



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

DEPARTMENT OF PHYSICS AND ASTRONOMY "A. RIGHI"

SECOND CYCLE DEGREE

PHYSICS

Improving Electromagnetic Shower Reconstruction in LArTPC-based Neutrino Detectors: ICARUS and ArCS

Supervisor

Prof. Nicoletta Mauri

Co-supervisors

Dr. Angela Fava

Dr. Marco del Tutto

Defended by

Giulia Cicogna

March 2026

Academic Year 2024/2025

*A Valeria,
sempre.*

Abstract

Electromagnetic shower reconstruction represents one of the most challenging aspects of event reconstruction in neutrino liquid argon time projection chambers (LArTPCs). Shower fragmentation, overlap with hadronic activity, and ambiguities between electrons, photons, and secondary electromagnetic components directly affect neutrino physics measurements, including charged-current ν_e interactions and appearance searches. Improving electromagnetic shower reconstruction is therefore a central requirement for fully exploiting the physics potential of modern neutrino LArTPC experiments.

This thesis addresses this problem through two complementary approaches. The first acts at the detector-design level and is developed within the ArCS (Argon detector with Charge Separation) R&D program, which aims at studying the operation of a magnetized LArTPC in a test-beam environment. The presence of a magnetic field enables charge-sign identification and curvature-based momentum reconstruction, thus yielding additional handles for the reconstruction of electromagnetic topologies. The work presented here focuses on DAQ-related tests and validation of the inherited LArIAT acquisition chain in preparation for future ArCS data-taking operations.

The second approach acts at the pattern-recognition level within the ICARUS experiment, which is a large-scale LArTPC currently operating at Fermilab within the broader Short-Baseline Neutrino (SBN) program. In particular, this thesis presents the retraining and validation of NuGraph2, a graph neural network (GNN) algorithm performing hit-level background filtering and semantic classification. Following the transition of ICARUS to an updated upstream reconstruction based on two-dimensional (2D) signal deconvolution and deep neural network region of interest identification, a dedicated NuGraph2 retraining is required to ensure consistency with the modified hit distributions. The new training is performed using official beam Monte Carlo samples and dedicated datasets enriched in electromagnetic topologies. Finally, preliminary studies are presented on the impact of NuGraph2 outputs at the end of the reconstruction chain, within ν_e charged-current interactions, and, more specifically, compared to the standard reconstruction.

Contents

Abstract	i
Introduction	1
1 Neutrinos in the Standard Model and Beyond	3
1.1 An Overview on Neutrinos Physics	3
1.2 A Closer Look at the Standard Model	4
1.3 Neutrino Mass	6
1.3.1 Acquisition of Dirac Masses by Neutrinos	6
1.3.2 Neutrinos as Majorana Particles	7
1.3.3 Effective Description of Neutrino Mass Generation	8
1.4 Neutrino Oscillation Theory	8
1.4.1 Plane-wave Approximation Mixing	9
1.4.2 Two-Neutrino Mixing Example	11
1.5 Experimental Neutrino Anomalies at Short Baseline	12
1.5.1 Reactor Antineutrino Anomaly	13
1.5.2 Gallium Anomaly	14
1.5.3 Accelerator-based Appearance Anomalies: LSND and MiniBooNE	14
1.5.4 Global Fits and Open Questions	15
1.6 The Fermilab Short-Baseline Neutrino Program	16
2 LArTPCs: the ICARUS detector and R&D perspectives	23
2.1 Detection in LArTPCs and their Operation	23
2.1.1 Energy deposition, ionization and scintillation	24
2.1.2 Electron drift and transport in liquid argon	25
2.1.3 Impurities, electron lifetime and charge attenuation	27
2.2 The ICARUS Detector	30
2.2.1 The detector architecture	31
2.2.2 From hardware to performance	35
2.3 Motivation for R&D in LArTPC-based Neutrino Experiments	38
2.3.1 The ArCS Experiment	39

3	Event Reconstruction in the ICARUS LArTPC	45
3.1	Overview of the ICARUS Reconstruction Chain	45
3.1.1	Signal Processing	45
3.1.2	Three-Dimensional Clustering: Cluster3D	47
3.1.3	Pattern Recognition with Pandora	47
3.1.4	Two-Dimensional Deconvolution	50
3.1.5	Deep Neural Network ROI identification	52
3.1.6	NuGraph2: Graph Neural Network Reconstruction in ICARUS	53
4	NuGraph2 Training with the Updated ICARUS Reconstruction Chain	63
4.1	Motivation for Retraining	63
4.2	Training Sample Production	63
4.2.1	Official Beam Samples	63
4.3	Training Sample Characterization	66
4.3.1	Beam-Induced Samples: BNB vs NuMI	66
4.3.2	MPV/MPR High-Multiplicity Sample	69
4.4	Training Results and Pre-Processing Studies	71
4.4.1	Evaluation on the Test Sample	73
5	Electron Candidate Selection with NuGraph2	77
5.1	Results	82
	Conclusions	87
	Acknowledgements	89
A	DAQ Tests for ArCS	91
A.0.1	Scope and Strategy	91
A.0.2	Standalone Digitizer Test	92
A.0.3	Analysis and Restoration of the LArIAT DAQ Infrastructure	94
A.0.4	Persistent Zero-Fragment Condition	95
A.0.5	Development Environment and Portability	96
	References	97

Introduction

Neutrino experiments based on liquid argon time projection chambers (LArTPCs) provide high-resolution imaging and calorimetric information, enabling detailed reconstruction of interaction topologies in the GeV energy range. However, the extraction of physics observables from LArTPC data is intrinsically limited reconstruction-wise. Raw wire waveforms must be converted into calibrated hits and subsequently assembled into tracks, showers, vertices, and full interaction hierarchies. Among all reconstructed objects, electromagnetic (EM) showers constitute one of the most problematic topologies.

Electromagnetic showers are intrinsically diffuse and non-linear. Their development may involve early photon conversions, branching structures, and overlapping activity with hadronic fragments. In surface detectors such as ICARUS, additional cosmic-ray activity further complicates event topology. As a consequence, shower fragmentation, incomplete clustering, and ambiguities between electrons and photons are dominant limitations for $CC \nu_e$ measurements and, more generally, for precise neutrino interaction studies. Improving EM shower reconstruction is therefore a structural requirement for the success of current neutrino programs such as the Short-Baseline Neutrino (SBN) program at Fermilab, and becomes even more relevant in view of next-generation long-baseline experiments that will rely on LArTPC technology.

There are two broad strategies to mitigate the intrinsic limitations of electromagnetic reconstruction in LArTPCs.

The first strategy acts at the detector-design level. In a non-magnetized LArTPC, the charge sign of final-state leptons cannot be determined and momentum reconstruction relies primarily on range or multiple Coulomb scattering, which are both subject to limitations. Introducing a magnetic field would enable curvature-based momentum reconstruction and charge-sign identification, providing an additional handle on event reconstruction. This motivates R&D efforts toward magnetized LArTPCs. The ArCS (Argon detector with Charge Separation) prototype was developed in this context as a test-beam program aimed at evaluating the performance of a magnetized LArTPC. ArCS reuses the LArIAT (Liquid Argon In A Testbeam) detector and operates inside the Jolly Green Giant dipole magnet at the Fermilab Test Beam Facility.

Within ArCS, the contribution presented in this thesis focuses on DAQ-related tests

regarding the inherited LArIAT acquisition system. These studies were essential for the forthcoming ArCS data-taking campaigns.

The second strategy acts at the reconstruction and pattern-recognition level. Modern machine learning techniques are capable of exploiting complex correlations in detector observables that are difficult to encode in deterministic algorithms. In ICARUS, the standard reconstruction chain combines upstream signal processing with downstream pattern recognition based on the Pandora framework, where Pandora is a multi-purpose algorithm for dedicated pattern recognition tasks within LArTPC-based neutrino experiments. Recent upgrades introduced an updated upstream configuration based on two-dimensional signal deconvolution and deep neural network Region of Interest identification, improving charge stability and reducing orientation-dependent effects.

Since graph neural networks operate directly on reconstructed hits, changes in the upstream signal processing modify the feature space on which such algorithms rely. NuGraph2 is a graph neural network integrated into the ICARUS reconstruction framework that performs hit-level background filtering and semantic classification. As of now, its outputs are intended to complement, rather than replace, the classical Pandora reconstruction. In this thesis, NuGraph2 is retrained on samples reconstructed with the updated 2D signal processing chain to preserve internal consistency and performance. Dedicated training samples are produced and characterized. Preliminary studies are then performed to assess the impact of NuGraph2 outputs at the final reconstructed particle objects, namely Particle Flow Particle (PFP), in ν_e charged-current interactions. The structure of this thesis is as follows.

Chapter 1 introduces the theoretical framework of neutrino physics and motivates the need for an extensive short-baseline program such as SBN.

Chapter 2 presents the LArTPC technology and the ICARUS detector, and discusses the motivations for dedicated R&D activities, with particular focus on the ArCS magnetized LArTPC prototype.

Chapter 3 provides an overview of the ICARUS reconstruction chain, describing both classical algorithms and machine learning approaches, and introduces the NuGraph2 architecture.

Chapter 4 describes the retraining of NuGraph2 performed in the context of the updated 2D signal processing, including the preparation and characterization of the training samples, as well as its hit-level performance.

Chapter 5 presents preliminary studies of reconstructed PFP in ν_e charged-current interactions, comparing the standard Pandora reconstruction with the additional information provided by NuGraph2.

AI tools (e.g., ChatGPT and Grammarly) have been used for proofreading and editing.

Chapter 1

Neutrinos in the Standard Model and Beyond

1.1 An Overview on Neutrinos Physics

In 1930, Wolfgang Pauli proposed the existence of a neutral and elusive particle to address apparent violations of energy and angular momentum conservation emerging from nuclear beta decay. Pauli's proposal was sustained by the idea that the missing energy and angular momentum had been taken away by a hypothetical particle that had not yet been detected and that was eventually named "neutrino" by Enrico Fermi in order to distinguish it from the heavy neutron discovered by James Chadwick in 1932 [1]. This marked the start of neutrino physics.

Fermi proposed his theoretical framework soon after, in 1934, laying the foundations of the weak interaction. In his theory, the neutron, proton, electron, and neutrino participate in a point-like interaction governed by the Fermi constant, whose small value inherently reflects the weak nature of the force. This description was successful, but the physical existence of the neutrino was still to be certified.

The first direct evidence came two decades later, in 1956, when Frederick Reines and Clyde Cowan were able to detect antineutrinos emitted from a nuclear reactor at Savannah River through the inverse beta decay process, finally establishing the neutrino as a physical particle [2, 3].

In the following years, a series of experiments succeeded one another and revealed several distinctive features of neutrinos and weak interactions. The theoretical proposal of parity violation in weak interactions by Tsung-Dao Lee and Chen Ning Yang in response to the τ - θ puzzle led to the 1956 experiment conducted by Chien-Shiung Wu, whose measurements of beta decay on Co-60 demonstrated parity violation, a symmetry that is conserved in electromagnetic and strong interactions [4]. Shortly after, Maurice

Goldhaber showed that neutrinos are left-handed particles, meaning that their spin is anti-aligned with their momentum [5]. These results ultimately revealed the maximally parity-violating nature of the weak interaction.

As experiments advanced, theoretical developments had to keep pace. By the late 1960s, Abdus Salam, Sheldon Glashow, and Steven Weinberg managed to incorporate neutrinos in the Standard Model (SM) as fundamental fermions interacting via the massive W^\pm and Z^0 gauge bosons, thereby formulating the electroweak theory, a framework where neutrinos interact exclusively through their left-chiral components. It eventually became apparent that the weakness of the weak force at low energies is due to these massive mediating gauge bosons suppressing the interaction strength with respect to electromagnetic processes, rather than a small coupling constant.

In 1962, the muon neutrino was identified at the Brookhaven Alternating Gradient Synchrotron, while the tau neutrino was observed in 2000 by the DONUT experiment at Fermilab. Three neutrino species were found, and the existence of exactly three active neutrino species was confirmed at the Large Electron–Positron Collider in the early 1990s by studying the decay width of the Z^0 boson.

In the 1960s, Ray Davis’s Homestake experiment measured a flux of solar neutrinos significantly lower than theoretical predictions, a discrepancy which would later be known as the solar neutrino problem. In the 1990s, the Super-Kamiokande experiment observed a directional dependence in atmospheric neutrinos, with fewer neutrinos detected after they had traversed the Earth. The solar neutrino problem found its solution in the early 2000s at the Sudbury Neutrino Observatory, where it was discovered that electron neutrinos produced in the Sun changed into other flavors before reaching Earth [6].

Altogether, these findings hinted that neutrinos could change flavor while propagating and that the SM, in its original description of neutrinos, was far from complete. A turning point was marked, and neutrino oscillations became a central subject of experimental investigation, leading to the design of a new generation of dedicated experiments.

1.2 A Closer Look at the Standard Model

The SM of particle physics is a quantum field theory based on the symmetry group $G = SU(3)_C \times SU(2)_L \times U(1)_Y$ that describes the strong and electroweak interactions through local gauge invariance. All known elementary particles are encoded in it, including quarks, leptons, and their corresponding antiparticles, as well as the bosons that mediate the fundamental forces: gluons for the strong force, W and Z bosons for the weak force, and photons for the electromagnetic force. This framework also includes the Higgs boson, a scalar particle that was introduced to explain the genesis of particle masses via the Higgs mechanism, through which the $SU(2)_L \times U(1)_Y$ symmetry

spontaneously breaks at low energies, unifying weak and electromagnetic forces under the electroweak interaction.

Matter particles are organized into three distinct generations of quarks and leptons, each including one quark doublet and one lepton doublet under the $SU(2)_L$ symmetry, along with their corresponding right-handed singlets under $U(1)_Y$. For leptons, the left-handed doublets are $L_\ell = (\nu_\ell, \ell)_L$ with corresponding right-handed charged leptons ℓ_R , for $\ell = e, \mu, \tau$.

Neutrinos in the SM Neutrinos in the SM are described as neutral, weakly interacting leptons. Organized into doublets under $SU(2)_L$ symmetry with their corresponding charged leptons, they are described as massless and purely left-handed particles, such that

$$\psi_\ell = \begin{pmatrix} \nu_\ell \\ \ell^- \end{pmatrix}_L, \quad \ell = e, \mu, \tau. \quad (1.1)$$

This structure stems from the electroweak gauge symmetry $SU(2)_L \times U(1)_Y$, which governs their interactions via the W^\pm and Z^0 bosons.

Neutrino interactions are precisely described by the SM electroweak Lagrangian. The charged-current (CC) interaction involves the coupling of neutrinos to their charged lepton counterparts via the W^\pm bosons, giving

$$\mathcal{L}_{\text{CC}} = -\frac{g}{2\sqrt{2}} (j_{W,L}^\rho W_\rho + j_{W,L}^\rho \dagger W_\rho^\dagger), \quad (1.2)$$

where g is the $SU(2)_L$ gauge coupling. The leptonic charged current $j_{W,L}^\rho$ is defined as

$$j_{W,L}^\rho = 2 \sum_{\alpha=e,\mu,\tau} \bar{\nu}_\alpha \gamma^\rho \frac{1-\gamma^5}{2} \ell_\alpha, \quad (1.3)$$

where γ^ρ are the Dirac gamma matrices and $\frac{1-\gamma^5}{2} \equiv P_L$ is the left-handed chiral projector.

The neutral-current (NC) interaction instead governs neutrino self-interactions mediated by the Z^0 boson, for which

$$\mathcal{L}_{\text{NC}} = -\frac{g}{2 \cos \theta_W} j_{Z,\nu}^\rho Z_\rho, \quad (1.4)$$

where θ_W is the Weinberg angle. The neutral current is expressed as

$$j_{Z,\nu}^\rho = \sum_{\alpha=e,\mu,\tau} \bar{\nu}_\alpha \gamma^\rho \frac{1-\gamma^5}{2} \nu_\alpha. \quad (1.5)$$

NC interactions conserve lepton flavor, while CC interactions couple each neutrino flavor to its corresponding charged lepton. Flavor mixing effectively becomes relevant only once neutrino masses are introduced.

Masslessness of Neutrinos in the SM In the SM, neutrinos are massless because they lack right-handed components ν_R , which are necessary to generate Dirac mass terms through Yukawa interactions with the Higgs field. Moreover, the SM does not allow Majorana mass terms, since these would require either an extension of the field content or a violation of gauge invariance. For this reason, neutrino oscillations require an extension of the SM to explain flavour mixing, which in turn requires non-zero masses.

1.3 Neutrino Mass

The experimental observation of neutrino flavor oscillations requires an extension of the SM to address the generation of non-zero neutrino masses and allow mixing between flavor states.

1.3.1 Acquisition of Dirac Masses by Neutrinos

A first extension of the SM, referred to as the minimally extended SM, introduces right-handed neutrino fields $\nu_{i,R}$, which are singlets under $SU(3)_C \times SU(2)_L$ with null hypercharge. In this way, a Dirac neutrino mass arises via Yukawa interaction with the Higgs field through a process analogous to that for other fermions.

The Yukawa interaction responsible for Dirac neutrino masses is

$$\mathcal{L}_Y^\nu = - \sum_{\alpha,\beta} Y_{\alpha\beta}^\nu \overline{L_{\alpha L}} \tilde{\phi} \nu_{\beta R} + \text{h.c.}, \quad (1.6)$$

where $L_{\alpha L}$ is the left-handed lepton doublet and $\tilde{\phi} = i\sigma_2\phi^*$. After electroweak symmetry breaking, $\langle\phi\rangle = v/\sqrt{2}$, which generates Dirac mass terms

$$m_k = \frac{y_k^\nu v}{\sqrt{2}}. \quad (1.7)$$

Neutrino masses arise from $SU(2)_L \times U(1)_Y$ -invariant Yukawa interactions. The total lepton number is conserved, and neutrinos and antineutrinos remain distinct particles, since both the fields $\nu(x)$ corresponding to neutrinos ($L = +1$) and antineutrinos ($L = -1$) are included in this theory. However, to generate sub-eV neutrino masses, Yukawa couplings must be extremely small, orders of magnitude below those of charged fermions. Dirac neutrino masses require $y_k^\nu \sim 10^{-12}$ or smaller, while, for comparison, the top-quark Yukawa coupling is of order unity, $y_t \sim \mathcal{O}(1)$, while the electron Yukawa coupling is of order $y_e \sim 10^{-6}$. We would thus have a much stronger suppression of the Yukawa interaction with respect to the charged fermion sector. Neutrinos can also admit Majorana mass terms, and alternative mechanisms for neutrino mass generation can be explored.

1.3.2 Neutrinos as Majorana Particles

An alternative possibility for neutrino mass generation is that neutrinos are Majorana particles, a concept introduced by Ettore Majorana in 1937, in which neutrinos are identical to their own antiparticles. This characteristic is summarized through the Majorana condition, such that

$$\psi = \psi^c = C\bar{\psi}^T, \quad (1.8)$$

where ψ is the field operator of the particle, ψ^c is its charge-conjugated field, and C is the charge conjugation matrix.

Given the electrical neutrality of neutrinos, we can generate a Majorana mass term from solely left-handed fields, with

$$\mathcal{L}_M = -\frac{1}{2} \sum_{l', l} \bar{\nu}_{l'L} M_M^{l'l} \nu_{lL}^c + \text{h.c.} \quad (1.9)$$

Here, $\nu_{l'L}$ and ν_{lL} represent the left-handed neutrino fields, with l' and l denoting the neutrino flavors ($l, l' = e, \mu, \tau$). The matrix M_M is the Majorana mass matrix, which is a 3×3 complex symmetric matrix that is generally non-diagonal, such that $M_M = M_M^T$.

Here lies the major difference between the Majorana and the Dirac mass terms, since the latter does not require such a constraint, given that it involves both left-handed and right-handed fields.

The matrix M_M can be diagonalized by a unitary transformation,

$$M_M = U m U^T, \quad (1.10)$$

which defines the mixing between flavor and mass eigenstates.

The bilinear term in the Majorana mass term couples a neutrino field to its charge-conjugate counterpart, thus violating lepton number by two units $\Delta L = -2$, or equivalently $\Delta L = +2$ for the reverse process.

A direct consequence of $\Delta L = 2$ violation is neutrinoless double beta decay,

$$(A, Z) \rightarrow (A, Z + 2) + 2e^-, \quad (1.11)$$

whose observation would establish the Majorana nature of neutrinos, and whose decay rate depends on the effective Majorana mass

$$m_{\beta\beta} = \left| \sum_i U_{ei}^2 m_i \right|. \quad (1.12)$$

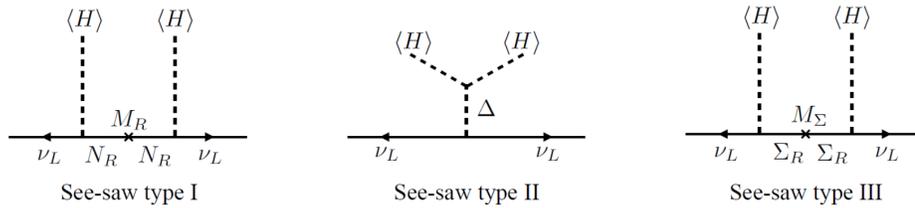


Figure 1.1: Tree-level realizations of the Weinberg operator in the three seesaw mechanisms. From left to right: Type-I seesaw, mediated by heavy fermion singlets N_R ; Type-II seesaw, mediated by a scalar triplet Δ ; and Type-III seesaw, mediated by fermion triplets Σ . Insertions of the Higgs vacuum expectation value $\langle H \rangle$ generate an effective Majorana mass term for light neutrinos after electroweak symmetry breaking [7].

1.3.3 Effective Description of Neutrino Mass Generation

At low energies, Majorana neutrino masses can be described by the dimension-five Weinberg operator,

$$\mathcal{L}_{\text{eff}} = -\frac{1}{\Lambda} \overline{(\tilde{\phi}^\dagger \psi_{lL})} X_{ll'} (\tilde{\phi}^\dagger \psi_{l'L})^c + \text{h.c.}, \quad (1.13)$$

which violates lepton number by two units. After electroweak symmetry breaking,

$$m_\nu \sim \frac{v^2}{\Lambda}. \quad (1.14)$$

The scale Λ suppresses the effective operator and is to be identified with the mass scale of the heavy mediators responsible for neutrino mass generation. Models in which light neutrino masses are suppressed by the ratio between the electroweak scale and a much larger mass scale are commonly referred to as seesaw mechanisms. Depending on the nature of the exchanged particles, three main models are distinguished:

- Type-I seesaw, involving heavy fermion singlets;
- Type-II seesaw, mediated by scalar triplets;
- Type-III seesaw, involving fermion triplets.

A schematic representation of the tree-level realization of the Weinberg operator in the three seesaw scenarios is shown in Fig. 1.1. Although these differ in their particle content, they all generate the same effective Weinberg operator at low energies and lead to the generation of Majorana masses.

1.4 Neutrino Oscillation Theory

After discussing how neutrino masses can arise by extending the SM, let us now turn to how these masses lead to the phenomenon of neutrino oscillations. Neutrinos ν_α with

flavor $\alpha = e, \mu, \tau$ are produced in CC weak interaction processes. In the presence of mixing, the leptonic charged current can be written as

$$j_{W,L}^\rho = 2 \sum_{\alpha=e,\mu,\tau} \bar{\nu}_{\alpha L} \gamma^\rho \ell_{\alpha L} = 2 \sum_{\alpha=e,\mu,\tau} \sum_k U_{\alpha k}^* \bar{\nu}_{kL} \gamma^\rho \ell_{\alpha L}. \quad (1.15)$$

Neutrinos are produced and detected as flavor eigenstates, while they propagate as mass eigenstates. As a result, the flavor states can be expressed as coherent superpositions of massive neutrino states through the mixing matrix as

$$|\nu_\alpha\rangle = \sum_k U_{\alpha k}^* |\nu_k\rangle. \quad (1.16)$$

Neutrino oscillations arise from the quantum mechanical phenomenon where neutrinos of a given flavor, ν_e , ν_μ , or ν_τ , oscillate into other flavors as they propagate through space. This behavior is due to the misalignment between the flavor eigenstates, which are defined by weak interactions, and the mass eigenstates, which define the neutrino's propagation in vacuum. These two bases are connected via the unitary matrix $U_{\alpha k}$ which appears in the charged-current interaction, and is the same mixing matrix introduced in the diagonalization of the neutrino mass term. It is identified with the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, with its mixing angles and possible CP-violating phases. The PMNS mixing matrix can be explicitly written as

$$U_{\alpha k} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \cdot \mathcal{P}, \quad (1.17)$$

where $c_{ij} = \cos \theta_{ij}$ and $s_{ij} = \sin \theta_{ij}$ represent the trigonometric functions of the three mixing angles θ_{12} , θ_{13} , and θ_{23} , while δ is the Dirac CP-violating phase. The term \mathcal{P} is a diagonal matrix of potential Majorana phases that becomes relevant only if neutrinos are Majorana particles since

$$\mathcal{P} = \text{diag}(1, e^{i\phi_2}, e^{i\phi_3}). \quad (1.18)$$

For Dirac neutrinos, the phases associated with \mathcal{P} can be removed by redefining the neutrino fields and have no physical significance. For Majorana neutrinos, instead, it is not the case due to the self-conjugate nature of Majorana fields. It is important to note that the Majorana phases contained in \mathcal{P} do not affect neutrino oscillation probabilities, which depend only on the mixing angles, the Dirac CP-violating phase, and the squared-mass differences.

1.4.1 Plane-wave Approximation Mixing

In the plane-wave approximation, the time evolution of a flavor state $|\nu_\alpha\rangle$ initially defined at $t = 0$ is given by

$$|\nu_\alpha(t)\rangle = \sum_k U_{\alpha k}^* e^{-iE_k t} |\nu_k\rangle, \quad (1.19)$$

such that $|\nu_\alpha(t=0)\rangle = |\nu_\alpha\rangle$. Following the convention introduced in Eq. (1.16), flavor states are expressed in terms of mass eigenstates with coefficients $U_{\alpha k}^*$. Therefore, the flavor state becomes

$$|\nu_\alpha(t)\rangle = \sum_{\beta=e,\mu,\tau} \left(\sum_k U_{\alpha k}^* e^{-iE_k t} U_{\beta k} \right) |\nu_\beta\rangle. \quad (1.20)$$

As a consequence, the flavor state $|\nu_\alpha(t)\rangle$, which at $t=0$ is a pure flavor state, evolves into a superposition of different flavor states at $t>0$, assuming the mixing matrix U is not diagonal, i.e., neutrinos are mixed. The transition amplitude for a neutrino created in a flavor state $|\nu_\alpha\rangle$ to be detected as $|\nu_\beta\rangle$ at time t is then

$$\mathcal{A}_{\alpha\rightarrow\beta}(t) = \langle\nu_\beta|\nu_\alpha(t)\rangle = \sum_k U_{\alpha k}^* U_{\beta k} e^{-iE_k t}. \quad (1.21)$$

The corresponding transition probability is

$$P_{\nu_\alpha\rightarrow\nu_\beta}(t) = |\mathcal{A}_{\alpha\rightarrow\beta}(t)|^2 = \sum_{k,j} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* e^{-i(E_k - E_j)t}. \quad (1.22)$$

For ultrarelativistic neutrinos, the energy of a mass eigenstate can be rewritten as

$$E_k \approx E + \frac{m_k^2}{2E}, \quad (1.23)$$

where E is the neutrino energy, approximated by $|\vec{p}|$, and m_k^2 is the squared mass of the eigenstate. The energy difference between two mass eigenstates is

$$E_k - E_j \approx \frac{\Delta m_{kj}^2}{2E}, \quad (1.24)$$

where $\Delta m_{kj}^2 = m_k^2 - m_j^2$ is the squared-mass difference. Substituting this into the transition probability and expressing it in terms of the source-detector distance L , since for ultrarelativistic neutrinos the propagation time can be identified with the source-detector distance, $t \simeq L$, in natural units, we get

$$P_{\nu_\alpha\rightarrow\nu_\beta}(L, E) = \sum_{k,j} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* \exp\left(-i \frac{\Delta m_{kj}^2 L}{2E}\right). \quad (1.25)$$

The oscillation phase is therefore defined as

$$\Phi_{kj} = -\frac{\Delta m_{kj}^2 L}{2E}. \quad (1.26)$$

Expanding the expression, the oscillation probability can be written as

$$P_{\nu_\alpha\rightarrow\nu_\beta}(L, E) = \sum_k |U_{\alpha k}|^2 |U_{\beta k}|^2 + 2 \operatorname{Re} \left(\sum_{k>j} U_{\alpha k}^* U_{\beta k} U_{\alpha j} U_{\beta j}^* \exp\left(-2\pi i \frac{L}{L_{kj}^{\text{osc}}}\right) \right), \quad (1.27)$$

with the oscillation length L_{kj}^{osc} being

$$L_{kj}^{\text{osc}} = \frac{4\pi E}{\Delta m_{kj}^2}. \quad (1.28)$$

Looking at Eq. (1.27), we see how oscillation probabilities depend on the mixing matrix elements and the squared-mass differences Δm_{kj}^2 . Finally, Eq. (1.27) represents the general transition probability, and different physical cases can be distinguished by considering either $\alpha = \beta$, corresponding to survival probabilities, or $\alpha \neq \beta$, corresponding to flavor transitions.

