

Searching for new physics with VAEs

Barry M. Dillon

November 16, 2020

Institute for Theoretical Physics
University of Heidelberg

UZH ML Workshop

in collab. with B. Bortolato, A. Smolkovic, J. F. Kamenik

UNIVERSITÄT
HEIDELBERG
Zukunft. Seit 1386.

Outline

1. Anomalies at the Large Hadron Collider
2. Finding anomalies with VAEs
3. Results on the LHC Olympics
4. Concluding remarks

1. Anomalies at the Large Hadron Collider

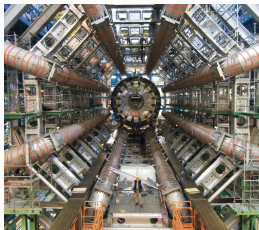
2. Finding anomalies with VAEs

3. Results on the LHC Olympics

4. Concluding remarks

The Large Hadron Collider

- Highest energy particle collider in the world
- Constructed by CERN between 1998-2008
- Collides protons at energies of 13-14 TeV producing new particles that are measured in the detectors



- The LHC will run for many more years yet
⇒ much more data to come

New physics searches

- Higgs boson discovery in 2012
- ... but no new physics discovery since
(composite Higgs, supersymmetry, ...)

New physics searches

- Higgs boson discovery in 2012
- ... but no new physics discovery since
(composite Higgs, supersymmetry, ...)

So, why? ...

- there's just no new physics there
- the statistics are too small
- or, we haven't performed the right search.

New physics searches

- Higgs boson discovery in 2012
- ... but no new physics discovery since (composite Higgs, supersymmetry, ...)

So, why? ...

- there's just no new physics there
- the statistics are too small
- or, we haven't performed the right search.

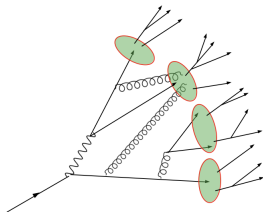
Can machine learning help?

- (Variational) AutoEncoders (← what I will get to soon)
- Classification Without Labels (CWoLa) [Collins, Howe, Nachman (2019)]
- Density estimation [Nachman, Shih (2020)]
- Latent Dirichlet Allocation [BMD, Faroughy, Kamenik (2019)]
- ...

Could new physics be hidden in jet substructure?

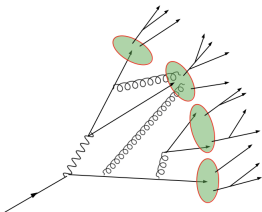
Could new physics be hidden in jet substructure?

Boosted jets

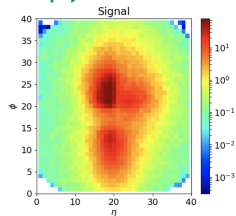
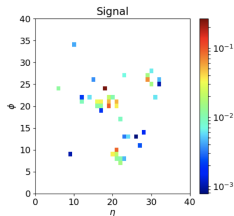


Could new physics be hidden in jet substructure?

Boosted jets

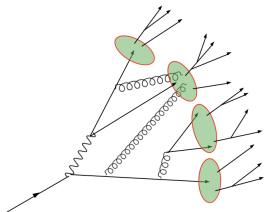


Boosted top jets

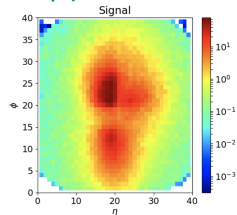
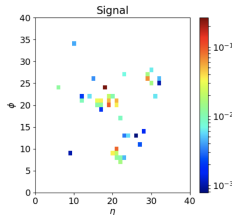


Could new physics be hidden in jet substructure?

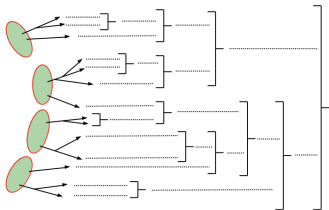
Boosted jets



Boosted top jets

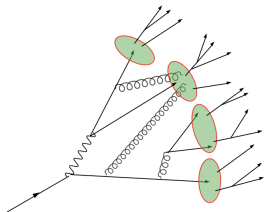


Sequential clustering histories

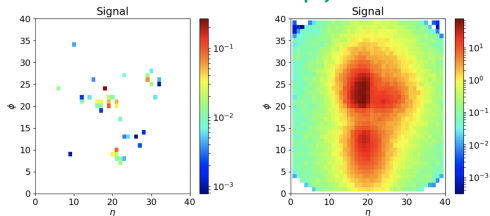


Could new physics be hidden in jet substructure?

