Efficient Monte-Carlo event generation for neutrino-nucleus exclusive cross sections

Mathias El Baz, University of Geneva
Natalie Jachowicz, Ashish Kumar Jha, Kajetan Niewczas,
Alexis Nikolakopoulos, Federico Sánchez-Nieto

MITP seminar
16/05/2025

Introduction

- A bit of context:
  - PhD at the University of Geneva,
    in collaboration with Ghent University.
- In this talk :
  - The development of a **general-purpose sampling scheme for exclusive neutrino-nucleus cross sections**.
  - The application of this method to the 1p1h Exclusive cross section using the **Mean Field model** of the Ghent group.

Motivations

- Neutrino-nucleus interaction modeling will limit sensitivity of next-generation oscillation experiments (T2HK, DUNE).
- Some detectors already observe low-momentum protons.

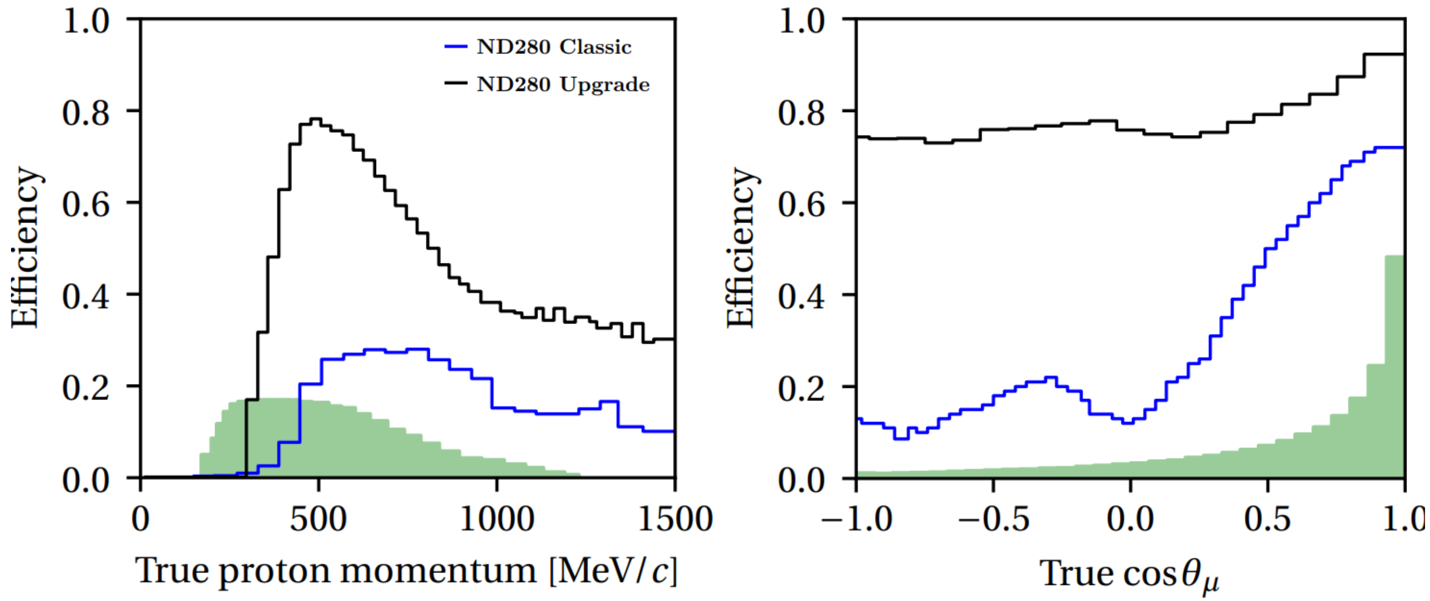

T2K's near detector upgrade allows to observe particles in a broader kinematic range

Implementations in MC generators

- Most of the implemented models provide only inclusive cross sections.
- Exclusive implementations rely on factorization
- E.g., NEUT uses this approach for QE interactions.

$$\frac{d^6\sigma_{\nu-\mathrm{Nucleus}}}{dE_\mu d\Omega_\mu d\Omega_N dE_N} \sim \mathrm{S}(\mathrm{E_m}, \mathrm{p_m}) \mathrm{L}_{\mu\nu} \mathrm{W}^{\mu\nu} \delta(\omega + M - E_m - E_p)$$

⇒ The outgoing nucleons are described as plane waves

⇒ This neglects final state interactions (interaction with other nucleons).

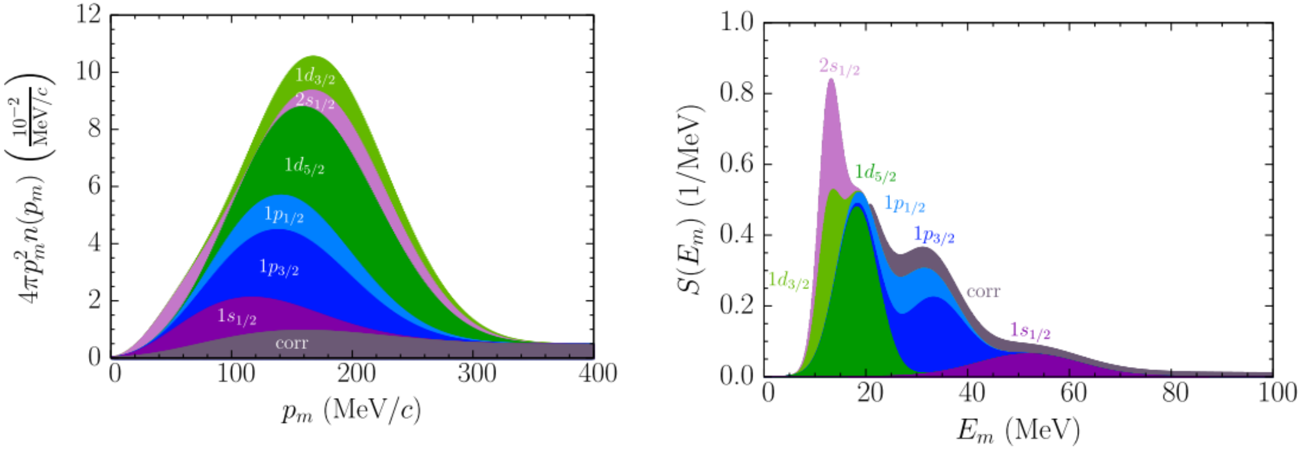

Argon spectral function (Omar Benhar)

Plane wave vs Distorted wave

- Outgoing nucleon state = Sum of distorted wave solutions to a nuclear potential (EDRMF, ROP…)
- Include elastic (real potential) and inelastic (imaginary or cascade) interactions.

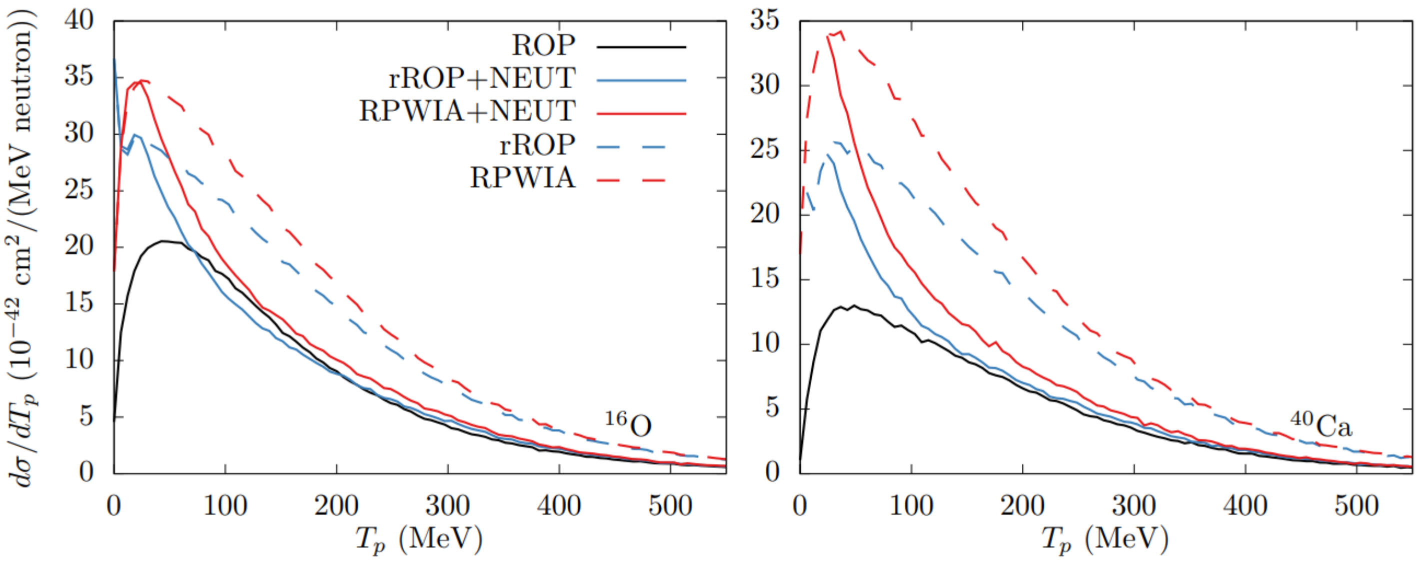

Cross section in terms of the leading protons kinetic energy averaged over the T2K flux (Alexis Nikolakopoulos)

Implementation of EDRMF in NEUT

- EDRMF and ROP in the next NEUT version by J. McKean et al.
- First fully exclusive implementation using the distorted wave approach.
- QE interactions for C12 and O16.

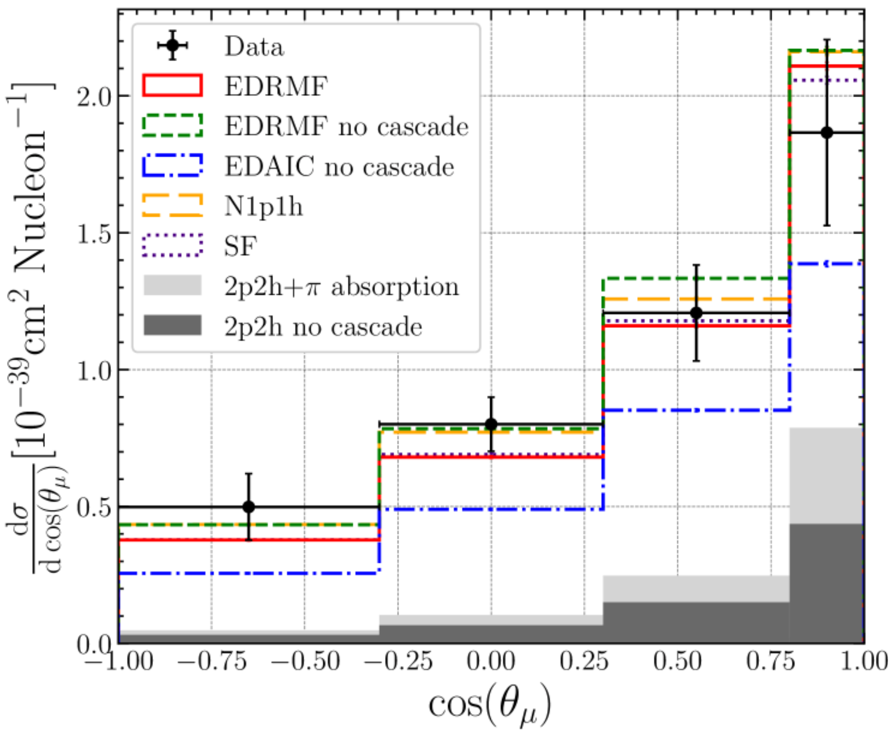

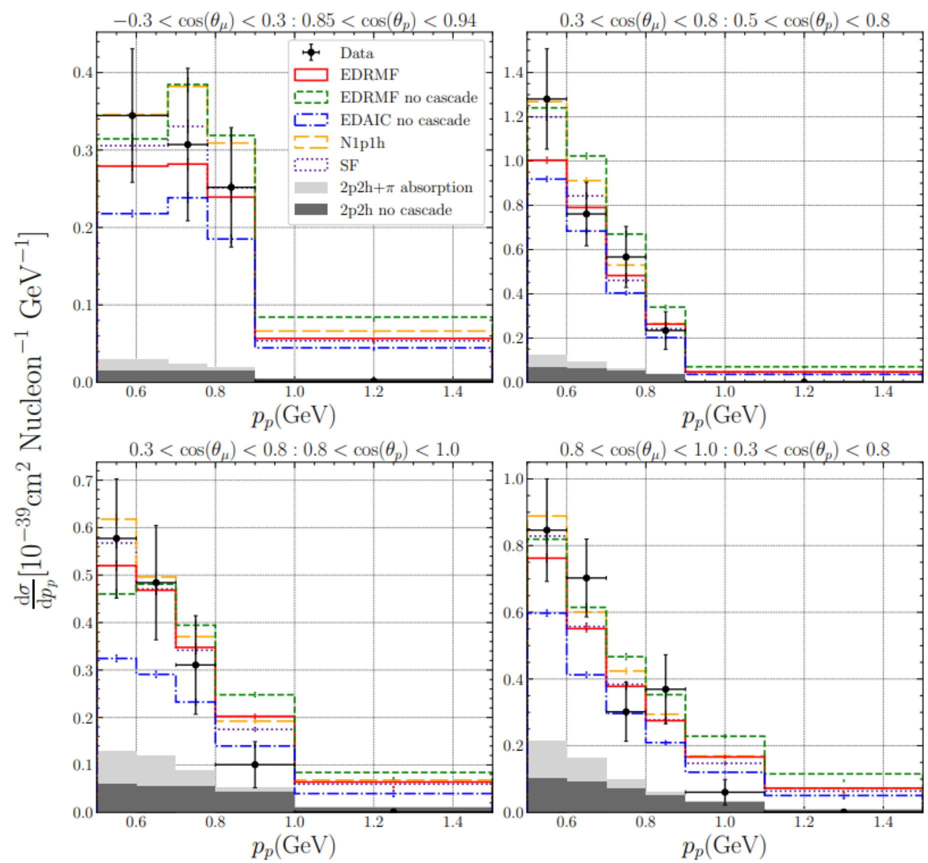

Cross section predictions on hydrocarbon with the new NEUT EDRMF implementation (Jake McKean). Other processes are described by an LFG

Computational cost of DW cross sections

Summation of a huge number of DW solutions
Low evaluation speed → Low sampling speed

↓

Current solution : Precomputing hadron response tables
Interpolation between the precomputed points and accept-reject

↓

Tables size scales exponentially with dimension

Accept-Reject sampling inefficient in high dimension

↓

Objective : Solve these issues

Hartree-Fock Mean field

- **Shell** model describing nucleons in a **mean nucleus potential**.
- **Non-relativistic** model
- Nucleon-nucleus potential derived iteratively via a **Hartree-Fock** procedure from an effective **Skyrme** nucleon-nucleon potential.

