

Online-compatible Unsupervised Non-resonant Anomaly Detection

Vinicius M. Mikuni, Benjamin Nachman, David Shih



Introduction

Looking for **new physics scenarios** that we don't know yet is challenging

- How does it look like?
- What is the underlying theory behind it?

For many years, **supervised** new physics searches were performed:

- Propose a well-motivated model
- Look for signatures
- If no deviation from the prediction, place upper limits

Kiv:2204.12413v1 [hep-ex] 26 Apr 2022



CMS-B2G-21-003



CERN-EP-2022-034
2022/04/27

Search for a massive scalar resonance decaying to a light scalar and a Higgs boson in the four b quarks final state with boosted topology

The CMS Collaboration

Abstract

We search for new massive scalar particles X and Y through the resonant process $X \rightarrow YH \rightarrow b\bar{b}b\bar{b}$, where H is the standard model Higgs boson. Data from CERN LHC proton-proton collisions are used, collected at a centre-of-mass energy of 13 TeV in 2016–2018 and corresponding to an integrated luminosity of 138 fb^{-1} . The search is performed in mass ranges of 0.9–4 TeV for X and 60–600 GeV for Y , where both Y and H are reconstructed as Lorentz-boosted single large-area jets. The results are interpreted in the context of the next-to-minimal supersymmetric standard model and also in an extension of the standard model with two additional singlet scalar fields.

The 95% confidence level upper limits for the production cross section vary between 0.1 and 150 fb depending on the X and Y masses, and represent a significant improvement over results from previous searches.



Anomaly detection

<https://iml-wg.github.io/HEPML-LivingReview/>

- Anomaly detection.

