


# Machine Learning in HEP

## Lecture III — Anomaly Detection




ChatGPT




UNIVERSITÀ  
DEGLI STUDI  
DI MILANO



BND School 2026 | Heimbach  
Ramon Winterhalder






## Lecture III


---

# Anomaly Detection


# Machine learning for collider physics



# Simulation or data-driven searches



(a) Signal sensitivity



(b) Background specificity

# Community interest in AD




## LHC Olympics

[Kasieczka et al: 2107.02821,  
2101.08320]




## ADC2021

[Govorkova et al: 2107.02157]



## Dark Machines

[Ostdiek et al: 2105.14027]



...and **many** papers:

### 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]

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



# Community interest in AD



## Machine Learning for Anomaly Detection in Particle Physics

Vasilis Belis<sup>1,\*</sup>, Patrick Odagiu<sup>1,\*</sup>, and Thea Klæboe Aarrestad<sup>1,\*</sup>

<sup>1</sup>Institute for Particle Physics and Astrophysics, ETH Zürich, 8093 Zürich, Switzerland

\*e-mail: vbelis@ethz.ch, podagiu@ethz.ch, thea.aarrestad@cern.ch

### ABSTRACT


The detection of out-of-distribution data points is a common task in particle physics. It is used for monitoring complex particle detectors or for identifying rare and unexpected events that may be indicative of new phenomena or physics beyond the Standard Model. Recent advances in Machine Learning for anomaly detection have encouraged the utilization of such techniques on particle physics problems. This review article provides an overview of the state-of-the-art techniques for anomaly detection in particle physics using machine learning. We discuss the challenges associated with anomaly detection in large and complex data sets, such as those produced by high-energy particle colliders, and highlight some of the successful applications of anomaly detection in particle physics experiments.

[2312.14190]

in] 20 Dec 2023

Dark Machines

[Ostdiek et al: 2105.14027]



Available on the CERN CDS information server

CMS PAS EXO-22-026

## CMS Physics Analysis Summary

Contact: cms-pag-conveners-exotica@cern.ch

2024/03/20

Model-agnostic search for dijet resonances with anomalous jet substructure in proton-proton collisions at  $\sqrt{s} = 13$  TeV

The CMS Collaboration

### Abstract

This note introduces a model-agnostic search for new physics in the dijet final state. Other than the requirement of a narrow dijet resonance with a mass in the range of 1800-6000 GeV, minimal additional assumptions are placed on the signal hypothesis. Search regions are obtained by utilizing multivariate machine learning methods to select jets with anomalous substructure. A collection of complementary anomaly detection methods – based on unsupervised, weakly-supervised and semi-supervised algorithms – are used in order to maximize the sensitivity to unknown new physics signatures. These algorithms are applied to data corresponding to an integrated luminosity of  $138 \text{ fb}^{-1}$ , recorded in the years 2016 to 2018 by the CMS experiment at the LHC, at a centre-of-mass energy of 13 TeV. No significant excesses above background expectation are seen, and exclusion limits are derived on the production cross section of benchmark signal models varying in resonance mass, jet mass and jet substructure. Many of these signatures have not previously been searched for at the LHC, making the limits reported on the corresponding benchmark models the first ever and the most stringent to date.


[CMS-PAS-EXO-22-026]



# Two types of anomaly detection


## Outlier Detection (non-resonant)

- Searching for unique and unexpected events
- In HEP, this (might) appear in the tails of dist.



## Overdensities (resonant)


- Analagous to traditional bump hunt



# Two types of anomaly detection


## Outlier Detection (non-resonant)

- Searching for unique and unexpected events
- In HEP, this (might) appear in the tails of dist.




## Overdensities (resonant)

- Analagous to traditional bump hunt




# Autoencoder for non-resonant AD



- AE trained on bg.

$$L = \frac{1}{N} \sum_i (\text{AE}(x_i) - x_i)^2$$




- Use  $L > L_c$  to cut interesting events

[Heimel et al: 1808.08979] [Farina et al: 1808.08992]

- 😊 Fully unsupervised
- 😞 Complexity bias [Finke et al: 2104.09051]
- 😞 not invariant under coordinate transformations [Kasieczka et al: 2209.06225]

