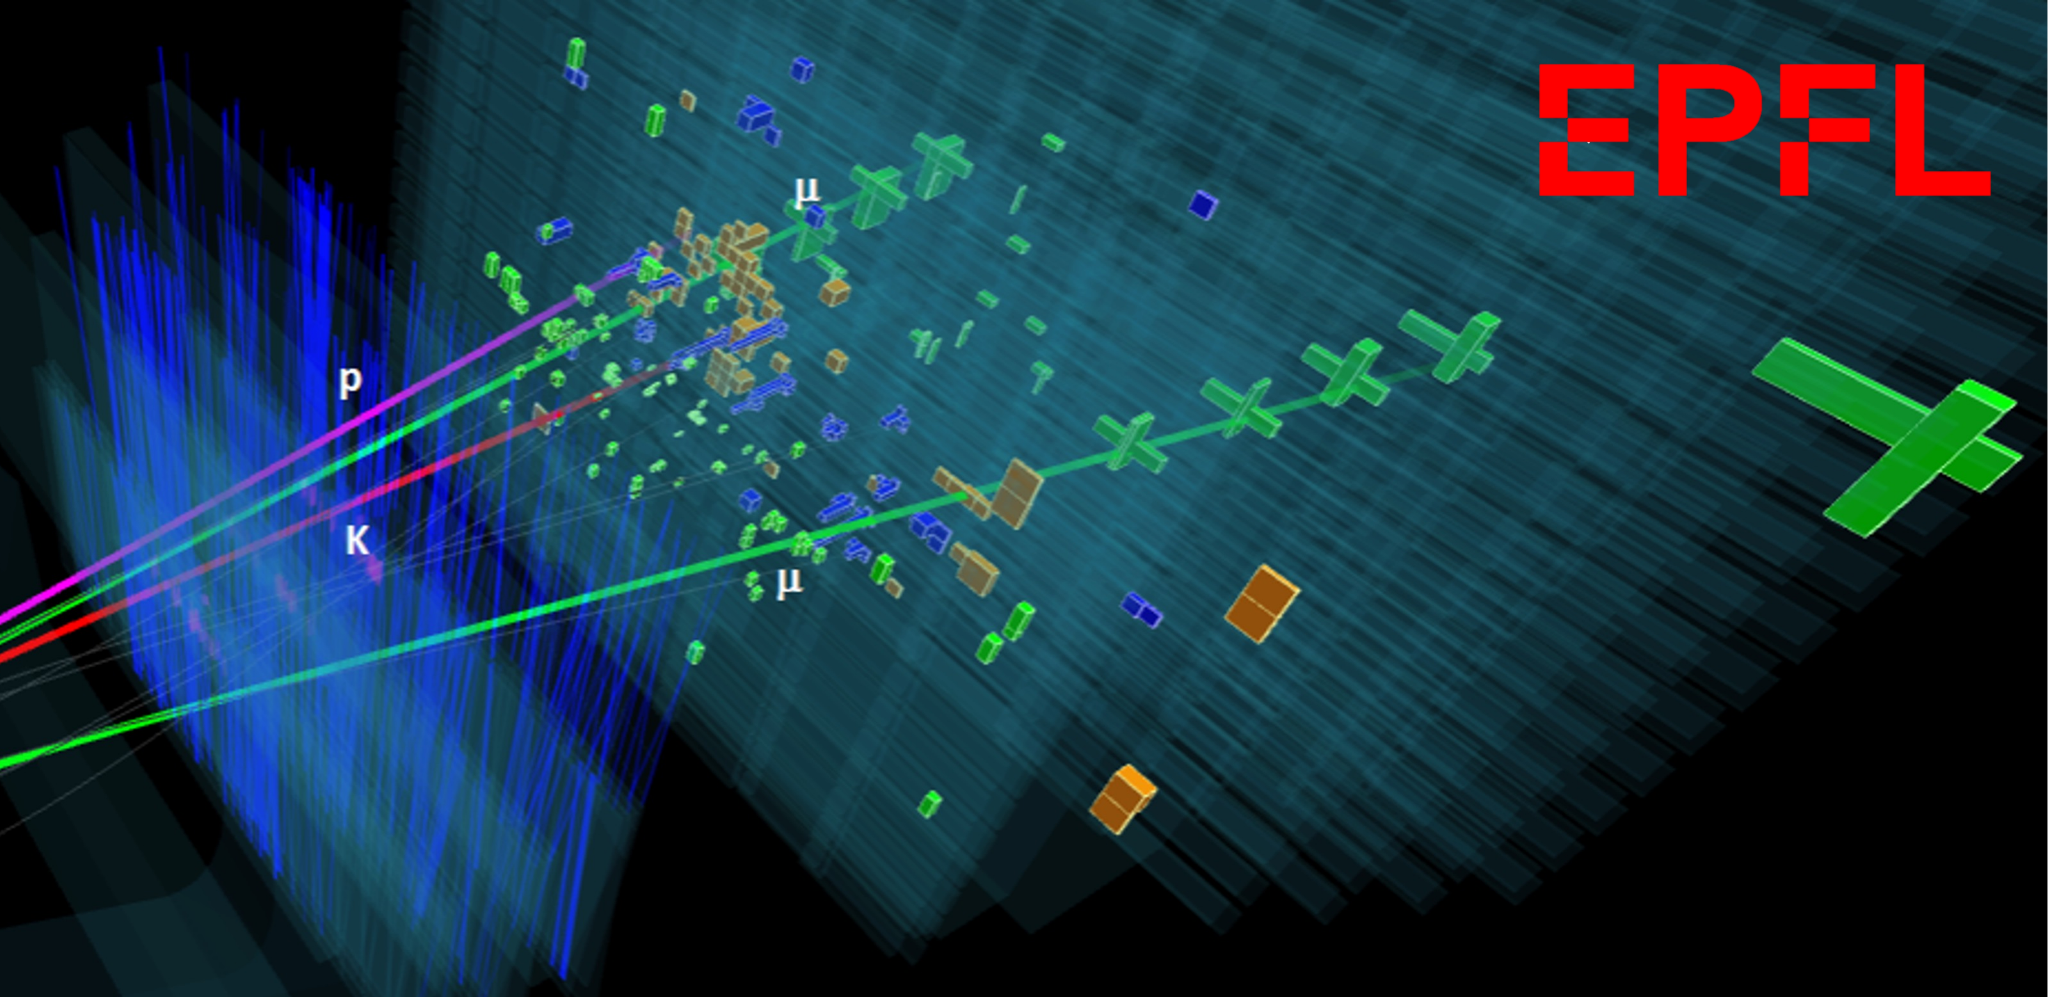

Anomaly detection at LHCb

Luca Hartman, on behalf of the LHCb collaboration

Two applications under development

In the trigger to select showers in the muon detector

LHCb-FIGURE-2024-015

- Making use of the full software trigger, running on GPU
- Remaining maximally model independent for various BSM models
- Reducing the number of trigger lines
- Reducing the difference between MC and data

In the control room to flag anomalous runs

arXiv:2405.15508

- Reducing the workload of the shifters
- Increasing the quality of the data

The LHCb detector

- **Forward** spectrometer for beauty and charm physics
  - Unique phase space region ($2 < \eta < 5$)
  - **Complementary** to ATLAS and CMS
- Designed for
  - High precision tracking and vertexing
    - Vertex locator, multiple tracking stations + magnet
  - Excellent particle identification
    - Two RICH detectors, EM- and HAD- calorimeters
    - Muon detector
- Used for BSM searches