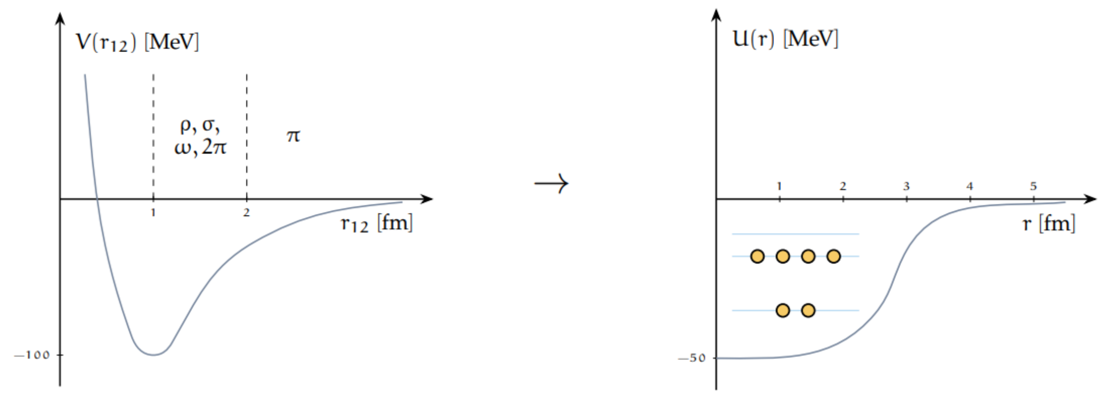

From Nucleon-nucleon potential to Nucleon-Nucleus potential (Kajetan Niewczas)

- Both initial and final nucleon state **solutions of the Schrödinger equation** with the Mean Field potential
- Include directly physical properties such as **Pauli Blocking**

1p1h Mean field framework

1p1h process $$\nu_\mu + {}^{A}_{Z}X \longrightarrow \mu^- + p + \left({}^{A-1}_{Z-1}X\right)^*$$

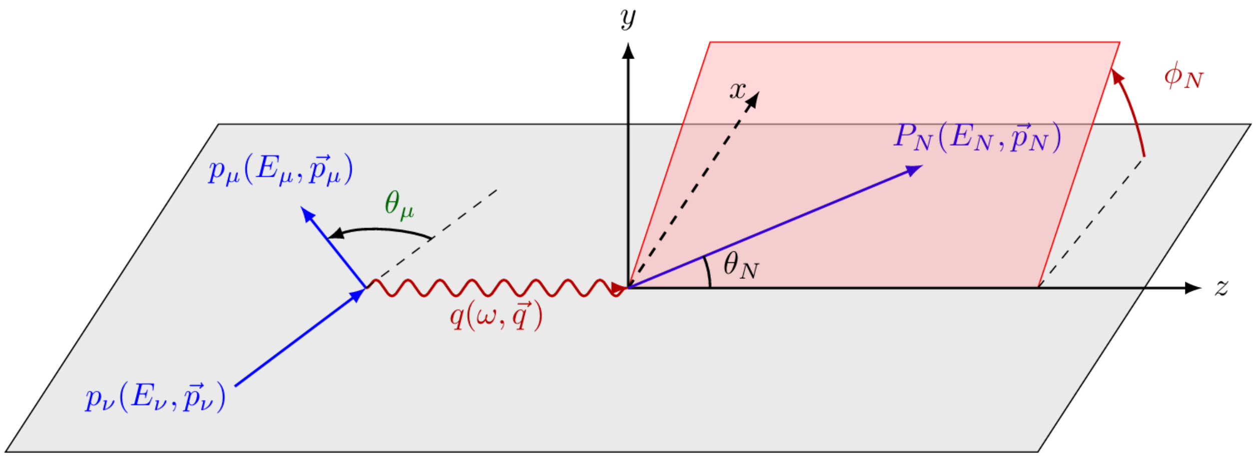

Differential cross section $$\frac{d^6\sigma(E_\nu)}{d\omega\, d\Omega_\mu\, dE_N\, d\Omega_N} \propto L_{\mu\nu} W^{\mu\nu}$$

Hadron tensor $$W^{\mu\nu} = \overline{\sum_{i,f}} \langle i|\mathcal{J}^{\mu\dagger}|f\rangle\langle f|\mathcal{J}^{\nu}|i\rangle \delta(E_N + E^*_{A-1} - M_A - \omega)$$

Dimension reduction

- **Invariance** by rotation around the momentum transfer axis
- We consider nucleon knockout to the continuum

$\Rightarrow$ Residual system can then only be left with an energy respecting :

$$\epsilon_\alpha = M_N + E^*_{A-1} - M_A$$

Initial energy for shell α

Excitation energy

$\Rightarrow$ Can be decomposed in a sum of partial cross sections :

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

The 1p1h exclusive cross section depends on **four** kinematic variables $(\theta_\mu, \omega, \theta_N, \phi_N)$ and is parametrized by a continuous parameter (E_ν) and a discrete or continuous parameter (ϵ_α).

Other processes

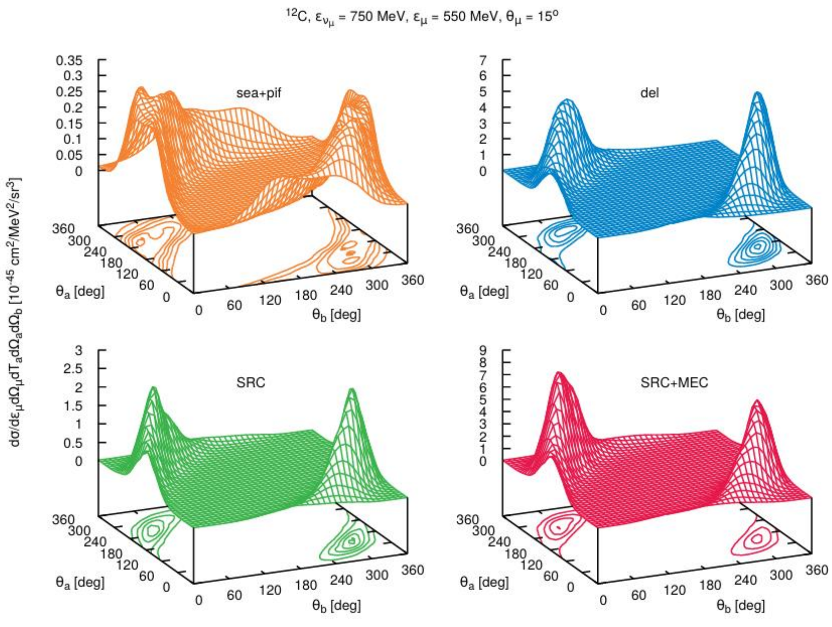

Exclusive 2p2h
cross section on Carbon
(Kajetan Niewczas)

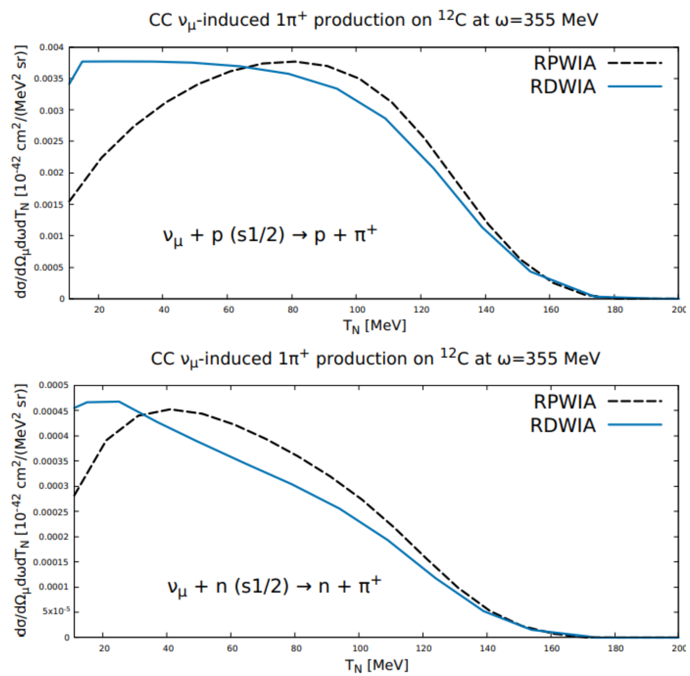

Exclusive SPP
cross section on Carbon
(Javier García-Marcos)

This work **focuses on the CCQE** process but **efforts are underway** to incorporate all processes in full complexity.

FACULTÉ DES SCIENCES

Key take-aways so far

- **We lack** of distorted wave-based cross sections in MC generators for exclusive predictions.
- We have two problems :
  1. **Computationally demanding**
  2. **Inefficient sampling high dimensional correlated distributions**.

$\Rightarrow$ Evaluation time slower for the 2p2h cross section !

Can **Machine Learning** work its dark magic to solve this problem?

UNIVERSITÉ DE GENÈVE
DEPARTEMENT DE PHYSIQUE NUCLEAIRE ET CORPUSCULAIRE

Normalizing flows

- **Construct a complex probability** transforming a simple base distribution (e.g. gaussian, uniform…)
- **Learnable diffeomorphism** of the probability space.
- Not just a simple fit of probability :
  - **Sampling is straightforward** : $\boldsymbol{x} = T(\boldsymbol{u}) \sim p_{\boldsymbol{x}}$ with $\boldsymbol{u} \sim p_{\boldsymbol{u}}$
  - **Evaluating the probability is also straightforward** :

$$p_{\boldsymbol{x}}(\boldsymbol{x}) = p_{\boldsymbol{u}}(\boldsymbol{u}) \,|\det\big(\mathrm{J}_T(\boldsymbol{u})\big)|^{-1}$$

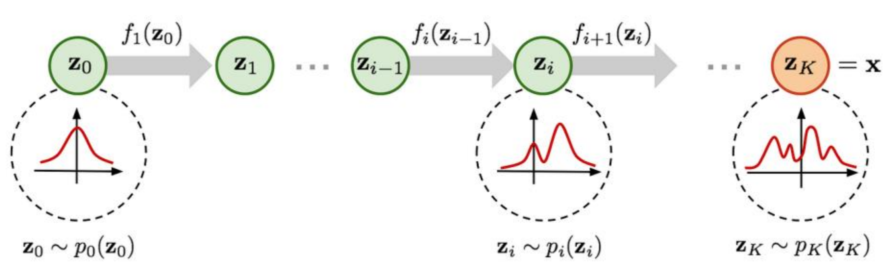

Event generation with NF in HEP

Event Generation with Normalizing Flows

Christina Gao,[1] Stefan Höche,[1] Joshua Isaacson,[1] Claudius Krause,[1] and Holger Schulz[2]
[1] *Fermi National Accelerator Laboratory, Batavia, IL, 60510, USA*
[2] *Department of Physics, University of Cincinnati, Cincinnati, OH 45219, USA*

We present a novel integrator based on normalizing flows which can be used to improve the unweighting efficiency of Monte Carlo event generators for collider physics simulations. In contrast to machine learning approaches based on surrogate models, our method generates the correct result even if the underlying neural networks are not optimally trained. We exemplify the new strategy using the example of Drell-Yan type processes at the LHC, both at leading and partially at next-to-leading order QCD.

