

The information content of jet quenching and machine learning assisted observable design

Yue Shi Lai,^a James Mulligan,^{a,b} Mateusz Płoskoń^a and Felix Ringer^{a,c,d}

^a*Nuclear Science Division, Lawrence Berkeley National Laboratory,
Berkeley, CA 94720, U.S.A.*

^b*Physics Department, University of California,
Berkeley, CA 94720, U.S.A.*

^c*C.N. Yang Institute for Theoretical Physics, Stony Brook University,
Stony Brook, NY 11794, U.S.A.*

^d*Department of Physics and Astronomy, Stony Brook University,
Stony Brook, NY 11794, U.S.A.*

E-mail: ylai@lbl.gov, james.mulligan@berkeley.edu, mploskon@lbl.gov,
felix.ringer@stonybrook.edu

ABSTRACT: Jets produced in high-energy heavy-ion collisions are modified compared to those in proton-proton collisions due to their interaction with the deconfined, strongly-coupled quark-gluon plasma (QGP). In this work, we employ machine learning techniques to identify important features that distinguish jets produced in heavy-ion collisions from jets produced in proton-proton collisions. We formulate the problem using binary classification and focus on leveraging machine learning in ways that inform theoretical calculations of jet modification: (i) we quantify the information content in terms of Infrared Collinear (IRC)-safety and in terms of hard vs. soft emissions, (ii) we identify optimally discriminating observables that are in principle calculable in perturbative QCD, and (iii) we assess the information loss due to the heavy-ion underlying event and background subtraction algorithms. We illustrate our methodology using Monte Carlo event generators, where we find that important information about jet quenching is contained not only in hard splittings but also in soft emissions and IRC-unsafe physics inside the jet. This information appears to be significantly reduced by the presence of the underlying event. We discuss the implications of this for the prospect of using jet quenching to extract properties of the QGP. Since the training labels are exactly known, this methodology can be used directly on experimental data without reliance on modeling. We outline a proposal for how such an experimental analysis can be carried out, and how it can guide future measurements.

KEYWORDS: Heavy Ion Phenomenology, Jets

ARXIV EPRINT: [2111.14589](https://arxiv.org/abs/2111.14589)

Contents

1	Introduction	1
2	Monte Carlo jet sample	4
3	Jet classification: pp vs. AA	5
3.1	IRC-safe vs. IRC-unsafe information content	6
3.1.1	Particle and Energy Flow Networks	6
3.2	Hard vs. soft information content	9
3.2.1	N -subjettiness basis	9
3.2.2	Energy Flow Polynomial basis	12
4	Observable design	14
4.1	Lasso regression	15
5	Information loss: the underlying event and background subtraction	19
6	Application to experimental data	22
7	Conclusions and outlook	23

1 Introduction

Jets are highly energetic and collimated groups of particles observed in the detectors of high-energy scattering experiments such as the Large Hadron Collider (LHC) and the Relativistic Heavy Ion Collider (RHIC). Jets reflect the underlying quark and gluon degrees of freedom of hard-scattering processes, in which the scattered partons fragment into a shower due to multiple soft and collinear emissions. These partons subsequently hadronize into the color neutral particles observed in detectors. Studies of the radiation pattern inside jets, known as jet substructure, aim to quantify and utilize jets to better understand quantum chromodynamics (QCD) and search for physics beyond the Standard Model [1–3]. Jets and their substructure have been studied both in proton-proton (pp) and heavy-ion (AA) collisions.

In high-energy heavy-ion collisions a state of matter known as the quark-gluon plasma (QGP) is formed, in which quarks and gluons are deconfined [4]. The QGP is believed to have populated the universe for most of the first few microseconds after the Big Bang [5], and today serves as a laboratory to study the emergence of a high temperature, strongly-coupled system from quarks and gluons in QCD. While several of its transport coefficients have been constrained by experimental measurements, the detailed properties of the QGP remain unknown. Information about the QGP can be obtained by comparing jets in

vacuum to their counterparts in heavy-ion collisions which have traversed the hot and dense nuclear matter. The modifications of the jets emerging after interacting with the QGP as compared to jets in pp collisions is known as *jet quenching*. In an ongoing effort to achieve a unified description of jet modification a significant experimental and theoretical effort has been made to measure and compute the medium-induced modifications of various jet observables [6–28].

Several guiding principles have been used to identify suitable observables in order to extract properties of the QGP. First, observables insensitive to collinear and soft emissions are given prominence — the so-called Infrared Collinear (IRC) safe observables. IRC safety (or by extension Sudakov safety [29–31]) ensures that they can be calculated in perturbation theory order-by-order and that they are less sensitive to experimental resolution effects. While nonperturbative effects can still be important, IRC safe observables do not require nonperturbative quantities such as parton-to-hadron fragmentation functions. Often observables that are under good theoretical control and those with a relatively simple factorization structure are preferred. This allows for precision computations to high perturbative accuracy and in several cases it allows for predictions on how nonperturbative contributions scale [32–39]. A second group of guiding principles for selecting an observable stems from experimental considerations where observables that can be measured well in the large fluctuating background in heavy-ion collisions are favored [27, 40, 41]. Third, observables are chosen based on their sensitivity to specific medium effects such as energy loss or p_T -broadening. This can be studied in the context of Bayesian inference (see e.g. ref. [42]), or by using machine learning to automate the design of observables that are most sensitive to model parameters such as the temperature profile of hydrodynamic simulations [43]. In this work, we explore a new guiding principle: the systematic quantification of the relative information content of quenched vs. vacuum jets.

In recent years machine learning techniques have been applied to studies in nuclear and high-energy physics [44–61]. In particular, several publications have applied machine learning methods to study jet quenching [43, 62–64]. Yet, one of the outstanding challenges in the application of machine learning to jet quenching is model dependence: supervised machine learning algorithms require training data, and when relying on Monte Carlo event generators to produce the training input, it is often unclear how to connect the results to jet quenching observed in nature.

In this work, we examine jet quenching as a classification problem — distinguishing jets in proton-proton collisions from jets in heavy-ion collisions, as illustrated in figure 1. We utilize machine learning to quantify the features and patterns that distinguish these two classes of jets. In the following we refer to the set of identified features as the *information content of jet quenching*. Rather than training the classifier purely as a black box, we focus on formulating the classifier in a way that allows us to dissect the machine-learned information in order to provide insights about jet quenching. In particular, we curate the input variables and the theoretical properties of the machine learning architecture in order to not only to estimate an upper bound on the information content in jet quenching but also to study how that information is distributed in terms of IRC safety and hard vs. soft physics. We follow this up by performing Lasso regression [65, 66] (a simple form of sym-

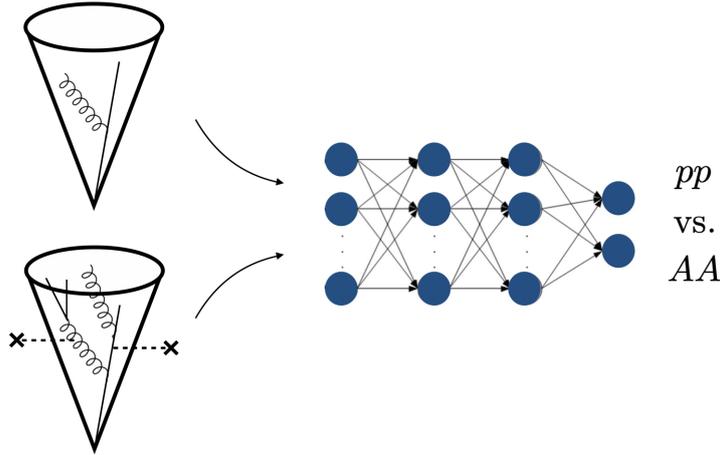


Figure 1. Schematic illustration of jets in pp (left) and heavy-ion AA (right) collisions. Interactions with the QGP can lead to a modification of the jet substructure. By training a binary classifier, the machine learns the relevant information that distinguishes jets in pp and AA collisions.

bolic regression) to design jet substructure observables that are maximally modified by the QGP. Strongly modified observables can provide guidance of how theoretical formulations of jet quenching should modify jet evolution. We identify single observables that are in principle calculable in perturbative QCD, by which we mean observables with closed-form expressions that obey IRC-safety or Sudakov safety and therefore can generally be calculated in the future with existing methods. This approach provides a handle between discrimination power and simplicity of the observable which allows for a bridge between machine-learned information and quantities that have the potential to be understood in perturbative QCD. With this method, we demonstrate that the discrimination power of our designed observables exceeds that of traditional observables already measured in heavy-ion collisions. We also quantify the information loss due to the large underlying event, one of the most significant challenges in jet quenching. For the first time, we study the information loss caused by background subtraction techniques that are traditionally employed in heavy-ion collisions.

To illustrate this methodology, we perform binary classification tasks using Monte Carlo event generators. We expect that the proposed analyses can and should ideally be carried out using event-by-event data from the experiments at the LHC and RHIC. These analyses can be performed either using uncorrected experimental data or using corrected full events [67]; these experimental considerations will be discussed further in section 6. Unlike other classification tasks in high-energy and nuclear physics, such as quark vs. gluon jet tagging, our goal is to classify two experimentally measurable distributions rather than to directly relate the measured jet properties to parton level quantities, and so there are no ambiguities concerning the labels of the classification task. One is therefore able to use fully supervised learning without reliance on modeling. We expect that the proposed data driven approach will serve as a guiding principle to construct observables in heavy-ion collisions and complement existing approaches in the literature. In addition, our approach can be

useful for the extraction of medium parameters by extending the existing set of observables to new ones which contain valuable additional information [68]. We also note that while we perform our analysis using only the particles inside identified inclusive jets, similar results can be obtained using full event information or other jet topologies. It can also be applied to (jets or events) in ep vs. eA collisions at the future Electron-Ion Collider (EIC) [69].

The remainder of our paper is organized as follows. In section 2, we discuss the event generation of the jet samples in pp and AA collisions and the treatment of the heavy-ion underlying event. In section 3, we present results examining IRC-safe information (section 3.1) and soft physics via a jet substructure basis (section 3.2) using different machine learning architectures. In section 4, we use a form of symbolic regression to design observables that approximate these classifiers. In section 5, we study the impact of the heavy-ion underlying event on jet classification. In section 6, we discuss experimental considerations which are relevant in order to perform the proposed analysis on actual data instead of parton shower simulations. We conclude and present an outlook in section 7.

2 Monte Carlo jet sample

To illustrate our methodology, we perform our analysis using jets from Monte Carlo event generators at $\sqrt{s_{\text{NN}}} = 5.02$ TeV. To generate vacuum jets, we use PYTHIA8 [70] with the Monash2013 tune [71]. To generate medium modified jets, we use JEWEL [72, 73] 2.2.0 with an initial temperature $T_i = 590$ MeV and initial quenching time $\tau_i = 0.4$ (formation time of the QGP), which provides an adequate description of a variety of jet quenching observables [74]. We do not include recoil particles in JEWEL, both for simplicity and because the physics of medium response is relatively poorly understood; we defer this to a future study using experimental data. In PYTHIA8, we disable multi-parton interactions (MPI) in order to match JEWEL.

In both cases, we reconstruct inclusive jets with the anti- k_T algorithm [75] and $R = 0.4$. We consider jets with transverse momentum p_T^{jet} in the range $100 < p_T^{\text{jet}} < 125$ GeV and pseudorapidity of $|\eta^{\text{jet}}| < 2$. We neglect particle identification (PID) and assign a pion mass assumption to all particles, leaving a study of the information conveyed by PID to future work. The p_T^{jet} range is determined from the final state particles, in both PYTHIA8 and JEWEL.

We choose to consider an inclusive jet sample because inclusive jets are the most straightforward to measure and theoretically well understood — the initial hard-scattering distribution is fixed by the quantities $(p_T^{\text{jet}}, \eta^{\text{jet}}, \sqrt{s_{\text{NN}}})$ and, hence, the allowed range of the initial parton transverse momentum \hat{p}_T^{jet} [76–79]. Our studies could be extended to disentangle medium quark/gluon fractions and jet-by-jet modifications using γ/Z -tagged jets and di-jets (which require a more complicated factorization to connect to the initial hard scattering distribution), which add sensitivity to the surface bias in heavy-ion collisions [80]), along with additional quark/gluon taggers or theoretical constraints.

In sections 3 and 4, we consider only the events generated by PYTHIA8 and JEWEL, which contain no underlying event stemming from particle production in uncorrelated scatterings in heavy-ion collisions. In section 5, we study the impact of the heavy-ion under-

lying event by combining the particles from an event simulated by PYTHIA8 or JEWEL together with the particles from a randomly drawn simulated background event. We generate a heavy-ion background based on a thermal model consisting of N particles drawn from a Gaussian with $\langle dN/d\eta \rangle \approx 2500$ and p_T sampled event-by-event from a Gamma distribution, $f_{\Gamma}(p_T; \alpha, \beta) \propto p_T^{\alpha-1} e^{-p_T/\beta}$ with $\alpha = 2$. We select $\beta = 0.4$, corresponding to $\langle p_T \rangle \approx 0.8$ GeV, in order to approximate the particle multiplicities, fluctuations, and transverse momentum observed in 0-10% central Pb-Pb data [41, 81].

After combining the PYTHIA8 or JEWEL particles together with the background particles, we make use of a background subtraction procedure to mitigate the influence of the large combinatorial background on reconstructed observables, as done in experimental analyses. We employ event-by-event constituent subtraction, which corrects the overall jet p_T and its substructure simultaneously by subtracting the underlying event energy constituent-by-constituent [82, 83]. In each event the subtraction algorithm distributes a uniform density of soft particles according to the average background density in the event, and subtracts their p_T from the original particles in the event within a maximum recombination rapidity-azimuth distance R_{\max} . We use a value of $R_{\max} = 0.25$, with other parameters set to their default values. We then perform jet reconstruction in the same manner as above, except that the p_T^{jet} range is selected based on the particles from PYTHIA8 or JEWEL alone, in order to allow comparison to studies without background.

