

# Equivariant, safe and sensitive — graph networks for new physics

Akanksha Bhardwaj <sup>a</sup>, Christoph Englert <sup>b</sup>, Wrishik Naskar <sup>b</sup>,  
Vishal S. Ngairangbam <sup>c</sup> and Michael Spannowsky <sup>c</sup>

<sup>a</sup>*Department of Physics, Oklahoma State University,  
Stillwater, OK, 74078, U.S.A.*

<sup>b</sup>*School of Physics and Astronomy, University of Glasgow,  
Glasgow G12 8QQ, United Kingdom*

<sup>c</sup>*Institute for Particle Physics Phenomenology, Department of Physics,  
Durham University, Durham DH1 3LE, United Kingdom*

*E-mail:* [akanksha.bhardwaj@okstate.edu](mailto:akanksha.bhardwaj@okstate.edu), [christoph.englert@glasgow.ac.uk](mailto:christoph.englert@glasgow.ac.uk),  
[w.naskar.1@research.gla.ac.uk](mailto:w.naskar.1@research.gla.ac.uk), [vishal.s.ngairangbam@durham.ac.uk](mailto:vishal.s.ngairangbam@durham.ac.uk),  
[michael.spannowsky@durham.ac.uk](mailto:michael.spannowsky@durham.ac.uk)

**ABSTRACT:** This study introduces a novel Graph Neural Network (GNN) architecture that leverages infrared and collinear (IRC) safety and equivariance to enhance the analysis of collider data for Beyond the Standard Model (BSM) discoveries. By integrating equivariance in the rapidity-azimuth plane with IRC-safe principles, our model significantly reduces computational overhead while ensuring theoretical consistency in identifying BSM scenarios amidst Quantum Chromodynamics backgrounds. The proposed GNN architecture demonstrates superior performance in tagging semi-visible jets, highlighting its potential as a robust tool for advancing BSM search strategies at high-energy colliders.

**KEYWORDS:** Jets and Jet Substructure, Dark Matter at Colliders

ARXIV EPRINT: [2402.12449](https://arxiv.org/abs/2402.12449)

---

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Inductive biases and QCD</b>	<b>2</b>
2.1	Infrared and collinear safety	2
2.2	Equivariance in rapidity-azimuth plane	3
<b>3</b>	<b>A physics application: tagging semi-visible jets</b>	<b>5</b>
3.1	Simulation details	6
3.2	Network architecture and training	8
3.3	Results	8
<b>4</b>	<b>Summary and conclusions</b>	<b>10</b>

---

## 1 Introduction

The application of machine learning algorithms to jet classification [1–29] provides an ideal environment to gauge the interplay between performance and interpretability. On the one hand, we need highly performant algorithms utilising the wealth of experimental data recorded at the Large Hadron Collider (LHC), while on the other, we want to ascertain the reason behind the algorithm’s efficiency. Architectures which are Lorentz equivariant [30–34], or are infra-red and collinear (IRC) safe [35–42] have been shown to enhance the physical biases of the algorithms with generally minimal loss, or at times, comparatively increased performance [31, 34] with respect to physics-opaque constructions. In this work, we combine IRC safety with Euclidean equivariance in the rapidity-azimuth plane. With well-known examples of jet-shape [43–50] variables, constructed as functions of distance in the rapidity-azimuth plane, this pushes the IRC-safe features closer to human-engineered QCD features without losing out on performance, even with two orders of magnitude lower number of model parameters.

Symmetry in a different sense, namely the degeneracy over nearly degenerate final states, is fundamental to the theoretical robustness of information that can be gained from phenomenological analyses. The Kinoshita-Lee-Nauenberg theorem [51, 52] guarantees cancellations of soft and (final-state) collinear divergences when algorithms are infrared (IR) and collinear safe; any departure from IRC-safe methods ultimately implies uncorrectable interpretation shortfalls. This technically elevates IRC safety to a comparable level as the space-time symmetries to perturbatively model and interpret scattering processes.

The recent surge in the application of machine learning for enhancing the sensitivity to new physics in available and projected particle physics data sets has led to a departure from observables that transparently reflect symmetry properties and IRC safety. Nonetheless, any such serious attempt should critically feature IRC safety whilst exploiting symmetries of the data sets in the corresponding architecture. This work aims to demonstrate that

this is not only feasible but directly applicable to the determination of BSM parameter regions in case a discovery can be claimed via these techniques. We construct an IRC-safe  $\mathbf{E}(2)$ -equivariant network and demonstrate the enhanced optimisation capability of such an algorithm compared to approaches that do not exploit symmetry. As most popular substructure variables are generally built out of  $\Delta R = \sqrt{(\Delta y)^2 + (\Delta\phi)^2}$ , our choice of the  $\mathbf{E}(2)$ -group preserves such a structure in the feature extraction process, thereby closely matching the QCD-intuitive picture of high-level variables. We then apply this approach to a classification of jet substructure analyses (we consider Hidden Valley Models [53–60] as a representative example), where IRC safety plays a critical role in obtaining a theoretically robust output score when the substructure becomes sparse or soft.

This work is organised as follows: in section 2, we discuss the inductive biases relevant to phenomenology at the LHC, namely IRC safety and equivariance in the rapidity-azimuthal plane. Then, in section 3, we apply this network to a relevant physics case where it discriminates between QCD jets and dark shower jets (semi-visible jets). We discuss the network architecture and its performance. Finally, we provide a summary and conclusion in section 4.

## 2 Inductive biases and QCD

The success of deep learning algorithms relies on reducing the search space of possible functions to a predetermined subset by imposing additional structures on the architecture. These inductive biases favour the extraction of relevant features which are not too specific (like high-level variables) but follow the intuition of the application domain. Here, we discuss the two inductive biases we impose on the Message Passing Neural Network utilised in our study: IRC safety and equivariance in the rapidity-azimuth plane.

### 2.1 Infrared and collinear safety

From a QCD perspective, infra-red and collinear safe feature extraction promises to provide better interpretability while not losing out on the high classification capabilities. Any message passing operation involves taking in the node features  $\mathbf{H}_i^{(l)}$  of each node  $i$  as the input and updating it to  $\mathbf{H}_i^{(l+1)}$ . As IRC safety is defined at the level of the whole jet and not for each constituents, these node features inevitably undergo a graph-readout operation for jet-level classification. A way to define an IRC-safe graph-representation is the “*conservation of collinearity*” of the node features in each representation indexed by  $l$  of the message passing operation which preserves the structure of the angular components in the limit of collinear emissions. In the following, each particle  $i$  has four vectors  $p_i^\mu = (z_i, \hat{\mathbf{p}}_i)$ , with  $z_i = p_T^i / \sum_k p_T^k$  and  $\hat{\mathbf{p}}_i$  denoting the angular vector. Mathematically, a node representation  $\mathbf{H}_i^{(l)}$  conserves collinearity, if for any particle  $q$ ,  $r$ , and  $s$ , we have

$$z_q = z_r + z_s \wedge \hat{\mathbf{p}}_q = \hat{\mathbf{p}}_r = \hat{\mathbf{p}}_s \implies \mathbf{H}_q^{(l)} = \mathbf{H}_r^{(l)} = \mathbf{H}_s^{(l)} \quad . \quad (2.1)$$

