CERN-FASER-CONF-2026-005
25 March 2026

First Measurement of Muon Neutrinos as a Function of Energy and Rapidity with FASER

FASER Collaboration

We present the first double-differential measurement of the muon neutrino flux as functions of neutrino energy E and rapidity y, and neutrino–nucleon cross section as a function of E; using proton–proton collision data at $\sqrt{s} = 13.6$ TeV recorded during LHC Run 3, corresponding to an integrated luminosity of $(186 \pm 4)\,\mathrm{fb}^{-1}$. Charged-current muon neutrino interactions are identified using the active electronic components of the FASER detector. After subtracting the expected backgrounds, 766.8 ± 29.6 interaction events are observed. We unfold the neutrino events into a fiducial volume corresponding to the massive target region preceding the FASER spectrometer. These results provide the first double-differential characterization of TeV-scale collider neutrinos in energy and rapidity and contribute to a detailed experimental mapping of the forward neutrino flux at the LHC.

© 2026 CERN for the benefit of the FASER Collaboration. Reproduction of this article or parts of it is allowed as specified in the CC-BY-4.0 license.

CONTENTS

1. Introduction 1

2. FASER Detector 2

3. Dataset and simulated samples 3

4. Event Selection and background rejection 4

5. Background Estimation 4

6. Signal Extraction 7

7. Systematic Uncertainties 11

8. Results 12

9. Summary and conclusion 17

References 17

1. INTRODUCTION

The Large Hadron Collider (LHC) produces an intense and highly collimated beam of neutrinos in the forward direction, originating predominantly from decays of pions, kaons, and heavy hadrons produced in proton–proton (pp) collisions. The study of these collider neutrinos provides a unique

opportunity to probe neutrino interactions at energies between 360 GeV and 6.3 TeV, a regime that bridges the gap between fixed–target accelerator measurements [1] and observations from astroparticle experiments [2]. Measurements in this energy range provide new input to neutrino cross section modelling, perturbative and non–perturbative quantum chromodynamics, and the understanding of high–energy atmospheric and astrophysical neutrino fluxes [3, 4].

Although the possibility of studying neutrinos produced at colliders was proposed already in 1984 [5], their experimental detection became feasible only with the installation of dedicated forward detectors at the LHC. In 2023, the FASER Collaboration reported the first direct observation of collider neutrinos at the LHC [6], subsequently confirmed by the SND@LHC Collaboration [7]. This was followed by the first measurement of neutrino interactions cross sections at TeV energies using the FASER emulsion detector [8], and the first individual measurement of the muon neutrino interaction cross section and flux as a function of neutrino energy and the rapidity using the FASER electronic detector [9, 10].

In this note, we present the first double-differential measurement of muon neutrino interactions as functions of neutrino rapidity y and lepton-number–over–energy $-L/E_\nu$, using the electronic components of the FASER detector. The analysis is based on data collected during LHC Run 3 in 2022, 2023 and 2024, corresponding to a total integrated luminosity of $(186 \pm 4)\,\text{fb}^{-1}$. Charged–current ν_μ interactions are reconstructed via the identification of a muon track in the spectrometer, with the muon transverse position serving as a proxy for the neutrino rapidity. At detector level, the reconstructed ratio of muon charge over momentum q_μ/p_μ is used and unfolded to the neutrino observable $-L/E_\nu$. After subtracting contributions from backgrounds, we unfold the neutrino events into a fiducial volume corresponding to the target regions of the FASER detector. The fiducial volume for this analysis is defined as a cylindrical region within the FASERν target, 100 mm in radius around the central axis of the tracking spectrometer.

2. FASER DETECTOR

The FASER detector, described in detail in Ref. [11], is located in the TI12 tunnel approximately 480 m downstream of the ATLAS interaction point (IP1) and about 5 m from the LHC proton beam line.

The FASER detector consists of several subsystems arranged approximately along the collision axis line-of-sight (LOS) from IP1. The FASER detector uses a Cartesian coordinate system with the z-axis pointing along the LOS away from IP1, the y-axis pointing vertically upwards, and the x-axis pointing horizontally towards the LHC machine. From upstream to downstream these include the VetoNu scintillator station, the FASERν neutrino detector, the interface tracker, a scintillator veto station, the timing scintillator station, the tracking spectrometer, the preshower scintillator system, and the electromagnetic calorimeter. Signals from the scintillator stations and the calorimeter provide the trigger for the detector. The FASERν detector is composed of 730 layers of tungsten plates interleaved with emulsion films and has a mass of approximately 1.1 tons. In this analysis, it serves as the target material for neutrino interactions, while the electronic detector components downstream are used to reconstruct the outgoing muon. During part of the 2024 data–taking period the FASERν detector was temporarily replaced by the *CaloNu* detector. This was a prototype neutrino detector consisting of layers of tungsten plates placed between four spare LHCb calorimeter modules [12]. Although the electronic readout of this detector is not used in this analysis, it acts as a passive target material. The total mass of this prototype detector was 0.68 metric tons, and the change in mass, material and mass distribution is accounted for in the detector simulation. For brevity, unless stated otherwise, we will use FASERν to refer to all data taking periods.

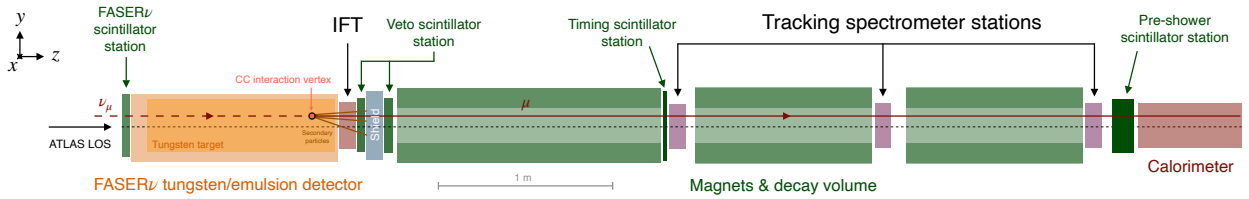


FIG. 1. Side view diagram of FASER with a charged–current muon neutrino interacting within the tungsten target.

Muon tracks are reconstructed using the tracking spectrometer, which consists of three tracking stations separated by three dipole magnets providing a magnetic field of 0.57 T [13]. Each tracking station contains three planes of eight double–sided ATLAS Semiconductor Tracker modules [14]. The magnetic field allows the measurement of the charge and momentum of charged particles traversing the detector and consequently the separation of neutrino and antineutrino events. The transverse acceptance of the spectrometer is defined by the 200 mm diameter magnet aperture. A schematic diagram a signal event interacting in the FASER detector is shown in Fig. 1.

3. DATASET AND SIMULATED SAMPLES

This analysis uses data collected during stable beam conditions in 2022, 2023, and 2024. The dataset corresponds to a total integrated luminosity of $(186 \pm 4)\,\text{fb}^{-1}$. The neutrino flux at FASER is simulated using the fast forward neutrino flux sim, which takes propagates a hadrons through the LHC and decays them to result in a flux of neutrinos [15], adjusted to the LHC Run 3 configuration [16]. Light hadron production in pp collisions is modelled using `EPOS-LHC` [17] as the baseline, with uncertainties estimated from an envelope of alternative hadronic interaction models including `SIBYLL 2.3d` and `QGSJET 2.04` [18–20]. Charm hadron production is modelled using `POWHEG` interfaced with `PYTHIA 8.3`, with uncertainties estimated by varying the factorization and renormalization scales [21–25].

The exact position of the LOS in the detector depends on the crossing angle of the the proton beams in IP1, which varies over time. To account for time–dependent variations of the LHC beam optics, the distribution of the half–crossing angle is taken directly from data provided by the LHC, and the simulated neutrino flux is shifted accordingly to reproduce the corresponding line-of-sight position at the FASER location. During the 2022 and 2023 data taking, the half-crossing angle varied between $-165\,\mu\text{rad}$ and $-135\,\mu\text{rad}$. During the 2024 data taking, the half-crossing angle varied between $150\,\mu\text{rad}$ and $160\,\mu\text{rad}$. This results in a vertical displacement of the LOS relative to the detector axis. To improve alignment with the LOS, the detector was mechanically lowered by $(12.3 \pm 0.2)\,\text{mm}$. This is also accounted for in the flux modelling.

Neutrino interactions in the detector are simulated using `GENIE 3.04.0` [26–28], with cross sections modelled using the Bodek–Yang approach [29–31]. The agreement of this model with alternative calculations in the relevant energy range is at the level of approximately 6 % [32].

Separate Monte Carlo (MC) samples are produced for the nominal FASERν geometry and for the CaloNu configuration used during part of 2024 data taking. The simulated samples correspond to an integrated luminosity of $10\,\text{ab}^{-1}$ for each production mode. Secondary particle propagation and detector response are simulated using `Geant4` [33] with the `FTFP_BERT` physics list [34, 35]. Alternative physics lists within `GEANT4` are used to assess uncertainties related to intranuclear cascade modelling and neutron transport.

