

Lawrence Berkeley National Laboratory

LBL Publications

Title

Improving generative model-based unfolding with Schrödinger bridges

Permalink

<https://escholarship.org/uc/item/9m2580c6>

Journal

Physical Review D, 109(7)

ISSN

2470-0010

Authors

Diefenbacher, Sascha

Liu, Guan-Horng

Mikuni, Vinicius

et al.

Publication Date

2024-04-01

DOI

10.1103/physrevd.109.076011

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Improving generative model-based unfolding with Schrödinger bridges

Sascha Diefenbacher,^{1,*} Guan-Horng Liu,^{2,†} Vinicius Mikuni^{3,‡}, Benjamin Nachman^{1,4,§} and Weili Nie^{5,||}

¹*Physics Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

²*Autonomous Control and Decision Systems Laboratory, Georgia Institute of Technology, Atlanta, Georgia 30332, USA*

³*National Energy Research Scientific Computing Center, Berkeley Lab, Berkeley, California 94720, USA*

⁴*Berkeley Institute for Data Science, University of California, Berkeley, California 94720, USA*

⁵*Machine Learning Research Group, NVIDIA Research*



(Received 22 September 2023; accepted 12 March 2024; published 11 April 2024)

Machine learning-based unfolding has enabled unbinned and high-dimensional differential cross section measurements. Two main approaches have emerged in this research area; one based on discriminative models and one based on generative models. The main advantage of discriminative models is that they learn a small correction to a starting simulation while generative models scale better to regions of phase space with little data. We propose to use Schrödinger bridges and diffusion models to create SBUNFOLD, an unfolding approach that combines the strengths of both discriminative and generative models. The key feature of SBUNFOLD is that its generative model maps one set of events into another without having to go through a known probability density as is the case for normalizing flows and standard diffusion models. We show that SBUNFOLD achieves excellent performance compared to state of the art methods on a synthetic $Z + \text{jets}$ dataset.

DOI: [10.1103/PhysRevD.109.076011](https://doi.org/10.1103/PhysRevD.109.076011)

I. INTRODUCTION

Correcting detector effects in spectra—called deconvolution or unfolding—is the central statistical task in differential cross section measurements in particle, nuclear, and astrophysics. Classical unfolding methods are based on histograms, which result in binned measurements in a small number of dimensions. These methods are usually a variation of regularized matrix inversion, where the (response) matrix describes migrations from particle level to detector level. A number of reviews covering binned unfolding methods can be found in Refs. [1–7] are the most widely used approaches.

Machine learning has the potential to revolutionize differential cross section measurements by enabling unbinned and high-dimensional measurements. A number of machine learning-based unfolding techniques have been proposed [8–22] (see also Ref. [23] for an overview) and

the OMNIFOLD method [12,16] has recently been applied to studies of hadronic final states with data from H1 [24–27], LHCb [28], CMS [29], and STAR [30].

State-of-the-art machine learning-based unfolding methods fall into two categories; discriminative and generative. Discriminative models learn corrections to a starting simulation while generative models enable sampling from the unfolded result. A key benefit of discriminative models is that if the starting simulation is close to the unfolded result, only a small correction must be approximated. In contrast, generative models need to map random numbers, often drawn from a Gaussian distribution, into the data. When the data are far from a standard normal distribution, then generative models have to learn a big correction. A key benefit of generative models is that their performance does not degrade with limited data statistics. In contrast, discriminative models suffer when there is little data for learning density ratios. We propose a new machine learning-based unfolding method that integrates the best of both approaches. Our technique is called SBUNFOLD and is based on generative models, using a technique called a Schrödinger bridge, but only learns a small correction to the starting simulation.

This paper is organized as follows. Section III introduces Schrödinger bridges and Sec. IV describes how this can be used for unfolding. Numerical results from the OMNIFOLD public $Z + \text{jets}$ dataset are presented in Sec. V. The paper ends with conclusions and outlook in Sec. VI.

*sdiefenbacher@lbl.gov

†ghliu@gatech.edu

‡vmikuni@lbl.gov

§bpnachman@lbl.gov

||wnie@nvidia.com

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. 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.

II. UNBINNED UNFOLDING

Let X represent¹ an event at detector-level and Z represent the same event at particle level. The goal of unfolding is to infer the most likely density $p_{\text{data}}(z)$ using simulated pairs (Z, X) and observations x from data. In classical approaches, X and Z are discretized and unfolding proceeds via regularized matrix inversion to approximate the maximum likelihood solution. The likelihood is a product of Poisson probability mass functions, although most measurements do not directly maximize this likelihood. For example, one of the most common classical unfolding algorithms is called Lucy-Richardson deconvolution [31,32] (also known as iterative Bayesian unfolding [33]) which uses an expectation-maximization (EM) algorithm to converge to a maximum likelihood estimator.

In the unbinned case, the likelihood is not known. One solution is to use the EM algorithm, which is at the core of two maximum-likelihood machine learning approaches. The first is OMNIFOLD, starting from events with an initial set of event-level weights $\omega_0(x)$ and $\nu_0(x)$ proceeds as follows²:

$$\begin{aligned} E \text{ step: } \omega_{i+1}(x) &= p_{\text{data}}(x)/\tilde{p}_{\text{sim}}(x) \\ \tilde{p}_{\text{sim}}(x) &\equiv \int dz p_{\text{sim}}(x, z)\nu_i(z); \\ M \text{ step: } \nu_{i+1}(z) &= \tilde{p}_{\text{sim}}(z)/p_{\text{sim}}(z) \\ \tilde{p}_{\text{sim}}(z) &\equiv \int dx p_{\text{sim}}(x, z)\omega_{i+1}(x). \end{aligned}$$

Both the expectation (E) and maximization (M) steps are achieved in practice by training classifiers and interpreting the resulting score as the target likelihood ratio, e.g., take samples from p_{data} and from \tilde{p}_{sim} . (samples from p_{sim} weighted by ν) and train a classifier to distinguish them for the E step. The final result is the set of simulated events weighted by ν with probability density $p_{\text{sim}}(z)\nu(z)$ in the continuum limit. An alternative approach to CINN called ICINN [18] is instead based on generative models:

$$\begin{aligned} E \text{ step: } p_i(z|x) &\propto p_{\text{sim}}(x|z)\nu_i(z); \\ M \text{ step: } \nu_{i+1}(z) &= \check{p}_{\text{sim}}(z)/p_{\text{sim}}(z) \\ \check{p}_{\text{sim}}(z) &\equiv \int dx p_i(z|x)p_{\text{data}}(x). \end{aligned}$$

The E step of the ICINN is achieved by training a generative model (a normalizing flow [34]) to emulate $p_{\text{sim}}(x|z)\nu_i(z)$ while the M step uses a classifier as in OMNIFOLD. The foundation for ICINN was laid in Ref. [14], which focused on the E step, as have other papers using normalizing flows [15], diffusion models [21], and generative adversarial networks [9,13]. Comparing both the E and M steps of OMNIFOLD and the ICINN, one can see that they are formally the same and thus the EM proof in Ref. [12] applies to both

approaches. While formally both algorithms achieve the same EM algorithm step by step, they have complementary strengths in practice. One of OMNIFOLD strengths is that it starts from an existing simulation and if that simulation is close to nature, then the neural networks only have to learn a small correction. In contrast,³ the normalizing flows in ICINN have to map a known probability density (e.g., a multidimensional Gaussian) to the data and these can be quite different. However, the E step of OMNIFOLD trains directly on data and thus its performance degrades when there are not many events. In contrast, the E step of the ICINN is trained using simulation and so can better handle the case of fewer events.

We propose an approach called SBUNFOLD that incorporates strengths of both OMNIFOLD and the ICINN. We employ a technique called a Schrödinger bridge (SB) using a diffusion model to learn the generative model in the ICINN workflow. In contrast to a normalizing flow or standard diffusion model, a Schrödinger bridge learns to map one dataset into another without needing to know the probability density of one of the datasets. Thus, the Schrödinger bridge should ideally learn a small correction while also preserving the E step learning with simulation and not data. The latter property means that SBUNFOLD should outperform OMNIFOLD when there are few events in data. Since OMNIFOLD and the ICINN differ only in the E step, we focus exclusively on the E step of the first iteration for all methods and will refer to ICINN as CINN for disambiguation.