We use a sample of approximately 2×10^5 jets, balanced evenly between pp and AA . These jet samples will serve as input to several different machine learning architectures described in the following sections. We reserve 20% of the training sample as a validation set, and an additional 20% as a test set. The size of the sample is determined by studying when the classification performance approximately saturates as the statistics are increased (noting that our model-dependent study is not meant to achieve sub-percent level optimization), and corresponds roughly to the number of jets recorded with an integrated luminosity of $\mathcal{O}(1 \text{ fb}^{-1})$ at the LHC energies. In all cases, we report classification metrics on the test set, and ensure compatible performance with the training and validation sets. In general, it is sufficient to train a per-instance (jet-by-jet) instead of a per-ensemble classifier since collider events are independently and identically distributed [84].

The performance of a classifier is typically assessed by analyzing the receiver operating characteristic (ROC) curve and the area under the ROC curve (AUC). The ROC curve shows the cumulative distribution functions of the true positive rate vs. the false positive rate as the decision threshold is varied. In our case we define “positive” to refer to AA jets. A random classifier follows a diagonal line with $\text{AUC} = 0.5$ and the better a classifier is, the closer the curve is to the upper left edge of the plot, with a perfect classifier having $\text{AUC} = 1$.

3 Jet classification: pp vs. AA

The information content of jet quenching can be studied as a binary classification problem of jets in pp vs. AA collisions. This classification task can be tackled with machine learning techniques: by evaluating the performance of the trained classifiers, we are able to quantify how much information the machine has learned to distinguish the two samples of pp and AA

jets. We will explore two sets of machine learning architectures: the first to quantify the IRC-safe vs. IRC-unsafe information, and the second to quantify the information content in hard vs. soft physics. The latter approach will be the starting point for the machine-learning assisted observable design discussed in section 4 below. In this section, we consider only the jets generated by PYTHIA8 and JEWEL, without any background present from the heavy-ion underlying event. While the results in this section are presented using Monte Carlo event generators, they can be performed on either uncorrected experimental data or corrected full events; we outline a proposal for how such an experimental analysis can be carried out in section 6.

First, we will use a classifier based on deep sets [85–87] or Particle Flow Networks (PFNs) [88] which take as input sets of jet-by-jet four vectors. Since PFNs include the four vectors of all soft and collinear particles, they can make use of IRC unsafe information. Second, we compare the PFN-based classifier to results based on Energy Flow Networks (EFNs) which are an IRC-safe analogue of PFNs [88]. The difference in performance of these two classifiers can be considered as a measure of the IRC unsafe information content of quenched jets.

To examine the information content in hard vs. soft physics, we use a complete set of IRC safe observables which spans the phase space of a jet’s substructure. We explore two complete sets. The first is based on the approach developed in refs. [89–91] using the N -subjettiness observables [92, 93], which we feed as input to a Dense Neural Network (DNN). See also refs. [94–97]. The second complete set of observables we explore is based on Energy Flow Polynomials (EFPs) [98]. EFPs constitute a linear basis of jet substructure observables and, correspondingly, we train a linear classifier. We then study the discrimination power of the classifier as a function of the number of basis elements, and examine when the discrimination power saturates, indicating that all the relevant information is captured.

3.1 IRC-safe vs. IRC-unsafe information content

We start by reviewing permutation invariant neural network architectures. We consider both the IRC-unsafe PFNs and the IRC-safe EFNs [88]. We then train PFN/EFN-based classifiers on the pp vs. AA jet samples introduced above.

3.1.1 Particle and Energy Flow Networks

Deep sets [85–87] allow for a permutation invariant representation of data and naturally accommodate data sets with variable length. They are well suited for collider physics applications where deep sets can be trained directly on event-by-event particle four vectors. In the context of collider physics, deep sets were introduced in ref. [88] and referred to as PFNs, and subsequently studied with related architectures [99–105].

Consider a classifier represented by the function f , which depends on the four momenta p_i ($i = 1, \dots, M$) of M particles in the jet. In order to avoid that the classifier learns an unphysical ordering, it needs to be permutation invariant with respect to the per-event input variables, i.e. $f(p_1, \dots, p_M) = f(p_{\pi(1)}, \dots, p_{\pi(M)})$, where π is the permutation

operator. Following ref. [85], we can approximate the function f as

$$f(p_1, \dots, p_M) = F\left(\sum_{i=1}^M \Phi(p_i)\right), \quad (3.1)$$

where the functions Φ, F represent neural networks. The function Φ maps the momenta p_i to a d -dimensional latent space, where d can be chosen to be sufficiently large such that it approximates the function f arbitrarily well. We then sum over the functions $\Phi(p_i)$, in order to ensure that the result is permutation invariant. Finally, the function F is a map from the d -dimensional latent space to the scalar value of the function f from which the loss is calculated.

To separate the IRC safe information from the IRC unsafe information, we consider the EFNs of ref. [88]. IRC safety can be built into the permutation invariant neural network in eq. (3.1) as follows. Every particle inside a jet can be written in terms of its transverse momentum p_{T_i} (relative to the beam axis), rapidity y_i and azimuthal angle ϕ_i . We normalize the transverse momenta to the total jet transverse momentum and obtain the dimensionless momentum fractions $z_i = p_{T_i} / \sum_j p_{T_j}$. In addition, we introduce a 2-component vector which contains the angular variables $\hat{p}_i = (y_i, \phi_i)$. The corresponding permutation invariant and IRC safe neural networks, \tilde{f} , can then be written as

$$\tilde{f}(p_1, \dots, p_M) = F\left(\sum_{i=1}^M z_i \Phi(\hat{p}_i)\right). \quad (3.2)$$

Due to the weighting of Φ with the momentum fraction z_i , the resulting expression is IRC safe [88].

We parametrize the functions Φ and F in eqs. (3.1) and (3.2) in terms of DNNs, using the `EnergyFlow` package [88] with `Keras` [106]/`TensorFlow` [107]. For Φ we use two hidden layers with 100 nodes each and a latent space dimension of $d = 256$. For F we include three layers with 100 nodes each. For each dense layer we use the ReLU activation function [108] and we use the softmax activation function for the final output layer of the classifier. We train the neural networks using the Adam optimizer [109] with learning rates ranging from 10^{-3} to 10^{-4} . We use the binary cross entropy loss function [110], and train for 10 epochs with a batch size of 500. We find no significant changes in performance when changing the size of the layers, latent space dimension, learning rate, and batch size by factors of 2-10.

For each reconstructed jet, we record the transverse momentum, rapidity and azimuthal angle (p_{T_i}, y_i, ϕ_i) of each particle i inside the jet. Following ref. [88], we perform a preprocessing step to simplify the training process. We rescale the transverse momenta of each particle inside the jet with the total transverse momentum of the observed jet. In addition, we center the rapidity and azimuthal angles of the particles in the jet with respect to the jet direction. The jet axis is determined using the E -scheme [111]. Here we only consider PFNs without PID and we leave a more detailed exploration for future work. We benchmark our setup using the quark- vs. gluon-jet data set provided in ref. [112] as well as our own generated quark and gluon samples with PYTHIA8, finding compatible results with ref. [88].

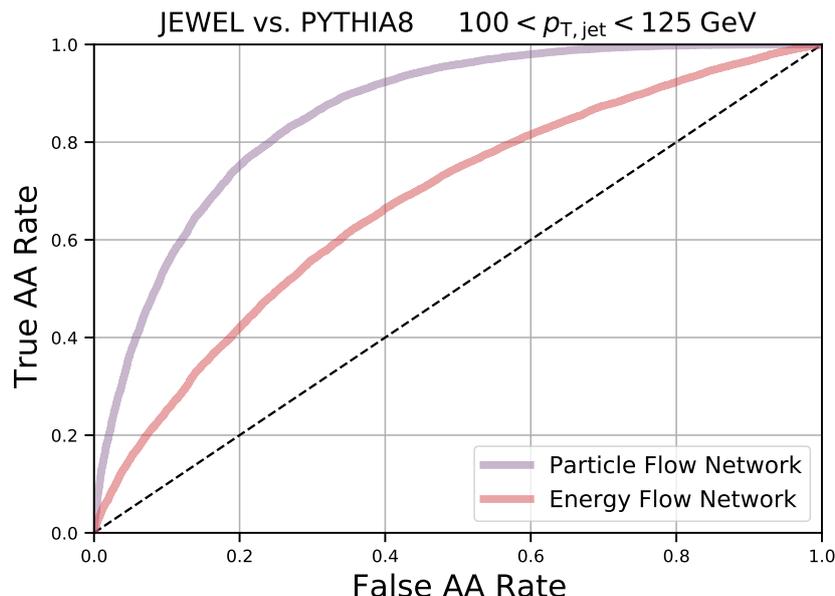


Figure 2. Classification performance of pp vs. AA jets quantified in terms of ROC curves using IRC-unsafe PFNs and IRC-safe EFNs. The jet samples in pp and AA collisions are obtained from Pythia 8 [70] and Jewel [72, 73].

Figure 2 shows the ROC curve for pp vs. AA jets using the PFNs and EFNs. The AUC is 0.860 for the PFN and 0.675 for the EFN. Since PFNs can efficiently make use of all the available information, we use them as a benchmark for the other classification techniques discussed below.

The difference in the classification performance between PFNs and EFNs can be considered as a measure of the IRC unsafe information contained in jet quenching. We observe that the classifier based on IRC-unsafe PFNs significantly outperforms the one based on IRC-safe EFNs. While our findings here are model dependent, this result indicates that a significant part of the information content of jet quenching is in the IRC unsafe physics. This striking difference between PFNs and EFNs was not observed for quark/gluon jet tagging in ref. [88]. Our findings suggest that in order to advance our understanding of the quenching process, it will be valuable to measure jet substructure observables in heavy-ion collisions that go beyond IRC-safety. If model parameters such as \hat{q} were extracted only using IRC-safe quantities such as the jet mass, a large amount of additional information which is encoded in the modified jet substructure in heavy-ion collisions would be missed. Examples of IRC-unsafe observables that are calculable within perturbative QCD include Sudakov safe observables [29–31], such as the groomed momentum sharing fraction [30, 113], ratios of jet angularities [29, 31], and the jet pull [114, 115]; or hadronic observables, which require the introduction of nonperturbative components, such as the transverse momentum spectra of hadrons [116] or the jet charge [117, 118]. Our results can therefore provide guidance for future efforts to explore which IRC-unsafe observables complement existing results and allow us to extract the maximum information content of jet quenching.

3.2 Hard vs. soft information content

In this section, we consider a different data representation and machine learning setup based on complete sets of IRC safe observables. In order to examine the distribution of information, we train classifiers using IRC safe observables as input and we explore when the discrimination power of the classifier saturates as more observables are included.

The number of jet substructure elements required to saturate the information content depends in general on the jet substructure basis employed. We consider two different choices. In section 3.2.1, we employ the N -subjettiness basis introduced in ref. [89]. The N -subjettiness basis offers the benefit of minimally describing the available phase space, and vanishes when N exceeds the particle multiplicity. In section 3.2.2, we consider EFPs which were introduced in ref. [98] as a linear IRC safe basis of jet substructure. Through their linear nature, EFPs are well suited for machine learning assisted observable design discussed in section 4. Individual EFPs and certain linear combinations are calculable in perturbative QCD. See ref. [119] for recent progress in this direction.

3.2.1 N -subjettiness basis

We start by reviewing the N -subjettiness observables [92–94, 96] which provide a complete and minimal basis of the M -body phase space of emissions inside the jet [89]. The N -subjettiness observables $\tau_N^{(\beta)}$ are defined as follows. First, we identify N axes inside the jet using the exclusive k_T algorithm [120, 121] with the E -recombination scheme [111]. The N -subjettiness variables measure the radiation in the direction of these axes:

$$\tau_N^{(\beta)} = \frac{1}{p_T^{\text{jet}}} \sum_{i \in \text{Jet}} p_{Ti} \min \left\{ R_{1i}^\beta, R_{2i}^\beta, \dots, R_{Ni}^\beta \right\}, \tag{3.3}$$

where p_{Ti} is the transverse momentum of particle i inside the jet and R_{ji} is its distance in the $\eta - \phi$ plane with respect to the identified axes j . The exponent $\beta > 0$ is a tuneable parameter. For a given N and β , the N -subjettiness observables output a single number per jet which quantifies the radiation pattern inside the jet. We note that for $\leq N$ particles inside the jet, the N -subjettiness observables vanish. We compute the N -subjettiness observables using the implementation in the FastJet Contrib software package [122]. Figure 3 shows the N -subjettiness distributions for several different N and β values for jets in both pp and AA collisions (diagonal panels) and their pairwise correlation (off-diagonal panels).

Following ref. [89], we consider $3M - 4$ N -subjettiness observables per jet which completely specify the M -body phase space. For example, the 2-body phase space is $3 \cdot 2 - 4 = 2$ dimensional and the measurement of two 1-subjettiness variables with different β values, e.g. $\tau_1^{(1)}$ and $\tau_1^{(2)}$, is sufficient to completely specify the kinematics of the two emissions inside the jet, equivalent to the longitudinal momentum fractions $z, 1 - z$ and the relative angle θ of the two emissions. For each additional emission, three phase space dimensions are added which are fully specified by measuring three additional N -subjettiness observables.

