UNIVERSIDAD AUTÓNOMA DE SAN LUIS POTOSÍ

INSTITUTO DE FÍSICA

Studies of $\pi^+ \rightarrow e^+ +$ invisible topologies in the NA62 experiment

TESIS

QUE PARA OBTENER EL GRADO DE:

DOCTOR EN CIENCIAS (FÍSICA)

PRESENTA:

M.C. Alejandro Briano Olvera

ASESOR DE TESIS

Dr. Jürgen Engelfried

San Luis Potosí, S.L.P.

August 26, 2026

Abstract

A search for Heavy Neutral Leptons (HNL) in the decay $\pi^+ \to e^+ N$ is presented using the data set collected by the NA62 experiment in the period 2017-2024. The analysis scans 63 distinct HNL mass hypotheses in the range 95–126 MeV/c^2, assuming a lifetime exceeding 50 ns. The observed data are consistent with the background-only hypothesis across the entire mass range, resulting in upper limits on the mixing matrix element $|U_{e4}|^2$ at the 90% confidence level of the order of 10^{-8}. The same analysis framework was applied to perform studies of the R_π ratio $(\Gamma(\pi_{e2})/\Gamma(\pi_{\mu 2}))$ using the common event sample, where both decay modes are identified via their characteristic decay signatures. This work provides a statistical benchmark of the measurement precision; the computed ratio, considering only statistical uncertainties, establishes the intrinsic statistical resolution of the dataset pending future systematic evaluation.

Acknowledgements

This work would not have been possible without the invaluable direction and endless patience of **Dr. Jürgen Engelfried** throughout the research process.

I also want to give a huge thanks to Nora Estrada for her advice and support, and for always letting me take a week off at the end when scheduling shifts (though, honestly, I just used that time to catch up on sleep).

To all my Mexican colleagues—from undergrads to grad students—Akbar, Alan, Kevs, Erick, Didi, Tom, Erika, Tona, Emilio, Maria, and Fer: thanks for hanging out with me and making this whole journey so much more fun.

To L.E.S.D. José Limón Castillo, for his work in the High Energy Lab, and to I.E. Luz del Carmen Nuche Garza, the lab technician, for all their help.

Finally, a massive thank you to Maricela, the grad program secretary, and Gustavo, who runs the Physics Institute library, for everything they did to help out.

This work was supported by the 'Fondo Sectorial de Investigación para la Educación SEP/CONACyT', project 242139, and by the project 'Participación de México en la Frontera de Física de Altas Energías en el CERN', CONACyT Frontier Projects, number 2042.

Contents

Introduction

The discovery of neutrino oscillations, the phenomenon where neutrinos change flavor as they travel and thus must have mass, is the only direct experimental evidence of physics beyond the Standard Model. The fundamental nature of this mass remains unknown, as neutrinos could be Dirac particles with distinct antiparticles or Majorana particles where the neutrino is its own antiparticle. The Type I seesaw mechanism offers a theoretical explanation by introducing heavy right handed neutrinos across a wide range of mass scales. While traditionally associated with energy scales far above the electroweak scale, such as the Grand Unification scale, the framework also supports low mass realizations like the Neutrino Minimal Standard Model (νMSM), where Heavy Neutral Leptons (HNLs) could exist in the keV to GeV range, making them detectable in experiments studying pion decays.

Leptonic pion decays offer a clean environment to search for these particles. The Standard Model process where a pion decays to a positron and a neutrino is heavily suppressed by helicity conservation. This makes the decay highly sensitive to new physics, such as the production of a massive HNL. Additionally, the ratio of decay widths R_π, comparing the decay to an electron versus a muon, serves as a strict test of lepton universality, the principle that electroweak interactions are identical for all lepton flavors. Any deviation in this ratio would indicate new physics, while agreement helps validate the experimental approach.

This thesis presents an analysis of pion decays using data from the NA62 experiment at CERN. The experiment uses an unseparated secondary hadron beam with a mean momentum of 75 GeV/c, composed mainly of pions (70%), protons (23%), and kaons (6%). This high pion flux allows for precise studies of rare decays. The work has two main objectives. The primary goal is a peak search for HNLs in the pion to positron plus HNL decay channel across 63 distinct mass hypotheses in the range of 95 to 126 MeV/c^2, assuming a lifetime exceeding 50 ns. The results are used to set upper limits on the mixing parameter $|U_e4|^2$ based on the observed data. Complementing this search, the same analysis framework is applied to the pion to electron neutrino and pion to muon neutrino channels to conduct studies of the R_π ratio. These studies evaluate the statistical precision achievable with the current dataset, providing a benchmark for the intrinsic resolution of the measurement as a test of lepton universality, pending future systematic evaluations. Following a description of the theoretical framework and the NA62 experimental apparatus, the thesis details the HNL search strategy and results, followed by the methodology and statistical findings of the R_π analysis.

Chapter 1

Theoretical Framework and Motivation

1.1 The Nature of Neutrino Mass

1.1.1 Neutrinos in the Standard Model

The Standard Model (SM) is a gauge theory based on the local symmetry group $SU(3)_C \times SU(2)_L \times U(1)_Y$, describing the strong, weak, and electromagnetic interactions of elementary particles. Spontaneous symmetry breaking reduces this framework to the subgroup $SU(3)_C \times U(1)_{\mathrm{EM}}$. The electromagnetic charge is defined by:

$$Q_{\mathrm{EM}} = T_{L3} + Y$$

where T_{L3} is the third generator of weak isospin spanning the $SU(2)_L$ group, and Y is the weak hypercharge.

Neutrinos are assigned as color singlets under $SU(3)_C$ and carry no electromagnetic charge under $U(1)_{\mathrm{EM}}$. They are components of the left-chiral lepton doublets:

$$L_{L\ell} = \begin{pmatrix} \nu_\ell \\ \ell \end{pmatrix}_L$$

Here, the index $\ell = e, \mu, \tau$ denotes the three generations of charged leptons. The left-chiral fields are isolated using the projection operator $P_L = \frac{1}{2}(1 - \gamma^5)$. Neutrinos that are part of these doublets are classified as active neutrinos.

The requirement of local $SU(2)_L$ gauge invariance fixes the mathematical structure of the weak interactions, defining the charged-current (CC) interactions between neutrinos and their corresponding charged leptons and the neutral-current (NC) interactions among the neutrinos themselves as:

$$-\mathcal{L}_{\mathrm{CC}} = \frac{g}{\sqrt{2}} \sum_\ell \bar{\nu}_{L\ell} \gamma^\mu \ell_L W_\mu^+ + \text{h.c.}$$

$$-\mathcal{L}_{\mathrm{NC}} = \frac{g}{2\cos\theta_W} \sum_\ell \bar{\nu}_{L\ell} \gamma^\mu \nu_{L\ell} Z_\mu^0$$

where g represents the $SU(2)_L$ gauge coupling constant, and θ_W represents the Weinberg angle.

The neutral-current interaction determines the partial decay width of the Z boson into light neutrino states ($m_\nu \leq m_Z/2$). High-precision measurements constrain the number of active neutrino generations to [1]:

$$N_\nu = 2.9963 \pm 0.0074$$

Consequently, any extensions preserving standard electroweak couplings are restricted to three light active neutrino states. Sterile neutrinos are defined as fields that carry no gauge quantum numbers under the full SM group. The minimal SM contains no such sterile states. The field content of the SM yields an accidental global symmetry group:

$$G_{\rm SM}^{\rm global} = U(1)_B \times U(1)_{L_e} \times U(1)_{L_\mu} \times U(1)_{L_\tau}$$

where $U(1)_B$ is the baryon number conservation symmetry, and $U(1)_{L_\ell}$ represents the individual conservation laws for lepton flavors. Total lepton number ($L = L_e + L_\mu + L_\tau$) is preserved as a subgroup of $G_{\rm SM}^{\rm global}$.

Fermion masses arise when left-handed and right-handed particles interact with the Higgs field through specific coupling strengths. For charged leptons, this interaction is described by:

$$-\mathcal{L}_{\text{Yukawa, lep}} = Y_{ij}^\ell \bar{L}_{Li} \phi E_{Rj} + \text{h.c.}$$

where Y_{ij} are the Yukawa couplings setting the interaction strength, L_{Li} is the left-chiral lepton doublet (containing the neutrino and charged lepton), ϕ is the Higgs doublet, and E_{Rj} is the right-chiral charged lepton singlet.

Spontaneous symmetry breaking yields the mass matrices for the charged leptons:

$$m_{ij}^\ell = Y_{ij}^\ell \frac{v}{\sqrt{2}}$$

where v is the vacuum expectation value of the Higgs field. Because right-chiral neutrino singlets are omitted from the SM, an equivalent gauge invariant Yukawa term for the neutrino sector cannot be formed.

In principle, a neutrino mass term could be generated at the loop level. Within the framework of the SM particle content, the only possible neutrino mass term that can be constructed is the bilinear $\bar{L}_L L_L^c$, where $L_L^c = C\bar{L}_L^T$ represents the charge conjugated field and C represents the charge conjugation matrix. However, this term is strictly forbidden in the SM because it violates the total lepton symmetry by two units ($\Delta L = 2$). Consequently, it cannot be induced by loop corrections, as doing so would break this accidental global symmetry of the model. Furthermore, because $U(1)_{B-L}$ is a non-anomalous subgroup of $G_{\rm SM}^{\rm global}$, the $\bar{L}_L L_L^c$ bilinear cannot be induced by non-perturbative corrections either, since it explicitly breaks $B - L$ symmetry.

Neutrinos are strictly massless within the minimal Standard Model framework. A non-zero neutrino mass requires extensions to the particle content or gauge structure of the theory.

1.1.2 Experimental Evidence for Mass and Oscillations

1.1.2.1 Experimental Evidence

The foundational experimental evidence for non-zero neutrino masses stems from a persistent discrepancy between the theoretical predictions of astrophysical neutrino flux models and direct laboratory observations. The first empirical indication of this anomaly was established in the late 1960s by the Homestake experiment [2]. The primary objective of this experiment was to measure the flux of electron neutrinos (ν_e) produced by nuclear fusion processes within the solar core, as quantified by the Standard Solar Model (SSM) [3]. Utilizing a radiochemical detector filled with perchloroethylene (C_2Cl_4), the experiment successfully detected solar neutrinos via the charged-current capture reaction:

$$^{37}\text{Cl} + \nu_e \rightarrow {}^{37}\text{Ar} + e^-$$

However, the measured rate of neutrino interactions was consistently approximately one-third of the flux predicted by the theoretical model [4]. This significant deficit between experimental measurement and astrophysical theory became widely known as the solar neutrino problem.

For decades, it remained unclear whether this deficit was due to an incomplete understanding of solar core physics, systematic experimental errors, or unknown neutrino properties. The ambiguity was definitively resolved in 2001 by the Sudbury Neutrino Observatory (SNO) [5]. SNO utilized a heavy-water (D_2O) Cherenkov detector, which allowed for the simultaneous measurement of two distinct channels: a charged-current (CC) interaction sensitive exclusively to electron neutrinos, and a neutral-current (NC) interaction sensitive to all active neutrino flavors (ν_e, ν_μ, ν_τ).

The CC measurement confirmed the expected deficit of ν_e, while the NC measurement demonstrated that the total flux of all active neutrino flavors was in precise agreement with the Standard Solar Model predictions. This provided unambiguous proof that electron neutrinos produced in the core of the Sun were actively transitioning into muon and tau neutrino flavors during transit to Earth.

Parallel to the resolution of the solar neutrino problem, a complementary anomaly was observed in the study of cosmic-ray interactions within the Earth's atmosphere. Primary cosmic rays striking the upper atmosphere produce cascades of secondary mesons (primarily pions and kaons) that decay into leptons. Kinematically, this decay chain yields an expected ratio of approximately two atmospheric muon neutrinos to every one electron neutrino:

$$\frac{\nu_\mu + \bar{\nu}_\mu}{\nu_e + \bar{\nu}_e} \approx 2$$

In 1998, the Super-Kamiokande experiment reported a definitive anomaly in this ratio [6]. Using a large underground water Cherenkov detector, Super-Kamiokande measured the atmospheric neutrino flux as a function of the zenith angle. While the flux of electron neutrinos matched isotropic expectations, the flux of upward-going muon neutrinos which traverse the entire diameter of the Earth before reaching the detector was significantly depleted relative to the downward-going muon neutrinos, which travel only a few kilometers through the local atmosphere.

This asymmetry conclusively demonstrated that the missing ν_μ were disappearing as a function of their propagation baseline, pointing directly to flavor transitions into a third, unmeasured state (ν_τ). Together, the solar and atmospheric deficits firmly established that neutrino flavors are not conserved quantities during long-baseline propagation, providing the first definitive evidence of physics beyond the minimal Standard Model.

1.1.2.2 Neutrino Oscillations and the PMNS Matrix

To reconcile the flavor changing phenomena observed by experiments like SNO and Super-Kamiokande, the Standard Model must be extended to decouple the fields responsible for weak gauge interactions from those responsible for free space propagation. This requires treating the flavor eigenstates which dictate neutrino production and detection as linear superpositions of distinct mass eigenstates.

Neutrino Oscillations

Neutrinos are produced via the charged-current weak interaction in association with a charged lepton of specific flavor $\ell \in \{e, \mu, \tau\}$. A neutrino of flavor eigenstate $|\nu_\ell\rangle$ can be expressed as a unitary transformation of the physical mass eigenstates $|\nu_i\rangle$ possessing definite mass eigenvalues m_i ($i = 1, 2, 3$):

$$|\nu_\ell\rangle = \sum_i U^*_{\ell i} |\nu_i\rangle$$

where U is a unitary mixing matrix.

Assuming a plane-wave approximation for ultra-relativistic neutrinos propagating through a vacuum, the time evolution of a mass eigenstate $|\nu_i(t)\rangle$ is governed by the time-dependent Schrödinger equation, resulting in:

$$|\nu_i(t)\rangle = e^{-iE_i t} |\nu_i(0)\rangle$$

Here, $E_i = \sqrt{p_i^2 + m_i^2}$ represents the energy eigenvalue of the i-th mass state with momentum p_i. In the relativistic limit ($p_i \simeq p_j \equiv p$), the energy can be expanded via a Taylor series to first order:

$$E_i \simeq p_i + \frac{m_i^2}{2p_i} \simeq E + \frac{m_i^2}{2E}$$

where $E \simeq p$ is the average neutrino energy. After traveling a distance L ($L \simeq ct$ for relativistic neutrinos), the state vector evolves as:

$$|\nu_\ell(L)\rangle = \sum_i U^*_{\ell i} e^{-i \frac{m_i^2 L}{2E}} |\nu_i\rangle$$

By applying the inverse unitary relation $|\nu_i\rangle = \sum_{\ell'} U_{\ell' i} |\nu_{\ell'}\rangle$, the transition amplitude $\mathcal{A}_{\ell \to \ell'}$ for a neutrino of flavor ℓ to be detected as flavor ℓ' after traveling a distance L is given by the inner product:

$$\mathcal{A}_{\ell \to \ell'} = \langle \nu_{\ell'} | \nu_\ell(L) \rangle = \sum_i U^*_{\ell i} U_{\ell' i} e^{-i \frac{m_i^2 L}{2E}}$$

The physical transition probability $P_{\ell \to \ell'}$ is the squared modulus of this amplitude ($|\mathcal{A}_{\ell \to \ell'}|^2$). Expanding this expression isolates the distinct oscillatory interference terms:

$$P_{\ell \to \ell'} = \delta_{\ell \ell'} - 4 \sum_{i>j} \mathrm{Re}\left(U^*_{\ell i} U_{\ell' i} U_{\ell j} U^*_{\ell' j}\right) \sin^2\left(\frac{\Delta m^2_{ij} L}{4E}\right) + 2 \sum_{i>j} \mathrm{Im}\left(U^*_{\ell i} U_{\ell' i} U_{\ell j} U^*_{\ell' j}\right) \sin\left(\frac{\Delta m^2_{ij} L}{2E}\right)$$

where $\Delta m^2_{ij} \equiv m_i^2 - m_j^2$ represents the mass-squared splitting between the respective mass eigenstates. This equation demonstrates that neutrino flavor transitions cannot occur unless two conditions are simultaneously satisfied: there must be kinematic mixing between the bases ($U \neq \mathbb{I}$), and at least two mass eigenstates must be non-degenerate ($\Delta m^2_{ij} \neq 0$).

PMNS Matrix

The 3×3 unitary mixing matrix U is designated as the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix. A general 3×3 unitary matrix contains nine independent physical parameters: three mixing angles and six phases. However, five of these phases can be absorbed by rephasing the left-handed lepton fields without altering the physical observables, leaving a single, irreducible Dirac CP-violating phase (δ_{CP}).

The standard parameterization adopted by the Particle Data Group (PDG) decomposes the PMNS matrix into a product of three independent spatial rotations, each corresponding to an experimental mixing sector:

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

Where $c_{ij} \equiv \cos\theta_{ij}$ and $s_{ij} \equiv \sin\theta_{ij}$ represent the cosines and sines of the mixing angles θ_{12}, θ_{23}, and θ_{13}. Explicitly multiplying these matrices gives the standard form:

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

Note that while the standard PMNS parameterization assumes a Dirac nature, standard flavor oscillations are fundamentally insensitive to whether neutrinos are Dirac or Majorana fields. If neutrinos are fundamentally Majorana particles rather than Dirac particles, the lepton fields cannot undergo the same arbitrary rephasing without destroying the real Majorana mass terms. Consequently, two additional physical phases remain unabsorbed, and the matrix must be modified by a diagonal phase factor:

$$U_{\mathrm{Majorana}} = U \times \mathrm{diag}\left(1, e^{i\eta_1/2}, e^{i\eta_2/2}\right)$$

where η_1 and η_2 are the Majorana CP-violating phases. Notably, because these Majorana phases factor out linearly when computing the transition amplitude $\mathcal{A}_{\ell\to\ell'}$, they cancel entirely in the calculation of $P_{\ell\to\ell'}$. Thus, standard flavor oscillation probabilities remain entirely insensitive to the underlying Dirac or Majorana neutrino nature, a physical distinction that will be formalized in Section 1.1.4.

1.1.3 The Neutrino Mass Hierarchy

While the observation of flavor transitions conclusively demonstrates that neutrinos possess non-zero masses, the preservation of the oscillation probability equations restricts experimental sensitivity exclusively to the differences of the squared masses ($\Delta m^2_{ij} \equiv m_i^2 - m_j^2$). Consequently, oscillation data alone cannot determine the absolute mass scale of the neutrino eigenstates (m_1, m_2, m_3), nor can it completely resolve the relative ordering of these states.

Global fits of solar [7] and reactor neutrino [8][9] data establish a definitive positive sign for the solar mass splitting, $\Delta m^2_{21} = (7.60 \pm 0.17) \times 10^{-5}$ eV^2 which mathematically mandates that $m_2 > m_1$.

However, long-baseline atmospheric and accelerator experiments only constrain the absolute magnitude of the atmospheric mass splitting, $|\Delta m^2_{31}| \approx 2.5 \times 10^{-3}$ eV^2 [10]. Because the sign of Δm^2_{31} remains undetermined, two distinct physical scenarios emerge for the spectrum of the mass eigenstates:

- Normal Ordering (NO): Governed by a positive sign configuration ($\Delta m^2_{31} > 0$). In this scenario, the mass eigenstates follow a hierarchical progression where the two closest mass states are the lightest, yielding $m_1 < m_2 < m_3$.
- Inverted Ordering (IO): Governed by a negative sign configuration ($\Delta m^2_{31} < 0$). In this scenario, a solitary light mass state sits decoupled below two heavily degenerate, heavier mass states, yielding $m_3 < m_1 < m_2$.

Resolving this sign ambiguity represents a critical milestone in modern particle physics, as the underlying ordering fundamentally alters the parameter space of cosmic structure formation and dictates the sensitivity limits required for next-generation lepton-number-violation probes.

Beyond the ordering of the states, the absolute mass scale of the lightest eigenstate (m_{lightest}) represents an independent boundary condition that oscillations leave entirely unconstrained ($m_{\text{lightest}} \to 0$ remains mathematically viable within oscillation frameworks). To bound the absolute scale, non-oscillatory kinematic probes must be perform. The most stringent direct laboratory constraint on the neutrino mass comes from high-precision spectroscopy of the molecular tritium beta-decay endpoint:

$$\mathrm{T}_2 \to {}^3\mathrm{HeT}^+ + e^- + \bar{\nu}_e$$

This process probes the effective electron anti-neutrino mass, defined as:

$$m_\nu = \sqrt{\sum_{i=1}^{3} |U_{ei}|^2 m_i^2}$$

By analyzing the spectral distortion near the endpoint, the KATRIN collaboration has established an upper bound of $m_\nu < 0.45\,\text{eV}$ (90% C.L.) [11], firmly restricting the absolute mass eigenvalues to the sub-eV regime.

1.1.4 The Fundamental Nature of Neutrino Mass: Dirac vs. Majorana

While the physical mass spectrum (m_1, m_2, m_3) and its ordering configurations are definitively constrained by the oscillation metrics detailed in Section 1.1.2, the fundamental nature of the neutrino mass term remains an open question. Because neutrinos are the only fundamentally neutral fermions in the theory, lacking both color and electric charge, their mass terms can be accommodated via two distinct field-theoretic formulations: Dirac mass terms, which conserve lepton number and require distinct antiparticles, or Majorana mass terms, which violate lepton number and imply the neutrino is its own antiparticle.

1.1.4.1 Dirac Mass Formulation

To generate neutrino masses in a manner identical to the charged leptons defined in Equation 1.1, the minimal Standard Model particle spectrum must be extended by introducing three right-handed neutrino fields ν_{Ri}, which transform as total singlets under the $SU(3)_C \times SU(2)_L \times U(1)_Y$ gauge group. With these fields present, a new leptonic Yukawa Lagrangian can be constructed using the conjugate Higgs doublet ($\tilde{\phi} \equiv i\tau_2\phi^*$):

$$-\mathcal{L}_{\text{Yukawa},\nu} = Y^\nu_{ij}\bar{L}_{Li}\tilde{\phi}\nu_{Rj} + \text{h.c.}$$

Following spontaneous symmetry breaking, this operator yields an explicit Dirac mass term:

$$-\mathcal{L}_{\mathrm{Dirac}} = m^D_{ij}\bar{\nu}_{Li}\nu_{Rj} + \mathrm{h.c.}$$

where the Dirac mass matrix is defined by $m^D_{ij} = Y^\nu_{ij}\frac{v}{\sqrt{2}}$.

The primary theoretical objection to a purely Dirac paradigm is the exceptional fine-tuning required of the coupling eigenvalues. To reconcile this framework with the absolute sub-eV mass scale established by kinematic endpoint limits (Section 1.2.3), the Yukawa couplings must be restricted to the regime of $Y^\nu \sim \mathcal{O}(10^{-12})$. This is roughly six orders of magnitude smaller than the coupling of the electron, representing an unnatural hierarchy in the fermion sector.

Phenomenologically, however, $\mathcal{L}_{\mathrm{Dirac}}$ remains invariant under a global $U(1)$ phase transformation applied to all lepton fields, ensuring that total Lepton Number (L) is a strictly conserved global quantum number.

1.1.4.2 Majorana Mass Formulation

Because the right-handed fields ν_R carry no gauge charges of any kind, they are not mathematically confined to couple exclusively to the left-handed sector via the Higgs field. A gauge-invariant Majorana mass term can be written using only the right-handed field and its charge-conjugate field ($\nu_R^c \equiv C\bar{\nu}_R^T$, where C is the charge conjugation matrix):

$$-\mathcal{L}_{\mathrm{R}} = \frac{1}{2}M_{Rij}\bar{\nu}^c_{Ri}\nu_{Rj} + \mathrm{h.c.}$$

Here, M_R is a symmetric mass matrix. Crucially, because M_R does not originate from electroweak symmetry breaking, its mass scale is not bound to the electroweak vacuum expectation value v and can be arbitrarily large.

Similarly, one can construct an effective Majorana mass term using only the active left-handed fields via a higher-dimensional dimension-5 operator (the Weinberg operator):

$$-\mathcal{L}_{\mathrm{L}} = \frac{1}{2}m^L_{ij}\bar{\nu}_{Li}\nu^c_{Lj} + \mathrm{h.c.}$$

A physical Majorana fermion is defined by the reality condition $\nu = \nu^c$, meaning the particle is identical to its own antiparticle. Because these bilinear operators couple fields to their own charge conjugates, they explicitly break the global $U(1)_L$ symmetry by two units ($\Delta L = 2$). Total lepton number is thus violated, and the distinction between matter and antimatter in the neutrino sector is eliminated.

1.1.4.3 Phenomenological Implications: Neutrinoless Double Beta Decay

The definitive experimental avenue to distinguish between the Dirac and Majorana paradigms is the search for neutrinoless double beta decay ($0\nu\beta\beta$). In certain even-even nuclei where standard single beta decay is kinematically forbidden, a second-order weak process can occur. If neutrinos are Majorana particles, the virtual neutrino emitted at one weak vertex can be absorbed as an antineutrino at the second vertex, resulting in a collective nuclear transition that violates lepton number conservation by two units:

$$(Z, A) \to (Z + 2, A) + 2e^-$$

The inverse half-life of this process scales quadratically with the effective Majorana mass parameter $\langle m_{\beta\beta}\rangle$:

$$\langle m_{\beta\beta}\rangle = \left|\sum_{i=1}^{3} U_{ei}^2 m_i\right| = \left|c_{12}^2 c_{13}^2 m_1 + s_{12}^2 c_{13}^2 m_2 e^{i\eta_1} + s_{13}^2 m_3 e^{i(\eta_2 - 2\delta_{\mathrm{CP}})}\right|$$

where η_1 and η_2 are the unabsorbed Majorana CP-violating phases defined in Section 1.1.2.

Because this parameter depends directly on the absolute mass eigenvalues (m_i), its allowed physical boundaries are heavily dictated by the mass ordering configurations established in Section 1.2.3:

Normal Ordering (NO): The alignment of the mass states allows for potential destructive interference driven by the phases η_1 and η_2, meaning $\langle m_{\beta\beta}\rangle$ can vanish entirely even if neutrinos are Majorana particles.

Inverted Ordering (IO): The electron flavor state is dominated by the two heaviest, nearly degenerate mass states ($m_1 \approx m_2$). Destructive interference cannot fully eliminate the amplitude, establishing a strict lower experimental bound of $\langle m_{\beta\beta}\rangle \gtrsim 15$ MeV (subject to nuclear matrix element uncertainties).

1.1.5 Beyond the Standard Model Mass Generation: The Seesaw Mechanism and the νMSM

The Type-I Seesaw mechanism[12] represents a conservative extension to the Standard Model capable of generating non-zero neutrino masses. By introducing right-handed singlet neutrino fields (ν_{Ri}), the framework accounts for the smallness of light neutrino masses through a specific matrix geometry.

Crucially, from a field-theoretic standpoint, the Lagrangian leaves the mass scale of these right-handed states completely unfixed; it is a free parameter of the theory. Depending on the underlying empirical anomalies one seeks to resolve simultaneously, different mass scales for the heavy states emerge as distinct, parallel theoretical paradigms

1.1.5.1 Mathematical Formulation and Eigenvalue Derivation

To trace how the physical mass states are generated, the underlying matrix algebra can be illustrated using a simplified system containing one active left-handed neutrino flavor (ν_L) and one heavy right-handed singlet flavor (ν_R). In the flavor basis $\begin{pmatrix}\nu_L & \nu_R^c\end{pmatrix}^T$, the complete, symmetric neutrino mass matrix $\mathcal{M}$ takes the block form:

$$\mathcal{M} = \begin{pmatrix} 0 & m_D \\ m_D & M_R \end{pmatrix}$$

where $m_D = Y^\nu \frac{v}{\sqrt{2}}$ is the Dirac mass term generated by standard electroweak symmetry breaking via the Higgs vacuum expectation value v, and M_R is the explicit Majorana mass term for the right-handed field. Because ν_R carries no gauge charges, M_R is a gauge-invariant bare mass term decoupled from the electroweak scale.

To extract the physical mass eigenvalues (λ), we solve the characteristic equation $\det(\mathcal{M}-\lambda I) = 0$:

$$\det \begin{pmatrix} -\lambda & m_D \\ m_D & M_R - \lambda \end{pmatrix} = 0 \implies \lambda^2 - \lambda M_R - m_D^2 = 0$$

Applying the quadratic formula yields the exact physical mass eigenvalues of the system:

$$\lambda_{1,2} = \frac{M_R \pm \sqrt{M_R^2 + 4m_D^2}}{2}$$

Under the Seesaw parameter configuration, the Majorana mass scale is assumed to be significantly larger than the Dirac mass scale ($M_R \gg m_D$). Factoring M_R^2 out of the radical allows the expression to be simplified via a first-order Taylor expansion, $\sqrt{1+x} \approx 1 + \frac{1}{2}x$:

$$\lambda_{1,2} = \frac{M_R}{2}\left[1 \pm \sqrt{1 + \left(\frac{2m_D}{M_R}\right)^2}\right] \approx \frac{M_R}{2}\left[1 \pm \left(1 + \frac{2m_D^2}{M_R^2}\right)\right]$$

This expansion yields the two decoupled physical mass eigenvalues of the theory:

$$\lambda_1 \approx -\frac{m_D^2}{M_R}$$

$$\lambda_2 \approx M_R + \frac{m_D^2}{M_R} \approx M_R$$

The negative sign in the light eigenvalue λ_1 represents a parity phase factor that can be absorbed by a chiral rotation of the field, leaving the physical light neutrino mass positive: $m_\nu = m_D^2/M_R$. The heavy mass eigenstate scales directly with the bare mass parameter: $m_N \approx M_R$. Generalizing this geometry to three active and three sterile fields expands $\mathcal{M}$ into a 6×6 matrix, dictating that the active-sterile mixing matrix element Θ is defined as:

$$\Theta \approx m_D M_R^{-1}$$

While this single-flavor framework demonstrates the core algebraic mechanics of the Seesaw expansion, a realistic extension requires the introduction of multiple right-handed singlet fields. To accommodate the two distinct mass-squared splittings observed in solar and atmospheric flavor oscillations, a minimum of two right-handed fields ($n \geq 2$) is mathematically mandatory to generate at least two non-zero light mass eigenvalues.

When expanding the theory to a three-generation architecture ($n = 3$), the physical motivation splits depending on the chosen paradigm. In high-scale frameworks, a three-generation allocation is typically dictated by SO(10) Grand Unified Theories, which group all fermions of a given family into a single multiplet, thereby restoring generational symmetry between the lepton and quark sectors. Conversely, in the low-scale νMSM framework, a three-generation configuration represents the absolute minimal particle spectrum required to satisfy all cosmological boundary conditions simultaneously. While two of the heavy states are fully consumed by generating active neutrino masses and driving baryogenesis, the third state is liberated to act as a viable candidate for dark matter.

High-Scale Realization: The Grand Unification Paradigm

The first paradigm is a top-down approach motivated by Grand Unification (GUT). If one assumes that the neutrino Dirac mass matrix m_D is generated by the same Yukawa coupling hierarchies that govern the quarks and charged leptons ($m_D \sim \mathcal{O}(100 \text{ GeV})$), then to suppress the light eigenvalues m_ν down to the sub-eV scale $\mathcal{O}(0.1 \text{ eV})$, the free parameter M_R must sit near the GUT scale:

$$M_R \sim \mathcal{O}(10^{14} \text{ GeV})$$

This scale configuration attributes the smallness of neutrino mass to the existence of high-energy physics near the Planck scale. However, this configuration possesses a definitive phenomenological consequence. Because $\Theta \approx m_D M_R^{-1}$, the active-sterile mixing parameter scales down to $|\Theta|^2 \sim 10^{-24}$. This extremely small value suppresses direct production and decay rates at low energies, meaning the heavy states are effectively decoupled from direct laboratory observation.

Low-Scale Realization: The Cosmic Minimality Paradigm

The second paradigm is a bottom-up approach motivated by cosmic minimality exploring whether the parameters of the Standard Model can be extended to solve multiple empirical anomalies simultaneously without introducing new high energy scales.

As formalized in the Neutrino Minimal Standard Model (νMSM) [13], if the free parameter M_R is assigned to the MeV-GeV range, the light neutrino masses are kept small by tracking correspondingly smaller Yukawa couplings ($Y_\nu \sim \mathcal{O}(10^{-8})$). This low-scale configuration is motivated by its ability to resolve cosmological anomalies that the high-scale paradigm leaves unaddressed:

- Baryogenesis: The coherent oscillations of two nearly degenerate MeV/GeV-scale heavy states (N_2, N_3) in the early universe generate the CP-violating asymmetries required to explain the observed baryon asymmetry of the Universe via leptogenesis.
- Dark Matter: The remaining light singlet state (N_1) can naturally reside at the keV scale, providing a stable candidate for Warm Dark Matter.

Empirical Scope

Because the exact choice of nature for the free parameter M_R is not fixed by any gauge principle, the Seesaw mechanism must be investigated across its entire parameter space. High-scale configurations are constrained indirectly through cosmological data, lepton flavor violation limits, and neutrinoless double beta decay searches.

Concurrently, the low-scale paradigm represents an alternative realization. Because an MeV-scale heavy mass state sits within the kinematically open windows of standard meson decays, high-precision laboratory searches can directly probe this low-scale frontier, testing the unmapped parameters of the neutrino mass sector.

1.2 Leptonic Pion Decays

1.2.1 The π^+ Meson in the Standard Model

The pion was theoretically predicted by Hideki Yukawa in 1935 as the quantum mediator of the short-range strong nuclear force responsible for binding protons and neutrons inside atomic nuclei

[14]. The physical particle was subsequently observed experimentally within cosmic ray interactions in 1947 [15].

The π^+ meson is a valence quark-antiquark bound state, specifically composed of an up quark and a down antiquark ($u\bar{d}$). This explicit flavor combination governs its strong interactions via Quantum Chromodynamics (QCD) and results in a net macroscopically observable electric charge of $+1e$.

The internal dynamics of the valence constituents dictate a rigid set of space-time and internal quantum numbers that define the particle as a pseudoscalar meson:

- Spin ($J = 0$): The intrinsic spins of the constituent up quark and down antiquark are aligned anti-parallel to one another in their ground state. This configuration results in a net total intrinsic angular momentum of zero, categorizing the π^+ as a boson.
- Parity ($P = -1$): Because a fermion and an antifermion possess opposite intrinsic parities, and the spatial orbital angular momentum between them is zero ($L = 0$), the collective wave function behaves asymmetrically under spatial inversion.
- Isospin ($I = 1, I_z = +1$): Within the framework of hadronic flavor symmetries, the π^+ represents the positive projection of the basic light-meson isospin triplet, meaning it behaves symmetrically to its neutral and negative counterparts under the strong nuclear force.

The rest mass of the π^+ meson is approximately 139.57 MeV. This exceptionally small mass scale relative to other hadronic states (such as the proton or the ρ meson) is a direct consequence of Chiral Symmetry Breaking in QCD.

In the theoretical limit where the bare masses of the up and down quarks vanish entirely, the QCD Lagrangian exhibits an exact global chiral symmetry. The non-perturbative QCD vacuum spontaneously breaks this symmetry, which would naturally generate a strictly massless Goldstone boson. Because the physical light quarks possess a small but non-zero bare mass, the chiral symmetry is only approximate. This explicit breaking turns the physical π^+ meson into a pseudo-Goldstone boson, explaining its uniquely suppressed rest mass compared to non-Goldstone composite nuclear states.

1.2.2 Derivation of Leptonic Decay Widths and the Helicity Suppression Mechanism

The transition of a charged pion into a lepton-neutrino pair is governed by the coupling between hadronic states and the electroweak sectors of the Standard Model. Evaluating this process requires mapping the high-energy gauge boson exchange to a low-energy effective field theory. This framework allows the analytical derivation of the tree-level partial decay width $\Gamma(\pi^+ \to \ell^+ \nu_\ell)$ for a generic lepton flavor and isolates the specific mass dependencies that dictate the helicity suppression mechanism.