The survival probability, where $\alpha = \beta$, takes a simplified form since the quartic terms are real and reads as

$$P_{\nu_\alpha \rightarrow \nu_\alpha}(L, E) = 1 - 4 \sum_{k>j} |U_{\alpha k}|^2 |U_{\alpha j}|^2 \sin^2 \left(\frac{\Delta m_{kj}^2 L}{4E} \right). \quad (1.29)$$

On the other hand, the transition probability for $\alpha \neq \beta$, representing flavor change, k , includes both the oscillatory and interference terms. This distinction is at the basis of appearance and disappearance experiments, respectively. Appearance experiments measure $P(\nu_\alpha \rightarrow \nu_\beta)$, detecting flavor transitions, while disappearance experiments focus on $P(\nu_\alpha \rightarrow \nu_\alpha) = 1 - \sum_{\beta \neq \alpha} P(\nu_\alpha \rightarrow \nu_\beta)$, quantifying the survival of the initial flavor. Finally, neutrino oscillations require the mass eigenstates to remain coherent. If coherence is lost, for example, due to limited detector resolution or uncertainties in the propagation distance, the oscillations average out, leaving a constant transition probability. Moreover, probability conservation ensures unity when summing the probabilities for all possible final states. Oscillation experiments are thus a direct probe of squared-mass differences and mixing angles, while remaining insensitive to absolute neutrino masses and Majorana phases.

1.4.2 Two-Neutrino Mixing Example

Let us consider the illustrative case of two neutrino flavors, ν_e and ν_μ , which provides a useful approximation in many experimental situations. In this scenario, the PMNS mixing matrix reduces to a 2×2 unitary matrix, becoming

$$U = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}, \quad (1.30)$$

where θ is the mixing angle. In this framework, an electron neutrino ν_e at time $t = 0$ is described as a superposition of the mass eigenstates ν_1 and ν_2 through the relation

$$\begin{pmatrix} \nu_e \\ \nu_\mu \end{pmatrix} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix}. \quad (1.31)$$

As the neutrino propagates, the interference between these mass eigenstates gives rise to oscillations between ν_e and ν_μ flavors.

The survival probability of ν_e , $P_{e \rightarrow e}$, is

$$P_{e \rightarrow e} = 1 - \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2 L}{4E}\right), \quad (1.32)$$

where $\Delta m^2 = m_2^2 - m_1^2$ is the squared-mass difference, L is the distance traveled by the neutrino, and E is its energy. Similarly, the transition probability, $P_{e \rightarrow \mu}$, becomes

$$P_{e \rightarrow \mu} = \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2 L}{4E}\right). \quad (1.33)$$

The oscillation length in this case is given by

$$L^{\text{osc}} = \frac{4\pi E}{\Delta m^2}. \quad (1.34)$$

For distances $L \ll L^{\text{osc}}$, the transition probability is very small as the oscillation phase is negligible. For $L \gg L^{\text{osc}}$ instead, oscillations become very rapid and are typically averaged out due to finite experimental resolution. The sensitivity of an experiment to the squared-mass difference Δm^2 ultimately depends on the ratio L/E , where L is the baseline length and E , as previously, is the neutrino energy. Maximal sensitivity occurs when $\Delta m^2 L/2E \sim 1$. For this reason, neutrino oscillation experiments need to be tuned with appropriate values of L/E to probe specific ranges of Δm^2 . Fig.1.2 illustrates this dependence by showing the oscillation probabilities as a function of L/E .

Finally, the two-neutrino effective mixing matrix does not include any phase terms, as there are no CP or T violations in this scenario. As a result, the transition probabilities are symmetric and equal for neutrinos and antineutrinos, as well as for direct and inverted channels

$$P_{\nu_\alpha \rightarrow \nu_\beta}(L, E) = P_{\nu_\beta \rightarrow \nu_\alpha}(L, E) = P_{\bar{\nu}_\alpha \rightarrow \bar{\nu}_\beta}(L, E) = P_{\bar{\nu}_\beta \rightarrow \bar{\nu}_\alpha}(L, E). \quad (1.35)$$

1.5 Experimental Neutrino Anomalies at Short Baseline

The past decades have seen significant advancement in the study of neutrino oscillations that has led to a mostly consistent three-flavor description of solar, atmospheric, reactor, and long-baseline accelerator neutrino data.

Nonetheless, a series of anomalous results has recently emerged in the short baseline experiment data, covering measurements of reactor antineutrinos, source calibrations,

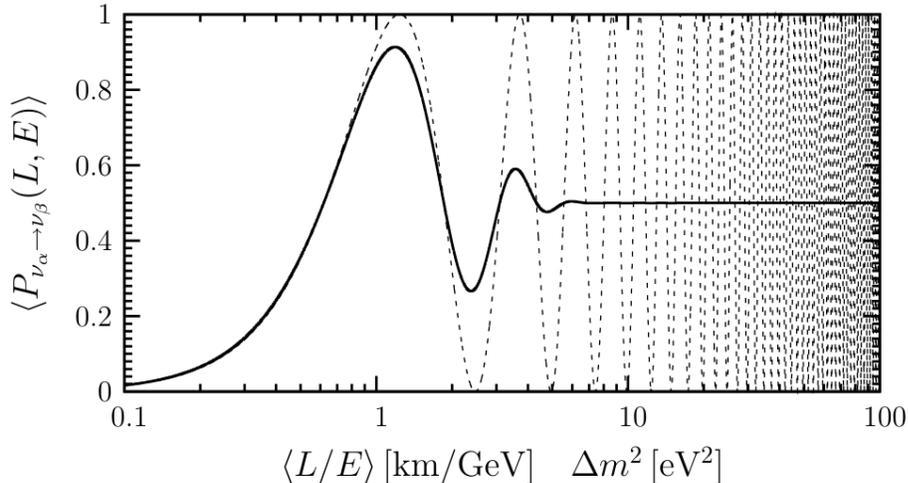


Figure 1.2: Transition probability $P_{\nu_\alpha \rightarrow \nu_\beta}$ for $\sin^2(2\theta) = 1$ as a function of the mean ratio $\langle L/E \rangle$ (in km/GeV), for a fixed value of Δm^2 (in eV^2). The dashed curve represents the unaveraged oscillation probability, while the solid curve shows the probability averaged over a Gaussian L/E distribution. The averaging shows the suppression of oscillatory behavior at large L/E due to finite experimental resolution. [8].

as well as the search for the appearance of known flavors via accelerators. When interpreted in terms of neutrino oscillations, the observed anomalies point to a common mass-squared splitting of order

$$\Delta m^2 \gtrsim 1 \text{ eV}^2, \quad (1.36)$$

significantly larger than the solar and atmospheric mass-squared differences,

$$\Delta m_{21}^2 \simeq 7.4 \times 10^{-5} \text{ eV}^2, \quad |\Delta m_{31}^2| \simeq 2.5 \times 10^{-3} \text{ eV}^2. \quad (1.37)$$

Because only two independent mass-squared differences are possible in the three-flavor scheme, it is necessary to introduce at least one more neutrino mass state participating in neutrino oscillations to affect the mass scale, and, due to precision in electroweak constraints on active neutrinos, it would have to be sterile.

1.5.1 Reactor Antineutrino Anomaly

Electron antineutrinos are continuously produced within nuclear reactors through β decays of neutron-rich fission fragments originating mainly from the isotopes ^{235}U , ^{238}U , ^{239}Pu and ^{241}Pu . These antineutrinos, which typically are in the MeV scale, are emitted isotropically from the reactor core and detected via the inverse beta decay process,

$$\bar{\nu}_e + p \rightarrow e^+ + n. \quad (1.38)$$

This process is characterized by a distinctive delayed-coincidence signature, with a prompt positron signal followed by the delayed capture of the thermalized neutron.

In 2011, a revised β -conversion method increased the predicted $\bar{\nu}_e$ flux [9, 10]. When compared with short-baseline reactor data, this led to a systematic deficit of about 5-6%, known as the Reactor Antineutrino Anomaly. Within sterile neutrino models, this deficit can be interpreted as $\bar{\nu}_e$ disappearance driven by $\Delta m^2 \gtrsim 1 \text{ eV}^2$, and producing rapid oscillations that average out at short baselines.

More recent measurements of the antineutrino yield as a function of reactor fuel composition suggest that part of the deficit may originate from inaccuracies in flux predictions [11]. For this reason, the statistical significance of the anomaly has decreased, although it remains unresolved.

1.5.2 Gallium Anomaly

Additional indications of electron neutrino disappearance at short baselines were reported by radiochemical gallium experiments originally designed to measure the solar neutrino flux. The GALLEX and SAGE experiments employed large volumes of liquid ${}^{71}\text{Ga}$ as a detection medium and observed the CC interaction

$$\nu_e + {}^{71}\text{Ga} \rightarrow {}^{71}\text{Ge} + e^-, \quad (1.39)$$

where the produced ${}^{71}\text{Ge}$ atoms were chemically extracted and counted. Intense artificial neutrino sources, such as ${}^{51}\text{Cr}$ and ${}^{37}\text{Ar}$, were placed inside the detectors for calibration. These sources emit monoenergetic electron neutrinos with energies of a few hundred keV, and measurements were taken at source-detector distances of approximately one meter. Both experiments observed interaction rates that were systematically lower than expectations, with a combined deficit of about 10-15% [12, 13].

Similarly to the reactor-based case, the Gallium Anomaly can be interpreted as ν_e disappearance driven by a mass-squared splitting $\Delta m^2 \gtrsim 1 \text{ eV}^2$.

Recently, the BEST experiment tested the sterile neutrino hypothesis by employing a segmented gallium detector exposed to a ${}^{51}\text{Cr}$ source [14], and observations were consistent with the Gallium Anomaly.

1.5.3 Accelerator-based Appearance Anomalies: LSND and MiniBooNE

Short-baseline accelerator experiments are sensitive to $\nu_\mu \rightarrow \nu_e$ appearance, a transition that is otherwise strongly suppressed within the standard three-flavor framework at short distances. The first evidence for such an anomaly was reported by the LSND experiment, which employed pion decay to produce a neutrino source with a well-characterized $\bar{\nu}_\mu$ flux, a mean energy of approximately 30 MeV, and negligible intrinsic

$\bar{\nu}_e$ contamination [15]. An excess of $\bar{\nu}_e$ events was eventually observed at a baseline of about 30 m, consistent with $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations driven by a mass-squared splitting of order $\Delta m^2 \sim 1 \text{ eV}^2$.

The MiniBooNE experiment was designed to test the LSND result using a different experimental configuration. The neutrino beam had energies of order 1 GeV and a baseline of 540 m, thus having a comparable sensitivity in L/E [16]. MiniBooNE observed an excess of electron-like events at low reconstructed energies in both neutrino and antineutrino running modes. However, the interpretation of this excess is complicated by the limited ability of the employed Cherenkov detector to distinguish electrons from photons and by the absence of a clear oscillation pattern.

LSND and MiniBooNE anomalies point to a similar oscillation scale. Nonetheless, the absence of corresponding signal in ν_μ and ν_e disappearance experiments constrains the possibility of their findings within a minimal sterile neutrino model.

1.5.4 Global Fits and Open Questions

The short-baseline anomalies discussed above have motivated extensions of the standard three-flavor oscillation framework. The minimal scenario capable of introducing an additional oscillation scale at $\Delta m^2 \sim \mathcal{O}(1 \text{ eV}^2)$ is the so-called $3 + 1$ model, in which a single sterile neutrino state is added to the three active neutrinos.

Within this context, the flavor eigenstates are related to four mass eigenstates through a 4×4 mixing matrix. At short baselines, oscillations driven by the solar and atmospheric mass-squared differences are negligible, and the oscillation probabilities are dominated by the largest splitting Δm_{41}^2 . The electron neutrino survival probability can be written as

$$P(\nu_e \rightarrow \nu_e) \simeq 1 - 4|U_{e4}|^2(1 - |U_{e4}|^2) \sin^2\left(\frac{\Delta m_{41}^2 L}{4E}\right), \quad (1.40)$$

while the appearance probability in the muon-to-electron channel is given by

$$P(\nu_\mu \rightarrow \nu_e) \simeq 4|U_{e4}|^2|U_{\mu4}|^2 \sin^2\left(\frac{\Delta m_{41}^2 L}{4E}\right). \quad (1.41)$$

Ultimately, disappearance experiments constrain $|U_{e4}|^2$, while appearance searches probe the product $|U_{e4}|^2|U_{\mu4}|^2$. Any short-baseline $\nu_\mu \rightarrow \nu_e$ appearance signal, therefore, implies nonzero disappearance effects in both the ν_e and ν_μ channels. Global analyses show that the mixing parameters required to reproduce the LSND and MiniBooNE excesses predict disappearance signals that are not observed in other experiments, leading to significant tension within the minimal $3 + 1$ framework. When appearance and disappearance data are combined, no single parameter region provides a fully consistent description of all available results, as illustrated in Fig. 1.3. While these findings do not exclude sterile neutrinos, they highlight the limitations of the simplest extension of

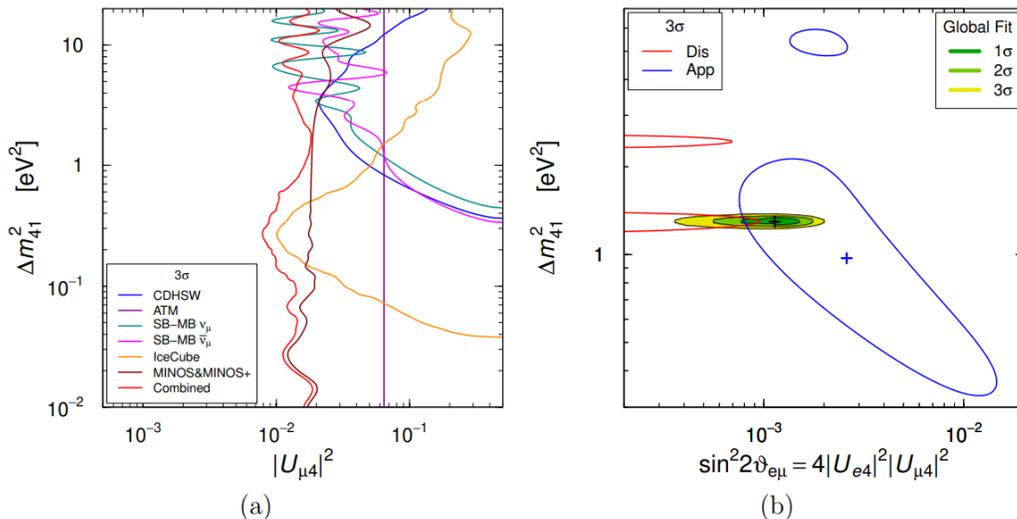


Figure 1.3: Short-baseline ν_μ disappearance results and global fits in the $(\sin^2 2\theta_{\mu e}, \Delta m_{41}^2)$ plane. (a) Exclusion limits from ν_μ disappearance experiments. (b) Comparison of appearance-only (App) and disappearance (Dis) constraints with the combined global fit. In panel (a), and for the Dis curves in panel (b), the region to the right is excluded at 3σ , while the App curves in panel (b) indicate the 3σ allowed regions [17].

the oscillation framework and motivate further experimental tests and more complex theoretical scenarios.

1.6 The Fermilab Short-Baseline Neutrino Program

Since a consistent picture has not yet emerged from the experimental point of view, the Fermi National Accelerator Laboratory (Fermilab) established the Short-Baseline Neutrino (SBN) program to provide a decisive test of the sterile-neutrino hypothesis. The strategy relies on several detectors placed at different baselines along the same beamline, the Booster Neutrino Beam (BNB), all employing the same liquid argon time projection chamber (LArTPC) technology. This latter has fine-grained tracking and precise calorimetric capabilities that directly address the dominant experimental ambiguities affecting the MiniBooNE low-energy excess.

The SBN configuration enables a near-far comparison in which the near detector(s) constrain the unoscillated neutrino spectrum, while the far detector probes possible oscillation-induced distortions. Using the same detection technique at both locations also suppresses common systematic effects related to flux modeling, neutrino-argon cross sections, and detector response. Along the BNB line, the near detector is the Short-Baseline Near Detector (SBND), while ICARUS-T600 serves as the far detector

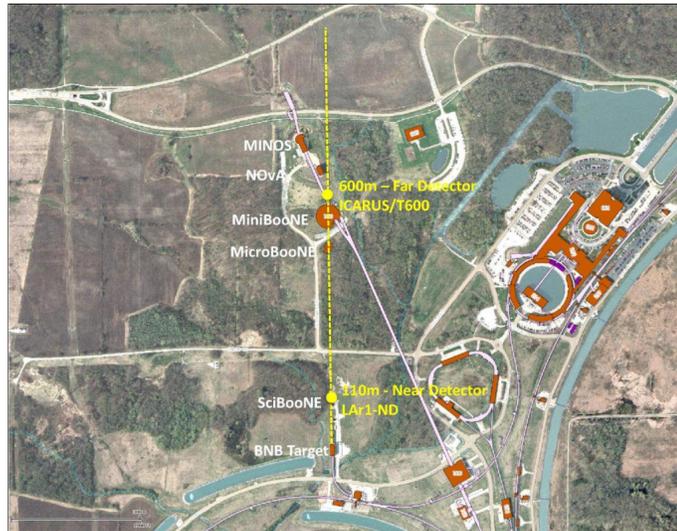


Figure 1.4: Schematic layout of the SBN detectors along the Booster Neutrino Beam (BNB), indicating the near (SBND), intermediate (MicroBooNE), and far (ICARUS-T600) locations [18].

Table 1.1: Main characteristics of the SBN LArTPC detectors along the BNB [18].

Detector	Distance from BNB target [m]	Total LAr mass [t]	Active LAr mass [t]
SBND	110	220	112
MicroBooNE	470	170	89
ICARUS-T600	600	760	476

(FD-SBN). MicroBooNE, located at an intermediate baseline, has completed data-taking and is no longer operating. As shown in Fig. 1.4, the three LArTPC detectors are at baselines of approximately 110 m (SBND), 470 m (MicroBooNE), and 600 m (ICARUS-T600) from the BNB target. Table 1.1 summarizes the corresponding total and active liquid-argon masses for each detector.

The oscillation searches are carried out in both disappearance and appearance channels, in particular $\nu_\mu \rightarrow \nu_\mu$ and $\nu_\mu \rightarrow \nu_e$, with a projected sensitivity that covers the parameter space suggested by the remaining anomalies at the $> 5\sigma$ level, as illustrated in Fig. 1.5. Beyond oscillations, SBN also supports a broader physics program, including precision measurements of ν -Ar interactions in an energy range relevant for long-baseline experiments, as well as searches for sub-GeV dark matter and other beyond-the-Standard-Model signatures.

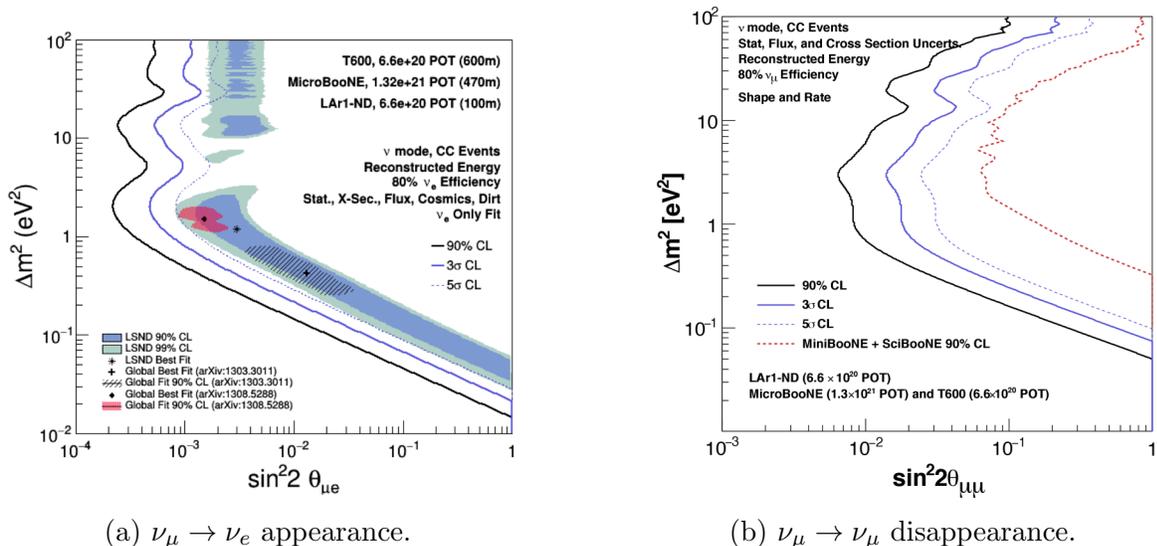


Figure 1.5: Projected SBN sensitivity to sterile-neutrino oscillations in appearance and disappearance channels in the $(\Delta m^2, \sin^2 2\theta)$ parameter space. The contours show the expected confidence-level reach for the SBN configuration. From Ref. [18].

Booster Neutrino Beam (BNB). The SBN detectors are exposed to the Booster Neutrino Beam (BNB), where protons of kinetic energy of 8 GeV from the Fermilab Booster strike a beryllium target to produce secondary hadrons, mainly pions. A single toroidal magnetic horn (operated with $\mathcal{O}(10^2)$ kA) current pulses, synchronized with proton delivery) focuses charged mesons and selects the beam mode by choosing the horn polarity. The focused secondaries then enter an air-filled decay region (about 50 m long, radius ~ 0.91 m), where most pions decay into muons and ν_μ ; hadrons that do not decay are absorbed at the end of the tunnel by a concrete/steel absorber.

The beam is delivered in short spills of $\sim 1.6 \mu\text{s}$, typically with $\sim 5 \times 10^{12}$ protons on target per spill, which is advantageous for selecting in-time activity in coincidence with the spill window. In neutrino mode, the resulting flux is dominated by ν_μ (roughly at the $\sim 90\%$ level), while the intrinsic ν_e contamination is small (about 0.5% below ~ 1.5 GeV), largely originating from the $\pi \rightarrow \mu \rightarrow e$ decay chain. The predicted flux spectra at the SBN detector locations are shown in Fig. 1.6, as obtained from beamline simulations constrained by external hadron-production measurements.

Neutrinos at the Main Injector (NuMI) In addition to the BNB exposure, ICARUS also records neutrino interactions from the Neutrinos at the Main Injector (NuMI) beam. NuMI is produced by steering 120 GeV protons onto a graphite target, generating secondary hadrons that are focused and sign-selected by two magnetic horns and then allowed to decay in a long decay region. The remaining hadrons are stopped in an absorber, while a downstream rock volume ranges out the residual muon component. For ICARUS, this provides a complementary neutrino sample with a different energy

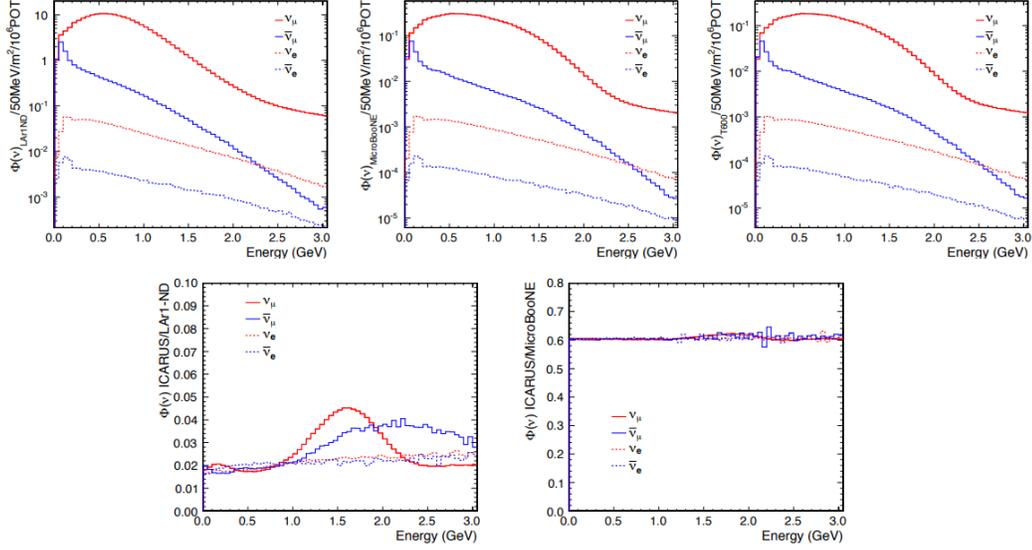


Figure 1.6: Predicted Booster Neutrino Beam (BNB) flux at the three SBN detector locations. Here “LAr1-ND” denotes the near detector, i.e. the Short-Baseline Near Detector (SBND); the other sites are MicroBooNE and ICARUS-T600. The top row shows the energy spectra for ν_μ , $\bar{\nu}_\mu$, ν_e , and $\bar{\nu}_e$; the bottom row shows the corresponding flux ratios between sites, highlighting the baseline dependence and the change in spectral shape due to the different viewing angles of the decay region. From Ref. [18].

scale and flavor content with respect to BNB, and it is available in both neutrino- and antineutrino-focusing configurations (Forward Horn Current, FHC, and Reverse Horn Current, RHC).

ICARUS is located off the NuMI beam axis by about 6° , which modifies the observed spectrum and increases the relative electron-neutrino contribution compared to an on-axis exposure. At this off-axis location, ICARUS collects a large sample of ν_e interactions in the $\sim 0\text{--}3$ GeV energy range. In terms of composition, the NuMI sample is still dominated by ν_μ from π/K decays, with a few-percent wrong-sign $\bar{\nu}_\mu$ component and a smaller intrinsic ν_e contribution from subdominant decay chains; the off-axis enhancement of ν_e makes this sample particularly useful for validating electromagnetic-shower reconstruction and for analyses involving ν_e final states. The expected flavor-separated fluxes at the ICARUS location are shown in Fig. 1.7 for FHC and RHC running configurations.

Thanks to the simultaneous exposure to BNB and to the off-axis NuMI flux, SBN will collect very large neutrino interaction samples on argon. As an order-of-magnitude estimate, the expected statistics are $\mathcal{O}(10^6)$ events per year in the near detector from the sub-GeV BNB spectrum, and $\mathcal{O}(10^5)$ events per year in ICARUS from the higher-energy off-axis NuMI sample.

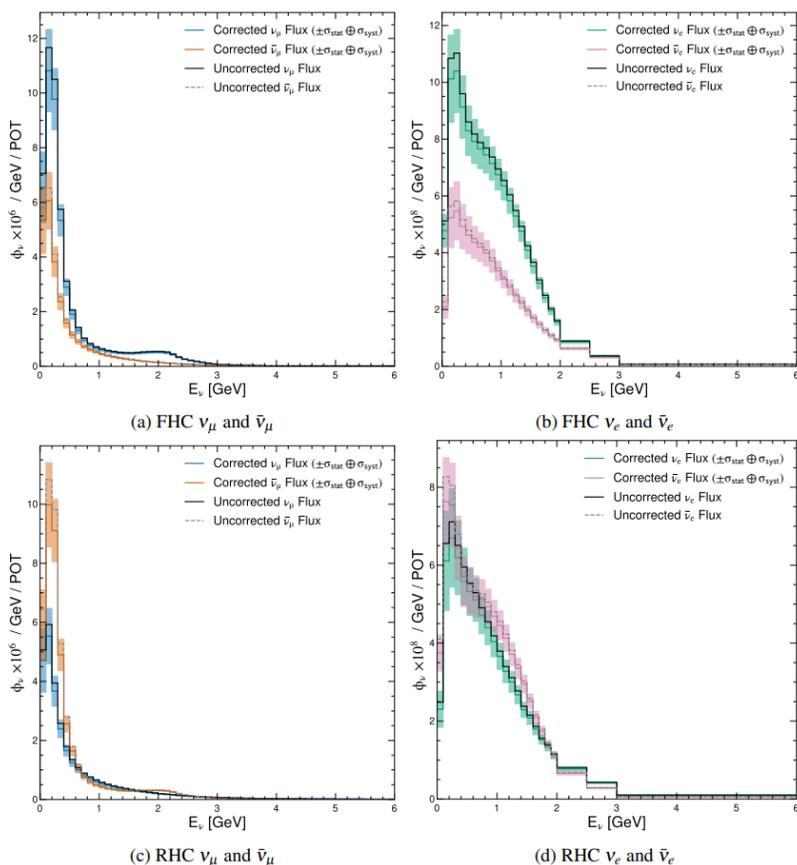


Figure 1.7: Predicted NuMI neutrino flux at the ICARUS location for forward horn current (FHC, ν -focusing) and reverse horn current (RHC, $\bar{\nu}$ -focusing) running. The spectra are shown separately for ν_μ , $\bar{\nu}_\mu$, ν_e , and $\bar{\nu}_e$. From Ref. [19].