# Autoencoder for non-resonant AD



SciPost Physics

Submission

## A Normalized Autoencoder for LHC Triggers

Barry M. Dillon<sup>1</sup>, Luigi Favaro<sup>1</sup>, Tilman Plehn<sup>1</sup>, Peter Sorrenson<sup>2</sup>, and Michael Krämer<sup>3</sup>

<sup>1</sup> Institut für Theoretische Physik, Universität Heidelberg, Germany

<sup>2</sup> Heidelberg Collaboratory for Image Processing, Universität Heidelberg, Germany

<sup>3</sup> Institute for Theoretical Particle Physics and Cosmology (TTK), RWTH Aachen University, Germany

June 23, 2023

[2206.14225]

SciPost Physics

Submission


## Anomalies, Representations, and Self-Supervision

Barry M. Dillon, Luigi Favaro, Friedrich Feiden, Tanmoy Modak, Tilman Plehn

Institut für Theoretische Physik, Universität Heidelberg, Germany

January 13, 2023

[2301.04660]



- Use  $L > L_C$  to cut interesting events

[Heimel et al: 1808.08979] [Farina et al: 1808.08992]

😊 Fully unsupervised


😞 Complexity bias [Finke et al: 2104.09051]

😞 not invariant under coordinate transformations [Kasieczka et al: 2209.06225]

# Two types of anomaly detection


## Outlier Detection (non-resonant)

- Searching for unique and unexpected events
- In HEP, this (might) appear in the tails of dist.




## Overdensities (resonant)

- Analagous to traditional bump hunt



# Resonant AD as a search strategy



# Resonant AD as a search strategy

## Neyman-Pearson Lemma

$$R = \frac{p_{\text{data}}(x)}{p_{\text{bg}}(x)}$$

Optimal  
hypothesis test

❖ Idealized anomaly detector (IAD)

❖ Best you can do **if...**  
...you know  $p_{\text{data}}$  and  $p_{\text{bg}}$

ML

❖ Use  $R$  as cut discriminant  
 $\rightarrow R > R_c$

Sideband

Signal Region

Sideband

Goal: observe new  
physics **signal...**

...above the **SM  
background**

SM Template

Data

$M$

# How to get the optimal test statistic?

16

## Classifier

If we have samples from **data** and **SM background**...

...an **optimal classifier** yields

$$f(x) = \frac{p_{\text{data}}(x)}{p_{\text{data}}(x) + p_{\text{bg}}(x)}$$

❖ Get  $x \sim p_{\text{data}}$  and  $x \sim p_{\text{bg}}$  from **MC simulations**

❖ Estimate samples from **data**:

$$x \sim p_{\text{data}}(x | \text{SR})$$

$$x \sim p_{\text{data}}(x | \text{SB}) \approx p_{\text{bg}}(x)$$

$$R = \frac{p_{\text{data}}(x)}{p_{\text{bg}}(x)}$$

## Density estimator

Instead of learning the **likelihood ratio** directly...

...use a **density estimator** to learn

$$p_{\omega}(x | \text{SR}) \simeq p_{\text{data}}(x | \text{SR})$$

$$p_{\omega}(x | \text{SB}) \simeq p_{\text{bg}}(x)$$

❖ Then **calculate**  $R$  directly from the individual likelihoods

# Example I

---

## CWoLa Hunting

# Reminder — Classification Problem


**Goal:** learn the signal to background ratio

An optimal classifier yields the likelihood ratio

$$R_{\text{optimal}} = \frac{f(x)}{1 - f(x)} = \frac{p_{\text{sig}}(x)}{p_{\text{bg}}(x)}$$

⊕ Can be approximated with a **supervised classifier (ML)**

⊖ Labels **are not available** in experimental data



# Classification without labels (CWoLa)

Two **mixed datasets** with signal fractions  $w_i$

$$p_i(x) = w_i p_{\text{sig}}(x) + (1 - w_i) p_{\text{bg}}(x)$$


Classifier gives likelihood ratio

$$R_{\text{mixed}} = \frac{w_1 R_{\text{optimal}}(x) + (1 - w_1)}{w_2 R_{\text{optimal}}(x) + (1 - w_2)}$$


