Exhaustive neural importance sampling applied to Monte Carlo event generation

Sebastian Pina-Otey^{1,2,*}, Federico Sánchez³, Thorsten Lux², and Vicens Gaitan¹

¹*Aplicaciones en Informática Avanzada (AIA), Sant Cugat del Vallès (Barcelona) 08172, Spain*

²*Institut de Física d'Altes Energies (IFAE)—Barcelona Institute of Science and Technology (BIST), Bellaterra (Barcelona) 08193, Spain*

³*University of Geneva, Section de Physique, DPNC, Geneva 1205, Switzerland*

(Received 27 May 2020; accepted 2 July 2020; published 16 July 2020)

The generation of accurate neutrino-nucleus cross section models needed for neutrino oscillation experiments requires simultaneously the description of many degrees of freedom and precise calculations to model nuclear responses. The detailed calculation of complete models makes the Monte Carlo generators slow and impractical. We present exhaustive neural importance sampling, a method based on normalizing flows to find a suitable proposal density for rejection sampling automatically and efficiently, and discuss how this technique solves common issues of the rejection algorithm.

DOI: [10.1103/PhysRevD.102.013003](https://doi.org/10.1103/PhysRevD.102.013003)

I. MOTIVATION

In modern science and engineering disciplines, the generation of random samples from a probability density function to obtain data sets or compute expectation values has become an essential tool. These theoretical models can be described by a target probability density function $p(\mathbf{x})$. Ideally, to generate samples following $p(\mathbf{x})$, the inverse transformation method is used. To perform the inverse transformation, the cumulative probability has to be calculated and the inverse to this function has to be found. Numerical methods have to be applied to obtain the Monte Carlo (MC) samples when this is not feasible computationally. This is especially true for high-dimensional spaces, where the integrals required to find such inverse transformation become analytically challenging.

A standard numerical method to obtain such data set is to perform a Markov Chain Monte Carlo (MCMC) algorithm [1], which provides good results for expected value calculations. Compared to other methods, it has the advantage that, in general, it requires very little calibration, and high dimensions can be broken down into conditional smaller dimension densities [2]. However, the MCMC method produces samples that form a correlated sequence. Also, the convergence of the samples' chain to the target density cannot be guaranteed for all possible models.

Another standard algorithm to produce MC samples is the acceptance-rejection or simply rejection sampling [3–6], which produces i.i.d. (independent and identically distributed) samples from the target density via an auxiliary proposal function. The proposal has to satisfy being a density which can both be sampled from and evaluated efficiently, as well as being as close to the target density as possible. The main disadvantages of the method are as follows [7]:

- (1) Designing the proposal function close to a particular target density can be very costly in human time.
- (2) If a generic proposal function is taken, such as a uniform distribution over the domain, the algorithm is usually very inefficient.
- (3) The sampling efficiency decreases rapidly with the number of dimensions.

Ideally, to avoid these inconveniences, one would like to have a method to find a proposal function that adapts to a given target density automatically. This would solve simultaneously the human time cost as well as the inefficiency of generic proposal densities.

An approach of the usage of normalizing flows to find a suitable proposal for a given target density has been suggested previously as neural importance sampling (NIS) [8], focused on the integration of functions via importance sampling [9]. Normalizing flows provide an expressive family of parametrized density functions $q_\phi(\mathbf{x})$ through neural networks, by defining a differentiable and invertible transformation from a base distribution to a target distribution, allowing to evaluate and sample from complex distributions by transforming simpler ones. The concept of integrating via importance sampling with normalizing flows for High Energy Physics (HEP) has been explored in

*pinas@aia.es

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

other works to obtain top-quark pair production and gluon-induced multijet production [10] or to simulate collider experimental observables for the Large Hadron Collider [11].

In this work, we further explore the possibility of utilizing normalizing flows to find a proposal function for a given target density to perform rejection sampling for MC samples, and analyze its viability through the following points:

- (i) We discuss the importance of adding an additional density (background) to the target one to assure the coverage of the whole phase space when performing rejection sampling.
- (ii) We define a two-phase training scheme for the normalizing flow to boost initial inefficiency in the optimization when adjusting the initialized random density toward the target one.
- (iii) We measure the performance of the method and argue for relaxing the rejection sampling constant factor k to improve largely the efficiency of acceptance while quantifying the error committed in doing this approximation via the concept of coverage.

Considering the proposed algorithm covers the whole domain of interest by modifying NIS with the background density, we denote this method by exhaustive neural importance sampling (ENIS).

We apply the above algorithm to an HEP problem, in the form of the charged current quasielastic (CCQE) cross-section for antineutrinos interactions with nuclei, performing in-depth analysis and discussion of the efficiency of the method. Neutrino-nucleus cross section modeling is one of the main sources of systematic uncertainties in neutrino oscillations measurements [12–14]. Cross section models are either analytically simple but describe the experimental data poorly or involving complex numerical computations, normally related to the description of the nucleus, that imposes limitations in their MC implementation. New tendencies in the field also call for a fully exclusive description of the interaction adding complexity to the calculations. The analytical model utilized in this paper is simple, but it is a realistic one and a good reference to demonstrate the capabilities of the proposed method to generate neutrino-nucleus cross sections efficiently. We will show that ENIS opens the possibility to incorporate efficiently complex theoretical models in the existing MC models enhancing the physics reach of running and future neutrino oscillation experiments.

ENIS algorithm may be used beyond the scope of neutrino physics. Further applications to be evaluated in detail in the future are particle/nuclear physics experiments, detector responses for medical physics, engineering studies, or theoretical modeling. In general, it could be applied to any Monte Carlo simulation that is limited by the algorithm's speed, such as for importance sampling to provide fast Monte Carlo with sufficient accuracy (i.e., fast

detector simulation, design studies, minimum bias background simulations, etc.). Additionally, the technique may help model developers extract expected values from their theoretical predictions in realistic conditions by including simple detector effects in models, such as effects of detector acceptance cuts, impact of model degrees of freedom on the predictions, or uncertainty propagation.

II. FRAMEWORK

In this section, we will describe background and framework needed for the rest of the paper. Section II A explains the physical model of charged current quasielastic neutrino interaction we will apply ENIS to in Sec. IV. As a summary and to introduce our notation, Sec. II B overviews the rejection sampling algorithm. Finally, in Sec. II C, we make a modest introduction to normalizing flows, focusing on the implementation of neural spline flows.

A. Model of charged current quasielastic antineutrinos interactions with nuclei

The charged current quasielastic (CCQE) is a basic model of neutrino interactions that might be expressed in simple formulas. The CCQE model has many advantages during this exploratory work, as it can be implemented in a simple software function, while at the same time it is also a realistic environment to understand the implications of modeling cross sections with the proposed methodology. The selected model to describe CCQE is the well-established Smith-Monith [15]. The nucleon momentum distribution follows a relativistic Fermi gas (noninteracting nucleons in a nuclear potential well) with a 0.225 GeV/c Fermi level. The model includes the Pauli blocking effect, preventing the creation of final state nucleons below the nucleus Fermi level. The model can be applied both to neutrino and antineutrino interactions. Antineutrinos are selected for this study due to the vector axial current cancellation imposing more stringent conditions at the edges of the kinematic phase space. The model includes the following degrees of freedom generated by the Monte Carlo model: the neutrino energy, the $\mu^\pm$ momentum and angle, and the target nucleon Fermi momentum. Contrary to other MC implementations, the neutrino energy is not a fixed input value but it is generated by the algorithm to add complexity to the calculations and to check the capabilities of the calculations to reproduce the cross section as a function of the neutrino energy. The implementation of this model for fixed energy value is also possible. The basic kinematic distributions obtained with this model will be discussed in Sec. IV.