MicroBooNE A MicroBooNE analysis was recently published exploiting data from two independent accelerator neutrino beams, the Booster Neutrino Beam (BNB) and the off-axis NuMI beam, characterized by different intrinsic electron-flavor contents [20].

This configuration allows a simultaneous and coherent probe of $\nu_\mu \rightarrow \nu_e$ appearance and ν_e disappearance, breaking the degeneracy that affects single-beam short-baseline searches within the $3 + 1$ sterile neutrino framework.

No evidence for non-standard flavor oscillations is observed. The results exclude, at the 95% confidence level, the region of the $(\Delta m_{41}^2, \sin^2 2\theta_{\mu e})$ parameter space capable of explaining the LSND and MiniBooNE anomalies through a single light sterile neutrino, and significantly constrain the parameter space favored by gallium experiments. The data are consistent with the three-flavor oscillation hypothesis, with no indication of sterile-induced oscillations. These results show that the LSND and MiniBooNE excesses cannot be simultaneously accommodated within the minimal $3+1$ oscillation framework once appearance and disappearance channels are tested with sufficient experimental

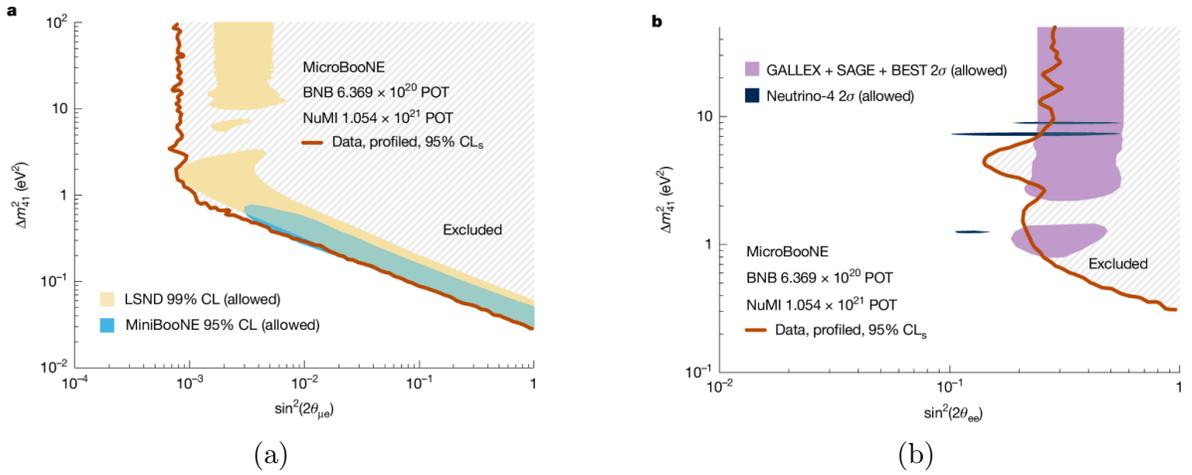


Figure 1.8: Constraints on sterile neutrino oscillation parameters from the MicroBooNE two-beam analysis [20]. The red curves show the 95% CL_s exclusion limits obtained from MicroBooNE data. Panel (a) shows constraints in the $(\Delta m_{41}^2, \sin^2 2\theta_{\mu e})$ plane. The yellow and blue regions correspond to the allowed parameter space by LSND and MiniBooNE. Panel (b) displays constraints in the $(\Delta m_{41}^2, \sin^2 2\theta_{ee})$ plane. The purple and blue regions are the favored parameter space for gallium experiments and the Neutrino-4 result. The MicroBooNE data rule out regions to the right of the exclusion curves.

control. These considerations further motivate the SBN program at Fermilab to finally provide a conclusive test of the eV-scale sterile neutrino scenarios.

Chapter 2

LArTPCs: the ICARUS detector and R&D perspectives

2.1 Detection in LArTPCs and their Operation

Having outlined the experimental motivation and the near-far strategy of the SBN program, let us now zoom in on the detector technology that enables such a measurement. The LArTPC is an imaging particle detector that uses ultra-pure liquid argon as both the interaction target and the tracking medium [21]. Liquid argon is a particularly convenient detection medium because it is relatively dense ($\rho \simeq 1.4 \text{ g/cm}^3$) and easy to source at an industrial scale, since argon makes up about $\sim 0.93\%$ of dry air. Moreover, it is chemically inert and can be operated as a stable cryogenic liquid at $\sim 87 \text{ K}$ (around 1 atm), thus making it well suited for large-mass detectors. Useful reference values for liquid argon are collected in Table 2.1.

Building on early proposals, LArTPCs have evolved into a well-established technology for precision neutrino physics. In particular, the ICARUS (Imaging Cosmic And Rare Underground Signals) program, which was originally proposed by C. Rubbia [22], was fundamental in developing and demonstrating this technology at a large scale. At a glance, LArTPCs' detection capabilities provide a fully active and homogeneous volume with continuous sensitivity, millimetre-scale three-dimensional imaging, and calorimetric information from local energy deposition dE/dx . As discussed in Ref. [23], these features enable bubble-chamber-like event reconstruction with intrinsic background rejection based on topology and dE/dx .

Table 2.1: Relevant LAr properties for LArTPC operation [24].

Quantity	Condition	Value
Density (liquid)	at 87.3 K, 1.013 bar	1.395 g/cm ³
Argon abundance in dry air		0.934 %
Boiling point	at 1.013 bar	87.3 K
Ionization work function		23.6 eV
Scintillation yield (no field)	at 128 nm	$\sim 4.0 \times 10^4$ photons/MeV
Scintillation decay times	fast / slow components	~ 6 ns / ~ 1.6 μ s
Electron drift velocity	$E = 500$ V/cm, $T \simeq 89$ K	1.55 mm/ μ s
Longitudinal diffusion coefficient	$T \simeq 89$ K	4.8 cm ² /s

2.1.1 Energy deposition, ionization and scintillation

A LArTPC operates with an electric field applied across the active liquid argon volume, which is constrained by a cathode plane on one side and an anode plane on the other. A high-voltage bias between them generates a nearly uniform drift field. As a charged particle traverses liquid argon, it loses energy mainly through inelastic interactions with the medium, producing both excitation and ionization of argon atoms. The deposited energy is therefore converted into two measurable signals: free ionization electrons and scintillation photons (see Fig. 2.1). These channels are produced simultaneously and are at the basis of LArTPC detection.

Ionization Ionization creates electron-ion pairs along the particle trajectory, with the electrons toward the anode. The average energy required to free one electron in liquid argon is the ionization work function, $W_{ion} = 23.6$ eV (Table 2.1), corresponding to an ideal yield of $\sim 4.2 \times 10^4$ electrons per 1 MeV of deposited energy, before accounting for recombination and attachment. The spatial distribution of the ionization charge closely follows the local energy deposition profile, dE/dx . At the anode, the drifting charge typically crosses one or more induction wire planes before being finally collected on a collection plane. The induction planes are biased to be nearly transparent, so that electrons pass through while inducing a transient current on the wires, resulting in waveforms that are therefore typically bipolar. On the other hand, the collection plane absorbs the electrons on the wires, producing a unipolar signal whose time integral is proportional to the collected charge.

Scintillation At the same time, a fraction of the deposited energy excites argon atoms to Ar^* states. Through collisions with surrounding ground-state atoms, the excitation energy is rapidly converted into excited dimers (Ar_2^* , or excimers) via self-trapping processes. Excimers are also produced through recombination, where an Ar^+

ion can capture a free electron and, after interactions with neighbouring atoms, form an excited dimer (see Fig. 2.2). The Ar_2^* excimer exists in two spin configurations, the singlet ($^1\Sigma_u^+$) and triplet ($^3\Sigma_u^+$) states. Their radiative de-excitation produces vacuum-ultraviolet (VUV) scintillation light with characteristic wavelength $\lambda \simeq 128$ nm (the LAr M-band). The singlet component is prompt, with a decay time of about 2 ns to 6 ns (fast component), while the triplet component is much slower, with a decay time typically in the range 1100 ns to 1600 ns (slow component). The relative fraction of fast and slow scintillation depends on the local ionization density of the track and, indirectly, on recombination.

Recombination In liquid argon, the ionization and scintillation yields are coupled by recombination. At higher ionization density (large dE/dx), a larger fraction of the initial electron-ion pairs recombines, reducing the free charge yield while increasing the light yield. Increasing the drift field suppresses recombination by separating charges more efficiently, therefore leading to a larger collected charge and a reduced scintillation yield. As a consequence, the relationship between deposited energy and collected charge is not exactly linear and must be corrected for recombination effects, especially for highly ionizing particles.

This effect can be described by a semi-empirical parametrization:

$$Q = \frac{A Q_0}{1 + (k/E) \langle dE/dx \rangle}, \quad (2.1)$$

where Q is the collected charge, $Q_0 = (1/W_{\text{ion}}) dE/dx$ is the initial ionization charge, E is the electric field, and A and k are parameters determined from data. No recombination model provides a universal description over the full range of ionization densities, in particular in regions of high dE/dx and near the end of particle tracks, where effects such as δ -ray production become relevant. For this reason, the conversion from dQ/dx to dE/dx in modern LArTPCs must rely on empirical calibration using reference samples, such as minimum-ionizing particles, stopping tracks, and cosmic rays.

The collected charge dQ/dx provides calorimetric and tracking information, while scintillation light is used to estimate the absolute event time t_0 , thus contributing to triggering and background rejection.

2.1.2 Electron drift and transport in liquid argon

The free electrons produced by ionization are transported toward the readout planes under the action of a uniform electric field, whose uniformity is essential, given that spatial distortions from field non-uniformities directly affect reconstruction through biases. For example, when considering detectors operated at the surface or near, a source of field distortions is the accumulation of positive ions produced by intense activity of cosmic rays. Because ions drift orders of magnitude more slowly than electrons, due to their much lower mobility, they can build up in the liquid and generate a space charge

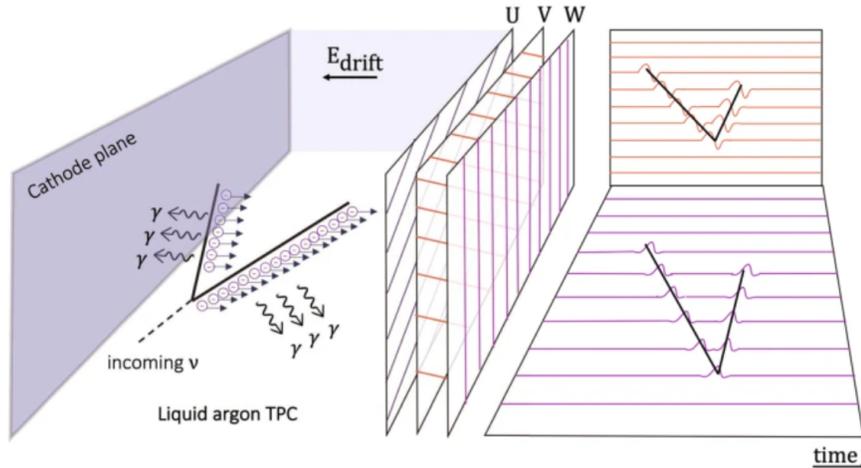


Figure 2.1: Schematic operation of a LArTPC. Charged particles produced in a neutrino interaction ionize and excite the argon. Ionization electrons drift in a uniform electric field toward the anode wire planes, where signals are recorded in induction (bipolar signals) and collection (unipolar signals) views. From Ref. [25].

that perturbs the nominally uniform drift field. These bias the reconstructed positions and therefore need to be mitigated through calibration and field-mapping [23].

Typical drift field values in LArTPCs range from 200 V/cm to 500 V/cm, and they are chosen as a compromise between charge collection efficiency, recombination suppression, and high-voltage stability. Under these conditions, electrons drift with an average velocity v_D that depends on the applied electric field and the temperature of the liquid argon. At a temperature of approximately 87 K, corresponding to the boiling point of liquid argon at atmospheric pressure, the drift velocity reaches values of about 1.6 mm/ μ s for an electric field of 500 V/cm [27]. The drift time t_{drift} for electrons produced at a distance L from the anode is therefore given by

$$t_{\text{drift}} = \frac{L}{v_D}. \quad (2.2)$$

The measured arrival time of the charge, once referred to the interaction time t_0 given by scintillation light, is used to reconstruct the position along the drift direction.

Diffusion Effects During their transport, drifting electrons undergo diffusion due to thermal motion and interactions with the medium, leading to a gradual spatial spreading of the charge cloud both longitudinally, along the drift direction, and transversely. The longitudinal diffusion can be expressed in terms of a diffusion coefficient D_L , with the variance of the charge distribution increasing as

$$\sigma_L^2 = 2D_L t_{\text{drift}}. \quad (2.3)$$

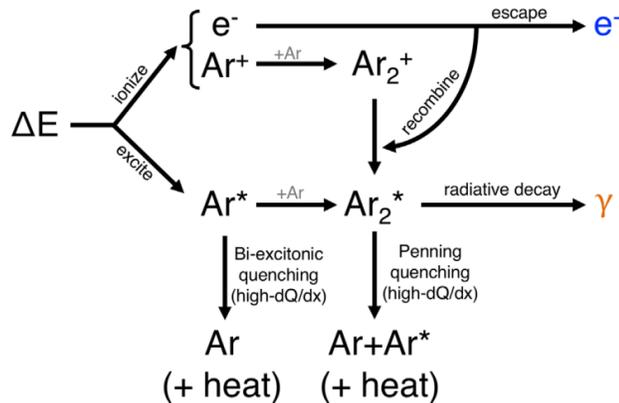


Figure 2.2: Schematic formation of scintillation in liquid argon. Energy deposition ΔE produces both ionization ($Ar^+ + e^-$) and excitation (Ar^*). Excited dimers Ar_2^* are formed either directly from Ar^* (self-trapping) or via recombination of ionization electrons, and de-excite radiatively emitting VUV photons. At high ionization density, additional non-radiative quenching channels (e.g. bi-excitonic processes) reduce the light yield. From Ref. [26].

A similar relation holds for the transverse diffusion, described by a coefficient D_T .

The transverse diffusion is related to the mobility $\mu(E)$ through the Einstein-Smoluchowski relation,

$$D_T = \frac{kT}{e} \mu(E), \quad (2.4)$$

while the longitudinal diffusion has an additional contribution from the field dependence of the mobility,

$$D_L \simeq \frac{kT}{e} \left(\mu(E) + |E| \frac{\partial \mu(E)}{\partial |E|} \right). \quad (2.5)$$

This naturally implies that $D_L < D_T$ in liquid argon, since μ decreases with increasing $|E|$. For typical operating conditions ($E \simeq 500$ V/cm), electron diffusion leads to a transverse charge spread of order $\sigma_T \sim 1$ mm after a 1 m drift [23]. Similarly, the longitudinal diffusion coefficient in liquid argon is of order $0.4 \text{ mm}^2/\text{ms}$ [24], resulting in a longitudinal spread of comparable magnitude over drift times of order ~ 1 ms. Both effects are smaller than the typical wire pitch in large LArTPCs, yet they define an intrinsic limit on the achievable spatial resolution, especially for long drift distances.

2.1.3 Impurities, electron lifetime and charge attenuation

Drifting electrons can be lost during transport through liquid argon due to attachment to electronegative impurities. The most relevant contaminants for LArTPC operation are oxygen and water, which capture free electrons and lead to an attenuation of the ionization charge as a function of drift time [28, 29].

The impact of impurities on charge transport is expressed in terms of the electron lifetime τ_e , defined as the mean survival time of drifting electrons before attachment to impurities. The electron lifetime can be expressed as

$$\tau_e = \frac{1}{k \cdot [\text{impurity}]}, \quad (2.6)$$

where k is the electron attachment rate constant and $[\text{impurity}]$ is the concentration of electronegative contaminants in the liquid argon.

Often, the impurity concentration is expressed as an oxygen-equivalent level, since O_2 and H_2O dominate electron attachment in many operating conditions. A widely used rule is that an oxygen-equivalent contamination of $\mathcal{O}(0.1)$ ppb corresponds to electron lifetimes of a few milliseconds, which is sufficient to keep charge attenuation small over meter-scale drifts.

For an ionization charge drifting for a time t_{drift} , the ratio between the collected charge at the anode and the produced charge can be expressed as

$$\frac{Q_{\text{anode}}}{Q_0} = \exp\left(-\frac{t_{\text{drift}}}{\tau_e}\right), \quad (2.7)$$

where Q_0 represents the charge produced at the point of ionization, and Q_{anode} is the charge collected at the readout plane. τ_e needs to be effectively monitored using both purity monitors and track-based methods. Purity monitors are small drift cells that generate a known electron charge (e.g., via photoemission) and measure the fraction collected after a fixed drift time, giving a real-time estimate of the lifetime (see ICARUS' example in Fig. 2.3). At the same time, τ_e can be extracted from the attenuation of through-going tracks, such as cosmic muons, by fitting the measured charge as a function of drift time across the full active volume. Reliable charge collection requires the electron lifetime to be significantly longer than the maximum drift time, $\tau_e \gg t_{\text{drift}}$, such that charge losses due to attachment remain negligible over the full drift distance. A common design target is to keep the maximum attenuation at the cathode below the $\sim 10\%$ level, which is equivalent to requiring $\tau_e \gtrsim 10 t_{\text{drift,max}}$. During early ICARUS-T600 operation, a maximum electron lifetime of ~ 1.8 ms was achieved, representing a significant milestone for large-scale long-drift LArTPCs at the time [30]. In a later operation, ICARUS itself demonstrated higher lifetimes, eventually exceeding 15 ms [31].

In addition to charge attenuation, impurities can also affect the scintillation light yield through quenching processes. Nitrogen (N_2), in particular, is known to efficiently quench argon scintillation by introducing non-radiative de-excitation channels for the excimer states [32]. This effect mainly suppresses the slow scintillation component associated with the long-lived triplet state, whose extended lifetime increases the probability of collisional quenching.

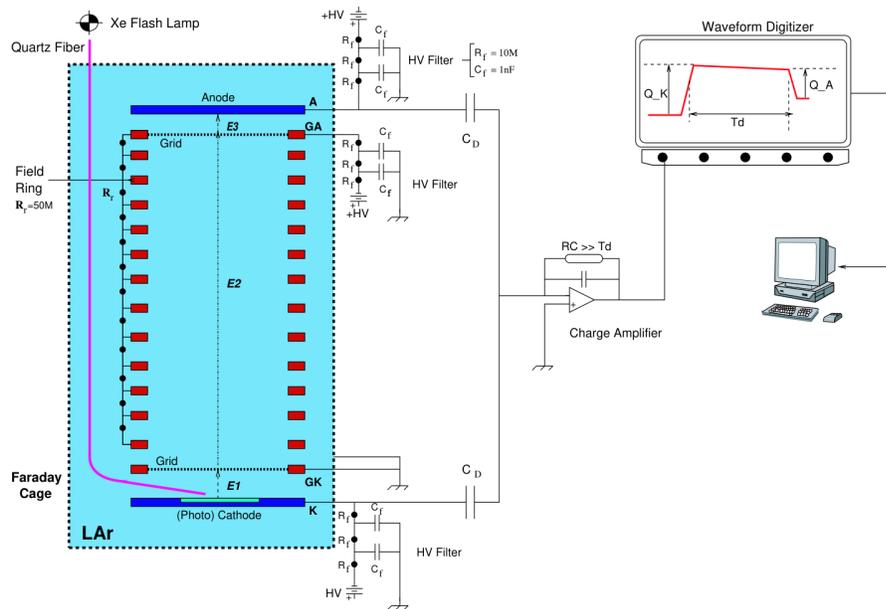


Figure 2.3: Scheme of a purity monitor for the electron lifetime measurement [24].

For typical LArTPC operation, efficient charge collection requires ultra-pure liquid argon with oxygen-equivalent impurity concentrations at or below the 0.1 ppb level. Nitrogen, on the other hand, does not significantly impact charge transport. Nonetheless, its concentration must be kept below the part-per-million level, typically below 1 ppm, to preserve the scintillation light yield and timing performance [32, 33].

These conditions are sustained by the cryostats designed for LArTPCs, which typically employ vacuum-insulated structures with double walls for large detectors (of the order of hundreds of tons of LAr) or foam-insulated ones for even larger mass applications. The initial phase usually involves purging the cryostat with argon to eliminate air, followed by purification techniques that target residual oxygen and water using materials such as molecular sieves or activated copper on alumina. The argon must also be continuously filtered, and purity monitors must be used.

Notes on signal formation and digitization The signals recorded in a LArTPC result from the motion of ionization electrons in the electric field near the anode and from the response of the readout electronics. The ionization charge is processed by low-noise front-end electronics, which in modern LArTPCs are typically operated at cryogenic temperature to reduce capacitance and electronic noise. The front-end amplifiers convert the collected charge into a voltage signal with a configurable gain, typically expressed in units of mV/fC. For example, a minimum-ionizing particle deposits about 2.1 MeV/cm in liquid argon, corresponding to an ionization charge of the order of a few femtocoulombs per wire after recombination and attenuation effects.

The shaped analog signals are digitized by analog-to-digital converters (ADCs) at a fixed sampling rate, producing discrete waveform samples expressed in ADC counts.

The measured waveform on each channel reflects the convolution of the ionization charge distribution with the detector response. In the induction planes, the signal is dominated by the motion of electrons in the weighting field and is therefore bipolar with approximately zero net area, while in the collection plane, the electrons are absorbed and the signal is unipolar, with an area proportional to the collected charge. As a result, signal processing is characterized by a response-function deconvolution to recover a charge estimate that is proportional to the deposited ionization [23].

The digitized waveforms are the lowest-level input to the signal processing and reconstruction chain.

Finally, the intrinsic spatial granularity of the charge readout is determined by the anode segmentation, e.g., wire pitch or pixel size. Considering a segmented readout pitch p and assuming a uniform charge distribution within one single element, the corresponding geometric position resolution is $\sigma_x \simeq p/\sqrt{12}$. Along the drift direction, the position resolution is determined by the sampling period Δt and the drift velocity v_D , such that the discretization scale is $\Delta x \simeq v_D \Delta t$. This means that the achievable resolution is further limited by charge diffusion and electronic shaping. For instance, at $E = 500 \text{ V/cm}$ ($v_D \simeq 1.6 \text{ mm}/\mu\text{s}$), a sampling period of $\Delta t = 0.5 \mu\text{s}$ corresponds to $\Delta x \sim 0.8 \text{ mm}$.

2.2 The ICARUS Detector

The ICARUS-T600 detector is a large-scale LArTPC which was originally conceived within the ICARUS program as a demonstrator of long-drift LArTPCs at an argon mass of hundreds of tons. It was first commissioned at CERN and then operated underground at the Laboratori Nazionali del Gran Sasso (LNGS) during the CERN Neutrinos to Gran Sasso (CNGS) program (2010-2012), where it collected both neutrino interactions and cosmic-ray data. After the end of CNGS, ICARUS underwent major refurbishment and was repurposed as the far detector of the Fermilab SBN program. This refurbishment campaign mainly needed to address the challenging near-surface deployment of the detector, with an intense cosmic-ray activity both in-time (during the beam spill) and out-of-time (during the full drift interval). For this reason, an upgraded photon-detection system and an external Cosmic Ray Tagger (CRT), together with a passive overburden, were added to control cosmogenic backgrounds [34]. From the point of view of the detector concept, ICARUS is a modular LArTPC whose components trace the classical LArTPC architecture previously seen.

2.2.1 The detector architecture

ICARUS-T600 consists of two identical, adjacent cryostat modules, historically denoted as T300 half-modules, that share a common warm vessel, as is illustrated in Fig. 2.4. Each cryostat module holds two TPCs, and the total LAr mass is about 760 t, with an active mass of 476 t in the SBN configuration. This corresponds to an overall active volume of $\sim 3.4 \times 10^2 \text{ m}^3$. The modular layout was originally driven by construction, transport, and installation constraints at LNGS, while still reaching a large target mass [24]. A schematic summary of the main detector parameters, mainly concerning TPC geometry and operating conditions, is provided in Table 2.2.

Table 2.2: Selected ICARUS-T600 parameters (TPC geometry and nominal operating conditions). Values are collected from Ref. [24].

Quantity	Value
Number of cryostat modules	2 (T300 half-modules)
Number of TPCs	4 (two per half-module, sharing a central cathode)
Sensitive volume per TPC	$\sim 85 \text{ m}^3$
Total / active LAr mass	$\sim 760 \text{ t} / \sim 476 \text{ t}$
Cathode-to-collection distance	$\sim 1.50 \text{ m}$
Maximum drift length (cold)	1482 mm
Nominal drift field	500 V/cm
Nominal cathode HV	$\sim 75 \text{ kV}$
Maximum drift time (nominal)	$\sim 1 \text{ ms}$
Wire planes per TPC	3 (2 induction + 1 collection)
Wire pitch	3 mm
Wire plane spacing	3 mm
Wire orientations	$0^\circ, +60^\circ, -60^\circ$
Wire diameter	150 μm
Total number of wires	53 248
Wires per TPC	13312
Photon detectors (SBN)	360 PMTs (90 per TPC)

Cryostat and internal mechanical structure

Each T300 cryostat is an aluminium cold vessel surrounded by thermal insulation and placed within a warm mechanical structure. The cold vessel consists of $\sim 150 \text{ mm}$ -thick aluminium panels, with external aluminium skins and an aluminium honeycomb core. In one half-module, the LN_2 cooling pipes are directly inside the honeycomb structure, and they provide the cooling circuit for the cryostat walls.

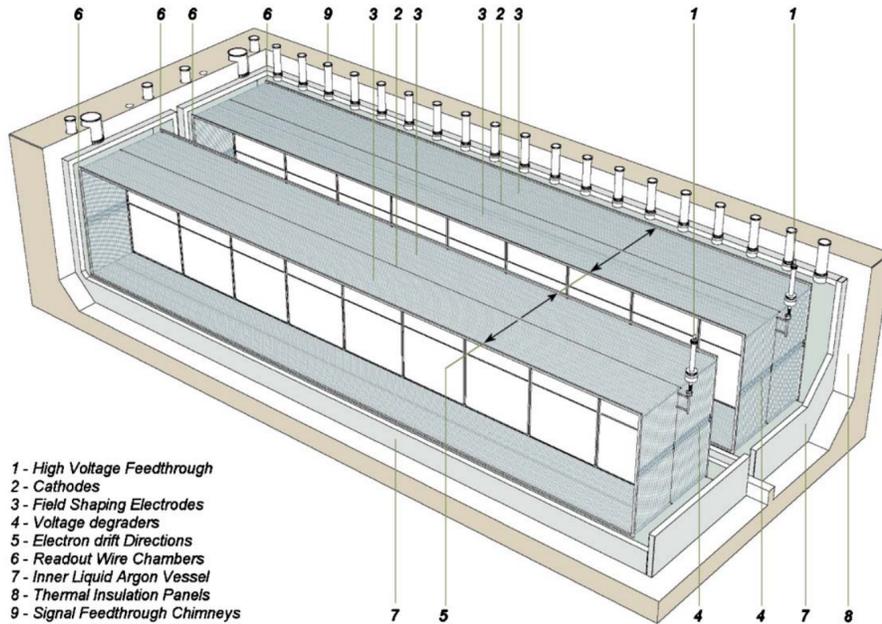


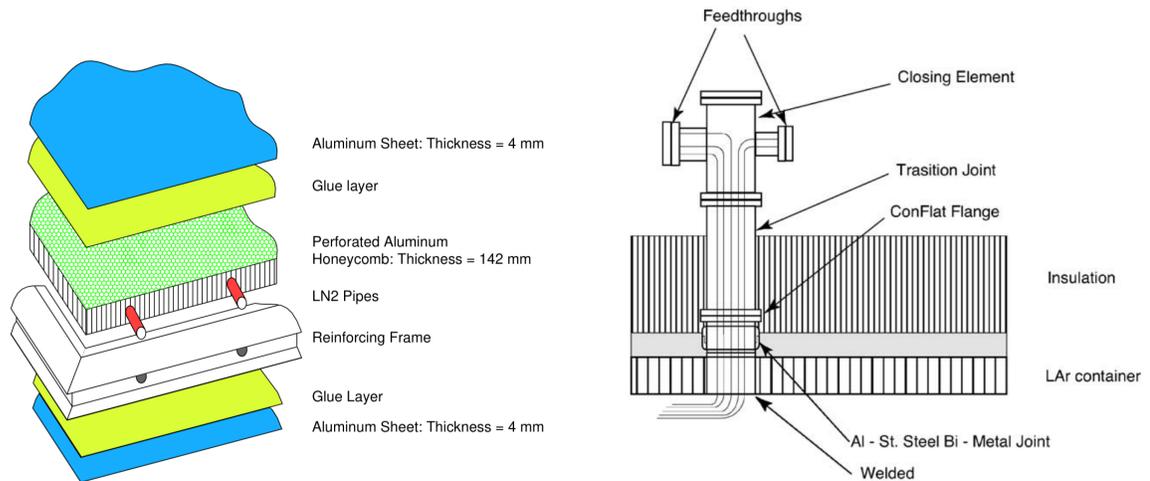
Figure 2.4: Schematic illustration of the ICARUS T600 detector with its main components [24].

