Published by IOP Publishing for Sissa Medialab

Received: *October 2, 2026*
Accepted: *October 23, 2026*
Published: *November 16, 2026*

Goodness-of-fit for multi-distribution neutrino cross-section measurements with shared events

F. Martínez López[a,*] **L. Cooper-Troendle**[b] **S. Gardiner**[c] **P. Green**[d] **and T. Mohayai**[a]

[a] *Indiana University, Bloomington, IN 47405, U.S.A.*
[b] *University of Pittsburgh, Pittsburgh, PA 15260, U.S.A.*
[c] *Fermi National Accelerator Laboratory, Batavia, IL 60510, U.S.A.*
[d] *University of Oxford, Oxford, OX1 3RH, United Kingdom*

E-mail: frmart@iu.edu

Abstract: When multiple differential cross-section measurements are extracted from a common event sample, the same events contribute simultaneously to multiple distributions. This event-sharing structure imposes exact linear constraints among the bin counts, reducing the effective dimensionality of the measurement below the total number of bins. Since the covariance matrix obeys these inter-distribution constraints, it contains a rank deficiency. While limited numerical precision may inadvertently restore invertibility, the χ^2 contributions along constrained dimensions will be arbitrary, yielding unphysically inflated χ^2 values in global goodness-of-fit tests. We present the range-projected χ^2, a test statistic that restricts the goodness-of-fit test to the subspace carrying independent statistical information, yielding a χ^2 with $N_{\text{bins}} - N_{\text{null}}$ degrees of freedom, where N_{null} is the number of independent constraints. We show that this rank deficiency is a predictable consequence of the event-sharing structure, and that N_{null} decomposes into a structural contribution determined a priori from the binning geometry and a kinematic contribution that depends on the phase-space occupancy. The method is validated with an analytical toy model and a simulated neutrino-argon cross-section measurement including unfolding and multi-source systematic uncertainties.

Keywords: Analysis and statistical methods; Neutrino detectors

ArXiv ePrint: 2609.24796

*Corresponding author.

© 2026 The Author(s). Published by IOP Publishing Ltd on behalf of Sissa Medialab. Original content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

https://doi.org/10.1088/1748-0221/21/11/P11017

2026 JINST 21 P11017

Contents

1 Introduction **1**

2 Constraints from shared events **3**
- 2.1 Statistical covariance from shared events 3
- 2.2 Combination matrix and null space 4
- 2.3 Determination of the structural null space 5
- 2.4 Effect of systematic uncertainties 7

3 Range-projected χ^2 test statistic **9**
- 3.1 Test statistic definition 10
- 3.2 Statistical validity 11
- 3.3 Relation to alternative approaches 11

4 Toy example **13**
- 4.1 Model and measurement setup 13
- 4.2 Null-space structure 14
- 4.3 Goodness-of-fit validation 14
- 4.4 Parameter estimation 17

5 Event generator study **18**
- 5.1 Model and measurement setup 18
- 5.2 Systematic uncertainties and unfolding 19
- 5.3 Results 19

6 Limitations and scope **21**

7 Summary and conclusions **23**

A Structural combination matrix algorithm **24**

1 Introduction

Modern neutrino cross-section literature is experiencing an increase of published measurements reporting multiple differential distributions extracted from the same selected event sample. In these analyses, the same events contribute to multiple reported distributions, leading to strong correlations between them. Examples include MicroBooNE measurements of cross sections for final states with and without protons [1, 2] and multi-differential QE-like cross section measurements in kinematic imbalance variables [3, 4]. Other relevant examples include the multiple final-state kinematic projections from T2K [5, 6], and double-differential inclusive results from NOvA [7–9]. Separate but related efforts have produced simultaneous measurements across different nuclear targets [10–12] and combined neutrino-antineutrino results [13], where correlations arise from shared systematic uncertainties rather than shared events.

2026 JINST 21 P11017

When multiple distributions are extracted from a common event sample, the event-sharing structure imposes exact linear constraints among the bin counts. If every selected event populates one bin in each distribution, then the total number of events in each distribution is exactly identical. More generally, the pattern of shared bin occupancy across distributions determines a set of constraints that reduce the effective dimensionality of the measurement below the total bin count. These constraints are a property of the measurement binning itself, independent of any particular statistical treatment.

The blockwise unfolding framework of ref. [14] provides a rigorous method for computing the full cross-distribution covariance matrices in such measurements. In this approach, the bins are organized into blocks, with each block corresponding to one reported differential cross-section distribution. Each bin block is unfolded independently using any standard method such as iterative D'Agostini unfolding [15] or Wiener-SVD [16]. Separately, the full covariance matrix, including correlations between bins from different distributions, is constructed in reconstructed space and propagated to truth variables using a global unfolding matrix that consists of the unfolding matrices from each block in a block-diagonal format. The statistical covariance between bins in different blocks is determined by the number of events they share. The constraints mentioned above manifest as a rank deficiency in the statistical and systematic covariance matrices that respect the same event sharing patterns, as certain directions in bin space carry identically zero variance. The number of these degenerate directions is predictable from the binning geometry.

A natural use of these combined measurements is computing global goodness-of-fit metrics, such as a χ^2 statistic, that test a model against all distributions simultaneously. Evaluating each distribution separately discards the correlations between blocks, both statistical fluctuations and systematic uncertainties correlated across distributions; a global test uses them to sharpen the comparison. Such global metrics are essential for interaction model tuning efforts [17–19], where measurements of multiple kinematic distributions provide complementary sensitivity to different aspects of the underlying physics. The event-sharing constraints, however, fix the data exactly, and a forward-folded prediction only approximately, along these directions. The data-model difference there is therefore zero in principle, and a global χ^2 should receive no contribution from them. In practice it does. Independently of how the covariance matrix is built, finite numerical precision keeps it from being exactly singular, so it can still be inverted, while any systematic treatment that breaks the sharing lifts the constrained directions to spurious small variances. The residual along these directions, left small but nonzero by imperfect forward-folding, is then divided by a variance that is itself zero or spuriously small, an ill-defined ratio that numerical noise renders arbitrarily large. The result inflates the χ^2, and even if the offending contributions were removed by hand the test would be left with the wrong number of degrees of freedom unless the constraints are properly identified and excluded.

We present a method to correctly compute a global goodness-of-fit test statistic for multi-distribution measurements with shared events. The approach identifies the constrained directions from the bin-combination structure of the measurement and projects it onto the subspace orthogonal to the constrained directions; we refer to the resulting test statistic as the range-projected χ^2. Our contribution is twofold: we show that the rank deficiency of the covariance is a predictable consequence of the shared-event structure, rather than an incidental numerical degeneracy, and we provide a prescription to compute the number of constrained directions, with a structural component fixed a priori by the binning geometry and a kinematic component set by the phase-space occupancy. The statistical

2026 JINST 21 P11017

foundation of the projected statistic is the standard treatment of quadratic forms with singular covariance matrices [20], a connection we make explicit while contrasting our approach with related methods.

The paper is organized as follows. Section 2 describes the constraints imposed by the shared-event structure; section 3 constructs the range-projected test statistic and relates it to existing approaches; section 4 and section 5 validate the method on a toy model and a simulated measurement; and section 6 discusses the scope and limitations of the method.

2 Constraints from shared events

Consider a neutrino cross-section analysis simultaneously measuring B different flux-integrated differential cross section distributions, each with N_b true bins. If the measurement is extracted from a common event sample then each selected event will populate exactly one bin in each distribution it belongs to. Recent MicroBooNE analyses [2, 21, 22] have adopted the blockwise unfolding approach presented in ref. [14] that follows this approach to simultaneous cross-section measurements.

This method provides a systematic framework for reporting the correlations between the different extracted differential cross sections. It also gives a general prescription to construct statistical covariance matrices that respect the correlations between bins in different distributions or blocks. The full inter-block covariance matrix is first constructed in reconstructed space and then propagated through unfolding using a block-diagonal unfolding matrix formed from the direct sum of the individual unfolding matrices for each block. The result of the combined measurement is a vector of $N_{\text{bins}} = \sum_{b=1}^{B} N_b$ cross-section values with an $N_{\text{bins}} \times N_{\text{bins}}$ covariance matrix. The total covariance decomposes as $C_{\text{total}} = C_{\text{stat}} + C_{\text{syst}}$, where the statistical covariance component, computed from the shared-event structure, encodes the constraints described in the introduction.

2.1 Statistical covariance from shared events

Inter-block correlations in the blockwise unfolding framework can be understood by considering two overlapping bins X and Y, which may be partitioned into three non-overlapping regions u, v, and w. The events in the intersection $X \cap Y$ are assigned to bin v, while bin u contains the events in the relative complement of Y in X (i.e. $X \setminus Y$), and bin w corresponds to the relative complement of X in Y (i.e. $Y \setminus X$). Denoting the number of predicted events in each bin as n_i, and assuming these follow independent Poisson distributions, the covariance between the counts in bins X and Y is given by [14]

$$\begin{aligned} \text{Cov}(n_X, n_Y) &= \text{Cov}(n_u + n_v, n_v + n_w) \\ &= \text{Cov}(n_u, n_v) + \text{Cov}(n_u, n_w) + \text{Cov}(n_v, n_v) + \text{Cov}(n_v, n_w) \\ &= \text{Var}(n_v), \end{aligned} \tag{2.1}$$

as by definition bins u, v, and w have no events in common. This implies that the covariance between any two bins can be estimated from the variance of the events they share.

Monte Carlo (MC) events do not in general count as unity. It is standard practice to reweight event generator predictions to an alternative or tuned interaction model rather than simulate an entirely new sample, which would be computationally prohibitive [23, 24]. Each event therefore carries a weight, and the variance for a set of such events is estimated from the sum of their squared weights. This gives us a general prescription to calculate the elements of the statistical covariance matrix as

$$(C_{\text{stat}})_{ij} = \sum_{e \in (i \cap j)} w_e^2, \tag{2.2}$$

2026 JINST 21 P11017

where w_e is the weight of event e, and the sum runs over events belonging to the intersection of i and j. In the case of measured data the weights are 1, and the covariance element simplifies to the number of events shared between the bins n_{ij}.

Throughout, we restrict the discussion to measurements in which each block uses the same number and structure for the reconstructed and true bins (as is the case in ref. [21]). In this case the unfolding matrix of each block is square and invertible, so after propagation through the block-diagonal unfolding the covariance matrix in truth space has the same rank as its reconstructed-space counterpart. The rank has a direct interpretation in terms of the number of distinct event-sharing patterns across distributions, which we discuss next.

2.2 Combination matrix and null space

The previous discussion shows that the statistical covariance between any two bins is determined solely by the events they share. We now formalize this observation to derive an exact formula for the rank of C_{stat}, and therefore for the number of directions in bin space that carry zero statistical variance.

For each event e, we can define the indicator vector $\boldsymbol{\phi}_e \coloneqq (\phi_{e,i})_{i=1}^{N_{\text{bins}}}$ such that $\phi_{e,i} = 1$ if $i \in \text{bins}(e)$ and $\phi_{e,i} = 0$ if $i \notin \text{bins}(e)$, where $\text{bins}(e)$ is the set of bins event e populates across all distributions. The indicator vector records which bins an event populates across all distributions. For example, in a two-block measurement with 3 bins each, an event populating bin 2 of block 1 and bin 1 of block 2 has indicator vector $\boldsymbol{\phi}_e = (0, 1, 0, 1, 0, 0)$. The statistical covariance matrix can then be re-written in terms of indicator vectors as

$$C_{\text{stat}} = \sum_e w_e^2 \, \boldsymbol{\phi}_e \, \boldsymbol{\phi}_e^{\mathsf{T}}. \tag{2.3}$$

This outer-product decomposition makes the rank structure of C_{stat} clear. The rank of the statistical covariance matrix equals the dimension of $\text{span}\{\boldsymbol{\phi}_e\}$, which depends only on the distinct indicator vectors present, not on their weights.

