

Towards machine learning analytics for jet substructure

Gregor Kasieczka,^a Simone Marzani,^b Gregory Soyez^c and Giovanni Stagnitto^{d,e,f}

^a*Institut für Experimentalphysik, Universität Hamburg,
Luruper Chaussee 149, Hamburg D-22761, Germany*

^b*Dipartimento di Fisica, Università di Genova and INFN, Sezione di Genova,
Via Dodecaneso 33, Genoa 16146, Italy*

^c*Institut de Physique Théorique, Paris Saclay University, CEA Saclay,
CNRS Orme des Merisiers, Bat 774, Gif-sur-Yvette F-91191, France*

^d*Sorbonne Université, CNRS, Laboratoire de Physique Théorique et Hautes Énergies, LPTHE,
Paris F-75005, France*

^e*Université de Paris, LPTHE, Paris F-75005, France*

^f*Tif Lab, Dipartimento di Fisica, Università di Milano and INFN, Sezione di Milano,
Via Celoria 16, Milano I-20133, Italy*

E-mail: gregor.kasieczka@uni-hamburg.de, simone.marzani@ge.infn.it,
gregory.soyez@ipht.fr, giovanni.stagnitto@lpthe.jussieu.fr

ABSTRACT: The past few years have seen a rapid development of machine-learning algorithms. While surely augmenting performance, these complex tools are often treated as black-boxes and may impair our understanding of the physical processes under study. The aim of this paper is to move a first step into the direction of applying expert-knowledge in particle physics to calculate the optimal decision function and test whether it is achieved by standard training, thus making the aforementioned black-box more transparent. In particular, we consider the binary classification problem of discriminating quark-initiated jets from gluon-initiated ones. We construct a new version of the widely used N -subjettiness, which features a simpler theoretical behaviour than the original one, while maintaining, if not exceeding, the discrimination power. We input these new observables to the simplest possible neural network, i.e. the one made by a single neuron, or perceptron, and we analytically study the network behaviour at leading logarithmic accuracy. We are able to determine under which circumstances the perceptron achieves optimal performance. We also compare our analytic findings to an actual implementation of a perceptron and to a more realistic neural network and find very good agreement.

KEYWORDS: Jets, QCD Phenomenology

ARXIV EPRINT: [2007.04319](https://arxiv.org/abs/2007.04319)

Contents

1	Introduction	1
2	Quark/gluon tagging with N-subjettiness	3
2.1	Primary N -subjettiness	5
2.2	Primary N -subjettiness at leading-logarithmic accuracy	6
2.3	Fixed-coupling limit	8
3	Perceptron analytics	8
3.1	Generic considerations	8
3.2	Minimisation of the cost function	10
4	Perceptron numerics	14
5	Monte Carlo studies	16
6	Conclusions and outlook	22
A	Details of the analytic calculations	23
B	Analytic results for the cost function	25
B.1	Cross-entropy loss	25
B.2	Quadratic loss	26
C	Scaling of a_n in the log inputs case	26

1 Introduction

The CERN Large Hadron Collider (LHC) is continuously exploring the high-energy frontier. However, although well-motivated from a theoretical viewpoint, no conclusive evidence of new particles or interactions, and no significant cracks in the Standard Model (SM) of particle physics have been found. This challenges the particle physics community, theorists and experimenters alike, to find new ways to interrogate and interpret the data. In the context of analyses involving hadronic final states, jet substructure has emerged as an important tool for searches at the LHC. Consequently, a vibrant field of theoretical and experimental research has developed in the past ten years, producing a variety of studies and techniques [1]. Many grooming and tagging algorithms have been developed, successfully tested, and are currently used in experimental analyses [2–7].

The field of jet substructure is currently undergoing a revolution. The rapid development, within and outside academia, of machine-learning (ML) techniques is having a profound impact on particle physics in general and jet physics in particular. In the context of high-energy physics, ML classification algorithms are typically trained on a control sample, which could be either Monte Carlo pseudo-data or a high-purity dataset, and then applied to an unknown sample to classify its properties. These ideas have been exploited in particle physics for a long time. However, because of limitations on the algorithms' efficiency and on

computers' power, so-far algorithms were applied to relatively low-dimensional projections of the full radiation pattern that one wished to classify. Even so, such projections usually correspond to physically-motivated observables, such as, e.g. the jet mass, jet shapes, and therefore limitation in performance was mitigated with physics understanding. The current ML revolution moves away from low-dimensional projections and exploit deep neural network (NN) to perform classification. Early progress was made on supervised classification of known particles [8–12] to the point where powerful network architectures are now available [13–17] and the focus lies on improving the stability [18–22] of these approaches.

These techniques, developed by both theorists and experimentalists, have already been tested in the LHC environment and dedicated studies have been performed by the ATLAS and CMS collaborations [23–25]. Despite its rapid development and the unquestionable improvements that ML brings to particle physics, it is often met with a certain degree of suspicion, because physicists often feel uneasy about relying on black-box techniques. We believe that the situation is not too dissimilar to the one we faced a decade ago, at the birth of jet substructure. At the time there was a plethora of new tagging algorithms and a deeper understanding was reached only when theoretical studies based on QCD were performed, see e.g. [26–38].

A common expectation is that well-trained networks can achieve the performance of the likelihood ratio, i.e. they can learn a monotonic function of the likelihood ratio. In this paper, we are in the almost-unique position to analytically derive the weights corresponding to a cut on the likelihood ratio for a realistic problem and to test whether they agree with a network trained the usual way. Inspired by the successful approach of the past, we want to anchor this investigation in perturbative QCD, with a first-principle understanding approach, by looking at a situation where most studies can be performed analytically, and afterwards by investing how our findings in the simplest case compare to a more general setup.

For this current study, we focus on an issue that has received a great deal of both theoretical and experimental attention, namely the ability of disentangling jets that can be thought as having originated by the fragmentation of a high-energy quark (henceforth quark jets), from the ones originated from a gluon (henceforth gluon jets) [39]. In particular, in section 2, we introduce a variant of the well-known N -subjettiness variables [40] based on the primary Lund plane declustering tree [41]. This *primary* N -subjettiness \mathcal{T}_N is more amenable to an all-order QCD analysis, which we perform at leading logarithmic (LL) accuracy, while maintaining, if not improving, the quark/gluon discrimination power. This definition is such that, if we measure the N -subjettiness variables $\{\mathcal{T}_1 \dots \mathcal{T}_n\}$, then, at LL, a cut on the likelihood ratio simply corresponds to a cut on \mathcal{T}_n . With this definition, and within the stated accuracy, we are able to write simple analytic expressions for the cumulative distributions and for the so-called area under the curve (AUC), which measures the discriminating power of the observable, for any value of n .¹

We then move in section 3 to use the aforementioned primary N -subjettiness variables as input to a NN. Because we expect the optimal discriminant to be a cut on \mathcal{T}_n , we

¹Expressions for cumulative distributions can also be obtained for the standard N -subjettiness definition, see e.g. [29, 42], although their complexity increases with n to a point where other quantities, such as, for instance, the AUC are no longer tractable analytically.

can start by considering the simplest possible NN, namely a single neuron, or *perceptron*, with n inputs and one output. This opens up the interesting possibility of performing analytic calculations that describe the behaviour of the perceptron at LL accuracy in QCD and allow us to determine the optimal decision function. The ability to fully control the behaviour of a ML technique analytically from first-principles is the main novelty of this paper. In sections 4 and 5, we compare our theoretical findings with actual networks with architecture of increasing complexity using pseudo-data, which we generate either according to a leading-logarithmic distribution in QCD (section 3) or using a general-purpose Monte Carlo parton shower (section 5). Finally, we draw our conclusions in section 6.

2 Quark/gluon tagging with N -subjettiness

In the framework of perturbative QCD, final-state jets, i.e. collimated sprays of hadrons, are formed by the fragmentation of high-energetic partons, which emerge from the hard scattering. While the association of a single parton, say a quark or a gluon, to a final-state jet is an intrinsically ambiguous — if not ill-defined — operation, essentially because of higher-order corrections — it is however possible to employ operational definitions that associate the value of a measurable quantity to an enriched sample of quarks or gluons [39, 43]. Many such observables have been proposed in the literature, ranging from more traditional jet observables such as angularities [44] or energy-correlation functions [36], to observables that exploit more modern jet substructure techniques, such as the iterated soft-drop multiplicity [45] or the Les Houches multiplicity [46], to operational definitions inspired by data science [47, 48].

A general approach to quantitatively assess the power of quark/gluon discrimination was put forward in ref. [42]. The core idea of that approach is to measure multiple infra-red and collinear (IRC) safe observables on a single jet as a means to characterise the available emission phase space. The observables of choice were the N -subjettiness variables [40, 49] τ_N , where the measurement of the first n variables $\{\tau_1 \dots \tau_n\}$ is able to resolve n emissions in the jet.² The N -subjettiness shape is one of the most used jet substructure variables. As its name suggests, it aims to identify jets with an N -prong structure, taking inspiration from the event-shape N -jettiness [52]. In order to achieve this, a set of axes a_1, \dots, a_N is introduced and the N -subjettiness is defined as

$$\tau_N^{(\beta)} = \frac{1}{p_t R_0^\beta} \sum_{i \in \text{jet}} p_{ti} \min(\Delta_{ia_1}^\beta, \dots, \Delta_{ia_N}^\beta), \quad (2.1)$$

where β is a free parameter often set to unity, p_t and R_0 are the jet transverse momentum and radius respectively and $\Delta_{ia_j} = \sqrt{\Delta y_{ia_j}^2 + \Delta \phi_{ia_j}^2}$ is the distance between particle i and the axis a_j in the azimuth-rapidity plane. Note that the axes a_j can be defined in several ways, for a review see for instance ref. [1]. Crucially, IRC safety guarantees that

²There is indeed more information in a jet than what is captured by this set of N -subjettiness variables. One could study, for instance, an extended set of N -subjettiness variables [50], energy-flow polynomial [51], or counting observables, e.g. [45]. We thank Andrew Larkoski and Ben Nachman for comments on this point.

N -subjettiness distributions can be meaningfully calculated in perturbative QCD. These distributions have been the subject of several theoretical investigations [29, 30, 53] (see also [34, 54, 55] for theoretical studies on the related observable D_2). The N -subjettiness distributions were calculated explicitly at LL accuracy, i.e. in the limit in which all emissions are strongly ordered, in ref. [42] for the cases $n = 1, 2, 3$. Furthermore, the Authors of ref. [42] were able to calculate from first-principle familiar quantities that enter statistical analyses. We start our discussion with a brief recap of the elements in ref. [42] which are relevant for our current study.

The likelihood ratio, which according to the Neyman-Pearson lemma [56] provides the best single-variable discriminant, is simply given by the ratio of the probability distributions computed for background (gluons) and signal (quark) jets:

$$\mathcal{L}(\tau_1, \dots, \tau_n) = \frac{p_B}{p_S} = \frac{p_g(\tau_1, \dots, \tau_n)}{p_q(\tau_1, \dots, \tau_n)}, \quad (2.2)$$

where

$$p_i(\tau_1, \dots, \tau_n) = \frac{1}{\sigma_i} \frac{d\sigma_i}{d\tau_1 \dots d\tau_n}, \quad i = q, g. \quad (2.3)$$

In order to assess the discriminating power of an observable, Receiver Operating Characteristic (ROC) curves are often considered. These curves show the background (gluon) efficiency against the signal (quark) efficiency and are remarkably useful to directly compare the performance of different tools. Furthermore, note that ROC curves are invariant under monotonic rescaling of the inputs.

For a single-variable observables $V \equiv V(\tau_1, \dots, \tau_n)$, one can define the normalised cumulative distribution

$$\Sigma_i(v) = \int d\tau_1 \dots d\tau_n p_i(\tau_1, \dots, \tau_n) \Theta(V(\tau_1, \dots, \tau_n) < v). \quad (2.4)$$

The ROC curve is then obtained as

$$\text{ROC}(x) = \Sigma_g(\Sigma_q^{-1}(x)) \quad (2.5)$$

where x is the signal (quark) efficiency. The area under the ROC curve (AUC) can be used as a quantifier of the discrimination power, where $\text{AUC} = 0$ corresponds to perfect performance.³ For example, if one considers an IRC observable v that only resolves one emission, e.g. $v = \tau_1$, then the LL cumulative distributions for quarks and gluons simply differ by their colour factor, leading to the so-called Casimir scaling:

$$\text{ROC}(x) = (x)^{\frac{C_A}{C_F}}, \quad \text{AUC} = \int_0^1 dx x^{\frac{C_A}{C_F}} = \frac{C_F}{C_F + C_A}. \quad (2.6)$$

The fact that τ_n resolves n emissions inside the jet results in a rather complicated structure, even in the limit where the emissions are strongly ordered. This is because the

³Note that one could have alternatively defined the ROC curve as showing the background rejection against the signal efficiency i.e. $\text{ROC}(x) = 1 - \Sigma_g(\Sigma_q^{-1}(x))$. In this case the perfect performance would correspond to $\text{AUC} = 1$.

emissions that set the values of the observables τ_i , which are always gluons at LL accuracy, can either be primary emissions, i.e. they originate from the original hard parton which could be a quark or a gluon, or they can originate from subsequent gluon splittings. If we consider, for instance, the case of a jet initiated by a hard quark, one ends up with n contributions with colour factors $C_F^{n-i}C_A^i$, $i = 0, \dots, n - 1$. Furthermore, a rather intricate resummation structure emerges, i.e. a Sudakov form factor with both C_F and C_A contributions and depending on the complete tower of n emissions. It is clear that this intricate structure does not facilitate analytical calculations, especially because, in this study, we are not interested in considering p_q and p_g as final results, but rather, as inputs for subsequent calculations.

