Available on CMS information server

CMS DP -2024/121

The Compact Muon Solenoid Experiment

CMS Performance Note

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

20 November 2024

Model-Independent Real-Time Anomaly Detection at the CMS Level-1 Calorimeter Trigger with CICADA

CMS Collaboration

Abstract

In the search for new physics, real-time detection of anomalous events is critical for maximizing the discovery potential at the LHC. CICADA (Calorimeter Image Convolutional Anomaly Detection Algorithm) is a novel CMS trigger algorithm operating at the 40 MHz collision rate. By leveraging unsupervised deep learning techniques, CICADA aims to enable physics-model-independent trigger decisions, enhancing sensitivity to unanticipated signals. One of the key challenges is deploying such a system on resource-constrained hardware without compromising performance. This is addressed by utilizing knowledge distillation to replicate performance of an unsupervised anomaly detection model (teacher) in smaller supervised model (student) that maintains high detection sensitivity while significantly reducing the memory footprint and computational demands. The final compressed model is deployed on FPGAs, allowing CICADA to perform real-time decision-making while operating within the stringent constraints of the CMS trigger system during Run 3 data taking. This note details the training procedure of the CICADA anomaly detection algorithm and profiles the data that were used for it. The final model is evaluated on a mixture of real and simulated data. Furthermore, CICADA's behaviour in the Level-1 Trigger (L1T) is emulated, and comparisons with standard trigger algorithms are drawn. Lastly, a preliminary study based on 2024 proton-proton collision data is presented, indicating that CICADA is able to trigger on top quark pair production events. If not explicitly stated otherwise, the shown data were recorded in 2023.

Model-Independent Real-Time Anomaly Detection at the CMS Level-1 Calorimeter Trigger with CICADA

CMS Collaboration

cms-dpg-conveners-l1t@cern.ch
https://twiki.cern.ch/twiki/bin/view/CMSPublic/L1TriggerDPGResults

Abstract

In the search for new physics, real-time detection of anomalous events is critical for maximizing the discovery potential at the LHC. CICADA (Calorimeter Image Convolutional Anomaly Detection Algorithm) is a novel CMS trigger algorithm operating at the 40 MHz collision rate. By leveraging unsupervised deep learning techniques, CICADA aims to enable physics-model-independent trigger decisions, enhancing sensitivity to unanticipated signals. One of the key challenges is deploying such a system on resource-constrained hardware without compromising performance. This is addressed by utilizing knowledge distillation to replicate performance of an unsupervised anomaly detection model (teacher) in smaller supervised model (student) that maintains high detection sensitivity while significantly reducing the memory footprint and computational demands. The final compressed model is deployed on FPGAs, allowing CICADA to perform real-time decision-making while operating within the stringent constraints of the CMS trigger system during Run 3 data taking.

This note details the training procedure of the CICADA anomaly detection algorithm and profiles the data that were used for it. The final model is evaluated on a mixture of real and simulated data. Furthermore, CICADA's behaviour in the Level-1 Trigger (L1T) is emulated, and comparisons with standard trigger algorithms are drawn. Lastly, a preliminary study based on 2024 proton-proton collision data is presented, indicating that CICADA is able to trigger on top quark pair production events. If not explicitly stated otherwise, the shown data were recorded in 2023.

Introduction

The previous CMS Detector Performance note on CICADA [1] was published in November 2023. Since then, the machine learning models have undergone further fine-tuning. The latest student model has been deployed on an FPGA in the CMS L1T, recording data since October 2024. All results in this note reflect the state of the models in late 2024. Unlike the previous note, this one provides a detailed account of the training process, including profiling of training and validation data, loss curves, and insights into the teacher model. Additionally, the models' performance on simulated signals is analyzed more thoroughly, with ROC curves presented. CICADA's emulated trigger rates in relation to other L1T algorithms [2] are shown as well as a first look at proton-proton collision data recorded by CICADA in late 2024. Finally, we introduce initial results from a study that demonstrates the ability of CICADA to trigger on top quark pair production events.

