

Autoencoder-based Online Data Quality Monitoring for the CMS Electromagnetic Calorimeter

Abhirami Harilal, Kyungmin Park, Michael Andrews
and Manfred Paulini on behalf of the CMS Collaboration

Department of Physics, Carnegie Mellon University, Pittsburgh, PA 15213, U.S.A.

E-mail: aharilal@cern.ch

Abstract. The online Data Quality Monitoring system (DQM) of the CMS electromagnetic calorimeter (ECAL) is a crucial operational tool that allows ECAL experts to quickly identify, localize, and diagnose a broad range of detector issues that would otherwise hinder physics-quality data taking. Although the existing ECAL DQM system has been continuously updated to respond to new problems, it remains one step behind newer and unforeseen issues. Using unsupervised deep learning, a real-time autoencoder-based anomaly detection system is developed that is able to detect ECAL anomalies unseen in past data. After accounting for spatial variations in the response of the ECAL and the temporal evolution of anomalies, the new system is able to efficiently detect anomalies while maintaining an estimated false discovery rate between 10^{-2} to 10^{-4} , beating existing benchmarks by about two orders of magnitude. The real-world performance of the system is validated using anomalies found in 2018 and 2022 LHC collision data. Additionally, first results from deploying the autoencoder-based system in the CMS online DQM workflow for the ECAL barrel during Run 3 of the LHC are presented, showing its promising performance in detecting obscure issues that could have been missed in the existing DQM system.

1. Introduction

In a large-scale high energy physics experiment like CMS [1], ensuring timely detection of issues that could affect the detector performance and the quality of data taken is a major task that requires significant manpower and time. Presently, the CMS online DQM [2] consists of a set of histograms that are filled based on a first-pass analysis of a subset of data collected by the detector. These histograms are monitored continuously by a DQM *shifter* who reports on any apparent irregularities observed. Depending on the severity of the problem, various mitigation measures up to stopping of the data taking would be performed by the relevant experts. While this system has proven to be dependable, the changing running conditions and increasing collision rates, along with aging electronics, bring forth failure modes that are newer and harder to predict. Machine learning (ML) based approaches to anomaly detection in DQM have been adopted by previous efforts in CMS [3, 4]. In this paper, an unsupervised method of anomaly detection for the online DQM of the CMS ECAL is presented utilizing an autoencoder (AE) [5] on ECAL data processed as two-dimensional (2D) images. In a novel approach, correction strategies are implemented to account for spatial differences in the ECAL response as well as the time-dependent nature of anomalies in the detector. This system is deployed in the online DQM for the ECAL barrel during LHC Run 3 collisions, complementing

the existing DQM plots. First results from the AE-based anomaly detection system are reported indicating it to be a highly valuable diagnostic tool for ECAL experts involved in real-time data taking operations. All the figures shown here are from the approved CMS public results in Refs. [6, 7].

2. The CMS Electromagnetic Calorimeter and Data Quality Monitoring

The CMS ECAL [8] is a hermetic calorimeter which measures the energy, time, and position of photons, electrons and electromagnetic fraction of jets. It played a crucial role in the discovery of the Higgs Boson [9] as well as in the measurement of the Higgs properties. ECAL is made up of scintillating lead tungstate crystals arranged in a cylindrical central barrel (EB) section closed by two endcaps (EE+ and EE-).

The DQM system in ECAL mainly consists of two kinds of histograms: occupancy-style histograms as shown in Fig. 1(a), filled with real-time data of vital quantities, and quality-style histograms as shown in Fig. 1(b), which are drawn based on thresholds and rules applied on the quantity in the occupancy-style histograms. Quality-style histograms are in color-coded maps that are easy-to-interpret, so that it is possible to tell at a glance if something is wrong in the ECAL. The color code used is as follows: green for *good*, red for *bad*, brown for a *known problem*, and yellow for *no data*, which may or may not be an issue depending on the context. The histograms are often plotted at the granularity of a *trigger tower* (TT), which is defined as a set of 5x5 crystals. Each single square in Fig. 1 represents a TT, and 68 such TTs form a *supermodule* (SM) in the barrel, corresponding to the numbered rectangles in Fig. 1. In online DQM, these histograms are accumulated over a CMS data acquisition *run*, with each histogram plotted at every *lumisection* (LS), that corresponds to a time interval of about 23 seconds, over which the luminosity is considered to remain approximately constant.