As a consequence, we find convenient to introduce a variant of N -subjettiness that is sensitive, at LL accuracy, only to primary emissions, such that the distributions p_i are determined by strongly-ordered gluon emissions off the initial hard parton. We present this new observable in the next section.

2.1 Primary N -subjettiness

We define the new observable primary N -subjettiness as follows. Starting from a jet of radius R_0 , one first builds the list of primary Lund declusterings [41]:

1. Recluster the jet constituents with the Cambridge/Aachen algorithm [57, 58].
2. Iteratively undo the last step of the clustering $j \rightarrow j_1 + j_2$, with $p_{t1} > p_{t2}$. At step i ($i = 1, \dots, m$), define

$$\tilde{p}_{ti} = p_{t2} \quad \text{and} \quad \Delta_i = \sqrt{\Delta y_{12}^2 + \Delta\phi_{12}^2}. \tag{2.7}$$

Repeat the procedure with $j = j_1$, i.e. following the harder branch of the declustering.

3. When the declustering terminates, i.e. when j is no longer coming from a $j_1 + j_2$ clustering, define \tilde{p}_{t0} as the transverse momentum of j .

From the set of transverse momenta, we can define the momentum fractions

$$z_i = \frac{\tilde{p}_{ti}}{\sum_{i=0}^m \tilde{p}_{ti}} \quad i = 0, \dots, m, \tag{2.8}$$

where we note that the final hard momentum \tilde{p}_{t0} is included in the normalisation. This produces a set of values (z_i, Δ_i) , for $i = 1, \dots, m$, that we order such that $z_1\Delta_1^\beta \geq z_2\Delta_2^\beta \geq \dots \geq z_m\Delta_m^\beta$. The primary N -subjettiness is then defined as⁴

$$\mathcal{T}_N = \sum_{i=N}^m z_i \left(\frac{\Delta_i}{R_0} \right)^\beta. \tag{2.9}$$

Note that $\mathcal{T}_1 \geq \dots \geq \mathcal{T}_n$, like with the standard N -subjettiness τ_N . The primary N -subjettiness definition in eq. (2.9) is very similar to the standard N -subjettiness definition with the main difference that it is computed based on primary Lund declusterings. The definition of the momentum fractions z_i is such that $\sum_{i=0}^m z_i = 1$.

⁴Note that an alternative definition, equivalent at leading-logarithmic accuracy, but different beyond, would be to define $\mathcal{T}_N^{(\max)} = z_N \left(\frac{\Delta_N}{R_0} \right)^\beta$.

2.2 Primary N -subjettiness at leading-logarithmic accuracy

By construction, the leading-logarithmic (LL) expression for the n -dimensional differential distribution is obtained by considering strongly-ordered independent emissions off the originating hard parton (either a quark or a gluon). In this strongly-ordered limit, only the first term in the r.h.s. of eq. (2.9) should be kept, i.e. $\mathcal{T}_N \approx z_N(\Delta_N/R_0)^\beta$. In that context, the structure of the leading-logarithmic resummation is particularly trivial: one gets a factor for each of the n emissions associated with $\mathcal{T}_1, \dots, \mathcal{T}_n$ as well as a Sudakov factor vetoing any additional real emission with $z(\Delta/R_0)^\beta > \mathcal{T}_n$.⁵ This yields

$$p_i(\mathcal{T}_1, \dots, \mathcal{T}_n) = \left(\prod_{j=1}^n \frac{\mathcal{R}'(\mathcal{T}_j)}{\mathcal{T}_j} \right) (C_i)^n \exp[-C_i \mathcal{R}(\mathcal{T}_n)], \quad (2.10)$$

where $i = q, g$ refers to the flavour of the jet (with $C_q = C_F$ and $C_g = C_A$). The radiator \mathcal{R} , which and has been stripped of its colour factor, is computed in the soft-and-collinear-limit including running-coupling corrections, as appropriate for our LL accuracy:

$$\mathcal{R}(\mathcal{T}) = 2 \int_0^1 \frac{d\theta^2}{\theta^2} \int_0^1 \frac{dz}{z} \frac{\alpha_s(z\theta p_t R_0)}{2\pi} \Theta(z\theta^\beta > \mathcal{T}) = \frac{2}{\beta} \int_{\mathcal{T}}^1 \frac{d\theta^\beta}{\theta^\beta} \int_{\mathcal{T}/\theta^\beta}^1 \frac{dz}{z} \frac{\alpha_s(z\theta p_t R_0)}{\pi}. \quad (2.11)$$

We have also introduced $\mathcal{R}'(\mathcal{T}) = \frac{d\mathcal{R}(\mathcal{T})}{d \log(1/\mathcal{T})}$, where $\log x$ always denotes the natural logarithm of x .

An important observation is the following. From eq. (2.10) we note that the structure of the probability distributions at LL in QCD for primary definition of N -subjettiness is the *same* for quark and gluon jets except for the colour factor, C_F or C_A which appears as an overall factor in both the \mathcal{R}' pre-factors and the Sudakov exponent. (This is not case with the standard definition of N -subjettiness). Consequently, the likelihood ratio eq. (2.2) at LL becomes

$$\mathcal{L}^{\text{LL}} = \left(\frac{C_A}{C_F} \right)^n \exp[-(C_A - C_F) \mathcal{R}(\mathcal{T}_n)], \quad (2.12)$$

which is a monotonic function of $\mathcal{R}(\mathcal{T}_n)$, and hence at LL of \mathcal{T}_n , only. Therefore, at LL, a cut on the likelihood ratio is equivalent to a cut on \mathcal{T}_n . The remarkable simplicity of this result is the strongest motivation for introducing primary N -subjettiness. This observable is thus the ideal laboratory to study analytically how a neural network that takes the primary N -subjettiness variables as inputs performs. Due to the simplicity of the classifier — a cut on a single variable — we expect that even the simplest network, i.e. a single neuron, should be enough to achieve optimal classification performance, i.e. a perceptron should be able to learn a function which is monotonically related to likelihood ratio. Studying more complex architectures — especially including one or more hidden layers — would be very interesting but is not in the scope of this exploratory work.

Analytic studies of a perceptron will be the topic of section 3, but, before moving to that, let us derive a couple of results that allow us to establish primary N -subjettiness

⁵As usual, this is equivalent to saying that real and virtual emissions cancel each other for $z(\Delta/R_0)^\beta < \mathcal{T}_n$ and only virtual corrections contribute for $z(\Delta/R_0)^\beta > \mathcal{T}_n$. These virtual corrections trivially exponentiate.

as an appealing observable on its own, rather than just a shortcut to more tractable analytic results. Firstly, it is interesting to obtain an analytic expression for the cumulative distribution with a cut \mathcal{T} on \mathcal{T}_n . As we have just seen, this is equivalent to a cut on the likelihood. We have

$$\begin{aligned} \Sigma_i(\mathcal{T}_n < \mathcal{T}) &= \int_0^1 d\mathcal{T}_1 \int_0^{\mathcal{T}_1} d\mathcal{T}_2 \cdots \int_0^{\mathcal{T}_{n-2}} d\mathcal{T}_{n-1} \int_0^{\mathcal{T}_{n-1}} d\mathcal{T}_n p_i(\mathcal{T}_1, \dots, \mathcal{T}_n) \Theta(\mathcal{T}_n < \mathcal{T}) \\ &= \int_0^1 d\mathcal{T}_1 \int_0^{\mathcal{T}_1} d\mathcal{T}_2 \cdots \int_0^{\mathcal{T}_{n-2}} d\mathcal{T}_{n-1} \int_0^{\min[\mathcal{T}_{n-1}, \mathcal{T}]} d\mathcal{T}_n p_i(\mathcal{T}_1, \dots, \mathcal{T}_n), \end{aligned} \quad (2.13)$$

where $i = q, g$. The latter expression splits naturally in two terms: if $\mathcal{T}_{n-1} < \mathcal{T}$, we simply find $\Sigma_i(\mathcal{T}_{n-1} < \mathcal{T})$; if $\mathcal{T}_{n-1} > \mathcal{T}$, the exponential in (2.10) factors out and we obtain

$$e^{-C_i \mathcal{R}(\mathcal{T})} \int_{\mathcal{T}}^1 d\mathcal{T}_1 \int_{\mathcal{T}}^{\mathcal{T}_1} d\mathcal{T}_2 \cdots \int_{\mathcal{T}}^{\mathcal{T}_{n-2}} d\mathcal{T}_{n-1} p_i(\mathcal{T}_1, \dots, \mathcal{T}_n) = e^{-C_i \mathcal{R}(\mathcal{T})} \frac{C_i^{n-1} \mathcal{R}^{n-1}(\mathcal{T})}{(n-1)!}. \quad (2.14)$$

By induction we arrive at

$$\Sigma_i(\mathcal{T}_n < \mathcal{T}) = e^{-C_i \mathcal{R}(\mathcal{T})} \sum_{k=1}^n \frac{C_i^{k-1} \mathcal{R}^{k-1}(\mathcal{T})}{(k-1)!} = \frac{\Gamma(n, C_i \mathcal{R}(\mathcal{T}))}{\Gamma(n)}, \quad (2.15)$$

where $\Gamma(n, x)$ is the incomplete Gamma function.

It is also possible to find an analytic expression for the AUC. Exploiting the ROC curve definition in eq. (2.5), we can write the AUC as an integral of the quark and gluon distributions:

$$\begin{aligned} \text{AUC} &= \int_0^1 d\mathcal{T}_{1q} \int_0^{\mathcal{T}_{1q}} d\mathcal{T}_{2q} \cdots \int_0^{\mathcal{T}_{n-1,q}} d\mathcal{T}_{nq} \int_0^1 d\mathcal{T}_{1g} \int_0^{\mathcal{T}_{1g}} d\mathcal{T}_{2g} \cdots \int_0^{\mathcal{T}_{n-1,g}} d\mathcal{T}_{ng} \\ &\quad \cdot p_q(\mathcal{T}_{1q}, \mathcal{T}_{2q}, \dots, \mathcal{T}_{nq}) p_g(\mathcal{T}_{1g}, \mathcal{T}_{2g}, \dots, \mathcal{T}_{ng}) \Theta(\mathcal{T}_{nq} > \mathcal{T}_{ng}). \end{aligned} \quad (2.16)$$

The details of the integration are not all trivial and are thus given in appendix A. Here we just quote the final result:

$$\text{AUC} = 1 - \left(\frac{C_F C_A}{(C_F + C_A)^2} \right)^n \frac{\Gamma(2n)}{\Gamma(n)\Gamma(1+n)} {}_2F_1 \left(1, 2n, 1+n; \frac{C_A}{C_F + C_A} \right). \quad (2.17)$$

We can compare the values given by this expression with the ones computed in [42] for the standard definition of N -subjettiness.

$n = 1$	→	$\text{AUC}_{\text{standard}} = 0.308,$	$\text{AUC}_{\text{primary}} = 0.308$
$n = 2$	→	$\text{AUC}_{\text{standard}} = 0.256,$	$\text{AUC}_{\text{primary}} = 0.226$
$n = 3$	→	$\text{AUC}_{\text{standard}} = 0.231,$	$\text{AUC}_{\text{primary}} = 0.173.$

We conclude that, at least at LL, a cut on the primary N -subjettiness \mathcal{T}_n provides better quark/gluon discrimination power than a cut on the standard N -subjettiness. The comparison of the two different definitions when evaluated on Monte-Carlo generated pseudo-data will be discussed in section 5.