B. Rejection sampling

Rejection sampling is a well-known technique [3–7] to obtain MC samples from a target density $p(\mathbf{x})$ which can be evaluated (up to a constant), but cannot be sampled through

the inverse transform. It relies on an auxiliary proposal function $q(\mathbf{x})$, from which one should be able to sample from and evaluate efficiently. A constant $k > 0$ is introduced which has to satisfy that

$$k \cdot q(\mathbf{x}) \geq p(\mathbf{x}) \quad \forall \mathbf{x}: p(\mathbf{x}) > 0. \quad (1)$$

The resulting function $k \cdot q(\mathbf{x})$ is called the comparison function.

The procedure to sample from the target density is then as follows:

- (1) A sample $\mathbf{x}$ is generated following $q(\mathbf{x})$, $\mathbf{x} \sim q(\mathbf{x})$.
- (2) A random number u is generated uniformly in the range $[0, k \cdot q(\mathbf{x})]$, $u \sim \text{Unif}(0, k \cdot q(\mathbf{x}))$.
- (3) If u fulfills the condition $u \leq p(\mathbf{x})$, the sample is accepted; otherwise, it is rejected.

Additionally, if $p(\mathbf{x})$ is normalized, the probability that a sample is accepted is proportional to $p_{\text{acc}} \propto 1/k$, i.e., k gives an intuition of the number of tries until we obtain an accepted sample.

C. Neural density estimation using neural spline flows

A family of density functions $q_\phi(\mathbf{x})$ over the real D -dimensional space $\mathbb{R}^D$ parametrized by ϕ satisfies that $q_\phi(\mathbf{x}) \geq 0$ for all $\mathbf{x}, \phi$, and that $\int q_\phi(\mathbf{x}) d\mathbf{x} = 1$ for all ϕ . Normalizing flows are a mechanism of constructing such flexible probability density families $q_\phi(\mathbf{x})$ for continuous random variables $\mathbf{x} \in \mathbb{R}^D$. A comprehensive review on the topic can be found in [16], from which a brief summary will be shown in this section on how normalizing flows are defined, and how the parameters ϕ are obtained, together with a specific implementation, the neural spline flows (NSFs) [17].

Consider a random variable $\mathbf{u}$ defined over $\mathbb{R}^D$, with known probability density $p_{\mathbf{u}}(\mathbf{u})$. A normalizing flow characterizes itself by a transformation T from another density $p(\mathbf{x})$ of a random variable $\mathbf{x}$ defined also over $\mathbb{R}^D$, the target density, to this known density, via

$$\mathbf{u} = T(\mathbf{x}), \quad \text{with } \mathbf{x} \sim p(\mathbf{x}). \quad (2)$$

The density $p_{\mathbf{u}}(\mathbf{u})$ is known as base density and has to satisfy that it is easy to evaluate (e.g., a multivariate D -dimensional normal distribution, as will be chosen through this work, or a uniform distribution in dimension D). The transformation T has to be invertible, and both T and T^{-1} have to be differentiable, i.e., T defines a diffeomorphism over $\mathbb{R}^D$.

This allows us to evaluate the target density by evaluating the base density using the change of variables for density functions,

$$p(\mathbf{x}) = p_{\mathbf{u}}(T(\mathbf{x})) |\det J_T(\mathbf{x})|,$$

where the Jacobian $J_T(\mathbf{x})$ is a $D \times D$ matrix of the partial derivatives of the transformation T ,

$$J_T(\mathbf{x}) = \begin{bmatrix} \frac{\partial T_1}{\partial x_1} & \cdots & \frac{\partial T_1}{\partial x_D} \\ \vdots & \ddots & \vdots \\ \frac{\partial T_D}{\partial x_1} & \cdots & \frac{\partial T_D}{\partial x_D} \end{bmatrix}.$$

The transformation T in a normalizing flow is defined partially through a neural network with parameters ϕ , as will be described below, defining a density

$$q_\phi(\mathbf{x}) = p_{\mathbf{u}}(T_\phi(\mathbf{x})) |\det J_{T_\phi}(\mathbf{x})|. \quad (3)$$

The subindex of T_ϕ will be omitted in the following, simply denoting the transformation of the neural network by T .

If the transformation is flexible enough, the flow could be used to evaluate any continuous density in $\mathbb{R}^D$. In practice, however, the property that the composition of diffeomorphisms is a diffeomorphism is used, allowing to construct a complex transformation via composition of simpler transformations. Consider the transformation T as a composition of simpler T_k transformations,

$$T = T_K \circ \cdots \circ T_1.$$

Assuming $\mathbf{z}_0 = \mathbf{x}$ and $\mathbf{z}_K = \mathbf{u}$, the forward evaluation and Jacobian are

$$\mathbf{z}_k = T_k(\mathbf{z}_{k-1}), \quad k = 1 : K,$$

$$|J_T(\mathbf{x})| = \left| \prod_{k=1}^K J_{T_k}(\mathbf{z}_{k-1}) \right|.$$

These two computations (plus their inverse) are the building blocks of a normalizing flow [18]. Hence, to make a transformation efficient, both operations have to be efficient. From now on, we will focus on a simple transformation $\mathbf{u} = T(\mathbf{x})$, since constructing a flow from it is simply applying compositions.

To define a transformation satisfying both operations to be efficient, the transformation is broken down into autoregressive one-dimensional (1D) ones for each dimension of $\mathbb{R}^D$,

$$u_i = \tau(x_i; \mathbf{h}_i) \quad \text{with} \quad \mathbf{h}_i = c_i(\mathbf{x}_{<i}; \phi),$$

where u_i is the i th component of $\mathbf{u}$ and x_i the i th of $\mathbf{x}$. τ is the transformer, which is a one-dimensional diffeomorphism with respect to x_i with parameters $\mathbf{h}_i$. c_i is the i th conditioner, a neural network, which takes as input $\mathbf{x}_{<i} = (x_1, x_2, \dots, x_{i-1})$, i.e., the previous components of $\mathbf{x}$, and ϕ are the parameters of the neural network. The conditioner provides the parameters $\mathbf{h}_i$ of the

i th transformer of x_i depending on the previous components $\mathbf{x}_{<i}$, defining implicitly a conditional density over x_i with respect to $\mathbf{x}_{<i}$. The transformer is chosen to be a differentiable monotonic function, since then it satisfies the requirements to be a diffeomorphism. The transformer also satisfies that it makes the transformation easily computable in parallel and decomposing the transformation in one-dimensional autoregressive transformers allows the computation of the Jacobian to be trivial, because of its triangular shape. To compute the parameter $\mathbf{h}_i$ of each transformer, one would need to process a neural network with input $\mathbf{x}_{<i}$ for each component, a total of D times.

Masked autoregressive neural networks [19] enable us to compute all the conditional functions simultaneously in a single forward iteration of the neural network. This is done by masking out, with a binary matrix, the connections of the $\mathbf{h}_i$ th output with respect to all the components with index bigger or equal to i , $\geq i$, making it a function of the $<i$ components.

The transformer can be defined by any monotonic function, such as affine transformations [20], monotonic neural networks [21–23], sum-of-squares polynomials [24], or monotonic splines [8,17,25]. In this work, we will focus on a specific implementation of monotonic splines, the neural spline flows.

In their work on neural spline flows [17], Durkan *et al.* advocate for utilizing monotonic rational-quadratic splines as transformers τ , which are easily differentiable, more flexible than previous attempts of using polynomials for these transformers, since their Taylor-series expansion is infinite, and are analytically invertible.

Each monotonic rational-quadratic function in the splines is defined by a quotient of two quadratic polynomial. In particular, the splines map the interval $[-B, B]$ to $[-B, B]$, and outside of it the identity function is considered. The splines are parametrized following Gregory and Delbourgo [26], where K different rational-quadratic functions are used, with boundaries set by the pair of coordinates $\{(x^{(k)}, u^{(k)})\}_{k=0}^K$, known as knots of the spline and are the points where it passes through. Note that $(x^{(0)}, u^{(0)}) = (-B, -B)$ and $(x^{(K)}, u^{(K)}) = (B, B)$. Additionally, we need $K - 1$ intermediate positive derivative values, since the boundary points derivatives are set to 1 to match the identity function.

