



JEDI-net: a jet identification algorithm based on interaction networks

Eric A. Moreno¹, Olmo Cerri¹, Javier M. Duarte^{2,3,a}, Harvey B. Newman¹, Thong Q. Nguyen¹, Avikar Periwal¹, Maurizio Pierini⁴, Aidana Serikova¹, Maria Spiropulu¹, Jean-Roch Vlimant¹

¹ California Institute of Technology, Pasadena, CA 91125, USA

² Fermi National Accelerator Laboratory (FNAL), Batavia, IL 60510, USA

³ University of California San Diego, La Jolla, CA 92093, USA

⁴ European Center for Nuclear Research (CERN), 1211 Geneva, Switzerland

Received: 10 September 2019 / Accepted: 3 January 2020 / Published online: 25 January 2020

© The Author(s) 2020

Abstract We investigate the performance of a jet identification algorithm based on interaction networks (JEDI-net) to identify all-hadronic decays of high-momentum heavy particles produced at the LHC and distinguish them from ordinary jets originating from the hadronization of quarks and gluons. The jet dynamics are described as a set of one-to-one interactions between the jet constituents. Based on a representation learned from these interactions, the jet is associated to one of the considered categories. Unlike other architectures, the JEDI-net models achieve their performance without special handling of the sparse input jet representation, extensive pre-processing, particle ordering, or specific assumptions regarding the underlying detector geometry. The presented models give better results with less model parameters, offering interesting prospects for LHC applications.

1 Introduction

Jets are collimated cascades of particles produced at particle accelerators. Quarks and gluons originating from hadron collisions, such as the proton-proton collisions at the CERN Large Hadron Collider (LHC), generate a cascade of other particles (mainly other quarks or gluons) that then arrange themselves into hadrons. The stable and unstable hadrons' decay products are observed by large particle detectors, reconstructed by algorithms that combine the information from different detector components, and then clustered into jets, using physics-motivated sequential recombination algorithms such as those described in Refs. [1–3]. Jet identification, or *tagging*, algorithms are designed to identify the nature of the particle that initiated a given cascade, inferring

it from the collective features of the particles generated in the cascade.

Traditionally, jet tagging was meant to distinguish three classes of jets: light flavor quarks $q = u, d, s, c$, gluons g , or bottom quarks (b). At the LHC, due to the large collision energy, new jet topologies emerge. When heavy particles, e.g. W, Z , or Higgs (H) bosons or the top quark, are produced with large momentum and decay to all-quark final states, the resulting jets are contained in a small solid angle. A single jet emerges from the overlap of two (for bosons) or three (for the top quark) jets, as illustrated in Fig. 1. These jets are characterized by a large invariant mass (computed from the sum of the four-momenta of their constituents) and they differ from ordinary quark and gluon jets, due to their peculiar momentum flow around the jet axis.

Several techniques have been proposed to identify these jets by using physics-motivated quantities, collectively referred to as “jet substructure” variables. A review of the different techniques can be found in Ref. [4]. As discussed in the review, approaches based on deep learning (DL) have been extensively investigated (see also Sect. 2), processing sets of physics-motivated quantities with dense layers or raw data representations (e.g. jet images or particle feature lists) with more complex architectures (e.g. convolutional or recurrent networks).

In this work, we compare the typical performance of some of these approaches to what is achievable with a novel jet identification algorithm based on an interaction network (JEDI-net). Interaction networks [5] (INs) were designed to decompose complex systems into distinct objects and relations, and reason about their interactions and dynamics. One of the first uses of INs was to predict the evolution of physical systems under the influence of internal and external forces, for example, to emulate the effect of gravitational interactions

^a e-mail: javier.mauricio.duarte@cern.ch




Fig. 1 Pictorial representations of the different jet categories considered in this paper. Left: jets originating from quarks or gluons produce one cluster of particles, approximately cone-shaped, developing along the flight direction of the quark or gluon that started the cascade. Cen-

ter: when produced with large momentum, a heavy boson decaying to quarks would result in a single jet, made of 2 particle clusters (usually referred to as prongs). Right: a high-momentum $t \rightarrow Wb \rightarrow q\bar{q}b$ decay chain results in a jet composed of three prongs

in n -body systems. The n -body system is represented as a set of objects subject to one-on-one interactions. The n bodies are embedded in a graph and these one-on-one interaction functions, expressed as trainable neural networks, are used to predict the post-interaction status of the n -body system. We study whether this type of network generalizes to a novel context in high energy physics. In particular, we represent a jet as a set of particles, each of which is represented by its momentum and embedded as a vertex in a fully-connected graph. We use neural networks to learn a representation of each one-on-one particle *interaction*¹ in the jet, which we then use to define jet-related high-level features (HLFs). Based on these features, a classifier associates each jet to one of the five categories shown in Fig. 1.

For comparison, we consider other classifiers based on different architectures: a dense neural network (DNN) [6] receiving a set of jet-substructure quantities, a convolutional neural network (CNN) [7–9] receiving an image representation of the transverse momentum (p_T) flow in the jet,² and a recurrent neural network (RNN) with gated recurrent units [10] (GRUs), which process a list of particle features. These models can achieve state-of-the-art performance although they require additional ingredients: the DNN model requires processing the constituent particles to pre-compute HLFs, the GRU model assumes an ordering criterion for the input particle feature list, and the CNN model requires representing the jet as a rectangular, regular, pixelated image.

¹ Here, we refer to the abstract message-passing interaction represented by the edges of the graph and not the physical interactions due to quantum chromodynamics, which occur before the jet constituents emerge from the hadronization process.

² We use a Cartesian coordinate system with the z axis oriented along the beam axis, the x axis on the horizontal plane, and the y axis oriented upward. The x and y axes define the transverse plane, while the z axis identifies the longitudinal direction. The azimuthal angle ϕ is computed from the x axis. The polar angle θ is used to compute the pseudorapidity $\eta = -\log(\tan(\theta/2))$. We use natural units such that $c = \hbar = 1$ and we express energy in units of electronVolt (eV) and its prefix multipliers.

Any of these aspects can be handled in a reasonable way (e.g. one can use a jet clustering metric to order the particles), sometimes sacrificing some detector performance (e.g., with coarser image pixels than realistic tracking angular resolution, in the case of many models based on CNN). It is then worth exploring alternative solutions that could reach state-of-the-art performance without making these assumptions. In particular, it is interesting to consider architectures that directly takes as input jet constituents and are invariant for their permutation. This motivated the study of jet taggers based on recursive [11], graph networks [12, 13], and energy flow networks [14]. In this context, we aim to investigate the potential of INs.

This paper is structured as follows: we provide a list of related works in Sect. 2. In Sect. 3, we describe the utilized data set. The structure of the JEDI-net model is discussed in Sect. 4 together with the alternative architectures considered for comparison. Results are shown in Sect. 5. Sections 6 and 7 discuss what the JEDI-net learns when processing the graph and quantify the amount of resources needed by the tagger, respectively. We conclude with a discussion and outlook for this work in Sect. 8. “Appendix A” describes the design and optimization of the alternative models.

2 Related work

Jet tagging is one of the most popular LHC-related tasks to which DL solutions have been applied. Several classification algorithms have been studied in the context of jet tagging at the LHC [15–22] using DNNs, CNNs, or physics-inspired architectures. Recurrent and recursive layers have been used to construct jet classifiers starting from a list of reconstructed particle momenta [11–13]. Recently, these different approaches, applied to the specific case of top quark jet identification, have been compared in Ref. [23]. While many of these studies focus on data analysis, work is underway

to apply these algorithms in the early stages of LHC real-time event processing, i.e. the trigger system. For example, Ref. [24] focuses on converting these models into firmware for field programmable gate arrays (FPGAs) optimized for low latency (less than 1 μ s). If successful, such a program could allow for a more resource-efficient and effective event selection for future LHC runs.

Graph neural networks have also been considered as jet tagging algorithms [25, 26] as a way to circumvent the sparsity of image-based representations of jets. These approaches demonstrate remarkable categorization performance. Motivated by the early results of Ref. [25], graph networks have been also applied to other high energy physics tasks, such as event topology classification [27, 28], particle tracking in a collider detector [29], pileup subtraction at the LHC [30], and particle reconstruction in irregular calorimeters [31].

3 Data set description

This study is based on a data set consisting of simulated jets with an energy of $p_T \approx 1$ TeV, originating from light quarks q , gluons g , W and Z bosons, and top quarks produced in $\sqrt{s} = 13$ TeV proton-proton collisions. The data set was created using the configuration and parametric description of an LHC detector described in Refs. [24, 32], and is available on the Zenodo platform [33–36].

