

# QINNs: Quantum-Informed Neural Networks

---

Aritra Bal\* <sup>a</sup> Markus Klute <sup>a</sup> Benedikt Maier\* <sup>b</sup> Melik Oughton <sup>b,c</sup> Eric Pezone <sup>b</sup>  
Michael Spannowsky\* <sup>d</sup>

<sup>a</sup>*Institute of Experimental Particle Physics, Karlsruhe Institute of Technology, 76131 Karlsruhe, Germany*

<sup>b</sup>*Blackett Laboratory, Imperial College of Science, Technology and Medicine, London, SW7 2AZ, United Kingdom*

<sup>c</sup>*Department of Physics and Astronomy, University College London, London WC1E 6BT, United Kingdom*

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

*E-mail:* [aritra.bal@kit.edu](mailto:aritra.bal@kit.edu), [benedikt.maier@cern.ch](mailto:benedikt.maier@cern.ch),  
[michael.spannowsky@durham.ac.uk](mailto:michael.spannowsky@durham.ac.uk)

**ABSTRACT:** Classical deep neural networks can learn rich multi-particle correlations in collider data, but their inductive biases are rarely anchored in physics structure. We propose quantum-informed neural networks (QINNs), a general framework that brings quantum information concepts and quantum observables into purely classical models. While the framework is broad, in this paper, we study one concrete realisation that encodes each particle as a qubit and uses the Quantum Fisher Information Matrix (QFIM) as a compact, basis-independent summary of particle correlations. Using jet tagging as a case study, QFIMs act as lightweight embeddings in graph neural networks, increasing model expressivity and plasticity. The QFIM reveals distinct patterns for QCD and hadronic top jets that align with physical expectations. Thus, QINNs offer a practical, interpretable, and scalable route to quantum-informed analyses, that is, tomography, of particle collisions, particularly by enhancing well-established deep learning approaches.

---

\*Corresponding author.

---

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Subjet-correlations of boosted hadronic jets</b>	<b>3</b>
<b>3</b>	<b>The QFIM-only Classifier</b>	<b>5</b>
<b>4</b>	<b>Exploiting the constituent particle correlations with graphs</b>	<b>9</b>
4.1	Tomography-enhanced graph classifiers for top-tagging	10
4.2	Inference and Results: Tomography-enhanced Graph classifier	12
<b>5</b>	<b>Quantum-informed GNN</b>	<b>14</b>
<b>6</b>	<b>Summary and Conclusions</b>	<b>16</b>

---

## 1 Introduction

Jets, collimated cascades of hadrons originating from energetic partons, provide one of the most powerful microscopes with which the LHC explores quantum chromodynamics (QCD) and searches for new physics [1]. Their internal substructure encodes a wealth of information about the underlying hard process, colour flow, and hadronisation dynamics [2–4]. When the initiating parton is a top quark with TeV-scale transverse momentum, its three hadronic decay products overlap so tightly that they are reconstructed as a single top jet [5–7]. Disentangling the physics encoded in such highly correlated multi-particle radiation patterns is central to boosted-object tagging and precision top-quark measurement [8–11]. Yet, the complexity of these correlation structures resists simple factorized descriptions and often leaves modern machine-learning classifiers difficult to interpret physically [12–15].

Jet tomography seeks to overcome this challenge by treating the energy flow inside a jet as a tomographic object, much like medical imaging reconstructs structure from multiple projections [16]. We introduce quantum-informed neural networks (QINNs), a general framework that brings quantum information concepts and observables into purely classical models; in this paper, we instantiate it as quantum-informed jet tomography built around the QFIM. The QFIM provides a basis-independent measure of the dynamical sensitivity of a quantum state and thus naturally encodes correlations between jet constituents when events are embedded in a resource-efficient one-particle-one-qubit (1P1Q) representation [17, 18]. Importantly, we find that the QFIM extracted from ensembles of QCD and top jets exhibit strikingly different structures: QCD jets appear largely diagonal, reflecting quasi-independent splittings, whereas top jets show pronounced off-diagonal elements

characteristic of their three-pronged colour-connected topology. This makes the QFIM an interpretable observable for jet physics. More broadly, we use QINN as an umbrella term for methods that inject quantum structure into classical learning while remaining fully classical at training and inference. A QINN is specified by an encoding of event information into a quantum state, for example, 1P1Q or continuous-variable maps, a choice of quantum observables or metrics, for example, the QFIM, reduced density matrices, entanglement measures, or quantum distances, and an interface to classical models, for example inputs, edge or node embeddings, regularisers, or initialization priors. This work studies one concrete member of that family: a 1P1Q encoding with the QFIM used both as an interpretable observable and as guidance for optimisation.

Recent progress in jet tagging has largely resulted from deep learning methods that learn correlations directly from data. However, these models often remain opaque and can be sensitive to training choices. Alternative approaches, including energy-flow polynomials, correlation functions, and graph-based neural networks, aim to make the underlying structures more transparent. In parallel, the growing field of quantum machine learning has introduced representations such as amplitude and qubit encodings to capture relations in high-dimensional data [19, 20]. Our study positions itself at the intersection of these developments: by using QFIM as an interpretable, quantum-native probe, we connect quantum geometry with classical machine learning to expose and exploit the correlation patterns that distinguish different jet topologies. In this context, QFIM also links naturally to optimisation via its role as a quantum metric and in natural-gradient training, and connects to known trainability phenomena in variational circuits [21–24]. Thus, we first establish the role of QFIM as an untrained classifier, then demonstrate how it can inform and accelerate a simple neural network, and finally extend the framework to GNNs enriched with both QFIM-derived and kinematic features.

This paper makes three contributions. First, it formulates QINN as a general framework for quantum-informed design of classical models. Second, it develops a QFIM-based jet tomography instantiation that exposes class-dependent correlation geometry in an interpretable way. Third, it shows how QFIM-derived quantities act as lightweight graph embeddings and training priors that improve stability and performance in GNNs. Concretely, in Sec 2 we introduce the QFIM and review the kinematic anatomy of boosted hadronic top decays, set out the qubit map and detail the construction of the QFIM as a correlation matrix.

In Sec. 3, we show that the QFIM can serve directly as a classifier without training. By vectorising the averaged QFIM over event samples, we obtain class prototypes for QCD and top jets. Each new event is then scored according to its proximity to these prototypes, yielding a powerful discriminant even in its untrained form. Building on this, we convert the QFIM-derived feature vector into the input of a shallow neural network with a single trainable encoding layer and linear output. This model learns patterns that closely track the QFIM structure itself, and crucially, initialising the network with the QFIM pattern accelerates convergence and improves optimisation dynamics. In this way, the QFIM does not just separate classes, but also informs the training of neural networks.

Having established the QFIM as a tomographic probe and an informative prior for ma-

chine learning, we extend the framework to graphs and GNNs in Sec. 4. In this representation, jet constituents are nodes and the QFIM entries act as edge features, complemented by kinematic information stored on the nodes. As shown in Sec. 5, the resulting quantum-informed GNN captures both the geometry of quantum correlations and the kinematics of the jet, achieving excellent tagging performance and exhibiting faster and smoother training than purely kinematics-based models [14, 15]. Thus, the QFIM acts both as a physically interpretable observable and as a guiding principle for network design.

Finally, we summarise our conclusions in Sec. 6. Our study demonstrates that QINNs for jet tomography unify the interpretability of quantum geometry with the scalability of modern deep learning. By embedding jets into the 1P1Q formalism and exploiting the QFIM, we expose correlation structures that distinguish top jets from the QCD background, accelerate training through informed initialisation, and enhance GNN performance. This study outlines how to integrate quantum observables into collider analyses for classification tasks. It is compatible with near-term quantum hardware and outlines an approach for how to utilise hybrid classical–quantum analysis frameworks.

