



Unsupervised Quark/Gluon Jet Tagging With Poissonian Mixture Models

E. Alvarez^{1*}, M. Spannowsky^{2,3} and M. Szewc¹

¹ International Center for Advanced Studies (ICAS) and CONICET, UNSAM, San Martin, Argentina, ² Institute for Particle Physics Phenomenology, Durham University, Durham, United Kingdom, ³ Department of Physics, Durham University, Durham, United Kingdom

OPEN ACCESS

Edited by:

Jan Kieseler,
European Organization for Nuclear
Research (CERN), Switzerland

Reviewed by:

Dag Gillberg,
Carleton University, Canada
Maximilian Swiatlowski,
TRIUMF, Canada

*Correspondence:

E. Alvarez
sequi@unsam.edu.ar

Specialty section:

This article was submitted to
Big Data and AI in High Energy
Physics,
a section of the journal
Frontiers in Artificial Intelligence

Received: 11 January 2022

Accepted: 18 February 2022

Published: 17 March 2022

Citation:

Alvarez E, Spannowsky M and
Szewc M (2022) Unsupervised
Quark/Gluon Jet Tagging With
Poissonian Mixture Models.
Front. Artif. Intell. 5:852970.
doi: 10.3389/frai.2022.852970

The classification of jets induced by quarks or gluons is important for New Physics searches at high-energy colliders. However, available taggers usually rely on modeling the data through Monte Carlo simulations, which could veil intractable theoretical and systematical uncertainties. To significantly reduce biases, we propose an unsupervised learning algorithm that, given a sample of jets, can learn the SoftDrop Poissonian rates for quark- and gluon-initiated jets and their fractions. We extract the Maximum Likelihood Estimates for the mixture parameters and the posterior probability over them. We then construct a quark-gluon tagger and estimate its accuracy in actual data to be in the 0.65–0.7 range, below supervised algorithms but nevertheless competitive. We also show how relevant unsupervised metrics perform well, allowing for an unsupervised hyperparameter selection. Further, we find that this result is not affected by an angular smearing introduced to simulate detector effects for central jets. The presented unsupervised learning algorithm is simple; its result is interpretable and depends on very few assumptions.

Keywords: jets, QCD, unsupervised learning, inference, LHC

1. INTRODUCTION

As ongoing searches at the LHC have not succeeded in providing guidance on the nature of extensions of the Standard Model, unbiased event reconstruction and classification methods for new resonances and interactions have become increasingly important to ensure that the gathered data is exploited to its fullest extent (Aarrestad et al., 2021; Kasieczka et al., 2021). Unsupervised data-driven classification frameworks, often used as anomaly-detection methods, are comprehensive in scope, as their success does not hinge on how theoretically well-modeled signal or background processes are allowing for more signal-agnostic analyses (Andreassen et al., 2019, 2020; Choi et al., 2020; Dohi, 2020; Hajer et al., 2020; Nachman and Shih, 2020; Roy and Vijay, 2020; Caron et al., 2021; d’Agnoles et al., 2021).

Here, we apply the unsupervised-learning paradigm to the discrimination of jets induced by quarks or gluons. The so-called quark-gluon tagging of jets can be a very powerful method to separate signal from background processes. Important examples include the search for dark matter at colliders, where the dark matter candidates are required to recoil against a single hard jet (CMS coll., 2015), the measurement of Higgs boson couplings in the weak-boson fusion process (Dokshitzer et al., 1992; Rainwater et al., 1998) or the discovery of SUSY cascade decays involving squarks or gluinos (Bhattacharjee et al., 2017). Thus, a robust and reliable method to discriminate

between quark and gluon jets furthers the scientific success of the LHC programme in precision measurements and searches for new physics. Consequently, several approaches have been proposed to exploit the differences in the radiation profiles of quarks and gluons (Gallicchio and Schwartz, 2011; Larkoski et al., 2013, 2014b; Bhattacharjee et al., 2015; Ferreira de Lima et al., 2017; Kasieczka et al., 2019) and have been studied in data by ATLAS coll (2014) and CMS coll (2013).

The discrimination of quarks and gluons as incident particles for a jet poses a challenging task. Some of the best performing observables to classify quark/gluon jets are infrared and/or collinear (IRC) unsafe, e.g. the number of charged tracks of a jet. Thus, evaluating the classification performance of IRC unsafe observables from the first principles is an inherently difficult task. Instead, SoftDrop (Larkoski et al., 2014a) has been shown to achieve a high classification performance while maintaining IRC safety. Further, at leading-logarithmic accuracy, the SoftDrop multiplicity n_{SD} exhibits a Poisson-like scaling (Frye et al., 2017), allowing us to construct an entire data-driven unsupervised classifier based on a mixture model. Although IRC safety is in principle not necessary to construct an unsupervised tagger, the fact that we are able to know the leading-logarithmic behavior of the observable is what allows us to build a simple and interpretable probabilistic model. To build a tagger that discriminates between quark and gluon jets, we extract the Maximum Likelihood Estimate (MLE) and the posterior distributions for the rate of the Poissonians and the mixing proportions of the respective classes. We find such a tagger to have a high accuracy (≈ 0.7) while remaining insensitive to detector effects. In a second step, we augment this method by using Bayesian inference to obtain the full set of posterior distributions and correlations between the model parameters, which allows to calculate the probability of a jet being a quark or gluon jet. The latter method results in a robust tagger with even higher accuracy. Thus, this approach opens a novel avenue to analyse jet-rich final states at the LHC, thereby increasing the sensitivity in searches for new physics.

The structure of this article is as follows: In section 2, we present the datasets considered and describe the mixture model method for the discrimination of quark and gluon jets. In section 3, we discuss the performance and uncertainties of the MLE algorithm, detailing the viability of this algorithm in the presence of detector effects. We use Bayesian inference to obtain the full posterior probability density function in section 4. In section 5, we offer a summary and conclusions.

2. MIXTURE MODELS FOR QUARK- AND GLUON-JETS DATA

To showcase and benchmark our model performance in a way comparable to other algorithms, we have considered two datasets available in the literature, both considered initially in Komiske et al. (2019b). These two datasets contain quark and gluon jets after hadronization, and correspond to a set (Komiske et al., 2019a) generated with Pythia (Sjöstrand et al., 2015) and a set (Pathak et al., 2019) generated with Herwig (Bahr et al.,

2008; Bellm et al., 2016). The reason for using datasets from different generators is to verify that the algorithm is independent of the generator and from any specific tuning. As detailed in the documentation, the quark- and gluon-initiated jets are generated from $qg \rightarrow Z(\rightarrow \nu\bar{\nu}) + u/d/s$ and $q\bar{q} \rightarrow Z(\rightarrow \nu\bar{\nu}) + g$ processes in pp collisions at $\sqrt{s} = 14$ TeV. After hadronization, the jets are clustered using the anti- k_T algorithm with $R = 0.4$. For the sake of validation and comparing supervised and unsupervised metrics, we use the parton level information to define whether a jet is a true quark or a true gluon. This definition is known to be problematic and we emphasize that our model does not depend on these unphysical labels and could instead provide an operational definition of quark and gluon jets (Komiske et al., 2018b). In addition, there is a very strict selection cut and all the provided jets have transverse momentum $p_T \in [500.0, 550.0]$ GeV and rapidity $|y| < 1.7$. We detail the impact of these cuts on the tagging observable in the following paragraphs. Finally, the dataset is balanced with an equal number of quark and gluon jets.

As a tagging observable, we have considered the Iterative SoftDrop Multiplicity n_{SD} defined in Frye et al. (2017). Once defined the jet radius R used to cluster the constituents with the Cambridge-Aachen algorithm (Dokshitzer et al., 1997), n_{SD} has three hyperparameters z_{cut} , β , and θ_{cut} . The dependence on these hyperparameters and the classification performance on supervised tasks has been explored in Frye et al. (2017). In this work and in agreement with Frye et al. (2017) we will consider IRC safe parameter choices: $z_{cut} > 0$, $\beta < 0$ and $\theta_{cut} = 0$.

