CICADA: Anomaly Detection for New Physics Searches at the CMS Level-1 Trigger

Ho Fung Tsoi (Wisconsin) for the CICADA Team in the CMS Collaboration

US LUA 2023 Lightning Round Talk, Dec 12 – 15

CICADA team

CICADA: Calorimeter Image Convolutional Anomaly Detection Algorithm (AD at Calo)

Model/Implementation/Emulator/Firmware/Software/Hardware/Operations

- **University of Wisconsin-Madison**
  - Kiran Das, Sridhara Dasu, Tom Gorski, Alexander Savin, Varun Sharma, Ales Svetek, Jesra Tikalsky, Ho Fung Tsoi

- **Princeton University**
  - Pallabi Das, Kiley Kennedy, Andrew Loeliger, Luis Moreno, Isobel Ojalvo, Adrian Alan Pol

Background

CICADA: Calorimeter Image Convolutional Anomaly Detection Algorithm

What
- ML-based anomaly detection to search for rare/new physics in a model-independent way using low-level calo trigger info

Why
- Novel method to improve the current triggers

Where
- Calorimeter layer-1 trigger subsystem (CaloLayer-1) at the CMS Level-1 trigger

When
- Expecting Run 3 in 2024

Who
- Wisconsin/Princeton (members from the previous slide)

- https://cicada.web.cern.ch/
- CMS-DP-2023-086

Motivations

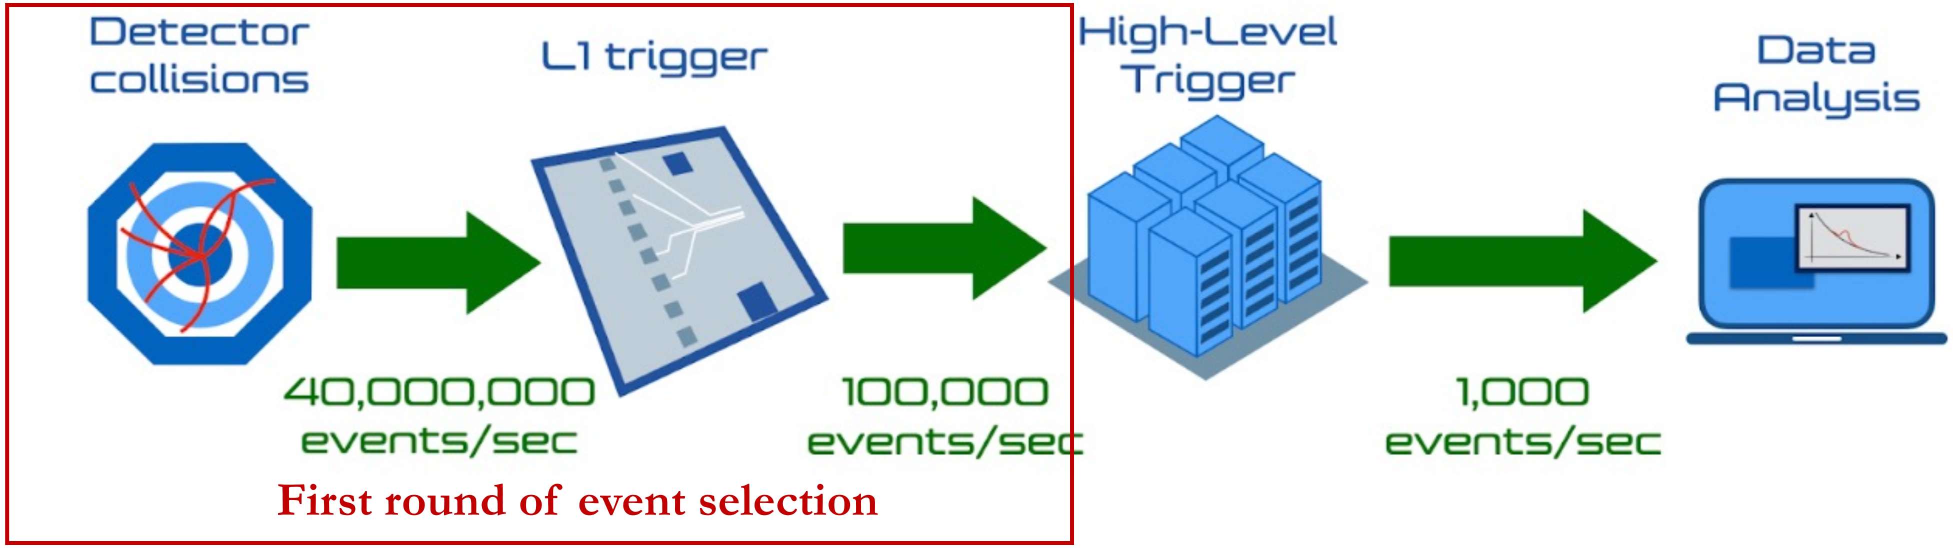

No BSM discovery at the LHC (yet!), three possibilities:

- New physics not possible at the current LHC energy scale
- Not enough data collected
- Maybe new physics already there, but we are looking at the wrong places or using wrong event selections

ML-based anomaly detection trigger

- Minimize human bias, completely data-driven
- ML can unearth unknown and complex correlation
- New physics searches in model-independent way

LHC in the language of anomaly detection

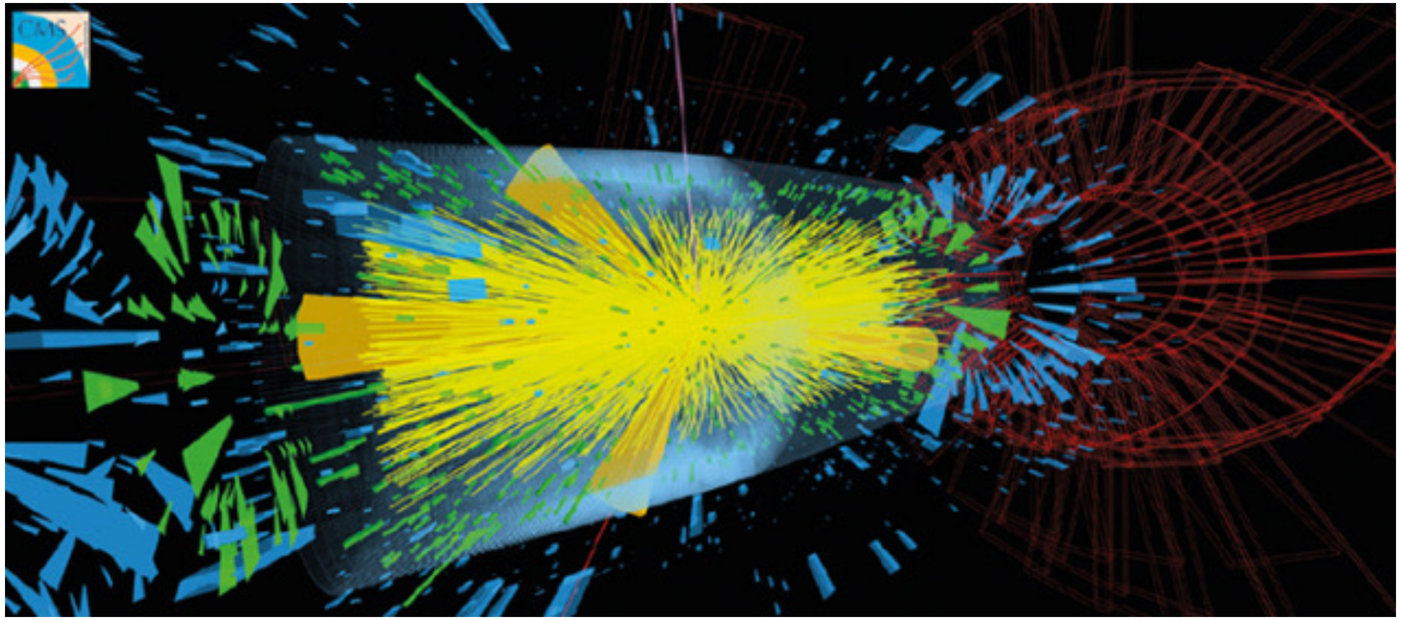

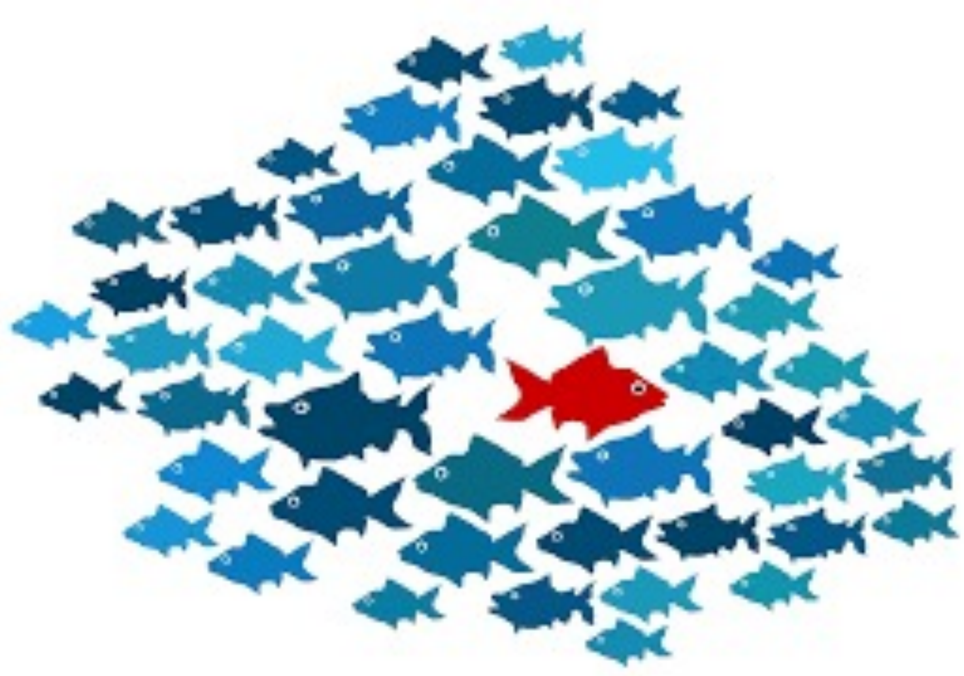

"Normal" events

- Whatever processes dominating the collisions
  - ➢ Large production cross sections at the LHC
  - ➢ Soft QCD…

"Anomalous" events

- Whatever processes deviating from the "normal"
  - ➢ SM with small cross sections and BSM
  - ➢ Top, Higgs, SUSY, …

Anomaly trigger vs. the traditional

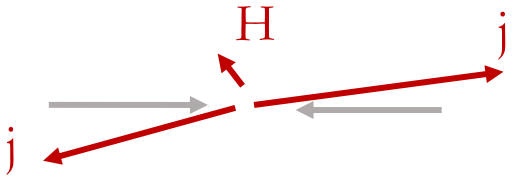

Model-dedicated trigger (e.g. VBF Higgs topology)

- Low rate and very sensitive to a particular signal class
- But phase space too limited, cannot be generalized