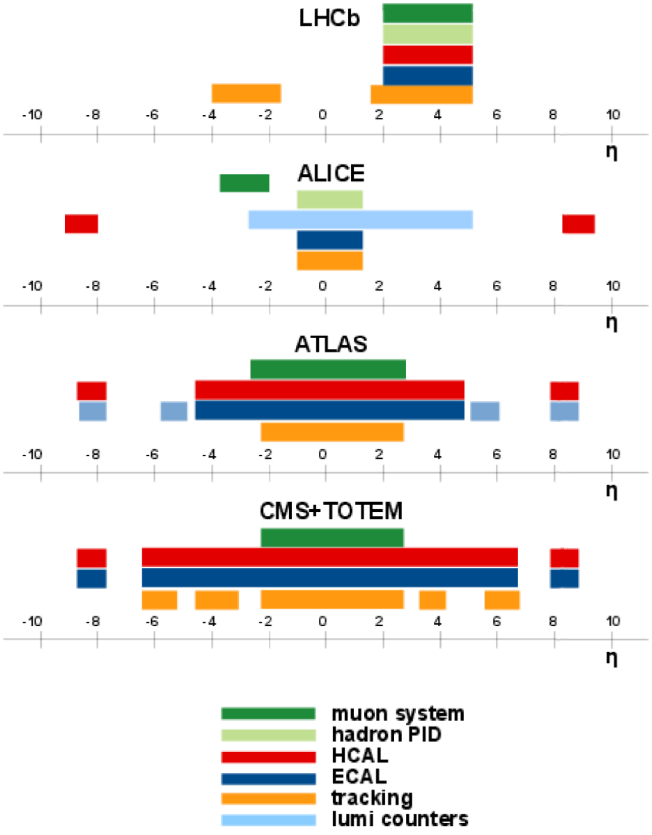

[LHCb results from proton-ion collisions](), L. M. Massacrier, 2015, 45th International Symposium on Multiparticle Dynamics

Unique detector signatures

- **Feebly interacting particles** appear in many BSM scenarios
  - Heavy Neutral Leptons (HNL) [1]
  - Axion like particles [2]
- Long lifetimes lead to **unique challenges and opportunities**
  - LLPs could decay beyond tracking stations
- We can use the **muon system** as a **sampling calorimeter**
  - Very rare signature in the SM
  - Similar searches by ATLAS [3] and CMS [4, 5]
  - Accepted/proposed dedicated future experiments [SHiP, MATHUSLA, and others]
  - LHCb could contribute in a short timescale [6]

LHCb muon detector

- Four multiwire proportional chambers (M2-M5)
- Three iron layers of each $4.8\lambda_I$ (80 cm of iron)
- Large decay volume
- But **not** designed for shower detection
  - No energy deposit measurements

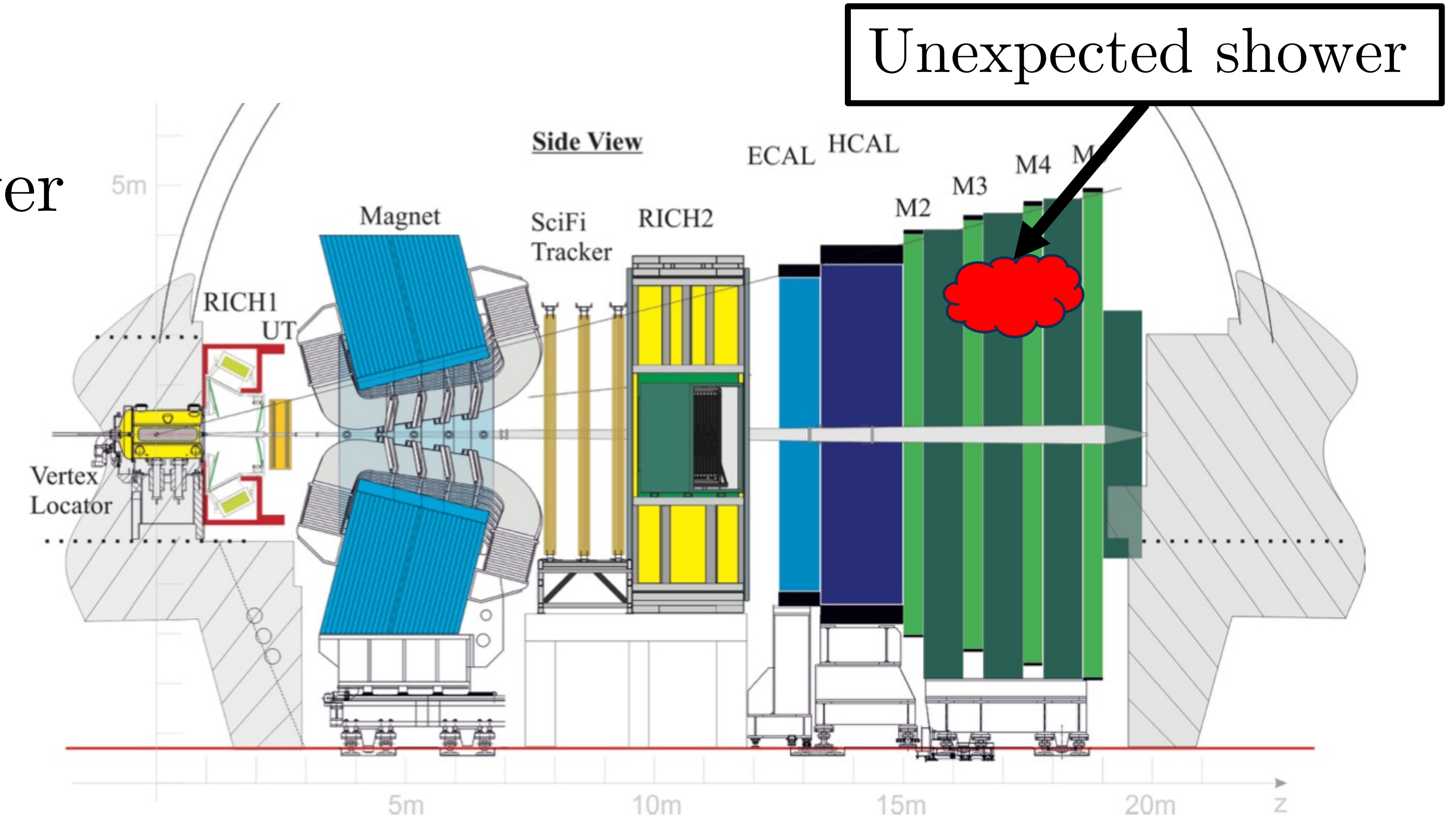

[LHCb reoptimized detector design and performance : Technical Design Report, LHCb Collaboration, 2003]

LHCb muon detector

- Four multiwire proportional chambers
- Three iron layers of each $4.8\lambda_I$ (80 cm of iron)
- Large decay volume
- But **not** designed for shower detection
  - No energy deposit measurements
- Very **clean** environment
  - First plane (M2) after $6.7\lambda_I$ of material