Simulated events are reconstructed using the FASER offline software framework, `CALYPSO`, based

on Athena from the ATLAS Collaboration [36], with track reconstruction performed using the ACTS framework [37]. The same reconstruction and event selection criteria are applied to recorded and simulated events.

4. EVENT SELECTION AND BACKGROUND REJECTION

This analysis selects charged–current (CC) muon neutrino interactions occurring in the FASERν tungsten target that produce a muon reconstructed in the FASER spectrometer. The selection criteria closely follow those used in the previous measurement of muon neutrino interactions as functions of energy and rapidity [9, 10].

To suppress background originating from muons entering the detector, events are required to have no or only a small signal in the front scintillator station. Neutrino interactions may nevertheless produce activity in this detector due to low energy neutrons released in by the de–excitation of the tungsten nuclei, this reduces the signal efficiency. To account for this we therefore apply the reduced VetoNu charge method described in Ref. [9]. This method exploits the timing difference between the upstream front VetoNu station and the downstream Veto station to ensure that charge deposits from neutrino–induced secondaries are not misidentified as incoming muons. It is required that the reduced VetoNu charge must be less than 30 pC. The selection further requires the measured signals in the downstream scintillator stations to be consistent with a minimum–ionizing particle traversing the detector. Specifically, the signal in the downstream Veto station must exceed 40 pC, the Timing station signal must be greater than 20 pC (with corrections applied for module gaps), and the Preshower scintillator signal must exceed 2.5 pC to ensure reliable muon trajectory reconstruction.

Muon tracks reconstructed in the spectrometer are required to satisfy several quality criteria. Tracks must contain at least 14 hits distributed across at least 7 silicon detector planes, have a good fit quality, and a momentum greater than 100 GeV to suppress background contamination. In addition, the track radius must be smaller than 95 mm in all spectrometer tracking stations, the extrapolated radius must be below 95 mm at the Veto station and below 120 mm at the FASERν front scintillator station, and the track angle with respect to the detector axis, θ, must be less than 25 mrad. The track reconstruction is seeded from the most downstream station, improving efficiency for neutrino interactions that produce a large number of hits in the upstream tracking stations.

Additional requirements suppress beam–related backgrounds and cosmic muons. Events must satisfy a valid trigger from the Veto, Timing, or Preshower scintillator systems and must originate from a colliding bunch crossing at the ATLAS interaction point. These criteria yield a high–purity sample of neutrino–induced muon tracks while minimizing contamination from background sources.

5. BACKGROUND ESTIMATION

Most particles produced at the interaction point are either deflected by the LHC magnets, or absorbed by the roughly 100 m of rock and concrete before reaching the FASER detector, resulting in a very low background environment. However, several sources of background can mimic the signal topology of a muon produced in a neutrino interaction within the target. In particular, high–energy muons originating upstream of the detector may evade the front veto system and produce a reconstructed track in the spectrometer. Additional contributions arise from non–signal neutrino interactions. We consider four classes of background: large–angle muons which miss VetoNu, neutral hadron backgrounds, neutrino–induced backgrounds, and beam induced and cosmic backgrounds.

TABLE 1. Summary of the expected geometric background yields for each analysis period. The 2022-23 and 2024 estimates are computed separately due to the shift in LHC optics and the corresponding change in muon flux.

Period	Yield
2022/2023 Pos. tracks	$0.000^{+0.013}_{-0.000}$
2022/2023 Neg. tracks	$0.024^{+0.011}_{-0.011}$
2024 Pos. tracks	$0.002^{+0.01}_{-0.002}$
2024 Neg. tracks	$0.008^{+0.11}_{-0.008}$

Most incoming muons are rejected by the front scintillator station. The inefficiency of each of its two scintillator layers has been measured to be below 10^{-7} using events that otherwise satisfy the event selection but have a reduced Veto charge above 30 pC in one layer. Assuming the two layers operate independently, the number of muons which pass undetected through both is of order 10^{-6}, which is negligible compared to other background contributions.

A potential background arises from large–angle muons that miss the front scintillator station but scatter within the tungsten of FASERν and enter the fiducial volume of the analysis. This contribution is referred to as the *geometric muon background*. To estimate it, we determine the probability for muons to scatter by sufficiently large angles in the tungsten target. A stand–alone GEANT4 simulation is used in which 10^9 muons, sampled from a uniform energy distribution, are propagated head–on through a 1 m block of tungsten, serving as a proxy for the FASERν target. From the simulated outgoing muons we extract the probability for a muon of a given energy to scatter by a given angle.

These scattering probabilities are applied to muons observed in a high–angle sideband. The geometric muon sideband is defined by inverting the track angle and radius requirements of the signal region, requiring $\theta > 25$ mrad and the extrapolated track radius at VetoNu to be greater than 120 mm. The radius requirement is inverted to increase the statistics of high–angle, high–momentum muons. For each bin of p_μ and θ in this sideband we determine the additional scattering angle required for the muon to enter the signal region, $\Delta\theta = \theta - 25$ mrad, and weight the observed events by the probability for a muon to scatter by at least this angle. Since the expected background contribution is very small, the resulting distribution is projected onto the p_μ axis to obtain the background shape.

The overall normalization of the geometric muon background is obtained from a simultaneous fit of this background component and the predicted sum of neutrino interactions from GENIE to data in a low–momentum sideband defined by inverting the requirement $p_\mu > 100$ GeV. The fit is performed separately for each muon charge and data–taking period. In general, the fits provide a good description of the data, with χ^2/ndof values between 0.33 and 1.25.

To estimate a systematic uncertainty on this background, the procedure is repeated using an alternative muon flux in which the geometric sideband is defined without requiring the absence of hits in the front scintillator station. This approach yields a poorer description of the data and overestimates the background, and is therefore taken as a conservative uncertainty estimate. As the geometric background is small and contributes primarily at low momentum, it is included in the two–dimensional fit only in lowest momentum q/p_μ bins, and is conservatively assumed to be independent of y. The expected geometric background yields for each muon charge and data–taking period are summarized in Table 1.

Another potential background arises when upstream muons interact with material before reaching the detector and produce secondary neutral hadrons that subsequently enter the detector volume. These neutral particles may occasionally generate tracks that resemble signal events. This

contribution is evaluated using simulated samples corresponding to approximately $2\,\mathrm{ab}^{-1}$ of integrated luminosity within a radial region of 150 mm. In most cases the parent muon is identified by the front scintillator station, or the neutral hadron is absorbed within the FASERν or CaloNu detector target without producing a reconstructable track. Consequently, the expected number of such background events passing the full event selection is below 10^{-3} and is therefore considered negligible.

Neutrino–induced backgrounds include ν_μ CC interactions occurring outside the fiducial target region, as well as charged–current interactions of electron and tau neutrinos and neutral–current interactions of all neutrino flavours. The latter are highly suppressed due to the high density of the tungsten target, which absorbs the scattering products before they reach the spectrometer. These contributions are evaluated using simulated samples generated with `GENIE` and are subtracted from the observed event yields. Possible backgrounds from cosmic rays and beam–related activity are studied using events recorded during non–colliding bunch crossings and are found to be negligible.

The expected background contributions together with the number of selected events are summarized in Tables 2 and 3.

TABLE 2. Number of simulated signal, background, and observed events for each reconstructed rapidity bin. The uncertainty on the simulated number of observed events includes statistical and systematic contributions. The geometric background is extremely small and so is omitted from the table, however it it still included for the statistical analysis.

Rapidity bin		[8.5, 9.0]	[9.0, 9.25]	[9.25, 9.5]	[9.5, 10]	[10, ∞)	Total
Simulation							
Signal	ν_μ CC (fid.)	149.0 ± 21.0	130.0 ± 19.0	109.0 ± 15.0	145.0 ± 18.0	129.0 ± 19.0	662.0 ± 84.6
Background	ν_μ CC (non-fid.)	11.2 ± 1.7	6.2 ± 1.2	5.2 ± 1.1	7.5 ± 1.5	6.6 ± 1.3	36.2 ± 6.1
	ν_e CC	1.4 ± 0.7	1.1 ± 0.7	1.0 ± 0.7	1.3 ± 0.6	1.3 ± 0.7	6.2 ± 3.2
	ν_τ CC	0.20 ± 0.21	0.18 ± 0.16	0.07 ± 0.18	0.12 ± 0.12	0.14 ± 0.13	0.7 ± 0.7
	ν NC	2.4 ± 0.7	2.2 ± 0.7	1.8 ± 0.5	2.5 ± 0.5	2.0 ± 0.5	11.1 ± 2.3
Total		163.8 ± 26.1	139.7 ± 22.1	117.5 ± 17.6	156.3 ± 21.2	139.4 ± 21.7	716.7 ± 84.9
Data							
Total		192.0 ± 13.9	157.0 ± 12.5	137.0 ± 11.7	185.0 ± 13.6	150.0 ± 12.2	821.0 ± 28.7

TABLE 3. Number of simulated signal, background, and observed events for each reconstructed muon q_μ/p'_μ bin. The uncertainty on the simulated number of observed events includes statistical and systematic contributions.

