Anomaly Detection in Particle Physics

Katherine Fraser

Department of Physics
Harvard University
and NSF IAIFI
kfraser@g.harvard.edu

LHCP 2023

Why anomaly detection?

Typical Searches

- Looking for a specific, physics motivated signal
- Maximum sensitivity for a specific model
- Not useful for other models

Anomaly Detection

- Goal is to be model agnostic
- Looking for deviations from background only
- Less sensitive to any specific model, but can look for multiple different models
- Can be at the event level, but not always (ex. jets)

Community Interest

There is substantial community interest, including through challenges:

LHC Olympics

[Kasieczka et al: 2107.02821, 2101.08320]

ADC2021

[Govorkova et al: 2107.02157]

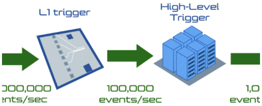

Dark Machines

[Ostdiek et al: 2105.14027]

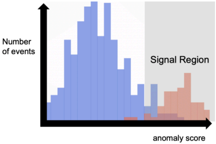

and many papers:

- 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]
  - 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

[https://iml-wg.github.io/HEPML-LivingReview/]

Two Types of Anomaly Detection

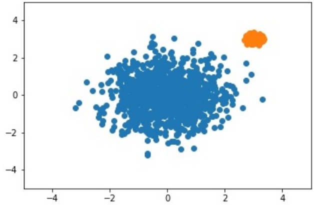

Outlier Detection
[Nonresonant]

- Searching for unique or unexpected events
- In HEP, this is the tails of distributions

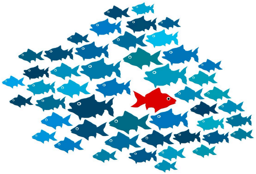

[1807.10261, 1808.08979, 1808.08992, 1811.10276, 1903.02032, 1912.10625, 2004.09360, 2006.05432, 2007.01850, 2007.15830, 2010.07940, 2102.08390, 2104.09051, 2105.07988, 2105.10427, 2105.09274, 2106.10164, 2108,03986, 2109.10919, 2110.06948, 2112.04958, 2203.01343,2206.14225, 2304.03836, …]

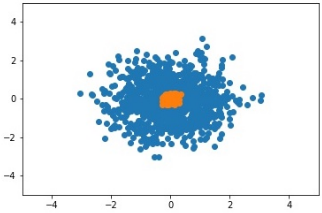

Finding Overdensities
[Resonant]

- Analagous to the traditional bump hunt

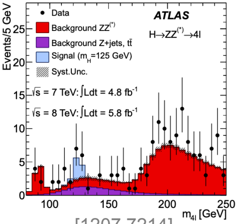

[1207.7214]

[1805.02664, 1806.02350, 1902.02634, 1912.12155, 2001.05001, 2001.04990, 2012.11638, 2106.10164, 2109.00546, 2202.00686, 2203.09470, 2208.05484, 2210.14924, 2212.11285, ….]

Two Types of Anomaly Detection

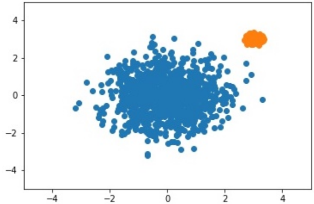

Outlier Detection
[Nonresonant]

- Searching for unique or unexpected events
- In HEP, this is the tails of distributions

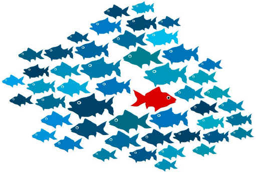

[1807.10261, 1808.08979, 1808.08992, 1811.10276, 1903.02032, 1912.10625, 2004.09360, 2006.05432, 2007.01850, 2007.15830, 2010.07940, 2102.08390, 2104.09051, 2105.07988, 2105.10427, 2105.09274, 2106.10164, 2108,03986, 2109.10919, 2110.06948, 2112.04958, 2203.01343,2206.14225, 2304.03836, …]