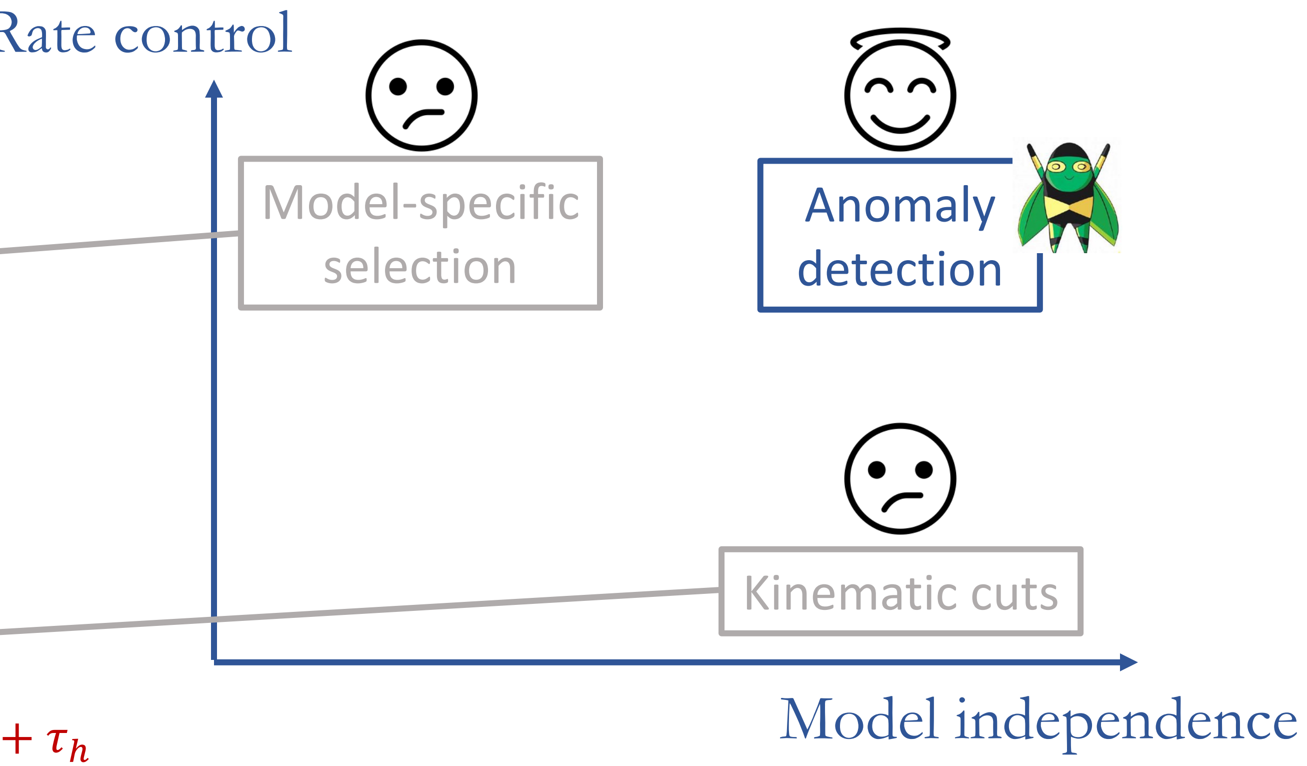

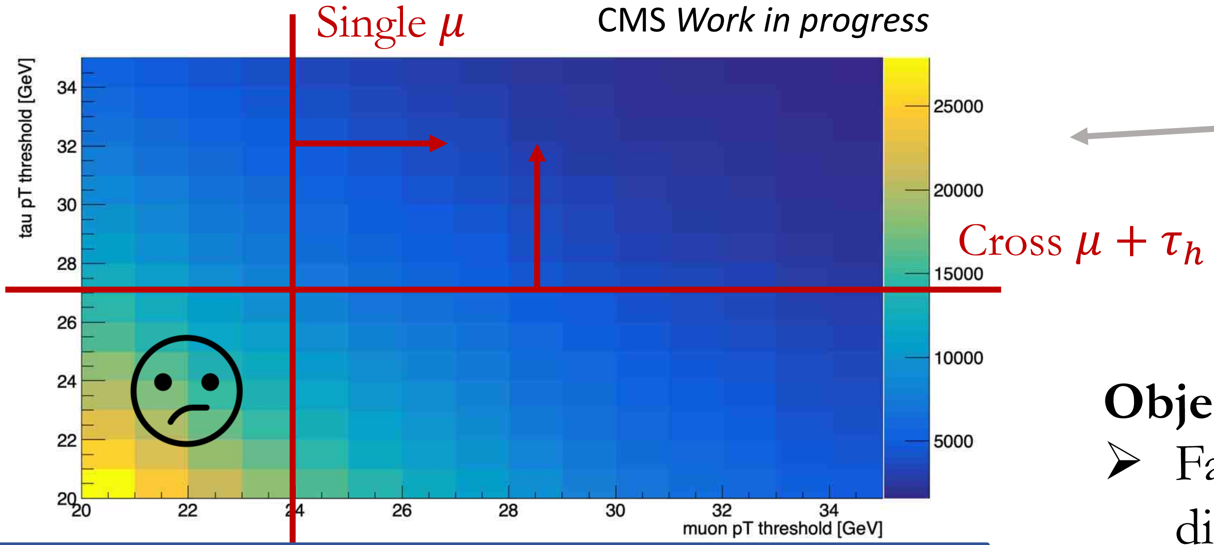

E.g. signal yield in the $\mathrm{h} \rightarrow aa \rightarrow \tau\tau\mathrm{bb}$ analysis with varying τ_μ and τ_h p_{T}, before and after L1 triggers

Object-based trigger (e.g. μ $p_{\mathrm{T}} > X$)

- Fairly model independent, can be used to trigger different signals
- But rejects lots of signals due to high thresholds

Input at CaloLayer-1

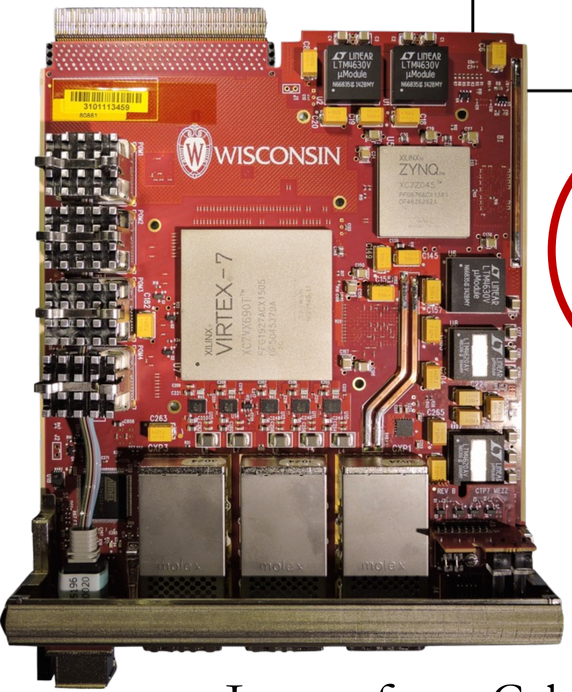

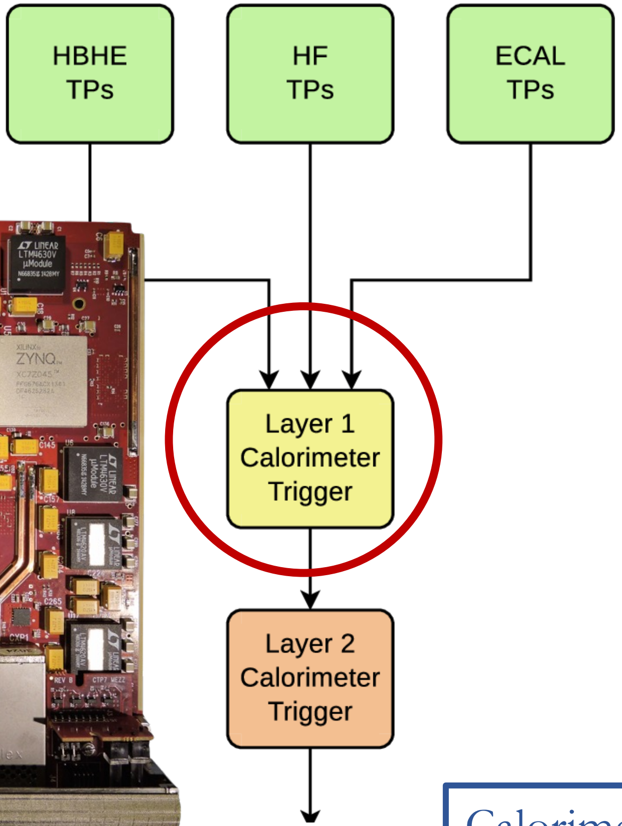

One region = 4x4 trigger towers

CMS-DP-2023-086

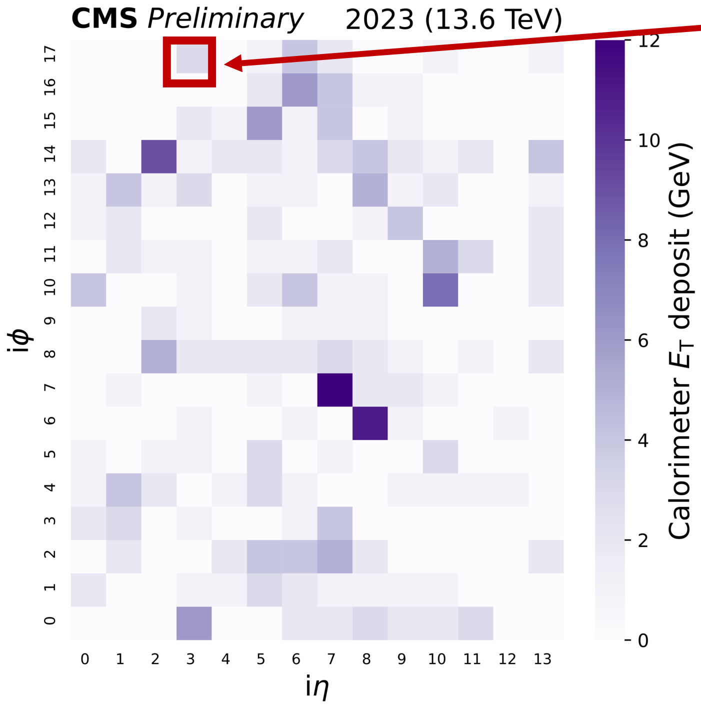

Calorimeter E_T deposit from one Zerobias event

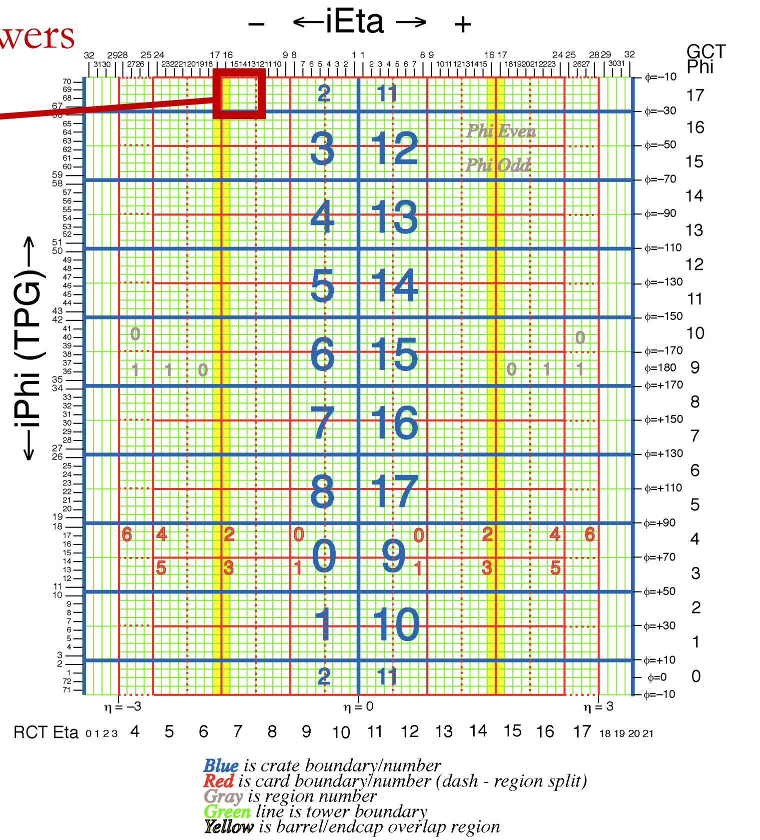

Calorimeter trigger tower-region map

Inputs from CaloLayer1

- $18\ \phi \times 14\ \eta = 252$ regions in total
- Each region contains energy deposits from both ECAL and HCAL
- Summary of the energy distribution profile within the region
- Low-level information not dependent on jet reconstruction etc.

