

Towards Universal Unfolding of Detector Effects in High-Energy Physics using Denoising Diffusion Probabilistic Models

Camila Pazos¹, Shuchin Aeron^{2,4}, Pierre-Hugues Beauchemin^{1,4}, Vincent Croft³,
Zhengyan Huan^{2,4} Martin Klassen¹ and Taritree Wongjirad^{1,4}

¹ Department of Physics and Astronomy, Tufts University, Medford, Massachusetts

² Department of Electrical and Computer Engineering, Tufts University, Medford,
Massachusetts

³ Leiden Institute for Advanced Computer Science LIACS, Leiden University, The Netherlands

⁴ The NSF AI Institute for Artificial Intelligence and Fundamental Interactions

camila.pazos@tufts.edu, shuchin@eecs.tufts.edu, hugo.beauchemin@tufts.edu,
vincent.croft@cern.ch, zhengyan.huan@tufts.edu, martin.klassen@tufts.edu,
taritree.wongjirad@tufts.edu,

Abstract

Correcting for detector effects in experimental data, particularly through unfolding, is critical for enabling precision measurements in high-energy physics. However, traditional unfolding methods face challenges in scalability, flexibility, and dependence on simulations. We introduce a novel approach to multidimensional object-wise unfolding using conditional Denoising Diffusion Probabilistic Models (cDDPM). Our method utilizes the cDDPM for a non-iterative, flexible posterior sampling approach, incorporating distribution moments as conditioning information, which exhibits a strong inductive bias that allows it to generalize to unseen physics processes without explicitly assuming the underlying distribution. Our results highlight the potential of this method as a step towards a “universal” unfolding tool that reduces dependence on truth-level assumptions, while enabling the unfolding of a wide range of measured distributions with improved adaptability and accuracy.

Copyright attribution to authors.

This work is a submission to SciPost Physics.

License information to appear upon publication.

Publication information to appear upon publication.

Received Date

Accepted Date

Published Date

1

2 Contents

3	1 Introduction	2
4	2 Methods	4
5	2.1 Our Unfolding Approach	4
6	2.2 Denoising Diffusion Probabilistic Models	5
7	2.3 Unfolding with cDDPMs: Toy Model	7
8	3 Application to Particle Physics Data	11
9	3.1 Setup and Implementation	11
10	3.2 Performance Evaluation	12
11	3.3 Computational Performance	15

12	4 Conclusion	15
13	A cDDPM Details	19
14	A.1 Loss Derivation	19
15	A.2 Model Parameters	21
16	B Datasets	22
17	B.1 Data Processing for cDDPMs	22
18	B.2 Physics Generation	23
19	B.3 Detector Smearing and Jet Matching	23
20	References	25
21	<hr/>	
22		

23 1 Introduction

24 Experimental data in high-energy physics (HEP) presents a distorted picture of the underlying
 25 physics processes due to detector effects. Unfolding is an inverse problem solved through sta-
 26 tistical techniques that aims to correct the detector distortions of the observed data to make
 27 inferences about the true underlying distribution of particle properties. This process is essen-
 28 tial for the validation of theories, constraining new physics models using experimental data,
 29 precision measurements, and comparison of experimental results between different experi-
 30 ments.

31 A standard approach to unfolding [1] begins with a predicted particle distribution $f_{\text{true}}(\mathbf{x})$,
 32 where \mathbf{x} represents the true particle-level kinematic properties, that characterizes the underly-
 33 ing physics process of interest, and a detailed detector simulation that describes how detector
 34 effects distort the particle property distributions. These distortions affect the reconstruction of
 35 kinematic quantities of particles incident to the detector, resulting in an altered observed par-
 36 ticle distribution $f_{\text{det}}(\mathbf{y})$, where \mathbf{y} represents the reconstructed detector-level kinematic prop-
 37 erties. This relationship between an observed distribution and its underlying kinematic prop-
 38 erties can be written as a Fredholm integral equation of the first kind,

$$f_{\text{det}}(\mathbf{y}) = \int d\mathbf{x} P(\mathbf{y}|\mathbf{x}) f_{\text{true}}(\mathbf{x}) \quad (1)$$

39 where $P(\mathbf{y}|\mathbf{x})$ is the conditional probability distribution describing the detector effects. Un-
 40 folding requires the inverse process $P(\mathbf{x}|\mathbf{y})$, which can be expressed with Bayes' theorem, as

$$P(\mathbf{x}|\mathbf{y}) = \frac{P(\mathbf{y}|\mathbf{x})f_{\text{true}}(\mathbf{x})}{f_{\text{det}}(\mathbf{y})}. \quad (2)$$

41 In this context, a detector dataset can be unfolded by sampling from the posterior $P(\mathbf{x}|\mathbf{y})$
 42 to recover the distribution $f_{\text{true}}(\mathbf{x})$. The detector effects $P(\mathbf{y}|\mathbf{x})$ are assumed to be the same for
 43 any physics process, and it is clear that the posterior $P(\mathbf{x}|\mathbf{y})$ depends on the prior distribution
 44 $f_{\text{true}}(\mathbf{x})$. Although one can sample from $f_{\text{true}}(\mathbf{x})$ through the use of particle generators, there is
 45 no guarantee that any particular assumed $f_{\text{true}}(\mathbf{x})$ accurately represents the underlying physics
 46 of the specific data to unfold. Consequently, unfolding results can be significantly influenced
 47 by the assumed underlying distribution, potentially introducing bias or limiting the method's

48 ability to detect unexpected phenomena - a challenge commonly known as the bias problem
49 in unfolding. While many traditional unfolding methods [2–4] attempt to address this bias
50 problem, they face several inherent limitations: they require binned histograms, cannot un-
51 fold multiple observables simultaneously, retain a potential residual bias towards an assumed
52 underlying distribution (introduced as a necessary trade-off to reduce large variances in the
53 unfolded distribution), and depend heavily on specific choices made by experimenters during
54 the data analysis process. These challenges reveal the difficulty in developing a *universal* un-
55 folder, which aims to remove detector effects from any set of measured data agnostic of the
56 process of interest with no bias towards any prior distribution and without constraints on the
57 final interpretation. This task of creating a widely applicable unfolding method is known as
58 the generalization problem.

59 **Related Work:** Various machine learning approaches have emerged in recent years to ad-
60 dress these challenges. These include re-weighting methods like OmniFold [5] [6], as well
61 as several generative approaches. Among the generative techniques are those using genera-
62 tive adversarial networks (GANs) [7,8], conditional invertible neural networks (cINN) [9–11],
63 and variational latent diffusion models (VLD) [12] [13]. Additionally, distribution mapping
64 techniques have been developed such as SBUnfold, which utilizes Schrödinger bridges [14],
65 and DiDi, a direct diffusion model [15]. For a comprehensive overview of these methods, see
66 the recent survey by [16]. Each new method has made further strides in unfolding and shown
67 the advantages in machine learning based approaches compared to traditional techniques. Ta-
68 ble 1 provides a comparative summary of these machine learning-based methods, traditional
69 techniques, and our proposed approach, highlighting some key characteristics and advantages.
70 While this overview is not exhaustive of all unfolding algorithms developed to date, it provides
71 a comprehensive portrait that allows situating our proposed approach within the landscape of
72 existing solutions to the unfolding problem.

73 **Objectives:** Our work seeks to overcome the limitations of traditional unfolding methods
74 while expanding upon the benefits offered by machine learning-based approaches. The pro-
75 posed approach builds upon the advantages of object-wise unfolding, a technique common in
76 machine learning-based unfolding methods, which reconstructs the properties of individual
77 particles or physics objects rather than operating on binned distributions. Through object-
78 wise unfolding, some of the challenges posed by traditional methods can be addressed: the
79 impact of the experimenter’s selections and cuts on the unfolded results can be minimized,
80 while underlying correlations between the unfolded distributions are preserved.

81 We first present a “dedicated” unfolders, an approach similar to many existing machine
82 learning-based methods, which learns and applies a specific posterior distribution for a par-
83 ticular physics process. This approach serves as an effective solution for well-understood pro-
84 cesses and provides a benchmark for our subsequent work. Building upon this foundation,
85 our aim is to develop a “generalizable” unfolders to handle a wide range of physics processes
86 and observables, including those not explicitly seen during training. This generalization ca-
87 pability is crucial for enhancing the method’s applicability across various physics scenarios,
88 while ideally avoiding dependence on specific physics generator models and. This amounts
89 to addressing both the bias and generalization problems in our solution to unfolding. Such a
90 method would enable the unfolding of distributions for a wide range of processes, including
91 those involving yet-undiscovered particles in new physics searches at high-energy colliders.

92 An effective new unfolding method should achieve an accuracy that falls within the typical
93 uncertainty range of measurements where unfolding is applied. For instance, the ATLAS collab-
94 oration’s [17] measurement of the W +jets differential cross-sections [18] obtained from data
95 resulting from proton-proton collisions at the Large Hadron Collider [19] provides a bench-

96 mark for the necessary level of precision. These results [20] demonstrate that a 10-15% total
 97 uncertainty is typical for energy-momentum related quantities, with approximately 3-5% at-
 98 tributable to unfolding. The goal is to achieve this level of accuracy while simultaneously
 99 preserving the benefits of object-wise unfolding, such as maintaining correlations between
 100 kinematic quantities, and offering generalization capabilities. With these objectives, we hope
 101 to contribute a more flexible, accurate, and widely applicable unfolding tool to the high-energy
 102 physics community.

103 **Our Contribution:** This work introduces a novel approach using conditional Denoising Dif-
 104 fusion Probabilistic Models (cDDPM) to unfold detector effects in HEP data. We demonstrate
 105 that a single cDDPM, trained on diverse particle data and incorporating statistical moments of
 106 various distributions, can serve as a “generalized” unfolders by performing multidimensional
 107 object-wise unfolding for multiple physics processes without explicit assumptions about the
 108 underlying distribution, thereby minimizing bias. Figure 1 illustrates the effectiveness of this
 109 approach in two key scenarios. Panel (a) shows an “unknown” process created by combining
 110 data from multiple known processes (40% $t\bar{t}$, 35% W +jets, and 25% leptoquark). Here, the
 111 generalized unfolders outperforms a “dedicated” unfolders, which is designed to unfold only a
 112 single specific physics process (in this case $t\bar{t}$, selected because it forms the largest component
 113 of the unknown process). Panel (b) provides further evidence of the generalized unfolders’s
 114 flexibility, demonstrating its ability to accurately unfold data from graviton production (gener-
 115 ated in the context of large extra-dimension scenarios [21]) accompanied by jets, a completely
 116 new physics process absent from the training phase. In both cases, the generalized unfolders
 117 achieves accuracy within typical LHC uncertainty budgets. The accuracy of the generalized
 118 approach illustrates its ability to handle previously unseen physics processes without assum-
 119 ing an underlying distribution. This flexibility demonstrated by the generalized unfolders is
 120 beneficial for new physics searches and studying processes not accurately modeled by cur-
 121 rent theories, providing an unfolding solution to the bulk of the data analyses performed at
 122 high-energy colliders. Section 3 provides a detailed analysis of these results.