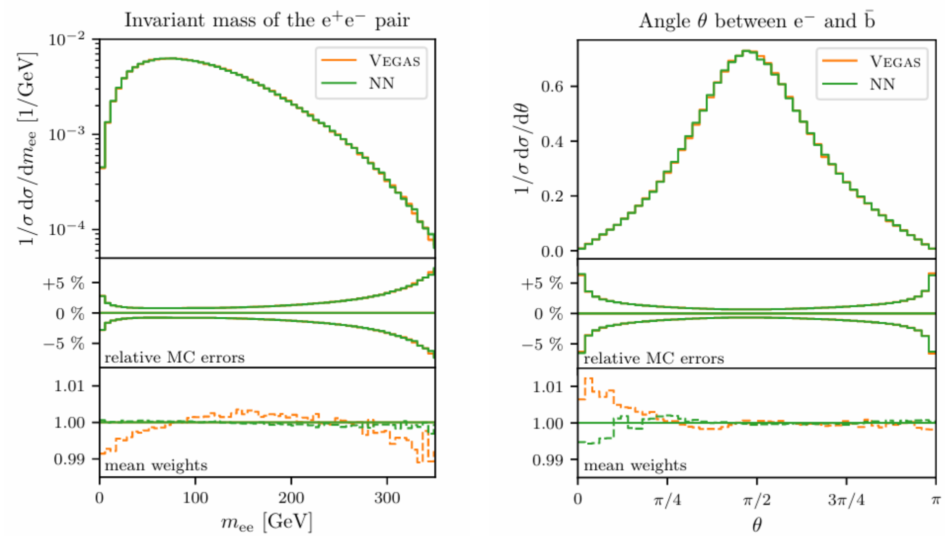

$pp \to W/Z + nj$
(Gao et al.)
2001.10028, PRD

Top-quark pair production
(Bothmann et al.)
https://arxiv.org/pdf/2203.07460

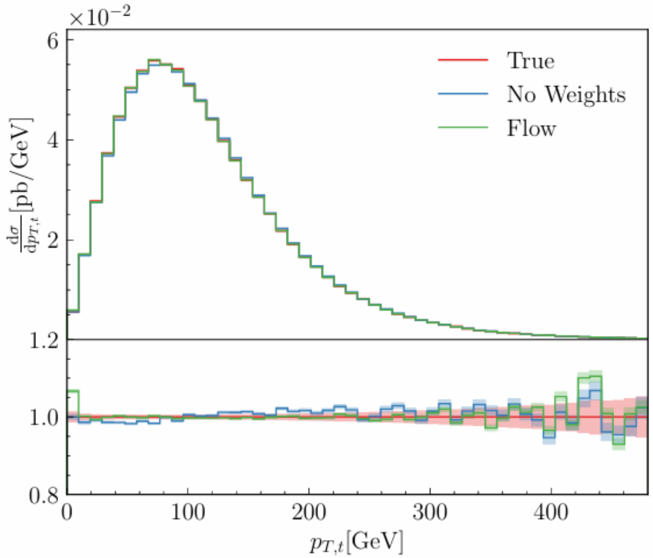

Top pair production events
(Stienen et al.)
https://scipost.org/SciPostPhys.10.2.038/pdf

Autoregressive flows

Computation of $|\det(J_T(\boldsymbol{u}))|^{-1}$ is expensive ($\sim O(D^3)$)

$$z'_i = \tau(z_i; \boldsymbol{h}_i) \quad \text{where} \quad \boldsymbol{h}_i = c_i(\mathbf{z}_{<i})$$

1D transformation

Conditioner "The NN"

$$J_{f_\phi}(\mathbf{z}) = \begin{bmatrix} \frac{\partial \tau}{\partial z_1}(z_1; \boldsymbol{h}_1) & & \mathbf{0} \\ & \ddots & \\ \mathbf{L}(\mathbf{z}) & & \frac{\partial \tau}{\partial z_D}(z_D; \boldsymbol{h}_D) \end{bmatrix}$$

Triangular Jacobian

$$\log\left|\det J_{f_\phi}(\mathbf{z})\right| = \log\left|\prod_{i=1}^{D} \frac{\partial \tau}{\partial z_i}(z_i; \boldsymbol{h}_i)\right| = \sum_{i=1}^{D} \log\left|\frac{\partial \tau}{\partial z_i}(z_i; \boldsymbol{h}_i)\right|$$

Autoregressive flows implementation

The NN implementation is **easy** !

It is a **standard neural network with masked connections**

What we want to know is the “distortion at a given point z”

$$z'_i = \tau(z_i; \boldsymbol{h}_i) \quad \text{where} \quad \boldsymbol{h}_i = c_i(\mathbf{z}_{<i})$$

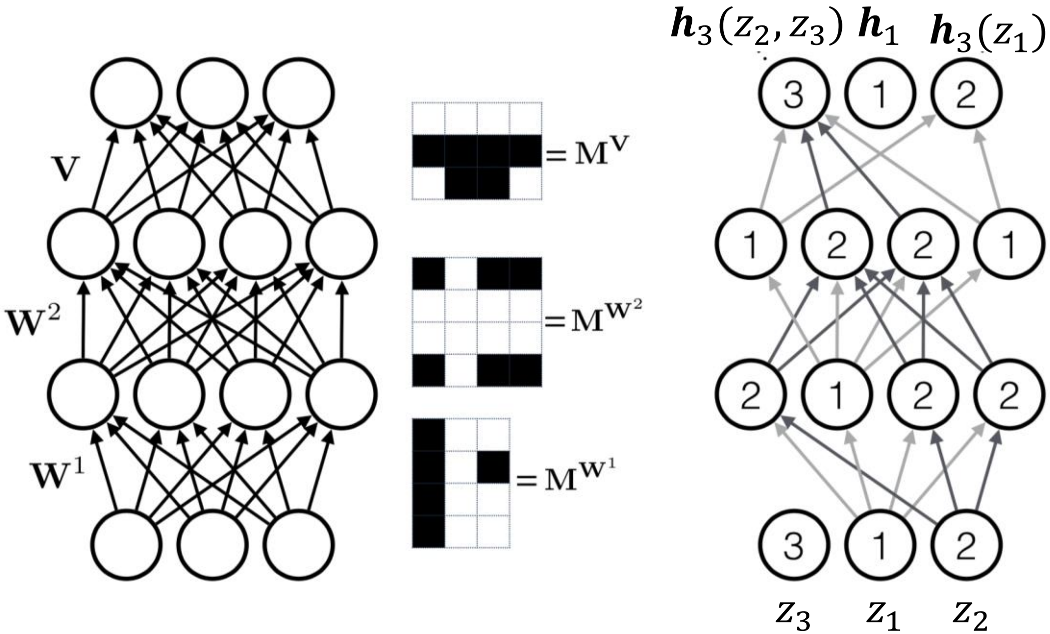

Masked Autoregressive Network

Rational quadratic Neural Spline Flows

$$z'_i = \tau(z_i; \boldsymbol{h}_i) \quad \text{where} \quad \boldsymbol{h}_i = c_i(\mathbf{z}_{<i})$$

- Each transformer is parametrized by a **monotonous spline made of piecewise rational quadratic functions**.
- One of the **most expressive** flows

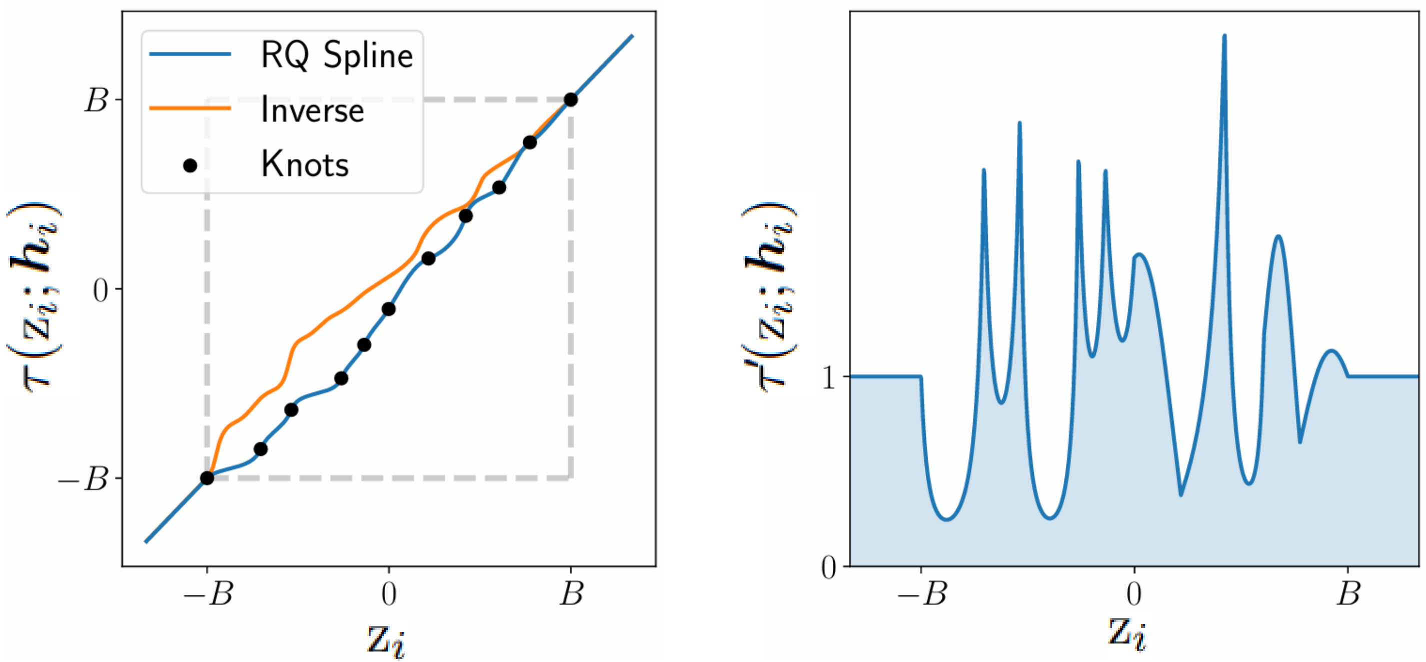

RQ-NSF and its derivative
(https://arxiv.org/pdf/1906.04032)

NF on non-Euclidean manifolds

- Differential cross sections depend on angular quantities
- Need to be defined on there natural spaces (circle, tori, sphere …)

Normalizing Flows on Tori and Spheres

Danilo Jimenez Rezende [* 1] **George Papamakarios** [* 1] **Sébastien Racanière** [* 1] **Michael S. Albergo** [2]
Gurtej Kanwar [3] **Phiala E. Shanahan** [3] **Kyle Cranmer** [2]

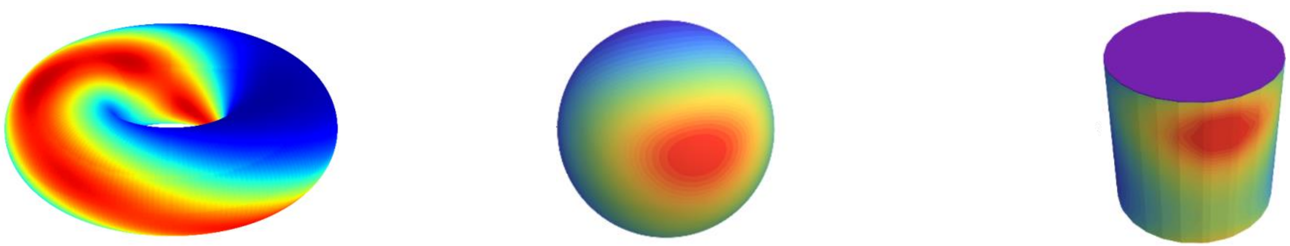

https://arxiv.org/pdf/2002.02428

- The two inclusive kinematics (ω, θ_μ) can be represented on a **cylinder** C and the two nucleon angles (θ_N, ϕ_N) on a **sphere** S^2

Manifold $\mathcal{M} = C \times S^2$

Energy Dependent Flows

- Neutrino event generators sample events based on the **conditional cross section given E_ν.**
- We mimic this by using **conditional NF** conditioned on E_ν.

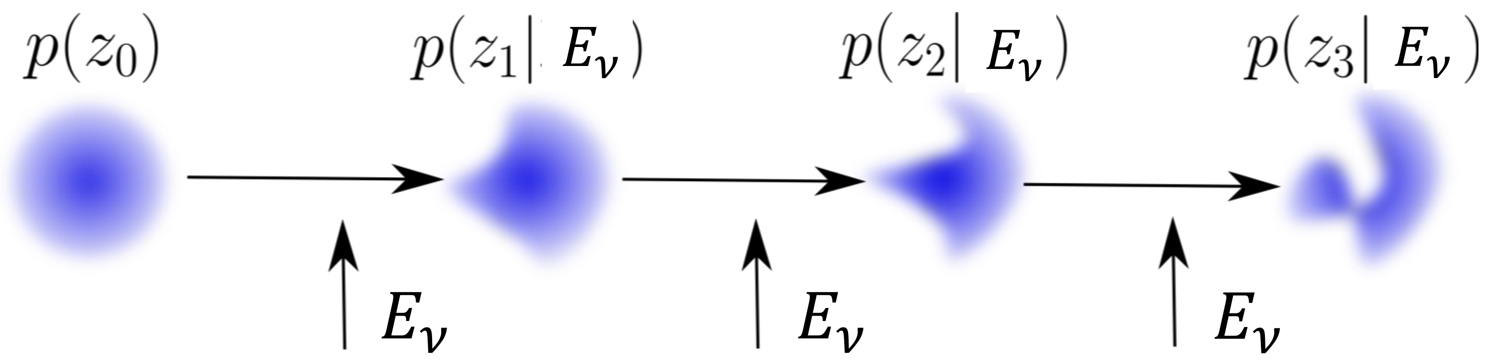

No need to retrain for different E_ν !

