

A guide to constraining effective field theories with machine learning

Johann Brehmer,¹ Kyle Cranmer,¹ Gilles Louppe,² and Juan Pavez³¹*New York University, New York 10003, New York, USA*²*University of Liège, 4000 Liège, Belgium*³*Federico Santa María Technical University, Valparaíso 2390123, Chile*

(Received 12 May 2018; published 12 September 2018)

We develop, discuss, and compare several inference techniques to constrain theory parameters in collider experiments. By harnessing the latent-space structure of particle physics processes, we extract extra information from the simulator. This augmented data can be used to train neural networks that precisely estimate the likelihood ratio. The new methods scale well to many observables and high-dimensional parameter spaces, do not require any approximations of the parton shower and detector response, and can be evaluated in microseconds. Using weak-boson-fusion Higgs production as an example process, we compare the performance of several techniques. The best results are found for likelihood ratio estimators trained with extra information about the score, the gradient of the log likelihood function with respect to the theory parameters. The score also provides sufficient statistics that contain all the information needed for inference in the neighborhood of the Standard Model. These methods enable us to put significantly stronger bounds on effective dimension-six operators than the traditional approach based on histograms. They also outperform generic machine learning methods that do not make use of the particle physics structure, demonstrating their potential to substantially improve the new physics reach of the Large Hadron Collider legacy results.

DOI: [10.1103/PhysRevD.98.052004](https://doi.org/10.1103/PhysRevD.98.052004)

I. INTRODUCTION

An important aspect of the legacy of the Large Hadron Collider (LHC) experiments will be precise constraints on indirect signatures of physics beyond the Standard Model (SM), parametrized for instance by the dimension-six operators of the Standard Model effective field theory (SMEFT). The relevant measurements can easily involve tens of different parameters that predict subtle kinematic signatures in the high-dimensional space of the data. Traditional analysis techniques do not scale well to this complex problem, motivating the development of more powerful techniques.

The analysis of high-energy-physics data is based on an impressive suite of simulation tools that model the hard interaction, parton shower, hadronization, and detector response. The community has invested a tremendous amount of effort into developing these tools, yielding the high-fidelity modeling of LHC data needed for precision measurements. Simulators such as PYTHIA [1] and GEANT4 [2] use Monte Carlo techniques to sample the

multitudinous paths through which a particular hard scattering might develop. A single event's path through the simulation can easily involve many millions of random variables. While Monte Carlo techniques can efficiently sample from the distributions implicitly defined by the simulators, it is not feasible to calculate the likelihood for a particular observation because doing so would require integrating over all the possible histories leading to that observation. Clearly it is infeasible to explicitly calculate a numerical integral over this enormous latent space. While this problem is ubiquitous in high energy physics, it is rarely acknowledged explicitly.

Traditionally, particle physicists have approached this problem by restricting the analysis to one or two well-motivated discriminating variables, discarding the information contained in the remaining observables. The probability density for the restricted set of discriminating variables is then estimated with explicit functions or nonparametric approaches such as template histograms, kernel density estimates, or Gaussian processes [3]. These low-dimensional density estimates are constructed and validated using Monte Carlo samples from the simulation. While well-chosen variables may yield precise bounds along individual directions of the parameter space, they often lead to weak constraints in other directions in the parameter space [4]. The sensitivity to multiple parameters can be substantially improved by using the fully differential

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.

cross section. This is the forte of the matrix element method [5–19] and optimal observables [20–22] techniques, which are based on the parton-level structure of a given process. Shower and event deconstruction [23–26] extend this approach to the parton shower. But all these methods still require some level of approximations on the parton shower and either neglect or crudely approximate the detector response. Moreover, even a simplified description of the detector effects requires the numerically expensive evaluation of complicated integrals for each observed event. None of these established approaches scales well to high-dimensional problems with many parameters and observables, such as the SMEFT measurements.

In recent years there has been increased appreciation that several real-world phenomena are better described by simulators that do not admit a tractable likelihood. This appears in fields as diverse as ecology, phylogenetics, epidemiology, cardiac simulators, quantum chemistry, and particle physics. Inference in this setting is often referred to as *likelihood-free inference*, where the inference strategy is restricted to samples generated from the simulator. Implicitly, these techniques aim to estimate the likelihood. A particularly ubiquitous technique is approximate Bayesian computation (ABC) [23–29]. ABC is closely related to the traditional template histogram and kernel density estimation approach used by physicists. More recently, approximate inference techniques based on machine learning and neural networks have been proposed [30–47]. All these techniques have in common that they only take into account simulated samples similar to the actual observables—they do not exploit the structure of the process that generates them.

We develop new *simulation-based inference* techniques that are tailored to the structure of particle physics processes. The key insight behind these methods is that we can extract more information than just samples from the simulations, and that this additional information can be used to efficiently train neural networks that precisely estimate likelihood ratios, the preferred test statistics for LHC measurements. These methods are designed for scalability to both high-dimensional parameter spaces as well as to many observables. They do not require any simplifying assumptions to the underlying physics: they support state-of-the-art event generators with parton shower, reducible and irreducible backgrounds, and full detector simulations. After an upfront training phase, they are very efficient to evaluate. Our tools directly provide an estimator for the likelihood ratio, an intuitive and easily interpretable quantity. Finally, limits derived from these tools with toy experiments have the reassuring property that even if they might not be optimal, they are never wrong, i.e., no points are said to be excluded that should not be excluded at a given confidence level.

In Ref. [48], the companion paper of this publication, we focus on the key ideas and sensitivity enabled by these techniques. Reference [49] presents the methods in a more

abstract setting. Here we describe the actual algorithms in detail, developing several different methods side by side. Given the number of discussed variations, this publication might have the look and feel of a review article and we present it as a guide to the interested practitioner. We focus on the main ideas and differences between the approaches and postpone many technical details until the Appendix.

We evaluate the performance of these different methods on a specific example problem, the measurement of two dimension-six operators in Higgs production in weak boson fusion (WBF) in the four-lepton mode at the LHC. For part of this analysis, we work in an idealized setting in which we can access the true likelihood function, providing us with a ground truth for the comparison of the different analysis methods. After establishing the precision of the likelihood ratio estimation, we turn towards the more physical question of how strongly the two operators can be constrained with the different techniques. We repeat the analysis with a simplified detector response where the ground-truth likelihood is no longer tractable.

We begin by laying out the problem in Sec. II: we summarize the effective field theory idea, list the challenges posed by EFT measurements, translate the problem from a physics perspective into the language of statistics, and discuss its important structural properties. We also set up the example process used throughout the rest of the paper. The description of the analysis methods is split in two parts: in Sec. III we define the different techniques to estimate the likelihood ratio, which includes most of the conceptual work presented here. Section IV then explains how to set limits on the EFT parameters based on these tools. In Sec. V we evaluate the performance of the different tools in our example process. Finally, in Sec. VI we summarize our findings and give recommendations for practitioners. The Appendix describes the different algorithms in more detail and provides additional results. The code and data used for this paper are available online at Ref. [50].

II. THE EFT MEASUREMENT PROBLEM

A. Effective field theory

Effective field theories (EFTs) [51–53] parametrize the effects of physics at an energy scale Λ on observables at smaller energies $E \ll \Lambda$ as a set of local operators. The form of these operators is fixed by the light particles and the symmetry structure of the theory and is entirely independent of the high-energy model. Systematically expanding the Lagrangian in $1/\Lambda$, equivalent to ordering the operators by their canonical dimension, leaves us with a finite set of operators weighted by Wilson coefficients that describe all possible new physics effects up to some order in E/Λ .

In the absence of new particles at the TeV scale, and assuming the symmetry structure of the SM, we can thus describe any new physics signature in LHC processes in terms of a set of higher-dimensional operators [54–59].

In this SM effective field theory (SMEFT), the leading effects beyond the SM come from 59 independent dimension-six operators \mathcal{O}_o with Wilson coefficients f_o ,

$$\mathcal{L}_{\text{D6}} = \mathcal{L}_{\text{SM}} + \sum_o \frac{f_o}{\Lambda^2} \mathcal{O}_o, \quad (1)$$

where the SM corresponds to all $f_o = 0$ and any measurement of a deviation hints at new physics.

The dimension-six Wilson coefficients are perfectly suited as an interface between experimental measurements and theory interpretations. They are largely model independent, can parametrize a wide range of observables, including novel kinematic features, and are theoretically consistent beyond tree level. On the technical side, dimension-six operators are implemented in standard Monte Carlo event generators [60], allowing us to generate predictions for rates and kinematic observables for any combination of Wilson coefficients. Measured values of f_o/Λ^2 can easily be translated to the parameters of specific models through well-established matching procedures [61]. All in all, SMEFT measurements will likely be a key part of the legacy of the LHC experiments [62].

Let us briefly comment on the question of EFT validity. A hierarchy of energy scales $E \ll \Lambda$ is the key assumption behind the EFT construction, but in a bottom-up approach the cutoff scale Λ cannot be known without additional model assumptions. From a measurement $f_o/\Lambda^2 \neq 0$ we can estimate the new physics scale Λ only by assuming a characteristic size of the new physics couplings $\sqrt{f_o}$, and compare it to the energy scale E of the experiment. It has been found that dimension-six operators often capture the dominant effects of new physics even when there is only a moderate scale separation $E \lesssim \Lambda$ [63]. All these concerns are not primarily of interest for the measurement of Wilson coefficients, but rather important for the interpretation of the results in specific UV theories.

B. Physics challenges and traditional methods

EFT measurements at the LHC face three fundamental challenges:

- (1) Individual scattering processes at the LHC are sensitive to several operators and require simultaneous inference over a multidimensional parameter space. While a naive parameter scan works well for one or two dimensions, it becomes prohibitively expensive for more than a few parameters.
- (2) Most operators introduce new coupling structures and predict nontrivial kinematic features. These do not translate one to one to traditional kinematic observables such as transverse momenta, invariant masses or angular correlations. An analysis based on only one kinematic variable typically cannot constrain the full parameter space efficiently. Instead, most of the operator effects only become fully

apparent when multiple such variables including their correlations are analyzed [4,64].

- (3) The likelihood function of the observables is intractable, making this the setting of “likelihood-free inference” or “simulator-based inference.” There are simulators for the high-energy interactions, the parton shower, and detector effects that can generate events samples for any theory parameter values, but they can only be run in the forward mode. Given a set of reconstruction-level observables, it is not possible to evaluate the likelihood of this observation given different theory parameters. The reason is that this likelihood includes the integral over all possible different parton shower histories and particle trajectories through the detector as a normalizing constant, which is infeasible to calculate in realistic situations. We will discuss this property in more detail in the following section.

The last two issues are typically addressed in one of three ways. Most commonly, a small set of discriminating variables (also referred to as summary statistics or engineered features) is handpicked for a given problem. The likelihood in this low-dimensional space is then estimated, for instance, by filling histograms from simulations. While well-chosen variables may lead to good constraints along individual directions of the parameter space, there are typically directions in the parameter space with limited sensitivity [4,64].

The matrix element method [5,6,8–15,17–19] or optimal observables [20–22] go beyond a few specific discriminating variables and use the matrix element for a particular process to estimate the likelihood ratio. While these techniques can be very powerful, they suffer from two serious limitations. The parton shower and detector response are either entirely neglected or approximated through *ad hoc* transfer function. Shower and event deconstruction [23–26] allow for the calculation of likelihood ratios at the level of the parton shower, but still rely on transfer functions to describe the detector response. Finally, even with such a simple description of the shower and detector, the evaluation of the likelihood ratio estimator requires the numerically expensive computation of large integrals for each observed event.

Finally, there is a class of generic methods for likelihood-free inference. For Bayesian inference, the best-known approach is approximate Bayesian computation (ABC) [23–28]. Similar to the histogram approach, it relies on the choice of appropriate low-dimensional summary statistics, which can severely limit the sensitivity of the analysis. Different techniques based on machine learning have been developed recently. In particle physics, the most common example is discriminative classifiers between two discrete hypotheses, such as a signal and a background process. This approach has recently been extended to parameter measurements [30,31]. More generally, many techniques based on the idea of using a classification model, such as neural

networks, for inference in the absence of a tractable likelihood function have been introduced in the machine learning community [32–47]. All of these methods only require samples of events trained according to different parameter points. They do not make use of the structure of the particle physics processes, and thus do not use all available information.

All of these methods come with a price. We develop new techniques that

- (i) are tailored to particle physics measurements and leverage their structural properties,
- (ii) scale well to high-dimensional parameter spaces,
- (iii) can accommodate many observables,
- (iv) capture the information in the fully differential cross sections, including all correlations between observables,
- (v) fully support state-of-the-art simulators with parton showers and full detector simulations, and
- (vi) are very efficient to evaluate after an upfront training phase.

C. Structural properties of EFT measurements

1. Particle-physics structure

One essential step to finding the optimal measurement strategy is identifying the structures and symmetries of the problem. Particle physics processes, in particular those described by effective field theories, typically have two key properties that we can exploit.

First, any high-energy particle physics process factorizes into the parton-level process, which contains the matrix element and in it the entire dependence on the EFT coefficients, and a residual part describing the parton shower and detector effects. In many plausible scenarios of new physics neither the strong interactions in the parton shower nor the electromagnetic and strong interactions in the detector are affected by the parameters of interest. The likelihood function can then be written as

$$p(x|\theta) = \int dz p(x, z|\theta) = \int dz p(x|z) p(z|\theta). \quad (2)$$

Here and in the following x are the actual observables after the shower, detector, and reconstruction; θ are the theory parameters of interest; and z are the parton-level momenta (a subset of the latent variables). Table I provides a dictionary of these and other important symbols that we use.

The first ingredient to this likelihood function is the distribution of parton-level four-momenta,

$$p(z|\theta) = \frac{1}{\sigma(\theta)} \frac{d\sigma(\theta)}{dz}, \quad (3)$$

where $\sigma(\theta)$ and $d\sigma(\theta)/dz$ are the total and differential cross sections, respectively. Crucially, this function is tractable:

TABLE I. Dictionary defining many symbols that appear in this paper. p, r, t symbols denote intractable quantities. The last three rows explain our conventions for indices.

Symbol	Physics meaning	Machine learning abstraction
x	Set of all observables	Features
v	One or two kinematic variables	Low-dimensional summary statistics/engineered feature
$z \equiv z_{\text{parton}}$	Parton-level four-momenta	Latent variables
z_{shower}	Parton shower trajectories	Latent variables
z_{detector}	Detector interactions	Latent variables
$z_{\text{all}} = (z_{\text{parton}}, z_{\text{shower}}, z_{\text{detector}})$	Full simulation history of event	All latent variables
θ	Theory parameters (Wilson coefficients)	Parameters of interest
$\hat{\theta}$	Best fit for theory parameters	Estimator for parameters of interest
$p(x \theta)$	Distributions of observables given theory parameters	Intractable likelihood
$p(z \theta)$	Parton-level distributions from matrix element	Tractable likelihood of latent variables
$p(x z)$	Effect of shower, detector, reconstruction	Intractable density defined through stochastic generative process
$r(x \theta_0, \theta_1)$	Likelihood ratio between hypotheses θ_0, θ_1 , see Eq. (11)	
$\hat{r}(x \theta_0, \theta_1)$	Estimator for likelihood ratio	
$t(x \theta)$	Score, see Eq. (14)	
$\hat{t}(x \theta)$	Estimator for score	
x_e, z_e	Event	Data point
θ_o	Wilson coefficient for one operator	Individual parameter of interest
$\theta_c, w_c(z), p_c(x)$	Morphing basis points, coefficients, densities, see Eq. (6)	

the matrix element and the parton density functions can be evaluated for arbitrary four-momenta z and parameter values θ . In practice this means that matrix-element codes such as MADGRAPH [60] cannot only be run in a forward, generative mode, but also define functions that return the squared matrix element for a given phase-space point z . Unfortunately, there is typically no user-friendly interface to these functions, so evaluating it requires some work.

Second, the conditional density $\mathbf{p}(x|z)$ describes the probabilistic evolution from the parton-level four-momenta to observable particle properties. While this symbol looks innocuous, it represents the full parton shower, the interaction of particles with the detector material, the sensor response and readout, and the reconstruction of observables. Different simulators such as PYTHIA [1], GEANT4 [2], or DELPHES [65] are often used to generate samples $\{x\} \sim \mathbf{p}(x|z)$ for given parton-level momenta z . This sampling involves the Monte Carlo integration over the possible shower histories and detector interactions,

$$\mathbf{p}(x|z) = \int d\mathbf{z}_{\text{detector}} \int d\mathbf{z}_{\text{shower}} \mathbf{p}(x|\mathbf{z}_{\text{detector}}) \times \mathbf{p}(z_{\text{detector}}|z_{\text{shower}}) \mathbf{p}(z_{\text{shower}}|z). \quad (4)$$

This enormous latent space can easily involve many millions of random numbers, and these integrals are clearly intractable, which we denote with the symbol \mathbf{p} . In other words, given a set of reconstruction-level observables x , we cannot calculate the likelihood function $\mathbf{p}(x|z)$ that describes the compatibility of parton-level momenta z with the observation. By extension, we also cannot evaluate $\mathbf{p}(x|\theta)$, the likelihood function of the theory parameters given the observation. The intractable integrals in Eq. (4) are the crux of the EFT measurement problem.

The factorization of Eq. (2) together with the tractability of the parton-level likelihood $p(z|\theta)$ is immensely important. We will refer to the combination of these two properties as *particle-physics structure*. The far-reaching consequences of this structure for EFT measurements will be the topic of Sec. III B. Many (but not all) of the inference strategies we discuss will rely on this condition.

Note that this Markov property holds even with reducible and irreducible backgrounds and when a matching scheme is used to combine different parton-level multiplicities. In these situations there may be different disjoint parts of z space, even with different dimensionalities, for instance when events with n and $n+1$ partons in the final state can lead to the same configuration of observed jets. The integral over z then has to be replaced with a sum over “ z_n spaces” and an integral over each z_n , but the logic remains unchanged.

2. Operator morphing

Effective field theories (and other parametrizations of indirect signatures of new physics) typically contribute a finite number of amplitudes to a given process, each of

which is multiplied by a function of the Wilson coefficients.¹ In this case the likelihood can be written as

$$p(z|\theta) = \sum_{c'} \tilde{w}_{c'}(\theta) f_{c'}(z), \quad (5)$$

where c' labels the different amplitude components, and the functions $f_{c'}(z)$ are not necessarily properly positive definite or normalized.

The simplest example is a process in which one SM amplitude $\mathcal{M}_0(z)$ interferes with one new physics amplitude $\mathcal{M}_{\text{BSM}}(z|\theta) = \theta \mathcal{M}_1(z)$, which scales linearly with a new physics parameter θ . The differential cross section, proportional to the squared matrix element, is then $d\sigma(z) \propto |\mathcal{M}_0(z)|^2 + 2\theta \text{Re} \mathcal{M}_0(z)^\dagger \mathcal{M}_1(z) + \theta^2 |\mathcal{M}_1(z)|^2$. There are three components, representing the SM, interference, and pure Beyond the Standard Model (BSM) terms, each with their own parameter dependence $\tilde{w}_{c'}(z)$ and momentum dependence $f_{c'}(z)$.

We can then pick a number of basis² parameter points θ_c equal to the number of components c' in Eq. (5). They can always be chosen such that the matrix $W_{cc'} = \tilde{w}_{c'}(\theta_c)$ is invertible, which allows us to rewrite (5) as a mixture model

$$p(z|\theta) = \sum_c w_c(\theta) p_c(z) \quad (6)$$

with weights $w_c(\theta) = \sum_{c'} \tilde{w}_{c'}(\theta) W_{cc'}^{-1}$ and (now properly normalized) basis densities $p_c(z) = p(z|\theta_c)$. The weights $w_c(\theta)$ depend on the choice of basis points and are analytically known. This “morphing” procedure therefore allows us to extract the full likelihood function $p(z|\theta)$ from a finite set of evaluations of basis densities $p_c(z)$.

Calculating the full statistical model through morphing requires the likelihood $p(z|\theta)$ to be tractable, which is true for parton-level momenta as argued above. However, the same trick can be applied even when the exact likelihood is intractable, but we can estimate it. For instance, the marginal distribution of any individual kinematic variable $v(x)$ can be reliably estimated through histograms or other density estimation techniques, even when shower and detector effects are taken into account. The morphing procedure then lets us evaluate the full conditional distribution $p(v|\theta)$ based on a finite number of Monte Carlo simulations [66].

Finally, note that Eq. (6) together with Eq. (2) imply

$$\mathbf{p}(x|\theta) = \sum_c w_c(\theta) \mathbf{p}_c(x), \quad (7)$$

¹Exceptions can arise for instance when particle masses or widths depend on the parameters of interest. But in an EFT setting one can expand these quantities in $1/\Lambda$, restoring the factorization.

²Note that the morphing basis points θ_c are unrelated to the choice of an operator basis for the effective field theory.

even if the likelihood function $p(x|\theta)$ and the components $p_c(x)$ are intractable. This will later allow us to impose the morphing structure on likelihood ratio estimators.

Not all EFT amplitudes satisfy the morphing structure in Eq. (5), so we discuss both measurement strategies that rely on and make use of this property as well as more general ones that do not require it to hold.

D. Explicit example

1. Weak-boson-fusion Higgs to four leptons

As an explicit example LHC process we consider Higgs production in weak boson fusion (WBF) with a decay of the Higgs into four leptons,

$$qq \rightarrow qqh \rightarrow qqZZ \rightarrow qq\ell^+\ell^-\ell^+\ell^- \quad (8)$$

with $\ell = e, \mu$, as shown in Fig. 1.

While this process is rare and is likely to only be observed during the high-luminosity run of the LHC, it has a few compelling features that make it a prime candidate to study the efficient extraction of information. First, the two jets from the quarks and in particular the four leptons can be reconstructed quite precisely in the LHC detectors. Even when assuming on-shell conditions and energy-momentum conservation, the final-state momenta span a 16-dimensional phase space, giving rise to many potentially informative observables.

Second, both the production of the Higgs boson in weak boson fusion as well as its decay into four leptons are highly sensitive to the effects of new physics in the Higgs-gauge sector. We parametrize these with dimension-six operators in the SMEFT, following the conventions of the Hagiwara-Ishihara-Szalapski-Zeppenfeld basis [58]. For simplicity, we limit our analysis to the two particularly relevant operators:

$$\begin{aligned} \mathcal{L} = \mathcal{L}_{\text{SM}} + \frac{f_W}{\Lambda^2} \frac{ig}{2} \underbrace{(D^\mu \phi)^\dagger \sigma^a D^\nu \phi W_{\mu\nu}^a}_{\mathcal{O}_W} \\ - \frac{f_{WW}}{\Lambda^2} \frac{g^2}{4} \underbrace{(\phi^\dagger \phi) W_{\mu\nu}^a W^{\mu\nu a}}_{\mathcal{O}_{WW}}. \end{aligned} \quad (9)$$




FIG. 1. Feynman diagram for Higgs production in weak boson fusion in the 4ℓ mode. The red dots show the Higgs-gauge interactions affected by the dimension-six operators of our analysis.

