## Fermilab

Fast Jet Tagging with MLP-Mixers on FPGAs

FERMILAB-PUB-25-0180-PPD

arXiv:2503.03103

This manuscript has been authored by Fermi Forward Discovery Group, LLC under Contract No. 89243024CSC000002 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

### Fast Jet Tagging with MLP-Mixers on FPGAs

# Chang Sun<sup>1,\*</sup>, Jennifer Ngadiuba<sup>2</sup>, Maurizio Pierini<sup>3</sup>, Maria Spiropulu<sup>1</sup>

<sup>1</sup>California Institute of Technology, Pasadena, CA, USA <sup>2</sup>Fermi National Accelerator Laboratory, Batavia, IL, USA <sup>3</sup>European Organization for Nuclear Research (CERN), Geneva, Switzerland

E-mail: chsun@cern.ch

Abstract. We explore the innovative use of MLP-Mixer models for real-time jet tagging and establish their feasibility on resource-constrained hardware like FPGAs. MLP-Mixers excel in processing sequences of jet constituents, achieving state-of-the-art performance on datasets mimicking Large Hadron Collider conditions. By using advanced optimization techniques such as High-Granularity Quantization and Distributed Arithmetic, we achieve unprecedented efficiency. These models match or surpass the accuracy of previous architectures, reduce hardware resource usage by up to 97%, double the throughput, and half the latency. Additionally, non-permutation-invariant architectures enable smart feature prioritization and efficient FPGA deployment, setting a new benchmark for machine learning in real-time data processing at particle colliders.

#### 1. Introduction

The CERN Large Hadron Collider (LHC) [1] generates vast quantities of data, producing hundreds of terabytes per second from proton-proton collisions occurring every 25 ns. The first reduction of this enormous data throughput is managed by the hardware-based level-1 trigger system (L1T), which filters events in real-time with stringent latency requirements of a few microseconds [2, 3]. This system, utilizing approximately 1,000 Field-Programmable Gate Arrays (FPGAs), ensures critical events are preserved for offline analysis while drastically reducing the data rate by two orders of magnitude, thereby managing the downstream bandwidth effectively. In this context, the algorithms' accuracy is crucial while the computational complexity of each algorithm must be kept minimal as the resources are scarce. The forthcoming High-Luminosity LHC (HL-LHC) [4] upgrade will significantly increase the data rate and complexity, posing challenges for real-time data processing. Machine learning techniques are being actively explored to enhance the speed and accuracy of the trigger algorithms [2, 3]; however, addressing strict resource and latency constraints without compromising performance present important challenges.

Jet tagging is a crucial task for physics analyses at the LHC, which enables the identification of rare events that probe fundamental questions about the universe, such as the existence of new particles or the verification of theoretical models. This task involves identifying the origin of jets, which are collimated sprays of particles produced in high-energy collisions, and associating them with specific underlying physics processes. While jet tagging is traditionally performed offline using large machine learning models and high-performance computing resources [5, 6, 7, 8, 9], extending this capability to the trigger system ensures that such valuable events are not discarded prematurely. However, this transition is highly challenging due to the stringent constraints of the trigger system. Offline models, designed for accuracy without hardware constraints, often exceed the memory, processing speed, and on-chip resource availability of FPGAs, necessitating significant architectural adaptations and optimizations to make them viable for real-time deployment.

MLP-Mixer [10] is an architecture consisting solely of Multi-layer Perception (MLP) layers, originally designed for vision tasks. For those applications, it splits the input into patches and processes them as a sequence of vectors. As the inputs for jet-tagging tasks at the trigger level are already sequences of particles constituents, MLP-Mixer offers a plausible structure for online jet tagging tasks. In this work, we explore and demonstrate the application of MLP-Mixers for jet-tagging, as well as their deployment on FPGAs for use in low-level trigger systems. We train the models with High-Granularity Quantization (HGQ) [11] to fit within FPGA constraints, and further optimizes their firmware with Distributed Arithmetic (DA) to minimize the latency and maximizing throughput. The firmware is generated and deployed using the hls4ml [12] framework with the Vitis HLS [13] backend, ensuring compatibility and efficiency on FPGA hardware.

This work makes the following contributions:

- We show that MLP-Mixer models achieve state-of-the-art accuracy on a jet tagging dataset generated to closely resemble realistic LHC data.
- We demonstrate the deployment of MLP-Mixer models on FPGAs, achieving low latency and high throughput, meeting the requirements of the low-level trigger system.
- We highlight that MLP-Mixer models outperform prior architectures in accuracy while using significantly fewer resources, achieving higher throughput, and lower latency.
- We emphasize the advantage of using a non-permutation-invariant model for jet tagging when leveraging HGQ, which allows for selective feature prioritization via heterogeneous activation bitwidths to minimize resource consumption.

#### 2. Background and Related Works

#### 2.1. Fast Jet tagging on FPGAs

Jet tagging is one of the most critical and well-studied tasks in high-energy physics experiments, with various machine learning algorithms proposed to address it [5, 6, 7, 8, 9, 14]. State-of-the-art models for jet tagging include those based on Graph Neural Networks (GNN), such as Particle Net [5] and PELLICAN [7], or Transformer architectures like Particle Transformer [6] and the Lorentz-Equivariant Geometric Algebra Transformers [8].

While these models demonstrate excellent performance for offline data processing, they are typically too large and computationally intensive for real-time deployment on FPGA-based trigger systems. Research efforts [15, 16, 17, 18, 19, 12] have focused on adapting machine learning models for FPGA deployment in real-time environments, showcasing their potential for ultra-low latency and high throughput inference. Notably, Ref. [16] demonstrated that quantization-aware training (QAT) could enable jet tagging with quantized neural networks on FPGAs, achieving competitive accuracy to the full precision models with significantly reduced resource usage. However, these efforts often relied on toy datasets with high-level jet features which are unavailable in real-time trigger systems, making them impractical for actual deployment.

For real-time jet tagging on FPGAs, models like Particle Net and Particle Transformer are quickly ruled out due to their size and complexity. More resourceefficient alternatives include JEDI-net [9], an efficient architecture based on Interaction Networks (IN) [20], and models using the Deep Sets (DS) architecture [21], are being explored. JEDI-net has been successfully implemented on FPGAs [22] with realistic latency and resource consumption, while the DS-based model achieved better latency and resource utilization than IN-based models under comparable hardwareaware constraints [23]. These advancements highlight the feasibility of deploying highlyperformant particle-based machine learning models on FPGAs for jet tagging, but further innovations are needed to balance performance, resource efficiency, and latency.