Training and sampling process

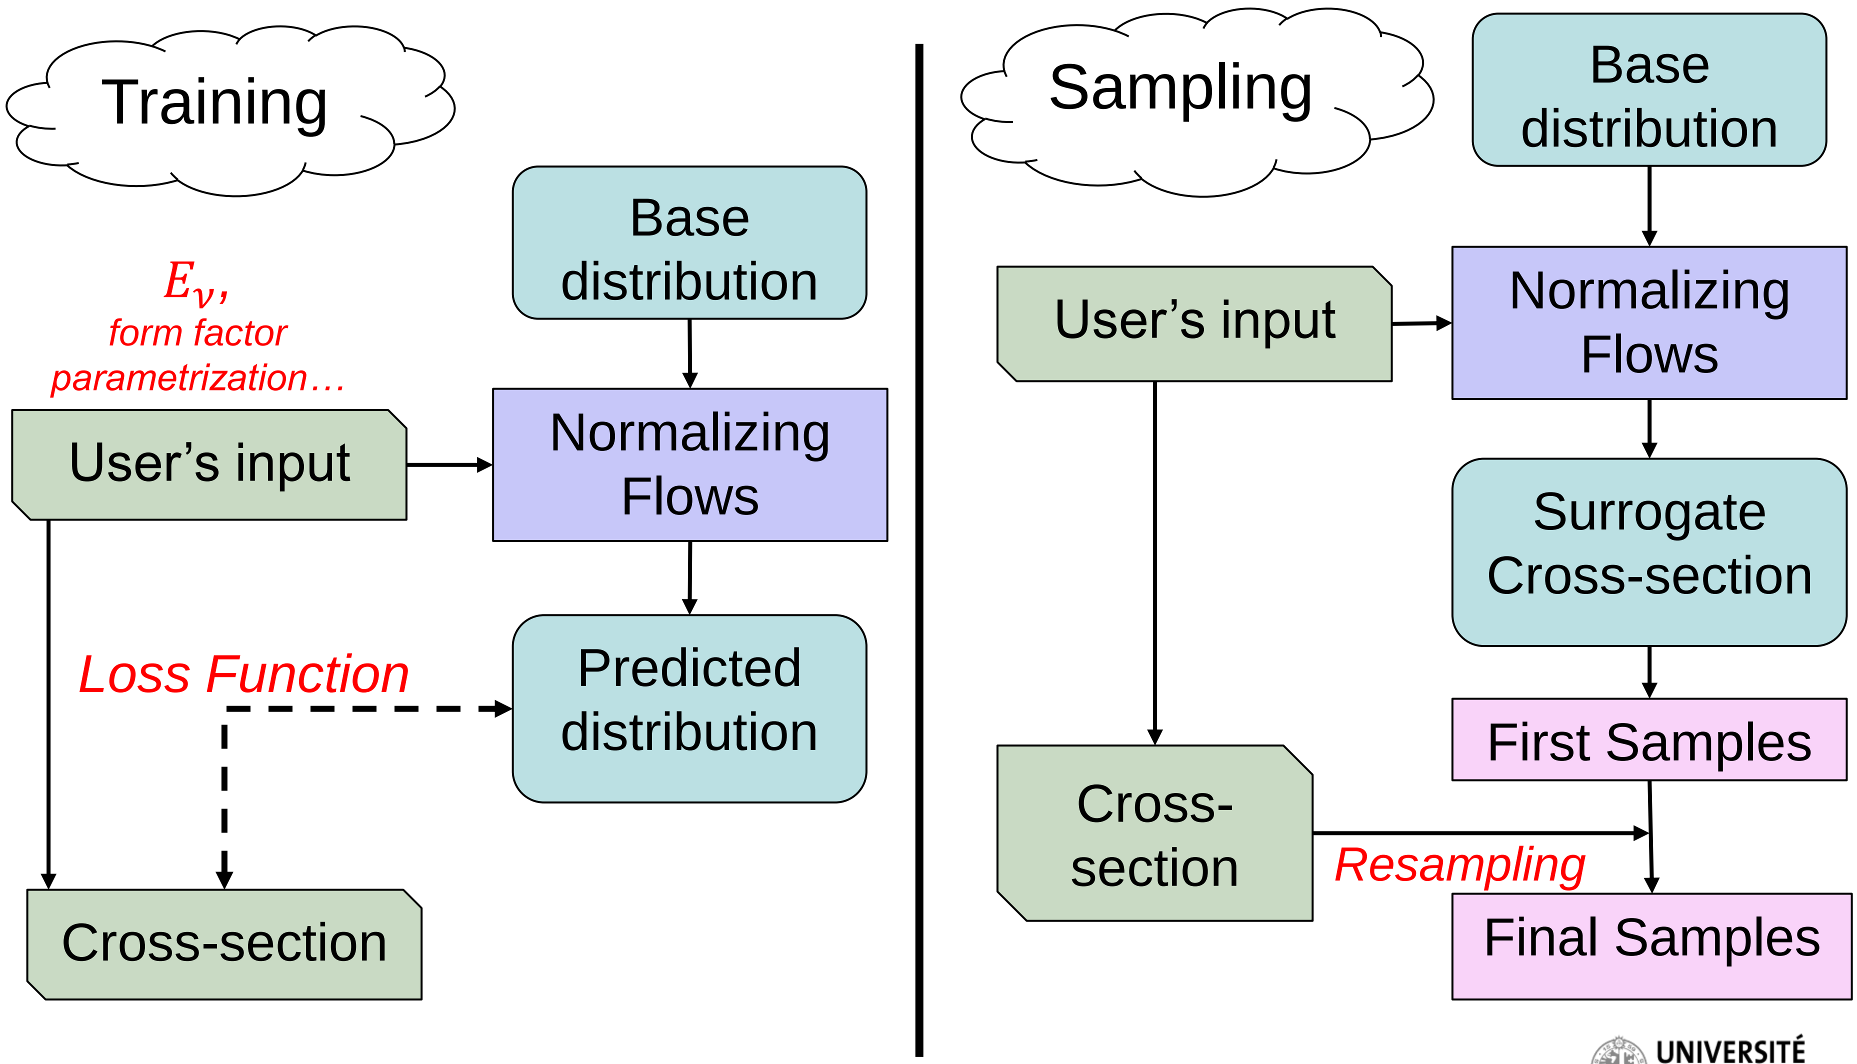

200 MeV, 1p shell

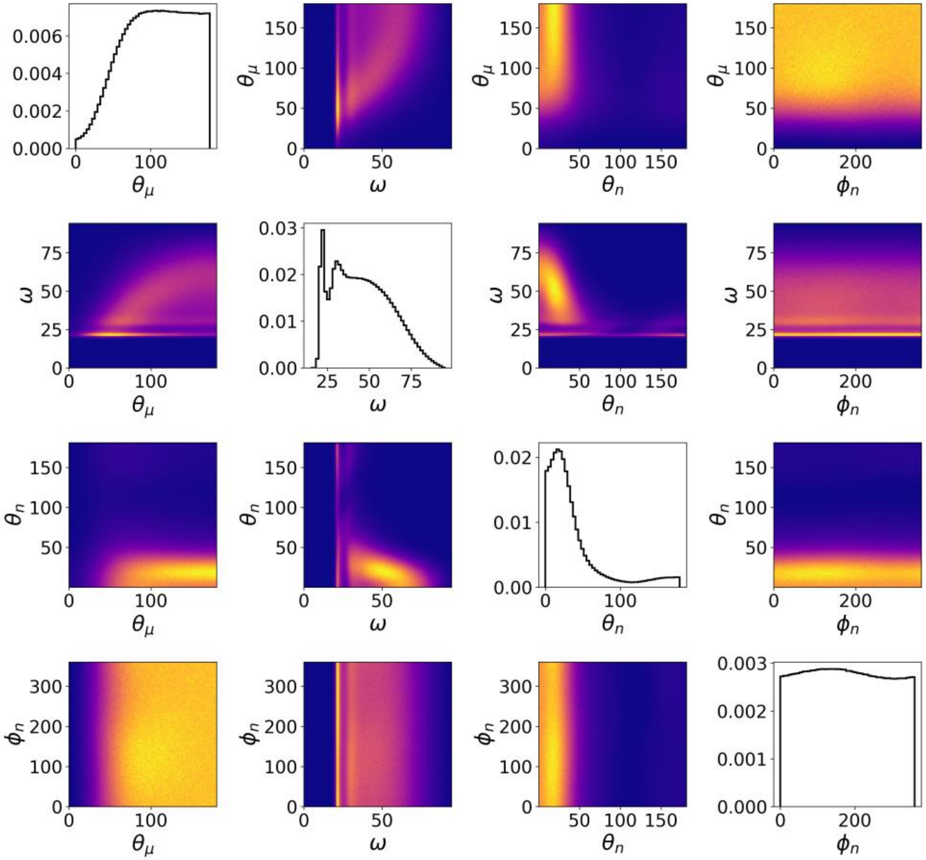

True Cross section

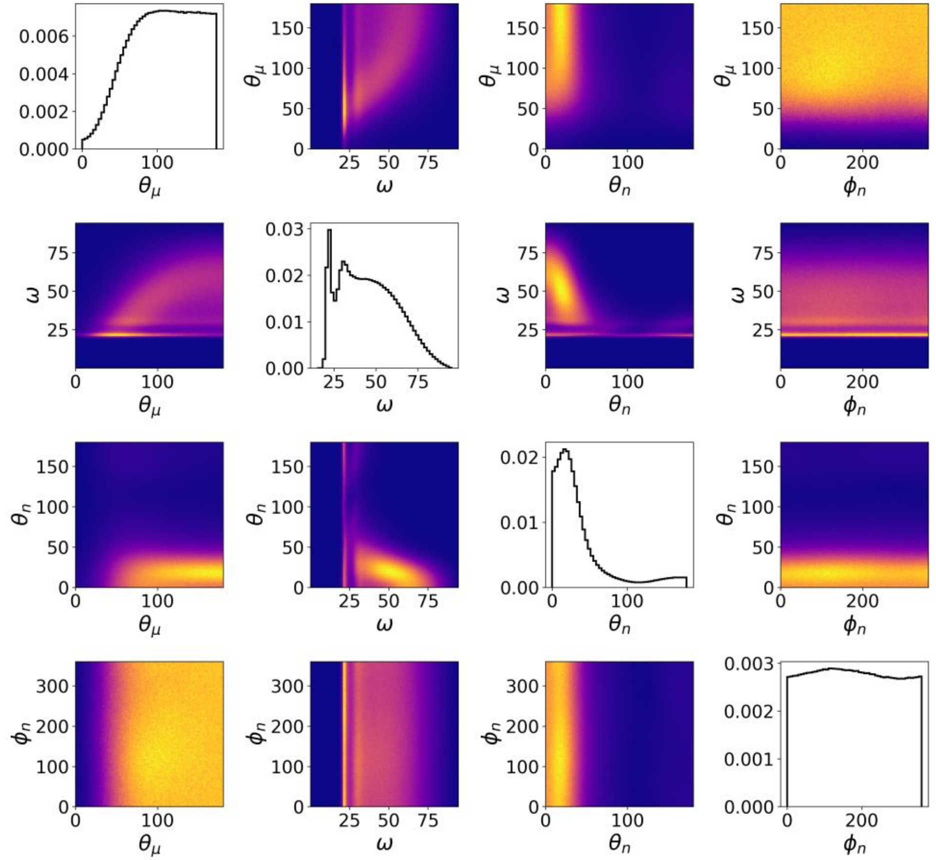

NF Surrogate Cross section

200 MeV, 1p shell

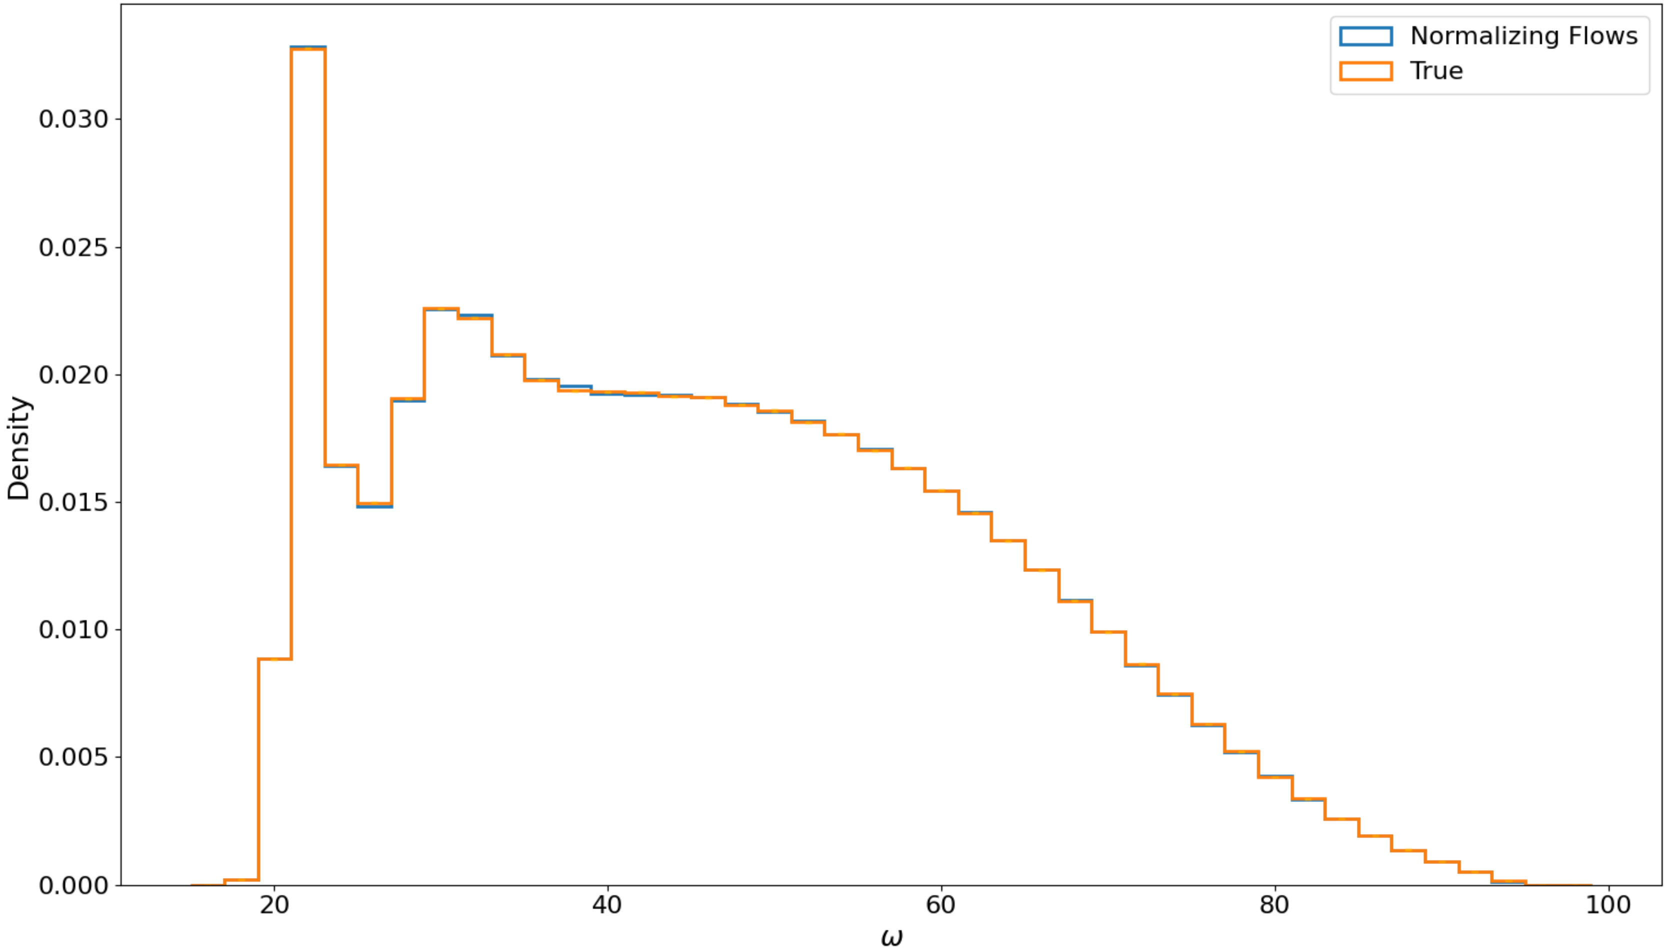

700 MeV, 1s shell

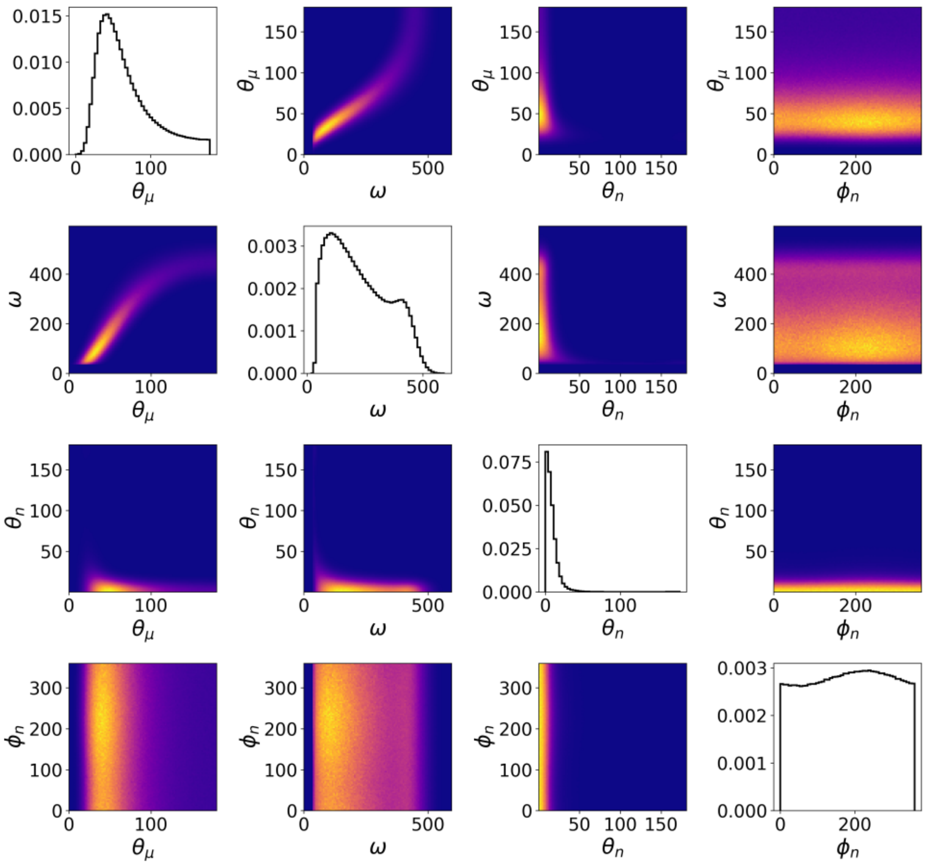

True Cross section

