

Optimizers for stabilizing likelihood-free inference

G. Bruno De Luca^{1,2,*} Benjamin Nachman^{3,4,†} Eva Silverstein^{1,2,‡} and Henry Zheng^{1,2,§}

¹Stanford Institute for Theoretical Physics, Stanford, California 94306, USA

²Department of Physics, Stanford University, Stanford, California 94306, USA

³Physics Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA

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



(Received 6 March 2025; accepted 21 October 2025; published 13 November 2025)

A growing number of applications in particle physics and beyond use neural networks as unbinned likelihood ratio estimators applied to real or simulated data. Precision requirements on the inference tasks demand a high-level of stability from these networks, which are affected by the stochastic nature of training. We show how physics concepts can be used to stabilize network training through a physics-inspired optimizer. In particular, the energy conserving descent (ECD) optimization framework uses classical Hamiltonian dynamics on the space of network parameters to reduce the dependence on the initial conditions while also stabilizing the result near the minimum of the loss function. We develop a version of this optimizer known as $\text{ECD}_{q=1}$, which has few free hyperparameters with limited ranges guided by physical reasoning. We apply $\text{ECD}_{q=1}$ to representative likelihood-ratio estimation tasks in particle physics and find on average that it out-performs the widely used Adam optimizer. We expect that ECD will be a useful tool for wide array of data-limited problems, where it is computationally expensive to exhaustively optimize hyperparameters and mitigate fluctuations with ensembling.

DOI: [10.1103/x88j-mv39](https://doi.org/10.1103/x88j-mv39)

I. INTRODUCTION

Neural networks (NN) are playing a growing role in nearly all aspects of particle physics and other areas of science. For most applications, the stochastic nature of NN training is a small nuisance that does not play a major role in determining performance. However, these fluctuations are a major challenge for neural simulation-based inference (NSBI) [1] and related tasks, where NNs are directly used as test statistics and for uncertainty quantification, going beyond usual classification or regression tasks. A typical setting is that the uncertainty from some component of the simulation is assessed by repeating the inference with the component shifted within its uncertainty. It is not possible to probe uncertainties smaller than the inherent stochasticity of the NN training. So even if the component is physically unrelated to the inference target, there will be an apparent uncertainty at a scale set by the NN stochasticity.

This is compounded for every nuisance parameter. Fluctuations at the 1% level can thus quickly become the limiting factor in a statistical analysis, thus reducing the potential gains from NSBI.¹ Measurements employing NSBI have addressed this challenge by ensembling many NNs [3–7]. Other experimental analyses using NNs for likelihood-ratio estimation, such as anomaly detection, also mitigate fluctuations by ensembling [8,9]. Ensembling works in principle, but is computationally challenging and in some cases, prohibitive.

A complementary approach to ensembling is hyperparameter optimization. The result can be stabilized with a well-chosen NN architecture and training regimen. The impact of activation and loss functions in NSBI accuracy and precision have been well studied [10–19], but the protocol for optimizing networks has been less explored. Our goal is to study the impact of the learning procedure on the NN accuracy and precision for particle physics-relevant problems. A key feature of NN optimization is differentiability—(stochastic) gradient descent is very efficient and so most learning methods are based on it. By thinking about NN optimization as a physical system, where the learned parameters are the degrees of freedom, we build on previous work to construct a new optimizer. This optimizer employs energy conserving descent (ECD) [20,21] and makes use of the concept that Hamiltonian dynamics give a

*Contact author: gbdeluca@stanford.edu

†Contact author: bpnachman@lbl.gov

‡Contact author: evas@stanford.edu

§Contact author: henryzheng@stanford.edu

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

¹This can depend on the structure of the data as in [2].

measure on parameter space that can systematically reduce spreading. The results with ECD should not depend strongly on the starting values for the weights, mitigating the need for carefully constructed initializations (possibly, through pretraining [22–26]). Furthermore, the ECD optimizer has few free parameters and physical reasoning provides guidance for their ranges.

This paper is organized as follows. Section II introduces the neural network optimization algorithms. In order to showcase the performance of the ECD optimizer, Sec. III presents numerical results with particle physics-relevant datasets. In particular, we consider two likelihood-ratio estimation problems. First, we consider simulation-to-data reweighting of the thrust distribution from e^+e^- collisions. Next, we consider a multidimensional example built from two simulations of hadronic jets produced in proton-proton collisions. In both cases, we quantify the accuracy and precision of likelihood-ratio estimation based on classifiers trained with the ECD optimizer compared to the widely used Adam optimizer [27]. The paper ends with conclusions and outlook in Sec. IV. Links to data and code are available following Sec. IV.

II. METHODS

The “learning” in a machine learning algorithm is another word for “optimization,” the problem of finding a good choice of parameters starting from a (usually) random initialization; this problem is old, hard, and of general importance, and various algorithms have been proposed, differing in the amount of information used about the target function. Part of the machine learning revolution has been the availability of efficient gradient computation via automatic differentiation, which has led to increased interest in gradient-based optimization methods. The standard approach uses gradient descentlike dynamics, such as “gradient descent with momentum” [28,29] and more refined versions like Adam [27], algorithmically discovered signed optimizers such as Lion [30], or more recent higher-order methods like Ademamix [31]. Physically, they correspond to friction-dominated motion on the loss landscape, which generically converges to local minima, and for which direct analytic control of the optimization dynamics is hard, in particular due to stochastic effects and the absence of conserved quantities. Below, we will use the Adam optimizer as a baseline for our method. This widely used optimizer maintains an exponential moving average of both the gradients and their element-wise squares, using the former as a momentum term and the latter to adaptively tune the step size independently for each parameter dimension.

Recently, a new class of physics-inspired optimization algorithms that do not rely on friction to converge has emerged, providing an immediate constant of motion (the energy) that can be used to analytically predict and steer the dynamics in favorable ways. This class of algorithms uses

both the gradient information *and* the current value of the function—information always available when using automatic differentiation, yet usually discarded by standard gradient-based methods. We will now turn to these.

A. ECD optimization: Tuneable measure on parameter space

The ECD framework for optimization [20,21] and the related sampler [32] leverages energy-conserving Hamiltonian dynamics to derive algorithms with several new theoretical handles on their dynamics and performance. This can lead to qualitative and quantitative improvements in practice. In this work, we develop a particular example of this optimization scheme and design it to parametrically improve upon a key data analysis problem of neural network stability from random initializations.

In general, these algorithms arise as a discretized form of energy-conserving classical Hamiltonian dynamics. The update rules are governed by a time-independent Hamiltonian $H(\mathbf{\Pi}, \mathbf{\Theta})$ with $\mathbf{\Theta}$ a d -dimensional vector of configuration variables (e.g., weights and biases of a neural network) and $\mathbf{\Pi}$ their conjugate momenta, and obtained by discretizing the equations of motion

$$\dot{\mathbf{\Theta}} = \partial_{\mathbf{\Pi}} H, \quad \dot{\mathbf{\Pi}} = -\partial_{\mathbf{\Theta}} H, \quad (1)$$

which conserve the total energy $H(\mathbf{\Pi}, \mathbf{\Theta}) = E$ throughout the dynamics. For the remainder of this paper, we take the Hamiltonian to depend on $\mathbf{\Pi}$ via its norm squared Π^2 .