Jets are clustered from individual reconstructed particles, using the anti- k_T algorithm [3, 37] with jet-size parameter $R = 0.8$. Three different jet representations are considered:

- A list of 16 HLFs, described in Ref. [24], given as input to a DNN. The 16 distributions are shown in Fig. 2 for the five jet classes.
- An image representation of the jet, derived by considering a square with pseudorapidity and azimuth distances $\Delta\eta = \Delta\phi = 2R$, centered along the jet axis. The image is binned into 100×100 pixels. Such a pixel size is comparable to the cell of a typical LHC electromagnetic calorimeter, but much coarser than the typical angular resolution of a tracking device for the p_T values relevant to this task. Each pixel is filled with the scalar sum of the p_T of the particles in that region. These images are obtained by considering the 150 highest- p_T constituents for each jet. This jet representation is used to train a CNN classifier. The average jet images for the five jet classes are shown in Fig. 3. For comparison, a randomly chosen set of images is shown in Fig. 4.
- A constituent list for up to 150 particles, in which each particle is represented by 16 features, computed from the particle four-momenta: the three Cartesian coordinates of the momentum (p_x , p_y , and p_z), the absolute energy E , p_T , the pseudorapidity η , the azimuthal angle ϕ , the

distance $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$ from the jet center, the relative energy $E^{\text{rel}} = E^{\text{particle}}/E^{\text{jet}}$ and relative transverse momentum $p_T^{\text{rel}} = p_T^{\text{particle}}/p_T^{\text{jet}}$ defined as the ratio of the particle quantity and the jet quantity, the relative coordinates $\eta^{\text{rel}} = \eta^{\text{particle}} - \eta^{\text{jet}}$ and $\phi^{\text{rel}} = \phi^{\text{particle}} - \phi^{\text{jet}}$ defined with respect to the jet axis, $\cos\theta$ and $\cos\theta^{\text{rel}}$ where $\theta^{\text{rel}} = \theta^{\text{particle}} - \theta^{\text{jet}}$ is defined with respect to the jet axis, and the relative η and ϕ coordinates of the particle after applying a proper Lorentz transformation (rotation) as described in Ref. [38]. Whenever less than 150 particles are reconstructed, the list is filled with zeros. The distributions of these features considering the 150 highest- p_T particles in the jet are shown in Fig. 5 for the five jet categories. This jet representation is used for a RNN with a GRU layer and for JEDI-net.

4 JEDI-net

In this work, we apply an IN [5] architecture to learn a representation of a given input graph (the set of constituents in a jet) and use it to accomplish a classification task (tagging the jet). One can see the IN architecture as a processing algorithm to learn a new representation of the initial input. This is done replacing a set of input features, describing each individual vertex of the graph, with a set of engineered features, specific of each vertex but whose values depend on the connection between the vertices in the graph.

The starting point consists of building a graph for each input jet. The N_O particles in the jet are represented by the vertices of the graph, fully interconnected through directional edges, for a total of $N_E = N_O \times (N_O - 1)$ edges. An example is shown in Fig. 6 for the case of a three-vertex graph. The vertices and edges are labeled for practical reasons, but the network architecture ensures that the labeling convention plays no role in creating the new representation.

Once the graph is built, a receiving matrix (R_R) and a sending matrix (R_S) are defined. Both matrices have dimensions $N_O \times N_E$. The element $(R_R)_{ij}$ is set to 1 when the i^{th} vertex receives the j^{th} edge and is 0 otherwise. Similarly, the element $(R_S)_{ij}$ is set to 1 when the i^{th} vertex sends the j^{th} edge and is 0 otherwise. In the case of the graph of Fig. 6, the two matrices take the form:

$$R_S = \begin{pmatrix} E_1 & E_2 & E_3 & E_4 & E_5 & E_6 \\ O_1 & 0 & 0 & 0 & 1 & 1 & 0 \\ O_2 & 1 & 0 & 0 & 0 & 0 & 1 \\ O_3 & 0 & 1 & 1 & 0 & 0 & 0 \end{pmatrix} \quad (1)$$

$$R_R = \begin{pmatrix} E_1 & E_2 & E_3 & E_4 & E_5 & E_6 \\ O_1 & 1 & 1 & 0 & 0 & 0 & 0 \\ O_2 & 0 & 0 & 1 & 1 & 0 & 0 \\ O_3 & 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix}. \quad (2)$$

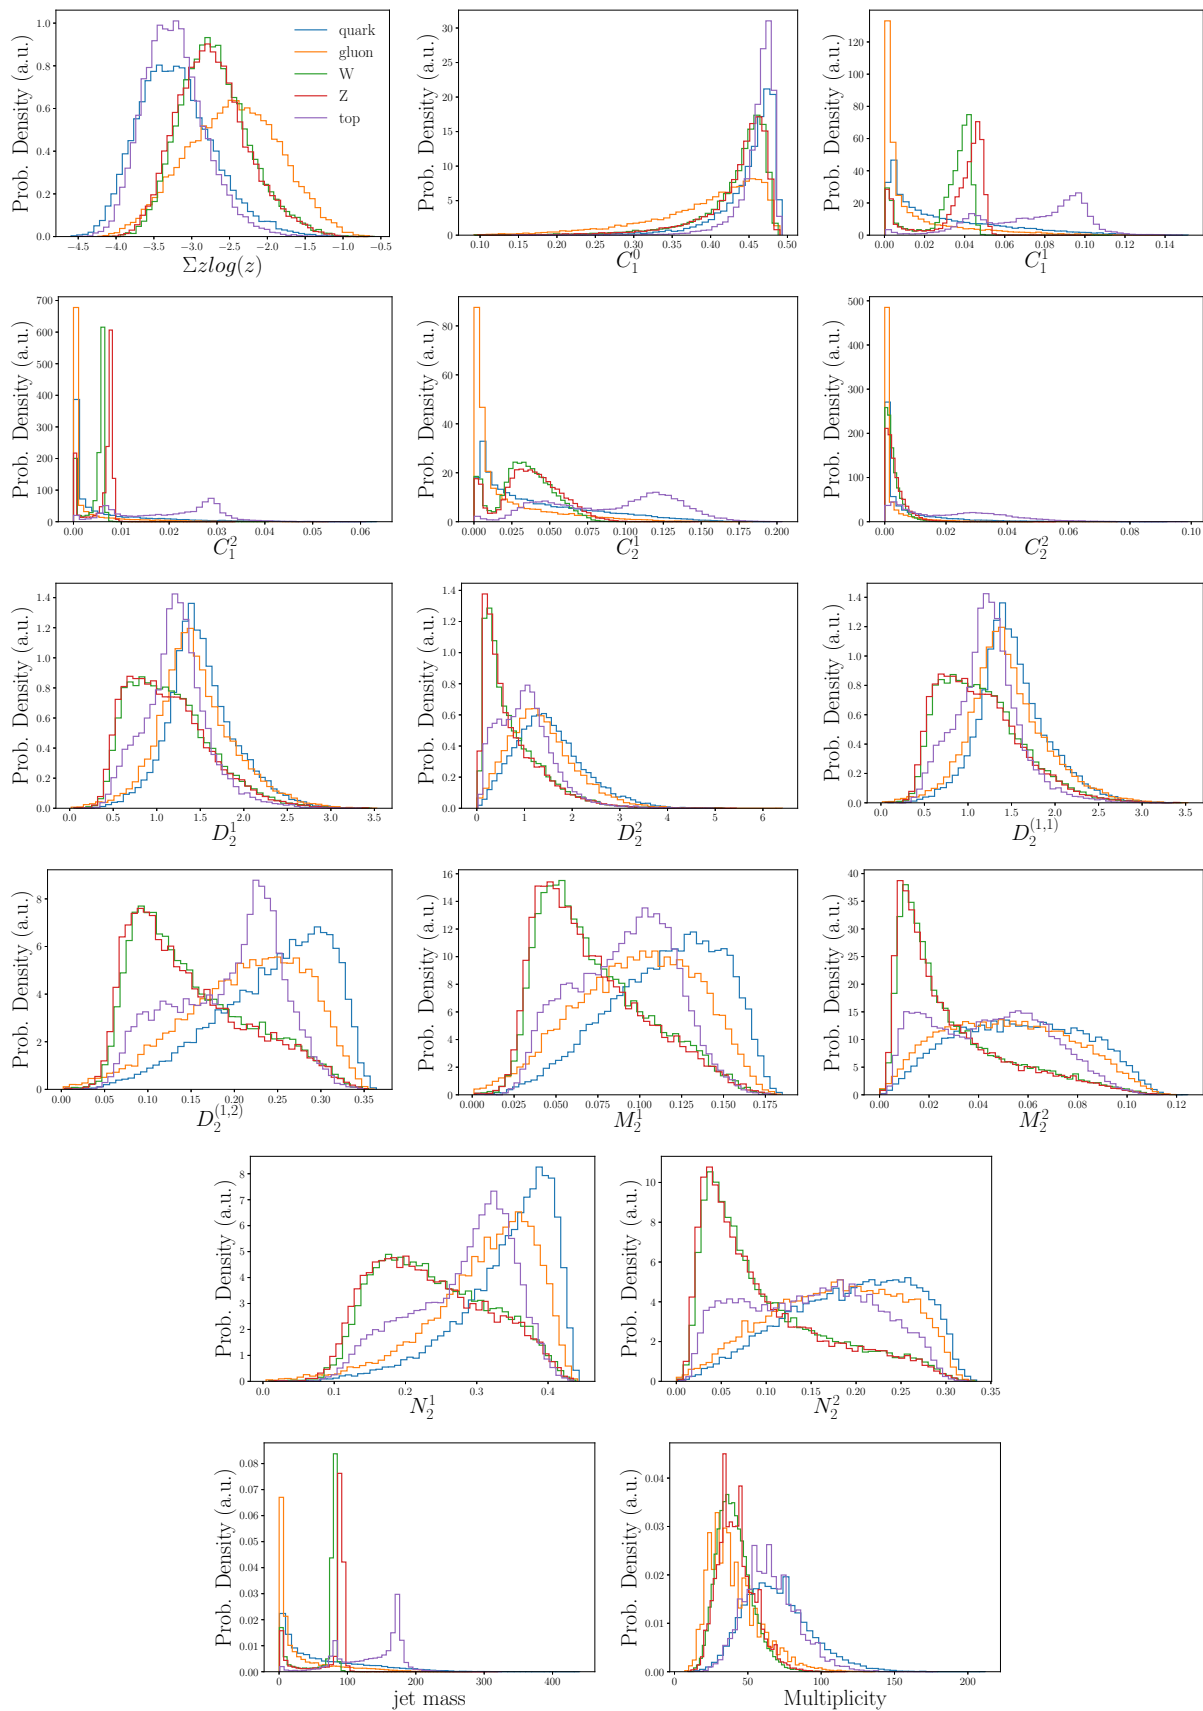


Fig. 2 Distributions of the 16 high-level features used in this study, described in Ref. [24]




Fig. 3 Average 100×100 images for the five jet classes considered in this study: q (top left), g (top center), W (top right), Z (bottom left), and top jets (bottom right). The temperature map represents the amount of

p_T collected in each cell of the image, measured in GeV and computed from the scalar sum of the p_T of the particles pointing to each cell

The input particle features are represented by an input matrix I . Each column of the matrix corresponds to one of the graph vertices, while the rows correspond to the P features used to represent each vertex. In our case, the vertices are the particles inside the jet, each represented by its array of features (i.e., the 16 features shown in Fig. 5). Therefore, the I matrix has dimensions $P \times N_O$.