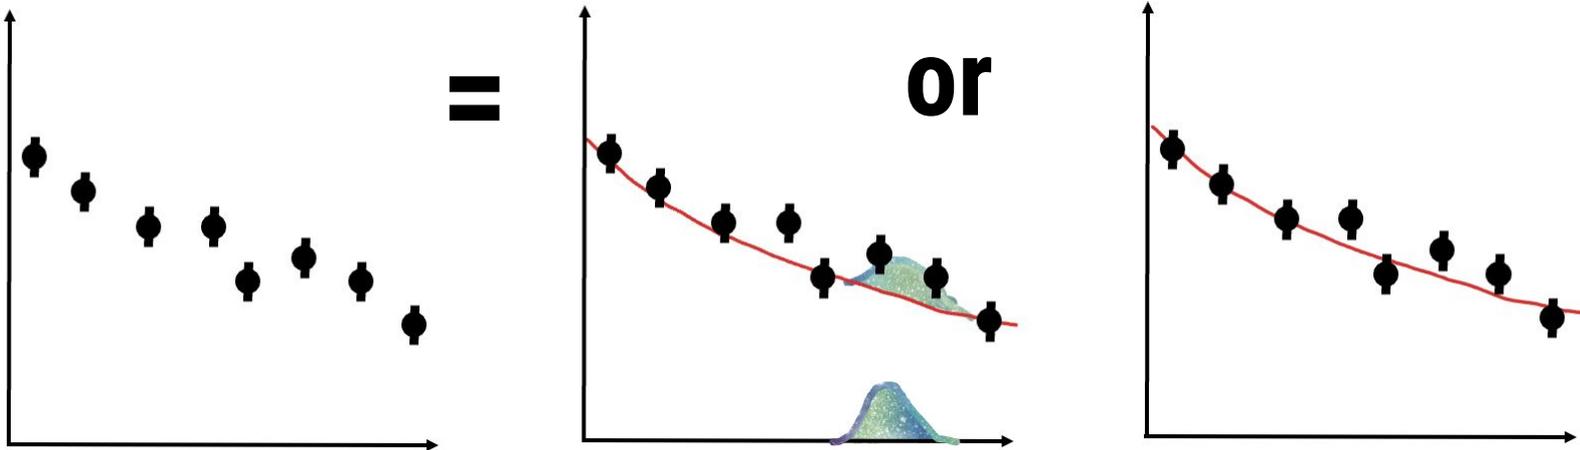
- Learning New Physics from a Machine [DOI]
 - Anomaly Detection for Resonant New Physics with Machine Learning [DOI]
 - Extending the search for new resonances with machine learning [DOI]
 - Learning Multivariate New Physics [DOI]
 - Searching for New Physics with Deep Autoencoders [DOI]
 - QCD or What? [DOI]
 - A robust anomaly finder based on autoencoder
 - Variational Autoencoders for New Physics Mining at the Large Hadron Collider [DOI]
 - Adversarially-trained autoencoders for robust unsupervised new physics searches [DOI]
 - Novelty Detection Meets Collider Physics [DOI]
 - Guiding New Physics Searches with Unsupervised Learning [DOI]
 - Does SUSY have friends? A new approach for LHC event analysis [DOI]
 - Nonparametric semisupervised classification for signal detection in high energy physics
 - Uncovering latent jet substructure [DOI]
 - Simulation Assisted Likelihood-free Anomaly Detection [DOI]
 - Anomaly Detection with Density Estimation [DOI]
 - A generic anti-QCD jet tagger [DOI]
-
- Transferability of Deep Learning Models in Searches for New Physics at Colliders [DOI]
 - Use of a Generalized Energy Mover's Distance in the Search for Rare Phenomena at Colliders [DOI]
 - Adversarially Learned Anomaly Detection on CMS Open Data: re-discovering the top quark [DOI]
 - Dijet resonance search with weak supervision using 13 TeV pp collisions in the ATLAS detector [DOI]
 - Learning the latent structure of collider events [DOI]
 - Finding New Physics without learning about it: Anomaly Detection as a tool for Searches at Colliders [DOI]
 - Tag N' Train: A Technique to Train Improved Classifiers on Unlabeled Data [DOI]
 - Variational Autoencoders for Anomalous Jet Tagging
 - Anomaly Awareness
 - Unsupervised Outlier Detection in Heavy-Ion Collisions
 - Decoding Dark Matter Substructure without Supervision
 - Mass Unspecific Supervised Tagging (MUST) for boosted jets [DOI]
 - Simulation-Assisted Decorrelation for Resonant Anomaly Detection
 - Anomaly Detection With Conditional Variational Autoencoders
-
- Unsupervised clustering for collider physics
 - Combining outlier analysis algorithms to identify new physics at the LHC
 - Quasi Anomalous Knowledge: Searching for new physics with embedded knowledge
 - Uncovering hidden patterns in collider events with Bayesian probabilistic models
 - Unsupervised in-distribution anomaly detection of new physics through conditional density estimation
 - The LHC Olympics 2020: A Community Challenge for Anomaly Detection in High Energy Physics
 - Model-Independent Detection of New Physics Signals Using Interpretable Semi-Supervised Classifier Tests
 - Topological Obstructions to Autoencoding
 - Unsupervised Event Classification with Graphs on Classical and Photonic Quantum Computers
 - Bump Hunting in Latent Space
 - Comparing Weak- and Unsupervised Methods for Resonant Anomaly Detection
 - Better Latent Spaces for Better Autoencoders
 - Autoencoders for unsupervised anomaly detection in high energy physics
 - Via Machinae: Searching for Stellar Streams using Unsupervised Machine Learning
 - Anomaly detection with Convolutional Graph Neural Networks
 - Anomalous Jet Identification via Sequence Modeling
 - The Dark Machines Anomaly Score Challenge: Benchmark Data and Model Independent Event Classification for the Large Hadron Collider
 - RanBox: Anomaly Detection in the Copula Space
 - Rare and Different: Anomaly Scores from a combination of likelihood and out-of-distribution models to detect new physics at the LHC
 - LHC physics dataset for unsupervised New Physics detection at 40 MHz
 - New Methods and Datasets for Group Anomaly Detection From Fundamental Physics
-
- Autoencoders on FPGAs for real-time, unsupervised new physics detection at 40 MHz at the Large Hadron Collider
 - Classifying Anomalies THrough Outer Density Estimation (CATHODE)
 - Deep Set Auto Encoders for Anomaly Detection in Particle Physics
 - Challenges for Unsupervised Anomaly Detection in Particle Physics
 - Improving Variational Autoencoders for New Physics Detection at the LHC with Normalizing Flows
 - Signal-agnostic dark matter searches in direct detection data with machine learning



Anomaly detection

What does a **successful** anomaly detector needs?