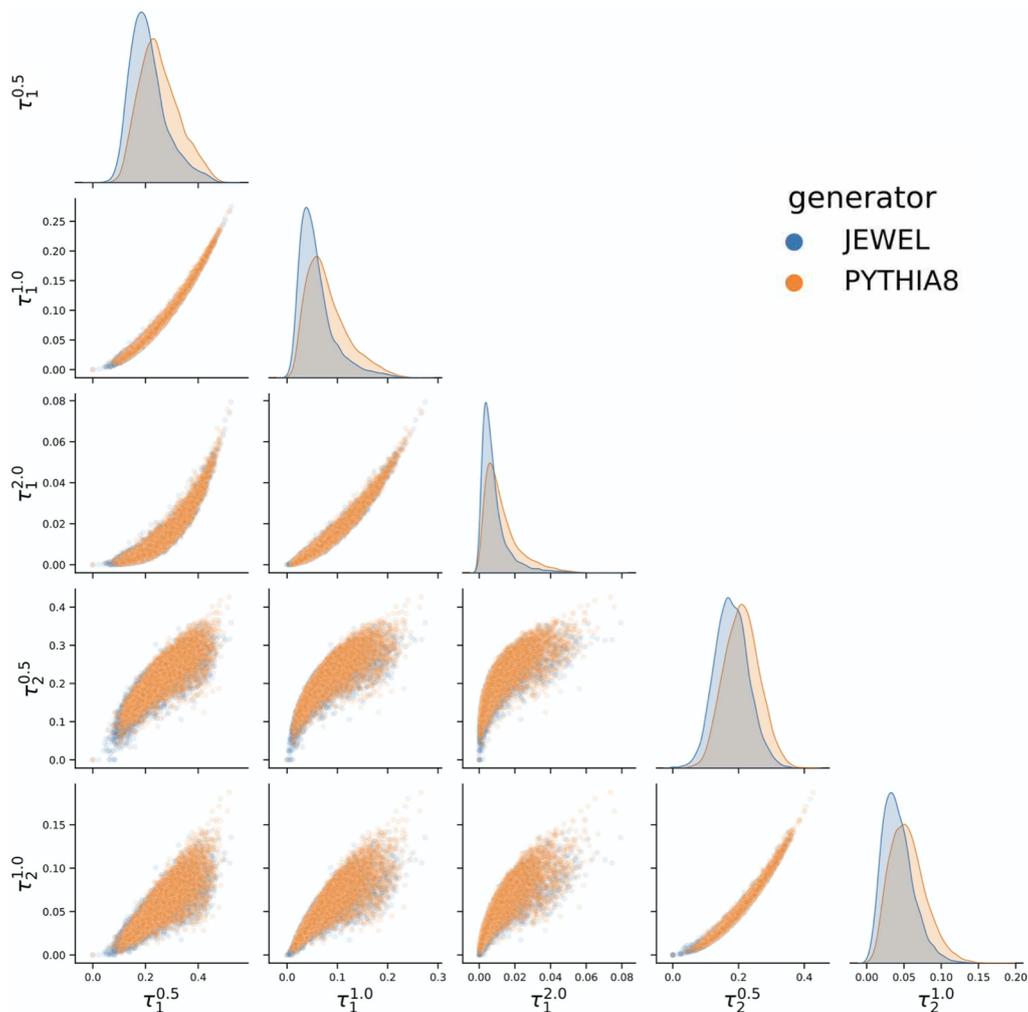


Figure 3. Scatter plot showing different N -subjettiness distributions (diagonal) and their pairwise correlations (off-diagonal panels) in pp and AA collisions without background. The pp and AA results shown here are obtained from Pythia 8 [70] and Jewel [72, 73], respectively.

We use the following observables to specify the M -body phase space [89]:

$$\begin{aligned}
 & \text{2-body : } \tau_1^{(1)}, \tau_1^{(2)}, \\
 & \text{3-body : } \tau_1^{(0.5)}, \tau_1^{(1)}, \tau_1^{(2)}, \tau_2^{(1)}, \tau_2^{(2)}, \\
 & \quad \vdots \\
 & \text{\textit{M}-body : } \tau_1^{(0.5)}, \tau_1^{(1)}, \tau_1^{(2)}, \dots, \tau_{M-2}^{(0.5)}, \tau_{M-2}^{(1)}, \tau_{M-2}^{(2)}, \tau_{M-1}^{(1)}, \tau_{M-1}^{(2)}. \tag{3.4}
 \end{aligned}$$

We note that the M -body phase space discussed here does not directly correspond to individual particles inside the jet but rather to “resolved emissions” inside the jet. This is analogous to the Soft Drop procedure [123] where a sufficiently hard emission passes the Soft Drop criterion. In an IRC safe representation of a jet, soft and collinear particles are clustered first leaving only the M most resolved emissions, which correspond to the $3M - 4$

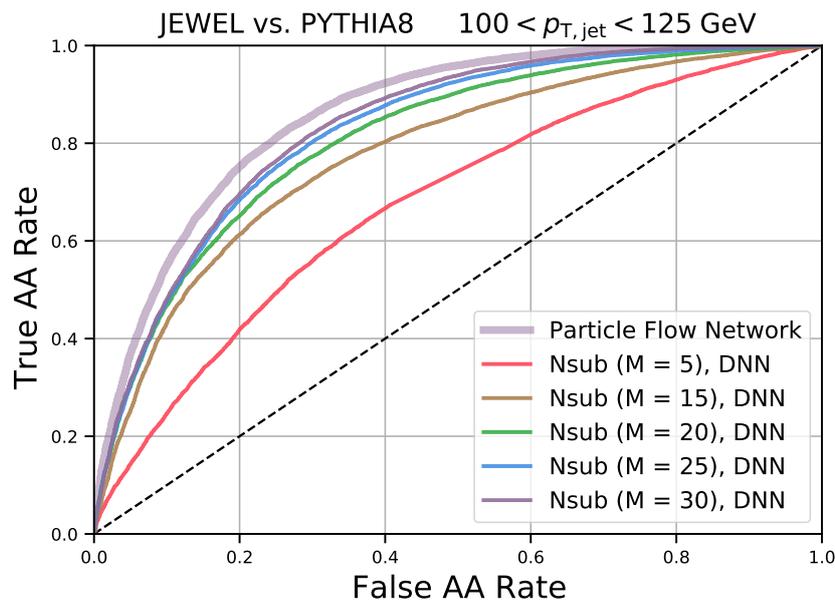


Figure 4. ROC curves for jets in pp vs. AA collisions using the N -subjettiness basis. For comparison we also show the result obtained using the classifier based on PFNs.

N -subjettiness observables passed to the machine learning algorithm. If M is sufficiently large, the kinematics of every particle can be fully reconstructed from the N -subjettiness observables. While the N -subjettiness observables are all IRC safe, nonperturbative effects can nevertheless play an important role. Especially the observables with large values of N are sensitive to very low energy scales. By varying the value of M we can control the sensitivity to soft and collinear emissions inside the jet.¹

We pass the $3M - 4$ N -subjettiness observables for each jet to the input layer of a fully connected DNN as the classifier. The resulting classifier is generally Sudakov safe [89]. We use 3 hidden layers with between 32-512 nodes, each with a ReLU activation function [108], followed by a sigmoid activation for the final output layer. We train the neural network on the N -subjettiness observables with the Adam optimizer [109] and a learning rate ranging from 0.01 to 0.001 and batch size 1000, with the binary cross entropy loss function of ref. [110]. We use Keras [106]/TensorFlow [107] for the implementation, and determine the number of nodes in each hidden layer and the learning rate using a hyperparameter optimization with the Hyperband algorithm [124] implemented in Keras Tuner [125]. To crosscheck the DNN setup, we compared our results to a Random Decision Forest and found quantitatively comparable results.

Figure 4 shows the ROC curves using the N -subjettiness basis where we include an increasing number of $3M - 4$ observables that fully specify the M -body phase space. For comparison, we also show the PFN result as in figure 2 above. While the relatively hard physics encapsulated by $M \leq 5$ shows a significant discrimination power, we observe that the discrimination power significantly improves as more observables are included. Even-

¹We thank Andrew Larkoski for clarifying discussions.

tually, the discrimination power starts to saturate at a relatively large value of $M \approx 30$. This indicates that soft emissions inside the jet play an important role and contain critical information about the interaction of hard probes with the QGP. The saturation at large M found here is in stark contrast to the discrimination of QCD vs. Z -boson jets, as the previous studies showed saturation of the information content with the $M = 4$ -body phase space, i.e. a rather limited number of observables captures the complete information [89]. This has been discussed in the context of entropy production, where multiple soft emissions are produced in the jet that carry little or no additional information about the initial hard-scattering Z -boson signal [126]. Unlike in the case of Z boson vs. QCD jets, however, for pp vs. AA jets we expect that there is significant information contained in the soft physics due to sensitivity to the surrounding medium in AA case. We note that this observation is generally in agreement with the large difference between PFNs and EFNs found in section 3.1. Our findings suggest that it will be necessary to measure new soft-sensitive jet substructure observables in heavy-ion collisions to fully make use of the available information recorded by the experimental collaborations. This information can be accessed by N -subjettiness observables for large values of N . We emphasize again that while the conclusions here are model-dependent, we are confident that a similar analysis can be performed with experimental data. In addition, we note that the studies here do not include the heavy-ion background, which poses a major obstacle in measurements of soft physics. We will discuss the impact of the heavy-ion underlying event in more detail in section 6.

3.2.2 Energy Flow Polynomial basis

EFPs were introduced in ref. [98] as an (over)complete linear basis of IRC-safe jet substructure observables. They are multi-particle correlators which can be indexed with multigraphs $G = (V, E)$ with V vertices and E edges. For a jet with M particles, the EFP $_G$ is defined as

$$\text{EFP}_G = \sum_{i_1=1}^M \cdots \sum_{i_V=1}^M z_{i_1} \cdots z_{i_V} \prod_{(k,l) \in E} \theta_{i_k i_l}. \quad (3.5)$$

Here z_i is a measure of the energy of particle i and θ_{ij} is a pairwise distance measure of particles i and j in the η - ϕ plane

$$z_i = \left(\frac{p_{Ti}}{\sum_j p_{Tj}} \right)^\kappa, \quad \theta_{ij} = (\Delta\eta_{ij} + \Delta\phi_{ij})^{\beta/2}, \quad (3.6)$$

where (κ, β) are free parameters. For $\kappa \neq 1$, the EFPs are IRC unsafe [56]. Here, we limit ourselves to the IRC-safe EFPs, and in particular $(\kappa, \beta) = (1, 0.5)$. The number of edges is the degree of the EFPs. Since the EFPs are a linear basis, any IRC-safe observable \mathcal{O} can be approximated as

$$\mathcal{O} \approx \sum_{G \in \mathcal{G}} c_G \text{EFP}_G, \quad (3.7)$$

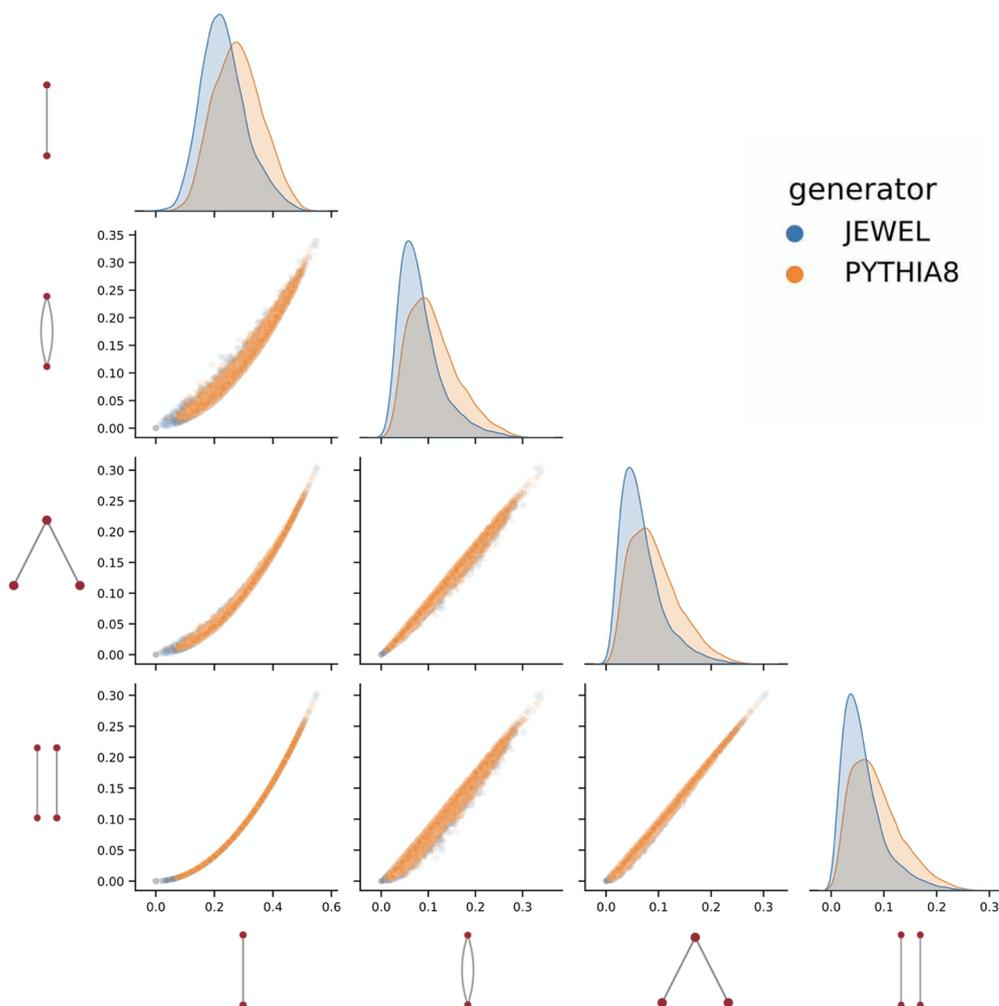


Figure 5. Scatter plot showing different EFP distributions (diagonal) and their pairwise correlations (off-diagonal panels) in pp and AA collisions without background. The pp and AA results shown here are obtained from Pythia 8 [70] and Jewel [72, 73], respectively.

with coefficients c_G and \mathcal{G} is a finite set of multigraphs. A simple example of the EFPs is

$$\frac{1}{2} \times \left(\text{diagram of two vertices connected by two curved lines} \right) = \frac{1}{2} \sum_{i_1=1}^M \sum_{i_2=1}^M z_{i_1} z_{i_2} \theta_{i_1 i_2}^2, \quad (3.8)$$

which, for $\beta = 1$, is equal to the jet mass $m_J^2/p_T^{\text{jet}2}$ up to power corrections in the jet radius. The distributions and correlations of several EFPs in pp and AA collisions are shown in figure 5. See refs. [98] for more details.

