

Fast and improved neutrino reconstruction in multineutrino final states with conditional normalizing flows

John Andrew Raine^{1,*} Matthew Leigh^{1,†} Knut Zoch^{1,2} and Tobias Golling¹

¹*Département de physique nucléaire et corpusculaire, University of Geneva, 1211 Geneva, Switzerland*

²*Laboratory for Particle Physics and Cosmology, Harvard University, Cambridge, 02138 Massachusetts, USA*

 (Received 20 July 2023; accepted 16 November 2023; published 11 January 2024)

In this work we introduce ν^2 -flows, an extension of the ν -flows method to final states containing multiple neutrinos. The architecture can natively scale for all combinations of object types and multiplicities in the final state for any desired neutrino multiplicities. In $t\bar{t}$ dilepton events, the momenta of both neutrinos and correlations between them are reconstructed more accurately than when using the most popular standard analytical techniques, and solutions are found for all events. Inference time is significantly faster than competing methods, and can be reduced further by evaluating in parallel on graphics processing units. We apply ν^2 -flows to $t\bar{t}$ dilepton events and show that the per-bin uncertainties in unfolded distributions is much closer to the limit of performance set by perfect neutrino reconstruction than standard techniques. For the chosen double differential observables ν^2 -flows results in improved statistical precision for each bin by a factor of 1.5 to 2 in comparison to the neutrino weighting method and up to a factor of four in comparison to the ellipse approach.

DOI: 10.1103/PhysRevD.109.012005

I. INTRODUCTION

At collider experiments in particle physics, such as those at the Large Hadron Collider (LHC) [1], beams of hadrons or leptons are accelerated to high energies and collided together. These collisions result in an array of particles which are studied by experiments comprising complex detectors surrounding the interaction points. General purpose detectors, such as ATLAS [2] and CMS [3], are designed to record and reconstruct nearly all stable particles predicted in the standard model of particle physics (SM). From these reconstructed stable particles, both precision measurements of the SM as well as searches for new phenomena beyond the SM (BSM) are performed.

Neutrinos, stable particles produced in many collisions, interact only through the electroweak force and traverse the detectors without leaving a trace. Their presence in collisions is inferred from a momentum imbalance in the transverse plane perpendicular to the beam axis. This imbalance, known as the missing transverse momentum \vec{p}_T^{miss} , is calculated from the negative vector sum of the

transverse momenta of all reconstructed objects in the transverse plane.

In order to reconstruct individual neutrinos, and thus fully reconstruct a single event, underlying assumptions need to be made on their origin and multiplicity. The \vec{p}_T^{miss} serves as a proxy for all unobserved particles in the collision, but does not indicate how many were present, or how the momentum should be shared in the case of multiple neutrinos. Furthermore, at hadron colliders there is no experimental equivalent for the missing longitudinal momentum.

Several approaches are used to reconstruct neutrinos using \vec{p}_T^{miss} and by setting constraints on the invariant masses of intermediate particles. In the case of top-quark pair production ($t\bar{t}$) in the semileptonic decay channel, where there is one lepton and one neutrino in the event, the neutrino momentum is typically reconstructed by solving the longitudinal momentum component p_z under the assumption that the invariant mass of the lepton-neutrino pair is exactly that of the W boson [4–12]. For the case of events with two neutrinos, such as $t\bar{t}$ production in the dileptonic channel, more complicated methods [13–15] are employed in order to resolve the share of momentum between the two neutrinos in the event [12,16–28]. These approaches still require the invariant mass of the two lepton-neutrino pairs to match the W boson mass, and in addition require the invariant mass of the lepton-neutrino-jet triplets to match the top-quark invariant mass.

In ν -flows [29] we use conditional normalizing flows [30–32] to learn the probability distribution of the neutrino

*john.raine@unige.ch
†matthew.leigh@unige.ch

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

momentum vector given the observed objects in an event for semileptonic $t\bar{t}$ events. From the learned conditional probability distribution, solutions can be sampled per event and the most probable solution can be determined from the learned likelihood of the solution. In this work, we extend ν -flows to more challenging final states with multiple neutrinos. We focus on $t\bar{t}$ final states in which both top quarks decay semileptonically. This results in an expected final state with exactly two oppositely charged leptons, at least two jets of which two should originate from b -quarks, and two neutrinos. In comparison to the single neutrino case, where the main challenge arises in recovering the longitudinal momentum component of the neutrino, events with multiple neutrinos have the additional complexity of how to share the total missing momentum vector between all neutrinos in the final state. We show that ν -flows yields improved reconstruction performance in comparison to standard approaches and demonstrate the direct benefit to the statistical precision in a simplified double differential $t\bar{t}$ dilepton analysis in observables such as the invariant mass of the $t\bar{t}$ system, $m_{t\bar{t}}$, and the angular separation of the two leptons, $\Delta\phi(\ell^+\ell^-)$. The repository¹ and data² used in this work are both made publicly available.

II. METHOD

A. ν^2 -flows

In ν -flows, conditional normalizing flows are used to capture the distribution of possible solutions for the neutrino momenta given the reconstructed momenta of observed objects in a collision. The overall model comprises two components, the event feature extraction, and the conditional normalizing flow. The event feature extraction learns a representative vector of the event from the collections of reconstructed objects, namely the jets and leptons, and \vec{p}_T^{miss} . The resulting vector is used as a conditional input to the normalizing flow, in order to learn the conditional density of all possible neutrino solutions from the training data.

In this work we extend the initial ν -flows architecture to predict multiple neutrinos and accommodate any number of leptons in addition to jets by using attention transformers [34]. We label this architecture as ν^2 -flows to distinguish it from the general method. In order to handle two neutrinos we double the dimensionality of the conditional normalizing flow (from three to six). The neutrinos are also always predicted in the same order for each event, with the momentum of the neutrino followed by the momentum of the antineutrino. When increasing the neutrino multiplicity further, the same procedure is used together with a predefined ordering for the neutrinos. The architecture of the normalizing flow is otherwise kept largely the same.

TABLE I. The different input observables used as inputs to the feature extraction network.

Category	Variables	Description
\vec{p}_T^{miss}	$p_x^{\text{miss}}, p_y^{\text{miss}}$	Missing transverse momentum 2-vector
Leptons	$p_x^\ell, p_y^\ell, p_z^\ell, \log E^\ell$ q^ℓ ℓ^{flav}	Lepton momentum 4-vector Lepton charge Whether lepton is an electron or muon
Jets	$p_x^j, p_y^j, p_z^j, \log E^j$ $i\text{is}B$	Jet momentum 4-vector Whether jet passes b -tagging criteria
Misc	$N_{\text{jets}}, N_{\text{bjets}}$	Jet and b -jet multiplicities in the event

The most substantial optimization has been performed on the feature extraction network. The feature extraction network attempts to produce a contextual vector, specific to each event, to guide the transformations within the normalizing flow. In the single lepton case, ν -flows uses an attention pooled deep set [35] to process the jets, with $\vec{p}_x^{\text{miss}}, \vec{p}_y^{\text{miss}}$, the lepton four momentum, and some event level information as extra conditional information. To extend ν^2 -flows to multiple leptons, we require a permutation invariant architecture that can accommodate a variable number of both jets and leptons, motivating the move to attention transformers.

To train ν^2 -flows, all jets and leptons are represented by their four-momentum vectors in the form $(p_x, p_y, p_z, \log E)$. Jets are assigned an additional binary decision on whether they are tagged as originating from a b -quark. Leptons are identified as being either an electron or a muon, as well as whether they had positive or negative charge. The target neutrino momenta are expressed as (p_x, p_y, p_z) for both the neutrino and antineutrino. The full set of inputs to the network are provided in Table I. The coordinates chosen to describe the input and target object kinematics were optimized in a grid search.

A schematic of the new architecture for ν^2 -flows is shown in Fig. 1, which makes use of attention transformers and object specific embedding networks. Initially, the jets, leptons, and \vec{p}_T^{miss} are all independently embedded into higher-dimensional space using object specific multilayer perceptrons (MLP).³ The embedded objects subsequently interact through a transformer encoder using several layers of multiheaded attention. Additional event information (Misc) containing object multiplicities, is injected into the network as conditional information by concatenating

¹<https://github.com/rodem-hep/nu2flows>.

²<https://zenodo.org/record/8113516> [33].

³Other final state objects such as photons and tau leptons can also be accommodated by embedding them additional MLPs for each particle type.

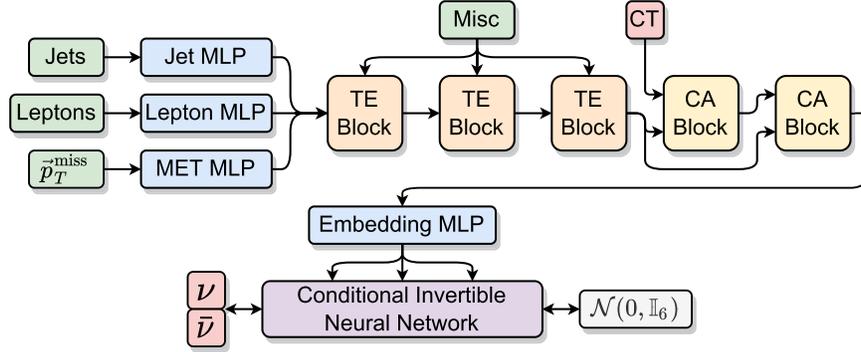


FIG. 1. A schematic of the ν^2 -flows network for learning the conditional likelihood of multiple neutrinos in the event. The network uses a transformer encoder with cross attention with a learnable class token (shown in the figure as TE, CA, and CT respectively) to embed an event representation for any multiplicity of physics objects. This operation is permutation invariant and can operate on any jet and lepton multiplicity. Each physics object has its own dedicated embedding network and additional event information (Misc) is used to condition the transformer encoder blocks. The representation vector is used to condition the transformation with the normalizing flow.

the vector to each token within the transformer encoder (TE) blocks. To obtain a single global vector from transformer as our conditioning vector for the normalizing flow, we perform repeated cross-attention (CA) with a learnable class token (CT), a common technique used in vision transformers [36].

The ν^2 -flows model in this paper comprises three transformer encoder blocks and two cross-attention blocks, each with an embedding dimension of 128 and 16 attention heads. All MLPs in the network have a single hidden layer with 256 neurons and use the LeakyReLU activation [37] and layer-normalization [38]. The output of the transformer is passed through an MLP to produce the context tensor for the flow with a dimension of 128. The invertible neural network employs 10 rational quadratic spline (RQS) coupling blocks [39] interspersed with LU-decomposed linear layers, implemented with the `nflows` package [40] and PyTorch v2.0 [41]. Each RQS has 10 bins with linear tail bounds outside ± 4 . The conditional normalizing flow is trained with the standard maximum likelihood estimation loss obtained through the change of variables formula, and transforms the input neutrino momenta to a standard multivariate normal distribution. The entire ν^2 -flows model, including the transformer, has around 600,000 trainable parameters.

We train the ν^2 -flows model for 100 epochs using the AdamW optimizer [42] with a learning rate cycling from 10^{-8} to 10^{-3} and back every 50 epochs. We use weight decay with a strength of 10^{-4} .

B. Reference methods

Several analytical techniques have been proposed to solve the reconstruction of the two neutrinos in dilepton $t\bar{t}$ events. Amongst these are neutrino weighting [13] (ν -weighting), an algebraic solution [14], and the ellipse method [15]. These have been successfully employed in a wide range of measurements at the Tevatron and

LHC [12,16–28], most notably ν -weighting. In this work we compare ν^2 -flows to the ν -weighting due to its common usage but also the ellipse method due to its reduced computation time.

