


The Compact Muon Solenoid Experiment

CMS Note

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland



28 August 2023 (v2, 12 December 2023)

Towards Real-Time Machine Learning Based Signal/Background Selection in the CMS Detector Using Quantized Neural Networks and Input Data Reduction

Arijana Burazin Misura, Josip Music, Marina Prvan, Damir Lelas

Abstract

To boost its discovery potential, the Large Hadron Collider (LHC) is being prepared for an extensive upgrade. The new phase, High Luminosity LHC (HL-LHC), will operate at luminosity (number proportional to the rate of collisions) increased by a factor of five. Such an increase in luminosity consequently will result in enormous amounts of generated data, the vast majority of which is uninteresting data or pile up (PU). HL-LHC detectors, including Compact Muon Solenoid (CMS), will thus have to rely on innovative technologies and methods to select, collect and analyze collisions data. In charge of data reduction at the early stages of data streaming is a Level 1 Trigger (L1T), the real-time event selection system based on information from calorimeters and muon detectors, with a decision time of around 12 microseconds. For the L1T method, we propose quantized neural network models deployed in targeted L1T devices, Field Programmable Gate Arrays (FPGAs), as a classifier between electromagnetic and pile-up/QCD showers. Traditional classifiers are based on cluster shapes and hand-crafted features, while the proposed quantized neural network uses raw detector data, that speeds up the classification process. Data reduction using selection and quantization additionally decreases model size retaining accuracy. The model execution requires less than 1 microsecond, making it a possible mechanism for real-time signal/background classification.

Towards Real-Time Machine Learning Based Signal/Background Selection in the CMS Detector Using Quantized Neural Networks and Input Data Reduction

Arijana Burazin Mišura, Josip Musić, Marina Prvan, and Damir Lelas

Faculty of Electrical Engineering, Mechanical Engineering and Naval Architecture,
University of Split

December 12, 2023

Abstract

The Large Hadron Collider (LHC) is being prepared for an extensive upgrade to boost its discovery potential. The new phase, High Luminosity LHC, will operate at luminosity (number proportional to the rate of collisions) increased by a factor of five. Consequently, such an increase in luminosity will result in enormous amounts of generated data, the vast majority of which is uninteresting data or pile up. High Luminosity LHC detectors, including Compact Muon Solenoid, will thus have to rely on innovative technologies and methods to select, collect, and analyze collision data. In charge of data reduction at the early stages of data streaming is a level 1 trigger (L1T), the real-time event selection system based on information from calorimeters and muon detectors, with a decision time of around 12 microseconds. For the L1T method, we propose quantized neural network models deployed in targeted L1T devices, field programmable gate arrays, as a classifier between electromagnetic and pile-up/quantum chromodynamics showers. Traditional classifiers are based on cluster shapes and hand-crafted features, while the planned quantized neural network uses raw detector data that speed up the classification process. Proposed data reduction methods additionally decrease the model size, retaining accuracy. The model execution requires less than 1 microsecond, making it a possible real-time signal/background classification mechanism.

Keywords: CMS; Level 1 Trigger; HGCAL; Quantized neural network; hls4ml; EM shower; Classification

1 Introduction

Exploration of the essential nature of space and discoveries of exotic particles would not be possible without particle accelerators, which are primary tools in high-energy physics (HEP). The world's largest particle accelerator, the Large Hadron Collider (LHC) [1], is built to understand the fundamental laws of nature and test different predictions of elementary particle physics. In 2012., the LHC experiments A Toroidal LHC Apparatus and Compact Muon Solenoid (CMS) [2] confirmed the existence of the Higgs boson through the discovery of the predicted fundamental particle. The current phase aims to fully characterize the Higgs boson properties and search for phenomena beyond the standard model.

In the LHC ring, high-energy proton collisions occur every 25 ns. To analyze the results of the collisions, four detectors surround collision points and record the obtained data. One of four LHC detectors is CMS, a cylindrically shaped general-purpose detector with several concentric layers of components (Figure 1). When interesting particles are produced as a result of the collision, they immediately decay into lighter, stable particles. As shown in Figure 1, decay products pass through detector layers and interact with them, allowing their direction and momentum to be measured. Among the detectable decay products, there are electrically charged leptons (electrons and muons) and particle jets (collimated streams of particles originating from quarks and gluons)[3].

This paper concentrates on the electromagnetic calorimeter sub-detector part, where particles such as electrons and photons are detected. Combining the data coming from different sub-detectors, an ultimate goal is to obtain a complete picture of the collision event and to enable the application of a wide range of physical analyses. However, the purpose of this paper is not the physics analysis but the input data reduction, as well as the implementation of the real-time model for signal versus background selection. Namely, the crucial part of the CMS dataflow is the event data selection in the early stage. The detector produces tens of terabytes of data per second, exceeding the available processing and storage resources. An event processing system – level 1 trigger (L1T) is running, deciding which events should be kept for further analysis and which should be discarded. The system is real-time, requiring a fast decision on the input data stream.




Figure 1: CMS detector overview [4]: sub-detector layers are designed to measure particles produced in proton-proton collisions (Layers: 1. Silicon Tracker, 2. Electromagnetic Calorimeter, 3. Hadron Calorimeter, 4. Superconducting Solenoid, 5. Iron return yoke interspersed with Muon chambers);

The decision will be even more demanding in the High Luminosity (HL) LHC phase because of the increased luminosity, the unprecedented radiation dose, and the vast majority of uninteresting additional proton-proton interactions referred to as pile-up (PU). The CMS is going through a major upgrade for the HL-LHC era, and one of the most important enhancements is the design of the High Granularity Calorimeter (HGCal) [5], the radiation-hard replacement of endcaps. It is a sampling calorimeter characterized by very high granularity using 6 million silicon and 400 thousand scintillator channels. Silicon cells are fine-segmented, radiation tolerant, and fast sufficient to mitigate PU. Hence, they will be used as active elements in high-radiation areas, and, to reduce cost, scintillator will be active element in lower-radiation areas (Figure 2). The HGCal consists of an electromagnetic part (CE-E) containing 14 layers of hexagonal silicon sensors, while the hadronic part (CE-H) consists of another 8 silicon layers.

At the moment when the experiment presented in the paper started, there were in total 50 layers foreseen in HGCal [6], as it is used in the experiment. In the recent design revision, the number of layers is reduced to 47 [7] but could also change till actual production.

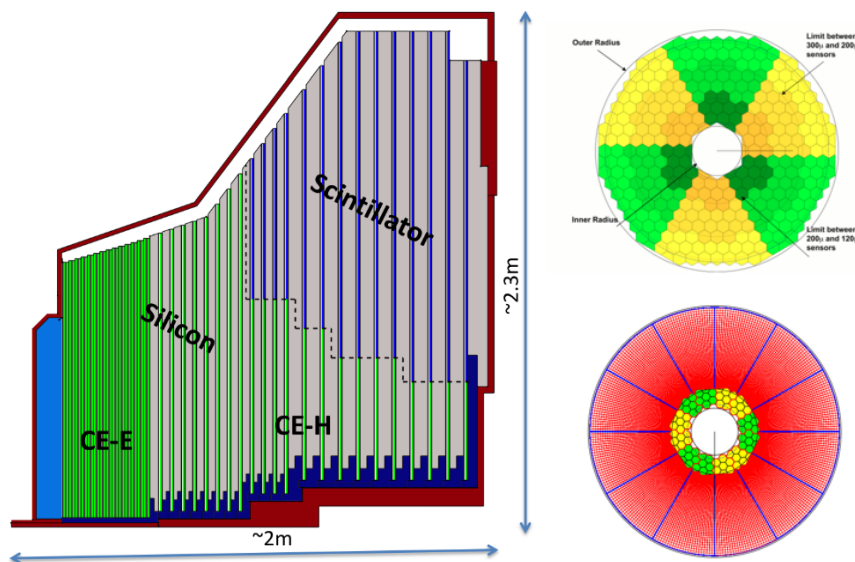


Figure 2: The HGCal endcap mechanical construction (left) and a single electromagnetic and hadronic calorimeter layer (right), adjusted from [4], [5].

An expected increase in the data rate for the HL-LHC requires advanced approaches that would enable efficient real-time analysis methods for the HGCal L1 trigger. Artificial intelligence, especially machine learning (ML), has become a standard tool in analyzing big data and is gradually becoming a common tool in HEP. The development of sophisticated ML algorithms, specifically convolutional neural networks (CNN), has been shown as a powerful tool in image recognition, solving many computer vision problems. Until recently, the main drawback of real-time implementation was the fact that running capable ML methods

requires time and powerful hardware. Advanced approaches have enabled the integration of ML methods in real-time processing tasks [8, 9, 10]. However, in restricted conditions under which the collection and primary selection of data take place inside the LHC detectors, the standard hardware and operating systems used for typical ML applications are not possible to use. Because of radiation exposure, restricted budget, the possibility of material damage, and the enormous data rate, a choice has been made on field-programmable gate arrays (FPGAs) to be used as a platform for L1 trigger algorithms. Compared to central processing units or graphic processing units, they are very efficient because of very high throughput, low latency, and low power budget. Recently, scientists began to explore using ML techniques, specifically neural networks (NN), on FPGAs to improve real-time event processing. This paper deals with the same topic and presents a case study for real-time signal/background classification of the HGCAL data using quantized neural networks deployed on FPGAs. The reason behind network quantization is an effort to reduce the network requirements since limited resources are available on the target FPGA, shared between multiple algorithms. Moreover, input data reduction is applied using selection and quantization procedures to additionally decrease the model size and reduce latency and memory requirements.

The paper is organized as follows. Section 2 describes the application of NN in HEP, particularly for triggering purposes in particle detectors. Section 3 presents the methodology of the conducted study: the process of the images data set is described, together with details of NN model optimization. The section ends with the hls4ml tool introduction. Section 4 illustrates the performance of the described models via several data refinement approaches and comparisons with baseline particle classification methods. The results of the implementation of selected models in FPGA using hls4ml are discussed at the end of the section. Section 5 concludes the paper, followed by the references used.

2 Previous Work

Because of their efficiency in processing large amounts of data, ML techniques have been used frequently in HEP data analysis. High computational requirements and execution time limited their use mainly to offline analysis. In particular, ML techniques such as NNs are applied at LHC for different purposes: event simulation and reconstruction [11], [12], event classification [13], [14], anomaly detection [15], [16], and in monitoring the supply current to ensure stable operation of the detector [17]. In doing so, different types of NNs are used. Fully connected neural networks in [18] search for new long-lived particles that decay into jets. Qasim et al. [19], using a graph neural network named GravNet, perform grouping, classification, and regression of energy and position. Graph neural networks have also been used to track charged particles [20]. Generative Adversarial Networks are often used as a substitute for computationally intensive parts of Monte Carlo simulations, such as the modeling of electromagnetic showers [21] and the reconstruction of jet images, which was shown by the authors in [22]. The group of authors in [23] based on the raw data obtained from the CMS calorimeter reconstructs, discusses, and simulates particles (electrons, photons, charged pions, and neutral pions) using various machine learning methods. For HGCAL sensors to function properly, in the [24], authors suggest a deep-learning-based pre-selection algorithm that fully automates visual inspection to ensure that sensors satisfy quality control criteria. Otherwise, defects and dust on a sensor surface can lead to sensor failures. ML is also used in other parts of the CMS experiment, in the Drift Tubes detector, for generating trigger primitives [25]. Guest et al. [26] provide a systematic survey on the application of deep learning in LHC physics.