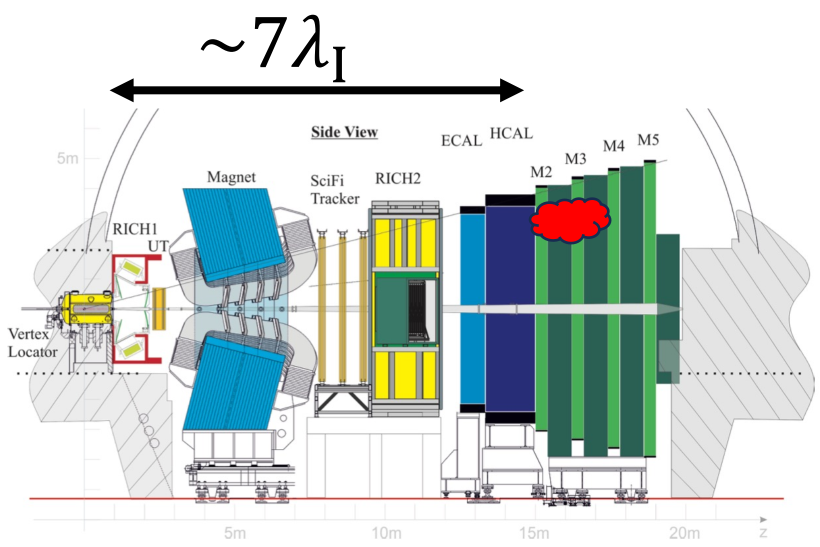

LHCb reoptimized detector design and performance : Technical Design Report, LHCb Collaboration, 2003

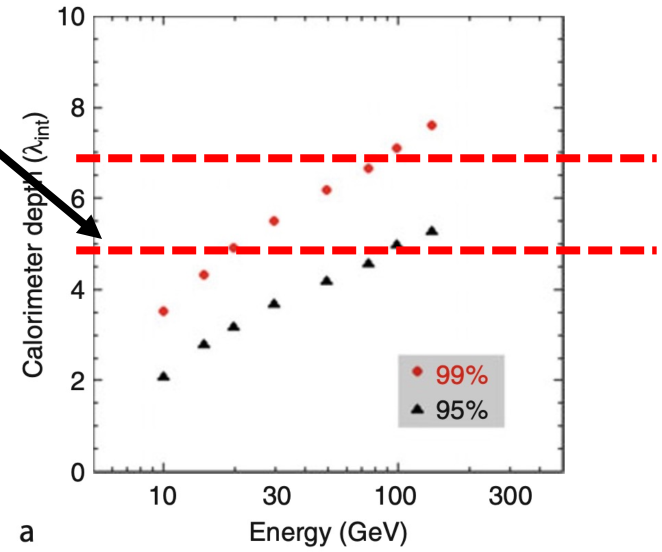

Handbook of Particle Detection and Imaging, R. Wigmans, pp 497–517

LHCb trigger

- Select events to save to disk
  - Run at 40 MHz
- **Full software** trigger at LHCb
  - Selecting **specific** signatures
- HLT1
  - ~99.9 – 99.99% background rejection
  - Running on **GPU** farm
    → fast neural network inference
- HLT2
  - Running on CPU
  - Partially saving event information

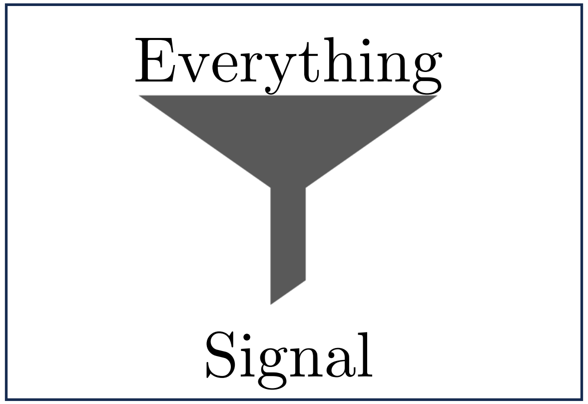

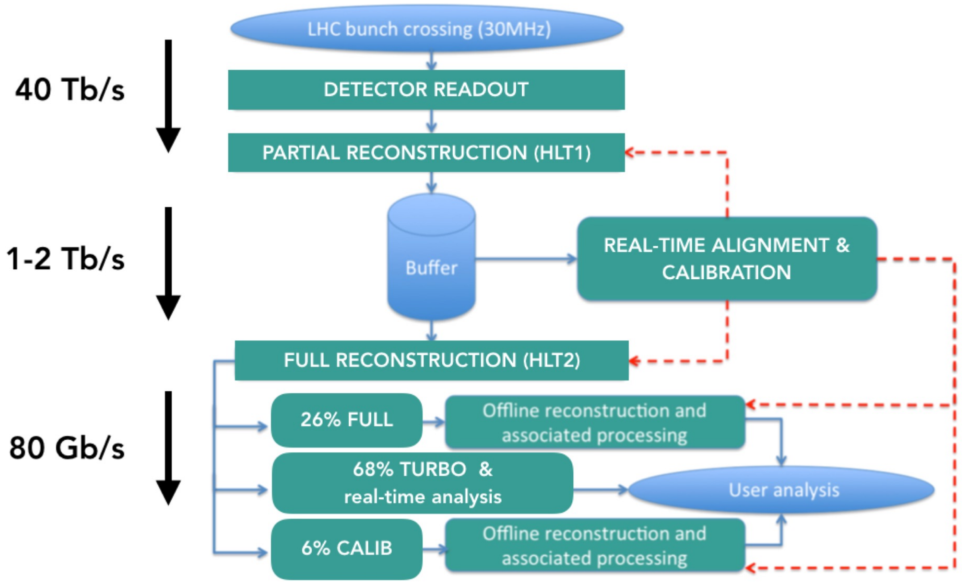

LHCB-TDR-021

Normalised autoencoders

- **Encoder** and **decoder** neural networks [7]
  - Information compression in the latent space
  - Train on the **background** to minimise the **reconstruction error**
- Add a **normalisation** to punish a too large reconstructible space
  - i.e., reconstruct well minimum bias events, and *only them*
  - Reconstructible = sufficiently low error
- Use **Monte Carlo sampling** to estimate the normalisation
  - Sampling probability related to the reconstruction error
- Train on unfiltered pp interactions
  - Evaluate efficiency on axion sample
    $H \to AA$, $A \to \tau^+\tau^-$, $\tau^\pm \to \pi^\pm\pi^\pm\pi^\mp\nu$, $m_A = 10$ GeV, $\tau_{\text{axion}} = 1$ ns
  - Only considering decays in muon detector

Does it work?

- The reconstruction error provides a very discriminant variable
  - Much larger for signal
- Similar/better than usual BDTs/NNs classifiers using signal samples
- Can be trained on data background only
  - No need for (MC) signal
  - No issues with MC-data differences

LHCb-FIGURE-2024-015

Including bin overflow

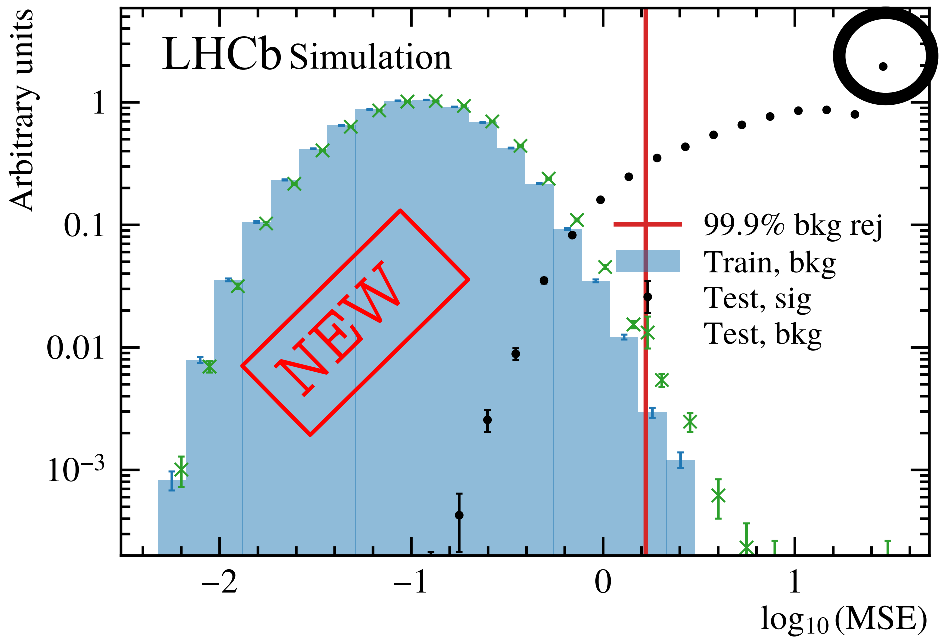

Sample	Efficiency [%]
Axion, 10 GeV	80.0 ± 0.5
HNL, 1.6 GeV	10.3 ± 0.3
HNL, 4 GeV	15.7 ± 0.3

How does efficiency vary?

