

The LHC Olympics 2020: A Community Challenge for Anomaly Detection in High Energy Physics

Benjamin Nachman

Lawrence Berkeley National Laboratory

bpnachman.com

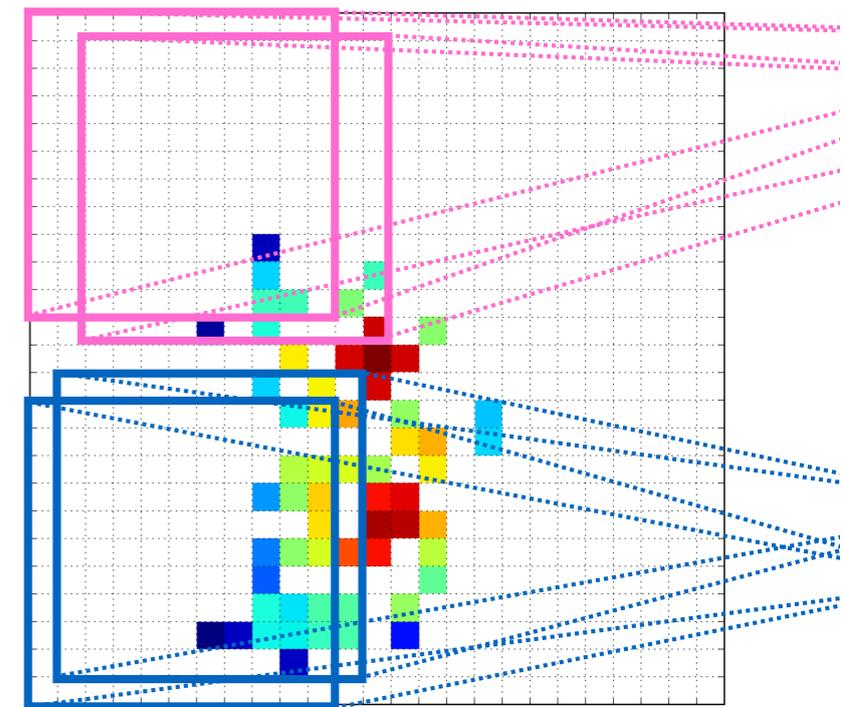
bpnachman@lbl.gov



@bpnachman



bnachman



LHC
Reinterpretation
Workshop
Feb. 18, 2021

The LHC Olympics 2020

A Community Challenge for Anomaly
Detection in High Energy Physics



Gregor Kasieczka (ed),¹ Benjamin Nachman (ed),^{2,3} David Shih (ed),⁴ Oz Amram,⁵ Anders Andreassen,⁶ Kees Benkendorfer,^{2,7} Blaz Bortolato,⁸ Gustaaf Brooijmans,⁹ Florencia Canelli,¹⁰ Jack H. Collins,¹¹ Biwei Dai,¹² Felipe F. De Freitas,¹³ Barry M. Dillon,^{8,14} Ioan-Mihail Dinu,⁵ Zhongtian Dong,¹⁵ Julien Donini,¹⁶ Javier Duarte,¹⁷ D. A. Faroughy,¹⁰ Julia Gonski,⁹ Philip Harris,¹⁸ Alan Kahn,⁹ Jernej F. Kamenik,^{8,19} Charanjit K. Khosa,^{20,30} Patrick Komiske,²¹ Luc Le Pottier,^{2,22} Pablo Martín-Ramiro,^{2,23} Andrej Matevc,^{8,19} Eric Metodiev,²¹ Vinicius Mikuni,¹⁰ Inês Ochoa,²⁴ Sang Eon Park,¹⁸ Maurizio Pierini,²⁵ Dylan Rankin,¹⁸ Veronica Sanz,^{20,26} Nilai Sarda,²⁷ Uroš Seljak,^{2,3,12} Aleks Smolkovic,⁸ George Stein,^{2,12} Cristina Mantilla Suarez,⁵ Manuel Szewc,²⁸ Jesse Thaler,²¹ Steven Tsan,¹⁷ Silviu-Marian Udrescu,¹⁸ Louis Vaslin,¹⁶ Jean-Roch Vlimant,²⁹ Daniel Williams,⁹ Mikael Yunus¹⁸

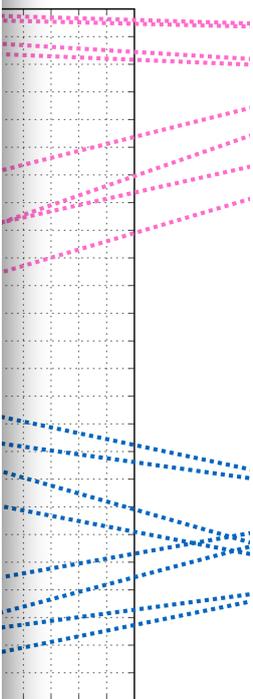
Lawre

bpnac

bpnachr

BE

BE



tion

21

(Brief) Motivation



Theoretical and **experimental** questions motivate a deep exploration **of the fundamental structure of nature**

Dark matter

Hierarchy problem

Strong CP

Flavor puzzles

Baryogenesis

Dark energy

We have performed thousands of hypothesis tests & have no significant evidence for physics beyond the Standard Model

Three possibilities



(Brief) Motivation



Theoretical and **experimental** questions motivate a deep exploration **of the fundamental structure of nature**

Dark matter

Hierarchy problem

Strong CP

Flavor puzzles

Baryogenesis

Dark energy

We have performed thousands of hypothesis tests & have no significant evidence for physics beyond the Standard Model

Three possibilities

(1) There is nothing new at LHC energies

(Brief) Motivation

5

Theoretical and **experimental** questions motivate a deep exploration **of the fundamental structure of nature**

Dark matter

Hierarchy problem

Strong CP

Flavor puzzles

Baryogenesis

Dark energy

We have performed thousands of hypothesis tests & have no significant evidence for physics beyond the Standard Model

Three possibilities

(1) There is nothing new at LHC energies

(2) Patience! (new physics is rare)

(Brief) Motivation



Theoretical and **experimental** questions motivate a deep exploration **of the fundamental structure of nature**

Dark matter

Hierarchy problem

Strong CP

Flavor puzzles

Baryogenesis

Dark energy

We have performed thousands of hypothesis tests & have no significant evidence for physics beyond the Standard Model

Three possibilities

(1) There is nothing new at LHC energies

(2) Patience! (new physics is rare)

(3) We are not looking in the right place

(Brief) Motivation



Theoretical and **experimental** questions motivate a deep exploration **of the fundamental structure of nature**