#### 2.2. hls4ml

hls4ml [12] is a framework that translates trained neural networks into C++ based highlevel synthesis (HLS) projects, which can be further compiled by HLS backends into description languages such as Verilog or VHDL. It supports a wide range of neural network architectures and HLS backends. When the input neural network is properly quantized and configured, hls4ml may generate firmware that is bit-accurate with the original model. This framework has been widely used for deploying neural networks on the hardware triggers at the LHC [16, 17, 15, 24, 25, 26, 27, 2, 3]. Notably, it has been employed for real-time data processing in the L1 trigger system of the CMS experiment, demonstrating its reliability and effectiveness [26, 27]. In this case, hls4ml has facilitated the implementation of autoencoders based anomaly detection models on FPGAs for physics-based applications, achieving remarkable success during LHC Run-3 data taking.

In this work, we leverage hls4ml with the HGQ [11] front-end and the Vitis HLS backend to generate the firmware for the MLP-Mixer models. This combination ensures bit-accuracy between the quantized python model and the generated firmware, while optimizing for resource efficiency on FPGAs.

#### 2.3. MLP-Mixer

MLP-Mixer [10] is a neural network architecture that exclusively uses Multi-layer Perception (MLP) layers, which is originally designed for vision tasks. For a vision task, the model first splits the image into patches, and treat the flattened vector from each patch as the feature embeddings. The core architecture of the networks involves a series of MLP layers applied along orthogonal directions to mix both spatial- and channel-wise information. Despite its simplicity, this approach has shown competitive performance in computer vision tasks. Although initially developed for image data, we found the MLP-Mixer architecture is well-suited for jet tagging tasks in high-energy physics. Since the input to the jet tagging model is already a sequence of vectors of particle features at L1 triggers, we directly apply the MLP-Mixer architecture on this data without additional modifications. The absence of complex mechanisms, such as selfattention in Transformers involving multiplication of variable-to-variable multiplication operations, simplifies the conversion of MLP-Mixer models with h1s4m1 and therefore their deployment on FPGAs and its optimization with DA.

#### 3. Setup

#### 3.1. Dataset

We use the hls4ml jet tagging dataset [28] for training and evaluation. This dataset comprises 620,000 jets in the training set and 260,000 jets in the validation set. The dataset is balanced across 5 classes: Jets originating from gluons (g), light quarks (q), W bosons (W), Z bosons (Z), and top quarks (t). Each jet contains up to 150 particles. Each particle is characterized by 16 kinematic features, as documented in Table 1. The distributions of particle count for each class are shown in Figure 1(a), with the average number of particles per jet varying from 38 to 67, depending on the class.

In alignment with trigger system constraints, only the top-N particles ordered by  $p_T$  are considered for real-time jet tagging tasks, with N varying up to a realistic constraint.

Furthermore, one may apply a typical preselection requirement of a minimum transverse momentum  $(p_T)$  of 2 GeV per particle to simulate the constraints of the trigger system at HL-LHC [30]. With this  $p_T$  threshold, the average particle count per jet decreases to a range of 31 to 49 as shown in Figure 1(b). If a jet contains fewer than N particles, we pad the missing features with zeros.

Je Du t	
Feature	Description
$p_x$	The $x$ component of the momentum of the particle.
$p_y$	The $y$ component of the momentum of the particle.
$p_z$	The $x$ component of the momentum of the particle.
E	The energy of the particle.
$E_{\rm rel}$	$E/E_{\rm jet}$ , relative energy of the particle.
$p_T$	Transverse momentum of the particle.
$p_{T_{\rm rel}}$	$p_T/p_{Tjet}$ , relative transverse momentum of the particle.
$\eta$	Pseudorapidity of the particle.
$\eta_{ m rel}$	$\eta - \eta_{\rm jet}$ , relative pseudorapidity of the particle.
$\eta_{ m rot}$	$\eta$ rotated as described in [29]
$\phi$	Azimuthal angle of the particle.
$\phi_{ m rel}$	$\phi - \phi_{\rm jet}$ , relative azimuthal angle of the particle.
$\phi_{ m rot}$	$\phi$ rotated as described in [29]
$\Delta R$	$\sqrt{(\eta - \eta_{\rm jet})^2 + (\phi - \phi_{\rm jet})^2}$ , $\Delta R$ between the particle and the jet.
$\cos  heta$	$\cos(\theta)$ , cosine of the polar angle of the particle.
$\cos \theta_{ m rel}$	$\cos(\theta - \theta_{\rm jet})$ , the cosine of relative polar angle of the particle.

**Table 1.** Description of the 16 kinematic features used for each particle in the hls4ml jet tagging dataset.

In this study, we use 558,000 jets from the training set for model training, reserving the remaining 62,000 jets for validation. The remaining 260,000 jets from the validation set are used exclusively for testing.

Currently, the JEDI-net models are considered the state-of-the-art for this dataset, and we use them as the baseline for comparison in our study.



**Figure I.** Distribution of the number of particles per jet in the jet tagging dataset. On the right, only particles with  $p_T \ge 2$  GeV are considered.

At the CMS Experiment [2], we expect the information of up to 128 particles to be available in the trigger system [30]. However, the number of particles required for downstream tasks, varies depending on the application. For instance, while it has been shown that a good performance in b-tagging requires only  $\sim 10$  particles [31], we have observed that the jet-tagging task on this dataset benefits from a larger number of particles in the input. Also, event-level tasks, such as missing transverse energy estimation or anomaly detection, may benefit from a larger input to represent the full event. To account for these variations, we train the models with 16, 32, 64, and 128 particles to cover the expected range of inputs.

#### 3.2. Model

We adopt the MLP-Mixer architecture proposed in [10] as the backbone of our model. However, since the original model is too large for practical FPGA implementation, we designed a downscaled version with only two mixer stages and a dense classification head. Nevertheless, we find that this simplified model is still capable to achieve stateof-art-performance. The exact architecture is depicted in Figure II. Following [22], we use all 16 particle features available in the dataset, and train the models with varying number of particles: 16, 32, 64, and 128.

As a baseline for comparison, we also train an MLP model with three hidden layers of size 64 with the same number of particles in the input. The MLP model is trained and synthesized with the same configuration and optimizations as the MLP-Mixer models.

In contrast to [22], we do not enforce permutation invariance on our models, and all inputs are ordered by  $p_T$ . This design choice is made for several reasons:

- In practical trigger systems, particles are often already sorted by  $p_T$ , to meet the requirements of parallel algorithms running on the same FPGA.
- Allowing the model to break permutation invariance enables it to learn which features to retain and which to discard, facilitating heterogeneous bitwidths optimization on its activations.
- While permutation invariance is ensured in previous works on the top-N elements, a preprocessing step is still required to extract the top-N elements by  $p_T$ . Depending on the exact implementation, the resource overhead of sorting the whole sequence and taking the top-N elements are likely to be comparable.