- Capable of **identifying anomalous events**
- Provide a description of non-anomalous events: **background estimation**

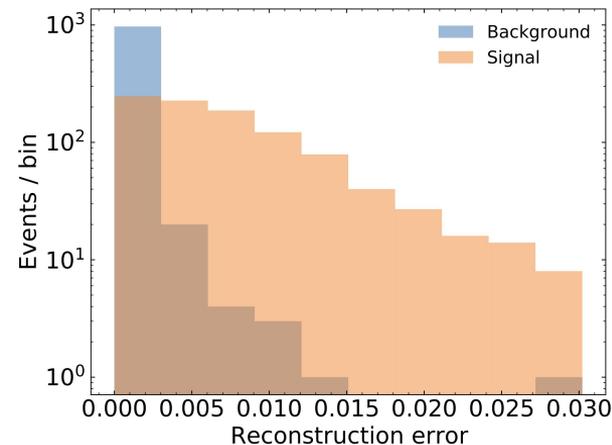
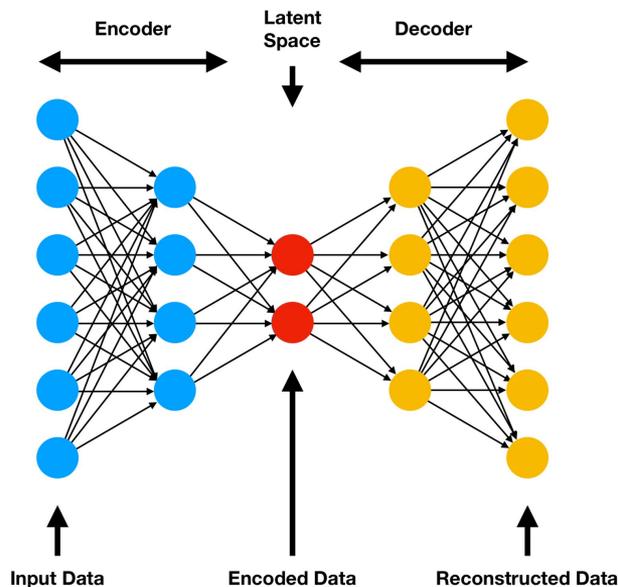




Autoencoders

Autoencoders are popular choices of **anomaly detection** algorithms

- Train the model with background-enriched data
- Encode the inputs to a low dimensional representation and try to decode it back to the input set
- Anomalous events** are often **poorly reconstructed** given low, if any, examples present during training
- What about background estimation?