1.2.2.1 The Electroweak Interaction and the Effective Fermi Lagrangian

Leptonic pion decay occurs via the exchange of a virtual W^+ gauge boson. The process originates from the electroweak interaction Lagrangian density, which dictates the coupling of the W boson to the left-chiral quark and lepton currents:

$$\mathcal{L}_{\text{int}} = -\frac{g}{\sqrt{2}} \left(J^{\mu}_{\text{leptonic}} + J^{\mu}_{\text{hadronic}} \right) W^{+}_{\mu} + \text{h.c.}$$

where g represents the $SU(2)_L$ gauge coupling constant. The weak currents select left-chiral fields via the projection operator $P_L = \frac{1}{2}(1-\gamma^5)$:

$$J^{\mu}_{\text{leptonic}} = \bar{\nu}_\ell \gamma^\mu P_L \ell = \frac{1}{2} \bar{\nu}_\ell \gamma^\mu (1-\gamma^5) \ell$$

$$J^{\mu}_{\text{hadronic}} = V_{ud} \bar{u} \gamma^\mu P_L d = \frac{1}{2} V_{ud} \bar{u} \gamma^\mu (1-\gamma^5) d$$

Because the mass of the mediating W boson ($m_W \approx 80.4$ GeV$/c^2$) is large relative to the energy scale of the pion ($m_\pi \approx 139.57$ MeV$/c^2$), the momentum transfer squared q^2 carried by the boson propagator satisfies the condition $q^2 \ll m_W^2$.In this low-energy regime, the heavy W boson field is integrated out, mapping the non-local propagator to a local, four-fermion contact interaction vertex. Absorbing the gauge coupling and the boson mass into the Fermi coupling constant ($G_F/\sqrt{2} = g^2/8m_W^2$) yields the Effective Fermi Lagrangian:

$$\mathcal{L}_{\text{eff}} = -\frac{G_F}{\sqrt{2}} V_{ud} \left[\bar{d} \gamma^\mu (1-\gamma^5) u \right] \left[\bar{\ell} \gamma_\mu (1-\gamma^5) \nu_\ell \right] + \text{h.c.}$$

The transition amplitude $\mathcal{M}$ factorizes into a perturbative leptonic matrix element and a non-perturbative hadronic matrix element:

$$\mathcal{M} = \langle \ell^+ \nu_\ell | \mathcal{L}_{\text{eff}} | \pi^+ \rangle = -\frac{G_F}{\sqrt{2}} V_{ud} \langle 0 | \bar{d} \gamma^\mu (1-\gamma^5) u | \pi^+(p_\pi) \rangle \left[\bar{u}(p_\nu) \gamma_\mu (1-\gamma^5) v(p_\ell) \right]$$

The hadronic matrix element is constrained by the Lorentz properties of the pion. The pion is a pseudoscalar ($J^P = 0^-$), which causes the vector component $\bar{d}\gamma^\mu u$ of the quark current to vanish under parity conservation. The remaining axial-vector component $\bar{d}\gamma^\mu\gamma^5 u$ is parameterized using the pion decay constant f_π and the pion four-momentum p^μ_π:

$$\langle 0 | \bar{d} \gamma^\mu \gamma^5 u | \pi^+(p_\pi) \rangle = i f_\pi p^\mu_\pi$$

Combining this parameterization with the leptonic current yields the complete tree-level matrix element:

$$\mathcal{M} = i \frac{G_F}{\sqrt{2}} V_{ud} f_\pi p^\mu_\pi \left[\bar{u}(p_\nu) \gamma_\mu (1-\gamma^5) v(p_\ell) \right]$$

1.2.2.2 Kinematic Matrix Reduction and Decay Width

Applying conservation of four-momentum ($p^\mu_\pi = p^\mu_\ell + p^\mu_\nu$) allows the distribution of the pion momentum across the leptonic fields:

$$\mathcal{M} = i \frac{G_F}{\sqrt{2}} V_{ud} f_\pi \left[\bar{u}(p_\nu) \not{p}_\nu (1-\gamma^5) v(p_\ell) + \bar{u}(p_\nu) \not{p}_\ell (1-\gamma^5) v(p_\ell) \right]$$

Enforcing the Dirac equations for the external states simplifies the expression. Assuming a massless neutrino ($\bar{u}(p_\nu)\not{p}_\nu = 0$), the first term vanishes. For the outgoing antilepton, the Dirac equation satisfies $\not{p}_\ell v(p_\ell) = -m_\ell v(p_\ell)$. Anticommuting $\not{p}_\ell$ with γ^5 isolates the lepton mass as a linear factor:

$$\mathcal{M} = -i\sqrt{2}G_F V_{ud} f_\pi m_\ell \left[\bar{u}(p_\nu) P_R v(p_\ell)\right]$$

where $P_R = \frac{1}{2}(1+\gamma^5)$ represents the right-chiral projection operator. The amplitude vanishes in the limit $m_\ell \to 0$.

Squaring the amplitude, summing over the unpolarized final spin states via trace theorems, and integrating over the Lorentz-invariant two-body phase space (dLips) generates the partial decay width:

$$\Gamma(\pi^+ \to \ell^+ \nu_\ell) = \frac{G_F^2 |V_{ud}|^2 f_\pi^2}{8\pi} m_\pi m_\ell^2 \left(1 - \frac{m_\ell^2}{m_\pi^2}\right)^2$$

1.2.2.3 The Helicity Suppression Mechanism

The functional dependence of the decay width on the squared mass of the lepton (m_ℓ^2) arises from the constraints that angular momentum conservation imposes on relativistic fields.

The pion is a spin-zero ($J = 0$) particle, requiring the total angular momentum projection (J_z) along the decay axis to equal zero. In the pion rest frame, the final-state leptons are emitted back-to-back with equal and opposite three-momentum ($\mathbf{p}_\nu = -\mathbf{p}_\ell$). To maintain $J_z = 0$, their individual spin projections must point in opposite directions along the axis of motion.The weak charged current couples exclusively to left-chiral projections. For a massless neutrino, chirality is identical to helicity, constraining the particle to a negative (left-handed) helicity state ($\mathbf{s}_\nu \cdot \hat{\mathbf{p}}_\nu = -1/2$). To satisfy $J_z = 0$, the antilepton ℓ^+ must also be emitted in a negative (left-handed) helicity state ($\mathbf{s}_\ell \cdot \hat{\mathbf{p}}_\ell = -1/2$).The $V - A$ weak operator requires the outgoing antiparticle to possess a right-handed chiral projection. For a massive particle, chirality and helicity states decouple, meaning a state of definite helicity is a linear combination of both chiral states. The probability P of finding the left-chiral component required for the interaction scales with the velocity β:

$$P \propto 1 - \beta = 1 - \frac{|\mathbf{p}|}{E}$$

Expanding this relation in the relativistic limit ($E \gg m_\ell$) isolates the mass dependence:

$$1 - \sqrt{1 - \frac{m_\ell^2}{E^2}} \approx \frac{m_\ell^2}{2E^2}$$

Because the total energy available to the system is bounded by the mass of the parent pion (m_π), the probability of the antilepton possessing the required chiral projection scales with its rest mass:

$$P \propto \left(\frac{m_\ell}{m_\pi}\right)^2$$

1.2.3 The Standard Model Ratio $R^\pi_{e/\mu}$ and Experimental Benchmarks

The analytical calculation of the partial decay widths in Section 2.2 provides the baseline tree-level ratio of the electronic and muonic decay channels. Evaluating the ratio eliminates structural hadronic parameters and enables a direct test of the Standard Model using high-precision experimental measurements.

1.2.3.1 Tree-Level Cancellation of Hadronic Parameters

The tree-level branching ratio $R_{e/\mu}^{\pi,\text{tree}}$ is defined as the ratio of the individual partial decay widths:

$$R_{e/\mu}^{\pi,\text{tree}} = \frac{\Gamma(\pi^+ \to e^+\nu_e)}{\Gamma(\pi^+ \to \mu^+\nu_\mu)} = \frac{\frac{G_F^2|V_{ud}|^2 f_\pi^2}{8\pi} m_\pi m_e^2 \left(1 - \frac{m_e^2}{m_\pi^2}\right)^2}{\frac{G_F^2|V_{ud}|^2 f_\pi^2}{8\pi} m_\pi m_\mu^2 \left(1 - \frac{m_\mu^2}{m_\pi^2}\right)^2}$$

The factors representing the Fermi coupling constant (G_F), the Cabibbo-Kobayashi-Maskawa matrix element (V_{ud}), and the non-perturbative hadronic parameterization via the pion decay constant (f_π) are identical in both channels. Consequently, these structural parameters cancel completely in the ratio, leaving an expression dependent exclusively on the physical masses of the particles:

$$R_{e/\mu}^{\pi,\text{tree}} = \left(\frac{m_e}{m_\mu}\right)^2 \left(\frac{m_\pi^2 - m_e^2}{m_\pi^2 - m_\mu^2}\right)^2$$

Substituting the established physical mass values ($m_e \approx 0.511$ MeV$/c^2$, $m_\mu \approx 105.66$ MeV$/c^2$, and $m_\pi \approx 139.57$ MeV$/c^2$) yields the numerical tree-level prediction:

$$R_{e/\mu}^{\pi,\text{tree}} \approx 1.28 \times 10^{-4}$$

1.2.3.2 Electromagnetic Radiative Corrections

To confront experimental data, the tree-level expression must be modified by higher-order virtual photon loops and real photon emission (inner bremsstrahlung):

$$R_{e/\mu}^{\pi} = R_{e/\mu}^{\pi,\text{tree}}(1 + \Delta_{\text{e.m.}})$$

The infrared divergences present in the virtual loops cancel exactly when summed with the phase-space integrated real photon emission rate. Due to the large mass difference between the electron and the muon, the electron channel is subject to large logarithmic corrections scaling as $\alpha \ln(m_\mu/m_e)$. The systematic application of these corrections shifts the quantitative prediction downwards by approximately 4%, establishing the standard reference framework value [16][17]:

$$R_{e/\mu}^{\pi,\text{SM}} = (1.2352 \pm 0.0002) \times 10^{-4}$$

The assigned theoretical uncertainty of 0.016% includes conservative bounds for uncalculated higher-order electroweak loops and structure-dependent effects.

1.2.3.3 Experimental Benchmarks

The Standard Model prediction is tested by measurements from experiments at meson accelerator facilities:

Experiment	Year	$R_{e/\mu}^{\text{exp}}$ Value	Precision
PSI [18]	1993	$[1.2346 \pm 0.0035(\text{stat}) \pm 0.0036(\text{syst})] \times 10^{-4}$	0.40%
TRIUMF [19]	1994	$[1.2265 \pm 0.0034(\text{stat}) \pm 0.0044(\text{syst})] \times 10^{-4}$	0.45%
PIENU [20]	2015	$[1.2344 \pm 0.0023(\text{stat}) \pm 0.0019(\text{syst})] \times 10^{-4}$	0.24%

These experimental results are consistent with the calculated Standard Model value of $R_{e/\mu}^{\pi,\mathrm{SM}} = (1.2352 \pm 0.0002) \times 10^{-4}$. However, the experimental uncertainty of the most precise measurement (0.24%) remains roughly fifteen times larger than the theoretical uncertainty (0.016%). This precision gap motivates next-generation experiments to further constrain potential non-Standard Model interactions.

1.3 General Phenomenological Probes of New Physics in $\pi^+ \to e^+ X$ Topologies

1.3.1 Probing HNLs: The Production Strategy

The experimental investigation of HNLs is fundamentally divided into two complementary frameworks: the decay strategy and the production strategy.

The decay strategy targets the downstream signatures of the heavy states. Instead of observing the creation event, this approach monitors a designated region for the visible SM particles resulting from the HNL decay. Because this method relies on the particle decaying within the observable region, its sensitivity is intrinsically linked to the total decay width (or lifetime) of the HNL. Consequently, this strategy provides constraints that depend on both the mixing strength and the probability of the particle decaying visibly within the acceptance of the detector.

Conversely, the production strategy isolates the initial creation vertex by reconstructing the missing four-momentum of the system. This approach is typically applied to the decays of light mesons (such as pions and kaons). In these events, the four-momentum of the parent particle and the four-momentum of the visible charged lepton are precisely measured. The difference gives the missing four-momentum carried away by the unseen neutral particle. In the SM, the invariant mass derived from this vector is consistent with zero, corresponding to a nearly massless neutrino. However, if a massive HNL is produced, the invariant mass of the missing four-momentum will peak at a distinct, non-zero value corresponding to the square of the HNL mass.

1.3.1.1 Analytical Expression for the Branching Ratio

If an HNL possesses a mass below the kinematic threshold of the parent meson ($m_N < m_{\pi^+} - m_e$), it can be directly produced via its mixing component with the active electron neutrino state. The branching ratio for the specific decay channel $\pi^+ \to e^+ N$ is formalized relative to the standard active mode as [21]:

$$\mathcal{B}(\pi^+ \to e^+ N) = \mathcal{B}(\pi^+ \to e^+ \nu_e) \cdot \rho_e \cdot |U_{eN}|^2$$

where $\mathcal{B}(\pi^+ \to e^+ \nu_e)$ represents the branching ratio within the SM framework, and $|U_{eN}|^2$ is the active-sterile mixing matrix element governing the electron-flavor coupling strength. The parameter ρ_e is a dimensionless kinematic phase-space and helicity correction factor defined as a function of the squared mass ratios:

$$\rho_e = \frac{(x+y) - (x-y)^2}{x(1-x)^2} \cdot \lambda^{1/2}(1, x, y)$$

where the normalized mass parameters are defined as $x \equiv m_e^2/m_\pi^2$ and $y \equiv m_N^2/m_\pi^2$. The function $\lambda(a, b, c) \equiv a^2 + b^2 + c^2 - 2ab - 2bc - 2ca$ is the kinematic Källén function, which in this rest-frame configuration reduces to:

$$\lambda(1, x, y) = 1 + x^2 + y^2 - 2x - 2y - 2xy$$

As derived explicitly in Section 2, the standard active decay channel $\pi^+ \to e^+\nu_e$ is severely suppressed within the SM due to the helicity constraints imposed by an axial-vector weak current acting on near-massless fermions.

Because current empirical bounds constrain the active-sterile mixing matrix element to exceptionally small values ($|U_{eN}|^2 \ll 1$), the total branching ratio $\mathcal{B}(\pi^+ \to e^+N)$ remains highly suppressed and sits multiple orders of magnitude below the SM baseline.

1.3.1.2 Current Limits on the HNL Mixing Parameter

The global parameter space for low-scale heavy neutral leptons is characterized by independent constraints across the individual active flavor sectors. Within the kinematic regime governed by the two-body direct production topology, the upper limits on active-sterile mixing are mapped as functions of the HNL mass m_N and the respective mixing elements $|U_{\ell N}|^2$.

The Electron-Flavor Frontier ($|U_{eN}|^2$). The upper limits constraining HNL coupling to the electron neutrino sector span from the lower kinematic bounds of pion decay up to the kaon mass threshold. This continuous mass spectrum is scanned via two primary experimental targets using the peak search methodology:

The Pion Mass Regime (20 MeV $\leq m_N \leq$ 140 MeV): This domain is governed by pion-at-rest spectroscopic searches. The early TRIUMF experiment established the first stringent constraints with an upper limit of $|U_{e4}|^2 \sim O(10^7)$ for masses 50–130 MeV [22]. The subsequent PIENU collaboration at TRIUMF significantly improved these bounds via high-precision scans of the positron energy spectrum, finding no statistically significant monochromatic peaks. These null results currently define the leading exclusion boundary in this window, pushing the sensitivity down to $|U_{e4}|^2 \sim O(10^{-8})$ over the mass range 60–125 MeV [23].

The Kaon Mass Regime (140 MeV $\leq m_N \leq$ 450 MeV): Beyond the pion kinematic threshold, the two-body peak-search topology shifts to charged kaon decay channels ($K^+ \to e^+N$). High-intensity in-flight measurements by the NA62 collaboration at CERN have probed this region, establishing upper limits for $|U_{e4}|^2$ at 90% CL of $\mathcal{O}(10^{-9})$ over the accessible heavy neutral lepton mass range of 144–462 MeV [24].

The Muon-Flavor Frontier ($|U_{\mu N}|^2$). The upper limits constraining HNL coupling to the muon neutrino sector span from the lower kinematic bounds of pion decay up to the kaon mass threshold. This continuous mass spectrum is scanned via two primary experimental targets using the peak search methodology: the two-body decays $\pi^+ \to \mu^+N$ and $K^+ \to \mu^+N$.

Direct peak-searches constrain this parameter space across complementary techniques. The PIENU experiment, utilizing stopped pions, sets a 90% CL upper limit of $|U_{\mu N}|^2 \sim \mathcal{O}(10^{-5})$ for masses 15–33 MeV [25]. In the higher-mass kaon window, the stopped-kaon experiment E949 [26] and the decay-in-flight experiment NA62 [27] reach sensitivities ranging from $|U_{\mu N}|^2 \sim \mathcal{O}(10^{-8})$ to $\mathcal{O}(10^{-9})$ over the mass range 175–384 MeV.

1.3.2 Global Electroweak Probes: The Role of Precision Measurements

1.3.2.1 Lepton Flavor Universality

In the Standard Model of particle physics, lepton flavor universality (LFU) is a fundamental consequence of the gauge sector. Under the electroweak $SU(2)_L \times U(1)_Y$ gauge symmetry, the gauge bosons couple with identical strength to all three generations of lepton fields: electron, muon, and tau.

Consequently, any differences observed in the decay rates of a charged pion into different lepton families do not arise from the underlying weak interaction coupling strength. Instead, they are determined entirely by the kinematics of the phase space and the physical masses of the participating leptons. Measuring the ratio of these decay rates, such as $R^{\pi}_{e/\mu}$, serves as a clean and precise probe of this gauge symmetry. Any deviation from the predicted ratio would point to a violation of this universal coupling.

1.3.2.2 Experimental and Phenomenological Convergence

The direct search for HNLs and the evaluation of the lepton flavor universality though the $R^{\pi}_{e/\mu}$ ratio are connected by a shared experimental signature. Both approaches investigate the same final state topology.

While the ratio study is evaluated as a test of standard electroweak couplings, it is inherently sensitive to the same Beyond the Standard Model parameters targeted by direct HNL searches. The presence of a heavy sterile neutrino field that mixes with the electron neutrino flavor would modify the total electronic decay rate. Therefore, if such a state exists within the permitted kinematic range, it would manifest as a global numerical shift in the measured value of $R^{\pi}_{e/\mu}$. This connects the two approaches into a complementary framework for constraining new physics.

Chapter 2

The NA62 experiment

NA62 is a fixed-target experiment located in the North Area of the Super Proton Synchrotron (SPS) at CERN (Figure 2.1), dedicated to high-precision measurements of rare kaon decays. The primary objective of the experiment is to determine the branching ratio of the ultra-rare decay $K^+ \to \pi^+ \nu \bar{\nu}$. Recent results from NA62 have reported a branching ratio of $BR(K^+ \to \pi^+ \nu \bar{\nu}) = (13.0^{+3.3}_{-3.0}) \times 10^{-11}$ [28], representing the smallest branching fraction measured to date, with a signal significance exceeding 5σ.

In addition to this primary goal, NA62 pursues an extensive parallel physics program. This includes the study of various other rare decay modes [29], searches for Lepton Flavor Violation (LFV) and Lepton Number Violation (LNV) [30], and direct searches for exotic particles [31]. The experimental program also incorporates advancements in neutrino tagging techniques [32] alongside searches for Heavy Neutral Lepton (HNL) production [24].

The following sections provide a general overview of the NA62 apparatus, including the beamline, subdetectors, and the trigger and data acquisition (TDAQ) system. For a more detailed description, see [33].

2.1 Beam line

A 400 GeV/c proton beam from the CERN Super Proton Synchrotron (SPS) is directed via the P42 beamline onto a beryllium target (T10), which has dimensions of 400 mm in length and 2 mm in diameter. This interaction produces an unseparated secondary beam composed of 70% π^+, 23% protons, and 6% K^+.

The NA62 beamline and detector configuration is illustrated in Fig. 2.2. A right-handed laboratory coordinate system is defined such that the beamline axis corresponds to the z-axis, with the origin located at the T10 target. Particles propagate along the positive z-direction. The y-axis is oriented vertically (positive upwards), and the x-axis is horizontal (positive to the left).

The T10 target is followed by a 950 mm water-cooled copper collimator (15 mm aperture) and a triplet of radiation-hard quadrupoles (Q1, Q2, Q3) providing large acceptance (± 2.7 mrad horizontally and ± 1.5 mrad vertically) at 75 GeV/c. An achromat (A1), consisting of four vertically-deflecting dipoles, selects a 75 GeV/c beam with a momentum spread of 1% rms. During this process, the beam passes through motorized beam-dump units (TAX1 and TAX2) to facilitate

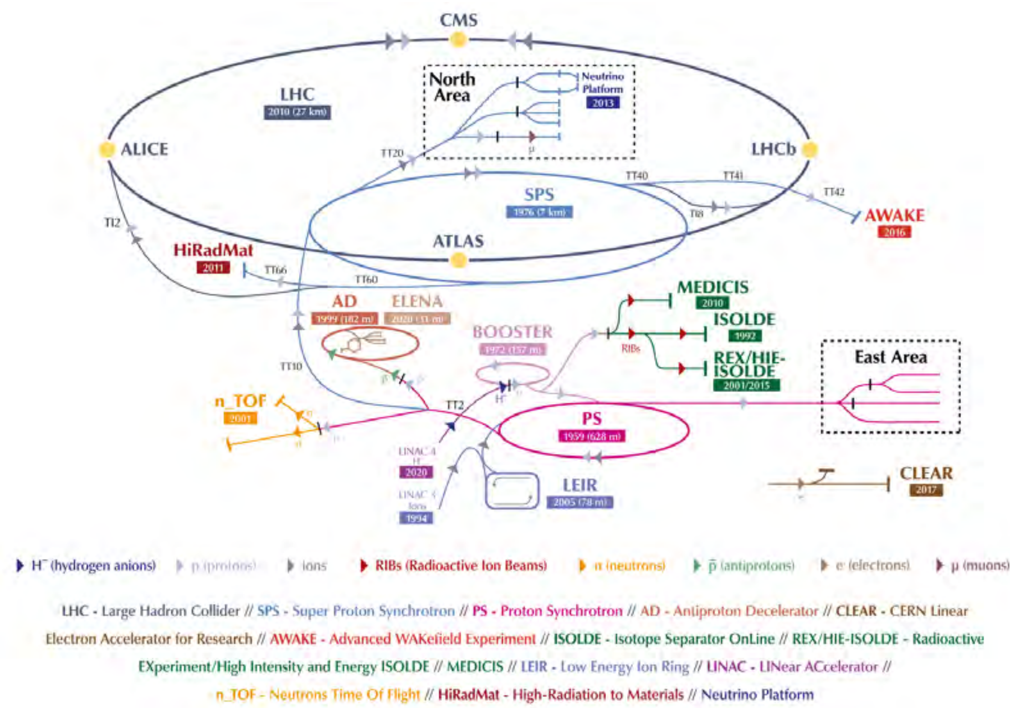

Figure 2.1: The CERN accelerator complex. The NA62 experiment is located in the region demarcated as the North Area. [34]

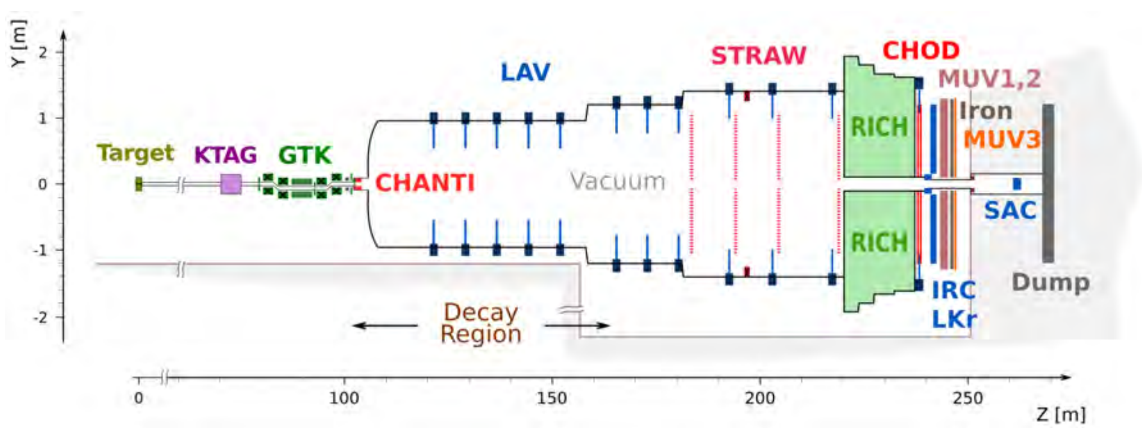

Figure 2.2: Schematic vertical view of the NA62 experimental setup. [33]

momentum selection and absorb the primary proton beam and unwanted secondary particles (Fig. 2.3).

Between TAX1 and TAX2, a tungsten radiator (up to 5 mm, 1.3 X_0) is introduced to degrade

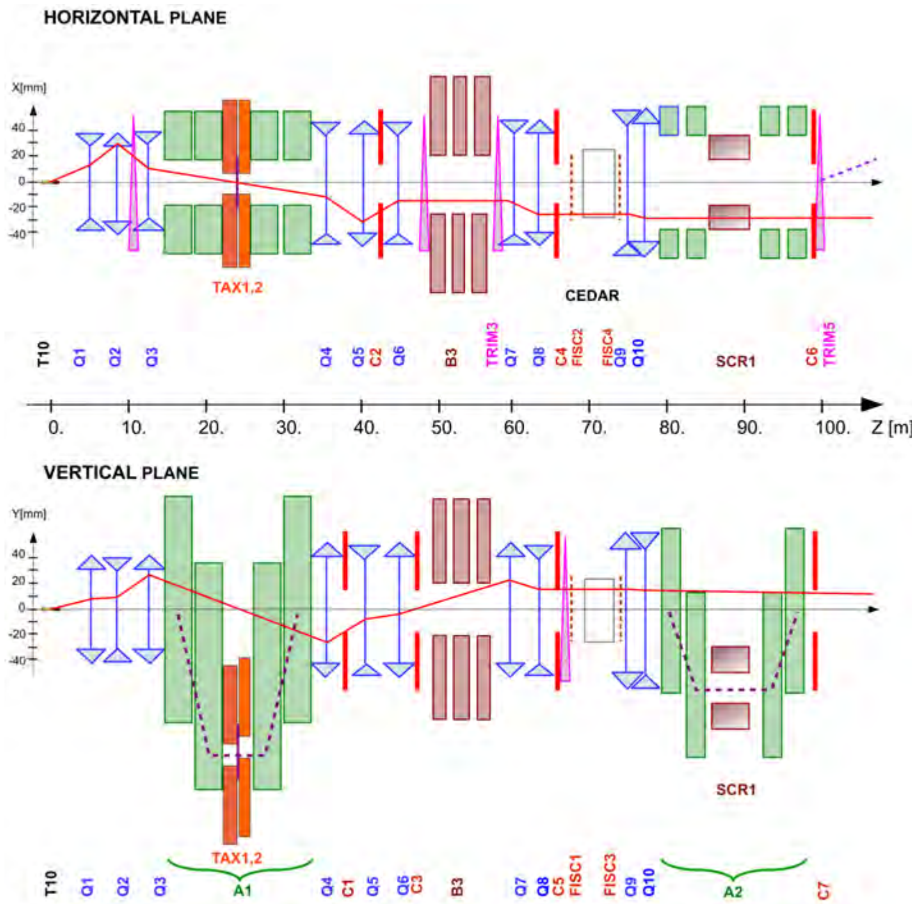

Figure 2.3: Schematic layout and optics of the high-intensity beam from the T10 target to the entrance of the decay region. [33]

positron energy via Bremsstrahlung while minimizing hadron scattering. A subsequent quadrupole triplet (Q4, Q5, Q6) and collimators (C1, C2, C3) redefine the beam acceptance and remove the momentum-degraded positrons.

The beam then traverses a 40 mm bore within three 2 m dipoles (B3) in iron plates; vertical magnetic fields sweep away muons, while steering dipoles (TRIM 2 and TRIM 3) compensate for stray fields. K^+ identification is performed by a differential Cherenkov counter (CEDAR) equipped with photodetectors (KTAG), preceded by quadrupoles (Q7, Q8) and cleaning collimators (C4, C5). Beam divergence is measured and tuned using two pairs of filament scintillators (FISC 1–4).

Quadrupoles (Q9, Q10) match the beam to the tracking and momentum-measurement stage, which comprises the GTK silicon pixel system (three stations, 2016–2018; four stations, from 2021 onwards). An achromat (A2) is positioned between GTK1 and GTK3, featuring a toroidal iron collimator (SCR1) to defocus muons (Fig. 2.4). GTK2 is located in a dispersive section ($\Delta Y = -60$ mm), while GTK3 (102.4 m from target) marks the decay region entrance, preceded by cleaning collimators (C6, C7).

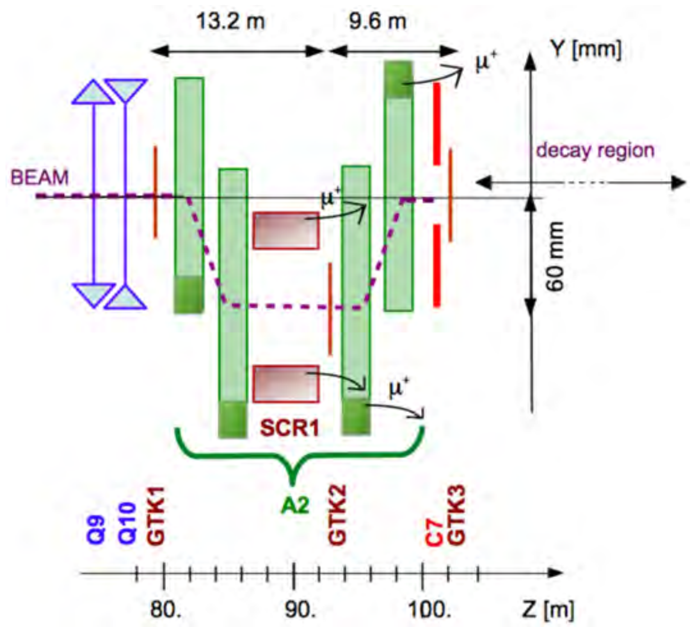

Figure 2.4: Schematic layout of the achromat [33]. GTK2 station was moved to be upstream of SCR1 from 2021 onwards.

A horizontal steering magnet (TRIM 5) deflects the beam by +1.2 mrad, which is subsequently countered by a −3.6 mrad deflection from the MNP33 dipole to direct the beam through the LKr calorimeter (Fig. 2.5).

The decay region occupies the first 75 m of a 117 m long, evacuated (10^{-6} mbar) tank, housing 11 LAV detectors and four STRAW chambers. The vessel diameter increases from 1.92 m to 2.8 m, utilizing telescopic flanges for detector removal. The downstream end features a 2 mm aluminum window separating the vacuum from the 17 m RICH counter, with a 168 mm diameter beam tube following the beam trajectory (Fig. 2.5).

The magnetic spectrometer includes two pairs of STRAW chambers on either side of the MNP33 dipole, which provides a 270 MeV/c horizontal kick (−3.6 mrad) to cross the axis 2.8 m downstream of the LKr center. After passing through FISC 5 and 6, the beam is deflected by −13.2 mrad via a BEND dipole to clear the SAC. The beam is finally absorbed in an iron/concrete dump, with intensity and profile monitored by a wire chamber and an ionization chamber.

2.2 Subdetectors

2.2.1 Kaon TAGger (KTAG)

Kaons constitute approximately 6% of the beam and are identified by the KTAG detector. The KTAG utilizes a CEDAR (Cherenkov Differential counter with Achromatic Ring focusing) type detector, which employs a gaseous radiator to produce Cherenkov light. The gas vessel is connected

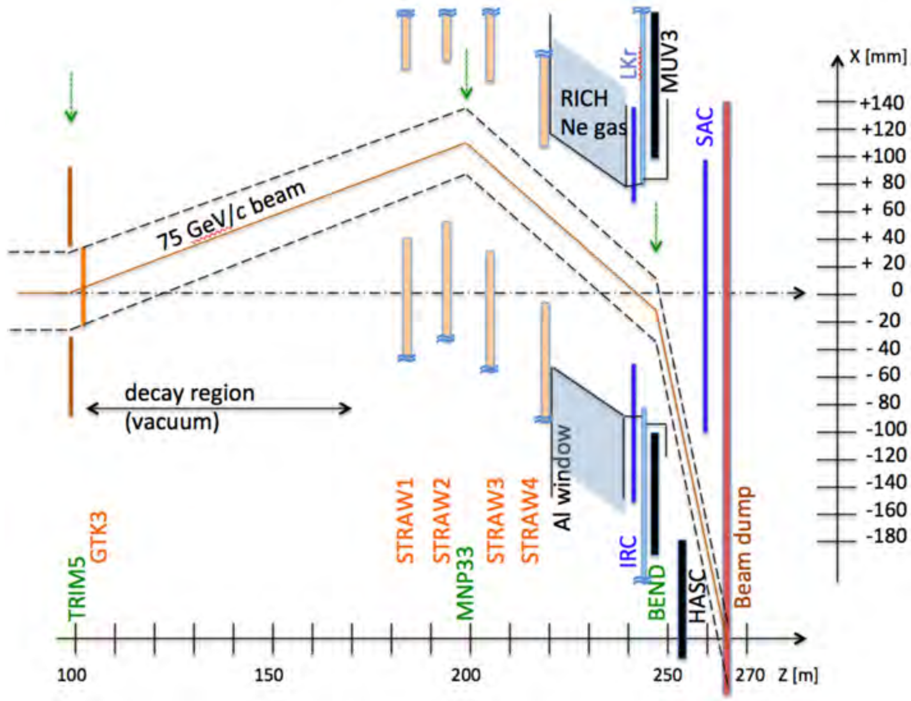

Figure 2.5: Beam line layout through the decay region and detectors in the (X, Z) plane. [33]

to the vacuum beam pipe at both ends and is isolated from the vacuum by two aluminium windows with thicknesses of 150 μm and 200 μm. Both the gas and the aluminium windows contribute to the total material budget in the beam path. The produced light is focused through eight quartz windows into an enclosed photon-detection system. Within this enclosure, eight spherical mirrors reflect the light by 90 degrees onto eight KTAG sectors (Fig. 2.6). Each sector consists of a Photomultiplier Tube (PMT) array comprising 48 PMTs, which are sensitive to wavelengths above 230 nm. The spherical mirrors are utilized to broaden the light spot and reduce the photon intensity incident on the PMTs.

Since the start of NA62 data-taking in 2016, KTAG has operated using a CEDAR-W configuration filled with nitrogen (N_2) at 1.71 bar. In this configuration, the CEDAR introduces 39×10^{-3} X_0 of material into the beam path, comprising 35×10^{-3} X_0 from the N_2 gas and 3.9×10^{-3} X_0 from the aluminium windows.

In 2023, KTAG underwent a major upgrade to a CEDAR-N type detector filled with H_2 at 3.85 bar (referred to as CEDAR-H). This configuration reduces the material budget in the beam path to 7.3×10^{-3} X_0, consisting of 3.4×10^{-3} X_0 from the H_2 gas and 3.9×10^{-3} X_0 from the aluminium windows [35].

The time resolutions for CEDAR-W and CEDAR-H are 71 ps and 66 ps, respectively. Based on coincidences in at least five sectors, the K^+ identification efficiency is 99.5% for CEDAR-W and 99.7% for CEDAR-H.

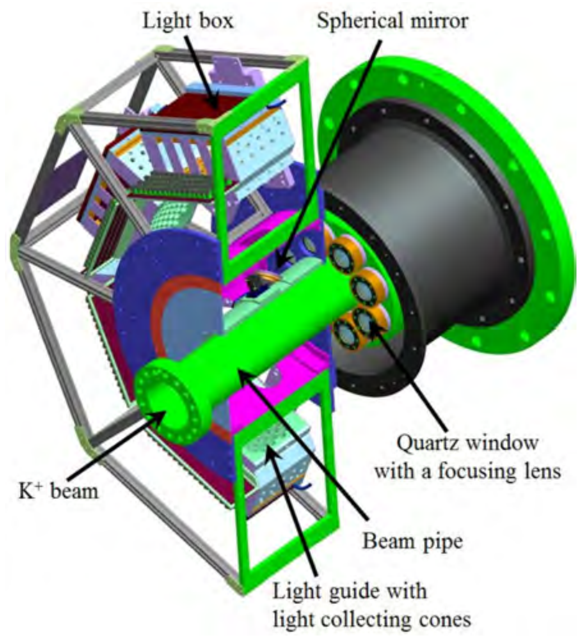

Figure 2.6: Illustration of the upstream section of the CEDAR and the KTAG. [33]

2.2.2 GigaTracKer (GTK)

The GTK, located within the vacuum pipe immediately upstream of the decay region, provides measurements of particle momentum, direction, and time with resolutions of 0.15 GeV/c, 16 μrad, and 100 ps, respectively.

