

Frequentist uncertainties on neural density ratios with $w_i f_i$ ensembles

Sean Benevedes^{*} and Jesse Thaler[†]

Center for Theoretical Physics—a Leinweber Institute, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA and The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, Cambridge, Massachusetts, USA



(Received 7 June 2025; accepted 3 September 2025; published 19 September 2025)

We introduce $w_i f_i$ ensembles as a novel framework to obtain asymptotic frequentist uncertainties on density ratios, with a particular focus on neural ratio estimation in the context of high-energy physics. When the density ratio of interest is a likelihood ratio conditioned on parameters, $w_i f_i$ ensembles can be used to perform simulation-based inference on those parameters. After training the basis functions $f_i(x)$, uncertainties on the weights w_i can be straightforwardly propagated to the estimated parameters without requiring extraneous bootstraps. To demonstrate this approach, we present an application in quantum chromodynamics at the Large Hadron Collider, using $w_i f_i$ ensembles to estimate the likelihood ratio between generated quark and gluon jets. We use this learned likelihood ratio to estimate the quark fraction in a synthetic mixed quark/gluon sample, showing that the resultant uncertainties empirically satisfy the desired coverage properties.

DOI: [10.1103/w28w-x5wh](https://doi.org/10.1103/w28w-x5wh)

I. INTRODUCTION

Machine learning is used ubiquitously in high-energy physics (HEP) for a wide breadth of applications, including jet tagging [1–6], anomaly detection [7–14], supervised searches for new physics [15–18], unfolding [19–28], and generative modeling [29–37]; see Ref. [38] for a more comprehensive bibliography. At the heart of many of these use cases is density ratio estimation (DRE), which is the task of inferring the ratio of two probability densities given samples from each density but not the functional form of the densities themselves. A paradigmatic example of DRE in HEP is for classification tasks like jet tagging: according to the Neyman-Pearson lemma [39], the optimal binary classifier is given by any monotonic function of the density ratio of the two likelihoods. The ratio of two densities can also be used to reweight samples from one density to obtain samples from the other, as is relevant for detector unfolding (see e.g., [19]). The primary application we consider here is for parameter estimation, where DRE is the foundation for various techniques in simulation-based inference (SBI) [40].

In practice, DRE is performed with a finite amount of data and compute, so the resultant point estimate of the density ratio will not be exactly equal to the true value. This discrepancy should be quantified as an uncertainty and propagated to downstream applications; see Ref. [41] for case studies where parameter estimation is unreliable due to misspecification of the density ratio. To the best of our knowledge, there are no existing frequentist methods that rigorously assess uncertainties on neural density ratios directly. Rather, the state of the art for SBI in HEP is to perform the Neyman construction [42], where one treats the estimated density ratio as an arbitrary test statistic and uses bootstrapping to determine its empirical distribution. While the Neyman construction produces statistically rigorous uncertainties for downstream parameter inference, it is computationally expensive due to the requisite bootstrapping, and it is not applicable when one requires a notion of uncertainty on the density ratio itself rather than on a downstream test statistic.

In this paper, we introduce $w_i f_i$ ensembles, a framework for DRE that uses model ensembles to produce statistically rigorous point estimates of, and uncertainties on, density ratios. To estimate a density ratio $r(x)$ of two distributions $n(x)$ and $d(x)$,

$$r(x) \equiv \frac{n(x)}{d(x)}, \quad (1)$$

we propose the model,

^{*}Contact author: seanmb@mit.edu

[†]Contact author: jthaler@mit.edu

$$\log \tilde{r}(x|w) = \sum_i w_i f_i(x), \quad (2)$$

where $\tilde{r}(x|w)$ is our parametrized model of $r(x)$, the $f_i(x)$ are neural networks (NNs) used to model the density ratio, and the w_i are weights inferred from training data (not to be confused with NN weights). We use results on M -estimators [43] to deduce asymptotic coverage properties for the confidence intervals of the w_i (and by extension, for confidence intervals on the density ratio itself). We also show how these uncertainties on the density ratio can be propagated using the Gong-Samaniego theorem [44] to obtain downstream uncertainties on parameters in the context of SBI.

Our results rely on two assumptions, both of which amount to conditions on the training of the $f_i(x)$. The first assumption is that the model is *well-specified*: that there exists a set of weights w_i^* for which $\tilde{r}(x|w^*) = r(x)$. The second assumption is that we can treat the $f_i(x)$ as fixed for the purposes of fitting the w_i . This is necessary in principle for our formulas [which do not include correlations between the $f_i(x)$ and the w_i] to be applicable, although empirically we have not found any evidence for such correlations. In practice, this can be achieved by using separate datasets for the purposes of training the $f_i(x)$ and fitting the w_i .

Intuitively, the ensemble weights w_i play the role of nuisance parameters in a collider physics analysis. If we only use a point estimate of a density ratio, e.g., directly from conventional NN training, then there is a resultant unquantifiable mismodeling of the density ratio (a systematic, in the collider physics parlance). Introducing and fitting the weights w_i trades off this unquantifiable mismodeling for an additional, quantifiable source of uncertainty (i.e., a statistical uncertainty). Note that $w_i f_i$ ensembles cannot solve the general challenge of model selection, and our approach is only as reliable as the assumed model form in Eq. (2).

The distinction between mismodeling and uncertainty is a crucial point:

- (i) *Mismodeling* (also known as model misspecification) is a mismatch between the true data-generating process and a particular model of that process. In the case of a parametrized model, mismodeling occurs when there is no set of parameters that reproduce the data-generating distribution. Mismodeling is generally unquantifiable, and collecting more data does not help.
- (ii) *Uncertainty* is the remaining degree of ignorance after performing statistical inference, assuming the absence of mismodeling. Uncertainty is generally quantifiable and systematically improvable by collecting more data.

Factors such as limited NN expressivity as well as stochasticity of NN initialization and training therefore potentially contribute to mismodeling, but are not uncertainties in this

framing. The aim of $w_i f_i$ ensembles is to trade off the unquantifiable mismodeling due to limited training data and training stochasticity in favor of quantifiable uncertainties. They do this by introducing a parametrized model which is more flexible than each individual basis function $f_i(x)$ and then quantifying the uncertainties through the introduced weights w_i .

To demonstrate our approach, we focus on one particular use case for DRE: parameter estimation. Parameter estimation is a central task in HEP, and a tremendous amount of work has gone into precise determinations of parameters like the strong-coupling constant α_s [45–50] and the top quark mass m_t [51–53]. Traditionally, these parameters are measured by collecting a large amount of high-dimensional data (e.g., collider events), reducing the dimensionality of the data by projecting onto a small number of handcrafted observables, and binning to estimate the likelihood conditioned on the parameter(s) of interest. However, these processing steps throw away some of the information in the data. SBI, using DRE, provides a means of fully exploiting this information by estimating the (full, high-dimensional) likelihood ratio. Our case studies will demonstrate how $w_i f_i$ ensembles can streamline the SBI process by quantifying the uncertainty on inferred parameters through an uncertainty estimate on the likelihood ratio directly.

Existing methods that use ensembles of NNs for uncertainty quantification, like deep ensembles [54], combine the outputs of the NNs through an unweighted average. In our case, the weights w_i assigned to each of the models in a $w_i f_i$ ensemble are free parameters determined from the training data. This allows us to formalize the idea that not all members of an ensemble will be of equal utility, and so they should not contribute equally. Separate from uncertainty quantification, NN ensembles are widely used to improve model robustness [55]. When combined with the Neyman construction to quantify uncertainties, though, it can be computationally daunting to train a bootstrapped ensemble of NN ensembles [56]. In our case, because the weights w_i can be treated just like nuisance parameters, we can use asymptotic formulas to obtain uncertainties without extraneous bootstraps. Of course, $w_i f_i$ ensembles are still an ensemble, and we will advocate for bootstrapping as an effective way to construct that ensemble; the point to emphasize is that we only need *one* ensemble to achieve both robustness and uncertainty quantification at the same time.

The remainder of the paper is organized as follows. In Sec. II, we review DRE, its connection with SBI, and the necessity of uncertainties in the inferred density ratio. Then, in Sec. III, we introduce $w_i f_i$ ensembles in detail, explaining how the density ratio and its uncertainties are inferred from data, and how these uncertainties are propagated for parameter estimation. In Sec. IV, we show empirically that $w_i f_i$ ensembles successfully perform DRE with well-calibrated uncertainties in a Gaussian toy

example where the density ratio is known. We then move to an application in the context of quantum chromodynamics (QCD) in Sec. V, showing that $w_i f_i$ ensembles can be used to estimate the quark fraction in a synthetic mixed sample of generated quark and gluon jets. We conclude in Sec. VI by discussing future directions. Additional statistical details appear in the appendices: Appendix A discusses the properties of our w_i estimator, Appendix B describes an alternative pseudolikelihood method for producing confidence intervals, and Appendix C studies the leading bias of the w_i estimator in the asymptotic limit. Implementations of $w_i f_i$ ensembles, and the code used in our experiments, can be found at <https://github.com/benevedes/wifi-ensembles>.

II. REVIEW OF DENSITY RATIO ESTIMATION

In this section, we review DRE and its applications to frequentist parameter estimation. In Sec. II A, we introduce the DRE task. Then, in Sec. II B, we review how DRE can be used to estimate likelihood ratios, enabling SBI. In Sec. II C, we discuss the uncertainties and mismodelings that arise in DRE, and their effects in the context of SBI. Finally, in Sec. II D, we describe existing frequentist approaches for estimating these uncertainties.

A. Problem statement

Consider two probability distributions $n(x)$ and $d(x)$ of an arbitrary random variable x , which we call the numerator and denominator distributions. Suppose further that it is possible to obtain samples drawn from each of these distributions, perhaps through experiment or through simulation, but that the functional forms of the distributions themselves are not known. The DRE task is then, given a dataset D consisting of samples drawn from each of the distributions, to estimate the ratio,

$$r(x) \equiv \frac{n(x)}{d(x)}. \quad (3)$$

Throughout, we take samples to be independent and identically distributed (i.i.d.).

B. Parameter estimation

Now, consider the parameter estimation task. Suppose that an experiment is performed to measure some set of parameters Θ . The data consist of phase space variables Φ collected by the experiment that are sampled according to some probability distribution $p(\Phi|\Theta^*)$, where Θ^* denotes the set of ground truth values of the Θ . To perform frequentist inference on Θ , we could perform a maximum likelihood estimate (MLE) analysis if we knew the family of distributions $p(\Phi|\Theta)$ as a function of Θ . However, it is not immediately obvious how to obtain the MLE in the situation where $p(\Phi|\Theta)$ is unknown, as is often the case in

collider physics; it is intractable to directly calculate the probability density over collider events.

DRE provides a way forward in this situation, under the assumption that one can sample (i.e., through simulation) the joint distribution $p(\Phi, \Theta)$.¹ The ability to sample from this joint distribution immediately implies that one can also sample from the product distribution $p(\Phi)p(\Theta)$; each of the marginals $p(\Phi)$ and $p(\Theta)$ can be sampled by sampling the joint and throwing away the information about the marginalized variables, and the product distribution can then be sampled by pairing a sample from $p(\Phi)$ with a sample from $p(\Theta)$. Then, we can perform DRE with $x = \{\Phi, \Theta\}$ using $p(\Phi, \Theta)$ as the numerator distribution and $p(\Phi)p(\Theta)$ as the denominator distribution,

$$\ell(\Phi|\Theta) \equiv \frac{p(\Phi, \Theta)}{p(\Phi)p(\Theta)} = \frac{p(\Phi|\Theta)}{p(\Phi)}. \quad (4)$$

This is then a frequentist procedure for so-called likelihood-free or simulation-based inference (SBI); this specific class of methods is called neural ratio estimation [40] if NNs are used to parametrize the estimated density ratio.

As opposed to traditional statistical methods in collider physics, which employ dimensionality reduction to a set of hand-selected summary statistics and binning to estimate likelihoods, SBI allows one to perform inference using the full phase space information of the data. This, in principle, allows for more sensitive measurements and probes for new physics.

C. Uncertainties in density ratio estimation

When we consider the DRE procedure in more detail, however, important complications arise. Practically, one encounters various sources of *mismodeling*, meaning a mismatch between

- (i) The true (but unknown) data-generating distribution as a function of parameters of interest $p(\Phi|\Theta)$, and
- (ii) The proxy distributions that we can access either through sampling or through evaluation as a function of data.

We similarly refer to a density ratio as being mismodeled if there is a mismatch between the ratio of the true data-generating distributions and the available proxy ratio.

We then use the term *uncertainty* to refer to our remaining ignorance of the true parameters Θ^* in the absence of mismodeling. Under mild assumptions like finite variance, uncertainty tends to zero as the amount of i.i.d. data collected tends to infinity, while errors induced

¹Note that this joint distribution implies the existence of a prior $p(\Theta)$. The inference procedures that we consider are frequentist in the sense that they are independent of this prior so long as the prior has support over all realizable values of Θ . We assume this support property throughout.

by mismodeling do not.² Moreover, when we produce confidence intervals, these address uncertainties but not mismodeling; their coverage guarantees apply only in the absence of mismodeling. In HEP parlance, our usage of mismodeling corresponds to unquantified systematic uncertainties, and our usage of uncertainties corresponds to quantified statistical and systematic uncertainties.

