Available on CMS information server

CMS DP -2023/086

The Compact Muon Solenoid Experiment

CMS Performance Note

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

05 November 2023

Level-1 Trigger Calorimeter Image Convolutional Anomaly Detection Algorithm

CMS Collaboration

Abstract

This performance note shows the implementation of the Calorimeter Image Convolutional Anomaly Detection Algorithm (CICADA), a Level-1 triggering algorithm designed to use fast, unsupervised machine learning techniques to trigger on anomalous events. Some of the score and rate performance of the algorithm as emulated on 2023 data are shown.

CICADA

CMS Collaboration,
cms-dpg-conveners-l1t@cern.ch

Introduction

- CICADA (**C**alorimeter **I**mage **C**onvolutional **A**nomaly **D**etection **A**lgorithm)
  - Fully autonomous ML algorithm for the selection of anomalous events in the Calorimeter Trigger in real-time
  - Utilizes a convolutional neural network auto-encoder to independently pick out topologies that are different from the vast majority of events seen at the LHC
  - Operates on 18×14 (iPhi×iEta) raw calorimeter energy deposits
  - The Mean Squared Error $\overline{(X_{Predicted} - X_i\ )^2}$ is used to measure the accuracy of the encoding/decoding reconstruction process
- The algorithm can run in hundreds of nanoseconds on a single Virtex-7 Field Programmable Gate Array (FPGA).
  - **Quantization** involves assigning the model's weights and activations to a lower bit-width representation.
  - **Pruning** deals with the problem of the optimal removal of parameters deemed unimportant via some heuristic.
  - **Knowledge distillation** is a method of training a smaller network, called the student, using soft labels generated by a larger network, called the teacher.
- The model is generated using the HLS4ML (High Level Synthesis for Machine Learning) package
- CICADA is emulated in a bit accurate fashion via compiling its model firmware for use inside CMS software.

Anomaly Detection with Auto-Encoders

CICADA is a 2D convolutional neural network designed to find events that show differences from the standard beam event. As input, CICADA uses the calorimeter region energies, represented by an order 2 tensor of 252 (18 by 14) transverse energy deposits in the CMS calorimeters. To find anomalous events in the detector, CICADA's neural network is trained to encode zero bias beam events into a space smaller than its input and then reconstruct the original tensor from this smaller representation (also called the "latent space"). This smaller latent space forces generalization to its training data, zero bias events. The quality of the overall reconstruction is measured with the Mean Squared Error (MSE) loss, which is the average of the squared error in reconstruction for each of the 252 individual energy deposits. The MSE is calculated as:

$$\mathrm{MSE} = \overline{(X_{Predicted} - X_{True})^2}$$

The crux of this approach is that the training is done on ~10^6 zero-bias events (which are primarily simple quark-gluon interactions) and this means it does not generalize well on rare events with exotic signatures. The MSE (also referred to as "anomaly score" for this purpose) will be higher for these rare signatures.