Dark matter

Hierarchy problem

Strong CP

Flavor puzzles

Baryogenesis

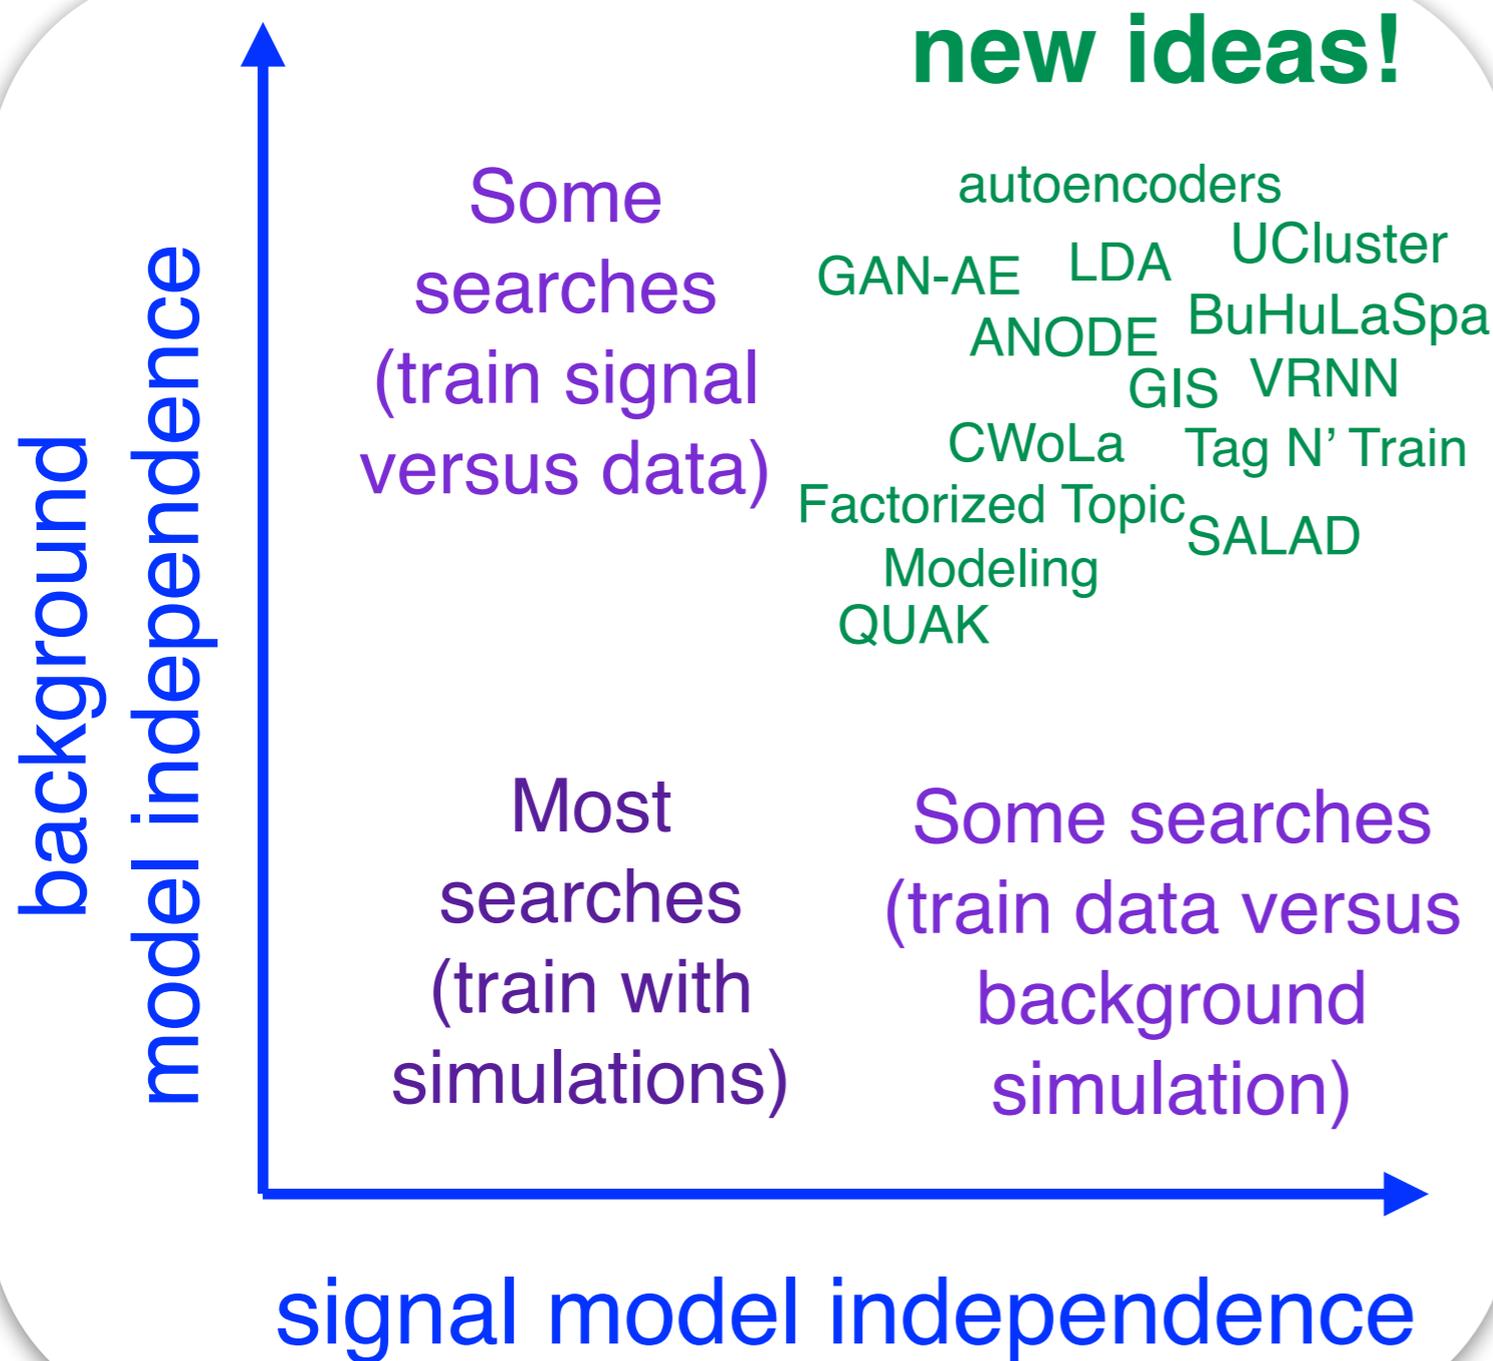
Dark energy

We have performed thousands of hypothesis tests & have no significant evidence for physics beyond the Standard Model

Three possibilities

This is what motivated this challenge!

(3) We are not looking in the right place



There are many new ideas that make use of modern machine learning

The goal is to learn **directly from data**, injecting as little bias as possible

*N.B. this is just for signal sensitivity - there is **also model dependence** for determining the background*

The Challenge



We provide a list of particle for each event (700 particles with the 3-vector of each particle)

1 dataset for R&D with labeled signal and background

3 black boxes with unlabeled data

The particle-level + detector-level simulation for background in the black boxes was modified for each dataset (think Pythia/Herwig, etc.)



Actually, all of the parameters are now public on Zenodo

The Challenge

10

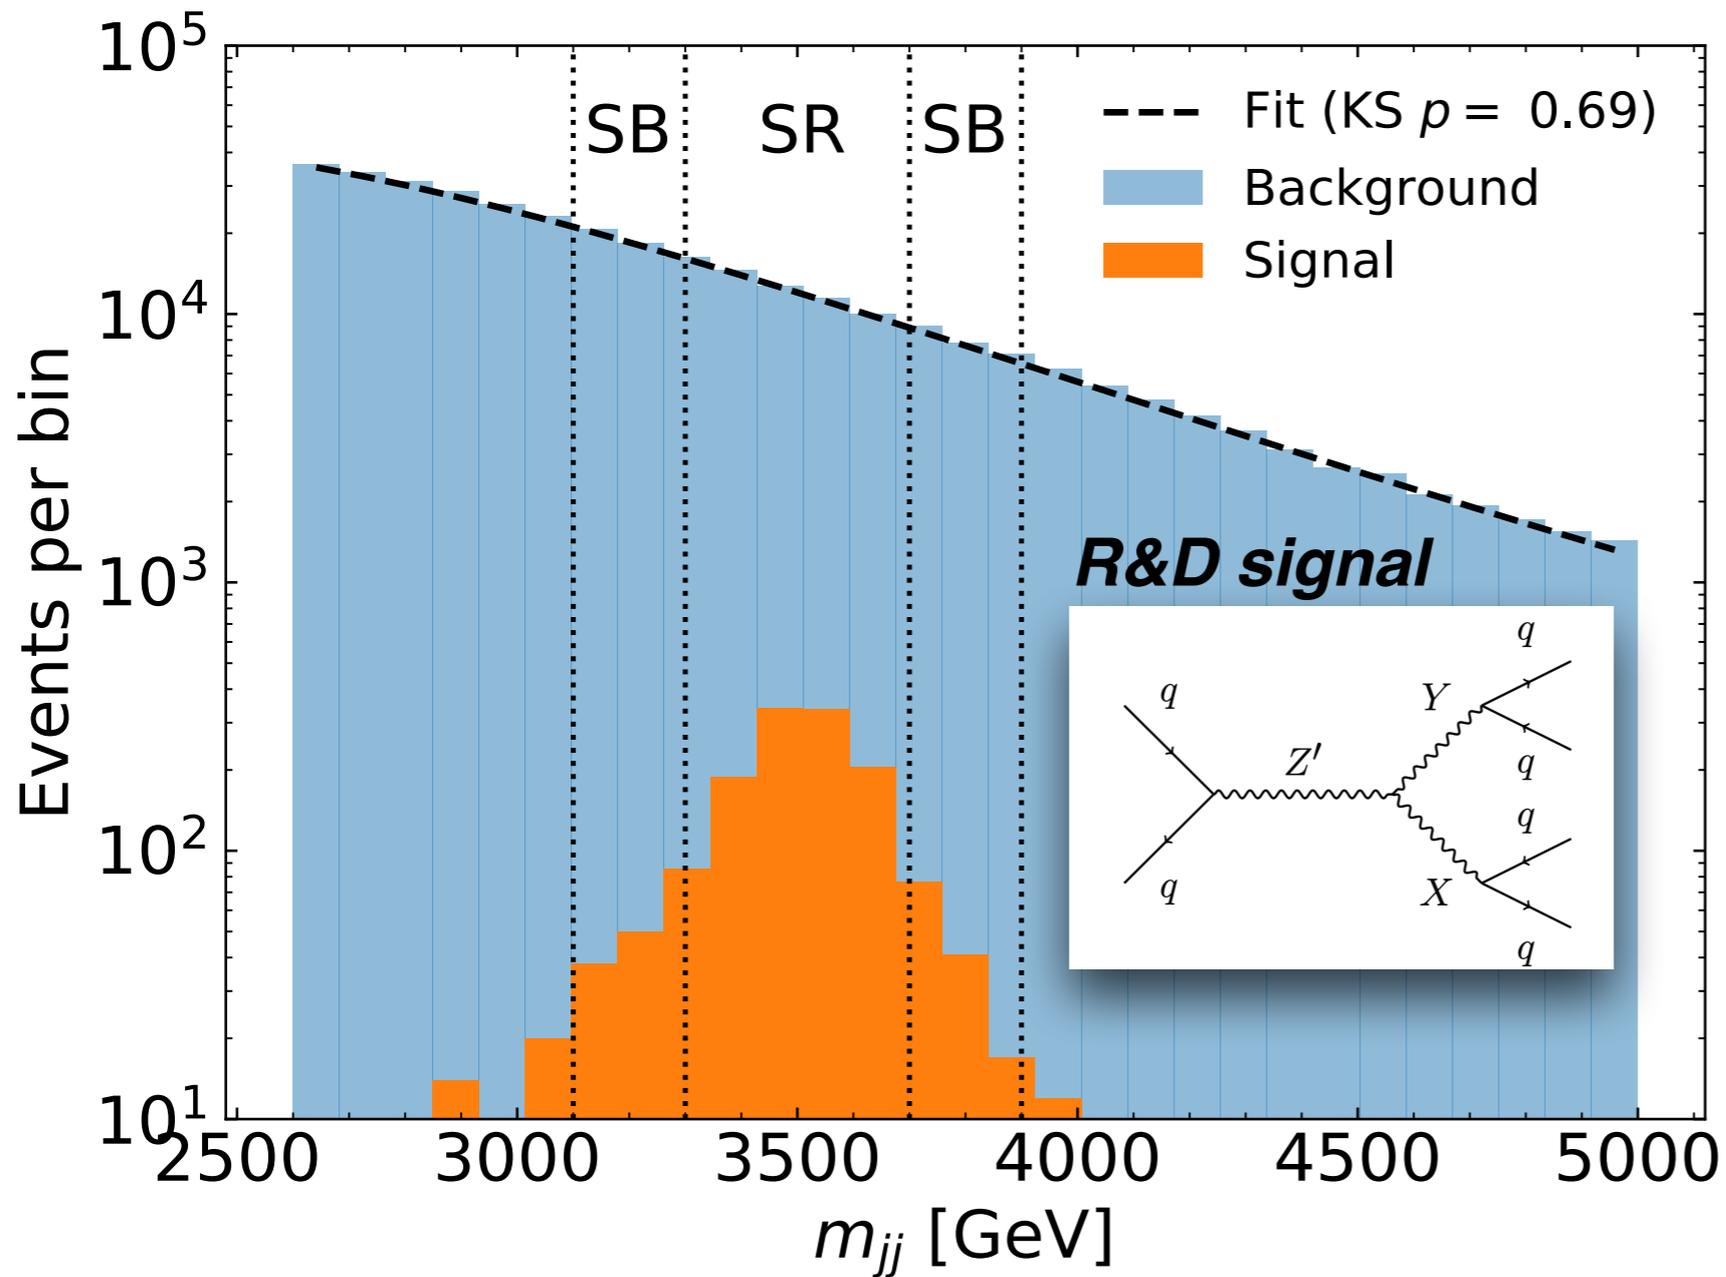
Section	Short Name	Method Type	Results Type
3.1	VRNN	Unsupervised	(i) (BB2,3) and (ii) (BB1)
3.2	ANODE	Unsupervised	(iii)
3.3	BuHuLaSpa	Unsupervised	(i) (BB2,3) and (ii) (BB1)
3.4	GAN-AE	Unsupervised	(i) (BB2-3) and (ii) (BB1)
3.5	GIS	Unsupervised	(i) (BB1)
3.6	LDA	Unsupervised	(i) (BB1-3)
3.7	PGA	Unsupervised	(ii) (BB1-2)
3.8	Reg. Likelihoods	Unsupervised	(iii)
3.9	UCluster	Unsupervised	(i) (BB2-3)
4.1	CWoLa	Weakly Supervised	(ii) (BB1-2)
4.2	CWoLa AE Compare	Weakly/Unsupervised	(iii)
4.3	Tag N' Train	Weakly Supervised	(i) (BB1-3)
4.4	SALAD	Weakly Supervised	(iii)
4.5	SA-CWoLa	Weakly Supervised	(iii)
5.1	Deep Ensemble	Semisupervised	(i) (BB1)
5.2	Factorized Topics	Semisupervised	(iii)
5.3	QUAK	Semisupervised	(i) (BB2,3) and (ii) (BB1)
5.4	LSTM	Semisupervised	(i) (BB1-3)

