OmniFold: A Method to Simultaneously Unfold All Observables

Anders Andreassen,^{1, 2, 3, *} Patrick T. Komiske,^{4, †} Eric M. Metodiev,^{4, ‡} Benjamin Nachman,^{2, §} and Jesse Thaler^{4, ¶}

¹Department of Physics, University of California, Berkeley, CA 94720, USA

²Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

³Google, Mountain View, CA 94043, USA

⁴Center for Theoretical Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A.

Collider data must be corrected for detector effects ("unfolded") to be compared with theoretical calculations and measurements from other experiments. Unfolding is traditionally done for individual, binned observables without including all information relevant for characterizing the detector response. We introduce OMNIFOLD, an unfolding method that iteratively reweights a simulated dataset, using machine learning to capitalize on all available information. Our approach is unbinned, works for arbitrarily high-dimensional data, and naturally incorporates information from the full phase space. We illustrate this technique on a realistic jet substructure example from the Large Hadron Collider and compare it to standard binned unfolding methods. This new paradigm enables the simultaneous measurement of all observables, including those not yet invented at the time of the analysis.

Measuring properties of particle collisions is a central goal of particle physics experiments, such as those at the Large Hadron Collider (LHC). Distributions of collider observables at truth-level can be compared with theoretical predictions as well as with measurements from other experiments. These comparisons are widely used to enhance our understanding of the Standard Model, tune parameters of Monte Carlo event generators, and enable precision searches for new physics. "Unfolding" is the process of obtaining these truth distributions (particlelevel) from measured information recorded by a detector (detector-level). The unfolding process ensures that measurements are independent of the specific experimental context, allowing for comparisons across different experiments and usage with the latest theoretical tools, even long after the original analysis is completed. Many unfolding techniques have been proposed and are currently in widespread use by experiments. See Refs. [1–4] for reviews and Refs. [5-7] for the most widely-used unfolding algorithms.

Current unfolding methods face three key challenges. First, all of the widely-used methods require the measured observables to be binned into histograms. This binning must be determined ahead of time and is often chosen manually. Second, because the measurements are binned, one can only unfold a small number of observables simultaneously. Multi-differential cross section measurements beyond two or three dimensions are simply not feasible. Finally, unfolding corrections for detector effects often do not take into account all possible auxiliary features that control the detector response. Even though the inputs to the unfolding can be calibrated, if the detector response depends on features that are not used directly in the unfolding, then the results will be suboptimal and potentially biased.

This letter introduces OMNIFOLD, a new approach that solves all three of these unfolding challenges. Detector-level quantities are iteratively unfolded, using machine learning to handle phase space of any dimensionality without requiring binning. Utilizing the full phase space information mitigates the problem of auxiliary features controlling the detector response. There have been previous proposals to use machine learning methods for unfolding [8-10] as well as proposals to perform unfolding without binning [9-12]. These proposals, however, are untenable in high dimensions and do not reduce to standard methods in the binned case. OMNI-FOLD naturally processes high-dimensional features, in the spirit of previous machine-learning-based reweighting strategies [13] (see also [14, 15]), and it reduces to well-established methods [5] in the binned case. We also introduce simpler versions of the procedure, using single or multiple observables, respectively named UNIFOLD and MULTIFOLD.¹

All unfolding methods require a trustable detector simulation to estimate the detector response. In the binned formulation, the folding equation can be written as $\mathbf{m} = \mathbf{R} \mathbf{t}$, where \mathbf{m} and \mathbf{t} are vectors of the measured detector-level and true particle-level histograms, respectively. \mathbf{R} is the "response matrix":

$$R_{ij} = \Pr(\text{measure } i \,|\, \text{truth is } j). \tag{1}$$

In general, **R** is not invertible, so that the unfolding problem has no unique solution, and methods attempt to achieve a useful solution in various ways. One of the most widely-used methods is Iterative Bayesian Unfolding (IBU) [5], also known as Richardson-Lucy deconvolution [17, 18]. Given a measured spectrum $m_i =$ $\Pr(\text{measure } i)$ and a prior spectrum $t_i^{(0)} = \Pr(\text{truth is } j)$,

¹ The name OMNIFOLD is taken from Emily Dickinson's poem *The Mountain Sat Upon the Plain* [16].

