The road to AI-based discoveries

Gregor Kasieczka
Email: gregor.kasieczka@uni-hamburg.de
in Gregor Kasieczka

Quantum100 ⊗ ai, 12.11.2025

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

CLUSTER OF EXCELLENCE
QUANTUM UNIVERSE

Outline

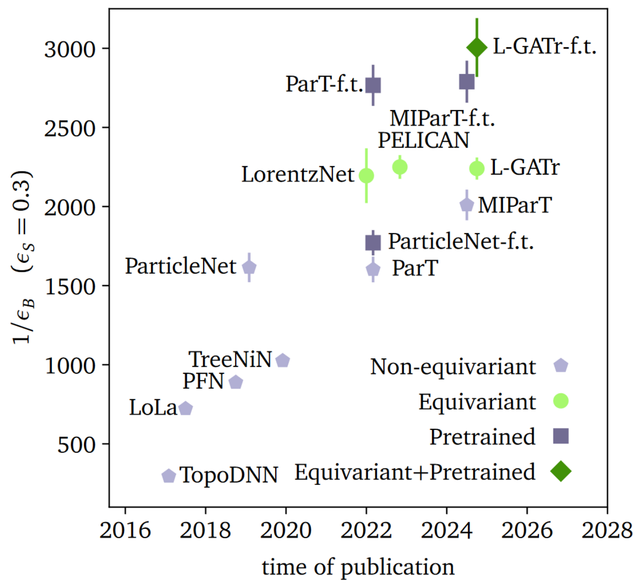

Tools for Discovery

Discovery Strategies

Autonomous Discovery

Increasing autonomy of AI systems

How to discover a new particle?

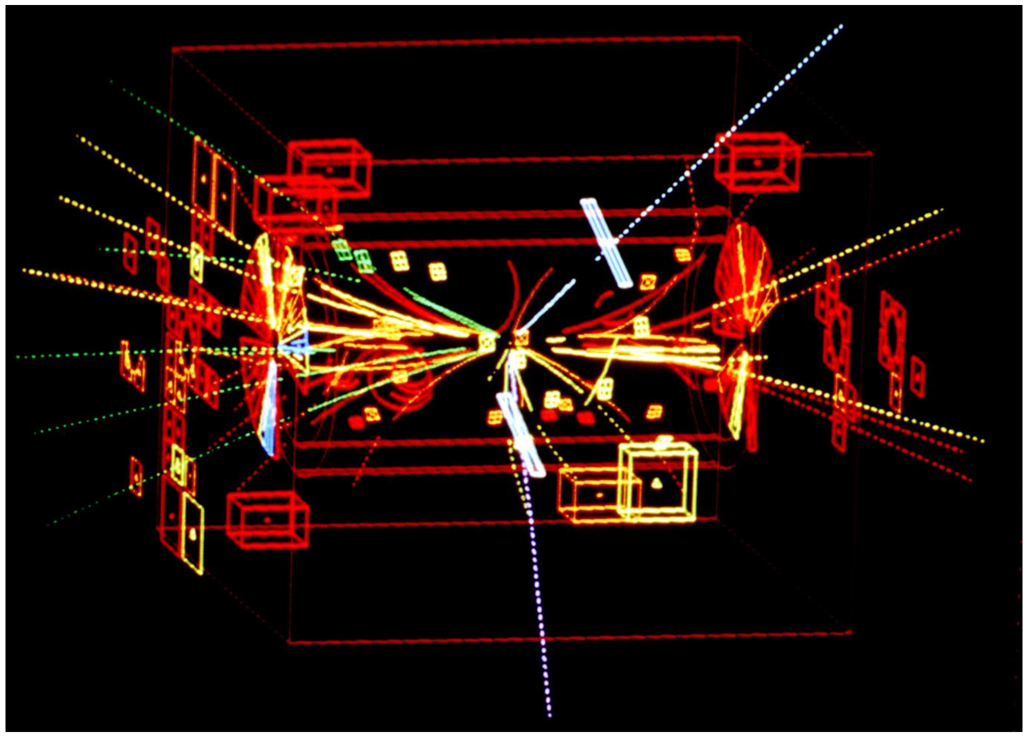

Event display from UA1 from 1983

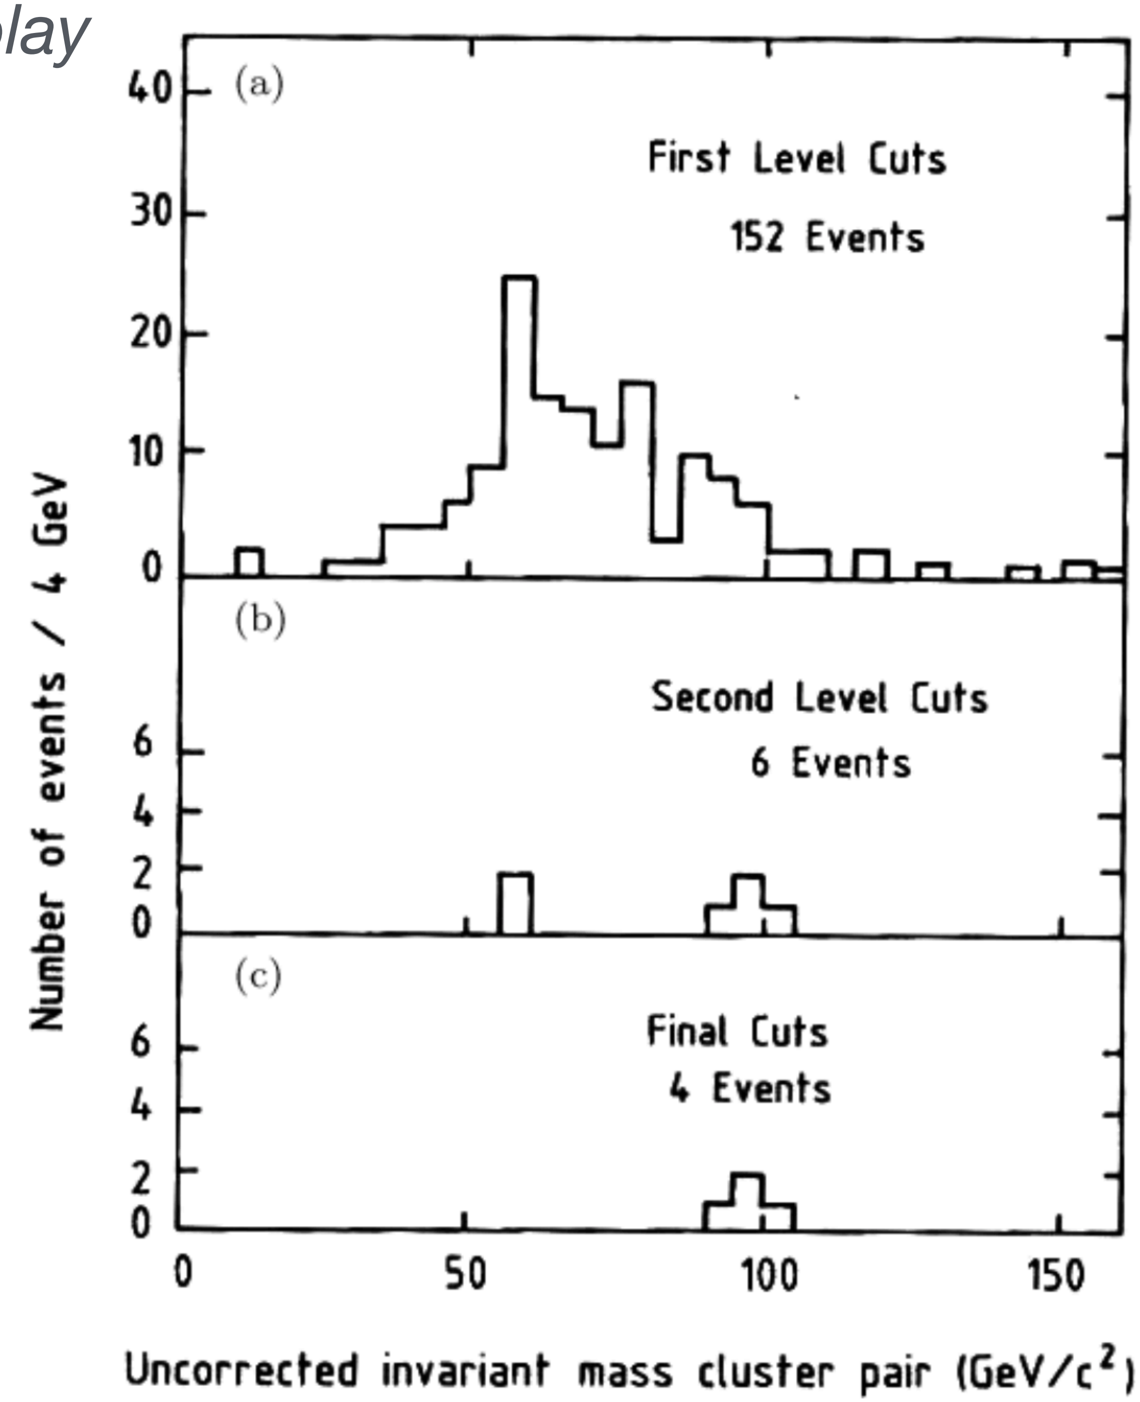

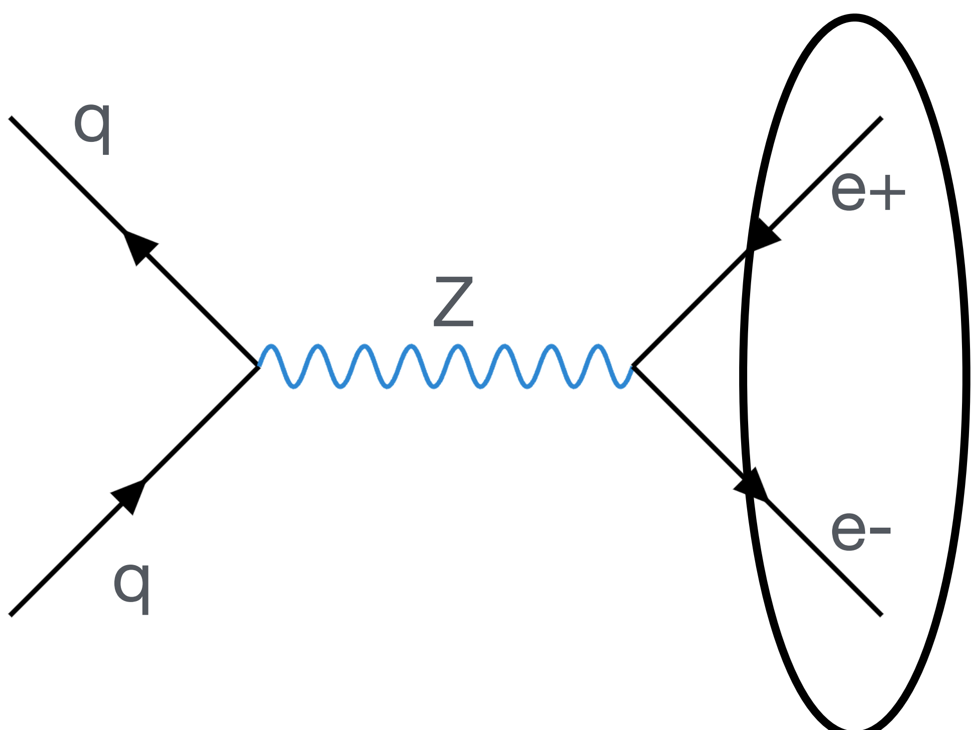

Calculate invariant mass of two particles to discover resonance

How to discover a new particle?