## 2 Subjet-correlations of boosted hadronic jets

We want to quantify, in a basis-independent manner, how much of the dynamical information present in a boosted hadronic top jet survives the mapping from classical four-vectors to a quantum state in the 1P1Q scheme. For this task, we employ the QFIM, a quantity that sits at the intersection of differential geometry and quantum estimation theory [17, 25]. On the manifold of pure states endowed with the Fubini–Study metric [26, 27], the QFIM acts as the local metric tensor; its components therefore measure the squared geodesic distance that the state  $|\psi(\boldsymbol{\theta})\rangle$  travels under an infinitesimal change of the circuit parameters  $\boldsymbol{\theta}$ . Large entries signal directions in parameter space along which the state moves rapidly, implying high sensitivity and good trainability. In contrast, vanishing entries identify *flat* directions that can give rise to barren plateaus. Formally, for a variational state  $|\psi(\boldsymbol{\theta})\rangle$  produced by the circuit shown in Figure 1, the QFIM is defined as:


$$F_{ij} = 4 \operatorname{Re} \left[ \langle \partial_i \psi(\boldsymbol{\theta}) | \partial_j \psi(\boldsymbol{\theta}) \rangle - \langle \partial_i \psi(\boldsymbol{\theta}) | \psi(\boldsymbol{\theta}) \rangle \langle \psi(\boldsymbol{\theta}) | \partial_j \psi(\boldsymbol{\theta}) \rangle \right], \quad (2.1)$$

where  $\partial_i \equiv \partial/\partial\theta_i$ .

In our circuit, Figure 1, every jet constituent is encoded on its own qubit. Thus, the input space is a tensor product of Hilbert spaces  $H_{IN} = |q_1\rangle \otimes |q_2\rangle \cdots \otimes |q_N\rangle$ , for a system comprising  $N$  jet constituents. The trainable parameters  $\boldsymbol{\theta} = (\theta_x, \theta_y, \theta_z)$  correspond to single-qubit rotations  $R_X(\theta_x)R_Y(\theta_y)R_Z(\theta_z)$  applied after an entangling ring of nearest-neighbor CNOT gates, exactly as described in Ref. [28].

Because the final quantum state factorises only without inter-subject correlations, the QFIM automatically picks up physical structures such as colour flow and the three-prong topology of top decays.

A geometric picture helps to relate the QFIM to the LHC observable under study. Consider two neighbouring parameter points  $\boldsymbol{\theta}$  and  $\boldsymbol{\theta} + d\boldsymbol{\theta}$ . The square of the Fubini–Study line element,  $ds^2 = \frac{1}{4} \sum_{ij} F_{ij} d\theta_i d\theta_j$ , quantifies the distinguishability of the corresponding



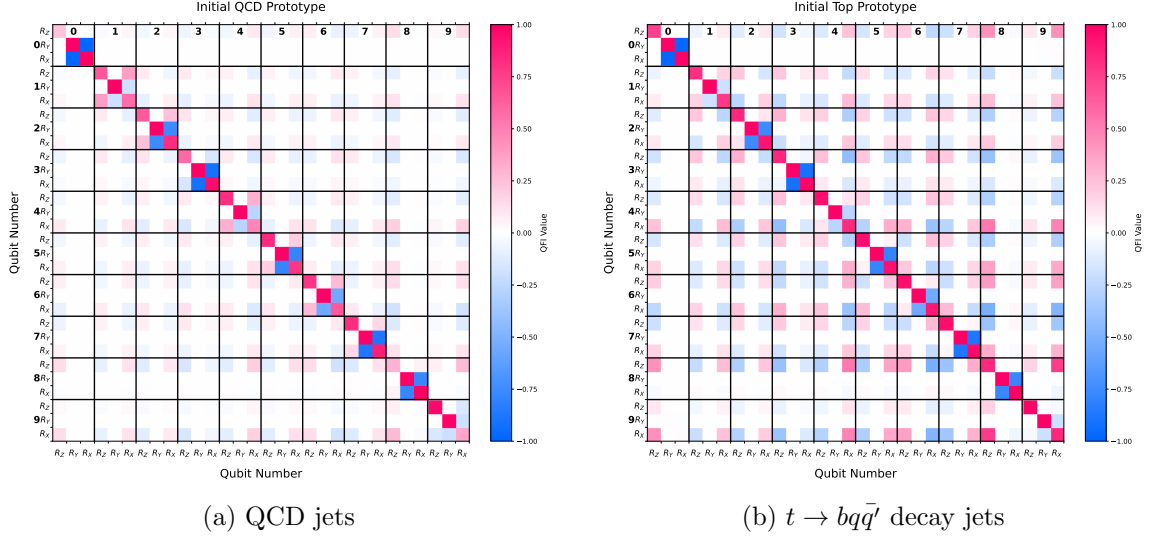
**Figure 1:** Circuit used to construct the VQC

quantum states. For jets originating from QCD splittings, whose radiation pattern is dominated by one soft/collinear branch, the state typically explores only a lower-dimensional sub-manifold of Hilbert space, leading to a QFIM whose off-diagonal entries are suppressed. In contrast, a boosted hadronic top produces three hard subjets tied together by colour connections. The required interference between three energetic qubits is reflected in sizable off-diagonal elements, especially between parameters that act on qubits representing different decay partons. In other words, the geometry encoded by  $F_{ij}$  acts as a "tomographic lens" that magnifies the characteristic triangular colour-flow pattern of top decays.

Figure 2 visualises the  $30 \times 30$  QFIM obtained by averaging over ensembles of one million simulated events, separately for ordinary QCD jets (left) and for jets matching a hadronic top hypothesis (right). For clarity, we omit the classical bias and global scale factor of the VQC, as they do not affect the quantum state itself. The heat maps corroborate the intuitive picture sketched above: QCD jets populate the diagonal stripes, indicating that each qubit evolves largely independently, whereas top jets exhibit pronounced off-diagonal structure, signalling entanglement between the three leading subjets and the softer, colour-connected radiation that surrounds them.

To facilitate a direct interpretation of inter-particle correlations, we reduce the full  $30 \times 30$  QFIM to a  $10 \times 10$  representation using the Frobenius norm  $\|C\|_F = \sqrt{\sum_{k,l} C_{kl}^2}$  for  $k \in [3i, 3i+3)$  and  $l \in [3j, 3j+3)$ . In this reduced QFIM, each matrix element  $C_{kl}^{\text{reduced}}$  provides a direct measure of the correlation strength between particles  $k$  and  $l$ , a consequence of the one particle-one qubit encoding scheme. The lines connecting particles in Figures 3c and 3d illustrate the intensity of these correlations, revealing that particles belonging to distinct prongs of the top jet exhibit stronger correlations in the QFIM latent space. The corresponding full QFIM for these representative events is shown in Figure 3a for the QCD jet and in Figure 3b for the top jet. This analysis demonstrates that the QFIM serves as a powerful diagnostic tool with inherent sensitivity to jet substructure.