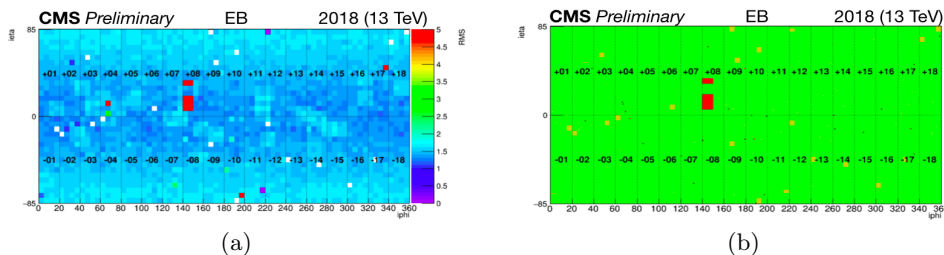


Figure 1: DQM plots for ECAL barrel: (a) occupancy-style and (b) quality-style histogram.

With increasing luminosity and harsher radiation environment, it becomes impossible to anticipate all failure modes in a complex detector like ECAL. Though there are multiple alarms set in the current DQM framework to catch various types of errors, they are prone to high false positives and are limited by having to define hard-coded rules for every possible detector geometry. To mitigate these issues, an automated approach to anomaly detection is explored using unsupervised ML.

3. Autoencoder-based Anomaly Detection System

3.1. Network and Anomaly Detection Strategy

An unsupervised anomaly detection method is developed using an AE trained on occupancy-style histograms from the ECAL DQM processed as 2D images. The AE built with a ResNet [10] architecture using PYTORCH [11] is trained with manually certified *good* data. The encoder network of the AE encodes the input image into a lower dimensional latent space, and the decoder network tries to reconstruct the original image from the encoded space. The goodness

of reconstruction is measured by the reconstruction loss \mathcal{L} , computed as the Squared Error between the input (x) and the AE-reconstructed output (x') as defined in Eq. 1:

$$\mathcal{L}(x, x') = |(x - x')|^2 \quad (1)$$

The network trained on good data reconstructs the nominal detector image well with minimal reconstruction loss. When fed with anomalous data, however, it fails to reconstruct the anomalies and gives a higher loss. The squared error on each tower is calculated and plotted as a 2D loss map, on which some post-processing steps are applied as explained in the next sections. A threshold is then derived from the anomalous loss values that can efficiently catch 99% of the anomalies.

3.2. Dataset, Training and Validation

Each input image to the AE is the digitized hit occupancy map from a single LS. The dataset used for training and validation is taken from the 2018 runs during LHC Run 2. It consists of 100 000 2D occupancy images, with training and validation dataset split in 9-to-1 ratio. In order to make the quality interpretation consistent across different run conditions, occupancy maps are normalized with respect to *pileup* (PU), which are additional proton-proton interactions within the same proton bunch crossing. After pre-processing with the PU correction, AE models are trained and validated separately for the barrel and each endcap.

In addition to the nominal validation, fake anomaly validation is performed, using the good images from the nominal validation set but with synthetic anomalies introduced in random towers. Three types of anomaly scenarios are studied:

- (i) Missing SM/sector: Entire SMs for the barrel and sectors for endcaps are randomly set to have zero occupancy values in each LS.
- (ii) Single zero occupancy tower: A single tower is set to have zero occupancy at random in each LS.
- (iii) Single hot tower: A single tower is set to be *hot*, or having higher-than-nominal occupancy.

3.3. Spatial Response Correction

In the ECAL, the occupancy close to the beam pipe tends to be higher. This is clearly visible in the average occupancy map shown in Fig. 2(a). This effect is reflected in the loss map shown in Fig. 2(b) in the case of a missing SM, where the towers in the higher $|\eta|$ region have higher loss than those in the lower $|\eta|$ region, for an anomaly affecting all the towers in the SM equally. In order to interpret the towers as equally anomalous across the SM, the loss map is normalized by the average occupancy map, as a spatial response correction, to obtain a uniform loss map as illustrated in Fig. 2(c).