This physical formalism gives us a handle on three important features of the dynamics:

- (i) The measure on configuration space,

$$p(\mathbf{\Theta}) = \int d^d \Pi \delta(H(\mathbf{\Pi}, \mathbf{\Theta}) - E) \\ = \frac{\Omega_{d-1} |\Pi|^{d-1}}{|\partial_{\mathbf{\Pi}} H|} = \frac{\Omega_{d-1} |\Pi|^{d-1}}{|d\mathbf{\Theta}/dt|}. \quad (2)$$

This determines the distribution of results after the chaotic mixing time (which we wish to be short). As we will review, we can readily obtain Hamiltonians designed so that this measure concentrates toward the bottom of the objective function. In terms of the objective/loss function $F(\mathbf{\Theta})$, we will in particular consider Hamiltonians with the property that

$$p(\mathbf{\Theta}) = (F(\mathbf{\Theta}) - F_0)^{-d\eta/2}, \quad (3)$$

where $F_0 \in \mathbb{R}$ and $\eta \in (0, \infty)$ are optimizer hyperparameters. Among the choices of η in this range, values for which $d\eta/2$ is an even integer have the property that the measure (3) does not require a branch choice if F falls below F_0 . The concentration in (3) will reduce spreading of results, helping to reduce the uncertainties in applying machine

learning to scientific data. Simple examples appeared in Refs. [20,21], including the case where the Hamiltonian describes a particle propagating in configuration space, with the particle having a mass that goes inversely to the η th power of the objective function. This particle slows down and hence concentrates measure toward the minimum of the objective function. In this work, we will consider an optimization version of a Hamiltonian developed in the sampling context in Refs. [32,33], laid out shortly in Sec. II A 1.

Two other features will aid us in characterizing the dynamics:

- (ii) The speed of propagation: Having stipulated that H is a function of Π^2 and Θ , we can obtain Π^2 as a function of E and Θ from the equation $H = E$. We then obtain an expression for the speed from (1),

$$\dot{\Theta}^2 = 4\Pi^2(\partial_{\Pi^2}H)^2|_{\Pi^2=\Pi^2(E,\Theta)}. \quad (4)$$

This was used in the case of a particular example in Ref. [21] to determine how to set hyperparameters so as to keep up the speed of evolution.

- (iii) The rate of mixing: We do not have an *a priori* calculation of the mixing time for an arbitrary objective function. To encourage short mixing time, we follow Ref. [20] and include an option for energy conserving stochastic bounces as implemented in smaller steps in Refs. [21,32]. It is possible to map the evolution to geodesic motion on a curved manifold, as discussed in Ref. [32] Appendix A, and generalized in the Appendix below, which can provide additional insight into the deviation of trajectories required for chaos. After the mixing time, the system becomes insensitive to the initial conditions, as desired.

As we will see in explicit experiments and demonstrations, this strategy will enable us to pin down a Hamiltonian and hyperparameters such that we can increase η and therefore parametrically improve on the concentration of results (3) while maintaining a short mixing time.

1. A particular example: The $q=1$ Hamiltonian

In the context of sampling, Ref. [32] introduced a particular family of Hamiltonians parametrized by a number q . For precision optimization, we will focus on the $q=1$ Hamiltonian from this family. Different choices of q within this class define distinct kinetic terms, which in turn lead to different dynamics in the parameter space of neural networks. As derived in detail in the Appendix [cf. Eq. (A6)], a key property of the $q=1$ case is that the velocity remains constant in magnitude throughout the evolution. This simplifying feature was previously explored for sampling in Ref. [33].

Given a C^1 real function $\mathcal{L}(\Theta)$, the Hamiltonian for $q=1$ is given by,

$$H(\Theta, \Pi) = |\Pi| - e^{-\frac{\mathcal{L}(\Theta)}{d-1}} + E \quad (5)$$

with dynamics (1),

$$\dot{\Theta}^i = \frac{\Pi^i}{|\Pi|} \equiv u^i, \quad (6)$$

$$\dot{\Pi}^i = \frac{\partial^i \mathcal{L}(\Theta)}{d-1} e^{-\frac{\mathcal{L}(\Theta)}{d-1}}, \quad (7)$$

$$|\Pi| = e^{-\frac{\mathcal{L}(\Theta)}{d-1}}. \quad (8)$$

We can relate $\mathcal{L}$ to the target F as

$$e^{-\mathcal{L}(\Theta)} \equiv (F(\Theta) - F_0)^{-\frac{\eta}{2}}, \quad (9)$$

$$\mathcal{L}(\Theta) \equiv \frac{\eta d}{2} \ln(F(\Theta) - F_0), \quad (10)$$

such that the marginal of the distribution $\delta(H(\Theta, \Pi) - E)$ over $d\Pi$, the momentum, gives us the desired target $p(\Theta)$, the distribution over model parameters,

$$\begin{aligned} p(\Theta) &\propto \int_{\mathbb{R}^d} \delta(H(\Theta, \Pi) - E) |\Pi|^{d-1} d\Pi \\ &= (F(\Theta) - F_0)^{-\frac{\eta d}{2}}. \end{aligned} \quad (11)$$

Our goal is to increase η so as to concentrate the results toward the bottom of the loss function, reducing the variance of results in the process.

In addition to F_0 and η , we have two additional hyperparameters, Δt and ν , corresponding to the learning rate and stepwise bounce angle; these appear in our discrete update rule at the end of this section. In Sec. II B, we will motivate relatively narrow priors on their values (as a function of model dimension).

It is important to note that the distribution (11) becomes exponentially concentrated around the maximum as we increase η , and hence more concentrated around the minimum of the loss F . In fact, we will work in a regime where this exponential effect is reduced to a power law, as follows. Consider the case where the minimal value of F , denoted $F_{\min}$, satisfies $\Delta F \equiv F_{\min} - F_0 > 0$. Then, within a quadratic basin $F(\Theta) \sim F_{\min} + F_2 \Theta^2$ the saddle point of the integral

$$\int d^d \Theta p(\Theta), \quad (12)$$

is given by [21]

$$\Theta_*^2 = |\Theta|_*^2 = \frac{(F_{\min} - F_0)(d-1)}{F_2(1+d(\eta-1))} \approx \frac{\Delta F}{F_2\eta}, \quad (13)$$

when $d \gg 1$ and $\eta > 1$. The distance of the loss to its minimal value $(F - F_{\min})|_* = F_2\Theta_*^2$ reduces like η^{-1} in this regime. This yields a weaker dependence on η than the original exponential, but still decreases with η . This milder regime features less extreme behavior in the trajectories as we will see shortly from the update rules and in the Appendix from a geometrical point of view. In Sec. II B we will motivate some simple default values for most of the hyperparameters in this regime, with the learning rate scanned in a limited way. This will prove effective compared to Adam, with its hyperparameters treated similarly.

The update rules for Θ and $\mathbf{u}$ in the continuum are,

$$\dot{\Theta} = \mathbf{u}, \quad (14)$$

$$\dot{\mathbf{u}} = \frac{-\eta d}{2(d-1)(F - F_0)} (\nabla F - (\mathbf{u} \cdot \nabla F)\mathbf{u}). \quad (15)$$

In the second formula we see a divergence in the update rule as $F \rightarrow F_0$, but in the regime $F_{\min} - F_0 > 0$ (say ~ 1), the behavior is less extreme and the saddle (13) still yields a parametrically decreasing loss as a function of η .