The detector configuration consisted of three silicon pixel stations during the 2016–2018 period and was upgraded to four stations in 2021, integrated with two pairs of dipole magnets forming an achromat (Fig. 2.7). Particle momentum is determined from the vertical displacement of the trajectory within the second station.

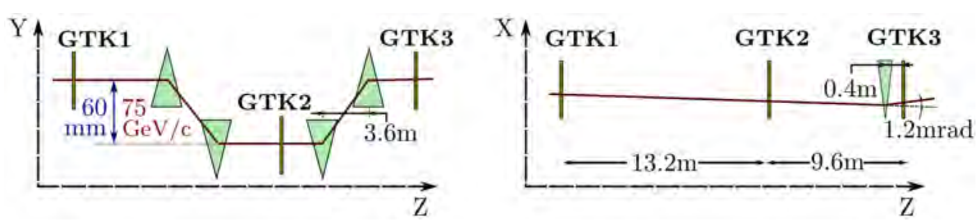

Figure 2.7: Vertical and horizontal schematic views of the GTK stations positioned within the achromat. [33]

Each station (Fig. 2.8) is a hybrid silicon detector comprising 18,000 pixels arranged in a 200×90 matrix, with each pixel measuring 300×300 μm^2. This arrangement covers a total area of 62.8×27 mm^2. The material budget for each station is maintained below 0.5% X_0, which corresponds to an equivalent thickness of approximately 500 μm of silicon.

Figure 2.8: Photographs of a GTK station. [33]

2.2.3 CHarged ANTI-coincidence detector (CHANTI)

The CHANTI detector provides rejection for background from inelastic interactions of the beam with the most downstream GTK station, GTK3.

The CHANTI is composed of six square hodoscope stations 300×300 mm^2 in cross section with a 95×65 mm^2 hole in the centre to leave room for the beam (Fig. A.5). The first station is placed 28 mm downstream of GTK3, and the distance between each station and the next one approximately doubles for successive stations, so that the angular region between 49 mrad and 1.34 rad is covered hermetically for particles generated on GTK3. The stations are made of scintillator bars of triangular cross section read out with fast wavelength-shifting (WLS) fibres coupled to silicon photomultipliers (SiPMs).

2.2.4 STRAW spectrometer (STRAW)

The straw spectrometer, located in the vacuum tank downstream of the decay region, reconstructs the trajectories and momenta of charged particles produced in beam decays. The track momentum resolution is $\frac{\sigma(p)}{p} = 0.30\% \oplus 0.005\% \cdot p$, where p is expressed in GeV/c. The spectrometer comprises four straw chambers and a large-aperture dipole magnet (MNP33) positioned between the second and third chambers (Fig. 2.5). This magnet provides an integrated field of 0.9 Tm, delivering a horizontal momentum kick of 270 MeV/c. The total radiation length of the spectrometer is 1.8% X_0.

Each straw chamber is composed of two modules. One module contains two views measuring X (0°), Y (90°) and the other module contains the U (−45°) and V (+45°) views (Fig. 2.10-left). The chamber active area is a circle of 2.1 m outer diameter centred on the longitudinal z-axis. Each view has a gap of about 12 cm without straws near the centre, such that, after overlaying the four views, an octagon shaped hole of 6 cm apothem is created for the beam passage. High detection redundancy is provided through a straw arrangement with four layers per view, which guarantees at least two hits per view, i.e. 8 –12 hits per track per chamber (Fig. 2.10-right). Due to the 12 cm gap without straws in each view, the number of hits per chamber is not evenly distributed over the detector surface.

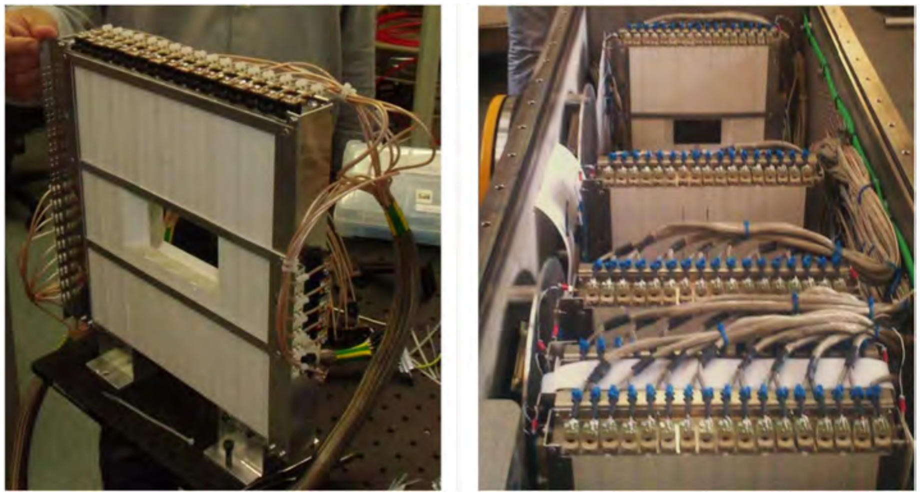

Figure 2.9: Left: A single CHANTI station. Right: The upstream part of the CHANTI assembly. [33]

Each chamber contains 1792 straws of 9.82 mm diameter and 2160 mm length. The gas inside the straws is a mixture of 70% Ar and 30% CO2 at atmospheric pressure. Each straw is made from 36 μm thick polyethylene terephthalate (PET), coated with 50 nm of copper and 20 nm of gold on the inside.

2.2.5 Large-Angle Veto system (LAV)

The LAV detectors provide full geometric coverage for photons originating from decays within the decay volume at angles ranging from 8.5 to 50 mrad with respect to the z-axis. The detector system consists of 12 stations: 11 are positioned along the vacuum volume, while the 12th station is located approximately 3 m upstream of the LKr calorimeter and operates in air.

Each LAV station is composed of modules arranged around the perimeter of the sensitive volume of the experiment, oriented radially to form an inward-facing ring. Multiple rings are utilized within each station to provide the depth required for efficient particle detection. The modules in successive rings are staggered in azimuth, with a longitudinal spacing of approximately 1 cm between rings (Fig. 2.11).

The detector modules consist of lead glass blocks containing approximately 75% lead oxide (PbO) by weight, with a density of $\rho = 5.5$ g/cm^3 and a radiation length of $X_0 = 1.50$ cm. Each block has a length of 37 cm, with a front face area of 10×10 cm^2 and a rear face area of 11×11 cm^2.

2.2.6 Liquid Krypton calorimeter (LKr)

The LKr is a quasi-homogeneous electromagnetic calorimeter designed for particle identification and photon detection within an angular range of 1 to 8.5 mrad. The detector consists of approximately 9,000 liters of liquid krypton maintained at 120 K, housed within a cryostat. Geometrically, the calorimeter extends from the beam pipe ($r \approx 8$ cm) to a radius of 128 cm, with a longitudinal depth

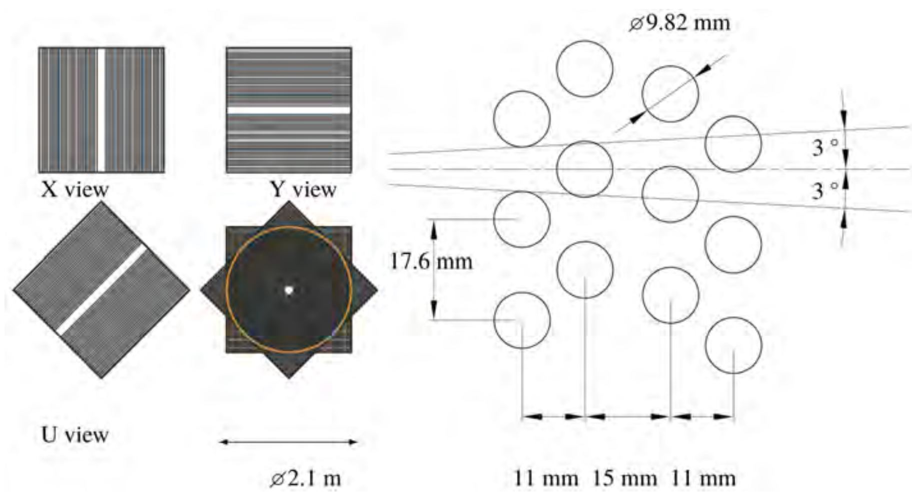

Figure 2.10: Left: Four views (X, Y, U, V) make up one straw chamber, each providing a measurement of one coordinate. Right: Straw geometry based on two double layers per view. [33]

of 127 cm, providing approximately $27X_0$ of radiation length.

The sensitive volume is partitioned into 13,248 longitudinal cells, each with a cross-section of approximately 2×2 cm^2. These cells are formed by Cu-Be electrodes aligned along the longitudinal axis of the experiment. To mitigate detection inefficiencies for particle showers occurring in close proximity to the anode, the electrodes feature a zig-zag geometry (Fig. 2.12). The energy resolution of the LKr calorimeter is given by:

$$\frac{\sigma(E)}{E} = \frac{4.8\%}{\sqrt{E}} \oplus \frac{11\%}{E} \oplus 0.9\%$$

2.2.7 Small-Angle Veto system (SAV)

The small-angle veto system provides hermeticity for photons emitted at angles down to zero degrees with respect to the z-axis. It consists of two detectors: the intermediate-ring calorimeter (IRC) and the Small-Angle Calorimeter (SAC). Both are shashlyk calorimeters, with lead and plastic-scintillator plates traversed by wavelength-shifting (WLS) fibres.

The SAC is located within the beam vacuum chamber near the end of the beam pipe, immediately upstream of the beam dump. The detector comprises 70 lead plates and 70 injection-molded plastic scintillator plates, each with a thickness of 1.5 mm and transverse dimensions of 205×205 mm^2, providing a total depth of 19 X_0 (Fig. 2.13-left).

The IRC is configured as an eccentric cylinder surrounding the beam pipe, positioned upstream of the LKr calorimeter. Centered on the z-axis, the detector features an outer diameter of 290 mm and a central bore of 120 mm. The IRC is divided into two longitudinal modules, upstream and downstream, with depths of 89 mm and 154 mm, respectively, separated by a 40 mm gap. The inner diameter of the downstream module is 2.2 mm larger than that of the upstream module.

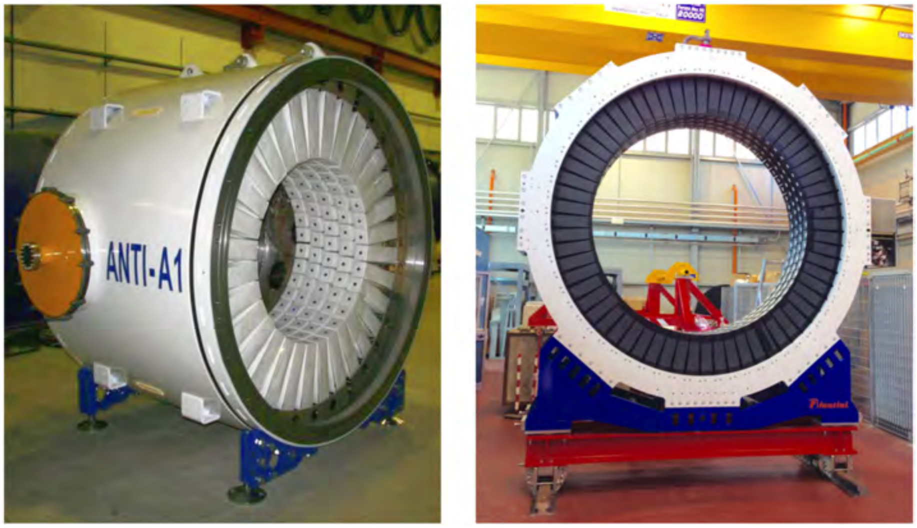

Figure 2.11: Photographs of the LAV1 (left) and LAV12 (right) stations prior to beam line installation. [33]

These modules comprise 25 and 45 ring-shaped layers, respectively; each layer consists of a 1.5 mm thick lead absorber interleaved with a 1.5 mm thick scintillator plate, providing a total depth of 19 X_0 for both modules (Fig. 2.13).

2.2.8 Ring Imaging CHerenkov counter (RICH)

The RICH detector measures charged particle times with a resolution of 70 ps, provides the trigger reference time and is used for particle identification.

The RICH radiator is a 17.5 m long cylindrical vessel made of ferro-pearlitic steel (Fig. 2.14) and filled with neon gas. The vessel consists of four sections of gradually decreasing diameter and different lengths. At the upstream end the vessel is about 4.2 m wide to accommodate the photomultiplier flanges outside the active area of the detector. The diameter of the last vessel element is 3.2 m, which is sufficient to house the mirrors and their support panel.

The active area of the detector extends to a radial distance of 1.1 m from the beam axis at the RICH entrance and to 1.4 m at the exit window. The entrance and exit windows have a conical shape and are made of aluminium with a thickness of 2 and 4 mm respectively. The entrance window is the only separator between the decay vacuum volume and the RICH radiator gas volume. A lightweight aluminium tube, connected to the decay-tank vacuum, passes centrally through the vessel.

The mirror layout consists of a mosaic of 20 spherical mirrors is used to reflect the Cherenkov light cone into a ring on the photomultiplier (PM) array in the mirror focal plane. The mirrors are divided into two spherical surfaces: one with the centre of curvature to the left and one to the right of the beam pipe. The total reflective surface exceeds 6 m^2 . The mirrors have a nominal radius of curvature of 34 m and hence a focal length of 17 m. The mosaic includes 18 mirrors of regular hexagonal shape (350 mm side) and two half mirrors. The two latter ones are used in the centre

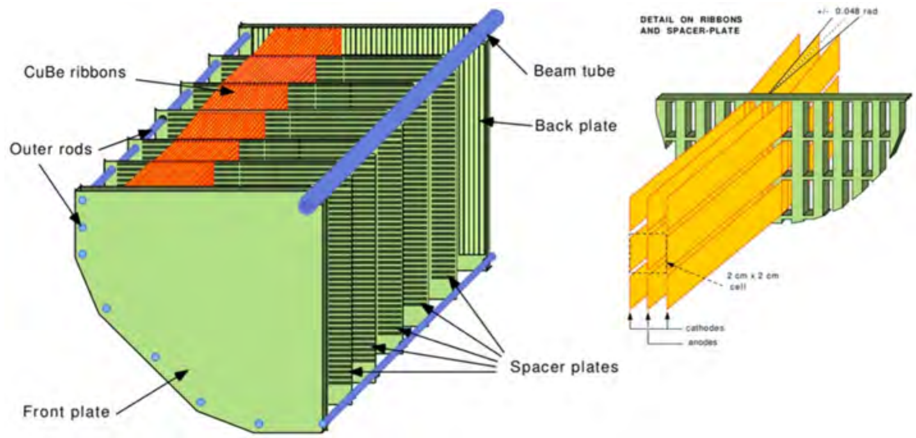

Figure 2.12: Left: Schematic representation of a single calorimeter quadrant. Right: Detailed illustration of the calorimeter cells. [33]

Figure 2.13: Left: SAC detector. Right: IRC detector. [33]

and have a circular opening to accommodate the beam pipe. The mirrors are made from 25 mm thick glass substrate coated with aluminium. A thin dielectric film is added for protection and to improve the reflectivity.

2.2.9 Charged particle HODoscopes (CHODs)

The charged particle hodoscope system consists of scintillators covering the lateral acceptance downstream of the RICH and upstream of the LKr calorimeter, bounded by the LAV12 detector inner radius (1070 mm) and the IRC detector outer radius (145 mm). The primary function of this system is to provide an input for the L0 trigger (Section 2.3.1) upon the detection of at least one charged particle, achieving a temporal resolution of 200 ps.

The system consists of the NA48-CHOD detector from the former kaon experiment NA48 and the newly constructed CHOD detector optimized for the high intensity conditions of NA62. The NA48-CHOD and the CHOD are located, respectively, downstream and upstream of the LAV12 detector, about 700 mm apart in the longitudinal direction.

The NA48-CHOD consists of two consecutive planes made of 64 vertical and 64 horizontal plastic scintillator slabs of 20 mm (0.10 X_0) thickness. The 128 counters are assembled into 4 quadrants

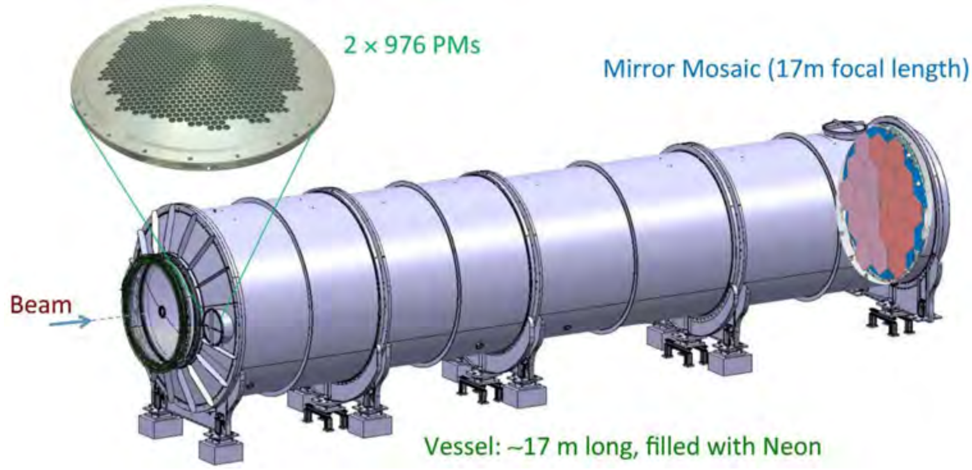

Figure 2.14: Layout of the RICH detector. A detailed view of one PM disk is shown on the left, while the mirror mosaic is seen through the neon container on the right. [33]

of 16 slabs in each plane. Slab lengths vary from 1210 mm (inner counters) to 600 mm (outer counters) forming an octagon of 1210 mm apothem. The slab widths are 65 mm in the central region close to the beam, where the particle flux is higher, and 99 mm in the outer region. The radius of the central hole crossed by the beam pipe is 128 mm (Fig. 2.15-left).

The CHOD active area is an array of 152 plastic scintillator tiles of 30 mm thickness covering an annulus with inner (outer) radii of 140 mm (1070 mm). The tiles are 108 mm high (except for 12 tiles near the external edge); laterally most tiles are either 134 mm or 268 mm wide. The tile centres are spaced vertically by 107 mm, resulting in a 1 mm overlap. The total thickness of the detector in the active area is 0.13 X_0 (Fig. 2.15-right).

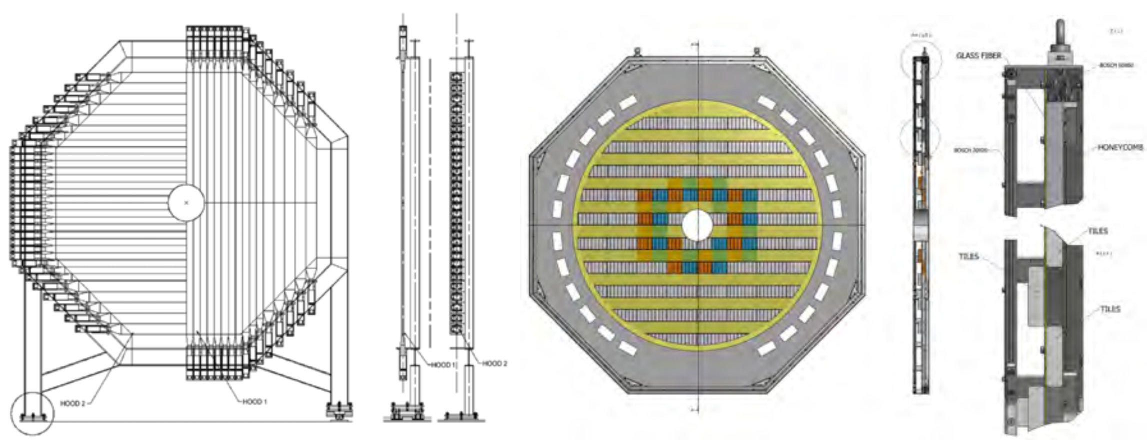

Figure 2.15: Left: NA48-CHOD detector. Right: CHOD detector. [33]

2.2.10 Hadron calorimeter (MUV1, MUV2)

The hadron calorimeter is a sampling calorimeter made from alternate layers of iron and scintillator corresponding to about 8 interaction lengths. The calorimeter is divided into two independent detectors: the front detector MUV1 and the back detector MUV2.

The MUV1 detector consists of 24 layers of 26.8 mm thick steel plates, corresponding to a total thickness of 4.1 interaction lengths (including the scintillating material). The innermost 22 layers have outer dimensions of 2700×2600 mm^2 while the first and the last layer are larger to cover the readout fibres and to serve as support of the whole structure. The distance between consecutive iron layers is 12 mm.

The MUV2 detector consists of 24 iron plates of 25 mm thickness, each followed by a layer of plastic scintillators.

2.2.11 Fast muon veto (MUV3)

The MUV3 detector, located downstream of the hadron calorimeter behind a 80 cm thick iron wall, provides fast L0 trigger signals and is used for muon identification; it detects charged particles traversing the whole calorimeter system (LKr, MUV1,2 and the iron wall) with a total thickness of over 14 interaction lengths. The MUV3 has a transverse size of 2640×2640 mm^2 and is built from 50 mm thick scintillator tiles, including 140 regular tiles of 220×220 mm^2 transverse dimensions and 8 smaller tiles adjacent to the beam pipe, as required by the high particle rate near the beam (Fig. 2.16).

2.2.12 Peripheral muon veto (MUV0)

The MUV0 detector is a scintillator hodoscope designed to detect π^- emitted in $K^+ \to \pi^+\pi^-\pi^+$ decays with momenta below 10 GeV/c, deflected towards positive X by the spectrometer magnet (which adds to the deflection of the K^+ beam by TRIM5 magnet), and leaving the lateral acceptance near the RICH. The MUV0 is mounted on the downstream flange of the RICH. Its active scintillator area of 1.4×1.4 m^2 covers the periphery of the lateral acceptance (1.545 m $< X <$ 2.945 m, $|Y| <$ 0.7 m), and consists of two layers of 48 plastic scintillator tiles with dimensions of $200 \times 200 \times 10$ mm^3.

2.2.13 HAdronic Sampling Calorimeter (HASC)

The HASC detector is used for the detection of π^+ emitted in $K^+ \to \pi^+\pi^-\pi^+$ decays with momentum above 50 GeV/c and propagating through the beam holes in the centres of the straw chambers. The detector is located downstream of MUV3 and the BEND dipole magnet (Fig. 2.5) which sweeps these pions out of the K^+ beam towards negative X. The detector covers the lateral acceptance of -0.48 m $< X < -0.18$ m, $|Y| < 0.15$ m. A second HASC2 station has been installed in 2021, mirroring the original HASC station, resulting in a symmetric detector at the end of the beamline.

2.2.14 Upgraded detector upstream layout

In order to reduce the background contamination to the measurement of the $K^+ \to \pi^+\nu\bar{\nu}$ decay, a modification of the beam layout upstream of the decay volume was made for the data taking period starting in 2021. The main changes were: a CEDAR upgrade for the KTAG (Section 2.2.1);

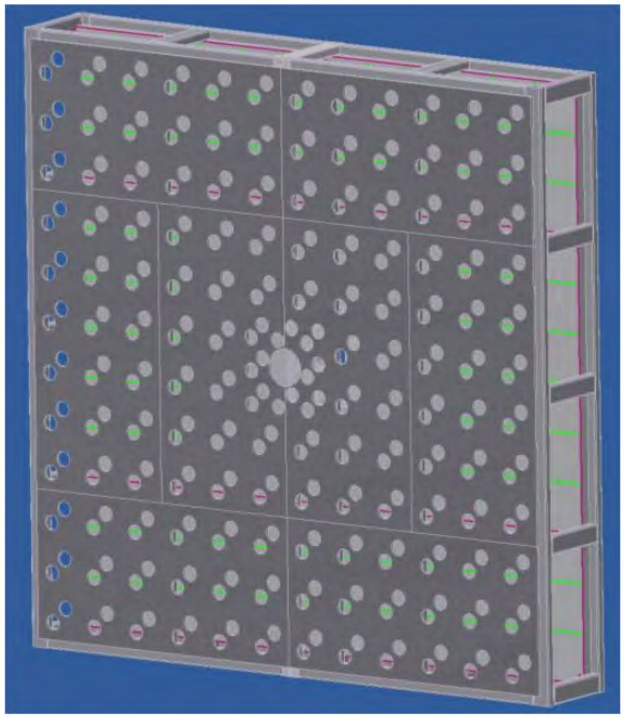

Figure 2.16: Layout of the MUV3 module. [36]

an optimized achromat; a 4th GTK station (GTK0), placed next to GTK1 (Section 2.2.2); a new VetoCounter around the beam pipe before and after the collimator and a new hodoscope (Fig. 2.17).

The VetoCounter was designed to significantly suppress upstream backgrounds from kaon decays occurring upstream of the fiducial volume, specifically between the GTK2 and final collimator. The VetoCounter is formed from three parallel planes of scintillating tile detectors: two before the final collimator, (VC1,2) separated by a 25 mm lead sheet, and one after the final collimator (VC3).

The Anti0 hodoscope was designed mainly to reduce the background present in dump mode searches. The detector consists of two planes of scintillating tiles readout by silicon photo-multipliers. It was installed in 2021 at the entrance of the fiducial volume.

2.3 Trigger and Data Acquisition system (TDAQ)

The NA62 trigger system operates in two successive stages. The first stage, the low-level hardware trigger (L0), utilizes custom programmable electronics to process data from a subset of detectors and generate an L0 decision. The second stage, the high-level trigger (L1), is implemented as software running on a dedicated data acquisition (DAQ) computing farm. Each stage is designed to reduce the event rate by approximately one order of magnitude, decreasing the initial K^+ decay rate of $\sim$10 MHz to a final storage rate of approximately 100 kHz.

Event selection is organized into trigger lines, each consisting of a sequence of L0 and L1 conditions tailored to specific event topologies. These conditions are implemented as either requirements or vetoes. A more detailed description is provided in [37].

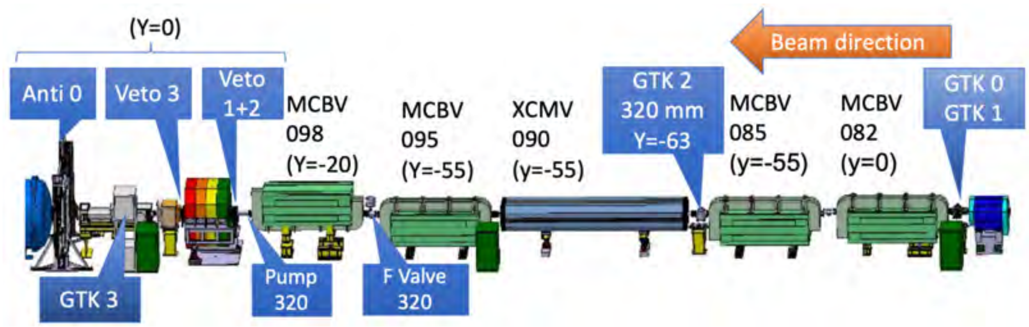

Figure 2.17: Schematic layout of the beamline since 2021. [38]

2.3.1 The low-level trigger

The L0 trigger incorporates inputs from five detectors. The RICH, CHOD, NA48-CHOD, and MUV3 detectors contribute via specialized firmware integrated into the readout hardware, which provides access to individual readout channels. The LKr calorimeter participates through a dedicated L0 calorimeter trigger system that processes aggregated information from groups of calorimeter readout modules.

2.3.2 The high-level trigger

The online computing framework comprises both the DAQ and HLT software components. This framework is deployed on a DAQ farm consisting of 30 computing nodes, where each node independently executes both software layers on a subset of the L0-accepted events.

Chapter 3

Search for HNL production in the $\pi^+ \to e^+N$

3.1 Experimental signature

The experimental objective of this analysis is the search for heavy neutral lepton (N) production via in-flight two-body leptonic decay of positive pions, defined by the process $\pi^+ \to e^+N$. The absolute kinematic phase space available for this search is fundamentally restricted by the mass difference between the parent pion and the downstream positron, dictating a strict upper bound on the HNL mass profile: $m_N < m_{\pi^+} - m_e = 139$ MeV/c^2. Within this allowed kinematic window, the framework assumes a proper lifetime for the sterile state exceeding $\tau_N > 50$ ns. Given the high-intensity secondary hadron beam optimized at a nominal momentum of 75 GeV/c, the produced HNLs experience an immense relativistic Lorentz boost of $\mathcal{O}(500)$. This kinematic dilation scales the characteristic laboratory decay length ($\lambda = \gamma\beta c\tau_N$) beyond 8 km, rendering the HNL effectively stable over the physical tracking dimensions of the experimental apparatus.

Consequently, the HNL does not undergo secondary decays into Standard Model (SM) particles within the active detector volume, operating as an entirely invisible state that manifests purely as missing four-momentum. The experimental signature of the signal process within the apparatus is thus uniquely defined by a single, fully reconstructed downstream positron track originating from a decay vertex within the decay region, spatially and temporally matched to an incoming beam track, with an absence of concurrent activity in the photon veto system.

This specific configuration shares an identical experimental topology with the SM electronic leptonic decay of the pion, $\pi^+ \to e^+\nu_e$ (referred to as the π_{e2} channel), where the active SM neutrino similarly escapes the apparatus undetected. Because of this topological equivalence, the search procedure measures the $\pi^+ \to e^+N$ decay rate with respect to the π_{e2} baseline channel, establishing a framework that cancels out major absolute systemic uncertainties.

To resolve the experimental degeneracy between these overlapping topologies and isolate the signal from the dominant SM background, the analysis employs the squared missing mass, m^2_{miss}, as the primary discriminating kinematic observable. This quantity is defined via the four-momentum of the undetected system, P_{miss}. Let P_π and P_e denote the four-momenta of the initial π^+ state and the downstream final-state positron, respectively. The missing four-momentum is defined as

$P_{\text{miss}} = P_\pi - P_e$. As a Lorentz vector, the squared norm of P_{miss} is a kinematic invariant. Specifically, for any four-momentum $P = (E, \vec{p})$, the square is given by:

$$P^2 = E^2 - |\vec{p}|^2 = m^2$$

Thus, $m^2_{\text{miss}} = P^2_{\text{miss}}$ provides a frame-independent observable to characterize the kinematics of the undetected system.

However, a major hardware constraint of the unseparated beam must be accounted for at this stage: while the beam comprises pions (70%) and kaons (6%), the upstream KTAG detector is optimized exclusively to tag kaons. Consequently, its tagging information is insensitive to the pion flux and is ignored at the analysis level.

Because the exact identity of the specific beam particle decaying in flight cannot be uniquely determined by upstream tagging, the primary kinematic observable is computed twice for every reconstructed event. Two parallel squared missing mass hypotheses are formulated by constructing the parent four-momentum using the three-momentum measured by the GTK combined with distinct mass hypotheses, and the positron four-momentum derived from the three-momentum measured with the STRAW:

$$m^2_{\text{miss}}(\pi) = (P_{\pi^+} - P_e)^2 \quad \text{and} \quad m^2_{\text{miss}}(K) = (P_{K^+} - P_e)^2$$

Here, the four-momenta P_π and P_K are constructed using the common three-momentum vector $\vec{p}_{\text{GTK}}$ measured by the GTK, assigned to the pion mass (m_{π^+}) and kaon mass (m_{K^+}) hypotheses, respectively:

$$P_\pi = \left(\sqrt{|\vec{p}_{\text{GTK}}|^2 + m^2_{\pi^+}},\, \vec{p}_{\text{GTK}}\right) \quad \text{and} \quad P_K = \left(\sqrt{|\vec{p}_{\text{GTK}}|^2 + m^2_{K^+}},\, \vec{p}_{\text{GTK}}\right)$$

The positron four-momentum P_e is similarly constructed from the three-momentum $\vec{p}_e$ measured by the STRAW and the positron mass m_e:

$$P_e = \left(\sqrt{|\vec{p}_e|^2 + m^2_e},\, \vec{p}_e\right)$$

Under the correct pion hypothesis, the SM neutrino signal concentrates into a sharp peak centered at $m^2_{\text{miss}}(\pi) = 0$. In contrast, the targeted signal process would manifest as a distinct peak at a positive, non-zero position, $m^2_{\text{miss}}(\pi) = m^2_N > 0$, corresponding to the squared mass of the HNL.

3.2 Data and Monte Carlo Samples

The analysis relies on two primary inputs: experimental data recorded by the physical detector and simulated templates generated to model both hypothetical signals and Standard Model background processes. This section details the specific data set employed, the initial filtering criteria applied to the raw data, and the Monte Carlo (MC) samples used to construct the theoretical models for comparison.

Table 3.1: `NA62RECO` versions used for each data group

Run1			**Run2**					
Sample	Reco	Spills	Sanple	Reco	Spills	Sample	Reco	Spills
2017A	v3.8.9	82k	2021A	v3.7.5	42k	2023A	v3.9.2	60k
2017B	v3.8.9	151k	2021B	v3.7.5	27k	2023B	v3.9.2	58k
2017C	v3.8.9	28k	2021C	v3.5.3	34k	2023C	v3.9.2	92k
2018A	v3.8.4	17k	2021D	v3.5.3	7k	2023D	v3.9.2	34k
2018B	v3.8.4	105k	2021E	v3.5.3	35k	2023E	v3.9.2	88k
2018C	v3.8.4	18k	2022A	v3.6.5	49k	2023F	v3.9.2	68k
2018D	v3.8.4	105k	2022B	v3.6.5	108k	2024A	v3.9.11	87k
2018E	v3.8.4	77k	2022C	v3.7.5	110k	2024B	v3.9.11	92k
2018F	v3.8.4	71k	2022D	v3.7.5	78k	2024C	v3.9.11	79k
2018G	v3.8.4	94k	2022E	v3.7.5	58k	2024D	v3.9.11	62k
2018H	v3.8.4	38k				2024E	v3.9.11	135k
						2024F	v3.9.15	112k

3.2.1 The Experimental Data Sample and Preliminary Selection

The data samples used for the analysis were recorded in 2017–2018 (Run1) and 2021–2024 (Run2), at a typical beam intensity of (2–4)$\times 10^{12}$ protons per spill corresponding to (330–580) $\times 10^6$ beam particles per second at the FV entrance, and a mean K^+ and π^+ decay rate in the FV of few 10^6 /s. The data samples of each year are divided in subsamples Table 3.1 shows all the subsamples used with their respective amount of spills recorded.

Prior to the application of the definitive analysis-level selection criteria, the dataset undergoes a minimal positron pre-selection using the REStricted POSitron (RESPOS) filter. While the core logic of this filter remains consistent across all data, slight variations in geometric and fiducial constraints exist due to differences in the reconstruction versions used for specific data-taking periods.

For the 2017-2018 and 2023–2024 datasets, the pre-selection requires geometric acceptance within the Liquid Krypton (LKr) calorimeter and restricts the nominal beam-axis vertex to a z-range between 102 and 180 m. In contrast, for the 2021-2022 datasets, the explicit geometric acceptance check for the LKr calorimeter is omitted, and the lower bound of the longitudinal vertex constraint is extended upstream to 90 m ($90 < z < 180$ m).

Across all periods, the selection universally requires full geometric acceptance within all four STRAW spectrometer stations, alongside consistent kinematic and quality criteria: a positive charge ($q = +1$), momentum below $70, \mathrm{GeV}/c$, an energy-to-momentum ratio of $E/p > 0.85$, and a track fit quality of $\chi^2 < 40$.

These minor variations in the pre-selection stage do not affect the final physics results. All datasets are subsequently subjected to the same definitive analysis-level selection criteria, which impose stricter, uniform constraints that supersede the initial filtering differences, ensuring a consistent and homogeneous event selection across the entire combined sample.

This initial filter provides a minimal positron track pre-selection.

3.2.2 Monte Carlo Samples

Monte Carlo (MC) simulations of particle interactions with the detector and its response, produced on the NA62 GRID using a software package based on the Geant4 toolkit [39], were used for both Run 1 and Run 2 data periods. In addition, accidental activity is simulated. The specific MC production versions and corresponding reconstruction (Reco) software releases employed for each sample are summarized in Table 3.2.

