

# Weakly supervised classification in high energy physics

---

Lucio Mwinmaarong Dery,<sup>a</sup> Benjamin Nachman,<sup>b</sup> Francesco Rubbo<sup>c</sup>  
and Ariel Schwartzman<sup>c</sup>

<sup>a</sup>*Physics Department, Stanford University,  
Stanford, CA, 94305, U.S.A.*

<sup>b</sup>*Physics Division, Lawrence Berkeley National Laboratory,  
1 Cyclotron Rd, Berkeley, CA, 94720, U.S.A.*

<sup>c</sup>*SLAC National Accelerator Laboratory, Stanford University,  
2575 Sand Hill Rd, Menlo Park, CA, 94025, U.S.A.*

*E-mail:* [ldery@slac.stanford.edu](mailto:ldery@slac.stanford.edu), [bnachman@cern.ch](mailto:bnachman@cern.ch),  
[rubbo@slac.stanford.edu](mailto:rubbo@slac.stanford.edu), [sch@slac.stanford.edu](mailto:sch@slac.stanford.edu)

**ABSTRACT:** As machine learning algorithms become increasingly sophisticated to exploit subtle features of the data, they often become more dependent on simulations. This paper presents a new approach called *weakly supervised classification* in which class proportions are the only input into the machine learning algorithm. Using one of the most challenging binary classification tasks in high energy physics — quark versus gluon tagging — we show that weakly supervised classification can match the performance of fully supervised algorithms. Furthermore, by design, the new algorithm is insensitive to any mis-modeling of discriminating features in the data by the simulation. Weakly supervised classification is a general procedure that can be applied to a wide variety of learning problems to boost performance and robustness when detailed simulations are not reliable or not available.

**KEYWORDS:** Jets

**ARXIV EPRINT:** [1702.00414](https://arxiv.org/abs/1702.00414)

---

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Weakly supervised classification</b>	<b>2</b>
<b>3</b>	<b>Example: quark and gluon jet discrimination</b>	<b>5</b>
<b>4</b>	<b>Conclusions</b>	<b>7</b>

---

## 1 Introduction

With the increasing complexity of theoretical models and computing power, many scientific projects increasingly rely on simulations to design analysis techniques. This is especially true for high energy particle physics, where high fidelity Monte Carlo (MC) simulation is used to model physical processes at distances ranging from  $10^{-25}$  meters all the way to the macroscopic dimensions of detectors. However, just as models have become more complex, analysis techniques have also become more sophisticated. Numerous, often subtle, features of events are combined using powerful supervised learning algorithms trained on large simulated (labeled) datasets. Despite these advances, there is no guarantee that techniques highly optimized in simulation are also optimal in nature. For example, it is well-known that comparisons of tagging efficiencies between data and simulation result in large ( $\gtrsim 10\%$ - $30\%$ ) corrections (*scale factors*) for  $b$ -tagging [1, 2], quark/gluon tagging [3, 4], boosted  $W$  tagging [5, 6], and high  $p_T$  top quark tagging [7, 8]. Non-unity correction factors show that the classifiers trained on simulation are necessarily sub-optimal when applied to data. It is sometimes the case that one knows the proportions of each class better than the properties of each class that are useful for classification. *Weakly supervised classification* is a new machine learning paradigm for classification where training is performed directly on (unlabeled) data.

The task of training a classifier on multidimensional data based only on class proportions is highly under-constrained. Neural network training is already a non-convex problem, but removing all local information about class labels significantly increases the difficulty of optimization. However, the field of multi-instance learning (MIL) [9] has shown that local information is not necessarily needed for classification.<sup>1</sup> The setup of MIL is a series of sets (‘bags’) of individual instances without individual labels. Consider the task of distinguishing two classes, called  $A$  and  $B$ . For the training set, it is known if a bag contains at least one instance of class  $A$ . The algorithm is then optimized to identify the presence of at least one instance of class  $A$  in an unseen bag. Recent work has extended this procedure to identify the class of individual instances, still only training on bag-level

---

<sup>1</sup>See ref. [10] for a review of recent work in MIL.

labels [11]. In this paper we make *supervision even weaker* in that bag-labels are only known on average. In particular, all that is known to the training is the expected fraction of class  $A$  in any particular bag. This paradigm is also referred to as Learning with Label Proportions (LLP) [12].

