CICADA: Real-Time Anomaly Detection with Event Images in the CMS Level-1 Trigger

Kiley Kennedy, Princeton University, on behalf of the CICADA Team

AD4HEP Workshop, Nevis Labs, Columbia University

June 16, 2025

Title image: DALL-E

Introduction

CMS Level 1 Trigger system discards 99.75% of collision events produced at the LHC

Introduction

CMS Level 1 Trigger system discards 99.75% of collision events produced at the LHC

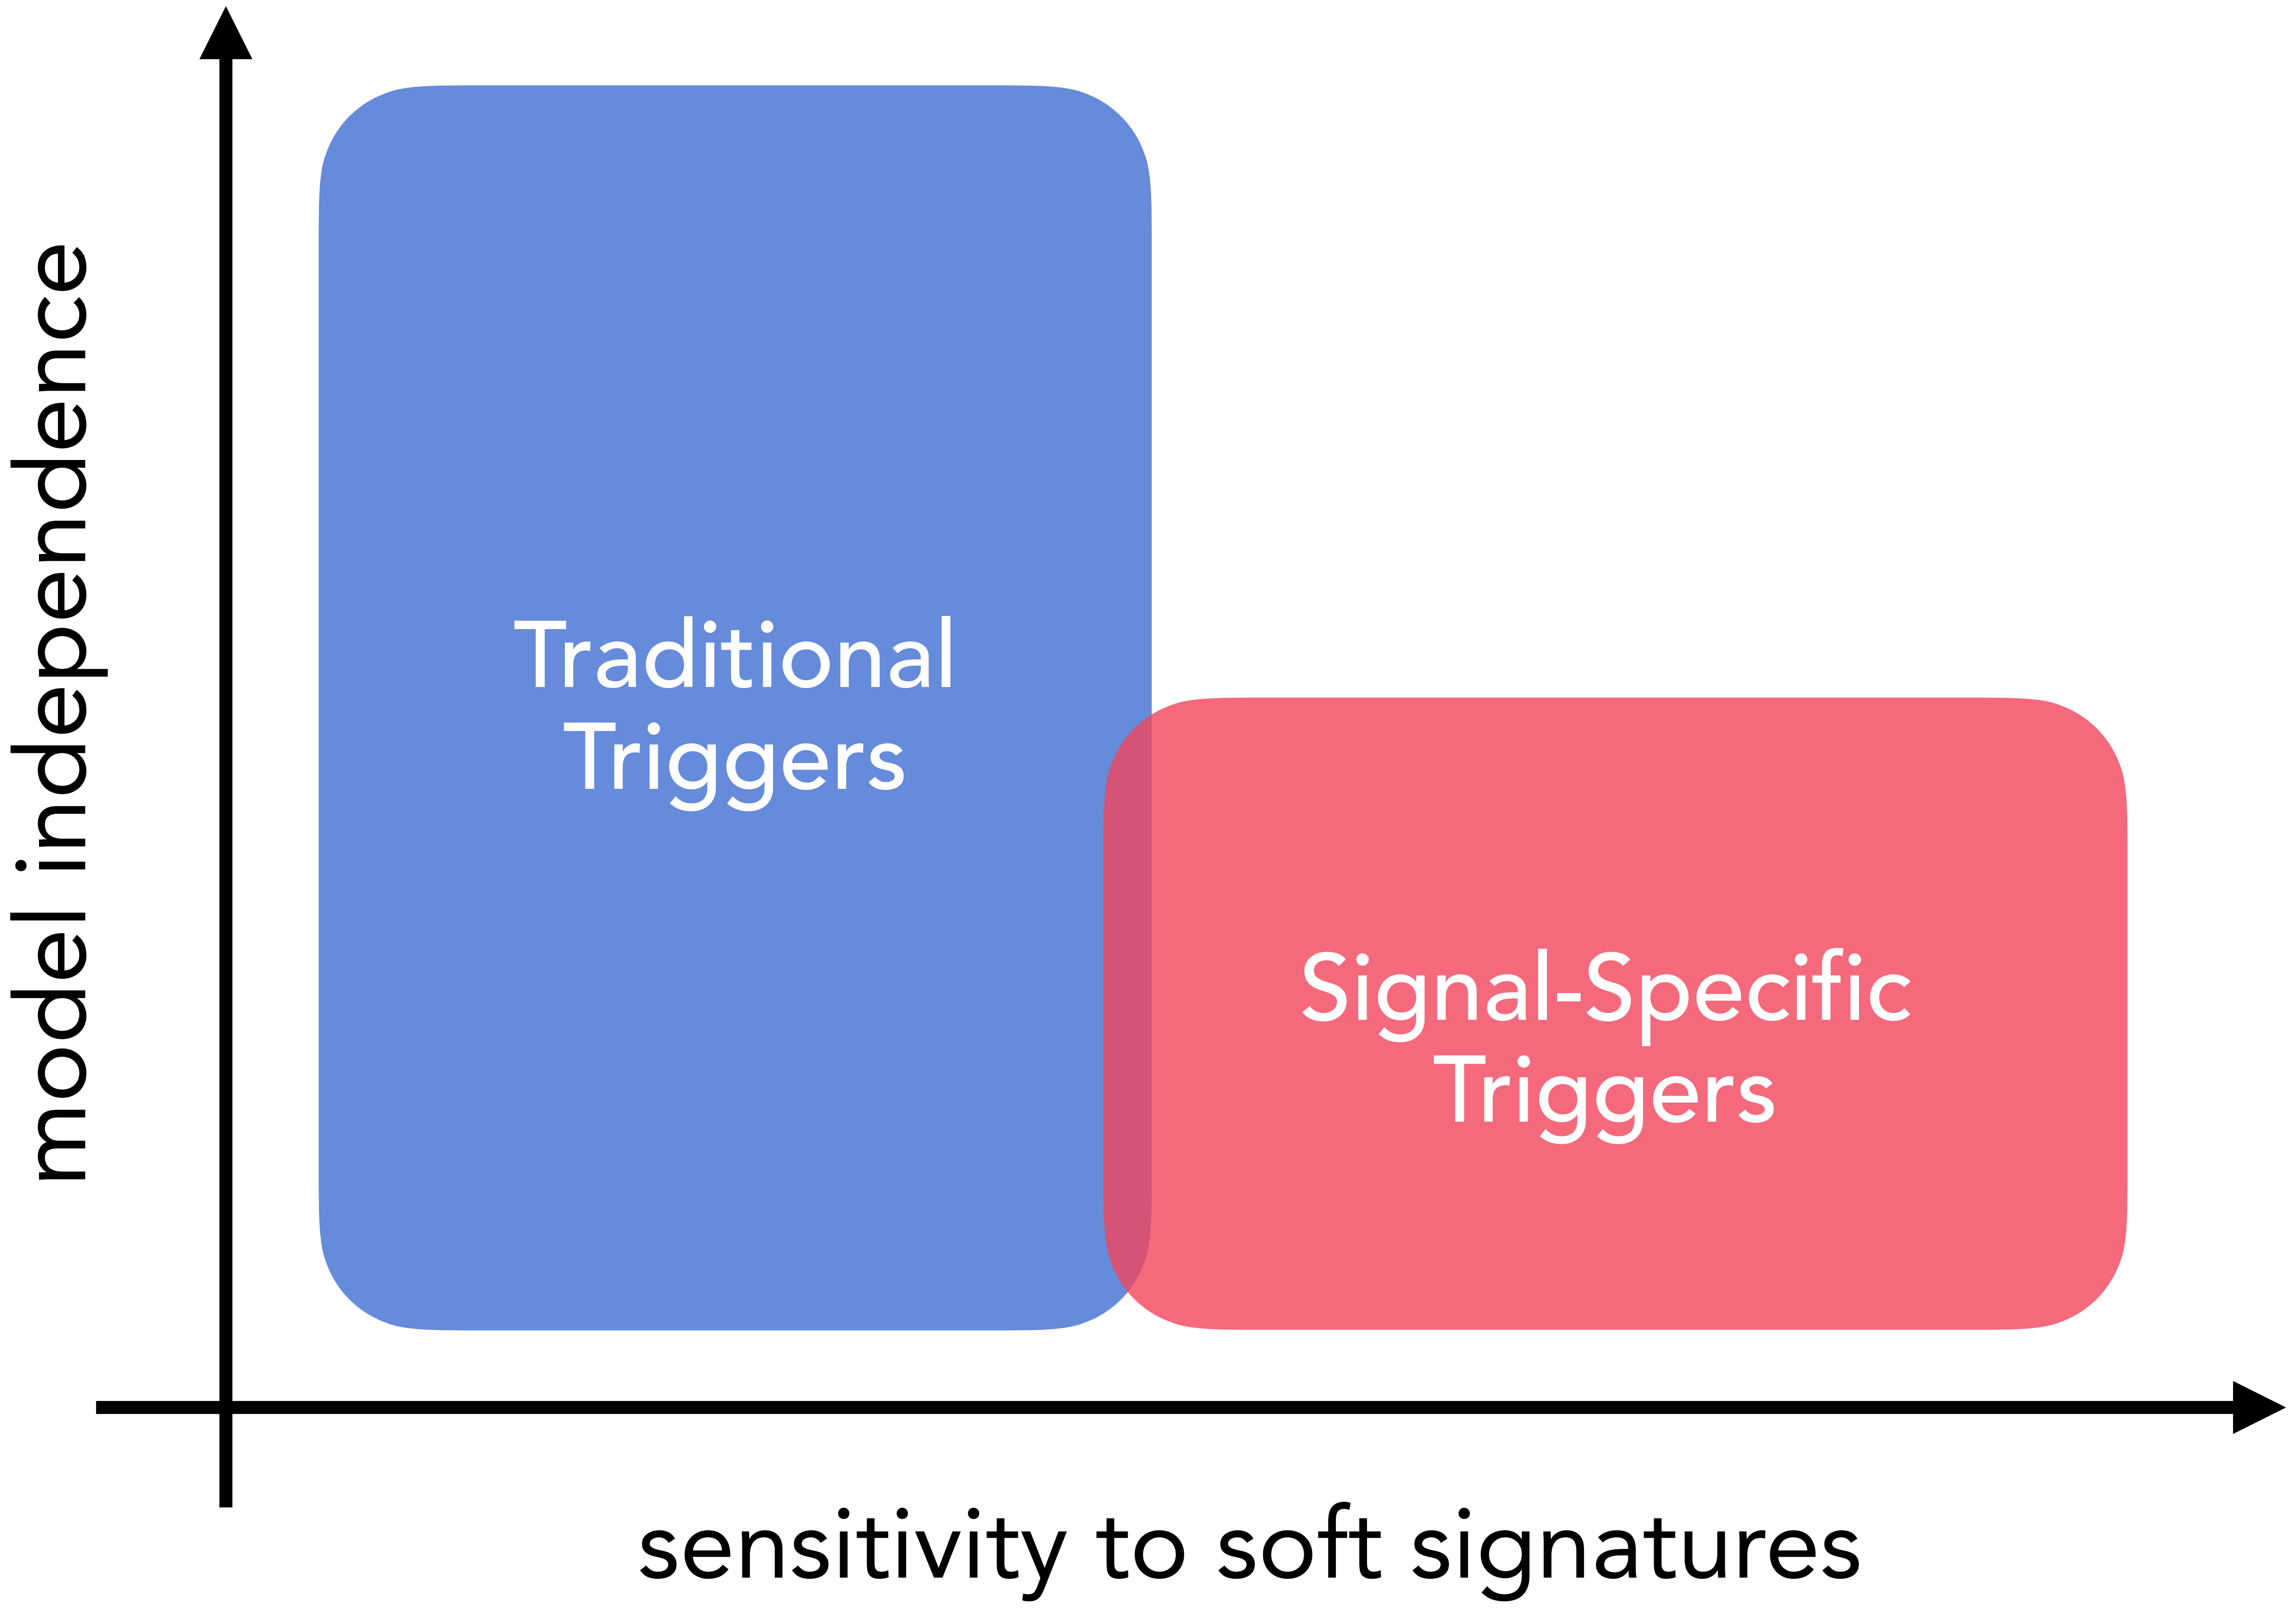

Introduction

CMS Level 1 Trigger system discards 99.75% of collision events produced at the LHC

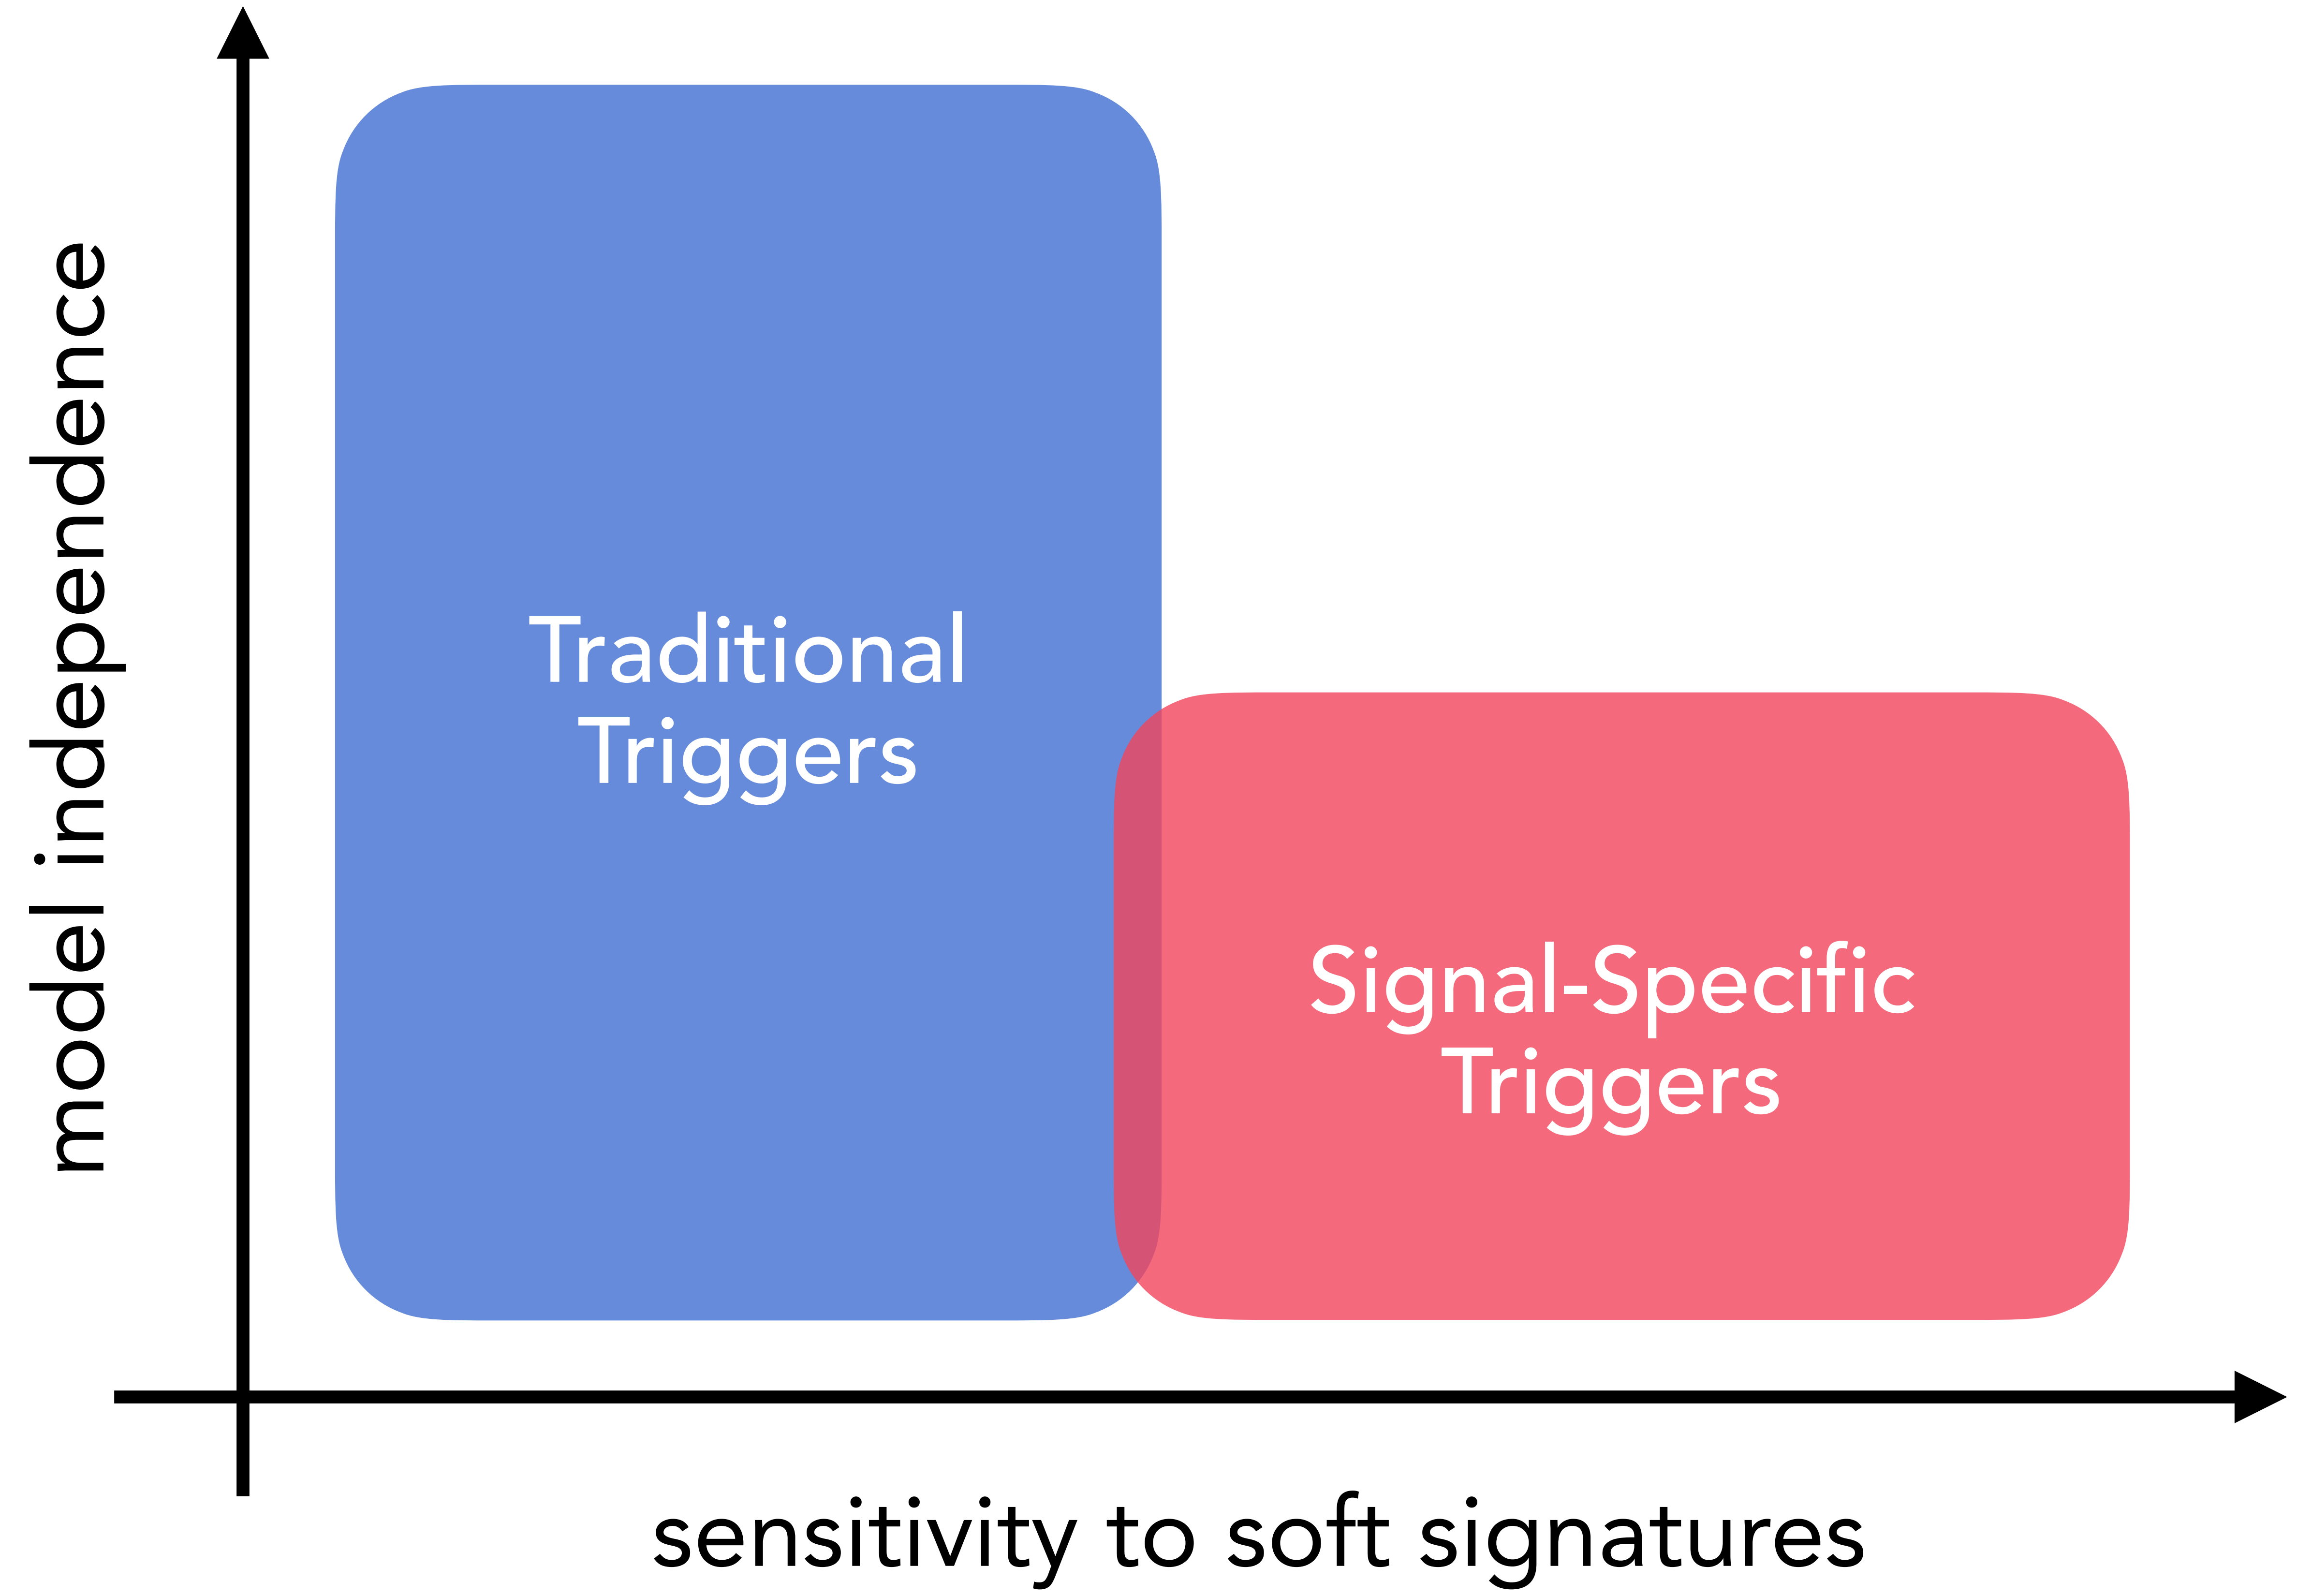

What if we miss new physics because we aren't saving the right events?

Introduction

CMS Level 1 Trigger system discards 99.75% of collision events produced at the LHC

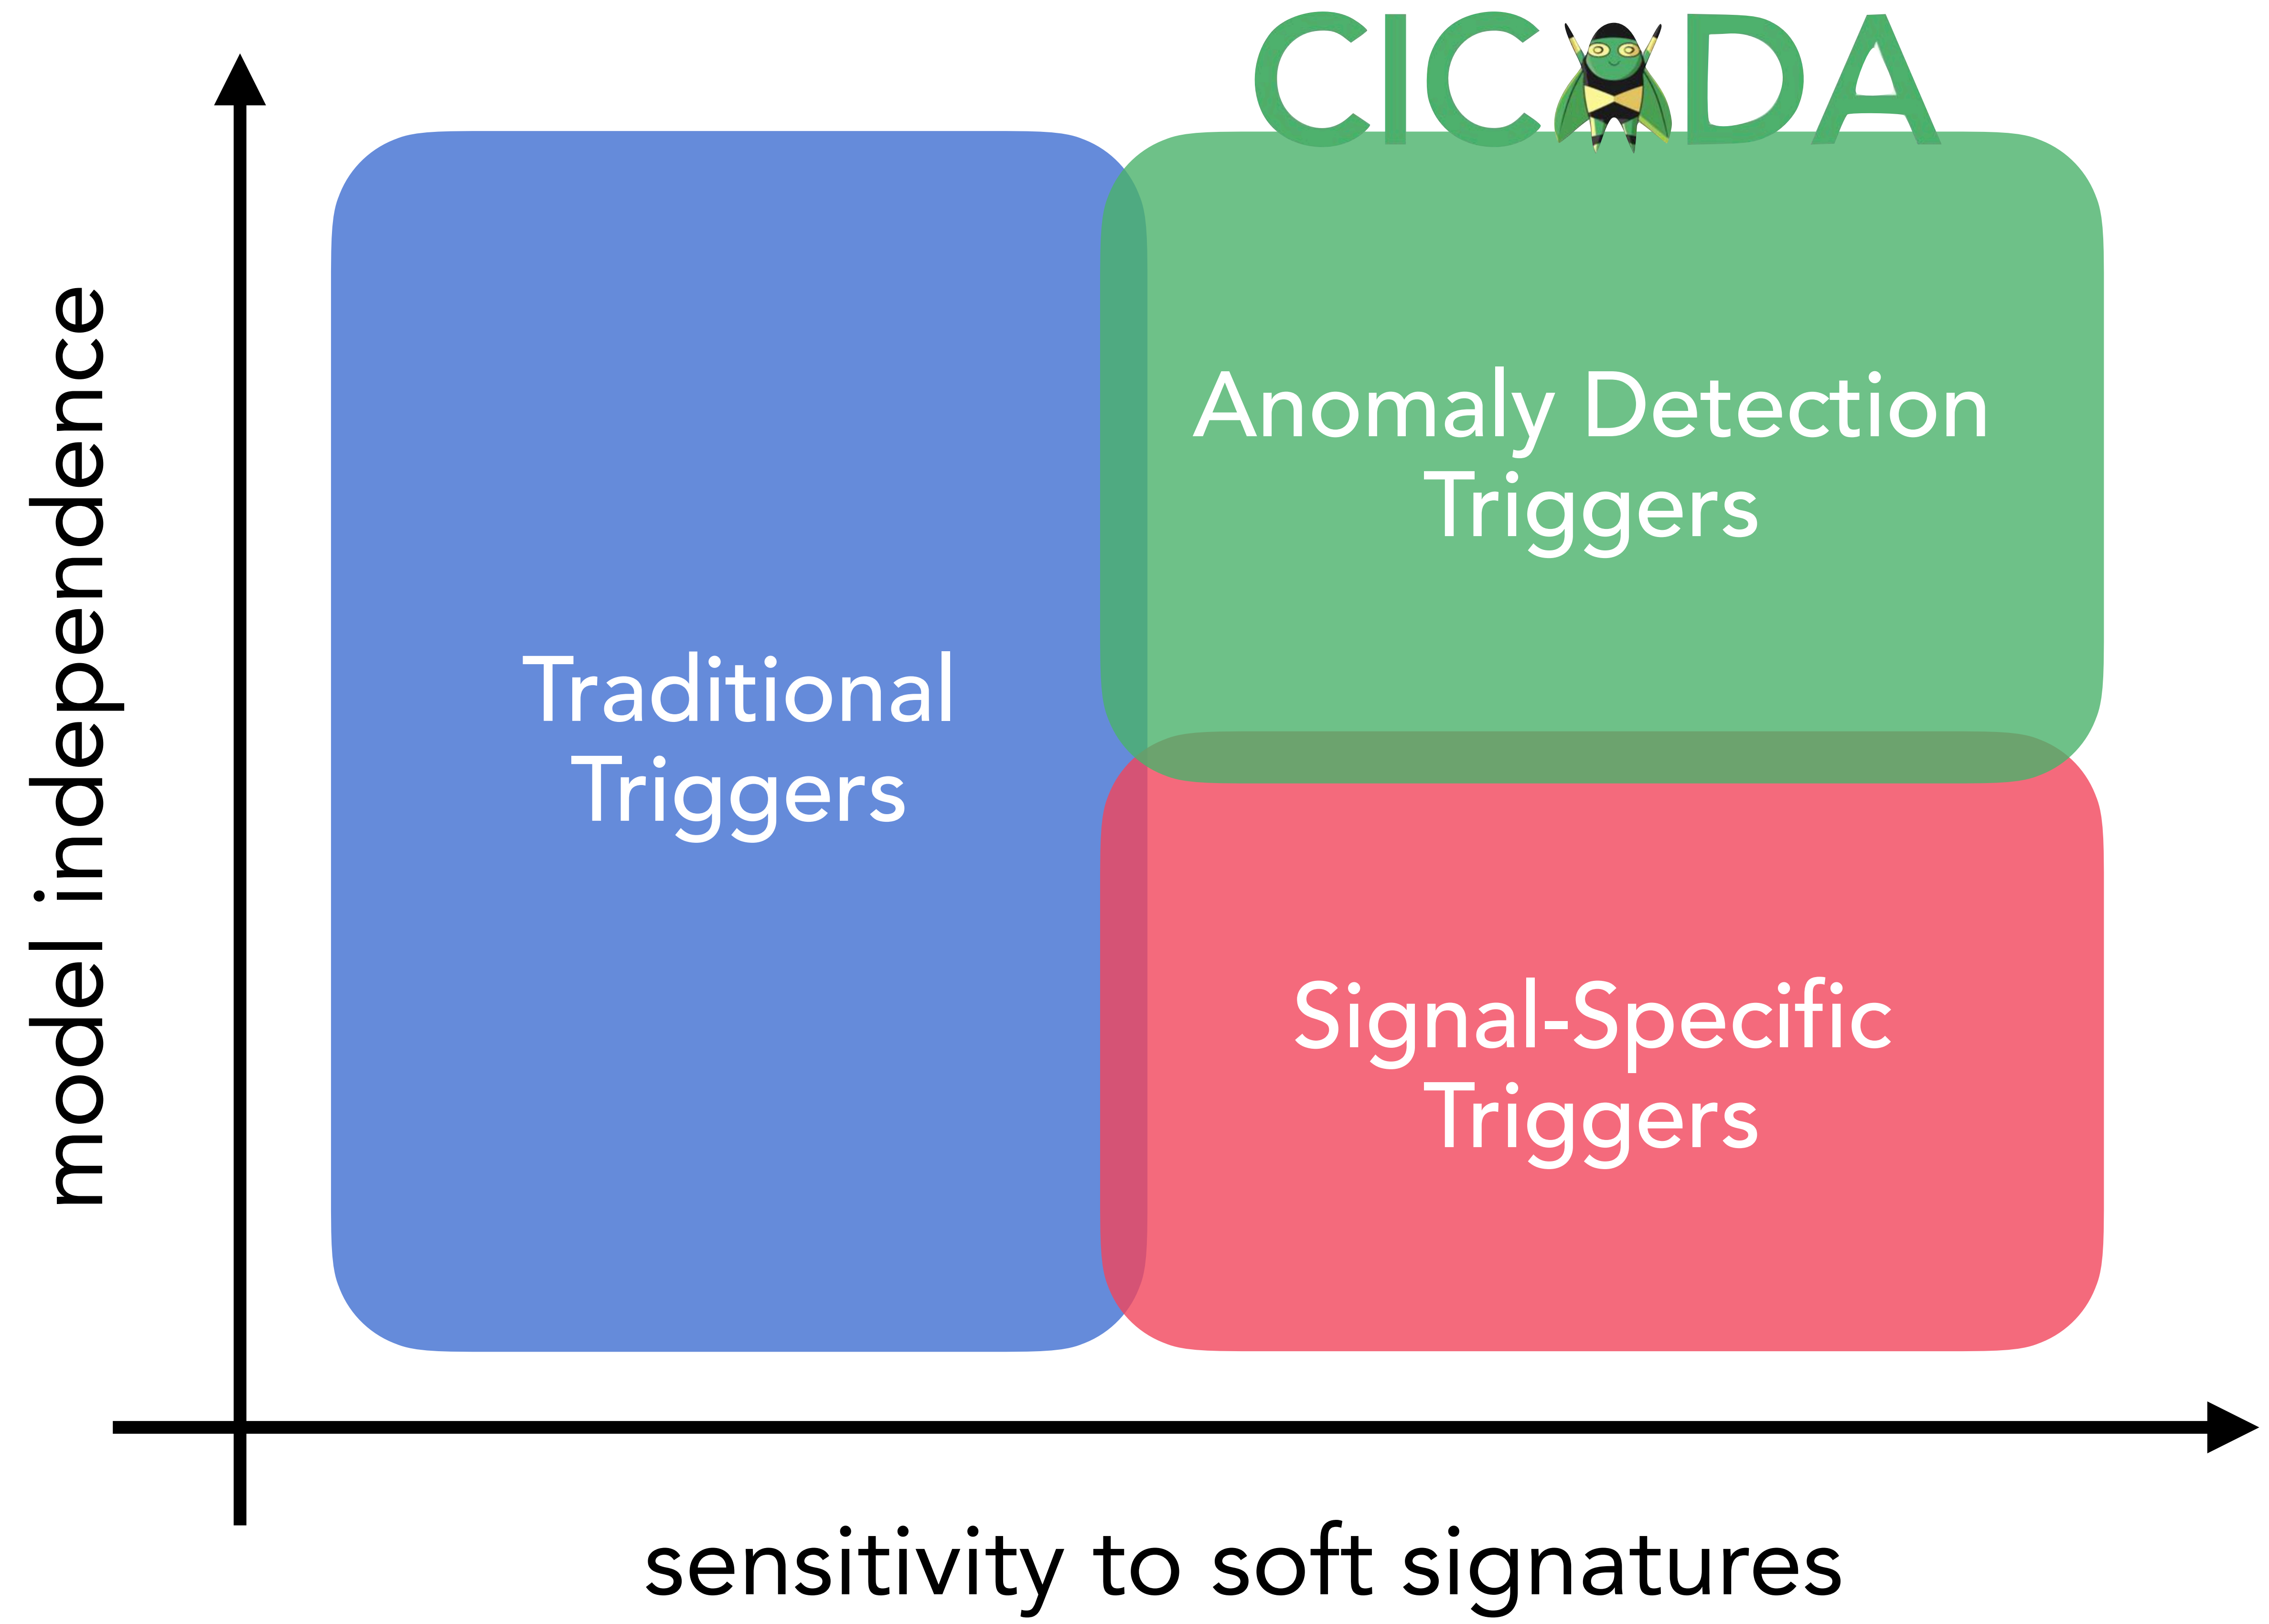

What if we miss new physics because we aren't saving the right events?

CICADA is one of two anomaly detection algorithms currently online in the CMS trigger