I. ν -weighting

In ν -weighting the kinematic properties of the neutrino are extracted from the identified leptons, jets, and missing transverse momentum in the event. For our implementation we follow the prescription described in Ref. [27]. Constraints on neutrino solutions using the invariant mass of the top quark and W boson

$$(\ell_{1,2} + \nu_{1,2})^2 = m_w^2 = (80.38 \text{ GeV})^2, \quad (1)$$

$$(\ell_{1,2} + \nu_{1,2} + b_{1,2})^2 = m_t^2 = (172.5 \text{ GeV})^2, \quad (2)$$

are applied, where $\ell_{1,2}$, $\nu_{1,2}$, and $b_{1,2}$ represent the four-momenta of the charged leptons, neutrinos, and b -tagged jets. However, this is not enough to fully constrain the kinematics of the neutrinos. Therefore, the neutrino and antineutrino rapidities (η^ν and $\eta^{\bar{\nu}}$) are individually hypothesized and tested. For each pair of values for η^ν and $\eta^{\bar{\nu}}$, we solve the mass equations in Eq. (1), yielding two possible solutions for the full pair of neutrino kinematics. Each solution produces an inferred missing transverse momentum $\vec{p}_T^{\nu\bar{\nu}}$ vector which can then be compared to the observed \vec{p}_T^{miss} . This comparison defines a weight

$$w = \exp\left(-\frac{\|\vec{p}_T^{\text{miss}} - \vec{p}_T^{\nu\bar{\nu}}\|_2^2}{2\sigma^2}\right), \quad (3)$$

where σ is a fixed resolution scale related to the \vec{p}_T^{miss} reconstruction in the detector. In ν -weighting the hypothesis that maximizes w is chosen as the correct solution. To find all solutions we perform a grid search of η^ν and $\eta^{\bar{\nu}}$ with values ranging from -5 to 5 with a step

size of 0.2. For each of these we also need to test all combinations of assigning b -tagged jets to each of the b -quarks from the $t\bar{t}$ decay. This represents a very costly computation scan that considers only discrete η values.

Another significant drawback of this method is that despite the large number of neutrino solutions being tested, it is still possible that the constraint systems are not solvable. This can be due to mismeasurement or misassignment of the leptons, jets, or \vec{p}_T^{miss} . Alternatively, this can arise from the masses of either the top quarks or W bosons in the event deviating from the nominal values. Therefore, to increase the success rate of this method we also iterate over different values of m_t from 171 GeV to 174 GeV with a step size of 0.5 GeV. This further increases the computational requirements, but increases the efficiency of finding a solution for each event.

After selecting the solution with the highest w , any solution which results in $m_{t\bar{t}} < 300$ GeV or where either of the two reconstructed top quarks have negative energy is rejected. The ν -weighting method is unable to find a valid solution on the nominal dataset around 5% of the time.

2. Ellipse method

The ellipse method is derived from a geometric approach to analytically constrain neutrino kinematics explicitly in processes where top quarks decay into leptons and neutrinos [15]. For a single neutrino, its momentum can be calculated as a function of the 4-momenta of the b -quark and the charged lepton, the W boson mass, and the top quark mass. The solution set for this function defines the surface of an ellipse. By combining this information with the observed \vec{p}_T^{miss} , the solution set collapses to a unique value. For events with two neutrinos in the final state, the method is extended to calculate the solution for neutrino pairs which are most likely to have produced the observed \vec{p}_T^{miss} .

We use the implementation from the authors of the ellipse method.⁴ To solve the b -jet combinatorics we use a simple minimum ΔR matching between the leptons and the b -jets in the event. To reduce the combinations we only take the two leading jets in p_T -passing b -tagging criteria. If the ΔR matching yields no solutions for the neutrino kinematics using the ellipse method, the opposite association is tested.

The drawbacks of this approach is that it requires accurate matching between each lepton and the associated b -jet in the event. Furthermore, like ν -weighting, it requires one to make hard assumptions on the mass of the W boson and top quark. Finally, it is possible that the method can yield no solutions, just like ν -weighting. The implementation used in this work fails to find solutions in 22% of $t\bar{t}$ dilepton events. In comparison to ν -weighting, the ellipse

method requires much less computational resources per event.

III. DATASET

In this work, ν^2 -flows is applied to simulated $t\bar{t}$ events where both top quarks decay semileptonically, resulting in a final state with exactly two leptons (ℓ),⁵ two neutrinos (ν) and two jets initiated by b -hadrons (b -jets). Additional jets arise from initial and final state radiation. All events are simulated in proton-proton collisions at a center-of-mass energy of $\sqrt{s} = 13$ TeV. Two different samples are generated, each using a different generator for the hard interactions in the matrix element.

In the nominal sample, hard interactions are simulated using MadGraph5_aMC@NLO [43] (v3.1.0). The top-quark mass m_t is set to 173 GeV for all events. All events are interfaced to PYTHIA8 [44] (v8.243) for the parton shower and hadronization, using the NNPDF2.3LO PDF set [45] with $\alpha_S(m_Z) = 0.130$ using the LHAPDF [46] framework.

In the alternative sample, both the hard interactions and parton shower are simulated with PYTHIA8 (v8.307) with the Monash tuned set of parameters [47] at leading-order accuracy. The same PDF set is used as for the nominal sample.

The detector response is simulated using DELPHES [48] (v3.4.2) with a parametrization similar to the response of the ATLAS detector [2]. Jets are reconstructed using the anti- k_r clustering algorithm [49] with a radius parameter of $R = 0.4$ using the FastJet package [50]. All jets are required to have a transverse momentum $p_T > 25$ GeV and fall within $|\eta| < 2.5$. A b -tagging working point corresponding to 70% inclusive signal efficiency is used to identify jets as originating from b -hadrons. Up to 10 jets are selected per event, ordered in descending p_T .

Events are required to have at least two b -tagged jets, and two oppositely charged leptons, each with $p_T > 15$ GeV and $|\eta| < 2.5$. Truth association of jets to the b -quarks in the $t\bar{t}$ hard scatter is performed using a ΔR matching, with partons matched to jets within $\Delta R < 0.4$. Events where multiple partons are matched to the same jet are removed. Truth association of the lepton to the parent top quark is performed assuming there is no charge misidentification, and the true neutrino momenta are taken directly from simulation.

In total there are 1.02 million events in the nominal sample and 1.4 million events in the alternative sample passing all selection requirements. 940,000 (970,000) events from the nominal (alternative) sample are used to train the network, with 80,000 nominal samples used for evaluation.

⁴Implementation available at github.com/betchart/analytic-nu.

⁵Leptons is used to designate either electrons (e) or muons (μ).

IV. RESULTS

A. Neutrino reconstruction

The first measure of performance is to validate that ν^2 -flows is able to correctly reconstruct the momenta and relative positions of the neutrino pair. To establish a baseline for the impact of neutrino reconstruction in all distributions, we define ν -truth to be the case where the truth neutrinos are reconstructed perfectly, but all other objects in the event remain the same. This defines the ground truth of the target distributions and is also the upper limit in performance for any neutrino reconstruction

approach. We compare the reconstruction performance with ν^2 -flows to the ν -weighting and ellipse methods.

The individual neutrino kinematics and the angular separation between the two neutrinos are shown in Figs. 2 and 3. Here we can see that both ν -weighting and ellipse overestimate the neutrino transverse momenta and energy and tend to prefer central neutrinos with a visible peak at $\eta = 0$. The neutrino pair is also predominantly predicted to be back to back in ϕ , with the opening angle between them showing a large degree of tension with the ground truth. In all neutrino kinematic distributions, ν^2 -flows is able to

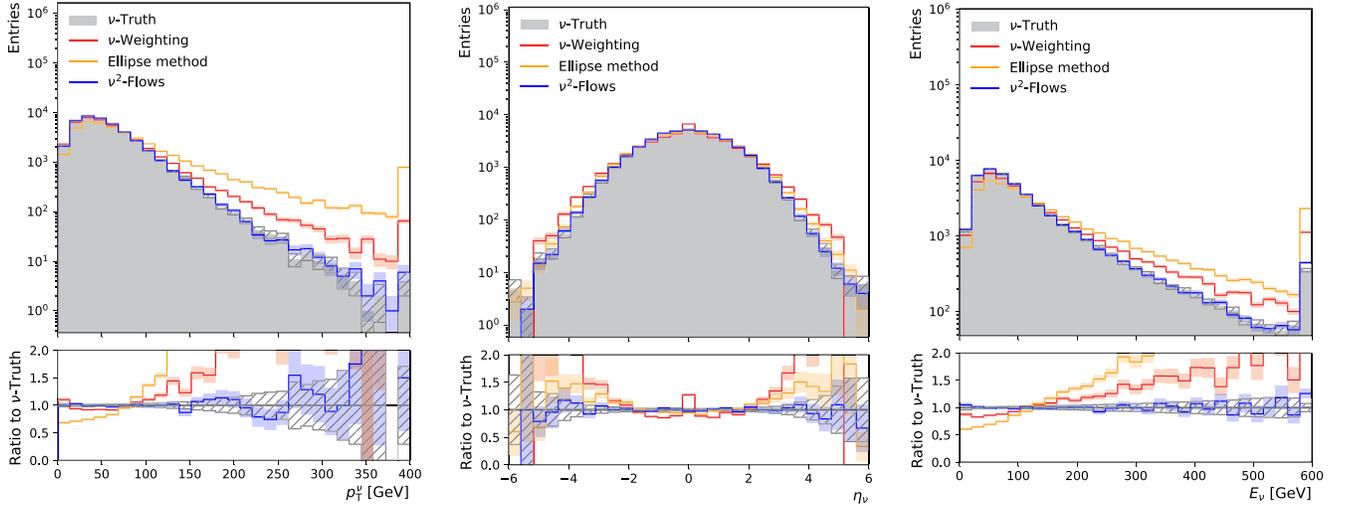


FIG. 2. The kinematics of the reconstructed (anti)neutrinos for the three reconstruction methods and ν -truth (shaded gray). The hashed areas represent statistical uncertainties in the ν -truth prediction.

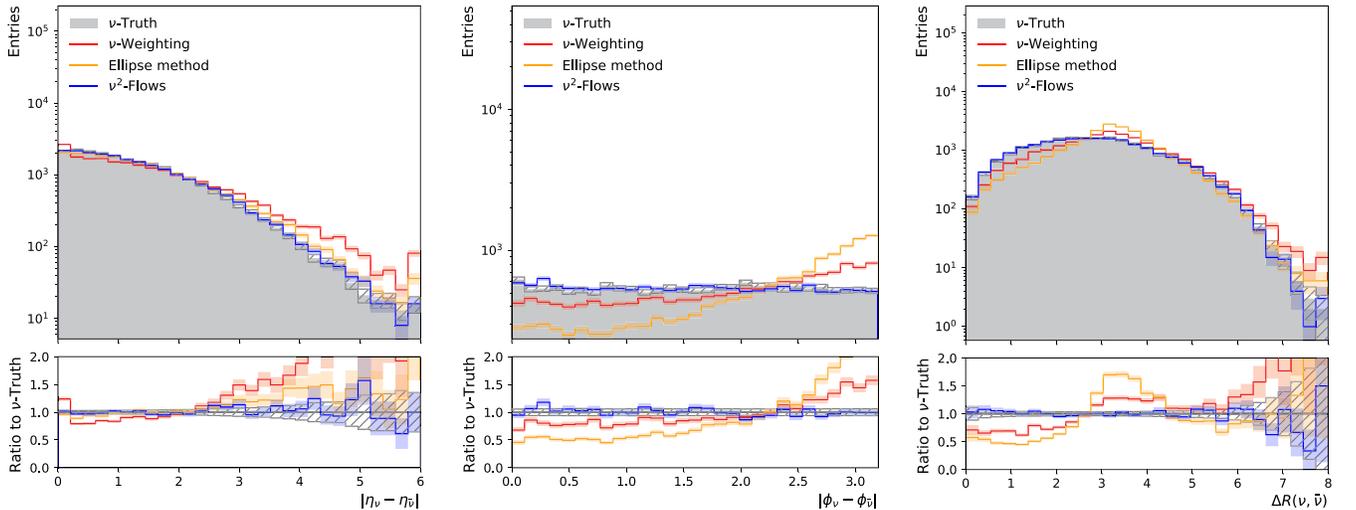


FIG. 3. The angular separation in η and ϕ between the reconstructed neutrino pair per event for the three reconstruction methods and ν -truth (shaded gray). The hashed areas represent statistical uncertainties in the ν -truth prediction.