High energy quarks and gluons produced in reactions at the Large Hadron Collider (LHC) result in collimated streams of particles traveling at nearly the speed of light, known as *jets*. One of the most challenging classification tasks in high energy physics is to distinguish quark-induced jets from gluon-induced jets based on their radiation pattern. There is an extensive literature exploring discriminating observables [13]. However, standard quark and gluon discriminants are known to be poorly modeled by state-of-the-art simulations [3, 14]. Despite this, the fraction of quark and gluon jets in a given sample is often well-known. At a fixed order in perturbation theory, the probability for an out-going parton to be a quark or a gluon depends on well-known parton distribution functions and matrix elements. Therefore, quark versus gluon discrimination is well-suited for weakly supervised classification and is therefore the main example used later in this paper.

This paper is organized as follows. Section 2 formally introduces weakly supervised classification and describes how it is applied in practice. The technique is illustrated using quark versus gluon jet tagging in section 3. The paper ends in section 4 with some concluding remarks. The source code implemented to produce the results presented in this paper is available at <https://doi.org/10.5281/zenodo.322813>.

## 2 Weakly supervised classification

Given a set of data originating from two classes labeled 0 and 1, the goal of classification is to construct a function  $f : \mathbb{R}^n \rightarrow \{0, 1\}$ , where  $n$  is the dimensionality of the feature space used to discriminate the two classes. In the traditional classification paradigm of fully supervised training, the function  $f_{\text{full}}$  is built by minimizing a loss function like the following:

$$f_{\text{full}} = \operatorname{argmin}_{f': \mathbb{R}^n \rightarrow \{0,1\}} \sum_{i=1}^N \ell(f'(x_i) - t_i), \quad (2.1)$$

where  $N$  is the number of labeled data available for training,  $\ell$  is a loss function with  $\lim_{x \rightarrow 0} \ell(x) = 0$ , and  $t_i$  is the true label of example  $i$ . A common loss function is the squared error. In order to provide flexibility and stability, one often modifies the original problem to take  $f : \mathbb{R}^n \rightarrow [0, 1]$ . The ideal classifier that one tries to approximate with eq. (2.1) is based on the likelihood ratio  $p(\vec{x}|0)/p(\vec{x}|1)$ , where  $p(\vec{x}|i)$  is the  $n$ -dimensional probability density for the feature vector  $\vec{x}$  for the class  $i \in \{0, 1\}$ . Ideally,  $f(x)$  would be 1 when the likelihood ratio is maximum and 0 when it is minimum. *Weakly supervised classification* is a new paradigm in which instead of knowing the  $t_i$ , all that is known is the proportion of events in either class:  $y = \sum_i t_i/N$ . Thus, the weakly supervised  $f_{\text{weak}}$  is given by

$$f_{\text{weak}} = \operatorname{argmin}_{f': \mathbb{R}^n \rightarrow [0,1]} \ell \left( \sum_{i=1}^N \frac{f'(x_i)}{N} - y \right). \quad (2.2)$$

The argument of eq. (2.2) is non-convex, with many minima. In particular, the trivial solution  $f'(x) = y$  results in a loss of zero. However, using multiple batches of data with different proportions  $y_k$  is sufficient to collapse the solution space, so long as the distribution  $p(\vec{x}|i; k) = p(\vec{x}|i)$ , i.e. the distribution of the discriminating features for a particular class is the same in every batch  $k$ . To build intuition for why there is any hope to solve this problem, consider a case where there are two batches  $A$  and  $B$  with proportions  $y_A$  and  $y_B$ . Consider an  $n$ -dimensional histogram where the  $i^{\text{th}}$  dimension captures a discretized version of the  $i^{\text{th}}$  discriminating feature. If the  $i^{\text{th}}$  dimension has  $m_i$  bins, then the total number of bins in the histogram is  $M = \sum_{i=1}^n m_i$ . One can always rearrange bins so that instead of an  $n$ -dimensional histogram with  $m_i$  bins in the  $i^{\text{th}}$  dimension, there is a one-dimensional histogram with  $M$  bins. As visualizing high dimensional histograms can be cumbersome, let  $h_A$  be one-dimensional histograms with  $M$  bins for the batch  $A$  and  $h_B$  be the corresponding histogram for batch  $B$ . Then, for each bin  $i$ , one can write