Autoencoder model

Model architecture: calo input → encoder → latent space → decoder → reconstructed input

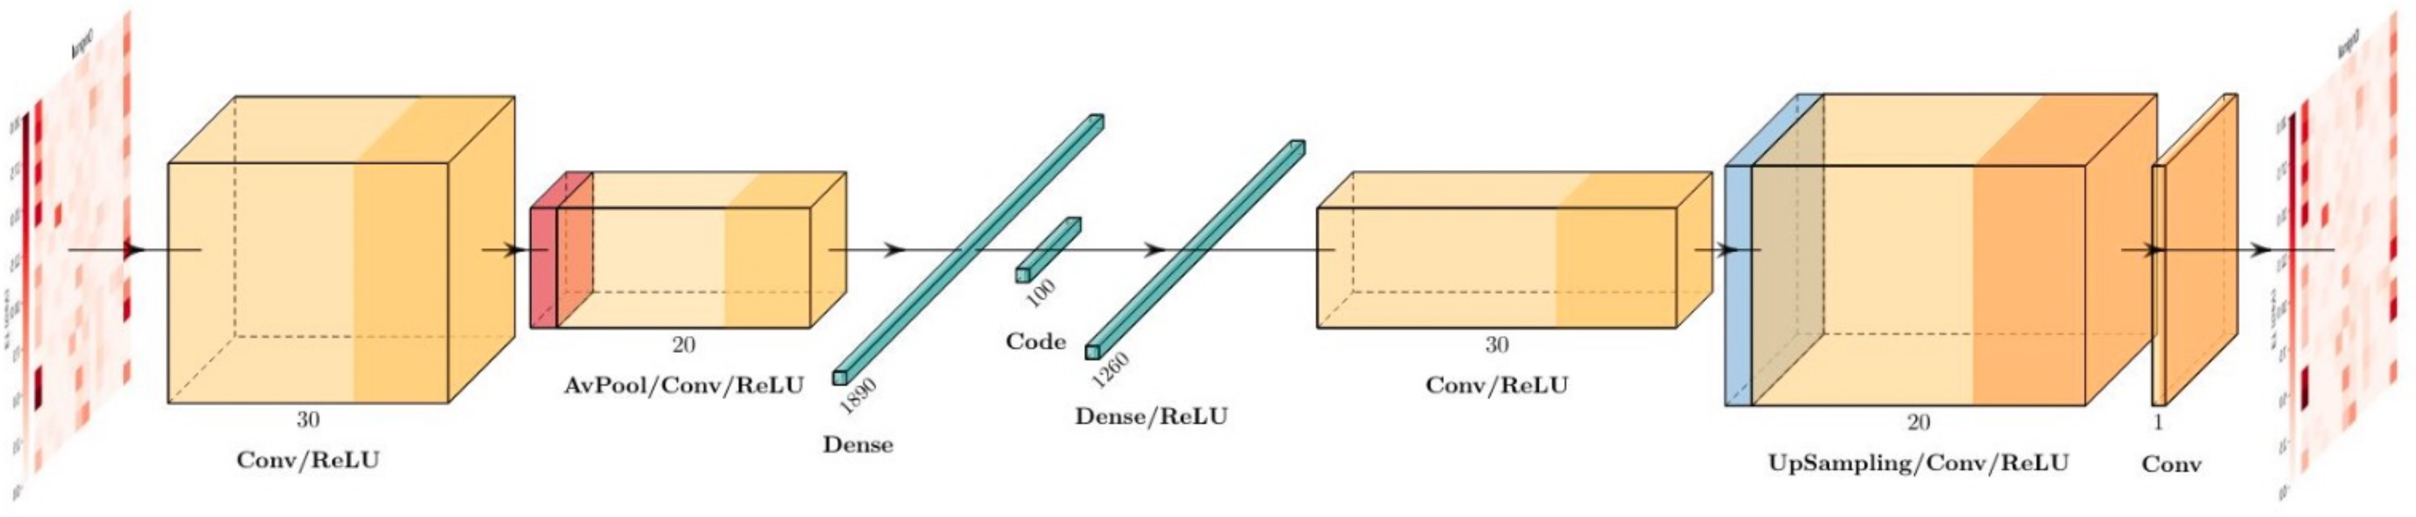

CICADA project: Calorimeter Image Convolutional Anomaly Detection Algorithm

➢ https://cicada.web.cern.ch/

Autoencoder-based anomaly detection

- Input is 2D tensor from the calorimeter region energy information
- Encoder and decoder are convolutional neural networks
- Unsupervised learning: train only on ZeroBias data to learn input reconstruction

Model expectation: reconstruction

ZeroBias data

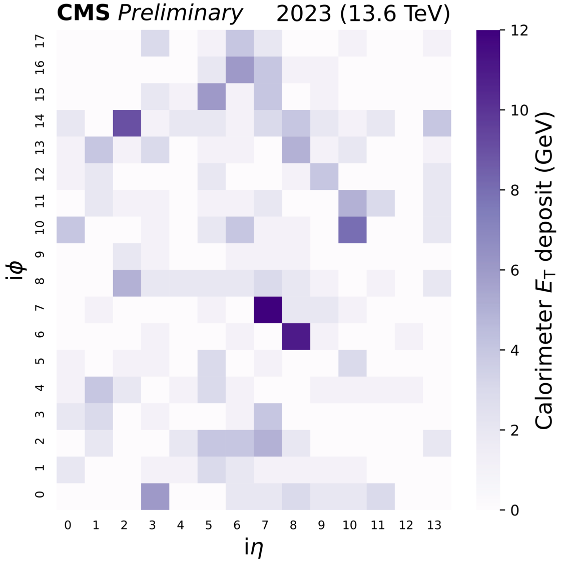

Input

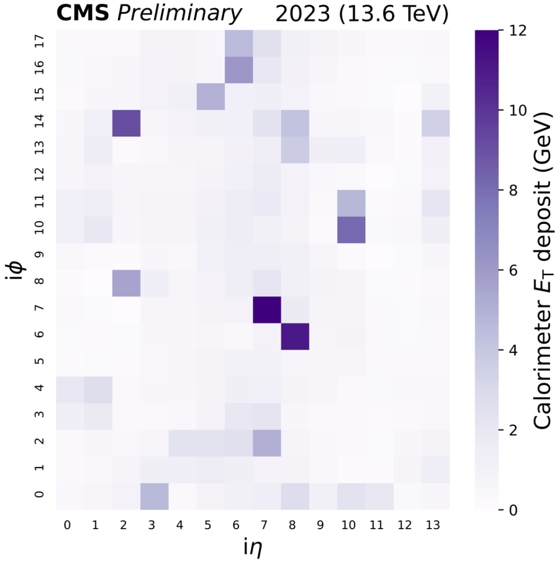

Reconstruction

BSM MC signal

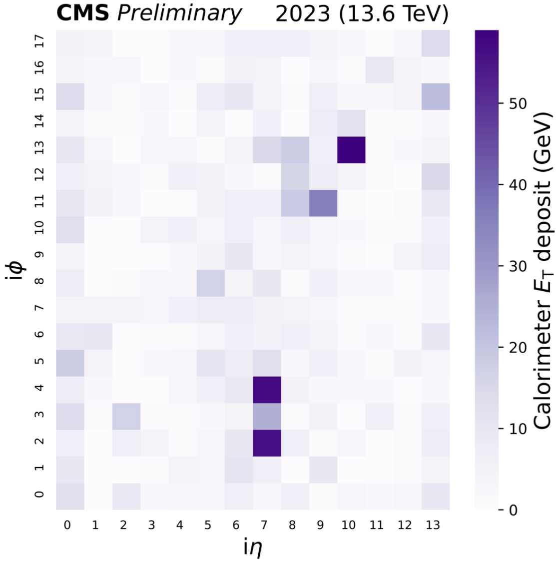

Input

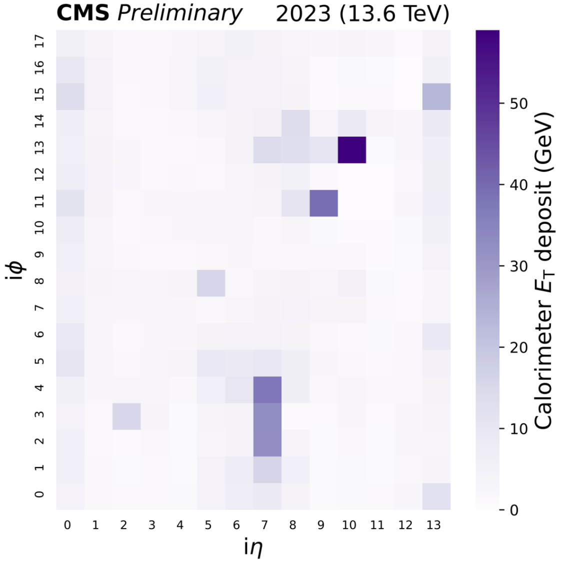

Reconstruction

CMS-DP-2023-086

Expectation:

- **Good reconstruction** on normal events (ZeroBias used for training)
- **Bad reconstruction** on anything else such as rare SM or BSM signals (never seen in training)

Goal:

- Use metric like the mean-squared error MSE(input, output) as anomaly score to trigger on anomalous events

Score distributions and ROCs

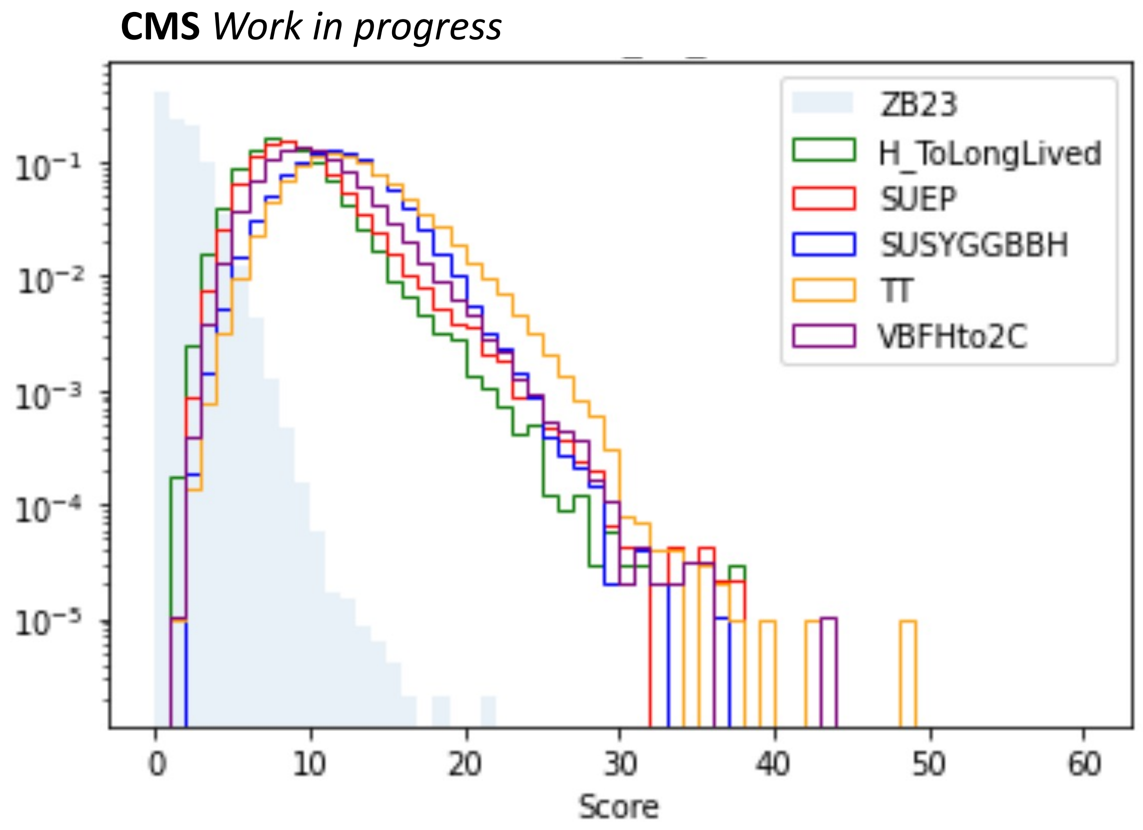

Anomaly score distribution

- ZeroBias data have low anomaly score
- Rare SM and BSM processes have high anomaly score