In the discrete case, using a first order integrator, the update rules are as follows,

$$\mathbf{u} \leftarrow \frac{\mathbf{u} + \nu \mathbf{z}}{|\mathbf{u} + \nu \mathbf{z}|} \approx \mathbf{u} + \nu \mathbf{z}, \quad \text{for } \nu \ll 1, \quad (16)$$

$$\mathbf{u} \leftarrow \mathbf{u} + \Delta t \dot{\mathbf{u}}, \quad (17)$$

$$\Theta \leftarrow \Theta + \Delta t \dot{\Theta}. \quad (18)$$

The velocity $\mathbf{u}$ is first updated with a bounce, $\mathbf{u} \rightarrow \mathbf{u} + \nu \mathbf{z}$ where $\mathbf{z}$ is drawn from a d -dimensional normal distribution $\mathbf{z} \sim \mathcal{N}(0, 1)$ before updating with the dynamical equations of motion. The bounce angle ν , when turned on as an option, is introduced to add chaos to the dynamics in order to ensure ergodicity. We will call (16)–(18) the $\text{ECD}_{q=1}$ optimizer.

2. Adam in physical language

It can be instructive to compare the update rules (16)–(18) with the corresponding update rules for Adam [27], which we will use as a baseline. Before doing so, let us start with the simpler example of gradient descent with momentum (GDM) [28,29]. Its update rules can be written as

$$\Pi \leftarrow \beta \Pi - \nabla F(\Theta), \quad (19)$$

$$\Theta \leftarrow \Theta + \alpha \Pi, \quad (20)$$

where $\beta \in (0, 1)$ is called the “momentum” hyperparameter (not to be confused with Π) and α is the learning rate.

Notice that we can rewrite the term proportional to Π on the right-hand side of (19) as $\Pi - f \Delta t \Pi$, with $f \equiv (1 - \beta)(\Delta t)^{-1}$ playing the role of a friction coefficient—i.e., the coefficient of a negative additive term in the update for the momentum, proportional to the momentum itself. Friction is how GDM converges: setting $f = 0$ (i.e. $\beta = 1$) removes the friction and ruins convergence.

Adam improves on this scheme by adaptively tuning the learning rate for each individual direction. Specifically, the Adam update rules can be written as [27]

$$\Pi \leftarrow \frac{\beta_1}{1 - \beta_1^t} \Pi + \frac{1 - \beta_1}{1 - \beta_1^t} \nabla F(\Theta), \quad (21)$$

$$\tau \leftarrow \frac{\beta_2}{1 - \beta_2^t} \tau + \frac{1 - \beta_2}{1 - \beta_2^t} (\nabla F(\Theta))^2, \quad (22)$$

$$\Theta \leftarrow \Theta - \alpha \frac{\Pi}{\sqrt{\tau + \varepsilon}}, \quad (23)$$

where the vector operations $(\nabla F(\Theta))^2$ and $\frac{\Pi}{\sqrt{\tau + \varepsilon}}$ are performed element-wise, and $\varepsilon > 0$ is a regularization coefficient. In addition to the friction term, Adam adaptively changes the learning rate for each component by keeping track of the square of the gradient along that direction. This improves performance at the cost of memory: for a d -dimensional problem, Adam needs to store $3d$ variables, as opposed to $2d$ for GDM and ECD. Along these lines, the recent Adammix optimizer [31] has found further improvements by keeping track of even more independent momentum buffers, again at the expense of memory usage. It would be interesting to see whether the performance of ECD algorithms can be further improved by similarly keeping more momentum buffers to adaptively tune the learning rate.² However, as we will see next, for the class of problems studied in this work, ECD outperforms Adam even with a lower memory usage.

B. Hyperparameter defaults and ranges

The $\text{ECD}_{q=1}$ optimizer has a space of hyperparameters traced out by the learning rate Δt , the loss offset F_0 , the bounce amplitude ν , and the measure concentration parameter η . These parameters have nontrivial relations to each other and to the dimensionality d , range $\Theta_{\max}$ of the configuration space, and step count budget n_{step} . In this section, we will explain these relations and define appropriately scaled versions of them. Finally, we will motivate a simple choice of default values and priors, which will prove to be effective in the experiments in Sec. III.

Before proceeding, we note that two broad regimes of interest have been identified in neural network training [34]. One is the traditional regime where overfitting is a

²We thank Vasudev Shyam for discussions on this point.

limiting factor. This applies to the experiments in the current work, where early stopping is useful. The second is the more recently defined scaling regime, applicable for example to data-dominated large language model pretraining. The parameter count, data quantity, hyperparameters, and relations among them may differ for these two types of problems. We focus on the first regime.

First, for problems where the nominal objective is $F = 0$, in order to work in the regime leading to (13) as described in the previous section, we set

$$F_0|_{\text{default}} = -1. \quad (24)$$

In particular, this mitigates against very strong exponential factors in the evolution on phase space. In situations where the objective $F_{\min}$ is approximately known and different from zero, we can set $-\Delta F|_{\text{default}} = F_0|_{\text{default}} - F_{\min} = -1$. Then, given the formula (13) for the typical set of evolution, we aim to take a step $\Delta\Theta \lesssim \Theta_*$ to resolve that scale. Since the velocity is constant, we have $\Delta\Theta = \Delta t$. Defining a rescaled learning rate

$$\Delta\hat{t} = \Delta t\sqrt{\eta}, \quad (25)$$

we set a small range of learning rates

$$\Delta\hat{t}|_{\text{default}} \in (0.1, 1.0) \Rightarrow \Delta\Theta \in (0.1, 1.0)\Theta_*. \quad (26)$$

Our experiments below will exhibit the utility of both orders of magnitude 0.1, 1.0 for the rescaled learning rate.

As derived in Ref. [32], the bounce angle scales like $\nu\sqrt{d}$; we define a scaled value $\hat{\nu} = \nu\sqrt{d}$. Since chaos is generic, one might expect a typical system not to require these bounces, although experiments on language models that we have done along with previous experiments in [20,32]³ show the utility of bounces to improve the mixing time. In the particle physics problems we study in the rest of this paper, we will find that no bounces are needed, so we can set $\hat{\nu} = 0$. We find the results robust to a range of small values of $\hat{\nu}$, so we recommend

$$\hat{\nu} \in (0, 10^{-6}), \quad (27)$$

on a problem of the scale of the experiments in this work. In the sampling setting [32], the need to accurately sample the typical set of the target distribution motivates a prescription of one full bounce per orbit of that locus. More generally, one may consider a value of $\hat{\nu}$ such that one full bounce occurs over a chosen range of motion in Θ space, which may depend on the dimensionality and geometry of the network. The optimization problem does not require fine sampling of a detailed distribution, so it may be less

sensitive to $\hat{\nu}$, as we have found in the experiments below, but more generally nonzero $\hat{\nu}$ can be a useful option.