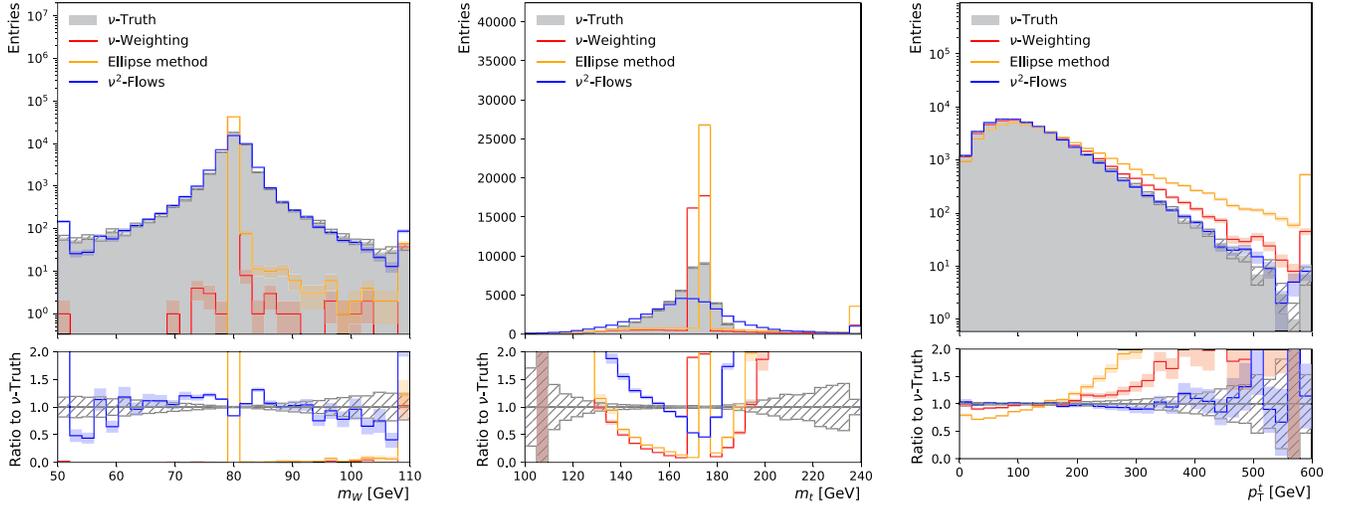


FIG. 4. The reconstructed invariant mass of W bosons (left) and top quarks (middle), as well as the top quark p_T (right) when using the three neutrino reconstruction methods in comparison to ν -truth (shaded gray).

reproduce the ground truth accurately, with slight discrepancies only visible in the low statistic tails.

From the reconstructed neutrinos it is also possible to reconstruct the W bosons, top quarks and the full $t\bar{t}$ system in the event. Figure 4 shows the reconstructed invariant mass of the W boson and top quark. Here we always use the perfect association of jets and leptons to the two top quarks. The strong bias to the top quark and W boson masses used in Eq. (1) for the ν -weighting and ellipse methods are clearly visible. In comparison, ν^2 -flows follows the full underlying target distribution for the W boson mass and better captures the distribution in the reconstructed top quark mass, but with a lower resolution than ν -truth.

However, due to the dependence on the masses of the two particles for ν -weighting and ellipse it is more

interesting to compare the reconstructed kinematics of the individual top quarks and the $t\bar{t}$ pair. The reconstructed top quark p_T is also compared for the three methods in Fig. 4. The true distribution is reproduced with ν^2 -flows and up to approximately 200 GeV by ν -weighting. However, ellipse tends to reconstruct top quarks with a harder p_T . In Fig. 5 the invariant mass, p_T and the rapidity of the reconstructed $t\bar{t}$ system $y_{t\bar{t}}$ are shown. Although discrepancies are observed at low $m_{t\bar{t}}$ values, ν^2 -flows is able to closely reproduce the kinematics of the $t\bar{t}$ system much better than ν -weighting and ellipse. The $p_T^{t\bar{t}}$ and $y_{t\bar{t}}$ distributions are well reconstructed with ν -weighting but there is an overestimate in the tail of the $m_{t\bar{t}}$ distribution, whereas ellipse shows poor modeling in all three observables.

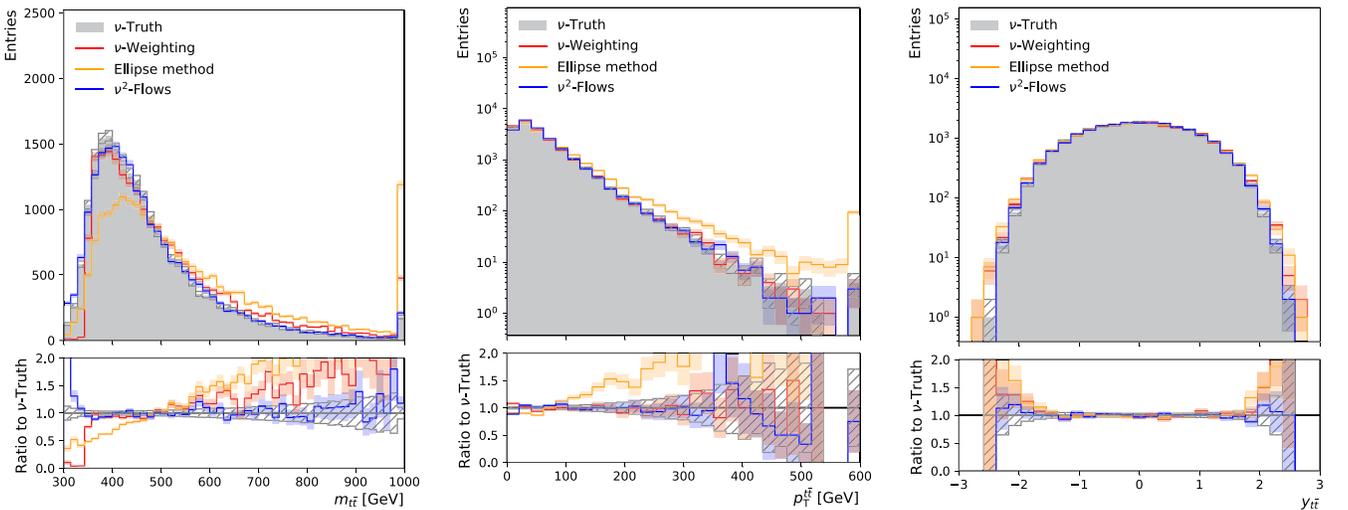


FIG. 5. The invariant mass, p_T , and rapidity of the reconstructed $t\bar{t}$ system when using the three neutrino reconstruction methods in comparison to ν -truth (shaded gray).

TABLE II. Required time for single event inference using ν^2 -flows. Times representative of using a single core of an AMD EPYC 7742 2.25 GHz CPU and a NVIDIA[®] RTX 3080 graphics card.

Resource	Batch size	Time/event [ms]
CPU	1	71.4
GPU	1	33.3
	1000	0.03

As also performed in Ref. [29], we perform a cross check on the benefit of using the normalizing flow in the ν^2 -flows architecture. We train the ν^2 -flows architecture but without the flow and predict the two neutrino momenta directly, establishing a simple machine learning baseline. The performance achieved is substantially worse with strong biases in all neutrino kinematics and resulting event level distributions (see the Appendix).

One of the main drawbacks of ν -weighting, and why ellipse is often considered despite the reduced performance, is the computational resources required. In comparison, ν -flows requires only a single forward pass for each event. The typical inference times on a CPU for single event inference are around 70 ms, with the computation time decreasing substantially with parallelized execution on a GPU as summarized in Table II.

B. Unfolding analysis

In order to evaluate the downstream impact of the improved neutrino reconstruction from ν -flows, we follow the unfolding analysis performed in Ref. [26], where a double differential cross section measurement is performed to measure the spin correlation in $t\bar{t}$ events, by measuring the invariant mass of the $t\bar{t}$ system $m_{t\bar{t}}$ and the angular separation between the two leptons $\Delta\phi(\ell^+\ell^-)$. Reconstruction of the two neutrinos is crucial in order to fully reconstruct the $t\bar{t}$ system, which in Ref. [26] is performed using ν -weighting. To benchmark our model, we replace the neutrino reconstruction with the result from ν^2 -flows. In addition to ν -weighting we also compare the performance to the ellipse method due to its reduced computational complexity.

TABLE III. Kinematic observables of the reconstructed $t\bar{t}$ system studied for an unfolding analysis in dilepton events.

Observables	
$m_{t\bar{t}}$	Invariant mass of $t\bar{t}$ system
$\Delta\phi(\ell^+\ell^-)$	Separation in ϕ between the two leptons
p_T^t	Transverse momentum of the top quark
$p_T^{\bar{t}}$	Transverse momentum of the $t\bar{t}$ system
$y_{t\bar{t}}$	Rapidity of the $t\bar{t}$ system

TABLE IV. Bin edges used for each of the kinematic observables of the reconstructed $t\bar{t}$ system studied for two-dimensional unfolding analyses.

Observable	Bin edges
$m_{t\bar{t}}$	[0, 400, 500, 800, inf] GeV
$\Delta\phi(\ell^+\ell^-)$	[0.0, 0.25, 0.5, 0.75, 1.0] rad/ π
p_T^t	[0, 75, 125, 175, inf] GeV
$p_T^{\bar{t}}$	[0, 70, 140, 200, inf] GeV
$y_{t\bar{t}}$	[-inf, -1.0, 0.0, 1.0, inf]

We focus on the reconstruction of individual observables dependent on the neutrino kinematics as well as the statistical precision of the unfolded distributions. In addition to $\Delta\phi(\ell^+\ell^-)$, we look at other observables in conjunction with $m_{t\bar{t}}$, motivated by the distributions measured in Ref. [51]. These observables are described in Table III and the corresponding bin edges for the double differential unfolding and corresponding response matrices are shown in Table IV.

All distributions are compared to ν -truth, and the total uncertainty in each bin after unfolding is calculated with respect to the optimal performance achieved when using ν -truth. The correct jet and lepton association is used for both top quarks in order to remove the effects arising from matching inefficiencies. We perform the unfolding using the singular value decomposition (SVD) method [52] with a regularization factor of 7, using the implementation provided in `Roounfold` [53]. The regularization factor was optimized for the ν -truth distributions for a reduced χ^2 value closest in agreement to one for the four double differential distributions. In all cases we only consider the $t\bar{t}$ process and ignore the impact of background estimation and subtraction in the methods.

Although more modern machine learning approaches for unfolding are an active area of research [54–59], we leave their study and application to dilepton $t\bar{t}$ events to future studies, in particular the study of unbinned multidimensional unfolding with the reconstructed neutrino kinematics.

The response matrix using each of the neutrino reconstruction methods is shown in Fig. 6 for the two-dimensional binning in $m_{t\bar{t}}$ and $\Delta\phi(\ell^+\ell^-)$ and in Fig. 7 for $m_{t\bar{t}}$ and p_T^t . In the ideal case only the main diagonal would contain entries, however due to inefficiencies in the neutrino reconstruction methods as well as detector resolution effects, off diagonal elements are unavoidable. In both cases it is clear that using ν^2 -flows to reconstruct the neutrino pair results in a more diagonal response matrix than the other two approaches. This is quantified by the trace fraction of each matrix.

Although the trace fraction can give a good measure of which method is performing best, the off diagonal elements still contribute to the unfolded distributions. To quantitatively assess the true impact of using each method, the

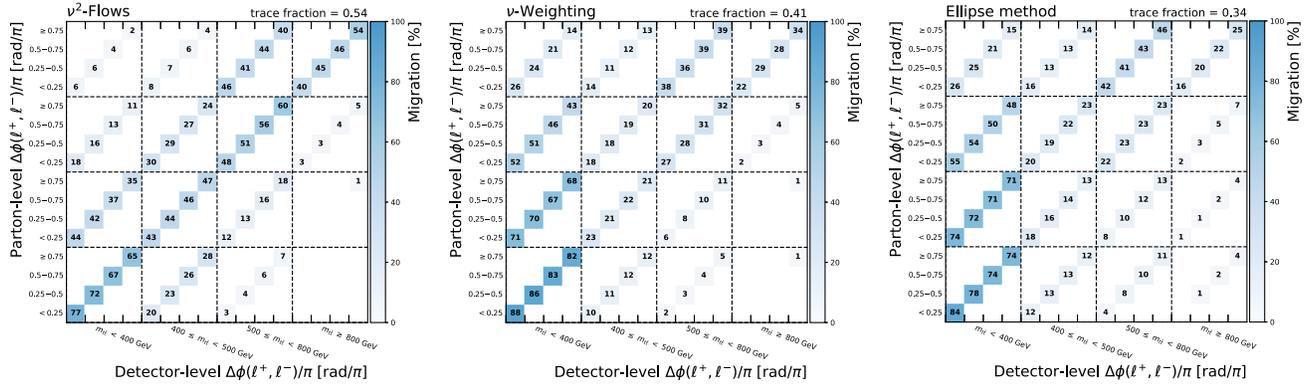


FIG. 6. Binned response matrices for the double-differential measurement of $m_{\bar{t}\bar{t}}$ and $\Delta\phi(\ell^+\ell^-)$ when using each of the three methods for neutrino reconstruction. The binning is symmetric for both the parton and detector level observables, however the $m_{\bar{t}\bar{t}}$ bins are labeled on the x -axis with the $\Delta\phi(\ell^+\ell^-)$ bins labeled on the y -axis. The trace fraction is calculated for each method for a simple quantitative comparison and is 0.73 when using ν -truth.