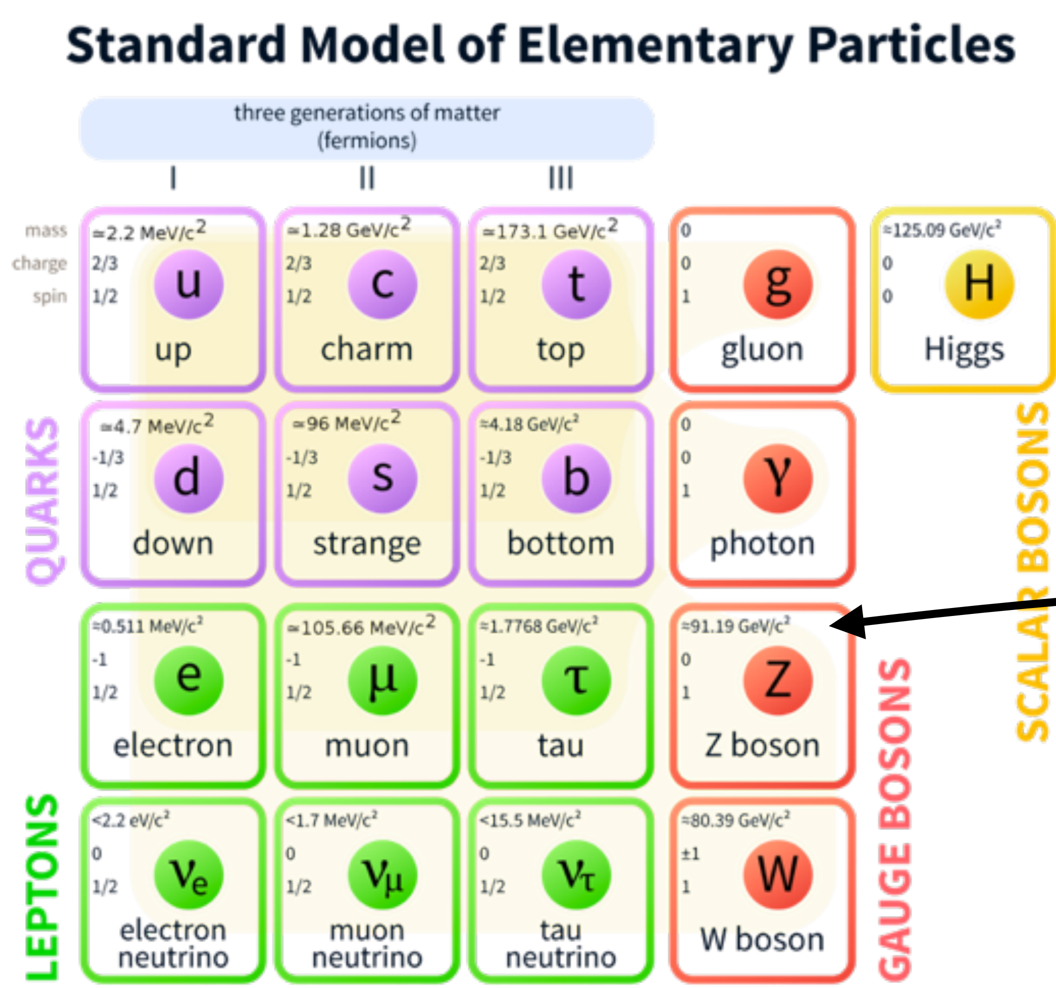

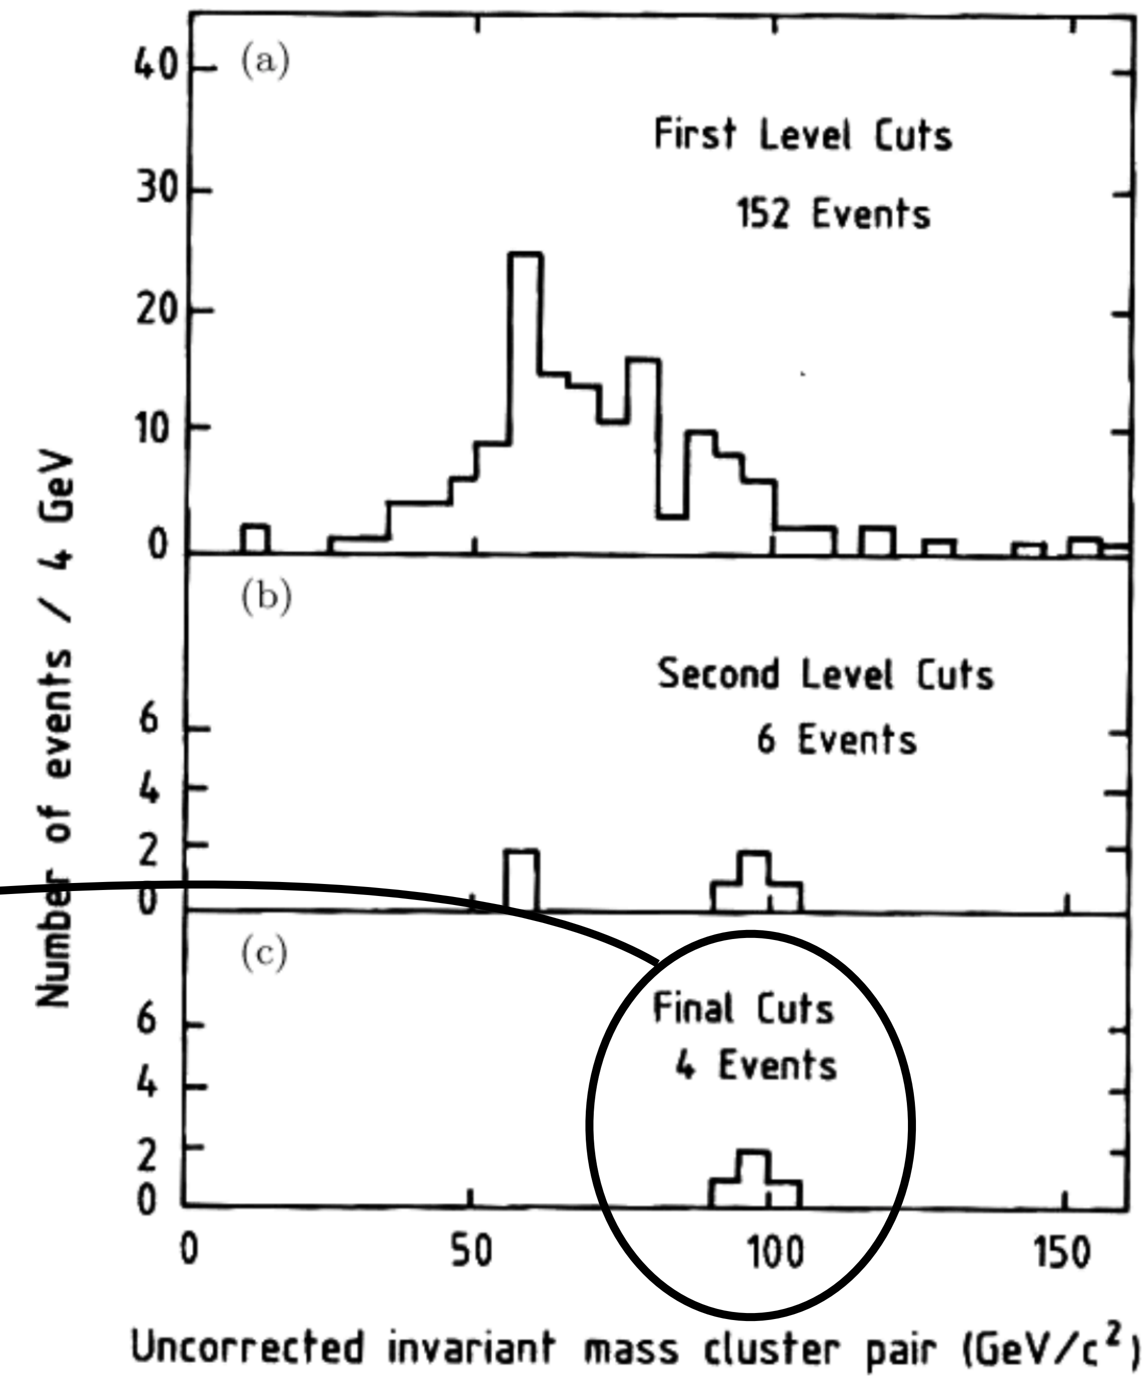

Nobel price for discovery of W and Z boson in 1984 (Rubbia and van der Meer)

How to discover a new particle?

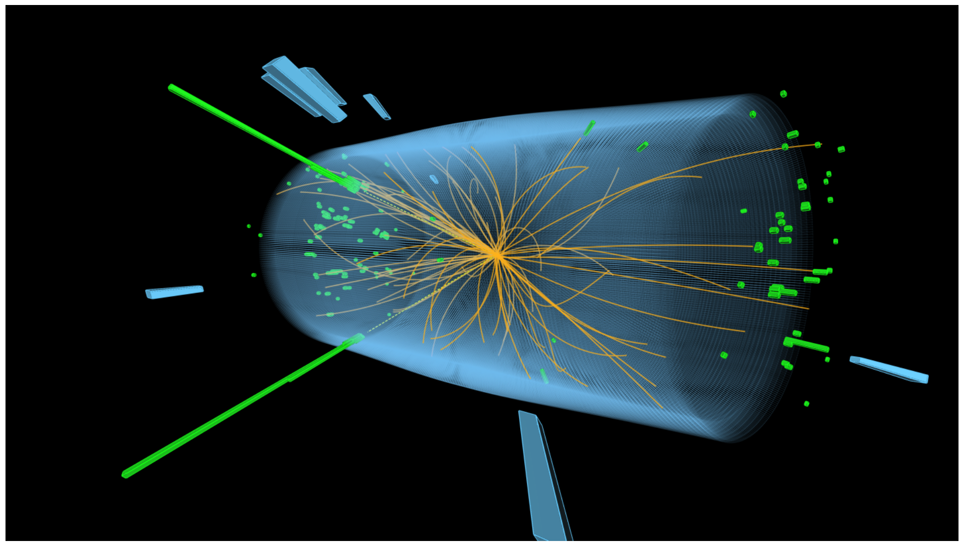

Event display from CMS from 2012

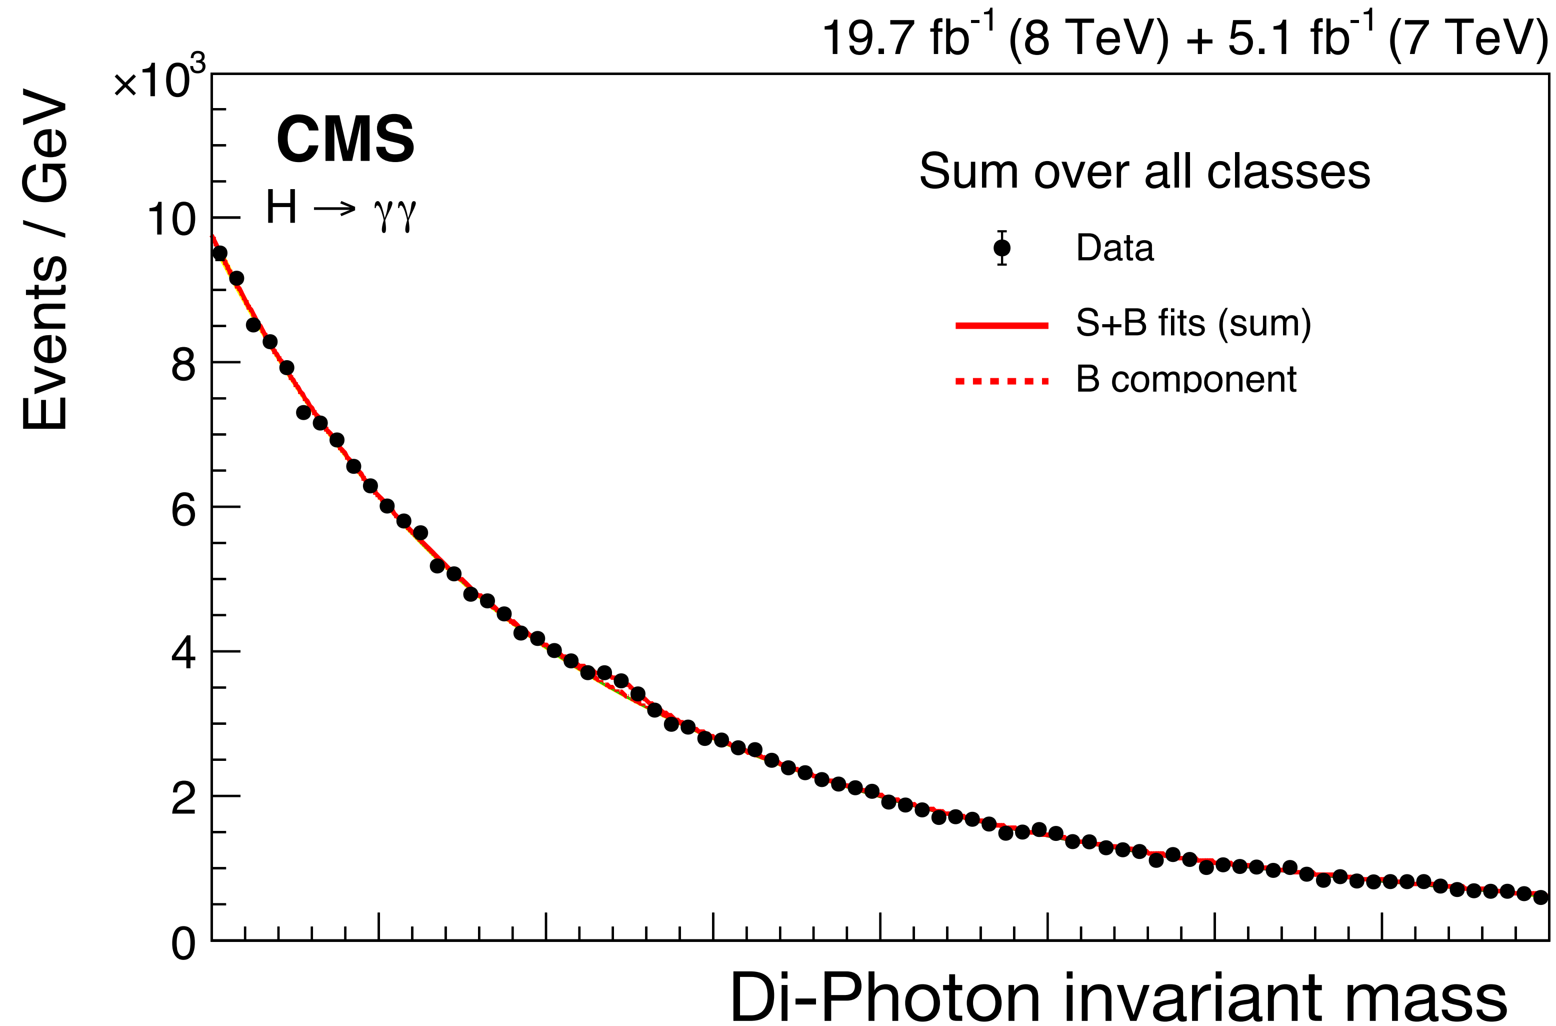

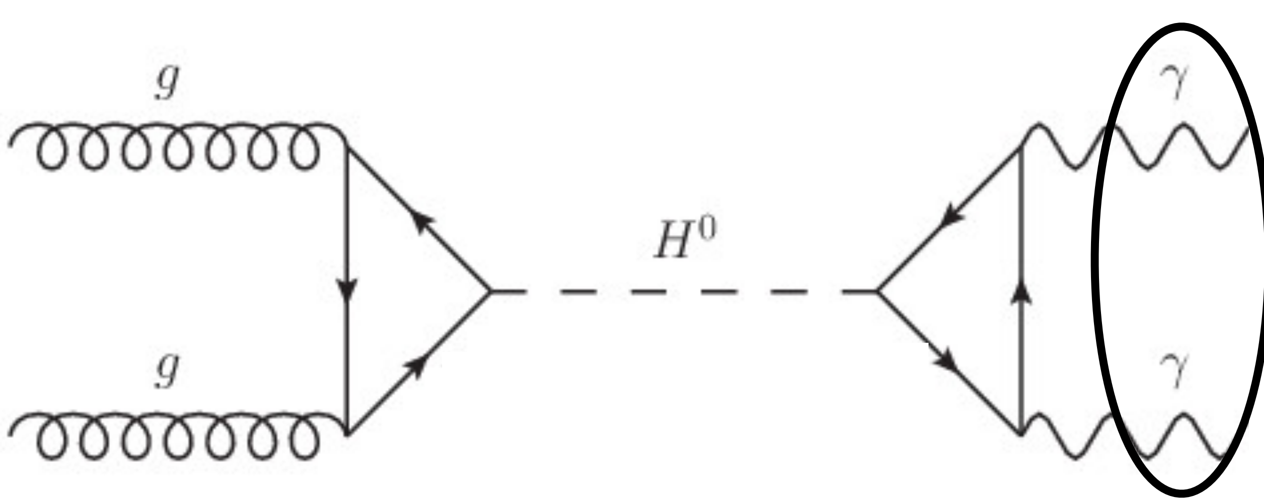

How to discover a new particle?

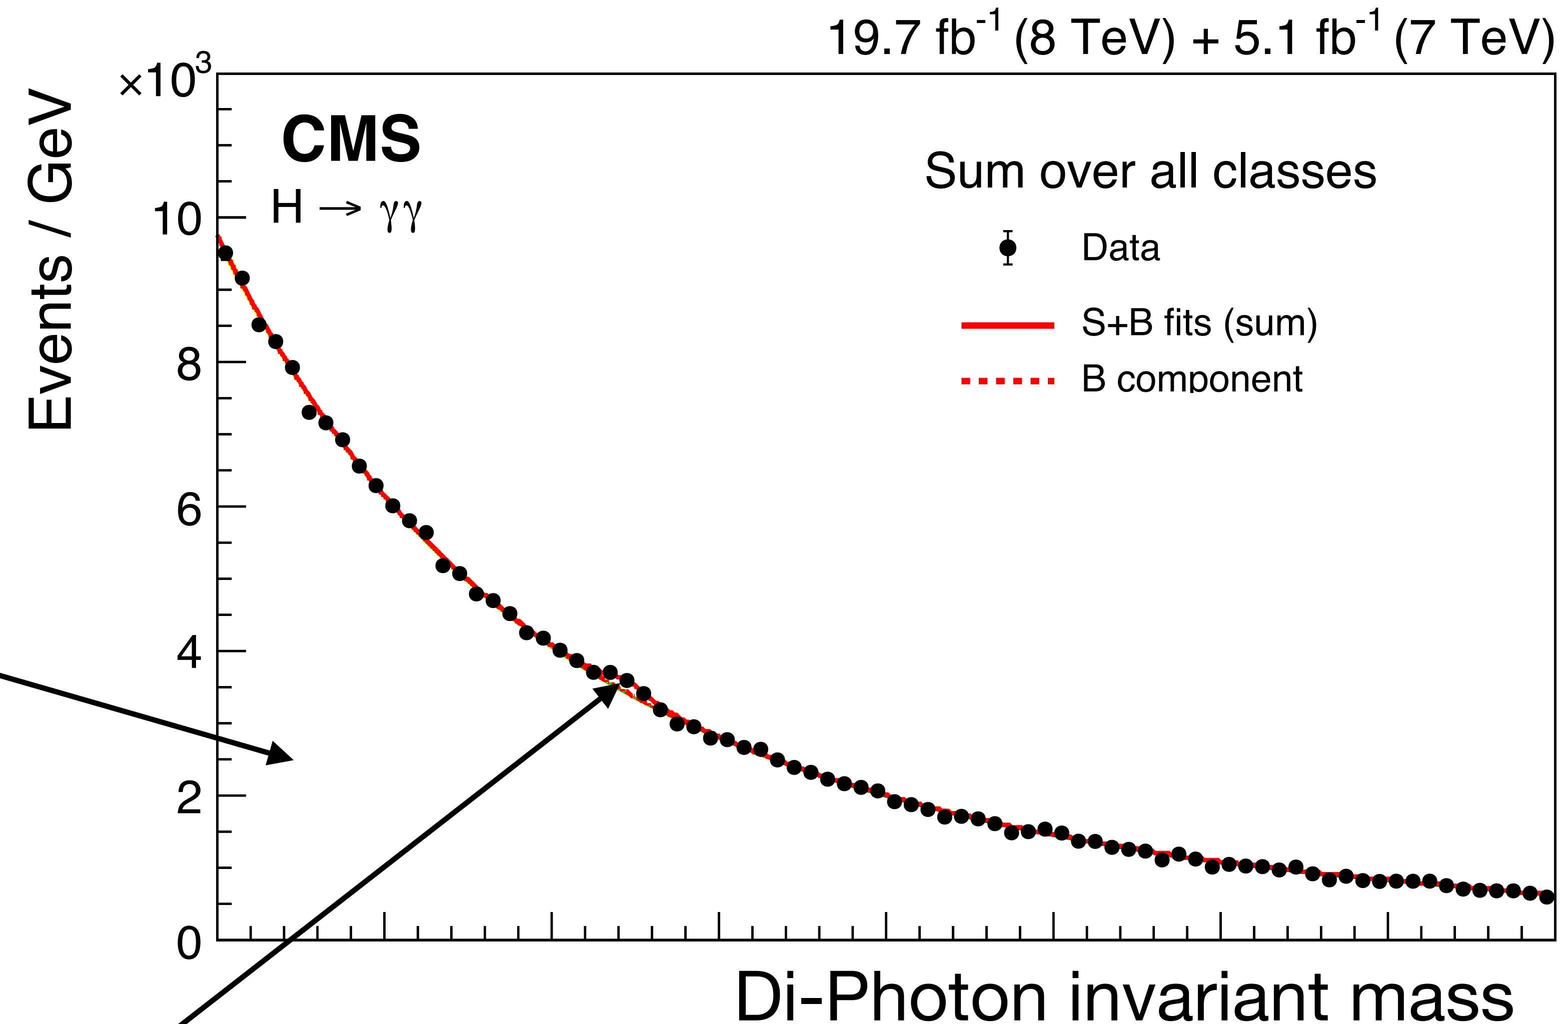

Most events are ***background***

With a ***tiny signal*** *on top*

How to discover a new particle?