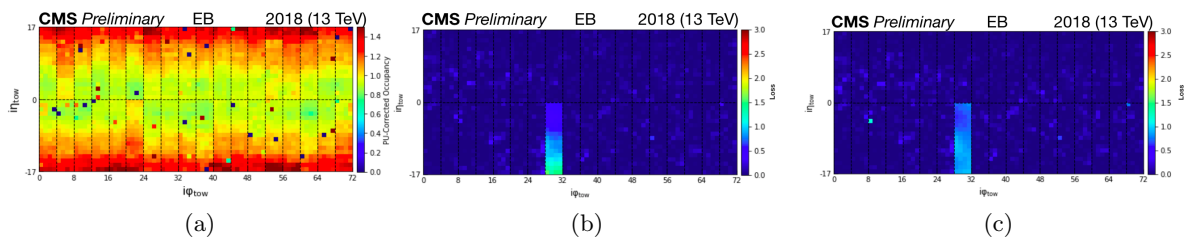


Figure 2: (a) Average occupancy map of EB. Loss maps for missing SM scenario (b) before and (c) after spatial response correction.

3.4. Time correction

A correction to further reduce false positives is implemented by utilizing the time-dependent nature of the anomalies in the detector. Unlike random fluctuations, real anomalies would persist with time. Figure 3 describes the time correction strategy, where loss maps that have been corrected for spatial effects from three consecutive LSs are multiplied with one another at a tower level to boost the anomaly and suppress the fluctuations.

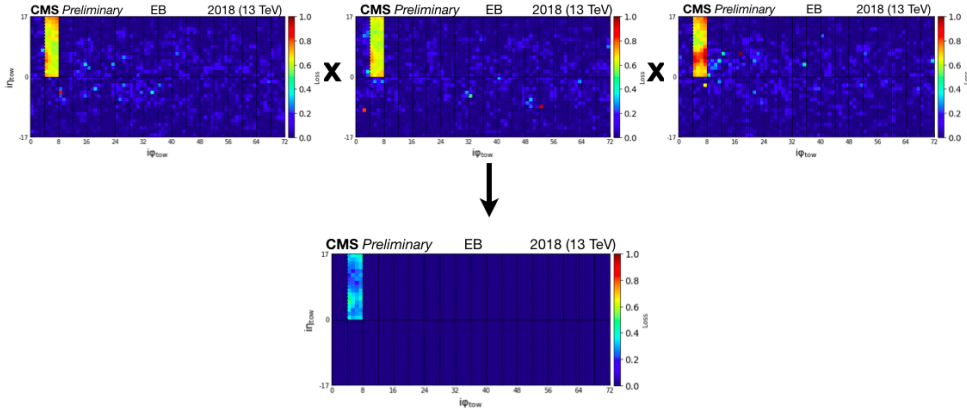


Figure 3: Time correction strategy: loss maps from three consecutive LSs (top panel) are multiplied at a tower level to make a time-corrected loss map (bottom).

4. Results

An anomaly tagging threshold is chosen based on the validation set with fake anomalies, such that 99% of the anomalies are detected. The performance of the AE is measured using the metric False Discovery Rate (FDR) defined as:

$$\text{FDR} = \frac{\text{Number of good towers above the anomaly threshold}}{\text{Number of good and bad towers above the anomaly threshold}} \quad (2)$$

The lower the FDR, lesser the false detection and better the performance of the AE.

4.1. Testing on Fake Anomalies

Table 1 summarizes the FDR values from each anomaly scenario. For the barrel and the endcaps, single zero occupancy towers are most challenging to detect. Spatial correction improves the AE performance, with a much greater effect for the endcaps. This is due to the presence of a larger effective gradient in the occupancy values across the endcap region compared to the barrel. Further improvement in the FDRs by an order of magnitude is achieved with the time correction. Final performance scores after applying all post-processing are around 5% for the zero occupancy tower scenario and sub-percent for the missing SM (for the barrel) and hot tower scenarios. Remaining false positives have likely come from real anomalous towers that have fallen into the manually certified good dataset. Based on the performance for each scenario, a single anomaly tagging threshold is chosen for each sub-detector part, such that it can efficiently tag all anomalies.

4.2. Testing on Real Anomalies

The AE model is further tested using real anomalous data from 2018 and 2022 LHC runs as illustrated in Fig. 4. Figures 4(a) and (b) respectively show occupancy plots from Run 2 with

Table 1: Summary of FDR at 99% anomaly detection in fake anomaly validation.