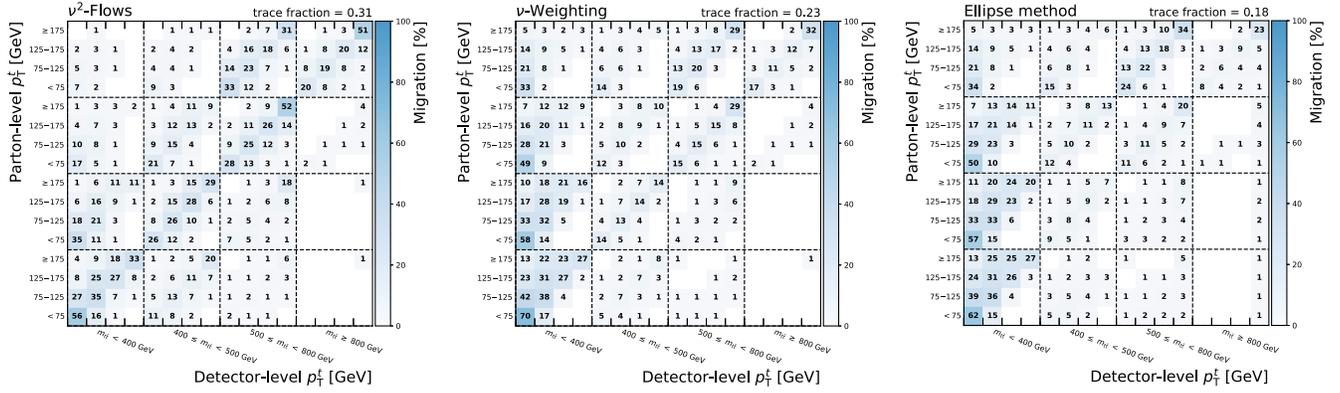


FIG. 7. Binned response matrices for the double differential measurement of $m_{\bar{t}\bar{t}}$ and p_T^{ℓ} when using each of the three methods for neutrino reconstruction. The binning is symmetric for both the parton and detector level observables, however the $m_{\bar{t}\bar{t}}$ bins are labeled on the x -axis with the p_T^{ℓ} bins labeled on the y -axis. The trace fraction is calculated for each method for a simple quantitative comparison and is 0.62 when using ν -truth.

response matrices are inverted using SVD and the overall uncertainties for each bin at parton level are calculated.

Figure 8 shows the relative statistical uncertainty for each method with respect to ν -truth. The values are obtained for each bin of the unfolded distributions using SVD with the chosen level of regularization. For the individual bins of the four double differential distributions the uncertainties are typically a factor of 1.5 to two times smaller when using ν^2 -flows compared to ν -weighting, and up to four times smaller in comparison with ellipse.

C. Robustness to training sample

In comparison to the standard analytical approaches, ν -flows is trained on a specific sample of Monte Carlo simulated events. This could introduce a performance dependence on the sample used for training, which may not be optimal for all generators. It should be noted that the

same model is used for all events and, just like the analytical approaches, is independent of which samples it is applied to. However, if ν -flows has learned sample specific effects this can lead to a suboptimal performance or even unusable levels of performance when applied to other samples.

To study the impact of this effect we train ν^2 -flows using the alternative $t\bar{t}$ dilepton sample [ν^2 -flows (PYTHIA8)] and use it to reconstruct the neutrinos for the nominal $t\bar{t}$ sample. We compare the reconstructed kinematic distributions as well as the uncertainties in each bin of the unfolded distributions.

Negligible differences are observed in the reconstructed neutrino kinematics, though the difference can clearly be seen for the reconstructed W boson mass and a slight difference is also seen for the reconstructed top quark invariant mass.

Some small differences are also observed in the tails of the reconstructed top quark and $t\bar{t}$ properties in Fig. 9, however the performance is still substantially improved in

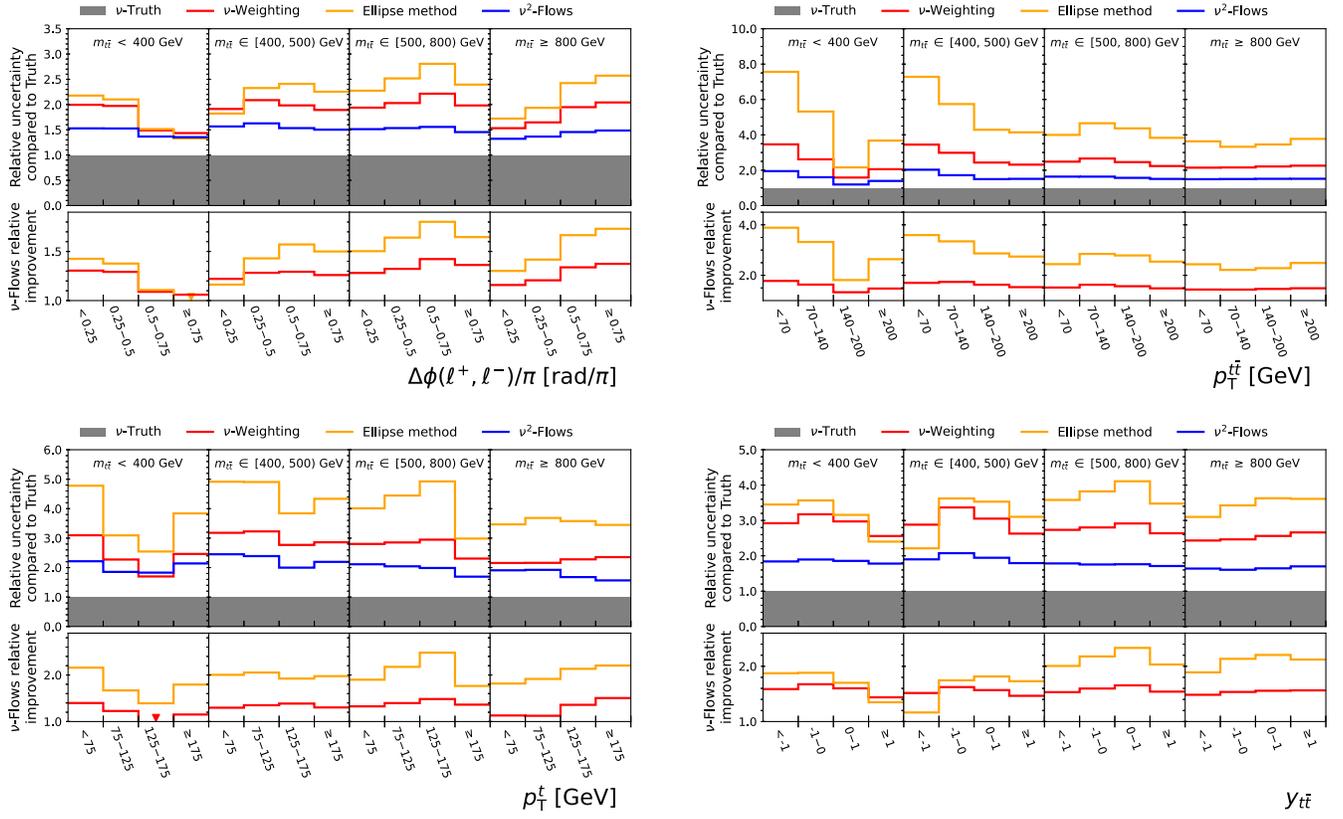


FIG. 8. The relative statistical uncertainty in each unfolded bin of the two-dimensional distributions for $m_{\ell\bar{\ell}}$ and $\Delta\phi(\ell^+\ell^-)$ (top left), $p_T^{\ell\bar{\ell}}$ (top right), $p_T^{\ell\bar{\ell}}$ (bottom left), and $y_{\ell\bar{\ell}}$ (bottom right). Comparing the three reconstruction methods with respect to ν -truth (upper pad), and the relative improvement of ν -flows with respect to ν -weighting and ellipse (lower pad).

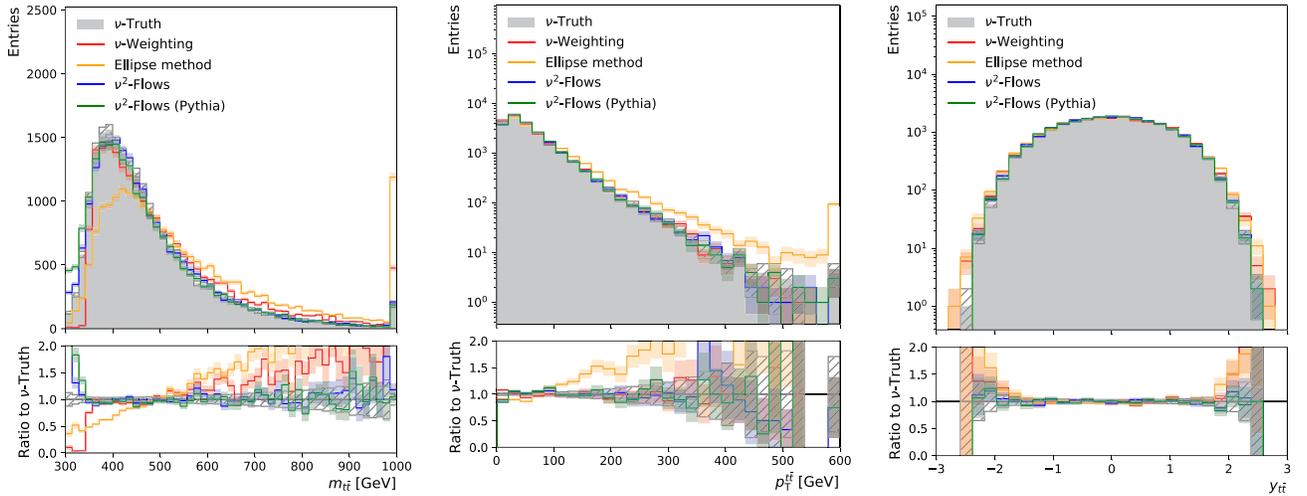


FIG. 9. The invariant mass, p_T , and rapidity of the reconstructed $t\bar{t}$ system when using ν^2 -flows trained on the nominal or alternative sample, in comparison to the two baseline approaches and ν -truth (shaded grey).

comparison to ν -weighting and ellipse. The response matrix for ν^2 -flows (PYTHIA8) for the double differential distribution in $m_{\ell\bar{\ell}}$ and $p_T^{\ell\bar{\ell}}$ is shown in Fig. 10. These differences translate to a very slight change in the statistical

precision in each bin after performing the unfolding (see the Appendix).

As another test of robustness we study how well ν^2 -flows performs in comparison the ν -weighting and ellipse at

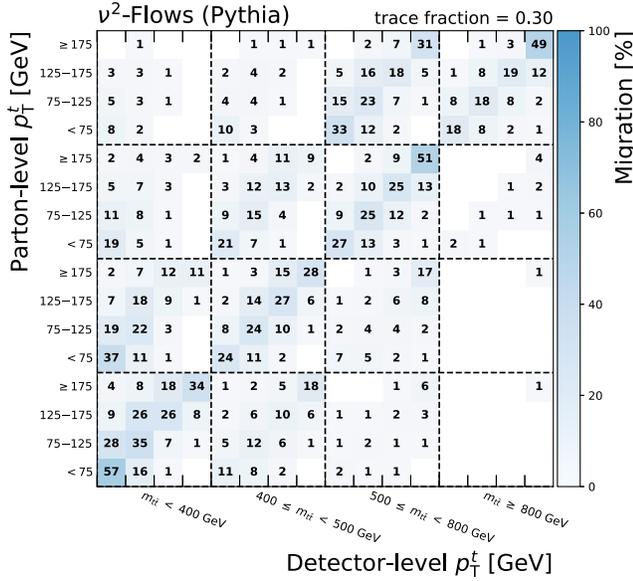


FIG. 10. Binned response matrices for the double differential measurement of $m_{t\bar{t}}$ and p_T^t when using ν^2 -flows (PYTHIA8) but evaluating on the nominal $t\bar{t}$ sample. The binning is symmetric for both the parton and detector level observables, however the $m_{t\bar{t}}$ bins are labeled on the x -axis with the p_T^t bins labeled on the y -axis. The trace fraction is calculated for each method for a simple quantitative comparison and is 0.62 when using ν -truth and 0.31 for the nominal ν^2 -flows.

reconstructing $t\bar{t}$ events which have been simulated with additional initial-state radiation (ISR). This sample has 280,000 events and is otherwise the same as the nominal sample, however substantial differences are observed in the jet multiplicity distributions and underlying kinematics of the top-quark and $t\bar{t}$ pair. Most notably the distribution of

p_T^t is much harder, due to the recoil against the additional ISR.