The initial use of ML methods was mainly based on manual tuning/finding high-level features, requiring more detailed knowledge of particle decay phenomenology. Before the development of deep learning, methods based on high-level features (which required preprocessing of the data) achieved better results than those using raw data. Baldi in [3] found that in some classification problems in HEP, shallow NNs using low-level data achieve almost the same performance as those using tuned features. It is shown that deep learning techniques based on low-level data can discover the insight contained in high-level features. Cogan et al. [27] recognized that the projection of the calorimeter structure, which is present in almost all detectors used in HEP, is similar to image pixels. This way of data representation allowed physicists to use new tools in image processing, such as CNNs. Andrews et al. [13] were among the first to present an approach based on the classification of images obtained on simulated CMS detector data. The paper studies the decay of the Standard Model Higgs boson to two photons using the 2012 CMS Open Data, simulated using the Geant4 simulator. The images were generated using data on the deposited energy of the particles without any processing related to the type of particles. Andrews later applied the same technique in the classification of various types of particles: the classification of quarks and gluons [28] and boosted top quarks [29]. All research is done on CMS Open data, collision, and simulated data that was recorded in older experiments and is now publicly available to enable the most successful collaboration of CMS collaboration with the ML community. Today, CNNs are widely used in HEP because of their efficiency in image processing and pattern recognition [29][30][31].

The idea of using NNs to increase trigger efficiency was introduced for the first time in 1990. in [32]. Four hundred fifty samples representing the 8x8 area around the cell in the calorimeter with the largest energy deposit were used to train the feed-forward network in the Collider Detector experiment at Fermilab. There was a huge gap between training and testing accuracy, and the plan was to implement NN in hardware. At that time, large-scale NNs implemented in silicon began to appear, and with latency times on the order of 1 ms, they were suitable for the trigger systems of the time. Since then, NNs have been continuously used in triggers, but high resource requirements (including energy constraints) did not allow their real-time application. In high-level CMS triggers, with no such resource constraints, NNs are used for different purposes, for example, for track seed filtering [33] or to label jets [34]. In CMS L1T, because of the limited size, required latency, and radiation exposure, the traditional approach to NNs is not acceptable. For this reason, work has begun on new approaches that enable NN use in limited resource environments.

The reduction of memory requirements while maintaining the NN accuracy was attempted to be achieved by quantization of the network. The first very efficient approaches to aggressive NN quantization were given in 2015. when Courbariaux et al. [35] presented BinaryConnect, an NN that uses binary weight values while maintaining model accuracy. Rastegari et al. in [36] present XNOR-Net, a binary CNN, by evaluating it on the ImageNet data set. Two different approaches were tested: in the first one, only the weights are binarized (Binary-Weight-Networks), while in the other (XNOR-Networks), the input to the network is also binarized. In both approaches, a memory saving of $\sim 32x$ is achieved, and the binary weights allow performing convolutions without using the multiplication operation, which in the first case results in a speedup of $\sim 2x$, and in the second, even $\sim 58x$, compared to the model where the network and data were used in complete precision. The first approaches to network quantization were performed so that the network is trained and then quantized, called post-training quantization, a technique that mainly suffers from a significant loss of accuracy. In the Quantization Aware Training approach, quantization is performed during training, which generally gives less accuracy drop compared to the previous technique. The QKeras [37] and Larq [38] libraries are designed as a quantization extension of the Keras API. In the paper, QKeras enables the approximation of NN weights, biases, and activation functions with low bit values that significantly reduce the network size.

Therefore, the development of such tools and the emergence of ML-compatible FPGA devices have opened up new possibilities for implementing classification networks on L1T. The paper investigates the possibility of additional resource reduction using data selection and quantization. Applied methods enable the usage of the proposed quantized models in classification within the set time limit of 12 μs .

3 Materials and Methods

3.1 Data set

Simulated data are used to design and optimize the detector's geometry, material composition, and readout electronics and to test the performance of the suggested classifiers. They are extremely valuable during the stages of detector development when real data are not yet available. CMS Software Components (CMSSW) is a collection of software needed to simulate CMS detector event data so that scientists can perform different types of analysis. It produces Monte Carlo simulation events, where primary physics processes are generated by programs such as Pythia [39]. A detailed simulation of the CMS detector uses Geant4 [40], a tool that simulates particle interactions with the detector material. For this study, it is necessary to generate images to match the HGAL structure. Hence, the data set used in the research is created based on data obtained from CMSSW version 12.1.0.

3.1.1 Process and Simulation

An electromagnetic (EM) shower is chosen for the data that represents a signal. It is obtained by simulating electrons with the transverse momentum (pt) between 2 and 200 GeV, where pt is a component of momentum perpendicular to the beam line. A minimum bias (neutrino) simulation with a number of PU events of 200 is used to create the background samples. The paper also considers another type of background, quantum chromodynamics (QCD) jets, sprays of particles produced by the hadronization of quarks and gluons. QCD jets background is obtained by simulating the QCD sample with pt from 50 to 80, with a PU value of 200. Electromagnetic showers, produced by electrons and photons, are processes of interest because decays of interesting heavy particles, like Higgs boson, can produce electrons or photons, and it is necessary to isolate them from the often generated, physically irrelevant, PU or QCD jets.

3.1.2 Event Selection

For EM shower, three-dimensional (3D) clusters close to generated electrons and positrons ($\Delta R < 0.2$) are considered, where angular distance is calculated as $\Delta R = \sqrt{(\eta_{cluster} - \eta_{genPart})^2 + (\phi_{cluster} - \phi_{genPart})^2}$, where η is the pseudorapidity, the spatial coordinate describing the angle of a particle relative to the beam axis, while ϕ is an angle that describes the rotation of a particle's trajectory in the plane perpendicular to the beamline. The requirement for the cluster's pt is that it has to be higher than 10 GeV. If there are more 3D clusters close to the generated particle, the clusters whose pt is between 80% and 120% of the pt of the belonging generated particle are selected, and the one with the highest pt value is picked. For PU, 3D clusters with pt higher than 5 GeV, which are not close to generated photons, electrons, and positrons ($\Delta R > 0.2$), are chosen. For QCD jets, 3D clusters close ($\Delta R < 0.4$) to genjets with $pt > 30$ GeV are selected. Genjets are jets produced from generator-level Monte Carlo particles, and they contain information about fractions from different types of generated particles (charged hadrons). Using the above description, 54 000 samples are generated: 18 000 EM samples, 18 000 PU samples, and 18 000 QCD jets samples. In this study, only low-level features are used: pt , position (x,y,z) , layer number, pseudorapidity (η), azimuthal angle (ϕ), and *particle ID* given by the generator. The generated events are split into training (68%), validation (17%), and testing (15%) data.

3.1.3 Image Formation

As shown in [27], energy deposited in calorimeter cells can be treated as pixels of an image, which allows the usage of powerful image-processing techniques. Also, this approach has a significant advantage since additional preprocessing needed for traditional methods for this kind of classification problem, like random forest (RF) or boosted decision trees (BDT), is avoided. Each HGCALE layer is represented as a separate image, resulting in 36 two-dimensional (2D) images. Considering that the first 14 layers belong to the CE-E part of the endcap and the last 22 to CE-H, the classifier can deal with a complete 36-layer image or take only the CE-E or CE-H part.

Passing through the calorimeter layers, the particles/decay products deposit energy in sensor cells grouped into trigger cells (TC). Trigger cells with the highest energies are grouped into 2D clusters within a single layer. Then, individual 2D clusters are connected into 3D ones, which gives complete 3D information about the shower (Figure 3a). If the targeted image size is smaller than an individual 2D cluster range, not all trigger cells participate in such image creation. It is shown in (Figure 3b), where all TCs are bordered depending on whether they participate in image generation (black border) or not (red border).

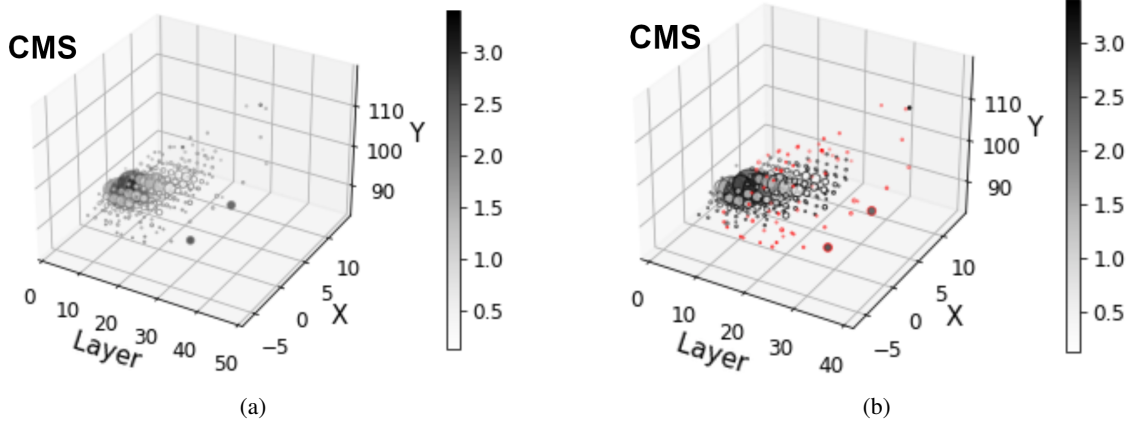


Figure 3: (a) 3D presentation of EM cluster where dot size and color are proportional to deposited pt . (b) the same cluster; cluster points participating in 5x5x36 image generation are bordered black, and the others have red borders.

Subsequently, from the 3D cluster data obtained from CMSSW, virtual images are created by applying the following steps: the line between the center of mass of the 3D cluster and the center of the detector represents the deposition axis. The intersection point of the axis with each individual layer is the center of the region of interest (ROI) on that layer (Figure 4a). Afterward, the range along the coordinate axes is calculated according to the center depending on the target size of the ROI. Finally, a 2D histogram of 5x5 bins, with a bin size $2 \times 2 \text{ cm}^2$ (adjusted to the size of the trigger cell), is applied to summarize all the energies (Figure 4b). The study [41] has shown that with a similar data set, window size enlargement does not significantly impact model accuracy.




Figure 4: Example of layer 13 image generation: **(a)** the projection of deposited particle energies **(b)** the resulting layer image.