The source of mismodeling that we aim to reduce in this paper (and replace with quantified uncertainties) arises from the DRE itself. In particular, we do not address any mismodeling that could arise if the training data are imperfectly sampled from n and d , i.e., mismodeling due to flawed simulation. We consider DRE procedures which model the density ratio with NNs; these networks are trained with loss functions which, asymptotically, admit the density ratio as the objective function. However, in practice, the NN only approximately learns this objective function. This is due to a variety of factors: the finite amount of data that the NN is trained on does not contain enough information to uniquely identify the ground truth density ratio, the training procedure does not perfectly exploit the information that is in the data, and a fixed, finite-width NN architecture is generically not expressive enough to capture the true density ratio.

The learned density ratio $\hat{\ell}(\Phi|\Theta)$ is then only approximately equal to $\ell(\Phi|\Theta)$, and the estimator for Θ obtained via maximization of $\hat{\ell}(\Phi|\Theta)$ is not the MLE that would have been obtained from the ground truth $\ell(\Phi|\Theta)$. Then, in general, this estimator does not enjoy the various nice properties of the MLE; in particular, it generally is not consistent [i.e., converge in probability to Θ^* as the number of samples drawn from $p(\Phi|\Theta^*)$ goes to infinity], and the calculation of its variance through asymptotic formulas is generically incorrect.³

²Later, we refer to NN training as a potential source of mismodeling. This mismodeling generically decreases as the amount of NN training data increases, which might at first glance seem contradictory with the previous statement. The resolution of this tension is that, in the contexts we describe, the training procedure is separate from the inference procedure. When we refer to mismodeling due to NN training, we have in mind mismodeling at inference time. At that point, the NNs should be regarded as fixed, so this source of mismodeling does not decrease as the size of the inference dataset increases.

³One can gain intuition for these effects in a toy example. Consider a normal distribution with unknown ground truth mean μ and variance σ^2 . Suppose a devious rival scientist tells you under the guise of collaboration that the true value of the mean is a false value μ_{wrong} , and it only remains to measure σ^2 . The MLE for the variance, $\hat{\sigma}^2 = \frac{1}{N} \sum_i (x_i - \mu_{\text{wrong}})^2$, is then biased: $E[\hat{\sigma}^2] = (\mu - \mu_{\text{wrong}})^2 + \sigma^2$. This arises because the normal distribution with mean μ_{wrong} mismodels the data, so downstream inference becomes unreliable. In this analogy, μ corresponds to the density ratio being estimated and σ corresponds to a parameter of interest. $w_i f_i$ ensembles then address this mismodeling by estimating the density ratio with appropriate uncertainties, and propagating these uncertainties for inferring the parameter(s) of interest.

These discrepancies between estimated and true density ratios can be ruinous for SBI. As shown in Ref. [41], existing methods for DRE can lead to underestimated uncertainties and overconfident intervals for inferred parameters. Therefore, to use DRE for parameter estimation, one needs not only a point estimate of the density ratio of interest, but also a notion of uncertainty on the estimate.

D. Existing work

Various approaches have been proposed to estimate uncertainties on NN outputs, both in DRE and in other contexts. The bulk of the machine learning literature has investigated Bayesian approaches with NNs: conceptually, instead of conventional NN training, one instead starts with a prior on network weights. Training then corresponds to a Bayesian update to obtain the weight posteriors (see e.g., Ref. [57] for a recent review of Bayesian methods). Once these posteriors are known, uncertainties on the network outputs can be estimated by sampling weight configurations from the posterior.

In practice, due to the computational difficulties involved in exact Bayesian inference, a variety of methods are used to approximate these posteriors, like sampling methods [58] and variational inference [59]. Various other techniques like dropout [60,61] and repulsive ensembles [62] can also be motivated as performing approximate Bayesian inference. These Bayesian techniques work well in many domains, but in the collider physics context, their prior dependence is undesirable.

Frequentist methods are less well-explored. Inference through a direct MLE analysis on NN weights using asymptotic formulas is infeasible for various reasons, including the difficulty of finding the global maximum of the likelihood as a function of the weights, the computational expense in computing and inverting the high-dimensional Hessian of the likelihood with respect to the weights, and the fact that one is rarely actually in an asymptotic regime for NNs due to their massive number of parameters.

In the face of these difficulties, the state of the art in collider physics is to use pseudoexperiments to perform the Neyman construction [42] (see Ref. [56] for a recent tour de force implementation of these methods). For example, for SBI, one still performs DRE to obtain an approximate likelihood ratio test statistic. If this test statistic were precisely the likelihood ratio, one could use asymptotic results like Wilks's theorem [63] to determine its distribution and perform a hypothesis test. Since this is not the case, one instead computes the test statistic on a large number of simulations, each generated with a fixed, counterfactual value of Θ^* , to determine this distribution pseudoempirically.

The Neyman construction enjoys the significant advantage that it provides reliable frequentist uncertainties under very general conditions: the quality of the trained NN determines the size of the resultant confidence intervals but not

their validity. However, it is computationally expensive: the number of simulations required is generically quite large. In particular, since it relies on having simulated samples which reasonably sample the possible values of Θ , this method suffers from a curse of dimensionality as the dimension of Θ grows. These difficulties can be improved, but not eliminated, using methods like those proposed in Refs. [64,65].

Another framework, similar to $w_i f_i$ ensembles, was developed and explored for DRE in Ref. [66] in the context of kernel methods. To the best of our knowledge, these methods have not previously been applied in collider physics or used with trained basis functions. Our approach of $w_i f_i$ ensembles, on the other hand, use trained NNs for the f_i . We also note that there are some interesting parallels between $w_i f_i$ ensembles and methods like AdaBoost [67], as it has been shown in Ref. [68] that there are deep connections between AdaBoost and the MLE of the ensemble weights.

III. INTRODUCING $w_i f_i$ ENSEMBLES

In order to quantify uncertainties on a density ratio $r(x)$, we introduce $w_i f_i$ ensembles: a framework that leverages ensembles of NNs to perform DRE with native uncertainties. To model the density ratio $r(x)$, we introduce the model,

$$\log \tilde{r}(x|w) = w_i f_i(x), \quad (5)$$

where the w_i are learned weights to be inferred from data, the $f_i(x)$ are trained NNs, and we employ Einstein notation to suppress the sum over the index i (which we take to run from 1 to M , the size of the ensemble).

We will see that uncertainties on the w_i can be assessed from data. Under the assumptions that the model is well-specified, i.e., that the f_i are chosen such that there exists a true set of weights w_i^* for which $\tilde{r}(x|w^*) = r(x)$, and that the $f_i(x)$ are fixed for the purposes of fitting the w_i , these uncertainties then immediately correspond to uncertainties on the density ratio. The well-specified assumption is, in effect, an assumption that the ansatz in Eq. (5) for the log density ratio does not suffer from mismodeling.

It is worth emphasizing that, in general, reducing mismodeling and estimating uncertainties are conceptually distinct goals. In the context of $w_i f_i$ ensembles, though, these goals are linked, since the weights w_i not only increase the flexibility of the model but also facilitate the extraction of confidence intervals. Of course, as mentioned in Sec. II C, $w_i f_i$ ensembles cannot reduce mismodeling or estimate uncertainties associated with flaws in the training data.

DRE with $w_i f_i$ ensembles consists of the following steps, which provides a rough outline for the rest of this section:

- (1) Train M networks $f_i(x)$ on a training dataset D_{train} (Sec. III A).
- (2) Freeze the f_i and infer the values of the best fit weights $\hat{w}_i$ on a distinct dataset D_{fit} by using a suitable optimization objective (Sec. III B).

- (3) Compute the covariance matrix $C_{ij} \equiv \text{Cov}[\hat{w}_i, \hat{w}_j]$ using analytic formulas (Sec. III C).
- (4) Propagate these uncertainties, either to the density ratio of interest itself or to downstream applications (Sec. III D).

The crux of $w_i f_i$ ensembles is in steps 2 and 3, while step 1 is relatively standard and step 4 is application specific. Finally, in Sec. III E we present the specific parameter estimation task, mixture fraction estimation, that we consider in the later empirical case studies.

A. Training the $f_i(x)$

The intuition behind $w_i f_i$ ensembles is to capture the uncertainty in density ratios by parametrizing them with the ansatz Eq. (5), using parametric methods to estimate the uncertainty on the $\hat{w}_i$, and then propagating these uncertainties to the density ratio itself (and to downstream applications). Clearly, if the density ratio of interest is to be written in the form Eq. (5), an arbitrary set of f_i does not suffice.

In particular, we emphasize that the asymptotic formulas which we derive are formally only applicable when the model is well-specified, where $\log r = w_i^* f_i$ for some w_i^* . Our chief objective then is to obtain a basis of functions f_i such that the vector space they span contains $\log r$.⁴

The optimal strategy to build such a basis in practice is not obvious. For our empirical case studies, we consider strategies where the f_i are individually trained to estimate the density ratio of interest.⁵ The two strategies that we consider are

- (i) *Partition*: Partition the training data into M disjoint groups and use each of these groups to train one of the f_i with Eq. (8).
- (ii) *Bootstrap*: Train each of the f_i networks with Eq. (8) on a bootstrapped resample of the training data.

The motivation behind this selection of strategies comes from the context of uniformly weighted ensembles where, as shown in Ref. [69], correlations between ensemble members introduce unnecessarily large variances which do not go to zero as the size of the ensemble becomes large.

These results do not carry over to $w_i f_i$ ensembles directly,

⁴In reality, it is exceptionally unlikely that this vector space actually contains $\log r$ for any realistic choices of training procedure and ensemble size M . For our inferences to be valid in practice, this condition is sufficient but not necessary; rather than requiring the total absence of mismodeling, it is in practice enough for the unquantifiable errors induced by mismodeling to be subdominant to the quantifiable uncertainties.

⁵We also tried a strategy where one preferred network is trained with Eq. (8) and the remaining networks are trained to be decorrelated from the preferred network and from each other. In our empirical testing, this strategy did not perform better than strategies which train the networks on equal footing, and the training was significantly more computationally intensive than the other methods.

as the $\hat{w}_i$ are themselves dependent on the f_i , but they motivate us to consider diverse f_i .

In addition to being diverse, the f_i should also be well-trained estimators of their objective functions. To see this, consider that a simple sufficient condition for the model to be well-specified is just for one of the f_i to be equal to the log density ratio. Intuitively, then, it is more plausible for the $w_i f_i$ functional form to be a well-specified model of the log density ratio of interest if the f_i are already individually close to learning this log density ratio. This motivates using standard best practices to train the f_i , from using an architecture well-suited to the data to using well-tuned hyperparameters.

The bootstrap and partition protocols both reduce correlations between the networks by training them on different training datasets, but represent different tradeoffs between diversity and individual performance of the f_i . Bootstrap allows each of the f_i to be trained with more data than partition, but partition totally eliminates the correlations between the f_i induced by the training data whereas bootstrap only reduces it. One can easily conceive of other strategies, e.g., bootstrap but using subsamples rather than resampling the whole dataset to accelerate training. We restrict ourselves to partition and bootstrap for simplicity, but it would be interesting to explore training methodologies more exhaustively.

For the case studies in Secs. IV and V, we will also consider a naive ensemble protocol as a baseline method for the sake of comparison. For naive ensemble, the f_i are trained identically to those of bootstrap, but the w_i are taken to be fixed to $1/M$ rather than fit to data. This means that naive ensemble does not provide a notion of uncertainty on density ratios, and no such uncertainty is propagated to downstream parameter estimation. In the large M limit, naive ensembles converge to a dataset-dependent average estimator, which will not generically be equal to the true density ratio.

Throughout, we use M to denote the number of trained f_i . However, we find empirically that it is also helpful to add one additional $f_0(x) = 1$; the corresponding w_0 then has the interpretation of satisfying a normalization constraint on the density ratio that we will encounter in Eq. (6). The index i then runs over $M + 1$ values. We include this additional f_0 in all experiments throughout, except for the naive ensemble protocol where the weights are fixed and not fit to data (alternatively, this is equivalent to fixing $w_0 = 0$).

In all cases, we use one dataset D_{train} to train the $f_i(x)$ and another dataset D_{fit} to fit the w_i . This is because our results for coverage on the w_i (and therefore on density ratios and downstream inferences) assume that the $f_i(x)$ are fixed functions, and formally they cannot be treated as fixed if they are determined through a training procedure which is dependent on the data used to fit the w_i . One way of seeing that the $f_i(x)$ must be fixed is to note that the true values of the weights w_i^* are themselves dependent on the $f_i(x)$;

therefore, one cannot even talk about frequentist coverage on the w_i unless the $f_i(x)$ are fixed. In principle, then, we expect that using the same dataset to train the $f_i(x)$ and to fit the w_i would neglect correlations between the $f_i(x)$ and the w_i , and so could result in confidence intervals which undercover. That said, we have performed experiments using the same dataset to train the $f_i(x)$ and fit the w_i , and we have not observed this effect in practice.