From a practical standpoint, computing the full QFIM without block-diagonal approximations is essential. Only the complete matrix captures the subtle parameter inter-



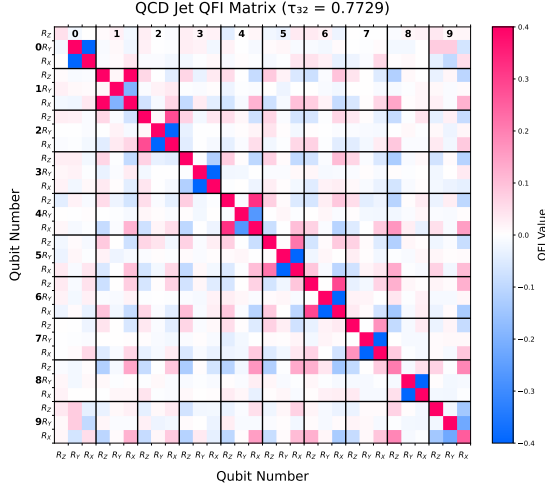
**Figure 2:** QFIM for the 30-parameter Variational Quantum Classifier circuit, implemented in PENNYLANE [29]. Left: QFIM averaged over  $10^6$  QCD jets. Right: QFIM averaged over  $10^6$  jets arising from a top quark jet. The X and Y axes show the parameters corresponding to the trainable rotations applied to each qubit used in the 1P1Q approach, where each jet constituent particle is represented by its own qubit. The matrix elements represent parameter correlations within the quantum circuit, with diagonal elements indicating individual parameter sensitivities and off-diagonal elements capturing inter-parameter dependencies.

dependencies that differentiate a three-body decay from a quasi-two-body QCD splitting, and only the complete matrix can therefore inform whether training will suffer from barren plateaus or instead benefit from rich gradients. In subsequent sections, we will exploit precisely these quantum-encoded correlations to inform the training of a GNN that excels at separating boosted top jets from the QCD background of light quark- or gluon-initiated jets.

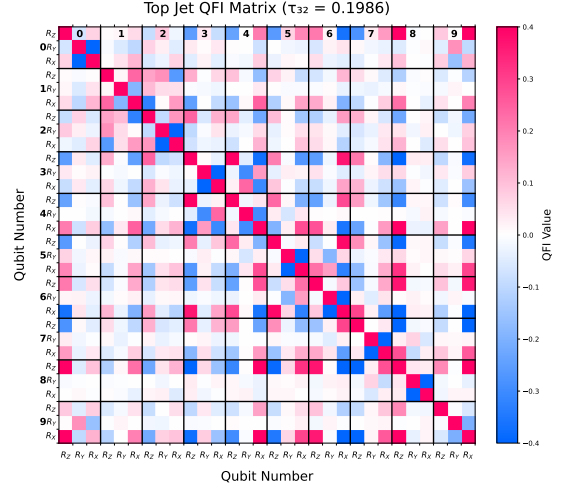
### 3 The QFIM-only Classifier

The QFIM, as computed from the Variational Quantum Classifier (VQC) of Ref. [28], provides an alternative representation of a jet that encodes essentially the same kinematic information contained in the VQC inputs. This observation suggests that the QFIM should, by itself, possess sufficient discriminatory power to distinguish between light quark/gluon-initiated jets and jets originating from the hadronic decay of top quarks. To validate this hypothesis, we begin by constructing a distance-based classifier that leverages the QFIM structure without requiring any training procedure.

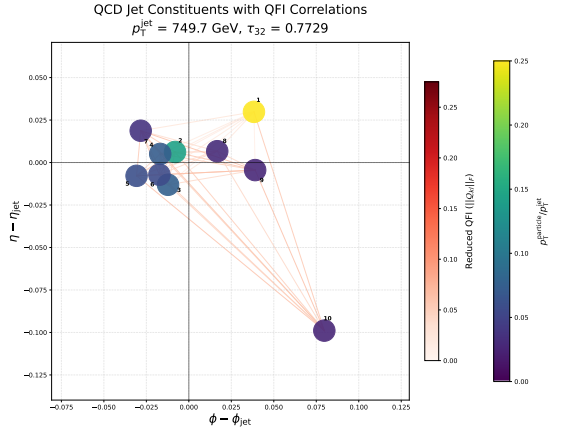
The  $N \times N$  QFIM, where  $N = 3 \times n_{\text{particles}}$ , is fully symmetric by construction. It therefore contains  $D = N(N+1)/2$  independent elements, these being the diagonal entries and either the upper or lower triangular components. Thus, each event with  $N$  particles yields a  $3N \times 3N$  QFIM  $F$  arranged as  $N \times N$  blocks  $F^{(ij)} \in \mathbb{R}^{3 \times 3}$  that quantify the



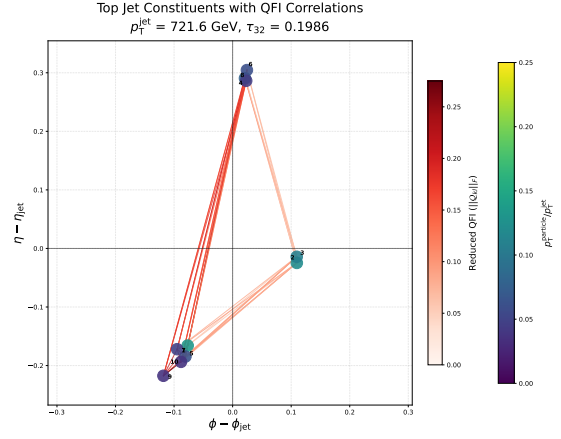
(a) QFIM for a representative QCD jet.



(b) QFIM for a representative top jet.



(c) QCD jet constituents in the  $\eta$ - $\phi$  plane, with the color of the circle representing the fraction of jet  $p_T$  carried by that constituent.



(d)  $t \rightarrow bq\bar{q}'$  jet constituents shown in the  $\eta$ - $\phi$  plane, with the difference in magnitude apparent when compared to the plot on the left.

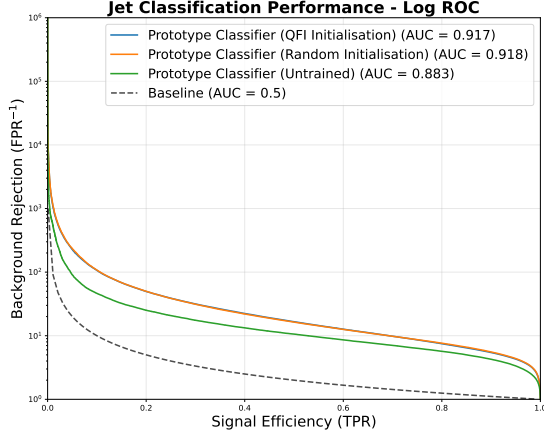
**Figure 3:** QFIMs and spatial constituent distributions for representative QCD (left) and top quark decay (right) jets, chosen using appropriate cuts on the  $n$ -subjettiness ratio  $\tau_3/\tau_2$  [30], from the  $p_T$  bin (700, 750) GeV. Top row: Full  $30 \times 30$  QFIMs with values clipped to the range  $[-0.4, 0.4]$ . Bottom row: Spatial distributions of jet constituents in the  $\eta$ - $\phi$  plane, with lines connecting constituents representing the strength of correlations derived from the reduced QFIM using the Frobenius norm, and the colour of each circle indicating the fraction of jet  $p_T$  carried by that constituent.

joint sensitivity of particles  $i$  and  $j$ , see Figs. 3a and 3b. For visualisation, we compress each block to a single edge weight  $C_{ij} = \|F^{(ij)}\|_F$  and draw the constituents in the  $\eta, \phi$  plane with an edge whose colour intensity and thickness scale with  $C_{ij}$ , while node colour reflects momentum fraction, Figs. 3c and 3d. Bright off-diagonal structure in the matrix corresponds to strong cross-prong edges in the graph. QCD jets are dominated by near




diagonal entries, Fig. 3a, indicating correlations confined within one prong and short-range links among collinear emissions; hadronic top jets exhibit pronounced off-diagonal blocks that tie the  $W \rightarrow q\bar{q}'$  pair and the  $b$  to the  $W$  system, producing a triangular network of edges across the three subjets, Fig. 3b. In this way, the QFIM provides an interpretable bridge between quantum geometry and jet kinematics, and the reduced  $C_{ij}$  graph visualises the correlation patterns that guide message passing in our GNNs of Secs. 4 and 5.