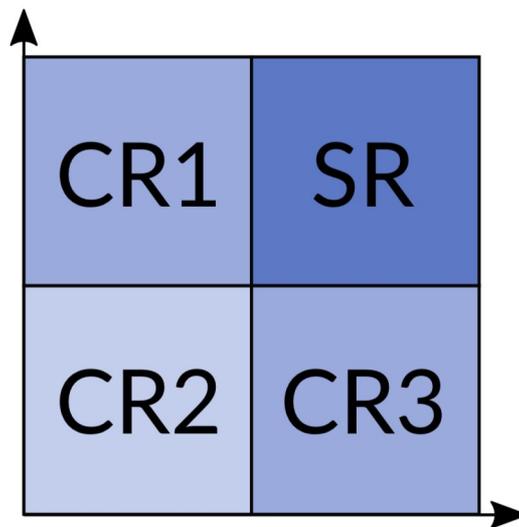




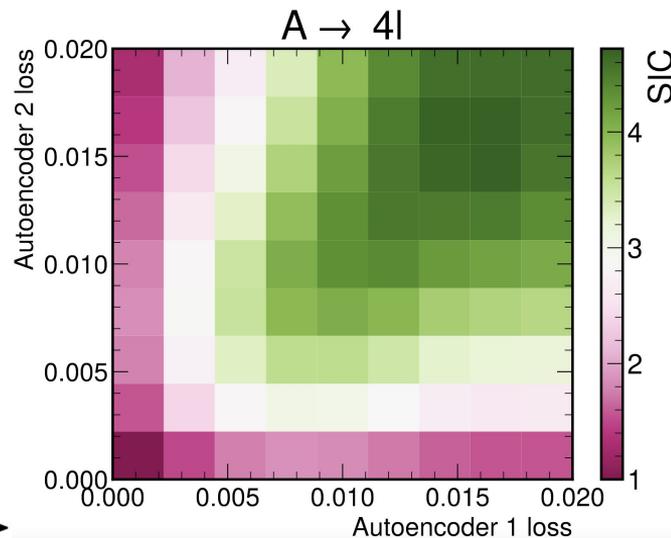
Autoencoders

ABCD method is a popular choice of data-driven background estimation

- Requires 2 **background-independent** distributions
- Both** distributions should provide **signal sensitivity** to avoid contamination
- Background in the signal-enriched region is described by the other background-dominated regions



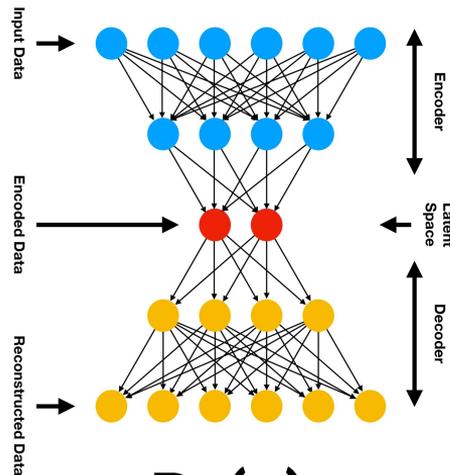
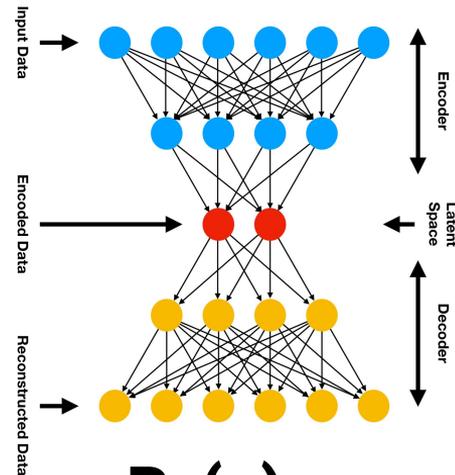
$$SR = CR1 * CR3 / CR2$$





Decorrelated autoencoders

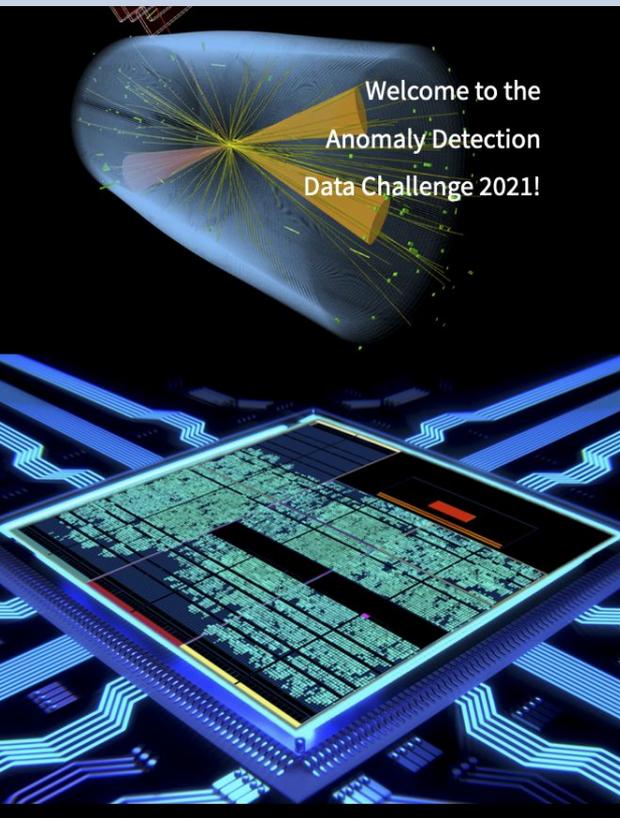
- Use the **reconstruction loss** of each autoencoder to define thresholds for the ABCD method
- Enforce the decorrelation between loss functions using the distance correlation (**DisCo¹**) loss