q_μ/p'_μ bin		$\left[-\frac{1}{100}, -\frac{1}{300}\right]$	$\left[-\frac{1}{300}, -\frac{1}{600}\right]$	$\left[-\frac{1}{600}, -\frac{1}{1000}\right]$	$\left[-\frac{1}{1000}, \frac{1}{1000}\right]$	$\left[\frac{1}{1000}, \frac{1}{300}\right]$	$\left[\frac{1}{300}, \frac{1}{100}\right]$	Total
Simulation								
Signal	ν_μ CC (fid.)	78.0 ± 18	138.0 ± 27	109.0 ± 17	180.0 ± 40	108.0 ± 24	44.0 ± 13	662.0 ± 85.1
Background	ν_μ CC (non-fid.)	7.8 ± 1.3	7.6 ± 1.6	4.3 ± 1.0	4.5 ± 1.4	6.0 ± 1.3	6.6 ± 1.7	36.2 ± 6.5
	ν_e CC	1.7 ± 0.8	0.6 ± 0.4	0.16 ± 0.10	0.19 ± 0.18	1.4 ± 0.6	2.2 ± 1.4	6.2 ± 3.2
	ν_τ CC	0.17 ± 0.19	0.10 ± 0.20	0.04 ± 0.09	0.09 ± 0.08	0.12 ± 0.22	0.19 ± 0.13	0.7 ± 0.7
	ν NC	3.6 ± 0.8	1.3 ± 0.4	0.47 ± 0.20	0.26 ± 0.22	1.9 ± 0.6	3.6 ± 0.8	11.1 ± 2.4
Total		91.6 ± 20.2	147.5 ± 29.2	114.3 ± 18.5	189.3 ± 43.1	117.5 ± 26.0	56.5 ± 14.8	716.7 ± 85.5
Data								
Total		113 ± 10.6	216 ± 14.7	167 ± 12.9	146 ± 12.1	118 ± 10.9	61 ± 7.8	821 ± 28.7

6. SIGNAL EXTRACTION

We apply a linear scaling to the reconstructed muon momentum, defining a corrected momentum $p'_\mu = p_\mu/0.8$. The factor of 0.8 reflects that muons passing the event selection retain on average about 80 % of the incident neutrino energy. At the detector level the observable used in the analysis is the reconstructed muon charge-over-momentum, q_μ/p'_μ. At the truth level the corresponding neutrino quantity used in the unfolding is the lepton-number–over–energy $-L/E_\nu$, where $L = +1$ (-1) for ν_μ $(\bar{\nu}_\mu)$. The migration matrix therefore relates the reconstructed muon observable q_μ/p'_μ to the neutrino observable $-L/E_\nu$.

We define six bins in the ratio of the muon charge and corrected momentum, q_μ/p'_μ, with bin edges

$$\left[-\frac{1}{100}, -\frac{1}{300}, -\frac{1}{600}, -\frac{1}{1000}, \frac{1}{1000}, \frac{1}{300}, \frac{1}{100}\right] \text{GeV}^{-1}. \tag{1}$$

This binning yields a comparable expected number of events in each bin and separates muons originating from neutrinos and antineutrinos up to a momentum of 1 TeV. Above this threshold the charge identification becomes more challenging, and muons and anti–muons with higher momentum are therefore combined into a single bin.

The rapidity y of the neutrino, which is identical to the pseudorapidity for massless particles, is defined in terms of the polar angle as

$$y = -\ln\tan\left(\frac{\theta}{2}\right). \tag{2}$$

The polar angle θ is defined as the angle between the neutrino momentum direction and the line of sight, i.e. the direction connecting the ATLAS interaction point (IP1) to the centre of the FASER detector. Since the neutrino is not directly observed, its direction is approximated by the line connecting the ATLAS interaction point to the neutrino interaction vertex within the detector. The position of the interaction vertex is estimated by extrapolating the reconstructed muon track back to the centre of the FASERν detector. The polar angle θ is then computed as the angle between the vectors pointing from the interaction point to the reconstructed vertex and from the interaction point to the detector center, and the neutrino rapidity is inferred accordingly. The rapidity is analysed in five bins with edges defined as

$$\Big[8.5, 9, 9.25, 9.5, 10, \infty\Big). \tag{3}$$

The annular rapidity regions with respect to the centre of the detector for the data taking periods of 2022-2023 and 2024 are shown in Fig. 2, along with the average LOS position during the corresponding running periods. The change in the LOS over time is accounted for when evaluating the rapidity.

The number of reconstructed neutrino interactions in the analysis bins is unfolded to the number of true charged–current muon neutrino interactions in the fiducial volume through a binned likelihood fit that incorporates detector response, background estimates and systematic uncertainties. The expected number of reconstructed muon events in bin k is given by

$$n_{\mu,k} = \sum_l M_{kl} \cdot \epsilon_l \cdot n_{\nu,l} + n_{\text{bkg},k}, \tag{4}$$

where M_{kl} is the migration matrix relating truth bin l to reconstructed bin k, ϵ_l is the total

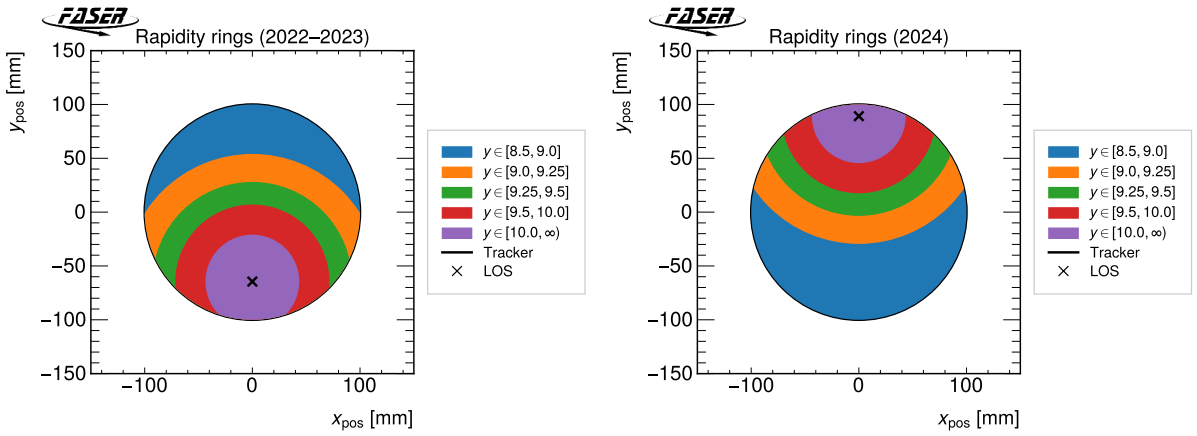


FIG. 2. Rapidity bin rings projected into the transverse (x, y) plane at the centre of the FASERν detector for 2022–2023 (left) and 2024 (right) data–taking periods. Each rapidity bin corresponds to an annular region centred on the line-of-sight position. The position of the LOS changed between 2023 and 2024 due to the change in crossing angle at IP1.

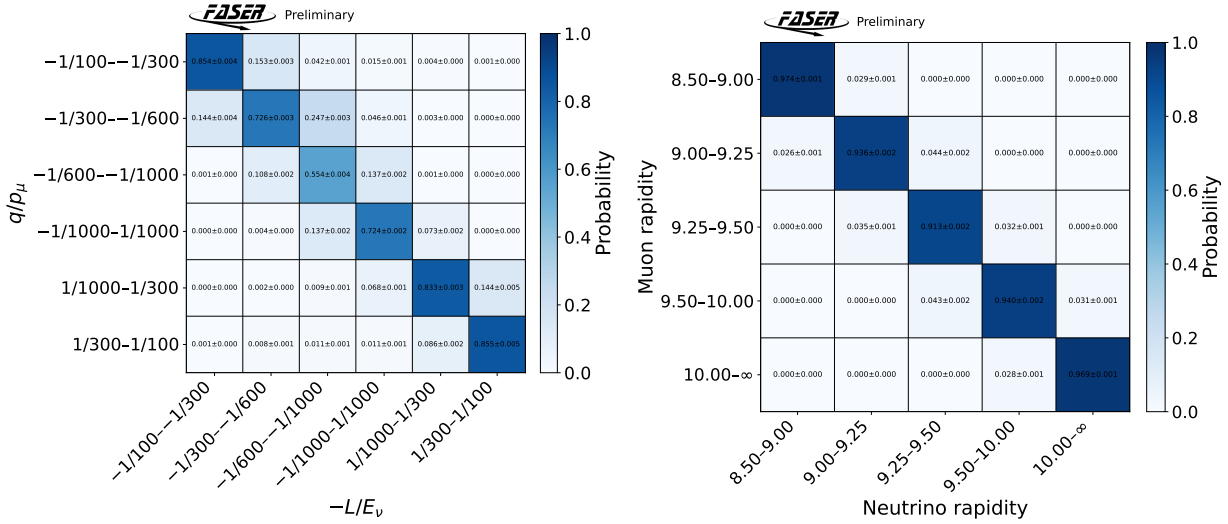


FIG. 3. Migration matrices used in the unfolding procedure for the one-dimensional measurements. Left: migration matrix mapping reconstructed muon q_μ/p'_μ to the true neutrino observable $-L/E_\nu$. Right: migration matrix mapping reconstructed muon rapidity to the true neutrino rapidity.

efficiency for a true signal event in truth bin l to be reconstructed and selected, and $n_{\mathrm{bkg},k}$ is the expected background contribution in reconstructed bin k from geometric muons and non-signal neutrino interactions, including charged-current interactions occurring outside the fiducial volume.

