five pair kinematic features were used as classification variables: (a) the pair invariant mass, (b) the pair energy, (c) the pair transverse momentum, (d) the pair aperture angle and (e) the pair-to-beam angle.

A single classifier was trained on all signal models (*pooled training*) to provide a single final selection algorithm that would not be biased towards any specific mass. Since the training datasets are not homogeneous in number nor in spacing in the M_N range, a weight was applied to signal dataset during training:

$$w_i = \frac{1}{N_i^{\text{tree}}} \cdot \frac{\log m_{i+1} - \log m_{i-1}}{\log m_{\text{final}} - \log m_{\text{initial}}}, \quad (6.2.1)$$

where N_i^{tree} is the number of events in the sample for m_i . The logarithmic spacing term compensates for the approximately log-distributed mass points (e.g. 50–100, 200–400 MeV), providing uniform weighting across the full mass range. As the different samples are added in TMVA sequentially, a random sampling is performed when splitting the training and test sets. A standard split into 80% for *training* and 20% for *test* was chosen.

Owing to their widespread use in signal-background classification in HEP, fast training performance and resistance to *overtraining* through *tree pruning* [166], the choice of classifier models was focused on *Boosted Decision Trees* (BDT) [167]. A set of five BDT models was trained and tested on the pooled sample, providing a limited span over the main TMVA configuration parameters:

- **MaxDepth**: the maximum depth of the BDT decision tree (default: 3),
- **nCuts**: the number of grid points in variable range used for finding the optimal cut (default: 20),

- the boosting type used by the algorithm (default: Adaptive Boost [168]),
- the use of **Bagging**, where a random subset of the events (50% in the corresponding model) is used for growing the trees in each iteration.

Three additional classifiers, a Multi-Layer Perceptron, a Fisher discriminant and a Projective Likelihood estimator were also trained with their default TMVA configurations to provide a reference for the BDTs. A summary of the tested models, which include the default TMVA BDT configuration is provided in Table 6.2.

Tested BDT models	MaxDepth	nCuts	BaggedBoost	BoostType
kBDT_def	3	20	no	AdaBoost
kBDT_Grad	3	50	no	GradBoost
kBDT_depth4	4	50	no	AdaBoost
kBDT_Grad_depth4	4	50	no	GradBoost
kBDT_bagging	3	50	yes	AdaBoost
Reference classifiers				
kLikelihood	Projective likelihood classifier			
kFisher	Fisher discriminant classifier			
kMLP	Multi-Layer Perceptron model			

Table 6.2: Configuration of the BDT models and reference classifiers tested for the analysis. Only the parameters that were varied compared to the default TMVA configuration are listed. In models with a GradBoost boost type, the **MinNodeSize** and **Shrinkage** parameters are set to 2.5% and 0.1 respectively (compared to the default 5% and 1.0).

Performance

A key figure for evaluating the performance of a trained classifier is its ROC curve, which plots the *signal efficiency* (ϵ_S) versus the *background rejection* ($1 - \epsilon_B$) as the score threshold varies. Since both axes lie in the $[0, 1]$ range, the ROC of an optimal classifier will have an integral, or *Area Under the Curve* (AUC), approaching 1, with random guessing yielding $AUC \sim 0.5$. The AUCs of the considered classifier methods are plotted in fig. 6.1: the BDT classifiers show the best performance, with values > 0.97 . All models reach AUC values above 0.9, consistent with a correct and stable training.

Despite being a valid metric for the global model performance, the AUC does not guarantee the highest discovery potential. Consequently, the final model choice was made by maximising the expected significance on the test sample for a benchmark model, as described in sect. 6.3.2.

Checks for overtraining

A basic indicator of overtraining is the discrepancy between the training and testing efficiency curves. When the number of degrees of freedom is too small relative to the model parameters, the model may appear to perform artificially well on the training set compared to the independent test set. An example of the efficiency




Figure 6.1: AUCs of the tested classifiers. `kBDT_Grad_depth4` and `kBDT_depth4` feature the better values cross the model range. The BDT models are found to be the best performing category.

curves behaviour is given in fig. 6.2 for the `kBDT_def` configuration. The training and test efficiencies for both signal and background are well aligned, indicating an unbiased training.




Figure 6.2: Efficiency curves as a function of the classifier score for the `kBDT_def` configuration trained on the $M_N = 106$ MeV signal model. The curves for the training and test sets are consistent for both signal and background.

As a further check, the residual Δ_{AUC} was computed as the difference between the AUCs of the training and test set: significantly higher training AUC compared to the

test AUC would indicate overtraining. The residual, normalized to the test set AUC, is reported for all considered models in fig. 6.3. The residual remains below a factor of $\sim 0.1\%$ of the test AUC for all classifiers, indicating no significant overtraining.




Figure 6.3: Maximum AUC residuals relative to AUC_{test} over the mass range for the tested classifiers. The maximum discrepancy of $\sim 0.05\%$, found for kBDT_Grad_depth4, do not indicate any significant overtraining in any of the configurations.

6.3 Estimation of the discovery potential

The estimation of the discovery potential for the dark neutrino models under study adopted a counting experiment approach, with the significance for each model Z_N being determined by the numbers S_N and B of signal and background events passing the full set of selection cuts. The model and working point of the TMVA classifier used for the last stage of the selection was thus chosen to optimise the significance according to a global metric. This section will first present the procedure for determining the optimal selection cuts, followed by the choice of classifier. The resulting full analysis pipeline will thus be used to estimate the discovery potentials for the considered dark neutrino models.

6.3.1 Score optimisation

Having assessed the training performance of the classifiers, a working point for each model, i.e. a classifier score threshold t , must be determined to provide the selection cut used to estimate Z_N . A standard choice for the working point is the score t^* that maximises the test-set significance Z_{test} , computed with the asymptotic formula in [169]:

$$Z_{test}(t) \equiv \sqrt{2 \left[(S_{test}(t) + B_{test}(t)) \ln \left(1 + \frac{S_{test}(t)}{B_{test}(t)} \right) - S_{test}(t) \right]}, \quad (6.3.1)$$

with $t^* \equiv \arg \max_t Z_{test}(t)$. The discovery significance is then computed on the independent *target* dataset using the same expression evaluated at the fixed threshold, $Z_{tgt}(t^*)$. The expected number of signal and background events S_{test} and B_{test} corresponding to the test set $\mathcal{D}_{test}$ is determined from the corresponding score distributions (see fig. 6.4) for the desired exposure E_Z as:

$$S_{test}(t) = \frac{E_Z}{E_{prod}^S \cdot f_{test}} \cdot \sum_{i \in \mathcal{D}_{test}^S} \mathbf{1}\{t_i \geq t\}$$

and

$$B_{test}(t) = \frac{E_Z}{E_{prod}^B \cdot f_{test}} \cdot \sum_{i \in \mathcal{D}_{test}^B} \mathbf{1}\{t_i \geq t\}, \quad (6.3.2)$$

where the sums are the numbers $N_{test}^S(t)$ and $N_{test}^B(t)$ of test set events with scores above the t threshold, E_{prod}^S and E_{prod}^B are the exposures of the signal and background datasets and f_{test} is the split fraction for the test set.




Figure 6.4: Test set response distributions for kBDT_def at $M_N = 106$ MeV. The signal and background histograms are not normalised: the additional factor $E_Z/(E_{prod} \cdot f_{test})$ is applied when computing Z_{test} .

The exposures $E_{prod} \cdot f_{test}$ corresponding to the test-set samples for all signal models and for the background are shown in fig. 6.5. Given the factor of $\gtrsim 10^4$ between the background and signal exposures, the significance is expected to be maximised at a t^* score corresponding to a maximal background rejection, $1 - \epsilon_B \simeq 1$.

In the present analysis, $Z_{test}(t)$ is evaluated on a discrete grid of thresholds by binning the signal and background score distributions, allowing to retrieve t^* from the maximum of the resulting curve, as fig. 6.6 illustrates. As $Z_{test}(t)$ is defined to be 0 for $B_{test}(t) = 0$, the sharp drop of the significance beyond t^* is consistent with the maximum corresponding to a maximal background rejection, as noted above. With additional statistics, a finer binning of the response distributions will be possible, enabling a more precise characterisation of the significance peak.




Figure 6.5: Equivalent exposures for the test-set samples of signal and background. The $f_{test} = 0.2$ split factor is applied to the original production exposure. The factor $\gtrsim 10^4$ between the background and signal exposures leads to a maximal significance for $1 - \epsilon_B \simeq 1$.




Figure 6.6: Test set significance $Z_{test}(t)$ computed from the response distributions for kBDT_def at $M_N = 106$ MeV at a reference 10^{21} POT exposure. $Z_{test}(t)$ is set to 0 for $B_{test}(t) = 0$. The maximum significance corresponds to $1 - \epsilon_B \simeq 1$.

In order to keep the selection cut model-independent, a single global working point was defined for each one of the classifiers. The significance $Z_{test}^m(t)$ was computed from each mass model subset of the response distributions, and a global t_{glob}^* was chosen to maximise the minimum $Z_{test}^m(t)$ across the model range. The dependence on t of the significance for each mass model is shown in fig. 6.7 for kBDT_def: the maximum for $Z_{test}^m(t)$ was found to be reached for values $\sim t_{glob}^*$ in all mass models. The global working point t_{glob}^* is thus close to optimal for all mass models.

Although all signal samples were generated at a common coupling $|U_{\mu 4}|^2$, the




Figure 6.7: Test set significance $Z_{test}^m(t)$ as a function of the score cut for all masses for kBDT_def trained with pooled signal samples. Binning on the y-axis corresponds to the separate model masses. The maximum significance is reached for $t \simeq 0.26$ consistently across the mass range.

expected signal yield S_{test} scales linearly with $|U_{\mu 4}|^2$ through the interaction cross section. In principle, this could cause the optimal threshold t^* to depend on the assumed coupling. However, as shown in sect. 6.3.3 for the selected classifier, t^* is found to remain constant across the tested coupling range.

6.3.2 Model selection

The final choice of model was based on the working point test-set significance $Z_{test}^m(t^*)$ reached for a benchmark dark neutrino mass, $M_N = 106$ MeV. The significances reached on the benchmark mass sample by each classifier model are shown in fig. 6.8: kBDT_Grad_depth4 was found to be the best performing. Overall, BDT models are shown to perform better than the reference alternative classifiers. It must be noted that, since the test set has been used to optimise the scores t_{glob}^* , the significances listed for the ranking will, in general, overestimate those reached on the final *target* sample.

The test-set signal efficiencies ϵ_S for kBDT_Grad_depth4 at the working point t_{glob}^* for all signal samples were obtained from the corresponding response distributions, and are shown in fig. 6.9 together with the test-set background rejection $(1 - \epsilon_B)$. The t_{glob}^* score is found to correspond to the near-complete rejection of background, consistently with the significantly higher equivalent exposures of the signal samples compared to the background. A reduction of ϵ_S down to values of $\sim 40\%$ is found for masses $M_N \gtrsim 100$ MeV.

The relative importance of the input kinematic variables for the chosen classifier was also checked: for BDTs in TMVA, this is expressed as the weighted separation

gain (based on the Gini index) accumulated over all splits using the given variable [166, 170]. The variables are ranked in Table 6.3 in order of importance: the pair-beam and opening angles are found to be the most discriminating, followed by the invariant mass. The lower importance of the pair energy and transverse momentum may be related to their variation across the mass range, as shown in sect. 5.2.3.




Figure 6.8: Optimised test set significances $Z_{\text{test}}^m(t_{\text{glob}}^*)$ obtained on the $M_N = 106$ MeV sample for the tested classifiers. The best significance is reached by kBDT_Grad_depth4.




Figure 6.9: Test-set efficiencies for the selected kBDT_Grad_depth4 classifier at the working point t_{glob}^* for all signal samples. The test-set background rejection is overlaid.

Variable	Variable importance
Pair-beam angle θ_{beam}	0.27
Pair opening angle θ_{ee}	0.25
Invariant mass M_{ee}	0.22
Pair transverse momentum P_T	0.16
Pair energy E_{ee}	0.11

Table 6.3: Classification variables ranked by their importance coefficient for the selected kBDT_Grad_depth4 model.