3.2.2.1 HNL Signal Samples

Signal MC samples for the process $\pi^+ \to e^+ N$ were generated using a multi-mass production strategy. This approach simulates the entire HNL mass spectrum (80–130 MeV/c^2) with a mass step size of 1 MeV/c^2 within a unified computational run. Consistent with the stability regime, the HNL is configured with an infinite lifetime, treating it as a stable particle that exits the detector as pure missing four-momentum. These samples are generated in the Standard decay region configuration. This methodology ensures a continuous and high-statistics mapping of the signal response across the full kinematic range without the need for independent generation campaigns for each mass hypothesis.

3.2.2.2 Normalization

A dedicated Monte Carlo sample of the Standard Model decay $\pi^+ \to e^+\nu_e$ provides the reference baseline. Generated in the Standard decay region, this process shares an identical experimental topology with the signal hypothesis a single positron track and missing four-momentum differing only in the mass of the invisible particle. This topological equivalence allows for the direct cancellation of systematic uncertainties related to to detector acceptance and reconstruction efficiency.

3.2.2.3 Background Samples

The decay $K^+ \to e^+\nu_e$ (K_{e2}): A dedicated Monte Carlo sample of the Standard Model decay $K^+ \to e^+\nu_e$ is employed. Generated in the Standard decay region, this process shares the same signal topology; since the exact identity of the decaying beam particle cannot be uniquely determined, the sample serves as the essential normalization for the kaon background component.

The decay chain $\pi^+/K^+ \to \mu^+ \to e^+$: Monte Carlo simulations model the sequential decay chain $\pi^+/K^+ \to \mu^+\nu_\mu$ followed by $\mu^+ \to e^+\nu_e\bar{\nu}_\mu$ across two distinct spatial configurations: the Upstream region (69.175–102.425 m), where the parent meson decays prior to the fiducial volume, and the Standard region (102.425–180.0 m), where the decay occurs within the fiducial volume. To enhance statistical precision for the rare muon-decay-in-flight component, all four samples employ a capped muon lifetime technique that biases generation toward shorter proper times, specifically limiting the simulated proper lifetime to a small fraction of the mean muon lifetime ($\tau_\mu = 2.2$ μs), while applying a corresponding weight correction to maintain physical accuracy. The specific vertex definitions and integrated lifetime caps for each sample are detailed below:

Upstream $\pi^+ \to \mu^+\nu_\mu$ ($\pi^{ups}_{\mu 2}$): Parent π^+ decays in the region $Z \in [69.175, 102.425]$ m, with the simulated muon proper lifetime capped at 0.07% of τ_μ.

Table 3.2: Monte Carlo samples used for the analysis with their respective `NA62MC` and `NA62RECO` versions.

Decay Mode	**Run1** MC-RECO	**Run2** MC-RECO
$\pi^+ \to e^+\nu$	v3.3.0-v3.7.3	v3.7.6-v3.7.6
$\pi^+ \to \mu^+\nu$, $\mu^+ \to e^+\nu\bar{\nu}$	v3.5.1-v3.7.3	v3.8.9-v3.8.9
$\pi^+ \to \mu^+\nu$ upstream, $\mu^+ \to e^+\nu\bar{\nu}$	v3.5.1-v3.7.3	v3.8.9-v3.8.9
$K^+ \to e^+\nu$	v3.9.9-v3.9.9	v3.7.6-v3.7.6
$K^+ \to \mu^+\nu$, $\mu^+ \to e^+\nu\bar{\nu}$	v3.9.9-v3.9.9	v3.8.9-v3.8.9
$K^+ \to \mu^+\nu$ upstream, $\mu^+ \to e^+\nu\bar{\nu}$	v3.5.1-v3.7.3	v3.8.9-v3.8.9
HNL ($\pi^+ \to e^+N$)	v3.9.13-v3.9.13	v3.9.13-v3.9.13
MuonHalo	v2.1.5	

Standard $\pi^+ \to \mu^+\nu_\mu$ ($\pi_{\mu 2}$): Parent π^+ decays inside the fiducial volume $Z \in [102.425, 180.0]$ m, with the simulated muon proper lifetime capped at 0.06% of τ_μ.

Upstream $K^+ \to \mu^+\nu_\mu$ ($K_{\mu 2}^{ups}$): Parent K^+ decays in the region $Z \in [69.175, 102.425]$ m, with the simulated muon proper lifetime capped at 0.89% of τ_μ.

Standard $K^+ \to \mu^+\nu_\mu$ ($K_{\mu 2}$): Parent K^+ decays inside the fiducial volume $Z \in [102.425, 180.0]$ m, with the simulated muon proper lifetime capped at 0.73% of τ_μ.

Muon Halo: This sample models prompt muons produced at the target or via early interactions upstream of the KTAG, which travel almost parallel with the beam and decay in flight ($\mu^+ \to e^+\nu_e\bar{\nu}_\mu$) within the detector acceptance. Unlike the π/K samples, the spatial definition for this background is based on the halo trajectory rather than a specific fiducial decay vertex. Consistent with the other background samples, this simulation utilizes a capped muon lifetime technique to enhance the statistical precision of decays occurring within the active volume. The capped muon lifetime, $\tau_{\text{cap}}(p)$, is momentum-dependent, while applying a corresponding weight correction to maintain physical accuracy.

3.3 Event selection and Background suppression

The primary experimental objective of the event selection is the isolation of a high-purity sample conforming to the two-body leptonic decay topology $\pi^+ \to e^+N$, while suppressing Standard Model background processes by multiple orders of magnitude. The selection targets a final state characterized by a single, fully reconstructed positron track originating from a valid decay vertex within the fiducial volume, with no unassociated detector activity. This section details the trigger employed, the offline cuts, the particle identification (PID) criteria, the active veto constraints and the matching between the deownstream particle and the parent beam particle that define the signal selection.

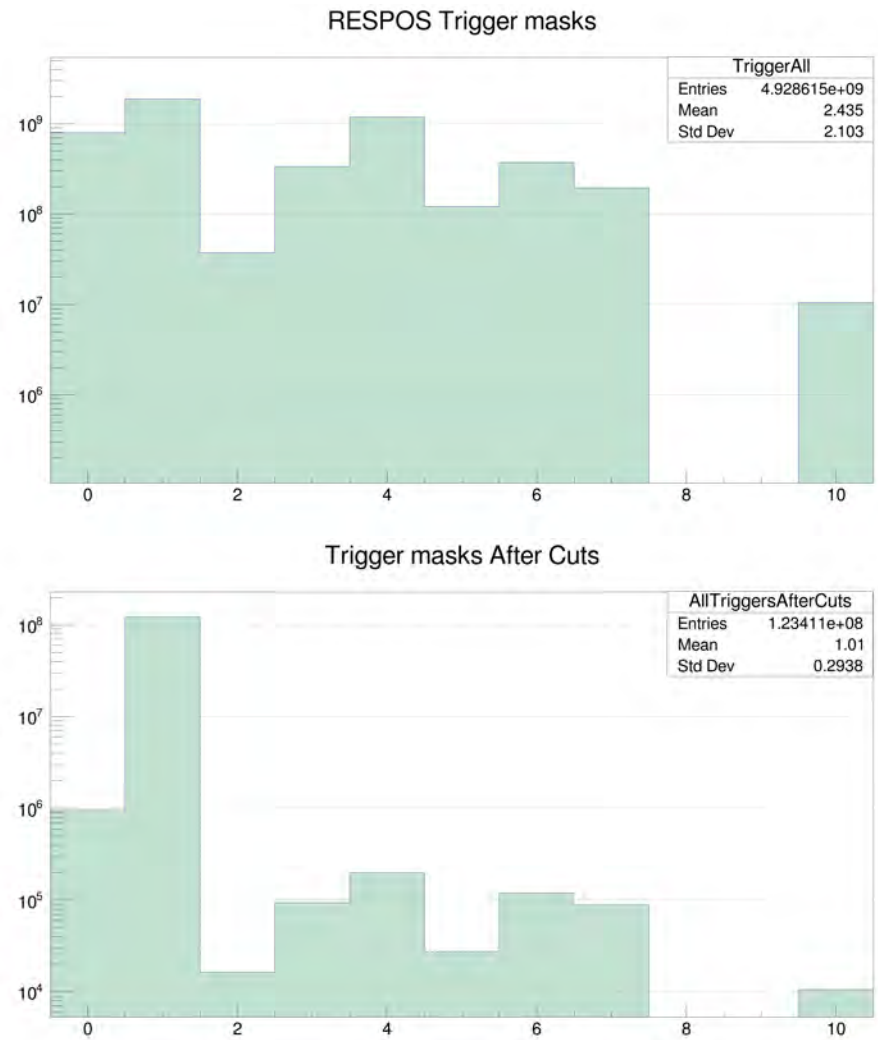

Figure 3.1: Number of events per trigger line in RESPOS filter before applying selection (top) and after applying selection (bottom) for Run1 data. Trigger line 1 represents the PNN trigger.

3.3.1 Trigger Conditions

The analysis utilizes the main trigger line dedicated to the $K^+ \to \pi^+\nu\bar{\nu}$ measurement (PNN trigger). This selection is motivated by the event yield distribution across all available trigger lines in the dataset processed by the REStricted POSitron (RESPOS) filter. As demonstrated in Figure 3.1 which display the number of events per trigger line before and after the downstream particle selection criteria, the PNN line provides the overwhelming majority of the surviving statistics. Consequently, restricting the analysis to this line maximizes the sensitivity for the HNL search in the $\pi^+ \to e^+N$

channel.

The PNN trigger operates via a two-level hierarchy:

Level-0 (L0) Hardware Trigger. The L0 decision requires a specific coincidence of signals to identify a candidate positron while rejecting high-rate backgrounds:

- **RICH:** At least two signals are required to ensure charged particle identification.
- **CHOD:** At least one signal in any quadrant is required, with a no signals on diagonally-opposite quadrants and a requirement of fewer than five total signals to reject multi-track events.
- **MUV3:** A veto condition (no signals) is applied to reject muons.
- **LKr Calorimeter:** An energy deposit between 5 GeV and 30 GeV is required for 2017–2018 data, extended to 40 GeV for the 2021–2024 dataset.

Level-1 (L1) Software Trigger. Events passing L0 are further filtered by L1 conditions:

- **KTAG:** A kaon signal must be identified within a 5 ns window of the L0 trigger time. Approximately one-third of beam pions are accepted at this stage due to accidental time-coincidences with a beam kaon, a critical factor for the π^+ decay search.
- **STRAW:** The trigger requires a positively charged track with momentum $p < 50$ GeV/c, compatible with originating from a beam particle decay within the Fiducial Volume (FV).
- **LAV:** A veto is applied requiring fewer than two signals in stations LAV2–11 to suppress events with extra photons.

For Monte Carlo (MC) samples, the response of the L0 trigger conditions is emulated to ensure consistency with the data acquisition logic.

3.3.2 Downstream Particle

The reconstruction of the downstream particle aims to isolate a single positron candidate consistent with the $\pi^+ \to e^+N$ decay topology. The selection strategy proceeds sequentially to suppress dominant backgrounds, starting with the identification of a high-quality, isolated track in the STRAW spectrometer. This candidate is then subjected to particle identification criteria using the LKr and RICH detectors to confirm its positron hypothesis and establish a precise reference time. Finally, the event is required to be absent of any additional electromagnetic activity in the LKr, LAV, and SAV detectors within a narrow time window, ensuring the absence of accompanying photons.

3.3.2.1 Good Track Criteria

The good track selection identifies a single, high-quality track per event. The track must carry a positive charge ($q = +1$) and possess hits in all four STRAW spectrometer chambers with full geometric acceptance. The momentum is restricted to the range $5 < p < 30\,\mathrm{GeV}/c$, directly reflecting the L0 LKr energy trigger conditions. While these L0 conditions differed between the

2017–2018 and 2021–2024 data taking periods, this variation has no impact on the sensitivity for the HNL search in the region of interest. Additional constraints limit the distance of closest approach to the beam axis to less than 25 mm and the reconstructed vertex position to between 100 and 180 m. Furthermore, the track is vetoed if it is associated with any other STRAW track in a reconstructed vertex, this ensures excluding multi-body decays.

This definition also incorporates the previously outlined RESPOS conditions. Including these criteria serves as an initial filtering step analogous to the one applied to real data, ensuring that Monte Carlo simulations undergo the same fundamental acceptance requirements before any subsequent, stricter analysis cuts are applied.

Consistent with the signal topology of a beam particle decaying into a single charged particle and an invisible neutrino, the criteria strictly requires exactly one good track per event; events containing zero or multiple good tracks are discarded.

3.3.2.2 Particle Identification (PID)

The identification of the downstream track relies on a redundant combination of measurements from the LKr electromagnetic calorimeter and the RICH detector. To be classified as a positron, the track must be associated with an energy cluster in the LKr calorimeter such that the ratio of deposited energy to measured momentum falls within the range $0.9 < E/p < 1.1$. This specific interval isolates the electromagnetic shower profile characteristic of positrons while rejecting pions and muons.

Complementing the calorimetric measurement, a successful association in the RICH detector is required, where the reconstructed Cherenkov ring must be consistent with a positron as the most likely hypothesis. The quality of this identification is quantified by the maximum probability assigned to any non-positron hypothesis (e.g., π^+ or μ^+). As illustrated in Figure 3.2, the distribution of this variable for selected events accumulates sharply at zero, confirming that alternative particle hypotheses are strongly disfavored and ensuring high purity in the positron sample. Furthermore, the RICH detector provides a precise time measurement for the positron track (t_{track}), which is critical for coincidence checks.

3.3.3 Photon Veto

To ensure the absence of additional photon activity, the selection vetoes extra energy deposits in the calorimeters relative to the positron candidate identified in the previous step. Using the precise arrival time from the RICH detector (t_{track}) as a reference, the following conditions are applied:

- **LKr Cluster Veto:** No extra energy clusters are allowed in the Liquid Krypton (LKr) calorimeter within 4 nanoseconds of the positron time, excluding the single cluster matched to the positron track. Figure 3.3 shows the time difference distribution for these extra clusters before the cut is applied. A prominent peak appears at zero time difference due to unrelated activity; the 4-nanosecond window is chosen to fully reject this peak.
- **LAV Veto:** No signals are allowed in the Large Angle Veto detectors within 5 nanoseconds.
- **SAV Veto:** No signals are allowed in the Small Angle Veto detectors (IRC and SAC) within 5 nanoseconds.

Events containing any extra energy deposits in these calorimeters within the specified windows are rejected.

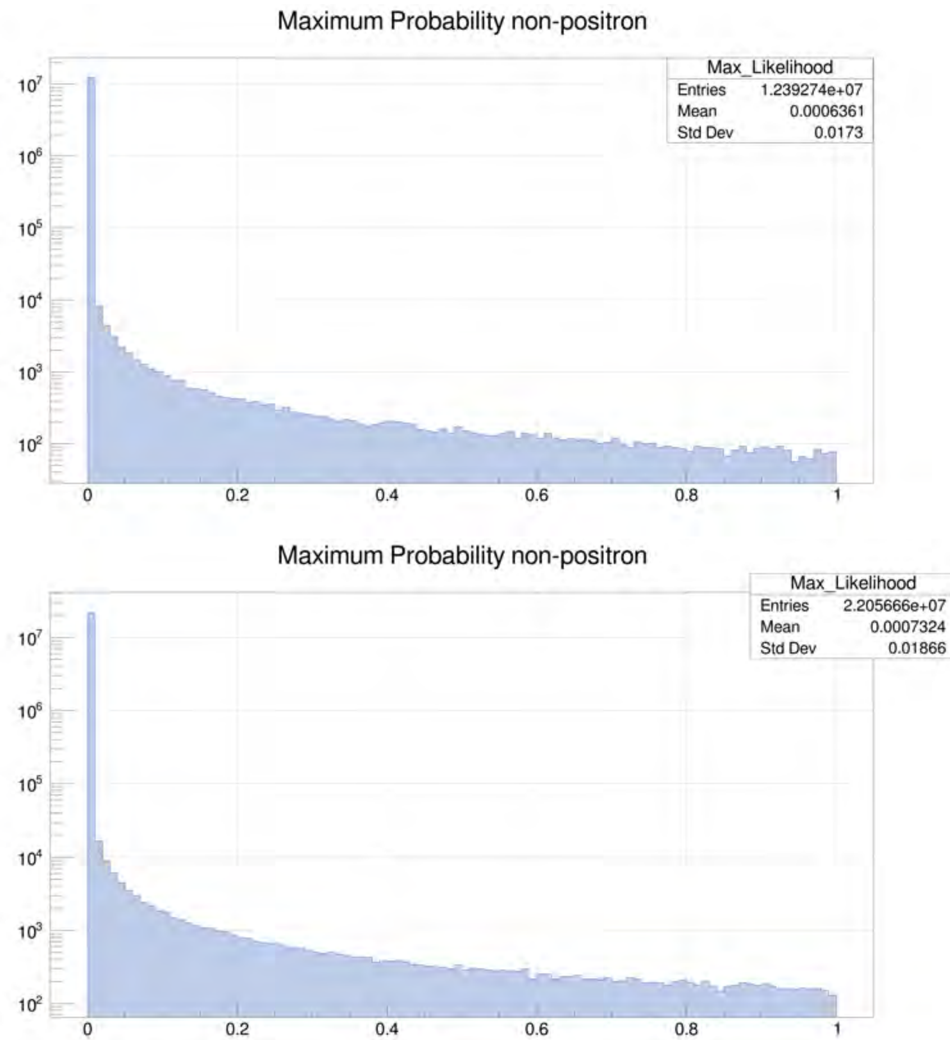

Figure 3.2: Maximum RICH probability among non-positron hypotheses for the group 2018E (top) and 2022B (bottom)

3.3.4 Downstream Positron and Beam Perticle Track Matching

To associate a positron track with its corresponding reconstructed GTK track, the offline matching algorithm (`SpectrometerGigaTrackerMatching`) is applied based on timing and the Blue Field-corrected Closest Distance of Approach (CDA).

For a GTK track to be successfully matched, it must first satisfy basic quality and kinematic filters: it must be consistent with the nominal beam profile in momentum and direction, have a good track fit quality (χ^2), and fall within a 0.5 ns timing window relative to the reference time.

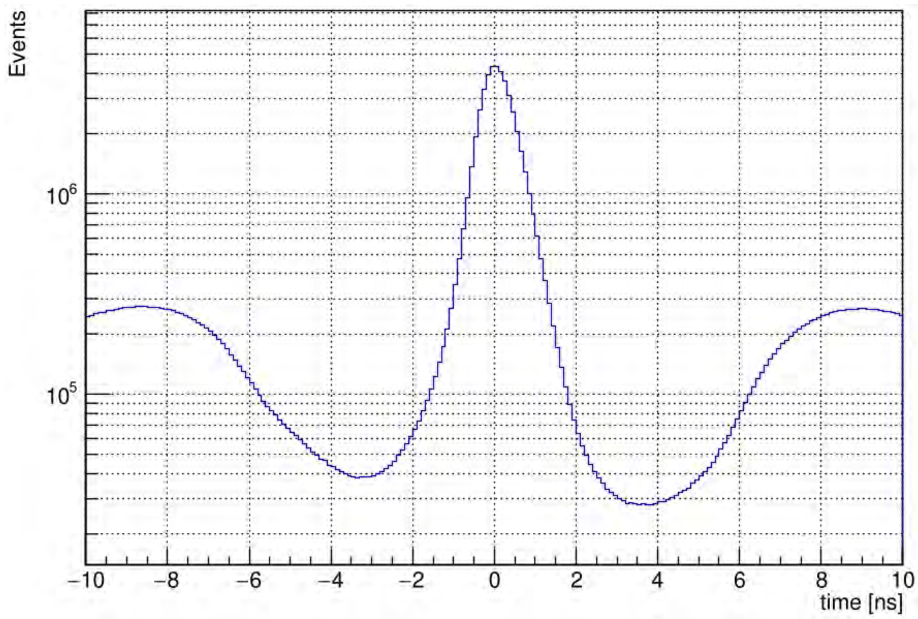

Figure 3.3: Time difference between the track and the LKr energy deposits not spatially associated with the track before the event selection.

For surviving candidates, spatial compatibility is determined by propagating both the GTK and STRAW trajectories through the magnetic field upstream of the decay volume to calculate their CDA. The algorithm then evaluates a joint spatio-temporal discriminant, $D(\Delta t, \text{CDA})$, defined as the product of complementary cumulative probabilities for the timing agreement and spatial separation. If multiple GTK candidates satisfy the quality conditions and 0.5 ns timing window for the STRAW track, the algorithm selects the best candidate the one that gives the highest value of $D(\Delta t, \text{CDA})$.

Finally, to ensure a clean sample of matched events, the selected best-matched candidate pair must satisfy $D > 0.01$ and CDA < 3 mm. This 3 mm limit was chosen to maximize the search sensitivity, defined by the ratio $N_S/\sqrt{N_B}$. Here, N_S represents the number of observed π_{e2} decays (which share the same topology as the signal) measured in the peak region of the missing mass spectrum, while N_B represents the background level estimated from a separate region dominated by background fluctuations. The optimization scan showing this figure of merit as a function of the CDA threshold is presented in Fig. 3.3.3, where the maximum sensitivity is clearly observed at the selected value. Additionally, the reconstructed vertex position calculated from this best match is required to lie within the fiducial decay region between 105 m and 170 m.

3.3.5 Overview of Selected Distributions

To illustrate the kinematic features of the events passing the full selection chain, the squared missing mass $m^2_{miss}(\pi)$ distributions are displayed for the data and key MC samples. This variable, separates the different event topologies.

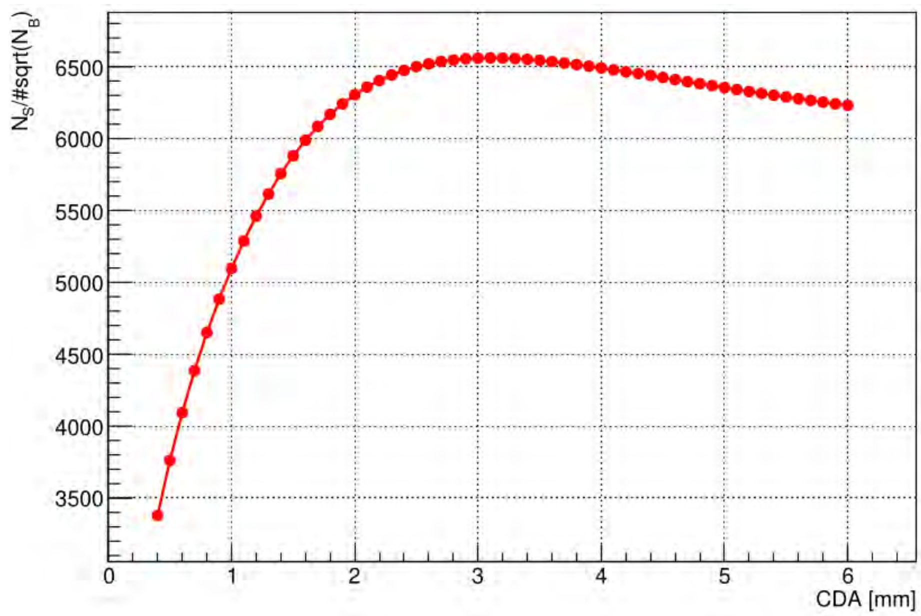

Figure 3.4: Optimization of the Closest Distance of Approach (CDA) cut. The figure of merit, defined as $N_S/\sqrt{N_B}$, is scanned over the range 0.4–6 mm. Sensitivity peaks at CDA = 3 mm, representing the optimal balance between keeping signal events and rejecting background noise.

Figure 3.3.5 shows the distribution for the 2018E data sample. The spectrum is dominated by a sharp peak at $m^2_{miss} = 0$, with a visible population of events extending into the positive region.

Figure 3.3.5 displays the normalization channel ($\pi^+ \to e^+\nu_e$) from MC simulation. The distribution forms a narrow peak centered at zero, corresponding to the two-body decay kinematics with a massless neutrino.

Figure 3.3.5 presents the signal hypothesis for a Heavy Neutral Lepton with mass $m_N = 100\,\mathrm{MeV}/c^2$. The distribution appears as a distinct peak shifted to positive m^2_{miss} values, reflecting the non-zero mass of the hypothetical particle.

Figure 3.3.5 shows the $\pi^+ \to \mu^+\nu_\mu$ sample from MC simulation. In this case, the downstream decay of the muon produces a positron that satisfies the selection criteria. The resulting distribution is broad and located entirely in the positive m^2_{miss} region, distinct from the peak at zero.

3.3.6 Exploratory Criteria and Supplemental Selection Studies

In addition to the final selection criteria, several alternative approaches were investigated to further reduce background events. These studies examined kinematic variables and additional detector responses as part of an iterative process to improve signal purity.

These methods were ultimately not included in the final analysis. Some introduced random veto effects, which reduced the overall selection efficiency. Others showed no clear separation between signal and background distributions, rendering them ineffective for discrimination.

A description of these studies, is provided in Appendix A.

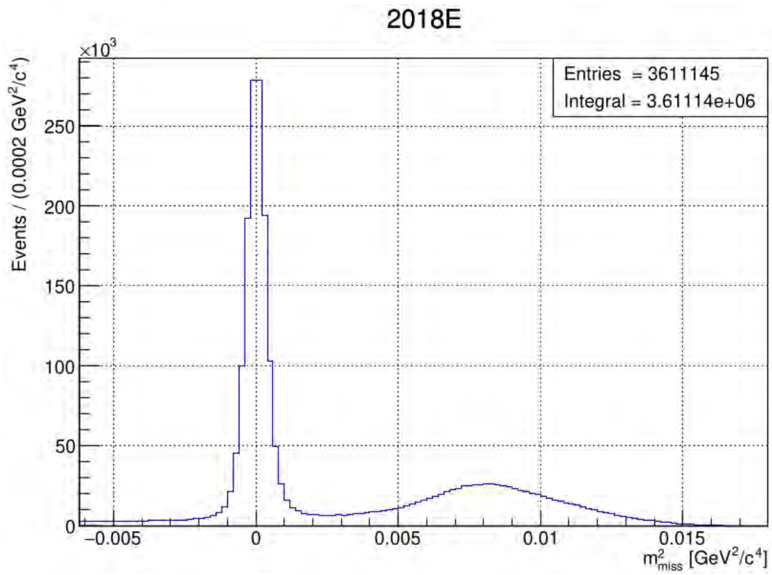

Figure 3.5: Squared missing mass $m^2_{miss}(\pi)$ distribution for the 2018E data sample after full event selection. The spectrum is dominated by the $\pi^+ \to e^+\nu_e$ peak at zero, with visible contributions from background processes in the positive region.

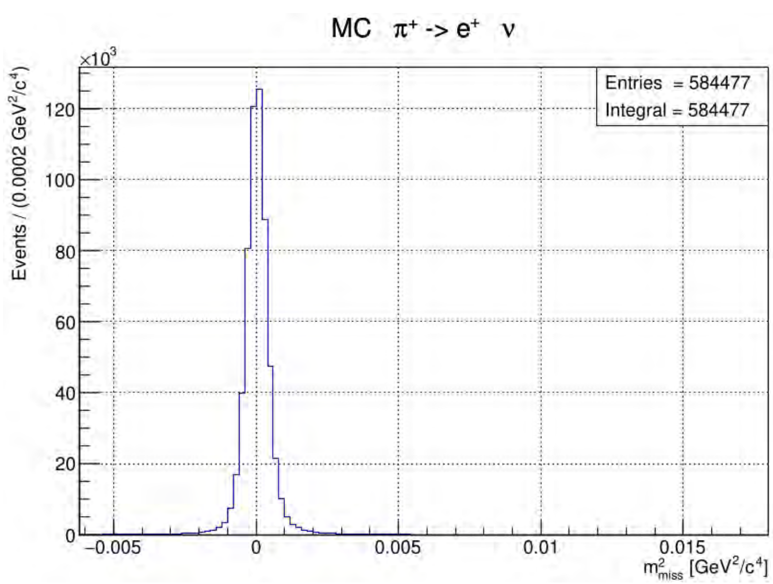

Figure 3.6: Monte Carlo simulation of the normalization channel $\pi^+ \to e + \nu_e$. The distribution shows a narrow peak centered at $m^2_{miss}(\pi) = 0$, characteristic of the two-body decay kinematics with a massless neutrino.

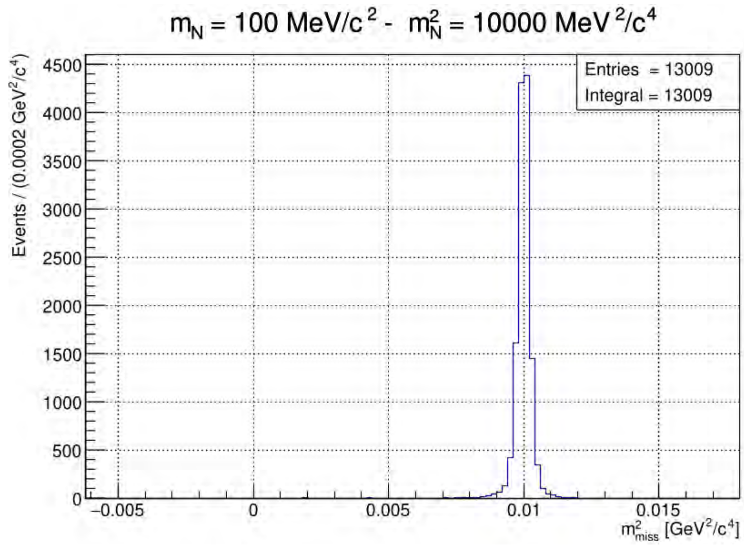

Figure 3.7: Signal hypothesis distribution for a Heavy Neutral Lepton (HNL) with mass $m_N = 100$,MeV/ c. The distinct peak shifted to positive $m^2_{miss}(\pi)$ values reflects the non-zero mass of the hypothetical particle.

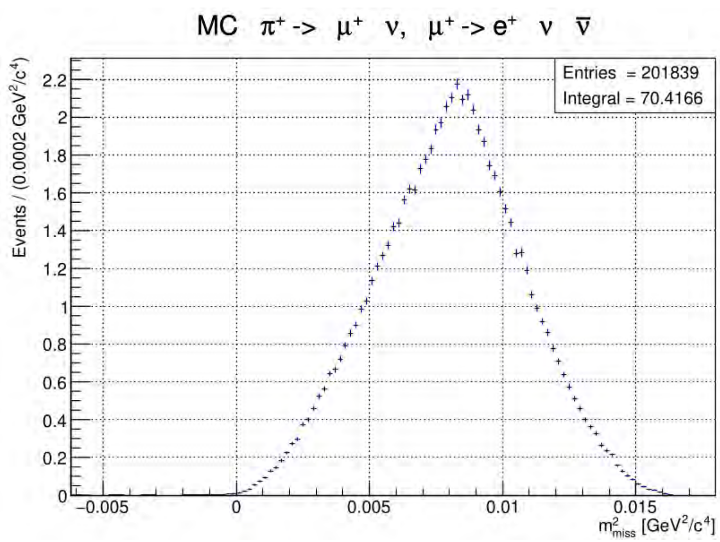

Figure 3.8: Monte Carlo simulation of the $\pi^+ \to \mu^+\nu_\mu$ background channel, where the muon decays downstream ($\mu^+ \to e^+\nu_e\bar{\nu_\mu}$). The resulting broad distribution lies entirely in the positive $m^2_{miss}(\pi)$ region. Note the difference between the selected event count and the integral value; this discrepancy arises from the muon capped lifetime technique, where events are assigned weights to correctly model the decay.

3.4 Normalization

The sensitivity for HNLs in $\pi^+ \to e^+ N$ decays relies fundamentally on the precise determination of the total number of pion decays (N_π) within the fiducial volume. The analysis utilizes the Standard Model leptonic decay $\pi^+ \to e^+\nu_e$ (π_{e2}) as a normalization channel, leveraging its identical topological signature to the signal.

A normalization region is defined as the interval $|m^2_{\text{miss}}(\pi)| < 0.002\ \text{GeV}^2/c^4$ centered on the π_{e2} peak. However, this region is not purely composed of pion decays; a kaon component kinematically overlaps with the pion normalization region when reconstructed under the pion hypothesis. This kaon contamination biases the raw event count in the pion normalization region. To correct for this, the total kaon flux must be independently determined to quantify and subtract the background component. This is achieved by defining a separate normalization region on the kaon missing mass squared spectrum around the $K^+ \to e^+\nu_e$ (K_{e2}) peak, defined as $|m^2_{\text{miss}}(K)| < 0.01\,\text{GeV}^2/c^4$.

3.4.1 Kaon Flux Determination

To account for the kaon contamination in the pion normalization region, the total kaon flux is determined independently using the $m^2_{\text{miss}}(K)$ spectrum. By selecting a clean signal region dominated by the $K^+ \to e^+\nu_e$ (K_{e2}) peak, the total number of kaon decays (N_{K^+}) is extracted separately from the pion sample. The yield is calculated as:

$$N_{K^+} = \frac{N^{K^+}_{SM}}{\sum_i A^K_i \cdot \mathcal{B}^K_i} \tag{3.1}$$

The inputs to N_{K^+} computation are displayed in table 3.3.

Table 3.3: Inputs to the computation of the number of kaon decays

	Ke2	
	Run1	Run2
$N^{K^+}_{SR}$	3 654 507	7 659 890
$A^{K^+}_{e^+}$	$(3.962 \pm 0.006) \times 10^{-2}$	$(3.667 \pm 0.006) \times 10^{-2}$
$\mathcal{B}(K^+ \to e^+\nu)$	$(1.582 \pm 0.007) \times 10^{-5}$	$(1.582 \pm 0.007) \times 10^{-5}$
$A^{K^+}_{\mu^+}$	$(7.68 \pm 1.41) \times 10^{-10}$	$(1.01 \pm 0.11) \times 10^{-9}$
$\mathcal{B}(K^+ \to \mu^+\nu)$	0.6356 ± 0.0011	0.6356 ± 0.0011

To visualize the input data for the flux determination, Figures 3.9 and 3.10 display the squared missing mass ($m^2_{\text{miss}}(K)$) spectra for the 2017–2018 (Run 1), 2021–2024 (Run 2), and the combined dataset. Each panel shows the full distribution of data events overlaid with the stacked Monte Carlo simulation (histograms). The normalization region used to extract the yield $N^{K^+}_{SM}$ is explicitly delimited by vertical arrows in each plot, centered on the K_{e2} peak.

The uncertainty on N_{K^+} is derived using standard error propagation for independent variables.

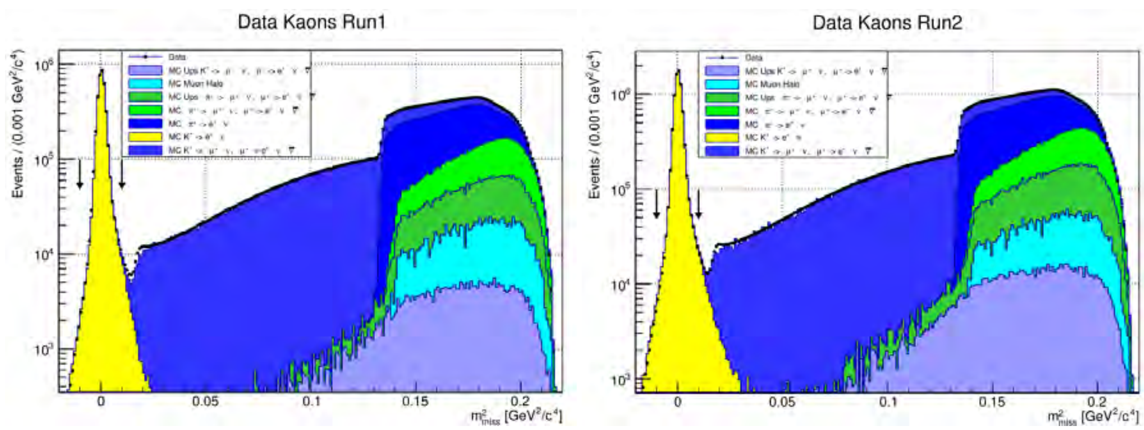

Figure 3.9: $m^2_{miss}(K)$ spectrum for data and MC for Run1 (left), and Run2 (right). The black arrows defined the normalization region.