Figure 11 shows the reconstructed distributions for the three approaches. As before, ν^2 -flows exhibits very good agreement across the majority of the kinematic phase space, despite having been optimized for a different sample.

D. Measuring the top-quark mass

In contrast to ν -weighting and ellipse, which both assume values for the top-quark mass in the neutrino reconstruction, ν^2 -flows only implicitly learns this relation from the training data. To test whether ν^2 -flows has sensitivity to the underlying top-quark mass, we use the default ν^2 -flows model to evaluate events in which the truth m_t has been changed to either 171 GeV or 175 GeV. These samples each have 160,000 events and are otherwise the same as the nominal sample.

We compare the reconstructed top-quark mass using the four momenta of the lepton, b -quark, and neutrino (m_t) to the invariant mass using only of the lepton and b -quark from the same top decay ($m_{b\ell}$).

The distributions are shown in Fig. 12. Despite training purely on events with the nominal top quark mass, the reconstructed distribution using ν^2 -flows is sensitive to the difference in truth m_t . The separation between the three templates is similar to using $m_{b\ell}$, however for ν^2 -flows the difference is more prominent in the bulk of the distribution. This sensitivity could be improved by training ν^2 -flows on samples with a range of values for m_t and parametrizing the network, however this would introduce additional computational complexity to the method. Another benefit of ν^2 -flows is also demonstrated, with a smoother templates

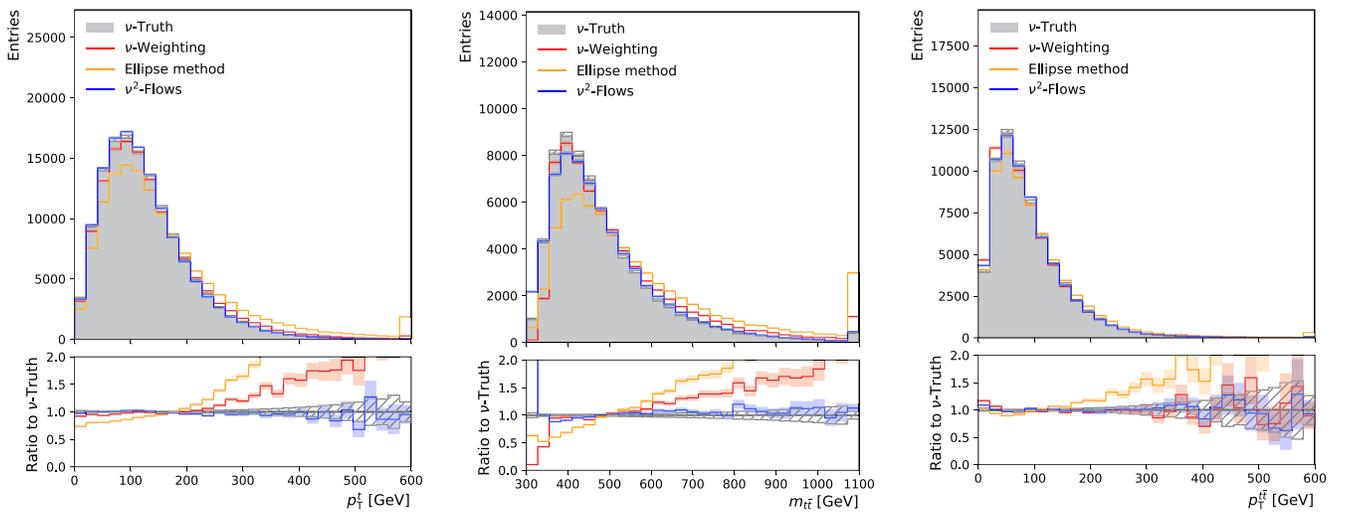


FIG. 11. The reconstructed top quark p_T , and the invariant mass and p_T of the reconstructed $t\bar{t}$ system for an independent sample with additional initial state radiation when using ν^2 -flows trained on the nominal sample in comparison to the two baseline approaches and ν -truth (shaded gray).

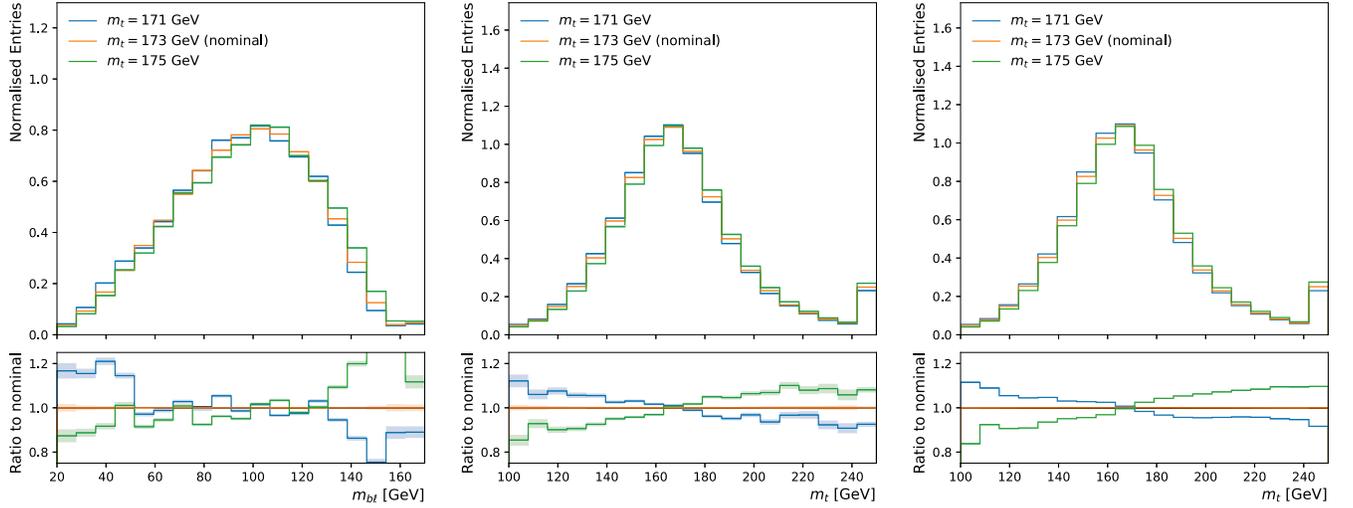


FIG. 12. The reconstructed top-quark mass using just from the lepton and b -quark from the top-quark decay ($m_{b\ell}$, left), the full invariant mass using ν^2 -flows to reconstruct the neutrinos from the top quarks (m_t , middle), and a smoothed template for m_t obtained by sampling 256 solutions for each event with ν^2 -flows (right). The statistical uncertainty on the distributions arising from the training dataset has not been propagated into the final statistical uncertainty.

constructed by sampling multiple neutrino solutions for each event.

V. CONCLUSIONS

With ν^2 -flows we have built upon the success of the ν -flows method to employ conditional normalizing flows to reconstruct the momentum vectors of multiple neutrinos in a single event. In comparison to other commonly used approaches, ν^2 -flows is able to reconstruct both neutrinos without enforcing strong constraints on reconstructed particle masses in order to find solutions in an underconstrained system. This translates to a reduced bias in the reconstruction of neutrinos, without a preference for back-to-back neutrinos, and with a more accurate reconstruction of the kinematics of individual top quarks and the full $t\bar{t}$ system. The reconstructed neutrinos can be used directly or potentially combined with other machine learning approaches which aim to reconstruct the underlying hard scatter event [60–64]. The generalized architecture introduced in ν^2 -flows has been designed to be easy to extend to any neutrino multiplicity, and does not place restrictions on the multiplicities of reconstructed objects, or how they are combined to extract information from the event.

In applying ν^2 -flows to dilepton $t\bar{t}$ events we achieve significant improvements in the statistical precision of unfolded distributions in comparison to the most commonly used analytical approaches for a wide range of distributions of interest. These improvements can lead to an increased sensitivity in the measurement of top-quark spin correlation and entanglement at the LHC. Furthermore, in contrast to other approaches, ν^2 -flows has solutions for all events and the fast single event inference below 75 ms on a single computing core. As each sample is associated with a

probability from the transformation under the normalizing flow, this could also provide opportunities for separating $t\bar{t}$ events from background processes, similar to the weight in ν -weighting. This could be done using the probability of the single solution per event, or, as multiple solutions can be sampled for each event, the highest possible probability value for an event.

ACKNOWLEDGMENTS

The authors would like to thank Chris Pollard for support producing the PYTHIA8 standalone samples. The authors acknowledge funding through the SNSF Sinergia grant “Robust Deep Density Models for High-Energy Particle Physics and Solar Flare Analysis (RODEM)” with Funding No. CRSII5_193716, and the SNSF project Grant No. 200020_212127 “At the two upgrade frontiers: machine learning and the ITk Pixel detector.” They would also like to acknowledge individual funding acquired through the Swiss Government Excellence Scholarships for Foreign Scholars and the Feodor Lynen Research Fellowship from the Alexander von Humboldt foundation.

J. A. R. and M. L. contributed equally to this work.

APPENDIX: ADDITIONAL MATERIAL

1. Unfolding efficiencies

The relative uncertainty for each bin in the unfolded distributions is compared in Table V for all methods with respect to ν -truth, including both ν^2 -flows and ν -flows (PYTHIA8). The efficiency of finding a solution with the ν -weighting and Ellipse methods for each bin of the double differential distributions are summarized in Table VI.

TABLE V. Relative uncertainty in each bin of the respective unfolded double differential distributions for each neutrino reconstruction method with respect to the uncertainty when using ν -truth. The bins are ordered first by increasing $m_{\bar{t}\bar{t}}$ followed by the second variable, with vertical dividers indicating the bin edges in $m_{\bar{t}\bar{t}}$. The method with the smallest relative increase in uncertainty in comparison to ν -truth is highlighted in bold.

Observables		Method	Relative uncertainty to ν -truth (per bin)															
$m_{\bar{t}\bar{t}}$	$y_{\bar{t}\bar{t}}$	ν -weighting	2.9	3.2	3.0	2.6	2.9	3.4	3.1	2.6	2.7	2.8	2.9	2.6	2.4	2.5	2.6	2.7
		Ellipse	3.5	3.6	3.2	2.4	2.2	3.6	3.5	3.1	3.6	3.8	4.1	3.5	3.1	3.4	3.6	3.6
		ν^2 -flows	1.8	1.9	1.9	1.8	1.9	2.1	1.9	1.8	1.8	1.8	1.8	1.7	1.6	1.6	1.6	1.7
		ν^2 -flows (PYTHIA8)	1.9	1.9	1.9	1.8	1.9	2.1	2.0	1.8	1.8	1.8	1.8	1.7	1.7	1.6	1.7	1.8
$p_{\bar{T}}^{\bar{t}\bar{t}}$	ν -weighting	ν -weighting	3.5	2.6	1.6	2.1	3.4	3.0	2.4	2.3	2.5	2.7	2.5	2.2	2.1	2.2	2.2	2.3
		Ellipse	7.6	5.3	2.2	3.7	7.3	5.7	4.3	4.1	4.0	4.7	4.4	3.8	3.6	3.3	3.5	3.8
		ν^2 -flows	1.9	1.6	1.2	1.4	2.0	1.7	1.5	1.5	1.6	1.6	1.6	1.5	1.5	1.5	1.5	1.5
		ν^2 -flows (PYTHIA8)	2.0	1.6	1.2	1.4	2.0	1.7	1.5	1.5	1.7	1.6	1.6	1.5	1.5	1.5	1.5	1.5
$p_{\bar{T}}^t$	ν -weighting	ν -weighting	3.1	2.3	1.7	2.5	3.2	3.2	2.8	2.9	2.8	2.9	2.9	2.3	2.2	2.2	2.3	2.4
		Ellipse	4.8	3.1	2.5	3.8	4.9	4.9	3.8	4.3	4.0	4.4	4.9	3.0	3.5	3.7	3.6	3.4
		ν^2 -flows	2.2	1.9	1.8	2.1	2.5	2.4	2.0	2.2	2.1	2.0	2.0	1.7	1.9	1.9	1.7	1.6
		ν^2 -flows (PYTHIA8)	2.3	1.9	1.9	2.2	2.5	2.4	2.0	2.3	2.2	2.1	2.0	1.7	2.0	2.0	1.7	1.6
$\Delta\phi(\ell^+\ell^-)$	ν -weighting	ν -weighting	2.0	2.0	1.5	1.4	1.9	2.1	2.0	1.9	1.9	2.0	2.2	2.0	1.5	1.6	1.9	2.0
		Ellipse	2.2	2.1	1.5	1.3	1.8	2.3	2.4	2.3	2.3	2.5	2.8	2.4	1.7	1.9	2.4	2.6
		ν^2 -flows	1.5	1.5	1.4	1.4	1.6	1.6	1.5	1.5	1.5	1.5	1.6	1.5	1.3	1.4	1.5	1.5
		ν^2 -flows (PYTHIA8)	1.5	1.5	1.4	1.4	1.6	1.6	1.5	1.5	1.5	1.5	1.6	1.5	1.3	1.4	1.5	1.5

TABLE VI. Efficiency for finding a solution in each bin of the respective unfolded double differential distributions with ν -weighting and Ellipse. The bins are ordered first by increasing $m_{\bar{t}\bar{t}}$ followed by the second variable, with vertical dividers indicating the bin edges in $m_{\bar{t}\bar{t}}$. The efficiency of ν^2 -flows in all bins is 100%.

