

2024

Is Infrared-Collinear Safe Information All You Need for Jet Classification?

Dimitrios Athanasakos
Stony Brook University

Andrew J. Larkoski
University of California

James Mulligan
University of California

Mateusz Płoskoń
Lawrence Berkeley National Laboratory

Felix Ringer
Old Dominion University, fringer@odu.edu

Follow this and additional works at: https://digitalcommons.odu.edu/physics_fac_pubs



Part of the [Artificial Intelligence and Robotics Commons](#), and the [Nuclear Commons](#)

Original Publication Citation

Athanasakos, D., Larkoski, A. J., Mulligan, J., Płoskoń, M., & Ringer, F. (2024). Is infrared-collinear safe information all you need for jet classification? *Journal of High Energy Physics*, 2024(7), 1-28, Article 257. [https://doi.org/10.1007/JHEP07\(2024\)257](https://doi.org/10.1007/JHEP07(2024)257)

This Article is brought to you for free and open access by the Physics at ODU Digital Commons. It has been accepted for inclusion in Physics Faculty Publications by an authorized administrator of ODU Digital Commons. For more information, please contact digitalcommons@odu.edu.

Is infrared-collinear safe information all you need for jet classification?

Dimitrios Athanasakos ^{a,b}, Andrew J. Larkoski ^c, James Mulligan ^{d,e},
Mateusz Płoskoń ^d and Felix Ringer ^{a,b,f,g}

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

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

^c*Department of Physics and Astronomy, University of California, Los Angeles, CA 90095, U.S.A.*

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

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

^f*Department of Physics, Old Dominion University, Norfolk, VA 23529, U.S.A.*

^g*Thomas Jefferson National Accelerator Facility, Newport News, VA 23606, U.S.A.*

E-mail: dimitrios.athanasakos@stonybrook.edu, larkoski@ucla.edu,
james.mulligan@berkeley.edu, mploskon@lbl.gov, fmringer@jlab.org

ABSTRACT: Machine learning-based jet classifiers are able to achieve impressive tagging performance in a variety of applications in high-energy and nuclear physics. However, it remains unclear in many cases which aspects of jets give rise to this discriminating power, and whether jet observables that are tractable in perturbative QCD such as those obeying infrared-collinear (IRC) safety serve as sufficient inputs. In this article, we introduce a new classifier, Jet Flow Networks (JFNs), in an effort to address the question of whether IRC unsafe information provides additional discriminating power in jet classification. JFNs are permutation-invariant neural networks (deep sets) that take as input the kinematic information of reconstructed subjects. The subject radius and a cut on the subject's transverse momenta serve as tunable hyperparameters enabling a controllable sensitivity to soft emissions and nonperturbative effects. We demonstrate the performance of JFNs for quark vs. gluon and Z vs. QCD jet tagging. For small subject radii and transverse momentum cuts, the performance of JFNs is equivalent to the IRC-unsafe Particle Flow Networks (PFNs), demonstrating that infrared-collinear unsafe information is not necessary to achieve strong discrimination for both cases. As the subject radius is increased, the performance of the JFNs remains essentially unchanged until physical thresholds that we identify are crossed. For relatively large subject radii, we show that the JFNs may offer an increased model independence with a modest tradeoff in performance compared to classifiers that use the full particle information of the jet. These results shed new light on how machines learn patterns in high-energy physics data.

KEYWORDS: Automation, Jets and Jet Substructure

ARXIV EPRINT: [2305.08979](https://arxiv.org/abs/2305.08979)

Contents

1	Introduction	1
2	The subjet basis	5
3	Jet Flow Networks (JFNs): deep sets of subjets	6
4	Data sets	8
5	JFN performance: gapless jet classification	9
5.1	Collinear safety	9
5.2	Soft safety	11
6	Learning physical scales	13
7	Performance vs. generalizability	18
8	Conclusions	21

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) [1–3]. Jets arise from the fragmentation of highly energetic quarks and gluons, which themselves can arise from the decay of unstable particles such as the Higgs boson. Classifying the origins of jets, such as quark vs. gluon initiated jets, QCD vs. boosted Z/W jets, and QCD vs. boosted top jets [4–24] is crucial to disentangle the various processes occurring at collider experiments and perform searches for physics beyond the Standard Model.

Jet classification algorithms have been developed based on multivariate combinations of jet substructure observables as well as using machine learning methods. Machine learning based jet classifiers significantly outperform traditional multivariate jet taggers that utilize a limited number of observables, since they are able to leverage the full information in the jet. However, machine learning based classifiers often have the drawback that they are not calculable by analytical methods. Efforts to address this have been an active area of research, such as enforcing Infrared-Collinear (IRC) safety [25] in the network architecture [26–30] or by finding optimal ways to reduce the amount of information provided as input to neural networks [31–35]. IRC safety is a frequently used guiding principle in high-energy physics for constructing suitable observables. It is often stated as: “An observable is IRC safe if it is insensitive to infinitesimally soft or exactly collinear emissions” [36]. It is a defining feature of observables that allows for calculations in QCD where divergences cancel at each order in perturbation theory [1, 37]. We note one has to distinguish between the IRC safety of an observable, i.e. the calculability or traceability of an observable within perturbative QCD,

and the question about the relevance of nonperturbative effects. While both questions are important, we will primarily focus on IRC safety throughout this work.

In order to increase the interpretability of machine learning based classifiers, a complete IRC-safe basis of jet substructure observables was introduced in refs. [10, 38, 39] based on N -subjettiness observables [40, 41]. These observables capture the momentum and relative angles of emissions inside the jet. The set of N -subjettiness observables is then used as input to a machine learning algorithm for jet classification. While the complete basis of IRC-safe observables is large ($3M - 4$ for M particles in the jet), it was found that the performance of classifiers saturates quickly with a relatively small number of observables. Another set of observables, Energy Flow Polynomials (EFPs), was developed as a linear and IRC-safe basis of jet substructure observables in ref. [26].

Interestingly, it was found that while the performance of classifiers based on complete sets of observables saturates, in most cases there remains a performance gap between classifiers with IRC-safe inputs (Sudakov safe classifiers) and IRC-unsafe classifiers that make use of the full information content of the particles inside the jet. Examples of such IRC-unsafe classifiers include architectures based on deep sets [11], point clouds [12, 42] and transformers [43]. This performance gap has been observed for a variety of jet classification tasks, including QCD vs. W/Z and H jets [35], and pp vs. AA jets [31]. For quark vs. gluon tagging, it was found that the IRC-safe EFPs can match the performance of PFNs when only momentum information of particles in the jet is considered [11]. Several efforts have been made to quantify the gap, with the aim to gain new insights into fundamental QCD dynamics. There are several possible explanations for the observed performance gap:

- IRC-unsafe classifiers may be able to make use of the very soft information content of jets, which is difficult to access with IRC-safe observables.
- IRC-unsafe classifiers such as PFNs take as input the exact position information of the particles inside the jet, whereas IRC-safe observables can only capture the information of relative distances. It is possible that existing machine learning algorithms can make more efficient use of position information.
- The specific form of the IRC-safe observables may not be optimal for classification tasks and there may be other sets of observables that could perform better.

With this question in mind, we introduce in this work a new machine learning-based jet classifier, Jet Flow Networks (JFNs),¹ which take as input the energy and position of reclustered subjects instead of individual particles. JFNs allow for soft and collinear emissions to be clustered into subjects making the input IRC-safe and the resulting classifier generally Sudakov safe [44, 45]. Sudakov safe observables are not IRC-safe, and as such are not defined at any order in the strong coupling α_s and yet have finite cross sections when all-orders effects are included. This happens because the Sudakov form factor regulates the infrared divergences. An example of such an observable is the ratio formed from two different angularities [46] measured on the same jet. However, different than the N -subjettiness or the EFP basis of observables, position information is used instead of having (indirectly) access only to relative

¹In analogy to Particle Flow Networks (PFNs) [11].

distances between emissions (or subjets) inside the jet. We note that we do not consider quark flavor tagging in this work which requires IRC-unsafe information, see e.g. refs. [47–49]. JFNs are closely related to Particle Flow Networks (PFNs) [11] and Energy Flow Networks (EFNs) [11], which will be elaborated on in section 3. In the limit of a vanishing subjet radius, where every subjet contains only a single hadron, and in the limit of a vanishingly small cut on the transverse momenta of the subjets, JFNs are identical to PFNs. The radius of the reclustered subjets in JFNs, as well as the cut on the transverse momenta, can be used to dial in nonperturbative information allowing for a smooth transition to IRC-unsafe classifiers. As such, JFNs complement the existing family of permutation-invariant networks in particle physics. As for PFNs and EFNs, we will utilize machine learning algorithms for JFNs based on a permutation invariant deep set architecture [11, 50–52].

In this paper, we will explicitly study classification tasks of quark vs. gluon jets and jets from QCD vs. jets from boosted hadronic decays of Z bosons. The particular IRC (un)safety of the likelihood ratios for these tasks has been studied previously. It is expected that the likelihood ratio for quark vs. gluon jet discrimination is IRC safe [29], which has been validated in previous machine learning studies [27, 28, 53]. The argument for IRC safety is that the N -body phase space can be spanned by additive, IRC safe observables [54] and so probability distributions for quarks and gluons will in general take the form of a Sudakov factor near the infrared regions of phase space. The rate of suppression of emissions is controlled by the corresponding color Casimirs, and, because gluons carry more color than quarks, the likelihood ratio itself vanishes at all infrared boundaries. Therefore, all fixed-order divergences are mapped to the same value of the likelihood ratio, namely 0, and so the classifier is IRC safe. This explains why EFPs, which are an IRC-safe classifier, can match the performance of the IRC-unsafe PFNs. By contrast, the likelihood ratio for QCD jets vs. Z jets is expected to only be Sudakov safe, as optimal observables for general one-vs. two-prong discrimination take the form of the ratio of IRC safe observables [40, 55–57]. Ratios of IRC safe observables are in general themselves not IRC safe [58].

To separately study soft and collinear safety, we consider both a finite subjet radius (collinear safety) and a low momentum cut p_T^{soft} on the transverse momentum of the subjets (soft safety). The main result of our work will be to show that the JFNs based on IRC-safe input achieve the same classification performance as PFNs. The JFNs represent the first example of a classifier based on IRC-safe inputs that achieve equivalent performance on several classification tasks as its IRC-unsafe counterpart. In addition, the machine learning architecture is equivalent to PFNs, which allows for one-to-one comparisons. The exact value of the subjet radius where the PFN performance is matched depends on the classification task at hand. Therefore, different than the classifiers based on complete IRC-safe sets of observables, JFNs constitute a “gapless” classifier indicating that the very soft aspects of jets are in fact not relevant for typical classification tasks at collider experiments. This answers in part the question about the features that are relevant for the performance of classifiers in high-energy physics. Throughout this work, PFNs are taken as a reference, but other permutation invariant classifiers such as GNNs, transformers (equivalent to a fully connected graph), and point clouds could equally well be trained on particles or reclustered subjets.