Calorimeter Image Convolutional Anomaly Detection Algorithm

The CMS Level 1 Trigger

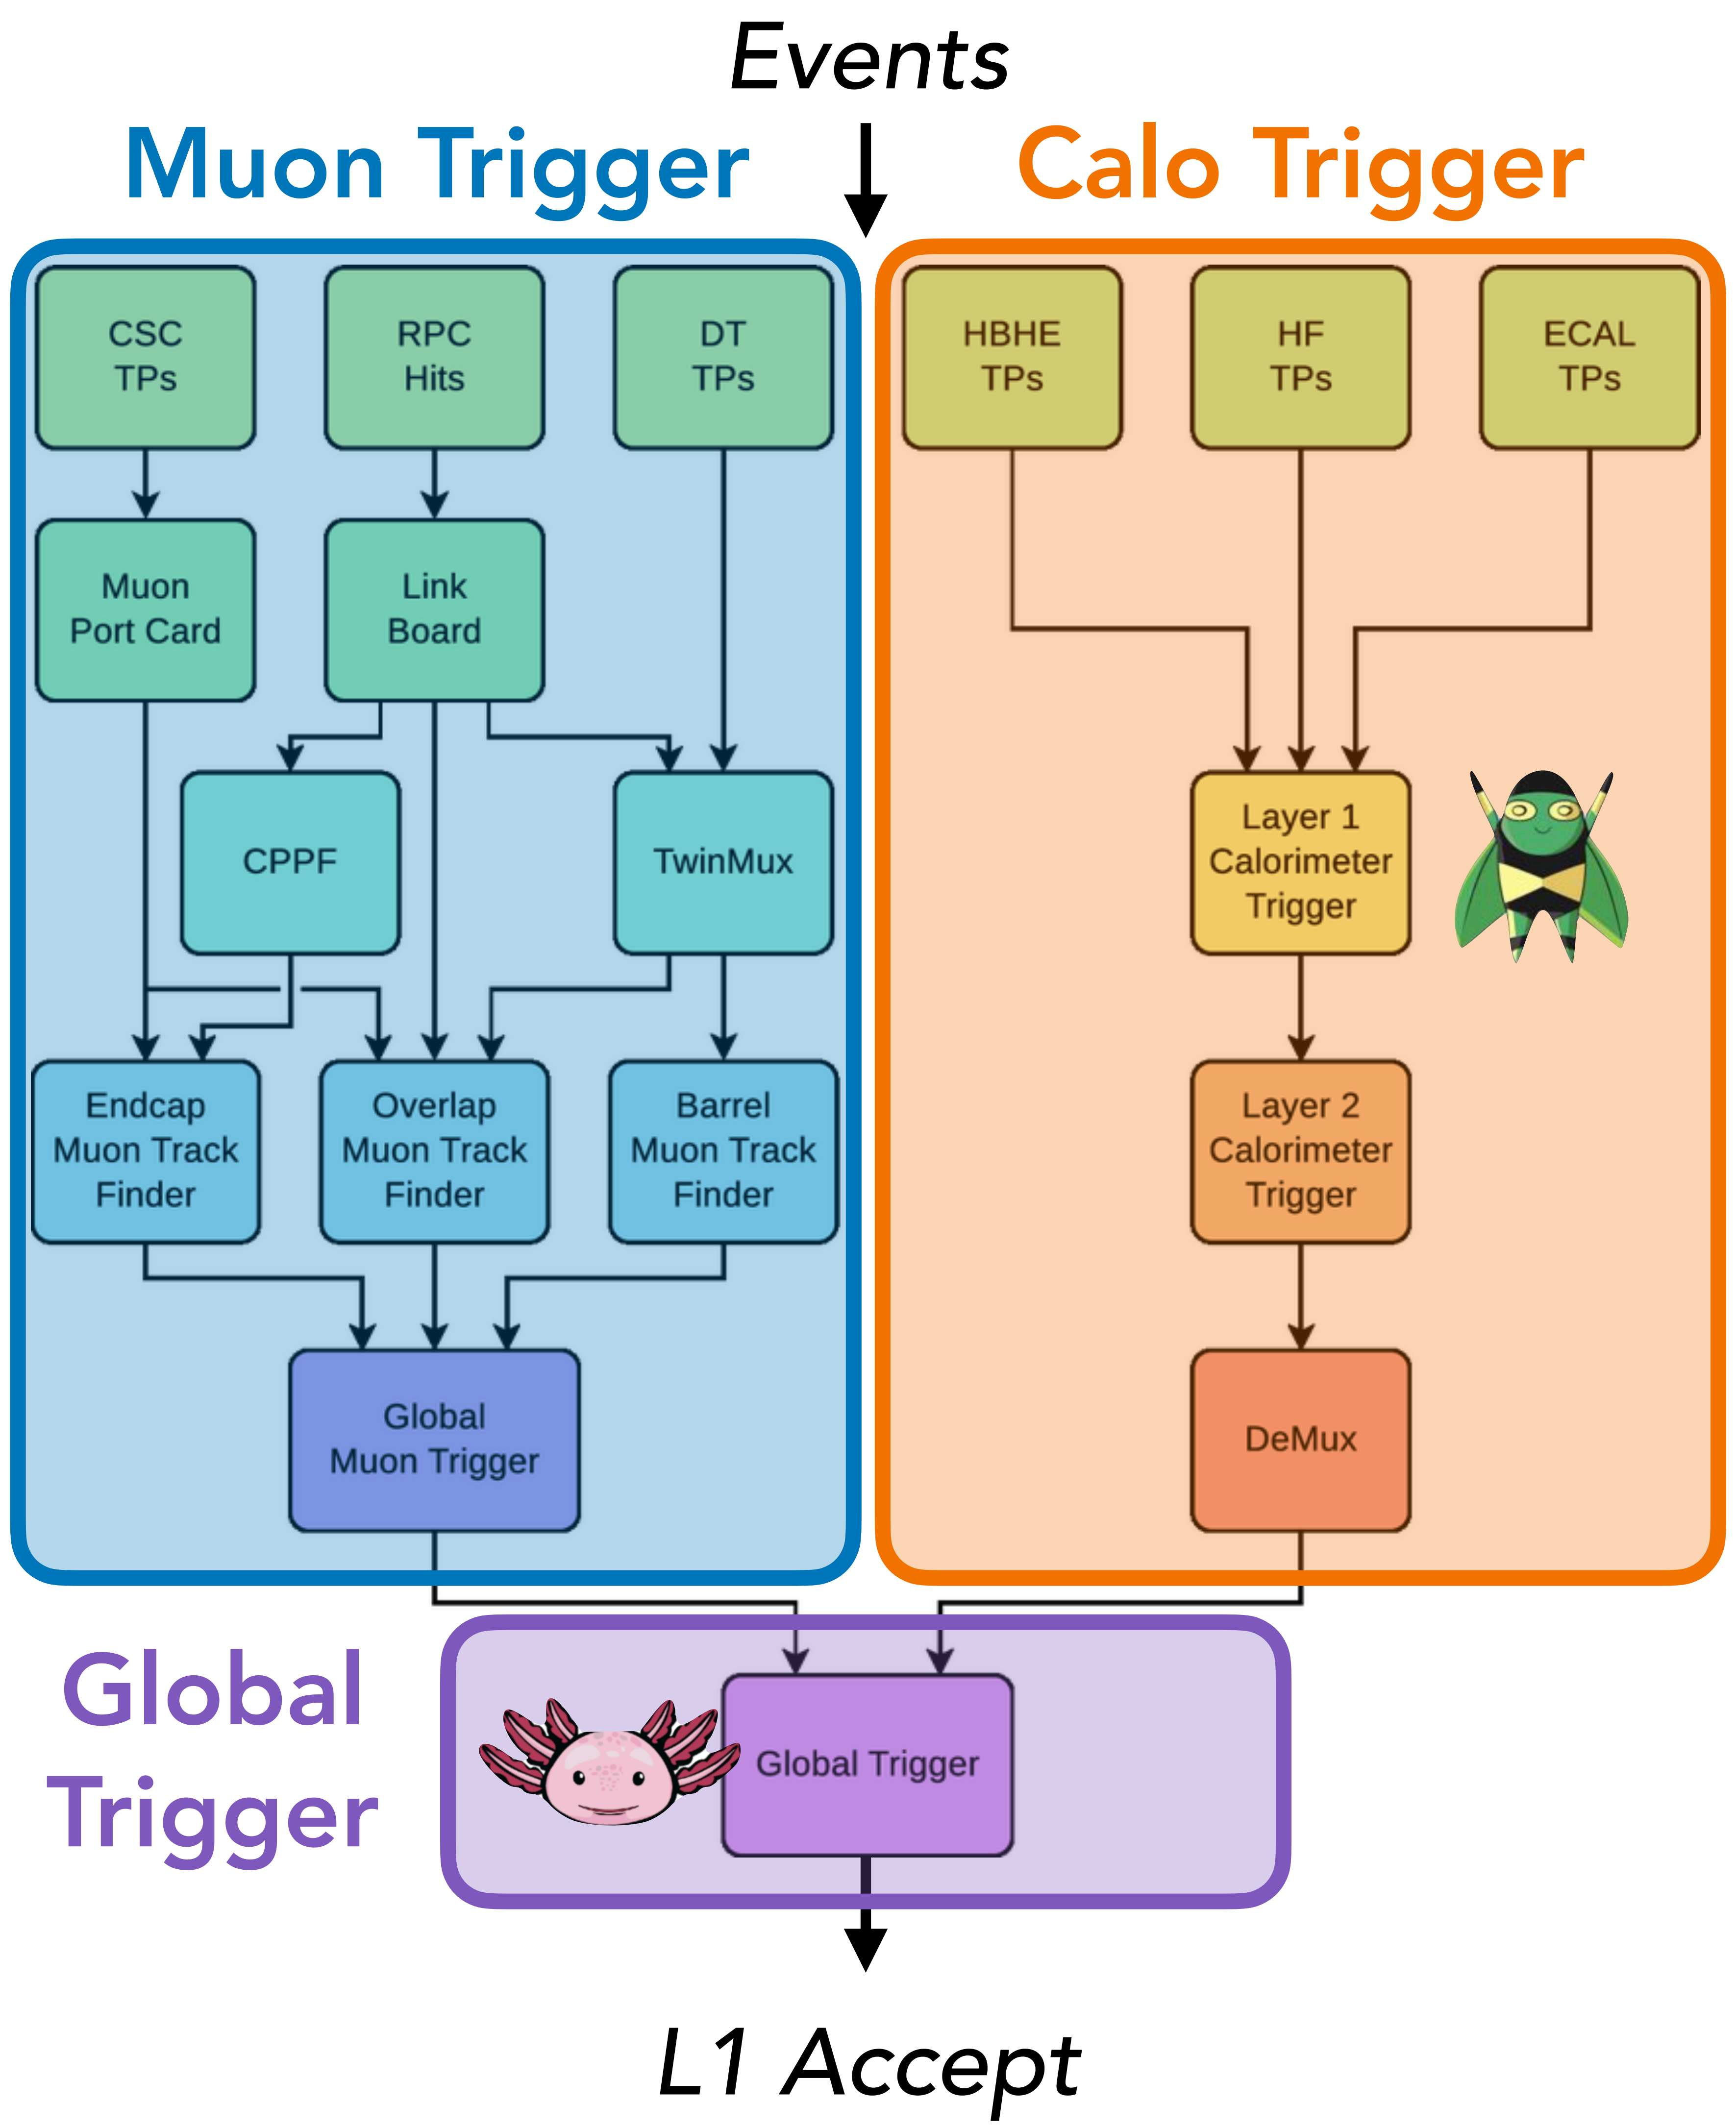

The CMS Level 1 Trigger

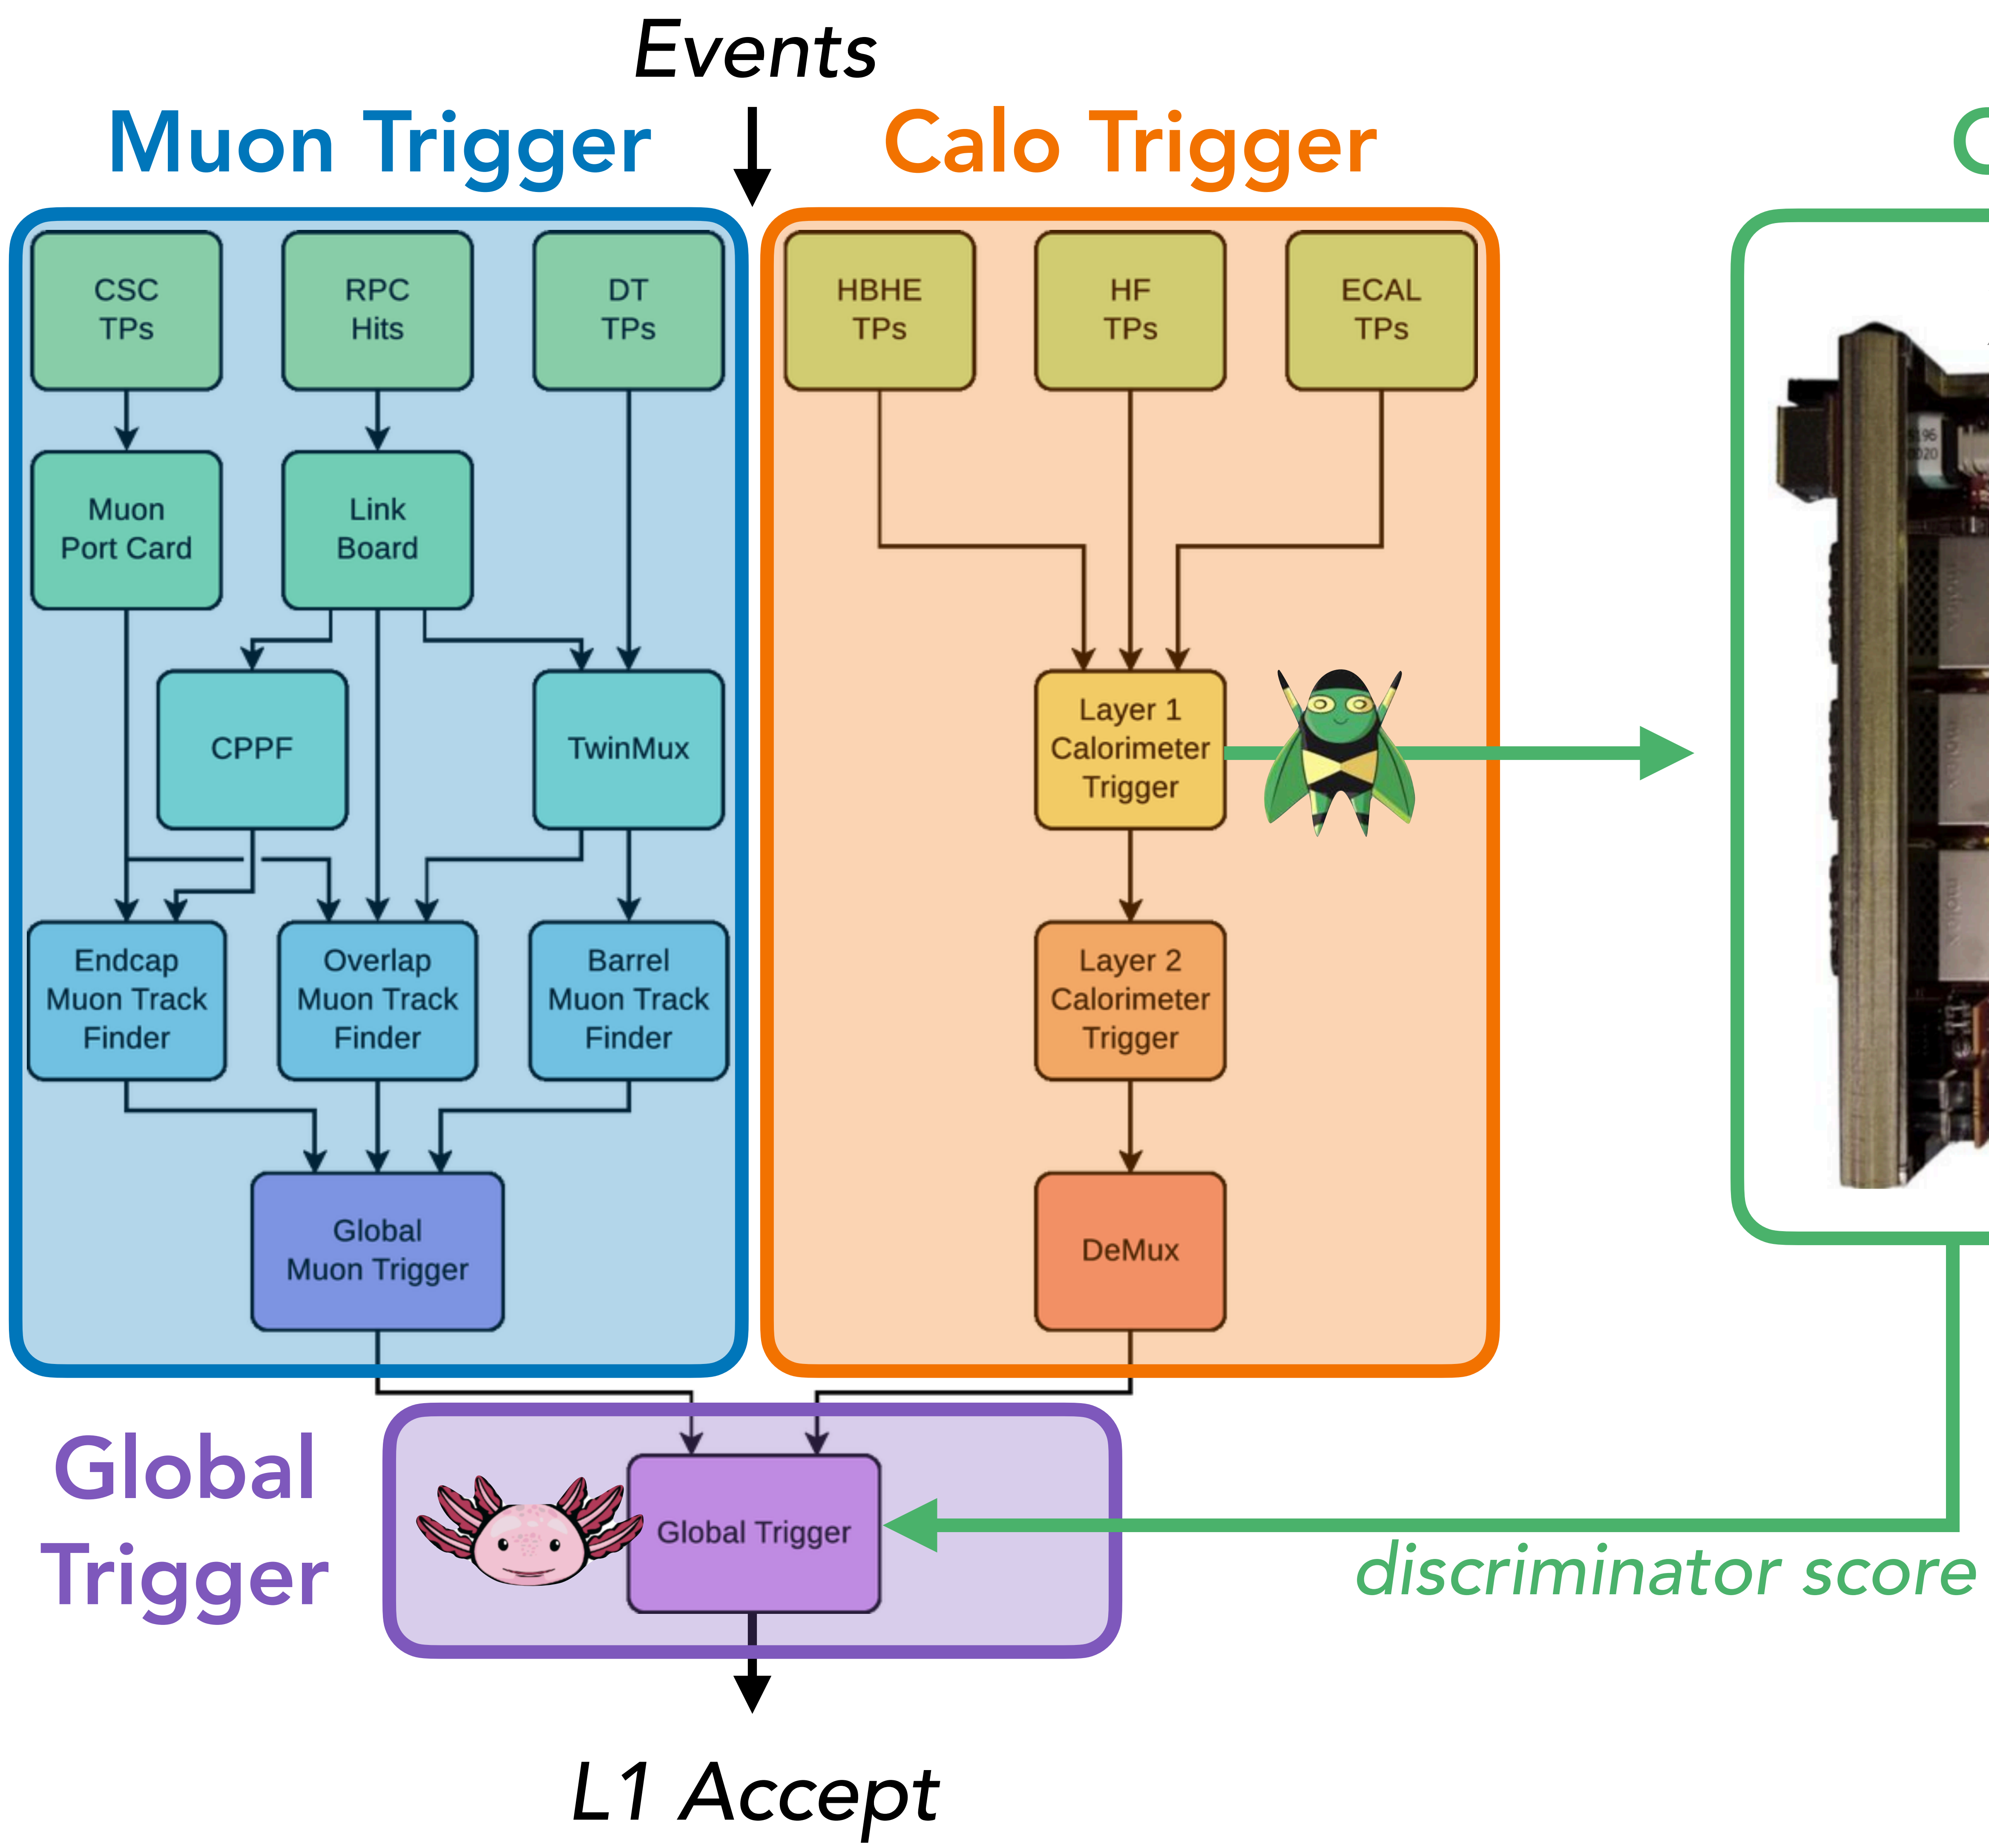

CICADA uses calorimeter system inputs

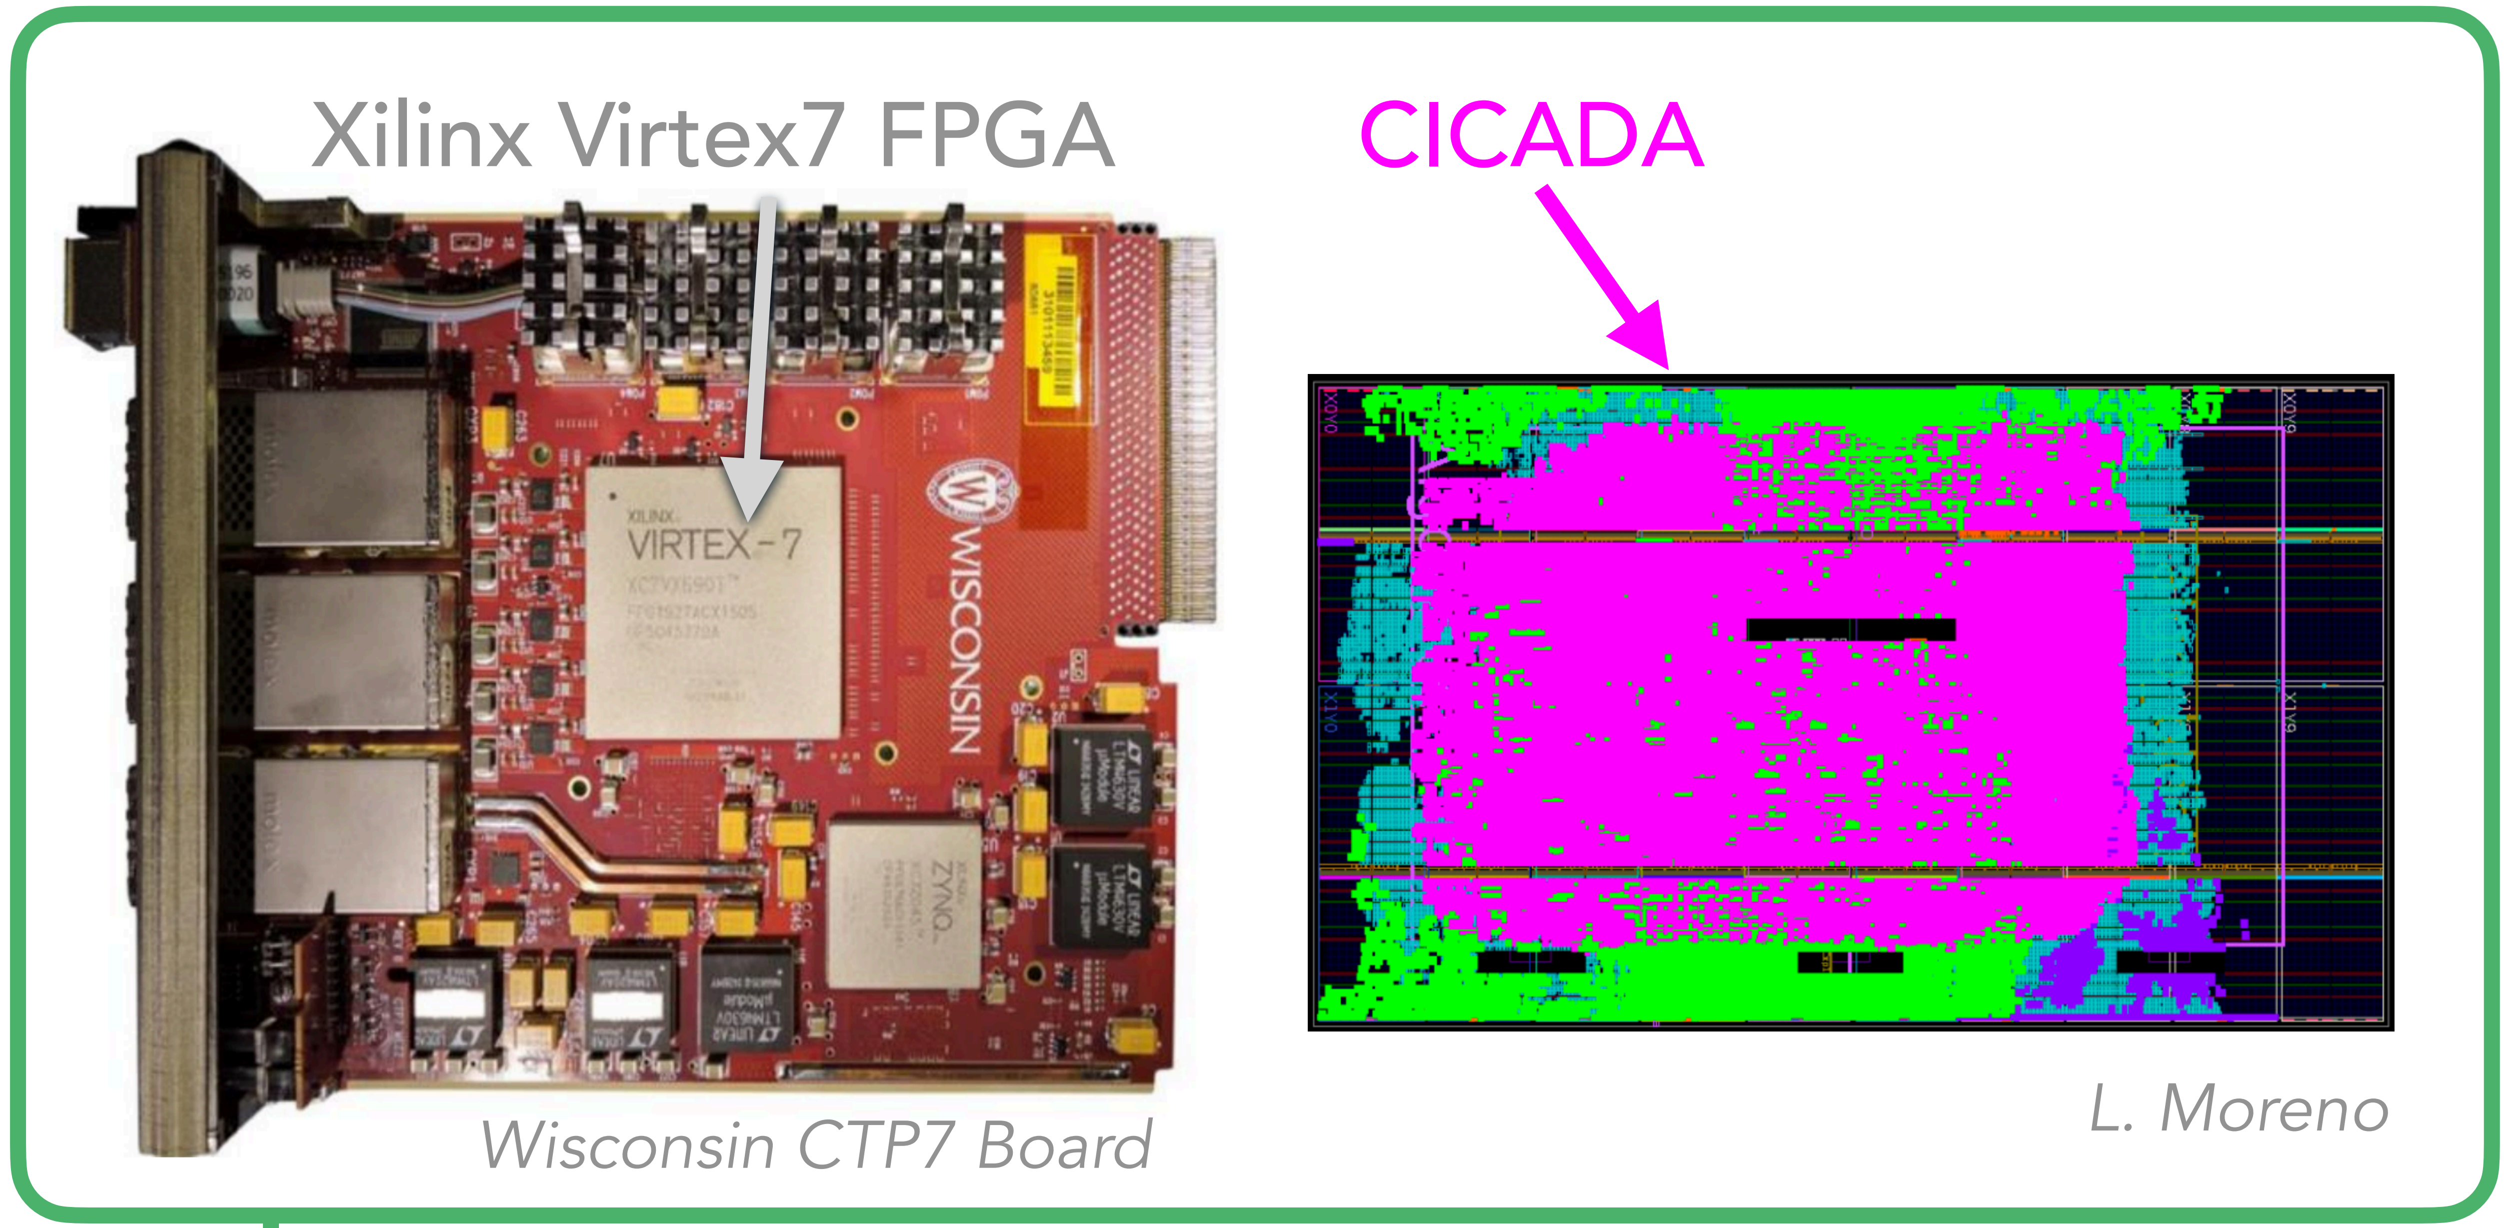

The CMS Level 1 Trigger

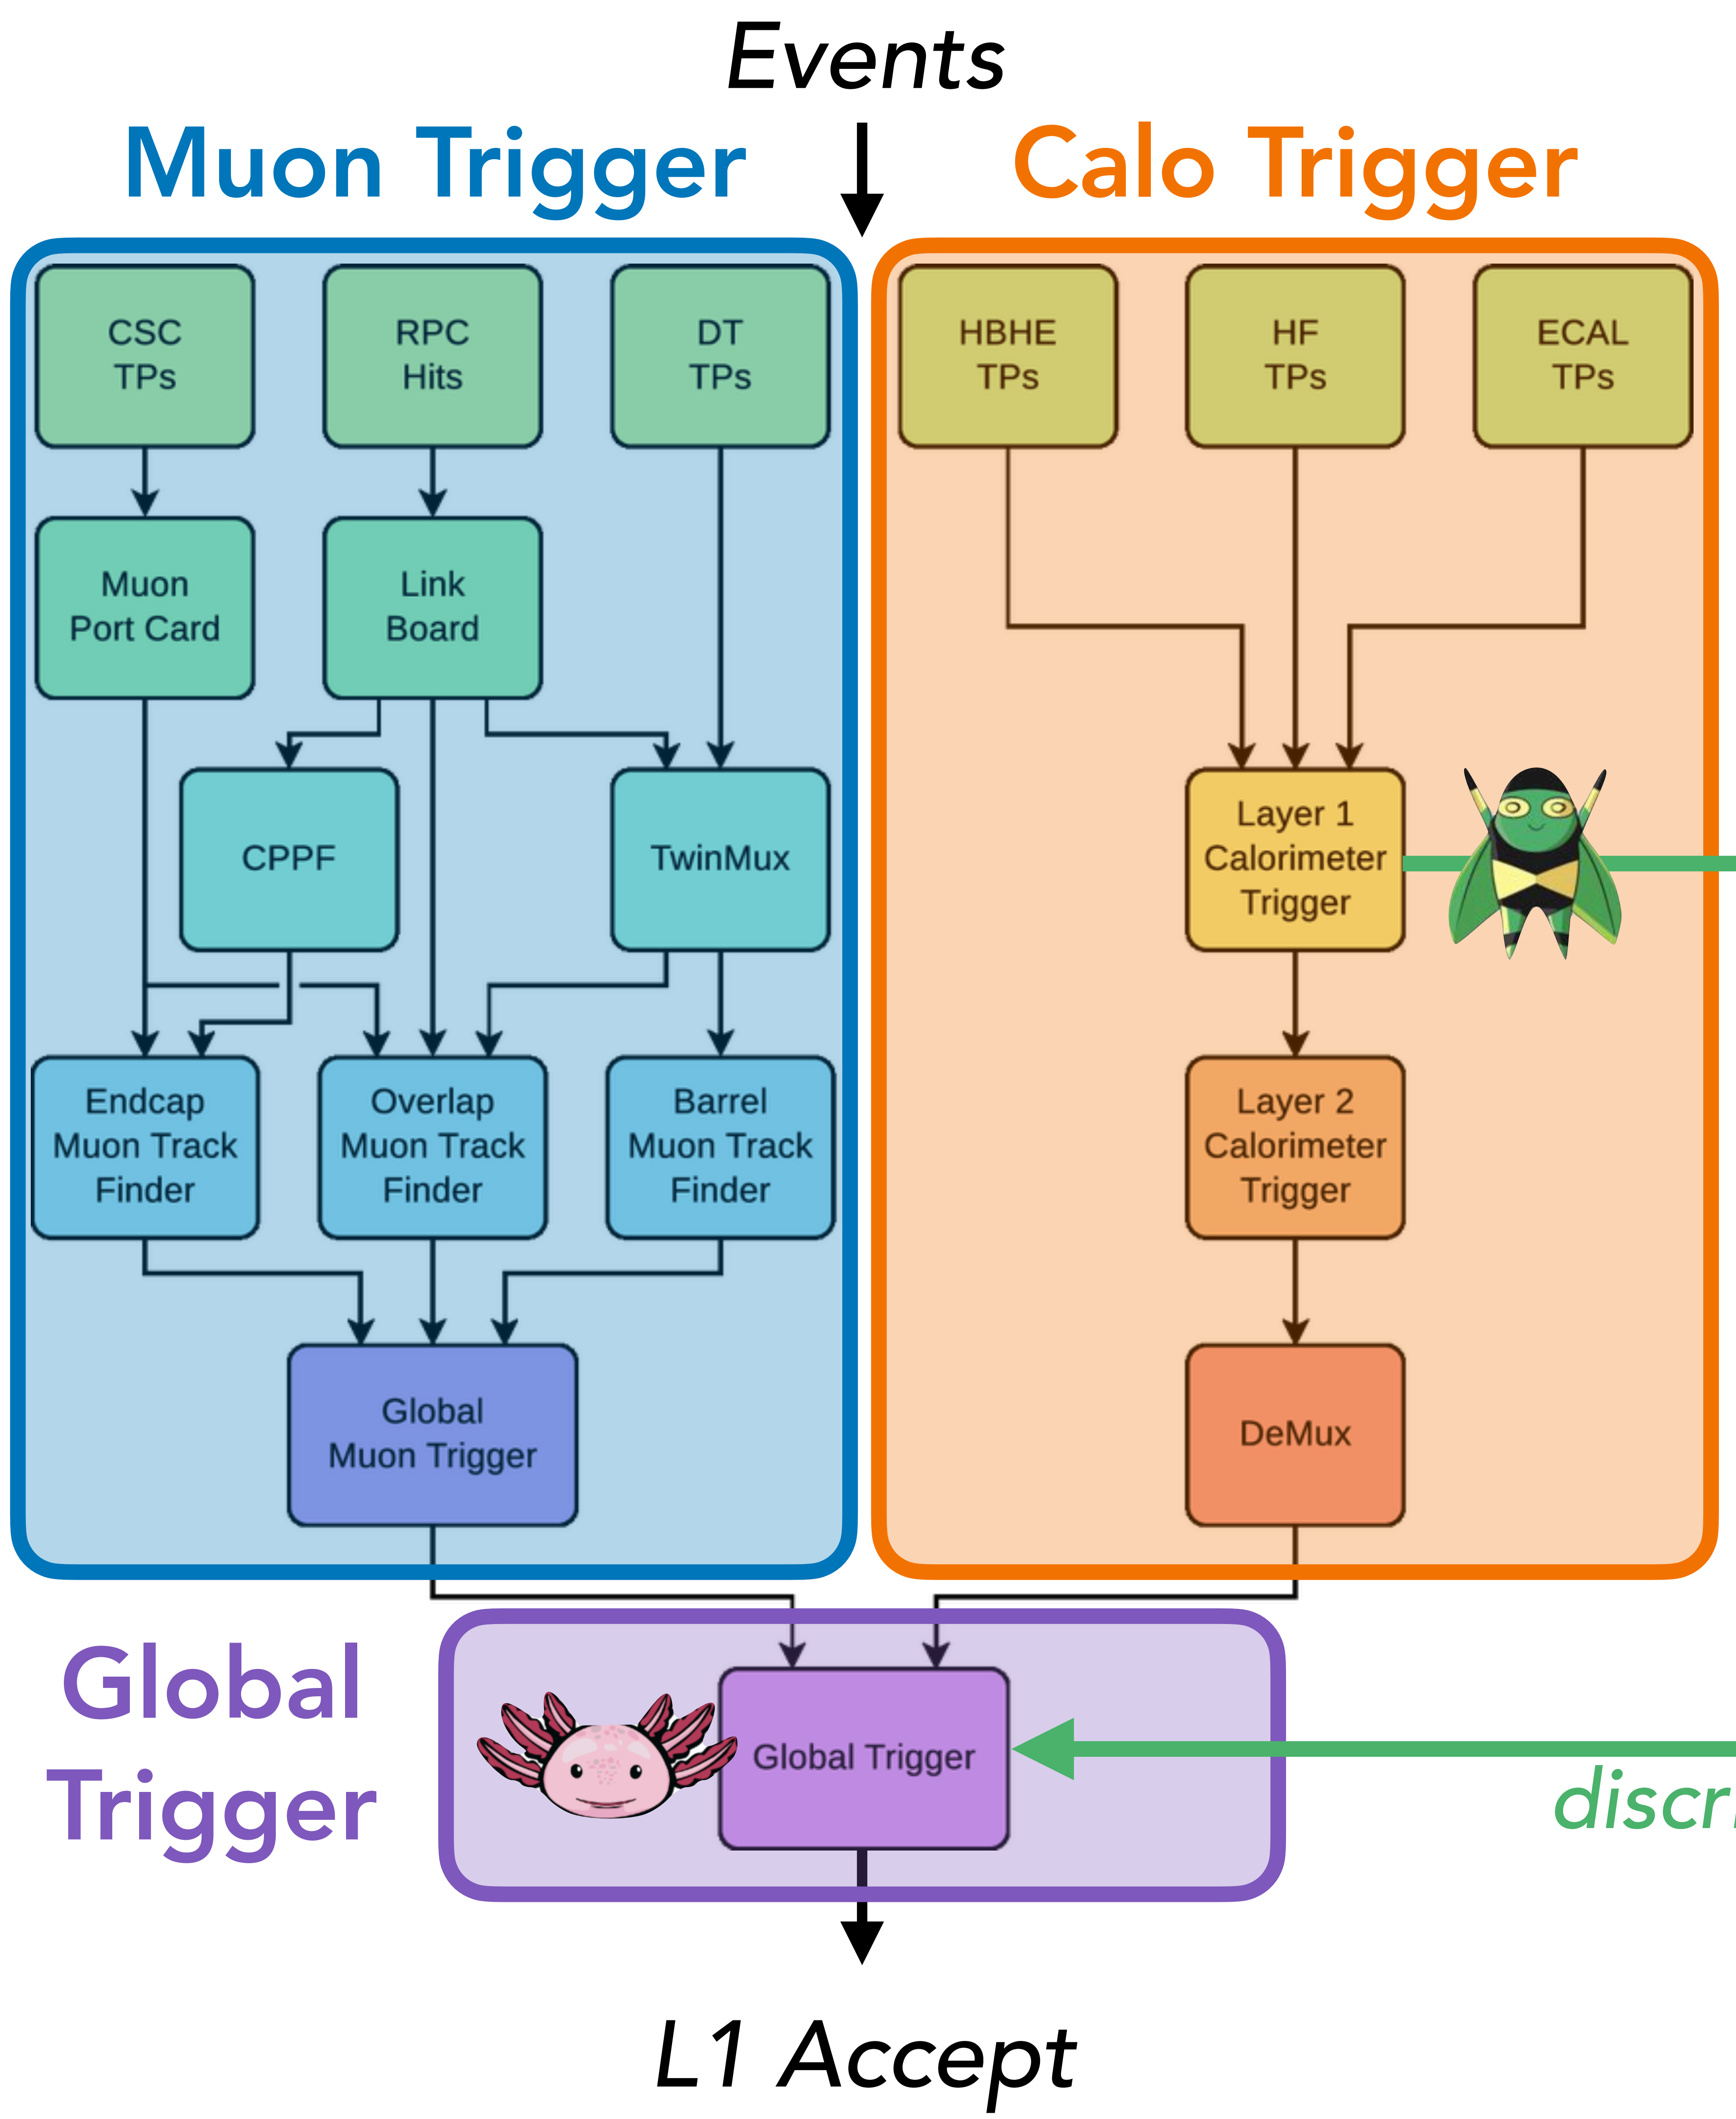

CICADA uses calorimeter system inputs

Wisconsin CTP7 Board

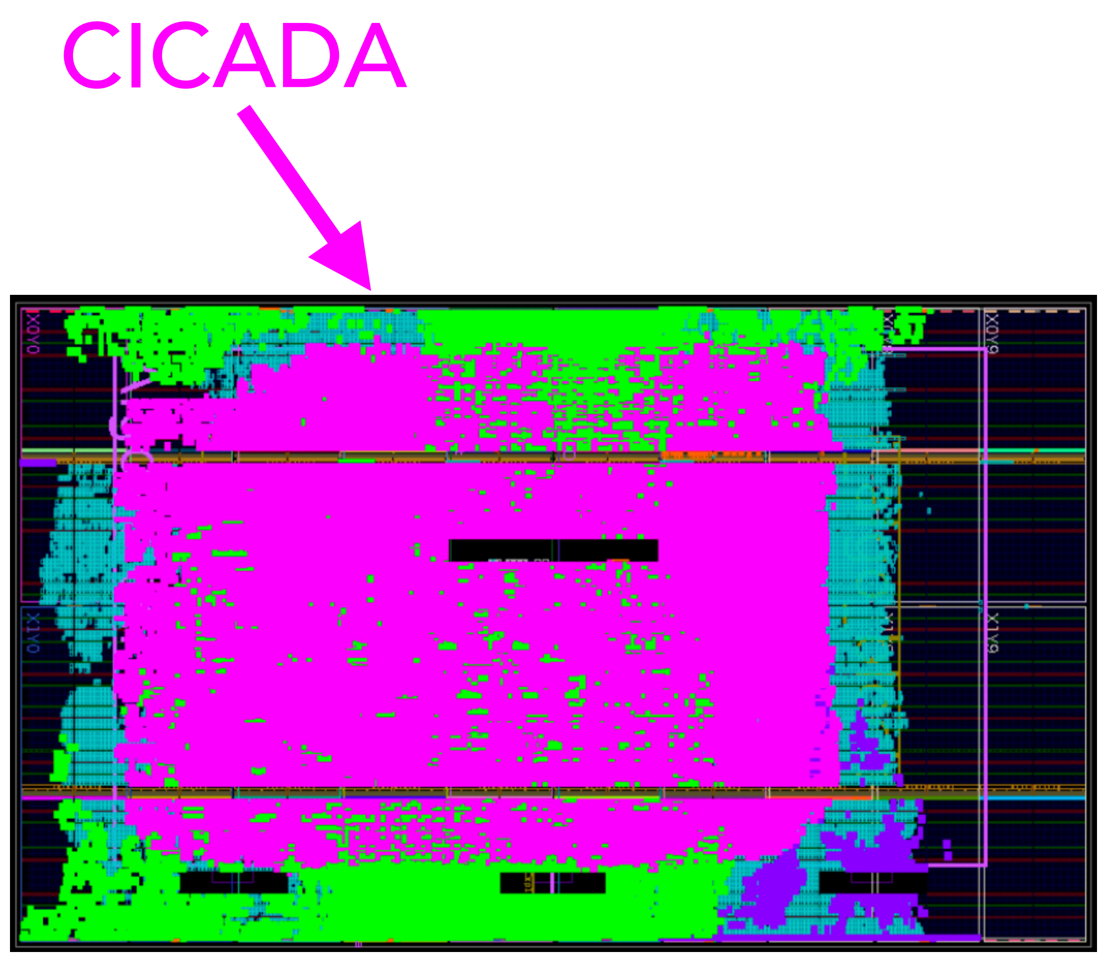

L. Moreno

discriminator score

Key challenges

- Model size (deploy on FPGA) → O(1-10k) params
- Model speed (fit into L1 latency) → O(μs) latency

Inputs to CICADA

Event-level calorimeter image 18ϕ x 14η *regions* (252 pixels)