In this  $D$ -dimensional feature space, we compute two mean representative vectors, one for each jet class, which we designate as the class prototypes. For each sample in the dataset, we then calculate the squared  $L_2$  distances to these class prototypes, denoted as  $d_{\text{QCD}}$  and  $d_{\text{top}}$ , and define a score metric as  $S = d_{\text{QCD}} - d_{\text{top}}$ . This score is evidently large and positive when the jet in question lies in close proximity to the top jet prototype in feature space, and conversely becomes large and negative for jets similar to the QCD prototype. As such,  $S$  naturally serves as a classification score that can be directly converted to a probability via the sigmoid function, enabling the calculation of standard performance metrics such as accuracy and the construction of receiver operating characteristic (ROC) curves.



(a) ROC curves comparing the performance of QFIM-initialised and randomly initialised neural network classifiers, both achieving a roughly similar level of performance that is better than the untrained baseline.



(b) Loss evolution during training for QFIM-initialised (stable, rapid convergence) versus randomly initialised (unstable, slow convergence) neural network classifiers.

**Figure 4:** Performance comparison of QFIM-initialised versus randomly initialised NN-based classifiers, with the untrained classifier serving as a baseline.

This naive classifier achieves an area under the curve (AUC) score of 0.883 without any training, as shown in Figure 4a. This performance is nearly identical to that of the VQC, thereby validating our assertion that the QFIM effectively encodes the learning outcome of the aforementioned VQC.

Having established the efficacy of the QFIM as a classifier representation, we proceed to demonstrate that the QFIM is already close to being the ideal classifier in our  $D$ -dimensional feature space, which constitutes a fully quantum representation of the jet, and



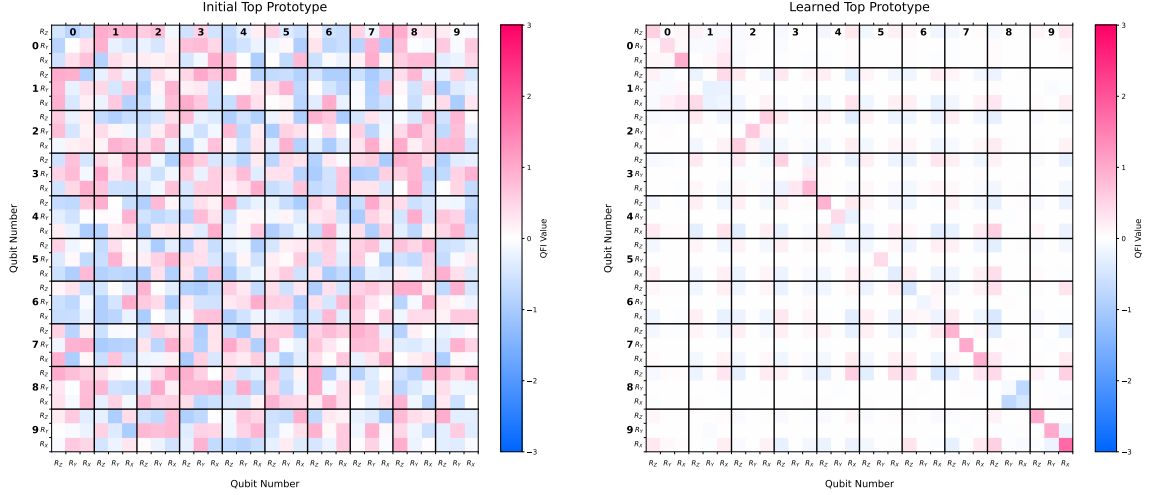
that utilising it as an initialisation point leads to substantially faster convergence and more stable loss evolution during training. To this end, we employ a shallow single-layer neural network conditioned on the individual jet QFIMs to learn the optimal class prototypes that best separate the two jet classes. Classification is performed using a distance-based approach analogous to that described above. The objective function is chosen to be the margin ranking loss [31] with a margin value of 1.0, with targets appropriately calculated from the truth labels (0 for QCD jets and 1 for top jets).

We consider two initialisation strategies for the class prototypes. In the first instance, the prototypes are initialised to the mean QFIMs for the two classes, the same values used in the naive classifier. For the baseline comparison, we initialise the class prototypes by randomly sampling from a uniform probability distribution in the interval  $(-1, 1)$ , thereby maintaining initial parity with the range of the QFIM, which is bounded within  $(-1, 1)$ . In both cases, however, the learned class prototypes are allowed to be unbounded during training, enabling the network to explore the full parameter space.

The results for these two initialisation strategies are presented in Figure 4. As shown in Figure 4a, both networks attain relatively similar classification performance, achieving an AUC score of 0.917, a modest but noticeable improvement over the untrained case described previously. However, the convergence behaviour differs starkly between the two approaches, as illustrated in Figure 4b. The QFIM-initialised network reaches loss values close to the final converged value after just one epoch and exhibits stable loss evolution throughout training. In contrast, the randomly initialised network requires more than 20 epochs to attain loss values in the same region, and moreover, experiences unstable loss dynamics during this initial training regime. Only after approximately 30 epochs does the loss stabilise and begin slowly converging toward the level achieved by the QFIM-initialised classifier. The network is updated for each batch of 256 samples. For a training dataset of size  $10^6$ , the network parameters are updated approximately 4000 times per epoch.

To further illustrate this point, we visualise the randomly initialised class prototype and the final learned class prototype in Figure 5 for the case of top jets. Remarkably, despite the random initialisation, the optimal class prototypes that best separate the two classes are very similar in structure to the mean QFIMs computed initially. This observation shows that the QFIM is indeed close to being the ideal separator between the two jet classes in this feature space.

The trainable neural networks described above are implemented in PyTorch. Optimisation is performed using the ADAM optimiser with an initial learning rate of 0.01, which is allowed to decay by a factor of 2 for every five epochs during which no improvement in the validation AUC is observed (defined as an improvement of less than 0.001). Early stopping is triggered when no such improvement in the validation AUC occurs over a period of 30 consecutive epochs. For all experiments, the training dataset comprises 1 million jets, the validation dataset (applicable only to the trainable case) consists of 200,000 jets, and the test dataset contains 1.2 million jets. In all cases, the datasets are equally divided between the two jet classes to ensure balanced training and evaluation.



(a) Randomly initialised  $t \rightarrow bq\bar{q}'$  class prototype, drawn from the uniform probability distribution in  $[0, 1]$ .

(b) Final learned  $t \rightarrow bq\bar{q}'$  class prototype, with the values allowed to be unbounded, so as to explore the entire possible phase space.

**Figure 5:** Comparison of the randomly initialised and learned final class prototypes for  $t \rightarrow bq\bar{q}'$  jets, demonstrating convergence to a structure similar to the mean QFIM. The range of the colour bars is increased so as to account for the unbounded range.

#### 4 Exploiting the constituent particle correlations with graphs

A GNN [32] is a neural network architecture that performs computation directly on graph structures, operating on vertices and edges instead of regular grids like images or sequences. Each vertex  $v_i$  carries a feature vector  $h_i^{(0)}$  (here the kinematic attributes of a single subjet, for example its rapidity, azimuth and transverse momentum). In contrast, each edge  $(i, j)$  is endowed with a weight or feature vector  $e_{ij}$  that influences how strongly the two vertices interact. The core operation of the network is message passing: at layer  $l$  every vertex gathers *messages* sent along the incident edges, aggregates them, and combines the result with its own state to produce an updated feature vector  $h_i^{(l+1)}$ . In algebraic shorthand, a broad class of GNNs can be described by the fundamental graph message-passing equation:

$$h_i^{(l+1)} = \phi^{(l)} \left( h_i^{(l)}, \text{AGG}_{j \in \mathcal{N}(i)} \psi^{(l)}(h_i^{(l)}, h_j^{(l)}, e_{ij}) \right), \quad (4.1)$$

where  $\psi^{(l)}$  is a learnable function that builds a message from the states of the two endpoints (these being the vertices themselves) and their connecting edge. AGG is a permutation-invariant reduction such as a sum or mean, and  $\phi^{(l)}$  is a learnable nonlinear update. Stacking several such layers lets information flow non-locally across the graph, and a final readout function, which in practice is also a relatively shallow neural network with a single output, maps the set of vertex states to an event-level score. In this case, this is the probability that the jet arose from a hadronic top decay rather than the QCD background.

