New Physics Searches at the LHC through Event-based Anomaly Detection and Development of ADFilter Web-tool

Wasikul Islam[*1], **Sergei Chekanov**[2], and **Nicholas Luongo**[2]

[1]Department of Physics, University of Wisconsin, Madison, WI 53706 USA
[2]HEP Division, Argonne National Laboratory, USA

(Presented at the 32nd International Symposium on Lepton Photon Interactions at High Energies, 25-29 August 2025 at Madison, Wisconsin.)

Abstract

This work presents advancements in model-agnostic searches for new physics at the Large Hadron Collider (LHC) through the application of event-based anomaly detection techniques utilizing unsupervised machine learning. We discuss the advantages of anomaly detection approach, as demonstrated in a recent ATLAS analysis, and introduce ADFilter, a web-based tool designed to process collision events using autoencoders based on deep unsupervised neural networks. ADFilter calculates loss distributions for input events, aiding in determining the degree to which events can be considered anomalous. Real-life examples are provided to demonstrate how the tool can be used to reinterpret existing LHC results, with the goal of significantly improving exclusion limits. Furthermore, we present a comparative study between anomaly detection and supervised machine learning techniques, using the search for heavy resonances decaying into two or more Higgs bosons as a representative case to demonstrate the application and effectiveness of these methods.

Preprint: HEP-ANL-201033

1 Introduction

The search for physics beyond the Standard Model (BSM) at the LHC has traditionally relied on targeted, model-driven analyses that optimize sensitivity for specific decay topologies. While powerful, such searches may lose coverage for unconventional, unexpected, or poorly modeled signatures. This limitation has motivated the development of *model-agnostic* strategies whose aim is to identify unusual events directly from data without relying on explicit signal hypotheses.

Event-level anomaly detection (AD), implemented using unsupervised autoencoders trained on data, has emerged as a promising framework in this direction. Unlike supervised approaches that require labeled signal samples, anomaly detection relies solely on Standard-Model–like events for training, thus providing a bias-minimized method for uncovering anomalous topologies. These techniques learn the high-dimensional correlations among reconstructed observables and assign an "anomaly score" derived from the reconstruction loss or latent-space structure.

A recent ATLAS study demonstrated the first full Run 2 anomaly-detection search based on deep autoencoders trained using 1287 observables constructed from jets, leptons, photons, and missing transverse momentum [1]. The outcome revealed both strong background modeling and enhanced sensitivity to a wide variety of possible new-physics topologies. Motivated by these developments, a publicly accessible web-based reinterpretation framework, ADFilter, has been developed, which applies the ATLAS autoencoder configuration to user-defined simulated

[*]Speaker at Lepton Photon 2025 conference.

samples. This work summarizes the anomaly-detection methodology, describes the ADFILTER framework, and presents comparative studies between anomaly detection and supervised classifiers in the context of di-Higgs resonance searches.

2 Event-Level Anomaly Detection in ATLAS

The ATLAS anomaly-detection analysis [1] employs a deep autoencoder trained on approximately 1% of Run 2 proton–proton collision data. Events are represented using the *Rapidity–Mass Matrix* (RMM), a 36×36 structured array (containing 1287 entries, excluding 9 invariant mass terms) encoding pairwise rapidity differences, invariant masses, and transverse momenta of reconstructed objects. This representation provides a dense summary of the full event topology while maintaining a uniform architecture for all final states.

The autoencoder compresses the 1287-dimensional RMM input into a 200-dimensional latent space before reconstructing it back to the original dimension. Events with large reconstruction loss are classified as anomalous. Thresholds in loss space corresponding to effective signal cross sections of roughly 10 pb, 1 pb, and 0.1 pb define three anomaly regions that bracket increasingly rare topologies.

Using these regions, the analysis performs bump-hunt searches in nine invariant-mass variables: m_{jj}, m_{jb}, m_{bb}, $m_{j\ell}$, $m_{j\gamma}$, m_{be}, $m_{b\mu}$,$m_{j\mu}$,$m_{j\gamma}$. Across all final states considered, no statistically significant deviations from the Standard Model expectation were observed. Localized fluctuations (such as 2.9σ near 4.8 TeV and 2.8σ near 1.2 TeV) were found to be consistent with the look-elsewhere effect. Nevertheless, the use of anomaly regions substantially improved sensitivity in the sub-TeV regime, with limits strengthened by factors of 2–3 relative to earlier inclusive strategies. These findings highlight the capability of event-level anomaly detection to complement targeted searches.

