



NexTGen
Next Generation Triggers



UNIVERSITÀ
DEGLI STUDI
DI PADOVA



Finanziato
dall'Unione europea
NextGenerationEU

REAL-TIME ANOMALY DETECTION IN THE CMS LEVEL-1 TRIGGER WITH AXOL1TL

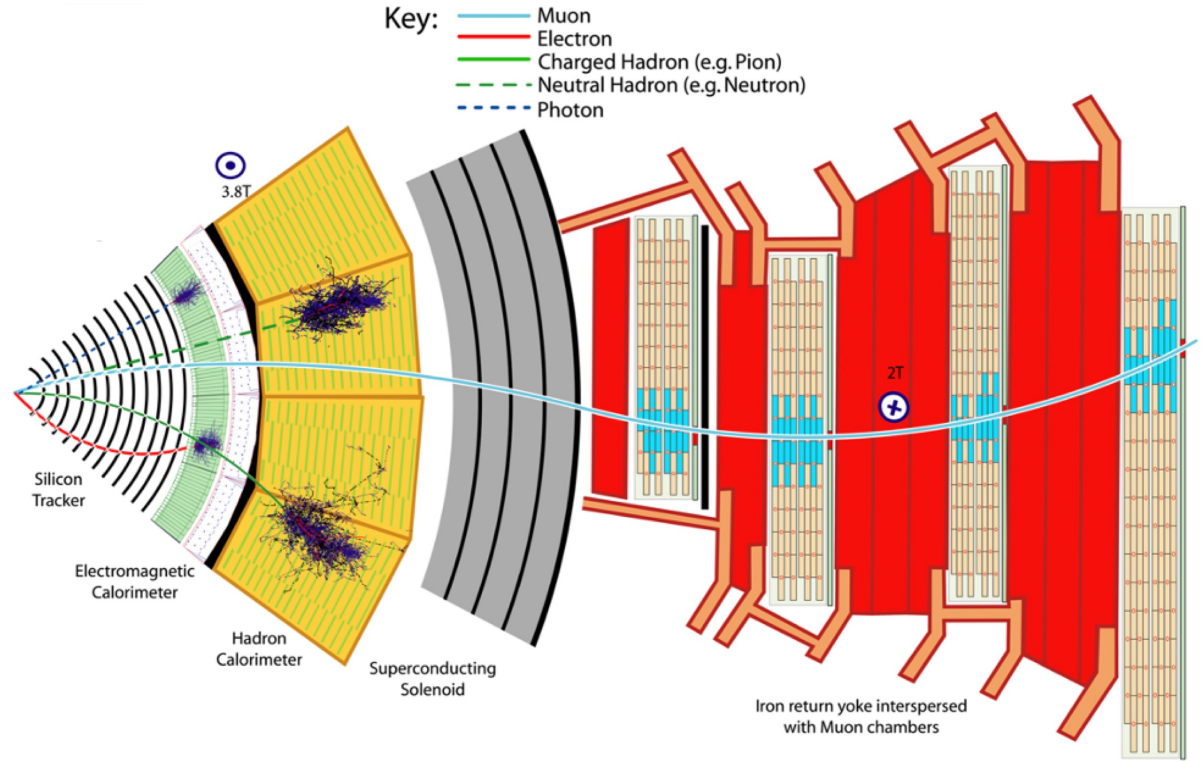

Fast Machine Learning for Science Conference 2025
ETH Zurich, September 1-5th 2025

Sabrina Giorgetti^{1,2,3,*} on behalf of the CMS collaboration

¹ Università Degli Studi di Padova , ² INFN sezione di Padova

³ CERN, *sabrina.giorgetti@cern.ch


THE CMS EXPERIMENT



THE CMS EXPERIMENT

TWO-TIER TRIGGER SYSTEM


- Proton-proton collisions at the LHC event's rate of 40 MHz.
- CMS raw data throughput ~ 30 TB/s: data reduction is needed.
- Two-tier trigger system:
 - Level-1 Trigger (L1T): hardware system
 - High-Level Trigger (HLT): software based
 - Event selection via configurable trigger menu (set of algorithms) based on reconstructed physics objects.



Why AD for triggering?

