

# UC Berkeley

## UC Berkeley Previously Published Works

### Title

Neural networks for full phase-space reweighting and parameter tuning

### Permalink

<https://escholarship.org/uc/item/9dq4j55f>

### Journal

Physical Review D, 101(9)

### ISSN

2470-0010

### Authors

Andreassen, Anders

Nachman, Benjamin

### Publication Date

2020-05-01

### DOI

10.1103/physrevd.101.091901

Peer reviewed

**Neural networks for full phase-space reweighting and parameter tuning**Anders Andreassen<sup>1,2,\*</sup> and Benjamin Nachman<sup>2,†</sup><sup>1</sup>*Department of Physics, University of California, Berkeley, California 94720, USA*<sup>2</sup>*Physics Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

(Received 17 August 2019; accepted 22 April 2020; published 12 May 2020)

Precise scientific analysis in collider-based particle physics is possible because of complex simulations that connect fundamental theories to observable quantities. The significant computational cost of these programs limits the scope, precision, and accuracy of Standard Model measurements and searches for new phenomena. We therefore introduce *Deep neural networks using Classification for Tuning and Reweighting* (DCTR), a neural network-based approach to reweight and fit simulations using all kinematic and flavor information—the full phase space. DCTR can perform tasks that are currently not possible with existing methods, such as implementing uncertainties on nonperturbative models without rerunning the simulation. The core idea behind the new approach is to exploit powerful high-dimensional classifiers to reweight phase space as well as to identify the best parameters for describing data. Numerical examples from  $e^+e^- \rightarrow$  jets demonstrate the fidelity of these methods for simulation parameters that have a big and broad impact on phase space as well as those that have a minimal and/or localized impact. The high fidelity of the full phase-space reweighting enables a new paradigm for simulations, parameter tuning, and model systematic uncertainties across particle physics and possibly beyond.

DOI: [10.1103/PhysRevD.101.091901](https://doi.org/10.1103/PhysRevD.101.091901)**I. INTRODUCTION**

In collider-based high-energy physics, parton-, particle-, and detector-level Monte Carlo (MC) simulation programs enable scientific inference by connecting fundamental theories to observable quantities. However, these tools are often computationally slow and emulate probability distributions that are analytically intractable. This has resulted in three key simulation challenges for particle physics: (1) an insufficient number of simulated events, (2) unaccounted for biases from simulation parameters, and (3) the inability to utilize all kinematic and flavor information (“the full phase space”) for parameter tuning.

A variety of approaches have been proposed to address the above challenges. The two existing solutions to (1) are to use more [1–3] and/or faster computers or accelerators [4,5] or to build fast surrogate models (“fast simulation”). Machine learning tools hold great promise for augmenting [6] or replacing [7–20] current fast detector simulation approaches, but are not yet precise enough to match the full, physics-based detector simulators that are often the limiting factor in the overall software pipeline. Deep

learning methods to circumvent expensive simulations for hypothesis testing were studied in the context of effective field theory fits [21–23]; related ideas will be useful also for reweighting. The only solution for (2) aside from generating a large set of simulations or interpolating between bins of low-dimensional histograms [24] is to assign event weights for parameter variations. Currently, this is only possible for a small number of perturbative parameters in parton shower programs [25–27] and for parton distribution functions [28,29]. Semiautomated procedures exist for tuning parton shower models [24,30], but the format of the existing public data means that these algorithms are restricted to a set of mostly one-dimensional inputs that must be assumed to be independent. The variational method proposed in Ref. [31] has been demonstrated with high-dimensional data, but utilizes a minimax optimization technique and requires running the simulator many times during training.

This paper introduces *Deep neural networks using Classification for Tuning and Reweighting* (DCTR, pronounced “doctor”), a new approach to solve all three computational challenges. In particular, deep neural network-based classifiers are used to (continuously) reweight one particle-level simulation into another and additionally use the full phase space to fit parameters within a given model. When the nominal particle-level sample has a corresponding detector-level simulation, then this procedure produces a new detector-level sample as well. Nondeep machine learning tools have been used in the past for discrete reweighting [18,32–34] with a small

\*andersja@berkeley.edu

†bpnachman@lbl.gov

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 SCOAP<sup>3</sup>.

number of observables. Deep-learning-based discrete weighting was considered in [18,35] and continuous single observable reweightings were presented in [36]. The reweighting presented here combines a full phase-space deep learning architecture [37] with parametrization [38,39] to fully morph one simulation into another. This approach can naturally accommodate high-dimensional feature spaces and many model parameters. In addition to reweighting, we show how DCTR can be used with a differentiable reweighting function (such as the one just mentioned) to optimize simulation parameters. Fitting parameters based on the parametrized classifiers was proposed in Ref. [38]; here, the fitting procedure uses a classifier to construct the loss function, which can readily incorporate all of the information from the (potentially high-dimensional) input features and be optimized using standard deep learning tools.