The final $\text{ECD}_{q=1}$ hyperparameter is η , which provides our parametric improvement in concentration of results. At the level of the measure formula (3), the concentration improves monotonically with η . Experiments show strong improvement as a function of η , but it breaks down at some point. We can understand and estimate this as follows, to prescribe the maximal useful value of η . In the realization of ECD given in Sec. II A 1, the speed on configuration space is fixed at unity. Consider a problem subject to a given budget of update steps n_{steps} , and with a range in configuration space of order $\Theta_{\max}$, which will be of order the distance from the initial Θ to the typical set at Θ_* for a reasonably direct trajectory. In this situation, the variables are related as

$$n_{\text{steps}}\Delta\hat{t}\Theta_* \sim n_{\text{steps}}\frac{\Delta\hat{t}}{\sqrt{\eta F_2}} \sim \Theta_{\max}. \quad (28)$$

Consider an initialization such that $\Theta_{\max} \sim \sqrt{fan_{\text{in}}}$, with fan_{in} the number of weights that feed into a given neuron, which is roughly speaking the width of the network. We then expect the maximal useful η to be of order $n_{\text{steps}}^2/fan_{\text{in}}$. With a particular network architecture containing layers of different width, one can use this logic to derive a more precise value of η . This goes into the values of η we use as in the discussion around equation (32) below.

In the next section, we will analyze two (sets of) scientific problems, comparing the performance of $\text{ECD}_{q=1}$ with these default hyperparameters and learning rate range to the default values of Adam, allowing it a scan over its learning rate. One can also perform a larger hyperparameter scan for both, but this adds computational cost and the $\text{ECD}_{q=1}$ priors defined in this section will prove effective compared to Adam.

III. NUMERICAL RESULTS

A key application of precision neural networks in particle physics is the reweighting of one dataset to match another. Often, one of the datasets is from a real detector and the other is from a simulated detector. This is the setting we will focus on for the following two examples.

A. Electron-positron collisions

First, we consider a one-dimensional example using e^+e^- collisions at the Z pole from the ALEPH experiment at large electron-positron collider [35]. The corresponding data and simulations (with Pythia 6 [36] and Geant 3 [37]) are the same as prepared by Ref. [38]. Each event is characterized by a set of outgoing particles, each with a four-vector and other information (such as electric charge). Here, we examine the thrust (T) distribution [39], which interpolates between dijets as $T \rightarrow 1$ and uniformly distributed particles as $T \rightarrow 0$. The input thrust distributions are shown in Fig. 1.

³In some situations, this may reflect symmetries in the problem or large number in the problem that weakens chaotic interactions.




FIG. 1. Residual ALEPH likelihood ratios trained with $\text{ECD}_{q=1}$, Adam after hyperparameter tuning, and default Adam hyperparameters compared to the truth ratio estimated from the histogram directly. The MAEs for the three estimates are $1.48(29) \times 10^{-3}$ for $\text{ECD}_{q=1}$, $2.59(27) \times 10^{-3}$ for Adam, and $3.20(87) \times 10^{-3}$ for default (Def.) Adam.

The simulated thrust distribution is reweighted to match the measured thrust distribution by training a neural network classifier to distinguish the two datasets from each other. Repurposing neural network classifiers as reweighters has been extensively used in particle physics for nearly a decade [10], as the starting point for a number of downstream inference tasks. The neural network is composed of three, fully connected layers, with 50, 100, and 50 nodes, respectively. Intermediate layers use the rectified linear unit (ReLU) activation function. Each hidden layer is followed by a dropout layer [40] with dropout probability $p = 5\%$. The output is passed through a sigmoid and the models are trained using the binary cross entropy loss. We did not extensively optimize these hyperparameters. Our models trained using $\text{ECD}_{q=1}$ and Adam have the same setup. For many scientific machine learning applications, the noise resulting from a singular model initialization can be problematic and can even dominate the data noise. Thus, we are often lead to taking an average over an ensemble of networks to reduce the model noise. However, such a procedure is computationally expensive especially if one needs to additionally perform some hyperparameter optimization.

For $\text{ECD}_{q=1}$ we fix $\nu = 0$, $\eta \sim n_{\text{steps}}^2 / fan_{\text{in}}$, and $F_0 = F_{\text{min}} - 1$ and perform a small scan over $\Delta t \in [0.1, 0.5, 1, 2]$. For Adam we sample 100 points of $(\Delta t, \beta_1, \beta_2)$ from the distribution,

$$\begin{aligned} \Delta t &\sim 10^{-U(1.5)}, \\ \beta_1 &\sim 1 - 10^{-U(0.5,3)}, \\ \beta_2 &\sim 1 - 10^{-U(0.5,4)}, \end{aligned} \quad (29)$$




FIG. 2. Corner plot of the normalizing flow dataset from [19]. The distributions of the “Monte Carlo” and “data” datasets match those generated by the trained normalizing flows.

where U denotes the uniform distribution. This use of 100 tries compared to a scan of four values for ECD (having fixed the other hyperparameters to defaults) gives Adam many more chances to succeed. In the experiments in the next section, we will instead scan Adam learning rates similarly to the small scan for ECD . The results of both pipelines demonstrate an advantage for ECD .

Ultimately, a particular advantage of $\text{ECD}_{q=1}$ over Adam is the reliable performance of its default parameters and relatively narrow priors, as described in Sec. II B. With default hyperparameter initialization, the theory behind $\text{ECD}_{q=1}$ predicts that it should reduce the noise and bias of each model initialization compared to Adam.

With finite samples of data and simulation, our goal is to reduce the spread of optimizer results below the level of the statistical uncertainties. We train $N = 10$ classifiers with random initializations and compute the mean absolute error (MAE) to assess the performance. The MAE is averaged over these N classifiers. For each classifier, the MAE is defined as the data-density-weighted, bin-wise⁴ absolute error between the likelihood ratio from the neural network and the histogram ratio, summed over the bins from Fig. 1 with $T \leq 0.3$. The uncertainty on the MAE is estimated from the standard deviation of MAE values across the N classifiers.

⁴We approximate the density on the NN side using the bin center. This could be improved in the future by taking into account the within-bin distribution.

To compare the performance of a given ensemble of classifiers, we introduce the weighted MAE, which is given by,

$$MAE \equiv \frac{1}{MN} \sum_{n=1}^N \sum_{i=1}^M w_i |\tilde{\mathcal{L}}_i^n - \hat{\mathcal{L}}_i| \quad (30)$$

where w_i is the weighting of each thrust bin by the fraction of samples within the bin, $\tilde{\mathcal{L}}_i^n$ is the n th classifier's prediction on the T_i thrust bin, $\hat{\mathcal{L}}_i$ is the ground truth of the T_i thrust bin, M is the number of thrust bins, and N is the number of classifiers in the ensemble. To focus on regions of the data where thrust bins contain sufficient samples, we additionally restrict our computation of the MAE to thrust bins $T \leq 0.3$.

We also compute the mean two sample Kolmogorov-Smirnov statistic for the classifier's prediction and the ground truth,

$$MKS \equiv \frac{1}{N} KS(\tilde{F}^n, F) \quad (31)$$

where KS is the Kolmogorov-Smirnov statistic between the empirical distribution function of the n th classifier's prediction $\tilde{F}^n$ and the ground truth F and $F = CDF(\mathcal{L}p(\theta_1)) = CDF(p(\theta_0))$. To obtain the classifier's empirical distribution function, we approximate it as $\tilde{F}^n \approx CDF(\tilde{\mathcal{L}}_i^n p(\theta_1))$, in order to work with a probability distribution rather than a ratio of probabilities.

