CLARIPHY AI COLLABORATION MEETING · UNIVERSITY OF WISCONSIN-MADISON

AIDA-Scout

AI-Inclusive Discovery of Anomalies in Scouting — model-independent anomaly detection at 40 MHz in the CMS Level-1 Scouting system

Keith Ulmer[1], Natalie Bruhwiler[1], **Abhijith Gandrakota**[2], Oz Amram[2], Jennifer Ngadiuba[2], Javier Duarte[3], Daniel Diaz[3], Melissa Quinnan[3], Ellison Scheuller[3], Petar Maksimovic[4], Xinlong Liu[4]

[1] University of Colorado Boulder · [2] Fermilab · [3] University of California San Diego · [4] Johns Hopkins University · with AMD

SUMMARY

The CMS trigger keeps roughly 1 in 10^4 collisions, selected by a menu designed around known signatures. Level-1 Scouting reads out every bunch crossing at 40 MHz with reduced-content objects. AIDA-Scout is an unsupervised analysis pipeline for that stream: two decorrelated anomaly scores, an ABCD background estimate from data, and an interpretability step that maps selected anomalies to new L1 seeds for full-readout follow-up.

Motivation

- L1 + HLT reduce 40 MHz to ~10 kHz. Rejected events are lost; there is no offline recovery.
- The menu is optimised for benchmark BSM models. Signatures outside them may be rejected online and never reach an analysis.
- L1 Scouting records every crossing (40 MHz, ~750 GB/s) with coarse L1 objects. The analysis must run online at full rate.
- AIDA-Scout: unsupervised anomaly detection, data-driven background estimation and interpretation, running online on the scouting stream.

