

Machine Learning Algorithms for b -Jet Tagging at the ATLAS Experiment

Michela Paganini
on behalf of the ATLAS Collaboration

Yale University, Department of Physics, 217 Prospect Street, New Haven, CT 06511-8499

E-mail: michela.paganini@cern.ch

Abstract. The separation of b -quark initiated jets from those coming from lighter quark flavors (b -tagging) is a fundamental tool for the ATLAS physics program at the CERN Large Hadron Collider. The most powerful b -tagging algorithms combine information from low-level taggers, exploiting reconstructed track and vertex information, into machine learning classifiers. The potential of modern deep learning techniques is explored using simulated events, and compared to that achievable from more traditional classifiers such as boosted decision trees.

1. Introduction

With the increase in the expected integrated luminosity delivered in the remaining years of Run 2 by the Large Hadron Collider [1] (LHC), a growing emphasis is put on algorithmic improvements at the reconstruction level [2] to leverage the full sensitivity potential for measurements and searches of interest to the ATLAS [3] physics program.

The task of identifying jets that contain b -hadrons and separating them from jets initiated by lighter quark flavors is of particular significance due to the prominence of b -hadrons in final states of interesting physics processes, and to the abundant background light-flavored-jet production at hadron colliders such as the LHC. This task, commonly referred to as b -tagging, is broadly applied to both precise standard model measurements and searches for signatures of new physics phenomena. b -tagging can be cast as a classification problem with the objective of correctly assigning jet flavors using the trajectories of charged particles (tracks) reconstructed




Figure 1: The set of b -tagging algorithms currently in use in the ATLAS experiment at the LHC includes five low-level taggers (IP3D, RNNIP, SMT, SV1, and JETFITTER), as well as two high level taggers (MV2, and DL1) which take as inputs combinations of the outputs of the low-level taggers.

in the inner detector [4], and then associated to jets reconstructed from clusters of energy in the electromagnetic [5] and hadronic [6] calorimeters.

We present the current hierarchy of algorithms used for b -tagging in ATLAS (Figure 1), and outline the latest Monte Carlo-based optimization results [7], with specific focus on deep learning-powered improvements, which include the use of fully-connected and recurrent neural networks. These are the algorithms to be used to analyze the data recorded in 2017-2018.

2. Heavy hadron topology

The typical heavy hadron topology presents one or more vertices that are displaced from the hard-scatter interaction point. The longer lifetimes of heavy-flavor hadrons correspond to macroscopic decay lengths that can be resolved by the ATLAS vertexing detectors. The hadron's decay is governed by its heavy quark decay chain. A b -hadron will usually decay through a cascade to a c -hadron, strange hadron, etc. Therefore, in a b -initiated event, the goal is to reconstruct not only the primary vertex whence the event initiates, but also a secondary, and possibly tertiary vertex along the decay cascade. Finally, the heavy-flavor hadrons' non-negligible branching ratio to semi-leptonic decays ($\approx 21\%$) suggests using the presence of soft muons within jets as yet another discriminating feature for flavor tagging.

Jets used for b -tagging are reconstructed using the anti- k_t algorithm [8] with radius $R = 0.4$, and pass the following selection criteria [7]:

- $p_T^{\text{jet}} > 20 \text{ GeV}$
- $|\eta^{\text{jet}}| > 2.5$
- if $|\eta^{\text{jet}}| < 2.4$ and $p_T^{\text{jet}} < 60 \text{ GeV}$: $\text{JVT}^{\text{jet}} > 0.59$

where JVT is the output of the Jet Vertex Tagger algorithm [9] used to suppress jets from pile-up interactions. An additional overlap removal requirement is imposed to discard electrons reconstructed as jets.

Tracks are associated to the nearest jet based on an angular separation criterion that varies as a function of the jet p_T [10].