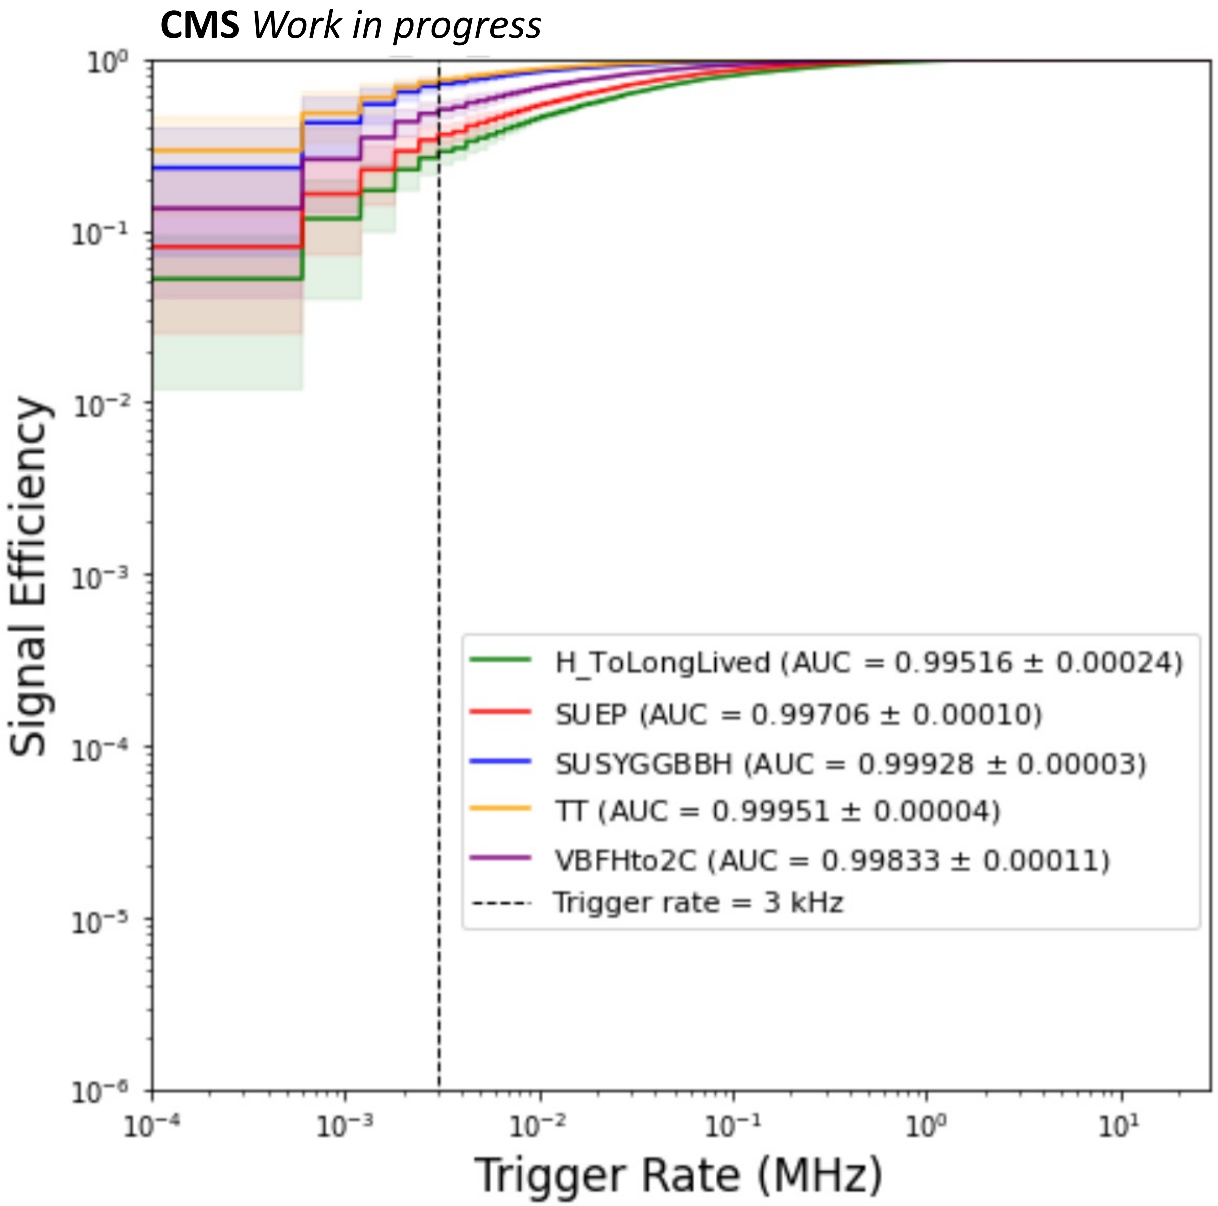

ROCs of a range of MC signals vs ZeroBias

Stability as a trigger

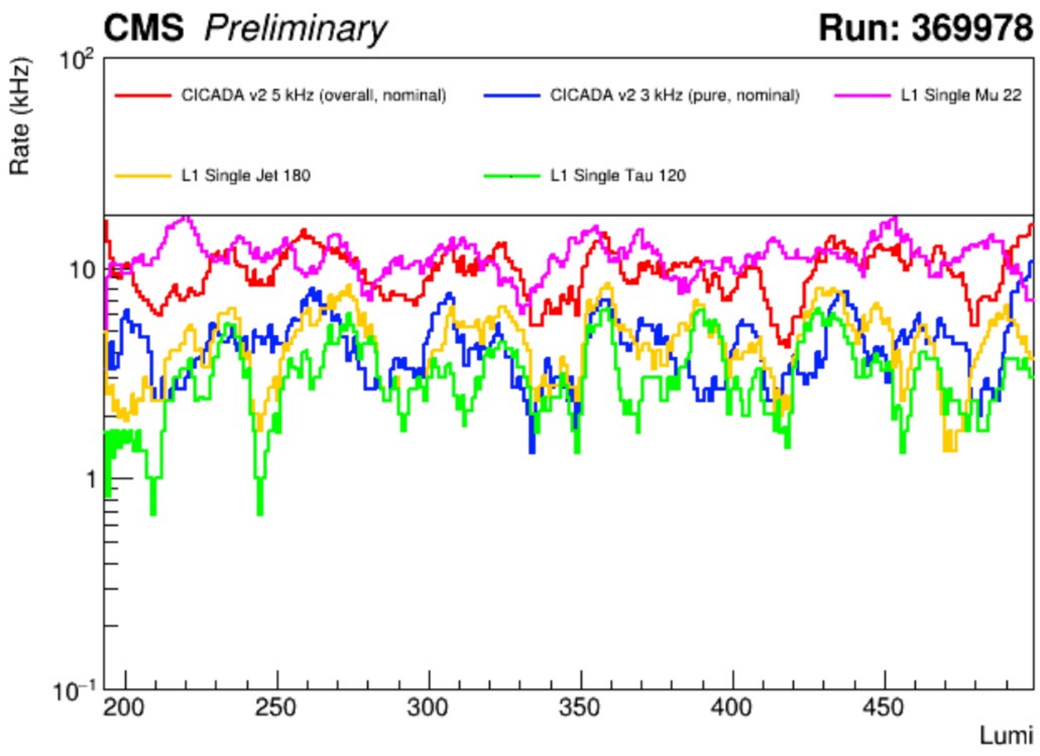

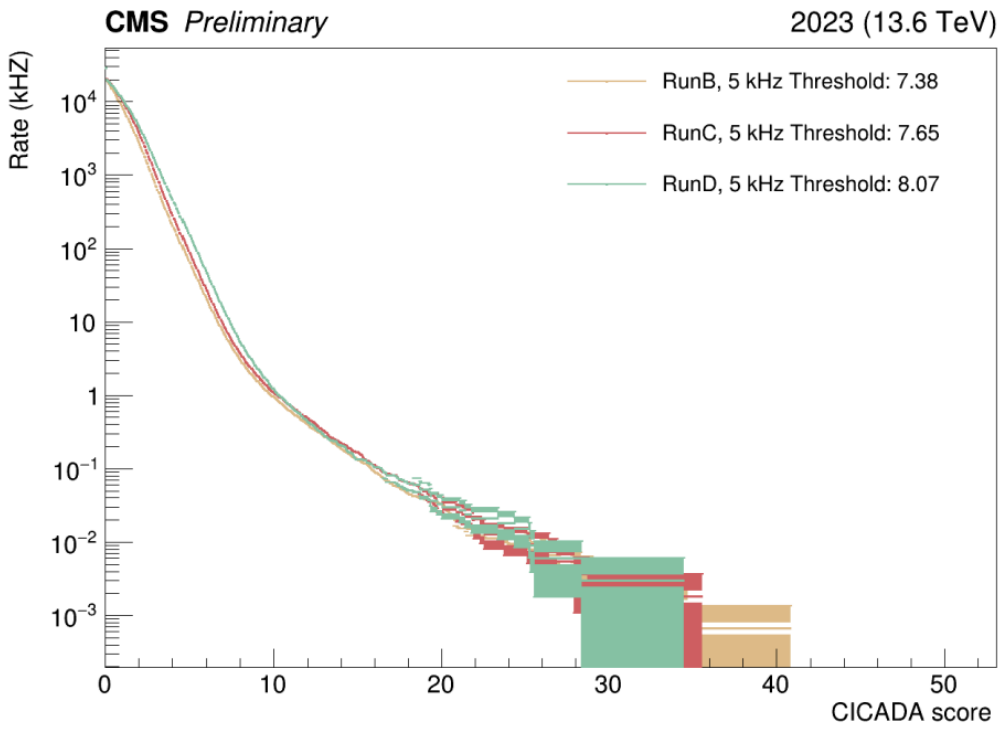

CMS-DP-2023-086

- Stable trigger rate over a run
- Stability verified and comparable to existing standard L1 seeds

- Stable trigger rate across different runs
- Flexible trigger: tunable threshold available to adapt changing run conditions

Summary

- **CICADA**
  - Novel trigger at L1 using ML-based anomaly detection
  - Search for rare/new physics in one place
  - Model-independent

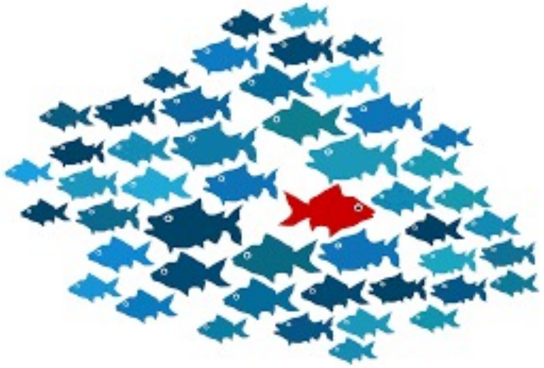

- **Demonstrated physics sensitivity**
  - Sensitive to a wide range of rare SM and BSM signals
  - Favoring high object multiplicity final states in general
  - Does catch more signals that are otherwise rejected by the current triggers

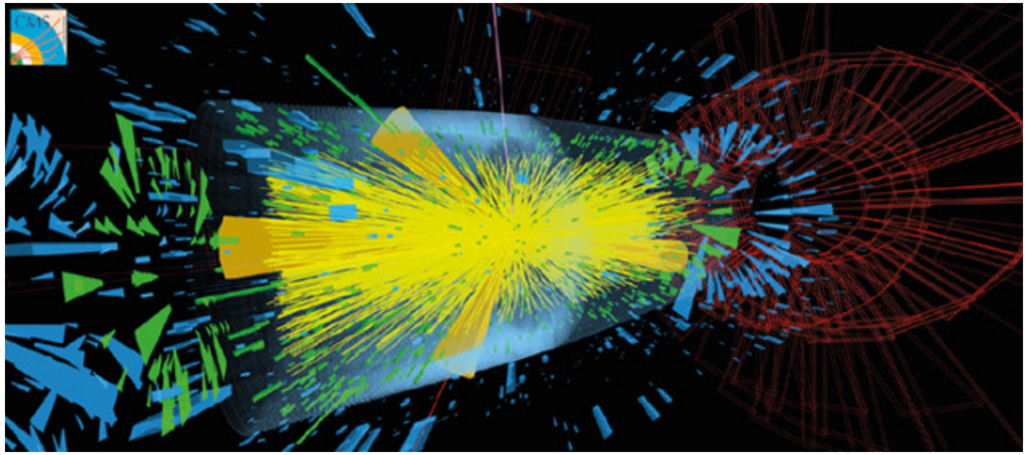

- **Demonstrated trigger capacity**
  - Tunable threshold for flexible rate control
  - Stable rates over a run and across different runs

- **Looking forward to the production starting from 2024**

Thank you for your attention!