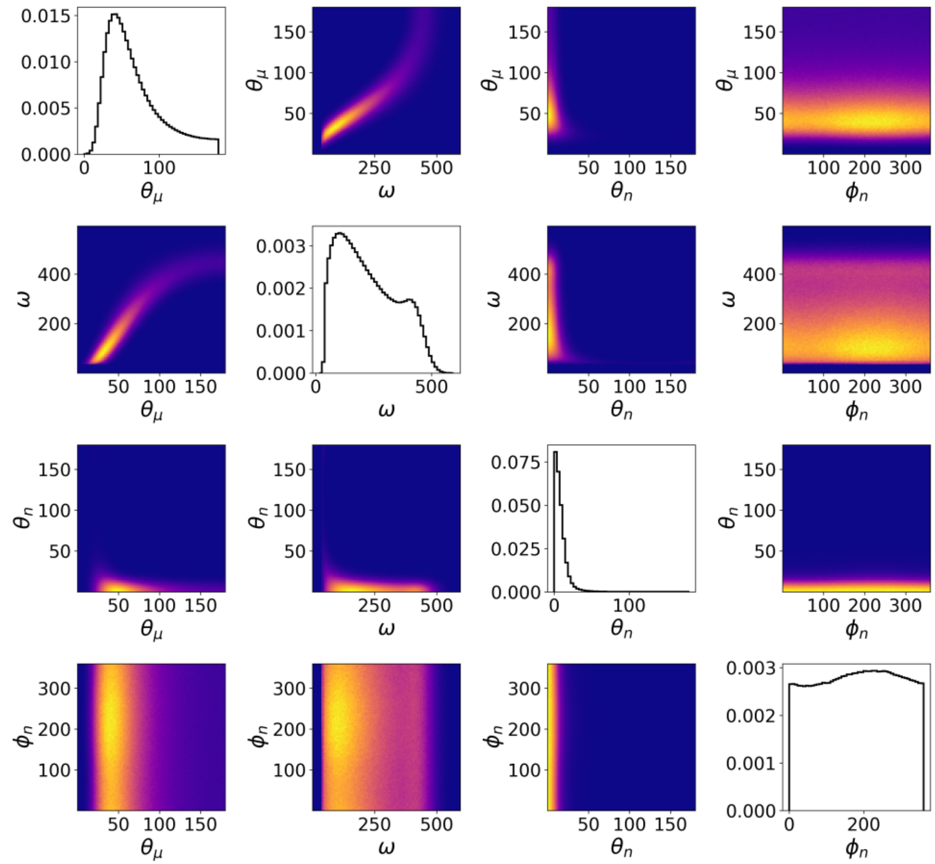

NF Surrogate Cross section

700 MeV, 1s shell

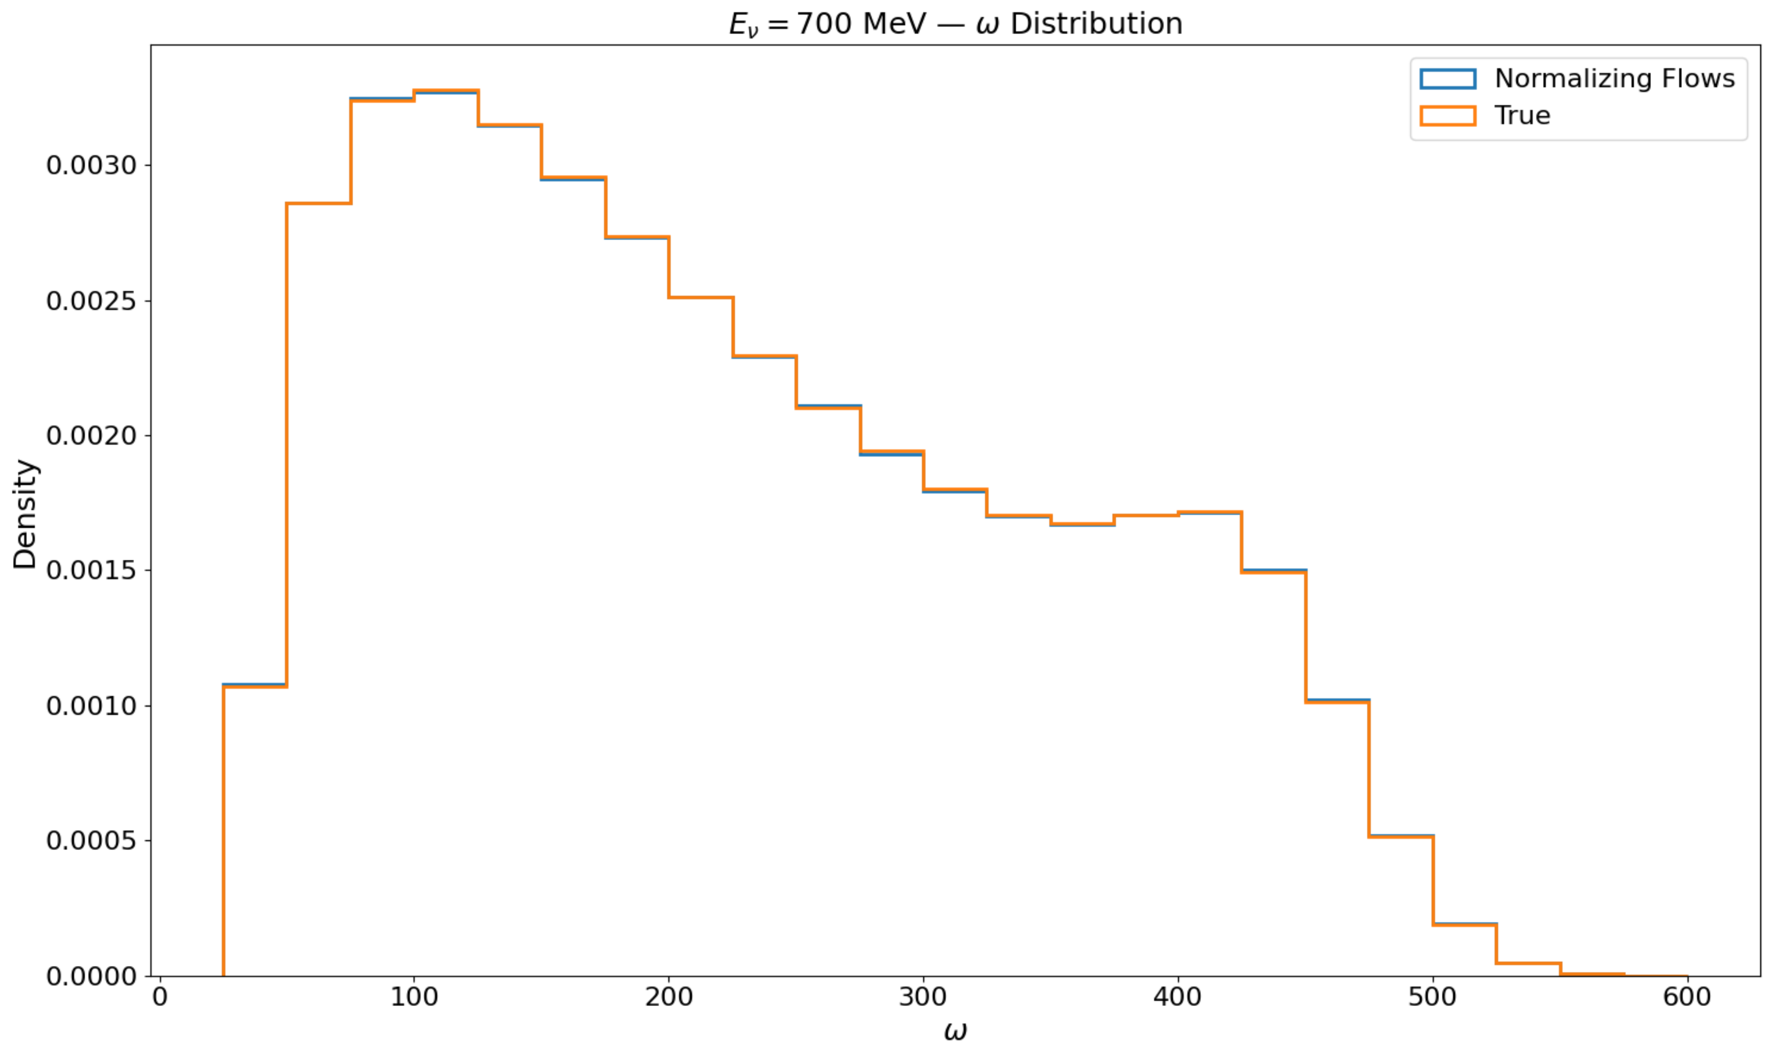

Improving unweighting efficiency

1 million events, both shells, flat flux $E_\nu = [200, 1000]\ \mathrm{MeV}$

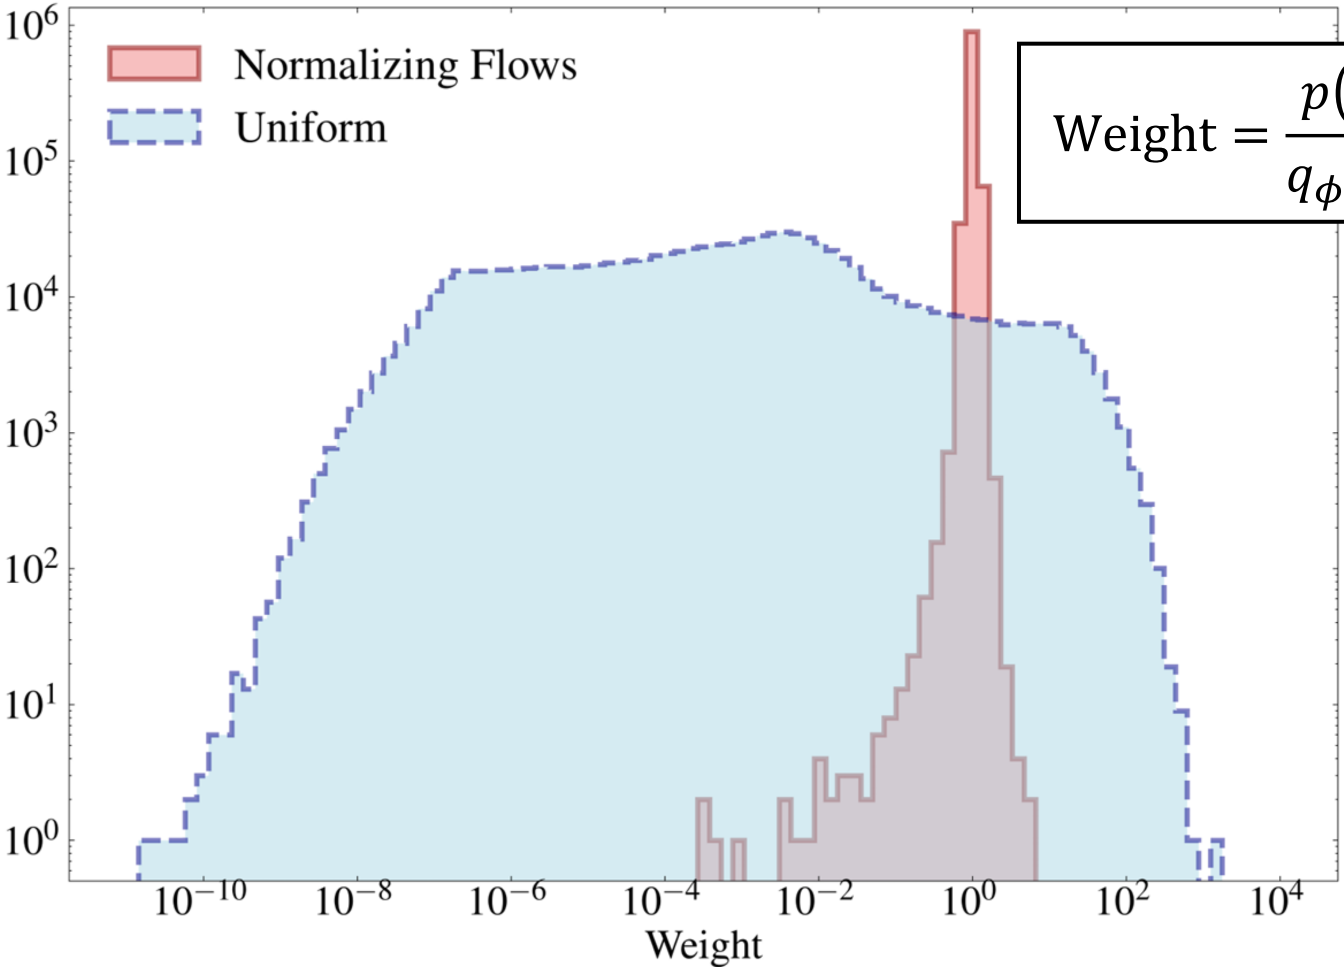

$$\text{Weight} = \frac{p(\theta_\mu, \omega, \theta_N, \varphi_N | E_\nu, \alpha)}{q_\phi(\theta_\mu, \omega, \theta_N, \varphi_N | E_\nu, \alpha)}$$

Improving unweighting efficiency

2 million events, both shells, flat flux $E_\nu = [200, 1000]\ \mathrm{MeV}$

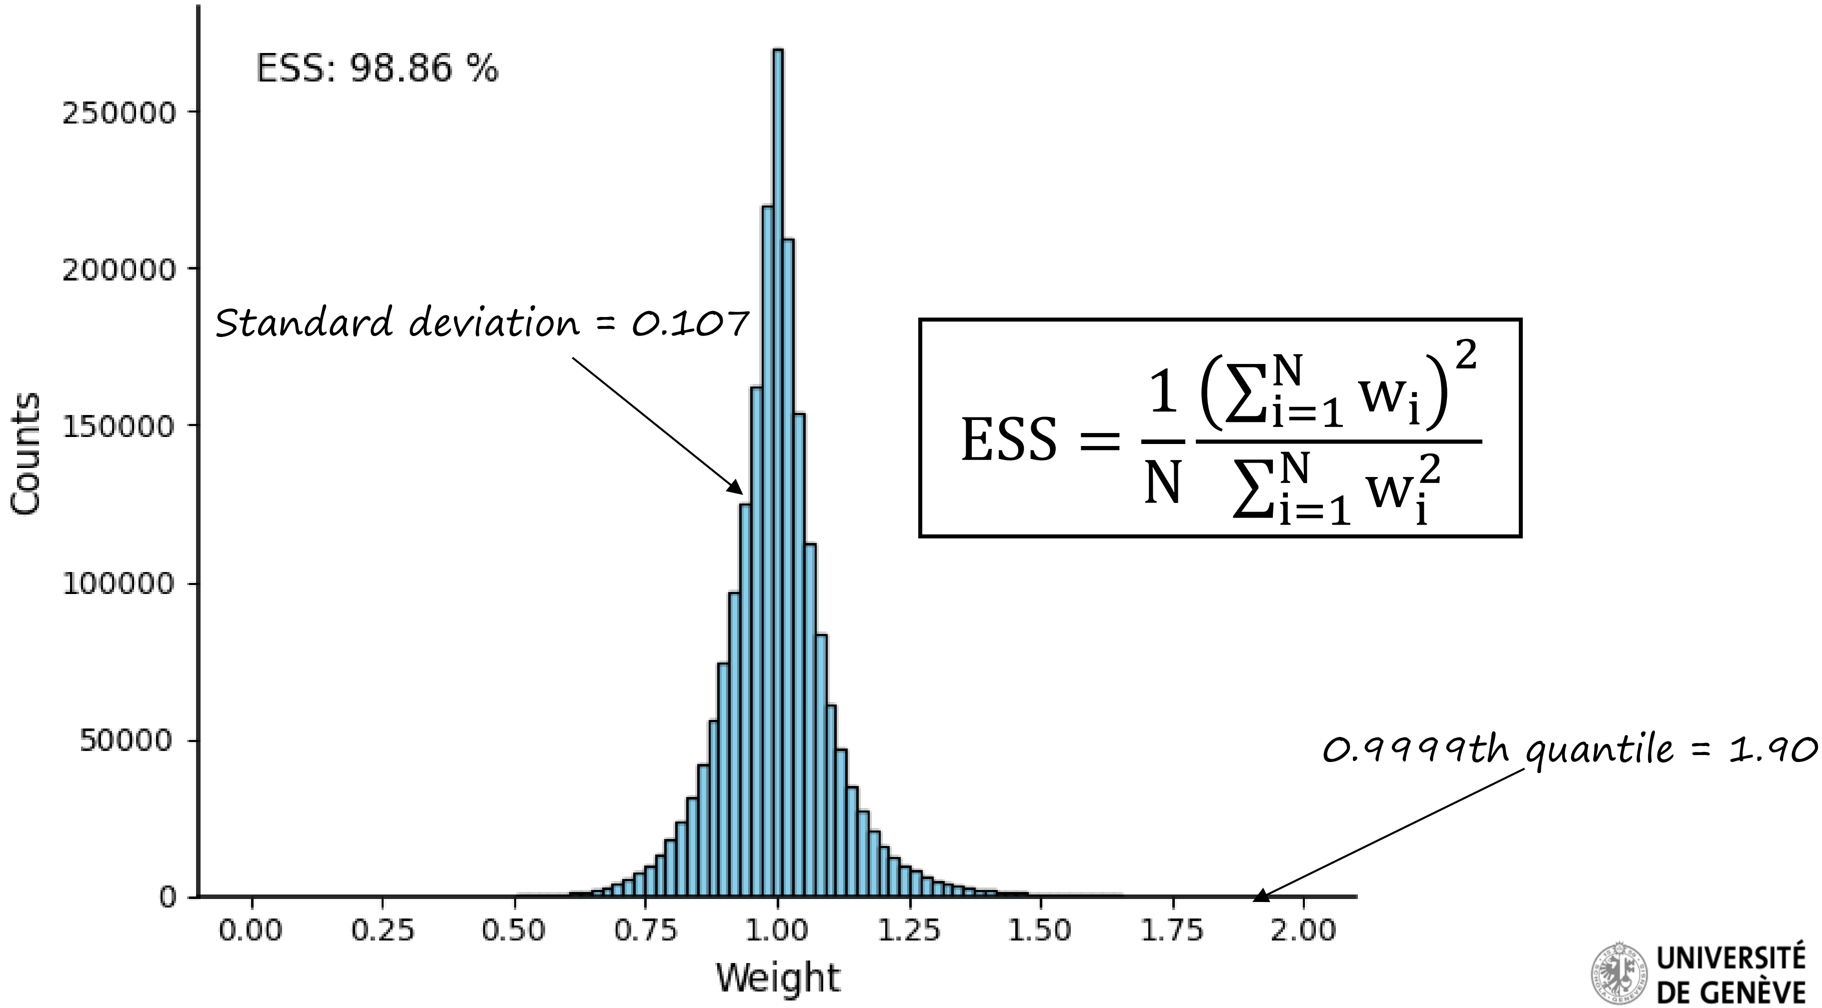

$$\mathrm{ESS} = \frac{1}{\mathrm{N}} \frac{\left(\sum_{\mathrm{i}=1}^{\mathrm{N}} \mathrm{w_i}\right)^2}{\sum_{\mathrm{i}=1}^{\mathrm{N}} \mathrm{w_i^2}}$$