Observables		Method	Efficiency at finding a solution (per bin) [%]															
$m_{\bar{t}\bar{t}}$	$y_{\bar{t}\bar{t}}$	ν -weighting	97	98	97	97	96	96	97	96	93	94	94	94	89	90	89	89
		Ellipse	91	92	92	91	84	82	84	82	67	68	68	70	58	54	51	50
$p_{\bar{T}}^{\bar{t}\bar{t}}$	ν -weighting	ν -weighting	99	95	91	89	98	95	90	84	97	92	89	82	94	88	81	81
		Ellipse	96	86	74	59	88	79	66	50	75	66	52	38	60	52	34	37
$p_{\bar{T}}^t$	ν -weighting	ν -weighting	98	96	92	86	97	97	96	91	95	95	95	93	96	89	90	89
		Ellipse	95	88	72	51	89	86	80	63	77	74	72	63	61	56	55	51
$\Delta\phi(\ell^+\ell^-)$	ν -weighting	ν -weighting	96	97	98	99	96	96	97	97	93	94	94	94	90	90	91	88
		Ellipse	90	92	92	94	84	84	83	82	76	74	68	64	57	57	57	48

As ν^2 -flows is able to find a solution for all events, a value lower than 100% indicates where additional statistical sensitivity can be achieved with ν^2 -flows. This is particularly notable with increasing values of $m_{\bar{t}\bar{t}}$, but also for higher values of $p_{\bar{T}}^{\bar{t}\bar{t}}$. This is especially the case in comparison to Ellipse, with some bins having more than 50% of events without solutions, however, in certain unfolding bins 10–20% of events also do not have solutions when using ν -weighting.

2. Direct regression

Strong biases are introduced when replacing the normalizing flow with a naïve feed-forward network to predict the neutrino momenta for each event directly. This network, ν^2 -FF, has the same architecture as the feature extraction network in ν^2 -flows, but with a dense output layer for the

neutrino momenta, and is trained with an L2 loss. Figure 13 shows the reconstructed invariant mass of the top quark and $\bar{t}\bar{t}$ pair when using ν^2 -FF to reconstruct the neutrinos in comparison to the other reconstruction approaches. A strong negative bias can be seen in both distributions.

3. Origin of improvement over ν -weighting

In the ν -weighting a weight is used to find the best solution, as defined in Eq. (3), which can also be interpreted as how good a solution is. The improvement from ν^2 -flows could arise from the events with low values of w or no solutions, or come from all events regardless of how well ν -weighting performs.

To investigate this we compare the reconstruction performance for events where ν -weighting has either a good solution or a poor solution. We define good

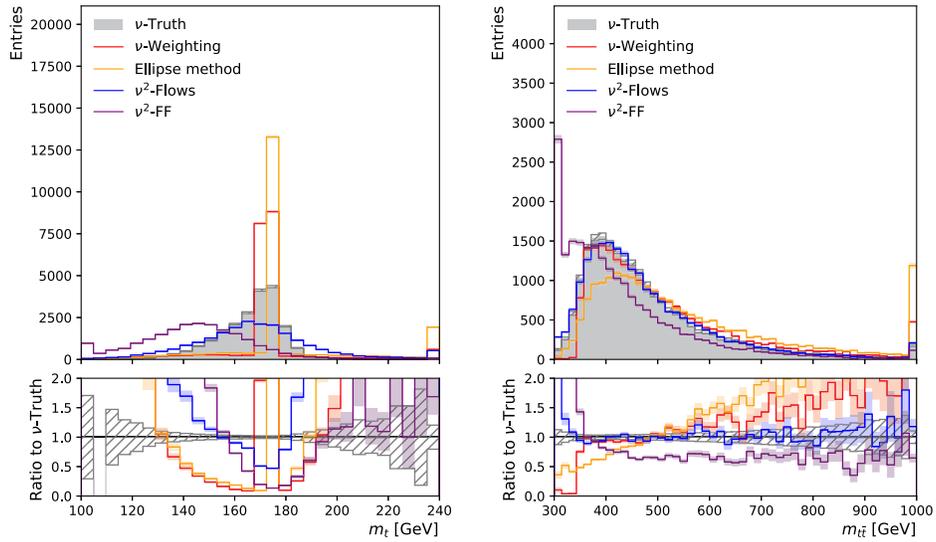


FIG. 13. The invariant masses of the reconstructed top quark and $t\bar{t}$ system when using the three neutrino reconstruction methods discussed in the paper, as well as a feed-forward regression model ν^2 -FF, in comparison to ν -truth (shaded gray). This plot highlights large negative bias induced by the feed-forward model.

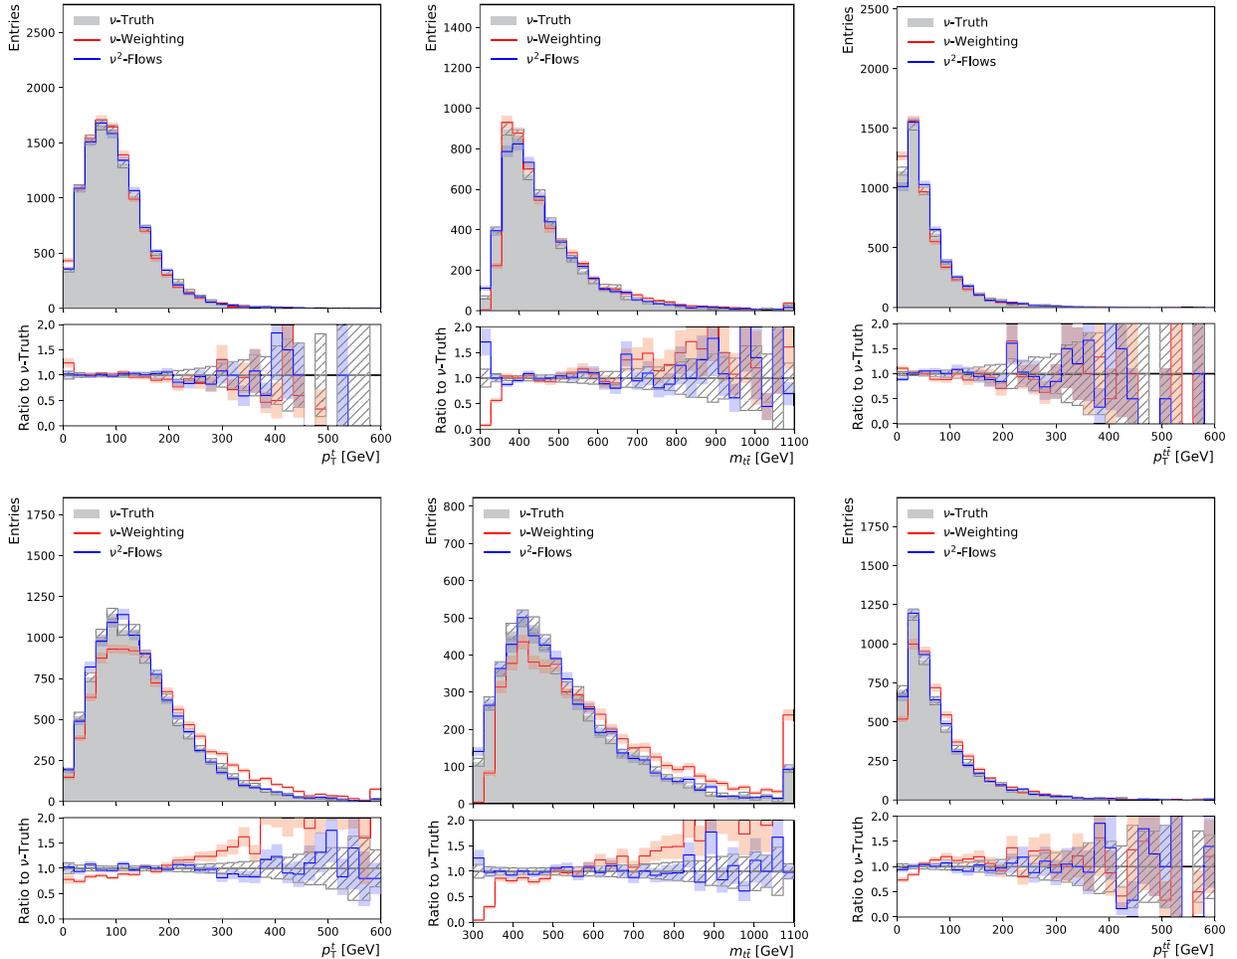


FIG. 14. The reconstructed top quark p_T , and the invariant mass and p_T of the reconstructed $t\bar{t}$ system when using ν^2 -flows and ν -weighting in comparison to the two baseline approaches and ν -truth (shaded gray) for events where ν -weighting returns $w > 0.95$ (top) and $w < 0.8$ (bottom).

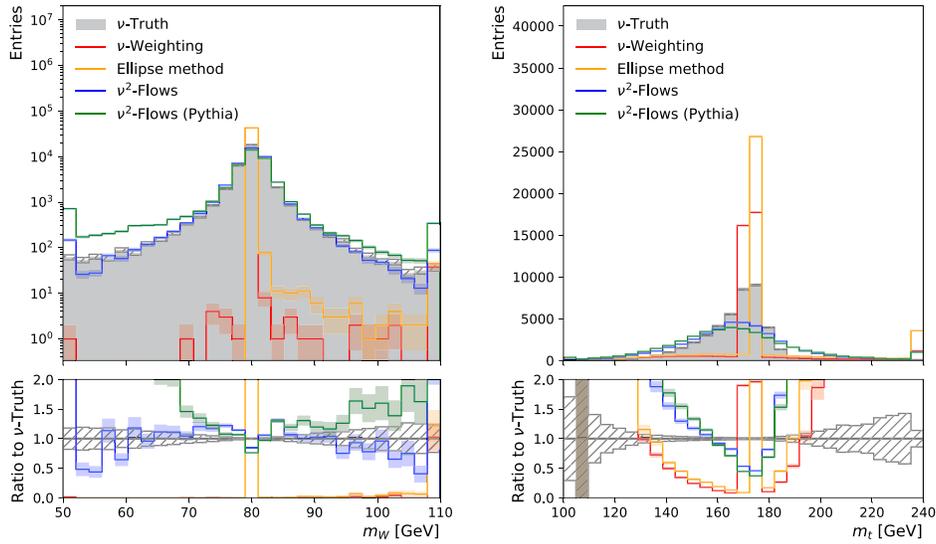


FIG. 15. The reconstructed invariant mass of the W boson and top quarks when using the three neutrino reconstruction methods in comparison to ν -truth (shaded grey) as well as the alternative ν^2 -flows (PYTHIA8) model.