Since EFPs constitute a linear basis of jet substructure observables, we make use of a linear classifier. We train a linear discriminant on the pp vs. AA jet data set using EFPs up to degree 7. We use an implementation with the EnergyFlow package [88] and

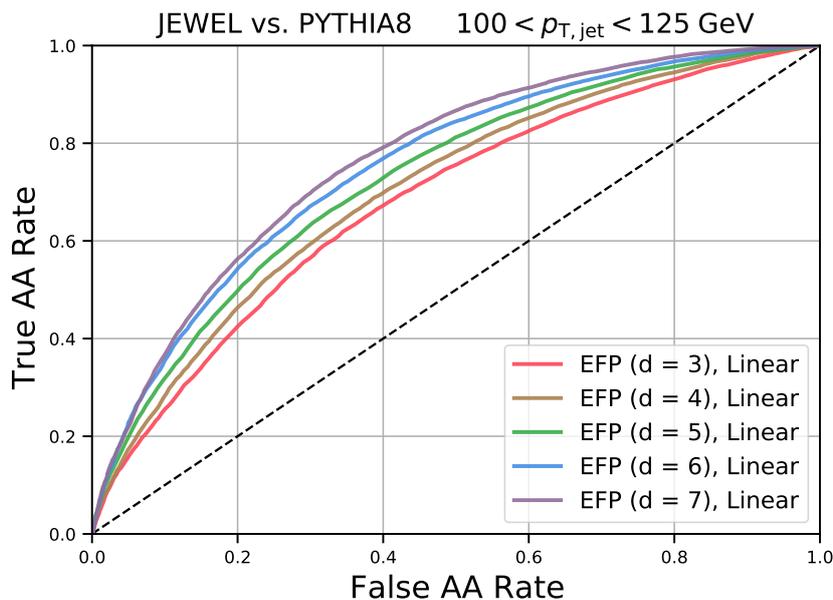


Figure 6. ROC curves for jets in pp vs. AA collisions using the EFP basis up to degree 7.

`scikit-learn` [127]. Since we train a linear classifier rather than a DNN, the result is not directly comparable to the N -subjettiness DNN, but rather is complementary to it.

The corresponding ROC curves are shown in figure 6. We find that as we increase the EFP dimension, the classification power increases. This is similar to the result with the N -subjettiness basis in the sense that a large number of EFPs needs to be included indicating that significant information is contained in the soft physics of jets. In fact, for EFPs up to degree 7, which corresponds to 1000 observables, we do not yet observe a saturation of the discriminating power. This corroborates our results in the previous section. We note again that the performance of the different classifiers here depends on the choice of the Monte Carlo model generator. Lastly, we note that there is no unique definition of the amount of IRC-unsafe information. For example, while both the EFNs in section 3.1 and the linear model of EFPs here have only access to IRC-safe information, they represent different architectures and one may be trainable more efficiently than the other (leading to a better performance) depending on classification task.

4 Observable design

The machine learning algorithms described in section 3 deliver powerful classification ability, and themselves can be used as observables that can be measured experimentally and calculated in Monte Carlo event generators. However, there are limits to their ability to provide insights about jet quenching since they are generally not amenable to analytical calculations in pQCD. In order to address this, we propose to use machine learning to assist in *observable design*, by approximating the more powerful classifiers with simpler models that are in principle calculable in perturbative QCD. In particular, we focus on observables with closed-form expressions that obey IRC-safety or Sudakov safety and therefore have the

potential to be calculated in the future with existing methods. To do this, we use symbolic regression to identify observables which are optimized to balance discrimination power and simplicity. We focus on identifying single observables that maximize discriminating power, i.e. using machine learning to identify the maximally modified jet substructure observable. Currently, it is unknown whether the traditional jet substructure observables that have been previously measured are close to maximally modified or not; the identification of maximally modified observables can provide guidance on what observables should be measured to elucidate the physics of jet quenching. Once this machine learning observable design is done, this observable can then be measured experimentally and in principle calculated theoretically using existing methods. The observable design procedure can be performed either with Monte Carlo generators in order to produce an ansatz observable that can then be measured with experimental data, or directly with experimental data, which will be discussed in section 6. In the future, this machine learning approach can be extended to consider the maximum discriminating power from a set of observables, which can ultimately guide the entire program of jet substructure measurements and calculations, and motivate new input to global analyses used to extract medium properties.

4.1 Lasso regression

In order to select dominant features of our classifier, we use a simple form of symbolic regression known as Lasso (least absolute shrinkage and selection operator) regression [65, 66]. Lasso is a linear regression model, although unlike regular linear or ridge regression, Lasso employs an ℓ_1 regularization, which causes an increasing number of coefficients to vanish as the regularization parameter is increased. In this way, Lasso is able to perform feature selection. The loss function of the Lasso regression is given by

$$\min_{\vec{w}} \|\vec{y} - T\vec{w}\|_2^2 + \lambda \|\vec{w}\|_1. \tag{4.1}$$

Here $\|\cdot\|_p$ denotes the ℓ_p norm, \vec{y} represents the true labels, T contains the input training data, \vec{w} contains the weights and λ is a regularization parameter which determines the level of sparsity of the obtained solution. For $\lambda = 0$, we recover the regular linear regression. Different than the ridge regression (ℓ_2 regularization), the Lasso regression forces the sum of the absolute values of the coefficients in $\vec{\beta}$ to be less than a certain value.

To formulate our Lasso classifier, we construct an observable that is a product of N -subjettiness variables, which is generally Sudakov safe. We choose the following functional form [90, 91]:

$$\mathcal{O}_{N\text{-sub}}^{\text{ML}} = \left(\tau_1^{(0.5)}\right)^{a_{1,1}} \left(\tau_1^{(1)}\right)^{a_{1,2}} \left(\tau_1^{(2)}\right)^{a_{1,3}} \cdots \left(\tau_{M-1}^{(1)}\right)^{a_{M-1,1}} \left(\tau_{M-1}^{(2)}\right)^{a_{M-1,2}}. \tag{4.2}$$

After taking the logarithm of $\mathcal{O}_{N\text{-sub}}^{\text{ML}}$, we can determine the $3M - 4$ exponents $a_{i,j}$ as the weights of the Lasso regression.

By constructing a single observable, we provide a feasible proposal for an experimental measurement as opposed to, for example, the full set of $3M - 4$ individual observables. By measuring the product observable that is most strongly modified, one can capture the leading effect of the modification of jets in heavy-ion collisions. Moreover, choosing a Sudakov

safe product form is motivated by our previous findings that IRC-unsafe observables add valuable information. This approach can be extended from a single observable to study a minimal set of observables that can achieve approximately equal discriminating power to the full set of $3M - 4$ observables. Note that despite the fact that the regression procedure requires $3M - 4$ observables as input, when performing this procedure on experimental data one need not perform $3M - 4$ corrected cross-section measurements. Rather, one can use detector-level information about these $3M - 4$ observables to construct an *approximate maximally discriminating observable*, which can then be measured as a single corrected cross-section. Experimental considerations will be discussed further in section 6.

The regularization parameter λ , which can take any positive value, provides a handle to balance the performance of the classifier with the simplicity of the resulting observable. When λ is small, a product observable with strong classification performance but many terms will be found, and as λ is increased, a product observable with decreased classification performance but fewer terms will be found. That is, larger values of λ generally decrease the number of terms, and pick out the “most important” N -subjettiness observables satisfying that regularization condition. When λ becomes sufficiently large the classifier will saturate and select only one term. As λ is decreased, additional terms will be added. Note that the N -subjettiness terms selected by the regression procedure at one choice of λ need not be a subset of those at a different λ . Moreover, even if they are a subset, there is nothing that requires that the “extra” higher-order terms to be less important — for example the 2-term observable could be more than twice as discriminating as the 1-term observable. The convergence of the Lasso regression can be slow for a large parameter space which is why we limit ourselves here to a relatively small number of input observables. For several values of λ , we find the following observables without background for $M = 15$ in our Monte Carlo model studies:

$$\lambda = 0.5 : \quad \mathcal{O}_{N\text{-sub}}^{\text{ML}} = \tau_{14}^{(1)}, \quad (4.3)$$

$$\lambda = 0.1 : \quad \mathcal{O}_{N\text{-sub}}^{\text{ML}} = \left(\tau_{10}^{(1)}\right)^{0.071} \left(\tau_{11}^{(1)}\right)^{0.157} \left(\tau_{14}^{(1)}\right)^{0.649} \tau_{14}^{(2)}, \quad (4.4)$$

$$\lambda = 0.01 : \quad \mathcal{O}_{N\text{-sub}}^{\text{ML}} = \left(\tau_2^{(0.5)}\right)^{0.608} \left(\tau_4^{(2)}\right)^{-0.186} \times \dots \times \tau_{14}^{(2)} \quad (23 \text{ terms}). \quad (4.5)$$

Since we can rescale the exponents by an overall factor without changing the performance of the classifier, we choose the exponent of the rightmost factor in eq. 4.2, in this case $\tau_{14}^{(\beta)}$, as 1 for all values of λ .

We find that the Lasso regression generally prefers large values of N . For sufficiently large values of λ , we find that the Lasso regression always picks only one observable which turns out to be one of the N -subjettiness observables with the largest allowed value of N . When λ is lowered gradually, the Lasso regression adds additional N -subjettiness observables with intermediate values of N . If we further lower λ , the Lasso regression eventually adds observables with small values of N . This is generally in line with our previous observations that significant information about jet quenching is in the soft physics. While large- N N -subjettiness are not simple to calculate, our results suggest motivation to extend perturbative QCD calculations to increasingly large N . Interpreting the exact

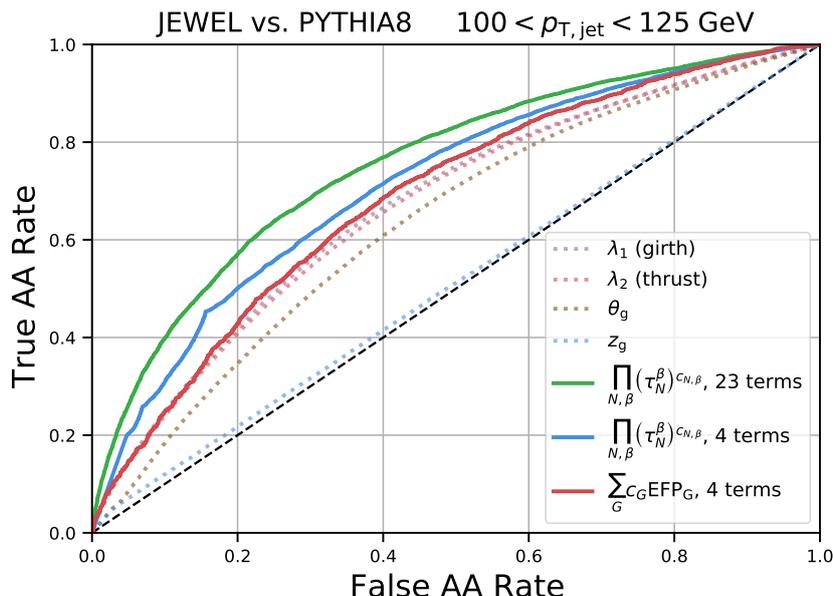


Figure 7. ROC curves for the Lasso regression using the N -subjettiness basis and EFPs. For comparison we also show the result for typical observables in heavy-ion collisions.

jet quenching physics message from these product observables will require more theoretical work — here we focus mainly on the important step that these are in principle calculable, unlike the high-dimensional machine learning inputs. We emphasize that these results are model-dependent, and the procedure should be repeated on experimental data in order to obtain nature’s maximally discriminating single observables. We also compared our results to regression algorithms with an ℓ_0 regularization (NP-hard) and found similar results. The exact set of observables which is chosen can vary somewhat when we run the Lasso regression with a different random seed. This is likely due to the existence of multiple local minima and the shared information between different N -subjettiness observables, see also figure 3. We also find that the discrimination power increases only gradually as we allow for more observables to be included in the product observable. This observation is in line with the findings of ref. [128]. We expect that more general symbolic regression techniques as developed for example in ref. [129] can also be adapted to identify suitable jet substructure observables which we plan to explore in the future work. See also ref. [130].

