

Fermilab

Efficient Monte Carlo Event Generation for Neutrino-Nucleus
Exclusive Cross Sections

FERMILAB-PUB-25-0107-V

arXiv:2502.14452

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.

Efficient Monte Carlo Event Generation for Neutrino-Nucleus Exclusive Cross Sections

Mathias El Baz^{1,*}, Federico Sánchez¹, Natalie Jachowicz², Kajetan

Niewczas², Ashish Kumar Jha², and Alexis Nikolakopoulos³

¹*Département de Physique Nucléaire et Corpusculaire, Université de Genève, Geneva, Switzerland*

²*Ghent University, Ghent, Belgium and*

³*Fermi National Accelerator Laboratory, Batavia, Illinois, USA*

(Dated: February 21, 2025)

Modern neutrino-nucleus cross section predictions need to incorporate sophisticated nuclear models to achieve greater predictive precision. However, the computational complexity of these advanced models often limits their practicality for experimental analyses. To address this challenge, we introduce a new Monte Carlo method utilizing Normalizing Flows to generate surrogate cross sections that closely approximate those of the original model while significantly reducing computational overhead. As a case study, we built a Monte Carlo event generator for the neutrino-nucleus cross section model developed by the Ghent group. This model employs a Hartree-Fock procedure to establish a quantum mechanical framework in which both the bound and scattering nucleon states are solutions to the mean-field nuclear potential. The surrogate cross sections generated by our method demonstrate excellent accuracy with a relative effective sample size of more than 98.4%, providing a computationally efficient alternative to traditional Monte Carlo sampling methods for differential cross sections.

I. INTRODUCTION

The study of neutrino interactions with nuclei is critical for reducing systematic uncertainties in modern neutrino oscillation experiments. Among the most important processes is Charged Current Quasi-Elastic (CCQE) scattering, where a neutrino interacts with one nucleon inside a nucleus, emitting a charged lepton while ejecting one nucleon. This interaction forms the backbone of event reconstruction in neutrino oscillation experiments like T2K (Tokai to Kamioka experiment [1]) and the new Hyper-Kamiokande water Cherenkov detector [2]. However, the theoretical modeling of CCQE interactions is complex due to the need to account for the nuclear environment in which the nucleons are bound originally [3].

Moreover, these experiments often use neutrino beams that span a broad energy spectrum—from a few hundred MeV to several GeV—making Monte Carlo (MC) event generators indispensable. MC generators, such as NEUT [4], GENIE [5] and NuWro [6], efficiently sample different interaction processes over many possible final states, enabling exhaustive modeling across this wide energy range. Such modeling is essential for extracting oscillation parameters and minimizing systematic uncertainties in neutrino physics.

A variety of interaction models have been implemented in MC generators to describe neutrino-nucleus interactions. However, most of these models are limited to providing only inclusive cross sections and therefore cannot reliably predict final-state hadron kinematics. As a consequence, the latter are often generated in an ad hoc and unrealistic manner [7, 8].

The only approaches currently used in neutrino generators that systematically provide final-state nucleon kinematics are based on the factorization obtained in the plane-wave impulse approximation (PWIA). In this picture, one assumes no final-state interactions, so the process can be factorized into (i) the primary neutrino-nucleon interaction and (ii) a nuclear model describing the bound nucleons. However, these approaches lack the effects of the nuclear medium on the outgoing nucleon.

In the simplest instance of PWIA, the nucleus is modeled as a Relativistic Fermi Gas (RFG) [9], in which nucleons are treated as fermions uniformly occupying momentum states up to the Fermi momentum in a constant potential. This approach has been remarkably successful in describing general properties of intermediate-energy processes, especially in inclusive electron scattering. Yet, it neglects important nuclear features such as the shell structure and nucleon-nucleon correlations.

A more refined description of the nuclear medium still within the PWIA is provided by the Spectral Function (SF) model [10], which yields a probability distribution $S(p, E)$ for finding a nucleon with momentum p and removal energy E based on $(e, e'p)$ data and theoretical inputs. Although the SF model captures the spread in binding energies and some nucleon-nucleus correlations, its essential reliance on the PWIA means that correlations essential to accurately predict the final-state hadron kinematics cannot be fully taken into account.

While these approximations were adequate when the focus was largely on lepton kinematics, the increased precision of upcoming accelerator-based long-baseline neutrino oscillation experiments will enable much more detailed measurements of outgoing hadron kinematics. To achieve the few-percent level of uncertainty on neutrino interactions required by these next-generation neutrino oscillation studies, cross section models must go beyond

* mathias.elbaz@unige.ch

PWIA to include final-state interactions.

To obtain exclusive differential cross sections, a theoretical framework that accurately describes the hadron system is required. One such approach is based on the distorted-wave impulse approximation (DWIA) in the mean field approximation [11, 12]. Mean field based models describe nucleons within an averaged nuclear potential generated by all other nucleons. This naturally leads to a shell structure, where each nucleon in the initial state is described by a bound wave function. Meanwhile, the distorted-wave treatment goes beyond the plane-wave description by also considering the final-state interactions of the outgoing nucleon: instead of emerging as a free particle, the outgoing nucleon propagates through the same mean field potential and undergoes distortions due to its interactions with the residual nucleus.

Despite their advantages, mean-field-based cross section calculations are computationally expensive compared to the factorized PWIA. The slow evaluation speed of these cross sections translates into slow sampling rates when traditional methods like accept-reject algorithms are used. To mitigate this, some current solutions [13–16] involve precomputing the hadron tensor components into heavy multidimensional grids, known as hadron tables [7, 17, 18]. This method allows to compute the cross section much faster, and make an accept-reject sampling possible. However, this approach has significant drawbacks: it lacks flexibility, requiring recalculation for any change in theoretical parameters, and seem infeasible for multi-nucleon knockout exclusive cross sections due to the exponential scaling of table size with dimensionality.

Recent advancements in Artificial Intelligence, particularly in Deep Learning, offer promising solutions to these challenges. In this work, we propose an alternative sampling method based on Normalizing Flows (NF) [19]. Normalizing Flows transform a simple base distribution, such as a Gaussian, into a complex target distribution through a series of learnable diffeomorphisms. This enables efficient sampling and evaluation of the target distribution—in this case, the neutrino-nucleus cross section. In this work, we demonstrate the strong potential of NF to build a Monte Carlo event generator for neutrino-induced cross sections. For our proof-of-principle we use the Hartree-Fock mean-field exclusive $1p1h$ (1-particle-1-hole) cross section formalism developed by the Ghent Group [12, 20, 21]. These cross sections exhibit all the features needed to demonstrate that our approach can pave the way for faster and more flexible event generation in neutrino physics.

The structure of this article is as follows: we first review the theoretical framework behind CCQE interactions within the mean-field approach adopted by the Ghent Group. Following this, we introduce the Normalizing Flows technique and outline how it is integrated into the Monte Carlo event generation process. Finally, we present results demonstrating the efficiency and accuracy of the generator.