III. SCHRÖDINGER BRIDGE AND THE CONNECTION WITH DIFFUSION MODELS

Diffusion models have become popular choices for generative modeling due to their large degree of flexibility, stable training, and often competitive results compared to other approaches. The core idea is to design a time-dependent perturbation process that slowly perturbs data towards a tractable noise distribution. The goal of the generative model is to then reverse this process, starting from a noise sample and denoising towards new data observations. From a fixed choice of stochastic differential equation (SDE),

$$d\mathbf{x} = f(\mathbf{x}, t)dt + g(t)d\mathbf{w}, \quad (1)$$

described by parameters $f(\mathbf{x}, t) \in \mathbb{R}^d$ and $g(t) \in \mathbb{R}$, the evolution over time of the data observation $\mathbf{x} \in \mathbb{R}^d$ is determined. The same initial data point can undergo different paths due to the additional stochastic term

¹Upper case letters denote random variables and lower case letters correspond to realizations of those random variables. We will use X/x and Z/z to refer to both real and synthetic data; probability densities are labeled as ‘sim.’ or ‘data’ to distinguish the two.

²This integral format is a continuum limit representation. In practice, the integrals are replaced with sums over examples.

³This is true for normalizing flows and traditional diffusion models. Generative Adversarial Networks (GANs) [35] do not require a tractable noise distribution, but are less stable. Reference [17] used a physical latent space for a variational autoencoder (VAE), but VAEs are not state of the art in generative modeling.

identified by the Wiener process, or Brownian motion, $w(t) \in \mathbb{R}^d$, often sampled from a normal distribution with the same dimension as the data. The reverse process follows the reverse SDE equation that reads,

$$d\mathbf{x} = [f(\mathbf{x}, t) - g(t)^2 \nabla \log p(\mathbf{x}, t)]dt + g(t)d\bar{w}. \quad (2)$$

The only unknown term in the Eq. (2) is the score function, $\nabla \log p(\mathbf{x}, t)$, which can be approximated using denoising score matching [36] and minimizing the loss function,

$$\mathcal{L}_{SGM} = \frac{1}{2} \mathbb{E}_{\mathbf{x}_t, t} \|\epsilon_\theta(\mathbf{x}_t, t) - \sigma_t \nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t | \mathbf{x})\|_2^2, \quad (3)$$

where a neural network with trainable parameters θ is optimized to approximate the score function of data that have been perturbed by a Gaussian distribution with time-dependent parameters α_t and σ_t . For this Gaussian perturbation, $q(\mathbf{x}_t | \mathbf{x}) = \mathcal{N}(\mathbf{x}_t, \alpha_t \mathbf{x}, \sigma_t^2 \mathbf{I})$, requiring $f(\mathbf{x}, t)$ to be affine with respect \mathbf{x} and resulting in a Gaussian noise at the end of the diffusion process. For an arbitrary noise distribution, the corresponding choice of $f(\mathbf{x}, t)$ cannot be easily identified, often restricting the family of distributions one is allowed to choose for the diffusion process. A more general framework that allows the mapping between general distributions corresponds to the solution of the SB problem. Initially proposed by Erwin Schrödinger [37], the problem concerned the inference of the trajectories from particles undergoing a diffusion process where experimental observations of the particle's trajectory are available only at specific time values, fixing the boundary conditions of the problem. The connection with diffusion generative models becomes more clear by considering the following forward and backward SDEs,

$$d\mathbf{x} = [f(\mathbf{x}, t) + g(t)^2 \nabla \log \Psi(\mathbf{x}, t)]dt + g(t)dw, \quad (4a)$$

$$d\mathbf{x} = [f(\mathbf{x}, t) - g(t)^2 \nabla \log \hat{\Psi}(\mathbf{x}, t)]dt + g(t)d\bar{w}. \quad (4b)$$

The wave functions Ψ and $\hat{\Psi}$ satisfy

$$\frac{\partial \Psi(\mathbf{x}, t)}{\partial t} = -\nabla \Psi(\mathbf{x}, t) \cdot f(\mathbf{x}, t) - \frac{1}{2} g(t)^2 \Delta \Psi(\mathbf{x}, t),$$

$$\frac{\partial \hat{\Psi}(\mathbf{x}, t)}{\partial t} = -\nabla \cdot (\hat{\Psi}(\mathbf{x}, t) f(\mathbf{x}, t)) + \frac{1}{2} g(t)^2 \Delta \hat{\Psi}(\mathbf{x}, t),$$

with $\Psi(\mathbf{x}, 0) \hat{\Psi}(\mathbf{x}, 0) = p_A(\mathbf{x})$ and $\Psi(\mathbf{x}, 1) \hat{\Psi}(\mathbf{x}, 1) = p_B(\mathbf{x})$ for densities $p_A(\mathbf{x})$ and $p_B(\mathbf{x})$. Compared to Eq. (2), the presence of an additional nonlinear term to the drift function, $g(t)^2 \nabla \log \hat{\Psi}(\mathbf{x}, t)$, enables the diffusion between densities that are not necessarily represented by a standard normal distributions at the end of the diffusion process. Additionally, $\nabla \log \hat{\Psi}(\mathbf{x}, t)$ no longer represents the score function of the perturbed data, but is related to it since

$$\Psi(\mathbf{x}, t) \hat{\Psi}(\mathbf{x}, t) = q(\mathbf{x}, t), \quad (5)$$

hence,

$$\nabla \log \Psi(\mathbf{x}, t) + \nabla \log \hat{\Psi}(\mathbf{x}, t) = \nabla \log q(\mathbf{x}, t). \quad (6)$$

While the equations describing the general SB problem have similarities with the standard framework for diffusion generative models, a general strategy to solve the problem is not immediately obvious. The authors of Ref. [38] propose a tractable solution named I²SB, by considering the presence of pairs of observations in the dataset such that $p(\mathbf{x}_a, \mathbf{x}_b) = p_A(\mathbf{x}_a) p_B(\mathbf{x}_b | \mathbf{x}_a)$. This assumption also holds for the unfolding methodology where pairs of particle collisions before and after detector interactions are always available in the simulation. With this approximation, the authors of Ref. [38] have shown that by setting the linear drift term $f(\mathbf{x}, t) := 0$, the posterior of Eq. (4), $q(\mathbf{x} | \mathbf{x}_a, \mathbf{x}_b)$, has the analytic form,

$$q(\mathbf{x} | \mathbf{x}_a, \mathbf{x}_b) = \mathcal{N}(\mathbf{x}_t; \mu_t(\mathbf{x}_a, \mathbf{x}_b), \Sigma_t), \quad (7)$$

$$\mu_t = \frac{\bar{\sigma}_t^2}{\bar{\sigma}_t^2 + \sigma_t^2} \mathbf{x}_a + \frac{\sigma_t^2}{\bar{\sigma}_t^2 + \sigma_t^2} \mathbf{x}_b, \quad \Sigma_t = \frac{\sigma_t^2 \bar{\sigma}_t^2}{\bar{\sigma}_t^2 + \sigma_t^2} \cdot I, \quad (8)$$

with $\sigma_t^2 := \int_0^t g^2(\tau) d\tau$ and $\bar{\sigma}_t^2 := \int_t^1 g^2(\tau) d\tau$. From this expression, given pairs $(\mathbf{x}_a, \mathbf{x}_b)$, we can directly determine $\mathbf{x}_t = \mu_t + \Sigma_t \epsilon$, $\epsilon \sim \mathcal{N}(0, 1)^d$ for any time step t . The loss function is then identified similar to Eq. (3) as

$$\mathcal{L}_{I^2SB} = \frac{1}{2} \mathbb{E}_{\mathbf{x}_t, t} \left\| \epsilon_\theta(\mathbf{x}_t, t) - \frac{\mathbf{x}_t - \mathbf{x}_0}{\sigma_t} \right\|_2^2, \quad (9)$$

where the right term of the loss function approximates the score function of the backward drift $\nabla \log \hat{\Psi}(\mathbf{x}, t)$, which can then be used during sampling to transport samples from p_B to p_A .