IBU proceeds iteratively according to the equation:

$$t_j^{(n)} = \sum_i \Pr(\text{truth is } j | \text{measure } i) \times \Pr(\text{measure } i)$$
$$= \sum_i \frac{R_{ij} t_j^{(n-1)}}{\sum_k R_{ik} t_k^{(n-1)}} \times m_i, \qquad (2)$$

where n is the iteration number.

OMNIFOLD uses machine learning to generalize Eq. (2) to the unbinned, full phase space. A key concept for this approach is the likelihood ratio:

$$L[(w,X),(w',X')](x) = \frac{p_{(w,X)}(x)}{p_{(w',X')}(x)},$$
(3)

where $p_{(w,X)}$ is the probability density of x estimated from empirical weights w and samples X. The function L[(w,X),(w',X')](x) can be approximated using a classifier trained to distinguish (w,X) from (w',X'). This property has been successfully exploited using neural networks for full phase-space Monte Carlo reweighting and parameter estimation [13, 19–22]. Here, we use neural network classifiers to iteratively reweight the particleand detector-level Monte Carlo weights, resulting in an unfolding procedure.

The OMNIFOLD technique is illustrated in Fig. 1. Intuitively, synthetic detector-level events ("simulation") are reweighted to match experimental data ("data"), and then the reweighted synthetic events, now evaluated at particle-level ("generation"), are further reweighted to estimate the true particle-level information ("truth"). The starting point is a synthetic Monte Carlo dataset composed of pairs (t, m), where each particle-level event t is pushed through the detector simulation to obtain a detector-level event m. Particle-level events have initial weights $\nu_0(t)$, and when t is pushed to m, these become detector-level weights $\nu_0^{\text{push}}(m) = \nu_0(t)$. OMNIFOLD iterates the following steps:

1.
$$\omega_n(m) = \nu_{n-1}^{\text{push}}(m) L[(1, \text{Data}), (\nu_{n-1}^{\text{push}}, \text{Sim.})](m),$$

2. $\nu_n(t) = \nu_{n-1}(t) L[(\omega_n^{\text{pull}}, \text{Gen.}), (\nu_{n-1}, \text{Gen.})](t).$

The first step yields new detector-level weights $\omega_n(m)$, which are pulled back to particle-level weights $\omega_n^{\text{pull}}(t) = \omega_n(m)$ using the same synthetic pairs (t,m). Note that ν^{push} and ω^{pull} are not, strictly speaking, functions because of the multi-valued nature of the detector simulation. The second step ensures that ν_n is a valid weighting function of the particle-level quantities.

Assuming $\nu_0(t) = 1$, in the first iteration Step 1 learns $\omega_1(m) = p_{\text{Data}}(m)/p_{\text{Sim.}}(m)$, which is pulled back to the particle-level weights $\omega_1^{\text{pull}}(t)$. Step 2 simply converts the per-instance weights $\omega_1^{\text{pull}}(t)$ to a valid particle-level weighting function $\nu_1(t)$. After one iteration, the new induced truth is:

$$\nu_{1}(t) p_{\text{Gen.}}(t) = \int dm' \, p_{\text{Gen.}|\text{Sim.}}(t|m') \, p_{\text{Data}}(m'). \quad (4)$$



FIG. 1. An illustration of OMNIFOLD, applied to a set of synthetic and natural data. As a first step, starting from prior weights ν_0 , the detector-level synthetic data ("simulation") is reweighted to match the detector-level natural data (simply "data"). These weights ω_1 are pulled back to induce weights on the particle-level synthetic data ("generation"). As a second step, the initial generation is reweighted to match the new weighted generation. The resulting weights ν_1 are pushed forward to induce a new simulation, and the process is iterated.