 $R_1(x)$

 $R_2(x)$

$$L[f_1, f_2, g_1, g_2] = \sum_i R_1(x_i)^2 + \sum_i R_2(x_i)^2 + \lambda \text{DisCo}^2[R_1(X), R_2(X)]$$



ADC2021 dataset



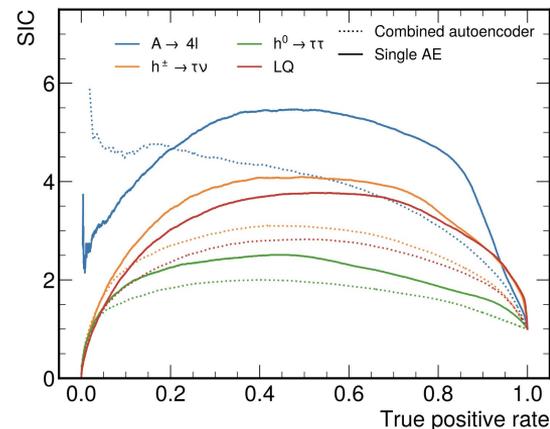
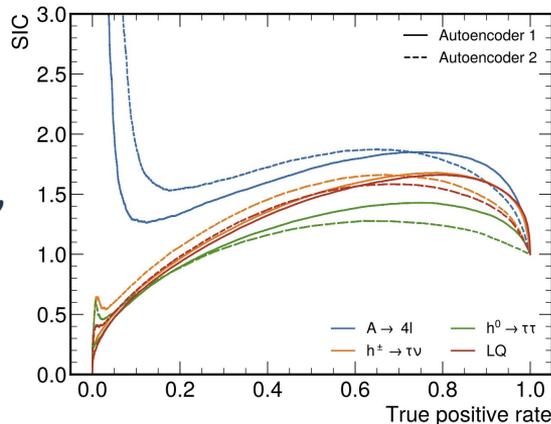
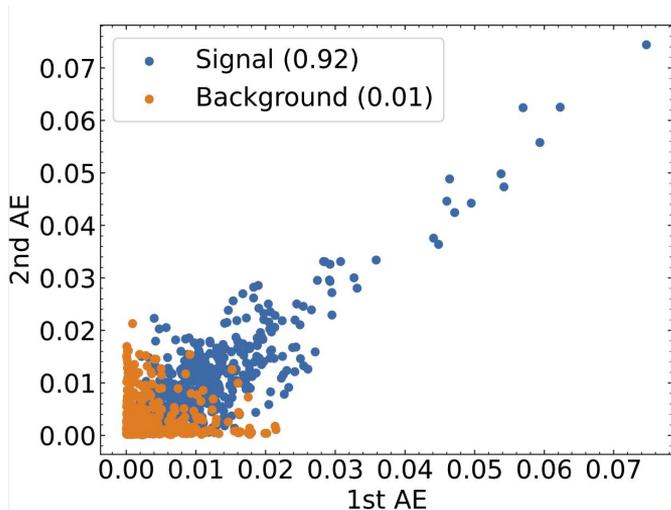
- Test the idea in a realistic setting: **anomaly detection at trigger level**
- **Goal:** Create algorithms that can trigger anomalous events that would otherwise be thrown away
- Dataset consists of a cocktail of **Standard Model processes** passing a **single lepton trigger**
- Momenta of leading **4 leptons** and **10 jets** are saved and used as inputs to the autoencoder
- **No invariant mass information used**
- **Train on background events** and **evaluate** over different new **physics scenarios** to test the performance