2.3 Fixed-coupling limit

In the previous section we have obtained analytic expressions for primary N -subjettiness distributions that are valid at LL, including running-coupling effects. Henceforth, for sake of simplicity, we are going to consider the fixed coupling limit of eq. (2.11). In this limit, the probability distributions for quarks and gluons eq. (2.10) is

$$p_i(\mathcal{T}_1, \dots, \mathcal{T}_n) = \left(\frac{\alpha_s}{\pi\beta}\right)^n (2C_i)^n \prod_{j=1}^n \left(\frac{\log(1/\mathcal{T}_j)}{\mathcal{T}_j}\right) \exp\left[-\frac{\alpha_s}{\pi\beta} C_i \log^2 \mathcal{T}_n\right], \quad (2.18)$$

and the likelihood ratio eq. (2.12) consequently becomes

$$\mathcal{L}^{\text{LL-f.c.}} = \left(\frac{C_A}{C_F}\right)^n \exp\left[-\frac{C_A - C_F}{\beta} \frac{\alpha_s}{\pi} \log^2 \mathcal{T}_n\right]. \quad (2.19)$$

In order to further simplify our notation, we can also reabsorb the factor $\alpha_s/(\pi\beta)$ in a redefinition of the variables \mathcal{T}_i , or, equivalently, we can define the colour factors C_F and C_A in units of $\alpha_s/(\pi\beta)$. We will do the latter, introducing $\tilde{C}_i = \frac{\alpha_s C_i}{\pi\beta}$. We will also express eq. (2.18) in terms of different variables $x_i = x_i(\mathcal{T}_i)$. We therefore rewrite eq. (2.18) as

$$p_i(x_1, \dots, x_n) = r'(x_1) \cdots r'(x_n) \left(\tilde{C}_i^n \exp\left[-\tilde{C}_i r(x_n)\right]\right). \quad (2.20)$$

In particular, we will consider the following three choices:

$$\text{log square: } x_i(\mathcal{T}_i) = L_i \equiv \log^2 \mathcal{T}_i, \quad r(L_n) = L_n, \quad r'(L_i) = 1, \quad (2.21)$$

$$\text{log: } x_i(\mathcal{T}_i) = l_i \equiv \log(1/\mathcal{T}_i), \quad r(l_n) = l_n^2, \quad r'(l_i) = 2l_i, \quad (2.22)$$

$$\text{linear: } x_i(\mathcal{T}_i) = \mathcal{T}_i, \quad r(\mathcal{T}_n) = \log^2 \mathcal{T}_n, \quad r'(\mathcal{T}_i) = 2 \frac{\log(1/\mathcal{T}_i)}{\mathcal{T}_i}. \quad (2.23)$$

3 Perceptron analytics


3.1 Generic considerations

In this section we will investigate the simplest neural network, namely the *perceptron*, where the input layer is directly linked to the output through a single neuron, as depicted in figure 1a. Analytically, this means that the network output y is a function of a weighted linear combination of the inputs:

$$y = f(\vec{x} \cdot \vec{a} + b) \quad (3.1)$$

with \vec{x} vector of the input variables, while \vec{a} is the vector of the *weights* of the neuron and b is a *bias*. The function f , called *activation function*, allows us to introduce an element of non-linearity in the network. In this paper we will focus on the *sigmoid*, which is widely adopted in the ML community. The sigmoid is defined as

$$f(x) \equiv \sigma(x) = \frac{1}{1 + e^{-x}}, \quad (3.2)$$



(a) The perceptron consists of n input units and one output $y = f(\vec{a} \cdot \vec{x} + b)$. (b) The sigmoid activation function, eq. (3.2).

Figure 1. Sketch of a perceptron (left) and of the sigmoid activation function (right).

and it is shown in figure 1b. Alternative functional forms are possible. For instance, in this context, one could replace the sigmoid with a rectified version (*hard-sigmoid*):

$$f(x) = \max(0, \min(x, 1)). \quad (3.3)$$

The neural network learns the best choice of the weights by minimising the *cost* function, which quantifies the difference between the expected value of the network \hat{y} and the predicted value y (given by eq. (3.1)). In this study we are focusing on the issue of quark versus gluon discrimination, which is an example of a binary classification problem, where we have $\hat{y} = 0$ (associated with quarks) or $\hat{y} = 1$ (associated with gluons). In this context, the *cross-entropy* loss is one of the most common choice for the loss function

$$C(y, \hat{y}) = -(1 - \hat{y}) \log(1 - y) - \hat{y} \log(y). \quad (3.4)$$

This is what we will use for this paper. However, many of the results we obtain also apply to other loss functions, such as, for instance, the *quadratic* loss:

$$C(y, \hat{y}) = (y - \hat{y})^2. \quad (3.5)$$

In order to train the NN, one usually starts with a so-called training sample, i.e. a collection of input vectors $\{\vec{x}_i\}$, each labelled as a quark jet or as a gluon jet. If we have a training sample of $2N$ inputs, equally divided between signal and background labels, we can write the cost function as:

$$\tilde{C}(\vec{a}, b) = \frac{1}{2N} \sum_{i=1}^N \left[C\left(f(\vec{x}_i^{(q)} \cdot \vec{a} + b), 0\right) + C\left(f(\vec{x}_i^{(g)} \cdot \vec{a} + b), 1\right) \right]. \quad (3.6)$$

The input variables $\vec{x}^{(i)}$ are generated according to a probability distribution $p_i(\vec{x})$, with p_q (p_g) being the probability distribution of the inputs \vec{x} for quark (gluon) jets. If the training sample is large, as it usually is, we can rewrite the above equation in the continuous limit:

$$\tilde{C}(\vec{a}, b) = \frac{1}{2} \int d\vec{x} \left[p_q(\vec{x}) C(f(\vec{x} \cdot \vec{a} + b), 0) + p_g(\vec{x}) C(f(\vec{x} \cdot \vec{a} + b), 1) \right]. \quad (3.7)$$

In a general classification problem, the probability distributions of the inputs are unknown. However, in the context of QCD studies, we can exploit expert-knowledge: if we choose the input variables \vec{x} as IRC safe observables, we can apply the powerful machinery of perturbative quantum field theory to determine these distributions at a well-defined and, in principle, systematically improvable accuracy.

In what follows, we will use the primary N -subjettiness variables $\{\mathcal{J}_i\}$ as inputs for our perceptron. Using the results obtained in section 2, we will evaluate eq. (3.7) using probability distributions calculated at LL accuracy. We will then study the global minimum of the cost function. We will focus on the sigmoid activation function, eq. (3.2), and the cross-entropy loss, eq. (3.4), although analogous results can be obtained for the quadratic loss, eq. (3.5). More specifically, our goal is to establish whether the set of weights \vec{a} and bias b that minimises eq. (3.7) corresponds to a cut on \mathcal{J}_n , which corresponds to the likelihood ratio at LL accuracy, cf. eq. (2.12). This corresponds to checking $a_1 = \dots = a_{n-1} = 0$.⁶ Note that the values of a_n and b are not fixed by our LL analysis, even though the network training will converge to some fixed values, corresponding to the minimum of the cost function. We will show that the ability of the perceptron to find the likelihood ratio crucially depends on the functional form of the input variables. We shall consider three cases, all based on the primary N -subjettiness, namely \mathcal{J}_i , $\log \mathcal{J}_i$ and $\log^2 \mathcal{J}_i$.

3.2 Minimisation of the cost function

In order to find the extrema of the cost function eq. (3.7), we consider the partial derivatives with respect to a generic weight a_i or b . With a simple application of the chain rule, we find a set of $n + 1$ simultaneous equations

$$\begin{cases} \frac{\partial \tilde{C}}{\partial a_i} = \frac{1}{2} \int d\vec{x} x_i f'(\vec{x} \cdot \vec{a} + b) \left[p_q(\vec{x}) C'(f(\vec{x} \cdot \vec{a} + b), 0) + p_g(\vec{x}) C'(f(\vec{x} \cdot \vec{a} + b), 1) \right] = 0, \\ \frac{\partial \tilde{C}}{\partial b} = \frac{1}{2} \int d\vec{x} f'(\vec{x} \cdot \vec{a} + b) \left[p_q(\vec{x}) C'(f(\vec{x} \cdot \vec{a} + b), 0) + p_g(\vec{x}) C'(f(\vec{x} \cdot \vec{a} + b), 1) \right] = 0, \end{cases} \quad (3.8)$$

where $i = 1, \dots, n$ and the prime indicates the derivative of a function with respect to its (first) argument.

In general, in order to solve the above system of simultaneous equations, we have to explicitly compute the n -dimensional integral in eq. (3.8). However, in our case, it is possible to find directly a solution at the integrand level. To this purpose, we observe that when the equality

$$\frac{p_q(\vec{x})}{p_g(\vec{x})} = -\frac{C'(f(\vec{x} \cdot \vec{a} + b), 1)}{C'(f(\vec{x} \cdot \vec{a} + b), 0)} \quad (3.9)$$

is satisfied for any value of \vec{x} , then the system of equations is fulfilled. Since the cross-entropy loss, eq. (3.4), satisfies $C'(y, 1) = -\frac{1}{y}$ and $C'(y, 0) = -C'(1 - y, 1)$, eq. (3.9) can

⁶We expect that these coefficients will receive corrections which are suppressed by powers of α_s . These corrections are negligible in the strict leading-logarithmic approximation, where only dominant contributions are kept.

be simplified to

$$\frac{p_q(\vec{x})}{p_g(\vec{x})} = \frac{1 - f(\vec{x} \cdot \vec{a} + b)}{f(\vec{x} \cdot \vec{a} + b)}. \quad (3.10)$$

We note that this result also holds, for instance, in the case of the quadratic loss.

More specifically, we want to compute eq. (3.8) for our probabilities p_i given by eq. (2.20). We explicitly consider 3 cases differing only by what is used as inputs to the perceptron: “log-square inputs”, eq. (2.21), “log inputs”, eq. (2.22), and “linear inputs”, eq. (2.23). In all three cases, we use the sigmoid activation function and the cross-entropy loss. Given the following identities for the sigmoid function:

$$1 - \sigma(x) = \sigma(-x) \quad \text{and} \quad \sigma'(x) = \sigma(x) \sigma(-x), \quad (3.11)$$

and the properties of the cross-entropy loss, eq. (3.8) may be rewritten as:

$$\left\{ \begin{array}{l} \frac{\partial \tilde{\mathcal{C}}}{\partial a_i} = \frac{1}{2} \int d\vec{x} x_i r'(x_1) \cdots r'(x_n) \\ \quad \left[\tilde{C}_F^n e^{-\tilde{C}_F r(x_n)} \sigma(\vec{x} \cdot \vec{a} + b) - \tilde{C}_A^n e^{-\tilde{C}_A r(x_n)} \sigma(-\vec{x} \cdot \vec{a} - b) \right] = 0, \\ \frac{\partial \tilde{\mathcal{C}}}{\partial b} = \frac{1}{2} \int d\vec{x} r'(x_1) \cdots r'(x_n) \\ \quad \left[\tilde{C}_F^n e^{-\tilde{C}_F r(x_n)} \sigma(\vec{x} \cdot \vec{a} + b) - \tilde{C}_A^n e^{-\tilde{C}_A r(x_n)} \sigma(-\vec{x} \cdot \vec{a} - b) \right] = 0. \end{array} \right. \quad (3.12)$$

Log-square inputs. Let us start by considering $L_i = \log^2 \mathcal{T}_i$ as inputs to the perceptron. In this case the probability distributions for quarks and gluons in the fixed-coupling limit are given by eq. (2.21). This is a very lucky scenario because we can determine the minimum at the integrand level. Indeed, the condition eq. (3.10) gives the constraint

$$\left(\frac{\tilde{C}_F}{\tilde{C}_A} \right)^n \exp \left[-(\tilde{C}_F - \tilde{C}_A) L_n \right] = \exp \left[-\vec{a} \cdot \vec{L} - b \right], \quad (3.13)$$

leading to the following solution

$$a_1 = \cdots = a_{n-1} = 0, \quad a_n = \tilde{C}_F - \tilde{C}_A, \quad b = n \log \left(\frac{\tilde{C}_A}{\tilde{C}_F} \right). \quad (3.14)$$

Hence, for the log-square inputs, the minimum of the cost function does agree with the optimal cut on \mathcal{T}_n dictated by the likelihood.

Note that, as $\tilde{C}_F < \tilde{C}_A$, the weight a_n is negative. This is expected, since the sigmoid function is monotonic and we have mapped the gluon (quark) sample to output 1 (0), see eq. (3.6), whereas the gluon sample has larger \mathcal{T}_i and thus smaller L_i : the negative sign of a_n restores the proper ordering between inputs and output of the perceptron.

If we restrict ourselves to the case $n = 2$, it is also possible to explicitly perform the integrals in eq. (3.7) and arrive at an analytic expression for the cost function. We report this calculation in appendix B.

Log inputs. We now turn our attention to logarithmic inputs $l = -\log \mathcal{J}_i$. In this case we are not able to determine the position of the minimum at the integrand level and we are forced to use the explicit constraints from eq. (3.12). In particular, we want to check if the likelihood condition $a_1 = \dots = a_{n-1} = 0$ is still a solution of the system of this system of $n + 1$ equations. To this purpose, we use eq. (3.12) with the probability distribution given by eq. (2.22), then we set $a_1 = \dots = a_{n-1} = 0$, thus explicitly obtaining

$$\begin{cases} \frac{\partial \tilde{C}}{\partial a_i} = 2^{n-1} \int_0^\infty dl_n l_n I_i(l_n) \left[\frac{\tilde{C}_F^n \exp(-\tilde{C}_F l_n^2)}{1 + \exp(-a_n l_n - b)} - \frac{\tilde{C}_A^n \exp(-\tilde{C}_A l_n^2)}{1 + \exp(a_n l_n + b)} \right] = 0, \\ \frac{\partial \tilde{C}}{\partial b} = 2^{n-1} \int_0^\infty dl_n l_n I_0(l_n) \left[\frac{\tilde{C}_F^n \exp(-\tilde{C}_F l_n^2)}{1 + \exp(-a_n l_n - b)} - \frac{\tilde{C}_A^n \exp(-\tilde{C}_A l_n^2)}{1 + \exp(a_n l_n + b)} \right] = 0, \end{cases} \quad (3.15)$$

where $I_i(l_n)$ is the result of the integration over l_1, \dots, l_{n-1}

$$I_i(l_n) = \int_0^{l_n} dl_{n-1} l_{n-1} \dots \int_0^{l_2} dl_1 l_1 v_i, \quad \text{with } v_0 = 1, v_i = l_i. \quad (3.16)$$

Up to an irrelevant overall multiplicative constant, these integrals are easily found to be


$$I_0(l_n) \propto l_n^{2n-2}, \quad I_i(l_n) \propto l_n^{2n-1} \quad \text{for } i = 1, \dots, n. \quad (3.17)$$

Replacing this result in eq. (3.15), we see that all of the derivatives with respect to a_i , $i = 1, \dots, n$ give rise to the same equation, and thus the system of $n + 1$ simultaneous equations reduces to a system of just two independent equations, for any n . These two equations correspond to two lines in the (a_n, b) plane. If these two lines never cross, then the system has no solution, the minimum of the cost function is not at $a_1 = \dots = a_{n-1} = 0$ and thus the perceptron is not able to correctly reproduce the likelihood. If instead these lines meet at some (\bar{a}_n, \bar{b}) , then the minimum of the cost function does correspond to the likelihood ratio. Despite numerous attempts, we have not been able to perform the final integration over l_n in eq. (3.15) analytically. However, we can perform the integration numerically and plot the result in the (a_n, b) plane. This is done in figure 2a, where, without loss of generality, we have concentrated on the case $n = 2$. It is clear from the plot that the two curves do meet in a point and hence the perceptron is able to find the correct minimum, i.e. the one dictated by the likelihood with $a_1 = 0$.