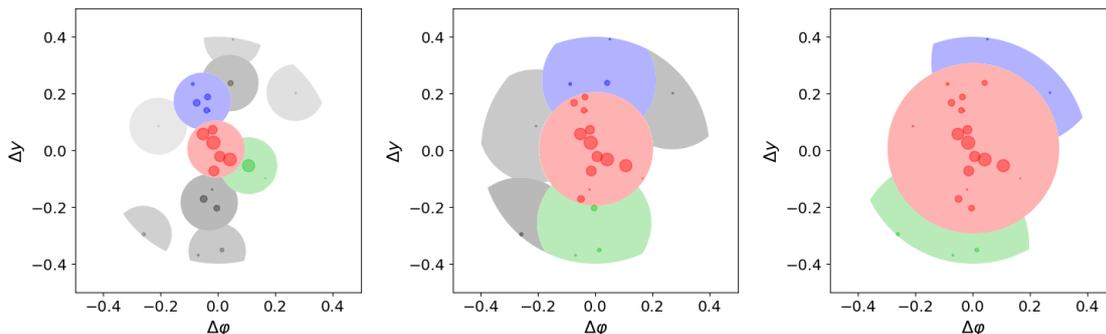


Figure 1. Illustration of a QCD jet with $p_T = 100$ GeV and radius parameter $R = 0.4$ reclustered into subjects for subjet radii $r = 0.1$ (left), $r = 0.2$ (middle), and $r = 0.3$ (right). We use the inclusive anti- k_T algorithm to identify the initial jet and the subjects. Particles are represented by small filled circles with radii proportional to the particle’s transverse momentum in the Δy vs. $\Delta\varphi$ plane, where $\Delta\varphi = \varphi^{\text{particle|subjet}} - \varphi^{\text{jet}}$ is the azimuthal angle with respect to the jet axis and $\Delta y = \Delta y^{\text{particle|subjet}} - \Delta y^{\text{jet}}$ is rapidity distance to the jet axis. Subjects are shown with larger colored areas where red marks the leading subjet, green marks the second leading subjet, blue marks the third leading jet, and shades of gray represent subjects with lower longitudinal momentum fraction $z = p_T^{\text{subjet}}/p_T$ with intensity proportional to z .

In addition to shedding light on the role of IRC-safe information, JFNs allow for new insights into the physics of jet tagging and may lead to various future applications at high-energy collider experiments. By studying the performance of the JFNs as a function of the subjet radius and the jet transverse momentum, we are able to identify the relevant physical scales of different classification tasks. For example, for Z vs. QCD-jet tagging, we find that the subjet scale $p_T r$ is sensitive the opening angle between the boosted hadronic decay products of the Z -boson. Second, we explore the generalization capability of JFNs to unseen data, which is crucial when deploying a classifier trained on simulations to experimental data. Due to the clustering of collinear and soft emissions into subjects, the resulting JFNs are relatively insensitive to the detailed modeling of the infrared (IR) physics that is often poorly understood. This raises the possibility to use JFNs to trade performance for generalizability by adjusting the number of reconstructed subjects. Lastly, we expect that subjects can be measured well in heavy-ion collisions despite the large fluctuating background. See ref. [59] for recent measurements of the energy spectrum of inclusive and leading subjects by the ALICE Collaboration. See also ref. [60].

The remainder of this paper is organized as follows. In section 2, we introduce the subjet basis and discuss differences between inclusive and exclusive subjet reconstruction algorithms. In section 3, we introduce the permutation invariant machine learning algorithms that take the kinematic information of the reconstructed subjects as input and in section 4, we briefly discuss the data sets used for different classification tasks used in this work. In section 5, we present numerical results for the classification performance of JFNs for quark vs. gluon and Z vs. QCD jets. In particular, we show that JFNs match the PFN performance for a finite subjet radius and a finite cut on the transverse momenta of the subjects in both test cases. Based on these results, we describe in section 6 that the machine learning algorithm is sensitive to different physical scales, which it can effectively learn. In section 7, we investigate

the tradeoff between performance and generalizability of the JFNs. In section 8, we draw conclusions and present an outlook.

2 The subjet basis

In this section, we describe the reconstruction of subjets that will serve as the input to the machine learning classifier. The initial jet is identified using the anti- k_T algorithm [61] and jet radius parameter R . In order to utilize the substructure of jets, we then recluster the jet constituents into subjets. We consider two approaches for the subjet reconstruction: inclusive anti- k_T subjets and the exclusive k_T subjets [62, 63]. Note that we do not consider exclusive anti- k_T jets as this clustering procedure does not lead to a “physical” clustering tree. In both cases, soft and collinear emissions are first clustered into subjets. For a finite subjet radius and an additional cut on the subjet’s transverse momentum p_T^{soft} , the input to the classifier is IRC safe. In this sense, subjets serve as a useful tool for throttling or controlling the input data to the machine in a way that is theoretically interpretable in perturbative QCD.

First, we consider inclusive subjets reconstructed with the anti- k_T algorithm using the E -recombination scheme [64] and a fixed subjet radius $r < R$. This approach fixes the maximally allowed size of the reconstructed subjets but the number of subjets varies for each jet. We illustrate the distribution of subjets in the η - ϕ plane for three different subjet radii in figure 1. As r is increased, the central subjet contains a large fraction of particles. Second, we consider subjets reconstructed with the exclusive k_T algorithm and the E -recombination scheme. Particles are clustered with the k_T algorithm until a fixed number of subjets N is obtained. Different than in the case of inclusive subjets, the number of identified subjets is fixed but their size varies jet-by-jet. The N subjets span the full information content of the N most resolved emissions inside the jet analogous to the N -subjettiness basis developed in refs. [10, 38, 39]. An alternative approach to the exclusive k_T algorithm is to identify subjets with the XCone algorithm [65]. We leave the exploration of this algorithm for future work. By taking the small- r (inclusive subjets) or large- N limit (exclusive subjets), as well as $p_T^{\text{soft}} \rightarrow 0$, we can study the transition to the nonperturbative regime where eventually, every subjet only contains a single hadron.

To illustrate the qualitative differences between the two reconstruction methods discussed above, we show as an example the longitudinal momentum distributions of subjets $z = p_T^{\text{subjet}}/p_T$ in figure 2 separately for quark and gluon jets. The details on the data generation are presented in section 4. Here p_T denotes the initial jet transverse momentum and p_T^{subjet} the longitudinal subjet momentum using either the inclusive or exclusive reconstruction method. As an example, we choose $N = 30$ for the exclusive reconstruction of subjets and $r = 0.02$ for inclusive subjets, which yields a comparable average number of 30 subjets. We observe that the two methods lead to qualitatively different spectra. The inclusive subjet spectrum exhibits a peak (quarks) or plateau (gluons) for intermediate to large values of z . In contrast, the spectrum for exclusive subjets only peaks at small values of z and falls off steeply for $z \rightarrow 1$. This is due to the fact that for exclusive clustering the k_T algorithm is used, where soft hadrons are clustered first. Only at the end hard emissions are combined, making it unlikely to find a subjet with $z \rightarrow 1$ for a fixed value of N . We note that for inclusive subjets, the z -distributions are qualitatively the same for both the anti- k_T and k_T

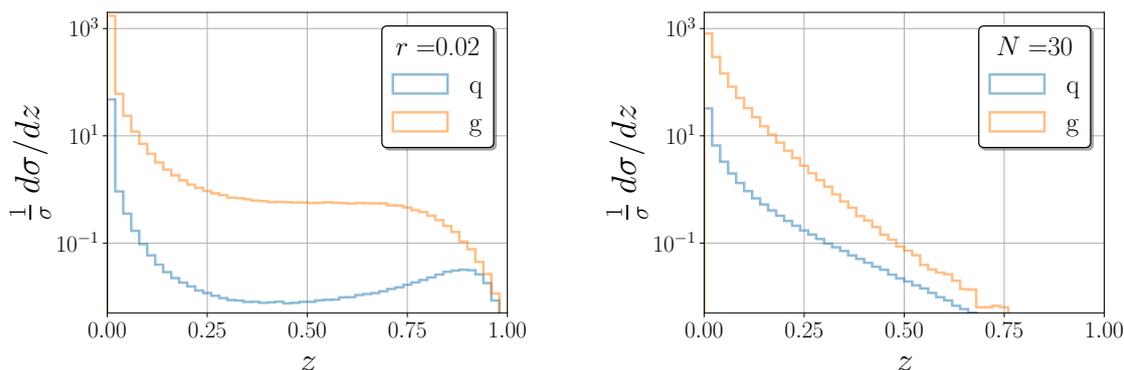


Figure 2. The longitudinal momentum distribution of inclusive subjects $z = p_T^{\text{subject}}/p_T$ originating from either a quark (blue) or a gluon (orange) jet simulated with PYTHIA [66]. The jets have $p_T = [500, 550]$ GeV with an average particle multiplicity of 43. We show the distributions for inclusive anti- k_T subject clustering with $r = 0.02$ (left) corresponding to an average of approximately 30 subjects, and for exclusive k_T clustering with a fixed number of $N = 30$ subjects (right). See also section 4 for more details.

algorithms. The longitudinal momentum spectrum for inclusive subjects was calculated within perturbative QCD up to next-to-leading logarithmic (NLL) accuracy. See refs. [67–70]. This close connection to first-principles calculations may allow for an increased understanding of machine learning algorithms in QCD.

From the identified subjects, the kinematic information (z_i, η_i, ϕ_i) of the each subject is used as input to the classifiers discussed below. In the limit that $r \rightarrow 0$ (inclusive subjects) or $N \rightarrow \infty$ (exclusive subjects), the subject basis becomes equivalent to the set of particle four-vectors of the jet, and the classifier can make use of the full information content of the jet. The subject basis therefore provides a means to limit the information supplied to the classifier, by using $r > 0$ or $N < \infty$.

3 Jet Flow Networks (JFNs): deep sets of subjects

In this section, we describe the permutation invariant neural networks that use the kinematic information of subjects as input to perform binary classification tasks. As introduced above, we refer to the machine learning architecture and the pre-processing step of clustering particles into subjects as JFNs.

The reconstructed subjects discussed in the previous section do not have an inherent ordering. Therefore, permutation-invariant neural networks are a natural choice to perform classification tasks that take as input the kinematic information of subjects. In refs. [50–52] deep sets were introduced as a permutation invariant neural network. In the context of particle physics deep sets were first discussed in ref. [11] as Particle Flow Networks (PFNs) that take as input the information of individual particles. A permutation invariant classifier f , which takes as input the subject four-momenta p_i satisfies $f(p_1, \dots, p_N) = f(p_{\pi(1)}, \dots, p_{\pi(N)})$.

	PFN [11]	JFN	EFN [11]
Input	particle 3-momenta	subset 3-momenta	particle 3-momenta
Classifier	IRC unsafe	Sudakov safe	IRC safe

Table 1. Overview of different classifiers based on permutation invariant neural networks.

Here π denotes the permutation operator. Following ref. [50], we can write the classifier f as

$$f(p_1, \dots, p_n) = F\left(\sum_{i=1}^N \Phi_i(p_i)\right), \tag{3.1}$$

where F , Φ are neural networks and, as an intermediate step, we sum over all reconstructed subsets N . The first neural network $\Phi : \mathbb{R}^4 \rightarrow \mathbb{R}^l$ takes as input the individual subset four momenta and maps it to an l -dimensional latent space. For massless subsets, we can write the individual four vectors in terms of (z_i, η_i, ϕ_i) . Here z_i is the subset's longitudinal momentum fraction, see figure 2, and (η_i, ϕ_i) denote its coordinates in the rapidity-azimuth plane. We note that further information can be included in the per-subset mapping such as the jet mass or the jet charge [71, 72], analogous to e.g. particle identification (PID) for PFNs. We leave quantitative studies of the impact of these additional features for future work. The summation in eq. (3.1) ensures that the classifier f is invariant under permutations of the input variables. The second neural network $F : \mathbb{R}^l \rightarrow \mathbb{R}$ is a map from the latent space where the summation operation is performed to the final classification score. Note that the classifier architecture in eq. (3.1) can accommodate both a fixed number N of subsets (exclusive subsets) and input with variable length (inclusive subsets).