BB = black box; (i) = blinded, (ii) = unblinded

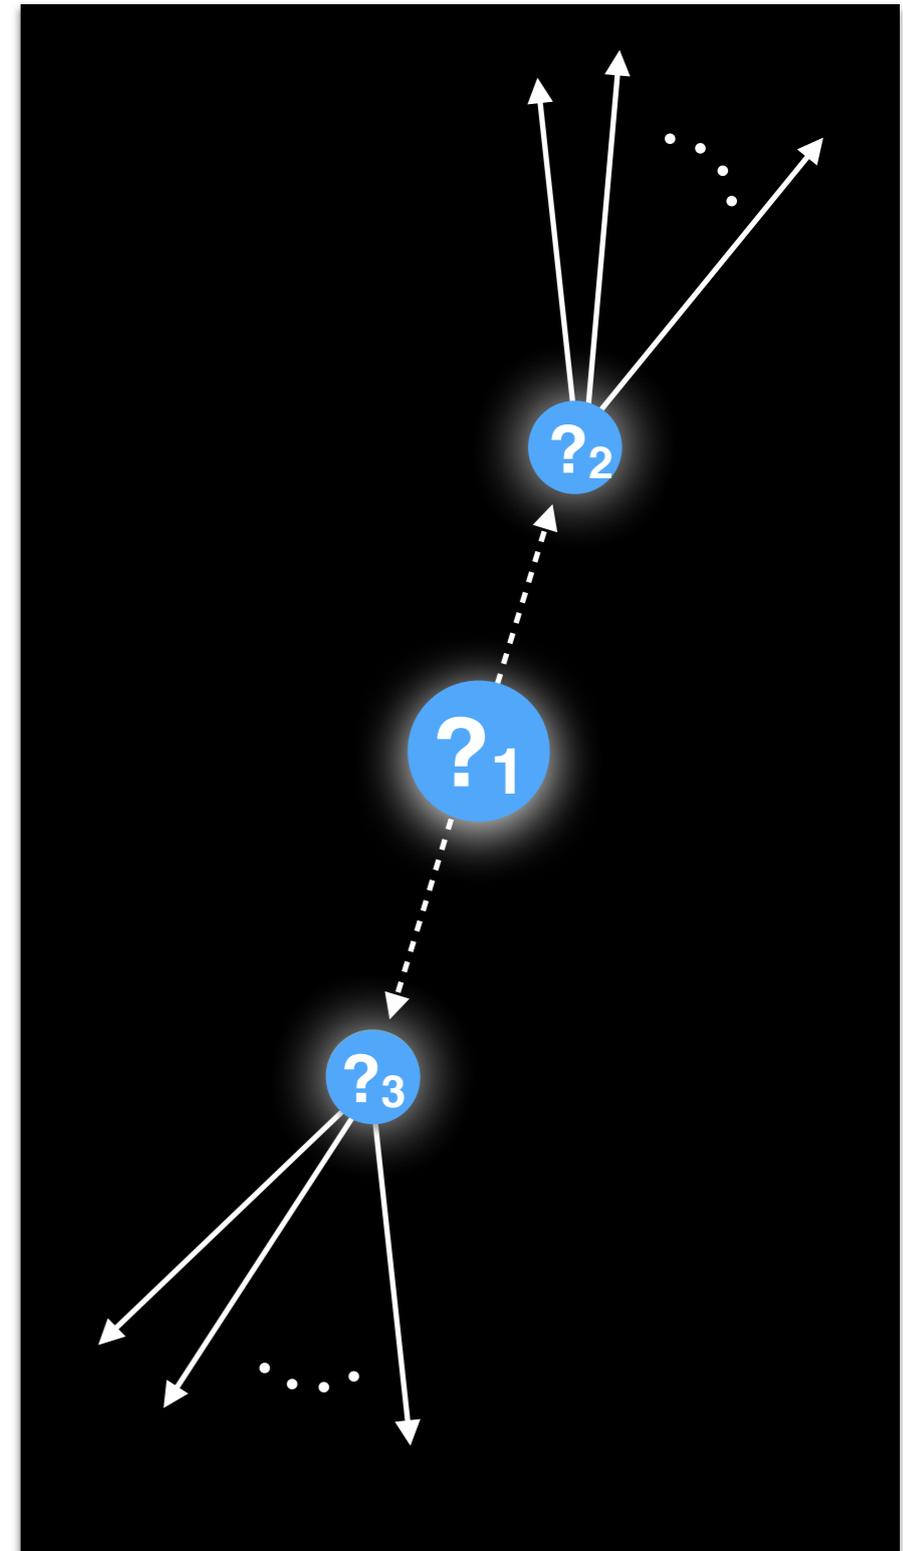
(I'll explain what supervision is momentarily)

The dataset

11



Dijet final state (allow for data-driven background + complex final state).



Timeline

12

April 4, 2019: R&D data released

Nov. 19, 2019: Black boxes released

Jan. 16, 2020: Winter Olympics + BB1 unblinded

July 16-17, 2020: Summer Olympics + BB2/3 unblinded

July 20, 2020: Invite contributions to community paper

Oct. 1, 2020: Contributions due

Jan. 20, 2021: Paper posted to arXiv

I don't have time to cover all of them - please see the paper for details! I'll just highlight some general ideas.

Section	Short Name	Method Type	Results Type
3.1	VRNN	Unsupervised	(i) (BB2,3) and (ii) (BB1)
3.2	ANODE	Unsupervised	(iii)
3.3	BuHuLaSpa	Unsupervised	(i) (BB2,3) and (ii) (BB1)
3.4	GAN-AE	Unsupervised	(i) (BB2-3) and (ii) (BB1)
3.5	GIS	Unsupervised	(i) (BB1)
3.6	LDA	Unsupervised	(i) (BB1-3)
3.7	PGA	Unsupervised	(ii) (BB1-2)
3.8	Reg. Likelihoods	Unsupervised	(iii)
3.9	UCluster	Unsupervised	(i) (BB2-3)
4.1	CWoLa	Weakly Supervised	(ii) (BB1-2)
4.2	CWoLa AE Compare	Weakly/Unsupervised	(iii)
4.3	Tag N' Train	Weakly Supervised	(i) (BB1-3)
4.4	SALAD	Weakly Supervised	(iii)
4.5	SA-CWoLa	Weakly Supervised	(iii)
5.1	Deep Ensemble	Semisupervised	(i) (BB1)
5.2	Factorized Topics	Semisupervised	(iii)
5.3	QUAK	Semisupervised	(i) (BB2,3) and (ii) (BB1)
5.4	LSTM	Semisupervised	(i) (BB1-3)

BB = black box; (i) = blinded, (ii) = unblinded

Supervision refers to the type of label information provided to the ML during training.

- Unsupervised** = no labels
- Weakly-supervised** = noisy labels
- Semi-supervised** = partial labels
- Supervised** = full label information

Section	Short Name	Method Type	Results Type
3.1	VRNN	Unsupervised	(i) (BB2,3) and (ii) (BB1)
3.2	ANODE	Unsupervised	(iii)
3.3	BuHuLaSpa	Unsupervised	(i) (BB2,3) and (ii) (BB1)
3.4	GAN-AE	Unsupervised	(i) (BB2-3) and (ii) (BB1)
3.5	GIS	Unsupervised	(i) (BB1)
3.6	LDA	Unsupervised	(i) (BB1-3)
3.7	PGA	Unsupervised	(ii) (BB1-2)
3.8	Reg. Likelihoods	Unsupervised	(iii)
3.9	UCluster	Unsupervised	(i) (BB2-3)
4.1	CWoLa	Weakly Supervised	(ii) (BB1-2)
4.2	CWoLa AE Compare	Weakly/Unsupervised	(iii)
4.3	Tag N' Train	Weakly Supervised	(i) (BB1-3)
4.4	SALAD	Weakly Supervised	(iii)
4.5	SA-CWoLa	Weakly Supervised	(iii)
5.1	Deep Ensemble	Semisupervised	(i) (BB1)
5.2	Factorized Topics	Semisupervised	(iii)
5.3	QUAK	Semisupervised	(i) (BB2,3) and (ii) (BB1)
5.4	LSTM	Semisupervised	(i) (BB1-3)

These categories are not exact and the boundaries are not rigid!

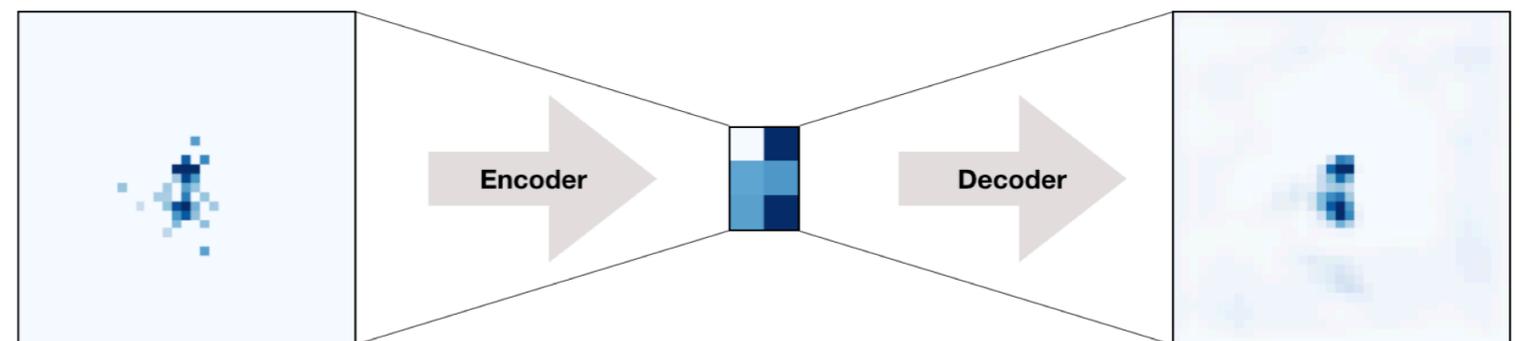
*N.B. Not everyone agrees on the boundary between semi-supervised and weakly supervised.

Solutions: Unsupervised

15

Unsupervised = no labels

Typically, the goal of these methods is to look for events with low $p(\text{background})$



Section	Short Name	Method Type	Results Type
3.1	VRNN	Unsupervised	(i) (BB2,3) and (ii) (BB1)
3.2	ANODE	Unsupervised	(iii)
3.3	BuHuLaSpa	Unsupervised	(i) (BB2,3) and (ii) (BB1)
3.4	GAN-AE	Unsupervised	(i) (BB2-3) and (ii) (BB1)
3.5	GIS	Unsupervised	(i) (BB1)
3.6	LDA	Unsupervised	(i) (BB1-3)
3.7	PGA	Unsupervised	(ii) (BB1-2)
3.8	Reg. Likelihoods	Unsupervised	(iii)
3.9	UCluster	Unsupervised	(i) (BB2-3)
4.1	CWoLa	Weakly Supervised	(ii) (BB1-2)
4.2	CWoLa AE Compare	Weakly/Unsupervised	(iii)
4.3	Tag N' Train	Weakly Supervised	(i) (BB1-3)
4.4	SALAD	Weakly Supervised	(iii)
4.5	SA-CWoLa	Weakly Supervised	(iii)
5.1	Deep Ensemble	Semisupervised	(i) (BB1)
5.2	Factorized Topics	Semisupervised	(iii)
5.3	QUAK	Semisupervised	(i) (BB2,3) and (ii) (BB1)
5.4	LSTM	Semisupervised	(i) (BB1-3)

One strategy (autoencoders) is to try to compress events and then uncompress them. When $x = \text{uncompress}(\text{compress}(x))$, then x probably has low $p(x)$.

M. Farina, Y. Nakai, D. Shih, 1808.08992; T. HeimeI, G. Kasieczka, T. Plehn, J. Thompson, 1808.08979; + many more

Solutions: Weakly-supervised

Weakly-supervised = noisy labels

Typically, the goal of these methods is to look for events with high $p(\text{possibly signal-enriched})/p(\text{possibly signal-depleted})$

e.g. Classification Without Labels (CWoLa), events in a signal region are labeled “signal” and events in a sideband are labeled “background”. These labels are “noisy” but a classifier trained with them can detect the presence of a signal.