$$h_{A,i} = y_A h_{1,i} + (1 - y_A) h_{0,i} \tag{2.3}$$

$$h_{B,i} = y_B h_{1,i} + (1 - y_B) h_{0,i}, \tag{2.4}$$

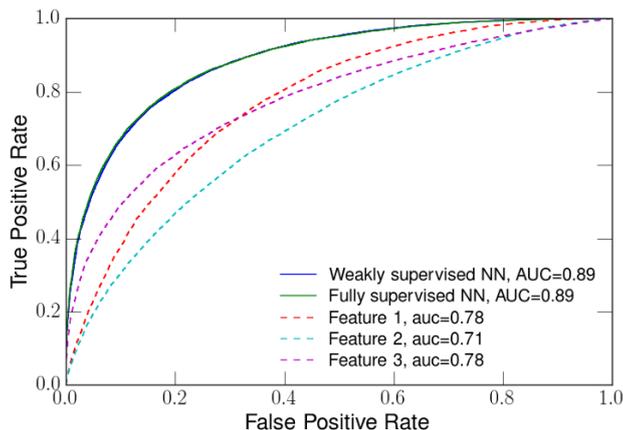
where  $h_{X,i}$  is the content of the  $i^{\text{th}}$  bin of the histogram  $h_X$ . Except for contrived scenarios, eq. (2.3) will have a unique solution for  $h_{0,i}$  and  $h_{1,i}$ , which are discretized versions of the probability densities  $p(\vec{x}|0)$  and  $p(\vec{x}|1)$ . One can then form an (approximately) optimal classifier from the ratio of histograms with bin contents  $h_{0,i}/h_{1,i}$ . If the number of dimensions is large, one can add a further step to use machine learning to approximate the optimal classifier from  $h_{0,i}$  and  $h_{1,i}$ . As a result, the problem is completely solvable. Weakly supervised training combines the classification step with the first step and does so without binning. Solving eq. (2.3) ‘by-hand’ is intractable when  $n$  is relatively large or the number of examples is relatively small. It is also complicated when there are more than two batches (over-constrained). These challenges are all naturally handled by the all-in-one machine learning approach of weakly supervised classification, as illustrated below.

In the weakly supervised training used in the following examples,  $f'$  in eq. (2.2) is parametrized as a three-layer neural network with three inputs, a hidden layer with 30 neurons, and a sigmoid output. We use the Adam optimizer [15] in Keras [16] with a learning rate of 0.009 and train for 25 iterations. As reference, we consider a traditional classifier

$$f_{\text{full}} = \operatorname{argmin}_{f': \mathbb{R}^n \rightarrow [0,1]} \ell(f'(x_i) - t_i), \tag{2.5}$$

where  $t_i$  labels the individual instances and  $f'$  is parametrized as a three-layer neural network with three inputs, a hidden layer with 10 neurons, and a sigmoid output. Minimization is performed with stochastic gradient descent in Keras with a learning rate of 0.01 run for 40 iterations. For each training, both networks are initialized with random weights, following a normal distribution.

Figure 1 shows the weakly supervised classifier performance when training with 9 subsets of data with proportions between 0.2 and 0.4 compared with that of the fully supervised one. Three features, labeled 1 – 3 are constructed so that the distribution of feature  $i$  given class  $j$  follows a normal distribution with mean  $\mu_{ij}$  and standard deviation



**Figure 1.** Receiver Operator Characteristic (ROC) curves for instance classification, using three individual features and then combined using a fully supervised network and the weakly supervised classifier. One performance metric is the Area Under the Curve (AUC) which is the integral of the ROC curve.

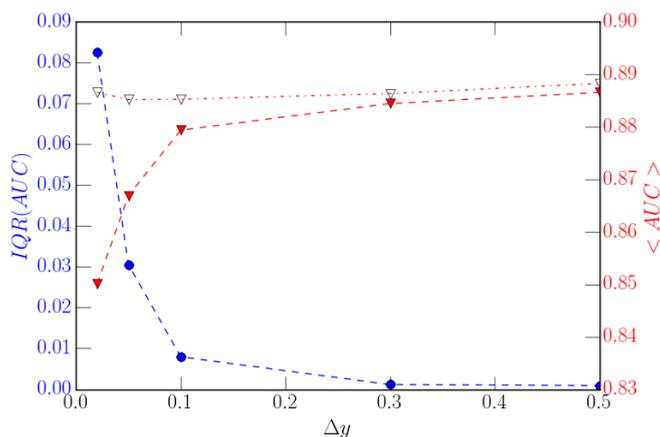
Feature	$\mu_0$	$\sigma_0$	$\mu_1$	$\sigma_1$
1	26	8	18	7
2	0.09	0.04	0.06	0.04
3	0.28	0.04	0.23	0.05

**Table 1.** Mean ( $\mu$ ) and standard deviation ( $\sigma$ ) values of the normal distributions for class 0 and 1 of each feature.