123 2 Methods

124 2.1 Our Unfolding Approach

125 We seek an approach that will enhance the inductive bias of the unfolding method to improve
 126 generalization to cover various posteriors pertaining to different physics data distributions,
 127 while avoiding systematically favoring any particular prior distribution. From the a priori
 128 information used in the formulation of the solution to the unfolding problem (Equation 2), it
 129 can be seen that the posteriors for two different physics processes i and j , where the detector
 130 effects are independent of the process, are related by a ratio of the probability density functions
 131 of each process,

$$\frac{P_i(\mathbf{x}|\mathbf{y})}{P_j(\mathbf{x}|\mathbf{y})} = \frac{f_{\text{true}}^i(\mathbf{x}) f_{\text{det}}^j(\mathbf{y})}{f_{\text{det}}^i(\mathbf{y}) f_{\text{true}}^j(\mathbf{x})}. \quad (3)$$

132 This relationship indicates that if a posterior for a given physics process can be learned,
 133 then distributional information about $f_{\text{true}}(\mathbf{x})$ and $f_{\text{det}}(\mathbf{y})$ could be used to estimate unseen
 134 posteriors. Specifically, the first moments of these distributions can be utilized as key features.
 135 This approach of using summary statistics like moments is particularly advantageous, as it
 136 allows for generalization without being overly sensitive to details and insignificant fluctuations

Table 1: Comparison of unfolding techniques and their key characteristics. The table presents traditional and machine learning-based approaches, with our proposed methods (dedicated and generalizable cDDPM) highlighted in green and orange, respectively. The “Type” column indicates the fundamental algorithmic structure of each method. “Posterior Estimation” describes whether the solution is obtained iteratively, non-iteratively, or partially, where partial refers to methods that estimate only certain components of the posterior. “Event-wise” indicates methods that unfold individual particles or objects without binning, allowing event-level information to be reconstructed from the unfolded results. “Tuneable Regularization” indicates whether the method implements adjustable bias-variance trade-offs in its unfolding solution. “Generalizable” indicates whether the method is designed to estimate a wide range of unseen posteriors. While not exhaustive, these characteristics provide a framework for comparing different approaches to the unfolding problem.

	Method	Type	Posterior Estimation	Event-Wise? [†]	Multi-dimensional?	Tuneable Regularization?	Generalizable? ^{*,*}
Traditional	IBU [2]	Bin-by-bin correction	Iterative	No	Limited	Yes	No
	SVD Unfolding [4]	Matrix inversion	Partial	No	Limited	Yes	No
	TUnfold [3]	Matrix inversion	Non-iterative	No	No	Yes	No
ML-Based	OmniFold [5, 6]	Re-weighting	Iterative	Yes	Yes	Yes	NA
	GANs [7, 8]	Generative	Non-iterative	Yes	Yes	No	No
	cINN [9–11]	Generative	Non-iterative	Yes	Yes	No	No
	VLD [12, 13]	Generative	Non-iterative	Yes	Yes	No	No
	SBUnfold [14]	Distribution mapping	Non-iterative	Yes	Yes	No	No
	DiDi [15]	Distribution mapping	Non-iterative	Yes	Yes	No	No
	Dedicated cDDPM	Generative	Non-iterative	Yes	Yes	No	No
	Generalizable cDDPM	Generative	Non-iterative	Yes	Yes	No	Yes

137 in the distributions. By making use of the first moments of the detector data distribution as
 138 conditionals, a more flexible folder can be created that is not strictly tied to a selected prior
 139 distribution, and enables interpolation and extrapolation to unseen posteriors based on the
 140 provided moments. Consequently, this unfolding tool gains the ability to handle a wider range
 141 of physics processes and enhances the generalization capabilities, making it a more versatile
 142 tool for unfolding in various high energy physics applications. It would therefore provide an
 143 unfolding solution addressing both the bias and generalization problems.

144 2.2 Denoising Diffusion Probabilistic Models

145 In learning systems, the challenge of generalization through inductive bias is central, as any
 146 system must have some bias beyond the training instances to make the inductive leap necessary
 147 to classify unseen cases [22]. Our proposed unfolding approach calls for a flexible generative
 148 model to address this challenge, and denoising diffusion probabilistic models (DDPMs) [23]
 149 lend themselves naturally to this task. DDPMs are designed to create new content based on
 150 training data, making them well-suited for these generalization needs. DDPMs can be trained
 151 to model a data distribution through a reversible generative process, which can be condi-
 152 tioned directly on the detector data values and on the moments of the distribution $f_{\text{det}}(\mathbf{y})$.
 153 This learned process provides a natural way to sample from $P(\mathbf{x}|\mathbf{y})$ for unfolding. We will
 154 first describe DDPMs in a general context before discussing their specific application to our
 155 unfolding problem.

156 **Unconditional DDPM** The standard unconditional DDPM [23] consists of two parts. First
 157 is a forward process (or diffusion process) $q(\mathbf{x}_t|\mathbf{x}_{t-1})$ which is fixed to a Markov chain that
 158 gradually adds Gaussian noise (following a variance schedule β_1, \dots, β_T) to data samples from

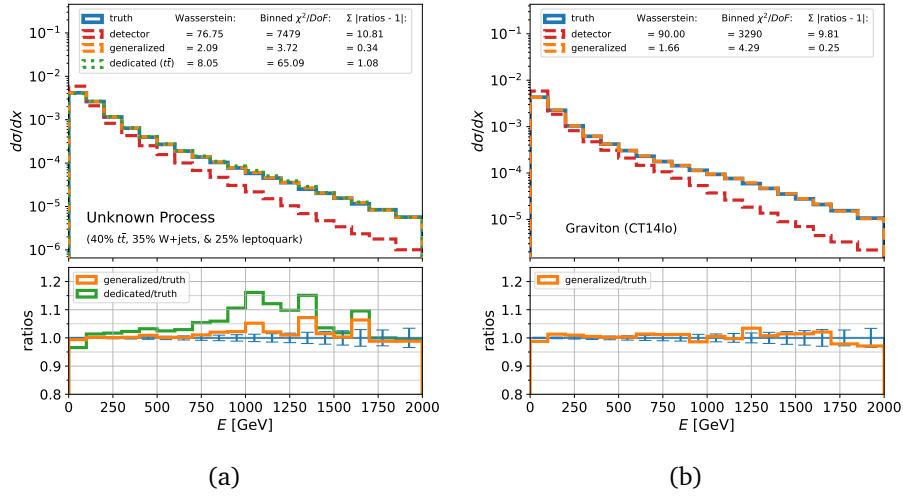


Figure 1: Unfolding results from the data-driven detector smearing using the generalized cDDPM unfolders. Panel (a) shows performance on data from an "unknown" physics process combining multiple processes. The generalized unfolded (orange) demonstrates superior performance compared to the dedicated unfolded (green), which was trained assuming a specific physics process. Panel (b) shows the generalized unfolded successfully handling data from graviton production accompanied by jets, a new physics process completely absent from the training data. The accuracy of the generalized approach in both scenarios illustrates its ability to handle previously unseen physics processes without assuming an underlying distribution.

159 a known initial distribution,

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) := \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I}). \quad (4)$$

160 Second is a learned reverse process (or denoising process) $p_\theta(\mathbf{x}_{0:T})$ parameterized by
 161 θ . The reverse process is also a Markov chain with learned Gaussian transitions starting at
 162 $p(\mathbf{x}_T) = \mathcal{N}(\mathbf{x}_T; \mathbf{0}, \mathbf{I})$,

$$p_\theta(\mathbf{x}_{0:T}) := p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) \quad (5)$$

$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) := \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_\theta(t, \mathbf{x}_t), \sigma_t^2 \mathbf{I}). \quad (6)$$

163 By learning to reverse the forward diffusion process, the model learns meaningful latent
 164 representations of the underlying data and is able to remove noise from data to generate new
 165 samples from the associated data distribution. This type of generative model has natural ap-
 166 plications in high energy physics, for example generating data samples from known particle
 167 distributions. However, to be used in unfolding the process must be altered so that the de-
 168 noising procedure is dependent on the observed detector data, \mathbf{y} . This dependence on the
 169 observed data is crucial because the goal of unfolding is to reconstruct the true particle-level
 170 properties from the observed detector-level data, necessitating a direct link between the de-
 171 noising process and the specific detector measurements. This can be achieved by incorporating
 172 conditioning methods to the DDPM.

173 **Conditional DDPM** Conditioning methods for DDPMs can either use conditions to guide
 174 unconditional DDPMs in the reverse process [24], or they can incorporate direct conditions
 175 to the learned reverse process. While guided diffusion methods have had great success in
 176 image synthesis [25], direct conditioning provides a framework that is particularly useful in
 177 unfolding since it allows for a more explicit and precise incorporation of the detector-level
 178 data into the unfolding process, enabling the model to learn a direct mapping between the
 179 observed detector measurements and the true particle properties.

180 We implement a conditional DDPM (cDDPM) for unfolding that keeps the original uncon-
 181 ditional forward process and introduces a simple, direct conditioning on the input \mathbf{y} to the
 182 reverse process,

$$p_{\theta}(\mathbf{x}_{0:T}|\mathbf{y}) := p(\mathbf{x}_T|\mathbf{y}) \prod_{t=1}^T p_{\theta}(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{y}). \quad (7)$$

183 Similar to an unconditional DDPM, this reverse process has the same functional form as the
 184 forward process and can be expressed as a Gaussian transition with a learned mean $\boldsymbol{\mu}_{\theta}$ and a
 185 fixed variance at each timestep σ_t^2 ,

$$p_{\theta}(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{y}) := \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_{\theta}(t, \mathbf{x}_t, \mathbf{y}), \sigma_t^2 \mathbf{I}). \quad (8)$$

186 This conditioned reverse process learns to model the posterior probability $P(\mathbf{x}|\mathbf{y})$ through its
 187 Gaussian transitions. The reverse process $p_{\theta}(\mathbf{x}_{0:T}|\mathbf{y})$ is parameterized by θ , where θ represents
 188 the learnable parameters of the model, such as the weights and biases of a neural network.
 189 This process learns to remove the noise introduced during the forward process to recover the
 190 target \mathbf{x} by conditioning directly on the input \mathbf{y} .

191 Training optimizes the parameters θ to maximize the likelihood of accurately estimating
 192 the noise ϵ that should be removed at each timestep in order to denoise \mathbf{x}_t given the con-
 193 dition \mathbf{y} . Similar to the unconditional DDPM, the Gaussian nature of these transitions and a
 194 reparametrization of the mean can be used to simplify the loss function to the mean squared
 195 error (MSE) between the noise ϵ added at each timestep during the forward process and the
 196 noise $\epsilon_{\theta}(\mathbf{x}_t, \mathbf{y})$ predicted by the model given the noisy sample \mathbf{x}_t at timestep t and the condition
 197 \mathbf{y} :

$$L(\theta) = \mathbb{E}_{t, \epsilon, \mathbf{x}_t, \mathbf{y}} \left[\left\| \epsilon - \epsilon_{\theta}(t, \mathbf{x}_t, \mathbf{y}) \right\|^2 \right]. \quad (9)$$

198 A detailed derivation of this loss can be found in A.1. This approach can be compared to
 199 the commonly used guided conditioning method, where the model estimates the noise with
 200 a weighted combination of the conditional prediction and the unconditional prediction as
 201 $\tilde{\epsilon}_{\theta}(\mathbf{x}_t, \mathbf{y}) = (1 + w) \epsilon_{\theta}(\mathbf{x}_t, \mathbf{y}) - w \epsilon_{\theta}(\mathbf{x}_t)$ [26]. In the guided approach, the use of the uncondi-
 202 tional prediction $\epsilon_{\theta}(\mathbf{x}_t)$ introduces a bias towards the underlying distribution used in training.
 203 This bias is undesirable for the unfolding task, as it is important to minimize any assumptions
 204 about or dependence on the underlying distribution. The cDDPM approach mitigates this bias
 205 by setting the guidance weight $w = 0$, relying solely on the conditional prediction. In this case,
 206 sampling would be done purely according to the learned conditional distribution $p_{\theta}(\mathbf{x}_t|\mathbf{y})$. Al-
 207 though the learned conditional probability implicitly depends on the prior, sampling from the
 208 cDDPM does not require explicitly evaluating the prior distribution $p(\mathbf{x}_t)$ over the data space.
 209 This makes the cDDPM a promising choice for applications like unfolding where the prior is
 210 unknown or difficult to model.

211 2.3 Unfolding with cDDPMs: Toy Model

212 Proof-of-concept is first demonstrated using a toy model with non-physics data. The method
 213 will then be tested with generated physics data, the results of which are discussed in Section

214 3.