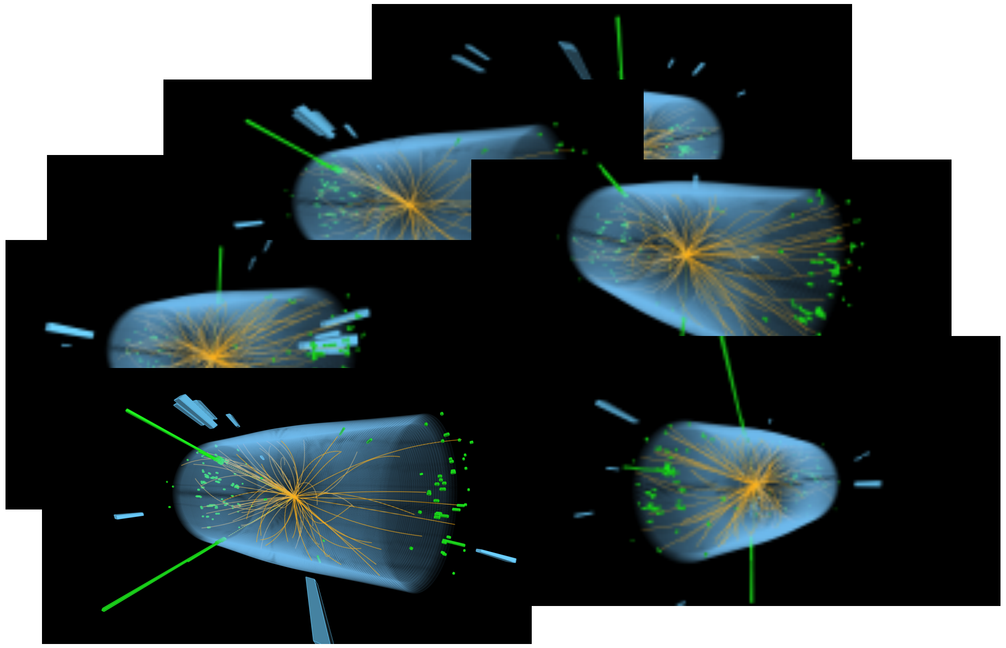

Divide into 25 categories according to purity, relative amount of signal using BDTs

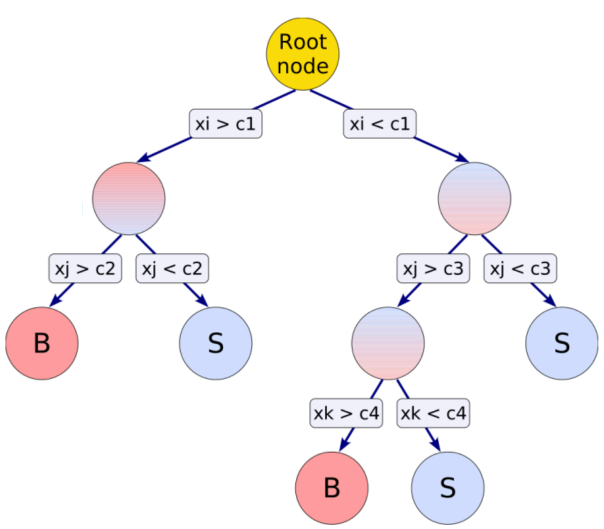

Weight events according to excepted purity of category

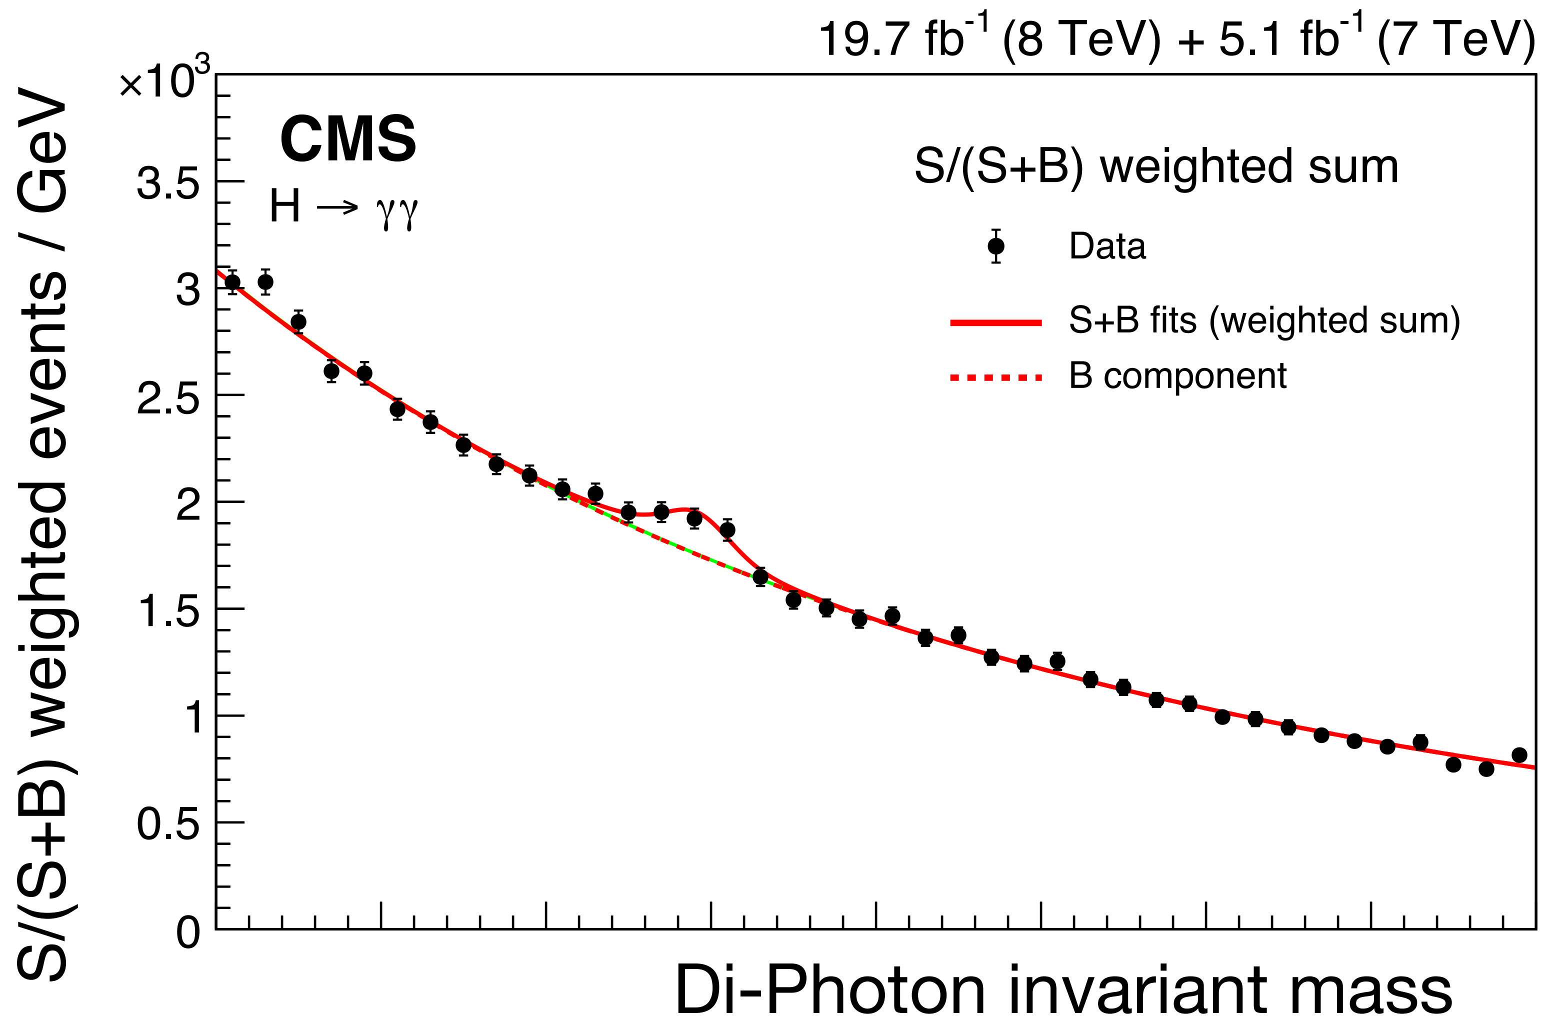

Nobel price for Higgs mechanism 2013 (Higgs and Englert)

How to discover a new particle?

Assume we know what we are looking for: e.g. a heavier version of the Z boson

We don't know the mass of the resonant particle

But we assume it decays to e.g. a pairs of top quarks

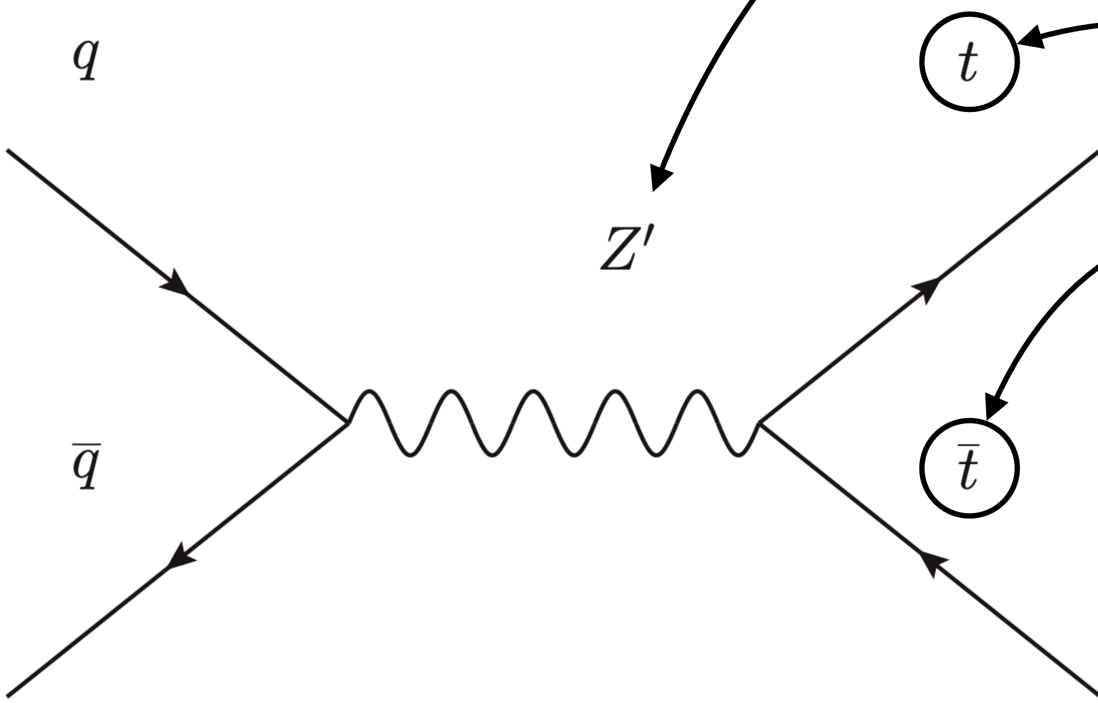

How to discover a new particle?