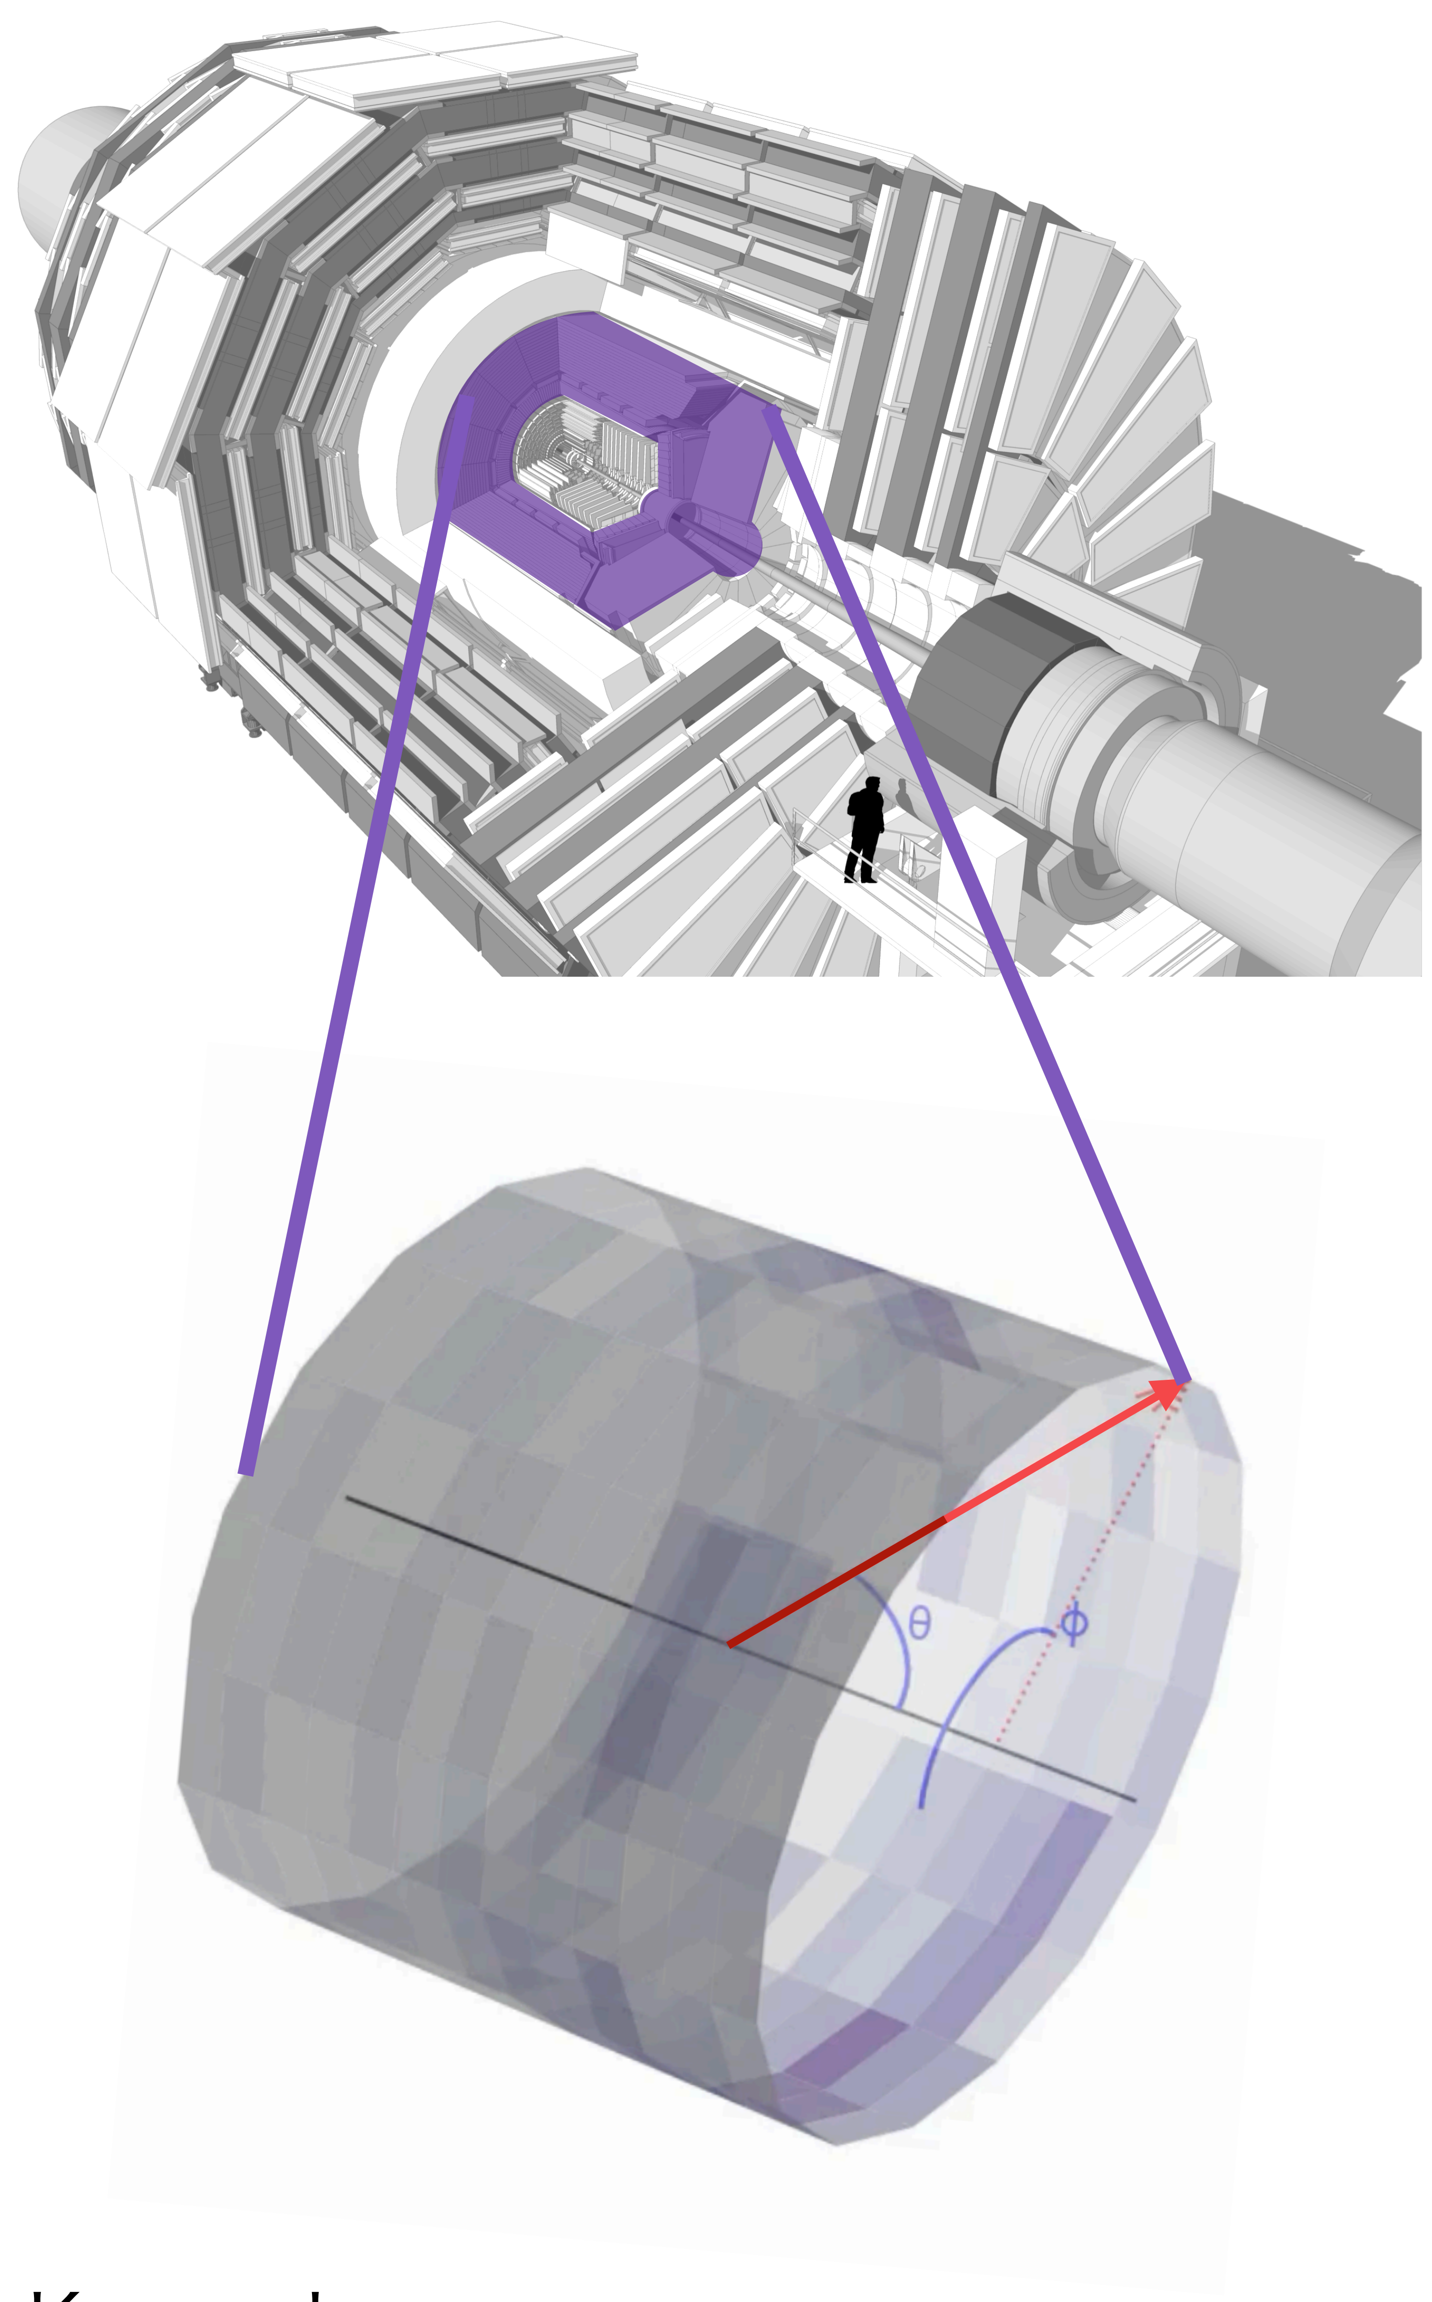

Inputs to CICADA

Event-level calorimeter image 18ϕ x 14η *regions* (252 pixels)

- Coarse information: each region represents a sum of 4x4 trigger towers, includes both ECAL and HCAL energy deposits
- Cluster-independent: low-level information before any other object algorithms or selections

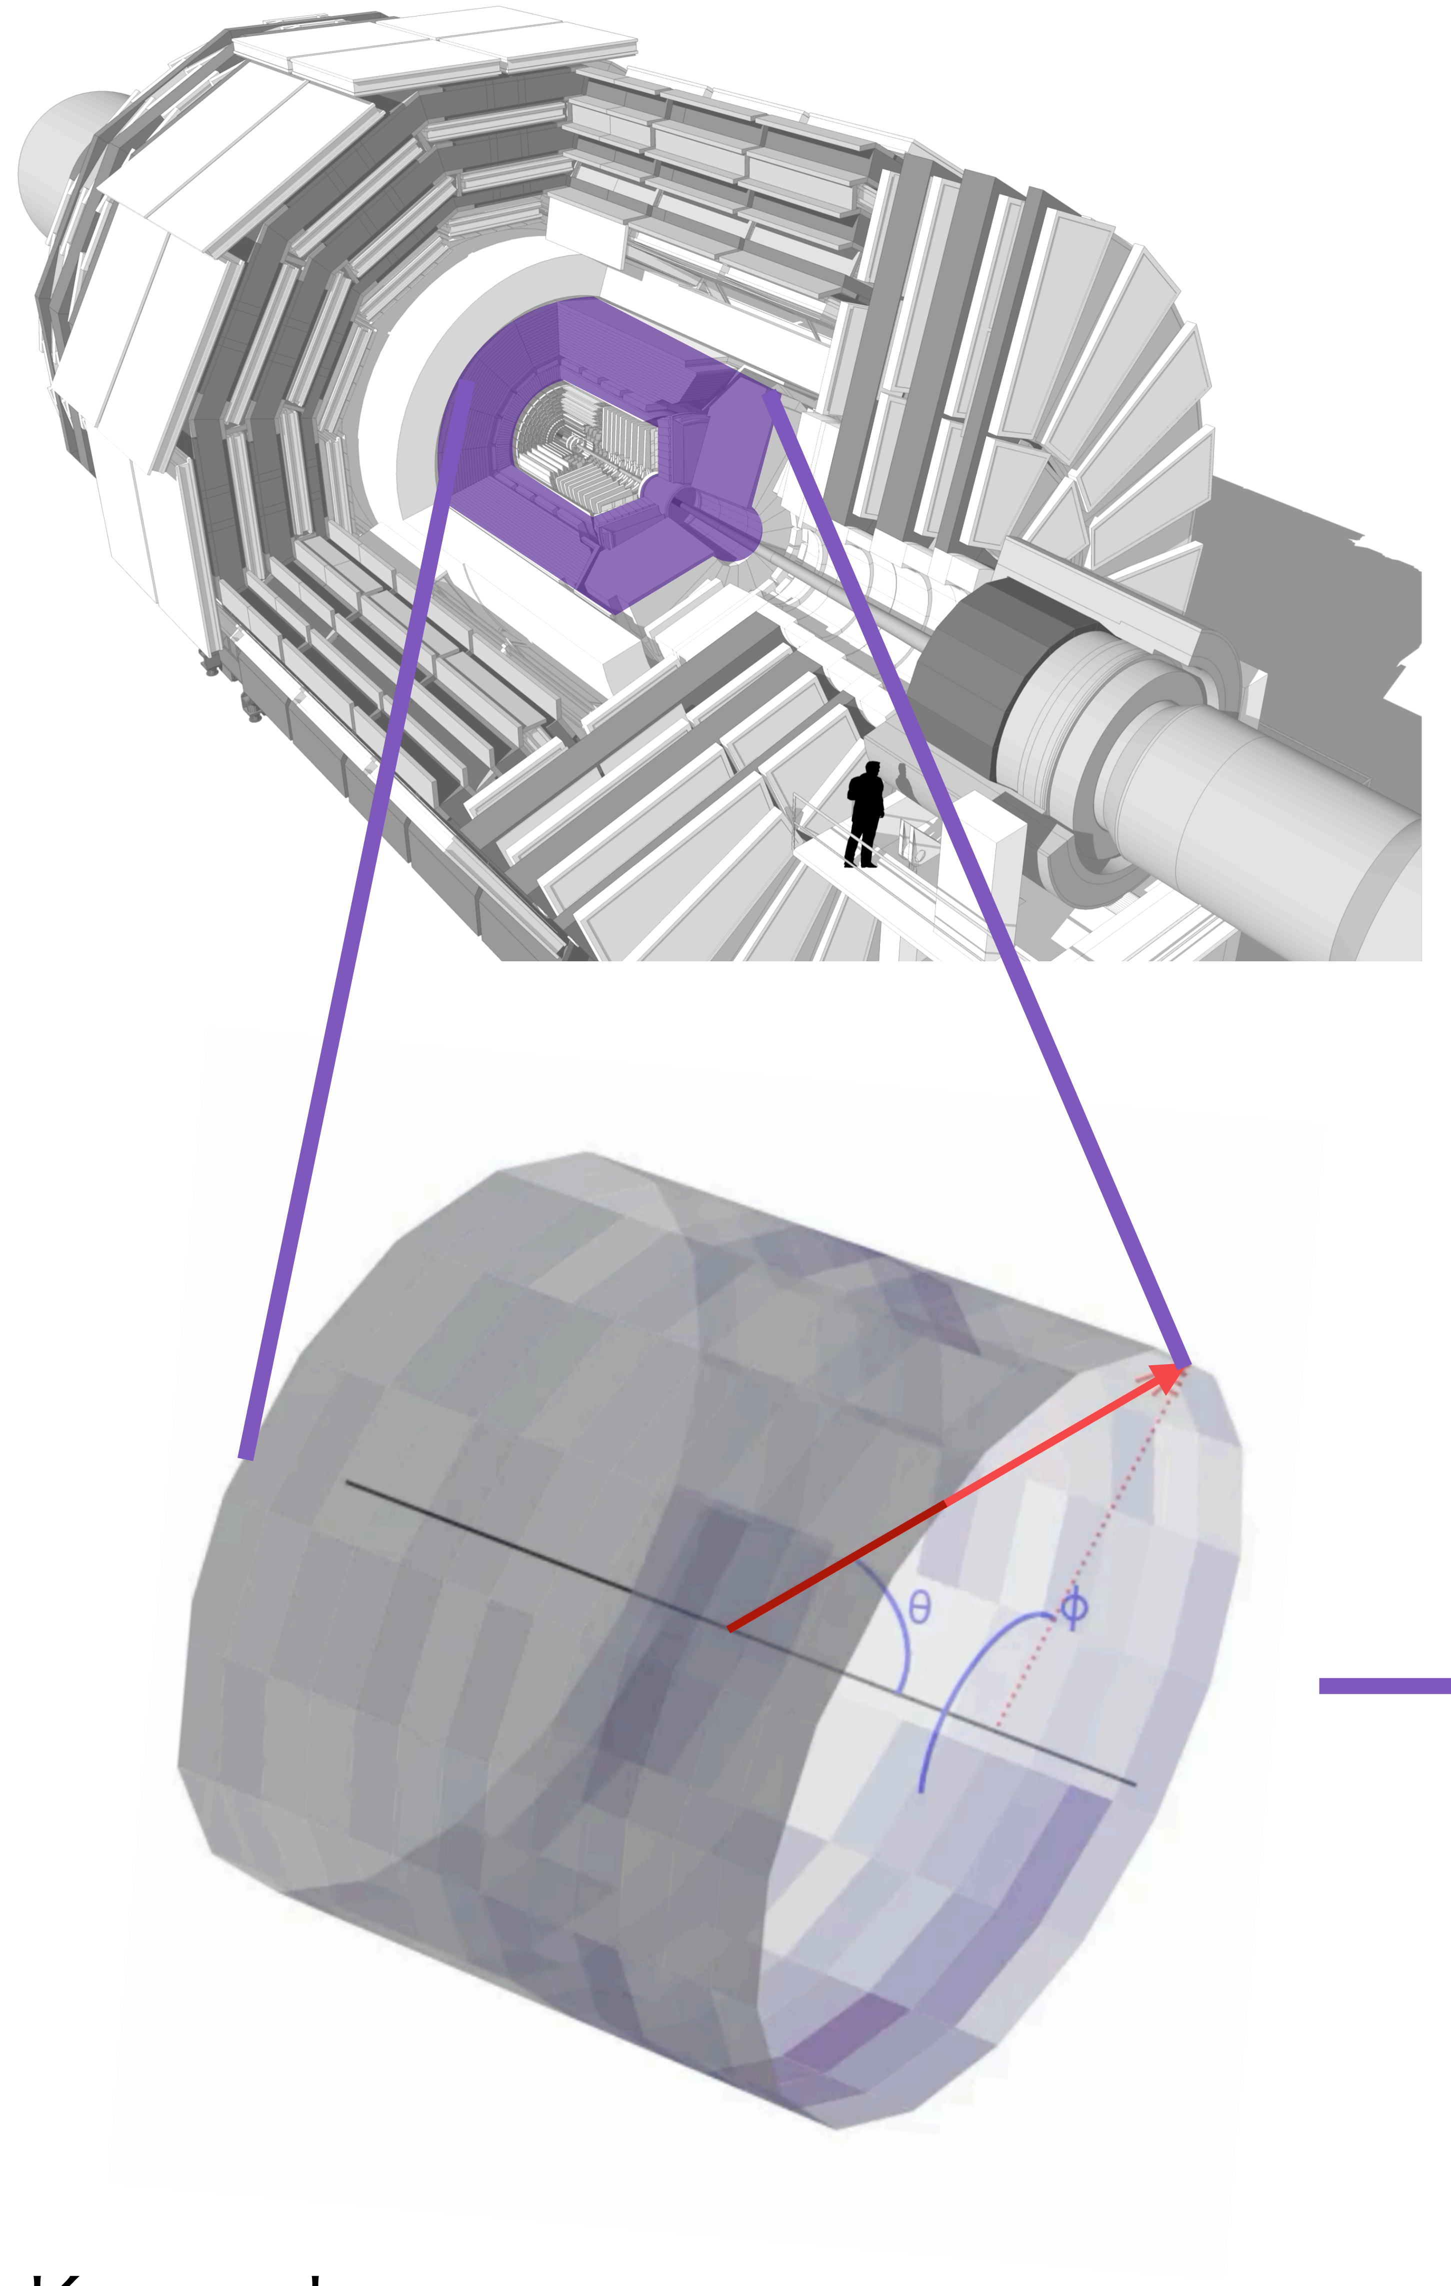

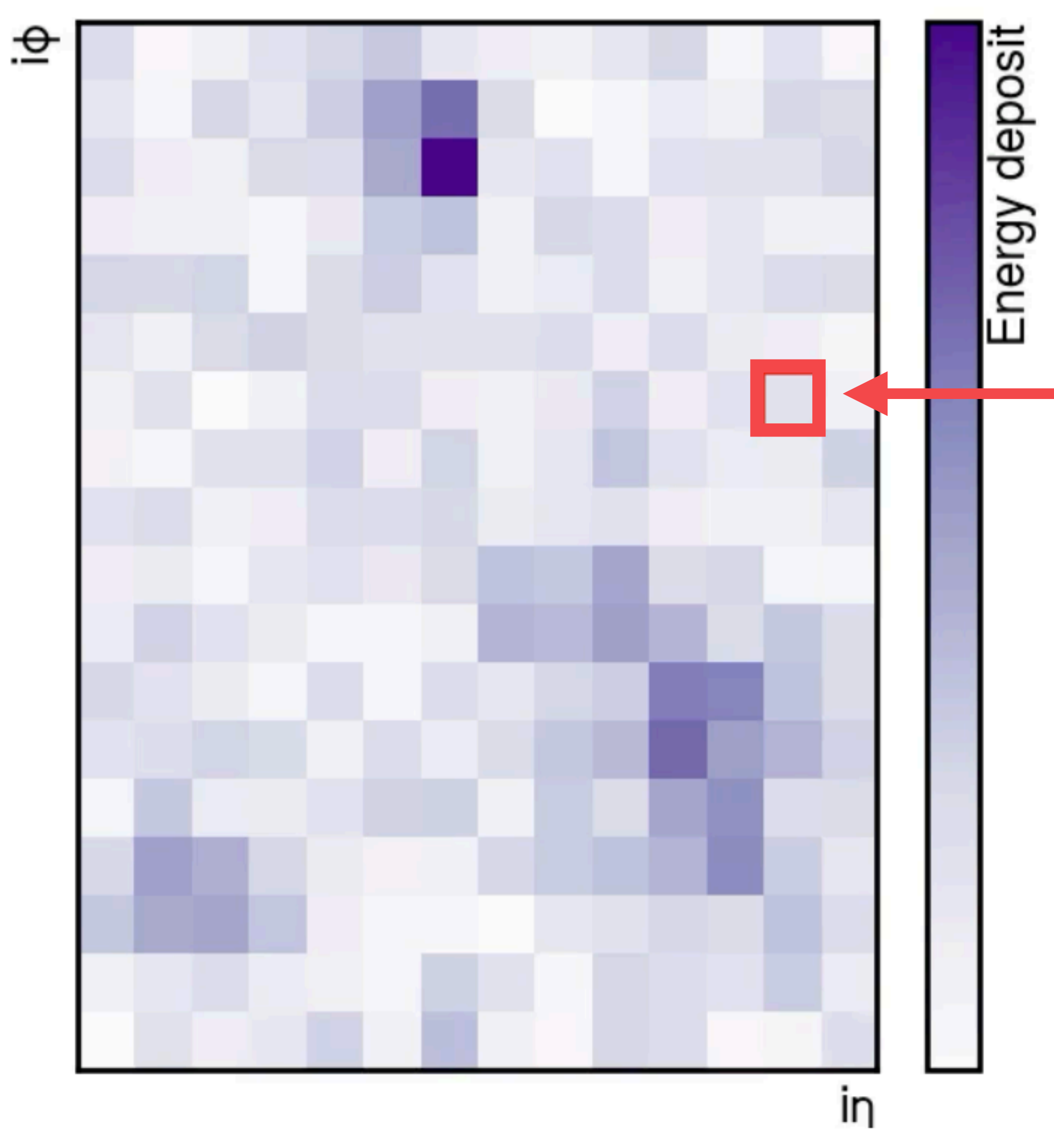

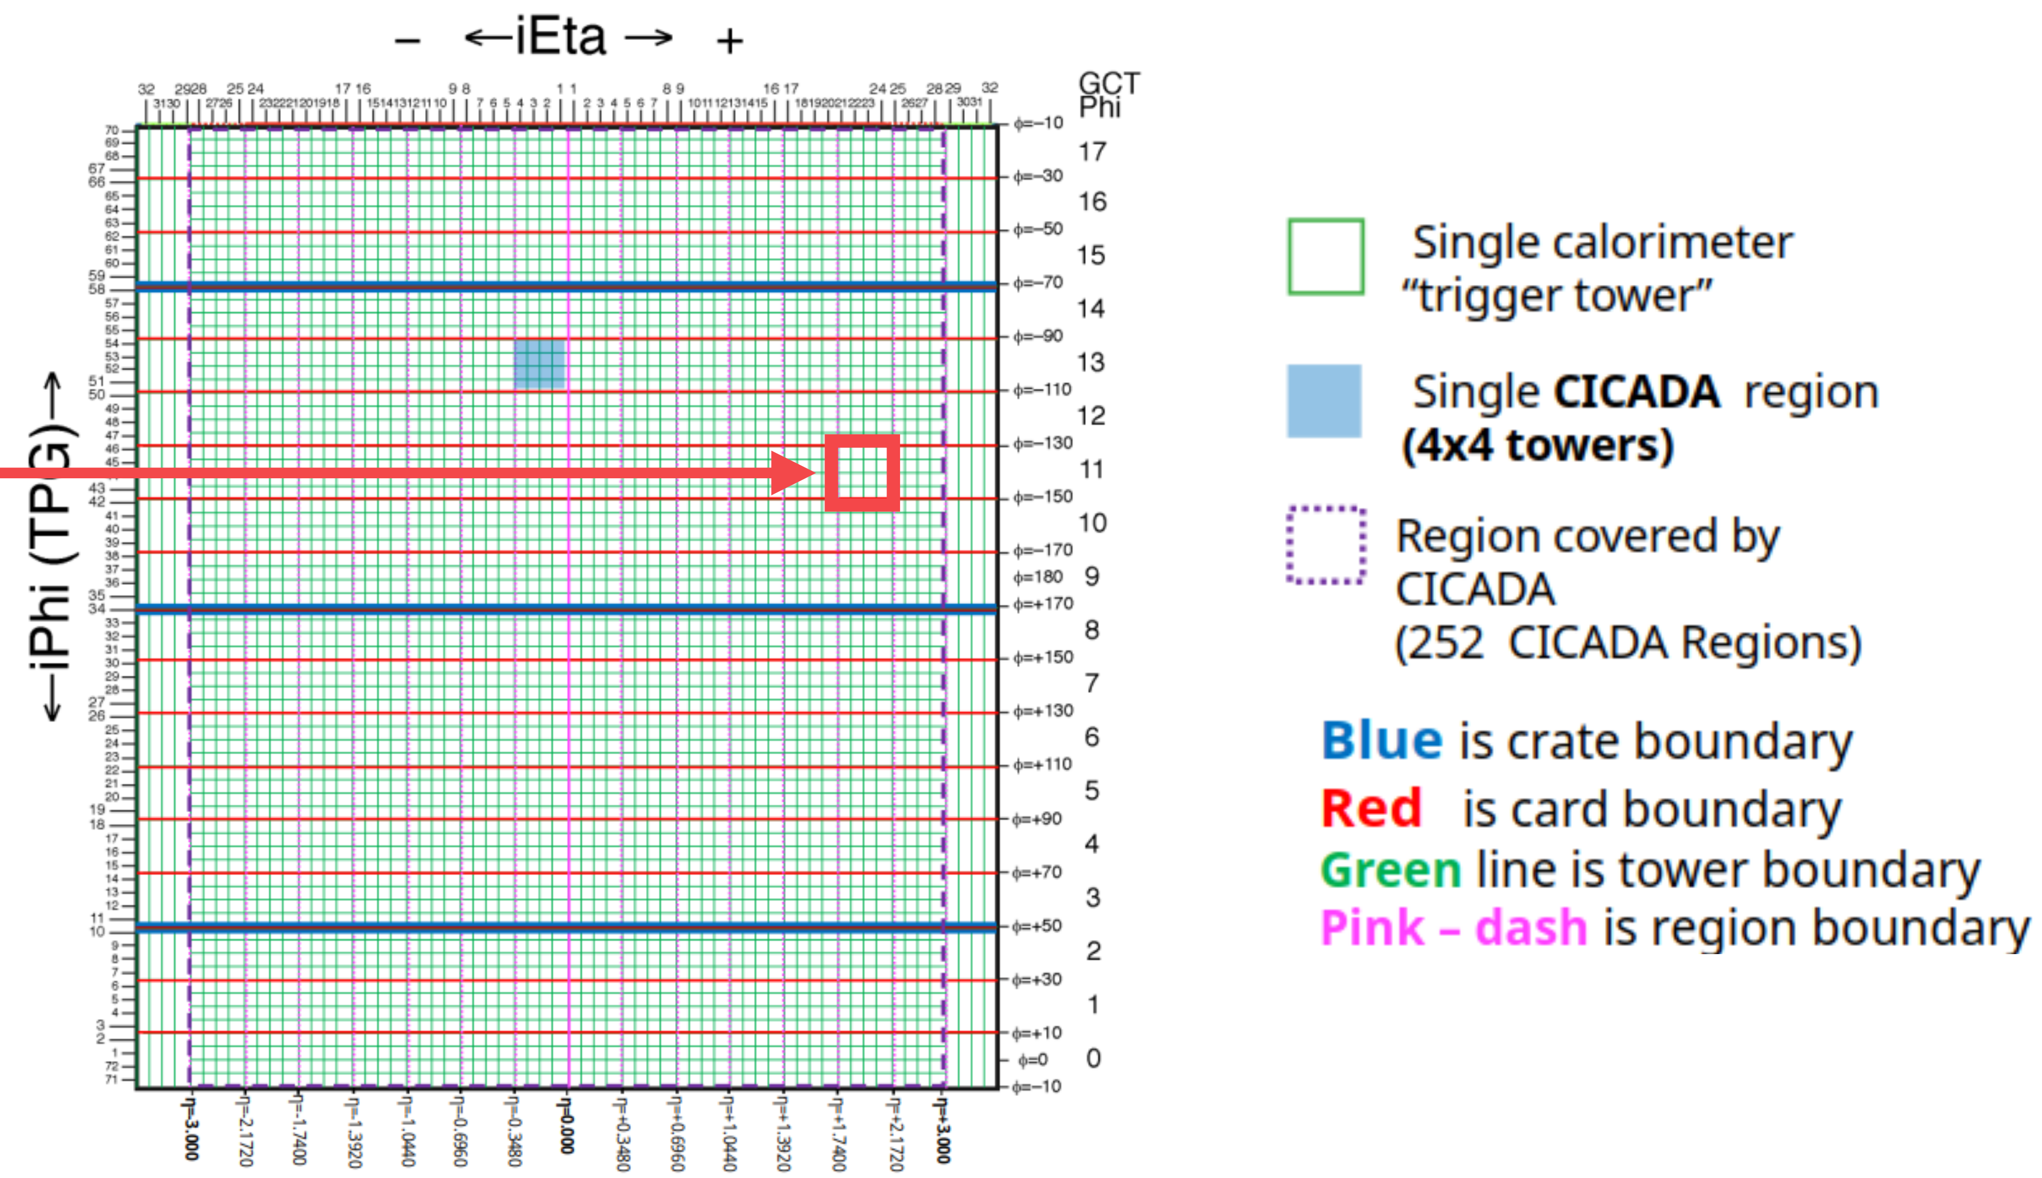

CICADA Model + Training

- Event-level calorimeter image → classic input to CNN
- Training: "Zero Bias" data (random selection of events)

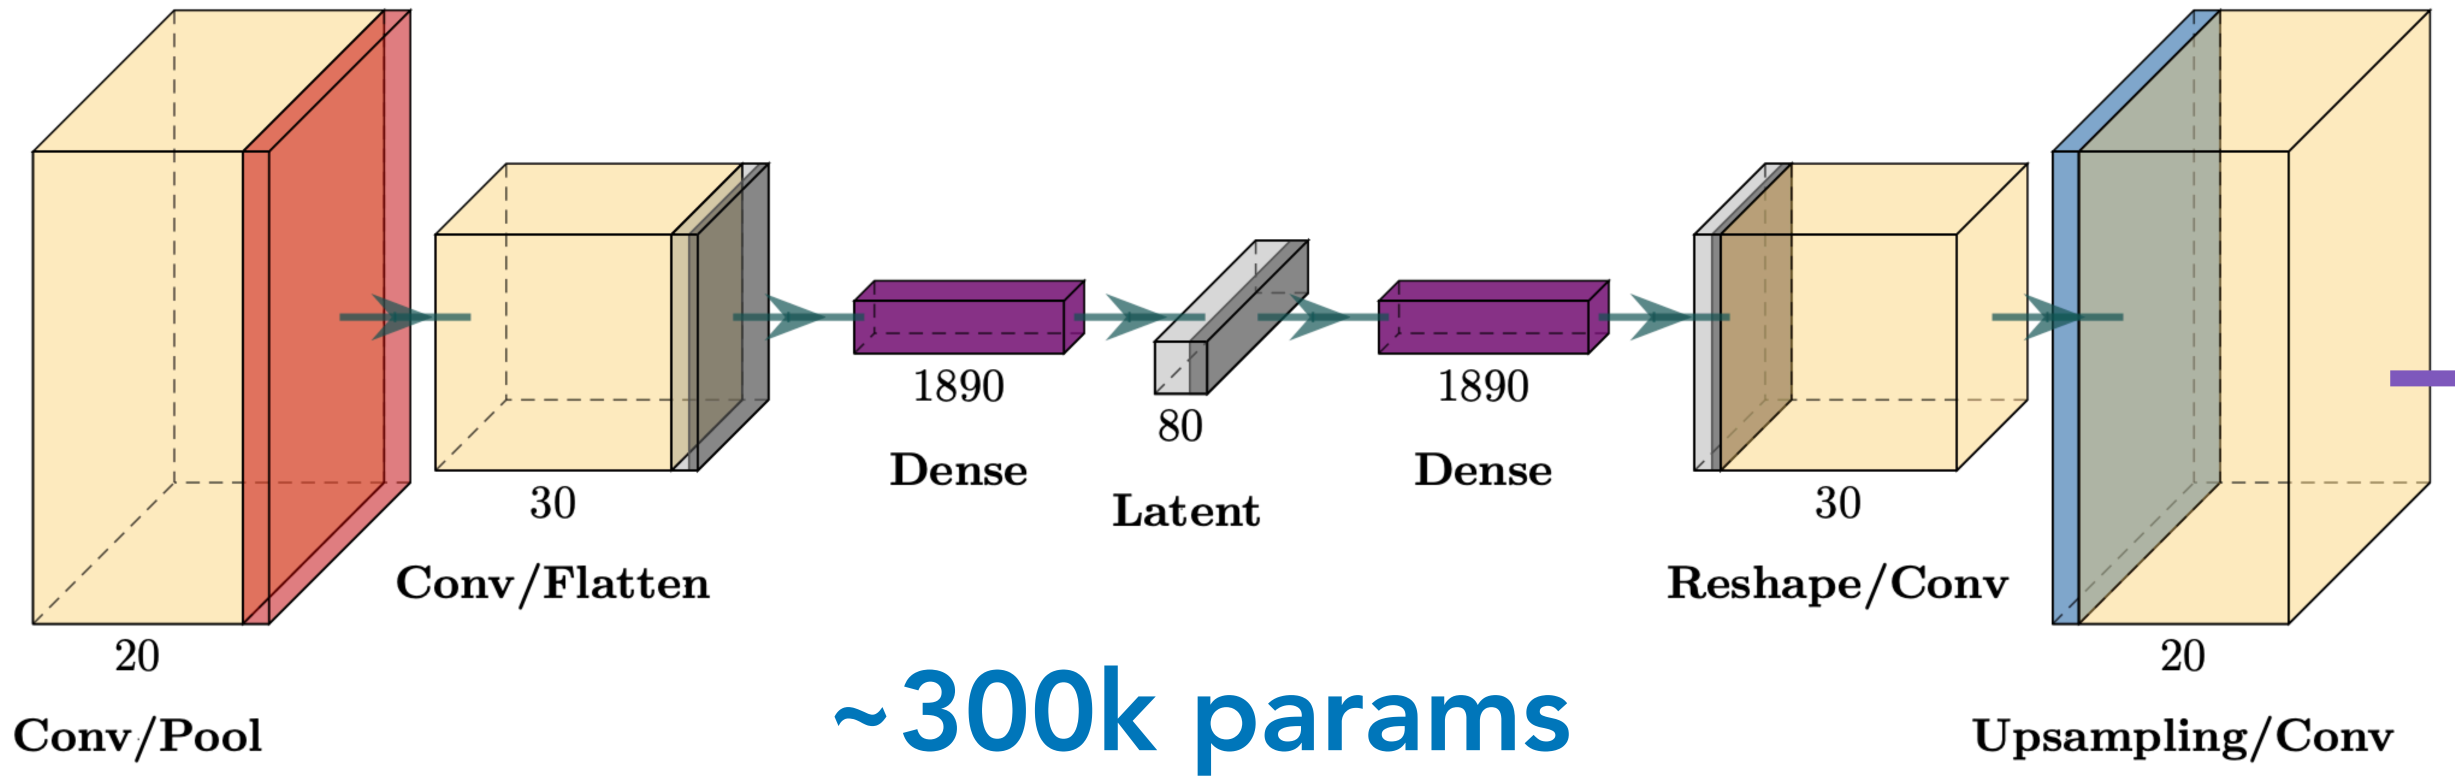

~300k params

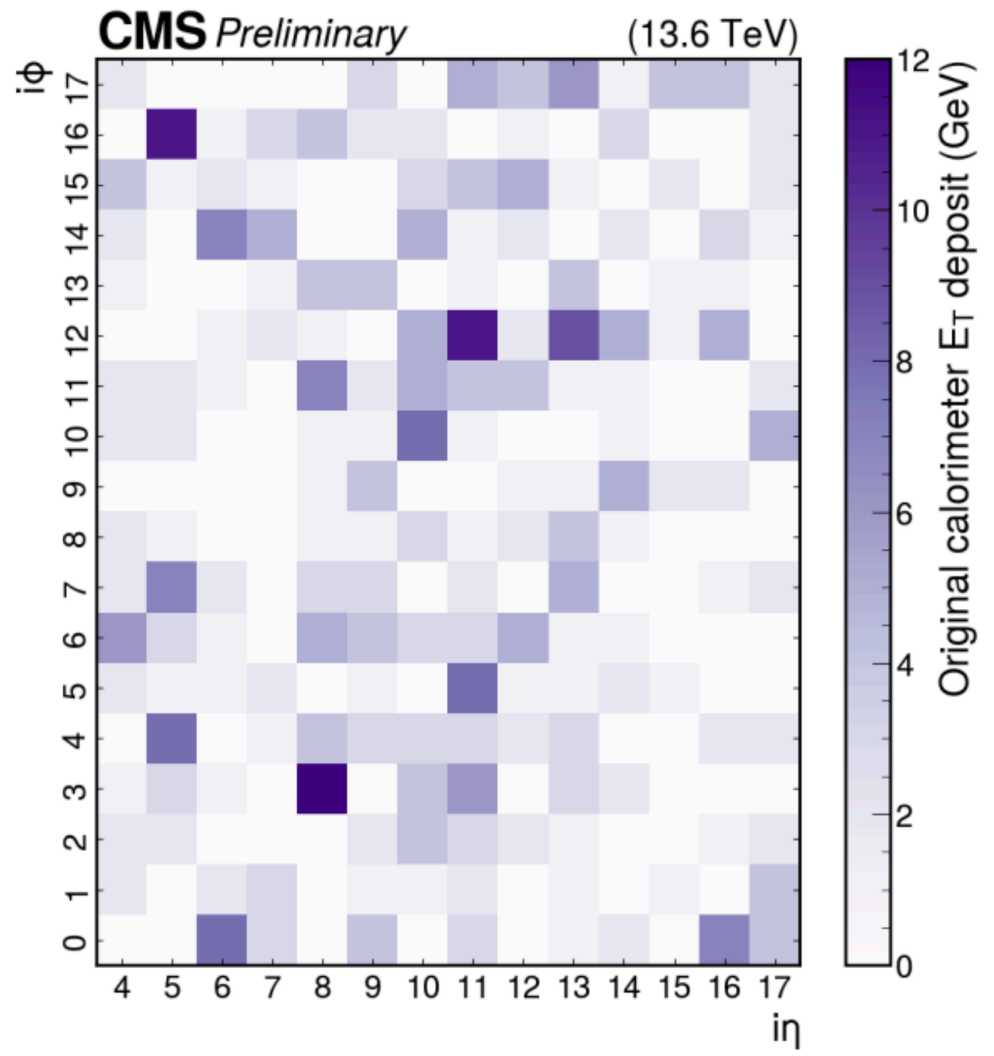

Anomaly score = reconstruction loss

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

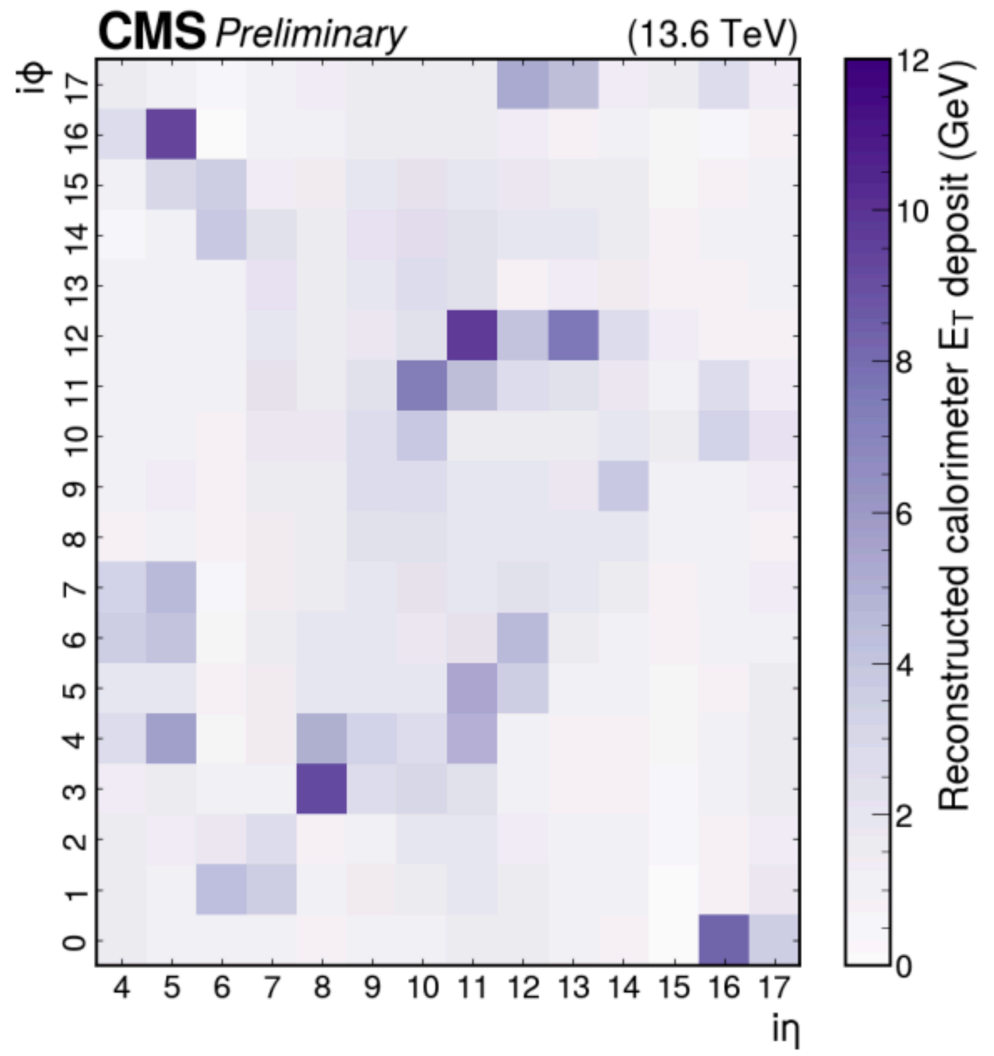

CICADA Model + Training

- Event-level calorimeter image → classic input to CNN
- Training: "Zero Bias" data (random selection of events)