The resulting MAE values are $1.48(29) \times 10^{-3}$ for $ECD_{q=1}$ and $2.59(27) \times 10^{-3}$ for Adam. The default $ECD_{q=1}$ trained classifier outperforms both the default Adam classifier and the scanned Adam classifier (with $\Delta t = 10^{-5}$, $\beta_1 = 0.997$, $\beta_2 = 0.998$). The KS values are $4.97(99) \times 10^{-4}$ for $ECD_{q=1}$ and $73(14) \times 10^{-5}$ for the scanned Adam classifier. We note also that the default $ECD_{q=1}$ hyperparameters ($F_0 = -1$, $\nu = 0$, $\eta = 10^6$) with $\Delta \hat{t} = 0.1$ achieved the lowest MAE and KS compared to the other scanned configurations, substantiating the robustness of the theory of our choice of default parametrization. In comparison, the default Adam (with $\Delta t = 10^{-3}$, $\beta_1 = 0.9$, $\beta_2 = 0.999$) achieves an MAE value of $3.20(87) \times 10^{-3}$ and a KS value of $6.38(77) \times 10^{-4}$. We find also that the deviations for $ECD_{q=1}$ are small compared to the statistical uncertainties, which is not true in all bins for Adam as shown in Fig 1. We did not explore the effect of different binning strategies in this work: while we expect these MAE numbers to be slightly dependent on the binning choice, we do not expect the qualitative behavior to be affected.

Our default hyperparameter initialization takes into account the dimensionality of the neural network. To test the robustness of the default initialization to scaling dimensionality, we perform the same tests on two larger

TABLE I. MAE's achieved by $ECD_{q=1}$ and Adam on the Aleph dataset when the dimensionality of the NN is scaled by a factor of 4 widthwise and depthwise.

	$ECD_{q=1}$	Adam
Width extended	1.58(18) $\times 10^{-3}$	$3.23(19) \times 10^{-3}$
Depth extended	1.87(58) $\times 10^{-3}$	$2.72(42) \times 10^{-3}$

NN models, one with an extended depth and one with extended width. In both versions, the model size is scaled by a factor of 4. We find in Table I that classifiers trained with $ECD_{q=1}$ are more robust to width and depth scaling compared to Adam.

B. Example with a higher-dimensional physical parameter space

As a second example, we consider the task of simultaneously reweighting many jet substructure observables from proton-proton collisions [41]. The “data” are based on a simulation with Herwig 7.1.5 [42–44] while the “Monte Carlo” simulation are based on a dataset constructed with Pythia 8.243 [45–47]. Both datasets use the Delphes 3.4.2 fast detector simulation [48]. Using the same setup as in Ref. [19], we consider a subset of six observables: jet mass m , constituent multiplicity M , jet width w , jet mass after Soft Drop grooming [49] $\ln \rho = \ln(m_{SD}^2/p_T^2)$, n -subjettiness ratio [50,51] $\tau_{21} = \tau_2^{(\beta=1)}/\tau_1^{(\beta=1)}$ and transverse momentum p_T . Normalizing flows [52] were trained on the “Monte Carlo” and “data” datasets in order to build surrogate models for which the true event-by-event probability density is known. We build datasets for the classifier by sampling 10^7 events from the normalizing flows with probability $p = 0.5$ of sampling from either the “Monte Carlo” flow or the “data” flow, and the test dataset is generated by sampling 10^5 events.⁵ See Fig. 2.

The classifier NNs are constructed as fully connected networks with three layers of 64, 128, and 64 nodes, respectively. As in the previous section, ReLU activations are used for hidden layers and a dropout of 10% provides regularization. The particular architecture here informs the choice of η as anticipated above in Sec. II B. For each layer, we have an initialization $\Theta_i \sim 1/\sqrt{fan_{in}}$. Each set of weights consists of a total of $fan_{in} \times fan_{out}$ parameters. Let us denote the largest width in the problem as $n = 128$ in the present architecture. In this network we have two sets of weights with dimension $n/2 \times n = 64 \times 128$. We have

$$\Theta_{\max} = \sqrt{\sum_{i=1}^d \Theta_i^2} = \sqrt{\sum_{\alpha l} \frac{1}{n} + \sum_{j\beta} \frac{1}{n/2}} = \sqrt{\frac{3}{2}} n, \quad (32)$$

⁵The trained flows are taken from the code base of Ref. [19], <https://github.com/shahzarizvi/reweighting-schemes>.

TABLE II. List of parametrizations of $f = \phi(z)$ for MLC and BCE loss used to train classifiers. Φ is the cumulative distribution function (CDF) of a normal distribution.

BCE	MLC
$\phi = \sigma(z)$	$\phi = \text{ReLU}(z)$
$\phi = \Phi(z)$	$\phi = z^2$
$\phi = \arctan(z)$	$\phi = e^z$

where in this formula Θ_i^2 refers to the variance of the i th parameter obtained from the initialization, and where $\alpha, \beta = 1, \dots, n/2$ and $I, J = 1, \dots, n$. In choosing η in our experiments, we apply the estimate $\eta \sim n_{\text{steps}}^2 \Delta \hat{t}^2 / \Theta_{\text{max}}^2 \sim 2n_{\text{steps}}^2 / 3n$. This is indicated with a $\sim$ due to the uncertainty in the Hessian scale F_2 in (28), here taken to be of order 1.

The final layer activation will depend on the type of loss and the choice of parametrization. Following Ref. [19], we study a variety of options. For each option, the learning rate Δt was scanned in the values $\Delta \hat{t} \in [0.1, 0.5, 1, 2]$ for $\text{ECD}_{q=1}$ and $\Delta t \in [10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}]$ for Adam. In this regard, we gave Adam more choices since as discussed in Sec. II B we need only scan ECD over the values 0.1, 1. Additionally, we implemented early stopping during our training with a patience of ten epochs on validation loss to reduce overfitting and compute.

For each optimizer, we perform our experiments over six types of classifiers categorized by the loss type and final layer activation type. The two types of loss functionals used to train our classifiers are the maximum likelihood classifier (MLC) loss [53,54] and the binary cross entropy (BCE) loss. The MLC and BCE losses are parametrized as,

$$\mathcal{L}_{\text{MLC}}[f] = \frac{1}{N} \sum_{i=1}^N y_i \ln f(x_i) + (1 - y_i) \ln(1 - f(x_i)), \quad (33)$$

$$\mathcal{L}_{\text{BCE}}[f] = \frac{1}{N} \sum_{i=1}^N y_i \ln f(x_i) + (1 - y_i) \ln(1 - f(x_i)), \quad (34)$$

where $y_i \in \{0, 1\}$ indicating whether the event was drawn from the “Monte Carlo” flow or the “data” flow, respectively, f is the NN, and x_i are the input features. Denoting $f = \phi(z)$ where z is the final layer preactivation, for the BCE loss any function $\phi: \mathbb{R} \rightarrow (0, 1)$ is sufficient, while for the MLC loss any function $\phi: \mathbb{R} \rightarrow (0, \infty)$ is sufficient. We have chosen three different realizations of ϕ for each type of loss as shown in Table II.