The explicit n -dependence of the coefficients a_n and b_n can be found numerically by solving the system of equations. Furthermore, it is possible to obtain analytically the scaling of a_n with respect to the colour factors, as detailed in appendix C. We find that, once we have factored out $\tilde{C}_F^{-1/2}$, the resulting coefficient only depends on the ratio of colour factors:

$$\frac{a_n}{\sqrt{\tilde{C}_F}} = F \left(\frac{\tilde{C}_A}{\tilde{C}_F} \right), \quad (3.18)$$

where F is a function that we have not determined.



(a) Log inputs.

(b) Linear inputs.

Figure 2. Solutions of eq. (3.12), in the case $n = 2$ and with $a_1 = 0$, plotted as points in the $(a_2, b/a_2)$ plane, for log (left) and linear (right) inputs.

Linear inputs. We now move to consider linear inputs of the perceptron, i.e. the variables \mathcal{J}_i directly. We follow the same logic as in the case of the logarithmic inputs, namely we want to check whether there is a minimum of the cost function satisfying $a_1 = \dots = a_{n-1} = 0$. Following the same steps as before, we have

$$\begin{cases} \frac{\partial \tilde{C}}{\partial a_i} = 2^{n-1} \int_0^1 \frac{d\mathcal{J}_n}{\mathcal{J}_n} \log \frac{1}{\mathcal{J}_n} I_i(\mathcal{J}_n) \left[\frac{\tilde{C}_F^n \exp(-\tilde{C}_F \log^2 \mathcal{J}_n)}{1 + \exp(-a_n \mathcal{J}_n - b)} - \frac{\tilde{C}_A^n \exp(-\tilde{C}_A \log^2 \mathcal{J}_n)}{1 + \exp(a_n \mathcal{J}_n + b)} \right] = 0, \\ \frac{\partial \tilde{C}}{\partial b} = 2^{n-1} \int_0^1 \frac{d\mathcal{J}_n}{\mathcal{J}_n} \log \frac{1}{\mathcal{J}_n} I_0(\mathcal{J}_n) \left[\frac{\tilde{C}_F^n \exp(-\tilde{C}_F \log^2 \mathcal{J}_n)}{1 + \exp(-a_n \mathcal{J}_n - b)} - \frac{\tilde{C}_A^n \exp(-\tilde{C}_A \log^2 \mathcal{J}_n)}{1 + \exp(a_n \mathcal{J}_n + b)} \right] = 0, \end{cases} \quad (3.19)$$

with

$$I_i(\mathcal{J}_n) = \int_{\mathcal{J}_n}^1 \frac{d\mathcal{J}_{n-1}}{\mathcal{J}_{n-1}} \log \frac{1}{\mathcal{J}_{n-1}} \dots \int_{\mathcal{J}_2}^1 \frac{d\mathcal{J}_1}{\mathcal{J}_1} \log \frac{1}{\mathcal{J}_1} w_i, \quad \text{with } w_0 = 1, w_i = \mathcal{J}_i. \quad (3.20)$$

There is a crucial difference between the integrals in eq. (3.20) and the corresponding ones in the case of the logarithmic inputs, eq. (3.16): the cases with $1 \leq i \leq n$ lead to n different constraints. For sake of simplicity, let us consider $n = 2$. We have

$$I_i(\mathcal{J}_2) = \int_{\mathcal{J}_2}^1 \frac{d\mathcal{J}_1}{\mathcal{J}_1} \log \frac{1}{\mathcal{J}_1} v_i = \begin{cases} \frac{1}{2} \log^2 \mathcal{J}_2, & i = 0, \\ 1 - \mathcal{J}_2 + \mathcal{J}_2 \log \mathcal{J}_2, & i = 1, \\ \frac{1}{2} \mathcal{J}_2 \log^2 \mathcal{J}_2, & i = 2. \end{cases} \quad (3.21)$$

Thus, all the three equations appearing in (3.19), i.e. $i = 0, 1, 2$ provide independent conditions. The system has solutions if the corresponding three curves in the (a_2, b) plane meet in one point. By numerically performing the integrations, we can check whether this happens or not. This is done in figure 2b where it is clear that the three curves do not intersect at a common point and hence the perceptron is unable to find the correct minimum dictated by the likelihood.

Let us summarise the findings of this section. Working in the leading logarithmic approximation of QCD, we have analytically studied the behaviour of a perceptron with sigmoid activation and cross-entropy cost function, in the context of a binary classification problem. We have explicitly considered three variants of the primary N -subjettiness inputs: squared logarithms, logarithms and linear inputs. In the first two cases the minimum of the cost function captures the expected behaviour from the likelihood ratio, i.e. $a_1 = \dots = a_{n-1} = 0$. This does not happen with linear inputs, although the configuration $a_1 = \dots = a_{n-1} = 0$ is within the reach of the network. This is most likely due to the fact that the simple perceptron with linear inputs struggles at correctly learning the probability distributions which are intrinsically logarithmic. We expect that a more complex network would be needed in this case and this is shown explicitly in section 5.

4 Perceptron numerics

In this section we validate our analytic findings with an actual implementation of a perceptron. In practice, we have used a simple implementation based on [59], with a learning rate of 0.1 and a mini-batch size of 32.⁷ Because our first-principle analysis has been developed at LL accuracy, in order to numerically test the perceptron performance we generate a sample of pseudo-data according to the QCD LL distribution for quark and gluon jets. We consider the three different input variants also used in the analytic study, namely square logarithms, logarithms and the linear version of the N -subjettiness inputs. In order to train our single-neuron network we use a sample of 1M events in the first two cases and 16M the for linear inputs, unless otherwise stated. Furthermore, we perform our study as a function of number of N -subjettiness variables. Specifically, when we quote a given value of n , we imply that all \mathcal{T}_i with $i \leq n$ have been used as inputs. The results of this study are collected in figure 3. Each plot shows the value of the network weights a_i and b after training, i.e. at the minimum of the cost function that has been found by the network through back-propagation and gradient descent. The plot on the left is for log-square inputs, the one in the middle for log inputs and the one on the right for linear inputs. The values of the weights determined by the network are shown as circles (for a_i with $i < n$), squared (for a_n) and triangles (for b), with the respective numerical uncertainties. In the case of log-square and log inputs, we also show the behaviour obtained from our analytic calculations in section 3. We find perfect agreement. In particular, the minimum of the cost function found by the network is consistent with $a_1 = \dots = a_{n-1} = 0$. This does not happen in the case of linear inputs, although the discrepancy is tiny.

It is interesting to investigate whether the theoretical issues for the linear inputs have some visible effect on the network performance. In order to do so, we first perform a study of the perceptron convergence as a function of the training sample size N_{train} . As before, we repeat this study for each of the input variants previously introduced. We also consider two different values of n : for the case $n = 3$ we build our inputs from $\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3$, while for $n = 5$, we also include \mathcal{T}_4 and \mathcal{T}_5 . In figure 4 we plot the cost function as a function of the training sample size. Figure 4 is obtained by training the perceptron with a

⁷We have also cross-checked our results against a standard PyTorch implementation.

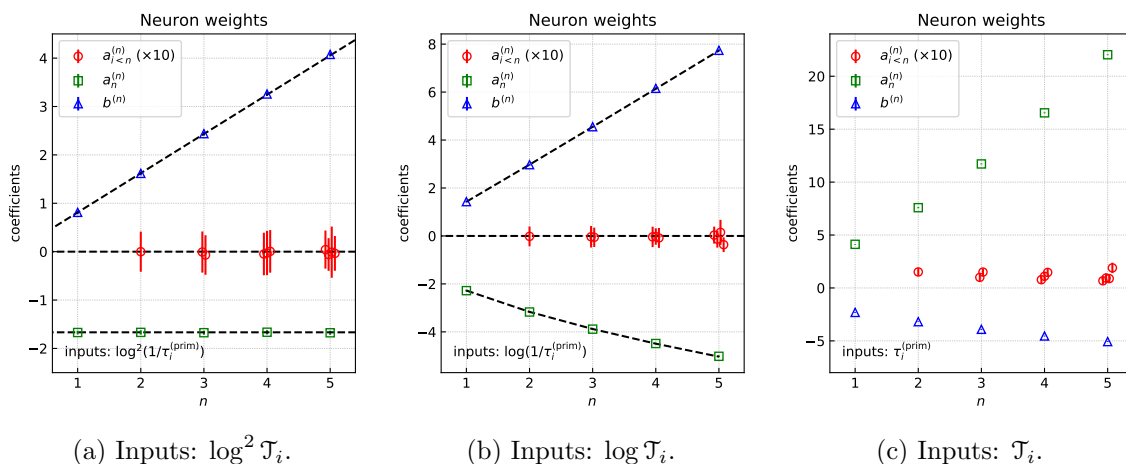


Figure 3. Perceptron parameters after training. When available, expected analytic results are shown as dashed lines.

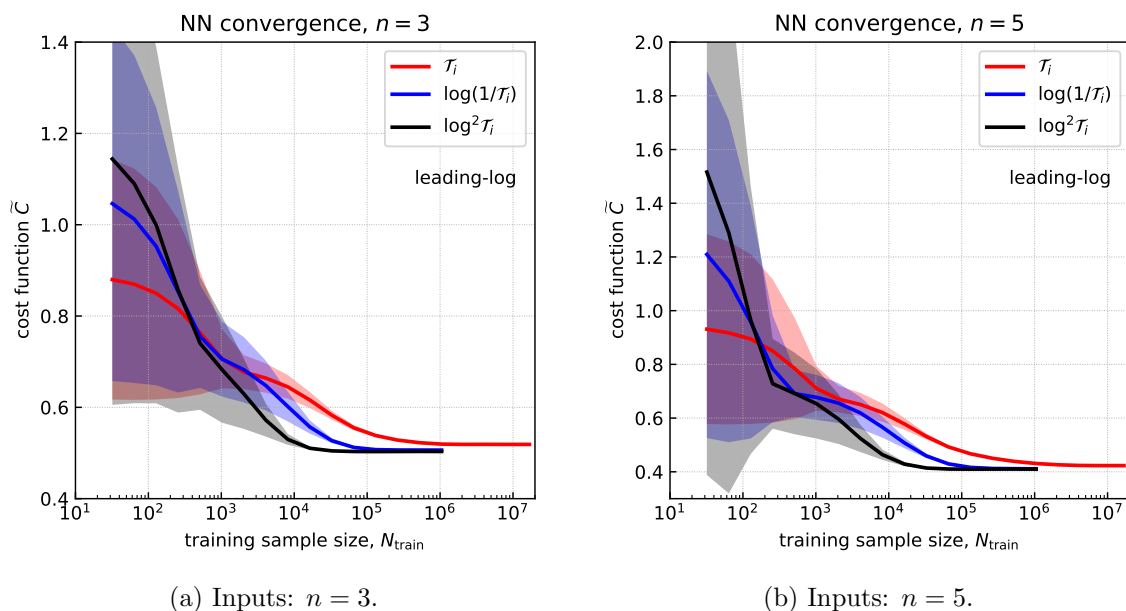


Figure 4. Convergence of the network as a function of the training sample size.

progressively increasing sample of pseudo-data generated on the fly according to the QCD LL distribution. At fixed values of N_{train} , the cost function \tilde{C} of the trained network is then evaluated on a test sample of fixed dimension. This procedure is iterated a number of times, and at the end, for each N_{train} , we take the average and standard deviation of the different runs as a measure of the central value and uncertainty. The plots clearly show that the convergence with linear inputs is slower than in the cases of log-squares and logs, exposing the fact that the single-neuron network struggles to learn intrinsically logarithmic distributions with linear inputs.