As a result of the described image generation procedure, 3D images with a shape $5 \times 5 \times 36$ are generated. The 5×5 parameter represents the window size, while the third value indicates the number of displayed layers (i.e., depth). Since every second layer is used for triggering in the electromagnetic part of the calorimeter, it is described with 14 images (for better visualization, 14 layers of absorbers are shown with the same pixel value as the previous layer), while the hadronic part consists of the remaining 22 layers, one for each layer. Figure 5 represents the longitudinal profile for EM (5a), PU (5b), and QCD showers (5c). It is clear that compared with EM showers (on average), PU and QCD start showering earlier, with a peak reached before layer 10 and a higher energy deposit in the hadronic part. Observing the distribution of the total energy of




Figure 5: Histogram representing average longitudinal profile for **(a)** EM shower, **(b)** PU, **(c)** QCD.

EM clusters with PU and QCD clusters, shown in Figure 6, it is evident that EM clusters, on average, have significantly higher energy than background clusters.

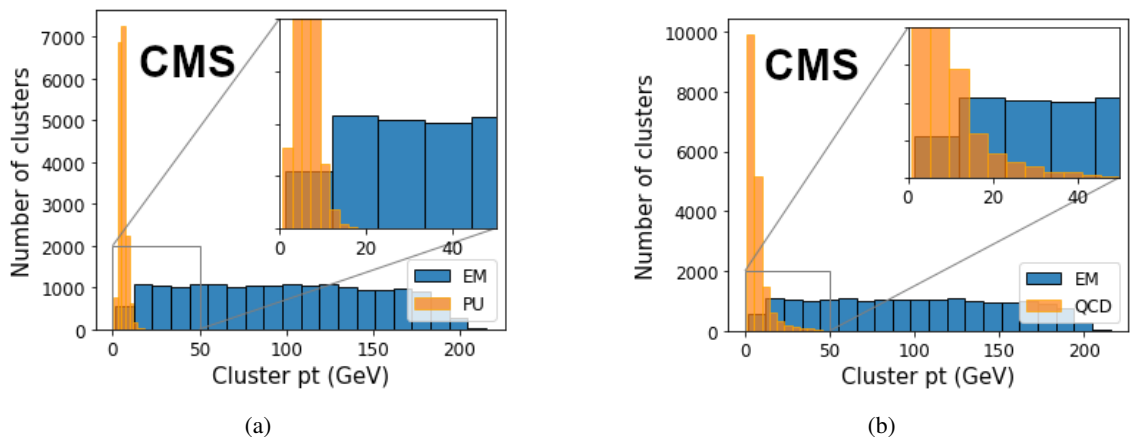


Figure 6: Comparison of average EM and **(a)** PU, **(b)** QCD clusters energie.

The models are tested on full HGCAL image (Figure 7a), but for some of the tested approaches, like reduction of the longitudinal profile, the HGCAL image is split into CE-E and CE-H image parts. CE-E image has a shape $5 \times 5 \times 14$ (Figure 7b) and shows only the electromagnetic part of the calorimeter, while CE-H image with a shape $5 \times 5 \times 22$ represents a hadronic part.

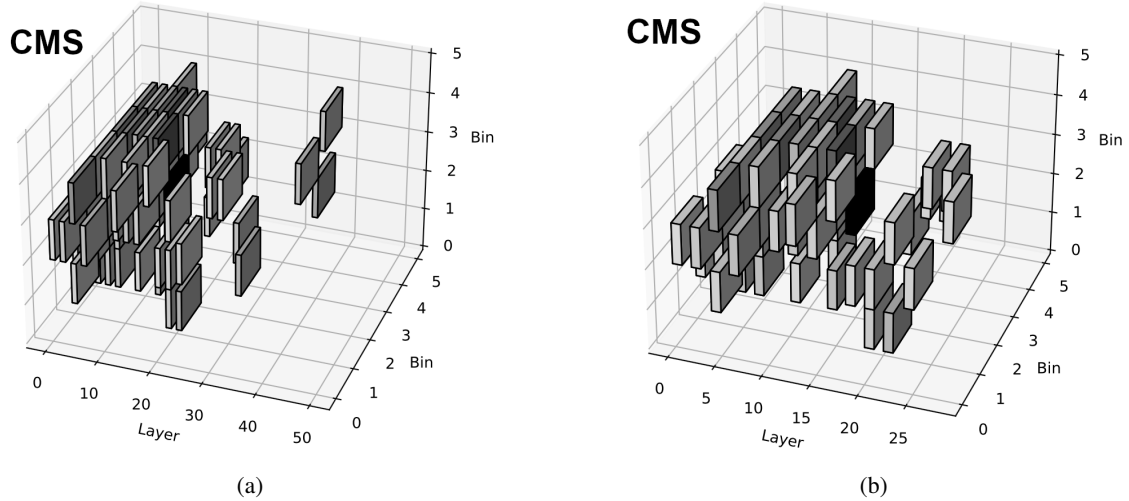


Figure 7: HGCAL image example, covering (a) complete HGCAL (b) just CE-E layers.

3.1.4 Image Preprocessing

Pixel intensities can have large variations between different images; quite often, they span a few orders of magnitude. In addition, there is a significant number of outliers, especially in the PU data set. The feature normalization is applied to improve the numerical stability of model(s) and speed up the training process. The standard procedure to handle outliers is to remove them. But in our case study, outliers are the highest energy points, and deleting them would mean losing vital information. Some standard approaches in the presence of outliers are to use (quantile-based) capping [42] or robust scaler normalization [43]. Tests are done with both techniques, and they show similar results. Due to simplicity, capping will be applied, shown in more detail in section 3.2.

The selection and quantization of input data are also considered to reduce the model's memory footprint, enabling its implementation on FPGAs and the execution of event selection in the set time frame (12.5μ). Two types of data quantizers are created: linear, where the quantization levels are evenly distributed, and non-linear, where the relationship between the quantization levels is logarithmic.

3.2 Data refinement

In addition to the data set containing complete HGCAL images, given in full precision, different ways of data refinement are considered. Although the size of the images is small (5×5), a high number of channels (HGCAL layers) results in a high number of pixels - features. Depending on the studied number of layers, it is 350 (for the CE-E part) or 900 for the complete HGCAL. Therefore, although the new detector structure enables analysis of a detailed collision image, the possibility of dimensionality reduction (the process of reducing the number of features) is examined. In the case where image data are observed, the number of pixels can be considered as the number of features; each pixel represents one feature.

Two main approaches in dealing with high-dimensional data are feature extraction (FE) and feature selection (FS). FE creates new features combining the existing ones, while FS identifies the relevant features and removes the redundant and irrelevant features. In [44], the authors review FE methods that reduce processing time while providing higher recognition accuracy. Although FE is the preferred approach in image processing, because of time and resource constraints, both approaches are used to scale down the dimensionality of the data without performance degradation. The experiment uses ad-hoc solutions suitable for our specific physically motivated classification problem. Considering the low computational time, two techniques are applied: filtering out the irrelevant pixel values (FS approach) and combining/summing particular layers (FE approach). Further reduction of resource consumption is achieved by data quantization. Data refinement is conducted in two steps: first, data selection, followed by data quantization.

3.2.1 Selection

In the first step, the selection is applied to full-precision data using one of the rules:

1. **Keep energies above the threshold.**

Here, a fixed threshold value of (0.1, 0.5, 1, 2) GeV is applied, all values greater than the threshold are kept, and others are set to 0. The smallest energy values are assumed to have minimal impact

on determining if a cluster represents a signal or background. The reason behind the small values proposed for the thresholds is that TC energies are much smaller than cluster pt .

2. Keep layer maximum (single sample per layer).

With this approach, just the maximum value in each layer is kept, while other values are set to 0. This approach assumes that the highest pt deposits have a decisive role in shower classification.

3. Capping.

Data capping is a method in which the maximum value of a feature is set to a specific value. It is a standard technique to treat outliers and is often a necessary step in executing the model in FPGA. A preliminary test with models implemented in FPGA using hls4ml has shown a drop of the accuracy of even 10% in dealing with no-capped data (compared with QKeras model accuracy).

4. Reduction of longitudinal profile.

Similar to the approach presented in [45] where detector images consist of three subdetector channels: one each for CE-E, CE-H, and one for the reconstructed tracks, the following data refinement method is based on summing pt /pixels from particular detector layers. By summing the values of individual layers, it is possible to considerably reduce the image (channels) size, drastically decreasing the NN model's size. Two different solutions are examined:

- Considering the fact that EM shower leaves pt deposits mainly in the CE-E part of the HGAL, all CE-H layers are projected in one plane, reducing the number of CE-H layers from 22 to 1. This approach is named E+Hf, and the same logic with the CE-H part is also applied in the next solution.
- Here, the fact that EM starts showering in the earlier detector layers and reaches the peak between layers 10 and 15 is used. Respectively, the CE-E is divided into three parts: layers 1-5, where the EM shower starts; layers 6-15, where the most pt is deposited; and layers 16-28, shown in Figure 8a. Each part is summed (Figure 8b), resulting in 3 5x5 CE-E images, and the CE-H part is reduced, creating a 5x5 CE-H image. This approach is referred to as 3Ef+Hf.

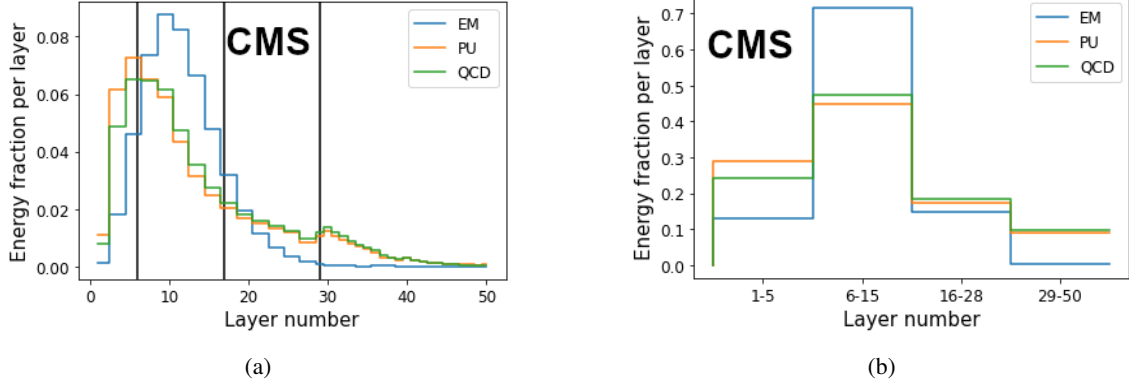


Figure 8: Distribution of pt across HGAL layers for EM, PU and QCD showers used in 3Ef+Hf approach (a) 3 CE-E + CE-H layers split (b) summed pt in grouped layers.

3.2.2 Quantization

The reason behind the data quantization is it enables additional limitations of resource usage. Data are quantized in three ways, using custom quantizers:

1. Selection bit(s).

Here, all data values that have passed the previously applied selection steps have a value of 1. Using this approach, FPGA implementation requires just 1 bit for input data layer presentation.