II. NEUTRINO-NUCLEUS EXCLUSIVE CCQE CROSS SECTION IN THE MEAN FIELD FRAMEWORK

The aim of this study is to demonstrate the ability to efficiently and accurately generate events for single nucleon knockout distributed according to the Hartree-Fock calculations developed by the Ghent group. While this model can predict both single and multiple nucleon knockout, providing a Monte Carlo event generator for all processes in their full complexity lies beyond the scope of this work. Future steps for handling the other processes are discussed in Section V. Here, we focus on $1p1h$ interactions due to a single-nucleon operator. The goal of this section is to introduce the relevant kinematic structure of the cross section.

The process is the following :

$$\nu_\mu + \frac{A}{Z}X \longrightarrow \mu^- + p + \left(\frac{A-1}{Z-1}X\right)^*, \quad (1)$$

where:

- ν_μ is the incoming muon neutrino,
- $\frac{A}{Z}X$ represents the target nucleus with atomic number Z and mass number A ,
- μ^- is the outgoing charged lepton,
- p is the ejected proton,
- $\left(\frac{A-1}{Z-1}X\right)^*$ is the residual nucleus in an excited state due to the interaction.

The kinematics are displayed in Figure 1, here the target nucleus is at rest, and the coordinate system is chosen so that the momentum transfer \vec{q} defines the z -axis.




FIG. 1: $1p1h$ kinematics in the laboratory frame with the z axis along the momentum transfer

The interaction between the neutrino and a nucleon within the nucleus by boson exchange is described using the contraction of the lepton and hadron tensors, denoted as $L_{\mu\nu}$ and $W^{\mu\nu}$, respectively :

$$\frac{d^6\sigma(E_\nu)}{d\omega d\Omega_\mu dE_N d\Omega_N} \propto L_{\mu\nu}W^{\mu\nu}.$$

where $\omega = E_\nu - E_\mu$ is the energy transfer, E_N the outgoing nucleon energy, and Ω the solid angles for the muon

and proton. E_N is the energy of the outgoing nucleon. The lepton tensor is independent of the nuclear model, and is calculated in electroweak theory. As the cross section is not differential in neutrino energy, event generators typically treat E_ν as an external input, modeling the conditional cross section for a fixed neutrino energy.

Clearly, the cross section is invariant with respect to rotations around the z -axis. This allows samples to be uniformly rotated by a ϕ angle in the range $[0, 2\pi]$ without altering the physical results. Consequently, the variables that need to be sampled at fixed E_ν are $(\omega, \theta_\mu, E_N, \Omega_N)$.

We make an additional simplification to reduce the number of independent variables which is specific to the mean-field model. The hadron tensor $W^{\mu\nu}$ is a bilinear product of the hadron currents with an energy conserving delta function

$$W^{\mu\nu} = \sum_{i,f} \langle i | \mathcal{J}^{\mu\dagger} | f \rangle \langle f | \mathcal{J}^\nu | i \rangle \delta(E_N + E_{A-1}^* - M_A - \omega), \quad (2)$$

where $|i\rangle$ and $|f\rangle$ are the initial and final hadron states. The summation is understood to run over all relevant final states and average over the initial states. \mathcal{J}^ν is the nuclear current operator.

The nuclear states $|i\rangle$ and $|f\rangle$ are obtained as Slater determinants, composed of nucleon solutions to the Schrödinger equation with the mean-field potential $V_{\text{MF}}^{(\alpha)}$:

$$[\hat{T}^{(\alpha)}(r) + V_{\text{MF}}^{(\alpha)}(r)] \phi_\alpha(r) = \epsilon_\alpha \phi_\alpha(r) \quad (3)$$

with the kinetic energy operator being :

$$\hat{T}^{(\alpha)}(r) = -\frac{\hbar^2}{2} \left[\frac{1}{r^2} \frac{d}{dr} \left(\frac{r^2}{m^*(r)} \frac{d}{dr} \right) - \frac{\ell_\alpha(\ell_\alpha + 1)}{m^*(r)r^2} \right]$$

with α denoting the set of quantum numbers identifying the single-nucleon state and $m^*(r)$ the effective mass (since the Hartree-Fock Hamiltonian is non-local) [22]. The mean-field potential is derived via a Hartree-Fock procedure [23], starting from an effective Skyrme nucleon-nucleon potential [24] and extending it to a nucleon-nucleus potential.

The final state $|f\rangle$ is generally expressed as a sum of orthogonal continuum solutions to the same equation with the correct asymptotic behavior. The complexity of nuclear modeling is embedded in the summation over all possible final states, which constitutes the primary computational bottleneck due to the necessity of computing and summing over a large number of distorted wave solutions of Equation 3.

The ground state $|i\rangle$ is described as a Slater determinant of the A lowest energy single-particle states of the mean field potential. The states have quantum numbers $(\epsilon, n, l, j, m_j, i_z)$, i.e. respectively energy, principal quantum number, orbital and total angular momentum, projection of total angular momentum and isospin. Due to spherical symmetry, the states within a shell characterized by quantum numbers $\alpha = (\epsilon_\alpha, n_\alpha, l_\alpha, j_\alpha, i_z)$

are energy-degenerate. Since we are considering nucleon knockout to the continuum, the residual system can then only be left in states with excitation energy given by the single-particle energies of the shells $\epsilon_\alpha = M_N + E_{A-1}^* - M_A$. Using this to integrate the delta function, and summing over m_j the cross section decomposes into a sum of partial cross sections,

$$\frac{d^4\sigma(E_\nu)}{d\omega d\theta_\mu d\Omega_N} = \sum_\alpha (2j_\alpha + 1) \frac{d^4\sigma_\alpha(E_\nu)}{d\omega d\theta_\mu d\Omega_N}, \quad (4)$$

where the sum runs over the shells with occupation $(2j_\alpha + 1)$. Each partial cross section produces a different spectrum of final-state nucleon energies determined by α .

We thus generate samples of the four independent continuous variables $(\theta_\mu, \omega, \theta_N, \phi_N)$ for each shell separately, with E_ν as an input.

Among these variables, only the energy transfer ω has a more complex marginal distribution, particularly at low values where nuclear effects play a significant role. In contrast, the remaining variables are angles which are characterized by fixed support and exhibit relatively smooth distributions.

We lastly note here, that the dependence of the cross section on the azimuthal angle of the hadron plane, as shown in Fig. 1, can always be decomposed into 5 simple functions [25, 26]. For nucleon knockout the dependence on the nucleon's azimuthal angle can be written as

$$\begin{aligned} \frac{d^4\sigma_\alpha(E_\nu)}{d\omega d\theta_\mu d\Omega_N} &= \frac{d^3\sigma_\alpha(E_\nu)}{d\omega d\theta_\mu d\cos\theta_N} \\ &+ B \cos\phi_N + C \cos 2\phi_N \\ &+ D \sin\phi_N + E \sin 2\phi_N, \end{aligned} \quad (5)$$

in terms of the angle integrated cross section and the functions B, C, D, E that only depend on $E_\nu, \omega, \theta_\mu, \theta_N$ and the shell. This decomposition can be exploited to sample the ϕ_N dependence [27], and hence reduce the number of non-trivial independent variables. For this work however, we sample the full kinematics directly from the partial cross sections of Eq. (4).

