Likelihood-free inference of experimental neutrino oscillations using neural spline flows

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

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

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

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

(Received 25 February 2020; accepted 19 May 2020; published 2 June 2020)

In machine learning, likelihood-free inference refers to the task of performing such analysis driven by data instead of an analytical expression. We discuss the application of neural spline flows, a neural density estimation algorithm, to the likelihood-free inference problem of the measurement of neutrino oscillation parameters in long baseline neutrino experiments. A method adapted to physics parameter inference is developed and applied to the case of the disappearance muon neutrino analysis at the T2K experiment.

DOI: [10.1103/PhysRevD.101.113001](https://doi.org/10.1103/PhysRevD.101.113001)

I. INTRODUCTION

First indications of neutrino oscillations using atmospheric neutrinos were presented by the Super-Kamiokande experiment in Japan in 1998 [1]. The confirmation with solar neutrinos was performed in 2002 by the SNO experiment [2] and the check of the atmospheric oscillation with a human-made beam experiments at K2K [3] for atmospheric oscillations in 2006 and at KamLAND [4] for solar neutrino oscillations in 2004. After eight very productive years, the oscillation phenomena were experimentally well established. Neutrino oscillation phenomena are the first, and so far the only, evidence of neutrinos having mass. Following the initial experiments, a successful set of measurements confirmed the oscillation picture and improved the understanding of the Pontecorvo-Maki-Nakagawa-Sakata neutrino flavour mixing matrix parameters [5]. By 2010, the measurement of a nonzero θ_{13} angle by T2K [6], followed by the precision measurement from Daya Bay [7], opened the potential existence of leptonic charge-parity (CP) symmetry violation (i.e., particles behaving differently from the antiparticles). As of today, all mixing angles have been measured [8], including the two mass differences between the three mass neutrino eigenstates. The remaining mixing matrix parameter to be measured are the imaginary phase responsible of the CP violation and the sign of one of the two neutrino mass splitting, which determines the so-called hierarchy.

Both measurements are at reach for current and near future oscillation experiments.

Statistical methods used in the oscillation analysis so far comprise both frequentist and Bayesian approaches [9,10]. There are, however, some limitations to these methods. Both approaches are based as of today on a binned likelihood algorithm which might limit the sensitivity of the experiment and some of them impose a Gaussian dependency in some of the nuisance parameters affecting the precision of the results and correctness of the evaluated uncertainties. Additionally, they require very intensive CPU processing time, which is a limiting factor that reduces the flexibility of the statistical analysis and checks, and introduces strong constraints on the delivery of the results. These limitations are derived from the intrinsic difficulties of the statistical data analysis that are depicted in Sec. II, using the T2K experiment as a reference example.

In this paper, we propose an alternative statistical method to overcome some of the limitations of the current methods in use. The proposed procedure is based on an unbinned likelihood inference using neural density estimators. This method has the potential of being accurate, fast, and to reduce the possible bias due to the intrinsic binning in other approaches. The neural spline flows (NSF), the implementation of neural density estimators we chose, have also some advantages since the Gaussian generator intrinsic to the method will facilitate the introduction of experimental errors in the distributions. We will discuss in this paper the basic concepts of the method and show the potential with a simplified example.

II. PROBLEM DEFINITION AND PHYSICAL SIMULATOR

Neutrino oscillation experiments search for the modification of the flavor content of a neutrino beam traveling 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.

vacuum or matter for a certain distance. Beams are normally characterized at a near site, where the neutrino energy spectrum and flavor composition are not yet altered by oscillations. The same beam is sampled after a certain flight distance L . The change on the flavor composition can be determined in two different ways which are as follows:

- (i) The neutrino flavor disappearance ($P(\nu_\alpha \rightarrow \nu_\alpha)$) experiments search for the disappearance of a certain neutrino flavor as function of the neutrino energy. The disappearance produces both a reduction in the flux of neutrinos of a given flavor and the distortion of the neutrino energy spectra that is observed in the distribution of the measured quantities. For example, T2K uses the muon momentum (p_μ) and the angle with respect to the neutrino direction (θ_μ).
- (ii) The neutrino flavor appearance ($P(\nu_\alpha \rightarrow \nu_\beta)$, $\alpha \neq \beta$) experiments search for the appearance of a neutrino flavor that is normally suppressed in the original neutrino flux. In T2K, this new flavor is the electron neutrino. The dependency of the oscillation with the neutrino energy is inferred from the momentum and the angle with respect to the neutrino direction of the electron ejected in the interaction of neutrinos with matter.

The neutrino flavor is determined by the flavor of the charged lepton (muon, electron or tauon) produced in charged current interactions of neutrinos with the nuclei. For the current analysis, we will concentrate on the disappearance phenomenon (i).

In a synthetic way, the experimental number of observed neutrinos with observed properties ($\vec{\theta}_\nu^{\text{reco}}$) can be described by

$$N_{\text{evts}}^{\text{near}}(\vec{\theta}_\nu^{\text{reco}}) = \int \sigma(E_\nu) \phi^{\text{near}}(E_\nu) P_{\text{near}}(\vec{\theta}_\nu^{\text{reco}}|E_\nu) dE_\nu + \text{Back}_{\text{near}}(\vec{\theta}_\nu^{\text{reco}})$$

for the near detector and

$$N_{\text{evts}}^{\text{far}}(\vec{\theta}_\nu^{\text{reco}}) = \int \sigma(E_\nu) \phi^{\text{far}}(E_\nu) P_{\text{far}}(\vec{\theta}_\nu^{\text{reco}}|E_\nu) P_{\text{osc}}(E_\nu) dE_\nu + \text{Back}_{\text{far}}(\vec{\theta}_\nu^{\text{reco}})$$

for the far detector. The number of observed neutrinos depends on the cross section ($\sigma(E_\nu)$), the neutrino flux ($\phi^{\text{far,near}}(E_\nu)$), the probability of observing the experimentally accessible quantities ($\vec{\theta}_\nu^{\text{reco}}$) given a neutrino energy ($P_{\text{far,near}}(\vec{\theta}_\nu^{\text{reco}}|E_\nu)$), the oscillation probability ($P_{\text{osc}}(E_\nu)$), and the backgrounds observed in the detectors ($\text{Back}_{\text{near,far}}(\vec{\theta}_\nu^{\text{reco}})$).