The I matrix is processed by the IN in a series of steps, represented in Fig. 7. The I matrix is multiplied by the R_R and R_S matrices and the two resulting matrices are then concatenated to form the B matrix, having dimension $2P \times N_E$:

$$B = \begin{pmatrix} I \times R_R \\ I \times R_S \end{pmatrix}. \quad (3)$$

Each column of the B matrix represents an edge, i.e. a particle-to-particle interaction. The $2P$ elements of each column are the features of the sending and receiving vertices for that edge. Using this information, a D_E -dimensional hidden representation of the interaction edge is created through a trainable function $f_R : \mathbb{R}^{2P} \mapsto \mathbb{R}^{D_E}$. This gives a matrix E with dimensions $D_E \times N_E$. The cumulative effects of the interactions received by a given vertex are gathered by summing the D_E hidden features over the edges arriving to it. This is done by computing $\bar{E} = ER_R^T$ with dimensions

$D_E \times N_O$, which is then appended to the initial input matrix I :

$$C = \begin{pmatrix} I \\ \bar{E} \end{pmatrix}. \quad (4)$$

At this stage, each column of the C matrix represents a constituent in the jet, expressed as a $(P + D_E)$ -dimensional feature vector, containing the P input features and the D_E hidden features representing the combined effect of the interactions with all the connected particles. A trainable function $f_O : \mathbb{R}^{P+D_E} \mapsto \mathbb{R}^{D_O}$ is used to build a post-interaction representation of each jet constituent. The function f_O is applied to each column of C to build the post-interaction matrix O with dimensions $D_O \times N_O$.

A final classifier ϕ_C takes as input the elements of the O matrix and returns the probability for that jet to belong to each of the five categories. This is done in two ways: (i) in one case, we define the quantities $\bar{O}_i = \sum_j O_{ij}$, where j is the index of the vertex in the graph (the particle, in our case), and the $i \in [0, D_E]$ index runs across the D_E outputs of the f_O function. The \bar{O} quantities are used as input to $\phi_C : \mathbb{R}^{D_O} \mapsto \mathbb{R}^N$. This choice allows to preserve the independence of the architecture on the labeling convention adopted to build the I , R_R , and R_S matrices, at the

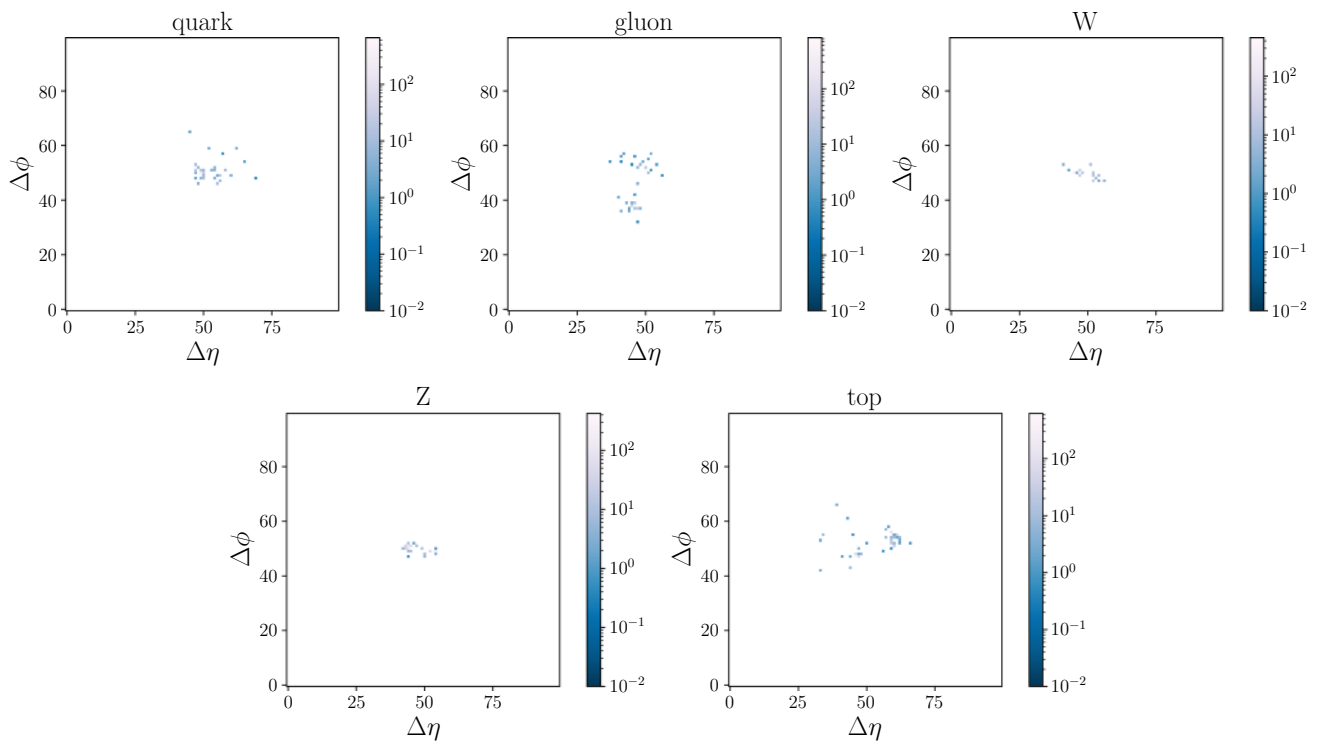


Fig. 4 Example of 100×100 images for the five jet classes considered in this study: q (top-left), g (top-right), W (center-left), Z (center-right), and top jets (bottom). The temperature map represents the amount of

p_T collected in each cell of the image, measured in GeV and computed from the scalar sum of the p_T of the particles pointing to each cell

cost of losing some discriminating information in the summation. (ii) Alternatively, the ϕ_C matrix is defined directly from the $D_O \times N_O$ elements of the O matrix, flattened into a one-dimensional array. The full information from O is preserved, but ϕ_C assumes an ordering of the N_O input objects. In our case, we rank the input particles in descending order by p_T .

The trainable functions f_O , f_R , and ϕ_C consist of three DNNs. Each of them has two hidden layers, the first (second) having N_n^1 ($N_n^2 = \lfloor N_n^1/2 \rfloor$) neurons. The model is implemented in PYTORCH [39] and trained using an NVIDIA GTX1080 GPU. The training (validation) data set consists of 630,000 (240,000) examples, while 10,000 events are used for testing purposes.

The architecture of the three trainable functions is determined by minimizing the loss function through a Bayesian optimization, using the GPyOPT library [40], based on GPy [41]. We consider the following hyperparameters:

- The number of output neurons of the f_R network, D_E (between 4 and 14).
- The number of output neurons of the f_O network, D_O (between 4 and 14).
- The number of neurons N_n^1 in the first hidden layer of the f_O , f_R , and ϕ_C network (between 5 and 50).

- The activation function for the hidden and output layers of the f_R network: ReLU [42], ELU [43], or SELU [44] functions.
- The activation function for the hidden and output layers of the f_O network: ReLU, ELU, or SELU.
- The activation function for the hidden layers of the ϕ_C network: ReLU, ELU, or SELU.
- The optimizer algorithm: Adam [45] or AdaDelta [46].

In addition, the output neurons of the ϕ_C network are activated by a softmax function. A learning rate of 10^{-4} is used. For a given network architecture, the network parameters are optimized by minimizing the categorical cross entropy. The Bayesian optimization is repeated four times. In each case, the input particles are ordered by descending p_T value and the first 30, 50, 100, or 150 particles are considered. The parameter optimization is performed on the training data set, while the loss for the Bayesian optimization is estimated on the validation data set.

Tables 1 and 2 summarize the result of the Bayesian optimization for the JEDI-net architecture with and without the sum over the columns of the O matrix, respectively. The best result of each case, highlighted in bold, is used as a reference for the rest of the paper.