At truth level, using the kinematic properties of the jet under investigation and the hadrons in the event, the logic of the devised labeling scheme goes as follows:

```

if ( $\Delta R(\text{jet}, B_{\text{hadron}}) < 0.3$ ) & ( $p_{T_{B \text{ hadron}}} > 5 \text{ GeV}$ ) then
  jet flavor  $\leftarrow b$ 
else if ( $\Delta R(\text{jet}, C_{\text{hadron}}) < 0.3$ ) & ( $p_{T_{C \text{ hadron}}} > 5 \text{ GeV}$ ) then
  jet flavor  $\leftarrow c$ 
else if ( $\Delta R(\text{jet}, \tau_{\text{lepton}}) < 0.3$ ) & ( $p_{T_{\tau \text{ lepton}}} > 5 \text{ GeV}$ ) then
  jet flavor  $\leftarrow \tau$ 
else
  jet flavor  $\leftarrow \text{light}$ 
end if

```

where $\Delta R(a, b) \equiv \sqrt{(\phi_a - \phi_b)^2 + (\eta_a - \eta_b)^2}$ ¹ and p_T identifies the component of the momentum that is perpendicular to the beam direction.

3. Low-level taggers

Given the unique topology of heavy hadron decays (Section 2), ATLAS deploys a suite of low-level b -tagging algorithms that aim at exploiting different discriminative aspects of b -quark jets.

¹ ATLAS uses a right-handed coordinate system with its origin at the nominal interaction point in the centre of the detector and the z -axis along the beam pipe. The x -axis points from the interaction point to the centre of the LHC ring, and the y -axis points upwards. Cylindrical coordinates (r, ϕ) are used in the transverse plane, ϕ being the azimuthal angle around the z -axis. The pseudorapidity is defined in terms of the polar angle θ as $\eta = -\ln \tan(\theta/2)$.

3.1. Secondary vertex-based b -tagging algorithms

Physically grounded algorithms have been designed to look for displaced vertices. ATLAS uses an inclusive secondary vertex algorithm, SV1 [11], which looks for displaced tracks and groups them in a secondary vertex, and a multi-vertex finding algorithm, JETFITTER [12], which reconstructs the topology and fits the entire decay chain along the hadron line of flight.

3.2. Impact parameter-based b -tagging algorithms

Another low-level b -tagging approach exploits the impact parameter properties of tracks and how they differ in the heavy flavor and light flavor scenarios. The goal is to compute the likelihood of tracks associated to jets originating at the primary vertex.

We define the transverse impact parameter, d_0 , as the distance of closest approach of the track to the primary vertex in the $r - \phi$ plane (the plane perpendicular to the beam direction \hat{z}), and the longitudinal impact parameter, z_0 , as the distance from the primary vertex to the intersection of the track projection onto the longitudinal plane with the beam axis [13]. The sign of the impact parameters indicates whether the track crosses the jet axis in front of or behind the primary vertex with respect to the jet direction of flight. While tracks in light jets have impact parameter significances that are close to zero, those generated from b -hadron decays and consistent with the secondary vertex hypothesis are more likely to deviate from zero [13].

The tracks used for impact parameter-based tagging satisfy:

- $p_T^{\text{track}} > 1 \text{ GeV}$
- $|d_0| < 1 \text{ mm}$ and $|z_0 \sin \theta| < 1.5 \text{ mm}$
- $N_{\text{Si}}^{\text{hits}} \geq 7$ and $N_{\text{Si}}^{\text{holes}} \leq 2$ with $N_{\text{pixel}}^{\text{holes}} \leq 1$

where a hole is defined as a missing expected hit and the subscripts on N identify different inner detector components.

3.2.1. IP3D Traditionally, impact parameter tagging uses a binned 2D likelihood method in different track grade categories using the transverse and longitudinal impact parameter significances. The IP3D tagger is based on a log likelihood ratio discriminant, computed – assuming no correlation among tracks – as the following sum of per-track contributions:

$$\text{IP3D LLR} = \sum_{i=1}^N \log \frac{p_{b_i}}{p_{u_i}} \quad (1)$$

where N is the number of tracks associated to the jet, and p_b, p_u are the template PDFs for the b and light flavor hypotheses.