Since all weights satisfy $w_e > 0$, no outer-product term in eq. (2.3) vanishes, and the column space of C_{stat} is exactly $\text{span}\{\boldsymbol{\phi}_e\}$. We define the bin-combination matrix Ω as the matrix whose rows are the unique indicator vectors among $\{\boldsymbol{\phi}_e\}$. Each row of Ω represents a distinct pattern of bin occupancy across the blocks. By construction the row space of Ω equals $\text{span}\{\boldsymbol{\phi}_e\}$, and therefore

$$\text{rank}(C_{\text{stat}}) = \text{rank}(\Omega). \tag{2.4}$$

The number of degenerate directions is then

$$N_{\text{null}} = N_{\text{bins}} - \text{rank}(\Omega). \tag{2.5}$$

Physically, each of the N_{null} directions corresponds to a linear combination of bins along which the measurement has identically zero variance. Such a combination is fixed by the event-sharing structure rather than measured, so the measurement has no statistical power to constrain it. Importantly, both Ω and N_{null} are directly computable from the MC event-to-bin mapping alone, before constructing any covariance matrix, confirming that the constraints are a property of the measurement rather than the statistical treatment. In practice the mapping is read off from the MC sample, so the resulting count reflects not only the binning geometry but also the regions of phase space that the sample populates. We separate these two contributions in section 2.3.

2026 JINST 21 P11017

A vector $\mathbf{x}$ is in the null space of C_{stat} if and only if, for every event e

$$\sum_{i \in \text{bins}(e)} x_i = 0, \tag{2.6}$$

or equivalently $\Omega\mathbf{x} = 0$. Each unique indicator vector acts as a linear constraint on the null-space vectors, requiring that the components of $\mathbf{x}$ sum to zero over the bins populated by the different events. The null space directions can be divided into two physically distinct categories: structural constraints that follow from the block topology alone, and kinematic constraints that depend on which regions of phase space are populated.

2.3 Determination of the structural null space

The combination matrix Ω presented previously is constructed from MC events by determining which bins each event populates and collecting the unique indicator vectors. However, the indicator vectors that can appear are constrained by two distinct mechanisms. The first is purely structural: the geometry of the binning scheme forbids certain bin combinations. The second is kinematic: correlations in model predictions between the different variables render some combinations unreachable even though they are geometrically allowed. These constraints are imposed by the underlying physics effects, mainly the differential cross section and the detector acceptance. In the following, we show that the structural constraints can be determined a priori, without generating MC events, from the bin edges used.

Our approach is to identify the geometrically allowed bin combinations directly from the bin edges. The key is to subdivide the binning finely enough that each resulting region lies within a single bin of every block, so that the allowed combinations can be read off directly. The structural null space follows from these combinations, as we develop below.

The variables a measurement is binned over form the set $\{v_1, \ldots v_D\}$. Each block b uses a subset of these variables, $\text{vars}(b) \subseteq \{v_1, \ldots v_D\}$. Restricting the discussion to combinations of single- and double-differential measurements, we have $|\text{vars}(b)| \in \{1, 2\}$ for every b. For each variable v used in a given block b, the binning scheme defines a set of bin edges $E_v^{(b)}$ that partition the range of v into intervals. In the case of a two-dimensional block, the edges in the binning variable may be slice-dependent, but nevertheless these can be collected into a single set of edges per variable across slices. For each variable $v \in \bigcup_b \text{vars}(b)$, we can define the refined edges

$$\mathcal{E}_v = \bigcup_{b\,:\,v \in \text{vars}(b)} E_v^{(b)}, \tag{2.7}$$

with $M_v = |\mathcal{E}_v| - 1$ denoting the number of refined bins for variable v. Since $\mathcal{E}_v$ contains all the bin edges from every block that uses that variable, each refined bin is contained entirely within one original bin of every such block.

The common refinement of all blocks is the Cartesian product of the refined bins across all variables, denoted by $\mathcal{R}$. Each cell $c \in \mathcal{R}$ specifies a refined bin index for every variable in the measurement. For every block b, there exists a unique mapping $\beta_b(c)$ between each cell c and an original bin in said block. The elements of the indicator vector $\boldsymbol{\phi}_c := (\phi_{c,i})_{i=1}^{N_{\text{bins}}}$ associated to cell c can be written as

$$\phi_{c,i} = \begin{cases} 1 & \text{if } i = \beta_b(c) \text{ for any block } b, \\ 0 & \text{otherwise.} \end{cases} \tag{2.8}$$

2026 JINST 21 P11017

If block b does not use variable v then β_b is independent of the refined bin index in v. Therefore, multiple cells may produce identical indicator vectors.

The structural combination matrix Ω_{struct} is defined as the matrix whose rows are the unique indicator vectors from the cells of the common refinement $\{\boldsymbol{\phi}_c : c \in \mathcal{R}\}$. By construction, every indicator vector that appears in the Ω matrix obtained from the MC events is also present in Ω_{struct}, therefore

$$\text{rank}(\Omega) \leq \text{rank}(\Omega_{\text{struct}}) \leq N_{\text{bins}}. \tag{2.9}$$

This hierarchy motivates a decomposition of the total number of degenerate directions

$$\begin{aligned} N_{\text{null}} &= N_{\text{bins}} - \text{rank}(\Omega_{\text{struct}}) + \text{rank}(\Omega_{\text{struct}}) - \text{rank}(\Omega) \\ &= N_{\text{null}}^{\text{struct}} + N_{\text{null}}^{\text{kinem}}. \end{aligned} \tag{2.10}$$

The first term, $N_{\text{null}}^{\text{struct}}$, counts null directions that arise purely from the binning geometry and is computable a priori from the bin edges alone. The second term, $N_{\text{null}}^{\text{kinem}}$, accounts for additional null directions introduced by phase-space restrictions that exclude geometrically allowed combinations and can only be determined from Monte Carlo. For example, in a measurement of muon momentum and scattering angle, the common refinement may include cells for high-momentum backward-going muons. These cells are allowed by the binning and therefore appear in Ω_{struct}. If they are empty because of the interaction kinematics or detector acceptance, their indicator vectors do not appear in Ω, producing additional kinematic null directions. A description of the algorithm used to compute Ω_{struct} is presented in appendix A.

The presence of kinematic null directions indicates that certain regions of the combined bin space are empty, which may signal that the binning extends beyond the kinematic reach of the measurement or that it is finer than the available statistics can support. In such cases, merging or removing the offending bins may be preferable to relying on a projection to handle the resulting constraints. The structural null directions, by contrast, are an unavoidable consequence of measuring multiple distributions from shared events. While their number depends on the binning, as discussed below, they cannot be removed entirely by any choice of bin edges.

The structure of the null space of Ω_{struct} depends on how the bin edges for the different blocks overlap. In the simplest case, when blocks do not share variables, every combination of one bin per block is allowed. The resulting null space is $(B-1)$-dimensional, corresponding to the $B-1$ independent constraints that require all block totals to be equal. When blocks share one or more variables, the allowed combinations are restricted and the dimension of the null space can be different from $B-1$. The number of null directions depends on how shared variable bin edges are aligned between blocks. Aligned slice edges decouple the measurement into different sub-problems, adding additional sum rules and increasing $N_{\text{null}}^{\text{struct}}$. On the other hand, misaligned edges couple adjacent distributions, potentially reducing the number of constraints. In the general case, the interplay between shared variables, bin edge alignment, and block topology makes a closed-form expression impractical, but the computation of $\text{rank}(\Omega_{\text{struct}})$ from the bin edge specifications is straightforward.

The rank deficiency and null space discussed so far are properties of the reconstructed-space covariance, and, under the restriction to matched reconstructed and true binning adopted, to the unfolded covariance as well. If a block were instead unfolded from a larger number of reconstructed bins to a smaller number of true bins, the rank of the unfolded covariance would be bounded by the number

2026 JINST 21 P11017

of true bins, and the null space of the unfolded measurement could in principle be empty even though the covariance matrix is rank deficient in reconstructed space. We do not consider this case further here.

The structural and kinematic components of the null space respond differently to systematic uncertainties. Whether these cause the rank deficiency of C_{stat} to persist, partially resolve, or fully resolve in the total covariance C_{total} depends on the nature of the systematic treatment, which we discuss next.

2.4 Effect of systematic uncertainties

Previously, we established that the event-sharing structure imposes $N_{\text{null}} = N_{\text{bins}} - \text{rank}(\Omega)$ exact linear constraints on the bin counts, which manifest as a rank deficiency of C_{stat}. By default, these same constraints apply to the construction of covariance matrices to describe various systematic effects, causing the same rank deficiencies. However, there is one notable exception in the case of computing the cross section covariance matrix, which we will discuss later. For now, we will focus on the treatment of other, more straightfordward systematic effects.

A systematic variation s produces a modified vector of expected events $\mathbf{n}_s$, whose departure from the central-value prediction can be expressed in terms of a shift vector

$$\boldsymbol{\xi}_s = \mathbf{n}_s - \mathbf{n}_{\text{CV}}. \tag{2.11}$$

The different contributions to C_{syst} are constructed from the outer product of these shift vectors. A null direction $\mathbf{x}$ of C_{stat} remains a null direction of C_{total} if and only if $\mathbf{x} \cdot \boldsymbol{\xi}_s = 0$ for every systematic variation in the uncertainty budget. We will now demonstrate that for a typical systematic variation, the shift vectors have zero projections onto the structural and kinematic null spaces.

Consider first the simplest case, where a systematic s acts by reweighting the reconstructed event counts directly. Each event e is assigned a modified weight $w_e^{(s)}$ while preserving its bin assignments. The shift vector can then be written as

$$\boldsymbol{\xi}_s = \sum_e (w_e^{(s)} - w_e)\, \boldsymbol{\phi}_e. \tag{2.12}$$

For any null vector $\mathbf{x}$ which satisfies the null-space condition of eq. (2.6), the projection of $\mathbf{x}$ onto the shift vector gives

$$\begin{aligned} \mathbf{x} \cdot \boldsymbol{\xi}_s &= \sum_e (w_e^{(s)} - w_e)\, (\boldsymbol{\phi}_e \cdot \mathbf{x}) \\ &= \sum_e (w_e^{(s)} - w_e) \sum_{i=1}^{N_{\text{bins}}} \phi_{e,i}\, x_i \\ &= 0, \end{aligned} \tag{2.13}$$

since the inner sum vanishes for every event. Direct reweighting therefore preserves the entire null space, adding no rank to C_{total}. Flux uncertainties, which modify the neutrino energy spectrum through per-event weights, are the classic example. When cross-section model variations and hadronic reinteraction uncertainties are implemented as per-event reweighting, they behave identically.

Systematic uncertainties can also be evaluated by directly comparing reconstructed event counts from independent MC samples simulated with nominal and varied parameters. Detector systematic uncertainties, which usually require independent samples with modified simulation parameters that

2026 JINST 21 P11017

cannot be captured by event reweighting, are the typical example where this prescription is applied. In this case, the shift vector takes the form

$$\boldsymbol{\xi}_s = \sum_{e'} w_{e'}^{(s)} \, \boldsymbol{\phi}_{e'}^{(s)} - \sum_{e} w_e \, \boldsymbol{\phi}_e, \tag{2.14}$$

where the first sum runs over the events in the varied-model sample and the second sum over the nominal sample. Since each event in both samples populates exactly one bin per block, their difference projects to zero onto the structural null directions by an analogous argument to that used for the direct reweighting. Kinematic null directions, however, can be lifted if the variation scatters events into reconstructed-bin combinations absent in the nominal sample. In that case, new indicator vectors appear that are not rows of Ω, potentially breaking the corresponding constraints. Therefore, despite their different treatment, detector systematic uncertainties also yield the same structural null directions determined by the binning scheme.

In contrast to these methods is a qualitatively different approach centered on varying the response matrix, commonly used in the calculation of cross-section systematic uncertainties. In this approach, adopted as the standard prescription in MicroBooNE for neutrino interaction uncertainties [14], a modified response matrix Δ'_s, mapping true bins to reconstructed bins under the given variation, is constructed for each model parameter of interest. This matrix is obtained either by reweighting the events used to build the nominal response matrix or from an independent MC sample with varied simulation parameters. In this case, the predicted reconstructed counts for systematic variation s are

$$\mathbf{n}_s = \Delta'_s \, \boldsymbol{\varphi}_{\mathrm{CV}}, \tag{2.15}$$

where $\boldsymbol{\varphi}_{\mathrm{CV}}$ is the central-value truth prediction. The shift vector is then

$$\boldsymbol{\xi}_s = (\Delta'_s - \Delta) \, \boldsymbol{\varphi}_{\mathrm{CV}}. \tag{2.16}$$

Since Δ is block-diagonal, the difference $\Delta'_s - \Delta$ acts independently on each block, projecting the nominal truth vector through block-specific changes in efficiency and migration. The per-event coupling that makes direct reweighting respect all null directions is absent in this prescription. When the weight of an event changes, it affects simultaneously every block it populates; the response-matrix method, in contrast, processes each block independently through its own sub-matrix. The projection of $\boldsymbol{\xi}_s$ onto any null vector, whether structural or kinematic, therefore involves differences in how each block responds to the same truth-level prediction, and in general is nonzero. Response-matrix variations thus break both types of null directions, regardless of how Δ'_s is obtained.