The entries  $C_{kl}$  of the QFIM provide a natural choice for the edge feature vectors  $e_{ij}$  in the graph message function of Eq. 4.1. This direct correspondence, in the form

of a sub-matrix of the full QFIM, leverages its role in quantifying parameter estimation sensitivity to establish meaningful connectivity patterns in the GNN architecture. A large value indicates that two constituents are kinematically and dynamically linked, and vice versa. By interpreting every jet constituent particle as a vertex and every non-zero  $C_{kl}$  as the feature vector of an edge between  $v_i$  and  $v_j$ , the QFIM can play a role similar to the adjacency matrix  $\mathbf{A}$  of the graph, optionally normalised or thresholded to control sparsity. Since the circuit of the VQC first introduced in Ref. [28] uses three trainable parameters per qubit, whereas  $C_{kl} \forall k \in [3i, 3i+3), l \in [3j, 3j+3)$  represents a  $3 \times 3$  sub-matrix of the QFIM, corresponding to the two qubits  $i$  and  $j$ . During the first message-passing step, the vector quantity  $\mathbf{e} = C_{ij}$  modulates how strongly vertex  $j$ 's features affect vertex  $i$ 's update; a large positive value allows a substantial flow of information, while a small value effectively decouples the pair.

Once this graph representation is fixed, training proceeds exactly like other supervised classification tasks. Simulated or labelled real events provide examples of top-quark jets and QCD jets. For each event, the QFIM controls the graph's geometry and the interactions between nodes, the constituent kinematics furnish the initial node features, and the network parameters are adjusted to minimise a loss such as the cross-entropy between the predicted label and the truth. During inference, the process is repeated without gradient updates: construct the matrix for the unseen jet, feed the graph through the trained GNN, and obtain a probability that the event is top-like. Because the correlation structure of a boosted hadronic top, containing three hard subjets with characteristic angular separations and energy sharing, differs systematically from that of a QCD jet, the message-passing layers learn discriminative patterns that capture precisely the physics encoded in the QFIM. In this way, the visually compelling heat-map we began with is elevated from an exploratory diagnostic to the central mathematical object that drives a robust, fully differentiable tagging algorithm.

#### 4.1 Tomography-enhanced graph classifiers for top-tagging

The QFIM effectively captures the correlations between jet constituent particles in a manner that enhances both the performance and convergence properties of GNNs. Remarkably, this beneficial effect can be observed without the computational overhead of actually training the network and updating its parameters. Instead, the QFIM can be directly incorporated into the information stored at graph vertices and utilised in the graph message passing through appropriately defined mathematical transformations of the input variables, specifically the jet constituent kinematics.

This approach enables the construction of a powerful yet untrained graph-based architecture that operates solely on the jet constituent particle kinematics and the QFIM to perform inference using an appropriate distance metric such as the Mahalanobis distance [33]. The novelty of the approach proposed here lies in the inclusion of the QFIM during the graph construction phase, which, as demonstrated in the following section, substantially enhances the discriminative power between jets originating from hadronic top quark decays and the overwhelming QCD background. This improvement is quantified

through a noticeable increase in standard classification metrics, such as the AUC score and overall classification accuracy.

We adopt the input features established in the 1P1Q approach, utilising the kinematic variables associated with each jet constituent. We utilise jet samples from the JETCLASS dataset [34], which contains up to 100 million jets, divided into ten classes. For the purpose of this study, we require only the jets initiated by a light quark or gluon class, or from a  $t \rightarrow bq\bar{q}'$  decay. Each jet in turn is represented by up to the ten highest- $p_T$  constituents, thereby defining each jet as a graph containing at most 10 vertices. Since all vertices are nominally connected through directed edges, this configuration yields a maximum of  $10^2 = 100$  edges per graph. The edge features, which determine the extent of interaction between vertices along these connections, are governed by the QFIM as mentioned in the previous paragraph. The input features for each constituent are therefore defined as relative kinematic variables with respect to the parent jet:

$$\begin{aligned} p_{T,i}^{(\text{rel})} &= \frac{p_{T,i}}{p_{T,\text{jet}}} \\ \eta_i^{(\text{rel})} &= \eta_i - \eta_{\text{jet}} \\ \phi_i^{(\text{rel})} &= \phi_i - \phi_{\text{jet}} \end{aligned} \quad (4.2)$$

The initial state of the vertices comprising the graph is not directly defined by the aforementioned input variables, but rather by the following mathematical transformations of these variables, specifically designed to maximise the separation between the two classes:

$$\begin{pmatrix} v_1 \\ v_2 \\ v_3 \\ v_4 \end{pmatrix} = \begin{pmatrix} \frac{\Delta p_{T,ij}}{p_{T,i}^{(\text{rel})} + p_{T,j}^{(\text{rel})} + \epsilon} \\ \frac{\Delta \eta_{ij}}{(1 + \Delta R_{ij})^3} \\ \frac{\Delta \phi_{ij}}{(1 + \Delta R_{ij})^3} \\ \Delta R_{ij} \end{pmatrix} \quad (4.3)$$

The  $\Delta p_T$ ,  $\Delta \eta$ , and  $\Delta \phi$  terms are calculated using the relative kinematic variables defined in Eq. 4.2. However, the nature of these transformations ensures that the specific choice of reference frame does not affect the discriminative power. The  $\epsilon$  term, chosen to have a small value of  $10^{-4}$ , in the denominator of  $v_1$  ensures stability and prevents divergences during inference. Lastly,  $\Delta R_{ij} = \sqrt{(\Delta \eta_{ij})^2 + (\Delta \phi_{ij})^2}$  represents the standard distance metric between two constituents in the  $\eta$ - $\phi$  plane.

A graph utilising these four features alone serves as our non-enhanced baseline classifier. Our approach proposes to increase the vertex feature dimensionality from 4 to 7 by incorporating additional features constructed through quantum information enhancement:

$$\begin{pmatrix} v_5 \\ v_6 \\ v_7 \end{pmatrix} = \mathbf{e} \cdot \begin{pmatrix} v_1 \\ v_2 \\ v_3 \end{pmatrix} \quad (4.4)$$

Here,  $\mathbf{e}$  is a  $3 \times 3 = C_{kl}$  is the sub-matrix of the complete  $30 \times 30$  QFIM (for a jet represented using 10 particles), as described in the previous section. It accurately captures

the quantum correlations between the particles at vertices  $i$  and  $j$  that are connected by the edge  $e_{ij}$ . This approach, therefore, constitutes our quantum information-enhanced graph classifier, which will be demonstrated to exhibit clear superiority over the classical benchmark.

The function  $\psi^{(l)}$  of Eq. 4.1 can therefore be viewed as a mathematical transform that converts the vertex features defined in Equation 4.2 and the QFIM sub-matrix  $\mathbf{e}$ , whose components serve as the edge features, to the 7-dimensional vector  $\mathbf{v} = (v_1, v_2, \dots, v_7)^T$ . Excluding the last three components reduces the graph classifier to our classical benchmark, enabling a direct performance comparison. The aggregation step, represented by the function AGG, is chosen to be the permutation-invariant sum of all vertex feature vectors. This aggregation step constitutes the core of graph message passing, wherein messages are collected at each vertex from all other vertices in the graph, ensuring that the information from the entire jet substructure is incorporated into each constituent’s representation.