The cold vessel dimensions are approximately $4.2 \times 3.9 \times 19.9 \text{ m}^3$ externally. The corresponding inner dimensions of the LAr volume instead are

$$L_I \simeq 19.6 \text{ m}, \quad W_I \simeq 3.6 \text{ m}, \quad H_I \simeq 3.9 \text{ m},$$

giving an inner volume of $V_I \simeq 275 \text{ m}^3$ per half-module. Alongside ensuring an ultra-pure LAr, the cryostat integrates the interfaces needed for operation, namely cryogenic plumbing, recirculation and purification lines, high-voltage delivery, signal feedthroughs for the TPC wires, and instrumentation (e.g., purity monitors, temperature and level probes). In particular, cabling and services are accessed through the top flanges, which host vacuum-tight feedthroughs used to route TPC wire signals to the external front-end electronics and to connect the internal instrumentation and services, with the corresponding routing passing through stainless-steel chimneys crossing the insulation. Different from more recent LArTPCs, ICARUS' front-end electronics are placed outside of the cryostats. This makes them more accessible, though at the cost of increased input capacitance from long signal cables. Details about thermal insulation and feedthrough are shown in Fig. 2.5.

Inside each cryostat, a stainless-steel mechanical frame supports the TPC wire planes, the central cathode, and the field-shaping electrodes. The internal structure is designed to preserve the wire geometry over time, accounting for the thermal contraction between the aluminum vessel and the stainless-steel structure.



(a) Aluminium sandwich panel concept with honeycomb core and integrated LN₂ cooling pipes (for one T300 half-module).

(b) Example of a cryostat penetration: feedthroughs through the insulation and Al and stainless-steel transition joint.

Figure 2.5: Cryostat construction details in ICARUS, from Ref. [24].

TPC layout, drift volumes and field cage

Each T300 half-module contains two identical TPCs (see Fig. 2.6), denoted as left and right drift regions. These are separated by a common central cathode plane. Therefore, the detector has four total independent drift volumes (one per TPC), each with its own high-voltage distribution chain for the field-shaping system. Since ionization electrons drift from the cathode toward the readout anode planes, each TPC images exactly one half of the argon volume. The cathode-to-anode distance also defines the maximum drift path, nominally ~ 1.5 m, and the effective drift length at LAr temperature is 1482 mm due to thermal contraction. At $E_D = 500$ V/cm, this corresponds to a maximum drift time of about ~ 0.95 ms, which is consistent with the expected electron drift velocity in LAr at nominal conditions. The drift field is obtained with a cathode bias of about 75 kV, and the same drift length and field are used in the SBN configuration, corresponding to a maximum drift time around 0.96 ms [34].

Field uniformity is achieved through a field cage made of field-shaping electrodes connected by a resistive divider, which generates a quasi-linear potential gradient from cathode to anode. In ICARUS, the electrodes are arranged in a so-called "race-track" configuration [24]. The nominal drift field of 500 V/cm over the ~ 1.5 m drift distance. The cathode itself is composed of perforated stainless-steel panels with an optical transparency of about 58%. This design, besides satisfying mechanical constraints, allows a significant fraction of scintillation light to pass between the two drift regions, improving the overall light-collection efficiency.

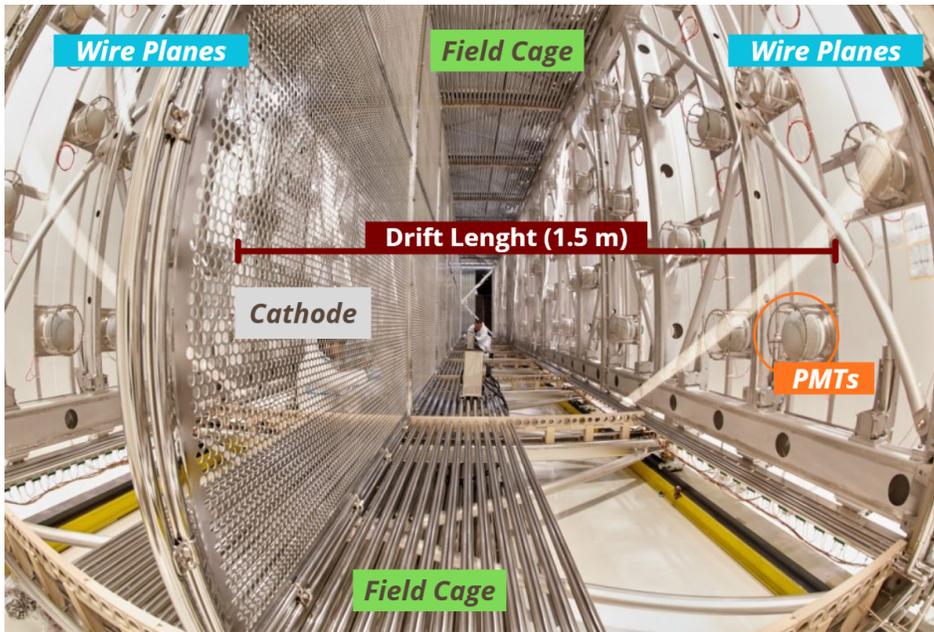


Figure 2.6: An inside look at the T300 half-module.

Anode wire planes and readout geometry

The anode of each TPC consists of three parallel, vertical wire planes, with a common active surface of about $17.95 \times 3.16\text{m}^2$, and with a spacing of 3 mm in between planes. The three planes give two induction views and one collection view, with wire orientations defined with respect to the horizontal direction as follows:

- (a) Induction-1 with 0° wires (horizontal),
- (b) Induction-2 with $+60^\circ$ wires,
- (c) Collection with -60° wires.

The wire pitch is 3 mm for all planes (normal to the wire direction), and the wires are made of stainless steel with a diameter of about $150\ \mu\text{m}$. The wire lengths depend on the plane and on the position within the frame: for the 0° plane the wires are up to $\sim 9.42\text{m}$ long, while for the $\pm 60^\circ$ planes the typical length is $\sim 3.77\text{m}$, with shorter wires at the corners down to $\sim 0.49\text{m}$. Mechanically, the wires are tensioned to limit gravitational and electrostatic pull. Typical values are $\sim 12\text{N}$ (and $\sim 5\text{N}$ for the longest horizontal wires) [24].

In the four TPCs of the detector, there are a total of $N_{\text{wires}} = 53,248$ readout wires, with each TPC chamber hosting 13,312 wires. The wire capacitance is of the order of 20-21 pF/m for the three planes. Typical bias values for ICARUS are approximately -220V , 0V , and $+280\text{V}$ for the Induction-1, Induction-2, and Collection planes, respectively. Each plane is a 2D projection in the (wire coordinate, drift time) space.

Finally, the readout electronics operate with 12-bit digitization at 2.5 MHz, covering the full ~ 1 ms drift window. Signal calibration is performed via injected test pulses, enabling the conversion from ADC counts to collected charge, as discussed in Sec. 2.1.3.

Photon detection and timing Scintillation light at $\lambda \simeq 128$ nm gives the absolute time reference of the interaction, t_0 , which is essential to convert wire time into an absolute drift coordinate and thus reconstruct events in three dimensions.

In the SBN configuration, ICARUS operates near the surface, where the short BNB spill (1.6 μ s) must be identified within a drift window of order ~ 1 ms, and an efficient optical triggering is necessary to suppress out-of-spill cosmic activity. For this purpose, ICARUS employs 360 eight-inch PMTs (Hamamatsu R5912-MOD), with 90 units per TPC, installed behind the anode planes. Since PMTs are not directly sensitive to 128 nm VUV light, wavelength shifting is implemented using TPB-coated acrylic plates [34].

2.2.2 From hardware to performance

Recalling the detector principles discussed so far, we can now follow the full chain that leads from the raw signal to particle identification in a LArTPC, and examine how this is realized in the case of ICARUS.

The information carried by the drifting ionization electrons is first recorded as digitized waveforms on the anode wires (i). These waveforms represent the time evolution of the induced or collected signal on each channel. After signal processing and hit finding (ii), the waveforms are reduced to localized charge deposits, referred to as hits, which constitute the fundamental reconstructed objects of the detector.

By combining the wire coordinate with the drift time (iii), each hit is assigned a three-dimensional position inside the active volume. In the collection view, the signal is unipolar and its time integral is proportional to the charge reaching the wire. Once liquid-argon specific effects such as recombination and electron attachment are taken into account (iv), the measured charge can be converted into deposited energy.

Each reconstructed hit gives both spatial and calorimetric information within a small effective voxel of the detector, enabling tracking, dE/dx reconstruction, and ultimately particle identification.

More specifically, the intrinsic spatial granularity of ICARUS is determined by the readout segmentation and the drift sampling. The wire pitch is 3 mm, corresponding to a transverse position quantization of approximately

$$\sigma_{\perp} \simeq \frac{3 \text{ mm}}{\sqrt{12}} \approx 0.9 \text{ mm}. \quad (2.8)$$

Along the drift direction, the sampling frequency of 2.5 MHz yields a spatial step of

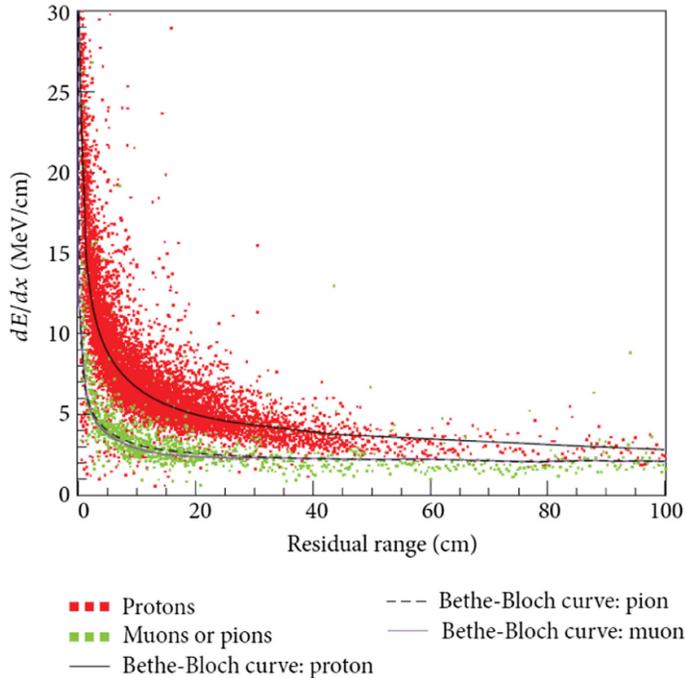


Figure 2.7: dE/dx as a function of residual range in ICARUS-T600 (LNGS) for identified protons and for muons/pions. The residual range is defined as the distance from the measurement point to the reconstructed track endpoint. From Ref. [23].

about 0.6 mm at nominal drift velocity, but due to electron diffusion during drift, the effective spatial resolution in all coordinates remains at the millimeter scale.

Tracking and particle identification. For sufficiently long tracks (typically longer than ~ 30 cm), the reconstructed direction achieves angular resolutions of order a few degrees [23]. When particles are fully contained, the dE/dx as a function of residual range is a powerful particle identification tool.

For example, protons exhibit the characteristic Bragg peak near the end of their range, while muons and pions are close to minimum-ionizing behaviour over most of their trajectory. As seen in Fig. 2.7, an efficient separation is clearly visible in the measured dE/dx versus residual range distributions, allowing proton-muon/pion discrimination for contained tracks. The same method can be extended to other hadrons if a sufficient track length is available.

Energy reconstruction. Energy reconstruction in ICARUS ultimately depends on the topology and containment of the final state.

In the case of fully contained charged particles, the resolution mainly depends on how precisely the track length and the stopping point are reconstructed, as the energy can

be reconstructed from the range. For stopping muons, a momentum resolution of a few percent is achieved. The precision can be further improved by including calorimetric information along the track and accounting for effects such as δ -ray emission and radiative energy losses.

For electromagnetic and hadronic showers, calorimetry information is required. The shower energy is obtained by summing the calibrated collected charge over the reconstructed cluster. For electromagnetic showers at GeV energies, the energy resolution can be parametrized as

$$\frac{\sigma_E}{E} \simeq 1\% \oplus \frac{3\%}{\sqrt{E(\text{GeV})}}. \quad (2.9)$$

At lower energies, below the critical energy of liquid argon (~ 30 MeV), studies of Michel electrons yield

$$\frac{\sigma(E)}{E} \simeq \frac{0.11}{\sqrt{E(\text{MeV})}} \oplus 0.02, \quad (2.10)$$

once bremsstrahlung photons are properly accounted for.

Hadronic showers are intrinsically more complex due to nuclear effects and invisible energy components; resolutions of order $\sim 30\%/\sqrt{E}$ are typically expected.

Momentum reconstruction via multiple Coulomb scattering. When muon tracks are not fully contained, range-based reconstruction is not applicable. In the absence of a magnetic field, ICARUS estimates the momentum using multiple Coulomb scattering (MCS).

The angular deviations accumulated along a segment of length L are inversely related to the particle momentum p , and can be approximated by

$$\psi \simeq \frac{13.6 \text{ MeV}/c}{\beta p} \sqrt{\frac{L}{X_0}} \left(1 + 0.038 \ln \frac{L}{X_0} \right), \quad (2.11)$$

where X_0 is the radiation length. By comparing the observed angular deviations with expectations, an estimate can be found using a likelihood-based approach. For tracks of order a few meters, typical resolutions are at the level of $\sim 15\text{-}25\%$ in the GeV range.

Overall, the performance of ICARUS follows directly from its 3D reconstruction and calorimetry, which enable accurate tracking, dE/dx -based particle identification, and energy reconstruction, with MCS providing momentum information when tracks are not fully contained.

2.3 Motivation for R&D in LArTPC-based Neutrino Experiments

LArTPC technology has reached a high level of maturity and has proven to be effective for neutrino detection, as first established by the ICARUS experiment. Despite its fine granularity and calorimetric capabilities, several challenges remain. While some limitations can be mitigated through improved reconstruction algorithms, others are intrinsic to the detector concept and ultimately affect particle identification and pattern recognition. A notable example is the separation of electrons and photons, which relies on local dE/dx measurements at the start of electromagnetic showers. This handle is intrinsically fragile, since photons do not ionize directly but may convert into an e^+e^- pair within a few millimeters of the interaction vertex. When this conversion occurs on a length scale comparable to the intrinsic mm-scale granularity set by readout segmentation and diffusion, the two prongs can overlap and mimic a single electron ionization.

Beyond reconstruction challenges, LArTPCs face intrinsic limitations that directly impact the physics reach of neutrino experiments, as they are typically operated without a magnetic field. In the absence of curvature-based momentum measurements, the charge sign of final-state leptons cannot be determined. Momentum reconstruction relies on range, restricted to fully contained tracks, or on MCS, whose resolution is limited and depends on track length. As a consequence, kinematic constraints are weakened and opposite-charge leptonic final states (e.g. e^- versus e^+ , or μ^- versus μ^+) cannot be discriminated.

Physics motivation for a magnetized LArTPC Introducing a magnetic field in a LArTPC would directly address these intrinsic limitations by enabling curvature-based momentum reconstruction and charge-sign identification. This additional information would strengthen kinematic reconstruction and improve pattern recognition in complex topologies.

Charge identification would be essential in appearance analyses, where photon-induced backgrounds such as $\pi^0 \rightarrow \gamma\gamma$ can mimic electromagnetic activity, and in configurations where neutrino and antineutrino samples must be separated due to wrong-sign contamination, such as the separation of ν_e and $\bar{\nu}_e$. Charge-sign information is also important for studies involving ν_τ interactions, where identifying the sign of visible decay products provides an additional handle to constrain the parent lepton charge and suppress background contamination.

A magnetized LArTPC would also extend the discovery potential in high-intensity neutrino beams by enabling sensitivity to signatures that are otherwise difficult to access. An example is trident-like final states,

$$\nu + A \rightarrow \nu + A + \ell^+ + \ell^-,$$

which arise in the SM and in several BSM scenarios. In such events, charge identification is needed for the opposite-sign lepton pair, while curvature-based reconstruction improves the separation of close-by tracks, particularly in boosted configurations with small opening angles.

Overall, these considerations motivate R&D efforts aimed at understanding the advantages and limitations of operating LArTPCs in a magnetic field. It is in this context that the ArCS (Argon detector with Charge Separation) prototype was developed.

2.3.1 The ArCS Experiment

The ArCS experiment was conceived as an R&D program focused on assessing the operation of a magnetized LArTPC under controlled conditions, namely a test beam environment. ArCS aims to evaluate reconstruction performance metrics that are otherwise inaccessible in standard LArTPCs. The primary goals are

- (a) the demonstration of charge-sign identification for charged leptons, in particular electrons and positrons;
- (b) the reconstruction of particle momentum from track curvature;
- (c) the minimum magnetic-field strength required to achieve useful charge discrimination and curvature-based momentum reconstruction.

To achieve these goals, ArCS deliberately builds on an existing LArTPC design rather than building a prototype from scratch. The experiment reuses the LArIAT (Liquid Argon In A Testbeam) detector and its readout chain, which has had a successful operation as a test beam LArTPC. The ArCS detector will be operated at the Fermilab Test Beam Facility (FTBF), where a well-characterized secondary beam impinges on a Cu target, thus generating a tertiary beam. Additionally, the FTBF houses multiple external beamline instruments, such as wire chambers, to aid particle identification. Finally, the LArTPC is placed inside the Jolly Green Giant (JGG) dipole magnet, which provides a vertical magnetic field of up to ~ 0.7 T. The field is approximately orthogonal to the electric drift field, limiting Lorentz-angle effects on drifting electrons while inducing measurable curvature of charged particle tracks along the drift direction, where spatial resolution is optimal.

Operational timeline The ArCS program has progressed through several stages of hardware preparation and integration. In 2022-2023, a dedicated cryogenic plant was designed, and a new cryostat was constructed. In 2024, the TPC was refurbished, the cold and warm electronics were fully tested, and the detector was installed inside the cryostat. Subsequent vacuum tests were performed to validate the cryogenic system, and the detector is currently installed inside the JGG magnet at the FTBF. The final commissioning phase, scheduled for 2026, will include full cryogenic installation and

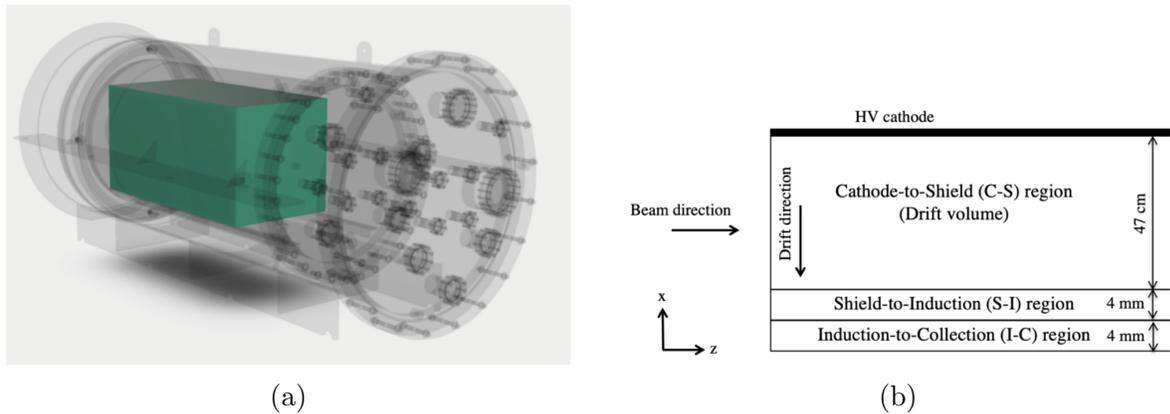


Figure 2.8: (a) Modified ArCS cryostat design (the green box is the TPC). (b) Schematic of the three drift regions: cathode-shield (C-S), shield-induction (S-I), and induction-collection (I-C). Adapted from Ref. [35].

initial data-taking with cosmic rays. Once the FTBF beam becomes available, the test beam run will take place.

LArIAT TPC layout ArCS reuses the LArIAT TPC, which features an active volume of $47\text{ cm} \times 40\text{ cm} \times 90\text{ cm}$, enclosed by a cathode and a rectangular field cage. The field cage is constructed from copper-clad G10 panels and consists of 1 cm-wide horizontal copper strips separated by 1 cm gaps. The strips wrap around the four walls of the drift volume and are connected through a resistor chain acting as a voltage divider, with an effective resistance of $250\text{ M}\Omega$ between adjacent strips. This configuration ensures a uniform electric field across the active volume, which is essential to preserve the geometry of the ionization track during electron drift.

As seen in Fig. 2.8, the TPC comprises three drift regions: the main drift volume between the high-voltage cathode and the shield plane (i), the region between the shield and induction planes (ii), and the region between the induction and collection planes (iii). The wire planes are arranged such that the shield-plane wires are vertical, while the induction and collection wires are oriented at $\pm 60^\circ$ with respect to the horizontal. The wire pitch is 4 mm, and a total of 480 channels from the induction and collection planes are instrumented and digitized.

Front-end and readout electronics The front-end electronics chain (see Fig. 2.9) closely follows the LArIAT design and consists of 480 analog channels. The first amplification stage is provided by ten 48-channel cold motherboards (CMBs), mounted directly on the TPC frame. Each CMB hosts three custom 16-channel LArASIC chips operating at cryogenic temperature, which amplify charge close to the signal source, minimizing noise and optimizing the signal-to-noise ratio. The LArASICs are typically operated at a gain of 25 mV/fC , corresponding to an output of approximately 90 mV

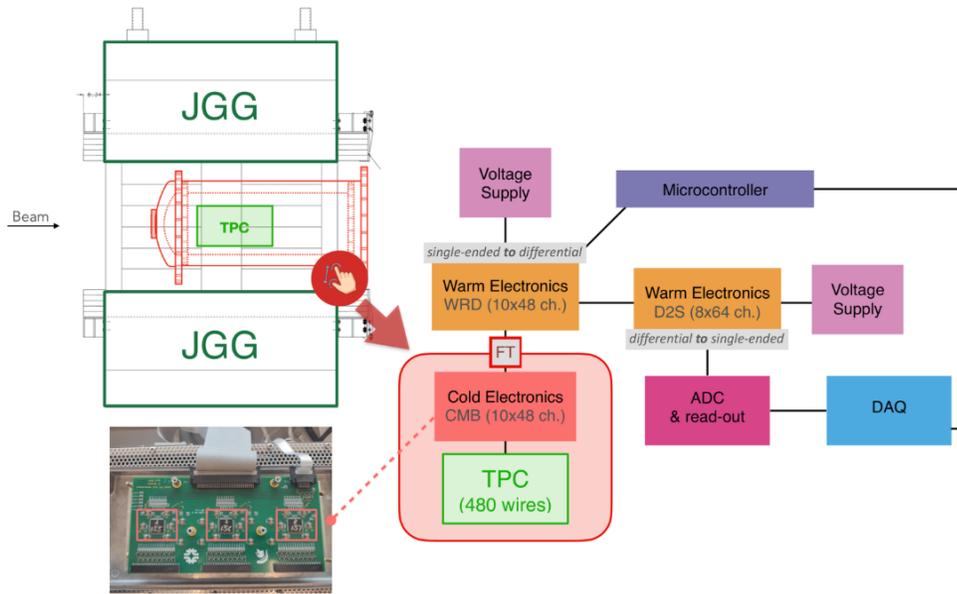


Figure 2.9: A walkthrough of the ArCS LArIAT-heredited TPC and electronics.

for a 3.5 fC input signal, which is characteristic of a minimum-ionizing particle in liquid argon.

Signals are routed through dedicated feedthroughs to the warm electronics stage. Then, ten 48-channel warm receiver and driver (WRD) boards convert single-ended signals into differential signals, reducing common-mode noise and accommodating ground reference differences. The differential signals are then processed by differential-to-single-ended (D2S) boards, each hosting up to 64 channels, before being digitized by CAEN V1740 modules.

As shown in Fig.2.10, the expected FTBF beam spectrum is within the $\mathcal{O}(100 \text{ MeV})$ region, thus overlapping with the typical energy range of primary leptons produced in BNB and off-axis NuMI interactions. In this regime, the simulated collection plane views in Fig. 2.11 clearly show that the curvature induced by a vertical magnetic field allows e^- and e^+ tracks to be distinguished at the topological level, providing a first indication of the achievable charge discrimination capability.

Regarding ArCS, I have recently contributed to DAQ-related activities in preparation for beam operation, as described in the Appendix A.

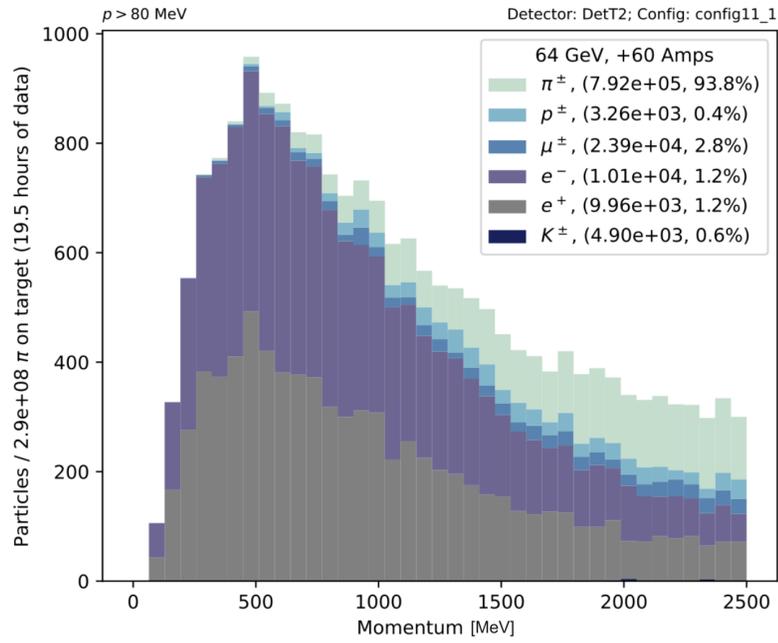


Figure 2.10: Momentum spectrum expected at FTBF. The spectrum was simulated with G4Beamline; electrons and positrons populate the $\mathcal{O}(100 \text{ MeV})$ region.