Having this in mind, the conditioner given by the neural network outputs a vector $\mathbf{h} = [\mathbf{h}^w, \mathbf{h}^h, \mathbf{h}^d]$ of dimension $(3 \times K - 1)$ for the transformer τ , $c_i(\mathbf{x}_{<i}; \phi) = \mathbf{h}_i$. $\mathbf{h}^w$ and $\mathbf{h}^h$ give the width and height of the K bins, while $\mathbf{h}^d$ is the positive derivative at the intermediate $(K - 1)$ knots.

Stacking up many of these transformations, a highly flexible neural density estimator, the NSF, can be build, which satisfies the following:

- (1) It is easy to sample from $q_\phi(\mathbf{x})$ using the inverse transform T^{-1} in Eq. (2) by sampling $\mathbf{u} \sim p_{\mathbf{u}}(\mathbf{u})$.

- (2) Equation (3) allows to evaluate the densities $q_\phi(\mathbf{x})$ of these samples when generating them in an efficient way.

This density estimator will be the one utilized during this work.

III. METHODOLOGY

With the framework introduced in Sec. II, we are now in a position to define the ENIS method and the different metrics we will use to measure its performance.

We start in Sec. III A by showing the objective function to be minimized by the NSF to adjust its proposal function $q_\phi(\mathbf{x})$ to the target density $p(\mathbf{x})$. Then, in Sec. III B, we discuss the importance of adding background noise to both ensure coverage of the whole phase space of $p(\mathbf{x})$ and to boost the initial phase of the training of ENIS. The exact training scheme is then shown in Sec. III C, differentiating the warm-up phase from the iterative phase. Finally, in Sec. III D, the performance metrics are introduced, explaining the concept of coverage and effective sample size when considering a more relaxed condition on the rejection constant k .

A. Optimizing the parameters of NSF

Consider a target probability density function $p(\mathbf{x})$ which can be evaluated for all $\mathbf{x}$ but from which we are unable to generate samples directly through an analytical inverse transform. If we could approximate this target density by our neural density estimator $q_\phi(\mathbf{x})$, then we could exactly sample from the target density using rejection sampling, since we can both sample and evaluate from $q_\phi(\mathbf{x})$.

To obtain the parameters ϕ of $q_\phi(\mathbf{x})$ given a density $p(\mathbf{x})$ which can be evaluated, we want to minimize the Kullback-Leibler divergence (KL divergence) [27] between both distributions, which is ≥ 0 and only equal to zero if both distributions match,

$$D_{\text{KL}}(p(\mathbf{x}) \| q_\phi(\mathbf{x})) = \int p(\mathbf{x}) \log \left(\frac{p(\mathbf{x})}{q_\phi(\mathbf{x})} \right) d\mathbf{x}. \quad (4)$$

When minimizing with respect to ϕ , the KL divergence is simplified to

$$\begin{aligned} & \underset{\phi}{\operatorname{argmin}} D_{\text{KL}}(p(\mathbf{x}) \| q_\phi(\mathbf{x})) \\ &= \underset{\phi}{\operatorname{argmin}} \int p(\mathbf{x}) \log \left(\frac{p(\mathbf{x})}{q_\phi(\mathbf{x})} \right) d\mathbf{x} \\ &= \underset{\phi}{\operatorname{argmin}} - \int p(\mathbf{x}) \log q_\phi(\mathbf{x}) d\mathbf{x} \\ &= \underset{\phi}{\operatorname{argmax}} \int p(\mathbf{x}) \log q_\phi(\mathbf{x}) d\mathbf{x} \\ &= \underset{\phi}{\operatorname{argmax}} \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [\log q_\phi(\mathbf{x})]. \end{aligned}$$

The last expression could be approximated numerically if we could sample $\mathbf{x} \sim p(\mathbf{x})$, since it corresponds to approximating an expected value of a function we can evaluate, $\log q_\phi(\mathbf{x})$, and is equal to maximizing the log-likelihood of these samples,

$$L = \frac{1}{N} \sum_{i=1}^N \log(q_\phi(\mathbf{x}_i)) \quad \text{with} \quad \mathbf{x}_i \sim p(\mathbf{x}). \quad (5)$$

Müller *et al.* [8] propose a solution for computing the gradient with respect to ϕ for this maximization problem. They suggest using importance sampling [9] for this particular expected value,

$$\begin{aligned} \nabla_\phi \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [\log q_\phi(\mathbf{x})] &= \int p(\mathbf{x}) \nabla_\phi \log q_\phi(\mathbf{x}) d\mathbf{x} \\ &= \int q_\phi(\mathbf{x}) \frac{p(\mathbf{x})}{q_\phi(\mathbf{x})} \nabla_\phi \log q_\phi(\mathbf{x}) d\mathbf{x} \\ &= \mathbb{E}_{\mathbf{x} \sim q_\phi(\mathbf{x})} [w(\mathbf{x}) \nabla_\phi \log q_\phi(\mathbf{x})] \\ &\approx \frac{1}{N} \sum_{i=1}^N w(\mathbf{x}_i) \nabla_\phi \log q_\phi(\mathbf{x}_i), \end{aligned} \quad (6)$$

with $\mathbf{x}_i \sim q_\phi(\mathbf{x})$ and the weights defined as $w(\mathbf{x}) = p(\mathbf{x})/q_\phi(\mathbf{x})$. Notice how we only need to be able to evaluate $p(\mathbf{x})$ to compute this quantity. With this gradient, we are able to minimize the KL divergence in Eq. (4) if the support of $q_\phi(\mathbf{x})$ (i.e., the domain where the function is nonzero) contains the support of $p(\mathbf{x})$ to perform the importance sampling of Eq. (6) correctly. Notice that, in order to properly optimize the parameters ϕ , $p(\mathbf{x})$ does not need to be normalized, since this simply changes the magnitude of the gradient, but not its direction. The lack of proper normalization can be properly handled by standard neural network optimizers such as Adam [28].

The method described by Eq. (6) implies an iterative way of optimizing $q_\phi(\mathbf{x})$ with the following steps:

- (1) A batch of $\mathbf{x}$ is generated according to the current state of the neural network, $q_\phi(\mathbf{x})$.
- (2) Using this batch, the parameters ϕ of the neural network are optimized via the gradient of Eq. (6).
- (3) This updated neural network then generates the next batch.

B. Relevance of background noise

As briefly discussed in Sec. II, in order to optimize the neural network following Eq. (6), the gradient is only correctly computed if the support of $q_\phi(\mathbf{x})$ contains the one of $p(\mathbf{x})$. Moreover, if we want to use $q_\phi(\mathbf{x})$ as our proposal function to sample from $p(\mathbf{x})$ via rejection sampling, this also has to hold.

To ensure the proper $p(\mathbf{x})$ support, we introduce the concept of a background density function, $p_{\text{bg}}(\mathbf{x})$. In HEP,

as in many other scientific areas, the density is restricted to a certain domain of $\mathbf{x} \in \mathbb{R}^D$, e.g., the cosine has to be in $[-1, 1]$, the magnitude of the momentum in an experiment has to be positive and has a maximum value of p_{max} , there are constraints in the conservation of energy and momentum, etc. Hence, $p_{\text{bg}}(\mathbf{x})$ should be a density that has a support beyond these phase-space boundaries. In what follows, a uniform distribution will be considered, with limits in each dimension according to the phase space of that coordinate. The selection of the functional form of the $p_{\text{bg}}(\mathbf{x})$ is arbitrary and it can be selected to adapt to the requirements of each project.