$\sigma_{ij}$ . For reference, the values of  $\mu_{ij}$  and  $\sigma_{ij}$  used for the example shown in figure 1 are in table. 1. Both the traditional and weakly supervised classifiers have the same Receiver Operator Characteristic (ROC) and thus have identical classification performance. Note that the loss for weakly supervised classification is symmetric with respect to swapping the class assignment, therefore the classifier output for a given training can give higher values for class 0, while for a different training it would give higher value for class 1.

As with any machine learning algorithm with inherent randomness, the performance of a weakly supervised classifier has a stochastic component. This is quantified by retraining the same network many times with different random number seeds in each iteration. The interquartile range (IQR) over the Area Under the Curve (AUC) values for each training is a measure of the spread due to the inherent randomness. Figure 2 shows the AUC IQR for the toy example with one proportion fixed to 0.2 and the second proportion scanned from 0.2 to 0.7. The stability improves as the difference between the class proportions increases. In addition to the performance varying less as the proportions are further apart, the overall performance quantified by the median AUC (denoted by  $\langle \text{AUC} \rangle$ ) also improves (increases). The improvement in the median AUC is not as dramatic as the reduction in the AUC IQR, but it does suggest that it is (slightly) easier for the machine learning algorithm when the proportions are very different.<sup>2</sup> This makes sense in the context of the

<sup>2</sup>Even when the proportions are within few percents, stable performance can be achieved if multiple



**Figure 2.** Median (solid triangles) and interquartile range (solid dots) of the AUC as a function of the difference in proportions  $\Delta y$  between the two subsets of the training sample. The proportion corresponding to one subset is fixed to 0.2, while the other varies. For each point the AUC is computed 100 times on the same test set with different trainings, each performed with a random weight initialization. The maximum AUC for each point is also shown (hollow triangles).

two-step intuition-building paradigm given above: the algorithm can spend more attention on the classification task if it is *easier* to extract the class distributions. *Easier* here refers to the fact that the solution to eq. (2.3) has a Jacobian factor with  $(\Delta y)^{-1}$ , which becomes numerically unstable as  $\Delta y$  becomes small.

### 3 Example: quark and gluon jet discrimination

Due to the strength of the strong force, there is a plethora of gluon jets produced at the LHC. However, many processes result in mostly quark jets. Prominent examples include the identification of hadronically decaying  $W$  bosons [17, 18], jets associated with vector boson fusion [19–21], and multi-quarks resulting from supersymmetry [22]. The references given here are the small number of public results that mention quark/gluon tagging, but there many more analyses that would benefit from a tagger if a robust technique existed.

The weakly supervised classification strategy is particularly useful for quark/gluon tagging because the fraction of quark jets for a particular set of events is well-known from parton distribution functions (PDF) and matrix element (ME) calculations while useful discriminating features have not been computed to high accuracy and simulations often mis-model the data. In particular, the fraction of out-going quarks and gluons depends on the convolution of PDF and ME calculations where the state-of-the-art is automated next-to-leading order [23–25] with many independent input and validation measurements spanning LEP, SLD, Tevatron, and the LHC. In contrast, fragmentation is modeled with a perturbative parton shower that only resums the leading logarithms [26–28] and hadronization models that are not derived from first principles (non-perturbative phenomenological models) [29, 30]. There is a lot of powerful discrimination information inside the jet radi-

---

( $> 2$ ) subsets with different proportions can be used for training.