However, in a b -hadron decay, several charged particles with large impact parameters can emerge from the secondary (or tertiary) vertex. In this scenario, the tracks impact parameters are intrinsically correlated: once a track with large impact parameter is found, finding a second track with large impact parameter becomes more likely. On the other hand, such a correlation should not be observed if there is no displaced decay, which is often the case for light-flavor jets. IP3D’s treatment of tracks as independent limits its power as its formulation (Eq. 1) disregards known track correlations.

3.2.2. RNNIP The traditional IP3D tagger has been augmented by the introduction of RNNIP [14], a Recurrent Neural Network-based b -tagging classifier.

Recurrent Neural Networks (RNNs) are generally used to learn sequence-based dependencies for arbitrary-length input sequences. The fundamental unit is a cell that holds an internal state vector. To allow information to persist, RNNs use loops that feed the output of the

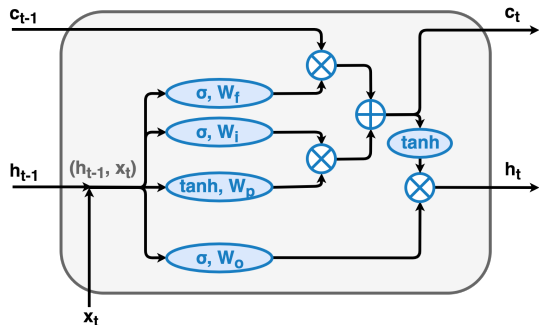


Figure 2: Anatomy of an LSTM cell with output, input, and forget gates.

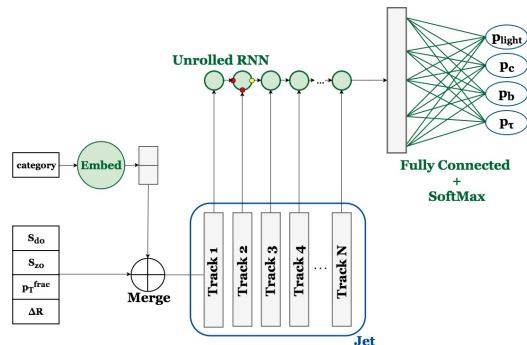


Figure 3: Diagram of the RNNIP tagger from input (left) to output (right).

cell back into the cell itself at the following step. Much of the recent success of RNNs comes from the formulation of Long Short-Term Memory [15] (LSTM) units and later variants such as Gated Recurrent Units [16] (GRUs). These modifications at the cell level mitigate issues related to vanishing and exploding gradients, to improve the knowledge persistence of long-term dependencies. They do so with different internal gating mechanisms that modify the cell state in order to regulate the relative importance of long-term versus short-term information, allowing the unit to learn when to read, write and reset the memory. The inner structure of an LSTM unit is depicted in Figure 2, while Figure 3 shows its usage in the context of the tagger.

This architecture is the natural candidate for inputs in sequence format. RNNIP represents jets as a sequence of tracks ordered by absolute transverse impact parameter significance, as track ordering studies suggest. Each track is described by a vector of features: transverse and longitudinal impact parameter significances, jet p_T fraction carried by the track, distance between the track and the jet axis, and a learned 2D embedding of the track quality.

This four-class tagger predicts a jet flavor probability associated with each flavor (b , c , τ , and light), which could be used to select a class using a maximum a posteriori estimation. For the purpose of calibrating multiple efficiency points, this information is instead combined into the discriminant

$$\text{RNNIP}(f_c, f_\tau) = \log \left(\frac{p_b}{f_c \cdot p_c + f_\tau \cdot p_\tau + (1 - f_c - f_\tau) \cdot p_u} \right) \quad (2)$$

where each background importance can be varied by varying their fraction (f_c , f_τ) after training.

3.3. Muon-based b -tagging algorithm

The Soft Muon Tagger [7] (SMT) relies on information from the reconstructed muons from semi-leptonic decays of heavy-flavor hadrons. SMT provides orthogonal information to impact parameter- and vertex-based taggers. This tagger uses a group of 6 hand-engineered variables and combines them in a boosted decision tree to obtain a discriminant.