⊕ Monotonic function

→ optimal on mixed = optimal on pure sample

→ Basis of **weak supervised classification**



# Supervised versus IAD




$$R_{\text{IAD}} = \frac{p_{\text{data}}(x)}{p_{\text{bg}}(x)}$$
$$= \epsilon R_{\text{supervised}} + (1 - \epsilon)$$

$$R_{\text{supervised}} = \frac{p_{\text{sig}}(x)}{p_{\text{bg}}(x)}$$

# CWoLa Hunting

21



Resonant  
observable

$$m_{jj} = m_{Z'} > m_X, m_Y$$


Other  
features

$$x = \{m_X, m_Y, \Delta m_j, \tau_{21}^{(1)}, \tau_{21}^{(2)}\}$$

$$p_{\text{bg}}(x | m_{jj} \in \text{SR}) \approx p_{\text{bg}}(x | m_{jj} \in \text{SB}) \approx p_{\text{bg}}(x)$$

LHC Olympics


[Kasieczka et al: 2107.02821,  
2101.08320]



[1902.02634]

# CWoLa Hunting

22



**Resonant  
observable**

$$m_{jj} = m_{Z'} > m_X, m_Y$$


**Other  
features**

$$x = \{m_X, m_Y, \Delta m_j, \tau_{21}^{(1)}, \tau_{21}^{(2)}\}$$

$$p_{\text{bg}}(x | m_{jj} \in \text{SR}) \approx p_{\text{bg}}(x | m_{jj} \in \text{SB}) \approx p_{\text{bg}}(x)$$


**CWoLa Likelihood estimate**

$$R_{\text{CWoLa}} = \frac{p_{\text{data}}(x | \text{SR})}{p_{\text{bg}}(x | \text{SB})}$$



# CWoLa Hunting

22



**Resonant  
observable**

$$m_{jj} = m_{Z'} > m_X, m_Y$$


**Other  
features**

$$x = \{m_X, m_Y, \Delta m_j, \tau_{21}^{(1)}, \tau_{21}^{(2)}\}$$

$$p_{\text{bg}}(x | m_{jj} \in \text{SR}) \approx p_{\text{bg}}(x | m_{jj} \in \text{SB}) \approx p_{\text{bg}}(x)$$

**CWoLa Likelihood estimate**

$$R_{\text{CWoLa}} = \frac{p_{\text{data}}(x | \text{SR})}{p_{\text{bg}}(x | \text{SB})} \approx \frac{p_{\text{data}}(x | \text{SR})}{p_{\text{bg}}(x | \text{SR})}$$



**Can we do better?**

## Example II

---

# Anomaly detection with Density Estimation (ANODE)


## CWoLa Likelihood estimate

$$R_{\text{CWoLa}} = \frac{p_{\text{data}}(x | \text{SR})}{p_{\text{bg}}(x | \text{SB})}$$

## The ANODE method

$p_{\omega_0}(x | m) \simeq p_{\text{bg}}(x | m)$ 
← NF
Trained in  $m \in \text{SB}$

$p_{\omega_1}(x | m) \simeq p_{\text{data}}(x | m)$ 
← NF
Trained in  $m \in \text{SR}$



## CWoLa Likelihood estimate

$$R_{\text{CWoLa}} = \frac{p_{\text{data}}(x | \text{SR})}{p_{\text{bg}}(x | \text{SB})}$$




## ANODE Likelihood estimate

$$R_{\text{ANODE}} = \frac{p_{\omega_1}(x | \text{SR})}{p_{\omega_0}(x | \text{SR})} \simeq \frac{p_{\text{data}}(x | \text{SR})}{p_{\text{bg}}(x | \text{SR})}$$

## The ANODE method

$$p_{\omega_0}(x | m) \simeq p_{\text{bg}}(x | m) \quad \text{Trained in } m \in \text{SB}$$

$$p_{\omega_1}(x | m) \simeq p_{\text{data}}(x | m) \quad \text{Trained in } m \in \text{SR}$$



**Are we already happy?**

# CWoLA versus ANODE

27

## CWoLa Likelihood estimate

$$R_{\text{CWoLa}} = \frac{p_{\text{data}}(x | \text{SR})}{p_{\text{bg}}(x | \text{SB})}$$

Pros and cons:

- ⊕ Classification is easy and precise
- ⊖ Sensitive to correlations between  $m_{jj}$  and other features  $x$

[1902.02634]



## ANODE Likelihood estimate

$$R_{\text{ANODE}} = \frac{p_{\omega_1}(x | \text{SR})}{p_{\omega_0}(x | \text{SR})}$$

Pros and cons:

- ⊕ Robust against correlations
- ⊖ Less powerful and sensitive than classification

[2001.04990]

**Can we get the best of both worlds?**

## Example III

---

### Classifying Anomalies Through Outer Density Estimation (CATHODE)

# Best of both worlds — CATHODE

30

## The CATHODE method

$$p_{\omega_0}(x|m) \simeq p_{\text{bg}}(x|m) \quad \text{Trained in } m \in \text{SB}$$

~~$$p_{\omega_1}(x|m) \simeq p_{\text{data}}(x|m)$$~~

1. Interpolate **SM background template** to SR and sample:


$$\hat{x}_{\text{bg}} \sim p_{\omega_0}(x|m \in \text{SR}) \simeq p_{\text{bg}}(x|\text{SR})$$

2. Then **train classifier** between

$$\hat{x}_{\text{bg}} \text{ and } x \sim p_{\text{data}}(x|\text{SR}) \text{ as in } \text{CWoLA}$$

## CATHODE Likelihood estimate


$$R_{\text{CATHODE}} = \frac{p_{\text{data}}(x|\text{SR})}{p_{\omega_0}(x|\text{SR})} \simeq \frac{p_{\text{data}}(x|\text{SR})}{p_{\text{bg}}(x|\text{SR})}$$



**How do they compare?**



# How to quantify improvement?

32



Statistical significance:  $\frac{S}{\sqrt{B}}$   $\xrightarrow{\text{AD}}$   $\frac{S \cdot \epsilon_S}{\sqrt{B \cdot \epsilon_B}} = \frac{S}{\sqrt{B}} \cdot \boxed{\frac{\epsilon_S}{\sqrt{\epsilon_B}}}$   $\xleftarrow{\text{Improvement factor}}$





# Results — Comparison



**Are there other ways?**



# ML techniques to construct SM template

35



# ML techniques to construct SM template

35



## The Interplay of Machine Learning-based Resonant Anomaly Detection Methods

Tobias Golling,<sup>a</sup> Gregor Kasieczka,<sup>b</sup> Claudius Krause,<sup>c</sup> Radha Mastandrea,<sup>d,e</sup> Benjamin Nachman,<sup>e,f</sup> John Andrew Raine,<sup>a</sup> Debajyoti Sengupta,<sup>a</sup> David Shih,<sup>g</sup> and Manuel Sommerhalder<sup>b</sup>

<sup>a</sup>Département de physique nucléaire et corpusculaire, Université de Genève, 1211 Genève, Switzerland

<sup>b</sup>Institut für Experimentalphysik, Universität Hamburg, 22761 Hamburg, Germany

<sup>c</sup>Institut für Theoretische Physik, Universität Heidelberg, 69120 Heidelberg, Germany

<sup>d</sup>Department of Physics, University of California, Berkeley, CA 94720, USA

<sup>e</sup>Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA


<sup>f</sup>Berkeley Institute for Data Science, University of California, Berkeley, CA 94720, USA

<sup>g</sup>NHETC, Dept. of Physics and Astronomy, Rutgers University, Piscataway, NJ 08854, USA

E-mail: [tobias.golling@unige.ch](mailto:tobias.golling@unige.ch), [gregor.kasieczka@uni-hamburg.de](mailto:gregor.kasieczka@uni-hamburg.de),  
[claudius.krause@thphys.uni-heidelberg.de](mailto:claudius.krause@thphys.uni-heidelberg.de), [rmastand@berkeley.edu](mailto:rmastand@berkeley.edu), [bpnachman@lbl.gov](mailto:bpnachman@lbl.gov),  
[john.raine@unige.ch](mailto:john.raine@unige.ch), [debajyoti.sengupta@unige.ch](mailto:debajyoti.sengupta@unige.ch), [shih@physics.rutgers.edu](mailto:shih@physics.rutgers.edu),  
[manuel.sommerhalder@uni-hamburg.de](mailto:manuel.sommerhalder@uni-hamburg.de)

**ABSTRACT:** Machine learning-based anomaly detection (AD) methods are promising tools for extending the coverage of searches for physics beyond the Standard Model (BSM). One class of AD methods that has received significant attention is resonant anomaly detection, where the BSM physics is assumed to be localized in at least one known variable. While there have been many methods proposed to identify such a BSM signal that make use of simulated or detected data in different ways, there has not yet been a study of the methods' complementarity. To this end, we address two questions. First, in the absence of any signal, do different methods pick the same events as signal-like? If not, then we can significantly reduce the false-positive rate by comparing different methods on the same dataset. Second, if there is a signal, are different methods fully correlated? Even if their maximum performance is the same, since we do not know how much signal is present, it may be beneficial to combine approaches. Using the Large Hadron Collider (LHC) Olympics dataset, we provide quantitative answers to these questions. We find that there are significant gains possible by combining multiple methods, which will strengthen the search program at the LHC and beyond.