The choice of a well-known tagging observable allows us to perform unsupervised quark-gluon discrimination by considering interpretable mixture models (Štěpánek et al., 2015; Metodiev et al., 2017; Komiske et al., 2018a; Metodiev and Thaler, 2018; Dillon et al., 2019, 2020, 2021; Alvarez et al., 2020, 2021; Graziani et al., 2021), where we think of the measurement of N jets and their n_{SD} values as originating from underlying themes, which we would ideally match with quark and gluon jets. In a probabilistic modeling framework, we want to obtain the underlying quark and gluon distributions from the observed data $X = \{n_{SD}^{(i)}, i = 1, \dots, N\}$, which has a likelihood function

$$p(X) = \prod_{i=1}^N p(n_{SD}^{(i)}) = \prod_{i=1}^N \sum_{k=\{q,g\}} \pi_k p(n_{SD}^{(i)} | k), \quad (1)$$

where k is the jet class or theme. The mixing fraction π_k denotes the probability of sampling a jet from theme k and $p(n_{SD}^{(i)} | k)$ is the n_{SD} probability mass functions conditioned on which theme the jet belongs to. In principle, the number of themes is a hyperparameter of the model and could be optimized with some criteria (see, e.g., Celeux et al. (2018) for a review on different methods to select the number of themes). In this work, we consider only two themes that we identify with quark and gluon jets. This choice is based on physical grounds. As we will detail in the following paragraphs, for a sufficiently small p_T range we only expect two themes. The final ingredient to build the probabilistic model is the specification of $p(n_{SD}^{(i)} | k)$. To model these probability mass functions, we make use of the fact that at

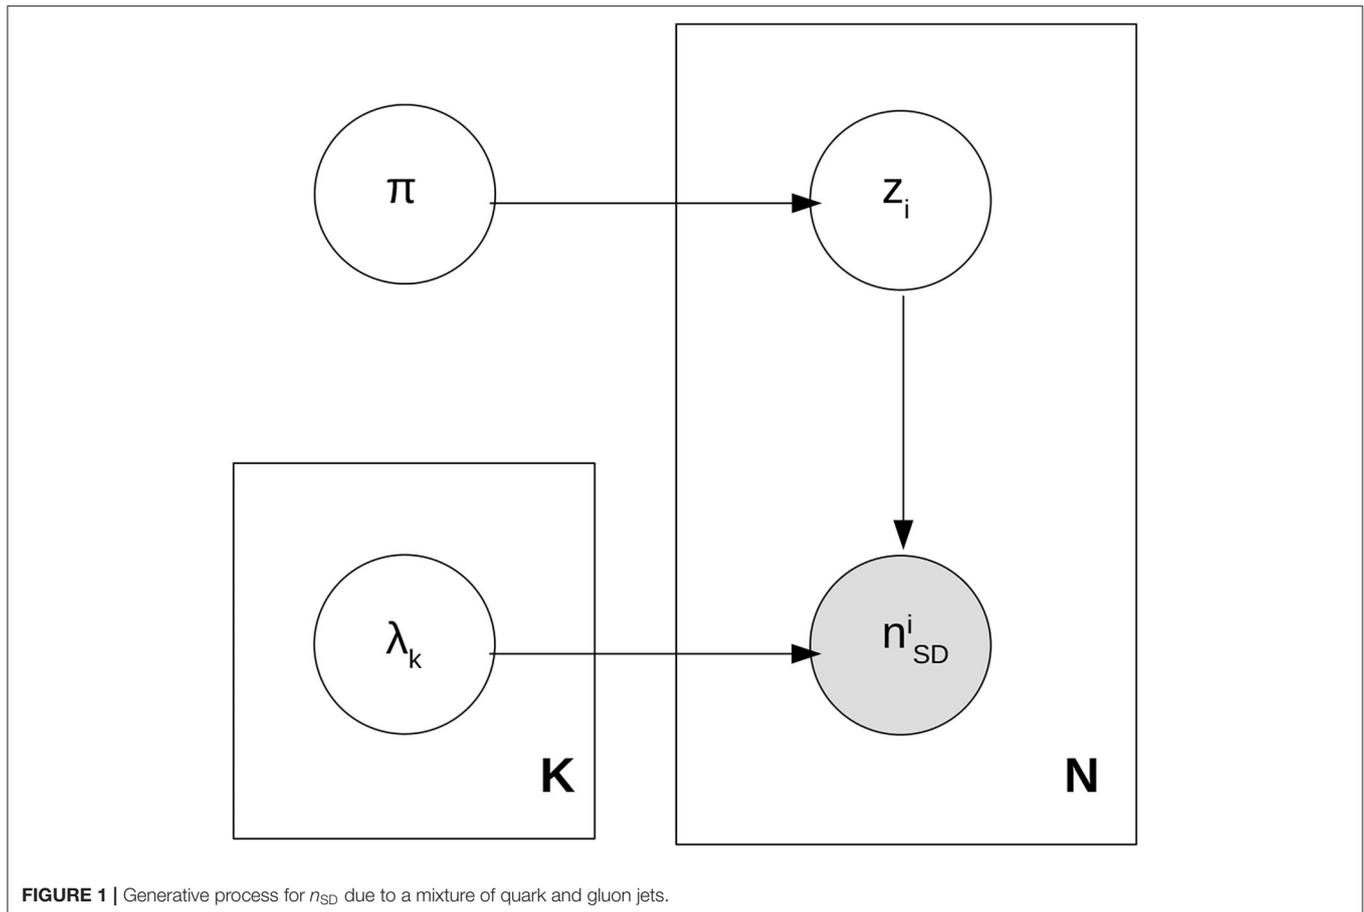


FIGURE 1 | Generative process for n_{SD} due to a mixture of quark and gluon jets.

leading logarithmic (LL) order n_{SD} is Poisson distributed (Frye et al., 2017):

$$p(X) = \prod_{i=1}^N \sum_{k=\{q,g\}} \pi_k \text{Poisson}(n_{SD}^{(i)}; \lambda_k), \quad (2)$$

where λ_k is the Poisson rate for each theme, which fixes the mean and variance of the n_{SD} distribution. The departure of the Poisson hypothesis by NLL corrections and Non Perturbative effects can then be seen by examining the variance to mean ratio for each class of jets. We see that deviations from the Poissonian behavior are parameter-dependent, in agreement with previous results detailed in Frye et al. (2017). Furthermore, we see that the deviations are more substantial for quark jets than for gluon jets. This is enhanced by the fact that quark jets usually have smaller n_{SD} values than gluon jets.

The behavior of n_{SD} is also dependent on the kinematics of the jet. For the samples considered, the limited p_T and $|y|$ ranges ensure that all quark- and gluon-initiated jets follow the same respective n_{SD} distributions. In a more realistic implementation of this model where the p_T of the jets populates a much wider range, the model implementation should be modified to account for the variation of the n_{SD} distribution with p_T . A straightforward strategy is to bin the p_T distribution and infer the mixture model parameters in each bin, effectively conditioning

the Poisson rates and the mixture fractions on the p_T of the jets populating such region. The p_T dependence also implies that the discriminating power of the mixture model will depend on the p_T of the jet as is usually the case for most quark/gluon classification methods.

The likelihood in Equation (2) describes how for a given value of the mixing fractions $\pi_{q,g}$ and the Poisson rates $\lambda_{q,g}$, each jet is sampled or generated. This is called a generative process and it is often useful to represent it as a plaque diagram (Bishop, 2013). The corresponding plaque diagram to Equation (2) can be seen in Figure 1.

In the figure, we have introduced a hidden or latent class variable, the theme assignment z , which dictates whether the generated jet is a quark or a gluon jet. This class assignment is necessary to think of the likelihood as a generative process and it is useful when performing inference and when building a probabilistic jet classifier. Having defined the probabilistic model and the relevant parameters π and λ , finding the underlying themes becomes synonymous with finding the posterior probabilities for π and λ . These posterior probabilities can then be used to build a quark/gluon classifier. Instead of tackling the Bayesian Inference problem head-on, one can first obtain point estimates for π and λ . Because we consider a statistically significant dataset and a straightforward model, at this stage we consider the Maximum Likelihood Estimates (MLE)

of π and λ instead of Maximum A Posterior (MAP) estimates. These estimates can be obtained easily through Expectation-Maximization (EM) or with Stochastic Variational Inference through dedicated software such as the Pyro package (Bingham et al., 2018; Phan et al., 2019). We need to be careful when estimating the point parameters as they can suffer from mode degeneracy and mode collapse. The former occurs due to the permutation symmetry of the classes and can be fixed by requiring that $\lambda_q < \lambda_g$ as dictated by basic principles. The latter occurs when one theme is emptied of samples. Because we only consider two classes, collapse is avoided for good hyperparameter choices because of the multimodality of the data distribution.

With the MLE point estimation of π^{MLE} and λ^{MLE} , we can construct a probabilistic jet classifier by computing the assignment probabilities or responsibilities,

$$p(z = \text{quark} | n_{\text{SD}}, \pi^{\text{MLE}}, \lambda^{\text{MLE}}) = \frac{\pi_q^{\text{MLE}} \text{Poisson}(n_{\text{SD}}, \lambda_q^{\text{MLE}})}{\sum_{k=\{q,g\}} \pi_k^{\text{MLE}} \text{Poisson}(n_{\text{SD}}, \lambda_k^{\text{MLE}})}, \quad (3)$$

with $p(z = \text{gluon}) = 1 - p(z = \text{quark})$. The classifier is obtained by selecting a threshold $0 \leq c \leq 1.0$ and labeling any jet with $p(z = \text{quark} | n_{\text{SD}}, \pi^{\text{MLE}}, \lambda^{\text{MLE}}) \geq c$ as a quark jet. This classifier has a clear probabilistic justification and it is interpretable, which is a considerable asset for an unsupervised task.