Backup

Public space

Spaces · cicada-project/**cicada-demo** · like 0 · Running

Calo Deposits

17,0,0,1,0,2,17,8,0,0,0,0,0,14
0,1,1,0,0,1,0,0,0,0,5,0,0,1
0,0,0,3,14,0,0,0,1,0,0,0,0,1
1,2,1,0,8,3,1,0,2,0,0,0,1,0
0,0,0,0,7,2,0,11,0,0,0,0,1,1
0,0,0,0,0,0,13,1,0,0,0,1,0,0
0,0,0,2,17,0,0,4,0,1,0,1,0,0
0,0,0,1,0,2,0,1,0,2,2,1,0,2
1,0,21,0,65,1,4,0,3,0,0,1,1,1
0,0,36,0,0,3,0,9,0,22,0,2,0,0
14,1,0,2,1,0,0,0,10,0,1,1,8,0
0,0,0,0,0,10,2,0,0,1,0,0,4,2
0,0,0,0,0,1,0,0,0,1,0,0,4,0
0,0,0,0,2,0,0,3,0,1,0,0,0,0
1,0,0,4,3,0,0,0,0,0,1,1,1,4
0,0,30,0,4,39,1,0,4,0,2,2,1,0
0,0,0,1,0,4,0,4,0,0,4,39,0,0
2,6,0,1,5,3,0,0,1,0,1,0,1,4

CICADA Anomaly Score for CICADA v1

35.90625

CICADA Anomaly Score for CICADA v2

71.03125

Calorimeter Input | Saliency Map for CICADAv1 | Saliency Map for CICADAv2

Plot

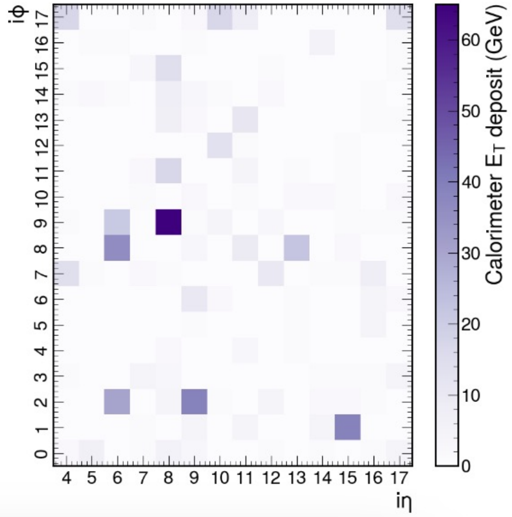

Generate random input

Do CICADA inference

CICADA

https://cicada.web.cern.ch/

Challenges at L1

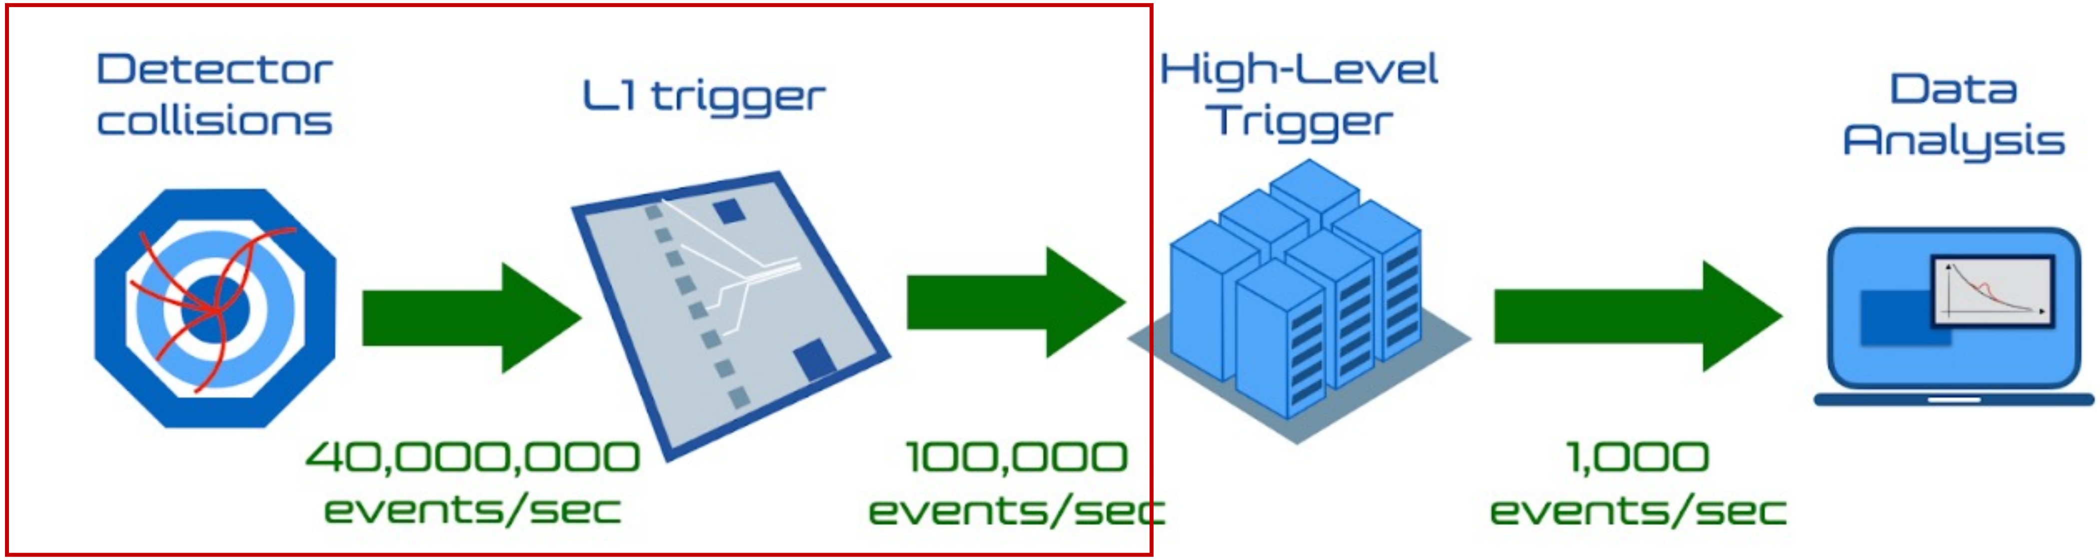

Fast trigger decision to fit the buffering system

- Extremely low latency < O(1) μs

Running on high-speed custom hardware at the edge

- Extremely tight resources from a single FPGA board

Data rate reduction from 40 MHz to 100 kHz

- Extremely low background rate < O(10^{-4}) demanded

→ Solved by model compression tricks to downscale the model size while preserving model performance

Naïve autoencoder model

Layer (type)	Output Shape	Param #
input (InputLayer)	[(None, 18, 14, 1)]	0
conv2d_1 (Conv2D)	(None, 18, 14, 20)	200
relu_1 (Activation)	(None, 18, 14, 20)	0
pool_1 (AveragePooling2D)	(None, 9, 7, 20)	0
conv2d_2 (Conv2D)	(None, 9, 7, 30)	5430
relu_2 (Activation)	(None, 9, 7, 30)	0
flatten (Flatten)	(None, 1890)	0
latent (Dense)	(None, 80)	151280
dense (Dense)	(None, 1890)	153090
reshape2 (Reshape)	(None, 9, 7, 30)	0
relu_3 (Activation)	(None, 9, 7, 30)	0
conv2d_3 (Conv2D)	(None, 9, 7, 30)	8130
relu_4 (Activation)	(None, 9, 7, 30)	0
upsampling (UpSampling2D)	(None, 18, 14, 30)	0
conv2d_4 (Conv2D)	(None, 18, 14, 20)	5420
relu_5 (Activation)	(None, 18, 14, 20)	0
output (Conv2D)	(None, 18, 14, 1)	181

Total params: 323,731
Trainable params: 323,731
Non-trainable params: 0

Encoder (compressor)

Latent space (compressed input)

Decoder (decompressor)

Despite good performance, model size of 300 k parameters can certainly fail the L1 constraints…

Model compression

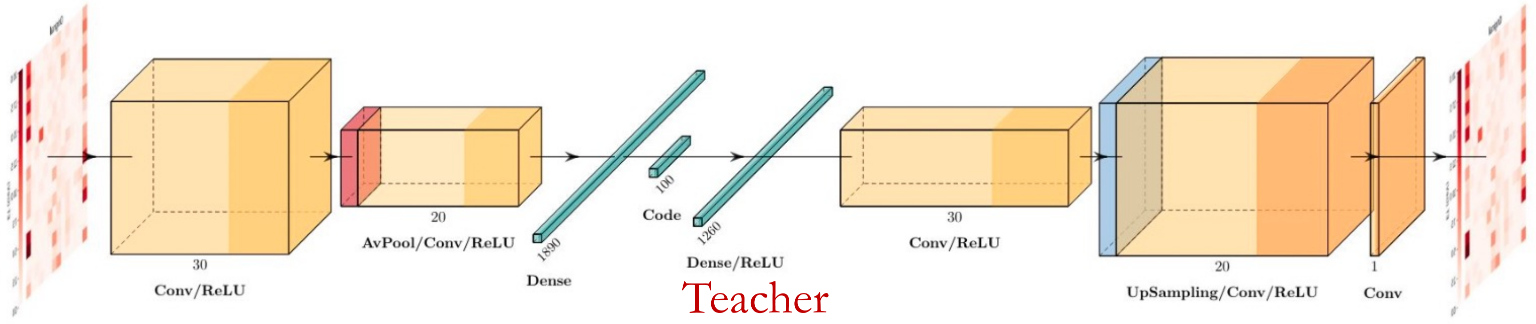

Teacher

Compute MSE(input, output) → Anomaly score

Compare with teacher's score

Score regression

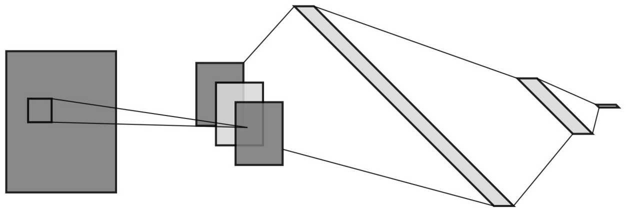

Student

→ hls4ml

QKeras github

hls4ml github

Knowledge Distillation

- Train a smaller model (student) under the guidance of a bigger model (teacher)
- The student can directly learn to regress MSE from teacher outputs

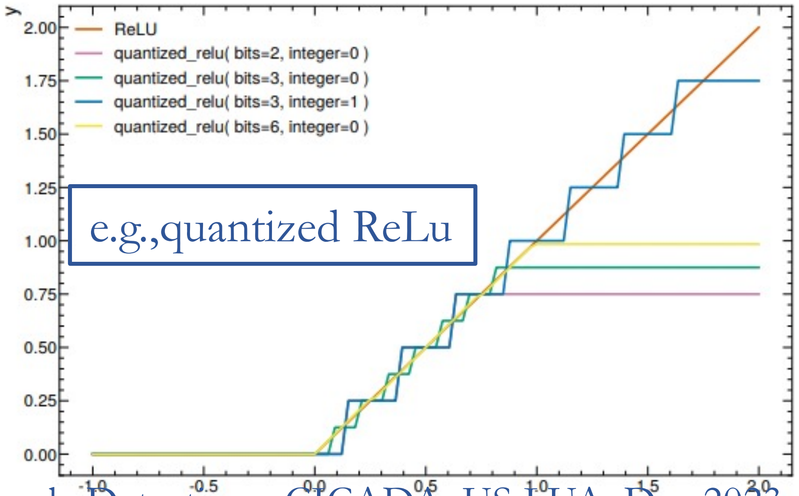

e.g.,quantized ReLu

Quantization-aware training (QKeras)

- Model weights quantized to fixed precision
- ➢ Train a quantized model rather than quantize a trained model