Once this condition is satisfied, it is straightforward to define an IRC safe graph representation as

$$\mathbf{G}_l = \sum_i z_i \mathbf{H}_i^{(l)} \quad . \quad (2.2)$$

For implementing IRC safety in Energy-weighted Message Passing Network (EMPN) [37], one defines an IRC-safe prescription for defining the graph structure. Once this is satisfied for any node  $i$  in the graph, the message-passing operation

$$\mathbf{H}_i^{(l+1)} = \sum_{j \in \mathcal{N}[i]} \omega_j^{(\mathcal{N}[i])} \hat{\Phi}^{(l+1)}(\mathbf{H}_i^{(l)}, \mathbf{H}_j^{(l)}) \quad (2.3)$$

where the scope-dependent factors are an analogue of  $z_i$ , defined as

$$\omega_j^{(\mathcal{N}[i])} = \frac{p_T^j}{\sum_{k \in \mathcal{N}[i]} p_T^k} \quad ,$$

recursively conserves collinearity of the node features if the initial input  $\mathbf{H}_i^{(0)}$  follows the same. For the input, one can use the relative coordinates in the rapidity-azimuth plane from the jet axis,  $\mathbf{H}_i^{(0)} = (\Delta y_{iJ}, \Delta \phi_{iJ})$  satisfies the condition trivially. The function  $\hat{\Phi}^{(l)}$ , in general, can be any well-behaved neural network which we take to be an edge convolution network [61]. The neighbourhood  $\mathcal{N}[i]$  is constructed by imposing the step function  $\Theta(\Delta R_{ij} < R_0)$  for each particle  $i$ ,  $R_0$  being a parameter that determines the density of the constructed graph. The graph features (by taking a graph readout as defined in eq. 2.2) for  $l > 0$  can be concatenated up to the total number of such operations  $L$ , which is a hyperparameter. For classification, a downstream network takes the graph representation as an input and gives a classification score.

## 2.2 Equivariance in rapidity-azimuth plane

Equivariance is at the core of the powerful learning ability shown by modern deep-learning algorithms such as Convolutional Neural Networks (CNNs) and Graph Neural Networks (GNNs). They reduce the possible set of functions that a neural network can approximate by assuming symmetries of the input data. For instance, CNNs assume translational symmetry in  $\mathbb{R}^2$  while GNNs assume permutation symmetry of the nodes. Equivariance enforces group algebraic structures in the hidden representations in a neural network. This reduces the set of possible functions that a neural network can approximate to those which follow the group equivariant property.

Mathematically, for a group  $\mathcal{G}$  with an element denoted as  $g$ , let  $\mathbf{X}$  and  $\mathbf{Y}$  be two sets with  $\mathbf{x} \in \mathbf{X}$  and  $\mathbf{y} \in \mathbf{Y}$  that admit group actions<sup>1</sup>  $\mathbf{T}_{\mathbf{X}}(g, \mathbf{x})$  and  $\mathbf{T}_{\mathbf{Y}}(g, \mathbf{y})$ , respectively. A function  $f : \mathbf{X} \rightarrow \mathbf{Y}$  is  $\mathcal{G}$ -equivariant if

$$f(\mathbf{T}_{\mathbf{X}}(g, \mathbf{x})) = \mathbf{T}_{\mathbf{Y}}(g, f(\mathbf{x})) \quad . \quad (2.4)$$

An equivariant neural network [62] approximates such a function  $f$  for a given group, and  $f$  is generally non-linear. From an approximation perspective, the neural network being able to learn only those functions that follow eq. (2.4) rather than any general function necessarily reduces the expressive power. However, the reduction generally results in higher efficiency of the optimisation process in terms of network complexity, on top of increasing the interpretability as the particular group  $\mathcal{G}$  is determined from domain knowledge. If the

---

<sup>1</sup>An action of a group  $\mathcal{G}$  is defined as  $\mathbf{T}_{\mathbf{X}} : \mathcal{G} \times \mathbf{X} \rightarrow \mathbf{X}$  such that for any  $\mathbf{x} \in \mathbf{X}$ , we have  $\mathbf{T}_{\mathbf{X}}(e, \mathbf{x}) = \mathbf{x}$  for the neutral element  $e \in \mathcal{G}$  and  $\mathbf{T}_{\mathbf{X}}(g_1, \mathbf{T}_{\mathbf{X}}(g_2, \mathbf{x})) = \mathbf{T}_{\mathbf{X}}(g_1 g_2, \mathbf{x})$ , for all  $g_1, g_2 \in \mathcal{G}$ .

action  $\mathbf{T}_Y$  is trivial, i.e.,  $\mathbf{T}_Y(g, \mathbf{y}) = \mathbf{y} \forall g \in \mathcal{G}$ , then the function  $f$  is said to be  $\mathcal{G}$ -invariant. GNNs are permutation equivariant in the node update stage, while for graph-level purposes, the global readout renders the graph representation permutation invariant.

For our purpose, the group which preserves the separation  $\Delta R$  in the  $(y, \phi)$  plane is the Euclidean group  $\mathbf{E}(2)$ . To incorporate  $\mathbf{E}(2)$ -equivariance in an EMPN, we modify the general  $\mathbf{E}(n)$ -equivariant message-passing operation of reference [63]. To define  $\mathbf{E}(n)$  equivariance in this form, one segregates the node features  $\mathbf{H}_i^{(l)}$ , into a scalar and a vector representation. Denoting the former as  $\mathbf{h}_i^{(l)}$  and the latter as  $\mathbf{x}_i^{(l)}$ , the  $\mathbf{E}(n)$  equivariant operation is given as

$$\mathbf{m}_{ij}^{(l+1)} = \Phi_e^{(l+1)}(\mathbf{h}_i^{(l)}, \mathbf{h}_j^{(l)}, |\mathbf{x}_i^{(l)} - \mathbf{x}_j^{(l)}|) \quad , \quad (2.5a)$$

$$\mathbf{x}_i^{(l+1)} = \mathbf{x}_i^{(l)} + \sum_{j \neq i} (\mathbf{x}_i^{(l)} - \mathbf{x}_j^{(l)}) \Phi_x^{(l+1)}(\mathbf{m}_{ij}^{(l+1)}) \quad , \quad (2.5b)$$

$$\mathbf{m}_i^{(l+1)} = \sum_{j \neq i} \mathbf{m}_{ij}^{(l+1)} \quad , \quad (2.5c)$$

$$\mathbf{h}_i^{(l+1)} = \Phi_h(\mathbf{h}_i^{(l)}, \mathbf{m}_i^{(l+1)}) \quad , \quad (2.5d)$$

where  $\Phi_e^{(l+1)}$ ,  $\Phi_x^{(l+1)}$ , and  $\Phi_h^{(l+1)}$  are multilayer perceptrons (MLPs), with  $\Phi_x^{(l+1)}$  giving a one-dimensional output with a sigmoid activation, thereby being interpreted as a scalar weight in the aggregation. The restriction  $j \neq i$  in the aggregation steps can be relaxed as the equivariance is satisfied in its absence. Moreover, as the neighbourhood construction modifies the message functions ( $\Phi_e^{(l+1)}$  and  $\Phi_x^{(l+1)}$ ) by a product with a step function  $\Theta(\Delta R_{ij} < R_0)$  which is itself  $\mathbf{E}(2)$  invariant, we can incorporate a local aggregation scheme without impeding global equivariance. We note that modifying eq. (2.3) to