For this particular graph classifier, only a single message-passing layer is necessary to achieve optimal performance. Following the aggregation step, we employ graph pooling, which reduces the entire graph to a single representative vector by combining information from all vertices. Specifically, we utilise the MAX pooling operation after the graph messages have been aggregated, selecting the maximum value across each feature dimension. The final output from this graph architecture is a 7-dimensional vector, maintaining the same dimensionality as the input feature space. Crucially, the conventional non-linear transformation step that would typically be present in a standard GNN to convert the output into interpretable probabilities is not required here. This is because our entire framework represents a non-trainable approach, and the inference procedure utilising an appropriate distance metric will now be defined in the next section.

## 4.2 Inference and Results: Tomography-enhanced Graph classifier

Our input dataset consists of jets initiated either by light quarks or gluons (constituting the background class) or by the hadronic decay of top quarks ( $t \rightarrow bq\bar{q}$ ), which serves as our signal class. We initially use a training dataset comprising 800,000 jets, equally divided between the two classes. These jets are processed through the untrained graph architecture to compute exactly two class means: a pair of 7-dimensional vectors that serve as the archetypal representations of the signal and background classes, respectively.

In the subsequent step, we utilise a separately held-out dataset consisting of 1 million jets, again equally divided between both classes, which serves as our test dataset for inference. Classification is then performed using the Mahalanobis distance as the discriminating metric, enabling us to assign each jet a class probability based on whether the output jet vector lies closer in the 7-dimensional hyperspace to the signal or background class archetype. The Mahalanobis distance is defined as:

$$d^2(\mathbf{x}) = (\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu}) \quad (4.5)$$

where  $\mathbf{x}$  represents the 7-dimensional output vector from the graph classifier for a given jet, the terms  $\boldsymbol{\mu}$  denote the mean vector for a particular class (either signal or background),

and  $\Sigma^{-1}$  the inverse covariance matrix, both being computed from the training dataset.

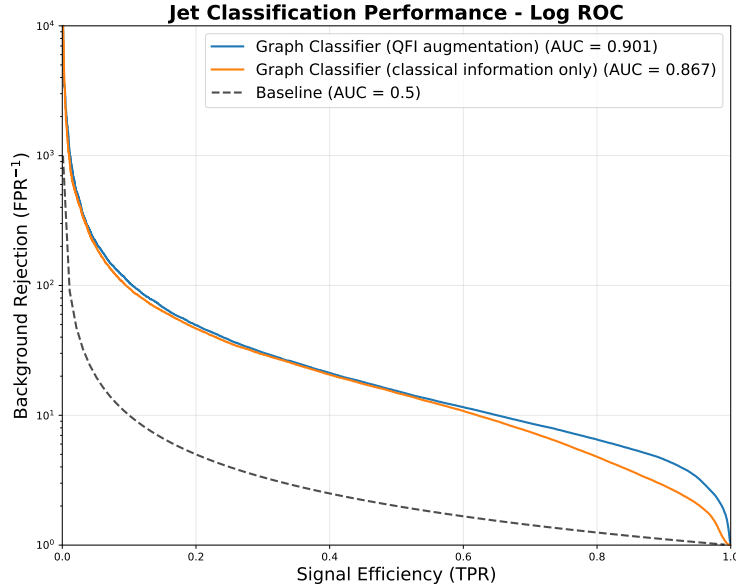
For each output jet vector, we calculate the squared Mahalanobis distances to both the mean signal ( $d_{\text{top}}^2$ ) and the mean background ( $d_{\text{QCD}}^2$ ) vectors. An additional normalization factor  $N = d_{\text{QCD}}^2 + d_{\text{top}}^2 + \epsilon$  is subsequently computed, where  $\epsilon = 10^{-8}$  is chosen for numerical stability and to prevent divergences. Class probabilities are then assigned to each jet according to:

$$P_{\text{QCD}} = \frac{d_{\text{QCD}}^2}{N} \quad (4.6)$$

$$P_{\text{top}} = \frac{d_{\text{top}}^2}{N} \quad (4.7)$$

This prescription ensures that class probabilities always sum to unity. Subsequently, classification becomes straightforward: jets lying closer to the mean top vector are classified as signal, and those closer to the background mean are classified accordingly. The proximity of a jet to its assigned class mean vector directly correlates with the confidence in that classification.

Finally, we construct the receiver operating characteristic (ROC) curve by calculating the true and false positive rates across various probability thresholds, and compute the area under this curve (AUC), which serves as a robust metric of classification performance. The comparison between the quantum-enhanced and non-enhanced versions of the graph classifier is presented in Figure 6, demonstrating that the AUC for the quantum-enhanced version clearly exceeds that of the classifier utilising purely kinematic information, thereby validating the efficacy of our proposed approach.



**Figure 6:** AUC: classification performance of untrained graph classifiers with and without QFIM-based enhancement


## 5 Quantum-informed GNN

The generalisation from an untrained graph classifier to a trained GNN is relatively straightforward and unlocks powerful capabilities that elevate classification performance to levels unattainable through simple mathematical transformations alone. The function  $\psi^{(l)}$  of Eq. 4.1 is now replaced by a convolutional neural network [35, 36] that accepts the same inputs as defined in Eq. 4.2. An additional optimisation we employ for more efficient utilisation of the input information is the creation of a compressed representation of the QFIM sub-matrix  $\mathbf{e}$  present along each edge. This compression is achieved by passing the flattened  $3 \times 3$  sub-matrix  $\mathbf{e}$  through a trainable multi-layer perceptron (MLP) [37] with three hidden layers [16, 8, 4] and a final 3-dimensional output. This 3-dimensional embedding is subsequently used as input to  $\psi^{(l)}$ .

The architecture of the convolutional neural network  $\psi^{(l)}$  is optimized through a hyperparameter search and configured as follows: a CONV1D operation with four output channels and kernel size 5, followed by an MLP with four hidden layers [64, 32, 16, 12] and a final output dimension of 3, matching the input dimensionality. The preservation of dimensionality between input and output enables the implementation of residual connections within the graph structure. Residual connections facilitate gradient flow during training and allow the network to learn identity mappings when beneficial, thereby enhancing both training stability and model expressivity. The complete message function is defined as:

$$\psi^{(l)} = \text{MLP} \left( \text{CONV1D} \left( h_i, h_j, \vec{\lambda} \cdot \text{MLP}(e_{ij}) \right) \right) \quad (5.1)$$

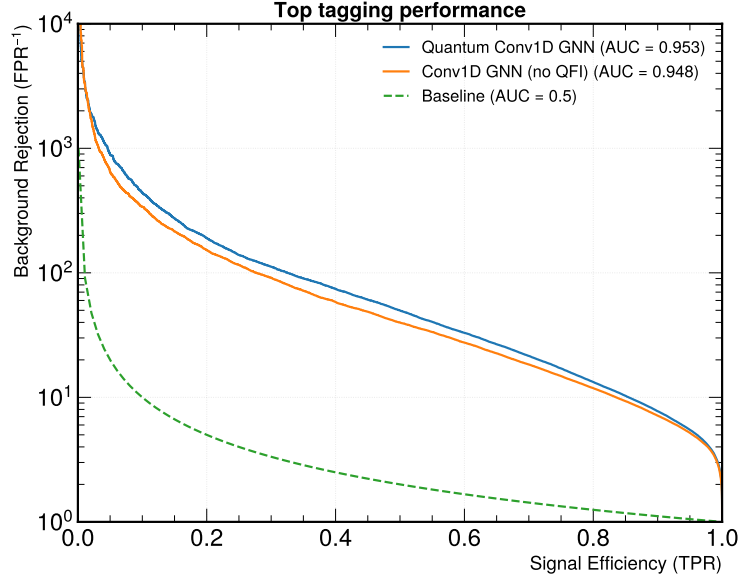
where  $\vec{\lambda}$  is a learnable temperature vector that allows the QFIM to influence the message passing operation. The network employs a total of 6 such message passing layers to capture complex multi-scale interactions that define the jet substructure. A schematic of the complete quantum information-enhanced GNN is shown in Figure 7.



