Universität Hamburg

DER FORSCHUNG | DER LEHRE | DER BILDUNG

PUNCH 4NFDI

Anomaly detection with ML in HEP

Anna Hallin, Inst. for Experimental physics, University of Hamburg; PUNCH TA3-WP3

PUNCH4NFDI TA5 - XFEL Joint Workshop on Machine Learning and Data Processing on FPGAs

2023-06-15

ML in HEP: a growing field

- Inspire search for ("machine learning" or "deep learning" or "neural") and (hep-ex or hep-ph or hep-th)
- 420 papers in 2022, 170 so far in 2023

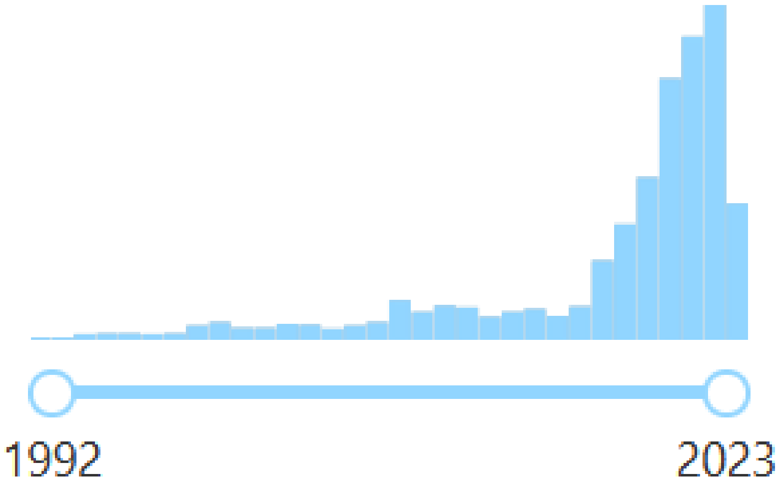

- Learning New Physics from a Machine [DOI]
- Anomaly Detection for Resonant New Physics with Machine Learning [DOI]
- Extending the search for new resonances with machine learning [DOI]
- Learning Multivariate New Physics [DOI]
- Searching for New Physics with Deep Autoencoders [DOI]
- QCD or What? [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
- The Data-Directed Paradigm for BSM searches
- 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 from mass unspecific jet tagging
- A method to challenge symmetries in data with self-supervised learning
- Stressed GANs snag desserts, a.k.a Spotting Symmetry Violation with Symmetric Functions
- Online-compatible Unsupervised Non-resonant Anomaly Detection
- Event-based anomaly detection for new physics searches at the LHC using machine learning
- Learning New Physics from an Imperfect Machine
- Autoencoders for Semivisible Jet Detection
- Anomaly detection in high-energy physics using a quantum autoencoder
- Creating Simple, Interpretable Anomaly Detectors for New Physics in Jet Substructure
- Taming modeling uncertainties with Mass Unspecific Supervised Tagging
- What's Anomalous in LHC Jets?
- Quantum Anomaly Detection for Collider Physics
- Self-supervised Anomaly Detection for New Physics
- Data-directed search for new physics based on symmetries of the SM [DOI]
- CURTAINs for your Sliding Window: Constructing Unobserved Regions by Transforming Adjacent Intervals
- Learning new physics efficiently with nonparametric methods
- ''Flux+Mutability'': A Conditional Generative Approach to One-Class Classification and Anomaly Detection
- Boosting mono-jet searches with model-agnostic machine learning [DOI]
- Event Generation and Density Estimation with Surjective Normalizing Flows
- A Normalized Autoencoder for LHC Triggers
- Mixture-of-theories Training: Can We Find New Physics and Anomalies Better by Mixing Physical Theories?
- Neural Embedding: Learning the Embedding of the Manifold of Physics Data
- Null Hypothesis Test for Anomaly Detection
- Resonant anomaly detection without background sculpting
- Anomaly Detection under Coordinate Transformations
- Quantum-probabilistic Hamiltonian learning for generative modelling \& anomaly detection
- Efficiently Moving Instead of Reweighting Collider Events with Machine Learning
- Nanosecond anomaly detection with decision trees for high energy physics and real-time application to exotic Higgs decays
- The Mass-ive Issue: Anomaly Detection in Jet Physics

Anomaly detection section in the HEP-ML Living Review

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

What are anomalies?

- An outlier is an observation which deviates so much from the other observations as to arouse suspicions that it was generated by a different mechanism.[2]
- Anomalies are instances or collections of data that occur very rarely in the data set and whose features differ significantly from most of the data.
- An outlier is an observation (or subset of observations) which appears to be inconsistent with the remainder of that set of data.[3]
- An anomaly is a point or collection of points that is relatively distant from other points in multi-dimensional space of features.
- Anomalies are patterns in data that do not conform to a well defined notion of normal behaviour.[1]
- Let T be observations from a univariate Gaussian distribution and O a point from T. Then the z-score for O is greater than a pre-selected threshold if and only if O is an outlier.

Wikipedia: Anomaly detection

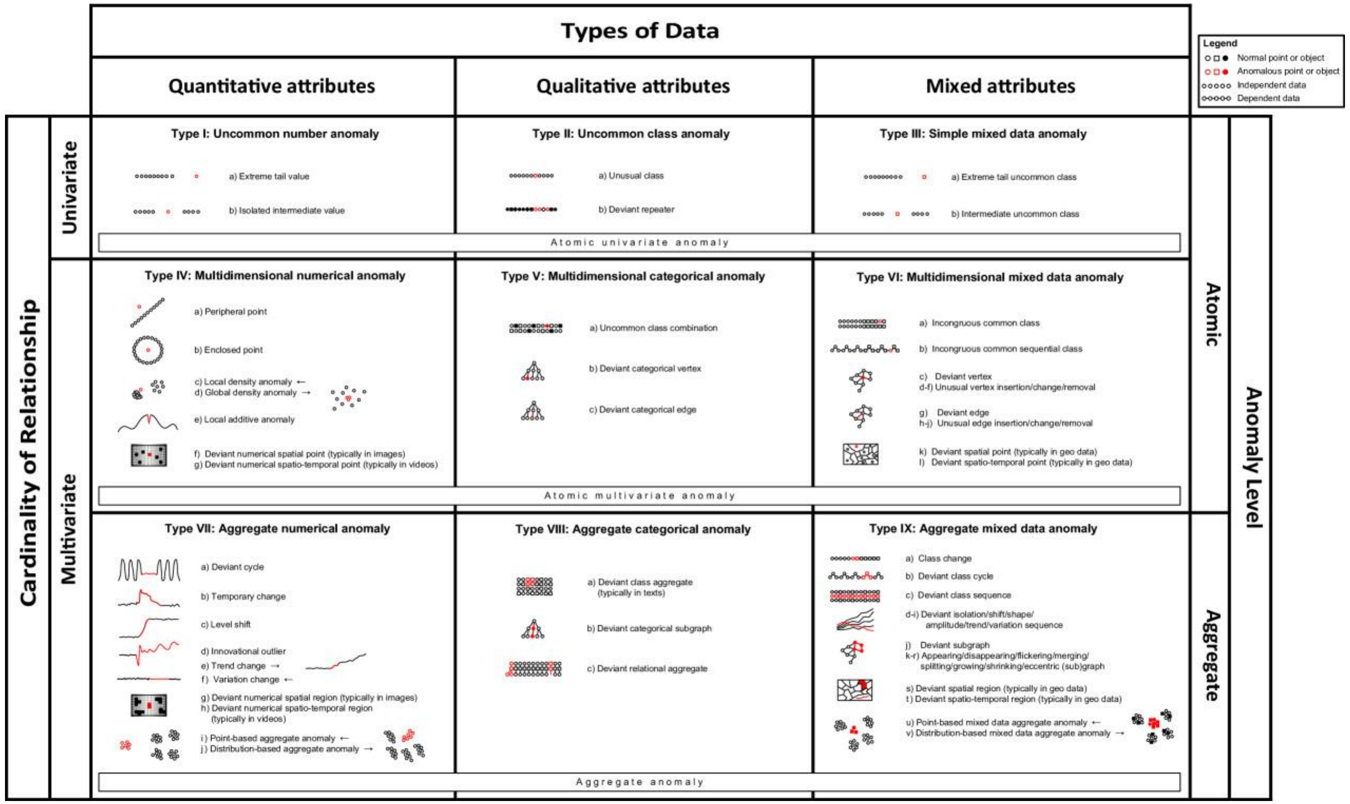

R. Foorthuis, On the nature and types of anomalies: a review of deviations in data

Outliers, or point anomalies

Individual examples are far away from the regular distribution

Examples:

- Detector malfunctions
- Background-free search

Group anomalies

Individual examples in themselves are normal/uninteresting, but as a group they form an interesting entity.