6.3.3 Discovery potential

Having determined the best-performing classifier and a corresponding optimal working point t^* , the discovery potential of SAND for the dark neutrino models under study was determined by running the full analysis pipeline on $\sim 5 \times 10^5$ event *target datasets* for the background and all signal models. The final target $Z_{tgt}(t^*)$ significance was then obtained for the given exposure E_Z with an expression analogous to eqn. (6.3.1):

$$Z_{tgt}(t^*) \equiv \sqrt{2 \left[(S_{tgt}(t^*) + B_{tgt}(t^*)) \ln \left(1 + \frac{S_{tgt}(t^*)}{B_{tgt}(t^*)} \right) - S_{tgt}(t^*) \right]}, \quad (6.3.3)$$

where the selection cut score has been fixed to t^* and $S_{tgt}(t^*)$, $B_{tgt}(t^*)$ are given by:

$$S_{tgt}(t^*) = \frac{E_Z}{E_{prod}^S} \cdot \sum_{i \in \mathcal{D}_{tgt}^S} \mathbf{1}\{t_i \geq t^*\}, \quad B_{tgt}(t^*) = \frac{E_Z}{E_{prod}^B} \cdot \sum_{i \in \mathcal{D}_{tgt}^B} \mathbf{1}\{t_i \geq t^*\}. \quad (6.3.4)$$

Rescaling of $|U_{\mu 4}|^2$

In order to determine the discovery potential curve in the $M_N - |U_{\mu 4}|^2$ plane for a given value of the significance Z_{curve} , the coupling was rescaled from the standard $|U_{\mu 4}|_{prod}^2 = 2 \times 10^{-10}$ value of the productions to match Z_{curve} . Referring to sect. 6.1, only $S_{tgt}(t)$ will depend on $|U_{\mu 4}|^2$, allowing to determine the target coupling numerically by bracketing the solution of eqn. (6.3.3) [171].

As mentioned in sect. 6.3.1, the test-set signal yield S_{test} will introduce a dependence on $U^2 \equiv |U_{\mu 4}|^2$ for t^* . The behaviour was quantified on the selected kBDT_Grad_depth4 model by rescaling $S_{test}(t)$ by a factor U^2/U_{prod}^2 for 50 couplings distributed logarithmically between 10^{-14} and 10^{-4} , and extracting t^* according to sect. 6.3.1. The t^* score was found to be constant over the range: while the constant value is a clear artefact of the finite spacing in t , it implies a negligible variation across the entire range.

Decay length correction

An additional dependence of the significance on U^2 is introduced by the dark neutrino decay rate in eqn. (4.2.5). Because of the $\Gamma_{N \rightarrow Z' + \nu} \propto U^2 (1 - U^2)$ dependence, the

flight lengths increase for smaller couplings compared to the production value and may exceed the size of the detector. To estimate this effect for each dark neutrino mass M_N , the proper dark neutrino flight lengths λ_M^{prod} obtained in sect. 5.2.2 were used. The proper lengths at the coupling U^2 was thus obtained via scaling by the production coupling:

$$\lambda_M(U^2) \simeq \lambda_M^{prod} \cdot \frac{U_{prod}^2}{U^2}. \quad (6.3.5)$$

This allowed to compute the mean dark neutrino decay probability within a characteristic detector length l_{det} in each signal production sample:

$$P_M(l_{det}, U^2) \simeq \frac{1}{N} \sum_i \left(1 - \exp \left[-\frac{l_{det}}{\gamma_i \lambda_M(U^2)} \right] \right). \quad (6.3.6)$$

A length $l_{det} = 1$ m, corresponding to roughly one third of the tracker length, was chosen, and the decay probabilities were computed for all masses and 100 couplings logarithmically spaced between 10^{-14} and 10^{-9} : the results are shown in fig. 6.10. The probabilities for the $M_N = 35$ MeV model, exhibiting the longest flight lengths, are highlighted: while $P(l_{det}) \simeq 1$ for all masses at the production coupling $U_{prod}^2 = 2 \times 10^{-10}$, a probability of $\sim 90\%$ for $M_N = 35$ MeV is found for $U^2 \simeq 3 \times 10^{-11}$.




Figure 6.10: Top: Mean probability for the dark neutrino to decay within $l_{det} = 1$ m as a function of the coupling U^2 in the range 10^{-14} and 10^{-9} , for all masses considered. Bottom: highlight of the $M_N = 35$ MeV case, which exhibits the longest decay lengths. A probability of $\sim 90\%$ is reached for $|U_{\mu 4}|^2 \simeq 3 \times 10^{-11}$.

The resulting probabilities were used to define an acceptance-related scaling factor

$$A_\lambda(M, U^2) \equiv \frac{P_M(l_{det}, U^2)}{P_M(l_{det}, U_{prod}^2)} \quad (6.3.7)$$

that was computed, for all masses, on a grid of 100 couplings logarithmically spaced between 10^{-14} and 10^{-9} . The resulting matrix was applied to the signal yield $S_{tgt}(t)$ when rescaling $|U_{\mu 4}|^2$, using a linear interpolation for $A_\lambda(U^2)$.

While the present scaling factor is an approximate value based on a single characteristic length, a more thorough parametrisation considering the tracker geometry and the distribution of event vertices within it may be performed.

Discovery potential curves

The curves for the discovery potential in the $M_N - |U_{\mu 4}|^2$ plane at $Z = 1.645$, corresponding to a one-sided 95% CL, for the dark neutrino models under study are shown in fig. 6.11. The results are presented for exposures of 1.11×10^{21} and 5.55×10^{21} POT (FHC), corresponding to approximately one and five years of data taking at SAND. The current reference 95% CL excluded region obtained by MicroBooNE [87] for the same models is shaded. The $M_N < M_{Z'} = 30$ MeV region for the present choice of dark photon mass cannot be currently simulated in GENIE. An improvement over the current limits is found across the tested mass range, with sensitivities to couplings down to $\sim 10^{-11}$ found in the 50 to 200 MeV mass range. The slight reduction in sensitivity at the lower end of the mass range ($M_N \lesssim 60$ MeV), may be related to the reduced signal efficiency of the topological selection cut in sect. 5.3.1, whereas the reduction at $M_N \gtrsim 200$ MeV is driven by the lower event rate (see sect. 5.1.1).

While the present result is limited to samples with a span in M_N due to constraints in the available computing time, the pipeline may be extended in the future to additional models, covering a range of $M_{Z'}$ values.




Figure 6.11: Discovery potential at $Z = 1.645$ (corresponding to a one-sided 95% CL) in SAND for the $35 \text{ MeV} \leq M_N \leq 1 \text{ GeV}$, $M_{Z'} = 30 \text{ MeV}$ dark neutrino models under study, at exposures of 1.11×10^{21} and 5.55×10^{21} POT (FHC), corresponding to one and five years of data taking, respectively. The corresponding MicroBooNE 95% CL excluded region is highlighted in orange [87]. The off-shell Z' decay region of the parameter space ($M_N \leq M_{Z'} = 30 \text{ MeV}$) cannot currently be simulated in GENIE and is highlighted in gray.

6.4 Future prospects

To provide an initial estimate of the SAND sensitivity to dark neutrino models, the simulation and analysis pipeline developed in this work was adapted to the current state of the SAND software framework, which is still under active development. This required a number of simplifications and approximations, namely the use of Monte Carlo truth at multiple stages of the chain and the application of an un-optimised reconstruction algorithm that only made use of the tracker information. Several aspects of the analysis methodology may thus be improved in the future to provide a more realistic estimate of the discovery potential:

- The development of event reconstruction and vertexing algorithms for SAND would eliminate the need for Monte Carlo information in the initial topological cut, allowing a more realistic selection of signal and background. Particle identi-

fication may also be performed algorithmically, instead of relying on the Monte Carlo truth.

- The simulation of the DUNE beam spill structure would allow to introduce the effect of pile-up in the selection of signatures.
- A detailed parametrisation of the detector acceptance for each signal model could be performed, including event generation in the external ECAL and magnet components.
- An optimised track reconstruction algorithm for e^+e^- would improve the separation of signal and background kinematic variables. Either a dedicated algorithm for e^+e^- reconstruction, or the **sandreco** Kalman filter implementation may be employed, introducing correction factors for Bremsstrahlung. Said reconstruction algorithm should furthermore take digitised signals as input, instead of relying on the smearing of the edep-sim hits.
- Information from the ECAL, using the **sandreco** clustering algorithm, may be used to further constrain the energy of electrons and positrons exiting the tracker. Furthermore, the detection of Bremsstrahlung photons exiting the tracker would be possible, enabling corrections to the reconstructed e^+e^- tracks.

While the present work has been carried out with an ad-hoc pipeline, future analyses would benefit from a further integration with the SAND reconstruction software and the wider DUNE Near Detector framework, allowing an extension of the current results to the other ND detectors.

In addition to the improvements to the pipeline discussed above, a comprehensive dark neutrino analysis in SAND may benefit from several additions to the DarkNeutrino GENIE module:

- The inclusion of scattering on ^1H , which, as shown in sect. 5.1.1, has a contribution smaller by $\mathcal{O}(10)$ compared to the other target masses in the inner volume.
- The modelling of dark neutrino decays to off-shell Z' , which, owing to their long decay lengths ($c\tau \gg 10\text{ cm}$), could produce e^+e^- pairs in the tracker from interactions occurring in the denser outer material (ECAL, magnet yoke and cavern), thereby improving the signal-to-background ratio.
- The modelling of $Z' \rightarrow \mu^+\mu^-$ decays, which would feature a clearer signature compared to e^+e^- , as well as being optimal for reconstruction in the tracker.

An improved analysis chain for the study of dark neutrinos may finally be adapted to a wider range of rare-event and BSM searches with similar di-lepton signatures, including the broader spectrum of HNL models as well as Neutrino tridents [106].

Conclusions

In addition to its wide-ranging program of neutrino oscillation physics, the next-generation DUNE experiment will leverage its high-intensity neutrino beam, long baseline and advanced detectors at the Near Detector and Far Detector complex to conduct an ample set of precision measurements and rare event searches, from the detection of supernova neutrinos, to highly accurate neutrino cross-section measurements and searches for Beyond the Standard Model (BSM) physics. The ND complex detectors will be especially competitive in the search for beam-produced BSM signatures, owing to the unprecedented neutrino flux of the LBNF beam and the complementary use of liquid argon and gas detector technologies, which will enable high resolutions in vertexing, momentum reconstruction and charge measurements.

Among the broad range of BSM physics possible at the DUNE-ND, searches for weakly-interacting long-lived particles such as Heavy Neutral Leptons (HNLs) and hidden sector portals have garnered a significant attention over the past years, as they could be produced in the beam and undergo in-flight decays to standard model particles within the fiducial volume of the ND detectors. A particularly interesting HNL scenario for searches at the near detectors of long-baseline experiments such as DUNE-ND, is that of Dark Neutrinos, in which HNLs are produced in the scattering of an active flavour state on a nuclear target, and can decay to standard di-lepton pairs, by the exchange of a dark $U(1)'$ gauge boson (or *dark photon* Z').

In this work, a first simulation and analysis pipeline for the search of dark neutrino events in one of the DUNE-ND detectors, SAND, was developed with the aim of determining their discovery potential. The optimal momentum and charge reconstruction performances offered by the SAND gas-detector inner tracker and magnetic field were leveraged to allow the identification of signatures of dark neutrino decays to e^+e^- pairs against the background from standard neutrino interactions. A dark neutrino event generator compatible with the standard DUNE-ND simulation framework was corrected and integrated in the analysis chain, enabling the full use of the SAND analysis software and, potentially, the extension of the analysis to the other ND detectors.