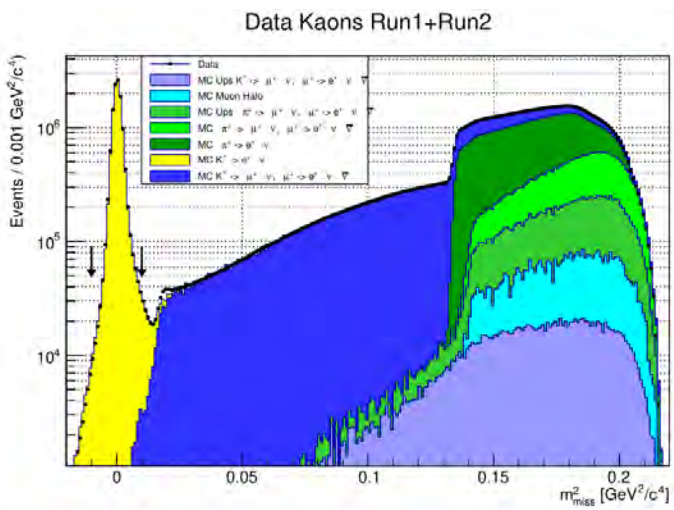

Figure 3.10: $m^2_{miss}(K)$ spectrum for data and MC for the full data set Run1+Run2. The black arrows defined the normalization region used for Run1 and Run2.

For a function $f(x_i)$, the error is defined as:

$$\Delta f = \sqrt{\sum_i \left(\frac{\partial f}{\partial x_i} \cdot \Delta x_i \right)^2}. \tag{3.2}$$

Applying this to Eq. 3.1, the relative uncertainty on the number of kaon decays is computed as:

$$\Delta N_{K^+} = N_{K^+} \sqrt{\frac{1}{N_{SM}^{K+}} + \frac{1}{(\sum A_i^K \mathcal{B}_i^K)^2} \sum_i \left((A_i^K)^2 (\Delta \mathcal{B}_i^K)^2 + (\mathcal{B}_i^K)^2 (\Delta A_i^K)^2 \right)}, \tag{3.3}$$

where $\Delta N_{SR}^{K^+} = \sqrt{N_{SR}^{K^+}}$ is the statistical error on the yield, and ΔA_i^K and $\Delta \mathcal{B}_i^K$ are the uncertainties on the acceptance and branching fractions, respectively. The branching fractions and their errors are taken from the Particle Data Group (PDG) [10].

3.4.1.1 Resolution Study and Window Validation

The choice of the normalization window width (6σ) is validated by comparing the squared missing mass resolution between Data and Monte Carlo (MC) simulations. Table 3.4 summarizes the K_{e2} peak resolutions for Run1 and Run2. The MC resolution differs from data by 2% and 3.6% for Run1 and Run2, respectively, confirming that the simulation accurately models the detector response.

Table 3.4: Squared missing mass resolutions (GeV^2/c^4) on the K_{e2} peak for data and MC.

Sample	**Run1**	**Run2**
Data	1.544×10^{-3}	1.575×10^{-3}
MC	1.512×10^{-3}	1.518×10^{-3}

3.4.1.2 Systematic Uncertainty Evaluation

The systematic uncertainty is evaluated by testing the stability of N_{K^+} with respect to variations in the half-width of the $m^2_{\text{miss}}(K)$ signal window. The window size is varied from $0.003\,\text{GeV}^2/c^4$ ($\sim 2\sigma$) to $0.016\,\text{GeV}^2/c^4$ ($\sim 10\sigma$), where the upper limit corresponds to the onset of $K^+ \to \mu^+\nu_\mu$ ($K_{\mu 2}$) dominance.

For Run2, the maximum variation of N_{K^+} with respect to the nominal window ($0.01\,\text{GeV}^2/c^4$) is observed at a half-width of $0.0045\,\text{GeV}^2/c^4$ and is assigned as the systematic error. For Run1, N_{K^+} exhibits larger instability due to discrepancies in Data/MC agreement, visible as an increasing ratio in the wider windows (see Fig. 3.11). Consequently, the systematic uncertainty for Run1 is defined by the maximum variation observed across the full range from $0.0045\,\text{GeV}^2/c^4$ to $0.016\,\text{GeV}^2/c^4$. The resulting statistical and systematic uncertainties are summarized in Table 3.5.

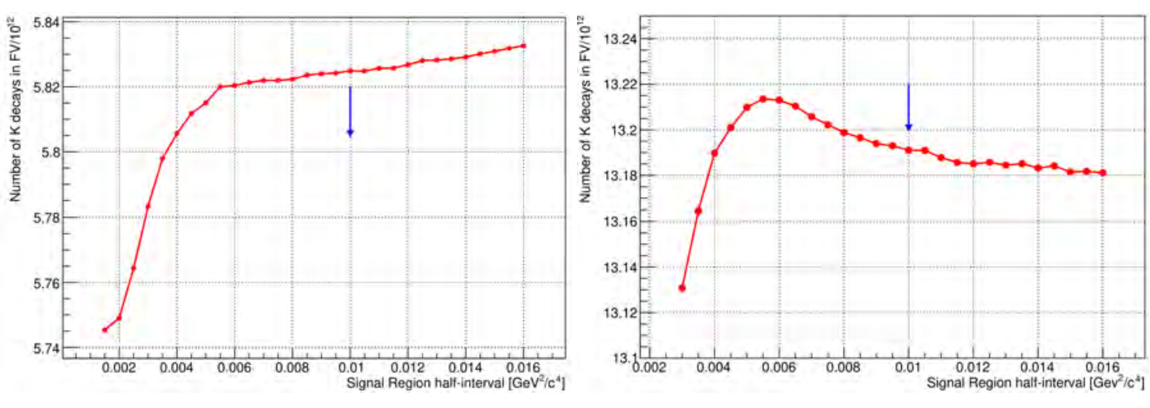

Figure 3.11: Stability checks of N_{K^+} with respect to the variation of the signal region for Run1 (left) and Run2 (right)

Table 3.5: Summary of statistical and systematic uncertainties on N_{K^+} for all data sets.

Error Type	**Run1**	**Run2**	**Run1+Run2**
Statistical	0.47%	0.47%	0.36%
Systematic	0.36%	0.25%	0.20%
Total	0.59%	0.53%	0.41%

3.4.1.3 Final Kaon Flux

The determined number of kaon decays for each period is:

$$N_{K^+}^{\text{Run1}} = (5.82 \pm 0.03) \times 10^{12},$$

$$N_{K^+}^{\text{Run2}} = (13.19 \pm 0.07) \times 10^{12}.$$

Combining both periods, the total kaon flux used for normalization is:

$$N_{K^+}^{\text{Total}} = (19.02 \pm 0.08) \times 10^{12},$$

where the quoted uncertainty includes both statistical and systematic components. This value serves as the primary normalization factor to quantify the kaon contamination in the $m^2_{\text{miss}}(\pi^+)$ sample.

3.4.2 Pion Flux Determination

Using the total kaon flux N_{K^+} derived in Sec. 3.4.1, the expected number of kaon events leaking into the pion normalization region is calculated. This contamination, coming from $K^+ \to \mu^+\nu_\mu$ decays and then muon decaying-in-flight through $\mu^+ \to e^+\nu_e\bar{\nu_\mu}$, is estimated as:

$$N^K_{\text{Bckg}} = \sum_i N^K_i = \sum_i \left(N_{K^+} \cdot \mathcal{B}^K_i \cdot A^{K\to\pi}_i\right) \tag{3.4}$$

where $A^{K\to\pi}_i$ represents the acceptance for a kaon decay mode i to pass the pion selection criteria and reconstruct within the pion normalization region ($|m^2_{\text{miss}}(\pi)| < 0.002\,\text{GeV}^2/c^4$).

The purified number of pion decays (N_{π^+}) is obtained by subtracting this estimated contamination from the total observed yield in the pion normalziation region ($N^{\pi^+}_{SM}$):

$$N_{\pi^+} = \frac{N^{\pi^+}_{SM} - N^K_{\text{Bckg}}}{\sum_j A^\pi_j \cdot \mathcal{B}^\pi_j} \tag{3.5}$$

The inputs to the N_{π^+} computation are displayed in table 3.6.

To visualize the inputs for the pion flux determination, Figures 3.12 and 3.13 display the squared missing mass ($m^2_{\text{miss}}(\pi)$) spectra for the 2017–2018 (Run 1), 2021–2024 (Run 2), and the combined dataset. Each panel shows the full distribution of data events overlaid with the stacked Monte Carlo simulation (histograms), which includes the dominant π_{e2} signal and the estimated contamination. The plot features two distinct marked regions: the normalization region, delimited by black arrows centered on the π_{e2} peak, which is used to extract the value $N^{\pi^+}_{SM}$; and the HNL search region, delimited by red arrows at higher mass values, where the signal search is performed.

Table 3.6: Inputs to the computation of the number of pion decays

Pie2		
	Run1	Run2
$N_{SM}^{\pi^+}$	13 496 089	33 515 907
N_{Km2}	478 687 $\pm$ 3 428	1 046 200 $\pm$ 9 118
N_{Km2ups}	37 674 $\pm$ 1 281	145 723 $\pm$ 3 732
$A_{e^+}^{\pi^+}$	$(6.070 \pm 0.008) \times 10^{-2}$	$(5.394 \pm 0.007) \times 10^{-2}$
$\mathcal{B}(\pi^+ \to e^+\nu)$	$(1.230 \pm 0.004) \times 10^{-4}$	$(1.230 \pm 0.004) \times 10^{-4}$
$A_{\mu^+}^{\pi^+}$	$(3.35 \pm 0.06) \times 10^{-8}$	$(3.40 \pm 0.06) \times 10^{-8}$
$A_{\mu^+(ups)}^{\pi^+}$	$(3.68 \pm 0.13) \times 10^{-8}$	$(4.69 \pm 0.14) \times 10^{-8}$
$\mathcal{B}(\pi^+ \to \mu^+\nu)$	$(99.98770 \pm 0.00004)\%$	$(99.98770 \pm 0.00004)\%$

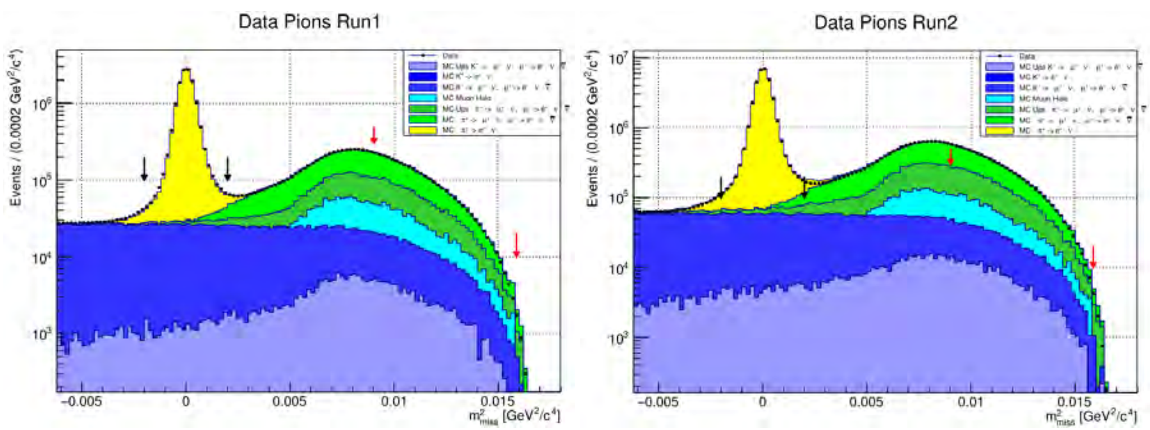

Figure 3.12: $m^2_{miss}(\pi)$ spectrum for Run1 (left), and Run2 (right). The black arrows define the normalization region and the red arrows define the HNL search region.

To account for uncertainties from the input parameters, the error on N_{π^+} is derived using the standard propagation formula for independent variables (Eq. 3.2), analogous to the kaon case. The total uncertainty is expressed as:

$$\Delta N_{\pi^+} = \frac{N_{\pi^+}}{N_{SM}^{\pi^+} - N_{\text{Bckg}}^{K}} \sqrt{N_{SM}^{\pi^+} + \sum_i (\Delta N_i^K)^2 + (N_{\pi^+})^2 \sum_j \left((A_j^\pi)^2 (\Delta \mathcal{B}_j^\pi)^2 + (\mathcal{B}_j^\pi)^2 (\Delta A_j^\pi)^2 \right)} \quad (3.6)$$

where $\Delta N_{SR}^{\pi^+} = \sqrt{N_{SR}^{\pi^+}}$ is the statistical error on the raw yield. The uncertainty on the kaon leakage term, ΔN_i^K, propagates the errors from the kaon flux, branching ratios, and cross-acceptance:

$$\Delta N_i^K = N_i^K \sqrt{\left(\frac{\Delta N_{K^+}}{N_{K^+}}\right)^2 + \left(\frac{\Delta \mathcal{B}_i^K}{\mathcal{B}_i^K}\right)^2 + \left(\frac{\Delta A_i^{K\to\pi}}{A_i^{K\to\pi}}\right)^2} \quad (3.7)$$

Here, ΔN_{K^+} is the total uncertainty on the kaon flux determined in Sec. 3.4.1.

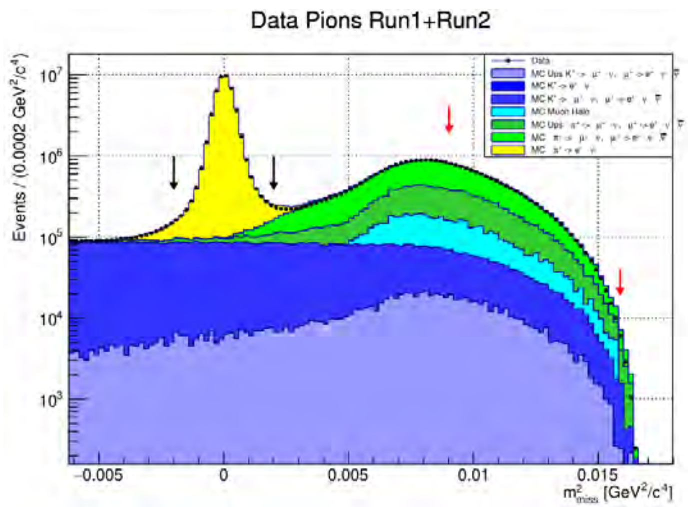

Figure 3.13: $m^2_{miss}(\pi)$ spectrum for data and MC for the full data set Run1+Run2. The black arrows define the normalization region used for Run1 and Run2, and the red arrows define the HNL search region.

3.4.2.1 Resolution Study and Normalization Window

The validity of the pion normalization window is confirmed by comparing the squared missing mass resolution between Data and Monte Carlo (MC) simulations. Table 3.7 summarizes the π_{e2} peak resolutions for Run1 and Run2. The MC resolution agrees with data within 0.96% for Run1 and 0.22% for Run2, demonstrating excellent modeling of the detector response.

Table 3.7: Squared missing mass resolutions (GeV^2/c^4) on the π_{e2} peak for data and MC.

Sample	**Run1**	**Run2**
Data	3.322×10^{-4}	3.321×10^{-4}
MC	3.290×10^{-4}	3.329×10^{-4}

Based on these resolutions ($\sigma \approx 3.32 \times 10^{-4}\,\mathrm{GeV}^2/c^4$), the normalization window is defined as $|m^2_{\mathrm{miss}}(\pi)| < 0.002\,\mathrm{GeV}^2/c^4$, corresponding to approximately 6σ. This window maximizes the statistical precision of the flux measurement while minimizing contamination from radiative tails and other background sources.

3.4.2.2 Systematic Uncertainty Evaluation

The systematic uncertainty on N_{π^+} is evaluated by testing the stability of the result with respect to variations in the half-width of the $m^2_{\mathrm{miss}}(\pi)$ normalization window. The window size is varied from $0.0004\,\mathrm{GeV}^2/c^4$ ($\sim 1.2\sigma$) to $0.0035\,\mathrm{GeV}^2/c^4$ ($\sim 10\sigma$), where the upper limit corresponds to the onset of significant Muon Halo contamination (see Fig. 3.14).

The systematic error is defined as the maximum variation of N_{π^+} observed in the range $0.001\,\mathrm{GeV}^2/c^4$ ($\sim 3\sigma$) to $0.0035\,\mathrm{GeV}^2/c^4$ with respect to the nominal window ($0.002\,\mathrm{GeV}^2/c^4$). The resulting statistical and systematic uncertainties are summarized in Table 3.8. Notably, the systematic uncertainty dominates the total error budget, reflecting the sensitivity of the pion yield to the background subtraction and window definition.

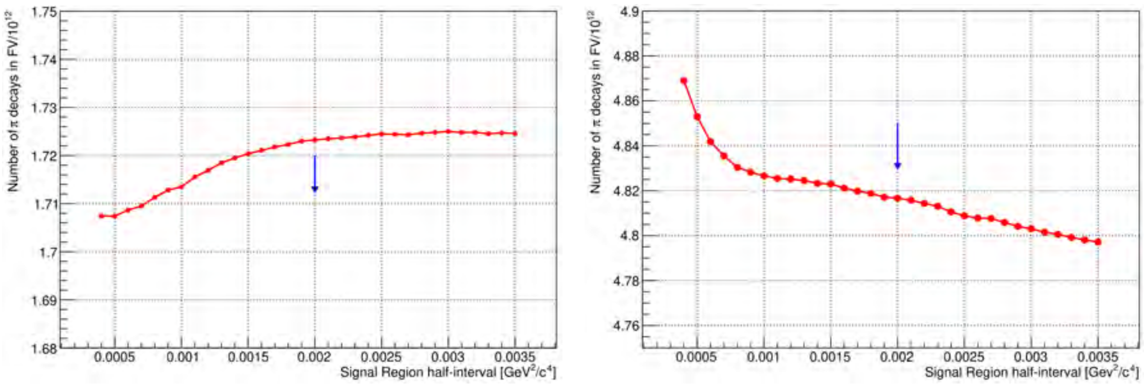

Figure 3.14: Stability checks of N_{π^+} with respect to the variation of the signal region for Run1 (left) and Run2 (right)

Table 3.8: Summary of statistical and systematic uncertainties on N_{π^+} for all data sets.

Error Type	**Run1**	**Run2**	**Run1+Run2**
Statistical	0.13%	0.14%	0.11%
Systematic	0.72%	0.61%	0.49%
Total	0.73%	0.62%	0.50%

Final Pion Flux

The determined number of pion decays in the fiducial volume for each period is:

$$N_{\pi^+}^{\mathrm{Run1}} = (1.72 \pm 0.01) \times 10^{12},$$

$$N_{\pi^+}^{\mathrm{Run2}} = (4.81 \pm 0.03) \times 10^{12}.$$

Combining both periods, the total pion flux used for the HNL search normalization is:

$$N_{\pi^+}^{\mathrm{Total}} = (6.54 \pm 0.03) \times 10^{12},$$

where the quoted uncertainty includes both statistical and systematic components.

3.4.3 Summary of Flux Determination

The computation of the pion and kaon fluxes, N_{π^+} and N_{K^+}, was derived directly from the normalization channels. The agreement between the observed data and the Monte Carlo simulation in these regions validates the accuracy of the flux values. With N_{π^+} and N_{K^+} now determined and fixed, they serve as normalization constants for the background in the signal region, with N_{π^+} specifically providing the reference scale against which the HNL signal is measured.

3.5 Background Estimation and Process Characterization

As shown in the stacked missing mass spectra in Section 3.4, the total background is built from separate pieces, each scaled according to its parent particle type. These simulations are linked directly to the measured pion (N_π) and kaon (N_K) fluxes.

Mathematically, the total expected background in any given bin, $N_{\text{bkg}}(m^2_{\text{miss}})$, is simply the sum of all the individual simulated components, each multiplied by its own scaling factor Z_i:

$$N_{\text{bkg}}(m^2_{\text{miss}}) = \sum_i Z_i \cdot N_{i\text{-MC}}(m^2_{\text{miss}})$$

In this equation, $N_{i\text{-MC}}$ is the raw number of simulated events for process i that passed the selection, and Z_i is the specific factor that scales this count to match the physical flux in the data. This ensures that the height of every colored background is fixed by real measurements. The following sections explain the physics behind these components, starting with the active processes that create a smooth background in the search region (95 MeV/c^2 $< m_N <$ 126 MeV/c^2), followed by those that are kinematically pushed out of this region.

3.5.1 Active Continuous Backgrounds in the Search Region

3.5.1.1 Pion Muon-Decay-in-Flight Chain ($\pi^+ \to \mu^+\nu_\mu \to e^+\nu_e\bar{\nu}_\mu\nu_\mu$)

The Physics Mechanism This process is the primary background in the search region. It begins with the dominant pion decay into a muon and a neutrino ($\pi_{\mu 2}$), occurring in either the Upstream or Standard regions. The resulting muon then decays in flight within the fiducial volume, producing a positron and two additional neutrinos. The final state consists of a single positron and three invisible neutrinos.

The Selection Survival These events pass the selection criteria naturally. The RICH detector identifies the daughter particle as a positron, and the Liquid Krypton (LKr) calorimeter confirms its identity by measuring an energy-to-momentum ratio (E/p) consistent with unity. Since the three accompanying neutrinos escape detection, the event mimics the signal topology of a single track with missing momentum.

Distribution Shape While a two-body decay produces a sharp peak at a specific missing mass value, this background involves a three-body muon decay. The three escaping neutrinos share the available energy in a continuous range of ways. When analyzed under the assumption of a two-body decay, this variable energy sharing smears the missing mass calculation. The result is a broad, smooth continuum rather than a peak, creating a flat background floor across the signal region.

Normalization The Upstream and Standard Monte Carlo samples are normalized independently using the scaling factor Z:

$$Z = \frac{N_\pi}{N_{\text{gen}}/\mathcal{B}(\pi_{\mu 2})}$$

This factor scales each sample to the physical kaon flux determined in the previous section.

3.5.1.2 Kaon Muon-Decay-in-Flight Chain ($K^+ \to \mu^+\nu_\mu \to e^+\nu_e\bar{\nu}_\mu\nu_\mu$)

The Physics Mechanism This process represents a background arising from the dominant kaon decay mode ($K_{\mu 2}$). The parent kaon decays into a muon and a neutrino in either the Upstream or Standard regions. The secondary muon subsequently decays in flight within the fiducial volume, yielding a positron and two additional neutrinos. The final state comprises a single positron and three invisible neutrinos.

The Selection Survival Since the beam is unseparated and the upstream tagging system does not identify pions, the specific identity of the parent particle remains unknown for any given event. Consequently, these events are reconstructed under the pion hypothesis alongside the signal candidates. The daughter positron is correctly identified by the RICH detector and satisfies the $E/p \approx 1$ requirement in the Liquid Krypton (LKr) calorimeter, allowing the event to pass the single-track selection criteria.

Distribution Shape When reconstructed assuming a pion parent, the kinematic mismatch between the true kaon mass and the assumed pion mass, combined with the continuous energy sharing among the three final-state neutrinos, results in a broad, smooth continuum distribution. This shape contributes to the background across the missing mass spectrum.

Normalization The Upstream and Standard Monte Carlo samples are normalized independently using the scaling factor Z:

$$Z = \frac{N_K}{N_{\text{gen}}/\mathcal{B}(K_{\mu 2})}$$

This factor scales each sample to the physical kaon flux determined in the previous section.

3.5.1.3 Muon Halo

The Physics Mechanism This background originates from high-energy muons produced at the target or via early interactions upstream of the KTAG detector. These muons travel nearly parallel to the beam axis, forming a halo around the core beam. Within the fiducial volume, a fraction of these halo muons undergo decay in flight ($\mu^+ \to e^+\nu_e\bar{\nu}_\mu$), producing a positron and two neutrinos.

The Selection Survival The decay positron is a genuine lepton that satisfies the RICH identification criteria and the $E/p \approx 1$ requirement in the Liquid Krypton (LKr) calorimeter. Since the accompanying neutrinos escape detection, the event topology matches the signal signature of a single positron track with missing four-momentum. The lack of an incoming beam track match for halo muons is mitigated by the high beam intensity and accidental overlaps, allowing some events to pass the temporal and spatial matching requirements.

Distribution Shape The decay in flight of the halo muon is a three-body process. The continuous energy sharing among the positron and the two neutrinos results in a broad, smooth continuum distribution in the squared missing mass spectrum. When reconstructed under the pion hypothesis, this background contributes a non-resonant floor across the signal region, indistinguishable in shape from the in-flight muon decays originating from pion or kaon parents.

Normalization An empirical scaling factor was applied to the Muon Halo Monte Carlo sample to ensure an accurate description of the data. Given the difficulty in computing the absolute halo flux from first principles, this a priori factor was determined by subtracting the contributions of all other known backgrounds from the data, attributing the remaining residual to the halo process. This approach ensures the Monte Carlo correctly reproduces the shape and magnitude of the observed background. It is important to note that while this normalization establishes a consistent background model, the final signal extraction employs a purely data-driven peak search scan, which identifies local excesses over this smooth distribution independent of the absolute scale of the halo component.

3.5.2 Kinematically Suppressed Components

3.5.2.1 The Decay $\pi^+ \to e^+\nu_e$ (π_{e2})

The Physics Mechanism This process is the Standard Model two-body leptonic decay of the pion, producing a single positron and an electron neutrino.

The Selection Survival This channel presents the exact experimental signature as the HNL signal: a single reconstructed positron track originating from the decay volume with no additional activity. The positron is identified by the RICH detector and satisfies the $E/p \approx 1$ requirement in the Liquid Krypton (LKr) calorimeter.

Distribution Shape As a two-body decay, the kinematics are fully constrained, resulting in a sharp, narrow peak centered at $m^2_{\text{miss}} = 0$. This distribution is kinematically suppressed in the signal region ($m^2_{\text{miss}} > 0$), contributing only through the detector resolution tails that smear the peak slightly into the positive mass range.

Normalization The sample is normalized using the scaling factor $Z_{\pi e2}$:

$$Z_{\pi e2} = \frac{N_\pi}{N_{\text{gen}}/\mathcal{B}(\pi_{e2})}$$

This factor scales the Monte Carlo prediction to the physical pion flux determined previously.

3.5.2.2 The Decay $K^+ \to e^+\nu_e$ (K_{e2})

The Physics Mechanism This process is the Standard Model leptonic decay of the kaon, producing a single positron and an electron neutrino.

The Selection Survival Since the beam is unseparated and the upstream tagging system does not identify pions, the specific identity of the parent particle remains unknown. The positron is identified by the RICH detector and satisfies the $E/p \approx 1$ requirement in the Liquid Krypton (LKr) calorimeter, presenting the exact experimental signature as the HNL signal.

Distribution Shape When reconstructed with the pion mass hypothesis, the significant mass difference between the true kaon parent and the assumed pion parent shifts the entire kinematic distribution into the negative squared missing mass region ($m^2_{\text{miss}} < 0$). This background manifests as a broad distribution grouped at negative values, distinct from any resonant peak. It is the most displaced background component in the negative spectrum and is kinematically suppressed in the signal region ($m^2_{\text{miss}} > 0$), contributing only via the far tails of the detector resolution.

Normalization The sample is normalized using the scaling factor Z_{Ke2}:

$$Z_{Ke2} = \frac{N_K}{N_{\text{gen}}/\mathcal{B}(K_{e2})}$$

This factor scales the Monte Carlo prediction to the physical kaon flux determined previously.

3.5.3 Background Summary

The estimated background contributions for the full data set, including their statistical uncertainties, are summarized in Table 3.9 and visualized in the stacked missing mass spectrum of Fig. 3.13. The data covers the region $m^2_{\text{miss}} > 0.0036\ \text{GeV}^2/c^4$, which encompasses the HNL search window.

The muon-decay-in-flight chains dominate the background, accounting for approximately 90% of the total expected events. Specifically, the pion decay chains ($\pi_{\mu 2}$) contribute the largest share ($\approx$ 21.6 million events), followed by the kaon decay chains ($K_{\mu 2}$) with $\approx$ 3.5 million events. The Muon Halo component provides a significant contribution ($\approx$ 2.5 million events) but carries a larger relative statistical uncertainty due to the empirical normalization method. In contrast, the kinematically suppressed π_{e2} decay contributes negligibly ($\approx 39,000$ events, $< 0.2\%$), confined to the low-mass tail.

The total expected background yield is 27.60$\pm$0.19 million events, which is in excellent agreement with the 27.47 million events observed in the data. This close correspondence, with a difference of less than 0.5%, confirms that the summed Monte Carlo model accurately reproduces the total rate and composition of the smooth continuum, validating its use as a stable baseline for the local peak search.

Table 3.9: Backgrounds with MC simulations in the region $m^2_{miss} > 0.0036\,\text{GeV}^2/c^4$.

Decay mode	Branching fraction	Background in the region $m^2_{miss} > 0.036\,\text{GeV}^2/c^4$
$\pi^+ \to e^+\nu$	$(1.230 \pm 0.004) \times 10^{-4}$	39 302 $\pm$ 1 516
$\pi^+ \to \mu^+\nu$, $\mu^+ \to e^+\nu\bar{\nu}$	$(99.98770 \pm 0.00004) \times 10^{-2}$	13 635 430 $\pm$ 123 025
$\pi^+ \to \mu^+\nu$ upstream, $\mu^+ \to e^+\nu\bar{\nu}$	$(99.98770 \pm 0.00004) \times 10^{-2}$	7 931 500 $\pm$ 79 311
$K^+ \to \mu^+\nu$, $\mu^+ \to e^+\nu\bar{\nu}$	$(63.56 \pm 0.11) \times 10^{-2}$	2 428 598 $\pm$ 13 250
$K^+ \to \mu^+\nu$ upstream, $\mu^+ \to e^+\nu\bar{\nu}$	$(63.56 \pm 0.11) \times 10^{-2}$	1 026 784 $\pm$ 9 602
Muon Halo, $\mu^+ \to e^+\nu\bar{\nu}$		2 542 845 $\pm$ 1 917 219
Total		27 604 470 $\pm$ 1 922 867
Data events		27 471 715

3.6 Mass Scan Strategy

Having established that the cumulative Standard Model background forms a smooth, non-resonant continuum across the missing mass spectrum, the analysis shifts from global process validation to a localized search for narrow resonant anomalies. Since the heavy neutral lepton mass m_N is a free parameter, a statistical peak search is performed across a fine set of mass hypotheses. This section details the experimental resolutions and acceptances that define the search range, and describes the data-driven sideband fitting procedure used to estimate local background expectations directly from the data.

3.6.1 HNL Acceptances and Resolution

The search strategy relies on two mass-dependent parameters: the signal acceptance $A(m_N)$ and the missing mass resolution $\sigma_{m^2}(m_N)$. These define the signal window geometry and the expected yield across the scanned mass range.

Squared Missing Mass Resolution $\sigma_{m^2}(m_N)$: The resolution determines the statistical width of a potential HNL signal peak and is quantified as the standard deviation (σ_{m^2}) of the reconstructed m^2_{miss} distribution. Figure 3.15 illustrates this for a single benchmark mass hypothesis in Run 1 and Run 2 conditions. The distribution is constructed from the difference between the true and reconstructed missing mass squared, which centers the spectrum at zero. A Gaussian function is fitted to this central region to extract σ_{m^2}. The resulting resolution values, summarized in Figure 3.16, define the local squared resolution width for each point in the mass scan.

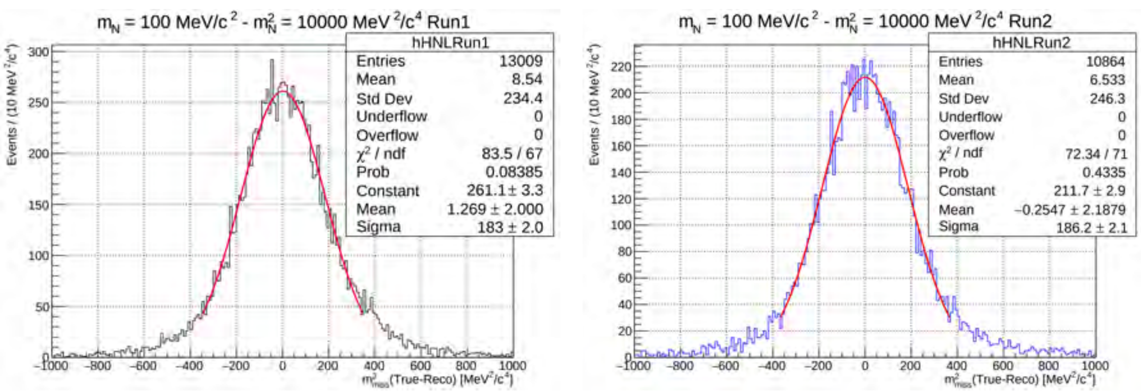

Figure 3.15: Distribution of the squared missing mass, m^2_{miss}(True-Reco), for the HNL mass hypothesis $m_N = 100$ MeV/c. The spectra are shown for both Run 1 (left) and Run 2 (right) conditions, centered at zero. The solid lines represent Gaussian fits to the central region, from which the resolution σ_{m^2} is extracted.

Signal Selection Acceptance $A(m_N)$: The acceptance represents the efficiency of detecting an HNL signal, accounting for both the detector geometry and the specific selection criteria. Because the detector conditions and beam intensity differed between the 2017–2018 (Run 1) and 2021–2024 (Run 2) data-taking periods, separate acceptance curves $A_{\text{Run1}}(m_N)$ and $A_{\text{Run2}}(m_N)$ were

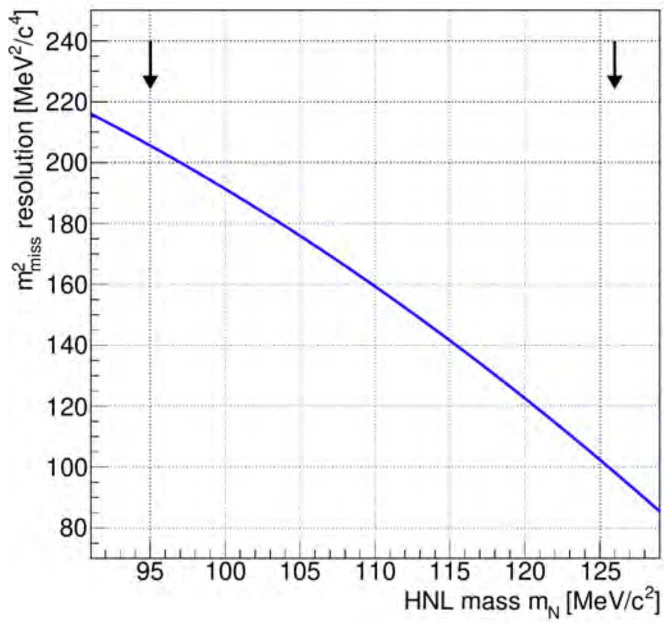

Figure 3.16: Squared missing mass resolution σ_{m^2} evaluated from simulations as functions of HNL mass. The boundaries of the HNL search region are indicated by the vertical arrows.

computed using simulated samples matched to each specific period conditions. The final acceptance curve used for the combined analysis is a weighted average of these two components, where the weights correspond to the fractional contribution of each run to the total pion flux N_{π^+}:

$$A_{\text{combined}}(m_N) = f_{\text{Run1}} \cdot A_{\text{Run1}}(m_N) + f_{\text{Run2}} \cdot A_{\text{Run2}}(m_N) \tag{3.8}$$

where f_{Run1} and f_{Run2} are the proportions of N_{π^+} contributed by Run 1 and Run 2, respectively.

To calculate the acceptance for each period, simulated events must satisfy the signal window criterion $|m^2_{\text{miss}}(\text{true}) - m^2_{\text{miss}}(\text{Reco})| < 1.5\sigma_{m^2_N}$. The acceptance is then defined as:

$$A_N = \frac{N(\text{Sel, in Signal Window})}{N(\text{FV})} \tag{3.9}$$

where $N(\text{Sel, in Signal Window})$ is the number of events passing the full selection and signal window cut, and $N(\text{FV})$ is the total number of simulated events generated within the fiducial volume ($105\,\text{m} < z_{\text{true}} < 180\,\text{m}$). Figure 3.17 displays the individual acceptance curves for Run 1 and Run 2, alongside the final weighted acceptance curve for the full dataset across the search range $95 \leq m_N \leq 126$ MeV/c^2.