In practice, whether or not the null space of C_{stat} survives in C_{total}, including these directions in a χ^2 test via standard matrix inversion produces inflated values with the wrong effective number of degrees of freedom. Along these directions the measurement carries no statistical information, yet the systematic model, in particular the cross section systematic, lifts their variance to a small nonzero value. The inverse covariance then weights them by the reciprocal of that small value. Any residual between data and prediction that projects onto these directions is amplified, producing a contribution to the χ^2 far in excess of one per degree of freedom, and not based in any physical data-model difference. We make this mechanism precise in section 4.3, where the excess is shown to be a noncentrality concentrated on the lifted directions. A valid goodness-of-fit test must therefore be restricted to the subspace where the measurement carries genuine sensitivity, which is the role of the projection method developed in the following section.

2026 JINST 21 P11017

3 Range-projected χ^2 test statistic

Before constructing the test statistic, we must establish why restricting it to the non-null subspace discards no physical information. Along a null direction $\mathbf{x}$ of C_{stat} the event-sharing constraint $\Omega\mathbf{x} = 0$ holds identically. The same selected events enter the bin combinations on both sides, so the data value along $\mathbf{x}$ is fixed by event conservation rather than measured, and a truth-level prediction respects the same constraint. The measurement therefore has no statistical power along $\mathbf{x}$, and no physical model can produce a genuine difference there. Forward folding a prediction through the detector response matrix does not in general preserve the constraint exactly, so a folded prediction can acquire a small non-zero component along $\mathbf{x}$. However, this residual is an artifact of the imperfect forward-folding rather than a real disagreement, and carries no statistical meaning, as the covariance assigns the direction no genuine variance. Discarding these directions removes no sensitivity, only the ill-defined contributions identified in section 2.4. This holds exactly to the extent that the event-sharing structure is exact; when sharing is only approximate the null directions acquire small but nonzero variance, a regime we return to in section 6. Under the exact-sharing assumption, then, it is safe to formulate the global goodness-of-fit test statistic in the subspace orthogonal to the constrained directions, which we now construct.

The constraints are most naturally identified in reconstructed space, where the combination matrix is constructed. The non-null subspace can be obtained decomposing $\Omega^\mathsf{T}\Omega$. This matrix is a square $N_{\text{bins}} \times N_{\text{bins}}$ matrix with the same column space and null space as Ω itself, making it a more compact representation of the same information. We can write its eigendecomposition as

$$\Omega^\mathsf{T}\Omega = V\,S\,V^\mathsf{T}, \tag{3.1}$$

where $S = \text{diag}(s_1, \ldots, s_{N_{\text{bins}}})$ contains the eigenvalues in decreasing order and V is the orthogonal matrix of eigenvectors. Of these eigenvalues, $r = \text{rank}(\Omega)$ are nonzero. Partitioning $V = [V_r \mid \bar{V}_r]$, with V_r collecting the eigenvectors associated with the nonzero eigenvalues and $\bar{V}_r$ those with zero eigenvalues, the columns of V_r span the subspace of reconstructed bin space that is not constrained by the event-sharing structure while those of $\bar{V}_r$ span the null space. The orthogonal projector onto this subspace is

$$P_{\text{reco}} = V_r\,V_r^\mathsf{T}. \tag{3.2}$$

Note that P_{reco} is identical to the projector onto the range of C_{stat} as defined in eq. (2.3), since C_{stat} and $\Omega^\mathsf{T}\Omega$ share the same column space by construction. The combination matrix formulation makes explicit the fact that the projector is determined by the measurement geometry.

Cross-section measurements are typically reported in unfolded space, and so it is natural to assume this is the space where goodness-of-fit tests should be evaluated. The blockwise unfolding matrix U, which is block-diagonal and invertible, maps reconstructed-space vectors onto unfolded space, which is truth space in the case of an unregularized measurement. The non-null basis vectors propagate as $W_r = U\,V_r$. Since U is not orthogonal in general, the columns of W_r are not orthonormal, and the projector onto their span takes the general form

$$P_{\text{true}} = W_r\,(W_r^\mathsf{T}\,W_r)^{-1}\,W_r^\mathsf{T}. \tag{3.3}$$

Since U is invertible, W_r has the same rank as V_r and the number of constrained directions is preserved. However, as shown next, the ordering of projection and unfolding matters as the commutation relation

2026 JINST 21 P11017

$U\,P_{\text{reco}} = P_{\text{true}}\,U$ generally does not hold. For this reason we choose to define the test statistic in reconstructed space, where the constraints originate.

In practice, limited numerical precision causes the computation of eigenvalues of $\Omega^{\mathsf{T}}\Omega$ that should be zero to yield small nonzero values. In many cases, these can be identified by setting a threshold relative to the largest eigenvalue (e.g., $s_i/s_1 < 10^{-10}$). So long as physical eigenvalues do not span too many orders of magnitude and remain sufficiently far from zero, a clear spectral gap of several orders of magnitude can be observed between physical eigenvalues and numerical noise, making the identification unambiguous.

3.1 Test statistic definition

The goodness-of-fit test can be evaluated in reconstructed space, where the residual vector is given by $\boldsymbol{\delta} = \mathbf{d} - \Delta\,\boldsymbol{\varphi}_{\text{pred}}$, where $\mathbf{d}$ is the background-subtracted data, Δ is the detector response matrix, and $\boldsymbol{\varphi}_{\text{pred}}$ is a truth-level model prediction. The range-projected χ^2 is defined as

$$\chi^2_{\text{proj}} = \boldsymbol{\delta}^{\mathsf{T}}\,(P_{\text{reco}}\,C_{\text{total}}\,P_{\text{reco}})^{+}\,\boldsymbol{\delta}, \tag{3.4}$$

with the superscript $^{+}$ denoting the Moore-Penrose pseudoinverse [25].

The pseudoinverse can be avoided entirely by reducing the problem to r dimensions. One can verify that the pseudoinverse of $P_{\text{reco}}\,C_{\text{total}}\,P_{\text{reco}}$ takes the closed form

$$(P_{\text{reco}}\,C_{\text{total}}\,P_{\text{reco}})^{+} = V_r\,(V_r^{\mathsf{T}}\,C_{\text{total}}\,V_r)^{-1}\,V_r^{\mathsf{T}}. \tag{3.5}$$

Substituting into eq. (3.4) and defining the reduced residual $\tilde{\boldsymbol{\delta}} = V_r^{\mathsf{T}}\,\boldsymbol{\delta}$ and the reduced covariance matrix $\tilde{C}_{\text{total}} = V_r^{\mathsf{T}}\,C_{\text{total}}\,V_r$, the test statistic becomes

$$\chi^2_{\text{proj}} = \tilde{\boldsymbol{\delta}}^{\mathsf{T}}\,\tilde{C}^{-1}_{\text{total}}\,\tilde{\boldsymbol{\delta}}, \tag{3.6}$$

which involves only a standard inversion of the $r \times r$ matrix $\tilde{C}_{\text{total}}$. Under the null hypothesis, χ^2_{proj} follows a χ^2 distribution with $r = N_{\text{bins}} - N_{\text{null}}$ degrees of freedom.

Without the projection, the χ^2 test statistic takes the same numerical value whether computed in reconstructed or unfolded space, as the bias in the measurement introduced by a regularized unfolding cancel identically when compared to model predictions that are left-multiplied by the regularization matrix $A_C \equiv U\,\Delta$ [16]. The introduction of the projector breaks this equivalence when the projection is performed after unfolding. In unfolded-space the residual is defined as $\mathbf{r} = \hat{\boldsymbol{\varphi}} - A_C\,\boldsymbol{\varphi}_{\text{pred}}$, where $\hat{\boldsymbol{\varphi}} = U\mathbf{d}$ is the unfolded background-subtracted data. Using the propagated basis $W_r = U\,V_r$, one can construct an analogous test statistic in unfolded space as

$$\chi^2_{\text{unf}} = \tilde{\mathbf{r}}^{\mathsf{T}}\,\tilde{V}^{-1}_{\text{total}}\,\tilde{\mathbf{r}}, \tag{3.7}$$

where $V_{\text{total}} = U\,C_{\text{total}}\,U^{\mathsf{T}}$ is the total covariance matrix in unfolded space, and $\tilde{\mathbf{r}} = W_r^{\mathsf{T}}\,\mathbf{r}$ and $\tilde{V}_{\text{total}} = W_r^{\mathsf{T}}\,V_{\text{total}}\,W_r$ are the reduced residual and reduced covariance matrix in unfolded space. This quadratic form is not numerically equal to eq. (3.6). Expressing it in terms of reconstructed-space quantities yields

$$\chi^2_{\text{unf}} = \boldsymbol{\delta}^{\mathsf{T}}\,U^{\mathsf{T}}U\,V_r\,(V_r^{\mathsf{T}}\,U^{\mathsf{T}}U\,C_{\text{total}}\,U^{\mathsf{T}}U\,V_r)^{-1}\,V_r^{\mathsf{T}}\,U^{\mathsf{T}}U\,\boldsymbol{\delta}, \tag{3.8}$$

which reduces to eq. (3.6) only when $U^{\mathsf{T}}U = I$, i.e. when the unfolding matrix is orthogonal.

2026 JINST 21 P11017

The χ^2 equivalence is restored if the projection is applied before unfolding. Projecting the reconstructed-space residual and covariance first, then propagating through the unfolding matrix, gives

$$\begin{aligned}(U\,P_{\text{reco}}\,\boldsymbol{\delta})^{\mathsf{T}}\,(U\,P_{\text{reco}}\,C_{\text{total}}\,P_{\text{reco}}\,U^{\mathsf{T}})^{+}\,(U\,P_{\text{reco}}\,\boldsymbol{\delta}) &= \tilde{\boldsymbol{\delta}}^{\mathsf{T}}\,W_r^{\mathsf{T}}\,(W_r\,\tilde{C}_{\text{total}}\,W_r^{\mathsf{T}})^{+}\,W_r\,\tilde{\boldsymbol{\delta}} \\ &= \tilde{\boldsymbol{\delta}}^{\mathsf{T}}\,G\,(G\,\tilde{C}_{\text{total}}\,G)^{-1}\,G\,\tilde{\boldsymbol{\delta}} \\ &= \chi^2_{\text{proj}},\end{aligned} \tag{3.9}$$

where the symmetric matrix $G = W_r^{\mathsf{T}}\,W_r$ acts as the unfolding matrix of the reduced system, and cancels similarly to the unprojected case. In the approach of eq. (3.8), the full residual, including the null-space components, is first rotated by U into directions that are no longer orthogonal to the range. The unfolded-space projector then picks up this null-space contamination. Projecting first in reconstructed space avoids this by removing the null directions before unfolding can mix them with the physical subspace.

For this reason, eq. (3.6) is the recommended definition of the range-projected χ^2.

3.2 Statistical validity

Under the null hypothesis, the residual $\boldsymbol{\delta}$ is drawn from a multivariate Gaussian distribution with mean zero and covariance C_{total}. The reduced residual $\tilde{\boldsymbol{\delta}} = V_r^{\mathsf{T}}\,\boldsymbol{\delta}$ is a linear transformation of a Gaussian vector, therefore Gaussian itself with mean zero and covariance $\tilde{C}_{\text{total}} = V_r^{\mathsf{T}}\,C_{\text{total}}\,V_r$. Provided that $\tilde{C}_{\text{total}}$ is positive definite, standard theory of Gaussian quadratic forms guarantees that $\tilde{\boldsymbol{\delta}}^{\mathsf{T}}\tilde{C}_{\text{total}}^{-1}\tilde{\boldsymbol{\delta}} \sim \chi^2(r)$. The positive definiteness of $\tilde{C}_{\text{total}}$ requires C_{total} to have nonzero variance along every direction in the range of the projector P_{reco}, which is generally satisfied. This can be verified by checking that all eigenvalues of $\tilde{C}_{\text{total}}$ are strictly positive.