We refer to the deep set classifier based on subsets in eq. (3.1) as JFNs. The JFNs are a family of classifiers due to the dependence on the continuous parameter r in the case of inclusive clustering or on N in the case of exclusive clustering, in which case the clustering is performed until N subsets remain. Since the JFN takes subset information as input, the resulting classifier is generally Sudakov safe [45]. We summarize the different aspects of permutation invariant network architectures based on deep sets in table 1. Since the JFNs are Sudakov safe, they constitute an intermediate point between IRC-unsafe PFNs and IRC-safe EFNs. In the limit of $r \rightarrow 0$ (inclusive subsets) or the large- N limit (exclusive subsets) and $p_T^{\text{soft}} \rightarrow 0$, we recover the PFN classifier.

We note that one can think of the input to the classifier as a multi-differential cross section or multi-variate probability distribution. To avoid for example regions of negative cross sections, the corresponding perturbative calculation requires the all-order resummation of sufficiently many logarithmic corrections. The resummed calculation then needs to be matched to fixed-order calculations. The classifier performs a highly nontrivial marginalization over the multi-variate probability distribution. Since all order resummations are required, the final result of the classifier is Sudakov but not necessarily IRC safe.

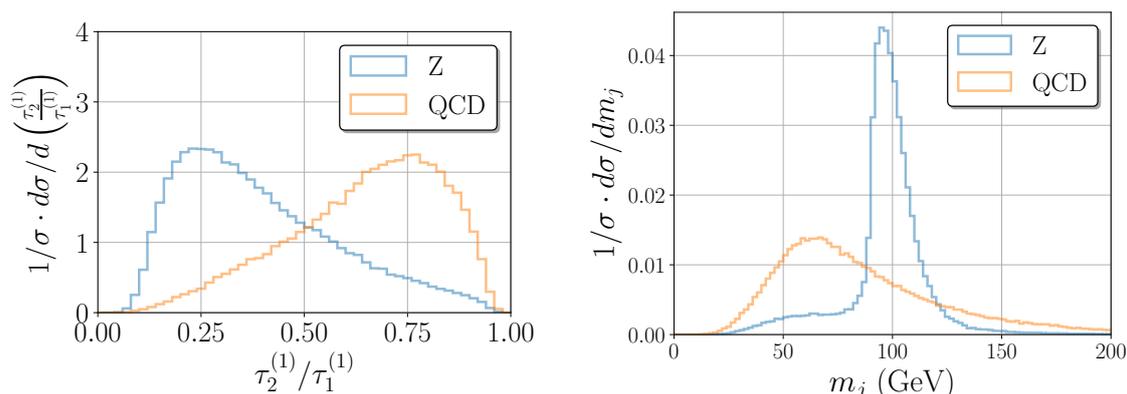


Figure 3. N -subjettiness ratio $\tau_2^{(1)}/\tau_1^{(1)}$ (left) and the jet mass distribution (right) for QCD and Z jets with $p_T = [500, 550]$ GeV. The N -subjettiness axes were identified using the one pass k_T clustering algorithm.

4 Data sets

In this work, we will consider JFNs for two exemplary binary classification tasks in high-energy physics. First, we consider quark vs. gluon jet classification and, second, Z vs. QCD jet classification. For the quark vs. gluon case, we make use of the data set in ref. [73], which consists of 2M jets with transverse momentum $p_T = [500, 550]$ GeV, rapidity $|\eta| < 1.7$, jet radius parameter $R = 0.4$, and center-of-mass energy $\sqrt{s} = 14$ TeV. We will make use of both the data set generated with PYTHIA [66] and HERWIG [74]. In order to explore the dependence on the jet transverse momentum, we also generate two additional data sets consisting of 500k jets each with transverse momentum $p_T = [300, 350]$ GeV and $[1000, 1050]$ GeV, respectively. The underlying processes are: $q\bar{q} \rightarrow Z(\rightarrow \nu\bar{\nu}) + g$ and $q\bar{q} \rightarrow Z(\rightarrow \nu\bar{\nu}) + (uds)$ analogous to ref. [73]. For the Z vs. QCD-jet case, we generate 500k jets for three different bins of jet transverse momentum, $[300, 350]$ GeV, $[500, 550]$ GeV and $[1000, 1100]$ GeV with a jet mass $m_j = [45, 135]$ GeV. The radius parameter is $R = 0.8$, the rapidity cut is $|\eta| < 1.7$ and the samples are generated using PYTHIA at $\sqrt{s} = 14$ TeV. Jets arising from Z bosons are identified by requiring that the leading Z boson is in the catchment area of the jet as extracted from the kinematics of the events at the particle level before hadronization with a Z -jet distance from the jet axis less than $R/2$. A similar tagging procedure is performed to differentiate between quark and gluon jets in the QCD sample. The tag is based on the leading parton within the catchment area of the jet. However, to strengthen the parton-jet association, we use parton-level kinematics injected into the hadron-level event using so-called ghost particles ($p_T = 10^{-5}$ GeV) that do not affect the jet reconstruction but allow for efficient tagging after the jet finding step.

The substructure of QCD jets is generally single-pronged, whereas the decay products of a Z cause the corresponding jets to have two prongs. The ratio of N -subjettiness observables [40, 41, 75, 76] is sensitive to the number of prongs inside a jet. In order to define the N -subjettiness, a given number of N axes are identified inside the jet using the exclusive k_T algorithm. The N -subjettiness variables $\tau_N^{(\beta)}$ measure the radiation along these

axes and are defined as

$$\tau_N^{(\beta)} = \frac{1}{p_T} \sum_{i \in \text{jet}} p_{Ti} \min \{ R_{1i}^\beta, R_{2i}^\beta, \dots, R_{Ni}^\beta \}. \quad (4.1)$$

Here the p_{Ti} of each particle i is weighted by its distance R_{ji} to the closest axis j raised to the power $\beta > 0$, which is a tunable parameter. For jets that are more like a single-prong jet, the variable τ_2 will peak at smaller values compared to τ_1 , whereas for two-prong like jets the variable τ_2 takes similar values compared to τ_1 (by construction, $\tau_{n+1} \leq \tau_n$). To illustrate the qualitative differences between QCD jets and Z jets, we plot in figure 3, the result for the ratio $\tau_2^{(1)}/\tau_1^{(1)}$, which shows the expected separation of the two jet samples (left panel), and the jet mass m_j distribution for QCD and Z jets (right panel). For all classification tasks, the training/validation/test split is 80%/10%/10%.

5 JFN performance: gapless jet classification

In this section, we will present results for the performance of the JFN classifier. In the two subsections, we separately study the JFN dependence on the subjet radius r and the transverse momentum cutoff p_T^{soft} , which are associated with collinear and soft safety, respectively. We will observe that the maximum performance is obtained for finite values r and p_T^{soft} .

5.1 Collinear safety

We start by studying JFNs for different values of the subjet radius r . In the limit $r \rightarrow 0$, every subjet contains only a single hadron, and the PFN performance is recovered. Throughout this section, we consider $p_T^{\text{soft}} = 0$ GeV. In the next subsection, we consider the JFN performance for finite values of the transverse momentum cutoff. We consider two exemplary binary classification tasks in high-energy physics: quark vs. gluon and Z vs. QCD jet identification.

In order to implement the permutation-invariant neural networks, we parametrize the functions Φ and F in eq. (3.1) in terms of DNNs, using the **EnergyFlow** package [11] with **Keras** [77]/**TensorFlow** [78]. 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 [79] and we use the softmax activation function for the final output layer of the classifier. We train the neural networks using the Adam optimizer [80] with learning rates ranging from 10^{-3} to 10^{-4} . We use the binary cross entropy loss function [81] and train for 60 epochs with a batch size of 256 and a patience parameter of 8 for early stopping. We find no significant changes in performance when changing the size or number of the layers, latent space dimension, learning rate, and batch size by factors of 2-5. Following ref. [11], we perform a preprocessing step to simplify the training process: we use the rescaled momentum fractions z_i and center the rapidity and azimuthal angles η_i, ϕ_i of the particles in the jet with respect to the jet direction.

We quantify the performance of the different classifiers in terms of the Receiver Operator Characteristic (ROC) curve and the Area Under the Curve (AUC) of ROC curve. The ROC curve is the cumulative distribution function of the true positive rate vs. the false positive rate of a classifier as the decision threshold is varied. The AUC takes values between 0.5 and 1, where 0.5 (1) corresponds to a random (perfect) binary classifier. We estimate the

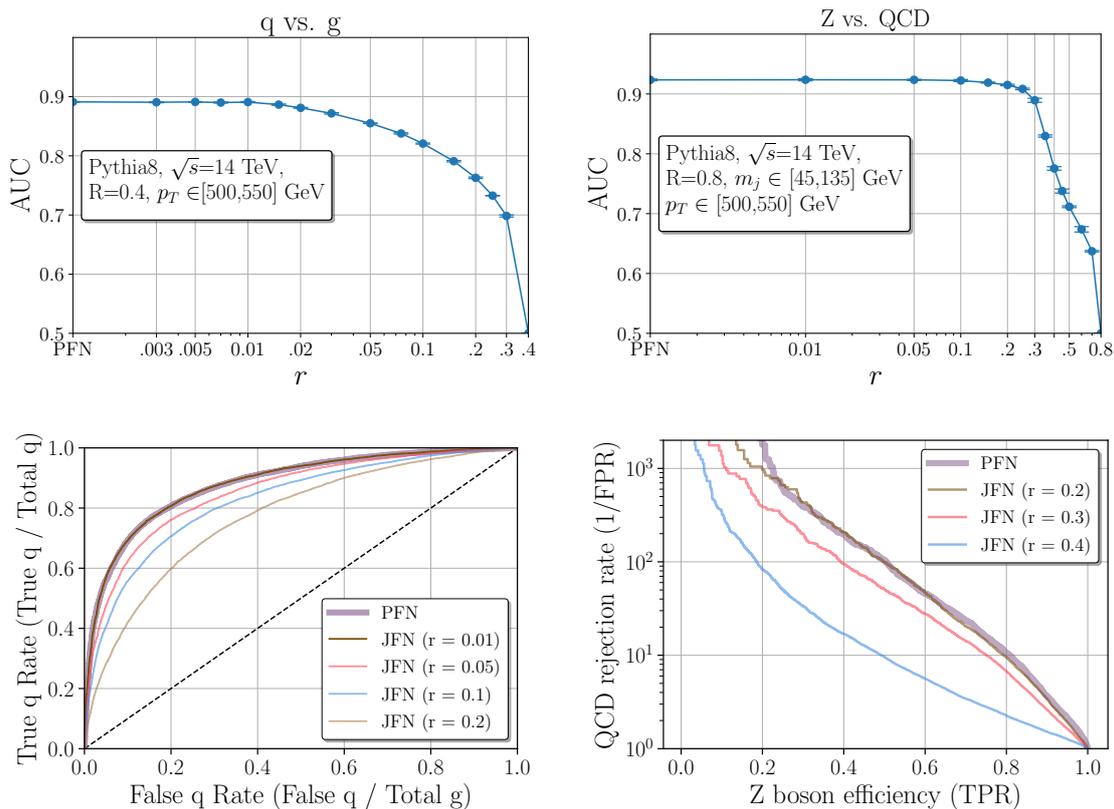


Figure 4. Top panel: AUC for quark vs. gluon (left) and Z vs. QCD (right) jet tagging using JFNs with different values of the (inclusive) subjet radius r . The PFN classifier is shown for reference at the leftmost value of r . Bottom panel: ROC curves for quark vs. gluon (left) and Z vs. QCD jet tagging using JFNs and the PFN with different values of the (inclusive) subjet radius r for the same datasets as the upper panel.

statistical uncertainty of the AUC by training the deep sets four times for each choice of the subjet radius r and using the standard deviation as the uncertainty.