The experimental challenge comes from inferring the neutrino energy, E_ν , given the experimental observable ($\vec{\theta}_\nu^{\text{reco}}$). The term $P_{\text{far,near}}(\vec{\theta}_\nu^{\text{reco}}|E_\nu)$ encapsulates not only the

detector resolution, but also the neutrino-nucleus cross section model predictions and uncertainties. Other difficulties raise from the limited knowledge ($\approx 9\%$ in the latest T2K results[11]) of the neutrino flux ($\phi^{\text{far,near}}(E_\nu)$) and of neutrino cross sections as function of the energy ($\sigma(E_\nu)$). The background terms ($\text{Back}_{\text{far,near}}(\vec{\theta}_\nu^{\text{reco}})$) are normally relevant ($\approx 20\%$ in T2K) [12], and they subsequently depend on the neutrino-nucleus cross sections in a nontrivial manner.

Both the frequentist and the Bayesian statistical approaches [9,10] utilize the near detector data to predict the probability density function at the far detector in the absence of oscillations,

$$f(\vec{\theta}_\nu^{\text{reco}}|E_\nu) = \sigma(E_\nu) \phi^{\text{far}}(E_\nu) P_{\text{far}}(\vec{\theta}_\nu^{\text{reco}}|E_\nu).$$

Once determined, the conditional probability density function $f(\vec{\theta}_\nu^{\text{reco}}|E_\nu)$ can be used to determine the oscillation parameters ($P_{\text{osc}}(E_\nu)$) by comparing it to the far detector events. Most of the experimental effort is actually devoted to the determination of this conditional probability which depends also on a large number of hidden and correlated parameters describing uncertainties in detector performances, cross section models, and the neutrino flux. Hidden parameters are marginalized or profiled in the analysis, providing the experimental result for oscillation parameters as (the posterior in the case of Bayesian approaches) probability maps.

In the particular case of the T2K experiment, the experimentally accessible observables ($\vec{\theta}_\nu^{\text{reco}}$) are the momentum and direction of the μ lepton ($p_\mu^{\text{reco}}, \theta_\mu^{\text{reco}}$). Near and far detectors are able to provide also the kinematic of pions (charged and neutral) and protons, or the total released energy in the interaction, but we will ignore these capabilities to simplify the discussion. The μ lepton is produced at the interaction of the neutrino with the target nucleus. In this case, the probability density function ($f(\vec{\theta}_\nu^{\text{reco}}|E_\nu)$) can be simplified to $p(p_\mu^{\text{reco}}, \theta_\mu^{\text{reco}}|E_\nu)$. The near detector of the experiment measures the neutrino flux and tunes the model of neutrino-nucleus interaction providing the estimation of the probability density function $p(p_\mu^{\text{reco}}, \theta_\mu^{\text{reco}}, E_\nu) = p(p_\mu^{\text{reco}}, \theta_\mu^{\text{reco}}|E_\nu) p(E_\nu)$ together with the expected number of interactions in the far detector in the absence of oscillations. The near detector provides also a dependency with free parameters in the model and a full error covariance matrix relating all of them. To simplify the exercise, we ignore the error covariance matrix in this study and assume that the probability $p(p_\mu^{\text{reco}}, \theta_\mu^{\text{reco}}, E_\nu)$ is implicitly known to the experiment through simulations.

The neutrino oscillation disappearance probability can be approached by the simplified two-flavor [8] oscillation. The disappearance probability as a function of the initial energy of the neutrino E_ν is

$$p_{\text{osc}}(E_\nu, \theta_{\text{mix}}, \Delta m^2) = \sin^2(2\theta_{\text{mix}}) \sin^2\left(1.27 \frac{\Delta m^2 295}{E_\nu}\right), \quad (1)$$

where 295 is the distance in kilometers between the near and far sites in the T2K experiment and 1.27 is a scaling parameter to adjust the oscillation phase to distance in kilometers. E_ν is the neutrino energy in GeV, θ_{mix} the mixing angle of the two flavours, and Δm^2 the difference in mass of the two mass eigenstates in eV^2 . θ_{mix} and Δm^2 are the parameters governing the oscillations.

Although the problem was simplified in the following analysis to make a proof of concept, in Sec. V we outline how the methodology could be applied to the full complex analysis.

A. Physical simulator

We have simplified the problem to demonstrate the viability of the proposed method to determine the oscillation parameters using neural spline flows. Event samples are generated using the NEUT [13] Monte Carlo (MC) event generator model that describes the interactions of neutrinos with nuclei. We also use a realistic neutrino flux energy spectrum provided by the T2K Collaboration [14]. With both inputs, we generate charge current quasielastic events (CCQE). CCQE is the most probable reaction at T2K energies, and the one dominating the statistical sensitivity of the experiment, where the neutrino transforms into a muon exchanging a neutron into a proton ($\nu + n \rightarrow \mu + p$). To simplify, we ignore other reaction channels and potential backgrounds, and also detector effects. The generator provides n-tuples of events weighted according to their probability as function of neutrino energy and angle and momentum of the muon.