Samples - Overview

Sample	# Events
2023 Zero Bias Run B (recorded on April 30th, 2023)	179988
2023 Ephemeral Zero Bias Run C (recorded on May 6th and 7th, 2023)	301669
2023 Zero Bias Run C (recorded on May 8th, 2023)	174997
2023 Zero Bias Run D (recorded on June 30th and July 4th, 2023)	174420
Soft Unclustered Energy Patterns (SUEP)	100000
Higgs to Long-Lived Particles (HtoLongLived)	40000
Vector Boson Fusion Higgs to Charm Pair (VBFHto2C)	98985
Top Quark Pair Production (TT)	100067
Supersymmetry Gluino-Gluino to Bottom-Quark Pair and Higgs Production (SUSYGGBBH)	98899
Outliers	20000
2023 Zero Bias (high pileup, only for evaluation) (recorded on July 7th, 2023)	441813

Table 1: Overview of the considered samples and corresponding sample sizes. The teacher autoencoder is trained exclusively on the Zero Bias (ZB) data samples in the upper half of the table. The training of the student model also uses the simulated outliers, however, without labeling them as such. Deliberately intended as an out-of-distribution dataset, the outliers are a random mixture of various simulated signal samples from a 2018 campaign. All other simulated samples are used only for the performance evaluation. All simulated samples were generated using next-to-leading order (NLO) matrix elements. Specifically, the VBFHto2C and TT samples utilized Powheg, while the remaining samples were generated using Pythia 8. For parton showering, Pythia 8 was consistently applied across all samples. Additionally, the NNPDF 3.0 NNLO parton density function (PDF) was used for all simulations. A more detailed list of samples can be found in the Appendix.

CICADA Inputs - Preface

CICADA processes inputs from the cylindrical electromagnetic and hadronic calorimeters, which are unrolled along ϕ to create a 2D, image-like grid. In this setup, energy deposits across 4x4 calo towers are summed to form an 18x14 pixel grid in the ($i\phi$, $i\eta$) domain. The 18 $i\phi$ bins (indexed 0-17) are uniformly spaced from $\phi = -\pi$ to $\phi = \pi$, while the 14 $i\eta$ bins (indexed 4-17) span from $\eta = -3$ to $\eta = 3$, symmetrically around 0. The first two and last two $i\eta$ bins are wider than the 10 equally spaced central bins, as outlined in Table 5. In this configuration, convolutional neural networks treat the inputs as ordered, image-like data, abstracting away their physical coordinates.

Axis	Bin Range	Physical Bin Edges	Bin Width
$i\phi$	0-8	$[0, 0.349, ..., \pi]$	0.349
	9-17	$[-\pi, -0.279, ..., 0]$	0.349
$i\eta$	4	$[-3.0, -2.172]$	0.828
	5	$[-2.172, -1.74]$	0.432
	6-15	$[-1.74, -1.392, ..., 1.74]$	0.348
	16	$[1.74, 2.172]$	0.432
	17	$[2.172, 3.0]$	0.828

Table 2: Schema explaining the relation between the ($i\phi$, $i\eta$) binning used for the CICADA inputs and the physical CMS default coordinates (ϕ, η).