Examples:

- New physics searches, e.g. resonances
- Excess in time series

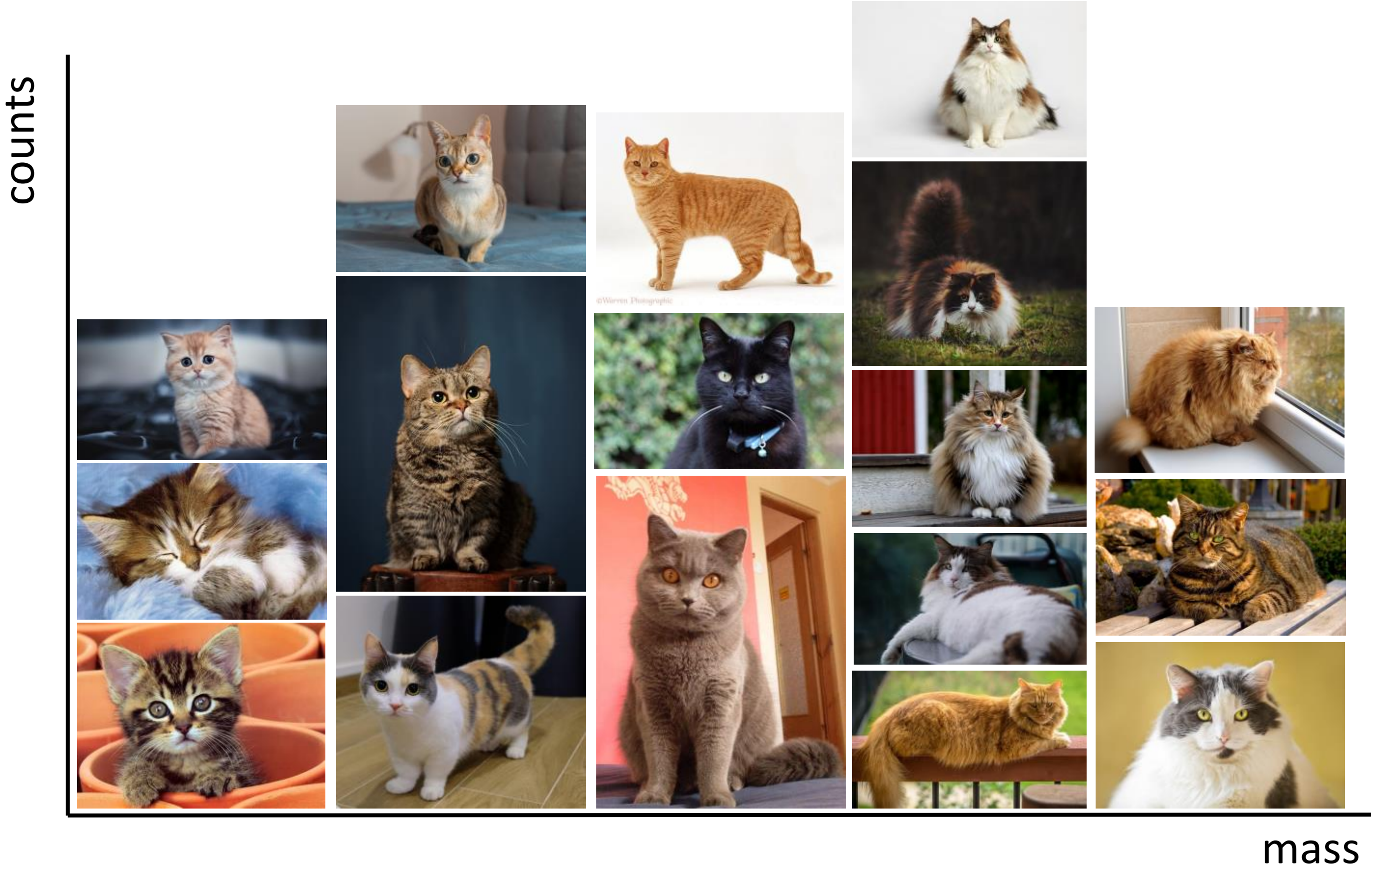

New physics searches at the LHC

- We know there must be new physics beyond the Standard Model
- New physics could show up as anomalies in the data
- The experiments at the LHC have set many limits on popular models of new physics, but so far no discoveries

smbc-comics.com

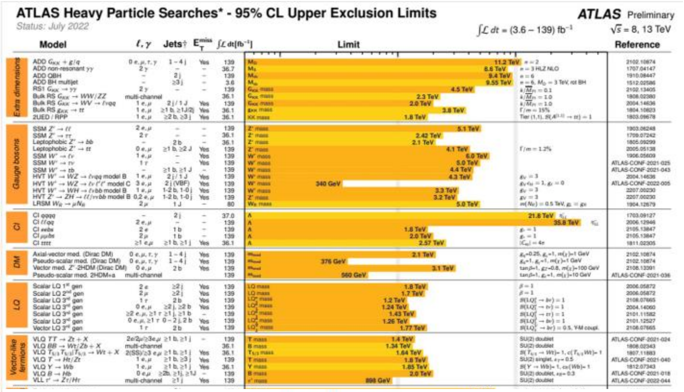

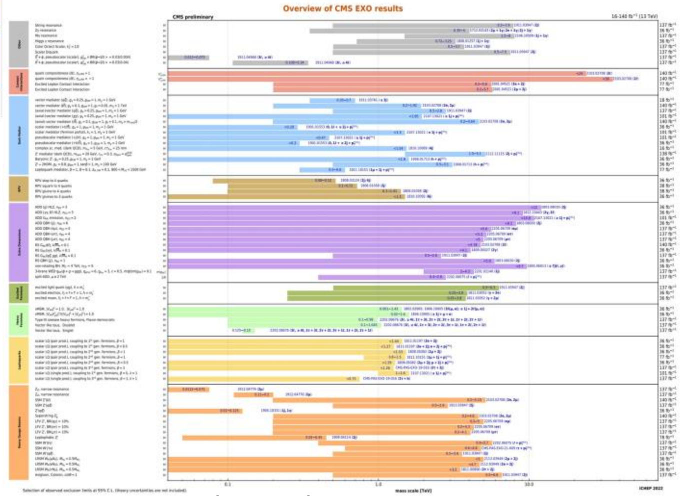

atlaspo.cern.ch/public/summary_plots
twiki.cern.ch/twiki/bin/view/CMSPublic/SummaryPlotsEXO13TeV

Anomaly detection methods I

- Most searches at the LHC:
  - Simulate signal and background
  - Optimize search strategy based on this
  - Maximally model dependent

- Some searches:
  - Signal clear but background messy
  - Simulate signal
  - Compare against data

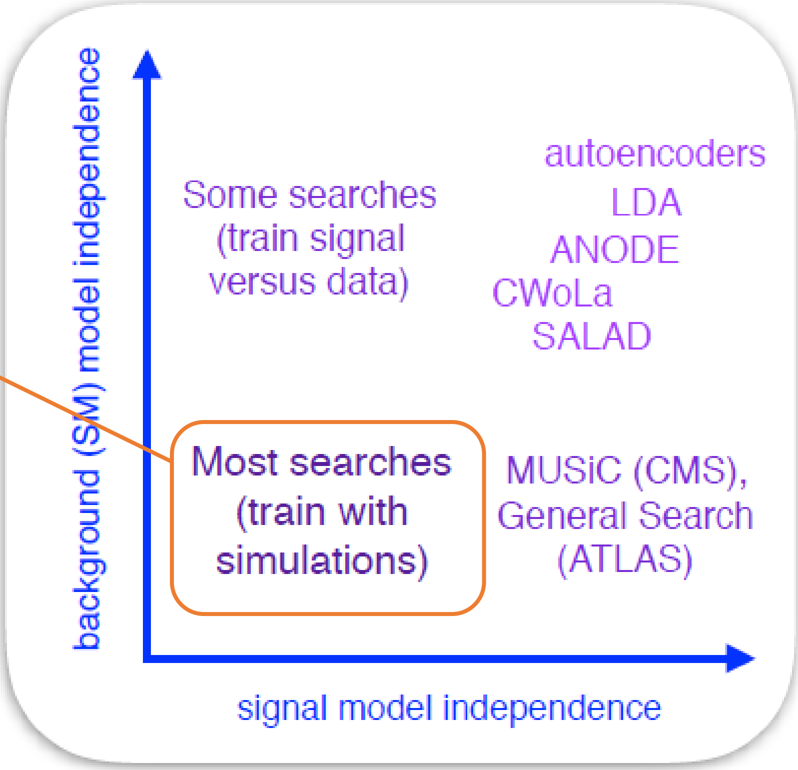

2001.04990

Anomaly detection methods I

- Most searches at the LHC:
  - Simulate signal and background
  - Optimize search strategy based on this
  - Maximally model dependent