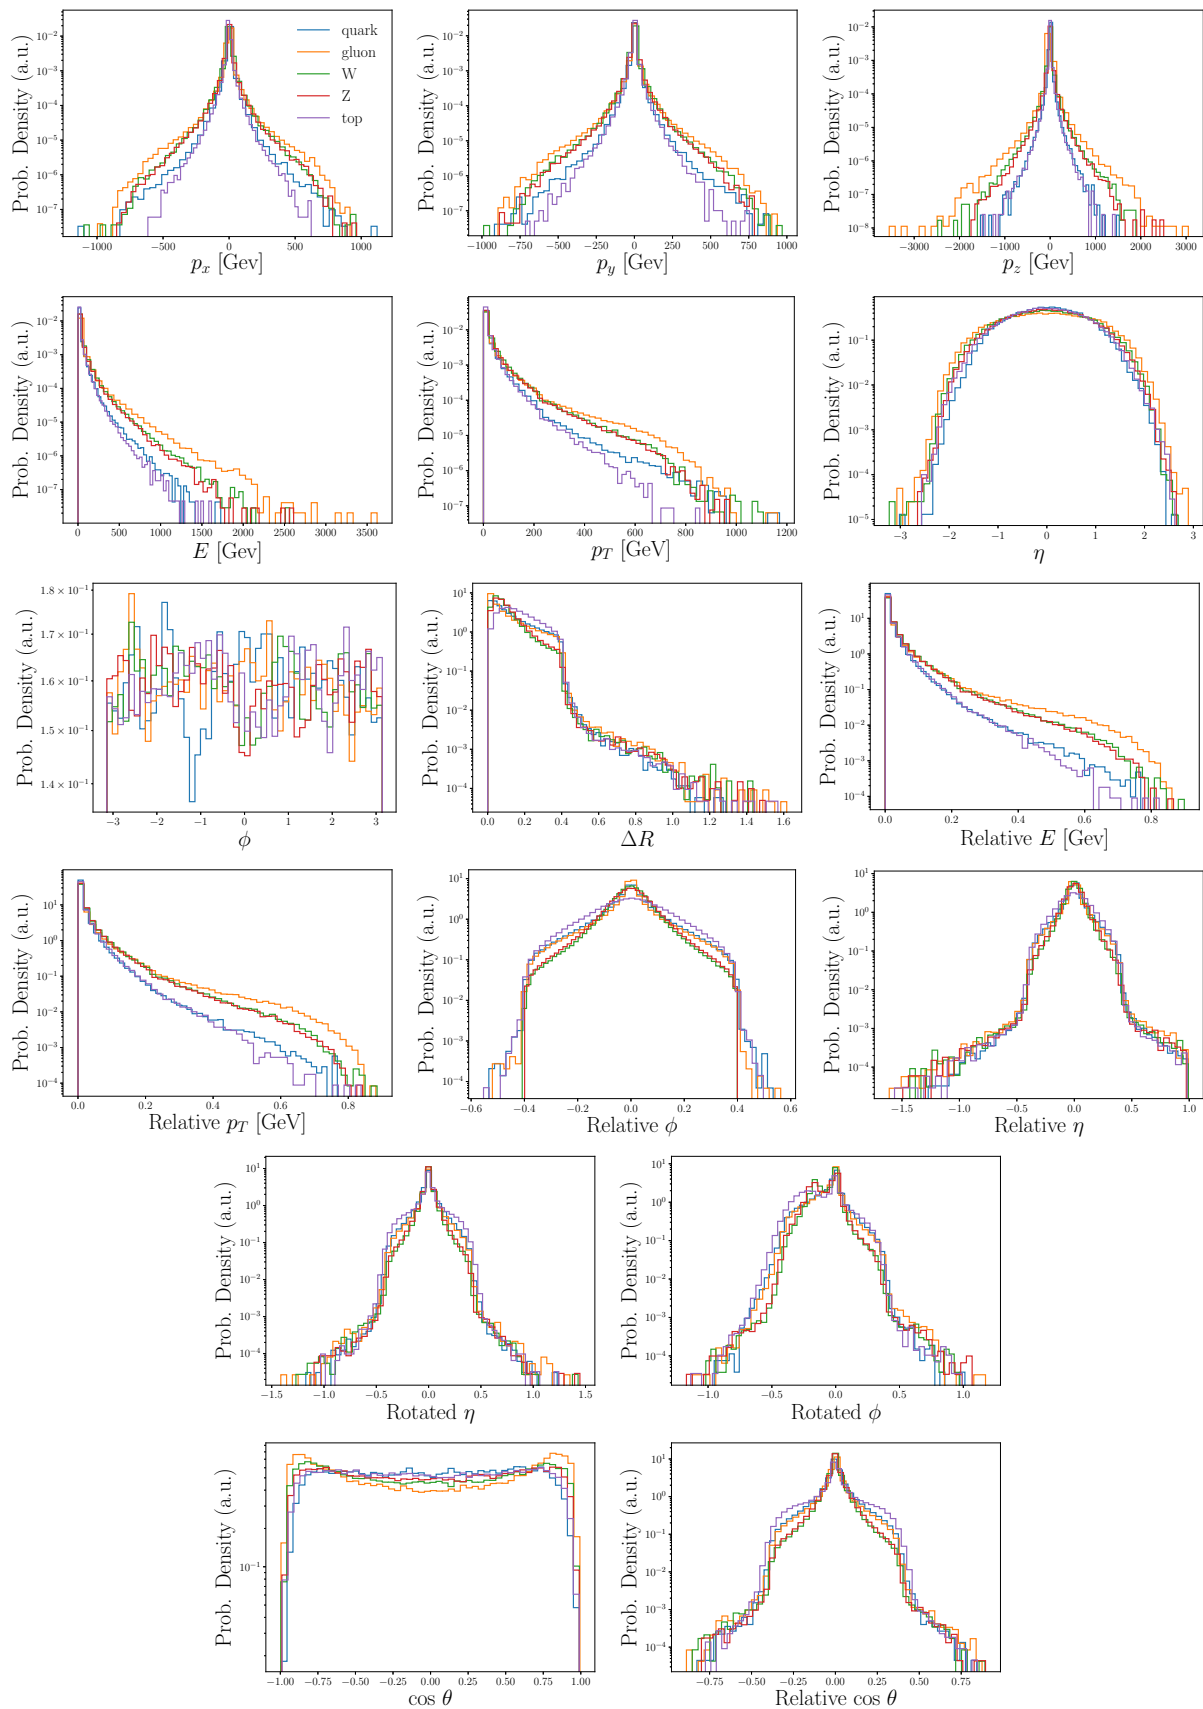


Fig. 5 Distributions of kinematic features described in the text for the 150 highest- p_T particles in each jet

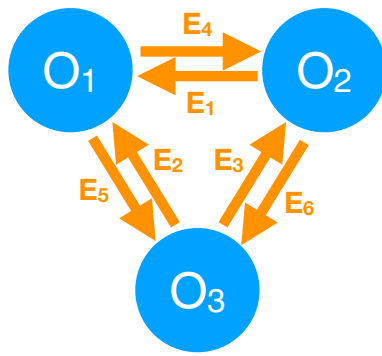


Fig. 6 An example graph with three fully connected vertices and the corresponding six edges

For comparison, three alternative models are trained on the three different representations of the same data set described in Sect. 3: a DNN model taking as input a list of HLFs, a CNN model processing jet images, and a recurrent model

applying GRUs on the same input list used for JEDI-net. The three benchmark models are optimized through a Bayesian optimization procedure, as done for the INs. Details of these optimizations and the resulting best models are discussed in “Appendix A”.

5 Results

Figure 8 shows the receiver operating characteristic (ROC) curves obtained for the optimized JEDI-net tagger in each of the five jet categories, compared to the corresponding curves for the DNN, CNN, and GRU alternative models. The curves are derived by fixing the network architectures to the optimal values based on Table 2 and “Appendix A” and performing a k -fold cross-validation training, with $k = 10$. The solid lines represent the average ROC curve, while the shaded bands quantify the ± 1 RMS dispersion. The area under the curve