Section	Short Name	Method Type	Results Type
3.1	VRNN	Unsupervised	(i) (BB2,3) and (ii) (BB1)
3.2	ANODE	Unsupervised	(iii)
3.3	BuHuLaSpa	Unsupervised	(i) (BB2,3) and (ii) (BB1)
3.4	GAN-AE	Unsupervised	(i) (BB2-3) and (ii) (BB1)
3.5	GIS	Unsupervised	(i) (BB1)
3.6	LDA	Unsupervised	(i) (BB1-3)
3.7	PGA	Unsupervised	(ii) (BB1-2)
3.8	Reg. Likelihoods	Unsupervised	(iii)
3.9	UCluster	Unsupervised	(i) (BB2-3)
4.1	CWoLa	Weakly Supervised	(ii) (BB1-2)
4.2	CWoLa AE Compare	Weakly/Unsupervised	(iii)
4.3	Tag N' Train	Weakly Supervised	(i) (BB1-3)
4.4	SALAD	Weakly Supervised	(iii)
4.5	SA-CWoLa	Weakly Supervised	(iii)
5.1	Deep Ensemble	Semisupervised	(i) (BB1)
5.2	Factorized Topics	Semisupervised	(iii)
5.3	QUAK	Semisupervised	(i) (BB2,3) and (ii) (BB1)
5.4	LSTM	Semisupervised	(i) (BB1-3)

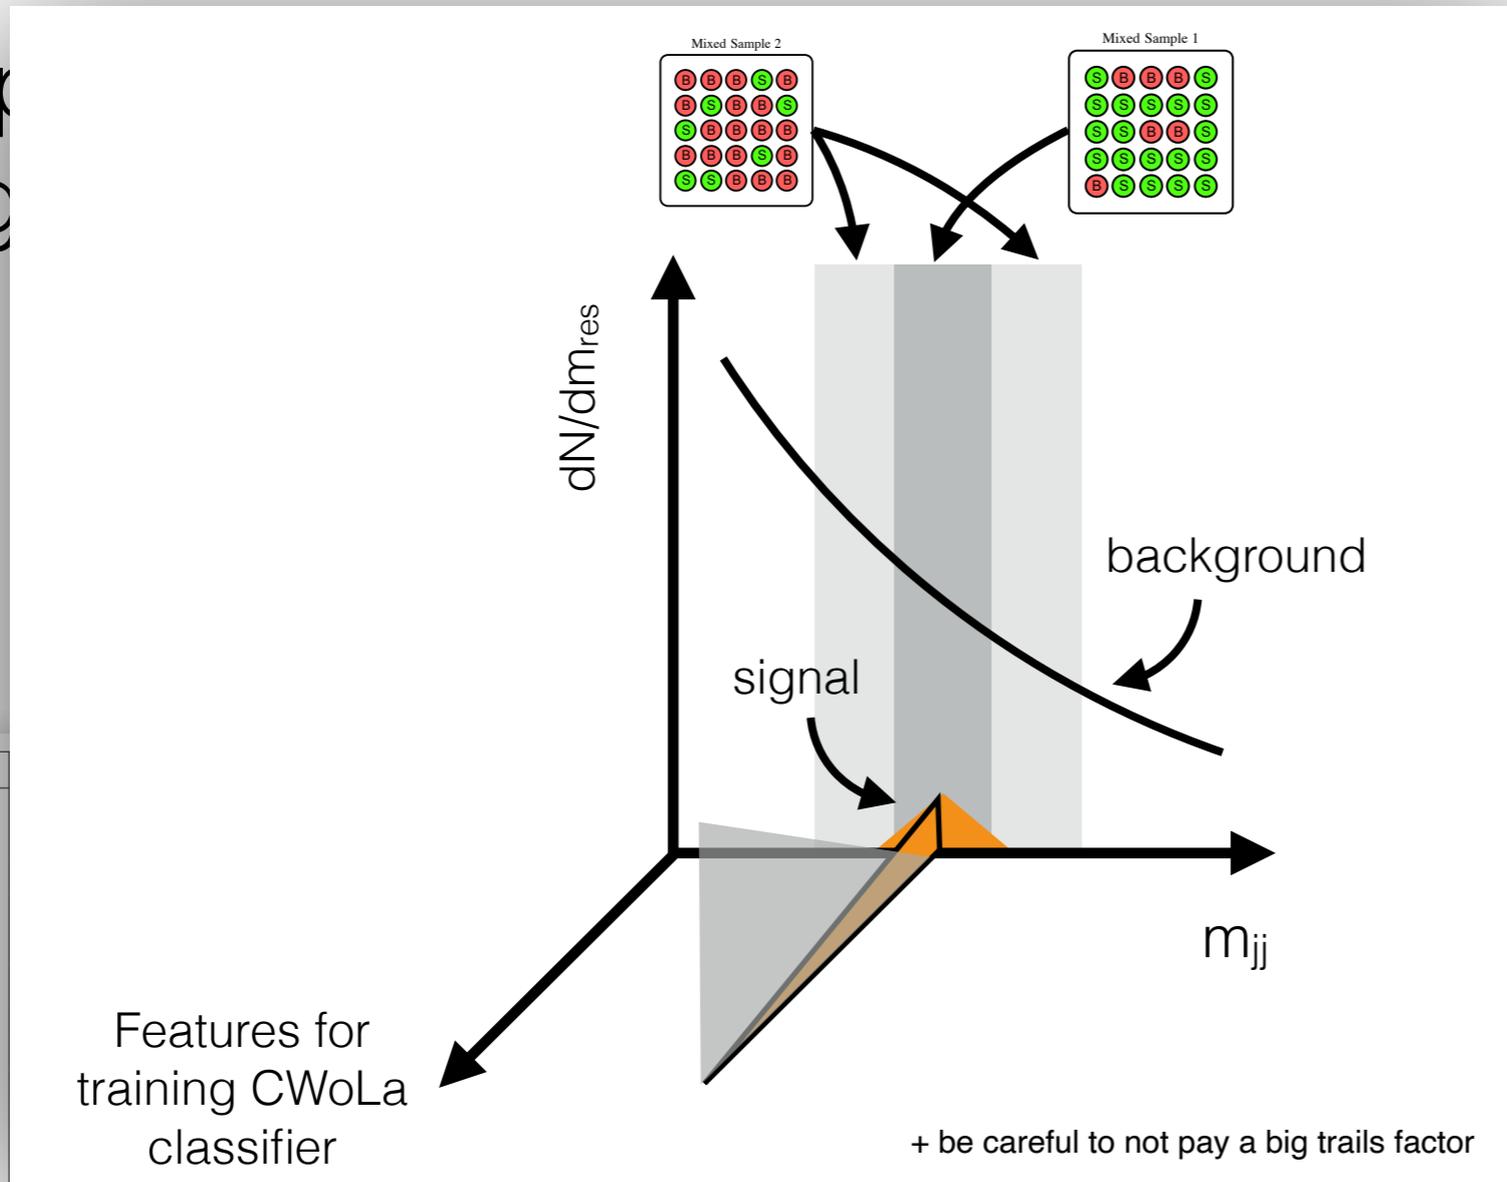
Solutions: Weakly-supervised

Weakly-supervised = noisy labels

Typ
high

to look for events with
possibly signal-depleted)

Classification Without Labels
events in a signal region are
" and events in a sideband
background". These labels
but a classifier trained with
ct the presence of a signal.



Section

- 3.1
- 3.2
- 3.3
- 3.4
- 3.5
- 3.6
- 3.7
- 3.8
- 3.9
- 4.1
- 4.2
- 4.3
- 4.4
- 4.5
- 5.1
- 5.2
- 5.3
- 5.4

Tag N' Train	Weakly Supervised	(i) (BB1-3)
SALAD	Weakly Supervised	(iii)
SA-CWoLa	Weakly Supervised	(iii)
Deep Ensemble	Semisupervised	(i) (BB1)
Factorized Topics	Semisupervised	(iii)
QUAK	Semisupervised	(i) (BB2,3) and (ii) (BB1)
LSTM	Semisupervised	(i) (BB1-3)

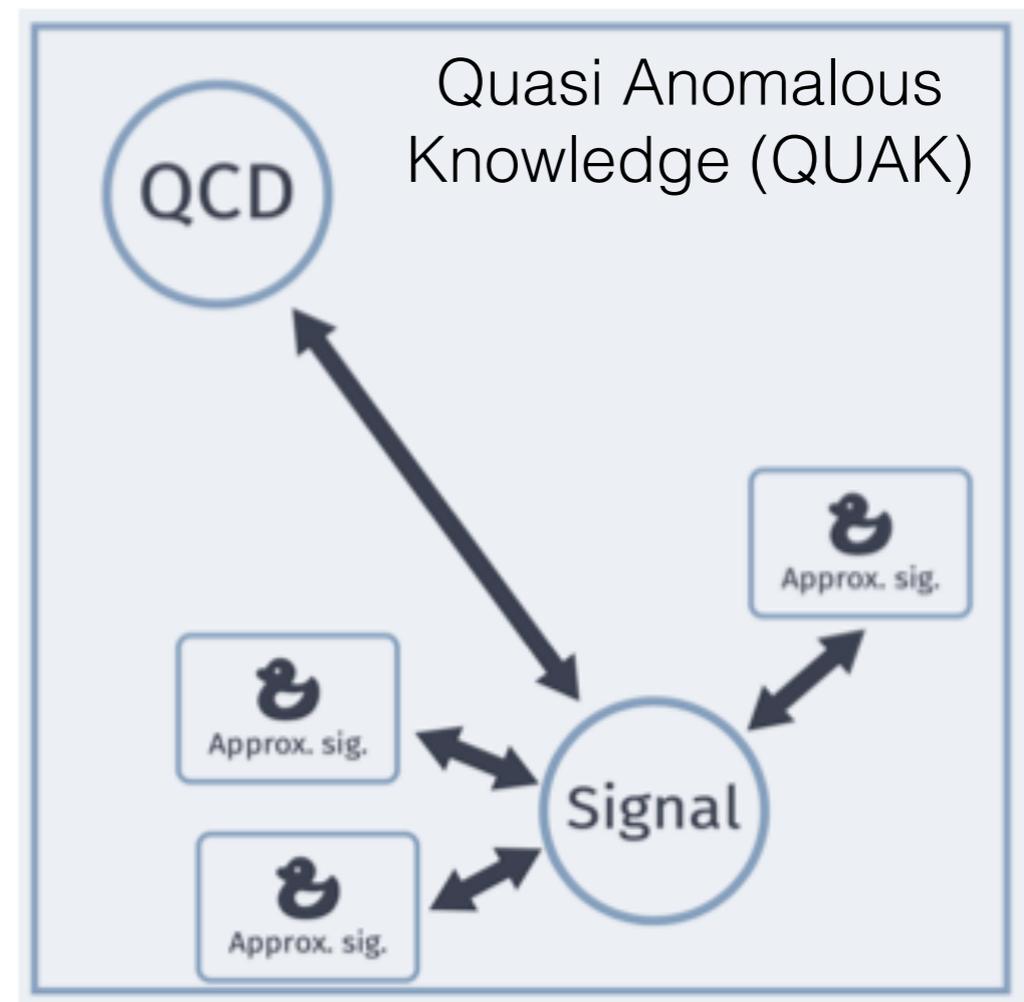
Solutions: Semi-supervised

Semi-supervised = partial labels

Typically, these methods use some signal simulations to build signal sensitivity

(We did not give bonus points for the best acronyms !)