215 In both the toy model and physics results, the Wasserstein-1 distance [27] is used to mea-
 216 sure the success of the proposed unfolding algorithm. This metric quantifies the discrepancy
 217 between two distributions, quantifying how closely the unfolded distribution matches the tar-
 218 get (truth) distribution, as well as the difference between the detector data and the true under-
 219 lying distribution. The objects to unfold are characterized by multiple kinematic properties,
 220 necessitating a multidimensional representation of each object. Unlike traditional methods
 221 that unfold each quantity separately thereby losing correlations, the cDDPM algorithm pre-
 222 serves these correlations by simultaneously unfolding the full multidimensional phase space.
 223 The 1D Wasserstein distances for individual quantities are provided (labeled "Wasserstein" in
 224 figures), as well as multidimensional Wasserstein distances for the complete set of variables
 225 defining physics objects (reported in tables).

226 In addition to the Wasserstein distance, two other metrics are employed to evaluate the
 227 unfolding performance. The first is the chi-squared per degree of freedom (χ^2/DoF) on the
 228 binned distribution, denoted as "Binned χ^2/DoF " in the figures. This metric assesses the agree-
 229 ment between the unfolded and true distributions while accounting for statistical fluctuations
 230 in each bin. The second metric is the sum of the absolute values of (ratios - 1), where the ratios
 231 are calculated as the detector or unfolded distributions divided by the truth distribution. This
 232 metric, labeled as " $\sum |\text{ratios} - 1|$ " in the figures, provides a measure of the overall deviation
 233 from the true distribution across all bins.

234 In the figures, the binning of the distributions is done after unfolding to provide a simple
 235 visual representation of the results. From the unfolded results, any choice of binning can
 236 be used, allowing the data to be presented in various ways and adapted for specific physics
 237 analyses.

238 This study focuses on QCD jets, which are narrow streams of hadrons produced by quark
 239 or gluon hadronization in high-energy particle collisions. Jets are typical object signatures in
 240 HEP data that provide information about the fundamental interaction of nature that leads to
 241 their production. Multiple toy model jet datasets are designed, each representing a distinct
 242 physics process, with each dataset independently distorted by detector effects. Each jet is
 243 characterized by a 4-vector containing kinematic information: transverse momentum (p_T),
 244 pseudorapidity (η), azimuthal angle (ϕ), and energy (E). To emulate realistic physics data,
 245 these parameters are each sampled from specific distributions. The particle p_T is sampled from
 246 an exponential distribution $f(x; 1/\beta) = (1/\beta) \exp(-x/\beta)$, reflecting the typical exponential
 247 behavior observed in p_T distributions in particle physics. The azimuthal angle ϕ is sampled
 248 uniformly from the range $[-\pi, \pi]$, while the pseudorapidity η follows a Gaussian distribution
 249 with $\mu = 0$ and $\sigma = 2$. Assuming massless jets for simplicity, the energy is calculated as
 250 $E = p_T \cosh \eta$. These components for the "truth-level" jet vector (\mathbf{x}), which is then processed
 251 through a detector-like smearing framework, producing a "detector-level" jet quantities (\mathbf{y})
 252 that mimics the particle interactions within an actual detector. For a comprehensive description
 253 of the detector smearing, please refer to B.3.

254 **Part 1: Dedicated Unfolder** We first consider how to setup a *dedicated* cDDPM unfold-
 255 (without use of the distributional moments) that can achieve multidimensional object-wise
 256 unfolding for a single physics process. A cDDPM can be trained with data pairs (\mathbf{x}, \mathbf{y}) as input
 257 to learn the posterior distribution $P(\mathbf{x}|\mathbf{y})$. To unfold, the detector data \mathbf{y} are given as input and
 258 the cDDPM acts as a posterior sampler of $P(\mathbf{x}|\mathbf{y})$. This dedicated unfold-er relies on learning a
 259 specific posterior distribution, which implicitly incorporates information about the prior dis-
 260 tribution of the training data. While not explicitly using the prior, this approach results in an
 261 unfold-er that is more tailored to the particular distribution represented in the training data.
 262 This makes the dedicated unfold-er particularly well-suited for scenarios involving well-known

263 distributions, where the implicit bias towards the training prior is acceptable or even desirable.
 264 To validate this approach, two test cases are presented that probe different aspects of the
 265 cDDPM dedicated unfold. In case (1), the ability of the cDDPM to learn a posterior $P(\mathbf{x}|\mathbf{y})$
 266 given a dataset of pairs $\{\mathbf{x}, \mathbf{y}\}$ is tested, where both the training and test datasets have the
 267 same prior and posterior distributions (Figure 2a). Case (2), depicted in Figure 2b, evaluates
 268 the unfolding accuracy when the training and test datasets have different underlying true dis-
 269 tributions (priors) but the exact same posterior $P(\mathbf{x}|\mathbf{y})$. Since the cDDPM does not explicitly
 270 evaluate the prior distribution of the training dataset, it can sample from the posterior distri-
 271 bution without an imposed bias towards underlying characteristics of the prior of the training
 272 data. The successful unfolding in each case validates the cDDPM formulation, showing that is
 273 able to learn the posterior $P(\mathbf{x}|\mathbf{y})$ with minimal influence from the specific shape of the training
 274 distribution, a crucial feature for developing our unbiased generalized unfold.

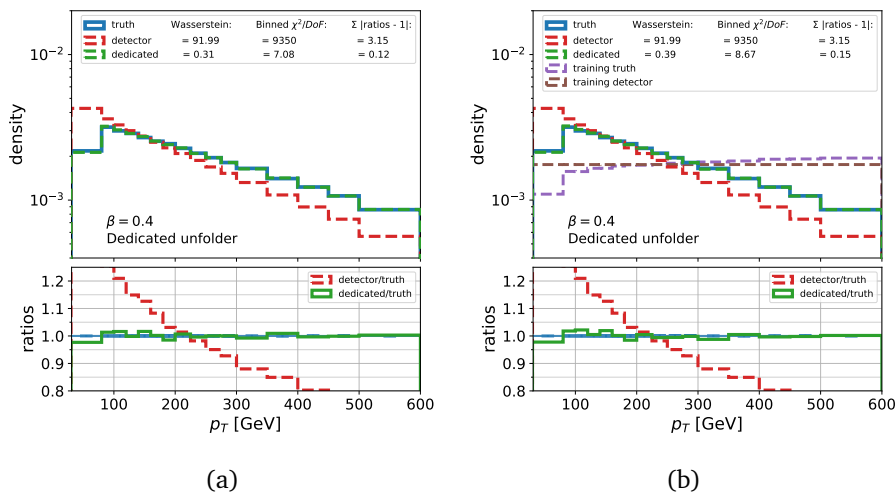


Figure 2: Unfolding results for toy-model data using a cDDPM dedicated unfold. On the left (a) are the results for case (1) that tests the ability of the cDDPM to learn a posterior $P(\mathbf{x}|\mathbf{y})$ given a dataset of pairs $\{\mathbf{x}, \mathbf{y}\}$. On the right (b) are the results for case (2) which aims to unfold the same test dataset (red) as (a), but the cDDPM unfold is trained using an alternative training dataset (purple and brown) that is created such that the posterior $P(\mathbf{x}|\mathbf{y})$ is the same as test data, while the priors (blue vs purple) are different.

275 **Part 2: Generalized Unfolder** Since physics generators cannot perfectly emulate real physics
 276 processes, it is important to minimize assumptions and bias towards specific generator mod-
 277 els. Although the dedicated unfold exhibits reduced bias towards the underlying distribution
 278 used in training, the learned posterior by definition contains information about the prior. Con-
 279 sequently, the posterior will not be the same across different data distributions subjected to
 280 the same detector effects. This means the generalization power of the dedicated unfold is
 281 limited, as it is strictly tied to one specific posterior, and it may not successfully unfold data for
 282 which the observed distribution widely differs from those assumed in the training phase of the
 283 algorithm. Here, the aim is to develop a *generalized* cDDPM unfold capable of handling a
 284 broader range of posteriors, thus enabling the unfolding of data from diverse physics processes
 285 with minimal bias towards generated physics distributions.

286 To achieve this, the model's inductive bias is enhanced by incorporating distributional mo-
 287 ments into the approach, allowing for interpolation and extrapolation to unseen posteriors as
 288 discussed in Section 2.1. This strategy is implemented by expanding the training dataset to

289 include data pairs (\mathbf{x}, \mathbf{y}) from multiple different distributions. For each dataset, six moments
 290 of the p_T distribution are computed: the first raw moment (mean) $\mu = \frac{1}{N} \sum_{i=1}^N p_{T,i}$, followed
 291 by the 2nd through 6th central moments calculated as $\mu_k = \frac{1}{N} \sum_{i=1}^N (p_{T,i} - \mu)^k$ for $k = 2, \dots, 6$.
 292 These moments are process-specific – they are calculated once for each physics process using
 293 the full set of events for that process. The moments are then appended to both the truth-
 294 level (\mathbf{x}) and detector-level (\mathbf{y}) vectors of every jet from that process. This means that while
 295 the kinematic components of \mathbf{x} and \mathbf{y} vary jet-by-jet, all jets from the same physics process
 296 carry identical moment values that characterize that process’s p_T distribution. The resulting
 297 augmented vectors contain both the jet-specific kinematic information and these process-level
 298 distributional features that help distinguish between different physics processes.

299 The moments of the p_T distribution are chosen because they are characteristic features
 300 for different physics processes, with the raw first moment capturing the average scale of the
 301 process and the higher central moments characterizing the shape and asymmetry of the distri-
 302 bution. Tests were performed incorporating moments from all components of the jet kinematic
 303 vector (p_T, η, ϕ, E) , but this resulted in degraded performance. This degradation likely stems
 304 from the inclusion of moments from components like η and ϕ distributions that remain similar
 305 across different physics processes, effectively introducing noise to the conditioning informa-
 306 tion rather than providing discriminating features. Additional tests comparing the use of four
 307 versus six moments showed improved performance with six moments, leading to this choice
 308 for the final implementation, though the possibility of including even higher moments was not
 309 investigated.

310 Two approaches were explored for incorporating the distributional moments into the un-
 311 folding process. The first method appends the moments only to the detector-level vector \mathbf{y} ,
 312 using them purely as conditioning information. The second approach includes the moments in
 313 both the truth-level and detector-level vectors $(\mathbf{x}$ and \mathbf{y} , respectively), allowing the truth-level
 314 moments to participate in the denoising process. While both approaches showed promising
 315 results, the latter demonstrated marginally better performance and was therefore adopted for
 316 all results presented in this work.

317 With this chosen approach, the truth-level vector \mathbf{x} and detector-level vector \mathbf{y} are now
 318 redefined to include these distributional moments, creating augmented jet vectors that en-
 319 compass both the original kinematic information and the newly added moment data. It is
 320 important to note that while these moments are unfolded along with the jet kinematic com-
 321 ponents, they primarily serve as conditioning information and are not part of the final output,
 322 being discarded after the unfolding process. By training with these diverse augmented data
 323 pairs (\mathbf{x}, \mathbf{y}) , the cDDPM is enabled to represent multiple posteriors corresponding to the dis-
 324 tributions in the expanded training dataset, distinguishable through the added distributional
 325 information provided by the moments (more details on the training dataset are provided in
 326 **B.1**).

327 To evaluate the efficacy of the generalized unfold approach, a series of tests are con-
 328 ducted using a cDDPM trained on an expanded dataset. This dataset incorporates four distinct
 329 p_T distributions: a uniform distribution and exponential functions with $\beta = 0.7, 0.3$, and 0.07 ,
 330 all augmented with their respective distributional moments. The model’s performance is then
 331 assessed on test datasets sampled from previously unseen p_T distributions ($\beta = 0.4, 0.2$, and
 332 0.06). Figure 3 presents these results, demonstrating the unfold’s ability to generalize across
 333 different distributions.

334 To further investigate the impact of including the distributional moments in the data vector,
 335 additional experiments are performed. Figure 4 illustrates two critical test cases: (a) unfolding
 336 without including any moments in the training or test datasets, and (b) unfolding with “fake”
 337 moments (random numbers) assigned to the distributions. The stark contrast in performance
 338 between these cases and the primary results highlights the crucial role that the distributional

339 moments play in achieving a truly generalized unfold. These results demonstrate the strong
 340 potential use of the cDDPM for generalized unfolding in HEP physics analyses.