To illustrate the performance of the different observables, we show the corresponding ROC curves in figure 7. As expected, solutions with more terms (lower values of λ) lead to better discriminating power at the expense of an increased complexity of the observable. For comparison, we also plot the ROC curves of four observables which have already been measured in heavy-ion collisions: the soft drop momentum sharing fraction z_g [28, 30, 113, 131], the groomed jet radius [28, 123, 132–134], jet thrust [135, 136] and jet girth [137]. We observe that the machine-learned product observables outperform these observables. The distributions of these observables are shown in figure 8. Note that the N -subjettiness product observable distribution exhibits a large difference between JEWEL and PYTHIA8 for $\mathcal{O}_{N\text{-sub}}^{\text{ML}} = 0$. This indicates that a significant amount of the classification

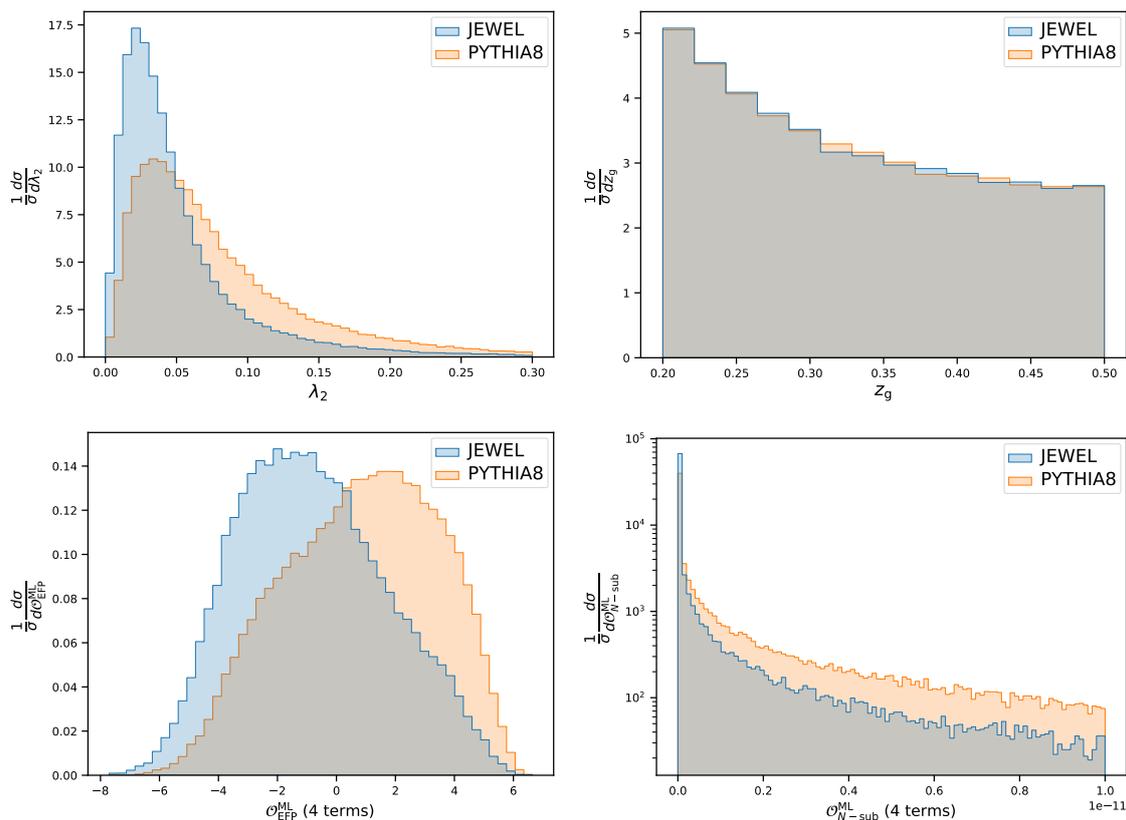


Figure 8. Distributions of observables in pp and AA collisions which have already been measured by experimental collaborations and examples of the machine-learned observables using the N -subjettiness and EFP basis.

power is related to the multiplicity of the jet, since τ_N^β vanishes when the jet constituent multiplicity is $< N$. This aspect of the modification is highly model-dependent, however, and the exact $\mathcal{O}_{N\text{-sub}}^{\text{ML}}$ found here should be taken for illustrative purposes only, and should be determined using experimental data.

In addition to an N -subjettiness product observable ansatz, we consider a linear combination of EFPs. As discussed above, EFPs form a linear basis of jet substructure observables making them ideally suited for a Lasso regression [98]. Here we preprocess the input data to zero mean and unit variance. Analogous to the previous section, we determine the coefficients c_G of

$$\mathcal{O}_{\text{EFP}}^{\text{ML}} = \sum_{G \in \mathcal{G}} c_G \text{EFP}_G, \tag{4.6}$$

except that now we use a Lasso regression instead of a “regular” linear regression which will pick out only the most important terms. As an example, we use EFPs up to degree 4

and the Lasso regression with $\lambda = 0.001$, for which we find a very stable result:

$$\begin{array}{c} \bullet \\ | \\ \bullet \end{array} + 3.54 \times \begin{array}{c} \bullet \\ / \quad \backslash \\ \bullet \quad \bullet \end{array} + 1.72 \times \begin{array}{c} \bullet \\ / \quad \backslash \quad / \quad \backslash \\ \bullet \quad \bullet \quad \bullet \quad \bullet \end{array} - 3.82 \times \begin{array}{c} \bullet \quad \bullet \\ | \quad | \\ \bullet \quad \bullet \end{array} . \quad (4.7)$$

The corresponding ROC curve and the distribution of this ML-learned observable are shown in figures 7, 8, respectively. We find that despite the simplicity of the machine-learned EFP observable, it outperforms the other “traditional” observables. The intriguing aspect of observables which involve a relatively small number of EFPs, as in eq. (4.7), are that they are generally calculable within perturbative QCD, see ref. [119] for recent progress. Moreover, we note that some combinations of EFPs are easier to calculate directly compared to individual EFPs [98]. The same holds for certain products of N -subjettiness observables [89]. We leave the identification of the subspace of observables that is relatively simple to calculate within perturbative QCD and the combination with ML techniques like the Lasso regression for future work.

5 Information loss: the underlying event and background subtraction

The large, fluctuating underlying event produced by the QGP causes notorious experimental and theoretical challenges in heavy-ion collisions — in particular, by limiting which observables can be reliably measured. Typically, background subtraction procedures are applied in order to mitigate this problem. Systematic uncertainties associated with the subtraction are estimated in order to adequately capture the lack of exact knowledge of which particles arise from the underlying event, and which from the jet.

From the perspective of information content, this presents two distinct mechanisms by which the information in jet quenching can be lost. First, the fluctuating underlying event can be viewed as a source of noise. One cannot distinguish particles arising from underlying event from those correlated to the jet, and so to the extent that the noise distribution overlaps with the signal distribution, the ability to distinguish the two is irrecoverably reduced. Second, background subtraction algorithms themselves can cause information loss. Since background subtraction inherently involves removal of particles from the jet, and one does not have exact knowledge of which particles arise from the underlying event, this procedure strictly results in information loss.

The jet classification methods used in section 3 can be used to evaluate the magnitude of each of these contributions. Within the context of the parton shower models considered, we assess the overall impact of the underlying event on the jet classification performance by comparing a PFN trained only on the hard jet particles to a PFN trained on the combination of jet and background particles (after performing constituent subtraction, as described in section 2). In order to perform a controlled comparison, in both cases we select jets based on the p_T^{jet} of the particles from PYTHIA8 or JEWEL alone. Figure 9 shows that there is a dramatic decrease in the classification power due to the presence

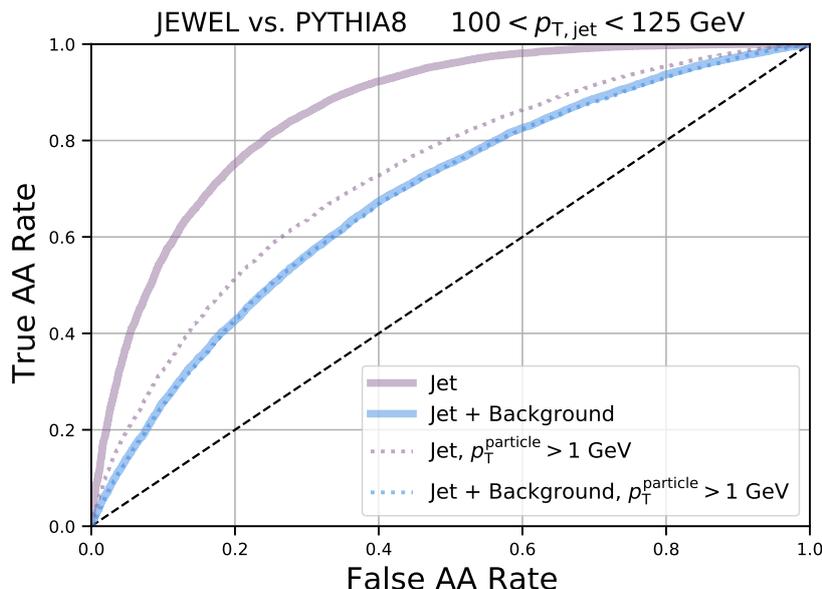


Figure 9. ROC curves for PFNs trained with (i) PYTHIA8/JEWEL jets, (ii) jets clustered from a combination of PYTHIA8/JEWEL events with a thermal background, with event-wide constituent subtraction applied ($R_{\text{max}} = 0.25$), (iii) PYTHIA8/JEWEL jets only considering jet constituents with $p_T > 1 \text{ GeV}$, and (iv) jets clustered from a combination of PYTHIA8/JEWEL events with a thermal background, only considering jet constituents with $p_T > 1 \text{ GeV}$, with event-wide constituent subtraction applied ($R_{\text{max}} = 0.25$).

of the underlying event. We also plot PFNs trained on jet particles with $p_T > 1 \text{ GeV}$. Comparing the ROC curves with and without this requirement, we find that in the case without background, a large discrimination power resides in the soft physics — whereas in the case with background, the presence of soft information makes no difference. That is, in the presence of background, sufficiently soft discrimination is no longer useful — and the discrimination is dominated by hard physics. This observation presents a delicate challenge for the study of jet quenching: soft information is crucial to maximize discrimination between quenched and unquenched jets, yet the fluctuating underlying event fundamentally prevents much of this information from being accessed.

In order to disentangle how much of this information loss is due to inherent noise of the fluctuating background, and how much is due to the background subtraction algorithm, we compare the classification performance before and after background subtraction is performed. To do this, we fix the jet axis using the hard jet in order to avoid distortion of the axis by the fluctuating background. We consider three different jet populations, each defined by a cone of particles within rapidity-azimuth distance of $\Delta R = 0.4$ from the common axis. The three different particle selections we consider are:

1. Hard jet particles only
2. Hard jet particles and background particles, with constituent subtraction applied
3. Hard jet particles and background particles, without background subtraction

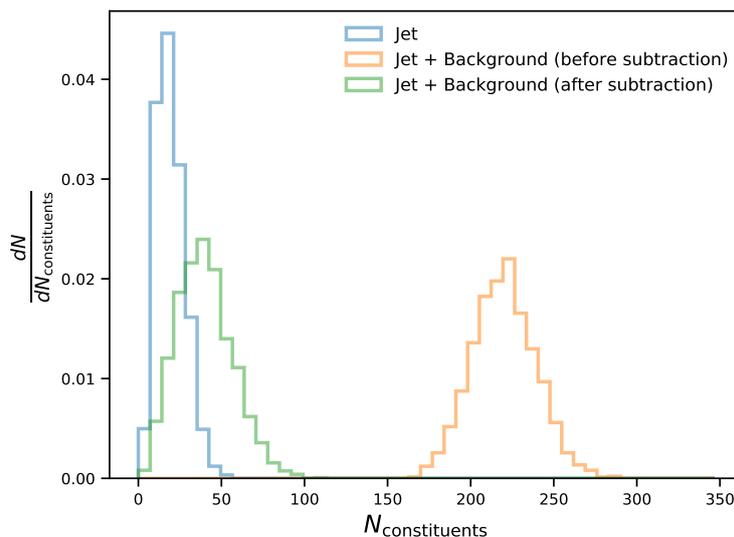


Figure 10. Distributions of particle multiplicities in the training jet population.

To illustrate the differences between these three selections, note that the background subtraction procedure can cause a dramatic change in the particle multiplicity of the jet. Figure 10 shows the distribution of multiplicities of jet constituents before and after background subtraction, for the nominal constituent subtraction configuration with $R_{\max} = 0.25$. Application of the background subtraction procedure brings the constituent multiplicity distribution much closer to the distribution in the absence of background. The dramatic change, however, emphasizes the possibility that the background subtraction procedure may cause significant information to be lost. Given the increased number of particles before background subtraction, when training the PFN before background subtraction, we double the number and size of all layers, as well as the dimension of the latent space, which we observe to give an improvement to the PFN performance (the case after background subtraction and the hard jet are unaffected).

Figure 11 shows the ROC curve for PFNs trained before and after background subtraction (for two values of R_{\max}), with comparison to the hard jet for reference. The ROC curve before background subtraction is higher than that after background subtraction, indicating that, within our model studies, the constituent subtraction procedure does remove a small but significant amount of information. As the background subtraction becomes stronger, going from $R_{\max} = 0.25$ to $R_{\max} = 1.0$, we observe additional loss of information. This constitutes the first study quantifying the information loss of background subtraction algorithms. We propose that the ROC curve should be used as a metric to quantify the information loss, and can be used to tune optimal parameters in the algorithm — including machine learning based background subtraction [138].