For the one–dimensional measurements in $-L/E_\nu$ and y, the indices k and l label bins of the corresponding reconstructed and truth observables. In the $-L/E_\nu$ measurement, the reconstructed bins correspond to the muon observable $q_\mu/p'\mu$, while the truth bins correspond to the neutrino observable $-L/E\nu$. In the rapidity measurement, the reconstructed and truth bins correspond to the muon and neutrino rapidities, respectively.

For the two–dimensional measurement, the same formalism is used, with the indices k and l labelling two–dimensional bins. In this case, reconstructed bins correspond to $(y, q_\mu/p'\mu)$, while truth bins correspond to $(y, -L/E\nu)$. The migration matrices used in the analysis are shown in

Figs. 3 and 4.

The likelihood function used in the fit is defined as

$$\mathcal{L} = \prod_k \mathcal{P}\left(N_k \mid n_{\mu,k}\right) \cdot \prod_j \mathcal{G}j, \tag{5}$$

where the Poisson terms $\mathcal{P}$ describe the probability of observing N_k events in bin k given the expected yield $n\mu, k$, and the Gaussian terms $\mathcal{G}_j$ constrain nuisance parameters associated with backgrounds and systematic uncertainties.

The likelihood fit is implemented using `iminuit` package [38]. The best–fit values of the neutrino interaction yields $n_{\nu,l}$ are obtained by minimizing $-2\log\mathcal{L}$, and uncertainties are extracted using the profile likelihood method.

The neutrino flux in each analysis bin l is determined from the fitted number of neutrino interactions as

$$\Phi_l = \frac{n_{\nu,l}}{\sigma_{\text{eff},l} \cdot A_{\text{eff},l} \cdot n_t \cdot \int \mathcal{L}, dt}, \tag{6}$$

where $\sigma_{\text{eff},l}$ is the effective neutrino cross section averaged over the neutrino energy spectrum in bin l, $A_{\text{eff},l}$ is the effective fiducial area for that bin, n_t is the target nucleon density, and $\int \mathcal{L}dt$ is the integrated luminosity.

This expression applies to both the one–dimensional and two–dimensional measurements. In the two–dimensional case, the index l labels bins in $(y, -L/E_\nu)$, and the effective cross section and fiducial area are evaluated separately for each bin. The same procedure is used to estimate $\sigma_{\text{eff},l}$ and propagate its uncertainties.

The interpretation of the number of neutrino interactions as a cross-section is done by assuming a simulated model for the neutrino flux $\Phi_{\text{sim},l}$. The equation for the cross-section is given by

$$\sigma_l \;=\; \frac{n_{\nu,l}}{\Phi_{\text{sim},l} \cdot \; A_{\text{eff},l} \cdot \; n_t \cdot \; \int \mathcal{L}\, dt}. \tag{7}$$