Furthermore, with the same set up, we can study the actual performance of the network using ROC curves. To highlight deviations from the ideal case dictated by the likelihood

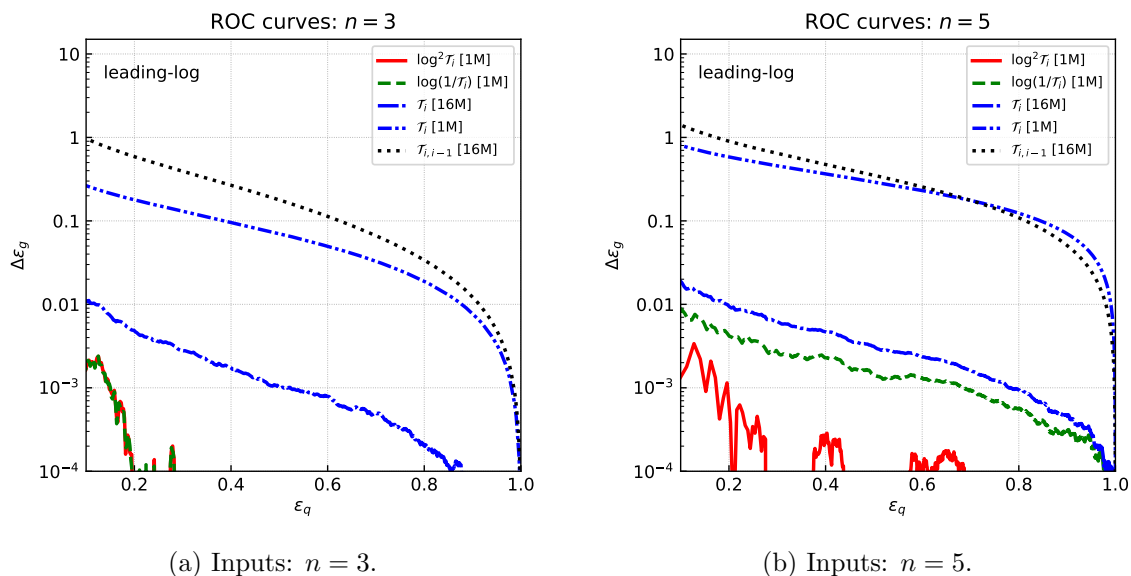


Figure 5. ROC curves obtained from a trained perceptron with different inputs. Note that, in order to highlight deviations from optimal performance, we show eq. (4.1). Thus, the closer these curves are to zero, the better the performance.

ratio, the plots in figure 4 show

$$\Delta\epsilon_g = \frac{\text{ROC}_{NN}}{\text{ROC}_{\text{lik}}} - 1, \quad (4.1)$$

where ROC_{NN} is the network ROC curve, while ROC_{lik} is the ROC curve that corresponds to a cut on \mathcal{T}_n . We perform this study for two different sizes of the training sample: 1 M and 16 M. We know from figure 4 that the former is enough to achieve convergence in the case of log-square and log inputs but not for linear inputs. Thus, we expect to see larger differences in this case. This is indeed confirmed by the plots in figure 5. The dash-dot-dotted curve in blue, corresponding to linear inputs and a training done on a sample of 1 M events, indicates that the difference with respect the optimal case is of order 10% for quark efficiencies $\epsilon_q < 0.6$ in the case of $n = 3$ and exceeds 30% in the same region for $n = 5$. If the training sample size is increased to 16 M, the perceptron with linear inputs still performs worse than the other choices, although the difference after training is now small and never exceeds 1%. For comparison, we also show the results for N -subjettiness ratios as inputs. In this case, for $n = 3$, we include \mathcal{T}_1 , $\mathcal{T}_{21} = \frac{\mathcal{T}_2}{\mathcal{T}_1}$ and $\mathcal{T}_{32} = \frac{\mathcal{T}_3}{\mathcal{T}_2}$. In this case, the expected ideal cut on \mathcal{T}_n cannot be reproduced by any choice of the perceptron weights a_i and bias b . The perceptron performance in this case is consequently much worse even after a training on our largest sample.

5 Monte Carlo studies

In the previous section, we have validated the results of section 3, using the same setting as for our analytical calculations, namely a single-neuron network fed with primary N -subjettiness variables (or functions thereof) distributed according to QCD predictions at

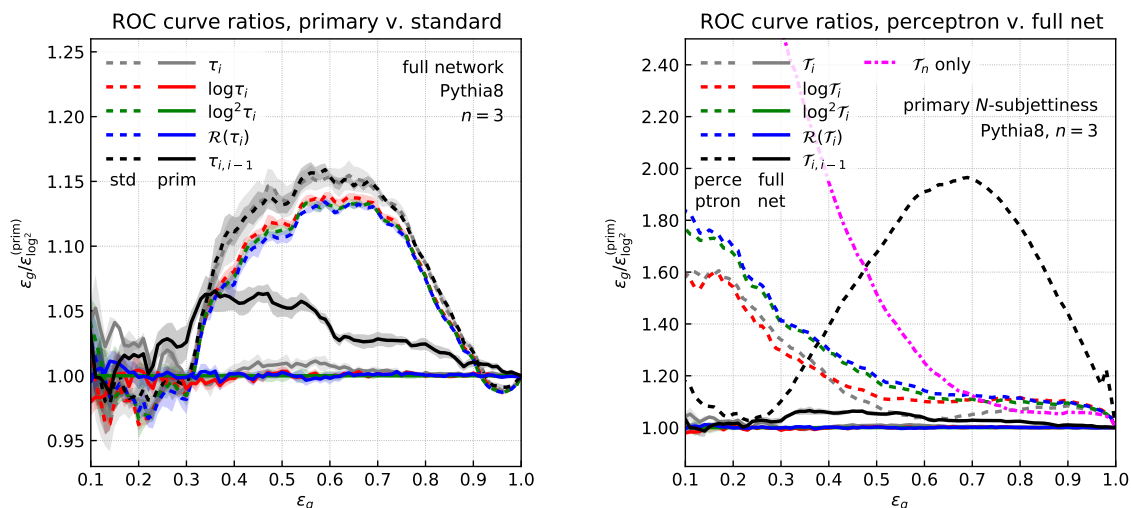
leading-logarithmic accuracy. This setup is somehow oversimplified, and one may wonder how our results compare in terms of performance and convergence to a fully-fledged neural network trained on Monte Carlo pseudo-data. In addition, we have adopted the primary N -subjettiness definition, whose nice analytical properties naturally justify its use in our theoretical studies. However, this alternative definition so far has been compared to the standard definition only in terms of AUC at LL accuracy. Even though the purpose of this work is not a detailed comparison of the two definitions, we need to make sure that primary N -subjettiness performs sensibly as a quark/gluon tagger.

The purpose of this section is to address some of these concerns, by extending the setup of section 4 to a more realistic scenario, both in terms of the generation of pseudo-data and of the network architecture employed in the analysis. First, we generate pseudo-data with Pythia 8.230 [60], including hadronisation and multi-parton interactions. We simulate dijet events $qq \rightarrow qq$ and $gg \rightarrow gg$ and we cluster jets with the anti- k_t algorithm [61], as implemented in FastJet [62], with jet radius $R_0 = 0.4$. We keep jets with $p_t > 2$ TeV and $|y| < 4$. We then compute τ_N and \mathcal{T}_N (i.e. both N -subjettiness definitions) up to the desired $N = n$ for each jet. We set the N -subjettiness parameter β to 1. For the standard N -subjettiness, we use the reference axes obtained from the exclusive k_t axis with winner-takes-all [63] recombination and a one-pass minimisation [49]. In addition to linear, log-square and log functional forms and N -subjettiness ratios already adopted in the previous section, we also consider a running-coupling-like input, obtained by evaluating the radiator $\mathcal{R}(\tau_i)$ of eq. (2.11) with a one-loop running coupling.⁸

The samples thus obtained are used to train either a single perceptron (as in the previous section) or a full neural network. The perceptron is trained using the same implementation based on ref. [59], training over 50 epochs, with a learning rate of 0.5. The *full network* uses a fully connected architecture implemented in PyTorch 1.4.0. The number of inputs equals the number n of input N -subjettiness values (between 2 and 7). There are three hidden layers with 128 nodes per layer and a LeakyReLU [64] activation function with negative slope of 0.01 is applied following each of these layers. We train the network using binary cross-entropy for a maximum of 50 epochs using the Adam [65] optimiser with a learning rate of 0.005 and a mini-batch size of 32 events. During training, the area under curve is evaluated after each epoch on a statistically independent validation set and training terminates if there is no improvement for 5 consecutive epochs [66]. The best model on the validation set is then selected and evaluated on another independent test set for which results are reported. Data are split 6 : 2 : 2 between training, validation, and testing. No extensive optimisation of hyperparameters has been performed and training usually terminates due to the early stopping criterion before reaching the maximally possible number of epochs.

We start with a comparison of the primary and the standard definition of N -subjettiness in terms of network performance. In figure 6a we plot the ROC curves for the primary (solid lines) and standard (dashed line) definitions of N -subjettiness, and for

⁸We set $\alpha_s(p_t R_0) = 0.09$, and we regulate the divergence at the Landau pole by freezing the coupling at a scale k_t such that $2\alpha_s(p_t R)\beta_0 \log(p_t R_0/k_t) = 0.99$.



(a) Performance comparison between the standard N -subjettiness (dashed lines) and the primary N -subjettiness (solid lines). All results are obtained using a full neural network.

(b) Performance comparison between a single perceptron (dashed lines) and a full neural network (solid lines). All results are obtained using primary N -subjettiness.

Figure 6. ROC curves obtained after training on a Pythia8 sample, using the N -subjettiness variables, with $n = 3$. To better highlight differences the vertical axes show the gluon rate normalised by what is obtained using $\log^2 \mathcal{T}_i$ inputs for primary N -subjettiness trained on a full neural network. The different colours correspond to different inputs to the neural network.

different choices of the input functional form (different colours). The gluon efficiencies are normalised by the central value obtained with primary N -subjettiness with $\log^2(\mathcal{T}_i)$ inputs trained on a full NN. The central values and the error bands are calculated by taking the average and standard deviation of five different runs with randomly initialised weights.

We see that the performance of networks trained with the standard definition is worse by 10–15% compared to the primary definition for mid-values of the quark efficiency, $0.4 \lesssim \varepsilon_q \lesssim 0.8$ and comparable elsewhere, except for a small region at large quark efficiency, $\varepsilon_q \gtrsim 0.9$. Even though the benefit of using primary N -subjettiness is not as large as one could have expected based on our leading-logarithmic calculations, we still observe a performance improvement. We note also that at very small quark efficiency non-perturbative effects have a sizeable effect, invalidating our arguments purely based on a perturbative calculation. Furthermore, large ε_q correspond to the regime where the cut on N -subjettiness is no longer small and our arguments based on the resummation of large logarithms no longer apply, and one should instead consider the full structure of hard matrix elements.⁹ That said, figure 6a shows that for the most of the range in quark efficiency, a better discrimination is reached if we adopt the primary- N -subjettiness definition.

It is also interesting to compare the results obtained with different inputs in figure 6a. Although one should expect that a full NN should ultimately reach the performance of

⁹Note that in that region, our Monte-Carlo pseudo-data, based on Pythia8, does not contain all the details of the physics in any case.


the likelihood independently on the choice of inputs, our practical setup still shows a mild dependence on the choice of inputs. The key point is that the favoured inputs agree with our analytic expectations from section 3 with logarithmic inputs ($\log \tau_i$, $\log^2 \tau_i$ and $\mathcal{R}(\tau_i)$) showing an equally-good performance, the linear inputs only marginally worse and the ratio inputs showing a (few percent) worse discriminating power.¹⁰ Even though the convergence of the neural network is more delicate for the case of ratio inputs (and, to a lesser extent, for linear inputs), a performance similar to the optimal one reported in figure 6a for logarithmic inputs could be achieved with a careful optimisation of the hyperparameters.

We now move to figure 6b, where we compare the ROC curves obtained with a full neural network to those of a simple perceptron. In this plot we select the primary N -subjettiness definition and we display results for the usual input choices. The solid lines in figure 6b coincide with the ones in figure 6a. First, we observe that a perceptron trained with Monte Carlo pseudo-data performs worse compared to the full NN for all the considered input types. This is not surprising as the arguments in section 3 are based on a simplified, leading-logarithmic, approach and subleading corrections can be expected to come with additional complexity that a single perceptron would fail to capture. It is actually remarkable that for $\varepsilon_q \gtrsim 0.6$ a single perceptron typically gives performances which are only 10% worse than the full network. This is not true if we consider N -subjettiness ratio inputs, in which case the perceptron performance is sizeably worse than the performance we get by using the full NN. This is in agreement with the observation made towards the end of section 4, namely that a more complex network architecture is able to learn the correct weights for the ratio inputs, while the perceptron is not.

The behaviour of the other input types is relatively similar, although, quite surprisingly, linear inputs tend to give a slightly better performance than logarithmic ones. The origin of this is not totally clear. One can argue that the difference in performance between linear and logarithmic inputs was small at leading-logarithmic accuracy and that several effects — subleading logarithmic effects, hard matrix-element corrections, non-perturbative effects, . . . — can lift this approximate degeneracy. One should also keep in mind that our leading-logarithmic calculation is valid at small-enough quark/gluon efficiencies, while, as the efficiency grows, we become sensitive to contributions from other kinematic domains. We also note that if we increase the jet p_t , e.g. if we look at 20 TeV jets at 100 TeV pp collider, the difference between logarithmic and linear inputs becomes again comparable, which is likely a sign that the phase-space at larger p_t has an increased contribution from the resummation region. We finally note that a simple cut on \mathcal{T}_n , which is optimal at LL, gives comparable results to the perceptron at large quark efficiencies $\varepsilon_q \gtrsim 0.7$, but performs rather worse otherwise.

Next, we study how the convergence of the network is affected by the choice of inputs. In this context, convergence can mean two different things. Firstly, we can speak about convergence of the network training (as we did in figure 4). This is shown in figure 7a, for the area under the ROC curve (AUC) as a function of the number of training epochs for a single perceptron. For this plot, we have used primary N -subjettiness and (mini-

¹⁰The choice of logarithmic inputs was also found to be helpful, on a phenomenological basis, in ref. [41].