The background density $p_{\text{bg}}(\mathbf{x})$ will be used for two tasks:

- (1) Improve initial training: At the beginning of the training, we cannot assure that the support of $q_\phi(\mathbf{x})$ contains the one of $p(\mathbf{x})$. Hence, instead of using $q_\phi(\mathbf{x})$ for the importance sampling of Eq. (6), $p_{\text{bg}}(\mathbf{x})$ will be used during the warm-up phase of the training. The distribution of the weight function $w(\mathbf{x}) = p(\mathbf{x})/p_{\text{bg}}(\mathbf{x})$ might span several orders of magnitude, but this way we ensure the full support of $p(\mathbf{x})$. This strategy gives a better approximation than the one obtained by the randomly initialized neural network $q_\phi(\mathbf{x})$ at the start of the training.
- (2) Ensure exhaustive coverage of the phase space: The target density $p_{\text{target}}(\mathbf{x})$ that the neural network will learn will be constructed as a linear combination of the true target density $p(\mathbf{x})$ and the background $p_{\text{bg}}(\mathbf{x})$,

$$p_{\text{target}}(\mathbf{x}) = (1 - \alpha) \cdot p(\mathbf{x}) + \alpha \cdot p_{\text{bg}}(\mathbf{x}), \quad (7)$$

with $\alpha \in (0, 1)$. This implementation adds a certain percentage α of background noise to the target density, spreading it over all the domain of the background density, allowing to properly apply the methods rejection and importance sampling with $q_\phi(\mathbf{x})$ as the proposal function, covering exhaustively the phase space. Experimentally, we have found good compromise with $\alpha = 0.05$.

Optimizing $q_\phi(\mathbf{x})$ to match $p_{\text{target}}(\mathbf{x})$ of Eq. (7) instead of $p(\mathbf{x})$ will make the proposal $q_\phi(\mathbf{x})$ slightly worse for rejection/importance sampling. By performing the optimization to $p(\mathbf{x})$ directly in an iterative way, as explained at the end of the last section, some regions of the phase space might disappear for future samplings. These regions are located normally close to the boundaries of sampled volume. Having a constant background noise prevents these losses from appearing, as the neural network has to also learn to generate this noise, covering properly the required phase-space volume. We will discuss the impact of the background term on the method performance in Sec. IV.

C. ENIS training scheme of the proposal function

The training procedure to obtain $q_\phi(\mathbf{x})$ from $p(\mathbf{x})$ following ENIS consists of two phases.

- (1) Warm-up phase
 - (a) Sample $\mathbf{x}_p \sim p_{\text{bg}}(\mathbf{x})$ and compute their weights $w_p(\mathbf{x}_p) = p(\mathbf{x}_p)/p_{\text{bg}}(\mathbf{x}_p)$.
 - (b) Sample background $\mathbf{x}_{\text{bg}} \sim p_{\text{bg}}(\mathbf{x})$ with associated weights $w_{\text{bg}}(\mathbf{x}_{\text{bg}}) = C_{w_{\text{bg}}} \cdot p_{\text{bg}}(\mathbf{x}_{\text{bg}})$, where $C_{w_{\text{bg}}} = \frac{\alpha}{1-\alpha} \frac{\langle w_p(\mathbf{x}_p) \rangle}{\langle p_{\text{bg}}(\mathbf{x}_{\text{bg}}) \rangle}$.
 - (c) Optimize the parameters of $q_\phi(\mathbf{x})$ via Eq. (6) using $\mathbf{x} = \{\mathbf{x}_p, \mathbf{x}_{\text{bg}}\}$ with weights $w(\mathbf{x}) = \{w_p(\mathbf{x}_p), w_{\text{bg}}(\mathbf{x}_{\text{bg}})\}$.
- (2) Iterative phase
 - (a) Sample $\mathbf{x}_q \sim q_\phi(\mathbf{x})$ and compute their weights $w_q(\mathbf{x}_q) = p(\mathbf{x}_q)/q_\phi(\mathbf{x}_q)$.
 - (b) Sample background $\mathbf{x}_{\text{bg}} \sim p_{\text{bg}}(\mathbf{x})$ with associated weights $w_{\text{bg}}(\mathbf{x}_{\text{bg}}) = C'_{w_{\text{bg}}} p_{\text{bg}}(\mathbf{x}_{\text{bg}})$, where $C'_{w_{\text{bg}}} = \frac{\alpha}{1-\alpha} \frac{\langle w_q(\mathbf{x}_q) \rangle}{\langle p_{\text{bg}}(\mathbf{x}_{\text{bg}}) \rangle}$.
 - (c) Optimize the parameters of $q_\phi(\mathbf{x})$ via Eq. (6) using $\mathbf{x} = \{\mathbf{x}_q, \mathbf{x}_{\text{bg}}\}$ with weights $w(\mathbf{x}) = \{w_q(\mathbf{x}_q), w_{\text{bg}}(\mathbf{x}_{\text{bg}})\}$.

Figure 1 depicts a flow diagram of the training method for ENIS, showing on the left block the warm-up phase, while

on the right block the iterative phase. The phase transition from warm-up to iterative phase is chosen heuristically. In our particular implementation of Sec. IV A, we have chosen the warm-up phase to comprise 20% of the total training iterations.

Steps 1 (b) and 2 (b) allow the method to add background following Eq. (7) to construct $p_{\text{target}}(\mathbf{x})$ even if $p(\mathbf{x})$ is not normalized.

D. Measuring the performance of the proposal function

The proposed method is not to use $q_\phi(\mathbf{x})$ as a direct approximation of $p(\mathbf{x})$, but as proposal function to perform either rejection (Sec. II B) or importance sampling [9]. This allows for the methods to correct any deviation in the neural network modeling of the exact density while utilizing its proximity to such density.

We use the learned probability density function $q_\phi(\mathbf{x})$ to generate samples via rejection sampling (see Sec. II B), which, in HEP, is of high interest and costly via standard procedures. The parameter k of the rejection algorithm has to be estimated empirically. Consider n samples $\{\mathbf{x}_i\}_{i=1}^n$ generated with the proposal function $\mathbf{x} \sim q(\mathbf{x})$, with weights $w(\mathbf{x}_i) = p(\mathbf{x})/q(\mathbf{x})$, satisfying the following:

FIG. 1. Exhaustive neural importance sampling flow diagram. The warm-up phase is depicted on the left block and the iterative phase on the right block. See text for a detail description.

(i) The average of the weights is

$$\langle w \rangle = \frac{1}{N} \sum_{i=1}^N w(\mathbf{x}_i) \approx \int q(\mathbf{x}) w(\mathbf{x}) d\mathbf{x} = C,$$

where C is the normalization of the density $p(\mathbf{x})$, i.e., its volume.

(ii) $k_{\max}$, the smallest constant $k > 0$ such that the inequality of Eq. (1) holds, is equal to $(\max w(\mathbf{x}_i))^{-1}$.

In real conditions, the parameter k can be relaxed. Instead of choosing the maximum value among the empirically computed weight distribution, it can be taken as the inverse of the Q quantile of these weights, w_Q , denoted by k_Q ,

$$k_Q = (Q \text{ quantile}(w))^{-1} = w_Q^{-1}. \quad (8)$$

This is equivalent of clipping the weights' maximum value to the Q quantile of w , capping the desired density function $p(\mathbf{x})$ we are generating using these weights for the rejection. The new weights $w'(\mathbf{x})$ are simply

$$w'(\mathbf{x}_i) = \begin{cases} w(\mathbf{x}_i) & \text{if } w(\mathbf{x}_i) \leq w_Q \\ w_Q & \text{if } w(\mathbf{x}_i) > w_Q \end{cases}.$$

The ratio of volume with respect to the original density $p(\mathbf{x})$ we are maintaining by clipping the weights this way defines the coverage we have of the rejection sampling and is equal to

$$\text{Coverage} = \frac{\sum_{i=1}^N w'(\mathbf{x}_i)}{\sum_{i=1}^N w(\mathbf{x}_i)}. \quad (9)$$

This allows us to quantify the error we are committing when choosing a quantile over the maximum of weights when defining a constant k for rejection sampling.