III. EVENT GENERATOR USING NORMALIZING FLOWS

A. Definition of Normalizing Flows

Normalizing Flows (NF) are a class of machine learning models used to transform simple probability distributions into complex ones through a sequence of invertible and differentiable transformations of the probability space. They are particularly useful for efficiently sampling from high-dimensional and non-Gaussian distributions, making them highly suited for generating events from complex interaction models, such as neutrino-nucleus interactions.

The core idea behind NF is to start with a simple base distribution, typically a multidimensional Gaussian or uniform distribution, and apply a series of transformations that map this distribution into the more complex target distribution of interest. Mathematically, a typical NF transformation consists of a sequence of K diffeomorphisms of the probability space, each parameterized by a neural network, that progressively transform the base distribution $p_Z(z)$ into the target distribution $p_X(x)$. If $z \sim p_Z(z)$ is a sample from the base distribution, the final output $x = f_K \circ f_{K-1} \circ \dots \circ f_1(z)$ represents a sample from the target distribution $p_X(x)$.

To compute the likelihood $p_X(x)$, we use the change of variables formula:

$$p_X(x) = p_Z(z) \left| \det \frac{\partial f^{-1}}{\partial x} \right|,$$

where $z = f^{-1}(x)$ is the pre-image of x under the inverse transformation, and $\det \frac{\partial f^{-1}}{\partial x}$ is the Jacobian determinant of the transformation. One can see the main advantage of NF for event generation: they allow to both sample and evaluate the probability of an event at the same time in an efficient and accurate way.

B. Autoregressive flows for event generation

NF have been successfully applied to event generation in high-energy physics. Gao et al. [28] demonstrated that NF can enhance the efficiency of Monte Carlo event generation, particularly in complex scenarios such as Drell-Yan processes at the Large Hadron Collider. Stienen and Verheyen [29] also explored the application of NF to top quark pair production. These works have shown that NF significantly improve unweighting efficiency compared to traditional methods. Building on these results, we adapt and extend the use of NF to neutrino-nucleus interactions, aiming to provide an efficient event generator for complex neutrino-nucleus cross sections.

Normalizing flows rely on invertible transformations that allow for efficient application of the change-of-variable formula. Two types of architectures show the best performances: autoregressive flows and coupling layers. Both approaches exploit triangular (or block-triangular) Jacobians to ensure computational efficiency, with complexity scaling linearly with the dimensionality of the probability space. However, the choice between autoregressive flows and coupling layer-based flows ultimately depends on the application.

Autoregressive flows such as inverse autoregressive flow (IAF) [30] or masked autoregressive flow (MAF) [31] are D times slower to invert than to evaluate, where D is the dimension of the probability space. On the other hand, flows based on coupling layers, such as NICE [32] or RealNVP [33], have an analytic one-pass inverse. Sampling and evaluating the density at the same time requires to evaluate both the forward and inverse transformations.

Therefore, coupling layers allow to both evaluate and sample a density fast. However, coupling layer-based flows are generally less expressive than autoregressive flows. Therefore, we chose to use autoregressive flows due to their higher expressiveness. A more recent comparative study between coupling and autoregressive flows by Cocco et al. [34] indicates that autoregressive flows stand out both in terms of accuracy and training speed. We will later also demonstrate that autoregressive flows offer sufficient evaluation and sampling speeds for the dimensionality of typical CCQE events.

In this study, we employ MAF, which use masked feed-forward networks to parameterize the NF transformations while enforcing their autoregressive property. The masking mechanism involves applying a binary mask to the weight matrices of the network, effectively setting specific connections to zero. This ensures that the output for a given dimension depends only on its predecessors in the ordering, preserving the conditional dependency property required for autoregressive modeling. A permutation of the dimensions after each autoregressive transformation ensures that all dimensions are treated equally.

C. Circular Rational-Quadratic Neural Spline Flows

In addition to using autoregressive flows to implement our NF, we need to define the parametrization of the flow transformation. In this study, we use Rational Quadratic Neural Spline Flows (RQ-NSF) introduced by Durkan et al. [35], which effectively models complex, non-Gaussian target distributions by utilizing piecewise rational quadratic transformations, as parametrized by Gregory and Delbourgo [36]. RQ-NSF are very expressive due to their infinite Taylor-series expansion while being defined by a small number of parameters. RQ-NSF are widely used in recent applications of NF for their state-of-the-art expressiveness. They have already been applied in the T2K collaboration for tasks such as modeling simple cross sections [37] and posterior density estimation for near-detector fits [38].

However, standard RQ-NSF, defined on Euclidean probability spaces, are suboptimal for modeling periodic variables, such as angular distributions. These distributions often exhibit sharp cut at boundaries when projected onto flat spaces, despite maintaining periodic boundary conditions. To address this, it is essential to model angles within their natural phase spaces—circles for angles and spheres for solid angles. In this work, we adopt the extension of RQ-NSF to non-Euclidean probability spaces, such as tori and spheres, as developed by Rezende et al. [39]. By modifying Rational Quadratic Spline Flows to ensure periodicity in certain dimensions, we can appropriately model the periodic distributions through differentiable transformations of the probability space. In our specific $1p1h$ cross section model, the

first two kinematics, θ_μ and ω , can be represented on a cylindrical manifold due to the periodicity of θ_μ . The remaining two variables, θ_N and ϕ_N , are naturally represented on a sphere. Therefore, the natural phase space \mathcal{M} where the cross section is defined is a cylinder time a sphere.

D. Energy-dependent flows

In neutrino–nucleus interactions, the differential cross section depends on both the kinematic variables and the incoming neutrino energy E_ν . Neutrino event generators, such as NEUT [4], typically take E_ν as an input parameter and generate samples weighted by the conditional cross section given E_ν , rather than modeling E_ν as an extra dimension. Consequently, we condition the NF directly on E_ν to mimic the functioning.

We model a family of probability distributions continuously parameterized by E_ν . This approach ensures that the NF captures the underlying dependence of the cross section on energy without explicitly including E_ν as part of the flow’s dimensional space. Additionally, this eliminates the need to retrain the NF for different values of E_ν .

For the specific case of $1p1h$ interactions, the goal is to model the conditional probability

$$p(\theta_\mu, \omega, \theta_N, \phi_N | E_\nu, \alpha),$$

over a wide range of E_ν . During the transformation, the flow modifies the kinematic variables ($\theta_\mu, \omega, \theta_N, \phi_N$) in a way that captures the energy dependence implicitly. Here, we model each initial shell α separately.