## II. METHODS

The first ingredient to the full phase-space reweighting procedure is a prescription to derive event weights. Consider two simulations that describe the same phase space  $\Omega$  and are described by probability densities  $p_0(x)$  and  $p_1(x)$ , for  $x \in \Omega$ . Assuming that  $p_0$  and  $p_1$  have the same support,<sup>1</sup> the function  $w(x) = p_0(x)/p_1(x)$  is the ideal per-event weight to morph the second simulation into the first one. A well-known key observation is that  $w$  can be well approximated by training a machine learning classifier to distinguish the two simulations. For example, let  $f(x)$  be a neural network and trained with the binary cross-entropy loss:

$$\text{loss}(f(x)) = -\sum_{i \in \mathbf{0}} \log f(x_i) - \sum_{i \in \mathbf{1}} \log(1 - f(x_i)), \quad (1)$$

where  $\mathbf{0}$  and  $\mathbf{1}$  represent sets of examples from the two simulations. Then a well-known result is that,  $f(x)/(1 - f(x)) \approx p_0(x)/p_1(x)$ . The benefit of parametrizing  $f$  as a neural network is that deep learning can readily analyze all of  $\Omega$ , which was not possible with shallow learning attempts with a similar statistical foundation. The closest attempt to a full phase-space approach directly tried to learn  $p_i(x)$  using the full kinematic (i.e., nonflavor) part of  $\Omega$  [35,40], but this is much harder than learning the ratio.

An important reweighting scenario is when the two simulations are from the same simulation program, but with a subset of different model parameters,  $\theta$ . For example, when model uncertainties are evaluated, one may want to transform  $p_\theta(x)$  into  $p_{\theta+\delta_\theta}(x)$ . When these uncertainties

are profiled in a fit, it is important that the transformation procedure be able to continuously interpolate between model parameters. The neural network reweighting approximation can be extended to this continuous case by adding  $\theta$  as a feature [38,39]:  $f(x, \theta)$ . In the examples presented below, the training data are generated with a uniform distribution in  $\theta$ , but this probability density can be optimized per application and can even be discrete.

Even though generators have many parameters that must be fit to data, gradient methods cannot be used directly with the models as the full phase space they produce is not usually differentiable (or at least the derivative is intractable) with respect to their model parameters. Surrogate generative models built from neural networks can be used for gradient-based parameter fitting, but may not have sufficient quality to be reliable. Reweighting is a robust alternative to surrogate generative models. A neural network-based continuous reweighting function is essentially a differentiable (in model parameters) version of the original simulator and can be used to perform inference on the parameters themselves. This is especially powerful for particle-level parameter tuning to data where one sample with a computational expensive full detector simulation can be continuously reweighted to other parameter points with the same detector model at no extra simulation cost.

An ideal loss function used to fit model parameters makes use of the full observable phase space. Typical metrics such as the  $\chi^2$  between histogram approximations to probability densities become impractical when  $\Omega$  is high dimensional because the number samples needed to populate the histogram bins grows exponentially with the number of features. As described above, classifiers are powerful tools for accessing all of the available information. Therefore, one can use a classifier *for the loss*. When a classifier trained to distinguish some  $\theta_0$  from a  $\theta_1$  performs poorly, then the two samples are close. As in Eq. (1), the bold notation represents sets of examples from two simulations, in this case distinguished by their value of  $\theta$ . While using classification to quantify differences between event samples has been used for anomaly detection [41–43], we are unaware of an example where it is used for parameter fitting. The idea of using the classifier loss as a metric is similar to the minimax strategy in generative adversarial networks [44], only in this context the generative part is a reweighter and is trained independently. A more elegant way of implementing this approach is to fit unknown parameters to the values that minimize the nominal classifier loss. In particular, suppose that a reweighter neural network  $f$  is trained as described above. Such a function will satisfy

$$f(x, \theta) = \operatorname{argmax}_{f': \Omega \rightarrow \mathbb{R}} \sum_{i \in \theta_0} \log f'(x_i, \theta) + \sum_{i \in \theta} \log(1 - f'(x_i, \theta)) \quad (2)$$

<sup>1</sup>In most physical applications, this is always the case. If there are regions where  $p_0(x)/p_1(x)$  is far from unity, one can add a regularization parameter to the training to mitigate large weights, which may significantly reduce the statistical power of the reweighted dataset. We found that this works well, but was unnecessary for the examples presented in this paper.