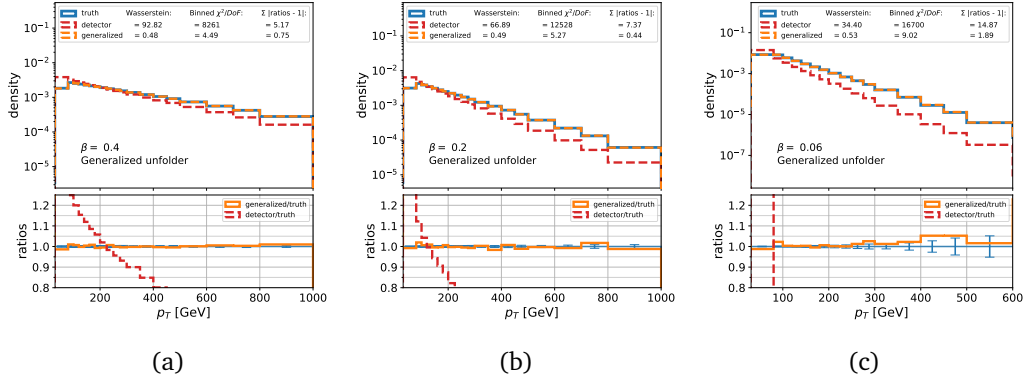


Figure 3: Unfolding results for toy-model data using a cDDPM generalized unfold. The unfold is trained on an expanded dataset including four distinct p_T distributions (uniform and exponential with $\beta = 0.7, 0.3$, and 0.07). Panels show unfolding performance on previously unseen p_T distributions: (a) $\beta = 0.4$, (b) $\beta = 0.2$, and (c) $\beta = 0.06$. This shows the unfold’s ability to generalize across different distributions, with the unfolded results (orange) closely matching the truth-level data (blue). Error bars indicate statistical uncertainties after binning.

341 3 Application to Particle Physics Data

342 3.1 Setup and Implementation

343 The approach is now tested on simulated particle physics data. The objective is to test whether
 344 the method can succeed in accurately unfolding physically relevant observable distributions,
 345 rather than simple functions, with more complex detector effect models. Using the PYTHIA
 346 event generator [28], jet datasets for various physics processes ($t\bar{t}$, W +jets, Z +jets, dijet, and
 347 leptoquark) are generated under different theoretical modelings of the processes (details of
 348 these synthetic datasets can be found in B.2). In this context, the jet kinematic information
 349 is defined with a vector that includes the transverse momentum (p_T), pseudorapidity (η),
 350 azimuthal angle (ϕ), and its 4-momentum (E, p_x, p_y, p_z).

351 These jet vectors are defined both at truth-level as \mathbf{x} and detector-level as \mathbf{y} , with one vec-
 352 tor pair (\mathbf{x}, \mathbf{y}) corresponding to each individual jet in an event. The generated truth-level jets
 353 were passed through two different detector simulation frameworks to simulate particle inter-
 354 actions within an LHC detector. The detector simulations used were DELPHES [29] with the
 355 standard configuration for the CMS detector [30], and another detector smearing framework
 356 developed using an analytical data-driven approximation for the p_T , η , and ϕ resolutions from
 357 results published by the ATLAS collaboration [31] (more details in B.3). DELPHES provides a
 358 comprehensive detector simulation that takes into account the full detector geometry and its
 359 impact on particle reconstruction, while the data-driven detector smearing focuses on resolu-
 360 tion effects, allowing testing of the unfolding success under more drastic detector smearing.
 361 Separate generalized unfolders were trained for the datasets from each detector-effects frame-
 362 work, but the implementation and application of the cDDPM methodology remained identical
 363 in both cases.

364 Using the same approach described for the toy model, both dedicated and generalized cD-

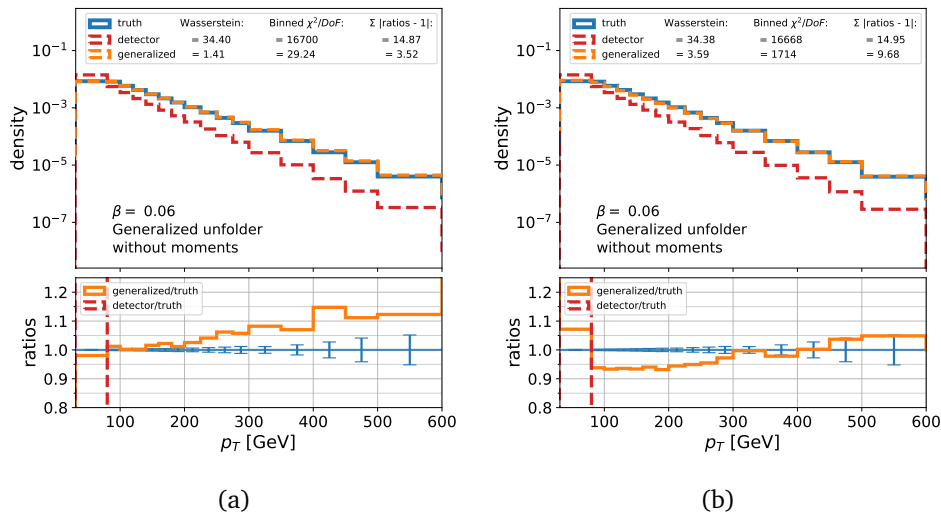


Figure 4: Investigation of the role of distributional moments in unfolding performance. Using the same distributions as Figure 3 (a uniform distribution as well as exponential distributions with $\beta = 0.7, 0.3, 0.07$) for training and $\beta = 0.06$ for testing, two test cases are shown: (a) unfolding without including any distributional moments in either training or test datasets, and (b) unfolding with random numbers assigned as "fake" moments. The degraded performance in both cases, compared to Figure 3c, demonstrates the crucial role that true distributional moments play in achieving successful generalized unfolding.

365 DPM unfolders are trained. The dedicated unfolders are trained using data pairs (\mathbf{x}, \mathbf{y}) , excluding
 366 the distributional moments. In contrast, the generalized unfolders are trained on multiple simu-
 367 lated physics processes (detailed in B.2). For the data-driven detector smearing framework,
 368 the training dataset includes 18 different physics processes, while for the DELPHES frame-
 369 work, 6 different processes were available for training. For each process, the training dataset
 370 incorporates the first 6 central moments of the p_T distribution, appending these moments to
 371 the corresponding truth-level and detector-level data vectors of that distribution. All results
 372 presented in this section use these trained generalized unfolders (one for each detector simu-
 373 lation framework), with all test cases shown being from datasets that were excluded from the
 374 training. A key distinction between the generalized and dedicated unfolders lies in their learn-
 375 ing outcomes: the generalized unfolders learn to model multiple posteriors from the diverse
 376 physics processes in its training data, whereas the dedicated unfolders capture only a single
 377 posterior represented by its specific training set. This difference allows us to use the dedicated
 378 unfolders as a performance benchmark, against which we can evaluate the effectiveness of the
 379 generalized unfolders.

380 3.2 Performance Evaluation

381 Figure 1 showcases two critical test cases that demonstrate the versatility of the generalized
 382 unfolders. Panel (a) presents results from an "unknown" process dataset, created by combining
 383 jets from multiple sources: 40% $t\bar{t}$, 35% W +jets, and 25% leptoquark test datasets. While this
 384 "unknown" dataset is constructed from known physics processes (though none were included
 385 in the training data), their combination produces a unique prior distribution. The moments
 386 used for conditioning are calculated from the combined dataset as a whole, presenting the
 387 generalized unfolders with previously unseen distributional characteristics. This scenario rep-
 388 represents the optimal use case for the generalized unfolders, as it simulates a situation where

389 the underlying physics is not fully known or understood, such as in new physics searches. To
390 demonstrate this, the generalized unfolders is compared against a dedicated unfolded trained
391 on $t\bar{t}$ data, chosen as it represents the majority component (40%) of the unknown process.
392 The generalized unfolded demonstrates superior performance when unfolding this unknown
393 process, effectively adapting to the mixed nature of the data without prior knowledge of its
394 composition. In contrast, the dedicated unfolded, constrained by its assumption of a $t\bar{t}$ -like
395 posterior, shows reduced accuracy. This comparison underscores the generalized unfolded's
396 potential in scenarios involving new or unexpected physics processes, where the underlying
397 distribution may deviate significantly from known models. Panel (b) further demonstrates the
398 generalized unfolded's performance, unfolding data from graviton production in the context
399 of large extra-dimension scenarios [21], accompanied by jets. This process, which features
400 distinctly different physics signatures from Standard Model processes, was completely absent
401 from the training data yet is accurately unfolded by our method. Here too, the generalized
402 unfolded demonstrates accuracy in reconstructing the true distributions, effectively adapting
403 to entirely new physics processes without prior knowledge of their underlying physics.

404 Figures 5 and 6 illustrate unfolding results for the data-driven detector smearing and
405 DELPHES detector simulation, respectively. The plots show various jet kinematics across dif-
406 ferent generated physics datasets that are not included in the training data, showcasing the
407 generalized unfolded's versatility. While the generalized unfolded's advantage is expected for
408 unknown processes, we also aim for comparable performance to dedicated unfolded on known
409 processes. The results demonstrate that the performance between the dedicated and general-
410 ized unfolded is indeed comparable. The statistical measures of agreement between the un-
411 folded and true distributions slightly favor the dedicated unfolded, particularly in the binned
412 χ^2/DoF metric, due to small fluctuations in the first three bins of the distributions where
413 the statistical uncertainty is negligible. This is particularly evident in the W+jets unfolding
414 (Figure 5c), where the generalized unfolded exhibits nearly perfect agreement with the true
415 distribution in the low- p_T region, achieving better performance than the dedicated unfolded
416 in these bins, despite showing some deviations in the tail of the distribution. When consid-
417 ered in the context of total systematic uncertainties typical in real measurement processes,
418 the observed deviations between dedicated and generalized unfolded would contribute neg-
419 ligibly to the overall measurement uncertainty due to unfolding. Therefore, the generalized
420 unfolded achieves a similar level of accuracy to the dedicated unfolded on known processes,
421 while maintaining the advantage of being able to unfold unknown processes.

422 To validate our framework's effectiveness we compare both unfolded across various test
423 datasets, and Table 2 presents the resulting multidimensional Wasserstein distances to their
424 true distributions. These quantitative results support the earlier conclusions, showing compa-
425 rable performance between the generalized and dedicated unfolded across all test processes.