For validation, we compute the usual supervised metrics: the accuracy obtained by assigning classes using the probabilistic working point $c = 0.5$ chosen because we have a binary classification problem and a probabilistic algorithm, the mistag rate at 50% signal efficiency $\epsilon_g^{-1}(\epsilon_q = 50\%)$ and the Area-Under-Curve (AUC). The accuracy is defined as the number of fraction of well classified samples, $\epsilon_{q,g}$ are the fraction of well classified quark/gluon jets and the AUC is the integral of the Receiver Operating Characteristic (ROC) curve $\epsilon_q(\epsilon_g)$ with a higher AUC usually signaling a higher overall performance. However, because we are interested in an unsupervised classifier trained directly on data, we also define unsupervised metrics. These metrics need to be correlated with the unseen accuracy so as to substitute it as a measure of performance in a fully data-driven implementation of the model. In an unsupervised metric we measure how consistent is the learned model with the measured data. We investigate two metrics that encode such consistency:

$$d_H(p, q) = \frac{1}{\sqrt{2}} \sqrt{\sum_{n_{\text{SD}}=0}^{\infty} (\sqrt{p(n_{\text{SD}})} - \sqrt{q(n_{\text{SD}})})^2}$$

$$\text{KL}(p||q) = - \sum_{n_{\text{SD}}=0}^{\infty} p(n_{\text{SD}}) \text{Ln} \left(\frac{q(n_{\text{SD}})}{p(n_{\text{SD}})} \right), \quad (4)$$

where d_H is the Hellinger distance (Deza and Deza, 2009) and KL is the Kullback-Leibler divergence between the learned data density and the measured data density. The latter can be interpreted as the amount of information needed to approximate samples that follow the distribution p with samples generated by a model q . In this article, p will be the measured data density obtained by the n_{SD} frequencies

and q will be the posterior predictive distribution $q(n_{\text{SD}}) = \sum_{k=\{q,g\}} \pi_k^{\text{MLE}} \text{Poisson}(n_{\text{SD}}, \lambda_k^{\text{MLE}})$. Other metrics such as the Energy Mover's Distance (Komiske et al., 2019c) could also be applied. We emphasize that this takes advantage of the fact that we are learning more than a classifier, as we are modeling the data density itself and the underlying processes that generate it. If we can match the learned models to quark and gluon jets, it means we can understand the data beyond merely a good discriminator.

In section 3, we apply this model to the two quark and gluon datasets (Komiske et al., 2019a), (Pathak et al., 2019) and obtain the different MLE point parameters and derived metrics for other choices of SoftDrop hyperparameters. We then go beyond the point estimate calculation by introducing priors for π and λ and obtain the corresponding posterior through numerical Bayesian inference in section 4. These priors can encode our theoretical domain knowledge, such as the LL estimates of λ_q and λ_g and also regularize our model and thus avoid mode degeneracy and mode collapse.

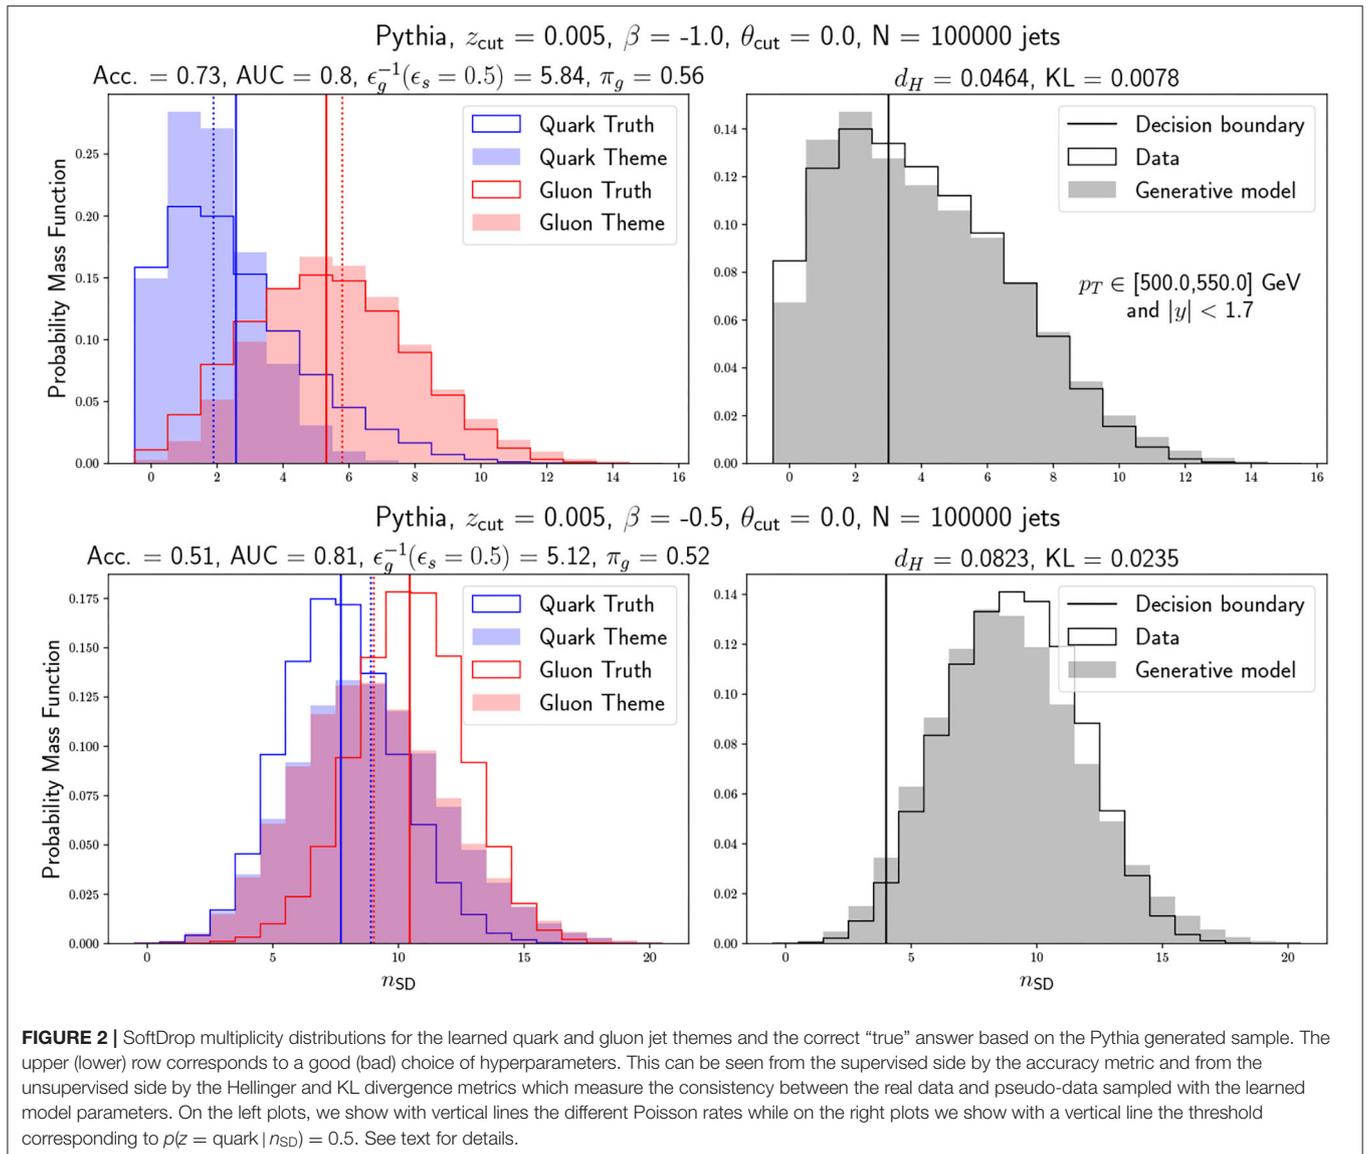
3. RESULTS

Having detailed the data and our model in section 2, we proceed to obtain point estimates for the parameters of the mixture model. In section 3.1, we study the model performance at the generator level for the two generator choices available, and we include a brief study of detector levels in section 3.2. All results are reported on a test set which was separated from the train set prior to the model training.

3.1. Model Performance at Generator Level

As detailed in section 2, we model the n_{SD} distribution as originating from a mixture of two Poissonians, which we aim to identify with gluon and quark jets (or, should we want to get rid of perturbative definitions, to operationally define gluon or quark enriched samples). For each choice of hyperparameters, we obtain the Maximum Likelihood Estimates (MLE) of the rates of the Poissonians, λ_g^{MLE} and λ_q^{MLE} , and the mixing fraction between the two π_g^{MLE} . We define the gluon theme as the theme with the larger rate, as oriented by the perturbative calculations. We obtain the MLE of the parameters with the help of the Pyro package (Bingham et al., 2018; Phan et al., 2019), which we have verified to coincide with the results obtained through Expectation-Maximization but provide us with a more flexible framework that can incorporate additional features to the generative model and optimize the code appropriately.