for all  $\theta$ . Note that the  $f'$  in the first sum takes the parameter  $\theta$  and not  $\theta_0$ , otherwise the discrimination task would be trivial. Now, suppose there is a new sample  $\theta_1$  where  $\theta_1$  is unknown (for instance,  $\theta_1$  are collider data). The claim is that if  $\theta^*$  is chosen as

$$\theta^* = \operatorname{argmax}_{\theta'} \sum_{i \in \theta_0} \log f(x_i, \theta') + \sum_{i \in \theta_1} \log(1 - f(x_i, \theta')) \quad (3)$$

then  $\theta^* = \theta_1$ . As  $f$  minimizes the cross-entropy loss for any  $\theta$  [Eq. (2)],

$$\begin{aligned} & \sum_{i \in \theta_0} \log f(x_i, \theta_1) + \sum_{i \in \theta_1} \log(1 - f(x_i, \theta_1)) \\ & \geq \sum_{i \in \theta_0} \log f(x_i, \theta^*) + \sum_{i \in \theta_1} \log(1 - f(x_i, \theta^*)) \end{aligned} \quad (4)$$

must hold. However, the converse must also be true since  $\theta^*$  minimizes the cross-entropy loss as well and therefore,  $\theta^* = \theta_1$ . Since  $f$  is differentiable, Eq. (3) can be solved using standard gradient-based methods. While Eq. (3) performs the fit on the same particle-level phase space as the reweighting, it can be readily extended to do the fitting (via the classification loss) at detector level while the reweighting can be performed at particle level using one fully detector-simulated event sample.

The last ingredient of DCTR is a suitable neural network architecture that can effectively capture all the salient features of  $\Omega$ . A natural tool for this task is the particle flow network (PFN) [37] built on the deep sets framework [45]. While many deep learning architectures incorporate the symmetries and structure of high-energy physics events [8,35,46–56], PFNs are particularly effective because they can operate on variable-length sets of particles and respect the quantum-mechanically induced permutation invariance of particle labels. These networks can also readily incorporate nonkinematic information such as particle flavor. A particle flow network is a composition of two neural networks  $F$  and  $\Phi$ :  $f(\{p_i\}) = F(\sum_{i=1}^n \Phi(p_i))$ , where  $p_i$  is the set of features that belongs to particle  $i$  (momentum and flavor) as well as  $\theta$ . The function  $\Phi$  embeds the input particles into an  $\ell$ -dimensional *latent space* and  $F$  is a simple  $\mathbb{R}^\ell \mapsto \mathbb{R}$  neural network. References [37,45] proved that this structure is sufficiently flexible to approximate any function and in practice,  $\ell \sim \mathcal{O}(10)$ .

### III. RESULTS

To illustrate the potential of DCTR, full phase-space reweighting and parameter tuning is performed on a sample of generated events from the PYTHIA 8.230 [57,58] event generator. Particle-level  $e^+e^- \rightarrow Z \rightarrow$  dijet events with about 100 particles in each event are clustered into jets using the anti- $k_r$  clustering algorithm [59] ( $R = 0.8$ ) with FASTJET 3.0.3 [60,61]. The jets are presented to the neural network for training, with each jet constituent represented

by  $(p_T, \eta, \phi, \text{particle type}, \theta)$ , where  $\theta$  is the parameter in Eq. (2). One million events were generated for each set of PYTHIA parameters. In addition to a default parameter set using the Monash tune [62], three separate samples were generated with uniformly sampled (per event) `TimeShower:alphaSvalue`, `StringZ:aLund`, and `StringFlav:probStoUD` in the ranges [0.10, 0.18], [0.50, 0.90], and [0.10, 0.30], respectively. We also generated one sample where all three parameters were simultaneously uniformly sampled. These parameters were chosen because they represent both perturbative and non-perturbative physical effects and the ranges are similar to those studied in Ref. [30]. The Monash values of the three parameters are 0.1365, 0.68, and 0.217, respectively.

The reweighting and fitting was found to work well without any hyperparameter modifications from Ref. [37]. In particular,  $\Phi$  has two hidden layers with  $\ell = 128$  and  $F$  is composed of three hidden layers and two output nodes for binary classification, and all the hidden layers have 100 nodes. The activation function used for all layers is ReLU with the exception of the classification output which uses softmax. All models were implemented in KERAS [63] with the TensorFlow backend [64] and trained using the cross-entropy loss with the Adam [65] optimizer for 50 epochs, using early stopping with patience 10, with batch size 1000. Each training set contained  $8 \times 10^5$  training and  $10^5$  validation jets of each class. Training time was 10–15 min for each model (20 sec per epoch) on an NVIDIA GeForce GTX 1080. The code and data for this paper are available at Ref. [66] and [67], respectively.

As a first test of DCTR, a single parameter (`TimeShower:alphaSvalue`) is reweighted using the full phase space of the generated jets. Results for discrete and continuous reweighting from a varied parameter to the nominal sample are presented in Fig. 1. The entire phase space is reweighted, but is too high dimensional to visualize. Instead, three histograms of physically relevant one-dimensional observables are presented: the number of particles inside the jet (multiplicity), an  $n$ -subjettiness ratio  $\tau_{32}$  [68,69], and a four-point energy correlation function [70]  $\text{ECF}(N=4, \beta=4)$ . By definition,  $\tau_{32} = \tau_3/\tau_2$  where  $\tau_n = \sum_{i \in \text{jet}} p_{T,i} \min_{j=1 \dots n} \{\Delta R(i, j)\}$  for axis  $j$ ; likewise,  $\text{ECF}(N, \beta)$  is the sum over all quadruples inside the jet weighted by the product of the momenta and the product of all opening angles raised to the power  $\beta$ . The large values of  $n$ ,  $N$ , and  $\beta$  are used to expose complex features with a nontrivial dependence on all particles inside the jet. Many more observables were studied, but these are representative.