This is a continuous version of IBU from Eq. (2), where the sum has been promoted to a full phase-space integral. After n iterations, the unfolded distribution is:

$$p_{\text{unfolded}}^{(n)}(t) = \nu_n(t) \, p_{\text{Gen.}}(t). \tag{5}$$

The unfolded result can be presented either as a set of generated events $\{t\}$ with weights $\{\nu_n(t)\}$ (and uncertainties) or, more compactly, as the learned weighting function ν_n and instructions for sampling from $p_{\text{Gen.}}$.

To demonstrate the versatility and power of OMNI-FOLD, we perform a proof-of-concept study relevant for the LHC. Specifically, we unfold the full radiation pattern (i.e. full phase space) of jets, which are collimated sprays of particles arising from the fragmentation and hadronization of high-energy quarks and gluons. Jets are an ideal environment in which to benchmark unfolding techniques, since detector effects often account for a significant portion of the experimental measurement uncertainties for many jet substructure observables [23]. With the radiation pattern unfolded, one can obtain the unfolded distribution of *any* observable using Eq. (5). Hence, this procedure can be viewed as simultaneously unfolding all observables.

Our study is based on proton-proton collisions generated at $\sqrt{s} = 14$ TeV with the default tune of HER-WIG 7.1.5 [24–26] and Tune 26 [27] of PYTHIA 8.243 [28– 30] in order to study a challenging setting where the "natural" and "synthetic" distributions are substantially dif-

3



FIG. 2. The unfolding results for six jet substructure observables, using HERWIG 7.1.5 ("Data"/"Truth") and PYTHIA 8.243 tune 26 ("Sim."/"Gen."), unfolded with OMNIFOLD and compared to IBU. OMNIFOLD matches or exceeds the unfolding performance of IBU on all of these observables. We emphasize that OMNIFOLD is a single general unfolding procedure, whereas unfolding with IBU must be done observable by observable. Statistical uncertainties are shown only in the ratio panel.

ferent. As a proxy for detector effects and a full detector simulation, we use the DELPHES 3.4.2 [31] fast simulation of the CMS detector, which uses particle flow reconstruction. Jets with radius parameter R = 0.4 are clustered using either all particle flow objects (detector-level) or stable non-neutrino truth particles (particle-level) with the anti- k_T algorithm [32] implemented in FASTJET 3.3.2 [33, 34]. One of the simulations (HERWIG) plays the role of "data"/"truth", while the other (PYTHIA) is used to derive the unfolding corrections. To reduce acceptance effects, the leading jets are studied in events with a Z boson with transverse momentum $p_T^Z > 200$ GeV. After applying the selections, we obtain approximately 1.6 million events from each generator.

Any suitable machine learning architecture can be used for OMNIFOLD. For this study, we use Particle Flow Networks (PFNs) [35, 36] to process jets in their natural representation as sets of particles. Intuitively, PFNs learn and processes a set of additive observables via $PFN(\{p_i\}_{i=1}^M) = F\left(\sum_{i=1}^M \Phi(p_i)\right)$ for an event with Mparticles p_i , where F and Φ are parameterized by fullyconnected networks. We specify the particles by their transverse momentum p_T , rapidity y, azimuthal angle ϕ , and particle identification code [37], restricted to the experimentally-accessible information (PFN-Ex [35]) at detector-level. To define separate models for Step 1 and Step 2, we use the PFN architecture and training parameters of Ref. [35] with latent space dimension $\ell = 256$, implemented in the ENERGYFLOW Python package [38]. Neural networks are trained with Keras [39] and Tensor-Flow [40] using the Adam [41] optimization algorithm. The models are randomly initialized in the first iteration and subsequently warm-started using the model from the previous iteration. 20% of the events are reserved as a validation set during training.

To investigate the unfolding performance, we consider six widely-used jet substructure observables [42]. The first four are jet mass m, constituent multiplicity M, the N-subjettiness ratio $\tau_{21} = \tau_2^{(\beta=1)}/\tau_1^{(\beta=1)}$ [43, 44], and the jet width w (implemented as $\tau_1^{(\beta=1)}$). Since jet grooming [45–49] is of recent interest, we also show the jet mass $\ln \rho = \ln m_{\rm SD}^2/p_T^2$ and momentum fraction z_g after Soft Drop grooming [48, 49] with $z_{\rm cut} = 0.1$ and $\beta = 0$. Several of these observables are computed with the help of FASTJET CONTRIB 1.042 [50].