NEW

- Naturally more performant for larger showers
  - Correlated with the energy
- Flat efficiency curve as a function of r_{xy}
  - Distance to the beam pipe
  - Worse near the beam → larger background
- Increased efficiency for decays before active layers
  - Shower starts in the shields
  - Distance until its maximal size

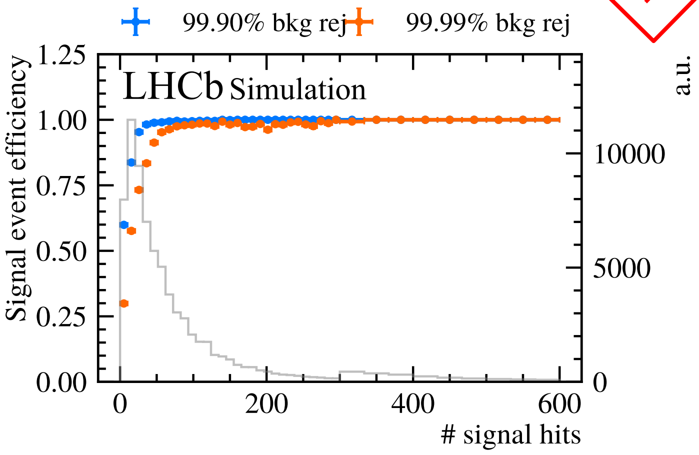

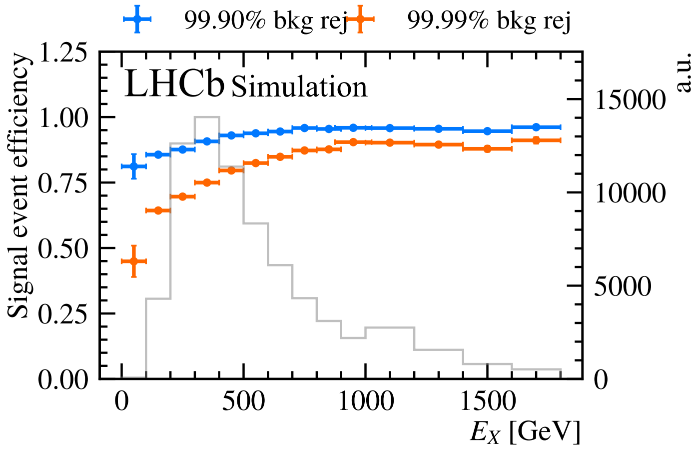

[LHCb-FIGURE-2024-015](#)

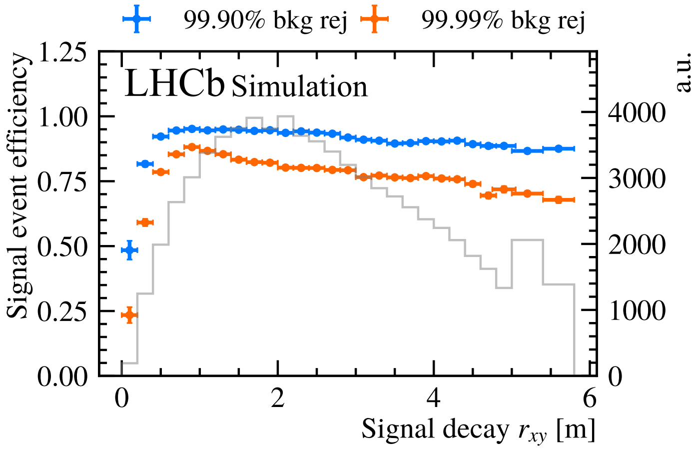

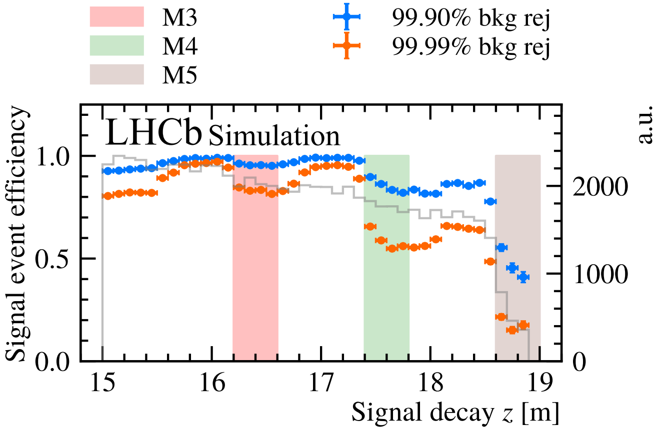

Data quality monitoring R&D

Proof of concept, not yet implemented in LHCb

- Many (non-expert) shifters required in the LHCb control room
  - Costly and with limited accuracy
  - Shifter rotation leads to variations in judgements
- Rewards based on human feedback
  - Data quality easy to spot for humans
  - Hard to manually provide updated references for all the histograms
- Operational regime changes over time
  - Model needs to adapt constantly
- Two contexts:
  - *Offline*: All labels available
  - *Online:* Shifter does not label all histograms

Human-in-the-loop Reinforcement Learning for Data Quality Monitoring in Particle Physics Experiments, O. J. Parra *et al.*, 2024, arXiv:2405.15508

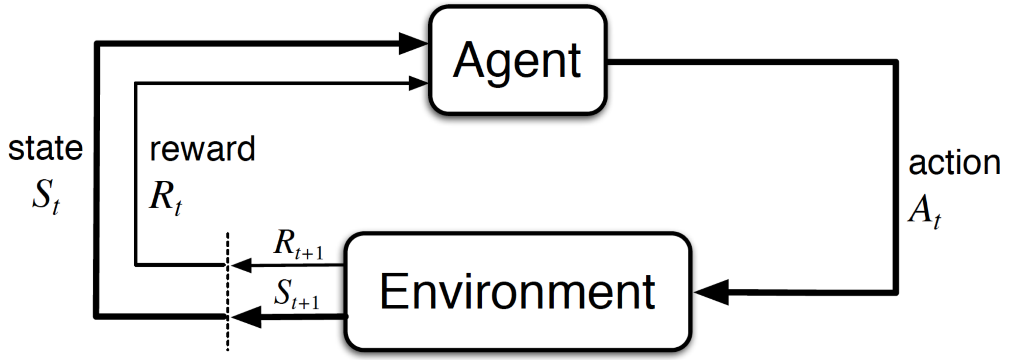

Reinforcement Learning: An introduction, R. S. Sutton & A. G. Barto, The MIT Press, 2015