For convenience, we rescale the Wilson coefficients to the dimensionless parameters of interest

$$\theta = \left(\frac{f_W v^2}{\Lambda^2}, \frac{f_{WW} v^2}{\Lambda^2} \right)^T, \quad (10)$$

where $v = 246$ GeV is the electroweak vacuum expectation value. As alluded to above, the validity range of the EFT cannot be determined in a model-independent way. For moderately weakly to moderately strongly coupled underlying new physics models, one would naively expect $|f_o| \lesssim \mathcal{O}(1)$ and the EFT description to be useful in the range $E \approx v \lesssim \Lambda$, or $-1 \lesssim \theta_o \lesssim 1$. This is the parameter range we analyze in this paper.

The interference between the Standard Model amplitudes and the dimension-six operators leads to an intricate relation between the observables and parameters in this process, which has been studied extensively. The precise measurement of the momenta of the four leptons provides access to a range of angular correlations that fully characterize the $h \rightarrow ZZ$ decay [10,67]. These variables are sensitive to the effects of dimension-six operators. But the momentum flow p through the decay vertex is limited by the Higgs mass, and the relative effects of these dimension-six operators are suppressed by a factor p^2/Λ^2 . On the other hand, the Higgs production through two off-shell gauge bosons with potentially high virtuality does not suffer from this suppression. The properties of the two jets recoiling against them are highly sensitive to operator effects in this vertex [68–71].

In Fig. 2 we show example distributions of two particularly informative observables, the transverse momentum of the leading (higher- p_T) jet $p_{T,j1}$, and the azimuthal angle between the two jets, $\Delta\phi_{jj}$. The two quantities are sensitive to different directions in parameter space. Note also that the interference between the different amplitudes can give rise to nontrivial effects. The size of the dimension-six amplitudes grows with momentum transfer, which is strongly correlated with the transverse momentum of the leading jet. If the interference of new-physics amplitudes with the SM diagrams is destructive, this can drive the total amplitude through zero [63]. The jet momentum distribution then dips and rises again with higher energies, as seen in the red curve in the left panel of Fig. 2. Such depleted regions of low probability can lead to very small or large likelihood ratios and potentially pose a challenge to inference methods.

By analyzing the Fisher information in these distributions, it is possible to compare the discrimination power in these two observables to the information contained in the full multivariate distribution or to the information in the total rate. It turns out that the full multivariate distribution $p(z|\theta)$ contains significantly more information than the one-dimensional and two-dimensional marginal distributions of any standard kinematic variables [4]. The total




FIG. 2. Kinematic distributions in our example process for three example parameter points. We assume an idealized detector response to be discussed in Sec. IID 2. Left: Transverse momentum of the leading (higher- p_T) jet, a variable strongly correlated with the momentum transfer in the process. The dip around 350 GeV is a consequence of the amplitude being driven through zero, as discussed in the text. Right: Separation in azimuthal angle between the two jets.

rate is found to carry much less information on the two operators, in particular when systematic uncertainties on the cross sections are taken into account. In this study we therefore only analyze the kinematic distributions for a fixed number of observed events.

2. Sample generation

Already in the sample generation we can make use of the structural properties of the process discussed in Sec. IIC. The amplitude of this process factorizes into a sum of parameter-dependent factors times phase-space-dependent amplitudes, as given in Eq. (5). The effect of the operators \mathcal{O}_W and \mathcal{O}_{WW} on the total Higgs width breaks this decomposition, but this effect is tiny and in practice irrelevant when compared to the experimental resolution. The likelihood function of this process

therefore follows the mixture model in Eq. (6) to good approximation, and the weights $w_c(\theta)$ can be calculated. Since the parton-level likelihood function is tractable, we can reconstruct the entire likelihood function $p(z|\theta)$ based on a finite number of simulator runs, as described in Sec. IIC 2.

To this end, we first generate a parton-level sample $\{z_e\}$ of 5.5×10^6 events with MADGRAPH 5 [60] and its add-on MADMAX [72–74], using the setup described in Ref. [4]. With MADMAX we can evaluate the likelihood $p(z_e|\theta_c)$ for all events z_e and for 15 different basis parameter points θ_c . Calculating the morphing weights $w_c(\theta)$ finally gives us the true parton-level likelihood function $p(z_e|\theta)$ for each generated phase-space point z_e .

In Fig. 3 we show the basis points θ_c and two of the morphing weights $w_c(\theta)$ with their dependence on θ .




FIG. 3. Basis points θ_c and some of the morphing weights $w_c(\theta)$ for our example process. Each panel shows the morphing weight of one of the components c as a function of parameter space. The weights of the remaining 13 components (not shown) follow qualitatively similar patterns. The dots show the position of the basis points θ_c ; the big black dot denotes the basis point corresponding to the morphing weight shown in that panel. Away from the morphing basis points, the morphing weights can easily reach $\mathcal{O}(100)$, with large cancellations between different components.

In some corners of parameter space the weights easily reach up to $|w_c| \lesssim \mathcal{O}(100)$, and there are large cancellations between positive and negative weights. This will pose a challenge for the numerical stability of every inference algorithm that directly uses the morphing structure of the process, as we will discuss later. Other basis choices have led to comparable or larger morphing weights.

Parton shower and detector effects smear the observed particle properties x with respect to the parton-level momenta z and make the likelihood function in Eq. (2) intractable. We develop inference methods that can be applied exactly in this case and that do not require any simplifying assumptions on the shower and detector response. However, in this realistic scenario we cannot evaluate their performance by comparing them to the true likelihood ratio. We therefore test them first on an idealized scenario in which the four-momenta, flavor, and charges of the leptons, and the momenta of the partons, can be measured exactly, $p(x|z) \approx \delta(x - z)$. In this approximation we can evaluate the likelihood $p(x|\theta)$.

After establishing the performance of the various algorithms in this idealized setup, we will analyze the effect of parton shower and detector simulation on the results. We generate an approximate detector-level sample by drawing events from a smearing distribution $p(x|z)$ conditional on the parton-level momenta z . This smearing function is loosely motivated by the performance of the LHC experiments and is defined in Appendix A.

III. LIKELIHOOD RATIO ESTIMATION

According to the Neyman-Pearson lemma, the likelihood ratio

$$r(x|\theta_0, \theta_1) \equiv \frac{p(x|\theta_0)}{p(x|\theta_1)} = \frac{\int dz p(x, z|\theta_0)}{\int dz p(x, z|\theta_1)} \quad (11)$$

is the most powerful test statistic to discriminate between two hypotheses θ_0 and θ_1 . Unfortunately, the integral over the latent space z makes the likelihood function $p(x|\theta)$ as well as the likelihood ratio $r(x|\theta_0, \theta_1)$ intractable. The first and crucial stage of all our EFT measurement strategies is therefore the construction of a *likelihood ratio estimator* $\hat{r}(x|\theta_0, \theta_1)$ that is as close to the true $r(x|\theta_0, \theta_1)$ as possible and thus maximizes the discrimination power between θ_0 and θ_1 .

This estimation problem has several different aspects that we try to disentangle as much as possible. The first choice is the overall structure of the likelihood ratio estimator and its dependence on the theory parameters θ . We discuss this in Sec. III A. Section III B analyzes what information is available and useful to construct (train) the estimators for a given process. Here we will introduce the main ideas that harness the structure of the EFT to increase the information that is used in the training process.

These basic concepts are combined into concrete strategies for the estimation of the likelihood ratio in Sec. III C. After training the estimators, there is an optional additional calibration stage, which we introduce in Sec. III D. Section III E describes the technical implementation of these strategies in terms of neural networks. Finally, we discuss the challenges that the different algorithms face in Sec. III F and introduce diagnostic tools for the uncertainties.

A. Modeling likelihood ratios

1. Likelihood ratios

There are different approaches to the structure of this estimator, in particular to the dependence on the theory parameters θ :

Point by point (PbP).—A common strategy is to scan the parameter space, randomly or in a grid. To reduce the complexity of the scan one can keep the denominator θ_1 fixed, while scanning only θ_0 . Likelihood ratios with other denominators can be extracted trivially as $\hat{r}(x|\theta_0, \theta_2) = \hat{r}(x|\theta_0, \theta_1)/\hat{r}(x|\theta_2, \theta_1)$. Instead of a single reference value θ_1 , we can also use a composite reference hypothesis $p(x|\theta_1) \rightarrow p_{\text{ref}}(x) = \int d\theta_1 \pi(\theta_1) p(x|\theta_1)$ with some prior $\pi(\theta_1)$. This can reduce the regions in feature space with small reference likelihood $p(x|\theta_1)$ and improve the numerical stability.

For each pair (θ_0, θ_1) separately, the likelihood ratio $\hat{r}(x|\theta_0, \theta_1)$ as a function of x is estimated. Only the final results are interpolated between the scanned values of θ_0 .

This approach is particularly simple, but discards all information about the structure and smoothness of the parameter space. For high-dimensional parameter spaces, the parameter scan can become prohibitively expensive. The final interpolation may introduce additional uncertainties.

Agnostic parametrized estimators.—Alternatively we can train one estimator as the full model $\hat{r}(x|\theta_0, \theta_1)$ as a function of both x and the parameter combination (θ_0, θ_1) [30,75]. A modification is again to leave θ_1 at a fixed reference value [or fixed composite reference hypothesis with a prior $\pi(\theta_1)$] and only learn the dependence on x and θ_0 .

This parametrized approach leaves it to the estimator to learn the typically smooth dependence of the likelihood ratio on the physics parameters and does not require any interpolation in the end. There are no assumptions on the form of the dependence of the likelihood on the ratios.

Morphing-aware estimators.—For problems that satisfy the morphing condition of Eq. (6) and thus also Eq. (7), we can impose this structure and the explicit knowledge of the weights $w_c(\theta)$ onto the estimator. Again, one option is to keep the denominator fixed at a reference value (or composite reference hypothesis), leading to




FIG. 4. Score vector as a function of kinematic observables in our example process. Left: First component of the score vector, representing the relative change of the likelihood with respect to small changes in \mathcal{O}_W direction. Right: Second component of the score vector, representing the relative change of the likelihood with respect to small changes in \mathcal{O}_{WW} direction. In both panels, the axes show two important kinematic variables. We find that the score vector is clearly correlated with these two variables.

$$\hat{r}(x|\theta_0, \theta_1) = \sum_c w_c(\theta_0) \hat{r}_c(x), \quad (12)$$

where the basis estimators $\hat{r}_c(x) = \hat{r}(x|\theta_c, \theta_1)$ only depend on x .

Alternatively, we can decompose both the numerator and denominator distributions to find [30,76]

$$\hat{r}(x|\theta_0, \theta_1) = \sum_c \left[\sum_{c'} \frac{w_{c'}(\theta_1)}{w_{c'}(\theta_0)} \hat{r}_{c',c}(x) \right]^{-1} \quad (13)$$

with pairwise estimators $\hat{r}_{c',c}(x) = \hat{r}(x|\theta_{c'}, \theta_c)$.

2. Score and local model

One remarkably powerful quantity is the score, defined as the relative tangent vector

$$\mathbf{t}(x|\theta) = \nabla_{\theta} \log p(x|\theta)|_{\theta_0}. \quad (14)$$

It quantifies the relative change of the likelihood under infinitesimal changes in parameter space and can be seen as a local equivalent of the likelihood ratio.

In a small patch around θ_0 in which we can approximate $\mathbf{t}(x|\theta)$ as independent of θ , Eq. (14) is solved by the *local model*

$$p_{\text{local}}(x|\theta) = \frac{1}{Z(\theta)} p(\mathbf{t}(x|\theta_0)|\theta_0) \exp[\mathbf{t}(x|\theta_0) \cdot (\theta - \theta_0)] \quad (15)$$

with a normalization factor $Z(\theta)$. The local model is in the exponential family. Note that the $\mathbf{t}(x|\theta_0)$ are the sufficient statistics for $p_{\text{local}}(x|\theta)$. This is significant: if we can estimate the vector-valued function $\mathbf{t}(x|\theta_0)$ (with one component per parameter of interest) of the high-dimensional x , we can reduce the dimensionality of our

space dramatically without losing any information, at least in the local model approximation [77].

In fact, ignoring the normalization factors and in the local model the likelihood ratio between θ_0 and θ_1 only depends on the scalar product between the score and $\theta_0 - \theta_1$, which will allow us to take this dimensionality reduction one step further and compress high-dimensional data x into a scalar without loss of power.

In our example process, we are interested in the Wilson coefficients of two dimension-six operators. The score vector therefore has two components. In Fig. 4 we show the relation between these two score components and two informative kinematic variables, the jet p_T and the azimuthal angle between the two jets, $\Delta\phi$. We find that the score vector is very closely related with these two kinematic quantities, but the relation is not quite one to one. Larger energy transfer, measured as larger jet p_T , increases the typical size of the score vector. The \mathcal{O}_{WW} component of the score is particularly sensitive to the angular correlation variable, in agreement with detailed studies of this process [4].

B. Available information and its usefulness

1. General likelihood-free case

All measurement strategies have in common that the estimator $\hat{r}(x|\theta_0, \theta_1)$ is learned from data provided by Monte Carlo simulations (the stochastic generative process). In the most general likelihood-free scenario, we can only generate samples of events $\{x_e\}$ with $x_e \sim p(x|\theta)$ through the simulator, and base an estimator $\hat{r}(x|\theta_0, \theta_1)$ on these generated samples.

One strategy [30] is based on training a classifier with decision function $\hat{s}(x)$ between two equal-sized samples $\{x_e\} \sim p(x|\theta_0)$, labeled $y_e = 0$, and $\{x_e\} \sim p(x|\theta_1)$, labeled $y_e = 1$. The cross-entropy loss functional




FIG. 5. Illustration of some key concepts with a one-dimensional Gaussian toy example. Left: Classifiers trained to distinguish two sets of events generated from different hypotheses (green dots) converge to an optimal decision function $s(x|\theta_0, \theta_1)$ (in red) given in Eq. (17). This lets us extract the likelihood ratio. Right: Regression on the joint likelihood ratios $r(x_e, z_e|\theta_0, \theta_1)$ of the simulated events (green dots) converges to the likelihood ratio $r(x|\theta_0, \theta_1)$ (red line).

$$L[\hat{s}] = -\frac{1}{N} \sum_e (y_e \log \hat{s}(x_e) + (1 - y_e) \log(1 - \hat{s}(x_e))) \quad (16)$$

is minimized by the optimal decision function

$$s(x|\theta_0, \theta_1) = \frac{p(x|\theta_1)}{p(x|\theta_0) + p(x|\theta_1)}. \quad (17)$$

From the decision function $\hat{s}(x)$ of a classifier we can therefore extract an estimator for the likelihood ratio as

$$\hat{r}(x|\theta_0, \theta_1) = \frac{1 - \hat{s}(x|\theta_0, \theta_1)}{\hat{s}(x|\theta_0, \theta_1)}. \quad (18)$$

This idea, sometimes called the likelihood ratio trick, is visualized in the left panel of Fig. 5.

As pointed out in Ref. [30], we can use the weaker assumption of any loss functional that is minimized by a decision function $s(x)$ that is a strictly monotonic function of the likelihood ratio. The underlying reason is that the likelihood ratio is invariant under any transformation $s(x)$ with this property. In practice, the output of any such classifier can be brought closer to the form of Eq. (17) through a calibration procedure, which we will discuss in Sec. III D.

2. Particle-physics structure

As we have argued in Sec. II C, particle physics processes have a specific structure that allow us to extract additional information. Most processes satisfy the factorization of Eq. (2) with a tractable parton-level likelihood $p(z|\theta)$. The generators do not only provide samples $\{x_e\}$, but also the corresponding parton-level momenta (latent variables) $\{z_e\}$ with $(x_e, z_e) \sim p(x, z|\theta_0)$. By evaluating the matrix elements at the generated momenta z_e for different hypotheses θ_0 and θ_1 , we can extract the parton-level

likelihood ratio $p(z_e|\theta_0)/p(z_e|\theta_1)$. Since the distribution of x is conditionally independent of the theory parameters, this is the same as the joint likelihood ratio

$$\begin{aligned} r(x_e, z_{\text{all } e}|\theta_0, \theta_1) &= \frac{p(x_e, z_{\text{detector } e}, z_{\text{shower } e}, z_e|\theta_0)}{p(x_e, z_{\text{detector } e}, z_{\text{shower } e}, z_e|\theta_1)} \\ &= \frac{p(x_e|z_{\text{detector } e})p(z_{\text{detector } e}|z_{\text{shower } e})p(z_{\text{shower } e}|z_e)p(z_e|\theta_0)}{p(x_e|z_{\text{detector } e})p(z_{\text{detector } e}|z_{\text{shower } e})p(z_{\text{shower } e}|z_e)p(z_e|\theta_1)} \\ &= \frac{p(z_e|\theta_0)}{p(z_e|\theta_1)}. \end{aligned} \quad (19)$$

So while we cannot directly evaluate the likelihood ratio at the level of measured observables $r(x|\theta_0, \theta_1)$, we can calculate the likelihood ratio for a generated event conditional on the latent parton-level momenta.

The same is true for the score, i.e., the tangent vectors or relative change of the (log) likelihood under infinitesimal changes of the parameters of interest. While the score $t(x_e|\theta_0) = \nabla_\theta \log p(x|\theta)|_{\theta_0}$ is intractable, we can extract the joint score

$$\begin{aligned} t(x_e, z_{\text{all } e}|\theta_0) &= \nabla_\theta \log p(x_e, z_{\text{detector } e}, z_{\text{shower } e}, z_e|\theta_0) \\ &= \frac{p(x_e|z_{\text{detector } e})p(z_{\text{detector } e}|z_{\text{shower } e})p(z_{\text{shower } e}|z_e)}{p(x_e|z_{\text{detector } e})p(z_{\text{detector } e}|z_{\text{shower } e})p(z_{\text{shower } e}|z_e)} \\ &\quad \times \frac{\nabla_\theta p(z_e|\theta)}{p(z_e|\theta)} \Big|_{\theta_0} \\ &= \frac{\nabla_\theta p(z_e|\theta)}{p(z_e|\theta)} \Big|_{\theta_0} \end{aligned} \quad (20)$$

from the simulator. Again, all intractable parts of the likelihood cancel. We visualize the score in Fig. 6 and




FIG. 6. Illustration of some key concepts with a one-dimensional Gaussian toy example. Left: Probability density functions for different values of θ and the scores $t(x_e, z_e|\theta)$ at generated events (x_e, z_e) . These tangent vectors measure the relative change of the density under infinitesimal changes of θ . Right: Dependence of $\log p(x|\theta)$ on θ for fixed $x = 4$. The arrows again show the (tractable) scores $t(x_e, z_e|\theta)$.

all available information on the generated samples in Fig. 7. It is worth repeating that we are not making any simplifying approximations about the process here, these statements are valid with reducible backgrounds, for state-of-the-art generators including higher-order matrix elements, matching of matrix element and parton shower, and with full detector simulations.

But how does the availability of the *joint* likelihood ratio $r(x, z|\theta)$ and score $t(x, z|\theta)$ (which depend on the latent parton-level momenta z) help us to estimate the likelihood ratio $r(x|\theta)$, which is the one we are interested in?

Consider the L^2 squared loss functional for functions $\hat{g}(x)$ that only depend on x , but which are trying to approximate a function $g(x, z)$,

$$\begin{aligned}
 L[\hat{g}(x)] &= \int dx dz \mathbf{p}(x, z|\theta) |g(x, z) - \hat{g}(x)|^2 \\
 &= \underbrace{\int dx [\hat{g}^2(x) \int dz \mathbf{p}(x, z|\theta) - 2\hat{g}(x) \int dz \mathbf{p}(x, z|\theta) g(x, z) + \int dz \mathbf{p}(x, z|\theta) g^2(x, z)]}_{F(x)}. \quad (21)
 \end{aligned}$$




FIG. 7. Illustration of some key concepts with a one-dimensional Gaussian toy example. Left: Full statistical model $\log r(x|\theta, \theta_1)$ that we are trying to estimate. Right: Available information at the generated events (x_e, z_e) . The dots mark the joint likelihood ratios $\log r(x_e, z_e|\theta_0, \theta_1)$, the arrows the scores $t(x_e, z_e|\theta_0, \theta_1)$.

Via calculus of variations we find that the function $g^*(x)$ that extremizes $L[\hat{g}]$ is given by [49]

$$0 = \frac{\delta F}{\delta \hat{g}} \Big|_{g^*} = 2\hat{g} \underbrace{\int d\mathbf{z} \mathbf{p}(x, z|\theta)}_{=\mathbf{p}(x|\theta)} - 2 \int d\mathbf{z} \mathbf{p}(x, z|\theta) g(x, z), \quad (22)$$

therefore

$$g^*(x) = \frac{1}{\mathbf{p}(x|\theta)} \int d\mathbf{z} \mathbf{p}(x, z|\theta) g(x, z). \quad (23)$$

We can make use of this general property in our problem in two ways. Identifying $g(x_e, z_e)$ with the joint likelihood ratios $r(x_e, z_{\text{all},e}|\theta_0, \theta_1)$ (which we can calculate) and $\theta = \theta_1$, we find

$$g^*(x) = \frac{1}{\mathbf{p}(x|\theta_1)} \int d\mathbf{z} \mathbf{p}(x, z|\theta_1) \frac{\mathbf{p}(x, z|\theta_0)}{\mathbf{p}(x, z|\theta_1)} = r(x|\theta_0, \theta_1). \quad (24)$$

By minimizing the squared loss

$$\begin{aligned} L[\hat{r}(x|\theta_0, \theta_1)] \\ = \frac{1}{N} \sum_{(x_e, z_e) \sim \mathbf{p}(x, z|\theta_1)} |r(x_e, z_{\text{all},e}|\theta_0, \theta_1) - \hat{r}(x_e|\theta_0, \theta_1)|^2 \end{aligned} \quad (25)$$

of a sufficiently expressive function $\hat{r}(x|\theta_0, \theta_1)$, we can therefore regress on the true likelihood ratio [49]. This is illustrated in the right panel of Fig. 5. Note that to get the correct minimum, the events (x_e, z_e) have to be sampled according to the denominator hypothesis θ_1 .

We can also identify $g(x_e, z_e)$ in Eq. (22) with the scores $t(x_e, z_{\text{all},e}|\theta)$, which can also be extracted from the generator. In this case,

$$g^*(x) = \frac{1}{\mathbf{p}(x|\theta)} \int d\mathbf{z} \nabla_{\theta} \mathbf{p}(x, z|\theta) = \mathbf{t}(x|\theta). \quad (26)$$

Thus minimizing

$$L[\hat{t}(x|\theta)] = \frac{1}{N} \sum_{(x_e, z_e) \sim \mathbf{p}(x, z|\theta)} |t(x_e, z_{\text{all},e}|\theta) - \hat{t}(x_e|\theta)|^2 \quad (27)$$

TABLE II. Availability of different quantities from the generative process in the most general likelihood-free setup vs in the particle-physics scenario with the structure given in Eq. (2). Asterisks (*) denote quantities that are not immediately available, but can be regressed from the corresponding joint quantity, as shown in Sec. III B.