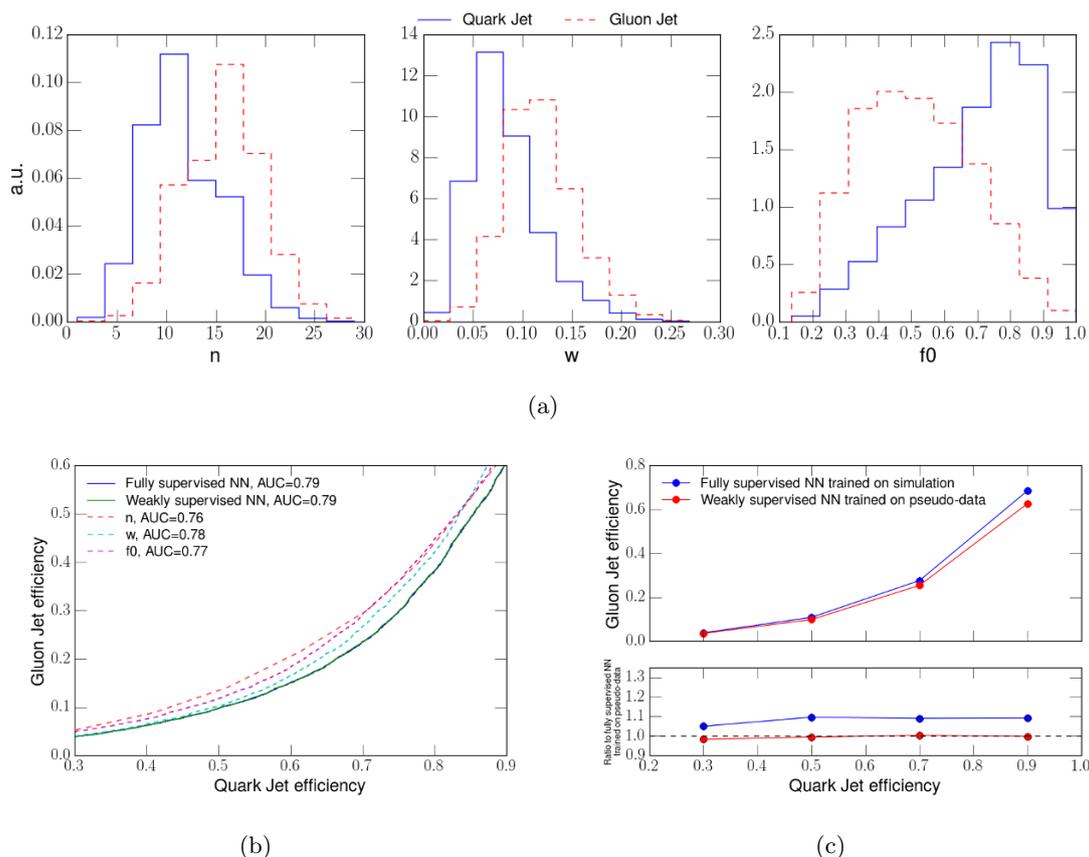
ation pattern,<sup>3</sup> so it is of great interest to use this information without heavily relying on simulation for fragmentation modeling.

To illustrate this concrete example, quark and gluon jets are simulated and a weakly supervised classifier is trained on the generated event sample. Unlike real data, in the simulated sample, we also know per-event labels which are used to additionally train a fully supervised classifier. Events with  $2 \rightarrow 2$  quark-gluon scattering (*dijet* events) are simulated using the Pythia 8.18 [34] event generator. Jets are clustered using the anti- $k_t$  algorithm [35] with distance parameter  $R = 0.4$  via the FastJet 3.1.3 [36] package. Jets are classified as quark- or gluon-initiated by considering the type of the highest energy quark or gluon in the full generator event record that is inside a 0.3 radius of the jet axis. For simplicity, one transverse momentum range is considered:  $45 \text{ GeV} < p_T < 55 \text{ GeV}$ . Additionally, there is a pseudo-rapidity requirement that mimics the usual detector acceptance for charged particle tracking:  $|\eta| < 2.1$ . Heuristically, gluons have twice as much strong-force charge as quark jets, resulting in more constituents and a broader radiation pattern. More precisely, at leading logarithm the ratio of the average gluon and quark jet constituent multiplicity is equal to the ratio of the QCD color-factor (*Casimir*) associated with quarks ( $C_F$ ) and gluons ( $C_A$ ):  $C_A/C_F = 9/4 \sim 2$  [13, 37]. In fact, due to Casimir scaling, the ROC curve for any  $p_T$  and angle-weighted moments of the jet radiation pattern is set by  $C_A/C_F$  [38]. The following variables are useful for quark/gluon discrimination: the number of jet constituents  $n$ , the first radial moment in  $p_T$  (*jet width*)  $w$ , and the fraction of the jet  $p_T$  carried by the leading anti- $k_T$   $R = 0.1$  subjet  $f_0$ . The constituents  $i$  considered for computing  $n$  and  $w$  are the hadrons in the jet with  $p_T > 500 \text{ MeV}$ .