For the sake of simplicity, we take all of the f_i to have identical architectures and hyperparameters. This choice may be suboptimal, as it has been shown (e.g., in Ref. [70]) that varying hyperparameters of each member in traditional ensembles can lead to increased diversity and better performance in that context. We leave exploration of such training strategies to future work.

B. Finding $\hat{w}_i$

Suppose that a set of M functions $f_i(x)$ has been obtained through one of the aforementioned methods and frozen. To estimate the values of the w_i , we want some analogue of the MLE. However, we must account for two constraints following from proper normalization of n and d :

$$1 = \int dx d(x) e^{w_i f_i(x)} = \int dx n(x) e^{-w_i f_i(x)}. \quad (6)$$

As has been shown in Ref. [71], normalization constraints of this kind can be imposed through a Lagrange multiplier, and solving for this Lagrange multiplier for one of the constraints yields the maximum likelihood classifier (MLC) loss first considered in Refs. [72,73].

The maximum likelihood objective in conjunction with one of these normalization constraints (i.e., the MLC loss up to a minus sign) is already sufficient to yield the density ratio of interest asymptotically. However, to respect the symmetry of the problem under which n and d are interchanged and $w_i f_i \rightarrow -w_i f_i$, we introduce the following symmetrized MLC optimization objective:

$$\begin{aligned} \mathcal{L}_{\text{symMLC}} = & \int dx [-w_i f_i(x) n(x) + d(x) (e^{w_i f_i(x)} - 1)] \\ & + \int dx [w_i f_i(x) d(x) + n(x) (e^{-w_i f_i(x)} - 1)]. \end{aligned} \quad (7)$$

It can be checked that the density ratio of interest minimizes Eq. (7); thus, as long as the model is well-specified so that there exists a w_i^* for which $w_i^* f_i$ is precisely this density ratio, w_i^* minimizes the loss.

Specializing to a sample of N_n i.i.d. draws from n and N_d i.i.d. draws from d , this loss can be estimated on the dataset D_{fit} as

$$\begin{aligned} \mathcal{L}_{\text{samp}} = & \langle -w_i f_i(x) + (e^{-w_i f_i(x)} - 1) \rangle_n \\ & + \langle w_i f_i(x) + (e^{w_i f_i(x)} - 1) \rangle_d, \end{aligned} \quad (8)$$

where

$$\langle \cdot \rangle_n \equiv \frac{1}{N_n} \sum_{x \sim n}^{N_n} (\cdot), \quad (9)$$

and similarly for d . N_n and N_d need not be the same, but they are taken to have the same asymptotic power counting. We then define the $\hat{w}_i$ as

$$\hat{w}_i \equiv \operatorname{argmin}_{w_i} \mathcal{L}_{\text{samp}}. \quad (10)$$

As we show in Appendix A, under the assumption that the model is well-specified so that there exists a w_i^* for which $w_i^* f_i = \log r$, $\hat{w}_i$ is an asymptotically unbiased estimator for the w_i^* . Equation (8) is convex in the w_i , so the optimization is tractable and can be performed using Newton's method, for example.

C. Estimating C_{ij}

As we discuss further in Appendix A, $\hat{w}_i$ is an M -estimator [43], which is a class of estimators that arise when the objective function is a sample average. Asymptotically, M -estimators are Gaussian distributed, and the covariance matrix C_{ij} for the $\hat{w}_i$ distribution can be estimated by the sandwich formula,

$$C^{ij} = V^{ik} U_{kl} V^{lj}, \quad (11)$$

where

$$V_{ik} = \left. \frac{\partial^2 \mathcal{L}_{\text{samp}}(w)}{\partial w_i \partial w_k} \right|_{w=w^*} \quad (12)$$

is the Hessian matrix of the loss with respect to the w_i evaluated at w_i^* , $V^{ij} = V_{ij}^{-1}$ is the matrix inverse, and

$$U_{kl} = \operatorname{Cov} \left[\left. \frac{\partial \mathcal{L}_{\text{samp}}}{\partial w_k}, \frac{\partial \mathcal{L}_{\text{samp}}}{\partial w_l} \right] \right|_{w=w^*} \quad (13)$$

is the covariance matrix of the analogue of the score vector. The derivatives can be calculated analytically due to the known, simple dependence of the model (and therefore of the loss) on the w_i ,

$$\frac{\partial \mathcal{L}_{\text{samp}}}{\partial w_i} = \langle a_i(x) \rangle_n + \langle b_i(x) \rangle_d, \quad (14)$$

$$\begin{aligned} \frac{\partial^2 \mathcal{L}_{\text{samp}}}{\partial w_i \partial w_j} &= \langle f_i(x) f_j(x) e^{-w_k f_k(x)} \rangle_n \\ &\quad + \langle f_i(x) f_j(x) e^{w_k f_k(x)} \rangle_d, \end{aligned} \quad (15)$$

where we have introduced the notation,

$$a_i(x) = -f_i(x)(1 + e^{-w_j f_j(x)}), \quad (16)$$

$$b_i(x) = f_i(x)(1 + e^{w_j f_j(x)}). \quad (17)$$

V_{ik} can then be consistently estimated at this asymptotic order with the substitution $w_i^* \rightarrow \hat{w}_i$.

U_{kl} can also be estimated with the same substitution by using the assumption that the data is i.i.d., so the score vector decomposes into two i.i.d. sums of contributions. The covariance of these sums can then be estimated using the empirical covariance of the summands. In particular, we can estimate

$$\begin{aligned} U_{kl} &\sim \frac{1}{N_n} (\langle a_k(x) a_l(x) \rangle_n - \langle a_k(x) \rangle_n \langle a_l(x) \rangle_n) \\ &\quad + \frac{1}{N_d} (\langle b_k(x) b_l(x) \rangle_d - \langle b_k(x) \rangle_d \langle b_l(x) \rangle_d). \end{aligned} \quad (18)$$

Therefore, asymptotically, the covariance matrix C_{ij} can be consistently estimated from D_{fit} using the sandwich estimator Eq. (11). The corresponding uncertainty in the density ratio, estimated as

$$\log \hat{r}(x|\hat{w}) = \hat{w}_i f_i(x), \quad (19)$$

can then be read off as

$$\sigma^2(x) = f_i(x) C_{ij} f_j(x). \quad (20)$$

It can be seen from the form of Eq. (11) that $C_{ij} \sim O(N^{-1})$, like in the case of the MLE, and this scaling then holds for Eq. (20) as well.

D. Parameter estimation with $w_i f_i$ ensembles

With uncertainties on the $\hat{w}_i$ and on the density ratio in hand, we now move to the final step: leveraging the inferred density ratio and its uncertainties to perform downstream parameter estimation.

Suppose we have performed an experiment and obtained a dataset D_{exp} of samples drawn from $p(\Phi|\Theta^*)$, where Θ^* is the true value of Θ to be inferred from the data and Φ are the observed phase space variables. Suppose also that we have used the methods described above with $x = \{\Phi, \Theta\}$ to model $\log \ell(\Phi|\Theta) \equiv \frac{p(\Phi|\Theta)}{p(\Phi)}$ as $\log \tilde{\ell}(\Phi|\Theta, w) = w_i f_i(\Phi, \Theta)$, and we have obtained an estimate $\log \hat{\ell}(\Phi|\Theta, \hat{w}) \equiv \log \tilde{\ell}(\Phi|\Theta, \hat{w}) = \hat{w}_i f_i(\Phi, \Theta)$.

If it were the case that $\log \hat{\ell} = \log \ell$ exactly, then the estimator $\hat{\Theta}$ defined as

$$\hat{\Theta} \equiv \operatorname{argmax}_{\Theta} \log \hat{\ell}(\{\Phi\}|\Theta, \hat{w}), \quad (21)$$

would be the MLE for the Θ (where we have used $\{\Phi\}$ to denote that $\log \hat{\ell}$ is to be evaluated on the whole dataset, i.e., by summing over i.i.d. data). However, since it is generically the case that $\log \hat{\ell} \neq \log \ell$, $\hat{\Theta}$ is not actually the MLE, and asymptotic formulas for the variance of the MLE

underestimate the variance of $\hat{\Theta}$ due to neglecting the stochasticity of $\log \tilde{\mathcal{L}}$ itself.

We can account for this additional uncertainty using the Gong-Samaniego theorem [44]. In the terminology of Gong and Samaniego, $\tilde{\mathcal{L}}$ is a pseudolikelihood and $\hat{\Theta}$ is a pseudo-MLE. Suppose for simplicity that Θ is one-dimensional, i.e., that there is one parameter of interest θ . Then, under mild regularity conditions about the existence and boundedness of derivatives of $\log \tilde{\mathcal{L}}(\{\Phi\}|\theta, w)$, they show that $\hat{\theta}$ is asymptotically normally distributed with mean θ^* and variance σ_{GS}^2 , which can be consistently estimated by⁶

$$\sigma_{\text{GS}}^2 = \sigma_{\text{MLE}}^2 (1 + \sigma_{\text{MLE}}^2 A_i C_{ij} A_j), \quad (22)$$

where

$$\sigma_{\text{MLE}}^2 \equiv - \frac{\partial^2 \log \tilde{\mathcal{L}}(\{\Phi\}|\theta, \hat{w})}{\partial^2 \theta} \Big|_{\theta=\hat{\theta}}, \quad (23)$$

$$A_i \equiv \frac{\partial^2 \log \tilde{\mathcal{L}}(\{\Phi\}|\theta, w)}{\partial \theta \partial w_i} \Big|_{w=\hat{w}, \theta=\hat{\theta}}, \quad (24)$$

and C_{ij} is, as before, the covariance of the weights which can be estimated from D_{fit} .

Putting the pieces together, to perform SBI with $w_i f_i$ ensembles, one must first train a set of neural networks f_i with the dataset D_{train} , use the dataset D_{fit} to infer the weights $\hat{w}_i$ using Eq. (10), and then estimate the covariance matrix C_{ij} with Eq. (11). The estimated density ratio and uncertainties can then be used to obtain the pseudo-MLE $\hat{\theta}$ on D_{exp} , and to estimate its uncertainty σ_{GS}^2 using Eq. (22). As long as the choice of the f_i is such that the mismatching of the $w_i f_i$ parametrization of the density ratio is negligible, this procedure then yields asymptotically valid frequentist confidence intervals (or regions, in the multivariate case) for θ .

E. Mixture fraction estimation

The algorithms presented above provide a general framework to perform SBI with $w_i f_i$ ensembles. The case studies we consider in Secs. IV and V take a special form, though, where the likelihood $p(\Phi|\theta)$ depends in a known, analytic way on the parameter θ . Dividing by an appropriate reference distribution, we can write a likelihood ratio suitable for an MLE of θ as

$$L(\Phi|\theta) \equiv \frac{p(\Phi|\theta)}{p_{\text{ref}}(\Phi)} = F(r(\Phi), \theta), \quad (25)$$

where F is a function that we know analytically, and $r(\Phi)$ is a density ratio to be estimated from data but which only depends on Φ and not on θ . Modeling $r(\Phi)$ as $\tilde{r}(\Phi|w)$, and defining $\tilde{L}(\Phi|\theta, w) = F(\tilde{r}(\Phi|w), \theta)$, we can still perform inference for θ using Eqs. (21) and (22) under the substitution $\tilde{\mathcal{L}} \rightarrow \tilde{L}$. This approach allows us to take the derivatives defined in Eq. (23) analytically, without recourse to autodifferentiation or methods like the Gaussian ansatz proposed in Refs. [74,75].

Concretely, consider the so-called *mixture fraction task*: estimate the parameter κ given a dataset D_{mix} of size N_{mix} drawn i.i.d. from the mixture model,

$$p(x|\kappa) = \kappa n(x) + (1 - \kappa) d(x), \quad (26)$$

where κ is the mixture fraction and takes values between 0 and 1, and n and d are two probability distributions. Then, since

$$\log L(x|\kappa) \equiv \log \frac{p(x|\kappa)}{d(x)} = \log(\kappa r(x) + (1 - \kappa)), \quad (27)$$

with $p_{\text{ref}}(x)$ taken to be $d(x)$, knowledge of the ratio $r(x)$ [again defined as n/d according to Eq. (1)] is sufficient to extract an MLE for κ . Note that in this setup, $x = \Phi$ does *not* depend on the parameter κ .