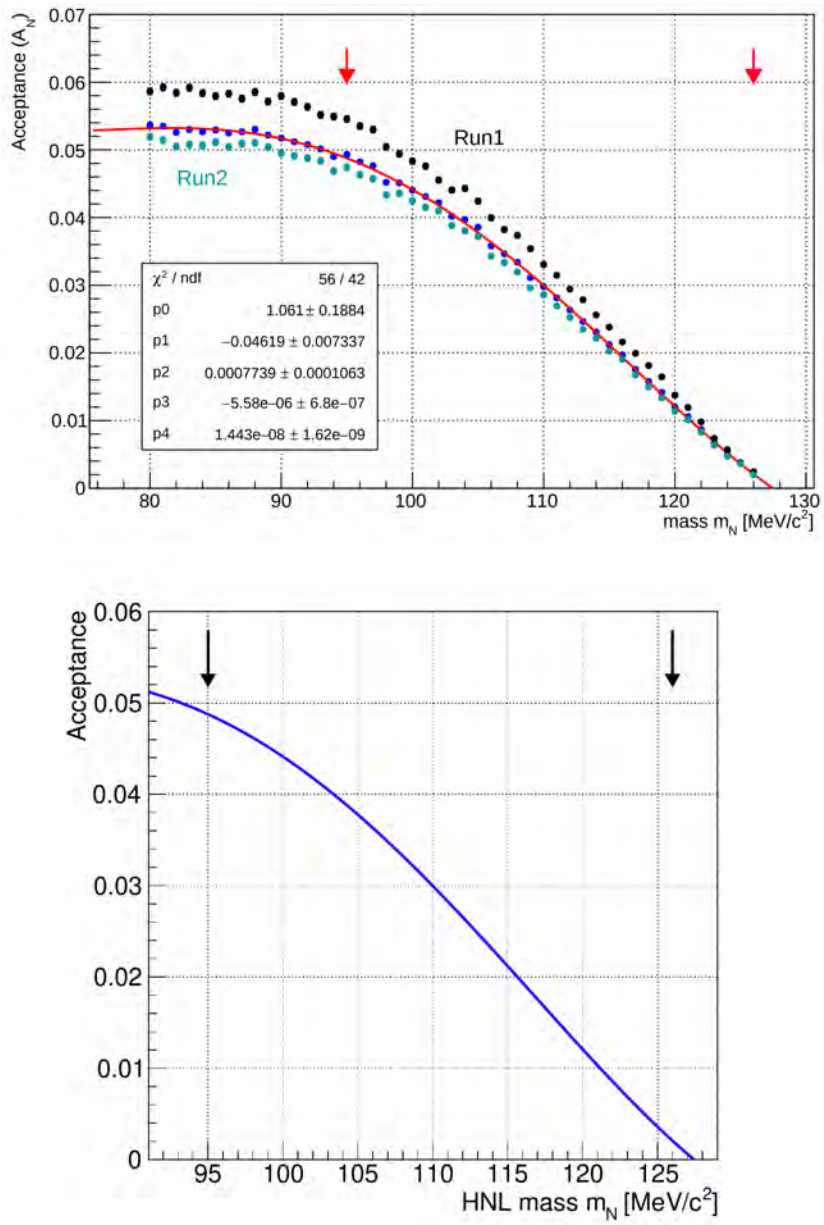

Figure 3.17: Acceptance A_N derived from simulations as a function of HNL mass. Individual components for Run 1 and Run 2, their combination, and a polynomial fit to the combined data (Top). Final parameterized acceptance combining Run 1 and Run 2 (Bottom). Vertical arrows indicate the boundaries of the HNL search region.

3.6.2 Signal Window

The HNL search is conducted across a series of discrete mass hypotheses, m_N, spanning the designated search interval. To ensure complete coverage of the phase space without loss of sensitivity, the step size between adjacent mass hypotheses is strictly coupled to the tracking resolution:

$$\Delta m_N = 0.75\,\sigma_{m^2}(m_N)$$

Applying this step constraint across the 95 MeV/c^2 $\leq m_N \leq$ 126 MeV/c^2 search region yields a total of 63 discrete mass hypotheses. For each tested mass hypothesis m_N, a localized Signal Window is centered symmetrically around the nominal expectation value. The signal window region is defined as:

$$\left|m^2_{\text{miss}}(\pi) - m^2_N\right| < 1.5\sigma_{m^2}(m_N)$$

Events falling within this designated signal window are counted as observed candidates (N_{obs}) for that specific hypothesis.

3.6.3 Data-Driven Background Estimation Using Sidebands

To estimate the expected background (N_{exp}) inside the signal window without relying on simulation shapes, a data-driven interpolation method is used. For each of the 63 mass points tested, the data immediately surrounding the signal region is analyzed to predict what lies inside.

First, a masked region is defined that exactly matches the signal window ($\pm 1.5\sigma_{m^2}$). The data inside this zone is hidden. To facilitate this procedure, a dedicated histogram is constructed for every mass hypothesis with a bin width equal to $0.75\sigma_{m^2}(m_N)$. This specific binning ensures that the signal window covers exactly four central bins, with the tested mass hypothesis m_N falling precisely at the center boundary between the two innermost bins.

The sidebands are then defined as the data regions immediately outside this masked zone, extending from $1.5\sigma_{m^2}$ out to $9.0\sigma_{m^2}$

$$1.5\sigma_{m^2}(m_N) < \left|m^2_{\text{miss}}(\pi) - m^2_N\right| < 9.0\sigma_{m^2}(m_N)$$

These sideband boundaries are adjusted if necessary to stay within the global valid mass range of the detector. The estimation proceeds in three steps:

- **Fit:** The visible data points in the sidebands are fitted with a third-order polynomial function:

  $$f(m^2_{\text{miss}}) = a_0 + a_1 m^2_{\text{miss}} + a_2 (m^2_{\text{miss}})^2 + a_3 (m^2_{\text{miss}})^3$$

  Figure 3.18 illustrates this process for a representative mass hypothesis, showing the data points in the sideband regions, the blinded signal window in the center, and the resulting third-order polynomial fit curve.

- **Interpolation:** This fitted curve is mathematically extended across the blinded signal region.

- **Integration:** The area under this interpolated curve within the $\pm 1.5\sigma_{m^2}$ window gives the expected background count:

  $$N_{\text{exp}} = \int_{m^2_N - 1.5\sigma_{m^2}}^{m^2_N + 1.5\sigma_{m^2}} f(m^2_{\text{miss}})\, d(m^2_{\text{miss}})$$

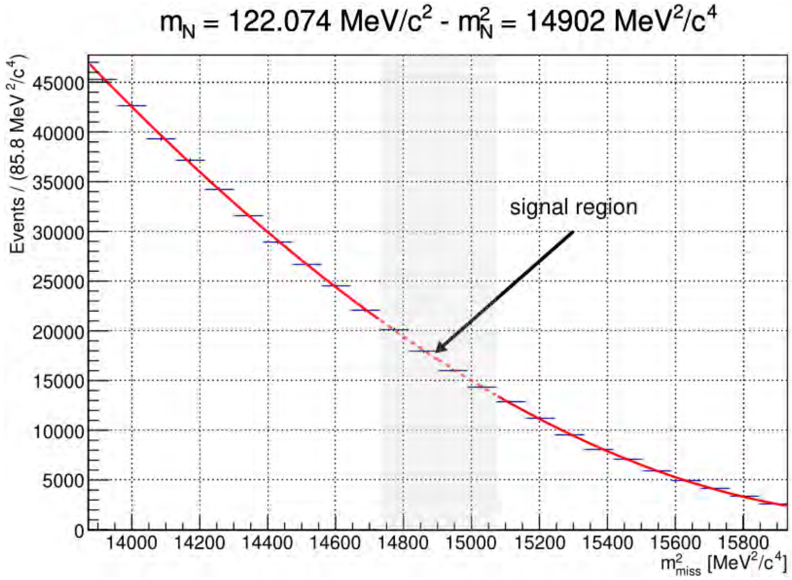

Figure 3.18: Example of the data-driven background estimation for a mass hypothesis. The central region (shaded) corresponds to the blinded signal window ($\pm 1.5\sigma_{m^2}$), which spans exactly four histogram bins with width $0.75\sigma_{m^2}$. The points in the sideband regions ($1.5\sigma_{m^2}$ to $9.0\sigma_{m^2}$) are fitted with a third-order polynomial (red line), which is then integrated over the masked region to estimate the expected background.

The total uncertainty on the background estimate, δN_{exp}, combines both statistical and systematic components added in quadrature. The statistical part comes directly from the errors on the fit parameters. The systematic part accounts for the choice of fit model and sideband definition; it is evaluated as the difference between the nominal N_{exp} (from the third-order fit) and values obtained from:

- Fits using a fourth-order polynomial.
- Fits using slightly wider sidebands (extending to $8.25\sigma_{m^2}$ and $9.75\sigma_{m^2}$).

Figure 3.19 displays the total δN_{exp} for each mass hypothesis, broken down by these contributions. This rigorous process provides the triad of values (N_{obs}, N_{exp}, δN_{exp}) required for the final statistical limit calculation.

3.7 Statistical Inference Method

The sideband fitting method in Section 3.6 calculates the observed events (N_{obs}), expected background (N_{exp}), and systematic uncertainties (δN_{exp}) for each of the 63 tested mass hypotheses. To interpret these values as a statement on the existence of a heavy neutral lepton, a formal statistical framework is applied, as detailed below.

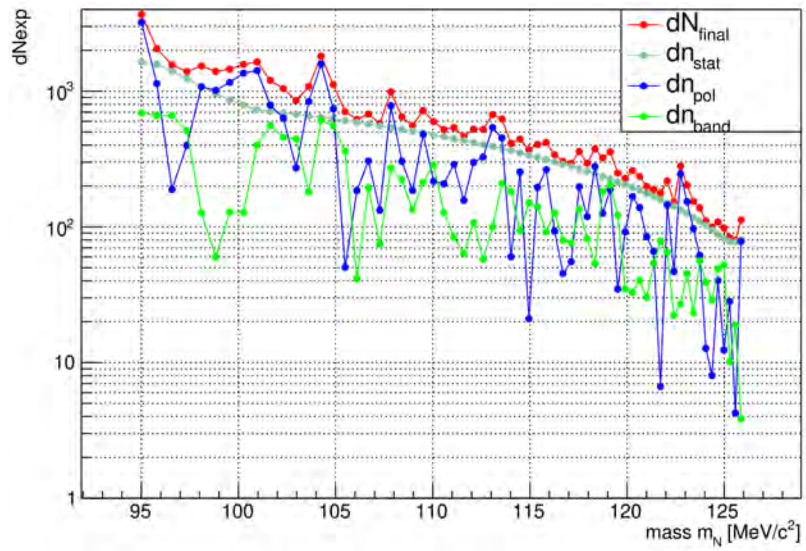

Figure 3.19: Breakdown of contributions to the total expected uncertainty δN_{exp}. Each component represents a distinct source of systematic or statistical error.

3.7.1 The CLs Method

The CL_s method, originally developed for the LEP Higgs searches in the late 1990s [40][41][42], defines the confidence level for signal exclusion as the ratio of two cumulative probabilities: $CL_s \equiv CL_{s+b}/CL_b$. In this framework, $CL_{s+b} = P(X \leq X_{obs} \mid s + b)$ represents the confidence in the signal-plus-background hypothesis (the probability that the test statistic X is less signal-like than observed), while $CL_b = P(X \leq X_{obs} \mid b)$ represents the confidence in the background-only hypothesis. By dividing the former by the latter, the method effectively re-normalizes the exclusion probability to penalize regions where the experiment lacks sensitivity (i.e., where the background and signal distributions overlap significantly), thereby preventing the exclusion of signal hypotheses that the experiment is incapable of distinguishing from background fluctuations. This approach ensures conservative coverage and avoids the counter-intuitive result of excluding a signal simply because the observed data fluctuated below the expected background.

The computation determines the compatibility of a hypothesized signal strength μ with a fixed observed event count N_{obs}. It proceeds by deriving a test statistic from the fundamental Poisson likelihoods and evaluating its cumulative probability under two competing hypotheses.

3.7.1.1 Derivation of the Test Statistic

The foundation of the method is the likelihood ratio, which compares the probability of the data under the signal-plus-background hypothesis (H_{s+b}) to the background-only hypothesis (H_b).

For a single channel counting experiment, the probability of observing N events follows a Poisson distribution. The likelihood functions for the two hypotheses are:

Signal+Background (H_{s+b}): The expected rate is $\mu + b$.

$$\mathcal{L}_{s+b} = P(N \mid \mu + b) = \frac{(\mu + b)^N e^{-(\mu+b)}}{N!}$$

Background-only (H_b): The expected rate is b.

$$\mathcal{L}_b = P(N \mid b) = \frac{b^N e^{-b}}{N!}$$

The likelihood ratio λ is the quotient of these two probabilities:

$$\lambda = \frac{\mathcal{L}_{s+b}}{\mathcal{L}_b} = \left(\frac{\mu + b}{b}\right)^N e^{-\mu}$$

To construct the test statistic X, we take the natural logarithm of the likelihood ratio. The log-likelihood ratio is:

$$X = \ln(\lambda) = N \cdot \ln\left(\frac{\mu + b}{b}\right) - \mu$$

This value X quantifies how "signal-like" the observation is.

3.7.1.2 The Ideal Case: Fixed Background

If the background b is known perfectly ($\sigma_b = 0$), the confidence levels are defined as the cumulative probabilities that the test statistic for a random experiment is less than or equal to the observed test statistic X_{obs}.

Let X_{s+b} be the test statistic computed for a random count drawn from the signal-plus-background hypothesis, and X_b be the statistic for a count drawn from the background-only hypothesis. The confidence levels are:

CL_{s+b}: The probability that the signal-plus-background hypothesis yields a test statistic less than or equal to the observed one.

$$CL_{s+b} = P\left(X_{s+b} \leq X_{obs} \mid \mu + b\right)$$

CL_b: The probability that the background-only hypothesis yields a test statistic less than or equal to the observed one.

$$CL_b = P\left(X_b \leq X_{obs} \mid b\right)$$

The Upper Limit is the value μ_{up} on the signal at a confidence level of α:

$$CL_s(\mu_{up}) = \frac{P\left(X_{s+b} \leq X_{obs} \mid \mu_{up} + b\right)}{P\left(X_b \leq X_{obs} \mid b\right)} = 1 - \alpha$$

3.7.1.3 The Realistic Case: Convolution of Likelihoods

In practice, the background expectation b has an uncertainty σ_b, modeled as a Gaussian distribution. The true background b' is a nuisance parameter. To account for this, we define the likelihoods as the convolution of the Poisson probability for the fixed observed count N and the Gaussian uncertainty density.

The likelihood for the Signal+Background hypothesis is:

$$\mathcal{L}_{s+b}(N \mid \mu, b) = \int_0^{\infty} \left[\frac{(\mu + b')^N e^{-(\mu+b')}}{N!} \right] \cdot \frac{1}{\sqrt{2\pi}\sigma_b} e^{-\frac{(b'-b)^2}{2\sigma_b^2}} \, db'$$

The likelihood for the Background-only hypothesis is:

$$\mathcal{L}_b(N \mid b) = \int_0^{\infty} \left[\frac{(b')^N e^{-b'}}{N!} \right] \cdot \frac{1}{\sqrt{2\pi}\sigma_b} e^{-\frac{(b'-b)^2}{2\sigma_b^2}} \, db'$$

In these equations, the term in the square brackets is the standard Poisson likelihood for observing N events given a specific true background b', and the second term is the Gaussian probability density for that b'. The integral averages the likelihood over all possible true background values, creating a single likelihood function that fully incorporates the systematic uncertainty. The Confidence Levels (CL_{s+b} and CL_b) used in the final ratio are then the cumulative probabilities derived from these convolved distributions.

3.7.1.4 Implementation Strategy: Monte Carlo Integration

The NA62 implementation [43] obtains the numerical upper limit by scanning the signal strength parameter over a range of values. For each tested value of μ, the convolution integrals are evaluated independently via a Monte Carlo integration procedure that replaces analytical calculation with massive statistical sampling. The scan iterates until the condition $CL_s(\mu) = 1 - \alpha$ (with $\alpha = 0.9$ for a 90% Confidence Level) is satisfied, defining the corresponding as the upper limit. The evaluation for each step proceeds as follows:

1. Sampling the Nuisance Parameter: Instead of integrating over the Gaussian curve analytically, the algorithm draws a large number of random background values b'_i from the Gaussian distribution $\mathcal{N}(b, \sigma_b)$. Each draw represents a possible "true" background consistent with the systematic uncertainty.

2. Generating Pseudo-Experiments: For each sampled b'_i, the algorithm generates random pseudo-experiments (toys) n_{toy} from the corresponding Poisson distributions $\mu + b'_i$ or b'_i.

3. Evaluating the Statistic: For each sampled background b'_i, the algorithm generates two distinct pseudo-experiment counts: n_{s+b} from the signal-plus-background Poisson distribution and n_b from the background-only Poisson distribution. It then calculates the test statistic for the signal hypothesis using n_{s+b} and for the background hypothesis using n_b, comparing each against the observed statistic X_{obs} derived from the fixed data count N_{obs} (all computed with the same b'_i).

4. Counting Successes: The algorithm performs two separate checks:

- It verifies if the signal-plus-background statistic satisfies $X_{s+b} \leq X_{obs}$, this is mathematically equivalent to checking $n_{s+b} \leq N_{obs}$. If true, the CL_{s+b} counter is incremented.
- It verifies if the background-only statistic satisfies $X_b \leq X_{obs}$. Similarly, this simplifies to checking $n_b \leq N_{obs}$. If true, the CL_b counter is incremented.

5. Forming the Ratio: After hundreds of thousands of iterations, the accumulated counts serve as estimators for the integrals. The final CL_s is simply the ratio of the total count of compatible signal+background scenarios to the total count of compatible background-only scenarios.

This approach naturally handles the convolution by averaging the results over many random fluctuations of the background, the final ratio converges to the value of the integral, providing a rigorous limit that fully incorporates uncertainties on the background.

3.7.2 Results of the CLs method

The final outcome of the CL_s procedure is summarized in Figure 3.20, which displays the 90% confidence level upper limits on the number of signal events, N_S^{UL}, as a function of the HNL mass m_N. The plot presents three key components derived from the analysis of the 63 mass hypotheses:

- Observed Limit: The solid red line represents the upper limit on the signal yield (N_S^{UL}) calculated using the actual observed event counts (N_{obs}) from the data.
- Expected Limit: The dashed black line indicates the median expected upper limit under the background-only hypothesis, representing the sensitivity of the experiment in the absence of a signal.
- Brazilian Bands: The green and yellow bands surrounding the expected limit correspond to the $\pm 1\sigma$ and $\pm 2\sigma$ quantiles of the expected limit distribution, respectively. These bands visualize the statistical fluctuations expected from the background model.

These limits on the signal event count represent the final statistical constraint derived from the data, serving as the foundation for the physical interpretation discussed in the following sections.

Here is the revised section explicitly showing the equality for the number of signal events (N_S) through both the Branching Ratio and the mixing parameter $|U_{e4}|^2$, along with the references to the SES plots.

3.8 Single Event Sensitivity

With the total number of pion decays N_{π^+} determined in Section 3.4 and the signal acceptance A_N derived from simulation, the experiment's sensitivity is quantified via the Single Event Sensitivity (SES). The SES is defined as the branching ratio (or mixing parameter value) corresponding to the expectation of exactly one signal event ($N_S = 1$).

The expected number of signal events, N_S, can be expressed equivalently through the branching ratio $\mathcal{B}(\pi^+ \to e^+ N)$ or the mixing parameter $|U_{e4}|^2$:

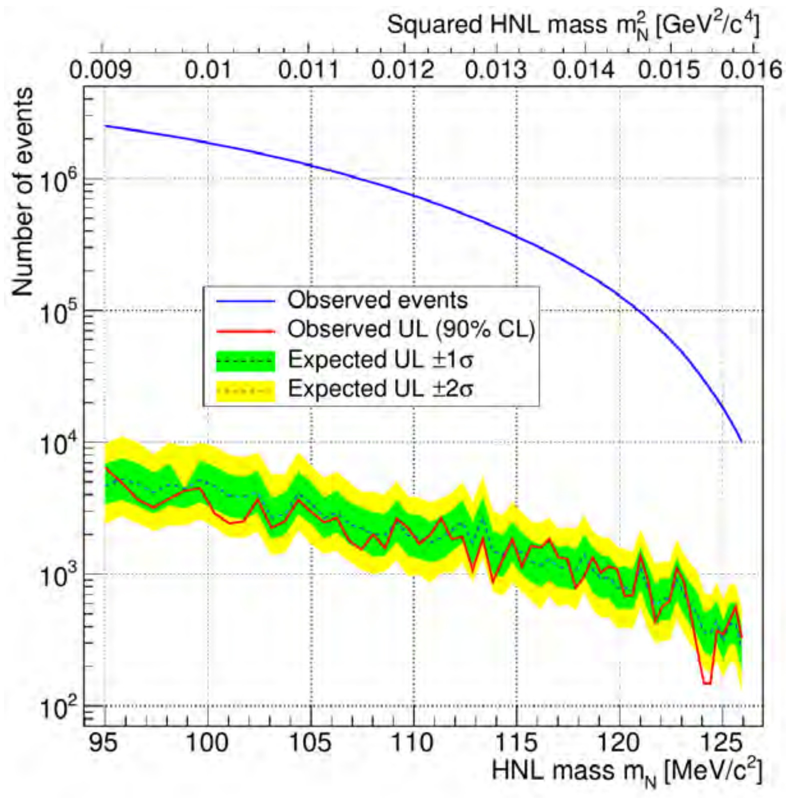

Figure 3.20: Number of observed events N_{obs}, observed and expected upper limits at 90 % CL for the number of $\pi^+ \to e^+ N$ events, and the expected $\pm 1\sigma$ and $\pm 2\sigma$ bands in the background-only hypothesis for each HNL mass value considered..

$$N_S = \frac{\mathcal{B}(\pi^+ \to e^+ N)}{\mathcal{B}_{\text{SES}}} = \frac{|U_{e4}|^2}{|U_{e4}|^2_{\text{SES}}} \tag{3.10}$$

where the SES terms are defined as:

$$\mathcal{B}_{\text{SES}} = \frac{1}{N_{\pi^+} \cdot A_N} \tag{3.11}$$

$$|U_{e4}|^2_{\text{SES}} = \frac{\mathcal{B}_{\text{SES}}}{\mathcal{B}(\pi^+ \to e^+ \nu_e) \cdot \rho(m_N)} \tag{3.12}$$

Here, $\mathcal{B}(\pi^+ \to e^+ \nu_e)$ is the Standard Model branching ratio and $\rho(m_N)$ is the kinematic phase space factor.

Figure 3.21 displays the $\mathcal{B}_{\text{SES}}$ and $|U_{e4}|^2_{\text{SES}}$ curves as a function of the HNL mass. These metrics provide the direct conversion factors: multiplying the upper limits on the signal count (N_S^{UL}) obtained in the statistical inference by these SES values yields the final exclusion limits on the branching ratio and mixing parameter.

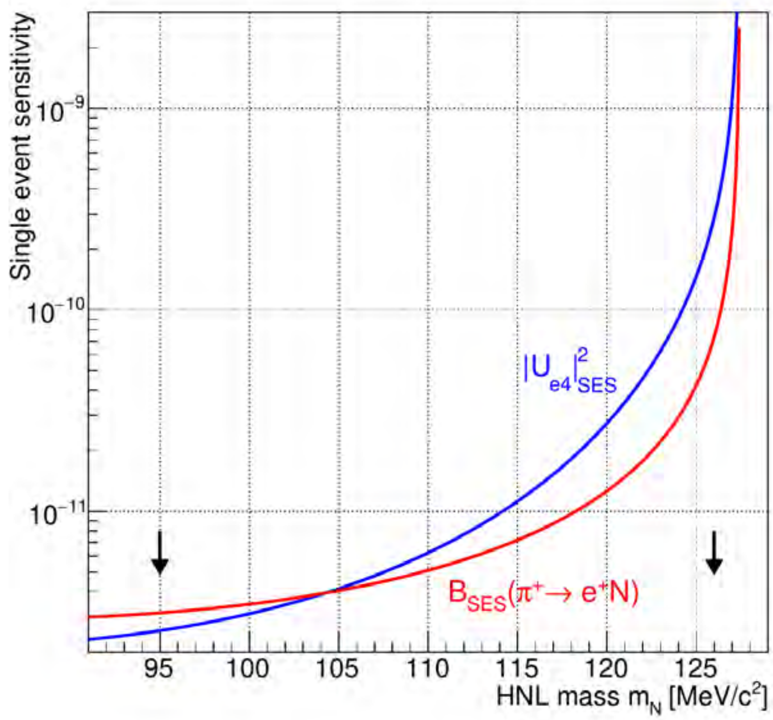

Figure 3.21: Single event sensitivities $\mathcal{B}_{\mathrm{SES}}(\pi^+ \to e^+N)$ and $|U_{e4}|^2_{\mathrm{SES}}$ as functions of the HNL mass.

3.9 Results and Discussion

A search for Heavy Neutral Lepton (HNL) production in $\pi^+ \to e^+N$ decays has been performed using the data collected by the NA62 experiment at CERN in 2017–2024. The statistical analysis of the 63 mass hypotheses reveals no significant excess of events in the data compared to the background-only expectation. Across the entire scanned mass range ($95 \leq m_N \leq 126$ MeV/c^2), the observed upper limits on the signal (N_S^{UL}) remain consistent with the expected limits, with no deviation exceeding 3σ. This agreement confirms that the data are compatible with the Standard Model background hypothesis.

Consequently, upper limits for the mixing parameter $|U_{e4}|^2$ have been established at the 10^{-8} level over the HNL mass range 95–126 MeV/c^2, assuming the mean lifetime exceeds 50 ns.

Figure 3.22 displays the upper limit at 90% CL on the branching ratio $\mathcal{B}(\pi^+ \to e^+N)$ as a function of the HNL mass m_N.

Figure 3.23 presents the Upper limit at 90% in the $|U_{e4}|^2$ versus HNL mass compared to previous results from the PIENU experiment [23]. The limits obtained in this study are comparable to, or more stringent than, those reported by PIENU for the overlapping mass region covered in this search.

Furthermore, Figure 3.24 compares the new pion decay limits with existing NA62 results from kaon decays ($K^+ \to e^+N$ and $K^+ \to \mu^+N$) [24][27]. While the kaon search covers higher masses, this analysis extends the NA62 limit reach to lower masses (95–126 MeV/c^2). An improved sensitivity of this analysis in terms of $|U_{e4}|^2$ and $|U_{\mu 4}|^2$ is expected including future NA62 data.

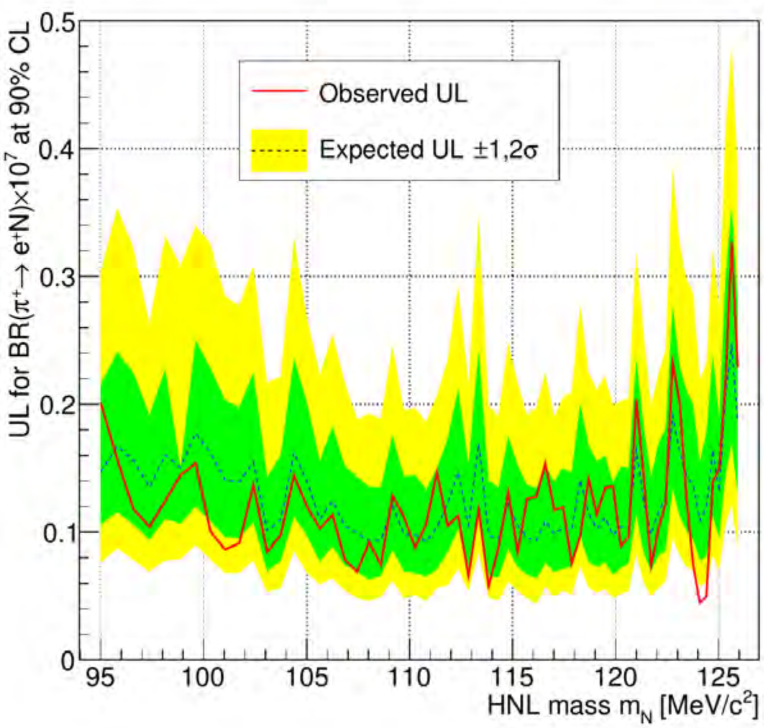

Figure 3.22: Observed and expected upper limits at 90% CL of PieN branching ratio in each HNL mass hypothesis.

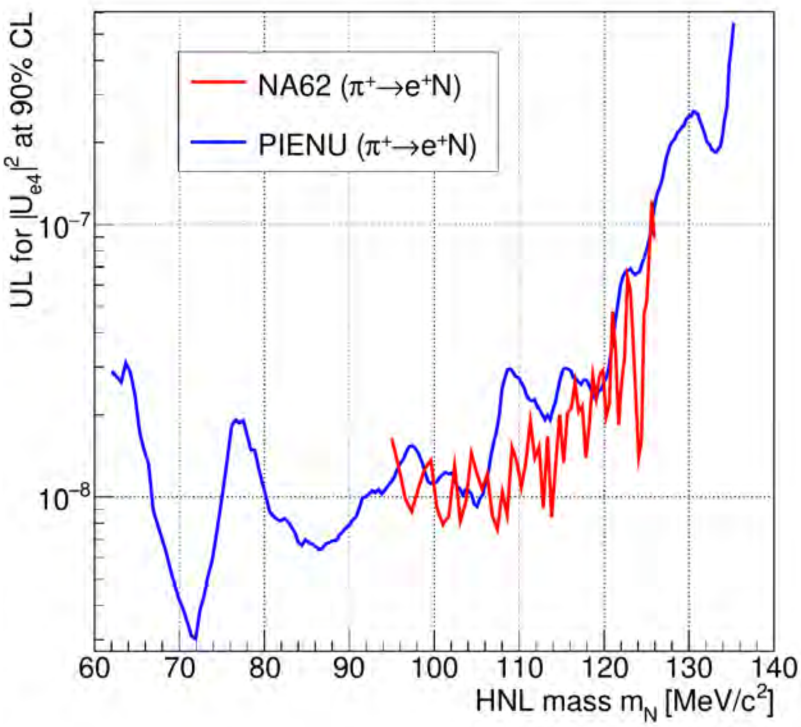

Figure 3.23: NA62 upper limits of the mixing parameter $|U_{e4}|^2$ at 90% CL compared with the PIENU upper limits.

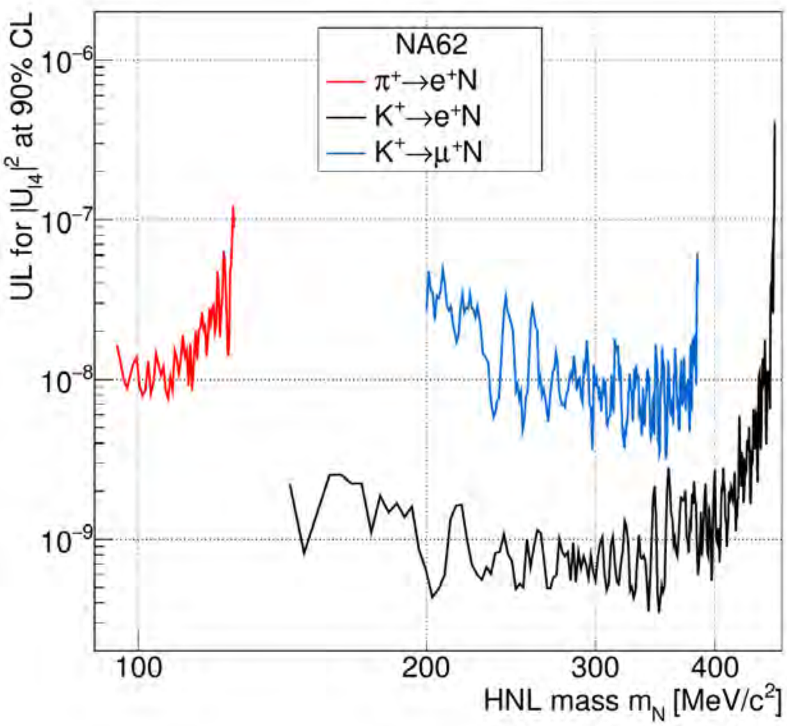

Figure 3.24: NA62 upper limits of the mixing parameter $|U_{i4}|^2$ (i=e, μ) at 90% CL. Upper limits shown in logarithmic scale on X and Y axes.

Chapter 4

Statistical Measurement of R_π

4.1 Analysis Strategy and Connection to HNL Search

This chapter presents a measurement of the ratio $R_\pi = \Gamma(\pi^+ \to e^+\nu_e)/\Gamma(\pi^+ \to \mu^+\nu_\mu)$ using the Run 1 dataset of the NA62 experiment. The primary goal of this analysis is to obtain the absolute number of π_{e2} and $\pi_{\mu 2}$ decays. The analysis strategy is built directly upon the framework established in the preceding chapter for the Heavy Neutral Lepton (HNL) search. Consequently, the event selection criteria and the reconstruction of the missing mass squared (m^2_{miss}) under both kaon and pion hypotheses are identical to those previously described in Section 3.3.

A critical input to this computation is the total decay kaon flux. As the dataset, MC samples and selection are shared with the HNL search, the value of N_K derived in Section 3.4 via the normalization region $|m^2_{miss}(\pi)| < 0.01\,\mathrm{GeV}^2/c^4$ is adopted here without modification. This serves the purpose of subtracting kaon contamination in the pion sample.

While the foundational selection remains unchanged, the objective of the Signal Region (SR), defined as the range $|m^2_{miss}(\pi)| < 0.001\,\mathrm{GeV}^2/c^4$, undergoes a fundamental shift. In the HNL analysis, this region (referred to as the normalization region) was utilized to obtain the total pion flux using the well-defined peak produced by the π_{e2} and its known branching ratio $BR(\pi_{e2})$. In this study, the π_{e2} counting and background subtraction must be accounted for without relying on flux normalization, as the number of decays is the unknown quantity to be measured. Crucially, a non-negligible fraction of $\pi_{\mu 2}$ decays also passes the selection and leaks into the SR. This contamination necessitates a new method to determine the number of $\pi_{\mu 2}$ decays in the SR without relying on external branching fraction values.

To achieve the measurement of the $\pi_{\mu 2}$ component, a Control Region (CR) is introduced in the range $0.004 < m^2_{miss}(\pi) < 0.016\,\mathrm{GeV}^2/c^4$. As this region is populated not only by a broad spectrum of $\pi_{\mu 2}$ decays in flight (where the secondary muon decays $\mu^+ \to e^+\nu_e\bar{\nu}_\mu$), but also by significant contributions from upstream $\pi_{\mu 2}$ decays and muon halo events. While these processes mimic the positron signature of the signal, they exhibit distinct kinematic properties due to the additional missing neutrinos or different origin vertices. The core logic of this analysis is to utilize the CR to extract the total $\pi_{\mu 2}$ decays directly from the data via a statistical fit capable of decomposing these overlapping components; this measurement simultaneously constrains the fraction of $\pi_{\mu 2}$ decays leaking into the SR. This dual purpose allows for a precise background subtraction and a self-contained computation of R_π using the same event sample as the HNL search.

The ratio R_π is thus determined by the ratio of the total number of π_{e2} to $\pi_{\mu 2}$ decays, which are obtained by correcting the measured yields in the Signal and Control Regions by their respective acceptances (A):

$$R_\pi = \frac{N^{\text{total}}_{\pi_{e2}}}{N^{\text{total}}_{\pi_{\mu 2}}} = \frac{N^{\text{SR}}_{\pi_{e2}}/A^{\text{SR}}_{\pi_{e2}}}{N^{\text{CR}}_{\pi_{\mu 2}}/A^{\text{CR}}_{\pi_{\mu 2}}} \tag{4.1}$$

where $N^{\text{SR}}_{\pi_{e2}}$ and $N^{\text{CR}}_{\pi_{\mu 2}}$ represent the background-subtracted event counts extracted from the fit in the Signal and Control Regions, respectively, and $A^{\text{SR}}_{\pi_{e2}}$ and $A^{\text{CR}}_{\pi_{\mu 2}}$ are the corresponding geometric and kinematic acceptances derived from simulation.

The remainder of this chapter details the subtraction of residual kaon contamination, the decomposition of the Control Region components, and the final statistical extraction of the ratio.

4.2 The Kaon-Subtracted Spectrum

Before the pion decay components can be separated, the raw missing mass squared spectrum ($m^2_{miss}(\pi)$) must be purified of kaon contamination.