- Some searches:
  - Signal clear but background messy
  - Simulate signal
  - Compare against data

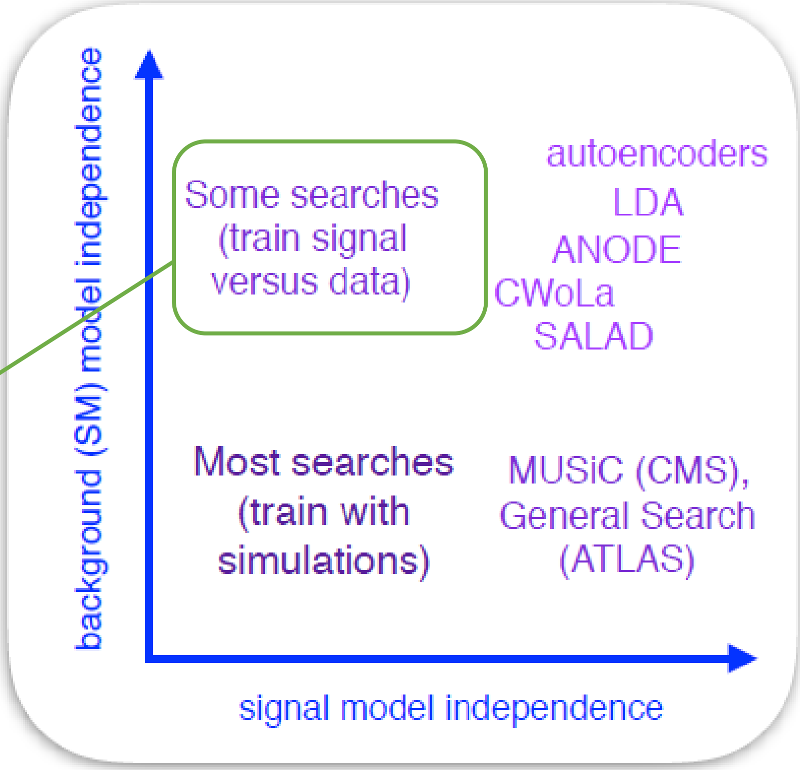

2001.04990

Model space is big

- A search at the LHC is a long process. Can we cover all possible models with dedicated searches? Probably not.
- We want to complement the dedicated searches with more model agnostic searches.

Image generated by midjourney.com

Anomaly detection methods II

- Signal model independent:
  - Simulate background
  - Directly compare data vs simulation over several final states

- Model ”agnostic” or ”independent”:
  - Does not assume a specific signal model
  - Does not rely on simulations
  - Does make assumptions about signal shape or optimal features to search in

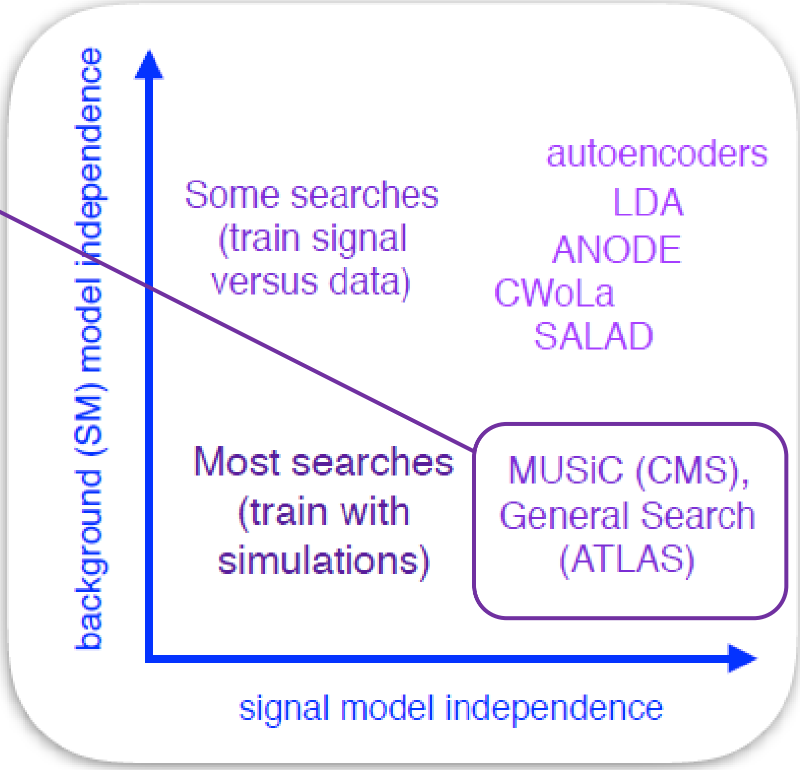

2001.04990

Anomaly detection methods II

- Signal model independent:
  - Simulate background
  - Directly compare data vs simulation over several final states
- Model ”agnostic” or ”independent”:
  - Does not assume a specific signal model
  - Does not rely on simulations
  - Does make assumptions about signal shape or optimal features to search in

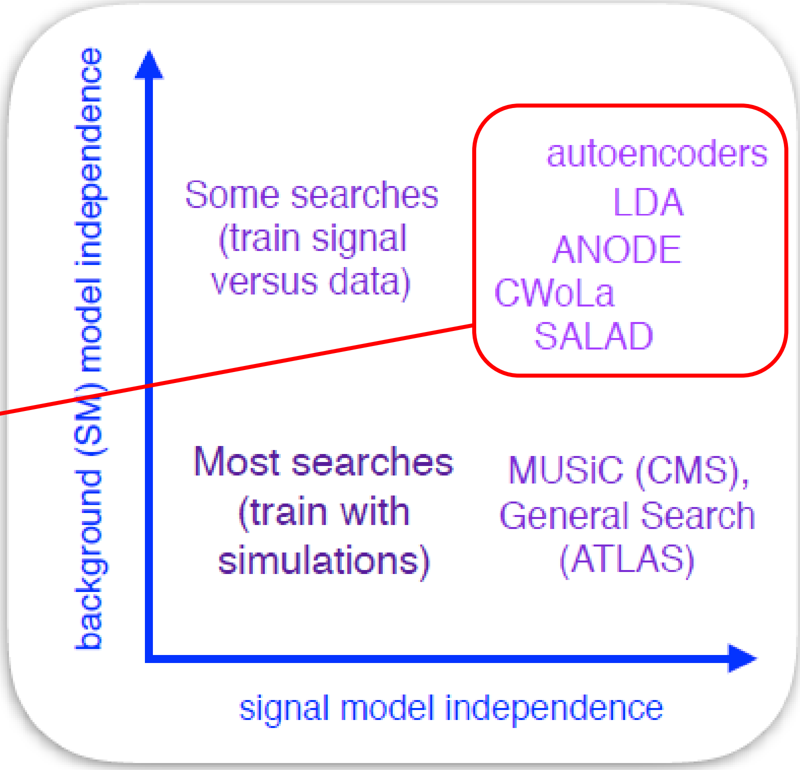

2001.04990

Anomaly detection methods II

- Signal model independent:
  - Simulate background
  - Directly compare data vs simulation over several final states

- Model ”agnostic” or ”independent”:
  - Does not assume a specific signal model
  - Does not rely on simulations
  - Does make assumptions about signal shape or optimal features to search in

Will focus on this type in the rest of the talk

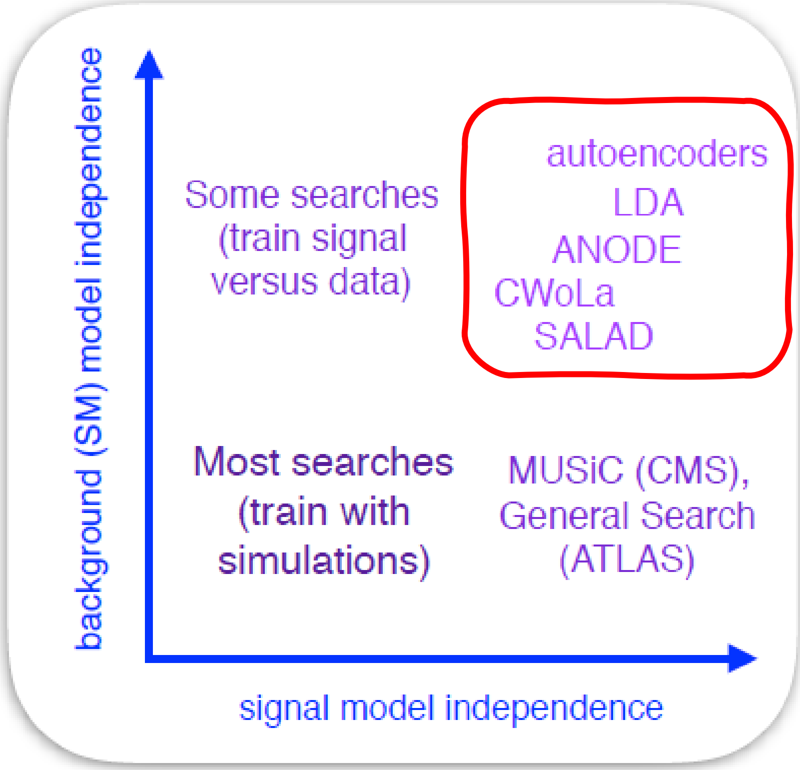

2001.04990

Supervision or not?