Teacher Reconstruction Example - Background

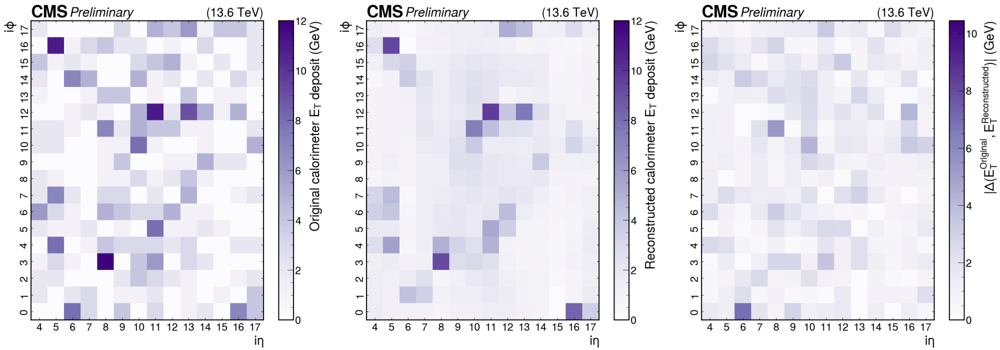

Figure 1: Reconstruction of a ZB event by the teacher model. From left to right: randomly chosen input event, resulting reconstructed output, and region-wise absolute error between the two. The mean squared error across all regions is $MSE = 2.57$ in this case.

Teacher Reconstruction Example - Signal

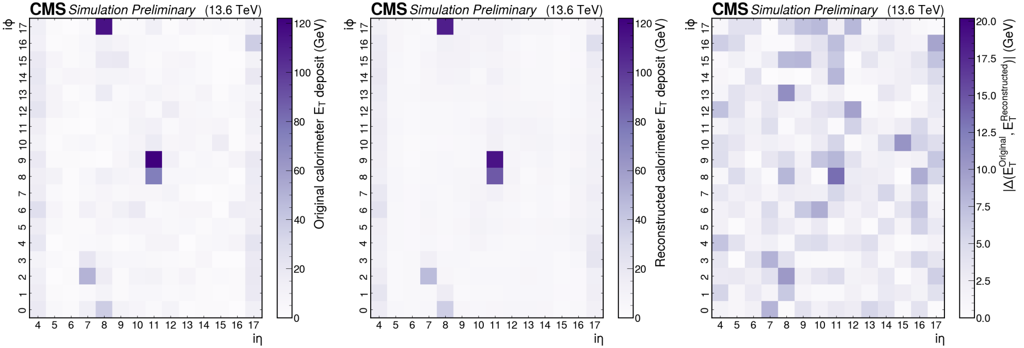

Figure 2: Reconstruction of a signal event (SUSYGGBBH) by the teacher model. From left to right: randomly chosen input event, resulting reconstructed output, and region-wise absolute error between the two. The mean squared error across all regions is $MSE = 14.89$ in this case.

Training History & Anomaly Score - Preface

An autoencoder is trained on ZB data using an unsupervised learning approach that enables it to learn a compressed representation of the calorimeter image data. The loss function for the autoencoder is the mean squared error (MSE) between the original input and the reconstructed output, calculated by summing the squared differences between the original and reconstructed E_T map over all 252 pixels:

$$\frac{1}{252}\sum_{i=1}^{252}\left(E_{T,i}^{original} - E_{T,i}^{reconstructed}\right)^2.$$

The model learns to best reconstruct common patterns in ZB data, with the MSE serving as an indicator of how anomalous an event is. However, the autoencoder is too large and slow for deployment in the L1 Trigger system. To address this, knowledge distillation is used to transfer the autoencoder's knowledge (teacher) to a smaller model (student), which directly learns to map the input image to the anomaly score, defined as

$$32 \cdot \log MSE$$

in a supervised regression task with a mean absolute error loss function. Experience from other teacher-student frameworks suggests that simultaneous learning is beneficial. After every teacher training epoch, a mix of ZB and outlier data is inferred with the teacher model, and the resulting anomaly scores are used to train the student for ten epochs. This approach results in the student undergoing ten times more training epochs than the teacher. As the teacher improves at reconstructing ZB data, its loss decreases, while the student, learning to match the teacher's performance, maintains a relatively flat loss curve despite improving as can be seen in Fig 3.