STANDARD TRIGGERING

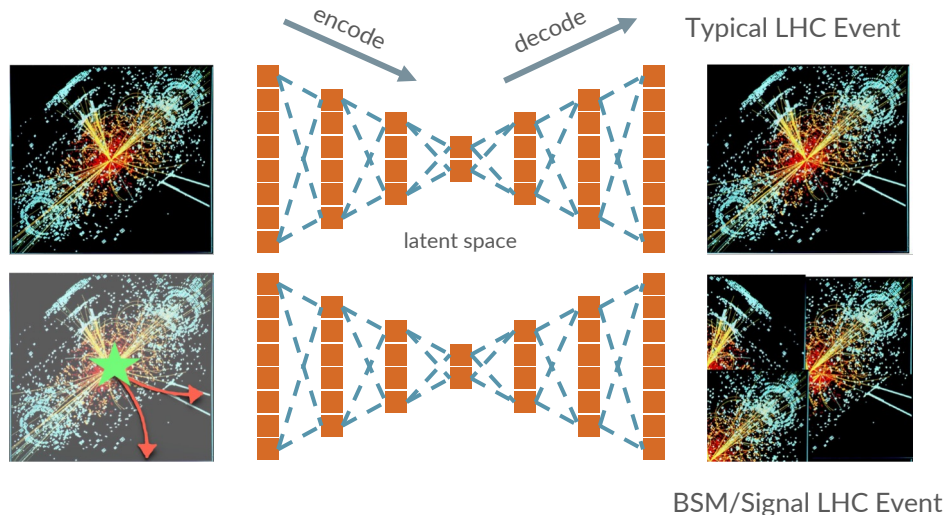
- Cuts on object kinematics (e.g. p_T, η, ϕ).
- Assumes known signatures.



WHAT IF NEW PHYSICS IT'S HIDDEN ELSEWHERE?

ANOMALY DETECTION




- Model-independent approach
- Look for the unknown!



Where to deploy AD in CMS ?


ANOMALY DETECTION IN CMS

- Over 99% of the events are rejected at each stage of the two-tier CMS trigger system.




DETECTOR COLLISIONS

40 MHz




LEVEL-1 TRIGGER

100 kHz



HIGH-LEVEL TRIGGER

O (1 kHz)






OFFLINE ANALYSIS

Where to deploy AD in CMS ?


ANOMALY DETECTION IN CMS

- Over 99% of the events are rejected at each stage of the two-tier CMS trigger system.

» Want to work at L1 before any selection is made!



40 MHz



100 kHz



O (1 kHz)