ADC2021 dataset

New physics benchmarks

- Neutral scalar boson (**A**), 50 GeV \rightarrow 4 l
- Leptoquark (**LQ**), 80 GeV \rightarrow b τ
- Scalar boson (**h⁰**), 60 GeV \rightarrow $\tau \tau$
- Charged scalar boson (**h[±]**), 60 GeV \rightarrow $\tau \nu$



SIC = Significance improvement characteristic:
 $\text{tpr}/\sqrt{\text{fpr}}$

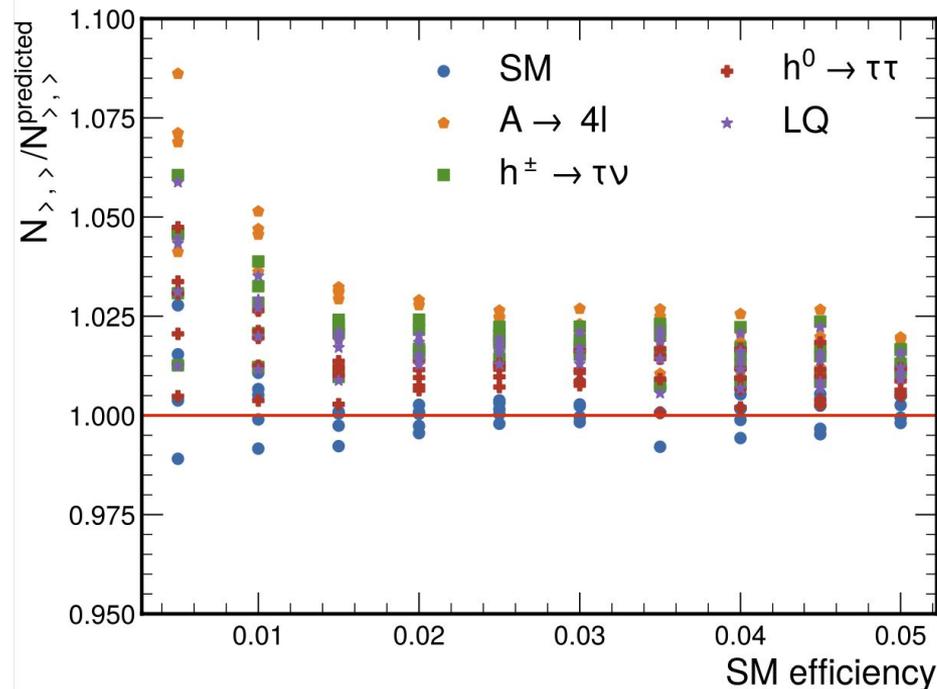


ADC2021 dataset

Calculate the **background** in the signal enriched region using the **ABCD method**

- **Non-closure test:** compare **real number** of events with **predicted background**
- Different threshold choices resulting in different results
- Nevertheless, samples with **new physics scenarios** consistently having **more events** than predicted

Spread in the y-axis represents the results when different selection thresholds are used