A selection criterion for the signatures of dark neutrino decay e^+e^- pairs in the SAND tracker was defined, and the resulting background from standard interactions (mainly NC-interactions with π^0 production) was assessed. The selection was performed using the event Monte Carlo information, with a more realistic treatment requiring vertexing and event reconstruction algorithms for SAND. A preliminary track fitting algorithm for the signature electrons and positrons was implemented in order to reconstruct the kinematics of the signature pairs. The algorithm, based on an helicoidal track model, was prone to underestimating e^+/e^- momenta, as it did

not include corrections for radiative energy losses. A resolution of $\sim 9\%$ was found on the single electron bending plane momentum, and the resulting smearing in the kinematics rendered background rejection through simple selection cuts ineffective. In order to overcome the current limitations in the e^+e^- tracking, a Boosted Decision Tree classifier (BDT) was applied to the reconstructed pair kinematic variables to perform the final signal–background rejection step. A pipeline for the training and testing of classifiers with the standard TMVA library was developed, and a set of BDT configurations was trained on the range of dark neutrino model samples considered in the analysis. The discovery significance reached on the test-set was used as a metric for determining the optimal BDT working point and configuration.

The full analysis pipeline thus defined was applied to interaction samples from signal models spanning a range of dark neutrino masses M_N . The significances obtained for the samples were used to determine the discovery potential at each tested M_N as a function of the dark neutrino coupling $|U_{\mu 4}|^2$. The current experimental exclusion limits are improved across the entire explored M_N range after one year of exposure (equivalent to 1.11×10^{21} Protons on Target), with the exclusion of couplings down to $|U_{\mu 4}|^2 \simeq 10^{-11}$ possible at a 95% C.L at $M_N \simeq 100$ MeV.

At present, the framework for response simulation and reconstruction in SAND is still undergoing active development, so that the analysis pipeline developed hitherto made use of Monte Carlo information at several points, namely in the selection of event signatures and in the track reconstruction algorithm. Multiple improvements could be made in the future on the selection of signatures, potentially matching the tracker signature with information from GRAIN and ECAL. A reconstruction algorithm based on digitised hits and including corrections for radiative losses in e^+/e^- would also be crucial to improve the signal–background separation by way of the pair kinematics. The introduction of additional processes in the generator, such as decays to off-shell Z' and $Z' \rightarrow \mu^+\mu^-$, would allow the scope of the analysis to be expanded.

It may also be noted that the analysis pipeline, while developed to study dark neutrino models, may be extended to other comparable rare-event and BSM searches, an example being Neutrino Tridents, which share the same di-lepton signature with a missing vertex as dark neutrinos.

Appendix A

Tracker response simulations with Garfield++

A key aspect of the **sandreco** simulation software (see sect. 4.1.4) is the implementation of *digitisation* algorithms, which convert the energy deposited in each detector volume into ADC and TDC values. A full modelling of the SAND tracker performance requires a realistic simulation of ionisation signals and readout electronics. This appendix introduces the current *fast digitisation* algorithm for the SAND tracker and describes the ongoing development of a realistic digitisation based on *lookup tables* produced with the Garfield++ toolkit [144]. So far, the development has focused on the drift chamber option in order to provide feedback for the prototype R&D, and this design is assumed throughout the appendix.

A.0.1 Fast digitisation

The present version of **sandreco** features a fast-digitisation algorithm, which determines the closest wire to each EdepSim segment (see sect. 4.1) and computes its TDC value as:

$$t \equiv T_{hit} + T_{drift} + T_{wire}, \quad (\text{A.0.1})$$

where T_{hit} is the time at which the hit is generated, T_{drift} is the drift time from the point of closest approach to the wire computed assuming a constant drift velocity of 50 $\mu\text{m}/\text{ns}$, and T_{wire} is the signal propagation time in the wire, assuming a velocity of 200 mm/ns . The measured TDC for a given wire is the shortest time among the hits associated to that wire. The fast digitisation process is illustrated in fig. A.1.

A.0.2 Response simulation with Garfield++

Being the standard toolkit for the modelling and response simulation of gas detectors, Garfield++ was used to guide the design of the drift-chamber tracker and to provide a reference during the data analysis of the small and medium scale prototypes. A three-dimensional geometry of the drift-chamber planes (see sect. 3.1.4) was initially implemented, using the Nearly Exact Boundary Element (neBEM) field solver available in Garfield++. The model was used to optimise the cell configuration and gas mixture to be used in the prototypes by evaluating the electron transport properties,




Figure A.1: Scheme of the fast digitisation algorithm for the tracker. The drift time is computed from the point of closest approach of the hit to the wire, assuming $v_{drift} = 50 \mu\text{m/ns}$. The wire TDC corresponds to the shortest time among the associated hits.

such as the drift velocity v_{drift} , the electron drift time t_{drift} and the gas gain. The optimisation targeted a gas gain of 10^4 – 10^5 and an approximately uniform drift velocity of $v_{drift} \simeq 50 \mu\text{m/ns}$ along the wire plane. These parameters were evaluated throughout the cell by sampling the geometry with a sufficiently high resolution (typically $100 \mu\text{m}$), yielding the *lookup tables* (LUTs) in fig. A.2.




Figure A.2: Lookup tables obtained for a $1.4 \times 1.2 \text{ cm}^2$ cell with $V_{field} = -1.7 \text{ kV}$ and $V_{mylar} = -1.6 \text{ kV}$: the drift velocity for the chosen Ar/CO₂ (85%/15%) gas mixture is pictured in (a), and the corresponding electron drift time is shown in (b). The region intersecting the field wires is not sampled.

Based on these results, a $1.4 \times 1.2 \text{ cm}^2$ cell using an Ar/CO₂ (85%/15%) gas mixture at near-atmospheric pressure was chosen for the medium-scale prototype. Electrode configurations with a grounded sense wire and both field wires and mylar

planes biased in the -1.5 kV to -1.7 kV range were found to meet the requirements and were therefore adopted during the prototype beam test [122].

The model was also used to study the boundary effects produced by the gas-gaps between the mylar plane electrodes and their effect on the drift field at different wire orientations. Wire angles of 0° , $\pm 5^\circ$ and $\pm 10^\circ$ with respect to the direction perpendicular to the gas gaps were tested, and the effects on the cell field were found to be negligible.

A.0.3 Realistic digitisation with lookup tables

Garfield++ features a realistic modelling of the ionisation from charged-particle tracks in detector geometries as well as the simulation of the resulting induced-current signals on sensor electrodes and of the detector readout response. However, fully integrating the toolkit in the SAND simulation framework is not practical due to the high computational cost of full track simulations and the introduction of additional, undesired dependencies.

An approach based on the lookup tables discussed above is being pursued instead: the induced-current signals of single-electrons are sampled throughout the cell, so that a trajectory ionisation signal can be constructed by summing the single waveforms. Given an EdepSim hit-segment within the cell:

1. The number of ionisation electrons corresponding to the hit energy deposit E_{hit} is computed as $N_e = E_{hit}/W_{gas}$, W_{gas} being the average energy to create an electron-ion pair.
2. The N_e electrons are distributed uniformly between the initial and final points of the hit-segment x_{hit}^{start} and x_{hit}^{stop} , and the corresponding induced current signals are extracted from the LUT and summed.

For a given event, the induced signals of all hits associated to a given cell are summed, including an offset T_{wire}^{hit} due to the propagation time in the wire. The combined signal is then convoluted with the response function of the readout, yielding the output voltage waveform, and allowing to compute the ADC and TDC values according to the readout specifications. The simulation code allows a flexible implementation and selection of the response of different readout systems. A scheme of the digitisation workflow is shown in fig. A.3.

Since the digitisation is performed at a later stage compared to the event generation and particle propagation with EdepSim (see sect. 4.1), the same events can be digitised with multiple LUTs and readouts, allowing to quickly compare the effect of different gas mixtures, voltage configurations and readout systems.

Simulation of a small-scale drift chamber prototype

An initial implementation of the LUT-based simulation was developed for a toy model of a small-scale drift chamber prototype [122], consisting of three drift planes with an active area of $30 \times 30\text{ cm}$ and $2 \times 1\text{ cm}^2$ cells, and using with an Ar/CO₂ (85%/15%) gas mixture at near-atmospheric pressure. The aim of the prototype, which operated with




Figure A.3: Scheme of the full digitisation algorithm for the drift chamber tracker, showing the combined induced current signal and the output voltage waveform from a VTX charge pre-amplifier [172], used in data-taking of the medium scale prototype.

cosmic rays, was to validate the chamber layout concept and provide a cross-check for the simulations.

Samples of cosmic ray muon tracks were simulated in the Garfield++ three-dimensional cell model discussed above, and the position and size of the ionisation clusters were saved. The corresponding induced signals were then constructed by sampling the single-electron currents from cell LUT, accounting for the cluster multiplicity. The resulting signals were compared with those simulated directly in Garfield++ and were found to be in good agreement. The induced current signals were convoluted to the response function of the VTX charge pre-amplifier [172] used for the prototype readout.

The simulated waveforms were processed consistently the prototype signals, except for an additional coherent-noise filtering step applied only to the latter. A peak-finding algorithm was used to obtain T_{drift} as the 10% constant fraction time of the waveform peak, allowing the retrieval of the distance-drift time curve shown in fig. A.4 for a cell with $V_{sense} = 1.6 \text{ kV}$, $V_{field} = -100 \text{ V}$ and $V_{mylar} = 0 \text{ V}$. The induced charge was estimated by integrating over a 200 ns window, starting from the rising edge of the peak and applying a scaling factor of $1/50 \Omega$ corresponding to the resistance of the digitiser channel termination. Figure A.5 shows a comparison between the estimated charge distributions of the simulated and prototype waveforms: the observed discrepancy was traced back to an underestimation of the prototype integrals due to the application of the coherent-noise filter

In addition to the comparison with the prototype data, it was possible to use the toy model to validate the assumption of a uniform distribution of electrons for the




Figure A.4: Distance-time curve of a $2 \times 1 \text{ cm}^2$ cell with $V_{sense} = 1.6 \text{ kV}$, $V_{field} = -100 \text{ V}$ and $V_{mylar} = 0 \text{ V}$ obtained from a sample of $\sim 4 \times 10^4$ events in the small-scale prototype model. The drift time is estimated as the peak $T_{10\%}$.




Figure A.5: Comparison between the estimated charge distributions of the simulated and small-scale prototype waveforms for a cell with $V_{sense} = 1.6 \text{ kV}$, $V_{field} = -100 \text{ V}$ and $V_{mylar} = 0 \text{ V}$. A cut on the wire-distance at $d_W < 1.12 \text{ cm}$, corresponding to the maximum distance within the $2 \times 1 \text{ cm}^2$ cell, is applied.

realistic digitisation in **sandreco**. For a sample of 5×10^4 muon tracks generated in Garfield++, equivalent waveforms were constructed by uniformly distributing the total number of electrons in the ionisation clusters along the particle track. A $1.4 \times 1.2 \text{ cm}^2$ cell was simulated, with $V_{sense} = 0 \text{ V}$, $V_{field} = -1.8 \text{ kV}$ and $V_{mylar} = -1.6 \text{ kV}$.

The resulting $T_{10\%}$ and total waveform integrals were compared with those obtained from the original Garfield++ waveforms, as shown in fig. A.6. Both quantities were found to be in good agreement across the full range, supporting the validity of the uniform-distribution approximation.




Figure A.6: Comparison between the $T_{10\%}$ (a) and total waveform integral (b) of cosmic ray muon waveforms simulated in Garfield++ and constructed with a uniform electron distribution along the particle track. A $1.4 \times 1.2 \text{ cm}^2$ cell with $V_{sense} = 0 \text{ V}$, $V_{field} = -1.8 \text{ kV}$ and $V_{mylar} = -1.6 \text{ kV}$ was used.

Response simulation for the medium-scale drift-chamber prototype


The full digitisation algorithm, starting from hit-segments, has been implemented within an already existing Geant4 simulation of the medium-scale drift chamber prototype introduced in sect. 3.1.4. The Geant4 simulation features a detailed detector geometry and allows the generation of samples of 200 GeV muons according to the conditions of beam test at SPS [122].