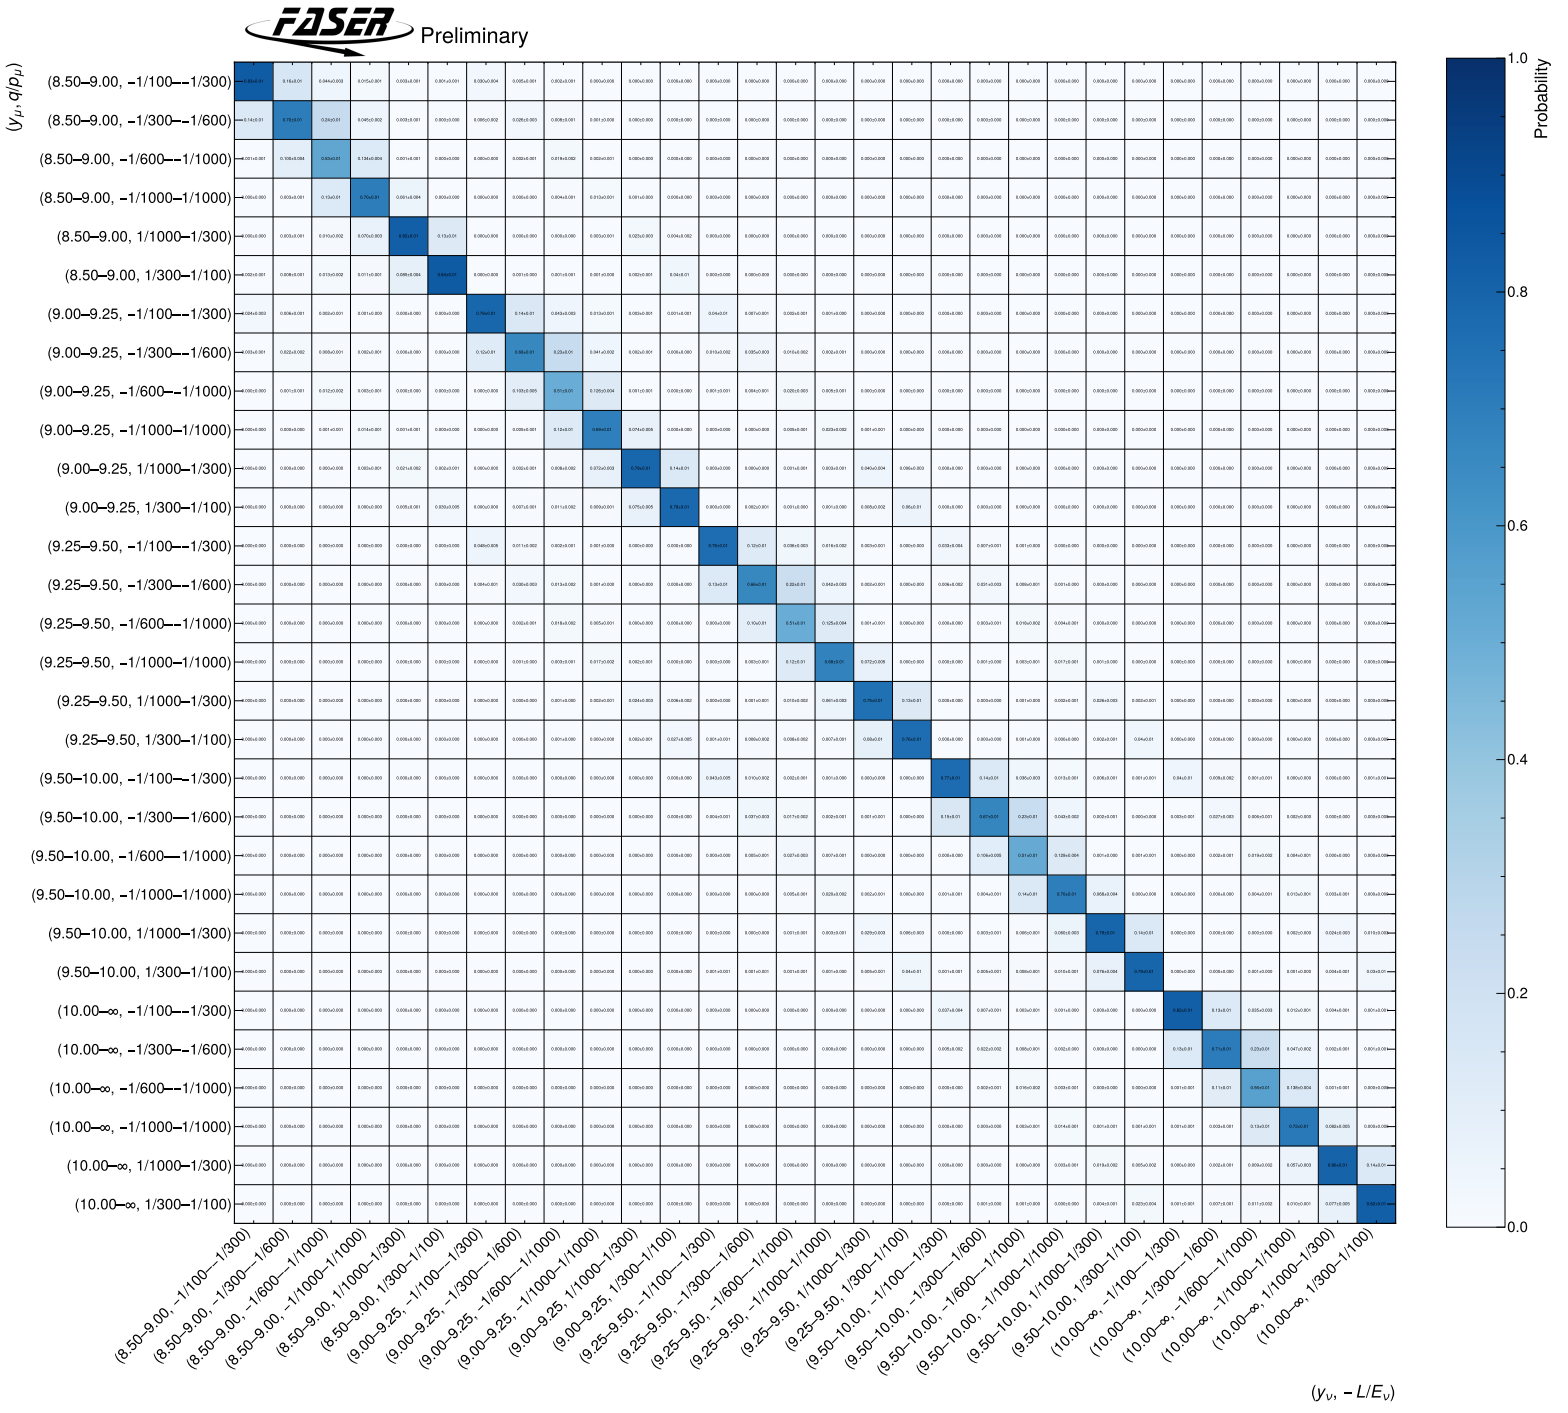


FIG. 4. Two-dimensional migration matrix mapping reconstructed $(y, q_\mu/p'_\mu)$ to true $(y, -L/E_\nu)$, used in the double-differential unfolding.

7. SYSTEMATIC UNCERTAINTIES

We evaluate uncertainties on the efficiency, the migration matrix and the predicted number of neutrino interactions. Systematic uncertainties affecting the efficiency and migration matrix are incorporated into the likelihood fit as nuisance parameters. The uncertainties on the number of expected neutrino interactions are not used in the fit but only for comparison with the data.

Flux modeling uncertainties arise from differences in neutrino flux predictions obtained from various forward flux MC generators [16]. We use different MC generators to simulate neutrino interactions originating from light hadrons decays (pions and kaons) and charm decays. The uncertainties for these two sources are evaluated separately and combined in quadrature. For light hadron-induced neutrinos, we take the prediction from `EPOS-LHC` as the nominal value and estimate the uncertainty as the maximum deviation between `EPOS-LHC` and other models: `SIBYLL`, `QGSJET`, and `Pythia` [39] in each rapidity bin. For charm-induced neutrinos, we use the prediction from `POWHEG+Pythia8` as the nominal value and define the uncertainty by varying the renormalization and factorization scales by a factor of two. The flux modelling uncertainty also affects the number of observed events outside the fiducial volume.

Uncertainties from `GENIE` and `Geant4` originate from differences in neutrino interaction modelling and detector simulations. We take the prediction from `GENIE G18_02a` tune, combined with the `Geant4 FTFP_BERT` physics list and the `G4NDL` neutron transport model, as the nominal value. The uncertainty is estimated as the maximum deviation between this nominal prediction and any of the considered alternative `GENIE` or `Geant4` tunes in each analysis bin. Systematic uncertainties from detector misalignment, which can induce q_μ/p'_μ migration and charge misidentification, are evaluated by reconstructing the nominal `GENIE` samples with modified alignment parameters. The detector geometry is shifted according to the unbiased residuals measured in data, applying correlated module displacements corresponding to $\pm 1\sigma$ of the Gaussian width of the residual distributions. The resulting shifts, typically between -5 and $5, \mu$m, together with the associated z-axis rotations, are propagated through the full reconstruction chain. The alignment uncertainty is defined as the largest deviation between the predictions obtained with the shifted geometries and the nominal prediction. The luminosity uncertainty is determined from the ATLAS luminosity measurement. The most recent luminosity determination reports a relative uncertainty of 2.2% [40–42]. Since the neutrino flux prediction scales linearly with luminosity, this uncertainty propagates directly to the expected number of neutrino interactions.

The mass of the neutrino target affects the expected number of neutrino interactions. A 1% relative mass uncertainty is assigned to the 2022–2023 data and the 2024 period without CaloNu, while a 5% uncertainty is assigned to the 2024 CaloNu period. Since the final analysis combines events from all periods, the mass uncertainty is computed as a luminosity–weighted combination of the period–dependent uncertainties.

An additional systematic uncertainty is assigned to account for possible mismodelling of the event selection efficiency between data and simulation. This uncertainty is evaluated by varying the selection criteria described in Sec. 4 within reasonable ranges motivated by data–simulation comparisons and detector performance. The resulting variations in the signal efficiency are propagated to the final observables, and the combined effect of all selection variations is treated as a single "event selection" systematic uncertainty, correlated across analysis bins.

Another potential source of uncertainty can occur due to the early triggering of the detector which can occur when one of the Veto scintillators becomes saturated from particle showers of close-by neutrino interactions. To estimate this uncertainty we multiply the simulated neutrino events by the probability that the event triggered early. The early trigger probabilities are determined from data and depend on the Veto station charge. Separate uncertainties are derived for Veto station 1 and station 2 and are found to be $\lesssim 1\%$ across all bins of q_μ/p'_μ and y. The resulting systematic un-

certainties on the efficiency and on the number of neutrino interactions are summarized in Tables 4 to 7. The same sources of systematic uncertainty are included in the two-dimensional measurement, relating reconstructed $(y, q_\mu/p'_\mu)$ bins to truth $(y, -L/E_\nu)$ bins, and are treated identically within the likelihood fit.

TABLE 4. Nominal signal efficiency in each rapidity bin and its associated absolute uncertainties, expressed in percentage points. The total uncertainty is obtained by adding the individual contributions in quadrature.

	$8.5 < y < 9.0$	$9.0 < y < 9.25$	$9.25 < y < 9.5$	$9.5 < y < 10.0$	$10.0 < y < \infty$
Nominal efficiency [%]	21.5	26.4	27.8	27.5	26.1
Event selection	1.5	1.6	1.7	1.5	1.5
Flux modeling	0.6	0.8	0.8	1.1	0.8
Geant4×GENIE tunes	0.4	0.2	0.4	0.6	0.3
Early trigger	0.2	0.3	0.3	0.3	0.3
Statistical MC	0.2	0.2	0.2	0.2	0.2
Alignment	< 0.1	< 0.1	< 0.1	< 0.1	< 0.1
Total uncertainty	1.7	1.8	1.9	2.0	1.7

TABLE 5. Nominal number of neutrino interactions in each rapidity bin and the corresponding absolute uncertainties, expressed in number of events. The total uncertainty is obtained by adding the individual contributions in quadrature.

	$8.5 < y < 9.0$	$9.0 < y < 9.25$	$9.25 < y < 9.5$	$9.5 < y < 10.0$	$10.0 < y < \infty$
Nominal [events]	690.4	489.1	394.4	526.2	496.7
Flux modelling	96.4	57.8	46.6	60.5	76.5
Luminosity	14.4	10.3	8.3	11.1	10.6
Mass	13.3	9.5	7.6	10.2	9.6
Geant4×GENIE tunes	9.4	4.7	5.4	5.8	8.7
Statistical MC	2.1	1.8	1.6	1.8	1.8
Total uncertainty	98.8	59.7	48.2	62.6	78.3

TABLE 6. Nominal signal efficiency in each $-L/E_\nu$ bin and its associated absolute uncertainties, expressed in percentage points. The total uncertainty is obtained by adding the individual contributions in quadrature.

	$[-\frac{1}{100}, -\frac{1}{300}]$	$[-\frac{1}{300}, -\frac{1}{600}]$	$[-\frac{1}{600}, -\frac{1}{1000}]$	$[-\frac{1}{1000}, \frac{1}{1000}]$	$[\frac{1}{1000}, \frac{1}{300}]$	$[\frac{1}{300}, \frac{1}{100}]$
Nominal efficiency [%]	13.0	22.6	27.6	30.5	36.6	22.2
Event selection	1.0	1.3	1.6	2.0	2.1	1.8
Flux modelling	0.3	0.4	0.5	0.9	0.4	0.5
Geant4×GENIE tunes	0.1	0.2	0.5	0.7	0.5	0.3
Early trigger	0.1	0.2	0.3	0.4	0.3	0.3
Statistical MC	0.1	0.2	0.2	0.2	0.3	0.3
Alignment	< 0.1	< 0.1	< 0.1	0.1	< 0.1	< 0.1
Total uncertainty	1.1	1.4	1.8	2.3	2.3	1.9

8. RESULTS

The unfolded number of muon neutrino interactions is shown in Figs. 5 and 6 as a function of the neutrino observables $-L/E_\nu$ and rapidity y. The results are compared to the simulation predictions from the nominal neutrino flux model. The extracted number of neutrino interactions in each bin is obtained from the likelihood fit described in Sec. 6.