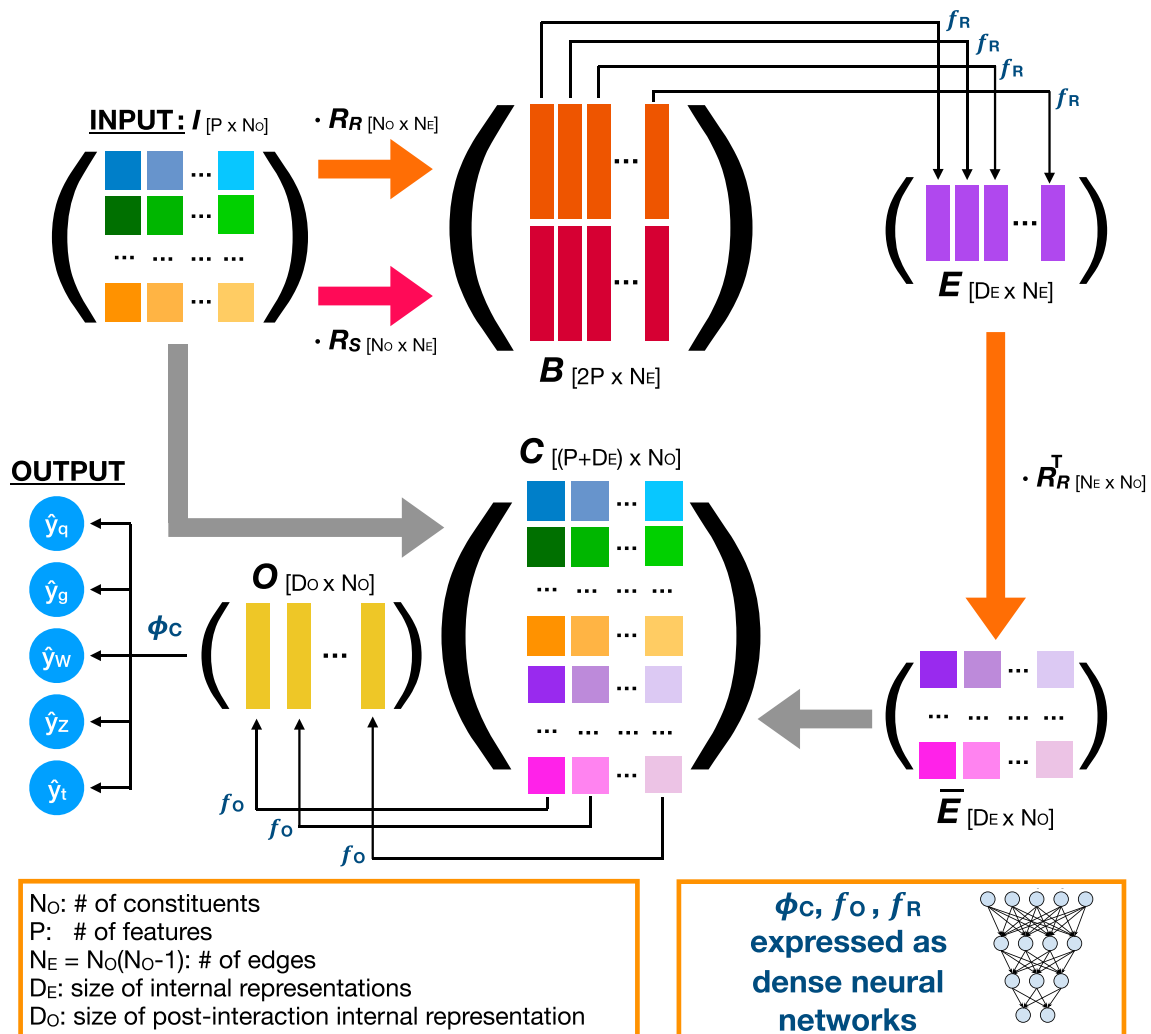


Fig. 7 A flowchart illustrating the interaction network scheme

Table 1 Optimal JEDI-net hyperparameter setting for different input data sets, when the summed \bar{O}_i quantities are given as input to the ϕ_C network. The best result, obtained when considering up to 150 particles per jet, is highlighted in bold

Hyperparameter	Number of jet constituents			
	30	50	100	150
N_n^1	6	50	30	50
D_E	8	12	4	14
D_O	6	14	4	10
f_R activation	ReLU	ReLU	SELU	SELU
f_O activation	ELU	ReLU	ReLU	SELU
ϕ_C activation	ELU	SELU	SELU	SELU
Optimizer	Adam	Adam	Adam	Adam
Optimized loss	0.84	0.58	0.62	0.55

Table 2 Optimal JEDI-net hyperparameter setting for different input data sets, when all the O_{ij} elements are given as input to the ϕ_C network. The best result, obtained when considering up to 100 particles per jet, is highlighted in bold

Hyperparameter	Number of jet constituents			
	30	50	100	150
N_n^1	50	50	30	10
D_E	12	12	10	4
D_O	6	14	10	14
f_R activation	ReLU	ELU	ELU	SELU
f_O activation	SELU	SELU	ELU	SELU
ϕ_C activation	SELU	ELU	ELU	SELU
Optimizer	Adam	Adam	Adam	Adam
Optimized loss	0.63	0.57	0.56	0.62

(AUC) values, reported in the figure, allow for a comparison of the performance of the different taggers.

The algorithm's tagging performance is quantified computing the true positive rate (TPR) values for two given reference false positive rate (FPR) values (10% and 1%). The comparison of the TPR values gives an assessment of the tagging performance in a realistic use case, typical of an LHC analysis. Table 3 shows the corresponding FPR values for the optimized JEDI-net taggers, compared to the corresponding values for the benchmark models. The largest TPR value for each class is highlighted in bold. As shown in Fig. 8 and Table 3, the two JEDI-net models outperform the other architectures in almost all cases. The only notable exception is the tight working point of the top-jet tagger, for which the DNN model gives a TPR higher by about 2%, while the CNN and GRU models give much worse performance.

The TPR values for the two JEDI-net models are within 1%. The only exception is observed for the tight working points of the W and Z taggers, for which the model using the \bar{O} sums shows a drop in TPR of $\sim 4\%$. In this respect,

the model using summed \bar{O} features is preferable (despite this small TPR loss), given the reduced model complexity (see Sect. 7) and its independence on the labeling convention for the particles embedded in the graph and for the edges connecting them.

6 What did JEDI-net learn?

In order to characterize the information learned by JEDI-net, we consider the \bar{O} sums across the N_O vertices of the graph (see Sect. 4) and we study their correlations to physics motivated quantities, typically used when exploiting jet substructure in a search. We consider the HLF quantities used for the DNN model and the N -subjettiness variables $\tau_N^{(\beta)}$ [47], computed with angular exponent $\beta = 1, 2$.

Not all the \bar{O} sums exhibit an obvious correlation with the considered quantities, i.e., the network engineers high-level features that encode other information than what is used, for instance, in the DNN model.

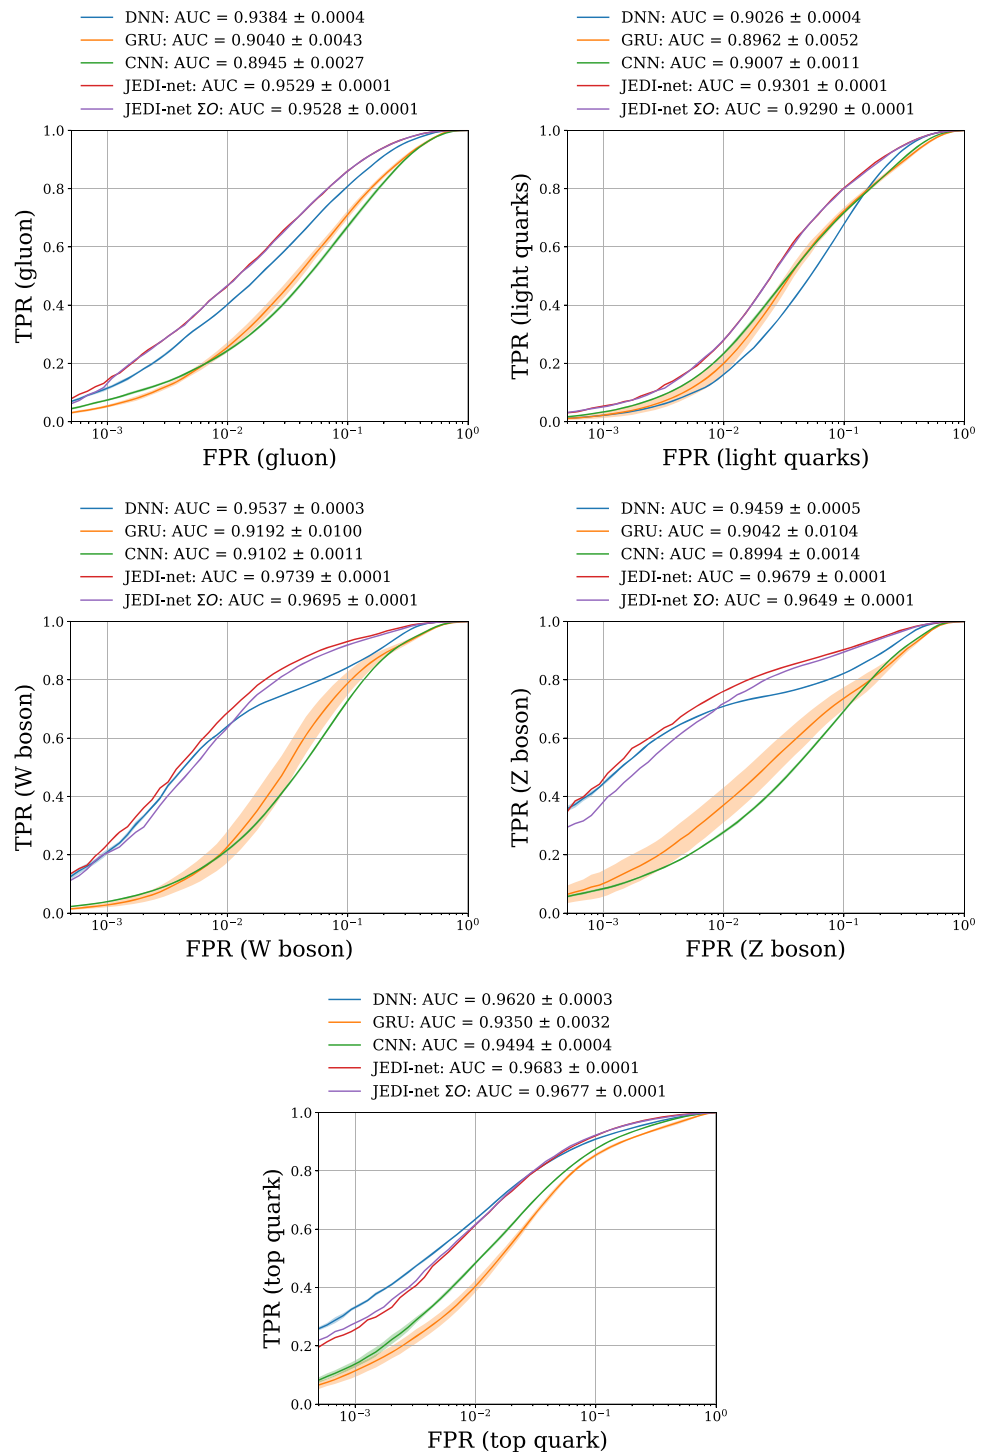
Nevertheless, some interesting correlation pattern between the physics motivated quantities and the \bar{O}_i sums is observed. The most relevant examples are given in Fig. 9, where the 2D histograms and the corresponding linear correlation coefficient (ρ) are shown. The correlation between \bar{O}_1 and the particle multiplicity in the jet is not completely unexpected. As long as the O quantities aggregated across the graph have the same order of magnitude, the corresponding sum \bar{O} would be proportional to jet-constituent multiplicity.