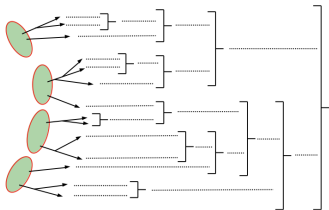
Boosted jets



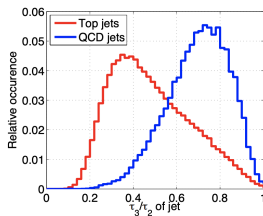
Boosted top jets



Sequential clustering histories



Global jet observables



Could new physics be hidden in jet substructure?

Proposed scenario:

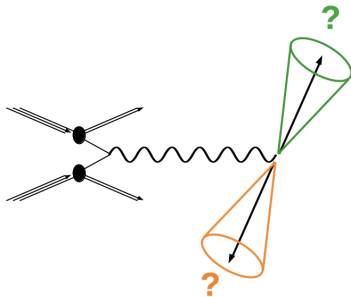
- Unknown heavy resonance 'A' decaying to di-jets

$$pp \rightarrow A \rightarrow B C$$

- 'B' and/or 'C' are also unknown

Goal:

- Unsupervised classification of these events from **background events**



Could new physics be hidden in jet substructure?

Proposed scenario:

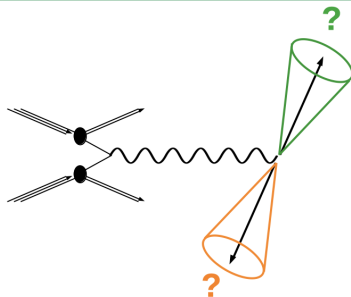
- Unknown heavy resonance 'A' decaying to di-jets

$$pp \rightarrow A \rightarrow B C$$

- 'B' and/or 'C' are also unknown

Goal:

- Unsupervised classification of these events from **background events**



Observables:

- High-level, e.g. jet masses, N-subjettiness, ...
 - ... but we don't know what to look for
- ⇒ Use neural networks to find anomalies in the data.

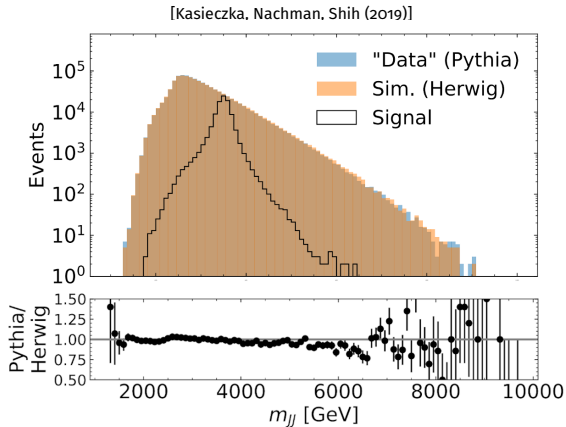
1. Anomalies at the Large Hadron Collider

2. Finding anomalies with VAEs

3. Results on the LHC Olympics

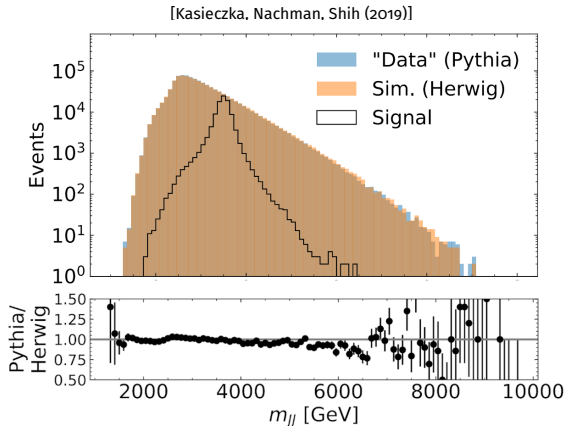
4. Concluding remarks

Bump hunting



- We expect anomalies to be localised in **invariant mass**
- We need to look for **bumps** in the invariant mass spectrum

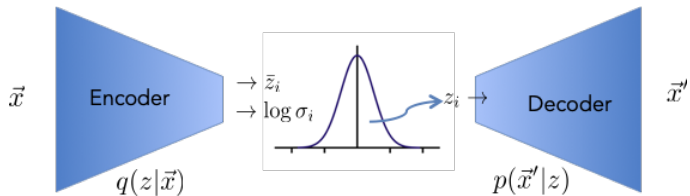
Bump hunting



- We expect anomalies to be localised in **invariant mass**
- We need to look for **bumps** in the invariant mass spectrum
- ... **We need a classifier to improve signal-to-background ratio**
- The classifier needs to be invariant-mass independent!

