Anomaly Detection in the CMS L1 Trigger

Melissa Quinnan

A3D3 HEP Monthly Meeting, February 19, 2024

CMS L1 Trigger

- L1 trigger rejects 99.75% of LHC events
  - Constrained by low latency of 50 ns and low resource utilization on FPGAs
- What if we are missing new physics because we did not design the right trigger?

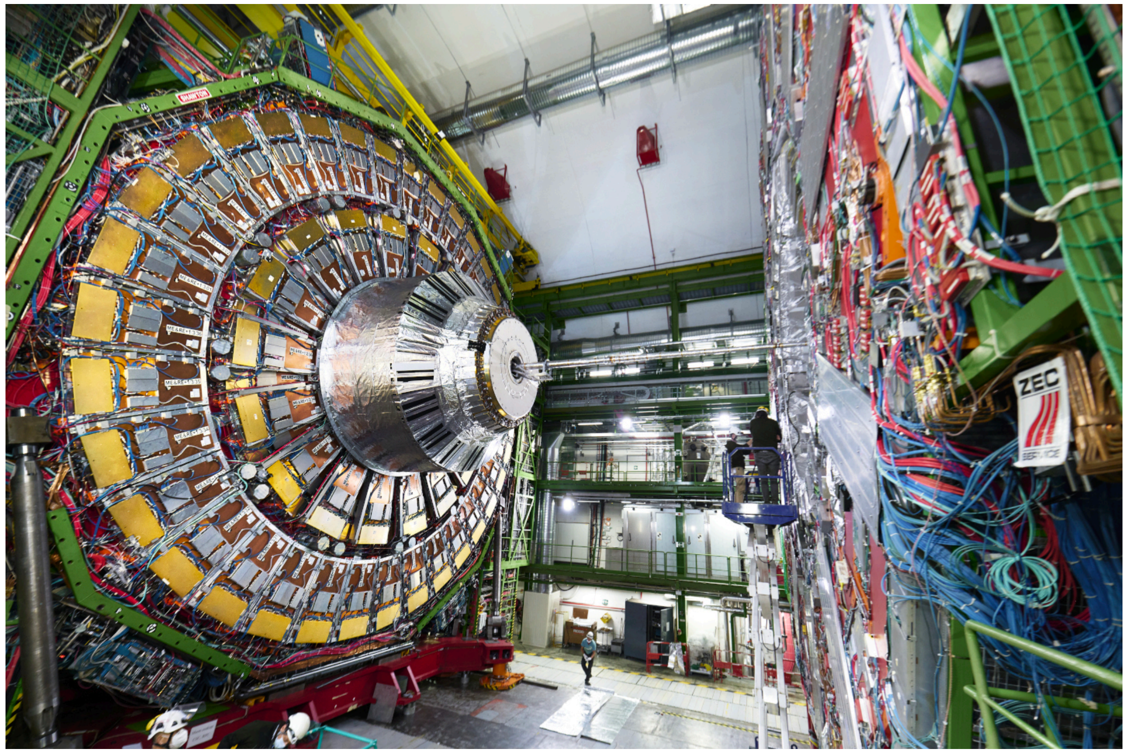

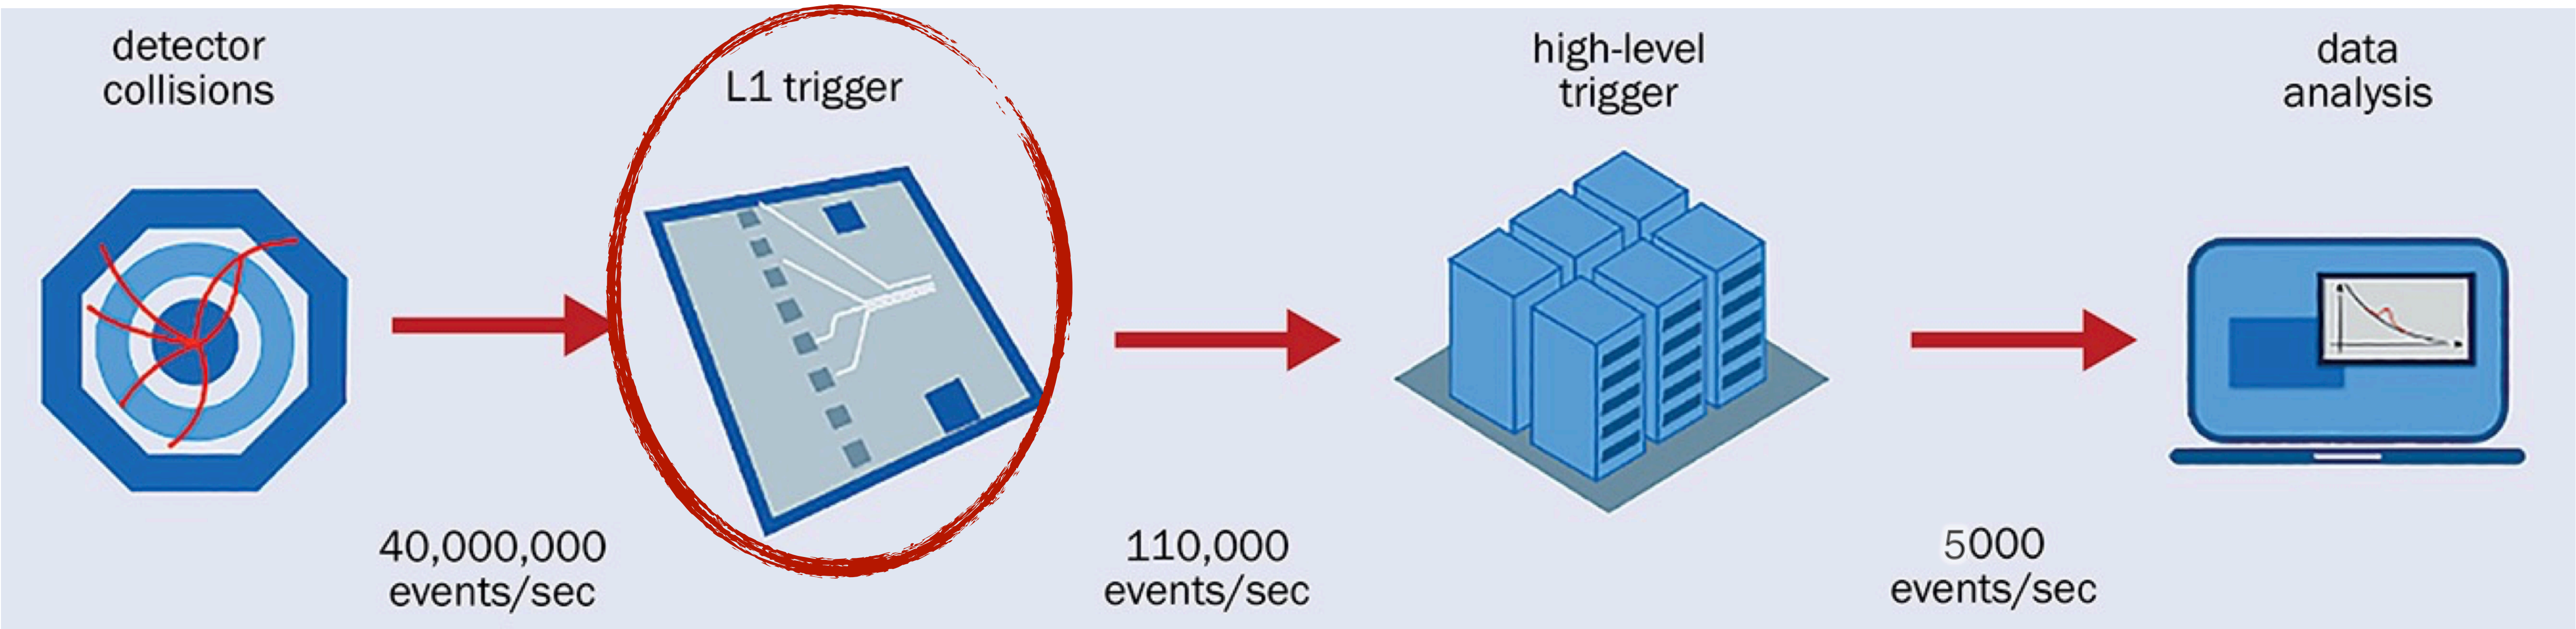

Anomaly Detection in the Trigger

- Signal agnostic - Applicable to signatures for which we have not had the foresight or person-power to design specific triggers
- High sensitivity – Can improve signal efficiency to signatures that are limited by L1 trigger bandwidth