3 The ADFilter Web Tool

The ADFILTER framework [2, 3] was designed to provide the high-energy physics community with a practical mechanism for evaluating anomaly scores for arbitrary simulated models. The tool reproduces the ATLAS anomaly-detection pipeline by implementing the same RMM-based feature construction and the same publicly available trained autoencoder model used in Ref. [1].

3.1 Workflow

The ADFILTER workflow proceeds as follows:

1. Users upload simulated events generated in LHE, HEPMC, ProMC, or Delphes ROOT formats at the ADFilter web interface [2].
2. Events are preprocessed and reconstructed into the 1287-dimensional RMM feature set.
3. The pre-trained autoencoder evaluates each event to compute a reconstruction-loss value.
4. The anomaly-score distribution is produced, and the fraction of events falling into the ATLAS anomaly regions is reported.
5. From these results, the *anomaly-detection acceptance* A is computed as

$$A = \frac{N_{\mathrm{AR}}}{N_{\mathrm{tot}}},$$

where N_{AR} is the number of events inside a chosen anomaly region and N_{tot} is the total number of generated events.

The acceptance A plays a critical role in reinterpretation of ATLAS limits. The model-independent ATLAS anomaly-detection search sets limits on $\sigma \times B \times A \times \epsilon$, where ϵ denotes detector and selection efficiency. For theorists wishing to compare their BSM models with

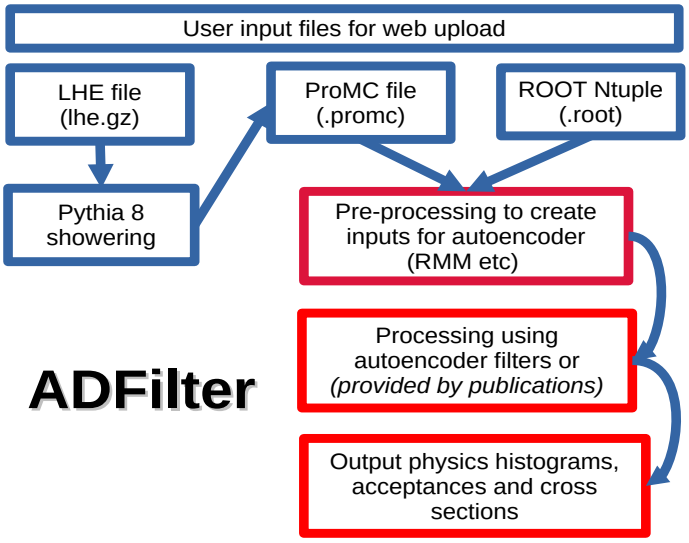

Figure 1: A schematic diagram showing the workflow of the ADFilter.

these limits, knowledge of A is essential. However, A cannot be determined without running the trained autoencoder used in Ref. [1]. This requires access to the network weights, RMM construction, and ML infrastructure—elements that are nontrivial for many groups to reproduce.

ADFilter removes this barrier by providing a direct evaluation of A for arbitrary signal hypotheses. A small acceptance ($A \ll 1$) indicates that the BSM model is too similar to the SM background to populate the anomaly regions, while a large acceptance ($A \sim 1$) implies that the model strongly populates anomalous parts of phase space. Once A is known, users can rescale the limits to obtain updated constraints on $\sigma \times B$ for their specific model. This provides a fast and robust route for model reinterpretation using the ATLAS anomaly-detection results.

3.2 Examples and Use Cases

Investigations using ADFilter demonstrate its usefulness across a range of benchmark new-physics models [3]. Examples include:

- A 2 TeV radion scenario, where an enhanced density of high-loss events reflects the more complex jet structure of radion decays.
- Heavy vector-boson models (W' and Z'), where elevated anomaly scores correspond to the expected kinematic tails of high-mass resonances.
- Charged-Higgs production (tbH^+), where boosted topologies lead to characteristic shapes in the anomaly-loss spectrum and modified exclusion limits.