As the dataset and event selection are identical to the HNL search, the normalization factors for the Kaon MC samples are the same as from the HNL analysis. For this analysis, the scaled kaon templates representing the kaon decays reconstructed under the pion mass hypothesis is subtracted bin-by-bin from the observed Run1 data. This yields the kaon-subtracted spectrum on a bin-by-bin basis:

$$N_{\text{clean}}(i) = N_{\text{data}}(i) - N_K(i)$$

where $N_K(i)$ is the content of bin i in the already-normalized kaon Monte Carlo histogram. This operation is applied individually to every bin i falling within the Control Region (CR) and the Signal Region (SR):

$$\begin{aligned} N_{\text{clean}}(i) &= N_{\text{data}}(i) - N_K(i), \quad \forall i \in \text{CR}, \\ N_{\text{clean}}(j) &= N_{\text{data}}(j) - N_K(j), \quad \forall j \in \text{SR}. \end{aligned}$$

The total event counts for each region are then obtained by summing these bin-wise cleaned values over their respective ranges:

$$\begin{aligned} N_{\text{CR, clean}} &= \sum_{i \in \text{CR}} N_{\text{clean}}(i) = \sum_{i \in \text{CR}} \left(N_{\text{data}}(i) - N_K(i)\right), \\ N_{\text{SR, clean}} &= \sum_{j \in \text{SR}} N_{\text{clean}}(j) = \sum_{j \in \text{SR}} \left(N_{\text{data}}(j) - N_K(j)\right). \end{aligned}$$

Finally, by separating the sums, we arrive at the working equations for the total yields:

$$N_{\text{CR, clean}} = N_{\text{CR, data}} - N^{\text{CR}}_K, \tag{4.2}$$
$$N_{\text{SR, clean}} = N_{\text{SR, data}} - N^{\text{SR}}_K, \tag{4.3}$$

where $N_{\text{CR, data}} = \sum_{i\in\text{CR}} N_{\text{data}}(i)$ and $N_K^{\text{CR}} = \sum_{i\in\text{CR}} N_K(i)$ (and similarly for the Signal Region). This explicit progression from bin-wise subtraction to regional integration ensures the statistical treatment remains consistent throughout the analysis.

The subtraction is illustrated in Figure 4.1. The top panel displays the raw $m^2_{miss}(\pi)$ spectrum from data overlaid with the scaled kaon Monte Carlo template. The bottom panel shows the resulting kaon-subtracted spectrum, where the broad background component is removed, alongside the raw data. The vertical dashed lines indicate the Signal (red) and Control (blue) Regions.

4.2.1 Statistical Uncertainty Treatment

A critical aspect of this subtraction is the correct propagation of statistical uncertainties, particularly because the Monte Carlo samples utilize event weights to match the data distribution. The variance of the cleaned spectrum in each bin is the sum of the variances of the independent data and kaon background terms:

$$\sigma^2_{\text{clean}}(i) = \sigma^2_{\text{data}}(i) + \sigma^2_K(i)$$

The data variance follows Poisson statistics, $\sigma^2_{\text{data}}(i) = N_{\text{data}}(i)$. The kaon variance term, $\sigma^2_K(i)$, accounts for the weighted nature of the simulation and the contribution of multiple decay channels. It is defined as the sum over all kaon decay channels k, where Z_k is the scaling factor for channel k and $w_{j,k}$ are the event weights:

$$\sigma^2_K(i) = \sum_k Z_k^2 \left(\sum_{j\in\text{bin } i} w^2_{j,k} \right)$$

To obtain the total uncertainty for the integrated yields in the Control Region (CR) and Signal Region (SR), we sum the variances of all individual bins. We define the total kaon variance terms, $\sigma^2_{N_K^{\text{CR}}}$ and $\sigma^2_{N_K^{\text{SR}}}$, as the sum of these channel-specific contributions over all bins in each region:

$$\sigma^2_{N_K^{\text{CR}}} = \sum_{i\in\text{CR}} \sigma^2_K(i), \quad \sigma^2_{N_K^{\text{SR}}} = \sum_{j\in\text{SR}} \sigma^2_K(j)$$

Consequently, the total statistical variances for the cleaned yields are expressed compactly as:

$$\sigma^2_{N_{\text{CR, clean}}} = N_{\text{CR, data}} + \sigma^2_{N_K^{\text{CR}}}, \tag{4.4}$$

$$\sigma^2_{N_{\text{SR, clean}}} = N_{\text{SR, data}} + \sigma^2_{N_K^{\text{SR}}}. \tag{4.5}$$

This formulation ensures that the reduced statistical power of weighted Monte Carlo samples is correctly propagated into the final error budget. The resulting standard deviations ($\sigma = \sqrt{\sigma^2}$) are assigned to the cleaned histogram bins and regional totals. Consistent with the scope of this feasibility study, the systematic uncertainty on the normalization factors Z_k is neglected.

4.3 Decomposition of the Control Region

With the kaon contamination removed, the analysis focuses on the Control Region. The primary purpose of this region is to isolate the $\pi_{\mu 2}$ decay component, which provides the value required for the R_π calculation.

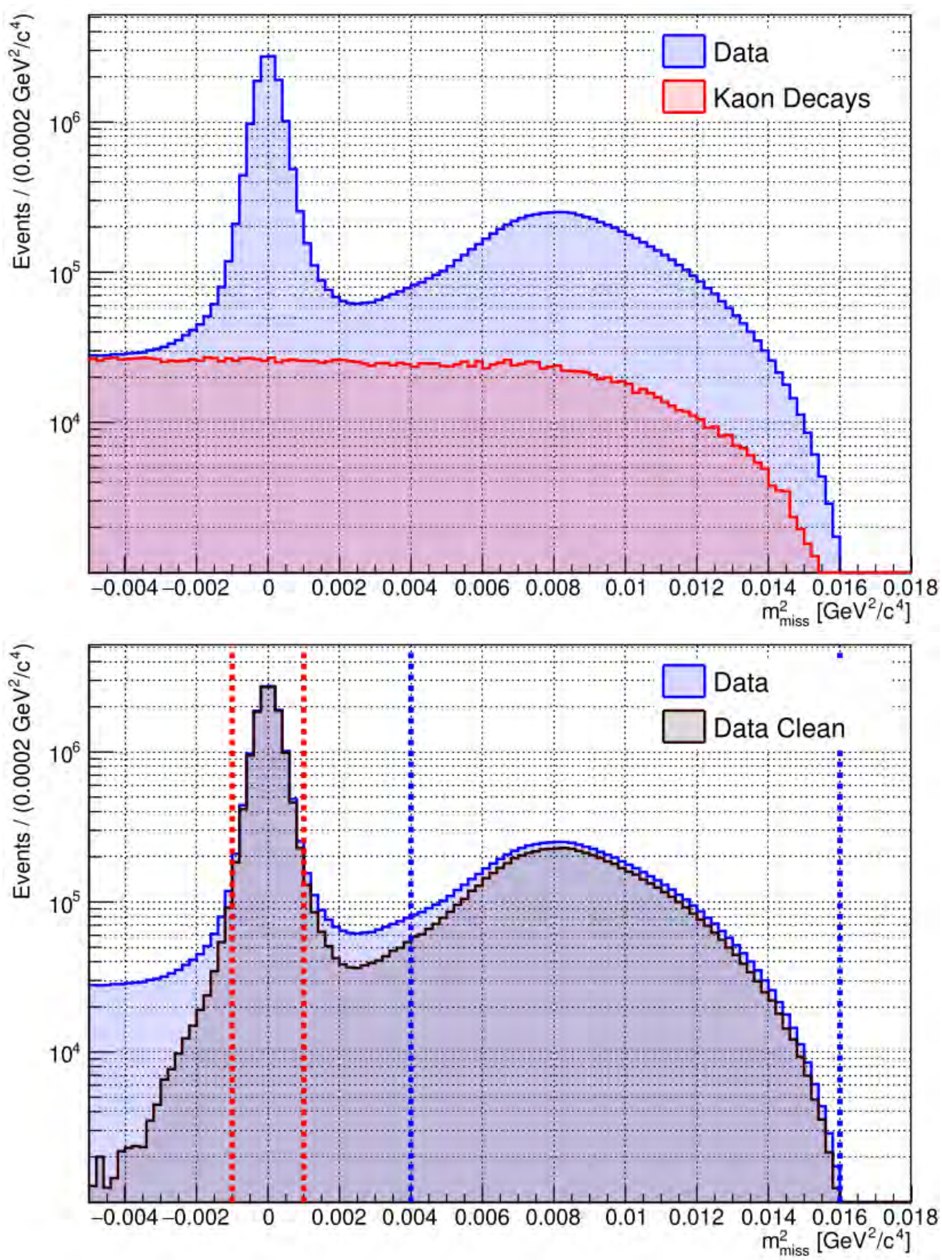

Figure 4.1: Raw missing mass squared spectrum overlaid with the scaled kaon simulation (top) and the resulting kaon-subtracted spectrum with the background removed (bottom). Vertical dashed lines indicate the Signal (red) and Control (blue) regions.

As established in the previous chapter, the kaon-free spectrum in this region consists of a superposition of three distinct physical processes: $\pi_{\mu 2}$ decays, Upstream $\pi_{\mu 2}$ decays, and the Muon Halo background. Since these processes overlap significantly in the m^2_{miss} spectrum, a simple counting method is insufficient to separate the $\pi_{\mu 2}$ signal from the backgrounds. Instead, a template fit is employed to disentangle their respective contributions.

To enable this fit, the Monte Carlo histograms for each component are converted into normalized Probability Density Functions (PDFs) over the Control Region. This step ensures that the resulting fit parameters directly correspond to the physical number of events for each process, allowing for a precise extraction of the $\pi_{\mu 2}$ yield needed for the final ratio.

4.4 Statistical Extraction via NLL Fit

To separate the overlapping physical processes contributing to the missing-mass squared (m^2_{miss}) spectrum in the Control Region (CR), a multi-template extended binned Negative Log-Likelihood (NLL) fit is performed. This approach relies on the observed data to directly determine the component fractions of the decay components.

4.4.1 Template Normalization and PDF Construction

The fit range is defined over the m^2_{miss} CR region, divided into N_{bins} discrete bins. Three distinct MC process templates are utilized:

- $\pi_{\mu 2}$ decays: $T_{\pi_{\mu 2}}$, representing $\pi^+ \to \mu^+ \nu$ decays followed by $\mu^+ \to e^+ \nu \bar{\nu}$ decays.
- Upstream $\pi_{\mu 2}$ decays: $T_{\pi^{\text{ups}}_{\mu 2}}$, representing $\pi^+ \to \mu^+ \nu$ upstream decays followed by $\mu^+ \to e^+ \nu \bar{\nu}$ decays.
- Muon Halo: T_{MH}, representing the Muon Halo component.

For each process $k \in \{\pi_{\mu 2}, \pi^{\text{ups}}_{\mu 2}, \text{MH}\}$, a normalized Probability Density Function (PDF) $P_k(i)$ is constructed at bin i by normalizing the template counts over the fit region:

$$P_k(i) = \frac{h_k(i)}{\sum_{j=1}^{N_{\text{bins}}} h_k(j)}$$

where $h_k(i)$ is the MC count in bin i. By construction, $\sum_{i=1}^{N_{\text{bins}}} P_k(i) = 1.0$ for all processes.

Figure 4.2 shows the normalized PDF shapes of the signal and background components before the fit was performed.

4.4.2 Parametrization and Predicted Model Events

The total number of observed data events in the fit window is denoted as $N_{\text{total}} = \sum_{i=1}^{N_{\text{bins}}} y_i$, where y_i is the measured data count in bin i.

The model is parametrized using two independent fractions: $f_{\pi_{\mu 2}}$ (the fraction of $\pi_{\mu 2}$ events) and $f_{\pi^{\text{ups}}_{\mu 2}}$ (the fraction of Upstream $\pi_{\mu 2}$ events). Total probability conservation requires that the remaining fraction belongs to the Muon Halo:

$$f_{\text{MH}} = 1.0 - f_{\pi_{\mu 2}} - f_{\pi^{\text{ups}}_{\mu 2}}$$

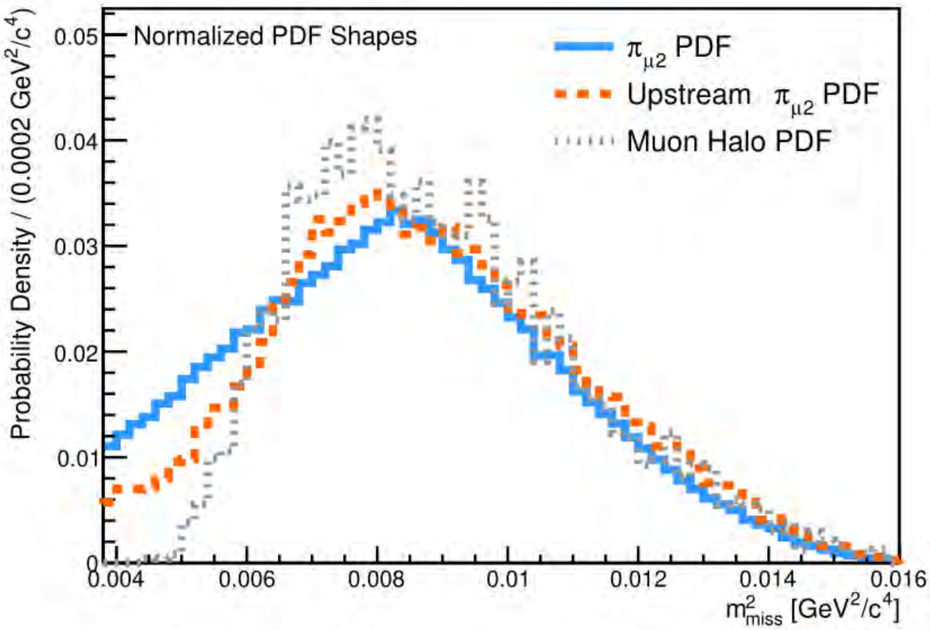

Figure 4.2: Pre-fit normalized template shapes in the Control Region.

The expected event count μ_i in bin i as a function of the parameter vector $\boldsymbol{\theta} = (f_{\pi_{\mu 2}}, f_{\pi^{\text{ups}}_{\mu 2}})^T$ is modeled as:

$$\mu_i(\boldsymbol{\theta}) = N_{\text{total}} \cdot \left[f_{\pi_{\mu 2}} \cdot P_{\pi_{\mu 2}}(i) + f_{\pi^{\text{ups}}_{\mu 2}} \cdot P_{\pi^{\text{ups}}_{\mu 2}}(i) + (1.0 - f_{\pi_{\mu 2}} - f_{\pi^{\text{ups}}_{\mu 2}}) \cdot P_{\text{MH}}(i) \right]$$

To ensure numerical stability during minimization, a smooth quadratic penalty function $Q(\boldsymbol{\theta})$ is added to prevent unphysical parameter space exploration ($f_{\text{MH}} < 0$):

$$Q(\boldsymbol{\theta}) = \begin{cases} 0 & \text{if } f_{\text{MH}} \geq 0 \\ \lambda \cdot (f_{\text{MH}})^2 & \text{if } f_{\text{MH}} < 0 \end{cases}$$

where $\lambda = 10^{10}$ is a large penalty multiplier.

4.4.3 Function Minimization

Assuming Poisson statistics in each bin, the likelihood $\mathcal{L}(\boldsymbol{\theta})$ of observing the dataset given the expected yields $\mu_i(\boldsymbol{\theta})$ is:

$$\mathcal{L}(\boldsymbol{\theta}) = \prod_{i=1}^{N_{\text{bins}}} \frac{\mu_i(\boldsymbol{\theta})^{y_i} e^{-\mu_i(\boldsymbol{\theta})}}{y_i!}$$

The minimization is executed by minimizing the Negative Log-Likelihood (NLL $= -\ln \mathcal{L}$) using the ROOT framework [44](version 6.32.08):

$$\mathrm{NLL}(\boldsymbol{\theta}) = \sum_{i=1}^{N_{\mathrm{bins}}} \left[\mu_i(\boldsymbol{\theta}) - y_i \ln\left(\mu_i(\boldsymbol{\theta})\right)\right] + Q(\boldsymbol{\theta})$$

Constant factorial terms ($\ln y_i!$) independent of $\boldsymbol{\theta}$ are omitted from the function.

4.4.4 Error Matrix and Covariance Propagation

At the likelihood minimum $\hat{\boldsymbol{\theta}} = (\hat{f}_{\pi_{\mu 2}}, \hat{f}_{\pi_{\mu 2}^{\mathrm{ups}}})^T$, the full covariance matrix $\boldsymbol{V}$ is:

$$\boldsymbol{V} = \begin{pmatrix} \sigma^2_{f_{\pi_{\mu 2}}} & \mathrm{Cov}(f_{\pi_{\mu 2}}, f_{\pi_{\mu 2}^{\mathrm{ups}}}) \\ \mathrm{Cov}(f_{\pi_{\mu 2}}, f_{\pi_{\mu 2}^{\mathrm{ups}}}) & \sigma^2_{f_{\pi_{\mu 2}^{\mathrm{ups}}}} \end{pmatrix}$$

Due to the kinematic shape similarity between $T_{\pi_{\mu 2}}$ and $T_{\pi_{\mu 2}^{\mathrm{ups}}}$, the two parameters exhibit a strong negative correlation coefficient:

$$r(f_{\pi_{\mu 2}}, f_{\pi_{\mu 2}^{\mathrm{ups}}}) = \frac{\mathrm{Cov}(f_{\pi_{\mu 2}}, f_{\pi_{\mu 2}^{\mathrm{ups}}})}{\sigma_{f_{\pi_{\mu 2}}} \cdot \sigma_{f_{\pi_{\mu 2}^{\mathrm{ups}}}}} = -0.9335$$

When computing the combined $\pi_{\mu 2}$ and upstream $\pi_{\mu 2}$ fraction $f_{\mu_s} = f_{\pi_{\mu 2}} + f_{\pi_{\mu 2}^{\mathrm{ups}}}$, the complete parameter variance propagation equation is used:

$$\sigma^2_{f_{\mu_s}} = \sigma^2_{f_{\pi_{\mu 2}}} + \sigma^2_{f_{\pi_{\mu 2}^{\mathrm{ups}}}} + 2 \cdot \mathrm{Cov}(f_{\pi_{\mu 2}}, f_{\pi_{\mu 2}^{\mathrm{ups}}})$$

The negative off-diagonal covariance term ($2 \cdot \mathrm{Cov}$) accounts for parameter trade-offs, preventing overestimation of the combined background uncertainty.

4.4.5 Goodness-of-Fit and Bin-by-Bin Residuals

To evaluate fit performance while accounting for both Poisson data fluctuations and finite Monte Carlo template statistics, the combined bin variance $\sigma^2_{i,\mathrm{tot}}$ is computed as:

$$\sigma^2_{i,\mathrm{tot}} = y_i + \left(N_{\mathrm{total}} \hat{f}_{\pi_{\mu 2}} \cdot \delta P_{\pi_{\mu 2}}(i)\right)^2 + \left(N_{\mathrm{total}} \hat{f}_{\pi_{\mu 2}^{\mathrm{ups}}} \cdot \delta P_{\pi_{\mu 2}^{\mathrm{ups}}}(i)\right)^2 + \left(N_{\mathrm{total}} \hat{f}_{\mathrm{MH}} \cdot \delta P_{\mathrm{MH}}(i)\right)^2$$

where $\delta P_k(i)$ represents the statistical error of template k in bin i.

The total chi-squared statistic over the N_{bins} histogram bins is evaluated as:

$$\chi^2 = \sum_{i=1}^{N_{\mathrm{bins}}} \frac{\left(y_i - \mu_i(\hat{\boldsymbol{\theta}})\right)^2}{\sigma^2_{i,\mathrm{tot}}}$$

For N_{bins} bins and $p = 2$ fitted parameters ($f_{\pi_{\mu 2}}$ and $f_{\pi_{\mu 2}^{\mathrm{ups}}}$), the number of degrees of freedom is $\mathrm{ndf} = N_{\mathrm{bins}} - 2$. The pull value P_i for each bin i is calculated as:

$$P_i = \frac{y_i - \mu_i(\hat{\boldsymbol{\theta}})}{\sigma_{i,\mathrm{tot}}}$$

A successful fit is confirmed if the pull values P_i follow a standard normal distribution $\mathcal{N}(0, 1)$ (mean = 0, width = 1) and the reduced chi-squared satisfies $\chi^2/\mathrm{ndf} \approx 1$.

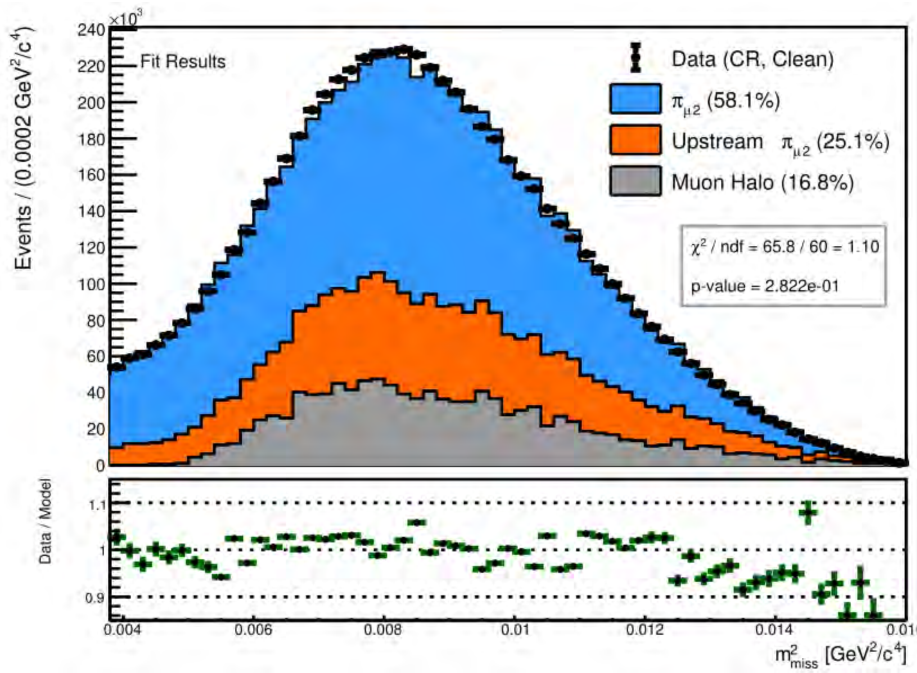

Figure 4.3: Post-fit overlay of the observed data (points) and the model prediction (colored stacked components). The bottom plot shows the data-to-model ratio, which remains within $\pm 10\%$. Goodness-of-fit statistics embedded in the plot indicate a reduced chi-squared of 1.10 and a p-value of 0.28.

4.4.6 Fit Results

The comparison between the observed data and the model is shown in Figure 4.3. The data points follow the model prediction closely across the range, with differences staying within $\pm 10\%$. The plot includes statistical tests that yield a reduced chi-squared value of 1.10 and a p-value of 0.28. These numbers indicate that the small differences seen are normal random variations and that the model describes the data well.

Figure 4.4 shows the pull for each bin. These differences vary within expected limits ($\pm 3\sigma$ standard deviations).

Figure 4.5 displays the relationship between the two main parameters determined by the fit. The results show a strong negative correlation ($r = -0.9335$), meaning that if one parameter value increases, the other tends to decrease. This strong link helps to reduce the overall uncertainty in the final background estimate.

4.5 Signal Region Background Extraction

Following the determination of the CR region composition ($\hat{f}_{\pi_{\mu 2}}, \hat{f}_{\pi_{\mu 2}^{\mathrm{ups}}}$), the yields are projected into the SR region using MC-derived acceptance transfer ratios. This step isolates the $\pi^+ \to e^+ \nu_e$

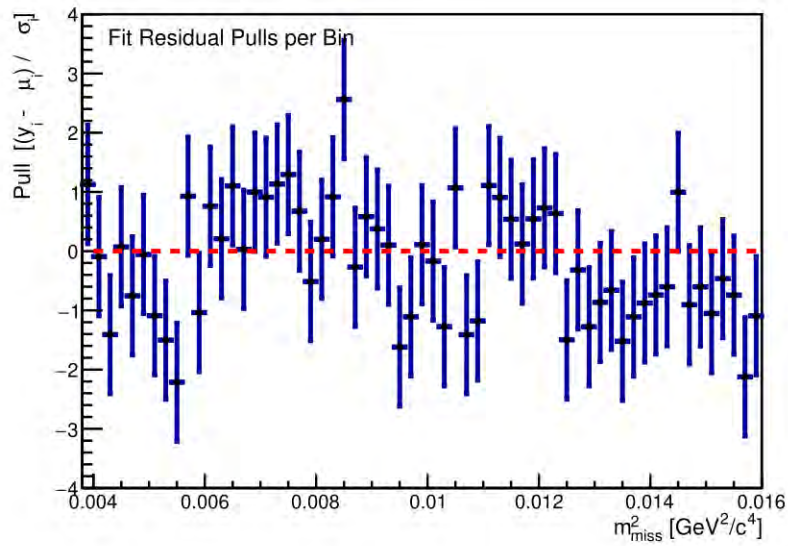

Figure 4.4: Distribution of residual pulls. The residuals fluctuate within $\pm 3\sigma$ standard deviations with no visible patterns or localized structures.

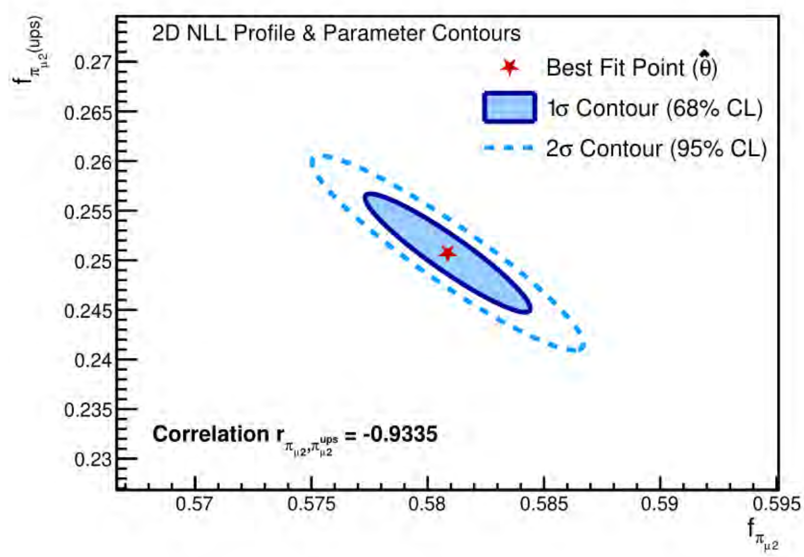

Figure 4.5: Two-dimensional confidence contours for the fitted parameters at 1σ (68% CL) and 2σ (95% CL) levels. The best-fit point is marked at the center of the narrow, tilted ellipse, which indicates a strong negative correlation ($r = -0.9335$) between the parameters.

signal from $\pi_{\mu 2}$ and upstream $\pi_{\mu 2}$ contamination without relying on pion flux normalizations.

4.5.1 Acceptance Transfer Ratios and Background Projection

For a given physical process k, the selection acceptance ratio α_k mapping events from the CR into the SR is defined as:

$$\alpha_k = \frac{A_k^{\mathrm{SR}}}{A_k^{\mathrm{CR}}} = \frac{N_{k,\mathrm{MC}}^{\mathrm{SR}}}{N_{k,\mathrm{MC}}^{\mathrm{CR}}}.$$

Because the Muon Halo is localized strictly within the CR, its transfer factor to the signal region is identically zero ($\alpha_{\mathrm{MH}} = 0$). Consequently, the transferred background contaminating the SR consists exclusively of $\pi_{\mu 2}$ and $\pi_{\mu 2}^{\mathrm{ups}}$ events:

$$N_{\mu_s}^{\mathrm{SR}} = N_{\pi_{\mu 2}}^{\mathrm{SR}} + N_{\pi_{\mu 2}^{\mathrm{ups}}}^{\mathrm{SR}} = N_{\mathrm{CR,\ clean}} \left(\hat{f}_{\pi_{\mu 2}} \alpha_{\pi_{\mu 2}} + \hat{f}_{\pi_{\mu 2}^{\mathrm{ups}}} \alpha_{\pi_{\mu 2}^{\mathrm{ups}}} \right).$$

4.5.2 Uncertainty Propagation on Transferred Background

The uncertainty on the transferred background yield $\sigma^2_{N_{\mu_s}^{\mathrm{SR}}}$ propagates both the uncertainty in the cleaned CR yield and the covariance from the likelihood fit:

$$\begin{aligned}\sigma^2_{N_{\mu_s}^{\mathrm{SR}}} = &\left(\hat{f}_{\pi_{\mu 2}} \alpha_{\pi_{\mu 2}} + \hat{f}_{\pi_{\mu 2}^{\mathrm{ups}}} \alpha_{\pi_{\mu 2}^{\mathrm{ups}}} \right)^2 \sigma^2_{N_{\mathrm{CR,\ clean}}} \\ &+ N^2_{\mathrm{CR,\ clean}} \left[\alpha^2_{\pi_{\mu 2}} \sigma^2_{f_{\pi_{\mu 2}}} + \alpha^2_{\pi_{\mu 2}^{\mathrm{ups}}} \sigma^2_{f_{\pi_{\mu 2}^{\mathrm{ups}}}} + 2\alpha_{\pi_{\mu 2}} \alpha_{\pi_{\mu 2}^{\mathrm{ups}}} \mathrm{Cov}(f_{\pi_{\mu 2}}, f_{\pi_{\mu 2}^{\mathrm{ups}}}) \right]\end{aligned}$$

The strong negative correlation between the parameter fractions ($\mathrm{Cov}(f_{\pi_{\mu 2}}, f_{\pi_{\mu 2}^{\mathrm{ups}}}) < 0$) suppresses the second bracketed term, significantly tightening the overall background uncertainty transferred into the signal region.

4.6 Final R_π Ratio Inputs

4.6.1 Experimental R_π Equation

The experimental ratio of decay rates $R_\pi = \Gamma(\pi^+ \to e^+\nu_e)/\Gamma(\pi^+ \to \mu^+\nu_\mu)$ is evaluated by combining the background extracted signal region yield with the pure fiducial normalization denominator:

$$R_\pi = \frac{N_{\pi e2}^{\mathrm{total}}}{N_{\pi\mu 2}^{\mathrm{total}}} = \frac{\left(N_{\mathrm{SR,\ data}} - N_K^{\mathrm{SR}} - N_{\mathrm{CR,\ clean}} \left[\hat{f}_{\pi_{\mu 2}} \alpha_{\pi_{\mu 2}} + \hat{f}_{\pi_{\mu 2}^{\mathrm{ups}}} \alpha_{\pi_{\mu 2}^{\mathrm{ups}}} \right] \right) / A_{\pi e2}^{\mathrm{SR}}}{\left(N_{\mathrm{CR,\ clean}} \cdot \hat{f}_{\pi_{\mu 2}} \right) / A_{\pi_{\mu 2}}^{\mathrm{CR}}}, \tag{4.6}$$

where $A_{\pi e2}^{\mathrm{SR}}$ is the selection acceptance for $\pi^+ \to e^+\nu_e$ decays occurring in the SR region, and $A_{\pi_{\mu 2}}^{\mathrm{CR}}$ is the selection acceptance for $\pi^+ \to \mu^+\nu_\mu$ decays in the CR region.

4.6.2 Statistical Uncertainty Components

Statistical Uncertainty on $N_{\pi_{e2}}^{\textbf{total}}$

The statistical uncertainty of total π_{e2} decays $N_{\pi e2}^{\text{total}}$ is determined by the uncertainties on $N_{\text{SR, clean}}$ and the transferred background error:

$$\sigma^2_{N_{\pi e2}^{\text{total}}} = \frac{1}{\left(A_{\pi e2}^{\text{SR}}\right)^2} \left[\sigma^2_{N_{\text{SR, clean}}} + \sigma^2_{N_{\mu_s}^{\text{SR}}}\right]. \tag{4.7}$$

Statistical Uncertainty on $N_{\pi_{\mu 2}}^{\textbf{total}}$

The statistical uncertainty of total $\pi_{\mu 2}$ decays $N_{\pi\mu 2}^{\text{total}}$ is determined by the uncertainties on $N_{\text{CR, clean}}$ and the fitted fraction $\hat{f}_{\pi_{\mu 2}}$:

$$\sigma^2_{N_{\pi\mu 2}^{\text{total}}} = \frac{1}{\left(A_{\pi\mu 2}^{\text{CR}}\right)^2} \left[\hat{f}^2_{\pi_{\mu 2}} \cdot \sigma^2_{N_{\text{CR, clean}}} + N^2_{\text{CR, clean}} \cdot \sigma^2_{f_{\pi_{\mu 2}}}\right]. \tag{4.8}$$

Cross Covariance Factor

Because both the numerator and denominator ($N_{\pi e2}^{\text{total}}$ and $N_{\pi\mu 2}^{\text{total}}$) depend explicitly on the cleaned control region count $N_{\text{CR, clean}}$ and the fit fraction $\hat{f}_{\pi_{\mu 2}}$, with the numerator additionally depending on $\hat{f}_{\pi_{\mu 2}^{\text{ups}}}$, their cross covariance $\text{Cov}\left(N_{\pi e2}^{\text{total}}, N_{\pi\mu 2}^{\text{total}}\right)$ is evaluated via first order error propagation over these shared variables:

$$\begin{aligned}\text{Cov}\left(N_{\pi e2}^{\text{total}}, N_{\pi\mu 2}^{\text{total}}\right) = \left(\frac{\partial N_{\pi e2}^{\text{total}}}{\partial f_{\pi_{\mu 2}}}\right)\left(\frac{\partial N_{\pi\mu 2}^{\text{total}}}{\partial f_{\pi_{\mu 2}}}\right)\sigma^2_{f_{\pi_{\mu 2}}} + \left(\frac{\partial N_{\pi e2}^{\text{total}}}{\partial f_{\pi_{\mu 2}^{\text{ups}}}}\right)\left(\frac{\partial N_{\pi\mu 2}^{\text{total}}}{\partial f_{\pi_{\mu 2}}}\right)\text{Cov}\left(f_{\pi_{\mu 2}}, f_{\pi_{\mu 2}^{\text{ups}}}\right) \\ + \left(\frac{\partial N_{\pi e2}^{\text{total}}}{\partial N_{\text{CR, clean}}}\right)\left(\frac{\partial N_{\pi\mu 2}^{\text{total}}}{\partial N_{\text{CR, clean}}}\right)\sigma^2_{N_{\text{CR, clean}}}.\end{aligned}$$

where the partial derivatives of the numerator yield $N_{\pi e2}^{\text{total}}$ and denominator yield $N_{\pi\mu 2}^{\text{total}}$ with respect to each fit fraction and control region count are defined explicitly as:

$$\frac{\partial N_{\pi e2}^{\text{total}}}{\partial f_{\pi_{\mu 2}}} = -\frac{N_{\text{CR, clean}} \cdot \alpha_{\pi_{\mu 2}}}{A_{\pi e2}^{\text{SR}}}, \qquad \frac{\partial N_{\pi\mu 2}^{\text{total}}}{\partial f_{\pi_{\mu 2}}} = +\frac{N_{\text{CR, clean}}}{A_{\pi\mu 2}^{\text{CR}}},$$

$$\frac{\partial N_{\pi e2}^{\text{total}}}{\partial f_{\pi_{\mu 2}^{\text{ups}}}} = -\frac{N_{\text{CR, clean}} \cdot \alpha_{\pi_{\mu 2}^{\text{ups}}}}{A_{\pi e2}^{\text{SR}}}, \qquad \frac{\partial N_{\pi\mu 2}^{\text{total}}}{\partial f_{\pi_{\mu 2}^{\text{ups}}}} = 0,$$

$$\frac{\partial N_{\pi e2}^{\text{total}}}{\partial N_{\text{CR, clean}}} = -\frac{\hat{f}_{\pi_{\mu 2}}\alpha_{\pi_{\mu 2}} + \hat{f}_{\pi_{\mu 2}^{\text{ups}}}\alpha_{\pi_{\mu 2}^{\text{ups}}}}{A_{\pi e2}^{\text{SR}}}, \qquad \frac{\partial N_{\pi\mu 2}^{\text{total}}}{\partial N_{\text{CR, clean}}} = +\frac{\hat{f}_{\pi_{\mu 2}}}{A_{\pi\mu 2}^{\text{CR}}}.$$

Substituting these partial derivatives into the first-order expansion yields the final explicit equation for the cross covariance:

$$\begin{aligned}\mathrm{Cov}\left(N^{\text{total}}_{\pi e2}, N^{\text{total}}_{\pi\mu 2}\right) = \frac{N^2_{\text{CR, clean}} \cdot \hat{f}_{\pi_{\mu 2}}}{A^{\text{SR}}_{\pi e2} \cdot A^{\text{CR}}_{\pi\mu 2}} \left[-\alpha_{\pi_{\mu 2}} \sigma^2_{f_{\pi_{\mu 2}}} - \alpha_{\pi^{\text{ups}}_{\mu 2}} \mathrm{Cov}\left(f_{\pi_{\mu 2}}, f_{\pi^{\text{ups}}_{\mu 2}}\right)\right] \\ - \frac{N_{\text{CR, clean}} \cdot \hat{f}_{\pi_{\mu 2}} \left(\hat{f}_{\pi_{\mu 2}} \alpha_{\pi_{\mu 2}} + \hat{f}_{\pi^{\text{ups}}_{\mu 2}} \alpha_{\pi^{\text{ups}}_{\mu 2}}\right)}{A^{\text{SR}}_{\pi e2} \cdot A^{\text{CR}}_{\pi\mu 2}} \sigma^2_{N_{\text{CR, clean}}}.\end{aligned} \tag{4.9}$$

Statistical Uncertainty on R_π

The total statistical uncertainty $\sigma^2_{R_\pi}$ is constructed by combining numerator uncertainty (Eq. 4.7), denominator uncertainty (Eq. 4.8), and covariance (Eq. 4.9):

$$\sigma^2_{R_\pi} = R^2_\pi \left[\frac{\sigma^2_{N^{\text{total}}_{\pi e2}}}{\left(N^{\text{total}}_{\pi e2}\right)^2} + \frac{\sigma^2_{N^{\text{total}}_{\pi\mu 2}}}{\left(N^{\text{total}}_{\pi\mu 2}\right)^2} - 2\frac{\mathrm{Cov}\left(N^{\text{total}}_{\pi e2}, N^{\text{total}}_{\pi\mu 2}\right)}{N^{\text{total}}_{\pi e2} \cdot N^{\text{total}}_{\pi\mu 2}}\right]. \tag{4.10}$$

4.6.3 Summary of Extracted Parameters

The R_π ratio is extracted by combining the observed data counts, background estimates, and acceptance factors across both the Signal Region (SR) and Control Region (CR).