426 To further evaluate the generalized unfolded's capabilities, a test is conducted using $t\bar{t}$
427 datasets generated with different theoretical modeling from varying settings in PYTHIA. Ded-
428 icated unfolded are trained on each of these $t\bar{t}$ variants, except for one "unseen" variant. The
429 variants used for training the dedicated unfolded employed different PDFs (CTEQ6L1 and
430 NNPDF23) with the default PYTHIA parton showers, while the unseen variant used CT14lo
431 PDF with the Vincia parton shower [32]. The performance of these dedicated unfolded and
432 the generalized unfolded is then compared in unfolding the unseen $t\bar{t}$ variant. Figure 7 illus-
433 trates the results of this test. As shown in Figure 7, the generalized unfolded accurately un-
434 folds the data from the unseen $t\bar{t}$ variant, and in some metrics outperforms the dedicated unfolded
435 trained on the other $t\bar{t}$ variants. This result demonstrates the generalized unfolded's ability
436 to capture subtle differences between posterior distributions arising from different generator
437 settings, not just large variations across different physics processes. Such capability suggests
438 that the generalized unfolded could be a valuable tool in refining models for generating known

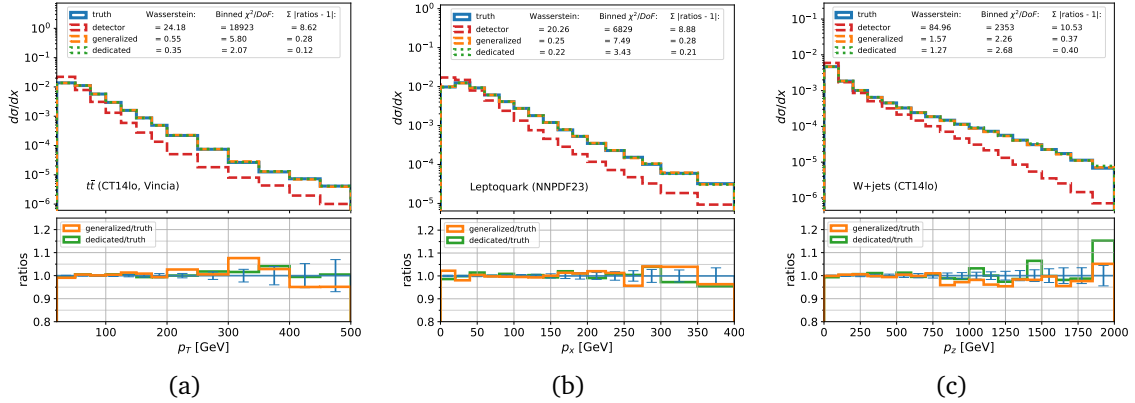


Figure 5: Unfolding results of the data-driven detector smearing of jet vector components for three different physics processes: (a) $t\bar{t}$, (b) leptoquark, and (c) W +jets. For each process, results from the generalized cDDPM unfolders (orange) are compared against a process-specific dedicated unfolded (green), where each dedicated unfolded was trained exclusively data from its corresponding process, while the generalized unfolded was trained on a diverse dataset excluding all three test processes. In all cases, the unfolding accuracy is within the expected uncertainty budget typical of experimental measurements of these distributions. Error bars indicate statistical uncertainties.

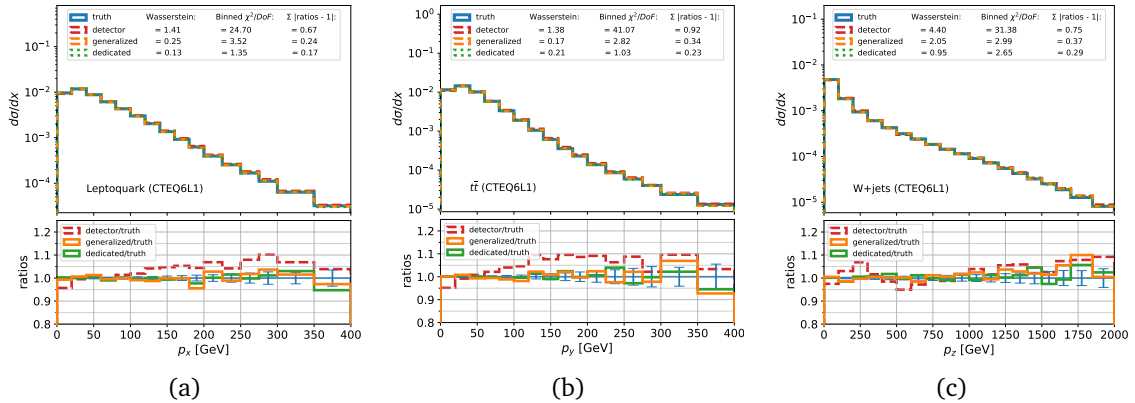


Figure 6: Unfolding results of the DELPHES detector simulation of jet vector components for three different physics processes: (a) leptoquark, (b) $t\bar{t}$, and (c) W +jets. For each process, results from the generalized cDDPM unfolded (orange) are compared against a process-specific dedicated unfolded (green). Each dedicated unfolded was trained exclusively on its corresponding physics process, while the generalized unfolded was trained on a diverse dataset excluding all three test processes. In all cases, the unfolding accuracy is within the expected uncertainty budget typical of experimental measurements of these distributions. Error bars indicate statistical uncertainties.

Table 2: Comparison of Wasserstein distances for detector-level data and unfolded results using the data-driven detector smearing, corresponding to the processes shown in Figures 5 and 6. Values are shown for both generalized and dedicated unfolders across different physics processes.

Detector	Process	Multidimensional Wasserstein Distances		
		Detector	Generalized	Dedicated
Data-driven	“Unknown”	28.20	0.74	2.68
	Graviton (CT14lo)	31.35	0.64	N/A
	$t\bar{t}$ (CT14lo, Vincia)	26.43	0.34	0.35
	Leptoquark (NNPDF23)	32.42	0.26	0.30
	W+Jets (CT14lo)	31.09	0.54	0.52
DELPHES	$t\bar{t}$ (CTEQ6L1)	1.49	0.21	0.20
	Leptoquark (CTEQ6L1)	1.51	0.27	0.15
	W+Jets (CTEQ6L1)	2.07	0.62	0.34

439 physics processes, as it can adapt to nuanced variations in the underlying distributions without
 440 requiring specific training on each variant of PDF or parton shower model.

441 In Figure 8, the model’s efficacy is further demonstrated with two tests: (1) reconstructing
 442 jet mass from unfolded results, indicating well-preserved correlations among jet vector com-
 443 ponents, and (2) reconstructing event-level observables from unfolded quantities, achieved
 444 by tracking event numbers through object-wise unfolding. The successful reconstruction of
 445 jet mass, which is not directly unfolded but derived from the unfolded jet vector compo-
 446 nents, showcases the method’s ability to maintain complex relationships between variables.
 447 This preservation of correlations allows for the calculation of various derived quantities post-
 448 unfolding, offering the option to construct new observables that are not explicitly part of the
 449 original unfolding process.

450 3.3 Computational Performance

451 The generalized folder demonstrates not only accuracy but also computational efficiency,
 452 making it a practical tool for large-scale physics analyses. The model used in these results,
 453 trained on a diverse dataset of 1.8 million jets, requires approximately 3 hours of training
 454 time on an NVIDIA A100 GPU. Once trained, the generalized folder can be applied rapidly,
 455 with the ability to unfold 1 million events in approximately 3 minutes. This speed is partic-
 456 ularly advantageous as the generalized folder does not require retraining for specific pro-
 457 cesses, unlike the dedicated folder and other machine-learning based unfolding methods.
 458 Consequently, it can provide fast, object-wise unfolding results across a wide range of physics
 459 analyses without incurring additional training overhead for each new process or dataset. For
 460 a more detailed discussion of the computational performance and parameters of the model,
 461 see A.2.

462 4 Conclusion

463 The results presented in this paper demonstrate the generalized cDDPM folder can success-
 464 fully unfold detector effects on particle jets from a variety of physics processes, including those
 465 not seen during training. The key feature of this method is its non-iterative and flexible pos-
 466 terior sampling approach, which exhibits a strong inductive bias allowing generalization to
 467 unseen processes without explicitly assuming the underlying physics distribution. The gen-

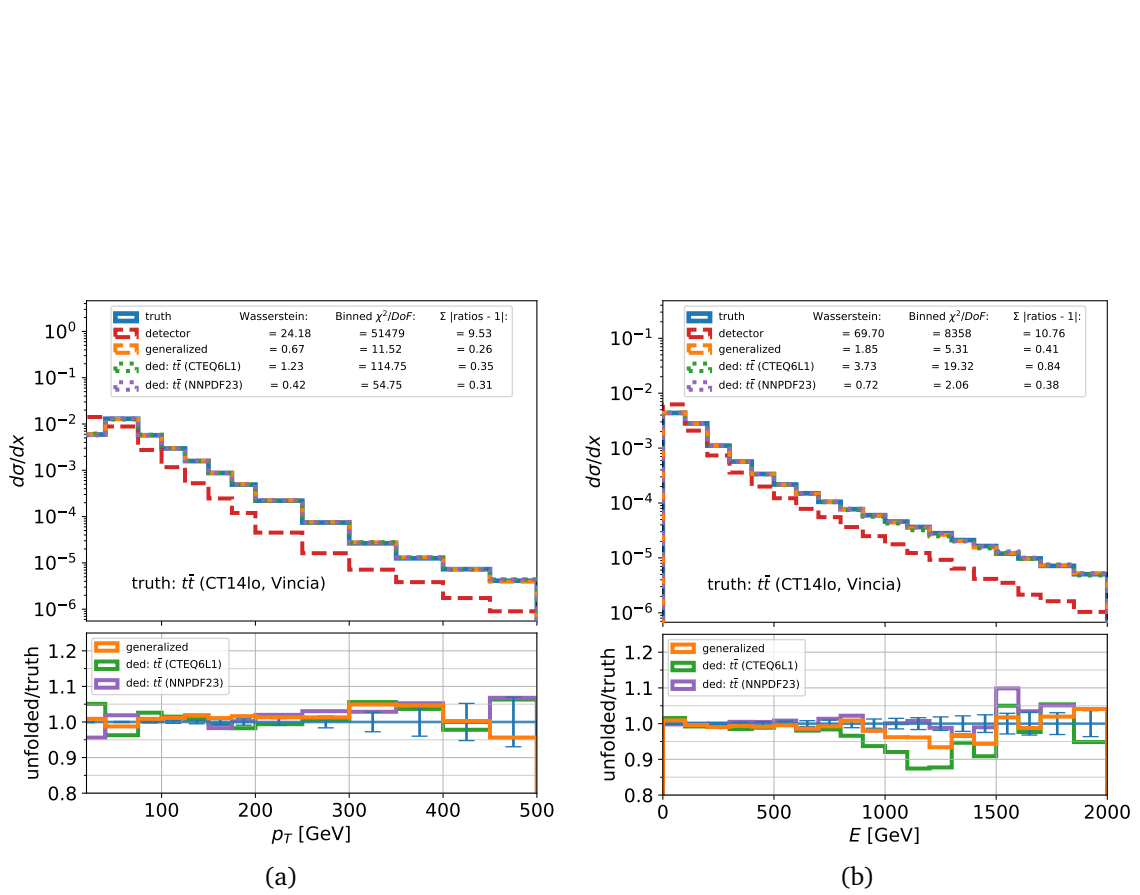


Figure 7: Comparison of unfolding performance for the data-driven detector smearing on an unseen $t\bar{t}$ variant generated using CT14lo PDF and the Vincia parton shower. The generalized unfolders (orange) demonstrates comparable performance to dedicated unfolders trained on other $t\bar{t}$ variants with different PDFs (CTEQ6L1 shown in green, NNPDF23 shown in purple). While these dedicated unfolders were each trained exclusively on their respective $t\bar{t}$ variant, the generalized unfolders were trained on a diverse dataset of multiple physics processes described in the text. Shown are the jet p_T (a) and E (b) distributions.