The strong correlation between the \bar{O}_4 and $\tau_1^{(\beta=2)}$ (with ρ values between 0.69 and 0.97, depending on the jet class) is much less expected. The τ_1^β quantities assume small values when the jet constituents can be arranged into a single *sub-jet* inside the jet. Aggregating information from the constituent momenta across the jet, the JEDI-net model based on the \bar{O} quantities learns to build a quantity very close to $\tau_1^{(\beta=2)}$. The last two rows of Fig. 9 show two intermediate cases: the correlation between \bar{O}_2 and $\tau_3^{(\beta=1)}$ and between \bar{O}_9 and $\tau_3^{(\beta=2)}$. The two \bar{O} sums considered are correlated to the corresponding substructure quantities, but with smaller (within 0.48 and 0.77) correlation coefficients.

7 Resource comparison

Table 4 shows a comparison of the computational resources needed by the different models discussed in this paper. The best-performing JEDI-net model has more than twice the number of trainable parameters than the DNN and GRU model, but approximately a factor of 6 less parameters than the CNN model. The JEDI-net model based on the summed \bar{O} features achieves comparable performance with about a

Fig. 8 ROC curves for JEDI-net and the three alternative models, computed for gluons (top-left), light quarks (top-right), W (center-left) and Z (center-right) bosons, and top quarks (bottom). The solid lines represent the average ROC curves derived from 10 k -fold trainings of each model. The shaded bands around the average lines are represent one standard deviation, computed with the same 10 k -fold trainings



factor of 4 less parameters, less than the DNN and GRU models. While being far from expensive in terms of number of parameters, the JEDI-net models are expensive in terms of the number of floating point operations (FLOP). The simple model based on \overline{O} sums, using as input a sequence of 150 particles, uses 458 MFLOP. The increase is mainly due to the scaling with the number of vertices in the graph. Many

of these operations are the $0 \times$ and $1 \times$ products involving the elements of the R_R and R_S matrices. The cost of these operations could be reduced with an IN implementation optimized for inference, e.g., through an efficient sparse-matrix representation.

In addition, we quote in Table 4 the average inference time on a GPU. The inference time is measured applying

Table 3 True positive rates (TPR) for the optimized JEDI-net taggers and the three alternative models (DNN, CNN, and GRU), corresponding to a false positive rate (FPR) of 10% (top) and 1% (bottom). The largest TPR value for each case is highlighted in bold

Jet category	DNN	GRU	CNN	JEDI-net	JEDI-net with $\sum O$
TPR for FPR = 10%					
Gluon	0.830 ± 0.002	0.740 ± 0.014	0.700 ± 0.008	0.878 ± 0.001	0.879 ± 0.001
Lght quarks	0.715 ± 0.002	0.746 ± 0.011	0.740 ± 0.003	0.822 ± 0.001	0.818 ± 0.001
W boson	0.855 ± 0.001	0.812 ± 0.035	0.760 ± 0.005	0.938 ± 0.001	0.927 ± 0.001
Z boson	0.833 ± 0.002	0.753 ± 0.036	0.721 ± 0.006	0.910 ± 0.001	0.903 ± 0.001
Top quark	0.917 ± 0.001	0.867 ± 0.006	0.889 ± 0.001	0.930 ± 0.001	0.931 ± 0.001
TPR for FPR = 1%					
Gluon	0.420 ± 0.002	0.273 ± 0.018	0.257 ± 0.005	0.485 ± 0.001	0.482 ± 0.001
Light quarks	0.178 ± 0.002	0.220 ± 0.037	0.254 ± 0.007	0.302 ± 0.001	0.301 ± 0.001
W boson	0.656 ± 0.002	0.249 ± 0.057	0.232 ± 0.006	0.704 ± 0.001	0.658 ± 0.001
Z boson	0.715 ± 0.001	0.386 ± 0.060	0.291 ± 0.005	0.769 ± 0.001	0.729 ± 0.001
Top quark	0.651 ± 0.003	0.426 ± 0.020	0.504 ± 0.005	0.633 ± 0.001	0.632 ± 0.001

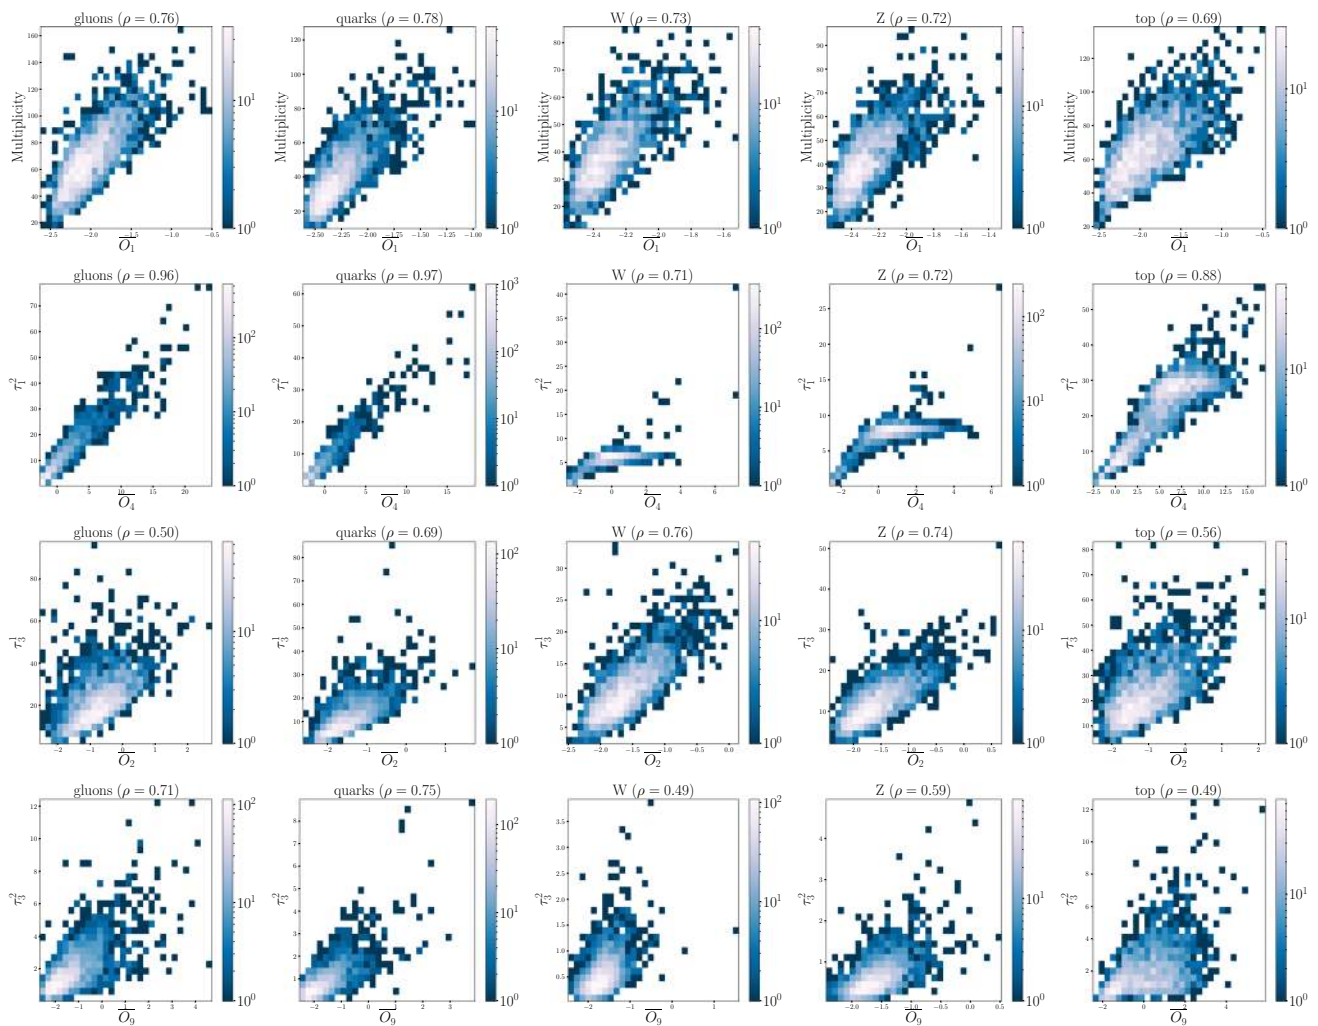
**Fig. 9** Two-dimensional distributions between (top to bottom) \overline{O}_1 and constituents multiplicity, \overline{O}_4 and $\tau_1^{(\beta=2)}$, \overline{O}_2 and $\tau_3^{(\beta=1)}$, \overline{O}_9 and $\tau_3^{(\beta=2)}$, for jets originating from (right to left) gluons, light flavor quarks, Wbosons, Z bosons, and top quarks. For each distribution, the linear correlation coefficient ρ is reported

Table 4 Resource comparison across models. The quoted number of parameters refers only to the trainable parameters for each model. The inference time is measured by applying the model to batches of 1000 events 100 times: the 50% median quantile is quoted as central value and the 10%-90% semi-distance is quoted as the uncertainty. The GPU used is an NVIDIA GTX 1080 with 8 GB memory, mounted on a commercial desktop with an Intel Xeon CPU, operating at a frequency of 2.60GHz. The tests were executed in PYTHON 3.7 with no other concurrent process running on the machine

Model	Number of parameters	Number of FLOP	Inference time/batch (ms)
DNN	14725	27 k	1.0 ± 0.2
CNN	205525	400 k	57.1 ± 0.5
GRU	15575	46 k	23.2 ± 0.6
JEDI-net	33625	116 M	121.2 ± 0.4
JEDI-net with $\sum O$	8767	458 M	402 ± 1

the model to 1000 events, as part of a PYTHON application based on TENSORFLOW [48]. To this end, the JEDI-net models, implemented and trained in PYTORCH, are exported to ONNX [49] and then loaded as TENSORFLOW graph. The quoted time includes loading the data, which occurs for the first inference and is different for different event representations, that is smaller for the JEDI-net models than for the CNN models. The GPU used is an NVIDIA GTX 1080 with 8 GB memory, mounted on a commercial desktop with an Intel Xeon CPU, operating at a frequency of 2.60 GHz. The tests were executed in PYTHON 3.7, with no other concurrent process running on the machine. Given the larger number of operations, the GPU inference time for the two IN models is much larger than for the other models.