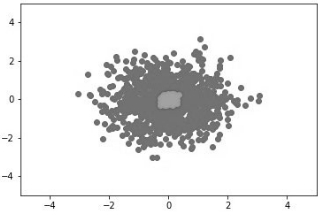

Finding Overdensities
[Resonant]

- Analagous to the traditional bump hunt

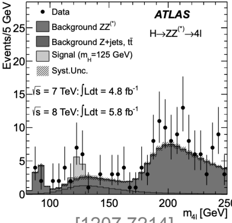

[1207.7214]

[1805.02664, 1806.02350, 1902.02634, 1912.12155, 2001.05001, 2001.04990, 2012.11638, 2106.10164, 2109.00546, 2202.00686, 2203.09470, 2208.05484, 2210.14924, 2212.11285, ….]

Autoencoders (AEs)

AEs work by learning compression to a latent space which preserves the original information.

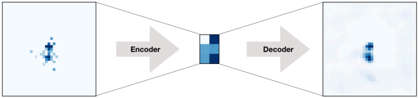

[Hajer et al: 1807.10261]
[Heimel et al: 1808.08979]
[Farina et al: 1808.08992]

The reconstruction fidelity gives an anomaly score.

Variational AEs (VAEs) add a stochastic component by having the decoder sample from latent space. There are multiple different choices for anomaly score.

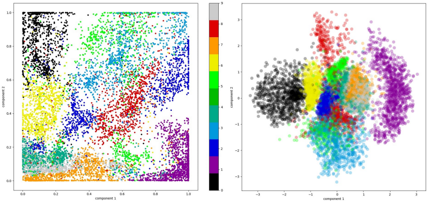

[Cerri et al: 1811.10276]

[Hajer et al: 1807.10261, Roy, Vijay: 1903.02032, Cheng et al: 2007.01850, Beekveld et al: 2010.07940, Batson et al: 2102.08380, Finke et al: 2104.09051, Govorkova et al: 2108.03986, Collins: 2109.10919, **Fraser et al: 2110.06948**, Ngairangbam et al: 2112.04958, Dillon et al: 2206.14225, Roche et al: 2304.03836,…]

Using Optimal Transport (OT)

- OT is a more physical alternative to look for outliers.
- OT is the minimum "effort" required to transform one event into another. Ex: Energy Movers Distance (EMD).

[Komiske et al: 1902.02346, 2004.04159]

- Can turn into an anomaly score by picking reference samples. Ex: average/medoid jets, reference events.

[Romano et al: 2004. 09360]
[Cai et al: 2008.08604]
[Fraser et al: 2110.06948]
[Buss et al: 2202.00686]

- For reference jets, correlated with VAE latent space distances.

[Fraser et al: 2110.06948]

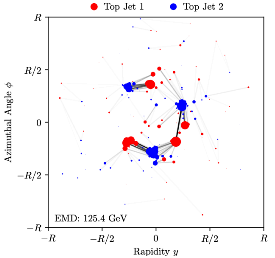

Example OT Plan
[Komiske et al: 1902.02346]

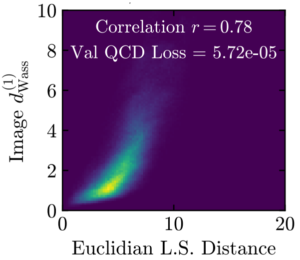

Correlation with VAE latent space
[Fraser et al: 2110.06948]

Challenges with Outlier Detection

1. It’s difficult to pick a metric to compare methods.
   [Ostdiek et al: 2105.14027]
2. Ideal optimization (input representation, architecture) is sensitive to signal. Signal sensitivity can be much weaker than supervised searches.
   [**Fraser et al: 2110.06948**]
   [Jawahar et al: 2110.08508]
3. Results are strongly dependent on background.
   [Finke et al: 2104. 09051]
4. Not invariant under feature space transformations.
   [Kasieczka et al: 2209.06225]
5. Unclear how to use selected events for analyses without a reliable background estimate.