The setup

- Human feedback
  - Flag data taking episodes (e.g. ~5 minutes) as normal or anomalous
  - Consider perfect or limited accuracy
- Reinforcement learning
  - *Predictor:* Classifies histogram as good/bad
  - *Checker*: Decides to call for feedback or not
  - Small multilayer perceptron
- Rewards:
  - *Predictor:* if correct/wrong
  - *Checker*: based on confidence on its decision

Toy results and next steps

- Generated toy samples
  - Same distribution for normal episodes
  - Varied distributions for the anomalies
- Algorithm performs beyond noisy labels
  - Also resistant to changes
  - Even when biased by the predictions

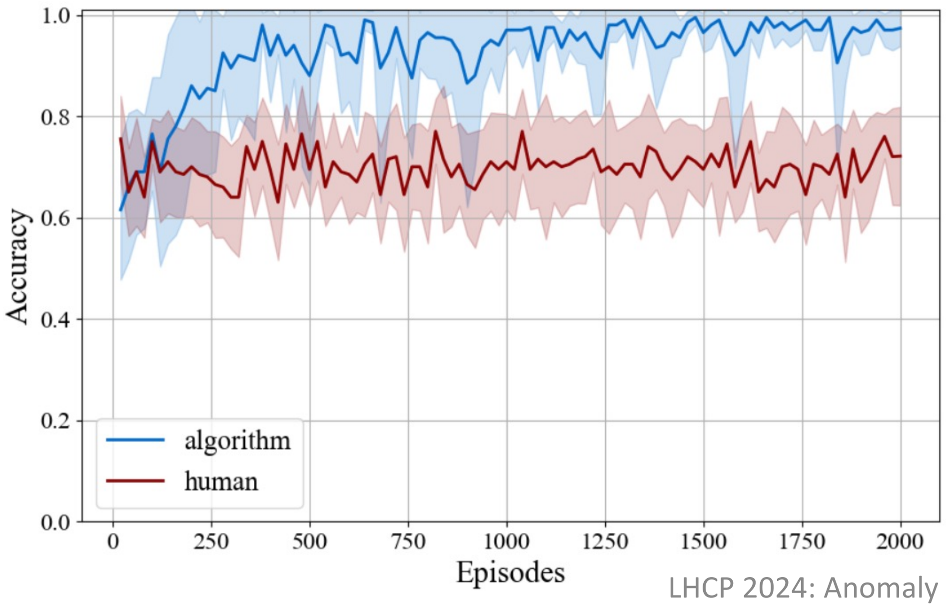

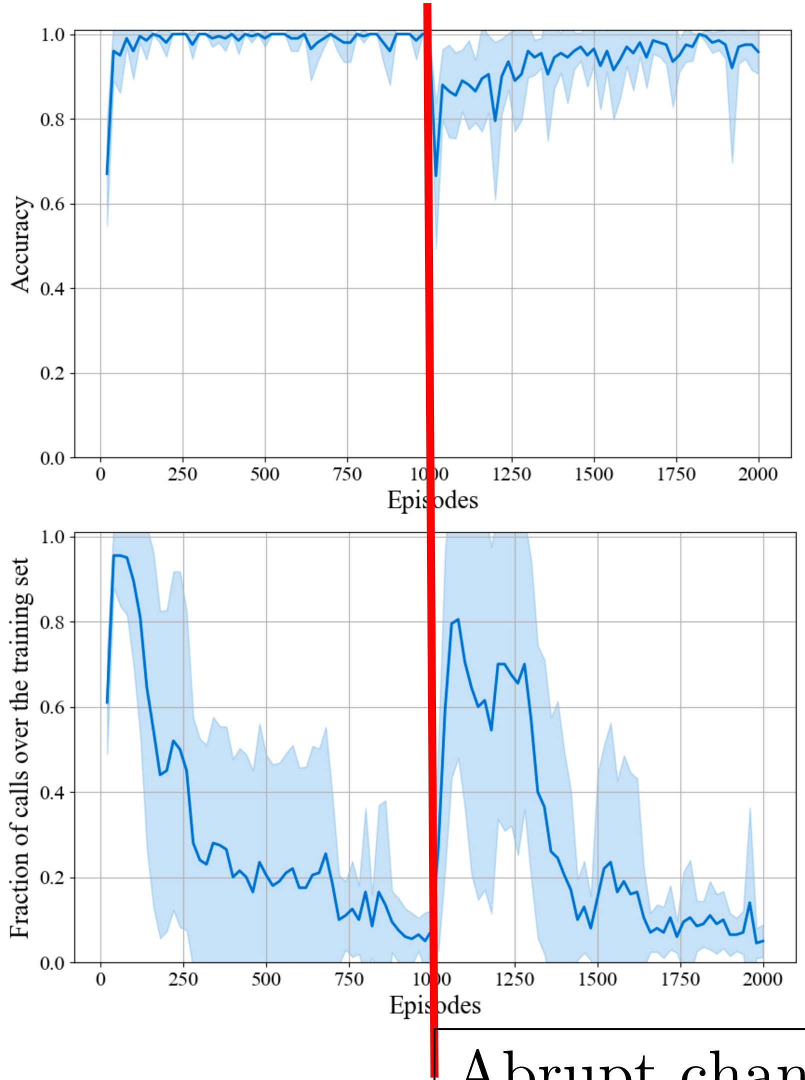

Abrupt change

arXiv:2405.15508

Summary

- Development of **first anomaly detection trigger in LHCb**
  - Making use of GPUs
  - Increasing sensitivity for LLPs

LHCb-FIGURE-2024-015

- **Comparable** (if not better) results than usual classifiers methods
- **Model independent** selection
  - Higher efficiencies
  - Fewer trigger lines, fewer models to develop and maintain
  - Fewer weights to store in memory and to infer
  - Not (as) limited by data-MC differences
- Development of anomaly detection for **data quality monitoring**
  - **Promising results** on toy study
  - Higher quality for less effort
  - Study application in LHCb

arXiv:2405.15508

References

- [1] Heavy neutral leptons - minimal and testable explanation for Beyond Standard Model phenomena, , K. Bondarenko, 2021, HEPHY seminar
- [2] Axion cosmology, D. J. E. Marsh, Physics Reports Volume 643, 2016
- [3] Search for events with a pair of displaced vertices from long-lived neutral particles decaying into hadronic jets in the ATLAS muon spectrometer in pp collisions at $\sqrt{s} = 13$ TeV, ATLAS Collaboration, 2022, Phys.Rev.D 106 (2022) 3, 032005
- [4] Search for Long-Lived Particles Decaying in the CMS End Cap Muon Detectors in Proton-Proton Collisions at $\sqrt{s} = 13$ TeV, CMS Collaboration, 2021, Phys.Rev.Lett. 127 (2021) 26, 261804
- [5] Search for long-lived heavy neutral leptons decaying in the CMS muon detectors in proton-proton collisions at $\sqrt{s} = 13$ TeV, CMS Collaboration, 2024, arXiv:2402.18658
- [6] Feebly interacting particles: Status and Perspectives (with special attention to LHCb), G. Lanfranchi, INFN e Laboratori Nazionali di Frascati, *Public LHCb meeting on Feebly Interacting Particles*, 2024
- [7] Autoencoder Under Normalization Constraints, S. Yoon *et al*, 2021, arXiv:2105.05735