#### 3.3. Training

All models are implemented and trained with HGQ [11] on TensorFLow 2.13. We use the Adam optimizer with an initial learning rate of  $5 \cdot 10^{-3}$ , along with the CosineDecayRestarts scheduler that resets the learning rate every 100 epochs. The full-precision models are trained for 500 epochs with a batch size of 512 with the categorical cross-entropy loss. The models are trained on a single NVIDIA GTX 1080 GPU with 8GB of memory. The training time for each model ranges between



Figure II. Architecture of the MLP-Mixer models used in this work. Each model consists of four MLP blocks with a single skip-connection. The input channel size match the number of per-particle input features. DenseBn represents a dense layer followed by a batch normalization layer during training, which are fused into a single layer during inference.

approximately 90 and 120 minutes<sup>‡</sup>, depending on the number of particles in the input. Only the models on the Pareto Frontier of the best validation accuracy and resource consumption estimated by Effective Bit Operations (EBOPs) [11] are selected and used to evaluate the performance on the test dataset.

#### 4. Quantization and Compression

Floating-point representations are commonly used for weights and activations in neural networks. However, such representations are not suitable for FPGA implementation due to their high computational complexity of multiplication and accumulation operations. To deploy models efficiently on FPGAs, we quantize all neural network parameters to fixed-point representations.

We adopt unified QAT and pruning provided by HGQ to make the model more compact and optimized for hardware deployment. Firmware generation is performed using hls4ml and Vitis HLS. DA optimization is performed using da4ml during firmware generation. The complete workflow for generating FPGA firmware is illustrated in Figure III.

#### 4.1. Optimization Techniques

We adopt the following techniques to optimize the model for FPGA deployment.

4.1.1. High Granularity Quantization We adopt the HGQ [11] method to train, quantize, and prune the MLP-Mixer models. HGQ is an advanced QAT method that optimizes the bitwidths of each individual parameter of the neural network with

<sup>‡</sup> Only as a reference. After code optimization, it was found that the training time can be reduced by a factor of 10.



Figure III. Workflow for training and deploying the MLP-Mixer models.

surrogate gradients on bitwidths. Comparing to traditional QAT methods that optimize the bitwidths at the layer level, such as QKeras [16] or Brevitas [32], HGQ provides a more fine-grained optimization that can achieve significant resource usage reduction for a fully unrolled firmware design while maintaining the accuracy of the original fullprecision model. As HGQ allow the bitwidths to reach zero, unstructured pruning is naturally provided. HGQ is a crucial component of our workflow, which ensures that the MLP-Mixer models can be successfully deployed on FPGAs, as synthesis becomes infeasible if the model is not sufficiently quantized.

4.1.2. Distributed Arithmetic DA [33, 34, 35, 36, 37] is a method that transforms a linear digital signal processing (DSP) problem with one constant term, such as finite impulse response filters, discrete sine transform, or discrete cosine transform, into a series of operations implemented without the need for dedicated multipliers. This typically involves decomposing constant multiplication operations into a sequence of additions/subtractions and bit-shifts. Optimizations can be applied to the DA implementation to reduce the number of operations and the resource usage. In this work, we adopt a variant of the constant matrix-vector multiplication (CMVM) algorithm proposed in [34] to optimize the CMVM operations in the MLP-Mixer models. An outline of the optimization is as follows:

- (i) Decompose all constants in the operand matrix using the canonical signed digit representation, and express the CMVM operation as a sequence of additions/subtractions and bit-shifts.
- (ii) Iteratively search for repeated two-term subexpressions and replace them with a single variable.
- (iii) Trace the necessary bitwidths for all intermediate variables and the final results to prevent overflows.

We use the da4ml [38] library as an extension to the hls4ml library to apply this optimization. By applying the stated DA optimization, we significantly reduce the latency and initiation interval of the models.

4.1.3. Fused Batch Normalization Since a dense layer performs a general affine transformation and a batch normalization (BN) layer applies a more restricted affine transformation, a BN layer immediately following a dense layer can be fused with the latter by scaling the kernel and shifting the bias accordingly. However, because the weights of the dense layer are expected to be highly quantized, post-training fusion may lead to suboptimal accuracy and/or resource consumption overheads. To address this, we use the fused DenseBatchnorm layer from the HGQ library, which fuses the BN layers into the dense layer during training. This ensures that quantization bitwidths are optimized for the fused layers. As a result, both bitwidths and model weights are tuned appropriately, preventing any degradation in performance or resource efficiency.

#### 4.2. Training

Following [11], we use a single training for each model to cover the entire Pareto frontier between accuracy and resource consumption. We initialize the models with a small  $\beta$  – a scalar parameter that encourages smaller models in a hardware resource-aware approach – and increase its value during training with an exponential scheduler. During training, we save all models that lie on the Pareto frontier between validation accuracy and the resource consumption estimated by the approximated Effective Bit-Operations (EBOPs) [11] provided by the HGQ framework. All other hyperparameters remain identical to those used in the full-precision training. Each model is trained for 5,000 epochs, requiring approximately 22 hours on a single NVIDIA GTX 1080 GPU. However, note that this lengthy training is only required for covering the entire Pareto frontier, and the training time can be significantly reduced if only one performing model is required. For instance, all best-performing quantized models depicted in this work are obtained within 400 epochs.

#### 5. Results

In this section, we first evaluate the performance of the full-precision MLP-Mixer models on the jet tagging dataset and compare them with the JEDI-net models [9] using different numbers of input particles. We then analyse the performance, resource consumption, latency, and throughput of the deployable firmware for both MLP-Mixer and MLP models on the same dataset. Additionally, we compare quantized models with their fullprecision counterparts and with previous studies to highlight the advantages of using MLP-Mixer-based architectures for this task. For fairness, the  $p_T \geq 2$  GeV selection was not applied to the input particles when comparing our results with [22] and [9].

Finally, we examine the features selected by the MLP-Mixer models and demonstrate the benefits of employing non-permutation-invariant architectures for realtime jet tagging.

#### 5.1. Full Precision Model

The performance of the full-precision MLP-Mixer models (MLPM-fp) and JEDI-net [9] models is summarized in Table 2. The four MLP-Mixer models take 16, 32, 64, and 128 particles as input, while the two most accurate JEDI-net variants use 100 and 150 input particles. To represent the input configuration, we use  $N_p^{\#\text{fea}} = \#\text{ptl}$ , where #fea denotes the number of features per particle and #ptl represents the maximum number of particles used as input. Both MLP-Mixer and JEDI-net models utilize all 16 particle features, as described in Section 3.1. For MLP-Mixer models, the top-N particles are selected and ordered by  $p_T$  before being fed into the model. In contrast, the JEDI-net models also take as input the top-N particles by  $p_T$  but do not rely on their specific order, as they utilize a permutation-invariant architecture.