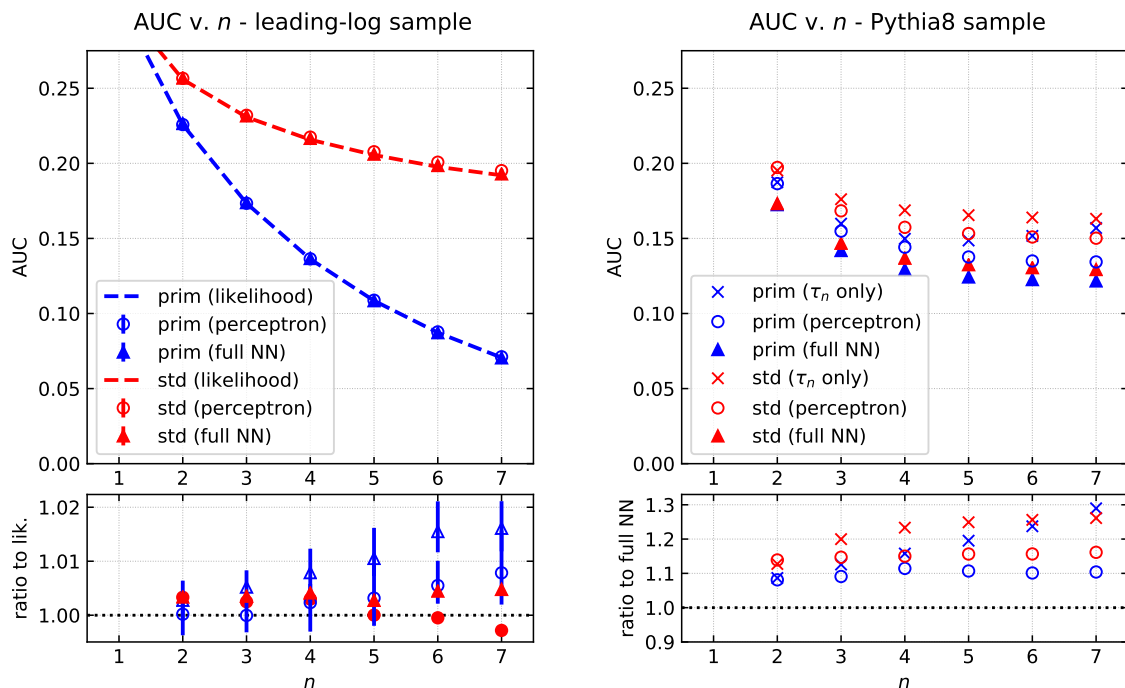
(a) AUC as a function of the number of epochs for a perceptron. (b) AUC as a function of the number of neurons in a single hidden layer. (c) AUC as a function of the number of hidden layers (with 128 neurons per layer).

Figure 7. AUC convergence study for the Pythia8 sample as a function of the network parameters: number of training epochs, network width and network depth. The different inputs are represented with the same colours as in figure 6.

batch) stochastic gradient descent with a learning rate $\lambda = 0.5$. The better performance obtained for linear inputs is confirmed in figure 7a. However, this happens at a cost in convergence speed: where logarithmic inputs show an almost complete convergence after a single epoch, linear inputs converge much slower.¹¹ We also see that using ratios of N -subjettiness as inputs yields a clearly worse AUC. Secondly, we can study convergence as a function of the size of the network which can be either the network width at fixed number of hidden layers, figure 7b, or the network depth at fixed number of neurons per hidden layer, figure 7c. These plots also show that logarithmic inputs converge (marginally) faster than linear inputs. Interestingly, after a full training (with 7 hidden layers and 128 neurons per hidden layer), we recover a situation where logarithmic inputs give a slightly better performance compared to linear inputs, as our analytic studies suggested.

As a final test, we study in figure 8 the AUC as a function of the number n of input N -subjettiness variables. Results are shown for inputs distributed according to the leading-logarithmic QCD distributions (as for our analytic calculations in section 3 and for our tests in section 4) in figure 8a as well as for a Pythia8 sample in figure 8b. We show the AUC for both the standard definition of N -subjettiness (red) and for our primary definition (blue), and in both cases we show the results of training either a single perceptron (open symbols) or a full NN (filled symbols). In the case of the Pythia8 sample, we also show the AUC corresponding to a simple cut on either \mathcal{T}_n or τ_n . For the case of the leading-logarithmic distributions, we also provide the expectations from the actual likelihood ratio (dashed lines). The latter is obtained from eq. (2.17) for the case of primary N -subjettiness and by numerically integrating the equivalent of eq. (2.16) for the standard definition. The lower panel shows the relative performance, normalised to the likelihood, if it is available, or to the full NN otherwise. These results confirm once more our

¹¹For lower learning rates, the fast convergence of logarithmic inputs is still observed, but linear inputs converge even slower.



(a) AUC as a function of the number of inputs n for a leading-logarithmic distribution. (b) AUC as a function of the number of inputs n for a Pythia8 distribution.

Figure 8. AUC as a function of the number of inputs n . We show results for both the standard (red) and primary (blue) definitions of N -subjettiness with training done either for a single perceptron (open circles) or for a full NN (filled triangles). In all cases, log-square inputs have been used. When available, results for the likelihood have been included (dashed lines) and used to normalise the results in the lower panel. When not available, the normalisation is to the full NN results. For the Pythia8 sample we also show the results of a cut on τ_n (or \mathcal{T}_n) only (crosses).

earlier findings. First, primary N -subjettiness performs better than the standard definition, although the difference, clearly visible in the leading-log distributions, is only marginal with the Pythia8 sample. Then, in the case of a leading-log distribution, the optimal performance of the primary N -subjettiness is already captured by a single perceptron, while for the case of standard N -subjettiness only the full NN is able to reach a performance on par with the likelihood expectation. For the Pythia8 sample, the training of the full NN leads a reduction of the AUC compared to a single perceptron for both N -subjettiness definitions. The improvement is however smaller for the primary definition ($\sim 10\%$) than for the standard one ($\sim 15\%$). This is likely a reminiscence of the behaviour observed at leading logarithmic accuracy. It is interesting to observe that the performance of a simple cut on \mathcal{T}_n , which is optimal at LL, is comparable to the perceptron at small n , while it degrades as n increases. This is most likely due to the fact that at large n we become more sensitive to non-perturbative corrections and consequently our LL approximation is no longer sufficient. Conversely, a cut on τ_n , while giving a performance comparable to that of the perceptron for $n = 2$, degrades more rapidly at larger n .

6 Conclusions and outlook

In this paper we have investigated the behaviour of a simple neural network made of just one neuron, in the context of a binary classification problem. The novelty of our study consists in exploiting the underlying theoretical understanding of the problem, in order to obtain firm analytical predictions about the perceptron performance.

In particular, we have addressed the question of distinguishing quark-initiated jets from gluon-initiated ones, exploiting a new variant of the N -subjettiness (IRC-safe) family of observables that we dubbed primary N -subjettiness. As the name suggests, these observables are only sensitive, to LL accuracy, to emissions of soft gluons off the original hard parton, while standard N -subjettiness also depends on secondary gluon splittings, already at LL. Thanks to the simple all-order behaviour of observables $\{\mathcal{T}_1 \dots \mathcal{T}_n\}$ we have been able to determine that the optimal discriminant at LL, i.e. one which is monotonically related to the likelihood ratio, is just a cut on \mathcal{T}_n . We have also been able to obtain analytic expressions for the area under the ROC curve and the area under the ROC curve which are standard figures of merit for a classifier's performance. Furthermore, we have found that, besides having nicer theoretical properties, \mathcal{T}_N typically outperforms τ_N as a quark/gluon tagger, in the range of quark efficiencies that is typically employed in phenomenological analyses.

The central part of this study has been the analytic study of the LL behaviour of a perceptron that takes primary N -subjettiness variables as inputs. We have considered a perceptron with a sigmoid as activation function and the cross-entropy as loss function. The main question we have raised is whether the values of the perceptron weights at the minimum of the cost function correspond to the aforementioned optimal LL classifier, i.e. a cut on the last primary N -subjettiness considered. We have been able to show analytically that this depends on the actual functional form of the inputs fed to the network. The perceptron is able to find the correct minimum if logarithms (or square logarithms) of the N -subjettiness are passed, but fails to do so with linear inputs. This reflects the fact that the LL distributions of the inputs are indeed logarithmic. These analytic results are fully confirmed by a numerical implementation of the perceptron, which has been trained on pseudo-data generated according to a LL distribution. Furthermore, we have also found that, in the case of linear inputs, the learning rate of the perceptron is substantially slower than for logarithmic inputs. As a by-product, we were also able to find, in a few cases, closed analytic expressions for the network cost function.

Finally, we have considered a more realistic framework for our analysis. Firstly, we have trained the perceptron pseudo-data generated with a general-purpose Monte Carlo parton shower, and we have obtained qualitative agreement with the analytic study. Secondly, we have observed that when considering a full neutral network, we obtain a comparable performance with all the variants of the inputs, with the possible exception of the N -subjettiness ratios, which requires more fine-tuning to be brought to the same performance.

Our work highlights in a quantitative way the positive role that first-principle understanding of the underlying physical phenomena has in classification problems that employs ML techniques. Even if a similar degree of performance is ultimately achieved with a full

NN regardless of the type of inputs, a knowledge of the underlying physics is helpful to build a simpler network less dependent on a careful optimisation of the hyperparameters. Furthermore, we stress that the use of theoretically well-behaved observables has allowed us to perform our analysis at a well-defined, and in principle improvable, accuracy. We believe that this is an important step towards the consistent use of ML in experimental measurements that are ultimately aimed at a detailed comparison with theoretical predictions.

Our current analysis suffers from the obvious limitations of being conducted at LL and for just a very simple network. We believe that going beyond LL accuracy for primary N -subjettiness is not only interesting, but desirable. The first reason would be to see if the simple resummation structure we see at LL accuracy survives at higher orders. Then, this very first analysis places primary N -subjettiness as a promising substructure observable. One is therefore tempted to also investigate its performance in the context of boosted vector boson or top tagging, although one might expect that these classification problems benefit from constraints beyond primary emissions. Furthermore, in the framework of our analytic study of ML techniques, it would be interesting to see if at a higher accuracy one is still able to find a simple combination of architecture, cost function and network inputs that guarantees, analytically, optimal performance. More generally, we look forward to future work where perturbative QCD helps designing better ML techniques with a comfortable degree of analytic control.

Acknowledgments

We thank the organisers of the BOOST 2019 workshop at MIT for a vibrant scientific environment that sparked the discussion about this project. We are all thankful to Matteo Cacciari for multiple interactions and interesting discussions. We thank Simone Caletti, Andrea Coccaro, Andrew Larkoski, Ben Nachman, and Jesse Thaler for helpful comments on the manuscript. SM and GiS wish to thank IPhT Saclay for hospitality during the course of this work. The work of SM is supported by Università di Genova under the curiosity-driven grant “Using jets to challenge the Standard Model of particle physics” and by the Italian Ministry of Research (MIUR) under grant PRIN 20172LNEEZ. GK acknowledges support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy — EXC 2121 “Quantum Universe” — 390833306. GrS’s work has been supported in part by the French Agence Nationale de la Recherche, under grant ANR-15-CE31-0016.

A Details of the analytic calculations

In this appendix we report details of the analytic calculation of the area under the ROC curve (AUC). We start from the generic definition of the AUC and we exploit the definition of the ROC curve given in eq. (2.5):

$$\text{AUC} = \int_0^1 dx \text{ROC}(x) = \int_0^1 dx \Sigma_{\mathcal{B}}(\Sigma_{\mathcal{S}}^{-1}(x)), \quad (\text{A.1})$$

where, for a given observable V , the cumulative distribution for signal or background as a function of a cut v reads as in eq. (2.4), which we rewrite as:

$$\Sigma_i(v) = \int dt p_i(\mathbf{t}) \Theta(V(\mathbf{t}) < v), \quad i = \mathcal{S}, \mathcal{B}. \quad (\text{A.2})$$

We now perform a change of variable from the efficiency x to the cut on the observable v , which results in the following Jacobian factor:

$$\frac{dx}{dv} = \frac{dx}{d\Sigma_{\mathcal{S}}^{-1}(x)} = \frac{d\Sigma_{\mathcal{S}}(v)}{dv}. \quad (\text{A.3})$$

Thus, we have

$$\text{AUC} = \int_0^1 dv \Sigma_{\mathcal{B}}(v) \frac{d\Sigma_{\mathcal{S}}(v)}{dv} \quad (\text{A.4})$$

$$= \int_0^1 dv \int dt_{\mathcal{B}} p_{\mathcal{B}}(\mathbf{t}_{\mathcal{B}}) \Theta(V(\mathbf{t}_{\mathcal{B}}) < v) \int dt_{\mathcal{S}} p_{\mathcal{S}}(\mathbf{t}_{\mathcal{S}}) \delta(V(\mathbf{t}_{\mathcal{S}}) - v) \quad (\text{A.5})$$

$$= \int dt_{\mathcal{B}} \int dt_{\mathcal{S}} p_{\mathcal{B}}(\mathbf{t}_{\mathcal{B}}) p_{\mathcal{S}}(\mathbf{t}_{\mathcal{S}}) \Theta(V(\mathbf{t}_{\mathcal{B}}) < v(\mathbf{t}_{\mathcal{S}})), \quad (\text{A.6})$$

which is the expression presented in eq. (2.16).