Training History

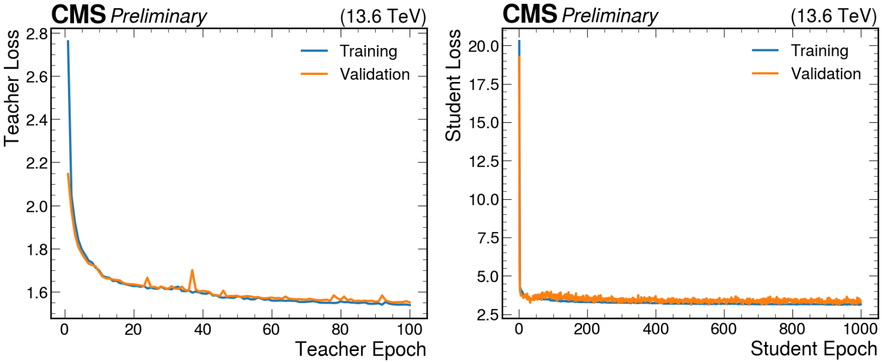

Figure 3: Training histories of teacher (left) and student (right). After each training epoch, the teacher is inferred to create new targets on which the student is trained for ten epochs.

Anomaly Score

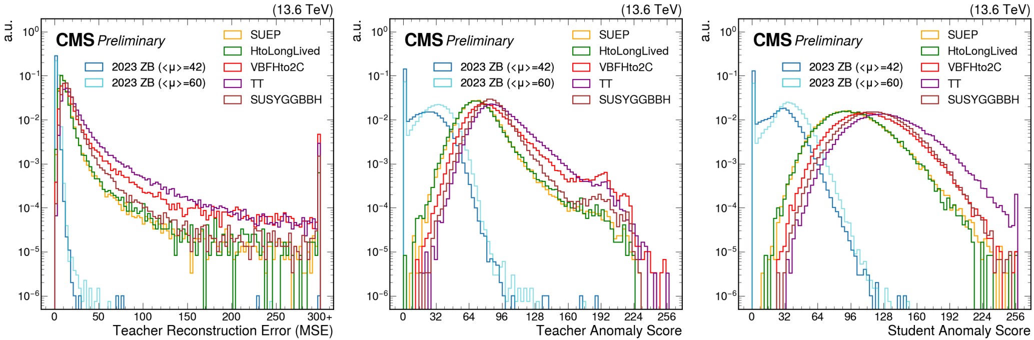

Figure 4: Teacher MSE distribution (left) and resulting anomaly scores (middle) for ZB data and various simulated signal samples as predicted by the teacher. The anomaly score is calculated as $s = \log MSE \cdot 32$ and subsequently quantized as 16 bits where 8 are used for the integer part, resulting in 65536 evenly spaced discrete values between 0 and 256. The plot on the right shows the anomaly scores as predicted by the student. In addition to the ZB test dataset consistent with the teacher's training data, the ZB dataset with higher pileup is also depicted to illustrate the effect on the anomaly score. The mean pileup of the simulated signals is 62.

ROC Curves - Preface

A ROC curve for a binary classifier plots the true positive rate (TPR) against the false positive rate (FPR). To generate this, we combine the entire ZB test dataset with the simulated signal for each process under consideration. The FPR is the fraction of ZB events exceeding a specific anomaly score threshold, while the TPR is the fraction of signal events above that threshold. In practice, we sort the dataset by anomaly score in descending order, iteratively calculating efficiencies as we encounter each ZB event to mark points on the ROC curve.

This process can be visualized as moving a lower anomaly score threshold from 256 leftward in Figure 4, recalculating the TPR at each ZB event, and plotting each point. The FPR can be converted into a trigger rate by multiplying it with the bunch crossing frequency for the data-taking period (approximately 28.61 MHz). With 332430 ZB events in the test set, each additional ZB event above the threshold increases the estimated trigger rate by about 86 Hz - determining the bin width in Fig. 7 and 8. While this introduces a large statistical uncertainty at lower trigger rates (4 ZB events to calculate the FPR in the first bin), the results show that the student consistently outperforms the teacher, as reflected in higher ROC areas under the curve (AUC). This improvement likely stems from the smaller model's capacity to compress training data into a more compact representation, enhancing generalization and signal-background separation.