III. METHODOLOGY

Evaluating the density of high-dimensional data has become an important task for unsupervised machine learning in recent years. Neural density estimation proposes a solution using neural networks by learning an estimation $q_\phi(\mathbf{x})$ of the exact *target density* $p(\mathbf{x})$ from samples $\mathbf{x} \sim p(\mathbf{x})$. In this work, this task is performed through neural spline flows, a specific implementation of the more general normalizing flows methods, presented in Sec. III A. We explain how the density estimation at the near detector is combined with the analytical formula for neutrino oscillation for the far detector to obtain the likelihood used to perform Bayesian inference for the T2K experiment in Sec. III B. Additionally, an alternative way of computing the posterior is explained in Sec. III C, based on the standard histogram approach but tweaked in order to attain the limit of unbinned likelihood, closely related to what is obtained by neural spline flows.

A. Neural density estimation using neural spline flows

Consider a simulator which can produce samples $\mathbf{x} \sim p(\mathbf{x})$ over the real D-dimensional space $\mathbb{R}^D$ from the target density. Consider also $q_\phi(\mathbf{x}) = q(\mathbf{x}; \phi)$ a flexible family of density functions parametrized by ϕ , i.e., $q_\phi(\mathbf{x}) \geq 0$ for all $\mathbf{x}, \phi$, and $\int q_\phi(\mathbf{x}) d\mathbf{x} = 1$ for all ϕ . Then, if we want to approximate $p(\mathbf{x})$ through $q_\phi(\mathbf{x})$, we optimize the parameters ϕ by maximizing the log-likelihood of the approximated density under simulated data from the target distribution,

$$L = \frac{1}{N} \sum_{i=1}^N \log(q_\phi(\mathbf{x}_i)) \quad \text{with} \quad \mathbf{x}_i \sim p(\mathbf{x}). \quad (2)$$

This objective function for the optimization procedure is equivalent to minimizing the Kullback-Leibler divergence between the target and approximated density,

$$D_{\text{KL}}(p(\mathbf{x}) \| q_\phi(\mathbf{x})) = \int p(\mathbf{x}) \log\left(\frac{p(\mathbf{x})}{q_\phi(\mathbf{x})}\right) d\mathbf{x}, \quad (3)$$

which is always non-negative and only equal to 0 if both densities are equivalent. Let us proof quickly this statement,

$$\begin{aligned} & \arg \min_{\phi} D_{\text{KL}}(p(\mathbf{x}) \| q_\phi(\mathbf{x})) \\ &= \arg \min_{\phi} \int p(\mathbf{x}) \log\left(\frac{p(\mathbf{x})}{q_\phi(\mathbf{x})}\right) d\mathbf{x} \\ &= \arg \min_{\phi} - \int p(\mathbf{x}) \log q_\phi(\mathbf{x}) d\mathbf{x} \\ &= \arg \max_{\phi} \int p(\mathbf{x}) \log q_\phi(\mathbf{x}) d\mathbf{x} \\ &\approx \arg \max_{\phi} \sum \log q_\phi(\mathbf{x}) \quad \text{with} \quad \mathbf{x} \sim p(\mathbf{x}), \end{aligned}$$

where in the last step we have approximated the expected value of $\log q_\phi(\mathbf{x})$ with respect to $p(\mathbf{x})$ by the finite expected value. Hence, maximizing the objective function, Eq. (2) is equivalent to minimizing the Kullback-Leibler divergence Eq. (3) between the distributions, i.e., approximating p by q_ϕ .

Normalizing flows are a mechanism of constructing such flexible probability density families $q_\phi(\mathbf{x})$ for continuous random variables. A comprehensive review on the topic can be found in [15], 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* [16].

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

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

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, 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 ,

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

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

$$q_{\phi}(\mathbf{x}) = p_{\mathbf{u}}(T_{\phi}(\mathbf{x})) |\det J_{T_{\phi}}(\mathbf{x})|. \quad (4)$$

The goal is to find the parameters ϕ to maximize Eq. (2) if $\mathbf{x} \sim p(\mathbf{x})$. 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 one to construct a complex transformation via composition of simpler transformations. Consider the transformation T as a composition of simpler T_k transformations,

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

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

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

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

These two computations (plus their inverse) are the building blocks of a normalizing flow [17]. Hence, to make a transformation efficient, both operations have to be efficient.

From now on forth, 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 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 *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 computational 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 [18] enable one 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 [19], monotonic neural networks [20–22], sum-of-squares polynomials [23], or monotonic splines [16,24,25]. In this work, we will focus on a specific implementation of monotonic splines, the neural spline flows.

In their work on neural spline flows [16], Durkan *et al.* advocate for utilizing monotonic rational-quadratic splines as transformers τ , which are easily differentiable, more flexible than previous attempts using polynomials for these transformers, since their Taylor-series expansion is infinite, and are analytically invertible.

The monotonic rational-quadratic transformers 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 $3K - 1$ 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 knots.

Stacking up many of these transformations, a highly flexible neural density estimator can be built and will be the one utilized during this work.