The performance metrics, including the area-under-curve (AUC) and the true positive rate (TPR) at false positive rate (FPR) of 1% and 10%, are presented in Table 2 for each class, with the best performance in each category highlighted in bold. The MLP-Mixer models consistently achieve higher AUC scores than the JEDI-net models across all classes. However, in terms of TPR at FPR=1% and FPR=10%, their performance varies slightly, with some classes showing marginal improvements while others experience minor reductions. Among all MLP-Mixer variants, the model with 64 input particles delivers the best overall performance across all categories.

The number of parameters and float-point operations (FLOP) for each model are listed in Table 3. For MLP-Mixer models, these values are computed after fusing batch normalization into dense layers. The table clearly shows that MLP-Mixer models require significantly fewer parameters and FLOPs compared to JEDI-net models, while still achieving comparable or superior performance.

Model	Input	Parameters	FLOPs
JEDI-net [9]	$N_p^{16} = 100$	$33,\!625$	$116M\S$
JEDI-net w/ $\Sigma O$ [9]	$N_p^{16} = 150$	8,767	458 M§
MLPM-fp	$N_p^{16} = 16$	2,465	21.6k
MLPM-fp	$N_p^{16} = 32$	3,265	50.5k
MLPM-fp	$N_p^{16} = 64$	6,401	133k
MLPM-fp	$N_p^{16} = 128$	$18,\!817$	396k

**Table 3.** The number of parameters and FLOPs for the MLP-Mixer and JEDI-net models. For MLP-Mixer models, values are computed after fusing batch normalization layers into dense layers.

§ The reported FLOPs may be overestimated, as certain matrix multiplications involving constant one-hot matrices are included in the FLOP calculations. However, these operations can be efficiently implemented as gather or lookup operations, significantly reducing computational complexity, as noted in [22].

16 features. The best performance in each category is highlighted in bold.								
N. L. L	JEDI-net [9]	JEDI-net	MLPM-fp	MLPM-fp	MLPM-fp	MLPM-fp		
Model		w/ $\Sigma O$ [9]						
	$N_p^{16} = 100$	$N_p^{16} = 150$	$N_p{}^{16} = 16$	$N_p^{16} = 32$	$N_p{}^{16} = 64$	$N_p^{16} = 128$		
	AUC (%)							
gluon	95.29	95.28	94.10	95.11	95.51	95.50		
light quarks	93.01	92.90	92.20	93.09	93.40	93.35		
W boson	97.39	96.95	96.09	97.51	97.84	97.83		
Z boson	96.79	96.49	95.38	97.11	97.50	97.47		
top quark	96.83	96.77	96.19	96.93	97.14	97.06		
	TPR at FPR=10%							
gluon	87.8	87.9	82.2	85.4	86.6	86.5		
light quarks	82.2	81.8	78.0	80.3	81.0	81.2		
W boson	93.8	92.7	89.5	93.1	93.9	93.9		
Z boson	91.0	90.3	86.5	91.2	92.2	92.1		
top quark	93.0	93.1	90.8	92.5	93.1	92.8		
	TPR at FPR=1%							
gluon	48.5	48.2	41.0	45.1	46.9	46.9		
light quarks	30.2	30.1	25.8	28.5	28.2	28.1		
W boson	70.4	65.8	51.8	70.4	75.4	75.3		
Z boson	76.9	72.9	66.5	77.7	80.9	80.7		
top quark	63.3	63.2	60.9	64.6	65.1	63.9		

**Table 2.** Performance comparison of the MLP-Mixer and JEDI-net models. The table reports the AUC and the TPR at FPR of 1% and 10%, for the different jet categories. The MLP-Mixer models are trained with 16, 32, 64, and 128 input particles, while the JEDI-net models use 100 and 150 input particles, where each particle is described by 16 features. The best performance in each category is highlighted in bold.

#### 5.2. Quantized Models

In this section we evaluate the performance, resource consumption, and latency of the quantized models optimized for deployment on FPGAs. For on-chip resource consumption, we consider the usage of Look-Up Tables (LUTs), digital signal processors (DSPs), Block RAMs (BRAMs), and flip-flops (FFs). All MLP-Mixer models presented in this work undergo HLS, logic synthesis, and place & route using Vitis HLS 2023.2 and Vivado 2023.2. Resource consumption is reported post place & route, ensuring accuracy in estimating the FPGA resource utilization. Timing convergence is verified and achieved for all models, unless explicitly stated otherwise. DA optimization using da4ml is applied for all models, and we use parallel I/O interfaces to the MLP models that were added to this study to ensure a fair comparison with simpler architectures.

The target FPGA is a Xilinx Virtex UltraScale+ XCVU9P (part number

 $<sup>\</sup>parallel {\rm Implemented \ using \ wire \ or \ register \ connections \ between \ layers, \ ensuring \ fully \ parallel \ data \ movement.}$ 

xcvu9p-f1ga2104-2L-e) operating at a 200 MHz clock frequency. To ensure bitaccurate consistency, we validate that the predictions from the quantized models obtained via the Python API precisely match those from the C-simulation of the HLS code across all samples in the test dataset. As a result, we do not differentiate between performance metrics obtained from Python and HLS simulations in subsequent evaluations. Latency and initiation interval¶ (II) are extracted from the HLS reports. Since the generated firmware does not include buffers or streaming interfaces, the reported latency and II can be considered accurate.



**Figure IV.** Accuracy vs. LUT usage and latency for the quantized MLP-Mixer and JEDI-net models, each trained with 16 features per particle. Dashed lines indicate the accuracy of the corresponding full-precision models.

Figures 4(a) and 4(b) illustrate the trade-off between accuracy, LUT usage, and latency for the quantized MLP-Mixer models, while Figure 5(a) and 5(b) present the same comparison for the quantized MLP models. Since our DA optimization involves explicit unrolling, both MLP-Mixer and MLP models operate without using any DSP slices. The dashed lines in the figures indicate the accuracy of the corresponding fullprecision models with the same number of input particles. For comparison, we also show the LUT usage, latency, and accuracy of the quantized JEDI-net models for  $N_p^1 6 = 30$ and 50 from Ref. [22].

The results demonstrate that MLP-Mixer models achieve a significantly better Pareto frontier between accuracy and resource consumption than JEDI-net models. All MLP-Mixer models operate within 100 ns latency, with the best-performing model achieving 81.64% accuracy at 70 ns latency. Additionally, they deliver a two-order-ofmagnitude higher throughput than JEDI-net models, while using fewer resources and achieving superior accuracy.

Although MLP models optimized with DA and HGQ achieve lower latency, their accuracy remains significantly lower than that of the MLP-Mixer models.