A weakly supervised classifier with one hidden layer of size 30 is trained by considering 12 bins of the distribution of the absolute difference in pseudorapidity between the two jets [39]. The proportion of quark initiated jets varies between 0.21 and 0.32. Figure 3 shows that, while the individual observables perform differently in the high or low gluon efficiency (true positive rate) regimes, their combination in a NN gives consistently better performance. The weakly supervised classifier matches the performance of the fully supervised NN, despite only knowing sample proportions instead of individual event labels. By construction the weakly supervised classifier is also robust against a realistic amount of mis-modeling in the input variables. This feature is tested by building a pseudo-data sample where the probability distributions of  $n$  and  $w$  are distorted in the training sample to emulate the difference in efficiency measured in ref. [3]. The study in ref. [3] found that a classifier extracted from simulation is more powerful than one extracted from the data. This is reflected in the results presented in the right plot of figure 3. When a fully supervised classifier is trained on a sample generated with the same distribution as the test sample (mimicking training and testing on simulation), it achieves a better performance than when trained on the original sample and tested on the distorted pseudo-data (mimicking training on simulation and testing on data). In contrast, the weakly supervised classifier can be trained directly on the distorted pseudo-data sample (representing the data) so is insensitive to the mismodeling of the input variables. This results in a 10%

---

<sup>3</sup>See refs. [3, 4, 13, 31, 32] and also ref. [33] for a summary.

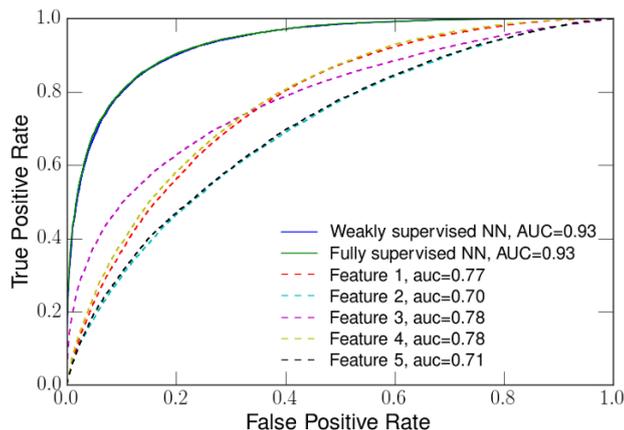


**Figure 3.** Distributions of the discriminating features  $n$ ,  $w$  and  $f_0$  (a) and comparison of ROC curves for quark/gluon jet discrimination using a fully supervised classifier or a weakly supervised classifier. In (b) the fully and weakly supervised classifiers are trained on identical simulated data and evaluated on a test sample drawn from the same population. The weakly supervised classifier matches the performance of the fully supervised one. The curves corresponding to the three input observables used as discriminant are shown as reference. In (c), the fully supervised classifier (blue line) is trained on a labeled simulated training sample. The weakly supervised classifier (red line) is trained on an unlabeled pseudo-data training sample. In both cases, the performance is evaluated on the same pseudo-data test sample. The ratios to the performance of a fully supervised classifier trained on a labeled pseudo-data sample are shown in the bottom pad.

bias from the standard procedure that is avoided by the weakly supervised classifier. Even larger differences may be expected from this and other classification tasks that utilize even more input features or are more mis-modeled. The weakly supervised classifier is robust and outperforms the standard supervised learning trained on simulation.

## 4 Conclusions

We have presented a new approach to classification with NN in cases where class proportions are known but individual labels are not readily available. This weakly supervised classification has broad applicability and has been demonstrated in one important discrim-



**Figure 4.** ROC curves for instance classification using five individual features and then combined using a fully supervised network and the weakly supervised classifier.

ination task in high energy physics: quark versus gluon jet tagging. In the quark/gluon and related contexts, weakly supervised classification provides a robust and powerful approach because it can be directly trained on examples from (unlabeled) data instead of (labeled, but unreliable) simulation. The examples presented so far have used a small number of input features to illustrate the ideas, but there is no algorithmic limitation on the number of features. Figure 4 is a simple extension of figure 1 with 5 features instead of 3; in future work, we will study the extension to many more features (tens to hundreds). This paper has laid the conceptual groundwork for this new tool that has started a new classification paradigm that can be applied to a wide variety of learning problems to boost performance and robustness when detailed simulations are not reliable or not available.

## Acknowledgments

This work is supported by the Stanford Data Science Initiative and by the US Department of Energy (DOE) under grant DE-AC02-76SF00515. We would like to thank Russell Stewart for useful discussions about label-free supervision strategies and non-convex optimization problems and Gilles Louppe for useful discussion about related work on learning from label proportions.