Reconstruction Qualities

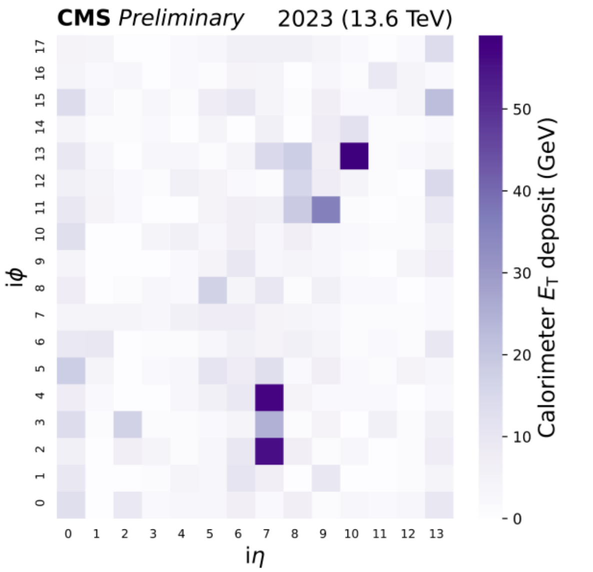

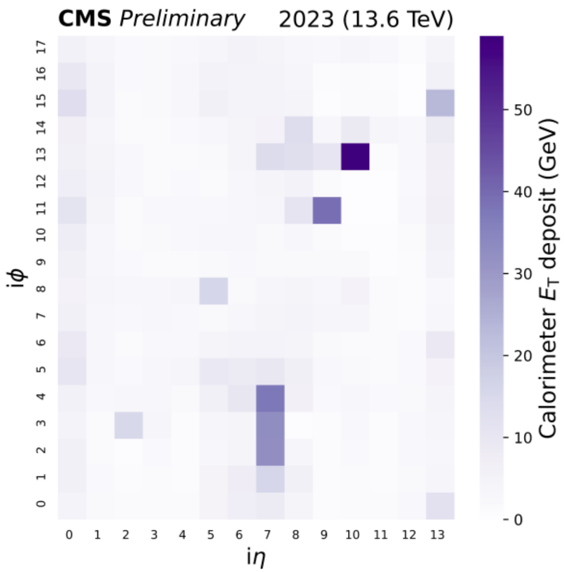

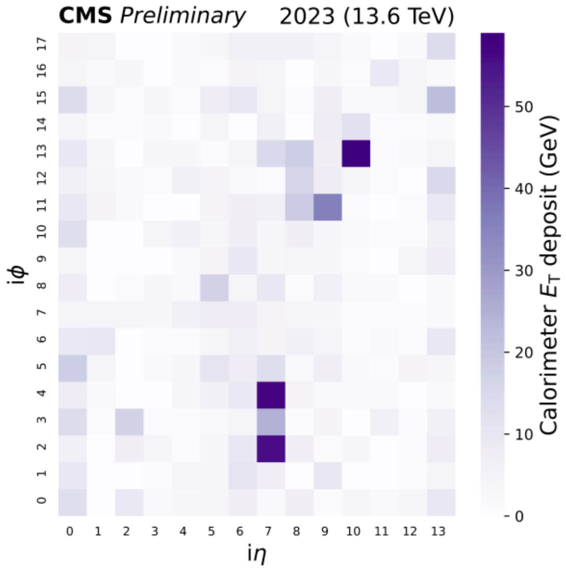

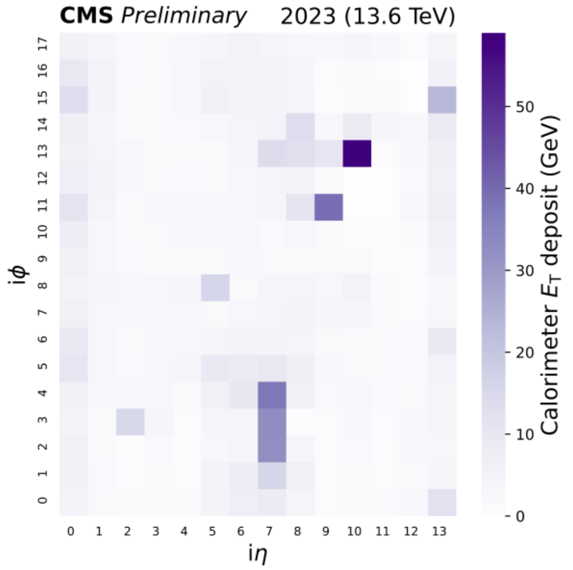

Shown here is a comparison of the teacher model ability to reconstruct a Zero Bias (ZB) beam event (original: far left, reconstructed: center left) versus a signal sample, Soft Unclustered Energy Patterns (SUEP) on the right (original: center right, reconstructed: far right). In general, the teacher model is better able to reconstruct the Zero Bias beam event as evidenced by a far lower loss (0.81) compared to the SUEP loss (14.21). This example shows how the CICADA anomaly detection mechanism works to find anomalies.