We now specialise to our case and we consider the probability distributions p_q and p_g for primary N -subjettiness at LL accuracy, which are given in eq. (2.10). We first integrate over quark variables and exploiting the result found in eq. (2.15), we obtain

$$\int_0^1 d\mathcal{T}_{1q} \cdots \int_0^{\mathcal{T}_{n-1,q}} d\mathcal{T}_{nq} p_q(\mathcal{T}_{1q}, \mathcal{T}_{2q}, \cdots, \mathcal{T}_{nq}) \Theta(\mathcal{T}_{nq} > \mathcal{T}_{ng}) = 1 - \frac{\Gamma(n, C_F \mathcal{R}(\mathcal{T}_{ng}))}{\Gamma(n)}. \quad (\text{A.7})$$

The integral over \mathcal{T}_{ng} can be performed using the following integration-by-parts identity:

$$\begin{aligned} & \int_0^{\mathcal{J}} d\mathcal{T}' C_A \frac{\mathcal{R}'(\mathcal{T}')}{\mathcal{T}'} e^{-C_A \mathcal{R}(\mathcal{T}')} \Gamma(n, C_F \mathcal{R}(\mathcal{T}')) \\ &= e^{-C_A \mathcal{R}(\mathcal{J})} \Gamma(n, C_F \mathcal{R}(\mathcal{J})) - \frac{C_F^n}{(C_F + C_A)^n} \Gamma(n, (C_F + C_A) \mathcal{R}(\mathcal{J})). \end{aligned} \quad (\text{A.8})$$

We see that iteratively the following kind of integrals appear:

$$\begin{aligned} & \int_0^{\mathcal{J}} d\mathcal{T}' C_A \frac{\mathcal{R}'(\mathcal{T}')}{\mathcal{T}'} \mathcal{R}(\mathcal{T}')^m \Gamma(n, (C_F + C_A) \mathcal{R}(\mathcal{T}')) \\ &= \frac{C_A}{1+m} \left(\frac{\Gamma(n+1+m, (C_F + C_A) \mathcal{R}(\mathcal{J}))}{(C_F + C_A)^{1+m}} - \mathcal{R}(\mathcal{J})^{1+m} \Gamma(n, (C_F + C_A) \mathcal{R}(\mathcal{J})) \right). \end{aligned} \quad (\text{A.9})$$

In the end we obtain:

$$\begin{aligned} & \int_0^1 d\mathcal{T}_{1g} \cdots \int_0^{\mathcal{T}_{n-1,g}} d\mathcal{T}_{ng} p_g(\mathcal{T}_{1g}, \mathcal{T}_{2g}, \cdots, \mathcal{T}_{ng}) \Gamma(n, C_F \mathcal{R}(\mathcal{T}_{ng})) \\ &= \Gamma(n) - \frac{C_F^n}{(C_F + C_A)^{2n-1}} \sum_{k=0}^{n-1} \frac{\Gamma(n+k)}{\Gamma(k+1)} C_A^k (C_F + C_A)^{n-1-k}. \end{aligned} \quad (\text{A.10})$$

The sum can be performed explicitly, leading to the result presented in eq. (2.17).

B Analytic results for the cost function

In section 3 we have looked for a minimum of the cost function, eq. (3.7), by computing from the very beginning the derivatives with respect to the NN weights. This procedure allowed us to determine analytically the position of the minimum for the case of log-square inputs. In the case of logarithmic and linear inputs, although we have not been able to perform all the integrals analytically, we have nonetheless reduced the problem of the determination of the position of the minimum of the cost function to a one-dimensional integration, which we have performed numerically.

An alternative approach would be to determine an analytic expression for the cost function, as a function of the NN weights and bias, which we would have to, in turn, minimise. As we shall see in this appendix, this approach is not as successful as the one presented in the main text. The integrals that appear in the determination of the cost function eq. (3.7) are rather challenging and we have been able to solve them only in a couple of simple cases. Namely, we limit ourselves to the case $n = 2$ with log-square inputs, by using the sigmoid as activation function, and we report results for the cross-entropy loss, eq. (3.4), and the quadratic loss, eq. (3.5). Even if the NN setup is minimal, the derivation of explicit expressions for the cost function may be instructive as they represent a first-principle determination of a NN behaviour and they could be valuable in the context of comparisons between experimental measurements that make use of machine-learning algorithms and theoretical predictions.

B.1 Cross-entropy loss

We start by considering the cost function eq. (3.7) with the cross-entropy loss. Explicitly:

$$\begin{aligned} \tilde{C}^{(\text{XE})}(a_1, a_2, b) = & \frac{1}{2} \int_0^\infty dL_2 \int_0^{L_2} dL_1 \left[\tilde{C}_F^2 e^{-\tilde{C}_F L_2} \left(-\log \left(\frac{1}{1 + e^{a_1 L_1 + a_2 L_2 + b}} \right) \right) \right. \\ & \left. + \tilde{C}_A^2 e^{-\tilde{C}_A L_2} \left(-\log \left(\frac{1}{1 + e^{-a_1 L_1 - a_2 L_2 - b}} \right) \right) \right]. \end{aligned} \tag{B.1}$$

Note that the replacement $\tilde{C}_F \leftrightarrow \tilde{C}_A$ is equivalent to $(a_1, a_2, b) \leftrightarrow (-a_1, -a_2, -b)$, due to the symmetries of the functions involved. We first observe that the integral over L_1 gives rise to a dilogarithm. In order to evaluate the integral over L_2 we make use of the following result:

$$\begin{aligned} \mathcal{I}^{(\text{XE})}(C_R, c, d) = & C_R^2 \int_0^\infty dx e^{-C_R x} \text{Li}_2 \left(-e^{cx+d} \right) \tag{B.2} \\ = & C_R \text{Li}_2 \left(-e^d \right) + \begin{cases} -c \log(1 + e^d) - \frac{c^2}{C_R} \left(1 - {}_2F_1 \left(1, -\frac{C_R}{c}, 1 - \frac{C_R}{c}; -e^d \right) \right) & (c < 0) \\ 0 & (c = 0) \\ -c \log(1 + e^d) - \frac{c^2}{C_R} {}_2F_1 \left(1, \frac{C_R}{c}, 1 + \frac{C_R}{c}; -e^{-d} \right) & (c > 0) \end{cases} \end{aligned}$$

The analytic expression that we obtain for the cost function is

$$\begin{aligned} \tilde{C}^{(\text{XE})}(a_1, a_2, b) = \frac{1}{2a_1} & \left[-\mathcal{I}^{(\text{XE})}(\tilde{C}_F, a_1 + a_2, b) - \mathcal{I}^{(\text{XE})}(\tilde{C}_F, a_2, b) \right. \\ & \left. + \mathcal{I}^{(\text{XE})}(\tilde{C}_A, -(a_1 + a_2), -b) - \mathcal{I}^{(\text{XE})}(\tilde{C}_A, -a_2, -b) \right]. \end{aligned} \quad (\text{B.3})$$

From section 3, we already know that the minimum is located at $a_1 = 0$. However, given the structure of the explicit result after integration, eq. (B.3), it is highly nontrivial to recover the position of the minimum analytically, due to the presence of a_1 both at denominator and in the arguments of hypergeometric function.

B.2 Quadratic loss

An analogous calculation can be performed in the case of the quadratic loss eq. (3.5). We have to calculate

$$\tilde{C}^{(x^2)}(a_1, a_2, b) = \frac{1}{2} \int_0^\infty dL_2 \int_0^{L_2} dL_1 \left[\frac{\tilde{C}_F^2 e^{-\tilde{C}_F L_2}}{(1 + e^{-a_1 L_1 - a_2 L_2 - b})^2} + \frac{\tilde{C}_A^2 e^{-\tilde{C}_A L_2}}{(1 + e^{a_1 L_1 + a_2 L_2 + b})^2} \right]. \quad (\text{B.4})$$

As in the case of the cross-entropy loss, the integral over L_1 is straightforward. We then make use of the following identity to perform the remaining integral over L_2 :

$$\begin{aligned} \mathcal{I}^{(x^2)}(C_R, c, d) &= C_R^2 \int_0^\infty dx e^{-C_R x} \left[\log(1 + e^{cx+d}) + \frac{1}{1 + e^{cx+d}} \right] \\ &= C_R \log(1 + e^d) + \begin{cases} c + (C_R - c) {}_2F_1\left(1, -\frac{C_R}{c}, 1 - \frac{C_R}{c}; -e^d\right) & (c < 0) \\ \frac{C_R}{1 + e^d} & (c = 0) \\ C_R + (c - C_R) {}_2F_1\left(1, \frac{C_R}{c}, 1 + \frac{C_R}{c}; -e^{-d}\right) & (c > 0). \end{cases} \end{aligned} \quad (\text{B.5})$$

We find

$$\begin{aligned} \tilde{C}^{(x^2)}(a_1, a_2, b) &= \frac{1}{2a_1} \left[\left(\mathcal{I}^{(x^2)}(\tilde{C}_F, a_1 + a_2, b) - \mathcal{I}^{(x^2)}(\tilde{C}_F, a_2, b) \right) \right. \\ & \quad \left. - \left(\mathcal{I}^{(x^2)}(\tilde{C}_A, -(a_1 + a_2), -b) - \mathcal{I}^{(x^2)}(\tilde{C}_A, -a_2, -b) \right) \right]. \end{aligned} \quad (\text{B.6})$$

As for eq. (B.3), the position of the minimum is hard to find analytically. To check whether we recover the result obtained in eq. (3.14), we can numerically look for the minimum of eqs. (B.3) and (B.6). For illustration purposes, we fix \tilde{C}_A and \tilde{C}_F to C_A and C_F respectively. In figure 9 we plot each cost function around the found minimum point. We see that the condition $a_1 = 0$ and the (negative) value of $a_2 = C_F - C_A \simeq -1.67$ are indeed confirmed.

C Scaling of a_n in the log inputs case

In this appendix we would like to study the scaling of a_n with respect to \tilde{C}_F and \tilde{C}_A . In order to simplify the notation, in this appendix we will rename $x \equiv l_n$ and $a \equiv a_n$.




Figure 9. Cost function as a function of a_1 and a_2 around the minimum point. b has been accordingly fixed to the value in eq. (3.14).

We first expand in series around $x = 0$ the sigmoid functions which appear in eq. (3.15):

$$\frac{1}{1 + e^{-ax-b}} = \frac{e^b}{1 + e^b} + \sum_{k=1}^{\infty} a^k f^{(k)}(e^b) x^k, \quad (\text{C.1})$$

$$-\frac{1}{1 + e^{ax+b}} = -\frac{1}{1 + e^b} + \sum_{k=1}^{\infty} a^k f^{(k)}(e^b) x^k, \quad (\text{C.2})$$

where:

$$f^{(k)}(e^b) = \frac{(-1)^{1+k} e^b}{\Gamma(1+k)(1+e^b)^{1+k}} \times \text{polynomial in } e^b \text{ with degree of } k-1. \quad (\text{C.3})$$

Note that the two series with index k appearing on the r.h.s. of eqs. (C.1)–(C.2) are the same. By substituting eqs. (C.1)–(C.2) under integration, eq. (3.15) becomes:

$$\int_0^{\infty} dx x^m \left[\tilde{C}_F^n e^{-\tilde{C}_F x^2} \left(\frac{e^b}{1 + e^b} + \sum_{k=1}^{\infty} a^k f^{(k)}(e^b) x^k \right) + \tilde{C}_A^m e^{-\tilde{C}_A x^2} \left(-\frac{1}{1 + e^b} + \sum_{k=1}^{\infty} a^k f^{(k)}(e^b) x^k \right) \right] = 0 \quad (\text{C.4})$$

with $m = 2n$ or $m = 2n - 1$. Given the following integration identity:

$$\int_0^{\infty} dx x^m e^{-cx^2} = \frac{1}{2\sqrt{c^{m+1}}} \Gamma\left(\frac{m+1}{2}\right), \quad c > 0 \quad (\text{C.5})$$

we obtain:

$$\frac{\Gamma\left(\frac{m+1}{2}\right)}{2(1 + e^b)} \left(\frac{\tilde{C}_F^n e^b}{\sqrt{\tilde{C}_F^{m+1}}} - \frac{\tilde{C}_A^n}{\sqrt{\tilde{C}_A^{m+1}}} \right) + \sum_{k=1}^{\infty} a^k f^{(k)}(e^b) \frac{\Gamma\left(\frac{m+1+k}{2}\right)}{2} \left(\frac{\tilde{C}_F^m}{\sqrt{\tilde{C}_F^{m+1+k}}} + \frac{\tilde{C}_A^m}{\sqrt{\tilde{C}_A^{m+1+k}}} \right) = 0. \quad (\text{C.6})$$

We now divide by $\sqrt{\tilde{C}_F^{2n-m-1}}$ and collect $1/\sqrt{\tilde{C}_F^k}$ in the second bracket:

$$\frac{\Gamma\left(\frac{m+1}{2}\right)}{2(1+e^b)} \left(e^b - \left(\frac{\tilde{C}_A}{\tilde{C}_F} \right)^{(2n-m-1)/2} \right) + \sum_{k=1}^{\infty} \left(\frac{a}{\sqrt{\tilde{C}_F}} \right)^k f^{(k)}(e^b) \frac{\Gamma\left(\frac{m+1+k}{2}\right)}{2} \left(1 + \left(\frac{\tilde{C}_A}{\tilde{C}_F} \right)^{(2n-m-1-k)/2} \right) = 0. \quad (\text{C.7})$$

This equation suggests that:

$$\frac{a_n}{\sqrt{\tilde{C}_F}} = F \left(\frac{\tilde{C}_A}{\tilde{C}_F} \right). \quad (\text{C.8})$$

which is the relation quoted in the main text.