The conditional probability is computed by normalizing the full 4D differential cross section for a given shell α by the total cross section :

$$p(\theta_\mu, \omega, \theta_N, \phi_N | E_\nu, \alpha) = \frac{d^4 \sigma_\alpha(E_\nu)}{d\omega d\theta_\mu d\theta_N d\phi_N \sigma_\alpha(E_\nu)}.$$

To model $\sigma_\alpha(E_\nu)$, a standard approach is to integrate the cross section over the other variables and fit the resulting function using a polynomial approximation:

$$\sigma_\alpha(E_\nu) \approx P_\alpha(E_\nu).$$

E. Exhaustive training of the flow-based event generator

To train NF, we typically optimize the model by minimizing the Kullback-Leibler Divergence (KL-D) between the true distribution p and the learned distribution q_θ (where θ refers to the neural net parametrization of the surrogate distribution). The KL-D measures the difference between two probability distributions, but it is inherently asymmetric, which means it leads to different

behaviors depending on whether we minimize the forward or reverse KL-D. Minimizing the forward KL-D, expressed as

$$D_{KL}(p||q_\theta) = \mathbb{E}_{x \sim p} \left[\log \frac{p(x)}{q_\theta(x)} \right],$$

can lead to mean-seeking behavior, where the learned distribution q_θ covers the support of the true distribution p but may overestimate the distribution tails. Conversely, minimizing the reverse KL-D, expressed as

$$D_{KL}(q_\theta||p) = \mathbb{E}_{x \sim q_\theta} \left[\log \frac{q_\theta(x)}{p(x)} \right],$$

can result in mode-seeking behavior, where q_θ focuses on some modes of p but may ignore other regions.

To balance these tendencies, we use the symmetric KL-D, also known as Jensen-Shannon metric, which is the average of these two:

$$D_s(p, q_\theta) = \frac{1}{2} (D_{KL}(p||q_\theta) + D_{KL}(q_\theta||p)).$$

However, direct computation of the KL-D is impractical for our use case, as it would involve sampling from the true distribution p , which is computationally and time expensive. Furthermore, sampling from the learned distribution q_θ to compute the loss function introduces instability. This instability arises because, in the early training stages, the learned distribution may deviate significantly from the target, leading to unreliable loss estimates.

A significant improvement in our training process was achieved through the application of importance sampling to compute the loss function [40]. Importance sampling employs a proposal distribution that is simple to sample from yet effectively spans the target distribution p . In our case, we used a uniform distribution u defined over the manifold \mathcal{M} , as defined in Section III C, and sampled energy values uniformly within the chosen training range. This enabled us to compute a Monte Carlo estimate of the symmetric KL-D :

$$D_s(p, q_\theta) \approx \frac{1}{Z} \sum_{i,j} \left[(p(x_i|E_j) - q_\theta(x_i|E_j)) \log \frac{p(x_i|E_j)}{q_\theta(x_i|E_j)} \right],$$

where Z is a normalization constant, and x_i and E_j are sampled uniformly from \mathcal{M} and the training energy range, respectively.

This approach ensures robust and stable training by exhaustively covering the manifold, preventing the learned distribution q_θ from collapsing into narrow regions or missing areas.

To contextualize, a related approach was proposed by Pina-Otey et al. [37] in their Exhaustive Neural Importance Sampling (ENIS) framework. They employed a

two-stage training process: an initial "warm-up" phase, during which the training points were uniformly sampled for 20% of the total training time, followed by direct sampling from the learned distribution to compute the loss function. The first stage allowed exhaustive exploration of the phase space, while the second stage accelerated convergence by focusing on regions where the model had already predicted probability density.

Although the two-stage strategy is appealing, our results indicate that maintaining uniform sampling throughout the training process is not only simpler but also sufficiently fast for the dimensionalities involved in the exclusive $1p1h$ processes. Furthermore, in scenarios where cross section calculations are computationally expensive, this approach becomes impractical, as it would require evaluating the cross section repeatedly during training for new samples of the predicted probability density. Exhaustivity remains a critical factor in cross section modeling, where the primary goal is to develop the most accurate and comprehensive surrogate for the cross section, rather than merely optimizing training time. However, for higher-dimensional processes like two-particle two-hole ($2p2h$) interactions, uniformly sampling the kinematics may become infeasible. In such cases, an alternative approach could involve training the model using samples drawn from an already trained model for the semi-exclusive cross section, where the kinematics of the sub-leading nucleon are marginalized. This semi-exclusive distribution could then serve as a proposal distribution for training a normalizing flow to model the fully exclusive cross section.

IV. PERFORMANCES OF THE FLOW-BASED EVENT GENERATOR

A. Detail of implementation and training

We trained a flow-based Monte Carlo event generator for the CCQE $1p1h$ interaction between a neutrino and a carbon nucleus (^{12}C). The carbon nucleus consists of two shells, $1s_{1/2}$ and $1p_{3/2}$, with respective separation energies of 36.20 MeV and 18.72 MeV. To enhance the flexibility of our event generation, we chose to model the interactions for each shell separately. This approach serves theoretical purposes, enabling the study of interactions with individual shells. A similar work can be applied to any nuclei, which would involve modeling more or less shells.

Additionally, we trained four separate models per shell, each corresponding to overlapping neutrino energy ranges in MeV: [180, 350], [300, 500], [450, 750], and [700, 1050]. Consequently, the event generator covers neutrino energies from 180 MeV to 1050 MeV. The overlapping regions between these energy ranges ensure smooth transitions between models. In these overlap regions, we blend the two overlapping models by sampling from and evaluating the second model with a probability that increases lin-

early from 0 at the start of the overlap range to 1 at the end. ^{12}C having 2 shells, the total number of models is therefore 8.

Concerning the implementation of a single model, we used an implementation of Autoregressive RQ-NSF based on the `nflows` Pytorch implementation of Durkan et al. [41]. Some modifications were made to accommodate the computation of the Symmetric Kullback-Leibler divergence using importance sampling.

We trained three types of models, each differing in size: a small model (S), a medium model (M), and a large model (L). All three models utilize 9 bins in their spline transformations. Each RQ-NSF is parameterized by a masked autoregressive network composed of three hidden layers, with a residual connection from the first to the third layer. Model S includes 256 hidden nodes per hidden layer and 10 RQ-NSFs, Model M uses 512 hidden nodes and 10 RQ-NSFs, and Model L employs 512 hidden nodes and 25 RQ-NSFs. The performance of these three models is evaluated in Table I. The base distribution is a four-dimensional probability density with independent dimensions. The three angular dimensions are modeled as uniform distributions, while the energy transfer dimension follows a normal distribution.

We used a relatively low learning rate of 0.0001 with a Cosine annealing scheduler with a minimum learning rate of 5×10^{-7} . We trained each model for 100000 epochs. Each epoch computes the Symmetric KL-D on a batch of 2^{14} samples with associated true cross section. 10 millions samples were used to train each model in total. These samples and their cross sections were pre-generated prior to training—a potentially time-intensive step depending on the cross section computation speed, but required only once. The θ_μ and θ_N distributions are unfolded to $[-\pi, \pi]$ to ensure 2π -periodicity. The training time for each model (on an NVIDIA RTX-4090 GPU) are given in Table I.