B. Neural spline flows applied to the T2K oscillation problem

In this subsection, we will explain how the likelihood for Bayesian inference is constructed. In the context of machine learning, likelihood-free inference, as stated in the abstract, refers to the task of performing such analysis when the densities are data driven, but no explicit likelihood function can be constructed. This is the case for the near detector, where we have data for the different magnitudes but no analytical density available. Therefore, a density estimation for the near detector is performed through NSF. This near detector density is then combined with the analytical formula for neutrino oscillation for the far detector in order to obtain the likelihood for the experiment. By doing so, we are combining the potential of NSF with expertise of the particular problem.

We start by estimating the explicit density of the expected energy spectrum E_ν of the neutrinos, together with the momentum p_μ and angle θ_μ of the measured muon without oscillations, obtaining $p(p_\mu, \theta_\mu, E_\nu)$, as measured by the near detector. This is done by learning the density using an NSF from the Monte Carlo data, generated as presented in Sec. II A. The base density $p_u(\mathbf{u})$ used for Eq. (4) is a three-dimensional standard normal distribution.

Having estimated the joint probability $p(p_\mu, \theta_\mu, E_\nu)$ of the initial distribution at the near detector, we need to construct the conditional density $p(p_\mu, \theta_\mu | \theta_{\text{mix}}, \Delta m^2)$ of the observed magnitudes given the oscillation parameters at the far detector in order to perform Bayesian inference. For this, we simply integrate the probability of not oscillating, $1 - p_{\text{osc}}$, using Eq. (1), over the energy spectrum of the joint distribution,

$$p(p_\mu, \theta_\mu | \theta_{\text{mix}}, \Delta m^2) = C(\theta_{\text{mix}}, \Delta m^2) \cdot \int p(p_\mu, \theta_\mu, E_\nu) (1 - p_{\text{osc}}(E_\nu, \theta_{\text{mix}}, \Delta m^2)) dE_\nu,$$

where $C(\theta_{\text{mix}}, \Delta m^2)$ is a constant of normalization computed after performing the integral. With this we have the probability of observing a single muon with momentum p_μ

and angle θ_μ after oscillating given the parameters θ_{mix} and Δm^2 .

In order to take into account the number of observed samples, the extended likelihood [27,28] is used, modifying the likelihood with a Poisson count term to consider the expected number of events for a given set of parameters and the actual observed number

$$L(\boldsymbol{\theta}) = \frac{[\mu(\boldsymbol{\theta})]^n}{n!} e^{-\mu(\boldsymbol{\theta})} \prod_{i=1}^n p(\mathbf{x}_i | \boldsymbol{\theta}).$$

In our case, the Poisson parameter $\mu(\boldsymbol{\theta})$ is obtained by integrating the possible oscillations over all the energy spectrum, scaled to the initial number of particles N_{ini} ,

$$\mu(\theta_{\text{mix}}, \Delta m^2) = N_{\text{ini}} \iiint p(p_\mu, \theta_\mu, E_\nu) \times (1 - p_{\text{osc}}(E_\nu, \theta_{\text{mix}}, \Delta m^2)) dE_\nu dp_\mu d\theta_\mu.$$

Hence, the extended likelihood we apply for the analysis is

$$L(\theta_{\text{mix}}, \Delta m^2) = \frac{[\mu(\theta_{\text{mix}}, \Delta m^2)]^n}{n!} e^{-\mu(\theta_{\text{mix}}, \Delta m^2)} \times \prod_{i=1}^n p(p_\mu^{(i)}, \theta_\mu^{(i)} | \theta_{\text{mix}}, \Delta m^2), \quad (5)$$

and the posterior for the parameters takes the form of

$$p(\theta_{\text{mix}}, \Delta m^2 | \{p_\mu^{(i)}, \theta_\mu^{(i)}\}_{i=1}^n) \propto L(\theta_{\text{mix}}, \Delta m^2) p(\theta_{\text{mix}}, \Delta m^2), \quad (6)$$

with $p(\theta_{\text{mix}}, \Delta m^2)$ the prior information before observing the events and $\{p_\mu^{(i)}, \theta_\mu^{(i)}\}_{i=1}^n$ the set of observed events.

C. Reference analysis using an approximate unbinned likelihood

The results of the experiments are validated using an approximate unbinned likelihood. The likelihood is computed following Eq. (5). To do so, event histograms ($M(p_\mu^i, \theta_\mu^j | E^k)$) binned in muon momentum and angle for a given a neutrino energy are generated from the simulated data described in Sec. II A. The oscillated probability is computed by reweighting the histogram content, using Eq. (1), as

$$M^{\text{osc}}(p_\mu^i, \theta_\mu^j | \theta_{\text{mix}}, \Delta m^2) = \sum_k M(p_\mu^i, \theta_\mu^j | E^k) p_{\text{osc}}(E^k, \theta_{\text{mix}}, \Delta m^2)$$

and is interpolated linearly to reduce the effects due to the coarse binning,

$$M^{\text{osc}}(p_\mu^i, \theta_\mu^j | \theta_{\text{mix}}, \Delta m^2) \rightarrow M^{\text{osc}}(p_\mu, \theta_\mu | \theta_{\text{mix}}, \Delta m^2).$$

The number of expected events ($\mu(\theta, \Delta m^2)$) is computed by summing the binned probability

$$\mu(\theta_{\text{mix}}, \Delta m^2) = \sum_{i,j} M^{\text{osc}}(p_\mu^i, \theta_\mu^j | \theta_{\text{mix}}, \Delta m^2).$$

The probability obtained by normalizing the oscillated maps takes the form

$$p(p_\mu, \theta_\mu | \theta, \Delta m^2) = \frac{M^{\text{osc}}(p_\mu, \theta_\mu | \theta_{\text{mix}}, \Delta m^2)}{\mu(\theta_{\text{mix}}, \Delta m^2)}.$$

The likelihood probability $L(\theta_{\text{mix}}, \Delta m^2)$ is finally computed for discrete values of θ_{mix} and Δm^2 . Those values are distributed in a grid identical to the one used for the NSF approach for proper comparison between both methods. We have tested the results with different number of initial bins in energy, momentum, and angle to find a good compromise between stability of the result and speed. We call this cross-check method in what follows the *hist* method and it will be used as a reference for the NSF calculations.