These results come with the caveat that there is a third source that can deteriorate the performance of the classifier in the presence of background: the performance of the PFN may eventually deteriorate due to the increased size of the training data. In particular, we note that it is possible for the ROC curve to improve after the background subtraction but

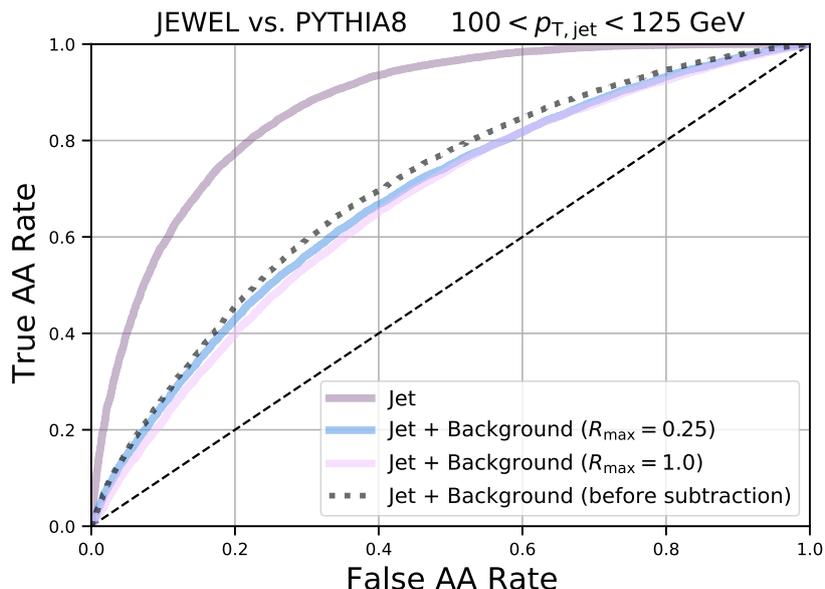


Figure 11. ROC curves comparing the performance with (i) Hard jet particles only, (ii) Hard jet particles and background particles, with constituent subtraction applied (for two different values of R_{max} , and (iii) Hard jet particles and background particles, without any background subtraction applied

only if the PFN is not a sufficiently good approximation [86] of the optimal classifier which is given by the likelihood ratio according to the Neyman-Pearson lemma [139]. Nevertheless, we can unambiguously see that, within our model studies, the background subtraction procedure (which *reduces* the size of the per-jet training data) overall deteriorates the performance of the PFN.

6 Application to experimental data

In order to assess the actual information content that distinguishes jets in AA collisions from those in pp , our analysis can be performed using the event-by-event data collected by the experimental collaborations at the LHC and RHIC. In this section we outline several considerations for how such an analysis can be carried out.

In experimental data, one cannot unambiguously separate the hard jet component from the QGP underlying event. Consequently, the quenched jets in heavy-ion collisions should be compared to a sample of proton-proton jets embedded into heavy-ion background, with or without background subtraction. These two samples define the two classes to be distinguished by the classifier, which we denote AA and $pp \oplus AA$, respectively.

The simplest approach to measure the ROC curve experimentally is to train the classifier on detector-level quantities: AA vs. $pp \oplus AA$, taking careful consideration of differences in detector effects in the two samples. Note that experimentally, one does not have access to the “hard jet” p_T^{jet} , unlike the studies we performed in section 5. Instead, one can only select jet populations based on the p_T^{jet} in the presence of background, for example after

background subtraction. In order to correct for the detector and the background fluctuation effects, one can apply corrections either after training (such as with a bin-by-bin correction) or before training (with an unfolding procedure to correct the inputs). Using traditional unfolding methods, however, one cannot easily build a response matrix to perform the correction, unless one performs an unfolding of the high-dimensional input space, such as a $3M - 4$ -dimensional set of N -subjettiness observables. With the recent prospect of machine learning based methods to unfold high dimensional events [67], this may be possible, even to the extent of unfolding full events and thereby enabling the training of a classifier directly on corrected particles.

There are several additional challenges in performing these measurements compared to the Monte Carlo studies presented above. First, the detector conditions between the proton-proton and heavy-ion data taking periods may be different — and the classifier will naively learn these differences. Second, in the $pp \oplus AA$ jet sample, one must ensure that only soft particles — and not hard jets — enter the distribution from the embedded heavy-ion event. Third, the size of the jet sample is limited by the available statistics recorded by the experiment, which in turn can limit the performance of the classifier. These challenges are each surmountable, and we are optimistic that such an analysis can be performed at the LHC.

We propose that each of the three complementary studies in sections 3–5 can be performed on experimental data:

- *Measuring the ROC curve.* The measured ROC curve can serve as an observable that can be compared to Monte Carlo event generators. Moreover, the distribution of information content with complete sets of jet substructure observables can provide a differential test of jet quenching models, to the extent that highly soft-sensitive observables, such as high- N N -subjettiness or high-dimension EFPs, can be reliably measured in the presence of the heavy-ion underlying event.
- *ML-assisted observable design.* Regardless of whether the classifier is trained on detector-level inputs or corrected inputs, symbolic regression can be used to identify approximate maximally discriminating observables. These identified observables can then be measured with traditional techniques: correcting for detector and background effects, and in principle comparing to jet quenching calculations.
- *Information content and background subtraction techniques.* The information loss caused by various background subtraction algorithms can be quantified by comparing classification performance before and after subtraction, and can be used to select and tune subtraction algorithms to minimize information loss.

7 Conclusions and outlook

We explored the use of machine learning classifiers to investigate the information content distinguishing jets in heavy-ion collisions from jets in proton-proton collisions. The labels for the fully supervised machine learning task here are unambiguous and we expect that

new insights about the nature of the quark-gluon plasma (QGP) can be obtained with our approach. As a proof of concept, we trained classifiers and designed new observables using jet samples generated from parton shower event generators. We expect that the developed techniques can be applied directly to event-by-event experimental data.

Within our model studies, we found that there is a significant ability to distinguish jets in the absence of the heavy-ion underlying event. By comparing Infrared Collinear (IRC) unsafe Particle Flow Networks to IRC safe Energy Flow Networks, we found that a substantial amount of this information resides in IRC-unsafe physics. By using complete sets of IRC safe N -subjettiness observables and Energy Flow Polynomials, we studied the convergence of the classifier as a function of the number of observables included in the training procedure, and found that the performance saturates only when a large number of observables are included. This demonstrates that a substantial amount of information is contained in soft emissions inside the jet, and is in stark contrast to e.g. QCD vs. Z -jet tagging where only a small number of observables is sufficient.

The ability to distinguish jets in heavy-ion collisions from jets in proton-proton collisions appears to substantially decrease in the presence of the heavy-ion underlying event. Our model studies indicate that this is primarily due to fluctuations inherent to the background, which lead to irrecoverable information loss, but that there is also a significant information loss due to commonly used background subtraction algorithms. We performed the first such study assessing the information loss of background subtraction, which can be used to tune and develop new background subtraction algorithms in the future to preserve the information content of the jet quenching process.

In order to directly connect our machine learning approach to theoretical calculations, we designed new observables using Lasso regression that maximize the discrimination power between jets in heavy-ion collisions and jets in proton-proton collisions and are in principle calculable in perturbative QCD. By finding observables that maximize a balance of being (i) highly discriminating, and (ii) in principle calculable, we provide guidance for the theoretical understanding of jet quenching. These studies can be extended in the future to not just explore maximally modified observables, but to identify sets of observables that offer the maximal ability to constrain QGP properties. Already, our observation that the soft physics of jets plays an important role provides motivation to study specifically soft-sensitive jet substructure observables — which may bridge the gap between traditional jet substructure measurements and the machine-learned classifiers. We expect these methods to inform future experimental analyses and theoretical focus.

It is essential that the studies presented here be performed on event-by-event experimental data, since parton shower event generators for jet quenching involve assumptions ranging from the role of medium response to the factorization structure of jet quenching. Since the training labels are exactly known, this serves as a case in which machine learning can be used directly on experimental data without reliance on modeling. The studies should be extended to different jet topologies, different kinematic ranges, and multiple jet radii. Such studies will present new opportunities to use jets to precisely constrain QGP properties, and can reveal the extent to which the highly discriminating soft information encoding jet quenching is irrecoverably lost due to the heavy-ion underlying event. Lastly,

we note that while our studies here focus specifically on jet substructure, we expect that similar techniques can be applied to full events in heavy-ion collisions or in electron-nucleus collisions at the future Electron-Ion Collider.

Data availability. The data sets used in this work can be found at:

https://zenodo.org/record/5758081#.Ya_PAS2cZTY.

We provide the particle four vectors and N -subjettiness values of jets in pp (Pythia 8 [70]) and AA (Jewel [72, 73]) collisions for the kinematics specified above. The data sets with and without the thermal background are included.

Acknowledgments

We would like to thank Raghav Elayavalli, Andrew Larkoski and Ben Nachman for helpful discussions and Ben Nachman for feedback on our manuscript. YL, JM, MP are supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, under the contract DE-AC02-05CH11231. YSL, MP and FR were supported by the LDRD Program of Lawrence Berkeley National Laboratory. FR is supported by the Simons Foundation under the Simons Bridge program for Postdoctoral Fellowships at SCGP and YITP, award number 815892 and the NSF, award number 1915093.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited. SCOAP³ supports the goals of the International Year of Basic Sciences for Sustainable Development.

References

- [1] A.J. Larkoski, I. Moult and B. Nachman, *Jet Substructure at the Large Hadron Collider: A Review of Recent Advances in Theory and Machine Learning*, *Phys. Rept.* **841** (2020) 1 [[arXiv:1709.04464](https://arxiv.org/abs/1709.04464)] [[INSPIRE](#)].
- [2] R. Kogler et al., *Jet Substructure at the Large Hadron Collider: Experimental Review*, *Rev. Mod. Phys.* **91** (2019) 045003 [[arXiv:1803.06991](https://arxiv.org/abs/1803.06991)] [[INSPIRE](#)].
- [3] S. Marzani, G. Soyez, and M. Spannowsky, *Looking inside jets. An introduction to jet substructure and boosted-object phenomenology*, *Lect. Notes Phys.* **958** (2019) 1.
- [4] W. Busza, K. Rajagopal and W. van der Schee, *Heavy Ion Collisions: The Big Picture, and the Big Questions*, *Ann. Rev. Nucl. Part. Sci.* **68** (2018) 339 [[arXiv:1802.04801](https://arxiv.org/abs/1802.04801)] [[INSPIRE](#)].
- [5] S. Weinberg, *The First Three Minutes. A Modern View of the Origin of the Universe*, Bantam Books (1977) [ISBN:978-0465024377].
- [6] M. Gyulassy and X.-n. Wang, *Multiple collisions and induced gluon Bremsstrahlung in QCD*, *Nucl. Phys. B* **420** (1994) 583 [[nuc1-th/9306003](https://arxiv.org/abs/nuc1-th/9306003)] [[INSPIRE](#)].
- [7] R. Baier, Y.L. Dokshitzer, A.H. Mueller, S. Peigne and D. Schiff, *Radiative energy loss and p_T broadening of high-energy partons in nuclei*, *Nucl. Phys. B* **484** (1997) 265 [[hep-ph/9608322](https://arxiv.org/abs/hep-ph/9608322)] [[INSPIRE](#)].

- [8] B.G. Zakharov, *Fully quantum treatment of the Landau-Pomeranchuk-Migdal effect in QED and QCD*, *JETP Lett.* **63** (1996) 952 [[hep-ph/9607440](#)] [[INSPIRE](#)].
- [9] M. Gyulassy, P. Levai and I. Vitev, *Reaction operator approach to nonAbelian energy loss*, *Nucl. Phys. B* **594** (2001) 371 [[nucl-th/0006010](#)] [[INSPIRE](#)].
- [10] X.-N. Wang and X.-f. Guo, *Multiple parton scattering in nuclei: Parton energy loss*, *Nucl. Phys. A* **696** (2001) 788 [[hep-ph/0102230](#)] [[INSPIRE](#)].
- [11] P.B. Arnold, G.D. Moore and L.G. Yaffe, *Photon and gluon emission in relativistic plasmas*, *JHEP* **06** (2002) 030 [[hep-ph/0204343](#)] [[INSPIRE](#)].
- [12] H. Liu, K. Rajagopal and U.A. Wiedemann, *Calculating the jet quenching parameter from AdS/CFT*, *Phys. Rev. Lett.* **97** (2006) 182301 [[hep-ph/0605178](#)] [[INSPIRE](#)].
- [13] G. Ovanessian and I. Vitev, *An effective theory for jet propagation in dense QCD matter: jet broadening and medium-induced bremsstrahlung*, *JHEP* **06** (2011) 080 [[arXiv:1103.1074](#)] [[INSPIRE](#)].
- [14] N. Armesto et al., *Comparison of Jet Quenching Formalisms for a Quark-Gluon Plasma 'Brick'*, *Phys. Rev. C* **86** (2012) 064904 [[arXiv:1106.1106](#)] [[INSPIRE](#)].
- [15] Y. Mehtar-Tani, J.G. Milhano and K. Tywoniuk, *Jet physics in heavy-ion collisions*, *Int. J. Mod. Phys. A* **28** (2013) 1340013 [[arXiv:1302.2579](#)] [[INSPIRE](#)].
- [16] JET collaboration, *Extracting the jet transport coefficient from jet quenching in high-energy heavy-ion collisions*, *Phys. Rev. C* **90** (2014) 014909 [[arXiv:1312.5003](#)] [[INSPIRE](#)].
- [17] J. Brewer, J.G. Milhano and J. Thaler, *Sorting out quenched jets*, *Phys. Rev. Lett.* **122** (2019) 222301 [[arXiv:1812.05111](#)] [[INSPIRE](#)].
- [18] J.-W. Qiu, F. Ringer, N. Sato and P. Zurita, *Factorization of jet cross sections in heavy-ion collisions*, *Phys. Rev. Lett.* **122** (2019) 252301 [[arXiv:1903.01993](#)] [[INSPIRE](#)].
- [19] JETSCAPE collaboration, *Determining the jet transport coefficient \hat{q} from inclusive hadron suppression measurements using Bayesian parameter estimation*, *Phys. Rev. C* **104** (2021) 024905 [[arXiv:2102.11337](#)] [[INSPIRE](#)].
- [20] P. Caucal, E. Iancu and G. Soyez, *Deciphering the z_g distribution in ultrarelativistic heavy ion collisions*, *JHEP* **10** (2019) 273 [[arXiv:1907.04866](#)] [[INSPIRE](#)].
- [21] C. Andres, L. Apolinário and F. Dominguez, *Medium-induced gluon radiation with full resummation of multiple scatterings for realistic parton-medium interactions*, *JHEP* **07** (2020) 114 [[arXiv:2002.01517](#)] [[INSPIRE](#)].
- [22] V. Vaidya and X. Yao, *Transverse momentum broadening of a jet in quark-gluon plasma: an open quantum system EFT*, *JHEP* **10** (2020) 024 [[arXiv:2004.11403](#)] [[INSPIRE](#)].
- [23] V. Vaidya, *Probing a dilute short lived Quark Gluon Plasma medium with jets*, [arXiv:2109.11568](#) [[INSPIRE](#)].
- [24] H. Clayton, M.D. Sievert and W.A. Horowitz, *Jet broadening in the opacity and twist expansions*, *Eur. Phys. J. C* **82** (2022) 437 [[arXiv:2110.14737](#)] [[INSPIRE](#)].
- [25] ATLAS collaboration, *Measurement of the nuclear modification factor for inclusive jets in Pb+Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV with the ATLAS detector*, *Phys. Lett. B* **790** (2019) 108 [[arXiv:1805.05635](#)] [[INSPIRE](#)].
- [26] ALICE collaboration, *Measurements of inclusive jet spectra in pp and central Pb-Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV*, *Phys. Rev. C* **101** (2020) 034911 [[arXiv:1909.09718](#)] [[INSPIRE](#)].