As we saw in Sec. III D, a consistent estimator $\hat{r}(x|\hat{w})$ dependent on parameters w_i can then be used in place of $r(x)$ at the cost of introducing the additional variance captured by Eq. (22). Since the dependence of the likelihood on κ and w_i takes a known analytic form, the requisite (estimates of) derivatives of the likelihood can be calculated analytically when $\log \hat{r} = \hat{w}_i f_i$ is a $w_i f_i$ ensemble,

$$\sigma_{\text{MLE}}^2 = \left(\sum_{a=1}^{N_{\text{mix}}} \left(\frac{e^{\hat{w}_i f_i(x_a)} - 1}{\hat{\kappa} e^{\hat{w}_j f_j(x_a)} + (1 - \hat{\kappa})} \right)^2 \right)^{-1},$$

$$A_i = \sum_{a=1}^{N_{\text{mix}}} \frac{f_i(x_a) e^{\hat{w}_j f_j(x_a)}}{(\hat{\kappa} e^{\hat{w}_k f_k(x_a)} + (1 - \hat{\kappa}))^2}, \quad (28)$$

where the sum is over elements of D_{mix} . Here, $\hat{\kappa}$ is the pseudo-MLE for κ obtained from Eq. (21) after plugging in the estimates $\hat{w}$,

$$\hat{\kappa} = \text{argmax}_{\kappa} \sum_{a=1}^{N_{\text{mix}}} \log(\kappa e^{\hat{w}_i f_i(x_a)} + (1 - \kappa)). \quad (29)$$

We empirically confirm the validity of these formulas in the following case studies.

⁶The usual statement of the Gong-Samaniego theorem includes an extra term proportional to the covariance of $\log \tilde{\mathcal{L}}(\{\Phi\}|\theta^*, w^*)$ with respect to $\hat{w}_i$. Since the former depends only on the inference data D_{exp} and the latter depends only on the data D_{fit} with which $\hat{w}_i$ was determined, this covariance is 0 in our case.

IV. GAUSSIAN EXAMPLE

In this section, we show how $w_i f_i$ ensembles can be used to solve a one-dimensional Gaussian toy problem. This simple problem, where the density ratio of interest is known exactly, allows us to check the coverage properties of $w_i f_i$ ensembles both for the density ratio itself and for a downstream SBI task. In Sec. IV A, we describe the Gaussian toy problem in detail and establish the training methodology for our networks. In Sec. IV B, we discuss how we evaluate coverage of confidence intervals. Then, in Sec. IV C, we present results for coverage on the density ratio itself. Finally, in Sec. IV D, we examine coverage on an estimated parameter in the context of the mixture fraction task.

A. Training methodology

Let the numerator distribution n be a one-dimensional Gaussian with mean μ and variance 1, and the denominator distribution d also be a Gaussian with the same variance but mean $-\mu$. The DRE task is then to estimate the log likelihood ratio of n and d , $\log r(x) = 2\mu x$. We also consider the mixture fraction task, introduced in detail in Sec. III E, to estimate the fraction κ of events drawn from n in a mixed population of events from n and d . The mixture fraction task gives us a window into the performance of $w_i f_i$ ensembles that does not require a known analytic form for the density ratio of interest, as relevant for the QCD case study considered in Sec. V.

We take the training sets D_n and D_d to each contain $N_n = N_d \equiv N = 25,000$ events, and we choose $\mu = 0.1$. We take the $f_i(x)$ to be feedforward NNs with one hidden layer and width 32, using leaky ReLU activation functions with leakiness 0.2. We find that this simple architecture is sufficiently expressive for this Gaussian toy problem. To train the f_i , we use the Adam [76] optimizer and early stopping with a patience of 10 with a validation dataset, also with N samples each from n and d . Throughout, we use PyTorch [77] to implement and train NNs.

B. Coverage test methodology

To evaluate the behavior of $w_i f_i$ ensembles, we perform empirical tests of their coverage properties. We emphasize that these (expensive) tests are only to validate the procedure, and they would not be needed to deploy $w_i f_i$ ensembles in practice. Since the coverage depends on the flexibility of the model, we sweep over the number of trained networks $M \in \{2, 4, 8, 16\}$.

For each of the protocols proposed in Sec. III A (partition and bootstrap) and each value of M , we train the networks f_i $N_{\text{trainings}}$ times starting from different random initializations. Then, for each one of these trainings, we perform N_{trials} trials to fit the weights w_i and extract confidence intervals on both $\log r$ and κ , where each trial consists of the following steps:

- (1) Generate a dataset of size $N = 25,000$.
- (2) Calculate $\hat{w}_i$ and C_{ij} according to Eqs. (10) and (11). We use the PYTORCH-MINIMIZE package [78] to perform full-batch optimization with Newton's method over the w_i .
- (3) Sample one x at which to evaluate $\log r$; we choose either n or d with probability 0.5 and then sample x from the chosen distribution.
- (4) Calculate and save the true $\log r(x) = 2\mu x$, the estimated $\log \hat{r}(x) = \hat{w}_i f_i(x)$, and the estimated variance $\sigma^2(x)$ of $\log \hat{r}(x)$ given by Eq. (20). The estimate and the variance can be used to obtain a $z\sigma$ confidence interval $[\hat{w}_i f_i(x) - z\sqrt{\sigma^2(x)}, \hat{w}_i f_i(x) + z\sqrt{\sigma^2(x)}]$.
- (5) For each value of κ in $\{0.01, 0.02, 0.05, 0.1, 0.2, 0.5\}$, generate a mixture dataset of size $N = 25,000$ by sampling each event from n with probability κ and from d with probability $(1 - \kappa)$.
- (6) Compute $\hat{\kappa}$ from Eq. (29) and σ_{GS}^2 from Eq. (22), using Eq. (28) to evaluate the latter. We use the MINIMIZE_SCALAR routine in SciPy [79] to perform the minimization and find $\hat{\kappa}$.
- (7) Construct $z\sigma$ confidence intervals as $[\hat{\kappa} - z\sqrt{\sigma_{\text{GS}}^2}, \hat{\kappa} + z\sqrt{\sigma_{\text{GS}}^2}]$.

In the asymptotic limit, assuming the model is well-specified, the predictions for $\log \hat{r}(\hat{\kappa})$ should be normally distributed around the true $\log r(\kappa)$ with variance equal to $\sigma^2(\sigma_{\text{GS}}^2)$. For each training, we can estimate the coverage c of the $z\sigma$ confidence intervals on a quantity q with estimator $\hat{q}$ and uncertainty σ_q^2 as a function of z as

$$c(z) = \frac{1}{N_{\text{trials}}} \sum_{i=1}^{N_{\text{trials}}} \Theta \left(z - \frac{|\hat{q} - q|}{\sqrt{\sigma_q^2(x_i)}} \right), \quad (30)$$

where Θ is the Heaviside step function, defined to be 1 when its argument is positive and 0 otherwise. Said in words, Eq. (30) is the empirical fraction of outcomes observed in N_{trials} that lie within the desired confidence interval. The nominal coverage $\bar{c}(z)$ then satisfies $\bar{c}(1) \approx 0.68$, $\bar{c}(2) \approx 0.95$, and so forth; it is calculated as $\bar{c}(z) = P_{\text{norm}}(z) - P_{\text{norm}}(-z)$, where P_{norm} is the cumulative distribution function of a one-dimensional standard normal distribution. Model misspecification and deviations from the asymptotic limit could cause departures from these nominal coverages. Comparing $c(z)$ to $\bar{c}(z)$ then provides a direct check of the coverage properties of confidence intervals. To smooth out fluctuations over training instances, we average $c(z)$ over the $N_{\text{trainings}}$, and we also report the variances in the coverage over these instances. For these experiments, we take $N_{\text{trainings}} = 10$ and $N_{\text{trials}} = 300$.

C. DRE results

We start with the coverage properties on the density ratio itself. In Fig. 1, we plot the coverage $c(z)$ of $z\sigma$ confidence




FIG. 1. For the Gaussian case study, coverage $c(z)$ of a $z\sigma$ confidence interval for the density ratio $\log r$ for the partition and bootstrap training protocols from Sec. III A. Each quadrant corresponds to a different ensemble size $M \in \{2, 4, 8, 16\}$, and each color corresponds to a different training methodology. Coverage is estimated through $N_{\text{trials}} = 300$ trials, the solid lines are means over $N_{\text{trainings}} = 10$ trainings, and the shaded regions correspond to the standard error over the trainings. The dashed line corresponds to ideal coverage, $c(z) = \bar{c}(z)$; the region above the dashed line corresponds to overcoverage, and the region below the dashed line corresponds to undercoverage. In each quadrant, the upper panel shows the coverage $c(z)$ and the lower panel shows the residual coverage $c(z) - \bar{c}(z)$.

intervals on $\log r$. The results for the partition protocol are shown in blue and bootstrap are in orange. The nominal coverage $\bar{c}(z)$ is shown by the black dashed line. Each quadrant of the plot shows a different value of M (i.e., a different size ensemble).

The first phenomenon that we can see in Fig. 1 is that the coverage improves as M grows. For small values of M , the confidence intervals tend to undercover, but as M grows they saturate the correct coverage properties. This is consistent with the interpretation that failure to achieve the nominal coverage comes from model misspecification; as M grows, linear combinations of the f_i become increasingly expressive, and model misspecification necessarily decreases. We note that it is not always the case that coverage improves with increasing M ; when M becomes very large, asymptotic formulas derived with fixed M and N going to infinity are no longer applicable. We show an

example of this phenomenon and describe how to correct for it with moderate values of M , in Appendix C.

We can also observe in Fig. 1 that the differences in coverage between the training protocols are largely insignificant. For $M = 2, 4$, and 16 , the difference between bootstrap and partition is so small that we cannot sharply distinguish their performance over $N_{\text{trainings}} = 10$ trainings. For $M = 8$, partition does clearly outperform bootstrap, although in absolute terms the mean difference in coverage is never large; it is always less than 0.05 . We also find that once the coverage has converged, i.e., in the $M = 16$ case, the average sizes of the resultant confidence intervals are essentially identical for both methods.

For concreteness, we explicitly show one estimator for $\log r(x)$ and its uncertainties in Fig. 2 with $M = 16$, trained using the bootstrap protocol. A word of caution: $\log \hat{r}(x) = \hat{w}_i f_i(x)$ and $\log \hat{r}(x') = \hat{w}_i f_i(x')$ are correlated, with




FIG. 2. Upper panel: analytic (black dashed) and estimated (blue solid) log likelihood ratio $\log r$ as a function of x . The blue shaded region shows the 1σ confidence interval for $\log r$ as a function of x . Note that the confidence intervals are correlated between x values. Lower panel: the residual between the estimated and true log likelihood ratios.

covariance equal to $f_i(x)C_{ij}f_j(x')$. This means that we do not necessarily expect 68% of the blue shaded region in Fig. 2 to cover the true $\log r$ for any particular instantiation of $\log \hat{r}$, as coverage for one value of x is correlated with coverage for other values. Rather, the relevant coverage property is that for each fixed x , 68% of intervals obtained in this way should contain the true value of $\log r$ for that x ; this cannot be read off of Fig. 2.

D. Mixture fraction results

We now move to the mixture fraction task, where the goal is to extract the value of κ . In Fig. 3, we show the coverage results for all training methods with fixed $M = 16$, for each value of $\kappa \in \{0.01, 0.02, 0.05, 0.1, 0.2, 0.5\}$. As with Fig. 1, the solid lines and shaded regions show the coverage (or coverage residual) of each of the training methods. The dot-dashed line labeled “naive ensemble”, described in Sec. III A, corresponds to training an ensemble as in the Bootstrap protocol, but using uniform $1/M$ ensemble weights instead of fitting the w_i . The uncertainties in this naive case are taken to be σ_{MLE}^2 , so uncertainty on the density ratio itself is not assessed.

We can see that both the partition and bootstrap protocols achieve close to nominal coverage for all examined values of κ . It is visible in the residuals that the coverage improves marginally as κ becomes closer to its midpoint, and at $\kappa = 0.5$ the curves are essentially indistinguishable from the ideal coverage curve. Both $w_i f_i$ ensemble protocols dramatically outperform naive ensembles, which under-cover on average for all examined values of κ . We conclude that $w_i f_i$ ensembles yield reliable confidence intervals, at least in the context of a one-dimensional toy problem where the density ratio is known.

V. QUARK/GLUON DISCRIMINATION

To evaluate the performance of $w_i f_i$ ensembles on a nontrivial task, we apply them to a paradigmatic task in QCD and jet substructure. In this section, we use $w_i f_i$ ensembles to estimate the likelihood ratio of quark and gluon jets, as well as to estimate the quark fraction κ in a mixed sample of quark and gluon jets. In Sec. V A, we explain the physics context, the dataset, and our methods in detail. We present the results of this case study in Sec. V B, and we visualize the inferred likelihood ratio by reweighting a jet substructure observable in Sec. V C.

A. Methodology

The goal is to estimate the likelihood ratio of quark and gluon jets, which by the Neyman-Pearson lemma [39] is monotonically related to the optimal classifier for quark/gluon discrimination. See Ref. [40] for more on this equivalence, Refs. [80–89] for a selection of work on quark/gluon discrimination, and Refs. [90,91] for reviews of jet substructure in general. We focus only on information that is infrared and collinear (IRC) safe, so strictly speaking the extracted likelihood ratio is only monotonically related to the optimal IRC-safe classifier.