We use again the MAE as a metric for the performance of our trained classifiers. The average is computed over 100 initializations and we use the multidimensional, unbinned true likelihood ratio (from the ratio of normalizing flows) event-by-event for the absolute difference.

As shown in Table III, $\text{ECD}_{q=1}$ achieves lower MAE on all parametrizations of MLC and BCE loss, with $\sigma(z)$ and

TABLE III. MAE for various loss parametrizations of MLC and BCE classifiers trained with $\text{ECD}_{q=1}$ and Adam over 100 independent initializations. The uncertainties are computed as the standard deviation of the mean absolute errors.

Loss	Param.	$\text{ECD}_{q=1}$	Def. Adam	$ z $ -score
MLC	$\text{ReLU}(z)$	0.407(5)	0.432(5)	3.54
	z^2	0.412(7)	0.434(5)	2.56
	$\exp(z)$	0.401(4)	0.402(6)	0.14
BCE	$\sigma(z)$	0.391(4)	0.398(5)	1.09
	$\Phi(z)$	0.391(4)	0.397(5)	0.94
	$\arctan(z)$	0.393(5)	0.399(5)	0.85

TABLE IV. MAE and hyperparameters used for $\text{ECD}_{q=1}$ runs reported in Table III. We have fixed $F_0 = -0.3$ for BCE loss as empirically we found $F_{\text{min}} \sim 0.68$ while for MLC loss $F_{\text{min}} \sim 0$. To fix η , we find that $n_{\text{steps}} = \text{epochs} \times \text{batch size} \sim 800$ such that $\eta \sim 2n_{\text{steps}}^2 / 3n \sim 1.6 \times 10^3$ according to Eq. (32).

Loss	Parametrization	MAE	η	ν	F_0	$\Delta \tilde{t}$
MLC	$\text{ReLU}(z)$	0.407(5)	1.6×10^3	0	-1	1
	z^2	0.412(7)	1.6×10^3	0	-1	1
	$\exp(z)$	0.401(4)	1.6×10^3	0	-1	1
BCE	$\sigma(z)$	0.391(4)	1.6×10^3	0	-0.3	1
	$\Phi(z)$	0.391(4)	1.6×10^3	0	-0.3	1
	$\arctan(z)$	0.393(5)	1.6×10^3	0	-0.3	1

$\arctan(z)$ giving the best overall MAE. For ECD, the hyperparameters are the defaults, with the learning rate within its narrow prior given in Sec. II B. Its learning rates here and in the ALEPH case are 1 and 0.1, respectively. For Adam, the best learning rate is its default 10^{-3} here, and was orders of magnitude lower (10^{-5}) for the ALEPH case. While the effect size is not as big as in the previous section, the ReLU and z^2 parametrization demonstrate a statistically significant improvement compared to Adam at the two-sigma level. in many cases. The $\text{ECD}_{q=1}$ advantage is with a narrower range of hyperparameters overall, including the experiments in this section and the last. We also report the hyperparameters for each experiment in Table IV.

IV. CONCLUSIONS AND OUTLOOK

In this paper, we have developed a new, physics-inspired optimizer, $\text{ECD}_{q=1}$ building from [20,21,32,33]. While this optimizer can be applied to any problem with gradients, we have focused on neural networks due to their large number of parameters and the large stochasticity from random initializations. In contrast to other gradient-based methods that rely on the analog of friction to converge to reasonable solutions, the $\text{ECD}_{q=1}$ optimizer conserves the analog of energy. This allows for different handles to tune the

optimization and provides for physical intuition about the values of hyperparameters.

While this new optimizer may have widespread use, we believe it will be particularly useful for particle physics. Machine learning for particle physics is entering a precision regime, where neural networks are being deployed to utilize all available information in complex events, but extreme precision requirements demand high-levels of robustness in network predictions. For two representative examples in collider physics, we showed that the $\text{ECD}_{q=1}$ optimizer performs as well as or (much) better than the widely used Adam optimizer, recognizing that in the cases of the exponential parametrization of the MLC loss and all parametrizations of BCE loss, ECD only achieves slightly better performance. Overall, ECD demonstrates greater robustness to such choices.

ACKNOWLEDGMENTS

We thank Matthew Avaylon, Anthony Badea, Patrick Komiske, Yen-Jie Lee, Eric Metodiev, and Jesse Thaler for extensive conversations about the ALEPH dataset. We also would like to thank Yonatan Kahn for very useful discussions of these topics of mutual interest and Alice Gatti and Vasudev Shyam for very helpful discussions and experiments on ECD optimization. This work is supported by the U.S. Department of Energy (DOE), Office of Science under Contract No. DE-AC02-05CH11231. This research is also supported by a Simons Investigator award and National Science Foundation Grant No. PHY-2310429.

DATA AVAILABILITY

The data that support the findings of this article are openly available [55].

APPENDIX: GEOMETRIC INTERPRETATION AND CURVATURE FORMULA

The Hamiltonian approach introduced in Sec. II A can also be understood in the language of Riemannian geometry, with the benefit of allowing the use of well-developed geometric techniques to understand the dynamics. In particular, curvature controls the spread of geodesics, thus influencing the mixing rate.

Upon choosing a Riemannian metric in the configuration space $\mathcal{M}$, the natural motion to consider is the geodesic motion. For simplicity, we will focus on isotropic conformally flat metrics

$$ds^2 = h(\Theta)^{2/d} \delta_{ij} d\Theta^i d\Theta^j, \quad (\text{A1})$$

which is equivalent to restricting to Hamiltonians depending on Π only through $|\Pi|$. When the geodesic motion is ergodic, trajectories are distributed in configuration space according to the Riemannian volume form $\sqrt{g} = h(\Theta)$, where the function $h(\Theta)$ is assumed to be non-negative.

More generally, for a generic observable $\mathcal{O}$, its expectation value on M agrees with its expectation value along the geodesic trajectory if the geodesic motion is ergodic. In formulas

$$\int_M \sqrt{g} \mathcal{O}(\Theta) d^d \Theta \propto \frac{1}{T} \int_0^T \mathcal{O}(\Theta(t)) dt, \quad (\text{A2})$$

when $\Theta(t)$ is a solution of the system of second order ordinary differential equation (ODEs)

$$h(\Theta) \frac{d}{dt} \left(h(\Theta)^{\frac{2}{d}} \frac{d\Theta^i}{dt} \right) = \frac{1}{d} \frac{\partial h(\Theta)}{\partial \Theta^i}. \quad (\text{A3})$$

Along solutions of this equation the geodesics have tangent vectors with constant unit norm, as measured with the metric (A1), implying $h(\Theta)^{2/d} (\frac{d\Theta}{dt})^2 = 1$. Choosing $h(\Theta) = \mu(\Theta)$ with μ as in Eq. (3) and prescribing a certain first order scheme to integrate the geodesic equation (A3) results in the $\text{ECD}_{q=0}$ optimization algorithm introduced in Ref. [21].