Quantity		General likelihood-free	Particle physics
Samples	$\{x_e\}$	✓	✓
Likelihood	$\mathbf{p}(x_e \theta)$		
Likelihood ratio	$r(x_e \theta_0, \theta_1)$		*
Score	$\mathbf{t}(x_e \theta)$		*
Latent state	$\{x_e, z_e\}$		✓
Joint likelihood	$\mathbf{p}(x_e, z_{\text{all},e} \theta)$		
Joint likelihood ratio	$r(x_e, z_{\text{all},e} \theta_0, \theta_1)$		✓
Joint score	$\mathbf{t}(x_e, z_{\text{all},e} \theta)$		✓

of a sufficiently expressive function $\hat{t}(x|\theta)$ allows us to regress on the score $\mathbf{t}(x|\theta)$ [49].³ Now the (x_e, z_e) have to be sampled according to θ . We summarize the availability of the (joint) likelihood, likelihood ratio, and score in the most general likelihood-free setup and in particle physics processes in Table II.

This is one of our key results and opens the door for powerful new inference methods. Particle physics processes involve the highly complex effects of parton shower, detector, and reconstruction, modeled by a generative process with a huge latent space and an intractable likelihood. Still, the specific structure of this class of processes allows us to calculate how much more or less likely a generated event becomes when we move in the parameter space of the theory. We have shown that by regressing on the joint likelihood ratios or scores extracted in this way, we can recover the actual likelihood ratio or score as a function of the observables.

C. Strategies

Let us now combine the estimator structure discussed in Sec. III A with the different quantities available during training discussed in Sec. III B and define our strategies to estimate the likelihood ratio. Here we restrict ourselves to an overview over the main ideas of the different approaches. A more detailed explanation and technical details can be found in Appendix B.

³A similar loss function (with a nonstandard use of the term “score”) was used in Ref. [78], though the derivative is taken with respect to x and, critically, the model did not involve marginalization over the latent variable z .

1. General likelihood-free case

Some approaches are designed for the most general likelihood-free scenario and only require the samples $\{x_e\}$ from the generator:

Histograms of observables.—The traditional approach to kinematic analyses relies on one or two kinematic variables $v(x)$, manually chosen for a given process and set of parameters. Densities $\hat{p}(v(x)|\theta)$ are estimated by filling histograms with generated samples, leading to the likelihood ratio

$$\hat{r}(x|\theta_0, \theta_1) = \frac{\hat{p}(v(x)|\theta_0)}{\hat{p}(v(x)|\theta_1)}. \quad (28)$$

We use this algorithm point by point in θ_0 , but a morphing-based setup is also possible (see Sec. II C 2). We discuss the histogram approach in more detail in Appendix B 1.

Approximate frequentist computation (AFC).—Approximate Bayesian computation (ABC) is currently the most widely used method for likelihood-free inference in a Bayesian setup. It allows to sample parameters from the intractable posterior, $\theta \sim p(\theta|x) = p(x|\theta)p(\theta)/p(x)$. Essentially, ABC relies on the approximation of the likelihood function through a rejection probability

$$p_{\text{rejection}}(x|\theta) = K_\epsilon(v(x), v(x_e)), \quad (29)$$

with $x_e \sim p(x|\theta)$, a kernel K_ϵ that depends on a bandwidth ϵ , and a sufficiently low-dimensional summary statistics $v(x)$.

Inference in particle physics is usually performed in a frequentist setup, so this sampling mechanism is not immediately useful. But we can define a frequentist analogue, which we call “approximate frequentist computation” (AFC). In analogy to the rejection probability in Eq. (29), we can define a kernel density estimate for the likelihood function as

$$\hat{p}(x|\theta) = \frac{1}{N} \sum_e K_\epsilon(v(x), v(x_e)). \quad (30)$$

The corresponding likelihood ratio estimator is

$$\hat{r}(x|\theta_0, \theta_1) = \frac{\hat{p}(x|\theta_0)}{\hat{p}(x|\theta_1)}. \quad (31)$$

We use this approach point by point in θ_0 with a fixed reference θ_1 . As summary statistics, we use subsets of kinematic variables, similar to the histogram approach. We give more details in Appendix B 2.

Calibrated classifiers (CARL⁴).—As discussed in Sec. III B 1, the decision function $\hat{s}(x|\theta_0, \theta_1)$ of a classifier trained to discriminate between samples generated according to θ_0 from θ_1 can be turned into an estimator for the likelihood ratio

$$\hat{r}(x|\theta_0, \theta_1) = \frac{1 - \hat{s}(x|\theta_0, \theta_1)}{\hat{s}(x|\theta_0, \theta_1)}. \quad (32)$$

This is illustrated in the left panel of Fig. 5.

If the classifier does not learn the optimal decision function of Eq. (17), but any monotonic function of the likelihood ratio, a calibration procedure can improve the performance significantly. We will discuss this in Sec. III D below.

We implement this strategy point by point in θ_0 , as an agnostic parametrized classifier $\hat{r}(x|\theta_0, \theta_1)$ that learns the dependence on both x and θ_0 , as well as a morphing-aware parametrized classifier. More details are given in Appendix B 3.

Neural conditional density estimators (NDE).—Several other methods for conditional density estimation have been proposed, often based on neural networks [32–38]. One particularly interesting class of methods for density estimation is based on the idea of expressing the target density as a sequence of invertible transformations applied to a simple initial density, such as a Gaussian [39–42, 47]. The density in the target space is then given by the Jacobian determinant of the transformation and the base density. A closely related and successful alternative are neural autoregressive models [43–46], which factorize the target density as a sequence of simpler conditional densities. Both classes of estimators are trained by maximizing the log likelihood.

We leave a detailed discussion of these techniques for particle physics problems as well as an implementation in our example process for future work.

2. Particle-physics structure

As we argued in Sec. III B, particle physics simulations let us extract the joint likelihood ratio $r(x_e, z_e|\theta_0, \theta_1)$ and the joint score $t(x_e, z_e|\theta_0, \theta_1)$, giving rise to strategies tailored to this class of problems:

Ratio regression (ROLR⁵).—We can directly regress the likelihood ratio $\hat{r}(x|\theta_0, \theta_1)$. As shown in the previous section, the squared error loss between a function $\hat{r}(x_e|\theta_0, \theta_1)$ and the available joint likelihood ratio $r(x_e, z_e|\theta_0, \theta_1)$ is minimized by the likelihood ratio $r(x|\theta_0, \theta_1)$, provided that the samples (x_e, z_e) are drawn according to θ_1 . Conversely, the squared error of $1/r(x_e, z_e|\theta_0, \theta_1)$ with $(x_e, z_e) \sim p(x, z|\theta_0)$ is also

⁴Calibrated ratios of likelihoods

⁵Regression on likelihood ratio

minimized by the likelihood ratio. We can combine these two terms into a combined loss function

$$L[\hat{r}(x|\theta_0, \theta_1)] = \frac{1}{N} \sum_{(x_e, z_e, y_e)} \left(y_e |r(x_e, z_e|\theta_0, \theta_1) - \hat{r}(x|\theta_0, \theta_1)|^2 + (1 - y_e) \left| \frac{1}{r(x_e, z_e|\theta_0, \theta_1)} - \frac{1}{\hat{r}(x|\theta_0, \theta_1)} \right|^2 \right) \quad (33)$$

with $y_e=0$ for events generated according to $(x_e, z_e) \sim \mathbf{p}(x, z|\theta_0)$ and $y_e = 1$ for $(x_e, z_e) \sim \mathbf{p}(x, z|\theta_1)$. The factors of y_e and $(1 - y_e)$ ensure the correct sampling for each part of the loss functional. We illustrate this approach in the right panel of Fig. 5.

This strategy is again implemented point by point in θ_0 , in an agnostic parametrized setup, as well as in a morphing-aware parametrized setup. We describe it in more detail in Appendix B 4.

CARL + score regression (CASCAL⁶).—The parametrized CARL strategy outlined above learns a classifier decision function $\hat{s}(x|\theta_0, \theta_1)$ as a function of θ_0 . If the classifier is realized with a differentiable architecture such as a neural network, we can calculate the gradient of this function and of the corresponding estimator for the likelihood ratio $\hat{r}(x|\theta_0, \theta_1)$ with respect to θ_0 and derive the estimated score

$$\begin{aligned} \hat{t}(x|\theta_0) &= \nabla_{\theta} \log \hat{r}(x|\theta, \theta_1)|_{\theta_0} \\ &= \nabla_{\theta} \log \frac{1 - \hat{s}(x|\theta_0, \theta_1)}{\hat{s}(x|\theta_0, \theta_1)} \Big|_{\theta_0}. \end{aligned} \quad (34)$$

If the estimator is perfect, we expect this estimated score to minimize the squared error with respect to the joint score data available from the simulator, following the arguments in Sec. III B.

We can turn this argument around and use the available score data during the training. Instead of training the classifier just by minimizing the cross entropy, we can instead simultaneously minimize the squared error on this derived score with respect to the true joint score $t(x, z|\theta_0, \theta_1)$. The combined loss function is given by

$$L[\hat{s}] = \frac{1}{N} \sum_e [y_e \log \hat{s}(x_e) + (1 - y_e) \log(1 - \hat{s}(x_e)) + \alpha(1 - y_e) |t(x_e, z_e|\theta_0) - \hat{t}(x_e|\theta_0)|^2] \quad (35)$$

with $\hat{t}(x|\theta_0)$ defined in Eq. (34) and a hyperparameter α that weights the two pieces of the loss function relative to each other. Again, $y_e = 0$ for events

generated according to $(x_e, z_e) \sim \mathbf{p}(x, z|\theta_0)$ and $y_e = 1$ for $(x_e, z_e) \sim \mathbf{p}(x, z|\theta_1)$, and the factors of y_e and $(1 - y_e)$ ensure the correct sampling for each part of the loss functional.

This strategy relies on the parametrized modeling of the likelihood ratio. We implement both an agnostic version as well as a morphing-aware model. See Appendix B 5 for more details.

Neural conditional density estimators + score (SCAN-DAL⁷).—In the same spirit as the CASCAL method, neural density estimators such as autoregressive flows can be augmented with score information. We have started to explore this class of algorithms in Ref. [49], but leave a detailed study and the application to particle physics for future work.

Ratio + score regression (RASCAL⁸).—The same trick works for the parametrized ROLR approach. If the regressor is implemented as a differentiable architecture such as a neural network, we can calculate the gradient of the parametrized estimator $\hat{r}(x|\theta_0, \theta_1)$ with respect to θ_0 and calculate the score

$$\hat{t}(x|\theta_0) = \nabla_{\theta} \log \hat{r}(x|\theta, \theta_1)|_{\theta_0}. \quad (36)$$

Instead of training just on the squared likelihood ratio error, we can minimize the combined loss

$$\begin{aligned} L[\hat{r}(x|\theta_0, \theta_1)] &= \frac{1}{N} \sum_{(x_e, z_e, y_e)} \left[y_e |r(x_e, z_e|\theta_0, \theta_1) - \hat{r}(x_e|\theta_0, \theta_1)|^2 + (1 - y_e) \left| \frac{1}{r(x_e, z_e|\theta_0, \theta_1)} - \frac{1}{\hat{r}(x_e|\theta_0, \theta_1)} \right|^2 + \alpha(1 - y_e) |t(x_e, z_e|\theta_0) - \hat{t}(x_e|\theta_0)|^2 \right] \end{aligned} \quad (37)$$

with $\hat{t}(x|\theta_0)$ defined in Eq. (36) and a hyperparameter α . The likelihood ratios and scores again provide complementary information as shown in Fig. 7.

Once more we experiment with both an agnostic parametrized model as well as a morphing-aware version.

This technique uses all the available data from the simulator that we discussed in Sec. III B to train an estimator of particularly high fidelity. It is essentially a machine-learning version of the matrix element method. It replaces computationally expensive numerical integrals with an upfront regression phase, after which the likelihood ratio can be evaluated very efficiently. Instead of manually specifying simplified smearing functions, the effect of parton shower and

⁶CARL and score approximate likelihood ratio

⁷Score and neural density approximate likelihood

⁸Ratio and score approximate likelihood ratio

detector is learned from full simulations. For more details on RASCAL, see Appendix B 6.

Local score regression and density estimation (SALLY⁹).—In the local model approximation discussed in Sec. III A 2, the score evaluated at some reference point θ_{score} is the sufficient statistics, carrying all the information on θ . A precisely estimated score vector (with one component per parameter of interest) is therefore the ideal summary statistics, at least in the neighborhood of the Standard Model or any other reference parameter point.

In the last section we argued that we can extract the joint score $t(x_e, z_e | \theta_{\text{score}})$ from the simulator. We showed that the squared error between a function $\hat{t}(x | \theta_{\text{score}})$ and the joint score is minimized by the intractable score $t(x | \theta_{\text{score}})$, as long as the events are sampled as $(x_e, z_e) \sim p(x, z | \theta_{\text{score}})$. We can thus use the augmented data to train an estimator $\hat{t}(x | \theta_{\text{score}})$ for the score at the reference point.

In a second step, we can then estimate the likelihood $\hat{p}(\hat{t}(x | \theta_{\text{score}}) | \theta)$ with histograms, Kernel density estimation (KDE), or any other density estimation technique, yielding the likelihood ratio estimator

$$\hat{r}(x | \theta_0, \theta_1) = \frac{\hat{p}(\hat{t}(x | \theta_{\text{score}}) | \theta_0)}{\hat{p}(\hat{t}(x | \theta_{\text{score}}) | \theta_1)}. \quad (38)$$

This particularly straightforward strategy is a machine-learning analogue of optimal observables that learns the effect of parton shower and detector from data. After an upfront regression phase, the analysis of an event only requires the evaluation of one estimator to draw conclusions about all parameters. See Appendix B 7 for more details.

Local score regression, compression to scalar, and density estimation (SALLINO¹⁰).—The SALLY technique compresses the information in a high-dimensional vector of observables x into a lower-dimensional estimated score vector. But for measurements in high-dimensional parameter spaces, density estimation in the estimated score space might still be computationally expensive. Fortunately, the local model of Eq. (15) motivates an even more dramatic dimensionality reduction to one dimension, independent of the number of parameters: Disregarding the normalization constants, the ratio $r(x | \theta_0, \theta_1)$ only depends on the scalar product between the score and $\theta_0 - \theta_1$.

Given the same score estimator $\hat{t}(x | \theta_{\text{score}})$ developed for the SALLY method, we can define the scalar function

$$\hat{h}(x | \theta_0, \theta_1) \equiv \hat{t}(x | \theta_{\text{SM}}) \cdot (\theta_0 - \theta_1). \quad (39)$$

Assuming a precisely trained score estimator, this scalar encapsulates all information on the likelihood ratio between θ_0 and θ_1 , at least in the local model approximation. The likelihood ratio can then be estimated as

$$\hat{r}(x | \theta_0, \theta_1) = \frac{\hat{p}(\hat{h}(x | \theta_0, \theta_1) | \theta_0)}{\hat{p}(\hat{h}(x | \theta_0, \theta_1) | \theta_1)}, \quad (40)$$

where the $\hat{p}(\hat{h})$ are simple univariate density estimators.

This method allows us to condense any high-dimensional observation x into a scalar function without losing sensitivity, at least in the local model approximation. It thus scales exceptionally well to problems with many theory parameters. We describe SALLINO in more detail in Appendix B 8.

Several other variants are possible, including other combinations of the loss functionals discussed above. We leave this for future work.¹¹

We summarize the different techniques in Table III. With this plethora of well-motivated analysis methods, the remaining key question for this paper is how well they do in practice, and which of them (if any) should be used. This will be the focus of Sec. V. In the next sections, we will first discuss some important additional aspects of these methods.

D. Calibration

While the likelihood ratio estimators described above work well in many cases, their performance can be further improved by an additional calibration step. Calibration takes place *after* the “raw” or uncalibrated estimators $\hat{r}_{\text{raw}}(x | \theta_0, \theta_1)$ have been trained. In general, it defines a function C with the aim that $\hat{r}_{\text{cal}} = C(\hat{r}_{\text{raw}})$ provides a better estimator of the true likelihood ratio than \hat{r}_{raw} . We consider two different approaches to defining this function C , which we call probability calibration and expectation calibration.

1. Probability calibration

Consider the CARL strategy, which trains a classifier with a decision function $s(x)$ as the basis for the likelihood

¹¹Not all initially promising strategies work in practice. We experimented with an alternative strategy based on the morphing structure in Eq. (6). Consider the case in which the training sample consists of a number of subsamples, each generated according to a morphing basis point θ_c . Then the morphing basis sample c used in the event generation is a latent variable that we can use instead of the parton-level momenta z to define joint likelihood ratios and joint scores. Regressing on these quantities should converge to the true likelihood ratio and score, in complete analogy to the discussion in Sec. III B. But the joint ratios and scores defined in this way span a huge range, and the densities of the different basis points are very similar. The convergence is therefore extremely slow and the results based on this method suffer from a huge variance.

⁹Score approximates likelihood locally

¹⁰Score approximates likelihood locally in one dimension

TABLE III. Overview over the discussed measurement strategies. The first three techniques can be applied in the general likelihood-free setup, they only require sets of generated samples $\{x_e\}$. The remaining five methods are tailored to the particle physics structure and require the availability of $r(x_e, z_e|\theta_0, \theta_1)$ or $t(x_e, z_e|\theta_0)$ from the generator, as discussed in Sec. III B. Brackets denote possible variations that we have not implemented for our example process. In the SALLY and SALLINO strategies, “estimator versions” refers to the density estimation step. In the loss function columns, “CE” stands for the cross entropy, “ML” for maximum likelihood, “ratio” for losses of the type $|r(x, z) - \hat{r}(x)|^2$, and “score” for terms such as $|t(x, z) - \hat{t}(x)|^2$.

Strategy	Estimator versions			Loss function				Asymptotically exact
	PbP	Param	Aware	CE	ML	Ratio	Score	
Histograms	✓		(✓)					
AFC	✓		(✓)					
CARL	✓	✓	✓	✓				✓
NDE	(✓)	(✓)	(✓)		✓			✓
ROLR	✓	✓	✓			✓		✓
CASCAL		✓	✓	✓			✓	✓
SCANDAL		(✓)	(✓)		✓		✓	✓
RASCAL		✓	✓			✓	✓	✓
SALLY	✓		(✓)				✓	
SALLINO	✓		(✓)				✓	

ratio estimation. Even if this classifier can separate the two classes of events from θ_0 and θ_1 well, its decision function $\hat{s}(x|\theta_0, \theta_1)$ might not have a direct probabilistic interpretation: it might be any approximately monotonic function of the likelihood ratio rather than the ideal solution given in Eq. (17). In this case, the CARL approach requires calibration, an additional transformation of the raw output $\hat{r}_{\text{raw}} = \hat{r}(x|\theta_0, \theta_1)$ into a calibrated decision function

$$\hat{r}_{\text{cal}} = C(\hat{r}_{\text{raw}}) = \frac{\hat{p}(\hat{r}_{\text{raw}}|\theta_0)}{\hat{p}(\hat{r}_{\text{raw}}|\theta_1)}, \quad (41)$$

where the densities $\hat{p}(\hat{r}_{\text{raw}}|\theta)$ are estimated through a univariate density estimation technique such as histograms. This calibration procedure does not only apply to classifiers, but to any other likelihood ratio estimation strategy.

2. Expectation calibration

Consider some likelihood ratio $r(x|\theta_0, \theta_1)$. The expectation value of this ratio assuming θ_1 to be true is given by

$$\mathbb{E}[r(x|\theta_0, \theta_1)|\theta_1] = \int dx p(x|\theta_1) \frac{p(x|\theta_0)}{p(x|\theta_1)} = 1. \quad (42)$$

A good estimator for the likelihood ratio should reproduce this property. We can numerically approximate this expectation value by evaluating $\hat{r}(x|\theta, \theta_1)$ on a sample $\{x_e\}$ of N events drawn according to θ_1 ,

$$\hat{R}(\theta) = \frac{1}{N} \sum_{x_e \sim \theta_1} \hat{r}(x_e|\theta, \theta_1) \approx 1. \quad (43)$$

If a likelihood ratio estimator $\hat{r}_{\text{raw}}(x|\theta, \theta_1)$ (which might be entirely uncalibrated or already probability calibrated)

does not satisfy this condition, we can calibrate it by rescaling it as

$$\hat{r}_{\text{cal}}(x|\theta, \theta_1) = \frac{\hat{r}_{\text{raw}}(x|\theta, \theta_1)}{\hat{R}_{\text{raw}}(\theta)}. \quad (44)$$

For a perfect estimator with $\hat{r}(x|\theta_0, \theta_1) = r(x|\theta_0, \theta_1)$, we can even calculate the variance of the numeric calculation of the expectation value in Eq. (43). We find

$$\text{var}[\hat{R}(\theta)] = \frac{1}{N} [\mathbb{E}[\hat{r}(x|\theta, \theta_1)|\theta] - 1], \quad (45)$$

where N is the number of events used to calculate the expectation value $R(\theta)$, and the expectation $\mathbb{E}[\hat{r}(x|\theta, \theta_1)|\theta]$ (under the *numerator* hypothesis) can be calculated numerically.

This calibration strategy can easily improve classifiers that are off by some θ -dependent factor. However, a few rare events x_e with large $\hat{r}(x_e|\theta, \theta_1)$ can dominate the expectation value. If these are misestimated, the expectation calibration can actually degrade the performance of the estimator on the bulk of the distribution with smaller $\hat{r}(x|\theta, \theta_1)$.

E. Implementation

1. Neural networks

With the exception of the simple histogram and AFC methods, all strategies rely on a classifier $\hat{s}(x|\theta_0, \theta_1)$, score regressor $\hat{t}(x|\theta_0)$, or ratio regressor $\hat{r}(x|\theta_0, \theta_1)$ that is being learnt from training data. For our explicit example, we implement these functions as fully connected neural networks:

- (i) In the point-by-point setup, the neural networks take the features x as input and models $\log \hat{r}(x|\theta_0, \theta_1)$. For ratio regression, this is exponentiated to yield the final output $\hat{r}(x|\theta_0, \theta_1)$. In the CARL strategy, the network output is transformed to a decision function

$$\hat{s}(x|\theta_0, \theta_1) = \frac{1}{1 + \hat{r}(x|\theta_0, \theta_1)}. \quad (46)$$

- (ii) In the agnostic parametrized setup, the neural networks take both the features x as well as the numerator parameter θ_0 as input and model $\log \hat{r}(x|\theta_0, \theta_1)$. In addition to the same subsequent transformations as in the point-by-point case, taking the gradient of the network output gives the estimator score.
- (iii) In the morphing-aware setup, the estimator takes both x and θ_0 as input. The features x are fed into a number of independent networks, one for each basis component c , which model the basis ratios $\log r_c(x)$. From θ_0 the estimator calculates the component weights $w_c(\theta_0)$ analytically. The components are then combined with Eq. (12). For the CARL approaches, the output is again transformed with Eq. (46), and for the score-based strategies the gradient of the output is calculated.

We visualize these architectures in Fig. 8.

All networks are implemented in shallow, regular, and deep versions with two, three, and five hidden layers of 100 units each and tanh activation functions. They are trained with the ADAM optimizer [79] over 50 epochs with early stopping and learning rate decay. We implement them in KERAS [80] with a TENSORFLOW [81] backend. Experiments modeling s rather than $\log r$, with different activation functions, adding dropout layers, or using other optimizers and learning rate schedules have led to a worse performance.