2. Uniform quantizer.

Min-max normalization is applied on capped data, multiplied by $2^n - 1$, where n is a chosen number of bit width. The result is rounded to the nearest integer. The proposed quantizer function is presented in Figure 9a.

$$Q(x) = \text{round}\left(\frac{x - \min}{\max - \min} \cdot (2^n - 1)\right) \quad (1)$$

3. **Non-uniform quantizer.** The quantizer is defined as follows

$$Q(x) = \begin{cases} \text{round}(\log_2(x)) & \text{if } x > 0 \\ -2 & \text{otherwise} \end{cases} \quad (2)$$

The next step is to rescale data (using a min-max scaler), so it does not contain negative values. The non-uniform quantizer function is presented in Figure 9b.

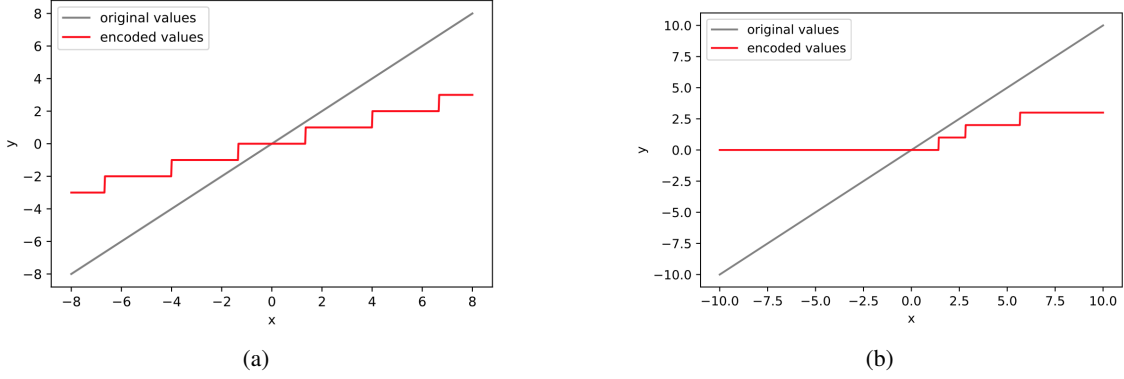


Figure 9: Quantizer function $y = Q(x)$: **(a)** in the case of a uniform quantizer, **(b)** in the case of non-uniform quantizer .

The overview of data refinement methods described in this section is given in Table 1.

Table 1: Data refinement methods.

Refinement methods	
Selection	Quantization
Keep energies above the threshold	Selection bit
Keep layer maximum	Uniform quantizer
Capping	Non-uniform quantizer
Reduction of longitudinal profile	

3.3 Model Architectures

Already in 1989, Cybanko [46] proved that standard multilayer feedforward networks could approximate any continuous function of n real variables to any desired accuracy. The same year, Hornik et al. [47] showed that multilayer feed-forward networks with as few as one hidden layer are universal approximators. In the last decade, state-of-the-art models like CoAtNets, ResNet, EfficientNet, YOLOv5, and many others have been developed. They achieve excellent results, but the high computational complexity of neural network-based classifiers and their requirements on the resources don't allow their usage in LIT, where power consumption is limited, as well as budget on available devices.

3.3.1 Model Size

The first strategy to reduce computational complexity is limiting the network size in the sense of several layers and neurons while taking care of model performance. In more complex computer vision tasks, it would lead to significantly worse performance. Still, it is shown in [41, 48] that deep multilayer perceptron (MLP) with just three layers can be successfully used in EM shower/PU classification. Unlike in the previous approaches, the new classification cases are added, with background represented as QCD events and a mixed background scenario with both QCD and neutrino PU events.

Two simple NN models are considered: MLP and CNN, but because of space constraints, just the last classification case for both models is analyzed in detail since it includes both types of backgrounds. A hyperparameter optimization software framework Optuna [49] is used to determine the models' structure. Optuna uses a tree-structured Parzen estimator sampler, a Bayesian optimization technique that models the search space. It uses the history of previously evaluated hyperparameter configurations to sample the following ones.

For the CNN, the search space contains 6-16 kernels for the Conv2D layer and 8-18 nodes in the following 1-2 dense layers, so far resulting in 1452 possible network combinations. For the MLP, the search space contains 2-3 layers, each with 4-16 nodes, which generates 2366 combinations. The proposed activation functions for intermediate layers for both models are ReLu and sigmoid, while the output layer uses softmax for the activation function. Search space also contains optimizers: Adam, RMSprop, and SGD, each with a range of learning rates, which greatly increases the number of possible combinations in search space on 8712 for CNN and 14 196 for MLP. When more architectures with similar accuracies are highly ranked, the one that minimizes the number of trainable parameters is chosen.

3.3.2 Model Quantization

The second strategy to reduce memory footprint and computational complexity is model quantization, a process of model transformation into an equivalent representation but using parameters and computations at a lower precision. According to [50], the usage of low-precision fixed integer values representation has the potential to reduce the memory footprint and latency by a factor of 16x. Despite being fast and very easy to use, the post-training quantization approach is not an option because it suffers from significant degradation in model accuracy in case of precision lower than 8 bits [51].

QKeras library enables quantization-aware training using a simple replacement for Keras layers, greatly simplifying the quantization process. QKeras allows heterogeneous quantization: it is possible to use different quantization levels and different quantized activation functions on each layer. The concept of a straight-through estimator presented in [52] is used, hence the forward pass applies the quantization functions and the backward pass assumes the quantization as the identity function to make the gradient differentiable [53]. After exhaustive research, it is decided to use the quantizer *quantized_bits* defined as follows

$$2^{int-b+1} \text{clip}(\text{round}(x * 2^{b-int-1}), -2^{b-1}, 2^{b-1} - 1), \quad (3)$$

for weights and kernels, where x =input, b =number of bits for the quantization, int =how many bits are to the left of the decimal point.

Quantized_relu is chosen as a quantized replacement for the ReLu activation function, used in the inner layers of models.

3.4 Evaluation

This section presents the metrics used to demonstrate the impact of quantization, model selection, and data refinement techniques on model performance.

3.4.1 Standard Metrics

Essential evaluation metrics usually used to illustrate the performance of the classifiers are:

- accuracy (ACC), percentage of correct classifications, calculated as

$$ACC = \frac{TP + TN}{TP + TN + FP + FN} \quad (4)$$

- sensitivity, recall, or true positive rate (TPR)

$$TPR = \frac{TP}{TP + FN} \quad (5)$$

- specificity, or true negative rate (TNR)

$$TNR = \frac{TN}{TN + FP} \quad (6)$$

- false negative rate (FNR)

$$FNR = \frac{FN}{FN + TP} = 1 - TPR \quad (7)$$

- false positive rate (FPR)

$$FPR = \frac{FP}{FP + TN} = 1 - TNR \quad (8)$$

- F1 score

$$F1 = 2 \frac{PPV * TPR}{PPV + TPR} \quad (9)$$

where TP stands for the number of correctly predicted positive classes, TN is the number of correctly predicted negative classes, FP is the number of samples incorrectly predicted as positive classes, and finally, FN is the number of samples incorrectly predicted as negative classes. In the case of class imbalance, accuracy is not a good metric. Instead, the $F1$ score is defined as the harmonic mean of precision (PPV, where $PPV = \frac{TP}{TP+FP}$) and recall, designed to work well on imbalanced data.

Also, it is a standard technique for summarizing classifier performance to produce a receiver operating characteristic curve (ROC), which represents the false positive rate (FPR) versus the true positive rate (TPR), and the corresponding area under the curve (AUC) is calculated. In the HEP, it is common to interpret the ROC curve in terms of the signal efficiency (true positive rate) vs. background rejection (true negative rate). Also, physicists are often interested in other metrics, like signal efficiency at some fixed background rejection.

There exist many metrics for classification models, but for particle detector triggers, it is essential to detect as many interesting events as possible in addition to accuracy. Therefore, the trigger algorithm should operate at a very low FNR. On the other hand, low FPR ensures staying within the available trigger bandwidth but does not impact the success of the analysis.

3.4.2 Classification threshold adjustment

Depending on the application, binary classification often needs to optimize specific metrics, such as the minimization of FNR or FPR (the same as maximizing TNR or TPR). Sometimes, the default threshold (0.5) may not represent an optimal interpretation of the predicted probabilities. It is the case when the predicted probabilities are not calibrated, which is a known problem for modern neural networks [54]. Also, when the cost of one type of misclassification is more important, changing the default decision threshold is one way to handle it. According to [55], misclassifying an actual positive example into a negative is often more expensive than an actual negative example into a positive. Thresholding is a cost-sensitive meta-learning method described in [56], based on the selection of the probability that minimizes the total misclassification cost on the training instances as the threshold for predicting testing instances. Therefore, when interpreting the predictions of a model, sometimes there is a need to change the default decision threshold of 0.5. If TPR and TNR have the same importance for the L1 trigger, one way of calculating the cut-off is using Youden's J statistic ([57]), defined as follows:

$$J = \text{Sensitivity} + \text{Specificity} - 1, \quad (10)$$

or simplified:

$$J = TPR - FPR. \quad (11)$$

Youden's index is often used in combination with ROC curve analysis. The index is calculated for all points of a ROC curve, and the optimal threshold with the largest J value is chosen. If both metrics, FNR and FPR, are not equally important, using the trial and error method helps find the threshold to lower the targeted metric.

3.5 Baseline Particle Classification Methods

As baseline methods to be compared to NN-based approaches, the simple cut-off and Random Forest (RF) are applied.

3.5.1 Cut-off

Considering the fact that average EM cluster pt is much higher than PU or QCD cluster pt as shown in Figure 6, the simplest approach to particle classification is to apply linear selection criteria: (pt) cut-off value can be used to accept or reject clusters as EM, as it is likely that the low pt cluster is background. The test is done with two cut-off values: the first is the maximum PU (QCD) cluster pt , and the second is the third quartile PU (QCD) cluster pt value.

3.5.2 Random Forest

Another baseline method is RF, a supervised ML algorithm that combines the output of multiple decision trees to reach a single result. Together with BDT, it is often used in particle classification, as it is shown in [58, 59]. Both RF and BDT usually use data sets containing well-known physical parameters as features. Because of time constraints needed for L1T selection, the same input as for NN models is used: raw detector data - HGAL images.

RF consists of a large number of decision trees that operate as an ensemble. Each individual tree in the RF declares a class prediction, and the class with the most votes becomes the model's prediction. In the RF

learning process, there are two major layers of randomness. The first random component is that RF uses a random (bootstrapped) sample of the original training data set for each individual decision tree. As a result, in the vast majority of cases, the training sets for the trees are different from one another, which reduces the correlation between trees and improves the generalization of the predictions. Another randomness is in a (random) selection of the features considered at each node. The method itself is introduced by Leo Breiman ([60]), and the detailed method analysis is available in [61]. RF’s simple usage and flexibility, together with the possibility to handle classification and regression problems, have contributed to its widespread usage. It is considered to be one of the best off-the-self-learning algorithms, requiring almost no tuning, although it enables fine control over the model that is learned.