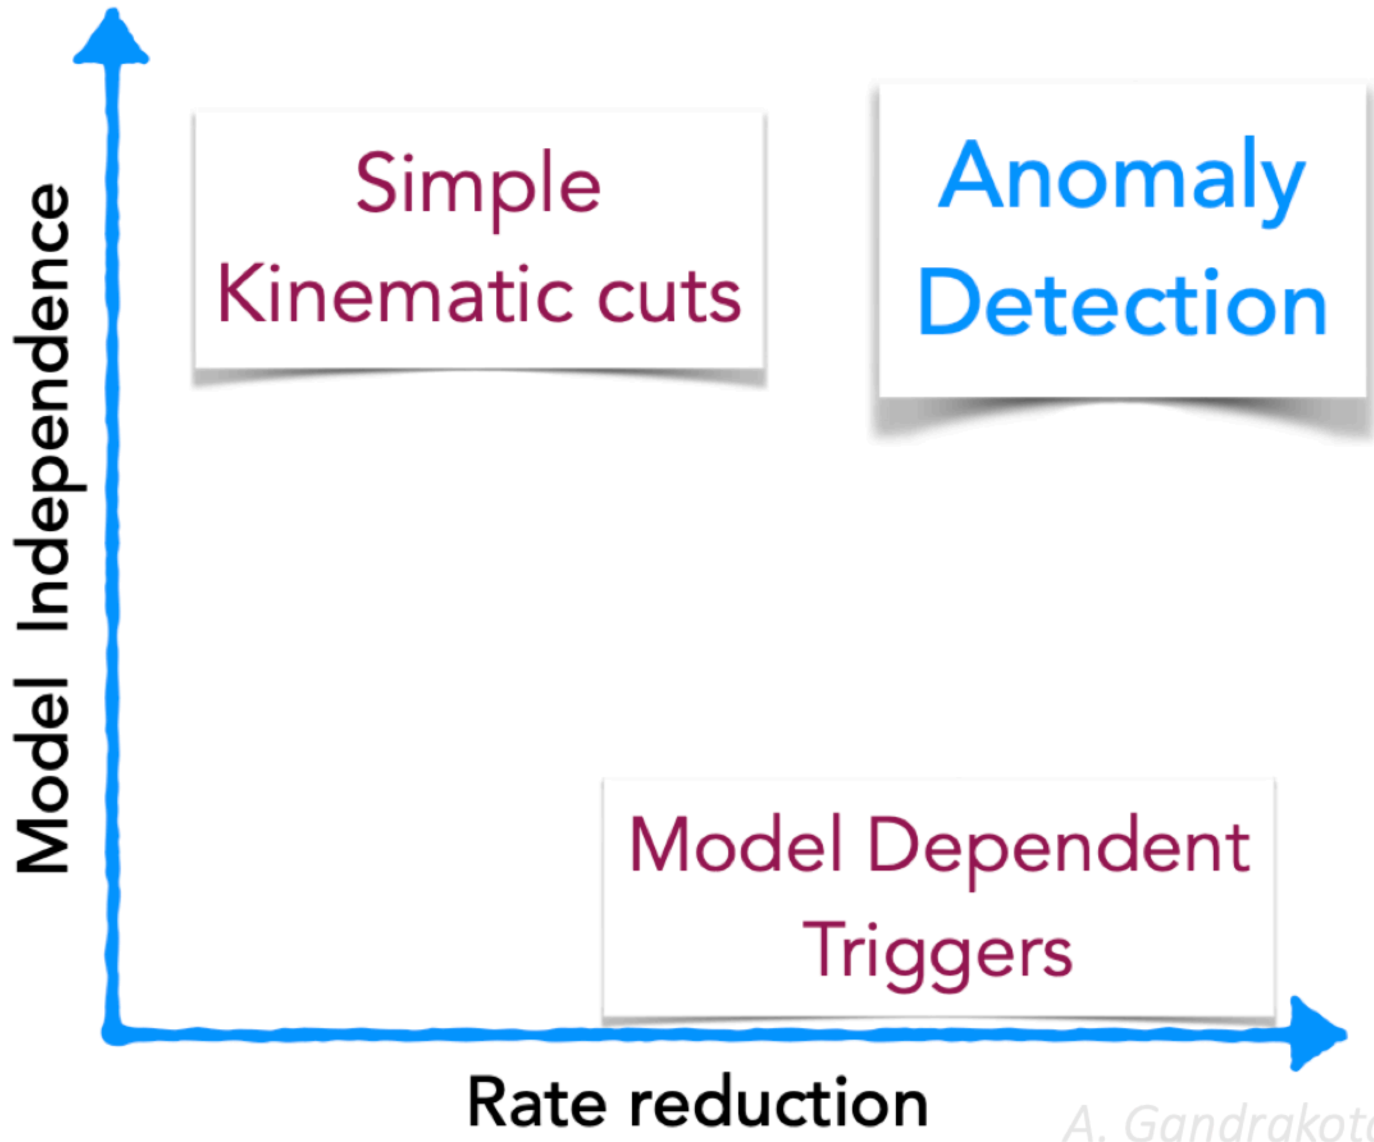

Autoencoders for Anomaly Detection

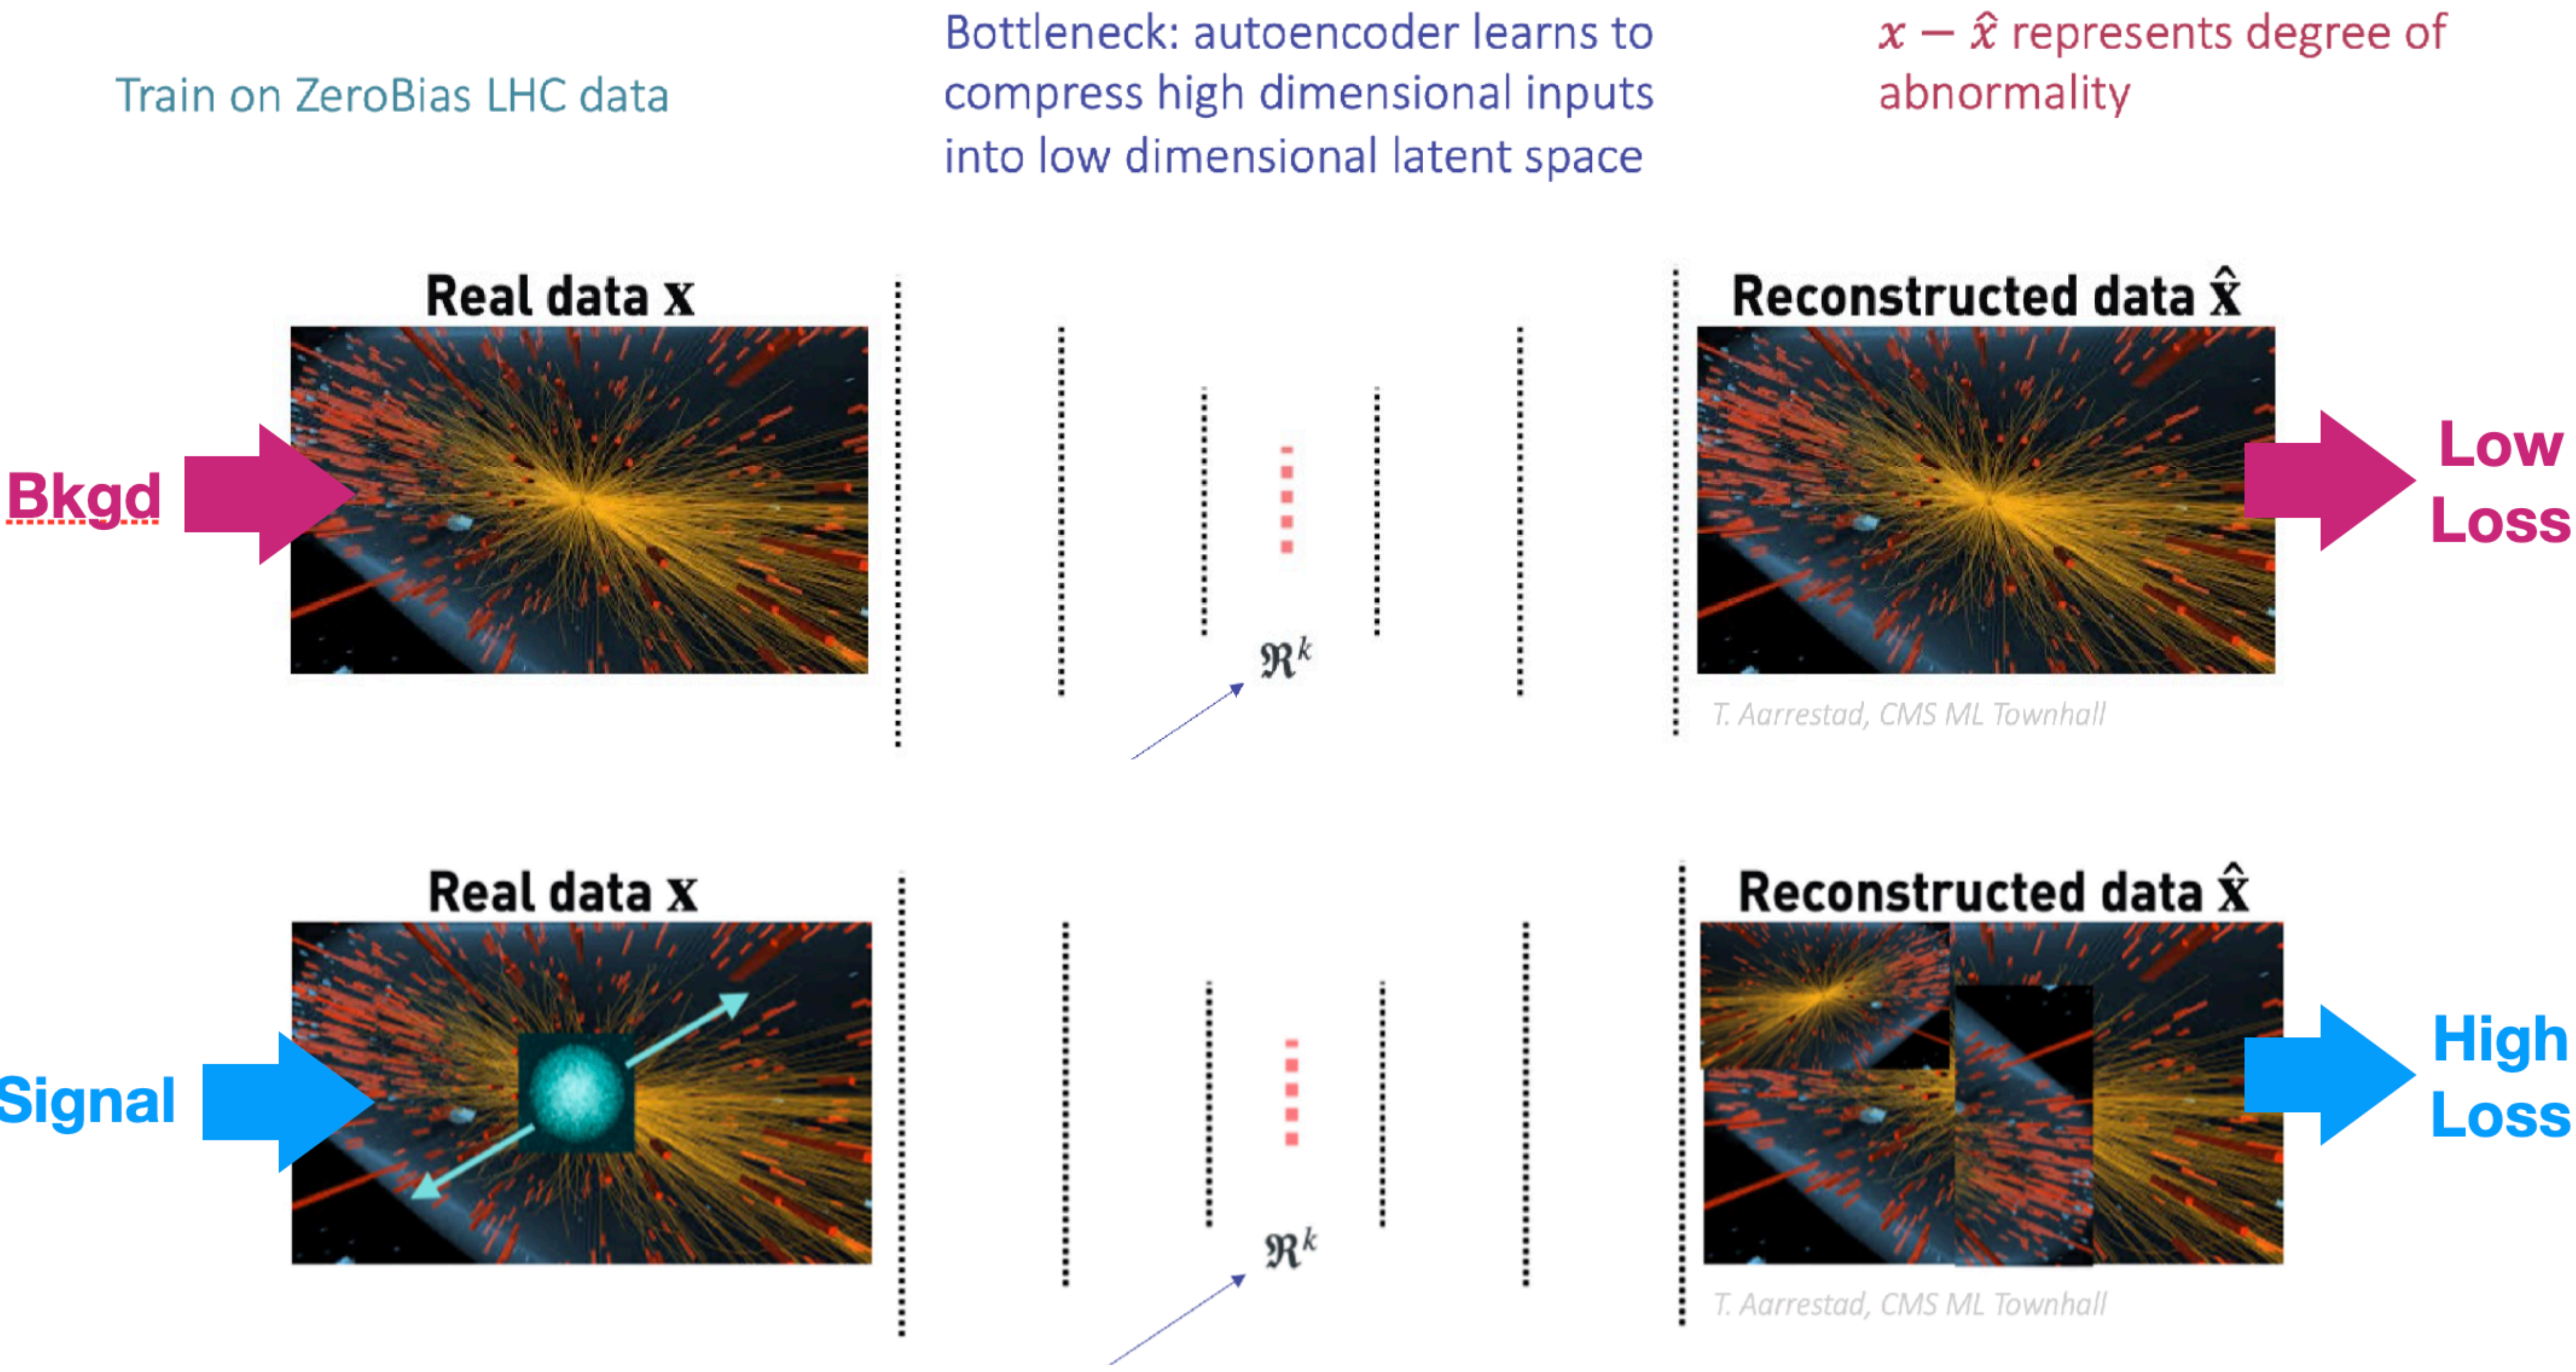

- Unsupervised ML, trained on data

Anomaly Detection at the L1 Trigger

- 2 autoencoder-based anomaly detection algorithms planned for current (Phase 1/Run 3) L1 trigger: AXOL1TL and CICADA

1) CICADA: Calorimeter Image Convolutional Anomaly Detection Algorithm for L1 CaloLayer1

2) AXOL1TL: "Anomaly eXtraction Online L1 Trigger Lightweight" for the L1 Global Trigger

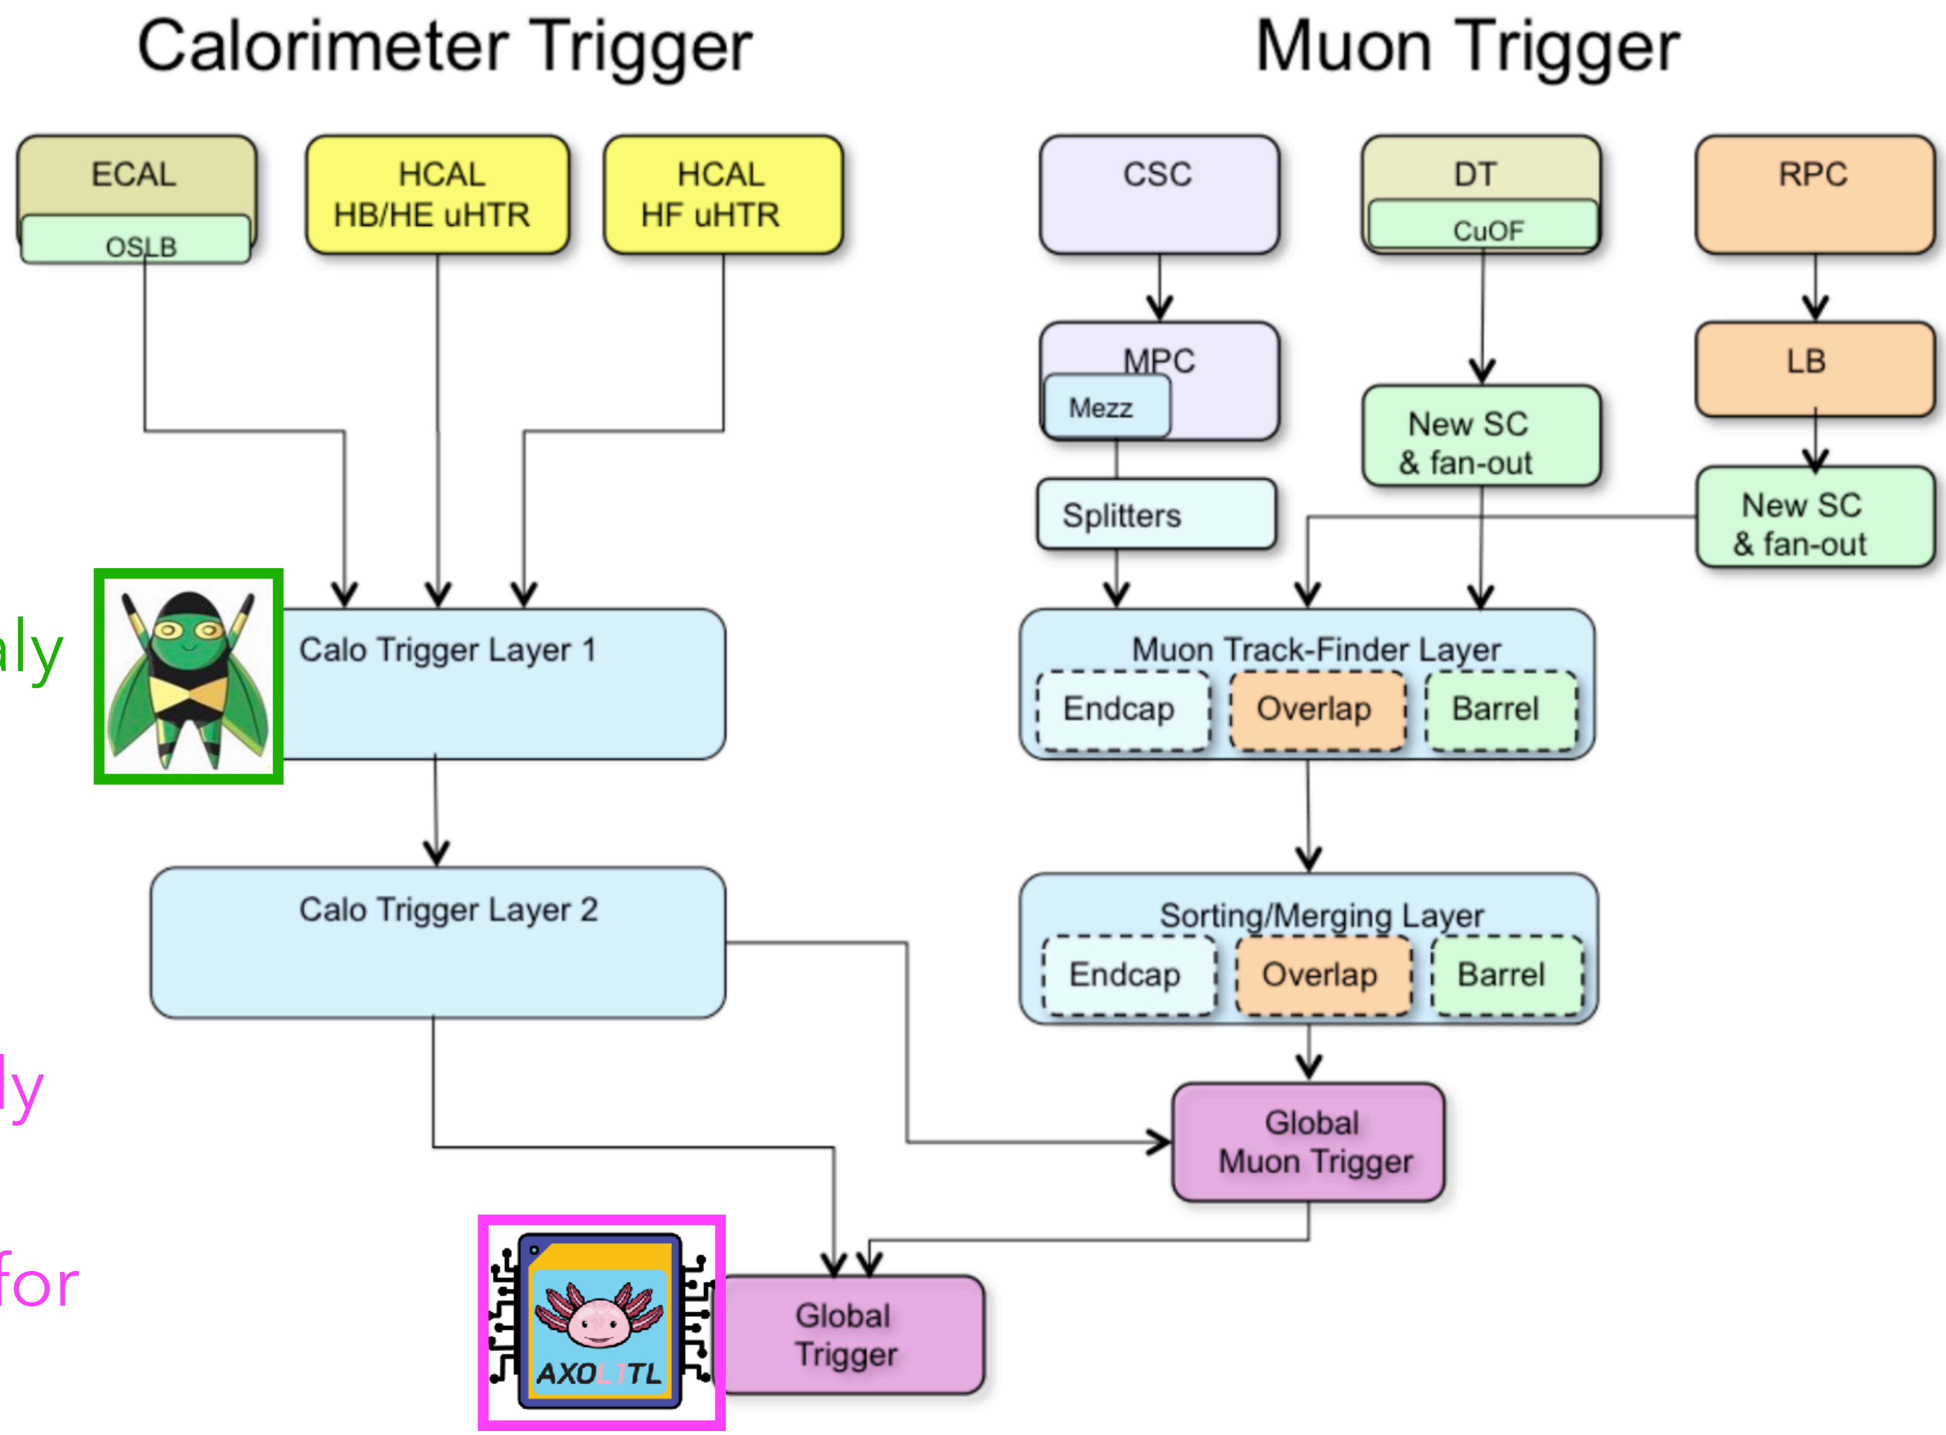

Anomaly Detection with AXOL1TL

- "Anomaly eXtraction Online L1 Trigger Lightweight" → Anomaly detection at L1 Global Trigger in Run 3
  - Autoencoder that inputs L1TGlobal trigger objects
  - On target to be deployed for start of collisions this year!

AXOL1TL Design

- **Variational autoencoder**:
  - Encodes input as a distribution over the latent space
  - Additional loss term regularizes latent space to be Gaussian
- **Inputs L1 trigger objects**: (pT, η, ϕ) of MET, up to 4 electron/photons, 4 muons, and 10 jets
- **Train on ZeroBias data** collected by CMS in 2023 at $\sqrt{s}$=13.6 TeV, 10.5 million events 50/50 training/testing

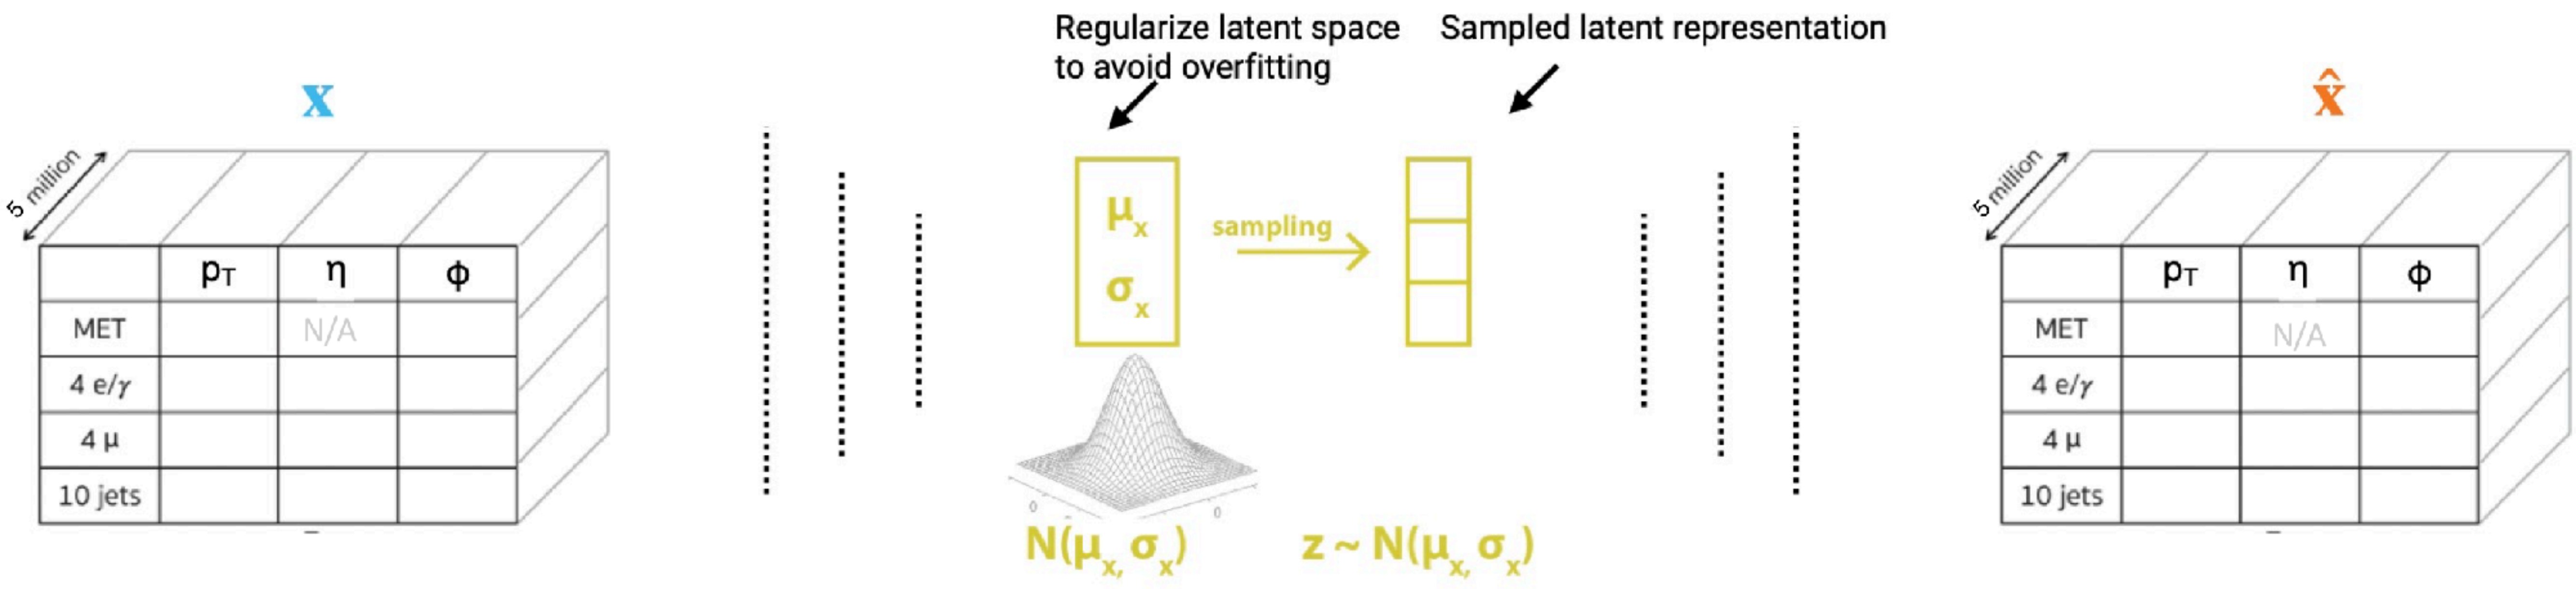

$$\text{loss} = ||x - \hat{x}||^2 + \mathrm{KL}[\, N(\mu_x, \sigma_x), N(0, I)\,]$$

T. Aarrestad

AXOL1TL Design

- Only deploy encoder half of the network, compute degree of abnormality from latent space directly → **Halves the network size and latency!**
- Small, fully connected network architecture
  - **Satisfies strict μGT requirements**: latency of 50ns, low resource utilization on FPGAs
  - **Different anomaly score thresholds can be used to target different trigger rates**