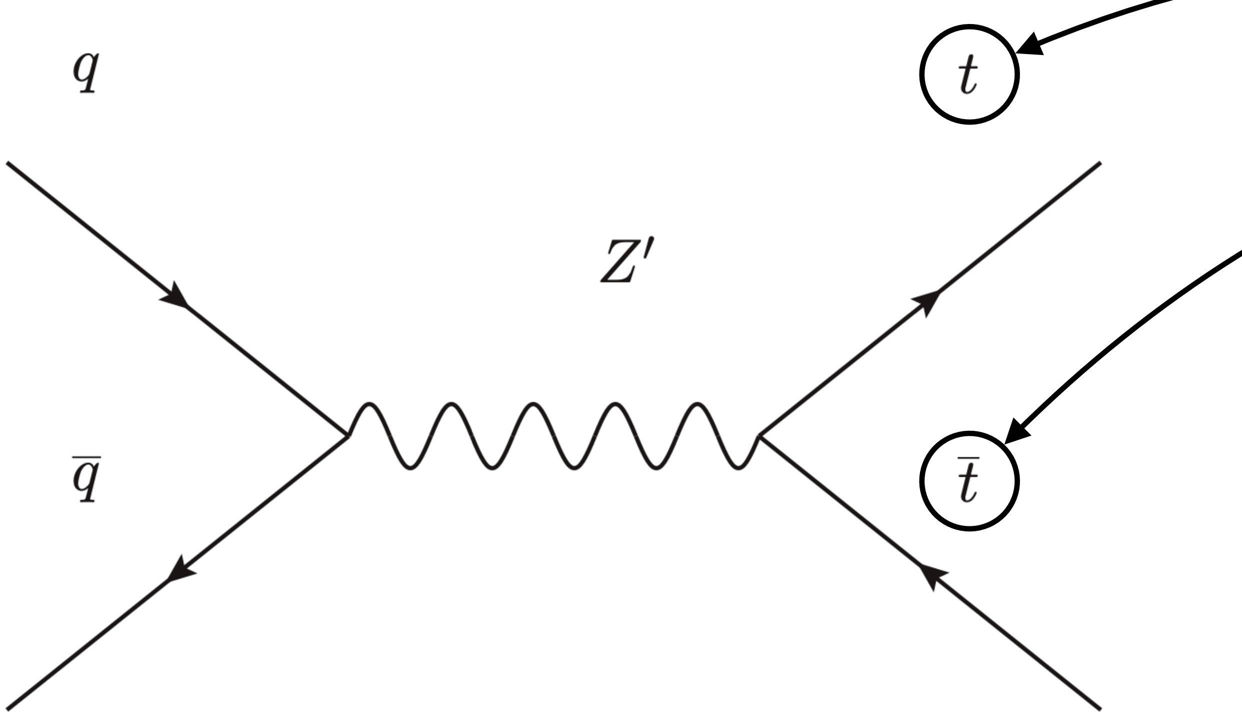

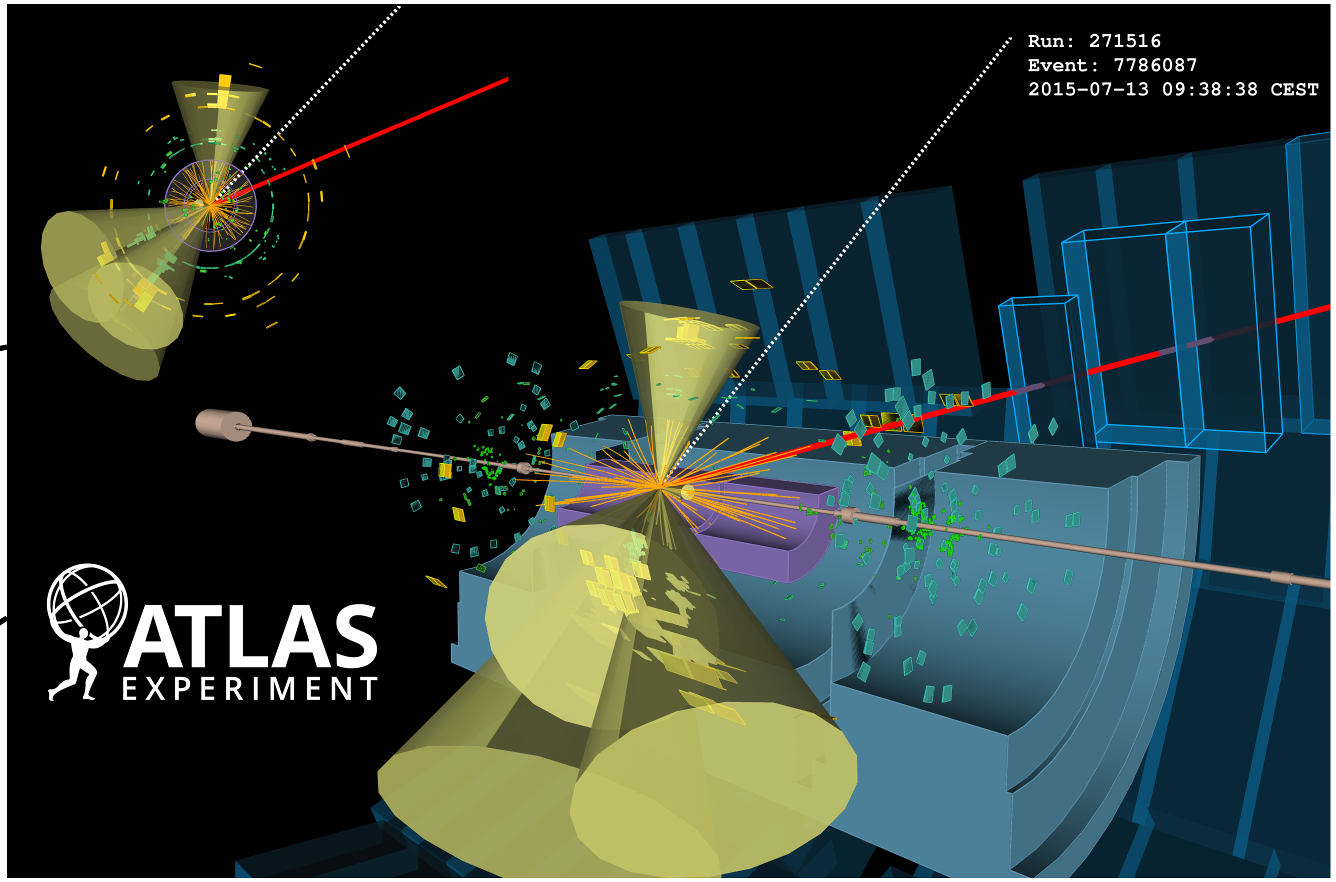

More complicated than electrons:
Use AI for top tagging

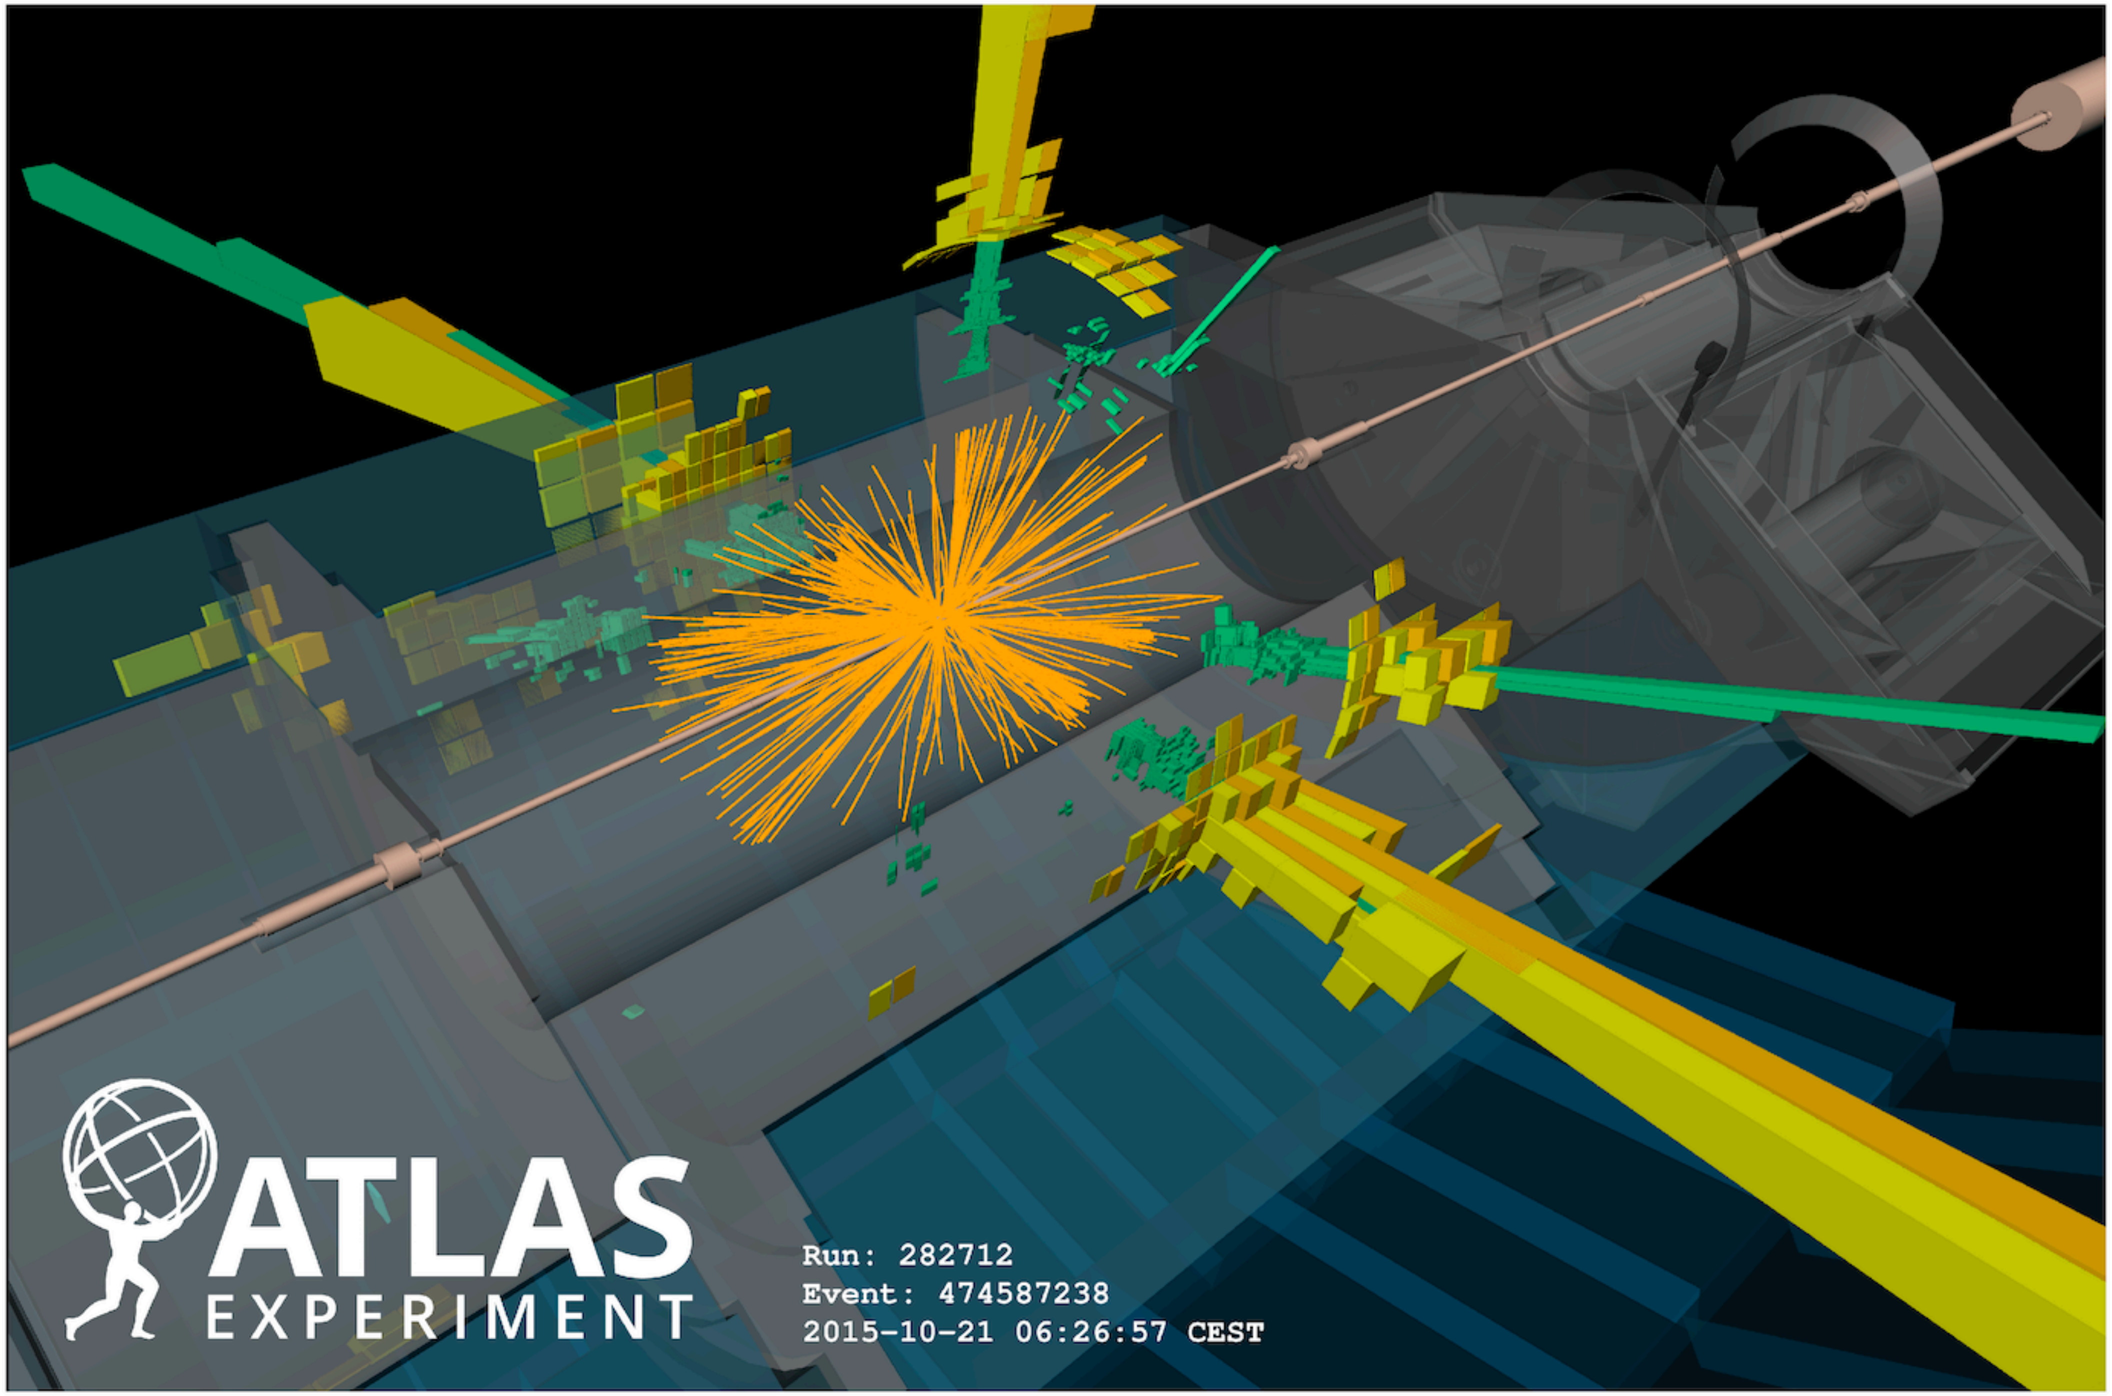

A jet is a
collimated shower of particles in the detector

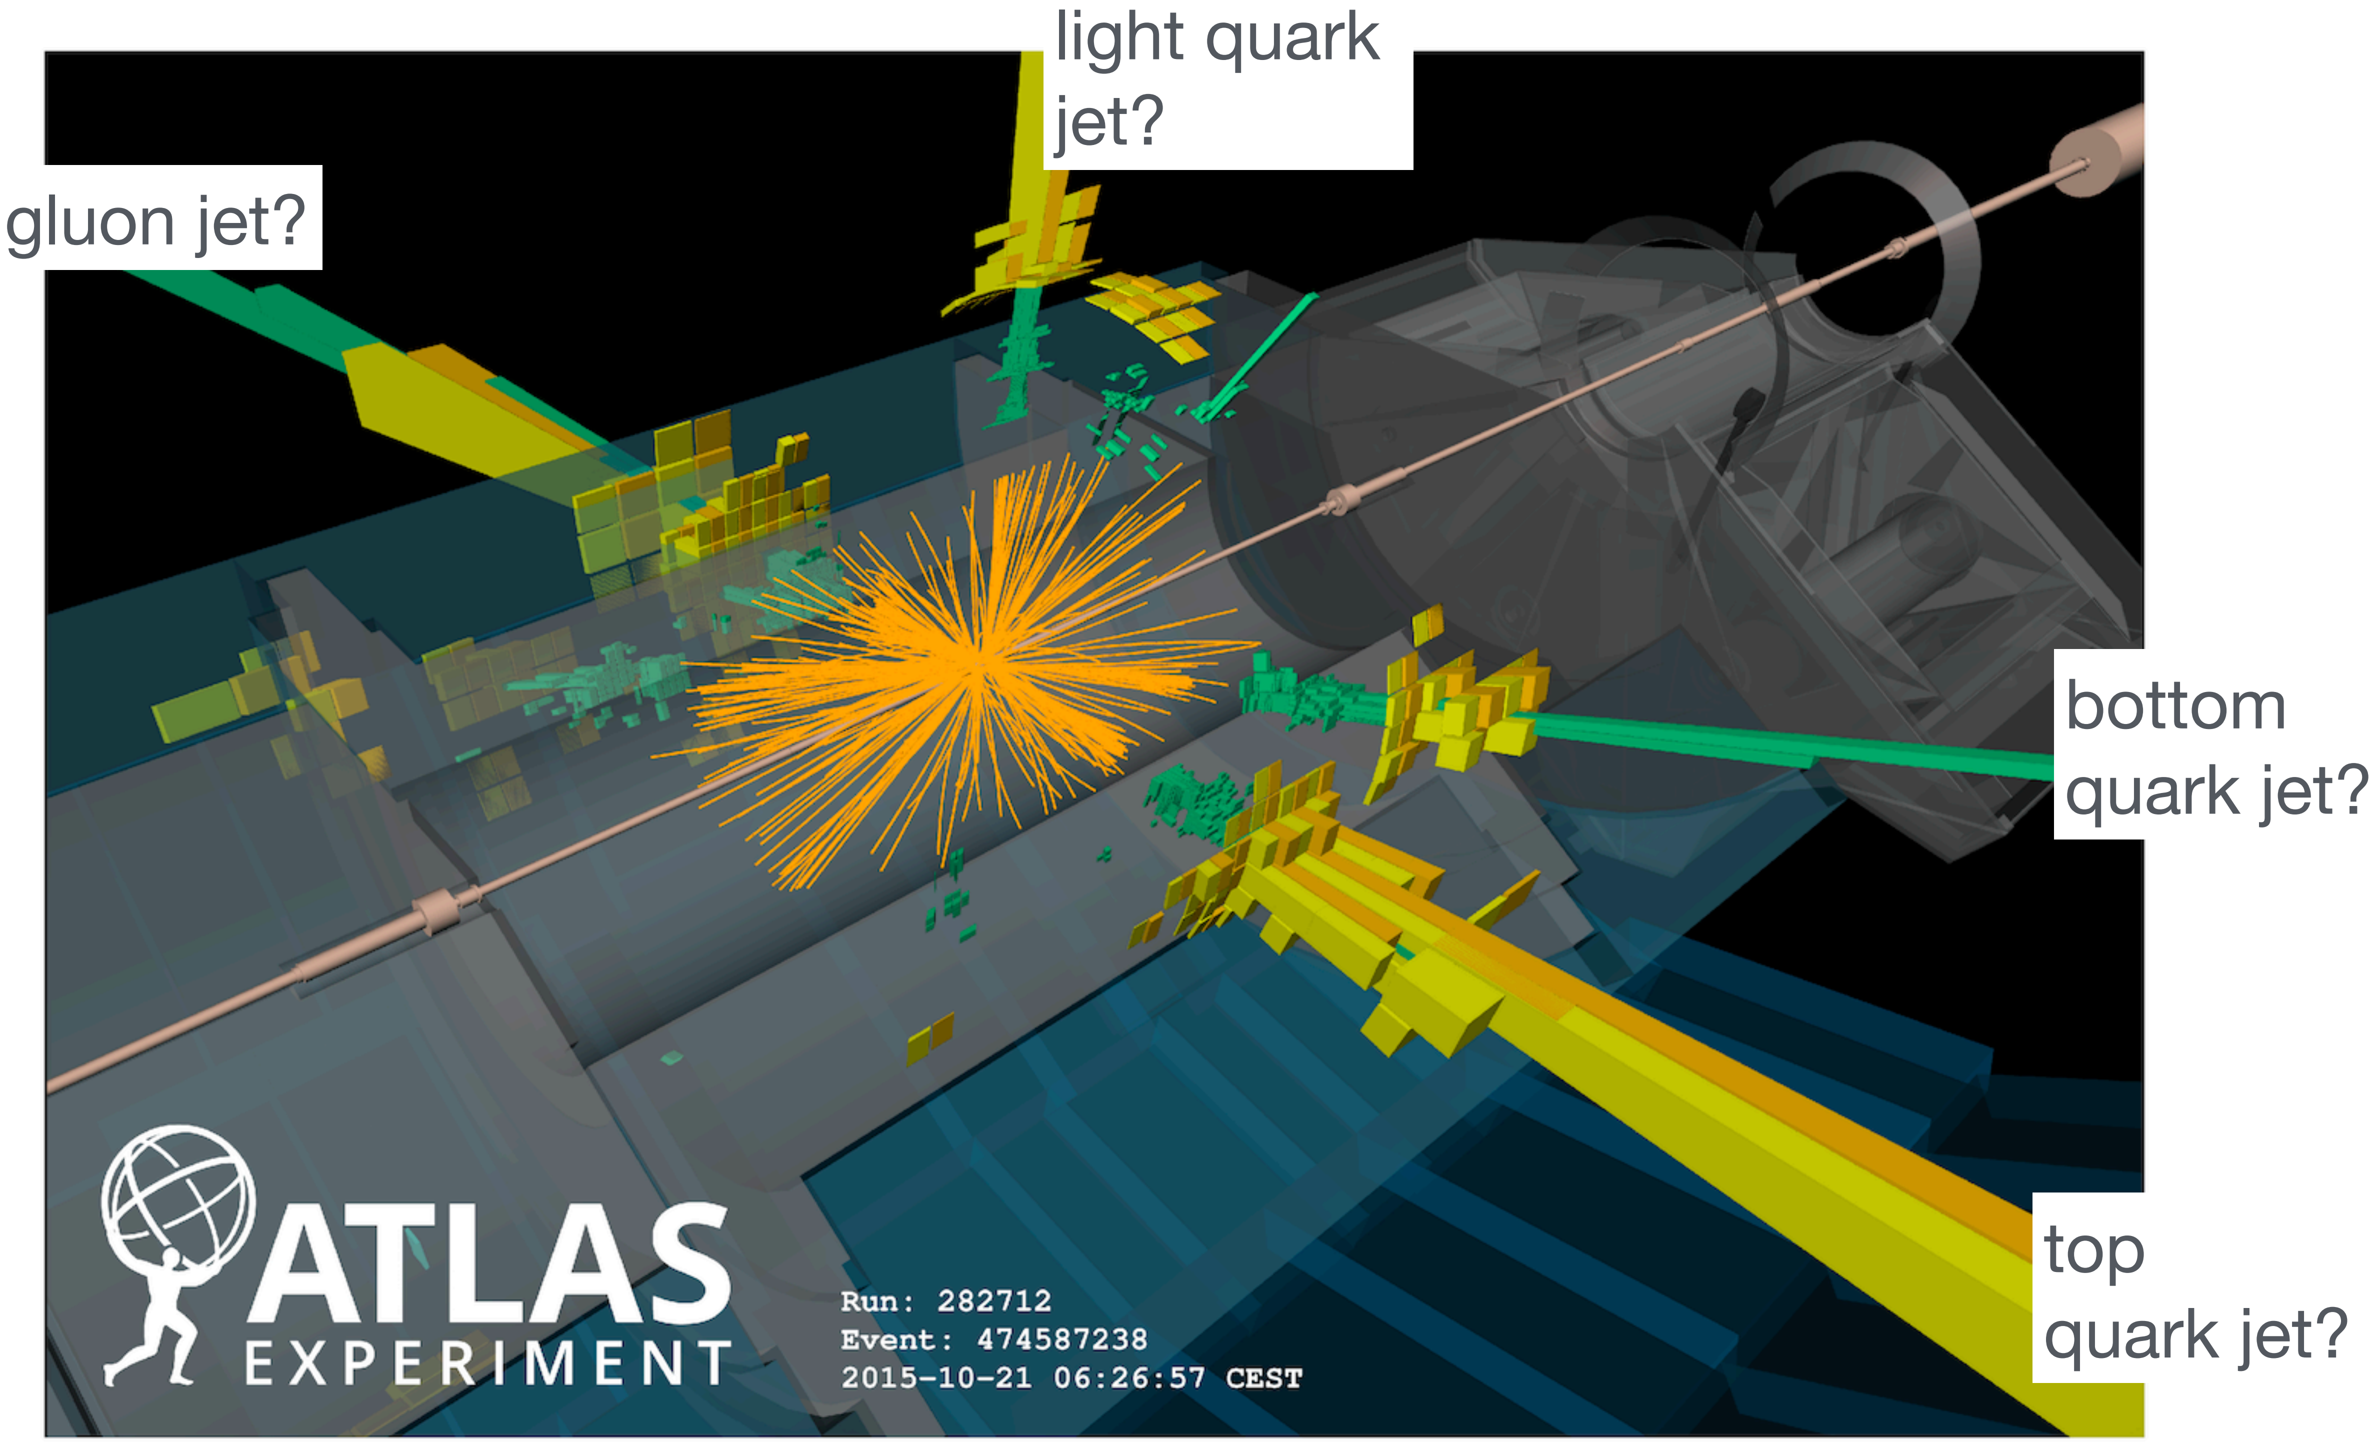

We want to know
which particle produced a jet

Top Quark Tagging