$$\mathbf{H}_i^{(l+1)} = \mathbf{F}^{(l+1)} \left( \mathbf{H}_i^{(l)}, \sum_{i \in \mathcal{N}[i]} \omega_j^{(\mathcal{N}[i])} \hat{\Phi}^{(l+1)}(\mathbf{H}_i^{(l)}, \mathbf{H}_j^{(l)}) \right) \quad , \quad (2.6)$$

with a well-behaved (to be learned) function  $\mathbf{F}^{(l+1)}$  still conserves collinearity of the updated node features. This is because the conservation of collinearity in the second argument already requires its conservation in the first argument  $\mathbf{H}_q^{(l)} = \mathbf{H}_r^{(l)} = \mathbf{H}_s^{(l)}$ . In general, by the definition of a function not being one-to-many, any function  $f(\mathbf{H}_i^{(1)}, \mathbf{H}_i^{(2)}, \dots, \mathbf{H}_i^{(n)})$  of  $n$  node representations  $\mathbf{H}_i^{(l)}$  of a single node feature  $i$  conserves collinearity, if each  $\mathbf{H}_i^{(l)}$  conserves collinearity. Equipped with this observation, it is straightforward to modify eq. (2.5) to satisfy IRC safety by incorporating the energy-weighted sum structure in the aggregation as

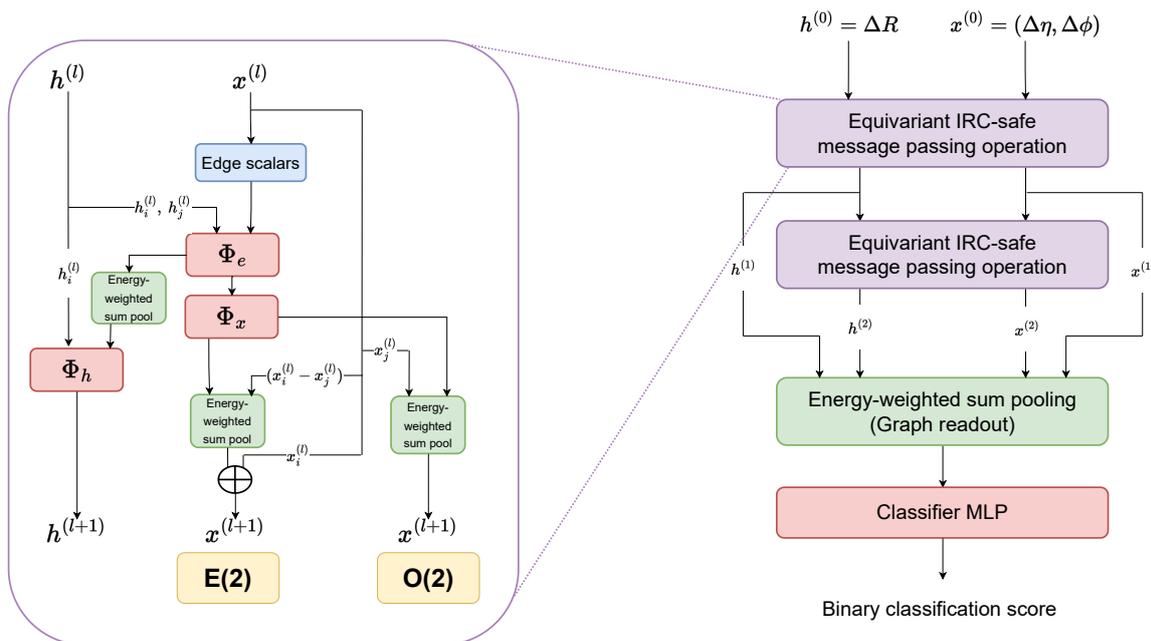
$$\mathbf{m}_{ij}^{(l+1)} = \Phi_e^{(l+1)}(\mathbf{h}_i^{(l)}, \mathbf{h}_j^{(l)}, |\mathbf{x}_i^{(l)} - \mathbf{x}_j^{(l)}|^2) \quad , \quad (2.7a)$$

$$\mathbf{x}_i^{(l+1)} = \mathbf{x}_i^{(l)} + \sum_{j \in \mathcal{N}[i]} \omega_j^{(\mathcal{N}[i])} (\mathbf{x}_i^{(l)} - \mathbf{x}_j^{(l)}) \Phi_x^{(l+1)}(\mathbf{m}_{ij}^{(l+1)}) \quad , \quad (2.7b)$$

$$\mathbf{m}_i^{(l+1)} = \sum_{j \in \mathcal{N}[i]} \omega_j^{(\mathcal{N}[i])} \mathbf{m}_{ij}^{(l+1)} \quad , \quad (2.7c)$$

$$\mathbf{h}_i^{(l+1)} = \Phi_h^{(l+1)}(\mathbf{h}_i^{(l)}, \mathbf{m}_i^{(l+1)}) \quad . \quad (2.7d)$$

Conservation of collinearity of the updated scalar features  $\mathbf{h}_i^{(l+1)}$ , and vector features  $\mathbf{x}_i^{(l+1)}$  follows from comparing their corresponding aggregation equations to eq. (2.6). We have



**Figure 1.** A schematic representation of  $\mathbf{E}(2)$ -EMPN and  $\mathbf{O}(2)$ -EMPN network architecture.

modified the input Euclidean norm to the squared value as we have self-loops, and the square root is not differentiable at zero.

The general  $\mathbf{E}(n)$  group consists of translations and rotations. While there are some instances where the translational symmetry is utilised by defining vector observables like jet pull [64–66], most substructure observables are confined to a rotationally symmetric definition by using only the scalar separation  $\Delta R$ . Therefore, we also study an  $\mathbf{O}(2)$ -equivariant architecture by forgoing translational equivariance. This can be done by modifying eqs. (2.7a) and (2.7b) to

$$\mathbf{m}_{ij}^{(l+1)} = \Phi_e^{(l+1)}(\mathbf{h}_i^{(l)}, \mathbf{h}_j^{(l)}, |\mathbf{x}_i^{(l)} - \mathbf{x}_j^{(l)}|^2, \mathbf{x}_i^{(l)} \cdot \mathbf{x}_j^{(l)}) \quad , \quad (2.8a)$$

$$\mathbf{x}_i^{(l+1)} = \sum_{j \in \mathcal{N}[i]} \omega_j^{(\mathcal{N}[i])} \mathbf{x}_j^{(l)} \Phi_x^{(l+1)}(\mathbf{m}_{ij}^{(l+1)}) \quad , \quad (2.8b)$$

and keeping eqs. (2.7c) and (2.7d) unchanged. We have added the inner product  $\mathbf{x}_i^{(l)} \cdot \mathbf{x}_j^{(l)}$  as it is invariant to rotations, which is not the case for translations.

Note that for the final IRC safe feature extraction, we need to perform an energy-weighted summed global readout as given in eq. (2.2) for both  $\mathbf{x}_i^{(l)}$  and  $\mathbf{h}_i^{(l)}$ . Moreover, the input scalars  $\mathbf{h}_i^{(0)}$  cannot contain  $z_i$  from the requirement of IRC safety, even though it is invariant under  $\mathbf{E}(2)$ . A schematic representation of the equivariant message passing is shown on the left of figure 1.