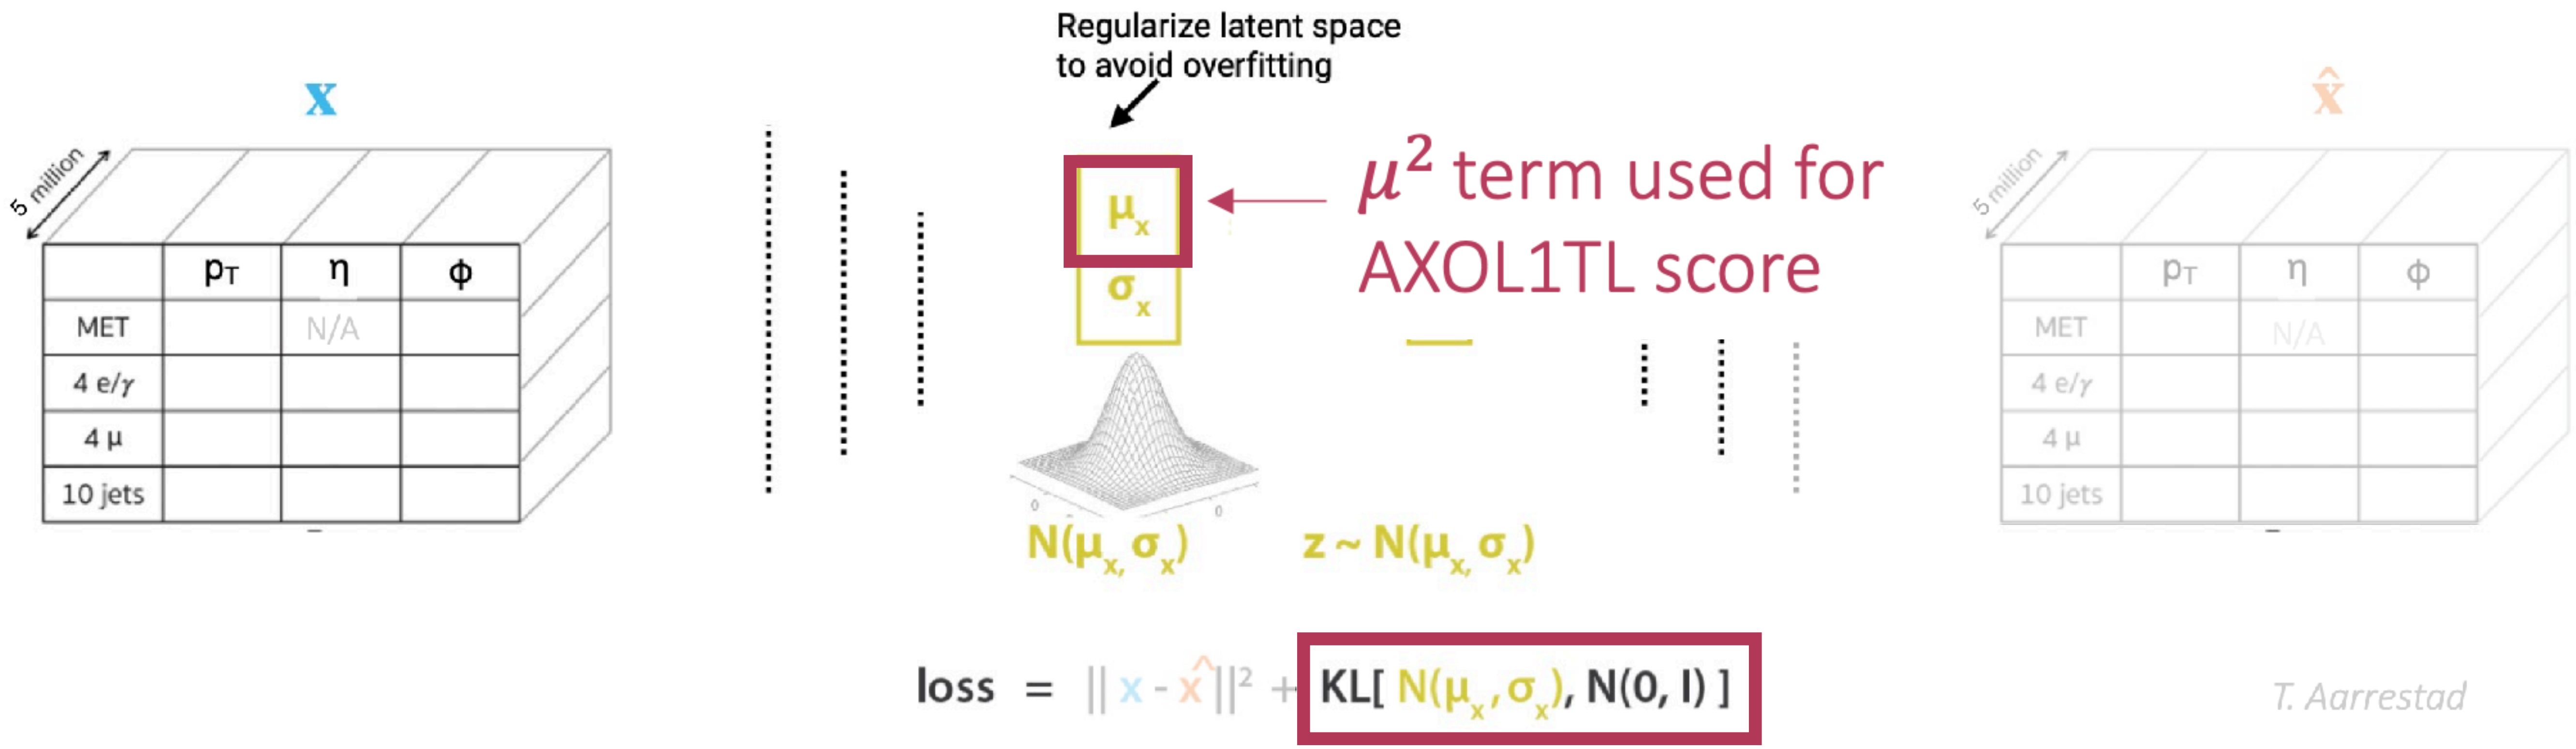

AXOL1TL Run 3 Implementation

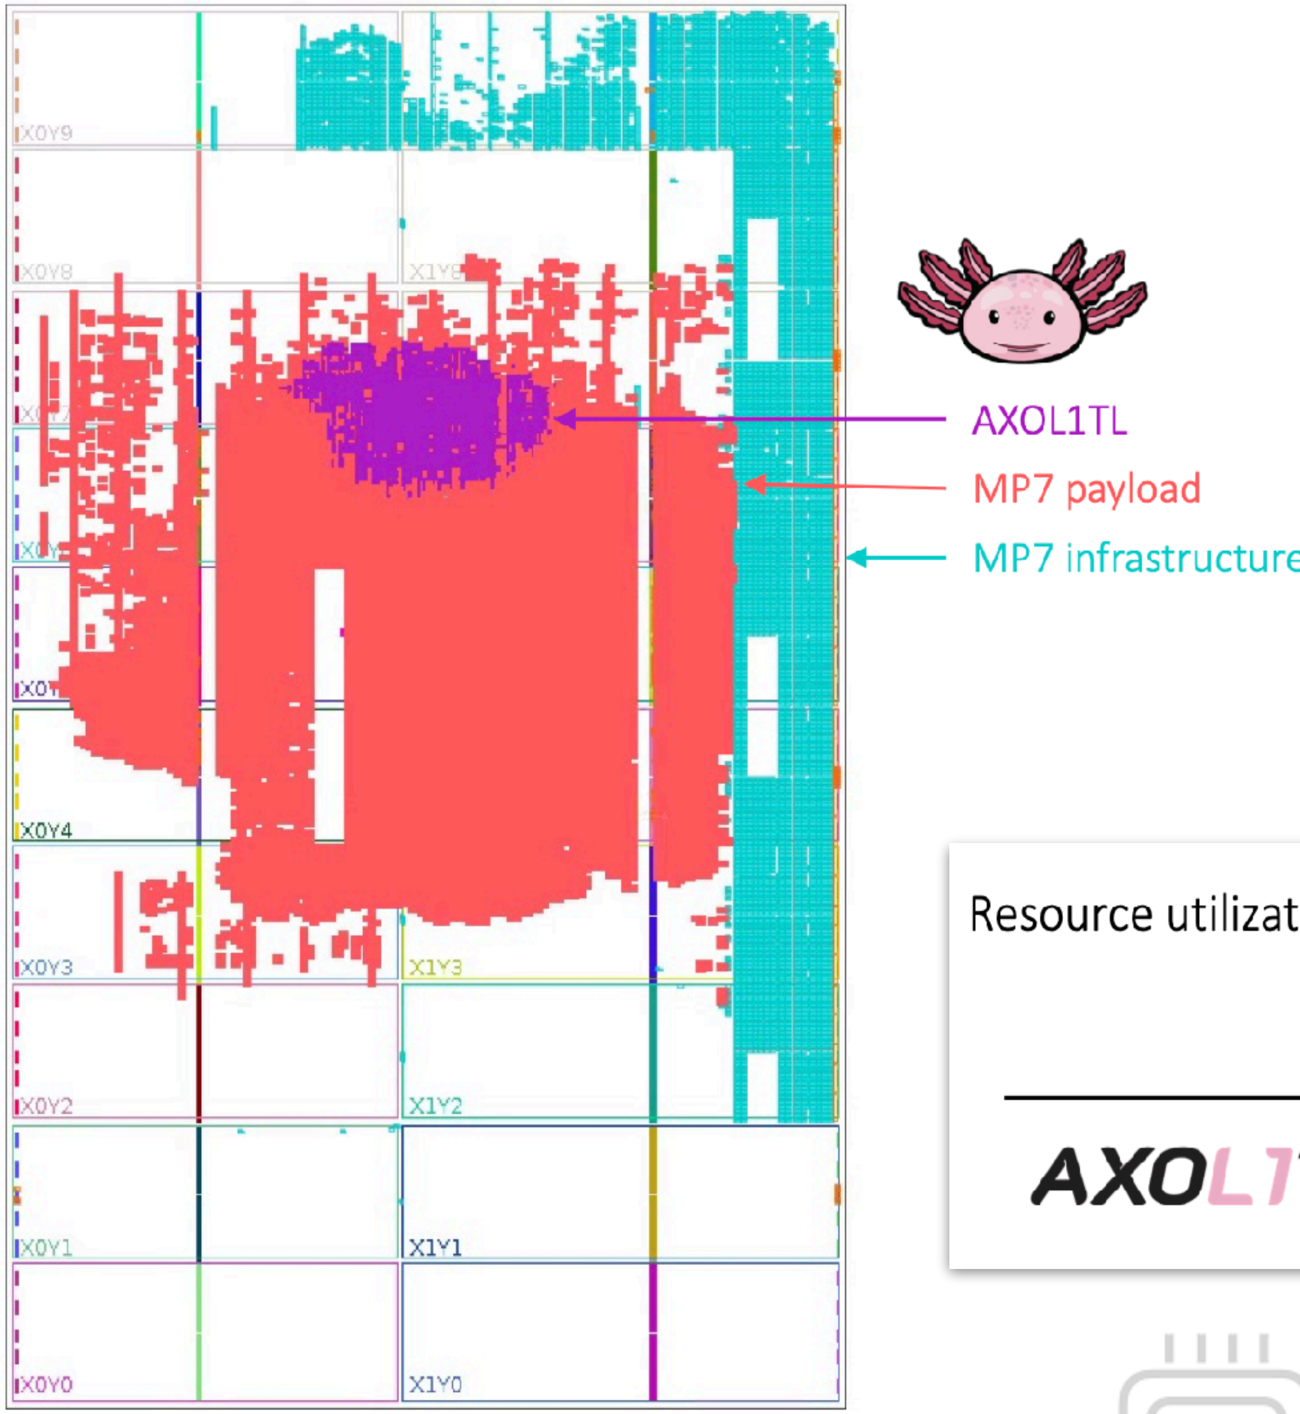

- Implemented on L1 Xilinx Virtex-7 XCVU9P FPGA
- **Satisfies strict μGT requirements**: latency of 50ns, low resource utilization on FPGAs

Resource utilization of Virtex-7 FPGA chip on Imperial College MP7 µGT board

	Latency	LUTs	FFs	DSPs	BRAMs
AXOL1TL	2 ticks 50 ns	2.1%	~0	0	0

hls4ml

Test Crate Implementation

- AXOL1TL **deployed in μGT test crate** in 2023 (DPG JIRA ticket)
  - 4 trigger seeds for different rate/score thresholds
  - ADT_80, ADT_400, ADT_4000, ADT_20000
  - Included in L1Menu_Collisions2022_v1_4_0_adt-d1.xml menu
- **Rates stable relative to other L1 triggers**