BACKUP

Simple BDT

- Hit counting not sufficiently discriminant for the trigger
- Improved efficiencies by using a χ^2/N test using the number of hits in N parts of the muon system
- Train a BDT as standard in HEP
  - Reference to compare the other models to
  - Train on MC background and signal (axion)
  - Good results on the axion, much poorer on the HNLs$\rightarrow eX$
  - Very large differences between data and MC
  - Very sensitive to data taking conditions

Sample	Axion	HNL 1.6 GeV	HNL 4 GeV
Sig. eff. @ 99.99% bkg rej	$(48.4 \pm 0.4)\%$	$(6.1 \pm 0.2)\%$	$(8.3 \pm 0.2)\%$

BDT and NN output

- Some overtraining in the tail of the background distributions
  - Very difficult to remove without sacrificing signal efficiencies
  - Very few events in that range

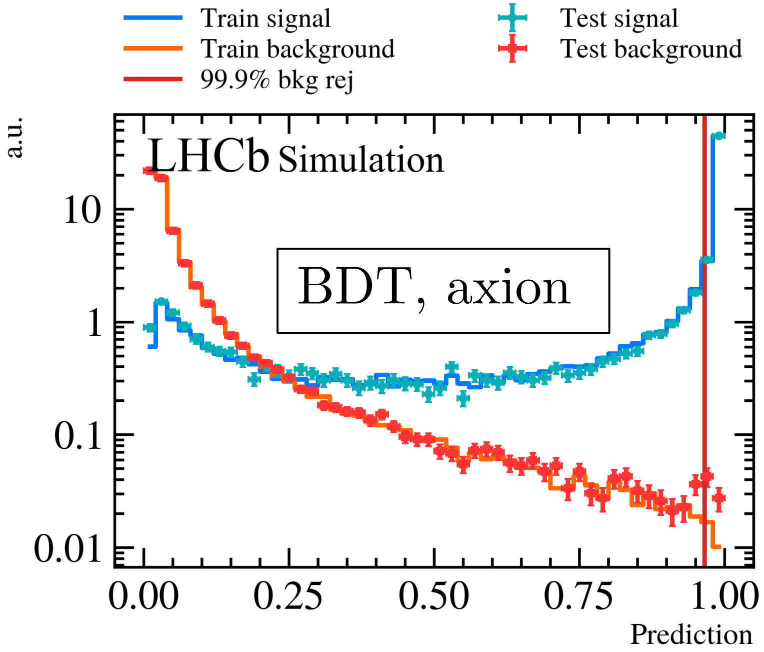

NEW

[LHCb-FIGURE-2024-015]

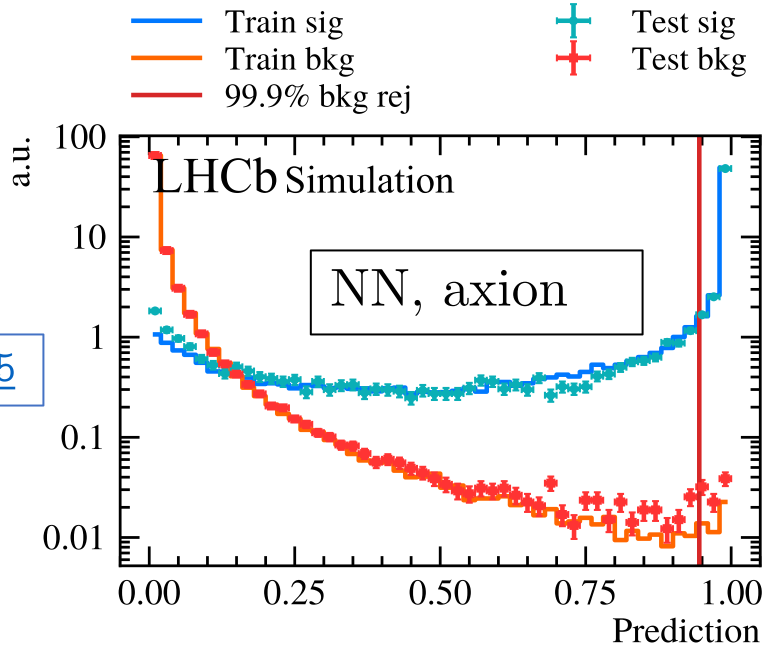

NN efficiency curves

- Similar trends as for the NAE
- More hits from the signal shower required
- Translates to larger energy
- Stronger dip for larger r_{xy}