### 3 A physics application: tagging semi-visible jets

The radiation pattern arising from perturbative QCD emissions is known to be approximately uniform in the rapidity-azimuth plane. This guides the design of most jet-shape observables [43] to use the relative distance between  $\Delta R$ , between pairs of particles [46, 48], or from

well-defined axes [44, 45]. The implicit assumption underlying such definitions is the invariance of the observables under the two-dimensional Euclidean group  $\mathbf{E}(2)$  in the rapidity-azimuth plane. It is, therefore, a natural extension of IRC-safe feature extraction to incorporate  $\mathbf{E}(2)$ -equivariance to better connect to known jet substructure observables in literature.

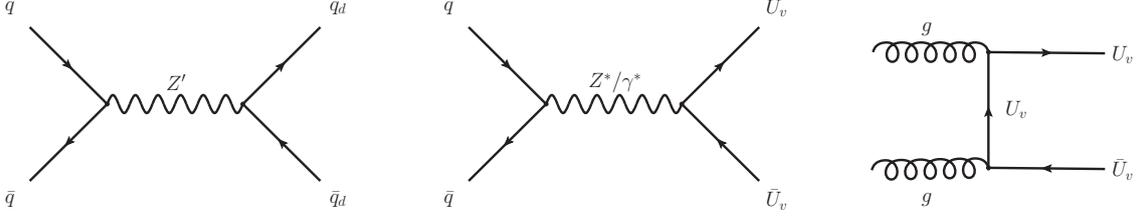
Although equivariance, in principle, would help segregate QCD radiation patterns with different  $n$ -prong structures, a particularly challenging case would be when the parton shower itself is altered with non-trivial physics. For example, a QCD-like dark shower in Hidden Valley models can produce a mixture of BSM and normal SM particles, generally modelled via interleaved parton showers. In such a case, the differences that arise from the QCD radiation are not in the presence of the hard prongs but rather at an intricate level in the relative distribution of the particle's position in the  $(y, \phi)$  plane, with possibly some missing patches left by undetectable stable or quasi-stable particles within the jet. Generally, the energy values would be similar to a one-prong QCD jet. At the same time, the relative position of the particles in the  $(y, \phi)$  plane would be the distinguishing feature in such semi-visible jets. Therefore,  $\mathbf{E}(2)$  equivariance would be more prominent in identifying such semi-visible jets.

Hidden Valley Models offer an intriguing perspective in particle physics, suggesting the existence of secluded, almost decoupled sectors that interact weakly with the SM. Such a scenario can be brought about by extending the SM gauge group with a new group  $G_v$ . The choices of the gauge group and new fermions and their relation to the SM particles offer a rich phenomenology with possibly distinct signatures like emerging jets or semi-visible jets not covered by traditional searches at the LHC, like monojet searches [67, 68]. In general, the SM particles remain neutral under the new group  $G_v$ , while there can be two classes of new particles: those charged under the SM gauge group and the  $G_v$  and those which solely have a  $G_v$  charge. Interactions between the new sector and SM particles can occur through intermediary mechanisms, such as a TeV scale  $Z'$  or through higher-dimensional operators and loops involving heavy particles possessing non-trivial  $G_{\text{SM}}$  and  $G_v$  charges. The case of a confined dark sector decoupled from SM in the lower energies is of particular interest.

### 3.1 Simulation details

Strongly coupled hidden-sector particles produced at the LHC could produce invisible dark mesons along with stable mesons. In such a case, novel LHC signatures such as semi-visible jets could be produced, almost indistinguishable from QCD backgrounds, therefore posing an exciting challenge for jet classification. This study considers resonant and non-resonant production of semi-visible jets. For the latter, we consider both Electroweak (EW) and QCD productions separately. We simulate these dark shower jets events at 14 TeV centre-of-mass energy proton-proton collisions using the Hidden Valley (HV) module of PYTHIA8.310 [69]. They are generated as:

- **Resonant production:** a new leptophobic vector boson  $Z'$  is produced as a BSM mediator, decaying to a pair of dark quarks. Events are generated setting with the process `HiddenValley:ffbar2Zv = on`. We set the mass of  $Z' = 1$  TeV with a width of 10 MeV. The  $Z'$  decays to a pair of dark quarks ( $q_d \bar{q}_d$ ), which has no SM charge but can decay to another dark sector particle and an SM quark. The subsequent showering and hadronisation produce dark mesons and SM particles.



**Figure 2.** Representative Feynman diagrams for the dark shower event generation processes.

- **Non-resonant EW production:** we produce a pair of  $U_v$  quarks using the process `HiddenValley:ffbar2UvUvbar = on`. These  $U_v$ 's are charged under the fundamental representation of the dark  $SU(N)$  and mirror the SM charges of the  $u$ -quark. It, therefore, radiates into both SM and dark sector particles.
- **Non-resonant QCD production:** the same  $U_v$  pair of quarks are produced using the process `HiddenValley:gg2UvUvbar = on`. Unlike the previous channel, which occurs via an s-channel colour singlet propagator, this channel has an SM colour connection between the initial and the final state lines, resembling QCD dijet productions more closely.

The fraction of stable to total dark mesons produced in the process is controlled by the `probVector` parameter in the HV Module. We have chosen `probVector=0.5` in our analysis. This choice is motivated by higher values producing events with low missing transverse energy, which is covered by dijet bump hunts. In comparison, lower values correspond to signatures where the missing transverse energy's  $\phi$  direction is away from the jet axes and, hence, are covered by more traditional dark matter searches.

All final state particles detectable at the LHC are clustered into microjets of radius  $R = 0.1$  using the anti- $k_t$  algorithm [70], mimicking the calorimeter resolution. Inclusive microjets with transverse momentum  $p_T > 1$  GeV are then clustered into jets of radius  $R = 0.8$  with the anti- $k_t$  algorithm. The jet clusterings are performed using the FASTJET (v3.4.1) [71] package. Events are required to have at least two jets within  $|\eta| < 3$ , the leading jet is required to have  $p_T > 150$  GeV, and other sub-leading jets are required to have  $p_T > 120$  GeV.

For the background, QCD dijet events are generated by setting `HardQCD:all = on`. To have better efficiency in the event selection, we set `PhaseSpace:pThatMin= 100`, which puts a lower bound of 100 GeV on the transverse momentum on each of the two final state legs in the  $2 \rightarrow 2$  processes for all classes. We generate 300k events passing the selection criteria for each class, which are used to create three binary classification datasets for each of the three signals vs the background. The leading jet's constituents in these events are used to construct jet graphs. We use  $R_0 = 0.5$  to construct the radius graph and extract  $\mathbf{h}_i^{(0)} = \Delta R_{iJ}$ , and  $\mathbf{x}_i^{(0)} = (\Delta y_{iJ}, \Delta \phi_{iJ})$  along with  $z_i$  and  $\omega_j^{(\mathcal{N}^{[i]})}$ . For the ordinary EMPN architecture, we use  $\mathbf{H}_i^{(0)} = \mathbf{x}_i^{(0)}$ . Since all three networks process the node features to extract relevant edge features as inputs to the learnable functions, no edge features are added. The combined datasets of the signal and background (consisting of 600k samples) are segregated into 60% training, 20% testing, and 20% validation datasets.