During sampling, standard recursive samplers like DDPM [39] can be used with the prediction of the denoised data $\mathbf{x}_0^\theta = \mathbf{x}_{n+1} - \sigma_{n+1} \epsilon_\theta(\mathbf{x}_{n+1}, n+1)$ at time step $n < N$ defined as

$$\mathbf{x}_n \sim q(\mathbf{x}_n | \mathbf{x}_0^\theta, \mathbf{x}_{n+1}), \quad \mathbf{x}_N \sim p_B. \quad (10)$$

The stochastic solver is able to produce different observations given the same exact inputs, which can then be used for different coverage tests regarding the validity of the generated outputs. A second option is to consider the deterministic case where the posterior distributions are effectively replaced by their means. This particular case is described by the solution of the following ODE:

$$d\mathbf{x}_t = \mathbf{v}_t(\mathbf{x}_t | \mathbf{x}_0) dt = \frac{\beta_t}{\sigma_t^2} (\mathbf{x}_t - \mathbf{x}_0) dt, \quad (11)$$

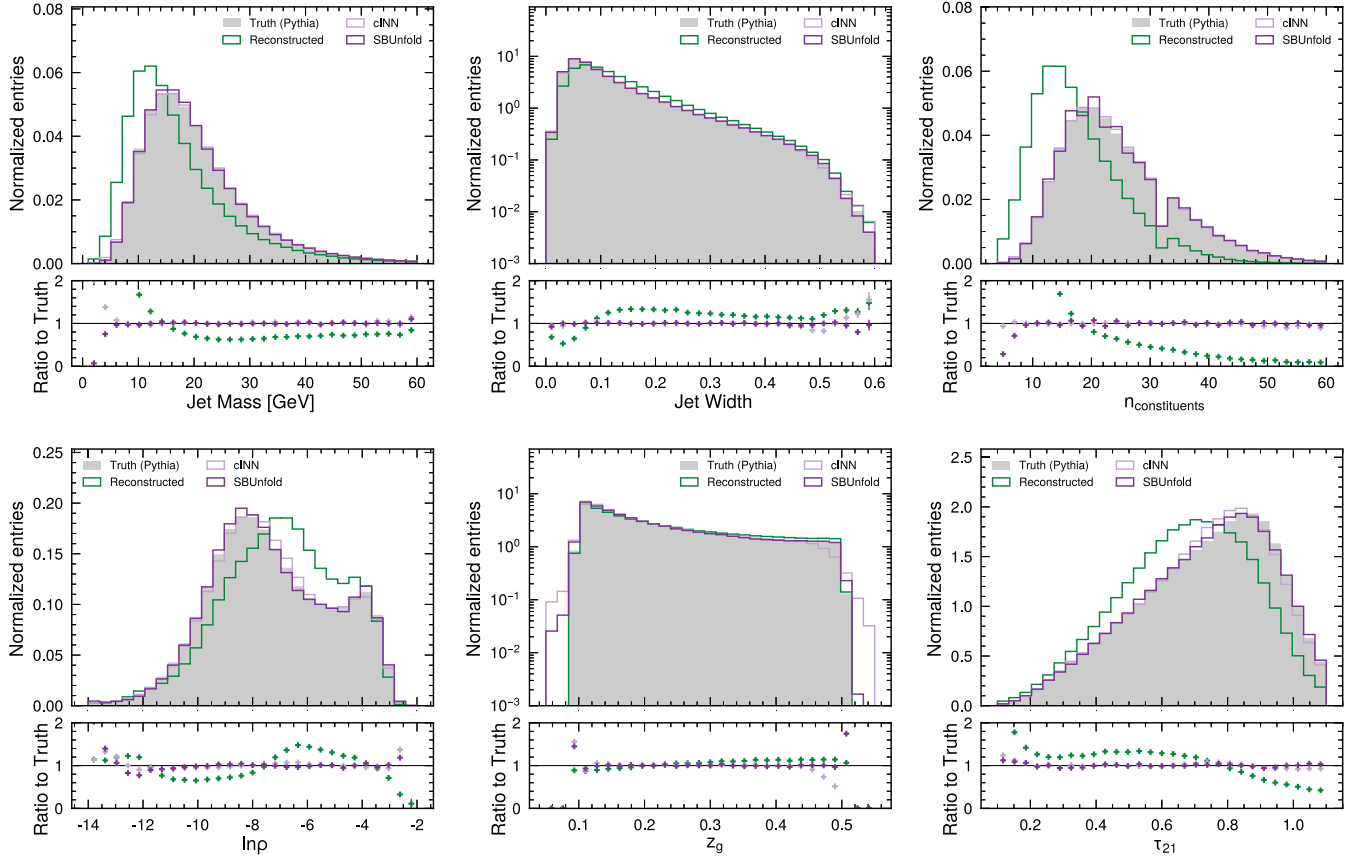


FIG. 1. Comparison between different unfolding algorithms for six different physics observables unfolded. All observables are unfolded simultaneously without binning, with histograms shown only for evaluation. Results are evaluated over 600,000 pseudodata points. Statistical uncertainties are shown only in the ratio panel. Pseudodata and simulation are described by PYTHIA.

which describes an optimal transport plan [40]. In our studies, we observe similar sample quality between the stochastic and deterministic settings and report the results based on the deterministic case for simplicity.

We set $f(\mathbf{x}, t) := 0$ and $g(t) = \sqrt{\beta(t)}$, with $\beta(t)$ the triangular function,

$$\beta(t) = \begin{cases} \beta_0 + 2(\beta_1 - \beta_0)t, & 0 \leq t < \frac{1}{2}, \\ \beta_1 - 2(\beta_1 - \beta_0)(t - \frac{1}{2}), & \frac{1}{2} \leq t \leq 1. \end{cases} \quad (12)$$

with $\beta_0 = 10^{-5}$ and $\beta_1 = 10^{-4}$.

IV. EXPERIMENTAL SETUP

We test SBUNFOLD using the public dataset from Ref. [12], available on Zenodo [41] and briefly summarized in the following. Proton-proton collisions producing a Z boson are generated at a center-of-mass energy of $\sqrt{s} = 14$ TeV. A nontrivial test of unfolding requires at least two datasets, one that acts as the ‘data’ and one that is the ‘simulation’. For the ‘data’, collisions are simulated with the default tune of Herwig 7.1.5 [42–44]. We only make use of the reconstructed events from Herwig as would

be the case with real data. The Herwig particle-level events are only used when evaluating the performance of different methods. For the ‘simulation’, events are simulated with Tune 26 [45] of PYTHIA 8.243 [46–48]. Detector distortions are simulated with DELPHES 3.4.2 [49] and the CMS tune that uses a particle flow reconstruction. In future work, we will investigate the performance of SBUNFOLD on the full phase space of particles (particle level) and particle flow objects (detector level). For this study, we focus on a fixed set of six observables that serve as the benchmark for fixed-dimension unfolding. These observables are computed from the substructure of the leading jet. The jets are clustered using all particle flow objects at detector level and all stable non-neutrino truth particles at particle level. They are defined by the anti- k_T algorithm [50] with radius parameter $R = 0.4$ as implemented in FasJet 3.3.2 [51,52]. The Z bosons are required to have $p_T > 200$ GeV in order to render acceptance effects negligible.⁴ Each dataset consists of about 1.6 million events.

The noise prediction model of SBUNFOLD is implemented using a fully-connected architecture incorporating multiple

⁴See Ref. [16] for how to include acceptance effects.