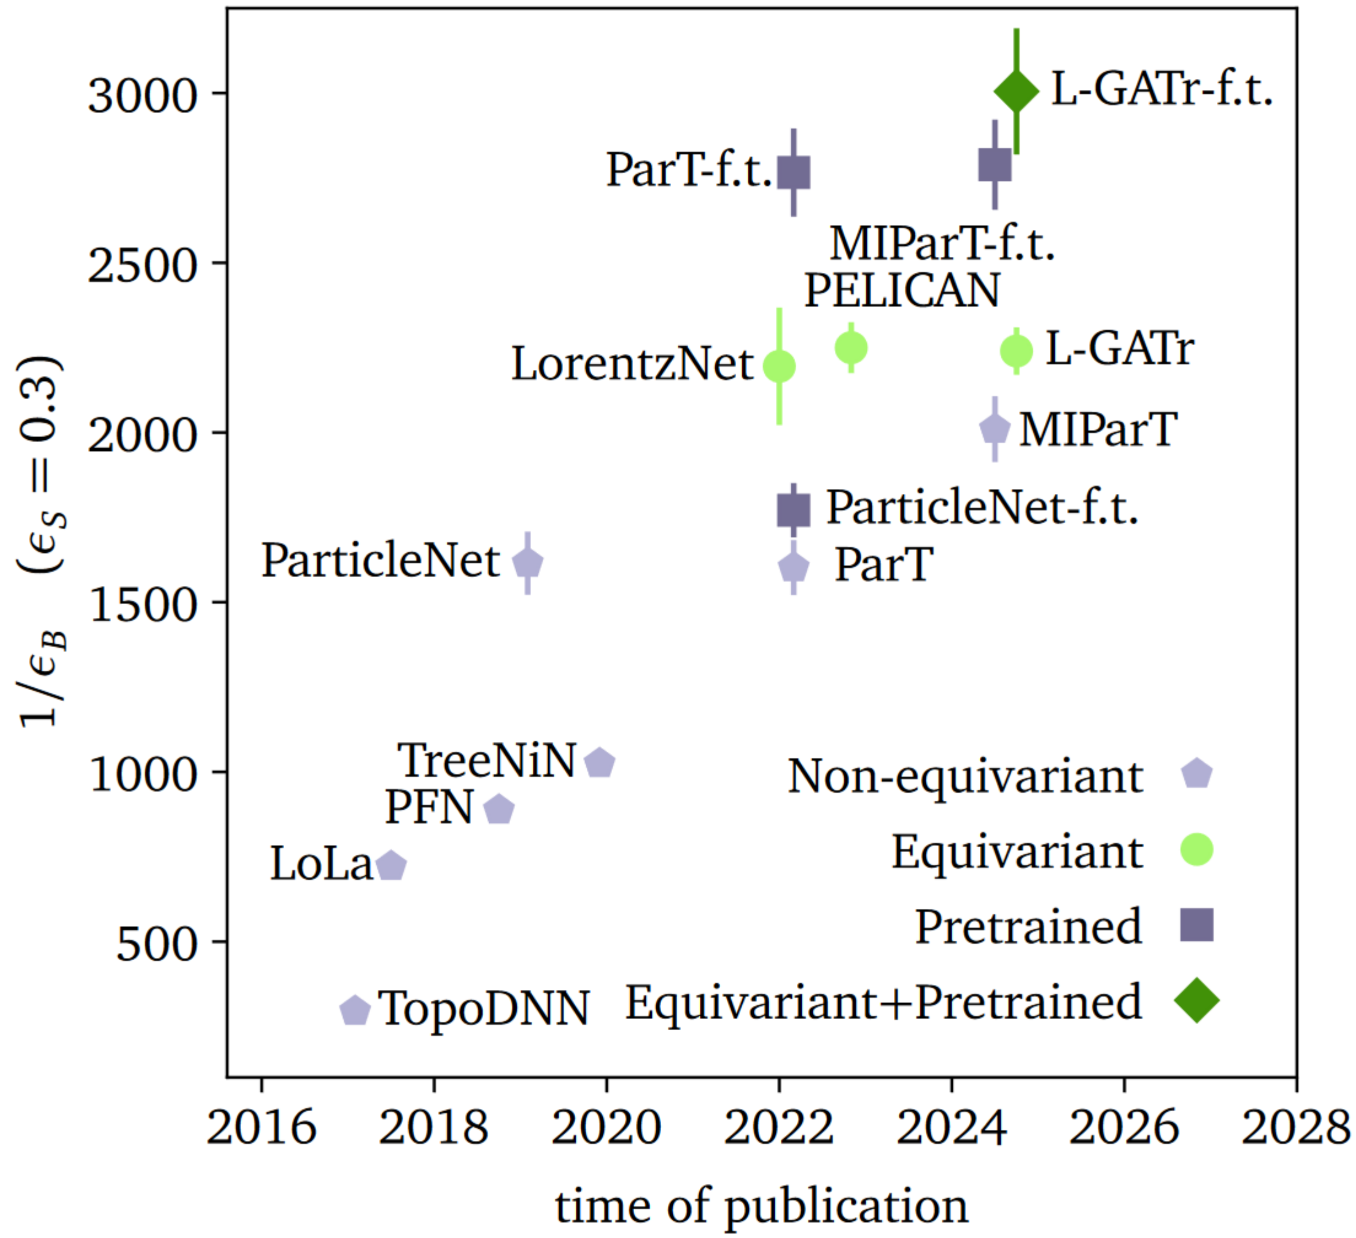

Order of magnitude improvement from ML

Modern tagging algorithms are widely used in searches for new particles

Assumptions, revisited

We don't know the mass and type of the resonant particle

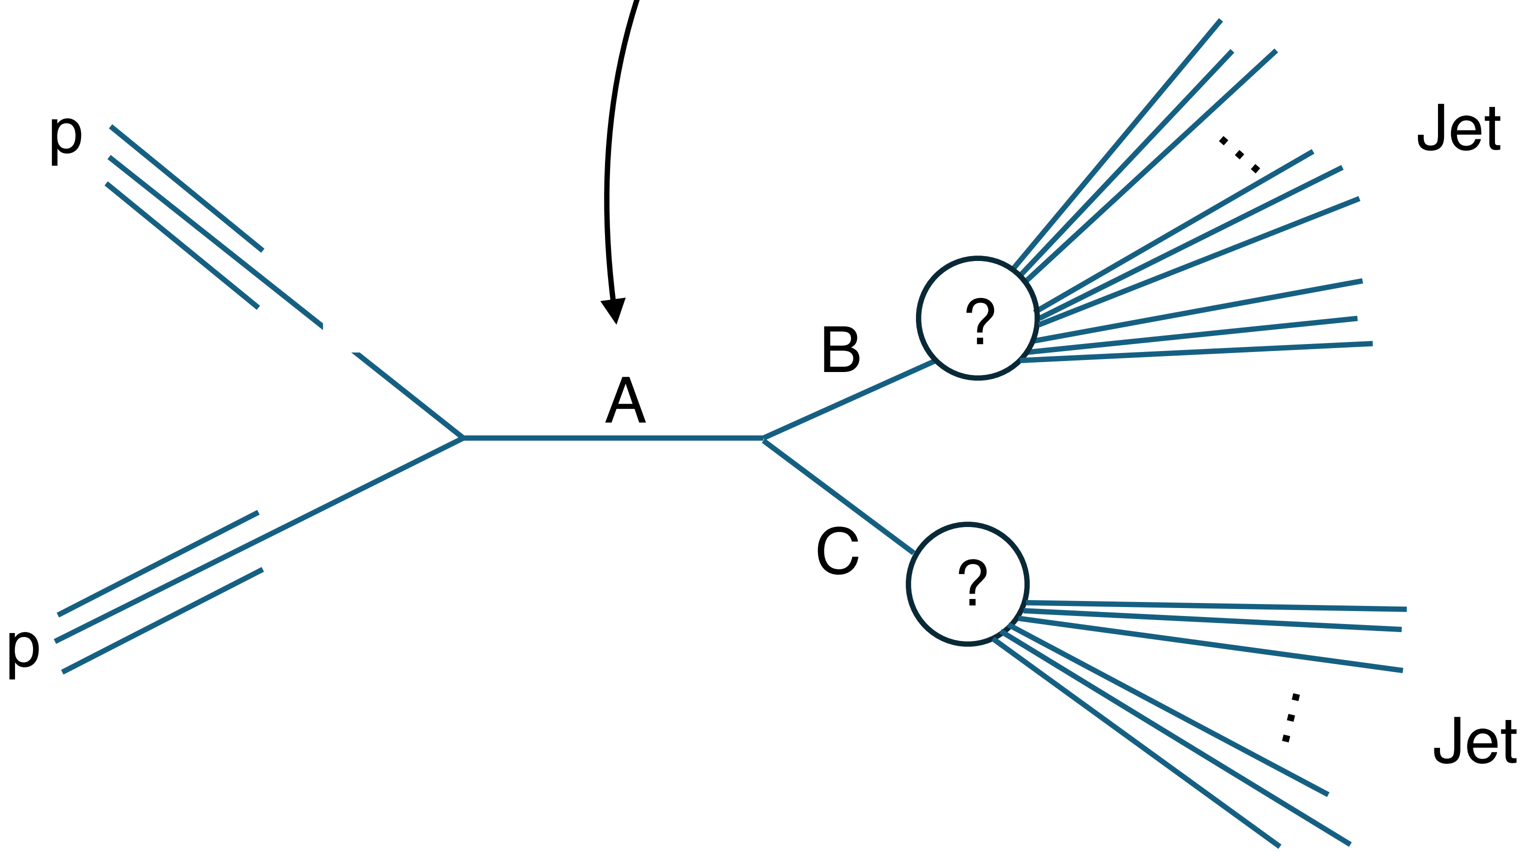

And we don't know what particles it decays to

(assume some jets)