An open-source machine learning library for the Python programming language named scikit-learn [62] is used for testing RF classifiers as it provides a robust implementation combining both algorithmic and code optimizations. To improve the accuracy of the RF classifier for particular situations, there are parameters that may be tuned. The grid search is used for performing hyper-parameter optimization, employing the next hyperparameter search space:

- *max_features*: maximum number of features random forest considers splitting a node, with available values 'auto', 'sqrt', 'log2'. The default value is 'auto'.
- *n_estimators*: the number of trees in the forest, with values 100, 200, 300. Increasing this hyperparameter generally improves the performance of the model. The drawback is it also increases the computational cost of training and predicting. The default value is 100.
- *max_depth*: maximum number of levels in each decision tree, using values 8, 16, 20. The default value is None, which means that nodes are expanded until all leaves are pure or until all leaves contain less than 2 samples.

An addition parameter, *criterion*, which measures the quality of the split, with values 'gini' and 'entropy', is added to the search space too, and the loss function which is optimizing is accuracy.

3.6 hls4ml

Recently, in 2018. Duarte et al. [31] presented hls4ml, a compiler that translates ML models into Register Transfer Level for FPGAs using the High-Level Synthesis (HLS) tool. hls4ml enables models’ firmware feasibility without comprehensive Verilog/VHDL experience, thus accelerating the development cycle. The hls4ml authors focus on the task of FPGA triggers of the ATLAS and CMS experiment, The models described in this study use Vivado (2020.1) for HLS synthesis with a Xilinx Kintex UltraScale FPGA (part number xcvu13p-fhgb2104-2-e) [63] as the target device, with the synthesis clock frequency set to 200 MHz. To achieve as low as possible latency while maintaining accuracy, different settings are considered depending on the benchmark model. In the FPGA, hls4ml operates with fixed-point arithmetic. By adjusting the fixed-point data type and using profiling tools, it is possible to lower resource consumption. The hls4ml profiling tool helps to decide appropriate model precision. Without profiling, low global precision settings that help to reduce the FPGA resource usage of a model may result in a loss of model performance if chosen inappropriately). Another hls4ml feature that helps to optimize an NN is the possibility to determine the parallelization of the calculations in each layer. It is done by configuring the parameter “reuse factor,” which determines the number of times a multiplier is used to do a computation. It is clear that a fully parallelized process results in low latency but requires more resources; respectively, latency and required resources are inversely proportional. After setting the specified parameters, Vivado HLS is used to synthesize the model.

4 Results and Discussion

The following section presents the structure of NN used in the experiment and explains the results of the model and data quantization and selection methods described in the previous section.

4.1 Model Architectures

Here, the architecture of used CNN and MLP models is presented, followed by the model quantization results.

4.1.1 Base NN Models

Optimization results for the hidden layers of CNN and MLP models for each of the classification cases are shown in Table 2. For example, the CNN model for EM vs. PU classification contains a convolutional layer

with 11 3x3 kernels followed by one dense layer with 12 nodes. It can be noticed that the optimizer gives quite a similar NN structure for MLP in all classification tasks, while the suggested CNN in EM vs. MIX (Figure 10) is somewhat bigger, it has an additional dense layer, which is expected because the classification problem is more complex.

Table 2: The results of model structure optimization for CNN and MLP for each classification case. The table describes just hidden layers.

case/model	CNN	MLP
EM vs. PU	conv1 layer: 11 kernels dense1 layer: 12 nodes	dense1 layer: 5 nodes dense2 layer: 13 nodes
EM vs. QCD	conv1 layer: 12 kernels dense1 layer: 10 nodes	dense1 layer: 4 nodes dense2 layer: 15 nodes
EM vs. MIX	conv1 layer: 6 kernels dense1 layer: 12 nodes dense2 layer: 13 nodes	dense1 layer: 8 nodes dense2 layer: 13 nodes

Activation functions are the same in all cases, inner layers use ReLu, and classification is done using softmax. For both models, the loss function is the binary cross-entropy, and the optimization is done using the Adam algorithm with a lr learning rate obtained by the optimizer and a batch size of 128. Learning rate values determined by the optimizer are in the range of 0.009 to 0.07, depending on the model and classification case. The early stopping is implemented if no progress is seen beyond 10 epochs. Both training and testing sets contain balanced samples of the classes in all classification tasks.




Figure 10: Vizualization of CNN (and QCNN) model used for EM vs. MIX classification.

The first group of tests is done on the full data set without any normalization, selection, or quantization. Input images represent the full HGCAL profile, all 36 layers.

The results of baseline NN methods for different types of EM vs. background classification are shown: for PU in Table 3, for QCD in Table 4, and for mixed background in Table 5

Table 3: The results for the base CNN and MLP model for EM vs. PU.

model	acc	prec	rec	f1	AUC	fpr	fnr
CNN	0.9861	0.9899	0.9822	0.9861	0.9951	0.0100	0.01781
MLP	0.9830	0.9891	0.9767	0.9829	0.9920	0.0107	0.0233

Table 4: The results for the base CNN and MLP model for EM vs. QCD.

model	acc	prec	rec	f1	AUC	fpr	fnr
CNN	0.9698	0.9756	0.9637	0.9696	0.9914	0.0241	0.0363
MLP	0.9639	0.9665	0.9611	0.9638	0.9884	0.0333	0.0389

It can be noticed that in all cases, CNN gives slightly better accuracy than MLP. Both models are better at classifying between EM and neutrino PU clusters than between EM and QCD jets. EM vs. MIX classifier,

Table 5: The results for the base CNN and MLP model for EM vs. MIX.

model	acc	prec	rec	f1	AUC	fpr	fnr
CNN	0.9731	0.9758	0.9704	0.9731	0.9899	0.0241	0.0296
MLP	0.9689	0.9665	0.9715	0.9690	0.9893	0.0337	0.0285

as expected, is somewhat better than EM vs. QCD, and worse than EM vs. PU. In the rest of the paper, the focus is put on EM vs. MIX classification problem. The models are minimized, making them suitable for FPGA implementation and playing an important role in L1 triggering. At the same time, the methods described in the previous section are used to improve accuracy and minimize FNR.

4.1.2 Models Quantization

To examine the effect of weights/biases/kernels quantization level on model accuracy, tests are done iterating over the number of bits for the quantization in the range from 2 to 16, and there is no significant difference in achieved results. A comparison of accuracy and FNR for different quantization levels given in Fig 11 shows that accuracy is very similar for all quantization levels, while FNR varies more. Therefore, given that the goal is model minimization in the sense of memory footprint and processing requirements, aggressive 2-bit quantization of model weights/biases/kernels is chosen to be applied for model quantization, even if FNR is higher than, for example, for 6-bit quantized models. In the rest of the paper 2-bit quantized CNN is referred to as QCNN, and accordingly, the quantized 2-bit version of MLP is named QMLP.

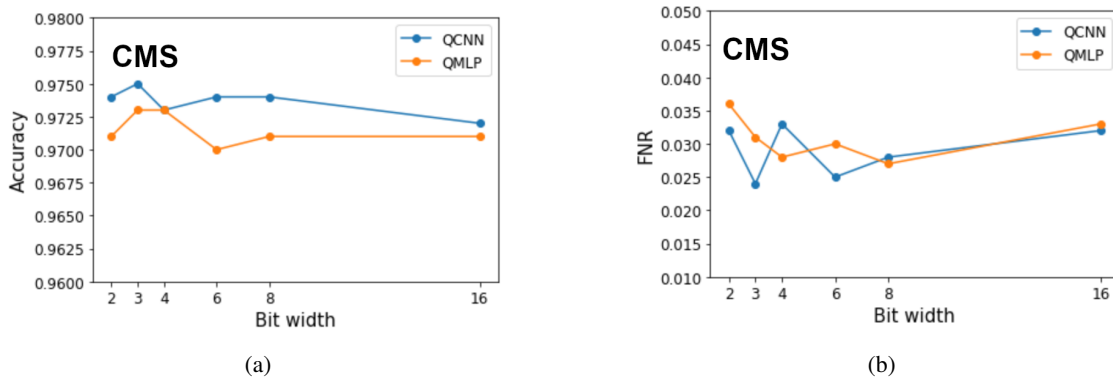


Figure 11: The effect of weights/biases/kernels quantization level on CNN and MLP model: (a) Accuracy (as a function of bit width), (b) FNR (as a function of bit width).

The performance of the 2-bit quantized version of CNN and MLP is summarized in Table 6. It is evident that the quantized models have the same discrimination power as the baseline architectures. Considering memory and processing limited devices on which the trigger algorithms will be implemented, from now on, the study is conducted just on quantized models. Comparing QCNN with QMLP, it is evident that both models provide very similar results. All metrics are slightly higher for QCNNs, confirming that convolutional layers ensure better performance in computer vision problems.

Table 6: The results for the 2-bit quantized CNN and MLP models.

model	acc	prec	rec	f1	AUC	fpr	fnr
QCNN	0.9737	0.9790	0.9681	0.9736	0.9875	0.0207	0.0319
QMLP	0.9713	0.9782	0.9641	0.9711	0.9812	0.0215	0.0359

4.2 Baseline Particle Classification Methods

4.2.1 Cut-off

The data set contains 18 000 samples for each of the three types of clusters, and 3628 of them, or 20.16%, are low pt EM clusters, respectively EM clusters with pt lower than maximum PU cluster pt that would be lost, using these criteria. Compared with QCD clusters, there are 8159, or 45.33% low pt clusters. On the other side, using the cut-off value lower than the maximum PU (QCD) cluster pt , the number of wrongly classified PU (QCD) clusters grows rapidly, resulting in unnecessary storage of useless data. In the case of

the third quartile cut-off value, 45 EM clusters (0.25 %) are lost in the case of the EM/PU classification and 87 (0.48 %) in the EM/QCD classification.

Therefore, simply accepting/rejecting clusters based on cluster pt value is not an acceptable option for the LIT selection method.

4.2.2 Random Forest

Using the same metrics as for NN models, the RF classifier with default settings gives the results presented in Table 7.

Table 7: The results of Random Forest classifier.

settings	acc	prec	rec	f1	AUC	fpr	fnr
default	0.9693	0.9731	0.9652	0.9691	0.9693	0.0267	0.0348
optimized	0.9694	0.9735	0.9652	0.9693	0.9694	0.0263	0.0348

The hyperparameters optimization gives the next combination of tuned hyperparameters: *criterion*: 'gini', *max_depth*: 40, *max_features*: 'auto', *n_estimators*: 200. The achieved metrics are almost the same as those obtained with the predefined settings, which is shown in Table 7. Even though the achieved accuracy is similar to the NN model's case, both FNR and FPR are slightly higher. The next attempt is to run a grid search to optimize the recall score, but with no improvement (maximizing recall, FNR is minimized).