Section	Short Name	Method Type	Results Type
3.1	VRNN	Unsupervised	(i) (BB2,3) and (ii) (BB1)
3.2	ANODE	Unsupervised	(iii)
3.3	BuHuLaSpa	Unsupervised	(i) (BB2,3) and (ii) (BB1)
3.4	GAN-AE	Unsupervised	(i) (BB2-3) and (ii) (BB1)
3.5	GIS	Unsupervised	(i) (BB1)
3.6	LDA	Unsupervised	(i) (BB1-3)
3.7	PGA	Unsupervised	(ii) (BB1-2)
3.8	Reg. Likelihoods	Unsupervised	(iii)
3.9	UCluster	Unsupervised	(i) (BB2-3)
4.1	CWoLa	Weakly Supervised	(ii) (BB1-2)
4.2	CWoLa AE Compare	Weakly/Unsupervised	(iii)
4.3	Tag N' Train	Weakly Supervised	(i) (BB1-3)
4.4	SALAD	Weakly Supervised	(iii)
4.5	SA-CWoLa	Weakly Supervised	(iii)
5.1	Deep Ensemble	Semisupervised	(i) (BB1)
5.2	Factorized Topics	Semisupervised	(iii)
5.3	QUAK	Semisupervised	(i) (BB2,3) and (ii) (BB1)
5.4	LSTM	Semisupervised	(i) (BB1-3)



Results on the blinded* challenge

19

Black Box 1

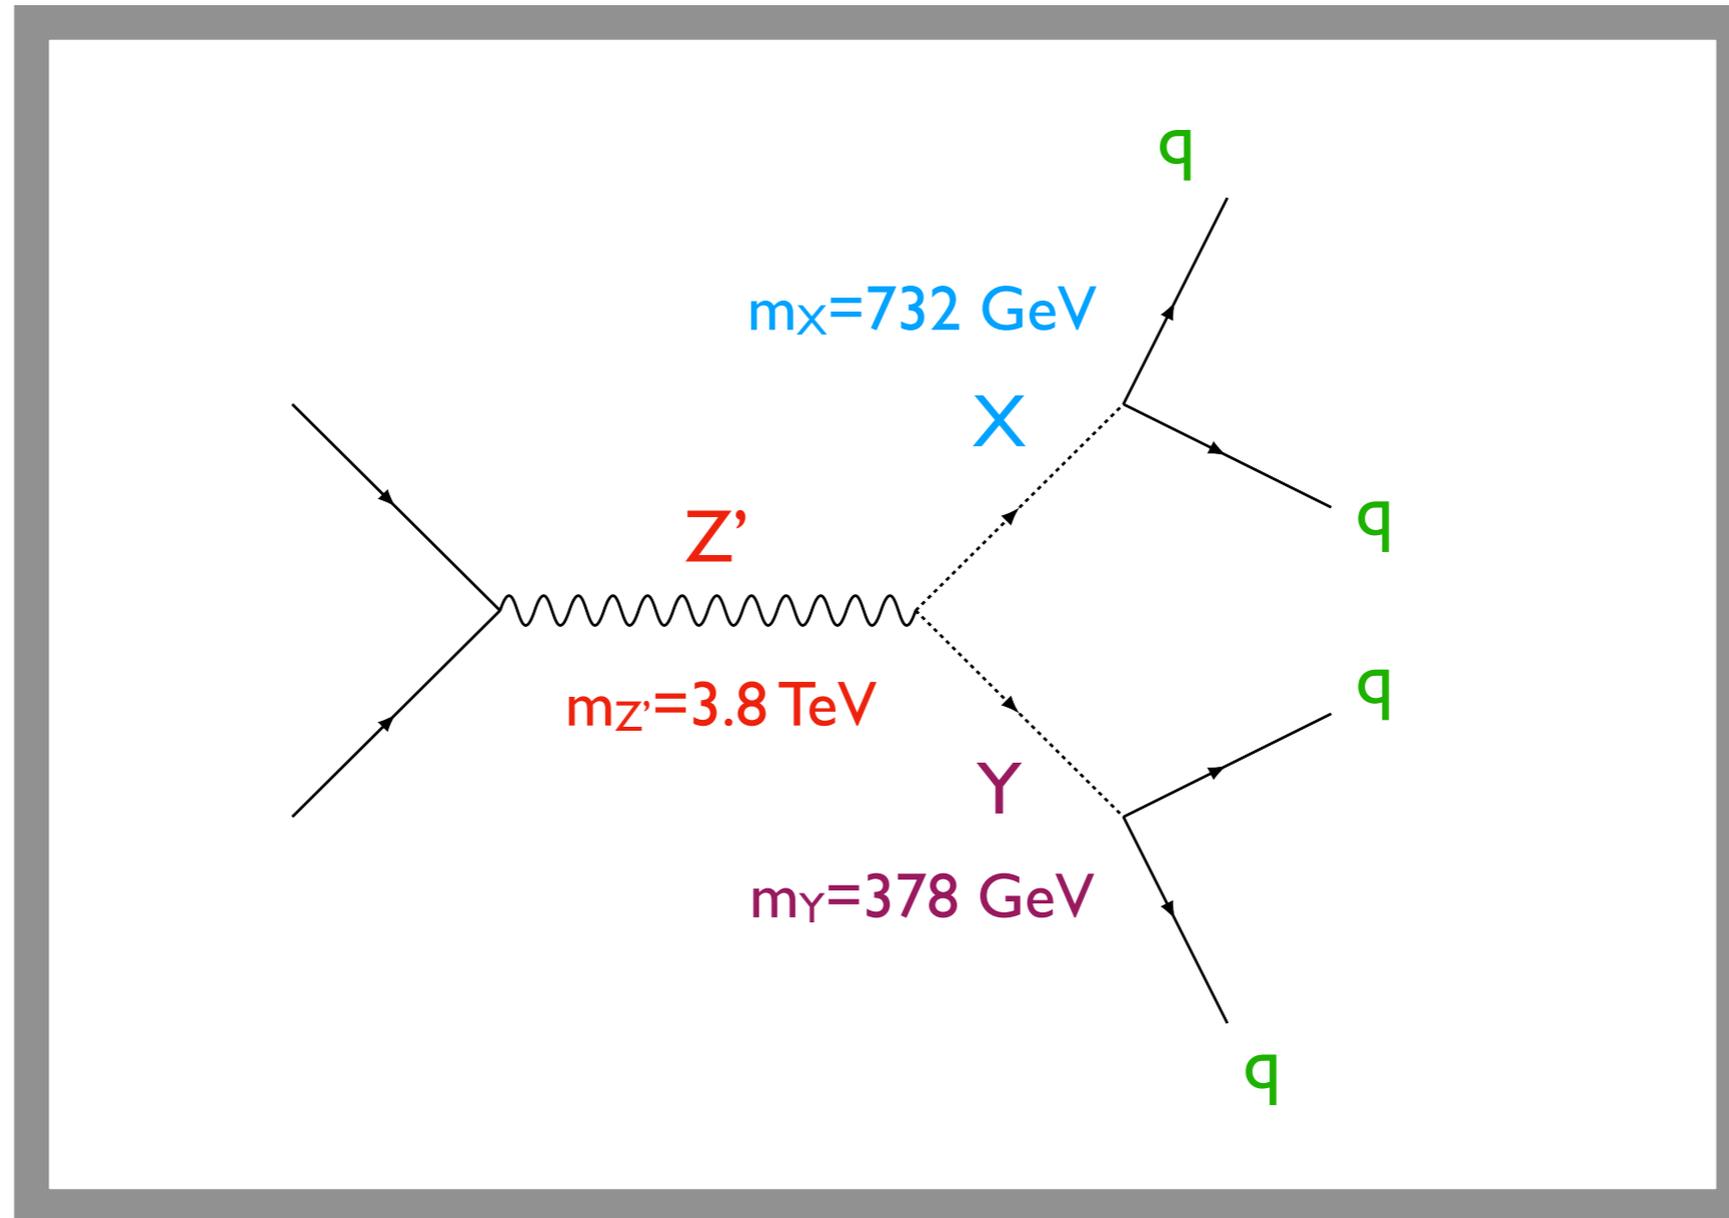


*There are a lot more results in the summary paper (!)