In the simulation, the Geant4 energy deposits are converted to a format equivalent to the EdepSim hit-segments, and grouped by drift cell according to the prototype channel layout, allowing the application of the workflow described in sect. A.0.3. A two-dimensional analytical geometry of the drift planes is implemented in Garfield++, enabling a fast generation of the lookup tables. LUTs are saved in HDF5 format to ensure code compatibility with the future updates to the SAND simulation framework. The selected LUT is specified at the digitisation stage, allowing the same events to be processed with different gas mixtures and voltage configurations.


The readout simulation is designed to be expandable, allowing the implementation of multiple readout systems and their rapid selection at runtime. In particular, the response of the TIGER ASIC [173], used during the beam test, is implemented. The TDC value is computed by convolving the induced current with the fast response function of the TIGER *time-branch* and applying a fixed threshold. The slower

response function of the TIGER *energy-branch* is used to compute estimated charge ADC from the peaking-time of the corresponding output waveform. The effect of the TIGER charge saturation at ~ 50 fC is included.

The distance-time curve and ADC distributions of a cell with $V_{sense} = 0$ V, $V_{field} = -1.3$ kV and $V_{mylar} = -1.3$ kV, obtained from a sample of 5×10^4 muons, are shown in fig. A.7. The TIGER saturation value of ~ 50 fC is visible in the ADC distribution.



(a)



(b)

Figure A.7: Distance-time curve (a) and ADC distribution (b) of a medium-scale prototype cell with $V_{sense} = 0$ V, $V_{field} = -1.3$ kV and $V_{mylar} = -1.3$ kV, obtained from a sample of 5×10^4 muons with the TIGER ASIC response simulation. The ~ 50 fC saturation charge is visible in the ADC distribution.

The response simulation that is being developed for the medium-scale prototype model will be validated with beam test data, and will serve as a template for the full digitisation implementation in the in the future updates of the **sandreco** simulation

code. While the current implementations were focused on the drift chamber design, in order to provide feedback to the prototype R&D, the production of Garfield++ lookup tables and the digitisation algorithm can be analogously applied to the straw tube design.

Appendix B

TMVA workflow

This appendix will introduce the basic workflow for the training, testing and application of TMVA multivariate classifier models, as used for the signal-background discrimination in Chap. 6.

B.0.1 Training and testing pipeline

The training and testing pipeline in TMVA is managed through the `Factory` object, which handles the loading of the training and test sets, books the analysis methods and subsequently calls for training and evaluation. This structure allows multiple algorithms to be trained and benchmarked on the same input samples. The `Factory` was instantiated as follows¹.

```
1 factory = ROOT.TMVA.Factory("<JobName>", outputFile, "<factoryConfig>")
```

Here, `<JobName>` specifies the name assigned to the output weight files, and `outputFile` is a pointer to a writable ROOT `TFile` that stores the training results. The third argument allows to specify the analysis type: in this study, `"AnalysisType:Classification"` was used to select a classification analysis. The handling of input datasets was performed through a `DataLoader` object, created alongside the corresponding `Factory`. The variables used in model training were first set:

```
1 # Define the DataLoader and set the output ROOT directory
2 loader = ROOT.TMVA.DataLoader("datasetName")
3 # Set the input variables from a list
4 for variable_name in training_variables_list:
5     loader.AddVariable(variable_name, "<variableType>")
```

The training and test ROOT `TTrees` were then specified for the signal and background and an event weight for each tree was applied, e.g.:

```
1 # Assign the training dataset for the signal
2 loader.AddSignalTree(signalTrainTree, weight, ROOT.TMVA.Types.kTraining)
```

¹The code examples use ROOT's Python language interface for TMVA, which provides access to TMVA's methods and workflows in Python and allows integration with standard Machine Learning libraries such as `scikit-learn`, `Keras`, `PyTorch` and `XGBoost` [174–177].

Once the samples were set, multiple training methods were booked through the `Factory` object, passing the model type as the first argument (e.g. `ROOT.TMVA.Types.kBDT` for BDT models) and a string with the model configuration options as the third:

```
1 # Book the TMVA model, setting its type, name and configuration
2 factory.BookMethod(loader, modelType, "<modelName>", "<modelConfiguration>")
```

As a last step, the training and testing of all booked methods was invoked, together with a preliminary assessment of the model performance on the test set:

```
1 factory.TrainAllMethods() # train the booked methods
2 factory.TestAllMethods() # apply the trained methods on the test set
3 factory.EvaluateAllMethods() # evaluate the performance
```

In the present analysis, the list of training variables and model configuration parameters was defined in an external configuration file, enabling a quick reconfiguration and consistent tracking of the different training experiments. This workflow allowed to train multiple algorithms on the same signal and background datasets.

B.0.2 Model output and application

The final output of the TMVA model training is an `xml weight file` containing trained model parameters, which is applied to the target datasets. The weight file is loaded in a `Reader` object, which, in a pipeline based on `RDataFrame`, can be used to apply the model on each sample event and evaluate the corresponding score:

```
1 # Load the model to Reader
2 tmva_model = ROOT.TMVA.Experimental.RReader("<weight_file_path>")
3 # Get the list of training variables
4 vars = list(tmva_model.GetVariableNames())
5 # Apply the model to each entry based on the training variables
6 target_dataframe = target_dataframe.Define(
7 "modelScore",
8 ROOT.TMVA.Experimental.Compute[len(vars), "<variableType>"](tmva_model),
9 vars)
```

A selection cut can thus be applied on the score value to select the signal events:

```
1 target_dataframe.Filter("modelScore[0] > <selectionCut>")
```

Since `Compute` returns a vector of scores (to support multiple output models such as multi-class classification), this pipeline selects the first element. The procedure through which the cut score value was optimised is discussed in sect. [6.3.1](#).

Bibliography

- [1] Carlo Giunti and Chung W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. 2007. ISBN: 978-0-19-850871-7. DOI: [10.1093/acprof:oso/9780198508717.001.0001](#).
- [2] Samoil Bilenky. *Introduction to the Physics of Massive and Mixed Neutrinos*. Vol. 947. Springer, 2018. ISBN: 978-3-319-74801-6, 978-3-319-74802-3. DOI: [10.1007/978-3-319-74802-3](#).
- [3] R. L. Workman et al. “Review of Particle Physics”. In: *PTEP* 2022 (2022), p. 083C01. DOI: [10.1093/ptep/ptac097](#).
- [4] Mark Thomson. *Modern particle physics*. New York: Cambridge University Press, Oct. 2013. ISBN: 978-1-107-03426-6, 978-1-139-52536-7. DOI: [10.1017/CBO9781139525367](#).
- [5] C. Hagedorn et al. “CP Violation in the Lepton Sector and Implications for Leptogenesis”. In: *Int. J. Mod. Phys. A* 33.05n06 (2018), p. 1842006. DOI: [10.1142/S0217751X1842006X](#). arXiv: [1711.02866 \[hep-ph\]](#).
- [6] L. Wolfenstein. “Neutrino Oscillations in Matter”. In: *Phys. Rev. D* 17 (1978), pp. 2369–2374. DOI: [10.1103/PhysRevD.17.2369](#).
- [7] B. T. Cleveland et al. “Measurement of the solar electron neutrino flux with the Homestake chlorine detector”. In: *Astrophys. J.* 496 (1998), pp. 505–526. DOI: [10.1086/305343](#).
- [8] 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](#).
- [9] M. Altmann et al. “Complete results for five years of GNO solar neutrino observations”. In: *Phys. Lett. B* 616 (2005), pp. 174–190. DOI: [10.1016/j.physletb.2005.04.068](#). arXiv: [hep-ex/0504037](#).
- [10] J. N. Abdurashitov et al. “Solar neutrino flux measurements by the Soviet-American Gallium Experiment (SAGE) for half the 22 year solar cycle”. In: *J. Exp. Theor. Phys.* 95 (2002), pp. 181–193. DOI: [10.1134/1.1506424](#). arXiv: [astro-ph/0204245](#).
- [11] Y. Fukuda et al. “Solar neutrino data covering solar cycle 22”. In: *Phys. Rev. Lett.* 77 (1996), pp. 1683–1686. DOI: [10.1103/PhysRevLett.77.1683](#).
- [12] J. Hosaka et al. “Solar neutrino measurements in super-Kamiokande-I”. In: *Phys. Rev. D* 73 (2006), p. 112001. DOI: [10.1103/PhysRevD.73.112001](#). arXiv: [hep-ex/0508053](#).

- [13] K. Abe et al. “Solar neutrino measurements using the full data period of Super-Kamiokande-IV”. In: *Phys. Rev. D* 109.9 (2024), p. 092001. DOI: [10.1103/PhysRevD.109.092001](#). arXiv: [2312.12907 \[hep-ex\]](#).
- [14] J. Boger et al. “The Sudbury neutrino observatory”. In: *Nucl. Instrum. Meth. A* 449 (2000), pp. 172–207. DOI: [10.1016/S0168-9002\(99\)01469-2](#). arXiv: [nucl-ex/9910016](#).
- [15] B. Aharmim et al. “Measurement of the ν_e and Total ^8B Solar Neutrino Fluxes with the Sudbury Neutrino Observatory Phase-III Data Set”. In: *Phys. Rev. C* 87.1 (2013), p. 015502. DOI: [10.1103/PhysRevC.87.015502](#). arXiv: [1107.2901 \[nucl-ex\]](#).
- [16] 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 \[hep-ex\]](#).
- [17] 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 \[nucl-ex\]](#).
- [18] G. Bellini et al. “Precision measurement of the ^7Be solar neutrino interaction rate in Borexino”. In: *Phys. Rev. Lett.* 107 (2011), p. 141302. DOI: [10.1103/PhysRevLett.107.141302](#). arXiv: [1104.1816 \[hep-ex\]](#).
- [19] B. Aharmim et al. “Electron energy spectra, fluxes, and day-night asymmetries of B-8 solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory”. In: *Phys. Rev. C* 72 (2005), p. 055502. DOI: [10.1103/PhysRevC.72.055502](#). arXiv: [nucl-ex/0502021](#).
- [20] S. Gariazzo et al. “Neutrino masses and their ordering: Global Data, Priors and Models”. In: *JCAP* 03 (2018), p. 011. DOI: [10.1088/1475-7516/2018/03/011](#). arXiv: [1801.04946 \[hep-ph\]](#).
- [21] Michele Maltoni and Alexei Yu. Smirnov. “Solar neutrinos and neutrino physics”. In: *Eur. Phys. J. A* 52.4 (2016), p. 87. DOI: [10.1140/epja/i2016-16087-0](#). arXiv: [1507.05287 \[hep-ph\]](#).
- [22] K. S. Hirata et al. “Experimental Study of the Atmospheric Neutrino Flux”. In: *Phys. Lett. B* 205 (1988). Ed. by J. Tran Thanh Van, p. 416. DOI: [10.1016/0370-2693\(88\)91690-5](#).
- [23] D. Casper et al. “Measurement of atmospheric neutrino composition with IMB-3”. In: *Phys. Rev. Lett.* 66 (1991), pp. 2561–2564. DOI: [10.1103/PhysRevLett.66.2561](#).
- [24] 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](#).
- [25] W. W. M. Allison et al. “Measurement of the L/E distributions of atmospheric neutrinos in Soudan 2 and their interpretation as neutrino oscillations”. In: *Phys. Rev. D* 68 (2003), p. 113004. DOI: [10.1103/PhysRevD.68.113004](#). arXiv: [hep-ex/0307069](#).