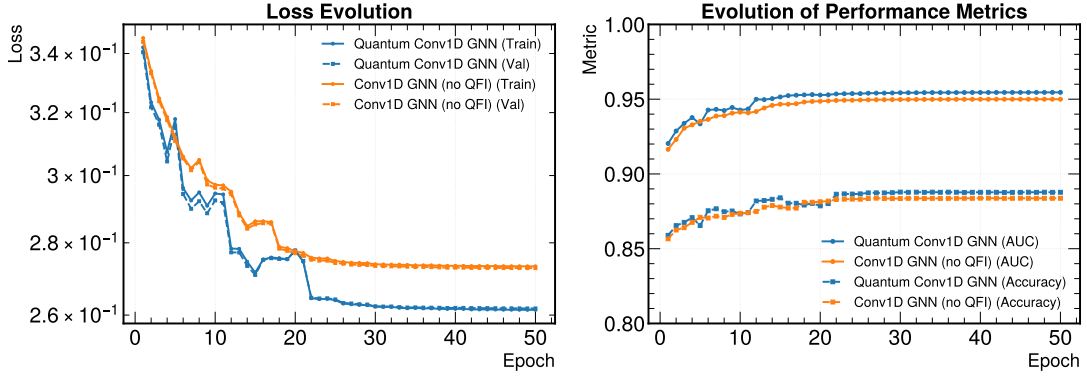
**Figure 7:** Schematic of the quantum-enhanced GNN



The classical benchmark is constructed by replacing  $\mathbf{e}$  with a null matrix, effectively removing all quantum information enhancement while maintaining identical network architecture. Following the message computation using the 1D convolutional network and embedding MLP, the graph aggregation and pooling steps remain identical to those employed in the untrained classifier: namely, a SUM aggregation followed by MAX pooling. Finally, an additional MLP, with four hidden layers [16, 12, 8, 4] followed by a softmax output layer, converts the aggregated and pooled representation into a directly interpretable probability score, following standard GNN practice. The network in both cases is trained using the Binary Cross Entropy (BCE) loss and the ADAM optimiser with a learning rate initially set to 0.005 and allowed to decay gradually, based on the improvement shown on the validation dataset. The training dataset comprises 1 million jets, whereas the validation dataset is of size 200,000 jets. The test dataset contains 1.2 million jets. In all three cases, the datasets are equally divided between the background (light quark- or gluon-initiated jets) and signal ( $t \rightarrow bq\bar{q}$  decay jets) classes.



**Figure 8:** AUC: classification performance of trained GNNs with and without QFIM-based enhancement



**Figure 9:** Evolution of metrics and loss function for both GNNs

## 6 Summary and Conclusions

Quantum-informed neural networks provide an interpretable and practically useful route to jet tomography. A one-particle-one-qubit embedding, paired with the Quantum Fisher Information Matrix, exposes physically meaningful correlations and propagates them into both simple baselines and scalable graph learning. Following the structure outlined in Sec. 1, the analysis proceeded in four steps with concrete findings at each stage. Collectively, our approach establishes a pipeline from quantum geometry to high-performing classifiers for collider data.

First, we focus on the tomographic premise. Encoding each of the leading constituents on its own qubit and computing the full, non-block-diagonal QFIM reveals class-dependent geometry in state space. Ensembles of QCD jets are characterised by QFIMs dominated by near-diagonal structure, consistent with largely quasi-independent splittings. In contrast, hadronic top jets exhibit pronounced off-diagonal elements that reflect three-pronged, colour-connected dynamics. Recovering these inter-dependencies required the complete QFIM rather than block diagonal approximations. This shows that the QFIM is a basis-independent, interpretable observable that magnifies the correlation patterns expected from the underlying colour flow.

Second, we tested whether this quantum geometric summary can separate classes without training. A prototype classifier that vectorises the per-event QFIM and scores jets by distance to class prototypes already achieves strong separation with an AUC of 0.883. Seeding a shallow network with the class mean QFIMs preserves the performance gain and, more importantly, stabilises and accelerates optimisation relative to random initialisation, with both initialisations converging to nearly identical, QFIM-like prototypes and reaching AUCs around 0.917. These observations show that the QFIM contain class information and also acts as an informative prior that improves optimisation dynamics by aligning the model with the intrinsic geometry of the data.

Third, we elevated the tomographic object into a fixed, non-trainable graph representation. Treating constituents as nodes and inserting QFIM sub-matrices as edge features creates an informed adjacency that guides message passing. With node features built from

simple relative kinematics, event-level inference via a Mahalanobis distance rule improves the AUC from 0.867 for a kinematics-only baseline to 0.901 when QFIM features are included. The construction uses up to ten constituents per jet and evaluates on held-out samples at the million-event scale, so the observed gain can be traced to the quantum geometric correlations themselves rather than to learned parameters.

Finally, we generalised to a trained, quantum-informed GNN. The message function combines a CONV1D operation on node pairs with a compact, learned embedding of the  $3 \times 3$  QFIM block attached to each edge, with residual connections used to maintain stable updates. Relative to an otherwise identical architecture with the QFIM input nullified, the quantum informed network improves top tagging AUC from 0.948 to 0.953 and converges faster and more smoothly in both loss and validation metrics. These results indicate that quantum encoded correlations help both accuracy and optimisation behaviour, while integrating naturally with standard graph learning practice.

Thus, the 1P1Q embedding and the QFIM reveal class-dependent correlation geometry, the untrained prototype confirms that the geometry alone separates classes and shapes the optimisation landscape, and the fixed graph demonstrates that QFIM augmentation sharpens classical summaries in a parameter-free setting. The trained GNN shows that quantum-informed edges translate into measurable performance gains and improved training dynamics at scale. In all cases, the quantum algorithms are compact and shallow, which preserves compatibility with near-term devices and keeps the approach attractive for hybrid classical quantum setups.

Beyond metrics, our approach shows that the QFIM provides an interpretable observable and a design prior for neural networks. As a quantum metric, the QFIM connects directly to sensitivity and natural gradient notions, so it offers a principled way to encode physics-aligned inductive bias into classical architectures. The empirical stability and faster convergence we observe are consistent with this view and suggest that quantum geometry can regularise optimisation while maintaining a clear physical interpretation of what the network learns.

We find that QINNs for jet tomography expose and exploit correlation structures in a way that is both interpretable and scalable. They improve classical baselines in fixed and trainable settings, and they remain compatible with near-term hardware. This combination of interpretability, performance, and practicality approaches a credible candidate for use in collider data analysis.

**Acknowledgement:** The authors acknowledge the support of Schmidt Sciences, the Alexander von Humboldt Foundation, and IPPP Durham, as well as the usage of computing resources on the **deeptthought** GPU cluster at the Institute of Experimental Particle Physics (ETP), KIT, and the RCS Cluster at Imperial College London.

## References

- [1] S. Marzani, G. Soyez, and M. Spannowsky, *Looking inside jets: an introduction to jet substructure and boosted-object phenomenology*, vol. 958. Springer, 2019. [[arXiv:1901.10342 \[hep-ph\]](#)].