DETECTOR COLLISIONS

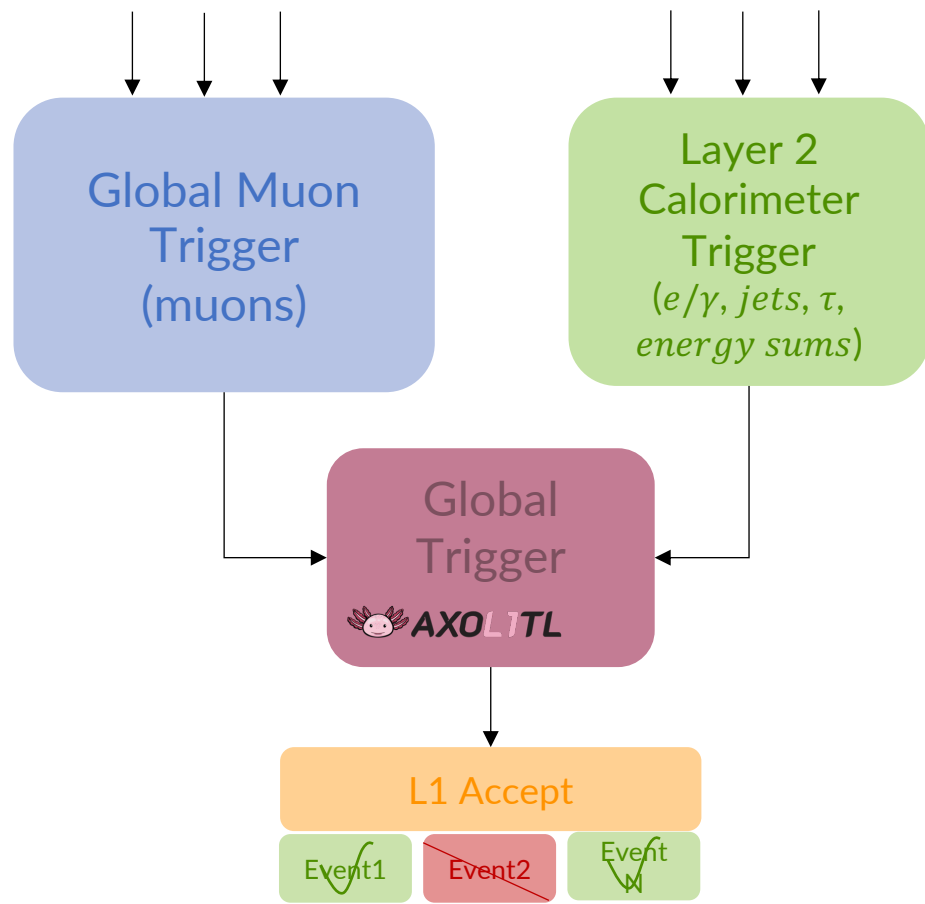
LEVEL-1 TRIGGER

HIGH-LEVEL TRIGGER


OFFLINE ANALYSIS

AD AT THE CMS L1T

- CMS L1T system performs a coarse reconstruction of physics objects.
- Within 4 μs , the L1T decides which events to accept.
- In the final stage, the Global Trigger (GT) runs algorithms with a **50 ns latency** → **ultra-low-latency inference**.
- Two AD algorithms in the GT:
 - AXOL1TL : an event-level trigger based on GT input objects.
 - CICADA : an event-level trigger based on Layer 1 calorimeter towers



AXOL1TL timeline




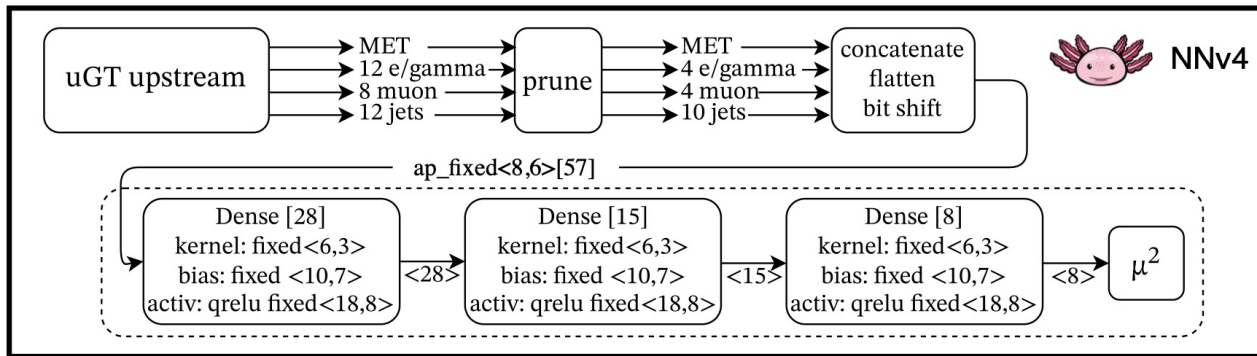
AXOL1TL

AXOL1TL (Anomaly Extraction Online L1 Trigger Lightweight) is an unsupervised variational autoencoder (VAE):

- Quantum-aware training on Zero Bias (unbiased trigger) dataset with inputs L1T objects:
(p_T, η, ϕ) of 4 e/ γ , 4 μ , and 10 jets and missing transverse energy (MET or p_T^{miss})

AXOL1TL V4

- Trained with a VAE loss combining a reconstruction term and the KL divergence.
- Deploy only the encoder, due to resource constraints.
- Anomaly score per event defined as the L2 norm of the 8-dimensional latent vector.