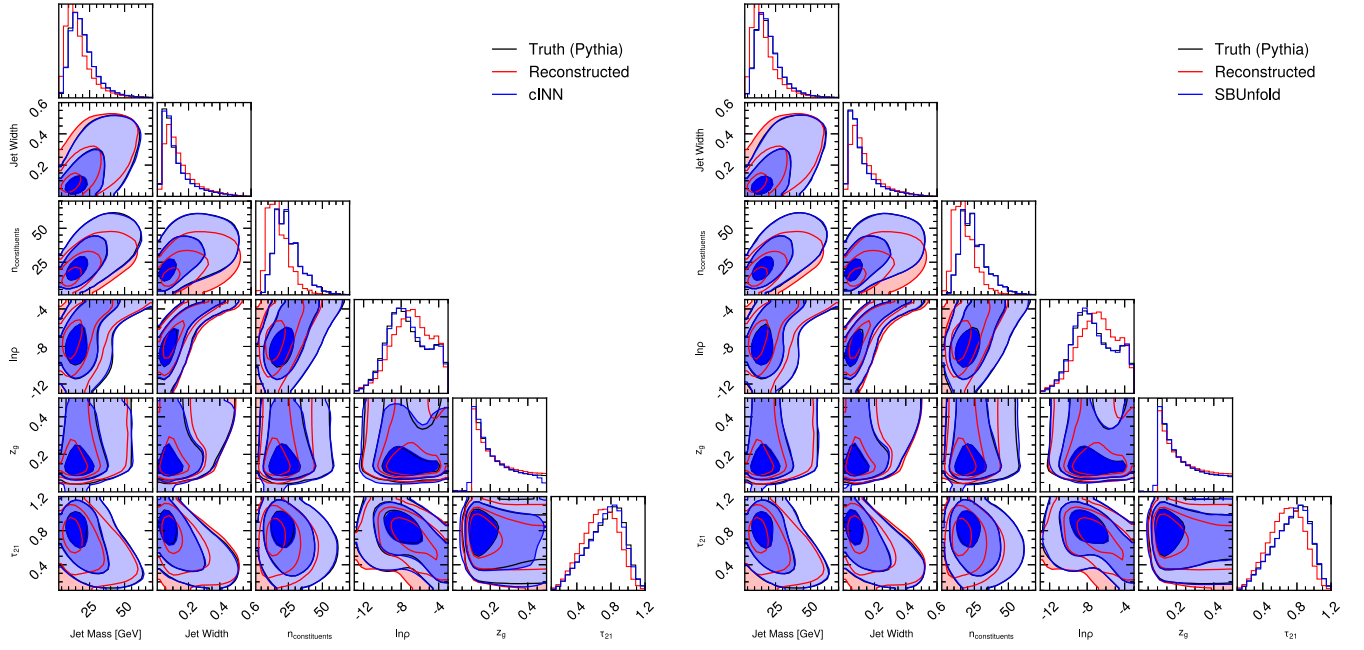


FIG. 2. Pairwise correlation plot between all unfolded variables using the CINN (left) or SBUNFOLD (right) algorithms compared to the expected distributions from PYTHIA events (truth).

skip connections. Specifically, the model employs four RESNET [53] blocks, where each residual layer is connected to the output of a single-layer network through a skip connection. The activation function used is LEAKYRULE [54] with a slope of $\alpha = 0.01$, and all layer sizes are set to 32. The implementation of the model is carried out using PyTorch [55]. The generation is carried out by starting from data observations at reconstruction level and using the SB to move towards the generator level events using 1000 time steps to improve accuracy. The CINN implementation uses the normalizing flow model presented in Ref. [14] while the first step of the first iteration of OMNIFOLD uses a fully-connected model with three hidden layers, RELU activation function, and node sizes set to 50, 150, 50, before the output layer with a single node and SIGMOID activation function. Both CINN and OMNIFOLD are implemented in TENSORFLOW [56]. For all models, the number of trainable parameters are kept similar at around 13 k. The initial learning rate is set to 10^{-3} and all models are trained for 30 epochs with batch size of 128. A total of 1M training events are used to train all models taken from the PYTHIA simulations. The pseudodata used to evaluate the performance of the unfolding is taken from HERWIG. Both SBUNFOLD and CINN do not have access to the data during training while OMNIFOLD uses a total of 2M training events (1M from PYTHIA and 1M from HERWIG). We investigate the impact of the pseudodata training size in Sec. V when determining the unfolded results. All trainings are carried out using the Perlmutter supercomputer with a single NVIDIA A100 GPU.

V. RESULTS

We evaluate the performance of SBUNFOLD compared to other unfolding methodologies by first looking at unfolded distributions where both pseudodata and samples used to

TABLE I. Comparison of the earth mover's distance (EMD) and triangular discriminator between different algorithms. EMD is calculated over unbinned distributions while triangular discriminator uses histograms as inputs. Uncertainties from EMD are derived using 100 bootstraps with replacement taken from the unfolded data. Uncertainties from triangular discriminator are taken from 100 histogram variations within the statistical uncertainty of the prediction. Results are evaluated using 600,000 pseudodata points sampled from PYTHIA. Quantities in bold represent the method with best performance.

Model	EMD($\times 10$)/Triangular discriminator($\times 10^3$)	
	CINN	SBUNFOLD
Jet mass	1.4 ± 0.2	0.70 ± 0.06
	0.27 ± 0.05	0.28 ± 0.04
Jet width	0.013 ± 0.002	0.0029 ± 0.0005
	0.22 ± 0.05	0.04 ± 0.01
N	2.3 ± 0.8	0.57 ± 0.04
	0.08 ± 0.03	0.84 ± 0.22
log ρ	1.1 ± 0.3	0.27 ± 0.01
	0.60 ± 0.18	0.61 ± 0.12
z_g	0.095 ± 0.003	0.009 ± 0.001
	7.5 ± 1.3	1.7 ± 0.3
τ_{21}	0.2 ± 0.1	0.016 ± 0.001
	0.60 ± 0.20	0.20 ± 0.06

TABLE II. Comparison of the earth mover’s distance (EMD) and triangular discriminator between different algorithms. EMD is calculated over unbinned distributions while triangular discriminator uses histograms as inputs. Uncertainties from EMD are derived using 100 bootstraps with replacement taken from the unfolded data. Uncertainties from triangular discriminator are taken from 100 histogram variations within the statistical uncertainty of the prediction. Results are evaluated using 600,000 pseudodata points sampled from Herwig. Quantities in bold represent the method with best performance.

Model	EMD($\times 10$) and Triangular discriminator($\times 10^3$)		
	OMNIFOLD Step 1	CINN	SBUNFOLD
Jet mass	6.1 ± 0.1	10.1 ± 1.2	9.0 ± 0.1
	1.5 ± 0.5	2.4 ± 0.7	3.2 ± 0.9
Jet width	0.06 ± 0.001	0.05 ± 0.003	0.02 ± 0.001
	1.2 ± 0.3	0.7 ± 0.2	0.17 ± 0.04
N	1.7 ± 0.1	6.1 ± 4.0	3.0 ± 0.1
	0.12 ± 0.03	0.12 ± 0.03	0.9 ± 0.2
$\log \rho$	1.35 ± 0.03	3.1 ± 2.1	0.4 ± 0.1
	1.1 ± 0.4	0.9 ± 0.2	0.8 ± 0.2
z_g	0.086 ± 0.001	0.3 ± 0.1	0.049 ± 0.001
	1.2 ± 0.3	9.4 ± 2.0	2.2 ± 0.4
τ_{21}	0.23 ± 0.02	0.7 ± 0.4	0.12 ± 0.02
	4.6 ± 1.4	3.5 ± 1.1	1.4 ± 0.45

train the different unfolding methods originate from the PYTHIA simulation. The results showing the unfolded distributions are shown in Fig. 1. In Appendix A, we also provide the comparison of the results obtained by a standard diffusion model. We investigate the same set of observables used in previous studies in this dataset. The first four observables are the jet mass m , constituent multiplicity M , the N -subjettiness ratio $\tau_{21} = \tau_2^{(\beta=1)}/\tau_1^{(\beta=1)}$ [57,58], and the jet width w [implemented as $\tau_1^{(\beta=1)}$]. The remaining two observables are the jet mass $\ln \rho = \ln m_{SD}^2/p_T^2$ and momentum fraction z_g after soft drop grooming [59,60] with $z_{\text{cut}} = 0.1$ and $\beta = 0$. Many of the observables are computed with FASTJET CONTRIB 1.042 [61].

We observe a good agreement with the expected distributions from PYTHIA at generator level and the responses of the generative models for unfolding. We omit the OMNIFOLD results since the weighting function becomes trivial when both data and simulation are statistically identical. We also investigate the distribution of the pairwise correlation between each pair of unfolded features to verify that SBUNFOLD is also capable of learning the correct correlations between distributions. The results are shown in Fig. 2.