<sup>¶</sup> Number of blocking clock cycles required for processing each jet.



**Figure V.** Accuracy vs. LUT usage and latency for the quantized MLP and JEDI-net models, each trained with 16 features per particle. Dashed lines indicate the accuracy of the corresponding full-precision models. Models marked with "×" markers successfully underwent HDL synthesis but failed timing closure during the place & route phase.

Table 4 presents a detailed comparison of selected quantized MLP-Mixer, MLP, and the JEDI-net models. In the table,  $MLPM_{max}$  refers to the MLP-Mixer model with the highest accuracy for a given number of input particles, while  $MLPM_{alt}$  represents an alternative model with slightly lower accuracy but significantly reduced resource consumption. Similarly, the best performing MLP model trained in this study is labeled as  $MLP_{max}$ . Across all configurations, MLP-Mixer models consistently outperform the JEDI-net models in terms of accuracy, LUT usage, latency, and throughput, achieving superior efficiency with a significant margin.

While the models saturate at different accuracy levels, we observe that the tradeoff curves for the MLP-Mixer and MLP models with varying number of input particles remain mostly consistent – particularly in the accuracy region below 78% – with the exception of the 16-particles variants. This suggests that relying on parameter count or FLOPs as a proxy for resource consumption in quantized models can be misleading, as these metrics scale strongly with input size, as shown in Table 3. To provide further insights, Table 5 presents the bitwidths for weights and activations in the MLP-Mixer models, as defined in [11]. The data clearly shows that as the number of input particles increases, both the bitwidths and (1–sparsity) of the weights and activations decrease, effectively compensating for the otherwise higher resource consumption.

Finally, the Receiver Operating Characteristic (RoC) curves for the full-precision and best performing quantized MLP-Mixer models with 16, 32, 64, and 128 input particles are shown in Figure VI. In all cases, the quantized models achieve comparable, but slightly degraded performance compared to the full-precision models.

**Table 4.** Comparison of quantized MLP-Mixer, MLP, and the JEDI-net models in terms of accuracy, resource consumption, and latency.  $MLPM_{max}$  refers to the MLP-Mixer model with the highest accuracy for a given number of input particles, while  $MLPM_{alt}$  represents an alternative model with slightly lower accuracy but significantly reduced resource usage. Similarly,  $MLP_{max}$  denotes the best-performing MLP model trained in this study. All models are trained on the hls4ml jet tagging dataset with 16 features for input particle.

Model	$N_{p}^{16}$	Accuracy (%)	Latency	DSP	LUT (k)	FF(k)	BRAM	II
JEDI-net U4 $[22]$	50	80.9	$130 \ (650 \ ns)$	8,945	855	201	25	110
JEDI-net U5 $\left[ 22\right]$	50	81.2	$181 \ (905 \ ns)$	8,986	815	189	37	150
JEDI-net J4 $[22]$	32	78.4	58 (290  ns)	8,776	865	138	37	30
JEDI-net J5 [22]	32	79.9	$181 \ (905 \ ns)$	$9,\!833$	911	158	37	150
MLPM <sub>max</sub>	16	77.5	15 (75  ns)	0	102	24	0	1
$\mathrm{MLPM}_{\mathrm{max}}$	32	80.7	$14 \ (70 \ {\rm ns})$	0	87	22	0	1
$\mathrm{MLPM}_{\mathrm{alt}}$	32	79.9	$13 \ (65 \ {\rm ns})$	0	42	11	0	1
$\mathrm{MLPM}_{\mathrm{max}}$	64	81.6	$14 \ (70 \ {\rm ns})$	0	126	32	0	1
$\mathrm{MLPM}_{\mathrm{alt}}$	64	81.3	$13 \ (65 \ {\rm ns})$	0	58	16	0	1
$\mathrm{MLPM}_{\mathrm{max}}$	128	81.3	$16 \; (80 \; \mathrm{ns})$	0	151	42	0	1
MLP <sub>max</sub>	16	77.3	6 (30  ns)	0	73	18	0	1
$\mathrm{MLP}_{\mathrm{max}}$	32	79.7	$7~(35~{\rm ns})$	0	103	23	0	1
$\mathrm{MLP}_{\mathrm{max}}$	64	79.8	6 (30  ns)	0	57	15	0	1
$\mathrm{MLP}_{\mathrm{max}}$	128	80.0	$7~(35~\mathrm{ns})$	0	113	29	0	1

**Table 5.** Average bitwidths (Avg. BW), average bitwidths of non-zero elements (Avg. BW (non-zero)), and sparsity for weights and activations of the  $MLPM_{max}$  models. Activation sparsity refers to the static sparsity after quantization, which is independent of input data during inference.

	$N_p^{-16}$	Weights			Activations			
Model		Avg. BW	Avg. BW	Sparsity	Avg. BW	Avg. BW	Sparsity	
			(non-zero)			(non-zero)		
$\mathrm{MLPM}_{\mathrm{max}}$	16	1.93	2.52	0.23	4.80	6.34	0.24	
$\mathrm{MLPM}_{\mathrm{max}}$	32	1.55	2.51	0.38	3.27	5.66	0.42	
$\mathrm{MLPM}_{\mathrm{alt}}$	32	0.60	2.27	0.73	1.78	4.82	0.63	
$\mathrm{MLPM}_{\mathrm{max}}$	64	1.06	2.44	0.57	2.55	5.19	0.51	
$\mathrm{MLPM}_{\mathrm{alt}}$	64	0.58	2.29	0.75	1.67	4.81	0.65	
$\mathrm{MLPM}_{\mathrm{max}}$	128	0.57	2.27	0.75	1.78	4.76	0.62	

#### 5.3. Feature Importance

In an HGQ-trained model, the bitwidth of activations can serve as a proxy for feature importance. When quantizing a value drawn from a smooth distribution, the quantization-induced change can be interpreted as additive noise affecting the original value. Since the integer bitwidth is adjusted to accommodate the maximum activation



(a) RoC for MLP-Mixers with  $N_p^{16} = 16$  (b) RoC for MLP-Mixers with  $N_p^{16} = 32$ 



(c) RoC for MLP-Mixers with  $N_p^{16} = 64$  (d) RoC for MLP-Mixers with  $N_p^{16} = 128$ 

Figure VI. Receiver Operating Characteristic (RoC) curves of the MLPM<sub>max</sub> models (solid lines) compared to their full precision counterpart (dashed lines) trained with 16, 32, 64, and 128 input particles. The Area Under the Curve (AUC) for each jet category is provided in the legend, formatted as "full-precision / quantized".

value, the quantization noise for a given value x is approximately proportional to  $\max x \cdot 2^{-\text{bitwidth}}$ . This suggests that bitwidth allocation reflects the model's tolerance for additive noise on a given feature when trading-off with resource usage: features assigned higher bitwidths are considered more critical, as the model minimizes quantization errors for them. However, this interpretation is less reliable when input values originate from highly discrete or narrowly distributed distributions compared to the magnitude of the quantization noises.