The validity of the range-projected χ^2 can also be understood from the general theory of quadratic forms of correlated normal variables. Theorem 9.2.2 of ref. [20] establishes that if $\mathbf{Y} \sim \mathcal{N}(0, \Sigma)$ with $\text{rank}(\Sigma) = k$, then $\mathbf{Y}^{\mathsf{T}}\Sigma^{+}\mathbf{Y} \sim \chi^2(k)$ for any choice of generalized inverse Σ^{+}. Assuming the null hypothesis, the projected residual $P_{\text{reco}}\boldsymbol{\delta}$ is normally distributed with covariance $P_{\text{reco}}C_{\text{total}}P_{\text{reco}}$, which has rank r by construction. The formulation of eq. (3.4) with the Moore-Penrose pseudoinverse is a specific instance of this general result. The key distinction from purely numerical approaches to handling singular covariance matrices is that the rank is not determined by an eigenvalue threshold applied to the covariance, but is predicted exactly via the combination matrix Ω from the structure of the measurement.

3.3 Relation to alternative approaches

As a statistical object, the range-projected χ^2 is an instance of the standard treatment of quadratic forms with singular covariance matrices described in section 3.2. It is a generalized inverse restricted to the range of the covariance, distributed as a χ^2 with degrees of freedom equal to the rank [20]. Dropping zero-variance directions from a χ^2 is common practice in fits with degenerate covariance matrices. The contribution of this work is not the statistic itself but the recognition that the rank deficiency arises specifically and predictably from the shared-event structure, together with a prescription to compute its dimension a priori from the binning geometry. We discuss here how this differs from related approaches.

The most direct alternative is to regularize the covariance matrix numerically, for instance by performing a singular value decomposition and discarding eigenvalues below a fixed threshold

2026 JINST 21 P11017

relative to the largest. This is reliable only when a clean gap separates the null eigenvalues from the physical ones, which a realistic measurement does not guarantee. The null eigenvalues are lifted to small but nonzero values by any covariance component that does not preserve the event sharing; in cross-section measurements typically the response-matrix treatment of the interaction uncertainties. These lifted directions need not be separated from the smallest physical eigenvalues by a clean gap, so a threshold must then be chosen to sit between them. Setting it too low retains the inflating directions discussed next in section 4.3, while setting it too high discards directions carrying real information. The combination-matrix approach avoids this choice. In this case the rank is fixed by Ω from the measurement geometry, independently of the systematics-contaminated spectrum of C_{total}. A threshold is still used to identify the null space, but it is applied to $\Omega^{\mathsf{T}}\Omega$, where the gap is clean by construction, rather than to C_{total} where systematics have filled it in.

The projection is distinct from, and complementary to, the regularization performed in unfolding. Methods such as Wiener-SVD [16] address the ill-posedness of inverting the response matrix, producing the regularization matrix A_C that is applied to truth-level predictions before comparison. They do not address the event-sharing rank deficiency, which is a property of the statistical covariance and survives into unfolded space regardless of the unfolding method. The two operate on different sources of degeneracy and are used together. The event generator study of section 5 unfolds each block with Wiener-SVD and applies the range projection on top.

A more complete alternative is to extract the cross section directly from a binned maximum likelihood fit, in which the truth-level bins are free parameters of interest and the systematic uncertainties enter as nuisance parameters with prior constraints (see, e.g., ref. [26]). The likelihood is evaluated in reconstructed space, where the inter-distribution correlations are carried by the shared events and the common nuisance parameters of the forward model. Therefore, the likelihood method never forms an inverted covariance over the bin counts, so the constrained directions are absorbed automatically and the goodness of fit can be assessed without confronting the rank deficiency at all. This approach is also more robust to the imperfect-sharing scenarios discussed in section 6, where the constraints are only approximate. Its requirements are correspondingly greater: it needs the full forward model and nuisance parameterization, which are generally available only within the originating group. In this regard, the range-projected χ^2 requires far less information.

For a group producing a blockwise-unfolded measurement, the range-projected χ^2 is a natural global goodness-of-fit statistic. It is computed directly from the covariance and response matrices already used in the extraction, and is evaluated in reconstructed space as recommended in section 3.1. It applies equally to an already-published measurement, as performed by generator comparison and tuning frameworks [27]. A standard data release reporting the unfolded data, unfolded covariance, and the block-diagonal unfolding matrix is sufficient to recover the reconstructed-space data and covariance by inversion; the reconstructed-space test then additionally requires the detector response matrix to forward-fold a model prediction. Where the response matrix is not available, the projection can instead be performed in unfolded space using the regularization matrix A_C, which is an increasingly standard release product. This unfolded-space test is less exact, as the non-orthogonal unfolding mixes a residual null-space contamination into the physical subspace, but it still removes the bulk of the inflation. The two approaches are complementary: the likelihood fit is more general when the complete model is available, while the projection provides a lightweight, a priori test from published results.

2026 JINST 21 P11017

4 Toy example

The formalism developed previously is now illustrated with a toy model that captures the essential features of a multi-distribution neutrino cross-section measurement: a two-dimensional underlying physics process, a simple detector model, and a multi-block binning scheme in which events are shared across distributions. The toy is designed to be minimal but non-trivial. A three-block structure produces a rank-deficient statistical covariance matrix, and a confidence-region coverage study demonstrates the consequences of ignoring or correctly treating the resulting null directions.

4.1 Model and measurement setup

The toy model mimics a muon neutrino charged-current quasi-elastic-like interaction, producing a muon characterized by its momentum p_μ and scattering angle θ_μ. The joint truth-level differential cross section is parametrized as

$$f(p_\mu, \cos\theta_\mu) = p_\mu^k \, e^{-k\, p_\mu/p_0} \, e^{\kappa \cos\theta_\mu} \, (1 + \alpha \, p_\mu \cos\theta_\mu), \tag{4.1}$$

where the four parameters (k, p_0, κ, α) control the shape of the momentum spectrum, the level of forward peaking, and the strength of the momentum-angle correlation, respectively. A central-value (CV) model with $(k, p_0, \kappa, \alpha) = (3.0, 0.2, 2.0, 0.3)$ is used to build the response matrix and estimate the covariance, while a "fake" model with $(k, p_0, \kappa, \alpha) = (3.3, 0.2, 1.8, 0.35)$ serves as the true underlying distribution from which pseudo-data are drawn. Truth events are generated via rejection sampling over $p_\mu \in [0.0, 1.2]\,\mathrm{GeV}/c$ and $\cos\theta_\mu \in [-1.0, 1.0]$.

A simple detector model introduces two different effects. Momentum is smeared multiplicatively, $p_{\mathrm{reco}} = p_{\mathrm{true}}(1 + \delta_p)$ with $\delta_p \sim \mathcal{N}(0, \sigma_p)$ and $\sigma_p = 0.05$, while the angle is smeared additively, $\theta_{\mathrm{reco}} = \theta_{\mathrm{true}} + \delta_\theta$ with $\delta_\theta \sim \mathcal{N}(0, \sigma_\theta)$ and $\sigma_\theta = 1^\circ$. A sigmoid efficiency function $\varepsilon(p_{\mathrm{reco}}) = [1 + e^{-(p_{\mathrm{reco}} - p_{\mathrm{thr}})/w}]^{-1}$ with threshold $p_{\mathrm{thr}} = 0.10\,\mathrm{GeV}/c$ and width $w = 0.02\,\mathrm{GeV}/c$ models the turn-on of the event acceptance at low momentum.

The analysis uses three bin blocks in reconstructed space: a p_μ marginal (5 bins), a $\cos\theta_\mu$ marginal (5 bins), and a two-dimensional distribution sliced in $\cos\theta_\mu$ (3 slices, each with 4 bins in p_μ), for a total of $N_{\mathrm{bins}} = 22$. The bin edges of the 2D block do not align exactly with those of the 1D marginals, which is the usual situation in real analyses. A CV sample of $N_{\mathrm{MC}} = 500{,}000$ truth events is generated, processed through the detector model, and used to construct the response matrix Δ. Pseudo-data samples are generated by drawing $N_{\mathrm{data}} = 5{,}000$ events from the fake model.

Systematic uncertainties are estimated following the response-matrix approach described in section 2.4. Interaction-model variations are obtained by reweighting the CV sample to six alternative model configurations, each producing a varied response matrix Δ_s. The shift vector $\boldsymbol{\xi}_s = (\Delta_s - \Delta)\,\boldsymbol{\varphi}$ captures the effect on the predicted reconstructed-space distribution. Detector variations are obtained by generating eight independent samples with perturbed detector parameters (σ_p, σ_θ, p_{thr}, and w), each producing a new response matrix from which a shift vector is computed similarly. The total systematic covariance matrix is $C_{\mathrm{syst}} = N_{\mathrm{var}}^{-1} \sum_s \boldsymbol{\xi}_s \boldsymbol{\xi}_s^{\mathsf{T}}$, scaled to data statistics. The statistical covariance is computed using eq. (2.2), assuming unit event weights, which captures the inter-block correlations from event sharing.

Figure 1 shows the predicted event distributions for the CV and fake models in all three blocks of the toy analysis, together with a pseudo-data sample distribution and the systematic uncertainty bands.

2026 JINST 21 P11017

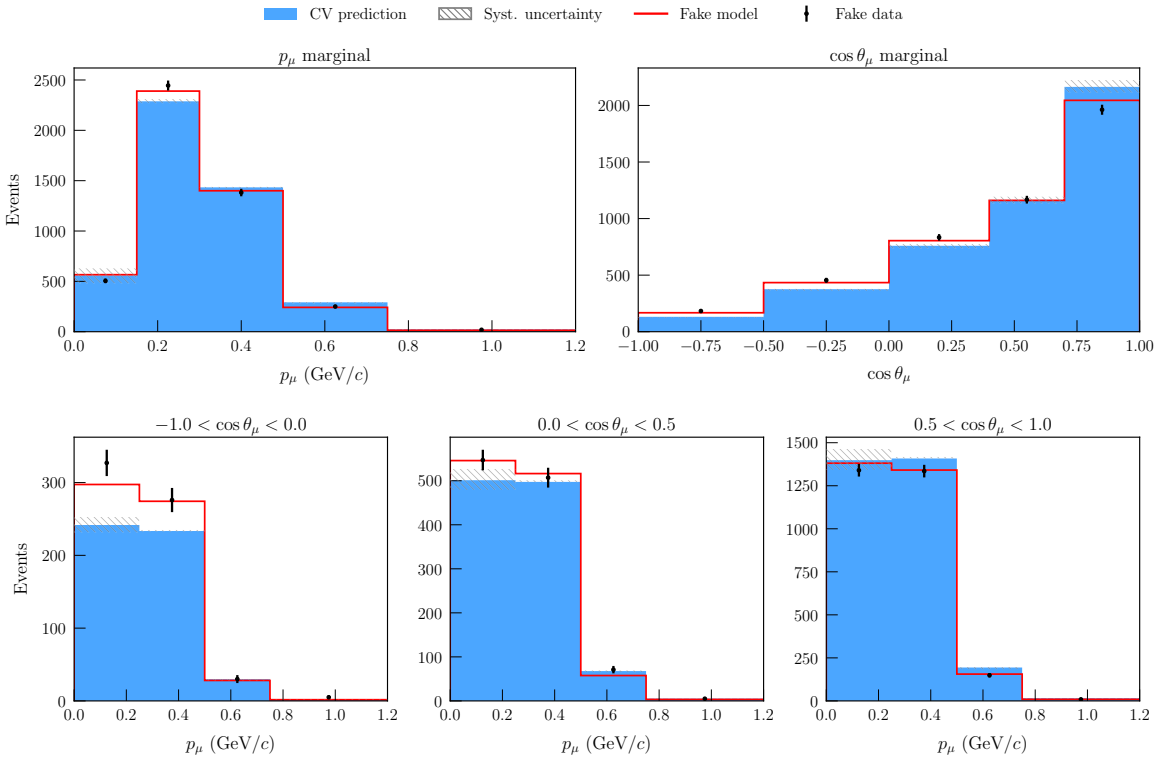


Figure 1. Predicted event distributions in the toy analysis blocks for the CV model (filled histograms) and the fake model (solid lines), together with a pseudo-data realization (points with statistical error bars). The shaded bands indicate the systematic uncertainty on the CV prediction. The top row shows the p_μ and $\cos\theta_\mu$ marginals; the bottom row shows the three $\cos\theta_\mu$ slices of the 2D block, each binned in p_μ.