**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] ATLAS collaboration, *Performance of  $b$ -jet identification in the ATLAS experiment*, 2016 *JINST* **11** P04008 [[arXiv:1512.01094](https://arxiv.org/abs/1512.01094)] [[INSPIRE](https://inspirehep.net/literature/1512094)].
- [2] CMS collaboration, *Identification of  $b$ -quark jets with the CMS experiment*, 2013 *JINST* **8** P04013 [[arXiv:1211.4462](https://arxiv.org/abs/1211.4462)] [[INSPIRE](https://inspirehep.net/literature/1211446)].

- [3] ATLAS collaboration, *Light-quark and gluon jet discrimination in pp collisions at  $\sqrt{s} = 7$  TeV with the ATLAS detector*, *Eur. Phys. J. C* **74** (2014) 3023 [[arXiv:1405.6583](#)] [[INSPIRE](#)].
- [4] CMS collaboration, *Performance of quark/gluon discrimination in 8 TeV pp data*, *CMS-PAS-JME-13-002* (2013).
- [5] ATLAS collaboration, *Identification of boosted, hadronically decaying W bosons and comparisons with ATLAS data taken at  $\sqrt{s} = 8$  TeV*, *Eur. Phys. J. C* **76** (2016) 154 [[arXiv:1510.05821](#)] [[INSPIRE](#)].
- [6] CMS collaboration, *Identification techniques for highly boosted W bosons that decay into hadrons*, *JHEP* **12** (2014) 017 [[arXiv:1410.4227](#)] [[INSPIRE](#)].
- [7] ATLAS collaboration, *Identification of high transverse momentum top quarks in pp collisions at  $\sqrt{s} = 8$  TeV with the ATLAS detector*, *JHEP* **06** (2016) 093 [[arXiv:1603.03127](#)] [[INSPIRE](#)].
- [8] CMS collaboration, *Boosted top jet tagging at CMS*, *CMS-PAS-JME-13-007* (2013).
- [9] T.G. Dietterich, R.H. Lathrop and T. Lozano-Pérez, *Solving the multiple instance problem with axis-parallel rectangles*, *Artif. Intell.* **89** (1997) 31.
- [10] J. Amores, *Multiple instance classification: Review, taxonomy and comparative study*, *Artif. Intell.* **201** (2013) 81.
- [11] D. Kotzias, M. Denil, N. de Freitas and P. Smyth, *From group to individual labels using deep features*, in the proceedings of the 21<sup>th</sup> *ACM SIGKDD international conference on knowledge discovery and data mining (KDD15)*, August 10–13, Sydney, Australia (2015).
- [12] G. Patrini, R. Nock, P. Rivera and T. Caetano, *(Almost) No label no cry*, in *Advances in Neural Information Processing Systems 27*, Z. Ghahramani et al. eds., Curran Associates Inc., U.S.A. (2014).
- [13] J. Gallicchio and M.D. Schwartz, *Quark and gluon tagging at the LHC*, *Phys. Rev. Lett.* **107** (2011) 172001 [[arXiv:1106.3076](#)] [[INSPIRE](#)].
- [14] J.R. Andersen et al., *Les Houches 2015: Physics at TeV Colliders Standard Model Working Group Report*, in the proceedings of the 9<sup>th</sup> *Les Houches Workshop on Physics at TeV Colliders (PhysTeV 2015)*, June 1–19, Les Houches, France (2016), [[arXiv:1605.04692](#)] [[INSPIRE](#)].
- [15] D.P. Kingma and J. Ba, *Adam: A method for stochastic optimization*, [[arXiv:1412.6980](#)].
- [16] F. Chollet, *Keras*, <https://github.com/fchollet/keras> (2015).
- [17] CMS collaboration, *V tagging observables and correlations*, *CMS-PAS-JME-14-002* (2014).
- [18] ATLAS collaboration, *Search for high-mass diboson resonances with boson-tagged jets in proton-proton collisions at  $\sqrt{s} = 8$  TeV with the ATLAS detector*, *JHEP* **12** (2015) 055 [[arXiv:1506.00962](#)] [[INSPIRE](#)].
- [19] CMS collaboration, *Search for the standard model Higgs boson produced through vector boson fusion and decaying to  $b\bar{b}$* , *Phys. Rev. D* **92** (2015) 032008 [[arXiv:1506.01010](#)] [[INSPIRE](#)].
- [20] CMS collaboration, *Measurement of electroweak production of two jets in association with a Z boson in proton-proton collisions at  $\sqrt{s} = 8$  TeV*, *Eur. Phys. J. C* **75** (2015) 66 [[arXiv:1410.3153](#)] [[INSPIRE](#)].
- [21] ATLAS collaboration, *Search for the Standard Model Higgs boson produced by vector-boson fusion and decaying to bottom quarks in  $\sqrt{s} = 8$  TeV pp collisions with the ATLAS detector*, *JHEP* **11** (2016) 112 [[arXiv:1606.02181](#)] [[INSPIRE](#)].