- [2] J. Gallicchio and M. D. Schwartz, “Seeing in color: Jet superstructure,” *Phys. Rev. Lett.* **105** (2010) 022001, [[arXiv:1001.5027 \[hep-ph\]](#)].
- [3] A. J. Larkoski, I. Moulton, and D. Neill, “Energy correlation functions for jet substructure,” *JHEP* **12** (2013) 009, [[arXiv:1307.0007 \[hep-ph\]](#)].
- [4] A. J. Larkoski, I. Moulton, and D. Neill, “Power counting to better jet observables,” *JHEP* **12** (2014) 009, [[arXiv:1402.2657 \[hep-ph\]](#)].
- [5] J. M. Butterworth, A. R. Davison, M. Rubin, and G. P. Salam, “Jet Substructure as a New Higgs Search Channel at the LHC,” *Phys. Rev. Lett.* **100** (2008) 242001, [[arXiv:0802.2470 \[hep-ph\]](#)].
- [6] D. E. Kaplan, K. Rehermann, M. D. Schwartz, and B. Tweedie, “Top tagging: A method for identifying boosted hadronically decaying top quarks,” *Phys. Rev. Lett.* **101** (2008) 142001, [[arXiv:0806.0848 \[hep-ph\]](#)].
- [7] T. Plehn, G. P. Salam, and M. Spannowsky, “Fat Jets for a Light Higgs,” *Phys. Rev. Lett.* **104** (2010) 111801, [[arXiv:0910.5472 \[hep-ph\]](#)].
- [8] D. E. Soper and M. Spannowsky, “Finding top quarks with shower deconstruction,” *Phys. Rev. D* **87** (2013) 054012, [[arXiv:1211.3140 \[hep-ph\]](#)].
- [9] D. E. Soper and M. Spannowsky, “Finding physics signals with event deconstruction,” *Phys. Rev. D* **89** no. 9, (2014) 094005, [[arXiv:1402.1189 \[hep-ph\]](#)].
- [10] A. J. Larkoski, I. Moulton, and D. Neill, “Jet Substructure at the Large Hadron Collider: A Review of Recent Advances,” *Phys. Rept.* **841** (2020) 1–63, [[arXiv:1709.04464 \[hep-ph\]](#)].
- [11] R. Kogler *et al.*, “Jet substructure at the large hadron collider: Experimental review,” *Rev. Mod. Phys.* **91** no. 4, (2019) 045003, [[arXiv:1803.06991 \[hep-ex\]](#)].
- [12] J. Cogan, M. Kagan, E. Strauss, and A. Schwarcz, “Jet-images: Computer vision inspired techniques for jet tagging,” *JHEP* **02** (2015) 118, [[arXiv:1411.2738 \[hep-ph\]](#)].
- [13] P. T. Komiske, E. M. Metodiev, and J. Thaler, “Energy flow networks: Deep sets for particle jets,” *JHEP* **01** (2019) 121, [[arXiv:1810.05165 \[hep-ph\]](#)].
- [14] H. Qu and L. Gouskos, “ParticleNet: Jet Tagging via Particle Clouds,” *Phys. Rev. D* **101** no. 5, (2020) 056019, [[arXiv:1902.08570 \[hep-ph\]](#)].
- [15] J. Shlomi, P. Battaglia, and J.-R. Vlimant, “Graph Neural Networks in Particle Physics,” *Mach. Learn.: Sci. Technol.* **2** no. 2, (2021) 021001, [[arXiv:2007.13681 \[hep-ex\]](#)].
- [16] F. A. Dreyer, G. P. Salam, and G. Soyez, “The Lund Jet Plane,” *JHEP* **12** (2018) 064, [[arXiv:1807.04758 \[hep-ph\]](#)].
- [17] S. L. Braunstein and C. M. Caves, “Statistical Distance and the Geometry of Quantum States,” *Phys. Rev. Lett.* **72** (1994) 3439–3443.
- [18] J. Liu, H. Yuan, X. Lu, and X. Wang, “Quantum Fisher Information Matrix and Multi-Parameter Estimation,” *J. Phys. A* **53** no. 2, (2020) 023001, [[arXiv:1907.08037 \[quant-ph\]](#)].
- [19] V. Havlíček *et al.*, “Supervised learning with quantum-enhanced feature spaces,” *Nature* **567** (2019) 209–212, [[arXiv:1804.11326 \[quant-ph\]](#)].
- [20] M. Schuld and N. Killoran, “Quantum Machine Learning in Feature Hilbert Spaces,” *Phys. Rev. Lett.* **122** (2019) 040504, [[arXiv:1803.07128 \[quant-ph\]](#)].

- [21] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, “Quantum natural gradient,” *Quantum* **4** (2020) 269, [[arXiv:1909.02108 \[quant-ph\]](#)].
- [22] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, “Barren plateaus in quantum neural network training landscapes,” *Nature Commun.* **9** (2018) 4812, [[arXiv:1803.11173 \[quant-ph\]](#)].
- [23] J. Preskill, “Quantum Computing in the NISQ era and beyond,” *Quantum* **2** (2018) 79, [[arXiv:1801.00862 \[quant-ph\]](#)].
- [24] A. Blance and M. Spannowsky, “Quantum Machine Learning for Particle Physics using a Variational Quantum Classifier,” *JHEP* **02** (2021) 212, [[arXiv:2010.07335 \[hep-ph\]](#)].
- [25] O. E. Barndorff-Nielsen and R. D. Gill, “Fisher information in quantum statistics,” arXiv preprint [quant-ph/9808009](#) (1998) .
- [26] G. Fubini, “Sulle metriche definite da una forma hermitiana,” *Atti del Reale Istituto Veneto di Scienze, Lettere ed Arti* **63** (1904) 501–513.
- [27] E. Study, “Kürzeste wege im komplexen gebiet,” *Mathematische Annalen* (1905) .
- [28] A. Bal, M. Klute, B. Maier, M. Oughton, E. Pezone, and M. Spannowsky, “One particle - one qubit: Particle physics data encoding for quantum machine learning,” *Phys. Rev. D* **112** (Oct, 2025) 076004. <https://link.aps.org/doi/10.1103/18y2-87vq>.
- [29] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, S. Ahmed, V. Ajith, M. S. Alam, G. Alonso-Linaje, B. AkashNarayanan, A. Asadi, *et al.*, “PennyLane: Automatic differentiation of hybrid quantum-classical computations.” Arxiv preprint [arxiv:1811.04968](#), 2018.
- [30] J. Thaler and K. Van Tilburg, “Identifying Boosted Objects with N-subjettiness,” *JHEP* **03** (2011) 015, [[arXiv:1011.2268 \[hep-ph\]](#)].
- [31] PyTorch Documentation, “`torch.nn.MarginRankingLoss`.” <https://docs.pytorch.org/docs/stable/generated/torch.nn.MarginRankingLoss.html>. Accessed: 2025-10-15.
- [32] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, “The Graph Neural Network model,” *IEEE Transactions on Neural Networks* **20** no. 1, (2008) 61–80.
- [33] P. C. Mahalanobis, “On the generalised distance in statistics,” *Sankhya A* **80** no. Suppl 1, (2018) 1–7. <https://doi.org/10.1007/s13171-019-00164-5>. Reprint of 1936 original.
- [34] H. Qu, C. Li, and S. Qian, “Jetclass: A large-scale dataset for deep learning in jet physics,” June, 2022. <https://doi.org/10.5281/zenodo.6619768>. Zenodo, version 1.0.0.
- [35] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel, “Handwritten digit recognition with a back-propagation network,” *Advances in Neural Information Processing Systems* **2** (1989) 396–404.
- [36] Schmidhuber, Jürgen, “Deep learning in neural networks: An overview,” *Neural Networks* **61** (2015) 85–117.
- [37] F. Rosenblatt, “The perceptron: A probabilistic model for information storage and organization in the brain,” *Psychological Review* **65** no. 6, (1958) 386–408.