The reweighted distributions are in excellent agreement with the target nominal distribution. Samples used to make the histograms shown for `TimeShower:alphaSvalue` values 0.1365 and 0.1600 were not used during training or validation. The fidelity of a continuous reweighting is quantified in the lower right plot of Fig. 1, which presents the  $\chi^2/\text{ndf}$  as a function of the initial  $\alpha_s$  parameter value.

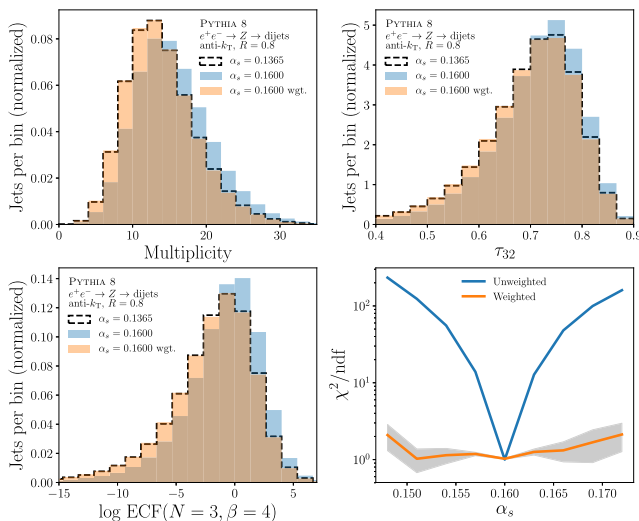


FIG. 1. The three histograms show the result before and after reweighting between two values of `TimeShower:alphaSvalue` =  $\alpha_s$  on different 1D observables. To quantify the quality of the reweighting, and to illustrate one trained model can continuously reweight for any parameter, we show the  $\chi^2/\text{ndf}$  for multiplicity as a function of  $\alpha_s$  in the lower right plot for reweighting to  $\alpha_s = 0.1600$ . For each value, we compare the  $\chi^2$  relative to  $\alpha_s = 0.1600$  before and after reweighting. Each  $\chi^2$  value is averaged over ten runs and the gray band marks the standard deviation, which is consistent with  $\chi^2/\text{ndf} \approx 1$ .

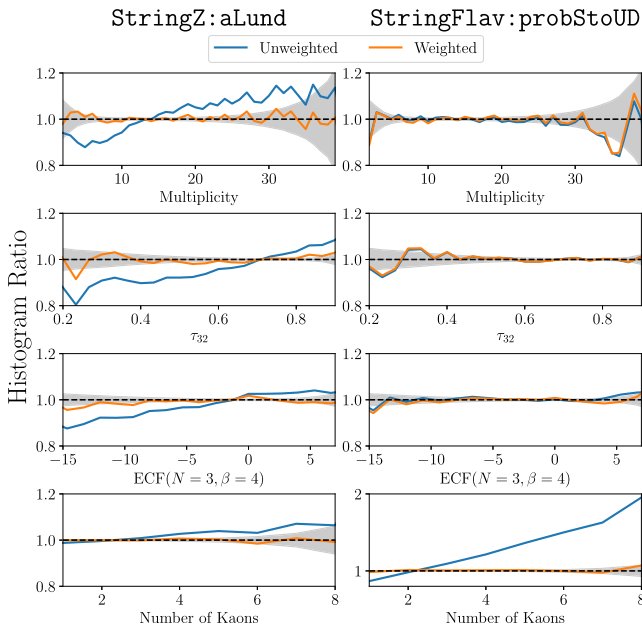


FIG. 2. Ratio of histograms from nominal distribution to sample generated with `StringZ:aLund` = 0.8 on the left and `StringFlav:probStoUD` = 0.275 on the right. Both unweighted and weighted histograms ratios are shown. The gray band indicates the statistical uncertainty from both the nominal and variation sample. After reweighting, the ratio only differs from 1 within the statistical uncertainty.