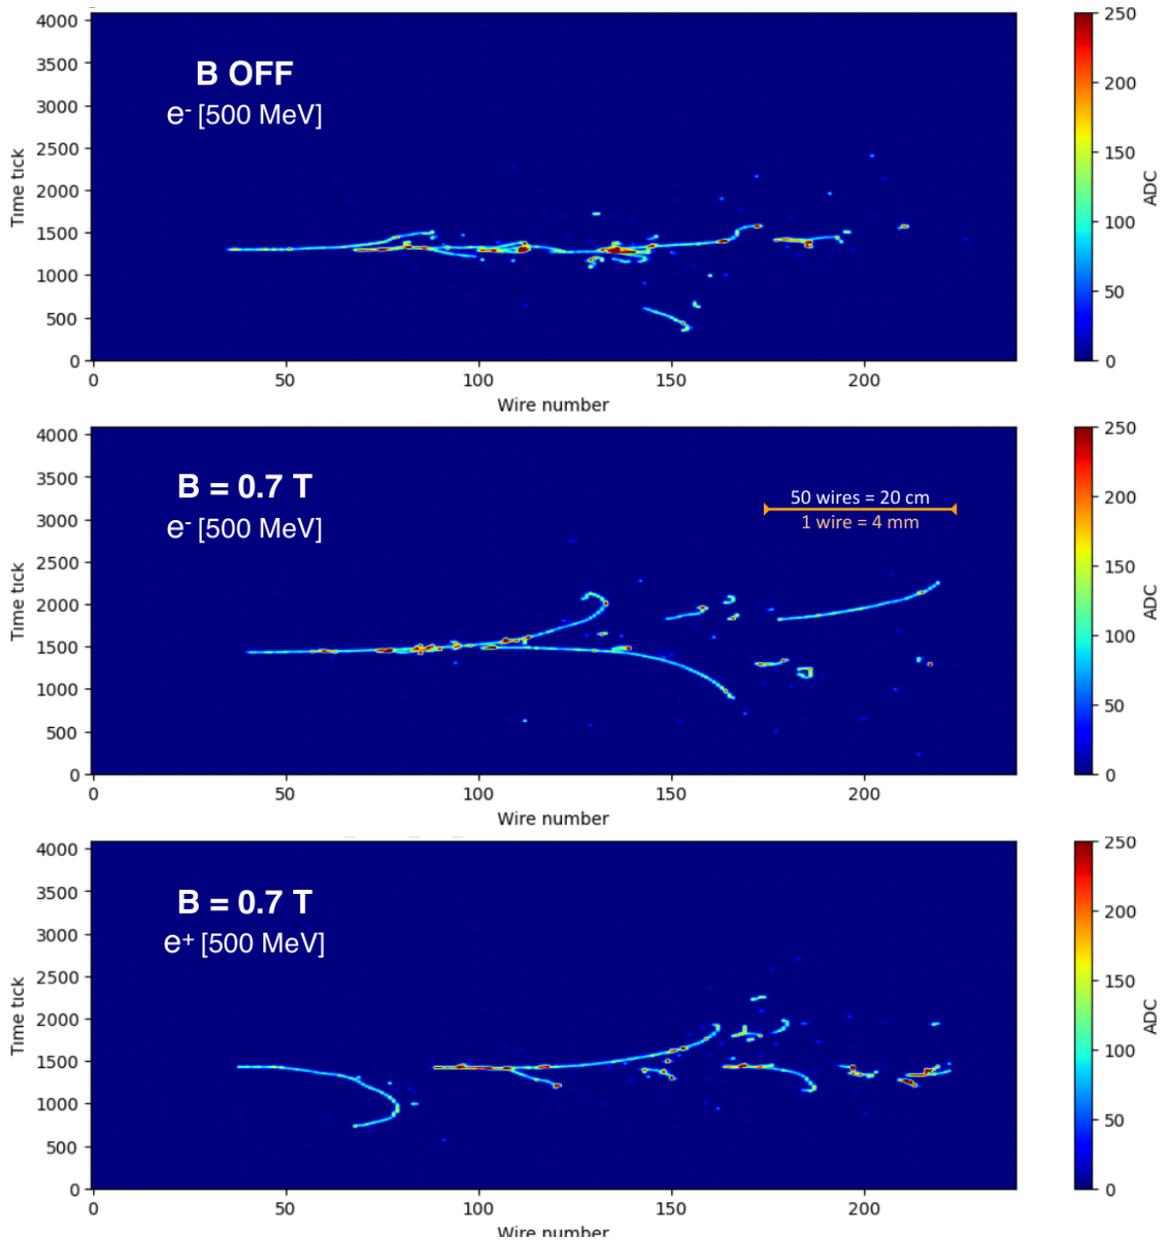


Figure 2.11: Collection plane views from a particle gun simulation using the LArIAT configuration. Top: e^- at 500 MeV with $B = 0$. Middle: e^- at 500 MeV with $B = 0.7$ T. Bottom: e^+ at 500 MeV with $B = 0.7$ T. The magnetic field is included at the Geant4 propagation step as a static vertical field.

Chapter 3

Event Reconstruction in the ICARUS LArTPC

LArTPC event reconstruction begins with the raw digitized waveforms recorded by the readout electronics. From these, the reconstruction chain aims to identify physically meaningful objects to be used in the actual physics analysis, such as three-dimensional interaction vertices, particle tracks, and electromagnetic showers. In ICARUS, the reconstruction workflow can be organized into three main stages:

- (a) **signal processing**, which converts raw waveforms into calibrated charge measurements;
- (b) **hit reconstruction** and three-dimensional **clustering**;
- (c) **pattern recognition** and **particle reconstruction**.

These steps are often grouped into two broader levels. The first is commonly referred to as Stage 0 or upstream reconstruction, and it includes signal processing and hit reconstruction (2D hits and 3D space points). The second level, instead, is Stage 1, or downstream reconstruction, which includes pattern recognition, particle reconstruction, and hierarchical interaction classification. A scheme of the Pandora reconstruction chain is shown in Fig. 3.2.

3.1 Overview of the ICARUS Reconstruction Chain

3.1.1 Signal Processing

The digitized waveforms are the convolution of three main contributions: (i) the ionization signal induced by drifting electrons $S_{\text{ion}}(t)$, (ii) the detector field response $R_{\text{field}}(t)$, and (iii) the electronics shaping response $R_{\text{elec}}(t)$. The measured waveform $M(t)$ can therefore be expressed as

$$M(t) = (S_{\text{ion}} * R_{\text{field}} * R_{\text{elec}})(t) + n(t), \quad (3.1)$$

where $n(t)$ represents noise contributions.

Defining the total detector response function as

$$R_{\text{tot}}(t) = (R_{\text{field}} * R_{\text{elec}})(t), \quad (3.2)$$

the measured signal can be written as

$$M(t) = (S_{\text{ion}} * R_{\text{tot}})(t) + n(t). \quad (3.3)$$

Before deconvolution, correlated noise components that arise due to common electronics and power supply pickup and that are shared among groups of adjacent wires are identified and removed through filtering algorithms to prevent noise amplification and improve the signal-to-noise (SNR) ratio. Then, a one-dimensional (1D) deconvolution is applied independently to each wire plane to correct for both the detector and electronics response. In frequency space, the convolution effectively becomes a product,

$$M(\omega) = S_{\text{ion}}(\omega) R_{\text{tot}}(\omega), \quad (3.4)$$

and the deconvolution corresponds to

$$S_{\text{ion}}(\omega) = \frac{M(\omega)}{R_{\text{tot}}(\omega)} F(\omega), \quad (3.5)$$

where $F(\omega)$ is a filter suppressing noise at frequencies where $R_{\text{tot}}(\omega)$ becomes small. The final output is a deconvolved waveform proportional to the ionization charge collected on each wire.

Then, a threshold-based Region of Interest (ROI) algorithm is applied to identify portions of the waveform that are likely to contain a signal. The thresholds are intentionally kept moderately low to maximize efficiency. Within each ROI, Gaussian fits are finally performed to extract discrete 2D hits. The integral of the fitted Gaussian is an estimate of the number of ionization electrons associated with the localized energy deposition. In this way, a 2D hit is defined as a localized charge deposition on a single wire and identified as a peak in the deconvolved waveform within a given ROI, thus representing the projection of a three-dimensional energy deposition onto one wire plane. Each reconstructed hit is passed down to the downstream reconstruction and is thereby characterized by: (i) the wire coordinate, (ii) the peak time (related to drift coordinate), (iii) the integrated charge (proportional to deposited energy), and (iv) the width of the Gaussian fit.

3.1.2 Three-Dimensional Clustering: Cluster3D

Since the hit-finding stage applies low thresholds, fake hits due to noise fluctuations may be reconstructed. To suppress these, ICARUS uses the `Cluster3D` algorithm, whose goal is that to match 2D hits from different wire planes into geometrically consistent 3D space points. A valid match must have (i) geometrical consistency, meaning that the corresponding wires intersect in space, and (ii) time compatibility across planes, within a tolerance that accounts for hit time width.

Each candidate combination is assigned a score that depends on time differences and hit extensions, and only hits associated with reconstructed 3D space points are propagated to the next stage.

3.1.3 Pattern Recognition with Pandora

The main pattern recognition framework employed in ICARUS is **Pandora**, a modular software toolkit widely used in liquid argon experiments and designed to reconstruct complete interaction topologies and hierarchies.

Pandora consists of a large set of configurable and independent algorithms, each targeting specific reconstruction tasks or topological features (e.g., clustering, vertex finding, track and shower reconstruction). The sequential execution of these modules progressively builds a consistent and hierarchical description of the event.

Two complementary reconstruction paths are executed:

- **Pandora Cosmic**, optimized for cosmic-ray muon reconstruction;
- **Pandora Nu**, optimized for neutrino interactions originating from the beam.

A first pass of Pandora Cosmic identifies clear cosmic-ray candidates based on topology, containment, direction with respect to the vertical axis, and timing compatibility with the beam spill window. In this way, hits that are associated with cosmic muons are removed before neutrino reconstruction proceeds.

Event Slicing After cosmic removal, the remaining hits are grouped into *slices*, which are collections of 3D clusters hypothesized to originate from a single physical interaction.

Multiple interactions may overlap within the same readout window, and the slicing stage aims to separate these into independent reconstruction units. Ideally, each slice contains all and only the hits belonging to the primary interaction and its full cascade of secondary particles and decays.

Flash Matching Since ICARUS is a surface detector, optical information is essential in distinguishing beam-induced neutrino interactions from out-of-time cosmic activity.

PMTs record scintillation light in the form of optical hits, which are grouped into optical flashes. For each reconstructed slice, a barycenter in space and time is computed from the associated TPC activity. This barycenter is compared to the position and timing of optical flashes.

A compatibility score is then computed based on spatial and temporal consistency, such that slices that match an in-beam flash are favored as neutrino candidates.

Neutrino Vertex Reconstruction Within each slice, Pandora Nu attempts to identify the 3D neutrino interaction vertex. Candidate vertices are generated from cluster endpoints across different planes. Unphysical candidates are removed, and the remaining ones are evaluated using a Boosted Decision Tree (BDT).

The BDT combines several classes of variables, such as calorimetric variables, to ensure consistency between cluster energy and vertex position, and beam direction variables, such that vertices consistent with the beam axis are favored. Global event quantities such as the total number of hits, the number of clusters, and longitudinal event extension are also considered. The candidate with the highest BDT score is selected as the reconstructed neutrino vertex.

Track and Shower Reconstruction After the interaction vertex has been identified within a slice, Pandora proceeds to reconstruct the individual particles associated with that vertex.

The fundamental reconstructed object in Pandora is the **Particle Flow Particle** (PFP), which represents a reconstructed particle candidate containing:

- (a) the collection of 3D hits associated with the particle;
- (b) its reconstructed spatial trajectory (track or shower axis);
- (c) its hierarchical relationship within the interaction tree (primary or secondary).

Therefore, a PFP is Pandora's complete representation of a particle hypothesis. The reconstruction of PFPs is based on topological pattern recognition, and it is divided into multiple stages. First, 2D clusters in each wire plane are built; then, these clusters are matched across planes to form consistent 3D objects.

Track reconstruction Track-like candidates are formed by matching 2D clusters across the three wire planes. Sliding linear fits are applied in each view to obtain a smooth parametrization of the cluster trajectory. A consistent 3D track is then constructed by requiring compatibility in drift time and wire intersection across planes. A χ^2 metric is used to resolve ambiguities and select the most consistent combination. This is particularly effective for MIPs such as muons and for contained protons.

Shower reconstruction Given the diffuse and non-linear topology of electromagnetic showers, candidate shower clusters are instead identified in each plane based on transverse spread and branching structure. These are then grouped into larger shower objects originating near the reconstructed vertex. Clusters reconstructed in two wire planes are used to determine a set of possible three-dimensional space point candidates. Then, they are projected onto the third plane, where compatible clusters are searched for within timing and geometrical tolerances. When consistency is achieved across all three planes, a 3D shower object is constructed.

Both track-like and shower-like 3D objects are stored as PFPs.

Track-Shower Classification and Scores Once a PFP has been geometrically reconstructed, Pandora assigns a classification score that quantifies how compatible the object is with a track-like or shower-like topology.

This classification is performed using a BDT trained on simulated neutrino interactions, and this classifier uses various reconstruction variables, including geometric variables describing straightness and elongation, Principal Component Analysis (PCA) variables, calorimetric variables related to charge deposition, and spatial variables such as distance to the reconstructed vertex. Straightness quantifies how compatible the hits are with a linear trajectory, and it is generally computed from deviations with respect to a best-fit line; elongation measures how extended the hit distribution is along a dominant spatial axis, and it is often estimated from the ratio of longitudinal to transverse spreads. PCA, instead, is a linear transformation that diagonalizes the covariance matrix of the 3D hit coordinates to identify the orthogonal principal axes along which the spatial variance of the hit distribution is maximized. The resulting eigenvalues are an indicator of the elongation of the charge deposition along each spatial direction.

The output of this classifier is the **track score**, a continuous variable between 0 and 1.

- Track score ≈ 1 indicates strong compatibility with a track-like topology;
- Track score ≈ 0 indicates shower-like behavior.

The track score does not determine how the object is reconstructed, but only classifies an already reconstructed PFP.

After classification, Pandora constructs the full interaction hierarchy: the neutrino candidate is defined as the parent particle, and PFPs originating directly from the reconstructed vertex are the primary daughters. Secondary particles, produced through further interactions or decays, are attached to their corresponding parent PFP via a stored parent identifier.

Particle Identification (PID). Particle identification is based on calorimetric information. For stopping tracks, the energy loss per unit length (dE/dx) is analyzed

as a function of the residual range. Unphysical dE/dx values are removed before this stage.

For each particle hypothesis, a χ^2 score is computed as

$$\text{PID} = \frac{1}{N_{\text{hits}}} \sum_{\text{hits}} \left[\frac{(dE/dx)_{\text{meas}} - (dE/dx)_{\text{theory}}}{\sigma_{dE/dx}} \right]^2. \quad (3.6)$$

Lower values represent a better compatibility with a given particle hypothesis.

Slice Classification After both reconstruction paths are completed, a final BDT is used to classify each slice as neutrino-like or cosmic-like.

The neutrino hypothesis variables include the number of reconstructed daughter particles, the total number of 3D hits, the vertex position, the direction of primary particles, and PCA observables near the vertex.

The cosmic hypothesis relies mainly on properties of the longest reconstructed track, such as vertical direction and hit fraction. The final output of this classifier is known as the neutrino score.

3.1.4 Two-Dimensional Deconvolution

The 1D deconvolution described in the previous section treats each wire independently, correcting the measured waveform using a response function that depends only on time. It is therefore assumed that the signal recorded on a wire originates exclusively from the charge drifting directly toward that wire.

In reality, the detector response in a LArTPC is intrinsically multi-wire, since the induced signals are not only on the closest wire but also on neighbouring wires. This is especially evident in induction planes where the bipolar response extends over several wires and time ticks. For this reason, the full detector response depends both on time and on the transverse wire coordinate.

To account for these correlations, a 2D deconvolution approach was developed and studied in MicroBooNE within the Wire-Cell framework.

In the 2D treatment, the measured signal is expressed as a convolution in both time and wire coordinate as

$$M(w, t) = \sum_{w'} \int R(w - w', t - t') S(w', t') dt', \quad (3.7)$$

where w is the wire index, $R(w - w', t - t')$ is the detector field response including inter-wire induction effects, and $S(w', t')$ is the original ionization signal.

By performing Fourier transforms along both the time and wire directions, the convolution $M(k, \omega)$ becomes a product in frequency space with the form

$$M(k, \omega) = R(k, \omega) S(k, \omega), \quad (3.8)$$

where k is the spatial frequency along the wire direction and ω the temporal frequency.

The ionization signal is then reconstructed through

$$S(k, \omega) = \frac{M(k, \omega)}{R(k, \omega)} F(k, \omega), \quad (3.9)$$

where $F(k, \omega)$ is a filter introduced to suppress noise amplification at frequencies where the response function becomes small.

Within the Wire-Cell signal processing chain, the 2D deconvolution reconstructs the original ionization charge distribution by removing the combined effects of the field response and electronics shaping, all the while suppressing noise and correcting the signal baseline.

This has a significant impact on the ROI since, in MicroBooNE, ROIs are identified using both:

- 1D deconvolution with response functions derived from data;
- 2D deconvolution with simulated field responses and low-frequency filtering.

Once the ROIs are identified, the 2D deconvolution output is further processed with an adaptive baseline (AB) correction. The AB technique computes a local baseline within each ROI by imposing continuity at the start and end points of the region and performing a linear interpolation between them. This removes residual low-frequency distortions that may remain after deconvolution and prevents bias in the reconstructed charge.

The ICARUS collaboration has now adopted the 2D signal processing to improve the upstream reconstruction. In particular, the reconstructed charge becomes less dependent on track orientation and induction-plane effects, and more consistent with the original ionization signal, thus leading to an improved stability of reconstructed observables that rely on calorimetric information. Given the recent introduction of the updated signal processing chain, extensive validation studies are ongoing to quantify its impact on reconstructed physics quantities.

3.1.5 Deep Neural Network ROI identification

As discussed in the previous sections, ROI identification is a central step in LArTPC signal processing. After deconvolution, waveform segments likely containing ionization charge are selected in order to suppress noise and restrict the subsequent hit reconstruction to physically meaningful regions.

In the traditional ICARUS reconstruction chain, ROI finding is performed using a threshold algorithm applied independently to each wire that relies on local signal and neglects any correlation between wire planes.

Together with the update of 2D deconvolution in the signal processing stage, a Deep Neural Network (DNN) approach to ROI identification was developed for the SBN program and for both SBND and ICARUS. DNN ROI operates upstream of Gaussian hit reconstruction and 3D clustering, before, for example, Cluster3D and Pandora.

Each wire plane is represented as a 2D image, with the wire number and time tick being the axes. The network’s task is to classify each pixel (w, t) as either belonging to an ROI or to background (noise).

The network has three input channels:

- (a) The deconvolved waveform on the target plane, after a low-pass filter to suppress high-frequency noise;
- (b) A binary mask for the two-plane time coincidences;
- (c) A binary mask for the three-plane time coincidences.

The coincidence masks introduce cross-plane constraints: ionization signals are correlated across planes, while noise fluctuations and detector artifacts are not. In this way, the network can distinguish real charge depositions from spurious signals more effectively than the previous amplitude threshold approach.

Several encoder–decoder architectures were explored, and for both SBND and ICARUS, the most efficient model was a U-ResNet architecture, which includes residual blocks within a U-Net structure. U-Net is, in brief, a convolutional neural network (CNN) architecture for image segmentation consisting of an encoder-decoder structure. The encoder reduces the spatial resolution of the input image while extracting high-level features, and the decoder uses these features to recover the original image resolution and produce a pixel classification map.

Compared to the traditional algorithm, DNN ROI leads to improved efficiency and purity performances, especially in the case of shallow-angle or elongated tracks and electromagnetic showers [36]. Finally, the ROI efficiency exhibits reduced dependence on track orientation, particularly in induction planes, and remains stable under detector variations such as changes in electron lifetime, gain calibration, and noise level.

3.1.6 NuGraph2: Graph Neural Network Reconstruction in ICARUS

The reconstruction described so far relies on deterministic pattern recognition algorithms and the BDTs within the Pandora framework. A complementary approach called NuGraph2 (NG2) and based on graph neural networks (GNNs) has been developed within the SBN program [37].

NG2 is a low-level reconstruction algorithm that directly acts on reconstructed TPC hits. Its purposes are mainly two: (i) to identify and reject background hits not associated with the primary neutrino interaction (filter task); (ii) to perform hit-level semantic classification into physically meaningful particle categories (semantic task).

Moreover, unlike CNNs, which require voxelized or image representations of detector data, NG2 works on the intrinsically sparse detector observables by representing each event as a graph, thereby avoiding discretization steps and preserving the natural structure of LArTPC data.

NG2 was originally trained on MicroBooNE simulation samples and later adapted and integrated into the ICARUS reconstruction framework.

Graph Construction

The starting point for NG2 is the collection of reconstructed Gaussian hits. Each hit is characterized by four features:

- wire coordinate;
- peak time (drift coordinate);
- integrated charge;
- RMS width of the Gaussian pulse.

For each readout plane (U, V, Y), hits are treated as nodes of an independent planar subgraph. Connectivity within each plane is established using Delaunay triangulation in the two-dimensional (w, t) space, where w is the wire coordinate and t is the drift time.

Given a set of points in the plane, the Delaunay triangulation generates a mesh of non-overlapping triangles such that the circumscribed circle of each triangle contains no other input point inside. This maximizes the minimum angle of the triangulation and prevents the formation of elongated and highly anisotropic triangles.

In this context, the triangulation defines the graph edges, such that two hits are connected if they share an edge in the Delaunay mesh. The resulting planar graph is sparse but fully connected within each plane. This configuration guarantees local connectivity between adjacent hits, automatic adaptation to non-uniform hit density, and inclusion

of both short and long edges. Furthermore, the graph is robust against gaps in the hit distributions, such as those induced by dead, non-efficient wires or a fragmented shower development.

The key strength in Delaunay triangulation is that it does not impose any fixed metric or radius to build relations among hits, thus making it particularly versatile and flexible.

NG2 also introduces **nexus nodes**, constructed from reconstructed 3D space points (e.g., from Cluster3D). A nexus node is connected to the planar hits that participate in the corresponding 3D match. Nexus nodes do not possess intrinsic input features, nor are they connected to one another. They are exclusively used as information-sharing hubs between planes.

The resulting structure is a heterogeneous graph composed of three planar subgraphs and a set of inter-plane nexus connections. This explicitly encodes both 2D local structure and 3D geometric consistency into the network topology.

The architecture of NG2 consists of an initial embedding stage (encoder), followed by multiple iterations of planar and nexus message-passing blocks, and final task-specific decoders. The underlying structure is represented in Fig. 3.1.

Encoder

The encoder transforms the raw hit observables into a higher-dimensional learned representation, referred to as a latent feature space. The input to the network is a set of reconstructed hits, each characterized by the four observables mentioned above. If the number of hits in the event is N_{nodes} and the number of input features per hit is $N_{\text{in}} = 4$, the initial input tensor has shape

$$x_{\text{in}} \in \mathbb{R}^{(N_{\text{nodes}}, N_{\text{in}})}. \quad (3.10)$$

NG2 employs a *categorical embedding* mechanism. Instead of assigning a single feature vector to each hit, the network maintains separate feature representations for each semantic class. Let N_{classes} be the number of semantic categories (MIP, HIP, shower, Michel, diffuse). The input tensor is copied along a categorical dimension, so that each semantic class is assigned its own set of features, thus leading to

$$x^c \in \mathbb{R}^{(N_{\text{nodes}}, N_{\text{classes}}, N_{\text{in}})}. \quad (3.11)$$

A categorical linear multi-layer perceptron (MLP), ϕ_e , is then applied independently to each category, producing the initial latent embedding

$$n_i^{(0)} = \phi_e(x_i^c), \quad (3.12)$$

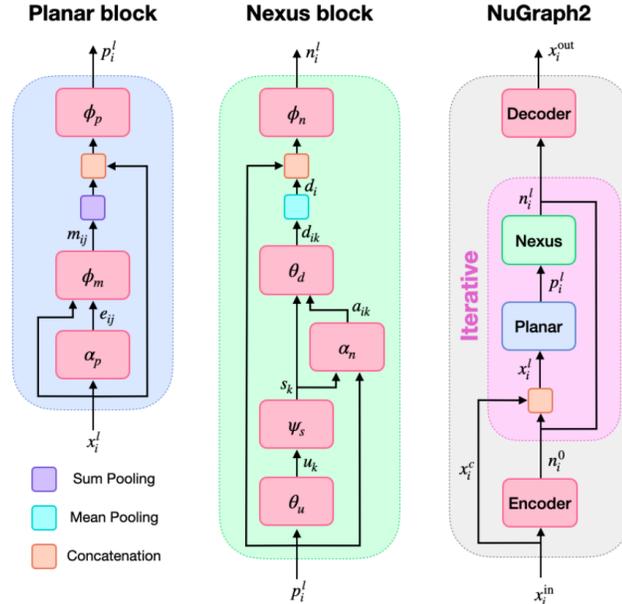


Figure 3.1: Schematic representation of the NuGraph2 architecture. The network operates iteratively and is composed of a planar block, a nexus block, and a final decoder. Starting from input features x_i^{in} , an encoder produces initial node embeddings n_i^0 . In the planar block, neighbouring nodes within the same TPC plane exchange information through message passing, producing updated planar embeddings $p_i^{(l)}$. In the nexus block, $p_i^{(l)}$ are aggregated into nexus representations and redistributed across planes, producing refined node embeddings $n_i^{(l)}$. After several iterations, the final embeddings x_i^{out} are processed by the decoder. From Ref. [37].

with

$$n^{(0)} \in \mathbb{R}^{(N_{\text{nodes}}, N_{\text{classes}}, N_{\text{features}})}. \quad (3.13)$$

The categorical linear structure implies that each semantic class evolves its own feature representation through independent weight matrices. This reduces the number of trainable parameters compared to a fully connected embedding over all classes simultaneously and allows the network to learn class-specific latent descriptions, as well as to exchange controlled information in later stages.

Iterative Message Passing

After encoding, NG2 performs multiple iterations of message passing. Each iteration l consists of two sequential components:

- a **planar block**, which propagates information within each individual wire plane;

- a **nexus block**, which exchanges information between planes using nexus nodes derived from reconstructed 3D space points.

At the beginning of iteration l , the planar node features from the previous step, $n_i^{(l-1)}$, are concatenated with the original input embedding x_i^c

$$x_i^{(l)} = \left[n_i^{(l-1)}, x_i^c \right]. \quad (3.14)$$

This skip connection ensures that low-level geometrical information is preserved throughout the network depth.

Planar block Within each wire plane, neighbouring hits are connected through edges defined by the Delaunay triangulation. For a given hit i , let $\mathcal{N}(i)$ denote the set of neighbouring hits in the same plane.

First, a categorical cross-attention mechanism α_p computes attention weights between connected nodes

$$e_{ij}^{(l)} = \alpha_p \left(x_i^{(l)}, x_j^{(l)} \right), \quad j \in \mathcal{N}(i). \quad (3.15)$$

The cross-attention operates along the categorical dimension: it produces a weight for each semantic class, enhancing messages from strongly activated categories and suppressing less relevant ones.

Edge messages are then constructed through a categorical linear transformation ϕ_m such that

$$m_{ij}^{(l)} = \phi_m \left(x_i^{(l)}, x_j^{(l)}, e_{ij}^{(l)} \right). \quad (3.16)$$

Messages from neighbouring hits are then aggregated via sum pooling through

$$m_i^{(l)} = \sum_{j \in \mathcal{N}(i)} m_{ij}^{(l)}. \quad (3.17)$$

The aggregated message is concatenated with the node’s current representation and passed through another categorical linear layer ϕ_p

$$p_i^{(l)} = \phi_p \left(m_i^{(l)}, x_i^{(l)} \right). \quad (3.18)$$

The planar block, therefore, propagates local topological and calorimetric information within each 2D view, and the network refines the representation of each hit using its planar neighbourhood.

Nexus block The nexus block introduces cross-plane consistency. Nexus nodes correspond to reconstructed 3D space points and connect hits from different planes that are geometrically compatible.

Let $\mathcal{X}(i)$ be the set of nexus nodes connected to planar node i . Planar features are first projected to the nexus space using a transformation θ_u such that

$$u_k^{(l)} = \theta_u(p_i^{(l)}), \quad k \in \mathcal{X}(i). \quad (3.19)$$

Since each nexus node is connected to at most one hit per plane, planar features are aggregated in the nexus node via concatenation across planes. A categorical linear transformation ψ_s then produces updated nexus features

$$s_k^{(l)} = \psi_s(u_k^{(l)}). \quad (3.20)$$

Information is subsequently propagated back to planar nodes. First, a categorical cross-attention mechanism α_n computes weights between planar and nexus nodes

$$a_{ik}^{(l)} = \alpha_n(p_i^{(l)}, s_k^{(l)}). \quad (3.21)$$

Messages from nexus nodes are then constructed via θ_d with

$$d_{ik}^{(l)} = \theta_d(s_k^{(l)}, a_{ik}^{(l)}). \quad (3.22)$$

These messages are finally aggregated using mean pooling

$$d_i^{(l)} = \frac{1}{|\mathcal{X}(i)|} \sum_{k \in \mathcal{X}(i)} d_{ik}^{(l)}. \quad (3.23)$$

Mean pooling is adopted in the nexus block to stabilize the scale of the aggregated message, since the number of connected nexus nodes may vary.

Finally, planar and nexus information are combined through a categorical linear transformation

$$n_i^{(l)} = \phi_n(p_i^{(l)}, d_i^{(l)}). \quad (3.24)$$

The updated embedding $n_i^{(l)}$ is then passed to the next iteration.

Through repeated planar and nexus message-passing steps, the network progressively builds representations that encode both local 2D topology and global 3D geometric consistency, thus improving semantic consistency across planes.

Decoders and Outputs

After L iterations of planar and nexus message passing, each hit i is represented by a latent embedding

$$n_i^{(L)} \in \mathbb{R}^{(N_{\text{classes}}, N_{\text{features}})}. \quad (3.25)$$

Two independent decoders are then applied to perform the two reconstruction tasks: background filtering (filter task) and semantic classification (semantic task).

Filter decoder The filter decoder performs a binary classification at the hit level, identifying whether a hit is compatible with the primary neutrino interaction or with background activity (e.g., electronics noise or cosmic-induced hits).

Since this task is agnostic to the semantic class information, the categorical dimension of the embedding is first flattened such that

$$\tilde{n}_i^{(L)} \in \mathbb{R}^{(N_{\text{classes}} \cdot N_{\text{features}})}. \quad (3.26)$$

A fully connected linear projection is then applied to produce a single scalar score per hit

$$f_i = \sigma\left(\phi_f(\tilde{n}_i^{(L)})\right), \quad (3.27)$$

where ϕ_f denotes the filter decoder network and σ is the sigmoid activation function. The output $f_i \in [0, 1]$ represents the probability that the hit belongs to the neutrino interaction.