IV. EXPERIMENTS

The methodology is tested on experiments of simulated neutrino oscillations according to the T2K experiment. For this, ten different sets of observed events are constructed (see the Appendix), with additional MC data used to fit both the NSF and construct the unbinned likelihood. The training of the NSF to fit the density of $(p_\mu, \theta_\mu, E_\nu)$ is shown and verified in Sec. IV A. Afterward, in Sec. IV B, the inference on the ten experiments are performed utilizing both NSF and the unbinned likelihood, discussing the findings of both methodologies and the possible bias introduced by them.

All events were generated as defined in Sec. II A, using the neutrino flux energy spectrum from the T2K Collaboration [14] to produce the energy of the incoming neutrino E_ν , and the NEUT event generator [13] to compute the momentum p_μ and angle θ_μ of the resulting muon, forming a triplet of $(p_\mu, \theta_\mu, E_\nu)$, describing an implicit density of these three magnitudes.

A. Training and validation of the NSF

In order to apply the methodology described in Sec. III B, we start by estimating the density $p(p_\mu, \theta_\mu, E_\nu)$ using a neural spline flow on ≈ 15 M MC events as the training set to fit the parameters of the flow. As is a standard procedure in machine learning, an additional set of ≈ 4 M events is used for validation of the model outside of the training set, i.e., to check the integrity of the models for data not seen to fit the parameters.

Because of the similarities regarding dimension of the data and number of samples, but with a simpler structure,

FIG. 1. Neural spline flow training log probability for estimating $p(p_\mu, \theta_\mu, E_\nu)$ for training (solid) and validation (dashed) sets, as shown by Eq. (2). For the validation, the log probability stabilizes during the training to converge to a certain value which depends on the architecture of the network.

we have chosen the hyperparameters almost as the POWER data set in Appendix B from [16], i.e., ADAM optimizer [29] with learning rate 0.0005, batch size 512, training steps 400k, flow steps 5, transform blocks 5, hidden features 128, and bins 8. Additionally, the learning rate was decreased during the training using a cosine scheduler to ensure stabilization at the end of the training procedure. As shown in Fig. 1, the validation set stabilizes at the end of the training and appears to converge for the selected architecture of the network.

To ensure that the transformation T of Eq. (4) was found properly by the NSF, consider the transformed data $\mathbf{u}$ for a new set of ≈ 1 M MC events, never seen before by the algorithm. If T is correctly approximated, $\mathbf{u}$ should follow a three-dimensional standard normal distribution, as is shown qualitatively in Fig. 2.

For this, an χ^2 test was performed over a binning of $50 \times 50 \times 50$ in the domain $[-3, 3]^3$ of the transformed data $\mathbf{u}$ to test the goodness-of-fit to a three-dimensional standard normal distribution, obtaining a p value of 0.3062. With this, we can assume that the transformed data $\mathbf{u}$ correspond to samples from such base distribution, justifying that the transformation T was properly found by the NSF, hence allowing to evaluate $p(p_\mu, \theta_\mu, E_\nu)$ accurately through it.

B. Inference results

To test the performance of obtaining the posterior according to Sec. III B, ten different observation sets were constructed, as explained in the Appendix, with five different mixing angles θ_{mix} and difference in mass squared Δm^2 .

For each of the five combinations of parameters, low and high statistics (number of observed events) experiments were performed. Low statistics are of the order of the real observed samples at T2K and used to assure its performance when a small number of events is dealt with. High statistics (2 orders of magnitude larger number of observed

FIG. 2. Two-dimensional histograms of samples from the initial distribution of $(p_\mu, \theta_\mu, E_\nu)$ (top) and the transformed data $\mathbf{u} = (u_1, u_2, u_3)$ (bottom) under T according to Eq. (4). If transformation T is approximated properly, $\mathbf{u}$ should follow a three-dimensional standard normal distribution, as is depicted qualitatively in this figure.

events compared to the usual expected number in the low statistics case) allow us to check the agreement between traditional binning methodology with a large number of very fine bins and the unbinned NSF posterior.

Table I shows ten experiments, with the number of observed samples, the true parameters, and the results using the NSF (Sec. III B). Additionally, a result using an approximate unbinned likelihood, denoted by hist (Sec. III C), is also displayed, which would correspond to the limit case when histograms can be performed with a large number of bins to behave like an unbinned estimation. Since Bayesian inference is used, the inference on the parameters describes a density function according to Eq. (6). The central value shown for each parameter is the one that maximizes the joint posterior density. The uncertainty is then computed by marginalizing in the two-dimensional density one of the parameters to obtain the one-dimensional one of the other and finding the interval such that for a $1 - \alpha$ confidence level (CL), $\alpha/2$ of the density is found on each side. This is done for both NSF posterior and hist posterior.

In general, the results of both methodologies agree, with slight fluctuations in the confidence levels. The difference could come from the NSF not learning perfectly the density of the points, from the interpolation done by the hist method introducing wrong approximations or from intrinsic biases, which will be discussed at the end of this subsection.