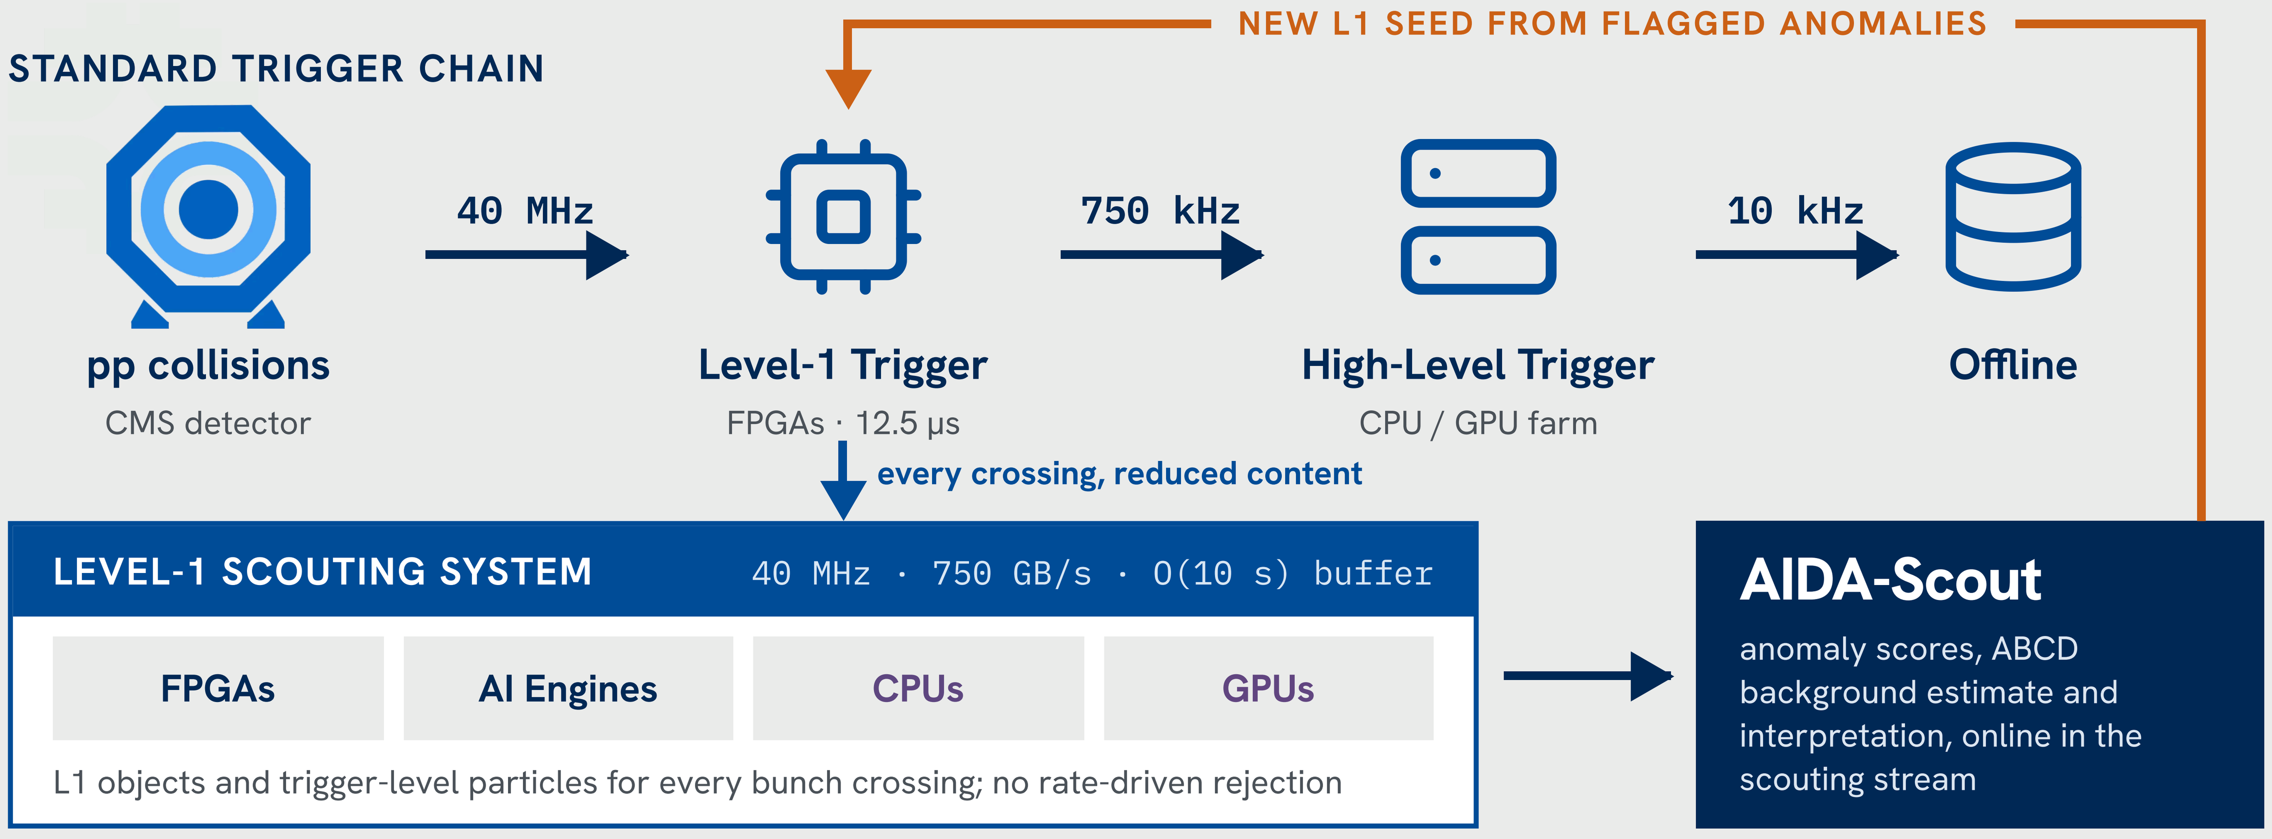

Method

1 Representation learning

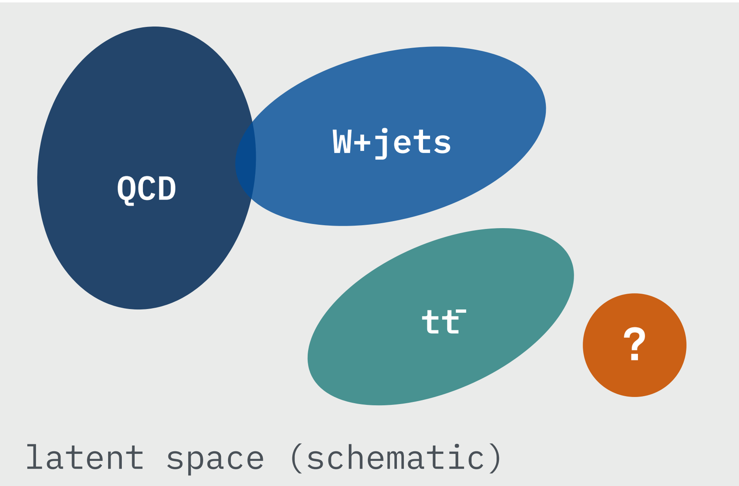

- Supervised contrastive encoder (SimCLR / VICReg) on trigger-level particles, trained on SM MC
- Standard Model (SM) processes cluster in distinctive regions, while Beyond the BSM physics stands as outliers.