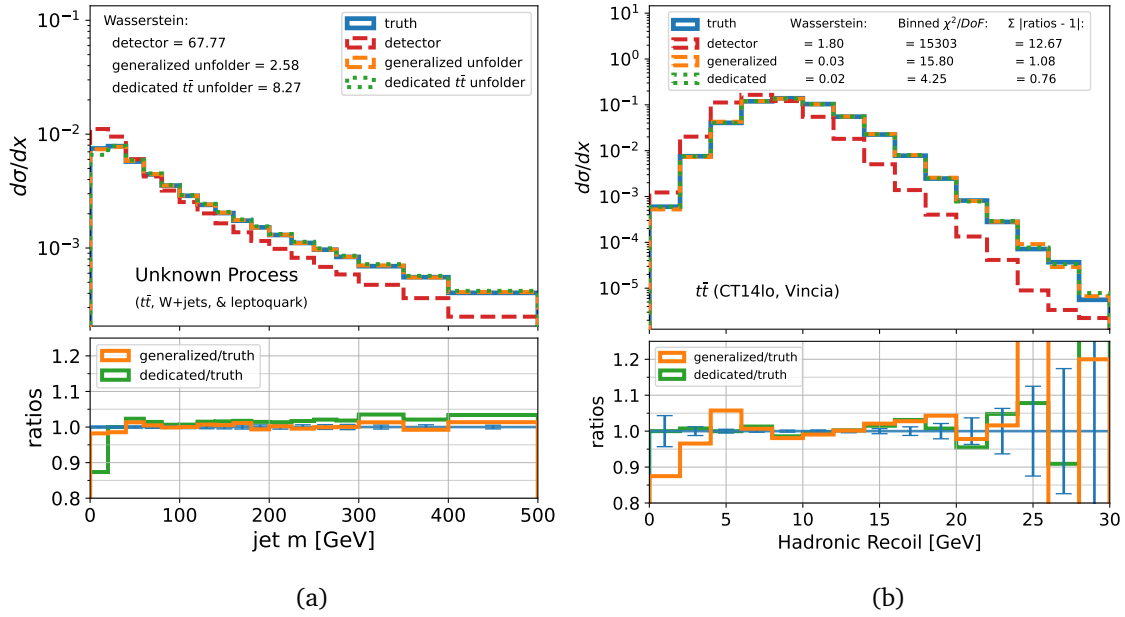


Figure 8: Demonstration of the unfolding method’s ability to preserve correlations and handle event-level quantities using the data-driven detector smearing. Reconstruction of jet mass (a), derived from the unfolded four-momentum components, showing that correlations between kinematic variables are well-preserved through the unfolding process. Reconstruction of the hadronic recoil (b), an event-level observable obtained by combining multiple unfolded jets within each event, demonstrating the method’s capability to handle event-level physics quantities. Error bars indicate statistical uncertainties.

468 eralized unfolders therefore provides a solution to the unfolding problem that addresses both
469 the bias and the generalization challenges, something other approaches to unfolding did not
470 attempt.

471 Additionally, the generalized unfolded is able to accurately reconstruct jet properties and
472 derived quantities like jet mass, even for processes absent from its training data. By preserving
473 correlations between jet vector components, it enables the construction of complex observables
474 post-unfolding, offering new possibilities in data analysis.

475 This approach also offers computational advantages. Once trained, it eliminates the need
476 for process-specific retraining, reducing computational overhead. The ability to unfold a mil-
477 lion events in approximately 3 minutes demonstrates its potential for efficient large-scale data
478 processing in high-energy physics experiments.

479 While the generalized approach excels in adaptability, dedicated unfolders remain valu-
480 able for scenarios requiring maximum precision on well-understood processes. The cDDPM
481 formulation can easily be adapted for either the generalized or dedicated unfolded, allowing for
482 flexibility in choosing the most appropriate approach based on specific analysis requirements
483 and the desired balance between adaptability and specialized accuracy.

484 Several open questions remain regarding the implementation of the conditioning on the
485 moments. These include optimal selection of priors and the number of moments required for
486 the best unfolding performance. Further investigation is needed to determine the extent of
487 the cDDPM's inductive bias and its tolerance to variations in the underlying physics processes.
488 Understanding these aspects will help refine the method and ensure its robustness across a
489 wide range of scenarios.

490 While this approach shows promise, key limitations are acknowledged. The current studies
491 were performed on QCD jets, and extending this method to other particle types is necessary for
492 its comprehensive application in data analysis. Addressing particles outside detector thresh-
493 olds and accounting for systematic and experimental uncertainties are crucial improvements
494 needed to fully realize the method's potential in practical applications. An important constraint
495 of the current implementation is that while correlations between object vector components are
496 preserved, the model lacks access to event-wise information which impacts the reconstruction
497 accuracy of certain event-level observables. These improvements and extensions are left for
498 future work.

499 To conclude, the results confirm the versatility of the generalized cDDPM unfolded across
500 diverse physics processes. This non-iterative and flexible posterior sampling approach exhibits
501 a strong inductive bias that allows the cDDPM to generalize to unseen processes without ex-
502 plicitly assuming the underlying distribution, setting it apart from other unfolding techniques
503 developed so far.

504 Acknowledgements

505 This work has been made possible thanks to the support of the Department of Energy Office of
506 Science through the Grant DE-SC0023964. Shuchin Aeron and Taritree Wonhijrad would also
507 like to acknowledge support by the National Science Foundation under Cooperative Agreement
508 PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions,
509 <http://iaifi.org/>).

510 **Funding information** This work was supported by the U.S. Department of Energy, Office of
511 Science, under Award Number DE-SC0023964.

512 A cDDPM Details

513 A.1 Loss Derivation

514 In the proposed cDDPM, the forward process is a Markov chain that gradually adds Gaussian
515 noise to the data according to a variance schedule β .

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) := \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I}) \quad (\text{A.1})$$

516 To recover the original sample from a Gaussian noise input, this process needs to be reversed.
517 This can be achieved through the use of a model p_θ which corresponds to the joint distribution
518 $p_\theta(x_{0:T}|y) = p_\theta(x_0, x_1, \dots, x_T|y)$, and it is defined as a Markov chain with learned Gaussian
519 transitions starting at $p(x_T|y) = \mathcal{N}(x_T; 0, I)$

$$p_\theta(\mathbf{x}_{0:T}|\mathbf{y}) := p(\mathbf{x}_T|\mathbf{y}) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{y}) \quad (\text{A.2})$$

520

$$p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{y}) := \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_\theta(t, \mathbf{x}_t, \mathbf{y}), \Sigma_\theta(t, \mathbf{x}_t, \mathbf{y})) \quad (\text{A.3})$$

521 where $\boldsymbol{\mu}_\theta$ represents the learned mean, and Σ_θ represents the learned covariance of the Gaus-
522 sian transitions, which vary with time step t :

$$\Sigma_\theta(t, \mathbf{x}_t, \mathbf{y}) = \sigma^2 \mathbf{I}, \quad \sigma^2 = \beta_t. \quad (\text{A.4})$$

523 Training involves learning the reverse Markovian transitions that maximize the likelihood of
524 the training samples, which is equivalent to minimizing the variational upper bound on the
525 negative log likelihood. This negative log likelihood can be expressed in terms of the Kullback-
526 Leibler (KL) divergence [33], a statistical measure of the difference between two probability
527 distributions P and Q :

$$D_{KL}(P||Q) = \sum_{x \in \mathcal{X}} P(x) \left(\log \frac{P(x)}{Q(x)} \right) \quad (\text{A.5})$$

528 Applying this, the variational bound on the negative log likelihood can be expressed as:

$$\begin{aligned} \mathbb{E}[-\log p_\theta(\mathbf{x}_0|\mathbf{y})] &\leq \mathbb{E}[-\log p_\theta(\mathbf{x}_0|\mathbf{y})] + D_{KL}(q(\mathbf{x}_{1:T}|\mathbf{x}_0)||p_\theta(\mathbf{x}_{1:T}|\mathbf{x}_0, \mathbf{y})) \\ &= \mathbb{E}[-\log p_\theta(\mathbf{x}_0|\mathbf{y})] + \mathbb{E}_q \left[\log \frac{q(\mathbf{x}_{1:T}|\mathbf{x}_0)}{p_\theta(\mathbf{x}_{1:T}|\mathbf{x}_0, \mathbf{y})} \right] \\ &= \mathbb{E}[-\log p_\theta(\mathbf{x}_0|\mathbf{y})] + \mathbb{E}_q \left[\log \frac{q(\mathbf{x}_{1:T}|\mathbf{x}_0)}{p_\theta(\mathbf{x}_{0:T}|\mathbf{y})/p_\theta(\mathbf{x}_0|\mathbf{y})} \right] \\ &= \mathbb{E}[-\log p_\theta(\mathbf{x}_0|\mathbf{y})] + \mathbb{E}_q \left[\log \frac{q(\mathbf{x}_{1:T}|\mathbf{x}_0)}{p_\theta(\mathbf{x}_{0:T}|\mathbf{y})} \right] + \mathbb{E}[\log p_\theta(\mathbf{x}_0|\mathbf{y})] \\ &= \mathbb{E}_q \left[-\log \frac{p_\theta(\mathbf{x}_{0:T}|\mathbf{y})}{q(\mathbf{x}_{1:T}|\mathbf{x}_0)} \right] \\ &= \mathbb{E}_q \left[-\log p(\mathbf{x}_T|\mathbf{y}) - \sum_{t \geq 1} \log \frac{p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{y})}{q(\mathbf{x}_t|\mathbf{x}_{t-1})} \right] := L \end{aligned} \quad (\text{A.6})$$

529 Following the similar derivation provided in [23], this loss can then be rewritten using the
530 KL-divergence

$$\begin{aligned}
L &= \mathbb{E}_q \left[-\log \frac{p_\theta(\mathbf{x}_{0:T}|\mathbf{y})}{q(\mathbf{x}_{1:T}|\mathbf{x}_0)} \right] \\
&= \mathbb{E}_q \left[-\log p(\mathbf{x}_T|\mathbf{y}) - \sum_{t \geq 1} \log \frac{p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{y})}{q(\mathbf{x}_t|\mathbf{x}_{t-1})} \right] \\
&= \mathbb{E}_q \left[-\log p(\mathbf{x}_T|\mathbf{y}) - \sum_{t > 1} \log \frac{p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{y})}{q(\mathbf{x}_t|\mathbf{x}_{t-1})} - \log \frac{p_\theta(\mathbf{x}_0|\mathbf{x}_1, \mathbf{y})}{q(\mathbf{x}_1|\mathbf{x}_0)} \right] \\
&= \mathbb{E}_q \left[-\log p(\mathbf{x}_T|\mathbf{y}) - \sum_{t > 1} \log \frac{p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{y})}{q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)} \cdot \frac{q(\mathbf{x}_{t-1}|\mathbf{x}_0)}{q(\mathbf{x}_t|\mathbf{x}_0)} - \log \frac{p_\theta(\mathbf{x}_0|\mathbf{x}_1, \mathbf{y})}{q(\mathbf{x}_1|\mathbf{x}_0)} \right] \\
&= \mathbb{E}_q \left[-\log \frac{p(\mathbf{x}_T|\mathbf{y})}{q(\mathbf{x}_T|\mathbf{x}_0)} - \sum_{t > 1} \log \frac{p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{y})}{q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)} - \log p_\theta(\mathbf{x}_0|\mathbf{x}_1, \mathbf{y}) \right] \\
&= \mathbb{E}_q \left[\underbrace{D_{KL}(q(\mathbf{x}_T|\mathbf{x}_0) \| p(\mathbf{x}_T|\mathbf{y}))}_{L_T} + \sum_{t > 1} \underbrace{D_{KL}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) \| p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{y}))}_{L_{1:T-1}} - \underbrace{\log p_\theta(\mathbf{x}_0|\mathbf{x}_1, \mathbf{y})}_{L_0} \right] \tag{A.7}
\end{aligned}$$

531 The term L_T is a constant, as it is the KL-divergence between two distributions of pure noise,
532 and the L_0 term is a final denoising step with no comparison to the forward process posteriors.
533 For the term $L_{1:T-1}$, the forward process posteriors can be written as

$$q(\mathbf{x}_{t+1}|\mathbf{x}_t, \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_{t+1}; \tilde{\boldsymbol{\mu}}_t(\mathbf{x}_t, \mathbf{x}_0), \tilde{\boldsymbol{\beta}}_t I) \tag{A.8}$$

534

$$\begin{aligned}
&\text{where } \tilde{\boldsymbol{\beta}}_t = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \boldsymbol{\beta}_t, \quad \bar{\alpha}_t = \prod_{s=1}^t (1 - \beta_s) \\
&\text{and } \tilde{\boldsymbol{\mu}}_t(\mathbf{x}_t, \mathbf{x}_0) = \left(\frac{\beta_t \sqrt{\bar{\alpha}_{t-1}}}{1 - \bar{\alpha}_t} \mathbf{x}_0 + \frac{\sqrt{\bar{\alpha}_t} (1 - \bar{\alpha}_{t-1})}{(1 - \bar{\alpha}_t)} \mathbf{x}_t \right). \tag{A.9}
\end{aligned}$$