→ x10 reduction in resources/latency

Model: teacher → student

Layer (type)	Output Shape	Param #
input (InputLayer)	[(None, 18, 14, 1)]	0
conv2d_1 (Conv2D)	(None, 18, 14, 20)	200
relu_1 (Activation)	(None, 18, 14, 20)	0
pool_1 (AveragePooling2D)	(None, 9, 7, 20)	0
conv2d_2 (Conv2D)	(None, 9, 7, 30)	5430
relu_2 (Activation)	(None, 9, 7, 30)	0
flatten (Flatten)	(None, 1890)	0
latent (Dense)	(None, 80)	151280
dense (Dense)	(None, 1890)	153090
reshape2 (Reshape)	(None, 9, 7, 30)	0
relu_3 (Activation)	(None, 9, 7, 30)	0
conv2d_3 (Conv2D)	(None, 9, 7, 30)	8130
relu_4 (Activation)	(None, 9, 7, 30)	0
upsampling (UpSampling2D)	(None, 18, 14, 30)	0
conv2d_4 (Conv2D)	(None, 18, 14, 20)	5420
relu_5 (Activation)	(None, 18, 14, 20)	0
output (Conv2D)	(None, 18, 14, 1)	181

Total params: 323,731
Trainable params: 323,731
Non-trainable params: 0

Teacher

Layer (type)	Output Shape	Param #
In (InputLayer)	[(None, 252)]	0
reshape (Reshape)	(None, 18, 14, 1)	0
conv (QConv2D)	(None, 8, 6, 3)	27
relu1 (QActivation)	(None, 8, 6, 3)	0
flatten (Flatten)	(None, 144)	0
dense1 (QDense)	(None, 20)	2880
relu2 (QActivation)	(None, 20)	0
output (QDense)	(None, 1)	20

Total params: 2,927
Trainable params: 2,927
Non-trainable params: 0

Student

- 300k parameters go down to 3k parameters, while preserving the performance
- Inference latency ~ 100 nanoseconds
- Computational resources fit to a single FPGA board by large margins