2 Two views, two scores

AD 1 · OBJECT LEVEL
Autoencoder on L1 jets, e, μ, γ, E_T^{miss}. Score: reconstruction loss.

SAME EVENT

AD 2 · PARTICLE LEVEL
Step-1 encoder on trigger-level particles. Score: distance to SM clusters.

- Two views of the same event at particle and object level, two complementary scores
- Signal region: displays anomalous behavior, with both scores above the threshold

3 Decorrelation and ABCD

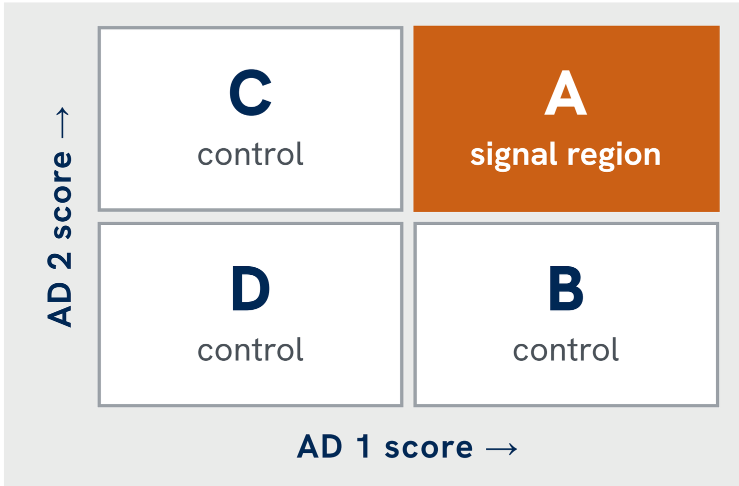

$$N_A^{bkg} = N_B \cdot N_C \,/\, N_D$$

- Independence established with nuisance-randomised distillation
- Background from data, requires very high closure at optimal selection

4 Interpretability and follow-up

Anomalous events — region A
↓
Attribution — what drives the score
↓
New L1 seed — for that signature
↓
Full-readout data, offline analysis

- Attribution methods inspired from the CMS anomalous-jet searches
- Interpretability framework to design traditional L1 Triggers to record anomalies for further follow up

Online Deployment

- Benchmark each stage (encoder, both ADs, scoring) for per-event latency, memory and sustained throughput
- Targets: FPGAs via `hls4ml`, AMD AI Engines via `AIE4ML`, and CPU/GPU nodes representative of the CMS L1 Scouting compute farm
- Map each stage to firmware or the heterogeneous farm, balancing model size against physics performance

Datasets

Collide-2V: 10^9 simulated pp events at HL-LHC conditions, each with an L1 view and an offline view. Used to train the contrastive encoder.

- CMS Open Data and simulation (> 2 PB): minimum-bias for rate studies, benchmark signals for efficiency
- Benchmark signals: SUEPs, RPV SUSY, exotic Higgs decays
- Input formats from the L1 Scouting demonstrator