4.3 Classification threshold adjustment

Two popular evaluation metrics to estimate the model performance, accuracy, and AUC are for QCNN and QMLP models above 97%, indicating good performance. Following the procedure described in section 3.4.2, for QCNN, the largest J statistic (0.947) is achieved using a threshold value of 0.939913, and for QMLP, the maximum J value (0.944) is for the threshold of 0.851953. However, this threshold value selection has resulted in better FPR results but increased FNR value for both models, as seen in Table 8.

On the other side, with the trial and error approach, using the threshold value of 0.1, FNR decreases to 0.02 for QCNN and to 0.03 for QMLP. Although the differences in achieved FNR are not high, it is important to remember that the volume of the data set on which the model is intended to be used is very high, and even small differences in percent result in high absolute values. For example, for QCNN, the difference between the best achieved FNR and FNR obtained for the default threshold is 0.0119 (1.19%). Applied to the data set with just 1 000 000 samples, it ends with 11 900 additionally recognized EM clusters, and any of them can have a significant role in some new physical discovery. The comparison of results obtained using different thresholds is given in Table 8.

Even though the obtained FNR (3.19% for QCNN and 3.59% for QMLP) and FPR (2.07% for QCNN and 2.15% for QMLP) are not high for the base quantized models with default threshold, using the thresholding technique they can be further optimized. It should be noted that the number of FP increases by changing the threshold to decrease the FN.

Table 8: The results for the QCNN and QMLP (gray background) with different threshold values.

threshold	acc	prec	rec	f1	AUC	fpr	fnr
0.1	0.9698	0.9604	0.9800	0.9701	0.9875	0.0404	0.0200
0.1	0.9689	0.9678	0.9700	0.9689	0.9812	0.0322	0.0300
0.5	0.9737	0.9790	0.9681	0.9736	0.9875	0.0207	0.0319
0.5	0.9713	0.9782	0.9641	0.9711	0.9812	0.0215	0.0359
0.939913	0.9744	0.9845	0.9641	0.9742	0.9875	0.0152	0.0359
0.851953	0.9717	0.9826	0.9604	0.9713	0.9812	0.0170	0.0396

It can be noticed there is no big difference in moving the threshold based on J-statistics or looking for minimization of FPR and FNR difference, but depending on needs, it is clear that it is possible to optimize targeted costs by choosing the appropriate threshold.

4.4 Data Refinement

This section describes the effect of different data refinement methods on NN model performances.

4.4.1 Selection

1. Keep energies above the threshold.

Before processing, in the training data set for the 5x5x36 sample, there are just 11.37% non-zero values. Using different threshold values, for described method, the percent of non-zero values decreases: for threshold value 0.1 GeV to 5.99%, for 0.5 GeV to 1.67%, for 1 GeV to 1.01%, and finally for threshold value 2 GeV to 0.6%. As it is visible from Table 9, as the threshold value increases, accuracy slightly drops, but there is an increase in FNR value for the threshold values higher than 0.5. Therefore, future tests are not done for those threshold values.

Table 9: The QCNN and QMLP (gray background) results with different thresholds applied on energie/pixel values.

threshold (GeV)	acc	prec	rec	f1	AUC	fpr	fnr
0.1	0.9722	0.9800	0.9641	0.9720	0.9844	0.0196	0.0359
0.1	0.9674	0.9647	0.9704	0.9675	0.9835	0.0356	0.0296
0.5	0.9706	0.9711	0.9700	0.9705	0.9851	0.0289	0.0300
0.5	0.9676	0.9685	0.9667	0.9676	0.9823	0.0315	0.0333
1	0.9693	0.9774	0.9607	0.9690	0.9826	0.0222	0.0393
1	0.9667	0.9684	0.9648	0.9666	0.9851	0.0315	0.0352
2	0.9578	0.9870	0.9278	0.9565	0.9657	0.0122	0.0722
2	0.9556	0.9858	0.9244	0.9541	0.9643	0.0133	0.0756

2. Keep layer maximum (single sample per layer).

Models, where just one value per image layer is kept, show no drop in classification precision compared with base quantized versions, according to results shown in Table 10

Table 10: The results for the keep layer maximum approach for QCNN and QMLP.

model	acc	prec	rec	f1	AUC	fpr	fnr
QCNN	0.9706	0.9725	0.9685	0.9705	0.9834	0.0274	0.0315
QMLP	0.9689	0.9713	0.9663	0.9688	0.9814	0.0285	0.0337

3. Capping.

Data capping is done by using different values (2, 4, 8, 16). For QMLP, there is no notable difference in metrics, but for QCNN, there is an improvement with FNR for all capping values (Table 11), no matter which capping value is chosen. The impact of the applied capping value on the model's accuracy is shown in Figure 12a and on FPR in Figure 12b.

Table 11: The results for the QCNN and QMLP (gray background) with different capping values.

c_value	acc	prec	rec	f1	AUC	fpr	fnr
2	0.9739	0.9723	0.9756	0.9739	0.9856	0.0278	0.0244
2	0.9678	0.9688	0.9667	0.9677	0.9811	0.0311	0.0333
4	0.9741	0.9737	0.9744	0.9741	0.9860	0.0263	0.0256
4	0.9704	0.9757	0.9648	0.9702	0.9836	0.0241	0.0352
8	0.9731	0.9674	0.9793	0.9733	0.9871	0.0330	0.0207
8	0.9713	0.9775	0.9648	0.9711	0.9811	0.0222	0.0352
16	0.9744	0.9741	0.9748	0.9745	0.9861	0.0259	0.0252
16	0.9702	0.9725	0.9678	0.9701	0.9824	0.0274	0.0322

4. Reduction of longitudinal profile.

Complete HGCAL image contains 5x5x36= 900 pixels per image, resulting in 2807 parameters for QCNN (respectively 7353 for QMLP). Here, QCNN is used as a classifier. As the result of reducing the longitudinal profile following procedures described in Section 4, the number of model parameters is drastically decreased. Described solutions are examined and results are shown in Table 12 :

- The QCNN model using a full HGCAL image has 2807 parameters, and with this approach (E+Hf), the number of parameters is reduced to 1673.

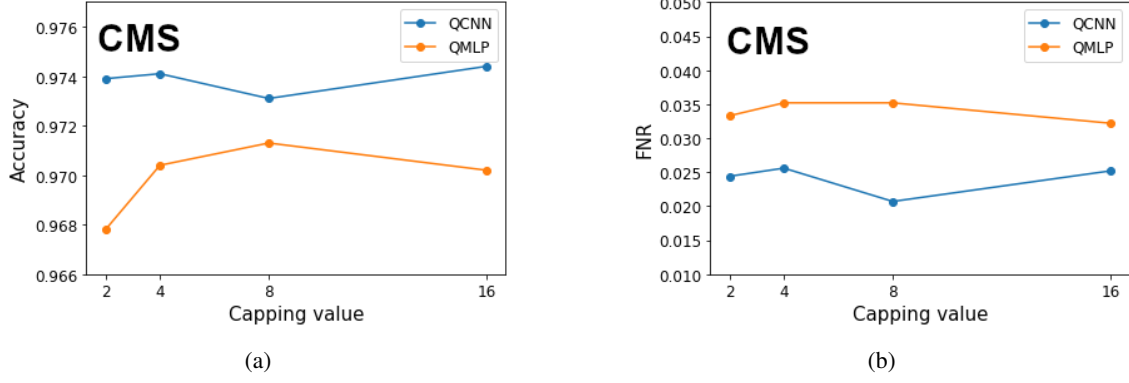


Figure 12: (a) Model accuracy and (b) FNR as a function of capping value.

- This approach (3Eh+Hf) replaces the full HGAL image (5x5x36) with a 5x5x4 image, decreasing the NN parameter numbers on 1079.

Table 12: The results for the reduction of longitudinal profile approach for QCNN.

model	acc	prec	rec	f1	AUC	fpr	fnr
E+Hf	0.9728	0.9691	0.9767	0.9729	0.9855	0.0311	0.0233
3Ef+Hf	0.9702	0.9686	0.9719	0.9702	0.9883	0.0315	0.0281

As it is clear from the above description, in approaches 1, 2, and 4, the HGAL layer structure image is retained. The global picture of particle decay is kept but simplified, emphasizing the importance of high pt points above very small pt deposits. Theoretically, instead of sending the complete image, just the position and value kept after the applied selection could be sent, notably cutting back the number of parameters in the NN's input layer. Using the approach 4, the number of layers can be drastically decreased, ultimately significantly reducing the number of neural network parameters.

4.4.2 Data Quantization

In this section, different data quantization approaches are presented.

1. Selection bit(s).

The combination of data thresholding and selection bit method would simplify/reduce the amount of data needed for acceptance/rejection decision even more. According to performance metrics shown in Table 13, for the threshold value 0.1, there is a drop of accuracy for both models, while a threshold value of 0.5 followed by selection bits method for QCNN gives results comparable with complete full precision data.

Table 13: The results of the selection bit method applied after thresholding, for the QCNN and QMLP (gray background) model.

threshold (GeV)	acc	prec	rec	f1	AUC	fpr	fnr
0.1	0.9611	0.9615	0.9607	0.9611	0.9832	0.0385	0.0393
0.1	0.9302	0.9029	0.9641	0.9325	0.9752	0.1037	0.0359
0.5	0.9659	0.9611	0.9711	0.9661	0.9784	0.0393	0.0289
0.5	0.9676	0.9664	0.9689	0.9676	0.9839	0.0337	0.0311

2. Uniform quantizer.

A mandatory step to use uniform quantization is capping. Otherwise, as there is a big range of values used for image presentation, most of them would become 0 after applied quantization. Fig 13 shows model accuracies and FNR depending on the capping value applied before uniform quantization.