CMS-DP-2023-079

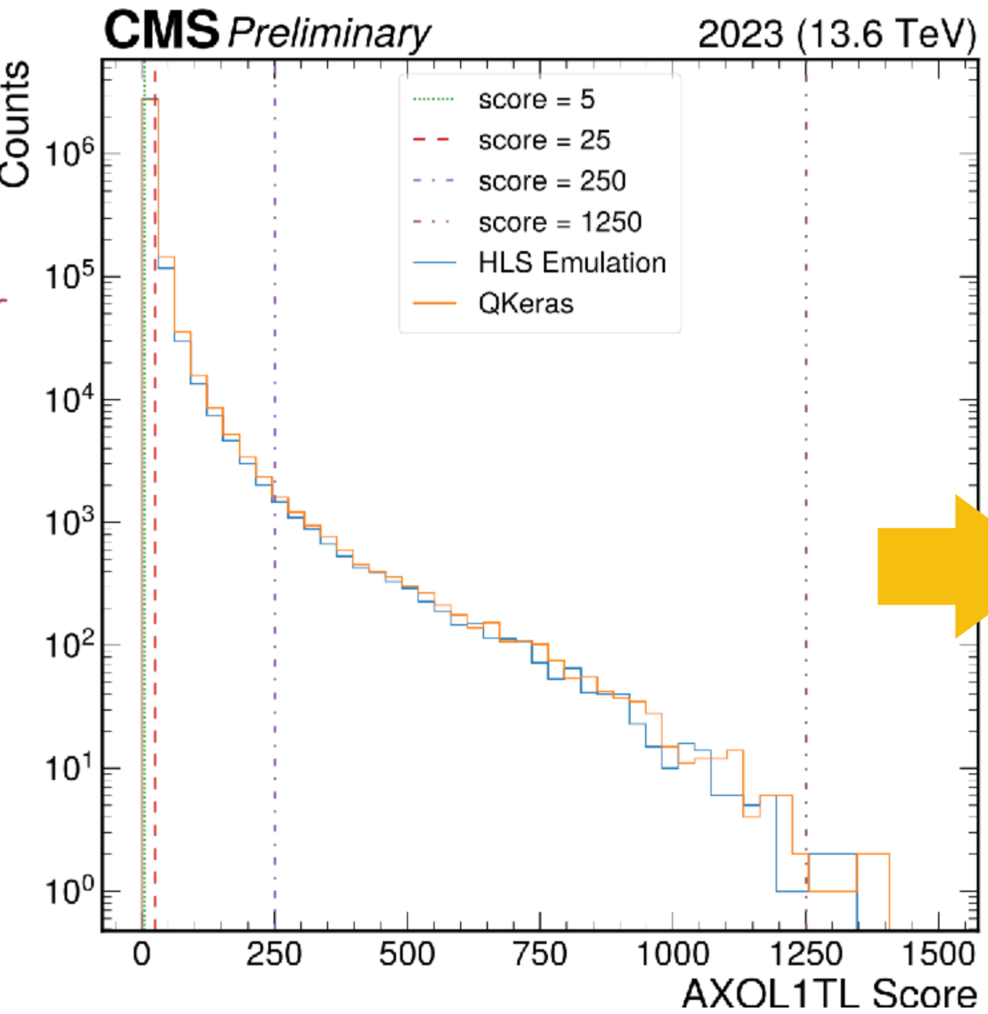

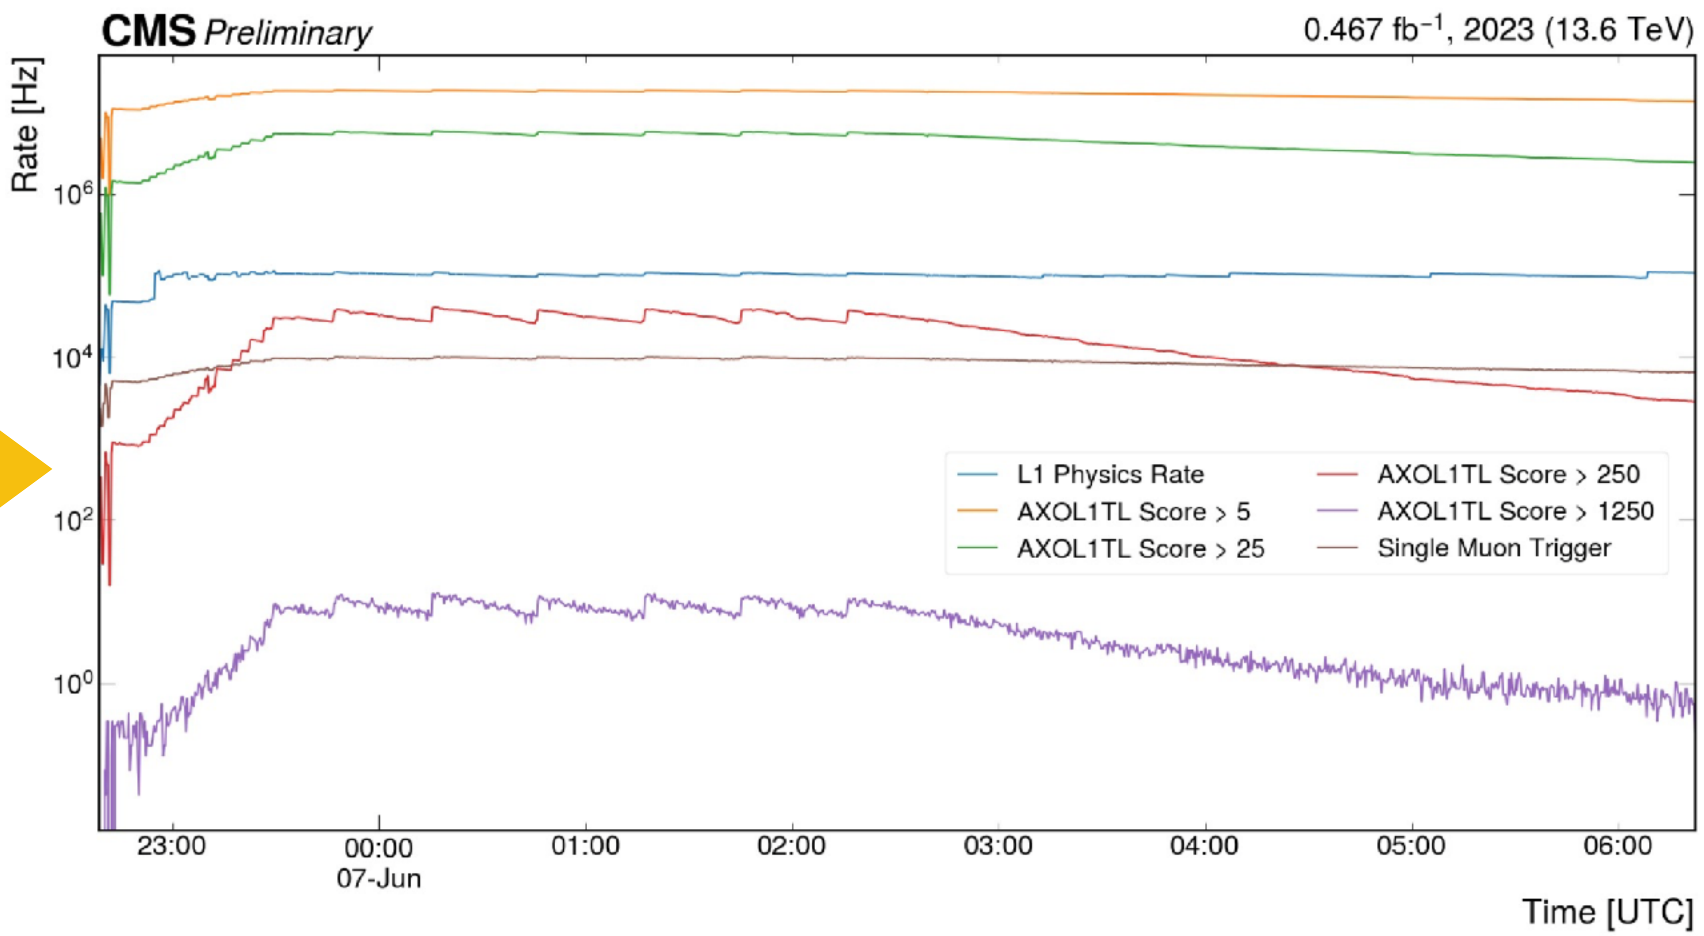

Physics Performance

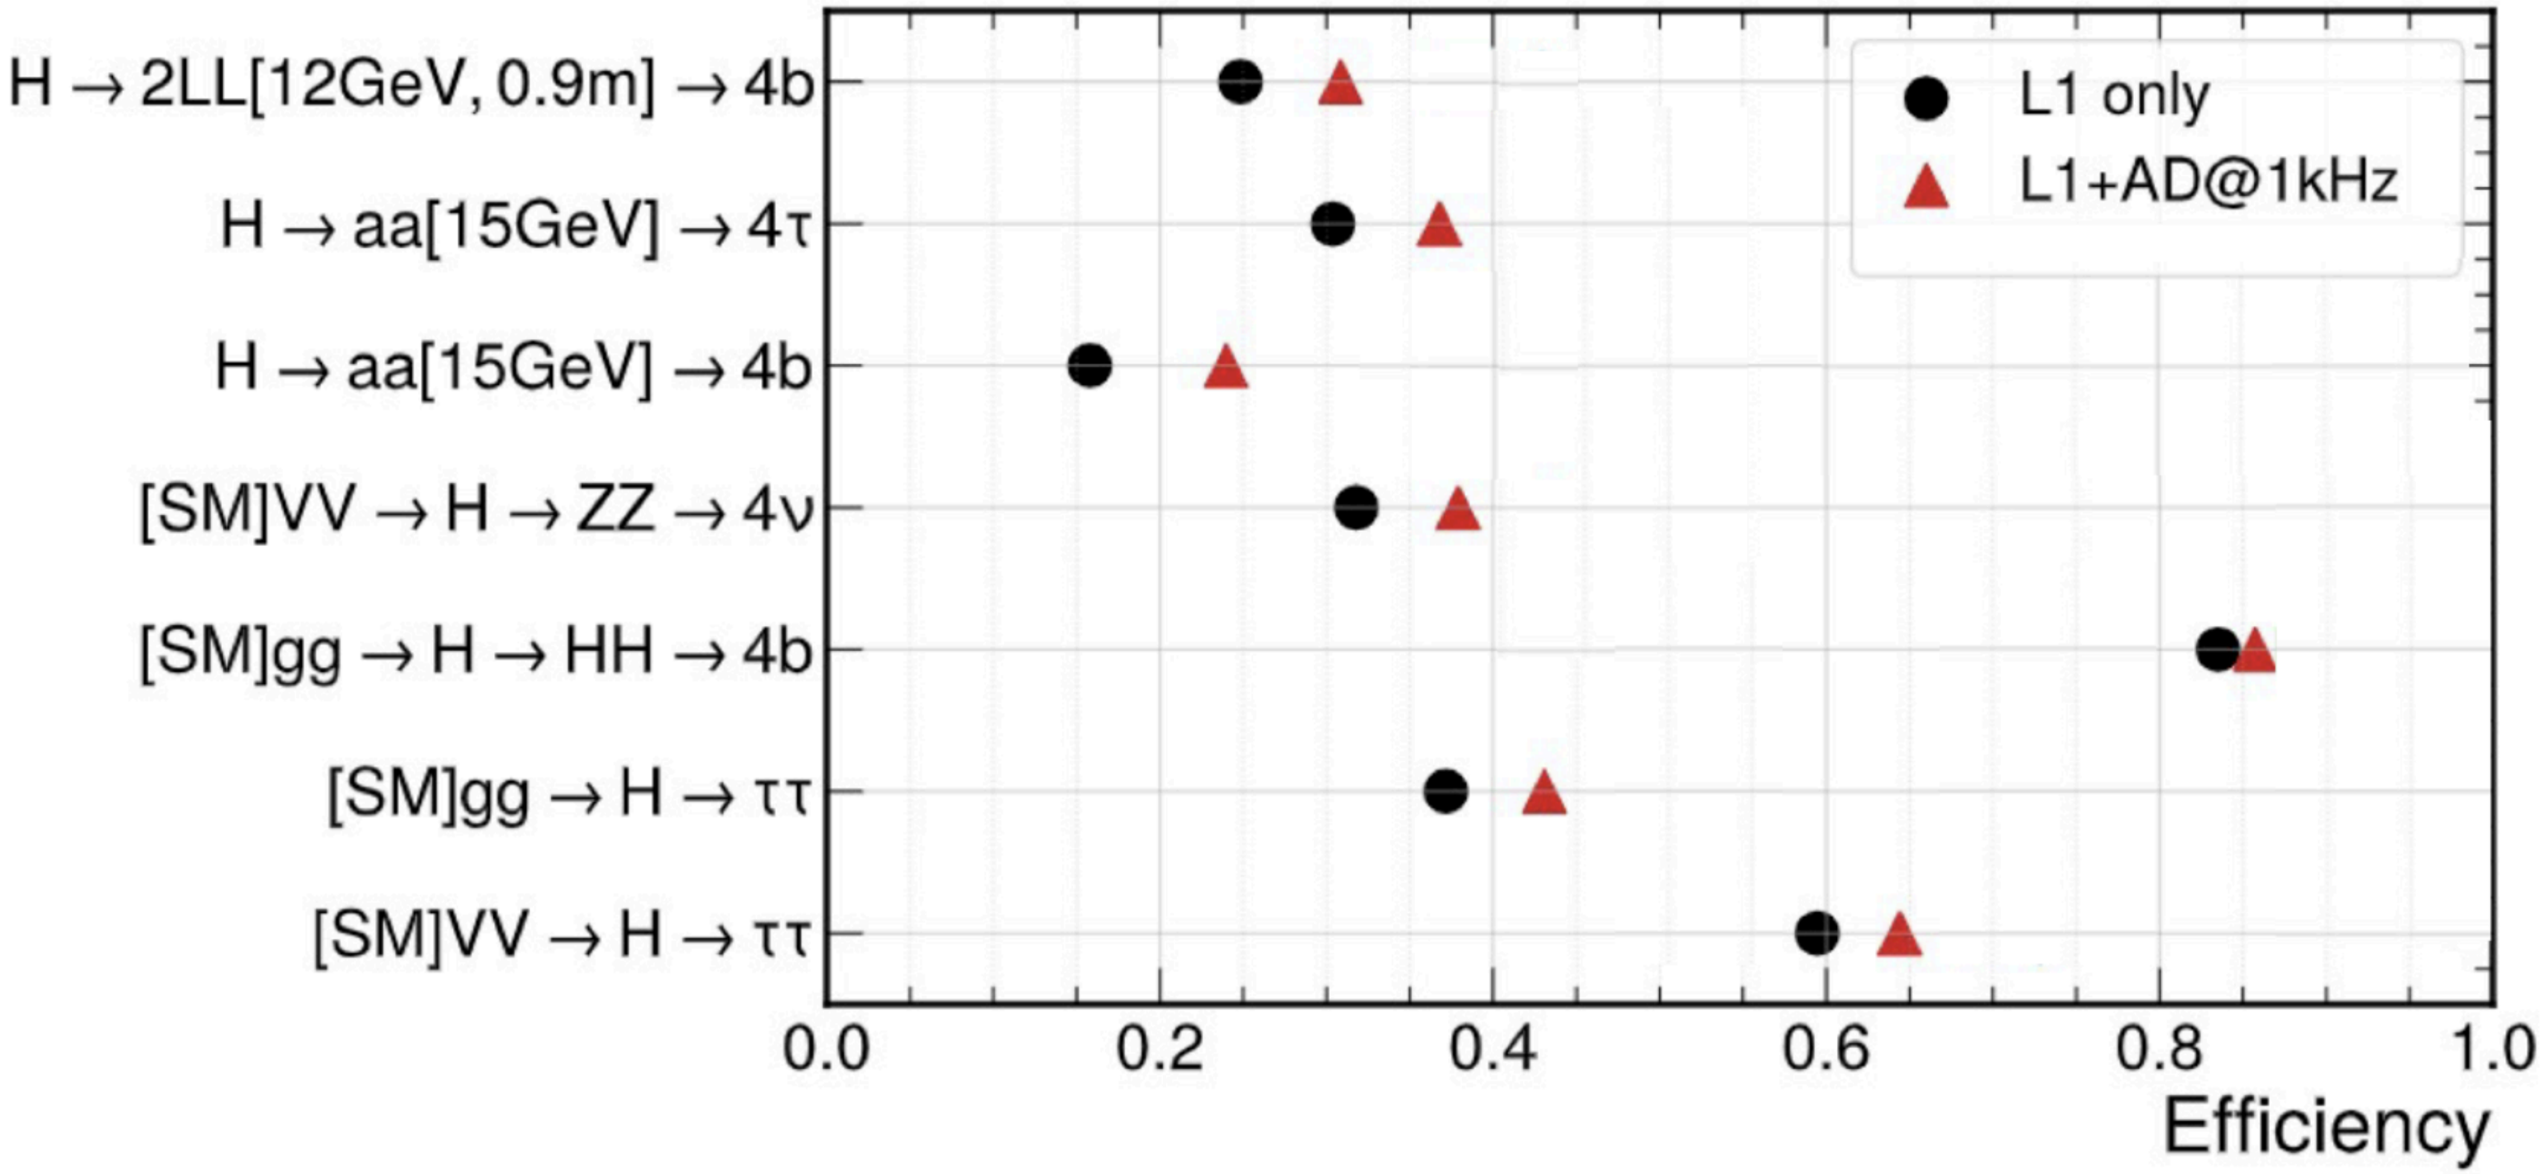

Improves signal efficiency at L1 by ~5%+ for a small rate of 1 kHz relative to L1 rate of ~110 kHz

Physics Performance

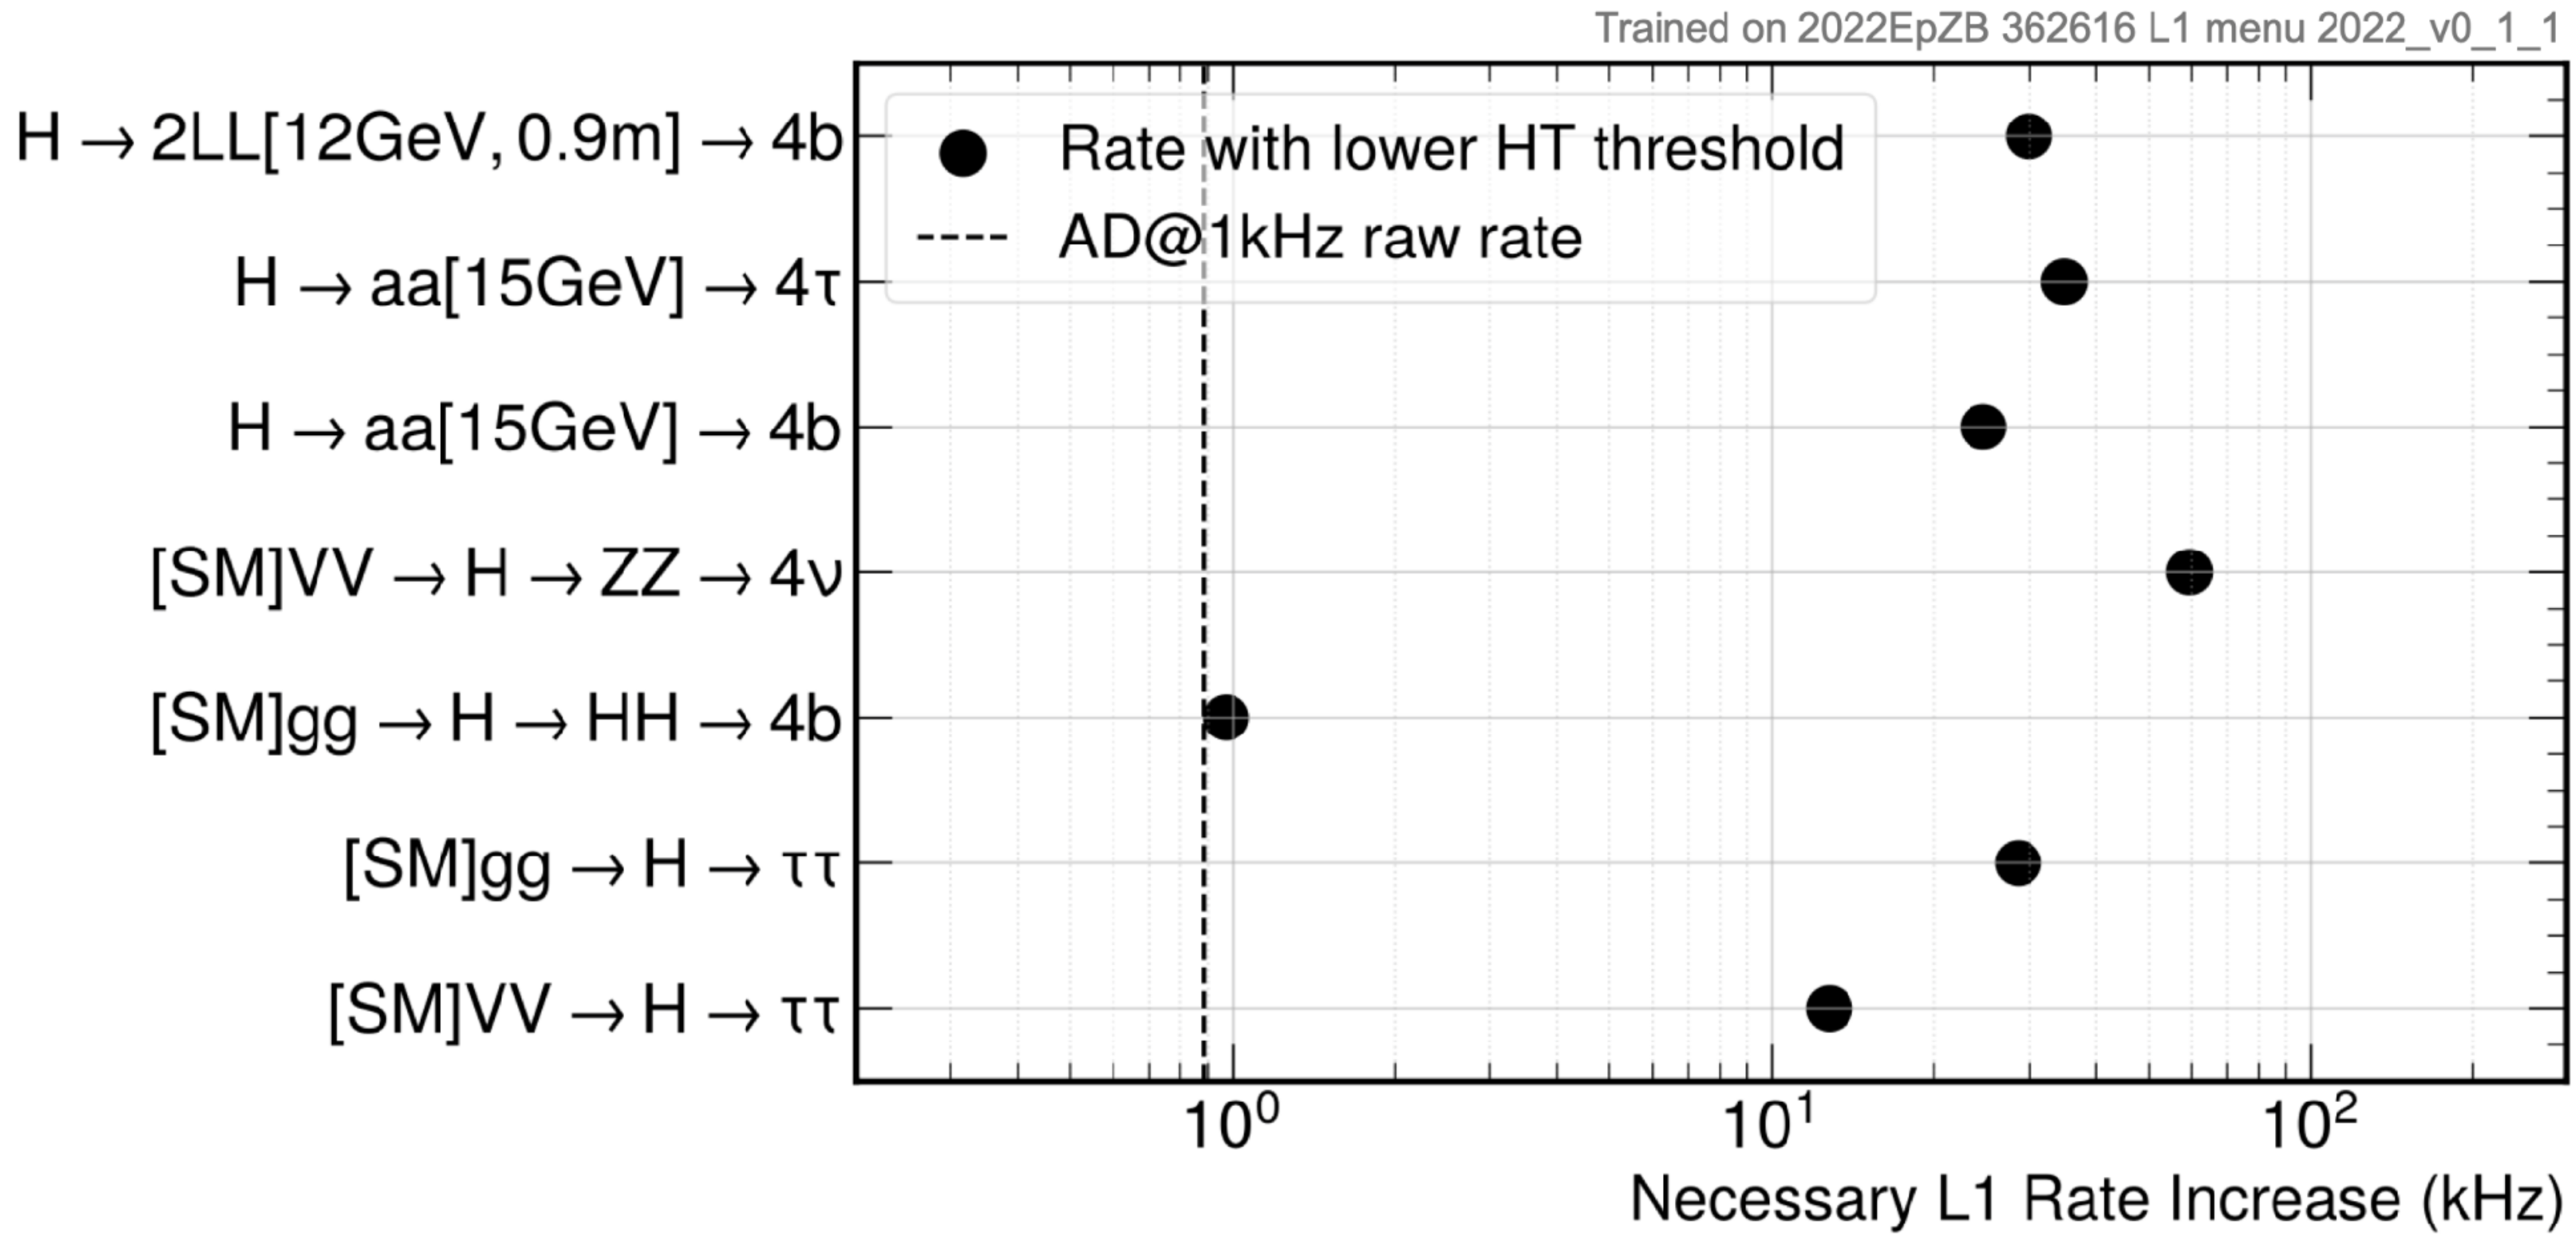

For same efficiency by lowering L1 H_T thresholds, L1 rate would have to increase 10-30x.