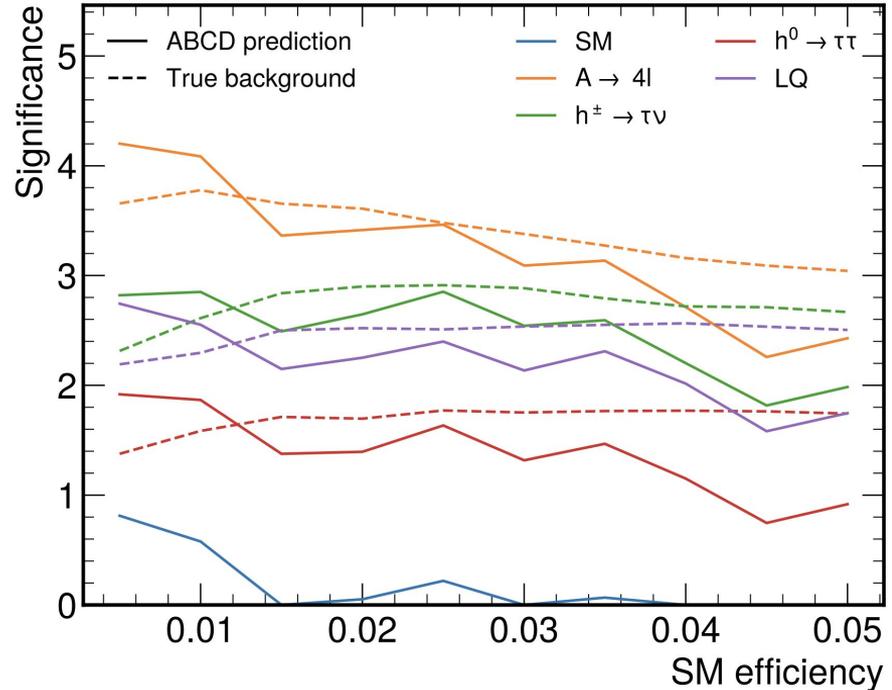




ADC2021 dataset

Quantify the difference in terms of **signal significance**

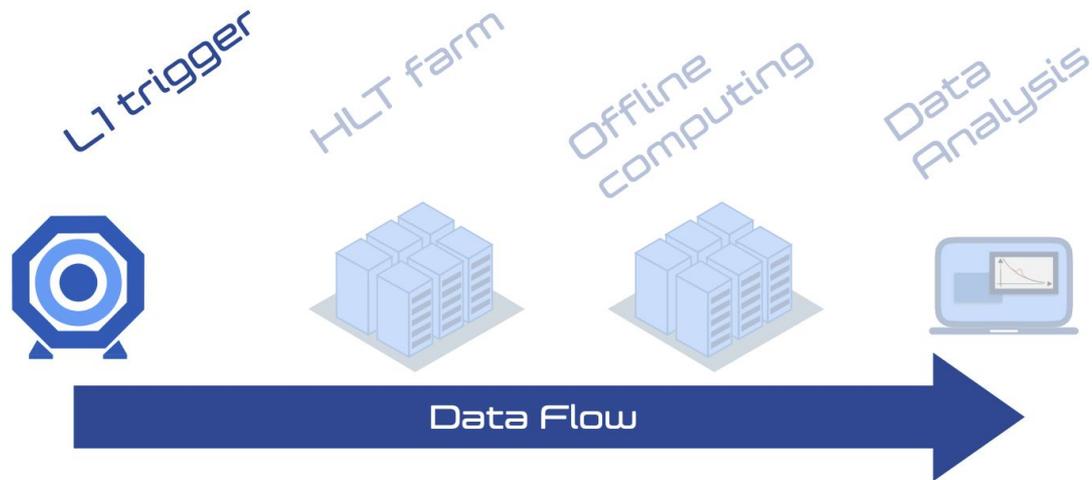
- **Less than 1 sigma** for sample **without NP** and **1-4 for different NP scenarios**
- **Signal contamination** in the sidebands can lead to incorrect significances: Corrections to background prediction for limit setting



Significance = $(N-B)/\sqrt{N}$, if $N > B$ and 0 otherwise



The LHC Big Data problem



- 40 MHz in / 100 KHz out
- ~ 500 KB / event
- Processing time: ~10 μ s
- Based on coarse local reconstructions
- FPGAs / Hardware implemented

- Tune the autoencoder thresholds to **save all events in the signal enriched region**
- **Prescale** the other 3 regions to determine the **background composition**
- Train events using **simulation** or **data** directly
 - ▷ Use data from a previous run or from a minimum bias trigger