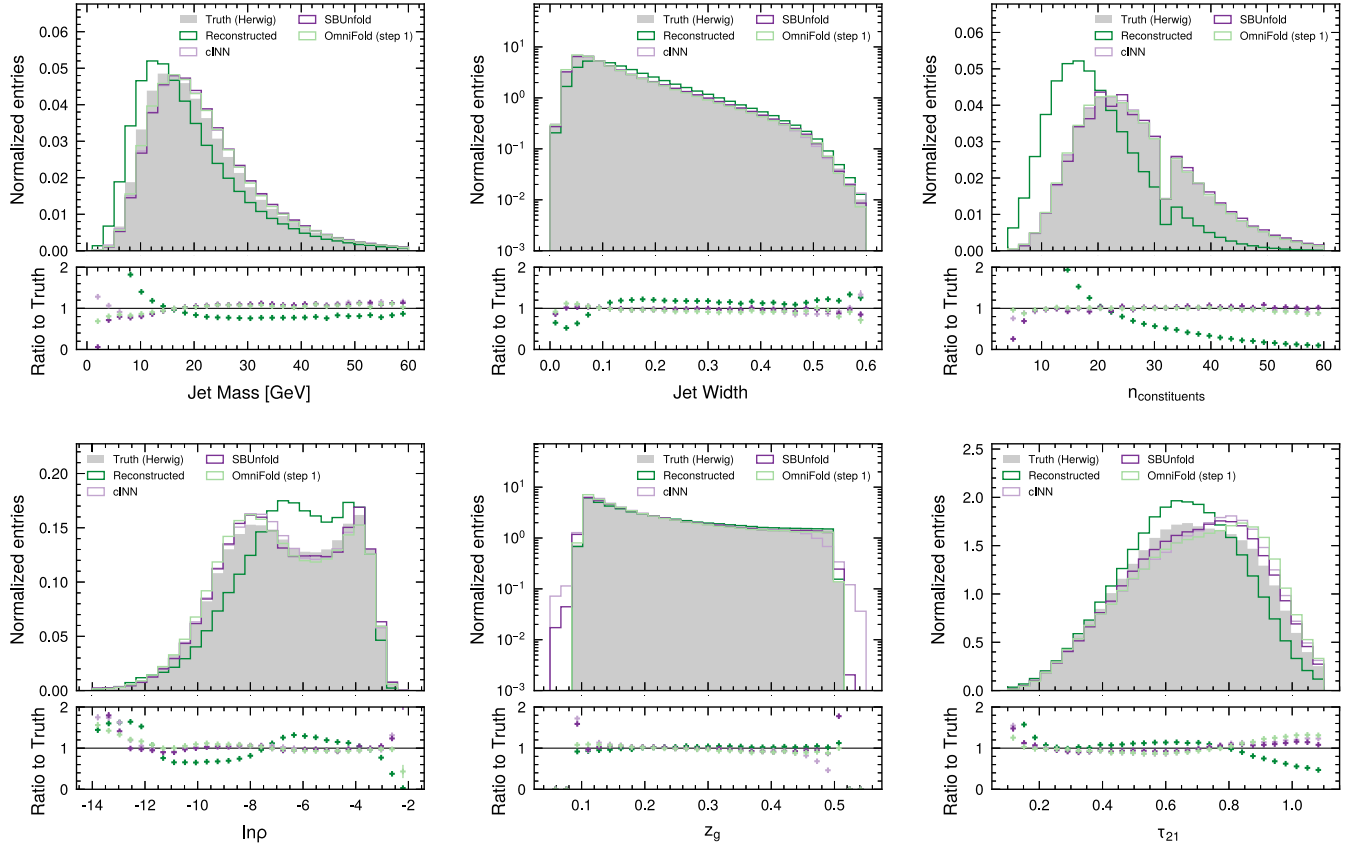


FIG. 3. Comparison between different unfolding algorithms for six different physics observables unfolded. All observables are unfolded simultaneously without binning, with histograms shown only for evaluation. Results are evaluated over 600,000 pseudodata points. Statistical uncertainties are shown only in the ratio panel. Pseudodata is taken from Herwig while PYTHIA is taken as the main simulator.

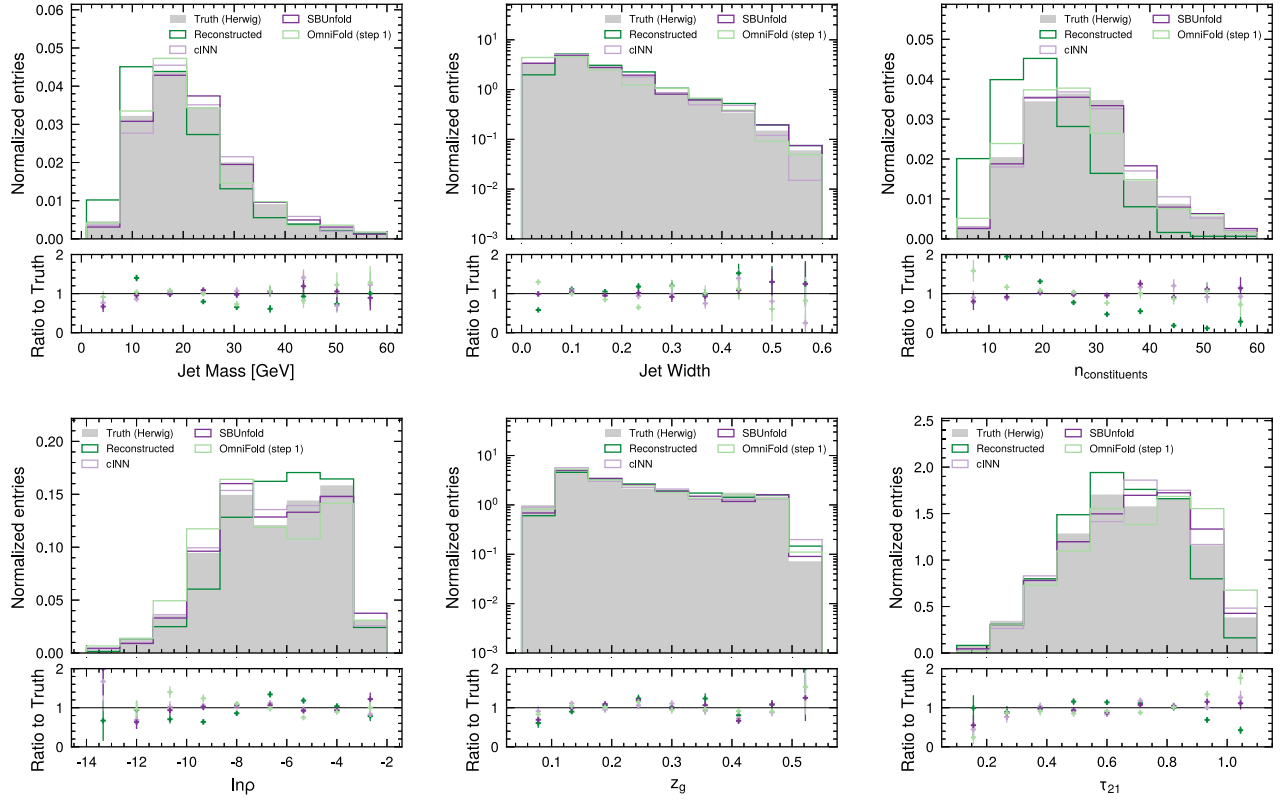


FIG. 4. Comparison between different unfolding algorithms for six different physics observables unfolded. All observables are unfolded simultaneously without binning, with histograms shown only for evaluation. Results are evaluated over 1,000 pseudodata points, the same number was used to train OMNIFOLD while other algorithms only rely on simulation. Statistical uncertainties are shown only in the ratio panel. Pseudodata is taken from Herwig while PYTHIA is taken as the main simulator.