For this case study, we will not be able to directly evaluate coverage for the likelihood ratio, as the true likelihood ratio is not explicitly known. We therefore only examine performance on the mixture fraction task, where κ is now the quark fraction of a mixed sample of quark and gluon jets. We use the quark/gluon dataset [92] included with the ENERGYFLOW package [93], which was generated using Pythia 8.226 [94]. This dataset contains 2 million total anti- k_t [95] jets clustered with jet radius $R = 0.4$ using FASTJET3.3.0 [96], restricted to transverse momenta with $p_T \in \{500, 550\}$ GeV.

For the f_i , we use energy flow networks (EFNs) [93], which are naturally suited to the point cloud format of collider events and which respect IRC safety. We note that IRC safety is motivated by, but insufficient for, perturbative tractability (see Ref. [97] for more details in the context of EFNs). As before, we implement EFNs in PYTORCH [77]. An EFN has a per-jet Φ network and a per-particle F network, and we take Φ to have two hidden layers and F to have three, with all hidden layers taken to have width 32 for simplicity. We use leaky ReLU activation functions with leakiness 0.2. We also take the latent dimensionality of Φ to be 32, since this is where Ref. [93] begins to observe performance saturation for the quark/gluon task.

We sweep over the number of networks $M \in \{4, 8, 16, 32\}$, and again consider the partition and bootstrap protocols, along with the naive ensemble protocol as a baseline. As before, we perform $N_{\text{trainings}} = 10$ rounds of training and $N_{\text{trials}} = 300$ rounds of fitting the $\hat{w}_i$ to training data and estimating mixture fractions κ at the same value as in Sec. IV: $\kappa \in \{0.01, 0.02, 0.05, 0.1, 0.2, 0.5\}$.


Gaussian Case Study $M = 16$, Coverage on Mixture Fraction κ 

FIG. 3. The same as Fig. 1, but now assessing coverage on the mixture fraction κ in the Gaussian case study with fixed $M = 16$. The six displays correspond to $\kappa \in \{0.01, 0.02, 0.05, 0.1, 0.2, 0.5\}$. Also shown in the dot-dashed line is the naive ensemble protocol with fixed $w_i = 1/M$ and no quantification of DRE uncertainties.

Since this case study involves a fixed dataset, we cannot generate new events on the fly as in the Gaussian example. To ensure that the datasets used for training and inference are totally disjoint, for each protocol and for each training, we randomly divide the total dataset in half to obtain two disjoint pools of events. We sample 20,000 jets from the first pool (so $N_n \approx N_d \approx 10,000$) to obtain the training set, which we use to train the f_i (again using early stopping with a patience of 10 using a disjoint validation set of the same size, also drawn from the first pool). For each trial, we sample a new dataset (of the same size) again from the first pool to fit the $\hat{w}$. Then, for each value of κ , we draw a mixture dataset of 10,000 events from the second pool to use for inference.

We use analytic formulas for the gradient and Hessian of Eq. (8) to fit the $\hat{w}$, using the MINIMIZE routine included in SciPy [79] to perform Newton's method. We again use the MINIMIZE_SCALAR routine to calculate $\hat{\kappa}$ and its confidence interval.

B. Mixture fraction results

In Fig. 4, we show the coverages for each value of the mixture fraction κ with $M = 32$, which is the value of M at

which performance saturates. In contrast with the behavior observed in Sec. IV, there is a stark dependence on training methodology here: Bootstrap clearly outperforms partition for all values of κ . This makes some intuitive sense; in this higher-dimensional setting, it seems plausible that linear combinations of very poorly trained functions (as one obtains after partitioning the data 32 times, for example) are insufficiently expressive to model the data generating process, while this behavior is masked in a low-dimensional toy problem. These results suggest that well-trained networks may become increasingly important to reach nominal coverage as the problem becomes increasingly high-dimensional. We have also checked that the coverage is not meaningfully different from the $\kappa = 0.01$ case for smaller values of κ .

For concreteness, we show 50 example predictions and intervals for $\kappa = 0.1$, $M \in \{4, 8, 16, 32\}$ using bootstrap in Fig. 5. These results demonstrate that even the smallest $w_i f_i$ ensembles we consider produce qualitatively sensible intervals (i.e., they are in the vicinity of the correct value of κ , and they cover a significantly nonzero proportion of the time), but only sufficiently large ensembles ($M = 32$ for




FIG. 4. The same as Fig. 3, but now assessing coverage for the mixture fraction κ in the quark/gluon case study with $M = 32$.




FIG. 5. For the quark/gluon case study, results from 50 mixture fraction predictions and intervals each for $M \in \{4, 8, 16, 32\}$. The rightmost panel corresponds to naive ensemble predictions, and the remaining panels are $w_i f_i$ ensembles trained with the bootstrap protocol. For each value of M , all of the predictions come from the same initial training but different mixture datasets for inference. The horizontal axis is arbitrary; each value on this axis corresponds to one prediction of (and interval on) κ . The vertical axis shows the values of κ lying within each interval. The dashed horizontal line corresponds to the true value $\kappa = 0.1$.

this dataset) have sufficient expressivity to produce quantitatively correct coverage properties.

We also show predictions from a naive ensemble with $M = 32$; in contrast with the $w_i f_i$ ensembles, these intervals clearly undercover and show a substantial bias. The magnitude and sign of this bias varies from training to training. The training depicted in Fig. 5 is representative: the magnitude of the bias in this training is close to the average bias we observe over the ten trainings that we performed.

The presence of this bias for the naive ensemble and its absence for the $w_i f_i$ ensembles shows that $w_i f_i$ ensembles not only allow us to characterize uncertainties but also yield a better point estimate of the density ratio than the naive ensemble. Indeed, we observe that for fixed $f_i(x)$, the average values over N_{trials} trials of the $\hat{w}_i$ vary widely and are not generically well-approximated as $1/M$. Specifically, we observe that the average $\hat{w}_i$ still approximately sum to 1, but that they are sometimes negative and span a range of magnitudes from $O(10^{-2})$ to $O(1)$.

C. DRE visualization

Since the quark/gluon likelihood ratio is jet-valued and therefore high dimensional, we cannot visualize it directly. That said, we can provide a visualization of the learned likelihood ratio that is somewhat in the spirit of Fig. 2. We consider the distribution of a particular observable, the jet angularity [98–101],

$$\lambda_\beta = \sum_i z_i \theta_i^\beta, \quad (31)$$

where the sum is over jet constituents, z_i is the transverse momentum fraction of constituent i (i.e., $p_{Ti}/\sum_j p_{Tj}$), $\theta_i = \sqrt{y_i^2 + \phi_i^2}$ is its angular distance from the jet axis (which we take here to be the p_T weighted average position of the jet), y_i is the rapidity, and ϕ_i is the azimuthal coordinate. We specialize to the $\beta = 2$ angularity, which is related to the jet mass.

We show the distribution of this angularity for quark and gluon jets in Fig. 6. In the top panel, the red and black curves correspond to the results for gluon and quark jets respectively, and the blue curve shows gluon jets reweighted by the learned quark/gluon likelihood ratio, using an $M = 32$ bootstrap $w_i f_i$ ensemble. In the bottom panel, the blue curve shows the ratio of this reweighted distribution to the quark distribution.

The uncertainty bands are 1σ confidence intervals including both the Poisson statistical uncertainty as well as the uncertainty in the learned density ratio from the $w_i f_i$ ensemble. We calculate the former as the square root of the sum of squared weights in each bin, and the latter by sampling 3000 values of the w_i from a Gaussian centered at $\hat{w}_i$ with covariance C_{ij} , and calculating the standard deviation over samples of the sum of weights in




FIG. 6. Distributions of the $\beta = 2$ angularities for quark and gluon jets in our simulated dataset. The top panel shows the distribution of the angularities themselves for quark (black) and gluon (red) jets with unit weights, as well as the distribution for gluons jet reweighted using the learned quark/gluon density ratio (blue). The displayed uncertainties correspond both to Poisson uncertainties (for all distributions) and to the uncertainties in the weights (for the reweighted distribution).

each bin.⁷ It can be shown that these uncertainties add in quadrature up to terms subleading either in C or the number of counts in each bin.

We observe good agreement between the quark curve and the gluon-reweighted-to-quark curve. Since the reweighting introduces correlations between the bins, quantitatively checking the coverage of the intervals would require reweighting an ensemble of datasets; one cannot simply count the proportion of bins with intervals that cover for any given dataset. However, the uncertainties are qualitatively reasonable, covering in most but not all bins.

VI. CONCLUSIONS

In this paper, we introduced $w_i f_i$ ensembles, a new method for estimating frequentist confidence intervals both for density ratio estimation tasks and for downstream parameter estimation. We derived asymptotic coverage guarantees for these confidence intervals under the assumption that the $w_i f_i$ model is well-specified. We then tested these intervals empirically, both in a Gaussian toy problem and in a nontrivial jet substructure application, and we found that the intervals can successfully achieve the desired coverage properties when the ensemble is sufficiently large. For completeness, we also derive the leading order bias of the Gong-Samaniego estimator in Appendix C, which is relevant for two-step inference procedures more broadly.

⁷These calculations are performed at the level of counts, since it is counts that are Poisson distributed; the unit conversion to a probability density is performed at the end.

We tested two different training protocols for $w_i f_i$ ensembles and found that the bootstrap approach gave the best overall performance, but there are many more possibilities to be explored. We did not consider ensembles of different types of models, or of models trained with different hyperparameters; it is plausible that these directions could yield more diverse ensembles and achieve satisfactory coverage with smaller ensemble sizes M . Moreover, many methods exist in the literature to obtain ensembles of models without the naive $O(M)$ computational scaling, like Refs. [102–104]. It would be interesting to examine $w_i f_i$ ensembles trained with efficient protocols such as these.

Since the $f_i(x)$ must be fixed for the inference on the w_i to be well defined, we have used separate datasets D_{train} and D_{fit} to train the $f_i(x)$ and to fit the w_i . In order to achieve coverage and maximize power (i.e., minimize the size of intervals) given a fixed budget of data, there is then an interesting tradeoff to make between allocating data to D_{train} and D_{fit} . In particular, the quantified uncertainties are only directly sensitive to the size of D_{fit} , but an insufficiently large D_{train} may result in poor basis functions $f_i(x)$ and cause model misspecification. Using the same dataset to train the $f_i(x)$ and to fit the w_i would remove this tradeoff and allow one to use all of the data both for training the $f_i(x)$ and for fitting the w_i , but in principle correlations between the w_i and the $f_i(x)$ must then be taken into account, and without accounting for these correlations we are unable to rigorously guarantee coverage when reusing the training dataset to fit the w_i . In the cases that we have checked, however, neglecting these correlations does not seem to affect the coverage properties of our estimators. It would be interesting to investigate these correlations further, and particularly to determine how to include their contributions to uncertainties and when it is safe to neglect them.

It would also be interesting to investigate how $w_i f_i$ ensembles can slot into existing workflows. In principle, if one already has a well-trained estimator of the density ratio, one can simply use this model as one of the f_i and train the remaining networks using some other, less computationally demanding protocol (or using other estimators that one may already have at hand). In the limit where training this powerful estimator takes a very large amount of computation relative to evaluating it, this could be used to obtain a notion of uncertainties with a comparatively small amount of additional compute. If one already has an ensemble of estimators, converting this preexisting ensemble to a $w_i f_i$ ensemble only requires the optimization in Eq. (10) to obtain the $\hat{w}_i$. Additionally, while the formulas for calculating uncertainties that we have presented throughout are only applicable in the asymptotic regime, the $w_i f_i$ ensemble ansatz could be combined with existing methods (like the Neyman construction) to yield small-sample uncertainties on density ratios.

While we proposed $w_i f_i$ ensembles in the context of extracting reliable confidence intervals, they may also be useful even if one is not interested in obtaining uncertainties. In particular, it seems plausible that a $w_i f_i$ ensemble may yield a better point estimate of the density ratio than a traditional ensemble; fitting the $\hat{w}_i$ provides a data-driven way to lend more credence to the members of the ensemble that are more helpful and ignore those that are not. This could be a useful supplement to existing techniques using ensembles to stabilize DRE, like those explored in Ref. [105].

We observe hints of improved stabilization in our experiments. In the context of the quark fraction task, we found that bootstrap not only yields confidence intervals with much better coverage properties than the (identically trained) naive ensemble, but it also provides point estimates for $\hat{\kappa}$ with a substantially smaller mean squared error. In other words, the performance difference between bootstrap and naive ensemble is not only due to the need to account for uncertainties in the density ratio to achieve reliable uncertainties on the parameter κ , but also because bootstrap does a substantially better job at estimating the density ratio (and therefore κ) in the first place. It would be interesting to explore this further, to determine how general this improvement might be.

Example implementations of $w_i f_i$ ensembles, and in particular the code used to generate the figures, can be found at Ref. [106].