- Supervised: labelled data
  - Example: anti-QCD tagging
    J. A. Aguilar-Saavedra, J. H. Collins, and R. K. Mishra, 1709.01087

Supervision or not?

- Supervised: labelled data
  - Example: anti-QCD tagging
    J. A. Aguilar-Saavedra, J. H. Collins, and R. K. Mishra, 1709.01087
- Unsupervised: unlabelled data
  - Example: autoencoders
    D. P. Kingma and M. Welling: Auto-encoding variational bayes; An introduction to variational autoencoders

Supervision or not?

- Supervised: labelled data
  - Example: anti-QCD tagging
    J. A. Aguilar-Saavedra, J. H. Collins, and R. K. Mishra, 1709.01087
- Unsupervised: unlabelled data
  - Example: autoencoders
    D. P. Kingma and M. Welling: Auto-encoding variational bayes; An introduction to variational autoencoders
- Weakly supervised: Imperfect or noisy labels
  - Example: Classification WithOut Labels (CWoLa)
    E. M. Metodiev, B. Nachman, and J. Thaler, Classfication without labels: Learning from mixed samples in high energy physics, 1708.02949
  - Example: Classifying Anomalies THrough Outer Density Estimation (CATHODE)
    A. Hallin, J. Isaacson, G. Kasieczka, C. Krause, B. Nachman, T. Quadfasel, M. Schlaffer, D. Shih, M. Sommerhalder, Classifying Anomalies THrough Outer Density Estimation (CATHODE), 2109.00546

Autoencoders

- Unsupervised
- Lossy encoder
- Reconstruct via decoder
- Measure reconstruction loss

- Idea: the reconstruction loss will be higher for unusual event types

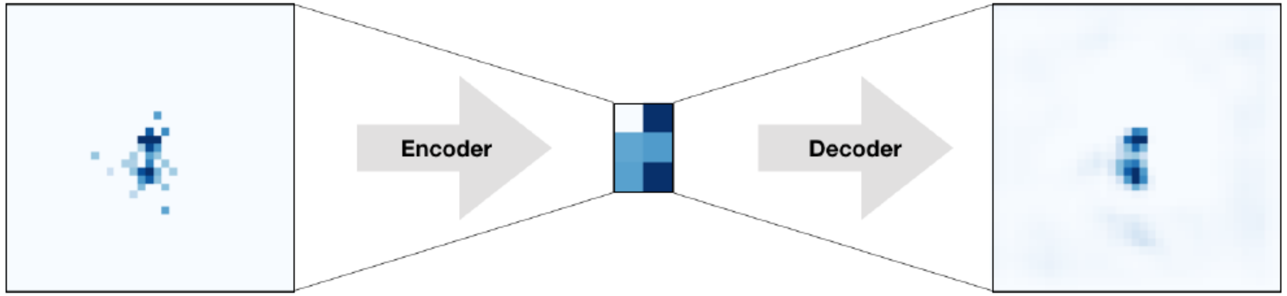

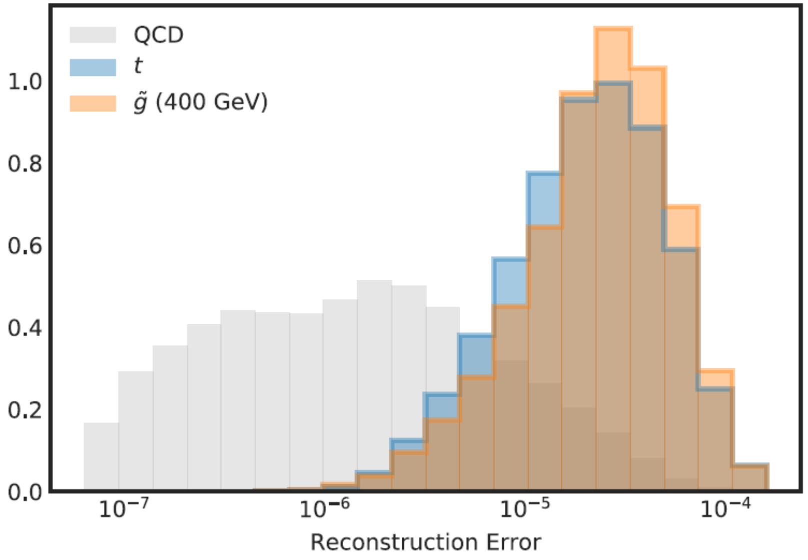

1808.08992

Autoencoders: complexity bias

- What if the signal has a less complex shape than the background?
- Reconstruction loss will not be a good measure

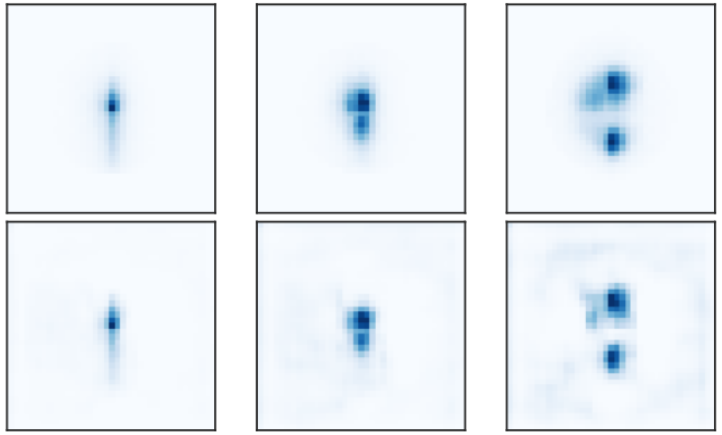

1808.08992

Autoencoders: complexity bias

- What if the signal has a less complex shape than the background?
- Reconstruction loss will not be a good measure

- 2105.05735 and 2206.14225 use *normalized* autoencoders, which helps avoid this issue.

  ”The reconstruction error is re-interpreted as an energy function, and defines a probabilistic model from an autoencoder. During maximum likelihood learning of NAE, outlier reconstruction is naturally suppressed by enforcing the normalization constraint”

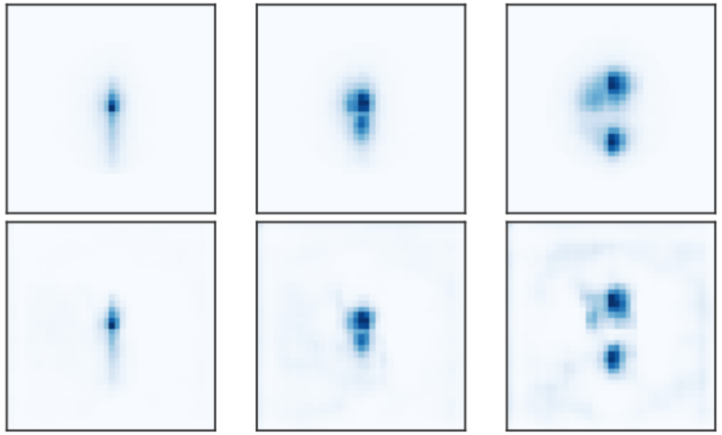

1808.08992

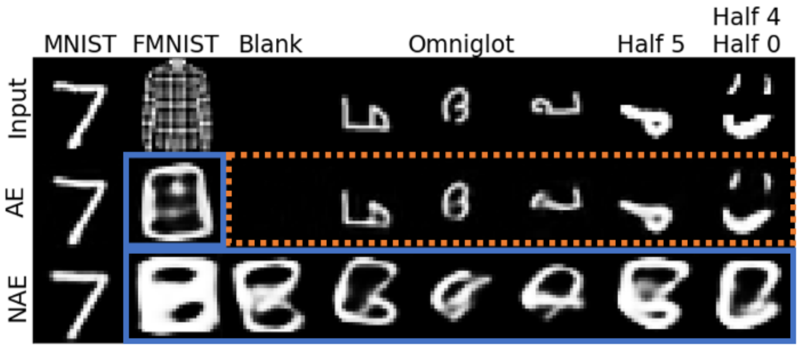

2105.05735

Weak supervision: mixed samples

- We do not know or assume the distribution of the signal
- Sample M_1 and M_2 both contain signal and background, but with different proportions
- Label M_1 and M_2
- Optimality in classifying M_1 vs M_2 → optimality in classifying S vs B
- Let M_1 = data, M_2 = background
- How do we find the background?