Variations in `TimeShower:alphaSvalue` modify many aspects of jet fragmentation and therefore it may be an easy parameter for the reweighting network to learn. In contrast, the hadronization parameters `StringZ:aLund` and `StringFlav:probStoUD` may be more difficult because the size of their effects on the phase space is small and/or localized. Figure 2 shows that despite these potential challenges, the reweighting procedure is able to effectively capture subtle and isolated modifications to the phase space. Variations in the `StringZ:aLund` parameter result in mostly percent-level differences in the presented distributions, which are corrected in the reweighted model. Modifying the `StringFlav:probStoUD` parameter only changes strange particles such as kaons, which highlights the importance of including flavor in the full phase-space network. Importantly, this model learns to only change the distributions related to strange particles, leaving other observables untouched. Simultaneously reweighting all three parameters also works well, but is more difficult to visualize. We also verified that DCTR works for  $pp$  MC simulations by reweighting from PYTHIA to HERWIG for both the quark and gluon samples taken from [37,71–73].

The well-trained DCTR model can now be used to demonstrate the potential for parameter tuning following Eq. (3). As a first step, Table I presents the result of a fit where the “data” are the same as the nominal, but with each parameter changed one at a time (each row is a separate fit). To illustrate the sensitivity to the randomness in the model initialization, each fit is performed ten times. This variation could be reduced with a more sophisticated neural network and/or more training data. For each of these one-dimensional fits, the fitted value is consistent with the target value within these statistical fluctuations from initialization, which are 1%–3%. As `TimeShower:alphaSvalue` has a bigger impact on the phase space, it is less sensitive to the initialization statistical fluctuations. For a fit with data, the statistical and systematic uncertainty could be determined with toys and even profiled, as is standard for parameter fitting. A separate test where a three-dimensional model is used to do these one-dimensional fits yielded the same accuracy and similar precision.

As a next step, the top part of Table II shows the result of a simultaneous fit to the three parameters. As with the one-dimensional fit, the fitted values are all statistically

TABLE I. Independent fit for simulation where one parameter was changed at a time. The reported numbers are the mean and standard deviation over ten runs with different model initializations.

Parameter	Target value	Fit value
<code>TimeShower:alphaSvalue</code>	0.1600	$0.1601 \pm 0.0018$
<code>StringZ:aLund</code>	0.8000	$0.7980 \pm 0.0257$
<code>StringFlav:probStoUD</code>	0.2750	$0.2754 \pm 0.0065$

TABLE II. Simultaneous fit for three parameters. The top row shows the results for the validation fit where we knew the target parameters, and the bottom row is the blinded fit. The reported numbers are the mean and standard deviation over 20 runs with different model initializations.

	Parameter	Target value	Fit value
Validation	TimeShower:alphaSvalue	0.1200	$0.1195 \pm 0.0022$
	StringZ:aLund	0.6000	$0.6276 \pm 0.0373$
	StringFlav:probStoUD	0.1200	$0.1203 \pm 0.0071$
Blinded	TimeShower:alphaSvalue	0.1700	$0.1707 \pm 0.0022$
	StringZ:aLund	0.7500	$0.7425 \pm 0.0453$
	StringFlav:probStoUD	0.1400	$0.1422 \pm 0.0065$

consistent with the target values. Interestingly, the sensitivity to the initialization statistical fluctuations is about the same for the three-dimensional fit as for the one-dimensional fits, providing confidence in the scaling to more parameters. In practice, the fitting procedure would be validated on a variety of simulations with known parameters, as just described. An illustration of the fit itself is shown in Fig. 3, where a two-dimensional slice through the likelihood landscape is presented and the fit execution demonstrated with markers and dashed lines. The broadness of the loss in the `StringZ:aLund` direction relative to the `TimeShower:alphaSvalue` one is a reflection of the significantly smaller impact of fragmentation function variations on the observable phase space compared with modifications to the final state shower strong coupling. After the validation, the model can be deployed on data, where the parameters are unknown. The lower part of Table II replicates this scenario, where the PYTHIA parameters were blinded during the fit. This closure test indicates that the method is robust to user bias.

One key challenge with any weighting method is statistical dilution. The MC statistical precision of the expectation value of an observable scales with

$1/\sqrt{\sum_i w_i^2}$ , for event weight  $w_i$ . When the  $w_i$  are far from unity, the uncertainty can be much larger than the unweighted uncertainty  $1/\sqrt{n}$ . For this reason, weighting methods are most useful when the initial and target samples have a similar phase space. This is usually the case for the systematic uncertainty variation as well as MC tuning examples presented in this paper.