The idea behind relaxing this constant k is that we will approximate wrongly only a small region of $p(\mathbf{x})$ with $q(\mathbf{x})$. In that small region, the ratio $p(\mathbf{x})/q(\mathbf{x})$ is large compared to the rest of the domain but still it is occupying a small volume of the density $p(\mathbf{x})$. This region can be ignored by relaxing k , making the overall ratio of $p(\mathbf{x})/(k \cdot q(\mathbf{x}))$ closer to 1 and improving drastically the rejection sampling at the cost of this small discrepancy which we are committing, quantified in Eq. (9).

As an additional qualitative measurement of the goodness of different proposals under different constants k , the effective sample size (ESS) will be used [29], which corresponds approximately to the number of independent samples drawn. The ESS for n samples of weights $\{w(\mathbf{x}_i)\}_{i=1}^n$ is defined as

$$N_{\text{ESS}} = \left(\sum_{i=1}^N w(\mathbf{x}_i) \right)^2 / \sum_{i=1}^N w(\mathbf{x}_i)^2. \quad (10)$$

This is a rule of thumb to obtain the number of independent samples. The closer N_{ESS} is to the number of samples n , the more uncorrelated the weighted samples are. If large weights appear, then the independence of the samples will be diminished, as a same sample gets represented many times. We define N_{ESS}/N as a rough estimate for the ratio of independence of the samples.

IV. MONTE CARLO GENERATION OF THE CCQE ANTINEUTRINO CROSS SECTION

We will now proceed to apply ENIS to the CCQE antineutrino cross section density. In Sec. IV A, we discuss how the training for the NSF was performed, describing the background we added to cover the phase space. We show qualitatively the obtained densities and compare them to the target one. After obtaining a suitable proposal, we discuss in depth the performance of the obtained result in Sec. IV B, comparing the ENIS proposal to a generic uniform one, demonstrating its potential while justifying the relaxation on the constant k for the rejection sampling.

A. Training

To find the proposal function $q_\phi(\mathbf{x})$ via NSF for the CCQE antineutrino interaction crosssection density, described in Sec. II A, we followed the training scheme from Sec. III C.

The background chosen is a uniform distribution, covering a range of $[0, 10]$ for the incoming neutrino energy E_ν (in GeV), $[0, 10]$ for the outgoing muon momentum p_μ (in GeV/c), $[0, \pi]$ for the angle of the outgoing muon θ_μ (in rad), and $[0, 0.225]$ for the target nucleon Fermi momentum p_{nucleon} (in GeV/c). These bounds were expanded by covering a slightly more extended domain, of an additional 2% at the beginning and end of each dimension, to assure that the physical boundaries are completely covered. This expanded background was added with a $\alpha = 0.05$ contribution to the cross section density in Eq. (7), as well as used during training for the warm-up phase.

As for the hyperparameters of the NSF, we have chosen the ADAM optimizer [28] with learning rate 0.0005, batch size 5000, training steps 400 000, 5 flow steps, 2 transform blocks, 32 hidden features, and 8 bins. This gives a total dimension of 37220 for the parameters ϕ of $q_\phi(\mathbf{x})$. This configuration for the NSF was chosen experimentally to have a relatively low number of parameters (one can have easily six million parameters instead of the ≈ 37000 we have) since a lower number speeds up the generation and evaluation of samples $\mathbf{x} \sim q_\phi(\mathbf{x})$. Additionally, the learning rate was decreased during the training using a cosine scheduler to ensure stabilization at the end of the training procedure.

The training consists in maximizing the log-likelihood of Eq. (5) by computing its gradient via Eq. (6) and is shown over the 400 000 iterations in Fig. 2. In the gray area, the

FIG. 2. Neural spline flow training log probability for estimating the modified CCQE cross section with background noise, following Eq. (5). In gray, the warm-up phase is performed, using $p_{bg}(\mathbf{x})$ to generate the weighted samples, while in the white area the current state of the NSF $q_\phi(\mathbf{x})$ is used. The log probability stabilizes during the training to converge to a certain value which depends on the expressiveness of the network and the normalization of the target density.

training is performed with samples of the background distribution $p_{bg}(\mathbf{x})$, while in the white area the samples of the training samples are generated by the current proposal distribution $q_\phi(\mathbf{x})$. Notice that since the samples are generated in real time during the training, there is no need to worry about possible overfitting of the parameters of the neural network, which is a common issue in many machine-learning applications. The values of Fig. 2 are computed every one thousand steps, for a batch of 200000 samples $\mathbf{x} \sim q_\phi(\mathbf{x})$. The log probability can be seen to converge at the end of the training, which is mainly due to the cosine scheduler, but also due to the saturation over the family of parametrized densities $q_\phi(\mathbf{x})$.

To have a visual representation, Fig. 3 shows the marginalized one-dimensional densities of the four cross section variables of the target density $p(\mathbf{x})$ (blue) vs the NSF proposal $q_\phi(\mathbf{x})$ (orange). The plots show a small discrepancy in each variable, but an overall agreement between the two densities. Aside from a mismodeling on the side of $q_\phi(\mathbf{x})$ in certain regions, the differences can also come from the fact that the NSF is learning a modified target density [Eq. (7)].

To assess qualitatively that the correlation between the variables are also captured by $q_\phi(\mathbf{x})$, Fig. 4 shows 2D histograms for both the real density $p(\mathbf{x})$ (left) and the proposal density $q_\phi(\mathbf{x})$ (right). Visually, an overall agreement can be seen. There is a slight discrepancy for high energy $p_{nucleon}$ values, where the attenuation indicates that for the NSF proposal function it is more spread due to the background noise $p_{bg}(\mathbf{x})$ it is also learning [Eq. (7)].

In what follows, the performance of the NSF proposal will be discussed in more quantitative ways and compared to a uniform proposal.

FIG. 3. $p(\mathbf{x})$ (blue) vs $q_\phi(\mathbf{x})$ (orange) one-dimensional normalized histograms of the marginalized CCQE cross section density for each of the variables. The plots show light discrepancy in each variable, but an overall agreement between the NSF proposal $q_\phi(\mathbf{x})$ and the CCQE cross-section density $p(\mathbf{x})$. Notice how the distribution of $q_\phi(\mathbf{x})$ is taken before performing rejection sampling on it.

B. Performance and discussion

In this section, we will focus on analyzing the performance of the proposal density obtained by the NSF while also comparing it to a uniform proposal density, $p_{Unif}(\mathbf{x})$, which in our case will be the same as $p_{bg}(\mathbf{x})$, defined in Sec. IV A.

We start by generating ten million samples from each proposal density and compute their associated weights. The proportion of samples with weight equal to zero is 5.56% for the NSF proposal, compared to the 98.03% for the uniform one. To understand the distribution of such weights, Fig. 5 shows the logarithmic scale of them (for the weights > 0), assuming the average of the weights is equal to 1. For the NSF $q_\phi(\mathbf{x})$ (left), all weights are concentrated around $\log_{10} w = 0$ with a small dispersion around it. Notice that there are only three weights in ten million slightly over one hundred. This shape justifies using not the maximum value of w to perform rejection sampling, but some quantile of it, as we will discuss below. Contrary, for the uniform distribution $p_{Unif}(\mathbf{x})$ (right), we can see that the spectrum of weights goes over 9 orders of magnitude. The mean for $\log_{10} w_{q_\phi}$ is 0.023 ± 0.040 , while for $\log_{10} w_{p_{Unif}}$ we obtain an average of 0.85 ± 0.88 , indicating a huge fluctuation in the magnitude of the weights.