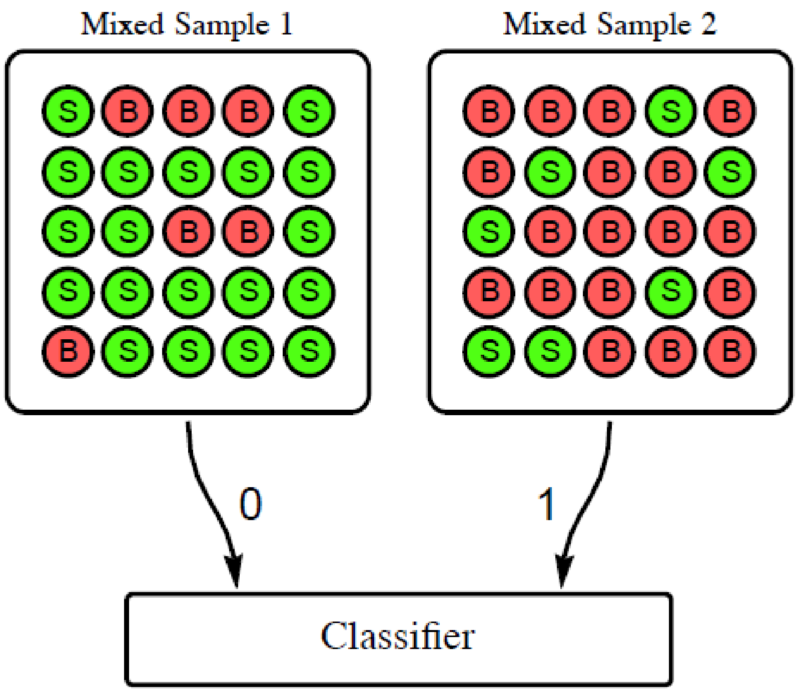

1708.02949

Weak supervision: CWoLa

- Scan over a variable where the signal is assumed to be localized, for example a mass.

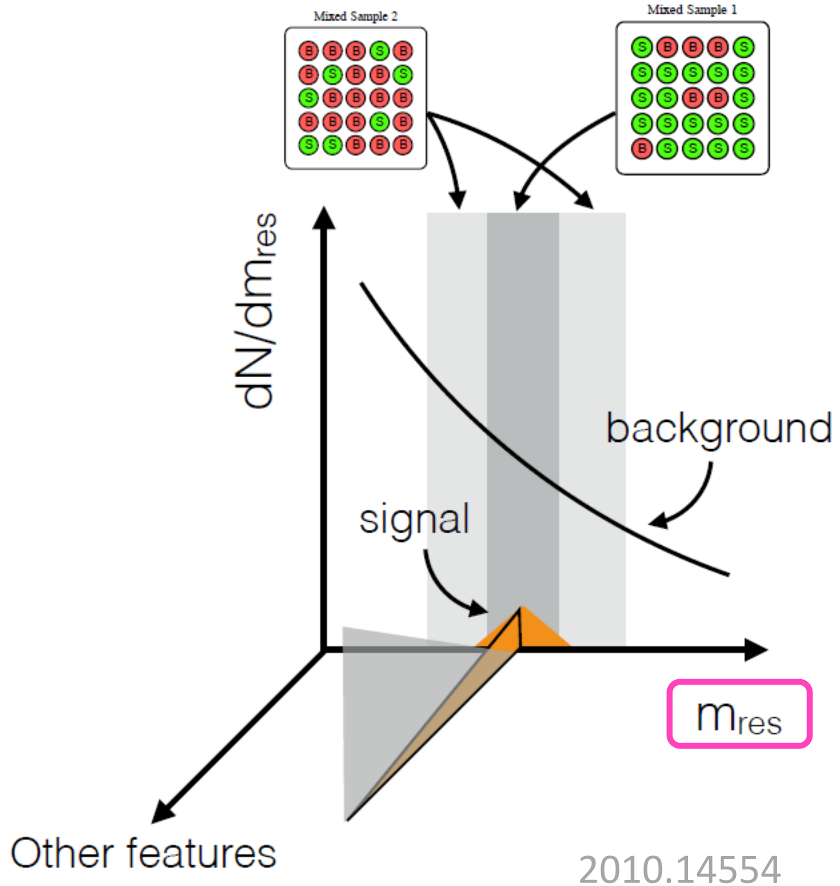

2010.14554

Weak supervision: CWoLa

- Scan over a variable where the signal is assumed to be localized, for example a mass.
- Assume $p_{bg,SR}(m, \mathbf{x}) = p_{data,SB}(m, \mathbf{x})$.

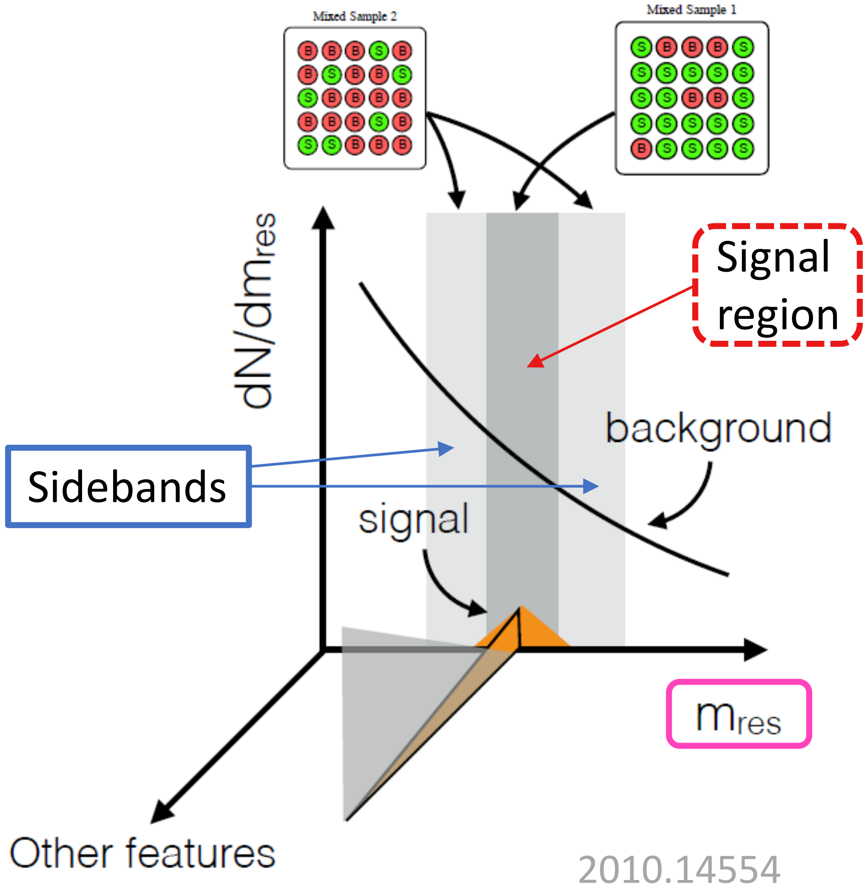

2010.14554

Weak supervision: CWoLa

- Scan over a variable where the signal is assumed to be localized, for example a mass.
- Assume $p_{bg,SR}(m, \mathbf{x}) = p_{data,SB}(m, \mathbf{x})$. Train a classifier *on the other features* to distinguish between events in the signal region and the sidebands.