Outlier detection has potential to be especially useful for triggering, so we would like to resolve these problems!

[Govorkova et al: 2107.02157, 2108.03986, Duarte et al: 2207.07958]

Solving (Some) Problems with Outlier Detection

- Weakly-supervised approaches using exposure to outliers/ potential signals (examples: OE-VAE, QUAK, OT with multiple samples) [Cheng et al: 2007.01850, Khosa, Sanz: 2007.14462, Park et al: 2011.03550, Gonski et al: 2108.13451, **Fraser et al: 2110.06948,** Caron et al: 2207.07631]
- Engineering better networks with less background dependence [Blance et al: 1905.10384] [Finke et al: 2104. 09051] [Dillon et al: 2206.14225]
- Picking smarter (self-supervised) representations [Buss et al: 2202.00686] [Park et al: 2208.05484] [Dillon et al: 2301.04660]
- Using multiple decorrelated AEs and the ABCD method to get a background estimate. [Mikuni et al: 2111.06417]

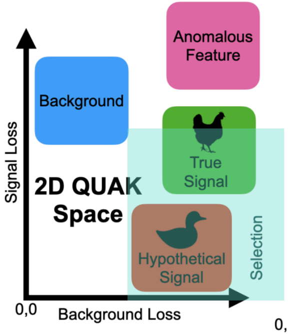

QUAK Space
[Park et al: 2011.03550]

Two Types of Anomaly Detection

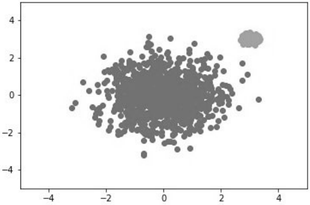

Outlier Detection

[Nonresonant]

- Searching for unique or unexpected events
- In HEP, this is the tails of distributions

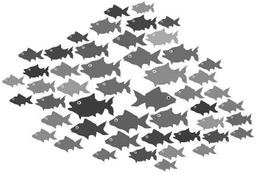

[1807.10261, 1808.08979, 1808.08992, 1811.10276, 1903.02032, 1912.10625, 2004.09360, 2006.05432, 2007.01850, 2007.15830, 2010.07940, 2102.08390, 2104.09051, 2105.07988, 2105.10427, 2105.09274, 2106.10164, 2108,03986, 2109.10919, 2110.06948, 2112.04958, 2203.01343,2206.14225, 2304.03836, …]

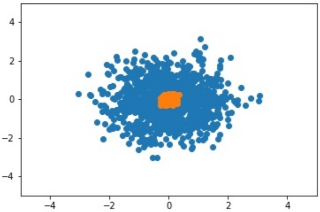

Finding Overdensities

[Resonant]

- Analagous to the traditional bump hunt

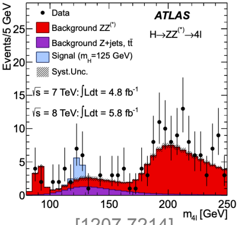

[1207.7214]

[1805.02664, 1806.02350, 1902.02634, 1912.12155, 2001.05001, 2001.04990, 2012.11638, 2106.10164, 2109.00546, 2202.00686, 2203.09470, 2208.05484, 2210.14924, 2212.11285, ….]

The CWOLA Bump Hunt

CWOLA is weakly supervised classification:

- Trained on two data samples with different signal fractions
- Classifier is also optimal for distinguishing signal vs background because optimal classifier is the likelihood ratio

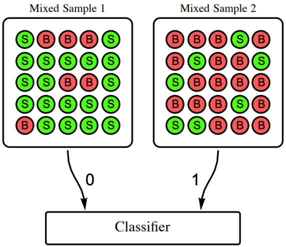

[Metodiev et al: 1708.02949]

CWOLA can also be used for a weakly supervised bump hunt:

- Train a classifier between signal region and side bands
- Apply a threshold cut on the classifier output and perform a bump hunt