Table 4.1 summarizes the complete set of input parameters, including the observed counts ($N^{\text{SR}}_{\text{data}}$), subtracted kaon background (N^{SR}_K), control region template fit fractions ($\hat{f}_{\pi_{\mu 2}}$, $\hat{f}_{\pi^{\text{ups}}_{\mu 2}}$), and their associated acceptance transfer ratios. From these inputs, the transferred $\pi_{\mu 2} + \pi_{\mu 2}$ (μ_s) background yield ($N^{\text{SR}}_{\mu_s}$), cleaned signal yield ($N^{\text{SR}}_{\pi e2}$), and total π_{e2} and $\pi_{\mu 2}$ decay yields ($N^{\text{total}}_{\pi e2}$ and $N^{\text{total}}_{\pi\mu 2}$) are evaluated alongside their respective statistical uncertainties and cross-covariance.

4.7 Results and Discussion

From the extracted values and acceptance corrections summarized in Section 4.6.3, the R_π ratio is measured to be:

$$R_\pi = (1.2376 \pm 0.0076_{\text{stat}}) \times 10^{-4},$$

corresponding to a relative statistical uncertainty of 0.61%.

The factors governing the statistical precision of R_π are expressed in Eq. 4.10. Their individual contributions to the relative variance budget are presented in Table 4.2.

The error budget reveals the following:

- **Signal Counting Precision:** With $N^{\text{SR}}_{\pi e2} = 1.2343 \times 10^7$ cleaned signal events in the Signal Region, the Poisson counting noise and background subtraction uncertainties contribute a relative variance of only 1.08×10^{-7}, representing just 0.29% of the total error budget.
- **Template Fit Dominance:** The overall statistical error is almost entirely driven by the denominator (100.96% of the total error), which is governed by the uncertainty in the template fit fraction $\hat{f}_{\text{FV}} = 0.5808 \pm 0.0035$ ($\sigma_{f_{\pi_{\mu 2}}}/\hat{f}_{\pi_{\mu 2}} = 0.61\%$).

Table 4.1: Input observables, background estimates, acceptances, and extracted total values entering R_π.

Category	Parameter / Quantity	$x \pm \delta x$ **(Stat.)**
Signal Region (SR)	SR Data Events ($N^{\text{SR}}_{\text{data}}$)	1.26414×10^7
	Background Subtracted K^+ (N^{SR}_{K})	261269 ± 1882
	Transferred μ_s Background ($N^{\text{SR}}_{\mu_s}$)	37546 ± 517
	Cleaned $\pi e2$ Signal ($N^{\text{SR}}_{\pi e2}$)	$(1.2343 \pm 0.0004) \times 10^7$
	Signal Region Acceptance ($A^{\text{SR}}_{\pi e2}$)	0.058569
Control Region (CR) & Fit	Cleaned CR Data ($N_{\text{CR, clean}}$)	$(6.6964 \pm 0.0032) \times 10^6$
	Template Fit Fraction $\hat{f}_{\pi_{\mu 2}}$	0.5808 ± 0.0035
	Template Fit Fraction $\hat{f}_{\pi^{\text{ups}}_{\mu 2}}$	0.2507 ± 0.0060
	Fit Covariance $\text{Cov}(\hat{f}_{\pi_{\mu 2}}, \hat{f}_{\pi^{\text{ups}}_{\mu 2}})$	-1.981207×10^{-5}
	Acceptance Transfer Ratio $\alpha_{\pi_{\mu 2}}$	0.003293
	Acceptance Transfer Ratio $\alpha_{\pi^{\text{ups}}_{\mu 2}}$	0.014734
	Normalization Acceptance ($A^{\text{CR}}_{\pi\mu 2}$)	2.0×10^{-6}
Extracted Total Yields	Total Fiducial Numerator ($N^{\text{total}}_{\pi e2}$)	$(2.1074 \pm 0.0007) \times 10^8$
	Total Fiducial Denominator ($N^{\text{total}}_{\pi\mu 2}$)	$(1.7028 \pm 0.0105) \times 10^{12}$
	Covariance $\text{Cov}\left(N^{\text{total}}_{\pi e2}, N^{\text{total}}_{\pi\mu 2}\right)$	8.364×10^{13}

Table 4.2: Decomposition of the relative statistical uncertainty budget for R_π.

Term	(σ_i^2/N_i^2)	(σ_i/N_i)	**Share of Total** $\sigma^2_{R_\pi}$
Numerator (SR Signal Stats)	1.07994×10^{-7}	0.03%	0.29%
Denominator (CR & Fit)	3.77096×10^{-5}	0.61%	100.96%
Cross-Covariance Correction	-4.66152×10^{-7}	—	-1.25%
Total Contribution (R_π)	$\mathbf{3.73514 \times 10^{-5}}$	**0.6112%**	**100.00%**

- **Cross-Covariance Offset:** The cross-covariance between the numerator transferred background and the denominator gives a negative term (-1.25%), providing a slight cancellation effect in the ratio variance.

In summary, the statistical precision of this R_π measurement for this analysis is not limited by signal region statistics (N_{SR}), but rather by the statistical uncertainty associated with the control region template fit fraction $\hat{f}_{\pi_{\mu 2}}$.

Appendix A

Supplementary Event Selection Studies

This appendix compiles the results of alternative selection criteria investigated during the analysis development. The exclusion of these criteria was based on observed limitations, including random veto effects and insufficient discrimination power between signal and background.

A.1 ANTI0 Studies

Studies were performed to evaluate the potential of the ANTI0 detector to suppress background in the Run 2 dataset. The analysis applied geometric and temporal association criteria between the positron track and hits in the ANTI0 detector. Events were rejected if a spatial association was identified and the time difference between the track and the detector hit satisfied $|t_{\text{track}} - t_{\text{Anti0}}| < \Delta t$, with time windows of $\Delta t = 2, 3, 4\,\text{ns}$ tested.

The effect of this requirement was evaluated by comparing event selection with and without the ANTI0 match. Figure A.1 shows the data distributions with and without the 4 ns time cut, along with the ratio between them. This ratio stays constant across the measured range. Similarly, scaled MC samples for the main HNL background processes were studied. As shown in Figure A.2, the ratio for these background samples is also constant over the same range.

The consistent reduction seen in both data and simulations shows that the ANTI0 criteria do not specifically target the background. Instead, the cut acts as a random veto, lowering the total number of events without improving the signal purity. For this reason, this criterion was not included in the final event selection.

A.2 Investigation of Kinematic Discriminants for $\pi_{\mu 2}$ Background

The primary background to $\pi^+ \to e^+ N$ decays arises from the $\pi^+ \to \mu^+ \nu_\mu$ decay, followed by the muon decaying in flight ($\mu^+ \to e^+ \nu_e \bar{\nu_\mu}$) within the fiducial volume. This decay chain represents the dominant background source for the HNL search.

Studies were performed to determine if this background could be reduced by applying cuts based on kinematic variables. Specifically, the analysis examined the relationship between the

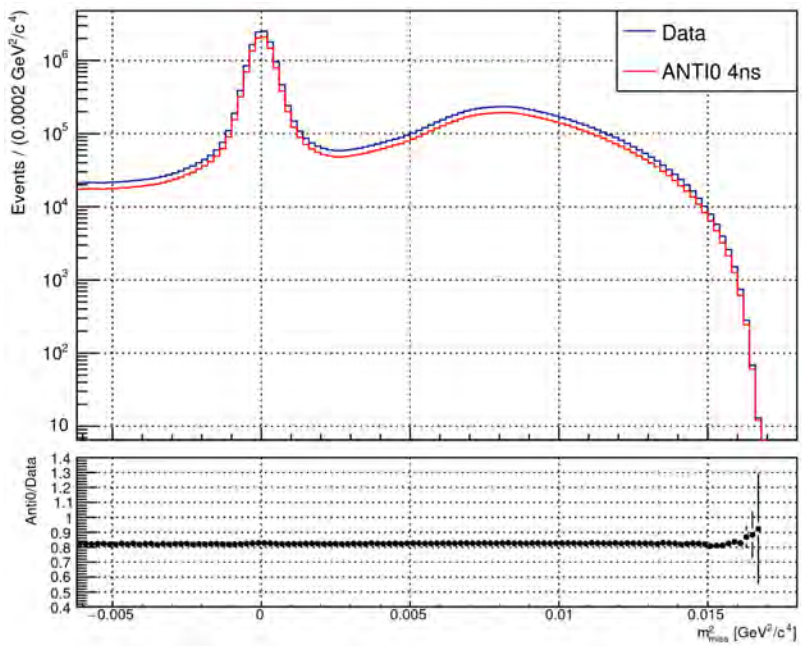

Figure A.1: $m^2_{miss}(\pi^+)$ distributions with ANTI0 cut applied in the time window $|t_{track} - t_{Anti0}| < 4\,\text{ns}$ (blue line) and without ANTI0 cut (red line) for 2021-2022 data. The ratio between the two histograms is also shown.

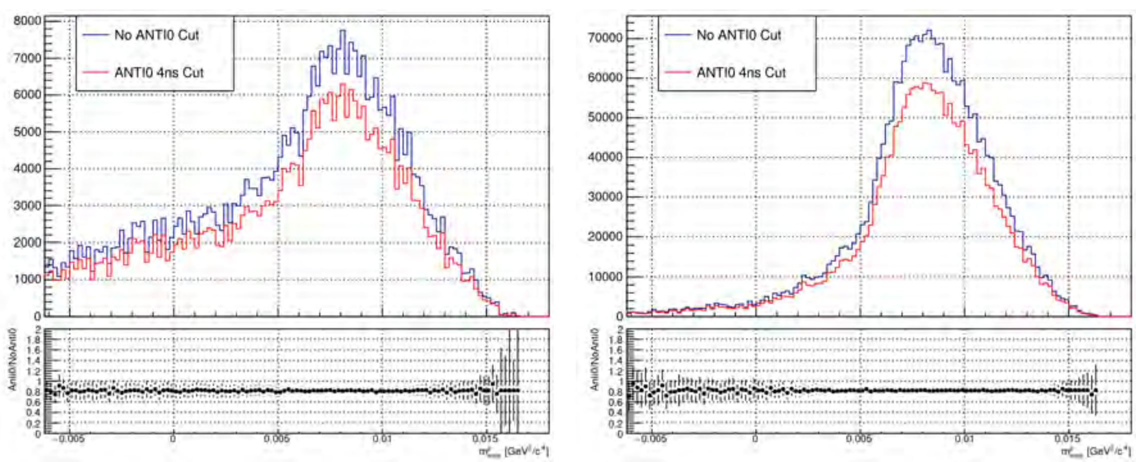

Figure A.2: MC $m^2_{miss}(\pi^+)$ distributions with ANTI0 cut applied in the time window $|t_{track} - t_{Anti0}| < 4\,\text{ns}$ and without ANTI0 cut for km2 upstream (left) and pim2 upstream (right) processes. The ratio between the two histograms is also shown.

emission angle in the mother particle rest frame and the z-position of the decay point, as well as the relationship between the emission angle and the positron energy. The goal was to identify regions where the background and signal distributions differed sufficiently to allow for separation.

However, no effective region suppression was found because the background distributions follow nearly the same pattern as the HNL signal. As shown in Figures A.3 and A.4, the shapes of these distributions are indistinguishable for both processes. Consequently, applying a cut on these variables would remove signal and background events at similar rates, negatively impacting the overall sensitivity of the HNL search without improving the signal purity. Therefore, these criteria were not included in the final selection.

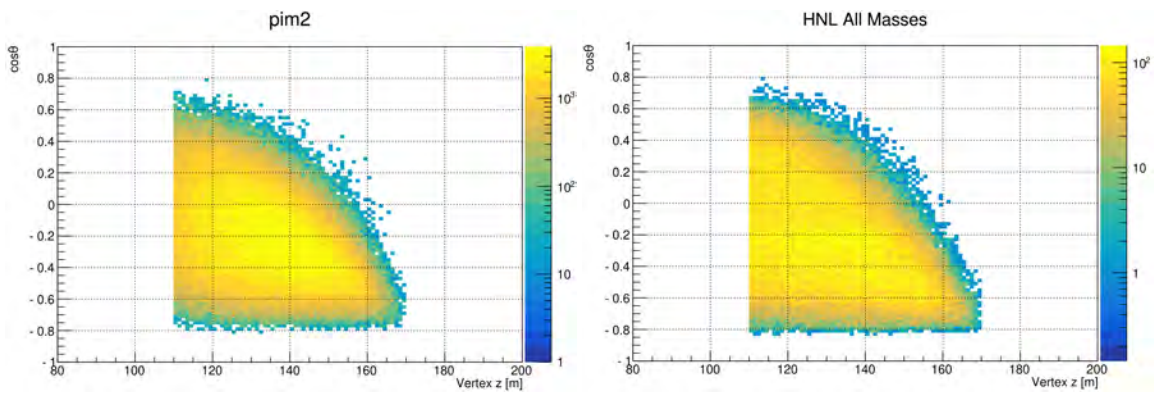

Figure A.3: Distribution of $\cos\theta$ vs. z_{vtx} where θ is the emission angle at the mother rest frame for $\pi^+ \to \mu^+\nu_\mu$ MC (left) and for all the HNL MC simulated masses added (right) after event selection.

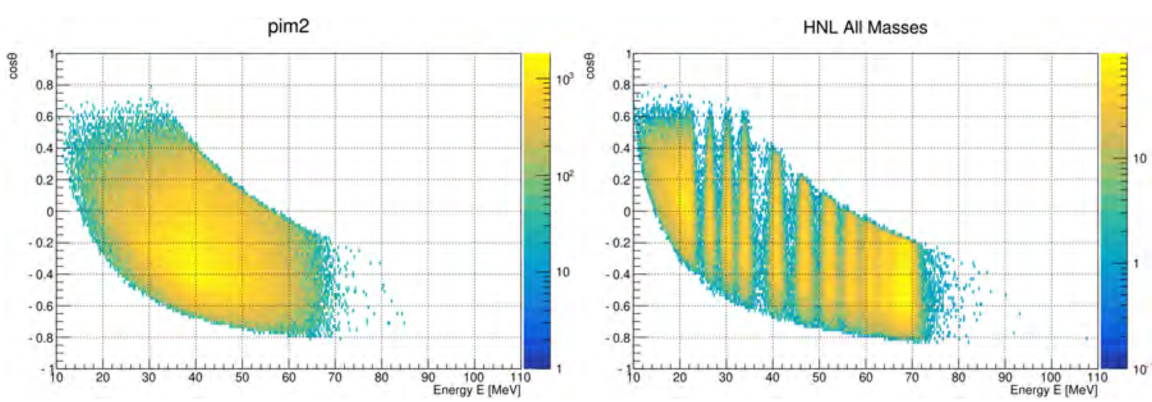

Figure A.4: Distribution of $\cos\theta$ vs. energy in the mother rest frame for Pim2 MC (left) and for all the HNL MC simulated masses added (right) after event selection.

A.3 CHANTI Studies for Upstream Background Rejection

The upstream $\pi^+ \to \mu^+\nu_\mu$ background consists of pions decaying into muons in the region before the fiducial volume, followed by the muon decaying in flight ($\mu^+ \to e^+\nu\bar{\nu}$) inside the fiducial volume. This process represents the second largest source of background for the $\pi^+ \to e^+N$ decay search.

Studies were performed using Run 1 data to determine if the CHANTI detector could suppress this background by applying a timing cut. The analysis required the time difference between the positron track and the CHANTI detector candidate to satisfy $|t_{\text{track}} - t_{\text{CHANTI}}| < \Delta t$, with time windows of $\Delta t = 2, 3, 4\,\text{ns}$ tested.

The results showed that the reduction in upstream $\pi^+ \to \mu^+\nu_\mu$ background events was nearly identical in percentage to the reduction in signal events. Figure A.5 compares the data distribution with the timing cut applied to the distribution with no cut, demonstrating that the cut removes events uniformly across the sample. This behavior confirms that the timing requirement acts as a random veto rather than a selective filter. Because no time window provided satisfactory separation between the background and the signal, this criterion was not included in the final event selection.

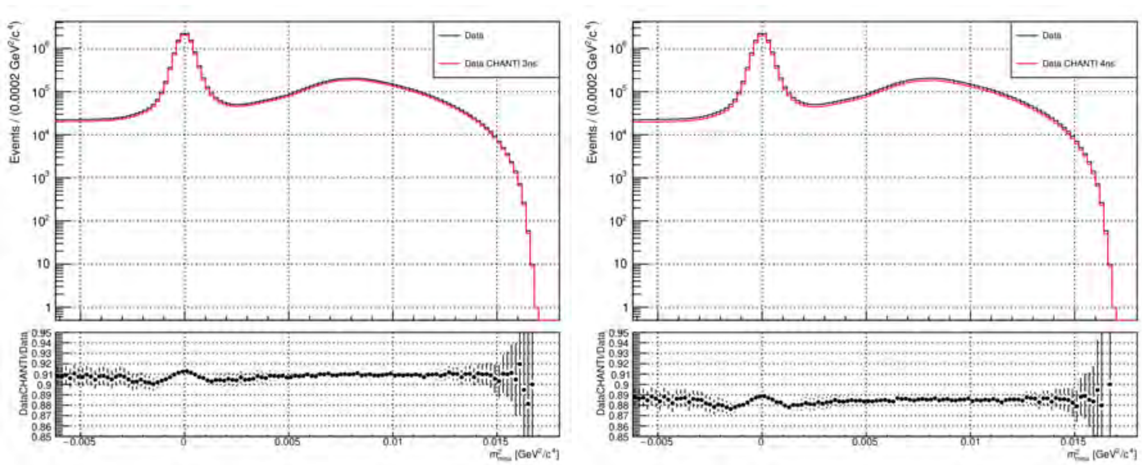

Figure A.5: $m^2_{miss}(\pi^+)$ distributions with CHANTI cut applied in the time window $|t_{track} - t_{Anti0}| <$ 3 ns (left) and in the time window $|t_{track} - t_{Anti0}| < 4$ ns (right) compared with the spectrum without CHANTI cut. The ratio between CHANTI cut and without cut is also shown.

A.4 Evaluation of Event Timestamp and Spill Structure Stability

The event timestamp records the time elapsed from the start of the beam burst to the occurrence of an event. For the 2021 data, the beam spill was divided into two components: the early spill (timestamp less than 2 s) and the main spill (timestamp equal to or greater than 2 s). This division was motivated by observations of anomalous beam behavior during the initial phase of the spill. Specifically, the 2021 beam extraction suffered from instability in the radiofrequency (RF) voltage settings, which caused intervals of anomalously high beam current within the first second of extraction.

Peak structures associated with this early instability appeared in several data groups, as illustrated in Figure A.6. To assess the impact of this background on the HNL search, the early and main spill components were compared. Figure A.7 displays the missing mass squared $m^2_{miss}(\pi)$ spectrum for both components. The comparison reveals that restricting the selection to the main spill reduces the event yield uniformly across the spectrum. The reduction affects the signal region and the background regions, including the dominant $\pi^+ \to e^+\nu$ decay, in the same proportion.

Because the unstable early spill does not contain a unique background component that can be isolated from the signal, excluding it provides no improvement in signal purity while unnecessarily discarding valid data. The cut acts as a global efficiency loss rather than a discriminatory filter. Consequently, no further reduction based on the timestamp was applied, and the full spill (including the early component) was retained for the final analysis, consistent with the approach used in Run 1.

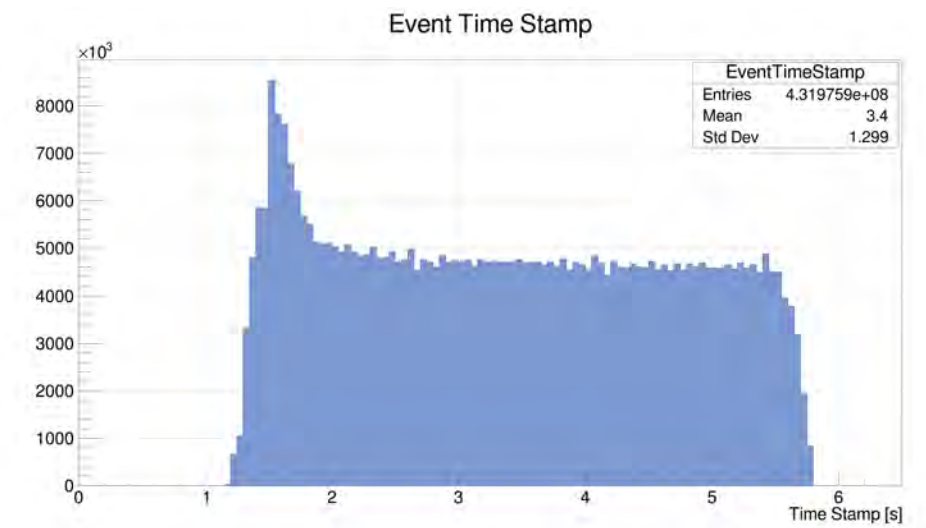

Figure A.6: Event time stamp for the data group 2021A

A.5 Evaluation of Data Groups 2016A and 2017D with Modified LKr Trigger Condition

The data groups collected in 2016A and 2017D differ from the rest of the Run 1 dataset solely in the LKr calorimeter trigger condition. While the standard $K^+ \to \pi^+\nu\bar{\nu}$ trigger line requires no energy deposit above 30 GeV ($\overline{\text{LKr30}}$), these specific periods operated with a stricter threshold of 20 GeV ($\overline{\text{LKr20}}$). This change impacts the acceptance for $\pi^+ \to e^+\nu$ decays, reducing it by approximately 50% compared to the standard configuration, resulting in an effective count of 1.38×10^{11} pion decays in the fiducial volume.

To evaluate the compatibility of these data with the HNL search, the positron momentum selection was restricted to the 5–20 GeV/c range, and MC simulations were emulated with the $\overline{\text{LKr20}}$ condition. Figure A.8 displays the missing mass squared (m^2_{miss}) spectrum for these periods, where the spectral shapes are consistent with the main dataset.

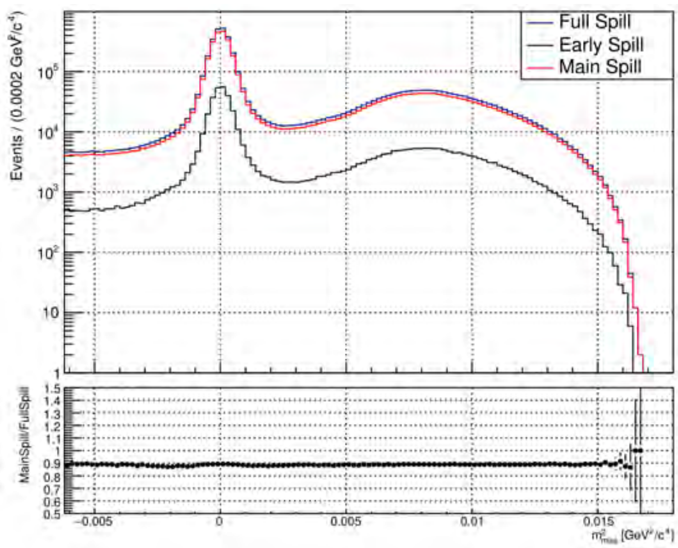

Figure A.7: $m^2_{miss}(\pi^+)$ distributions showing the components of the spill for 2021 data, in blue the full spill, in red the main spill and in black the early spill. The MainSpill/FullSpill ratio is also shown.

Despite this consistency, the 2016A and 2017D groups were excluded from the final combined dataset. The statistical gain from including them is minimal and insufficient to significantly improve the sensitivity for HNLs. Furthermore, their inclusion would require maintaining separate simulation samples with distinct trigger conditions, increasing complexity without commensurate benefit. Consequently, the analysis does not rely on this dataset.

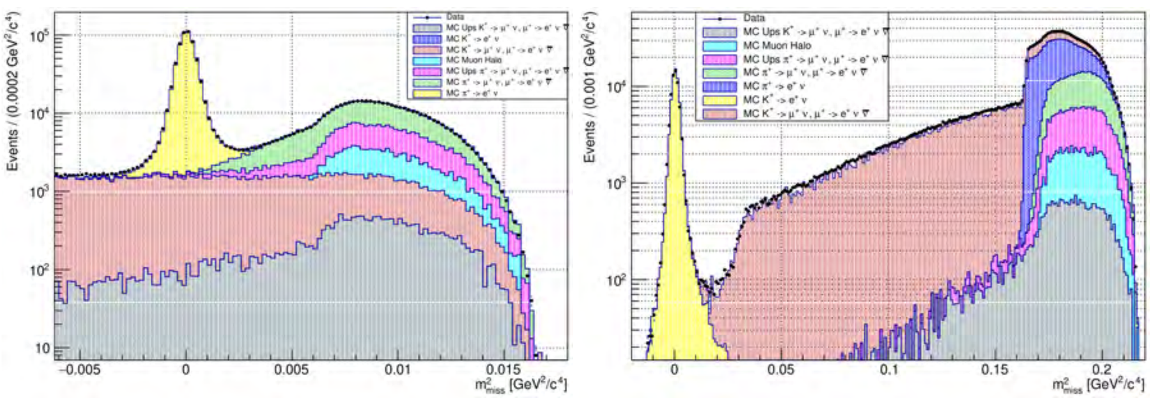

Figure A.8: m^2_{miss} spectrum for pion (left) and kaon (right) with the data sets 2016A and 2017D

Bibliography

[1] Patrick Janot **and** Stanisław Jadach. "Improved Bhabha cross section at LEP and the number of light neutrino species". **in** *Physics Letters B*: 803 (2020), **page** 135319. ISSN: 0370-2693. DOI: https://doi.org/10.1016/j.physletb.2020.135319. URL: https://www.sciencedirect.com/science/article/pii/S0370269320301234.

[2] Raymond Davis Jr, Don S Harmer **and** Kenneth C Hoffman. "Search for neutrinos from the sun". **in** *Physical Review Letters*: 20.21 (1968), **page** 1205.

[3] John N Bahcall **andothers**. "Standard solar models and the uncertainties in predicted capture rates of solar neutrinos". **in** *Reviews of Modern Physics*: 54.3 (1982), **page** 767.

[4] John N Bahcall, Neta A Bahcall **and** Giora Shaviv. "Present status of the theoretical predictions for the cl 37 solar-neutrino experiment". **in** *Physical Review Letters*: 20.21 (1968), **page** 1209.

[5] QR Ahmad **andothers**. "Measurement of the Rate of {nu}{sub e}+ d {yields} p+ p+ e {sup-} Interactions Produced by B". **in** *Physical Review Letters*: 87.7 (2001).

[6] Yoshiyuki Fukuda **andothers**. "Evidence for oscillation of atmospheric neutrinos". **in** *Physical review letters*: 81.8 (1998), **page** 1562.

[7] K Abe **andothers**. "Solar neutrino measurements using the full data period of Super-Kamiokande-IV". **in** *Physical review D*: 109.9 (2024), **page** 092001.

[8] M Abreu **andothers**. "Measurement of reactor antineutrino oscillation at SNO+". **in** *Physical review letters*: 135.12 (2025), **page** 121801.

[9] A. Allega **andothers**. "Initial measurement of reactor antineutrino oscillation at SNO+". **in** *The European Physical Journal C*: 85.1 (2025), **page** 17. DOI: 10.1140/epjc/s10052-024-13687-5. arXiv: 2405.19700 [hep-ex].

[10] F. Takahashi **andothers**. "Review of Particle Physics". **in** *Int. J. Mod. Phys. A*: 41 (2026), **page** 2630011. DOI: 10.1142/S0217751X26300115.

[11] KATRIN Collaboration† **andothers**. "Direct neutrino-mass measurement based on 259 days of KATRIN data". **in** *Science*: 388.6743 (2025), **pages** 180–185.

[12] Peter Minkowski. "$\mu \rightarrow$ eγ at a rate of one out of 109 muon decays?" **in** *Physics Letters B*: 67.4 (1977), **pages** 421–428.

[13] Takehiko Asaka **and** Mikhail Shaposhnikov. "The νMSM, dark matter and baryon asymmetry of the universe". **in** *Physics Letters B*: 620.1-2 (2005), **pages** 17–26.

[14] Hideki Yukawa. "On the interaction of elementary particles. I". **in** *Proceedings of the Physico-Mathematical Society of Japan. 3rd Series*: 17 (1935), **pages** 48–57.

[15] GD ROCHESTERDr **and** CC BUTLERDr. "Evidence for the existence of new unstable elementary particles". **in***Nature*: 160.4077 (1947), **pages** 855–857.

[16] Vincenzo Cirigliano **and** Ignasi Rosell. "π/K$\rightarrow$ e ν
e branching ratios to O (e 2 p 4) in chiral perturbation theory". **in***Journal of High Energy Physics*: 2007.10 (2007), **pages** 005–005.

[17] Douglas Bryman **andothers**. "Rare kaon and pion decays: Incisive probes for new physics beyond the standard model". **in***Annual Review of Nuclear and Particle Science*: 61.1 (2011), **pages** 331–354.

[18] Gerhard Czapek **andothers**. "Branching ratio for the rare pion decay into positron and neutrino". **in***Physical review letters*: 70.1 (1993), **page** 17.

[19] DI Britton **andothers**. "Measurement of the π+$\rightarrow$ e+ ν branching ratio". **in***Physical Review D*: 49.1 (1994), **page** 28.

[20] A Aguilar-Arevalo **andothers**. "Improved measurement of the $\pi\rightarrow$ e ν branching ratio". **in***Physical Review Letters*: 115.7 (2015), **page** 071801.

[21] Robert E Shrock. "New tests for and bounds on neutrino masses and lepton mixing". **in***Physics Letters B*: 96.1-2 (1980), **pages** 159–164.

[22] DI Britton **andothers**. "Improved search for massive neutrinos in π+$\rightarrow$ e+ ν decay". **in***Physical Review D*: 46.3 (1992), R885.

[23] A Aguilar-Arevalo **andothers**. "Improved search for heavy neutrinos in the decay $\pi\rightarrow$ e ν". **in***Physical Review D*: 97.7 (2018), **page** 072012.

[24] Eduardo Cortina Gil **andothers**. "Search for heavy neutral lepton production in K+ decays to positrons". **in***Phys. Lett. B*: 807 (2020), **page** 135599. DOI: 10.1016/j.physletb.2020.135599. arXiv: 2005.09575 [hep-ex].

[25] A Aguilar-Arevalo **andothers**. "Search for heavy neutrinos in $\pi\rightarrow\ \mu\nu$ decay". **in***Physics Letters B*: 798 (2019), **page** 134980.

[26] AV Artamonov **andothers**. "E949 Collaboration". **in***Phys. Rev. Lett*: 101 (2008), **page** 191802.

[27] E Cortina Gil **andothers**. "Search for K+ decays to a muon and invisible particles". **in***Physics Letters B*: 816 (2021), **page** 136259.

[28] Eduardo Cortina Gil **andothers**. "Observation of the $K^+ \rightarrow \pi^+\nu\overline{\nu}$ decay and measurement of its branching ratio". **in***JHEP*: 02 (2025), **page** 191. DOI: 10.1007/JHEP02(2025)191. arXiv: 2412.12015 [hep-ex].

[29] E Cortina Gil **andothers**. "A measurement of the $K^+ \rightarrow \pi^+\mu^+\mu^-$ decay". **in***Journal of High Energy Physics*: 2022.11 (2022), **page** 11.

[30] Eduardo Cortina Gil **andothers**. "Search for Lepton Number and Flavor Violation in K^+ and π^0 Decays". **in***Phys. Rev. Lett.*: 127.13 (2021), **page** 131802. DOI: 10.1103/PhysRevLett.127.131802. arXiv: 2105.06759 [hep-ex].

[31] Eduardo Cortina Gil **andothers**. "Search for Leptonic Decays of Dark Photons at NA62". **in***Phys. Rev. Lett.*: 133.11 (2024), **page** 111802. DOI: 10.1103/PhysRevLett.133.111802. arXiv: 2312.12055 [hep-ex].

[32] Eduardo Cortina Gil **andothers**. "First detection of a tagged neutrino in the NA62 experiment". **in***Phys. Lett. B*: 863 (2025), **page** 139345. DOI: 10.1016/j.physletb.2025.139345. arXiv: 2412.04033 [hep-ex].

[33] E Cortina Gil **andothers**. “The beam and detector of the NA62 experiment at CERN”. **in***Journal of Instrumentation*: 12.05 (2017), P05025–P05025.

[34] CERN. *The CERN accelerator complex, layout in 2022*. 2022. URL: https://cds.cern.ch/images/CERN-GRAPHICS-2022-001-1.

[35] Agni Bethani **andothers**. “Development of a new CEDAR for kaon identification at the NA62 experiment at CERN”. **in***Journal of Instrumentation*: 19.05 (2024), P05005.

[36] F Hahn **andothers**. “NA62: technical design document”. **in***NA62 Document*: (2010).

[37] Eduardo Cortina Gil **andothers**. “Performance of the NA62 trigger system”. **in***Journal of high energy physics*: 2023.3 (2023), **page** 122.

[38] Collaboration NA62. *2021 NA62 Status Report to the CERN SPSC*. techreport. 2021.

[39] John Allison **andothers**. “Recent developments in Geant4”. **in***Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment*: 835 (2016), **pages** 186–225.

[40] AL Read. *Optimal statistical analysis of search results based on the likelihood ratio and its application to the search for the MSM Higgs boson at sqrt (s)= 161 and 172 GeV*. techreport. 1997.

[41] Alexander L Read. *Modified frequentist analysis of search results (the $CL\_\{s\}$ method)*. techreport. CERN, 2000.

[42] Alexander L Read. “Presentation of search results: the CLs technique”. **in***Journal of Physics G: Nuclear and Particle Physics*: 28.10 (2002), **pages** 2693–2704.

[43] Evgueni Goudzovski. *Implementation of the CLs method for a single counting experiment*. techreport. 2019.

[44] Rene Brun **and** Fons Rademakers. “ROOT – An Object Oriented Data Analysis Framework”. **in***Nucl. Inst. & Meth. in Phys. Res. A*: 389 (1997), **pages** 81–86. DOI: 10.1016/S0168-9002(97)00048-X.