TABLE 7. Nominal number of neutrino interactions in each $-L/E_\nu$ bin and the corresponding absolute uncertainties, expressed in number of events. The total uncertainty is obtained by adding the individual contributions in quadrature.

	$[-\frac{1}{100}, -\frac{1}{300}]$	$[-\frac{1}{300}, -\frac{1}{600}]$	$[-\frac{1}{600}, -\frac{1}{1000}]$	$[-\frac{1}{1000}, \frac{1}{1000}]$	$[\frac{1}{1000}, \frac{1}{300}]$	$[\frac{1}{300}, \frac{1}{100}]$
Nominal [events]	456.7	544.6	428.1	723.5	283.9	159.4
Flux modelling	117.2	126.8	63.6	171.7	63.6	51.8
Luminosity	9.7	11.6	9.1	15.4	6.1	3.4
Mass	8.8	10.5	8.3	14.0	5.5	3.1
Geant4×GENIE tunes	6.9	6.6	5.9	10.0	3.5	0.9
Statistical MC	1.7	1.8	1.6	2.1	1.3	1.0
Total uncertainty	118.1	127.9	65.1	173.3	64.2	52.0

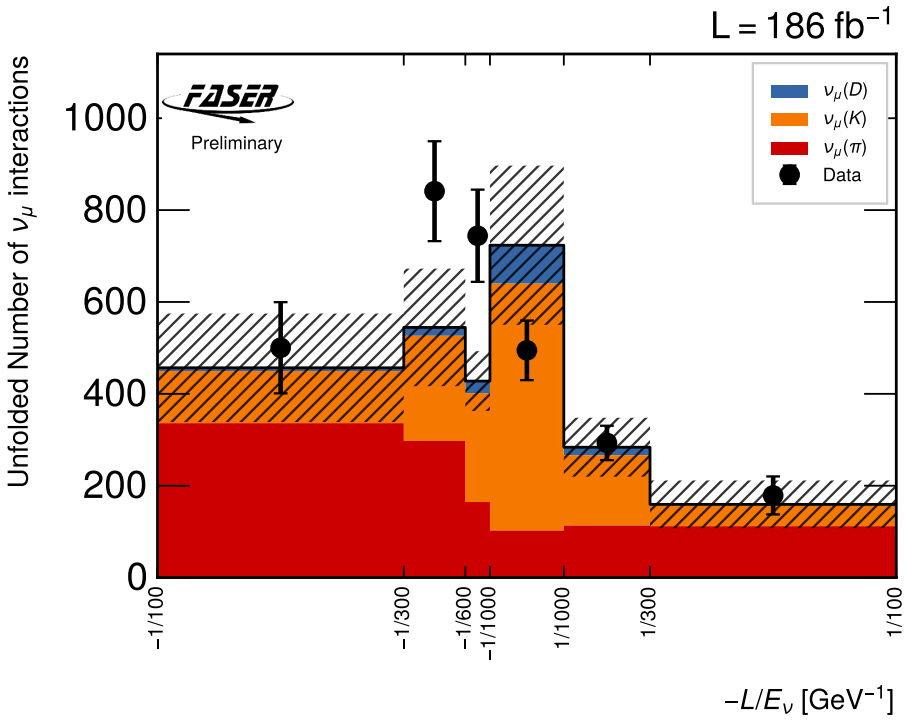


FIG. 5. Unfolded number of muon neutrino interactions as a function of the truth neutrino observable $-L/E_\nu$. The stacked histograms show the simulation prediction separated by neutrino production mode. The hatched bands represent the total uncertainty on the simulation prediction, including both statistical and systematic contributions. The black points represent the fitted interaction yields extracted from data through the unfolding from reconstructed muon q_μ/p'_μ.

After subtracting the expected backgrounds, we obtain

$$n_{\nu,\mathrm{obs}} = 766.8 \pm 28.7\,(\mathrm{stat.}) \pm 7.3\,(\mathrm{syst.})\,. \tag{8}$$

This result is consistent with the simulation expectation of 662.0 ± 84.6 within the uncertainties.

The corresponding total number of charged-current muon neutrino interactions occurring within the fiducial volume is

$$n_{\nu,\mathrm{fid}} = 3052.3 \pm 226.6 \tag{9}$$

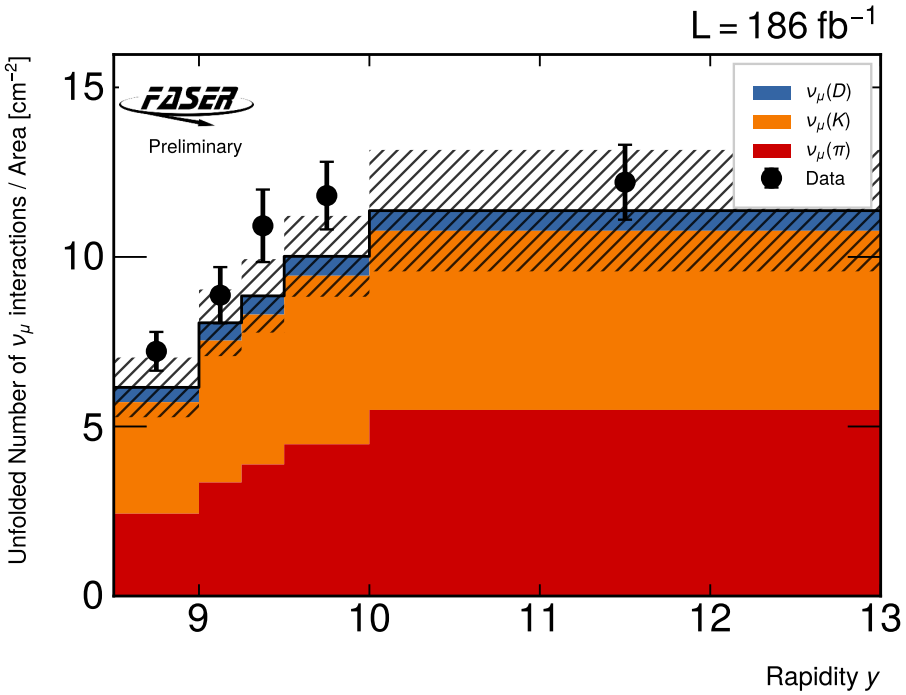


FIG. 6. Unfolded number of muon neutrino interactions as a function of neutrino rapidity y. The stacked histograms show the simulation prediction separated by neutrino production mode. The hatched bands represent the total uncertainty on the simulation prediction, including both statistical and systematic contributions. The black points represent the fitted interaction yields extracted from data.

The two-dimensional unfolded number of muon neutrino interactions and the neutrino flux extracted from the $(y, -L/E_\nu)$ analysis are shown in Figs. 7 and 8. The flux is obtained from the fitted number of neutrino interactions following the procedure described in Sec. 6. The measured flux is compared to the simulation prediction, separated into contributions from neutrinos originating from light meson decays and charm decays.

The one-dimensional cross-section extracted from the $-L/E$ analysis is shown in Fig. 9.

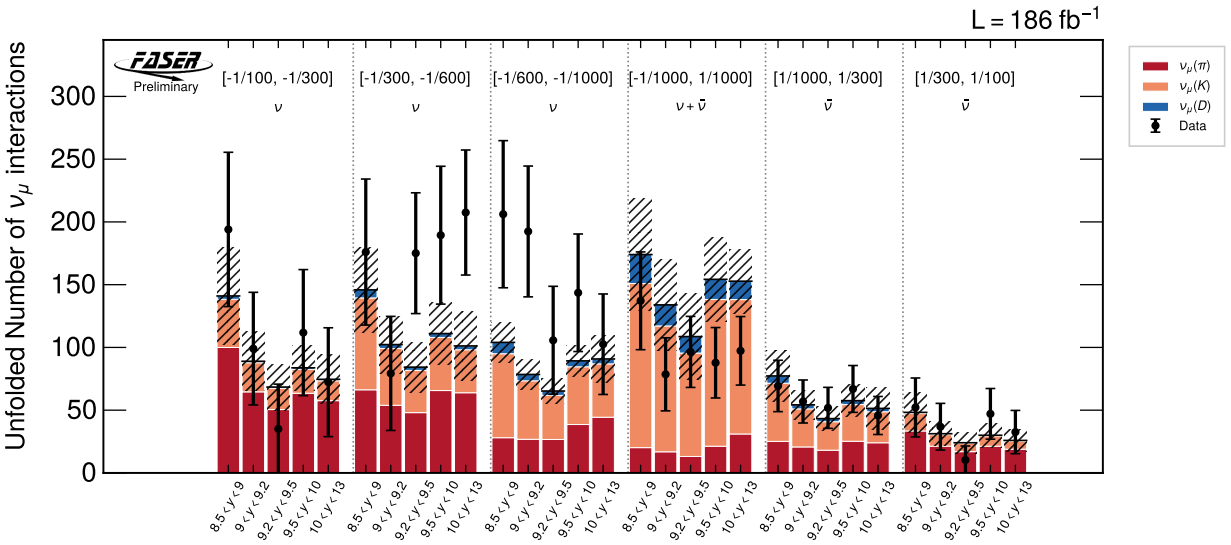


FIG. 7. The unfolded number of muon neutrino interactions measured in two-dimensional bins of neutrino rapidity y and lepton-number–over–energy $-L/E_\nu$. The horizontal axis shows rapidity bins, while the labels at the top of the figure indicate the corresponding $-L/E_\nu$ bin ranges. The stacked histograms show the simulation prediction separated by neutrino production mode. The hatched bands represent the total uncertainty on the simulation prediction, including both statistical and systematic contributions. The black points represent the fitted interaction yields extracted from data.