Overlap with other L1 triggers

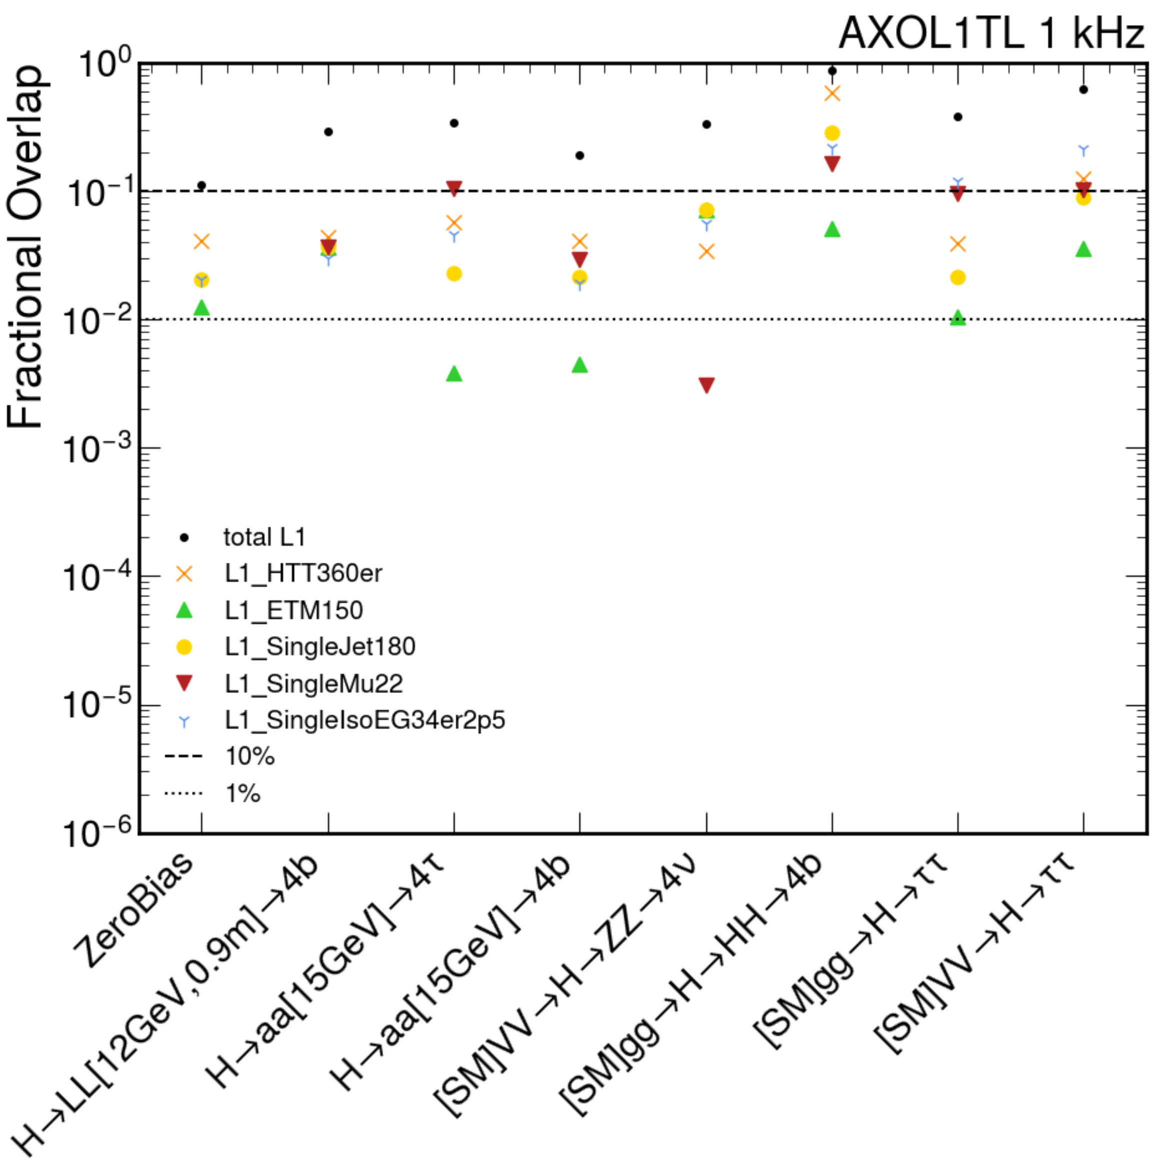

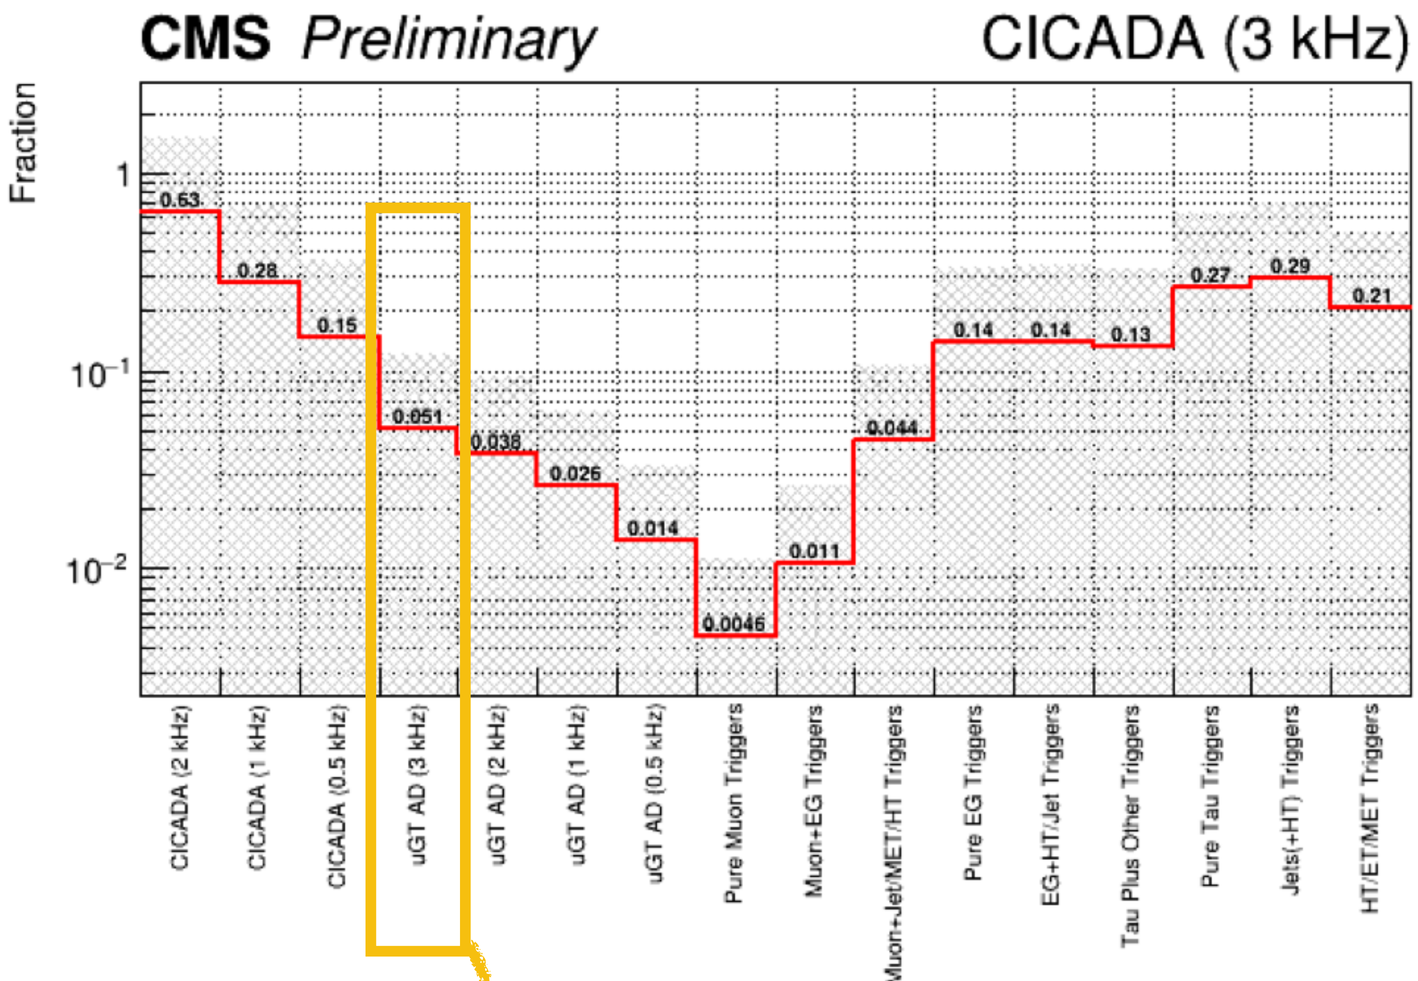

- **Low overlap with existing L1 seeds!**
- ~5% overlap between CICADA 3kHz / AXOL1TL 3kHz

Anomalous Event Example

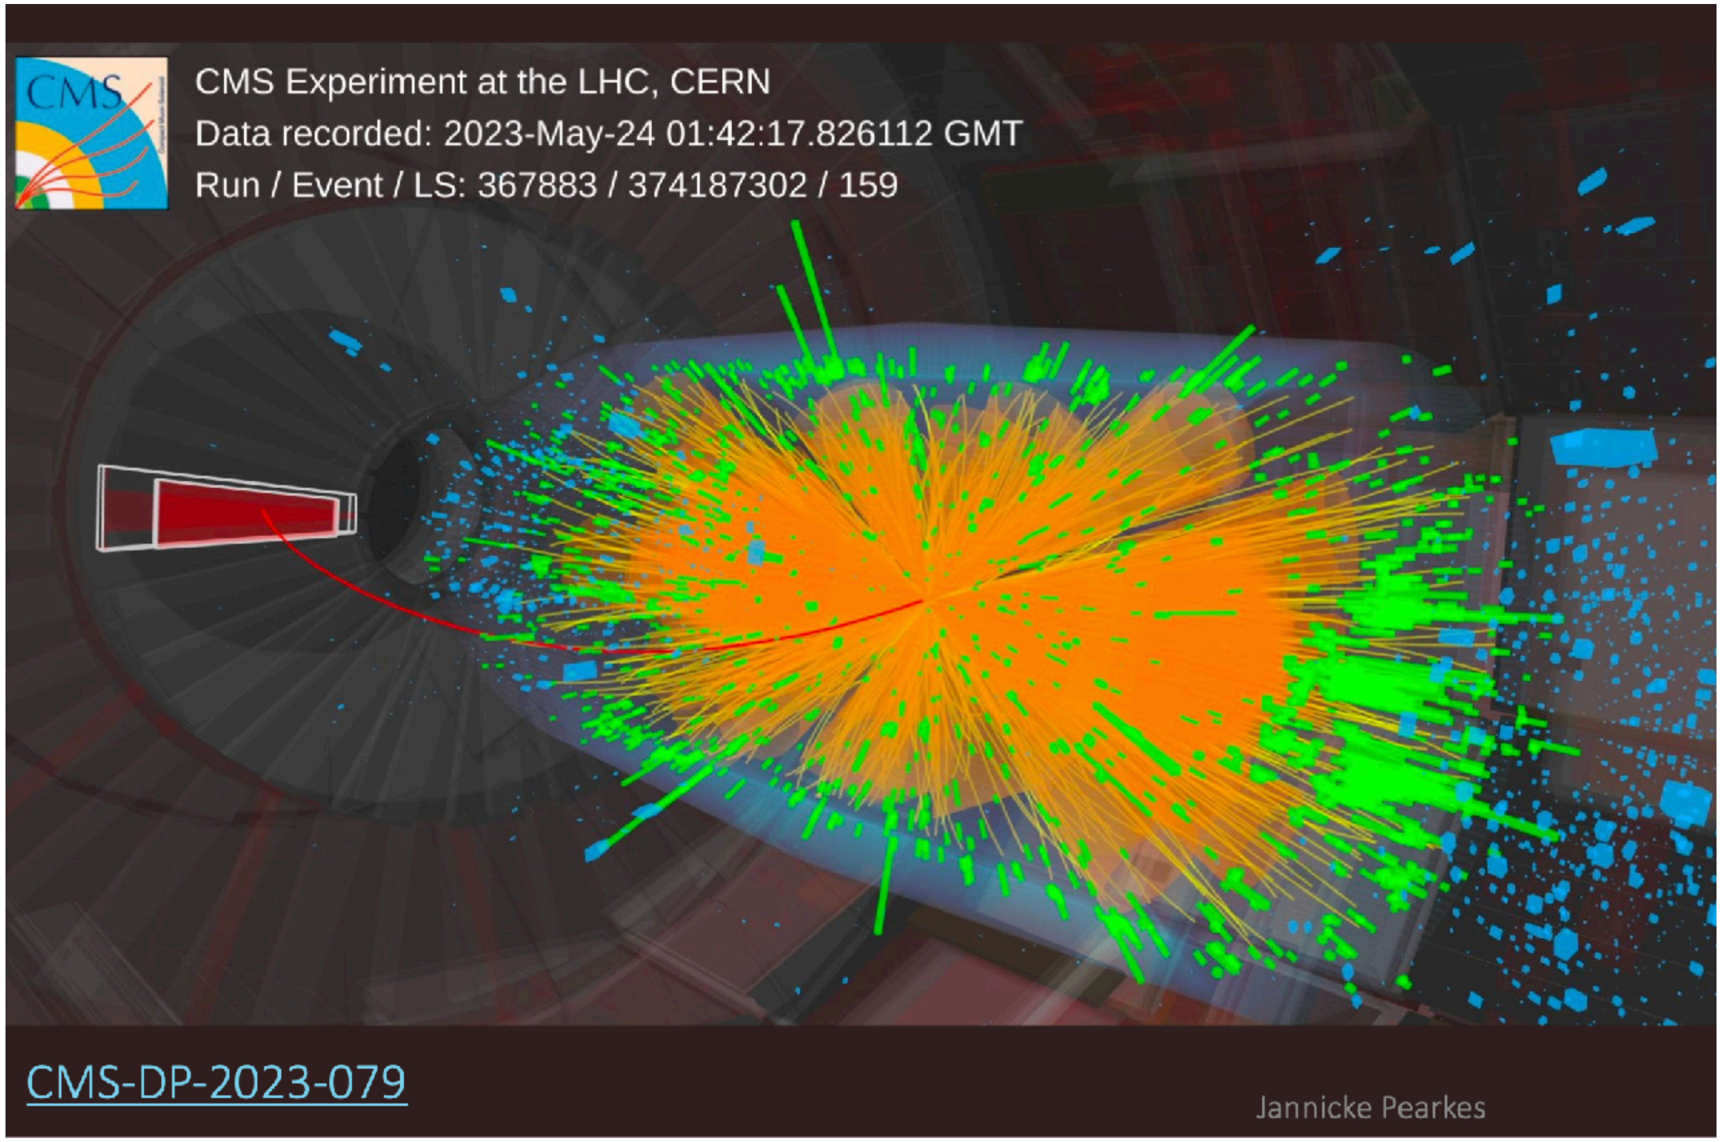

At Level 1:
- 12 jets (11 with E_T > 20 GeV)
- 1 muon with 3 GeV

Offline reconstruction:
- 7 PUPPI jets p_T > 15 GeV
- 1 muon
- 75 reconstructed vertices

- Anomalous event with the **highest AXOL1TL score that was not also triggered by the Level 1 menu in 2023**, Run 367883, Zero Bias data
- Busy event given the pile up profile of the Run 2-2018 training data and data collected in Run 3-2023 (see back-up for PU profiles)

Run 3 Strategy

- Plan to deploy AXOL1TL with low rate as soon as collisions start in 2024
  - Status: L1 trigger seeds and HLT strategy planned, fully integrated into cms-sw software, on target for deployment :)

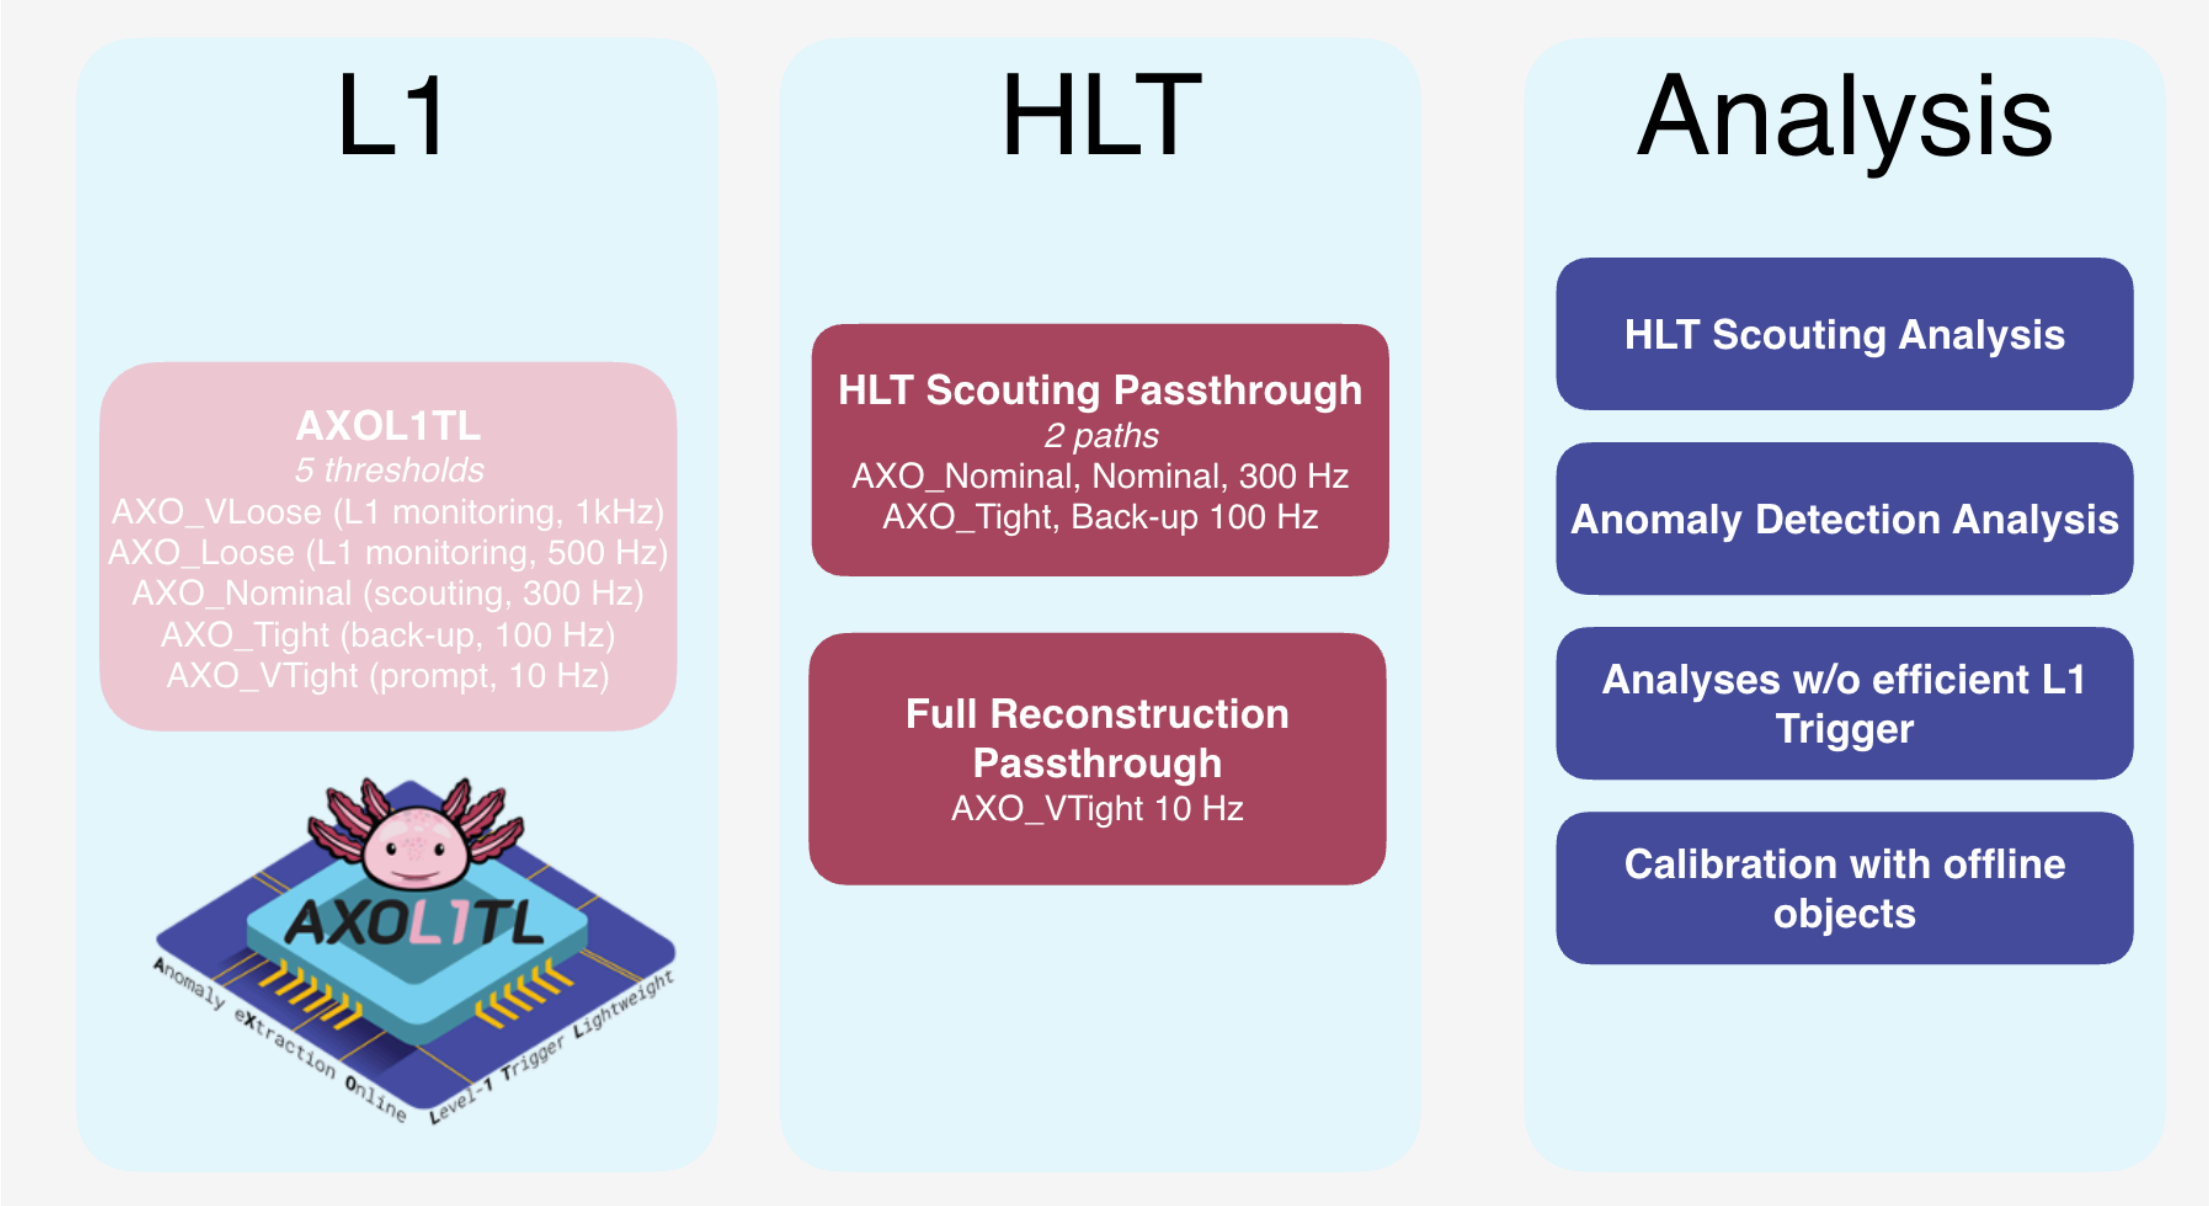

Jannicke Pearkes

AXOL1TL Summary & Next Steps

- **Run 3: → *Explore***
  - On track for deployment at start of collisions this year!
  - All required puzzle pieces in place (firmware, software, menu, HLT paths)
  - Currently adding final checks/ improvements (developing model version resistant to pileup, cleaner/improved software implementation)
- **Phase 2: → *Refine***
  - Run 3 data taking + Phase 2 upgrades will inform strategy and improvements!