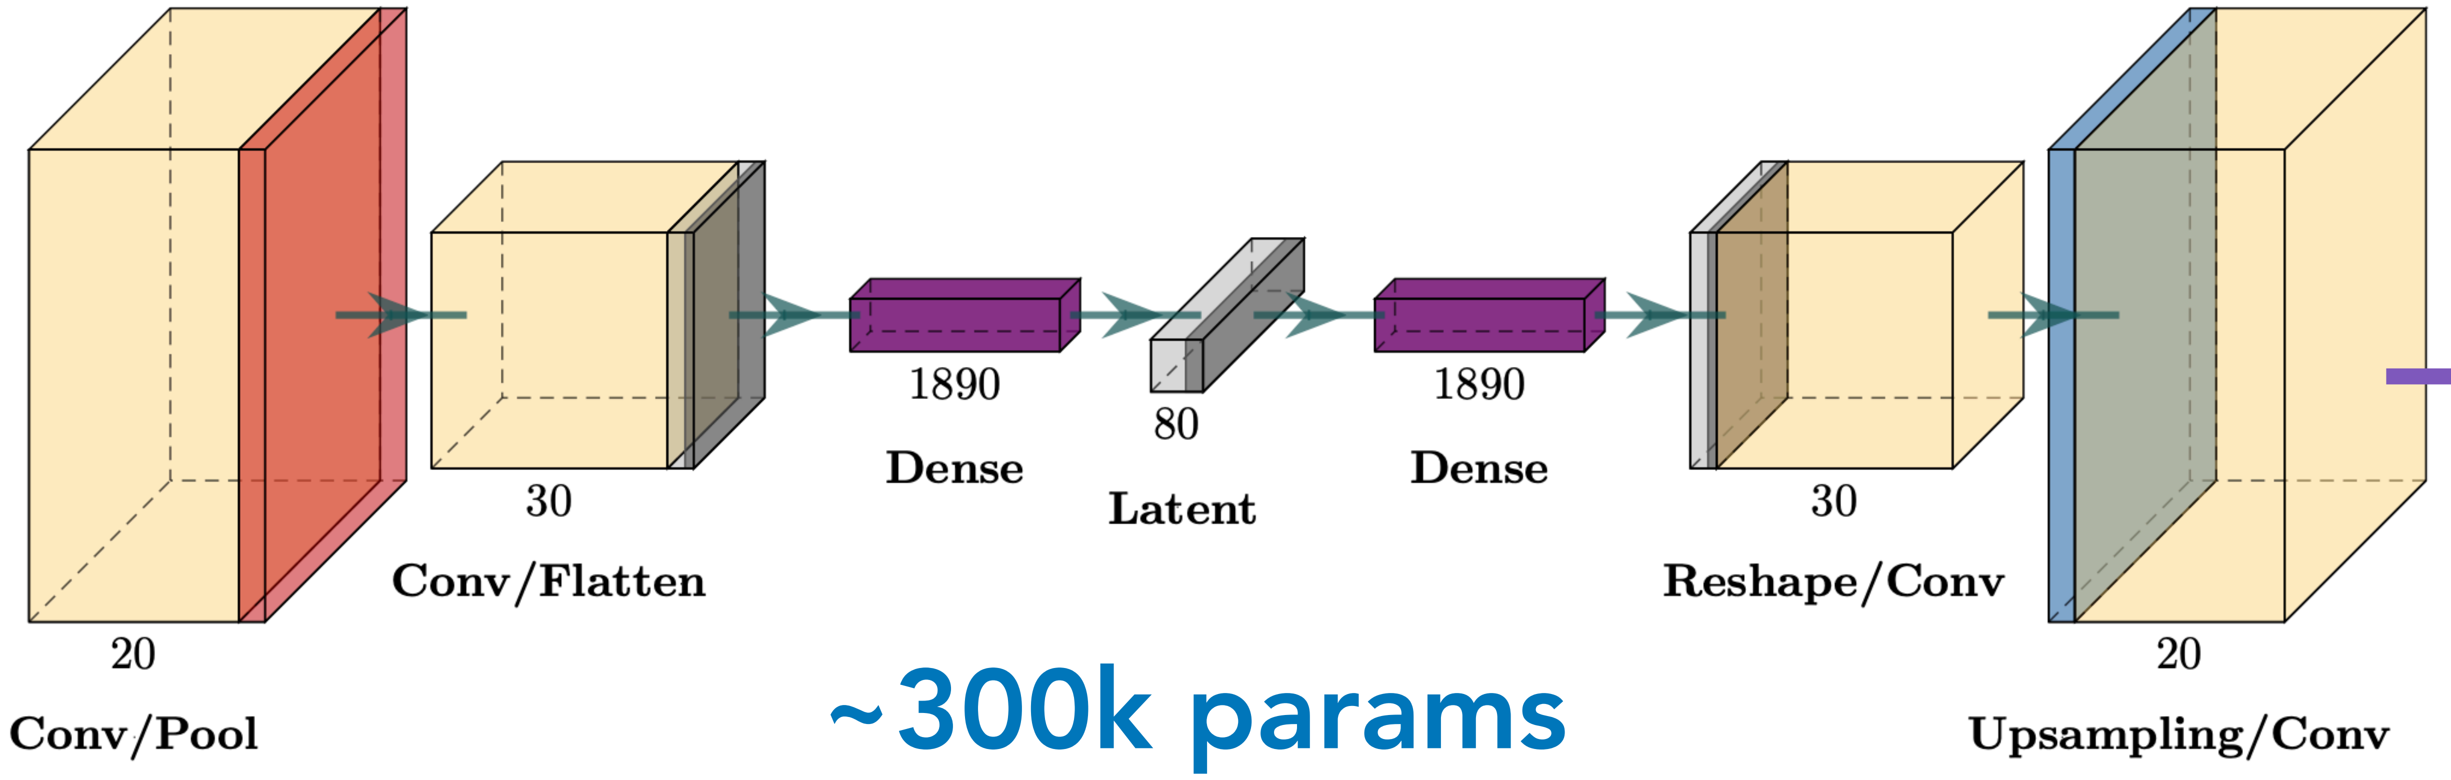

~300k params

Key challenge: model too large and slow for L1

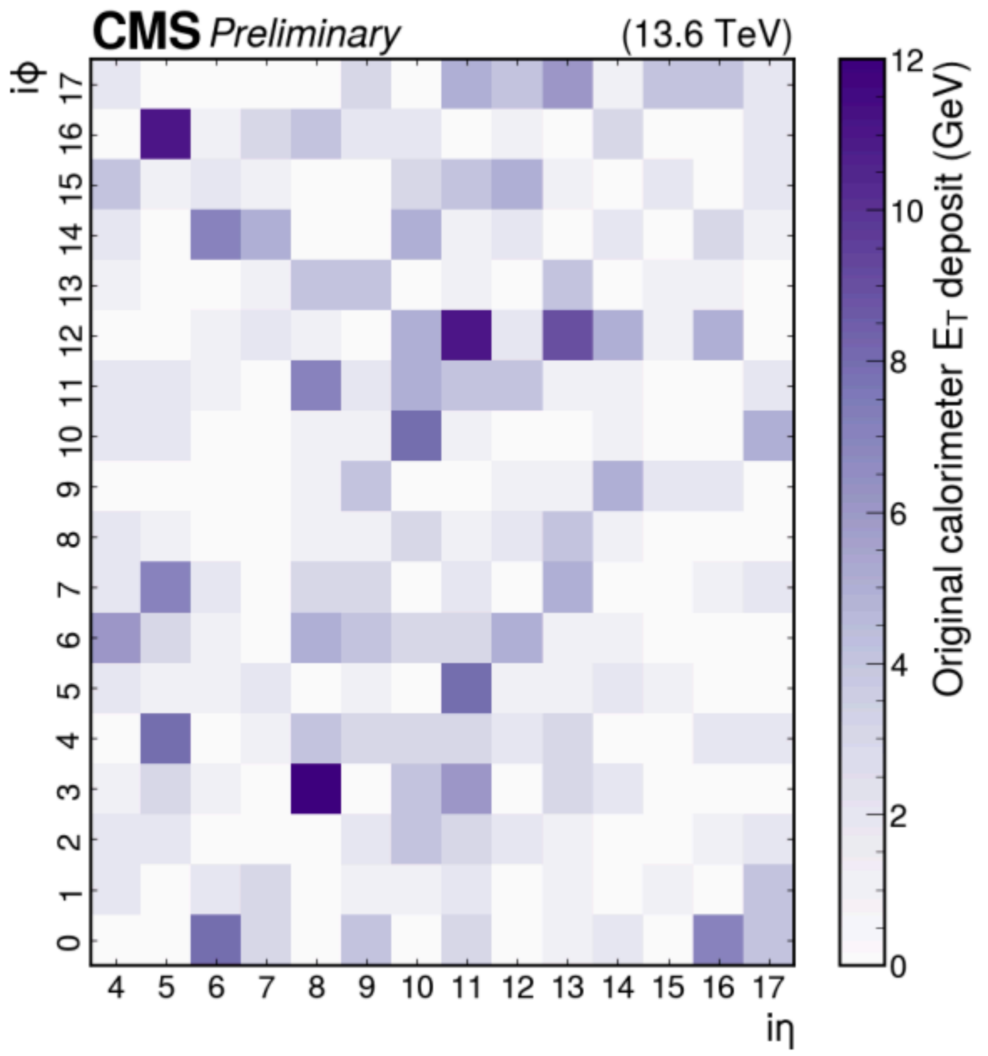

Anomaly score = reconstruction loss

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

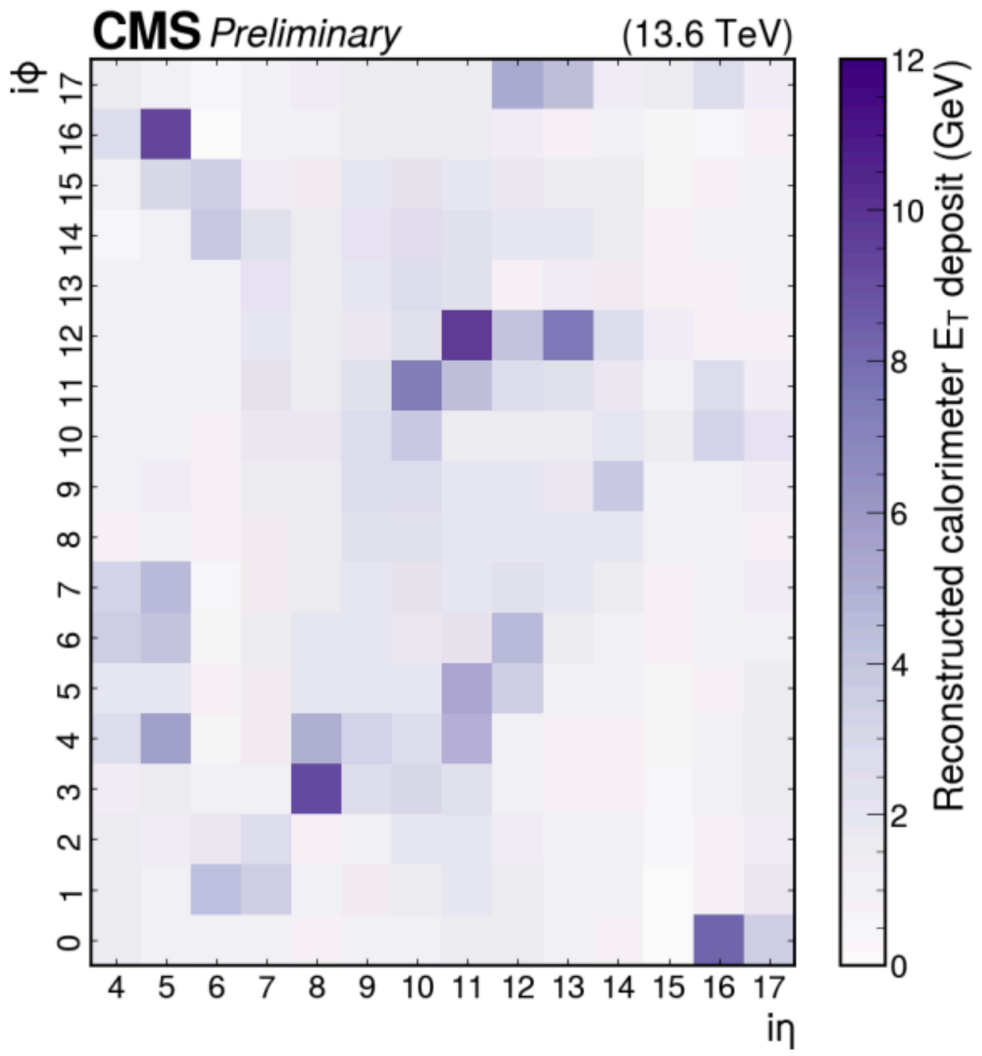

CICADA Model + Training

- Event-level calorimeter image → classic input to CNN
- Training: "Zero Bias" data (random selection of events)

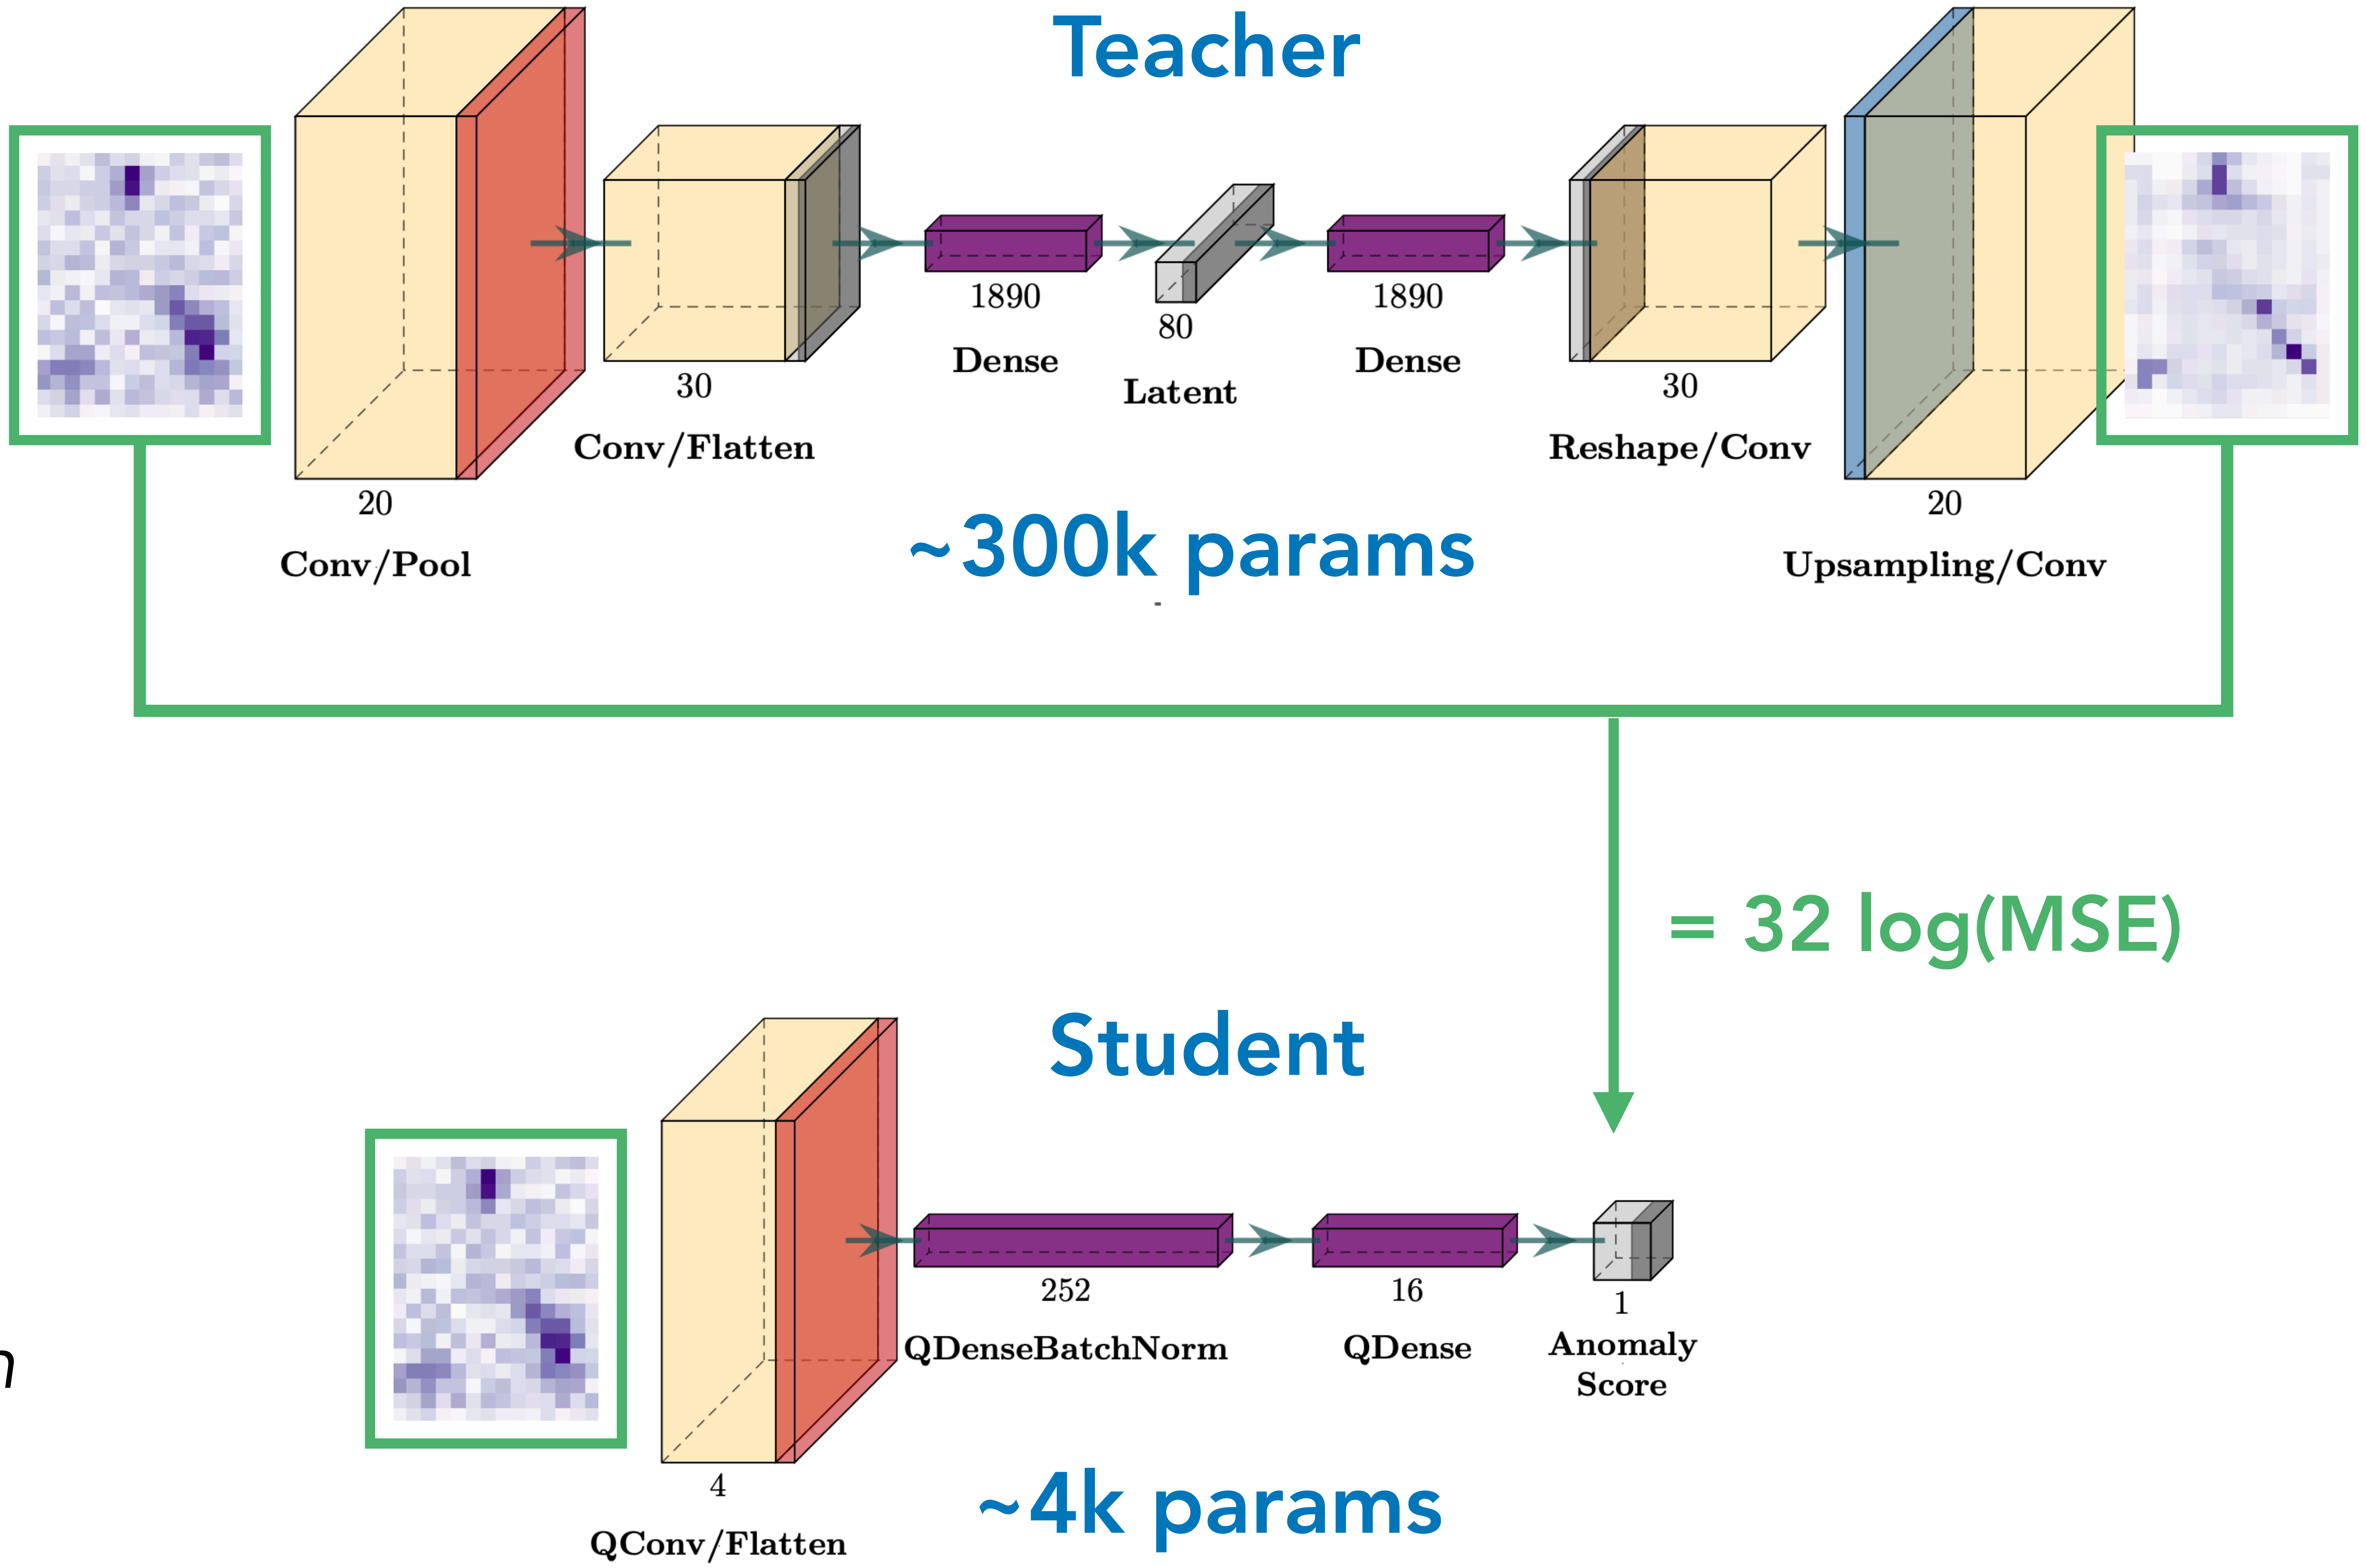

Key challenge: model too large and slow for L1

Solution: knowledge distillation

Train smaller model ("student") to learn score of the larger network ("teacher")

Implementation at Level 1

Train teacher with
tensorflow + keras

L. Gerlach

Implementation at Level 1

Train teacher with tensorflow + keras

Quantization-aware training for student with **QKeras**

- Model weights quantized to fixed precision → **10x reduction in resources/latency**
- Better to train a quantized model rather than quantize a trained model

Implementation at Level 1

Train teacher with tensorflow + keras

Translate trained NN into C++ with hls4ml

Quantization-aware training for student with **QKeras**

- Model weights quantized to fixed precision → **10x reduction in resources/latency**
- Better to train a quantized model rather than quantize a trained model

Implementation at Level 1

Train teacher with tensorflow + keras

Translate trained NN into C++ with hls4ml

L. Gerlach

Quantization-aware training for student with **QKeras**

Produce FPGA code from C++ with Xilinx HLS

- Model weights quantized to fixed precision → **10x reduction in resources/latency**
- Better to train a quantized model rather than quantize a trained model

Implementation at Level 1

Train teacher with tensorflow + keras

Translate trained NN into C++ with hls4ml

Deploy on AMD Vertex FPGA chip

L. Gerlach

Quantization-aware training for student with **QKeras**

Produce FPGA code from C++ with Xilinx HLS

- Model weights quantized to fixed precision → **10x reduction in resources/latency**
- Better to train a quantized model rather than quantize a trained model

Rates + Operations

<u>CICADA Fires on “Standard” and “New” Events</u>

- “Pure” refers to events *only* fired by CICADA
- Significant overlap with L1 menu (known “anomalies”) while saving new events

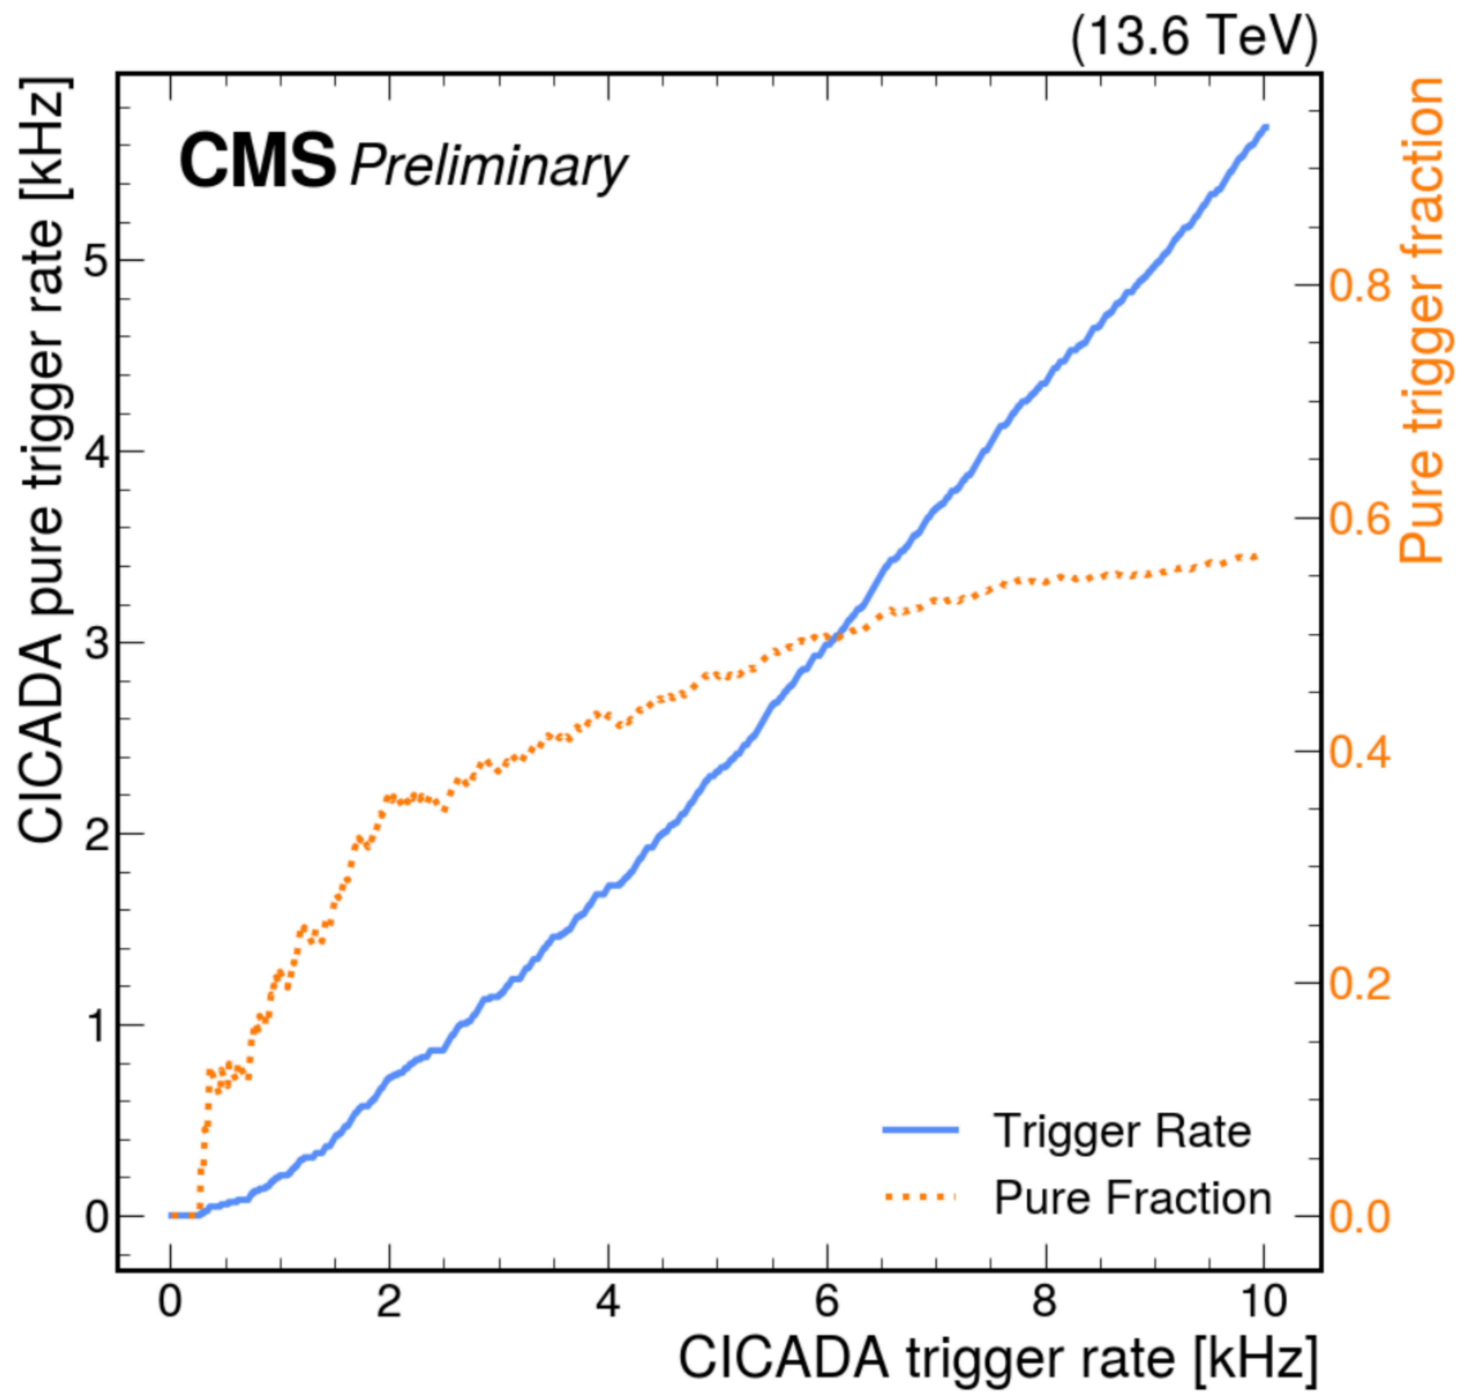

Rates + Operations

CICADA Fires on "Standard" and "New" Events

- "Pure" refers to events *only* fired by CICADA
- Significant overlap with L1 menu (known "anomalies") while saving new events

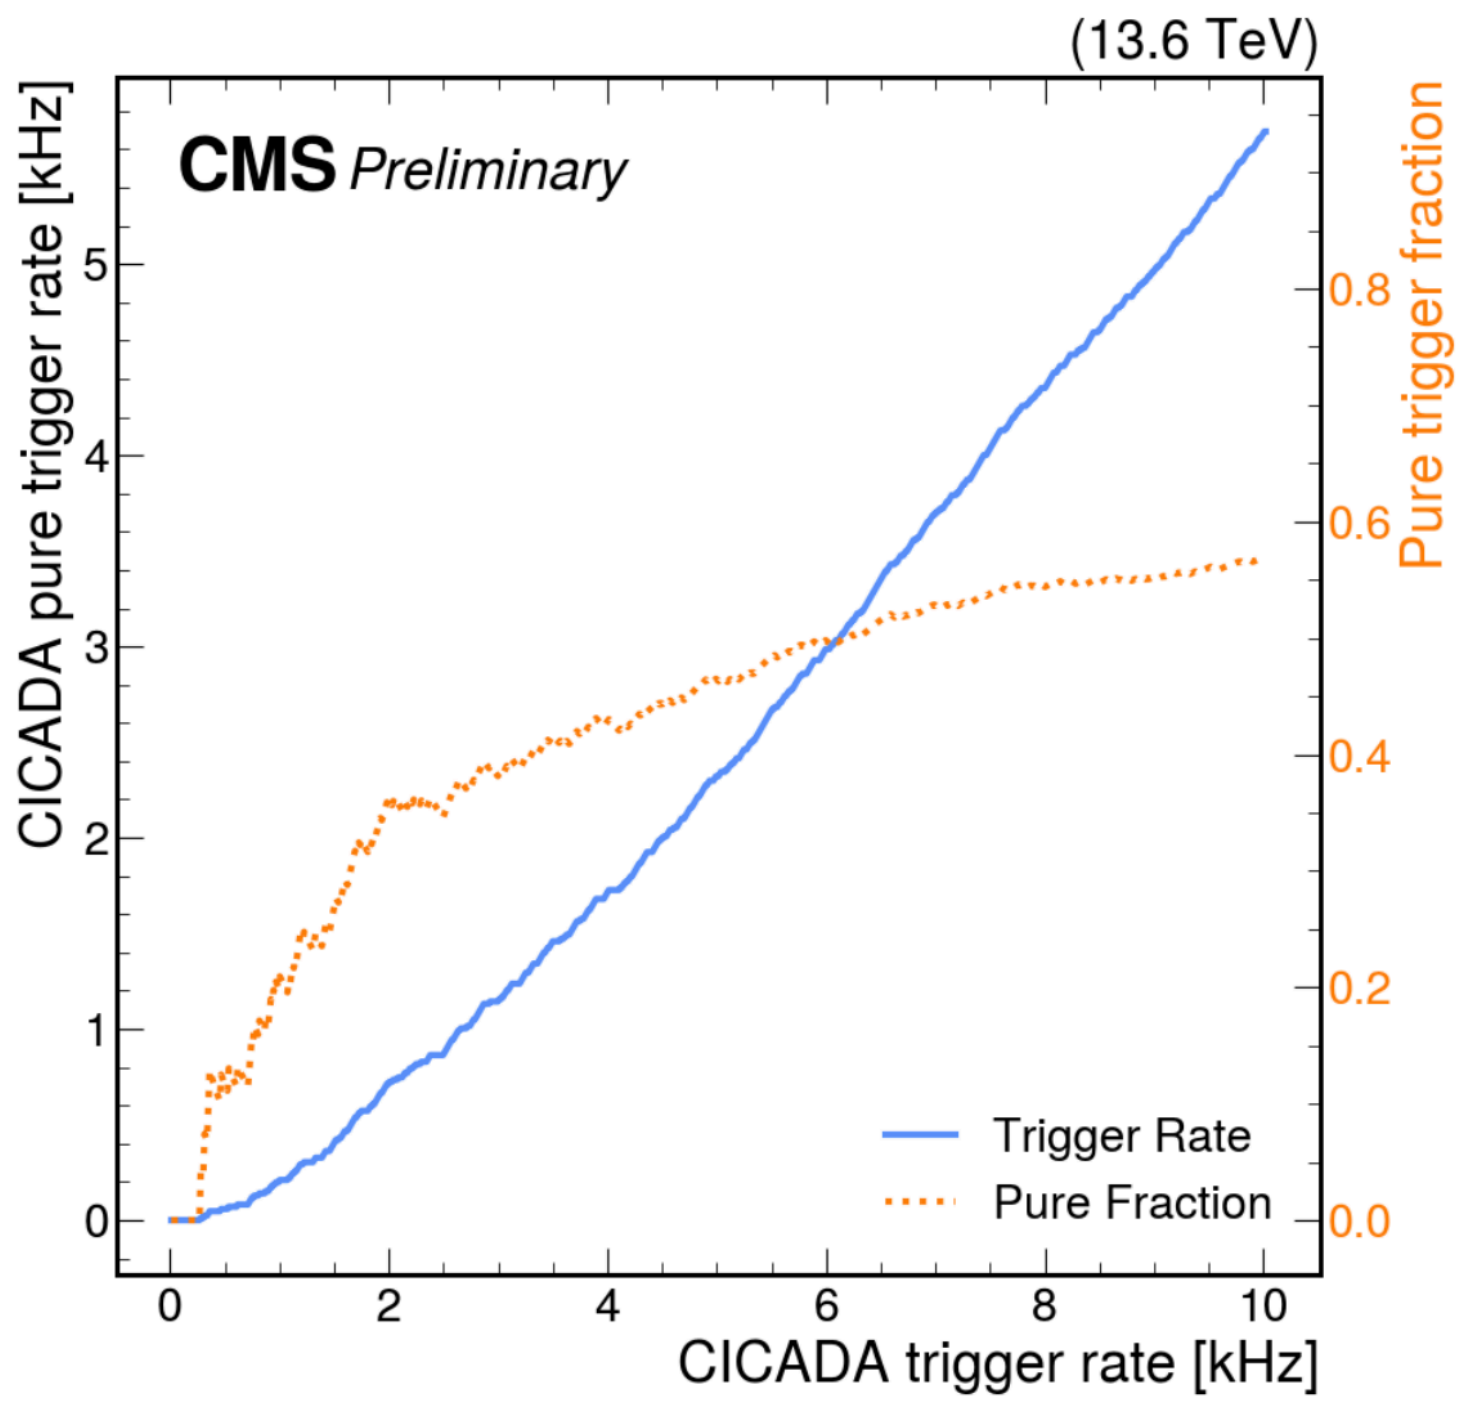

Observe Good Rate Stability

- CICADA rates relatively stable
- Some dependence on pileup and various calorimeter corrections (up to ~20% rate variation)