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.

References

- [1] S. Marzani, G. Soyez and M. Spannowsky, *Looking inside jets: an introduction to jet substructure and boosted-object phenomenology*, in *Lecture Notes in Physics* **958**, Springer (2019) [[arXiv:1901.10342](https://arxiv.org/abs/1901.10342)] [[INSPIRE](#)].
- [2] A. Abdesselam et al., *Boosted Objects: A Probe of Beyond the Standard Model Physics*, *Eur. Phys. J. C* **71** (2011) 1661 [[arXiv:1012.5412](https://arxiv.org/abs/1012.5412)] [[INSPIRE](#)].
- [3] A. Altheimer et al., *Jet Substructure at the Tevatron and LHC: New results, new tools, new benchmarks*, *J. Phys. G* **39** (2012) 063001 [[arXiv:1201.0008](https://arxiv.org/abs/1201.0008)] [[INSPIRE](#)].
- [4] A. Altheimer et al., *Boosted Objects and Jet Substructure at the LHC. Report of BOOST2012, held at IFIC Valencia, 23rd–27th of July 2012*, *Eur. Phys. J. C* **74** (2014) 2792 [[arXiv:1311.2708](https://arxiv.org/abs/1311.2708)] [[INSPIRE](#)].
- [5] D. Adams et al., *Towards an Understanding of the Correlations in Jet Substructure*, *Eur. Phys. J. C* **75** (2015) 409 [[arXiv:1504.00679](https://arxiv.org/abs/1504.00679)] [[INSPIRE](#)].
- [6] 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](#)].
- [7] L. Asquith 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](#)].
- [8] L. de Oliveira, M. Kagan, L. Mackey, B. Nachman and A. Schwartzman, *Jet-images — deep learning edition*, *JHEP* **07** (2016) 069 [[arXiv:1511.05190](https://arxiv.org/abs/1511.05190)] [[INSPIRE](#)].
- [9] P. Baldi, K. Bauer, C. Eng, P. Sadowski and D. Whiteson, *Jet Substructure Classification in High-Energy Physics with Deep Neural Networks*, *Phys. Rev. D* **93** (2016) 094034 [[arXiv:1603.09349](https://arxiv.org/abs/1603.09349)] [[INSPIRE](#)].

- [10] 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](#)].
- [11] L.G. Almeida, M. Backović, M. Cliche, S.J. Lee and M. Perelstein, *Playing Tag with ANN: Boosted Top Identification with Pattern Recognition*, *JHEP* **07** (2015) 086 [[arXiv:1501.05968](#)] [[INSPIRE](#)].
- [12] 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](#)].
- [13] G. Kasieczka et al., *The Machine Learning Landscape of Top Taggers*, *SciPost Phys.* **7** (2019) 014 [[arXiv:1902.09914](#)] [[INSPIRE](#)].
- [14] H. Qu and L. Gouskos, *ParticleNet: Jet Tagging via Particle Clouds*, *Phys. Rev. D* **101** (2020) 056019 [[arXiv:1902.08570](#)] [[INSPIRE](#)].
- [15] E.A. Moreno et al., *JEDI-net: a jet identification algorithm based on interaction networks*, *Eur. Phys. J. C* **80** (2020) 58 [[arXiv:1908.05318](#)] [[INSPIRE](#)].
- [16] CMS collaboration, *Performance of the DeepJet b tagging algorithm using 41.9/fb of data from proton-proton collisions at 13 TeV with Phase 1 CMS detector*, *CMS-DP-2018-058* (2018).
- [17] 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](#)].
- [18] G. Louppe, M. Kagan and K. Cranmer, *Learning to Pivot with Adversarial Networks*, in *Advances in Neural Information Processing Systems 30*, I. Guyon et al. eds., Curran Associates, Inc. (2017), pp. 981–990 [[arXiv:1611.01046](#)] [[INSPIRE](#)].
- [19] G. Kasieczka and D. Shih, *DisCo Fever: Robust Networks Through Distance Correlation*, [arXiv:2001.05310](#) [[INSPIRE](#)].
- [20] S. Bollweg, M. Haußmann, G. Kasieczka, M. Luchmann, T. Plehn and J. Thompson, *Deep-Learning Jets with Uncertainties and More*, *SciPost Phys.* **8** (2020) 006 [[arXiv:1904.10004](#)] [[INSPIRE](#)].
- [21] G. Kasieczka, M. Luchmann, F. Otterpohl and T. Plehn, *Per-Object Systematics using Deep-Learned Calibration*, [arXiv:2003.11099](#) [[INSPIRE](#)].
- [22] B. Nachman, *A guide for deploying Deep Learning in LHC searches: How to achieve optimality and account for uncertainty*, *SciPost Phys.* **8** (2020) 090 [[arXiv:1909.03081](#)] [[INSPIRE](#)].
- [23] ATLAS collaboration, *Performance of top-quark and W-boson tagging with ATLAS in Run 2 of the LHC*, *Eur. Phys. J. C* **79** (2019) 375 [[arXiv:1808.07858](#)] [[INSPIRE](#)].
- [24] ATLAS collaboration, *Performance of mass-decorrelated jet substructure observables for hadronic two-body decay tagging in ATLAS*, *ATL-PHYS-PUB-2018-014* (2018) [[INSPIRE](#)].
- [25] CMS collaboration, *Machine learning-based identification of highly Lorentz-boosted hadronically decaying particles at the CMS experiment*, *CMS-PAS-JME-18-002* (2019) [[INSPIRE](#)].
- [26] M. Dasgupta, A. Fregoso, S. Marzani and G.P. Salam, *Towards an understanding of jet substructure*, *JHEP* **09** (2013) 029 [[arXiv:1307.0007](#)] [[INSPIRE](#)].
- [27] M. Dasgupta, A. Fregoso, S. Marzani and A. Powling, *Jet substructure with analytical methods*, *Eur. Phys. J. C* **73** (2013) 2623 [[arXiv:1307.0013](#)] [[INSPIRE](#)].

- [28] A.J. Larkoski, S. Marzani, G. Soyez and J. Thaler, *Soft Drop*, *JHEP* **05** (2014) 146 [[arXiv:1402.2657](#)] [[INSPIRE](#)].
- [29] M. Dasgupta, L. Schunk and G. Soyez, *Jet shapes for boosted jet two-prong decays from first-principles*, *JHEP* **04** (2016) 166 [[arXiv:1512.00516](#)] [[INSPIRE](#)].
- [30] G.P. Salam, L. Schunk and G. Soyez, *Dichroic subjettness ratios to distinguish colour flows in boosted boson tagging*, *JHEP* **03** (2017) 022 [[arXiv:1612.03917](#)] [[INSPIRE](#)].
- [31] M. Dasgupta, A. Powling, L. Schunk and G. Soyez, *Improved jet substructure methods: Y-splitter and variants with grooming*, *JHEP* **12** (2016) 079 [[arXiv:1609.07149](#)] [[INSPIRE](#)].
- [32] M. Dasgupta, A. Powling and A. Siodmok, *On jet substructure methods for signal jets*, *JHEP* **08** (2015) 079 [[arXiv:1503.01088](#)] [[INSPIRE](#)].
- [33] A.J. Larkoski, I. Moult and D. Neill, *Power Counting to Better Jet Observables*, *JHEP* **12** (2014) 009 [[arXiv:1409.6298](#)] [[INSPIRE](#)].
- [34] A.J. Larkoski, I. Moult and D. Neill, *Analytic Boosted Boson Discrimination*, *JHEP* **05** (2016) 117 [[arXiv:1507.03018](#)] [[INSPIRE](#)].
- [35] A.J. Larkoski, *Improving the Understanding of Jet Grooming in Perturbation Theory*, [arXiv:2006.14680](#) [[INSPIRE](#)].
- [36] A.J. Larkoski, G.P. Salam and J. Thaler, *Energy Correlation Functions for Jet Substructure*, *JHEP* **06** (2013) 108 [[arXiv:1305.0007](#)] [[INSPIRE](#)].
- [37] 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](#)].
- [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] P. Gras et al., *Systematics of quark/gluon tagging*, *JHEP* **07** (2017) 091 [[arXiv:1704.03878](#)] [[INSPIRE](#)].
- [40] J. Thaler and K. Van Tilburg, *Identifying Boosted Objects with N-subjettiness*, *JHEP* **03** (2011) 015 [[arXiv:1011.2268](#)] [[INSPIRE](#)].
- [41] F.A. Dreyer, G.P. Salam and G. Soyez, *The Lund Jet Plane*, *JHEP* **12** (2018) 064 [[arXiv:1807.04758](#)] [[INSPIRE](#)].
- [42] A.J. Larkoski and E.M. Metodiev, *A Theory of Quark vs. Gluon Discrimination*, *JHEP* **10** (2019) 014 [[arXiv:1906.01639](#)] [[INSPIRE](#)].
- [43] J.R. Andersen et al., *Les Houches 2015: Physics at TeV Colliders Standard Model Working Group Report*, in proceedings of the *9th Les Houches Workshop on Physics at TeV Colliders (PhysTeV 2015)*, Les Houches, France, 1–19 June 2015, [arXiv:1605.04692](#) [[INSPIRE](#)].
- [44] A.J. Larkoski, J. Thaler and W.J. Waalewijn, *Gaining (Mutual) Information about Quark/Gluon Discrimination*, *JHEP* **11** (2014) 129 [[arXiv:1408.3122](#)] [[INSPIRE](#)].
- [45] 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](#)].
- [46] S. Amoroso et al., *Les Houches 2019: Physics at TeV Colliders: Standard Model Working Group Report*, in proceedings of the *11th Les Houches Workshop on Physics at TeV Colliders: PhysTeV Les Houches (PhysTeV 2019)*, Les Houches, France, 10–28 June 2019, [arXiv:2003.01700](#) [[INSPIRE](#)].

- [47] E.M. Metodiev and J. Thaler, *Jet Topics: Disentangling Quarks and Gluons at Colliders*, *Phys. Rev. Lett.* **120** (2018) 241602 [[arXiv:1802.00008](#)] [[INSPIRE](#)].
- [48] P.T. Komiske, E.M. Metodiev and J. Thaler, *An operational definition of quark and gluon jets*, *JHEP* **11** (2018) 059 [[arXiv:1809.01140](#)] [[INSPIRE](#)].
- [49] J. Thaler and K. Van Tilburg, *Maximizing Boosted Top Identification by Minimizing N -subjettiness*, *JHEP* **02** (2012) 093 [[arXiv:1108.2701](#)] [[INSPIRE](#)].
- [50] K. Datta and A.J. Larkoski, *How Much Information is in a Jet?*, *JHEP* **06** (2017) 073 [[arXiv:1704.08249](#)] [[INSPIRE](#)].
- [51] 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](#)].
- [52] 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](#)].
- [53] D. Napoletano and G. Soyez, *Computing N -subjettiness for boosted jets*, *JHEP* **12** (2018) 031 [[arXiv:1809.04602](#)] [[INSPIRE](#)].
- [54] A.J. Larkoski, I. Moult and D. Neill, *Factorization and Resummation for Groomed Multi-Prong Jet Shapes*, *JHEP* **02** (2018) 144 [[arXiv:1710.00014](#)] [[INSPIRE](#)].
- [55] I. Moult, B. Nachman and D. Neill, *Convolved Substructure: Analytically Decorrelating Jet Substructure Observables*, *JHEP* **05** (2018) 002 [[arXiv:1710.06859](#)] [[INSPIRE](#)].
- [56] J. Neyman and E.S. Pearson, *On the Problem of the Most Efficient Tests of Statistical Hypotheses*, *Phil. Trans. Roy. Soc. Lond. A* **231** (1933) 289.
- [57] Y.L. Dokshitzer, G.D. Leder, S. Moretti and B.R. Webber, *Better jet clustering algorithms*, *JHEP* **08** (1997) 001 [[hep-ph/9707323](#)] [[INSPIRE](#)].
- [58] M. Wobisch and T. Wengler, *Hadronization corrections to jet cross-sections in deep inelastic scattering*, in proceedings of the *Workshop on Monte Carlo Generators for HERA Physics (Plenary Starting Meeting)*, Hamburg, Germany, 27–30 April 1998, pp. 270–279 [[hep-ph/9907280](#)] [[INSPIRE](#)].
- [59] M.A. Nielsen, *Neural Networks and Deep Learning*, Determination Press (2015).
- [60] T. Sjöstrand et al., *An Introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159 [[arXiv:1410.3012](#)] [[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] M. Cacciari, G.P. Salam and G. Soyez, *FastJet User Manual*, *Eur. Phys. J. C* **72** (2012) 1896 [[arXiv:1111.6097](#)] [[INSPIRE](#)].
- [63] A.J. Larkoski, D. Neill and J. Thaler, *Jet Shapes with the Broadening Axis*, *JHEP* **04** (2014) 017 [[arXiv:1401.2158](#)] [[INSPIRE](#)].
- [64] A.L. Maas, A.Y. Hannun and A.Y. Ng, *Rectifier nonlinearities improve neural network acoustic models*, in proceedings of the *ICML Workshop on Deep Learning for Audio, Speech and Language Processing*, Atlanta, GA, U.S.A., 16 June 2013.
- [65] D.P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization*, [arXiv:1412.6980](#) [[INSPIRE](#)].
- [66] Y. Yao, L. Rosasco and A. Caponnetto, *On early stopping in gradient descent learning*, *Constr. Approx.* **26** (2007) 289.