First Results PRELIMINARY

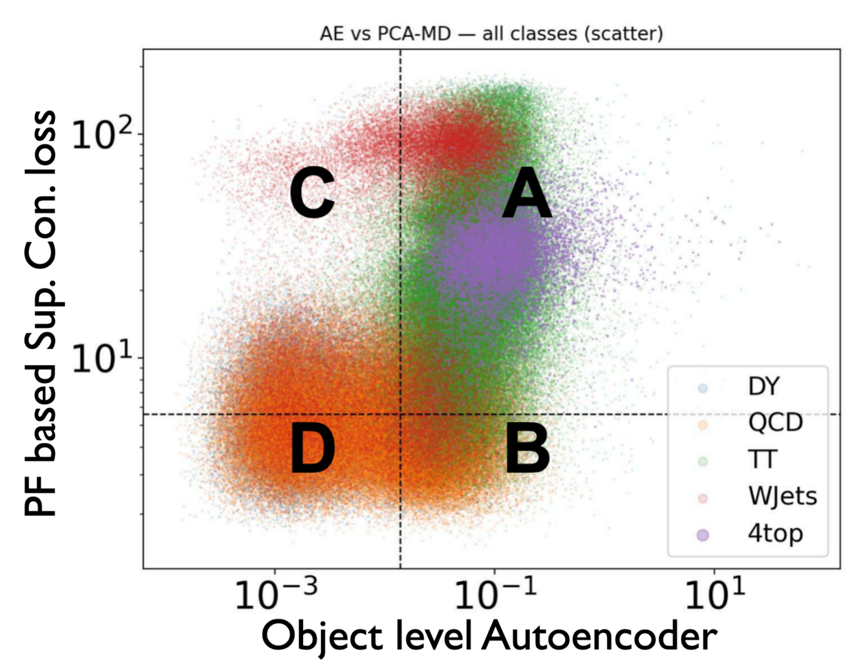

ABCD plane. AD 1 (object-level AE loss) vs. AD 2 (PF supervised-contrastive loss), SM MC. QCD and DY in D, W+jets in C, tt across B and A; rare tttt populates A as a BSM-like signature would.

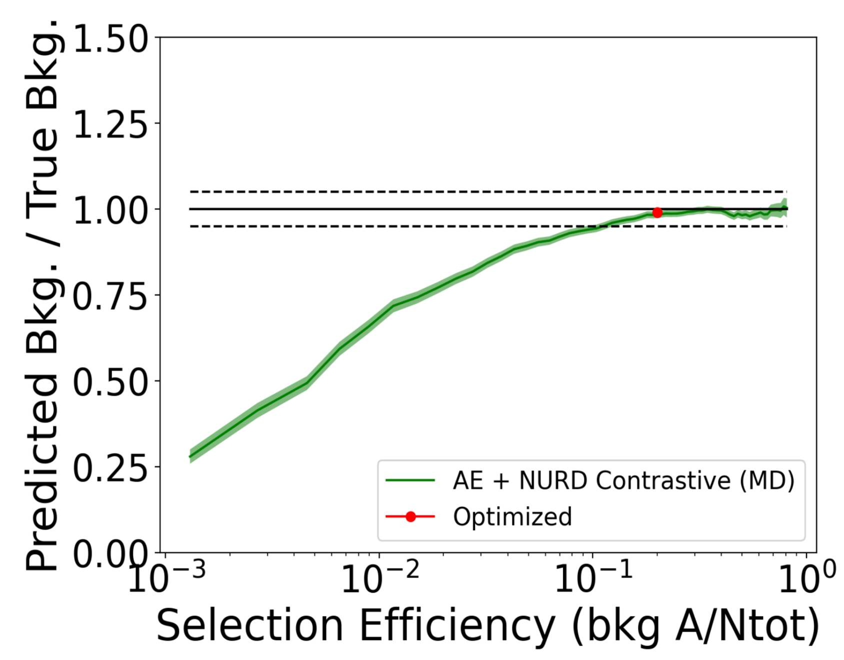

Closure. Predicted / true background in A vs. selection efficiency, AE + NuRD contrastive. Within 5% (dashed) down to ~10% efficiency; the optimised point is marked. Closure degrades at tighter cuts.

06 Milestones

M1 · MONTH 3
Data infrastructure
Scouting-format samples flowing end-to-end

M2 · MONTH 5
Representation learning
Encoder separates SM processes; gains on BSM benchmarks

M3 · MONTH 6
Background estimation
Decorrelated scores; ABCD closes in MC

M4 · MONTH 7
Online compatibility
Each stage mapped to hardware within the scouting budget

M5 · MONTH 9
End-to-end prototype
Scouting input → anomaly candidates with background estimate

M6 · MONTH 9
Interpretability
Attribution turns anomalies into a new L1 seed

FERMILAB-POSTER-26-0283-CMS

Supported by the U.S. Department of Energy Genesis Mission, DE-FOA-0003612 Phase I, Challenge 14-C: Unifying Physics from Quarks to the Cosmos. Contact: abhijith@fnal.gov