The unfolding performance of OMNIFOLD is shown in

	Observable					
Method	m	M	w	$\ln \rho$	$ au_{21}$	z_g
OmniFold	2.77	0.33	0.10	0.35	0.53	0.68
ManyFold	3.80	0.89	0.09	0.37	0.26	0.15
UniFold	8.82	1.46	0.15	0.59	1.11	0.59
IBU	9.31	1.51	0.11	0.71	1.10	0.37
Data	24.6	130	15.7	14.2	11.1	3.76
Generation	3.62	15	22.4	19	20.8	3.84

TABLE I. The unfolding performance of OMNIFOLD, MULTI-FOLD, and UNIFOLD on six jet substructure observables, compared to IBU. The performance is quantified by the triangular discriminator [53–55] $\Delta(p,q) = \frac{1}{2} \int d\lambda \frac{(p(\lambda)-q(\lambda))^2}{p(\lambda)+q(\lambda)}$ (×10³) between the unfolded and truth-level (binned) histograms. Also shown are the distances from data (no unfolding) and generation (the prior). The best unfolding method for each observable is shown in bold. All methods perform well, with OMNIFOLD providing consistently good performance.

Fig. 2 and compared to IBU, both with n = 5 iterations. We found little difference between n = 3 and n = 5, though OMNIFOLD exhibits a slight preference for more iterations. OMNIFOLD succeeds in simultaneously unfolding all of these observables, achieving performance comparable to or better than IBU applied to each observable individually. The mass is challenging for all methods as particle-type information is relevant at particle-level but is not fully known at detector-level, introducing additional prior dependence.

Statistical uncertainties from the prior distribution are shown in the bottom panels of Fig. 2, holding the unfolding procedure (i.e. response matrix and reweighting) fixed. For this proof-of-concept study, we do not show systematic uncertainties, though the procedure for deriving them is the same as for IBU. Non-closure and modeling uncertainties can be derived in the standard way by testing the procedure on different Monte Carlo samples and comparing the results to the known "truth" distributions. (We checked that OMNIFOLD satisfies technical closure when PYTHIA is unfolded to itself.) Experimental systematic uncertainties can be obtained by varying the relevant effects and repeating the unfolding procedure. Like other unfolding procedures, OMNIFOLD cannot improve the results in phase-space regions that are unconstrained by observed quantities. It can, however, improve the performance if the full phase space contains auxiliary features relevant for the detector response. To capitalize on this full phase-space approach, it is essential that the detector simulation properly describes these features and that systematic uncertainties are estimated using a high-dimensional approach [51, 52].

To highlight the flexibility of our unfolding framework, we study variations of OMNIFOLD, where the available information is varied by controlling the inputs:

• UNIFOLD: A single observable as input. This is an unbinned version of IBU.



FIG. 3. The correlation dimension of the space of jets, unfolded with OMNIFOLD. The unfolded results closely match the truth-level dimension over most of the energy range, tending toward the prior in the more difficult phase space region at low Q. Unfolding a complicated statistic such as the correlation dimension is challenging with standard methods.

- MULTIFOLD: Many observables as input. Here, we use the six jet substructure observables in Fig. 2 to derive the detector response.
- OMNIFOLD: The full event (or jet) as input, using the full phase space information.

The unfolding performance of each method on our six substructure observables is tabulated in Table I and compared to IBU. The UNIFOLD and MULTIFOLD implementations both use dense networks with three layers of one hundred nodes each and a two-node output layer. We see good unfolding performance across all methods, and even though OMNIFOLD is not directly trained on these six observables, it performs comparably to or better than MULTIFOLD. While the detector response depends on the jet rapidity, we checked that MULTIFOLD did not significantly benefit from including the rapidity, though doing so could be important in a real experimental context. In general, additional information can be included and the unfolding procedure can be repeated, with the final model chosen as the one with the best detector-level agreement with the data.