- [26] M. Ambrosio et al. “Matter effects in upward going muons and sterile neutrino oscillations”. In: *Phys. Lett. B* 517 (2001), pp. 59–66. DOI: [10.1016/S0370-2693\(01\)00992-3](https://doi.org/10.1016/S0370-2693(01)00992-3). arXiv: [hep-ex/0106049](https://arxiv.org/abs/hep-ex/0106049).
- [27] S. H. Ahn et al. “Detection of accelerator produced neutrinos at a distance of 250 km”. In: *Phys. Lett. B* 511 (2001), pp. 178–184. DOI: [10.1016/S0370-2693\(01\)00647-5](https://doi.org/10.1016/S0370-2693(01)00647-5). arXiv: [hep-ex/0103001](https://arxiv.org/abs/hep-ex/0103001).
- [28] 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](https://doi.org/10.1103/PhysRevLett.120.211801). arXiv: [1804.04912](https://arxiv.org/abs/1804.04912) [[hep-ex](#)].
- [29] 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](https://doi.org/10.1103/PhysRevD.99.032007). arXiv: [1901.05366](https://arxiv.org/abs/1901.05366) [[hep-ex](#)].
- [30] Y. Ashie et al. “A Measurement of atmospheric neutrino oscillation parameters by SUPER-KAMIOKANDE I”. In: *Phys. Rev. D* 71 (2005), p. 112005. DOI: [10.1103/PhysRevD.71.112005](https://doi.org/10.1103/PhysRevD.71.112005). arXiv: [hep-ex/0501064](https://arxiv.org/abs/hep-ex/0501064).
- [31] F. P. An et al. “Observation of electron-antineutrino disappearance at Daya Bay”. In: *Phys. Rev. Lett.* 108 (2012), p. 171803. DOI: [10.1103/PhysRevLett.108.171803](https://doi.org/10.1103/PhysRevLett.108.171803). arXiv: [1203.1669](https://arxiv.org/abs/1203.1669) [[hep-ex](#)].
- [32] J. K. Ahn et al. “Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment”. In: *Phys. Rev. Lett.* 108 (2012), p. 191802. DOI: [10.1103/PhysRevLett.108.191802](https://doi.org/10.1103/PhysRevLett.108.191802). arXiv: [1204.0626](https://arxiv.org/abs/1204.0626) [[hep-ex](#)].
- [33] Y. Abe et al. “Indication of Reactor $\bar{\nu}_e$ Disappearance in the Double Chooz Experiment”. In: *Phys. Rev. Lett.* 108 (2012), p. 131801. DOI: [10.1103/PhysRevLett.108.131801](https://doi.org/10.1103/PhysRevLett.108.131801). arXiv: [1112.6353](https://arxiv.org/abs/1112.6353) [[hep-ex](#)].
- [34] K. Hiraide et al. “Resolving θ_{23} degeneracy by accelerator and reactor neutrino oscillation experiments”. In: *Phys. Rev. D* 73 (2006), p. 093008. DOI: [10.1103/PhysRevD.73.093008](https://doi.org/10.1103/PhysRevD.73.093008). arXiv: [hep-ph/0601258](https://arxiv.org/abs/hep-ph/0601258).
- [35] Ivan Esteban et al. “NuFit-6.0: updated global analysis of three-flavor neutrino oscillations”. In: *JHEP* 12 (2024), p. 216. DOI: [10.1007/JHEP12\(2024\)216](https://doi.org/10.1007/JHEP12(2024)216). arXiv: [2410.05380](https://arxiv.org/abs/2410.05380) [[hep-ph](#)].
- [36] Lino Miramonti. “Neutrino Physics and Astrophysics with the JUNO Detector”. In: *Universe* 4.11 (2018). ISSN: 2218-1997. URL: <https://www.mdpi.com/2218-1997/4/11/126>.
- [37] Fengpeng An et al. “Neutrino Physics with JUNO”. In: *J. Phys. G* 43.3 (2016), p. 030401. DOI: [10.1088/0954-3889/43/3/030401](https://doi.org/10.1088/0954-3889/43/3/030401). arXiv: [1507.05613](https://arxiv.org/abs/1507.05613) [[physics.ins-det](#)].
- [38] Angel Abusleme et al. “JUNO physics and detector”. In: *Prog. Part. Nucl. Phys.* 123 (2022), p. 103927. DOI: [10.1016/j.ppnp.2021.103927](https://doi.org/10.1016/j.ppnp.2021.103927). arXiv: [2104.02565](https://arxiv.org/abs/2104.02565) [[hep-ex](#)].
- [39] 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](https://doi.org/10.1103/PhysRevD.106.032004). arXiv: [2108.08219](https://arxiv.org/abs/2108.08219) [[hep-ex](#)].

- [40] Babak 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](https://doi.org/10.1088/1748-0221/15/08/T08008). arXiv: [2002.02967](https://arxiv.org/abs/2002.02967) [[physics.ins-det](#)].
- [41] K. Abe et al. “Hyper-Kamiokande Design Report”. In: (May 2018). arXiv: [1805.04163](https://arxiv.org/abs/1805.04163) [[physics.ins-det](#)].
- [42] M. G. Aartsen et al. “Letter of Intent: The Precision IceCube Next Generation Upgrade (PINGU)”. In: (Jan. 2014). arXiv: [1401.2046](https://arxiv.org/abs/1401.2046) [[physics.ins-det](#)].
- [43] 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](https://doi.org/10.1088/0954-3899/43/8/084001). arXiv: [1601.07459](https://arxiv.org/abs/1601.07459) [[astro-ph.IM](#)].
- [44] Angel Abusleme et al. “First measurement of reactor neutrino oscillations at JUNO”. In: (Nov. 2025). arXiv: [2511.14593](https://arxiv.org/abs/2511.14593) [[hep-ex](#)].
- [45] X. Qian and P. Vogel. “Neutrino Mass Hierarchy”. In: *Prog. Part. Nucl. Phys.* 83 (2015), pp. 1–30. DOI: [10.1016/j.ppnp.2015.05.002](https://doi.org/10.1016/j.ppnp.2015.05.002). arXiv: [1505.01891](https://arxiv.org/abs/1505.01891) [[hep-ex](#)].
- [46] Tomasz Wachala. “Recent Results on the CP Violation Search in the Accelerator Neutrino Oscillations”. In: *Acta Phys. Polon. B* 48 (2017), p. 1969. DOI: [10.5506/APhysPolB.48.1969](https://doi.org/10.5506/APhysPolB.48.1969).
- [47] William Marciano and Zohreh Parsa. “Intense neutrino beams and leptonic CP violation”. In: *Nucl. Phys. B Proc. Suppl.* 221 (2011). Ed. by Geoffrey Mills et al., pp. 166–172. DOI: [10.1016/j.nuclphysbps.2011.03.114](https://doi.org/10.1016/j.nuclphysbps.2011.03.114). arXiv: [hep-ph/0610258](https://arxiv.org/abs/hep-ph/0610258).
- [48] Yago Herrera and Aldo Serenelli. *Standard Solar Models B23 / SF-III*. Version v1.2. Zenodo, Nov. 2023. DOI: [10.5281/zenodo.10822316](https://doi.org/10.5281/zenodo.10822316). URL: <https://doi.org/10.5281/zenodo.10822316>.
- [49] Asli M. Abdullahi et al. “The present and future status of heavy neutral leptons”. In: *J. Phys. G* 50.2 (2023), p. 020501. DOI: [10.1088/1361-6471/ac98f9](https://doi.org/10.1088/1361-6471/ac98f9). arXiv: [2203.08039](https://arxiv.org/abs/2203.08039) [[hep-ph](#)].
- [50] Peter Ballett, Tommaso Boschi and Silvia Pascoli. “Heavy Neutral Leptons from low-scale seesaws at the DUNE Near Detector”. In: *JHEP* 03 (2020), p. 111. DOI: [10.1007/JHEP03\(2020\)111](https://doi.org/10.1007/JHEP03(2020)111). arXiv: [1905.00284](https://arxiv.org/abs/1905.00284) [[hep-ph](#)].
- [51] Takehiko Asaka, Steve Blanchet and Mikhail Shaposhnikov. “The nuMSM, dark matter and neutrino masses”. In: *Phys. Lett. B* 631 (2005), pp. 151–156. DOI: [10.1016/j.physletb.2005.09.070](https://doi.org/10.1016/j.physletb.2005.09.070). arXiv: [hep-ph/0503065](https://arxiv.org/abs/hep-ph/0503065).
- [52] Apostolos Pilaftsis. “CP violation and baryogenesis due to heavy Majorana neutrinos”. In: *Phys. Rev. D* 56 (1997), pp. 5431–5451. DOI: [10.1103/PhysRevD.56.5431](https://doi.org/10.1103/PhysRevD.56.5431). arXiv: [hep-ph/9707235](https://arxiv.org/abs/hep-ph/9707235).
- [53] Peter Minkowski. “ $\rightarrow e$ at a rate of one out of 109 muon decays?” In: *Physics Letters B* 67.4 (1977), pp. 421–428. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(77\)90435-X](https://doi.org/10.1016/0370-2693(77)90435-X). URL: <https://www.sciencedirect.com/science/article/pii/037026937790435X>.

- [54] J. Schechter and J. W. F. Valle. “Neutrino masses in $SU(2) \otimes U(1)$ theories”. In: *Phys. Rev. D* 22 (9 Nov. 1980), pp. 2227–2235. DOI: [10.1103/PhysRevD.22.2227](https://doi.org/10.1103/PhysRevD.22.2227). URL: <https://link.aps.org/doi/10.1103/PhysRevD.22.2227>.
- [55] Robert Foot et al. “Seesaw Neutrino Masses Induced by a Triplet of Leptons”. In: *Z. Phys. C* 44 (1989), p. 441. DOI: [10.1007/BF01415558](https://doi.org/10.1007/BF01415558).
- [56] E. Cabrera, D. Cogollo and C.A. de S. Pires. “Naturally low scale type I seesaw mechanism and its viability in the 3-3-1 model with right-handed neutrinos”. In: *Nuclear Physics B* 996 (2023), p. 116372. ISSN: 0550-3213. DOI: <https://doi.org/10.1016/j.nuclphysb.2023.116372>. URL: <https://www.sciencedirect.com/science/article/pii/S0550321323003012>.
- [57] Enrique Fernández-Martínez et al. “Effective portals to heavy neutral leptons”. In: *JHEP* 09 (2023), p. 001. DOI: [10.1007/JHEP09\(2023\)001](https://doi.org/10.1007/JHEP09(2023)001). arXiv: [2304.06772](https://arxiv.org/abs/2304.06772) [hep-ph].
- [58] D. I. Britton et al. “Improved search for massive neutrinos in $\pi^+ \rightarrow e^+ \nu$ decay”. In: *Phys. Rev. D* 46 (1992), R885–R887. DOI: [10.1103/PhysRevD.46.R885](https://doi.org/10.1103/PhysRevD.46.R885).
- [59] A. Aguilar-Arevalo et al. “Improved search for heavy neutrinos in the decay $\pi \rightarrow e \nu$ ”. In: *Phys. Rev. D* 97.7 (2018), p. 072012. DOI: [10.1103/PhysRevD.97.072012](https://doi.org/10.1103/PhysRevD.97.072012). arXiv: [1712.03275](https://arxiv.org/abs/1712.03275) [hep-ex].
- [60] G. Bellini et al. “New limits on heavy sterile neutrino mixing in B8 decay obtained with the Borexino detector”. In: *Phys. Rev. D* 88.7 (2013), p. 072010. DOI: [10.1103/PhysRevD.88.072010](https://doi.org/10.1103/PhysRevD.88.072010). arXiv: [1311.5347](https://arxiv.org/abs/1311.5347) [hep-ex].
- [61] Eduardo Cortina Gil et al. “Search for heavy neutral lepton production in K^+ decays to positrons”. In: *Phys. Lett. B* 807 (2020), p. 135599. DOI: [10.1016/j.physletb.2020.135599](https://doi.org/10.1016/j.physletb.2020.135599). arXiv: [2005.09575](https://arxiv.org/abs/2005.09575) [hep-ex].
- [62] K. Abe et al. “Search for heavy neutrinos with the T2K near detector ND280”. In: *Phys. Rev. D* 100.5 (2019), p. 052006. DOI: [10.1103/PhysRevD.100.052006](https://doi.org/10.1103/PhysRevD.100.052006). arXiv: [1902.07598](https://arxiv.org/abs/1902.07598) [hep-ex].
- [63] Ryan Barouki, Giacomo Marocco and Subir Sarkar. “Blast from the past II: Constraints on heavy neutral leptons from the BEBC WA66 beam dump experiment”. In: *SciPost Phys.* 13 (2022), p. 118. DOI: [10.21468/SciPostPhys.13.5.118](https://doi.org/10.21468/SciPostPhys.13.5.118). arXiv: [2208.00416](https://arxiv.org/abs/2208.00416) [hep-ph].
- [64] P. Abreu et al. “Search for neutral heavy leptons produced in Z decays”. In: *Z. Phys. C* 74 (1997). [Erratum: *Z.Phys.C* 75, 580 (1997)], pp. 57–71. DOI: [10.1007/s002880050370](https://doi.org/10.1007/s002880050370).
- [65] Georges Aad et al. “Search for heavy neutral leptons in decays of W bosons produced in 13 TeV pp collisions using prompt and displaced signatures with the ATLAS detector”. In: *JHEP* 10 (2019), p. 265. DOI: [10.1007/JHEP10\(2019\)265](https://doi.org/10.1007/JHEP10(2019)265). arXiv: [1905.09787](https://arxiv.org/abs/1905.09787) [hep-ex].
- [66] Armen Tumasyan et al. “Search for long-lived heavy neutral leptons with displaced vertices in proton-proton collisions at $\sqrt{s}=13$ TeV”. In: *JHEP* 07 (2022), p. 081. DOI: [10.1007/JHEP07\(2022\)081](https://doi.org/10.1007/JHEP07(2022)081). arXiv: [2201.05578](https://arxiv.org/abs/2201.05578) [hep-ex].