#### IV. CONCLUSIONS

The empirical results demonstrate that DCTR is ready to be deployed for full phase-space tasks. The discrete reweighting could be used to generate new full-detector-simulated samples with a different particle-level simulation when at least one fully simulated sample exists. This could be particularly useful for systematic uncertainties computed using pairs of simulations (e.g., comparing PYTHIA and HERWIG) and for legacy data analysis in which the original detector simulation is no longer available [74]. Continuous reweighting will enable systematic parameter variations for uncertainty estimation that were not possible before (most parameters). Such variations can even be profiled during any statistical test that fits phase-space regions sensitive to the varied nuisance parameters. Finally, the full power of DCTR can be used for parameter tuning. Unlike traditional tuning which uses unfolded data that are usually one dimensional and without observable-observable correlations, a new paradigm is now possible where high-dimensional detector-level data can be used directly. The full power of the data can be utilized and all of the correlations are correctly accounted for in the fit. For the first time, this may allow for proper observable-observable covariance matrices (and thus correlated uncertainties) to be determined for simulation parameter values. While the numerical results in this paper have used collider-based particle physics examples, the techniques are directly applicable to any physics domain where tunable simulations are used for inference. Examples include cosmic ray physics (e.g., CORSIKA [75]), neutrino physics (e.g., GENIE [76]), dark matter direct detection (e.g., NEST [77]), astrophysics (e.g., LENSPOD [78]), cosmology (e.g., GADGET [79]), and nuclear physics (e.g., HIJING [80]). All of these opportunities illustrate the broad applicability of full phase-space reweighting and parameter tuning and the power

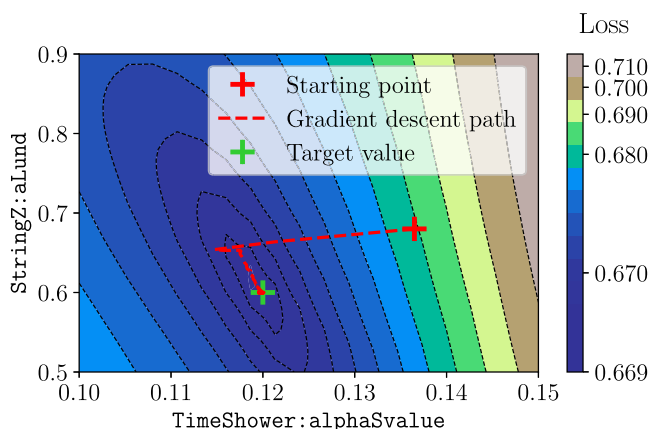


FIG. 3. Two-dimensional slice through the loss surface for the fit described in Table II. Markers indicate the starting point at nominal values, the gradient descent path, and the target values. From the starting point, gradient descent using Adam overshoots the minimum in its first two epochs before it converges to the target value.

of DCTR to extend the scope, precision, and accuracy of collider-based particle physics analyses and beyond.

### ACKNOWLEDGEMENTS

B.N. would like to thank Luke de Oliveira, Michela Paganini, Chase Shimmmin, and Paul Tipton for collaboration on an early stage prototype of this project. We thank Steve Mrenna and Peter Skands for helpful discussions about automated variations in PYTHIA. We also thank Kyle

Cranmer, Aviv Cukierman, Patrick Komiske, Eric Metodiev, and Jesse Thaler for lively discussions about deep learning-based reweighting. Additionally, we are grateful to Kyle Cranmer, Phil Ilten, and Jesse Thaler for feedback on the manuscript. This work was supported by the U.S. Department of Energy, Office of Science under Contract No. DE-AC02-05CH11231. Finally, we are grateful for the opportunity to use the Cori supercomputing resources at National Energy Research Scientific Center (NERSC).