535 Using this forward process posterior together with the reverse process posterior defined in
536 Equation A.3, a parametrization for $\boldsymbol{\mu}_\theta(\mathbf{x}_t, t, \mathbf{y})$ is introduced that aims to predict $\tilde{\boldsymbol{\mu}}_t(\mathbf{x}_t, \mathbf{x}_0)$.
537 With this the loss becomes

$$L_{t-1} = \mathbb{E} \left[\frac{1}{2\sigma_t^2} \|\tilde{\boldsymbol{\mu}}_t(\mathbf{x}_t, \mathbf{x}_0) - \boldsymbol{\mu}_\theta(t, \mathbf{x}_t, \mathbf{y})\|^2 \right] + C \tag{A.10}$$

538 where C is a constant, and $\tilde{\boldsymbol{\mu}}_t$ and $\boldsymbol{\mu}_\theta$ can be reparametrized using $\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}$
539 and reduced to

$$L_{t-1} = \mathbb{E}_{\boldsymbol{\epsilon}, \mathbf{x}_0, \mathbf{y}} \left[\frac{\beta_t^2}{2\sigma_t^2 \alpha_t (1 - \bar{\alpha}_t)} \|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_\theta(t, \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}, \mathbf{y})\|^2 \right]. \tag{A.11}$$

540 Finally we can write a simplified version of the loss with the terms differentiable in θ as

$$\begin{aligned}
L_{\text{simple}}(\theta) &= \mathbb{E}_{\boldsymbol{\epsilon}, \mathbf{x}_t, \mathbf{y}} \left[\|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_\theta(t, \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}, \mathbf{y})\|^2 \right] \\
&= \mathbb{E}_{\boldsymbol{\epsilon}, \mathbf{x}_t, \mathbf{y}} \left[\|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_\theta(t, \mathbf{x}_t, \mathbf{y})\|^2 \right]. \tag{A.12}
\end{aligned}$$

541 This derivation shows that in the cDDPM formulation, the task of learning a posterior
542 distribution reduces to minimizing a simple mean squared error between added and predicted
543 noise. This allows for estimation of the posterior without requiring explicit evaluation of the
544 prior distribution.

Algorithm 1 Conditional DDPM: Training

Input: dataset $\{\mathbf{x}_0, \mathbf{y}\}$, variance schedule β_1, \dots, β_T
 $t \leftarrow \text{Uniform}(\{1, \dots, T\})$ $\bar{\alpha}_t \leftarrow \prod_{s=1}^t (1 - \beta_s)$ $\epsilon \leftarrow \mathcal{N}(\mathbf{0}, \mathbf{I})$ **Repeat**a) $\mathbf{x}_t \leftarrow \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$ b) Calculate loss, $L = \|\epsilon - \epsilon_\theta(t, \mathbf{x}_t, \mathbf{y})\|^2$ c) Update θ via $\nabla_\theta L$ **Until** converged

Figure 9: The training procedure for the conditional DDPM unfolding model is presented. The algorithm trains on data samples $\{\mathbf{x}_0, \mathbf{y}\}$. In step (a) Gaussian noise ϵ is added to \mathbf{x}_0 over T timesteps according to the variance schedule. The model parameterized by θ is trained to estimate this added noise by observing the noisy states \mathbf{x}_t at a timestep t and the condition \mathbf{y} .

545 **A.2 Model Parameters**

546 During inference, the inputs are given to the denoising process are the vector \mathbf{y} and random
547 noise values $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. The denoising process removes noise from \mathbf{x}_T in T steps according
548 to the learned conditional distribution $p_\theta(\mathbf{x}_{0:T}|\mathbf{y})$. Pseudocode for the training and sampling
549 algorithms can be seen in Figures 9 and 10.

550 The cDDPM architecture consists of a Multi-Layer Perceptron (MLP), a feedforward neural
551 network, with approximately 1 million trainable parameters. It comprises three main compo-
552 nents: an initial linear layer with Gaussian Error Linear Unit (GELU) activation, which provides
553 smooth non-linear transformations, a time step embedding layer, and a series of linear layers
554 with GELU activations. The network takes as input the noised data and the time step. It first
555 processes the input through a 256-unit hidden layer, then adds a learned time step embed-
556 ding. This combined representation is passed through four 512-unit hidden layers, followed
557 by a 256-unit layer. Skip connections are employed between the input and output of the main
558 block. The final output layer predicts the noise at the given time step. Dropout (rate 0.01) is
559 applied after each linear layer to prevent overfitting during training.

560 The diffusion process employs a linear variance schedule over $T = 500$ time steps. The
561 schedule starts with an initial noise level $\beta_1 = 1e-4$ at the first step and increases linearly to β_T
562 $= 0.02$ at the final step. The model is trained using the Adam optimizer with an initial learning
563 rate of $3e-4$. To improve convergence and performance, a linear learning rate scheduler is
564 employed. It starts at the initial rate and linearly decreases to 1% of the initial rate ($3e-6$) by
565 the end of training.

566 The model is trained for 5000 epochs with a batch size of 2048. Using an NVIDIA A100
567 GPU, the training procedure on our full dataset or 1.8 million data points completes in approx-
568 imately 3 hours (more details on the training and test datasets in B). Once trained, the model
569 demonstrates efficient inference capabilities. Unfolding a dataset of 1 million data points takes
570 approximately 3 minutes on the A100 GPU, with processing time scaling linearly with the num-
571 ber of jets. Notably, this model functions as a generalizing unfold, eliminating the need for
572 retraining when applying it to various different datasets.

Algorithm 2 Conditional DDPM: Sampling

Input: detector-level data vector \mathbf{y} , variance schedule β_1, \dots, β_T
 $\mathbf{x}_T \leftarrow \mathcal{N}(\mathbf{0}, \mathbf{I})$ For $t = T, \dots, 1$ do

a) $\alpha_t \leftarrow 1 - \beta_t$, $\bar{\alpha}_t \leftarrow \prod_{s=1}^t \alpha_s$, $\sigma_t \leftarrow \sqrt{\beta_t}$

b) $\mathbf{z} \leftarrow \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $t > 1$, else $\mathbf{z} \leftarrow 0$

c) $\mathbf{x}_{t-1} \leftarrow \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}} \epsilon_\theta(t, \mathbf{x}_t, \mathbf{y}) \right) + \sigma_t \mathbf{z}$

Return \mathbf{x}_0

Figure 10: The trained conditional DDPM model serves as a posterior sampler, generating unfolded truth-level samples \mathbf{x}_0 given condition \mathbf{y} . Starting from pure noise \mathbf{x}_T , the conditioned reverse process denoises \mathbf{x}_t at each timestep by removing the estimated injected noise. Here $\sigma_t \equiv \sqrt{\beta_t}$ since this choice is optimal for a non-deterministic \mathbf{x}_0 .

573 **B Datasets**574 **B.1 Data Processing for cDDPMs**

575 The jets in both the toy model and physics datasets are limited to certain ranges for each of the
576 jet vector components, allowing a maximum p_T , p_x , p_y of 1000 GeV and a maximum p_z and
577 E of 4000 GeV, though these ranges can be extended by training with higher energy generated
578 physics datasets. The jet η range is between -4.4 and 4.4 (following standard detector limita-
579 tions), and the ϕ range between -3.5 and 3.5. The ϕ range is extended beyond $-\pi$ and π to
580 avoid discontinuity in training due to wrap-around values. The components of the jet vector
581 are divided by the respective maximum values so that the final vector components each range
582 between $[0,1]$ or $[-1, 1]$. As described in Section 2.3, the first 6 distributional moments of
583 the p_T distribution for each process are appended to the corresponding jet vectors. We choose
584 to use the moments of the p_T distribution because they serve as a distinguishing feature for
585 different physics processes. The p_T spectrum is particularly sensitive to the underlying physics
586 and provides valuable information for process discrimination. We experimented with includ-
587 ing moments from all components of the jet vector, but found that this approach led to reduced
588 performance.

589 The training dataset for the data-driven detector smearing is comprised of 1.8 million jets
590 generated from 18 different physics simulations described in Table 3. The training dataset for
591 the DELPHES detector simulation is comprised of 1.8 million jets generated from 6 different
592 physics simulations described in Table 4. The data pairs (\mathbf{x}, \mathbf{y}) are the input to the training
593 process. The test datasets are processed in the same way as the training dataset, limiting the
594 range and normalizing the range of the values. Additionally, a unique event identifier number,
595 which associates jets to their original event in the dataset, is included in each jet vector and
596 carried through the unfolding process (though not used as an input to the model) to enable
597 the reconstruction of event-level observables after unfolding. The jets in each of test datasets
598 pertain to only one process, with the exception of the “unknown process” dataset which is made
599 from a combination of 3 of the test datasets (40% $t\bar{t}$, 35% W +jets, and 25% leptoquarks) to
600 mimic a distribution from an unknown process that could not be easily unfolded with the

601 standard dedicated unfolding approach.

602 B.2 Physics Generation

603 Physics datasets are generated using PYTHIA 8.3 Monte-Carlo event generator. The simula-
 604 tions are run for proton-proton collisions at a center-of-mass energy of 13 TeV to emulate
 605 LHC physics interactions. The various physics processes used in this study are chosen for their
 606 high jet-production cross-section across a large jet energy range. The chosen processes are
 607 $t\bar{t}$, $(Z \rightarrow \mu\bar{\mu})+\text{jets}$, $(W \rightarrow \mu\bar{\nu})+\text{jets}$, dijets, and a new-physics process of leptoquarks. Note
 608 that each listed process typically includes multiple subprocesses (e.g., $t\bar{t}$ production occurs
 609 through both gluon-gluon and quark-antiquark collisions). For processes with unstable parti-
 610 cles, one particle per event decays to include at least one charged lepton, while other unstable
 611 particles decay hadronically. Each of these processes are run under multiple generator settings
 612 controlling the theoretical modeling of the underlying processes, with varying parton distri-
 613 bution functions (PDFs), parton shower models, and with an imposed phase space bias that
 614 increased the probability of generating events with high jet energies. For simulations where a
 615 phase space bias is applied, the events are sampled in the phase space as $(\hat{p}_T/p_T^{\text{ref}})^a$, where we
 616 set $p_T^{\text{ref}} = 100$ GeV and $a = 5$ such that events with a p_T over 100 GeV will be oversampled,
 617 increasing the event statistics in high-energy regions [28]. A list of the physics processes gen-
 618 erated for the data-driven detector smearing framework is shown in Table 3, and a list of the
 619 processes generated for the DELPHES detector simulation is shown in Table 4. Unless stated
 620 otherwise, the simulations are run with the PYTHIA simple parton shower model and no phase
 621 space bias.

622 B.3 Detector Smearing and Jet Matching

623 Our results present the unfolding of detector effects from two different detector smearing
 624 frameworks: DELPHES and a data-driven approach. DELPHES is a framework developed for
 625 the simulation of multipurpose detectors for physics studies [29]. Specifically, the DELPHES
 626 CMS configuration is frequently used as the detector simulation of choice in recent machine-
 627 learning based unfolding studies.

628 For the toy model studies presented in Section 2.3, the detector effects were simulated
 629 using the same resolution functions as the data-driven approach described below, applying
 630 the same smearing to the kinematic quantities p_T , ϕ , and η . However, in the toy model case,
 631 the jets were treated as massless particles for simplicity, with their energy calculated directly
 632 from the smeared momentum components.

633 To test the unfolding performance under more exaggerated detector effects, we develop a
 634 framework with a data-driven detector smearing using jet energy resolution results published
 635 by the ATLAS collaboration at a centre-of-mass energy of 8 TeV with an integrated luminosity
 636 of 20 fb^{-1} [31]. In this framework, the PYTHIA event generator is used to simulate truth-level
 637 particles, and the resulting partons are grouped into jets using the FastJet package [34]. The
 638 transverse momentum p_T , azimuthal angle ϕ , and pseudorapidity η of each truth-level jet is
 639 then smeared following an approximated ATLAS calibration and resolution functions. For the
 640 ϕ and η smearing, the effect is small since the angular resolution effects are proportional to
 641 the detector granularity. We assume that there is no angular shift and apply a smearing to ϕ
 642 and η by sampling from a Gaussian centered at the truth-level value and with a σ equal to
 643 the detector resolution for the particle. We apply a quadratic fit ($\sigma = a p_T^2 + b p_T + c$) to the
 644 calibration data presented in [31] to approximate the detector resolution in ϕ and η .