Since OMNIFOLD unfolds the full radiation pattern, it can be used to probe new, physically-interesting quantities that are challenging to unfold with existing methods. One example is the recently-proposed correlation (fractal) dimension of jets [56, 57], which is a function of the energy scale Q. This complicated statistic is defined by pairwise metric distances between jet radiation patterns, falling outside of the purview of single-observable unfolding techniques. Within our jet samples, we restrict to energetic jets with $p_T^{\text{jet}} > 500 \text{ GeV}$, boosted to the origin of the rapidity-azimuth plane, and with constituents rescaled to have p_T summing to 500 GeV. The correlation dimensions of these jets, both before and after applying OMNIFOLD, are shown in Fig. 3. The unfolded results match the true distribution over a wide range of Q values, with residual prior dependence seen in at low Q (i.e. the infrared) where jets have a higher dimensionality and detector effects have a larger impact, thus making the unfolding problem more difficult. More broadly, OMNI-FOLD opens the door to going beyond per-event collider observables towards more nuanced or intricate measurements of the data.

In conclusion, we have presented a potentially transformative unfolding paradigm based on iteratively reweighting a set of simulated events with machine learning. Our OMNIFOLD approach allows an entire dataset to be unfolded using all of the available information, avoiding the need for binning and restricting to single observables. We have demonstrated the power of this method in a (simulated) case of interest by unfolding the full radiation pattern of jets, paving the way for significant advances in jet substructure at the LHC. Our unfolding framework allowed us to go beyond per-event observables towards unfolding more complex dataset statistics, such as fractal dimensions of the space of jets. Going even further, (unsupervised) machine learning models may be trained directly at particle-level by using the unfolded and weighted dataset, which is a fascinating avenue for further exploration. These advances have broad applicability beyond particle physics in domains where deconvolution or unfolding is used, such as image-based measurements and quantum computation [58]. To enable future unfolding studies and developments, we have made our code and jet datasets publicly available [59, 60], including two additional tunes of PYTHIA beyond those presented here. Finally, our reweighting-based unfolding strategy allows for new observables to be measured long after the unfolding is carried out, which can significantly empower future public and archival collider data analyses [61].

We are grateful to Gabriel Collin, Christian Herwig, Matthew LeBlanc, Salvatore Rappoccio, Jennifer Roloff, Matthew Schwartz, Daniel Whiteson, and Mike Williams for helpful comments and discussions. The authors benefited from the hospitality of the Harvard Center for the Fundamental Laws of Nature. This work was partially completed at the Aspen Center for Physics, which is supported by National Science Foundation grant PHY-1607611. AA and BN were supported by the U.S. Department of Energy (DOE), Office of Science under contract DE-AC02-05CH11231. PTK, EMM, and JT were supported by the DOE Office of Nuclear Physics under Grant No. DE-SC0011090 and the DOE Office of High Energy Physics under Grant Nos. DE-SC0012567 and DE-SC0019128. JT is additionally supported by the Simons Foundation through a Simons Fellowship in Theoretical Physics. Cloud computing resources were provided through a Google Cloud allotment from the MIT Quest for Intelligence. BN would like to thank NVIDIA for providing Volta GPUs for neural network training.