For the sake of validation and understanding, we consider three supervised metrics: the accuracy using the probabilistic decision boundary $p(g | n_{\text{SD}}) = p(q | n_{\text{SD}})$, the inverse gluon mistag rate at 50% quark efficiency $\epsilon_g^{-1}(\epsilon_q = 50\%)$ and the AUC. Since we are dealing with an unsupervised algorithm, and as discussed above, we also apply the unsupervised metrics defined in section 2: the Hellinger distance (Deza and Deza, 2009) and the Kullback-Leibler divergence. These metrics compare the



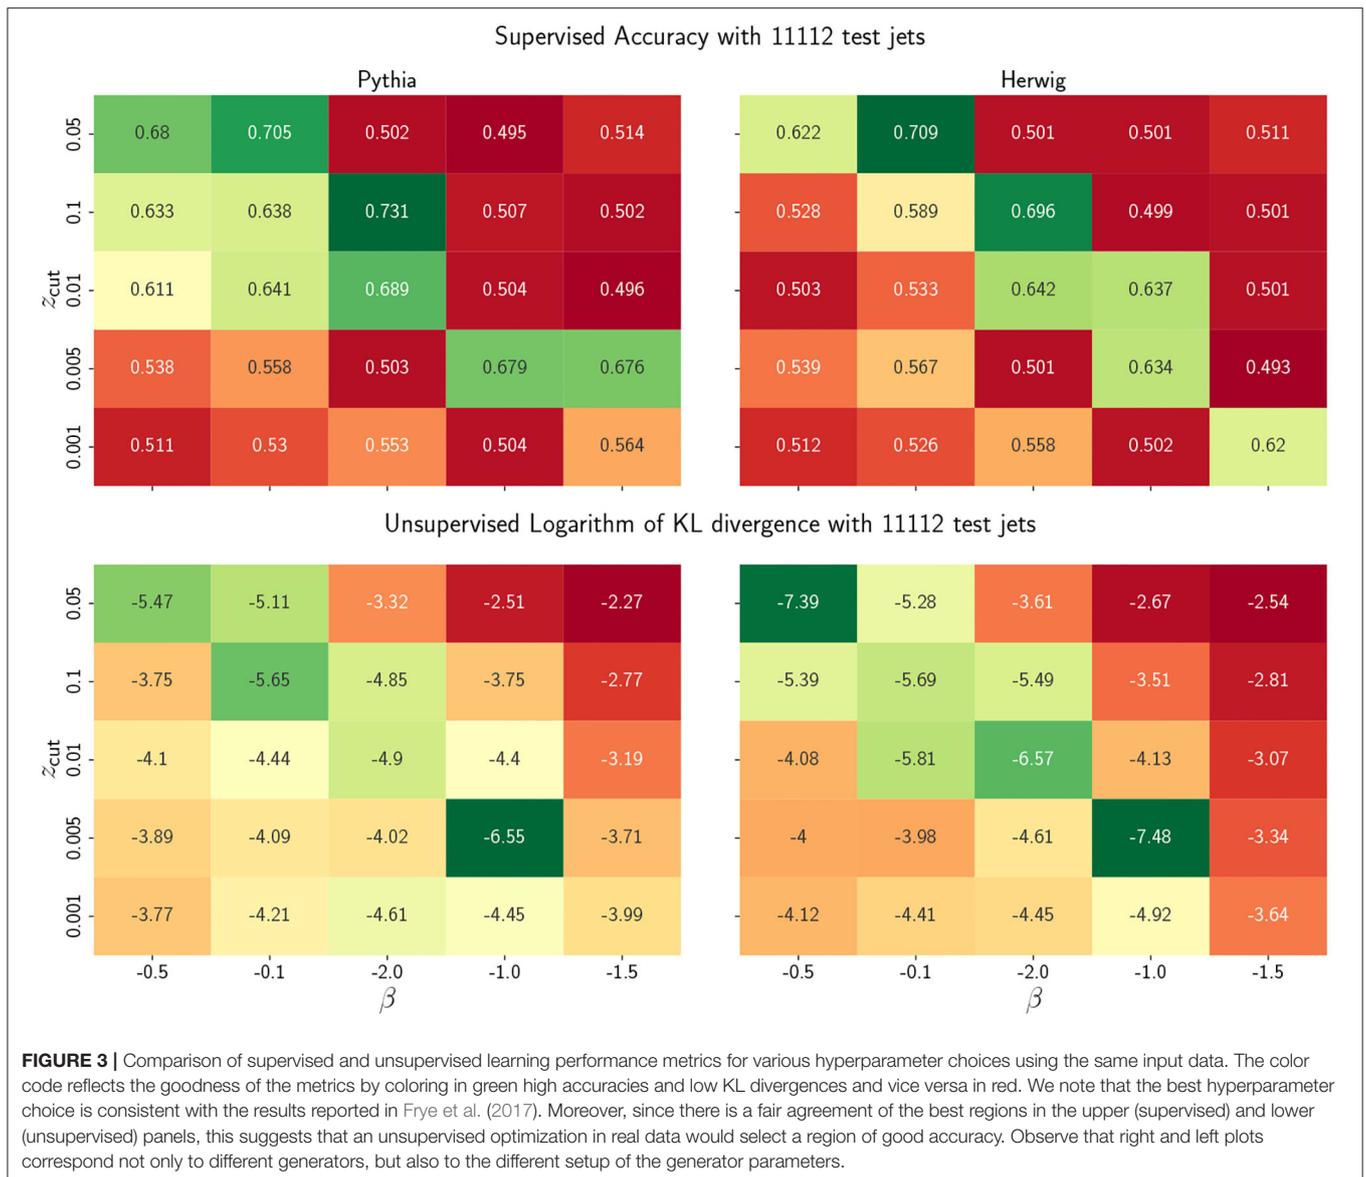
measured n_{SD} distribution (without any labels) to the learned total distribution

$$p(n_{\text{SD}} | \text{data}) = \pi_g^{\text{MLE}} \text{Poisson}(\lambda_g^{\text{MLE}}) + (1.0 - \pi_g^{\text{MLE}}) \text{Poisson}(\lambda_q^{\text{MLE}}). \quad (5)$$

We show two examples of the results we obtain for different SoftDrop hyperparameters in **Figure 2**. In the left column, we show the true underlying distributions and the two learned Poissonians, their respective means and Poissonian rates, and the supervised metrics. In the right column, we show the data distribution, the learned data distribution and the default decision boundary, along with the unsupervised data-driven metrics. In the top row, we show a good hyperparameter choice that leads to a data distribution that is well modeled by a mixture of Poissonians, and thus we obtain good supervised and unsupervised metrics. In the bottom row, we show a bad hyperparameter choice

leading to a data distribution that is not well modeled by a mixture of Poissonians, and thus we obtain mostly bad supervised and unsupervised metrics, with the exception of the AUC score.

The supervised metrics show that the accuracy and the AUC do not necessarily favour the same models. As shown in **Figure 2**, two very different cases can lead to high AUC, with the accuracy being able to reflect more the true performance of the model. This is due to the fact that the AUC is a more global metric which takes into account every possible threshold in $p(z = \text{quark} | n_{\text{SD}})$ including the default threshold used for computing the accuracy, $p(z = \text{quark} | n_{\text{SD}}) = 0.5$, and can be fooled by moving said threshold. Because the default threshold is theoretically well-motivated, as it takes full advantage of the probabilistic modeling to define a specific boundary between the two classes, it tends better to reflect the goodness of the modeling than the AUC. As we use probabilistic models for an unsupervised task, interpretability and consistency are important features to keep



in mind. In that sense, the accuracy is more aligned with the unsupervised metrics, which cannot be fooled by moving the decision threshold. The Hellinger distance and the KL divergence see whether the generated dataset is consistent with the measured dataset, taking advantage of the generative procedure.

As a next step, we scan the hyperparameter values to study the algorithm performance and how unsupervised metrics can assist us in having a good (unseen) supervised metric. We show the accuracy and the KL divergence for an array of hyperparameter values in **Figure 3**. We observe that the accuracy and the KL divergence have a fair agreement in qualifying a good model for a given SoftDrop parameter choice. Although their respective maximum and minimum do not match exactly, the regions of high accuracy coincide with the regions of low KL divergence. Therefore, we can trust that the accuracy will be increased for a reasonable parameter choice by inspecting the KL divergence and