4.2 Null-space structure

The structural combination matrix Ω_{struct} is constructed from the binning configuration of the three blocks, following the prescription from section 2.3. For this bin scheme, Ω_{struct} is rank 17, yielding 5 structural null directions. The combination matrix Ω computed from the MC event list also has rank 17, confirming that all 5 null directions are structural. No additional kinematic constraints arise for this choice of physics model and binning.

The eigenvalue spectrum of C_{stat} confirms this. It shows 5 eigenvalues that are numerically zero, separated from the 17 nonzero eigenvalues by a spectral gap of several orders of magnitude. Adding systematic uncertainties constructed via the response-matrix approach described in section 2.4 lifts the 5 null eigenvalues, as a consequence of the inconsistent folding through varied response matrices for each universe, making $C_{\text{total}} = C_{\text{stat}} + C_{\text{syst}}$ formally full rank. However, as shown in figure 2, the lifted eigenvalues remain considerably smaller than the bulk of the spectrum. Including them in the test statistic introduces spurious contributions to the χ^2, as demonstrated next.

4.3 Goodness-of-fit validation

The fidelity of the projected and unprojected χ^2 test statistics is assessed with a toy Monte Carlo exercise at the true hypothesis point. A high-statistics simulation of the fake model is used to estimate the event distribution on a fine reconstructed-space grid, defined by the union of all bin

2026 JINST 21 P11017

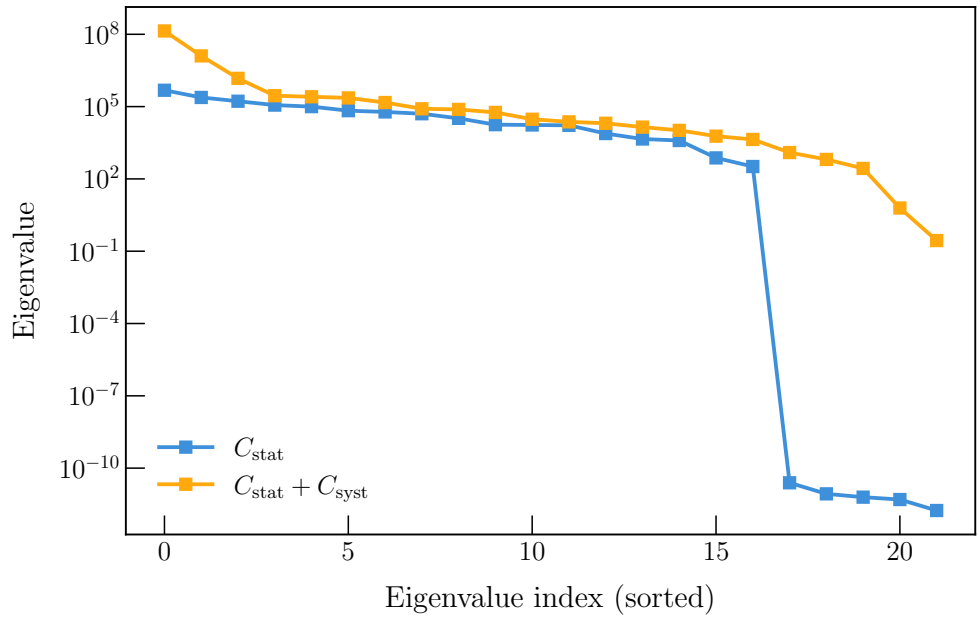


Figure 2. Eigenvalue spectra of the statistical covariance matrix C_{stat} (blue circles) and the total covariance matrix $C_{\text{total}} = C_{\text{stat}} + C_{\text{syst}}$ (orange squares), for the toy analysis, sorted in descending order.

edges in p_μ and $\cos\theta_\mu$ across the three blocks, corresponding to the common refinement described in section 2.3. Each toy experiment draws a Poisson-fluctuated total event count and distributes them across the fine grid via multinomial sampling. These events are then projected onto the three analysis blocks, preserving the inter-block correlations from event sharing. Systematic fluctuations are added as correlated Gaussian throws from the Cholesky decomposition of C_{syst}. The χ^2 is evaluated in reconstructed space, comparing each toy to the forward-folded prediction at the true hypothesis, for both the unprojected statistic, using the direct inverse of C_{total}, and the range-projected statistic of eq. (3.4).

Figure 3 shows the results from 10^6 toy experiments. The projected χ^2 statistic closely follows the expected $\chi^2(17)$ distribution, with a mean of 17.23 and a standard deviation of 6.01, close to the expected values of 17 and $\sqrt{2\times 17} \approx 5.83$. In contrast, the unprojected statistic is shifted to significantly higher values with a mean of 30.59, and a standard deviation of 8.96, far above the expectation for a $\chi^2(22)$ distribution of 22 and $\sqrt{2\times 22} \approx 6.63$. The right panel shows the corresponding p-value distributions. The projected statistic yields an approximately uniform distribution, with a false rejection rate of 5.7% at the 5% significance level. The unprojected statistic produces a false rejection rate of 32.4%, more than six times the expected value. An analyzer using the unprojected statistic for goodness-of-fit testing would falsely reject the correct model roughly one third of the time.

Writing the residual vector at the true hypothesis point as $\boldsymbol{\delta} = \boldsymbol{\delta}_{\text{rand}} + \mathbf{b}$, where $\boldsymbol{\delta}_{\text{rand}}$ has covariance C_{total} and $\mathbf{b} \equiv E[\boldsymbol{\delta}]$ is the deterministic data-prediction residual (estimated as the mean over the toy throws), and diagonalizing $C_{\text{total}} = \sum_j \lambda_j \mathbf{v}_j \mathbf{v}_j^{\mathsf{T}}$, the unprojected statistic decomposes as

$$\chi^2_{\text{unproj}} = \sum_j \frac{(\boldsymbol{\delta}\cdot\mathbf{v}_j)^2}{\lambda_j}, \tag{4.2}$$

2026 JINST 21 P11017

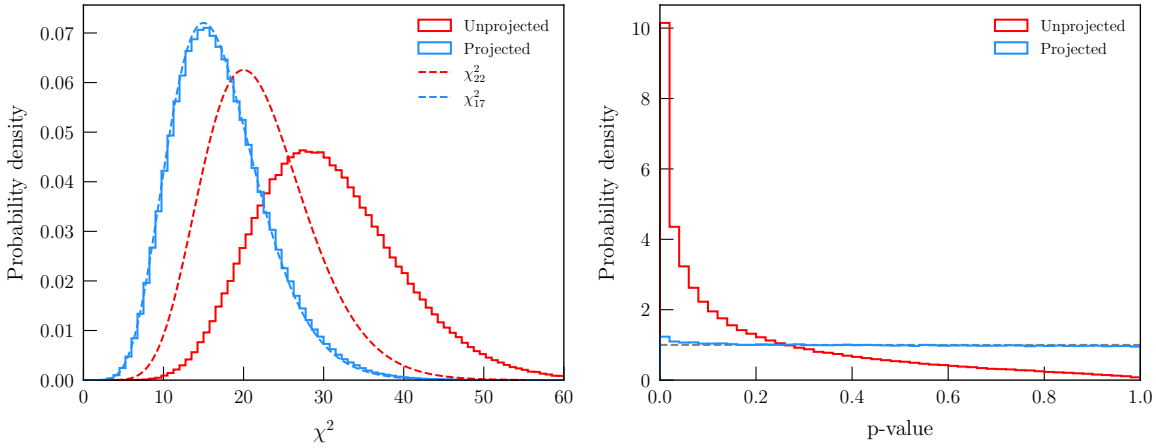


Figure 3. Goodness-of-fit validation for toy experiments generated at the true hypothesis. Left: distributions of the unprojected (red) and range-projected (blue) χ^2 statistics, compared with the expected $\chi^2(22)$ and $\chi^2(17)$ distributions (dashed curves). Right: corresponding p-value distributions, with the false rejection rates at the 5% significance level indicated.

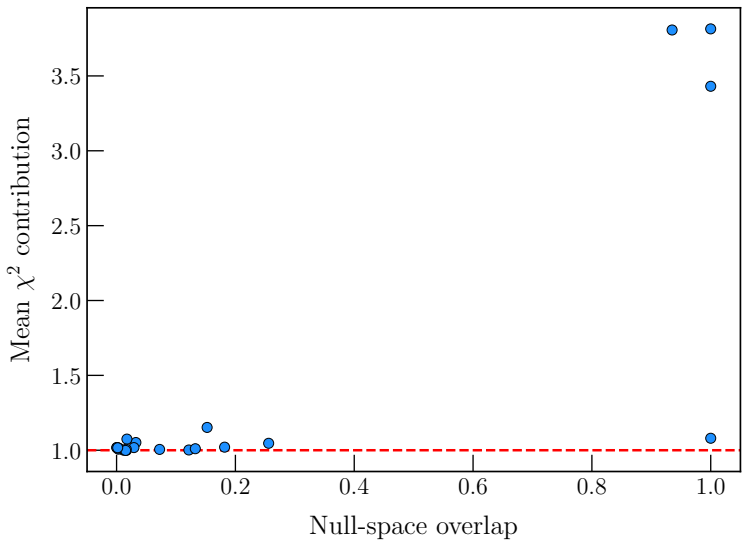


Figure 4. Mean contribution to the unprojected χ^2 from each eigendirection of C_{total}, averaged over the toy ensemble, versus the null-space overlap O_j from eq. (4.4).

with the expected value of each of these contributions given by

$$\mathrm{E}\left[(\boldsymbol{\delta}\cdot\mathbf{v}_j)^2/\lambda_j\right] = 1 + \frac{(\mathbf{b}\cdot\mathbf{v}_j)^2}{\lambda_j}. \tag{4.3}$$

Each eigendirection contributes as a noncentral $\chi^2(1)$ with noncentrality parameter $\psi_j = (\mathbf{b}\cdot\mathbf{v}_j)^2/\lambda_j$. On the directions overlapping the null space of C_{stat} the eigenvalue λ_j is the small lifted systematic variance, so any residual projecting onto them is amplified into a contribution far exceeding one.

To examine how the excess in the unprojected statistic is distributed across eigendirections, we quantify how much each eigenvector of C_{total} overlaps the null space of C_{stat}. Using the projector onto

2026 JINST 21 P11017

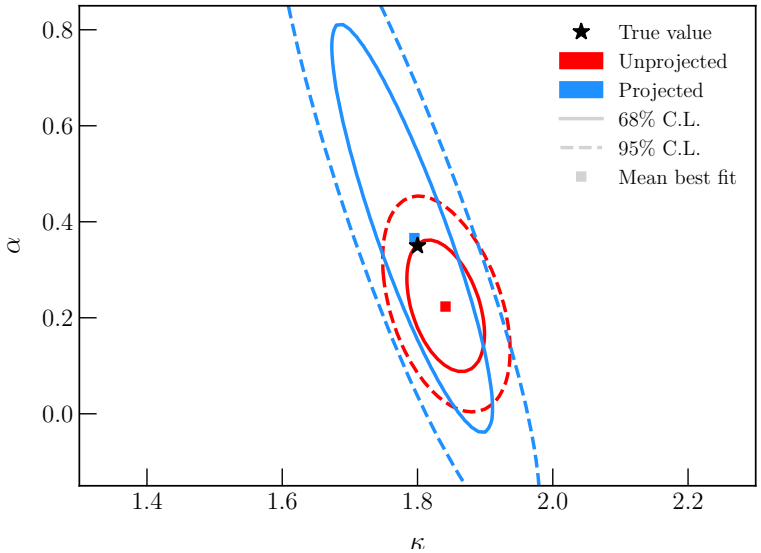


Figure 5. Confidence regions in the (κ, α) plane from the Asimov dataset, showing the 68% and 95% CL contours for the range-projected (blue) and unprojected (red) test statistics. The star marks the true parameter values; the squares indicate the mean best-fit points across 10^6 toys.

the non-null subspace P_{reco} from eq. (3.2), we define the overlap coefficient as

$$O_j = \mathbf{v}_j^{\mathsf{T}} \left(I - P_{\text{reco}}\right) \mathbf{v}_j, \tag{4.4}$$

which equals one for a direction lying entirely in the null space and zero for one in the range. Figure 4 shows the mean contributions to the χ^2 test statistic across the toy ensemble against O_j. The 17 non-null directions cluster around $O_j \approx 0$ with contributions consistent with unity. Across the ensemble the random part of the residual has the correct variance along every eigendirection, with the empirical variance matching λ_j to better than 1% on both the lifted and the range directions; the excess is therefore entirely noncentrality. The directions carrying the largest contributions are all those with substantial null-space overlap, where the small lifted eigenvalues amplify any residual projecting onto them. Summing over all directions, the noncentrality reproduces the observed excess, $\sum_j \psi_j \approx 8.6$, accounting for the unprojected mean of 30.6.