According to Fig 11, for a lower number of chosen bit widths (2-4) for data quantization, both QCNN and QMLP achieve better results for accuracy and FNR. Capping values 2 and 4 give better results

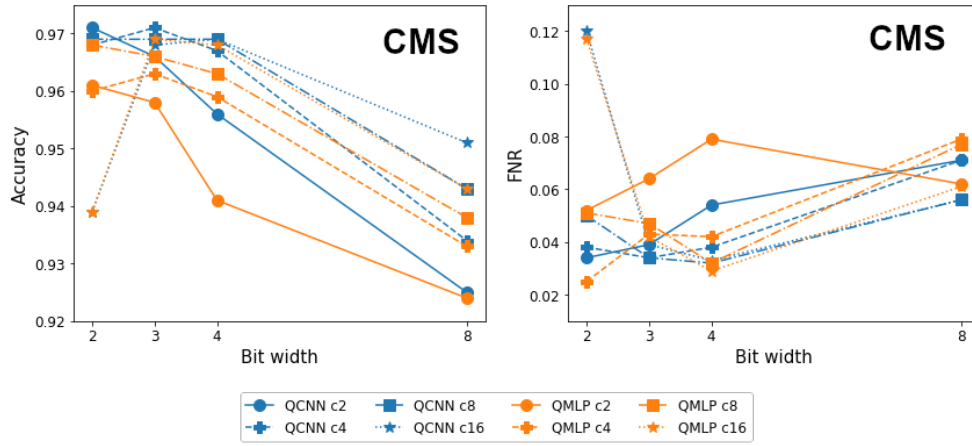


Figure 13: Model accuracy and FNR as a function of data bit width and capping value

Table 14: Uniform quantizer with different quantization levels, QCNN and QMLP (gray background) model.

bit width	acc	prec	rec	f1	AUC	fpr	fnr
2	0.9676	0.9727	0.9622	0.9674	0.9844	0.0270	0.0378
2	0.9598	0.9461	0.9752	0.9604	0.9825	0.0556	0.0248
3	0.9711	0.9757	0.9663	0.9710	0.9870	0.0241	0.0337
3	0.9633	0.9689	0.9574	0.9631	0.9816	0.0307	0.0426
4	0.9672	0.9723	0.9619	0.9670	0.9789	0.0274	0.0381
4	0.9593	0.9606	0.9578	0.9592	0.9791	0.0393	0.0422
8	0.9339	0.9383	0.9289	0.9336	0.9573	0.0611	0.0711
8	0.9330	0.9438	0.9207	0.9321	0.9635	0.0548	0.0793

than 8 and 16; hence 4 is applied as the capping value whenever uniform quantization is used, and results are given in Table 14.

According to Table 14, models with 2 and 3 input data bit width outperform models with wider data precision.

3. Non-uniform quantizer.

The models based on non-uniformly quantized data with results in Table 15 show somewhat lower performances than models based on 2-bit uniformly quantized input.

Table 15: Non uniform quantizer

model	acc	prec	rec	f1	AUC	fpr	fnr
QCNN	0.9685	0.9766	0.9600	0.9682	0.9860	0.0230	0.0400
QMLP	0.9678	0.9730	0.9622	0.9676	0.9834	0.0267	0.0378

4.5 hls4ml

Considering model accuracy and FNR, along with model size (in terms of the number of parameters and precision of input data), the first model implemented in hls4ml is QCNN with full precision data, with capping value 8, referred to as QCNN_FP. It is compared with QCNN with 2-bit quantized input data, from now on referred to as QCNN_U2. The last model implemented in hls4ml is QMLP with 2-bit quantized input data, named QMLP_U2E. It must be emphasized that because of some restrictions in Vivado/hls4ml on the number of parameters, the input data are images covering just the CE-E part of the HGCAL. The FPGA implementation results for each of the models are reported in Table 16. It is important to bring attention to the fact that, in this case, (hls4ml) models are not optimized in the sense that the precision is adjusted by layers. They use the precision defined at the model level, with settings necessary to preserve the accuracy of the QKeras model. It is clear that all models satisfy the required latency constraint ($< 12\mu s$). FPGA resource usage metrics show flip-flops (FF) employment in all cases is under 2.5%, and lookup tables (LUT) utilization is under 15%. For all models, the number of used digital signal processors (DSP) is 2, which, compared to available DSPs (3072), makes 0.065%.

Table 16: The results of FPGA implementation for different models.

Model	Time [ns]	Interval [Cycles]	Latency [μ s] (Cycles)	FF # (%)	LUT # (%)
QCNN_FP	4.374	30	0.170 (34)	19655 (2.27)	63874 (14.79)
QCNN_U2	4.37	30	0.170 (34)	20594 (2.38)	63261 (14.64)
QMLP_U2E*	4.320	5	0.085 (17)	1134 (0.13)	43401 (10.05)

Using low precision (settings) can help reduce the FPGA resource usage of a model but may result in loss of model performance if chosen inappropriately. In addition to setting the model precision, hls4ml allows per-layer optimization. The profiling tools in hls4ml help to decide the appropriate model/layer precision. Each layer in the QNN is evaluated using the given test data, and the distribution of values is shown with a box and whisker diagram. The grey shaded boxes show the range which can be represented with the used hls4ml configuration settings.

When the test is done with QCNN_UNIQ2, using default hls4ml precision $\langle 16, 6 \rangle$, there was a gap between QKeras (0.968148) and hls4ml (0.9409259) accuracy. The profiling tool has shown that first-layer output precision has to be extended from the default (Fig 14a). With models extended precision $\langle 18, 8 \rangle$, accuracy is improved (0.9679629).




Figure 14: hls4ml profiling, plots show examples of the distributions of the QCNN_UNIQ2 model layers: (a) before optimization, (b) after optimization.

To show the effect of input data quantization together with the possibility of per-layer optimization, the model QCNN_UNIQ2 is implemented in three different ways: 1st using default settings, 2nd optimizing hidden layers precision, and finally also optimizing input layer precision, as is shown in Fig 14b. Afterward, the resource usage for those three cases is compared, which is shown in Table 17. With per-layer optimization for model inner layers, FF # goes down 4.5% (from 20594 to 19670) and LUT # 3% (from 63261 to 61387) Compared with the model where the input layer is also adjusted, saving in FF # is 36.8% and 8.5% in LUT #. The other metrics (Time, Interval, and Latency) are the same in all three cases, corresponding to the default model implementation, shown in Table 16.

Table 17: The results of different types of optimizations for the QCNN-UNIQ2 model.

settings	FF #	FF %	LUT #	LUT %
default	20594	2.38	63261	14.64
hidden layers precision adjusted	19670	2.28	61387	14.21
input + hidden layers precision adjusted	13008	1.51	57884	13.40

5 Conclusion and future work

The paper describes a neural network (NN) model for classifying electromagnetic (EM) showers from the mixed background made of pile-up (PU) and quantum chromodynamics (QCD) jets at level 1 trigger within

a given time limit of a few μs . Recent advances have made it possible to adjust the NN models to limitations caused by the targeted hardware, field programmable gate arrays (FPGAs).

The first tests are done with a data set containing ideal signal (EM) images; in the data simulation step, it was chosen to generate EM showers without PU. The NN classifiers distinguishing between idealistic EM shower and PU achieve 99% accuracy. The EM shower with included PU is generated to get closer to a more realistic situation. As there is a big difference in average EM and PU clusters pt , which simplifies the classification process, another type of particle, QCD jets, is also added to background samples. Simple convolutional neural network and multi-layer perceptron models with signal/background classification accuracy between 96 and 99% and false negative rate (FNR) between 2 and 4%, depending on the background case, are developed. The achieved results could eventually be improved by adding physical features to input data, but it would be difficult to incorporate into level 1 trigger latency requirements. The rest of the paper considers the EM versus mixed background, One-vs-All classification problem. Using the QKeras quantization library, NN models are aggressively quantized, representing the weights and activations with 2 bits, reducing computational and memory costs (in inference time), with accuracy degradation of less than 1%. Further, different approaches to data refinement are presented. The methods that affect the range of input data, like capping, using appropriate values, can significantly lower FPR (for a quantized convolutional neural network with capping value 8, for 34%). Image size reduction methods also decrease the FNR value, lowering the number of NN parameters. The quantization methods, uniform quantization, and selection bit approach achieve performances comparable with a base model, with a large reduction in the precision required for the data storage. As shown, the presented data refinement methods have enabled further reduction of the model resource consumption. Using the hls4ml library, model examples are converted into the firmware to be implemented on FPGA. All implemented models can be executed with a latency lower than 1 μs and flip-flops and lookup tables utilization under 2.5%, respectively 15%. Using variances of implementation of the proposed quantized convolutional neural network, it is shown that model and data quantization enable additional savings of resources when implementing the model in FPGA: the number of used flip-flops is reduced by 36.8% and the number of lookup tables by 8.5%.

Future work could be based on searching for an efficient method for exploiting the sparse nature of calorimeter images.

Author Contributions: Conceptualization, A.B.M and J.M.; methodology, A.B.M., J.M. and M.P.; software, A.B.M.; validation, J.M. and M.P.; formal analysis, D.L.; investigation, A.B.M.; resources, A.B.M., J.M. and M.P.; data curation, D.L. and A.B.M.; writing—original draft preparation, A.B.M.; writing—review and editing, J.M. and M.P.; visualization, A.B.M; supervision, J.M. and M.P.; project administration, J.M., and D.L.. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Acknowledgments: We thank our HGICAL colleagues for their contributions and help with this project, especially Jean-Baptiste Sauvan (Laboratoire Leprince-Ringuet: Palaiseau) and HGICAL Reco+L1 algorithms team for their valuable comments and suggestions during the development of the experiment.

Abbreviations The following abbreviations are used in this manuscript:

3D	Three-Dimensional
ACC	Accuracy
AUC	Area Under the Curve
BDT	Boosted Decision Tree
CE-E	Electromagnetic Section
CE-H	Hadronic Section
CMS	Compact Muon Solenoid
CMSSW	CMS Software Components
CNN	Convolutional Neural Network
EM	Electromagnetic
FE	Feature Extraction
FN	False Negative
FNR	False Negative Rate
FP	False Positive
FPGA	Field Programmable Gate Array
FPR	False Positive Rate
FS	Feature Selection
HDL	Hardware Description Language
HEP	High Energy Physics
HGCAL	High Granularity Calorimeter
HL-LHC	High Luminosity LHC
HLS	High-Level Synthesis
LHC	Large Hadron Collider
L1T	Level 1 Trigger
PU	Pile Up
ML	Machine Learning
MLP	Multilayer Perceptron
NN	Neural Network
QCD	Quantum Chromodynamics
RF	Random Forrest
ROC	Receiver Operating Characteristic
ROI	Region of Interest
TC	Trigger Cell
TN	True Negative
TNR	True Negative Rate
TP	True Positive
TPR	True Positive Rate

References

- [1] G. Apollinari, Béjar Alonso I., O. Brüning, P. Fessia, M. Lamont, L. Rossi, and L. Tavian, editors. *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*. CERN Yellow Reports: Monographs. CERN, Geneva, 2020.
- [2] CMS Collaboration. The CMS experiment at the CERN LHC. *Journal of Instrumentation*, 3(08):S08004, 2008.
- [3] P. Baldi, P. Sadowski, and D. Whiteson. Searching for exotic particles in high-energy physics with deep learning. *Nature communications*, 5:4308, 07 2014.
- [4] P. Paulitsch. The silicon sensors for the high granularity calorimeter of CMS. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 978:164428, 07 2020.
- [5] CMS Collaboration. The Phase-2 Upgrade of the CMS Endcap Calorimeter. Technical report, CERN, Geneva, 2017.
- [6] M Noy. The CMS HGCAL silicon region architecture specification and optimisation. *JINST*, 17(03):C03010, 2022.
- [7] H. Gerwig. Engineering challenges in mechanics and electronics in the world’s first particle-flow calorimeter at a hadron collider: The CMS high-granularity calorimeter. *Nuclear Instruments and*