B. Comparison of the True cross section and Its Surrogate Projecting in 1D and 2D Spaces

In Monte Carlo event generators widely used in neutrino physics, such as NEUT, users can freely select any distribution of neutrino energy. As a result, our model must accurately represent all possible combinations of energy and, in our case, the interacting nucleon's shell. While it would be impractical to visually compare the true and surrogate cross sections for every such combination, Sections IVC and IVD provide a more quantitative and exhaustive evaluation of the surrogate cross section's performance. This section offers a more visual and qualitative evaluation for a specific energy and shell combination.

Figure 2 illustrates the comparison between the true cross section and its surrogate modeled using Normalizing Flows for a fixed neutrino energy of $E_\nu = 200$ MeV, with the initial bound nucleon in the $1p_{3/2}$ shell. This




FIG. 2: 100 million samples with weights from the true cross section (a) and 10 million weighted samples from the predicted surrogate cross section (b) for a neutrino energy of $E_\nu = 200$ MeV and the $1p_{3/2}$ shell. The diagonal plots represent the marginal distribution of the 4 kinematic variables ($\theta_\mu, \omega, \theta_N, \phi_N$). The off-diagonal plots correspond to the 2D histograms of 2 different kinematic variables.

specific combination was selected due to the pronounced influence of nuclear effects on its cross section. At low energy transfer, the cross section exhibits a pronounced dependence on shell structure effects, a core addition of theoretical nuclear shell models such as the Hartree-Fock Mean-Field approach. The angular distribution of the outgoing nucleon relative to the transferred momentum, θ_N , is strongly influenced by the properties of the initial bound state. Nevertheless, the NF seem to model accurately this distribution from the shapes of the four marginal distributions to the correlations between two kinematic variables.

C. Comparison of the True cross section and Its Surrogate Using Density Estimation

To compare two multidimensional distributions, one often simplifies the problem by projecting the densities onto one or two dimensions at a time and/or by overlapping their binned histograms. These approaches, while visual, come with inherent limitations. By reducing the dimensionality, dependencies that might exist in higher-dimensional spaces are inevitably smoothed out. Similarly, binning locally averages out the densities and therefore can obscure fine-scale variations, introducing artificial agreement in regions where the true distributions may differ subtly.

Therefore, one needs a more holistic approach to com-

paring multidimensional distributions. In this work, we evaluate directly the weight which is defined as the ratio between the true distribution p and its Normalizing Flows surrogate q_θ (where θ refers to the neural net parametrization of the surrogate distribution) :

$$w(\theta_\mu, \omega, \theta_N, \phi_N, E_\nu, \alpha) = \frac{p(\theta_\mu, \omega, \theta_N, \phi_N | E_\nu, \alpha)}{q_\theta(\theta_\mu, \omega, \theta_N, \phi_N | E_\nu, \alpha)}$$

where α refers to the bound nucleon shell.

The histogram of weights for the **L** model along with the histogram of weights for a uniform distribution defined on the manifold \mathcal{M} are shown in Figure 3. In the case of the surrogate cross section, the weights are tightly concentrated around 1, with a measured standard deviation of 10.7% and a 99.99th percentile value of 1.90. This represents a significant advantage when compared to a uniform distribution, where the weights span 13 orders of magnitude, leading to a large proportion of samples contributing minimally or not at all.

From these weights, we derive the Relative Effective Sample Size (RESS), which provides a quantitative measure of the quality of the surrogate distribution q_θ in approximating p . The RESS quantifies the proportion of effectively independent samples represented by the surrogate distribution when weighted by the importance




FIG. 3: Log-scale histogram of 1 million weights for both shells of ^{12}C under a flat neutrino energy flux from 200 to 1000 MeV, generated using the surrogate cross section (red line) and a uniform distribution (blue dashed line) defined on the manifold \mathcal{M} .

weights. It is defined as:

$$\text{RESS} = \frac{1}{N} \frac{\left(\sum_{i=1}^N w_i\right)^2}{\sum_{i=1}^N w_i^2},$$

where w_i are the weights for the samples $i = 1, \dots, N$.

A RESS close to 1 indicates that the weights w_i are similar in magnitude, reflecting that the surrogate distribution q_θ closely approximates the true distribution p . In contrast, a RESS close to 0 suggests that the weights are spread out, implying a significant mismatch between q_θ and p . The performance of the models is evaluated using 1 million samples drawn from q_θ across the full neutrino energy range, [200, 1000] MeV, and for both shells combined. The RESS values for the three model sizes are presented in Table I. These values are computed for two of neutrino fluxes: one with a flat distribution and another based on the truncated T2K near-detector flux [42].

D. Comparison Between the True cross section and Its Surrogate Through Their Respective Datasets

In neutrino experiments such as T2K, we cannot achieve exact estimates of the kinematics we measure or infer. As a result, evaluating a model based on a direct point-by-point comparison of its distribution with the true one, as discussed in Section IV C, may be overly conservative. Instead, the goal in this work is to construct an event generator capable of populating the kinematic phase space in a manner consistent with a non-biased accept-reject method. This must be achieved with fixed precision while accounting for the full dimensionality of the phase space.

Therefore, we compare four-dimensional binned histograms derived from:

1. Samples generated by the *surrogate cross section*, and
2. Samples generated by an *accept-reject algorithm* with a uniform proposal distribution defined on the manifold \mathcal{M} .

We assess how closely these histograms align with a high-statistics, non-biased *reference histogram* by computing the Multinomial Negative Log-Likelihood (MNLL):

$$\text{MNLL}(H^{\text{MC}}) = \frac{1}{N} \sum_{\text{bin}} \left[\log(H_{\text{bin}}^{\text{MC}}!) - H_{\text{bin}}^{\text{MC}} \log\left(\frac{H_{\text{bin}}^{\text{ref}}}{N}\right) \right],$$

where the summation runs over all bins in the histogram and:

- H^{MC} is the histogram from the Monte Carlo dataset (either surrogate or accept-reject),
- H^{ref} is the high-statistics, non-biased reference histogram rescaled to match the number of samples in H^{MC} ,
- $N = \sum_{\text{bin}} H_{\text{bin}}^{\text{MC}}$ is the total number of observed samples.

To obtain statistical distributions of the MNLL for both the surrogate model and the accept-reject method, we repeat the procedure 100 times per method. Figure 4 illustrates how the MNLL evolves with neutrino energy for both approaches. In this context, the unbiased accept-reject MNLL serves as the optimal performance benchmark for the surrogate cross section.

This comparison is carried out at fixed neutrino energies of 200, 300, 400, 500, 600, 700, 800, 900 and 1000 MeV, considering the $1s_{1/2}$ and $1p_{3/2}$ shells separately. For each energy and shell, θ_μ and θ_N are binned with a width of 10° , while ϕ_N has a bin width of 30° . For the energy transfer, the bin width is 10 MeV from the separation energy up to 100 MeV, and then 50 MeV until reaching the maximum energy transfer of $E_\nu - m_\mu$, where E_ν is the neutrino energy and $m_\mu = 105.66$ MeV is the muon rest mass. We choose a finer binning at lower energy transfer to capture the complex structure inherited from the shell modeling. Each histogram H^{MC} uses 100,000 samples, while the reference histogram is built from 10 million weighted samples.