[2307.11157]



**Can we do even better?**

## Example IV


---

# Residual ANODE (R-ANODE)

## The ANODE method

$$p_{\omega_0}(x|m) \simeq p_{\text{bg}}(x|m) \quad \text{Trained in } m \in \text{SB}$$

$$p_{\omega_1}(x|m) \simeq p_{\text{data}}(x|m) \quad \text{Trained in } m \in \text{SR}$$



## The ANODE method

$$p_{\omega_0}(x|m) \simeq p_{\text{bg}}(x|m) \quad \text{Trained in } m \in \text{SB}$$

~~$$p_{\omega_1}(x|m) \simeq p_{\text{data}}(x|m) \quad \text{Trained in } m \in \text{SR}$$~~



## The R-ANODE method

$$p_{\text{data}}(x|\text{SR}) = w p_{\text{sig}}(x|\text{SR}) + (1 - w) p_{\text{bg}}(x|\text{SR})$$

↑  
NF

$$p_{\omega_s}(x|m)$$

Trained in  $m \in \text{SR}$


↑  
NF

$$p_{\omega_0}(x|m)$$



Trained in  $m \in \text{SB}$

## R-ANODE Likelihood estimate



$$R_{\text{R-ANODE}} = \frac{p_{\omega_s}(x|\text{SR})}{p_{\omega_0}(x|\text{SR})} \simeq \frac{p_{\text{sig}}(x|\text{SR})}{p_{\text{bg}}(x|\text{SR})}$$



# Results — R-ANODE



# Results — R-ANODE



**More sources?**

# ML lecture notes



## Modern Machine Learning for LHC Physicists

Tilman Plehn<sup>a,\*</sup>, Anja Butter<sup>a,b</sup>, Barry Dillon<sup>a</sup>,  
Theo Heimel<sup>a</sup>, Claudius Krause<sup>c</sup>, and Ramon Winterhalder<sup>d</sup>

<sup>a</sup> Institut für Theoretische Physik, Universität Heidelberg, Germany

<sup>b</sup> LPNHE, Sorbonne Université, Université Paris Cité, CNRS/IN2P3, Paris, France

<sup>c</sup> HEPHY, Austrian Academy of Sciences. Vienna, Austria

<sup>d</sup> CP3, Université catholique de Louvain, Louvain-la-Neuve, Belgium

April 15, 2024

### Abstract

Modern machine learning is transforming particle physics fast, bullying its way into our numerical tool box. For young researchers it is crucial to stay on top of this development, which means applying cutting-edge methods and tools to the full range of LHC physics problems. These lecture notes lead students with basic knowledge of particle physics and significant enthusiasm for machine learning to relevant applications. They start with an LHC-specific motivation and a non-standard introduction to neural networks and then cover classification, unsupervised classification, generative networks, and inverse problems. Two themes defining much of the discussion are well-defined loss functions and uncertainty-aware networks. As part of the applications, the notes include some aspects of theoretical LHC physics. All examples are chosen from particle physics publications of the last few years.<sup>1</sup>

- Check lecture notes for more **details and applications**



iv:2211.01421v3 [hep-ph] 12 Apr 2024

CS:vi

# Summary and Outlook

## Take-home messages

- ML beneficial in **every step** of the **simulation** and **analysis chain**
- We find both **proof-of-concepts** as well as established use cases (→ **AD, MadNIS,...**)
- Interesting **interplay** between **physics** and **ML**
  - Physics provides **~infinite data** for ML
  - Physics requirements (**precision, symmetries,...**) **different** than industry applications

## Future exercises

- Full integration of ML-based methods into standard tools → **Taggers, MadGraph,...**
- Make everything run on **GPUs** and make it **differentiable**
- Foster deeper collaboration between **theory, experiment, and ML** community