NEW

LHCb-FIGURE-2024-015

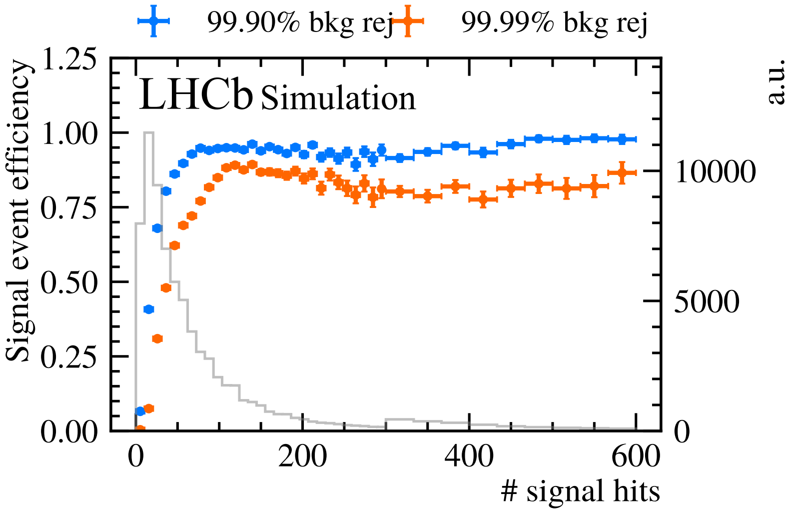

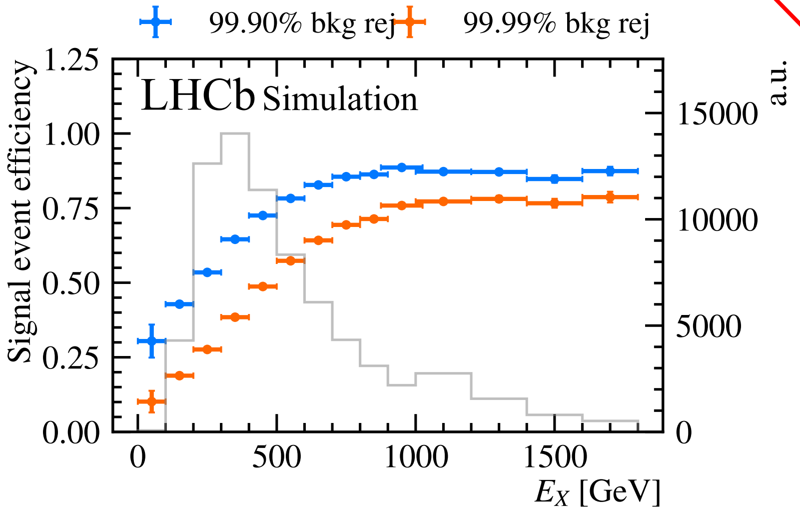

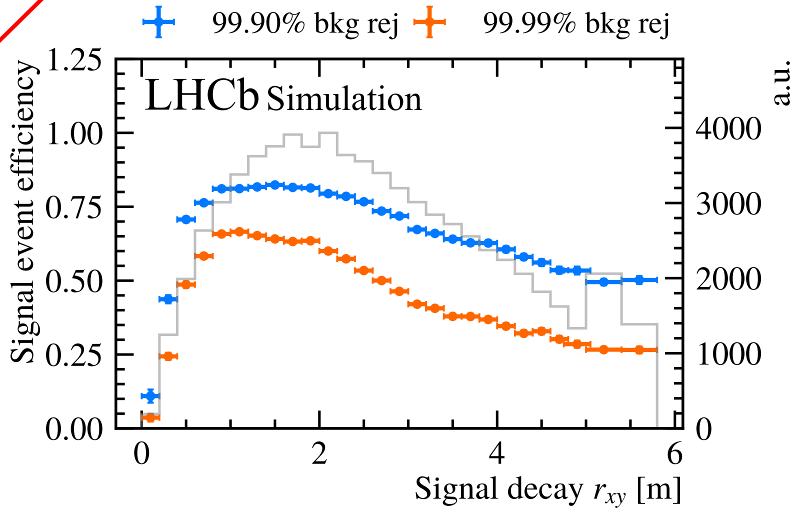

Samples and features

- Simulated samples
  - Mininum bias: unfiltered pp events
  - Axion $H \to AA$, $A \to \tau\tau$, $\tau \to \pi\pi\pi$, $m_A = 10$ GeV
  - HNLs $B_u \to N\mu$, $N \to eX$, $m = 1.6, 4$ GeV
    - X anything hadronic leading to a shower
- Require a shower within the muon detector
  - Shower caused by particle decay in the shields
- Number of hits per parts of the detector
  - Separate per station, region, quarter
  - Outermost region (4) split into three

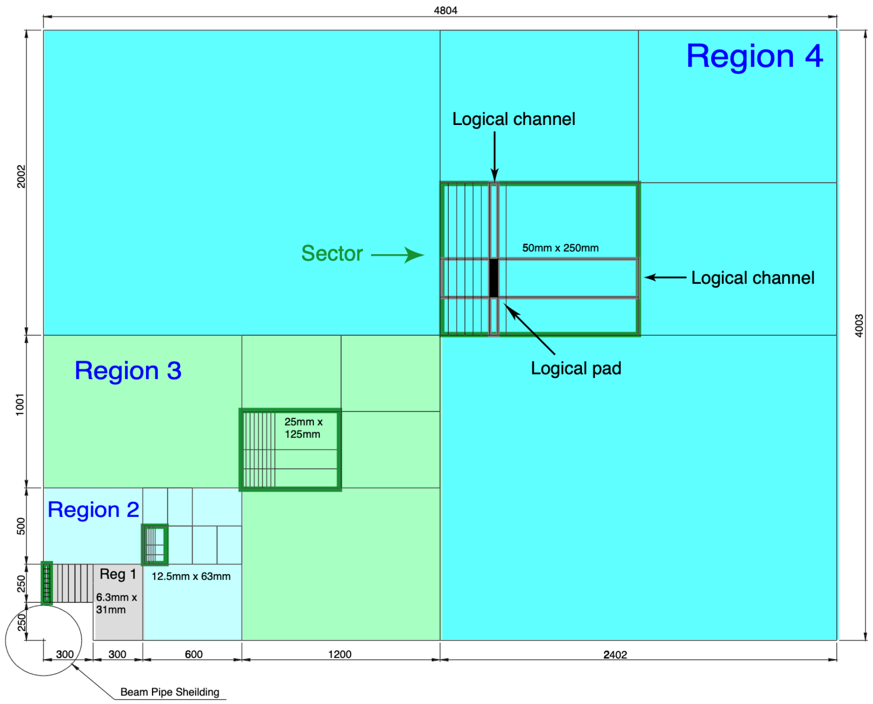

LHCb muon system : Technical Design Report, 2001

Autoencoders in a nutshell

- Two back-to-back neural networks
  - “Encoder”: reduces the dimensionality of the input
  - Bottleneck: dimension of the latent space
  - “Decoder”: reconstructs the original input from the latent space
- Train to minimise the reconstruction error
- Bottleneck reduces the “generalisation” to other types of events
  - Small error on background events
  - Large error on signal events

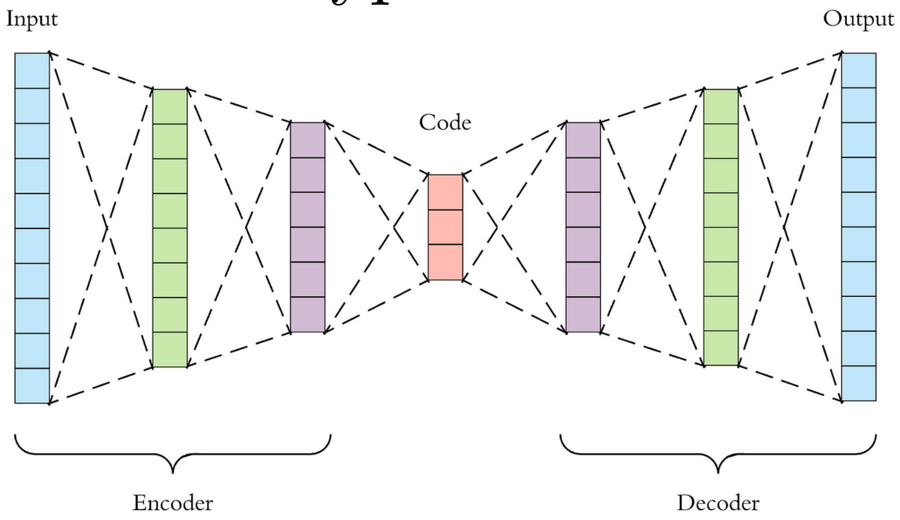

Applied Deep Learning - Part 3: Autoencoders, Arden Dertat

How well does it work?

LHCb-FIGURE-2024-015

- Clear struggle to reconstruct the input well for the signal
- Significant portion of the signal as well reconstructed as the background
- Make use of the HLT1 GPUs

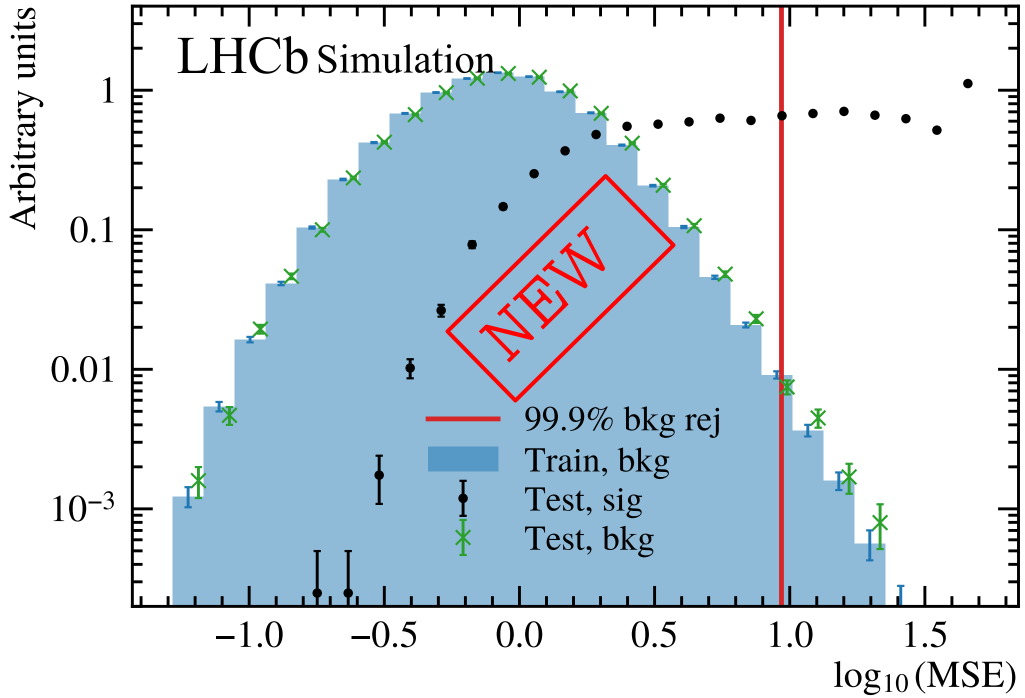

Sample	Axion	HNL 1.6 GeV	HNL 4 GeV
Sig. eff. @ 99.99% bkg rej	$(38.9 \pm 0.2)\%$	$(3.3 \pm 0.2)\%$	$(5.3 \pm 0.2)\%$

What is limiting the performance?