2. Training samples

Starting from the weighted event sample described in Sec. IID 2, we draw events (x_e, z_e) randomly with probabilities given by the corresponding $p(x_e, z_e|\theta)$. Due to the form of the likelihood $p(x, z|\theta)$ and due to technical limitations of our simulator, individual data points can carry large probabilities $p(x_e, z_e|\theta)$, leading to duplicate events in the training samples. However, we enforce that there is no duplication between training and evaluation samples, so this limitation can only degrade the performance.

For the point-by-point setup, we choose 100 values of θ_0 , five of which are fixed at the SM [$\theta_0 = (0, 0)$] and at the corners of the considered parameter space, with the remaining 95 chosen randomly with a flat prior over $-1 \leq \theta_o \leq 1$. For each of these training points we sample 250 000 events according to θ_0 and 250 000 events according to the reference hypothesis

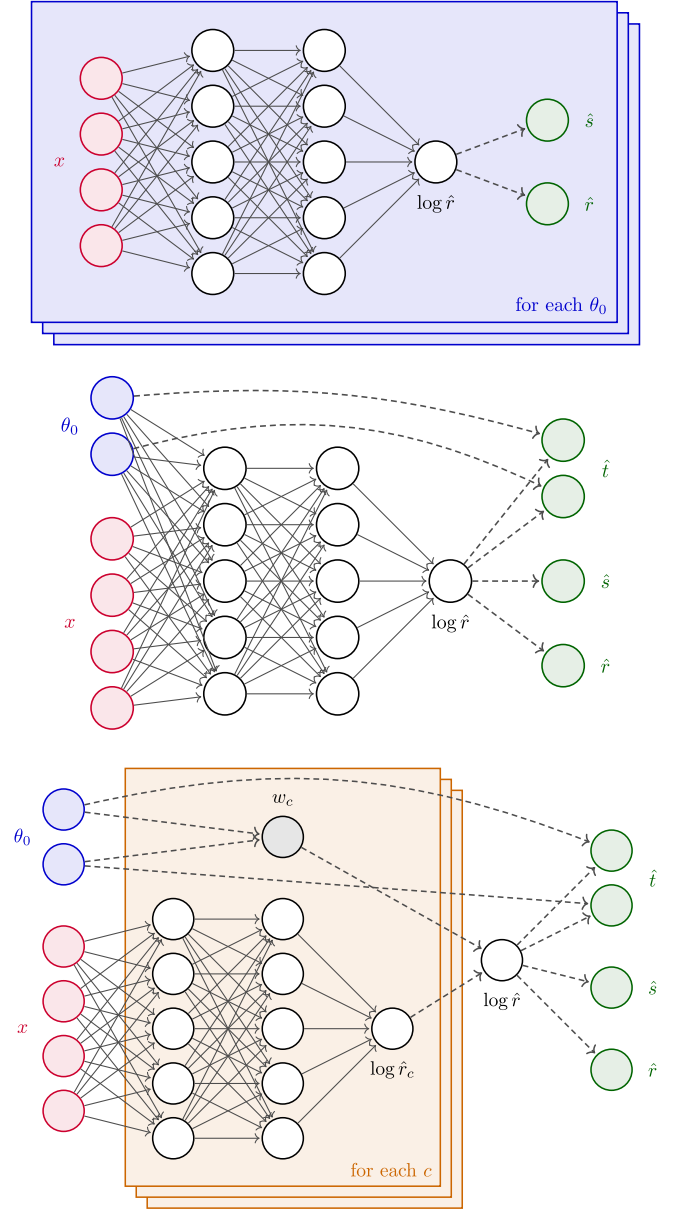


FIG. 8. Schematic neural network architectures for point-by-point (top), agnostic parametrized (middle), and morphing-aware parametrized (bottom) estimators. Solid lines denote dependencies with learnable weights, dashed lines show fixed functional dependencies.

$$\theta_1 = (0.393, 0.492)^T. \quad (47)$$

For the parametrized strategies, we compare three different training samples, each consisting of 10^7 events:

Baseline.—For 1000 values of θ_0 chosen randomly in θ space, we draw 5000 events according to θ_0 and 5000 events according to θ_1 given in Eq. (47).

Random θ .—In this sample, the value of θ_0 is drawn randomly independently for each event. Again we use a flat prior over $\theta_0 \in [-1, 1]^2$.

Morphing basis.—For each of the 15 basis hypotheses θ_i from the morphing procedure, we generate 333 000 events according to $\theta_0 = \theta_i$ and 333 000 according to θ_1 .

Finally, for the local score regression model we use a sample of 10^7 events drawn according to the SM.

Our evaluation sample consists of 50 000 events drawn according to the SM. We evaluate the likelihood ratio for these events for a total of 1016 values of θ_0 , 1000 of which are the same as those used in the baseline training sample. Again we fix θ_1 as in Eq. (47).

Each event is characterized by 42 features:

- (i) the energies, transverse momenta, azimuthal angles, and pseudorapidities of all six particles in the final state;
- (ii) the energies, transverse momenta, azimuthal angles, pseudorapidities, and invariant mass of the four-lepton system as well as the two-lepton systems that reconstruct the two Z bosons; and
- (iii) the invariant mass, separation in pseudorapidity, and separation in azimuthal angle of the dijet system.

The derived variables in the feature set help the neural networks pick up the relevant features faster, though we did not find that their choice affects the performance significantly.

3. Calibration and density estimation

We calibrate the classifiers for our example process with probability calibration as described in Sec. III D 1. We determine the calibration function $C(r)$ with isotonic regression [82], which constrains $C(r)$ to be monotonic. Experiments with other regression techniques based on histograms, kernel density estimation, and logistic regression did not lead to a better performance. We apply the calibration point by point in θ_0 . It is based on an additional event sample that is independent of the training and evaluation data. The same events are used to calibrate each value of θ , with an appropriate reweighting. This strategy to minimize variance is based on the availability of the parton-level likelihood function $p(z|\theta)$.

The techniques based on local score regression require the choice of a reference point to evaluate the score. For the EFT problem, the natural choice is $\theta_{\text{score}} = \theta_{\text{SM}} = (0, 0)^T$. In the SALLY approach, we perform the density estimation based on two-dimensional histograms of the estimated score at the SM, point by point in θ_0 . For the SALLINO technique, we use a one-dimensional histograms of $\hat{h}(x|\theta_{\text{SM}})$, point by point in θ_0 .

F. Challenges and diagnostics

1. Uncertainties

Even the most evolved and robust estimators will have some deviations from the true likelihood ratio, which should be taken into account in an analysis as an additional

modeling uncertainty. Most of the estimators developed above converge to the true likelihood ratio in the limit of infinite training and calibration samples. But with finite statistics, there are different sources of variance that affect some strategies more than others.

Consider the traditional histogram approach. In the point-by-point version, each separate estimator $\hat{r}(x|\theta_0, \theta_1)$ is trained on the small subset of the data generated from a specific value of θ_0 , so the variance from the finite size of the training data, i.e., the statistical uncertainty from the Monte Carlo simulation, is large. At θ_0 values between the training points, there are additional sources of uncertainty from the interpolation. On the other hand, morphing-aware histograms use all of the training data to make predictions at all points, and since the dependence on θ_0 is known, the interpolation is exact. But the large morphing weights $w_c(\theta_0)$ and the cancellations between them mean that even small fluctuations in the individual basis histograms can lead to huge errors on the combined estimator.

Similar patterns hold for the ML-based inference strategies. The point-by-point versions suffer from a larger variance due to small training samples at each point and interpolation uncertainties. The agnostic parametrized models have more statistics available, but have to learn the more complex full statistical model including the dependence on θ_0 . The morphing-aware versions make maximal use of the physics structure of the process and all the training data, but large morphing weights can dramatically increase the errors of the individual component estimators $\log \hat{r}_c(x)$. We demonstrate this for our example process in Fig. 9: in some regions of parameter space, in particular far away from the basis points, the errors on a morphing-aware estimator $\log \hat{r}(x|\theta, \theta_1)$ can easily be 100 times larger than the individual errors on the component estimators $\log \hat{r}_c(x)$. The θ_0 dependence on this error depends on the choice of the basis points; this uncertainty can thus be somewhat mitigated by optimizing the basis points or by combining several different bases.

2. Diagnostics

After discussing the sources of variance, let us now turn towards diagnostic tools that can help quantify the size of estimator errors and to assign a modeling uncertainty for the statistical analysis. These methods are generally closure tests: we can check the likelihood ratio estimators for some expected behavior, and use deviations either to correct the results (as in the calibration procedures described in Sec. III D), to define uncertainty bands, or to discard estimators altogether. We suggest the following tests:

Ensemble variance.—Repeatedly generating new training data (or bootstrapping the same training sample) and training the estimators again gives us an ensemble of predictions $\{\hat{r}_r(x|\theta_0, \theta_1)\}$. We can use the ensemble variance as a measure of uncertainty of the prediction




FIG. 9. Uncertainty $\Delta \log \hat{r}(x_e|\theta, \theta_1)$ of morphing-aware estimators due to uncertainties $\Delta \log \hat{r}_c(x_e|\theta, \theta_1)$ on the individual basis ratios as a function of θ . We fix θ_1 as in Eq. (47) and show one random event x_e , the results for other events are very similar. We assume iid Gaussian uncertainties on the $\log \hat{r}_c(x|\theta, \theta_1)$ and use Gaussian error propagation. The white dots show the position of the basis points θ_c . Small uncertainties in the individual basis estimators $\hat{r}_c(x_e|\theta, \theta_1)$ are significantly increased due to the large morphing weights and can lead to large errors of the combined estimator $\hat{r}(x_e|\theta, \theta_1)$.

that is due to the variance in the training data and random seeds used during the training.

Reference hypothesis variation.—Any estimated likelihood ratio between two hypotheses θ_A, θ_B

$$\hat{r}(x|\theta_A, \theta_B) = \frac{\hat{r}(x|\theta_A, \theta_1)}{\hat{r}(x|\theta_B, \theta_1)} \quad (48)$$

should be independent of the choice of the reference hypothesis θ_1 used in the estimator \hat{r} . Training several independent estimators with different values of θ_1 thus provides another check of the stability of the results [30].

Much like the renormalization and factorization scale variations that are ubiquitous in particle physics calculations, this technique does not have a proper statistical interpretation in terms of a likelihood function. We can still use it to qualitatively indicate the stability of the estimator under this hyperparameter change.

Ratio expectation.—As discussed above, the expectation value of the estimated likelihood ratio assuming the denominator hypothesis should be very close to one. We can numerically calculate this expectation value $\hat{R}(\theta)$, see Eq. (43). In Sec. III D 2 we argued that this expectation value can be used to calibrate the estimators, but that this calibration can actually decrease the performance in certain situations.

If expectation calibration is used, the calibration itself has a nonzero variance from the finite sample size used to calculate the expectation value $\hat{R}(\theta)$. As we

pointed out in Sec. III D 2, we can calculate this source of statistical uncertainty, at least under the assumption of a perfect estimator with $\hat{r}(x|\theta_0, \theta_1) = r(x|\theta_0, \theta_1)$. The result given in Eq. (45) provides us with a handle to calculate the statistical uncertainty of this calibration procedure from the finite size of the calibration sample. Note that for imperfect estimators, the variance of $R[\hat{R}]$ may be significantly larger.

Independent of whether expectation calibration is part of the estimator, the deviation of the expectation $\hat{R}(\theta)$ from one can serve as a diagnostic tool to check for mismodeling of the estimator. We can take $\log \hat{R}(\theta)$ as a measure of the uncertainty of $\log \hat{r}(x|\theta, \theta_1)$. As in the case of the reference hypothesis variation, there is no consistent statistical interpretation of this uncertainty, but this does not mean that it is useless as a closure test.

Reweighting distributions.—A good estimator $\hat{r}(x|\theta_0, \theta_1)$ should satisfy

$$\mathbf{p}(x|\theta_0) \approx \hat{r}(x|\theta_0, \theta_1) \mathbf{p}(x|\theta_1). \quad (49)$$

We cannot evaluate the $\mathbf{p}(x|\theta_0)$ to check this relation explicitly. However, we *can* sample events $\{x_e\}$ from them. This provides another diagnostic tool [30]: we can draw a first sample as $x_e \sim \mathbf{p}(x_e|\theta_0)$, and draw a second sample as $x_e \sim \mathbf{p}(x_e|\theta_1)$ and reweight it with $\hat{r}(x_e|\theta_0, \theta_1)$. For a good likelihood ratio estimator, the two samples should have similar distributions. This can easily be tested by training a discriminative classifier between the samples. If a classifier can distinguish between the sample from the first hypothesis and the sample drawn from the second hypothesis and reweighted with the estimated likelihood ratio, then $\hat{r}(x|\theta_0, \theta_1)$ is not a good approximation of the true likelihood ratio $r(x|\theta_0, \theta_1)$. Conversely, if the classifier cannot separate the two classes, the classifier is either not efficient, or the likelihood ratio is estimated well.

Note that passing these closure tests is not a guarantee for a good estimator of the likelihood ratio. In Sec. IV B we will discuss how we can nevertheless derive exclusion limits that are guaranteed to be statistically correct, i.e., that might not be optimal, but are never wrong.

In our example process, we will use a combination of the first two ideas of this list: we will create copies of estimators with independent training samples and random seeds during training, as well as with different choices of the reference hypothesis θ_1 , and analyze the median and envelope of the predictions.

IV. LIMIT SETTING

The final objective of any EFT analysis is exclusion limits on the parameters of interest at a given confidence level. These can be derived in one of two ways. The Neyman construction based on toy experiments provides a generic and fail-safe method; we will discuss it in Sec. IV B. But since the techniques developed in the previous section directly provide

an estimate for the likelihood ratio, we can alternatively apply existing statistical methods for likelihood ratios as test statistics. This much more efficient approach will be the topic of the following section.

A. Asymptotics

Consider the test statistics

$$\begin{aligned} q(\theta) &= -2 \sum_e \log r(x_e | \theta, \hat{\theta}) \\ &= -2 \sum_e (\log r(x_e | \theta, \theta_1) - \log r(x_e | \hat{\theta}, \theta_1)) \end{aligned} \quad (50)$$

for a fixed number N of observed events $\{x_e\}$ with the maximum-likelihood estimator

$$\hat{\theta} = \arg \max_{\theta} \sum_e \log r(x_e | \theta, \theta_1). \quad (51)$$

In the asymptotic limit, the distribution according to the null hypothesis, $p(q(\theta) | \theta)$, is given by a chi-squared distribution. The number of degrees of freedom (d.o.f.) k is equal to the number of parameters θ . This result by Wilks [83] allows us to translate an observed value $q_{\text{obs}}(\theta)$ directly to a p -value that measures the confidence with which θ can be excluded:

$$p_{\theta} \equiv \int_{q_{\text{obs}}(\theta)}^{\infty} dq p(q | \theta) = 1 - F_{\chi^2}(q_{\text{obs}}(\theta) | k), \quad (52)$$

where $F_{\chi^2}(x | k)$ is the cumulative distribution function of the chi-squared distribution with k d.o.f. In our example process $k = 2$, for which this simplifies to

$$p_{\theta} = \exp\left(-\frac{q_{\text{obs}}(\theta)}{2}\right). \quad (53)$$

In particle physics it is common practice to calculate “expected exclusion contours” by calculating the expected value of $q_{\text{obs}}(\theta)$ based on a large “Asimov” data set generated according to some θ' [84]. With Eq. (52) this value is then translated into an expected p -value.¹²

¹²For $k = 1$, this standard procedure reproduces the median expected p -value. Note however that this is not true anymore for more than one parameter of interest. In this case, the median expected p -value can be calculated based on a different, but not commonly used, procedure. It is based on the fact that the distribution of q according to an alternate hypothesis, $p(q(\theta) | \theta')$, is given by a noncentral chi-squared distribution [85]. In the asymptotic limit, the noncentrality parameter is equal to the expectation value $\mathbb{E}[q(\theta) | \theta']$ [84]. This allows us to calculate for instance the median expected $q(\theta)$ assuming some value θ' based on the Asimov data. Combining all the pieces, the median expected p -value p_{θ} with which θ can be excluded under the assumption that θ' is true is then given by $p_{\theta}^{\text{expected from } \theta'} = 1 - F_{\chi^2}(F_{\chi^2}^{-1}(\frac{1}{2} | k, \mathbb{E}[q(\theta) | \theta']) | k)$, where $F_{\chi^2}(x | k)$ is the cumulative distribution function for the chi-squared distribution with k d.o.f. and $F_{\chi^2}^{-1}(p | k, \Lambda)$ is the inverse cumulative distribution function for the noncentral chi-squared distribution with k d.o.f. and noncentrality parameter Λ .

In practice we cannot access the true likelihood ratio defined on the full observable space and thus also not $q(\theta)$. But if the error of an estimator $\hat{r}(x | \theta_0, \theta_1)$ compared to the true likelihood ratio is negligible, we can simply calculate

$$\hat{q}(\theta) = -2 \sum_e \log \hat{r}(x_e | \theta, \hat{\theta}) \quad (54)$$

with maximum likelihood estimator $\hat{\theta}$ also based on the estimated likelihood ratio. The p -value can then be read off directly from the estimator output, substituting \hat{q} for q in Eq. (52).

Under this assumption and in the asymptotic limit, constructing confidence intervals is thus remarkably simple and computationally cheap: after training an estimator $\hat{r}(x | \theta, \theta_1)$ as discussed in the previous section, the observed events $\{x_e\}$ are fed into the estimator for each value of θ on some parameter grid. From the results we can read off the maximum likelihood estimator $\hat{q}(\theta)$ and calculate the observed value of the test statistics $\hat{q}(\theta)$ for each θ . Equation (52) then translates these values to p -values, which can then be interpolated between the tested θ points to yield the final contours.

To check whether these asymptotic properties apply to a likelihood ratio estimator, we can use the diagnostic tools discussed in Sec. III F 2. In addition, we can explicitly check whether the distribution of $q(\theta)$ actually follows a chi-squared distribution by generating toy experiments for a few θ points. If it does, the asymptotic results are likely to apply at other points in parameter space as well. If the variance of the toy experiments is larger than expected from the chi-squared distribution, the residual variance may be taken as an error estimate on the estimator prediction.

B. Neyman construction

Rather than relying on the asymptotic properties of the likelihood ratio test, we can construct the distribution of a test statistic with toy experiments. This is computationally more expensive, but useful if the number of events is not in the asymptotic regime or if the uncertainty of the estimators cannot be reliably quantified. Constraints derived in this way are conservative: even if the likelihood ratio is estimated poorly, the resulting contours might not be optimal, but they are never wrong (at a specified confidence level).

A good choice for the test statistics is the estimated profile log likelihood ratio $\hat{q}(\theta)$ given in Eq. (54), which allows us to compare the distribution of the toy experiments directly to the asymptotic properties discussed in the previous section. However, its construction requires finding the maximum likelihood estimator for every toy experiment. This increases the necessary computation time substantially, especially in high-dimensional parameter spaces. An alternative test statistics is the estimated log likelihood ratio with respect to some fixed hypothesis, which need not be identical to the reference denominator θ_1

used in the likelihood ratio estimators. In the EFT approach, the natural choice is the estimated likelihood ratio with respect to the SM,

$$\hat{q}'(\theta) \equiv -2 \sum_e \log \hat{r}(x_e | \theta, \theta_{\text{SM}}). \quad (55)$$

Using this test statistic rather than the profile likelihood ratio defined in Eq. (50) is expected to lead to stronger constraints if the true value of θ is close to the SM point, as expected in the EFT approach, and less powerful bounds if the true value is substantially different from the SM.

In practice we can efficiently calculate the distribution of $\hat{q}^{(i)}(\theta)$ after n events by first calculating the distribution of $\hat{q}^{(i)}(\theta)$ for one event and convolving the result with itself $(n-1)$ times.

C. Nuisance parameters

The tools developed above also support nuisance parameters, for instance to model systematic uncertainties in the theory calculation or the detector model. One strategy is to train parametrized estimators on samples generated with different values of the nuisance parameters ν and let them learn the likelihood ratio

$$\hat{r}(x | \theta_0, \theta_1; \nu_0, \nu_1) \equiv \frac{\hat{p}(x | \theta_0; \nu_0)}{\hat{p}(x | \theta_1; \nu_1)} \quad (56)$$

with its dependence on the nuisance parameters. As test statistics we can then use the estimator version of the usual profile log likelihood ratio,

$$\hat{q}(\theta) = -2 \sum_e \log \left[r(x_e | \theta, \hat{\theta}; \hat{\nu}) \frac{q(\hat{\nu})}{q(\hat{\nu})} \right] \quad (57)$$

with constraint terms $q(\nu)$,

$$\hat{\nu} = \arg \max_{\nu} \sum_e \log \left[\hat{r}(x_e | \theta, \theta_1; \nu, \nu_1) \frac{q(\nu)}{q(\nu_1)} \right], \quad \text{and} \quad (58)$$

$$(\hat{\theta}, \hat{\nu}) = \arg \max_{(\theta, \nu)} \sum_e \log \left[\hat{r}(x_e | \theta, \theta_1; \nu, \nu_1) \frac{q(\nu)}{q(\nu_1)} \right]. \quad (59)$$

The profile log likelihood ratio has two advantages: it is pivotal, i.e., its value and its distribution do not depend on the value of the nuisance parameter, and it has the asymptotic properties discussed in Sec. IV A.

Similarly, we can train the score including nuisance parameters,

$$\hat{t}(x | \theta_0; \nu_0) = \nabla_{(\theta, \nu)} \log [\hat{p}(x | \theta; \nu) q(\nu)]|_{\theta_0, \nu_0}. \quad (60)$$

If the constraints $q(\nu)$ limit the nuisance parameters to a relatively small region around some ν_0 , i.e., a range in

which the shape of the likelihood function does not change significantly, the SALLY and SALLINO methods seem particularly appropriate.

Finally, an adversarial component in the training procedure lets us directly train pivotal estimators $\hat{r}(x | \theta_0, \theta_1)$, i.e., that do not depend on the value of the nuisance parameters [86]. Compared to learning the explicit dependence on ν , this can dramatically reduce the dimensionality of the parameter space as early as possible, and does not require manual profiling. However, the estimators will generally not converge to the profile likelihood ratio, so its asymptotic properties do not apply and limit setting requires the Neyman construction.

V. RESULTS

We now apply the analysis techniques to our example process of WBF Higgs production in the 4ℓ decay mode. We first study the idealized setup discussed in Sec. II D 2, in which we can assess the techniques by comparing their predictions to the true likelihood ratio. In Sec. V B we then calculate limits in a more realistic setup.

A. Idealized setup

1. Quality of likelihood ratio estimators

Table IV summarizes the performance of the different likelihood ratio estimators in the idealized setup. For 50 000 events $\{x_e\}$ drawn according to the SM, we evaluate the true likelihood ratio $r(x_e | \theta_0, \theta_1)$ as well as the estimated likelihood ratios $\hat{r}(x_e | \theta_0, \theta_1)$ for 1000 values of θ_0 sampled randomly in $[-1, 1]^2$. As a metric we use the expected mean squared error on the log likelihood ratio

$$\varepsilon[\hat{r}(x)] = \sum_{\theta_0} \pi(\theta_0) \frac{1}{N} \sum_e [(\log \hat{r}(x_e | \theta_0, \theta_1) - \log r(x_e | \theta_0, \theta_1))^2]. \quad (61)$$

The 1000 tested values of θ_0 are weighted with a Gaussian prior

$$\pi(\theta) = \frac{1}{Z} \mathcal{N}(\|\theta\| | 0, 2 \cdot 0.2^2) \quad (62)$$

with normalization factor Z such that $\sum_{\theta_0} \pi(\theta) = 1$. In addition we show the expected trimmed mean squared error, which truncates the top 5% and bottom 5% of events for each θ . This allows us to analyze the quality of the estimators for the bulk of the phase space without being dominated by a few outliers. In Table IV and in the figures of this section, we only show results for a default set of hyperparameters for each likelihood ratio estimators. These default setups are defined in Appendix B. Results for other hyperparameter choices are given in Appendix C.