CICADA Score Distributions

CICADA has an emulator compiled from its firmware which can be used to test the firmware model on genuine detector data accurately. All output from the emulator is bit accurate to the output of the firmware

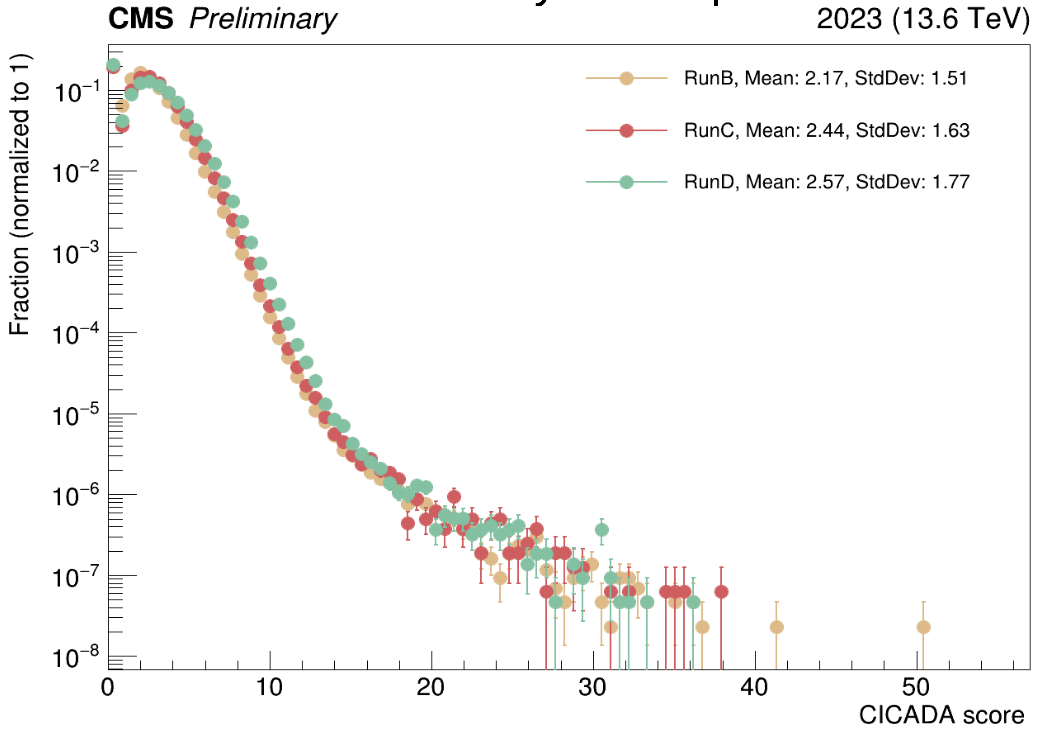

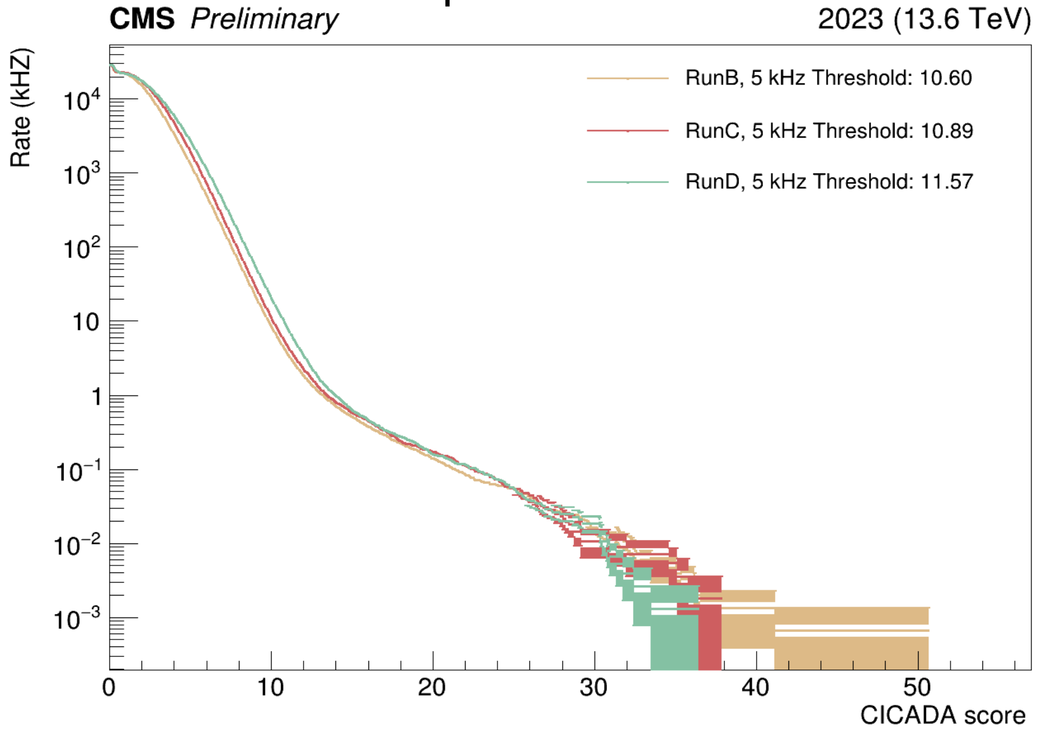

The plot on the left shows the CICADA score on different 2023 data taking periods (for events CICADA was not trained on) for zero-bias events. Run B corresponds to data taken between April-May 2023, Run C corresponds to data taken between May-June 2023, and Run D corresponds to data taken between June-August 2023. The plot on the right calculates a rate from the efficiencies obtained using these scores as a threshold. The score is stable between periods, and will only need re-training when there are changes in the detector conditions.

CICADA Pure Vs Overall Rate

On the right is a plot showing the correspondence of the overall CICADA rate versus the CICADA rate without any other un-prescaled trigger paths firing (the “pure” rate)

A useful approximation is that 5 kHz overall rate is ~3 kHz pure rate.

Worth noting is that as the overall rate is reduced in CICADA, the pure rate is reduced much quicker, i.e. the relative fraction of impurity increases with overall rate reduction.