4.4 Parameter estimation

The impact of using the projected test statistic on parameter estimation is assessed with a two-dimensional confidence-region study. We scan the two parameters (κ, α) that control the angular shape and the momentum-angle correlation, holding the momentum-spectrum parameters (k, p_0) fixed at their fake-model values, as these are well constrained by the p_μ marginal and largely decoupled from the null-space directions of interest. At each grid point the truth-level prediction is forward-folded through the nominal response matrix and compared to the pseudo-data, and the profiled test statistic $\Delta\chi^2 = \chi^2(\kappa, \alpha) - \chi^2_{\text{min}}$ is evaluated for both the unprojected and projected cases. Confidence regions are defined by the standard thresholds $\Delta\chi^2 < F^{-1}_{\chi^2_n}(\text{CL})$ for n scanned (or profiled) parameters, and their frequentist coverage is then measured directly from the toy ensemble.

Figure 5 shows the 68% and 95% confidence level contours from the Asimov dataset for both test statistics. The projected contours are correctly centered on the true parameter values, while the

2026 JINST 21 P11017

unprojected contours are shifted, with the 68% region excluding the true point. The unprojected best fit values are biased, particularly in α. The mean best-fit value across toys is $\hat{\alpha} = 0.22$, far from the true value of 0.35, while the projected statistic recovers $\hat{\alpha} = 0.37$. The bias in κ is milder ($\hat{\kappa} = 1.84$ for unprojected, $\hat{\kappa} = 1.80$ for projected, compared to the true 1.80). This indicates that α has greater overlap with the null directions.

The frequentist coverage confirms these observations. At the 68.3% confidence level, the projected statistic achieves nominal coverage in both the joint two-dimensional region (68.2%) and the one-dimensional profiled intervals (68.1% for κ, 67.9% for α). The unprojected statistic suffers severe undercoverage: 36.5% in the joint region, 44.6% for κ, and 34.1% for α. The undercoverage is worst for α, consistent with its stronger sensitivity to the null-space contamination. The projected contours are larger than the unprojected ones, reflecting the genuine sensitivity of the measurement once the ill-constrained directions are properly removed.

5 Event generator study

The toy study presented in the previous section validates the usage of the range-projected χ^2 test statistic in a controlled setting, with an analytic cross-section model and known input parameters. Now, we demonstrate that the same pathology arises, and the same fix applies, when working with a realistic measurement setup. In this case, we use a realistic neutrino flux prediction, neutrino event generators, cross-section systematic uncertainties, and unfolding strategy. The latter is a key addition over the toy study, as it allows us to verify numerically the equivalence between the results in reconstructed and unfolded space, established in section 3.1. In addition, we show that the projected test statistic recovers its discriminating power between generator predictions, correctly rejecting the wrong model while the unprojected statistic does not.

5.1 Model and measurement setup

The central-value model is built from GENIE v3.6.2 [28] using the `AR23_20i_00_000` tune on an ^{40}Ar target. The fake-data model, which plays the role of the unknown true cross section, is NuWro v25.11 [29] on the same target. Both generators are run with the SBND BNB ν_μ flux over the energy range 0–10 GeV. Since both predictions use the same flux, flux systematic uncertainties cancel in this comparison and are not considered; a measurement comparing against an external flux prediction would include them as an additional systematic.

Events are selected using a topological CC1μ1p0π definition applied identically to both generators. We require exactly one muon, exactly one proton above 40 MeV kinetic energy, and no charged or neutral mesons above 50 MeV in the final state. This topological selection captures CCQE, CCMEC, and CCRES events where the pion was absorbed in final-state interactions.

The detector model is the same as in the toy study. For each candidate event, we apply a Gaussian momentum and angular smearing ($\sigma_p = 5\%$, $\sigma_\theta = 1°$), and a sigmoid efficiency threshold at $p_{\text{thr}} = 100$ MeV with a width $w = 20$ MeV to the primary muon.

The analysis binning follows the same three-block scheme as the toy study: a p_μ marginal, a $\cos\theta_\mu$ marginal, and a two-dimensional block sliced in $\cos\theta_\mu$. In the present sample, however, the highest-momentum bin of the most backward $\cos\theta_\mu$ slice of the two-dimensional block receives no selected events, as the CC1μ1p0π topology has negligible contributions from high-momentum

2026 JINST 21 P11017

backward-going muons. This empty bin is merged with its neighbor, leaving $5 + 5 + (3 + 4 + 4) = 21$ bins across the three blocks.

One million GENIE events are generated and used to construct the response matrix and estimate the systematic uncertainties. One million NuWro events are generated separately; 5,000 of these are drawn as pseudo-data, with the remainder serving as the high-statistics truth prediction.

5.2 Systematic uncertainties and unfolding

Interaction-model systematic uncertainties are estimated using the GENIE Reweight framework [28], which provides per-event weights at $\pm 1\sigma$ for each systematic knob. Fifteen knobs are varied, spanning the dominant sources of uncertainty for the CC1μ1p0π topology: four z-expansion coefficients for the CCQE axial form factor [30], RPA screening and Pauli blocking corrections [31], the resonance axial and vector masses and overall normalization [32], the CCMEC normalization, two non-resonant single-pion background parameters, and three final-state interaction parameters controlling the pion mean free path, absorption fraction, and charge-exchange fraction [33]. The interaction-model systematic covariance matrices are constructed via the response-matrix approach described in section 2.4, applying the varied $\Delta_s^{'}$ on the fixed CV prediction. Detector uncertainties are estimated as before, using $\pm 1\sigma$ variations of the four detector parameters (σ_p, σ_θ, p_{thr}, w).

Each of the three distribution blocks is unfolded independently using the Wiener-SVD approach [16], with the response matrix and total data covariance computed block-by-block. The per-block unfolding matrices are assembled into the block-diagonal matrix U, and the regularization matrix A_C is also constructed from the per-block results. The full covariance is propagated through U as $V_{\text{total}} = U\, C_{\text{total}}\, U^{\mathsf{T}}$, preserving the cross-block correlations from event sharing. Truth-level predictions are regularized with A_C before comparison to the unfolded data.

5.3 Results

The structural combination matrix Ω_{struct}, built from the three-block binning scheme, has rank 17, yielding $N_{\text{null}}^{\text{struct}} = 21 - 17 = 4$ structural null directions. The full combination matrix Ω, computed from the GENIE event list, has rank 16 and therefore reveals one kinematic null direction, bringing the total to $N_{\text{null}} = 5$. This is distinct from the bin merging mentioned in section 5.1. In that case we merged a single bin of zero occupancy, whereas the kinematic null is a linear constraint among the populated bins. It arises because two cells of the common refinement corresponding to high-momentum backward-going muons are not populated by any selected events, a region where the CCQE-like cross section is strongly suppressed. The marginal bins themselves remain populated; only the fine-grained intersection of the high-momentum and backward-angle regions is empty. The projected test statistic therefore has $r = 16$ degrees of freedom.

Figure 6 shows the reconstructed-space event distributions in all three blocks. The GENIE CV prediction and its systematic uncertainty band are compared to the NuWro fake-data realization and the NuWro prediction. The GENIE systematic band does not fully cover the NuWro prediction in several bins, reflecting genuine model differences that the χ^2 test should be able to identify.

Figure 7 shows the corresponding unfolded distributions, with the GENIE and NuWro truth-level predictions smeared through the regularization matrix A_C. Per-block χ^2 and p-values, computed using only the block-diagonal covariance, are annotated in each panel. Comparisons between the unfolded NuWro fake data and the true NuWro model prediction yield p-values mostly above 0.32

2026 JINST 21 P11017

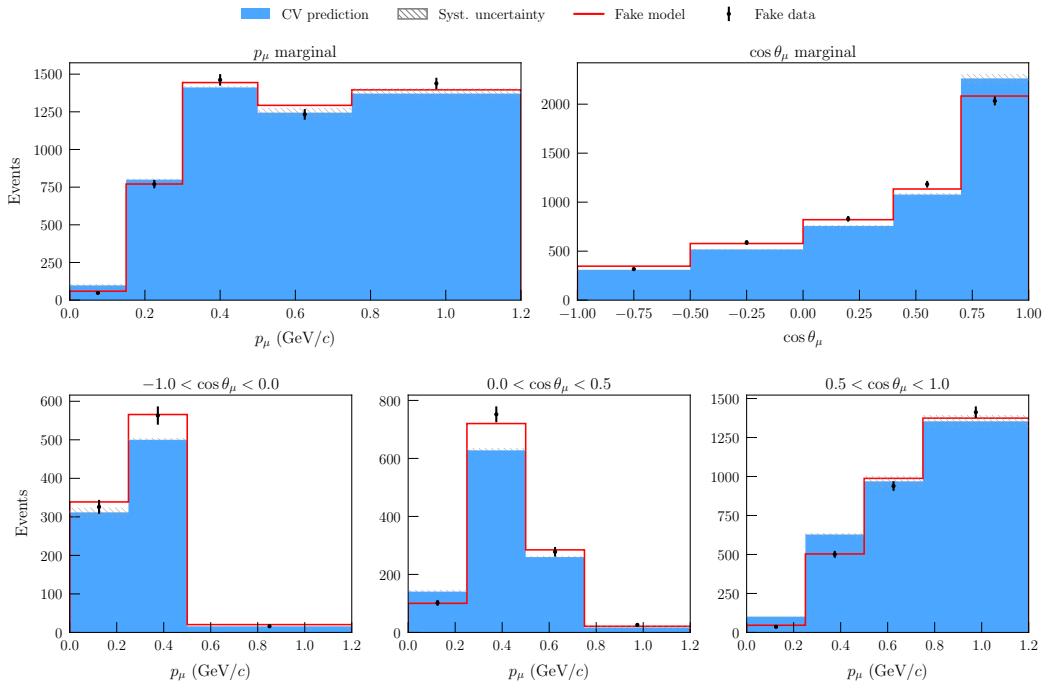


Figure 6. Reconstructed-space event distributions for the event generator study. The GENIE CV prediction (filled histograms) and systematic uncertainty bands (shaded) are compared to the NuWro prediction (red line) and a 5,000-event NuWro pseudo-data realization (points with statistical error bars). The top row shows the p_μ and $\cos\theta_\mu$ marginals; the bottom row shows the three $\cos\theta_\mu$ slices of the 2D block.

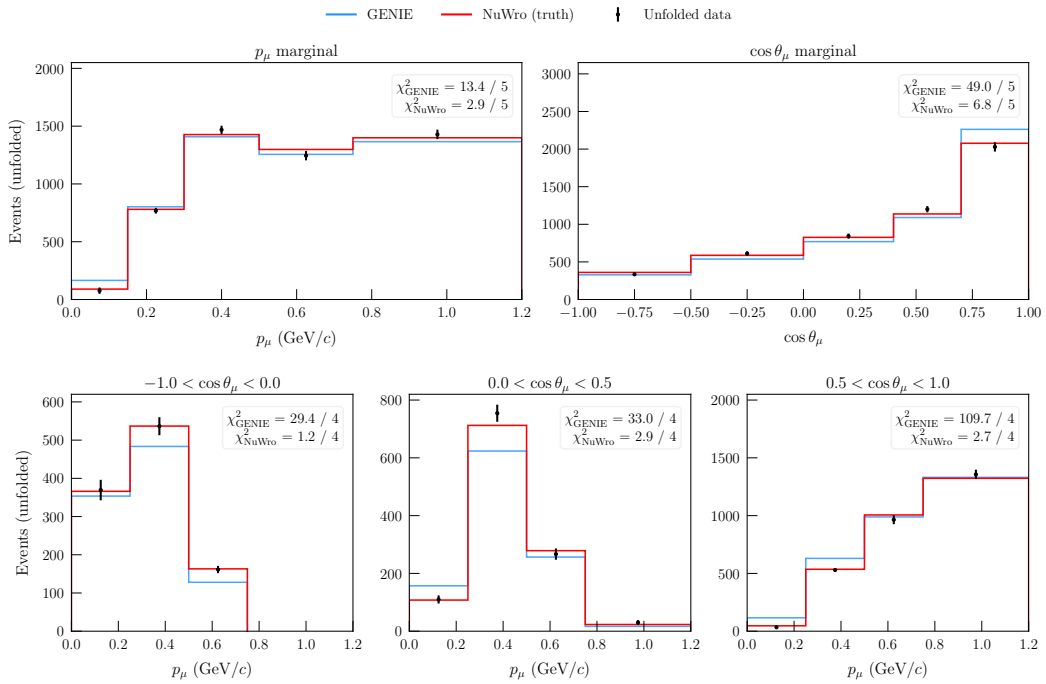


Figure 7. Unfolded event distributions for the event generator study. The GENIE (blue) and NuWro (red) truth-level predictions are regularized by the regularization matrix A_C for comparison with the unfolded fake data (points with propagated uncertainties). Per-block χ^2 values are annotated in each panel.