Results on the blinded challenge

20

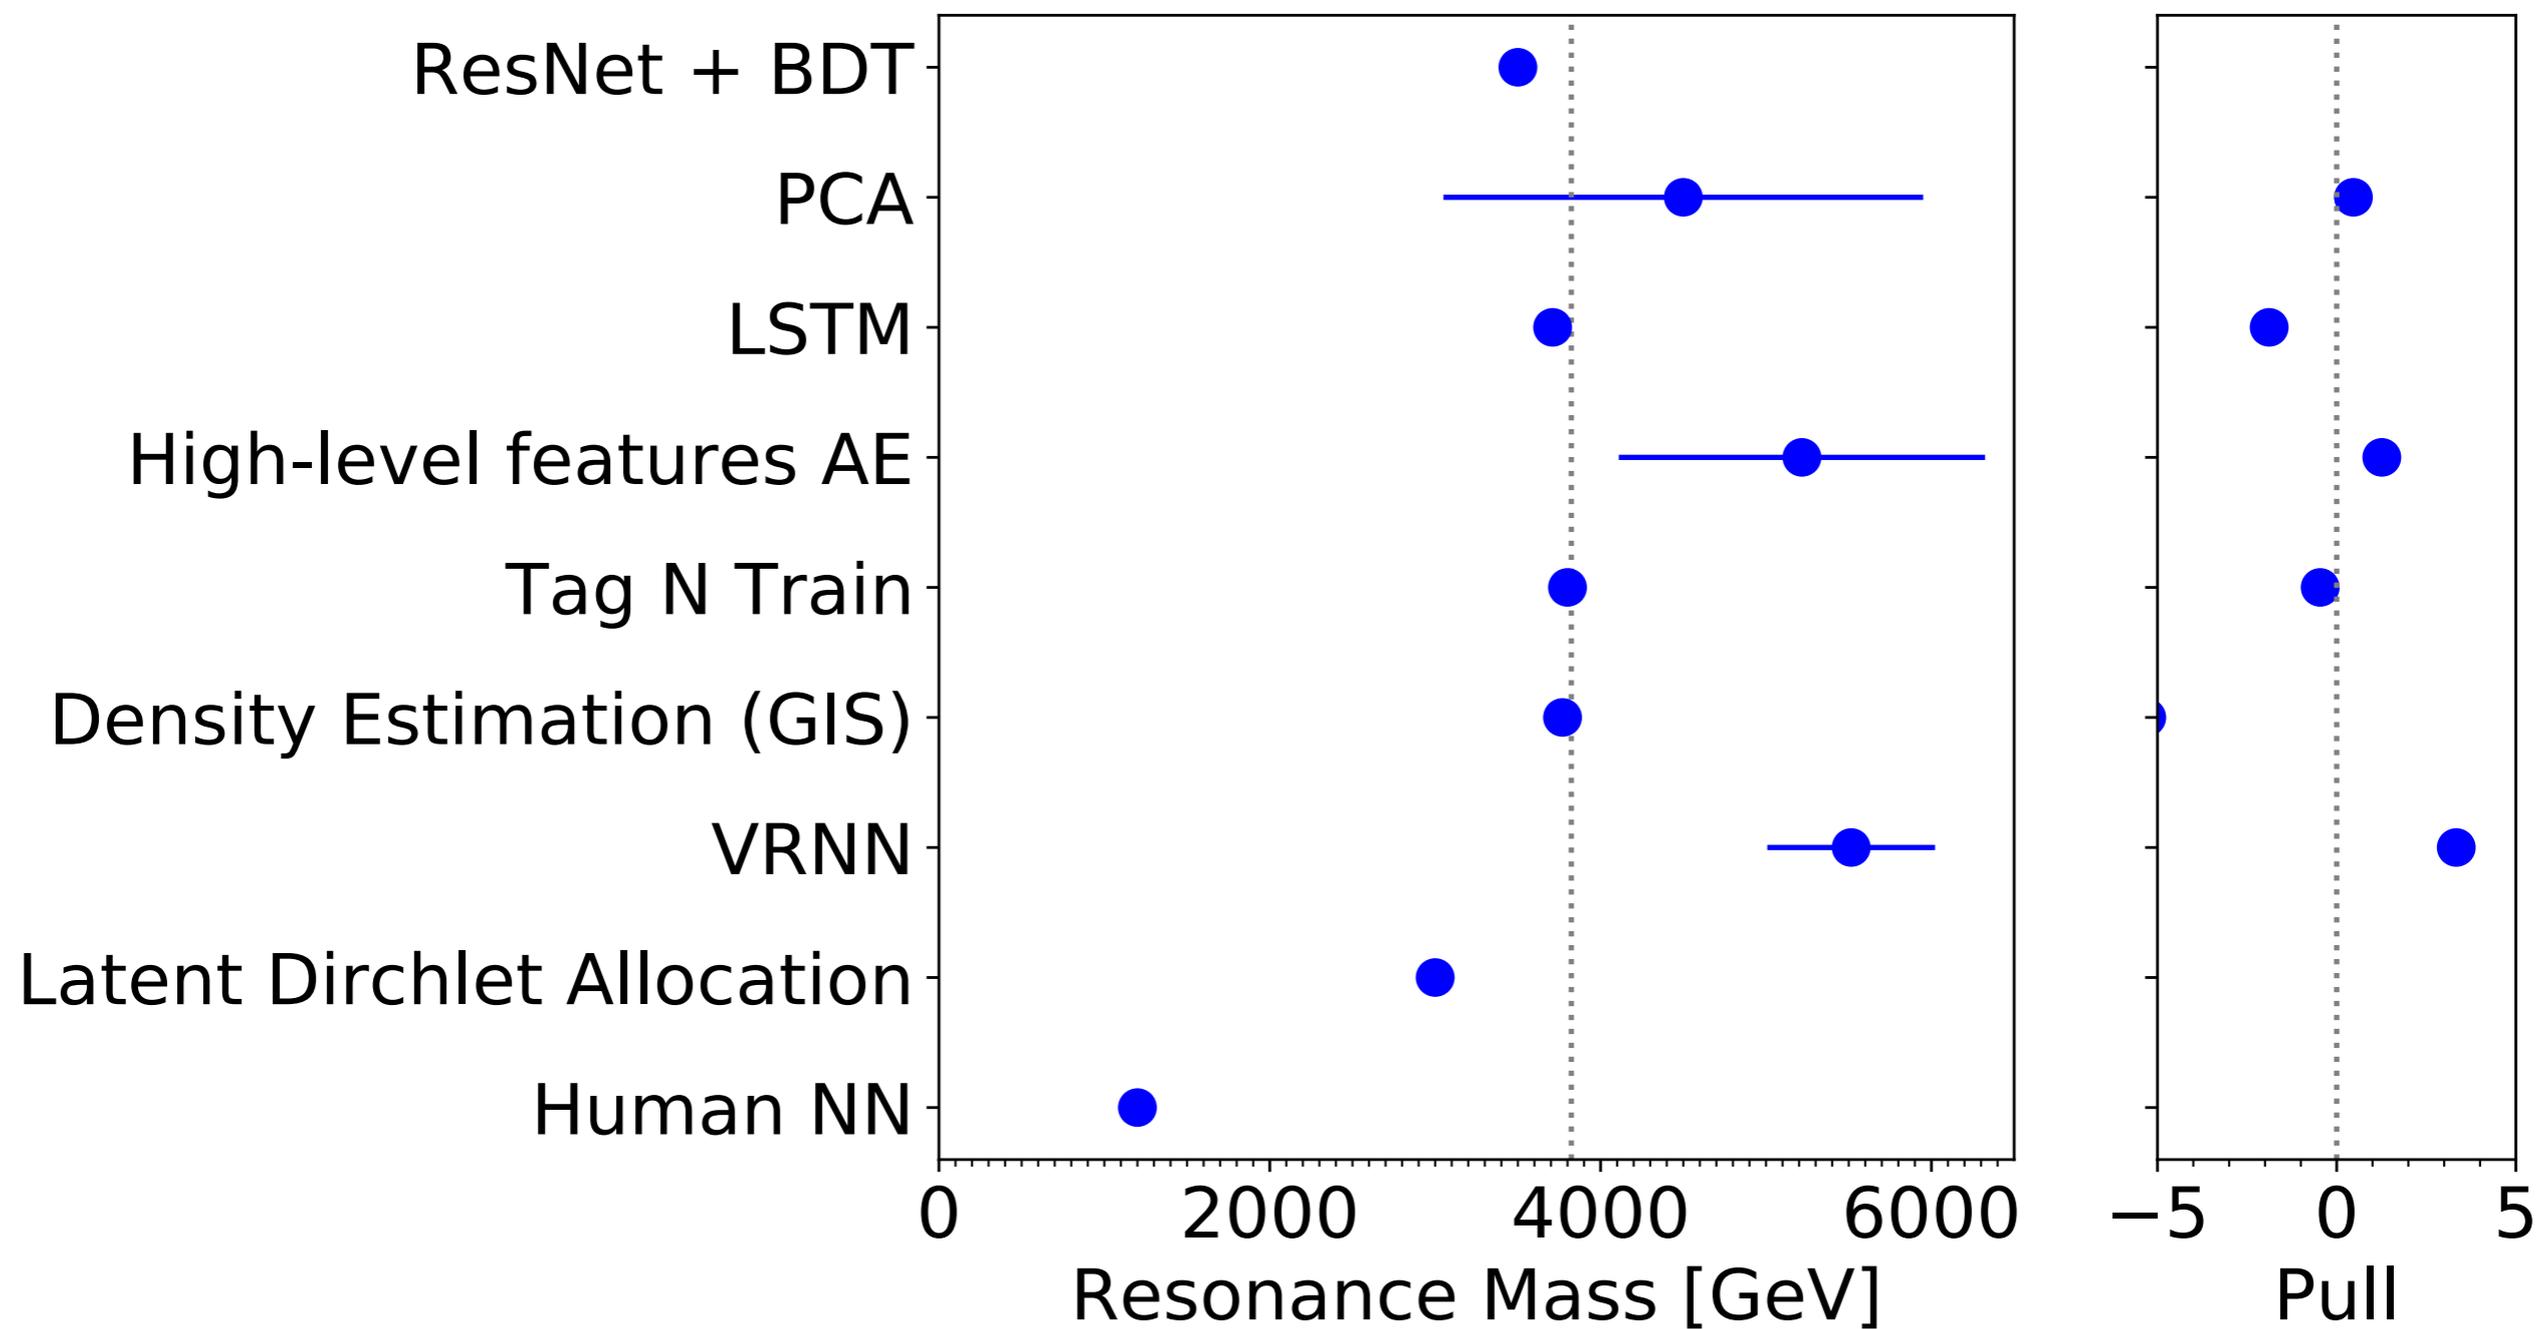
Black Box 1



834 events. Same topology as R&D dataset
(not known to participants)

Results on the blinded challenge

21



...and the winners are...

22

We refrained from ranking the results ... all methods are an important contribution to this growing research area.

However, two submissions clearly stood out:

Conditional density estimation for anomaly detection (GIS)

George Stein, Uroš Seljak, Biwei Dai, He Jia

Tag N' Train
(CWoLa + Autoencoder)

Oz Amram and Cristina Mantilla Suarez

Results on the blinded challenge

23

Black Box 2



Results on the blinded challenge

24

Black Box 2



Empty! Multiple teams still reported (fake) signals...