TABLE IV. Performance of the different likelihood ratio estimation techniques in our example process. The metrics shown are the expected mean squared error on the log likelihood ratio with and without trimming, as defined in the text. Checkmarks in the last column denote estimators shown in the following figures. Here we only give results based on default settings, which are defined in Appendix B. An extended list of results that covers more estimators is given in Appendix C.

Strategy	Setup	Expected Mean squared error (MSE)		Figures
		All	Trimmed	
Histogram	$p_{T,j1}, \Delta\phi_{jj}$	0.056	0.0106	✓
	$p_{T,j1}$	0.088	0.0230	
	$\Delta\phi_{jj}$	0.160	0.0433	
AFC	$p_{T,j1}, \Delta\phi_{jj}$	0.059	0.0091	
	$p_{T,j1}, m_{Z2}, m_{jj}, \Delta\eta_{jj}, \Delta\phi_{jj}$	0.078	0.0101	
CARL (PbP)	PbP	0.030	0.0111	Fig. 12
CARL (parametrized)	Baseline	0.012	0.0026	
	Random θ	0.012	0.0028	✓
CARL (morphing aware)	Baseline	0.076	0.0200	
	Random θ	0.086	0.0226	Fig. 12
	Morphing basis	0.156	0.0618	
ROLR (PbP)	PbP	0.005	0.0022	
ROLR (parametrized)	Baseline	0.003	0.0017	
	Random θ	0.003	0.0014	
ROLR (morphing aware)	Baseline	0.024	0.0063	
	Random θ	0.022	0.0052	
	Morphing basis	0.130	0.0485	
SALLY		0.013	0.0002	✓
SALLINO		0.021	0.0006	
CASCAL (parametrized)	Baseline	0.001	0.0002	✓
	Random θ	0.001	0.0002	
CASCAL (morphing aware)	Baseline	0.136	0.0427	
	Random θ	0.092	0.0268	
	Morphing basis	0.040	0.0081	
RASCAL (parametrized)	Baseline	0.001	0.0004	✓
	Random θ	0.001	0.0004	
RASCAL (morphing aware)	Baseline	0.125	0.0514	
	Random θ	0.132	0.0539	
	Morphing basis	0.031	0.0072	

The best results come from parametrized estimators that combine either a classifier decision function or ratio regression with regression on the score: the CASCAL and RASCAL strategies provide very accurate estimates of the log likelihood ratio. SALLY, parametrized ROLR, and parametrized CARL perform somewhat worse. For CARL and ROLR, parametrized estimators consistently perform better than the corresponding point-by-point versions. All these ML-based strategies significantly outperform the traditional one- or two-dimensional histograms and the approximate frequentist computation.

The morphing-aware versions of the parametrized estimators lead to a poor performance, comparable or worse than the two-dimensional histogram approach. As anticipated in Sec. III F, the large weight factors and the sizable cancellations between them blow up small errors on the estimation of the individual basis estimators $\hat{r}_i(x)$ to large errors on the combined estimator.

We find that probability calibration as discussed in Sec. III D 1 improves the results in almost all cases, in particular for the CARL method. An additional step of expectation calibration (see Sec. III D 2) after the probability calibration does not lead to a further improvement, and in fact often increases the variance of the estimator predictions. We therefore only use probability calibration for the results presented here.

The choice of the training sample is less critical, with nearly identical results between the baseline and random θ samples. For the CARL approach, shallow networks with two hidden layers perform better, while ROLR works best for three hidden layers and the score-based strategies benefit from a deeper network with five hidden layers.

In Fig. 10 we show scatter plots between the true and estimated likelihood ratios for a fixed hypothesis θ_0 . The likelihood ratio estimate from histograms of observables is widely spread around the true likelihood ratio, reflecting




FIG. 10. True vs estimated likelihood ratios for a benchmark hypothesis $\theta_0 = (-0.5, -0.5)^T$. Each dot corresponds to one event x_e . The CASCAL (right, red), RASCAL (right, orange), and SALLY (middle, blue) techniques can predict the likelihood ratio extremely accurately over the whole phase space. All new techniques clearly lead to more precise estimates than the traditional histogram approach (left, orange).

the loss of information from ignoring most directions in the observable space. CARL performs clearly better. ROLR and SALLY offer a further improvement. Again, the best results come from the CASCAL and RASCAL strategies, both giving predictions that are virtually one-to-one with the true likelihood ratio.

We go beyond a single benchmark point θ_0 in Fig. 11. This scatter plot compares true and estimated likelihood ratios for different values of θ_0 , taking the expectation value over x . We find that the CARL, ROLR, CASCAL, and RASCAL approaches converge to the correct likelihood ratio in this expectation value. For the SALLY and SALLINO techniques we find larger deviations, pointing towards the breakdown of the local model approximation. Much more obvious is the loss of information in the traditional histogram approach, which is clearly not asymptotically exact.

The point-by-point, agnostic parametrized, and morphing-aware versions of the CARL strategy are compared in Fig. 12. As expected from Table IV, the parametrized strategy performs better than the point-by-point version,

and both are clearly superior to the morphing-aware estimator.

2. Efficiency and speed

With infinite training data, many of the algorithms should converge to the true likelihood ratio. But generating training samples can be expensive, especially when a full detector simulation is used. An important question is therefore how much training data the different techniques require to perform well. In Fig. 13 we show the performance as a function of the training sample size.

The SALLY approach performs very well even with very little data. Its precision stops improving eventually, showing the limitations of the local model approximation. For the other methods we find that the more information a technique uses, the less training data points it requires. The RASCAL technique utilizes the most information from the simulator, leading to an exceptional performance with training samples of approximately 100 000 events. This is in contrast to the most general CARL method, which




FIG. 11. True vs estimated expected log likelihood ratio. Each dot corresponds to one value of θ_0 , where we take the expectation over $x \sim p(x|\theta_{SM})$. The new techniques are less biased than the histogram approach (left, orange).

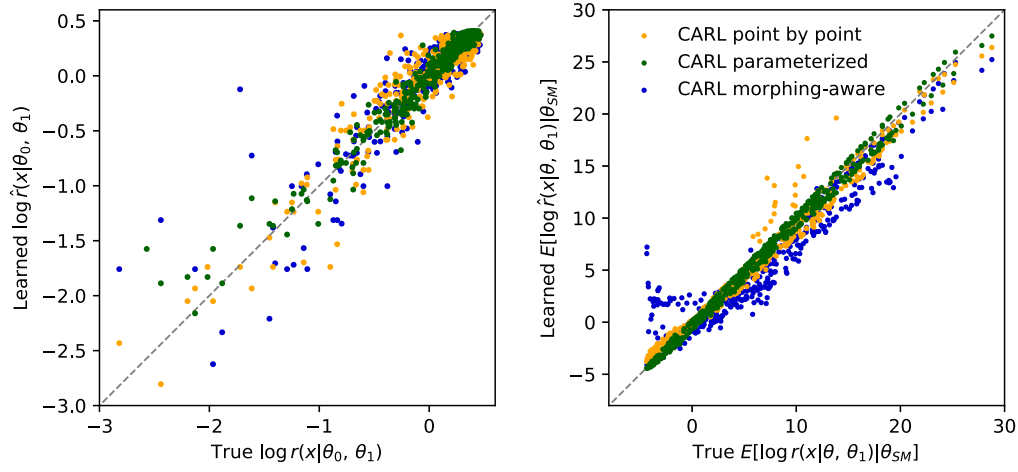


FIG. 12. Comparison of the point-by-point, parametrized, and morphing-aware versions of CARL. Top left: True vs estimated likelihood ratios for a benchmark hypothesis $\theta_0 = (-0.5, -0.5)^T$, as in Fig. 10. Each dot corresponds to one event x_e . Top right: True vs estimated expected log likelihood ratio, as in Fig. 11. Each dot corresponds to one value of θ_0 , where we take the expectation over $x \sim p(x|\theta_{SM})$. The parametrized estimator outperforms the point-by-point one and particularly the morphing-aware version.

does not use any of the extra information from the simulator and requires a 2 orders of magnitude larger training sample for a comparable performance.

In Fig. 14 we show the evolution of the likelihood estimation error and the cross entropy of the classification problem during the training of the parametrized estimators. For comparison, we also show the optimal metrics based on the true likelihood ratio, and the results of the two-dimensional histogram approach. Once again we see that either CASCAL or RASCAL leads to the best results. This result also holds true for the cross entropy, hinting that the

techniques we use to measure continuous parameters might also improve the power of estimators in discrete classification problems. Note that the CARL approach is more prone to overfitting than the others, visible as a significant difference between the metrics evaluated on the training and validation samples.

Equally important to the training efficiency is the computation time taken up by evaluating the likelihood ratio estimators $\hat{r}(x_e|\theta_0, \theta_1)$. We compare example evaluation times in Table V. The traditional histogram approach takes the shortest time. But all tested algorithms are very

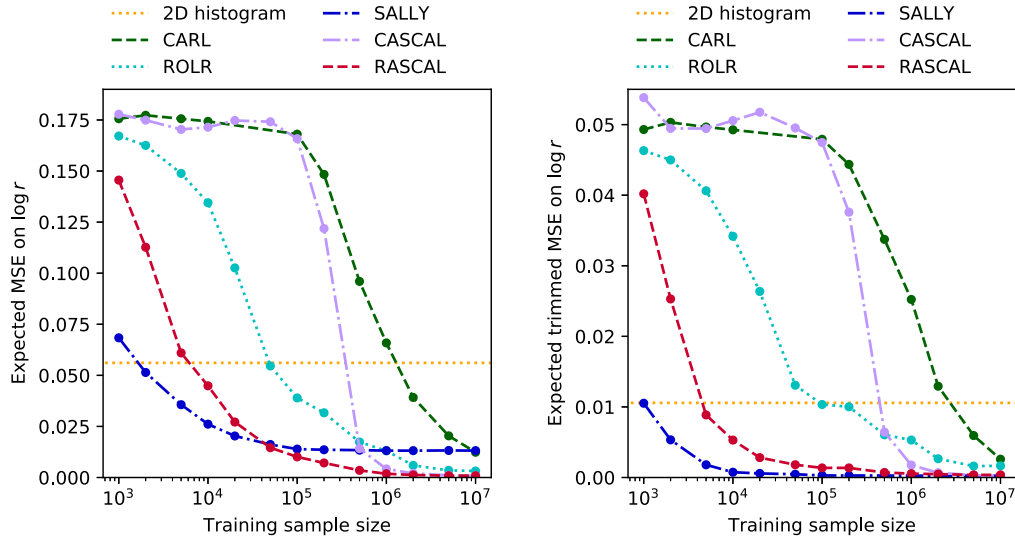


FIG. 13. Performance of the techniques as a function of the training sample size. As a metric, we show the mean squared error (left) and trimmed mean squared error on $\log r(r|\theta_0, \theta_1)$ weighted with a Gaussian prior, as discussed in the text. Note that we do not vary the size of the calibration data samples. The number of epochs is increased such that the number of epochs times the training sample size is constant; all other hyperparameters are kept constant. The SALLY method works well even with very little data, but plateaus eventually due to the limitations of the local model approximation. The other algorithms learn faster the more information from the simulator is used.

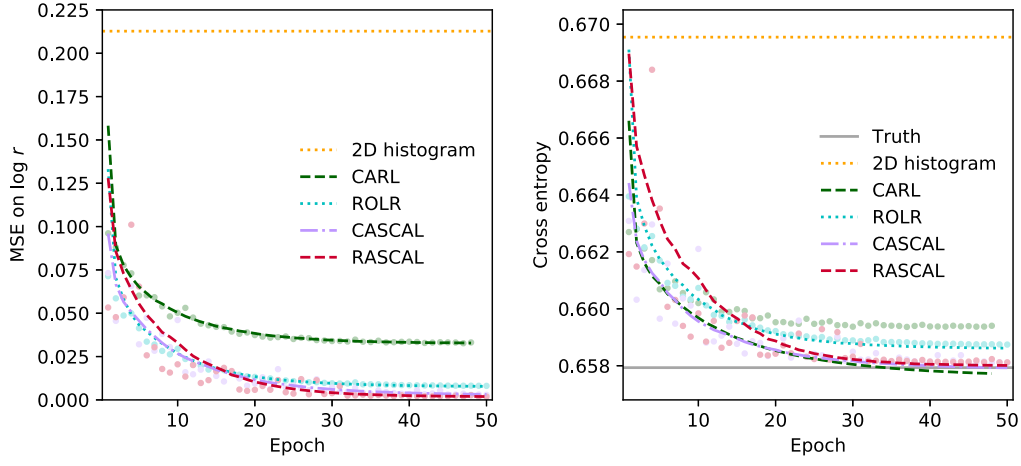


FIG. 14. Learning curve of the parametrized models. The solid lines show the metrics evaluated on the training sample, the dots indicate the performance on the validation sample. Note that these numbers are not comparable to the metrics in Table IV and Fig. 13, which are weighted with the prior in Eq. (62). These results also do not include calibration. Left: Mean squared error of $\log \hat{r}(x|\theta_0, \theta_1)$. Right: Binary cross entropy of the classification based on $\hat{s}(x|\theta_0, \theta_1)$ between the numerator and denominator samples. The solid grey line shows the “optimal” performance based on the true likelihood ratio. The CASCAL and RASCAL techniques converge to a performance close to the theoretic optimum. The CARL approach (green), based on minimizing the cross entropy, shows signs of overfitting. All machine-learning-based methods outperform traditional histograms (dashed orange).

fast: the likelihood ratio for fixed hypotheses (θ_0, θ_1) for 50 000 events $\{x_e\}$ can always be estimated in around one second or less. The local score regression method is particularly efficient, since the estimator $\hat{i}(x|\theta_{\text{score}}, \theta_1)$ has to be evaluated only once to estimate the likelihood ratio for any value of θ_0 . Only the comparably fast step of density estimation has to be repeated for each tested value of θ_0 .

So after investing some training time upfront, all the measurement strategies developed here can be evaluated on any events with very little computational cost and amortize

TABLE V. Computation times of evaluating $\hat{r}(x|\theta_0, \theta_1)$ in the different algorithms. We distinguish between steps that have to be calculated once per x and those which have to be repeated for every evaluated value of θ_0 . These numbers are from one run of our algorithms with default settings on the NYU HPC cluster on machines equipped with Intel Xeon E5-2690v4 2.6 GHz CPUs and NVIDIA P40 GPUs with 24 GB RAM, using a batch of 50000 events $\{x_e\}$, and taking the mean over 1017 values of θ_0 . The local score regression method and the traditional histogram method are particularly fast. But all techniques are many orders of magnitude faster to evaluate than the matrix element method or optimal observables.

Algorithm	Evaluation time [μ s]	
	per x_e	per x_e and θ_0
Histogram		0.2
CARL		19.7
SALLY	25.4	0.1
ROLR		19.7
CASCAL		25.1
RASCAL		21.7

quickly. While this is not the focus of our paper, note that this distinguishes our approaches from the matrix element method and optimal observable techniques. These well-established methods require the computationally expensive evaluation of complicated numerical integrals for every evaluation of the likelihood ratio estimator.

3. Physics results

The most important result of an EFT measurement is observed and expected exclusion contours, either based on asymptotics or toy experiments. In the asymptotic approach, the expected contours are determined just by the likelihood ratio evaluated on a large Asimov data set, as described in Sec. IV A. Figure 15 shows this expected log likelihood ratio in the SM after 36 events over a one-dimensional slice of the parameter space. In Fig. 16 we show the corresponding expected exclusion limits on the two Wilson coefficients. To estimate the robustness of the likelihood ratio estimators, each algorithm is run 5 times with different choices of the reference hypothesis; independent training, calibration, and evaluation samples; and independent random seeds during training. The lines show the median of the five replicas, while the shaded bands show the envelope. While this error band does not have a clear statistic interpretation, it does provide a diagnostic tool for the variance of the estimators.

A traditional histogram-based analysis of jet p_T and $\Delta\phi_{jj}$ leads to overly conservative results. It is interesting to note that this simple analysis works reasonably well in the region of parameter space with $f_W > 0$ and $f_{WW} > 0$, which is exactly the part of parameter space where informative high-energy events interfere mostly constructively with the SM




FIG. 15. Expected likelihood ratio with respect to the Standard Model along a one-dimensional slice of the parameter space. We assume 36 observed events and the SM to be true. For each estimator, we generate five sets of predictions with different reference hypotheses, independent data samples, and different random seeds. The lines show the median of this ensemble, the shaded error bands the envelope. All machine-learning-based methods reproduce the true likelihood function well, while the doubly differential histogram method underestimates the likelihood ratio in the region of negative Wilson coefficients.




FIG. 16. Expected exclusion contours based on asymptotics at 68% C.L. (innermost lines), 95% C.L., and 99.7% C.L. (outermost lines). We assume 36 observed events and the SM to be true. As test statistics, we use the profile likelihood ratio with respect to the maximum-likelihood estimator. For each estimator, we generate five sets of predictions with different reference hypotheses, independent data samples, and different random seeds. The lines show the median of this ensemble, the shaded error bands the envelope. The new techniques based on machine learning, in particular the CASCAL and RASCAL techniques, lead to expected exclusion contours very close to those based on the true likelihood ratio. An analysis of a doubly differential histogram leads to much weaker bounds.

amplitude. In the $f_W < 0$ region of parameter space, destructive interference dominates in the important regions of phase space with large momentum transfer. An extreme example is the “amplitude-through-zero” effect shown in the left panel of Fig. 2. Simple histograms with a rough binning generally lead to a poor estimation of the likelihood ratio in such complicated kinematic signatures.

We find that the new ML-based strategies allow us to place visibly tighter constraints on the Wilson coefficients than the doubly differential histogram. In particular the CARL + score and regression + score estimators lead to exclusion contours that are close to the contours based on the true likelihood ratio. In this analysis based on asymptotics, however, it is possible for the estimated contours to

be slightly too tight, wrongly marking parameter regions as excluded at a given confidence level. This problem can be mitigated by profiling over systematic uncertainties assigned to the likelihood ratio estimates.

Exclusion limits based on the Neyman construction do not suffer from this issue: contours derived in this way might be not optimal, but they are never wrong. We generate toy experiments to estimate the distribution of the likelihood ratio with respect to the SM for individual events. Repeatedly convolving this single-event distribution with itself, we find the distribution of the likelihood ratio after 36 observed events.

The expected corresponding expected exclusion limits are shown in Fig. 17. Indeed, errors in the likelihood ratio




FIG. 17. Expected exclusion contours based on the Neyman construction with toy experiments at 68% C.L. (innermost lines), 95% C.L., and 99.7% C.L. (outermost lines). We assume 36 observed events and the SM to be true. As test statistics, we use the likelihood ratio with respect to the SM. All machine-learning-based methods let us impose much tighter bounds on the Wilson coefficients than the traditional histogram approach (left, dotted orange). The Neyman construction guarantees statistically correct results: no contour based on estimators excludes parameter points that should not be excluded. The expected limits based on the CASCAL (right, lavender) and RASCAL (right, red) techniques are virtually indistinguishable from the true likelihood contours.

estimation never lead to undercoverage, i.e., the exclusion of points that should not be excluded based on the true likelihood ratio. Again we find that histograms of kinematic observables only allow us to place rather weak bounds on the Wilson coefficients. SALLY performs clearly better, with excellent performance close to the SM. Deviations from the optimal bounds become visible at the 2σ level, hinting at the breakdown of the local model approximation there. The best results come once more from the CASCAL and RASCAL methods. Both of these strategies yield exclusion bounds that are virtually indistinguishable from those based on the true likelihood ratio.

As a side note, a comparison of the expected contours based on asymptotics to those based on the Neyman construction shows the Neyman results to be tighter.

This reflects the different test statistics used in the two figures: in the asymptotics case, we use the profile likelihood ratio with respect to the maximum likelihood estimator, which itself fluctuates around the true value of θ (which in our case is assumed to be the SM). In the Neyman construction we use the likelihood ratio with respect to the SM, leading to tighter contours if the true value is in fact close to the SM, and weaker constraints if it is very different.

B. Detector effects

We have now established that the measurement strategy works very well in an idealized setup, where we can compare them to the true likelihood ratio. In a next step,




FIG. 18. Ratio of the estimated likelihood ratio $\hat{r}(x|\theta_0, \theta_1)$ to the *joint* likelihood ratio $r(x, z|\theta_0, \theta_1)$, which is conditional on the parton-level momenta and other latent variables. As a benchmark hypothesis we use $\theta_0 = (-0.5, -0.5)^T$, the events are drawn according to the SM. The spread common to all methods shows the effect of the smearing on the likelihood ratio. The additional spread in the histogram, CARL, and ROLR methods is due to a poorer performance of these techniques.




FIG. 19. Expected exclusion contours based on the Neyman construction with toy experiments at 68% C.L., 95% C.L., and 99.7% C.L. with smearing. We assume 36 observed events and the SM to be true. As test statistics, we use the likelihood ratio with respect to the SM. In the setup with smearing we cannot these results to the true likelihood contours. But since the Neyman construction is guaranteed to cover, these expected limits are correct. The new techniques, in particular CASCAL (right, dashed red) and RASCAL (right, dash-dotted orange), allow us to set much tighter bounds on the Wilson coefficients than a traditional histogram analysis (left, dotted orange).

we turn towards a setup with a rudimentary smearing function that models the effect of the parton shower and the detector response on the observables. In this setting, the true likelihood is intractable, so we cannot use it as a baseline to validate the predictions any more. But we can still discuss the relative ordering of the exclusion contours predicted by the different estimators.

Figure 18 shows the relation between the true *joint* likelihood ratio $r(x, z|\theta_0, \theta_1)$, which is conditional on the parton-level momenta z and other latent variables, to the estimated likelihood ratio $\hat{r}(x|\theta_0, \theta_1)$, which only depends on the observables x . We see that this relation is stochastically smeared out around 1. Recall that in the idealized scenario the best estimators described the true likelihood ratio perfectly, as shown in Fig. 10. This strongly suggests that the spread visible here is not due to errors of the likelihood ratio estimators, but rather shows the difference between the joint and true likelihood ratios, as illustrated in Fig. 5.

In Fig. 19 we show the expected exclusion contours based on the Neyman construction, which guarantees statistically correct results. The conclusions from the idealized setting are confirmed: a measurement based on the likelihood ratio estimators leads to robust bounds that are clearly more powerful than those based on a histogram. Once again, the CASCAL and RASCAL algorithms lead to the strongest limits.

VI. CONCLUSIONS

We have developed and analyzed a suite of new analysis techniques for measurements of continuous parameters in LHC experiments based on simulations and machine learning. Exploiting the structure of particle physics

processes, they extract additional information from the event generators, and use this information to train precise estimators for likelihood ratios.