4. High-level taggers

4.1. MV2

The recently improved MV2 tagger is a gradient boosted decision tree trained using the ROOT TMVA library [17] to perform a binary classification of jets into b and non- b . Three variants are investigated which differ in the quantity of input variables.

MV2: this baseline model uses 24 input variables from lower level taggers (IP2D, IP3D, SV1, JETFITTER) and kinematic properties. The full list is available in Ref. [18];

MV2Mu: adds the output of the SMT BDT to the list of inputs;

MV2MuRnn: adds the output of RNNIP to the list of inputs.

Importance sampling is performed to achieve a flat 2D distribution in p_T and η . The hyper-parameters of the model were optimized for the 2016 Run [13] to 1000 trees with maximum depth of 30, and minimum node size of 0.05% of the training sample.

4.2. DL1

DL1 [19] is deep neural network whose architecture is a mixture of fully-connected, maxout [20], and batch normalization [21] layers with rectified linear units [22] as activation functions. Dropout [23] is used as a stochastic regularization technique.

As with RNNIP, DL1 is trained using Keras [24], and incorporated into the ATLAS Athena code base [25] via the custom C++ network evaluator provided through the LWTNN library [26]. `rootpy` [27] and `root_numpy` [28] were used to convert the input file format.

The network is trained to minimize the cross-entropy loss using the Adam optimizer [29]. A thorough grid search is performed to select the network structure and hyper-parameters such as the number of hidden layers, number of nodes per layer, and the learning rate.

The three output nodes calculate the probabilities associated with each jet flavor (here: p_b , p_c , p_u) as highly non-linear functions of the input features. A final discriminant for b -tagging is obtained by combining the outputs into one tunable function of the fraction of c -jets in the background, f_c :

$$\text{DL1}(f_c) = \log \left(\frac{p_b}{f_c \cdot p_c + (1 - f_c) \cdot p_u} \right) \quad (3)$$

A similar equation can be derived for c -tagging, by swapping the b and c labels. Although the c -fraction can be varied depending on the background composition of each analysis, in order to facilitate the comparison with MV2, f_c is herein set to the natural occurring fraction in a typical $t\bar{t}$ sample, *i.e.* 7%. This parameter regulates the trade-off between c and light rejections at a given b tagging efficiency.

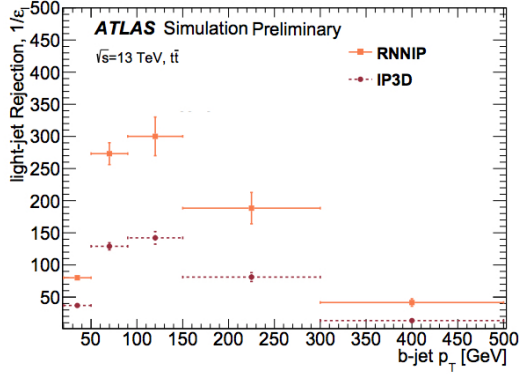
While DL1 and MV2 have been found to have similar performance, DL1 has emerged out of need for improved flexibility and power in future R&D cycles. The neural network can be easily trained with minimal standalone code; it is GPU-enabled; it is modular and can be effortlessly extended to include more input variables. It is also amenable to be trained using an adversarial loss to minimize data-Monte Carlo discrepancies [30], and it grants the possibility for future end-to-end trainings in conjunction with RNNIP.

5. Performance in Simulation

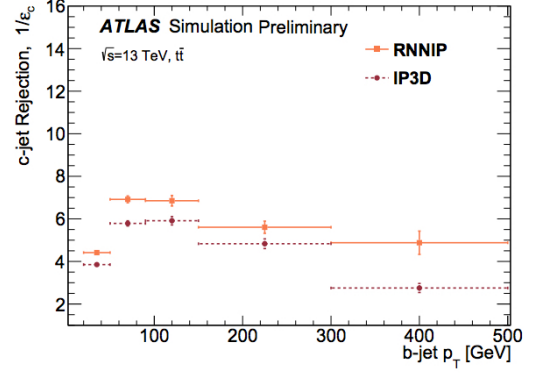
The high level taggers are trained on a newly introduced hybrid sample which combines $t\bar{t}$ with a broad Z' sample at $m_{Z'} = 4$ TeV to increase the representation of high p_T jets at training time.