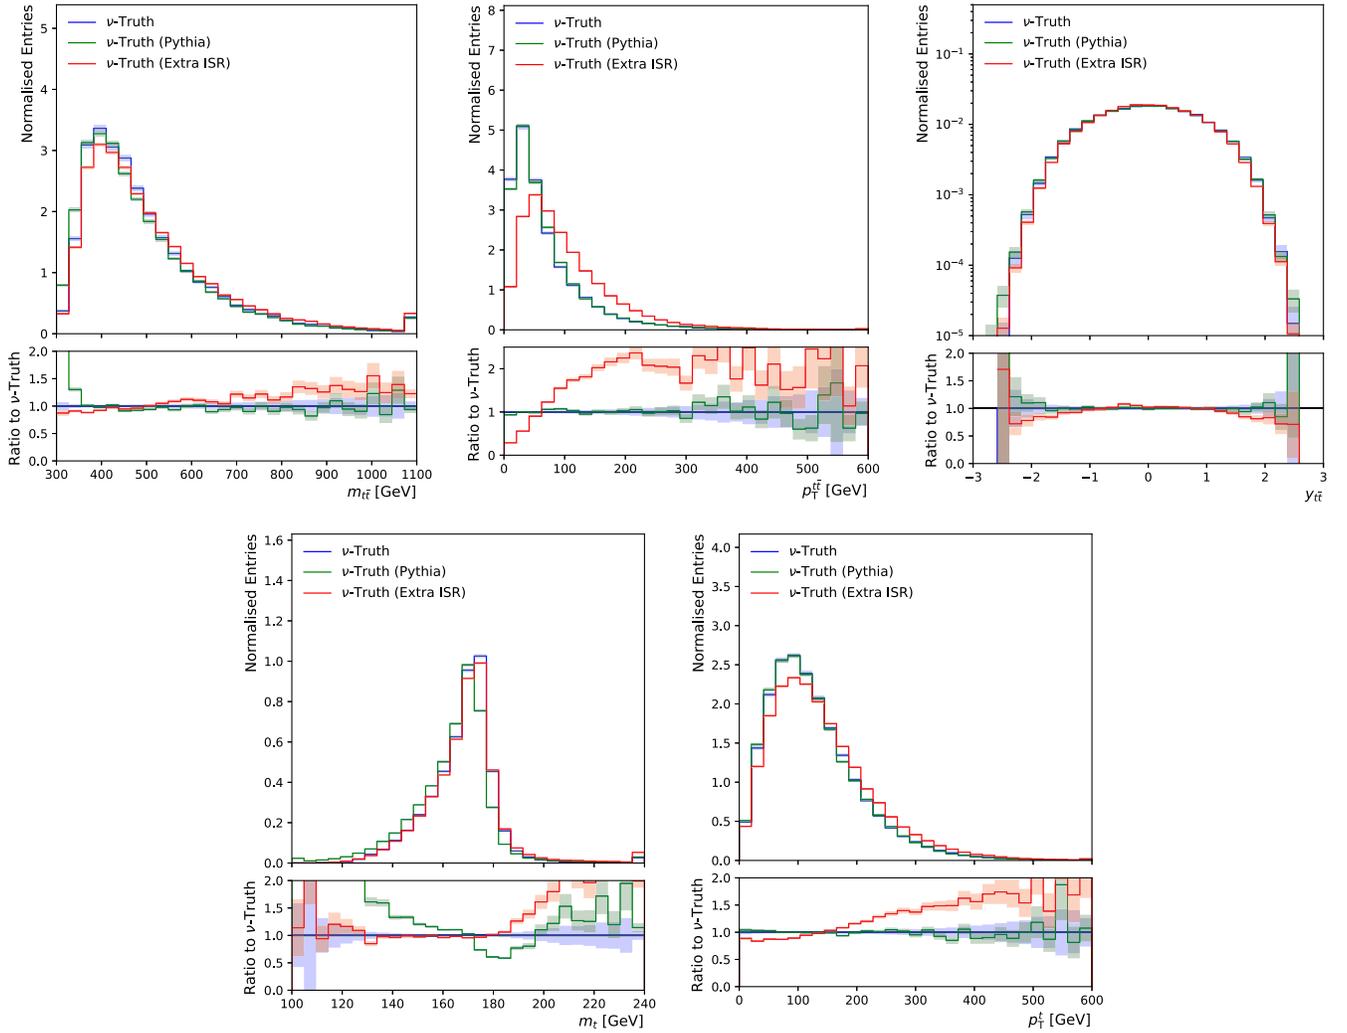
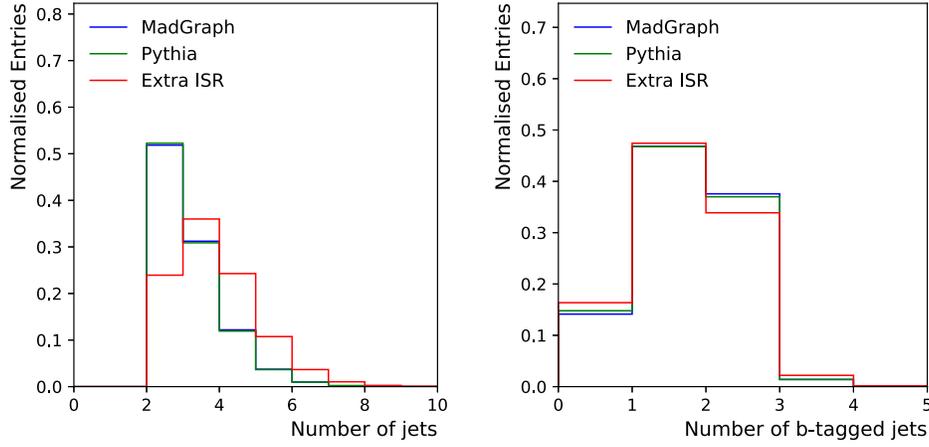


FIG. 16. The invariant mass, p_T , and rapidity of the reconstructed $t\bar{t}$ system (top row) and the invariant mass, and p_T of the reconstructed top quarks (bottom row) for ν -truth with the three independent simulated samples.

FIG. 17. The jet and b -jet multiplicities of the three independent simulated samples.

performance as values of $w > 0.9$ and poor performance as $w < 0.3$. The p_T^i , $p_T^{\bar{i}}$, and $m_{i\bar{i}}$ distributions with these selections are shown in Fig. 14. Although substantial improvement can be seen over ν -weighting in regions where ν -weighting has a low weight, ν^2 -flows still exhibits

improved agreement with ν -truth. It is equally encouraging that there is no noticeable difference in the quality of reconstruction with ν^2 -flows for events that can be considered well reconstructed or poorly reconstructed by ν -weighting.

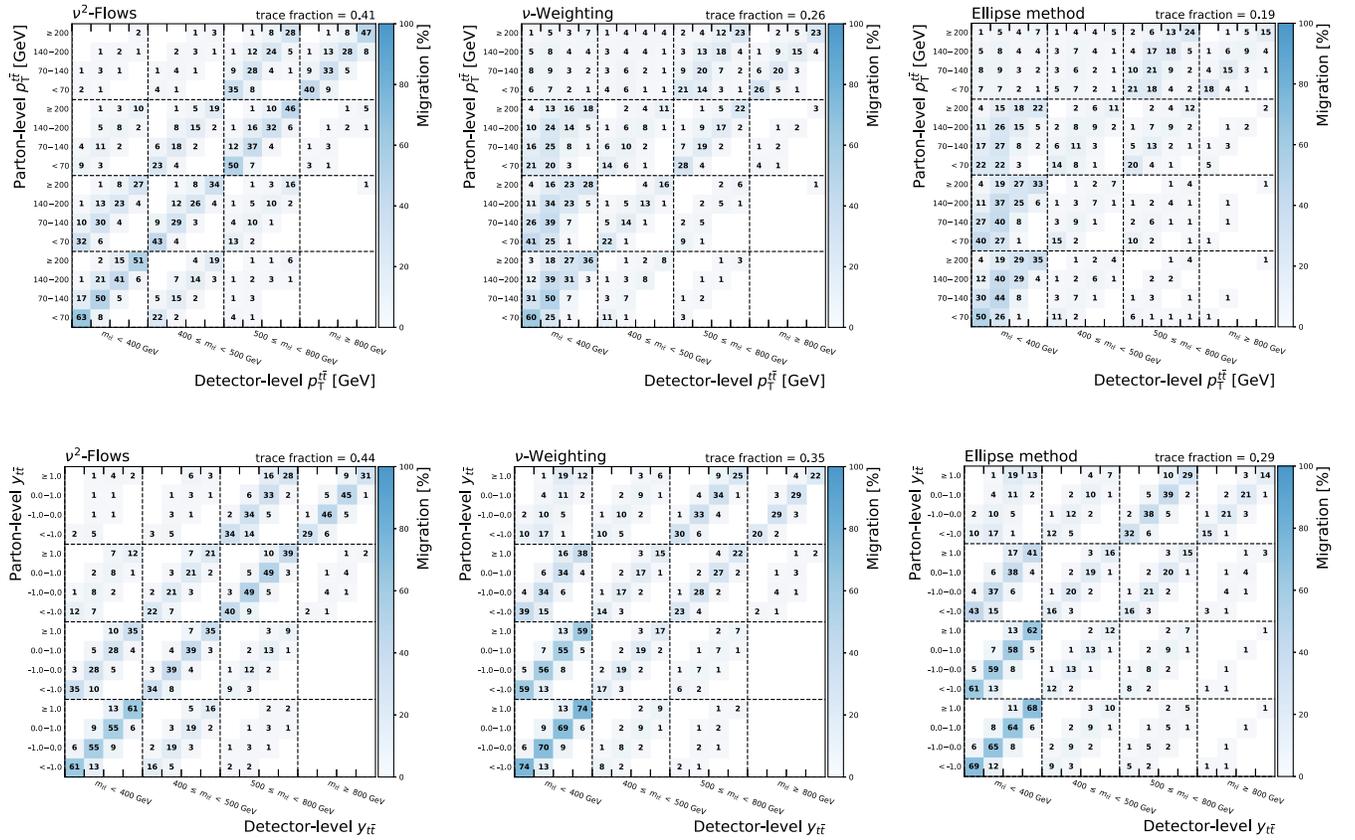


FIG. 18. Binned response matrices for the double differential measurement of $m_{i\bar{i}}$ and p_T^i (top) and $m_{i\bar{i}}$ and $y_{i\bar{i}}$ (bottom) when using each of the three methods for neutrino reconstruction. The binning is symmetric for both the parton and detector level observables, however the $m_{i\bar{i}}$ bins are labelled on the x -axis with the $y_{i\bar{i}}$ bins labelled on the y -axis. The trace fraction is calculated for each method for a simple quantitative comparison and is 0.70 when using ν -truth.

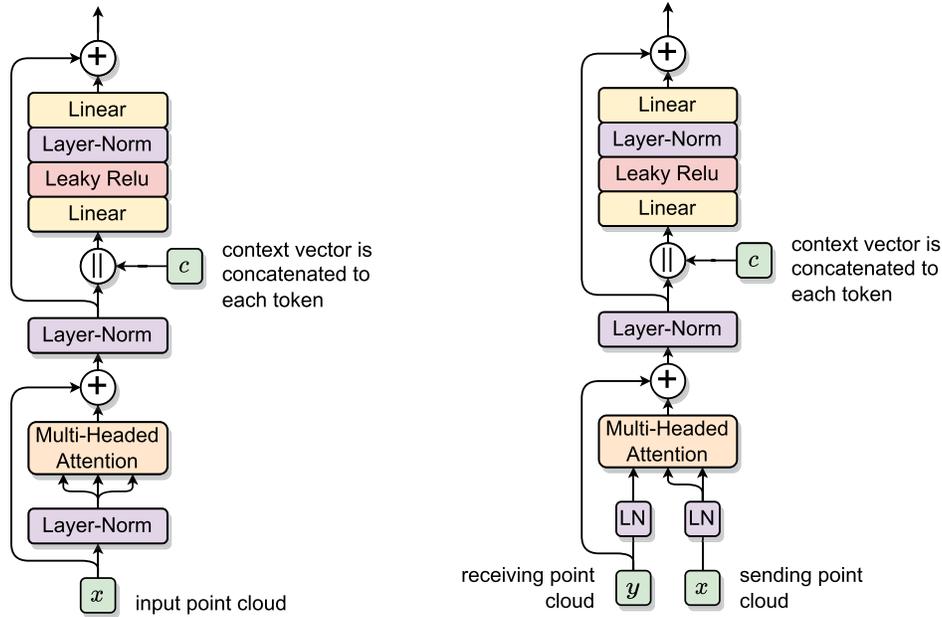


FIG. 19. Transformer encoder block (left) and cross-attention block (right) comprising multiheaded attention, layer normalization and simple linear layers. Residual connections are used after the multiheaded attention and linear operation. Conditional information is provided as context by concatenating it to each token before the linear layer.

4. Additional figures

The reconstructed top quark and W boson invariant masses are compared in Fig. 15. The ν -truth distributions for the nominal, alternative and extra ISR sample are compared in Fig. 16, and the jet and b -jet multiplicities for the three samples are compared in Fig. 17.

The response matrices for the double differential distributions in $m_{\bar{t}t}$ and p_T^t , and $m_{\bar{t}t}$ and $y_{\bar{t}t}$, are shown in Fig. 18.

Detailed schema of the transformer encoder block and cross-attention block used in ν^2 -flows are detailed in Fig. 19.