Figure 4 (left) shows the JFN results for q vs. g jet discrimination using inclusive subjet clustering and 2M jets for training, validation and testing. The top panel shows the AUC performance as we change the inclusive subjet radius r , and the bottom panel shows the ROC curves for several r . For comparison, we also show the PFN result. In the case of the AUC plot, we display the PFN classifier as the leftmost point on the r axis. As expected, for $r = R$ we obtain a random classifier since all particles are clustered into a single subjet that is identical to the original jet. Similarly, we find that for the smallest subjet radii r the performance of the PFN is recovered. Strikingly, however, the performance of the JFN does not diminish as r is increased for values of the subjet radius $r \lesssim 0.01$. At this critical r value, we have on average $n_{\text{subjets}}/n_{\text{hadrons}} \approx 0.75$. This observation is corroborated by the ROC curve in the lower panel, which shows that there is no performance loss in the JFN ($r = 0.01$) as compared to the PFN. This demonstrates that there is little-to-no information encoded in the very collinear emissions relevant for discriminating q vs. g jets, and suggests

Model	AUC q vs. g	Model	AUC Z vs. QCD
PFN	0.8912 ± 0.0005	PFN	0.9235 ± 0.0015
JFN ($r = 0.005$)	0.8911 ± 0.0002	JFN ($r = 0.01$)	0.9237 ± 0.0018
JFN ($r = 0.01$)	0.8904 ± 0.0009	JFN ($r = 0.05$)	0.9236 ± 0.0010
JFN ($r = 0.015$)	0.8865 ± 0.0011	JFN ($r = 0.1$)	0.9227 ± 0.0017
JFN ($r = 0.02$)	0.8812 ± 0.0008	JFN ($r = 0.15$)	0.9189 ± 0.0009
JFN ($r = 0.05$)	0.8550 ± 0.0004	JFN ($r = 0.2$)	0.9150 ± 0.0012
JFN ($r = 0.1$)	0.8207 ± 0.0009	JFN ($r = 0.3$)	0.7755 ± 0.0025
JFN ($r = 0.2$)	0.7629 ± 0.0013	JFN ($r = 0.4$)	0.7115 ± 0.0015

Table 2. Numerical values for the AUC for both classification tasks considered in this section.

that collinear safe inputs are sufficient for the purpose of q vs. g classification. In section 6, we will further discuss the physical interpretation of this critical r value.

Figure 4 (right) shows the analogous JFN results for Z vs. QCD jet classification using inclusive subjet clustering and 500k jets for training, validation and testing. We observe again that the JFNs smoothly converge to the result of the PFN. Different than for quark vs. gluon jet tagging, we can now choose a significantly larger subjet radius $r \lesssim 0.1$ without compromising the performance of the classifier. This is related to the fact that in this case, the boosted Z -boson decay products generally lead to a two-pronged jet substructure, whereas QCD (quark and gluon) jets exhibit a single-pronged jet substructure (see figure 3 as well as refs. [1, 3] for different observables and perturbative calculations that characterize the radiation patterns of QCD and boosted Z jets). In general, machine-learned classifiers can make use of more information than the one- vs. two-pronged structure inside these jets; for $r \sim 0.1$, a significant fraction of the hadrons inside the jets are clustered into subjets, $n_{\text{subjets}}/n_{\text{hadrons}} \approx 0.4$. However, due to the observed saturation up to $r \sim 0.1$, we conclude that the information contained in collinear emissions is significantly less relevant for this classification task compared to quark vs. gluon jet tagging. This is due to the physical scales that are relevant for the different jet classification tasks, which we will explore in more detail in section 6. For completeness, we list the numerical values for the AUC including uncertainties in table 2.

5.2 Soft safety

Next, we are going to study the sensitivity of the classifier to soft emissions by analyzing the JFN performance as a function of the cut on the subjet's transverse momenta p_T^{soft} . The results for the AUC are shown in figure 5 for both quark vs. gluon and Z vs. QCD jets. We studied the effects of the soft cut for two ranges of transverse momenta for the jet and using 500k jets as the dataset for each classification tasks. Again, we observe that the PFN performance is achieved for finite values of p_T^{soft} that are $\mathcal{O}(1 \text{ GeV})$. We observe only small differences for the two classification tasks. The critical value for p_T^{soft} where the performance reaches a plateau is independent of the total jet transverse momentum p_T within the displayed errors. This indicates that there is a soft scale below which no further information is added

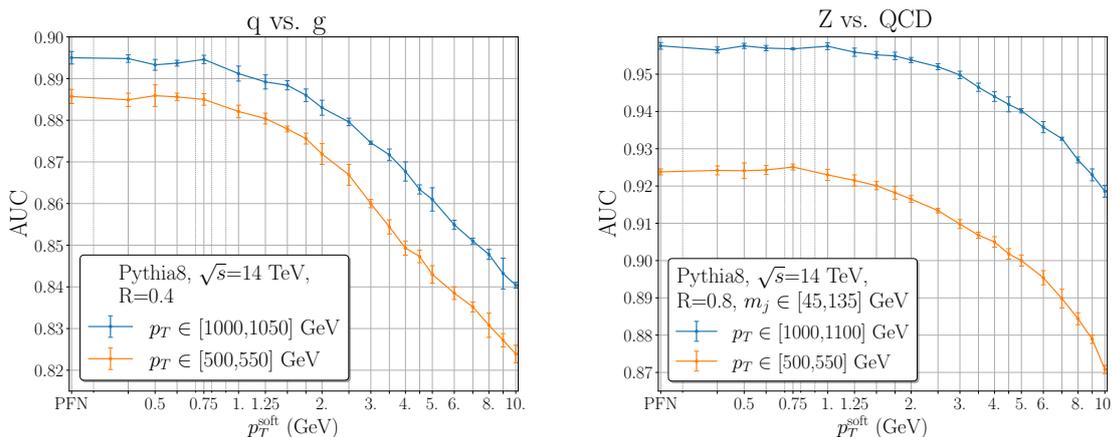


Figure 5. AUC of the JFN as a function of the cut on the subjet’s transverse momenta p_T^{soft} for quark vs. gluon (left) and Z vs. QCD (right) jet tagging.

that can improve the classification performance. We conclude that IRC-safe information is indeed sufficient to carry out the two jet classification tasks considered in this work, which is consistent with theoretical considerations in the literature [29, 82]. However, the relatively low values of p_T^{soft} that are needed to achieve the maximum performance indicate that nonperturbative effects are generally relevant. These observations may inform the design of suitable observables, which we leave for future work.

For both jet classification tasks, we have shown that for a range of subjet radii r and p_T^{soft} cuts, the JFN exhibits no significant difference in performance compared to the PFN. That is, the JFN classifier here is “gapless” in the sense that we smoothly approximate the PFN performance (for finite values of r and p_T^{soft}). The clustering of soft and collinear emissions into subjets does not affect the performance as long as r and p_T^{soft} are sufficiently small. This is in contrast to previous studies based on observables such as N -subjettiness variables, which exhibit a small but persistent performance gap to PFNs [7, 10, 11]. The JFN provides the first example of a classifier with IRC-safe inputs that achieves equivalent performance to the IRC-unsafe PFNs for several classification tasks. Our results are consistent with the intuitive expectation that very low-energy particles are essentially uncorrelated with the hard process and therefore do not provide relevant information for typical classification tasks in high-energy physics. The main question of our paper has thus been answered by these observations. At least for the two classification tasks considered here, we have found that IRC-safe information is sufficient to close the gap to IRC-unsafe classifiers. This was achieved by using the machine learning architecture and input type (momentum, position information) for both cases and by including subjet reclustering as a preprocessing step in the IRC-safe case. We note that our conclusions come with the following caveat. While we are going to identify relevant physical scales with the performance of the classifiers, it is possible that future advances in machine learning lead to more powerful algorithms that may require us to reduce the subjet radius r and cutoff p_T^{soft} to match the performance of IRC-unsafe classifiers.

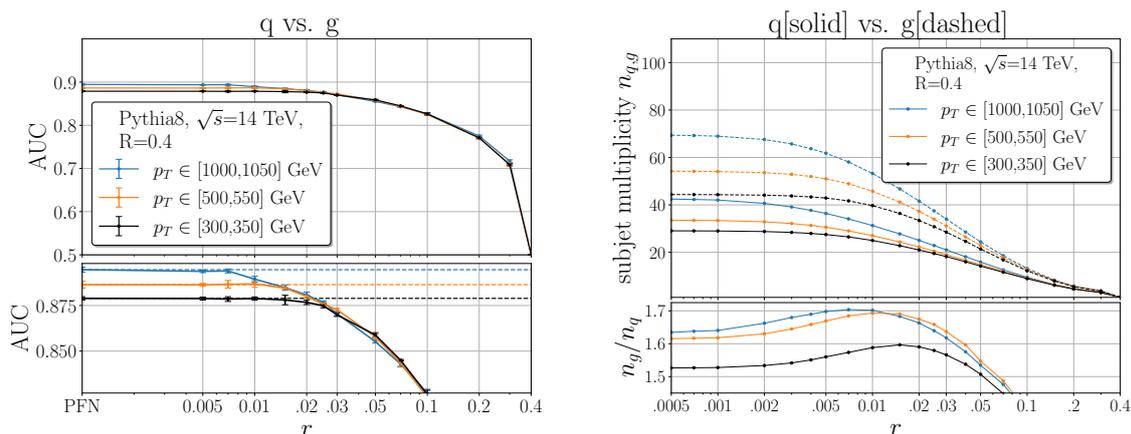


Figure 6. Left: AUC for quark vs. gluon jet classification for three different jet p_T intervals as a function of the subjet radius r . Upper right: average subjet multiplicity $n_{q,g}$ for quark (solid) and gluon (dashed) jets. Lower right: ratio of the average subjet multiplicities n_g/n_q for the three jet p_T intervals.

6 Learning physical scales

As discussed in the previous section, the performance of the JFNs based on IRC-safe input matches that of the IRC-unsafe PFNs for finite values of the subjet radius r where a significant fraction of hadrons is clustered into subjets. In this section, we quantify in more detail the onset of the drop in performance when the subjet radius crosses certain physical scales. Here we only focus on inclusive instead of exclusive subjet reconstructions since we are primarily interested in the physical scale associated with a fixed value of subjet radius r . In order to identify the physical scale associated with the classification tasks and study its scaling behavior, we are going to analyze the AUC for different bins of jet transverse momentum p_T . Throughout this section, we only focus on the subjet radius r and choose $p_T^{\text{soft}} = 0$ GeV since we already identified the relevant soft scale in the previous section.