### 3.2 Network architecture and training

The network analysis is carried out using the `PYTORCH-GEOMETRIC` [72] package. We consider an IRC-safe `EdgeConv` operation modified with the radius graph and the energy-weighted structure as an example of a non-equivariant architecture. The message passing operation for this architecture follows from eq. (2.3), with  $\hat{\Phi}$  being an MLP which takes the input of the form:

$$\mathbf{H}_i^{(l+1)} = \sum_{j \in \mathcal{N}[i]} \omega_j^{(\mathcal{N}[i])} \hat{\Phi}^{(l+1)}(\mathbf{H}_i^{(l)}, \mathbf{H}_j^{(l)} - \mathbf{H}_i^{(l)}) \quad . \quad (3.1)$$

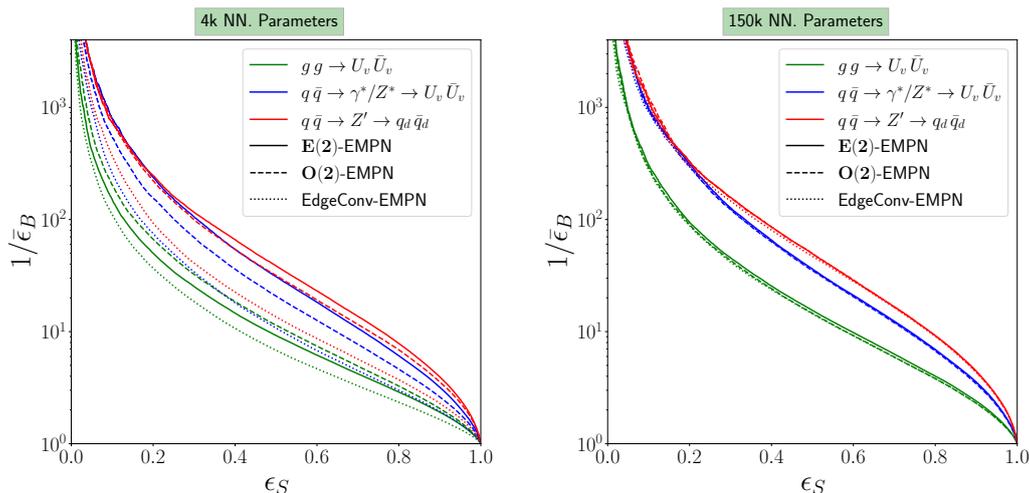
For ease of reading, we name the three architecture as **E(2)**-EMPn, **O(2)**-EMPn, and `EdgeConv`-EMPn. As one of the hallmarks of equivariant architectures is the efficiency of feature extraction with a low number of tunable parameters, we consider a small variant containing about 4k parameters and a large variant containing about 150k parameters. The **E(2)**-EMPn and **O(2)**-EMPn have the same structure of MLPs, and the difference arises from the edge input to  $\Phi_e^{(l)}$  and the aggregation of the vectors. Since the `EdgeConv`-EMPn has only one MLP, while the equivariant ones have three, we reduce the dimensionality of the hidden (scalar) node representations in the latter to keep the number of parameters comparable. We apply the message-passing operation twice for all architectures and extract the IRC-safe graph representation after each operation to feed to a classifier MLP. A schematic diagram for the equivariant ones is shown on the right side of figure 1.

For the small parameter size, the equivariant architectures have a 12-dimensional updated scalar node feature after each message operation, while the `EdgeConv`-EMPn has 16-dimensional updated node features. For  $l \in \{1, 2\}$ , the MLPs  $\Phi_e^{(l)}$ ,  $\Phi_x^{(l)}$ , and  $\Phi_h^{(l)}$  have two hidden layers with the same dimension as the updated scalar feature dimension. Similarly, the edge MLPs have two hidden layers with the same dimensions as the updated node features. An analogous architecture is repeated for the large parameter case, with the equivariant ones updating 80-dimensional scalar node features and the `EdgeConv`-EMPn updating 128-dimensional ones. All hidden layers have `ReLU` activation. Except for  $\Phi_x^{(l)}$ 's, which have a one-dimensional output with `Sigmoid` activation, we do not put any output activation for the other MLPs. The graph representation for the **E(2)**-EMPn and **O(2)**-EMPn cases is obtained from the two scalar node representations and two vector node representations.

In total, we have three binary classification datasets and six networks, thereby giving us eighteen training cases. We train the network from random initialisation five times for each of these cases with a batch size of 300 for 200 epochs. The networks are optimised with the `ADAM` [73] optimiser with an initial learning rate of  $10^{-3}$ . A decay-on-plateau criterion is utilised for the learning rate, triggering a decay by a factor of 0.5 if the validation loss remains stagnant for three consecutive epochs.

### 3.3 Results

We extract the prediction on the test dataset from the epoch with the lowest validation lost for each of the five training of the eighteen instances. To obtain a summary representation of the receiver-operator characteristics (ROC) curve, we extract the background acceptance  $\epsilon_B$ , at fixed signal efficiency  $\epsilon_S$  over the five training instances of training using a custom



**Figure 3.** The network performance is shown for two scenarios: one with a low number of neural network parameters ( $\approx 4k$ ) and another with a high number of parameters ( $\approx 150k$ ), for **E(2)**-EMPn (solid), **O(2)**-EMPn (dashed) and EdgeConv-EMPn (dotted) for the three cases of semi-visible jet tagger.

NN parameters	4k			150k		
	Process	$gg \rightarrow U_v \bar{U}_v$	$q\bar{q} \rightarrow \gamma^*/Z^* \rightarrow U_v \bar{U}_v$	$q\bar{q} \rightarrow Z' \rightarrow q_d \bar{q}_d$	$gg \rightarrow U_v \bar{U}_v$	$q\bar{q} \rightarrow \gamma^*/Z^* \rightarrow U_v \bar{U}_v$
<b>E(2)</b> -EMPn	$0.807 \pm 0.051$	$0.896 \pm 0.003$	$0.913 \pm 0.002$	$0.855 \pm 0.002$	$0.906 \pm 0.001$	$0.923 \pm 0.001$
<b>O(2)</b> -EMPn	$0.827 \pm 0.036$	$0.872 \pm 0.037$	$0.904 \pm 0.028$	$0.848 \pm 0.003$	$0.903 \pm 0.002$	$0.922 \pm 0.001$
EdgeConv-EMPn	$0.767 \pm 0.041$	$0.817 \pm 0.042$	$0.845 \pm 0.006$	$0.848 \pm 0.000$	$0.904 \pm 0.000$	$0.921 \pm 0.001$

**Table 1.** The network performance is illustrated for two scenarios: one with a low number of neural network parameters ( $\approx 4k$ ) and another with a high number of parameters ( $\approx 150k$ ), for the **E(2)**-EMPn, **O(2)**-EMPn, and EdgeConv-EMPn.

implementation.<sup>2</sup> Extracting the mean value of the background acceptance  $\bar{\epsilon}_B$ , over these five instances, we plot the ROC curve between  $1/\bar{\epsilon}_B$  and  $\epsilon_S$  in figure 3. The average of the area under the curve (AUC) and the standard deviation are listed in table 1. Before comparing the relative differences for the considered architectures, we see that the classification efficiency follows the usual physics intuition. The resonant production of a heavy colour singlet  $Z'$  sets a hard scale, which induces significant changes in the jets, while the non-resonant productions do not have such a hard scale. In the colour-singlet s-channel non-resonant production via the  $\gamma$  and  $Z$ , there is no colour flow from the initial quarks to the final state (even though the  $U_v$  quark carries an SM QCD charge). These jets are less QCD-like than the non-resonant production via gluons, where colour flows between the initial and final states and hence closely resembles QCD jet production.