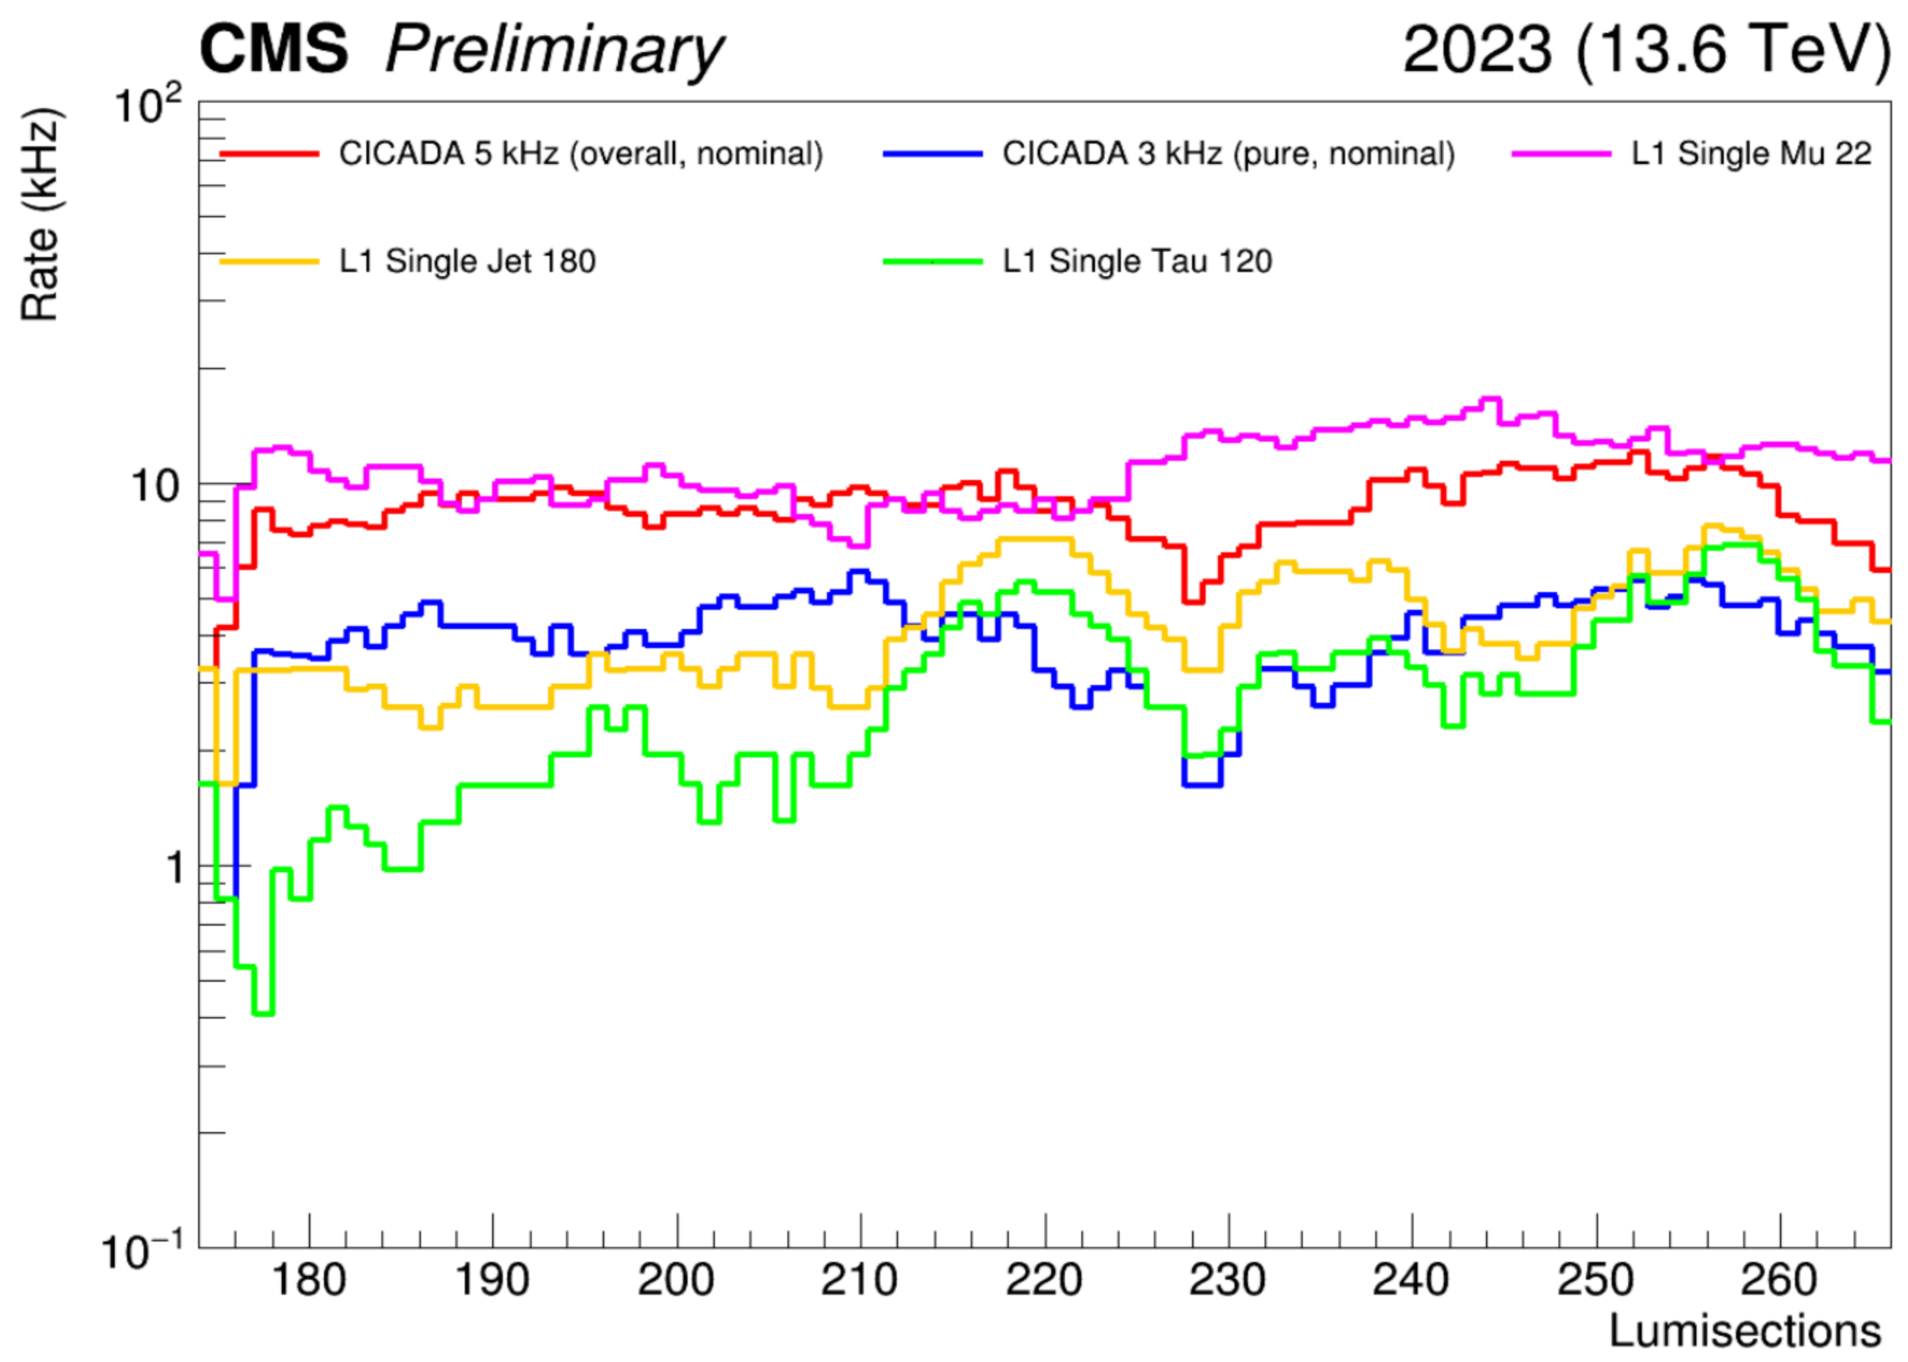

What Does a CICADA Event Look Like?

- Tends to accept high-multiplicity events, for any kind of object or combination of objects
- Tends to accept events with hard objects
- Anomaly score slightly correlated with event-level energy (HT)

CMS Event Only Fired by CICADA

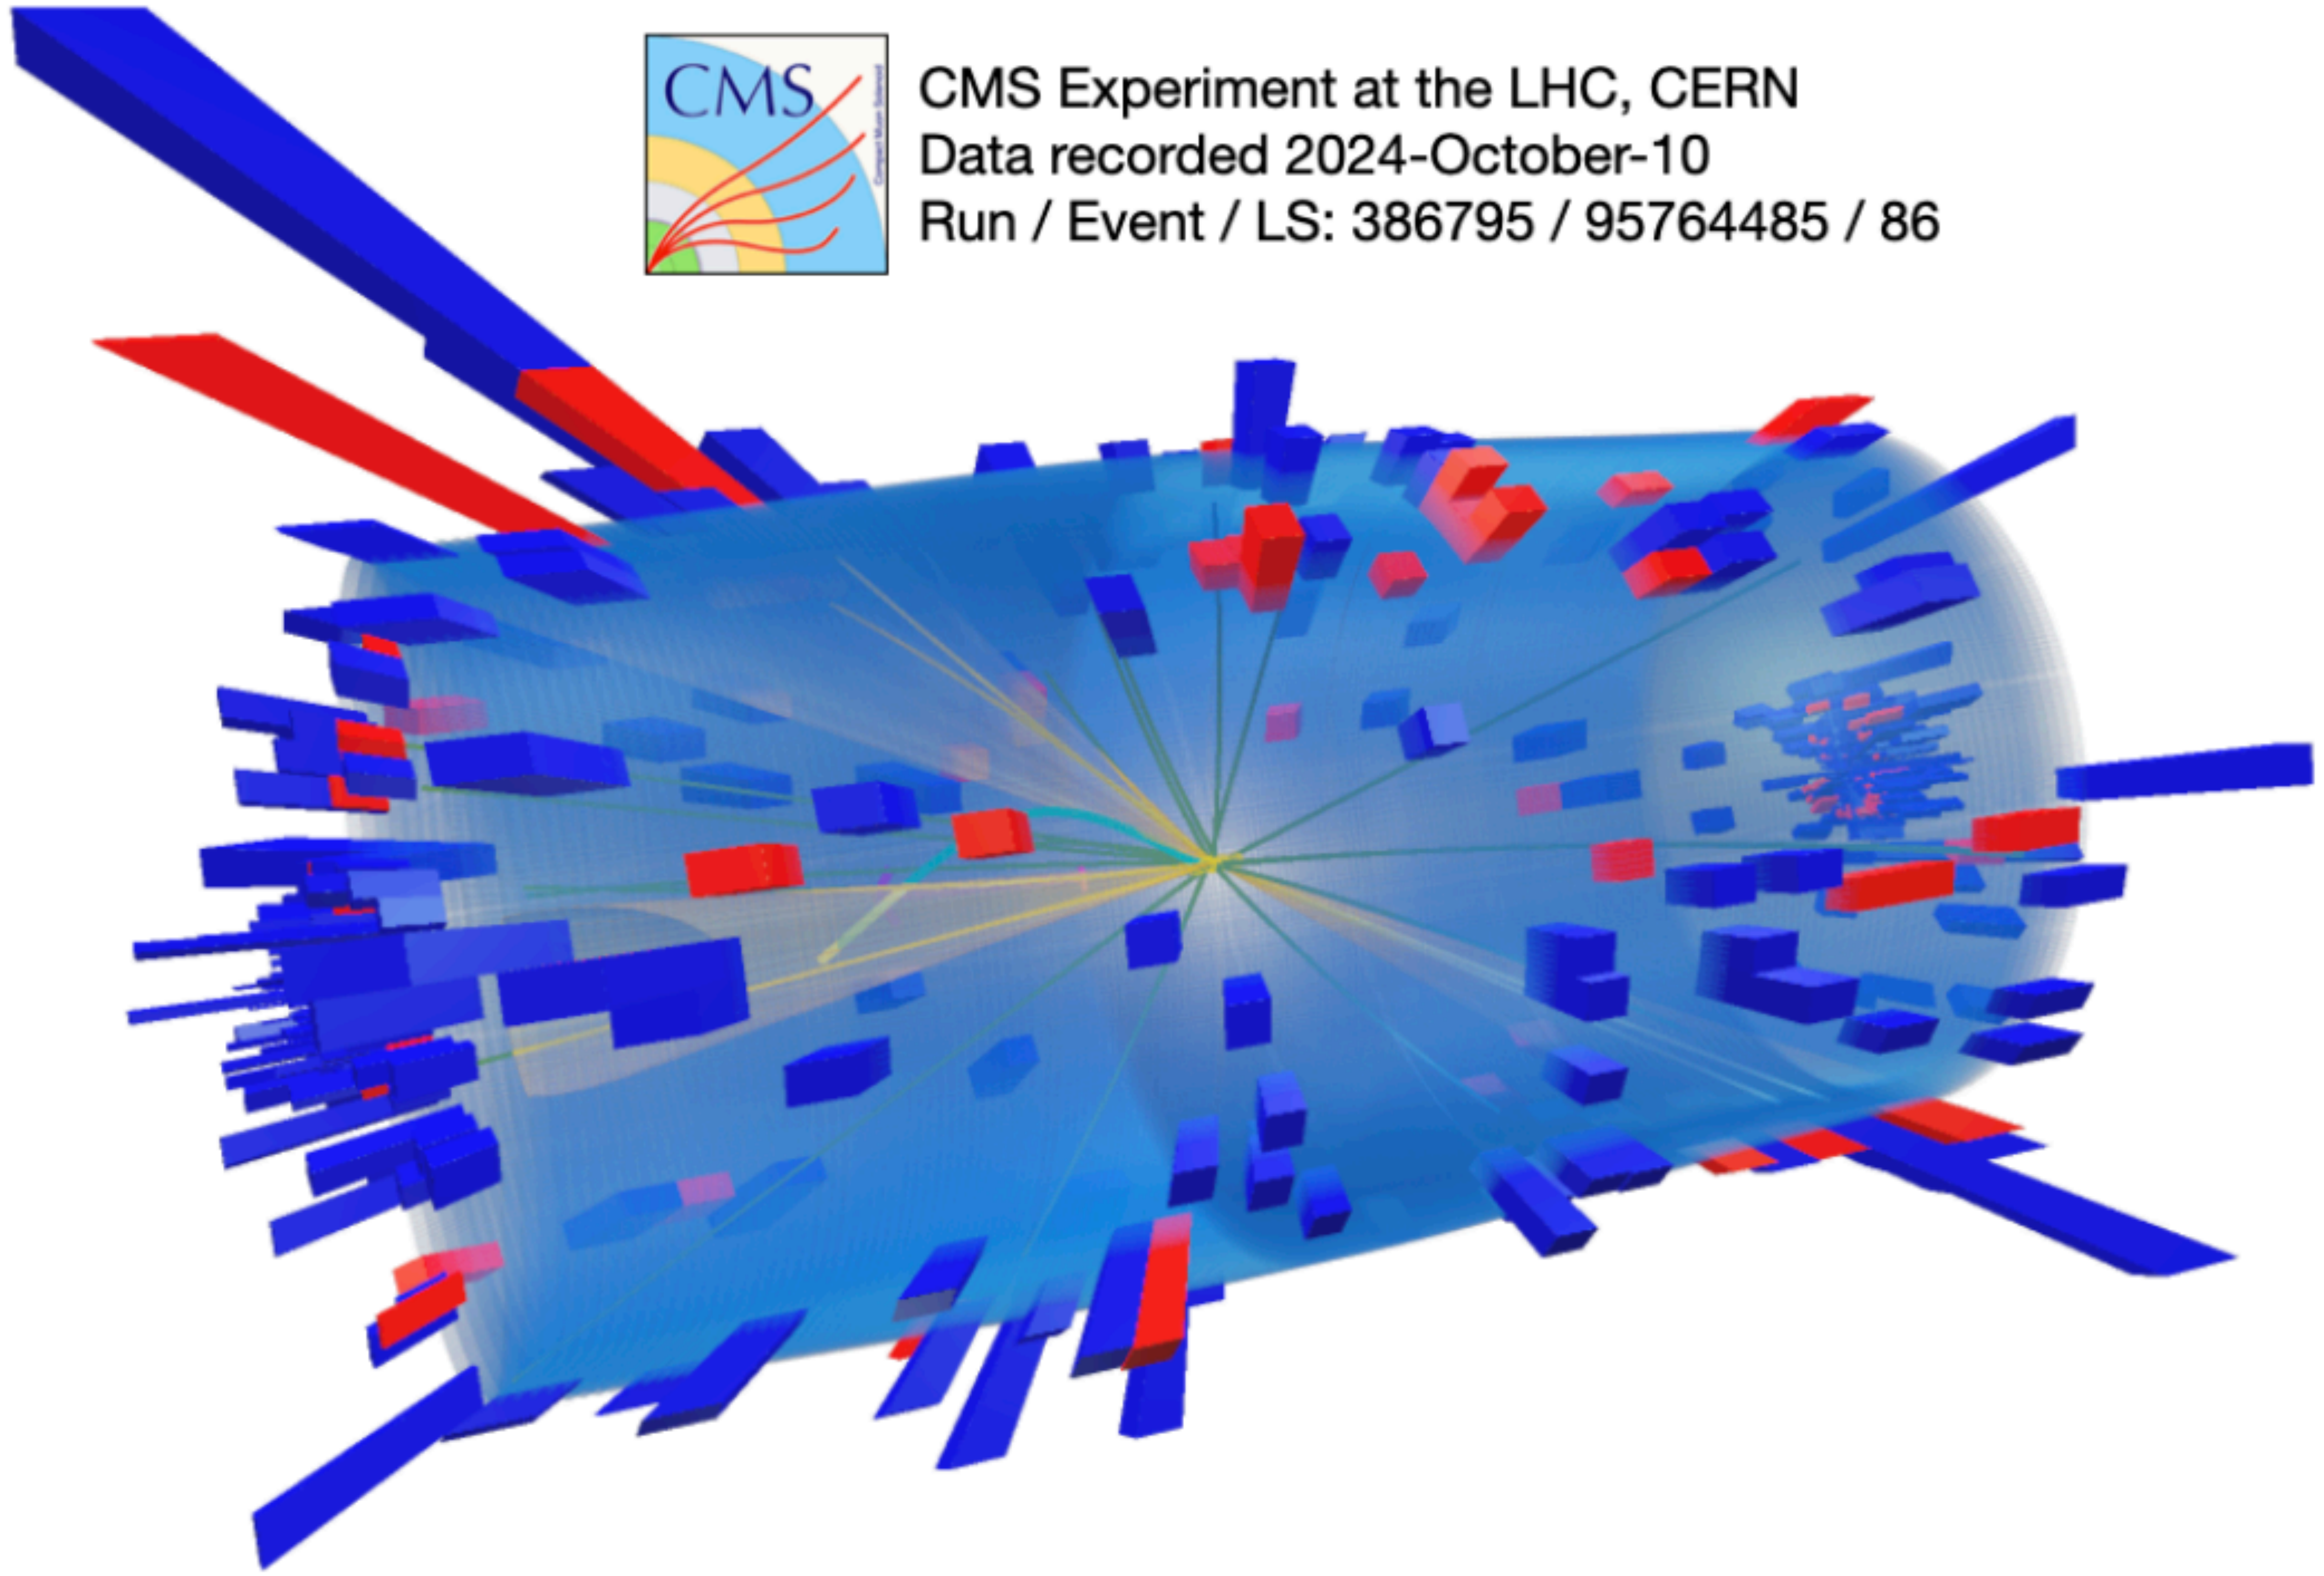

CICADA Score Distribution

- Excellent separation between **Zero Bias** background and variety of signal models
- Observe moderate pileup dependence (**low**, **high**)
  - For signal evaluation, need to carefully reweight number of primary vertices to desired distribution

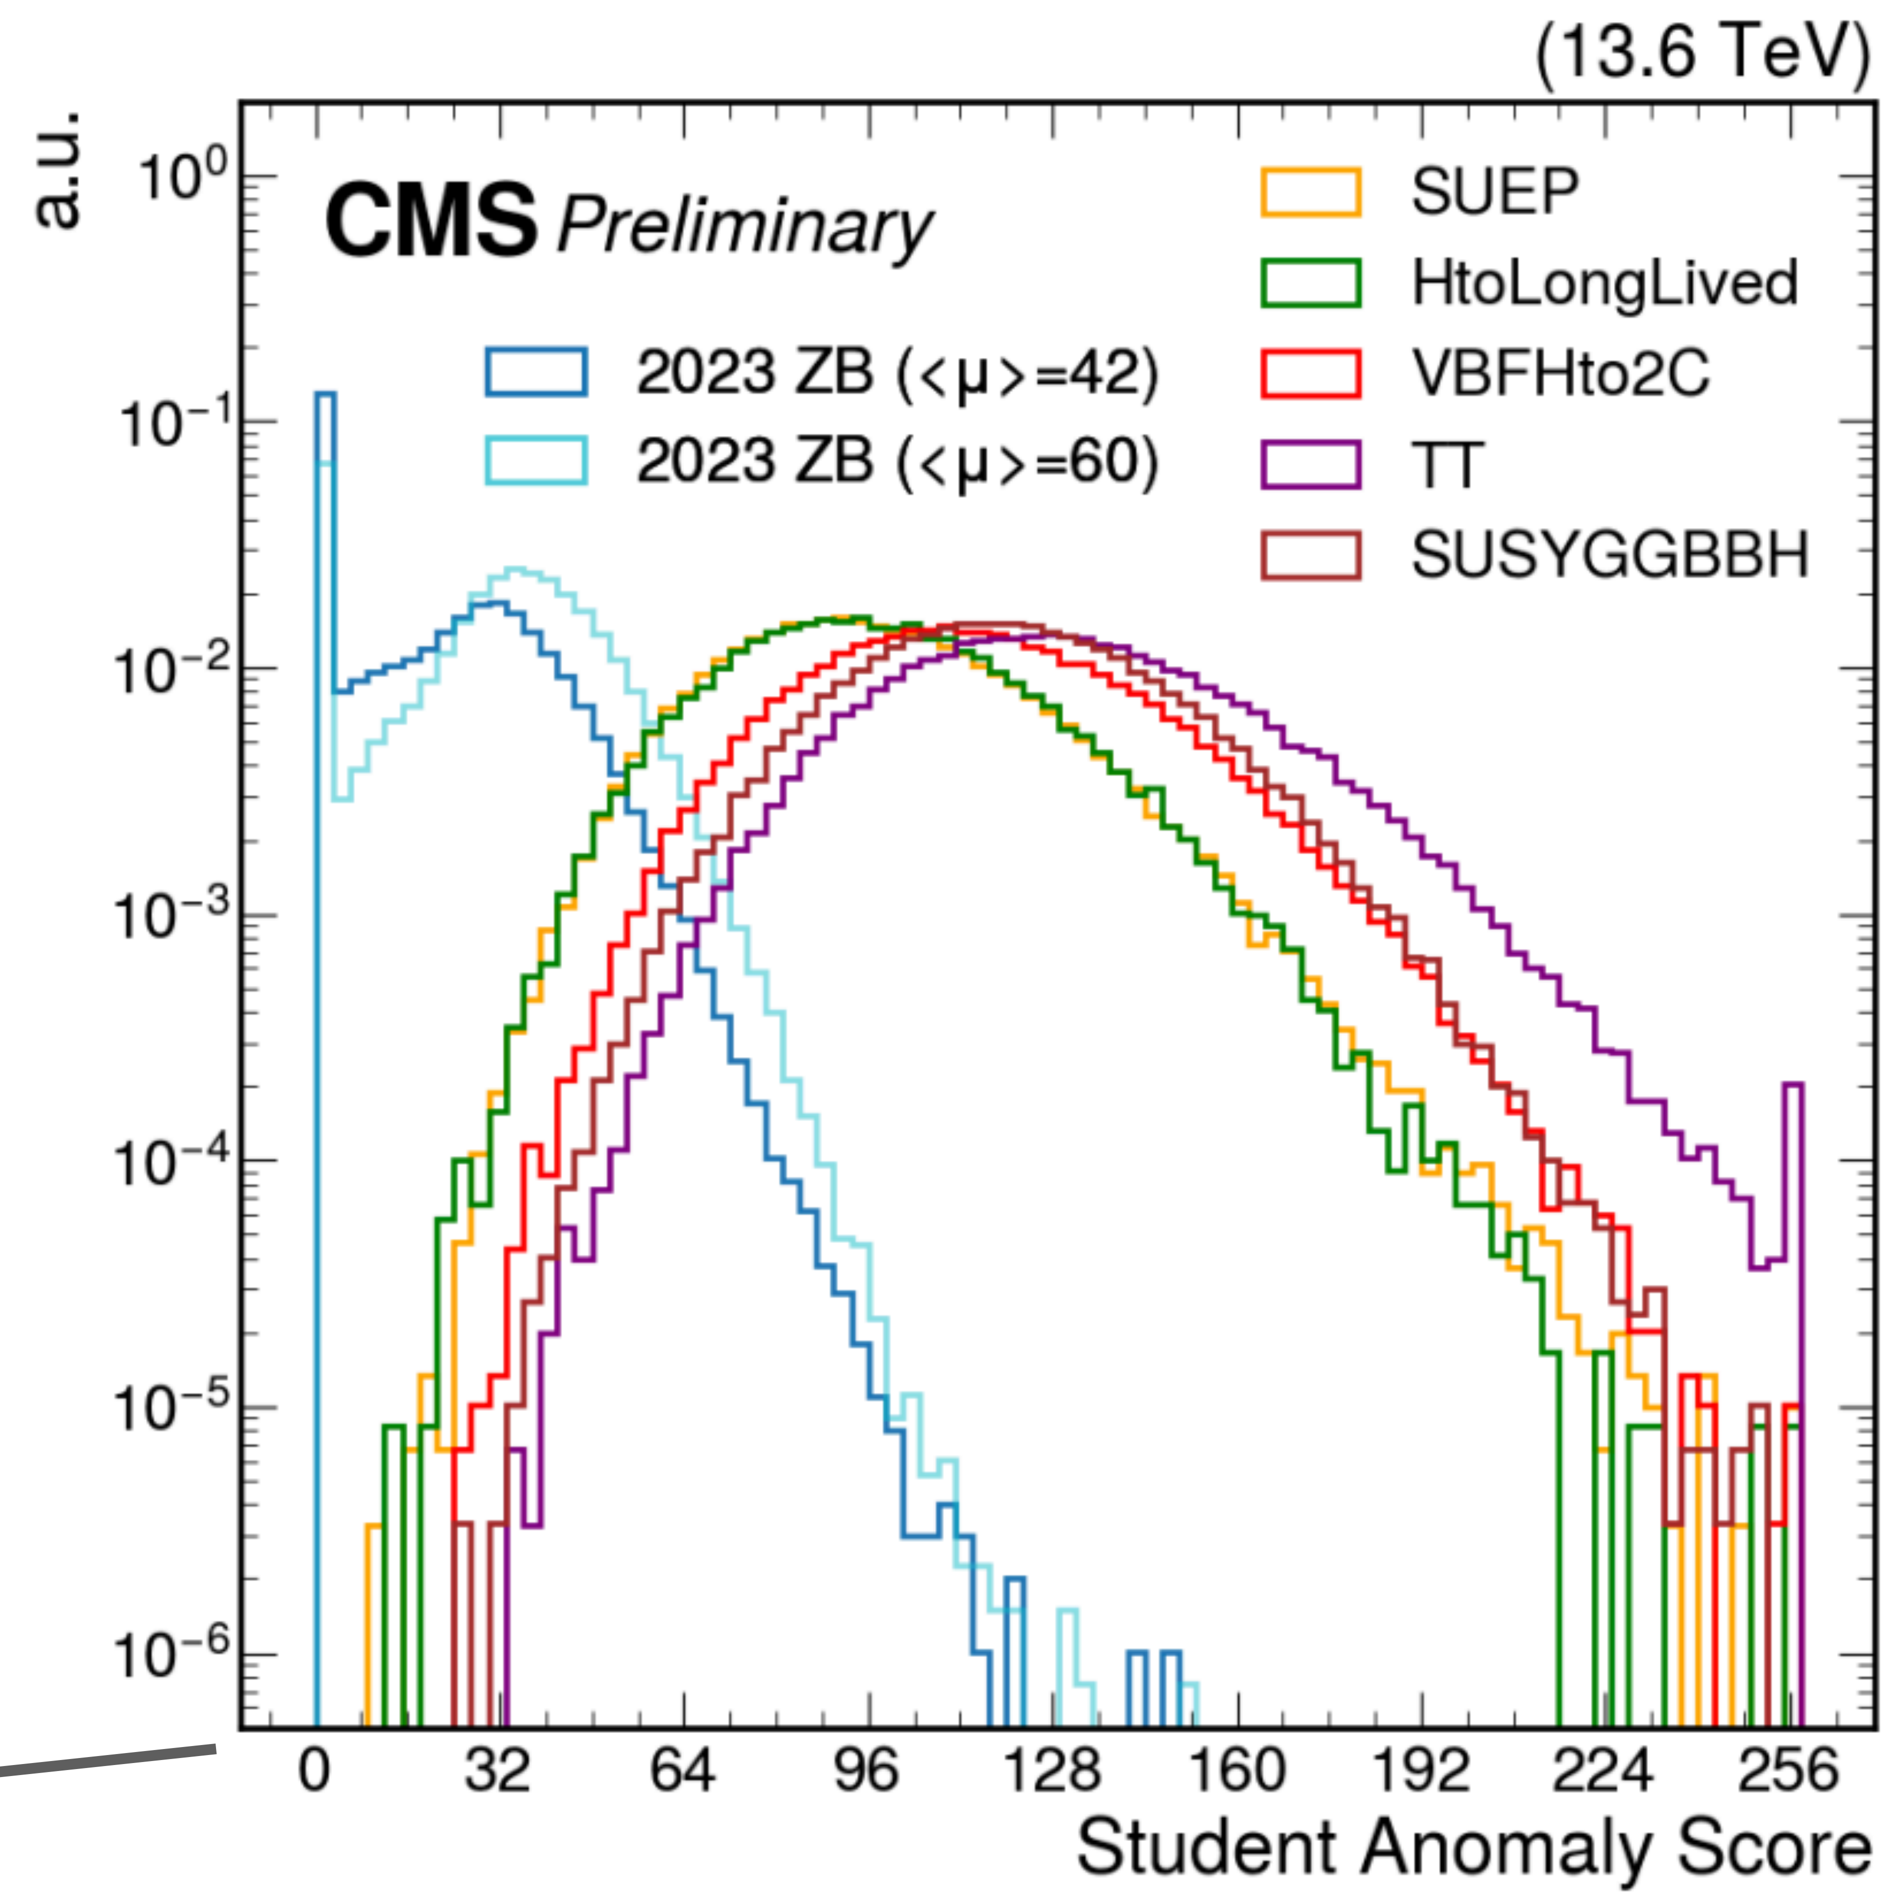

2^{16} discrete values in range $[0,2^8]$

CICADA Physics Performance

- Results given vs. trigger rate
  - "Realistic" range: 50 Hz - 10 kHz
- Student outperforms the teacher!
  - AD benefits from model compression
  - Teacher is "too good" at reconstructing events
- Most AUCs > 0.99
  - For more challenging PU, most AUCs still > 0.9