Figure 4 shows that RNNIP improves light rejection by a factor of 2 and c rejection by a factor of 1.2 compared to IP3D, and outperforms it across the entire p_T spectrum. A sizable fraction of b -jets is tagged correctly by only one of the two algorithms, where IP3D is more efficient at low track multiplicity and short decay length, while RNNIP is superior at high multiplicity.

The three MV2 versions are evaluated on exclusive $t\bar{t}$ and Z' samples; results from the latter are presented in Figure 5. Including muon information improves b -to-light jet separation, but the improvement vanishes at high p_T . On the other hand, b -to- c separation does not strongly benefit from the muon-in-jet information for efficiencies $> 50\%$ because of lepton presence in both b and c decays. MV2MuRNN further improves performance, especially in the high- p_T region where the RNNIP tagger is most performant. For example, when evaluating on Z' events (Figure 5), at $p_T^{\text{jet}} = 1$ TeV the RNN improves the light-jet rejection by a factor ~ 1.5 and the c -jet rejection by a factor of ~ 1.3 , compared to the baseline MV2 2017 configuration.

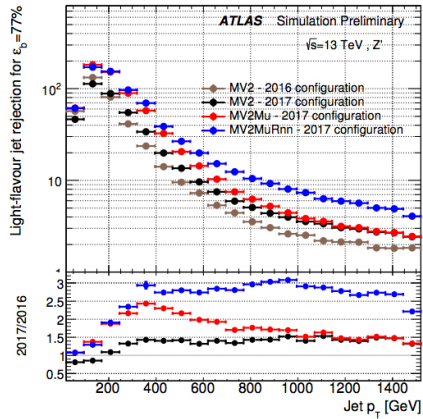


(a)




(b)

Figure 4: Performance comparison in bins of p_T , as it appeared in Ref. [7], between IP3D and RNNIP for light jet (a) and c -jet rejection (b) in simulated $t\bar{t}$ events at a constant b -efficiency of 70%.



(a)



(b)

Figure 5: Performance comparison in bins of p_T , as it appeared in Ref. [7], between MV2 in the 2016 configuration (brown), MV2 in the 2017 configuration (black), MV2Mu (red) and MV2MuRNN (blue) for light jet (a) and c -jet rejection (b) in simulated Z' events at a constant b -efficiency of 77%. The ratio plots use the 2016 MV2 configuration as baseline.

6. Conclusion

ATLAS developed new algorithms to identify b -jets using machine learning techniques. Their performance was investigated using Monte Carlo simulation of $p\bar{p}$ events with $\sqrt{s} = 13$ TeV in the ATLAS detector. These methods augment the already rich suite of b -tagging algorithms adopted by the ATLAS collaboration, by extracting complementary features and enhancing the overall flavor tagging performance.

ATLAS analysts can also expect an increased level of agreement between the train (Monte Carlo) and test (data) domains due to improvements in the tracking simulation with respect to the 2016 configuration, which reduces the performance losses associated with out-of-domain model application.

Acknowledgements

This work was supported in part by the Office of High Energy Physics of the U.S. Department of Energy under contract DE-FG02-92ER40704.