Changing the parametrization of the “time” coordinate along the geodesics by reparametrizing it as $t(\tau)$, the points $\Theta(\tau)$ visited during the evolution are not distributed to $h(\Theta)$ anymore, since the integral on the right hand side of (A2) would pick up a factor of $\frac{dt}{d\tau}$. If we do not correct for this factor, effectively we explore a different distribution $\mu \neq h$. Taking this into account, by this simple reparametrization argument we obtain that points collected along solutions of the system of second order ODEs

$$\mu(\Theta) \frac{d}{d\tau} \left(\mu(\Theta) h(\Theta)^{\frac{2}{d}-1} \frac{d\Theta^i}{d\tau} \right) = \frac{1}{d} \frac{\partial h(\Theta)}{\partial \Theta^i} \quad (\text{A4})$$

are distributed according to $\mu(\Theta)$ for any choice of $h(\Theta)$. The length of the tangent vector now is not constant, and we have $\mu(\Theta)^2 h(\Theta)^{\frac{2}{d}-2} (\frac{d\Theta}{d\tau})^2 = 1$. While $h(\Theta) = \mu(\Theta)$ reproduces (A3) and thus the $\text{ECD}_{q=0}$ optimization algorithm, more general choices are possible, resulting in different optimization algorithms when an integration scheme is prescribed.

A simple class is given by $h(\Theta)$ depending on $\mu(\Theta)$ via a power law. By parametrizing it as $h(\Theta) = \mu(\Theta)^{\frac{d}{d-q}}$, we get

$$\mu(\Theta)^{\frac{2q-d}{q-d}} \frac{d}{d\tau} \left(\mu(\Theta)^{\frac{q-2}{q-d}} \frac{d\Theta^i}{d\tau} \right) = \frac{1}{d-q} \frac{\partial \mu(\Theta)}{\partial \Theta^i}, \quad (\text{A5})$$

where now

$$\mu(\Theta)^{\frac{2(1-q)}{d-q}} \left(\frac{d\Theta}{d\tau} \right)^2 = 1. \quad (\text{A6})$$

Different choices of $q \geq 0$ produce new optimization algorithms ECD_q , upon identifying μ with the objective

F as in (3) and choosing an integration scheme. Connecting to the original physical dynamics, notice from (A5) that $q = 2$ results in simple Newtonian evolution on a potential related to μ ; $q \neq 2$ instead corresponds to noncanonical kinetic terms, with the same q parametrization introduced in Ref. [32] in context of sampling. From (A6) we see that for the particular choice $q = 1$ the speed is constant, as discussed from the Hamiltonian point of view in Sec. II A 1 for the $\text{ECD}_{q=1}$ algorithm.

For general q , the continuum update rules (14) generalize to

$$\begin{aligned}\dot{\Theta} &= \mathbf{u} \\ \dot{\mathbf{u}} &= \frac{-\eta d}{2(d-1)(F-F_0)} ((F-F_0)^{-d\eta\frac{q-1}{d-q}} \nabla F \\ &\quad + (q-2)(\mathbf{u} \cdot \nabla F) \mathbf{u}).\end{aligned}\quad (\text{A7})$$

The geometrical formulation allows one to use geometric quantities to understand the spread of the geodesics, and thus of the optimization trajectories. In particular, the

sectional curvature controls the rate at which different close-by geodesics spread (if negative) or converge (if positive). For the metric associated to general q , the sectional curvature on the plane i, j reads

$$\begin{aligned}K_{ij} &= (F-F_0)^{\frac{\eta d}{d-q}} \frac{d\eta}{2(d-q)} \left(\partial_i^2 \log(F-F_0) \right. \\ &\quad \left. + \partial_j^2 \log(F-F_0) - \frac{\eta d}{2(d-q)} \sum_{k \neq i,j} (\partial_k \log(F-F_0))^2 \right).\end{aligned}\quad (\text{A8})$$

The rate of divergence among close-by geodesics can be then obtained via the Jacobi equation, and is controlled by the curvature. We do not write the full expression here, but we notice that close to the bottom $F \sim F_{\min}$ the same exponential divergences that appear in the evolution equations (A7) and (7) also appear in the curvature. As in the discussion in Sec. II A 1, these are mitigated in the regime $\Delta F \equiv F_{\min} - F_0$ of order 1.

-
- [1] K. Cranmer, J. Brehmer, and G. Louppe, The frontier of simulation-based inference, *Proc. Natl. Acad. Sci. U.S.A.* **117**, 30055 (2020).
 - [2] Y. Kahn, S. Roy, J. Shelton, and V. Riki, Systematic uncertainties and data complexity in normalizing flows, *Neurips Workshop* (2024).
 - [3] V. Andreev *et al.* (H1 Collaboration), Measurement of lepton-jet correlation in deep-inelastic scattering with the H1 detector using machine learning for unfolding, *Phys. Rev. Lett.* **128**, 132002 (2022).
 - [4] V. Andreev *et al.* (H1 Collaboration), Unbinned deep learning jet substructure measurement in high Q2ep collisions at HERA, *Phys. Lett. B* **844**, 138101 (2023).
 - [5] G. Aad *et al.* (ATLAS Collaboration), Simultaneous unbinned differential cross-section measurement of twenty-four Z + jets kinematic observables with the ATLAS detector, *Phys. Rev. Lett.* **133**, 261803 (2024).
 - [6] CMS Collaboration, Measurement of event shapes in minimum bias events from pp collisions at 13 TeV, Technical Reports No. CMS-SMP-23-008; No. CERN-EP-2025-041, CERN, Geneva, 2024.
 - [7] G. Aad *et al.* (ATLAS Collaboration), Measurement of off-shell Higgs boson production in the $H^* \rightarrow ZZ \rightarrow 4\ell$ decay channel using a neural simulation-based inference technique in 13 TeV pp collisions with the ATLAS detector, *Rep. Prog. Phys.* **88**, 057803 (2025).
 - [8] G. Aad *et al.* (ATLAS Collaboration), Dijet resonance search with weak supervision using $\sqrt{s} = 13$ TeV pp collisions in the ATLAS detector, *Phys. Rev. Lett.* **125**, 131801 (2020).
 - [9] V. Chekhovsky *et al.* (CMS Collaboration), Model-agnostic search for dijet resonances with anomalous jet substructure in proton-proton collisions at $\sqrt{s} = 13$ TeV, *Rep. Prog. Phys.* **88**, 067802 (2025).
 - [10] K. Cranmer, J. Pavez, and G. Louppe, Approximating likelihood ratios with calibrated discriminative classifiers, *arXiv:1506.02169*.
 - [11] K. Kong, K. T. Matchev, S. Mrenna, and P. Shyamsundar, New machine learning techniques for simulation-based inference: InferenceStatic nets, kernel score estimation, and kernel likelihood ratio estimation, *arXiv:2210.01680*.
 - [12] J. Brehmer, K. Cranmer, G. Louppe, and J. Pavez, Constraining effective field theories with machine learning, *Phys. Rev. Lett.* **121**, 111801 (2018).
 - [13] J. Brehmer, K. Cranmer, G. Louppe, and J. Pavez, A guide to constraining effective field theories with machine learning, *Phys. Rev. D* **98**, 052004 (2018).
 - [14] R. T. D'Agnolo and A. Wulzer, Learning new physics from a machine, *Phys. Rev. D* **99**, 015014 (2019).
 - [15] B. Nachman and J. Thaler, Learning from many collider events at once, *Phys. Rev. D* **103**, 116013 (2021).
 - [16] M. Stoye, J. Brehmer, G. Louppe, J. Pavez, and K. Cranmer, Likelihood-free inference with an improved cross-entropy estimator, *arXiv:1808.00973*.
 - [17] G. V. Moustakides and K. Basioti, Training neural networks for likelihood/density ratio estimation, *arXiv:1911.00405*.
 - [18] N. Jeffrey and B. D. Wandelt, Evidence networks: Simple losses for fast, amortized, neural Bayesian model comparison, *Mach. Learn. Sci. Tech.* **5**, 015008 (2024).