We again observe a good agreement between the generated distributions and expected distributions from the simulation for both generative models, with SBUNFOLD showing improved description compared to the cINN for a few distributions such as z_g vs $\log \rho$ and z_g vs τ_{21} . The distribution of z_g shows a sharp cutoff at $z_g = 0$ which is hard to reproduce with generative models. Since the same sharp distribution is observed at reconstructed level events, SBUNFOLD can take advantage of the reconstructed prior when starting the diffusion process, contrary to the cINN that always requires samples from the standard normal distribution for the prior transformation. We also quantify the agreement between unfolded distributions with expected values derived from PYTHIA by calculating both the triangular discriminator metric [62–64] over histograms⁵ as presented in Fig. 1, as well as using the Earth mover’s distance (EMD) directly on the marginalized

⁵The triangular discriminator d is calculated for each bin i as $d = \frac{1}{2} \sum w_i \frac{(x_i - y_i)^2}{x_i + y_i}$, for histograms containing entries x_i and y_i counts and bin width w_i .

one-dimensional features generated by the unfolding methods. The results are listed in Table I.

For all distributions we observe significantly lower EMD values for SBUNFOLD compared to cINN with four out of six also showing lower values of the triangular discriminator. Uncertainties from the EMD calculation are derived from the standard deviation of 100 bootstraps with replacement taken from the unfolded data. We observe the EMD uncertainties from cINN to be often higher than SBUNFOLD, which points to improved stability from SBUNFOLD due the more informative prior.

Next, we keep the same unfolding methodology from SBUNFOLD and ICINN trained over PYTHIA samples but instead use reconstructed events from HERWIG as the data representative. We compare the unfolded results between the different distributions in Fig. 3.

We observe improved agreement in all unfolded distributions compared to the distributions of reconstructed events, while systematic shifts are observed in all unfolding methodologies for the jet mass and τ_{21} distributions. These features highlight the prior dependence of all models which cannot be excluded without the presence of the M step and

additional iterations. In Table II we calculate the EMD and triangular discriminator using the unfolded distributions with HERWIG as pseudodata.

Once again we observe a good performance of SBUNFOLD, achieving the lowest EMD values for four out of six distributions and lowest triangular discriminator values for three of the six. We also observe OMNIFOLD achieving similar performance for all observables, in particular OMNIFOLD shows improved performance for the particle multiplicity, which is not a continuous distribution and hence harder for generative models to determine precisely. We also study how SBUNFOLD corrects physics observables back to generator level events by calculating the migration matrix between reconstructed and unfolded observables. Results showing that SBUNFOLD learns to apply small corrections to reconstructed events are shown in Appendix B.

OMNIFOLD is also the only algorithm that uses the statistical power of the data to determine the unfolded distributions. We investigate how the unfolding results change when the available number of data entries is reduced to only 1,000 instead of the 600,000 thousands we used previously. For both SBUNFOLD and CINN this change only modifies the number of generated samples from the trained model while OMNIFOLD only has access to the 1,000 data examples during training while the number of PYTHIA samples is not changed and kept at 1,000,000. The unfolded results are shown in Fig. 4.

TABLE III. Comparison of the earth mover’s distance (EMD) and triangular discriminator between different algorithms. EMD is calculated over unbinned distributions while triangular discriminator uses histograms as inputs. Uncertainties from EMD are derived using 100 bootstraps with replacement taken from the unfolded data. Uncertainties from triangular discriminator are taken from 100 histogram variations within the statistical uncertainty of the prediction. Results are evaluated using 1,000 pseudodata points sampled from Herwig. Quantities in bold represent the method with best performance.

Model	EMD($\times 10$)/Triangular discriminator($\times 10^3$)		
	OMNIFOLD Step 1	CINN	SBUNFOLD
Jet mass	8.7 ± 1.8	9.2 ± 3.0	7.7 ± 2.5
	4.0 ± 2.5	3.6 ± 1.8	1.9 ± 1.2
Jet width	0.14 ± 0.02	0.07 ± 0.02	0.05 ± 0.02
	11 ± 4	3.2 ± 1.7	0.7 ± 0.9
N	12 ± 3	5.4 ± 1.3	5.8 ± 1.6
	6.7 ± 3.2	2.0 ± 1.4	2.2 ± 1.6
$\log \rho$	4.0 ± 0.8	1.6 ± 0.5	1.2 ± 0.3
	9.2 ± 4.0	2.1 ± 1.4	2.5 ± 1.5
z_g	0.08 ± 0.02	0.08 ± 0.03	0.06 ± 0.01
	1.9 ± 1.4	5.7 ± 2.5	6.6 ± 2.8
τ_{21}	0.4 ± 0.07	0.2 ± 0.05	0.1 ± 0.04
	12 ± 4	5.4 ± 2.6	2.8 ± 1.7

Due to the lower number of data entries we observe larger statistical fluctuations from the unfolded results which still shows a good agreement compared to the expected quantities. We evaluate further the unfolded results by calculating the EMD and triangular discriminator, listed in Table III.

The reported results for SBUNFOLD and CINN are less affected by the reduced data sizes, even though the uncertainties from EMD have greatly increased due to limited pseudodata examples. On the other hand, OMNIFOLD shows a worse degradation compared to previous results, with EMD values often disagreeing by more than two standard deviations.

VI. CONCLUSION AND OUTLOOK

In this paper we presented SBUNFOLD, an unfolding algorithm that uses a Schrödinger bridge (SB) to design a stochastic mapping between distributions. Since the SB allows the transport between arbitrary distributions,⁶ SBUNFOLD does not require a tractable prior similar to other generative models, but instead uses a diffusion process to directly denoise physics observables from detector effects. We have demonstrated the performance of SBUNFOLD using synthetic $Z + \text{jets}$ samples and compared with both state-of-the-art methods using generative models (CINN) and reweighting (OMNIFOLD). Using different metrics, we observe an excellent performance of SBUNFOLD. Compared to CINN, SBUNFOLD shows improved fidelity for distributions with sharp features and smaller uncertainties overall likely due to more informative priors from reconstructed level events. Compared to OMNIFOLD, SBUNFOLD shows a more robust performance against variations of the number of available data observations, showing promising results for cases where the amount of available data is reduced.

We have only investigated the expectation step of different unfolding algorithms, leaving the maximization step for future study. We notice that the maximization step based on classifiers used by ICINN is also applicable for SBUNFOLD, and would be interesting to compare how the unfolding results improve with subsequent iterations between both algorithms to reduce the dependence on the initial simulation parameters.

Finally, similar to OMNIFOLD, SBUNFOLD can be readily adapted to different data structures such as imagelike datasets. On the other hand, the diffusion process currently relies on the specific ordering of the features used during

⁶Our Schrödinger bridge is not a universal function approximator; it will map the distributions to each other, but may not preserve the conditional probability density [38]. Our numerical results indicate that this is not an issue and we expect that monotonic detector distortions should be in the class of functions that SBs are able to accommodate, but we leave further investigations to future work. We thank Jesse Thaler for useful discussions on this point.

unfolding [see Eq. (8)], making the application of SBUNFOLD with low-level features, such as particles, not trivial. Following studies will investigate how to accommodate permutation invariance within SBUNFOLD.

The code for this paper can be found at <https://github.com/ViniciusMikuni/SBUnfold>.

ACKNOWLEDGMENTS

We thank Daniel Whiteson, Tilman Plehn, and Jesse Thaler for thoughtful discussions and feedback on the manuscript. V. M., S. D., and B. N. are supported by the U.S. Department of Energy (DOE), Office of Science under Contract No. DE-AC02-05CH11231. G. H. L. would like to thank Evangelos Theodorou for helpful discussions. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231 using NERSC Award No. HEP-ERCAP0021099.

APPENDIX A: COMPARISON WITH STANDARD DIFFUSION MODELS

We investigate the benefits of using SBUNFOLD as opposed to a standard diffusion model. The diffusion model implementation follows the FPCD model proposed in [65] for jet kinematic generation. The backbone neural network architecture used is the same as the one used in SBUNFOLD. Similarly, the learning rate, number of training epochs, and training samples are kept the same between the two models. The main conceptual difference between the implementations is that SBUNFOLD starts from the reconstruction level events instead of a Gaussian prior used in FPCD. Similarly to the CINN implementation, the FPCD model is only conditioned on reconstruction level events to determine the unfolded response. Results of the unfolded distributions using PYTHIA both as pseudodata and simulation are shown in Fig. 5.