2026 JINST 21 P11017

Table 1. Global χ^2 values for the NuWro fake data compared to the GENIE and NuWro truth-level predictions. The projected test statistic has $N_{\text{bins}} - N_{\text{null}} = 16$ degrees of freedom; the unprojected has $N_{\text{bins}} = 21$.

Test statistic	GENIE	NuWro
Unprojected χ^2/ndof	252.4 / 21	121.9 / 21
Projected χ^2/ndof	154.4 / 16	15.7 / 16

and all above 0.05, confirming that the unfolding correctly recovers the true distribution within uncertainties at the per-block level

Table 1 summarizes the global χ^2 results. The unprojected statistic yields 121.9/21 for the NuWro prediction. This results in a disproportionate rejection of the correct model, entirely driven by the null-space contamination. The projected statistic gives 15.7/16, consistent with the expected $\chi^2(16)$ distribution and confirming that the method is correctly calibrated. This was further validated with 10^5 pseudo-experiments, which confirmed that the projected χ^2 distribution follows $\chi^2(16)$ while the unprojected distribution does not follow $\chi^2(21)$, consistent with the toy study results of section 4.3.

The projected statistic retains full discriminating power against the wrong model. The comparison to GENIE yields 154.4/16, a strong rejection indicating that the GENIE uncertainty budget cannot accommodate the NuWro prediction. Without the projection, both the correct and incorrect models are rejected, and the test provides no useful discrimination.

Finally, the projected χ^2 was verified to give identical numerical values whether computed in reconstructed or unfolded space, for both the GENIE and NuWro predictions. This confirms the equivalence established in section 3.1 and validates the practical implementation of the projection in the presence of unfolding.

6 Limitations and scope

The proposed method rests on the assumption that the event-sharing structure is exact. Every selected event populates exactly one bin in each distribution, so the constraints $\Omega\mathbf{x} = 0$ hold identically and the null space of C_{stat}, built according to eq. (2.2), coincides with that of Ω. Under this assumption the measurement has no statistical power along a null direction, so projecting it out discards no statistical information, as argued in section 3. Therefore, our recommendation is that multi-distribution measurements be designed to preserve exact event sharing across all blocks, so the constrained directions can be projected out cleanly. The remainder of this section examines what happens when this requirement is relaxed, how imperfect sharing manifests, and how it might be handled.

In real measurements, this assumption is sometimes only approximate. Distributions extracted from a common sample may nonetheless use different phase-space limits, apply distribution-dependent background subtractions, or define their signal components differently, so that the populations entering different blocks are not strictly identical. Consider the simplest constraint that two block totals be equal, encoded by the direction $\mathbf{x}$ with entries $+1$ on the bins of one block and -1 on those of another. Its statistical variance is given by

$$\mathbf{x}^{\mathsf{T}} C_{\text{stat}}\, \mathbf{x} = \sum_e w_e^2\, (\boldsymbol{\phi}_e \cdot \mathbf{x})^2 = \sum_{e \notin \text{both blocks}} w_e^2, \tag{6.1}$$

2026 JINST 21 P11017

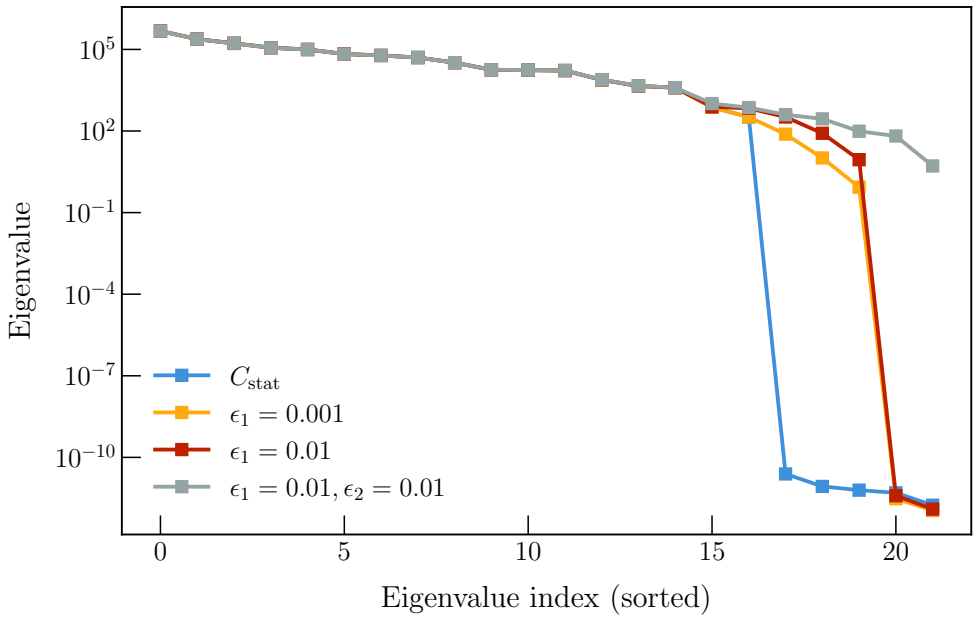


Figure 8. Eigenvalue spectra of the statistical covariance matrix C_{stat} for the toy analysis, sorted in descending order, as events are removed at random from one or two blocks, breaking the exact event-sharing structure. Fractions ϵ_1 and ϵ_2 of events are dropped from the p_μ and $\cos\theta_\mu$ marginal blocks, respectively.

which vanishes only when every event is shared. When a fraction ϵ of events is unshared, the variance is of order ϵN. In this case, the corresponding eigenvalue is lifted from zero into the body of the spectrum by the statistical model itself, before any systematic uncertainty is added. This is a qualitatively different situation from the inflation of section 4.3. There the variance along some directions was exactly zero in C_{stat} and lifted only by the systematic model to a value unrelated to any statistical power, so dividing a residual by it produced an unphysical contribution. A direction carrying a genuine variance ϵN from real event mismatch is not pathological, as the test weights it by its true variance and remains calibrated. The risk of imperfect sharing is therefore not renewed inflation but over-projection, i.e. discarding a direction that now carries a small but real amount of information.

The risk is confined to the a priori structural prescription, and it can be checked for. Three objects encode the null space under exact sharing, namely the combination matrices Ω_{struct} and Ω and the eigenspectrum of the statistical covariance matrix, and they respond differently when sharing is broken. The structural matrix Ω_{struct} depends only on the binning geometry and is blind to event occupancy entirely. When there is a genuine difference in phase-space acceptance between blocks, this can produce bin combinations absent from Ω_{struct}, breaking the hierarchy of eq. (2.9) so that $\text{rank}(\Omega)$ exceeds $\text{rank}(\Omega_{\text{struct}})$; the effect is visible in Ω but not in Ω_{struct}. The empirical combination matrix Ω depends on which bin combinations are occupied, but not on how many events carry each, so it is blind to a violation caused only by missing events in certain blocks. In that case we can have $\text{rank}(\Omega) = \text{rank}(\Omega_{\text{struct}})$ even though the affected direction has acquired genuine variance. Only the spectrum of C_{stat} reflects that variance, lifting the corresponding eigenvalue out of the null space.

These two failure modes are illustrated with the toy model of section 4. The first is demonstrated by extending the p_μ acceptance of the two-dimensional block beyond that of the p_μ marginal, so that events with $p_\mu > 1.2\,\text{GeV}/c$ populate the former but not the latter. Of the $\mathcal{O}(5 \times 10^5)$ generated

2026 JINST 21 P11017

events only seven fall in this region, yet their presence is enough to raise $\mathrm{rank}(\Omega)$ from 17 to 18, above $\mathrm{rank}(\Omega_{\mathrm{struct}}) = 17$, and to make the affected block total fall short of the others by exactly those seven events. The rank signature is thus sensitive to even a single forbidden bin combination.

The second is demonstrated by removing a fraction of events from one block at random, so that the shared-event structure is broken without altering the binning. Figure 8 shows the eigenvalue spectrum of C_{stat} as this fraction is increased independently for blocks 1 (p_μ marginal) and 2 ($\cos\theta_\mu$ marginal). The formerly null eigenvalues, exactly zero under perfect sharing, are lifted into the bulk of the spectrum, filling the spectral gap that otherwise identifies the null space cleanly. The pattern set is unchanged, so $\mathrm{rank}(\Omega) = \mathrm{rank}(\Omega_{\mathrm{struct}})$ throughout; only C_{stat} reflects the violation.

These complementary sensitivities provide a practical set of consistency checks rather than a reason to abandon the geometric construction. Before relying on the structural projection, one should verify that C_{stat} has the expected number of near-zero eigenvalues, separated from the bulk of the spectrum by a clean gap, that any difference between $\mathrm{rank}(\Omega)$ and $\mathrm{rank}(\Omega_{\mathrm{struct}})$ is understood, and that the measured block totals are consistent within their correlated statistical uncertainties. The spectral-gap check assumes C_{stat} is built from the sum of squared event weights following eq. (2.2), which assigns exactly zero variance to the constrained directions. Estimators that quantify the statistical uncertainty differently may populate these directions with small nonzero values and need not show a clean gap even under exact sharing; there, the estimator-independent comparison between $\mathrm{rank}(\Omega)$ and $\mathrm{rank}(\Omega_{\mathrm{struct}})$ is the reliable diagnostic. A deficit, $\mathrm{rank}(\Omega) < \mathrm{rank}(\Omega_{\mathrm{struct}})$, is expected when allowed bin combinations are simply unpopulated; this is the kinematic null case seen in section 5, handled correctly by the empirical Ω. An excess, $\mathrm{rank}(\Omega) > \mathrm{rank}(\Omega_{\mathrm{struct}})$, instead signals that events occupy structurally forbidden combinations, the hallmark of broken sharing. When the diagnostics are consistent with exact sharing, the a priori count of section 2 applies as derived. When they are not, the structural null space over-counts the constrained directions, and projecting onto it would discard information; this signals that the binning or selection definitions should be reconciled across blocks, or that a treatment retaining the affected directions, such as marginalizing over nuisance parameters describing the differences (section 3.3), is more appropriate.

7 Summary and conclusions

Multi-distribution cross-section measurements that share a common event sample contain exact linear constraints among their bin counts. When the statistical covariance accounts for the inter-distribution correlations, these constraints appear as a rank deficiency, and a global χ^2 that ignores them is ill-defined. The null directions either prevent inversion or, once lifted by systematic uncertainties, produce unphysically inflated values. We have presented the range-projected χ^2, a goodness-of-fit statistic that restricts the test to the subspace carrying independent statistical information, yielding a χ^2 with $N_{\mathrm{bins}} - N_{\mathrm{null}}$ degrees of freedom.

This work makes two key contributions. First, we have shown that this rank deficiency is a predictable consequence of the shared-event structure rather than an incidental numerical degeneracy. The constrained directions are exactly those along which the data carry no statistical power, so removing them discards no physical information. Second, we have provided a prescription to compute the constrained directions a priori from the binning geometry, decomposing N_{null} into a structural contribution fixed by the bin edges and a kinematic contribution set by the phase-space occupancy. Because the rank is determined from the measurement geometry rather than from the systematics-

2026 JINST 21 P11017

contaminated covariance spectrum, the method requires no eigenvalue threshold on the full covariance, where no clean separation between null and physical directions needs to exist.