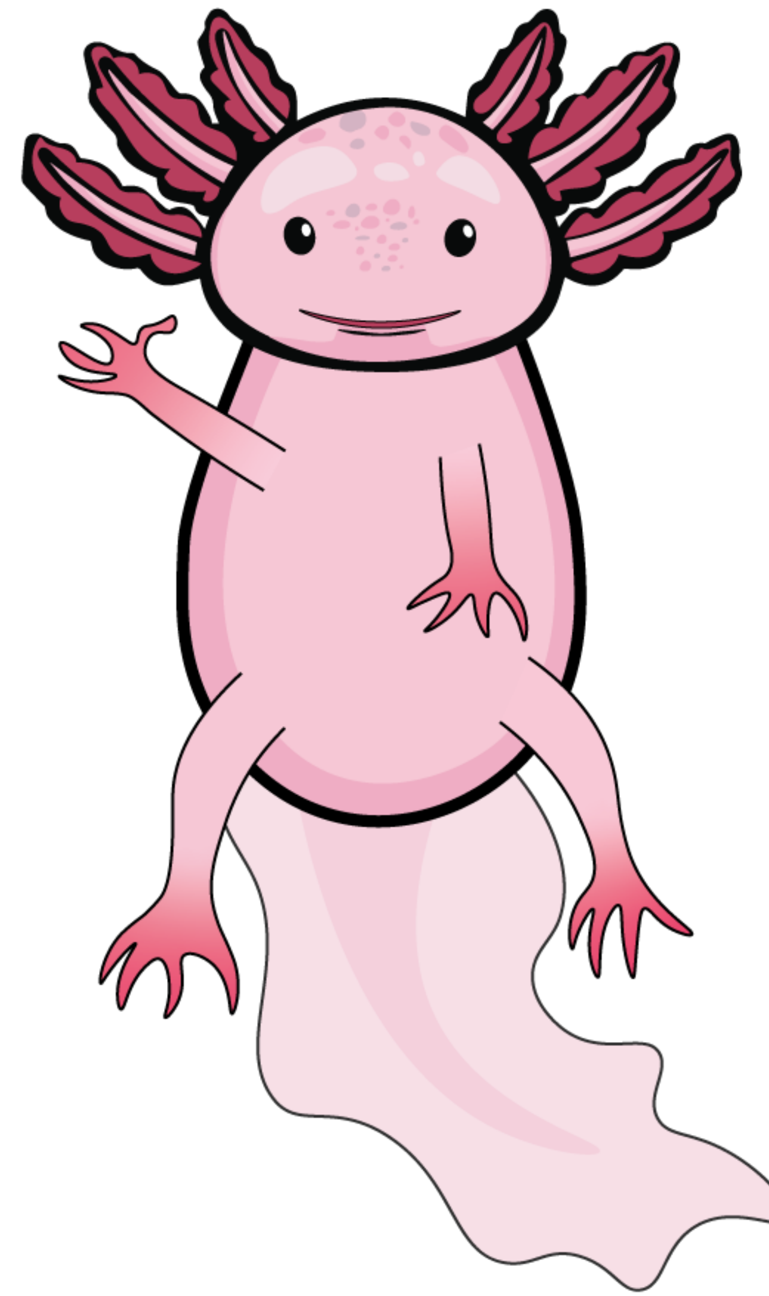

Anomaly Detection with CICADA

CICADA

Public Site

- "Calorimeter Image Convolutional Anomaly Detection Algorithm" → Anomaly detection at calorimeter layer-1 subsystem for L1 trigger
  - Autoencoder that takes in low-level calo trigger inputs
  - Use custom board deployed at L1 calo layer 1
  - Targeting deployment during Run 3 this year

WISCONSIN
UNIVERSITY OF WISCONSIN–MADISON

CICADA Design

- Inputs low level calorimeter energy deposit information from calo layer 1 (both ECAL and HCAL)
- $18\phi \times 14\eta = 252$ calo tower regions = ~2D summary of energy distribution profile for each region

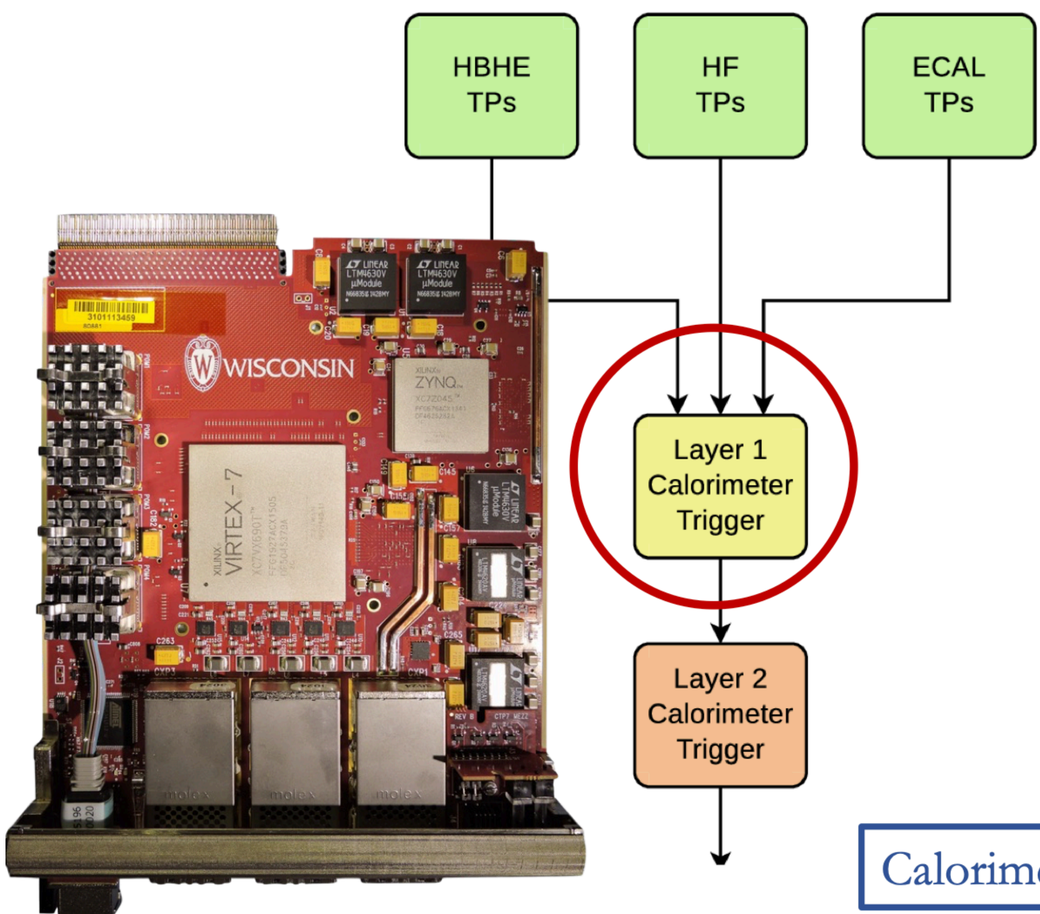

Inputs from CaloLayer1

One region = 4x4 trigger towers

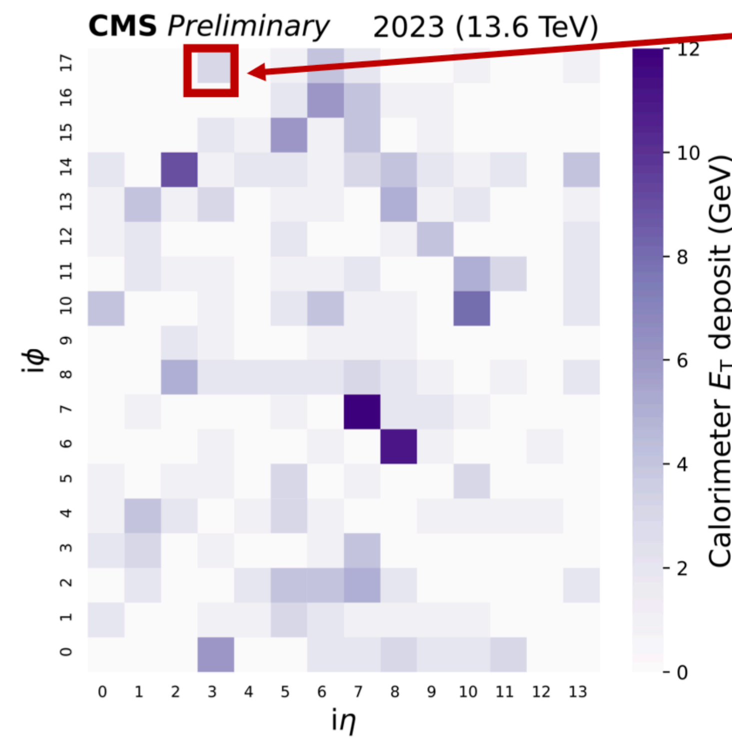

Calorimeter E_T deposit from one Zerobias event

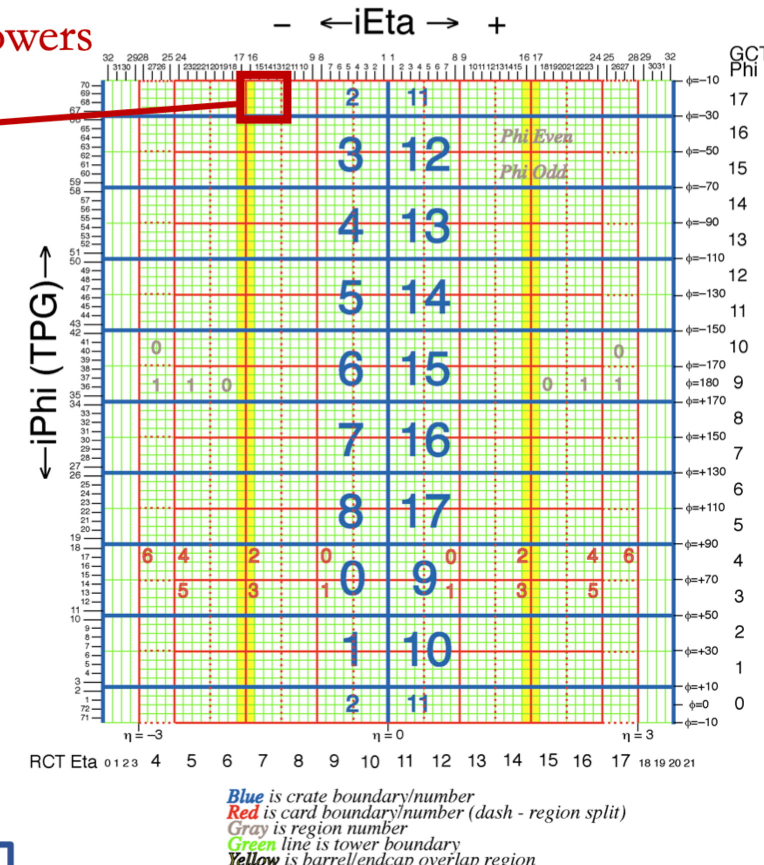

Calorimeter trigger tower-region map

Ho Fung Tsoi

CICADA Design

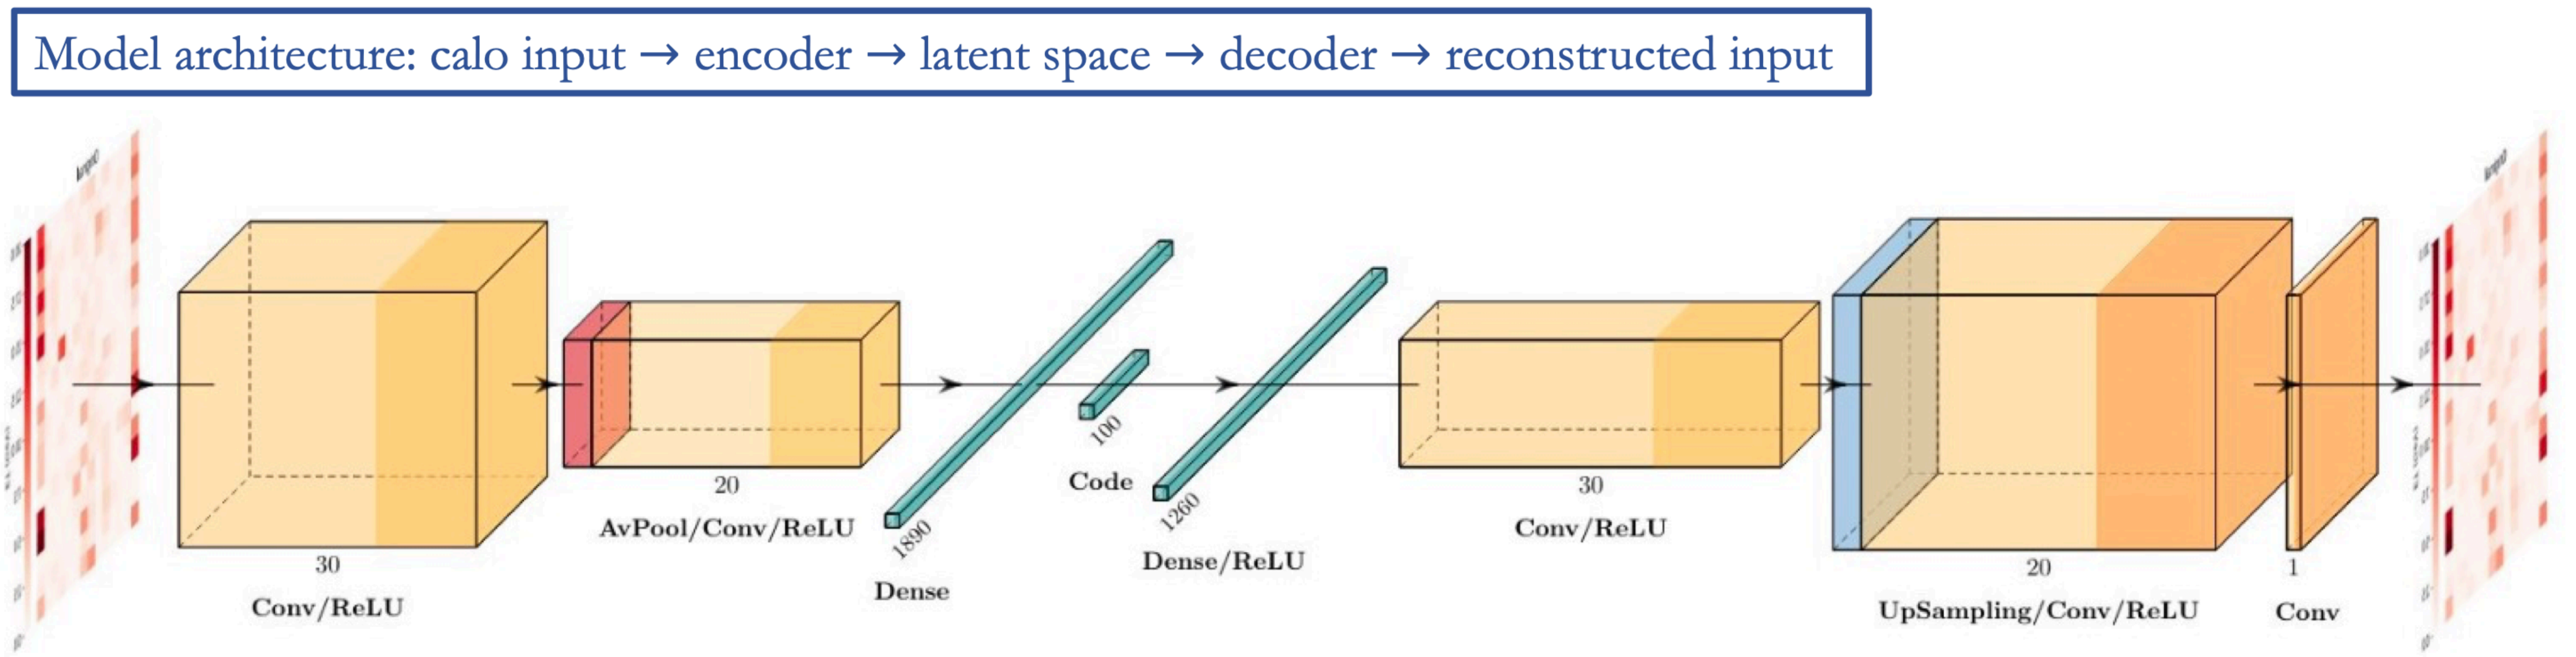

Ho Fung Tsoi

- Unsupervised autoencoder trained on ZeroBias data
  - Input 2D tensor calorimeter energy maps
  - Encoder and decoder are convolutional neural networks
  - Learn reconstruction of input energy maps
  - Use mean-squared error MSE(input, output) as anomaly score to trigger on anomalous events

CICADA Design

ZeroBias Data

BSM Signal

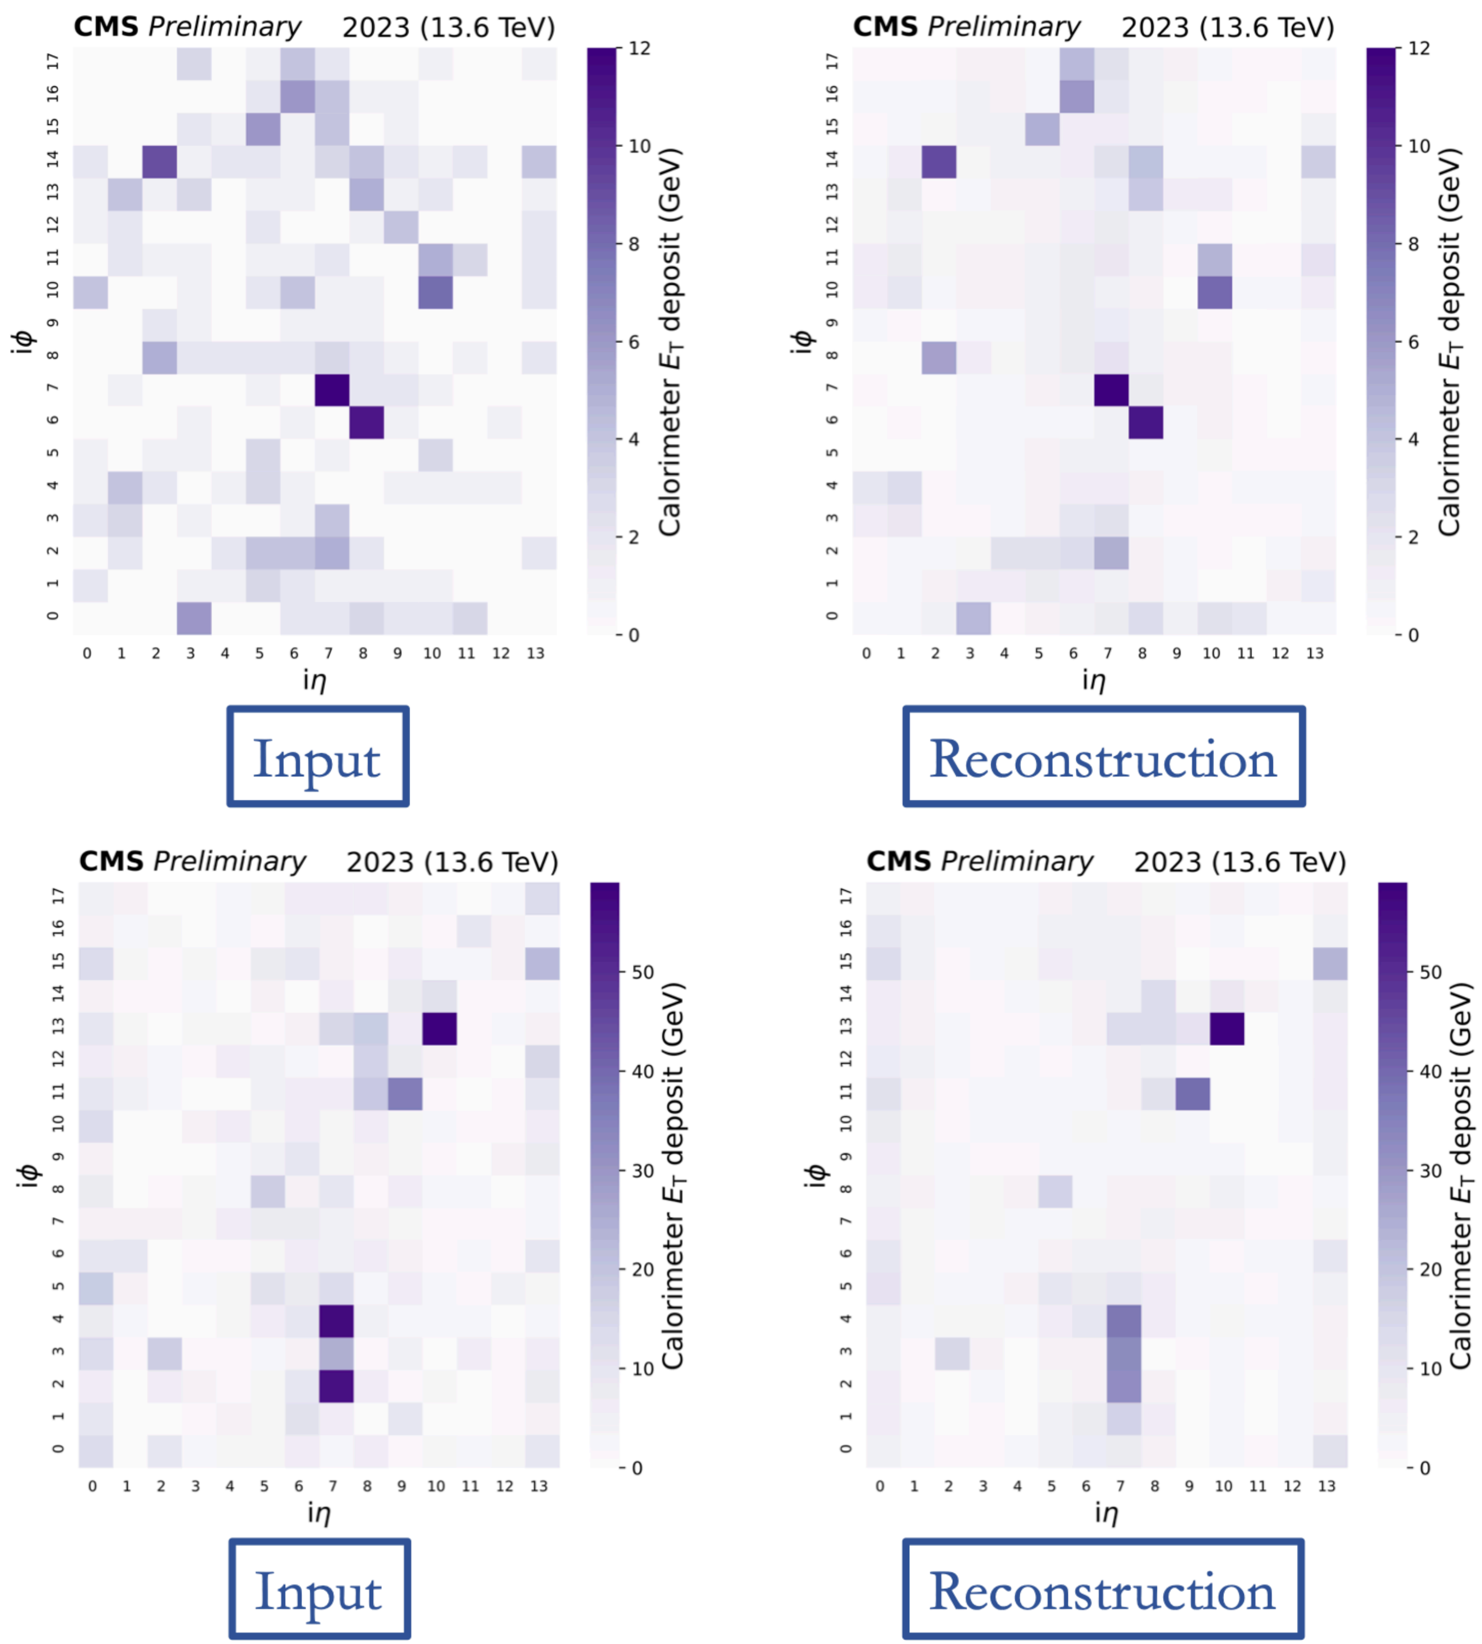

Good reconstruction

Bad reconstruction

Ho Fung Tsoi

Knowledge Distillation

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

$\boxed{\text{Total params: 323,731}}$
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