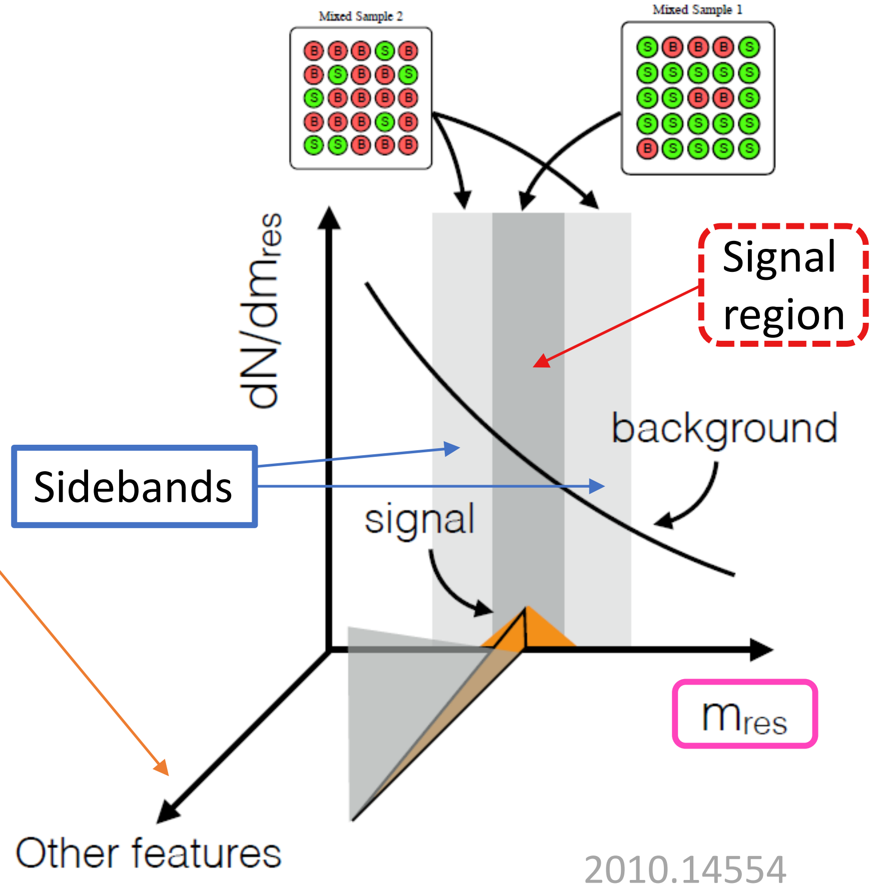

2010.14554

Weak supervision: CWoLa

- Scan over a variable where the signal is assumed to be localized, for example a mass.
- Assume $p_{bg,SR}(m,\mathbf{x}) = p_{data,SB}(m,\mathbf{x})$. Train a classifier *on the other features* to distinguish between events in the signal region and the sidebands.
- If the training features are correlated with the search feature, the approach collapses.

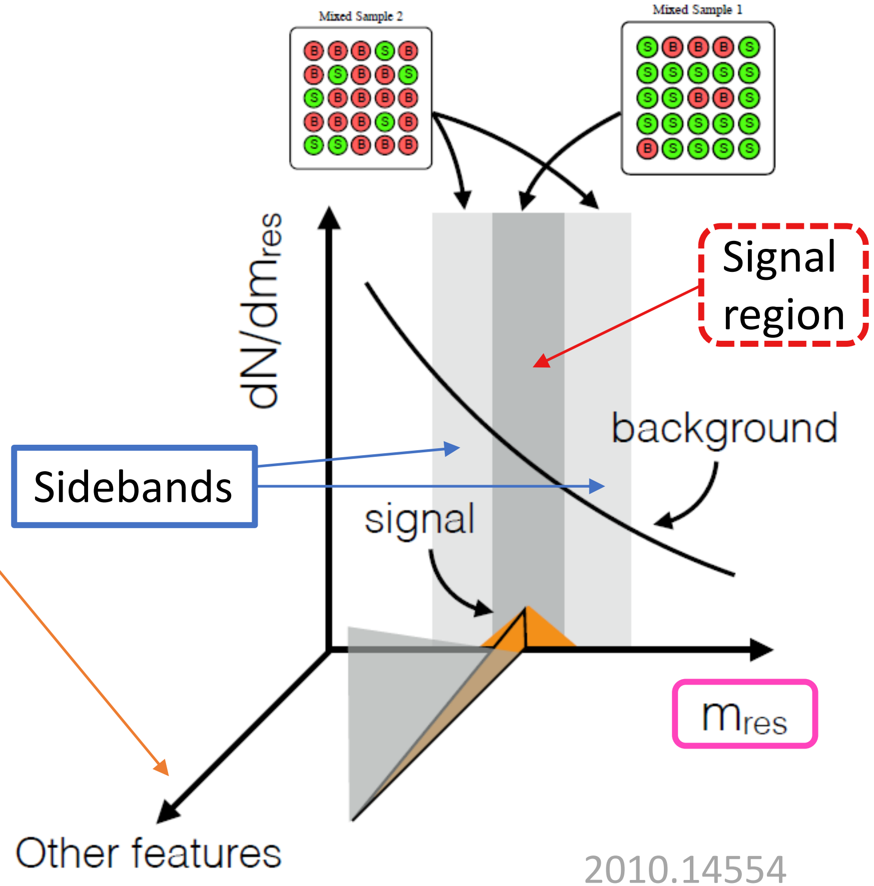

2010.14554

Weak supervision: CATHODE

- Define wide sidebands in the search variable and train a density estimator on auxiliary variables in this region.

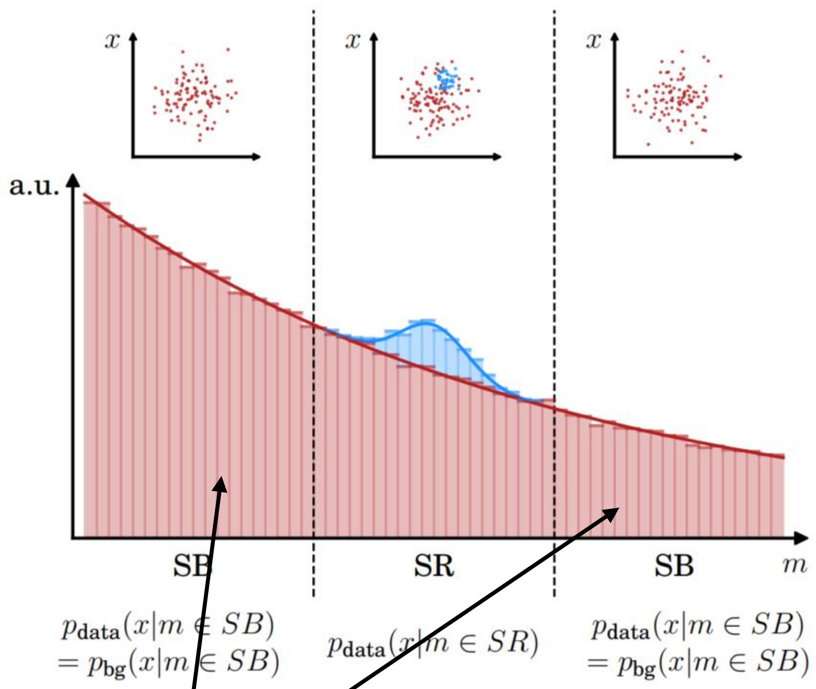

2109.00546

Train density estimator here

Weak supervision: CATHODE

- Define wide sidebands in the search variable and train a density estimator on auxiliary variables in this region.
- The density estimator in CATHODE is a normalizing flow. Normalizing flows learn an invertible map f between the feature space with distribution $p_X(\mathbf{x})$ and a latent space with distribution $p_Z(\mathbf{z})$, such that $\mathbf{z} = f(\mathbf{x})$.

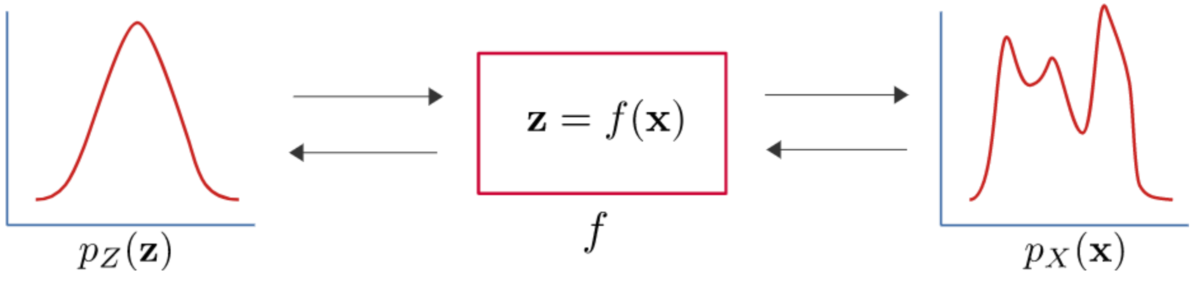

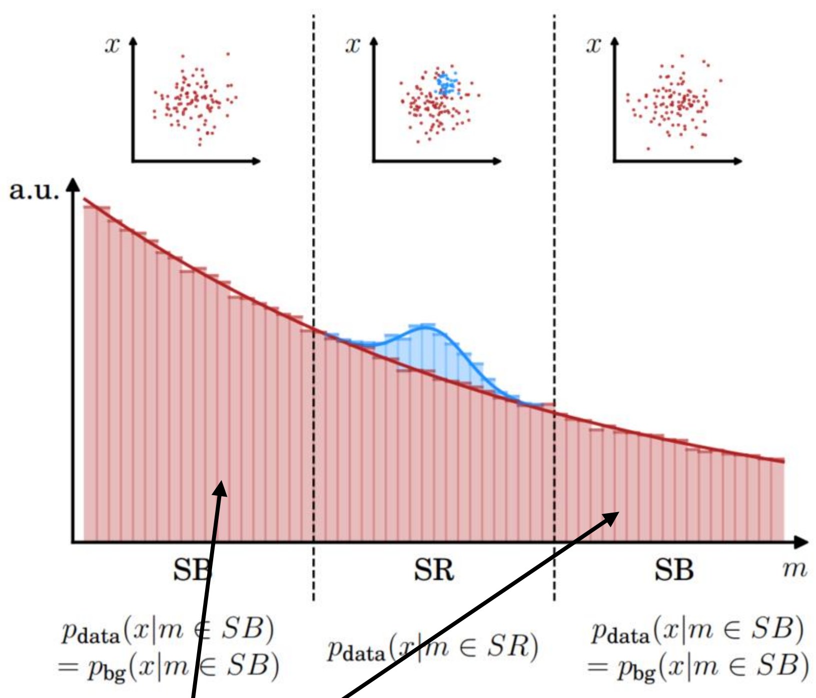

2109.00546

Train density estimator here

Weak supervision: CATHODE

- Define wide sidebands in the search variable and train a density estimator on auxiliary variables in this region.
- Use the learned probability distribution to sample ”background” in the SR.