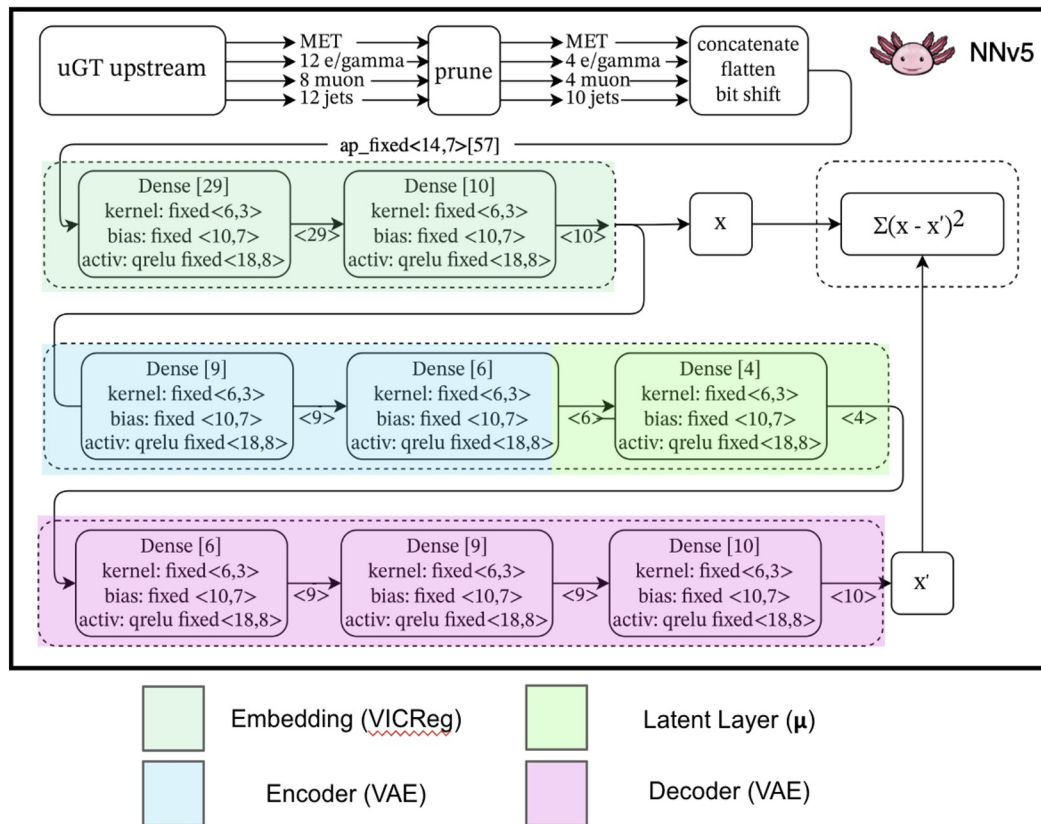


AXOL1TL V5

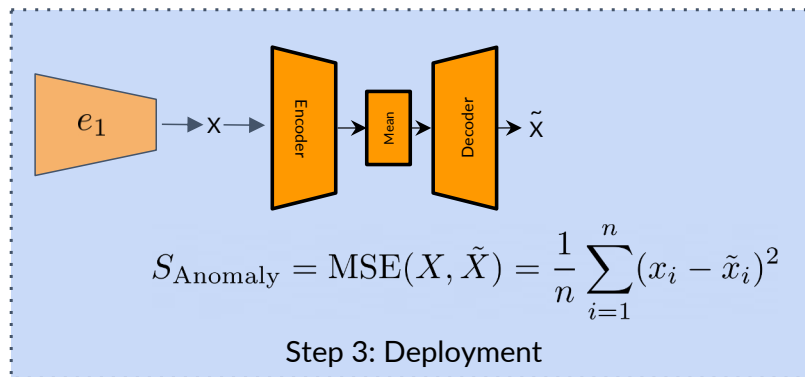
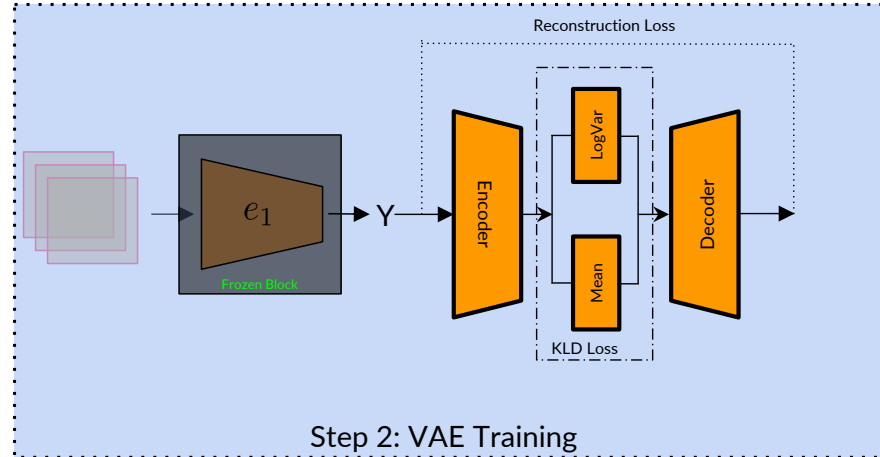

BEYOND BASELINE

For 2025 data-taking, a novel model architecture has been develop based on **contrastive learning**:


- **VICReg-trained** feature extractor stacked on top of a VAE.
- The **anomaly score** is now the **squared reconstruction error**.