Teacher

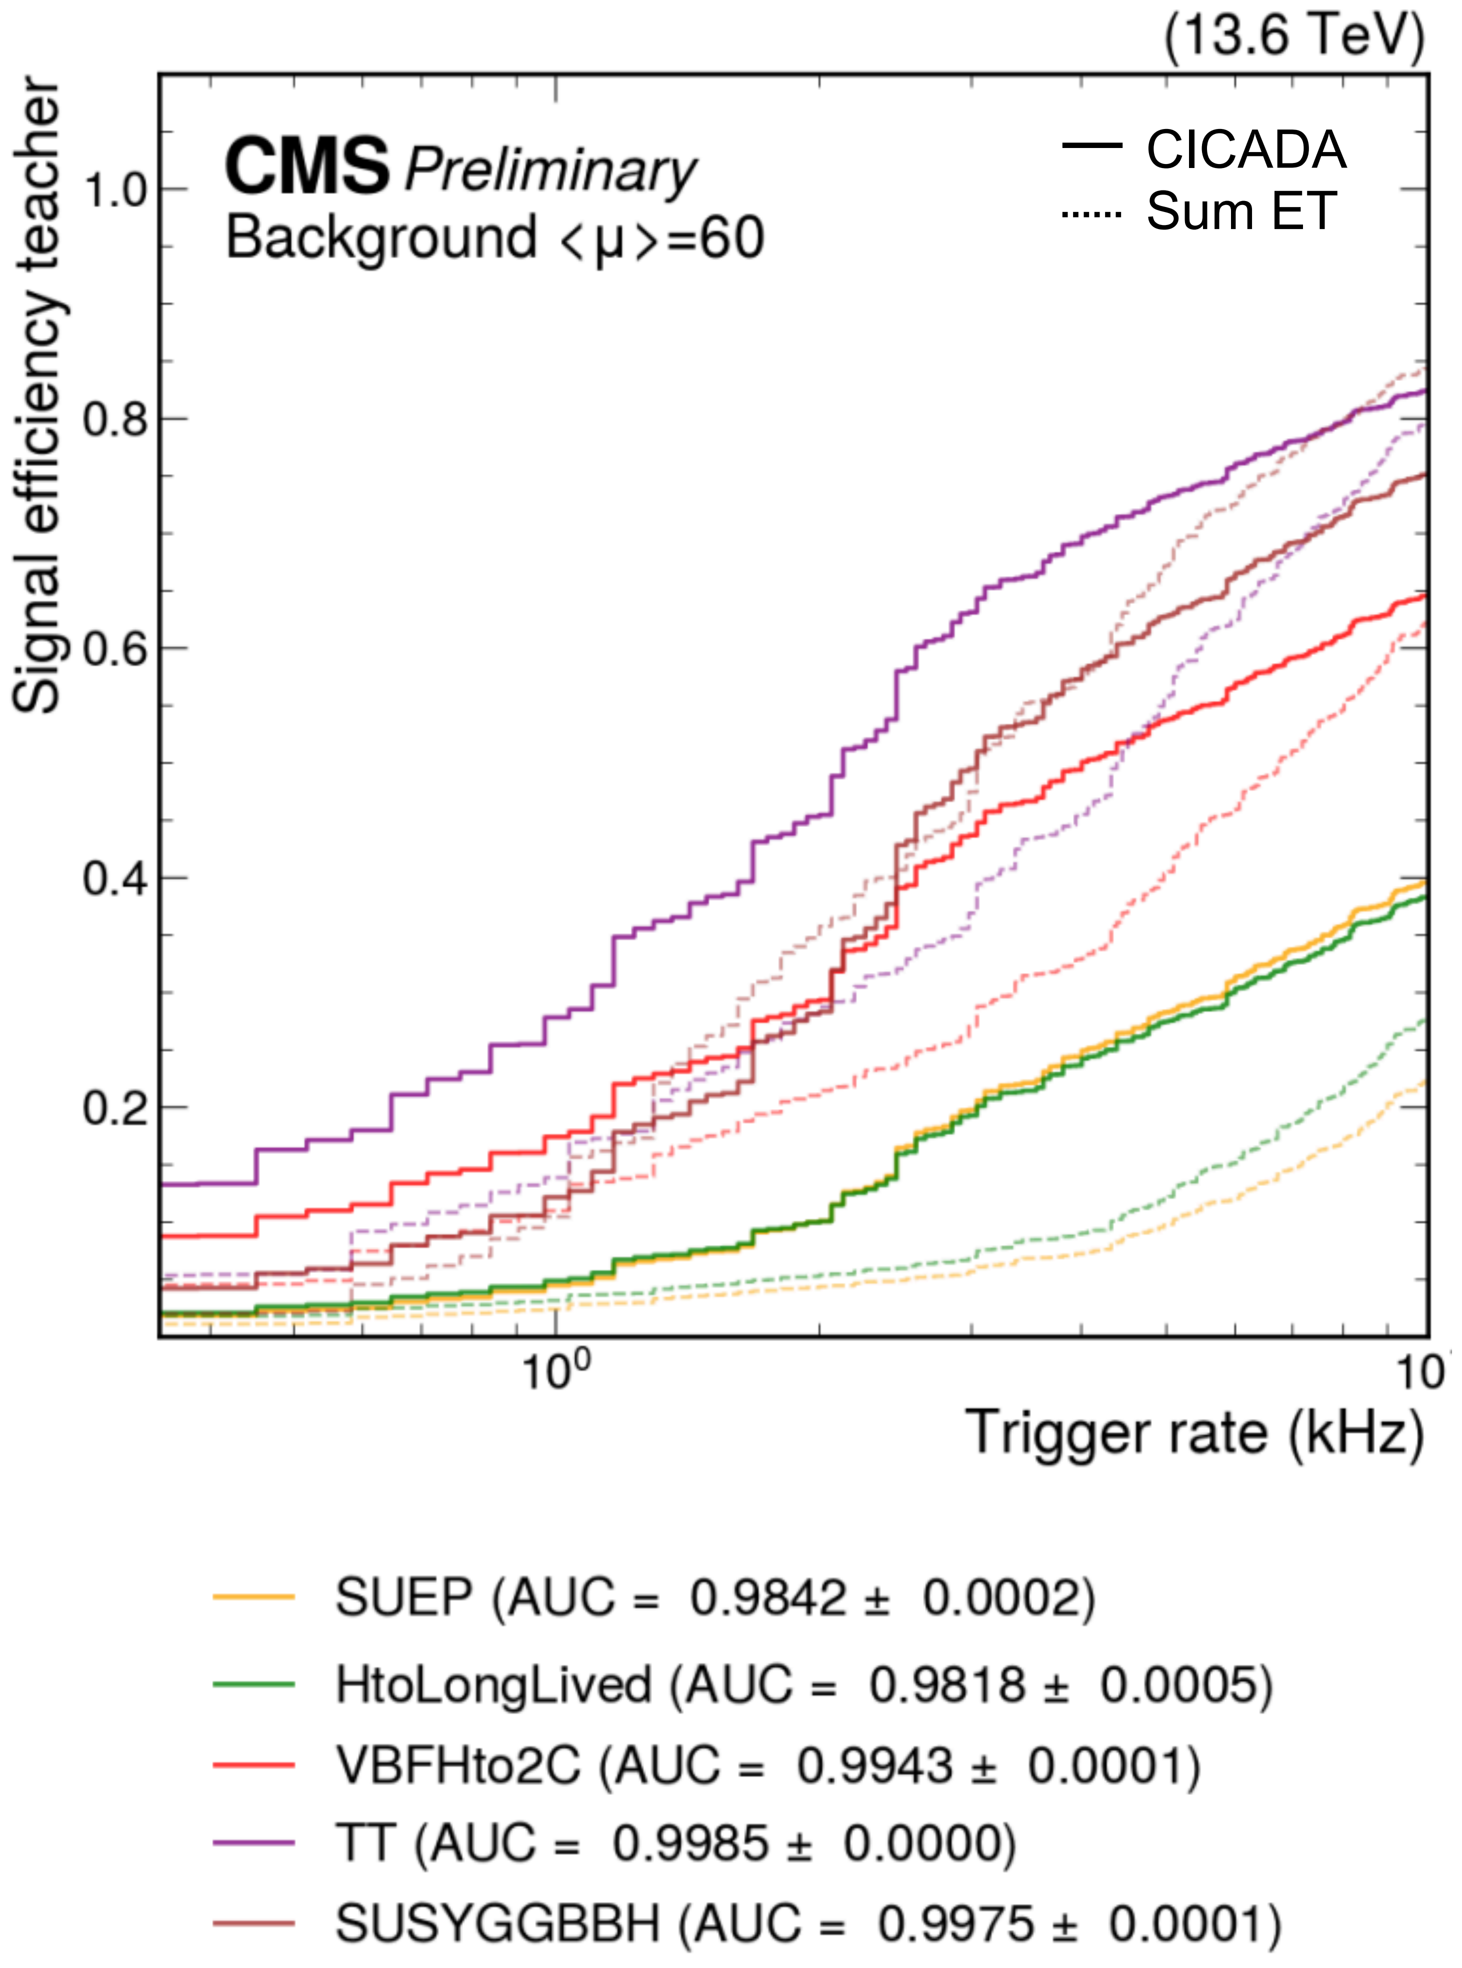

Student

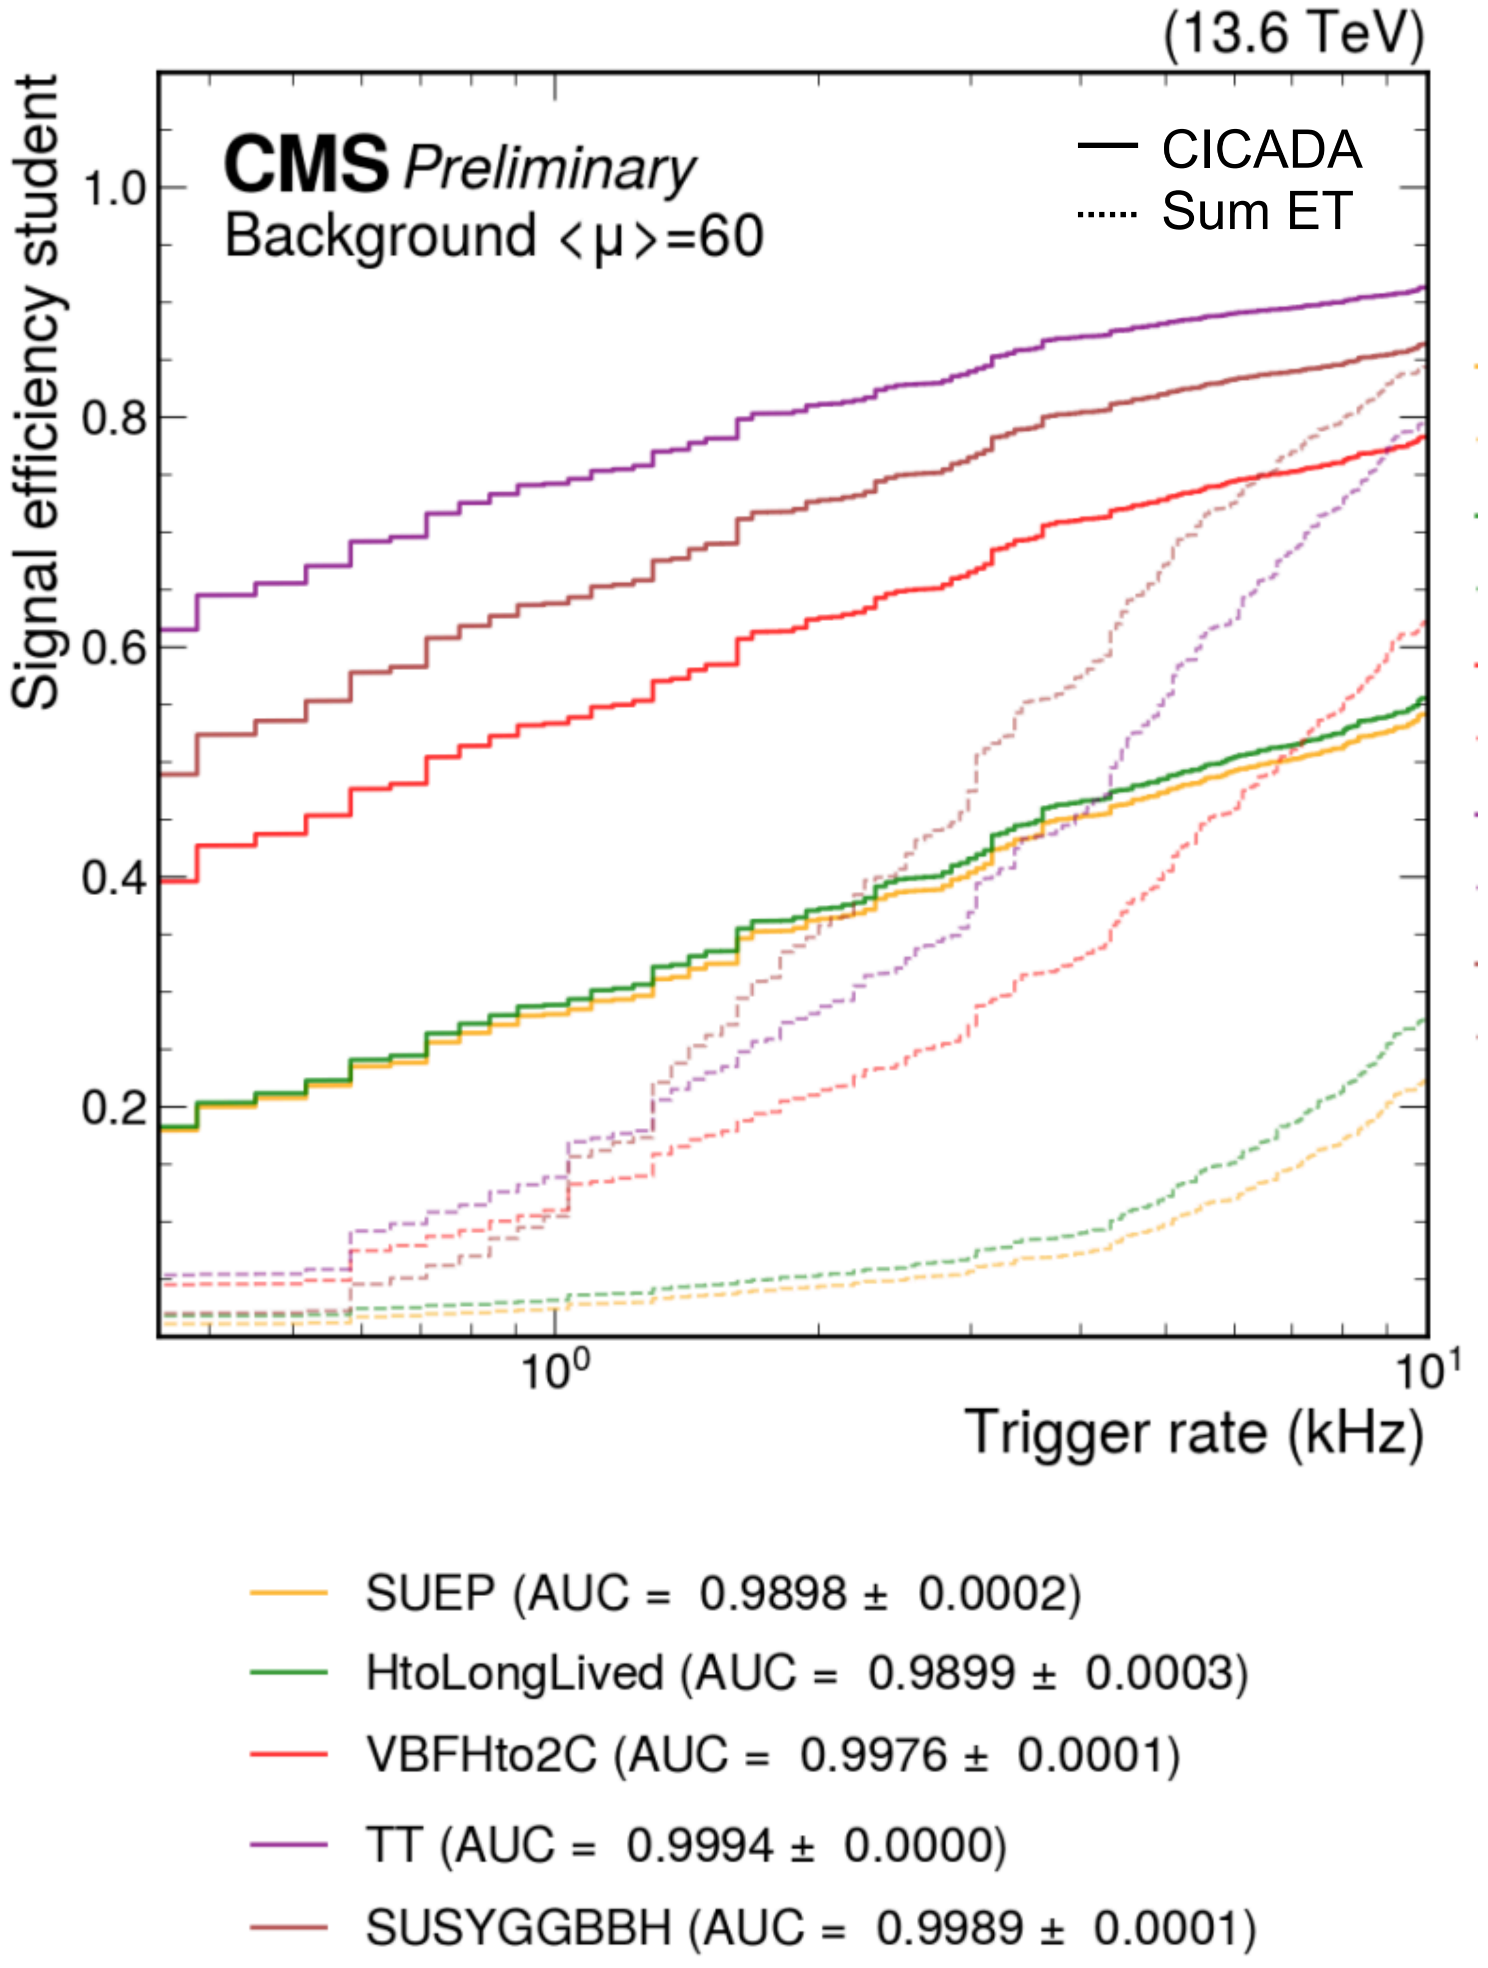

Validation with TTbar in Data

CICADA successfully selects semileptonic ttbar events – trigger on “interesting” physics in data

- Reconstruct hadronic top mass
  - Background = ZB data with no CICADA cut
  - Data = ZB data with CICADA score > 115

- Offline selection: ≥ 3 jets, ≥ 1 b-tagged jet

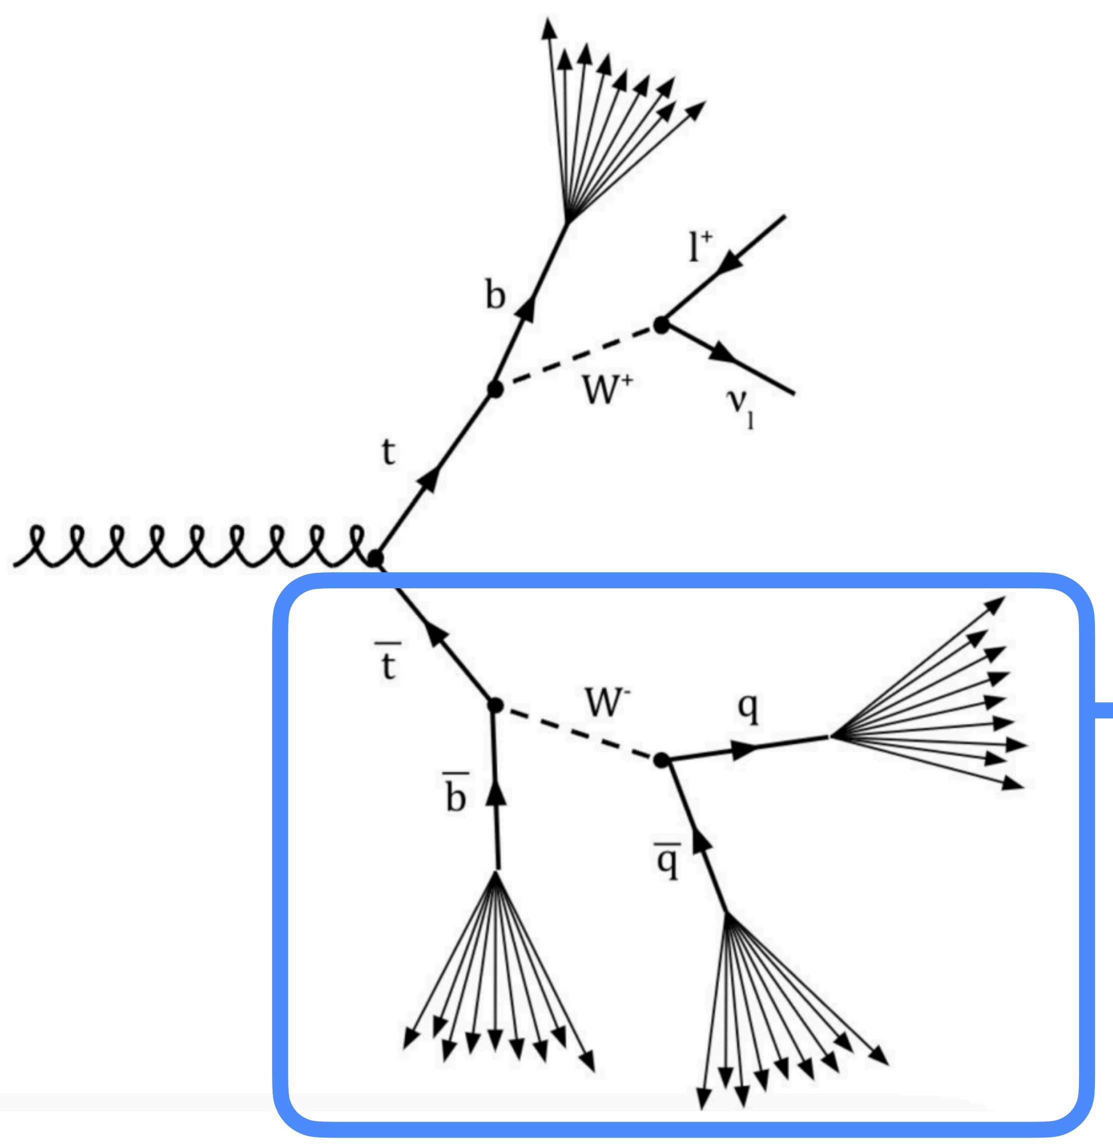

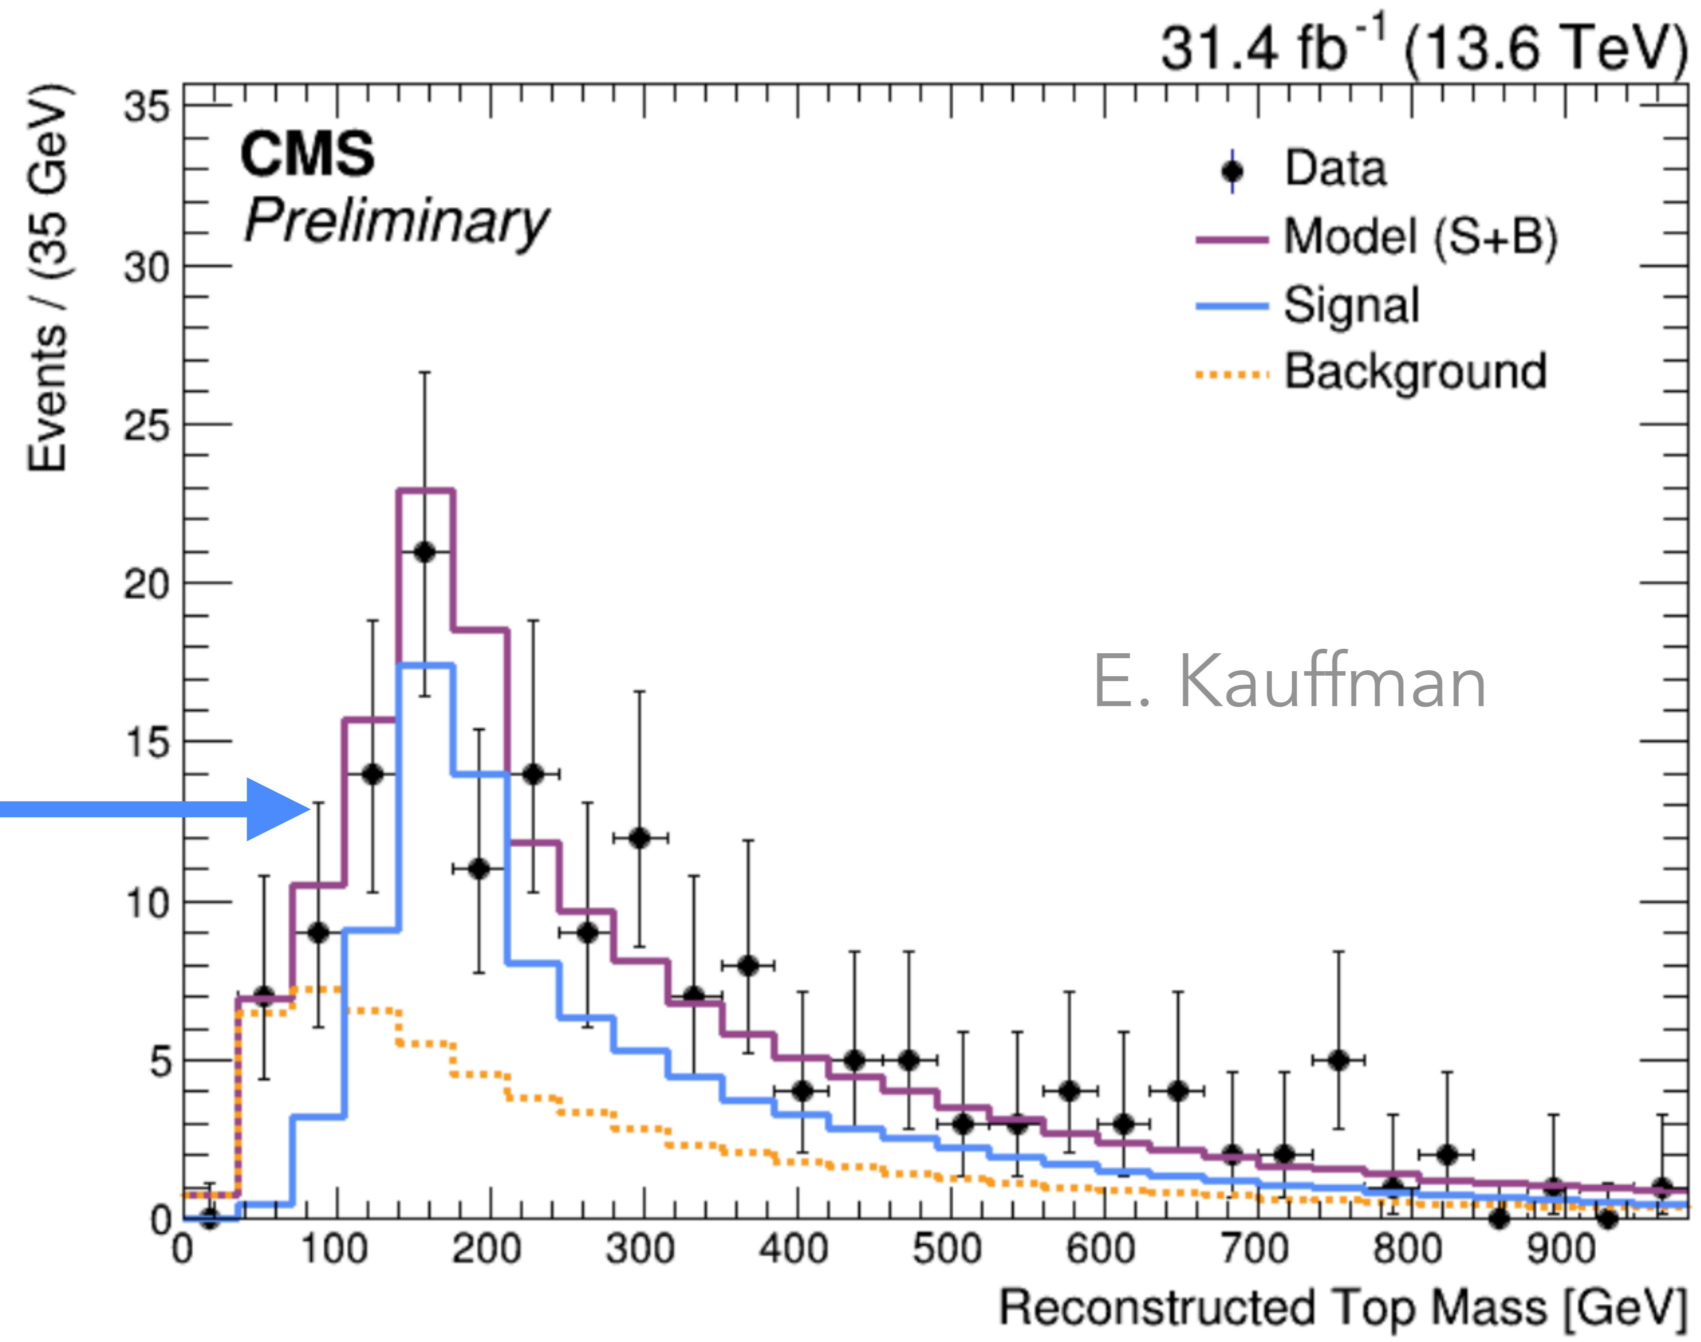

AD@L1 in the HL-LHC Era

- **Improved ID**: distinguish between ECAL/HCAL (low-level), particle flow (high level)
- **Higher granularity**: utilize full granularity of calorimeter at L1, including HCAL depth
- **Integrate tracking**: tracks available at L1 with CMS Outer Tracker upgrade

CMS HCAL Depth Segmentation

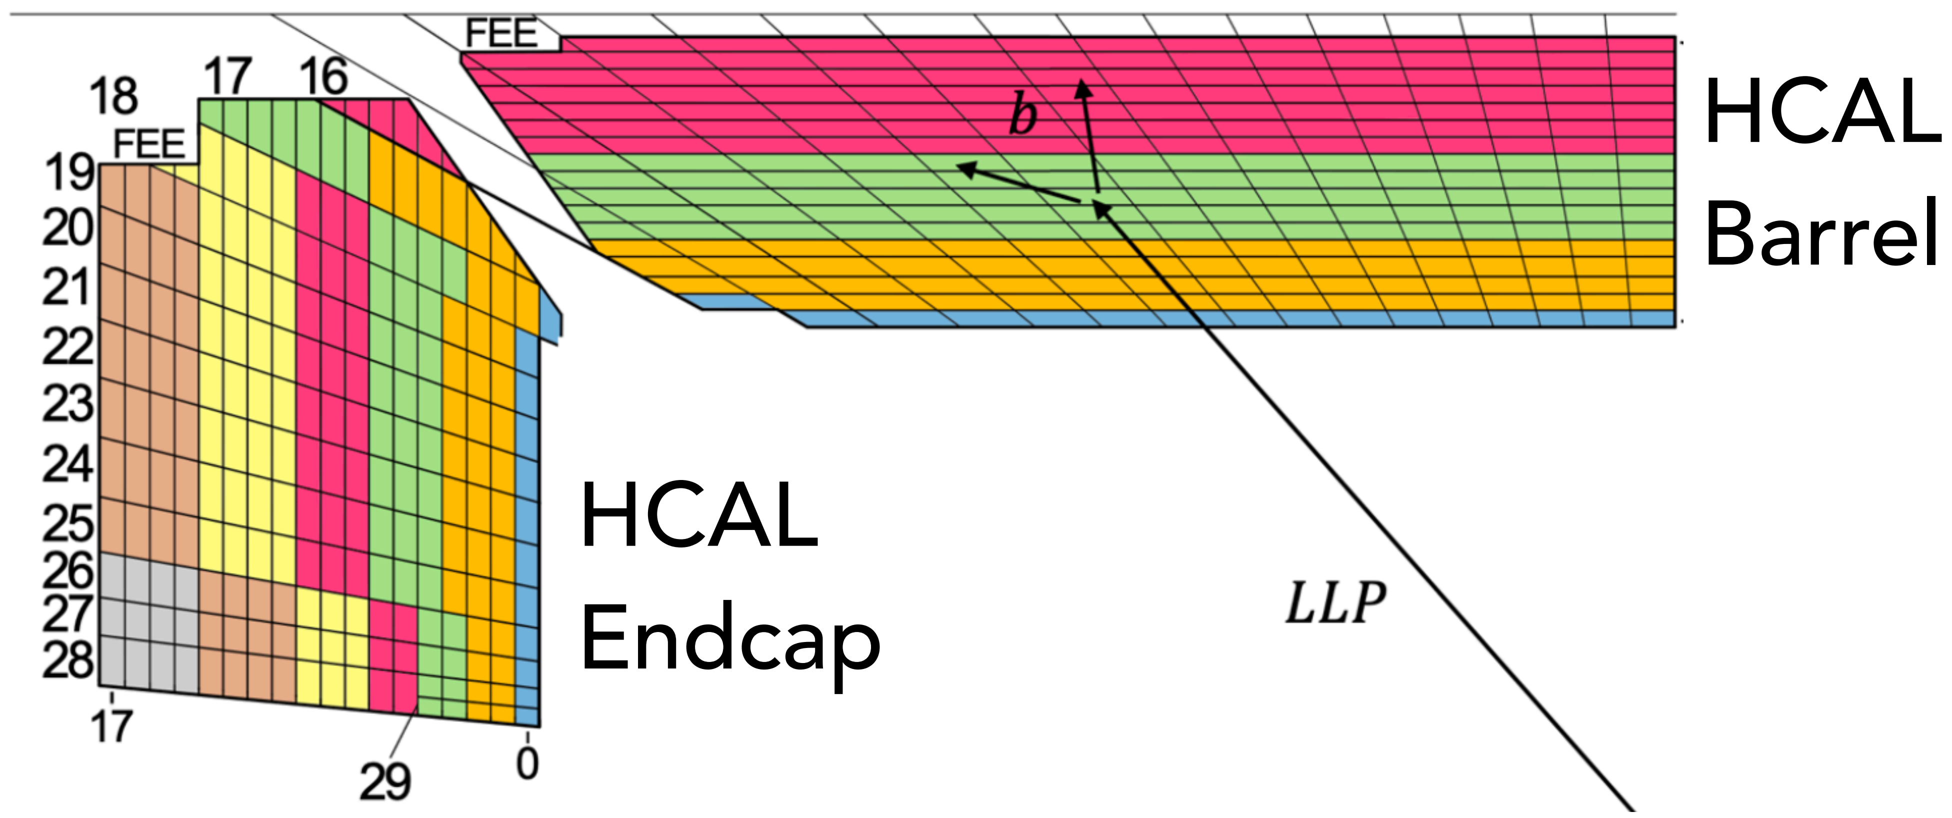

CMS Level-1 Trigger Updgrade

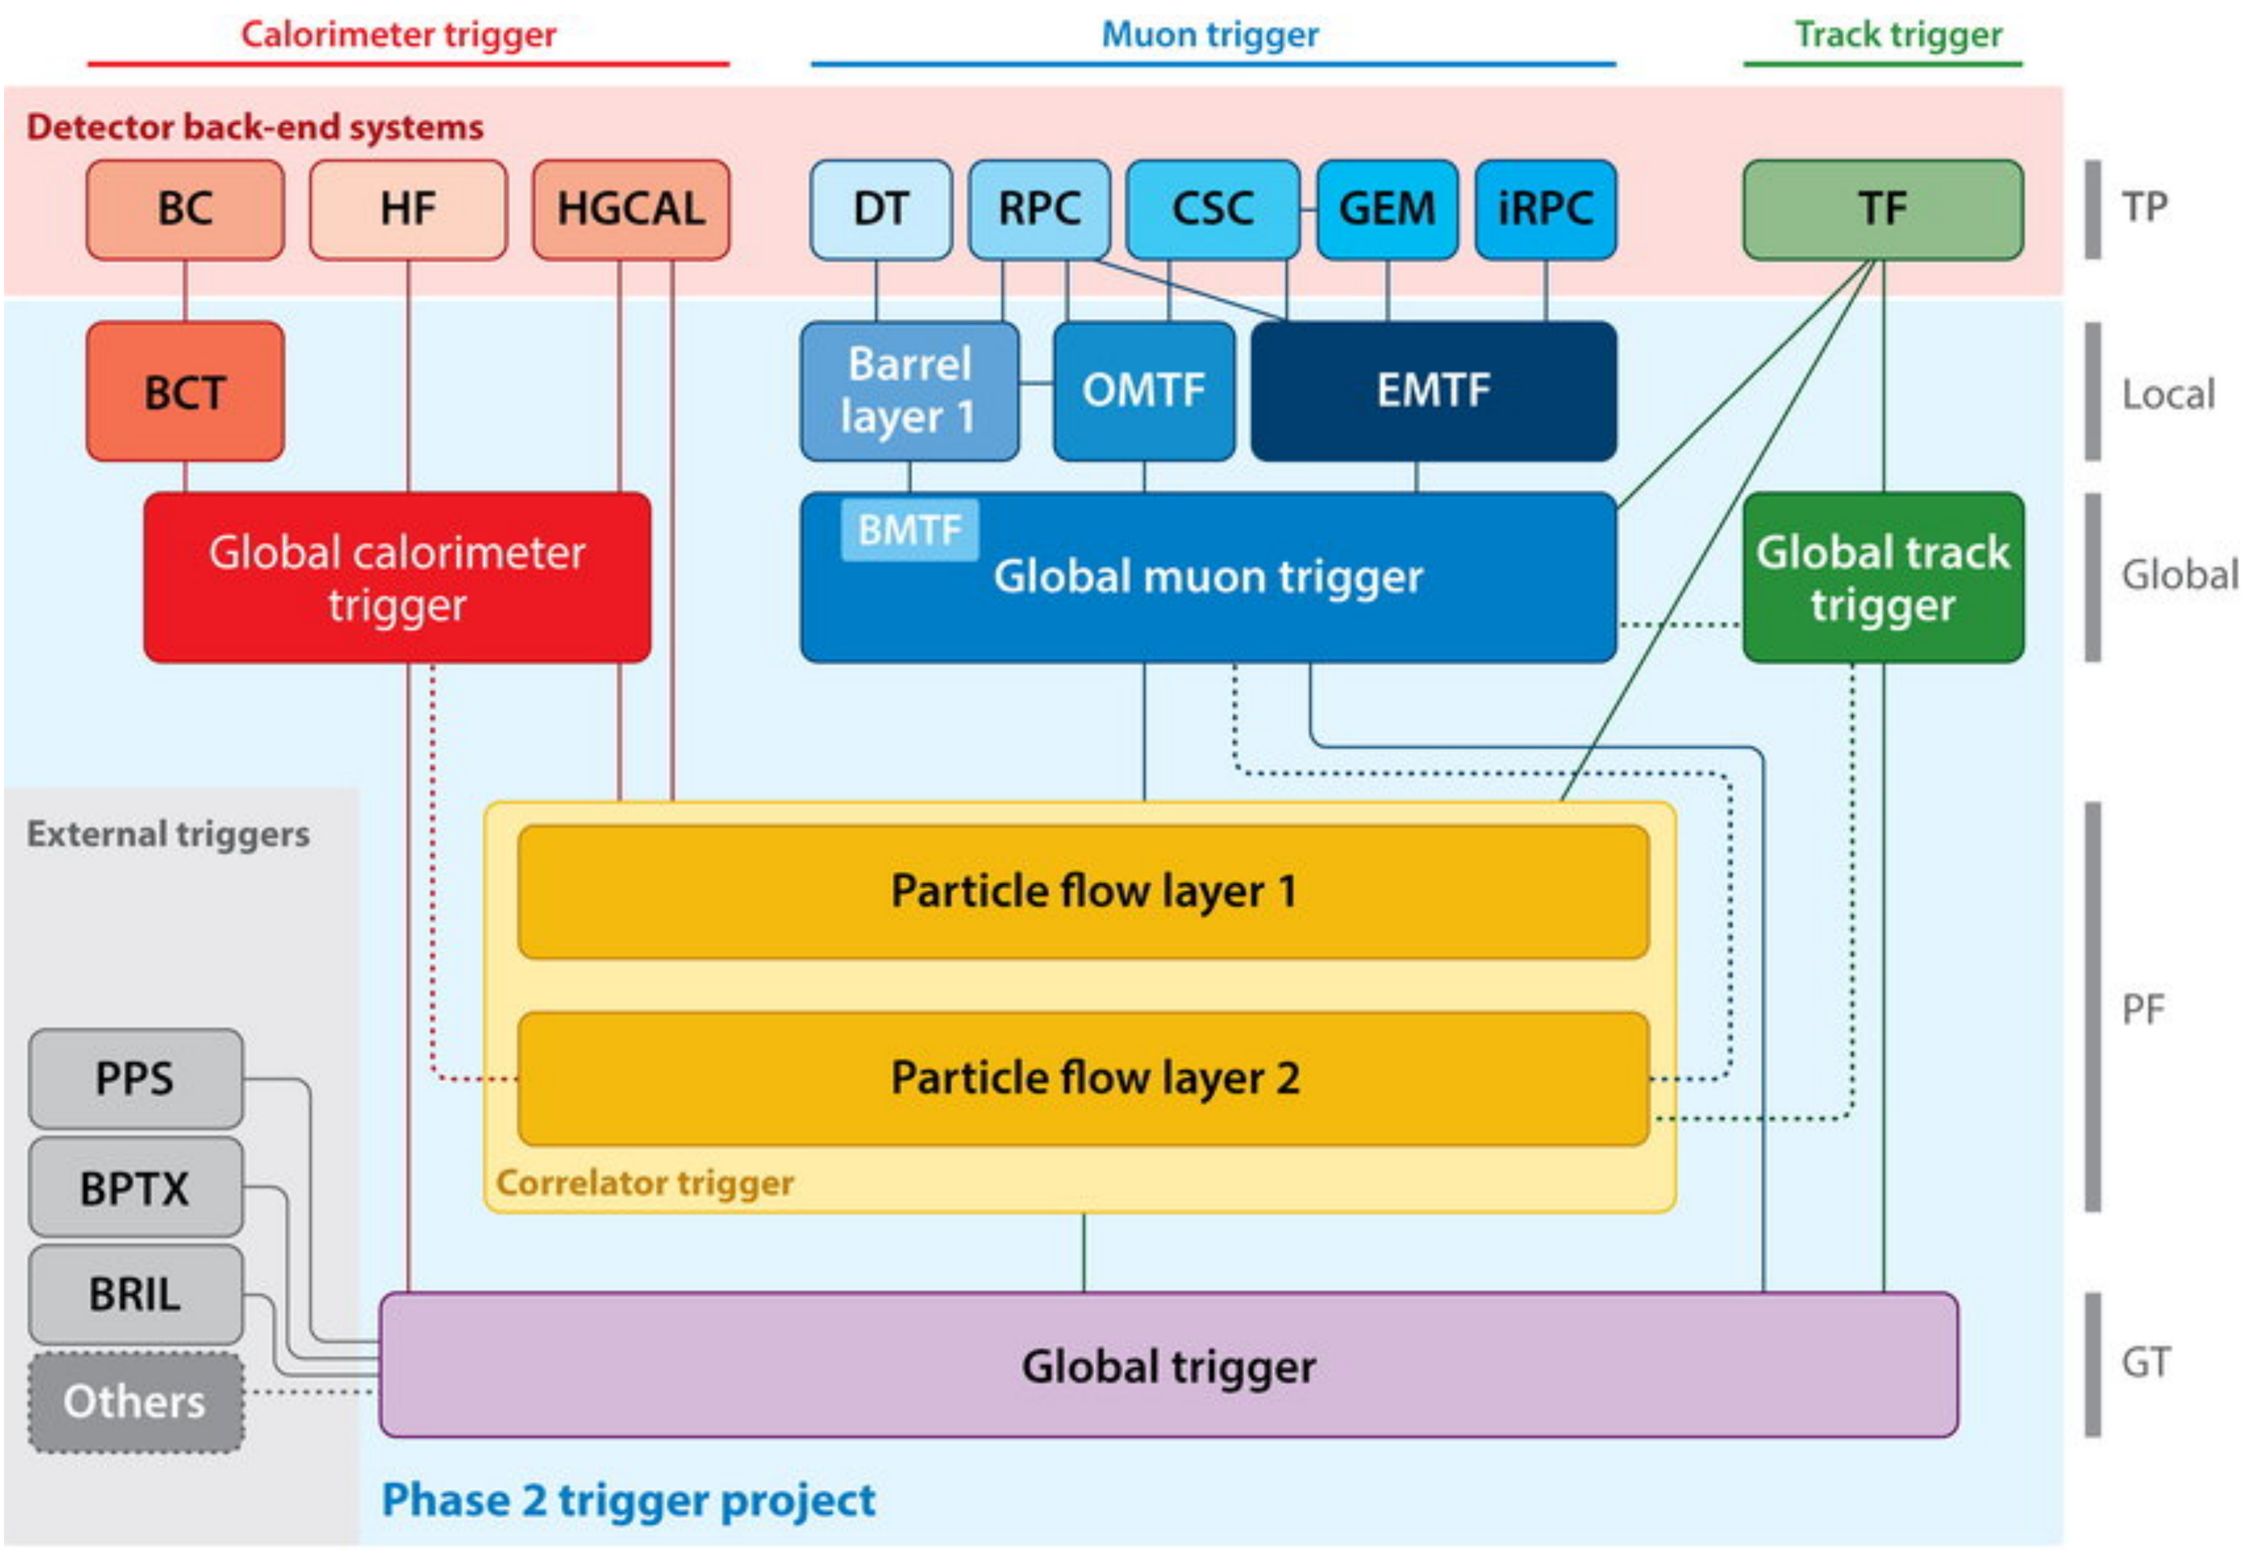

Conclusions + Outlook

- CICADA is the first Level-1 AD trigger using low-level inputs deployed at ATLAS or CMS
  - Leveraged **knowledge distillation** and **quantization-aware training** to contend with challenging operating conditions (small + fast model on FPGA)
  - Successfully commissioned in 2024
- Excellent discrimination between wide variety of signal models and zero bias background
- Smooth operations in 2025, especially given recent model updates to improve PU robustness
- Early stages of analysis plans with 2024 and 2025 CMS data!

Thank You!