To provide a measure of how the surrogate's MNLL values differ from those of the accept-reject method in average, we compute the averaged Z-Score (Z-S) across all energies for a given shell:

$$\text{Z-Score} = \frac{1}{N_e} \sum_{i=1}^{N_e} \left(\frac{\overline{\text{MNLL}(H_{E_\nu^{(i)}}^{\text{NF}})} - \overline{\text{MNLL}(H_{E_\nu^{(i)}}^{\text{AR}})}}{\sigma_{E_\nu^{(i)}}^{\text{AR}}} \right).$$

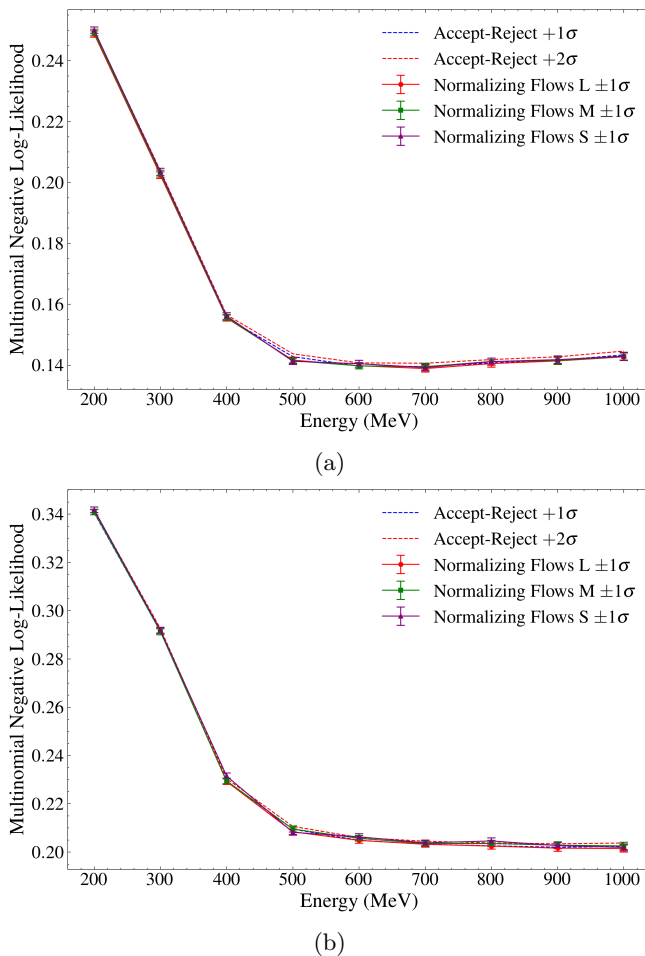


FIG. 4: Evolution of the Multinomial Negative Log-Likelihood with the neutrino energy for the $1s_{1/2}$ ((a)) and $1p_{3/2}$ ((b)) shell. Each plot shows both the MNLL for the for the three Normalizing Flows model size. The $+1\sigma$ and $+2\sigma$ upper bounds of the MNLL using the Accept-Reject datasets are given for comparison.

where $\overline{\text{MNLL}}(H_{E_\nu^{(i)}}^{\text{NF}})$ is the mean MNLL for Normalizing Flow datasets at energy $E_\nu^{(i)}$, $\overline{\text{MNLL}}(H_{E_\nu^{(i)}}^{\text{AR}})$ is the mean MNLL for accept-reject datasets at the same energy, and $\sigma_{E_\nu^{(i)}}^{\text{AR}}$ is the standard deviation of the accept-reject MNLL distribution at energy $E_\nu^{(i)}$. The Z-S for the three model sizes are given in Table I. The MNLL distributions for the surrogate and accept-reject histograms are statistically compatible within their Z-S = 1 range for the L model and within Z-S = 2 for the M and S models. In the case of the model size L , we find an average Z-S of 0.66σ for the $1s_{1/2}$ shell and 0.82σ for the $1p_{3/2}$ shell. In other words, two datasets of 100,000 samples respectively generated by accept-reject and by the Normalizing Flow surrogate are practically equivalent given the binning that we chose. Overall, these results demonstrate that the Normalizing Flow surrogate reproduces

the cross section with high fidelity in all tested combinations of shell and energy, while significantly reducing computational overhead compared to the accept-reject method.

E. Overall performance

The performances for the three model sizes are summarized in Table I. While accuracy improves slightly with increasing model size, this comes at the expense of reduced computational efficiency and larger model sizes. The sampling speeds are here given for a single neutrino energy and shell combination. As noted earlier, we prioritized modeling accuracy over sampling speed by not selecting the fastest RQ-NSF implementation. Nonetheless, even the slowest model, L , achieves a sampling rate of 10 million samples in under 12 minutes on an NVIDIA RTX-4090 GPU and in 1 hour 52 minutes on a single Intel Xeon CPU core, far surpassing the requirement of 10 million samples per day. This encourages us to consider this implementation as a viable approach for higher-dimensional processes, such as $2p2h$ or charged-current single-pion production.

Model	S	M	L
Size (MB)	200	719	1797
Number of flows	10	10	25
Number of hidden nodes	256	512	512
Training time (hour)	5	6	13
CPU Speed (sec / million samples)	141.8	235.6	669.3
GPU Speed (sec / million samples)	12.5	23.6	70.7
RESS, Flat Flux (%)	98.48	98.64	98.87
RESS, T2K Flux (%)	98.48	98.64	98.85
Z-Score ($1s_{1/2}$)	1.17	0.92	0.66
Z-Score ($1p_{3/2}$)	1.66	1.32	0.82

TABLE I: Comparison of model performance for different model sizes (S , M , L). The size is here given for the eight models (as discussed in Section IV A).

The results show that models with larger sizes provide higher accuracy, as indicated by a steady rise in RESS values from 98.48% (S) to 98.87% (L). However, this improvement is relatively small, suggesting that the smallest model might be sufficient when sampling speed or memory constraints outweigh minor accuracy gains. This trade-off between sampling speed and accuracy depends on how the surrogate cross section is utilized.

There are three ways to utilize the surrogate cross section:

1. Use the surrogate as a proposal distribution and apply an accept-reject algorithm to resample its outputs.
2. Sample from the surrogate and provide per-event weights by calling the true cross section once for

each sampled event.

3. *Sample directly from the surrogate without any further resampling.* This option is by far the fastest for cross sections that are computationally heavy.

Nevertheless, using the surrogate cross section for accept-reject sampling can be beneficial in some situations. In cases where the true cross section can be evaluated quickly, it may be most efficient to adopt the smallest model (**S**) and, if necessary, reweight its samples with an accept-reject algorithm. However, if evaluating the true cross section is time-consuming, the larger (**L**) model's higher accuracy can justify its additional computational and storage costs.