Figure VII presents the average input bitwidths for the four best-performing quantized MLP-Mixer models. These are obtained by adding and dividing by two the bitwidths at two key locations: (i) at the input to MLP1 and (ii) before the addition



Figure VII. Average input bitwidths for the best-performing quantized MLP-Mixer models with 16, 32, 64, and 128 input particles. The reported values are computed by summing the bitwidths at the input to MLP1 and one leg of the addition operation before MLP3, then dividing by two.

operation in MLP3 (refer to Figure II). From Figure VII, we observe that the models allocate more bits to particles carrying higher energy, particularly for the features E,  $E_{\rm rel}$ ,  $p_T$ , and  $p_{T\rm rel}$ , which are assigned higher bitwidths for the first ~ 40 particles. In contrast,  $\eta_{\rm rot}$ ,  $\phi_{\rm rot}$ ,  $\Delta R$  receive higher bitwidths across all particles, indicating their persistent importance in jet tagging. This variation in bitwidth allocation suggests that MLP-Mixer models actively discard less relevant features to optimize resource usage. Furthermore, it reinforces the advantage of breaking permutation invariance in jet tagging tasks. If the model were forced to maintain permutation invariance, bitwidths would need to remain constant across particle indices for each feature, leading to substantial resource overhead. This highlights the efficiency gains achieved by allowing the model to prioritize key features dynamically.

#### 5.4. Quantized model with $p_T$ , $\eta$ and $\phi$ inputs

In this section, we evaluate the performance of the quantized MLP-Mixer models trained using only  $p_T$ ,  $\eta$ , and  $\phi$  for each particle, rather than the full 16-feature set described in Table 1. This configuration reflects a more realistic experimental scenario as the derived kinematic features beyond momentum may not always be available in the hardware trigger systems. In addition to restricting input features, we apply a  $p_T \geq 2$  GeV selection to each particle, as suggested in Ref. [23], to better approximate detector behavior. Apart from adjusting the layer sizes to accommodate the reduced feature set (as shown in Figure II with n = 3), all other hyperparameters remain unchanged. Training and evaluation procedures follow the same methodology used for models trained



(a) Accuracy vs. LUT usage

(b) Accuracy vs. Latency

Figure VIII. Accuracy vs. LUT usage (left) and latency (right) of the quantized MLP-Mixer, Interaction Network (IN) [23], and Deep Sets (DS) [23] models trained using only  $p_T$ ,  $\eta$ , and  $\phi$  as input features per particle. Dashed lines indicate the accuracy of the corresponding full-precision models with the same number of input particles. Models marked with a "×" successfully underwent HDL synthesis but failed timing closure during the place & route phase. The error bars for the IN and DS models reflect ambiguity in the reported accuracy in the original work [23].



Figure IX. Accuracy vs. LUT usage (left) and latency (right) of the quantized MLP, Interaction Network (IN) [23], and Deep Sets (DS) [23] models trained with using only  $p_T$ ,  $\eta$ , and  $\phi$  as input features per particle. Dashed lines indicate the accuracy of the corresponding full-precision models with the same number of input particles. Models marked with a "×" successfully underwent HDL synthesis but failed timing closure during the place & route phase. The error bars for the IN and DS models reflect ambiguity in the reported accuracy in the original work [23].

with 16 features per particle, as detailed in Sections 5.1 and 5.2.

The performance of the quantized MLP-Mixer models trained with particle features

**Table 6.** Performance, latency, and resource consumption of the quantized MLP-Mixer, MLP, DS and IN models described by the accuracy and resource consumption. All models are trained on the hls4ml jet tagging dataset using only  $p_T$ ,  $\eta$ , and  $\phi$  as input features per particle. Only particles with  $p_T \geq 2$  GeV are included in the training and evaluation. The intervals for accuracies for the DS and IN models are due to the ambiguity in the original work [23], where the author did not report model accuracies after quantization and firmware generation.

Model	$N_p^{3}$	Accuracy (%)	Latency	DSP	LUT (k)	FF (k)	BRAM	II
DS [23]	16	62.5 - 69.4	23 (115  ns)	555	747	239	0	3
IN [23]	16	63.7 - 70.8	$36 \ (180 \ ns)$	5,362	1,388	594	0	3
DS [23]	32	68.3 - 75.9	$26 \ (130 \ ns)$	434	903	359	0	2
IN [23]	32	68.2 - 75.8	$41 \ (205 \ ns)$	$2,\!120$	1,162	761	0	3
MLPM <sub>max</sub>	16	71.7	$14 \ (70 \ ns)$	0	75	17	0	1
$\mathrm{MLPM}_{\mathrm{alt}}$	16	70.8	$10 \ (50 \ {\rm ns})$	0	40	8.4	0	1
$\mathrm{MLPM}_{\mathrm{max}}$	32	78.0	$13 \ (65 \ ns)$	0	63	15	0	1
$\mathrm{MLPM}_{\mathrm{alt}}$	32	76.1	$11 \ (55 \ ns)$	0	25	7.1	0	1
$\mathrm{MLPM}_{\mathrm{max}}$	64	79.7	$15 \ (75 \ ns)$	0	159	36	0	1
$\mathrm{MLPM}_{\mathrm{alt}}$	64	75.9	$13 \ (65 \ ns)$	0	26	7.9	0	1
$\mathrm{MLPM}_{\mathrm{max}}$	128	79.8	$15 \ (75 \ ns)$	0	83	21	0	1
MLP <sub>max</sub>	16	71.1	7 (35  ns)	0	69	14	0	1
$\mathrm{MLP}_{\mathrm{max}}$	32	75.6	$7~(35~\mathrm{ns})$	0	70	13	0	1
$\mathrm{MLP}_{\mathrm{max}}$	64	75.5	$7~(35~\mathrm{ns})$	0	43	8.7	0	1
$\mathrm{MLP}_{\mathrm{max}}$	128	76.7	$7~(35~{\rm ns})$	0	101	19	0	1

 $p_T$ ,  $\eta$ , and  $\phi$  is shown in Figures 8(a) and 8(b). The corresponding results for the quantized MLP models we trained are shown in Figures 9(a) and 9(b). These models are compared with the 8-bits quantized Interaction Network (IN) and Deep Sets (DS) models proposed in Ref. [23] where they were trained with the same inputs. For the IN and DS models, the error bars on the accuracy reflect ambiguities in the reported performance in the original work, where only full-precision accuracy was provided. The accuracy of their quantized counterparts was estimated to be at least 90% of their full-precision versions. As shown in the figures, the quantized MLP-Mixer models consistently outperform the IN and DS models in both accuracy and resource consumption.