Additionally, in order to visualize the agreement in two dimensions in Fig. 3, the highest posterior density (HPD) curves [30] of 68% and 95%, together with the best fit (highest posterior value) were computed for experiments 1 (top left), 2 (top right), 7 (bottom left), and 8 (bottom right). In both HPD regions, a clear overlap for low statistics and slight fluctuations on larger statistics can be observed. When comparing the difference in area size, the relative difference of the hist method with respect to NSF through the ten experiments is $3.1 \pm 1.9\%$, showing that the areas agree within $2 - \sigma$ on average.

TABLE I. Posterior inference of ten different experiments, alternating between low and high number of observed events. For the inferred parameters using NSF and the unbinned histogram approximation, hist, the 95% confidence level was computed using the one-dimensional marginalized densities of each parameter. θ_{mix} is given in rad and Δm^2 in $\times 10^{-3} \text{ eV}^2$.

Experiment #	N_{obs}	Parameter	True values	NSF 95% C.L.	Hist 95% C.L.
1	506	θ_{mix}	0.7594	$0.785^{+0.055}_{-0.056}$	$0.785^{+0.055}_{-0.056}$
		Δm^2	2.463	$2.440^{+0.091}_{-0.085}$	$2.446^{+0.092}_{-0.087}$
2	49672	θ_{mix}	0.7594	$0.751^{+0.074}_{-0.006}$	$0.754^{+0.069}_{-0.007}$
		Δm^2	2.463	$2.464^{+0.008}_{-0.012}$	$2.467^{+0.007}_{-0.012}$
3	532	θ_{mix}	0.7353	$0.71^{+0.18}_{-0.03}$	$0.71^{+0.18}_{-0.03}$
		Δm^2	2.463	$2.46^{+0.10}_{-0.11}$	$2.46^{+0.10}_{-0.11}$
4	49646	θ_{mix}	0.7353	$0.738^{+0.099}_{-0.006}$	$0.739^{+0.097}_{-0.006}$
		Δm^2	2.463	$2.458^{+0.009}_{-0.011}$	$2.461^{+0.009}_{-0.011}$
5	493	θ_{mix}	0.6847	$0.65^{+0.28}_{-0.02}$	$0.65^{+0.28}_{-0.02}$
		Δm^2	2.463	$2.554^{+0.120}_{-0.127}$	$2.563^{+0.122}_{-0.128}$
6	49665	θ_{mix}	0.6847	$0.682^{+0.210}_{-0.003}$	$0.683^{+0.208}_{-0.004}$
		Δm^2	2.463	$2.479^{+0.008}_{-0.014}$	$2.482^{+0.009}_{-0.014}$
7	506	θ_{mix}	0.7353	$0.73^{+0.14}_{-0.04}$	$0.73^{+0.14}_{-0.04}$
		Δm^2	2.363	$2.347^{+0.097}_{-0.095}$	$2.353^{+0.096}_{-0.099}$
8	50145	θ_{mix}	0.7353	$0.734^{+0.108}_{-0.006}$	$0.735^{+0.106}_{-0.006}$
		Δm^2	2.363	$2.365^{+0.010}_{-0.011}$	$2.368^{+0.009}_{-0.012}$
9	481	θ_{mix}	0.7353	$0.785^{+0.060}_{-0.061}$	$0.785^{+0.060}_{-0.061}$
		Δm^2	2.663	$2.620^{+0.086}_{-0.087}$	$2.626^{+0.087}_{-0.089}$
10	49710	θ_{mix}	0.7353	$0.741^{+0.094}_{-0.005}$	$0.743^{+0.089}_{-0.005}$
		Δm^2	2.663	$2.659^{+0.007}_{-0.011}$	$2.662^{+0.008}_{-0.010}$

In Fig. 3, one observes that the areas, even being similar in size, are slightly shifted one from another. Empirical experiments of computing the posterior using actual discrete binning show that, by using a denser binning, its posterior was shifting toward the NSF result. To measure the bias $(\hat{\theta} - \theta)$ of the estimated parameters $(\hat{\theta}_{\text{mix}}, \hat{\Delta m^2})$ by both methods, an observation of $N_{\text{obs}} \approx 500 \text{ k}$ events was used for the two sets of parameters used in Fig. 3. Results are summarized in Table II where $(\hat{\theta}_{\text{mix}}, \hat{\Delta m^2})$ are taken as the maximum value of the posterior probability obtained in each set of parameters. Table II shows that NSF have a significant smaller bias compared to the hist method. This explains the discrepancy in the plots, aside from justifying a better performance by the NSF method. The hist bias comes from binning and interpolation approximations, while for NSF the bias may come from the density

FIG. 3. HPD curves for both NSF and hist posteriors of experiments 1 (top left), 2 (top right), 7 (bottom left), and 8 (bottom right), together with best fit of each posterior. Red (blue) lines indicate 68% (dashed) and 95% (continuous) HPD curves for the NSF (hist) method. Orange x-crosses indicate the true parameter used to generate the observed events. Red stars (blue +crosses) indicate the best fit for the NSF (hist) method. A clear overlap can be found in experiments of low statistics (left) and a slight fluctuation on large statistics (right). Notice a change of scale in the high statistics plots.

estimation for the near detector not being perfect. In real experiments, this bias can be estimated using a representative MC data set, called the *Asimov data set*.

Both quantitative, Table I, and qualitative, Fig. 3, show that NSF indeed provide a tool to perform likelihood-free inference on physical simulators such as the one of the T2K experiment, in agreement with unbinned likelihood approaches as we have compared it to, but with less bias as shown in Table II.