Outline

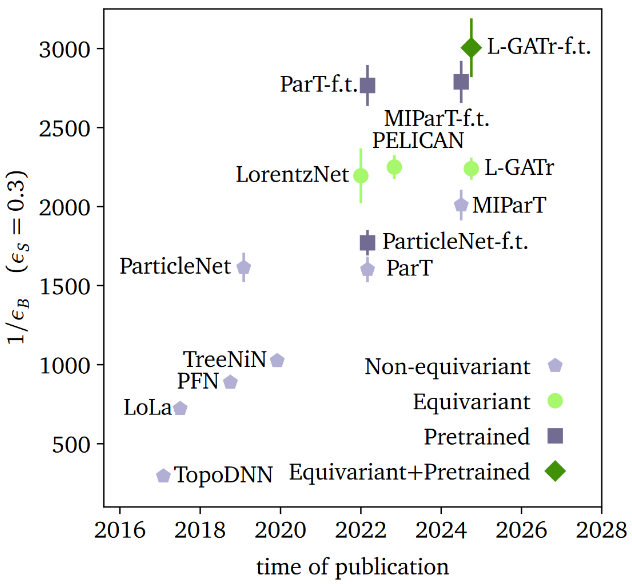

Tools for Discovery

Discovery Strategies

Autonomous Discovery

Order of magnitude (~x10) improvements in traditional search strategies with AI

Increasing autonomy of AI systems

Outline

Tools for Discovery

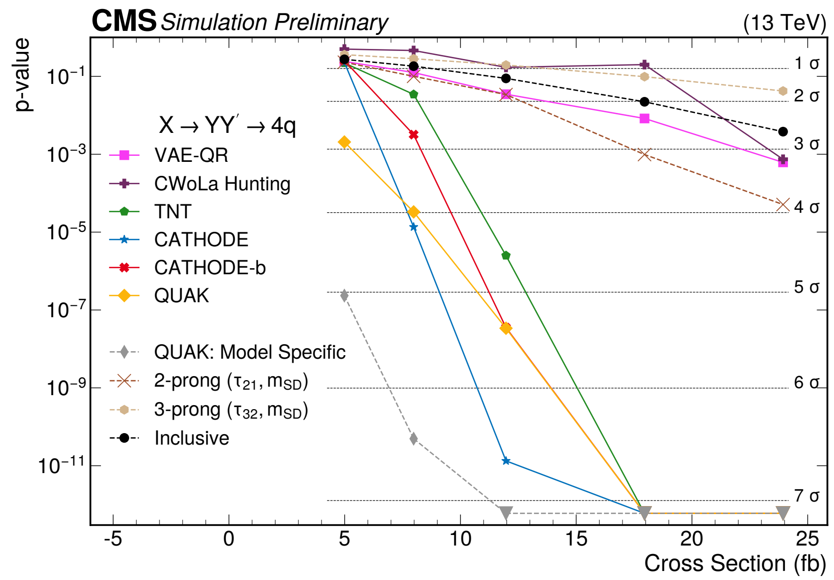

Discovery Strategies

Autonomous Discovery

Increasing autonomy of AI systems

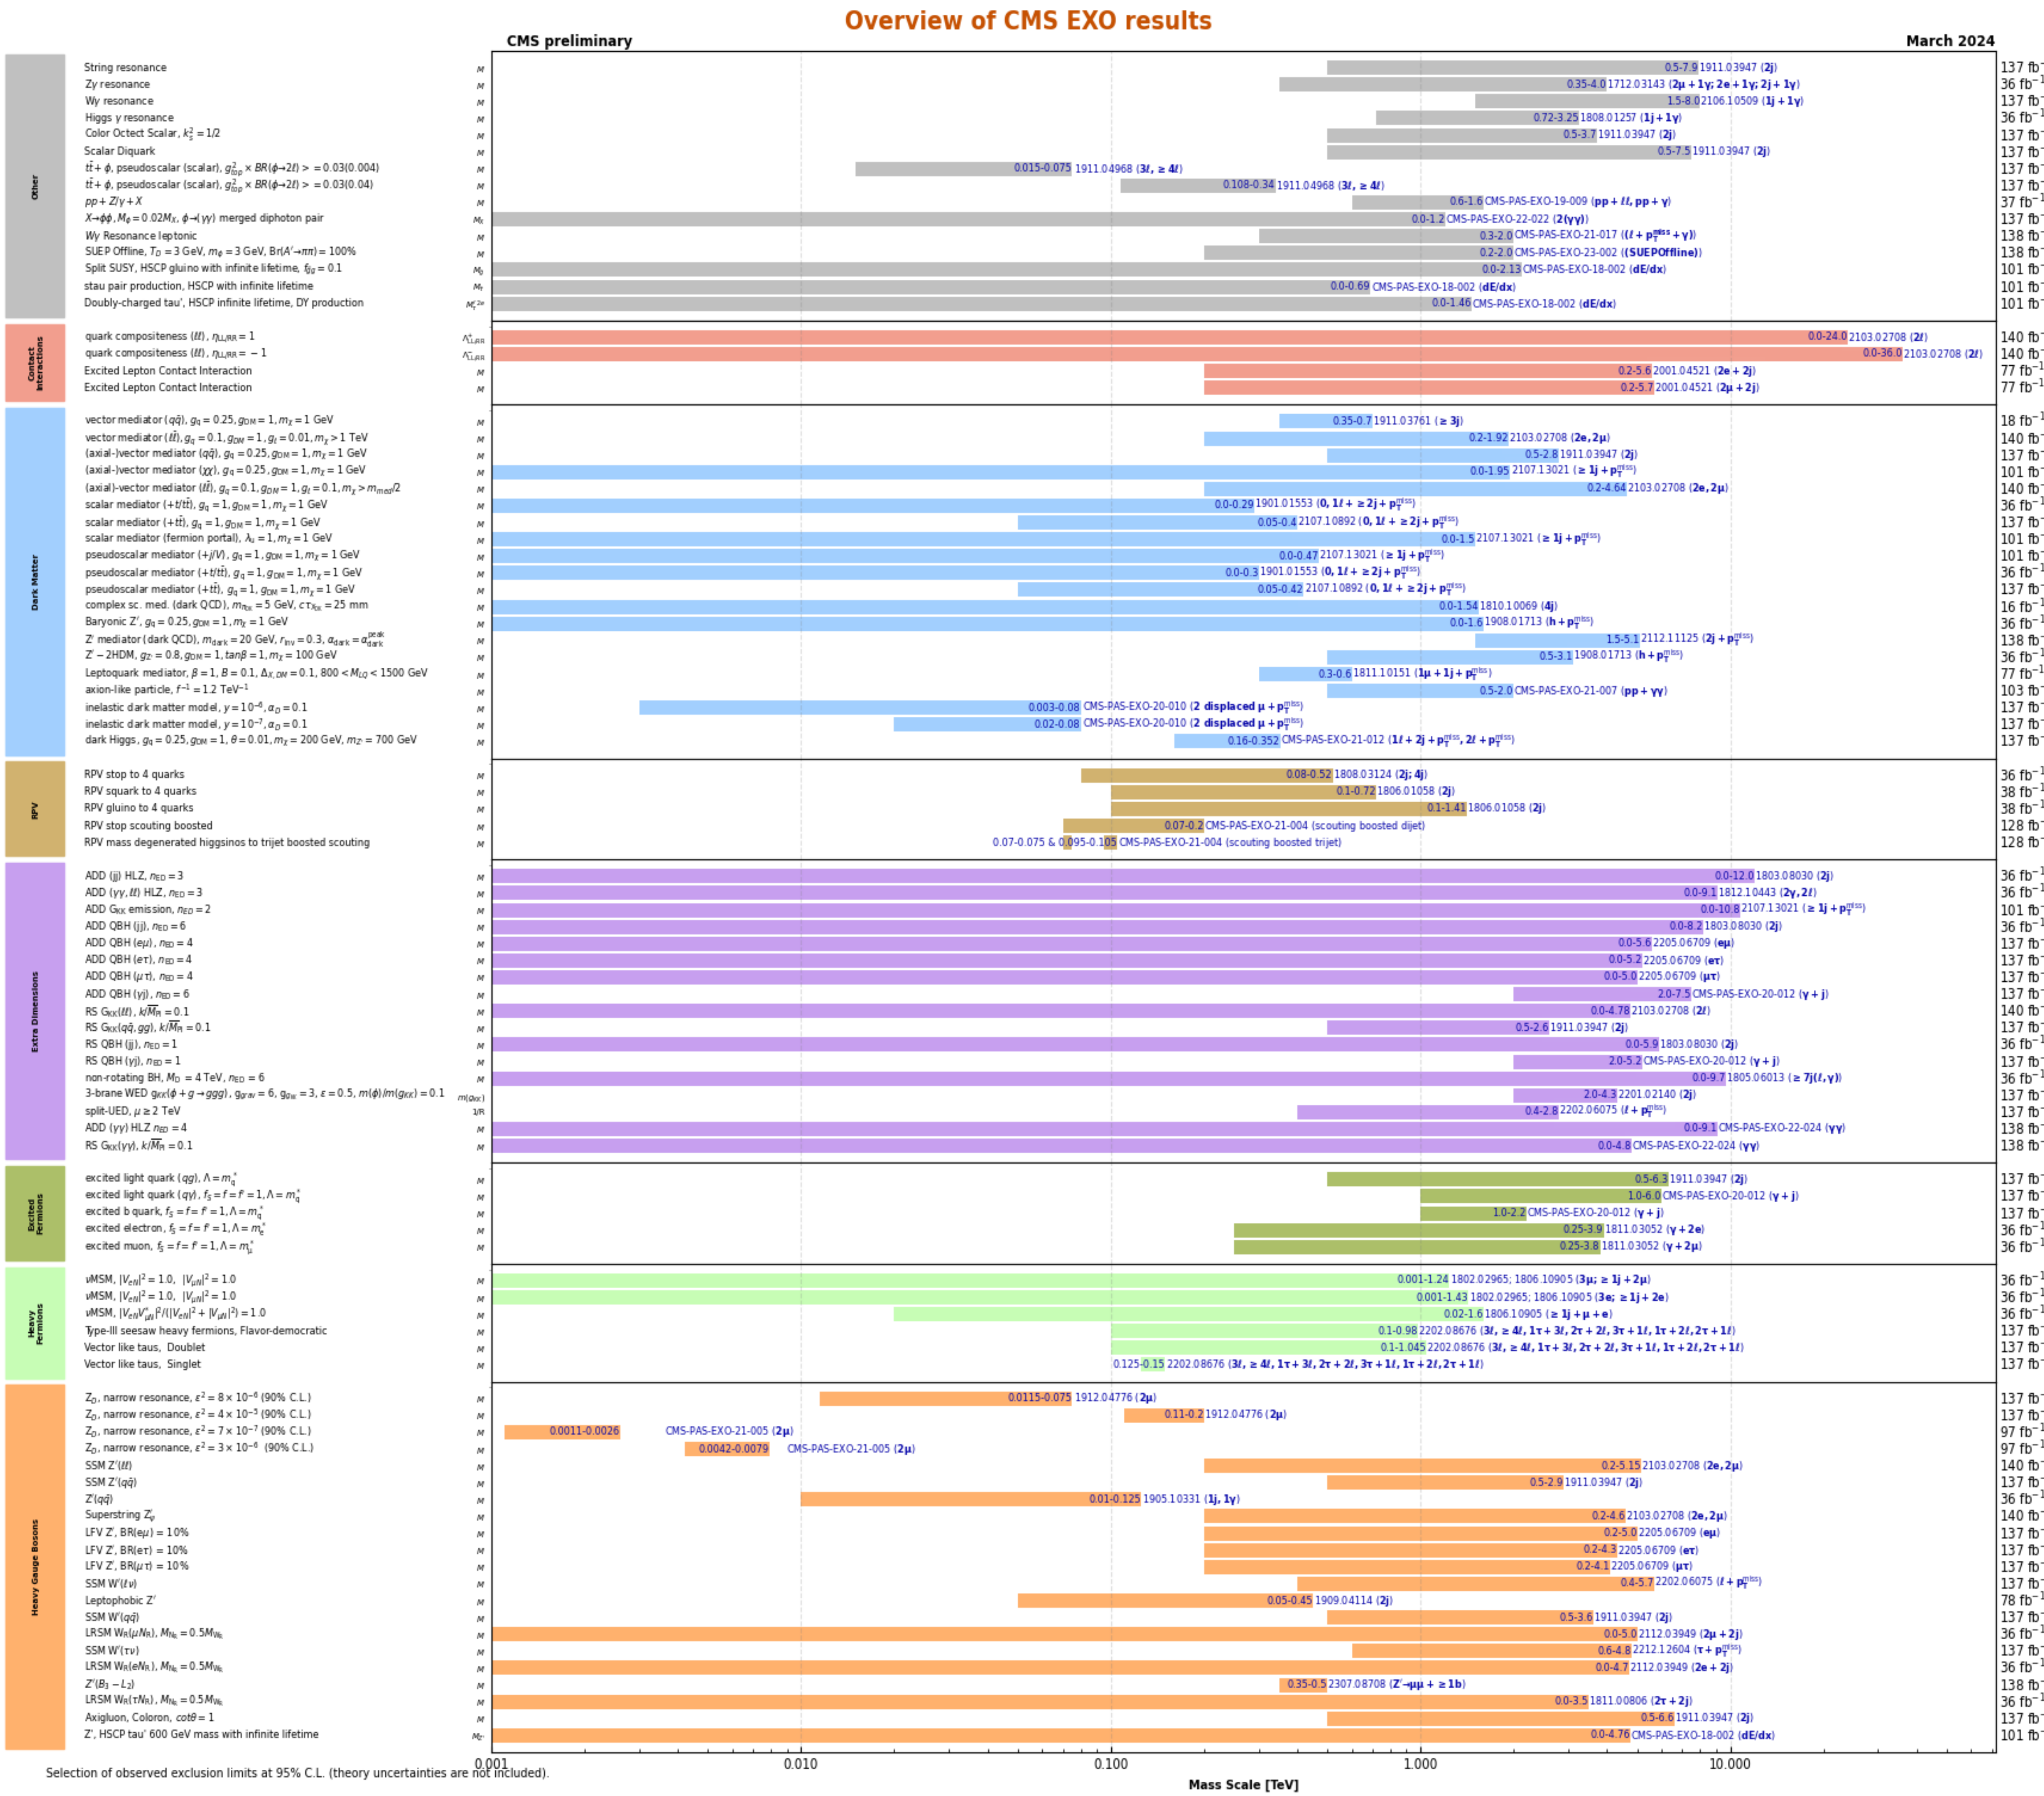

So far, no new physics in model-driven searches

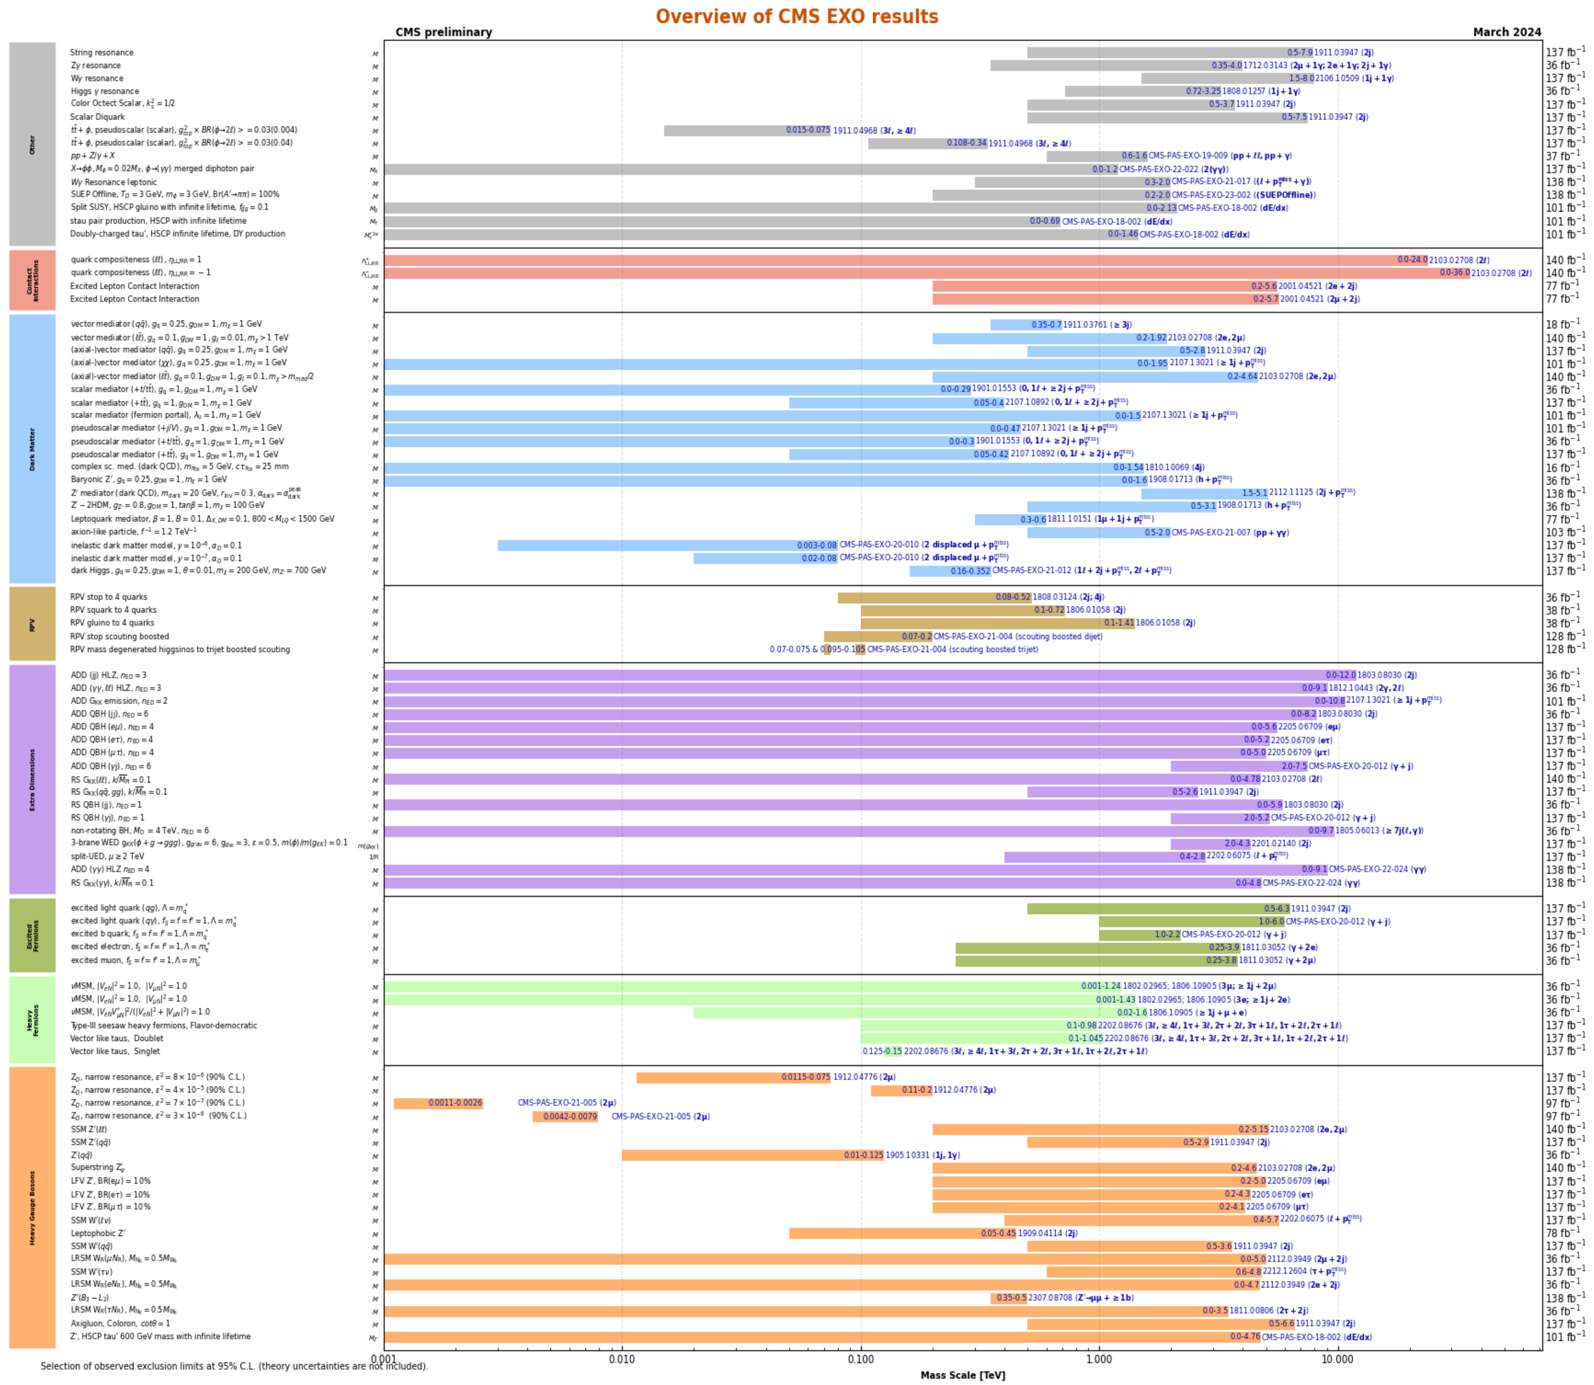

No new physics exists at LHC scales