- * ajandreassen@google.com
- † pkomiske@mit.edu
- [‡] metodiev@mit.edu [§] hppschman@lbl.gov
- § bpnachman@lbl.gov
- [¶] jthaler@mit.edu
- G. Cowan, "A survey of unfolding methods for particle physics," Conf. Proc. C0203181, 248 (2002).
- [2] Volker Blobel, "Unfolding Methods in Particle Physics," PHYSTAT2011 Proceedings, 240 (2011).
- [3] Volker Blobel, "Unfolding," Data Analysis in High Energy Physics, 187 (2013).
- [4] Rahul Balasubramanian, Lydia Brenner, Carsten Burgard, Glen Cowan, Vincent Croft, Wouter Verkerke, and Pim Verschuuren, "Statistical method and comparison of different unfolding techniques using RooFit," (2019), arXiv:1910.14654 [physics.data-an].
- [5] G. D'Agostini, "A Multidimensional unfolding method based on Bayes' theorem," Nucl. Instrum. Meth. A362, 487–498 (1995).
- [6] Andreas Hocker and Vakhtang Kartvelishvili, "SVD approach to data unfolding," Nucl. Instrum. Meth. A372, 469–481 (1996), arXiv:hep-ph/9509307 [hep-ph].
- [7] Stefan Schmitt, "TUnfold: an algorithm for correcting migration effects in high energy physics," JINST 7, T10003 (2012), arXiv:1205.6201 [physics.data-an].
- [8] Nikolai D. Gagunashvili, "Machine learning approach to inverse problem and unfolding procedure," (2010), arXiv:1004.2006 [physics.data-an].
- [9] Alexander Glazov, "Machine learning as an instrument for data unfolding," (2017), arXiv:1712.01814 [physics.data-an].
- [10] Kaustuv Datta, Deepak Kar, and Debarati Roy, "Unfolding with Generative Adversarial Networks," (2018), arXiv:1806.00433 [physics.data-an].
- [11] G. Zech and B. Aslan, "Binning-Free Unfolding Based on Monte Carlo Migration," Statistical problems in particle physics, astrophysics and cosmology. Proceedings, Conference, PHYSTAT 2003, Stanford, USA, September 8-11, 2003, eConf C030908, TUGT001 (2003), [,138(2003)].
- [12] L. Lindemann and Gunter Zech, "Unfolding by weighting Monte Carlo events," Nucl. Instrum. Meth. A354, 516– 521 (1995).
- [13] Anders Andreassen and Benjamin Nachman, "Neural Networks for Full Phase-space Reweighting and Parameter Tuning," (2019), arXiv:1907.08209 [hep-ph].
- [14] A. Rogozhnikov, "Reweighting with Boosted Decision Trees," Proceedings, 17th International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT 2016): Valparaiso, Chile, January 18-22, 2016, J. Phys. Conf. Ser. **762**, 012036 (2016), arXiv:1608.05806 [physics.data-an].
- [15] Anders Andreassen, Ilya Feige, Christopher Frye, and Matthew D. Schwartz, "JUNIPR: a Framework for Un-

6

supervised Machine Learning in Particle Physics," Eur. Phys. J. C79, 102 (2019), arXiv:1804.09720 [hep-ph].