In the upper left panel of figure 6, we show the AUC for q vs. g discrimination for three different bins of the jet transverse momentum. For comparison, we show analogous to figure 4 the PFN result as the left-most point ($r = 0.001$). For all three classification tasks we used 500k jets for training, validation and testing. First, we notice that at sufficiently large r all three AUC curves merge and the classification performance is (approximately) independent of the jet transverse momentum. This can be traced back to the approximate scale invariance of the QCD parton shower cascade. Second, we observe that in all three cases, the AUC reaches a plateau for finite values of the jet radius as r is decreased. In the lower left panel of figure 6, we show the results in the transition region in more detail. The onset of the plateau shifts to the left as the jet p_T is increased. For higher jet p_T , the jet constituents are more collimated leading to a smaller critical value r where the JFNs match the PFN performance. We can identify the following approximate scale where the AUC reaches a plateau and agrees with the PFN results:

$$p_T \cdot r \sim 5 \text{ GeV}. \quad (6.1)$$

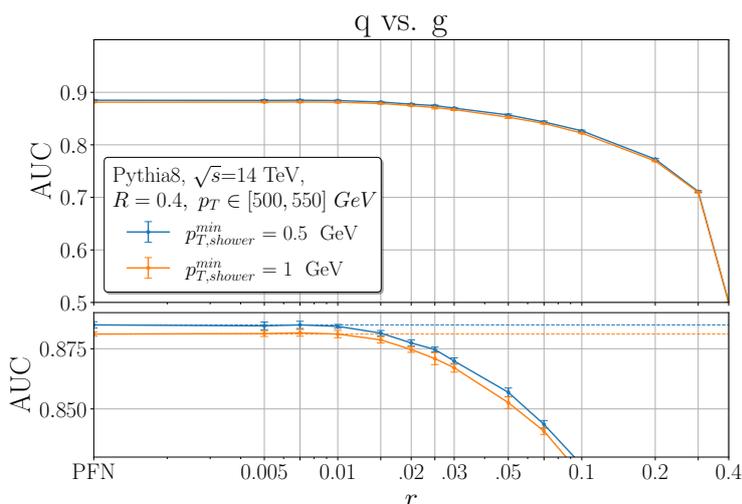


Figure 7. AUC for quark vs. gluon jet classification for two different choices of the perturbative shower cutoff parameter $p_{T,shower}^{\min}$ in PYTHIA as a function of r . The value of r where the maximum classification performance is achieved is independent of the shower cutoff within the displayed errors.

Since there is no additional physical scale in the quark vs. gluon jet classification task besides the jet p_T , the agreement with the PFN result is achieved for relatively low energy scales. However, we would like to stress that the identified scale is still in the perturbative regime. This suggests that because we only use particle momentum as our classifier’s input that non-perturbative symmetries, such as isospin, forbid useful information at hadron level for discrimination. This would be broken and discrimination could improve if information sensitive to flavor was measured, which has been established in some recent studies of the jet charge [48, 49].

Since the (IRC-unsafe) particle multiplicity is known to be a powerful discriminant for quark vs. gluon jet tagging [29, 83, 84], we study the relation of our results to the average subjet multiplicities $n_{q,g}$ as a function of the subjet radius r , which is shown in the upper right panel of figure 6. As expected, the subjet multiplicities for both quark and gluon jets increase as the subjet radius r is decreased. In the limit $r \rightarrow 0$, the subjet multiplicity smoothly asymptotes to the particle multiplicity. The expected value for the ratio of the particle multiplicities at leading order is $n_g/n_q = C_A/C_F = 9/4$ [85, 86]. In agreement with the discussion in the previous section, we notice that the subjet radius where the quark vs. gluon jet AUC reaches a plateau is larger than the r value at which the subjet multiplicity reaches the particle multiplicity by an order of magnitude. This confirms that matching the PFN performance with JFNs is a non-trivial result. As shown in the lower right panel of figure 6, we observe that the ratio n_g/n_q peaks at intermediate values of r , which is in the region of the better modeled perturbative physics [84, 87–89]. Interestingly, the location of the peaks is approximately the same as where the AUC for quark vs. gluon jet tagging reaches the plateau and agrees with the PFN result. This interesting correlation indicates that we can increase the subjet radius r without affecting the classification performance until the subjet multiplicity n_g/n_q starts to decrease.

Next, we investigate the dependence of our results on model parameters. In particular, we focus on the cutoff scale of the shower in PYTHIA at which the perturbative evolution ends and partons get converted to hadrons using the Lund string fragmentation model. The default value of this parameter is $p_{T\text{shower}}^{\text{min}} = 0.5 \text{ GeV}$. To assess the dependence on this parameter, we compare the default choice to $p_{T\text{shower}}^{\text{min}} = 1 \text{ GeV}$, which changes the transition from perturbative to non-perturbative splittings. The AUC results for the resulting classifiers are shown in figure 7 for quark vs. gluon jets. We find that the asymptotic value of the AUC is slightly different when a higher transition scale is used. This is expected since the hadronization module used in PYTHIA masks differences between quark and gluon jets. More importantly, we observe that indeed the plateau is reached for the same finite value of $r \sim 0.01$, i.e. the scale we identified is independent of the choice of this parameter in PYTHIA. This illustrates that while the detailed modeling of soft physics is needed to describe the asymptotic AUC of the classifier, we observe that its peak performance is determined by a finite value of the subjet radius.

Next, we consider QCD vs Z jet classification. The AUC for three jet transverse momentum intervals is shown in figure 8 as a function of the subjet radius r . In all three cases we used 500k jets for training, validation and testing. As the jet p_T is increased, the value of the subjet radius $r \sim 0.1 - 0.2$, where JFNs match the PFN performance is shifted to the left. This observation is generally consistent with quark vs. gluon jet tagging discussed above. We note that for the AUC curve with $p_T \in [300, 350] \text{ GeV}$ the choice of the jet radius R might start to play a role. The value of r where the performance reaches the plateau is roughly a factor of 10 higher compared to quark vs. gluon jet tagging. As already hinted at above, this is due to the presence of different physical scales. QCD jets do not have any additional intrinsic scales except for the hadronization scale. Instead, jets that contain the decay products of the boosted Z boson are sensitive to the Z -boson mass M_Z .

In order to gain further insights into the underlying physics, we are going to study the distribution of the opening angle of the two leading subjets θ_{12} . This variable is closely related to the 2-pronged structure of Z jets and serves as a useful discriminant since at leading order the Z -boson decays into a quark and anti-quark, which correspond to the two leading subjets at this order. The boosted Z decay products have an opening angle θ_Z , which is determined by M_Z and the jet transverse momentum p_T as

$$\theta_Z \sim \frac{2M_Z}{p_T}. \tag{6.2}$$

For higher p_T , the decay products are more boosted and θ_Z is smaller. See figure 9 for an illustration of the boosted Z -boson decay products clustered into subjets. If the subjet radius parameter is sufficiently small $r < \theta_Z/2$, the Z -boson decay products are clustered into separate subjets. Instead, for $r > \theta_Z$, they are merged into a single subjet. In the intermediate region, $r < \theta_Z < 2r$, they are identified as two separate subjets but the subjet catchment areas overlap.

In figure 10, we show the distributions of the opening angle θ_{12} between the first two leading subjets for both QCD and Z jets for different values of the subjet radius r . Here, θ_{12} corresponds to the geometric distance in the η - ϕ plane, i.e. without rescaling by the jet radius R . The left column of figure 10 shows the opening angles θ_{12} between the two leading hadrons,

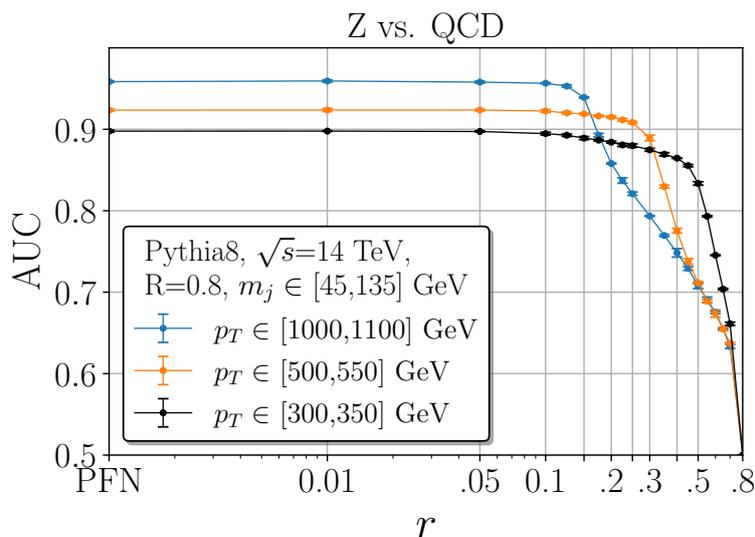


Figure 8. AUC for Z vs. QCD jets for three different jet p_T intervals as a function of the subjet radius r .

which corresponds to the subjet radius $r = 0$. The middle column shows the θ_{12} distributions for r values in the plateau region where the AUC of the JFNs matches the PFN result. The right column corresponds to higher r values where the AUC has dropped significantly, see figure 8 above. The top and bottom row correspond to two different jet p_T intervals as indicated in the figure. We observe that for both QCD and Z jets, the distributions are bounded from below by the chosen subjet radius, $\theta_{12} > r$, and the distributions vanish when the angle between the leading subjets reaches the jet radius $\theta_{12} \lesssim R = 0.8$. Due to collinear QCD emissions, the distribution of the angle between the two leading hadrons peaks at $\theta_{12} \sim 0$. As the subjet radius is increased, it peaks close to the lower bound $\theta_{12} \sim r$. Eventually, the θ_{12} distribution becomes broader for large values of r . Instead, both at hadron level and for r values in the plateau region of the AUC, the θ_{12} distribution of Z -jets has a two-peak structure. The left peak is due to QCD emissions and it occurs at the same θ_{12} value as the single peak of QCD jets. The second peak occurs around the opening angle of the Z decay products $\theta_{12} \sim \theta_Z$. The width of the peaks scales as $\sim 1/p_T$. When r is chosen in the plateau region of the AUC in figure 8, the JFN performance agrees with the PFN result. In this region, the two-peak structure of the Z -jet θ_{12} distribution can be clearly identified. The two-prong structure of Z jets is the most prominent feature that distinguishes the two jet samples and it is clearly resolved as long as r is sufficiently small. In this region, we found that the JFN performance is the same as the IRC-unsafe PFN result. While the location of the $\theta_{12} \sim \theta_Z$ peak is fixed, the peak due to QCD emissions moves to larger θ_{12} as r is increased. Eventually, the two peaks start to merge. This is illustrated in the right-most column of figure 10, which shows the θ_{12} distributions for r values where the JFN performance is significantly below its maximal value. The two peaks of the Z jets have merged into one peak and the distribution is very similar to QCD jets. In this case, the Z -decay products

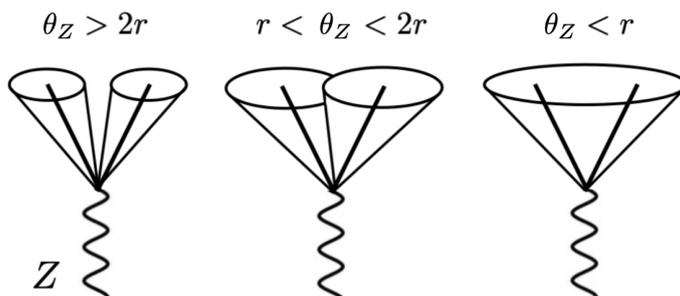


Figure 9. Reclustering of the Z -boson decay products into subjects with different radii.