$\boxed{\text{Total params: 2,927}}$
Trainable params: 2,927
Non-trainable params: 0

Student

Ho Fung Tsoi

- Use knowledge distillation to train a smaller "student" model to replicate "teacher" performance to fit strict L1T constraints
    - Student learns to directly regress MSE from teacher outputs
    - x10 reduction in resources/latency

CICADA Performance/Implementation

- Just like AXOL1TL, select threshold based on desired trigger rate
- Anomaly score distribution distinct for signals vs. ZeroBias data
- Rates stable relative to other L1 triggers

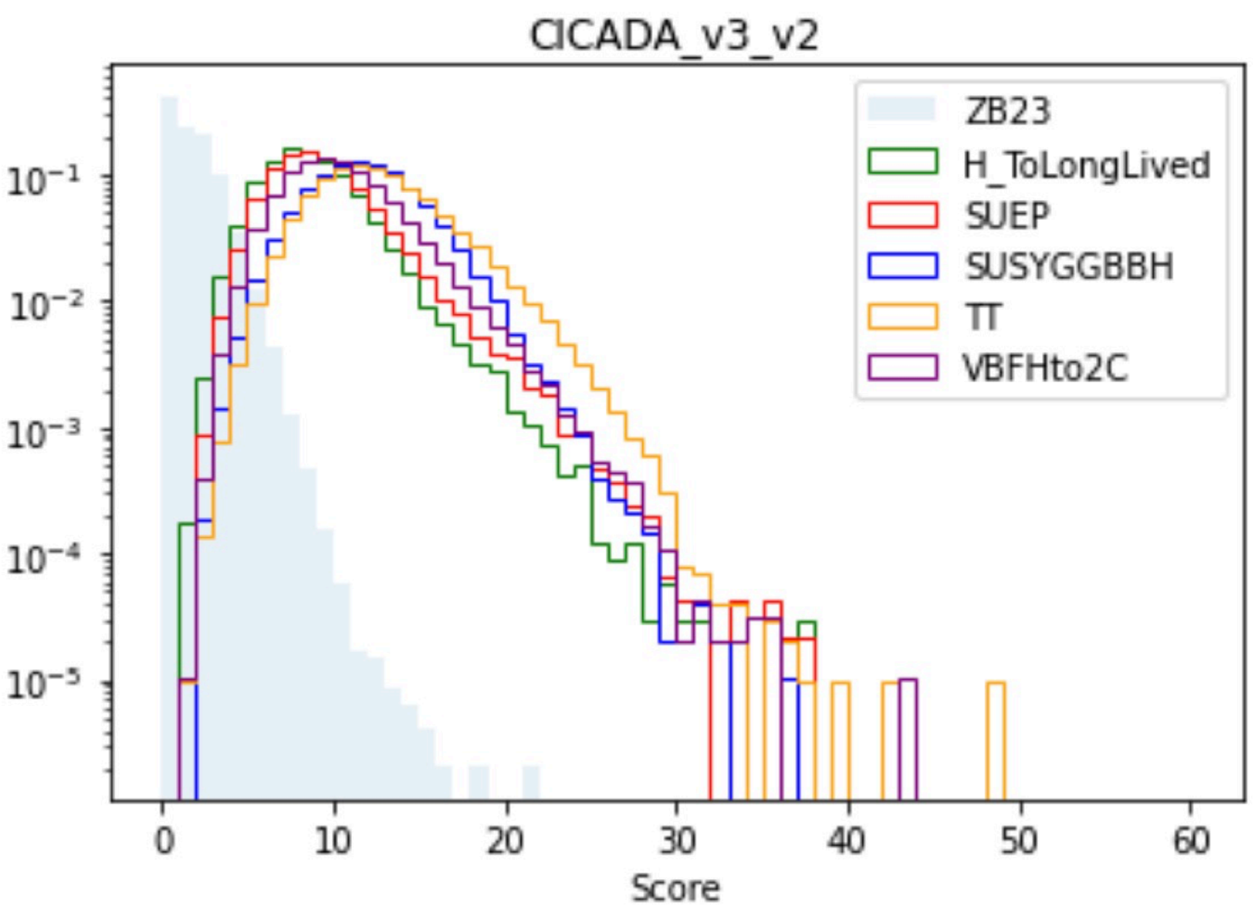

Anomaly score distribution

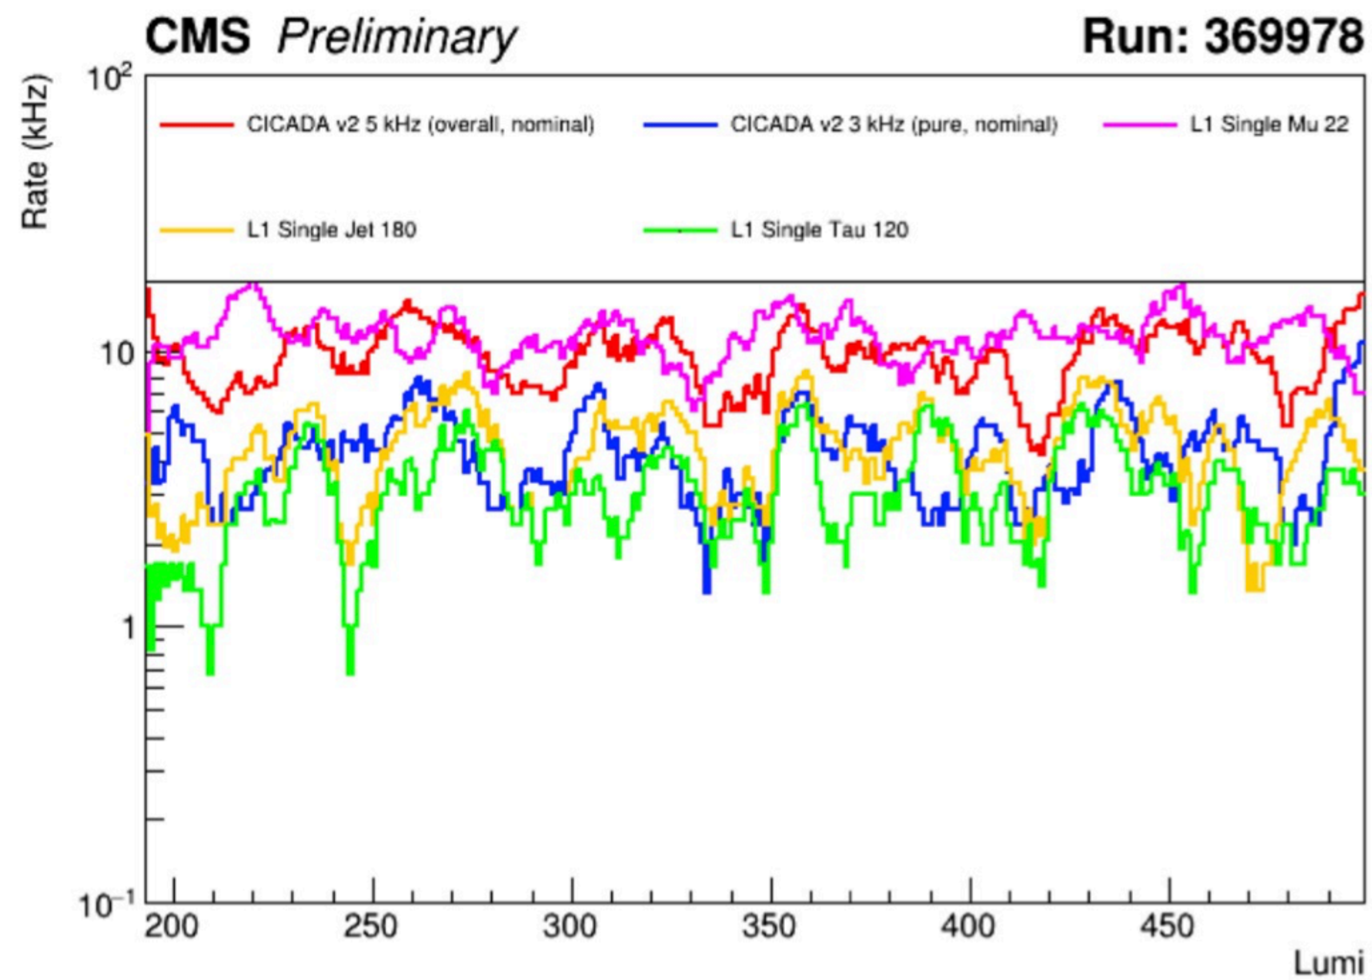

Physics Performance

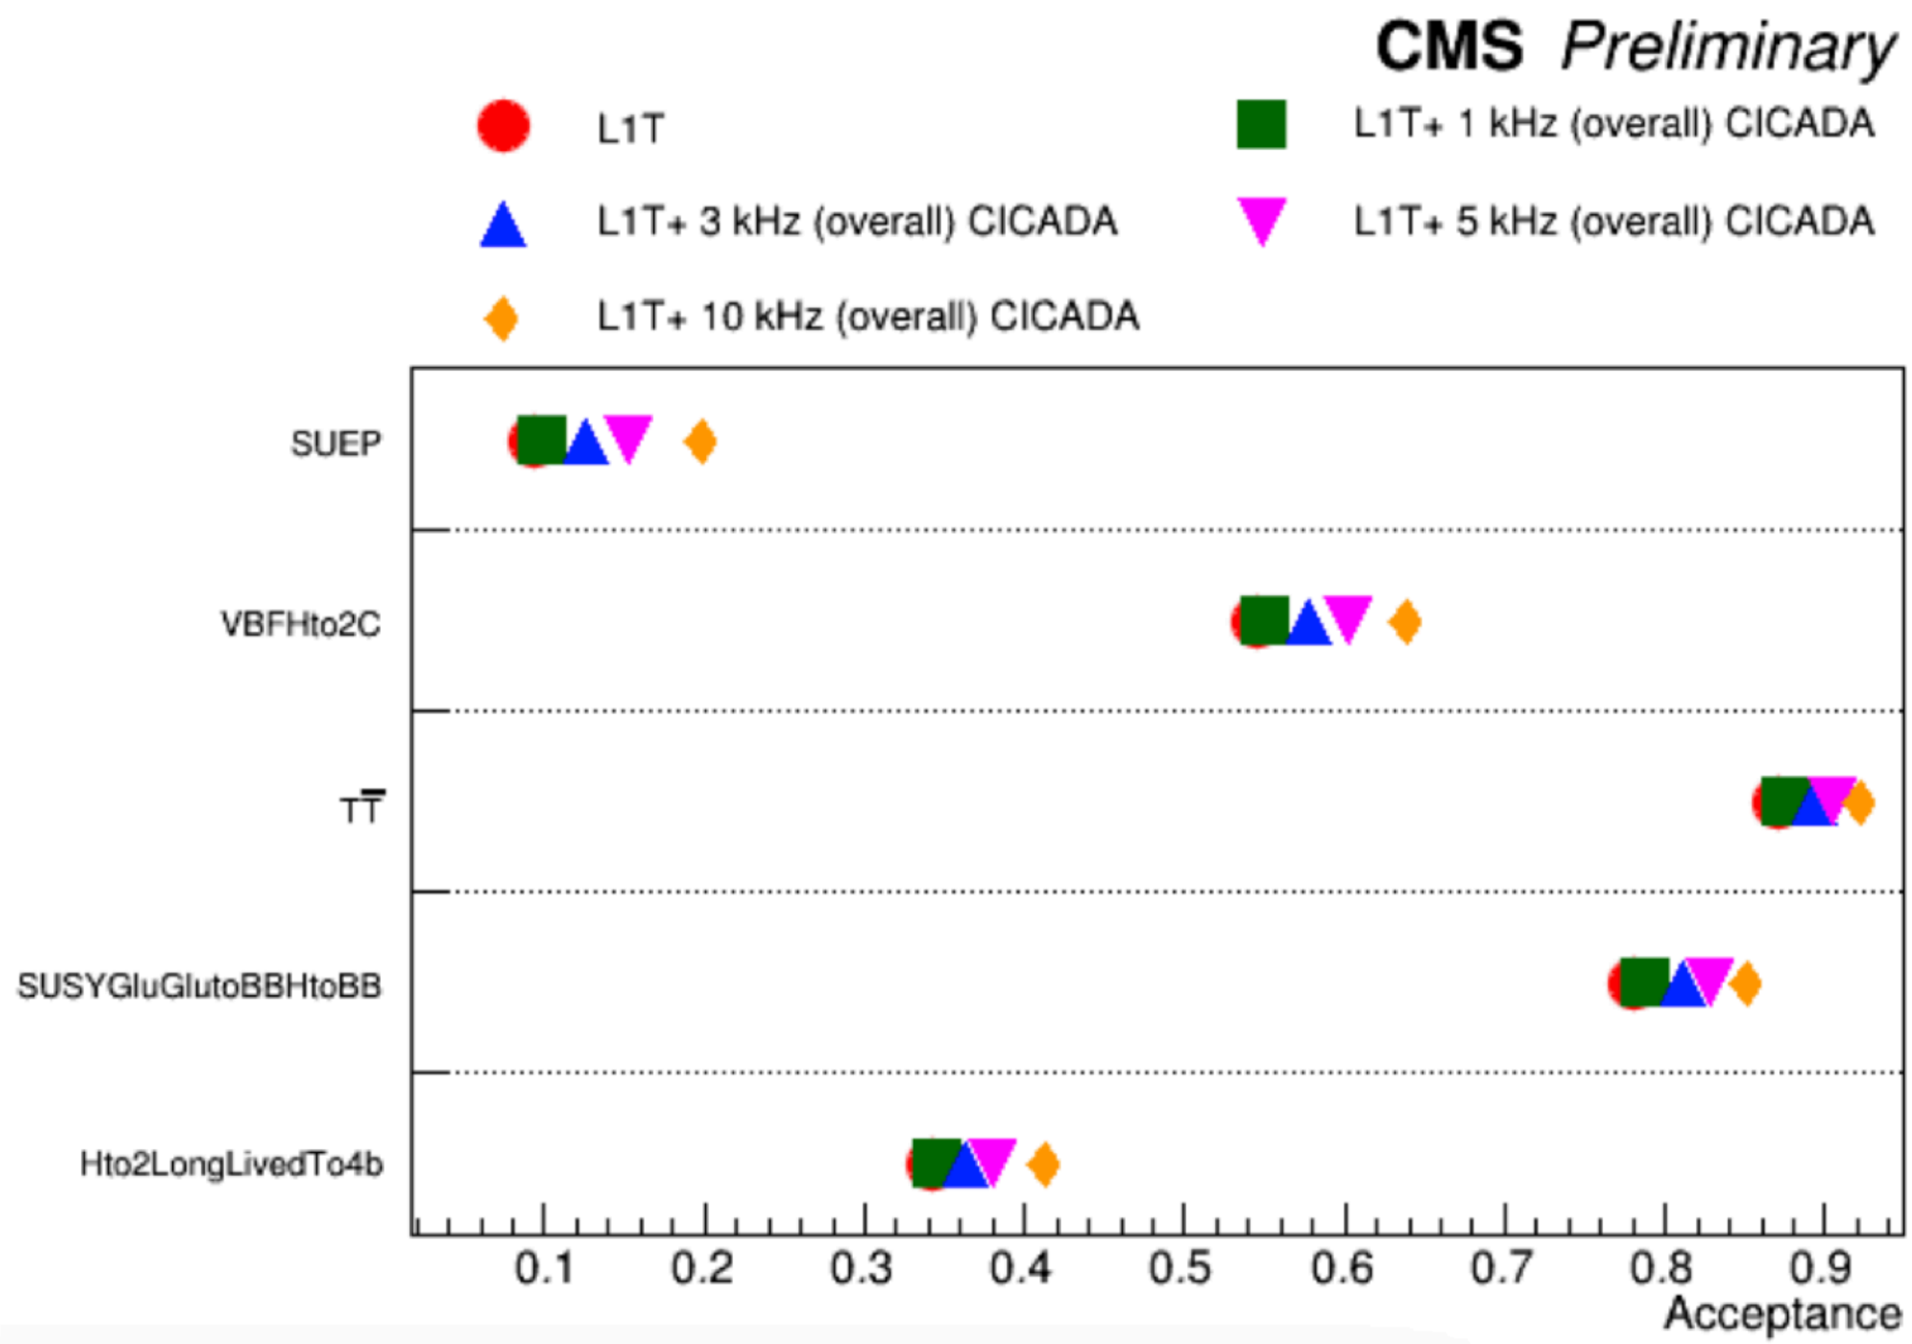

Improves signal efficiency at L1 for rate of 1,3,5, or 10 kHz relative to L1 rate of ~110 kHz

Ho Fung Tsoi

CICADA Summary & Next Steps

- On track for deployment in Run 3 this year
  - Last remaining necessary online software component done (SWATCH Dec1)
  - Firmware/Emulator/Software/Hardware developed and fully functioning
  - Menu and HLT paths in progress

Conclusions

- 2 autoencoder based anomaly detection algorithms planned for deployment in the CMS L1Trigger this year!
    - AXOL1TL for global trigger
    - CICADA for calo layer 1
- Experience of deploying these ML triggers+ the data we collect this year will inform the development of these and other triggers and anomaly detection strategies moving forward!
    - Defining roadmap for future hls4ml-driven ML algorithms to be incorporated in CMS trigger systems

Backup

Documentation

- J. Pearkes - **ML4Jets**, Nov 2023
  https://indico.cern.ch/event/1253794/contributions/5588638/
- C. Sun – **FastML Workshop**, Sept 2023
  https://indico.cern.ch/event/1283970/contributions/5554350/
- N. Zipper – **TWEPP**, Oct 2023
  https://indico.cern.ch/event/1255624/contributions/5444028/

Documentation

- Public website: https://cicada.web.cern.ch/
- CPAD Workshop, 7 Nov 2023 : https://indico.slac.stanford.edu/event/8288/contributions/7717/
-