Both ROC curve plots also show the performance of a 'Baseline' algorithm for reference. Its score is the squared sum of all energy deposits, $\sum_{i=1}^{252} E_{T,i}^2$, which was found to be a well-performing simplified classifier.

The following two roc curves use the 2023 ZB test dataset with an average pileup of 42 as background. The two roc curves after that deploy a ZB dataset with a higher mean pileup of 60 as background. Although the discrimination power between signal and background deteriorates with increasing background pileup, CICADA consistently outperforms the baseline algorithm. Now that the algorithm has been shown to successfully run online, pileup mitigation techniques are being studied.

ROC Curves Teacher

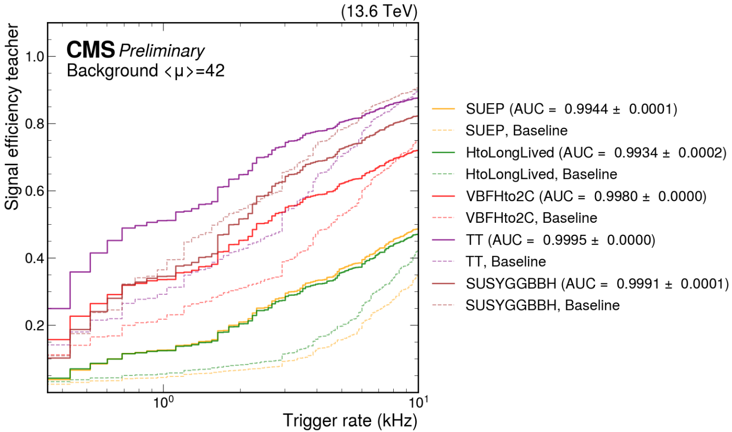

Figure 5: ROC curve of the teacher for various simulated signal samples. Squared sum of all energy deposits is deployed as score for the *Baseline* classifier. The 2023 ZB test dataset consistent with the teacher's training data and an average pileup of 42 was deployed as background.

ROC Curves Student

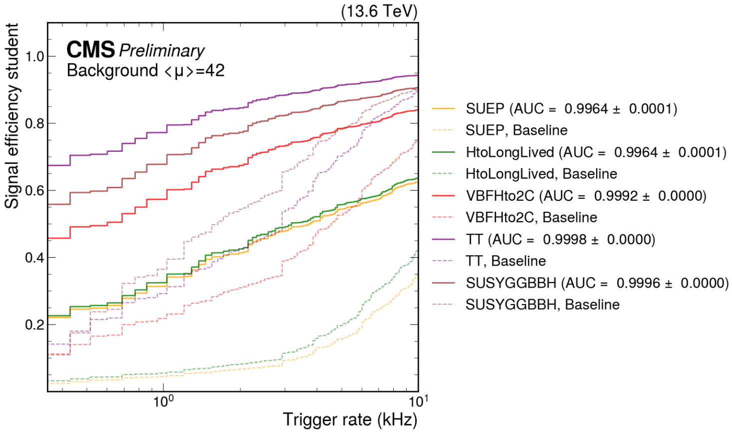

Figure 6: ROC curve of the student for various simulated signal samples. Squared sum of all energy deposits is deployed as score for the *Baseline* classifier. The 2023 ZB test dataset consistent with the teacher's training data and an average pileup of 42 was deployed as background.