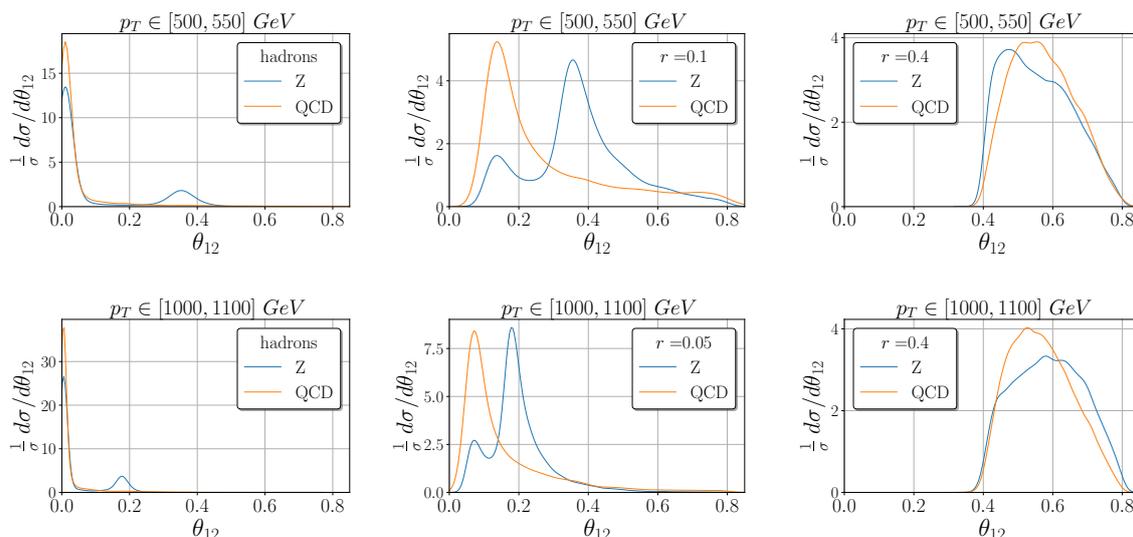


Figure 10. Distributions of the opening angle θ_{12} between the two leading subjects for both QCD and Z jets. We show the results for the two leading hadrons ($r = 0$, left column) and two representative r values (middle and right column). The upper and lower row correspond to two intervals of the jet transverse momentum p_T .

cannot be clearly resolved and the performance of the classifier deteriorates. At this scale, the classifier does not have access to the UV physics anymore and as such the performance for the $p_T = [1000, 1100]$ GeV jets matches the performance for the $p_T = [500, 550]$ GeV jets. By comparing the upper and lower row of figure 10, we observe that the location of the peaks is shifted to lower values for higher jet p_T . In addition, the width of the peaks is narrower $\sim 1/p_T$. This agrees with the observation that for higher jet p_T , the end of the AUC plateau in figure 8 is reached for smaller r values.

Another way of illustrating the importance of resolving the Z -boson decay products is by training deep sets using only the information of the first few leading subjects. In figure 11, we show the AUC for QCD vs Z jets as a function of the subject radius r for three classifiers. We compare the JFNs to deep sets trained on the kinematic information of only the first two or three leading subjects. As an example, we use the jet transverse momentum interval of $p_T = [500, 550]$ GeV. We observe that for large and intermediate values of the subject radius r ,

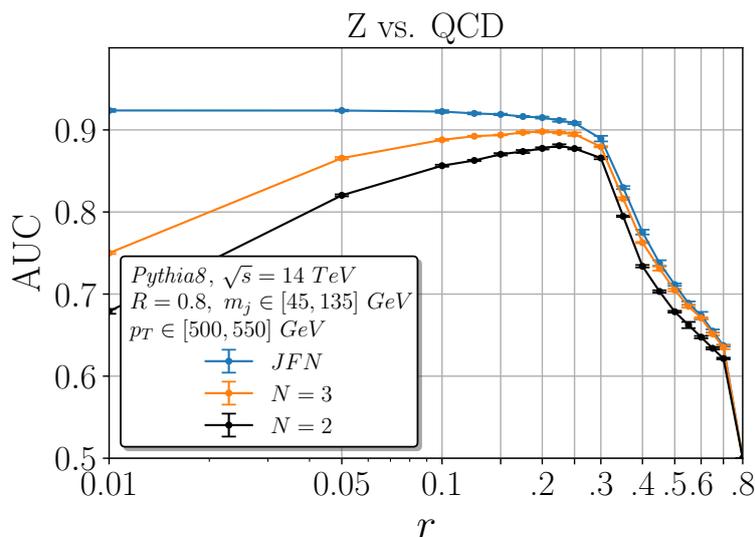


Figure 11. AUC of the JFNs for Z vs. QCD jets trained on the full information (inclusive subjets) compared to deep sets trained only on the two or three leading subjets.

the JFN performance is close to the deep sets trained on only two or three leading subjets. For small values of r , the leading two or three subjets do not contain enough information to match the JFN result. Especially, using the information of three leading subjets closely approximates the JFN performance down to a subjet radius of $r \sim 0.2$. The relevance of the third leading subjet corroborates the results of ref. [38], where the leading emission off the color dipole was identified as an important component for Z vs. QCD jet classification.

Analogous classification tasks where physical scales can likely be identified are light QCD vs. c or b -jets [90–92], QCD vs. Higgs [93] or QCD vs. top quark jets [14]. We leave the exploration of these topics for future work.

7 Performance vs. generalizability

Machine learning-based classifiers are often deployed in experimental analyses to tag jet topologies. A typical method is to train the classifier using fully supervised learning on precise theoretical simulations and apply it to experimental data [92, 94, 95]. However, this approach introduces model dependence as simulations do not perfectly match the actual data. In this section, we will explore some of the systematic uncertainties associated with this method. Other options that have been proposed include semi- or weakly-supervised techniques [96, 97], as well as data-driven methods [98].

When using fully supervised learning to develop classifiers, it is crucial to ensure that the model can generalize well to the unseen experimental data. For JFNs, soft and collinear particles are clustered into subjets making them less sensitive to the modeling of IR physics. Since it is generally challenging to model the very soft physics (both perturbative and nonperturbative) of collider events in universal Monte Carlo event generators, JFNs may have an advantage compared to PFNs in terms of generalizability. On the other hand, if too

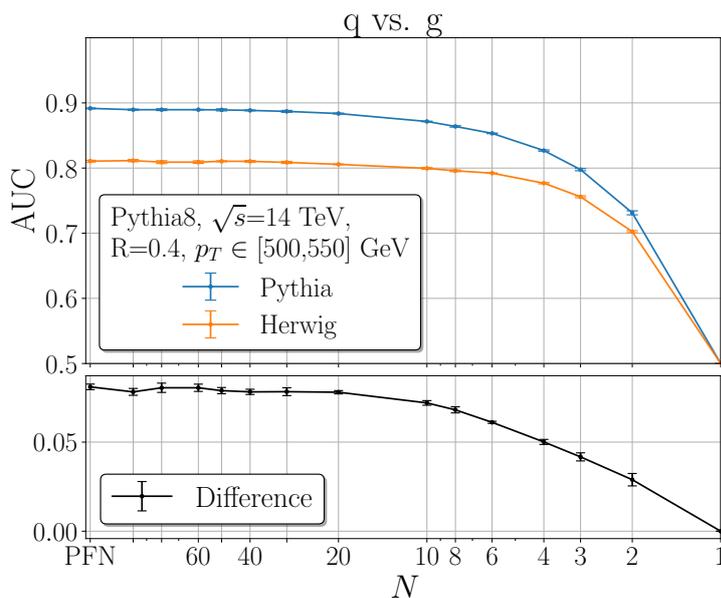


Figure 12. Classification performance for quark vs. gluon jets using JFNs and exclusive k_T clustered subjects plotted as a function of the number of subjects N . Upper panel: JFNs trained and tested on PYTHIA [66] (blue), JFNs trained on PYTHIA and tested on HERWIG [74] (orange). Lower panel: the difference in the performance of the two results.

many particles are clustered into few subjects, the overall performance deteriorates. In order to assess whether a classifier performs well on unseen data, we train PFNs and JFNs with different parameters on PYTHIA [66] (training + validation data set) and test on HERWIG [74] simulations. Here, HERWIG can be considered as a surrogate for experimental data. It has been often observed that Pythia and Herwig envelope many jet substructure observables, sometimes referred to as the “Pythia-Herwig sandwich”. See e.g. ref. [99]. We note that while the final results of both event generators are quite similar, the underlying physics of both the perturbative parton shower and the hadronization model can differ significantly. One generally expects that quark jets are quite similar in PYTHIA and HERWIG but the results for gluon jets tend to differ more significantly [100–102]. See also ref. [6], where PYTHIA and HERWIG studies were presented using Convolutional Neural Networks (CNNs). Moreover, in ref. [103] mixed HERWIG/PYTHIA samples were used together with a Bayesian Network in order to increase model robustness.

We consider quark vs. gluon jet tagging for $p_T = [500, 550]$ GeV using exclusive k_T clustering of the subjects that are taken as input to the machine learning algorithm. Figure 12 shows the AUC as a function of the number of the subjects N . Analogous to the previous figures we show the PFN result as the left-most point ($N = 140$). The upper panel shows the result for JFNs as a function of N trained on PYTHIA and tested on PYTHIA (blue) or HERWIG (orange). In both cases, we observe a plateau in classifier performance as the number of (exclusive) subjects is increased. Within the shown errors, we observe that the AUC in both cases reaches its maximum value for $N \sim 30$. As expected, there is a performance gap

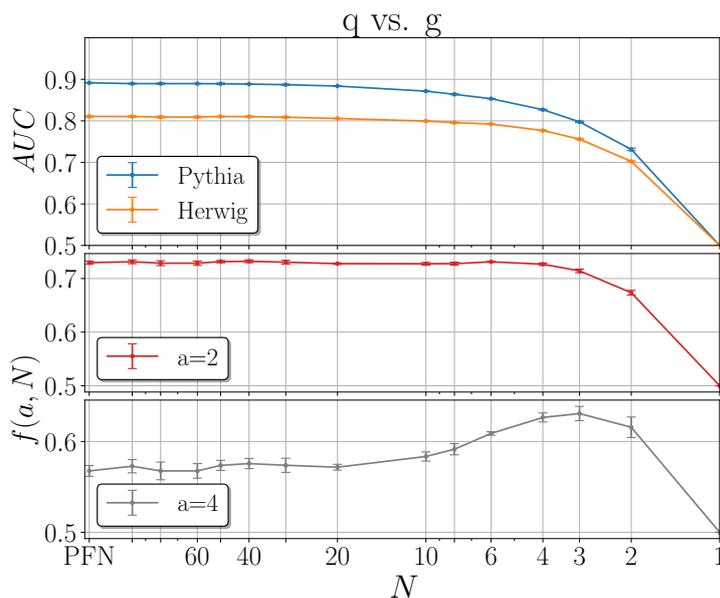


Figure 13. Upper panel: AUC of the JFNs trained on PYTHIA and tested on either PYTHIA or HERWIG plotted as a function of the number of (exclusive) subjects N , see also figure 12. Middle and lower panels: the objective function $f(a, N)$ defined in eq. (7.1), where a is a weighting factor between optimal performance and generalizability.

when testing the PYTHIA-trained classifier on quark vs. gluon jets generated with HERWIG compared to testing it on PYTHIA simulations. This observation is consistent with the results of refs. [6, 103]. However, we observe that the performance gap decreases as N decreases. To better visualize this aspect, we show in the lower panel the difference between the two AUC curves shown in the upper panel. The difference becomes smaller as N is increased indicating improved generalizability of the model. Our findings suggest that clustering particles into subjects can reduce the overall performance, but it also masks modeling uncertainties of the IR physics leading to more robust classifiers. Interestingly, we find that the difference between PYTHIA and HERWIG does not decrease for small N for Z vs. QCD jet classification (not shown). We do not observe an increased robustness using inclusive subject clustering.