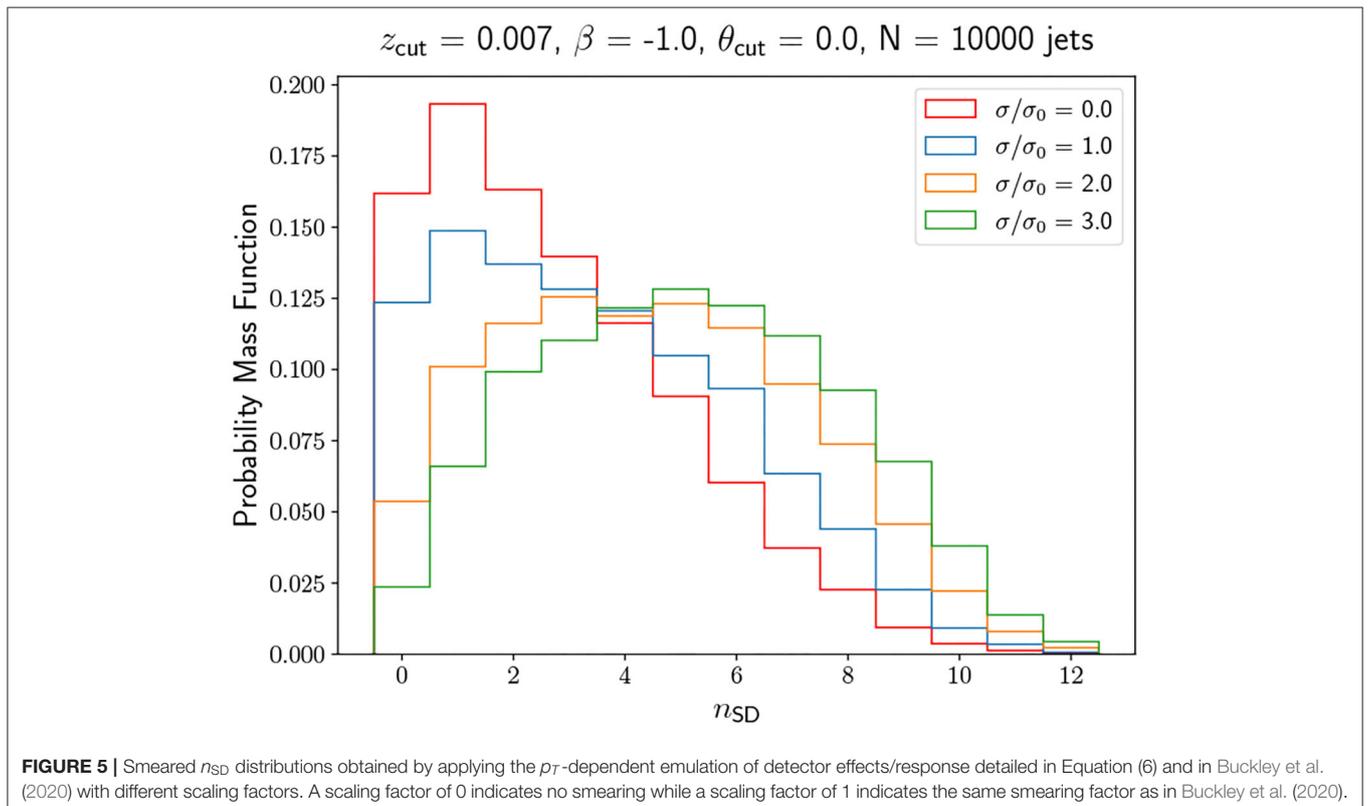
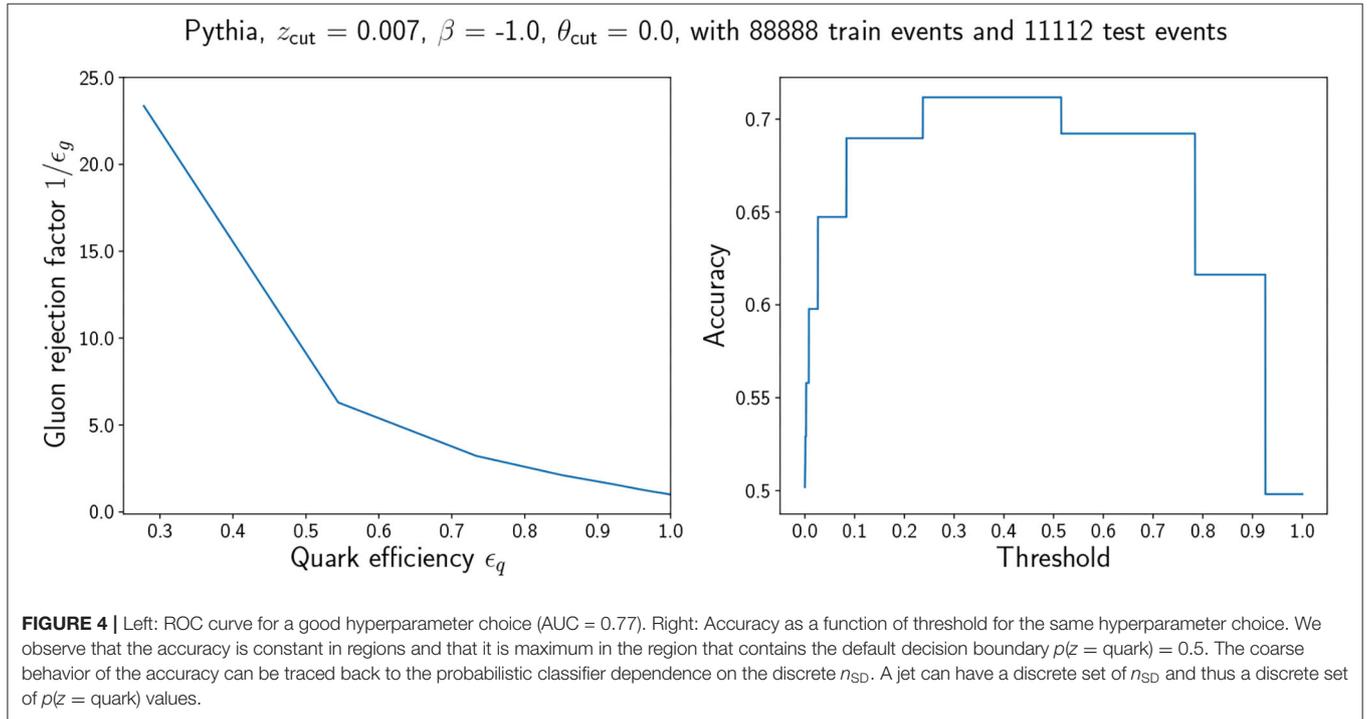
verifying that the obtained quark and gluon themes are suitable. As for the other metrics (not shown in the plot), we find that the Hellinger distance is consistent with the KL divergence and the mistag rate is consistent with the accuracy. The AUC presents the caveat discussed above and thus is less relevant for this study.

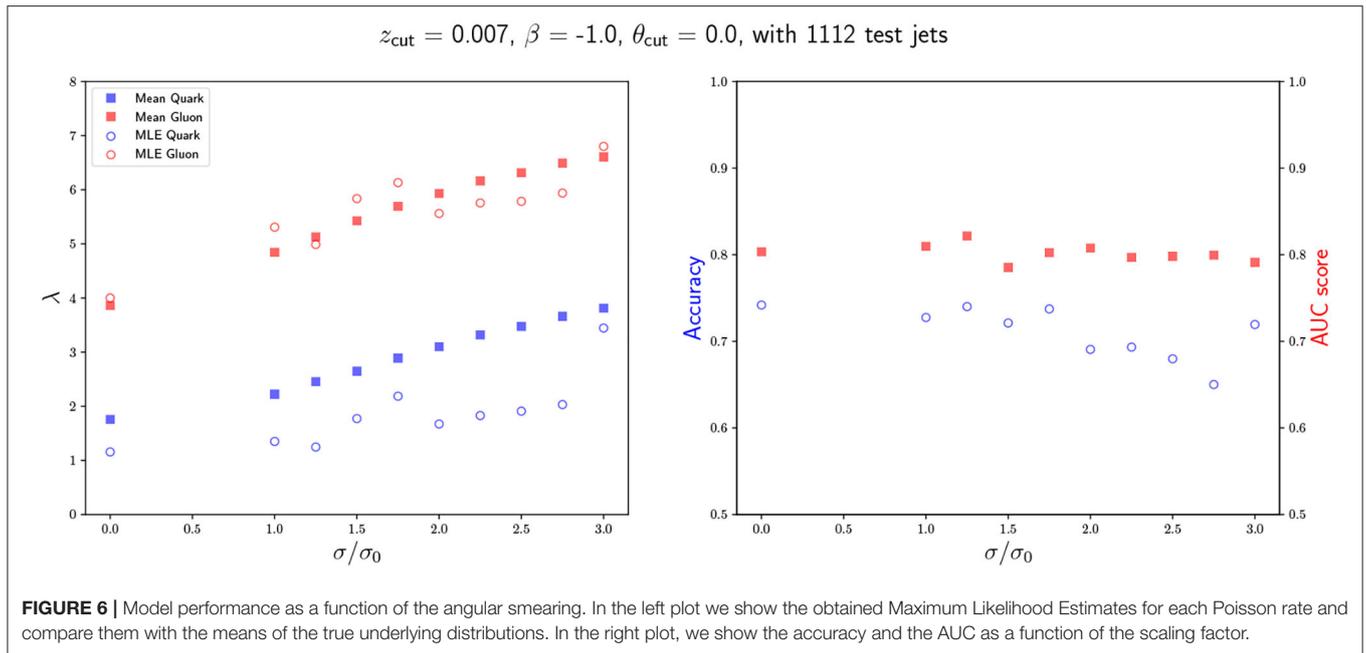
From the above results, we see that, by choosing the accuracy as the relevant metric, there is a significant overlap of the good regions in hyperparameter space according to the unsupervised and supervised metrics. This indicates that classification performance coincides with generative performance, and therefore opens the door for exploring a fully unsupervised approach where the quark/gluon tagging is defined by a relatively simple parameter scan—yielding an unsupervised, interpretable and simple model for classification.

To study the performance of the proposed unsupervised classifier in some more detail, we compute the ROC and

the accuracy as a function of the threshold c to which the classes are defined. We show in **Figure 4** both results for a good point in hyperparameter space. In a real case scenario,

one would only have access to the bottom panel in **Figure 3**, and choosing a point with small KL divergence would yield a tagger that for the threshold $p(z = \text{quark}) = 0.5$





has an accuracy of roughly ~ 0.65 – 0.73 (Pythia) and 0.62 – 0.70 (Herwig).