Adding to the current triggers

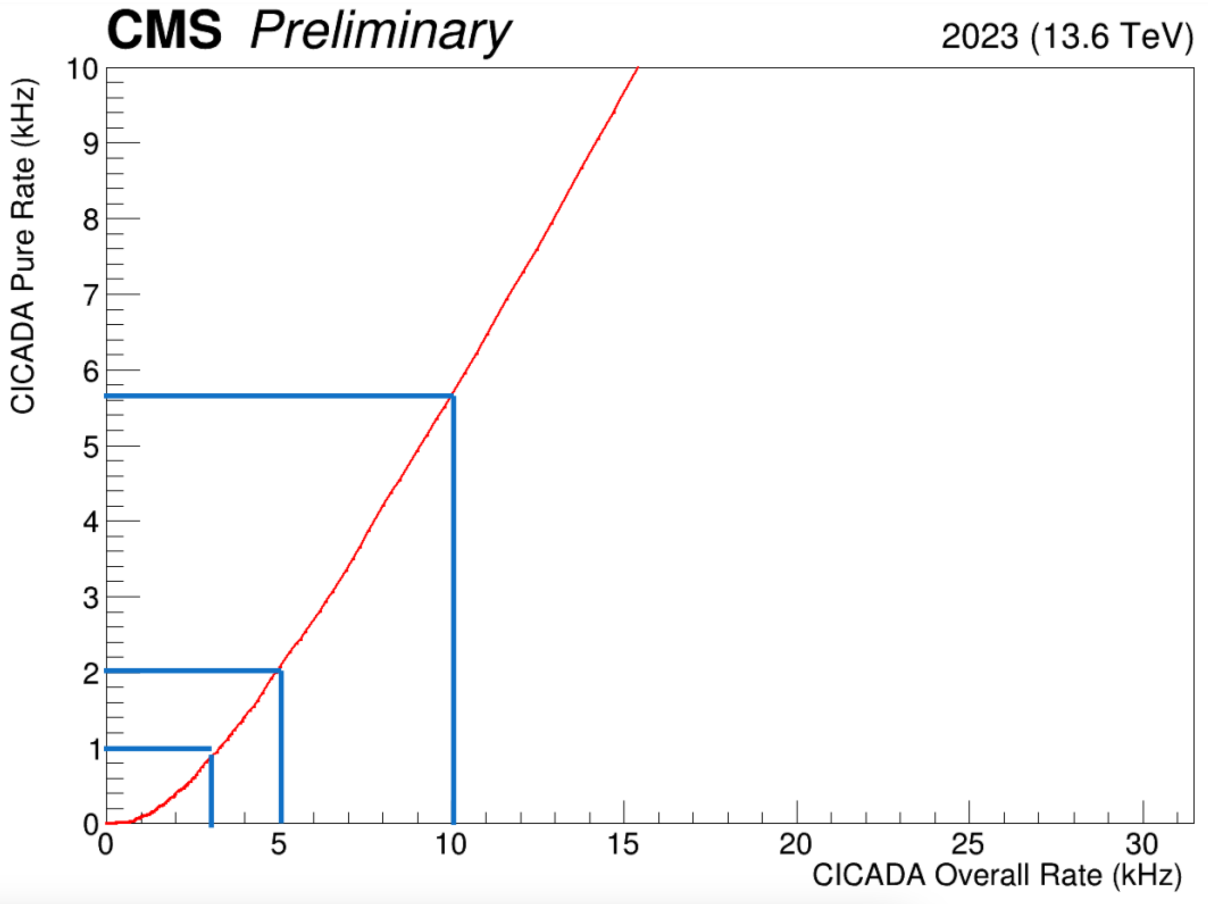

CMS-DP-2023-086

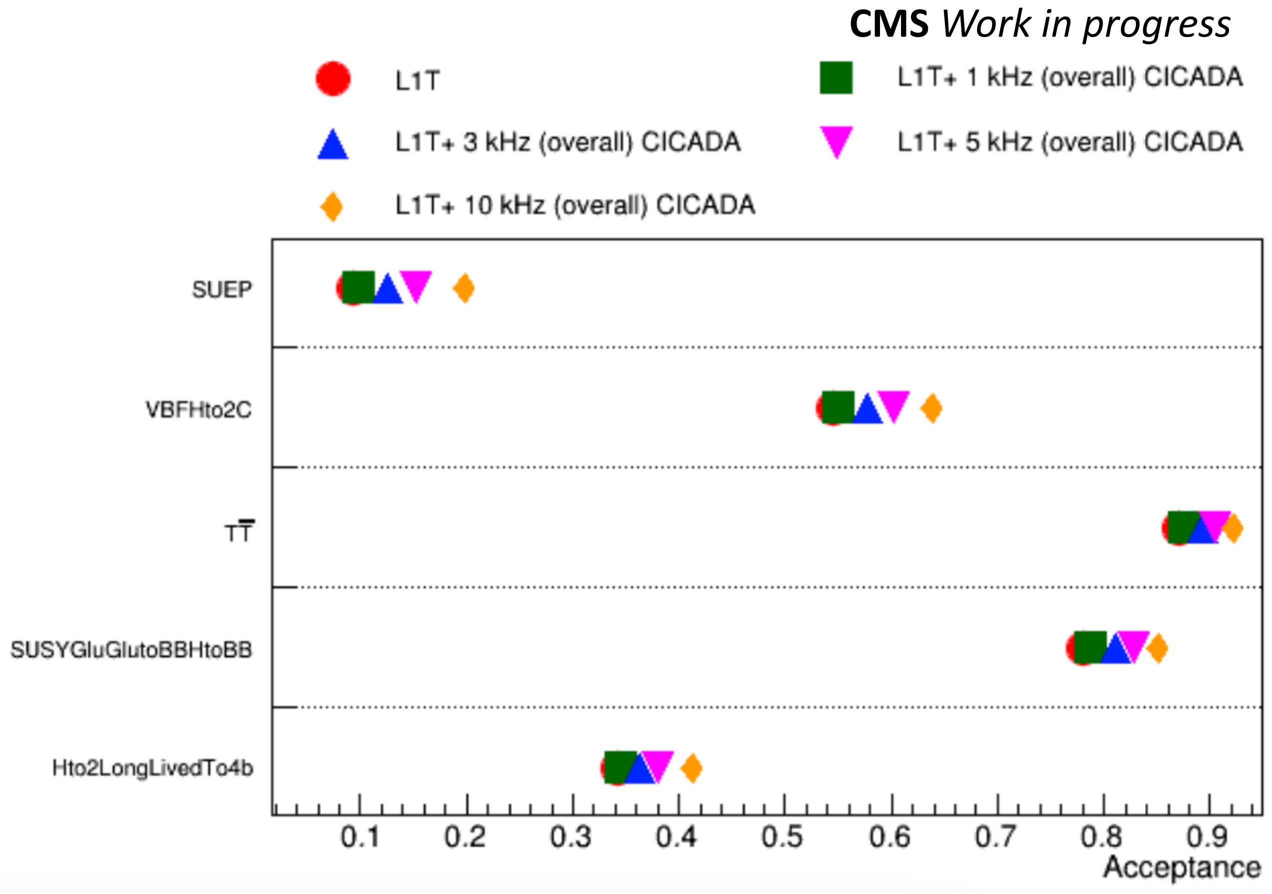

CICADA vs. HT trigger

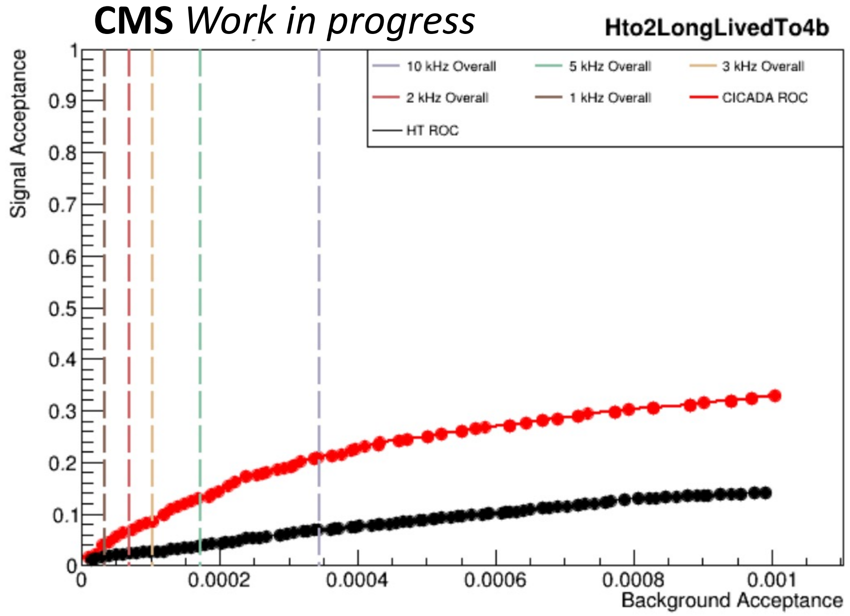

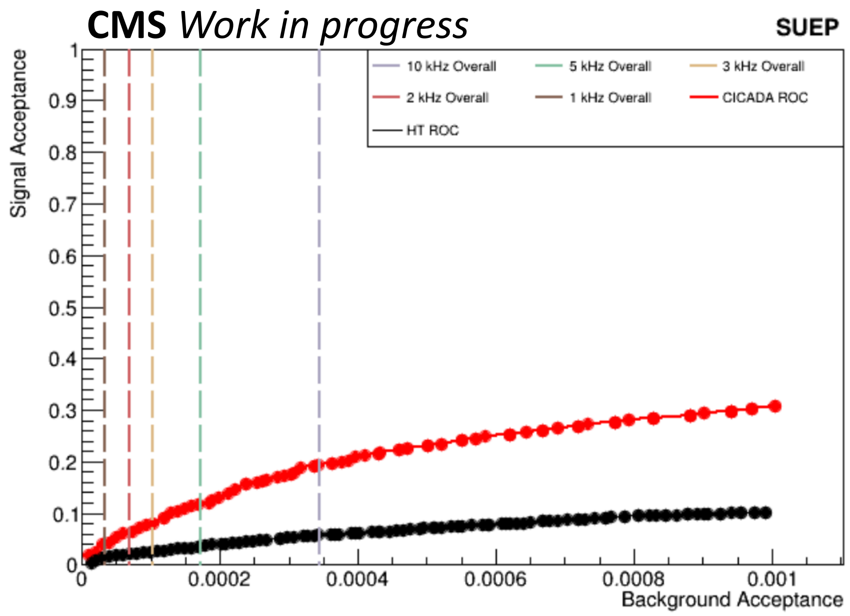

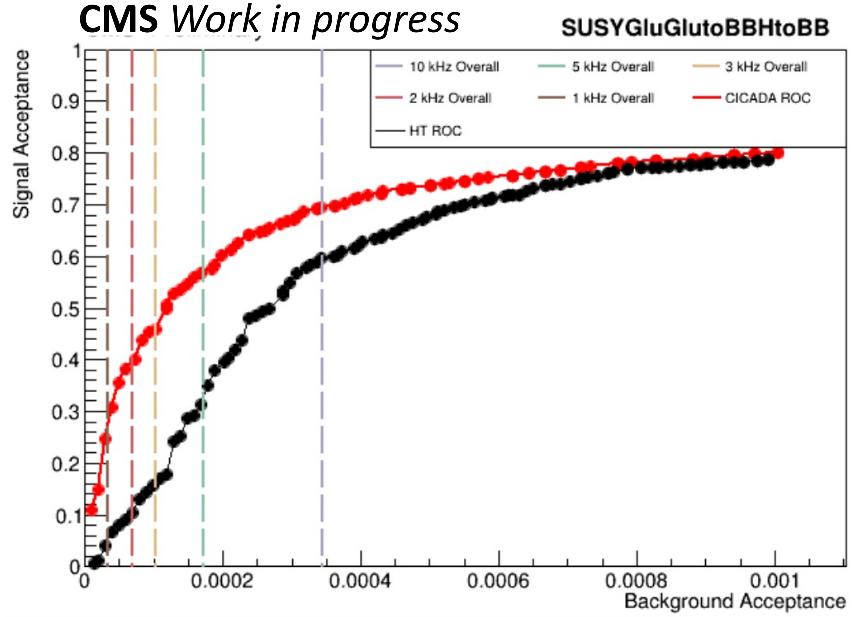

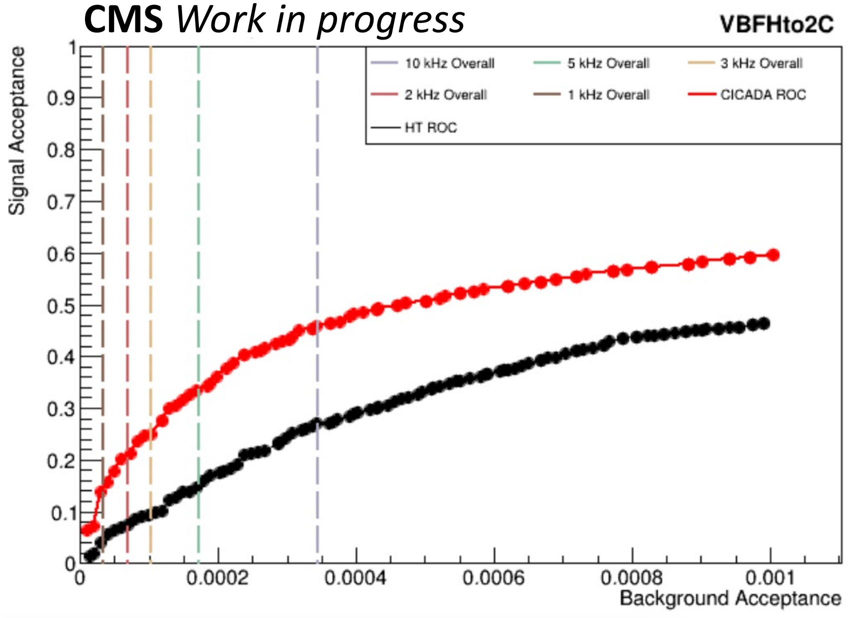

- **HT trigger** vs **CICADA**
- **Region of interest**
  - Bkg rate at $O(10^{-4})$, or
  - Trigger rate at O(1) kHz

Preserving signal shapes

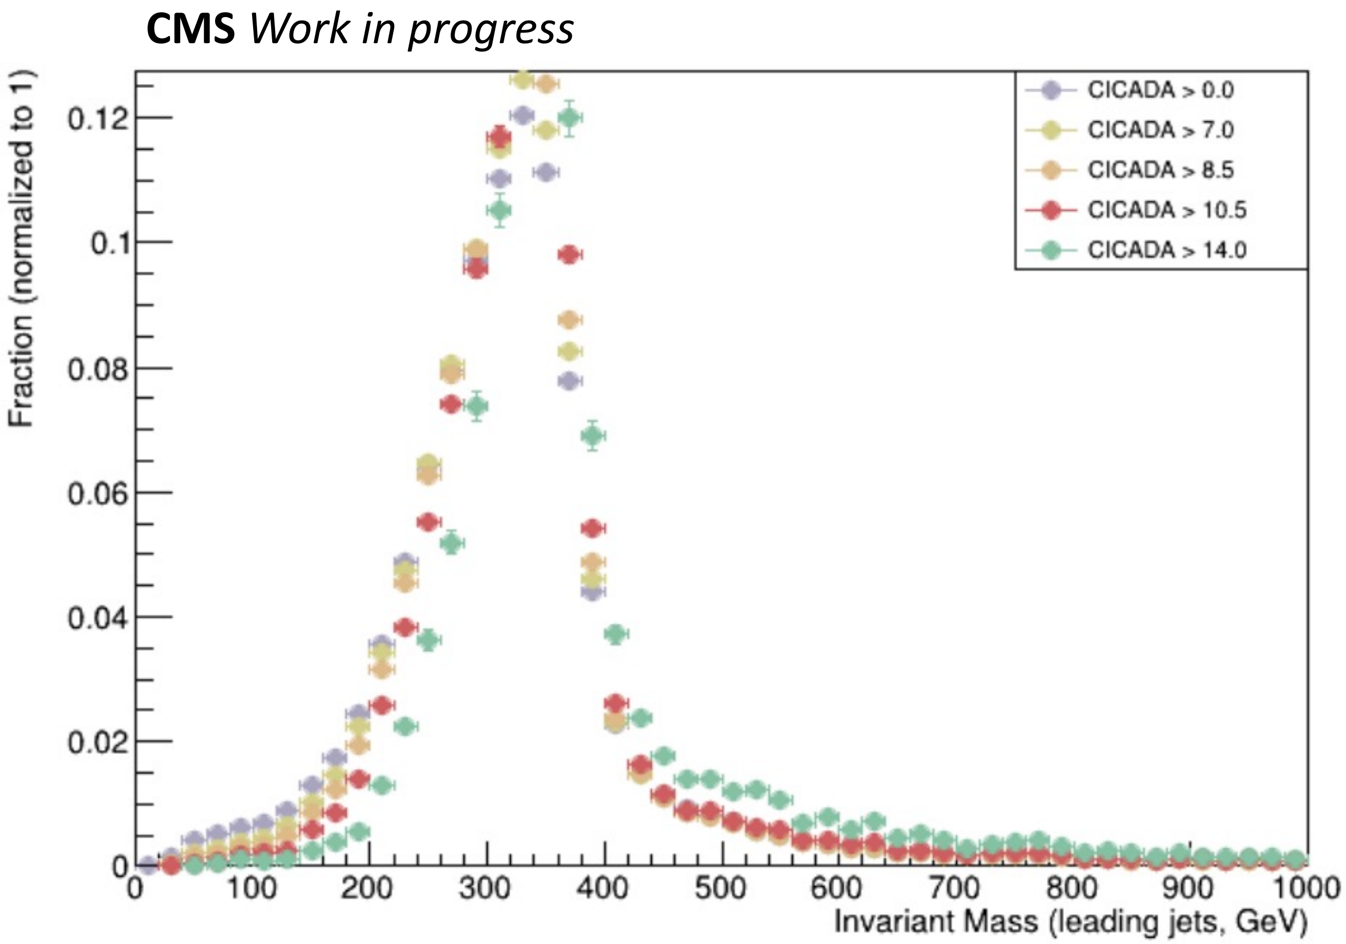

E.g.
/SUSYGluGlutoBBHtoBB_NarrowWidth_M-350_TuneCP5_13p6TeV_pythia8/Run3Winter23MiniAOD126X_mcRun3_2023_forPU65_v1-v2/MINIAODSIM

- **Invariant mass of 2 leading jets**

Some trigger objects (EGamma)

nObjects

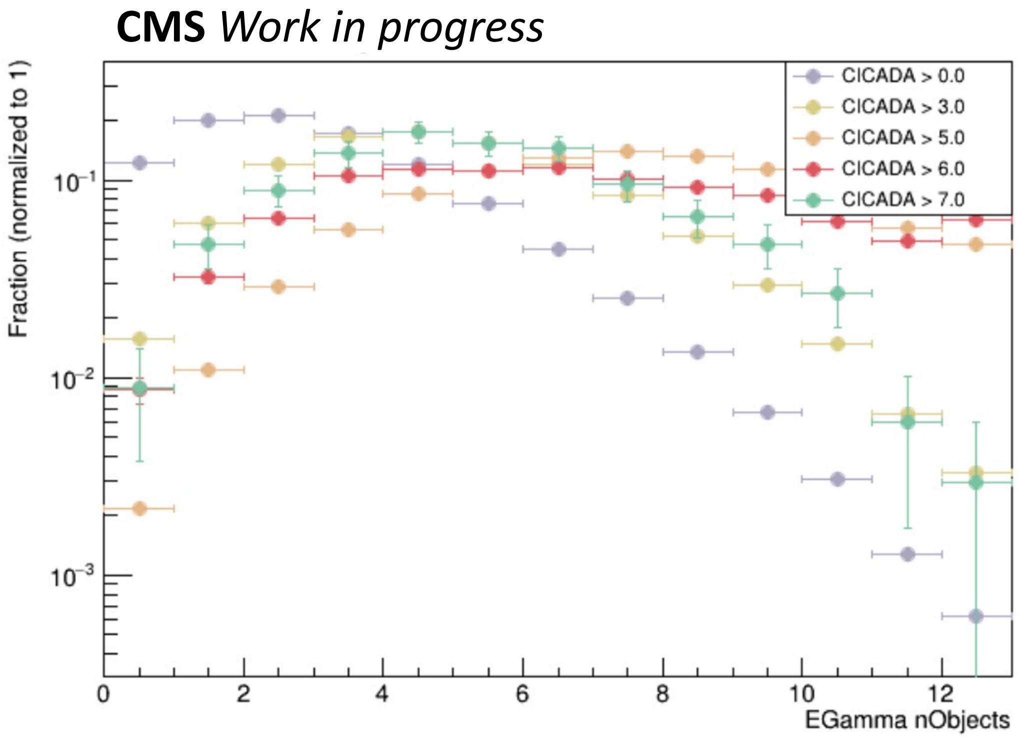

- CICADA has preference for more EG objects
  - 3 or more

Pt

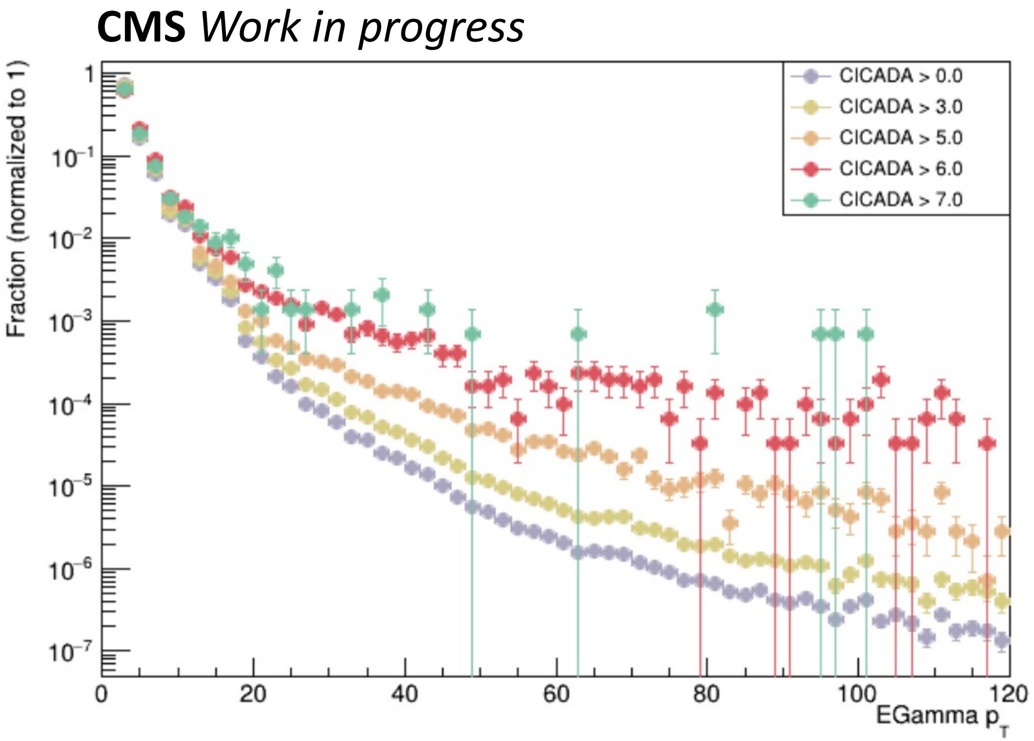

- Slight preference for higher Pt

- Still sensitive to the low Pt objects as well

Some trigger objects (Jet)