- [1] L. R. Evans and P. Bryant, LHC machine, *J. Instrum.* **3**, S08001 (2008).
- [2] ATLAS Collaboration, The ATLAS experiment at the CERN Large Hadron Collider, *J. Instrum.* **3**, S08003 (2008).
- [3] CMS Collaboration, The CMS experiment at the CERN LHC, *J. Instrum.* **3**, S08004 (2008).
- [4] M. Grossi, J. Novak, B. Kerševan, and D. Rebuffi, Comparing traditional and deep-learning techniques of kinematic reconstruction for polarization discrimination in vector boson scattering, *Eur. Phys. J. C* **80**, 1144 (2020).
- [5] J. Kvita, Study of methods of resolved top quark reconstruction in semileptonic $t\bar{t}$ decay, *Nucl. Instrum. Methods Phys. Res., Sect. A* **900**, 84 (2018).
- [6] ATLAS Collaboration, Measurements of top-quark pair differential and double-differential cross-sections in the $\ell + \text{jets}$ channel with pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector, *Eur. Phys. J. C* **79**, 1028 (2019); **80**, 1092(E) (2020).
- [7] CMS Collaboration, Measurement of differential $t\bar{t}$ production cross sections in the full kinematic range using lepton + jets events from proton-proton collisions at $\sqrt{s} = 13$ TeV, *Phys. Rev. D* **104**, 092013 (2021).
- [8] ATLAS Collaboration, Determination of the top-quark pole mass using $t\bar{t} + 1$ -jet events collected with the ATLAS experiment in 7 TeV pp collisions, *J. High Energy Phys.* **10** (2015) 121.
- [9] CMS Collaboration, Measurement of the top quark mass with lepton + jets final states using pp collisions at $\sqrt{s} = 13$ TeV, *Eur. Phys. J. C* **78**, 891 (2018); **82**, 323(E) (2022).
- [10] ATLAS Collaboration, Measurement of the top quark mass in the $t\bar{t} \rightarrow \text{lepton} + \text{jets}$ channel from $\sqrt{s} = 8$ TeV ATLAS data and combination with previous results, *Eur. Phys. J. C* **79**, 290 (2019).
- [11] ATLAS Collaboration, Measurement of the top-quark mass in $t\bar{t} + 1$ -jet events collected with the ATLAS detector in pp collisions at $\sqrt{s} = 8$ TeV, *J. High Energy Phys.* **11** (2019) 150.
- [12] ATLAS Collaboration, Evidence for the charge asymmetry in $pp \rightarrow t\bar{t}$ production at $\sqrt{s} = 13$ TeV with the ATLAS detector, *J. High Energy Phys.* **08** (2023) 077.

- [13] B. Abbott *et al.* (D0 Collaboration), Measurement of the top quark mass using dilepton events, *Phys. Rev. Lett.* **80**, 2063 (1998).
- [14] L. Sonnenschein, Algebraic approach to solve $t\bar{t}$ dilepton equations, *Phys. Rev. D* **72**, 095020 (2005).
- [15] B. A. Betchart, R. Demina, and A. Harel, Analytic solutions for neutrino momenta in decay of top quarks, *Nucl. Instrum. Methods Phys. Res., Sect. A* **736**, 169 (2014).
- [16] CMS Collaboration, Measurement of the $t\bar{t}$ production cross section and the top quark mass in the dilepton channel in pp collisions at $\sqrt{s} = 7$ TeV, *J. High Energy Phys.* **07** (2011) 049.
- [17] CMS Collaboration, Measurement of the top-quark mass in $t\bar{t}$ events with dilepton final states in pp collisions at $\sqrt{s} = 7$ TeV, *Eur. Phys. J. C* **72**, 2202 (2012).
- [18] ATLAS Collaboration, Measurement of top quark polarization in top-antitop events from proton-proton collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector, *Phys. Rev. Lett.* **111**, 232002 (2013).
- [19] CMS Collaboration, Measurements of the $t\bar{t}$ charge asymmetry using the dilepton decay channel in pp collisions at $\sqrt{s} = 7$ TeV, *J. High Energy Phys.* **04** (2014) 191.
- [20] ATLAS Collaboration, Measurements of spin correlation in top-antitop quark events from proton-proton collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector, *Phys. Rev. D* **90**, 112016 (2014).
- [21] ATLAS Collaboration, Measurement of the charge asymmetry in dileptonic decays of top quark pairs in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector, *J. High Energy Phys.* **05** (2015) 061.
- [22] D0 Collaboration, Precise measurement of the top quark mass in dilepton decays using optimized neutrino weighting, *Phys. Lett. B* **752**, 18 (2016).
- [23] CMS Collaboration, Measurements of $t\bar{t}$ -bar spin correlations and top quark polarization using dilepton final states in pp collisions at $\sqrt{s} = 8$ TeV, *Phys. Rev. D* **93**, 052007 (2016).
- [24] CMS Collaboration, Measurements of $t\bar{t}$ charge asymmetry using dilepton final states in pp collisions at $\sqrt{s} = 8$ TeV, *Phys. Lett. B* **760**, 365 (2016).
- [25] ATLAS Collaboration, Measurements of top quark spin observables in $t\bar{t}$ events using dilepton final states in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector, *J. High Energy Phys.* **03** (2016) 113.
- [26] ATLAS Collaboration, Measurements of top-quark pair differential cross-sections in the $e\mu$ channel in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector, *Eur. Phys. J. C* **77**, 292 (2017).
- [27] ATLAS Collaboration, Measurements of top-quark pair spin correlations in the $e\mu$ channel at $\sqrt{s} = 13$ TeV using pp collisions in the ATLAS detector, *Eur. Phys. J. C* **80**, 754 (2020).
- [28] CMS Collaboration, Measurement of the top quark polarization and $t\bar{t}$ spin correlations using dilepton final states in proton-proton collisions at $\sqrt{s} = 13$ TeV, *Phys. Rev. D* **100**, 072002 (2019).
- [29] M. Leigh, J. A. Raine, K. Zoch, and T. Golling, ν -flows: Conditional neutrino regression, *SciPost Phys.* **14**, 159 (2023).
- [30] E. G. Tabak and E. Vanden-Eijnden, Density estimation by dual ascent of the log-likelihood, *Commun. Math. Sci.* **8**, 217 (2010).
- [31] D. Rezende and S. Mohamed, Variational inference with normalizing flows, in *Proceedings of the 32nd International Conference on Machine Learning*, PMLR Vol. 37 (PMLR, Lille, France, 2015), pp. 1530–1538.
- [32] L. Ardizzone *et al.*, Guided image generation with conditional invertible neural networks, [arXiv:1907.02392](https://arxiv.org/abs/1907.02392).
- [33] J. A. Raine, M. Leigh, K. Zoch, L. Ehrke, D. Sengupta, and T. Golling, Dileptonic $t\bar{t}$ bar neutrino regression dataset, [10.5281/zenodo.8113516](https://zenodo.org/record/8113516) (2023).
- [34] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, Attention is all you need, in *Proceedings of Advances in Neural Information Processing Systems* (Curran Associates, Inc., Long Beach, USA, 2017), Vol. 30.
- [35] M. Zaheer *et al.*, Deep sets, in *Proceedings of Advances in Neural Information Processing Systems* (Curran Associates, Inc., Long Beach, USA, 2017), Vol. 30, pp. 3391–3401.
- [36] H. Touvron, M. Cord, A. Sablayrolles, G. Synnaeve, and H. Jégou, Going deeper with image transformers, in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (IEEE, 2021), pp. 32–42.
- [37] A. L. Maas, A. Y. Hannun, and A. Y. Ng, Rectifier nonlinearities improve neural network acoustic models, in *Proceedings of The International Conference on Machine Learning* (PMLR, Atlanta, USA, 2013), Vol. 28, p. 3.
- [38] J. L. Ba, J. R. Kiros, and G. E. Hinton, Layer normalization, [arXiv:1607.06450](https://arxiv.org/abs/1607.06450).
- [39] C. Durkan *et al.*, Neural spline flows, in *Proceedings of Advances in Neural Information Processing Systems* (Curran Associates, Inc., Vancouver, Canada, 2019), Vol. 32, pp. 7511–7522.
- [40] C. Durkan *et al.*, nflows: Normalizing flows in pyTorch (Zenodo, 2020), [10.5281/zenodo.4296287](https://zenodo.org/record/4296287).
- [41] A. Paszke *et al.*, PyTorch: An imperative style, high-performance deep learning library, in *Proceedings of Advances in Neural Information Processing Systems* (Curran Associates, Inc., Vancouver, Canada, 2017), Vol. 32, pp. 8024–8035.
- [42] I. Loshchilov and F. Hutter, Decoupled weight decay regularization, [arXiv:1711.05101](https://arxiv.org/abs/1711.05101).
- [43] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.-S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations, *J. High Energy Phys.* **07** (2014) 079.
- [44] T. Sjöstrand, S. Mrenna, and P. Skands, A brief introduction to PYTHIA 8.1, *Comput. Phys. Commun.* **178**, 852 (2008).
- [45] R. D. Ball *et al.*, Parton distributions with LHC data, *Nucl. Phys.* **B867**, 244 (2013).
- [46] A. Buckley, J. Ferrando, S. Lloyd, K. Nordström, B. Page, M. Rüfenacht, M. Schönherr, and G. Watt, LHAPDF6: Parton density access in the LHC precision era, *Eur. Phys. J. C* **75**, 132 (2015).
- [47] P. Skands, S. Carrazza, and J. Rojo, Tuning PYTHIA 8.1: The Monash 2013 Tune, *Eur. Phys. J. C* **74**, 3024 (2014).

- [48] 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, *J. High Energy Phys.* **02** (2014) 057.
- [49] M. Cacciari, G. P. Salam, and G. Soyez, The anti-kt jet clustering algorithm, *J. High Energy Phys.* **04** (2008) 063.
- [50] M. Cacciari, G. P. Salam, and G. Soyez, FastJet user manual, *Eur. Phys. J. C* **72**, 1 (2012).
- [51] CMS Collaboration, Measurement of double-differential cross sections for top quark pair production in pp collisions at $\sqrt{s} = 8$ TeV and impact on parton distribution functions, *Eur. Phys. J. C* **77**, 459 (2017).
- [52] A. Hocker and V. Kartvelishvili, SVD approach to data unfolding, *Nucl. Instrum. Methods Phys. Res., Sect. A* **372**, 469 (1996).
- [53] L. Brenner, R. Balasubramanian, C. Burgard, W. Verkerke, G. Cowan, P. Verschuur, and V. Croft, Comparison of unfolding methods using RooFitUnfold, *Int. J. Mod. Phys. A* **35**, 2050145 (2020).
- [54] N. D. Gagunashvili, Machine learning approach to inverse problem and unfolding procedure, [arXiv:1004.2006](https://arxiv.org/abs/1004.2006).
- [55] K. Datta, D. Kar, and D. Roy, Unfolding with generative adversarial networks, [arXiv:1806.00433](https://arxiv.org/abs/1806.00433).
- [56] 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).
- [57] M. Bellagente, A. Butter, G. Kasieczka, T. Plehn, A. Rousselet, R. Winterhalder, L. Ardizzone, and U. Köthe, Invertible networks or partons to detector and back again, *SciPost Phys.* **9**, 074 (2020).
- [58] M. Backes, A. Butter, M. Dunford, and B. Malaescu, An unfolding method based on conditional Invertible Neural Networks (cINN) using iterative training, [arXiv:2212.08674](https://arxiv.org/abs/2212.08674).
- [59] J. Chan and B. Nachman, Unbinned profiled unfolding, *Phys. Rev. D* **108**, 016002 (2023).
- [60] M. J. Fenton, A. Shmakov, T.-W. Ho, S.-C. Hsu, D. Whiteson, and P. Baldi, Permutationless many-jet event reconstruction with symmetry preserving attention networks, *Phys. Rev. D* **105**, 112008 (2022).
- [61] A. Shmakov, M. J. Fenton, T.-W. Ho, S.-C. Hsu, D. Whiteson, and P. Baldi, SPANet: Generalized permutationless set assignment for particle physics using symmetry preserving attention, *SciPost Phys.* **12**, 178 (2022).
- [62] A. Badea, W. J. Fawcett, J. Huth, T. J. Khoo, R. Poggi, and L. Lee, Solving combinatorial problems at particle colliders using machine learning, *Phys. Rev. D* **106**, 016001 (2022).
- [63] S. Qiu, S. Han, X. Ju, B. Nachman, and H. Wang, Parton labeling without matching: Unveiling emergent labelling capabilities in regression models, *Eur. Phys. J. C* **83**, 622 (2023).
- [64] L. Ehrke, J. A. Raine, K. Zoch, M. Guth, and T. Golling, Topological reconstruction of particle physics processes using graph neural networks, *Phys. Rev. D* **107**, 116019 (2023).