- [27] CMS collaboration, *First measurement of large area jet transverse momentum spectra in heavy-ion collisions*, *JHEP* **05** (2021) 284 [[arXiv:2102.13080](#)] [[INSPIRE](#)].
- [28] A LARGE ION COLLIDER EXPERIMENT, ALICE collaboration, *Measurement of the groomed jet radius and momentum splitting fraction in pp and Pb-Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV*, *Phys. Rev. Lett.* **128** (2022) 102001 [[arXiv:2107.12984](#)] [[INSPIRE](#)].
- [29] A.J. Larkoski and J. Thaler, *Unsafe but Calculable: Ratios of Angularities in Perturbative QCD*, *JHEP* **09** (2013) 137 [[arXiv:1307.1699](#)] [[INSPIRE](#)].
- [30] A.J. Larkoski, S. Marzani and J. Thaler, *Sudakov Safety in Perturbative QCD*, *Phys. Rev. D* **91** (2015) 111501 [[arXiv:1502.01719](#)] [[INSPIRE](#)].
- [31] M. Procura, W.J. Waalewijn and L. Zeune, *Joint resummation of two angularities at next-to-next-to-leading logarithmic order*, *JHEP* **10** (2018) 098 [[arXiv:1806.10622](#)] [[INSPIRE](#)].
- [32] Y.L. Dokshitzer and B.R. Webber, *Calculation of power corrections to hadronic event shapes*, *Phys. Lett. B* **352** (1995) 451 [[hep-ph/9504219](#)] [[INSPIRE](#)].
- [33] G.P. Korchemsky and G.F. Sterman, *Power corrections to event shapes and factorization*, *Nucl. Phys. B* **555** (1999) 335 [[hep-ph/9902341](#)] [[INSPIRE](#)].
- [34] C. Lee and G.F. Sterman, *Momentum Flow Correlations from Event Shapes: Factorized Soft Gluons and Soft-Collinear Effective Theory*, *Phys. Rev. D* **75** (2007) 014022 [[hep-ph/0611061](#)] [[INSPIRE](#)].
- [35] M. Dasgupta, L. Magnea and G.P. Salam, *Non-perturbative QCD effects in jets at hadron colliders*, *JHEP* **02** (2008) 055 [[arXiv:0712.3014](#)] [[INSPIRE](#)].
- [36] Y. Makris, D. Neill and V. Vaidya, *Probing Transverse-Momentum Dependent Evolution With Groomed Jets*, *JHEP* **07** (2018) 167 [[arXiv:1712.07653](#)] [[INSPIRE](#)].
- [37] A.H. Hoang, S. Mantry, A. Pathak and I.W. Stewart, *Nonperturbative Corrections to Soft Drop Jet Mass*, *JHEP* **12** (2019) 002 [[arXiv:1906.11843](#)] [[INSPIRE](#)].
- [38] P. Cal, D. Neill, F. Ringer and W.J. Waalewijn, *Calculating the angle between jet axes*, *JHEP* **04** (2020) 211 [[arXiv:1911.06840](#)] [[INSPIRE](#)].
- [39] H. Chen, I. Moulst, X. Zhang and H.X. Zhu, *Rethinking jets with energy correlators: Tracks, resummation, and analytic continuation*, *Phys. Rev. D* **102** (2020) 054012 [[arXiv:2004.11381](#)] [[INSPIRE](#)].
- [40] ALICE collaboration, *Measurement of jet quenching with semi-inclusive hadron-jet distributions in central Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, *JHEP* **09** (2015) 170 [[arXiv:1506.03984](#)] [[INSPIRE](#)].
- [41] J. Mulligan and M. Ploskon, *Identifying groomed jet splittings in heavy-ion collisions*, *Phys. Rev. C* **102** (2020) 044913 [[arXiv:2006.01812](#)] [[INSPIRE](#)].
- [42] JETSCAPE collaboration, *Multisystem Bayesian constraints on the transport coefficients of QCD matter*, *Phys. Rev. C* **103** (2021) 054904 [[arXiv:2011.01430](#)] [[INSPIRE](#)].
- [43] Y.S. Lai, *Automated Discovery of Jet Substructure Analyses*, [arXiv:1810.00835](#) [[INSPIRE](#)].
- [44] L. de Oliveira, M. Kagan, L. Mackey, B. Nachman and A. Schwartzman, *Jet-images — deep learning edition*, *JHEP* **07** (2016) 069 [[arXiv:1511.05190](#)] [[INSPIRE](#)].
- [45] P.T. Komiske, E.M. Metodiev and M.D. Schwartz, *Deep learning in color: towards automated quark/gluon jet discrimination*, *JHEP* **01** (2017) 110 [[arXiv:1612.01551](#)] [[INSPIRE](#)].

- [46] L.-G. Pang, K. Zhou, N. Su, H. Petersen, H. Stöcker and X.-N. Wang, *An equation-of-state-meter of quantum chromodynamics transition from deep learning*, *Nature Commun.* **9** (2018) 210 [[arXiv:1612.04262](#)] [[INSPIRE](#)].
- [47] J.W. Monk, *Deep Learning as a Parton Shower*, *JHEP* **12** (2018) 021 [[arXiv:1807.03685](#)] [[INSPIRE](#)].
- [48] A.J. Larkoski and E.M. Metodiev, *A Theory of Quark vs. Gluon Discrimination*, *JHEP* **10** (2019) 014 [[arXiv:1906.01639](#)] [[INSPIRE](#)].
- [49] G. Kasieczka, S. Marzani, G. Soyez and G. Stagnitto, *Towards Machine Learning Analytics for Jet Substructure*, *JHEP* **09** (2020) 195 [[arXiv:2007.04319](#)] [[INSPIRE](#)].
- [50] Y.S. Lai, D. Neill, M. Płoskoń and F. Ringer, *Explainable machine learning of the underlying physics of high-energy particle collisions*, *Phys. Lett. B* **829** (2022) 137055 [[arXiv:2012.06582](#)] [[INSPIRE](#)].
- [51] S. Bieringer et al., *Measuring QCD Splittings with Invertible Networks*, *SciPost Phys.* **10** (2021) 126 [[arXiv:2012.09873](#)] [[INSPIRE](#)].
- [52] E. Bernreuther, T. Finke, F. Kahlhoefer, M. Krämer and A. Mück, *Casting a graph net to catch dark showers*, *SciPost Phys.* **10** (2021) 046 [[arXiv:2006.08639](#)] [[INSPIRE](#)].
- [53] Y. Huang, L.-G. Pang, X. Luo and X.-N. Wang, *Probing criticality with deep learning in relativistic heavy-ion collisions*, *Phys. Lett. B* **827** (2022) 137001 [[arXiv:2107.11828](#)] [[INSPIRE](#)].
- [54] A. Andreassen, I. Feige, C. Frye and M.D. Schwartz, *JUNIPR: a Framework for Unsupervised Machine Learning in Particle Physics*, *Eur. Phys. J. C* **79** (2019) 102 [[arXiv:1804.09720](#)] [[INSPIRE](#)].
- [55] Y. Alanazi et al., *Machine learning-based event generator for electron-proton scattering*, [arXiv:2008.03151](#) [[INSPIRE](#)].
- [56] T. Faucett, J. Thaler and D. Whiteson, *Mapping Machine-Learned Physics into a Human-Readable Space*, *Phys. Rev. D* **103** (2021) 036020 [[arXiv:2010.11998](#)] [[INSPIRE](#)].
- [57] F.A. Dreyer and H. Qu, *Jet tagging in the Lund plane with graph networks*, *JHEP* **03** (2021) 052 [[arXiv:2012.08526](#)] [[INSPIRE](#)].
- [58] H1 collaboration, *Measurement of Lepton-Jet Correlation in Deep-Inelastic Scattering with the H1 Detector Using Machine Learning for Unfolding*, *Phys. Rev. Lett.* **128** (2022) 132002 [[arXiv:2108.12376](#)] [[INSPIRE](#)].
- [59] A. Ghosh, B. Nachman and D. Whiteson, *Uncertainty-aware machine learning for high energy physics*, *Phys. Rev. D* **104** (2021) 056026 [[arXiv:2105.08742](#)] [[INSPIRE](#)].
- [60] P. Konar, V.S. Ngairangbam and M. Spannowsky, *Energy-weighted message passing: an infra-red and collinear safe graph neural network algorithm*, *JHEP* **02** (2022) 060 [[arXiv:2109.14636](#)] [[INSPIRE](#)].
- [61] B.M. Dillon, G. Kasieczka, H. Otschlag, T. Plehn, P. Sorrenson and L. Vogel, *Symmetries, safety, and self-supervision*, *SciPost Phys.* **12** (2022) 188 [[arXiv:2108.04253](#)] [[INSPIRE](#)].
- [62] Y.-T. Chien and R. Kunnawalkam Elayavalli, *Probing heavy ion collisions using quark and gluon jet substructure*, [arXiv:1803.03589](#) [[INSPIRE](#)].
- [63] Y.-L. Du, D. Pablos and K. Tywoniuk, *Jet Tomography in Heavy-Ion Collisions with Deep Learning*, *Phys. Rev. Lett.* **128** (2022) 012301 [[arXiv:2106.11271](#)] [[INSPIRE](#)].

- [64] L. Apolinário, N.F. Castro, M. Crispim Romão, J.G. Milhano, R. Pedro and F.C.R. Peres, *Deep Learning for the classification of quenched jets*, *JHEP* **11** (2021) 219 [[arXiv:2106.08869](#)] [[INSPIRE](#)].
- [65] F. Santosa and W.W. Symes, *Linear inversion of band-limited reflection seismograms*, *SIAM J. Sci. Stat. Comput.* **7** (1986) 1307.
- [66] R. Tibshirani, *Regression Shrinkage and Selection via the Lasso*, *J. Roy. Stat. Soc. B* **58** (1996) 267.
- [67] 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** (2020) 182001 [[arXiv:1911.09107](#)] [[INSPIRE](#)].
- [68] JETSCAPE collaboration, *Determining the jet transport coefficient \hat{q} from inclusive hadron suppression measurements using Bayesian parameter estimation*, *Phys. Rev. C* **104** (2021) 024905 [[arXiv:2102.11337](#)] [[INSPIRE](#)].
- [69] R. Abdul Khalek et al., *Science Requirements and Detector Concepts for the Electron-Ion Collider: EIC Yellow Report*, *Nucl. Phys. A* **1026** (2022) 122447 [[arXiv:2103.05419](#)] [[INSPIRE](#)].
- [70] T. Sjöstrand, S. Mrenna and P.Z. Skands, *A Brief Introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852 [[arXiv:0710.3820](#)] [[INSPIRE](#)].
- [71] P. Skands, S. Carrazza and J. Rojo, *Tuning PYTHIA 8.1: the Monash 2013 Tune*, *Eur. Phys. J. C* **74** (2014) 3024 [[arXiv:1404.5630](#)] [[INSPIRE](#)].
- [72] K.C. Zapp, F. Krauss and U.A. Wiedemann, *A perturbative framework for jet quenching*, *JHEP* **03** (2013) 080 [[arXiv:1212.1599](#)] [[INSPIRE](#)].
- [73] K.C. Zapp, *JEWEL 2.0.0: directions for use*, *Eur. Phys. J. C* **74** (2014) 2762 [[arXiv:1311.0048](#)] [[INSPIRE](#)].
- [74] R. Kunnawalkam Elayavalli and K.C. Zapp, *Medium response in JEWEL and its impact on jet shape observables in heavy ion collisions*, *JHEP* **07** (2017) 141 [[arXiv:1707.01539](#)] [[INSPIRE](#)].
- [75] M. Cacciari, G.P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063 [[arXiv:0802.1189](#)] [[INSPIRE](#)].
- [76] M. Dasgupta, F. Dreyer, G.P. Salam and G. Soyez, *Small-radius jets to all orders in QCD*, *JHEP* **04** (2015) 039 [[arXiv:1411.5182](#)] [[INSPIRE](#)].
- [77] T. Kaufmann, A. Mukherjee and W. Vogelsang, *Hadron Fragmentation Inside Jets in Hadronic Collisions*, *Phys. Rev. D* **92** (2015) 054015 [*Erratum ibid.* **101** (2020) 079901] [[arXiv:1506.01415](#)] [[INSPIRE](#)].
- [78] Z.-B. Kang, F. Ringer and I. Vitev, *The semi-inclusive jet function in SCET and small radius resummation for inclusive jet production*, *JHEP* **10** (2016) 125 [[arXiv:1606.06732](#)] [[INSPIRE](#)].
- [79] L. Dai, C. Kim and A.K. Leibovich, *Fragmentation of a Jet with Small Radius*, *Phys. Rev. D* **94** (2016) 114023 [[arXiv:1606.07411](#)] [[INSPIRE](#)].
- [80] T. Renk, *Biased showers: A common conceptual framework for the interpretation of high- P_T observables in heavy-ion collisions*, *Phys. Rev. C* **88** (2013) 054902 [[arXiv:1212.0646](#)] [[INSPIRE](#)].