- [16] Emily Dickinson, "The Mountain Sat Upon the Plain," in *The Complete Poems of Emily Dickinson* (Little, Brown, and Company, 1924).
- [17] L. B. Lucy, "An iterative technique for the rectification of observed distributions," Astronomical Journal 79, 745 (1974).
- [18] William Hadley Richardson, "Bayesian-based iterative method of image restoration," J. Opt. Soc. Am. 62, 55– 59 (1972).
- [19] Kyle Cranmer, Juan Pavez, and Gilles Louppe, "Approximating Likelihood Ratios with Calibrated Discriminative Classifiers," (2015), arXiv:1506.02169 [stat.AP].
- [20] Johann Brehmer, Kyle Cranmer, Gilles Louppe, and Juan Pavez, "Constraining Effective Field Theories with Machine Learning," Phys. Rev. Lett. **121**, 111801 (2018), arXiv:1805.00013 [hep-ph].
- [21] Johann Brehmer, Kyle Cranmer, Gilles Louppe, and Juan Pavez, "A Guide to Constraining Effective Field Theories with Machine Learning," Phys. Rev. D98, 052004 (2018), arXiv:1805.00020 [hep-ph].
- [22] Johann Brehmer, Gilles Louppe, Juan Pavez, and Kyle Cranmer, "Mining gold from implicit models to improve likelihood-free inference," (2018), arXiv:1805.12244 [stat.ML].
- [23] Roman Kogler *et al.*, "Jet Substructure at the Large Hadron Collider: Experimental Review," (2018), arXiv:1803.06991 [hep-ex].
- [24] M. Bahr et al., "Herwig++ Physics and Manual," Eur. Phys. J. C58, 639–707 (2008), arXiv:0803.0883 [hep-ph].
- [25] Johannes Bellm et al., "Herwig 7.0/Herwig++ 3.0 release note," Eur. Phys. J. C76, 196 (2016), arXiv:1512.01178 [hep-ph].
- [26] Johannes Bellm *et al.*, "Herwig 7.1 Release Note," (2017), arXiv:1705.06919 [hep-ph].
- [27] ATLAS Run 1 Pythia8 tunes, Tech. Rep. ATL-PHYS-PUB-2014-021 (CERN, Geneva, 2014).
- [28] Torbjorn Sjöstrand, Stephen Mrenna, and Peter Z. Skands, "A Brief Introduction to PYTHIA 8.1," Comput. Phys. Commun. 178, 852–867 (2008), arXiv:0710.3820 [hep-ph].
- [29] Torbjorn Sjöstrand, Stephen Mrenna, and Peter Z. Skands, "PYTHIA 6.4 Physics and Manual," JHEP 05, 026 (2006), arXiv:hep-ph/0603175 [hep-ph].
- [30] Torbjorn Sjöstrand, Stefan Ask, Jesper R. Christiansen, Richard Corke, Nishita Desai, Philip Ilten, Stephen Mrenna, Stefan Prestel, Christine O. Rasmussen, and Peter Z. Skands, "An Introduction to PYTHIA 8.2," Comput. Phys. Commun. **191**, 159–177 (2015), arXiv:1410.3012 [hep-ph].
- [31] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lematre, A. Mertens, and M. Selvaggi (DELPHES 3), "DELPHES 3, A modular framework for fast simulation of a generic collider experiment," JHEP 02, 057 (2014), arXiv:1307.6346 [hep-ex].
- [32] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez, "The Anti-k(t) jet clustering algorithm," JHEP 04, 063 (2008), arXiv:0802.1189 [hep-ph].
- [33] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez, "FastJet User Manual," Eur. Phys. J. C72, 1896 (2012), arXiv:1111.6097 [hep-ph].
- [34] Matteo Cacciari and Gavin P. Salam, "Dispelling the N^3 myth for the k_t jet-finder," Phys. Lett. B641, 57–

61 (2006), arXiv:hep-ph/0512210 [hep-ph].

- [35] Patrick T. Komiske, Eric M. Metodiev, and Jesse Thaler, "Energy Flow Networks: Deep Sets for Particle Jets," JHEP 01, 121 (2019), arXiv:1810.05165 [hep-ph].
- [36] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabás Póczos, Ruslan Salakhutdinov, and Alexander J. Smola, "Deep sets," in Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA (2017) pp. 3391–3401.
- [37] M. Tanabashi *et al.* (Particle Data Group), "Review of Particle Physics," Phys. Rev. D98, 030001 (2018).
- [38] Patrick T. Komiske and Eric M. Metodiev, "EnergyFlow," (2019), https://energyflow.network/.
- [39] Francois Chollet, "Keras," https://github.com/ fchollet/keras (2017).
- [40] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al., "Tensorflow: A system for large-scale machine learning." in OSDI, Vol. 16 (2016) pp. 265–283.
- [41] Diederik Kingma and Jimmy Ba, "Adam: A method for stochastic optimization," (2014), arXiv:1412.6980 [cs].
- [42] Andrew J. Larkoski, Ian Moult, and Benjamin Nachman, "Jet Substructure at the Large Hadron Collider: A Review of Recent Advances in Theory and Machine Learning," (2017), arXiv:1709.04464 [hep-ph].
- [43] Jesse Thaler and Ken Van Tilburg, "Identifying Boosted Objects with N-subjettiness," JHEP 03, 015 (2011), arXiv:1011.2268 [hep-ph].
- [44] Jesse Thaler and Ken Van Tilburg, "Maximizing Boosted Top Identification by Minimizing N-subjettiness," JHEP 02, 093 (2012), arXiv:1108.2701 [hep-ph].
- [45] David Krohn, Jesse Thaler, and Lian-Tao Wang, "Jet Trimming," JHEP 02, 084 (2010), arXiv:0912.1342 [hepph].
- [46] Stephen D. Ellis, Christopher K. Vermilion, and Jonathan R. Walsh, "Recombination Algorithms and Jet Substructure: Pruning as a Tool for Heavy Particle Searches," Phys. Rev. D81, 094023 (2010), arXiv:0912.0033 [hep-ph].
- [47] Stephen D. Ellis, Christopher K. Vermilion, and Jonathan R. Walsh, "Techniques for improved heavy particle searches with jet substructure," Phys. Rev. D80, 051501 (2009), arXiv:0903.5081 [hep-ph].
- [48] Mrinal Dasgupta, Alessandro Fregoso, Simone Marzani, and Gavin P. Salam, "Towards an understanding of jet substructure," JHEP 09, 029 (2013), arXiv:1307.0007 [hep-ph].
- [49] Andrew J. Larkoski, Simone Marzani, Gregory Soyez, and Jesse Thaler, "Soft Drop," JHEP 05, 146 (2014), arXiv:1402.2657 [hep-ph].
- [50] "Fastjet contrib," https://fastjet.hepforge.org/ contrib/.
- [51] Benjamin Nachman and Chase Shimmin, "AI Safety for High Energy Physics," (2019), arXiv:1910.08606 [hepph].
- [52] Benjamin Nachman, "A guide for deploying Deep Learning in LHC searches: How to achieve optimality and account for uncertainty," (2019), arXiv:1909.03081 [hepph].
- [53] F. Topsoe, "Some inequalities for information divergence and related measures of discrimination," IEEE Transactions on Information Theory 46, 1602–1609 (2000).