- [67] M. Daum et al. “Search for admixtures of massive neutrinos in the decay $\pi^+ \rightarrow \mu^+ + \nu$ ”. In: *Phys. Rev. D* 36 (9 Nov. 1987), pp. 2624–2632. DOI: [10.1103/PhysRevD.36.2624](https://doi.org/10.1103/PhysRevD.36.2624). URL: <https://link.aps.org/doi/10.1103/PhysRevD.36.2624>.
- [68] A. Aguilar-Arevalo et al. “Search for heavy neutrinos in $\pi \rightarrow \mu \nu$ decay”. In: *Phys. Lett. B* 798 (2019), p. 134980. DOI: [10.1016/j.physletb.2019.134980](https://doi.org/10.1016/j.physletb.2019.134980). arXiv: [1904.03269](https://arxiv.org/abs/1904.03269) [hep-ex].
- [69] G. Bernardi et al. “Further limits on heavy neutrino couplings”. In: *Physics Letters B* 203.3 (1988), pp. 332–334. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(88\)90563-1](https://doi.org/10.1016/0370-2693(88)90563-1). URL: <https://www.sciencedirect.com/science/article/pii/0370269388905631>.
- [70] Kevin J. Kelly and Pedro A. N. Machado. “MicroBooNE experiment, NuMI absorber, and heavy neutral leptons”. In: *Physical Review D* 104.5 (Sept. 2021). ISSN: 2470-0029. DOI: [10.1103/physrevd.104.055015](https://doi.org/10.1103/physrevd.104.055015). URL: <http://dx.doi.org/10.1103/PhysRevD.104.055015>.
- [71] Carlos A. Argüelles, Nicolò Foppiani and Matheus Hostert. “Heavy neutral leptons below the kaon mass at hodoscopic neutrino detectors”. In: *Phys. Rev. D* 105.9 (2022), p. 095006. DOI: [10.1103/PhysRevD.105.095006](https://doi.org/10.1103/PhysRevD.105.095006). arXiv: [2109.03831](https://arxiv.org/abs/2109.03831) [hep-ph].
- [72] R. S. Hayano et al. “Heavy-Neutrino Search Using $K_{\mu 2}$ Decay”. In: *Phys. Rev. Lett.* 49 (18 Nov. 1982), pp. 1305–1309. DOI: [10.1103/PhysRevLett.49.1305](https://doi.org/10.1103/PhysRevLett.49.1305). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.49.1305>.
- [73] A. V. Artamonov et al. “Study of the decay $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ in the momentum region $140 < P_\pi < 199$ MeV/c”. In: *Phys. Rev. D* 79 (2009), p. 092004. DOI: [10.1103/PhysRevD.79.092004](https://doi.org/10.1103/PhysRevD.79.092004). arXiv: [0903.0030](https://arxiv.org/abs/0903.0030) [hep-ex].
- [74] Eduardo Cortina Gil et al. “Search for K^+ decays to a muon and invisible particles”. In: *Phys. Lett. B* 816 (2021), p. 136259. DOI: [10.1016/j.physletb.2021.136259](https://doi.org/10.1016/j.physletb.2021.136259). arXiv: [2101.12304](https://arxiv.org/abs/2101.12304) [hep-ex].
- [75] A.M. Cooper-Sarkar et al. “Search for heavy neutrino decays in the BEBC beam dump experiment”. In: *Physics Letters B* 160.1 (1985), pp. 207–211. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(85\)91493-5](https://doi.org/10.1016/0370-2693(85)91493-5). URL: <https://www.sciencedirect.com/science/article/pii/0370269385914935>.
- [76] A. Vaitaitis et al. “Search for neutral heavy leptons in a high-energy neutrino beam”. In: *Phys. Rev. Lett.* 83 (1999), pp. 4943–4946. DOI: [10.1103/PhysRevLett.83.4943](https://doi.org/10.1103/PhysRevLett.83.4943). arXiv: [hep-ex/9908011](https://arxiv.org/abs/hep-ex/9908011).
- [77] J. Dorenbosch et al. “A search for decays of heavy neutrinos in the mass range 0.5–2.8 GeV”. In: *Physics Letters B* 166.4 (1986), pp. 473–478. ISSN: 0370-2693. DOI: [https://doi.org/10.1016/0370-2693\(86\)91601-1](https://doi.org/10.1016/0370-2693(86)91601-1). URL: <https://www.sciencedirect.com/science/article/pii/0370269386916011>.
- [78] Georges Aad et al. “Search for Heavy Neutral Leptons in Decays of W Bosons Using a Dilepton Displaced Vertex in s=13 TeV pp Collisions with the ATLAS Detector”. In: *Phys. Rev. Lett.* 131.6 (2023), p. 061803. DOI: [10.1103/PhysRevLett.131.061803](https://doi.org/10.1103/PhysRevLett.131.061803). arXiv: [2204.11988](https://arxiv.org/abs/2204.11988) [hep-ex].

- [79] Ryan Plestid. “Luminous solar neutrinos II: Mass-mixing portals”. In: *Phys. Rev. D* 104 (2021). [Erratum: *Phys. Rev. D* 105, 099901 (2022)], p. 075028. DOI: [10.1103/PhysRevD.104.075028](https://doi.org/10.1103/PhysRevD.104.075028). arXiv: [2010.09523](https://arxiv.org/abs/2010.09523) [hep-ph].
- [80] R. Andrew Gustafson, Ryan Plestid and Ian M. Shoemaker. “Neutrino portals, terrestrial upscattering, and atmospheric neutrinos”. In: *Phys. Rev. D* 106.9 (2022), p. 095037. DOI: [10.1103/PhysRevD.106.095037](https://doi.org/10.1103/PhysRevD.106.095037). arXiv: [2205.02234](https://arxiv.org/abs/2205.02234) [hep-ph].
- [81] Iryna Boiarska et al. “Constraints from the CHARM experiment on heavy neutral leptons with tau mixing”. In: *Phys. Rev. D* 104.9 (2021), p. 095019. DOI: [10.1103/PhysRevD.104.095019](https://doi.org/10.1103/PhysRevD.104.095019). arXiv: [2107.14685](https://arxiv.org/abs/2107.14685) [hep-ph].
- [82] R. Acciarri et al. “New Constraints on Tau-Coupled Heavy Neutral Leptons with Masses $m_N=280-970$ MeV”. In: *Phys. Rev. Lett.* 127.12 (2021), p. 121801. DOI: [10.1103/PhysRevLett.127.121801](https://doi.org/10.1103/PhysRevLett.127.121801). arXiv: [2106.13684](https://arxiv.org/abs/2106.13684) [hep-ex].
- [83] J. P. Lees et al. “Search for heavy neutral leptons using tau lepton decays at BaBar”. In: *Phys. Rev. D* 107.5 (2023), p. 052009. DOI: [10.1103/PhysRevD.107.052009](https://doi.org/10.1103/PhysRevD.107.052009). arXiv: [2207.09575](https://arxiv.org/abs/2207.09575) [hep-ex].
- [84] *Heavy-Neutrino-Limits GitHub project*. 2025. URL: <https://github.com/mhostert/Heavy-Neutrino-Limits> (visited on 05/09/2025).
- [85] Peter Ballett, Matheus Hostert and Silvia Pascoli. “Dark Neutrinos and a Three Portal Connection to the Standard Model”. In: *Phys. Rev. D* 101.11 (2020), p. 115025. DOI: [10.1103/PhysRevD.101.115025](https://doi.org/10.1103/PhysRevD.101.115025). arXiv: [1903.07589](https://arxiv.org/abs/1903.07589) [hep-ph].
- [86] 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](https://doi.org/10.1103/PhysRevLett.121.221801). arXiv: [1805.12028](https://arxiv.org/abs/1805.12028) [hep-ex].
- [87] A. M. Abdullahi et al. “First Search for Dark Sector e^+e^- Explanations of the MiniBooNE Anomaly at MicroBooNE”. In: (Feb. 2025). arXiv: [2502.10900](https://arxiv.org/abs/2502.10900) [hep-ex].
- [88] Carlos A. Argüelles, Matheus Hostert and Yu-Dai Tsai. “Testing New Physics Explanations of the MiniBooNE Anomaly at Neutrino Scattering Experiments”. In: *Phys. Rev. Lett.* 123.26 (2019), p. 261801. DOI: [10.1103/PhysRevLett.123.261801](https://doi.org/10.1103/PhysRevLett.123.261801). arXiv: [1812.08768](https://arxiv.org/abs/1812.08768) [hep-ph].
- [89] Asli M. Abdullahi et al. “Panorama of new-physics explanations to the MiniBooNE excess”. In: *Phys. Rev. D* 111.3 (2025), p. 035028. DOI: [10.1103/PhysRevD.111.035028](https://doi.org/10.1103/PhysRevD.111.035028). arXiv: [2308.02543](https://arxiv.org/abs/2308.02543) [hep-ph].
- [90] Peter Ballett, Silvia Pascoli and Mark Ross-Lonergan. “U(1)’ mediated decays of heavy sterile neutrinos in MiniBooNE”. In: *Phys. Rev. D* 99 (2019), p. 071701. DOI: [10.1103/PhysRevD.99.071701](https://doi.org/10.1103/PhysRevD.99.071701). arXiv: [1808.02915](https://arxiv.org/abs/1808.02915) [hep-ph].
- [91] Enrico Bertuzzo et al. “Dark Neutrino Portal to Explain MiniBooNE excess”. In: *Phys. Rev. Lett.* 121.24 (2018), p. 241801. DOI: [10.1103/PhysRevLett.121.241801](https://doi.org/10.1103/PhysRevLett.121.241801). arXiv: [1807.09877](https://arxiv.org/abs/1807.09877) [hep-ph].