TABLE II. Bias computation for the posterior densities of the parameters in Fig. 3. The estimators ($\hat{\theta}_{\text{mix}}$, $\hat{\Delta m^2}$) are taken as the maximum value of the posterior obtained in each set of parameters for a high-statistical experiment ($N_{\text{obs}} \approx 500$ k), and the bias is defined as $\hat{\theta} - \theta$. NSF show a significant reduction of bias compared to the hist method. θ_{mix} is given in rad and Δm^2 in $\times 10^{-3}$ eV².

Parameter	True values	NSF bias	Hist bias
θ_{mix}	0.7594	0.0020	0.0087
Δm^2	2.463	-0.0012	0.0031
θ_{mix}	0.7353	-0.00050	0.00410
Δm^2	2.363	-0.00063	0.00675

V. MODELING SYSTEMATIC UNCERTAINTIES

A comprehensive description of the problem is beyond the scope of this paper, but we will sketch possible implementation alternatives. In experiments, near detector neutrino interactions data are used to constrain uncertainties in cross section models ($\hat{\phi}_{\text{xsect}}$), neutrino flux ($\hat{\phi}_{\text{flux}}$), and detector smearing and efficiency ($\hat{\phi}_{\text{det}}$). Comparing the experimental data from the near detector to the Monte Carlo model, experiments obtain the distribution for the parameters, $p(\hat{\phi}|\text{ND})$. The uncertainty parameters ($\hat{\phi}_i$) are applied to the far detector to predict the data distributions in the absence of oscillations. The transport of constrained uncertainties are done either by traditional covariance matrices or by a more sophisticated Markov Chain Monte Carlo that easily accounts for non-Gaussian probability distributions. The uncertainties are then marginalized or profiled to propagate the uncertainties to the oscillation analysis.

The proposed method can be used in different ways in this analysis framework. The simplest approach is to obtain the $p(\hat{\phi}|\text{ND})$ that includes all possible parameter correlations. In this case, the model provides at the same time a simple way to generate Monte Carlo to sample the distributions.

This method, using Gaussian base densities, will easily learn the nuisance parameters probability density function which is expected to be close to Gaussian. The next level of complexity is to learn, as we have done in this example, the probability density function but adding nuisance parameters, $p(p_\mu, \theta_\mu, E_\nu, \hat{\phi})$. This implementation will allow to perform unbinned likelihoods as described in this paper. The more complex and inclusive approach is to avoid the intermediate nuisance parameter density function ($p(\hat{\phi}|\text{ND})$) description and model both near and far detectors with a set of common uncertainties. The advantage of this final description is that all the analysis is carried out in a single fit avoiding the description of hundreds of uncertainties ($\hat{\phi}$).

VI. SUMMARY

In this work, we have presented the viability of a likelihood-free inference methodology through neural spline flows on a simplified neutrino oscillation problem at the T2K experiment. A brief introduction to normalizing flows as density estimators from data samples is performed, making emphasis on the particular implementation of neural spline flows. We developed a framework to use this estimation of the density from data taken at the near detector in order to perform inference of the oscillation parameters at the far detector for a simplified two-flavor neutrino problem, allowing to perform exact inference if the density is properly estimated. This method provides potential advantages over traditional binned histogram methods, specially when the statistics is low as is the case in the T2K experiment. An unbinned alternative formulated by interpolating the histogram method was constructed to check the results. Additionally, the integrity of the learned density was thoroughly verified through different statistical and empirical tests. The results obtained using the neural spline flow methodology and the unbinned likelihood methodology show results which are in agreement with each other in the estimation of the statistical errors. The alternative method is not refined enough and it shows larger bias in the estimated parameters. The results presented in this paper open new possibilities to use similar likelihood-free neural network inference for more complex statistical analyses.

ACKNOWLEDGMENTS

The authors have 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 in the preparation of this document.

APPENDIX: EXPERIMENT SIMULATION

In this appendix, we will describe how observed events (p_μ, θ_μ, E_ν) are generated for the different experiments. Because of statistical fluctuations due to an event oscillation with certain probability [following Eq. (1)], the exact number of observed events cannot be determined beforehand, but can be approximately be chosen. In order to generate the observations of a fixed set of parameters ($\theta_{\text{mix}}, \Delta m^2$), the procedure is the following:

- (1) Choose an approximated number of observed events N_{approx} .
- (2) Generate a MC event (p_μ, θ_μ, E_ν), make the event oscillate according to its energy E_ν , and accept it with probability given by Eq. (2), until accepting a total of N_{approx} events. This takes a stochastic number of N_{init} tries, which corresponds to the number of initial events before oscillating.
- (3) Generate new N_{init} number of MC events, and accept them with probability according to Eq. (2). The accepted samples form the observation set, with a number of N_{obs} samples, which is similar to N_{approx} .

To summarize, given a fixed set of parameters ($\theta_{\text{mix}}, \Delta m^2$), we determine an approximate initial number N_{init} of events before oscillating needed to generate N_{approx} oscillated observations. However, because this is a stochastic procedure and we stopped exactly when N_{approx} were generated, the procedure has to be repeated with a fixed number of N_{init} tries, giving us $N_{\text{obs}} \sim N_{\text{approx}}$ observed events.

Two kind of experiments were performed for five different set of parameters: one of low statistics (around 500 observations), with a number of the order of the real number of observed events in T2K, and one of high statistics (around 50k observations), to see how the algorithm behaves and compares to traditional methods when having a larger amount of data available.