ROC Curves Teacher (High Pileup Background)

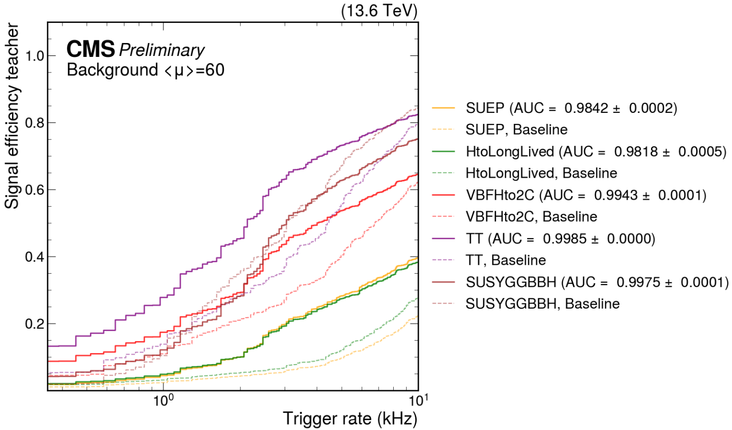

Figure 7: ROC curve of the teacher for various simulated signal samples. Squared sum of all energy deposits is deployed as score for the *Baseline* classifier. The deployed background sample is the 2023 ZB dataset with an average pileup of 60, which is higher than that of the training sample (42).

ROC Curves Student (High Pileup Background)

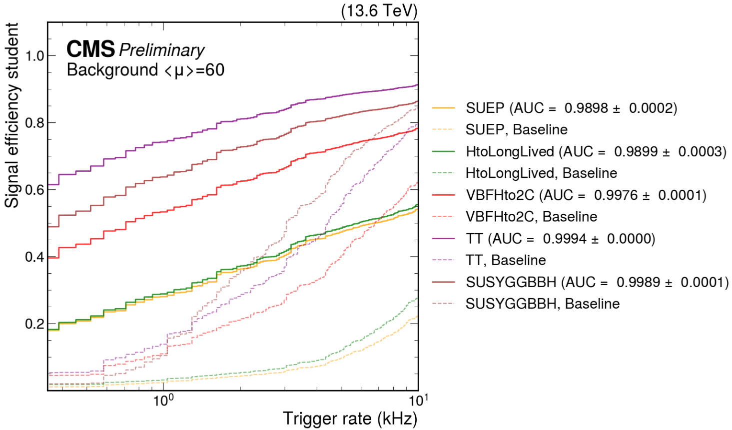

Figure 8: ROC curve of the student for various simulated signal samples. Squared sum of all energy deposits is deployed as score for the *Baseline* classifier. The deployed background sample is the 2023 ZB dataset with an average pileup of 60, which is higher than that of the training sample (42).

L1 Emulator - Score & Trigger Rates

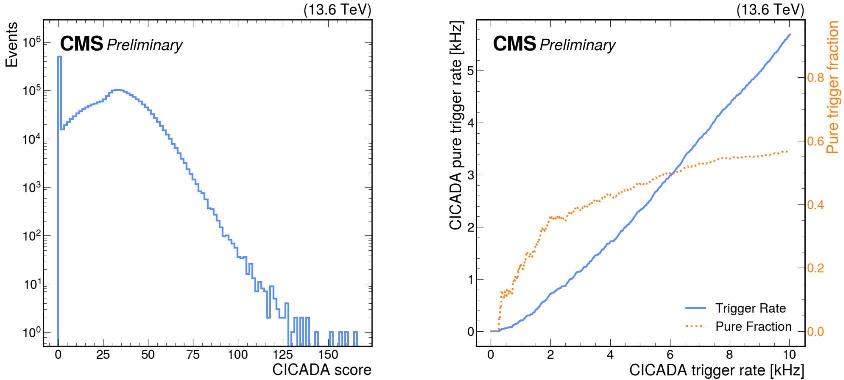

Figure 9: Emulated CICADA anomaly score (left) and L1 trigger rates (right) for 2023 ZB data. 'Pure' refers to events where no L1 bit except for CICADA has fired. As expected, the pure rate approaches 1 when loosening cut on the anomaly score. The plot also shows that operating CICADA at around 1 kHz would lead to roughly 200 Hz of novel events at L1.