For all distributions we observe a good agreement between the unfolded results and the pseudodata. In Table IV we calculate the EMD and triangular discriminator using

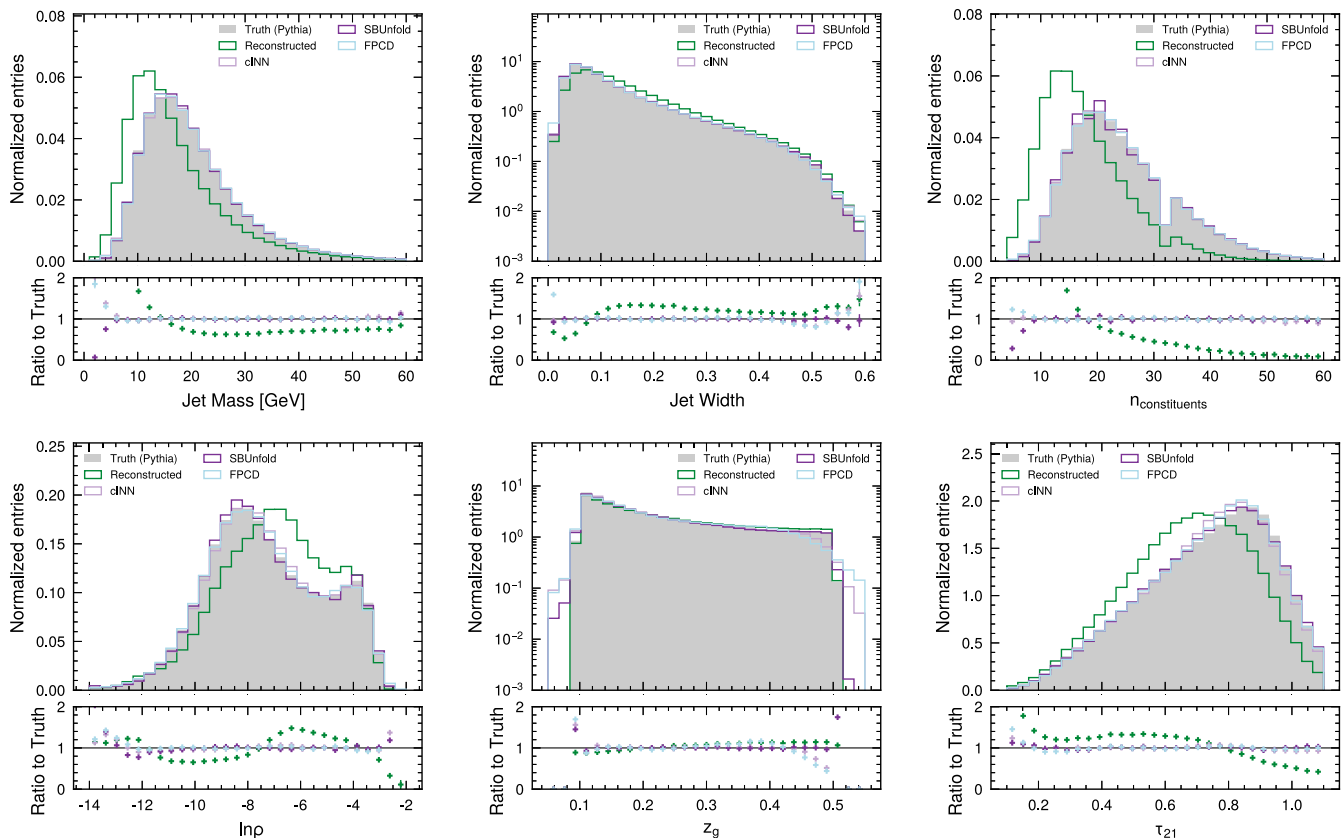


FIG. 5. Comparison between different unfolding algorithms for six different physics observables unfolded. Results are evaluated over 600,000 pseudodata points. Statistical uncertainties are shown only in the ratio panel. Pseudodata and simulation are described by PYTHIA.

TABLE IV. Comparison of the earth mover’s distance (EMD) and triangular discriminator between different unfolding methodologies. EMD is calculated over unbinned distributions while triangular discriminator uses histograms as inputs. Uncertainties from EMD are derived using 100 bootstraps with replacement taken from the unfolded data. Uncertainties from triangular discriminator are taken from 100 histogram variations within the statistical uncertainty of the prediction. Results are evaluated using 600,000 pseudodata points sampled from PYTHIA. Quantities in bold represent the method with best performance.

Model	EMD($\times 10$) and Triangular discriminator($\times 10^3$)		
	FPCD	CINN	SBUNFOLD
Jet mass	0.74 ± 0.08	1.4 ± 0.2	0.70 ± 0.06
	0.19 ± 0.04	0.27 ± 0.05	0.28 ± 0.04
Jet width	0.0087 ± 0.0006	0.013 ± 0.002	0.0029 ± 0.0005
	0.88 ± 0.08	0.22 ± 0.05	0.04 ± 0.01
N	0.81 ± 0.06	2.3 ± 0.8	0.57 ± 0.04
	0.14 ± 0.04	0.08 ± 0.03	0.84 ± 0.22
$\log \rho$	0.34 ± 0.01	1.1 ± 0.3	0.27 ± 0.01
	0.77 ± 0.16	0.60 ± 0.18	0.61 ± 0.12
z_g	0.035 ± 0.007	0.095 ± 0.003	0.009 ± 0.001
	12.4 ± 2.4	7.5 ± 1.3	1.7 ± 0.3
τ_{21}	0.024 ± 0.002	0.2 ± 0.1	0.016 ± 0.001
	0.32 ± 0.09	0.60 ± 0.20	0.20 ± 0.06

the unfolded distributions obtained by the different methodologies.

From the results we observe that FPCD often outperforms the CINN in the EMD metric while still not achieving the same level of performance as SBUNFOLD. These results indicate that standard diffusion models may be more expressive compared to the normalizing flow model implemented in the original CINN, but still not as precise as SBUNFOLD which leverages the more informative prior distribution from reconstruction level objects rather than the Gaussian prior required by both CINN and FPCD.

APPENDIX B: MIGRATION MATRIX OF UNFOLDED EVENTS

In this appendix we investigate how SBUNFOLD maps reconstructed level distributions back to generator level samples by calculating the corresponding migration matrix using Herwig as the pseudodata while the detector response is derived from the PYTHIA simulation. The distributions are shown in Fig. 6.

The migration matrix for all observables has a high fraction of events present close to the diagonal, consistent with SBUNFOLD learning to apply only a small correction to the inputs.




FIG. 6. Migration distribution between reconstructed level events and generator level distribution after unfolding with SBUNFOLD. Results are evaluated over 600,000 pseudodata points. Pseudodata is taken from Herwig while PYTHIA is taken as the main simulator.