The results of the performance test for rejection sampling are summarized in Table I, where we compare various quantities for the NSF $q_\phi(\mathbf{x})$ and uniform $p_{Unif}(\mathbf{x})$ proposal

FIG. 4. 2D histograms comparison of cross section density for the real cross section $p(\mathbf{x})$ (left) and the proposal density $q_\phi(\mathbf{x})$ (right). Visually, an overall agreement can be seen.

functions. For this, different quantiles for the constant k for the rejection method are used, following Eq. (8), relaxing its restriction as discussed in Sec. III D. The quantiles for k were chosen using the ten million weights computed for the previous discussion of the weight magnitudes, as well as the probability of acceptance, the coverage, and the effective sample size. We considered a case of sampling one million accepted samples via rejection sampling, where samples from the proposal were generated and checked for acceptance/rejection in parallel, in batches of 300000

samples. The purpose of the parallelization is to exploit the computational capacities of a graphics processing unit. We denoted each of these batches of generating and checking a cycle of the rejection sampling. The values in Table I, for each quantile value and a proposal function, are the following:

- (i) p_{accept} : Probability of accepting a single event, given by the average of $p(\mathbf{x})/(k \cdot q_\phi(\mathbf{x}))$. If $p(\mathbf{x})/(k \cdot q_\phi(\mathbf{x})) > 1$, it is taken as 1 for the computation.
- (ii) Cycles: Number of rejection sampling cycles of size 300 000 samples needed to obtain one million accepted samples,

$$\text{Cycles} = \left\lceil \frac{10^6}{p_{\text{accept}} \cdot 3 \times 10^5} \right\rceil. \quad (11)$$

- (iii) Time: Seconds it takes to compute these cycles and obtain one million accepted samples: $t_{\text{cycle}} \cdot \text{cycles}$.
- (iv) Coverage: Volume of the original density covering when taking k with a certain quantile [Eq. (8)], following Eq. (9).
- (v) N_{ESS}/N : The ratio of effective sample size over the total number of samples, quantifying an estimate of the ratio of independence of the events. This was computed for a sample size of ten million.

The probability of acceptance, p_{accept} , for the NSF is at least 1 order of magnitude higher than the one obtained from uniform sampling. Additionally, NSF grows rapidly toward $\sim 70\%$ acceptance while also covering $>99\%$ of the original density volume, as shown in the Coverage column.

FIG. 5. Logarithmic weight distribution for ten million samples from the NSF proposal $q_\phi(\mathbf{x})$ (left) vs the same number of samples from the uniform proposal $p_{\text{unif}}(\mathbf{x})$ (right). Notice that we are only computing the logarithm for weights >0 . For the NSF, all weights are concentrated around $\log_{10} w = 0$ with a small dispersion around it, while for the uniform distribution the spectrum of weights goes over 9 orders of magnitude.

TABLE I. Performance values for different quantile choices of k for rejection sampling, as discussed in Sec. III D, comparing both NSF $q_\phi(\mathbf{x})$ and uniform (Unif.) $p_{\text{Unif}}(\mathbf{x})$ proposal (Prop.) functions. For this exercise, one million samples were generated, performing rejection sampling in batches of 300 000 tries. The quantities are the probability of accepting a single sample p_{accept} , the number of rejection cycles (batches of 300 000) used to obtain one million accepted samples, the time it took in seconds to generate these accepted samples, the coverage of the target density for that particular quantile [Eq. (9)], and the ratio of effective sample size N_{ESS} [Eq. (10)] over the total number of samples N .

Quantile	Prop.	p_{accept}	Cycles	Time (s)	Coverage	N_{ESS}/N
1.00000	NSF	0.0051	649	201.822	1.0000	0.9140
	Unif.	0.0002	16199	47.886	1.0000	0.0016
0.99999	NSF	0.0633	53	16.482	0.9999	0.9242
	Unif.	0.0004	8056	23.814	0.9947	0.0017
0.99990	NSF	0.1623	21	6.530	0.9996	0.9284
	Unif.	0.0008	4240	12.534	0.9480	0.0020
0.99900	NSF	0.3590	10	3.110	0.9984	0.9338
	Unif.	0.0027	1217	3.598	0.6185	0.0045
0.99000	NSF	0.7187	5	1.555	0.9939	0.9400
	Unif.	0.0137	244	0.721	0.0818	0.0154
0.98500	NSF	0.7730	5	1.555	0.9927	0.9405
	Unif.	0.0171	195	0.576	0.0325	0.0179
0.98100	NSF	0.7968	5	1.555	0.9920	0.9408
	Unif.	0.0193	173	0.511	0.0039	0.0195

This is not the case for the uniform distribution, which, while being only 1 order of magnitude behind NSF with regards to acceptance, is missing a large volume of coverage of the original density.

The number of rejection sampling cycles needed to achieve the desired number of accepted samples is inversely proportional to p_{accept} , as shown in Eq. (11). In a cycle, the algorithm has to sample from the proposal and evaluate both the proposal and $p(\mathbf{x})$. For the NSF, the cycles get stalled when reaching a high percentage of acceptance, since the number of cycles has to be a whole number, which is equivalent to the whole number +1. Notice how the number of cycles for the NSF is 2 orders of magnitude smaller compared to the uniform one; however, the coverage of the uniform drops drastically when decreasing the quantile, and hence the quality of the samples. We will discuss more in depth in the Appendix.

When looking at the time it takes to obtain one million accepted samples, it is directly proportional to the cycles for each proposal. The main difference is that a cycle for the uniform proposal takes a fraction of the time of a cycle for the NSF. This is because sampling and evaluating for the NSF is heavy computationally compared to doing this task for a uniform distribution. As mentioned before, see the Appendix for a more in-depth discussion.

The coverage is the main quantity of measurement of the quality of the produced samples since it measures the volume conserved of the original distribution when performing rejection sampling with certain quantiles. For all the chosen quantiles, the NSF drops a volume $<1\%$, while for the uniform distribution, the loss is of $>5\%$ for quantile 0.9999, $>38\%$ for 0.999, and $>91\%$ for 0.99, which is unacceptable when trying to produce samples from the original distribution. For NSF, this level of performance when taking the above quantiles is expected, as in Fig. 5 we have seen that the upper tail of weights with large magnitudes is a small percentage of the whole distribution. However, for uniform distribution, the loss of coverage is caused by two facts: (i) 98.03% of the weights are zero, hence placing the whole distribution on a 1.97% of the weights. (ii) These weights, as seen in Fig. 5, span over many orders of magnitude, making a cut on the quantile of their distribution more noticeable, as will be discussed below.

To visualize the coverage and the regions missing by choosing a quantile $k_{Q0.999}$ and different proposal functions, Fig. 6 shows two-dimensional histogram representation of the marginalized coverage bin-to-bin, taking the variables in pairs, where each bin quantifies the coverage of that bin (i.e., the sum of weights in that bin after choosing a certain quantile over the sum of weights of those weights without clipping). On the left, the coverage for the NSF is presented and shows that only few regions of the phase space have values smaller than 1, and even in those regions the coverage has no noticeable discrepancies. On the right, the coverage of the uniform proposal is shown for the same quantile 0.999, marking clear regions where the coverage drops drastically to values close to zero.

When comparing both coverage regions in Fig. 6, a clear pattern can be seen for the uniform one, while it looks quite random for the NSF. This is because the coverage is related to the ratio $p(\mathbf{x})/q(\mathbf{x})$, with $q(\mathbf{x})$ the corresponding proposal density. For NSF, $q_\phi(\mathbf{x})$ has a shape very closely related to $p(\mathbf{x})$, as shown in Figs. 3 and 4, so the coverage would correspond to regions where the discrepancy is large, which has a noisy behavior. Contrary, for the uniform proposal, $p_{\text{Unif}}(\mathbf{x})$, this ratio is proportional to $p(\mathbf{x})$; hence, by clipping, we are doing so according to that particular shape, making the coverage less chaotic and more structured. This translates into making highly probable areas equally likely than others with less probability, affecting this exact group of regions as we will now analyze.