- [19] S. Rizvi, M. Pettee, and B. Nachman, Learning likelihood ratios with neural network classifiers, *J. High Energy Phys.* **02** (2024) 136.
- [20] G. B. De Luca and E. Silverstein, Born-infeld (BI) for AI: Energy-conserving descent (ECD) for optimization, in *Proceedings of the 39th International Conference on Machine Learning*, edited by K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, Vol. 162 of Proceedings of Machine Learning Research (PMLR, Baltimore, Maryland, USA, 2022), pp. 4918–4936.
- [21] G. B. De Luca, A. Gatti, and E. Silverstein, Improving energy conserving descent for machine learning: Theory and practice, [arXiv:2306.00352](https://arxiv.org/abs/2306.00352).
- [22] M. P. Kuchera, R. Ramanujan, J. Z. Taylor, R. R. Strauss, D. Bazin, J. Bradt, and R. Chen, Machine learning methods for track classification in the AT-TPC, *Nucl. Instrum. Methods Phys. Res., Sect. A* **940**, 156 (2019).
- [23] A. Chappell and L. H. Whitehead, Application of transfer learning to neutrino interaction classification, *Eur. Phys. J. C* **82**, 1099 (2022).
- [24] F. A. Dreyer, R. Grabarczyk, and P. F. Monni, Leveraging universality of jet taggers through transfer learning, *Eur. Phys. J. C* **82**, 564 (2022).
- [25] H. Beauchesne, Z.-E. Chen, and C.-W. Chiang, Improving the performance of weak supervision searches using transfer and meta-learning, *J. High Energy Phys.* **02** (2024) 138.
- [26] V. Mikuni and B. Nachman, Solving key challenges in collider physics with foundation models, *Phys. Rev. D* **111**, L051504 (2025).
- [27] D. P. Kingma, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [28] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, Learning representations by back-propagating errors, *Nature (London)* **323**, 533 (1986).
- [29] N. Qian, On the momentum term in gradient descent learning algorithms, *Neural Netw.* **12** **1**, 145 (1999).
- [30] X. Chen *et al.*, Symbolic discovery of optimization algorithms, *Adv. Neural Inf. Process. Syst.* **36**, 49205 (2023).
- [31] M. Pagliardini, P. Ablin, and D. Grangier, The ademamix optimizer: Better, faster, older, [arXiv:2409.03137](https://arxiv.org/abs/2409.03137).
- [32] J. Robnik, G. B. D. Luca, E. Silverstein, and U. Seljak, Microcanonical Hamiltonian Monte Carlo, *J. Mach. Learn. Res.* **24**, 1 (2023), <https://jmlr.org/papers/v24/22-1450.html>.
- [33] J. Robnik and U. Seljak, Fluctuation without dissipation: Microcanonical Langevin Monte Carlo, [arXiv:2303.18221](https://arxiv.org/abs/2303.18221).
- [34] L. Xiao, Rethinking conventional wisdom in machine learning: From generalization to scaling, [arXiv:2409.15156](https://arxiv.org/abs/2409.15156).
- [35] D. Decamp *et al.* (ALEPH Collaboration), ALEPH: A detector for electron-positron annihilations at LEP, *Nucl. Instrum. Methods Phys. Res., Sect. A* **294**, 121 (1990); **303**, 393(E) (1991).
- [36] T. Sjostrand, P. Eden, C. Friberg, L. Lonnblad, G. Miu, S. Mrenna, and E. Norrbin, High-energy physics event generation with Pythia 6.1, *Comput. Phys. Commun.* **135**, 238 (2001).
- [37] R. Brun, F. Bruyant, M. Maire, A. C. McPherson, and P. Zanmarini, GEANT3, Report No. CERN-DD-EE-84-1, <https://inspirehep.net/literature/252007>.
- [38] A. Badea, A. Baty, P. Chang, G. M. Innocenti, M. Maggi, C. McGinn, M. Peters, T.-A. Sheng, J. Thaler, and Y.-J. Lee, Measurements of two-particle correlations in e^+e^- collisions at 91 GeV with ALEPH archived data, *Phys. Rev. Lett.* **123**, 212002 (2019).
- [39] E. Farhi, A QCD test for jets, *Phys. Rev. Lett.* **39**, 1587 (1977).
- [40] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, *J. Mach. Learn. Res.* **15**, 1929 (2014).
- [41] 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).
- [42] M. Bähr *et al.*, Herwig++ physics and manual, *Eur. Phys. J. C* **58**, 639 (2008).
- [43] J. Bellm *et al.*, Herwig 7.0/Herwig++ 3.0 release note, *Eur. Phys. J. C* **76**, 196 (2016).
- [44] J. Bellm *et al.*, Herwig 7.1 Release Note, Technical Reports No. CERN-PH-TH-2017-109; No. CERN-TH-2017-109; No. MAN-HEP-2017-08; No. UWTHPH-2017-10; No. IFJ-PAN-IV-2017-7N; No. IKHEF-2017-026; No. HERWIG-2017-02; No. KA-TP-19-2017; No. MCNET-17-08; No. IPPP-17-40, 2017, <https://herwig.hepforge.org/>.
- [45] T. Sjöstrand, S. Mrenna, and P. Skands, A brief introduction to Pythia 8.1, *Comput. Phys. Commun.* **178**, 852 (2008).
- [46] T. Sjöstrand, S. Mrenna, and P. Skands, Pythia 6.4 physics and manual, *J. High Energy Phys.* **05** (2006) 026.
- [47] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands, An introduction to Pythia 8.2, *Comput. Phys. Commun.* **191**, 159 (2015).
- [48] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, M. Selvaggi (The DELPHES3 Collaboration), Delphes 3: A modular framework for fast simulation of a generic collider experiment, *J. High Energy Phys.* **02** (2014) 057.
- [49] A. J. Larkoski, S. Marzani, G. Soyez, and J. Thaler, Soft drop, *J. High Energy Phys.* **05** (2014) 146.
- [50] J. Thaler and K. Van Tilburg, Identifying boosted objects with N-subjettiness, *J. High Energy Phys.* **03** (2011) 015.
- [51] J. Thaler and K. Van Tilburg, Maximizing boosted top identification by minimizing N-subjettiness, *J. High Energy Phys.* **02** (2012) 093.
- [52] D. Jimenez Rezende and S. Mohamed, Variational inference with normalizing flows, [arXiv:1505.05770](https://arxiv.org/abs/1505.05770).
- [53] R. T. D’Agnolo and A. Wulzer, Learning new physics from a machine, *Phys. Rev. D* **99**, 015014 (2019).
- [54] B. Nachman and J. Thaler, Learning from many collider events at once, *Phys. Rev. D* **103**, 116013 (2021).
- [55] <https://github.com/zhenghenry/LikelihoodRatioStability/>.