We expect that the generalizability or robustness of machine learning-based classifiers will be useful for certain experimental applications where the trade-off between performance and generalizability needs to be considered. To illustrate this aspect, we introduce the objective function $f(a, N)$, defined as:

$$f(a, N) = \text{AUC}_{\text{Pythia}}(N) - a \cdot (\text{AUC}_{\text{Pythia}}(N) - \text{AUC}_{\text{Herwig}}(N)) , \quad (7.1)$$

where N is the number of exclusive subjects. Here the performance and generalizability are combined additively and a weighting factor $a > 0$ is introduced that allows us to increase/decrease the relevance of the two metrics. An optimization problem to find the optimal balance between performance (first term in eq. (7.1)) and generalizability (second term $\sim a$ in eq. (7.1)) can now be formulated as follows: for a given choice of the weighting

factor a , find the maximal value of the objective function $f(a, N)$. The optimal number of exclusive subjets is then given by $N_{\text{opt}} = \arg \max_N f(a, N)$. We plot $f(a, N)$ for two different values of a in figure 13 (middle and lower panels). We observe that as a is increased (the generalizability is weighted higher), the objective function peaks at an intermediate value of N . For example, for $a = 4$ we find $N_{\text{opt}} = 3$. While our objective function is constructed for illustration purposes, this result indicates that for certain experimental analyses that employ machine learning-based classifiers, it can be advantageous to use JFNs with a finite number of subjets to achieve the desired goals.

8 Conclusions

The classification of jets at collider experiments is relevant for a wide range of tasks in high-energy particle and nuclear physics. Over the past years, machine learning-based classifiers have been developed that can achieve impressive tagging performance. While machine learning generally outperforms traditional methods by efficiently making use of the full information content, it is often unclear where the performance difference is coming from. In particular, it had been unclear if classifiers based on infrared-collinear (IRC) information can match the performance of IRC-unsafe classifiers. IRC safety is primarily motivated by theoretical considerations ensuring that observables are tractable in perturbative QCD. In addition, it is expected that the very soft physics is uncorrelated to the hard partonic process making it unlikely to be the reason of the performance gap that has been observed between IRC-unsafe machine learning results and traditional IRC-safe observables.

In order to address these questions, we introduced in this work a new family of classifiers, the Jet Flow Networks (JFNs). Here, particles inside a jet are first clustered into subjets and their position and momentum are taken as input to a permutation-invariant neural network (deep set). The clustering of subjets with a certain radius and transverse momentum cut allows us to control the sensitivity to soft and collinear emissions making the input to the classifier IRC safe. As the subjet radius (collinear safety) and momentum cut (soft safety) vanish, we recover the IRC-unsafe Particle Flow Networks (PFNs). We investigated both inclusive and exclusive subjet clustering, which can lead to important differences depending on the application. As representative examples, we considered two classification tasks: quark vs. gluon and Z vs. QCD jet tagging. Interestingly, we observed that the JFN performance matches the IRC-unsafe PFN result for finite values of the subjet radius and the soft transverse momentum cut. This makes JFNs the first classifier based on IRC-safe input without a performance gap to their IRC-unsafe counterpart for several jet classification tasks. This observation answered the main question we aimed to address in this work and indeed IRC-safe information is sufficient for the jet classification tasks considered here. As the subjet radius is increased, the performance of the JFNs remains unchanged (and in agreement with the PFNs) until physical thresholds are crossed. For example, for quark vs. gluon jets this threshold is around 5 GeV, whereas for Z vs. QCD jets it is determined by the kinematics of the hadronic boosted decay products of the Z -boson. The analogous threshold for the soft momentum cut is $\mathcal{O}(1 \text{ GeV})$. This indicates that while the jet classification tasks are IRC safe, nonperturbative physics is generally relevant for jet classification. In addition, we found that JFNs may offer a decreased model dependence for certain classification tasks with

only a modest tradeoff in performance. This was illustrated in section 7 using a toy example and applications may depend on the type of experimental analysis that is carried out. This observation may lead to interesting applications of JFNs in collider phenomenology.

Our results shed new light on the information that machines learn in high-energy physics applications. As more powerful algorithms will be developed it will be interesting to revisit the question of the potential gap between classifiers based on IRC-safe and unsafe information. While more work is needed in this direction, our work represents an important step toward increasing the interpretability of machine learning methods in high-energy physics. In addition, we anticipate various applications of JFNs in heavy-ion collisions and the future Electron-Ion Collider [104].

Acknowledgments

We thank Giacinto Piacquadio for helpful discussions. DA is supported by the NSF Grant PHY2210533 and the Onassis Foundation. JM, MP are supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, under the contract DE-AC02-05CH11231. FR was supported by the Simons Foundation under the Simons Bridge program for Postdoctoral Fellowships at SCGP and YITP award number 815892; the NSF, award number 1915093; the DOE Contract No. DE-AC05-06OR23177, under which Jefferson Science Associates, LLC operates Jefferson Lab, and in part by the DOE, Office of Nuclear Physics, Early Career Program under contract No. DE-SC0024358. AL is supported in part by the UC Southern California Hub, with funding from the UC National Laboratories division of the University of California Office of the President. This research used resources of the National Energy Research Scientific Computing Center (NERSC), which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY4.0](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.

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](https://inspirehep.net/literature/1709044)].
- [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](https://inspirehep.net/literature/1803069)].
- [3] S. Marzani, G. Soyez and M. Spannowsky, *Looking inside jets: an introduction to jet substructure and boosted-object phenomenology*, Springer (2019) [[DOI:10.1007/978-3-030-15709-8](https://doi.org/10.1007/978-3-030-15709-8)] [[INSPIRE](https://inspirehep.net/literature/1709044)].
- [4] L. Lonnblad, C. Peterson and T. Rognvaldsson, *Finding gluon jets with a neural trigger*, *Phys. Rev. Lett.* **65** (1990) 1321 [[INSPIRE](https://inspirehep.net/literature/1709044)].
- [5] L. de Oliveira et al., *Jet-images — deep learning edition*, *JHEP* **07** (2016) 069 [[arXiv:1511.05190](https://arxiv.org/abs/1511.05190)] [[INSPIRE](https://inspirehep.net/literature/1709044)].

- [6] 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](#)].
- [7] 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](#)].
- [8] G. Kasieczka, T. Plehn, M. Russell and T. Schell, *Deep-learning top taggers or the end of QCD?*, *JHEP* **05** (2017) 006 [[arXiv:1701.08784](#)] [[INSPIRE](#)].
- [9] G. Louppe, K. Cho, C. Becot and K. Cranmer, *QCD-aware recursive neural networks for jet physics*, *JHEP* **01** (2019) 057 [[arXiv:1702.00748](#)] [[INSPIRE](#)].
- [10] K. Datta and A. Larkoski, *How much information is in a jet?*, *JHEP* **06** (2017) 073 [[arXiv:1704.08249](#)] [[INSPIRE](#)].
- [11] 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](#)].
- [12] H. Qu and L. Gouskos, *ParticleNet: jet tagging via particle clouds*, *Phys. Rev. D* **101** (2020) 056019 [[arXiv:1902.08570](#)] [[INSPIRE](#)].
- [13] G. Kasieczka, S. Marzani, G. Soyez and G. Stagnitto, *Towards machine learning analytics for jet substructure*, *JHEP* **09** (2020) 195 [[arXiv:2007.04319](#)] [[INSPIRE](#)].
- [14] F.A. Dreyer and H. Qu, *Jet tagging in the Lund plane with graph networks*, *JHEP* **03** (2021) 052 [[arXiv:2012.08526](#)] [[INSPIRE](#)].
- [15] A. Butter et al., *The machine learning landscape of top taggers*, *SciPost Phys.* **7** (2019) 014 [[arXiv:1902.09914](#)] [[INSPIRE](#)].
- [16] A. Butter, G. Kasieczka, T. Plehn and M. Russell, *Deep-learned top tagging with a Lorentz layer*, *SciPost Phys.* **5** (2018) 028 [[arXiv:1707.08966](#)] [[INSPIRE](#)].
- [17] Y.-C.J. Chen, C.-W. Chiang, G. Cottin and D. Shih, *Boosted W and Z tagging with jet charge and deep learning*, *Phys. Rev. D* **101** (2020) 053001 [[arXiv:1908.08256](#)] [[INSPIRE](#)].
- [18] J.Y. Araz and M. Spannowsky, *Combine and conquer: event reconstruction with Bayesian ensemble neural networks*, *JHEP* **04** (2021) 296 [[arXiv:2102.01078](#)] [[INSPIRE](#)].
- [19] S. Gong et al., *An efficient Lorentz equivariant graph neural network for jet tagging*, *JHEP* **07** (2022) 030 [[arXiv:2201.08187](#)] [[INSPIRE](#)].
- [20] M.D. Schwartz, *Modern machine learning and particle physics*, *Harvard Data Sci. Rev.* (2021) [[arXiv:2103.12226](#)] [[INSPIRE](#)].
- [21] A. Khot, M.S. Neubauer and A. Roy, *A detailed study of interpretability of deep neural network based top taggers*, *Mach. Learn. Sci. Tech.* **4** (2023) 035003 [[arXiv:2210.04371](#)] [[INSPIRE](#)].
- [22] J. Lin, M. Freytsis, I. Moutl and B. Nachman, *Boosting $H \rightarrow b\bar{b}$ with machine learning*, *JHEP* **10** (2018) 101 [[arXiv:1807.10768](#)] [[INSPIRE](#)].
- [23] C.K. Khosa and S. Marzani, *Higgs boson tagging with the Lund jet plane*, *Phys. Rev. D* **104** (2021) 055043 [[arXiv:2105.03989](#)] [[INSPIRE](#)].
- [24] Y. Nakai, D. Shih and S. Thomas, *Strange jet tagging*, [arXiv:2003.09517](#) [[INSPIRE](#)].
- [25] CTEQ collaboration, *Handbook of perturbative QCD: version 1.0*, *Rev. Mod. Phys.* **67** (1995) 157 [[INSPIRE](#)].
- [26] 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](#)].

- [27] A. Romero et al., *Safety of quark/gluon jet classification*, [arXiv:2103.09103](#) [INSPIRE].
- [28] 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].
- [29] A.J. Larkoski and E.M. Metodiev, *A theory of quark vs. gluon discrimination*, *JHEP* **10** (2019) 014 [[arXiv:1906.01639](#)] [INSPIRE].
- [30] S. Choi, S.J. Lee and M. Perelstein, *Infrared safety of a neural-net top tagging algorithm*, *JHEP* **02** (2019) 132 [[arXiv:1806.01263](#)] [INSPIRE].
- [31] Y.S. Lai, J. Mulligan, M. Płoskoń and F. Ringer, *The information content of jet quenching and machine learning assisted observable design*, *JHEP* **10** (2022) 011 [[arXiv:2111.14589](#)] [INSPIRE].
- [32] R. Das, G. Kasieczka and D. Shih, *Feature selection with distance correlation*, *Phys. Rev. D* **109** (2024) 054009 [[arXiv:2212.00046](#)] [INSPIRE].
- [33] T. Fausett, J. Thaler and D. Whiteson, *Mapping machine-learned physics into a human-readable space*, *Phys. Rev. D* **103** (2021) 036020 [[arXiv:2010.11998](#)] [INSPIRE].
- [34] A. Bogatskiy et al., *Lorentz group equivariant neural network for particle physics*, [arXiv:2006.04780](#) [INSPIRE].
- [35] Y. Lu et al., *Resolving extreme jet substructure*, *JHEP* **08** (2022) 046 [[arXiv:2202.00723](#)] [INSPIRE].
- [36] R.K. Ellis, W.J. Stirling and B.R. Webber, *QCD and collider physics*, Cambridge University Press, Cambridge, U.K. (2011) [[DOI:10.1017/CB09780511628788](#)] [INSPIRE].
- [37] S.D. Ellis and D.E. Soper, *Successive combination jet algorithm for hadron collisions*, *Phys. Rev. D* **48** (1993) 3160 [[hep-ph/9305266](#)] [INSPIRE].
- [38] K. Datta and A.J. Larkoski, *Novel jet observables from machine learning*, *JHEP* **03** (2018) 086 [[arXiv:1710.01305](#)] [INSPIRE].
- [39] 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].
- [40] J. Thaler and K. Van Tilburg, *Identifying boosted objects with N -subjettiness*, *JHEP* **03** (2011) 015 [[arXiv:1011.2268](#)] [INSPIRE].
- [41] J. Thaler and K. Van Tilburg, *Maximizing boosted top identification by minimizing N -subjettiness*, *JHEP* **02** (2012) 093 [[arXiv:1108.2701](#)] [INSPIRE].
- [42] V. Mikuni and F. Canelli, *ABCNet: an attention-based method for particle tagging*, *Eur. Phys. J. Plus* **135** (2020) 463 [[arXiv:2001.05311](#)] [INSPIRE].
- [43] H. Qu, C. Li and S. Qian, *Particle transformer for jet tagging*, [arXiv:2202.03772](#) [INSPIRE].
- [44] A.J. Larkoski, S. Marzani and J. Thaler, *Sudakov safety in perturbative QCD*, *Phys. Rev. D* **91** (2015) 111501 [[arXiv:1502.01719](#)] [INSPIRE].
- [45] A.J. Larkoski and J. Thaler, *Unsafe but calculable: ratios of angularities in perturbative QCD*, *JHEP* **09** (2013) 137 [[arXiv:1307.1699](#)] [INSPIRE].
- [46] C.F. Berger, T. Kucs and G.F. Sterman, *Event shape/energy flow correlations*, *Phys. Rev. D* **68** (2003) 014012 [[hep-ph/0303051](#)] [INSPIRE].
- [47] K. Fraser and M.D. Schwartz, *Jet charge and machine learning*, *JHEP* **10** (2018) 093 [[arXiv:1803.08066](#)] [INSPIRE].