Figure 6 gives us an overall picture of where the densities are wrongly estimated by choosing certain quantile, but it does not quantify or indicate the amount of error, that is, it is not telling us whether the coverage is poor in areas of small or high density. To answer this question, a multidimensional histogram over all four dimensions was performed, with 20 bins in each dimension. Then, for each bin, we compute the percentage of weight for a proposal q ,

FIG. 6. 2D histogram representation of the marginalized coverage bin-to-bin for NSF proposal (left) and for the uniform proposal (right) for k quantile = 0.999. The coverage of the NSF presents a negligible discrepancies in small areas, justifying the use of a quantile for k to improve acceptance and time, as shown in Table I. For the uniform proposal, the coverage presents an important size of the total area with significant low coverage, which is unacceptable when trying to perform rejection sampling from it.

$$\% w_q \text{ of bin} = \sum_{\mathbf{x} \in \text{bin}} w_q(\mathbf{x}) / \sum_{\mathbf{x}} w_q(\mathbf{x}),$$

which is equivalent to the percentage of density $p(\mathbf{x})$ in that bin and the coverage for the quantile $k_{Q0.999}$. Figure 7 shows a histogram of the number of bins according to their $\% w_q$ of bin vs their coverage. Notice how $\% w_q$ is presented on a logarithmic scale. For the NSF (Fig. 7 left),

FIG. 7. Representation of the number of bins according to their $\% w_q$ of bin vs their coverage for four-dimensional bins in $p(\mathbf{x})$, taking $k_{Q0.999}$. For the NSF (left), the coverage is close to one in most of the bins, and when it drops it is for low $\% w_q$. Contrary, for the uniform proposal (right), the coverage drops for high $\% w_q$, making the overall coverage way smaller than the one for NSF, as shown in Table I.

the regions of coverage visibly smaller than one are 2–4 orders of magnitude smaller in $\% w_q$ than the denser high $\% w_q$ region on the top right. This means that the areas being misrepresented by taking the quantile $k_{Q0.999}$ are relatively small. Also, most of the area with coverage < 1 are close to full coverage. Contrary, the uniform proposal (Fig. 7 right) shows the coverage dropping for high values of $\% w_q$, indicating that important regions of the original density are being trimmed down by choosing $k_{Q0.999}$. This observation is in agreement with the total coverage we are seeing in Table I.

Concerning the ratio of ESS, we can see that for NSF it is larger than 90%, giving highly uncorrelated events. For the uniform proposal, however, the ESS drops to the range of 0.16–1.95%, even for lower quantiles. The differences can be understood from Eq. (10) and by looking at the weight distribution for each proposal in Fig. 5. A large percentage of NSF weights have the same order of magnitude while for the uniform we go over a spectrum of 8 orders of magnitudes. Additionally, this ratio of ESS has to be considered for the area of the original distribution given by the coverage, where the uniform distribution poorly reproduces important regions of phase space by choosing smaller quantiles.

To summarize the analysis performed on the results of Table I, in the case of the NSF, by lowering the quantile, the probability of acceptance grows until reaching almost 80%, reducing the time to a 0.7% of the original one while maintaining coverage of over 0.99%. These scores allow us

to justify using a smaller quantile for rejection sampling to improve significantly the performance in time, while also quantifying the misrepresentation we are doing by lowering the constant k . On the contrary, for the uniform proposal, the analysis showed weakness when trying to utilize smaller quantiles, lowering the coverage by over 38% when using a quantile of only 0.999. Additionally, looking at the N_{ESS} and Fig. 5, we can see that most of the distribution is concentrated in a relatively small number of samples.

V. CONCLUSION

In this paper, we have presented ENIS, a framework to find accurate proposal density functions in the form of normalizing flows. This proposal density is subsequently used to perform rejection sampling on a target density to obtain MC data sets or compute expected values. We argue that ENIS solves the main issues associated with rejection sampling algorithms as described in the Introduction: (i) the training to find a good proposal is done automatically from the target density, with little configuration needed from the human point of view. (ii) Compared to generic proposal functions such as the uniform one, the normalizing flow adapts its density over the target one, getting rid of the inefficiencies which are usually on the downside of the method. (iii) The proposal function is generated based on a set of trivial normally distributed random numbers transformed through the flow, without any rejection method applied.

The performance of the method has been demonstrated and analyzed in a real case scenario, the CCQE cross section of the antineutrino interaction with a nucleus, where the density is four-dimensional. We have shown that, for the normalizing flow proposal, we can relax the condition in the constant k , used to construct the comparison function, boosting greatly the efficiency (up to $\approx 80\%$ of acceptance rate) while committing a very small error on the target density (less than 1%), bringing orders of magnitude of speed up in computing time compared to the same error committed by a uniform proposal. Additionally, we investigated the coverage of the generation method as a function of the constant k . We showed that the areas of the phase space where the error is committed are less relevant to the final result compared to the error in the uniform case.

The method can be used to generate fast MC samples in cases where the precision is less relevant versus the algorithm speed. High Energy Physics presents some of these examples such as extensive statistical studies based on “Asimov data sets” [30], fast detector simulations, or simply in fast studies for detector developments and designs. In those cases, the learned proposal function might be sufficiently precise and easy to generate with a set of simple normal random generators transformed through the flow.

Regarding its usage, ENIS brings the possibility of applying the same normalizing flow for rejection sampling of similar densities, e.g., densities coming from a model where the parameters are changed slightly, altering the overall density smoothly. The weight distribution of the ratio between target and proposal will be altered, but no additional training of the neural network would be needed regarding the theoretical model remains similar to the original one. This is a significant advantage compared to methods like MCMC, where one would have to rerun the complete algorithm to obtain samples from each of the different densities.

As a last remark, we have demonstrated the potential of ENIS for the four-dimensional CCQE interaction density. We believe this will only be more noticeable when applying it to higher dimensions, as the original paper of NSF [17] shows a remarkable performance on spaces of dimensions up to 63. The advantages will be also more obvious as the underlying cross section model becomes more complex and computationally involved.

ACKNOWLEDGMENTS

The authors received funding from the Spanish Ministerio de Economía y Competitividad (SEIDI-MINECO) under Grants No. FPA2016-77347-C2-2-P and No. SEV-2016-0588, from the Swiss National Foundation Grant No. 200021_85012 and acknowledge the support of the Industrial Doctorates Plan of the Secretariat of Universities and Research of the Department of Business and Knowledge of the Generalitat of Catalonia. The authors also are indebted to the Servei d’Estadística Aplicada from the Universitat Autònoma de Barcelona for their continuous support. The authors would also like to thank Durkan *et al.* for the code provided in their paper [25] which served as a basis to build ENIS, and Tobias Golling for the valuable discussion he brought to the table.

APPENDIX: MODEL COMPLEXITY DISCUSSION

We can see in Table I that although the number of rejection sampling cycles is remarkably lower for the NSF proposal, the time for obtaining one million samples is fairly similar. This is because the time for a cycle is the sum of the time in generating and evaluating the proposal plus the time it takes to evaluate $p(\mathbf{x})$, in this case the CCQE cross section. On average, for one rejection cycle of this experiment, generating + sampling for the NSF a batch of 300000 samples takes 0.31 seconds, for a uniform proposal it only takes 0.00079 seconds, and evaluating 300000 samples for the CCQE cross section model takes 0.0021 seconds. When summing up the time for the selected proposal plus the time of the model, we can see that the NSF takes a large fraction of the computational time (more than hundred times more than evaluating the cross section model), so overall a cycle for the NSF

proposal takes over a hundred cycles of uniform proposal. This is because the CCQE cross section model we have chosen for this study (Sec. II A) is relatively simple, especially compared to other applications in HEP. Additionally, we are still relatively low in the number of dimensions. As a thought experiment, we considered different cases for a similar set up: sampling ten million samples using batches of 300000 in each cycle, where the time of evaluating the model $p(\mathbf{x})$ is increased by a factor of t_{increase} .