Results on the blinded challenge

25

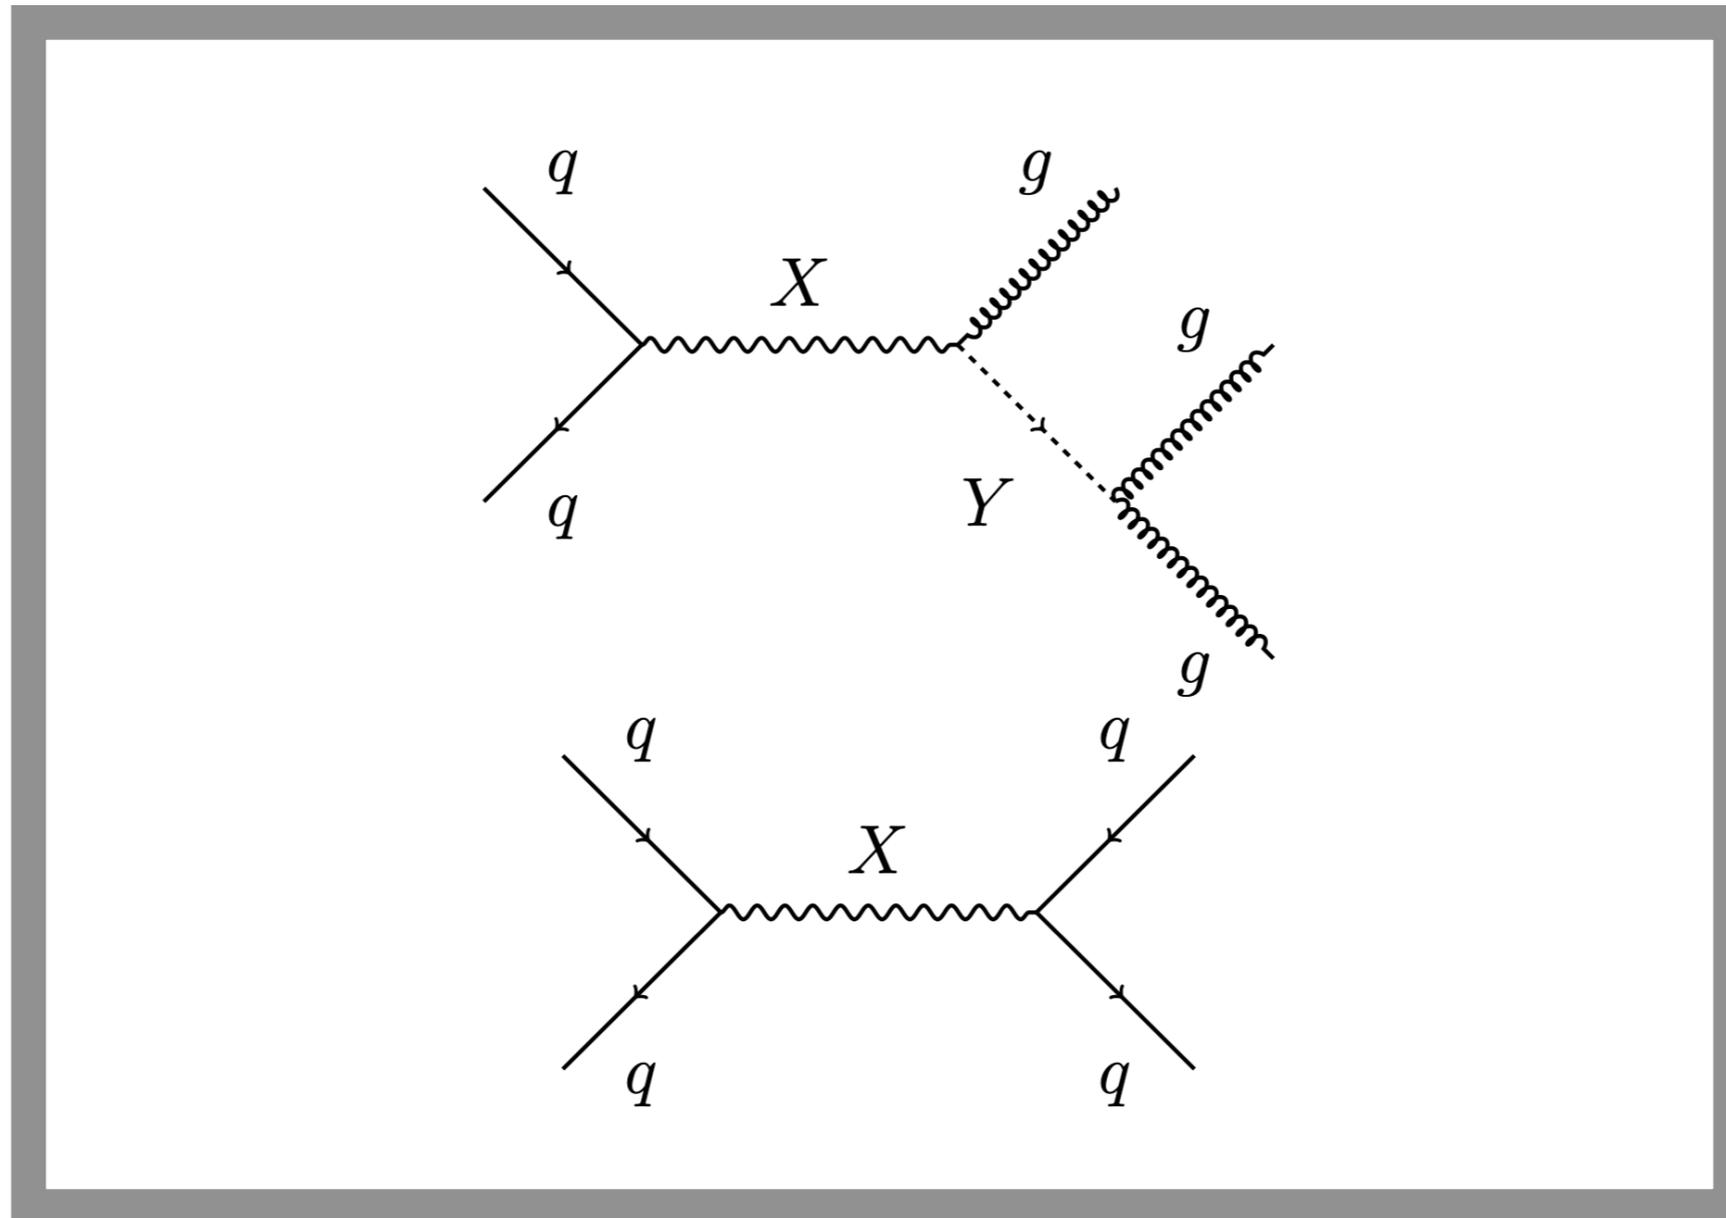
Black Box 3



Results on the blinded challenge

26

Black Box 3

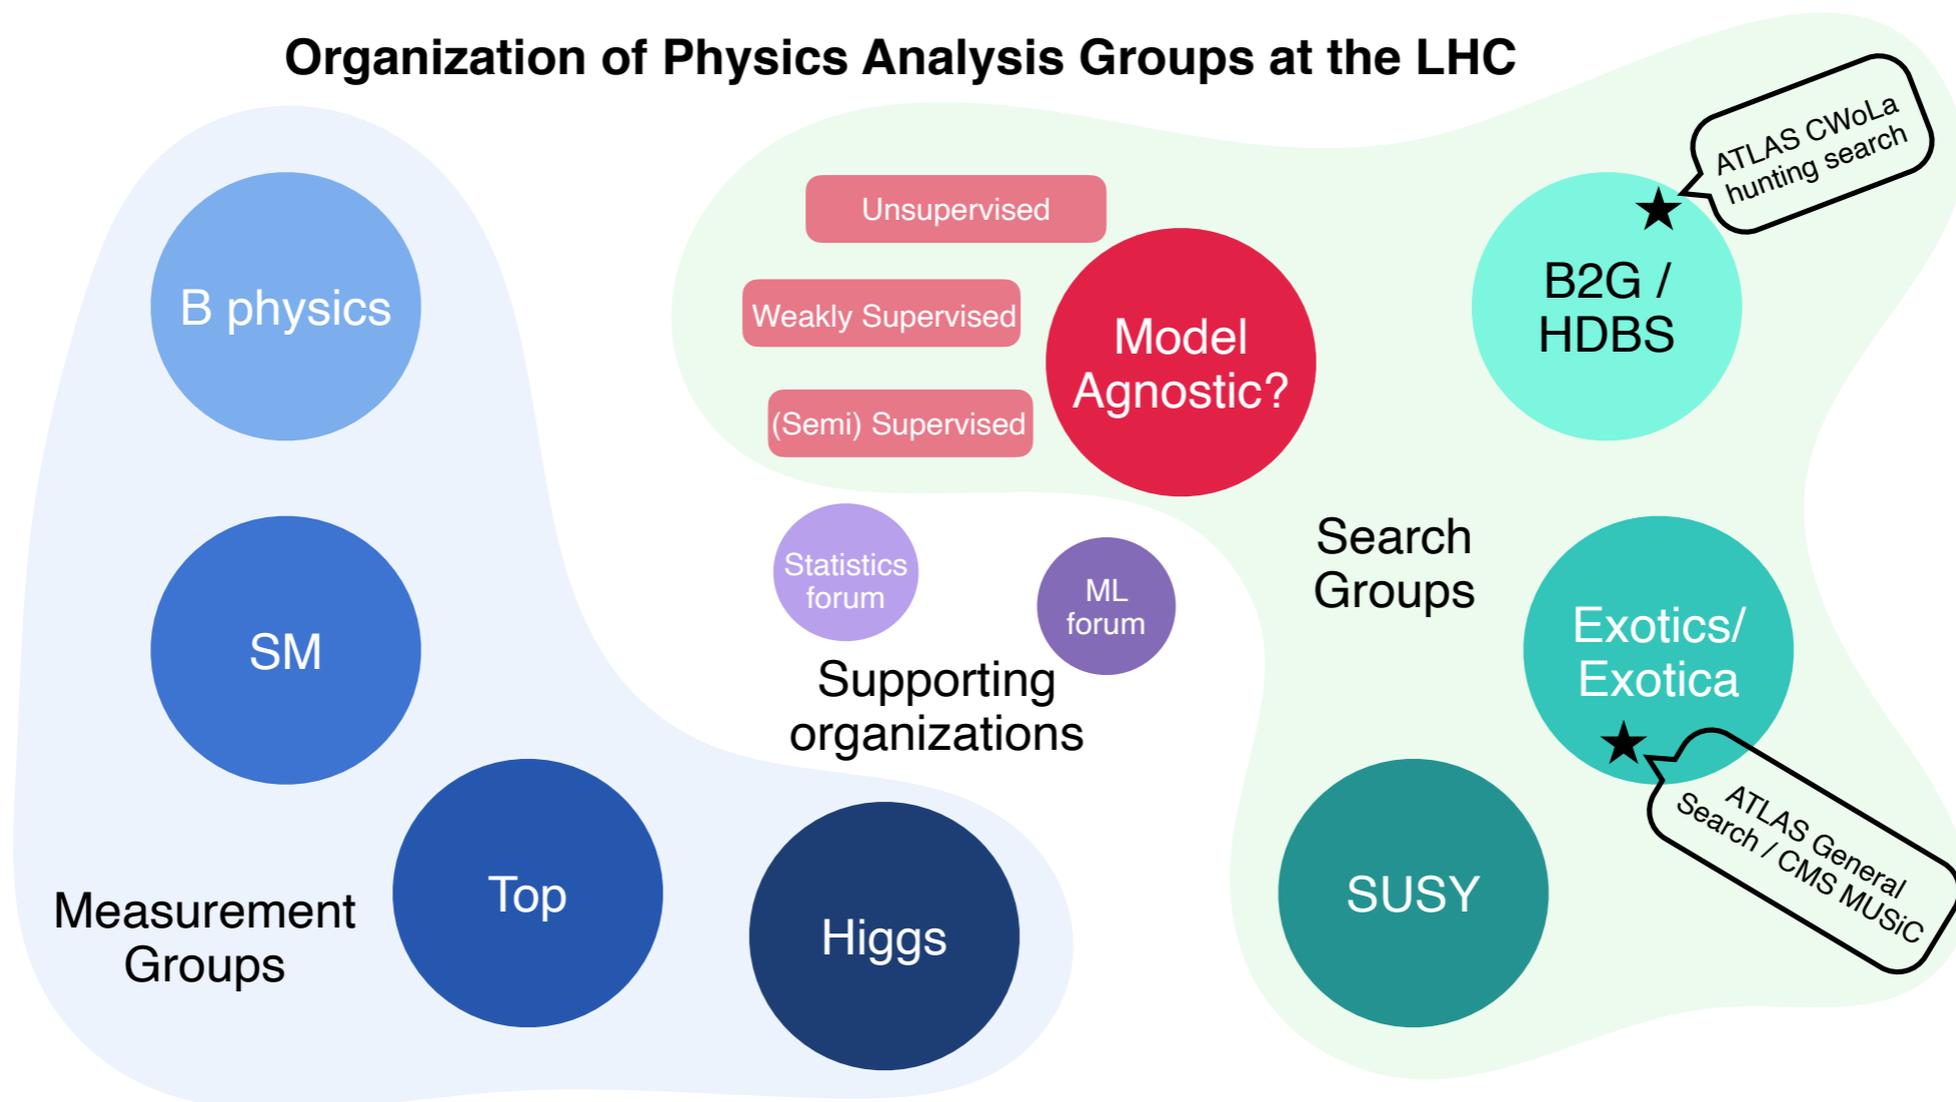


This one was difficult - there are **two decay modes**. Without finding both, the signal is not significant. **No one** convincingly found this signal (**more work** is needed!)