AXOL1TL V5




AXOL1TL on FPGA



- AXOL1TL fits within the latency and resources budget of the CMS GT boards Xilinx Virtex-7 FPGA.
- The AXOL1TL QKeras model is converted for FPGA using [hls4ml](#).
- Resource usage of V5 is comparable with V4 thanks to the optimizations with distributed arithmetic ([da4ml](#)) for V5.
 - » Check out Chang Sun's [tutorial](#) & [talk](#) for FASTML25!

Model	Latency	LUTs (% w.r.t GT)	FFs (% w.r.t GT)	DSPs
AXOL1TL V4	2 clocks, 50 ns	18961 (4.3%)	424 (0.05%)	0
AXOL1TL V5 (w/ da4ml)	2 clocks, 50 ns	22023 (5.0%)	627 (0.07%)	0
AXOL1TL V5 (w/o da4ml)	2 clocks, 50 ns	30583 (7.1%)	965 (0.11%)	0
GT	160 clocks, 4 us	433,200	866,400	-


Physics performance




- Scan of AXOL1TL anomaly score thresholds, corresponding to a total range, for SM and BSM simulated samples and calculate the efficiency.
- **V5 efficiency improved** compared to V4 and competitive with a standard trigger like H_T (scalar sum of jets p_T).
- Broad **sensitivity** of the unsupervised AXOL1TL approach **to a wide range of new physics scenarios**.

AXOL1TL anomaly score

- Set a threshold on the AXOL1TL anomaly score: define several L1 triggers, then passthrough to the CMS HLT Scouting system where objects are reconstructed online.
- Pure AXOL1TL events: events uniquely triggered by AXOL1TL and not by any other L1 trigger.
 - » Collect events that would not be triggered otherwise!







AXOL1TL V4 in 2024
Pure rate 200 Hz. Total rate 1.5 kHz



AXOL1TL V5 in 2025
Pure rate 250 Hz. Total rate 5.9 kHz

AXOL1TL results


- Pure AXOL1TL events in the energy sums spectrum: enhanced sensitivity to lower L1 H_T and p_T^{miss} values.
- AXOL1TL selects events with higher object multiplicity, particularly for jets and e/γ , than those recorded by the Zero Bias trigger (reduced rate unfiltered trigger).



AXOL1TL results

Use SM candles as validation of the AXOL1TL trigger.

AXOL1TL sees the $Z(\mu\mu)$ boson as well as other resonances in the di-muon spectrum!



Conclusions

SUMMARY

- Presented the first anomaly-detection-based trigger implemented in the CMS L1T.
 - Successful AXOL1TL data-taking in 2024 and 2025 with first analysis results!
- Showcased the latest model development for 2025 with improved performance.

WHAT'S NEXT

- Push the effort on end-to-end physics analyses using AXOL1TL data.
- Drive the development of anomaly detection for the Phase-2 CMS Level-1 Trigger at HL-LHC.
- AXOL1TL as a use case for dedicated MLOps infrastructure.
 - >> Join the *“MLOps Pipeline for Continuous Deployment of Machine Learning Algorithms for HEP”* session tomorrow!






CMS Experiment at the LHC, CERN

Data recorded: 2023-May-24 01:42:17.826112 GMT

Run / Event / LS: 367883 / 374187302 / 159



THANK YOU!