Scenario	Missing SM	Zero Occupancy Tower			Hot Tower		
	Barrel	Barrel	EE+	EE-	Barrel	EE+	EE-
AE no correction	3.6%	51%	86%	87%	2.8%	0.01%	0.00%
AE after spatial correction	3.1%	49%	13%	14%	2.9%	0.06%	0.05%
AE after spatial and time correction	0.13%	4.1%	5.6%	6.3%	0.00%	0.00%	0.00%

a missing SM EB-03, and with a ring of hot towers with a zero occupancy tower in the center in EB-01 and EB-18. Their corresponding final quality plots from the AE loss map identify the bad towers in red in Figures 4(d) and 4(e) respectively. For the endcaps, a real anomaly case from Run3 is illustrated in Fig. 4(c), with two zero occupancy towers and a hot tower around the edge. The AE quality plot in Fig. 4(f) spots the two zero occupancy towers correctly. The hot tower does not show up in the quality plot, as it is previously identified as problematic and therefore masked in the DQM system.

The performance of the AE on real-life anomalies demonstrates its ability to detect various kinds of anomalies using a single threshold, without the need for hard-coded rules based on the type of anomaly or the geometry, and emphasizes the power of unsupervised ML as an efficient, adaptable anomaly detection tool.

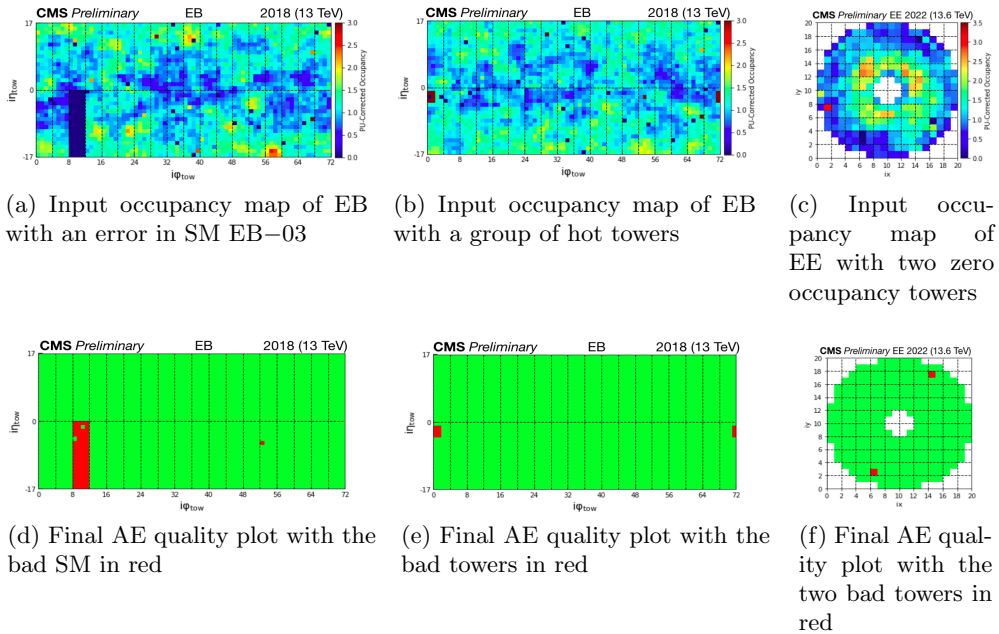


Figure 4: Input occupancy images with real anomalies (top) and corresponding AE quality plots (bottom) for the barrel and endcaps.

5. Deployment in Online DQM for LHC Run3

The AE-based anomaly detection system, named MLDQM, has been deployed in the ECAL online DQM workflow for the ECAL barrel since the start of Run 3 of the LHC. This introduces

a new ML quality plot from the AE in the ECAL DQM workspace (see Fig. 5(a)) complementing the existing plots. The model inference is done by exporting the trained PYTORCH model to ONNX [12] and run in CMS production using ONNX RUNTIME [13]. The deployment of the models for the endcaps is underway and will be ready for the 2023 runs.

It has been observed that the MLDQM is able to correctly identify consistently bad towers, as well as transient bad towers, which may point to deteriorating channels in the ECAL that are not easily detected by the existing DQM plots. Figure 5(a) shows the ML quality plot with two towers in EB+08 marked red: Tower 1 closest to $i\eta = 0$ and Tower 2 closest to $i\eta = 85$. Figure 5(b) shows the total digitized hit occupancy map accumulated over all LSs of the run in the online DQM. Here, Tower 1 is visible with very low occupancy compared to other towers, indicating that it is a persistent zero occupancy tower, while Tower 2 is faintly visible, indicating that it likely had zero occupancy in several LSs but not in all, thus pointing to a transient anomaly. This feature also shows up in the occupancy map in Fig. 5(c), produced offline by averaging over several runs in Run 3. The low occupancy of Tower 2 in this plot implies that the tower indeed had zero occupancy transiently for many LSs in these runs, and that it is likely not a stand-alone random occurrence. This tower could be pointing to channels on the way to becoming permanently faulty. This feature of the MLDQM which detects potentially deteriorating channels would be immensely helpful to detector experts monitoring the health of the detector. By keeping track of how often a particular tower is flagged as bad by the AE and defining a threshold on this frequency, experts can choose to mask the transient tower proactively.