References

- CMS-DP-2023-086
- CMS-DP-2024-121

CMS Trigger: Parking + Scouting

Data flow for a typical 2018 data-taking scenario

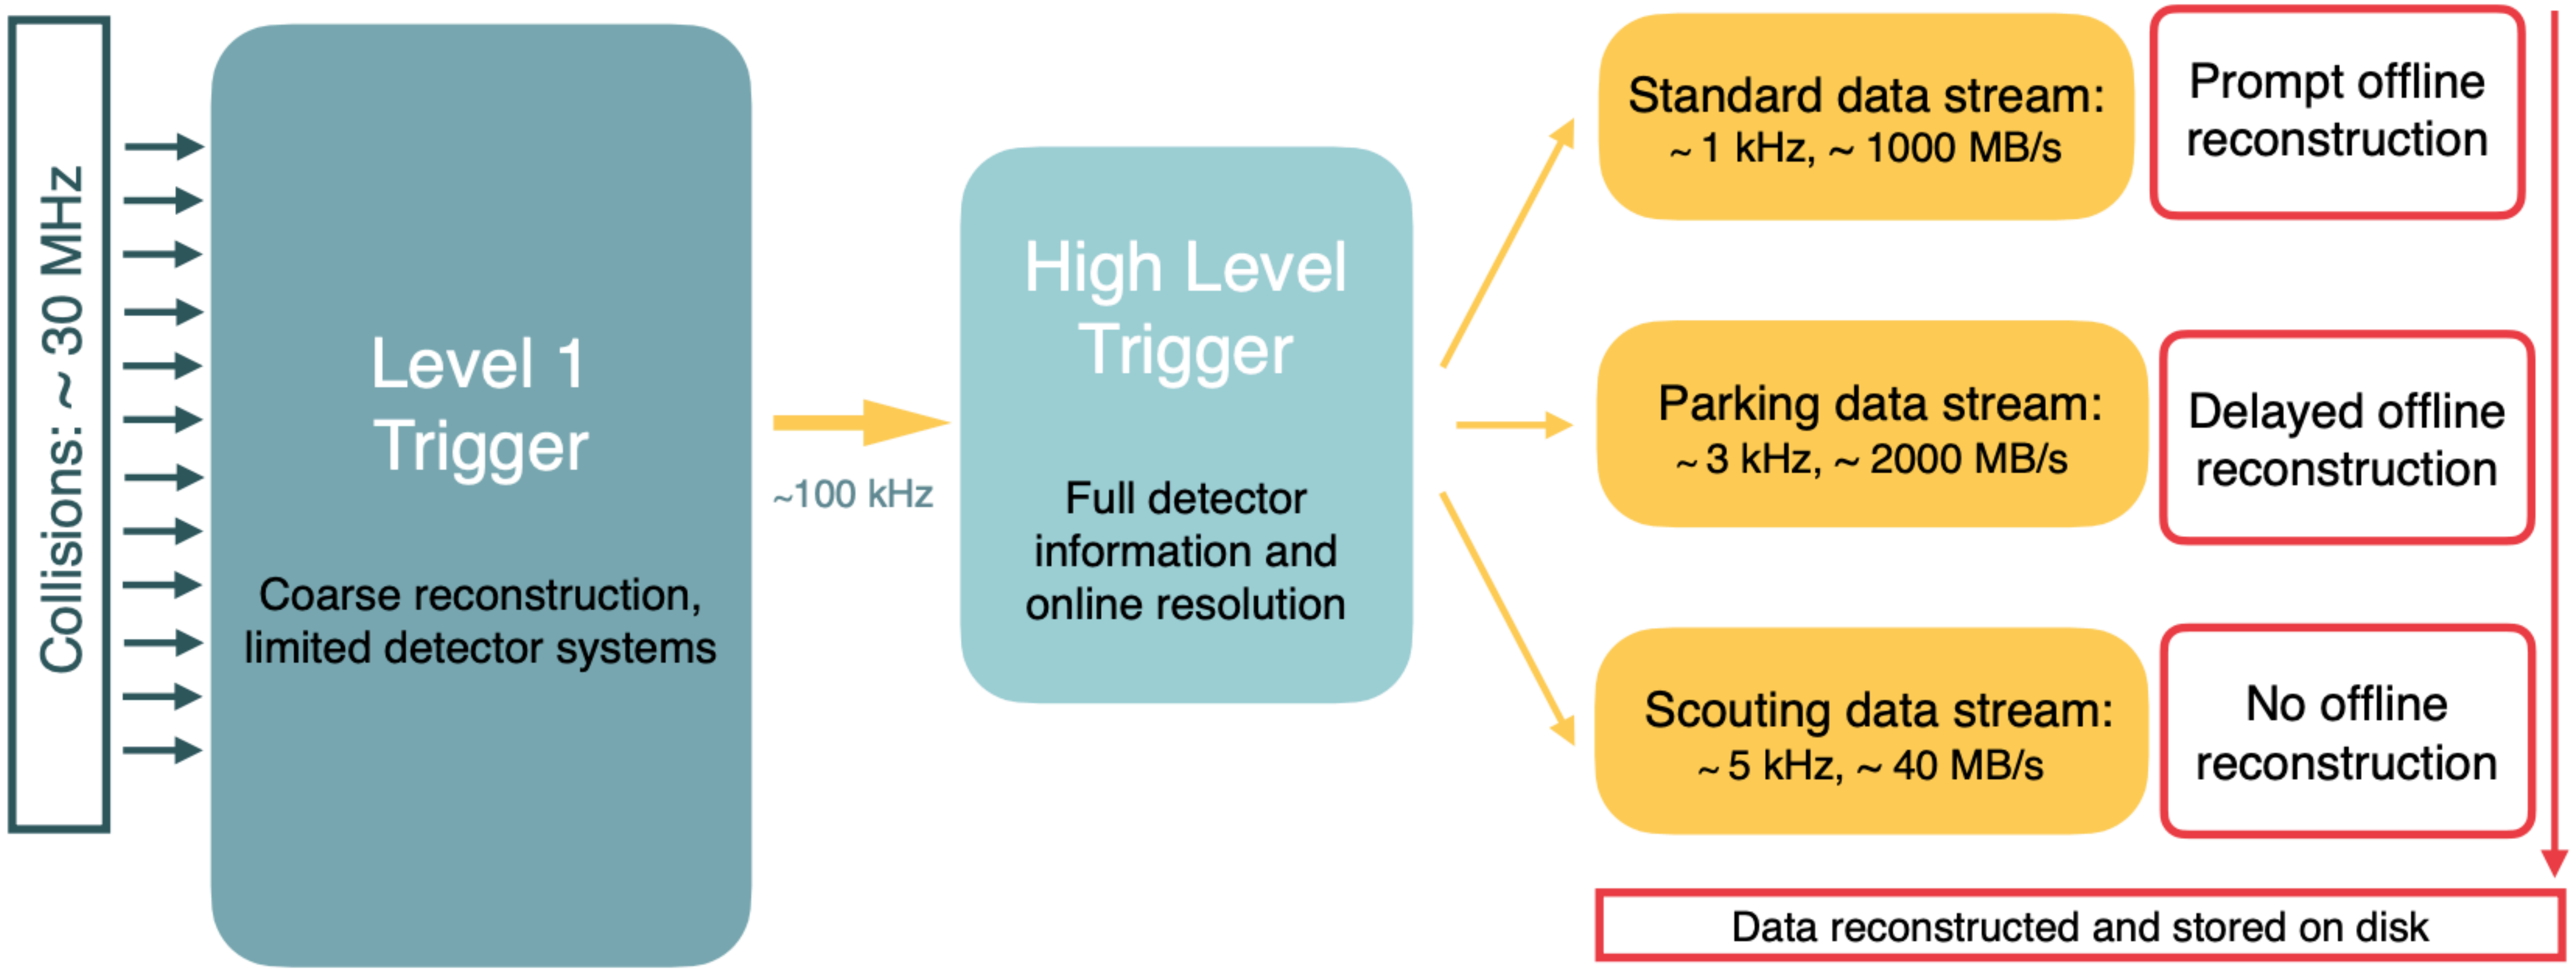

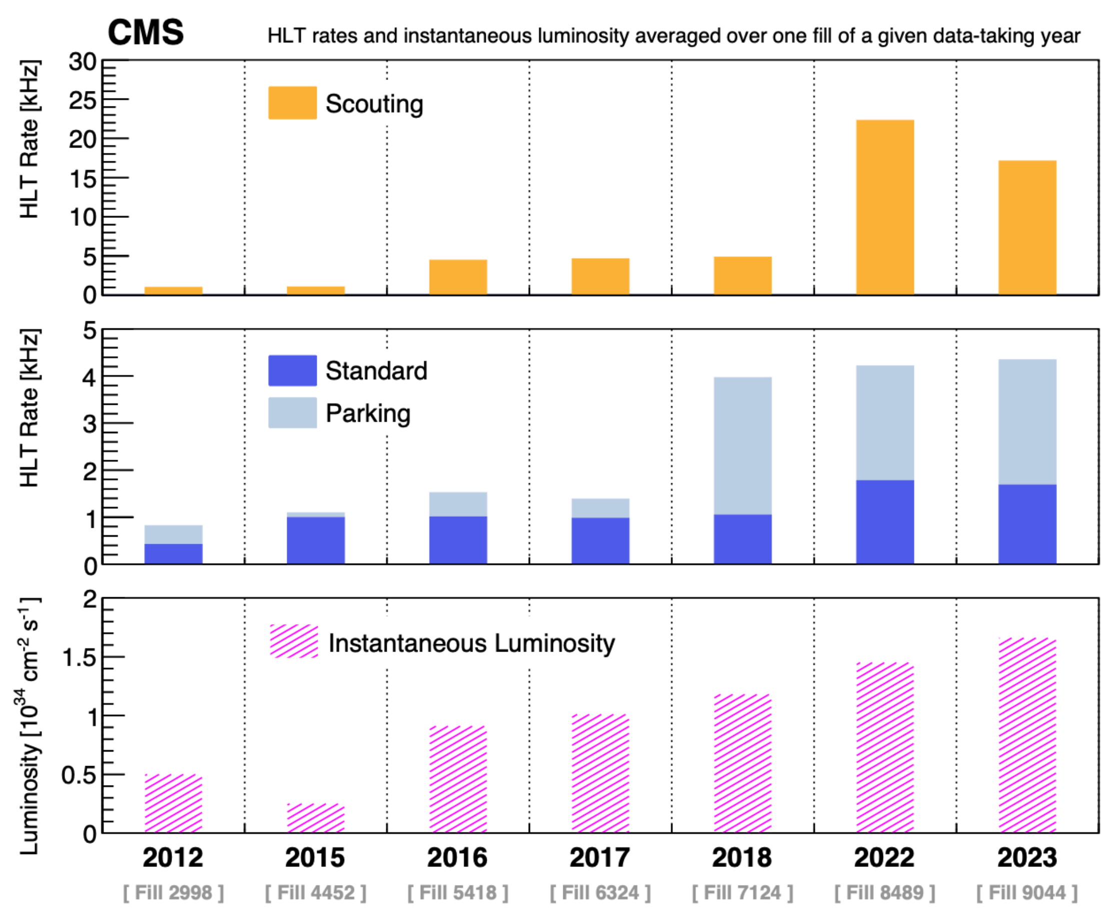

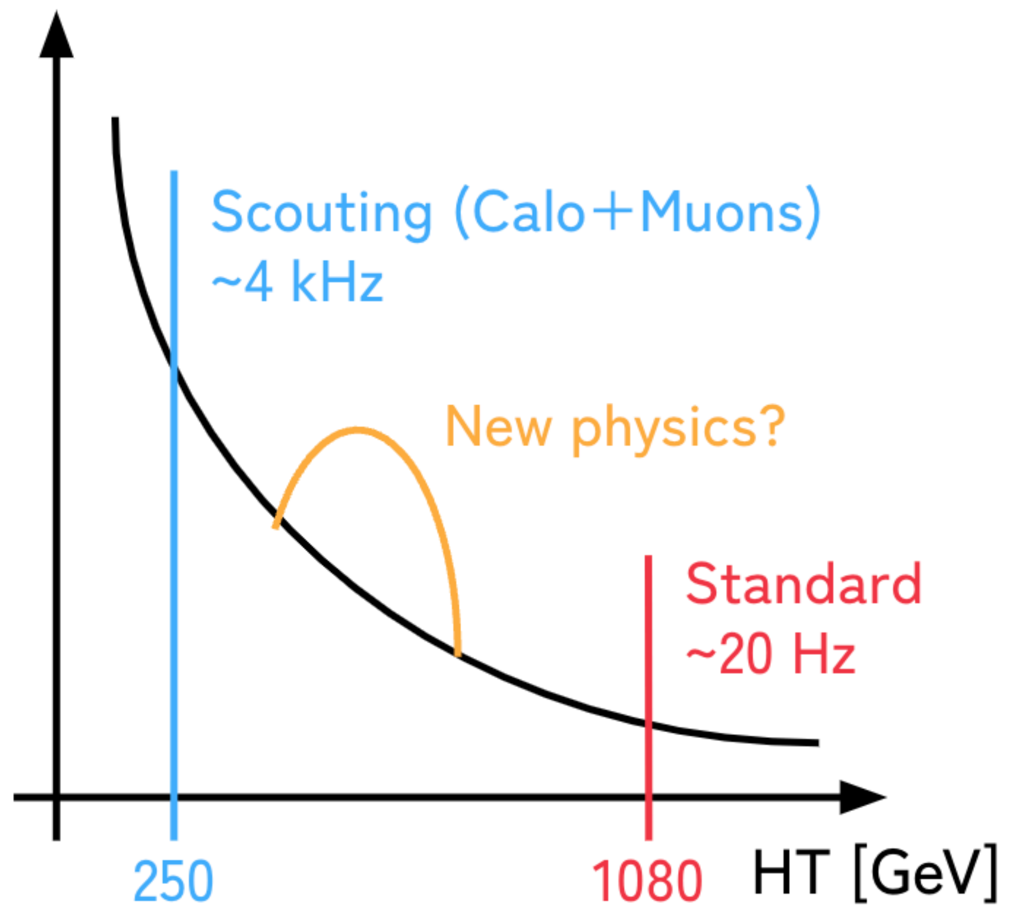

Anomaly Score

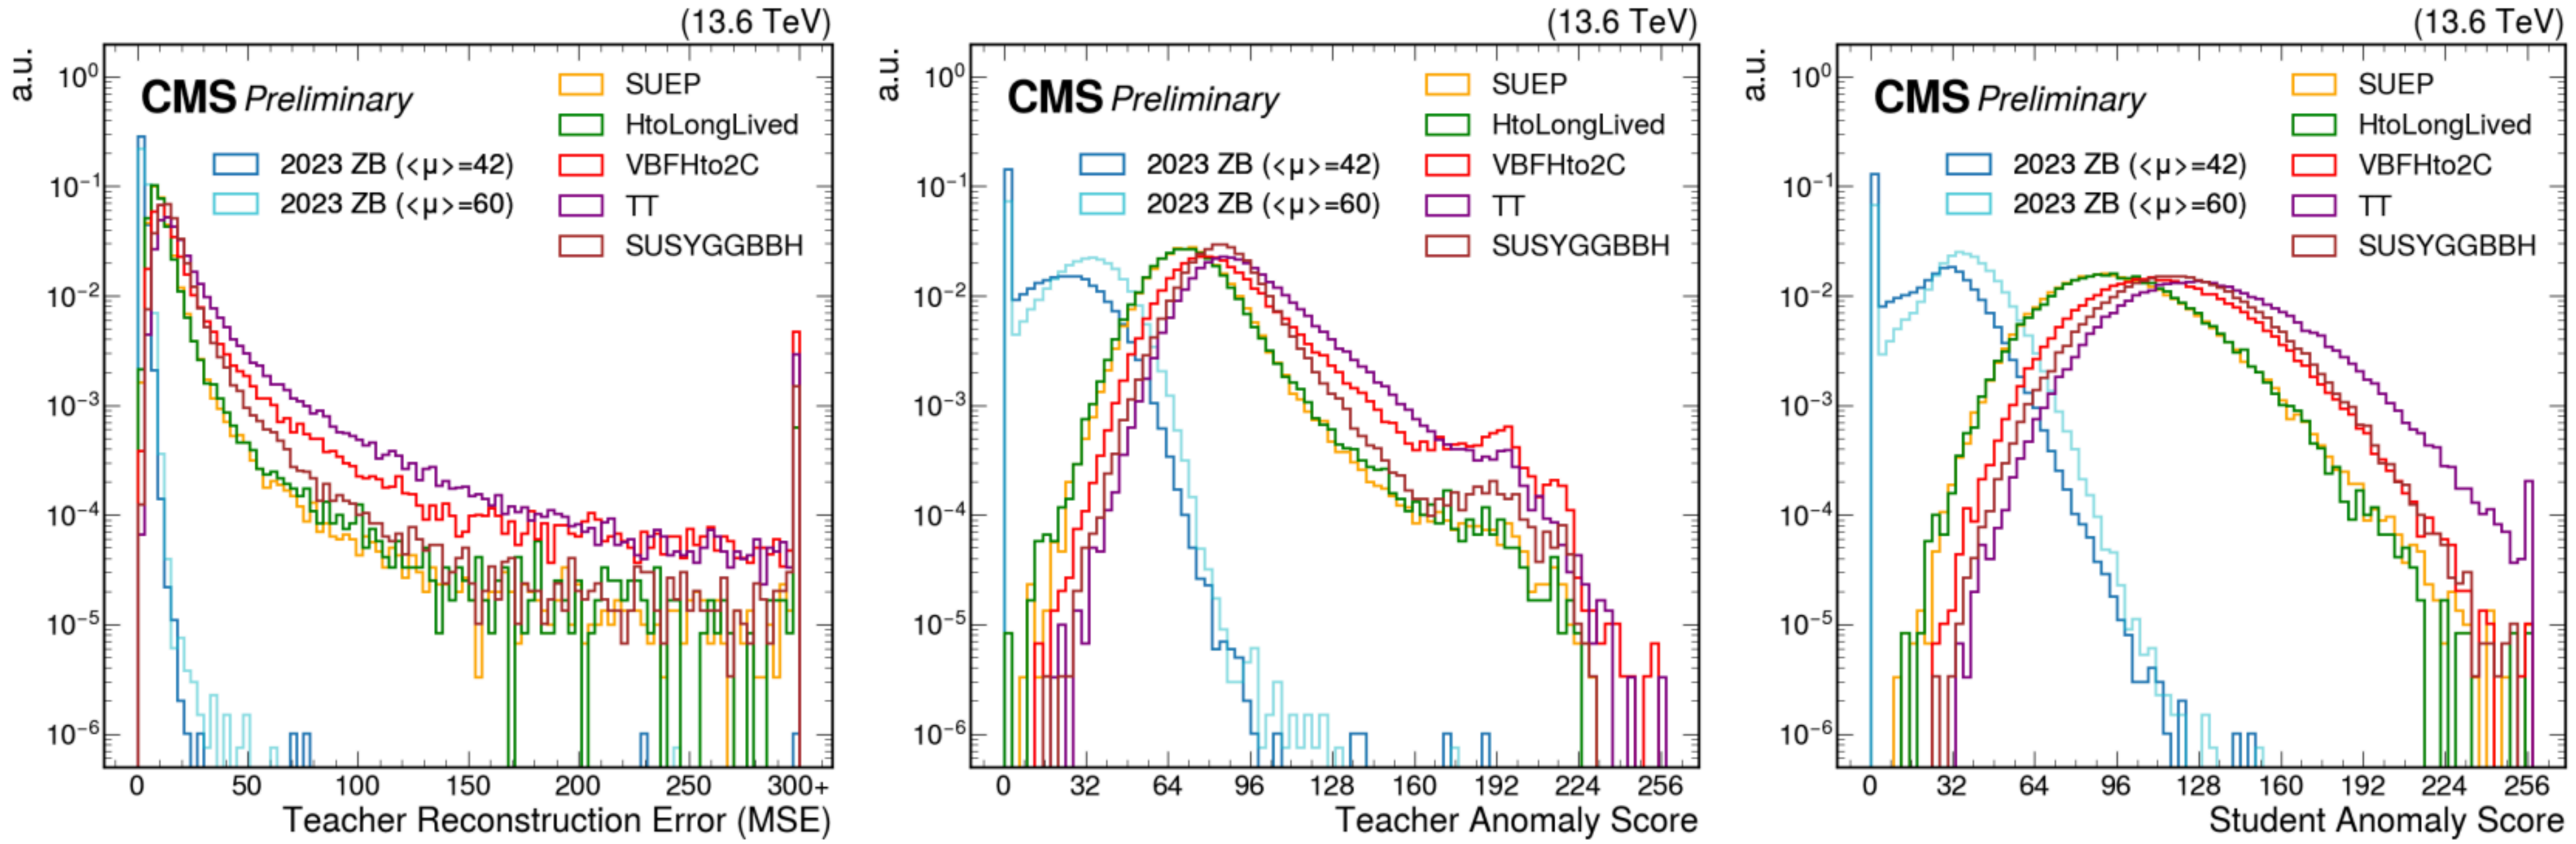

Figure 4: Teacher MSE distribution (left) and resulting anomaly scores (middle) for ZB data and various simulated signal samples as predicted by the teacher. The anomaly score is calculated as $s = \log MSE \cdot 32$ and subsequently quantized as 16 bits where 8 are used for the integer part, resulting 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.