CICADA

AD Encoder Implementation

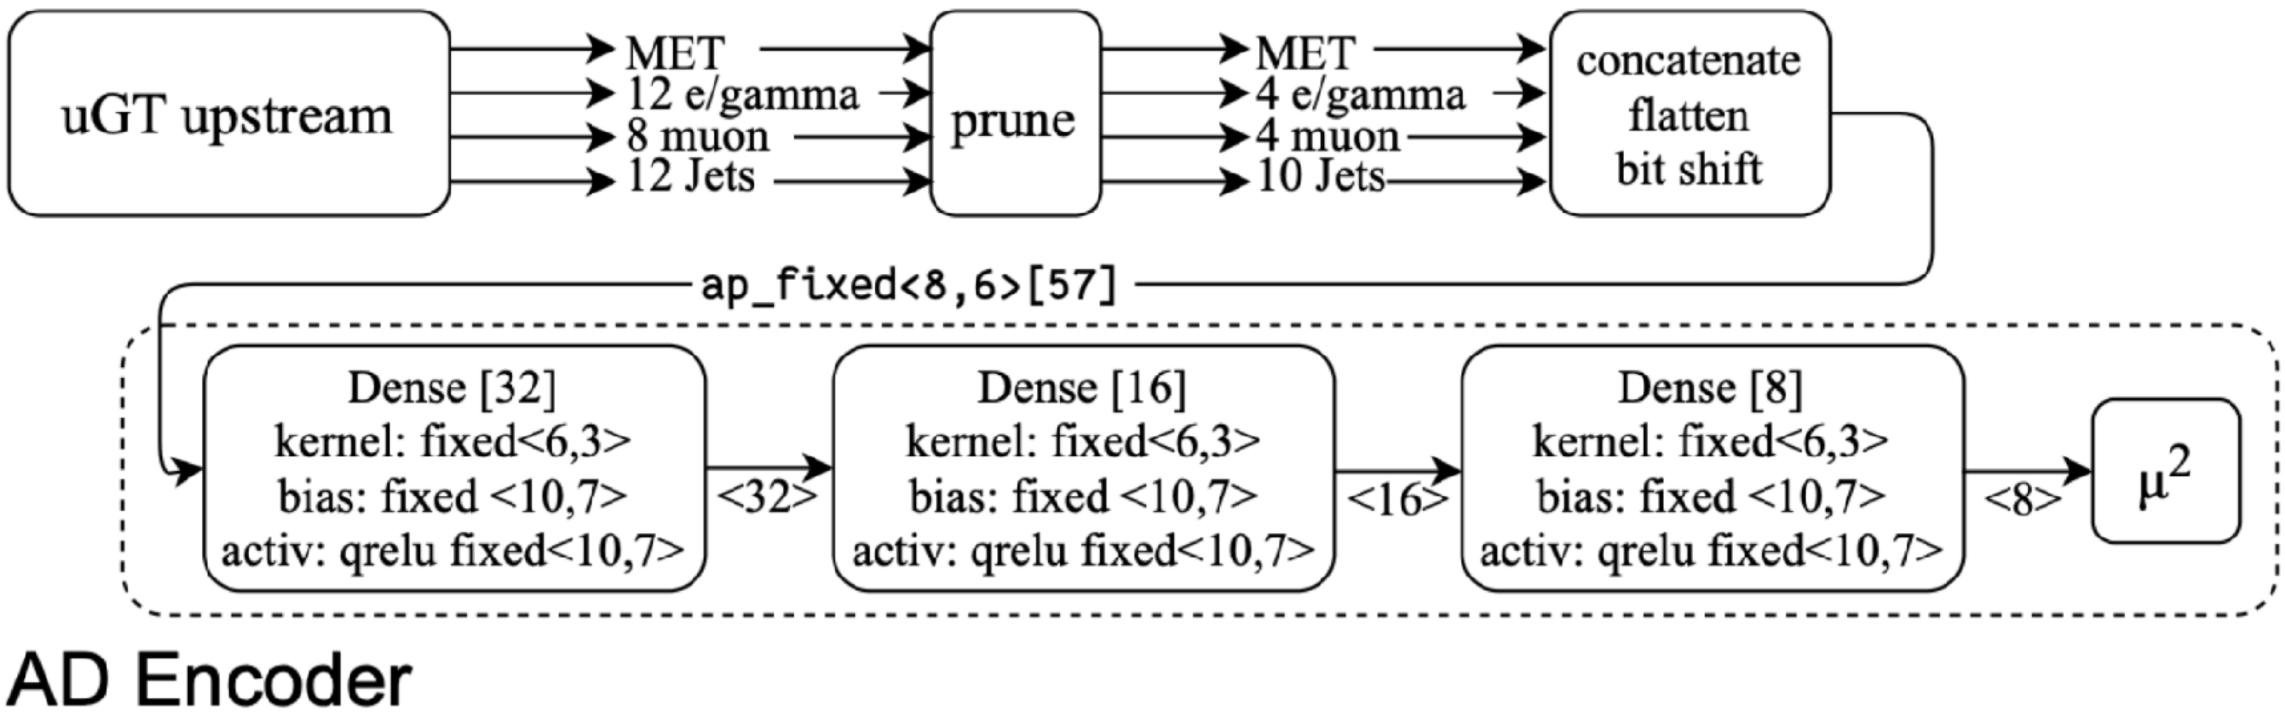

Run 3 µGT Trigger Design

- Preferred timing constraint of ~ 50 ns (Maximum ~ 125 ns)

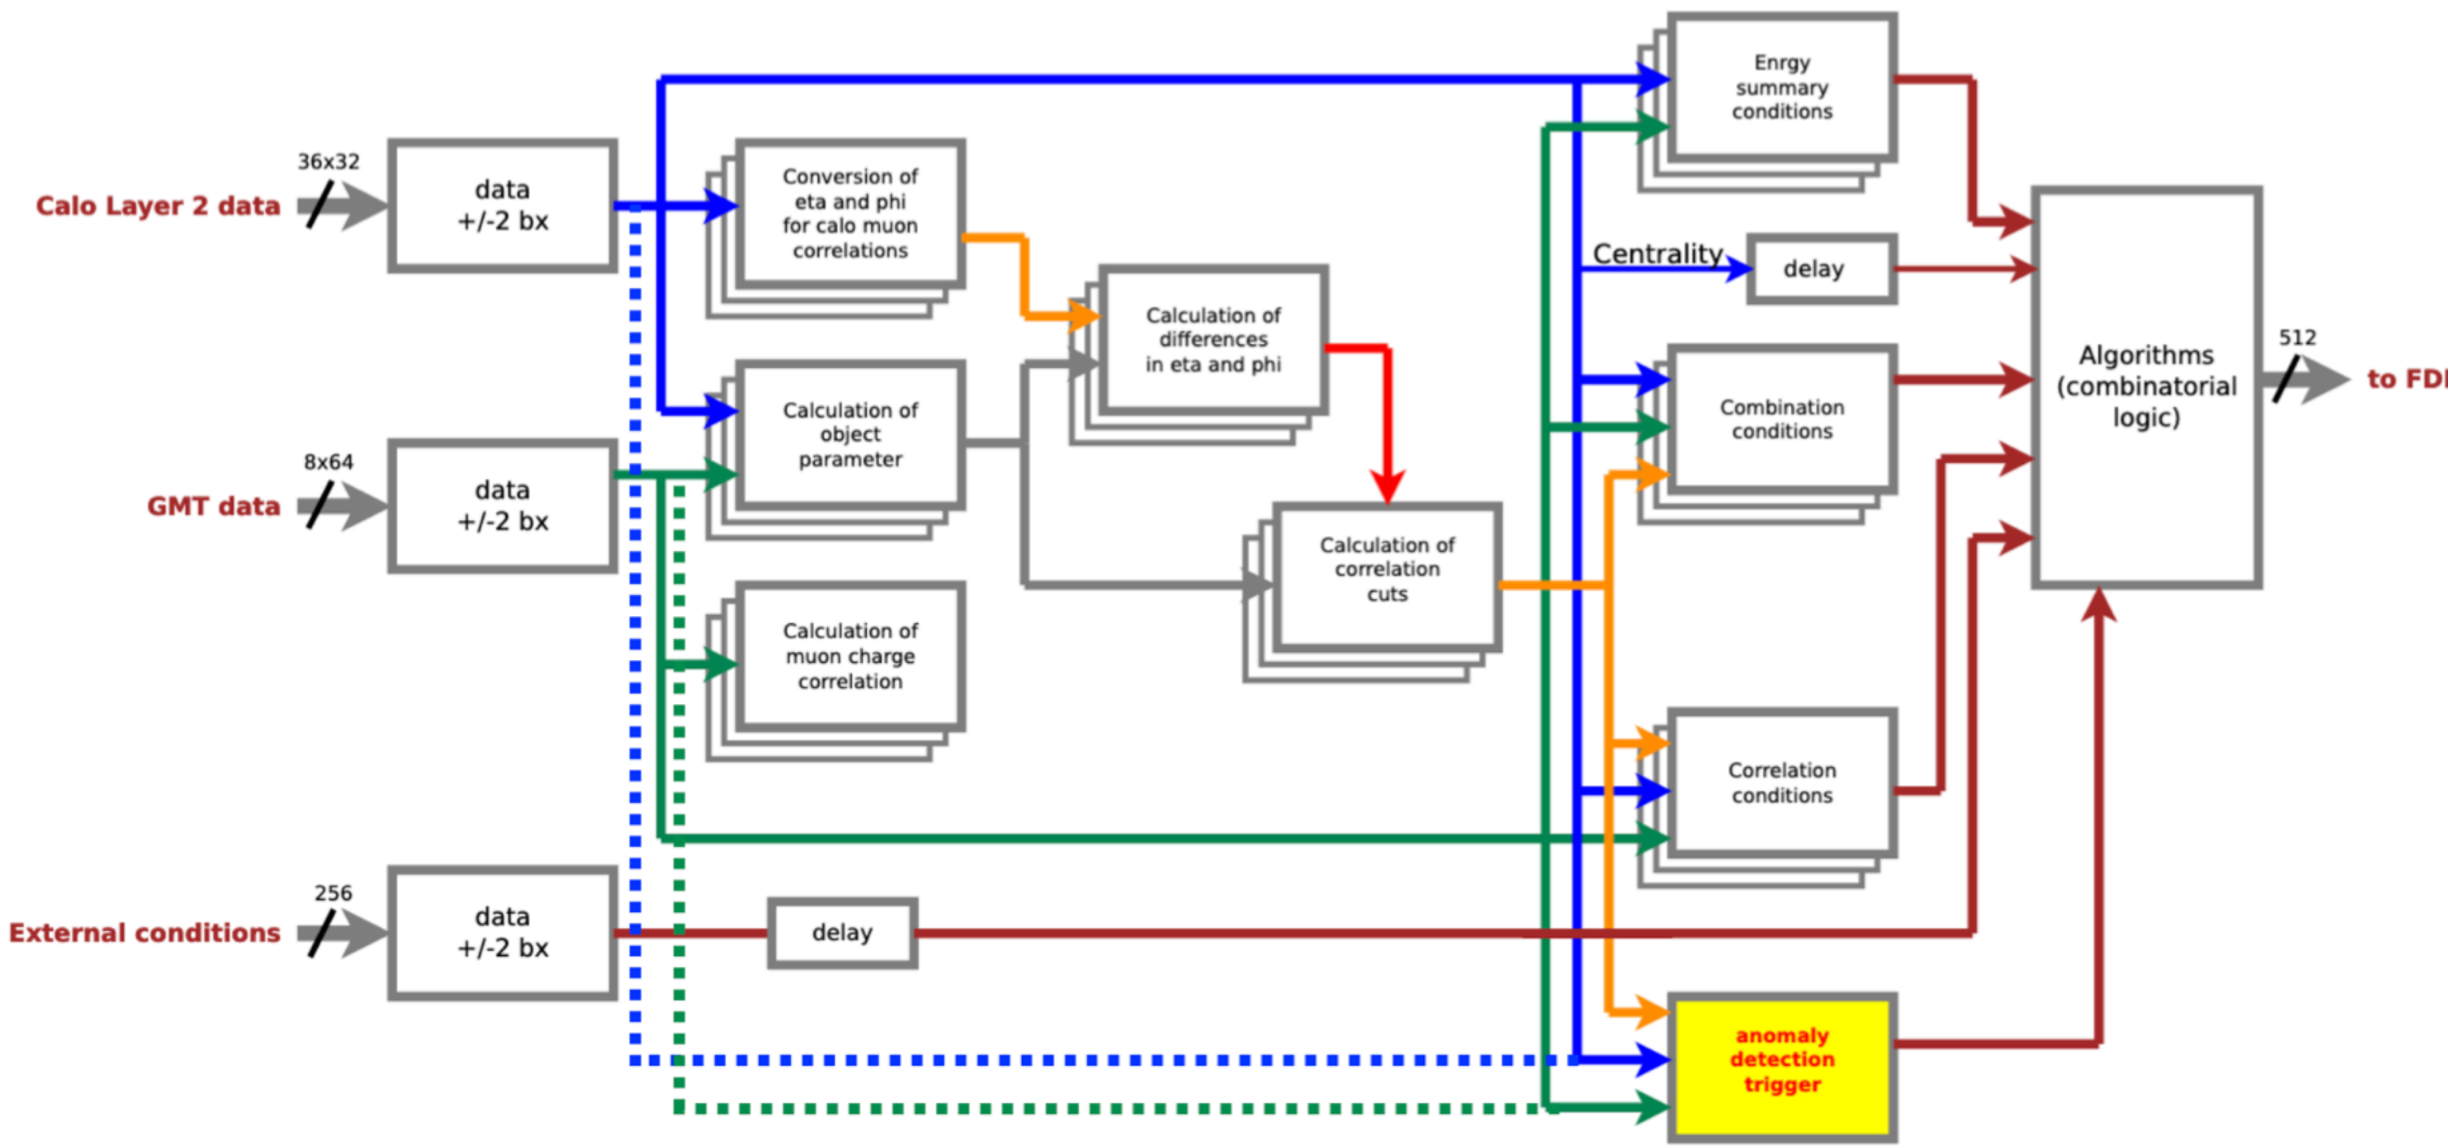

From Herbert Bergauer and Manfred Jeitler

$\sim 50ns$

Physics Performance

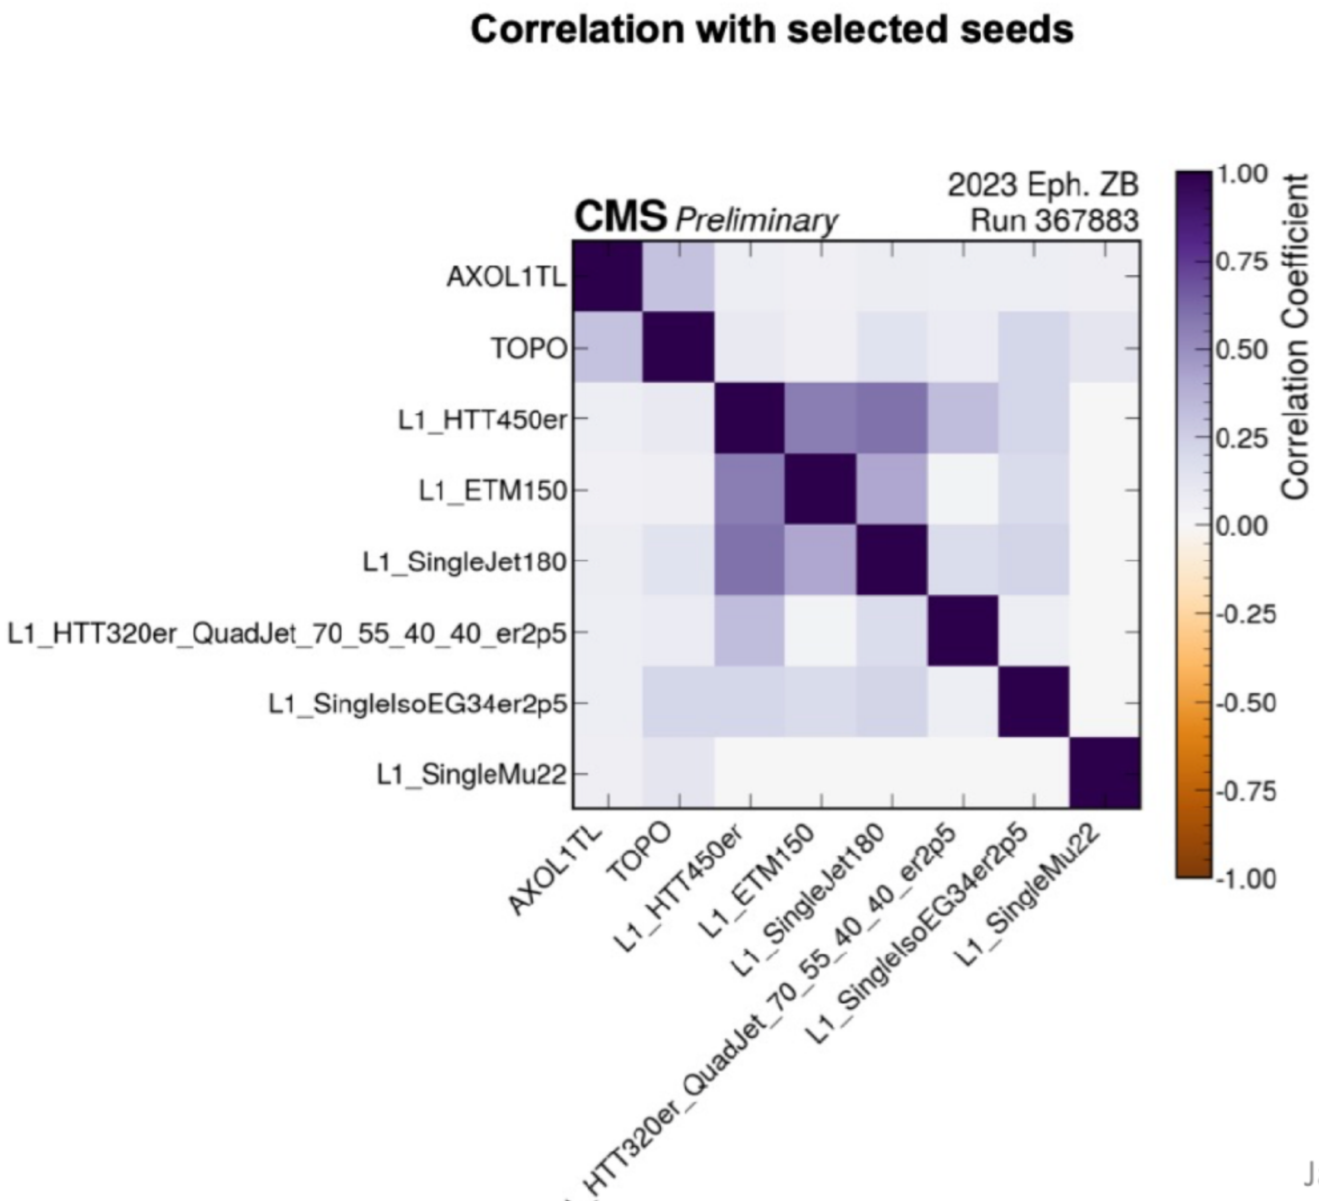

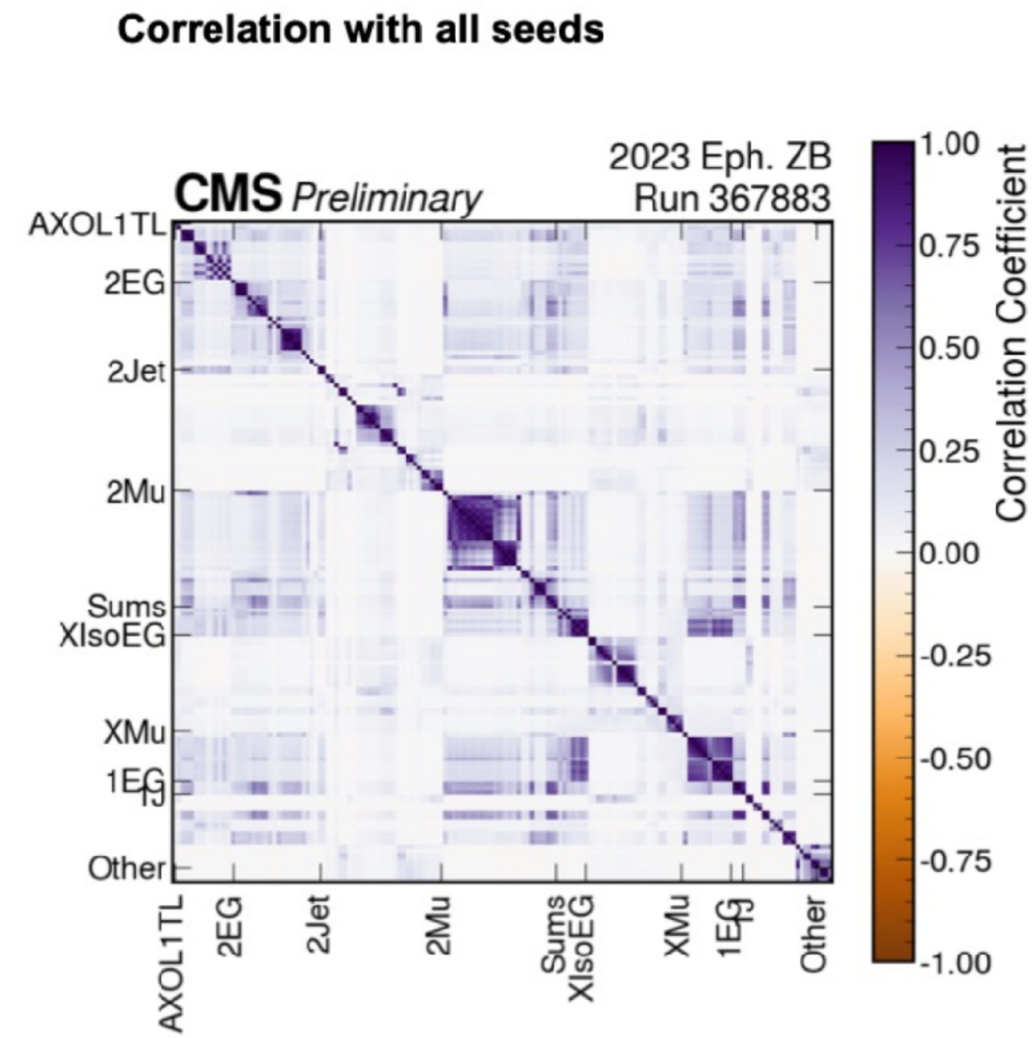

Jannicke Pearkes

- Correlations between continuous anomaly score (v2), continuous topo score, and L1 trigger bits on Zero Bias data
  - → **Low overlap with existing L1 seeds!**