We just need more data

We are testing the wrong models

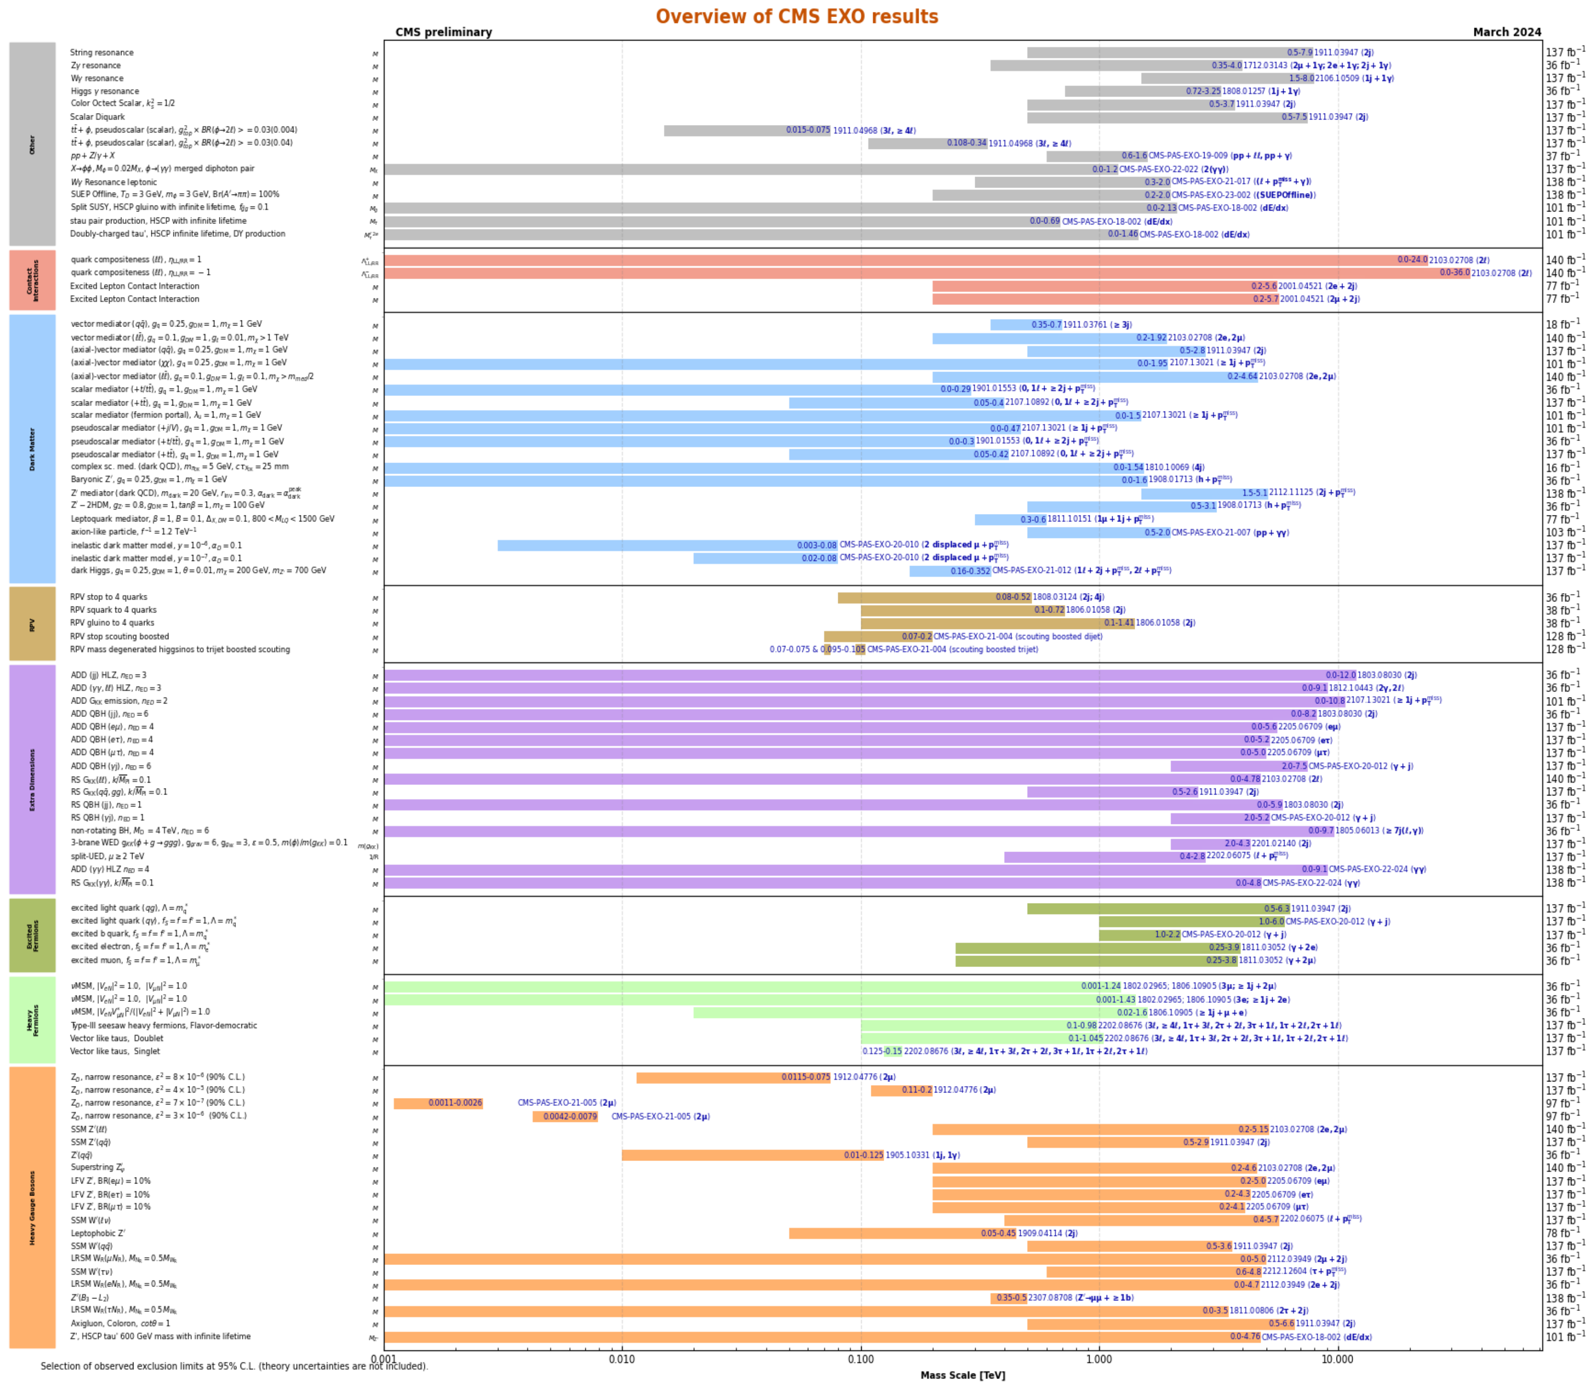

No new physics exists at LHC scales

We just need more data

We are testing the wrong models

Let's look for anomalies

Types of anomalies

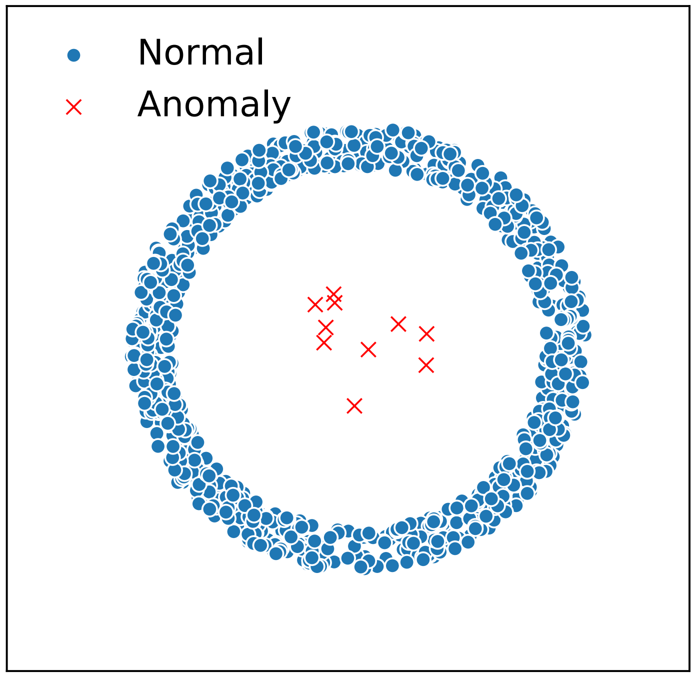

Outliers

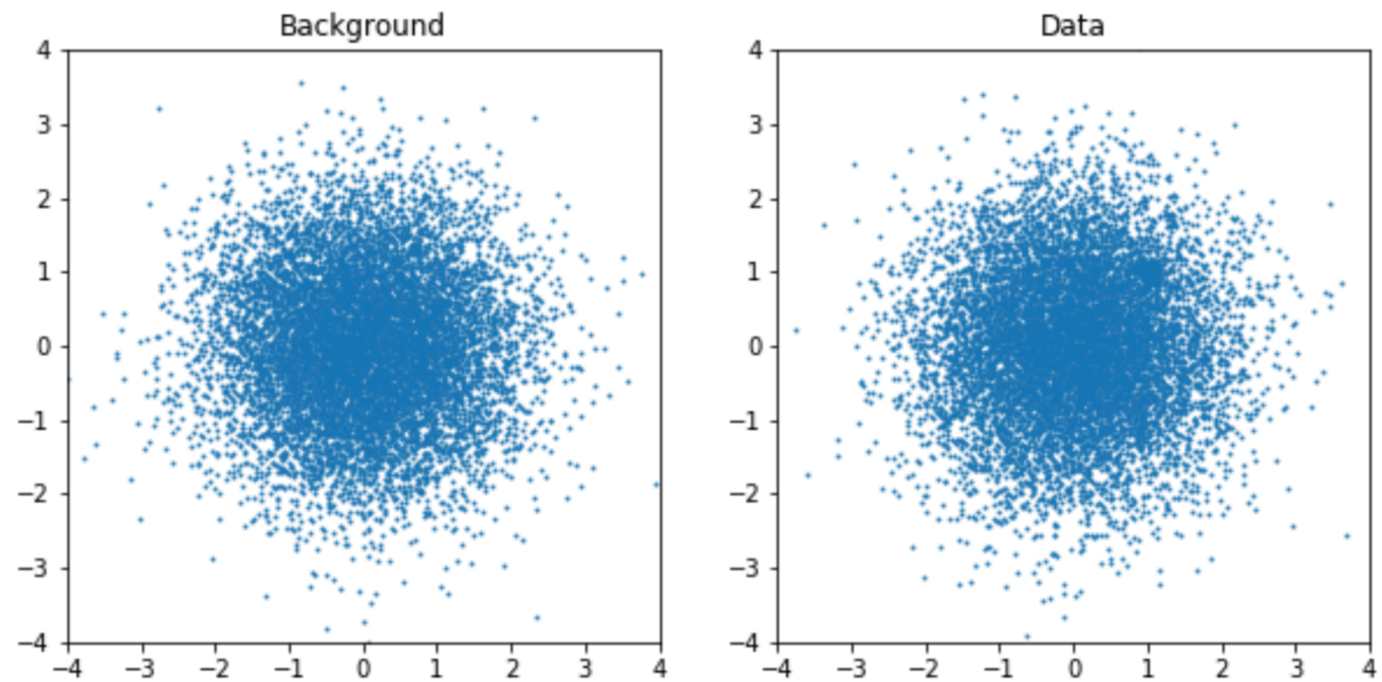

Overdensities

Image from D. Shih
at ACAT25

Types of anomalies

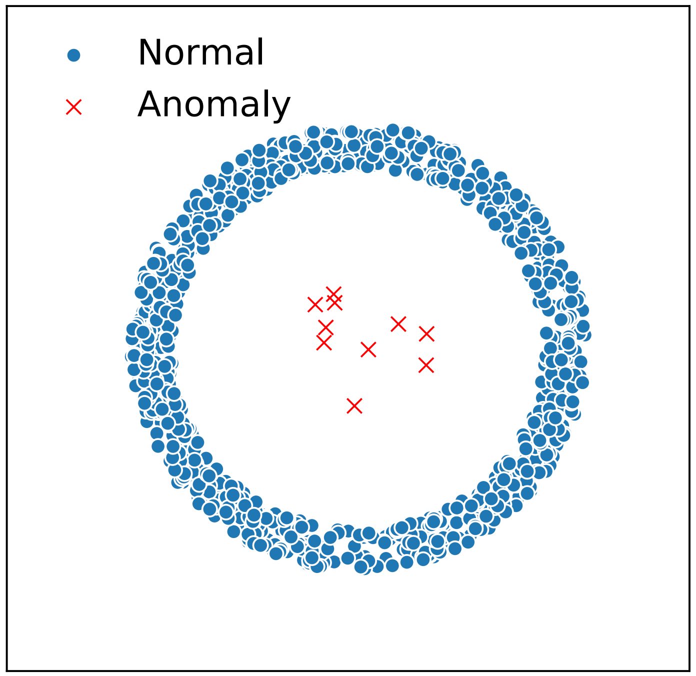

Outliers

Overdensities

Identify anomaly as events in region of low p(background)