Variational Autoencoders

Data for each event embedded in a vector \vec{x}_i

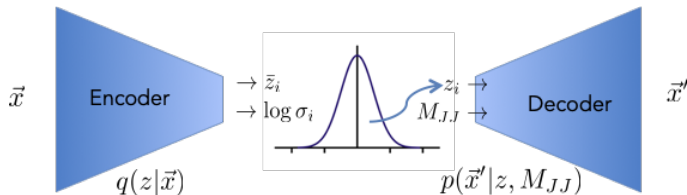


Two anomaly measures:

1. Distance from centre in latent space
2. Reconstruction error

Variational Autoencoders invariant mass latent dimension

Data for each event embedded in a vector \vec{x}_i



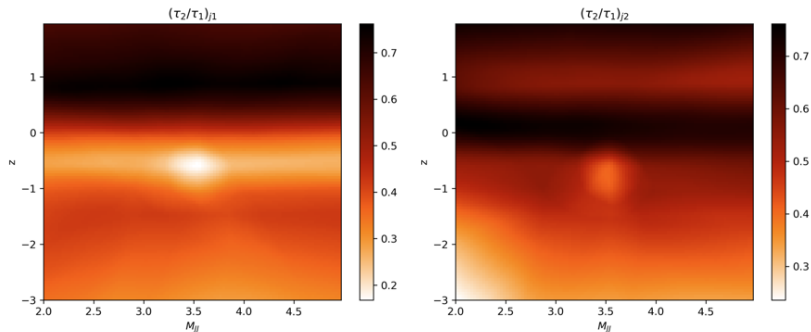
Two anomaly measures:

1. Distance from centre in latent space
2. Reconstruction error

The network can learn locally in the invariant mass.

An invariant mass latent dimension

Decoder output:



Through the latent space, we can explore what the networks learn about invariant mass

1. Anomalies at the Large Hadron Collider

2. Finding anomalies with VAEs

3. Results on the LHC Olympics

4. Concluding remarks

The LHC Olympics [Kasieczka, Nachman, Shih (2019)]

What is it?

A LHC anomaly challenge using di-jets

The LHC Olympics [Kasieczka, Nachman, Shih (2019)]

What is it?

A LHC anomaly challenge using di-jets

Test data

- 1M background events and 1k signal events
- Use to design and optimise techniques

The LHC Olympics [Kasieczka, Nachman, Shih (2019)]

What is it?

A LHC anomaly challenge using di-jets

Test data

- 1M background events and 1k signal events
- Use to design and optimise techniques

Blackboxes

- 1M unlabelled events possibly containing anomalies

Goals:

- Identify anomalies
- Characterise them as well as possible
i.e. how heavy are the jets, describe the physical processes, ...

Results: test data

Neural network architecture and optimisation:

Results: test data

Neural network architecture and optimisation:

- 1 latent dimension + an invariant mass latent dimension

Results: test data

Neural network architecture and optimisation:

- 1 latent dimension + an invariant mass latent dimension
- 2 layers of 100 nodes in encoder and decoder

Results: test data

Neural network architecture and optimisation:

- 1 latent dimension + an invariant mass latent dimension
- 2 layers of 100 nodes in encoder and decoder
- SeLu activations and MSE loss

Neural network architecture and optimisation:

- 1 latent dimension + an invariant mass latent dimension
- 2 layers of 100 nodes in encoder and decoder
- SeLu activations and MSE loss
- Adadelata optimiser

Results: test data

Neural network architecture and optimisation:

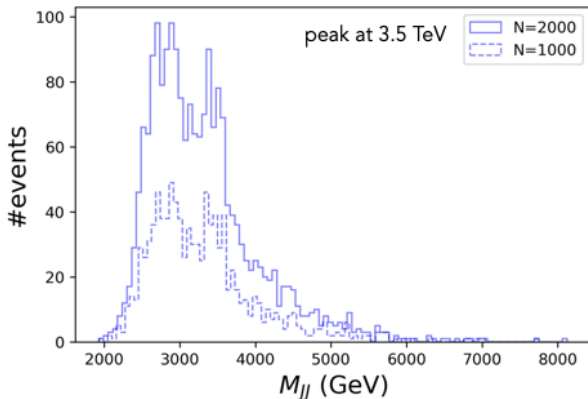
- 1 latent dimension + an invariant mass latent dimension
- 2 layers of 100 nodes in encoder and decoder
- SeLu activations and MSE loss
- Adadelata optimiser
- Train for 100 epochs, batch size of 10k events

Neural network architecture and optimisation:

- 1 latent dimension + an invariant mass latent dimension
- 2 layers of 100 nodes in encoder and decoder
- SeLu activations and MSE loss
- Adadelta optimiser
- Train for 100 epochs, batch size of 10k events
- Best performance correlated with minima in the mean log-variance in latent space

Results: test data

Now perform cuts, and study the invariant mass spectrum:

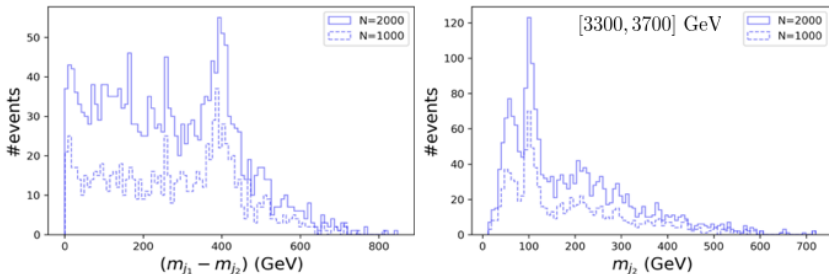


The VAE successfully finds the anomalous events localised at 3.5 TeV

Results: test data

Can we see characteristics of these anomalous events?