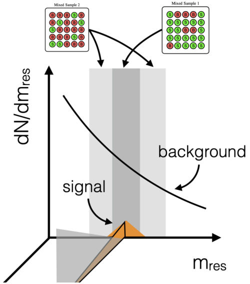

[Image: Ben Nachman Talk]
[Collins et al: 1902.02634]

Improving Unsupervised Bump Hunts

- ANODE: interpolates probability densities from sidebands to the signal-region & constructs likelihood ratio
  [Nachman, Shih: 2001.04990]
- CATHODE: samples from the background model in signal region after interpolating and estimates likelihood ratio with classifier
  [Hallin et al: 2109.00546]
- LaCATHODE: Use a flow to perform CATHODE in latent space
  [Hallin et al: 2210.14924]

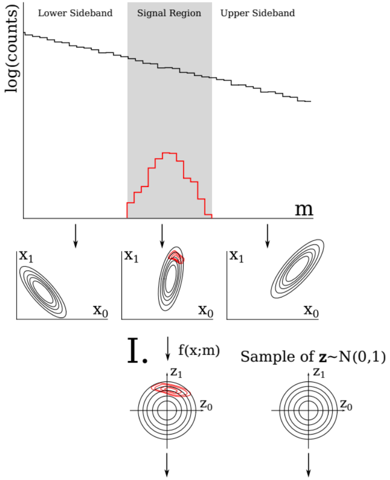

LaCATHODE
[Hallin et al: 2210.14924]

More Unsupervised Bump Hunts

- SALAD: Reweight simulation to match sidebands, then interpolate into the signal region and use a second classifier to get the likelihood ratio [Andreassen et al: 2001.05001]
- CURTAINS: Train an invertible neural network conditioned on mass to map between sidebands [Raine et al: 2203.09470]
- FETA: Map simulation to data in sidebands, then compare to SR data [Golling et al: 2212.11285]

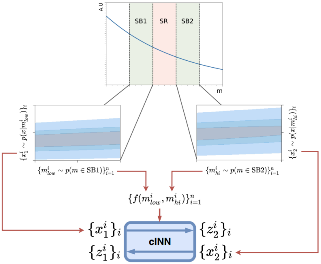

CURTAINS
[Raine et al: 2203.09470]

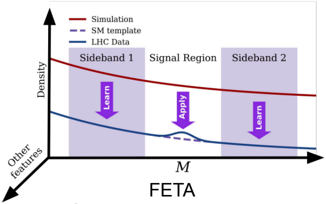

FETA
[Golling et al: 2212.11285]

Methods for Both Resonances and Tails

Some strategies can be used for both types of anomaly detection.

[D'Agnolo, Wulzer: 1806.02350]
[De Simone, Jacques: 1807.06038]

However, these are often strongly dependent on simulation because they are directly comparing to it

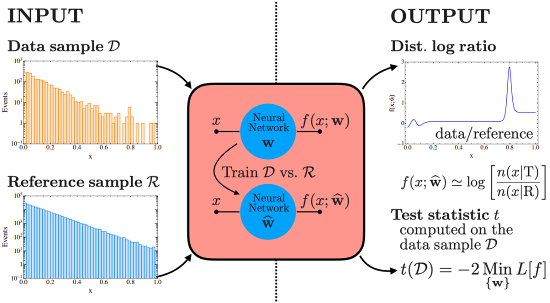

Learning New Physics from a Machine
[D'Agnolo, Wulzer: 1806.02350]

Summary

Anomaly detection can either search for resonant signals (overdensities) or non-resonant signals (outliers).

There are general challenges with outlier detection, though some of these challenges can be overcome with engineering. Outlier detection is potentially useful for triggering.

There are many methods for unsupervised bump hunts that are complementary for different data sets and resonances.

There is substantial ongoing work in anomaly detection, and its exciting to see it starting to be used in experimental results.

[ATLAS: 2005.02983, ATLAS-CONF-2022-045, ATLAS-CONF-2023-022, CMS-DP-2022-021, CMS-DP-2022-043….]