- [92] Pedro AN Machado, Ornella Palamara and David W Schmitz. “The Short-Baseline Neutrino Program at Fermilab”. In: *Ann. Rev. Nucl. Part. Sci.* 69 (2019), pp. 363–387. DOI: [10.1146/annurev-nucl-101917-020949](https://doi.org/10.1146/annurev-nucl-101917-020949). arXiv: [1903.04608](https://arxiv.org/abs/1903.04608) [hep-ex].
- [93] R. Acciarri et al. “Design and Construction of the MicroBooNE Detector”. In: *JINST* 12.02 (2017), P02017. DOI: [10.1088/1748-0221/12/02/P02017](https://doi.org/10.1088/1748-0221/12/02/P02017). arXiv: [1612.05824](https://arxiv.org/abs/1612.05824) [physics.ins-det].
- [94] R. Acciarri et al. “The Short-Baseline Near Detector at Fermilab: Input to the European Strategy for Particle Physics 2026 Update”. In: (Mar. 2025). arXiv: [2504.00245](https://arxiv.org/abs/2504.00245) [hep-ex].
- [95] P. Abratenko et al. “ICARUS at the Fermilab Short-Baseline Neutrino program: initial operation”. In: *Eur. Phys. J. C* 83.6 (2023), p. 467. DOI: [10.1140/epjc/s10052-023-11610-y](https://doi.org/10.1140/epjc/s10052-023-11610-y). arXiv: [2301.08634](https://arxiv.org/abs/2301.08634) [hep-ex].
- [96] 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](https://doi.org/10.1140/epjc/s10052-020-08456-z). arXiv: [2006.16043](https://arxiv.org/abs/2006.16043) [hep-ex].
- [97] Eduardo Cortina Gil et al. “The Beam and detector of the NA62 experiment at CERN”. In: *JINST* 12.05 (2017), P05025. DOI: [10.1088/1748-0221/12/05/P05025](https://doi.org/10.1088/1748-0221/12/05/P05025). arXiv: [1703.08501](https://arxiv.org/abs/1703.08501) [physics.ins-det].
- [98] C. Cazzaniga et al. “Probing the explanation of the muon (g-2) anomaly and thermal light dark matter with the semi-visible dark photon channel”. In: *Eur. Phys. J. C* 81.10 (2021), p. 959. DOI: [10.1140/epjc/s10052-021-09705-5](https://doi.org/10.1140/epjc/s10052-021-09705-5). arXiv: [2107.02021](https://arxiv.org/abs/2107.02021) [hep-ex].
- [99] M. Ball et al. “The PIP-II Conceptual Design Report”. In: (Mar. 2017). Ed. by Valeri Lebedev. DOI: [10.2172/1346823](https://doi.org/10.2172/1346823).
- [100] James Strait et al. “Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 3: Long-Baseline Neutrino Facility for DUNE June 24, 2015”. In: (Jan. 2016). arXiv: [1601.05823](https://arxiv.org/abs/1601.05823) [physics.ins-det].
- [101] Babak Abi et al. “Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics”. In: (Feb. 2020). arXiv: [2002.03005](https://arxiv.org/abs/2002.03005) [hep-ex].
- [102] Emile Lavaut. “Low energy calibration in DUNE far detector prototypes”. In: *59th Rencontres de Moriond on QCD and High Energy Interactions: Moriond QCD 2025*. May 2025. arXiv: [2505.18073](https://arxiv.org/abs/2505.18073) [hep-ex].
- [103] Babak Abi et al. “Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume IV: Far Detector Single-phase Technology”. In: *JINST* 15.08 (2020), T08010. DOI: [10.1088/1748-0221/15/08/T08010](https://doi.org/10.1088/1748-0221/15/08/T08010). arXiv: [2002.03010](https://arxiv.org/abs/2002.03010) [physics.ins-det].
- [104] Adam Abed Abud et al. “The DUNE Far Detector Vertical Drift Technology. Technical Design Report”. In: *JINST* 19.08 (2024), T08004. DOI: [10.1088/1748-0221/19/08/T08004](https://doi.org/10.1088/1748-0221/19/08/T08004). arXiv: [2312.03130](https://arxiv.org/abs/2312.03130) [hep-ex].

- [105] Sabrina Sacerdoti. “A LArTPC with Vertical Drift for the DUNE Far Detector”. In: *PoS NuFact2021* (2022), p. 173. DOI: [10.22323/1.402.0173](https://doi.org/10.22323/1.402.0173).
- [106] V. Hewes et al. “Deep Underground Neutrino Experiment (DUNE) Near Detector Conceptual Design Report”. In: *Instruments* 5.4 (2021), p. 31. DOI: [10.3390/instruments5040031](https://doi.org/10.3390/instruments5040031). arXiv: [2103.13910](https://arxiv.org/abs/2103.13910) [[physics.ins-det](#)].
- [107] L. Tortora *et al.* “The KLOE detector and physics program”. In: *Nuclear Physics B - Proceedings Supplements* 78.1 (1999). Advanced Technology and Particle Physics, pp. 157–162. ISSN: 0920-5632. DOI: [https://doi.org/10.1016/S0920-5632\(99\)00539-3](https://doi.org/10.1016/S0920-5632(99)00539-3).
- [108] M. Adinolfi *et al.* “The KLOE electromagnetic calorimeter”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 482.1 (2002), pp. 364–386. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(01\)01502-9](https://doi.org/10.1016/S0168-9002(01)01502-9). URL: <https://www.sciencedirect.com/science/article/pii/S0168900201015029>.
- [109] Jerome Gava et al. “A Dynamical collective calculation of supernova neutrino signals”. In: *Phys. Rev. Lett.* 103 (2009), p. 071101. DOI: [10.1103/PhysRevLett.103.071101](https://doi.org/10.1103/PhysRevLett.103.071101). arXiv: [0902.0317](https://arxiv.org/abs/0902.0317) [[hep-ph](#)].
- [110] Kate Scholberg. “Supernova Signatures of Neutrino Mass Ordering”. In: *J. Phys. G* 45.1 (2018), p. 014002. DOI: [10.1088/1361-6471/aa97be](https://doi.org/10.1088/1361-6471/aa97be). arXiv: [1707.06384](https://arxiv.org/abs/1707.06384) [[hep-ex](#)].
- [111] Francesco Capozzi et al. “DUNE as the Next-Generation Solar Neutrino Experiment”. In: *Phys. Rev. Lett.* 123.13 (2019), p. 131803. DOI: [10.1103/PhysRevLett.123.131803](https://doi.org/10.1103/PhysRevLett.123.131803). arXiv: [1808.08232](https://arxiv.org/abs/1808.08232) [[hep-ph](#)].
- [112] Klaes Møller et al. “Measuring the supernova unknowns at the next-generation neutrino telescopes through the diffuse neutrino background”. In: *JCAP* 05 (2018), p. 066. DOI: [10.1088/1475-7516/2018/05/066](https://doi.org/10.1088/1475-7516/2018/05/066). arXiv: [1804.03157](https://arxiv.org/abs/1804.03157) [[astro-ph.HE](#)].
- [113] Jim Alexander et al. “Dark Sectors 2016 Workshop: Community Report”. In: Aug. 2016. arXiv: [1608.08632](https://arxiv.org/abs/1608.08632) [[hep-ph](#)].
- [114] Valentina de Romeri, Kevin J. Kelly and Pedro A. N. Machado. “Hunting for light dark matter with DUNE PRISM”. In: *J. Phys. Conf. Ser.* 1468.1 (2020). Ed. by Masayuki Nakahata, p. 012061. DOI: [10.1088/1742-6596/1468/1/012061](https://doi.org/10.1088/1742-6596/1468/1/012061).
- [115] Lina Necib et al. “Boosted Dark Matter at Neutrino Experiments”. In: *Phys. Rev. D* 95.7 (2017), p. 075018. DOI: [10.1103/PhysRevD.95.075018](https://doi.org/10.1103/PhysRevD.95.075018). arXiv: [1610.03486](https://arxiv.org/abs/1610.03486) [[hep-ph](#)].
- [116] G. Adamov *et al.* *A proposal to enhance the DUNE Near-Detector Complex*. DUNE doc 1362. 2021. URL: <https://docs.dunescience.org/cgi-bin/private/ShowDocument?docid=13262>.
- [117] Valerio Pia. “Study of the calorimetric performance of a liquid Argon detector based on a novel optical imaging concept”. PhD thesis. amsdottorato, Alma Mater Studiorum - Università di Bologna, Bologna U., 2023. DOI: [10.48676/unibo/amsdottorato/10933](https://doi.org/10.48676/unibo/amsdottorato/10933).

- [118] Matteo Vicenzi. “A GRAIN of SAND for DUNE: Development of simulations and reconstruction algorithms for the liquid Argon target of the SAND detector in DUNE”. PhD thesis. Genoa U., 2023. DOI: [10.15167/vicenzi-matteo_phd2023-01-20](https://doi.org/10.15167/vicenzi-matteo_phd2023-01-20).
- [119] Valentina Cicero. “Study of the tracking performance of a liquid Argon detector based on a novel optical imaging concept”. PhD thesis. Alma Mater Studiorum - Università di Bologna, 2023. DOI: [10.48676/unibo/amsdottorato/10871..](https://doi.org/10.48676/unibo/amsdottorato/10871..)
- [120] H. Duyang et al. “A Novel Approach to Neutrino-Hydrogen Measurements”. In: (Sept. 2018). arXiv: [1809.08752](https://arxiv.org/abs/1809.08752) [[hep-ph](#)].
- [121] N. Doble et al. “The Upgrated muon beam at the SPS”. In: *Nucl. Instrum. Meth. A* 343 (1994), pp. 351–362. DOI: [10.1016/0168-9002\(94\)90212-7](https://doi.org/10.1016/0168-9002(94)90212-7).
- [122] Francesco Chiapponi. “Design, construction and operation of a Drift Chamber prototype for the inner tracker of SAND”. PhD thesis. amsdottorato, Alma Mater Studiorum - Università di Bologna, Bologna U., 2025.
- [123] L. Alvarez-Ruso et al. “NuSTEC White Paper: Status and challenges of neutrino–nucleus scattering”. In: *Prog. Part. Nucl. Phys.* 100 (2018), pp. 1–68. DOI: [10.1016/j.pnpnp.2018.01.006](https://doi.org/10.1016/j.pnpnp.2018.01.006). arXiv: [1706.03621](https://arxiv.org/abs/1706.03621) [[hep-ph](#)].
- [124] R. Acciarri et al. “Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE): Conceptual Design Report, Volume 2: The Physics Program for DUNE at LBNF”. In: (Dec. 2015). arXiv: [1512.06148](https://arxiv.org/abs/1512.06148) [[physics.ins-det](#)].
- [125] R. Petti. “Precision Measurements of Fundamental Interactions with (Anti)Neutrinos”. In: *27th International Workshop on Deep Inelastic Scattering and Related Subjects*. Oct. 2019. arXiv: [1910.05995](https://arxiv.org/abs/1910.05995) [[hep-ex](#)].
- [126] S. Alekhin, Sergey A. Kulagin and R. Petti. “Modeling lepton-nucleon inelastic scattering from high to low momentum transfer”. In: *AIP Conf. Proc.* 967.1 (2007). Ed. by Geralyn P. Zeller, Jorge G. Morfin and Flavio Cavanna, pp. 215–224. DOI: [10.1063/1.2834481](https://doi.org/10.1063/1.2834481). arXiv: [0710.0124](https://arxiv.org/abs/0710.0124) [[hep-ph](#)].
- [127] A. B. Arbuzov, D. Yu. Bardin and L. V. Kalinovskaya. “Radiative corrections to neutrino deep inelastic scattering revisited”. In: *JHEP* 06 (2005), p. 078. DOI: [10.1088/1126-6708/2005/06/078](https://doi.org/10.1088/1126-6708/2005/06/078). arXiv: [hep-ph/0407203](https://arxiv.org/abs/hep-ph/0407203).
- [128] S. A. Kulagin and R. Petti. “Neutrino inclusive inelastic scattering off nuclei”. In: *Phys. Rev. D* 76 (9 Nov. 2007), p. 094023. DOI: [10.1103/PhysRevD.76.094023](https://doi.org/10.1103/PhysRevD.76.094023). URL: <https://link.aps.org/doi/10.1103/PhysRevD.76.094023>.
- [129] William J. Marciano and Zohreh Parsa. “Neutrino electron scattering theory”. In: *J. Phys. G* 29 (2003), pp. 2629–2645. DOI: [10.1088/0954-3899/29/11/013](https://doi.org/10.1088/0954-3899/29/11/013). arXiv: [hep-ph/0403168](https://arxiv.org/abs/hep-ph/0403168).
- [130] Stephen L. Adler. “Tests of the Conserved Vector Current and Partially Conserved Axial-Vector Current Hypotheses in High-Energy Neutrino Reactions”. In: *Phys. Rev.* 135 (1964), B963–B966. DOI: [10.1103/PhysRev.135.B963](https://doi.org/10.1103/PhysRev.135.B963).
- [131] D. Allasia et al. “ Q^2 Dependence of the Proton and Neutron Structure Functions from Neutrino and anti-neutrinos Scattering in Deuterium”. In: *Z. Phys. C* 28 (1985), p. 321. DOI: [10.1007/BF01413595](https://doi.org/10.1007/BF01413595).