Autoencoders on FPGAs for real-time, unsupervised new physics detection at 40 MHz at the Large Hadron Collider

VIII. CONCLUSIONS

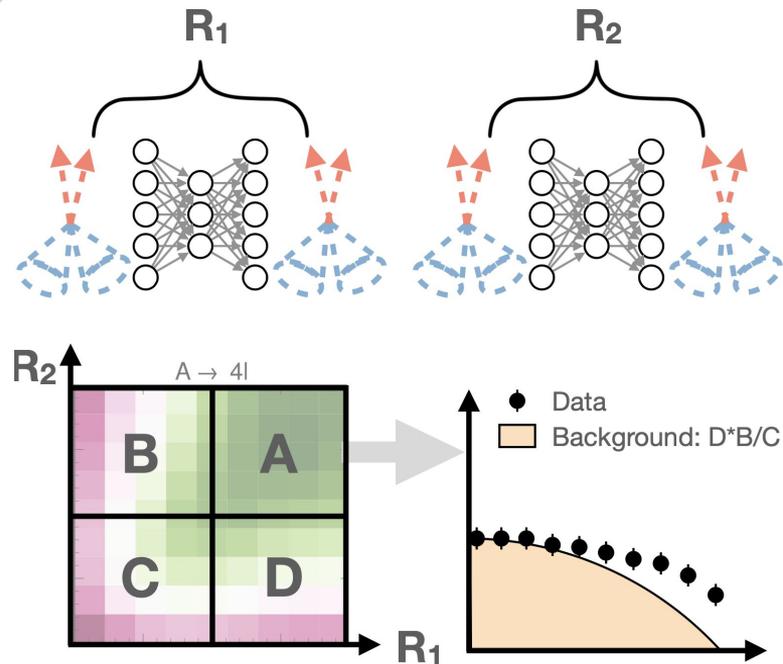
We discussed how to extend new physics detection strategies at the LHC with autoencoders deployed in the L1T infrastructure of the experiments. In particular, we show how one could deploy a deep neural network (DNN) or convolutional neural network (CNN) AE on a field-programmable gate array (FPGA) using the `hls4ml` library, within a $\mathcal{O}(1)\mu\text{s}$ latency and with small resource utilization once the model is quantized and pruned. We show that one can retain accuracy by compressing the model at training time. Moreover, we discuss different

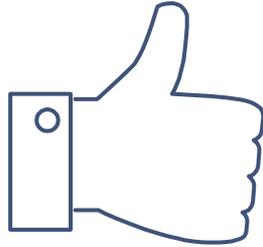
- Our model uses only fully connected layers: **demonstrated** to satisfy trigger budget constraints when running on FPGAs after pruning and compression
- **First complete** online compatible anomaly detection protocol to be proposed



Conclusions

- In this work we proposed an **online-compatible Unsupervised Non-resonant Anomaly detection** method
- We use **autoencoders** as **anomaly detectors** and enforce the **decorrelation** between reconstruction losses using the **DisCo** loss
- Background estimation using the ABCD method
 - **Non-closure** for samples containing **new physics** events
 - **Significances** up to **4** for initial signal **contaminations** of **0.1%**
- **Online compatibility:** Signal enriched region saved together with prescaled sidebands for background estimation
- Available on [Phys. Rev. D 105, 055006](#)
- Scripts to run the model available on [github](#)





THANKS!

Any questions?

Decorrelation function

- Given the output space of 2 neural networks F and G, the distance covariance is defined as

$$\begin{aligned} \text{dCov}^2[f, g] = & \langle |f - f'| \times |g - g'| \rangle \\ & + \langle |f - f'| \rangle \times \langle |g - g'| \rangle - 2 \langle |f - f'| \times |g - g''| \rangle \end{aligned}$$

- Where f and f' are sampled from F and $g, g',$ and g'' are sampled from G
- The correlation distance is then defined as

$$\text{dCorr}^2[f, g] = \frac{\text{dCov}^2[f, g]}{\text{dCov}[f, f] \text{dCov}[g, g]}.$$