For the large network on the right side of figure 3, the three architectures have virtually overlapping ROCs with a very minute decrease for the non-equivariant EdgeConv-EMPn. The values of the AUCs in table 1 follow the same trend. This suggests that with sufficiently large parametrisation, the EdgeConv-EMPn can catch the distinguishing features as well as

<sup>2</sup>This was verified against the SCIKIT-LEARN [74] implementation with AUCs being equal up to four decimal points.

the  $\mathbf{E}(2)$  and  $\mathbf{O}(2)$  equivariant ones, even though the extracted features are not equivariant. In stark contrast, there is a very noticeable dip in performance for the EdgeConv-EMPn of 4k parameters for all three signal cases compared to the two equivariant architectures, with the  $\mathbf{E}(2)$ -EMPn having the best classification power for the resonant and electroweak non-resonant processes, while the  $\mathbf{O}(2)$ -EMPn bettering the  $\mathbf{E}(2)$ -EMPn for the non-resonant production via gluons. Moreover, the best equivariant AUCs in the low parameter case are not that far off from their corresponding best ones in the large parameter cases, with the highest difference for the hardest case of distinguishing gluon-induced production of semi-visible jets. This clearly shows that the inductive bias of equivariance in the rapidity-azimuth plane helps optimise feature extraction and reduces the need for many parameters.

#### 4 Summary and conclusions

The future roadmap of particle physics crucially depends on maximising the discovery potential of the LHC and its high luminosity phase. Recent developments in optimising searches for new physics using machine learning have demonstrated the enormous opportunity created by less traditional approaches to data analysis. Interfacing such highly adapted strategies with theoretical predictions to make robust statements about parameter measurements, whether related to SM or BSM effects, highlights infrared and collinear safety as a central theme of neural network architecture design, e.g. through energy-weighted message passing. Adding such features to the neural network design can practically increase computational demands, which requires optimisation improvements. In this work, we have considered  $\mathbf{E}(2)$  equivariance as an avenue to remove computational overhead. In particular, we have considered jet-substructure-based analysis to show that an  $\mathbf{E}(2)$  equivariant Graph Neural Network optimises the network’s learning capability by directly exploiting the physical symmetry that the data set possesses. As most jet substructure observables deal with rotationally invariant quantities, we also considered  $\mathbf{O}(2)$ -equivariance which is a subgroup of  $\mathbf{E}(2)$ .

Such advances directly impact physics applications. To highlight this, we have considered a representative (but not exclusive) scenario of semi-visible jets. In theories underpinning this BSM signature, extracting, e.g., dark shower parameters critically depends on an IRC-safe implementation of the tagging algorithm as the model navigates between visibly hard and invisibly soft final states. This asks for the correct modelling of the QCD null hypothesis to draw theoretically consistent conclusions. The architecture presented in this work provides excellent signal vs. background discrimination with these desired properties. This opens up the possibility of employing the architectures proposed in this work as a new tool for BSM discovery in challenging QCD-dominated final states.

#### Acknowledgments

A.B. is supported by the U.S. Department of Energy under grant number DE-SC 0016013. C.E. is supported by the UK Science and Technology Facilities Council (STFC) under grant ST/X000605/1 and the Leverhulme Trust under RPG-2021-031. W.N. is funded by a University of Glasgow College of Science and Engineering Scholarship. V.S.N. and M.S. are

supported by the STFC under grant ST/P001246/1. Parts of the computation utilised the Param Vikram-1000 High Performance Computing Cluster at the Physical Research Laboratory.

**Open Access.** This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

## References

- [1] J. Cogan, M. Kagan, E. Strauss and A. Schwartzman, *Jet-Images: Computer Vision Inspired Techniques for Jet Tagging*, *JHEP* **02** (2015) 118 [[arXiv:1407.5675](https://arxiv.org/abs/1407.5675)] [[INSPIRE](#)].
- [2] L. de Oliveira et al., *Jet-images — deep learning edition*, *JHEP* **07** (2016) 069 [[arXiv:1511.05190](https://arxiv.org/abs/1511.05190)] [[INSPIRE](#)].
- [3] G. Kasieczka, T. Plehn, M. Russell and T. Schell, *Deep-learning Top Taggers or The End of QCD?*, *JHEP* **05** (2017) 006 [[arXiv:1701.08784](https://arxiv.org/abs/1701.08784)] [[INSPIRE](#)].
- [4] G. Louppe, K. Cho, C. Becot and K. Cranmer, *QCD-Aware Recursive Neural Networks for Jet Physics*, *JHEP* **01** (2019) 057 [[arXiv:1702.00748](https://arxiv.org/abs/1702.00748)] [[INSPIRE](#)].
- [5] F.A. Dreyer, G.P. Salam and G. Soyez, *The Lund Jet Plane*, *JHEP* **12** (2018) 064 [[arXiv:1807.04758](https://arxiv.org/abs/1807.04758)] [[INSPIRE](#)].
- [6] T. Heimel, G. Kasieczka, T. Plehn and J.M. Thompson, *QCD or What?*, *SciPost Phys.* **6** (2019) 030 [[arXiv:1808.08979](https://arxiv.org/abs/1808.08979)] [[INSPIRE](#)].
- [7] M. Farina, Y. Nakai and D. Shih, *Searching for New Physics with Deep Autoencoders*, *Phys. Rev. D* **101** (2020) 075021 [[arXiv:1808.08992](https://arxiv.org/abs/1808.08992)] [[INSPIRE](#)].
- [8] H. Qu and L. Gouskos, *ParticleNet: Jet Tagging via Particle Clouds*, *Phys. Rev. D* **101** (2020) 056019 [[arXiv:1902.08570](https://arxiv.org/abs/1902.08570)] [[INSPIRE](#)].
- [9] A. Butter et al., *The Machine Learning landscape of top taggers*, *SciPost Phys.* **7** (2019) 014 [[arXiv:1902.09914](https://arxiv.org/abs/1902.09914)] [[INSPIRE](#)].
- [10] E. Bernreuther et al., *Casting a graph net to catch dark showers*, *SciPost Phys.* **10** (2021) 046 [[arXiv:2006.08639](https://arxiv.org/abs/2006.08639)] [[INSPIRE](#)].
- [11] H. Qu, C. Li and S. Qian, *Particle Transformer for Jet Tagging*, in K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu and S. Sabato eds., *Proceedings of the 39th International Conference on Machine Learning, Proc. Mach. Learn. Res.* **162** (2022) 18281 [[arXiv:2202.03772](https://arxiv.org/abs/2202.03772)] [[INSPIRE](#)].
- [12] G. Kasieczka, S. Marzani, G. Soyez and G. Stagnitto, *Towards Machine Learning Analytics for Jet Substructure*, *JHEP* **09** (2020) 195 [[arXiv:2007.04319](https://arxiv.org/abs/2007.04319)] [[INSPIRE](#)].
- [13] F.A. Dreyer and H. Qu, *Jet tagging in the Lund plane with graph networks*, *JHEP* **03** (2021) 052 [[arXiv:2012.08526](https://arxiv.org/abs/2012.08526)] [[INSPIRE](#)].
- [14] C.K. Khosa and S. Marzani, *Higgs boson tagging with the Lund jet plane*, *Phys. Rev. D* **104** (2021) 055043 [[arXiv:2105.03989](https://arxiv.org/abs/2105.03989)] [[INSPIRE](#)].
- [15] F. Canelli et al., *Autoencoders for semivisible jet detection*, *JHEP* **02** (2022) 074 [[arXiv:2112.02864](https://arxiv.org/abs/2112.02864)] [[INSPIRE](#)].