- [132] Sergey Alekhin et al. “Strange and non-strange distributions from the collider data”. In: *PoS DIS2018* (2018), p. 008. DOI: [10.22323/1.316.0008](https://doi.org/10.22323/1.316.0008). arXiv: [1808.06871](https://arxiv.org/abs/1808.06871) [hep-ph].
- [133] 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](https://doi.org/10.1016/j.nima.2009.12.009). arXiv: [0905.2517](https://arxiv.org/abs/0905.2517) [hep-ph].
- [134] J. Allison et al. “Recent developments in Geant4”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 835 (2016), pp. 186–225. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2016.06.125>. URL: <https://www.sciencedirect.com/science/article/pii/S0168900216306957>.
- [135] Fons Rademakers et al. *root-project/root: v6.20/06*. Version v6-20-06. June 2020. DOI: [10.5281/zenodo.3895852](https://doi.org/10.5281/zenodo.3895852). URL: <https://doi.org/10.5281/zenodo.3895852>.
- [136] *DUNE-ND-GGD*. 2025. URL: <https://github.com/DUNE/dunendggd> (visited on 08/08/2025).
- [137] *Dune fluxes repository*. 2025. URL: <https://glaucus.crc.nd.edu/DUNEFluxes/> (visited on 28/08/2025).
- [138] Júlia 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](https://doi.org/10.1103/PhysRevD.104.072009). arXiv: [2104.09179](https://arxiv.org/abs/2104.09179) [hep-ph].
- [139] *Dune fluxes repository*. 2025. URL: <https://glaucus.crc.nd.edu/DUNEFluxes/OptimizedEngineeredNov2017/> (visited on 28/08/2025).
- [140] *GENIE Spline Data Releases*. 2025. URL: https://scisoft.fnal.gov/scisoft/packages/genie_xsec/ (visited on 18/08/2025).
- [141] *EdepSim GitHub page*. 2025. URL: <https://github.com/ClarkMcGrew/edep-sim> (visited on 18/08/2025).
- [142] *EDepReader baltig page*. 2025. URL: <https://baltig.infn.it/dune/edep-reader> (visited on 25/08/2025).
- [143] *sandreco repository*. 2025. URL: <https://github.com/DUNE/sandreco> (visited on 08/10/2025).
- [144] H. Schindler and Rob Veenhof. *Garfield++ website*. Available at <https://garfieldpp.web.cern.ch/garfieldpp> [Accessed: 15 Oct. 2025].
- [145] Tomasz Golan, Cezary Juszczak and Jan T. Sobczyk. “Effects of final-state interactions in neutrino-nucleus interactions”. In: *Phys. Rev. C* 86 (1 July 2012), p. 015505. DOI: [10.1103/PhysRevC.86.015505](https://doi.org/10.1103/PhysRevC.86.015505). URL: <https://link.aps.org/doi/10.1103/PhysRevC.86.015505>.
- [146] O. Buss et al. “Transport-theoretical Description of Nuclear Reactions”. In: *Phys. Rept.* 512 (2012), pp. 1–124. DOI: [10.1016/j.physrep.2011.12.001](https://doi.org/10.1016/j.physrep.2011.12.001). arXiv: [1106.1344](https://arxiv.org/abs/1106.1344) [hep-ph].

- [147] Asli M. Abdullahi et al. “DarkNews: A Python-based event generator for heavy neutral lepton production in neutrino-nucleus scattering”. In: *Comput. Phys. Commun.* 297 (2024), p. 109075. DOI: [10.1016/j.cpc.2023.109075](https://doi.org/10.1016/j.cpc.2023.109075). arXiv: [2207.04137](https://arxiv.org/abs/2207.04137) [hep-ph].
- [148] G. Peter Lepage. “Adaptive multidimensional integration: VEGAS enhanced”. In: *J. Comput. Phys.* 439 (2021), p. 110386. DOI: [10.1016/j.jcp.2021.110386](https://doi.org/10.1016/j.jcp.2021.110386). arXiv: [2009.05112](https://arxiv.org/abs/2009.05112) [physics.comp-ph].
- [149] James M. Cline et al. “Update on scalar singlet dark matter”. In: *Phys. Rev. D* 88 (2013). [Erratum: *Phys.Rev.D* 92, 039906 (2015)], p. 055025. DOI: [10.1103/PhysRevD.88.055025](https://doi.org/10.1103/PhysRevD.88.055025). arXiv: [1306.4710](https://arxiv.org/abs/1306.4710) [hep-ph].
- [150] G. Fricke et al. “Nuclear Ground State Charge Radii from Electromagnetic Interactions”. In: *Atom. Data Nucl. Data Tabl.* 60 (1995), pp. 177–285. DOI: [10.1006/adnd.1995.1007](https://doi.org/10.1006/adnd.1995.1007).
- [151] Nicholas W. Kamp et al. “Dipole-coupled heavy-neutral-lepton explanations of the MiniBooNE excess including constraints from MINERvA data”. In: *Phys. Rev. D* 107.5 (2023), p. 055009. DOI: [10.1103/PhysRevD.107.055009](https://doi.org/10.1103/PhysRevD.107.055009). arXiv: [2206.07100](https://arxiv.org/abs/2206.07100) [hep-ph].
- [152] The pandas development team. *pandas-dev/pandas: Pandas*. Version 2.3.2. Aug. 2025. DOI: [10.5281/zenodo.3509134](https://doi.org/10.5281/zenodo.3509134). URL: <https://doi.org/10.5281/zenodo.3509134>.
- [153] Wes McKinney. “Data Structures for Statistical Computing in Python”. In: *Proceedings of the 9th Python in Science Conference*. Ed. by Stéfan van der Walt and Jarrod Millman. 2010, pp. 56–61. DOI: [10.25080/Majora-92bf1922-00a](https://doi.org/10.25080/Majora-92bf1922-00a).
- [154] Charles R. Harris et al. “Array programming with NumPy”. In: *Nature* 585.7825 (Sept. 2020), pp. 357–362. DOI: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2). URL: <https://doi.org/10.1038/s41586-020-2649-2>.
- [155] Andy Buckley et al. “The HepMC3 event record library for Monte Carlo event generators”. In: *Comput. Phys. Commun.* 260 (2021), p. 107310. DOI: [10.1016/j.cpc.2020.107310](https://doi.org/10.1016/j.cpc.2020.107310). arXiv: [1912.08005](https://arxiv.org/abs/1912.08005) [hep-ph].
- [156] Iker Loïc de Icaza Astiz, Marco Roda and Pedro Machado. *DarkNeutrino module*. 2025. URL: <https://github.com/GENIE-MC/Generator/tree/master/src/Physics/DarkNeutrino> (visited on 26/08/2025).
- [157] J. Engel. “Nuclear form-factors for the scattering of weakly interacting massive particles”. In: *Phys. Lett. B* 264 (1991), pp. 114–119. DOI: [10.1016/0370-2693\(91\)90712-Y](https://doi.org/10.1016/0370-2693(91)90712-Y).
- [158] *Displaced primaries edep-sim fork*. 2025. URL: <https://github.com/AlessandroRuggeri/edep-sim> (visited on 26/08/2025).
- [159] Robert Piessens et al. “Guidelines for the Use of QUADPACK”. In: *Quadpack: A Subroutine Package for Automatic Integration*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1983, pp. 75–111. ISBN: 978-3-642-61786-7. DOI: [10.1007/978-3-642-61786-7_4](https://doi.org/10.1007/978-3-642-61786-7_4). URL: https://doi.org/10.1007/978-3-642-61786-7_4.

- [160] *GENIE DarkNeutrino integrator fork*. 2025. URL: <https://github.com/AlessandroRuggeri/Generator> (visited on 07/09/2025).
- [161] Gianfranco Ingrassia. “Study of Neutrino interactions on Hydrogen in the SAND detector of DUNE”. PhD thesis. alma, Mar. 2025. URL: <https://amsdottorato.unibo.it/id/eprint/11966/>.
- [162] Claus Grupen and Boris Schwartz. *Particle Detectors, Second Edition*. Cambridge University Press, 2008. ISBN: 978-1-009-40153-1, 978-1-009-40149-4, 978-1-009-40151-7. DOI: [10.1017/9781009401531](https://doi.org/10.1017/9781009401531).
- [163] *Fitter class (ROOT Documentation)*. Accessed 2025-09-25. 2025. URL: https://root.cern.ch/doc/master/classROOT_1_1Fit_1_1Fitter.html.
- [164] F. James, M. Winkler and ROOT Team. *Minuit2 Minimization Library (ROOT Documentation)*. Accessed 2025-09-25. 2025. URL: <https://root.cern.ch/doc/master/Minuit2Page.html>.
- [165] I. Kása. “A circle fitting procedure and its error analysis”. In: *IEEE Transactions on Instrumentation and Measurement* IM-25.1 (1976), pp. 8–14. DOI: [10.1109/TIM.1976.6312298](https://doi.org/10.1109/TIM.1976.6312298).
- [166] A. Hoecker et al. *TMVA - Toolkit for Multivariate Data Analysis*. 2009. arXiv: [physics/0703039](https://arxiv.org/abs/physics/0703039) [physics.data-an]. URL: <https://arxiv.org/abs/physics/0703039>.
- [167] Yann Coadou. “Boosted decision trees”. In: (Mar. 2022). DOI: [10.1142/9789811234033_0002](https://doi.org/10.1142/9789811234033_0002). arXiv: [2206.09645](https://arxiv.org/abs/2206.09645) [physics.data-an].
- [168] Yoav Freund and Robert E Schapire. “A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting”. In: *Journal of Computer and System Sciences* 55.1 (1997), pp. 119–139. ISSN: 0022-0000. DOI: <https://doi.org/10.1006/jcss.1997.1504>. URL: <https://www.sciencedirect.com/science/article/pii/S002200009791504X>.
- [169] Glen Cowan et al. “Asymptotic formulae for likelihood-based tests of new physics”. In: *Eur. Phys. J. C* 71 (2011). [Erratum: *Eur.Phys.J.C* 73, 2501 (2013)], p. 1554. DOI: [10.1140/epjc/s10052-011-1554-0](https://doi.org/10.1140/epjc/s10052-011-1554-0). arXiv: [1007.1727](https://arxiv.org/abs/1007.1727) [physics.data-an].
- [170] Leo Breiman et al. *Classification and Regression Trees*. Wadsworth, 1984. ISBN: 0-534-98053-8.
- [171] William H. Press et al. *Numerical Recipes: The Art of Scientific Computing (Third Edition)*. Cambridge University Press, 2007.
- [172] R.J. Yarema et al. “A high performance multi-channel preamplifier ASIC”. In: *IEEE Transactions on Nuclear Science* 39.4 (1992), pp. 742–746. DOI: [10.1109/23.159698](https://doi.org/10.1109/23.159698).
- [173] A. Rivetti et al. “TIGER: A front-end ASIC for timing and energy measurements with radiation detectors”. In: *Nucl. Instrum. Meth. A* 924 (2019), pp. 181–186. DOI: [10.1016/j.nima.2018.09.010](https://doi.org/10.1016/j.nima.2018.09.010).
- [174] F. Pedregosa et al. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830.

-
- [175] François Chollet et al. *Keras*. <https://keras.io>. 2015.
 - [176] Adam Paszke et al. “PyTorch: An Imperative Style, High-Performance Deep Learning Library”. In: (Dec. 2019). arXiv: [1912.01703](https://arxiv.org/abs/1912.01703) [[cs.LG](#)].
 - [177] Tianqi Chen and Carlos Guestrin. “XGBoost: A Scalable Tree Boosting System”. In: (Mar. 2016). DOI: [10.1145/2939672.2939785](https://doi.org/10.1145/2939672.2939785). arXiv: [1603.02754](https://arxiv.org/abs/1603.02754) [[cs.LG](#)].