Moreover, our Normalizing Flows approach to build a proposal distribution already achieves greater accuracy than the previous Normalizing Flows applied to neutrino-nucleus cross section modelling: Pina-Otey *et al.* [37] reported a RESS of 91.40%, which is significantly lower than even our smallest model (**S**). The concentration of weights around 1 as shown in Figure 3 highlights the efficiency of the surrogate cross section as a proposal distribution for an accept-reject algorithm. The low spread of weights ensures that most samples contribute effectively to the estimation of observables, minimizing the variance introduced by the importance weights. If the weight threshold is set at 1.90, corresponding to the 99.99th percentile of weights then only one in ten thousand samples surpasses this limit. The expected acceptance rate for the surrogate cross section with this weight cap, when used as a proposal distribution, is approximately $\frac{1}{1.90} \approx 53\%$.

Although the surrogate can serve as an efficient proposal distribution, Section IV D shows that a dataset of events sampled directly from the Normalizing Flow is practically equivalent to an unbiased accept-reject dataset at our target precision. Consequently, we recommend sampling directly from the surrogate whenever possible. This is particularly advantageous for complex, theory-driven cross sections requiring lengthy evaluations, since Normalizing Flows can reduce total sampling time from several days on a CPU (via accept-reject) to a few minutes on a GPU or CPU.

V. DISCUSSION

The aim of this study was to demonstrate the ability to efficiently and accurately generate $1p1h$ samples without relying on the full complexity of the Ghent model, which also accounts for two-body currents and the Continuum Random Phase Approximation (CRPA) [12]. These corrections are not expected to pose major challenges for the proposed method. Two-body currents, involving two-particle states, provide corrections to the impulse approximation, correcting the single nucleon knockout processes and giving rise to double nucleon knockout processes. The effect of two-body currents on the $1p1h$ cross section is relatively small and can be accounted for by fine-tuning

a model already trained within the impulse approximation.

Similarly, although including more detailed physics features in the cross section, from a numerical point of view the inclusion of CRPA correlations is expected to simplify the energy transfer distribution by broadening the sharp features at low energy, smoothing out the delta-like structures associated with the missing energy from shell modeling, hence facilitating event generation using normalizing flows.

The greater challenge arises in describing $2p2h$ processes due to the increased number of dimensions in the phase space. These processes correspond to two-nucleon knockout, and their fully exclusive differential cross-sections involve nine independent kinematic variables:

$$\frac{d^9\sigma(E_\nu)}{d\omega d\Omega_\mu d\Omega_{N_1} dE_{N_1} d\Omega_{N_2} dE_{N_2}} \propto L_{\mu\nu} W^{\mu\nu}.$$

As in the $1p1h$ case, we treat the neutrino energy as a conditional variable, the differential cross section does not depend on the muon's azimuthal angle ϕ , and we incoherently sum over the energy eigenstates of the target nucleons. Consequently, we must handle seven independent kinematic variables, three more than in the $1p1h$ scenario. The increased dimensionality should not affect the expressiveness of autoregressive Normalizing Flows as shown in [34]. However, a larger phase space requires more training samples, and this challenge is made even more difficult by the longer computation time needed for the $2p2h$ cross section.

In this context, our work on $1p1h$ modeling provides a solid foundation for extending the approach to semi-exclusive cross sections, where the kinematics of the sub-leading nucleon are integrated out. Training a model on semi-exclusive cross sections captures key features of the $2p2h$ process and allows for a more efficient training of fully exclusive $2p2h$ cross sections. Instead of using uniformly distributed samples, as done in this work, the training can be performed using samples drawn from the surrogate semi-exclusive cross section, reducing the number of required evaluations and improving training efficiency.

VI. CONCLUSION

In this work, we presented a Normalizing Flow-based surrogate for generating neutrino-nucleus scattering events with high accuracy and efficiency. By comparing our surrogate's output to that of an unbiased accept-reject method, we demonstrated that the two methods produce practically equivalent samples across a wide range of neutrino energies and nuclear shells. Even more importantly, we showed that the Normalizing Flow approach can generate millions of events in minutes on modern GPUs, outperforming traditional Monte Carlo techniques in terms of computational speed.

From a theoretical standpoint, such surrogates enable systematic studies of model variations and uncertainty quantification, giving insights into how different theoretical inputs manifest across multi-dimensional phase space.

Looking ahead, the flexibility of Normalizing Flows makes them well-suited to other processes beyond the $1p1h$ process studied here. In particular, higher-dimensional processes such as $2p2h$ or single-pion production can benefit from a similar surrogate approach. Currently, neutrino experiments like T2K rely on “Frankenstein” models that stitches together submodels based on different theoretical frameworks, each tuned to a specific process, and combined with ad-hoc relative pro-

portions. In contrast, a Normalizing Flow-based surrogate can make possible a unified sampling scheme for *all* processes within a single, theory-grounded cross section model, such as the Mean Field cross section model developed by the Ghent group. Such a tool could be crucial for future neutrino experiments, delivering fast, accurate, and cohesive event generation that seamlessly bridges the gap between the experimental and theoretical frontiers.

ACKNOWLEDGMENTS

The author acknowledges support from the Swiss National Foundation grants No. 200021E_213196.