3.2. Detector Effects

To study the sensitivity of our tagger to detector effects, we used the procedure outlined in section 6 in Buckley et al. (2020) and smeared the η, ϕ distribution of each jet constituent. We considered the same smearing as in Buckley et al. (2020), where they spread the η and ϕ values of a constituent by sampling Gaussian noise with mean zero and standard deviation given by:

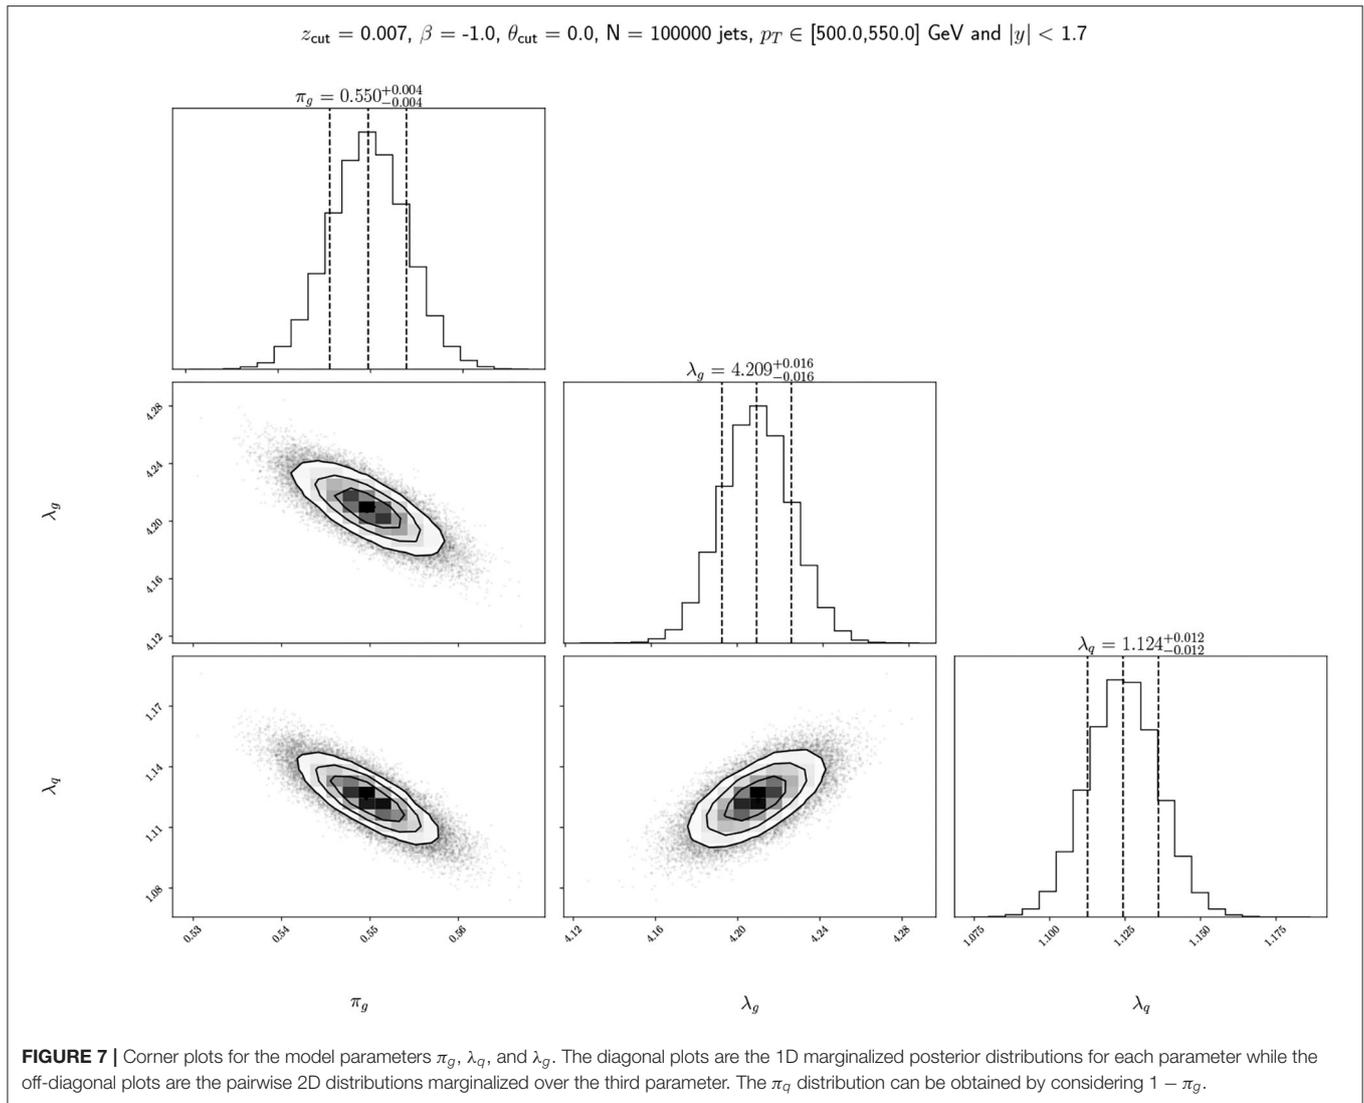
$$\sigma_0(p_T) = \frac{0.028}{1 + e^{(p_T - 25 \text{ GeV})/0.1 \text{ GeV}}}, \quad (6)$$

where p_T is the p_T of the constituent. We consider different smearing noise factors σ obtained by re-scaling σ_0 by a global multiplicative factor. The results are shown in Figures 5, 6. Although there is a difference in the MLE due to the change in the distributions, we find no significant alteration in the supervised metrics, and hence in the model performance. It seems that generator effects as simulated are not challenging the model. This may not be surprising since the model only relies on the assumptions that the integer value n_{SD} are composed of a mixture of approximately Poissonian distributions. In any case, a more realistic detector simulation should be implemented to verify this analysis, which includes modeling the energy response of various jet constituents, should be implemented. A different and interesting extension is to extend the $|y|$ range to include forward jets. For forward jets, the detector granularity changes and thus the n_{SD} becomes more dependent on $|y|$. It would then be necessary to introduce similar strategies as the ones detailed for dealing with a large p_T range. Finally, we should mention that potential pile-up issues would only

have a minor effect on the tagger, not only because n_{SD} is a robust observable as it discards soft emission, but also because we are keeping a small radius ($R = 0.4$) in the jet clustering algorithm.

4. BAYESIAN ANALYSIS

As a final study for the model, we perform Bayesian inference to obtain the full posterior probability density function over the parameters. Introducing uniform priors and performing numerical Bayesian inference, we obtain the posterior probabilities of π and λ . In order to achieve this goal, we employ the dedicated `emcee` package (Foreman-Mackey et al., 2013). We show the resulting corner plot for a justified hyperparameter choice in Figure 7. Because we have so many jets and we consider uniform priors, the inference is likelihood dominated with a prominent posterior peak in the MLE. However, one should not lose sight of the fact that the posterior distribution includes more information than the MAP point estimates since we can quantify the uncertainty of the π and λ estimation and their correlation. Suppose the generative model for the data is precise enough. In that case, this is a potentially useful application as one could establish a distance between the data-driven posterior distribution and the different Monte Carlo tunes one needs to consider to relate data with Standard Model predictions. In the specific case of the LL approximation for the n_{SD} distribution as Poissonians, we find by inspecting the MC labeled data that the agreement is not good enough to perform such a task (see Figure 8). However, the approximation is good enough to distinguish the quark from the gluon n_{SD} distributions, and therefore to create a good unsupervised classifier.



Another feature of Bayesian computation is that we can compute the probability of a given measurement n_{SD} belonging to class z integrated over the λ_g , λ_q , and π_g posterior distribution. Using our Monte Carlo samples, we calculate

$$p(z | n_{\text{SD}}, X) \approx \frac{1}{T} \sum_{t=1}^T p(z | n_{\text{SD}}, \pi_g^{(t)}, \lambda_g^{(t)}, \lambda_q^{(t)}) \quad (7)$$

where X represents the training dataset and t is the posterior sample index. We show this probability for both classes in **Figure 9**. Although the MLE dominates the likelihood because of our uniform priors and the amount of data, this probability is a more solid estimate when we only care about classifying samples as it considers all possible values of the underlying model parameters weighted by previous measurements through the posterior. The performance of this tagger using the decision threshold of $p(z = \text{quark}) = 0.5$ yields an accuracy of 0.71.