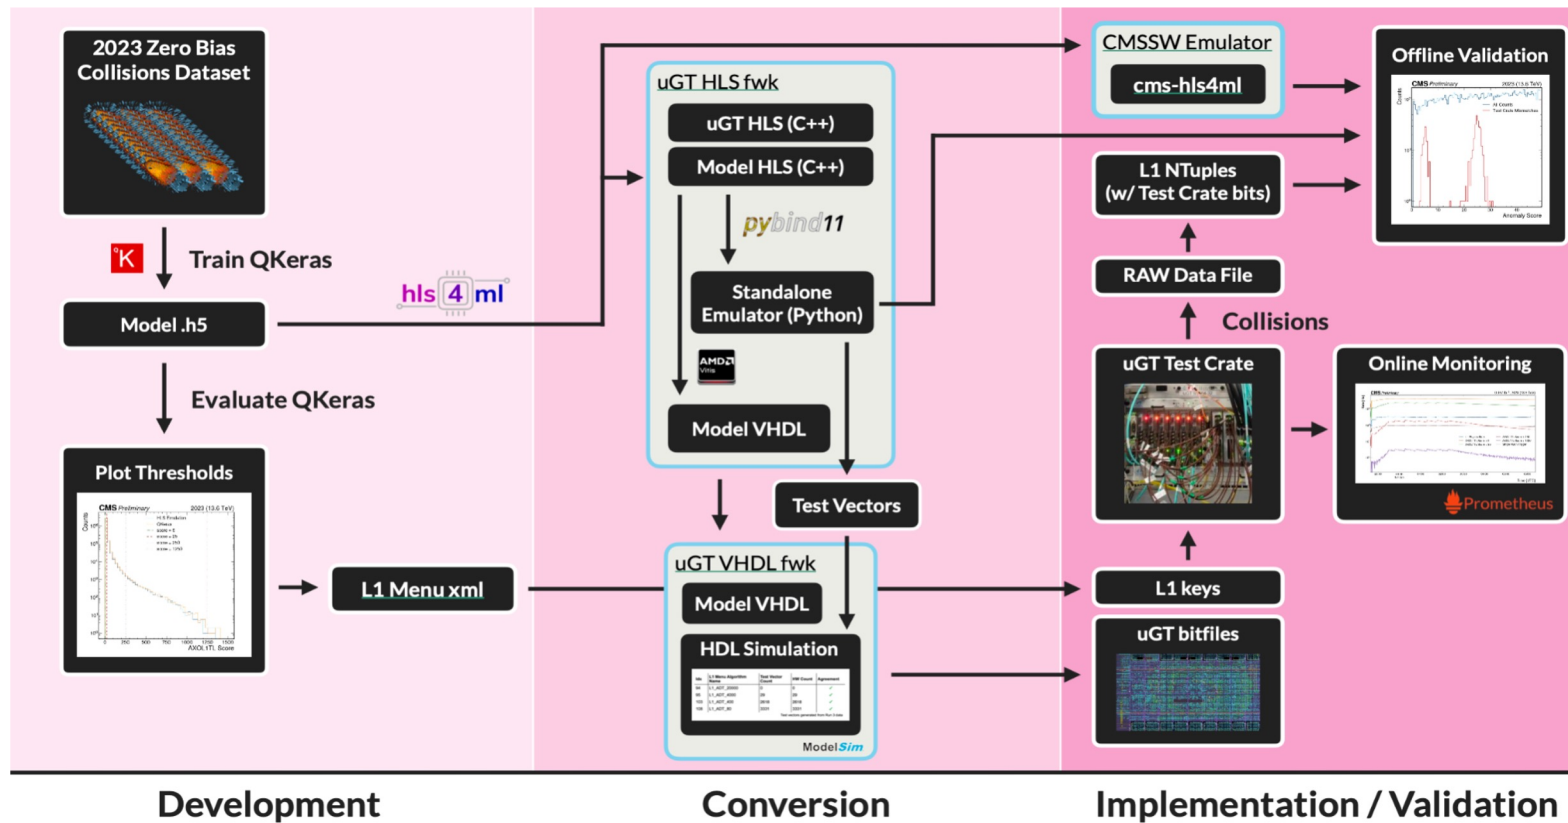


BACKUP


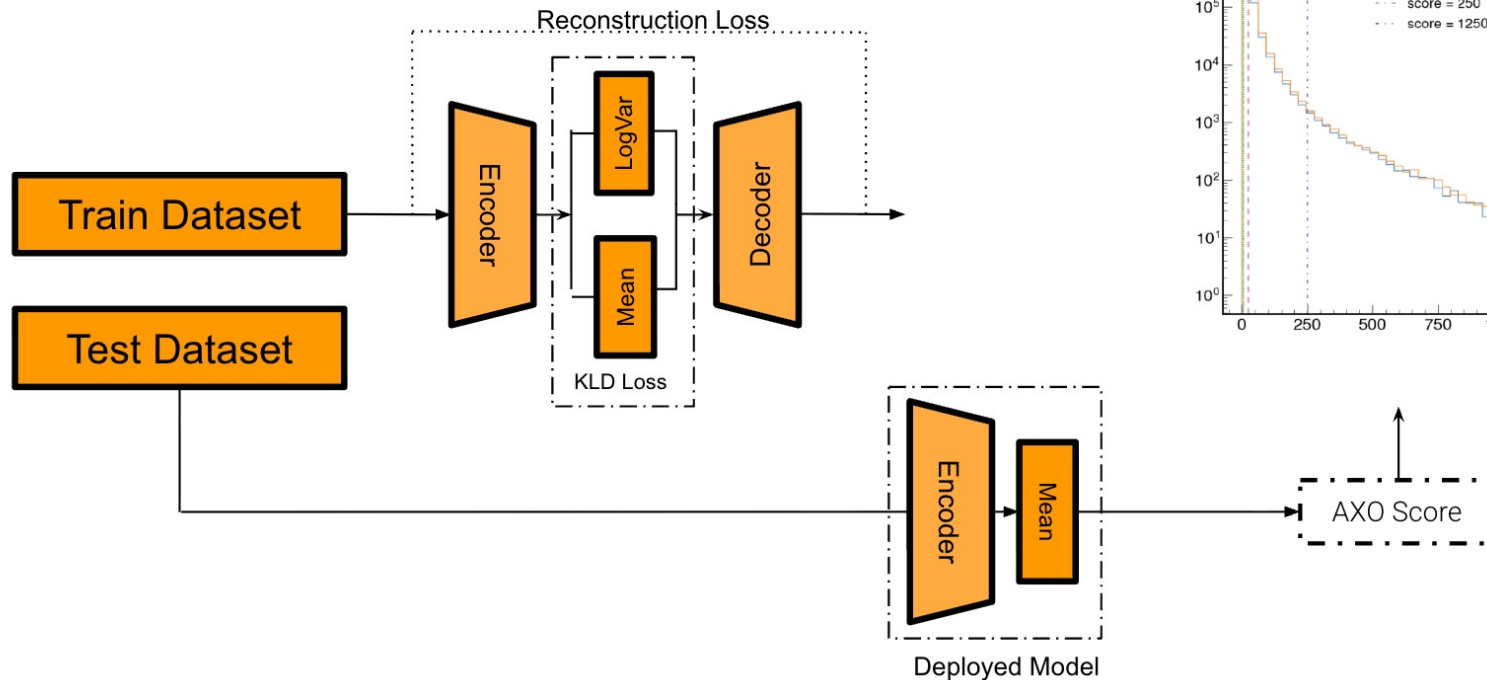
References

- [1] CMS Collaboration. (2008). *The CMS experiment at the CERN LHC*. JINST, 3, S08004. <https://doi.org/10.1088/1748-0221/3/08/S08004>
- [2] Jeitler, M., et al. (2007). *The Level-1 Global Trigger for the CMS Experiment at LHC*. JINST, 2, P01006. <https://doi.org/10.1088/1748-0221/2/01/P01006>
- [3] CMS Collaboration. (2023). *Anomaly Detection in the CMS Global Trigger Test Crate for Run 3*. CMS Detector Performance Summary, CMS-DP-2023-079. <https://cds.cern.ch/record/2876546>
- [4] CMS Collaboration. (2024). *2024 Data Collected with AXOL1TL Anomaly Detection at the CMS Level-1 Trigger*. CMS Detector Performance Summary, CMS-DP-2024-059. <https://cds.cern.ch/record/2904695>
- [5] Bardes, A., Ponce, J., & LeCun, Y. (2022). VICReg: Variance-Invariance-Covariance Regularization for Self-Supervised Learning. arXiv preprint arXiv:2105.04906. <https://arxiv.org/abs/2105.04906>