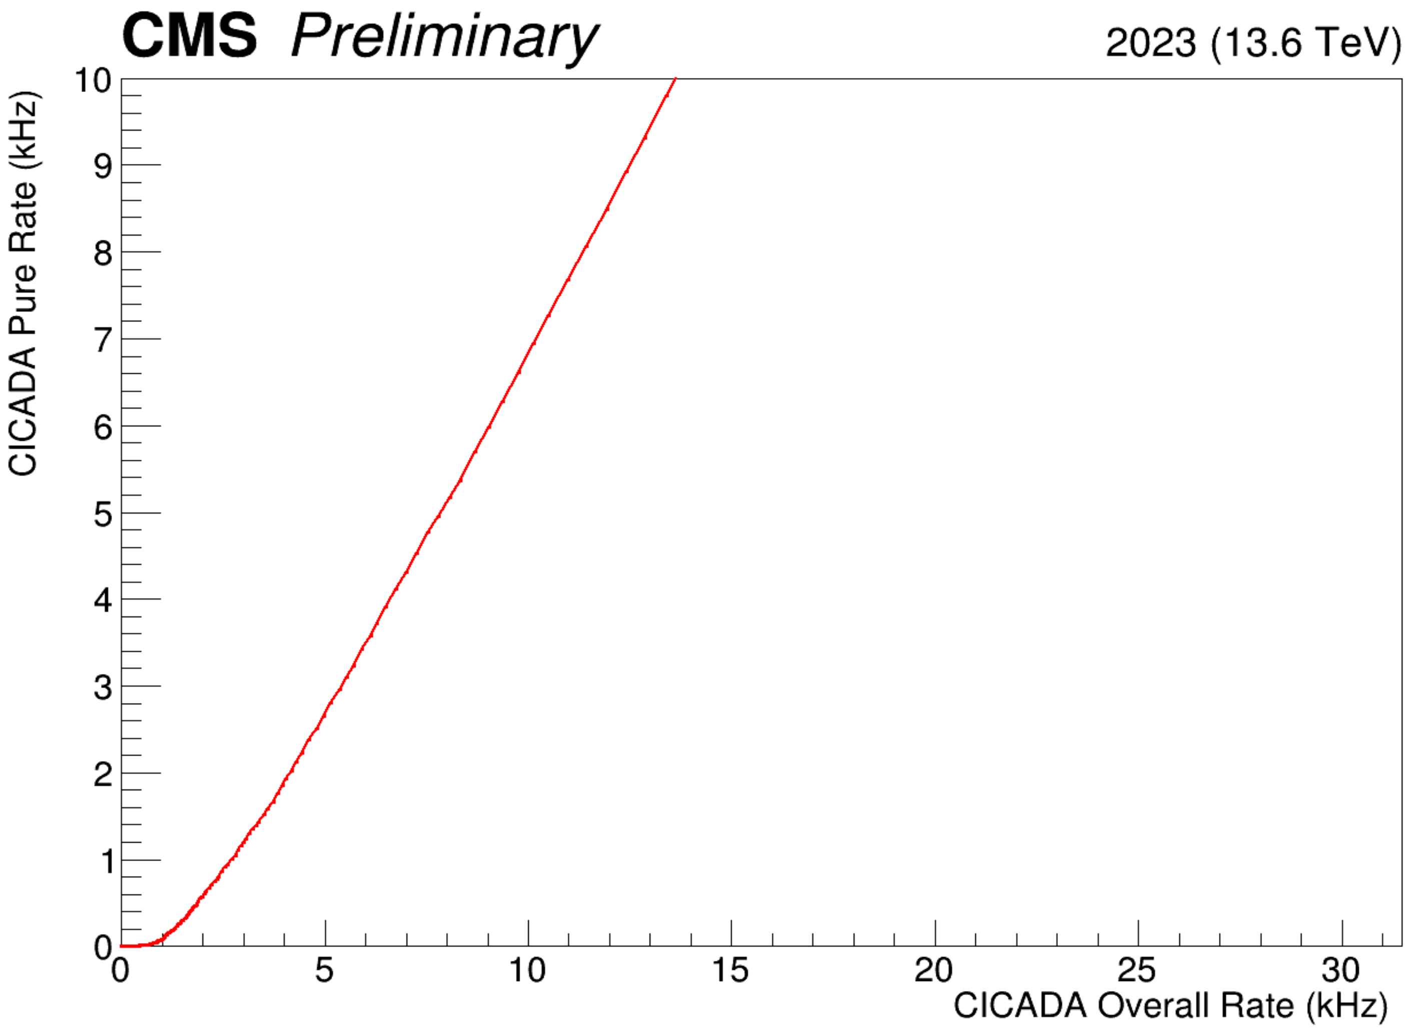

Emulated CICADA Run Performance

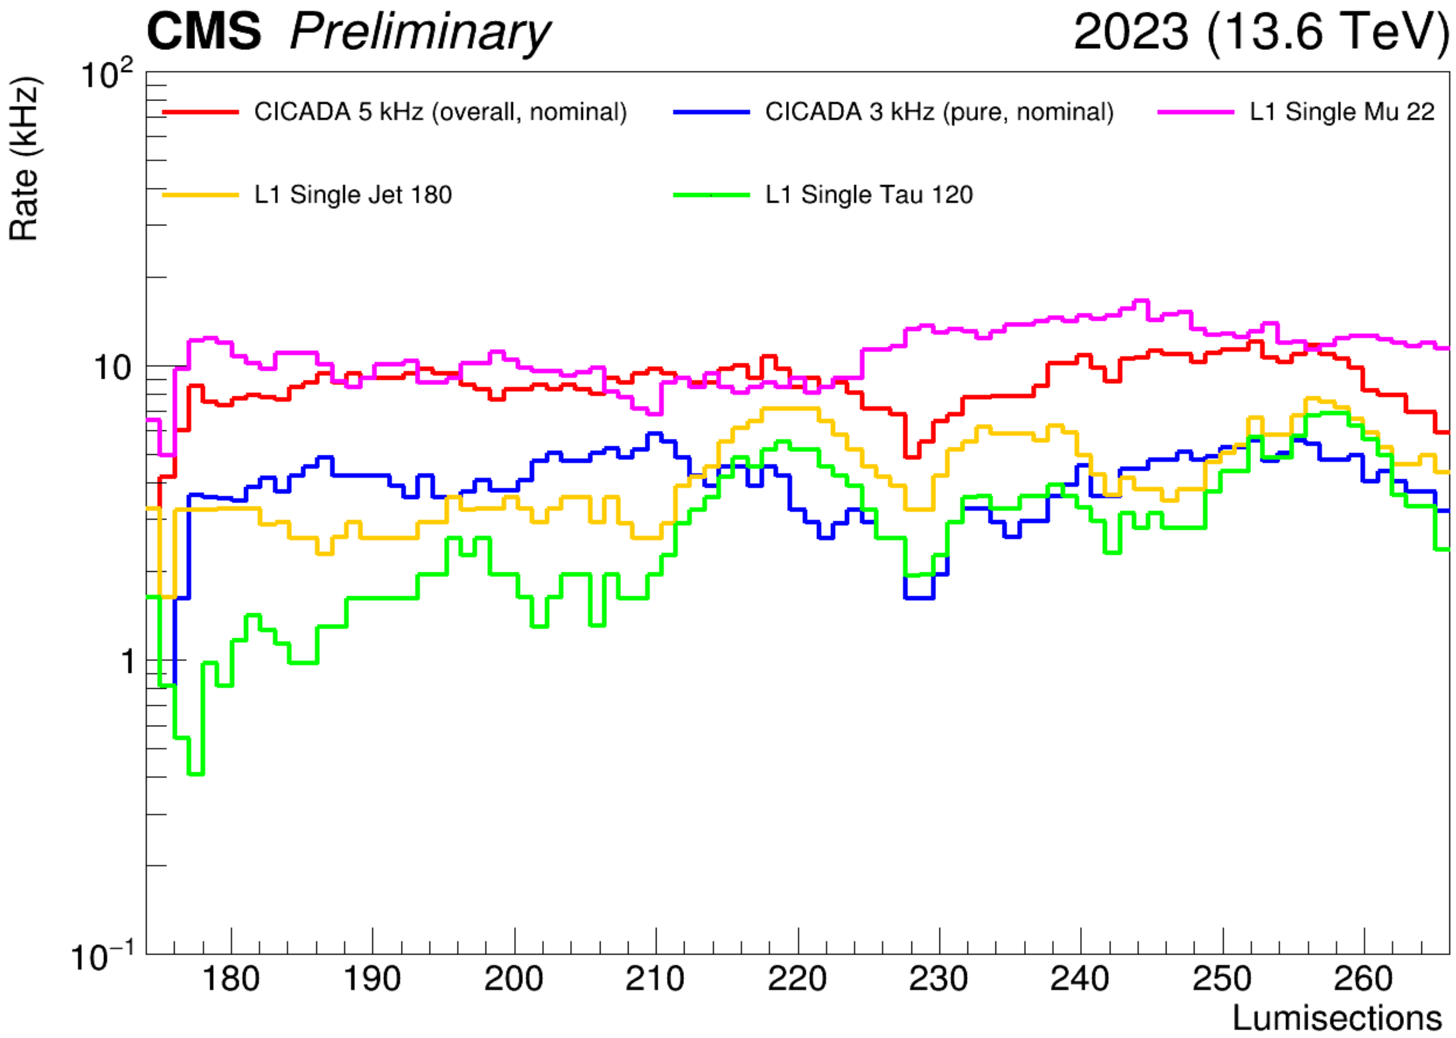

The accuracy of the emulator can also be used to demonstrate CICADA's performance inside of a run. The plot on the left shows CICADA rate for a 5 kHz overall rate threshold (red), and 3 kHz pure trigger (blue), compared against common unprescaled bits like SingleMu22 (pink, corresponds to triggering on a muon with transverse momentum of 22 GeV), SingleJet180 (yellow, corresponds to triggering on a jet of transverse energy 180 GeV) and SingleTau 120 (green, corresponds to triggering on a tau of transverse momentum 120 GeV). CICADA's performance is very similar to these simple benchmark unprescaled bits.

The data shown here was taken in July 2023.