ACKNOWLEDGMENTS

We would like to thank Brian Nord for inspiring conversations about uncertainty quantification, Aishik Ghosh for discussion about ensemble methods, and Prasanth Shyamsundar for discussion about correlations between the f_i and the w_i . We would also like to thank Rikab Gambhir and Benjamin Nachman for useful discussions and comments. SB and J. T. are supported by the National Science Foundation (NSF) under Cooperative Agreement No. PHY-2019786 (The NSF AI Institute for Artificial Intelligence and Fundamental Interactions <http://iaifi.org/>), and by the U.S. Department of Energy (DOE) Office of High Energy Physics under Grant No. DE-SC0012567. JT is additionally supported by the Simons Foundation through Investigator Grant No. 929241, and in part by Grant No. NSF PHY-2309135 to the Kavli Institute for Theoretical Physics (KITP). The computations in this paper were run on the FASRC Cannon cluster supported by the FAS Division of Science Research Computing Group at Harvard University.

DATA AVAILABILITY

The data that support the findings of this article are openly available [106].

APPENDIX A: PROPERTIES OF THE $\hat{w}$ ESTIMATOR

In this appendix, we expound on crucial properties of the estimated weights $\hat{w}_i$, including their normality and asymptotic covariance matrix, as needed for the discussion in Sec. III C. We assume throughout that the model is well-specified, meaning that there is some true value w_i^* of the w_i for which $w_i^* f_i = \log r$, for r the density ratio of interest.

Recall that the $\hat{w}_i$ are defined as the minimizers of the $\mathcal{L}_{\text{samp}}$ loss in Eq. (8). Because the loss is defined through a sample average, this estimator is then an M -estimator in the language of Ref. [43]. It is consistent (i.e., asymptotically unbiased), since in the large N limit $\mathcal{L}_{\text{samp}}$ converges to $\mathcal{L}_{\text{symMLC}}$ in Eq. (7), which is minimized by the w_i^* .

Defining the gradient with respect to the weights as $u_i(w) \equiv \partial_{w_i} \mathcal{L}_{\text{samp}}$, the minimum $\hat{w}$ satisfies

$$u_i(\hat{w}) = 0. \quad (\text{A1})$$

Performing a Taylor expansion of this expression around the true w^* , we have

$$0 = u_i(\hat{w}) = u_i(w^*) + (w_j^* - \hat{w}_j) u_{ij}(w^*) + O((w^* - \hat{w})^2), \quad (\text{A2})$$

where we have defined $u_{ij} \equiv \partial_{w_j} u_i$. Asymptotically we are justified in neglecting the quadratic and higher order terms due to the consistency of $\hat{w}$, so we can rearrange this expression as

$$\sqrt{N}(w_j^* - \hat{w}_j) = u^{ji}(w^*)(\sqrt{N}u_i(w^*)), \quad (\text{A3})$$

where u^{ji} is the matrix inverse of u_{ji} . The factors of N (i.e., the characteristic data size) have been inserted for later convenience, and we have restricted for simplicity to the case of balanced samples $N_n = N_d \equiv N$.

The central limit theorem ensures that $\sqrt{N}u_i(w^*)$ is normally distributed; it is easy to check that its mean is zero for our estimator, and we denote its covariance matrix as $\tilde{U}_{ik}$. This covariance matrix does not scale with N . On the other hand, due to the law of large numbers applied to $u_{ji}(w^*)$, $u^{ji}(w^*)$ asymptotes to its expectation value $V^{ji}(w^*)$, where $V^{ij} \equiv V_{ij}^{-1}$. Slutsky's lemma [107], which allows us to apply these limiting arguments independently to each factor in a product, then tells us that the left-hand side of Eq. (A3) is normally distributed as

$$\sqrt{N}(w_j^* - \hat{w}_j) \sim \mathcal{N}(0, V^{ik} \tilde{U}_{kl} V^{lm}). \quad (\text{A4})$$

Defining $U_{ij} = \tilde{U}_{ij}/N$ to match the notation in the main text, we have

$$(w_j^* - \hat{w}_j) \sim \mathcal{N}(0, V^{ik} U_{kl} V^{lm}). \quad (\text{A5})$$

Thus, the estimated weights are indeed normally distributed with covariance matrix given by Eq. (11), which as mentioned in Sec. III C can itself be consistently estimated from the data.

APPENDIX B: PSEUDO-LIKELIHOOD RATIO TEST

There are a variety of ways to construct confidence intervals in standard asymptotic likelihood theory. In the main text, we focused for simplicity on the analog of Wald intervals [108], which are obtained in the one-dimensional case as “central value plus or minus estimated standard deviation”.

In the collider physics community, however, it is common practice is to construct intervals using a likelihood ratio test. Specifically, it can be shown that (twice) the difference between the log-likelihood evaluated at the maximum likelihood estimator and evaluated at the true value of the parameters follows a χ^2 distribution, so this difference can be used as a test statistic to construct confidence intervals. Intervals constructed with the likelihood ratio test enjoy several theoretical advantages, notably including parametrization invariance.

As such, in this appendix, we review an analogous pseudolikelihood ratio test appropriate for the construction of confidence intervals on a single parameter θ . This test, and the general test appropriate for inference over multiple parameters, was proposed in Ref. [109]. In our empirical case studies, the intervals produced with the pseudolikelihood ratio test were essentially indistinguishable from the naive intervals we presented in the main text, but we provide information on this alternative approach for completeness.

Consider the Taylor expansion around the estimated parameter $\hat{\theta}$ of the log-pseudo-likelihood evaluated at $\hat{w}$ and the ground truth value θ^* ,

$$\begin{aligned} \log L(\Phi|\hat{w}, \theta^*) &\approx \log L(\Phi|\hat{w}, \hat{\theta}) \\ &+ \frac{1}{2}(\theta^* - \hat{\theta})^2 \left(\frac{\partial^2}{\partial \theta^2} \log L(\Phi|\hat{w}, \theta) \right) \Big|_{\theta=\hat{\theta}}, \end{aligned} \quad (\text{B1})$$

where the linear term vanishes since $\hat{\theta}$ is at a maximum of the pseudolikelihood, and the higher order terms can be neglected since they are subleading in the characteristic dataset size N .

The distribution of $\hat{\theta} - \theta^*$ is asymptotically Gaussian with mean 0 and variance σ_{GS}^2 given by Eq. (22), and the second derivative term can be replaced at this order by minus the inverse of σ_{MLE}^2 . This means that the test statistic T defined as

$$T \equiv \frac{2(\log L(\Phi|\hat{w}, \hat{\theta}) - \log L(\Phi|\hat{w}, \theta^*))}{1 + \sigma_{\text{MLE}}^2 A_i C_{ij} A_j} \quad (\text{B2})$$

is asymptotically χ^2 distributed with one degree of freedom (with A and C defined as in Secs. III D and III C). As expected, this reduces to the usual likelihood ratio test statistic in the limit where the uncertainty on the w goes to zero.

APPENDIX C: LEADING ASYMPTOTIC BIAS

The estimators for the weights $\hat{w}$ and fractions $\hat{\kappa}$ presented in Sec. III are asymptotically unbiased in the sense that their bias vanishes at the leading order in the asymptotic expansion, i.e., it is $o(N^{-1/2})$, but it does not vanish at higher orders. The leading contribution to the bias of these estimators is generically $O(N^{-1})$.

Naively, the squared bias should then be negligible compared to the variance for $N \gg 1$. This is borne out in most of our experimentation, but we find that this expectation can break down as the number of members of the ensemble M grows large. This is because the asymptotic expansion assumes that all other quantities in the problem, like parameter dimensionality, are held fixed and finite as N grows large, so that in this power counting $N \gg M^p$ for any fixed p .

When this assumption is only mildly violated, the asymptotic expansion does not totally break down, but higher order terms can become important. In our experiments, this manifests as a bias in $\hat{\kappa}$, which causes coverage to be nonmonotonic in M for a fixed N (i.e., as M gets big, coverage first gets better, then worse).

In this appendix, we quote the result for the leading order bias of an M -estimator, which immediately yields the bias on the $\hat{w}$. We also derive the leading order bias in the Gong-Samaniego estimator for $\hat{\kappa}$. We study an example where coverage of a $w_i f_i$ ensemble is adversely affected by the presence of a bias, that this bias is well-described by the leading order asymptotic bias, and that subtracting off an estimate of this bias largely restores coverage for $\hat{\kappa}$.

Formally, the next-to-leading order contribution to the variance is of the same order as the leading squared-bias, but we find that it is numerically less important for coverage in our experiments. As such, we leave the calculation of this contribution to the variance (which is substantially more involved than the calculation of the bias) to future work.

The study of higher-order contributions to the distribution of asymptotic estimators has a long history. The asymptotic bias of the MLE was first calculated in 1968 in Ref. [110], and the next-to-leading contribution to the variance was first calculated in Ref. [111]. There has been a large amount of subsequent work on so-called high-dimensional statistics, where the dimensionality of the parameters is allowed to grow with N ; see for example Ref. [112], which considers the context of high-dimensional logistic regression.

Within this large body of work, the leading order bias for an M -estimator has been studied in Refs. [113,114]. This

immediately allows us to estimate the bias of $\hat{w}_i$ as

$$\langle \hat{w}^i - w^{*i} \rangle = V^{ik} \left(V^{lm} \sum_x u_{s,kl} u_{s,m} + \frac{1}{2} V^{ln} V^{mp} \left(\sum_x u_{s,klm} \right) \left(\sum_x u_{s,n} u_{s,p} \right) \right), \quad (C1)$$

where the s subscript indicates that the quantity is evaluated on one sample, not the whole dataset, and other subscripts indicate differentiation with respect to the w_i . All quantities can be calculated with $w_i = \hat{w}_i$ at this order. Empirically, we find that this correction is negligible throughout. For the pseudolikelihood estimator for a parameter θ , where the likelihood is a function of parameters w_i and the pseudolikelihood is obtained with the plug-in estimators $\hat{w}_i$, we can Taylor expand the pseudoscore u to find

$$u(\hat{\theta}, \hat{w}) = u + (\hat{\theta} - \theta^*) u' + (\hat{w}_i - w_i^*) u_i + \frac{1}{2} (\hat{\theta} - \theta^*)^2 u'' + \frac{1}{2} (\hat{\theta} - \theta^*) (\hat{w}_i - w_i^*) u'_i + \frac{1}{2} (\hat{w}_i - w_i^*) (\hat{w}_j - w_j^*) u_{ij} + O(N^{-1/2}), \quad (C2)$$

where primes denote differentiation with respect to θ , subscripts mean differentiation with respect to the w_i , and the functions on the right-hand side are all evaluated at the true parameter values. By definition of the pseudo-MLE, the left-hand side vanishes, so at leading order we find

$$\hat{\theta} - \theta^* = \frac{-u - (\hat{w}_i - w_i^*) u_i}{u'}, \quad (C3)$$

and by the law of large numbers, at this order of approximation we can substitute the denominator with its expectation value. We now plug this into Eq. (C2), use that the left-hand side vanishes, and take an expectation value to obtain

$$\langle (\hat{\theta} - \theta^*) u' \rangle = -\frac{1}{2} \left[\frac{\langle u^2 u'' \rangle}{\langle u' \rangle^2} + \frac{C_{ij} \langle u_i u_j u'' \rangle}{\langle u' \rangle^2} - \frac{C_{ij} \langle u_i u'_j \rangle}{\langle u' \rangle} + C_{ij} \langle u_{ij} \rangle \right]. \quad (C4)$$

Furthermore, we have

$$\begin{aligned} \langle (\hat{\theta} - \theta^*) u' \rangle &= \langle \hat{\theta} - \theta^* \rangle \langle u' \rangle + \text{Cov}[\hat{\theta} - \theta^*, u'] \\ &= \langle \hat{\theta} - \theta^* \rangle \langle u' \rangle - \text{Cov} \left[\frac{u}{\langle u' \rangle}, u' \right] \\ &= \langle \hat{\theta} - \theta^* \rangle \langle u' \rangle - \frac{\langle u u' \rangle}{\langle u' \rangle}, \end{aligned} \quad (C5)$$

so

$$\langle \hat{\theta} - \theta^* \rangle = \frac{\langle uu' \rangle}{\langle u' \rangle^2} - \frac{1}{2\langle u' \rangle} \left[\frac{\langle u^2 u'' \rangle}{\langle u' \rangle^2} + \frac{C_{ij} \langle u_i u_j u'' \rangle}{\langle u' \rangle^2} - \frac{C_{ij} \langle u_i u'_j \rangle}{\langle u' \rangle} + C_{ij} \langle u_{ij} \rangle \right]. \quad (\text{C6})$$

We can now do some decompositions. Let $a, b, c \dots$ be sample indices, then for example we have

$$u = \frac{1}{N} \sum_a u_a, \quad (\text{C7})$$

and likewise for all derivatives of u . Therefore,

$$\begin{aligned} \langle u^2 u'' \rangle &= \left\langle \sum_{a,b,c} u_a u_b u''_c \right\rangle \\ &= N^2 \langle u_s^2 \rangle \langle u''_s \rangle + O(N), \end{aligned} \quad (\text{C8})$$

where the subscript s means that the function is evaluated on one sample rather than on the entire dataset, and the N^3 term vanishes because $\langle u_s \rangle = 0$. Similar logic applies to the other terms, so one obtains to the desired order that

$$\begin{aligned} N \langle \hat{\theta} - \theta^* \rangle &= \frac{\langle u_s u'_s \rangle}{\langle u'_s \rangle^2} - \frac{1}{2\langle u'_s \rangle} \left[\frac{\langle u_s^2 \rangle \langle u''_s \rangle}{\langle u'_s \rangle^2} + N C_{ij} \langle u_{s,ij} \rangle \right. \\ &\quad \left. + \frac{N C_{ij} \langle u_{s,i} \rangle \langle u_{s,j} \rangle \langle u''_s \rangle}{\langle u'_s \rangle^2} - \frac{N C_{ij} \langle u_{s,i} \rangle \langle u'_{s,j} \rangle}{\langle u'_s \rangle} \right]. \end{aligned} \quad (\text{C9})$$