- [16] F.A. Dreyer, G. Soyez and A. Takacs, *Quarks and gluons in the Lund plane*, *JHEP* **08** (2022) 177 [[arXiv:2112.09140](#)] [[INSPIRE](#)].
- [17] L. Cavallini et al., *Tagging the Higgs boson decay to bottom quarks with colour-sensitive observables and the Lund jet plane*, *Eur. Phys. J. C* **82** (2022) 493 [[arXiv:2112.09650](#)] [[INSPIRE](#)].
- [18] A. Butter, B.M. Dillon, T. Plehn and L. Vogel, *Performance versus resilience in modern quark-gluon tagging*, *SciPost Phys. Core* **6** (2023) 085 [[arXiv:2212.10493](#)] [[INSPIRE](#)].
- [19] B.M. Dillon et al., *A normalized autoencoder for LHC triggers*, *SciPost Phys. Core* **6** (2023) 074 [[arXiv:2206.14225](#)] [[INSPIRE](#)].
- [20] T. Faucett, S.-C. Hsu and D. Whiteson, *Learning to identify semi-visible jets*, *JHEP* **12** (2022) 132 [[arXiv:2208.10062](#)] [[INSPIRE](#)].
- [21] P. Keicher, *Machine Learning in Top Physics in the ATLAS and CMS Collaborations*, in the proceedings of the *15th International Workshop on Top Quark Physics*, (2023) [[arXiv:2301.09534](#)] [[INSPIRE](#)].
- [22] T. Finke, M. Krämer, A. Mück and J. Tönshoff, *Learning the language of QCD jets with transformers*, *JHEP* **06** (2023) 184 [[arXiv:2303.07364](#)] [[INSPIRE](#)].
- [23] C.-T. Lu et al., *Probing dark QCD sector through the Higgs portal with machine learning at the LHC*, *JHEP* **08** (2023) 187 [[arXiv:2304.03237](#)] [[INSPIRE](#)].
- [24] D. Bardhan, Y. Kats and N. Wunch, *Searching for dark jets with displaced vertices using weakly supervised machine learning*, *Phys. Rev. D* **108** (2023) 035036 [[arXiv:2305.04372](#)] [[INSPIRE](#)].
- [25] E.M. Metodiev, J. Thaler and R. Wynne, *Anomaly Detection in Collider Physics via Factorized Observables*, [arXiv:2312.00119](#) [[INSPIRE](#)].
- [26] L. Favaro et al., *Semi-visible jets, energy-based models, and self-supervision*, [arXiv:2312.03067](#) [[INSPIRE](#)].
- [27] A. Romero and D. Whiteson, *Jet Rotational Metrics*, [arXiv:2311.06686](#) [[INSPIRE](#)].
- [28] A. Furuichi, S.H. Lim and M.M. Nojiri, *Jet Classification Using High-Level Features from Anatomy of Top Jets*, [arXiv:2312.11760](#) [[INSPIRE](#)].
- [29] T. Gaertner and J. Reiten, *The simplicial substructure of jets*, [arXiv:2312.06596](#) [[INSPIRE](#)].
- [30] A. Bogatskiy et al., *Lorentz Group Equivariant Neural Network for Particle Physics*, [arXiv:2006.04780](#) [[INSPIRE](#)].
- [31] S. Gong et al., *An efficient Lorentz equivariant graph neural network for jet tagging*, *JHEP* **07** (2022) 030 [[arXiv:2201.08187](#)] [[INSPIRE](#)].
- [32] C. Li et al., *Does Lorentz-symmetric design boost network performance in jet physics?*, *Phys. Rev. D* **109** (2024) 056003 [[arXiv:2208.07814](#)] [[INSPIRE](#)].
- [33] Z. Hao, R. Kansal, J. Duarte and N. Chernyavskaya, *Lorentz group equivariant autoencoders*, *Eur. Phys. J. C* **83** (2023) 485 [[arXiv:2212.07347](#)] [[INSPIRE](#)].
- [34] A. Bogatskiy, T. Hoffman, D.W. Miller and J.T. Offermann, *PELICAN: Permutation Equivariant and Lorentz Invariant or Covariant Aggregator Network for Particle Physics*, [arXiv:2211.00454](#) [[INSPIRE](#)].
- [35] S. Choi, S.J. Lee and M. Perelstein, *Infrared Safety of a Neural-Net Top Tagging Algorithm*, *JHEP* **02** (2019) 132 [[arXiv:1806.01263](#)] [[INSPIRE](#)].