- [1] G. Cowan, *Conf. Proc. C* **0203181**, 248 (2002).
- [2] V. Blobel, *PHYSTAT2011 Proceedings* (CERN, 2011), 240, [10.5170/CERN-2011-006](https://arxiv.org/abs/10.5170/CERN-2011-006).
- [3] V. Blobel, *Data Analysis in High Energy Physics* (Wiley, New York, 2013), 187.
- [4] R. Balasubramanian, L. Brenner, C. Burgard, G. Cowan, V. Croft, W. Verkerke, and P. Verschuur, *Int. J. Mod. Phys. A* **35**, 2050145 (2020).
- [5] G. D'Agostini, *Nucl. Instrum. Methods Phys. Res., Sect. A* **362**, 487 (1995).
- [6] A. Hocker and V. Kartvelishvili, *Nucl. Instrum. Methods Phys. Res., Sect. A* **372**, 469 (1996).
- [7] S. Schmitt, *J. Instrum.* **7**, T10003 (2012).
- [8] A. Glazov, [arXiv:1712.01814](https://arxiv.org/abs/1712.01814).
- [9] K. Datta, D. Kar, and D. Roy, [arXiv:1806.00433](https://arxiv.org/abs/1806.00433).
- [10] M. Bunse, N. Piatkowski, T. Ruhe, W. Rhode, and K. Morik, in *5th International Conference on Data Science and Advanced Analytics (DSAA)* (IEEE, Turin, Italy, 2018), pp. 21–30.
- [11] T. Ruhe, T. Voigt, M. Wornowizki, M. Börner, W. Rhode, and K. Morik, Mining for spectra—the Dortmund spectrum estimation algorithm, in *Astronomical Data Analysis Software and Systems XXVI* edited by M. Molinaro, K. Shortridge, and F. Pasian Astronomical Society of the Pacific Conference Series, (2019) p. 394.
- [12] A. Andreassen, P. T. Komiske, E. M. Metodiev, B. Nachman, and J. Thaler, *Phys. Rev. Lett.* **124**, 182001 (2020).
- [13] M. Bellagente, A. Butter, G. Kasieczka, T. Plehn, and R. Winterhalder, *SciPost Phys.* **8**, 070 (2020).
- [14] M. Bellagente, A. Butter, G. Kasieczka, T. Plehn, A. Rousselot, and R. Winterhalder, *SciPost Phys.* **9**, 074 (2020).
- [15] M. Vandegar, M. Kagan, A. Wehenkel, and G. Louppe, [arXiv:2011.05836](https://arxiv.org/abs/2011.05836).
- [16] A. Andreassen, P. T. Komiske, E. M. Metodiev, B. Nachman, A. Suresh, and J. Thaler, [arXiv:2105.04448](https://arxiv.org/abs/2105.04448).
- [17] J. N. Howard, S. Mandt, D. Whiteson, and Y. Yang, *Sci. Rep.* **12**, 7567 (2022).
- [18] M. Backes, A. Butter, M. Dunford, and B. Malaescu, [arXiv:2212.08674](https://arxiv.org/abs/2212.08674).
- [19] M. Arratia, D. Britzger, O. Long, and B. Nachman, *J. Instrum.* **17**, P07009 (2022).
- [20] J. Chan and B. Nachman, *Phys. Rev. D* **108**, 016002 (2023).
- [21] A. Shmakov, K. Greif, M. Fenton, A. Ghosh, P. Baldi, and D. Whiteson, [arXiv:2305.10399](https://arxiv.org/abs/2305.10399).
- [22] T. Alghamdi *et al.*, *Phys. Rev. D* **108**, 094030 (2023).
- [23] M. Arratia *et al.*, *J. Instrum.* **17**, P01024 (2022).
- [24] V. Andreev *et al.* (H1 Collaboration), *Phys. Rev. Lett.* **128**, 132002 (2022).
- [25] H1 Collaboration, H1prelim-22-031 (2022), https://www-h1.desy.de/publications/H1preliminary.short_list.html.
- [26] V. Andreev *et al.* (H1 Collaboration), *Phys. Lett. B* **844**, 138101 (2023).
- [27] H1 Collaboration, H1prelim-23-031 (2023), https://www-h1.desy.de/publications/H1preliminary.short_list.html.
- [28] LHCb Collaboration, *Phys. Rev. D* **108**, L031103 (2023).
- [29] P. T. Komiske, S. Kryhin, and J. Thaler, *Phys. Rev. D* **106**, 094021 (2022).
- [30] Y. Song (STAR Collaboration), [arXiv:2307.07718](https://arxiv.org/abs/2307.07718).
- [31] L. B. Lucy, *Astron. J.* **79**, 745 (1974).
- [32] W. H. Richardson, *J. Opt. Soc. Am.* **62**, 55 (1972).
- [33] G. D'Agostini, *Nucl. Instrum. Methods Phys. Res., Sect. A* **362**, 487 (1995).
- [34] D. J. Rezende and S. Mohamed, *Proceedings of Machine Learning Research* **37** 1530 (2016).
- [35] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, [arXiv:1406.2661](https://arxiv.org/abs/1406.2661).
- [36] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, [arXiv:2011.13456](https://arxiv.org/abs/2011.13456).
- [37] E. Schrödinger, *Sitzungsber. Preuss. Akad. Wissen. Phys. Math. Klasse Sonderausgabe* **9**, 144 (1931).
- [38] G.-H. Liu, A. Vahdat, D.-A. Huang, E. A. Theodorou, W. Nie, and A. Anandkumar, [arXiv:2302.05872](https://arxiv.org/abs/2302.05872).
- [39] J. Ho, A. Jain, and P. Abbeel, in *Advances in Neural Information Processing Systems*, edited by H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin (Curran Associates, Inc., San Diego, CA, 2020), Vol. 33, pp. 6840–6851.
- [40] G. Peyré, M. Cuturi *et al.*, *Found. Trends Mach. Learn.* **11**, 355 (2019).
- [41] A. Andreassen, P. Komiske, E. Metodiev, B. Nachman, and J. Thaler, PYTHIA/Herwig + DELPHES jet datasets for OMNIFOLD unfolding, Zenodo, [10.5281/zenodo.3548091](https://zenodo.org/record/105281) (2019).
- [42] M. Bahr *et al.*, *Eur. Phys. J. C* **58**, 639 (2008).
- [43] J. Bellm *et al.*, *Eur. Phys. J. C* **76**, 196 (2016).
- [44] J. Bellm *et al.*, [arXiv:1705.06919](https://arxiv.org/abs/1705.06919).
- [45] The ATLAS Collaboration, ATLAS Run 1 Pythia8 tunes to 7 TeV data, Technical Report No. ATL-PHYS-PUB-2014-021, CERN, Geneva, 2014.
- [46] T. Sjöstrand, S. Mrenna, and P. Z. Skands, *Comput. Phys. Commun.* **178**, 852 (2008).
- [47] T. Sjöstrand, S. Mrenna, and P. Z. Skands, *J. High Energy Phys.* **05** (2006) 026.
- [48] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, *Comput. Phys. Commun.* **191**, 159 (2015).
- [49] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, and M. Selvaggi (DELPHES 3 Collaboration), *J. High Energy Phys.* **02** (2014) 057.
- [50] M. Cacciari, G. P. Salam, and G. Soyez, *J. High Energy Phys.* **04** (2008) 063.
- [51] M. Cacciari, G. P. Salam, and G. Soyez, *Eur. Phys. J. C* **72**, 1896 (2012).
- [52] M. Cacciari and G. P. Salam, *Phys. Lett. B* **641**, 57 (2006).
- [53] K. He, X. Zhang, S. Ren, and J. Sun, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (IEEE, San Juan, PR, 2016), pp. 770–778.
- [54] B. Xu, N. Wang, T. Chen, and M. Li, [arXiv:1505.00853](https://arxiv.org/abs/1505.00853).
- [55] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, *Adv. Neural Inf. Process. Syst.* **32** (2019).
- [56] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, in *OSDI* (2016), Vol. **16**, pp. 265–283, <https://www.tensorflow.org/about/bib>.

- [57] J. Thaler and K. Van Tilburg, *J. High Energy Phys.* **03** (2011) 015.
- [58] J. Thaler and K. Van Tilburg, *J. High Energy Phys.* **02** (2012) 093.
- [59] A. J. Larkoski, S. Marzani, G. Soyez, and J. Thaler, *J. High Energy Phys.* **05** (2014) 146.
- [60] M. Dasgupta, A. Fregoso, S. Marzani, and G. P. Salam, *J. High Energy Phys.* **09** (2013) 029.
- [61] Fastjet contrib, <https://fastjet.hepforge.org/contrib/>.
- [62] F. Topsoe, *IEEE Trans. Inf. Theory* **46**, 1602 (2000).
- [63] P. Gras, S. Höche, D. Kar, A. Larkoski, L. Lönnblad, S. Plätzer, A. Siódmok, P. Skands, G. Soyez, and J. Thaler, *J. High Energy Phys.* **07** (2017) 091.
- [64] S. Bright-Thonney and B. Nachman, *J. High Energy Phys.* **03** (2019) 098.
- [65] V. Mikuni, B. Nachman, and M. Pettee, *Phys. Rev. D* **108**, 036025 (2023).