- 
- [1] S. Ahn, J. Apostolakis, M. Asai, D. Brandt, G. Cooperman, G. Cosmo, A. Dotti, X. Dong, S. Yung Jun, and A. Nowak, *GEANT4-MT : bringing multi-threading into GEANT4 production* (2014), p. 04213, <https://doi.org/10.1051/snmc/201404213>.
- [2] S. Farrell, A. Dotti, M. Asai, P. Calafiura, and R. Monnard, [arXiv:1605.08371](https://arxiv.org/abs/1605.08371).
- [3] J. T. Childers, T. D. Uram, T. J. LeCompte, M. E. Papka, and D. P. Benjamin, *Comput. Phys. Commun.* **210**, 54 (2017).
- [4] O. Seiskari, J. Kommeri, and T. Niemi, [arXiv:1209.5235](https://arxiv.org/abs/1209.5235).
- [5] P. Canal, D. Elvira, S. Y. Jun, J. Kowalkowski, M. Paterno, and J. Apostolakis, *J. Phys. Conf. Ser.* **513**, 052013 (2014).
- [6] ATLAS Collaboration, CERN Report No. ATL-SOFT-PUB-2018-002, 2018.
- [7] M. Paganini, L. de Oliveira, and B. Nachman, *Phys. Rev. D* **97**, 014021 (2018).
- [8] L. de Oliveira, M. Paganini, and B. Nachman, *Comput. Softw. Big Sci.* **1**, 4 (2017).
- [9] M. Paganini, L. de Oliveira, and B. Nachman, *Phys. Rev. Lett.* **120**, 042003 (2018).
- [10] L. de Oliveira, M. Paganini, and B. Nachman, *J. Phys. Conf. Ser.* **1085**, 042017 (2018).
- [11] M. Erdmann, L. Geiger, J. Glombitza, and D. Schmidt, *Comput. Softw. Big Sci.* **2**, 4 (2018).
- [12] P. Musella and F. Pandolfi, *Comput. Softw. Big Sci.* **2**, 8 (2018).
- [13] M. Erdmann, J. Glombitza, and T. Quast, *Comput. Software Big Sci.* **3**, 4 (2019).
- [14] F. Carminati, A. Gheata, G. Khattak, P. Mendez Lorenzo, S. Sharan, and S. Vallecorsa, *J. Phys. Conf. Ser.* **1085**, 032016 (2018).
- [15] V. Chekalina, E. Orlova, F. Ratnikov, D. Ulyanov, A. Ustyuzhanin, and E. Zakharov, *EPJ Web Conf. Ser.* **214**, 02034 (2019).
- [16] B. Hashemi, N. Amin, K. Datta, D. Olivito, and M. Pierini, [arXiv:1901.05282](https://arxiv.org/abs/1901.05282).
- [17] R. Di Sipio, M. Faucci Giannelli, S. Ketabchi Haghighat, and S. Palazzo, *J. High Energy Phys.* **08** (2019) 110.
- [18] R. Aaij *et al.* (LHCb Collaboration), *Phys. Rev. Lett.* **119**, 062001 (2017).
- [19] J. W. Monk, *J. High Energy Phys.* **12** (2018) 021.
- [20] ATLAS Collaboration, CERN Report No. ATL-SOFT-PUB-2018-001, 2018.
- [21] J. Brehmer, K. Cranmer, G. Louppe, and J. Pavez, *Phys. Rev. Lett.* **121**, 111801 (2018).
- [22] J. Brehmer, K. Cranmer, G. Louppe, and J. Pavez, *Phys. Rev. D* **98**, 052004 (2018).
- [23] J. Brehmer, G. Louppe, J. Pavez, and K. Cranmer, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 5242 (2020).
- [24] A. Buckley, H. Hoeth, H. Lacker, H. Schulz, and J. E. von Seggern, *Eur. Phys. J. C* **65**, 331 (2010).
- [25] S. Mrenna and P. Skands, *Phys. Rev. D* **94**, 074005 (2016).
- [26] J. Bellm, S. Plätzer, P. Richardson, A. Siodmok, and S. Webster, *Phys. Rev. D* **94**, 034028 (2016).
- [27] E. Bothmann, M. Schonherr, and S. Schumann, *Eur. Phys. J. C* **76**, 590 (2016).
- [28] J. Butterworth *et al.*, *J. Phys. G* **43**, 023001 (2016).
- [29] A. Buckley, J. Ferrando, S. Lloyd, K. Nordstrom, B. Page, M. Ruefenacht, M. Schonherr, and G. Watt, *Eur. Phys. J. C* **75**, 132 (2015).
- [30] P. Ilten, M. Williams, and Y. Yang, *J. Instrum.* **12**, P04028 (2017).
- [31] G. Louppe, J. Hermans, and K. Cranmer, in *Proceedings of Machine Learning Research*, edited by K. Chaudhuri and M. Sugiyama (AISTATS, Naha, Japan, 2019), Vol. 89, pp. 1438–1447.
- [32] D. Martschei, M. Feindt, S. Honc, and J. Wagner-Kuhr, *J. Phys. Conf. Ser.* **368**, 012028 (2012).
- [33] A. Rogozhnikov, *J. Phys. Conf. Ser.* **762**, 012036 (2016).
- [34] M. Aaboud *et al.* (ATLAS Collaboration), *Phys. Rev. D* **98**, 092002 (2018).
- [35] A. Andreassen, I. Feige, C. Frye, and M. D. Schwartz, *Euro. Phys. J.* **79**, 102 (2019).
- [36] E. Bothmann and L. Debbio, *J. High Energy Phys.* **01** (2019) 033.
- [37] P. T. Komiske, E. M. Metodiev, and J. Thaler, *J. High Energy Phys.* **01** (2019) 121.
- [38] K. Cranmer, J. Pavez, and G. Louppe, [arXiv:1506.02169](https://arxiv.org/abs/1506.02169).
- [39] P. Baldi, K. Cranmer, T. Faucett, P. Sadowski, and D. Whiteson, *Eur. Phys. J. C* **76**, 235 (2016).
- [40] A. Andreassen, I. Feige, C. Frye, and M. D. Schwartz, *Phys. Rev. Lett.* **123**, 182001 (2019).
- [41] R. T. D’Agnolo and A. Wulzer, *Phys. Rev. D* **99**, 015014 (2019).
- [42] J. H. Collins, K. Howe, and B. Nachman, *Phys. Rev. D* **99**, 014038 (2019).