Semantic decoder The semantic decoder performs a multi-class classification of hits belonging to the interaction. In this case, the categorical structure of the embedding must be preserved, since each category maintains its own learned feature representation.

A categorical linear transformation ϕ_s is applied independently to each class embedding, producing semantic class logits such that

$$z_{i,c} = \phi_s\left(n_{i,c}^{(L)}\right), \quad (3.28)$$

where c runs over the N_{classes} semantic categories.

A softmax activation is finally applied along the categorical dimension,

$$p_{i,c} = \frac{\exp(z_{i,c})}{\sum_{c'} \exp(z_{i,c'})}, \quad (3.29)$$

with $p_{i,c}$ being the probability that hit i belongs to semantic class c .

Integration in the ICARUS Reconstruction Chain

NG2 has been integrated into the `LArSoft` (Liquid Argon Software) framework, which is the common software framework for simulation and reconstruction in liquid argon TPC experiments. It has also been integrated within `icaruscode`, the ICARUS software implementation built on top of `LArSoft`. In the current configuration:

- it runs after Pandora slicing and barycenter flash matching;
- only the slice closest to the triggering flash is processed;
- inference outputs are stored per hit as `anab::FeatureVector` objects;
- summary quantities are written into CAF (Common Analysis Format) files, which store high-level reconstructed variables for physics analyses.

The CAFs currently store:

- the fraction of slice hits passing the filter;
- for each reconstructed PFP, the fraction of hits assigned to each semantic category.

Inference time is of order 0.1 s-0.2 s per event on CPU, so the algorithm is compatible with production workflows.

Role within the Overall Reconstruction Strategy

NG2 is not intended to replace Pandora in ICARUS at this stage. Instead, it is intended as a complementary tool providing hit-level information that can be exploited downstream for multiple purposes. Potential applications include:

- improving slice purity using the filter output;
- aiding track-shower discrimination;
- enhancing shower completeness by recovering missing fragments;
- assisting particle identification using semantic fractions;
- improving vertex reconstruction.

In general, the integration of machine learning techniques within a classical reconstruction framework is a necessary development to address the intrinsic limitations of

LArTPC reconstruction, including track splitting, shower fragmentation, and ambiguities in particle identification, which are otherwise difficult to overcome using deterministic algorithms. In this context, NG2 is an additional layer of relational information that can help mitigate such limitations by operating directly at the upstream stages of reconstruction.

In ICARUS, the initial NG2 model was trained on samples reconstructed with 1D deconvolution. The most recent adoption of the 2D signal processing modifies both the charge response and the hit distributions, meaning that an updated training is required to preserve performance and internal consistency across the reconstruction chain.

This retraining becomes especially relevant as the collaboration increasingly explores NG2 as an integrated reconstruction component within the workflow. Moreover, recent developments have demonstrated growing interest in NG2 across the entire SBN program, and a porting of the algorithm to SBND is currently planned.

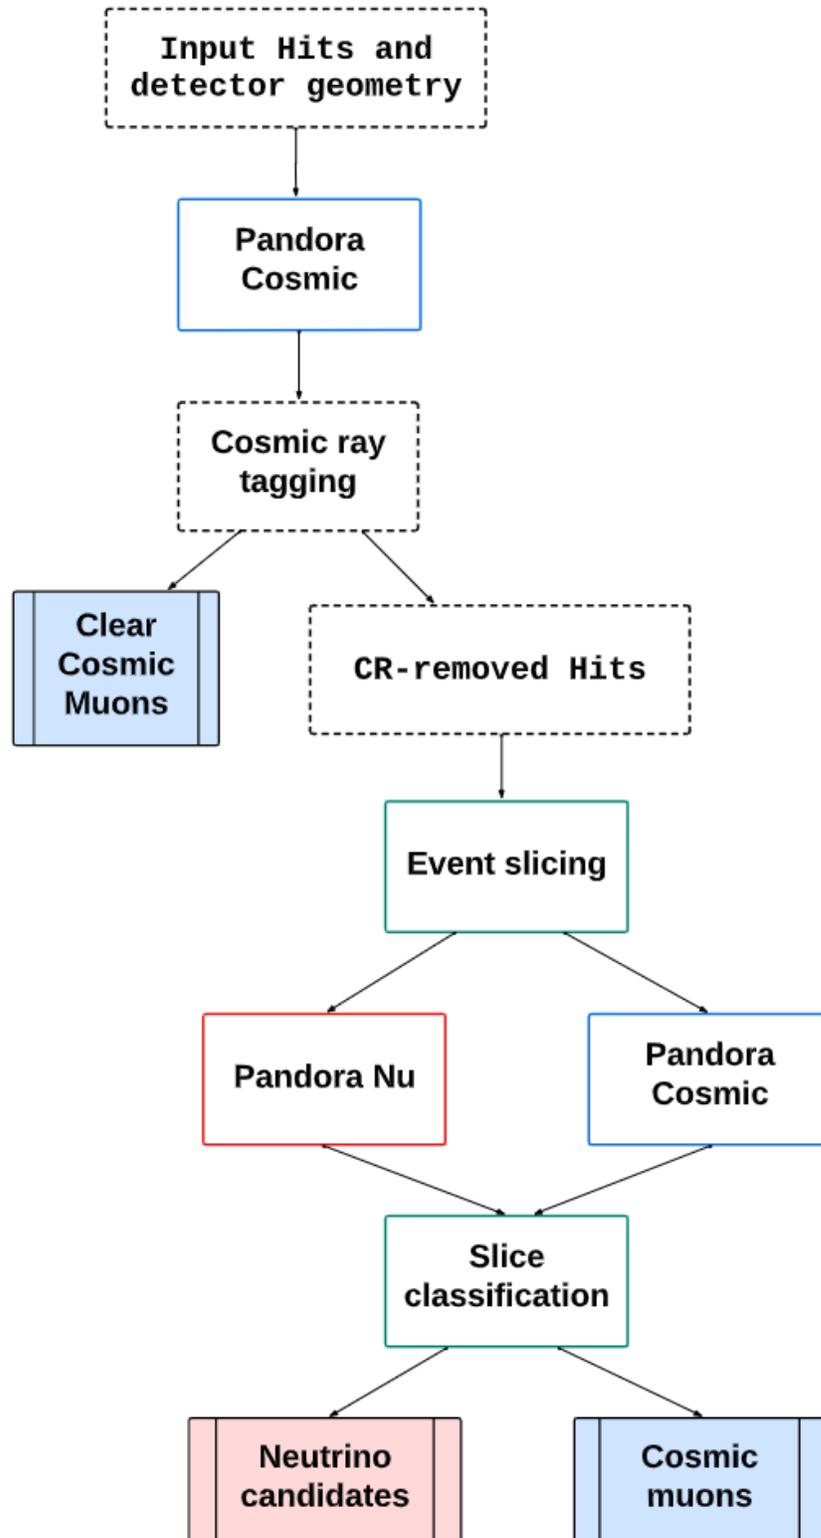


Figure 3.2: Pandora processing chain in a schematic view. Reconstruction steps that handled by outsiders (e.g., LArSoft) are indicated within a dashed line box. Colored boxes, instead, indicate outputs that are saved for later use.

Chapter 4

NuGraph2 Training with the Updated ICARUS Reconstruction Chain

4.1 Motivation for Retraining

The introduction of the updated ICARUS signal processing chain, based on 2D deconvolution and DNN-ROI identification, modifies the low-level representation of reconstructed charge deposits.

As discussed in Chapter 3, the 2D deconvolution corrects inter-wire induction effects and reduces the dependence of reconstructed charge on track orientation, while the DNN-based ROI identification improves signal efficiency and suppresses noise. Since NG2 operates directly on reconstructed Gaussian hits, its input feature space is intrinsically sensitive to any modifications made in upstream signal processing.

The previous ICARUS training of NG2 was performed on samples reconstructed with 1D deconvolution. The transition to 2D signal processing, therefore, actively shifts the hit-level distributions and disrupts NG2's functionality. To ensure internal consistency and performance stability of the reconstruction chain, a complete retraining of NG2 on samples processed with the updated upstream reconstruction was required.

4.2 Training Sample Production

4.2.1 Official Beam Samples

The core of the training dataset was produced through the official ICARUS Monte Carlo production campaigns at Fermilab. Three primary samples were generated:

- a **BNB nominal sample**, containing approximately 10^5 neutrino interactions simulated in the BNB configuration;
- a **NuMI nominal sample**, containing approximately 10^5 interactions corresponding to the off-axis NuMI flux;
- a **MPV/MPR high-multiplicity sample**, containing approximately 2×10^5 events specifically designed to enhance topological diversity.

Samples are produced per standard simulation production chain, thus starting from the neutrino event simulator GENIE that gets then propagated through matter with the largely used GEANT4 toolkit [38, 39]. The next steps include the detector simulation response and the Stage0 and Stage1 reconstruction steps. The nominal BNB and NuMI samples represent realistic beam-induced neutrino interactions and are characterized by their respective flux spectra and flavor compositions. The BNB sample is strongly dominated by ν_μ CC interactions, with subleading contributions from $\bar{\nu}_\mu$, ν_e , and $\bar{\nu}_e$. The NuMI sample, on the other hand, has a broader energy spectrum and an enhanced high-energy tail.

These beam samples ensure that the training dataset reflects realistic physics conditions which are relevant for ICARUS analyses.

MPV/MPR High-Multiplicity Sample In addition to beam-induced samples, a dedicated high-multiplicity dataset was included in the training.

The MPV (Multi-Particle Vertex) and MPR (Multi-Particle Rain) samples are synthetic event configurations developed within the SBN program to stress-test machine learning reconstruction algorithms in dense environments. The MPV configuration generates multiple particles that emerge from a common interaction vertex, thus displaying increased local topological complexity. The MPR configuration, instead, distributes multiple particles throughout the detector volume, producing spatially dense and potentially overlapping activity.

These samples are not intended to reproduce realistic beam fluxes, but are designed to expand the coverage of possible topologies encountered during training. Their inclusion serves several purposes, including the enhancement of particle multiplicity coverage, improved robustness against shower fragmentation and overlapping activity, and, most of all, mitigation of bias toward beam-like interaction topologies.

The truth information in this sample uses the ν_τ neutrino species, which allows a clear distinction from beam-like datasets, where ν_τ is absent.

It is important to underline that the multiplicity distribution of generator-level primary particles in the MPV/MPR sample is significantly broader than in the beam samples. This reflects its role as a stress-test dataset rather than a physics-accurate flux simulation.

Dedicated ν_e -Only Production As previously seen, electromagnetic shower reconstruction represents one of the most challenging aspects of LArTPC event reconstruction. Since the nominal beam samples are strongly dominated by ν_μ charged-current interactions, the fraction of ν_e -induced electromagnetic topologies is limited.

In order to increase the statistical representation of electromagnetic interactions, an additional independent production of approximately 5×10^4 ν_e -only events was performed for both BNB and NuMI flux configurations.

Due to computational constraints associated with full detector simulation, the detailed YZ field response simulation was disabled in this private production. In more detail, the YZ simulation accounts for the full transverse induction effects between adjacent wires and significantly increases the computational cost. Its removal reduced production time while still maintaining compatibility with the downstream 2D deconvolution and DNN ROI processing chain.

Given that NG2 operates on reconstructed hit-level observables after signal processing and clustering, this approximation does not change the structure of the input feature space in a way that could compromise the training consistency.

Reconstruction Chain and Data Format All samples, both official and privately produced, were reconstructed using the updated ICARUS configuration, including:

- 2D deconvolution;
- DNN-based ROI identification;
- Cluster3D space-point reconstruction;
- Pandora slicing.

It should be noted that, during training, only the slice containing the largest number of neutrino-induced hits was selected for each event. This choice ensures that the graph construction is performed on the slice most representative of the true neutrino interaction, minimizing contamination from residual cosmic or unrelated activity.

This differs from the NG2 inference configuration within the ICARUS reconstruction chain. In production, the algorithm is executed after slicing and flash matching, and only the slice associated with the triggering optical flash is processed.

For training purposes, the Stage1 reconstruction outputs were converted to the HDF5 file format. This conversion enables efficient data loading and batching within the Python-based training framework, which relies on graph data structures and the usage of GPUs for an accelerated optimization of the task.

Two-Cryostat Training Restoration The initial ICARUS implementation of NG2 training was restricted to a single cryostat. In this work, the full training chain was restored and validated for both ICARUS cryostats.

This required resolving inconsistencies in event indexing, hit-to-truth associations, and slice-level bookkeeping. Training on both cryostats is essential to ensure that the model does not develop biases related to asymmetric detector conditions. In addition, restricting the training to a single cryostat would unnecessarily reduce the available statistics, and thus limiting the representativeness of the dataset.

4.3 Training Sample Characterization

The training dataset combines beam-induced neutrino interactions from the BNB and NuMI fluxes with a dedicated high-multiplicity MPV/MPR production. The main physical and topological properties of each component are summarized below.

4.3.1 Beam-Induced Samples: BNB vs NuMI

Neutrino Flavor Composition

In Fig. 4.1 the true neutrino flavor composition for the BNB and NuMI samples is shown.

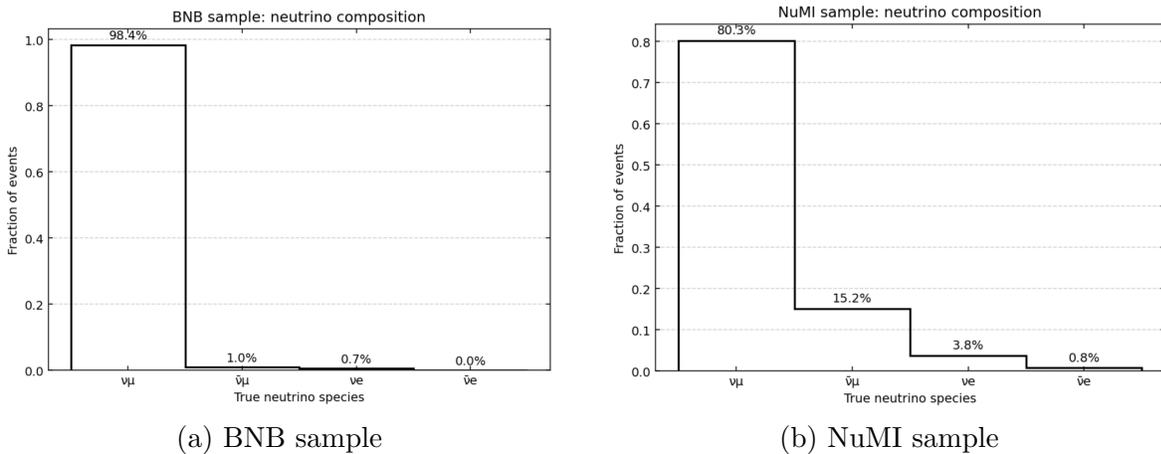


Figure 4.1: True neutrino flavor composition in the beam-induced samples.

The BNB sample is strongly dominated by ν_μ interactions (approximately 98%), with only minor $\bar{\nu}_\mu$ and intrinsic ν_e components.

The NuMI sample exhibits a broader composition, with roughly 80% ν_μ , 15% $\bar{\nu}_\mu$, and a visible ν_e fraction of about 4%. This reflects the distinct beam configuration and energy profile.

Neutrino Energy Spectra

The true neutrino energy distributions are shown in Figure 4.2.

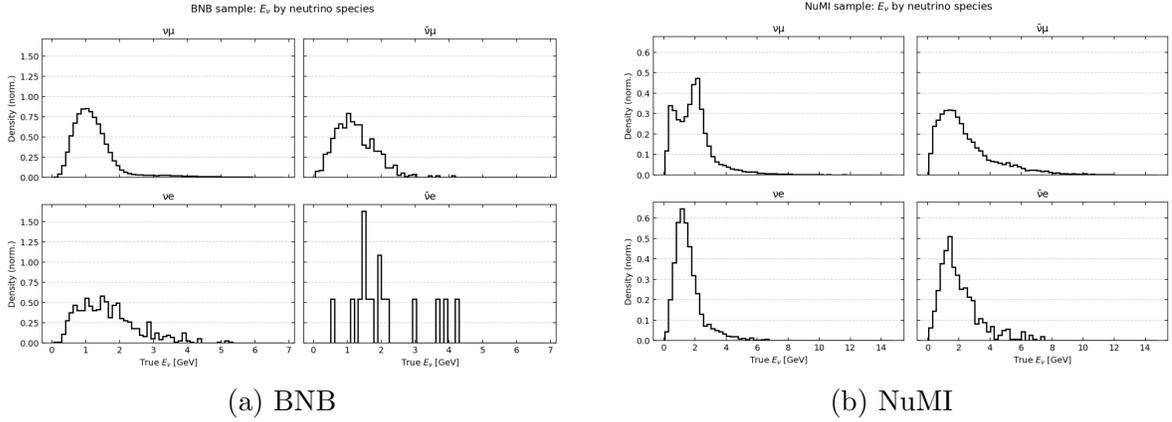


Figure 4.2: True neutrino energy distributions by species.

The BNB spectrum is concentrated below ~ 2 GeV, producing relatively soft final states. On the other hand, NuMI extends to significantly higher energies, with a long high-energy tail reaching beyond 10 GeV. This leads to increased secondary activity and more extended topologies.

Charged Lepton Energy (CC Only)

The charged lepton energy distributions are shown in Figure 4.3.

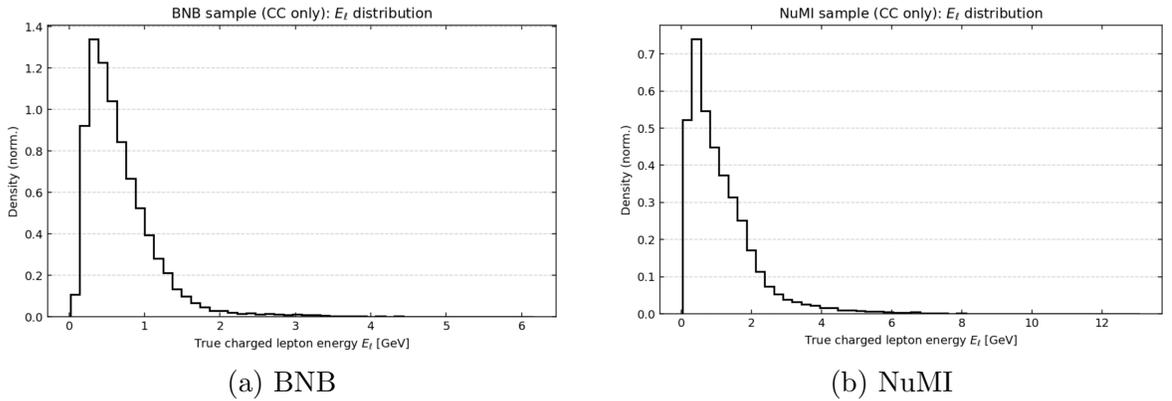


Figure 4.3: True charged lepton energy distributions (CC interactions).

The lepton kinematics follow the neutrino spectra, as BNB interactions predominantly produce low-energy muons, while NuMI exhibits broader distributions with visible high-energy tails.

Electromagnetic Activity

The electromagnetic multiplicity distributions are shown in Figure 4.4.

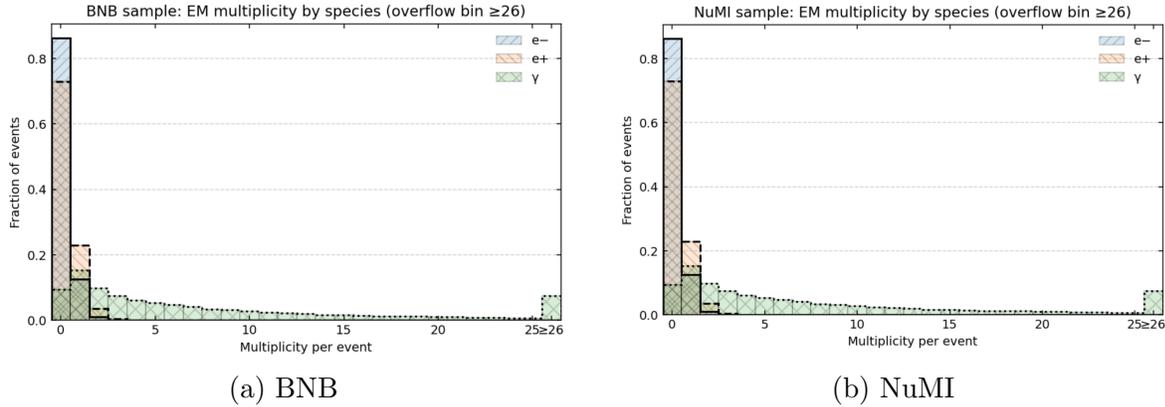
Figure 4.4: Electromagnetic multiplicity per event (overflow bin ≥ 26).

Table 4.1: Event-level electromagnetic activity in the BNB and NuMI samples.

	BNB	NuMI
Events with $\geq 1 \pi^0$	12.97%	18.15%
Resolved $\pi^0 \rightarrow \gamma\gamma$	8.14%	11.79%

Table 4.2: Primary and secondary fractions for electromagnetic particles in the BNB and NuMI samples.

EM species	BNB		NuMI	
	Primary	Secondary	Primary	Secondary
e^+	20.9%	79.1%	28.5%	71.5%
e^-	64.6%	35.4%	60.3%	39.7%
γ	90.2%	9.8%	91.7%	8.3%

Table 4.3: Dominant parent origin for electromagnetic particles (fractions normalized per EM species).

EM species	BNB	NuMI
e^+ from μ^+	99.6%	99.4%
e^- from μ^-	98.3%	96.0%
γ from hadronic origin	89.0%	90.7%

Electromagnetic activity plays a crucial role in the semantic reconstruction task, as electromagnetic showers are among the most challenging topologies for hit-level classification.

At the event level (Table 4.1), approximately 13% of BNB interactions contain at least one π^0 , with about 8% exhibiting a resolved $\pi^0 \rightarrow \gamma\gamma$ topology. In the NuMI sample, these fractions increase to 18% and 12%, respectively, thus reflecting the higher neutrino energies and the higher hadronic activity of this beam configuration.

At the particle level (Table 4.2), the electromagnetic component is largely dominated by secondary production mechanisms. Positrons are predominantly secondary in both samples (79% in BNB and 72% in NuMI), while electrons exhibit a mixed behavior, with roughly 60-65% classified as primary and the remainder produced in secondary interactions. Photons are mostly primary in terms of GEANT hierarchy ($\sim 90\%$ in both samples); however, the parent-origin analysis (Table 4.3) shows that the vast majority of electrons originate from muon decay, and that most photons are embedded within hadronic cascades rather than arising exclusively from clean π^0 decays.

These features imply that electromagnetic activity is rarely isolated. Instead, it is typically produced within complex hadronic environments, where overlapping showers, secondary conversions, and decay products coexist within the same slice. This naturally increases topological ambiguity at the hit level.

From a training perspective, this is particularly relevant. The model must learn to distinguish between primary electrons from ν_e -CC interactions, Michel electrons from muon decay, photons from π^0 decay, and diffuse electromagnetic activity generated within hadronic showers. These categories often share similar local hit patterns while differing in global topology and parent origin. This makes shower identification and semantic labeling one of the most demanding aspects of the reconstruction task.

While both beam-induced samples provide realistic interaction environments, the NuMI configuration amplifies these challenges due to its broader energy spectrum and higher multiplicity, leading to longer showers and greater particle overlap. Consequently, the inclusion of both BNB and NuMI samples is essential to expose the model to a wider spectrum of electromagnetic topologies and ensure robustness across different kinematic regimes.

4.3.2 MPV/MPR High-Multiplicity Sample

The MPV/MPR production is especially designed to enhance topological complexity. Although it does not reproduce a physical neutrino flux spectrum, it is tailored to generate neutrino-like multi-prong interaction topologies with elevated primary multiplicity and increased spatial overlap between final-state particles.

Primary Particle Multiplicity

The primary-particle multiplicity distribution is shown in Figure 4.5.

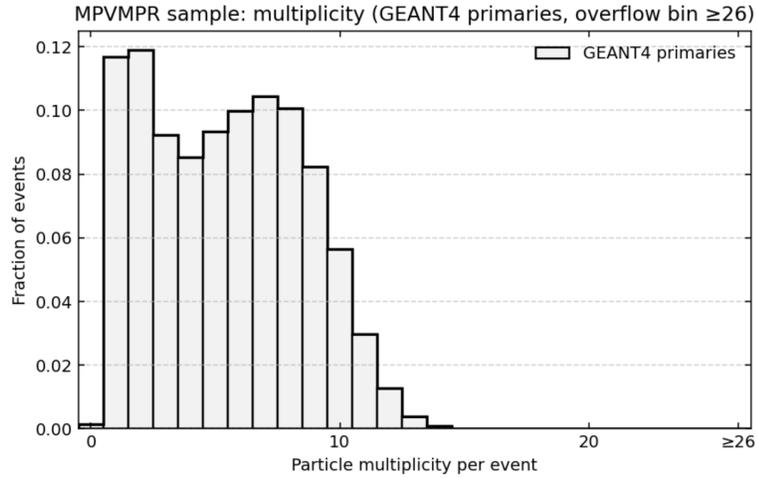


Figure 4.5: MPV/MPR sample: primary particle multiplicity (GEANT4 primaries).

The distribution is significantly broader than in the beam-induced samples, with events frequently containing several primary particles and a long, high-multiplicity tail. This leads to denser graph representations and increased semantic ambiguity during training.

Primary Lepton Energy

The primary lepton energy distribution is shown in Figure 4.6.

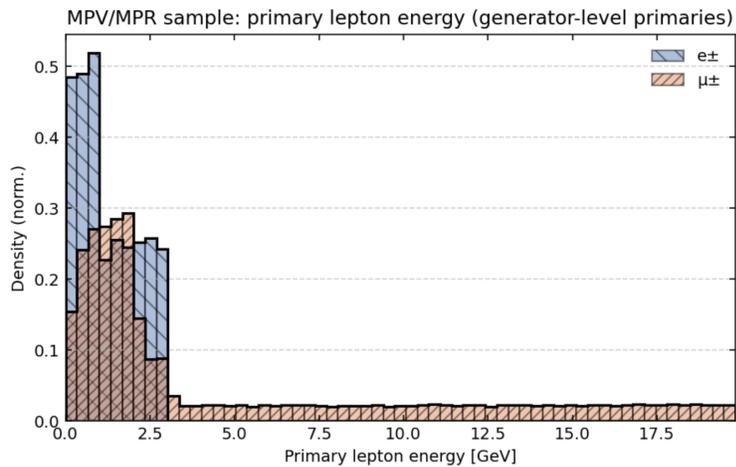


Figure 4.6: MPV/MPR sample: primary lepton energy (generator-level primaries).

Unlike the beam-induced samples, the energy spectrum is not constrained by a specific flux, thus covering a broad kinematic range.

Impact on Training The inclusion of MPV/MPR events increases graph density, spatial overlap between topologies, and message-passing complexity. Although this increases the intrinsic difficulty of the task, it improves model robustness and mitigates potential biases toward sparser beam-like configurations.

ν_e -Only Samples As previously said, dedicated ν_e -only BNB and NuMI productions were generated to improve the statistical coverage of electromagnetic shower topologies. Since these samples are artificially enriched in ν_e interactions and do not reflect the physical beam composition, their detailed distributions are not shown separately. Their only purpose is to improve EM representation during training.

4.4 Training Results and Pre-Processing Studies

The production of the training samples described in the previous section required a significant computational effort. The full simulation and reconstruction chain for the official and private productions extended over approximately three months. Given the cost associated with generating additional datasets, a dedicated study was performed to evaluate the performance of the NG2 filter task under different dataset preprocessing strategies before committing to the final training configuration.

In particular, the study focused on the behavior of the *filter decoder*, whose role is to identify hits belonging to the primary neutrino interaction while rejecting background activity.

During the early development phase, one of the proposed strategies was to use the filter not only to reject detector noise, but also to suppress activity originating from cosmic rays. In such a configuration, the filter would learn to distinguish neutrino-induced hits from both electronics noise and cosmic-induced charge depositions, effectively defining the ground-truth signal as only neutrino activity.

While conceptually appealing, this configuration significantly increases the difficulty of the classification task. Cosmic-ray interactions often produce topologies that are locally similar to neutrino-induced tracks, especially for minimum-ionizing particles traversing the detector volume. As a consequence, treating cosmic activity as background introduces a much more challenging decision boundary for the filter network.

To investigate this aspect, a first round of model evaluation was performed using a preliminary training campaign, hereafter referred to as the *first-pass training*. These models were trained on datasets composed of approximately 90% BNB and NuMI beam-induced interactions.

Three alternative dataset preprocessing configurations were tested, producing three corresponding models, labeled **Model A**, **Model B**, and **Model C**.

- **model A: cosmic-as-background configuration.** In this setup, the filter task

was trained to classify both pure detector noise and cosmic-ray induced activity as background. Therefore, the background class consisted of

$$\text{background} = \text{electronics noise} + \text{cosmic-induced hits}.$$

This configuration forces the filter to separate neutrino-induced activity from both instrumental noise and physically meaningful cosmic tracks.