A detailed comparison of selected MLP-Mixer, MLP, IN, and DS models is provided in Table 6. The notation follows Table 4. We exclude MLP-Mixer and MLP models that failed timing closure in the place & route phase from the comparison. Depending on the chosen MLP-Mixer model, we demonstrate that MLP-Mixers can outperform IN and DS models by orders of magnitude in resource efficiency while maintaining comparable or superior accuracy. For instance, the MLPM<sub>alt</sub> ( $N_p^3 = 32$ ) model achieves the following improvements over the best IN and DS models for each metric:

- 0.2% higher accuracy than the full precision models
- 97% lower LUT usage





(a) RoC for MLP-Mixers with  $N_p^3 = 16$ 

(b) RoC for MLP-Mixers with  $N_p^3 = 32$ 



(c) RoC for MLP-Mixers with  $N_p^3 = 64$  (d) RoC for MLP-Mixers with  $N_p^3 = 128$ 

**Figure X.** Receiver Operating Characteristic (RoC) curves of the MLPM<sub>max</sub> models (solid lines) compared to their full precision counterpart (dashed lines) trained with 16, 32, 64, and 128 input particles and using only  $p_T$ ,  $\eta$ , and  $\phi$  as input features per particle. The Area Under the Curve (AUC) for each jet category is provided in the legend, formatted as "full-precision / quantized".

- 97% lower FF usage
- 100% lower DSP usage
- 52% lower latency
- 100% higher throughput

Similar to models trained with 16 particles as input, the MLP models optimized with DA and HGQ achieve lower latency but are limited in accuracy compared to MLP-Mixer models. Nevertheless, our results demonstrate that when properly optimized, even MLP models can surpass permutation-invariant architectures (IN and DS models) in both resource efficiency and accuracy, as shown in Figure 9(a) and 9(b).

Finally, Figure X presents the ROC curves for the full-precision and best-performing quantized MLP-Mixer models trained with particle features  $p_T$ ,  $\eta$ , and  $\phi$ . The quantized models achieve comparable, albeit slightly lower, performance relative to their full-precision counterparts.

#### 6. Conclusions

In this work, we demonstrated that MLP-Mixer is a highly effective architecture for jet classification tasks using particle-level feature inputs. Our results show that MLP-Mixer models outperform JEDI-net models in accuracy, while requiring significantly fewer FLOPs, leading to faster inference times on CPUs and GPUs. Furthermore, we demonstrated that by leveraging HGQ and da4m1, MLP-Mixer models can be efficiently compressed and deployed on FPGAs. These models achieve a superior Pareto frontier between accuracy and resource consumption compared to previous models, while also delivering the highest accuracy, throughput, and lowest latency among the models evaluated.

Although MLP-Mixer models are not permutation invariant, we showed that they can dynamically discard less relevant features to optimize resource usage. Additionally, our analysis revealed that the models selectively allocate more bits to particles with higher  $p_T$ , suggesting that enforcing permutation invariance could introduce unnecessary resource overhead. Based on these findings, we argue that breaking permutation invariance can be beneficial for the jet classification task, particularly when the input data is known to be sorted by a physically meaningful quantity.

#### Author information

Correspondence and material requests should be emailed to C. Sun (chsun@cern.ch).

#### Data availability statement

The data and software required to reproduce this work can be found at https://doi.org/10.5281/zenodo.3602260 [28] and https://github.com/calad0i/HGQ-demos/tree/master/jet\_classifier\_large.

#### Acknowledgements

C.S. acknowledges partially supported by the NSF ACCESS Grant number PHY240298. C.S. and M.S. acknowledge partial support from the U.S. Department of Energy (DOE), Office of Science, Office of High Energy Physics grant DE-SC0011925. J.N., M.S., and C.S. are partially supported by the U.S. Department of Energy (DOE), Office of Science, Office of High Energy Physics "Designing efficient edge AI with physics phenomena" Project (DE-FOA-0002705). J.N. is partially supported by the AI2050 program at Schmidt Futures (Grant G-23-64934).

#### References

- The LHC Study Group 1995 The Large Hadron Collider, Conceptual Design Tech. rep. CERN/AC/95-05 (LHC) Geneva
- The CMS Collaboration (CMS) 2020 The Phase-2 Upgrade of the CMS Level-1 Trigger Tech. rep. CERN Geneva final version URL https://cds.cern.ch/record/2714892
- [3] The ATLAS Collaboration (ATLAS) 2017 Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System Tech. rep. CERN Geneva URL https://cds.cern.ch/record/2285584
- [4] Zurbano Fernandez I et al. 2020 CERN Yellow Reports: Monographs 10/2020
- [5] Qu H and Gouskos L 2020 Physical Review D 101 ISSN 2470-0029 URL http://dx.doi.org/10. 1103/PhysRevD.101.056019
- [6] Qu H, Li C and Qian S 2024 Particle transformer for jet tagging (Preprint 2202.03772) URL https://arxiv.org/abs/2202.03772
- Bogatskiy A, Hoffman T, Miller D W, Offermann J T and Liu X 2024 Journal of High Energy Physics 2024 113 ISSN 1029-8479 URL https://doi.org/10.1007/JHEP03(2024)113
- [8] Spinner J, Bresó V, de Haan P, Plehn T, Thaler J and Brehmer J 2024 Lorentz-equivariant geometric algebra transformers for high-energy physics (*Preprint* 2405.14806) URL https: //arxiv.org/abs/2405.14806
- [9] Moreno E A, Cerri O, Duarte J M, Newman H B, Nguyen T Q, Periwal A, Pierini M, Serikova A, Spiropulu M and Vlimant J R 2020 The European Physical Journal C 80 ISSN 1434-6052 URL http://dx.doi.org/10.1140/epjc/s10052-020-7608-4
- [10] Tolstikhin I, Houlsby N, Kolesnikov A, Beyer L, Zhai X, Unterthiner T, Yung J, Steiner A, Keysers D, Uszkoreit J, Lucic M and Dosovitskiy A 2021 Mlp-mixer: An all-mlp architecture for vision (*Preprint* 2105.01601) URL https://arxiv.org/abs/2105.01601
- [11] Chang S, Årrestad T, Lončar V, Ngadiuba J and Spiropulu M 2024 Gradient-based automatic per-weight mixed precision quantization for neural networks on-chip URL https://authors. library.caltech.edu/doi/10.7907/hq8jd-rhg30
- [12] Fahim F, Hawks B, Herwig C, Hirschauer J, Jindariani S, Tran N, Carloni L P, Guglielmo G D, Harris P C, Krupa J D, Rankin D, Valentin M B, Hester J, Luo Y, Mamish J, Orgrenci-Memik S, Aarrestad T, Javed H, Loncar V, Pierini M, Pol A A, Summers S, Duarte J M, Hauck S, Hsu S, Ngadiuba J, Liu M, Hoang D, Kreinar E and Wu Z 2021 CoRR abs/2103.05579 (Preprint 2103.05579) URL https://arxiv.org/abs/2103.05579
- [13] The Advanced Micro Devices, Inc 2023 Vitis high-level synthesis user guide (ug1399) URL https://docs.amd.com/r/2023.2-English/ug1399-vitis-hls
- [14] Komiske P T, Metodiev E M and Thaler J 2019 Journal of High Energy Physics 2019 ISSN 1029-8479 URL http://dx.doi.org/10.1007/JHEP01(2019)121
- [15] Sun C, Nakajima T, Mitsumori Y, Horii Y and Tomoto M 2023 Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 1045 167546 ISSN 0168-9002 URL http://dx.doi.org/10.1016/j.nima.2022.167546
- [16] Coelho C N, Kuusela A, Li S, Zhuang H, Ngadiuba J, Aarrestad T K, Loncar V, Pierini M, Pol A A and Summers S 2021 Nature Machine Intelligence 3 675–686 URL https://doi.org/10. 1038%2Fs42256-021-00356-5
- [17] Ngadiuba J, Loncar V, Pierini M, Summers S, Guglielmo G D, Duarte J, Harris P, Rankin D, Jindariani S, Liu M, Pedro K, Tran N, Kreinar E, Sagear S, Wu Z and Hoang D 2020 Machine Learning: Science and Technology 2 015001 URL https://dx.doi.org/10.1088/2632-2153/ aba042
- [18] Ioffe S and Szegedy C 2015 CoRR abs/1502.03167 (Preprint 1502.03167) URL http://arxiv. org/abs/1502.03167
- [19] Lou Q, Guo F, Kim M, Liu L and Jiang L 2020 Autoq: Automated kernel-wise neural network quantization International Conference on Learning Representations URL https:// openreview.net/forum?id=rygfnn4twS