That was a quick overview - please see the paper for many more details. Each section has the same format and is relatively brief (often further details are provided in standalone papers)

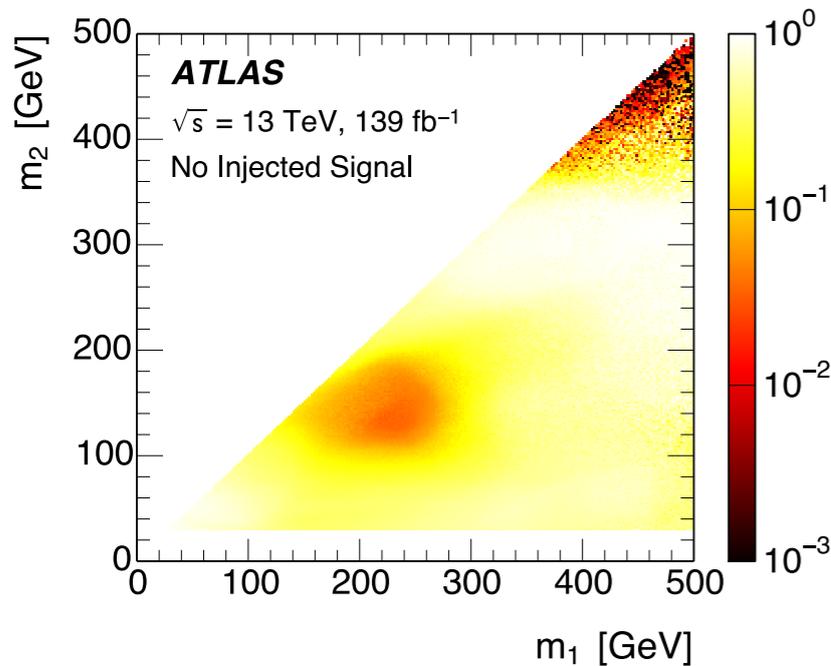
Section	Short Name	Method Type	Results Type
3.1	VRNN	Unsupervised	(i) (BB2,3) and (ii) (BB1)
3.2	ANODE	Unsupervised	(iii)
3.3	BuHuLaSpa	Unsupervised	(i) (BB2,3) and (ii) (BB1)
3.4	GAN-AE	Unsupervised	(i) (BB2-3) and (ii) (BB1)
3.5	GIS	Unsupervised	(i) (BB1)
3.6	LDA	Unsupervised	(i) (BB1-3)
3.7	PGA	Unsupervised	(ii) (BB1-2)
3.8	Reg. Likelihoods	Unsupervised	(iii)
3.9	UCluster	Unsupervised	(i) (BB2-3)
4.1	CWoLa	Weakly Supervised	(ii) (BB1-2)
4.2	CWoLa AE Compare	Weakly/Unsupervised	(iii)
4.3	Tag N' Train	Weakly Supervised	(i) (BB1-3)
4.4	SALAD	Weakly Supervised	(iii)
4.5	SA-CWoLa	Weakly Supervised	(iii)
5.1	Deep Ensemble	Semisupervised	(i) (BB1)
5.2	Factorized Topics	Semisupervised	(iii)
5.3	QUAK	Semisupervised	(i) (BB2,3) and (ii) (BB1)
5.4	LSTM	Semisupervised	(i) (BB1-3)

This challenge identified new methods and exposed gaps in coverage. It is clear that we will need **multiple approaches** to achieve broad coverage.



Outlook: the path to LHC data

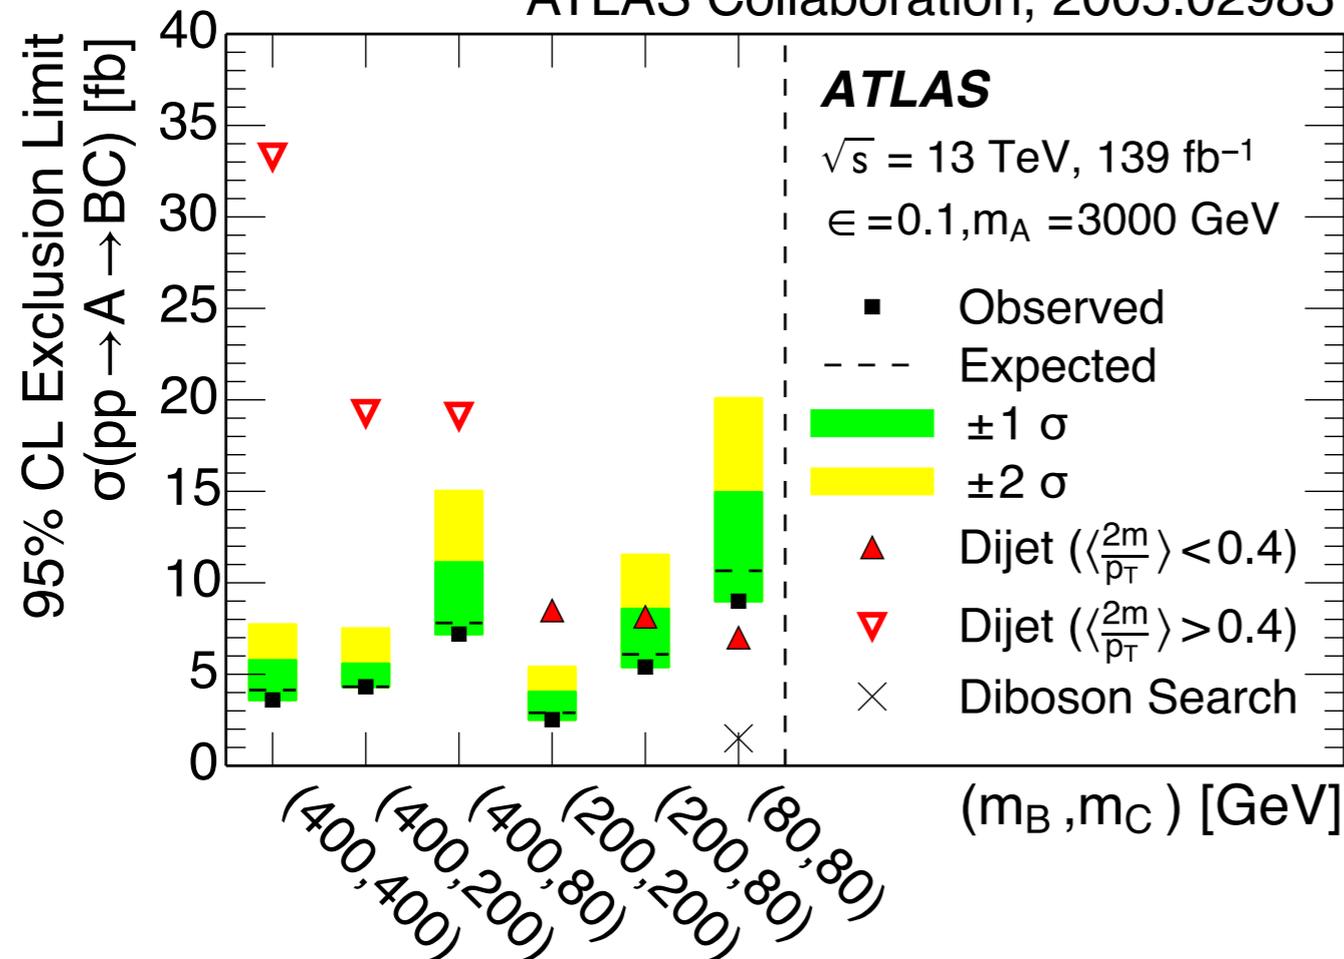
29



While we clearly need more R&D, it is also time to start deploying these methods on LHC data !

There are always new challenges with real data

ATLAS Collaboration, 2005.02983



I am proud to say that ATLAS recently completed a first analysis and I am hopeful that this is the start of many more to come!

A last thought: reinterpretation

30

After all, this is a reinterpretation workshop ...
these analyses are notoriously **difficult to reinterpret** as the “event selection” often depends on the data (unique to these analyses!).

For every signal region...

For every signal model...

For every signal cross section... *(probably want to repeat for low n)*

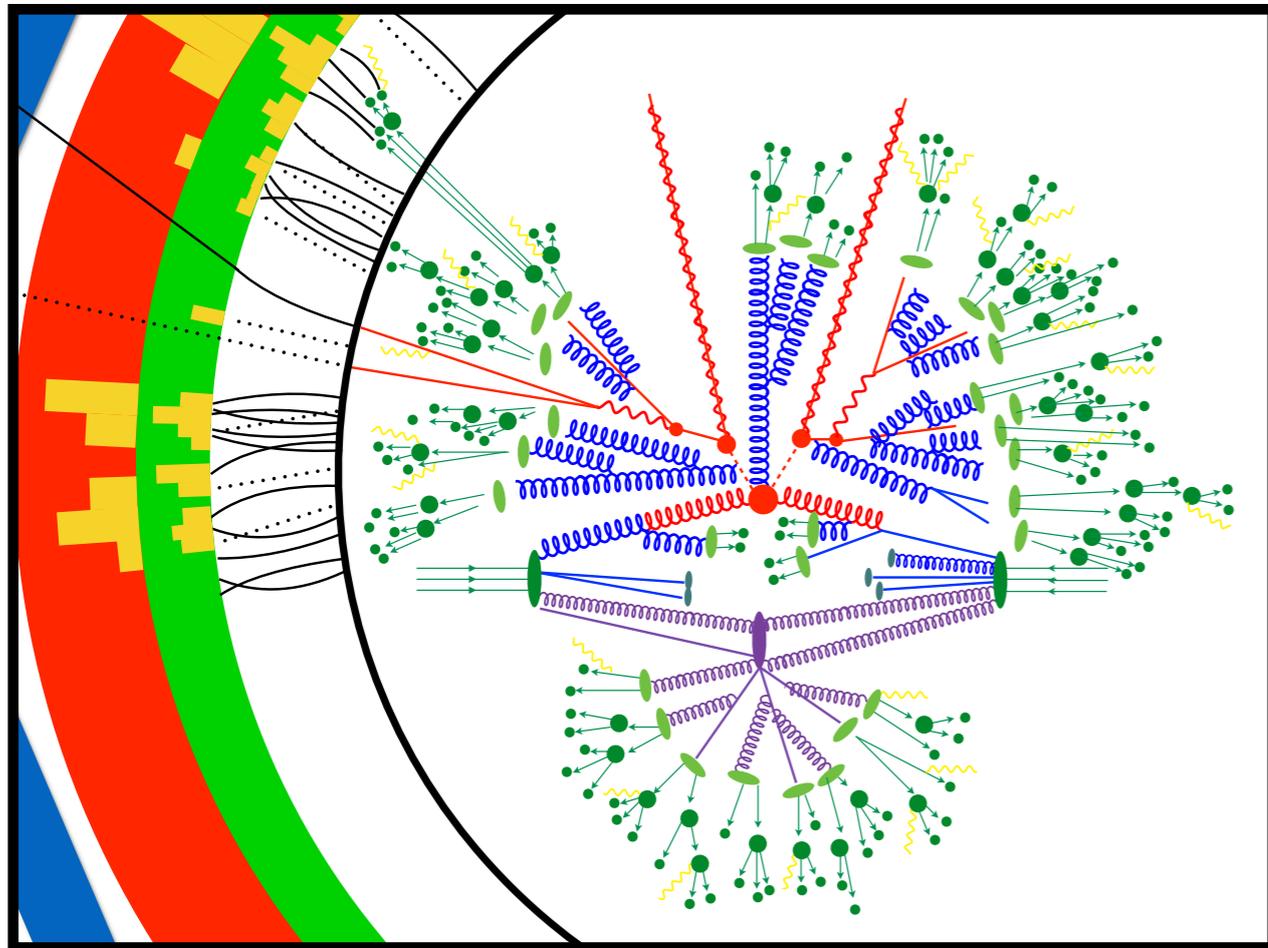
For every signal systematic uncertainty...

Inject signal, rerun entire analysis

*including whatever ensembling
is inside the training procedure!*

This is an interesting challenge we will have to address in the future.

It is critical that we complement the current search program with **model agnostic methods**



Machine learning provides many powerful solutions, which were prototyped on the **LHC Olympics dataset**

...and these tools will increase the LHC discovery potential!

Questions?