The origin of the spread of weights

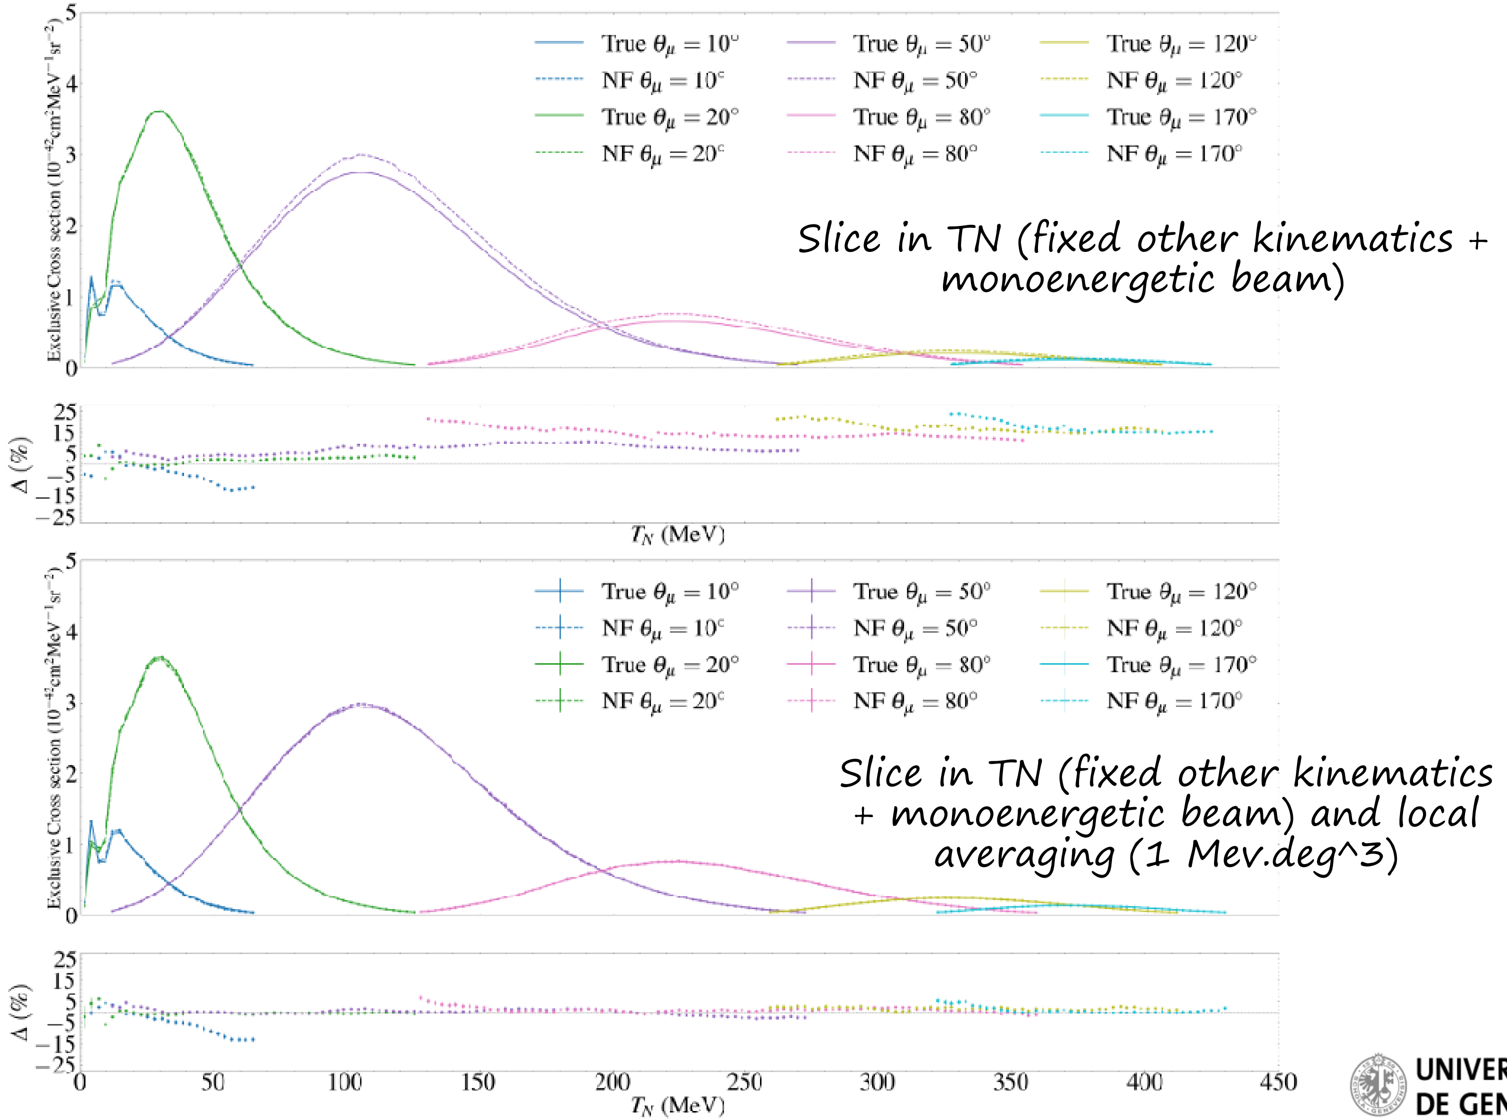

The origin of the spread of weights

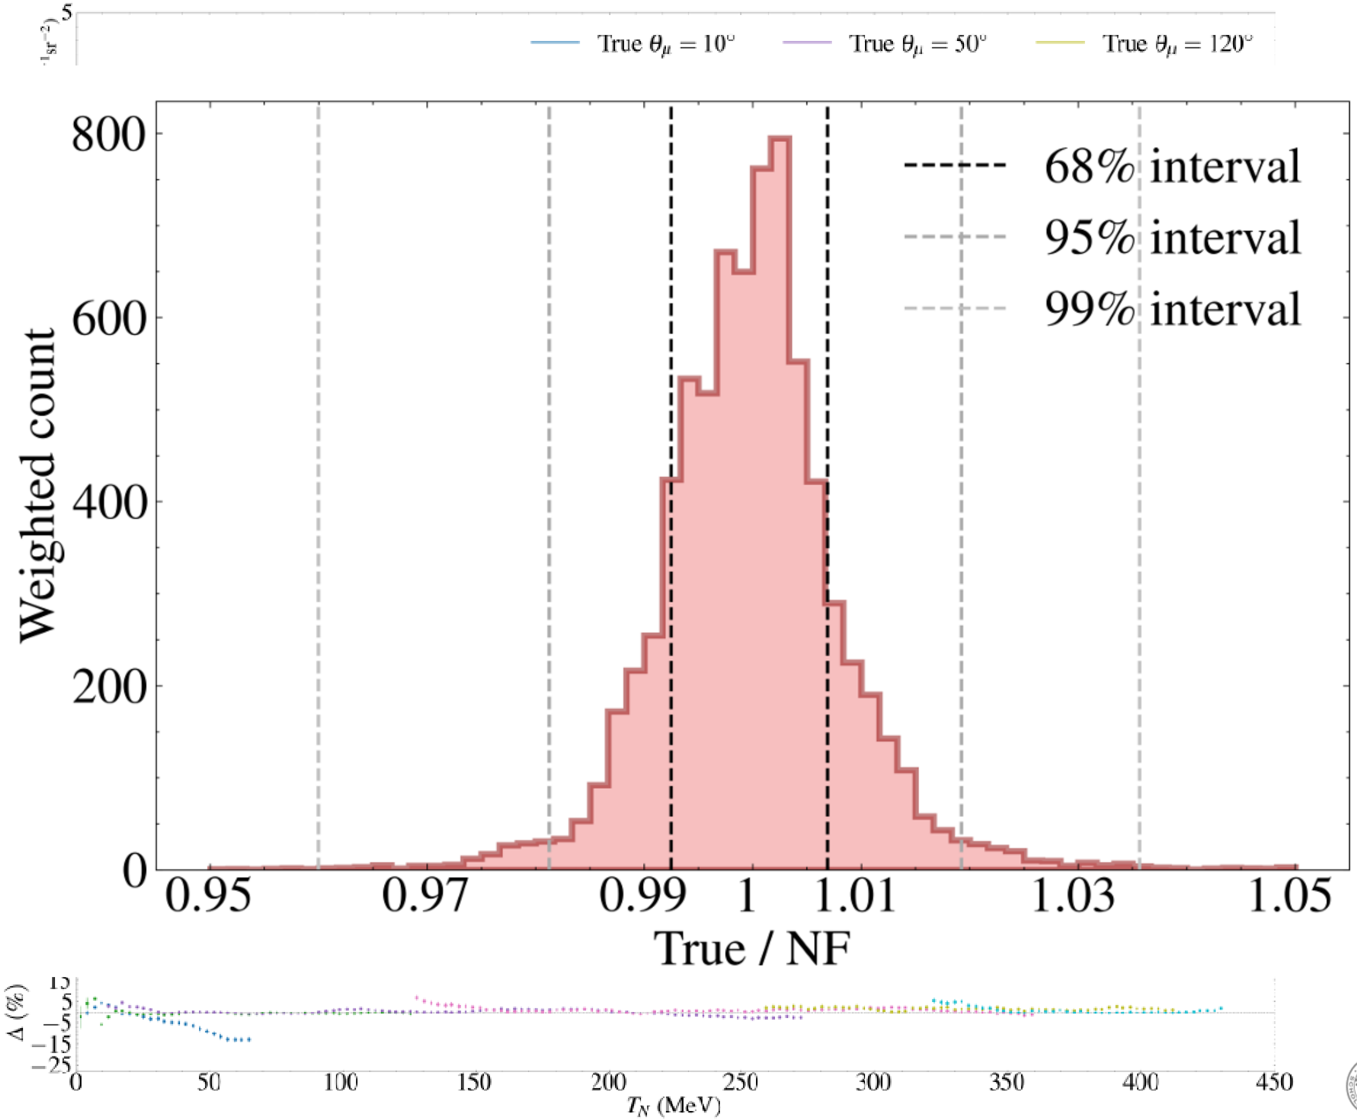

atics +

matics
d local
3)

The origin of the spread of weights

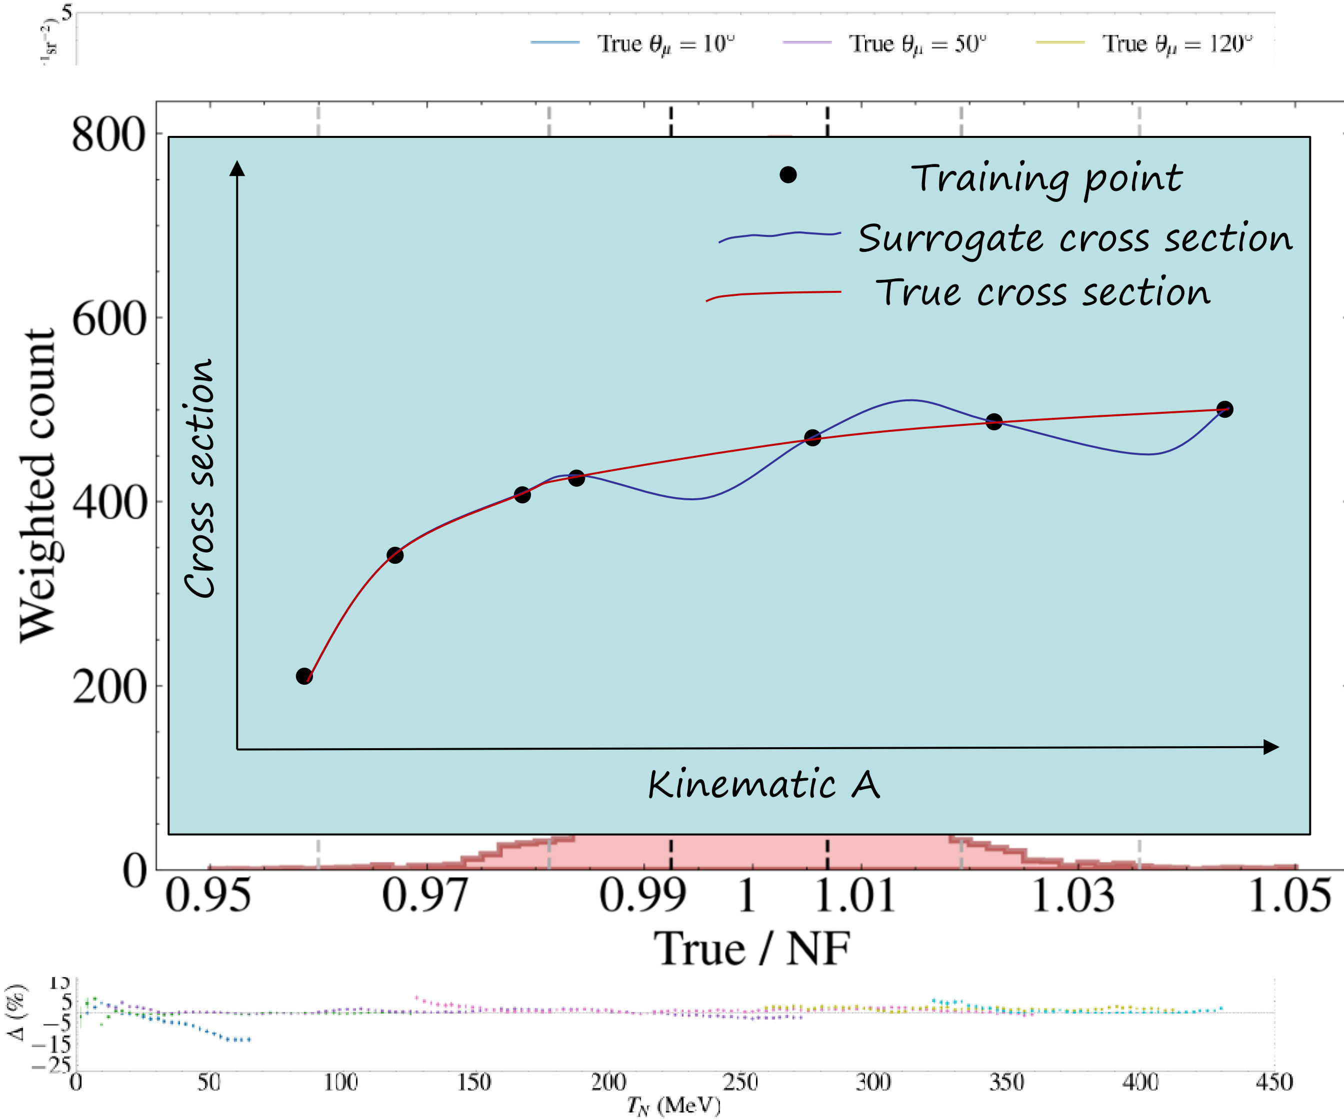

atics +

matics
l local
3)

Comparison with unbiased dataset

- Weight dist. : Point-by-point comparison may be too conservative
- **Binning smooths out local fluctuations**.
- Compare 4D histograms for different energy, shell combinations
  1. Histograms H_{NF}^{MC} from NF datasets (100k samples)
  2. Histograms H_{AR}^{MC} from Accept-reject datasets (100k samples)
  3. High-stat reference histogram H^{ref} (~500 unweighted samples/bin)

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

Multinomial Negative log-likelihood between H^{ref} and H^{MC} (NF or AR)

The MNNL for Accept-Reject datasets is the **optimal performance benchmark**.