- [20] Battaglia P W, Pascanu R, Lai M, Rezende D and Kavukcuoglu K 2016 Interaction networks for learning about objects, relations and physics (*Preprint* 1612.00222) URL https://arxiv.org/ abs/1612.00222
- [21] Zaheer M, Kottur S, Ravanbakhsh S, Poczos B, Salakhutdinov R and Smola A 2018 Deep sets (*Preprint* 1703.06114) URL https://arxiv.org/abs/1703.06114
- [22] Que Z, Fan H, Loo M, Li H, Blott M, Pierini M, Tapper A and Luk W 2024 ACM Transactions on Embedded Computing Systems 23 1–28 ISSN 1558-3465 URL http://dx.doi.org/10.1145/ 3640464
- [23] Odagiu P, Que Z, Duarte J, Haller J, Kasieczka G, Lobanov A, Loncar V, Luk W, Ngadiuba J, Pierini M, Rincke P, Seksaria A, Summers S, Sznajder A, Tapper A and Årrestad T K 2024 Machine Learning: Science and Technology 5 035017 ISSN 2632-2153 URL http://dx.doi. org/10.1088/2632-2153/ad5f10
- [24] Govorkova E, Puljak E, Aarrestad T, James T, Loncar V, Pierini M, Pol A A, Ghielmetti N, Graczyk M, Summers S, Ngadiuba J, Nguyen T Q, Duarte J and Wu Z 2021 Autoencoders on fpgas for real-time, unsupervised new physics detection at 40 mhz at the large hadron collider URL https://arxiv.org/abs/2108.03986
- [25] Bhattacherjee B, Konar P, Ngairangbam V S and Solanki P 2023 Llpnet: Graph autoencoder for triggering light long-lived particles at hl-lhc (*Preprint* 2308.13611) URL https://arxiv.org/ abs/2308.13611
- [26] The CMS Collaboration (CMS) 2024 2024 Data Collected with AXOL1TL Anomaly Detection at the CMS Level-1 Trigger URL https://cds.cern.ch/record/2904695
- [27] The CMS Collaboration (CMS) 2023 Level-1 Trigger Calorimeter Image Convolutional Anomaly Detection Algorithm URL https://cds.cern.ch/record/2879816
- [28] Pierini M, Duarte J M, Tran N and Freytsis M 2020 Hls4ml lhc jet dataset (150 particles) URL https://doi.org/10.5281/zenodo.3602260
- [29] Pearkes J, Fedorko W, Lister A and Gay C 2017 Jet constituents for deep neural network based top quark tagging (*Preprint* 1704.02124) URL https://arxiv.org/abs/1704.02124
- [30] Summers, Sioni, Bestintzanos, Ioannis and Petrucciani, Giovanni 2024 EPJ Web of Conf. 295 02024 URL https://doi.org/10.1051/epjconf/202429502024
- [31] The CMS Collaboration (CMS) 2022 Neural network-based algorithm for the identification of bottom quarks in the CMS Phase-2 Level-1 trigger URL https://cds.cern.ch/record/ 2814728?ln=en
- [32] Alessandro, Franco G, nickfraser, Umuroglu Y and vfdev 2021 Xilinx/brevitas: Release version 0.2.1 URL https://doi.org/10.5281/zenodo.4507794
- [33] Aksoy L, Costa E, Flores P and Monteiro J 2012 Multiplierless Design of Linear DSP Transforms pp 73–93
- [34] Hosangadi A, Fallah F and Kastner R 2005 Reducing hardware complexity of linear dsp systems by iteratively eliminating two-term common subexpressions *Proceedings of the ASP-DAC 2005*. Asia and South Pacific Design Automation Conference, 2005. vol 1 pp 523–528 Vol. 1
- [35] Nguyen H and Chattejee A 2000 IEEE Transactions on Very Large Scale Integration (VLSI) Systems 8 419–424
- [36] Voronenko Y and Püschel M 2007 ACM Trans. Algorithms 3 11-es ISSN 1549-6325 URL https://doi.org/10.1145/1240233.1240234
- [37] Mirzaei S, Hosangadi A and Kastner R 2006 Fpga implementation of high speed fir filters using add and shift method 2006 International Conference on Computer Design pp 308–313
- [38] Chang S 2024 da4ml 20241120 da4ml is an deploy-time optimization tool with distributed arithmetic for machine learning models. URL https://doi.org/10.5281/zenodo.14194660
- [39] Tange O 2023 Gnu parallel 20240122 ('frederik x') GNU Parallel is a general parallelizer to run multiple serial command line programs in parallel without changing them. URL https: //doi.org/10.5281/zenodo.10558745