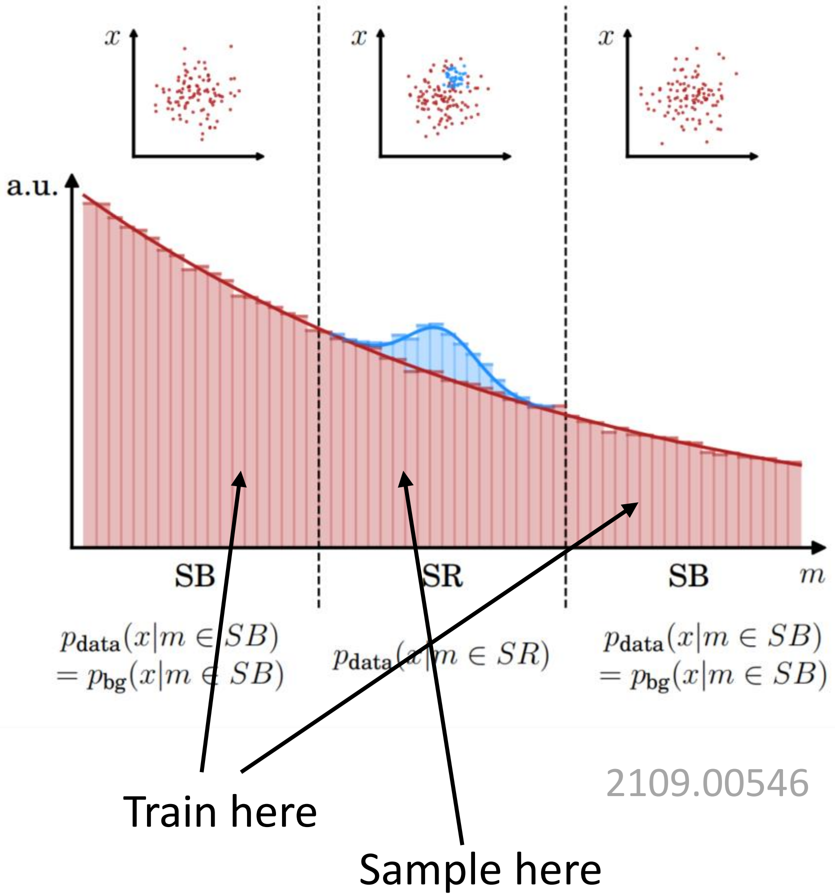

2109.00546

Weak supervision: CATHODE

- Define wide sidebands in the search variable and train a density estimator on auxiliary variables in this region.
- Use the learned probability distribution to sample ”background” in the SR.
- Train a classifier to distinguish between samples (”background”) and data.

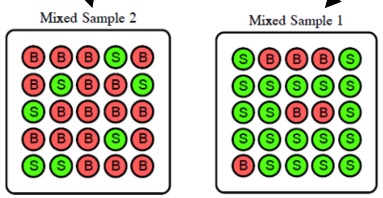

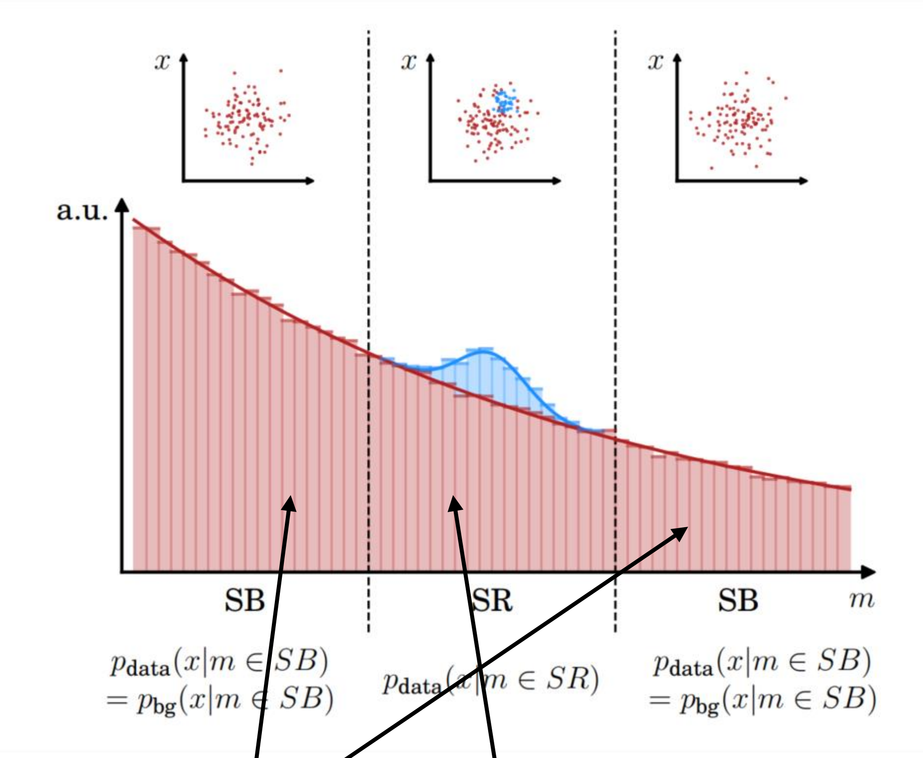

2109.00546

Train here

Sample here

Weak supervision: pitfalls

- Noisy features degrade performance
  - Need to be careful with feature selection
  - Hence not completely ”model independent”

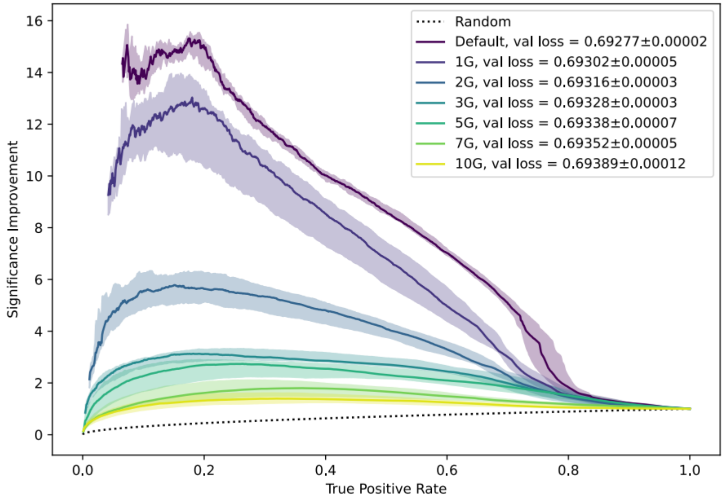

Weak supervision: pitfalls

- Noisy features degrade performance
  - Need to be careful with feature selection
  - Hence not completely "model independent"

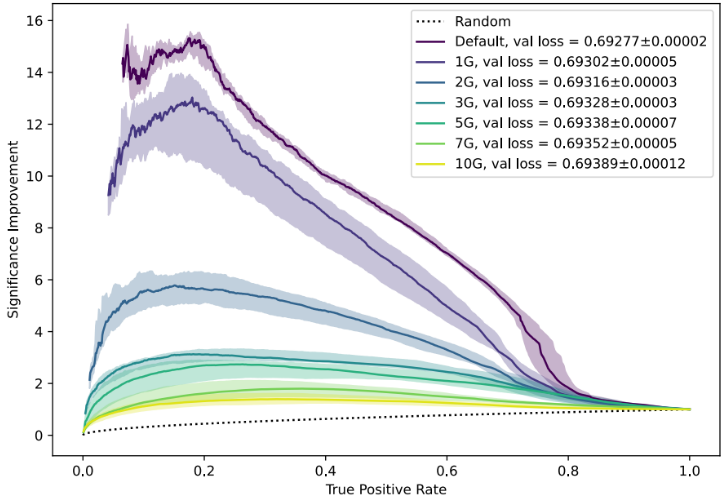

- Possible background sculpting
  - If there are strongly correlated features, this can happen after a cut on the anomaly score
  - CATHODE's performance restored when performing the classification in the latent space of the normalizing flow (LaCATHODE)