The current IN algorithm is costly to deploy in the online selection environment of a typical LHC experiment. A dedicated R&D effort is needed to reduce the resource consumption in a realistic environment in order to benefit from the improved accuracy that INs can achieve. For example, one could trade model accuracy for reduced resource needs by applying neural network pruning [50,51], reducing the numerical precision [52,53], and limiting the maximum number of particles in each jet representation.

8 Conclusions

This paper presents JEDI-net, a jet tagging algorithm based on interaction networks. Applied to a data set of jets from light-flavor quarks, gluons, vector bosons, and top quarks, this algorithm achieves better performance than models based on dense, convolutional, and recurrent neural networks, trained and optimized with the same procedure on

the same data set. As other graph networks, JEDI-net offers several practical advantages that make it particularly suitable for deployment in the data-processing workflows of LHC experiments: it can directly process the list of jet constituent features (e.g. particle four-momenta), it does not assume specific properties of the underlying detector geometry, and it is insensitive to any ordering principle applied to the input jet constituents. For these reasons, the implementation of this and other graph networks is an interesting prospect for future runs of the LHC. On the other hand, the current implementation of this model demands large computational resources and a large inference time, which make the use of these models problematic for real-time selection and calls for a dedicated program to optimize the model deployment on typical L1 and HLT environments.

The quantities engineered by one of the trained IN models exhibit interesting correlation patterns with some of the jet substructure quantities proposed in literature, showing that the model is capable of learning some of the relevant physics in the problem. On the other hand, some of the engineered quantities do not exhibit striking correlation patterns, implying the possibility of a non trivial insight to be gained by studying these quantities.

Acknowledgements We are grateful to Caltech and the Kavli Foundation for their support of undergraduate student research in cross-cutting areas of machine learning and domain sciences. We would also like to thank the Taylor W. Lawrence Research Fellowship and Mellon Mays Fellowship for supporting E. A. M. and making this research effort possible. This work was conducted at “iBanks,” the AI GPU cluster at Caltech. We acknowledge NVIDIA, SuperMicro and the Kavli Foundation for their support of “iBanks”. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement n° 772369) and is partially supported by the U.S. Department of Energy, Office of High Energy Physics Research under Caltech Contract No. DE-SC0011925. J. M. D. is supported by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

Data Availability Statement This manuscript has associated data in a data repository. [Authors’ comment: The data set was created using the configuration and parametric description of an LHC detector described in Ref. [24,32], and is available on the Zenodo platform [33–36].]

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.
Funded by SCOAP³.

Appendix

A Alternative models

The three benchmark models considered in this work are derived through a Bayesian optimization of their hyperparameters, performed using the GPyOPT library [40], based on GPy [41]. For each iteration, the training is performed using early stopping to prevent over-fitting and to allow a fair comparison between different configurations. The data set for training (validation) consists of 630,000 (240,000) jets, with 10,000 jets used for testing purposes. The loss for the Bayesian optimization is estimated on the validation data set. The CNN and GRU networks are trained on four different input data sets, obtained considering the first 30, 50, 100, or 150 highest- p_T jet constituents. The DNN model is trained on quantities computed from the full list of particles.

The DNN model consists on a multilayer perceptron, alternating dense layers to dropout layers. The optimal architecture is determined optimizing the following hyperparameters:

- Number of dense layers (N_{DL}) between 1 and 3.
- Number of neurons per dense layer (n_n): 10, 20, ..., 100.
- Activation functions for the dense layers: ReLU, ELU, or SELU.
- Dropout rate: Between 0.1 and 0.4.
- Batch size: 50, 100, 200, or 500.
- Optimization algorithm: Adam, Nadam [54], or AdaDelta.

The optimization process gives as output an optimal architecture with three hidden layers of 80 neurons each, activated by ELU functions. The best dropout rate is found to be 0.11, when a batch size of 50 and the Adam optimizer are used. This optimized network gives a loss of 0.66 and an accuracy of 0.76.

The CNN model consists of two-dimensional convolutional layers with batch normalization, followed by a set of dense layers. A 2×2 max pooling layer is applied after the first convolutional layer. The optimal architecture is derived optimizing the following hyperparameters:

- Number of convolutional layers N_{CL} between 1 and 3.
- Number of convolutional filters n_f in each layer (10, 15, 20, 25, or 30).
- Convolutional filter size: 3×3 , 5×5 , 7×7 , or 9×9 .
- Max pooling filter size: 2×2 , 3×3 , or 5×5 .
- Activation functions for the convolutional layers (ReLU, ELU, or SELU).
- Number of dense layers N_{DL} between 1 and 3.
- Number of neurons n_n per dense layer: 10, 20, ..., 60.
- Activation functions for the dense layers: ReLU, ELU, or SELU.

Table 5 Optimal CNN hyperparameter setting for different input data sets

Hyperparameter	Number of jet constituents			
	30	50	100	150
N_{CL}	3	1	1	3
n_f	20	10	30	30
Filter size	3×3	3×3	3×3	3×3
Max pooling size	2×2	5×5	5×5	2×2
Conv. activation	ReLU	ELU	ELU	ReLU
N_{DL}	2	3	3	3
n_n	60	50	60	60
Dense activation	SELU	ELU	ELU	ELU
Dropout	0.11	0.1	0.4	0.1
Batch size	200	500	100	50
Optimizer	Adam	Adam	Adam	Adam
Optimized loss	0.88	0.73	0.74	0.74
Optimized accuracy	0.67	0.74	0.74	0.74

The best configuration, used as a benchmark for comparison, is highlighted in bold

- Dropout rate: Between 0.1 and 0.4.
- Batch size: 50, 100, 200, or 500.
- Optimization algorithm: Adam, Nadam, or AdaDelta.

The stride of the convolutional filters is fixed to 1 and “same” padding is used. Table 5 shows the optimal sets of hyperparameter values, obtained for the four different data set representations. While the optimal networks are equivalent in performance, we select the network obtained for ≤ 50 constituents, because it has the smallest number of parameters.

The recurrent model consists of a GRU layer feeding a set of dense layers. The following hyperparameters are considered:

- Number of GRU units: 50, 100, 200, 300, 400, or 500.
- Activation functions for the GRU layers: ReLU, ELU, or SELU.
- Number of dense layers: Between 1 and 4.
- Number of neurons per dense layer: 10, 20, ..., 100.
- Activation functions for the dense layers: ReLU, ELU, or SELU.
- Dropout rate: between 0.1 and 0.4.
- Batch size: 50, 100, 200, or 500.
- Optimization algorithm: Adam, Nadam, or AdaDelta.

The best hyperparameter values are listed in Table 6. As for the CNN model, the best performance is obtained when the list of input particles is truncated at 50 elements.

Table 6 Optimal GRU hyperparameter settings for different input data sets

Hyperparameter	Number of jet constituents			
	30	50	100	150
n_u	100	50	200	50
N_{DL}	3	1	3	4
n_n	70	40	40	100
Dense activation	SELU	SELU	ReLU	ELU
Dropout	0.40	0.10	0.22	0.10
Batch size	500	500	500	500
Optimizer	Adam	Adam	Adam	AdaDelta
Optimized loss	0.78	0.71	0.78	0.85
Optimized accuracy	0.72	0.75	0.73	0.68

The best configuration, used as a benchmark for comparison, is highlighted in bold

B Performance on public top tagging data set

In this appendix, we retrain and evaluate the performance of JEDI-net on a public top tagging data set [19, 23] used to benchmark many neural networks architectures for the task of differentiating top quark jets from light quark jets. To select the hyperparameters of the model (with and without the sum over particles), we performed a Bayesian optimization. We scan N_n^1 from 16 to 256, D_E from 4 to 64, D_O from 4 to 64, ReLU, ELU, or SELU activation functions for f_R , f_O , and ϕ_C , and either the Adam or Adadelta optimizers with an initial learning rate of 10^{-3} . We report three metrics for the performance of the network on the top tagging data set: model accuracy, area under the ROC curve

Table 7 The optimized hyperparameters, number of trainable parameters, and performance metrics of the JEDI-net models on the top tagging data set. Performance metrics are evaluated on the test sample. We quote the area under the ROC curve (AUC), the accuracy, and the background rejection at a signal efficiency of 30%

Model	JEDI-net	JEDI-net with $\sum O$
Number of constituents	150	150
N_n^1	64	256
D_E	64	64
D_O	16	32
f_R activation	ReLU	SELU
f_O activation	SELU	ReLU
ϕ_C activation	ReLU	SELU
Optimizer	Adam	Adam
Number of parameters	169906	148962
Accuracy	0.9263	0.9300
AUC	0.9786	0.9807
$1/\epsilon_B(\epsilon_S = 30\%)$	590.4	774.6

(AUC), and background rejection power at a fixed signal efficiency of 30%, $1/\epsilon_B(\epsilon_S = 30\%)$. In Table 7, the accuracy, AUC, and $1/\epsilon_B(\epsilon_S = 30\%)$ values are listed for each model considered. The performance of JEDI-net compared to other models developed for this data set is approaching state-of-the-art [23].