Figure 2 shows, ADFilter showcases significant improvements in limits for for tbH^+ model using the anomaly detection technique. These studies demonstrate that ADFilter can strengthen exclusion limits and uncover parameter regions that are challenging for more traditional approaches.

4 Supervised Learning vs. Anomaly Detection for di-Higgs

Further investigations [5] explore the comparative behavior of supervised machine-learning classifiers and anomaly-detection models in di-Higgs resonance searches. The scenarios considered include $X \to HH$ and mixed $X \to SH$ decays (with $S \to HH$), which generate rich topologies containing multiple b-jets, boosted Higgs bosons, and high-p_T jets.

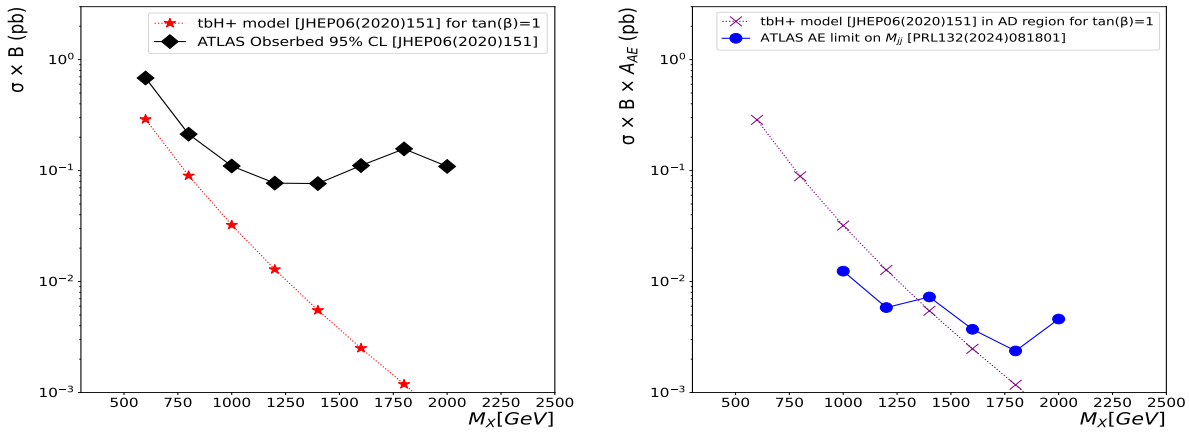

(a) Published LHC limits [4].

(b) Reinterpreted limits using Ref. [1] and ADFilter.

Figure 2: (a) Observed 95% credibility-level limits [4] on $\sigma \times B$ for tbH^+ production as a function of M_X, using dijet masses in events with at least one isolated lepton with $p_T > 60$ GeV. (b) ATLAS 95% limits for Gaussian signals from Ref. [1], rescaled using $(1/A_{\text{sel}} \times \varepsilon)$ following Ref. [4]. Magenta points show the cross section after the ADFilter correction. The anomaly-detection interpretation excludes the $\tan\beta = 1$ benchmark up to 1.35 TeV, whereas the original selection could not.

As Fig 3 shows, the supervised classifiers, trained on labeled signal and background samples, generally achieve the strongest separation when the assumed signal closely matches the training hypothesis. These networks perform particularly well for lower-mass resonances or for decay patterns that resemble documented signatures.

In contrast, anomaly-detection methods do not rely on explicit signal hypotheses and instead focus on deviations from the Standard Model event distribution. This confers an advantage in cases where the signal differs substantially from expected topologies or when the resonance mass becomes large enough that its decay products produce highly non-standard kinematic configurations.

Studies also show that the RMM representation captures characteristic features of di-Higgs events, such as dense substructure patterns from boosted Higgs bosons. These properties naturally influence the anomaly scores, while supervised classifiers benefit from the systematic presence of recurring features in the training data.

5 Outlook

Event-level anomaly detection has emerged as a versatile and model-agnostic strategy for exploring new-physics signals at the LHC. The ATLAS search demonstrates both competitive sensitivity and enhanced reach for unexpected signatures. The ADFilter tool makes such techniques broadly accessible, enabling rapid recasting studies for a variety of BSM scenarios.

Future improvements in ADFilter web tool may include the incorporation of additional neural network architectures, support for various detector formats, and the collection of trained anomaly detection models from different detectors and experiments. Combining supervised classifiers, anomaly-detection frameworks, and generative models may also provide a robust pipeline for next-generation searches for new physics across different particle physics experiments.