- [6] Kingma, D. P., & Welling, M. (2022). Auto-Encoding Variational Bayes. arXiv preprint arXiv:1312.6114. <https://arxiv.org/abs/1312.6114>
- [7] Duarte, J., et al. (2018). *Fast Inference of Deep Neural Networks in FPGAs for Particle Physics*. JINST, 13, P07027. <https://doi.org/10.1088/1748-0221/13/07/P07027>
- [8] Sun, C., Que, Z., Loncar, V., Luk, W., & Spiropulu, M. (2025). *da4ml: Distributed Arithmetic for Real-time Neural Networks on FPGAs*. arXiv:2507.04535. <https://arxiv.org/abs/2507.04535>
- [9] CMS Collaboration. (2024). *Enriching the Physics Program of the CMS Experiment via Data Scouting and Data Parking*. <https://arxiv.org/abs/2403.16134>
- [10] CMS Collaboration. (2025). *HLT Scouting performance in 2024: muon triggers and objects*. CMS Detector Performance Summary, CMS-DP-2025-025. <https://cds.cern.ch/record/2935659>
- [11] Noah Zipper, *Realtime Anomaly Detection in the CMS Experiment*, Fast Machine Learning for Science 2024.
- [12] Level-1 Trigger Public Performance Results, <https://twiki.cern.ch/twiki/bin/view/CMSPublic/L1TriggerDPGResults>