- **model B: neutrino-only filtering with cosmic hits excluded from the semantic loss.** In this configuration, the filter task was trained to reject only pure noise. Cosmic-ray hits were not treated as background for the filter and were excluded from the semantic classification loss in order to reduce class imbalance effects.
- **model C: neutrino-only filtering with cosmic hits included in the semantic loss.** This model follows the same filter definition as Model B, where only pure noise is considered background, but cosmic-ray hits are included in the semantic classification loss. This choice modifies the effective class balance during training while leaving the filter task unchanged.

In order to compare the filter behavior across these configurations, a threshold sweep study was performed on the filter output score. Since the models do not share the same definition of background (in particular Model A treats cosmic hits as noise), the evaluation was restricted to *pure detector noise* in order to provide a consistent comparison across models.

Figure 4.7 shows the signal efficiency and background rejection as a function of the filter threshold, where the threshold is applied to the binary filter score: a score close to unity is signal-like, while a score close to zero is background-like. The ideal situation is to have maximum noise rejection and signal survival for a score threshold close to one.

As expected, the performance of Model A is significantly worse in terms of noise rejection. This behavior reflects the increased difficulty of the classification task when cosmic activity is included in the background class. Cosmic-ray tracks share many local features with neutrino-induced tracks, making the discrimination problem intrinsically more complex for the network.

On the other hand, Models B and C exhibit very similar performance curves, indicating that the choice of including or excluding cosmic hits in the semantic loss has only a limited impact on the filter task itself.

Following internal discussions within the reconstruction development effort, it was concluded that the suppression of cosmic-ray activity should not be a primary objective of the NG2 filter at this stage. As a matter of fact, simultaneously reject cosmic-induced hits and detector noise worsens the overall filtering performance, at least in this training workflow.

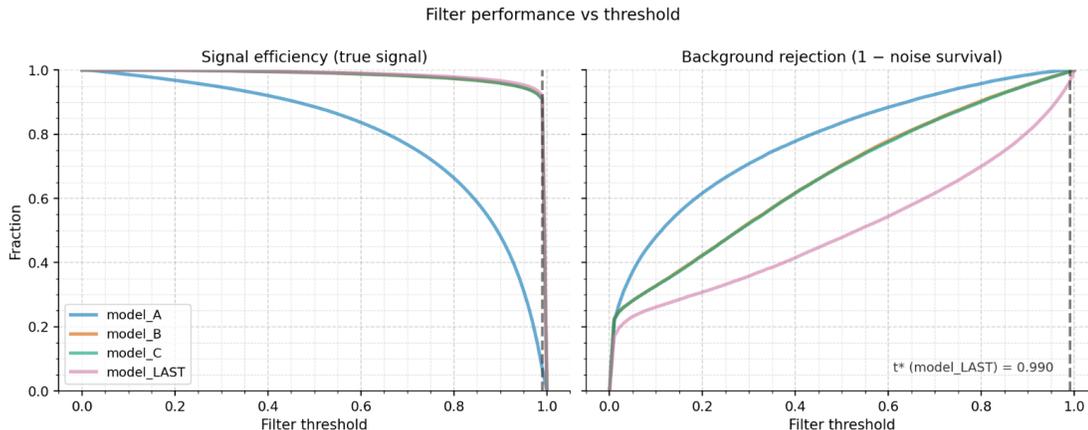


Figure 4.7: Filter performance as a function of the filter score threshold for the models evaluated during the first-pass training study. The left panel shows the signal efficiency for true neutrino hits, while the right panel shows the rejection of pure detector noise ($1 - \text{noise survival}$). Model A corresponds to the configuration where the filter was trained to reject both electronics noise and cosmic-induced hits, making the classification task significantly more difficult. Models B and C instead treat only pure detector noise as background, differing only in whether cosmic-ray hits are excluded from (B) or included in (C) the semantic loss. The curve labeled `model_LAST` corresponds to the final training configuration adopted after these studies.

For this reason, the final training strategy adopted for the updated NG2 model keeps the configuration in which the filter is optimized to reject only detector noise, leaving cosmic discrimination to other components of the reconstruction.

`model_LAST`, shown in Figure 4.7, corresponds to the final training configuration. In this setup, the filter is trained to reject only detector noise, while cosmic hits are excluded from the semantic loss. The model was trained using the full dataset described in the previous section and corresponds to the version integrated into the official `icaruscode` reconstruction chain. All subsequent results presented in this work therefore refer to this model.

4.4.1 Evaluation on the Test Sample

The performance of the final NG2 model (`model_LAST`) was evaluated on an independent test sample. Indeed, the full dataset used for training was partitioned into three subsets:

- **Training set** (90%), used to optimize the model parameters;
- **Validation set** (5%), used during training to monitor convergence and prevent overfitting;
- **Test set** (5%), used exclusively for the final performance evaluation.

The results presented in this section therefore correspond to the model predictions on the **test sample**, which is statistically independent from the data used during training and validation. This ensures that the reported metrics provide an unbiased estimate of the model’s generalization performance.

As described in Chapter 3, NG2 performs two distinct hit-level classification tasks:

- a **filter task**, which separates neutrino-induced hits (signal) from detector noise;
- a **semantic classification task**, which assigns each signal hit to one of the five semantic particle categories: MIP, HIP, shower, Michel, and diffuse.

The performance of both tasks is summarized through confusion matrices and derived classification metrics.

Precision and Recall Two commonly used metrics to evaluate classification performance are **precision** and **recall**.

For a given class, the **recall** (also known as efficiency or true positive rate) is defined as

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}, \quad (4.1)$$

where TP denotes the number of true positives and FN the number of false negatives. Recall therefore quantifies the fraction of true objects of a given class that are correctly identified by the model.

The **precision**, instead, is defined as

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad (4.2)$$

where FP denotes the number of false positives. Precision measures the fraction of objects assigned to a given class that are actually correct.

In the context of hit-level classification, recall reflects how efficiently the model recovers hits belonging to a given particle category, while precision quantifies the purity of the predicted class.

Semantic Classification Performance The confusion matrices for the semantic classification task are shown in Figures 4.8. These matrices summarize the fraction of hits assigned to each predicted semantic class.

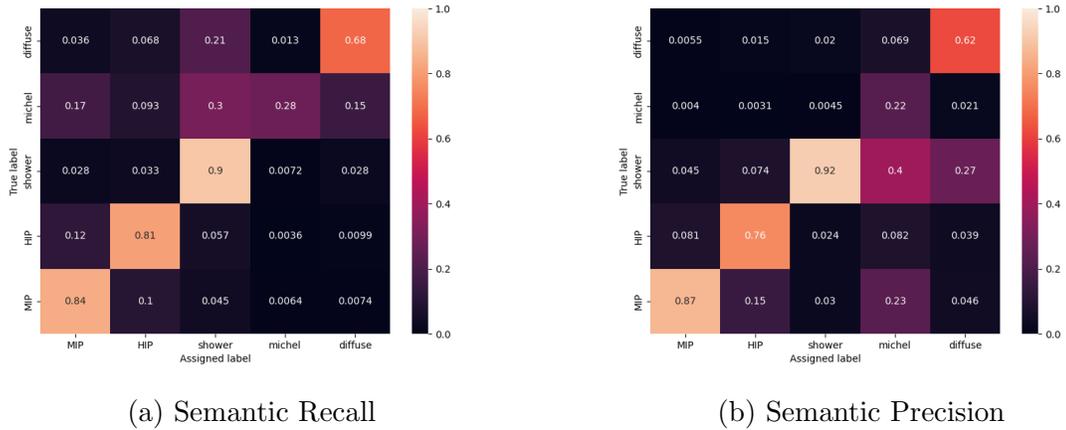


Figure 4.8: Normalized confusion matrix for the semantic classification task showing the recall and precision for each predicted class.

Filter Task Performance The performance of the filter decoder is summarized in Figures 4.9. These matrices quantify the ability of the network to correctly separate neutrino-induced hits from detector noise.

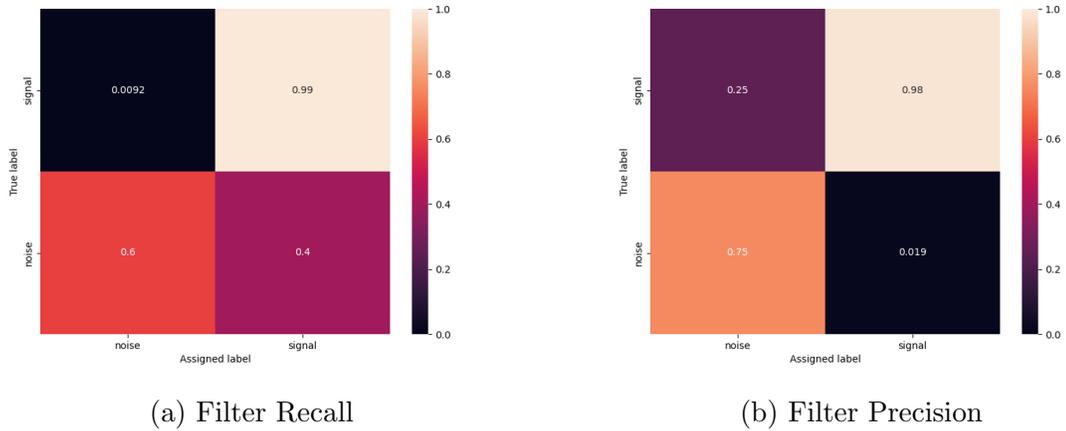


Figure 4.9: Normalized confusion matrix for the filter classification task showing the recall and precision for each predicted class.

The filter performance remains particularly challenging, likely due to the high event multiplicity present in the MPV/MPR samples. At the same time, it should also be considered within the 2D signal processing chain, there are less noise hits with respect to the standard 1D one, and, in other words, the model has less training ground in discriminating noise against signal.

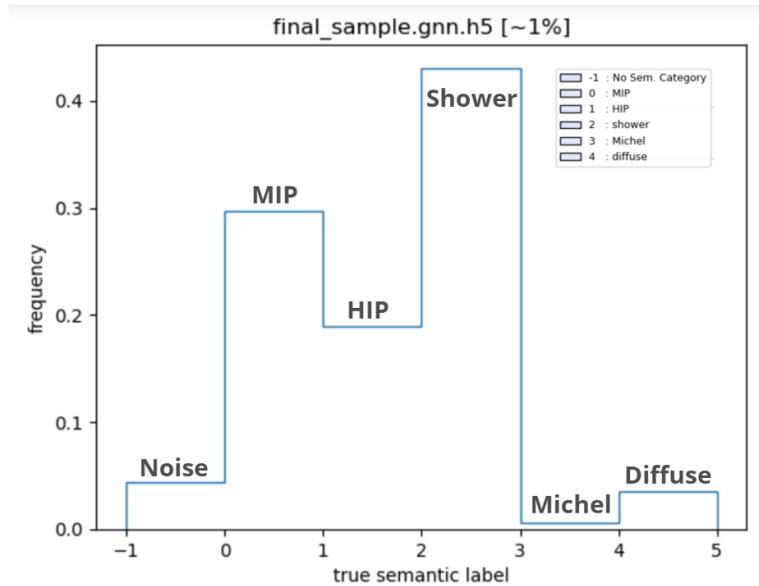


Figure 4.10: Distribution of the true semantic labels in the training dataset. The histogram shows the relative abundance of the five semantic hit categories (MIP, HIP, shower, Michel, diffuse), illustrating the intrinsic class imbalance present in the dataset. In particular, shower hits dominate the distribution due to the high electromagnetic activity in the MPV/MPR samples, while Michel and diffuse hits represent the least populated classes.

In fact, a perfectly balanced representation of each semantic hit category cannot be achieved in the training dataset, as shown in Fig. 4.10. In particular, the large electromagnetic activity in the MPV/MPR production leads to a strong dominance of shower hits, which are by far the most represented semantic class. Michel and diffuse hits are significantly less represented, as expected, thus resulting in a highly unbalanced class distribution and lower performance in discriminating those semantic classes.

Chapter 5

Electron Candidate Selection with NuGraph2

The purpose of this study is to assess whether the NG2 semantic classification outputs can be used to improve the quality of the electron candidate selected by the standard Pandora reconstruction in ICARUS. As seen in Chapter 3, in the nominal reconstruction chain, Pandora provides particle-flow objects (PFPs) and assigns to each of them a track-shower classification score (`trackScore`). NG2, instead, provides an independent, hit-driven semantic characterization of reconstructed activity, which is then summarised in the CAF as per-PFP semantic fractions (e.g., MIP/HIP/shower content).

Direct Pandora-NuGraph2 score comparison A first strategy considered for this study was to perform a fully parallel ν_e -like selection, applying the same event- and slice-level requirements (truth-matched slices, fiducial volume constraints, and comparable energy/quality cuts), and changing only the definition of the shower-like electron candidate. In this approach, the baseline candidate would be chosen either with the standard Pandora object-level classifier (via `trackScore`) or, alternatively, by using instead the NG2 per-PFP shower score (`shr_prob` / `shr_frac`), thus selecting the most shower-like PFP according to NG2.

After an initial validation on truth-defined PFPs, this direct replacement was discarded. Indeed, the NG2 shower score at PFP level was found to have no discriminating power between true electrons and non-electron objects, as it can be clearly seen in Fig. 5.1, where the distributions overlap almost completely, with indistinguishable central values. In the studied sample, the mean shower probability is $\langle \text{shr} \rangle = 0.13212$ for true electrons and $\langle \text{shr} \rangle = 0.13202$ for non-electrons, and the corresponding medians are $\tilde{\text{shr}} = 0.13237$ (electrons) and $\tilde{\text{shr}} = 0.13204$ (non-electrons). This behaviour is compatible with an approximately random guess, and is thus not a usable shower-non-shower separator at the PFP level. It should also be noted that this is an evident inconsistency with respect to the semantic performance at the hit-level, which was presented in Chapter 4, and it

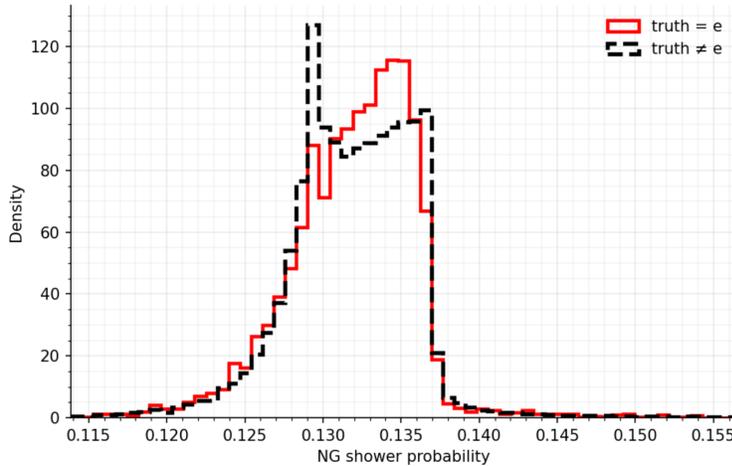


Figure 5.1: Density of NG2 shower probability score for non-e and e truth matched PFP.

could be related to the way that NG2 scores are translated to per-PFP information.

For this reason, the analysis does not rely on `shr_prob` as a primary selection variable. Instead, NG2 information is exploited through its most robust semantic categories: the MIP and HIP probabilities, directly sensitive to track-like and hadronic activity that can be used as a veto on the baseline Pandora candidate.

Data model and variable selection The analysis is based on reconstructed Monte Carlo CAF files produced within the ICARUS reconstruction framework. The sample corresponds to a simulation processed with the updated 2D signal processing chain, followed by the standard Stage 1 reconstruction and the subsequent NG2 inference step. The CAF files therefore contain both the conventional Pandora reconstruction outputs and the NG2 semantic summary quantities associated with each reconstructed PFP.

From these files two flat analysis tables were constructed:

- (a) a slice-level table (`slice.df`), used to define the neutrino interaction sample and its event-level context;
- (b) a PFP-level table (`pfp.df`), used to define candidate objects, truth matching, and NG2 semantic observables.

In the current ICARUS configuration, NG2 inference is executed only on the slice associated with the flash that best matches the beam trigger. Consequently, NG2 semantic information is available only for this selected slice. For consistency, the present study therefore restricts the analysis to slices for which both generator truth matching and NG2 semantic scores are defined.

The sample considered is a fraction of an official ICARUS Monte Carlo production that includes simulated beam neutrino interactions overlaid with cosmic-ray activity. Although the simulation includes cosmic overlay, the present study focuses exclusively on the reconstructed neutrino slice and does not explicitly use the overlay component.

After applying the requirements of truth-matched slices and the presence of valid NG2 outputs, the dataset used in this analysis contains a total of 74 706 reconstructed slices and 266 909 reconstructed PFP objects.

Signal slice definition The analysis is restricted to slices corresponding to true ν_e charged-current interactions occurring inside the detector fiducial volume. These slices constitute the *signal slice sample* used throughout the study.

All performance metrics are defined with respect to this signal sample. In particular, the total number of signal slices represents the denominator used for the computation of the overall reconstruction efficiency. This number corresponds to 185 slices, leading to a signal fraction with respect to the whole ν_μ -dominated sample of 0.2%.

Definition of the target electron In each signal slice, the reconstructed object corresponding to the primary electron is defined using truth-matching information at PFP level. A PFP is considered a true-electron candidate if its matched true PDG code satisfies $|\text{PDG}| = 11$.

Because electromagnetic showers may fragment into multiple reconstructed objects, several PFPs can be partially associated with the same true electron. To identify a unique reference object, the target electron PFP is defined as the one with the largest energy completeness,

$$\text{target} \equiv \underset{i \in \{|\text{PDG}_i|=11\}}{\text{argmax}} C_i^E, \quad (5.1)$$

where C^E is the truth-matching energy completeness, defined as

$$C^E = \frac{E_{\text{reco} \cap \text{true}}}{E_{\text{true}}}, \quad (5.2)$$

which is namely the fraction of the true particle energy that is associated to the reconstructed object through the truth-matching procedure.

Another truth-matching quantity used to characterize the reconstruction quality of electromagnetic showers is the hit completeness, or

$$C_{\text{hit}} = \frac{N_{\text{matched hits}}}{N_{\text{true hits}}}, \quad (5.3)$$

which measures the fraction of hits belonging to the true particle that are included in the reconstructed PFP.

Both the C_E and C_{hit} provide complementary information on the reconstruction performance: hit completeness is related to the spatial reconstruction of the shower, while energy completeness reflects how much of the particle energy is recovered.

Baseline Pandora electron candidate selection The baseline electron candidate is selected using only standard Pandora reconstruction information. Within each slice, shower-like PFPs are selected by requiring the Pandora track-shower discriminator to satisfy

$$\text{trackScore} < 0.5. \quad (5.4)$$

To suppress poorly reconstructed fragments and low-energy clusters, additional quality criteria are applied:

$$E_{\text{shw}} > 50 \text{ MeV}, \quad N_{\text{hits}}^{\text{coll}} \geq 20, \quad (5.5)$$

where E_{shw} is the reconstructed shower energy in the collection plane and $N_{\text{hits}}^{\text{coll}}$ is the number of collection plane hits.

If multiple PFPs satisfy these criteria, the Pandora electron candidate is defined as the shower with the largest reconstructed energy,

$$\text{candidate} \equiv \underset{i \in \text{shower-like}}{\text{argmax}} E_{\text{shw},i}. \quad (5.6)$$

This choice reflects the expectation that, in ν_e charged-current interactions, the primary electron typically carries the largest fraction of the electromagnetic energy in the event.

The target electron defined in Eq. (5.1) and the baseline candidate in Eq. (5.6) define the object-level comparison used throughout the remainder of this chapter.

Slice categories used in the performance evaluation In order to define the performance metrics used throughout this study, it is useful to distinguish several categories of slices within the selected sample.

First, we define the set of *signal slices*, denoted as N_{signal} , corresponding to all slices satisfying the ν_e charged-current condition and the fiducial-volume requirement discussed above.

Among these, we define the subset of *slices with a target electron*, N_{target} , corresponding to signal slices in which at least one reconstructed PFP is truth-matched to an electron ($|\text{PDG}| = 11$). In most cases this condition is satisfied, although reconstruction

failures or extreme shower fragmentation may prevent a reconstructed PFP from being associated with the true electron.

We then define the subset of *slices with a candidate*, N_{cand} , corresponding to slices in which the Pandora-based selection described in Eq. (5.6) identifies a valid electron candidate.

Finally, we consider the subset of slices where both a target electron and a candidate are present,

$$N_{\text{cand} \cap \text{target}}, \quad (5.7)$$

which represents the set of slices where a direct candidate-target comparison is possible.

Performance metrics The quality of the candidate selection is evaluated using several metrics defined using the slice categories introduced above.

The **candidate purity** measures the probability that the selected candidate corresponds to a true electron:

$$P = \frac{N_{\text{cand true electron}}}{N_{\text{selected candidates}}}. \quad (5.8)$$

The **conditional identification efficiency** evaluates the probability that the reconstructed candidate matches the target electron in slices where both objects are present:

$$\epsilon_{\text{ID}} = \frac{N_{\text{candidate=target}}}{N_{\text{cand} \cap \text{target}}}. \quad (5.9)$$

This metric isolates the performance of the candidate identification step independently of the probability of reconstructing a candidate.

Finally, the **overall reconstruction efficiency** measures the probability of correctly identifying the electron within the full signal sample:

$$\epsilon_{\text{overall}} = \frac{N_{\text{candidate=target}}}{N_{\text{signal}}}. \quad (5.10)$$

Unlike the conditional efficiency, this quantity includes both the probability of reconstructing a valid electron candidate and the probability of correctly identifying it. Statistical uncertainties on these binomial quantities are computed using Wilson confidence intervals at 68% confidence level.

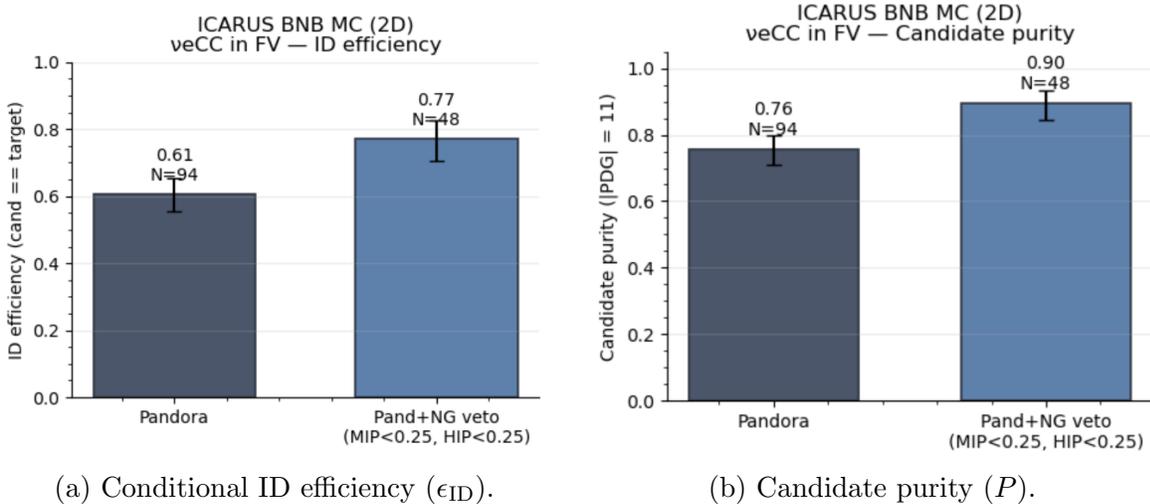


Figure 5.2: Left: conditional identification efficiency for the baseline Pandora electron candidate and after applying the NG2 semantic veto ($p_{\text{MIP}} < 0.25$ and $p_{\text{HIP}} < 0.25$), evaluated on ν_e CC signal slices with true vertex in the fiducial volume. The value of N reported above each bar corresponds to the denominator $N_{\text{cand} \cap \text{target}}$. Right: candidate purity for the same selections; here N corresponds to the number of selected candidates with defined truth match (N_{cand} with PDG_{cand} available). Error bars are binomial at 68% confidence level.

5.1 Results

This section summarizes the impact of the NG2 semantic veto on the quality of the baseline Pandora electron candidate. The comparison is performed between the nominal Pandora selection and the same selection after applying the NG2 veto based on the MIP and HIP probabilities, using the working point $p_{\text{MIP}} < 0.25$ and $p_{\text{HIP}} < 0.25$. This veto threshold was chosen among others as it was the one performing best.

Fig. 5.2a shows the conditional identification efficiency, defined as the fraction of slices in which the selected candidate matches the target electron in slices where both a candidate and a target exist (Eq. (5.1)). As said before, Uncertainties are computed using binomial confidence intervals at 68% confidence level.

Fig. 5.2, instead, summarizes the impact of the NG2 semantic veto on two complementary aspects of the electron-candidate selection. The left panel shows the *conditional identification efficiency*, defined as the fraction of slices in which the selected candidate matches the target electron in slices where both a candidate and a target exist (Eq. (5.1)). The right panel shows the *candidate purity*, defined as the fraction of selected candidates that are truth-matched to an electron.

The NG2 veto increases both the conditional identification efficiency and the fractions

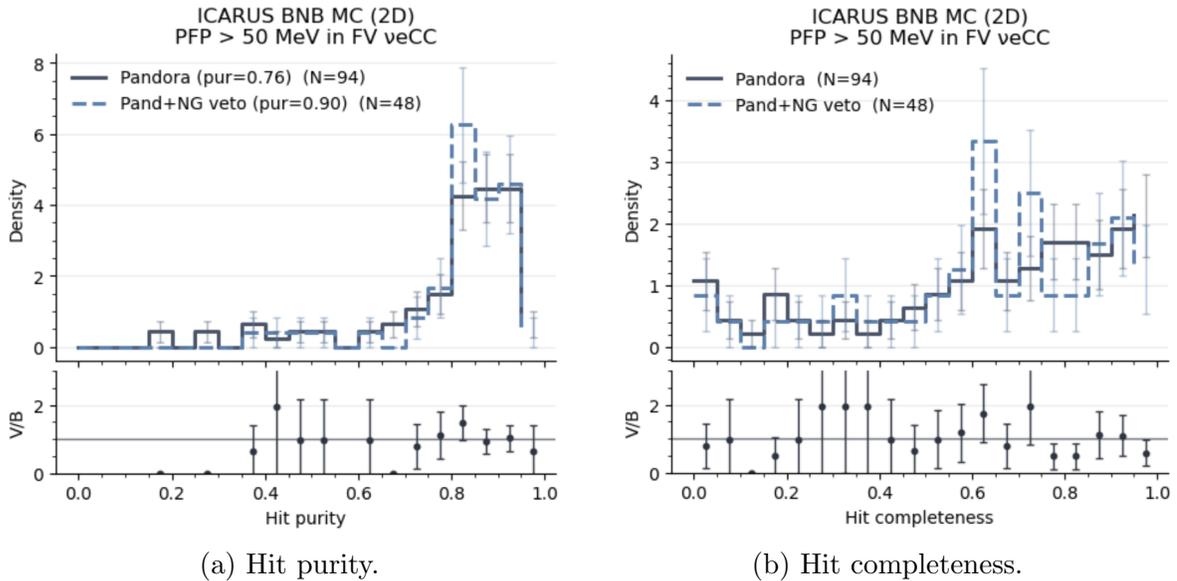


Figure 5.3: Distribution of the hit purity and completeness for the selected electron candidate in ν_e CC signal slices with true vertex in FV, comparing the baseline Pandora selection to the Pandora candidate after applying the NG2 semantic veto. The top panel shows normalized densities; the bottom panel shows the ratio (veto / baseline). The legend shows the candidate purity (truth electron fraction) for each selection.

of the candidates that are true electrons by removing track-like or hadronic candidates misidentified as electromagnetic showers. The cost is naturally reduced statistics due to the stricter selection.

Hit purity of the selected candidate Fig. 5.3a shows the distribution of the candidate hit purity, defined as the fraction of hits assigned to the candidate PFP that originate from its best-matched true particle. The comparison is performed between the baseline Pandora candidate and the same candidate after applying the NG2 semantic veto. The lower panel, instead, indicates the bin-by-bin ratio of the vetoed selection to the baseline.

The NG2 veto shifts the candidate population toward higher hit purity, which is consistent with the removal of track-like or hadronic PFPs that contaminate the Pandora shower candidate sample.

Hit completeness of the selected candidate Fig. 5.3b shows the distribution of the candidate hit completeness for the baseline Pandora selection and after applying the NG2 semantic veto. As previously, the lower panel reports the bin-by-bin ratio of the vetoed selection to the baseline.