References

- [1] Evans L and Bryant P 2008 *JINST* **3** S08001
- [2] Mitrevski J 2015 *Journal of Physics: Conference Series* **664** 072034
- [3] ATLAS Collaboration 2008 *JINST* **3** S08003
- [4] ATLAS Collaboration 2010 *The European Physical Journal C* **70** 787–821 ISSN 1434-6052
- [5] 1996 ATLAS liquid argon calorimeter: Technical design report Tech. Rep. CERN-LHCC-96-41 CERN Geneva
- [6] 1996 ATLAS tile calorimeter: Technical design report Tech. Rep. CERN-LHCC-96-42 CERN Geneva
- [7] 2017 Optimisation and performance studies of the ATLAS b -tagging algorithms for the 2017-18 LHC run Tech. Rep. ATL-PHYS-PUB-2017-013 CERN Geneva URL <http://cds.cern.ch/record/2273281>
- [8] Cacciari M, Salam G P and Soyez G 2008 *Journal of High Energy Physics* **2008** 063
- [9] 2014 Tagging and suppression of pileup jets with the ATLAS detector Tech. Rep. ATLAS-CONF-2014-018 CERN Geneva URL <https://cds.cern.ch/record/1700870>
- [10] ATLAS Collaboration 2016 *JINST* **11** P04008
- [11] 2017 Secondary vertex finding for jet flavour identification with the ATLAS detector Tech. Rep. ATL-PHYS-PUB-2017-011 CERN Geneva URL <http://cds.cern.ch/record/2270366>
- [12] Piacquadio G and Weiser C 2008 *Journal of Physics: Conference Series* **119** 032032
- [13] 2016 Optimisation of the ATLAS b -tagging performance for the 2016 LHC Run Tech. Rep. ATL-PHYS-PUB-2016-012 CERN Geneva URL <http://cds.cern.ch/record/2160731>
- [14] 2017 Identification of Jets Containing b -Hadrons with Recurrent Neural Networks at the ATLAS Experiment Tech. Rep. ATL-PHYS-PUB-2017-003 CERN Geneva URL <http://cds.cern.ch/record/2255226>
- [15] Hochreiter S and Schmidhuber J 1997 *Neural Comput.* **9** 1735–1780 ISSN 0899-7667
- [16] Chung J, Gulcehre C, Cho K and Bengio Y 2014 (*Preprint* [cs.NE/1412.3555](https://arxiv.org/abs/1412.3555))
- [17] Hoecker A, Speckmayer P, Stelzer J, Therhaag J, von Toerne E and Voss H 2007 *PoS ACAT* 040 (*Preprint* [physics/0703039](https://arxiv.org/abs/0703039))
- [18] 2015 Expected performance of the ATLAS b -tagging algorithms in Run-2 Tech. Rep. ATL-PHYS-PUB-2015-022 CERN Geneva URL <https://cds.cern.ch/record/2037697>
- [19] ATLAS Collaboration 2017 URL <http://cds.cern.ch/record/2274065>
- [20] Goodfellow I J, Warde-Farley D, Mirza M, Courville A and Bengio Y 2013 Maxout networks *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28* ICML'13 (JMLR.org) pp III-1319–III-1327
- [21] Ioffe S and Szegedy C 2015 Batch normalization: Accelerating deep network training by reducing internal covariate shift *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37* ICML'15 (JMLR.org) pp 448–456
- [22] Nair V and Hinton G E 2010 Rectified linear units improve restricted boltzmann machines *Proceedings of the 27th International Conference on International Conference on Machine Learning* ICML'10 (USA: Omnipress) pp 807–814 ISBN 978-1-60558-907-7
- [23] Srivastava N *et al.* 2014 *J. Mach. Learn. Res.* **15** 1929–1958 ISSN 1532-4435
- [24] Chollet F *et al.* 2015 Keras URL <https://github.com/fchollet/keras>
- [25] 2005 ATLAS Computing: technical design report URL <http://cds.cern.ch/record/837738>
- [26] Guest D H, Smith J W, Paganini M, Kagan M, Lanfermann M and Krasznahorkay A 2017 lwttn/lwttn: Version 2.4 DOI: [10.5281/zenodo.848256](https://doi.org/10.5281/zenodo.848256)
- [27] Dawe N, Waller P, Friis E K *et al.* 2015 rootpy: 0.8.0 DOI: [10.5281/zenodo.18897](https://doi.org/10.5281/zenodo.18897)
- [28] Dawe N, Ongmongkolkul P and Stark G 2017 *The Journal of Open Source Software* **2**
- [29] Kingma D P and Ba J 2014 (*Preprint* [cs.LG/1412.6980](https://arxiv.org/abs/1412.6980))
- [30] Ganin Y *et al.* 2015 (*Preprint* [stat.ML/1505.07818](https://arxiv.org/abs/1505.07818))