- [48] K. Lee et al., *Machine learning-based jet and event classification at the electron-ion collider with applications to hadron structure and spin physics*, *JHEP* **03** (2023) 085 [[arXiv:2210.06450](#)] [[INSPIRE](#)].
- [49] Z.-B. Kang, A.J. Larkoski and J. Yang, *Towards a nonperturbative formulation of the jet charge*, *Phys. Rev. Lett.* **130** (2023) 151901 [[arXiv:2301.09649](#)] [[INSPIRE](#)].
- [50] M. Zaheer et al., *Deep sets*, [arXiv:1703.06114](#) [[INSPIRE](#)].
- [51] E. Wagstaff et al., *On the limitations of representing functions on sets*, [arXiv:1901.09006](#).
- [52] B. Bloem-Reddy and Y.W. Teh, *Probabilistic symmetries and invariant neural networks*, [arXiv:1901.06082](#).
- [53] F.A. Dreyer, G. Soyez and A. Takacs, *Quarks and gluons in the Lund plane*, *JHEP* **08** (2022) 177 [[arXiv:2112.09140](#)] [[INSPIRE](#)].
- [54] A. Banfi, G.P. Salam and G. Zanderighi, *Principles of general final-state resummation and automated implementation*, *JHEP* **03** (2005) 073 [[hep-ph/0407286](#)] [[INSPIRE](#)].
- [55] L.G. Almeida et al., *Top jets at the LHC*, *Phys. Rev. D* **79** (2009) 074012 [[arXiv:0810.0934](#)] [[INSPIRE](#)].
- [56] A.J. Larkoski, G.P. Salam and J. Thaler, *Energy correlation functions for jet substructure*, *JHEP* **06** (2013) 108 [[arXiv:1305.0007](#)] [[INSPIRE](#)].
- [57] A.J. Larkoski, I. Moult and D. Neill, *Power counting to better jet observables*, *JHEP* **12** (2014) 009 [[arXiv:1409.6298](#)] [[INSPIRE](#)].
- [58] G. Soyez et al., *Pileup subtraction for jet shapes*, *Phys. Rev. Lett.* **110** (2013) 162001 [[arXiv:1211.2811](#)] [[INSPIRE](#)].
- [59] ALICE collaboration, *Measurement of inclusive and leading subjet fragmentation in pp and Pb-Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV*, *JHEP* **05** (2023) 245 [[arXiv:2204.10270](#)] [[INSPIRE](#)].
- [60] Y. Chen et al., *Jet energy spectrum and substructure in e^+e^- collisions at 91.2 GeV with ALEPH archived data*, *JHEP* **06** (2022) 008 [[arXiv:2111.09914](#)] [[INSPIRE](#)].
- [61] M. Cacciari, G.P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063 [[arXiv:0802.1189](#)] [[INSPIRE](#)].
- [62] 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](#)].
- [63] M. Cacciari, G.P. Salam and G. Soyez, *FastJet user manual*, *Eur. Phys. J. C* **72** (2012) 1896 [[arXiv:1111.6097](#)] [[INSPIRE](#)].
- [64] G.C. Blazey et al., *Run II jet physics*, in the proceedings of the *Physics at run II: QCD and weak boson physics workshop. Final general meeting*, (2000) [[hep-ex/0005012](#)] [[INSPIRE](#)].
- [65] I.W. Stewart et al., *XCone: N-jettiness as an exclusive cone jet algorithm*, *JHEP* **11** (2015) 072 [[arXiv:1508.01516](#)] [[INSPIRE](#)].
- [66] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159 [[arXiv:1410.3012](#)] [[INSPIRE](#)].
- [67] 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](#)].
- [68] 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](#)].

- [69] Z.-B. Kang, F. Ringer and W.J. Waalewijn, *The energy distribution of subjets and the jet shape*, *JHEP* **07** (2017) 064 [[arXiv:1705.05375](#)] [[INSPIRE](#)].
- [70] D. Neill, F. Ringer and N. Sato, *Leading jets and energy loss*, *JHEP* **07** (2021) 041 [[arXiv:2103.16573](#)] [[INSPIRE](#)].
- [71] W.J. Waalewijn, *Calculating the charge of a jet*, *Phys. Rev. D* **86** (2012) 094030 [[arXiv:1209.3019](#)] [[INSPIRE](#)].
- [72] 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](#)].
- [73] P. Komiske, E. Metodiev and J. Thaler, *Pythia8 quark and gluon jets for energy flow*, *Zenodo* (2019).
- [74] J. Bellm et al., *Herwig 7.2 release note*, *Eur. Phys. J. C* **80** (2020) 452 [[arXiv:1912.06509](#)] [[INSPIRE](#)].
- [75] 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](#)].
- [76] D. Napoletano and G. Soyez, *Computing N-subjettiness for boosted jets*, *JHEP* **12** (2018) 031 [[arXiv:1809.04602](#)] [[INSPIRE](#)].
- [77] F. Chollet et al., *Keras*, <https://github.com/fchollet/keras>, (2015).
- [78] M. Abadi et al., *TensorFlow: large-scale machine learning on heterogeneous distributed systems*, [arXiv:1603.04467](#).
- [79] 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., (2010), p. 807.
- [80] D.P. Kingma and J. Ba, *Adam: a method for stochastic optimization*, [arXiv:1412.6980](#) [[INSPIRE](#)].
- [81] D.R. Cox, *The regression analysis of binary sequences*, *J. Roy. Stat. Soc. B* **20** (1958) 215.
- [82] D. Neill and W.J. Waalewijn, *Entropy of a jet*, *Phys. Rev. Lett.* **123** (2019) 142001 [[arXiv:1811.01021](#)] [[INSPIRE](#)].
- [83] C. Frye, A.J. Larkoski, J. Thaler and K. Zhou, *Casimir meets Poisson: improved quark/gluon discrimination with counting observables*, *JHEP* **09** (2017) 083 [[arXiv:1704.06266](#)] [[INSPIRE](#)].
- [84] J. Gallicchio and M.D. Schwartz, *Quark and gluon tagging at the LHC*, *Phys. Rev. Lett.* **107** (2011) 172001 [[arXiv:1106.3076](#)] [[INSPIRE](#)].
- [85] S.J. Brodsky and J.F. Gunion, *Hadron multiplicity in color gauge theory models*, *Phys. Rev. Lett.* **37** (1976) 402 [[INSPIRE](#)].
- [86] K. Konishi, A. Ukawa and G. Veneziano, *A simple algorithm for QCD jets*, *Phys. Lett. B* **78** (1978) 243 [[INSPIRE](#)].
- [87] OPAL collaboration, *Multiplicity distributions of gluon and quark jets and tests of QCD analytic predictions*, *Eur. Phys. J. C* **1** (1998) 479 [[hep-ex/9708029](#)] [[INSPIRE](#)].
- [88] ALEPH collaboration, *Measurements of the structure of quark and gluon jets in hadronic Z decays*, *Eur. Phys. J. C* **17** (2000) 1 [[INSPIRE](#)].
- [89] DELPHI collaboration, *Investigation of the splitting of quark and gluon jets*, *Eur. Phys. J. C* **4** (1998) 1 [[INSPIRE](#)].

- [90] ATLAS collaboration, *ATLAS b-jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **79** (2019) 970 [[arXiv:1907.05120](#)] [[INSPIRE](#)].
- [91] ATLAS collaboration, *Measurement of the c-jet mistagging efficiency in $t\bar{t}$ events using pp collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS detector*, *Eur. Phys. J. C* **82** (2022) 95 [[arXiv:2109.10627](#)] [[INSPIRE](#)].
- [92] ATLAS collaboration, *Deep sets based neural networks for impact parameter flavour tagging in ATLAS*, [ATL-PHYS-PUB-2020-014](#), CERN, Geneva, Switzerland (2020).
- [93] C.K. Khosa and S. Marzani, *Higgs boson tagging with the Lund jet plane*, *Phys. Rev. D* **104** (2021) 055043 [[arXiv:2105.03989](#)] [[INSPIRE](#)].
- [94] ATLAS collaboration, *Constituent-based top-quark tagging with the ATLAS detector*, [ATL-PHYS-PUB-2022-039](#), CERN, Geneva, Switzerland (2022) [[INSPIRE](#)].
- [95] ATLAS collaboration, *Digluon tagging using $\sqrt{s} = 13$ TeV pp collisions in the ATLAS detector*, [ATL-PHYS-PUB-2021-027](#), CERN, Geneva, Switzerland (2021) [[INSPIRE](#)].
- [96] L.M. Dery, B. Nachman, F. Rubbo and A. Schwartzman, *Weakly supervised classification in high energy physics*, *JHEP* **05** (2017) 145 [[arXiv:1702.00414](#)] [[INSPIRE](#)].
- [97] E.M. Metodiev, B. Nachman and J. Thaler, *Classification without labels: learning from mixed samples in high energy physics*, *JHEP* **10** (2017) 174 [[arXiv:1708.02949](#)] [[INSPIRE](#)].
- [98] E.M. Metodiev and J. Thaler, *Jet topics: disentangling quarks and gluons at colliders*, *Phys. Rev. Lett.* **120** (2018) 241602 [[arXiv:1802.00008](#)] [[INSPIRE](#)].
- [99] R. Kogler et al., *Jet substructure at the Large Hadron Collider: experimental review*, *Rev. Mod. Phys.* **91** (2019) 045003 [[arXiv:1803.06991](#)] [[INSPIRE](#)].
- [100] 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](#)].
- [101] P. Gras et al., *Systematics of quark/gluon tagging*, *JHEP* **07** (2017) 091 [[arXiv:1704.03878](#)] [[INSPIRE](#)].
- [102] ATLAS collaboration, *Light-quark and gluon jet discrimination in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **74** (2014) 3023 [[arXiv:1405.6583](#)] [[INSPIRE](#)].
- [103] A. Butter, B.M. Dillon, T. Plehn and L. Vogel, *Performance versus resilience in modern quark-gluon tagging*, *SciPost Phys. Core* **6** (2023) 085 [[arXiv:2212.10493](#)] [[INSPIRE](#)].
- [104] 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](#)].