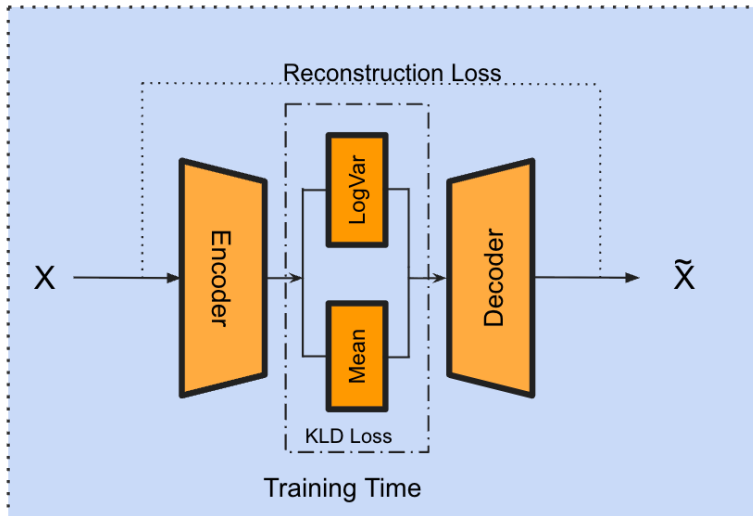
AXOL1TL Workflow



AXOL1TL V3 and V4

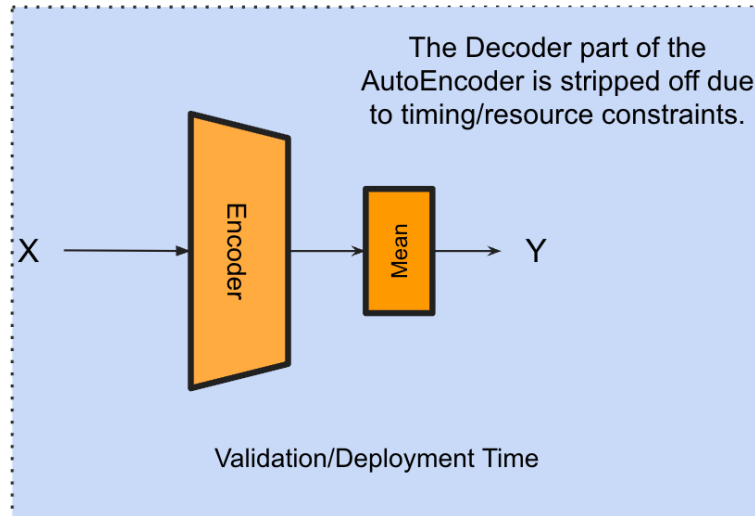


AXOL1TL V3 and V4



$$\epsilon \sim \mathcal{N}(0, 1)$$

$$z = \text{Mean} + \exp(0.5 \cdot \log \text{Var}) \cdot \epsilon$$





$$S_{\text{Anomaly}} = \frac{1}{n} \sum_{i=1}^n (Y_i)^2$$

Inputs: L1 trigger objects: (pT, η , ϕ) of MET, 4 electron/photons, 4 muons, 10 jets

AXOL1TL Validation

V5 & V4



V3