The parameters of the initial experiments 1 and 2 in Sec. IV B are chosen following the best fit value of $\sin^2 \theta_{23}$ and Δm^2 from [31]. The parameter θ_{mix} is computed as

$$\theta_{\text{mix}} = \frac{\pi}{2} - \arcsin \sqrt{\sin^2 \theta_{23}},$$

where we have used the fact that Eq. (2) is symmetric over $\theta_{\text{mix}} = \pi/4$ and fixed it in the interval $[0, \pi/4]$. For the best fit value, maximal mixing is almost achieved ($\sin 2\theta_{\text{mix}} = 0.9986$). Experiments 3 and 4 (5 and 6) follow the same procedure, but choosing $\sin^2 \theta_{23}$ within 1 (2) σ of the best fit. For experiments 7–10 the mixing angle has the value of experiments 3 and 4, but Δm^2 is changed in order to explore the behavior on the parameter space.

- [1] Y. Fukuda *et al.* (Super-Kamiokande Collaboration), *Phys. Rev. Lett.* **81**, 1562 (1998).
- [2] Q. R. Ahmad *et al.* (SNO Collaboration), *Phys. Rev. Lett.* **87**, 071301 (2001).
- [3] M. H. Ahn *et al.* (K2K Collaboration), *Phys. Rev. D* **74**, 072003 (2006).
- [4] T. Araki *et al.* (KamLAND Collaboration), *Phys. Rev. Lett.* **94**, 081801 (2005).
- [5] Z. Maki, M. Nakagawa, and S. Sakata, *Prog. Theor. Phys.* **28**, 870 (1962).
- [6] K. Abe *et al.* (T2K Collaboration), *Phys. Rev. Lett.* **107**, 041801 (2011).
- [7] F. P. An *et al.* (Daya Bay Collaboration), *Phys. Rev. Lett.* **108**, 171803 (2012).
- [8] M. Tanabashi *et al.* (Particle Data Group), *Phys. Rev. D* **98**, 030001 (2018).
- [9] K. Abe *et al.* (T2K Collaboration), *Phys. Rev. D* **96**, 092006 (2017); **98**, 019902(E) (2018).
- [10] M. A. Acero, *et al.* (NOvA Collaboration), *Phys. Rev. Lett.* **123**, 151803 (2019).
- [11] K. Abe *et al.* (T2K Collaboration), *Nature (London)* **580**, 339 (2020).
- [12] K. Abe *et al.* (T2K Collaboration), *Phys. Rev. D* **93**, 112012 (2016).
- [13] Y. Hayato, *Acta Phys. Pol. B* **40**, 2477 (2009), https://www.actaphys.uj.edu.pl/index_n.php?I=R&V=40&N=9#2477.
- [14] K. Abe, *et al.* (T2K Collaboration), *Phys. Rev. D* **87**, 012001 (2013).
- [15] G. Papamakarios, E. T. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, [arXiv:1912.02762](https://arxiv.org/abs/1912.02762).
- [16] 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.
- [17] D. J. Rezende and S. Mohamed, [arXiv:1505.05770](https://arxiv.org/abs/1505.05770).
- [18] M. Germain, K. Gregor, I. Murray, and H. Larochelle, in *Proceedings of the 32nd International Conference on Machine Learning, Proceedings of Machine Learning Research*, edited by F. Bach and D. Blei (PMLR, Lille, France, 2015), Vol. 37, pp. 881–889.
- [19] G. Papamakarios, I. Murray, and T. Pavlakou, 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.
- [20] C.-W. Huang, D. Krueger, A. Lacoste, and A. C. Courville, [arXiv:1804.00779](https://arxiv.org/abs/1804.00779).
- [21] N. D. Cao, I. Titov, and W. Aziz, Block neural autoregressive flow (2019), [arXiv:1904.04676](https://arxiv.org/abs/1904.04676).
- [22] 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.
- [23] P. Jaini, K. A. Selby, and Y. Yu, in *Proceedings of the 36th International Conference on Machine Learning, Proceedings of Machine Learning Research*, edited by Chaudhuri and R. Salakhutdinov (PMLR, Long Beach, 2019), Vol. 97, pp. 3009–3018.
- [24] T. Müller, B. McWilliams, F. Rousselle, M. Gross, and J. Novák, *ACM Trans. Graph.* **38**, 145 (2018).
- [25] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, [arXiv:1906.02145](https://arxiv.org/abs/1906.02145).
- [26] J. A. Gregory and R. Delbourgo, *IMA J. Numer. Anal.* **2**, 123 (1982), <https://academic.oup.com/imajna/article-pdf/2/2/123/2267745/2-2-123.pdf>.
- [27] M. Tanabashi, K. Hagiwara, K. Hikasa, K. Nakamura *et al.* (Particle Data Group), *Phys. Rev. D* **98**, 030001 (2018).
- [28] R. Barlow, *Nucl. Instrum. Methods Phys. Res., Sect. A* **297**, 496 (1990).
- [29] D. P. Kingma and J. Ba, CoRR abs/1412.6980 (2014), <https://arxiv.org/abs/1412.6980>.
- [30] M.-H. Chen and Q.-M. Shao, *J. Comput. Graph. Stat.* **8**, 69 (1999).
- [31] K. Abe, R. Akutsu, A. Ali, J. Amey, C. Andreopoulos, L. Anthony, M. Antonova, S. Aoki, A. Ariga, Y. Ashida, Y. Azuma, S. Ban, M. Barbi, G. J. Barker *et al.* (T2K Collaboration), *Phys. Rev. Lett.* **121**, 171802 (2018).