- Reconstructible space larger than just the only the background events distribution
- Signal events also well reconstructed → efficiency loss
- Ideally only reconstruct well the background
- Constrain the “size” of the reconstructible space

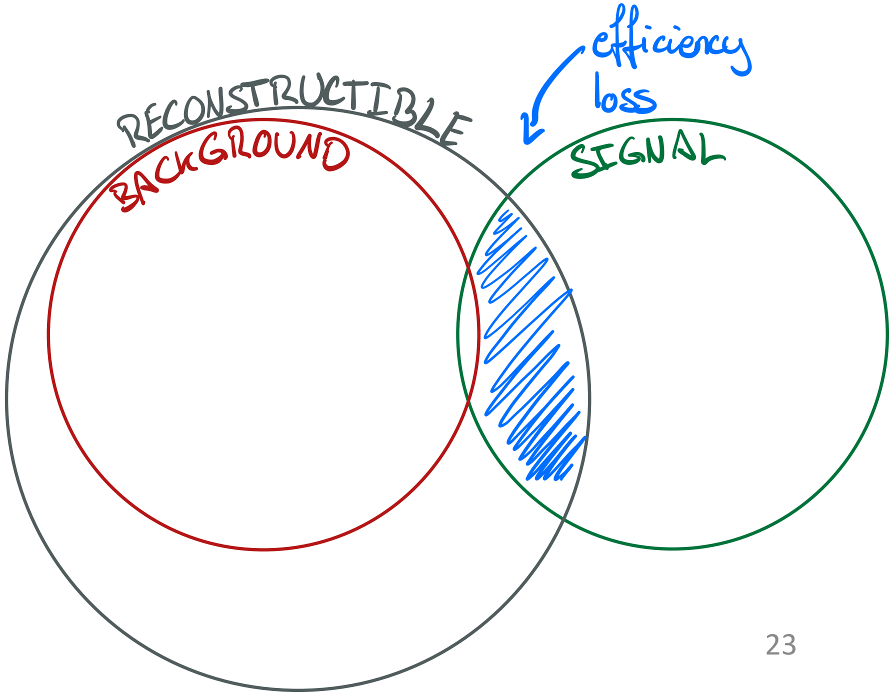

Can we improve using some signal?

- Autoencoders have limited adaptability
- Siamese neural networks
  - Trained on pairs of (background, background) and (background, signal)
  - Keep a set of reference events that could be updated
- Lower performance than NAE

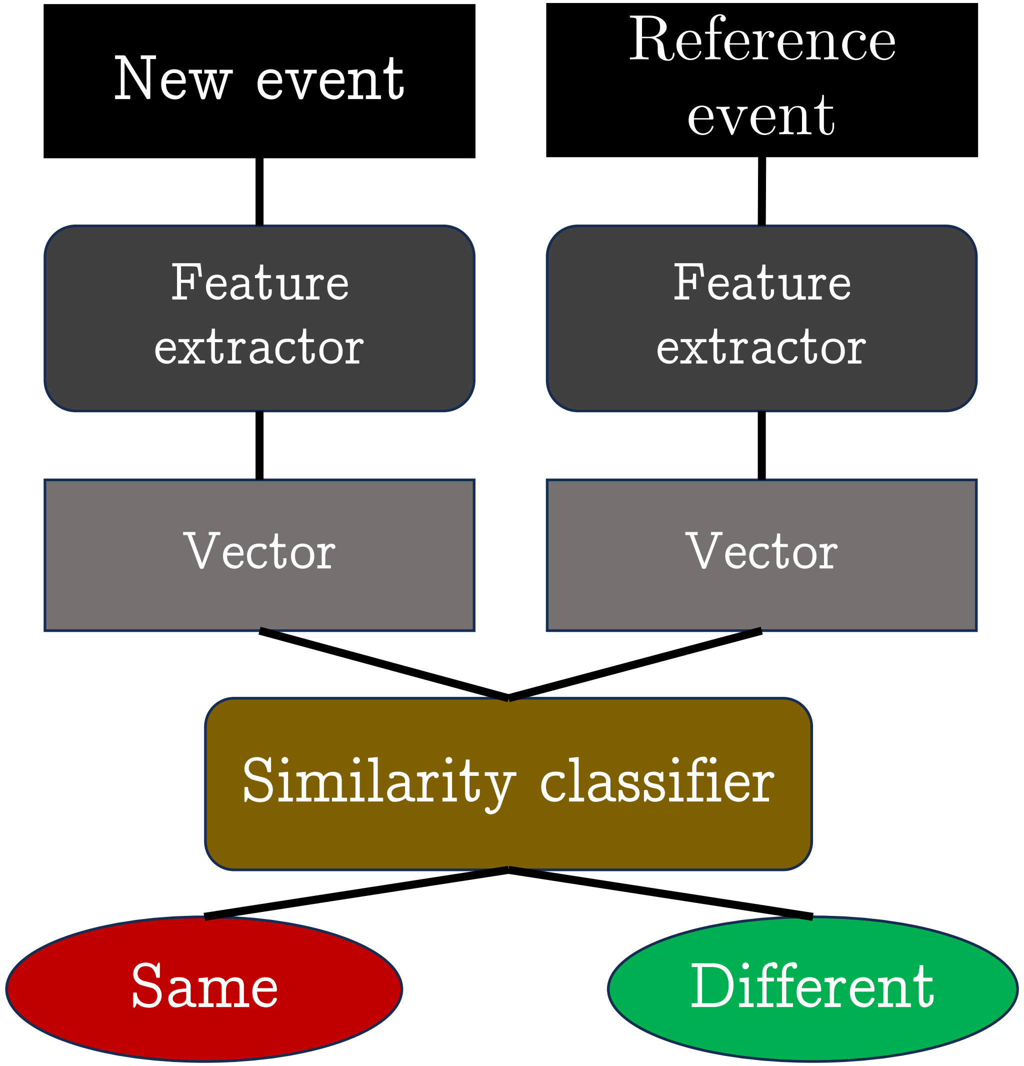

Sample	Efficiency [%]
Axion, 10 GeV	27.8 ± 0.4
HNL, 1.6 GeV	3.9 ± 0.2
HNL, 4 GeV	4.6 ± 0.2

LHCb-FIGURE-2024-015

Siamese NN: nHits dependence

- Requires many more hits from the signal shower to be efficiently selected
- Efficiency never reaches **100%**
- Further improvements could be possible
  - Contrastive loss function instead of a usual binary cross-entropy
  - Currently unbalanced samples
    - Background: $\sim 450k$ events
    - Signal: $\sim 80k$ events
    - Limited by inefficient signal generation

LHCb-FIGURE-2024-015