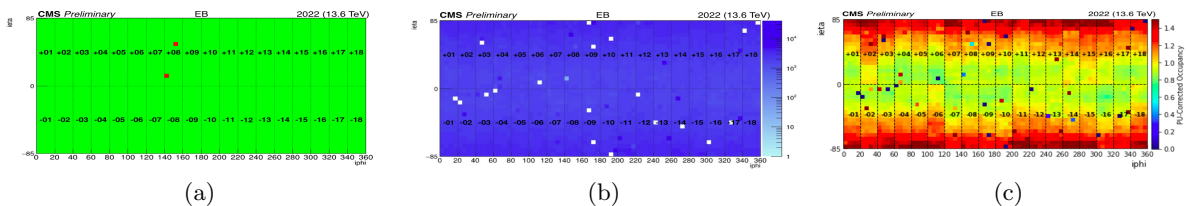


Figure 5: (a) ML quality plot accumulated over 9 LSs, from the MLDQM deployed in the online DQM workflow for EB from a 2022 run, showing two bad towers in red. (b) The digitized hit occupancy plot accumulated over all the LSs in the full run from the online DQM for the same run. (c) PU-corrected average occupancy plot over several runs in 2022.

6. Summary

A robust, deployment-ready AE-based anomaly detection and localization system has been developed for the CMS ECAL using unsupervised ML. An efficient pileup-based normalization strategy is applied on images made from the raw detector data, making quality interpretations consistent across changing experimental conditions. The application of novel techniques of spatial and time corrections gives an order of magnitude improvement in performance. The autoencoder-based system demonstrates efficient detection of anomalies of various degrees, shapes, and positions in the detector up to a tower-level granularity, using a single efficient threshold for each sub-detector part. The MLDQM system deployed in the online DQM workflow for the ECAL barrel performs well on real-time data from Run 3 by detecting anomalies as well as identifying potential degrading channels that goes undetected by the existing DQM plots. This ML-based DQM system is designed to complement and improve the existing DQM by helping detector experts make more accurate decisions and reduce false alarms. This autoencoder-based

anomaly detection system can be generalized and adapted to other particle physics experiments for data quality monitoring.

References

- [1] CMS Collaboration 2008 *JINST* **3** S08004
- [2] Azzolini V, van Besien B, Bugelskis D, Hreus T, Maeshima K, Fernandez Menendez J, Norkus A, Patrick J F, Rovere M and Schneider M A 2019 *EPJ Web Conf.* **214** 02003
- [3] Azzolini V, Andrews M, Cerminara G, Dev N, Jessop C, Marinelli N, Mudholkar T, Pierini M, Pol A and Vlimant J R (CMS) 2019 *EPJ Web Conf.* **214** 01007
- [4] Pol A A, Cerminara G, Germain C, Pierini M and Seth A 2019 *Comput. Softw. Big Sci.* **3** 3 (Preprint 1808.00911)
- [5] Hinton G and Salakhutdinov R 2006 *Science (New York, N.Y.)* **313** 504–7
- [6] 2022 URL <https://cds.cern.ch/record/2839738>
- [7] 2023 URL <https://cds.cern.ch/record/2850523>
- [8] CMS Collaboration 2010 *JINST* **5** T03010 (Preprint 0910.3423)
- [9] CMS Collaboration 2012 *Phys. Lett. B* **716** 30–61 (Preprint 1207.7235)
- [10] He K, Zhang X, Ren S and Sun J 2015 (Preprint 1512.03385)
- [11] Paszke A *et al.* 2019 Pytorch: An imperative style, high-performance deep learning library *Advances in Neural Information Processing Systems 32* (Curran Associates, Inc.) pp 8024–8035
- [12] Bai J, Lu F, Zhang K *et al.* 2019 Onnx: Open neural network exchange <https://github.com/onnx/onnx>
- [13] ONNX Runtime developers 2021 Onnx runtime <https://onnxruntime.ai/> version: x.y.z