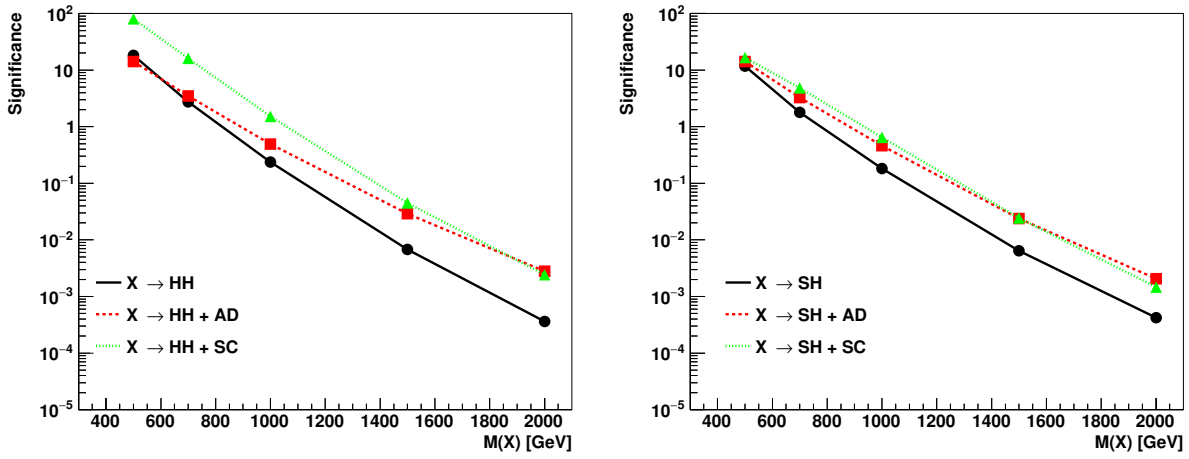

(a) $X \to HH$ (b) $X \to SH$

Figure 3: Comparison of significance values for (a) $X \to HH$ and (b) $X \to SH$ with $S \to HH$ for both supervised cluster method (SC) and unsupervised machine learning based anomaly detection(AD) method. The definitions of signal and background yields, as well as the treatment of uncertainties, follow Ref. [5].

Acknowledgments

We thank the organizers of the Lepton–Photon 2025 conference. WI is supported by the U.S. Department of Energy (DOE) grant DE-SC0017647. SC and NL are supported by UChicago Argonne, LLC, operator of Argonne National Laboratory, under Contract No. DE-AC02-06CH11357 with the U.S. DOE. Argonne's work is supported by the DOE Office of High Energy Physics under the same contract. We also acknowledge the computing resources provided by the Laboratory Computing Resource Center at Argonne National Laboratory.

References

[1] ATLAS Collaboration. "Search for New Phenomena in Two-Body Invariant Mass Distributions Using Unsupervised Machine Learning for Anomaly Detection at $\sqrt{s} = 13$ TeV with the ATLAS Detector". In: *Phys. Rev. Lett.* 132.8 (2024), p. 081801. DOI: 10.1103/PhysRevLett.132.081801.

[2] *ADFilter webpage*. URL: https://mc.hep.anl.gov/adfilter.

[3] Sergei V. Chekanov, Wasikul Islam, Rui Zhang, and Nicholas Luongo. "ADFilter—A Web Tool for New Physics Searches with Autoencoder-Based Anomaly Detection Using Deep Unsupervised Neural Networks". In: *Information* 16.4 (2025), p. 258. DOI: 10.3390/info16040258. arXiv: 2409.03065 [hep-ph].

[4] ATLAS Collaboration. "Search for dijet resonances in events with an isolated charged lepton using $\sqrt{s} = 13$ TeV proton-proton collision data collected by the ATLAS detector". In: *JHEP* 06 (2020), p. 151. DOI: 10.1007/JHEP06(2020)151. arXiv: 2002.11325 [hep-ex].

[5] Sergei V. Chekanov, Wasikul Islam, and Nicholas Luongo. "Enhancing Sensitivity for Di-Higgs Boson Searches Using Anomaly Detection and Supervised Machine Learning Techniques". In: (Apr. 2025). arXiv: 2504.12418 [hep-ph].