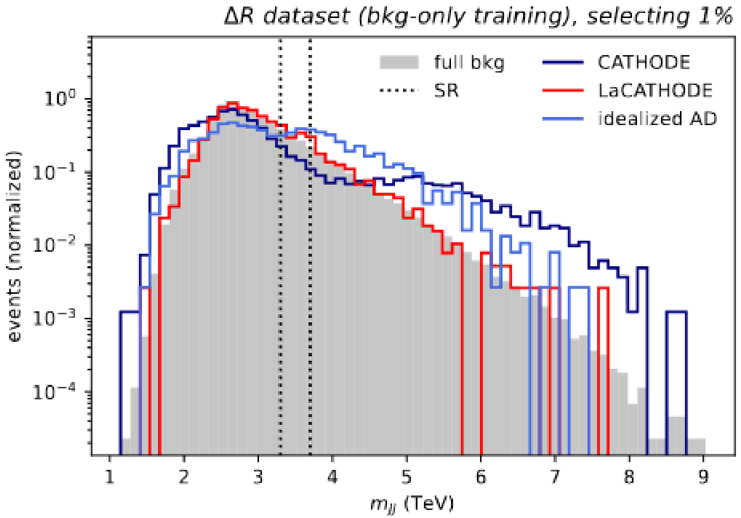

2210.14924

Example results from analyses

- ATLAS Dijet resonance search with weak supervision (CWoLa)
- 2D feature space (masses of the two jets): CWoLa correctly identifies the signal

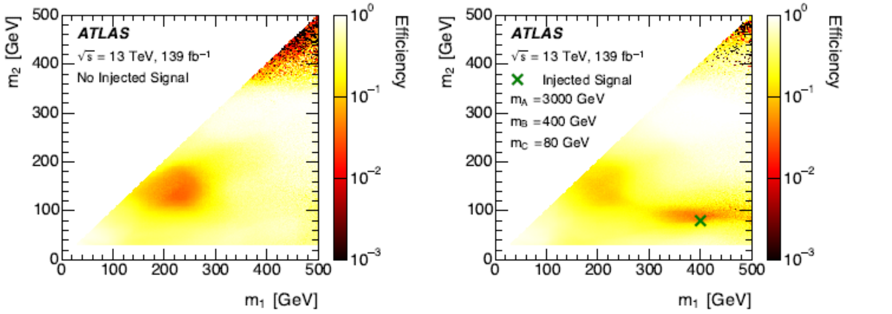

2005.02983

Example results from analyses

- ATLAS Dijet resonance search with weak supervision (CWoLa)
- 2D feature space (masses of the two jets): CWoLa correctly identifies the signal
- *"For certain masses, these limits are up to 10 times more sensitive than those obtained by the inclusive dijet search"*

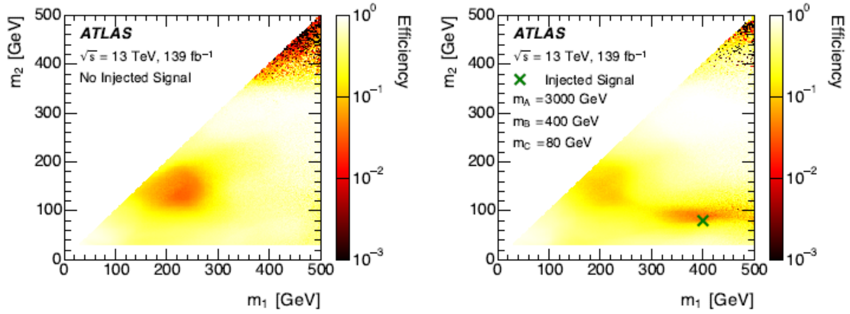

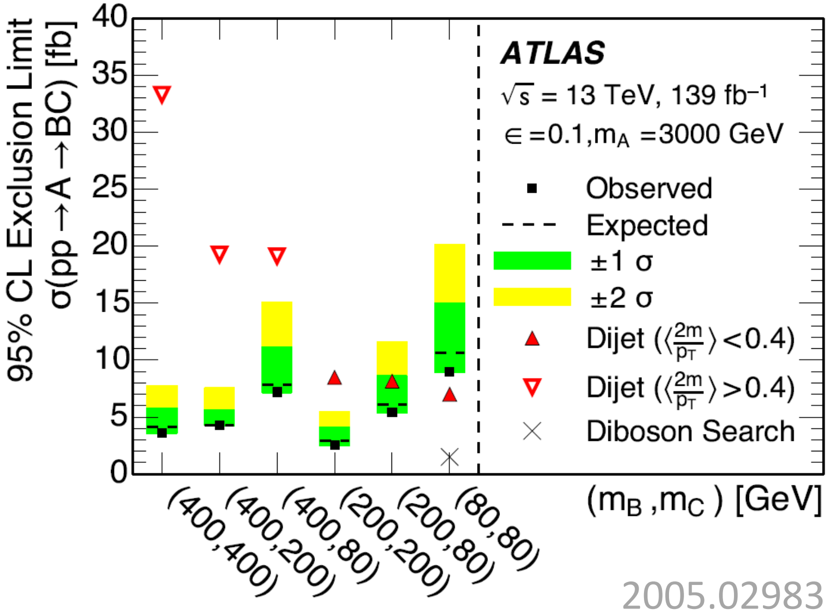

2005.02983

Example results from analyses

- CATHODE: particle physics analysis is currently being performed... stay tuned!

Example results from analyses

- CATHODE: particle physics analysis is currently being performed... stay tuned!
- We are also using it in space!
  - Stellar streams are streams of stars that the Milky Way has pulled from other stellar bodies
  - All stars in a stream travel in the same direction (background stars do not)
  - They all have the same origin → similar color and magnitude
  - They are anomalies!
  - Currently using CATHODE to find new streams.

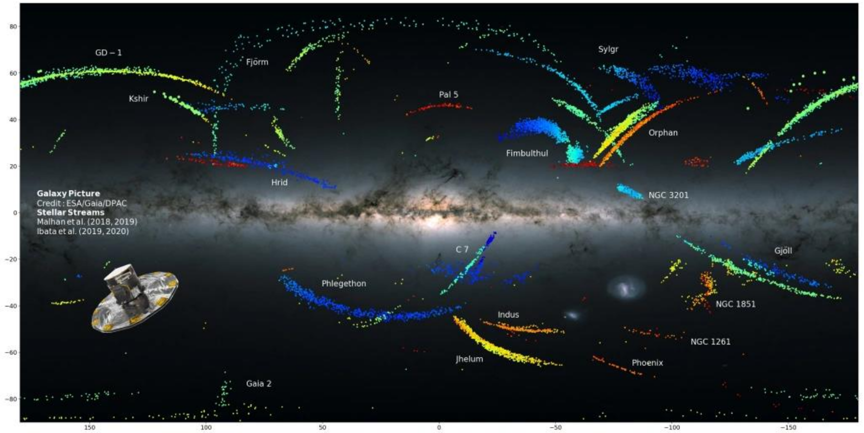

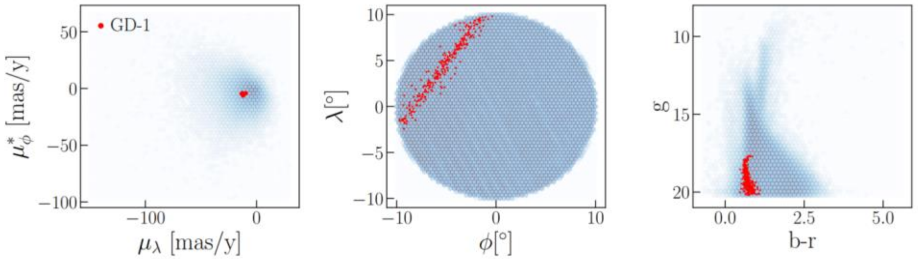

What about the trigger level?

- Active research field
- Challenge: needs to be fast

Summary and outlook

- There are many different types of anomalies
- Group anomalies are studied in the context of searches for new physics
- Anomaly detection can be more or less model dependent
- ML methods for Anomaly detection vary in the level of supervision
- Example of unsupervised approaches: autoencoders
- Examples of weak supervision: CWoLa, CATHODE
- Different methods are currently being implemented in real data analyses, results to come