-
- [1] K. Abe *et al.* (T2K), The T2K Experiment, Nucl. Instrum. Meth. A **659**, 106 (2011), arXiv:1106.1238 [physics.ins-det].
 - [2] H.-K. Proto-Collaboration, Hyper-kamiokande design report (2018), arXiv:1805.04163 [physics.ins-det].
 - [3] L. Alvarez-Ruso *et al.* (NuSTEC), NuSTEC White Paper: Status and challenges of neutrino–nucleus scattering, Prog. Part. Nucl. Phys. **100**, 1 (2018), arXiv:1706.03621 [hep-ph].
 - [4] Y. Hayato, A neutrino interaction simulation program library NEUT, Acta Phys. Polon. B **40**, 2477 (2009).
 - [5] C. Andreopoulos *et al.*, The GENIE Neutrino Monte Carlo Generator, Nucl. Instrum. Meth. A **614**, 87 (2010), arXiv:0905.2517 [hep-ph].
 - [6] J. Żmuda, K. Graczyk, C. Juszczak, and J. Sobczyk, **NuWro** monte carlo generator of neutrino interactions — first electron scattering results, Acta Physica Polonica B **46**, 2329 (2015).
 - [7] S. Dolan, G. Megias, and S. Bolognesi, Implementation of the susav2-meson exchange current $1p1h$ and $2p2h$ models in genie and analysis of nuclear effects in t2k measurements, Physical Review D **101**, 10.1103/physrevd.101.033003 (2020).
 - [8] A. Nikolakopoulos, S. Gardiner, A. Papadopoulou, S. Dolan, and R. González-Jiménez, From inclusive to semi-inclusive one-nucleon knockout in neutrino event generators (2023), arXiv:2302.12182 [nucl-th].
 - [9] A. V. Butkevich and S. P. Mikheyev, Test of fermi gas model and plane-wave impulse approximation against electron-nucleus scattering data, Phys. Rev. C **72**, 025501 (2005).
 - [10] O. Benhar, A. Fabrocini, S. Fantoni, and I. Sick, Spectral function of finite nuclei and scattering of gev electrons, Nuclear Physics A **579**, 493 (1994).
 - [11] P. Ring, Relativistic mean field theory in finite nuclei, Progress in Particle and Nuclear Physics **37**, 193 (1996).
 - [12] N. Jachowicz, K. Heyde, J. Ryckebusch, and S. Rombouts, Continuum random phase approximation approach to charged-current neutrino-nucleus scattering, Phys. Rev. C **65**, 025501 (2002).
 - [13] A. Nikolakopoulos, R. González-Jiménez, N. Jachowicz, K. Niewczas, F. Sánchez, and J. M. Udías, Benchmarking intranuclear cascade models for neutrino scattering with relativistic optical potentials, Physical Review C **105**, 10.1103/physrevc.105.054603 (2022).
 - [14] A. Nikolakopoulos, A. Ershova, R. González-Jiménez, J. Isaacson, A. M. Kelly, K. Niewczas, N. Rocco, and F. Sánchez, Final-state interactions in neutrino-induced proton knockout from argon in microboone, Physical Review C **110**, 10.1103/physrevc.110.054611 (2024).
 - [15] D. Douqa, S. Bordoni, L. Giannessi, F. Sánchez, C. Schloesser, and A. Nikolakopoulos, Superscaling variable and neutrino energy reconstruction from theoretical predictions to experimental limitations, Physical Review D **109**, 10.1103/physrevd.109.073001 (2024).
 - [16] R. González-Jiménez, M. B. Barbaro, J. A. Caballero, T. W. Donnelly, N. Jachowicz, G. D. Megias, K. Niewczas, A. Nikolakopoulos, J. W. Van Orden, and J. M. Udías, Neutrino energy reconstruction from semi-inclusive samples, Physical Review C **105**, 10.1103/physrevc.105.025502 (2022).
 - [17] S. Dolan, A. Nikolakopoulos, O. Page, S. Gardiner, N. Jachowicz, and V. Pandey, Implementation of the continuum random phase approximation model in the GENIE generator and an analysis of nuclear effects in low-energy transfer neutrino interactions, Phys. Rev. D **106**, 073001 (2022), arXiv:2110.14601 [hep-ex].
 - [18] J. Schwehr, D. Cherdack, and R. Gran, Genie implementation of ific valencia model for qe-like $2p2h$ neutrino-nucleus cross section (2017), arXiv:1601.02038 [hep-ph].
 - [19] D. J. Rezende and S. Mohamed, Variational inference with normalizing flows (2016), arXiv:1505.05770 [stat.ML].
 - [20] J. Ryckebusch, M. Vanderhaeghen, L. Machenil, and M. Waroquier, Effects of the final-state interaction in (γ, pn) and (γ, pp) processes, Nuclear Physics A **568**, 828 (1994).
 - [21] V. Van der Sluys, J. Ryckebusch, and M. Waroquier, Two-body currents in inclusive electron scattering, Phys. Rev. C **51**, 2664 (1995).
 - [22] D. Vautherin and D. M. Brink, Hartree-fock calculations with skyrme’s interaction. i. spherical nuclei, Phys. Rev. C **5**, 626 (1972).
 - [23] K. Heyde, *The Nuclear Shell Model*, Springer Series in Nuclear and Particle Physics (Springer Berlin Heidelberg, 2012).
 - [24] M. Waroquier, G. Wenes, and K. Heyde, An effective skyrme-type interaction for nuclear structure calcula-

- tions: (ii). excited-state properties, Nuclear Physics A **404**, 298 (1983).
- [25] O. Moreno, T. W. Donnelly, J. W. Van Orden, and W. P. Ford, Semi-inclusive charged-current neutrino-nucleus reactions, Phys. Rev. D **90**, 013014 (2014).
- [26] J. E. Sobczyk, E. Hernández, S. X. Nakamura, J. Nieves, and T. Sato, Angular distributions in electroweak pion production off nucleons: Odd parity hadron terms, strong relative phases, and model dependence, Phys. Rev. D **98**, 073001 (2018).
- [27] K. Niewczas, A. Nikolakopoulos, J. T. Sobczyk, N. Jachowicz, and R. González-Jiménez, Angular distributions in Monte Carlo event generation of weak single-pion production, Phys. Rev. D **103**, 053003 (2021), arXiv:2011.05269 [hep-ph].
- [28] C. Gao, S. Höche, J. Isaacson, C. Krause, and H. Schulz, Event Generation with Normalizing Flows, Phys. Rev. D **101**, 076002 (2020), arXiv:2001.10028 [hep-ph].
- [29] B. Stienen and R. Verheyen, Phase space sampling and inference from weighted events with autoregressive flows, SciPost Phys. **10**, 038 (2021).
- [30] D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling, Improving variational inference with inverse autoregressive flow (2017), arXiv:1606.04934 [cs.LG].
- [31] G. Papamakarios, T. Pavlakou, and I. Murray, Masked autoregressive flow for density estimation, in *Advances in Neural Information Processing Systems*, Vol. 30 (2017) pp. 2338–2347.
- [32] L. Dinh, D. Krueger, and Y. Bengio, Nice: Non-linear independent components estimation (2015), arXiv:1410.8516 [cs.LG].
- [33] L. Dinh, J. Sohl-Dickstein, and S. Bengio, Density estimation using real nvp (2017), arXiv:1605.08803 [cs.LG].
- [34] A. Coccaro, M. Letizia, H. Reyes-González, and R. Torre, Comparison of affine and rational quadratic spline coupling and autoregressive flows through robust statistical tests, Symmetry **16**, 942 (2024).
- [35] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, Neural spline flows, in *Advances in Neural Information Processing Systems*, Vol. 32 (2019) pp. 7511–7522.
- [36] R. Delbourgo and J. A. Gregory, Piecewise rational quadratic interpolation to monotonic data, IMA Journal of Numerical Analysis **3**, 141 (1983).
- [37] S. Pina-Otey, F. Sánchez, T. Lux, and V. Gaitan, Exhaustive neural importance sampling applied to monte carlo event generation, Phys. Rev. D **102**, 013003 (2020).
- [38] M. El Baz and F. Sánchez, Fast posterior probability sampling with normalizing flows and its applicability in bayesian analysis in particle physics, Phys. Rev. D **109**, 032008 (2024).
- [39] D. J. Rezende, G. Papamakarios, S. Racanière, M. S. Albergo, G. Kanwar, P. E. Shanahan, and K. Cranmer, Normalizing flows on tori and spheres (2020), arXiv:2002.02428 [stat.ML].
- [40] T. Müller, B. McWilliams, F. Rousselle, M. Gross, and J. Novák, Neural importance sampling (2019), arXiv:1808.03856 [cs.LG].
- [41] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, nflows: normalizing flows in PyTorch (2020).
- [42] T. T. collaboration (T2K Collaboration), Measurements of neutrino oscillation in appearance and disappearance channels by the t2k experiment with 6.6×10^{20} protons on target, Phys. Rev. D **91**, 072010 (2015).