Teacher Performance

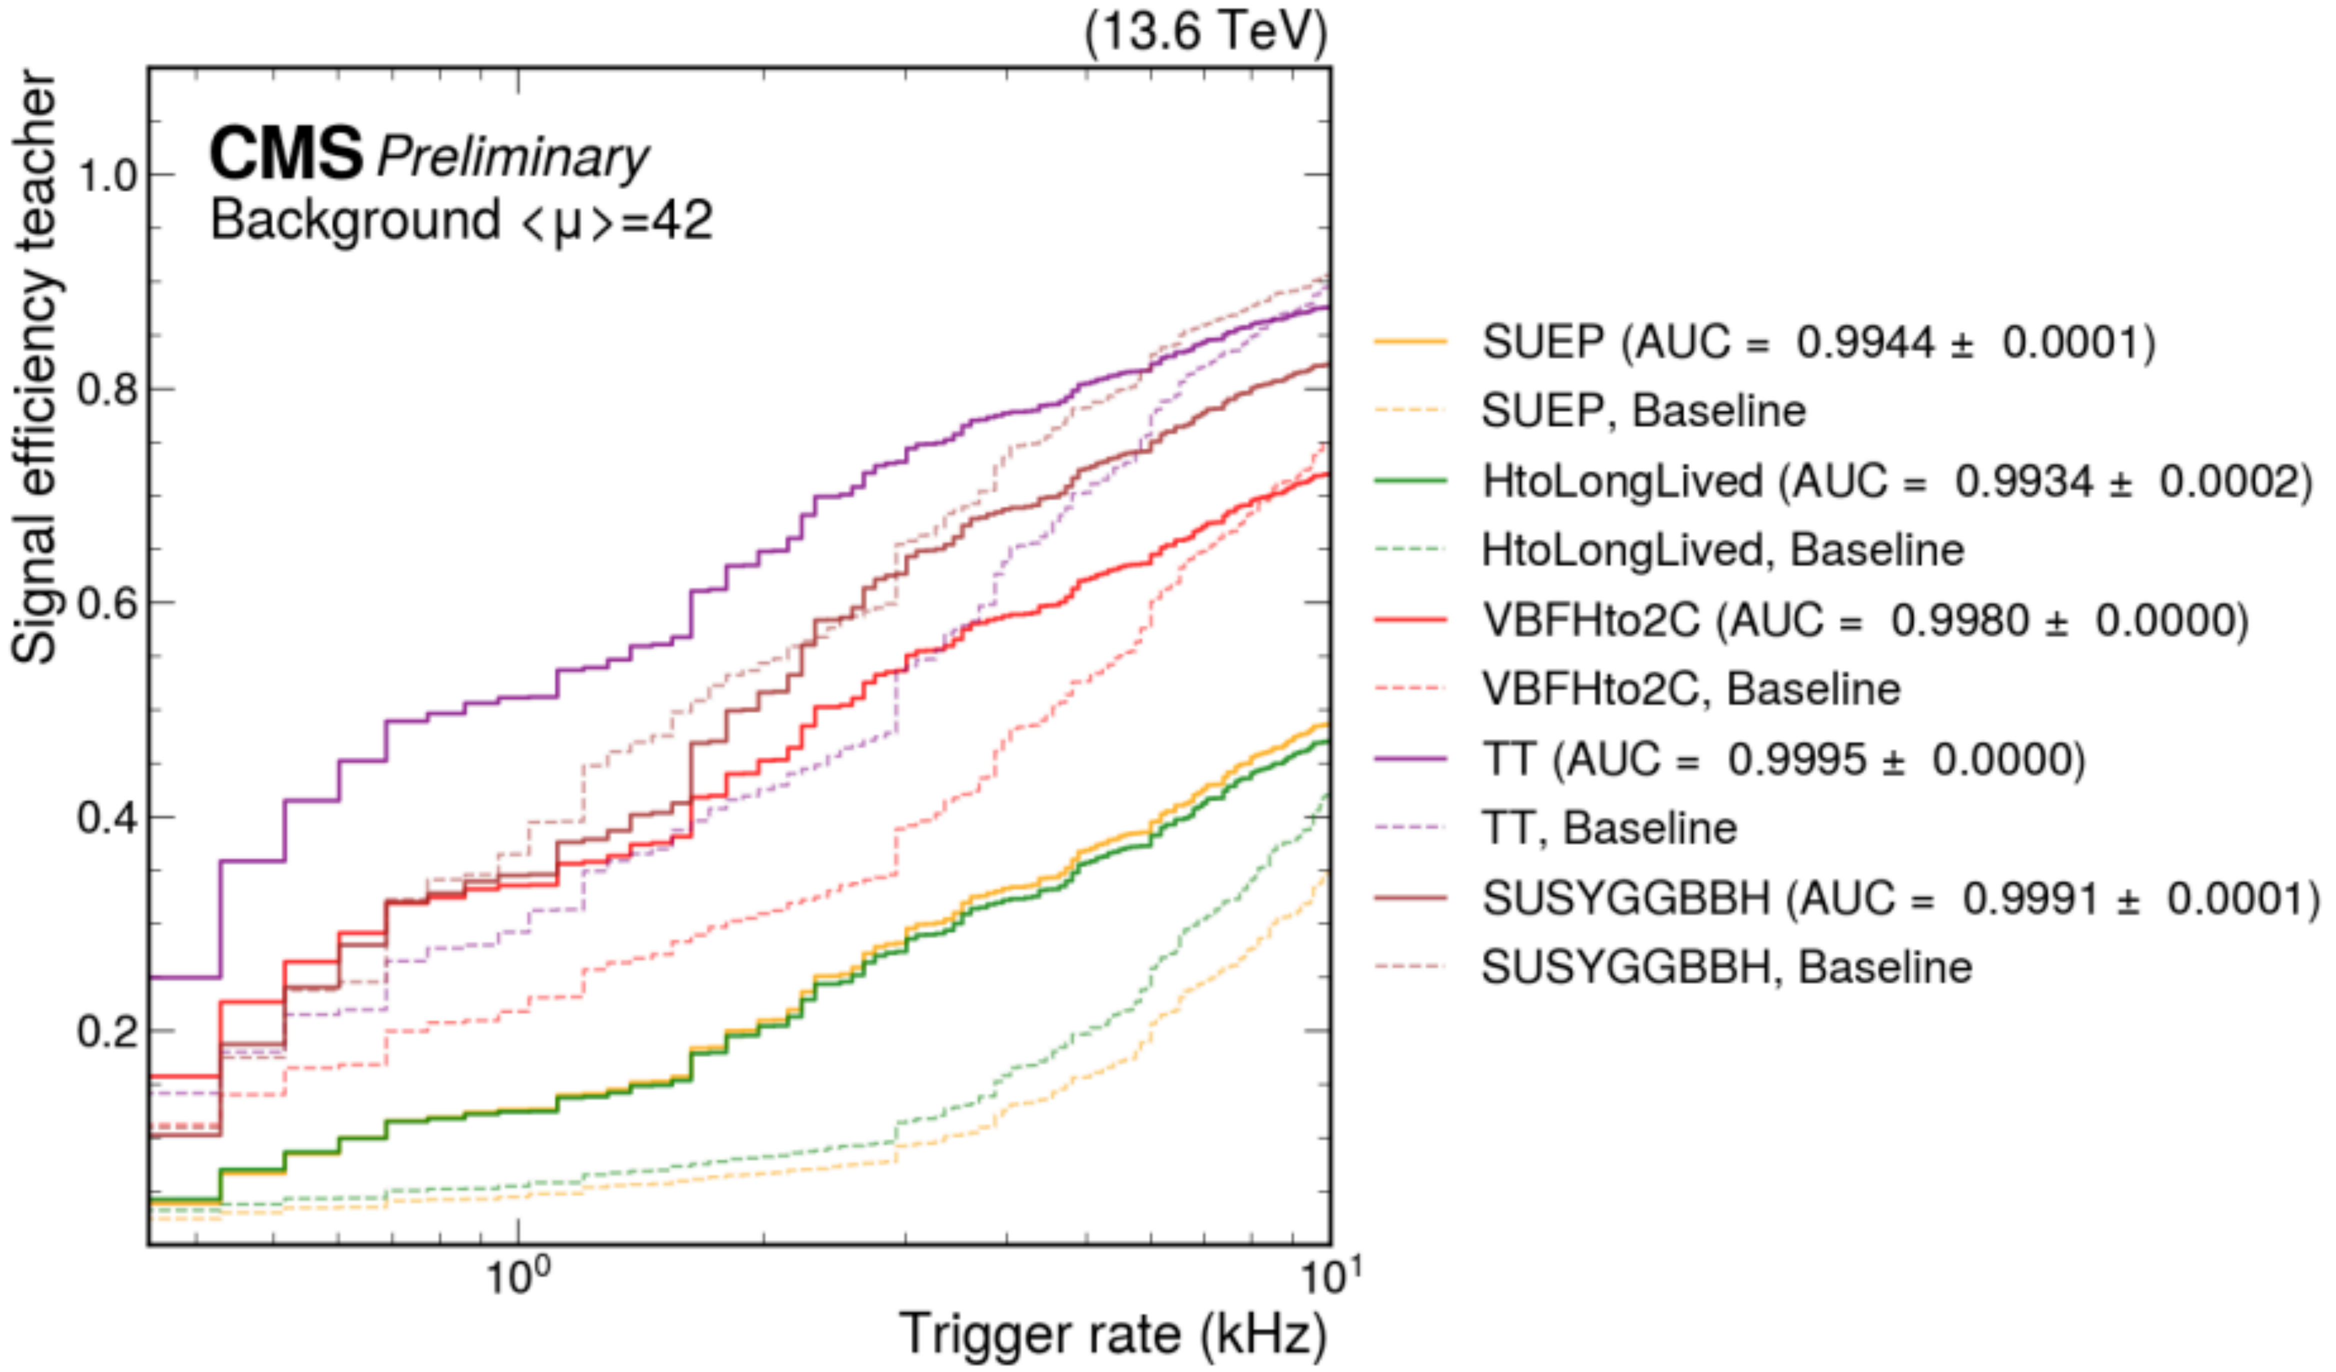

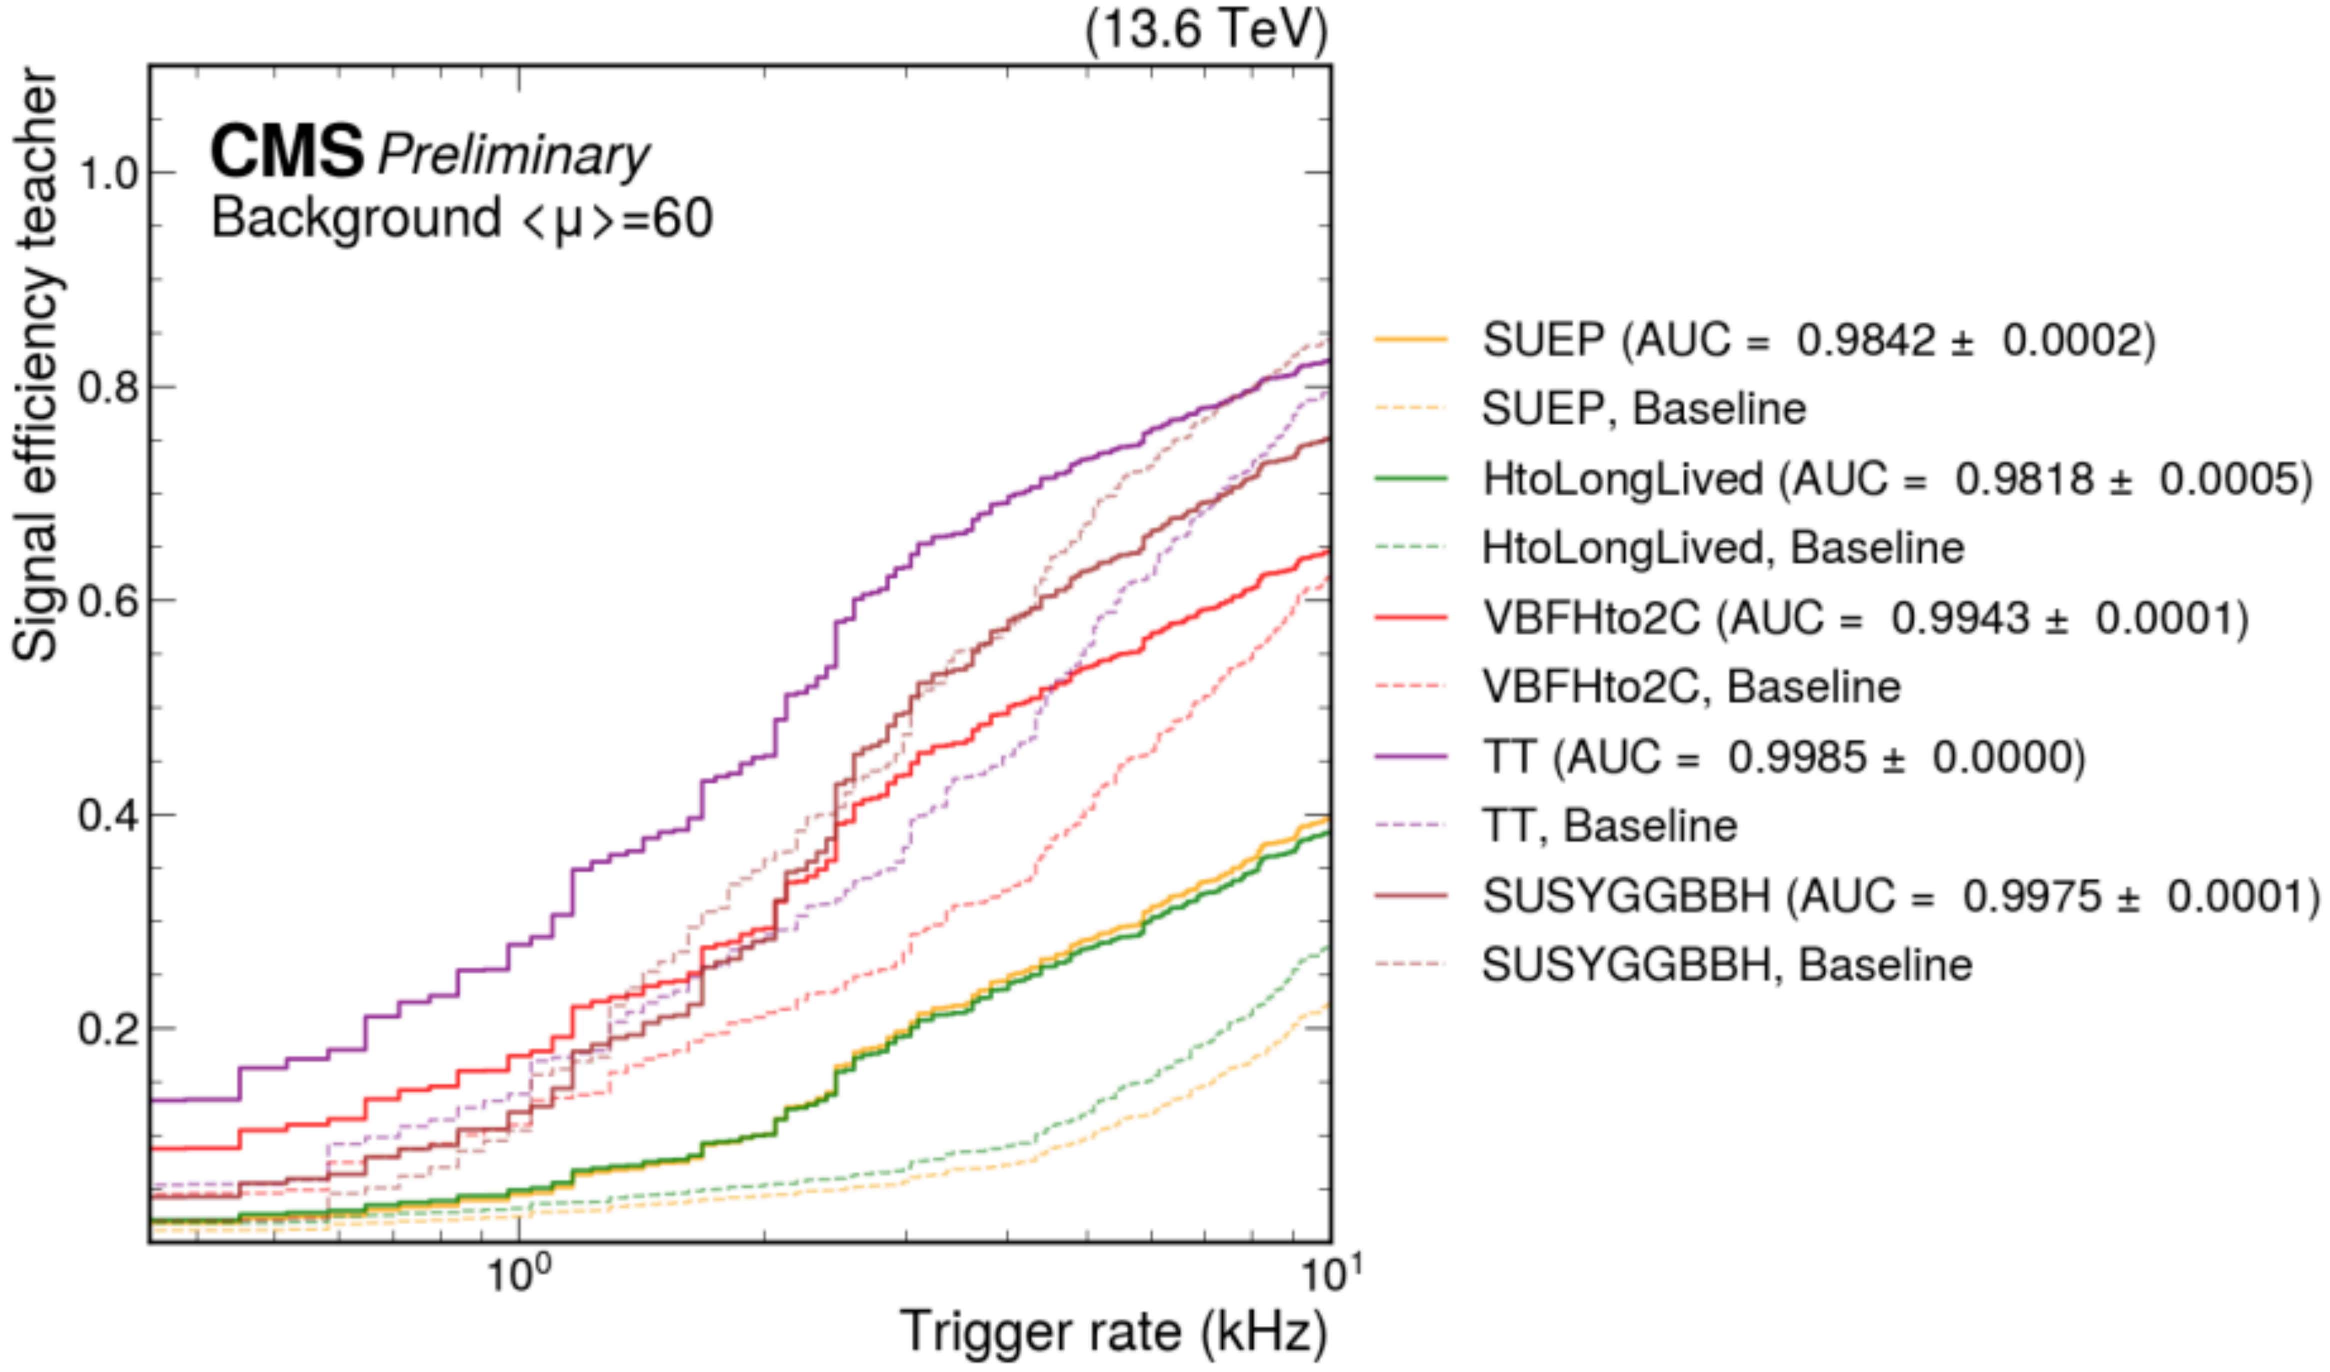

Student Performance

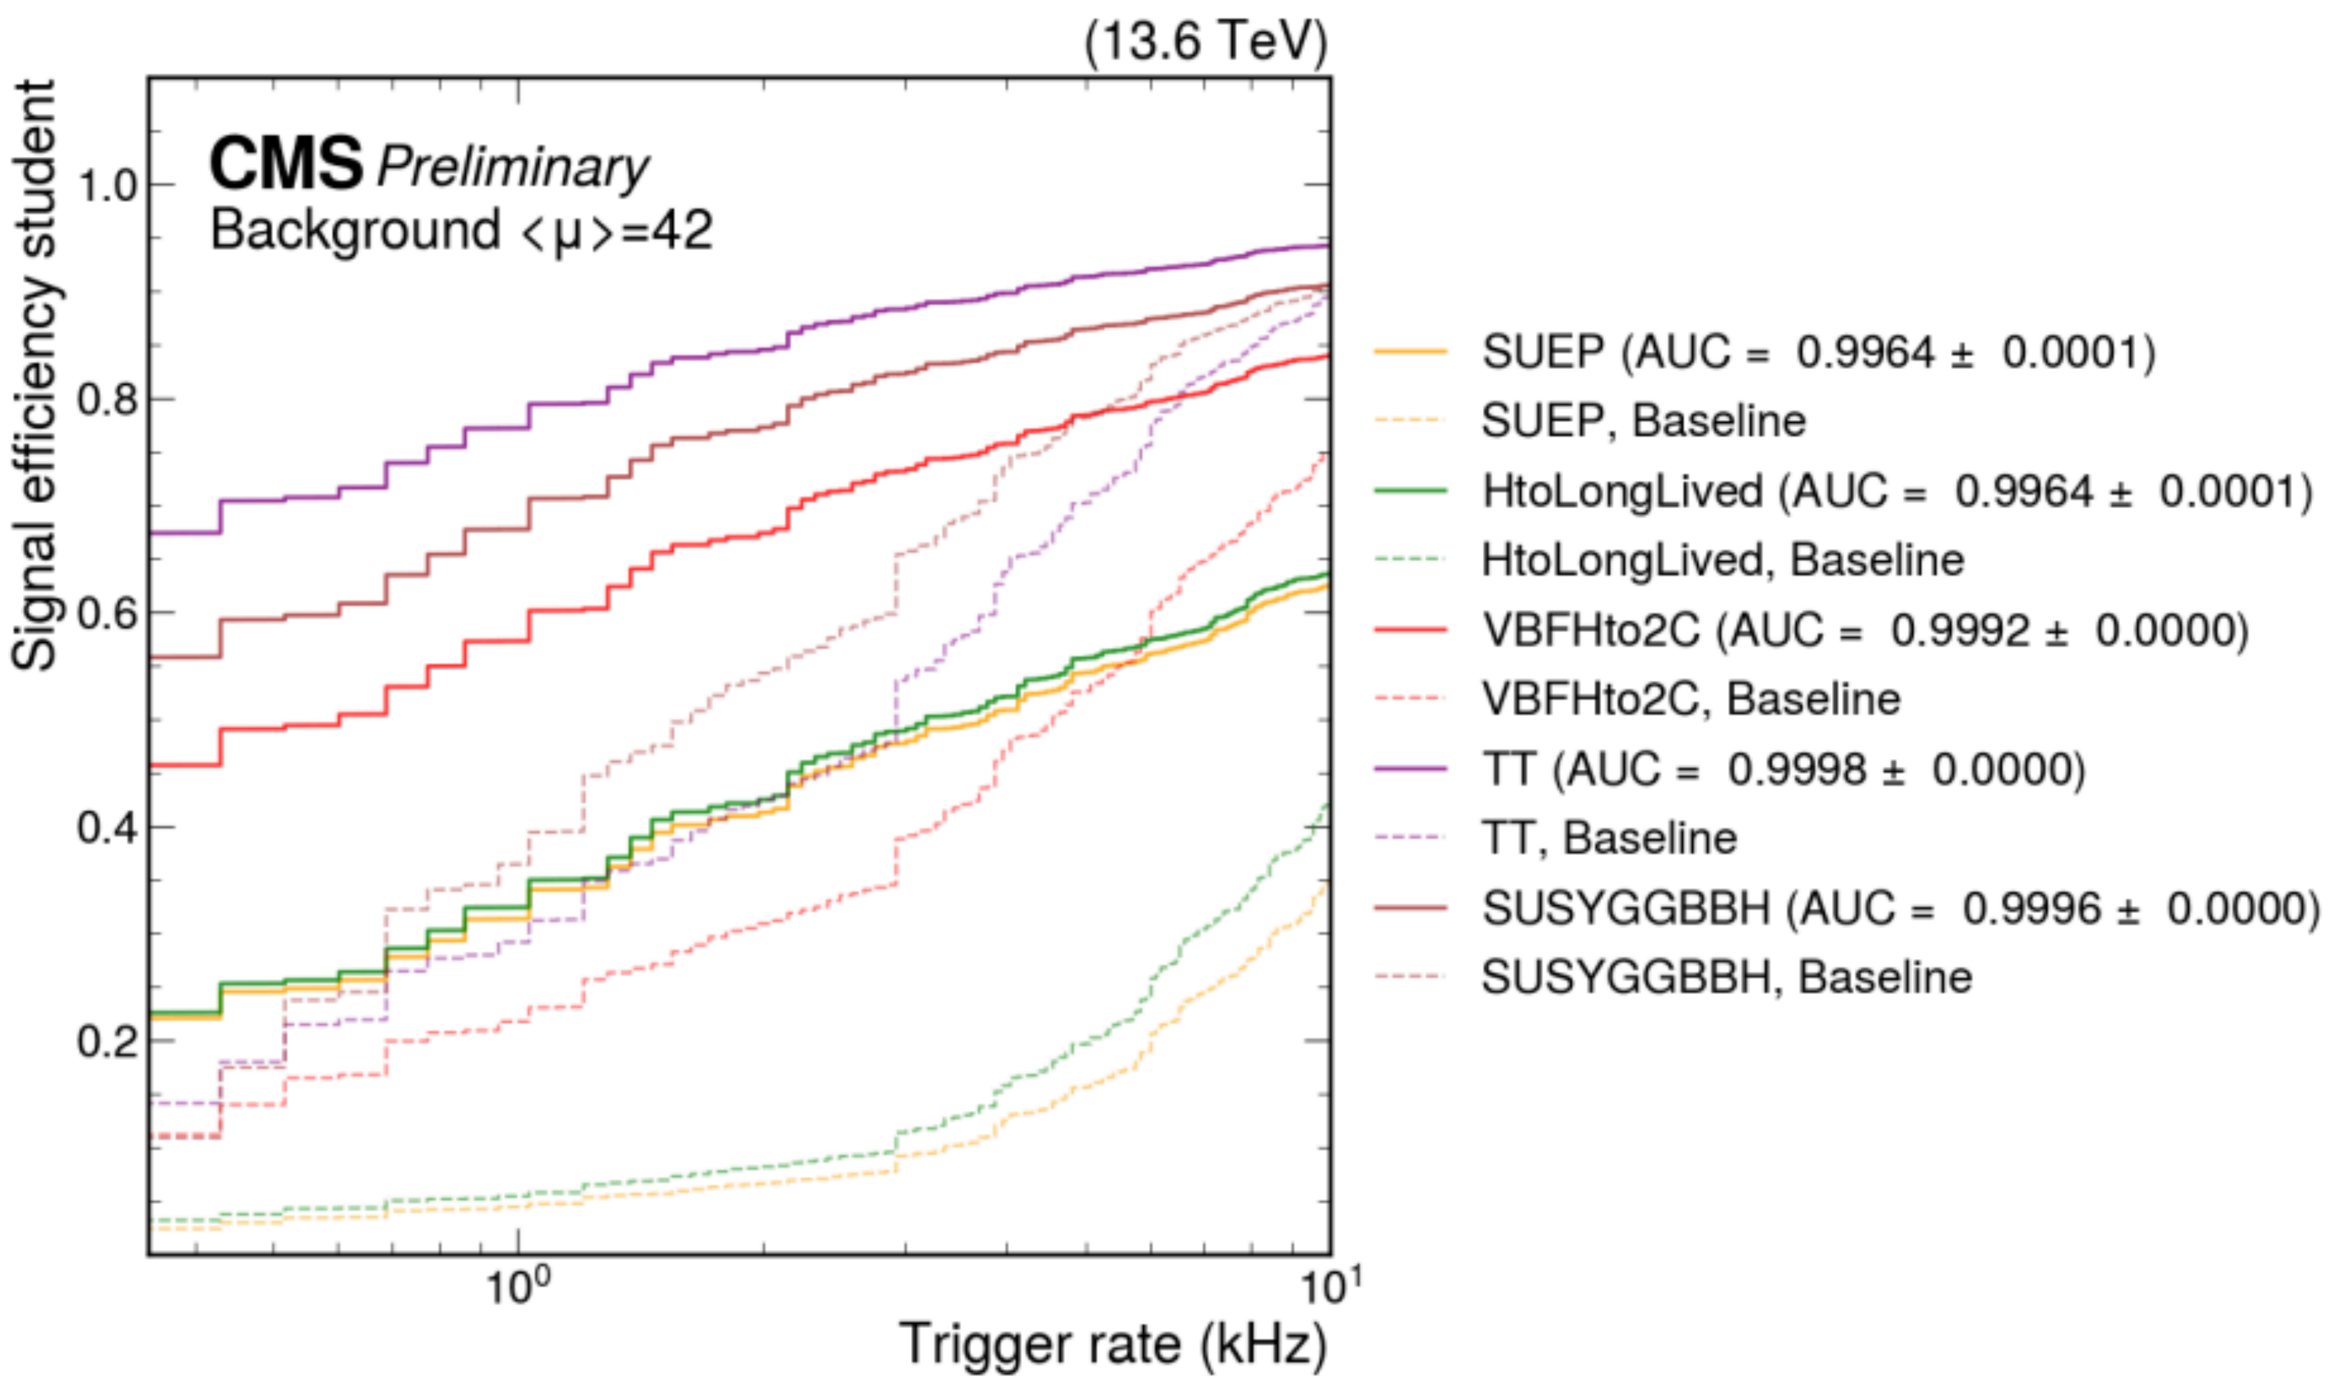

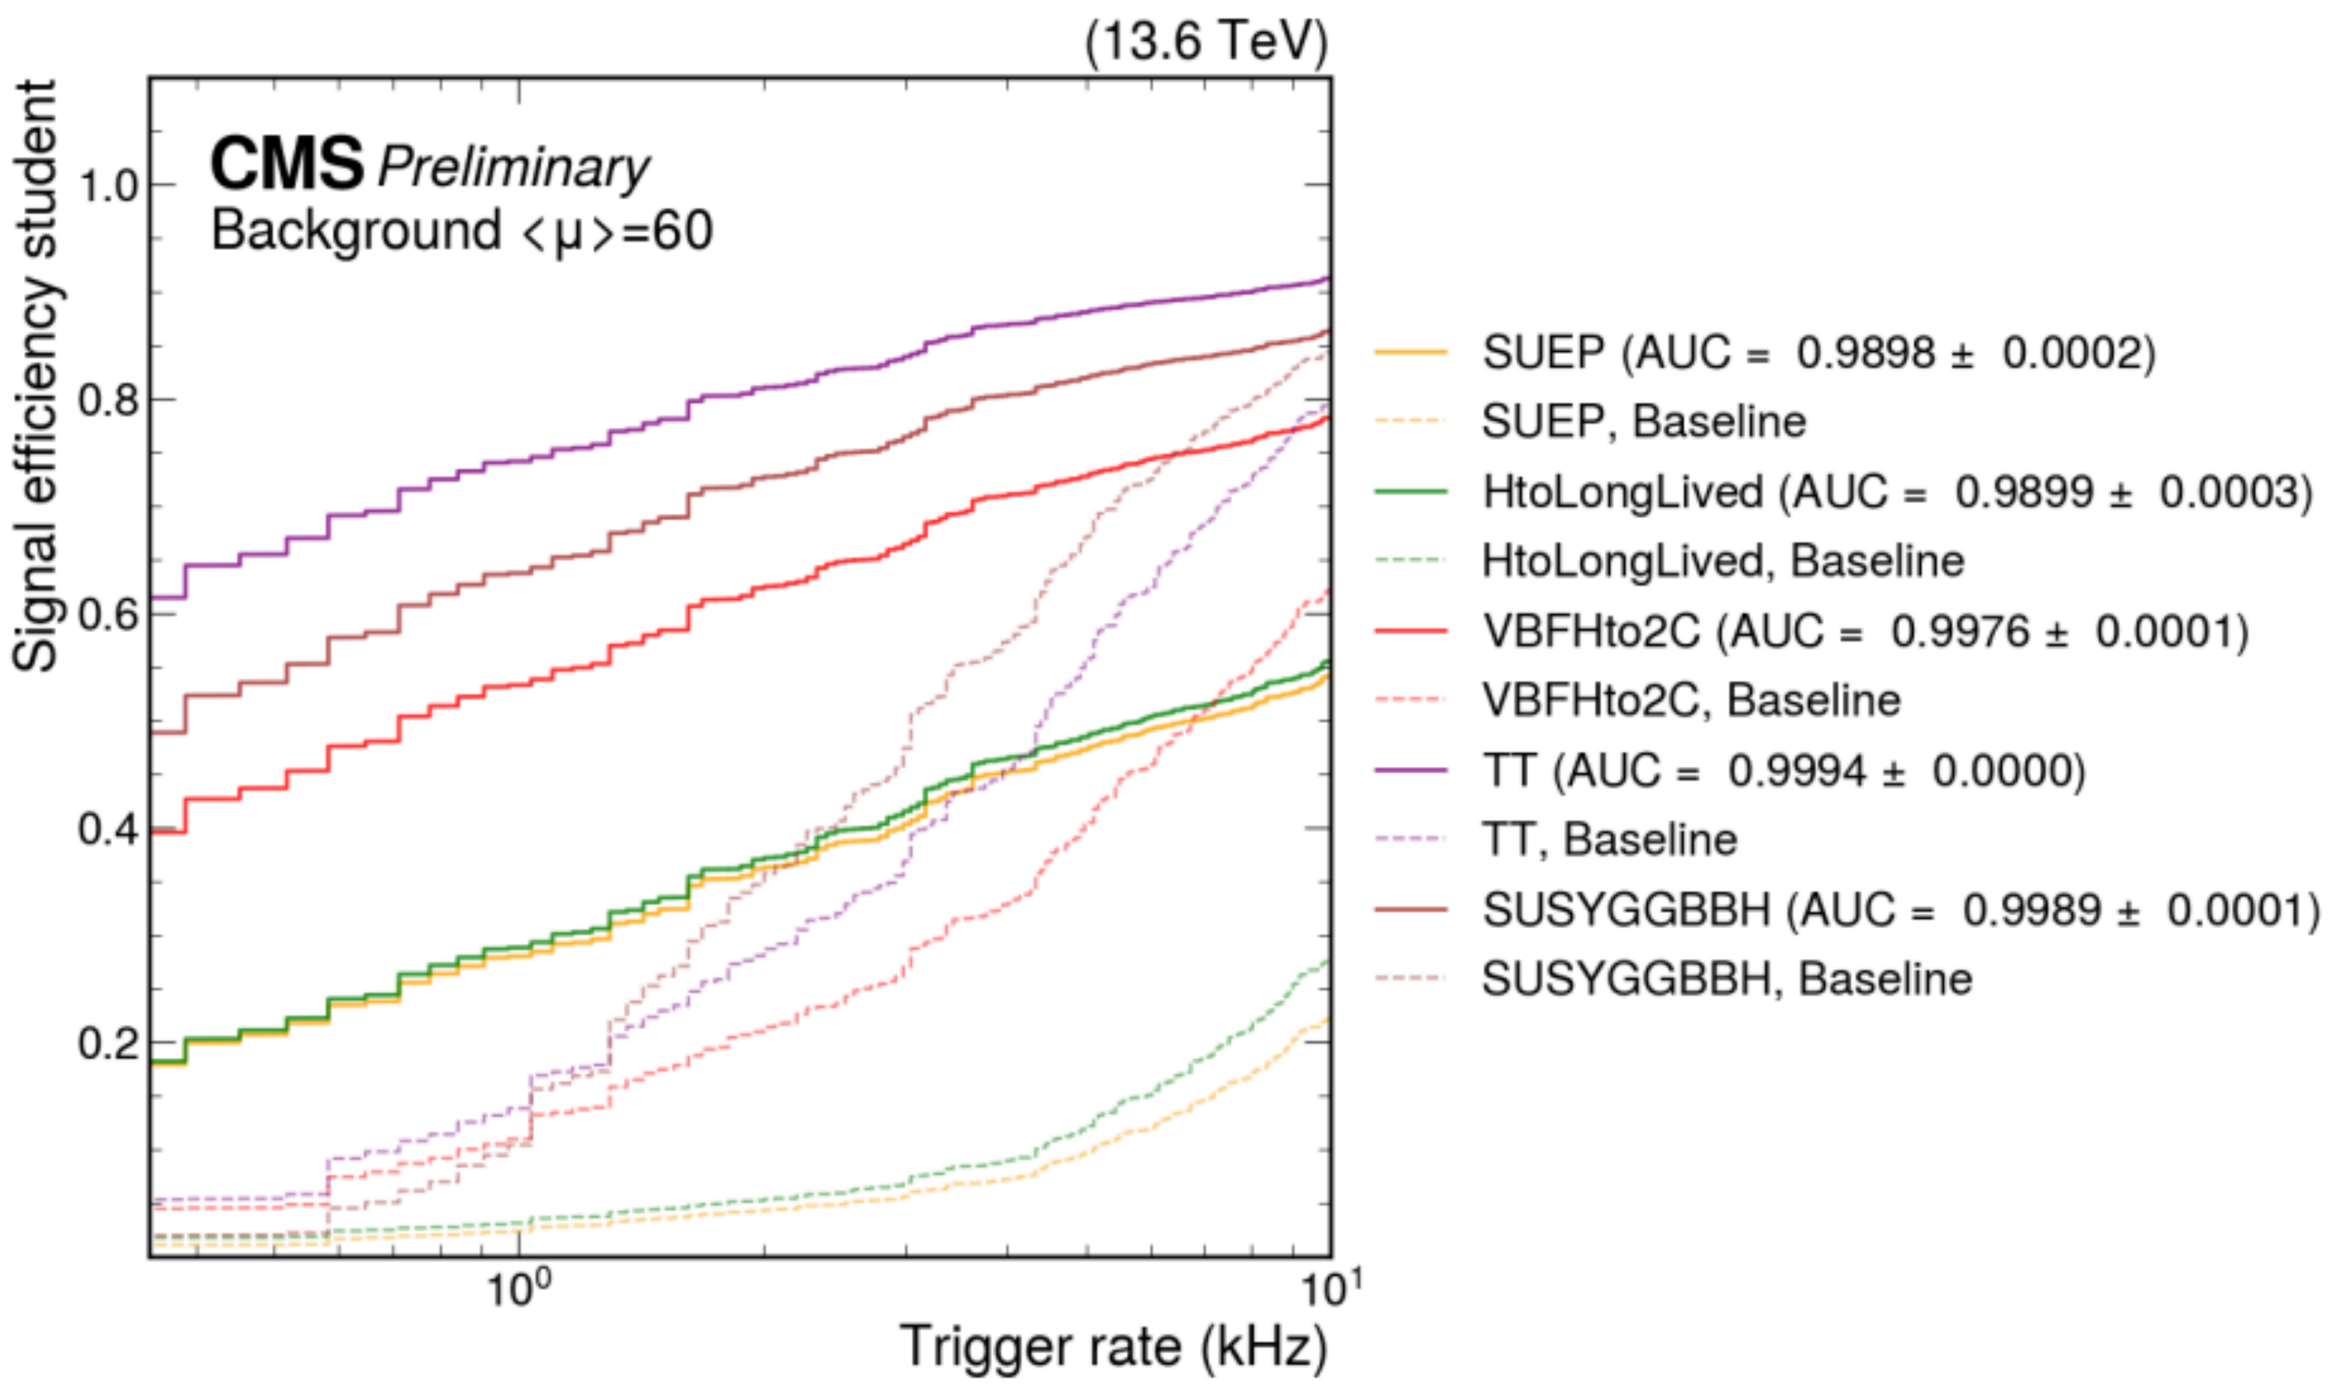

Samples + Examples

Training + Testing Samples

- Zero Bias: 2023BCD
- SUSY: pair-prod gluinos (M=350 GeV)
- H→SS→4b: m(S)-25 GeV, cτ=1500mm
- SUEP: H → dark pions → dark photons → SM pions

ZB Background: MSE = 2.57

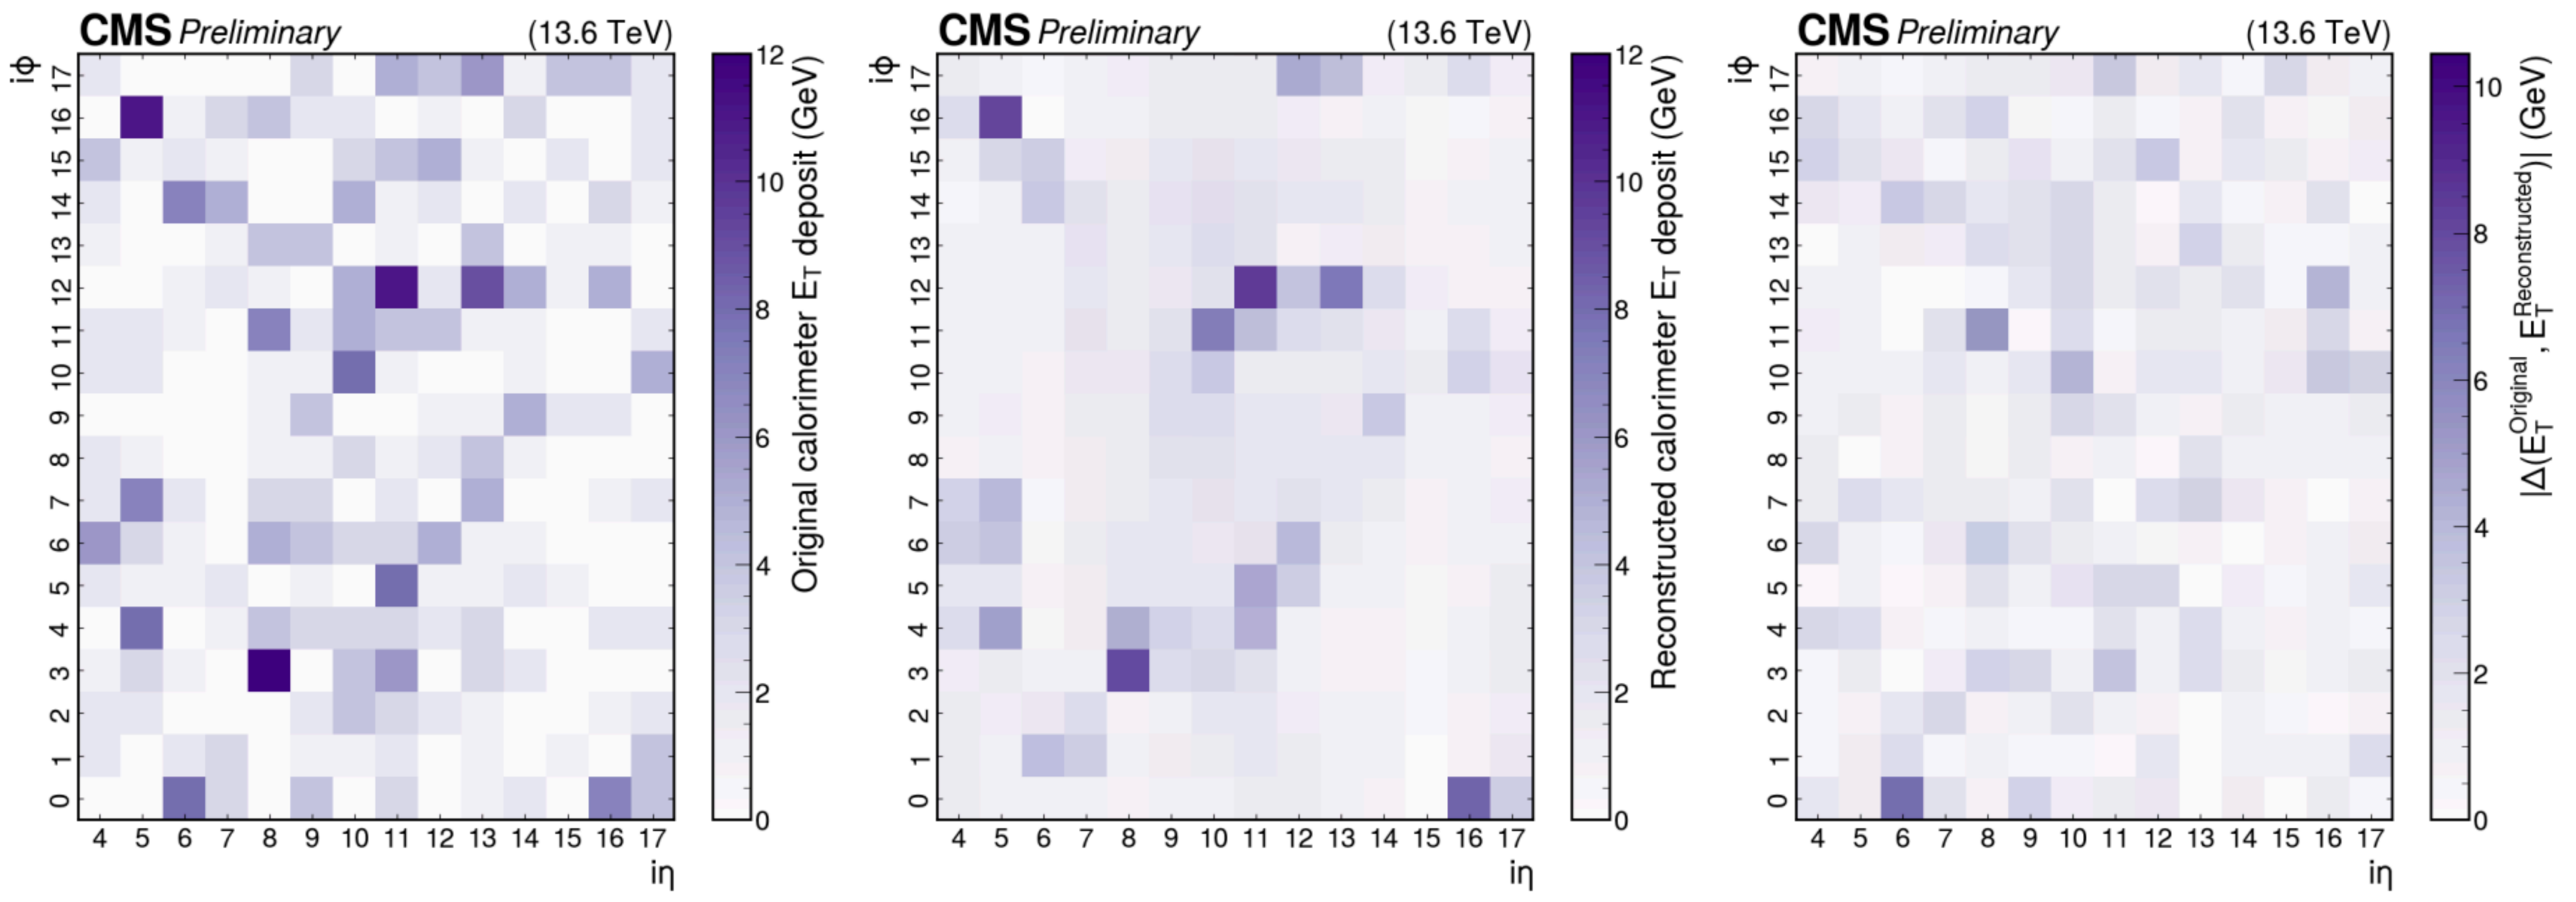

SUSY Signal: MSE = 14.89

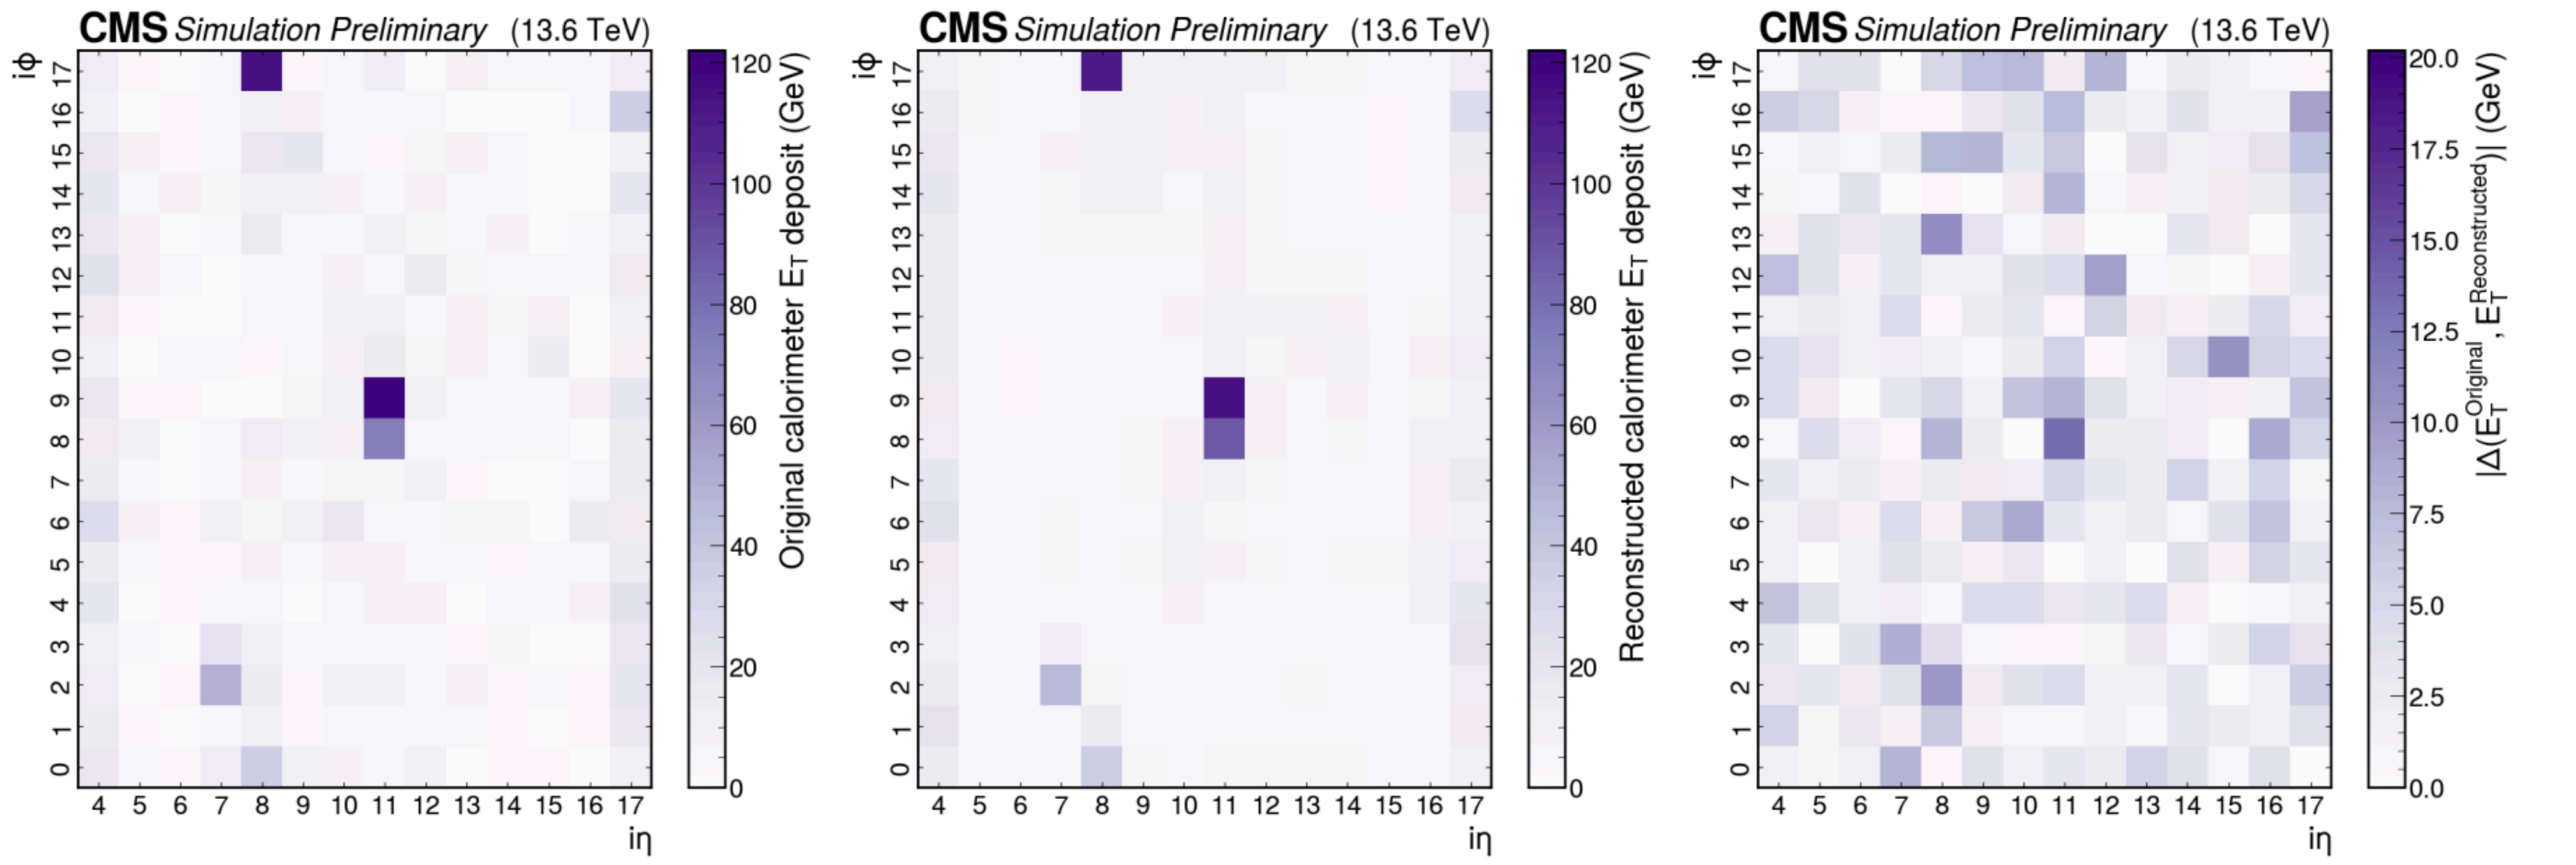

Training History

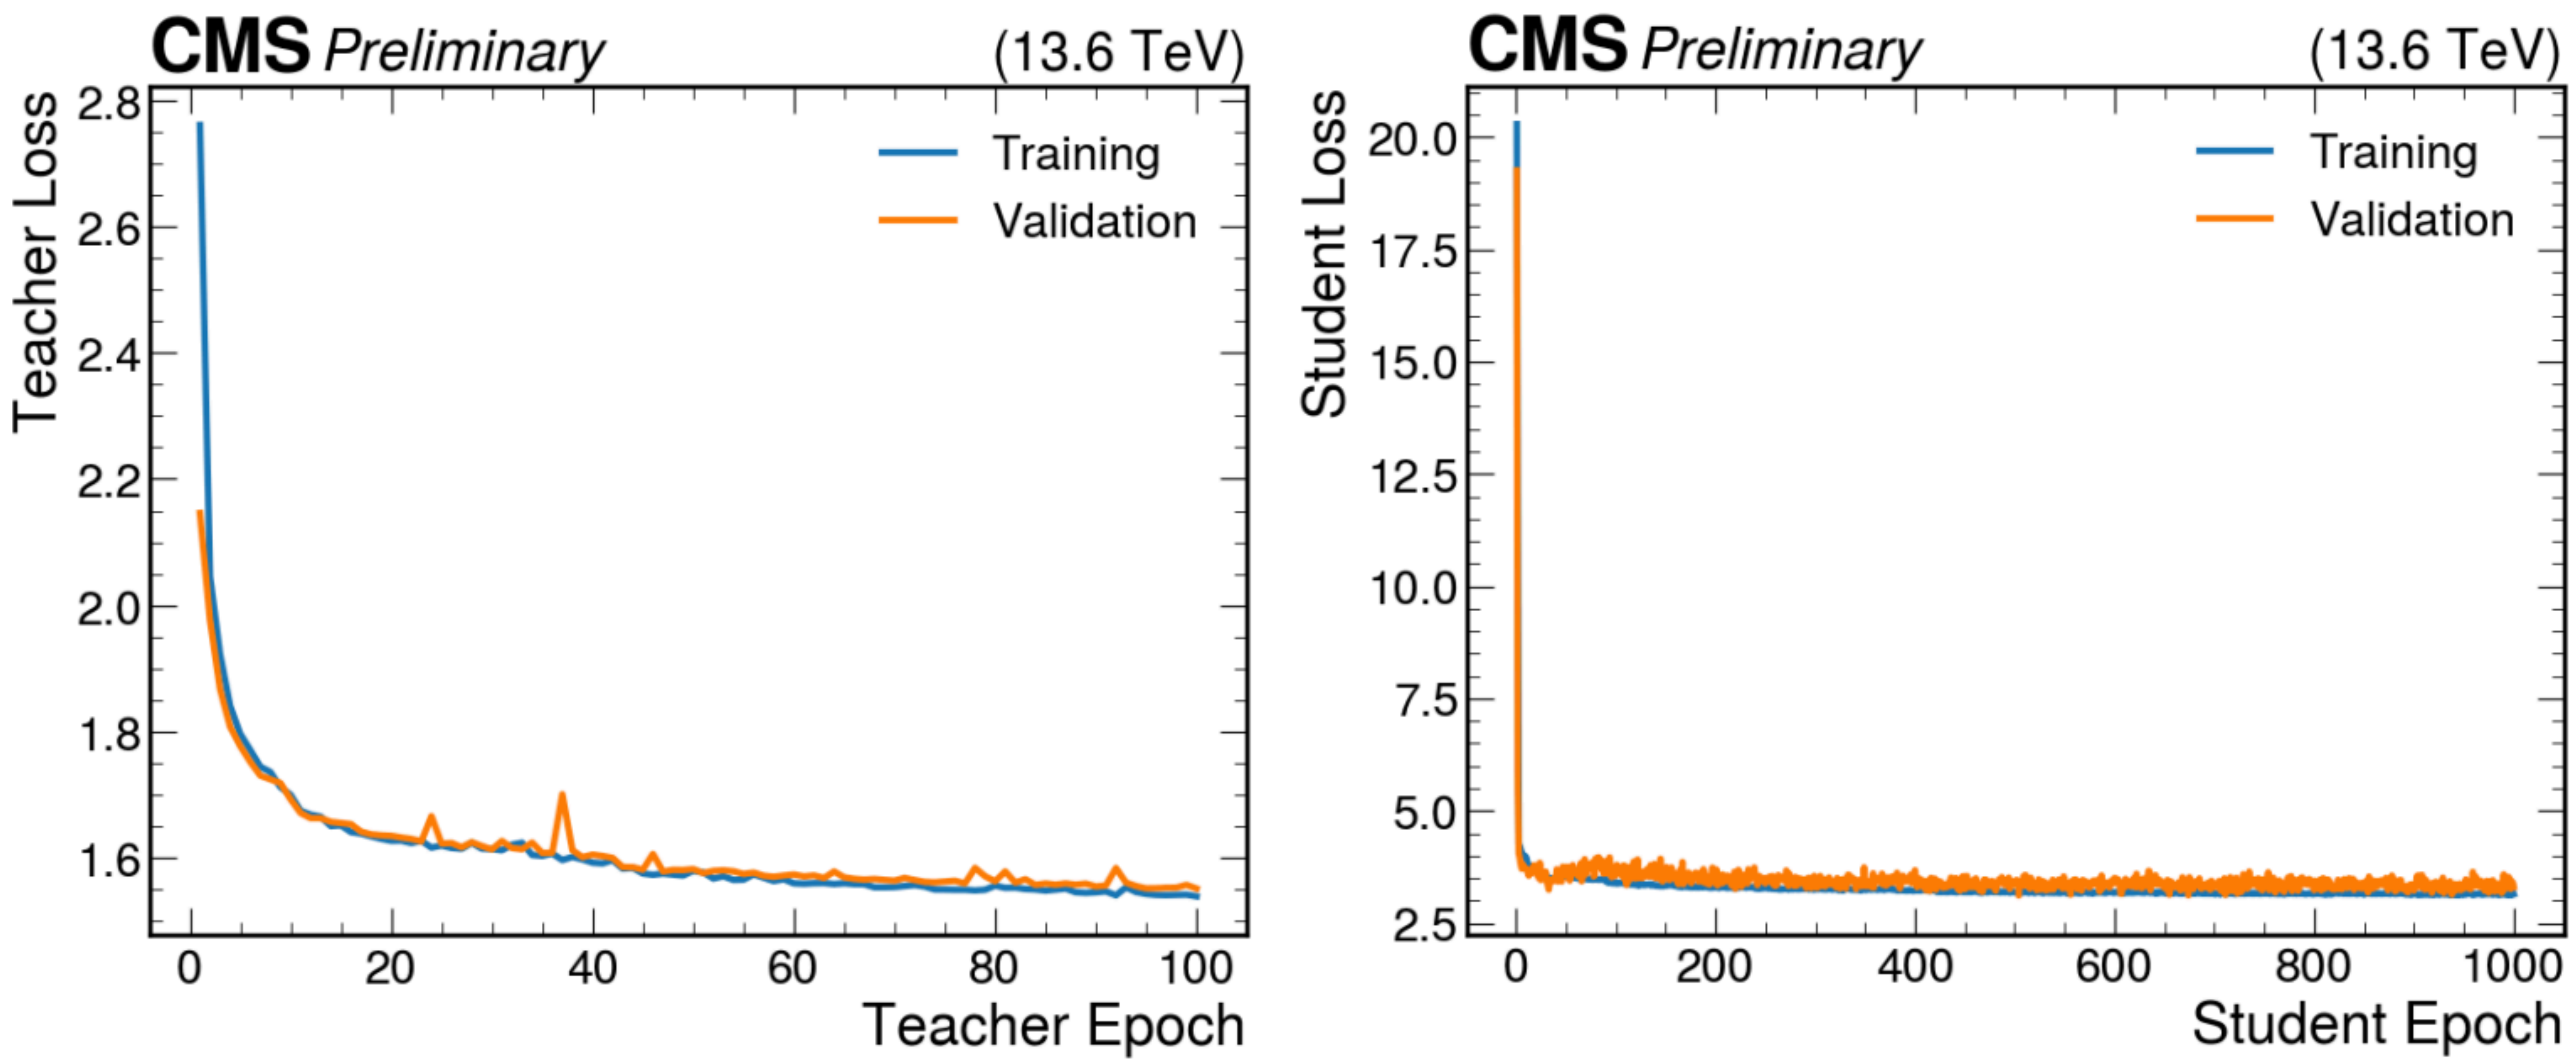

Figure 3: Training histories of teacher (left) and student (right). After each training epoch, the teacher is inferred to create new targets on which the student is trained for ten epochs.

TTBar Analysis

- 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!

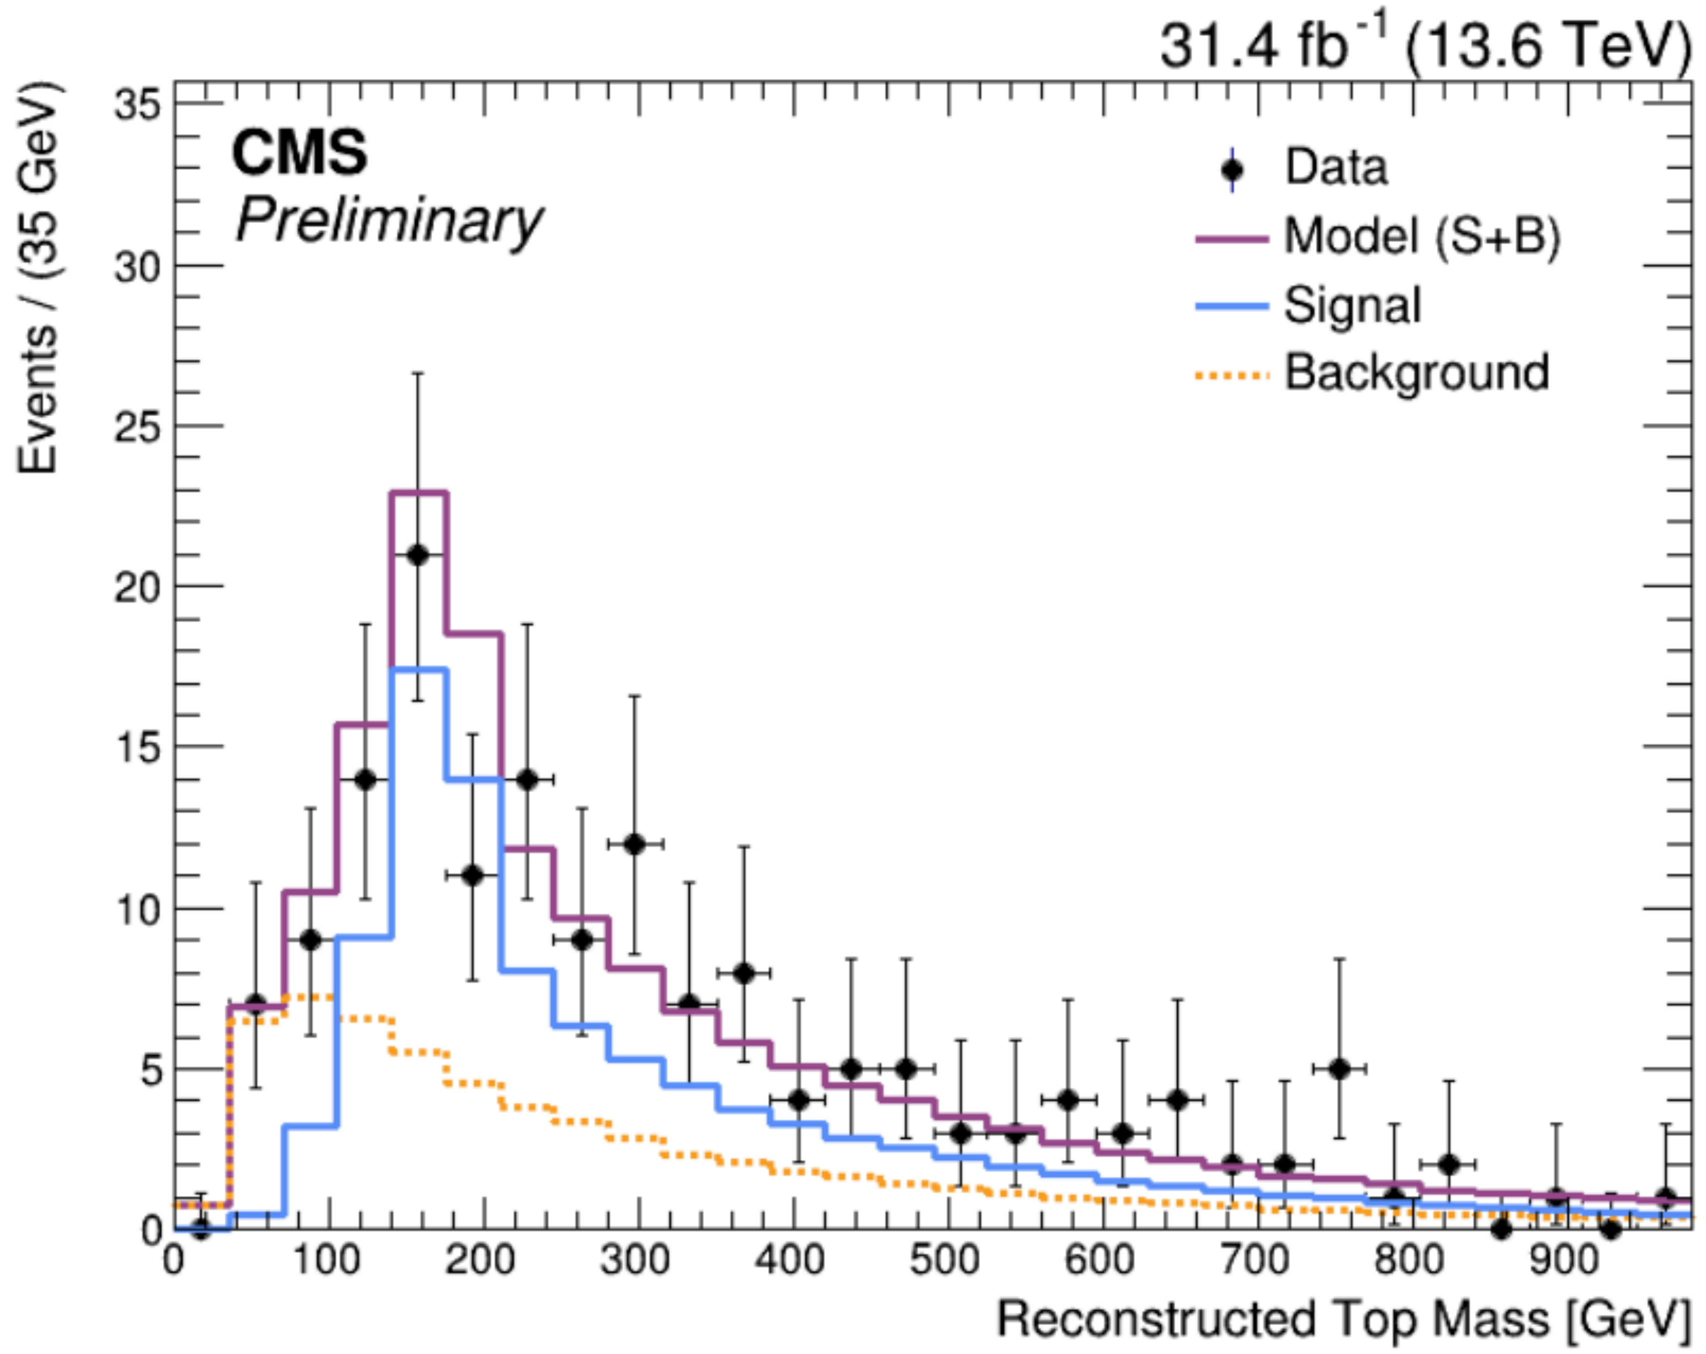

Figure 11: Post-fit distribution of the reconstructed top mass. In this case, a cut of CICADA score > 115 was used to define the 'Data' and 'Signal' shapes.