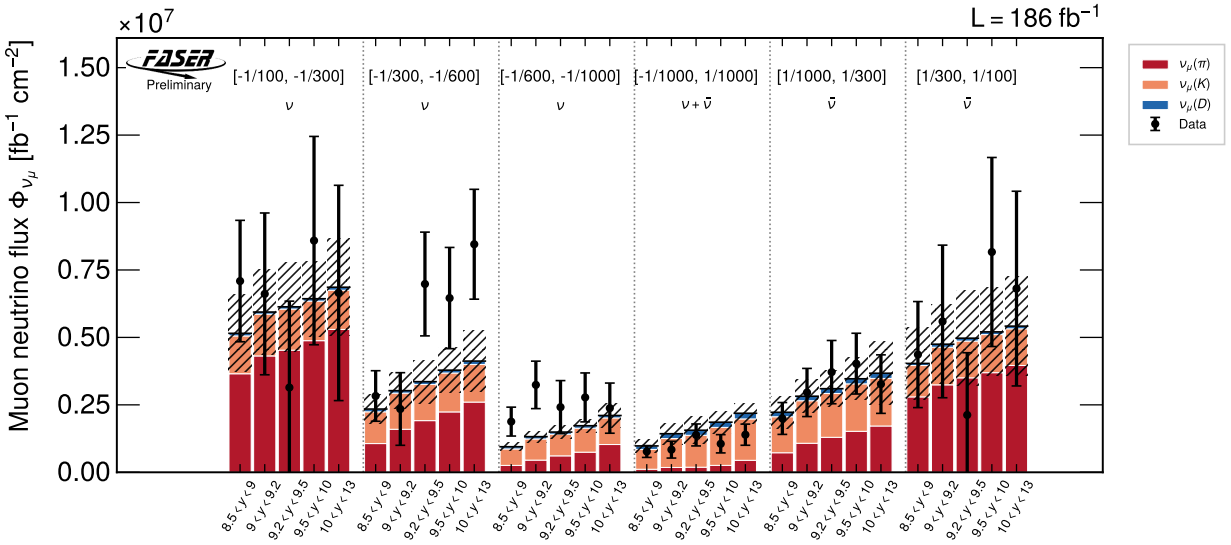


FIG. 8. Muon neutrino flux measured in two-dimensional bins of neutrino rapidity y and lepton-number–over–energy $-L/E_\nu$. The horizontal axis shows rapidity bins, while the labels at the top of the figure indicate the corresponding $-L/E_\nu$ bin ranges. The fitted flux values extracted from data are compared to the simulation prediction, shown as stacked contributions from different neutrino production modes. The hatched bands represent the total uncertainty on the simulation prediction, including both statistical and systematic contributions.

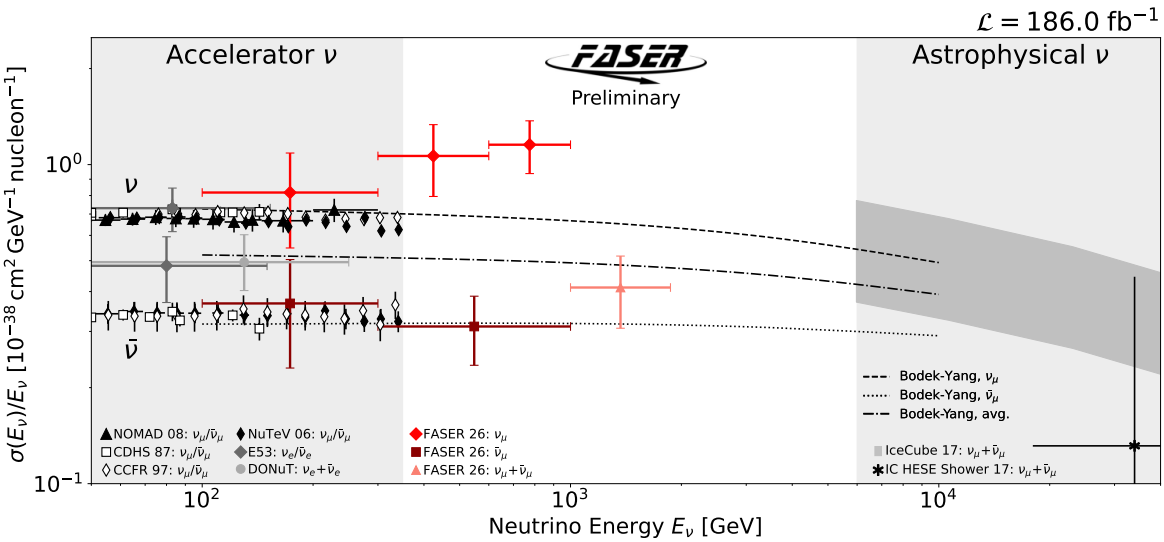


FIG. 9. Measured CC neutrino-nucleon cross sections divided by neutrino energy as a function of neutrino energy. The red markers show the results from this analysis. The cross section predicted by the Bodek Yang model is shown as dashed lines and previous measurements from fixed-target and astrophysical sources with gray/black markers [43–50].The horizontal error bars show the widths of the energy bins. The common high-energy bin contains neutrinos and anti-neutrinos with a momentum larger than 1 TeV.

9. SUMMARY AND CONCLUSION

This note presents the first double-differential measurement of the muon neutrino flux as a function of neutrino rapidity y and lepton-number–over–energy $-L/E_\nu$ in proton–proton collisions at the LHC. The analysis uses data collected during Run 3 at $\sqrt{s} = 13.6$ TeV, corresponding to an integrated luminosity of $(186 \pm 4)\,\mathrm{fb}^{-1}$. Muon neutrino candidates are selected using the electronic components of the FASER detector and reconstructed via charged-current interactions.

The signal neutrino interaction yields are extracted using a likelihood-based procedure that accounts for detector efficiency, migration effects, and systematic uncertainties. Summing over all analysis bins, we obtain a total fitted signal yield of

$$n_{\nu,\mathrm{obs}} = 766.8 \pm 28.7\,(\mathrm{stat.}) \pm 7.3\,(\mathrm{syst.})\,. \tag{10}$$

The CC neutrino-nucleon cross section is measured as a function of neutrino energy and is found to be consistent with the predictions of the Bodek Yang model. The extracted neutrino flux is measured as a function of both neutrino rapidity and lepton-number–over–energy, providing a two-dimensional characterization of the forward neutrino flux at the LHC. The results are found to be consistent with the predicted flux from light meson and charm decays within the experimental uncertainties. This measurement extends previous FASER results on neutrino cross sections and fluxes by providing the first double-differential determination of the muon neutrino flux in a collider environment, enabling a more detailed experimental mapping of TeV-scale neutrinos produced in proton–proton collisions.

[1] **Particle Data Group** Collaboration, "Review of Particle Physics," *Phys. Rev. D* **110** (8, 2024) 030001.

[2] **IceCube** Collaboration, "Measurement of the multi-TeV neutrino cross section with IceCube using Earth absorption," *Nature* **551** (2017) 596–600, arXiv:1711.08119 [hep-ex].

[3] **FASER** Collaboration, "Detecting and Studying High-Energy Collider Neutrinos with FASER at the LHC," *Eur. Phys. J. C* **80** (01, 2020) 61, arXiv:1908.02310 [hep-ex].

[4] J. L. Feng *et al.*, "The Forward Physics Facility at the High-Luminosity LHC," *J. Phys. G* **50** (2023) no. 3, 030501, arXiv:2203.05090 [hep-ex].

[5] A. De Rujula and R. Ruckl, "Neutrino and muon physics in the collider mode of future accelerators," in *SSC Workshop: Superconducting Super Collider Fixed Target Physics*, pp. 571–596. 5, 1984.

[6] **FASER** Collaboration, "First Direct Observation of Collider Neutrinos with FASER at the LHC," *Phys. Rev. Lett.* **131** (2023) no. 3, 031801, arXiv:2303.14185 [hep-ex].

[7] **SND@LHC** Collaboration, "Observation of Collider Muon Neutrinos with the SND@LHC Experiment," *Phys. Rev. Lett.* **131** (2023) no. 3, 031802, arXiv:2305.09383 [hep-ex].

[8] **FASER** Collaboration, "First Measurement of ν_e and ν_μ Interaction Cross Sections at the LHC with FASER's Emulsion Detector," *Phys. Rev. Lett.* **133** (7, 2024) 021802, arXiv:2403.12520 [hep-ex].