Our approach is designed for problems with large numbers of observables, where the likelihood function is not tractable and traditional methods based on individual kinematic variables often perform poorly. It scales well to high-dimensional parameter spaces such as that of effective field theories. The new methods do not require any approximations on the hard process, parton shower, or detector effects, and the likelihood ratio for any event and hypothesis pair can be evaluated in microseconds. These two properties set it apart from the matrix element method or optimal observables, which rely on crude approximations for the shower and detector and require the evaluation of typically very expensive integrals.

Using Higgs production in weak boson fusion in the four-lepton mode as a specific example process, we have evaluated the performance of the different methods and compared them to a classical analysis of the jet momentum and azimuthal angle between the tagging jets. We find that the new algorithms provide very precise estimates of arbitrary likelihood ratios. Using them as a test statistics allows us to impose significantly tighter constraints on the EFT coefficients than the traditional kinematic histograms.

Out of the several methods introduced and discussed in this paper, two stand out. The first, which we call SALLY, is designed for parameter regions close to the Standard Model:

- (1) As training data, the algorithm requires a sample of fully simulated events, each accompanied by the corresponding *joint score* at the SM: the relative change of the parton-level likelihood function of the parton-level momenta associated with this event under small changes of the theory parameters away

from the SM. This can be calculated by evaluating the squared matrix element at the same phase-space points for different theory parameters. We can thus extract this quantity from Monte Carlo generators such as MADGRAPH.

- (2) Regressing on this data, we train an estimator (for instance realized as a neural network) that takes as input an observation and returns the score at the SM. This function compresses the high-dimensional observable space into a vector with as many components as parameters of interest. If the parameter space is high dimensional, this can be even further compressed into the scalar product between the score vector and the difference between two parameter points.
- (3) The estimated score (or the scalar product between score and parameter difference) can then be treated like any set of observables in a traditional analysis. We can fill histograms of this quantity for different hypotheses, and calculate likelihood ratios from them.

There are two key ideas that underlie this strategy. First, note that the training data only consists of the *joint score*, which depends on the parton-level four-momenta of an event. But during the training the estimator converges to the actual score of the distribution of the observables, i.e., the relative change of the actual likelihood function under infinitesimal changes of the parameters. We have proven this powerful, yet surprisingly simple relation in this paper.

Second, close to the Standard Model (or any other reference parameter point), the score provides the sufficient statistics: it encapsulates all information on the local approximation of the statistical model. In other words, if the score is estimated well, the dimensionality reduction from high-dimensional observables into a low-dimensional vector does not lose any information on the parameters. The estimated score is a machine-learning version of the optimal observable idea, but requires neither approximations of the parton shower or detector treatment nor numerically expensive integrals.

As a matter of fact, the dimensionality reduction can be taken one step further. We have introduced the SALLINO technique that compresses the estimated score vector to a single scalar function, again without loss of power in the local approximation, and independent of the number of theory parameters.

In our example process, these simple and robust analysis strategies work remarkably well, especially close to the Standard Model. Deviations appear at the 2σ level, but even there it allows for much stronger constraints than a traditional analysis of kinematic variables. It requires significantly less data to train than the other discussed methods. Since the SALLINO method can compress any observation into a single number without losing much sensitivity, even for hundreds of theory parameters, this approach scales

exceptionally well to high-dimensional parameter spaces, as in the case of the SMEFT.

The second algorithm we want to highlight here is the RASCAL technique. Using even more information available from the simulator, it learns a parametrized likelihood ratio estimator: one function that takes both the observation and a theory parameter point as input and returns an estimate for the likelihood ratio between this point and a reference hypothesis given the observation. This estimator is constructed as follows:

- (1) Training this parametrized estimator requires data for many different values of the tested parameter point (the numerator in the likelihood ratio). For simplicity, the reference hypothesis (the denominator in the likelihood ratio) can be kept fixed. For each of these hypothesis pairs, event samples are generated according to the numerator and denominator hypothesis. In addition, we extract the joint likelihood ratio from the simulator: essentially the squared matrix element according to the numerator theory parameters divided by the squared matrix element according to the denominator hypothesis, evaluated at the generated parton-level momenta. Again, we also need the joint score, i.e., the relative change of the parton-level likelihood function under infinitesimal changes of the theory parameters. Both quantities can be extracted from matrix element codes.
- (2) A neural network models the estimated likelihood ratio as a function of both the observables and the value of the theory parameters (of the numerator in the likelihood ratio). We can calculate the gradient of the network output with respect to the theory parameter and thus also the estimated score. The network is trained by minimizing the squared error of the likelihood ratio plus the squared error of the score, in both cases with respect to the joint quantities extracted from the simulator.
- (3) After the training phase, the likelihood ratio can optionally be calibrated, for instance through isotonic regression.

This technique relies on a similar trick as the local score regression method: the likelihood ratio learned during the training converges to the true likelihood ratio, even though the joint ratio information in the training data is conditional on the parton-level momenta. The RASCAL method is among the best-performing methods of all analyzed techniques. It requires significantly smaller training samples than all other approaches, with the exception of SALLY and SALLINO. Expected exclusion limits derived in this way are virtually indistinguishable from those based on the true likelihood ratio.

On top of these two approaches, we have developed, analyzed, and compared several other methods. We refer the reader to the main part and the Appendix of this document, where all these algorithms are discussed in depth.

All tools developed here are suitable for large-scale LHC analyses. On the software side, only few modifications of existing tools are necessary. Most importantly, matrix-element generators should provide a user-friendly interface to calculate the squared matrix element for a given configuration of four-momenta and a given set of physics parameters. With such an interface, one could easily calculate the joint score and joint likelihood ratio data that is needed for the new algorithms. The training of the estimators is then straightforward, in particular for the SALLY and SALLINO methods. The limit setting follows established procedures, either based on the Neyman construction with toy experiments, or (since the tools provide direct estimates for the likelihood ratio) using asymptotic formulas.

While we have focused on the example of effective field theory measurements, these techniques equally apply to other measurements of continuous parameter in collider experiments as well as to a large class of problems outside of particle physics [49]. Some of the techniques can also be applied to improve the training of machine-learning-based classifiers. Finally, while we restricted our analysis to frequentist confidence intervals, as is common in particle physics, the same ideas can be used in a Bayesian setting.

All in all, we have presented a range of new inference techniques based on machine learning, which exploit the structure of particle physics processes to augment training data. They scale well to large-scale LHC analyses with many observables and high-dimensional parameter spaces. They do not require any approximations of the hard process, parton shower, or detector effects, and the likelihood ratio can be evaluated in microseconds. In an example analysis, these new techniques have demonstrated the potential to substantially improve the precision and new physics reach of the LHC legacy results.

ACKNOWLEDGMENTS

We would like to thank Cyril Becot and Lukas Heinrich, who contributed to this project at an early stage. We are grateful to Felix Kling, Tilman Plehn, and Peter Schichtel for providing the MADMAX code and helping us use it. K. C. wants to thank Centre for Cosmology, Particle Physics and Phenomenology (CP3) at UC Louvain for their hospitality. Finally, we would like to thank Atılım Güneş Baydin, Lydia Brenner, Joan Bruna, Kyunghyun Cho, Michael Gill, Ian Goodfellow, Daniela Huppenkothen, Hugo Larochelle, Yann LeCun, Fabio Maltoni, Jean-Michel Marin, Iain Murray, George Papamakarios, Duccio Pappadopulo, Dennis Prangle, Rajesh Ranganath, Dustin Tran, Rost Verkerke, Wouter Verkerke, Max Welling, and Richard Wilkinson for interesting discussions. J. B., K. C., and G. L. are grateful for the support of the Moore-Sloan data science environment at New York University (NYU). K. C. and G. L. were supported through the National Science Foundation (NSF) Grants No. ACI-1450310 and No. PHY-1505463. J. P. was partially supported by the Scientific and Technological Center of Valparaíso (CCTVal) under Fondecyt Grant No. BASAL FB0821. This work was supported in part through the NYU IT High Performance Computing resources, services, and staff expertise.

APPENDIX A: SIMPLIFIED DETECTOR DESCRIPTION

While most of our results are based on an idealized perfect measurement of parton-level momenta, we also consider a toy smearing representing the effect of parton shower and the detector. The total smearing function is given by

$$p(x|z) = \prod_{\ell \in \text{leptons}} p_{\ell}(x_{\ell}|z_{\ell}) \prod_{j \in \text{jets}} p_j(x_j|z_j). \quad (\text{A1})$$

Lepton momenta $x_{\ell} = (\hat{E}, \hat{p}_T, \hat{\eta}, \hat{\phi})$ are smeared by

$$p_{\ell}(\hat{E}, \hat{p}_T, \hat{\eta}, \hat{\phi}|E, p_T, \eta, \phi) = \mathcal{N}(\hat{p}_T|p_T, (3 \times 10^{-4} \text{ GeV}^{-1} p_T^2)^2) \cdot \delta(\hat{E} - E_0(\hat{p}_T, \hat{\eta}; m_{\ell})) \delta(\hat{\eta} - \eta) \delta(\hat{\phi} - \phi), \quad (\text{A2})$$

while the distribution of the jet properties depending on the quark momenta is given by

$$\begin{aligned} p_j(\hat{E}, \hat{p}_T, \hat{\eta}, \hat{\phi}|E, p_T, \eta, \phi) = & (\mathcal{N}(\hat{E}|a_0 + a_1\sqrt{E} + a_2E, (b_0 + b_1\sqrt{E} + b_2E)^2) \\ & + c\mathcal{N}(\hat{E}|d_0 + d_1\sqrt{E} + d_2E, (e_0 + e_1\sqrt{E} + e_2E)^2)) \\ & \cdot \delta(\hat{p}_T - p_{T0}(\hat{E}, \hat{\eta}; m_{\ell})) \mathcal{N}(\hat{\eta}|\eta, 0.1^2) \mathcal{N}(\hat{\phi}|\phi, 0.1^2). \end{aligned} \quad (\text{A3})$$

Here $\mathcal{N}(x|\mu, \sigma^2)$ is the Gaussian distribution with mean μ and variance σ^2 . The jet energy resolution parameters a_i, b_i, c, d_i , and e_i are based on the default settings of the jet transfer function in MADWEIGHT [87]. The functions $E_0(p_T, \eta, m)$ and $p_{T0}(E, \eta, m)$ refer to the energy and transverse momentum corresponding to an on-shell particle with mass m .

APPENDIX B: MODEL ALMANAC

In Sec. III A we developed different estimators for the likelihood ratio, focusing on the key ideas over technical details. Here we fill in the gaps, explain all strategies in a self-contained way, and document the settings we use for our example process. To facilitate their comparison, we describe all models in terms of a “training” and an

“evaluation” part, even if this language is not typically used e. g. for histograms.

1. Histograms of observables

Idea.—Most collider measurements are based on the number of events or the cross section of a process in a given phase-space region or on the differential cross section or distribution of one or at most a few kinematic

observables v . Typical choices are the reconstructed energies, momenta, angles, or invariant masses of particles. Choosing the right set of observables for a given measurement problem is all but trivial, but many processes have been studied extensively in the literature. Once this choice is made, this strategy is simple, fast, and intuitive. We illustrate the information used by this approach in the top panels of Fig. 20.

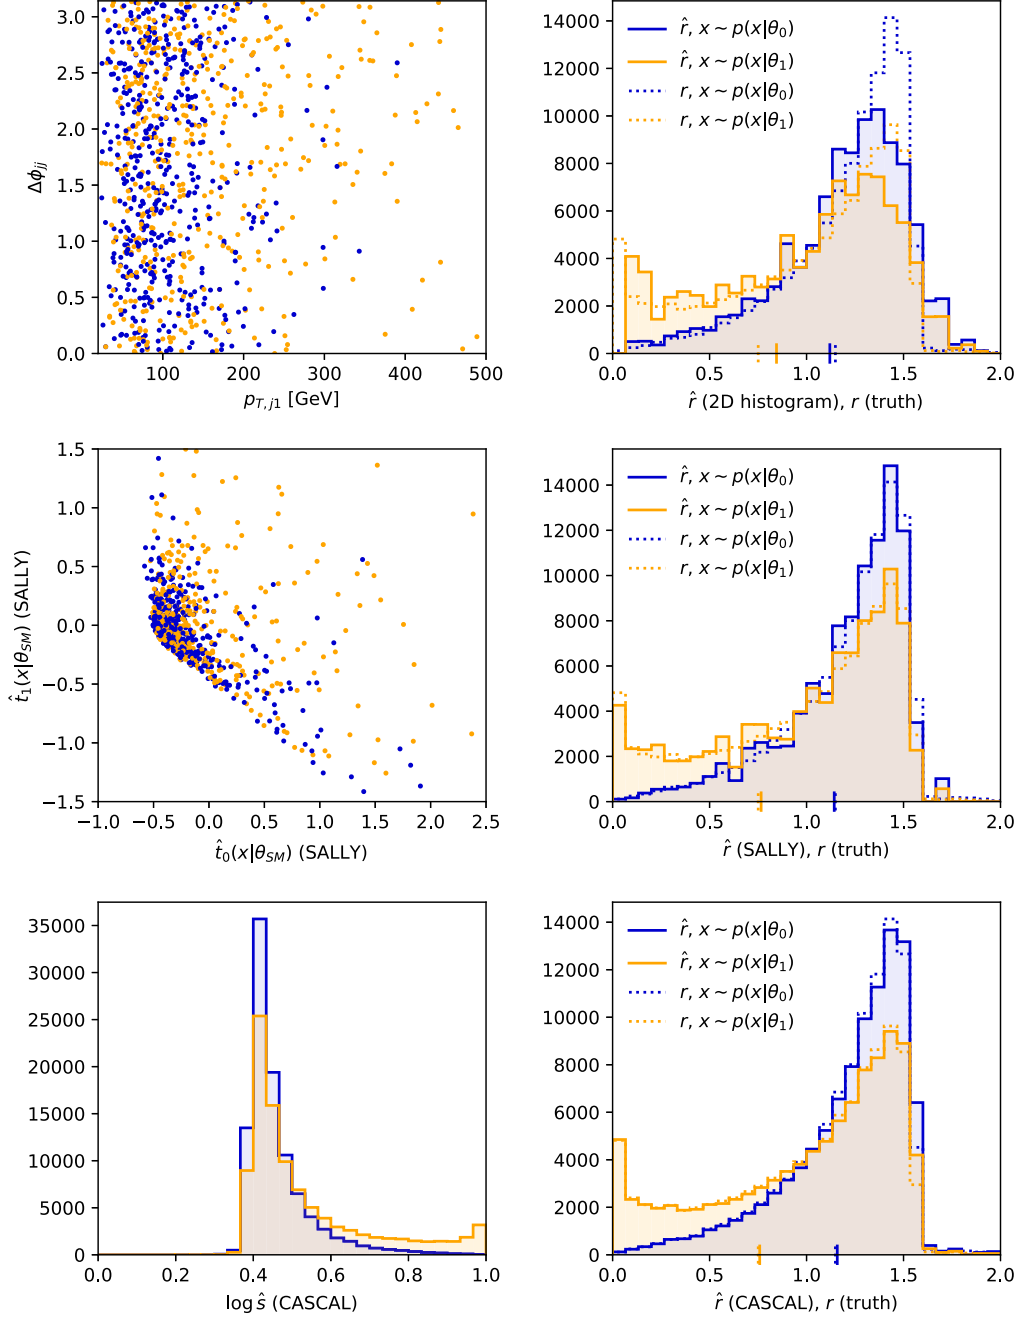


FIG. 20. Example distributions to illustrate the doubly differential histogram analysis (top), the SALLY technique (middle), and the CASCAL method (bottom). The left panels show the different spaces in which the densities and their ratios are estimated. On the right we show the corresponding distributions of the estimated ratio \hat{r} (solid) and compare them to the true likelihood ratio distributions (dotted). We use the benchmark $\theta_0 = (-0.5, -0.5)^T$ (blue) and θ_1 as in Eq. (47) (orange).

Requirements.—The histogram approach can be used in the general likelihood-free setting: it only requires a simulator that can generate samples $\{x\} \sim p(x|\theta)$.

Structure.—Histograms are most commonly used point by point in θ . If the problem has the morphing structure discussed in Sec. II C 2, they can also be applied in a morphing-aware parametrized way (we have not implemented this for our example process).

Training.—After generating samples for both the numerator and denominator hypotheses, the values of the chosen kinematic variables $v(x)$ are extracted, and binned into separate histograms for the two hypotheses.

Calibration.—With sufficient training data, histograms should be well calibrated, so we do not experiment with an additional calibration stage.

Evaluation.—To estimate the likelihood ratio between two hypotheses θ_0 and θ_1 for a given set of observables x , one has to extract the kinematic variables $v(x)$ and look up the corresponding bin contents in the histograms for θ_0 and θ_1 . Assuming equal binning for both histograms, the likelihood ratio is simply estimated as the ratio of bin contents.

Parameters.—The only parameters of this approach are the choices of kinematic observables and the histogram binning.

In our example process, we consider six different variants:

- (i) A one-dimensional histogram of the transverse momentum $p_{T,j1}$ of the leading (higher- p_T) jet with 80 bins.
- (ii) A one-dimensional histogram of the absolute value of the azimuthal angle $\Delta\phi_{jj}$ between the two jets with 20 bins.
- (iii) A “coarse” two-dimensional histogram of these two variables with ten bins in the $p_{T,j1}$ direction and five bins along $\Delta\phi_{jj}$.
- (iv) A “medium” two-dimensional histogram of these two variables with 20 bins in the $p_{T,j1}$ direction and ten bins along $\Delta\phi_{jj}$.
- (v) A “fine” two-dimensional histogram of these two variables with 30 bins in the $p_{T,j1}$ direction and 15 bins along $\Delta\phi_{jj}$.
- (vi) A “very fine” two-dimensional histogram of these two variables with 50 bins in the $p_{T,j1}$ direction and 20 bins along $\Delta\phi_{jj}$.
- (vii) An “asymmetric” two-dimensional histogram of these two variables with 50 bins in the $p_{T,j1}$ direction and five bins along $\Delta\phi_{jj}$.

For each pair (θ_0, θ_1) and each observable, the bin edges are chosen such that the same expected number of events according to θ_0 plus the expected number of events according to θ_1 is the same in each bin.

2. Approximate frequentist computation (AFC)

Idea.—Approximate Bayesian computation is a very common technique for likelihood-free inference in a

Bayesian setup. In its simplest form it keeps samples according to the rejection probability of Eq. (29). This amounts to an approximation of the likelihood function through kernel density estimation, which we can isolate from the ABC sampling mechanism and use in a frequentist setting. We call it approximate frequentist computation (AFC) to stress the relation to ABC. Just as ABC or the histogram approach, it requires the choice of a summary statistics $v(x)$, which in our example process we take to be a two-dimensional or five-dimensional subset of the kinematic variables.

Requirements.—AFC can be used in the general likelihood-free setting: it only requires a simulator that can generate samples $\{x\} \sim p(x|\theta)$.

Structure.—We use AFC point by point in θ . If the problem has the morphing structure discussed in Sec. II C 2, it can also be applied in a morphing-aware parametrized way.

Training.—For each event in the numerator and denominator training samples, the summary statistics $v(x)$ are calculated and saved.

Calibration.—AFC can be calibrated as any other technique on this list, but we left this for future work.

Evaluation.—The summary statistics $v(x)$ is extracted from the observation. For numerator and denominator hypothesis separately, the likelihood at this point is estimated with Eq. (30). The likelihood ratio estimate is then simply given by the ratio between the estimated numerator and denominator densities.

Parameters.—Just as for histograms, the choice of the summary statistics is the most important parameter. The performance of AFC also crucially depends on the kernel and bandwidth ε . Too small values for the bandwidth make large training samples necessary, too large values lead to an oversmoothing and loss of information.

In our example process, we consider two different variants:

- (i) A two-dimensional summary statistics space of the leading jet p_T and $\Delta\phi_{jj}$ (see above). Both variables are rescaled to zero mean and unit variance.
- (ii) A five-dimensional summary statistics space of the leading jet p_T , $\Delta\phi_{jj}$, the dijet invariant mass m_{jj} , the separation in pseudorapidity between the jets $\Delta\eta_{jj}$, and the invariant mass of the lighter (off-shell) reconstructed Z boson m_{Z2} . All variables are rescaled to zero mean and unit variance.

We use Gaussian kernels with bandwidths between 0.01 and 0.5.

3. Calibrated classifiers (CARL)

Idea.—CARL was developed in Ref. [30]. The authors showed that the likelihood ratio is invariant under any transformation that is monotonic with the likelihood ratio. In practice, this means that we can train a

classifier between two samples generated from the numerator and denominator hypotheses and turn the classifier decision function $\hat{s}(x)$ into an estimator for the likelihood ratio $\hat{r}(x)$. This relation between $\hat{s}(x)$ and $\hat{r}(x)$ can follow the ideal relation in Eqs. (18) and (46). But even if this relation does not hold, we can still extract a likelihood ratio estimator from the classifier output through probability calibration.

Requirements.—CARL can be used in the general likelihood-free setting: it only requires a simulator that can generate samples $\{x\} \sim p(x|\theta)$.

Structure.—CARL can be used either point by point, in an agnostic parametrized version, or [if the morphing condition in Eq. (6) holds] in a morphing-aware version. Figure 8 illustrates the structure of the estimator in these three cases.

Training.—A classifier with decision function $\hat{s}(x|\theta_0, \theta_1)$ is trained to discriminate between numerator (label $y = 0$) and denominator (label $y = 1$) samples by minimizing the binary cross entropy given in Eq. (16) (other loss functions are possible, but we have not experimented with them).

In the point-by-point version, the inputs to the classifiers are just the observables x , and the events in the numerator sample are generated according to one specific value θ_0 . In the parametrized versions of the estimator, the numerator training samples do not come from a single parameter θ_0 , but rather a combination of many different subsamples. In the agnostic parametrized setup, the value of θ_0 used in each event is then one of the inputs to the neural network. In the morphing-aware versions, it is used to calculate the weights $w_c(\theta_0)$ that multiply the different component networks $\hat{r}_c(x)$, as visualized in the bottom panel of Fig. 8.

Calibration.—In a next step, the classifier output is optionally calibrated as discussed in Sec. III D 1 using isotonic regression. The calibration curve is shown in the left panel of Fig. 21. We have also experimented with an additional step of expectation calibration, see Sec. III D 2.

Evaluation.—For a given x (and in the parametrized versions θ_0), the classifier decision function $\hat{s}(x|\theta_0, \theta_1)$ is evaluated. This is turned into a likelihood ratio estimator with the relation given in Eq. (18), and optionally calibrated.

Parameters.—The parameters of this approach are the architecture of the neural network, i.e., the number of layers and elements, the activation function, the optimizer used for training, its parameters, and optionally regularization terms.

For our example process we consider fully connected neural networks with two (“shallow”), three, or five (“deep”) layers of 100 neurons each and tanh activation functions. They are trained with the ADAM optimizer [79] over 50 epochs with early stopping and learning rate decay. Our default settings are given in Table VI. Experiments with different architectures, other activation functions, additional dropout layers, other optimizers, and different learning rate schedules yielded a worse performance.