645 In principle, a calorimeter cell measurement is an energy measurement, but since the jet
 646 calibration studies precisely measure the jet p_T resolution, we apply a shift and a smearing to
 647 the jet p_T instead. The jet p_T resolution can be expressed as $\sigma_{p_T} = p_T \sqrt{a/p_T^2 + b/p_T + c}$ and

Table 3: List of physics simulations generated for the data-driven detector smearing, along with the corresponding parton distribution functions (PDFs), parton shower models, phase space biases, and their inclusion in the training dataset. Simulations not included in the training dataset are used as test datasets.

Process	PDF with Parton Shower (Phase Space Bias)	In Training?
$t\bar{t}$	CT14lo	✓
	CT14lo (biased)	✓
	CT14lo with Vincia	
	NNPDF23_lo	✓
	CTEQ6L1	✓
	CTEQ6L1 (biased)	✓
Z+jets	CT14lo	✓
	CT14lo (biased)	✓
	NNPDF23_lo	✓
	CTEQ6L1	
	CTEQ6L1 (biased)	✓
W+jets	CT14lo	
	CT14lo (biased)	✓
	NNPDF23_lo	✓
	CTEQ6L1	✓
Dijets	CT14lo	✓
	CTEQ6L1	✓
	CTEQ6L1 (biased)	✓
Leptoquark	CT14lo	✓
	CT14lo (biased)	✓
	NNPDF23_lo	
	CTEQ6L1	✓

648 a fit of this function is applied to the jet calibration data (obtained for jets with $0 < |\eta| < 0.8$
649 for simplicity) to approximate this resolution. The jet p_T also has a calibration shift, which is
650 calculated from the data. The detector smeared p_T is then defined by sampling from a Gaus-
651 sian centered at the shifted p_T and with σ equal to the jet p_T resolution. Finally, the smeared
652 energy for each jet is calculated with $E = \sqrt{m^2 + |\vec{p}|^2}$ by fixing the mass of the particle m and
653 using the smeared p_T . This approach enables testing the unfolding algorithm's performance
654 under various detector resolution conditions, including unrealistically large smearing effects.
655 While each detector model requires retraining the unfolders, the algorithm's performance char-
656 acteristics remain consistent across different detector modelings.

Table 4: List of physics simulations generated for the DELPHES CMS detector simulation, along with the corresponding parton distribution functions (PDFs), phase space biases, and their inclusion in the training dataset. Simulations not included in the training dataset are used as test datasets.

Process	PDF (Phase Space Bias)	In Training?
$t\bar{t}$	CTEQ6L1	
	CTEQ6L1 (biased)	✓
Z+jets	CTEQ6L1	✓
	CTEQ6L1 (biased)	✓
W+jets	CTEQ6L1	
	CTEQ6L1 (biased)	✓
Dijets	CTEQ6L1	✓
	CTEQ6L1 (biased)	✓
Leptoquark	CTEQ6L1	

657 References

- 658 [1] V. Blobel, *An unfolding method for high energy physics experiments* (2002), [hep-ex/0208022](#).
659
- 660 [2] G. D’Agostini, *Improved iterative bayesian unfolding* (2010), [1010.0632](#).
- 661 [3] S. Schmitt, *Tunfold, an algorithm for correcting migration effects in high energy physics*, *Journal of Instrumentation* 7(10), T10003 (2012), [doi:10.1088/1748-0221/7/10/T10003](#).
662
663
- 664 [4] A. Höcker and V. Kartvelishvili, *Svd approach to data unfolding*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 372(3), 469 (1996), [doi:https://doi.org/10.1016/0168-9002\(95\)01478-0](#).
665
666
667
- 668 [5] A. Andreassen, P. T. Komiske, E. M. Metodiev, B. Nachman and J. Thaler, *Omnifold: A method to simultaneously unfold all observables*, *Phys. Rev. Lett.* 124, 182001 (2020), [doi:10.1103/PhysRevLett.124.182001](#).
669
670
- 671 [6] A. Andreassen, P. T. Komiske, E. M. Metodiev, B. Nachman, A. Suresh and J. Thaler, *Scaffolding simulations with deep learning for high-dimensional deconvolution* (2021), [2105.04448](#).
672
673
- 674 [7] M. Bellagente, A. Butter, G. Kasieczka, T. Plehn and R. Winterhalder, *How to gan away detector effects*, *SciPost Physics* 8(4) (2020), [doi:10.21468/scipostphys.8.4.070](#).
675
- 676 [8] K. Datta, D. Kar and D. Roy, *Unfolding with generative adversarial networks* (2018), [1806.00433](#).
677
- 678 [9] M. Backes, A. Butter, M. Dunford and B. Malaescu, *An unfolding method based on conditional invertible neural networks (cinn) using iterative training* (2024), [2212.08674](#).
679
- 680 [10] T. Heimel, N. Huetsch, R. Winterhalder, T. Plehn and A. Butter, *Precision-machine learning for the matrix element method* (2024), [2310.07752](#).
681

- 682 [11] M. Bellagente, A. Butter, G. Kasieczka, T. Plehn, A. Rousselot, R. Winterhalder, L. Ardizzone
683 and U. Köthe, *Invertible networks or partons to detector and back again*, SciPost
684 Physics **9**(5) (2020), doi:[10.21468/scipostphys.9.5.074](https://doi.org/10.21468/scipostphys.9.5.074).
- 685 [12] A. Shmakov, K. Greif, M. Fenton, A. Ghosh, P. Baldi and D. Whiteson, *End-to-end
686 latent variational diffusion models for inverse problems in high energy physics* (2023),
687 [2305.10399](https://arxiv.org/abs/2305.10399).
- 688 [13] A. Shmakov, K. Greif, M. J. Fenton, A. Ghosh, P. Baldi and D. Whiteson, *Full event particle-
689 level unfolding with variable-length latent variational diffusion* (2024), [2404.14332](https://arxiv.org/abs/2404.14332).
- 690 [14] S. Diefenbacher, G.-H. Liu, V. Mikuni, B. Nachman and W. Nie, *Improving generative
691 model-based unfolding with schrödinger bridges* (2023), [2308.12351](https://arxiv.org/abs/2308.12351).
- 692 [15] A. Butter, T. Jezo, M. Klasen, M. Kuschick, S. P. Schweitzer and T. Plehn, *Kicking it off(-
693 shell) with direct diffusion* (2024), [2311.17175](https://arxiv.org/abs/2311.17175).
- 694 [16] N. Huetsch, J. M. Villadamigo, A. Shmakov, S. Diefenbacher, V. Mikuni, T. Heimel, M. Fen-
695 ton, K. Greif, B. Nachman, D. Whiteson, A. Butter and T. Plehn, *The landscape of unfolding
696 with machine learning* (2024), [2404.18807](https://arxiv.org/abs/2404.18807).
- 697 [17] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, Journal
698 of Instrumentation **3**(08), S08003 (2008), doi:[10.1088/1748-0221/3/08/S08003](https://doi.org/10.1088/1748-0221/3/08/S08003).
- 699 [18] ATLAS Collaboration, *Measurement of differential cross sections and W^+/W^- cross-section
700 ratios for W boson production in association with jets at $\sqrt{s} = 8$ TeV with the ATLAS detector*,
701 Journal of High Energy Physics **2018**(5) (2018), doi:[10.1007/jhep05\(2018\)077](https://doi.org/10.1007/jhep05(2018)077).
- 702 [19] L. Evans and P. Bryant, *Lhc machine*, JINST **3**(08), S08001 (2008), doi:[10.1088/1748-
703 0221/3/08/S08001](https://doi.org/10.1088/1748-0221/3/08/S08001).
- 704 [20] ATLAS Collaboration, *Relative uncertainty for $W(\rightarrow e\nu) + \geq 1$ jets as a function of HT* ,
705 ATLAS Public Results (2023).
- 706 [21] S. Ask, I. Akin, L. Benucci, A. De Roeck, M. Goebel and J. Haller, *Real emission and virtual
707 exchange of gravitons and unparticles in pythia8*, Computer Physics Communications
708 **181**(9), 1593–1604 (2010), doi:[10.1016/j.cpc.2010.05.013](https://doi.org/10.1016/j.cpc.2010.05.013).
- 709 [22] T. M. Mitchell, *The need for biases in learning generalizations*, Technical Report CBM-TR-
710 117, Rutgers University, New Brunswick, NJ (1980).
- 711 [23] J. Ho, A. Jain and P. Abbeel, *Denoising diffusion probabilistic models* (2020), [2006.11239](https://arxiv.org/abs/2006.11239).
- 712 [24] J. Choi, S. Kim, Y. Jeong, Y. Gwon and S. Yoon, *Ilvr: Conditioning method for denoising
713 diffusion probabilistic models* (2021), [2108.02938](https://arxiv.org/abs/2108.02938).
- 714 [25] P. Dhariwal and A. Nichol, *Diffusion models beat gans on image synthesis* (2021), [2105.
715 05233](https://arxiv.org/abs/2105.05233).
- 716 [26] J. Ho and T. Salimans, *Classifier-free diffusion guidance* (2022), [2207.12598](https://arxiv.org/abs/2207.12598).
- 717 [27] C. Villani, *The wasserstein distances*, doi:[10.1007/978-3-540-71050-9_6](https://doi.org/10.1007/978-3-540-71050-9_6) (2009).
- 718 [28] C. Bierlich, S. Chakraborty, N. Desai, L. Gellersen, I. Helenius, P. Ilten, L. Lönnblad,
719 S. Mrenna, S. Prestel, C. T. Preuss, T. Sjöstrand, P. Skands *et al.*, *A comprehensive guide
720 to the physics and usage of pythia 8.3* (2022), [2203.11601](https://arxiv.org/abs/2203.11601).

- 721 [29] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens and M. Selvaggi, *Delphes 3: a modular framework for fast simulation of a generic collider experiment*,
722 Journal of High Energy Physics **2014**(2) (2014), doi:[10.1007/jhep02\(2014\)057](https://doi.org/10.1007/jhep02(2014)057).
723
- 724 [30] S. Chatrchyan *et al.*, *The CMS experiment at the CERN LHC*, JINST **3**, S08004 (2008),
725 doi:[10.1088/1748-0221/3/08/S08004](https://doi.org/10.1088/1748-0221/3/08/S08004).
- 726 [31] ATLAS Collaboration, *Determination of jet calibration and energy resolution in proton–proton collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector*, The European Physical
727 Journal C **80**(12) (2020), doi:[10.1140/epjc/s10052-020-08477-8](https://doi.org/10.1140/epjc/s10052-020-08477-8).
728
- 729 [32] W. T. Giele, L. Hartgring, D. A. Kosower, E. Laenen, A. J. Larkoski, J. J. Lopez-Villarejo,
730 M. Ritzmann and P. Skands, *The Vincia Parton Shower*, PoS DIS2013, 165 (2013),
731 doi:[10.22323/1.191.0165](https://doi.org/10.22323/1.191.0165), [1307.1060](https://doi.org/10.1307/1060).
- 732 [33] S. Kullback and R. A. Leibler, *On Information and Sufficiency*, The Annals of Mathematical
733 Statistics **22**(1), 79 (1951), doi:[10.1214/aoms/1177729694](https://doi.org/10.1214/aoms/1177729694).
- 734 [34] M. Cacciari, G. P. Salam and G. Soyez, *Fastjet user manual: (for version 3.0.2)*, The
735 European Physical Journal C **72**(3) (2012), doi:[10.1140/epjc/s10052-012-1896-2](https://doi.org/10.1140/epjc/s10052-012-1896-2).