All expectation values are $O(1)$, so the right-hand side is manifestly also $O(1)$. Since the right-hand side is evaluated at w^*, θ^* , we have $\langle u_s^2 \rangle = -\langle u'_s \rangle$, so

$$\begin{aligned} \langle \hat{\theta} - \theta^* \rangle &= \frac{\langle u_s u'_s \rangle}{N \langle u'_s \rangle^2} - \frac{1}{2N \langle u'_s \rangle} \left[-\frac{\langle u''_s \rangle}{\langle u'_s \rangle} + N C_{ij} \langle u_{s,ij} \rangle \right. \\ &\quad \left. + \frac{N C_{ij} \langle u_{s,i} \rangle \langle u_{s,j} \rangle \langle u''_s \rangle}{\langle u'_s \rangle^2} - \frac{N C_{ij} \langle u_{s,i} \rangle \langle u'_{s,j} \rangle}{\langle u'_s \rangle} \right] \\ &= \frac{\langle u_s u'_s \rangle}{N \langle u'_s \rangle^2} - \frac{1}{2\langle u'_s \rangle} \left[\langle u''_s \rangle \sigma_{\text{GS}}^2 \right. \\ &\quad \left. - \frac{C_{ij} \langle u_{s,i} \rangle \langle u'_{s,j} \rangle}{\langle u'_s \rangle} + C_{ij} \langle u_{s,ij} \rangle \right]. \end{aligned} \quad (\text{C10})$$

We can now define a bias corrected estimator $\hat{\theta}_{\text{BC}}$ by subtracting off the empirical estimate of Eq. (C10).

With this bias-corrected expression in hand, we return to the mixture fraction task in the Gaussian toy problem from Sec. IV. We use the bootstrap protocol with $M \in \{8, 16, 32\}$, using the same architecture and testing methodology described earlier. In Fig. 7, we show the distribution of naive estimate $\hat{\kappa}$ and bias-corrected version $\hat{\kappa}_{\text{BC}}$. It is immediately evident that $\hat{\kappa}$ is substantially biased for $M > 8$, but that the bias correction largely mitigates this for both $M = 16$ and $M = 32$.

In Fig. 8, we show the mixture fraction coverage for these configurations. We show the coverage both for $\hat{\kappa}$ and for $\hat{\kappa}_{\text{BC}}$, using in both cases the leading order expression for the variance given by Eq. (22). We see that the bias correction substantially improves the coverage for $M = 32$, where the bias becomes comparable to the spread of the estimator.

Gaussian Case Study, Uncorrected and Corrected Residuals on $\kappa = 0.1$




FIG. 7. Residuals histogram for the mixture fraction task in the Gaussian toy problem with $\kappa = 0.1$, using $M \in \{8, 16, 32\}$ networks in the {left, middle, right} panel. The networks are trained using the bootstrap protocol with (brown) and without (purple) the bias correction. The dashed black vertical line is at 0, the brown vertical line is the mean of the bias corrected predictions, and the purple vertical line is the mean of the uncorrected predictions; results are for ten trainings with 300 trials each.


Gaussian Case Study, Uncorrected and Corrected Coverage for $\kappa = 0.1$ 

FIG. 8. Coverage plot for the mixture fraction task in the Gaussian toy problem with $\kappa = 0.1$ using $M \in \{8, 16, 32\}$ networks trained with the bootstrap protocol with (brown) and without (purple) the bias correction. Results are shown for ten trainings with 300 trials each. The dashed black line in each panel shows the nominal coverage; the top panels show coverage, and the bottom panels show the difference between the coverage and the nominal coverage.

We observe empirically that the variance of $\hat{\kappa}_{BC}$ is about 15% larger than that of $\hat{\kappa}$ for $M = 32$; the variances of these estimators are identical at $O(N^{-1})$, the order at which we truncate, so this suggests that even better coverage could be

obtained by calculating the next-to-leading contribution to the variance. This would be a substantially more involved calculation, though, so we leave this to future work.

-
- [1] A. M. Sirunyan *et al.* (CMS Collaboration), Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV, *J. Instrum.* **13**, P05011 (2018).
 - [2] A. Butter *et al.*, The machine learning landscape of top taggers, *SciPost Phys.* **7**, 014 (2019).
 - [3] E. Bols, J. Kieseler, M. Verzetti, M. Stoye, and A. Stakia, Jet flavour classification using DeepJet, *J. Instrum.* **15**, P12012 (2020).
 - [4] H. Qu, C. Li, and S. Qian, Particle transformer for jet tagging, *arXiv:2202.03772*.
 - [5] S. Gong, Q. Meng, J. Zhang, H. Qu, C. Li, S. Qian, W. Du, Z.-M. Ma, and T.-Y. Liu, An efficient Lorentz equivariant graph neural network for jet tagging, *J. High Energy Phys.* **07** (2022) 030.
 - [6] A. Bogatskiy, T. Hoffman, D. W. Miller, J. T. Offermann, and X. Liu, Explainable equivariant neural networks for particle physics: PELICAN, *J. High Energy Phys.* **03** (2024) 113.
 - [7] G. Aad *et al.* (ATLAS Collaboration), Dijet resonance search with weak supervision using $\sqrt{s} = 13$ TeV *pp* collisions in the ATLAS detector, *Phys. Rev. Lett.* **125**, 131801 (2020).
 - [8] G. Kasieczka *et al.*, The LHC Olympics 2020 a community challenge for anomaly detection in high energy physics, *Rep. Prog. Phys.* **84**, 124201 (2021).
 - [9] T. Aarrestad *et al.*, The dark machines anomaly score challenge: Benchmark data and model independent event classification for the Large Hadron Collider, *SciPost Phys.* **12**, 043 (2022).
 - [10] E. M. Metodiev, J. Thaler, and R. Wynne, Anomaly detection in collider physics via factorized observables, *Phys. Rev. D* **110**, 055012 (2024).
 - [11] G. Aad *et al.* (ATLAS Collaboration), Anomaly detection search for new resonances decaying into a Higgs boson and a generic new particle *X* in hadronic final states using $\sqrt{s} = 13$ TeV *pp* collisions with the ATLAS detector, *Phys. Rev. D* **108**, 052009 (2023).
 - [12] G. Aad *et al.* (ATLAS Collaboration), Search for new phenomena in two-body invariant mass distributions using unsupervised machine learning for anomaly detection at

- $s = 13$ TeV with the ATLAS detector, *Phys. Rev. Lett.* **132**, 081801 (2024).
- [13] V. Chekhovsky *et al.* (CMS Collaboration), Model-agnostic search for dijet resonances with anomalous jet substructure in proton–proton collisions at $\sqrt{s} = 13$ TeV, *Rep. Prog. Phys.* **88**, 067802 (2025).
- [14] R. Gambhir, R. Mastandrea, B. Nachman, and J. Thaler, Isolating unisolated Upsilon's with anomaly detection in CMS open data, *Phys. Rev. Lett.* **135**, 021902 (2025).
- [15] G. Aad *et al.* (ATLAS Collaboration), Search for non-resonant Higgs boson pair production in the $b\bar{b}\ell\nu\ell\nu$ final state with the ATLAS detector in pp collisions at $\sqrt{s} = 13$ TeV, *Phys. Lett. B* **801**, 135145 (2020).
- [16] G. Aad *et al.* (ATLAS Collaboration), Search for Higgs boson decays into a Z boson and a light hadronically decaying resonance using 13 TeV pp collision data from the ATLAS detector, *Phys. Rev. Lett.* **125**, 221802 (2020).
- [17] G. Aad *et al.* (ATLAS Collaboration), Search for heavy resonances decaying into a pair of Z bosons in the $\ell^+\ell^-\ell'^+\ell'^-$ and $\ell^+\ell^-\nu\bar{\nu}$ final states using 139 fb $^{-1}$ of proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector, *Eur. Phys. J. C* **81**, 332 (2021).
- [18] A. Tumasyan *et al.* (CMS Collaboration), Search for a heavy Higgs boson decaying into two lighter Higgs bosons in the $\tau\tau b\bar{b}$ final state at 13 TeV, *J. High Energy Phys.* **11** (2021) 057.
- [19] A. Andreassen, P. T. Komiske, E. M. Metodiev, B. Nachman, and J. Thaler, OmniFold: A method to simultaneously unfold all observables, *Phys. Rev. Lett.* **124**, 182001 (2020).
- [20] M. Bellagente, A. Butter, G. Kasieczka, T. Plehn, and R. Winterhalder, How to GAN away detector effects, *SciPost Phys.* **8**, 070 (2020).
- [21] M. Bellagente, A. Butter, G. Kasieczka, T. Plehn, A. Rousselot, R. Winterhalder, L. Ardizzone, and U. Köthe, Invertible networks or partons to detector and back again, *SciPost Phys.* **9**, 074 (2020).
- [22] M. Vandegar, M. Kagan, A. Wehenkel, and G. Louppe, Neural empirical bayes: Source distribution estimation and its applications to simulation-based inference, *arXiv:2011.05836*.
- [23] A. Andreassen, P. T. Komiske, E. M. Metodiev, B. Nachman, A. Suresh, and J. Thaler, Scaffolding simulations with deep learning for high-dimensional deconvolution, in *Proceedings of the 9th International Conference on Learning Representations* (2021).
- [24] M. Arratia *et al.*, Publishing unbinned differential cross section results, *J. Instrum.* **17**, P01024 (2022).
- [25] J. N. Howard, S. Mandt, D. Whiteson, and Y. Yang, Learning to simulate high energy particle collisions from unlabeled data, *Sci. Rep.* **12**, 7567 (2022).
- [26] M. Backes, A. Butter, M. Dunford, and B. Malaescu, An unfolding method based on conditional invertible neural networks (cINN) using iterative training, *SciPost Phys. Core* **7**, 007 (2024).
- [27] A. Shmakov, K. Greif, M. Fenton, A. Ghosh, P. Baldi, and D. Whiteson, End-to-end latent variational diffusion models for inverse problems in high energy physics, *arXiv:2305.10399*.
- [28] A. Shmakov, K. Greif, M. J. Fenton, A. Ghosh, P. Baldi, and D. Whiteson, Full event particle-level unfolding with variable-length latent variational diffusion, *SciPost Phys.* **18**, 117 (2025).
- [29] S. Otten, S. Caron, W. de Swart, M. van Beekveld, L. Hendriks, C. van Leeuwen, D. Podareanu, R. Ruiz de Austri, and R. Verheyen, Event generation and statistical sampling for physics with deep generative models and a density information buffer, *Nat. Commun.* **12**, 2985 (2021).
- [30] B. Hashemi, N. Amin, K. Datta, D. Olivito, and M. Pierini, LHC analysis-specific datasets with generative adversarial networks, *arXiv:1901.05282*.
- [31] R. Di Sipio, M. Fucci Giannelli, S. Ketabchi Haghighat, and S. Palazzo, DijetGAN: A generative-adversarial network approach for the simulation of QCD Dijet events at the LHC, *J. High Energy Phys.* **08** (2019) 110.
- [32] A. Butter, T. Plehn, and R. Winterhalder, How to GAN LHC events, *SciPost Phys.* **7**, 075 (2019).
- [33] Y. Alanazi *et al.*, Simulation of electron-proton scattering events by a Feature-Augmented and Transformed Generative Adversarial Network (FAT-GAN), *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence Main Track*, 10.24963/ijcai.2021/293 (2020), pp. 2126–2132.
- [34] A. Butter and T. Plehn, Generative networks for LHC events, *arXiv:2008.08558*.
- [35] A. Butter, T. Heimel, S. Hummerich, T. Krebs, T. Plehn, A. Rousselot, and S. Vent, Generative networks for precision enthusiasts, *SciPost Phys.* **14**, 078 (2023).
- [36] E. Buhmann, G. Kasieczka, and J. Thaler, EPiC-GAN: Equivariant point cloud generation for particle jets, *SciPost Phys.* **15**, 130 (2023).
- [37] O. Amram *et al.*, CaloChallenge 2022: A community challenge for fast calorimeter simulation, *arXiv:2410.21611*.
- [38] M. Feickert and B. Nachman, A living review of machine learning for particle physics, *arXiv:2102.02770*.
- [39] J. Neyman and E. S. Pearson, On the problem of the most efficient tests of statistical hypotheses, *Phil. Trans. R. Soc. A* **231**, 289 (1933).
- [40] K. Cranmer, J. Pavez, and G. Louppe, Approximating likelihood ratios with calibrated discriminative classifiers, *arXiv:1506.02169*.
- [41] J. Hermans, A. Delaunoy, F. Rozet, A. Wehenkel, V. Begy, and G. Louppe, A trust crisis in simulation-based inference? Your posterior approximations can be unfaithful, *arXiv:2110.06581*.
- [42] J. Neyman, Outline of a theory of statistical estimation based on the classical theory of probability, *Phil. Trans. R. Soc. A* **236**, 333 (1937).
- [43] P. J. Huber, Robust estimation of a location parameter, *Ann. Math. Stat.* **35**, 73 (1964).
- [44] G. Gong and F. J. Samaniego, Pseudo maximum likelihood estimation: Theory and applications, *Ann. Stat.* **9**, 861 (1981).
- [45] S. Navas *et al.* (Particle Data Group), Review of particle physics, *Phys. Rev. D* **110**, 030001 (2024).
- [46] G. P. Salam, The strong coupling: A theoretical perspective, in *From My Vast Repertoire ...: Guido Altarelli's Legacy*, edited by A. Levy, S. Forte, and G. Ridolfi (World Scientific, Singapore, 2019), pp. 101–121.