- [43] J. H. Collins, K. Howe, and B. Nachman, *Phys. Rev. Lett.* **121**, 241803 (2018).
- [44] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, in *Proceedings of the 27th International Conference on Neural Information Processing Systems, NIPS'14* (MIT Press, Cambridge, MA, 2014), Vol. 2, pp. 2672–2680.
- [45] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Póczos, R. Salakhutdinov, and A. J. Smola, [arXiv:1703.06114](https://arxiv.org/abs/1703.06114).
- [46] L. de Oliveira, M. Kagan, L. Mackey, B. Nachman, and A. Schwartzman, *J. High Energy Phys.* **07** (2016) 069.
- [47] P. Baldi, K. Bauer, C. Eng, P. Sadowski, and D. Whiteson, *Phys. Rev. D* **93**, 094034 (2016).
- [48] J. Barnard, E. N. Dawe, M. J. Dolan, and N. Rajcic, *Phys. Rev. D* **95**, 014018 (2017).
- [49] P. T. Komiske, E. M. Metodiev, and M. D. Schwartz, *J. High Energy Phys.* **01** (2017) 110.
- [50] G. Louppe, K. Cho, C. Becot, and K. Cranmer, *J. High Energy Phys.* **01** (2019) 057.
- [51] H. Qu and L. Gouskos, *Phys. Rev. D* **101**, 056019 (2020).
- [52] A. Butter, G. Kasieczka, T. Plehn, and M. Russell, *SciPost Phys.* **5**, 028 (2018).
- [53] D. Guest, J. Collado, P. Baldi, S.-C. Hsu, G. Urban, and D. Whiteson, *Phys. Rev. D* **94**, 112002 (2016).
- [54] ATLAS Collaboration, CERN Report No. ATL-PHYS-PUB-2017-003, 2017.
- [55] CMS Collaboration, CERN Report No. CMS-DP-2017-049, 2017.
- [56] A. M. Sirunyan *et al.* (CMS Collaboration), *J. Instrum.* **13**, P05011 (2018).
- [57] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, *Comput. Phys. Commun.* **191**, 159 (2015).
- [58] T. Sjöstrand, S. Mrenna, and P. Z. Skands, *J. High Energy Phys.* **05** (2006) 026.
- [59] M. Cacciari, G. P. Salam, and G. Soyez, *J. High Energy Phys.* **04** (2008) 063.
- [60] M. Cacciari, G. P. Salam, and G. Soyez, *Eur. Phys. J. C* **72**, 1896 (2012).
- [61] M. Cacciari and G. P. Salam, *Phys. Lett. B* **641**, 57 (2006).
- [62] P. Skands, S. Carrazza, and J. Rojo, *Eur. Phys. J. C* **74**, 3024 (2014).
- [63] F. Chollet, KERAS, <https://github.com/fchollet/keras> (2017).
- [64] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, in *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI '16)*, November 24, 2016 Savannah, GA, USA (2016), Vol. 16, pp. 265–283.
- [65] D. Kingma and J. Ba, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [66] A. Andreassen and B. Nachman, <https://github.com/bnachman/DCTR>.
- [67] A. J. Andreassen and B. Nachman, DCTR: PYTHIA  $e^+e^- \rightarrow Z \rightarrow$  dijets datasets (2019), <https://doi.org/10.5281/zenodo.3518708>.
- [68] J. Thaler and K. Van Tilburg, *J. High Energy Phys.* **02** (2012) 093.
- [69] J. Thaler and K. Van Tilburg, *J. High Energy Phys.* **03** (2011) 015.
- [70] A. J. Larkoski, G. P. Salam, and J. Thaler, *J. High Energy Phys.* **06** (2013) 108.
- [71] A. Pathak, P. Komiske, E. Metodiev, and M. Schwartz, HERWIG7.1 quark and gluon jets (2019).
- [72] P. Komiske, E. Metodiev, and J. Thaler, PYTHIA8 quark and gluon jets for energy flow (2019).
- [73] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevD.101.091901>, for Algorithmic description of DCTR and  $pp$  reweighting examples.
- [74] A. Badea, A. Baty, P. Chang, G. M. Innocenti, M. Maggi, C. McGinn, M. Peters, T.-A. Sheng, J. Thaler, and Y.-J. Lee, *Phys. Rev. Lett.* **123**, 212002 (2019).
- [75] D. Heck and T. Pierog, CORSIKA an air shower simulation program (2019), <https://www.ikp.kit.edu/corsika/>.
- [76] GENIE Collaboration, Event generator & global analysis of neutrino scattering data (2019), <https://hep.ph.liv.ac.uk/~costasa/genie/>.
- [77] NEST Collaboration, Nobel element simulation technique (2019), <https://nest.physics.ucdavis.edu>.
- [78] T. E. Collett, *Astrophys. J.* **811**, 20 (2015).
- [79] V. Springel, *Mon. Not. R. Astron. Soc.* **364**, 1105 (2005).
- [80] G. Biro, G. Papp, G. G. Barnafoldi, D. Nagy, M. Gyulassy, P. Levai, X.-N. Wang, and B.-W. Zhang, *MDPI Proc.* **10**, 4 (2019).