[9] **FASER** Collaboration, R. Mammen Abraham *et al.*, "First Measurement of the Muon Neutrino Interaction Cross Section and Flux as a Function of Energy at the LHC with FASER," *Phys. Rev. Lett.* **134** (2025) no. 21, 211801, arXiv:2412.03186 [hep-ex].

[10] FASER Collaboration, "Measurement of the Muon Neutrino Flux as a Function of the Rapidity with FASER." https://cds.cern.ch/record/2927629/, 2025. CERN-FASER-CONF-2025-001.

[11] **FASER** Collaboration, "The FASER Detector," *J. Instrum.* **19** (5, 2024) P05066, arXiv:2207.11427 [physics.ins-det].

[12] **LHCb** Collaboration, *LHCb calorimeters*. Technical design report. LHCb. CERN, Geneva, 2000. https://cds.cern.ch/record/494264.

[13] **FASER** Collaboration, "The tracking detector of the FASER experiment," *Nucl. Instrum. Meth. A* **1034** (2022) 166825, arXiv:2112.01116 [physics.ins-det].

[14] A. Abdesselam *et al.*, "The barrel modules of the ATLAS semiconductor tracker," *Nucl. Instrum. Meth. A* **568** (2006) 642–671.

[15] F. Kling and L. J. Nevay, "Forward neutrino fluxes at the LHC," *Phys. Rev. D* **104** (2021) no. 11, 113008, arXiv:2105.08270 [hep-ph].

[16] **FASER** Collaboration, "Neutrino Rate Predictions for FASER," *Phys. Rev. D* **110** (7, 2024) 012009, arXiv:2402.13318 [hep-ex].

[17] T. Pierog *et al.*, "EPOS LHC: Test of collective hadronization with data measured at the CERN Large Hadron Collider," *Phys. Rev.* **C92** (2015) 034906, arXiv:1306.0121 [hep-ph].

[18] S. Ostapchenko, "Monte carlo treatment of hadronic interactions in enhanced Pomeron scheme: QGSJET-II model," *Phys. Rev. D* **83** (Jan, 2011) 014018. https://link.aps.org/doi/10.1103/PhysRevD.83.014018.

[19] F. Riehn, R. Engel, A. Fedynitch, T. K. Gaisser, and T. Stanev, "Hadronic interaction model Sibyll 2.3d and extensive air showers," *Phys. Rev. D* **102** (Sep, 2020) 063002. https://link.aps.org/doi/10.1103/PhysRevD.102.063002.

[20] M. Fieg, F. Kling, H. Schulz, and T. Sjöstrand, "Tuning pythia for forward physics experiments," *Phys. Rev. D* **109** (2024) no. 1, 016010, arXiv:2309.08604 [hep-ph].

[21] P. Nason, "A New method for combining NLO QCD with shower Monte Carlo algorithms," *JHEP* **11** (2004) 040, arXiv:hep-ph/0409146.

[22] S. Frixione *et al.*, "Matching NLO QCD computations with Parton Shower simulations: the POWHEG method," *JHEP* **11** (2007) 070, arXiv:0709.2092 [hep-ph].

[23] S. Alioli *et al.*, "A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX," *JHEP* **06** (2010) 043, arXiv:1002.2581 [hep-ph].

[24] C. Bierlich *et al.*, "A comprehensive guide to the physics and usage of PYTHIA 8.3," *SciPost Phys. Codeb.* **2022** (2022) 8, arXiv:2203.11601 [hep-ph].

[25] L. Buonocore *et al.*, "Predictions for Neutrinos and New Physics from Forward Heavy Hadron Production at the LHC," *Eur. Phys. J. C, Particles and Fields* **84** (4, 2024) 363, arXiv:2309.12793 [hep-ph].

[26] C. Andreopoulos *et al.*, "The GENIE Neutrino Monte Carlo Generator," *Nucl. Instrum. Meth.* **A614** (2010) 87–104, arXiv:0905.2517 [hep-ph].

[27] C. Andreopoulos *et al.*, "The GENIE Neutrino Monte Carlo Generator: Physics and User Manual," arXiv:1510.05494 [hep-ph].

[28] **GENIE** Collaboration, "Neutrino-nucleon cross-section model tuning in GENIE v3," *Phys. Rev. D* **104** (2021) no. 7, 072009, arXiv:2104.09179 [hep-ph].

[29] A. Bodek *et al.*, "Modeling deep inelastic cross-sections in the few GeV region," *Nucl. Phys. B Proc. Suppl.* **112** (2002) 70–76, arXiv:hep-ex/0203009.

[30] A. Bodek, I. Park, and U.-k. Yang, "Improved low Q^2 model for neutrino and electron nucleon cross sections in few GeV region," *Nucl. Phys. B Proc. Suppl.* **139** (2005) 113–118, arXiv:hep-ph/0411202.

[31] A. Bodek and U.-k. Yang, "Axial and Vector Structure Functions for Electron- and Neutrino-Nucleon Scattering Cross Sections at all Q^2 using Effective Leading order Parton Distribution Functions," arXiv:1011.6592 [hep-ph].

[32] **FASER** Collaboration, R. Mammen Abraham *et al.*, "Neutrino rate predictions for FASER," *Phys. Rev. D* **110** (2024) no. 1, 012009, arXiv:2402.13318 [hep-ex].

[33] **GEANT4** Collaboration, "GEANT4–a simulation toolkit," *Nucl. Instrum. Meth. A* **506** (2003) 250–303.

[34] D. Mancusi *et al.*, "Improving the description of proton-induced one-nucleon removal in intranuclear-cascade models," *Phys. Rev. C* **91** (2015) no. 3, 034602, arXiv:1411.4562 [nucl-th].

[35] E. Mendoza, D. Cano-Ott, T. Koi, and C. Guerrero, "New standard evaluated neutron cross section libraries for the geant4 code and first verification," *IEEE Transactions on Nuclear Science* **61** (2014) no. 4, 2357–2364.

[36] **ATLAS** Collaboration, "The ATLAS Collaboration Software and Firmware." ATL-SOFT-PUB-2021-001, 2021. https://cds.cern.ch/record/2767187.

[37] X. Ai *et al.*, "A Common Tracking Software Project," *Comput. Softw. Big Sci.* **6** (2022) no. 1, 8, arXiv:2106.13593 [physics.ins-det].

[38] F. James and M. Roos, "Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations," *Comput. Phys. Commun.* **10** (1975) 343–367.

[39] M. Fieg, F. Kling, H. Schulz, and T. Sjöstrand, "Tuning Pythia for Forward Physics Experiments." http://arxiv.org/abs/2309.08604.

[40] **ATLAS** Collaboration, "Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC," *Eur. Phys. J. C* **83** (10, 2023) 982, arXiv:2212.09379 [hep-ex].

[41] **ATLAS** Collaboration, "Preliminary analysis of the luminosity calibration of the ATLAS 13.6 TeV data recorded in 2022," CONF note, CERN, Geneva, 2023. https://cds.cern.ch/record/2853525.

[42] **ATLAS** Collaboration, "Preliminary analysis of the luminosity calibration for the ATLAS 13.6 TeV data recorded in 2023," CONF note, CERN, Geneva, 2024. http://cds.cern.ch/record/2900949.

[43] **IceCube** Collaboration, "Measurement of the multi-TeV neutrino cross section with IceCube using Earth absorption," *Nature* **551** (2017) 596–600, arXiv:1711.08119 [hep-ex].

[44] M. Bustamante and A. Connolly, "Extracting the Energy-Dependent Neutrino-Nucleon Cross Section above 10 TeV Using IceCube Showers," *Phys. Rev. Lett.* **122** (2019) no. 4, 041101, arXiv:1711.11043 [astro-ph.HE].

[45] **NOMAD** Collaboration, "A Precise measurement of the muon neutrino-nucleon inclusive charged current cross-section off an isoscalar target in the energy range 2.5 < E(nu) < 40-GeV by NOMAD," *Phys. Lett. B* **660** (2008) 19–25, arXiv:0711.1183 [hep-ex].

[46] J. P. Berge *et al.*, "Total Neutrino and Anti-neutrino Charged Current Cross-section Measurements in 100-GeV, 160-GeV and 200-GeV Narrow Band Beams," *Z. Phys. C* **35** (1987) 443.

[47] W. G. Seligman, *A Next-to-Leading Order QCD Analysis of Neutrino - Iron Structure Functions at the Tevatron*. PhD thesis, Nevis Labs, Columbia U., 1997.

[48] **NuTeV** Collaboration, "Precise measurement of neutrino and anti-neutrino differential cross sections," *Phys. Rev. D* **74** (2006) 012008, arXiv:hep-ex/0509010.

[49] **DONuT** Collaboration, K. Kodama *et al.*, "Final tau-neutrino results from the DONuT experiment," *Phys. Rev. D* **78** (2008) 052002, arXiv:0711.0728 [hep-ex].

[50] C. Baltay *et al.*, "$\nu_\mu - \nu_e$ Universality in Charged Current Neutrino Interactions," *Phys. Rev. D* **41** (1990) 2653.