Comparison with unbiased dataset

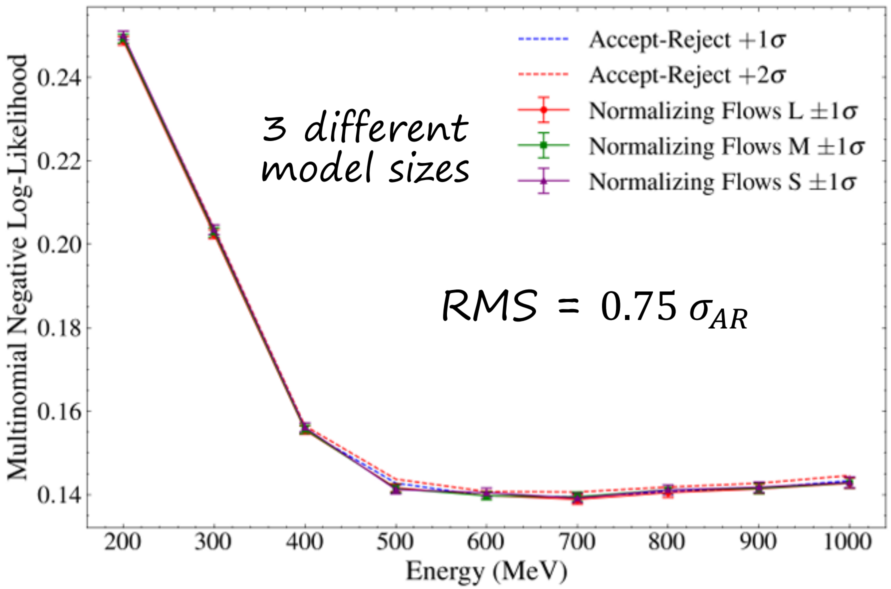

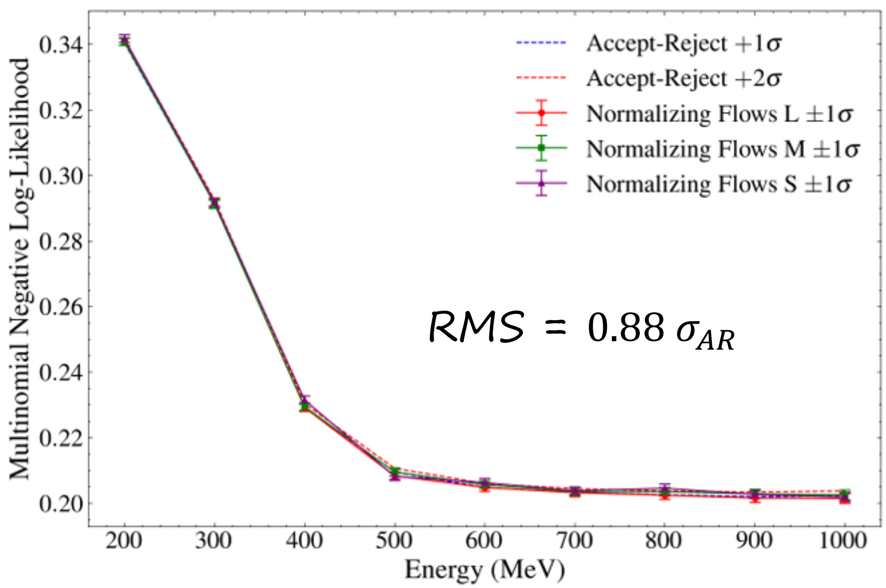

Evolution of the MNLL with the neutrino energy for 1s shell (left) and 1p shell (right)

Recap of performances

Model	S	M	L
Size (MB)	25	90	225
Number of flows	10	10	25
Number of hidden nodes	256	512	512
Training time (hour)	5	6	13
CPU Speed (sec / million samples)	141.8	235.6	669.3
GPU Speed (sec / million samples)	12.5	23.6	70.7
RESS, Flat Flux (%)	98.48	98.64	98.87
RESS, T2K Flux (%)	98.48	98.64	98.85
RMS_z $(1s_{1/2})$	1.36	1.01	0.75
RMS_z $(1p_{3/2})$	1.86	1.52	0.88

Recap of performances

Model	S	M	L
Size (MB)	25	90	225
Number of flows	10	10	25
Number of hidden nodes	256	512	512
Training time (hour)	5	6	13
CPU Speed (sec / million samples)	141.8	235.6	669.3
GPU Speed (sec / million samples)	12.5	23.6	70.7
RESS, Flat Flux (%)	98.48	98.64	98.87
RESS, T2K Flux (%)	98.48	98.64	98.85
RMS_z $(1s_{1/2})$	1.36	1.01	0.75
RMS_z $(1p_{3/2})$	1.86	1.52	0.88

How to use the modeled cross section?

Utilization*	As an importance sampler **	As a sample emulator
Million samples per day per CPU	0.53	129
Million samples per day per GPU	0.53	1222
Quality	Asymptotically unbiased	Asymptotically biased

* Assuming we want to sample the Ghent CCQE cross section
** Weight cap at the 0.9999^{th} quantile

I. Motivations
II. Exclusive cross sections using the HF Mean Field model
III. Efficient event generation with Normalizing Flows
IV. Application to the 1p1h Mean Field exclusive cross section
V. Conclusion and next steps

Summary

- We developed a **Normalizing Flow-based sampling method** for **exclusive neutrino-nucleus cross sections**

⇒ Suited for **theory-driven, high dimensional cross sections**

⇒ **Maximizes unweighting efficiency**

⇒ Serves as an **importance sampler** for unbiased analysis or as a **sample emulator** (though not asymptotically unbiased)

Adding more processes

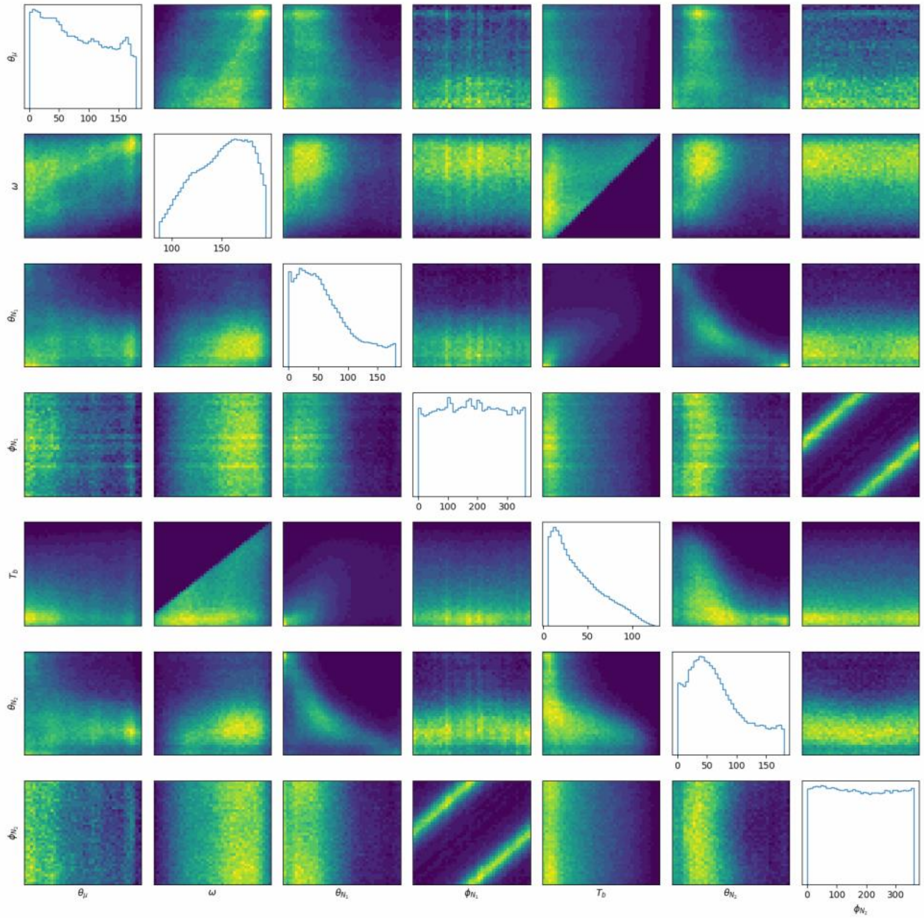

Samples of the trained model for the Mean Field 2p2h exclusive cross section
(pn) configuration with p in 1s shell

For $E_\nu \in [300, 1200]$ MeV

Dependance on the model parametrization

- We already learn energy-dependent NF

$\Rightarrow$ This can be **extended to other conditional parameters** !

$$q_{\phi}(\boldsymbol{x}|E_{\nu}, a_0, a_1, a_2 \dots)$$

- Learn **the dependence of the model on its parametrization**.

$\Rightarrow$ E.g., Axial form factor parameters :

$$F_A(q^2) = \sum_{k=0}^{k_{\max}} a_k z(q^2)^k$$

$\Rightarrow$ Width of the shell …

Conclusion of the conclusion

- A Normalizing Flow-based approach can **make “sampleable” Distorted waves exclusive cross-sections**

$\Rightarrow$ Enables a **unified sampling scheme within a single cross-section model** (e.g., predicted by Mean Field)

$\Rightarrow$ **Next important goal :** Implementation in a MC generator (and benefit from the cascade model), NEUT, ACHILLES… ?

Backup

Plane wave vs Distorted wave

- This neglects final state interactions (interaction with other nucleons).

$\Rightarrow$ Nuclear cascade can be added but no lepton correlation.

- Distorted-wave solutions to a nuclear potential (EDRMF, ROP…) include elastic (real potential) and inelastic (imaginary or cascade) interactions.

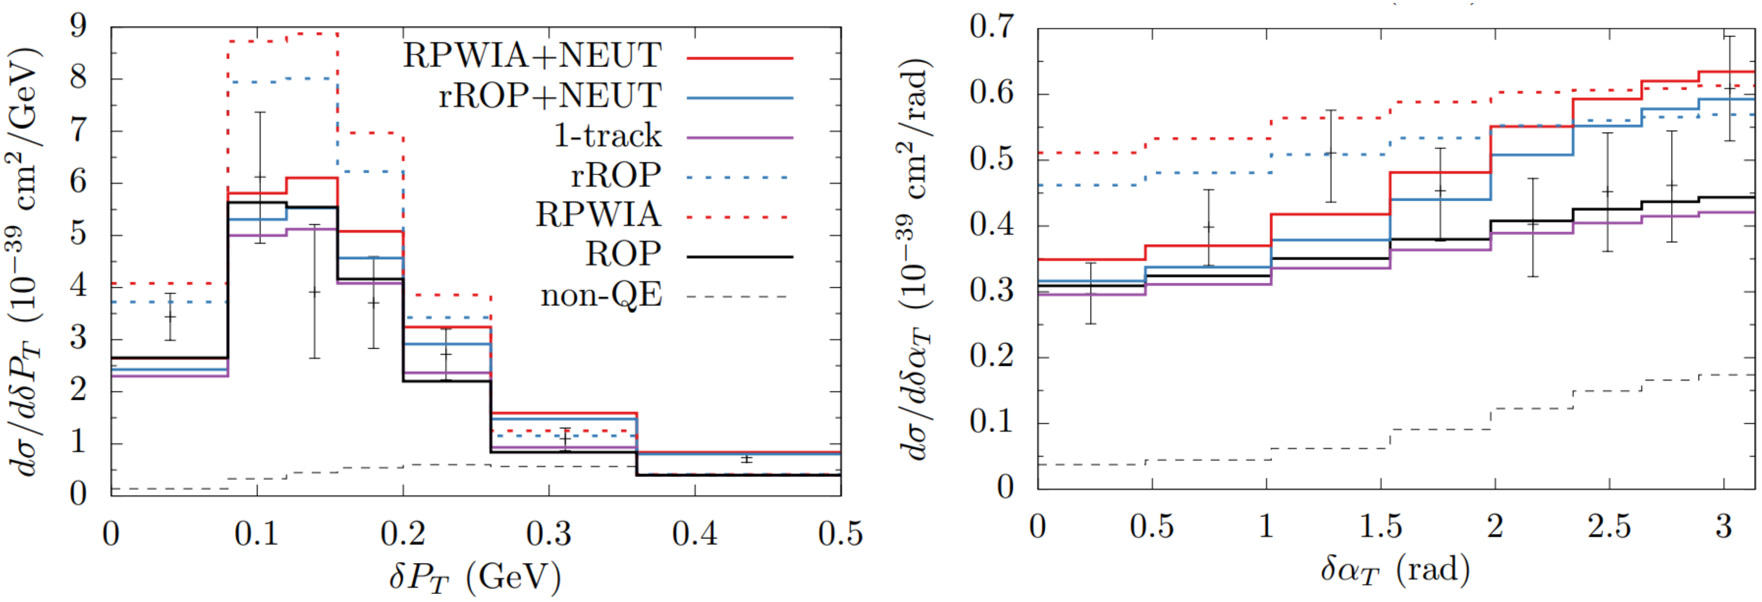

Cross section in terms of the TKI variables along with T2K data (Alexis Nikolakopoulos)