- [22] B. Bhattacharjee, S. Mukhopadhyay, M.M. Nojiri, Y. Sakaki and B.R. Webber, *Quark-gluon discrimination in the search for gluino pair production at the LHC*, *JHEP* **01** (2017) 044 [[arXiv:1609.08781](#)] [[INSPIRE](#)].
- [23] J. Rojo et al., *The PDF4LHC report on PDFs and LHC data: Results from Run I and preparation for Run II*, *J. Phys. G* **42** (2015) 103103 [[arXiv:1507.00556](#)] [[INSPIRE](#)].
- [24] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections and their matching to parton shower simulations*, *JHEP* **07** (2014) 079 [[arXiv:1405.0301](#)] [[INSPIRE](#)].
- [25] S. Alioli, P. Nason, C. Oleari and E. Re, *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*, *JHEP* **06** (2010) 043 [[arXiv:1002.2581](#)] [[INSPIRE](#)].
- [26] T. Sjöstrand, S. Mrenna and P.Z. Skands, *PYTHIA 6.4 physics and manual*, *JHEP* **05** (2006) 026 [[hep-ph/0603175](#)] [[INSPIRE](#)].
- [27] M. Bahr et al., *HERWIG++ physics and manual*, *Eur. Phys. J. C* **58** (2008) 639 [[arXiv:0803.0883](#)] [[INSPIRE](#)].
- [28] T. Gleisberg et al., *Event generation with SHERPA 1.1*, *JHEP* **02** (2009) 007 [[arXiv:0811.4622](#)] [[INSPIRE](#)].
- [29] B. Andersson, G. Gustafson, G. Ingelman and T. Sjöstrand, *Parton fragmentation and string dynamics*, *Phys. Rept.* **97** (1983) 31 [[INSPIRE](#)].
- [30] B.R. Webber, *A QCD model for jet fragmentation including soft gluon interference*, *Nucl. Phys. B* **238** (1984) 492 [[INSPIRE](#)].
- [31] ALEPH collaboration, D. Buskulic et al., *Quark and gluon jet properties in symmetric three jet events*, *Phys. Lett. B* **384** (1996) 353 [[INSPIRE](#)].
- [32] 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](#)].
- [33] G. Dissertori, I.G. Knowles and M. Schmelling, *High energy experiments and theory*, Clarendon Press, Oxford U.K. (2003).
- [34] T. Sjöstrand, S. Mrenna and P.Z. Skands, *A brief introduction to PYTHIA 8.1*, *Comput. Phys. Commun.* **178** (2008) 852 [[arXiv:0710.3820](#)] [[INSPIRE](#)].
- [35] M. Cacciari, G.P. Salam and G. Soyez, *The anti- $k_t$  jet clustering algorithm*, *JHEP* **04** (2008) 063 [[arXiv:0802.1189](#)] [[INSPIRE](#)].
- [36] M. Cacciari, G.P. Salam and G. Soyez, *FastJet user manual*, *Eur. Phys. J. C* **72** (2012) 1896 [[arXiv:1111.6097](#)] [[INSPIRE](#)].
- [37] PARTICLE DATA GROUP collaboration, K.A. Olive et al., *Review of particle physics*, *Chin. Phys. C* **38** (2014) 090001 [[INSPIRE](#)].
- [38] A.J. Larkoski, J. Thaler and W.J. Waalewijn, *Gaining (mutual) information about quark/gluon discrimination*, *JHEP* **11** (2014) 129 [[arXiv:1408.3122](#)] [[INSPIRE](#)].
- [39] ATLAS collaboration, *Measurement of the charged-particle multiplicity inside jets from  $\sqrt{s} = 8$  TeV pp collisions with the ATLAS detector*, *Eur. Phys. J. C* **76** (2016) 322 [[arXiv:1602.00988](#)] [[INSPIRE](#)].