4. Ratio regression (ROLR)

Idea.—Particle physics generators do not only provide sets of observables $\{x_e\}$, but also the corresponding parton-level momenta $\{z_e\}$. From matrix element codes such as MADGRAPH we can extract the squared matrix element $|\mathcal{M}|^2(z|\theta)$ given parton-level momenta z and theory parameter points θ . This allows us to calculate the joint likelihood ratio




FIG. 21. Calibration curves for different estimators, comparing the uncalibrated (raw) estimator to the estimator after probability calibration. The calibration curve for the truth prediction is a cross-check for consistency; we do not actually use calibration for the truth predictions. For the local score regression technique, we show the value of $\hat{h}(x|\theta_0, \theta_1)$ (essentially the log likelihood ratio in the local model) versus the estimated likelihood ratio after density estimation.

TABLE VI. Default settings for the analysis techniques. The neural network (NN) layers each have 100 units with tanh activation functions. The hyperparameter α multiplies the score squared error in the combined loss functions of Eqs. (35) and (37).

Strategy	NN layers	α	Calibration/density estimation
Histogram			Histogram
AFC			Gaussian KDE
CARL (PbP)	3		Isotonic probability calibration
CARL (parametrized)	2		Isotonic probability calibration
CARL (morphing aware)	2		Isotonic probability calibration
SALLY	5		Histogram
SALLINO	5		Histogram
ROLR (PbP)	3		Isotonic probability calibration
ROLR (parametrized)	3		Isotonic probability calibration
ROLR (morphing aware)	2		Isotonic probability calibration
CASCAL (parametrized)	5	5	Isotonic probability calibration
CASCAL (morphing aware)	2	5	Isotonic probability calibration
RASCAL (parametrized)	5	100	Isotonic probability calibration
RASCAL (morphing aware)	2	100	Isotonic probability calibration

$$r(x_e, z_e | \theta_0, \theta_1) = \frac{p(z_e | \theta_0)}{p(z_e | \theta_1)} = \frac{|\mathcal{M}|^2(z_e | \theta_0) \sigma(\theta_1)}{|\mathcal{M}|^2(z_e | \theta_1) \sigma(\theta_0)} \quad (\text{B1})$$

for any of the generated events.

In Sec. III B we have shown that regressing a function $\hat{r}(x)$ on the generated events $\{x_e\}$ and the corresponding joint likelihood ratios $r(x_e, z_e | \theta_0, \theta_1)$ will converge to

$$\hat{r}(x) \rightarrow r(x) = \frac{p(x | \theta_0)}{p(x | \theta_1)}, \quad (\text{B2})$$

provided that the events are sampled according to $x_e \sim p(x | \theta_1)$.

Requirements.—The ROLR technique requires a generator with access to the joint likelihood ratios $r(x_e, z_e | \theta_0, \theta_1)$. In the particle physics case, this means we have to be able to evaluate the squared matrix elements for given phase-space points and theory parameters.

Structure.—ROLR can be used either point by point, in an agnostic parametrized version, or [if the morphing condition in Eq. (6) holds] in a morphing-aware version. Figure 8 illustrates the structure of the estimator in these three cases.

Training.—The training phase is straightforward regression. It consists of minimizing the squared error loss between a flexible function $\hat{r}(x | \theta_0, \theta_1)$ (for instance a neural network) and the training data $\{x_e, r(x_e, z_e | \theta_0, \theta_1)\}$, which was generated according to θ_1 .

In the point-by-point version, the input to the regressor is just the observables x , and the ratio is between two fixed hypotheses θ_0 and θ_1 . In the

parametrized versions of the estimator, the ratios are based on various values θ_0 , while θ_1 is still kept fixed. In the agnostic parametrized setup, the value of θ_0 used in each event is then one of the inputs to the neural network. In the morphing-aware versions, it is used to calculate the weights $w_c(\theta_0)$ that multiply the different component networks $\hat{r}_c(x)$, as visualized in the bottom panel of Fig. 8.

In all cases we can slightly improve the structure by adding samples generated according to θ_0 to the training samples, regressing on $1/r$ instead of r on these events. The full loss functional is given in Eq. (33).

Calibration.—In a next step, the classifier output is optionally calibrated as discussed in Sec. III D 1 using isotonic regression. The calibration curve is shown in the middle panel of Fig. 21. We have also experimented with an additional step of expectation calibration, see Sec. III D 2.

Evaluation.—For a given x (and in the parametrized versions θ_0), the regressor $\hat{r}(x | \theta_0, \theta_1)$ is evaluated. The result is optionally calibrated.

Parameters.—The parameters of this approach are the architecture of the neural network, i.e., the number of layers and elements, the activation function, the optimizer used for training, its parameters, and optionally regularization terms.

For our example process we consider fully connected neural networks with two (shallow), three, or five (deep) layers of 100 neurons each and tanh activation functions. They are trained with the ADAM optimizer [79] over 50 epochs with early stopping and learning rate decay. Our default settings are given in Table VI. Experiments with different architectures,

other activation functions, additional dropout layers, other optimizers, and different learning rate schedules yielded a worse performance.

5. CARL + score regression (CASCAL)

Idea.—The parametrized CARL technique learns the full statistical model $\hat{r}(x|\theta_0, \theta_1)$, including the dependency on θ_0 . If it is realized as a differentiable classifier (such as a neural network), we can calculate the gradient of $\hat{r}(x|\theta_0, \theta_1)$ with respect to θ_0 , and thus the estimated score of this model. If the estimator is perfect, we expect this estimated score to minimize the squared error with respect to the joint score data available for the training data. This is based on the same argument as the local score regression technique; see Sec. III B for the proof.

We can turn this argument around and use the available score information during the training. To this end, we combine two terms in a combined loss function: the CARL-style cross entropy and the squared error between the estimated score and the joint score of the training data. These two pieces contain complementary information: the CARL part contains the information of the likelihood ratio for a fixed hypothesis comparison (θ_0, θ_1) , while the score part describes the relative change of the likelihood ratio under changes in θ_0 .

Requirements.—The CASCAL technique requires a generator with access to the joint score $t(x_e, z_e|\theta_0)$. In the particle physics case, this means we have to be able to evaluate the squared matrix elements for given phase-space points and theory parameters.

Structure.—Since the CASCAL method relies on the extraction of the estimated score from the estimator, it can only be used for parametrized estimators, either in an agnostic or morphing-aware version. The middle and bottom panels of Fig. 8 illustrate the structure of the estimator in these two cases.

Training.—A differentiable classifier with decision function $\hat{s}(x|\theta_0, \theta_1)$ is trained to discriminate between numerator (label $y = 0$) and denominator (label $y = 1$) samples, while the derived estimated score $\hat{r}(x|\theta_0)$ is compared to the joint score on the training samples generated from $y = 0$. The loss function that is minimized is thus a combination of the CARL-style cross entropy and the squared error on the score, weighted by a hyperparameter α . It is given in Eq. (35).

The numerator ($y = 0$) training samples do not come from a single parameter θ_0 , but rather a combination of many different subsamples. In the agnostic parametrized setup, the value of θ_0 used in each event is then one of the inputs to the neural network. In the morphing-aware versions, it is used to calculate the weights $w_c(\theta_0)$ that multiply the different

component networks $\hat{r}_c(x)$, as visualized in the bottom panel of Fig. 8.

Calibration.—In a next step, the classifier output is optionally calibrated as discussed in Sec. III D 1 using isotonic regression. The calibration curve is shown in the right panel of Fig. 21. We have also experimented with an additional step of expectation calibration, see Sec. III D 2.

Evaluation.—For a given x and θ_0 , the classifier decision function $\hat{s}(x|\theta_0, \theta_1)$ is evaluated. This is turned into a likelihood ratio estimator with the relation given in Eq. (18), and optionally calibrated.

Parameters.—The key hyperparameter of this technique is the factor α that weights the two terms in the loss function. Additional parameters set the architecture of the neural network, i.e., the number of layers and elements, the activation function, the optimizer used for training, its parameters, and optionally regularization terms.

For our example process we consider fully connected neural networks with two (shallow), three, or five (deep) layers of 100 neurons each and tanh activation functions. They are trained with the ADAM optimizer [79] over 50 epochs with early stopping and learning rate decay. Our default settings are given in Table VI. Experiments with different architectures, other activation functions, additional dropout layers, other optimizers, and different learning rate schedules yielded a worse performance.

6. Ratio + score regression (RASCAL)

Idea.—The parametrized ROLR technique learns the full statistical model $\hat{r}(x|\theta_0, \theta_1)$, including the dependency on θ_0 . If it is realized as a differentiable regressor, we can calculate the gradient of $\hat{r}(x|\theta_0, \theta_1)$ with respect to θ_0 , and thus the score of this model. If the estimator is perfect, we expect this estimated score to minimize the squared error with respect to the joint score data available for the training data.

We can turn this argument around and use the available likelihood ratio and score information during the training. To this end, we combine two terms in a combined loss function: the squared errors on the ratio and the score. These two pieces contain complementary information: the ratio regression part contains the information of the likelihood ratio for a fixed hypothesis comparison (θ_0, θ_1) , while the score part describes the relative change of the likelihood ratio under changes in θ_0 .

Requirements.—The RASCAL technique requires a generator with access to the joint likelihood ratio $r(x_e, z_e|\theta_0, \theta_1)$ and score $t(x_e, z_e|\theta_0)$. In the particle physics case, this means we have to be able to evaluate the squared matrix elements for given phase-space points and theory parameters.

Structure.—Since the RASCAL method relies on the extraction of the estimated score from the estimator, it can only be used for parametrized estimators, either in an agnostic or morphing-aware version. The middle and bottom panels of Fig. 8 illustrate the structure of the estimator in these two cases.

Training.—An estimator $\hat{r}(x|\theta_0, \theta_1)$ is trained through regression on the joint likelihood ratio, while the derived estimated score $\hat{t}(x|\theta_0)$ is compared to the joint score on the training samples generated from $y = 0$. The loss function that is minimized is thus a combination of the squared error on the ratio and the squared error on the score, weighted by a hyperparameter α . It is given in Eq. (35).

The numerator ($y = 0$) training samples do not come from a single parameter θ_0 , but rather a combination of many different subsamples. In the agnostic parametrized setup, the value of θ_0 used in each event is then one of the inputs to the neural network. In the morphing-aware versions, it is used to calculate the weights $w_c(\theta_0)$ that multiply the different component networks $\hat{r}_c(x)$, as visualized in the bottom panel of Fig. 8.

Calibration.—In a next step, the classifier output is optionally calibrated as discussed in Sec. III D 1 using isotonic regression. The calibration curve is shown in the right panel of Fig. 21. We have also experimented with an additional step of expectation calibration, see Sec. III D 2.

Evaluation.—For a given x and θ_0 , the estimator $\hat{r}(x|\theta_0, \theta_1)$ is evaluated and optionally calibrated.

Parameters.—The key hyperparameter of this technique is the factor α that weighs the two terms in the loss function. Additional parameters set the architecture of the neural network, i.e., the number of layers and elements, the activation function, the optimizer used for training, its parameters, and optionally regularization terms.

For our example process we consider fully connected neural networks with two (shallow), three, or five (deep) layers of 100 neurons each and tanh activation functions. They are trained with the ADAM optimizer [79] over 50 epochs with early stopping and learning rate decay. Our default settings are given in Table VI. Experiments with different architectures, other activation functions, additional dropout layers, other optimizers, and different learning rate schedules yielded a worse performance.

7. Local score regression and density estimation (SALLY)

Idea.—In Sec. III A 2 we introduced the score, the relative gradient of the likelihood with respect to the theory parameters. The score evaluated at some reference parameter point is the sufficient statistics of

the local approximation of the likelihood given in Eq. (15). In other words, we expect the score vector to be a set of “optimal observables” that includes all the information on the theory parameters, at least in the vicinity of the reference parameter point. If we can estimate the score from an observation, we can use it like any other set of observables. In particular, we can fill histograms of the score for any parameter point and thus estimate the likelihood ratio in score space.

To estimate the score, we again make use of the particle physics structure. Particle physics generators do not only provide sets of observables $\{x_e\}$, but also the corresponding parton-level momenta $\{z_e\}$. From matrix element codes such as MADGRAPH we can extract the squared matrix element $|\mathcal{M}|^2(z|\theta)$ given parton-level momenta z and theory parameter points θ . This allows us to calculate the joint score

$$\begin{aligned} t(x_e, z_e|\theta_0) &= \nabla_\theta \log p(z_e|\theta)|_{\theta_0} \\ &= \frac{\nabla_\theta |\mathcal{M}|^2(z_e|\theta_0)}{|\mathcal{M}|^2(z_e|\theta_0)} - \frac{\nabla_\theta \sigma(\theta_0)}{\sigma(\theta_0)} \end{aligned} \quad (\text{B3})$$

for any of the generated events. The derivatives in Eq. (B3) can always be evaluated numerically. If the process has the morphing structure of Eq. (6), one can alternatively calculate it from the morphing weights.

In Sec. III B we have shown that regressing a function $\hat{t}(x)$ on the generated events $\{x_e\}$ and the corresponding joint scores $t(x_e, z_e|\theta)$ will converge to

$$\hat{t}(x) \rightarrow \mathbf{t}(x) = \nabla_\theta \log p(x|\theta)|_{\theta_0}, \quad (\text{B4})$$

provided that the events are sampled according to $x_e \sim p(x|\theta_0)$.

This technique is illustrated in the middle panels of Fig. 20. The middle panel of Fig. 21 shows the relation between the scalar product of estimated score and $\theta_0 - \theta_1$ and the estimated likelihood ratio.

Requirements.—The SALLY technique requires a generator with access to the joint score $t(x_e, z_e|\theta_0)$. In the particle physics case, this means we have to be able to evaluate the squared matrix elements for given phase-space points and theory parameters.

Structure.—The technique consists of two separate steps: the score regression and the density estimation in the estimated score space. The score regression step is independent of the tested hypothesis and realized as a simple fully connected neural network. The subsequent density estimation is realized through multi-dimensional histograms, point by point in parameter space [if the morphing condition in Eq. (6) holds, a morphing-aware version is also possible].

Training.—The first part of the training is regression on the score (evaluated at some reference hypothesis). It consists of minimizing the squared error loss between

a flexible vector-valued function $\hat{t}(x|\theta_{\text{score}})$ (implemented for instance as a neural network) and the training data $\{x_e, t(x_e, z_e|\theta_{\text{score}})\}$, which was sampled according to θ_{score} .

The second step is density estimation in the estimated score space. We only consider histograms, but other density estimation techniques are also possible. For each value of θ_0 or θ_1 that is tested, we generate samples of events, estimate the corresponding score vectors, and fill a multidimensional histogram of the estimated score.

Calibration.—The density estimation step already calibrates the results, so we do not experiment with an additional calibration step.

Evaluation.—For a given observation x , the score regressor $\hat{t}(x)$ is evaluated. For each tested (θ_0, θ_1) pair, we then extract the corresponding bin contents from the numerator and denominator histograms, and calculate the estimated likelihood ratio with Eq. (38).

Parameters.—Both the score regression part and the subsequent density estimation have parameters. The first and most important choice is the reference hypothesis θ_{score} , at which the score is evaluated. For effective field theories the Standard Model is the natural choice, and we use it in our example process.

The score regression also depends on the hyperparameters of the neural network, i.e., the number of layers and elements, the activation function, the optimizer used for training, its parameters, and optionally regularization terms. For our example process we consider fully connected neural networks with two (shallow), three, or five (deep) layers of 100 neurons each and tanh activation functions. They are trained with the ADAM optimizer [79] over 50 epochs with early stopping and learning rate decay. Our default settings are given in Table VI. Experiments with different architectures, other activation functions, additional dropout layers, other optimizers, and different learning rate schedules yielded a worse performance.

The only parameter of the density estimation stage is the histogram binning. For our example process we consider two different variations:

- (i) Density estimation with a “fixed” binning, where the bin axes are aligned with the score components. We use 40 bins for each of the two score components.
- (ii) Density estimation with a “dynamic” binning, in which the bin axes are aligned with the $\theta_0 - \theta_1$ direction and the orthogonal one. We use 80 bins along the $\Delta\theta$ direction, which carries the relevant information in the local model approximation, and ten along the orthogonal vector.

For each pair (θ_0, θ_1) and each dimension, the bin edges are chosen such that the expected number of events according to θ_0 plus the expected number of events according to θ_1 is the same in each bin.

8. Local score regression, compression to scalar, and density estimation (SALLINO)

Idea.—In the proximity of the Standard Model (or any other reference parameter point), likelihood ratios only depend on the scalar product between the score and the difference between the numerator and denominator parameter points. If we can estimate the score from an observation, we can calculate this scalar product $\hat{h}(x|\theta_0, \theta_1)$, defined in Eq. (39), and use it like any other observable. In particular, we can fill histograms of \hat{h} for any parameter point and thus estimate the likelihood ratio in \hat{h} space.

To estimate the score, we once again exploit particle physics structure. Particle physics generators do not only provide sets of observables $\{x_e\}$, but also the corresponding parton-level momenta $\{z_e\}$. From matrix element codes such as MADGRAPH we can extract the squared matrix element $|\mathcal{M}|^2(z|\theta)$ given parton-level momenta z and theory parameter points θ . This allows us to calculate the joint score with Eq. (B3) for any of the generated events.

In Sec. III B we have shown that regressing a function $\hat{t}(x)$ on the generated events $\{x_e\}$ and the corresponding joint scores $t(x_e, z_e|\theta)$ will converge to $t(x)$, provided that the events are sampled according to $x_e \sim p(x|\theta_0)$.

Requirements.—The SALLINO technique requires a generator with access to the joint score $t(x_e, z_e|\theta_0)$. In the particle physics case, this means we have to be able to evaluate the squared matrix elements for given phase-space points and theory parameters.

Structure.—The technique consists of two separate steps: the score regression, and the density estimation in \hat{h} space. The score regression step is independent of the tested hypothesis and realized as a simple fully connected neural network. The subsequent density estimation in \hat{h} space is realized through one-dimensional histograms, point by point in parameter space [if the morphing condition in Eq. (6) holds, a morphing-aware version is also possible].

Training.—The first part of the training is regression on the score (evaluated at some reference hypothesis). It consists of minimizing the squared error loss between a flexible vector-valued function $\hat{t}(x|\theta_{\text{score}})$ (implemented for instance as a neural network) and the training data $\{x_e, t(x_e, z_e|\theta_{\text{score}})\}$, which was sampled according to θ_{score} .

The second step is density estimation in \hat{h} space. We only consider histograms, but other density estimation techniques are also possible. For each value of θ_0 or θ_1 that is tested, we generate samples of events, estimate the corresponding score vectors, calculate the scalar product in Eq. (39) to get $\hat{h}(x|\theta_0, \theta_1)$, and fill a one-dimensional histogram of this quantity.

Calibration.—The density estimation step already calibrates the results, so we do not experiment with an additional calibration step.

Evaluation.—For a given observation x , the score regressor $\hat{t}(x)$ is evaluated. For each tested (θ_0, θ_1) pair, we multiply it with $\theta_0 - \theta_1$ to get $\hat{h}(x|\theta_0, \theta_1)$, extract the corresponding bin contents from the numerator and denominator histograms, and calculate the estimated likelihood ratio with Eq. (40).

Parameters.—Both the score regression part and the subsequent density estimation have parameters. The first and most important choice is the reference hypothesis θ_{score} , at which the score is evaluated. For effective field theories the Standard Model is the natural choice, and we use it in our example process.

The score regression also depends on the hyperparameters of the neural network, i.e., the number of layers and elements, the activation function, the optimizer used for training, its parameters, and optionally regularization terms. For our example process we consider fully connected neural networks with two (shallow), three, or five (deep) layers of 100 neurons each and tanh activation functions. They are trained with the ADAM optimizer [79] over 50 epochs with early stopping and learning rate decay. Our default settings are given in Table VI. Experiments with different architectures, other activation functions, additional dropout layers, other optimizers, and different learning rate schedules yielded a worse performance.

The only parameter of the density estimation stage is the histogram binning. We use 100 bins. For each pair (θ_0, θ_1) and each dimension, the bin edges are chosen such that the same expected number of events according to θ_0 plus the expected number of events according to θ_1 is the same in each bin.

APPENDIX C: ADDITIONAL RESULTS

In Tables VII–XII we compare the performance of different versions of the likelihood ratio estimators. As metric we use the expected mean squared error on $\log r(x|\theta_0, \theta_1)$ as well as a trimmed version, as defined in Sec. VA. The estimators are an extended list of those given in Table IV, adding variations with different hyperparameter choices and the results for uncalibrated (raw) estimators. By default, we use neural networks with three hidden layers, the labels “shallow” and “deep” refer to two and five hidden layers, respectively. We highlight the versions of the estimators that were shown in the main part of this paper.

Because of the duality between density estimation and probabilistic classification [see Eqs. (18) and (46)], we can use all techniques to define classifiers. In Fig. 22 we show the ROC curves for two benchmark parameter points. Note how badly the two scenarios can be separated. This is not a shortcoming of the discrimination power of the classifiers, but due to the genuine overlap of the probability distributions, as can be seen from the identical ROC curve based on the true likelihood ratio.

TABLE VII. Comparison of techniques based on manually selected kinematic observables. The metrics shown are the expected mean squared error on the log likelihood ratio with and without trimming, as defined in the text. Checkmarks in the last two columns denote estimators included in Table IV and the figures in the main part of this paper, respectively.

Strategy	Setup	Expected MSE		Table IV	Figures
		All	Trimmed		
Histogram	$p_{T,j1}$	0.0879	0.0230	✓	
	$\Delta\phi_{jj}$	0.1595	0.0433	✓	
	2d (coarse binning)	0.0764	0.0117		
	2d (medium binning)	0.0630	0.0101		
	2d (fine binning)	0.0597	0.0115		
	2d (very fine binning)	0.0603	0.0153		
	2d (asymmetric binning)	0.0561	0.0106	✓	✓
AFC	2d, $\epsilon = 1$	0.1243	0.0257		
	2d, $\epsilon = 0.5$	0.0797	0.0144		
	2d, $\epsilon = 0.2$	0.0586	0.0091	✓	
	2d, $\epsilon = 0.1$	0.0732	0.0103		
	2d, $\epsilon = 0.05$	0.3961	0.0160		
	2d, $\epsilon = 0.02$	13.6816	0.0550		
	2d, $\epsilon = 0.01$	241.3264	0.2143		
	5d, $\epsilon = 1$	0.1252	0.0226		
	5d, $\epsilon = 0.5$	0.0779	0.0101	✓	
	5d, $\epsilon = 0.2$	0.0734	0.0128		
	5d, $\epsilon = 0.1$	0.9560	0.1833		
	5d, $\epsilon = 0.05$	38.1854	3.6658		
	5d, $\epsilon = 0.02$	2050.5289	57.0410		
	5d, $\epsilon = 0.01$	50024.7997	1668.8988		

TABLE VIII. Comparison of different versions of the CARL technique. The metrics shown are the expected mean squared error on the log likelihood ratio with and without trimming, as defined in the text. Checkmarks in the last two columns denote estimators included in Table IV and the figures in the main part of this paper, respectively.