The hit completeness distribution remains largely unchanged after the veto, since the

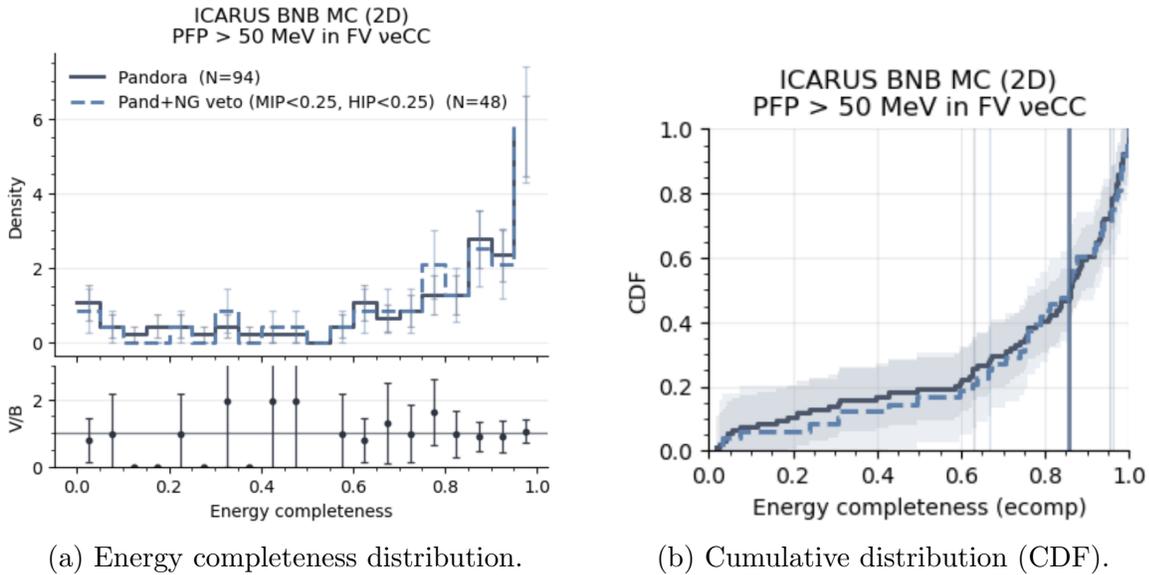


Figure 5.4: Candidate energy completeness (C^E) in ν_e CC signal slices with true vertex in FV, comparing the baseline Pandora selection to the Pandora candidate after applying the NG2 semantic veto ($p_{\text{MIP}} < 0.25$ and $p_{\text{HIP}} < 0.25$). Left: normalized density distribution with bin-by-bin ratio (veto / baseline). Right: CDF comparison (with the corresponding statistical band).

veto mainly removes wrong candidates rather than improving the containment of the reconstructed electromagnetic shower.

Energy completeness of the selected candidate Fig. 5.4 compares the baseline Pandora selection to the selection after the NG2 veto, showing both the differential distribution (with ratio panel) and the corresponding cumulative distribution function (CDF).

The energy completeness remains for the most part unchanged by the veto, since NG2 rejects wrong candidates and does not improve the reconstructed shower containment.

The net effect of the NG2 semantic veto is summarized in Fig. 5.5, which compares the baseline Pandora selection to the Pandora candidate after applying the NG2 veto ($p_{\text{MIP}} < 0.25$ and $p_{\text{HIP}} < 0.25$) across the main performance indicators introduced in Sec. 5.

The application of the NG2 semantic veto improves the quality of the selected electron candidate. In particular, the candidate purity increases from $0.755^{+0.042}_{-0.047}$ to $0.896^{+0.036}_{-0.053}$, while the conditional identification efficiency improves from $0.606^{+0.049}_{-0.051}$ to $0.771^{+0.055}_{-0.066}$. The veto retains $0.511^{+0.051}_{-0.052}$ of the baseline candidates and leads to an overall efficiency from $0.308^{+0.035}_{-0.033}$ to $0.200^{+0.031}_{-0.028}$.

This study represents a preliminary validation of the use of NG2 information at the

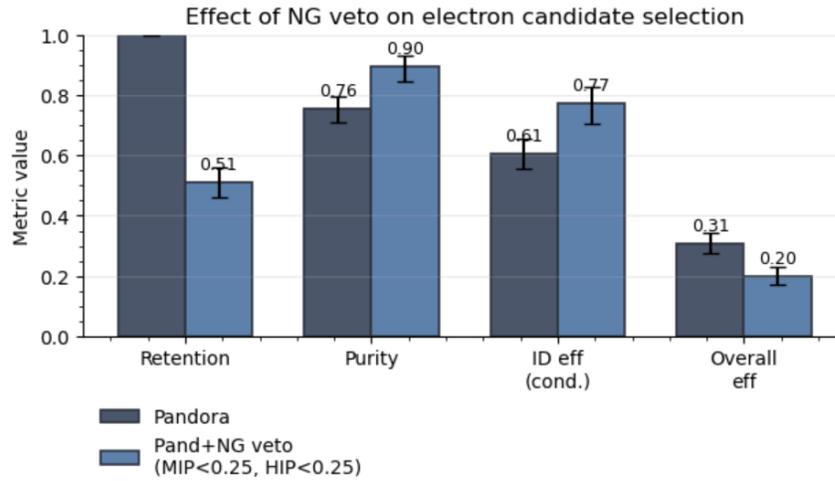


Figure 5.5: The figure compares the baseline Pandora selection to the Pandora candidate after applying the NG2 veto ($p_{\text{MIP}} < 0.25$ and $p_{\text{HIP}} < 0.25$). The metrics shown are: retention (fraction of baseline candidates surviving the veto), candidate purity, conditional identification efficiency, and overall efficiency (with denominator N_{signal}). Error bars correspond to Wilson binomial ones for confidence intervals at 68% confidence level.

end of the reconstruction chain, and, although limited to a simplified electron candidate selection, it aims to demonstrate the potential versatility of NG2 semantic outputs for improving electromagnetic shower reconstruction.

Conclusions

The work presented in this thesis addressed the problem of electromagnetic shower reconstruction in LArTPCs from two complementary perspectives, combining hardware and reconstruction-oriented developments.

On the instrumentation side, part of the work focused on the ArCS system and its data acquisition chain. In particular, the legacy LArIAT readout infrastructure was partially restored and adapted in order to prepare the system for future data-taking campaigns. This activity required recovering and validating elements of the existing DAQ chain and establishing a functioning baseline configuration for further developments.

The largest part of this thesis, however, concentrated on the reconstruction and machine-learning aspects of the ICARUS experiment. In this context, the new two-dimensional processing chain was restored and configured in order to enable training and evaluation of the NuGraph-based reconstruction framework. While functional, the workflow could be further optimized in the future, particularly with respect to processing and training times.

A detailed study of the training samples and model performance was carried out in order to identify a configuration providing stable and reliable training behaviour. The characterization of the input datasets and the resulting model outputs allowed the identification of a training processing configuration leading to the best performance during the training phase.

The trained NuGraph2 model was subsequently integrated into the ICARUS reconstruction chain and evaluated at analysis level. At present, the semantic score associated with electromagnetic showers does not provide a clear discrimination power when evaluated at the reconstructed PFP level. One possible explanation is that the model is trained and validated primarily using hit-level information, as reflected in the confusion matrices obtained during training, while the scores available in the CAF files correspond to PFP-level quantities. The procedure used to associate hit-level semantic information to reconstructed PFP objects may therefore dilute the discriminating power of the network output.

Further improvements could be achieved either by refining the method used to aggregate hit-level predictions into PFP-level scores, or by exploiting the semantic output

more directly within the reconstruction algorithms themselves. For instance, the semantic information provided by NuGraph2 could be directly incorporated into the Pandora reconstruction framework, for example by extending or retraining the existing track–shower discrimination BDT using the NuGraph2 semantic probabilities.

Despite these limitations, the application of a simple NuGraph-based veto on the Pandora electron candidate demonstrates that the semantic information already provides useful discrimination power and improves the quality of the selected electron candidate. The veto improves both the candidate purity and the conditional identification efficiency, indicating that the network is able to identify and suppress track-like or hadronic objects that are misidentified as electromagnetic showers by Pandora. As expected, the overall efficiency slightly decreases due to the stricter selection, and quantities related to shower containment, such as hit completeness and energy completeness, remain for the most part unchanged because the baseline reconstruction of the candidate is still performed by Pandora. In the ν_e CC signal slice sample (185 slices with true vertex in the fiducial volume), and using the definitions of target electron and candidate introduced in Sec. 5.1, the candidate purity increases from $0.755^{+0.042}_{-0.047}$ to $0.896^{+0.036}_{-0.053}$, while the conditional identification efficiency improves from $0.606^{+0.049}_{-0.051}$ to $0.771^{+0.055}_{-0.066}$. The veto retains $0.511^{+0.051}_{-0.052}$ of the baseline candidates and leads to a corresponding decrease of the overall reconstruction efficiency from $0.308^{+0.035}_{-0.033}$ to $0.200^{+0.031}_{-0.028}$.

This study therefore represents a preliminary validation of the use of NuGraph2 information at the end of the reconstruction chain. Although limited to a simplified electron candidate selection based on the Pandora reconstruction and evaluated through the purity and efficiency metrics defined in this work, it illustrates the potential versatility of NuGraph2 semantic outputs for improving electromagnetic shower reconstruction.

Overall, graph neural network approaches such as NuGraph2 can evidently provide valuable complementary information to traditional reconstruction algorithms in LArTPC experiments. Further developments are still needed to fully exploit their potential within the reconstruction chain, but the results obtained here indicate promising directions for improving electromagnetic shower reconstruction in current and future neutrino experiments, e.g., the current SBN program and future DUNE.

At the time of writing, steps are being taken within the ICARUS collaboration to integrate the trained NuGraph2 model into the `icaruscode` reconstruction framework. This step will allow the model to be tested and further developed by multiple collaborators, with the goal of evaluating how its semantic information can be most effectively exploited within the full reconstruction chain.

Acknowledgements

I would like to thank all those who made this thesis possible.

First of all, I would like to express my sincere gratitude to Angela Fava, Marco Del Tutto, and Giuseppe Cerati for their guidance during my time at Fermilab. Their patience, support, and ability to constantly encourage curiosity and critical thinking made this experience stimulating and rewarding.

I would also like to thank my thesis supervisor, Prof. Nicoletta Mauri, for her support throughout this work.

Finally, I would like to thank all the members of the ArCS working group, and in particular Geoff Savage for his guidance.

Desidero ringraziare la mia famiglia per la fiducia e la pazienza con cui mi hanno accompagnata durante tutto il mio percorso di studi. Un ringraziamento speciale va a Giovanni, che fin da quando ero piccola è sempre stato il mio punto di riferimento.

Ringrazio tutti gli amici che hanno condiviso con me questi anni, tra studio, entusiasmo e momenti di difficoltà.

Infine, desidero ringraziare i colleghi e amici del Collegio Superiore dell'Università di Bologna per questi anni insieme.

Appendix A

DAQ Tests for ArCS

A.0.1 Scope and Strategy

The ArCS data acquisition system inherits the general architecture of the LArIAT DAQ. Before beam operations, it is necessary to assess the status of the inherited chain and determine whether it can be reliably reused in the ArCS configuration.

The hardware and software components under study had not been operated since the conclusion of LArIAT data taking, and no a priori assumptions about the integrity of the acquisition chain could be made.

The basic system that was considered in this test study consists of:

- a DAQ server;
- a CAEN A3818 PCIe optical interface card installed in the server;
- optical transceivers linking the server to the VME crate;
- a Wiener VME crate hosting a total of eight digitizers;
- CAEN V1740 waveform digitizer boards.

The Wiener VME crate provides power distribution for the digitizer boards, while the CAEN A3818 is a PCIe optical interface card that handles high-speed data transfer between the VME crate and the DAQ server. In this configuration, the V1740 boards digitize incoming analog signals from the ArCS warm electronics and store the sampled waveforms in internal memory buffers. Upon readout, data are transmitted through the optical link to the A3818 PCIe card, then transferred to the DAQ server, where they are handled by the `artdaq` BoardReader process.

A step-by-step validation approach was used to separate possible hardware issues from DAQ software-related problems. The work proceeded through the following steps:

- (a) validation of the digitizer hardware and clock distribution;

-
- (b) analysis and restoration of the LArIAT `artdaq` control chain;
 - (c) investigation of fragment-generation failures;
 - (d) preliminary assessment of portability to AlmaLinux 9.

The hardware tests were performed independently of the `artdaq` infrastructure. Only after verifying that the front-end hardware was functioning correctly did I begin looking into restoring the LArIAT DAQ scripts.

A.0.2 Standalone Digitizer Test

Digitizer Board and Signal Path

The ArCS readout relies on CAEN V1740 waveform digitizers hosted in a VME crate. Each board samples analog input signals and stores digitized waveforms in on-board memory buffers. Data are transmitted to the DAQ computer through the optical link interface, where they are packaged into fragments by the BoardReader process. The first step was to verify that the minimal hardware acquisition chain was functioning correctly in standalone mode. Communication between the server and the V1740 board in standalone mode relies on the CAEN software stack, which includes the `VME Communication` library, the `CAENComm` interface, and the `CAENDigitizer` library. The `WaveDump` utility uses these libraries to configure the board, read out buffers, and retrieve waveform data. `CAENUpgraderGUI` was also used to check the board's internal FPGA and AD9510 status and to perform firmware upgrades. The AD9510 is the on-board clock distribution and PLL device responsible for generating and synchronizing the sampling clock used by the ADCs.

Baseline Acquisition

Initial tests were performed on a single digitizer board using the manufacturer's `WaveDump` utility. Baseline acquisition was verified with no external signal injected, and stable ADC counts and correct firmware operation were confirmed.

This test established that the board was operational at the most basic level and that data could be correctly read from the digitizer memory.

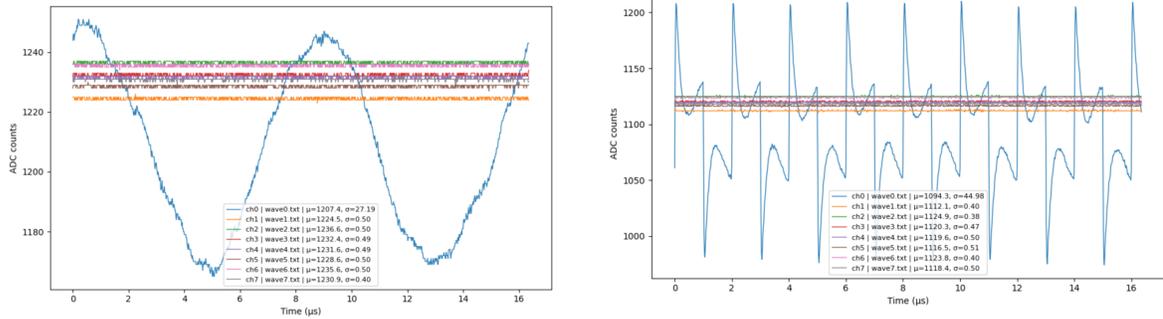
Controlled Signal Injection

To validate the full analog-to-digital response, known waveforms were injected into the digitizer inputs through an oscilloscope. The signal was injected into only one channel, the first (ch.0). The injected signals included sinusoidal signals at ~ 120 kHz and ~ 61 kHz, with an amplitude of ~ 0.25 V, as well as pulsed signals of ~ 1 μ s width and a repetition rate of ~ 0.7 MHz.

The acquired waveforms reproduced the expected frequencies and shapes in ADC

counts, thus demonstrating correct signal sampling and stability of the firmware configuration. Examples of the acquired waveforms are shown in Fig. A.1a and Fig. A.1b.

These tests were performed independently of `artdaq` and confirmed that the digitizer board and analog front-end were functioning as expected.



(a) Injected sinusoidal signal (~ 120 kHz).

(b) Injected pulsed signal.

Figure A.1: Examples of injected signals acquired with the V1740 in standalone mode, showing the expected waveform reconstruction and timing response.

Clock Distribution and EXT PLL Configuration

Multi-board operation requires proper clock synchronization. The V1740 boards allow external clock injection through the CLK-IN connector, with signal propagation via LVDS.

Direct CLK-IN injection was not initially possible due to the absence of a TTL-to-LVDS converter. Instead, the PLL on one single board was reprogrammed to enable external clock distribution and to create a daisy-chain configuration, also considering that the 8-board daisy-chain configuration is the one that LArIAT used for the final DAQ (see Fig. A.2).

A 50 MHz clock was first injected, which did not require reprogramming of the receiving board as outlined in the board manual. Subsequently, the clock frequency was changed to 60 MHz, and the receiving board PLL was reconfigured accordingly. In both cases, stable synchronization was achieved.

The clock distribution across boards is thus viable in the ArCS setup, and multi-board synchronization can be implemented.

Overall, the standalone digitizer tests confirmed that the digitizer hardware, firmware, and clock distribution are fully operational. For this reason, any subsequent data-taking failures were attributed to higher-level DAQ software or configuration issues rather than to front-end electronics not functioning.

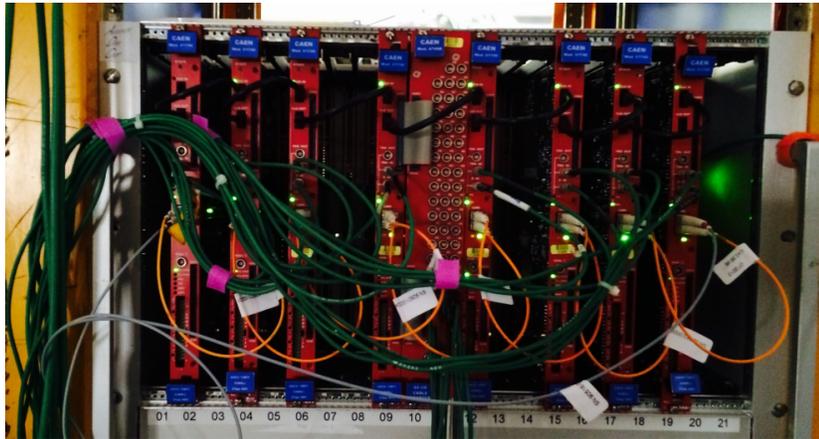


Figure A.2: LArIAT TPC readout Wiener VME crate showing the 8-board LVDS clock daisy chain (black cables) and optical fiber connections to the DAQ server (orange cables).

A.0.3 Analysis and Restoration of the LArIAT DAQ Infrastructure

Overview

The inherited LArIAT DAQ is organized in three conceptual layers:

- (a) a control layer, based on shell scripts;
- (b) an `artdaq` process layer;
- (c) a hardware readout layer responsible for fragment generation.

`artdaq` is a modular data acquisition framework developed at Fermilab that provides process management, inter-process communication, and event building, although it does not directly interface with hardware [40]. Instead, it relies on dedicated processes both to receive and assemble data fragments.

In the LArIAT configuration, the `artdaq` layer consists of three main processes:

- **BoardReader**: interfaces with the hardware readout code and receives data from the digitizers. It is the only process capable of originating data within the `artdaq` chain.
- **EventBuilder**: collects fragments from the BoardReader and assembles them into complete events.
- **Aggregator** (or DataLogger): receives complete events from the EventBuilder and writes them to disk.

Data are transmitted within `artdaq` as *fragments*, which are structured data blocks

containing detector information associated with a given trigger or spill, along with the metadata required for event building.

`artdaq` processes can start and complete state transitions correctly even in the absence of fragment production. If no fragments are generated at the hardware readout level, the EventBuilder and Aggregator simply remain idle without reporting fatal errors.

Recovery of the Control and Process Layers

The inherited system was initially non-functional due to misconfigurations and missing dependencies. To restore a stable running environment, the following steps were taken:

- restoration and verification of the processes startup script;
- validation of runtime configuration files generation and run-type injection;
- correction of environment variables required by `artdaq`;
- manual exposure of the `artdaq_mpich_plugin` to enable the process management tool;
- forced resolution of system-library compatibility issues required by `mpirun`.

After these corrections, the BoardReader, EventBuilder, and Aggregator processes could be launched, and all standard state transitions (`init`, `start`, `stop`, `shutdown`) completed successfully.

A.0.4 Persistent Zero-Fragment Condition

Despite the successful restoration of the control and process layers, no data fragments were produced during physics-mode or pedestal-mode runs. In particular:

- the BoardReader process remained active without crashing;
- the EventBuilder remained idle;
- no output `.root` files were written.

Given the architecture described above, this behavior is consistent with a failure in fragment generation rather than a failure in `artdaq` process control.

Possible causes include:

- absence of a valid external trigger (TRG-IN);
- misconfiguration of trigger or spill logic in the XML files;
- incomplete spill closure preventing fragment emission;
- mismatch in the expected number of fragments per spill.

Attempts were made to inject external triggers, to modify configuration files, and to enable software-triggered acquisition. However, no fragments were produced under the tested configurations.

At this stage, the most probable origin of the issue lies in the trigger-handling or spill-termination logic within the hardware readout layer.

A.0.5 Development Environment and Portability

In parallel with runtime debugging, an attempt was made to restore the original build environment of LArIAT, namely the `lariat-online` codebase. A dedicated setup script was written to define a clean development environment separate from the operational DAQ area.

Although the DAQ stack could be executed, the rebuild of `lariat-online` was not successful due to missing UPS products and incompatible qualifiers. Therefore, the current system can be run but cannot be cleanly rebuilt with its inherited configuration.

Finally, the planned ArCS DAQ computer is expected to operate under AlmaLinux 9. Unlike the legacy Scientific Linux 7 environment, AlmaLinux 9 does not provide native support for the UPS product management system on which the LArIAT stack depends.

While the overall `artdaq` architecture can, in principle, be preserved, the build and product-management layer would require substantial restructuring to ensure compatibility with AlmaLinux 9.

In summary, the digitizer hardware and clock distribution were validated independently of the DAQ software. The inherited LArIAT control and `artdaq` layers were successfully restored to a runnable state. However, fragment generation remains non-functional under the tested configurations, and the build infrastructure is not currently portable to modern operating systems. These preliminary studies summarize the current status of the inherited DAQ and the work still required to achieve stable ArCS beam operation.

Bibliography

- [1] Enrico Fermi. Tentativo di una teoria di raggi β . *Rendiconti dell'Accademia Nazionale dei Lincei*, 19:756–758, 1934.
- [2] F. Reines and C. L. Cowan. Detection of the free neutrino. *Phys. Rev.*, 92:830–831, Nov 1953.
- [3] Clyde L Cowan Jr, Frederick Reines, FB Harrison, HW Kruse, and AD McGuire. Detection of the free neutrino: a confirmation. *Science*, 124(3212):103–104, 1956.
- [4] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental test of parity conservation in beta decay. *Phys. Rev.*, 105:1413–1415, Feb 1957.
- [5] M. Goldhaber, L. Grodzins, and A. W. Sunyar. Helicity of neutrinos. *Phys. Rev.*, 109:1015–1017, Feb 1958.
- [6] A. Bellerive, J.R. Klein, A.B. McDonald, A.J. Noble, and A.W.P. Poon. The sudbury neutrino observatory. *Nuclear Physics B*, 908:30–51, 2016. Neutrino Oscillations: Celebrating the Nobel Prize in Physics 2015.
- [7] Sin Kyu Kang. Theoretical status of neutrino oscillation physics. In *Proceedings of the Conference*, page 002, 06 2020.
- [8] Carlo Giunti and Chung W. Kim. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, Oxford, UK, 2007.
- [9] G. Mention, M. Fechner, Th. Lasserre, Th. A. Mueller, D. Lhuillier, M. Cribier, and A. Letourneau. Reactor antineutrino anomaly. *Phys. Rev. D*, 83:073006, Apr 2011.
- [10] Patrick Huber. Determination of antineutrino spectra from nuclear reactors. *Phys. Rev. C*, 84:024617, Aug 2011.
- [11] F. P. *et al.* (Daya Bay Collaboration) An. Evolution of the reactor antineutrino flux and spectrum at daya bay. *Phys. Rev. Lett.*, 118:251801, Jun 2017.
- [12] GALLEX Collaboration. Results of the calibration of the gallex solar neutrino detector. *Phys. Lett. B*, 342:440–450, 1995.

-
- [13] SAGE Collaboration. Measurement of the response of a gallium metal solar neutrino experiment to neutrinos from a ^{37}Ar source. *Phys. Rev. C*, 80:015807, 2009.
- [14] BEST Collaboration. Results from the baksan experiment on sterile transitions (best). *Phys. Rev. Lett.*, 128:232501, 2022.
- [15] LSND Collaboration. Evidence for neutrino oscillations from the observation of anti-neutrino appearance. *Phys. Rev. D*, 64:112007, 2001.
- [16] MiniBooNE Collaboration. Significant excess of electronlike events in the mini-boone short-baseline neutrino experiment. *Phys. Rev. Lett.*, 121:221801, 2018.
- [17] Carlo Giunti and Thierry Lasserre. eV -scale sterile neutrinos. *Annual Review of Nuclear and Particle Science*, 69(1):163–190, October 2019.
- [18] R. Acciarri et al. A proposal for a three detector short-baseline neutrino oscillation program in the fermilab booster neutrino beam, 2015.
- [19] Anthony P. Wood. *The NuMI Neutrino Flux Prediction at ICARUS*. PhD thesis, Houston U., Houston U., 12 2024.
- [20] The MicroBooNE Collaboration. Search for light sterile neutrinos with two neutrino beams at microboone. *Nature*, 648:64–69, 2025.
- [21] S. Navas and et al. Review of particle physics. *Phys. Rev. D*, 110:030001, Aug 2024.
- [22] C. Rubbia. The Liquid Argon Time Projection Chamber: A New Concept for Neutrino Detectors. 5 1977.
- [23] W. M. Bonivento and F. Terranova. The science and technology of liquid argon detectors. *Rev. Mod. Phys.*, 96:045001, Oct 2024.
- [24] S. Amerio, , and et al. Design, construction and tests of the icarus t600 detector. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 527(3):329–410, 2004.
- [25] A. Abed et al. Abud. Neutrino interaction vertex reconstruction in dune with pandora deep learning. *The European Physical Journal C*, 85(6), June 2025.
- [26] W. et al. Foreman. Calorimetry for low-energy electrons using charge and light in liquid argon. *Physical Review D*, 101(1), January 2020.
- [27] W. Walkowiak. Drift velocity of free electrons in liquid argon. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 449(1):288–294, 2000.
- [28] R. Andrews, W. Jaskierny, H. Jostlein, C. Kendziora, S. Pordes, and T. Tope. A system to test the effects of materials on the electron drift lifetime in liquid argon

- and observations on the effect of water. *Nucl. Instrum. Meth. A*, 608:251–258, 2009.
- [29] Y. Li, C. Bromberg, M. Diwan, S. Kettell, S. Martynenko, X. Qian, V. Paolone, J. Stewart, C. Thorn, and C. Zhang. Parameterization of electron attachment rate constants for impurities in lartpc detectors. *Journal of Instrumentation*, 17(11):T11007, nov 2022.
- [30] S Amoruso et al. Analysis of the liquid argon purity in the icarus t600 tpc. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 516(1):68–79, 2004.
- [31] M. et al Antonello. Experimental observation of an extremely high electron lifetime with the icarus-t600 lar-tpc. *Journal of Instrumentation*, 9(12):P12006–P12006, December 2014.
- [32] Akira Hitachi, Tan Takahashi, Nobutaka Funayama, Kimiaki Masuda, Jun Kikuchi, and Tadayoshi Doke. Effect of ionization density on the time dependence of luminescence from liquid argon and xenon. *Phys. Rev. B*, 27:5279–5285, 1983.
- [33] P. Cennini and et al. Performance of a three-ton liquid argon time projection chamber. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 345(2):230–243, 1994.
- [34] Francesco Poppi. Status and perspective of icarus-t600 detector at the fermilab short-baseline neutrino program. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1072:170113, 2025.
- [35] R. et al. Acciarri. The liquid argon in a testbeam (lariat) experiment. *Journal of Instrumentation*, 15(04):P04026, apr 2020.
- [36] Avinay Bhat, Mun Jung Jung, Gray Putnam, and Haiwang Yu. Dnn-based signal processing for liquid argon time projection chambers, 2025.
- [37] A. Aurisano, V. Hewes, G. Cerati, J. Kowalkowski, C.S. Lee, W. Liao, D. Grzenda, K. Gumpula, and X. Zhang. Graph neural network for neutrino physics event reconstruction. *Physical Review D*, 110(3), August 2024.
- [38] C. Andreopoulos, A. Bell, D. Bhattacharya, F. Cavanna, J. Dobson, S. Dytman, H. Gallagher, P. Guzowski, R. Hatcher, P. Kehayias, A. Meregaglia, D. Naples, G. Pearce, A. Rubbia, M. Whalley, and T. Yang. The genie neutrino monte carlo generator. *Nuclear Instruments and Methods in Physics Research Section A*, 614:87–104, 2010.

- [39] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, et al. Geant4 – a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A*, 506:250–303, 2003.
- [40] Kurt Biery, Eric Flumerfelt, John Freeman, Wesley Ketchum, Gennadiy Lukhanin, and Ron Rechenmacher. artdaq: Daq software development made simple. *Journal of Physics: Conference Series*, 898(3):032013, oct 2017.