- [36] P.T. Komiske, E.M. Metodiev and J. Thaler, *Energy Flow Networks: Deep Sets for Particle Jets*, *JHEP* **01** (2019) 121 [[arXiv:1810.05165](#)] [[INSPIRE](#)].
- [37] P. Konar, V.S. Ngairangbam and M. Spannowsky, *Energy-weighted message passing: an infra-red and collinear safe graph neural network algorithm*, *JHEP* **02** (2022) 060 [[arXiv:2109.14636](#)] [[INSPIRE](#)].
- [38] O. Atkinson et al., *IRC-Safe Graph Autoencoder for Unsupervised Anomaly Detection*, *Front. Artif. Intell.* **5** (2022) 943135 [[arXiv:2204.12231](#)] [[INSPIRE](#)].
- [39] D. Athanasakos et al., *Is infrared-collinear safe information all you need for jet classification?*, [arXiv:2305.08979](#) [[INSPIRE](#)].
- [40] W. Shen, D. Wang and J.M. Yang, *Hierarchical high-point Energy Flow Network for jet tagging*, *JHEP* **09** (2023) 135 [[arXiv:2308.08300](#)] [[INSPIRE](#)].
- [41] P. Konar, V.S. Ngairangbam and M. Spannowsky, *Hypergraphs in LHC phenomenology — the next frontier of IRC-safe feature extraction*, *JHEP* **01** (2024) 113 [[arXiv:2309.17351](#)] [[INSPIRE](#)].
- [42] S. Bright-Thonney, B. Nachman and J. Thaler, *Safe but Incalculable: energy-weighting is not all you need*, [arXiv:2311.07652](#) [[INSPIRE](#)].
- [43] L.G. Almeida et al., *Substructure of high- $p_T$  Jets at the LHC*, *Phys. Rev. D* **79** (2009) 074017 [[arXiv:0807.0234](#)] [[INSPIRE](#)].
- [44] J. Thaler and K. Van Tilburg, *Identifying Boosted Objects with  $N$ -subjettiness*, *JHEP* **03** (2011) 015 [[arXiv:1011.2268](#)] [[INSPIRE](#)].
- [45] J. Thaler and K. Van Tilburg, *Maximizing Boosted Top Identification by Minimizing  $N$ -subjettiness*, *JHEP* **02** (2012) 093 [[arXiv:1108.2701](#)] [[INSPIRE](#)].
- [46] A.J. Larkoski, G.P. Salam and J. Thaler, *Energy Correlation Functions for Jet Substructure*, *JHEP* **06** (2013) 108 [[arXiv:1305.0007](#)] [[INSPIRE](#)].
- [47] A.J. Larkoski, J. Thaler and W.J. Waalewijn, *Gaining (Mutual) Information about Quark/Gluon Discrimination*, *JHEP* **11** (2014) 129 [[arXiv:1408.3122](#)] [[INSPIRE](#)].
- [48] I. Moult, L. Necib and J. Thaler, *New Angles on Energy Correlation Functions*, *JHEP* **12** (2016) 153 [[arXiv:1609.07483](#)] [[INSPIRE](#)].
- [49] M. Dasgupta, L. Schunk and G. Soyez, *Jet shapes for boosted jet two-prong decays from first-principles*, *JHEP* **04** (2016) 166 [[arXiv:1512.00516](#)] [[INSPIRE](#)].
- [50] S. Caletti et al., *Jet angularities in  $Z$ +jet production at the LHC*, *JHEP* **07** (2021) 076 [[arXiv:2104.06920](#)] [[INSPIRE](#)].
- [51] T. Kinoshita, *Mass singularities of Feynman amplitudes*, *J. Math. Phys.* **3** (1962) 650 [[INSPIRE](#)].
- [52] T.D. Lee and M. Nauenberg, *Degenerate Systems and Mass Singularities*, *Phys. Rev.* **133** (1964) B1549 [[INSPIRE](#)].
- [53] M.J. Strassler and K.M. Zurek, *Echoes of a hidden valley at hadron colliders*, *Phys. Lett. B* **651** (2007) 374 [[hep-ph/0604261](#)] [[INSPIRE](#)].
- [54] L. Carloni and T. Sjostrand, *Visible Effects of Invisible Hidden Valley Radiation*, *JHEP* **09** (2010) 105 [[arXiv:1006.2911](#)] [[INSPIRE](#)].
- [55] L. Carloni, J. Rathsman and T. Sjostrand, *Discerning Secluded Sector gauge structures*, *JHEP* **04** (2011) 091 [[arXiv:1102.3795](#)] [[INSPIRE](#)].
- [56] T. Cohen, M. Lisanti and H.K. Lou, *Semivisible Jets: Dark Matter Undercover at the LHC*, *Phys. Rev. Lett.* **115** (2015) 171804 [[arXiv:1503.00009](#)] [[INSPIRE](#)].

- [57] D. Kar and S. Sinha, *Exploring jet substructure in semi-visible jets*, *SciPost Phys.* **10** (2021) 084 [[arXiv:2007.11597](#)] [[INSPIRE](#)].
- [58] E. Bernreuther et al., *On the challenges of searching for GeV-scale long-lived particles at the LHC*, *JHEP* **04** (2021) 210 [[arXiv:2011.06604](#)] [[INSPIRE](#)].
- [59] J. Barron et al., *Unsupervised hadronic SUEP at the LHC*, *JHEP* **12** (2021) 129 [[arXiv:2107.12379](#)] [[INSPIRE](#)].
- [60] G. Albouy et al., *Theory, phenomenology, and experimental avenues for dark showers: a Snowmass 2021 report*, *Eur. Phys. J. C* **82** (2022) 1132 [[arXiv:2203.09503](#)] [[INSPIRE](#)].
- [61] Y. Wang et al., *Dynamic Graph CNN for Learning on Point Clouds*, [arXiv:1801.07829](#) [[INSPIRE](#)].
- [62] L.-H. Lim and B.J. Nelson, *What is... an Equivariant Neural Network?*, *Not. Amer. Math. Soc.* **70** (2023) 1.
- [63] V.G. Satorras, E. Hoogeboom and M. Welling, *E(n) Equivariant Graph Neural Networks*, in M. Meila and T. Zhang eds., *Proceedings of the 38th International Conference on Machine Learning, Proc. Mach. Learn. Res.* **139** (2021) 9323 [[arXiv:2102.09844](#)].
- [64] J. Gallicchio and M.D. Schwartz, *Seeing in Color: Jet Superstructure*, *Phys. Rev. Lett.* **105** (2010) 022001 [[arXiv:1001.5027](#)] [[INSPIRE](#)].
- [65] ATLAS collaboration, *Measurement of colour flow using jet-pull observables in  $t\bar{t}$  events with the ATLAS experiment at  $\sqrt{s} = 13$  TeV*, *Eur. Phys. J. C* **78** (2018) 847 [[arXiv:1805.02935](#)] [[INSPIRE](#)].
- [66] A. Larkoski, S. Marzani and C. Wu, *Safe Use of Jet Pull*, *JHEP* **01** (2020) 104 [[arXiv:1911.05090](#)] [[INSPIRE](#)].
- [67] CMS collaboration, *Search for new particles in events with energetic jets and large missing transverse momentum in proton-proton collisions at  $\sqrt{s} = 13$  TeV*, *JHEP* **11** (2021) 153 [[arXiv:2107.13021](#)] [[INSPIRE](#)].
- [68] ATLAS collaboration, *Search for new phenomena in final states with photons, jets and missing transverse momentum in pp collisions at  $\sqrt{s} = 13$  TeV with the ATLAS detector*, *JHEP* **07** (2023) 021 [[arXiv:2206.06012](#)] [[INSPIRE](#)].
- [69] C. Bierlich et al., *A comprehensive guide to the physics and usage of PYTHIA 8.3*, *SciPost Phys. Codeb.* **2022** (2022) 8 [[arXiv:2203.11601](#)] [[INSPIRE](#)].
- [70] M. Cacciari, G.P. Salam and G. Soyez, *The anti- $k_t$  jet clustering algorithm*, *JHEP* **04** (2008) 063 [[arXiv:0802.1189](#)] [[INSPIRE](#)].
- [71] M. Cacciari, G.P. Salam and G. Soyez, *FastJet User Manual*, *Eur. Phys. J. C* **72** (2012) 1896 [[arXiv:1111.6097](#)] [[INSPIRE](#)].
- [72] M. Fey and J.E. Lenssen, *Fast Graph Representation Learning with PyTorch Geometric*, [arXiv:1903.02428](#) [[INSPIRE](#)].
- [73] D.P. Kingma and J. Ba, *Adam: a Method for Stochastic Optimization*, [arXiv:1412.6980](#) [[INSPIRE](#)].
- [74] F. Pedregosa et al., *Scikit-learn: Machine Learning in Python*, *J. Machine Learning Res.* **12** (2011) 2825 [[arXiv:1201.0490](#)] [[INSPIRE](#)].