Strategy	Setup	Expected MSE		Table IV	Figures
		All	Trimmed		
CARL (PbP, raw)	PbP	0.0409	0.0213		
CARL (PbP, calibrated)	PbP	0.0301	0.0111	Fig. 12	
CARL (parametrized, raw)	Baseline	0.0157	0.0040		
	Baseline, shallow	0.0134	0.0035		
	Baseline, deep	0.0161	0.0038		
	Random θ	0.0148	0.0037		
	Random θ , shallow	0.0130	0.0037		
	Random θ , deep	0.0164	0.0038		
CARL (parametrized, calibrated)	Baseline	0.0156	0.0032		
	Baseline, shallow	0.0124	0.0026	✓	✓
	Baseline, deep	0.0160	0.0029		
	Random θ	0.0147	0.0029		
	Random θ , shallow	0.0122	0.0028	✓	
	Random θ , deep	0.0155	0.0029		
CARL (morphing aware, raw)	Baseline	0.1598	0.0350		
	Baseline, shallow	0.1483	0.0331		
	Random θ	0.1743	0.0429		
	Random θ , shallow	0.1520	0.0369		
	Morphing basis, shallow	10.1231	7.9314		
CARL (morphing aware, calibrated)	Baseline	0.1036	0.0282		
	Baseline, shallow	0.0762	0.0200	✓	Fig. 12
	Random θ	0.1076	0.0289		
	Random θ , shallow	0.0858	0.0226	✓	
	Morphing basis, shallow	0.1564	0.0618	✓	

TABLE IX. Comparison of different versions of the ROLR technique. The metrics shown are the expected mean squared error on the log likelihood ratio with and without trimming, as defined in the text. Checkmarks in the last two columns denote estimators included in Table IV and the figures in the main part of this paper, respectively.

Strategy	Setup	Expected MSE		Table IV	Figures
		All	Trimmed		
ROLR (PbP, raw)	PbP	0.0052	0.0023		
ROLR (PbP, calibrated)	PbP	0.0049	0.0022	✓	
ROLR (parametrized, raw)	Baseline	0.0034	0.0019		
	Baseline, shallow	0.0069	0.0037		
	Baseline, deep	0.0041	0.0022		
	Random θ	0.0034	0.0017		
	Random θ , shallow	0.0070	0.0036		
	Random θ , deep	0.0036	0.0017		
ROLR (parametrized, calibrated)	Baseline	0.0032	0.0017	✓	✓
	Baseline, shallow	0.0059	0.0030		
	Baseline, deep	0.0038	0.0019		
	Random θ	0.0030	0.0014	✓	
	Random θ , shallow	0.0060	0.0030		
	Random θ , deep	0.0034	0.0015		

(Table continued)

TABLE IX. (*Continued*)

Strategy	Setup	Expected MSE		Table IV	Figures
		All	Trimmed		
ROLR (morphing aware, raw)	Baseline	0.2029	0.1449		
	Baseline, shallow	0.1672	0.1305		
	Random θ	0.1908	0.1353		
	Random θ , shallow	0.1160	0.0755		
	Morphing basis, shallow	5.6668	3.8335		
ROLR (morphing aware, calibrated)	Baseline	0.0328	0.0088		
	Baseline, shallow	0.0243	0.0063	✓	
	Random θ	0.0321	0.0089		
	Random θ , shallow	0.0224	0.0052	✓	
	Morphing basis, shallow	0.1300	0.0485	✓	

TABLE X. Comparison of different versions of the SALLY and SALLINO techniques. The metrics shown are the expected mean squared error on the log likelihood ratio with and without trimming, as defined in the text. Checkmarks in the last two columns denote estimators included in Table IV and the figures in the main part of this paper, respectively.

Strategy	Setup	Expected MSE		Table IV	Figures
		All	Trimmed		
SALLY	Fixed 2D histogram	0.0174	0.0005		
	Fixed 2D histogram, shallow	0.0170	0.0005		
	Fixed 2D histogram, deep	0.0171	0.0005		
	Dynamic 2D histogram	0.0132	0.0003		
	Dynamic 2D histogram, shallow	0.0133	0.0003		
	Dynamic 2D histogram, deep	0.0132	0.0002	✓	✓
SALLINO	1D histogram	0.0213	0.0006		
	1D histogram, shallow	0.0215	0.0007		
	1D histogram, deep	0.0213	0.0006	✓	

TABLE XI. Comparison of different versions of the CASCAL technique. The metrics shown are the expected mean squared error on the log likelihood ratio with and without trimming, as defined in the text. Checkmarks in the last two columns denote estimators included in Table IV and the figures in the main part of this paper, respectively.

Strategy	Setup	Expected MSE		Table IV	Figures
		All	Trimmed		
CASCAL (parametrized, raw)	Baseline, $\alpha = 5$	0.0019	0.0004		
	Baseline, $\alpha = 5$, shallow	0.0037	0.0004		
	Baseline, $\alpha = 5$, deep	0.0010	0.0003		
	Baseline, $\alpha = 0.5$, deep	0.0017	0.0006		
	Baseline, $\alpha = 1$, deep	0.0014	0.0005		
	Baseline, $\alpha = 2$, deep	0.0017	0.0008		
	Baseline, $\alpha = 10$, deep	0.0013	0.0004		
	Baseline, $\alpha = 20$, deep	0.0016	0.0004		
	Baseline, $\alpha = 50$, deep	0.0024	0.0007		
	Random θ	0.0022	0.0006		
	Random θ , shallow	0.0038	0.0005		
	Random θ , deep	0.0010	0.0003		

(Table continued)

TABLE XI. (*Continued*)

Strategy	Setup	Expected MSE		Table IV	Figures
		All	Trimmed		
CASCAL (parametrized, calibrated)	Baseline, $\alpha = 5$	0.0012	0.0002		
	Baseline, $\alpha = 5$, shallow	0.0025	0.0003		
	Baseline, $\alpha = 5$, deep	0.0008	0.0002	✓	✓
	Baseline, $\alpha = 0.5$, deep	0.0013	0.0003		
	Baseline, $\alpha = 1$, deep	0.0011	0.0003		
	Baseline, $\alpha = 2$, deep	0.0010	0.0002		
	Baseline, $\alpha = 10$, deep	0.0010	0.0003		
	Baseline, $\alpha = 20$, deep	0.0011	0.0003		
	Baseline, $\alpha = 50$, deep	0.0016	0.0005		
	Random θ , $\alpha = 5$	0.0013	0.0003		
	Random θ , $\alpha = 5$, shallow	0.0027	0.0004		
	Random θ , $\alpha = 5$, deep	0.0009	0.0002	✓	
CASCAL (morphing aware, raw)	Baseline, $\alpha = 5$	0.1935	0.0810		
	Baseline, $\alpha = 5$, shallow	0.1870	0.0732		
	Random θ , $\alpha = 5$, shallow	0.1624	0.0643		
	Morphing basis, $\alpha = 5$, shallow	0.0707	0.0109		
CASCAL (morphing aware, calibrated)	Baseline, $\alpha = 5$	0.1408	0.0508		
	Baseline, $\alpha = 5$, shallow	0.1359	0.0427	✓	
	Random θ , $\alpha = 5$, shallow	0.0922	0.0268	✓	
	Morphing basis, $\alpha = 5$, shallow	0.0403	0.0081	✓	

TABLE XII. Comparison of different versions of the RASCAL technique. The metrics shown are the expected mean squared error on the log likelihood ratio with and without trimming, as defined in the text. Checkmarks in the last two columns denote estimators included in Table IV and the figures in the main part of this paper, respectively.

Strategy	Setup	Expected MSE		Table IV	Figures
		All	Trimmed		
RASCAL (parametrized, raw)	Baseline, $\alpha = 100$	0.0010	0.0003		
	Baseline, $\alpha = 100$, shallow	0.0025	0.0006		
	Baseline, $\alpha = 100$, deep	0.0009	0.0004		
	Baseline, $\alpha = 10$, deep	0.0011	0.0005		
	Baseline, $\alpha = 20$, deep	0.0009	0.0004		
	Baseline, $\alpha = 50$, deep	0.0009	0.0004		
	Baseline, $\alpha = 200$, deep	0.0009	0.0004		
	Baseline, $\alpha = 500$, deep	0.0011	0.0006		
	Baseline, $\alpha = 1000$, deep	0.0012	0.0007		
	Random θ	0.0011	0.0004		
	Random θ , shallow	0.0030	0.0010		
	Random θ , deep	0.0008	0.0004		

(Table continued)

TABLE XII. (Continued)

Strategy	Setup	Expected MSE		Table IV	Figures
		All	Trimmed		
RASCAL (parametrized, calibrated)	Baseline, $\alpha = 100$	0.0010	0.0003		
	Baseline, $\alpha = 100$, shallow	0.0021	0.0005		
	Baseline, $\alpha = 100$, deep	0.0009	0.0004	✓	✓
	Baseline, $\alpha = 10$, deep	0.0010	0.0004		
	Baseline, $\alpha = 20$, deep	0.0009	0.0004		
	Baseline, $\alpha = 50$, deep	0.0009	0.0004		
	Baseline, $\alpha = 200$, deep	0.0008	0.0004		
	Baseline, $\alpha = 500$, deep	0.0009	0.0005		
	Baseline, $\alpha = 1000$, deep	0.0012	0.0006		
	Random θ , $\alpha = 100$	0.0010	0.0004		
	Random θ , $\alpha = 100$, shallow	0.0025	0.0008		
	Random θ , $\alpha = 100$, deep	0.0008	0.0004	✓	
RASCAL (morphing aware, raw)	Baseline, $\alpha = 100$	0.2880	0.2024		
	Baseline, $\alpha = 100$, shallow	0.3569	0.2861		
	Random θ , $\alpha = 100$	0.2705	0.1825		
	Random θ , $\alpha = 100$, shallow	0.3243	0.2488		
	Morphing basis, $\alpha = 100$, shallow	0.1909	0.1673		
RASCAL (morphing aware, calibrated)	Baseline, $\alpha = 100$	0.1530	0.0673		
	Baseline, $\alpha = 100$, shallow	0.1250	0.0514	✓	
	Random θ , $\alpha = 100$	0.1358	0.0627		
	Random θ , $\alpha = 100$, shallow	0.1316	0.0539	✓	
	Morphing basis, shallow, $\alpha = 100$	0.0307	0.0072	✓	

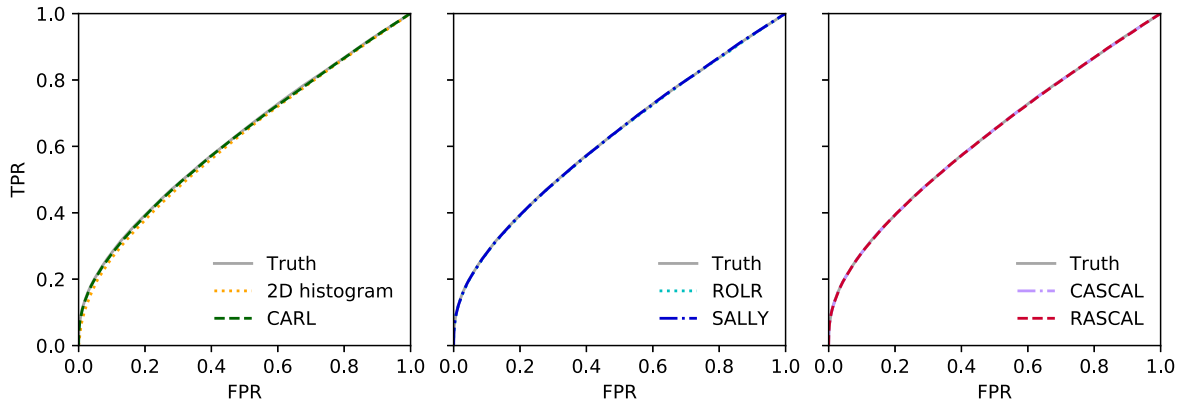


FIG. 22. Receiver operating characteristic (ROC) curves of true positive rates (TPR) vs false positive rates (FPR) for the classification between the benchmark scenarios $\theta_0 = (-0.5, -0.5)^T$ and θ_1 as in Eq. (47). The ROC Area under curve (AUC) based on the true likelihood is 0.6276. The results show how much the probability distributions for these two hypotheses overlap.

- [1] T. Sjostrand, S. Mrenna, and P. Z. Skands, *Comput. Phys. Commun.* **178**, 852 (2008).
- [2] S. Agostinelli *et al.* (GEANT4 Collaboration), *Nucl. Instrum. Methods Phys. Res., Sect. A* **506**, 250 (2003).
- [3] K. S. Cranmer, *Comput. Phys. Commun.* **136**, 198 (2001); K. Cranmer, G. Lewis, L. Moneta, A. Shibata, and W. Verkerke (ROOT Collaboration), Report No. CERN-OPEN-2012-016, 2012; M. Frate, K. Cranmer, S. Kalia, A. Vandenberg-Rodes, and D. Whiteson, [arXiv:1709.05681](#).
- [4] J. Brehmer, K. Cranmer, F. Kling, and T. Plehn, *Phys. Rev. D* **95**, 073002 (2017).
- [5] K. Kondo, *J. Phys. Soc. Jpn.* **57**, 4126 (1988).
- [6] V. M. Abazov *et al.* (D0 Collaboration), *Nature (London)* **429**, 638 (2004).
- [7] P. Artoisenet and O. Mattelaer, in *Proceedings, 2nd International Workshop on Prospects for Charged Higgs Discovery at Colliders (CHARGED 2008): Uppsala, Sweden, 2008* [*Proc. Sci.* CHARGED2008 (2008) 025].
- [8] Y. Gao, A. V. Gritsan, Z. Guo, K. Melnikov, M. Schulze, and N. V. Tran, *Phys. Rev. D* **81**, 075022 (2010).
- [9] J. Alwall, A. Freitas, and O. Mattelaer, *Phys. Rev. D* **83**, 074010 (2011).
- [10] S. Bolognesi, Y. Gao, A. V. Gritsan, K. Melnikov, M. Schulze, N. V. Tran, and A. Whitbeck, *Phys. Rev. D* **86**, 095031 (2012).
- [11] P. Avery *et al.*, *Phys. Rev. D* **87**, 055006 (2013).
- [12] J. R. Andersen, C. Englert, and M. Spannowsky, *Phys. Rev. D* **87**, 015019 (2013).
- [13] J. M. Campbell, R. K. Ellis, W. T. Giele, and C. Williams, *Phys. Rev. D* **87**, 073005 (2013).
- [14] P. Artoisenet, P. de Aquino, F. Maltoni, and O. Mattelaer, *Phys. Rev. Lett.* **111**, 091802 (2013).
- [15] J. S. Gainer, J. Lykken, K. T. Matchev, S. Mrenna, and M. Park, in *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, 2013*, [arXiv:1307.3546](#).
- [16] D. Schouten, A. DeAbreu, and B. Stelzer, *Comput. Phys. Commun.* **192**, 54 (2015).
- [17] T. Martini and P. Uwer, *J. High Energy Phys.* **09** (2015) 083.
- [18] A. V. Gritsan, R. Röntsch, M. Schulze, and M. Xiao, *Phys. Rev. D* **94**, 055023 (2016).
- [19] T. Martini and P. Uwer, *J. High Energy Phys.* **05** (2018) 141.
- [20] D. Atwood and A. Soni, *Phys. Rev. D* **45**, 2405 (1992).
- [21] M. Davier, L. Duflot, F. Le Diberder, and A. Rouge, *Phys. Lett. B* **306**, 411 (1993).
- [22] M. Diehl and O. Nachtmann, *Z. Phys. C* **62**, 397 (1994).
- [23] D. E. Soper and M. Spannowsky, *Phys. Rev. D* **84**, 074002 (2011).
- [24] D. E. Soper and M. Spannowsky, *Phys. Rev. D* **87**, 054012 (2013).
- [25] D. E. Soper and M. Spannowsky, *Phys. Rev. D* **89**, 094005 (2014).
- [26] C. Englert, O. Mattelaer, and M. Spannowsky, *Phys. Lett. B* **756**, 103 (2016).
- [27] S. A. Sisson and Y. Fan, *Likelihood-free MCMC* (Chapman and Hall/CRC, New York, 2011), p. 839.
- [28] J.-M. Marin, P. Pudlo, C. P. Robert, and R. J. Ryder, *Stat. Comput.*, **22**, 1167 (2012).
- [29] T. Charnock, G. Lavaux, and B. D. Wandelt, *Phys. Rev. D* **97**, 083004 (2018).
- [30] K. Cranmer, J. Pavez, and G. Louppe, [arXiv:1506.02169](#).
- [31] K. Cranmer and G. Louppe, J. Brief Ideas (2016).
- [32] Y. Fan, D. J. Nott, and S. A. Sisson, [arXiv:1212.1479](#).
- [33] G. Papamakarios and I. Murray, in *Advances in Neural Information Processing Systems 29*, edited by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Curran Associates, Inc., 2016), pp. 1028–1036.
- [34] B. Paige and F. Wood, [arXiv:1602.06701](#).
- [35] R. Dutta, J. Corander, S. Kaski, and M. U. Gutmann, [arXiv:1611.10242](#).
- [36] M. U. Gutmann, R. Dutta, S. Kaski, and J. Corander, *Stat. Comput.*, **28**, 411 (2018).
- [37] D. Tran, R. Ranganath, and D. M. Blei, [arXiv:1702.08896](#).
- [38] G. Louppe and K. Cranmer, [arXiv:1707.07113](#).
- [39] L. Dinh, D. Krueger, and Y. Bengio, [arXiv:1410.8516](#).
- [40] D. J. Rezende and S. Mohamed, [arXiv:1505.05770](#).
- [41] L. Dinh, J. Sohl-Dickstein, and S. Bengio, [arXiv:1605.08803](#).
- [42] G. Papamakarios, T. Pavlakou, and I. Murray, [arXiv:1705.07057](#).
- [43] B. Uria, M.-A. Côté, K. Gregor, I. Murray, and H. Larochelle, [arXiv:1605.02226](#).
- [44] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, [arXiv:1609.03499](#).
- [45] A. van den Oord, N. Kalchbrenner, O. Vinyals, L. Espeholt, A. Graves, and K. Kavukcuoglu, [arXiv:1606.05328](#).
- [46] A. van den Oord, N. Kalchbrenner, and K. Kavukcuoglu, [arXiv:1601.06759](#).
- [47] G. Papamakarios, D. C. Sterratt, and I. Murray, [arXiv:1805.07226](#).
- [48] J. Brehmer, K. Cranmer, G. Louppe, and J. Pavez, companion Letter, *Phys. Rev. Lett.* **121**, 111801 (2018).
- [49] J. Brehmer, G. Louppe, J. Pavez, and K. Cranmer, [arXiv:1805.12244](#).
- [50] J. Brehmer, K. Cranmer, G. Louppe, and J. Pavez, http://github.com/johannbrehmer/higgs_inference.
- [51] S. R. Coleman, J. Wess, and B. Zumino, *Phys. Rev.* **177**, 2239 (1969).
- [52] C. G. Callan, Jr., S. R. Coleman, J. Wess, and B. Zumino, *Phys. Rev.* **177**, 2247 (1969).
- [53] S. Weinberg, *Phys. Lett. B* **91**, 51 (1980).
- [54] C. J. C. Burges and H. J. Schnitzer, *Nucl. Phys.* **B228**, 464 (1983).
- [55] C. N. Leung, S. T. Love, and S. Rao, *Z. Phys. C* **31**, 433 (1986).
- [56] W. Buchmuller and D. Wyler, *Nucl. Phys.* **B268**, 621 (1986).
- [57] C. Arzt, M. B. Einhorn, and J. Wudka, *Nucl. Phys.* **B433**, 41 (1995).
- [58] K. Hagiwara, S. Ishihara, R. Szalapski, and D. Zeppenfeld, *Phys. Rev. D* **48**, 2182 (1993).
- [59] B. Grzadkowski, M. Iskrzynski, M. Misiak, and J. Rosiek, *J. High Energy Phys.* **10** (2010) 085.
- [60] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro, *J. High Energy Phys.* **07** (2014) 079.
- [61] B. Henning, X. Lu, and H. Murayama, *J. High Energy Phys.* **01** (2016) 023.

- [62] D. de Florian *et al.* (LHC Higgs Cross Section Working Group), [arXiv:1610.07922](#).
- [63] J. Brehmer, A. Freitas, D. Lopez-Val, and T. Plehn, *Phys. Rev. D* **93**, 075014 (2016).
- [64] J. Brehmer, F. Kling, T. Plehn, and T. M. P. Tait, *Phys. Rev. D* **97**, 095017 (2018).
- [65] J. de Favereau, C. Delaere, P. Demin, A. Giammanco, V. Lemaître, A. Mertens, and M. Selvaggi (DELPHES 3 Collaboration), *J. High Energy Phys.* **02** (2014) 057.
- [66] G. Aad *et al.* (ATLAS Collaboration), Report No. ATL-PHYS-PUB-2015-047, 2015.
- [67] J. R. Dell’Aquila and C. A. Nelson, *Phys. Rev. D* **33**, 80 (1986).
- [68] T. Plehn, D. L. Rainwater, and D. Zeppenfeld, *Phys. Rev. Lett.* **88**, 051801 (2002).
- [69] V. Hankele, G. Klamke, D. Zeppenfeld, and T. Figy, *Phys. Rev. D* **74**, 095001 (2006).
- [70] K. Hagiwara, Q. Li, and K. Mawatari, *J. High Energy Phys.* **07** (2009) 101.
- [71] C. Englert, D. Goncalves-Netto, K. Mawatari, and T. Plehn, *J. High Energy Phys.* **01** (2013) 148.
- [72] K. Cranmer and T. Plehn, *Eur. Phys. J. C* **51**, 415 (2007).
- [73] T. Plehn, P. Schichtel, and D. Wiegand, *Phys. Rev. D* **89**, 054002 (2014).
- [74] F. Kling, T. Plehn, and P. Schichtel, *Phys. Rev. D* **95**, 035026 (2017).
- [75] P. Baldi, K. Cranmer, T. Faucett, P. Sadowski, and D. Whiteson, *Eur. Phys. J. C* **76**, 235 (2016).
- [76] K. Cranmer, J. Pavez, G. Louppe, and W. K. Brooks, in *Proceedings, 17th International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT 2016): Valparaiso, Chile, 2016* [*J. Phys. Conf. Ser.* **762**, 012034 (2016)].
- [77] J. Alsing, B. Wandelt, and S. Feeney, [arXiv:1801.01497](#).
- [78] A. Hyvärinen, *J. Machine Learning Res.* **6**, 695 (2005).
- [79] D. P. Kingma and J. Ba, [arXiv:1412.6980](#).
- [80] F. Chollet *et al.*, Keras, <https://github.com/keras-team/keras>.
- [81] M. Abadi *et al.*, [arXiv:1603.04467](#).
- [82] J. B. Kruskal, *Psychometrika* **29**, 115 (1964).
- [83] S. S. Wilks, *Ann. Math. Stat.* **9**, 60 (1938).
- [84] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, *Eur. Phys. J. C* **71**, 1 (2011); **73**, 2501(E) (2013).
- [85] A. Wald, *Trans. Am. Math. Soc.* **54**, 426 (1943).
- [86] G. Louppe, M. Kagan, and K. Cranmer, [arXiv:1611.01046](#).
- [87] A. Mertens, in *Proceedings, 15th International Workshop on Advanced Computing and Analysis Techniques in Physics Research (ACAT 2013)* (2014), Vol. 523, p. 012028.