Table II shows the results for obtaining one million samples under these circumstances for the quantiles $k_{Q1.0}$, $k_{Q0.9999}$, and $k_{Q0.9990}$ with the same acceptance probability as in Table I and considers different hypothetical models with factors $t_{\text{increase}} \in \{1, 10, 100, 1000\}$ compared to the CCQE cross section. We show the time (in seconds) it takes to generate these ten million samples for the NSF, t_{NSF} , and for the uniform proposal, t_{Unif} , while also computing their ratio $t_{\text{Unif}}/t_{\text{NSF}}$. For models of $p(\mathbf{x})$ with larger computational complexity, we can see that even for the maximum quantile $k_{Q1.0}$ NSF is faster to obtain ten million accepted samples than using a uniform proposal, gaining a whole order of magnitude in time when the model is at least hundred times more complex. Choosing quantiles smaller than 1.0, for models which take a thousand times longer to compute compared to the CCQE cross section one, the NSF proposal generates these ten million samples over a hundred times faster than the uniform distribution. Even for simpler models with only hundred more computation complexity, the gain of using the NSF proposal is noticeable.

Regarding the number of parameters used by the NSF with increasing number of dimensions, when maintaining

TABLE II. Model complexity time comparison table for both NSF and uniform proposals to generate ten million samples via rejection sampling. Choosing k_Q as one of the below quantiles, we increase the time to evaluate the model $p(\mathbf{x})$ for each rejection cycle by a factor $t_{\text{increase}} \cdot t_{\text{NSF}}$ (t_{Unif}) is the time it takes to obtain ten million accepted samples, in seconds, for the NSF (uniform) proposal. $t_{\text{Unif}}/t_{\text{NSF}}$ is the ratio of the time between both proposals, showing that for more complex models than the CCQE cross section of this paper, NSF outperforms a uniform proposal by orders of magnitude.

Quantile	t_{increase}	t_{NSF}	t_{Unif}	$t_{\text{Unif}}/t_{\text{NSF}}$
1.0000	1	2017.26	471.64	0.23
	10	2141.10	3567.01	1.67
	100	3379.54	34520.76	10.21
	1000	15763.94	344058.17	21.83
0.9999	1	63.95	123.19	1.93
	10	67.87	928.73	13.68
	100	107.01	8984.12	83.96
	1000	498.39	89538.01	179.65
0.9990	1	28.87	35.36	1.22
	10	30.64	266.54	8.70
	100	48.31	2578.40	53.37
	1000	225.00	25696.93	114.21

the same architecture for the NSF, the number scales linearly with the dimension of the data. However, when considering higher dimensions, to enrich further the flexibility of the family of densities produced by the NSF to account for more complex correlations and dependencies between each dimension, the architecture might have to change, adding an increase larger than linear for the number of parameters.

-
- [1] W. K. Hastings, *Biometrika* **57**, 97 (1970).
 - [2] S. Geman and D. Geman, *IEEE Trans. Pattern Anal. Mach. Intell.* **6**, 721 (1984).
 - [3] G. Casella, C. P. Robert, and M. T. Wells, Generalized accept-reject sampling schemes, in *A Festschrift for Herman Rubin, Lecture Notes–Monograph Series* Vol. 45, edited by A. DasGupta (Institute of Mathematical Statistics, Beachwood, OH, 2004), pp. 342–347.
 - [4] L. Martino and J. Miguez, *Signal Processing* **90**, 2981 (2010).
 - [5] C. M. Bishop and N. M. Nasrabadi, *J. Electron. Imaging* **16**, 049901 (2007).
 - [6] J. von Neumann, in *Monte Carlo Method*, edited by A. S. Householder, G. E. Forsythe, and H. H. Germond, National Bureau of Standards Applied Mathematics Series Vol. 12 (US Government Printing Office, Washington, DC, 1951), Chap. 13, pp. 36–38.
 - [7] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods* (Springer, New York, 2004).
 - [8] T. Müller, B. McWilliams, F. Rousselle, M. Gross, and J. Novák, *ACM Trans. Graph.* **38**, 145 (2018).
 - [9] H. Kahn and A. W. Marshall, *J. Oper. Res. Soc. Am.* **1**, 263 (1953).
 - [10] E. Bothmann, T. Janen, M. Knobbe, T. Schmale, and S. Schumann, *SciPost Phys.* **8**, 069 (2020).
 - [11] C. Gao, S. Höche, J. Isaacson, C. Krause, and H. Schulz, *Phys. Rev. D* **101**, 076002 (2020).
 - [12] L. Alvarez-Ruso *et al.* (NuSTEC Collaboration), *Prog. Part. Nucl. Phys.* **100**, 1 (2018).
 - [13] K. Abe *et al.* (T2K Collaboration), *Nature (London)* **580**, 339 (2020).
 - [14] M. Acero *et al.* (NOvA Collaboration), *Phys. Rev. Lett.* **123**, 151803 (2019).
 - [15] R. Smith and E. Moniz, *Nucl. Phys.* **B43**, 605 (1972); **B101**, 547(E) (1975).

- [16] G. Papamakarios, E. T. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, [arXiv:1912.02762](#).
- [17] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., 2019), pp. 7511–7522, <http://papers.nips.cc/paper/8969-neural-spline-flows.pdf>.
- [18] D. J. Rezende and S. Mohamed, [arXiv:1505.05770](#).
- [19] M. Germain, K. Gregor, I. Murray, and H. Larochelle, in *Proceedings of the 32nd International Conference on Machine Learning*, edited by F. Bach and D. Blei, Proceedings of Machine Learning Research Vol. 37 (PMLR, Lille, France, 2015), pp. 881–889, <http://proceedings.mlr.press/v37/germain15.html>.
- [20] G. Papamakarios, I. Murray, and T. Pavlakou, Masked Autoregressive Flow for Density Estimation, in *Advances in Neural Information Processing Systems 30*, edited by I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Curran Associates, Inc., 2017), pp. 2338–2347, <http://papers.nips.cc/paper/6828-masked-autoregressive-flow-for-density-estimation.pdf>.
- [21] C.-W. Huang, D. Krueger, A. Lacoste, and A. C. Courville, [arXiv:1804.00779](#).
- [22] N. D. Cao, I. Titov, and W. Aziz, [arXiv:1904.04676](#).
- [23] A. Wehenkel and G. Louppe, in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., 2019), pp. 1545–1555, <http://papers.nips.cc/paper/8433-unconstrained-monotonic-neural-networks.pdf>.
- [24] P. Jaini, K. A. Selby, and Y. Yu, in *Proceedings of the 36th International Conference on Machine Learning*, edited by K. Chaudhuri and R. Salakhutdinov, Proceedings of Machine Learning Research Vol. 97 (PMLR, California, USA, 2019), pp. 3009–3018, <http://proceedings.mlr.press/v97/jaini19a.html>.
- [25] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, [arXiv:1906.02145](#).
- [26] J. A. Gregory and R. Delbourgo, *IMA J. Numer. Anal.* **2**, 123 (1982).
- [27] S. Kullback and R. A. Leibler, *Ann. Math. Stat.* **22**, 79 (1951).
- [28] D. P. Kingma and J. Ba, [arXiv:1412.6980](#).
- [29] J. S. Liu, *Stat. Comput.* **6**, 113 (1996).
- [30] The concept of the Asimov data is a representative date set, inspired by the short story Franchise by Isaac Asimov, in which the entire electorate is replaced by selecting the most representative voter.