2024 Data CICADA Event Display

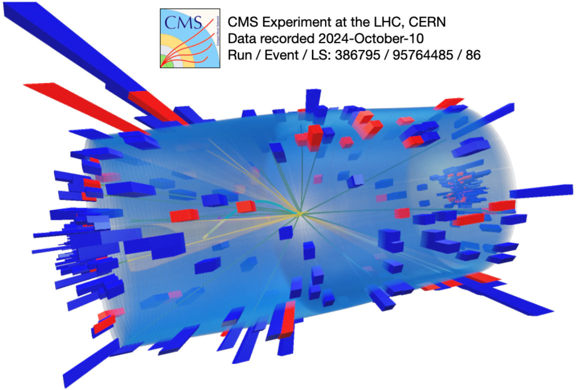

Figure 10: A proton-proton collision event recorded by CICADA on October 10th, 2024 during the CMS data-taking era 2024I. This displays an event, where CICADA was the only L1T algorithm that has fired. Run number 386795, lumi section 86, event number 95764485.

TTbar Analysis with CICADA - Event Selection & Mass Reconstruction

Analyze ZB data with following event selection:

- At least three jets
- At least one b-tagged jet[a]

Mass Reconstruction

- Examine all possible 3-jet combinations with at least 1 b-tagged jet in each event
- Calculate the invariant mass of the combination with the highest combined p_T

[a] btagDeepCSV > 0.5 [3]

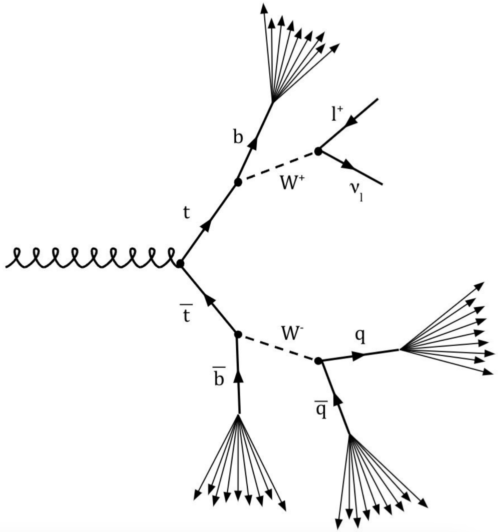

TTbar Analysis with CICADA - Fit Strategy

- Fit signal & background normalization in reconstructed top mass
- 'Data': ZB data with cut on CICADA score > 115
- 'Background': ZB data without cut on CICADA score
- 'Signal': Simulated TTbar data with cut on CICADA score > 115

Demonstration of CICADA's ability to trigger on interesting physics (ttbar). Not a cross section measurement!

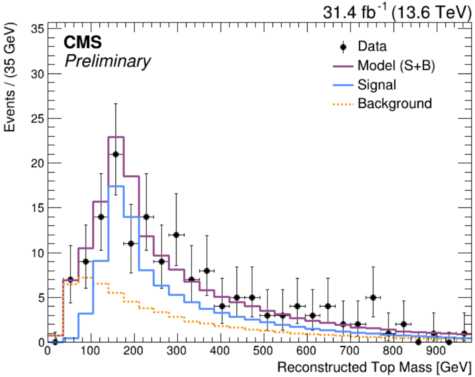

Figure 11: Post-fit distribution of the reconstructed top mass. In this case, a cut of CICADA score > 115 was used to define the 'Data' and 'Signal' shapes.

References I

[1] CMS Collaboration, "Level-1 Trigger Calorimeter Image Convolutional Anomaly Detection Algorithm", *https://twiki.cern.ch/twiki/bin/view/CMSPublic/CICADA2023* (2023).

[2] CMS Collaboration, "Performance of the CMS Level-1 trigger in proton-proton collisions at $\sqrt{s} =$ 13 TeV", *JINST* **15** (2020), no. 10, P10017, doi:10.1088/1748-0221/15/10/P10017, arXiv:2006.10165.

[3] CMS Collaboration, "Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV", *JINST* **13** (2018), no. 05, P05011, doi:10.1088/1748-0221/13/05/P05011, arXiv:1712.07158.