References

- Y.L. Dokshitzer, G.D. Leder, S. Moretti, B.R. Webber, JHEP **08**, 001 (1997). <https://doi.org/10.1088/1126-6708/1997/08/001>
- S. Catani, Y.L. Dokshitzer, M.H. Seymour, B.R. Webber, Nucl. Phys. B **406**, 187 (1993). [https://doi.org/10.1016/0550-3213\(93\)90166-M](https://doi.org/10.1016/0550-3213(93)90166-M)
- M. Cacciari, G.P. Salam, G. Soyez, JHEP **04**, 063 (2008). <https://doi.org/10.1088/1126-6708/2008/04/063>
- A.J. Larkoski, I. Moult, B. Nachman, (2017). [arXiv:1709.04464](https://arxiv.org/abs/1709.04464)
- P.W. Battaglia, R. Pascanu, M. Lai, D. Rezende, K. Kavukcuoglu, (2016). [arXiv:1612.00222](https://arxiv.org/abs/1612.00222)
- J. Orbach, Arch. Gen. Psychiatry **7**(3), 218 (1962). <https://doi.org/10.1001/archpsyc.1962.01720030064010>
- K. Fukushima, Biol. Cybern. **36**, 193 (1980)
- Y.L. Cun, B. Boser, J.S. Denker, R.E. Howard, W. Hubbard, L.D. Jackel, D. Henderson, Chap. *Handwritten Digit Recognition with a Back-propagation Network* (Morgan Kaufmann Publishers Inc., San Francisco, 1990), pp. 396–404
- A. Waibel, T. Hanazawa, G. Hinton, K. Shikano, K.J. Lang, Chap. *Phoneme Recognition Using Time-delay Neural Networks* (Morgan Kaufmann Publishers Inc., San Francisco, 1990), pp. 393–404
- K. Cho, B. van Merriënboer, D. Bahdanau, Y. Bengio, CoRR (2014). [arXiv:1409.1259](https://arxiv.org/abs/1409.1259)
- G. Louppe, K. Cho, C. Becot, K. Cranmer, JHEP **01**, 057 (2019). [https://doi.org/10.1007/JHEP01\(2019\)057](https://doi.org/10.1007/JHEP01(2019)057)
- S. Egan, W. Fedorko, A. Lister, J. Pearkes, C. Gay (2017). [arXiv:1711.09059](https://arxiv.org/abs/1711.09059)
- T. Cheng, Comput. Softw. Big Sci. **2**(1), 3 (2018). <https://doi.org/10.1007/s41781-018-0007-y>
- P.T. Komiske, E.M. Metodiev, J. Thaler, JHEP **01**, 121 (2019). [https://doi.org/10.1007/JHEP01\(2019\)121](https://doi.org/10.1007/JHEP01(2019)121)
- L. de Oliveira, M. Kagan, L. Mackey, B. Nachman, A. Schwartzman, JHEP **07**, 069 (2016). [https://doi.org/10.1007/JHEP07\(2016\)069](https://doi.org/10.1007/JHEP07(2016)069)
- D. Guest, J. Collado, P. Baldi, S.C. Hsu, G. Urban, D. Whiteson, Phys. Rev. D **94**(11), 112002 (2016). <https://doi.org/10.1103/PhysRevD.94.112002>
- S. Macaluso, D. Shih, JHEP **10**, 121 (2018). [https://doi.org/10.1007/JHEP10\(2018\)121](https://doi.org/10.1007/JHEP10(2018)121)
- K. Datta, A.J. Larkoski, JHEP **03**, 086 (2018). [https://doi.org/10.1007/JHEP03\(2018\)086](https://doi.org/10.1007/JHEP03(2018)086)
- A. Butter, G. Kasieczka, T. Plehn, M. Russell, SciPost Phys. **5**(3), 028 (2018). <https://doi.org/10.21468/SciPostPhys.5.3.028>
- G. Kasieczka, T. Plehn, M. Russell, T. Schell, JHEP **05**, 006 (2017). [https://doi.org/10.1007/JHEP05\(2017\)006](https://doi.org/10.1007/JHEP05(2017)006)
- P.T. Komiske, E.M. Metodiev, M.D. Schwartz, JHEP **01**, 110 (2017). [https://doi.org/10.1007/JHEP01\(2017\)110](https://doi.org/10.1007/JHEP01(2017)110)
- A. Schwartzman, M. Kagan, L. Mackey, B. Nachman, L. De Oliveira, J. Phys. Conf. Ser. **762**(1), 012035 (2016). <https://doi.org/10.1088/1742-6596/762/1/012035>
- A. Butter et al., SciPost Phys. **7**, 014 (2019). <https://doi.org/10.21468/SciPostPhys.7.1.014>
- J. Duarte et al., JINST **13**(07), P07027 (2018). <https://doi.org/10.1088/1748-0221/13/07/P07027>

25. I. Henrion, et al., In *Proceedings of the Deep Learning for Physical Sciences Workshop at NIPS* (2017). https://dl4physicsscience.github.io/files/nips_dlps_2017_29.pdf
26. H. Qu, L. Gouskos, (2019). [arXiv:1902.08570](https://arxiv.org/abs/1902.08570)
27. M. Abdughani, J. Ren, L. Wu, J.M. Yang, J. High Energy Phys. **2019**, (2019). [https://doi.org/10.1007/JHEP08\(2019\)055](https://doi.org/10.1007/JHEP08(2019)055)
28. N. Choma, F. Monti, L. Gerhardt, T. Palczewski, Z. Ronaghi, P. Prabhat, W. Bhimji, M. Bronstein, S. Klein, J. Bruna, 386–391 (2018). <https://doi.org/10.1109/ICMLA.2018.00064>
29. S. Farrell, et al., In *4th International Workshop Connecting The Dots 2018 (CTD2018) (Seattle, Washington, 2018 (20-22 Mar 2018))*
30. J. Arjona Martínez, O. Cerri, M. Pierini, M. Spiropulu, J.R. Vlimant, Eur. Phys. J. Plus **134**(7), 333 (2019). <https://doi.org/10.1140/epjp/i2019-12710-3>
31. S.R. Qasim, J. Kieseler, Y. Iiyama, M. Pierini, Eur. Phys. J. C **79**(7), 608 (2019). <https://doi.org/10.1140/epjc/s10052-019-7113-9>
32. E. Coleman, M. Freytsis, A. Hinzmann, M. Narain, J. Thaler, N. Tran, C. Vernieri, JINST **13**(01), T01003 (2018). <https://doi.org/10.1088/1748-0221/13/01/T01003>
33. J.M. Duarte, et al., *HLS4ML LHC Jet dataset (30 particles)* (Zenodo, 2020). <https://doi.org/10.5281/zenodo.3601436>
34. J.M. Duarte, et al., *HLS4ML LHC Jet dataset (50 particles)* (Zenodo, 2020). <https://doi.org/10.5281/zenodo.3601443>
35. J.M. Duarte, et al., *HLS4ML LHC Jet dataset (100 particles)* (Zenodo, 2020). <https://doi.org/10.5281/zenodo.3602254>
36. J.M. Duarte, et al., *HLS4ML LHC Jet dataset (150 particles)* (Zenodo, 2020). <https://doi.org/10.5281/zenodo.3602260>
37. M. Cacciari, G.P. Salam, G. Soyez, Eur. Phys. J. C **72**, 1896 (2012). <https://doi.org/10.1140/epjc/s10052-012-1896-2>
38. J. Pearkes, W. Fedorko, A. Lister, C. Gay, (2017). [arXiv:1704.02124](https://arxiv.org/abs/1704.02124)
39. A. Paszke, et al., In *NIPS-W* (2017). <https://openreview.net/pdf?id=BJJsrnfmCZ>
40. The GPyOpt authors. GPyOpt: A bayesian optimization framework in python (2016). <http://github.com/SheffieldML/GPyOpt>
41. GPy. GPy: A gaussian process framework in python (2012). <http://github.com/SheffieldML/GPy>
42. V. Nair, G.E. Hinton, In *Proceedings of ICML*, vol. 27 (2010), pp. 807–814
43. D. Clevert, T. Unterthiner, S. Hochreiter, CoRR (2015). [arXiv:1511.07289](https://arxiv.org/abs/1511.07289)
44. G. Klambauer, T. Unterthiner, A. Mayr, S. Hochreiter, CoRR (2017). [arXiv:1706.02515](https://arxiv.org/abs/1706.02515)
45. D.P. Kingma, J. Ba, CoRR (2014). [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)
46. M.D. Zeiler, CoRR (2012). [arXiv:1212.5701](https://arxiv.org/abs/1212.5701)
47. J. Thaler, K. Van Tilburg, JHEP **03**, 015 (2011). [https://doi.org/10.1007/JHEP03\(2011\)015](https://doi.org/10.1007/JHEP03(2011)015)
48. M. Abadi et al., TensorFlow: Large-scale machine learning on heterogeneous systems (2015). Software available from <https://tensorflow.org>
49. J. Bai, F. Lu, K. Zhang et al., Onnx: Open neural network exchange (2019). <https://github.com/onnx/onnx>
50. Y. LeCun, J.S. Denker, S.A. Solla, In *Advances in Neural Information Processing Systems 2*, ed. by D.S. Touretzky (Morgan-Kaufmann, 1990), pp. 598–605. <https://openreview.net/pdf?id=OM0jvwB8jlp57ZJjtNEZ>
51. S. Han, H. Mao, W.J. Dally, CoRR (2015). [arXiv:1510.00149](https://arxiv.org/abs/1510.00149)
52. Y. Cheng, D. Wang, P. Zhou, T. Zhang, CoRR (2017). [arXiv:1710.09282](https://arxiv.org/abs/1710.09282)
53. S. Gupta, A. Agrawal, K. Gopalakrishnan, P. Narayanan, CoRR (2015). [arXiv:1502.02551](https://arxiv.org/abs/1502.02551)
54. T. Dozat, In *ICLR Workshop* (2016)