Look at features of anomalous events in the signal region:

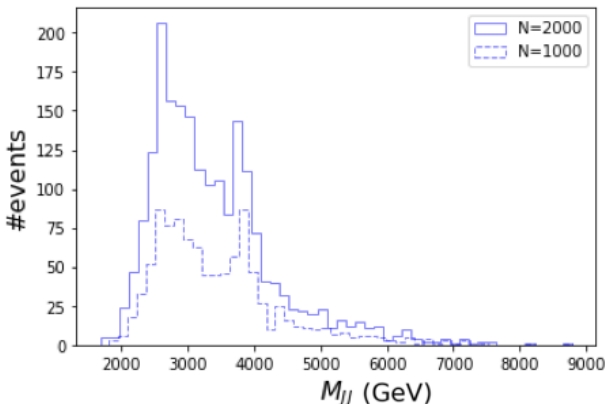


Anomalous events with jet masses of 500 GeV and 100 GeV

We can also do the same with other features.

Results: blackbox data

Here we don't have truth labels, so we run the same algorithm and study the invariant mass spectrum:

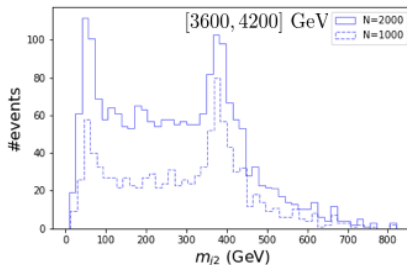
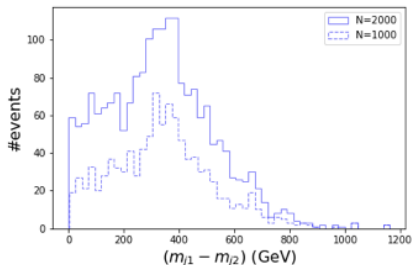


The VAE finds a **bump** localised at 3.8 TeV

Results: blackbox data

Is this bump actually a signal?

Let's look in the signal region and see if there are any features:



Anomalous events with jet masses of ~ 740 GeV and 370 GeV

We can also do the same with other features.

This agrees with the results from the challenge!

-
1. Anomalies at the Large Hadron Collider
 2. Finding anomalies with VAEs
 3. Results on the LHC Olympics
 4. Concluding remarks

Concluding remarks

1. Simple VAE with a 1D latent space was enough to find the anomalous events

Concluding remarks

1. Simple VAE with a 1D latent space was enough to find the anomalous events
2. Why?

Probably because of the high-level inputs we used
... we know *approximately* where to look

Concluding remarks

1. Simple VAE with a 1D latent space was enough to find the anomalous events
2. Why?
Probably because of the high-level inputs we used
... we know *approximately* where to look
3. What next?

Concluding remarks

1. Simple VAE with a 1D latent space was enough to find the anomalous events
2. Why?
Probably because of the high-level inputs we used
... we know *approximately* where to look
3. What next?
 - Understand the optimisation procedure
... we're not sure why this works

Concluding remarks

1. Simple VAE with a 1D latent space was enough to find the anomalous events
2. Why?
Probably because of the high-level inputs we used
... we know *approximately* where to look
3. What next?
 - Understand the optimisation procedure
... we're not sure why this works
 - Apply this technique to other types of signals,
perhaps with a richer substructure

Concluding remarks

1. Simple VAE with a 1D latent space was enough to find the anomalous events
2. Why?
Probably because of the high-level inputs we used
... we know *approximately* where to look
3. What next?
 - Understand the optimisation procedure
... we're not sure why this works
 - Apply this technique to other types of signals,
perhaps with a richer substructure
 - Explore the idea of embedding observables in the latent space

Concluding remarks

1. Simple VAE with a 1D latent space was enough to find the anomalous events
2. Why?
Probably because of the high-level inputs we used
... we know *approximately* where to look
3. What next?
 - Understand the optimisation procedure
... we're not sure why this works
 - Apply this technique to other types of signals,
perhaps with a richer substructure
 - Explore the idea of embedding observables in the latent space

Thanks for your time!