- [54] Philippe Gras, Stefan Hche, Deepak Kar, Andrew Larkoski, Leif Lnnblad, Simon Pltzer, Andrzej Sidmok, Peter Skands, Gregory Soyez, and Jesse Thaler, "Systematics of quark/gluon tagging," JHEP 07, 091 (2017), arXiv:1704.03878 [hep-ph].
- [55] Samuel Bright-Thonney and Benjamin Nachman, "Investigating the Topology Dependence of Quark and Gluon Jets," JHEP 03, 098 (2019), arXiv:1810.05653 [hep-ph].
- [56] Patrick T. Komiske, Eric M. Metodiev, and Jesse Thaler, "Metric Space of Collider Events," Phys. Rev. Lett. 123, 041801 (2019), arXiv:1902.02346 [hep-ph].
- [57] Patrick T. Komiske, Radha Mastandrea, Eric M. Metodiev, Preksha Naik, and Jesse Thaler, "Exploring the Space of Jets with CMS Open Data," (2019), arXiv:1908.08542 [hep-ph].
- [58] B. Nachman, M. Urbanek, W. A. de Jong, and C. Bauer, "Unfolding Quantum Computer Readout Noise," (2019), arXiv:1910.01969 [quant-ph].
- [59] Anders Andreassen, Patrick T. Komiske, Eric M. Metodiev, Benjamin Nachman, and Jesse Thaler, "Omnifold," https://github.com/ericmetodiev/OmniFold (2019).
- [60] Anders Andreassen, Patrick Komiske, Eric Metodiev, Benjamin Nachman, and Jesse Thaler, "Pythia/Herwig + Delphes Jet Datasets for OmniFold Unfolding," Zenodo (2019), 10.5281/zenodo.3548091.

[61] Xiaoli Chen, Sünje Dallmeier-Tiessen, Robin Dasler, Sebastian Feger, Pamfilos Fokianos, Jose Benito Gonzalez, Harri Hirvonsalo, Dinos Kousidis, Artemis Lavasa, Salvatore Mele, Diego Rodriguez Rodriguez, Tibor Šimko, Tim Smith, Ana Trisovic, Anna Trzcinska, Ioannis Tsanaktsidis, Markus Zimmermann, Kyle Cranmer, Lukas Heinrich, Gordon Watts, Michael Hildreth, Lara Lloret Iglesias, Kati Lassila-Perini, and Sebastian Neubert, "Open is not enough," Nature Physics 15, 113–119 (2019).

Emily Dickinson, #975

The Mountain sat upon the Plain In his tremendous Chair – His observation omnifold, His inquest, everywhere –

The Seasons played around his knees Like Children round a sire – Grandfather of the Days is He Of Dawn, the Ancestor –