5. DISCUSSION AND OUTLOOK

We have proposed an unsupervised data-driven learning algorithm to classify jets induced by quarks or gluons. The key of the method is to approximate that each class (quark and gluon) has a Poissonian distribution with a different rate for the jets' Soft Drop observable, n_{SD} . Therefore, the n_{SD} distribution of a sample of an unknown mixture of quark and gluon jets correspond to a mix of Poissonians. This observation, which is only for approximately constant jet p_T , allows to set up an unsupervised learning paradigm that can extract the Maximum Likelihood Estimate (MLE) and the posterior distributions for the rate of each Poissonian (λ_q and λ_g) and the fraction of each constituent in the sample ($\pi_{q,g}$) and thus, with this knowledge, one can create a tagger to discriminate quark and gluon induced jets. This is all achieved without relying on any Monte Carlo generator, nor any previous knowledge other than the assumption that n_{SD} is Poisson distributed for each class. We use the basic

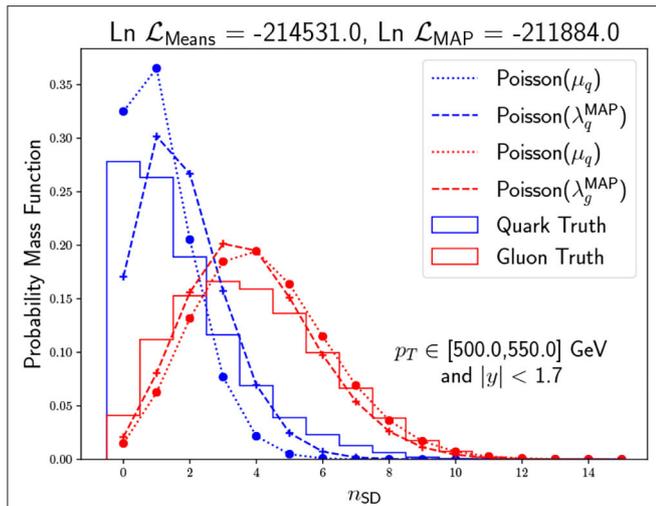


FIGURE 8 | n_{SD} distribution comparison between pseudo-data generated through MC (solid) and Poissonians estimates using as rates the mean of the data (dashed) and the Maximum a Posteriori MAP from the Bayesian inference (dashed). We see that the Poissonian approximations are good enough to distinguish quark from gluon, but there are slight differences when comparing each approximation to its corresponding data.

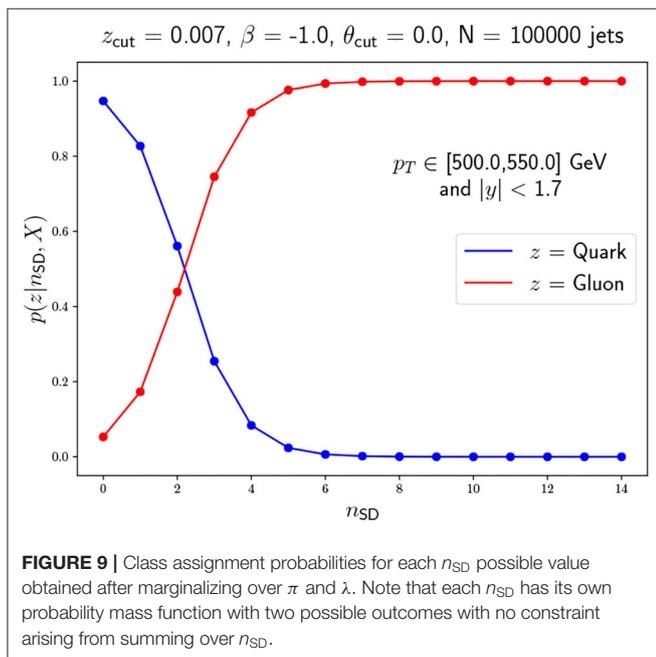


FIGURE 9 | Class assignment probabilities for each n_{SD} possible value obtained after marginalizing over π and λ . Note that each n_{SD} has its own probability mass function with two possible outcomes with no constraint arising from summing over n_{SD} .

principle knowledge that $\lambda_g > \lambda_q$ to assign the tagging of the reconstructed themes.

In the first part of the work, we have defined the generative process of the data according to the above hypothesis and obtained the MLE for the parameters using Stochastic Variational Inference and Expectation-Maximization techniques independently. We have then designed a quark-gluon tagger and discussed a method to find the best hyperparameters choice for

the SoftDrop algorithm that optimise the tagger accuracy. Since one cannot measure the accuracy in actual data because one does not have access to the labels, we have shown that minimizing the KL divergence between the real data and generated data sampled with the generative model improves the tagger accuracy. We have verified that the procedure works for different Monte Carlo with different tunes. One can expect that the described unsupervised tagger can have an accuracy in the range $\approx 0.65 - 0.70$.

We have performed a simple detector effect simulation by smearing the angular coordinate of each jet constituent, and we find that the tagger accuracy remains approximately the same. This is not surprising since, despite the detector effects, the n_{SD} observable is still a counting observable that may vary its value but still be Poissonian distributed with shifted rates. Therefore the whole machinery of the unsupervised algorithm works essentially the same.

In a second part of the article, we have performed a Bayesian inference on the parameters to extract the full posterior distribution and the correlation between the model parameters, namely λ_q , λ_g , and π_g . In particular, we have found that the Maximum a Posteriori (MAP) approximately coincides with the MLE of the parameters. Furthermore, we have found that although the reconstructed Poissons for each class does not match the labeled data within the posterior uncertainty, the classifier still works quite good. The reason for this is that, although we can see a slight departure of the approximation of the n_{SD} being Poissonian distributed, the two inferred Poissonians for quark and gluon still show a more pronounced difference between them than its corresponding labeled data.

With the posterior obtained through Bayesian inference, we have designed a quark-gluon tagger based on computing the probability of a jet being induced by either quark or gluon using all the observed data. This is a more robust tagger since it sees the posterior and hence the correlation between the parameters rather than the point MLE. With this tagger, we obtain an accuracy of 0.71.

There are potential improvements and limitations on the proposed algorithm. For example, suppose one could have a model for the n_{SD} that goes beyond the LL Poissonian approximation. In that case, one could modify the likelihood and obtain the posterior for the new likelihood parameters. Although we do not expect this to improve the tagger accuracy considerably, it could help tune a Monte Carlo using unsupervised learning. If one could have a reliable posterior for specific signal distribution, then one could check whether a Monte Carlo is compatible or not with it. Observe that, since Monte Carlo generators do not have a handle to set the value for each observable, having a prediction for some observable and its uncertainty provides the necessary information to check whether the Monte Carlo sampling is within the allowed regions defined by the posterior. On other aspects, we have performed simple modeling for the detector effects, which apparently would not affect the tagger performance. Further investigation in this direction would be helpful to find the actual limitations of the algorithm.

Finally, we should comment on the challenges that may arise when applying this algorithm in real data. A balanced

quark/gluon dataset is far from guaranteed. However, we have verified that the classification and generative powers of the model are robust against a change in the classes fractions up to a 80% in any class. There is also the possibility of sample contamination with, for example, charm- and bottom-quarks. If there is no need to disentangle charm- and bottom- from light-quarks, then no modification is needed as n_{SD} is mostly agnostic to quark flavor for relatively fixed jet kinematics. In particular, for jets with $p_T \gg 5$ GeV c - and b -jets are as massless as light-jets and they have a similar n_{SD} behavior. If b -tagging and c -tagging is needed, then the model should be extended by incorporating other observables which are sensitive to quark flavor, like the number of displaced vertices in the jet, before searching for three themes instead of two.

Current supervised algorithms to discriminate jets induced by quark or gluon have a non-negligible dependence on Monte Carlo and their tunes, which may hide some intractable systematic uncertainties or biases. Therefore, we find that proposing an unsupervised paradigm for quark-gluon determination is an appealing road that should be transited. In addition to being interpretable and straightforward, the

presented algorithm yields an accuracy in the 0.65–0.7 range, which is a good achievement for the small number of assumptions on which it relies.

DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author/s.

AUTHOR CONTRIBUTIONS

All authors listed have made a substantial, direct, and intellectual contribution to the work and approved it for publication.

ACKNOWLEDGMENTS

MS would like to thank the Jozef Stefan Institute for its warm hospitality during part of this work. We thank Referees for relevant and useful suggestions which were included in this revised version.