The method was validated in two settings. In an analytical toy model the projected statistic follows the expected $\chi^2(17)$ distribution, while the unprojected statistic is inflated to a mean of 30.6 and over-rejects the correct model in roughly one third of experiments; we traced the inflation to a noncentrality on the directions lifted out of the null space by systematic uncertainties. In a simulated neutrino-argon measurement with unfolding and multi-source systematic uncertainties, the projected statistic gives 15.7/16 for the correct model and 154.4/16 for an incorrect one, recovering both calibration and discriminating power, whereas the unprojected statistic rejects both models. The two studies share five null directions of differing composition, structural in the toy and a mixture of structural and kinematic in the simulated measurement. We note that the projected statistic is mildly over-dispersed relative to the ideal χ^2, with a false rejection rate of 5.7% at the nominal 5% level in the toy. This is a small residual effect that does not affect the conclusions.

The method assumes that the event-sharing structure is exact. For this reason we recommend that analyzers stick to exact event-sharing structures when possible. Departures from this assumption are detectable through the spectral gap of the statistical covariance, the comparison of the empirical and structural ranks, and the consistency of the measured block totals, and we have given a corresponding set of diagnostics and remedies. Where the full forward model and nuisance parameterization are available, a binned likelihood fit handles the constraints automatically and is more robust to imperfect sharing; the range-projected χ^2 instead requires only the reported cross-block covariance, the binning, and the detector response, making it suited to the goodness-of-fit testing, model tuning, and generator comparison that reuse published measurements.

As multi-distribution measurements with shared events become more common, a correctly calibrated global goodness-of-fit test becomes a routine requirement for their interpretation. The algorithm for constructing the structural combination matrix is given in appendix A, and a reference implementation can be found in ref. [34].

Acknowledgments

We thank Steve Pate-Morales, Andy Furmanski, Ben Bogart, and Nitish Nayak for valuable discussions that contributed to the development of this work, which originated from studies conducted within the MicroBooNE collaboration. P. Green is supported by a United Kingdom Research and Innovation (UKRI) Future Leaders Fellowship, MR/V022407/1. This manuscript has been authored by Fermi Forward Discovery Group, LLC under Contract No. 89243024CSC000002 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

A Structural combination matrix algorithm

We describe the algorithm implemented to construct the structural combination matrix Ω_{struct} from the bin edge specifications of the multi-block measurement. Each row of Ω_{struct} is a unique indicator vector encoding one structurally allowed combination of bins, one per block, that a single event can simultaneously populate. The rank of this matrix determines the number of null directions, as described in section 2.3.

2026 JINST 21 P11017

The construction proceeds in three stages. First, the kinematic variables used across all blocks are classified as shared (appearing in two or more blocks) or private (appearing in exactly one block). This distinction is key to the efficiency of the algorithm. Shared variables create correlations between blocks, while private variables contribute bins that are independent of all other blocks.

Second, a common refinement is computed for each shared variable by taking the union of all bin edges from every block that uses it. The algorithm then enumerates the cells of this shared-variable refinement. For each cell, the midpoint in each shared variable is used to identify which bin (or set of bins) in each block is compatible with that cell. When the binning of a given block is fully determined by the shared variables (e.g., a 1D block whose variable is shared, or a 2D block whose slice and bin variables are both shared) the cell maps to a single bin. When a block uses a private variable, the shared cell constrains only part of the binning, and the block contributes multiple candidate bins corresponding to the different private-variable choices. A 1D block whose variable is private contributes all of its bins; a 2D block with a shared slice variable and a private bin variable contributes all bins within the determined slice.

Third, the candidate bins from each block are combined. In the simplest case, the full Cartesian product of all candidate bins is expanded, and each element becomes a row of Ω_{struct}. However, when many blocks have private variables, this product can become very large. In that case, a representative subset of indicator vectors is used instead. A base combination is formed by selecting the first candidate bin from each block, and then the remaining candidates from each block are substituted one at a time while the others are held at their base values. This produces at most $1 + \sum_b (n_b - 1)$ rows per shared cell, where n_b is the number of candidate bins for block b. The representative set spans the same column space as the full Cartesian product, and therefore yields the same $\text{rank}(\Omega_{\text{struct}})$ and null-space structure.

To understand why, note that any indicator vector from the full product can be written as the sum of indicator vectors from the representative set. Denoting the base combination as $\mathbf{e}_0$ and the single-block variations as $\mathbf{e}_j^{(k)}$ (block k switched to its j-th candidate), any full-product row that selects candidate j_k in each block satisfies

$$\mathbf{e}_{j_1,\ldots,j_B} = \mathbf{e}_0 + \sum_k (\mathbf{e}_{j_k}^{(k)} - \mathbf{e}_0), \tag{A.1}$$

since the blocks occupy disjoint columns. The full product therefore lies in the span of the representative rows, and the rank is preserved.

When the measurement involves mutually exclusive event selections, blocks belonging to different selections share no events and are processed independently. The algorithm is applied separately to each selection group, and the resulting matrices are vertically concatenated.

The cost of the algorithm is dominated by the enumeration of shared-variable cells, which scales as $\prod_{v \in \text{shared}} M_v$ where M_v is the number of refined bins for shared variable v. Since only shared variables enter this product, measurements with many blocks but few shared variables, such as a set of 1D distributions in different kinematic variables, are handled efficiently regardless of the total number of bins.

Code Availability Statement. This article has associated code in a code repository. Code available at https://github.com/fmartinezlopez/GlobalChi2.

2026 JINST 21 P11017

References

[1] MicroBooNE collaboration, *First Simultaneous Measurement of Differential Muon-Neutrino Charged-Current Cross Sections on Argon for Final States with and Without Protons Using MicroBooNE Data*, *Phys. Rev. Lett.* **133** (2024) 041801 [arXiv:2402.19281].

[2] MicroBooNE collaboration, *Inclusive cross section measurements in final states with and without protons for charged-current ν_μ-Ar scattering in MicroBooNE*, *Phys. Rev. D* **110** (2024) 013006 [arXiv:2402.19216].

[3] MicroBooNE collaboration, *First Double-Differential Measurement of Kinematic Imbalance in Neutrino Interactions with the MicroBooNE Detector*, *Phys. Rev. Lett.* **131** (2023) 101802 [arXiv:2301.03706].

[4] MicroBooNE collaboration, *Multidifferential cross section measurements of ν_μ-argon quasielasticlike reactions with the MicroBooNE detector*, *Phys. Rev. D* **108** (2023) 053002 [arXiv:2301.03700].

[5] T2K collaboration, *Characterization of nuclear effects in muon-neutrino scattering on hydrocarbon with a measurement of final-state kinematics and correlations in charged-current pionless interactions at T2K*, *Phys. Rev. D* **98** (2018) 032003 [arXiv:1802.05078].

[6] T2K collaboration, *First T2K measurement of transverse kinematic imbalance in the muon-neutrino charged-current single-π^+ production channel containing at least one proton*, *Phys. Rev. D* **103** (2021) 112009 [arXiv:2102.03346].

[7] NOvA collaboration, *Measurement of the double-differential muon-neutrino charged-current inclusive cross section in the NOvA near detector*, *Phys. Rev. D* **107** (2023) 052011 [arXiv:2109.12220].

[8] NOvA collaboration, *Measurement of $d^2\sigma/d|\vec{q}|dE_{avail}$ in charged current ν_μ-nucleus interactions at $\langle E_\nu \rangle = 1.86$ GeV using the NOvA Near Detector*, *Phys. Rev. D* **111** (2025) 052009 [arXiv:2410.05526].

[9] NOvA collaboration, *Measurement of the double-differential cross section of muon-neutrino charged-current interactions with low hadronic energy in the NOvA Near Detector*, arXiv:2410.10222.

[10] T2K collaboration, *Simultaneous measurement of the muon neutrino charged-current cross section on oxygen and carbon without pions in the final state at T2K*, *Phys. Rev. D* **101** (2020) 112004 [arXiv:2004.05434].

[11] MINERvA collaboration, *Simultaneous Measurement of ν_μ Quasielasticlike Cross Sections on CH, C, H_2O, Fe, and Pb as a Function of Muon Kinematics at MINERvA*, *Phys. Rev. Lett.* **130** (2023) 161801 [arXiv:2301.02272].

[12] MINERvA collaboration, *Simultaneous Measurement of Muon Neutrino ν_μ Charged-Current Single π^+ Production in CH, C, H_2O, Fe, and Pb Targets in MINERvA*, *Phys. Rev. Lett.* **131** (2023) 011801 [arXiv:2209.07852].

[13] T2K collaboration, *First combined measurement of the muon neutrino and antineutrino charged-current cross section without pions in the final state at T2K*, *Phys. Rev. D* **101** (2020) 112001 [arXiv:2002.09323].

[14] S. Gardiner, *Mathematical methods for neutrino cross-section extraction*, arXiv:2401.04065.

[15] G. D'Agostini, *A Multidimensional unfolding method based on Bayes' theorem*, *Nucl. Instrum. Meth. A* **362** (1995) 487.

[16] W. Tang et al., *Data Unfolding with Wiener-SVD Method*, 2017 *JINST* **12** P10002 [arXiv:1705.03568].

[17] C. Wilkinson et al., *Testing charged current quasi-elastic and multinucleon interaction models in the NEUT neutrino interaction generator with published datasets from the MiniBooNE and MINERvA experiments*, *Phys. Rev. D* **93** (2016) 072010 [arXiv:1601.05592].

2026 JINST 21 P11017

[18] MINERvA collaboration, *Tuning the GENIE Pion Production Model with MINERvA Data*, *Phys. Rev. D* **100** (2019) 072005 [arXiv:1903.01558].

[19] GENIE collaboration, *Neutrino-nucleus CC0π cross-section tuning in GENIE v3*, *Phys. Rev. D* **106** (2022) 112001 [arXiv:2206.11050].

[20] C. Rao and S. Mitra, *Generalized Inverse of Matrices and Its Applications*, Probability and Statistics Series, Wiley (1971).

[21] MicroBooNE collaboration, *Measurement of single- and double-differential cross sections for mesonless charged-current muon neutrino interactions on argon with final-state protons using the MicroBooNE detector*, *Phys. Rev. D* **112** (2025) 112004 [arXiv:2403.19574].

[22] MicroBooNE collaboration, *First Double-Differential Cross Section Measurement of Neutral-Current π^0 Production in Neutrino-Argon Scattering in the MicroBooNE Detector*, *Phys. Rev. Lett.* **134** (2025) 161802 [arXiv:2404.10948].

[23] NOvA collaboration, *Adjusting neutrino interaction models and evaluating uncertainties using NOvA near detector data*, *Eur. Phys. J. C* **80** (2020) 1119 [arXiv:2006.08727].

[24] MicroBooNE collaboration, *New CC0π GENIE model tune for MicroBooNE*, *Phys. Rev. D* **105** (2022) 072001 [arXiv:2110.14028].

[25] R. Penrose, *A generalized inverse for matrices*, *Proc. Cambridge Phil. Soc.* **51** (1955) 406.

[26] T2K collaboration, *First measurement of muon neutrino charged-current interactions on hydrocarbon without pions in the final state using multiple detectors with correlated energy spectra at T2K*, *Phys. Rev. D* **108** (2023) 112009 [arXiv:2303.14228].

[27] P. Stowell et al., *NUISANCE: a neutrino cross-section generator tuning and comparison framework*, 2017 *JINST* **12** P01016 [arXiv:1612.07393].

[28] C. Andreopoulos et al., *The GENIE Neutrino Monte Carlo Generator*, *Nucl. Instrum. Meth. A* **614** (2010) 87 [arXiv:0905.2517].

[29] T. Golan, J.T. Sobczyk and J. Zmuda, *NuWro: the Wroclaw Monte Carlo Generator of Neutrino Interactions*, *Nucl. Phys. B Proc. Suppl.* **229–232** (2012) 499.

[30] A.S. Meyer, M. Betancourt, R. Gran and R.J. Hill, *Deuterium target data for precision neutrino-nucleus cross sections*, *Phys. Rev. D* **93** (2016) 113015 [arXiv:1603.03048].

[31] J. Nieves, I. Ruiz Simo and M.J. Vicente Vacas, *Inclusive Charged-Current Neutrino-Nucleus Reactions*, *Phys. Rev. C* **83** (2011) 045501 [arXiv:1102.2777].

[32] D. Rein and L.M. Sehgal, *Neutrino Excitation of Baryon Resonances and Single Pion Production*, *Annals Phys.* **133** (1981) 79.

[33] S. Dytman et al., *Comparison of validation methods of simulations for final state interactions in hadron production experiments*, *Phys. Rev. D* **104** (2021) 053006 [arXiv:2103.07535].

[34] Implementation of the structural combination matrix construction, https://github.com/fmartinezlopez/GlobalChi2.

2026 JINST 21 P11017