- [8] S. Summers et al. Fast inference of boosted decision trees in FPGAs for particle physics. *Journal of Instrumentation*, 15:P05026–P05026, 05 2020.
- [9] D. Acosta, A. Brinkerhoff, E. Busch, A. Carnes, I. Furic, S. Gleyzer, K. Kotov, J. Low, A. Madorsky, J. Rorie, B. Scurlock, and W. Shi. Boosted decision trees in the level-1 muon endcap trigger at CMS. *Journal of Physics: Conference Series*, 1085:042042, 09 2018.
- [10] Y. Iiyama et al. Distance-weighted graph neural networks on FPGAs for real-time particle reconstruction in high energy physics. *Frontiers in Big Data*, 3, 01 2021.
- [11] J. Pata, J. Duarte, F. Mokhtar, E. Wulff, J. Yoo, J. Vlimant, M. Pierini, and M. on behalf of the CMS Collaboration Girone. Machine learning for particle flow reconstruction at CMS. *Journal of Physics: Conference Series*, 2438(1):012100, feb 2023.
- [12] M. Touranakou, N. Chernyavskaya, J. Duarte, D. Gunopulos, R. Kansal, B. Orzari, M. Pierini, T. Tomei, and J. Vlimant. Particle-based fast jet simulation at the LHC with variational autoencoders. *Machine Learning: Science and Technology*, 3.
- [13] M. Andrews, M. Paulini, S. Gleyzer, and B. Poczos. Exploring end-to-end deep learning applications for event classification at CMS. *EPJ Web of Conferences*, 214:06031.
- [14] A. Aurisano, A. Radovic, D. Rocco, A. Himmel, M. D. Messier, E. Niner, G. Pawloski, F. Psihas, A. Sousa, and P. Vahle. A convolutional neural network neutrino event classifier. *Journal of Instrumentation*, 11(9):P09001–P09001.
- [15] J.H. Collins, K. Howe, and B. Nachman. Anomaly detection for resonant new physics with machine learning. *Physical Review Letters*, 121(24):241803.
- [16] A. Pol, V. Azzolini, G. Cerminara, F. Guio, G. Franzoni, M. Pierini, F. Siroky, and J. Vlimant. Anomaly detection using deep autoencoders for the assessment of the quality of the data acquired by the CMS experiment. *EPJ Web of Conferences*, 214:06008.
- [17] E. Shumka et al. Machine learning based tool for CMS RPC currents quality monitoring. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1054:168449, September 2023.
- [18] CMS Collaboration. A deep neural network to search for new long-lived particles decaying to jets. *Machine Learning: Science and Technology*, 1(3):035012.
- [19] S.R. Qasim, K. Long, J. Kieseler, M. Pierini, and R. Nawaz. Multi-particle reconstruction in the high granularity calorimeter using object condensation and graph neural networks. *EPJ Web of Conferences*, 251:03072.
- [20] A. Elabd, V. Razavimaleki, S.Y. Huang, J. Duarte, M. Atkinson, G. DeZoort, P. Elmer, S. Hauck, J.X. Hu, S.C. Hsu, B. Lai, M. Neubauer, I. Ojalvo, S. Thais, and M. Trahms. Graph neural networks for charged particle tracking on FPGAs. *Frontiers in Big Data*, 5:828666.
- [21] M. Paganini, L. de Oliveira, and B. Nachman. CaloGAN: Simulating 3d high energy particle showers in multi-layer electromagnetic calorimeters with generative adversarial networks. *Physical Review D*, 97(1):014021.
- [22] P. Musella and F. Pandolfi. Fast and accurate simulation of particle detectors using generative adversarial networks. *Computing and Software for Big Science*, 2(1):8.
- [23] D. Belayneh, F. Carminati, A. Farbin, B. Hooberman, G. Khattak, M. Liu, J. Liu, D. Olivito, V. Barin Pacela, M. Pierini, A. Schwing, M. Spiropulu, S. Vallecorsa, J.R. Vlimant, W. Wei, and M. Zhang. Calorimetry with deep learning: particle simulation and reconstruction for collider physics. *European Physical Journal C*, 80(7), 2020.
- [24] S. Gronroos, M. Pierini, and N. Chernyavskaya. Automated visual inspection of CMS HGCAL silicon sensor surface using an ensemble of a deep convolutional autoencoder and classifier.
- [25] CMS Collaboration. The analytical method algorithm for trigger primitives generation at the LHC drift tubes detector. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1049:168103.

- [26] D. Guest, K. Cranmer, and D. Whiteson. Deep learning and its application to LHC physics. *Annual Review of Nuclear and Particle Science*, 68(1):161–181.
- [27] J. Cogan, M. Kagan, E. Strauss, and A. Schwartzman. Jet-images: Computer vision inspired techniques for jet tagging. *Journal of High Energy Physics*, 2015, 07 2014.
- [28] M. Andrews, M. Paulini, S. Gleyzer, and B. Poczos. End-to-end physics event classification with CMS open data: Applying image-based deep learning to detector data for the direct classification of collision events at the LHC. *Computing and Software for Big Science*, 4(1):6.
- [29] M. Andrews, B. Burkle, Y. Chen, D. DiCroce, S. Gleyzer, U. Heintz, M. Narain, M. Paulini, N. Pervan, Y. Shafi, W. Sun, E. Usai, and K. Yang. End-to-end jet classification of boosted top quarks with the CMS open data. *Phys. Rev. D*, 105:052008, Mar 2022.
- [30] T. Aarrestad et al. Fast convolutional neural networks on FPGAs with hls4ml. *Machine Learning: Science and Technology*, 2(4):045015, 2021.
- [31] J. Alimena, Y. Iiyama, and J. Kieseler. Fast convolutional neural networks for identifying long-lived particles in a high-granularity calorimeter. *Journal of Instrumentation*, 15(12):P12006, 2020.
- [32] B. Denby, M. Campbell, F. Bedeschi, N. Chriss, C. Bowers, and F. Nesti. Neural networks for triggering. *IEEE Transactions on Nuclear Science*, 37(2):248–254, 1990.
- [33] A. Di Florio, F. Pantaleo, and A. Carta. Convolutional Neural Network for Track Seed Filtering at the CMS High-Level Trigger. *J. Phys.: Conf. Ser.*, 1085(4):042040, 2018.
- [34] A.M. Sirunyan et al. Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 tev. *Journal of Instrumentation*, 13:P05011–P05011, 05 2018.
- [35] M. Courbariaux, Y. Bengio, and J.P. David. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2, NIPS’15*, page 3123–3131, Cambridge, MA, USA, 2015. MIT Press.
- [36] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi. XNOR-net: ImageNet classification using binary convolutional neural networks, 10 2016.
- [37] Google. QKeras: a quantization deep learning library for tensorflow keras (software), 2019.
- [38] Plumerai. Larq: open-source deep learning library for training neural networks with extremely low precision weights and activations (software), 2023.
- [39] K. Pedro. Current and future performance of the CMS simulation. *EPJ Web of Conferences*, 214:02036, 01 2019.
- [40] Geant4 Collaboration. Geant4: Geant4 is a toolkit to create simulations of the passage of particles or radiation through matter. (software), 2006.
- [41] M. Prvan. *Algorithms for the Level-1 trigger with the HGAL calorimeter for the CMS HL-LHC upgrade*. PhD thesis, University of Split. Faculty of Electrical Engineering, Mechanical . . . , 2020.
- [42] M. Maleki Tehrani, N. Madani, and Xa. Emery. Capping and kriging grades with longtailed distributions. *Journal of the Southern African Institute of Mining and Metallurgy*, 114, 03 2014.
- [43] L.B.V. de Amorim, G.D.C. Cavalcanti, and R.M.O. Cruz. The choice of scaling technique matters for classification performance. *Applied Soft Computing*, 133:109924, 2023.
- [44] G. Kumar and P.K. Bhatia. A detailed review of feature extraction in image processing systems. In *2014 Fourth International Conference on Advanced Computing & Communication Technologies*, pages 5–12, 2014.
- [45] M. Andrews, J. Alison, S. An, B. Burkle, S. Gleyzer, M. Narain, M. Paulini, B. Poczos, and E. Usai. End-to-end jet classification of quarks and gluons with the CMS open data. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 977:164304, 2020.
- [46] G. Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems (MCSS)*, 2(4):303–314, 1989.

- [47] Hornik. K., Stinchcombe M.B., and White. H.L. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.
- [48] A.B. Mišura, J. Musić, J. Ožgović, and D. Lelas. Performance comparison of generic and quantized fully connected and convolutional neural networks for real- time signal/background classification. In *2022 International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, pages 1–6, 2022.
- [49] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '19, page 2623–2631, New York, NY, USA, 2019. Association for Computing Machinery.
- [50] A. Gholami, S. Kim, Z. Dong, Z.i Yao, M. Mahoney, and K. Keutzer. *A Survey of Quantization Methods for Efficient Neural Network Inference*, pages 291–326. 01 2022.
- [51] G. Shomron, F. Gabbay, S. Kurzum, and U. Weiser. Post-training sparsity-aware quantization, 2021.
- [52] Y. Bengio, N. Léonard, and A.C. Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *CoRR*, abs/1308.3432, 2013.
- [53] C. N. Coelho, A. Kuusela, S. Li, H. Zhuang, T. Aarrestad, V. Loncar, J. Ngadiuba, M. Pierini, A. Pol, and S. Summers. Automatic heterogeneous quantization of deep neural networks for low-latency inference on the edge for particle detectors. *Nature Mach. Intell.*, 3:675–686, 2021.
- [54] C. Guo, G. Pleiss, Y. Sun, and K.Q. Weinberger. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17, page 1321–1330. JMLR.org, 2017.
- [55] C. Sammut and G.I. Webb. *Encyclopedia of Machine Learning*. Springer Publishing Company, Incorporated, 1st edition, 2011.
- [56] V. Sheng and C. Ling. Thresholding for making classifiers cost sensitive. In *AAAI Conference on Artificial Intelligence*, volume 1, 01 2006.
- [57] W. J. Youden. Index for rating diagnostic tests. *Cancer*, 3(1):32–35, 1950.
- [58] T. Trzciński, L. Graczykowski, and M. Glinka. Using random forest classifier for particle identification in the alice experiment. In Piotr Kulczycki, Janusz Kacprzyk, Łászló T. Kóczy, Radko Mesiar, and Rafal Wisniewski, editors, *Information Technology, Systems Research, and Computational Physics*, pages 3–17, Cham, 2020. Springer International Publishing.
- [59] S. Kuzu. Random forest based multiclass classification approach for highly skewed particle data. *Journal of Scientific Computing*, 95, 02 2023.
- [60] L. Breiman. Random forests. *Machine Learning*, 45:5–32, 10 2001.
- [61] A. Cutler, D. Cutler, and J. Stevens. Random forests. *Machine Learning - ML*, 45:157–176, 01 2011.
- [62] F. Pedregosa et al. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [63] UltraScale and UltraScale+ FPGAs Packaging and Pinouts product specification, 2022.