REFERENCES

- Aarrestad, T., van Beekveld, M., Bona, M., Boveia, A., Caron, S., Davies, J., et al. (2021). The dark machines anomaly score challenge: benchmark data and model independent event classification for the large hadron collider. *Sci. Post Phys.* 12, 043. doi: 10.21468/SciPostPhys.12.1.043
- Alvarez, E., Dillon, B. M., Faroughy, D. A., Kamenik, J. F., Lamagna, F., and Szewc, M. (2021). Bayesian probabilistic modelling for four-tops at the LHC. doi: 10.48550/arXiv.2107.00668
- Alvarez, E., Lamagna, F., and Szewc, M. (2020). Topic model for four-top at the LHC. *JHEP* 1, 049. doi: 10.1007/JHEP01(2020)049
- Andreassen, A., Feige, I., Frye, C., and Schwartz, M. D. (2019). JUNIPR: a framework for unsupervised machine learning in particle physics. *Eur. Phys. J. C* 79, 102. doi: 10.1140/epjc/s10052-019-6607-9
- Andreassen, A., Nachman, B., and Shih, D. (2020). Simulation assisted likelihood-free anomaly detection. *Phys. Rev. D* 101, 095004. doi: 10.1103/PhysRevD.101.095004
- ATLAS coll. (2014). Light-quark and gluon jet discrimination in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector. *Eur. Phys. J. C* 74.8, 3023. doi: 10.1140/epjc/s10052-014-3023-z
- Bahr, M., Gieseke, S., and Gigg, M. A. (2008). Herwig++ physics and manual. *Eur. Phys. J. C* 58, 639–707. doi: 10.1140/epjc/s10052-008-0798-9
- Bellm, J., Gieseke, S., Grellscheid, D., Platzer, S., Rauch, M., Reuschle, C., et al. (2016). Herwig 7.0/herwig++ 3.0 release note. *Eur. Phys. J. C* 76, 196. doi: 10.1140/epjc/s10052-016-4018-8
- Bhattacharjee, B., Mukhopadhyay, S., Nojiri, M. M., Sakaki, Y., and Webber, B. R. (2015). Associated jet and subjet rates in light-quark and gluon jet discrimination. *JHEP* 4, 131. doi: 10.1007/JHEP04(2015)131
- Bhattacharjee, B., Mukhopadhyay, S., Nojiri, M. M., Sakaki, Y., and Webber, B. R. (2017). Quark-gluon discrimination in the search for gluino pair production at the LHC. *JHEP* 1, 044. doi: 10.1007/JHEP01(2017)044
- Bingham, E., Chen, J. P., Jankowiak, M., Obermeyer, F., Pradhan, N., Karaletsos, T., et al. (2018). Pyro: deep universal probabilistic programming. *J. Mach. Learn. Res.* 20, 973–978.
- Bishop, C. M. (2013). *Pattern Recognition and Machine Learning*. Heidelberg: Springer.
- Buckley, A., Kar, D., and Nordström, K. (2020). Fast simulation of detector effects in Rivet. *SciPost Phys.* 8, 025. doi: 10.21468/SciPostPhys.8.2.025
- Caron, S., Hendriks, L., and Verheyen, R. (2021). Rare and different: anomaly scores from a combination of likelihood and out-of-distribution models to detect new physics at the LHC. *SciPost Phys.* 12, 077. doi: 10.21468/SciPostPhys.12.2.077
- Celeux, G., Frühwirth-Schnatter, S., and Robert, C. (2018). “Model selection for mixture models—perspectives and strategies,” in *Handbook of Mixture Analysis*. CRC Press.
- Choi, S., Lim, J., and Oh, H. (2020). Data-driven estimation of background distribution through neural autoregressive flows. doi: 10.48550/arXiv.2008.03636
- CMS coll. (2013). *Performance of Quark/Gluon Discrimination in 8 TeV pp Data*. <https://cds.cern.ch/record/1599732?ln=es>
- CMS coll. (2015). Search for dark matter, extra dimensions, and unparticles in monojet events in proton–proton collisions at $\sqrt{s} = 8$ TeV. *Eur. Phys. J. C* 75.5, 235. doi: 10.1140/epjc/s10052-015-3451-4
- d’Agnolo, R. T., Grosso, G., Pierini, M., Wulzer, A., and Zanetti, M. (2021). Learning new physics from an imperfect machine. doi: 10.48550/arXiv.2111.13633
- Deza, M. M., and Deza, E. (2009). *Encyclopedia of Distances*. (Berlin: Springer), 1–583.
- Dillon, B. M., Faroughy, D. A., and Kamenik, J. F. (2019). Uncovering latent jet substructure. *Phys. Rev. D* 100, 056002. doi: 10.1103/PhysRevD.100.056002
- Dillon, B. M., Faroughy, D. A., Kamenik, J. F., and Szewc, M. (2020). Learning the latent structure of collider events. *JHEP* 10, 206. doi: 10.1007/JHEP10(2020)206
- Dillon, B. M., Faroughy, D. A., Kamenik, J. F., and Szewc, M. (2021). Learning latent jet structure. *Symmetry* 13, 1167. doi: 10.3390/sym13071167
- Dohi, K. (2020). Variational autoencoders for jet simulation. doi: 10.48550/arXiv.2009.04842
- Dokshitzer, Y. L., Khoze, V. A., and Sjostrand, T. (1992). Rapidity gaps in higgs production. *Phys. Lett. B* 274, 116–121.
- Dokshitzer, Y. L., Leder, G. D., Moretti, S., and Webber, B. R. (1997). Better jet clustering algorithms. *JHEP* 08, 001.
- Ferreira de Lima, D., Petrov, P., Soper, D., and Spannowsky, M. (2017). Quark-Gluon tagging with shower deconstruction: unearthing dark matter and Higgs couplings. *Phys. Rev. D* 95, 034001. doi: 10.1103/PhysRevD.95.034001
- Foreman-Mackey, D., Hogg, D. W., Lang, D., and Goodman, J. (2013). emcee: the mcmc hammer. *Publ. Astron. Soc. Pac.* 125, 306–312. doi: 10.1086/670067
- Frye, C., Larkoski, A. J., Thaler, J., and Zhou, K. (2017). Casimir meets poisson: improved quark/gluon discrimination with counting observables. *JHEP* 09, 083. doi: 10.1007/JHEP09(2017)083
- Galluccio, J., and Schwartz, M. D. (2011). Quark and Gluon tagging at the LHC. *Phys. Rev. Lett.* 107, 172001. doi: 10.1103/PhysRevLett.107.172001

- Graziani, G., Anderlini, L., Mariani, S., Franzoso, E., Pappalardo, L., and di Nezza, P. (2021). A neural-network-defined Gaussian mixture model for particle identification applied to the LHCb fixed-target programme. doi: 10.1088/1748-0221/17/02/P02018
- Hajer, J., Li, Y.-Y., Liu, T., and Wang, H. (2020). Novelty detection meets collider physics. *Phys. Rev. D* 101, 076015. doi: 10.1103/PhysRevD.101.076015
- Kasieczka, G., Kiefer, N., Plehn, T., and Thompson, J. M. (2019). Quark-Gluon tagging: machine learning vs detector. *SciPost Phys.* 6, 069. doi: 10.21468/SciPostPhys.6.6.069
- Kasieczka, G., Nachman, B., Shih, D., Amram, O., Andreassen, A., Benkendorfer, K., Bortolato, B., et al. (2021). The LHC olympics 2020: a community challenge for anomaly detection in high energy physics. *Rep. Prog. Phys [preprint]*. 84. doi: 10.1088/1361-6633/ac36b9
- Komiske, P., Metodiev, E., and Thaler, J. (2019a). Pythia8 quark and gluon jets for energy flow. *Zenodo*. doi: 10.5281/zenodo.3164691
- Komiske, P. T., Metodiev, E. M., Nachman, B., and Schwartz, M. D. (2018a). Learning to classify from impure samples with high-dimensional data. *Phys. Rev. D* 98, 011502. doi: 10.1103/PhysRevD.98.011502
- Komiske, P. T., Metodiev, E. M., and Thaler, J. (2018b). An operational definition of quark and gluon jets. *JHEP* 11, 059. doi: 10.1007/JHEP11(2018)059
- Komiske, P. T., Metodiev, E. M., and Thaler, J. (2019b). Energy flow networks: deep sets for particle jets. *JHEP*, 01:121. doi: 10.1007/JHEP01(2019)121
- Komiske, P. T., Metodiev, E. M., and Thaler, J. (2019c). Metric space of collider events. *Phys. Rev. Lett.* 123, 041801. doi: 10.1103/PhysRevLett.123.041801
- Larkoski, A. J., Marzani, S., Soyer, G., and Thaler, J. (2014a). Soft drop. *JHEP* 05, 146. doi: 10.1007/JHEP05(2014)146
- Larkoski, A. J., Salam, G. P., and Thaler, J. (2013). Energy correlation functions for jet substructure. *JHEP* 06, 108. doi: 10.1007/JHEP06(2013)108
- Larkoski, A. J., Thaler, J., and Waalewijn, W. J. (2014b). Gaining (mutual) information about quark/gluon discrimination. *JHEP* 11, 129. doi: 10.1007/JHEP11(2014)129
- Metodiev, E. M., Nachman, B., and Thaler, J. (2017). Classification without labels: learning from mixed samples in high energy physics. *JHEP* 10, 174. doi: 10.1007/JHEP10(2017)174
- Metodiev, E. M., and Thaler, J. (2018). Jet topics: disentangling quarks and gluons at colliders. *Phys. Rev. Lett.* 120, 241602. doi: 10.1103/PhysRevLett.120.241602
- Nachman, B., and Shih, D. (2020). Anomaly detection with density estimation. *Phys. Rev. D* 101, 075042. doi: 10.1103/PhysRevD.101.075042
- Pathak, A., Komiske, P., Metodiev, E., and Schwartz, M. (2019). Herwig7.1 quark and gluon jets. *Zenodo*. doi: 10.5281/zenodo.3066475
- Phan, D., Pradhan, N., and Jankowiak, M. (2019). Composable effects for flexible and accelerated probabilistic programming in numpyro. doi: 10.48550/arXiv.1912.11554
- Rainwater, D. L., Zeppenfeld, D., and Hagiwara, K. (1998). Searching for $H \rightarrow \tau^+ \tau^-$ in weak boson fusion at the CERN LHC. *Phys. Rev. D* 59, 014037. doi: 10.1103/PhysRevD.59.014037
- Roy, T. S., and Vijay, A. H. (2020). A robust anomaly finder based on autoencoders. Sjöstrand, T., Ask, S., Christiansen, J. R., Corke, R., Desai, N., Ilten, P., et al. (2015). An introduction to PYTHIA 8.2. *Comput. Phys. Commun.* 191, 159–177. doi: 10.1016/j.cpc.2015.01.024
- Štěpánek, M., Franc, J., and Kús, V. (2015). Modification of gaussian mixture models for data classification in high energy physics. *J. Phys. Conf. Series* 574, 012150. doi: 10.1088/1742-6596/574/1/012150

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2022 Alvarez, Spannowsky and Szewc. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.