nObjects

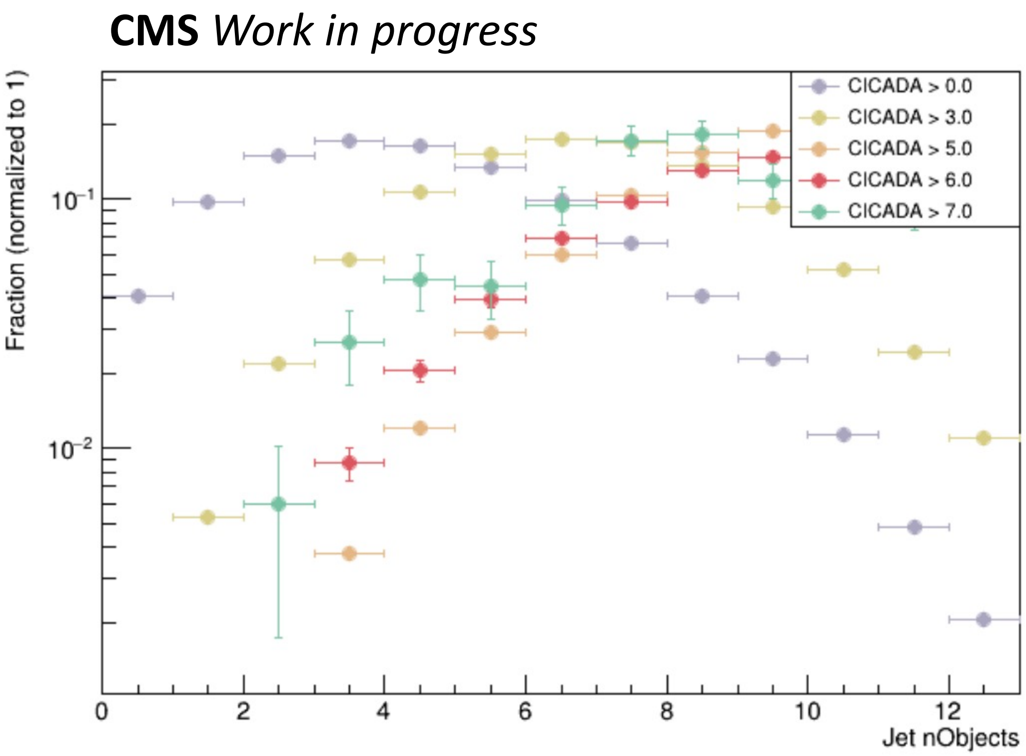

- CICADA's most consistent preference is for high jet multiplicity

Pt

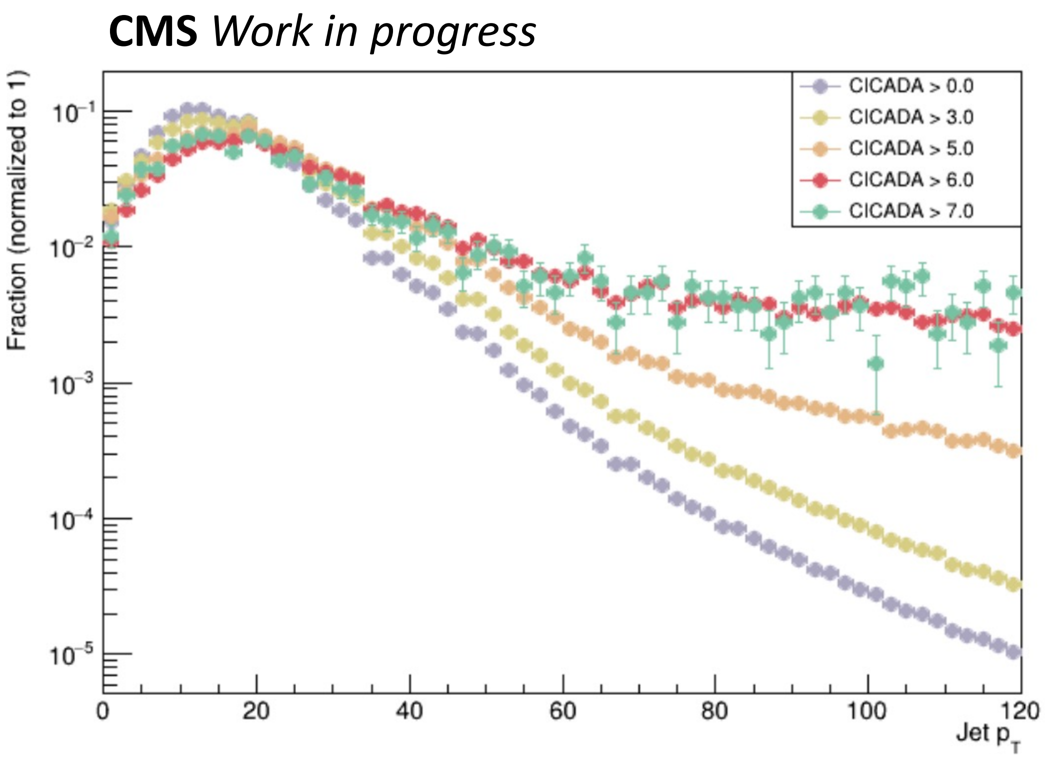

- Slight preference for higher Pt
  - Still sensitive to the low Pt objects as well

Some trigger object (Tau)

nObjects

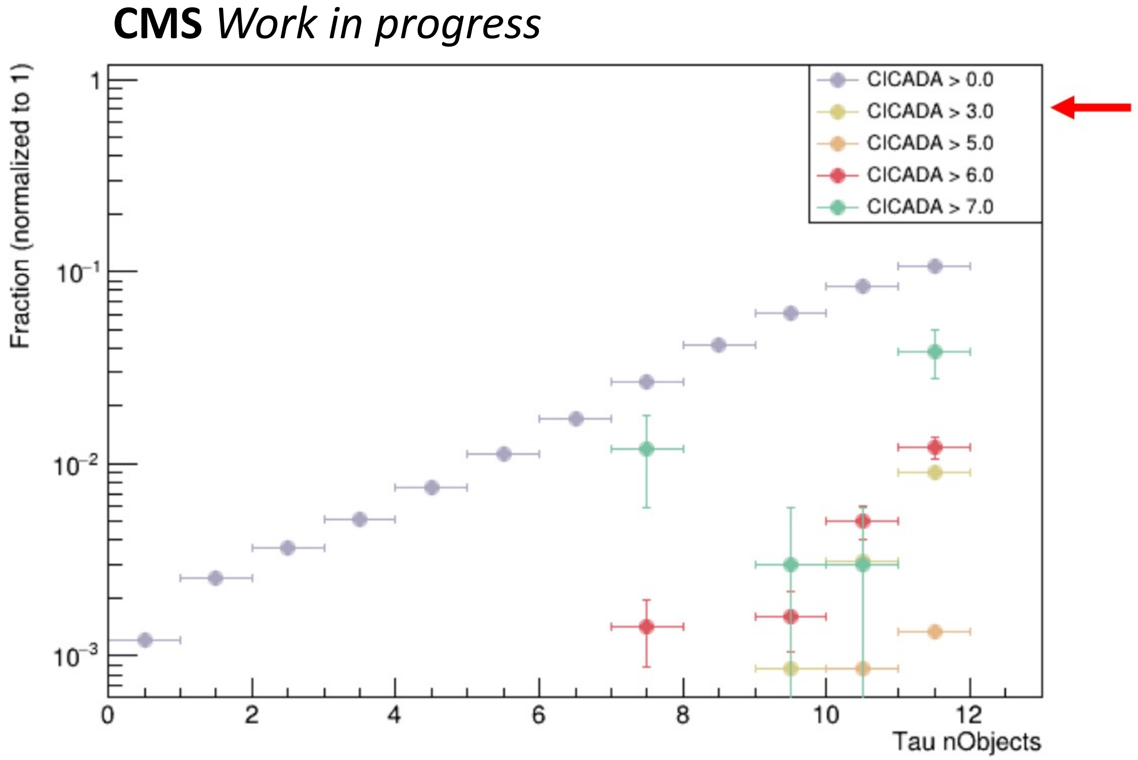

- If this looks a bit confusing, it's bad legend placement
- Nearly 97% of high score CICADA events have 12 trigger taus

Pt

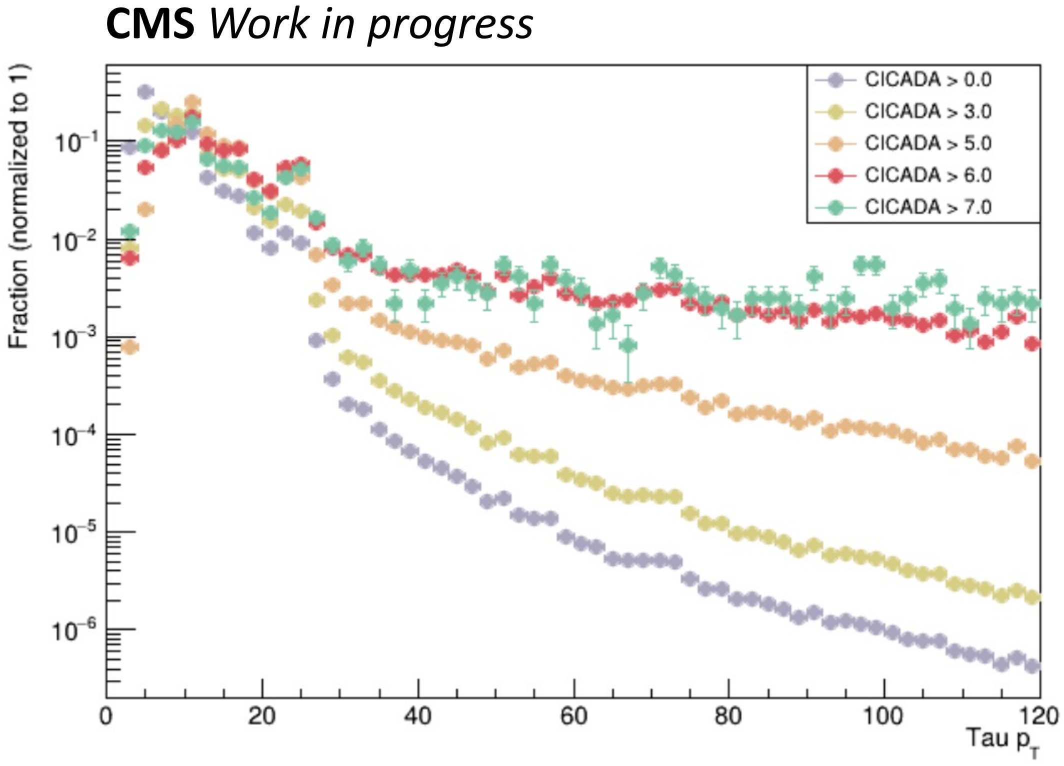

- Similar to other objects, there are slight preferences for higher Pt
  - Still sensitive to the low Pt objects as well