- [47] A. Pich, J. Rojo, R. Sommer, and A. Vairo, Determining the strong coupling: status and challenges, *Proc. Sci., Confinement2018* (**2018**) 035 [arXiv:1811.11801].
- [48] D. d'Enterria *et al.*, α_s (2019): Precision measurements of the QCD Coupling, arXiv:1907.01435.
- [49] A. Pich, Precision physics with inclusive QCD processes, *Prog. Part. Nucl. Phys.* **117**, 103846 (2021).
- [50] D. d'Enterria *et al.*, The strong coupling constant: State of the art and the decade ahead, *J. Phys. G* **51**, 090501 (2024).
- [51] P. Nason, The top mass in hadronic collisions, in *From My Vast Repertoire ...: Guido Altarelli's Legacy*, edited by A. Levy, S. Forte, and G. Ridolfi (World Scientific, Singapore, 2019) pp. 123–151.
- [52] A. H. Hoang, What is the top quark mass?, *Annu. Rev. Nucl. Part. Sci.* **70**, 225 (2020).
- [53] A. Hayrapetyan *et al.* (ATLAS Collaboration and CMS Collaboration), Combination of measurements of the top quark mass from data collected by the ATLAS and CMS experiments at $s = 7$ and 8 TeV, *Phys. Rev. Lett.* **132**, 261902 (2024).
- [54] B. Lakshminarayanan, A. Pritzel, and C. Blundell, Simple and scalable predictive uncertainty estimation using deep ensembles, arXiv:1612.01474.
- [55] T. G. Dietterich, Ensemble methods in machine learning, in *Multiple Classifier Systems* (Springer, Berlin, Heidelberg, 2000), pp. 1–15.
- [56] G. Aad *et al.* (ATLAS Collaboration), An implementation of neural simulation-based inference for parameter estimation in ATLAS, *Rep. Prog. Phys.* **88**, 067801 (2025).
- [57] J. Arbel, K. Pitas, M. Vladimirova, and V. Fortuin, A primer on bayesian neural networks: Review and debates, arXiv:2309.16314.
- [58] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, Weight uncertainty in neural networks, arXiv:1505.05424.
- [59] A. Graves, Practical variational inference for neural networks, in *Advances in Neural Information Processing Systems*, Vol. 24, edited by J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Weinberger (Curran Associates, Inc., Red Hook, NY, USA, 2011).
- [60] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, *J. Mach. Learn. Res.* **15**, 1929 (2014).
- [61] Y. Gal and Z. Ghahramani, Dropout as a bayesian approximation: Representing model uncertainty in deep learning, in *Proceedings of The 33rd International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 48, edited by M. F. Balcan and K. Q. Weinberger (PMLR, New York, 2016), pp. 1050–1059.
- [62] F. D'Angelo and V. Fortuin, Repulsive deep ensembles are bayesian, arXiv:2106.11642.
- [63] S. S. Wilks, The large-sample distribution of the likelihood ratio for testing composite hypotheses, *Ann. Math. Stat.* **9**, 60 (1938).
- [64] N. Dalmasso, R. Izbicki, and A. B. Lee, Confidence sets and hypothesis testing in a likelihood-free inference setting, arXiv:2002.10399.
- [65] N. Dalmasso, L. Masserano, D. Zhao, R. Izbicki, and A. B. Lee, Likelihood-free frequentist inference: Bridging classical statistics and machine learning for reliable simulator-based inference, *Electron. J. Stat.* **18**, 5045 (2024).
- [66] M. Sugiyama, T. Suzuki, S. Nakajima, H. Kashima, P. von Büna, and M. Kawanabe, Direct importance estimation for covariate shift adaptation, *Ann. Inst. Stat. Math.* **60**, 699 (2008).
- [67] Y. Freund and R. E. Schapire, A decision-theoretic generalization of on-line learning and an application to boosting, *J. Comput. Syst. Sci.* **55**, 119 (1997).
- [68] G. Lebanon and J. Lafferty, Boosting and maximum likelihood for exponential models, in *Advances in Neural Information Processing Systems*, Vol. 14, edited by T. Dietterich, S. Becker, and Z. Ghahramani (MIT Press, Cambridge, MA, 2001).
- [69] G. Louppe, Understanding random forests: From theory to practice, arXiv:1407.7502.
- [70] F. Wenzel, J. Snoek, D. Tran, and R. Jenatton, Hyperparameter ensembles for robustness and uncertainty quantification, arXiv:2006.13570.
- [71] B. Nachman and J. Thaler, Learning from many collider events at once, *Phys. Rev. D* **103**, 116013 (2021).
- [72] R. T. D'Agnolo and A. Wulzer, Learning new physics from a machine, *Phys. Rev. D* **99**, 015014 (2019).
- [73] R. T. D'Agnolo, G. Grosso, M. Pierini, A. Wulzer, and M. Zanetti, Learning multivariate new physics, *Eur. Phys. J. C* **81**, 89 (2021).
- [74] R. Gambhir, B. Nachman, and J. Thaler, Bias and priors in machine learning calibrations for high energy physics, *Phys. Rev. D* **106**, 036011 (2022).
- [75] R. Gambhir, B. Nachman, and J. Thaler, Learning uncertainties the frequentist way: Calibration and correlation in high energy physics, *Phys. Rev. Lett.* **129**, 082001 (2022).
- [76] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, arXiv:1412.6980.
- [77] A. Paszke *et al.*, PYTORCH: An imperative style, high-performance deep learning library, arXiv:1912.01703.
- [78] R. Feinman, PYTORCH-MINIMIZE (2022).
- [79] P. Virtanen *et al.* (SciPy 1.0 Contributors), SciPy 1.0: Fundamental algorithms for scientific computing in Python, *Nat. Methods* **17**, 261 (2020).
- [80] L. M. Jones, Tests for determining the parton ancestor of a hadron jet, *Phys. Rev. D* **39**, 2550 (1989).
- [81] Z. Fodor, How to see the differences between quark and gluon jets, *Phys. Rev. D* **41**, 1726 (1990).
- [82] L. Lonnblad, C. Peterson, and T. Rognvaldsson, Finding gluon jets with a neural trigger, *Phys. Rev. Lett.* **65**, 1321 (1990).
- [83] L. Lonnblad, C. Peterson, and T. Rognvaldsson, Using neural networks to identify jets, *Nucl. Phys. B* **349**, 675 (1991).
- [84] I. Csabai, F. Czako, and Z. Fodor, Quark and gluon jet separation using neural networks, *Phys. Rev. D* **44**, R1905 (1991).
- [85] L. Jones, Towards a systematic jet classification, *Phys. Rev. D* **42**, 811 (1990).
- [86] J. Pumplin, How to tell quark jets from gluon jets, *Phys. Rev. D* **44**, 2025 (1991).

- [87] P.D. Acton *et al.* (OPAL Collaboration), A study of differences between quark and gluon jets using vertex tagging of quark jets, *Z. Phys. C* **58**, 387 (1993).
- [88] J. Gallicchio and M.D. Schwartz, Quark and gluon jet substructure, *J. High Energy Phys.* **04** (2013) 090.
- [89] P. Gras, S. Höche, D. Kar, A. Larkoski, L. Lönnblad, S. Plätzer, A. Siódmok, P. Skands, G. Soyez, and J. Thaler, Systematics of quark/gluon tagging, *J. High Energy Phys.* **07** (2017) 091.
- [90] A. J. Larkoski, I. Moulton, and B. Nachman, Jet substructure at the Large Hadron Collider: A review of recent advances in theory and machine learning, *Phys. Rep.* **841**, 1 (2020).
- [91] R. Kogler *et al.*, Jet substructure at the Large Hadron Collider: Experimental review, *Rev. Mod. Phys.* **91**, 045003 (2019).
- [92] P. Komiske, E. Metodiev, and J. Thaler, Pythia8 quark and gluon jets for energy flow, Zenodo, [10.5281/zenodo.3164691](https://doi.org/10.5281/zenodo.3164691) (2019).
- [93] P. T. Komiske, E. M. Metodiev, and J. Thaler, Energy flow networks: Deep sets for particle jets, *J. High Energy Phys.* **01** (2019) 121.
- [94] C. Bierlich *et al.*, A comprehensive guide to the physics and usage of PYTHIA 8.3, *SciPost Phys. Codebases* **2022**, 8 (2022).
- [95] M. Cacciari, G. P. Salam, and G. Soyez, The anti- k_t jet clustering algorithm, *J. High Energy Phys.* **04** (2008) 063.
- [96] M. Cacciari, G. P. Salam, and G. Soyez, FASTJET user manual, *Eur. Phys. J. C* **72**, 1896 (2012).
- [97] S. Bright-Thonney, B. Nachman, and J. Thaler, Infrared-safe energy weighting does not guarantee small non-perturbative effects, *Phys. Rev. D* **110**, 014029 (2024).
- [98] C. F. Berger, T. Kucs, and G. F. Sterman, Event shape/energy flow correlations, *Phys. Rev. D* **68**, 014012 (2003).
- [99] L. G. Almeida, S. J. Lee, G. Perez, G. F. Sterman, I. Sung, and J. Virzi, Substructure of high- p_T jets at the LHC, *Phys. Rev. D* **79**, 074017 (2009).
- [100] S. D. Ellis, C. K. Vermilion, J. R. Walsh, A. Hornig, and C. Lee, Jet shapes and jet algorithms in SCET, *J. High Energy Phys.* **11** (2010) 101.
- [101] A. J. Larkoski, J. Thaler, and W. J. Waalewijn, Gaining (mutual) information about quark/gluon discrimination, *J. High Energy Phys.* **11** (2014) 129.
- [102] G. Huang, Y. Li, G. Pleiss, Z. Liu, J. E. Hopcroft, and K. Q. Weinberger, Snapshot ensembles: Train 1, get m for free, [arXiv:1704.00109](https://arxiv.org/abs/1704.00109).
- [103] T. Garipov, P. Izmailov, D. Podoprikin, D. Vetrov, and A. G. Wilson, Loss surfaces, mode connectivity, and fast ensembling of DNNs, [arXiv:1802.10026](https://arxiv.org/abs/1802.10026).
- [104] Y. Wen, D. Tran, and J. Ba, Batchensemble: An alternative approach to efficient ensemble and lifelong learning, [arXiv:2002.06715](https://arxiv.org/abs/2002.06715).
- [105] F. T. Acosta, T. Wamorkar, V. Mikuni, and B. Nachman, Stabilizing neural likelihood ratio estimation, [arXiv:2503.20753](https://arxiv.org/abs/2503.20753).
- [106] S. Benevedes and J. Thaler, $w_i f_i$ ensemble code, Github (2025).
- [107] E. Slutsky, Über stochastische asymptoten und grenzwerte (abdruck), *Tohoku Math. J.* **27**, 67 (1925).
- [108] A. Wald, Tests of statistical hypotheses concerning several parameters when the number of observations is large, *Trans. Am. Math. Soc.* **54**, 426 (1943).
- [109] Y. Chen and K.-Y. Liang, On the asymptotic behaviour of the pseudolikelihood ratio test statistic with boundary problems, *Biometrika* **97**, 603 (2010).
- [110] D. R. Cox and E. J. Snell, A general definition of residuals, *J. R. Stat. Soc. Ser. B* **30**, 248 (1968).
- [111] E. Naftali, First order bias and second order variance of the maximum likelihood estimator with application to multivariate Gaussian data and time delay and Doppler shift estimation, Ph.D. thesis, MIT, 2000.
- [112] P. Sur and E. J. Candès, A modern maximum-likelihood theory for high-dimensional logistic regression, *Proc. Natl. Acad. Sci. U.S.A.* **116**, 14516 (2019).
- [113] P. Rilstone, V. Srivastava, and A. Ullah, The second-order bias and mean squared error of nonlinear estimators, *J. Econ.* **75**, 369 (1996).
- [114] K. I. Kim, Higher order bias correcting moment equation for m-estimation and its higher order efficiency, *Econometrics* **4**, 48 (2016).