- [81] ALICE collaboration, *Measurement of Event Background Fluctuations for Charged Particle Jet Reconstruction in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, *JHEP* **03** (2012) 053 [[arXiv:1201.2423](#)] [[INSPIRE](#)].
- [82] P. Berta, M. Spousta, D.W. Miller and R. Leitner, *Particle-level pileup subtraction for jets and jet shapes*, *JHEP* **06** (2014) 092 [[arXiv:1403.3108](#)] [[INSPIRE](#)].
- [83] P. Berta, L. Masetti, D.W. Miller and M. Spousta, *Pileup and Underlying Event Mitigation with Iterative Constituent Subtraction*, *JHEP* **08** (2019) 175 [[arXiv:1905.03470](#)] [[INSPIRE](#)].
- [84] B. Nachman and J. Thaler, *Learning from many collider events at once*, *Phys. Rev. D* **103** (2021) 116013 [[arXiv:2101.07263](#)] [[INSPIRE](#)].
- [85] M. Zaheer, S. Kottur, S. Ravanbakhsh, B. Póczos, R. Salakhutdinov and A.J. Smola, *Deep sets*, [arXiv:1703.06114](#).
- [86] E. Wagstaff, F.B. Fuchs, M. Engelcke, I. Posner and M.A. Osborne, *On the limitations of representing functions on sets*, [arXiv:1901.09006](#).
- [87] B. Bloem-Reddy and Y. Teh, *Probabilistic symmetries and invariant neural networks*, *J. Mach. Learn. Res.* **21** (2020) 1 [[arXiv:1901.06082](#)].
- [88] P.T. Komiske, E.M. Metodiev and J. Thaler, *Energy Flow Networks: Deep Sets for Particle Jets*, *JHEP* **01** (2019) 121 [[arXiv:1810.05165](#)] [[INSPIRE](#)].
- [89] K. Datta and A. Larkoski, *How Much Information is in a Jet?*, *JHEP* **06** (2017) 073 [[arXiv:1704.08249](#)] [[INSPIRE](#)].
- [90] K. Datta and A.J. Larkoski, *Novel Jet Observables from Machine Learning*, *JHEP* **03** (2018) 086 [[arXiv:1710.01305](#)] [[INSPIRE](#)].
- [91] K. Datta, A. Larkoski and B. Nachman, *Automating the Construction of Jet Observables with Machine Learning*, *Phys. Rev. D* **100** (2019) 095016 [[arXiv:1902.07180](#)] [[INSPIRE](#)].
- [92] J. Thaler and K. Van Tilburg, *Identifying Boosted Objects with N -subjettiness*, *JHEP* **03** (2011) 015 [[arXiv:1011.2268](#)] [[INSPIRE](#)].
- [93] J. Thaler and K. Van Tilburg, *Maximizing Boosted Top Identification by Minimizing N -subjettiness*, *JHEP* **02** (2012) 093 [[arXiv:1108.2701](#)] [[INSPIRE](#)].
- [94] I.W. Stewart, F.J. Tackmann and W.J. Waalewijn, *N -Jettiness: An Inclusive Event Shape to Veto Jets*, *Phys. Rev. Lett.* **105** (2010) 092002 [[arXiv:1004.2489](#)] [[INSPIRE](#)].
- [95] A.J. Larkoski and I. Moulton, *The Singular Behavior of Jet Substructure Observables*, *Phys. Rev. D* **93** (2016) 014017 [[arXiv:1510.08459](#)] [[INSPIRE](#)].
- [96] D. Napoletano and G. Soyez, *Computing N -subjettiness for boosted jets*, *JHEP* **12** (2018) 031 [[arXiv:1809.04602](#)] [[INSPIRE](#)].
- [97] M. Dasgupta and J. Helliwell, *Investigating top tagging with Y_m -Splitter and N -subjettiness*, *JHEP* **10** (2021) 092 [[arXiv:2108.09317](#)] [[INSPIRE](#)].
- [98] P.T. Komiske, E.M. Metodiev and J. Thaler, *Energy flow polynomials: A complete linear basis for jet substructure*, *JHEP* **04** (2018) 013 [[arXiv:1712.07124](#)] [[INSPIRE](#)].
- [99] H. Qu and L. Gouskos, *ParticleNet: Jet Tagging via Particle Clouds*, *Phys. Rev. D* **101** (2020) 056019 [[arXiv:1902.08570](#)] [[INSPIRE](#)].
- [100] M.J. Dolan and A. Ore, *Equivariant Energy Flow Networks for Jet Tagging*, *Phys. Rev. D* **103** (2021) 074022 [[arXiv:2012.00964](#)] [[INSPIRE](#)].

- [101] ATLAS collaboration, *Deep Sets based Neural Networks for Impact Parameter Flavour Tagging in ATLAS*, [ATL-PHYS-PUB-2020-014](#) (2020).
- [102] N. Segol and Y. Lipman, *On universal equivariant set networks*, [arXiv:1910.02421](#).
- [103] A. Romero, D. Whiteson, M. Fenton, J. Collado and P. Baldi, *Safety of Quark/Gluon Jet Classification*, [arXiv:2103.09103](#) [INSPIRE].
- [104] V. Mikuni and F. Canelli, *Point cloud transformers applied to collider physics*, *Mach. Learn. Sci. Tech.* **2** (2021) 035027 [[arXiv:2102.05073](#)] [INSPIRE].
- [105] B. Ostdiek, *Deep Set Auto Encoders for Anomaly Detection in Particle Physics*, *SciPost Phys.* **12** (2022) 045 [[arXiv:2109.01695](#)] [INSPIRE].
- [106] F. Chollet et al., *Keras*, <https://github.com/fchollet/keras> (2015).
- [107] M. Abadi et al., *TensorFlow: Large-scale machine learning on heterogeneous systems*, software available from [tensorflow.org](https://www.tensorflow.org) (2015).
- [108] V. Nair and G.E. Hinton, *Rectified linear units improve restricted boltzmann machines*, in *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, J. Fürnkranz and T. Joachims eds., Haifa, Israel (2010), pg. 807.
- [109] D.P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, 12, 2014 [[arXiv:1412.6980](#)] [INSPIRE].
- [110] D.R. Cox, *The regression analysis of binary sequences*, *J. Roy. Stat. Soc. B* **20** (1958) 215.
- [111] G.C. Blazey et al., *Run II jet physics*, in *Physics at Run II: QCD and Weak Boson Physics Workshop: Final General Meeting*, Batavia, U.S.A. (2000), pp. 47 [[hep-ex/0005012](#)] [INSPIRE].
- [112] P. Komiske, E. Metodiev and J. Thaler, *Pythia8 quark and gluon jets for energy flow*, *Zenodo* (2019).
- [113] P. Cal, K. Lee, F. Ringer and W.J. Waalewijn, *The soft drop momentum sharing fraction z_g beyond leading-logarithmic accuracy*, *Phys. Lett. B* **833** (2022) 137390 [[arXiv:2106.04589](#)] [INSPIRE].
- [114] J. Gallicchio and M.D. Schwartz, *Seeing in Color: Jet Superstructure*, *Phys. Rev. Lett.* **105** (2010) 022001 [[arXiv:1001.5027](#)] [INSPIRE].
- [115] A.J. Larkoski, S. Marzani and C. Wu, *Theory Predictions for the Pull Angle*, *Phys. Rev. D* **99** (2019) 091502 [[arXiv:1903.02275](#)] [INSPIRE].
- [116] J.C. Collins, D.E. Soper and G.F. Sterman, *Factorization of Hard Processes in QCD*, *Adv. Ser. Direct. High Energy Phys.* **5** (1989) 1 [[hep-ph/0409313](#)] [INSPIRE].
- [117] D. Krohn, M.D. Schwartz, T. Lin and W.J. Waalewijn, *Jet Charge at the LHC*, *Phys. Rev. Lett.* **110** (2013) 212001 [[arXiv:1209.2421](#)] [INSPIRE].
- [118] W.J. Waalewijn, *Calculating the Charge of a Jet*, *Phys. Rev. D* **86** (2012) 094030 [[arXiv:1209.3019](#)] [INSPIRE].
- [119] P. Cal, J. Thaler and W.J. Waalewijn, *Power counting energy flow polynomials*, *JHEP* **09** (2022) 021 [[arXiv:2205.06818](#)] [INSPIRE].
- [120] S.D. Ellis and D.E. Soper, *Successive combination jet algorithm for hadron collisions*, *Phys. Rev. D* **48** (1993) 3160 [[hep-ph/9305266](#)] [INSPIRE].
- [121] S. Catani, Y.L. Dokshitzer, M.H. Seymour and B.R. Webber, *Longitudinally invariant K_t clustering algorithms for hadron hadron collisions*, *Nucl. Phys. B* **406** (1993) 187 [INSPIRE].

- [122] M. Cacciari, G.P. Salam and G. Soyez, *FastJet User Manual*, *Eur. Phys. J. C* **72** (2012) 1896 [[arXiv:1111.6097](#)] [[INSPIRE](#)].
- [123] A.J. Larkoski, S. Marzani, G. Soyez and J. Thaler, *Soft Drop*, *JHEP* **05** (2014) 146 [[arXiv:1402.2657](#)] [[INSPIRE](#)].
- [124] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh and A. Talwalkar, *Hyperband: A novel bandit-based approach to hyperparameter optimization*, *J. Mach. Learn. Res.* **18** (2017) 6765 [[arXiv:1603.06560](#)].
- [125] T. O'Malley et al., *Kerastuner*, <https://github.com/keras-team/keras-tuner> (2019).
- [126] D. Neill and W.J. Waalewijn, *Entropy of a Jet*, *Phys. Rev. Lett.* **123** (2019) 142001 [[arXiv:1811.01021](#)] [[INSPIRE](#)].
- [127] F. Pedregosa et al., *Scikit-learn: Machine learning in Python*, *J. Mach. Learn. Res.* **12** (2011) 2825
- [128] A.J. Larkoski, J. Thaler and W.J. Waalewijn, *Gaining (Mutual) Information about Quark/Gluon Discrimination*, *JHEP* **11** (2014) 129 [[arXiv:1408.3122](#)] [[INSPIRE](#)].
- [129] M. Cranmer et al., *Discovering Symbolic Models from Deep Learning with Inductive Biases*, [arXiv:2006.11287](#) [[INSPIRE](#)].
- [130] A. Butter, T. Plehn, N. Soybelman and J. Brehmer, *Back to the Formula — LHC Edition*, [arXiv:2109.10414](#) [[INSPIRE](#)].
- [131] CMS collaboration, *Measurement of the Splitting Function in pp and Pb-Pb Collisions at $\sqrt{s_{NN}} = 5.02$ TeV*, *Phys. Rev. Lett.* **120** (2018) 142302 [[arXiv:1708.09429](#)] [[INSPIRE](#)].
- [132] Z.-B. Kang, K. Lee, X. Liu, D. Neill and F. Ringer, *The soft drop groomed jet radius at NLL*, *JHEP* **02** (2020) 054 [[arXiv:1908.01783](#)] [[INSPIRE](#)].
- [133] ATLAS collaboration, *Measurement of soft-drop jet observables in pp collisions with the ATLAS detector at $\sqrt{s} = 13$ TeV*, *Phys. Rev. D* **101** (2020) 052007 [[arXiv:1912.09837](#)] [[INSPIRE](#)].
- [134] STAR collaboration, *Differential measurements of jet substructure and partonic energy loss in Au+Au collisions at $\sqrt{S_{NN}} = 200$ GeV*, *Phys. Rev. C* **105** (2022) 044906 [[arXiv:2109.09793](#)] [[INSPIRE](#)].
- [135] ALICE collaboration, *First measurement of jet mass in Pb-Pb and p-Pb collisions at the LHC*, *Phys. Lett. B* **776** (2018) 249 [[arXiv:1702.00804](#)] [[INSPIRE](#)].
- [136] CMS collaboration, *Measurement of the groomed jet mass in PbPb and pp collisions at $\sqrt{s_{NN}} = 5.02$ TeV*, *JHEP* **10** (2018) 161 [[arXiv:1805.05145](#)] [[INSPIRE](#)].
- [137] ALICE collaboration, *Medium modification of the shape of small-radius jets in central Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, *JHEP* **10** (2018) 139 [[arXiv:1807.06854](#)] [[INSPIRE](#)].
- [138] R. Haake and C. Loizides, *Machine Learning based jet momentum reconstruction in heavy-ion collisions*, *Phys. Rev. C* **99** (2019) 064904 [[arXiv:1810.06324](#)] [[INSPIRE](#)].
- [139] J. Neyman and E.S. Pearson, *On the Problem of the Most Efficient Tests of Statistical Hypotheses*, *Philos. Trans. R. Soc. Lond. A* **231** (1933) 289.