Types of anomalies

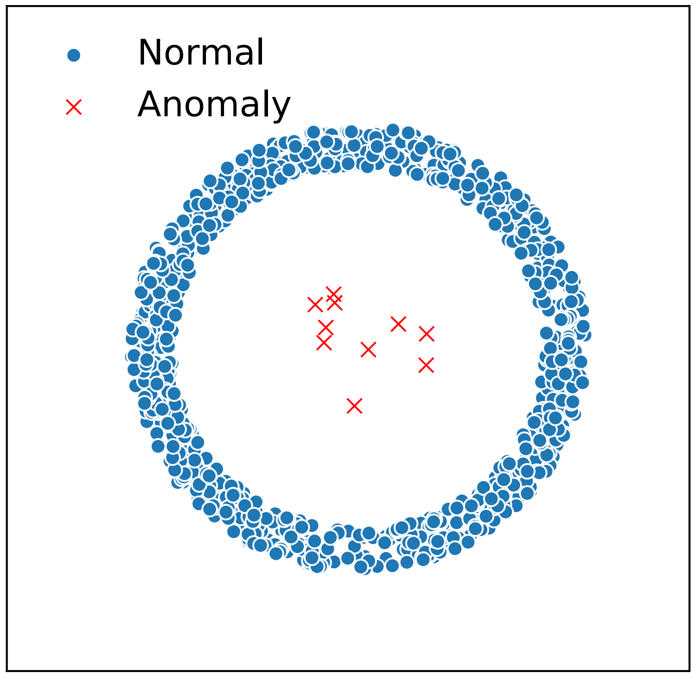

Outliers

Identify anomaly as events in region of low p(background)

Can identify with autoencoders*

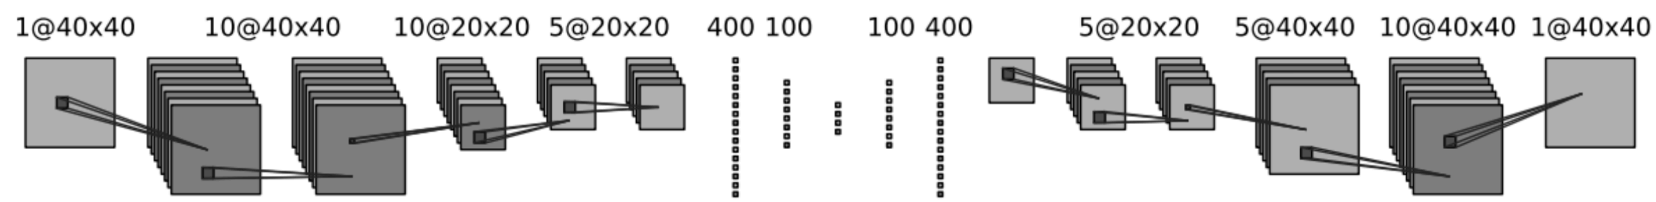

$$L = \frac{1}{N}\sum_{i=1}^{N}(x_i^{in} - x_i^{out})^2$$

Also works during data taking

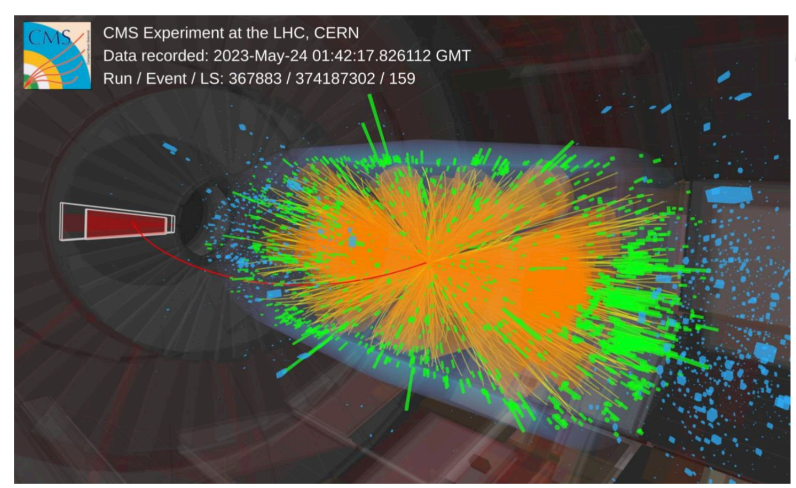

AXOL1TL

hls4ml

Farina et al 1808.08992; Heimel, **GK**, et al 1808.08979, CMS-DP-2024-059; See also Bortolato, Kamenik et al 2103.06595 &

*or extensions like NAE (2206.14225)

Types of anomalies

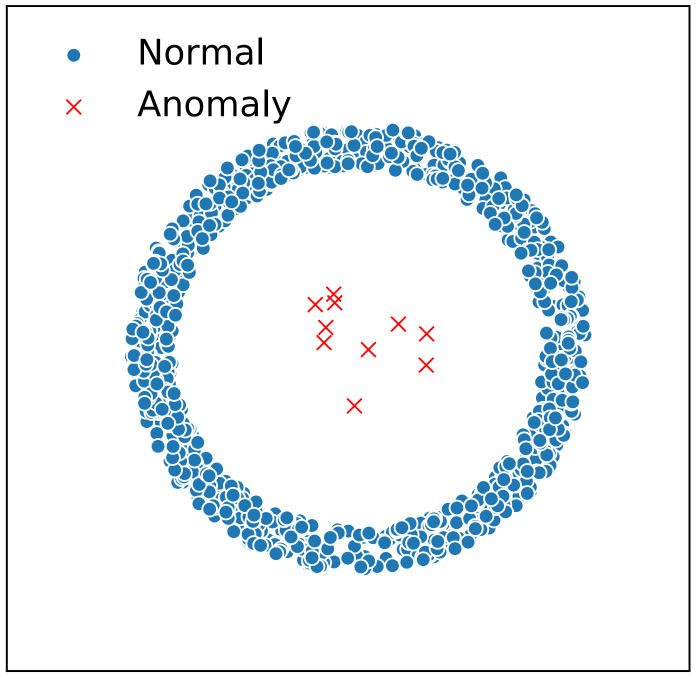

Outliers

Identify anomaly as events in region of low p(background)

No optimiality guarantees

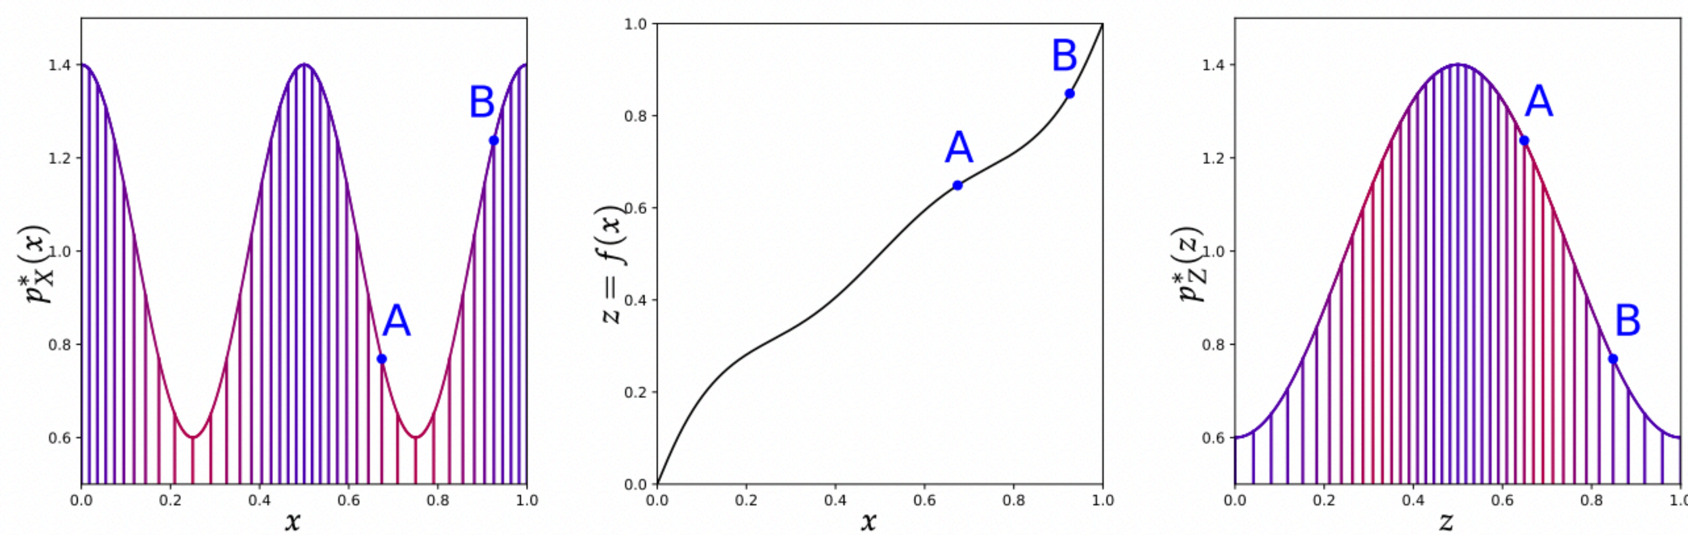

Lan, Dinh 2012.03808; **GK** et al 2209.06225

Types of anomalies

Outliers

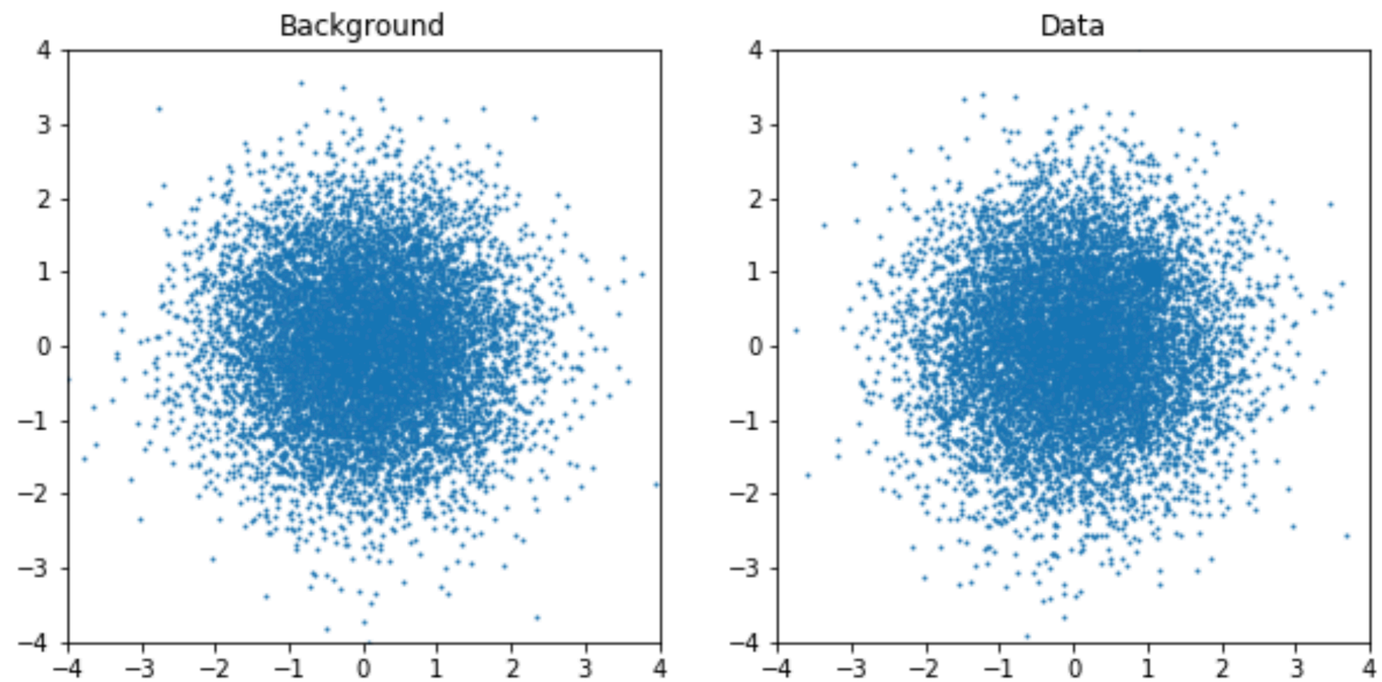

Overdensities

Identify anomaly as difference between p(background) and p(data)

Image from D. Shih at ACAT25

Types of anomalies

Outliers

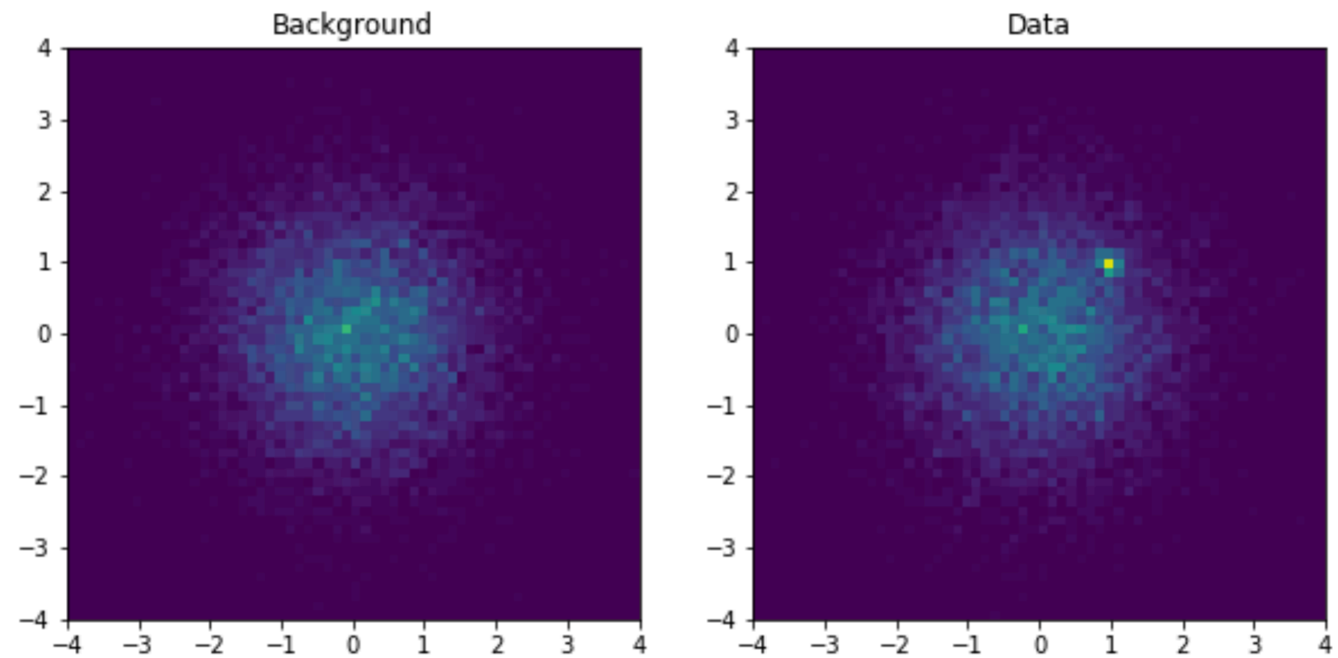

Overdensities

Identify anomaly as
difference between
p(background) and p(data)

Image from D. Shih
at ACAT25


Generation frame: 1024 x 1024

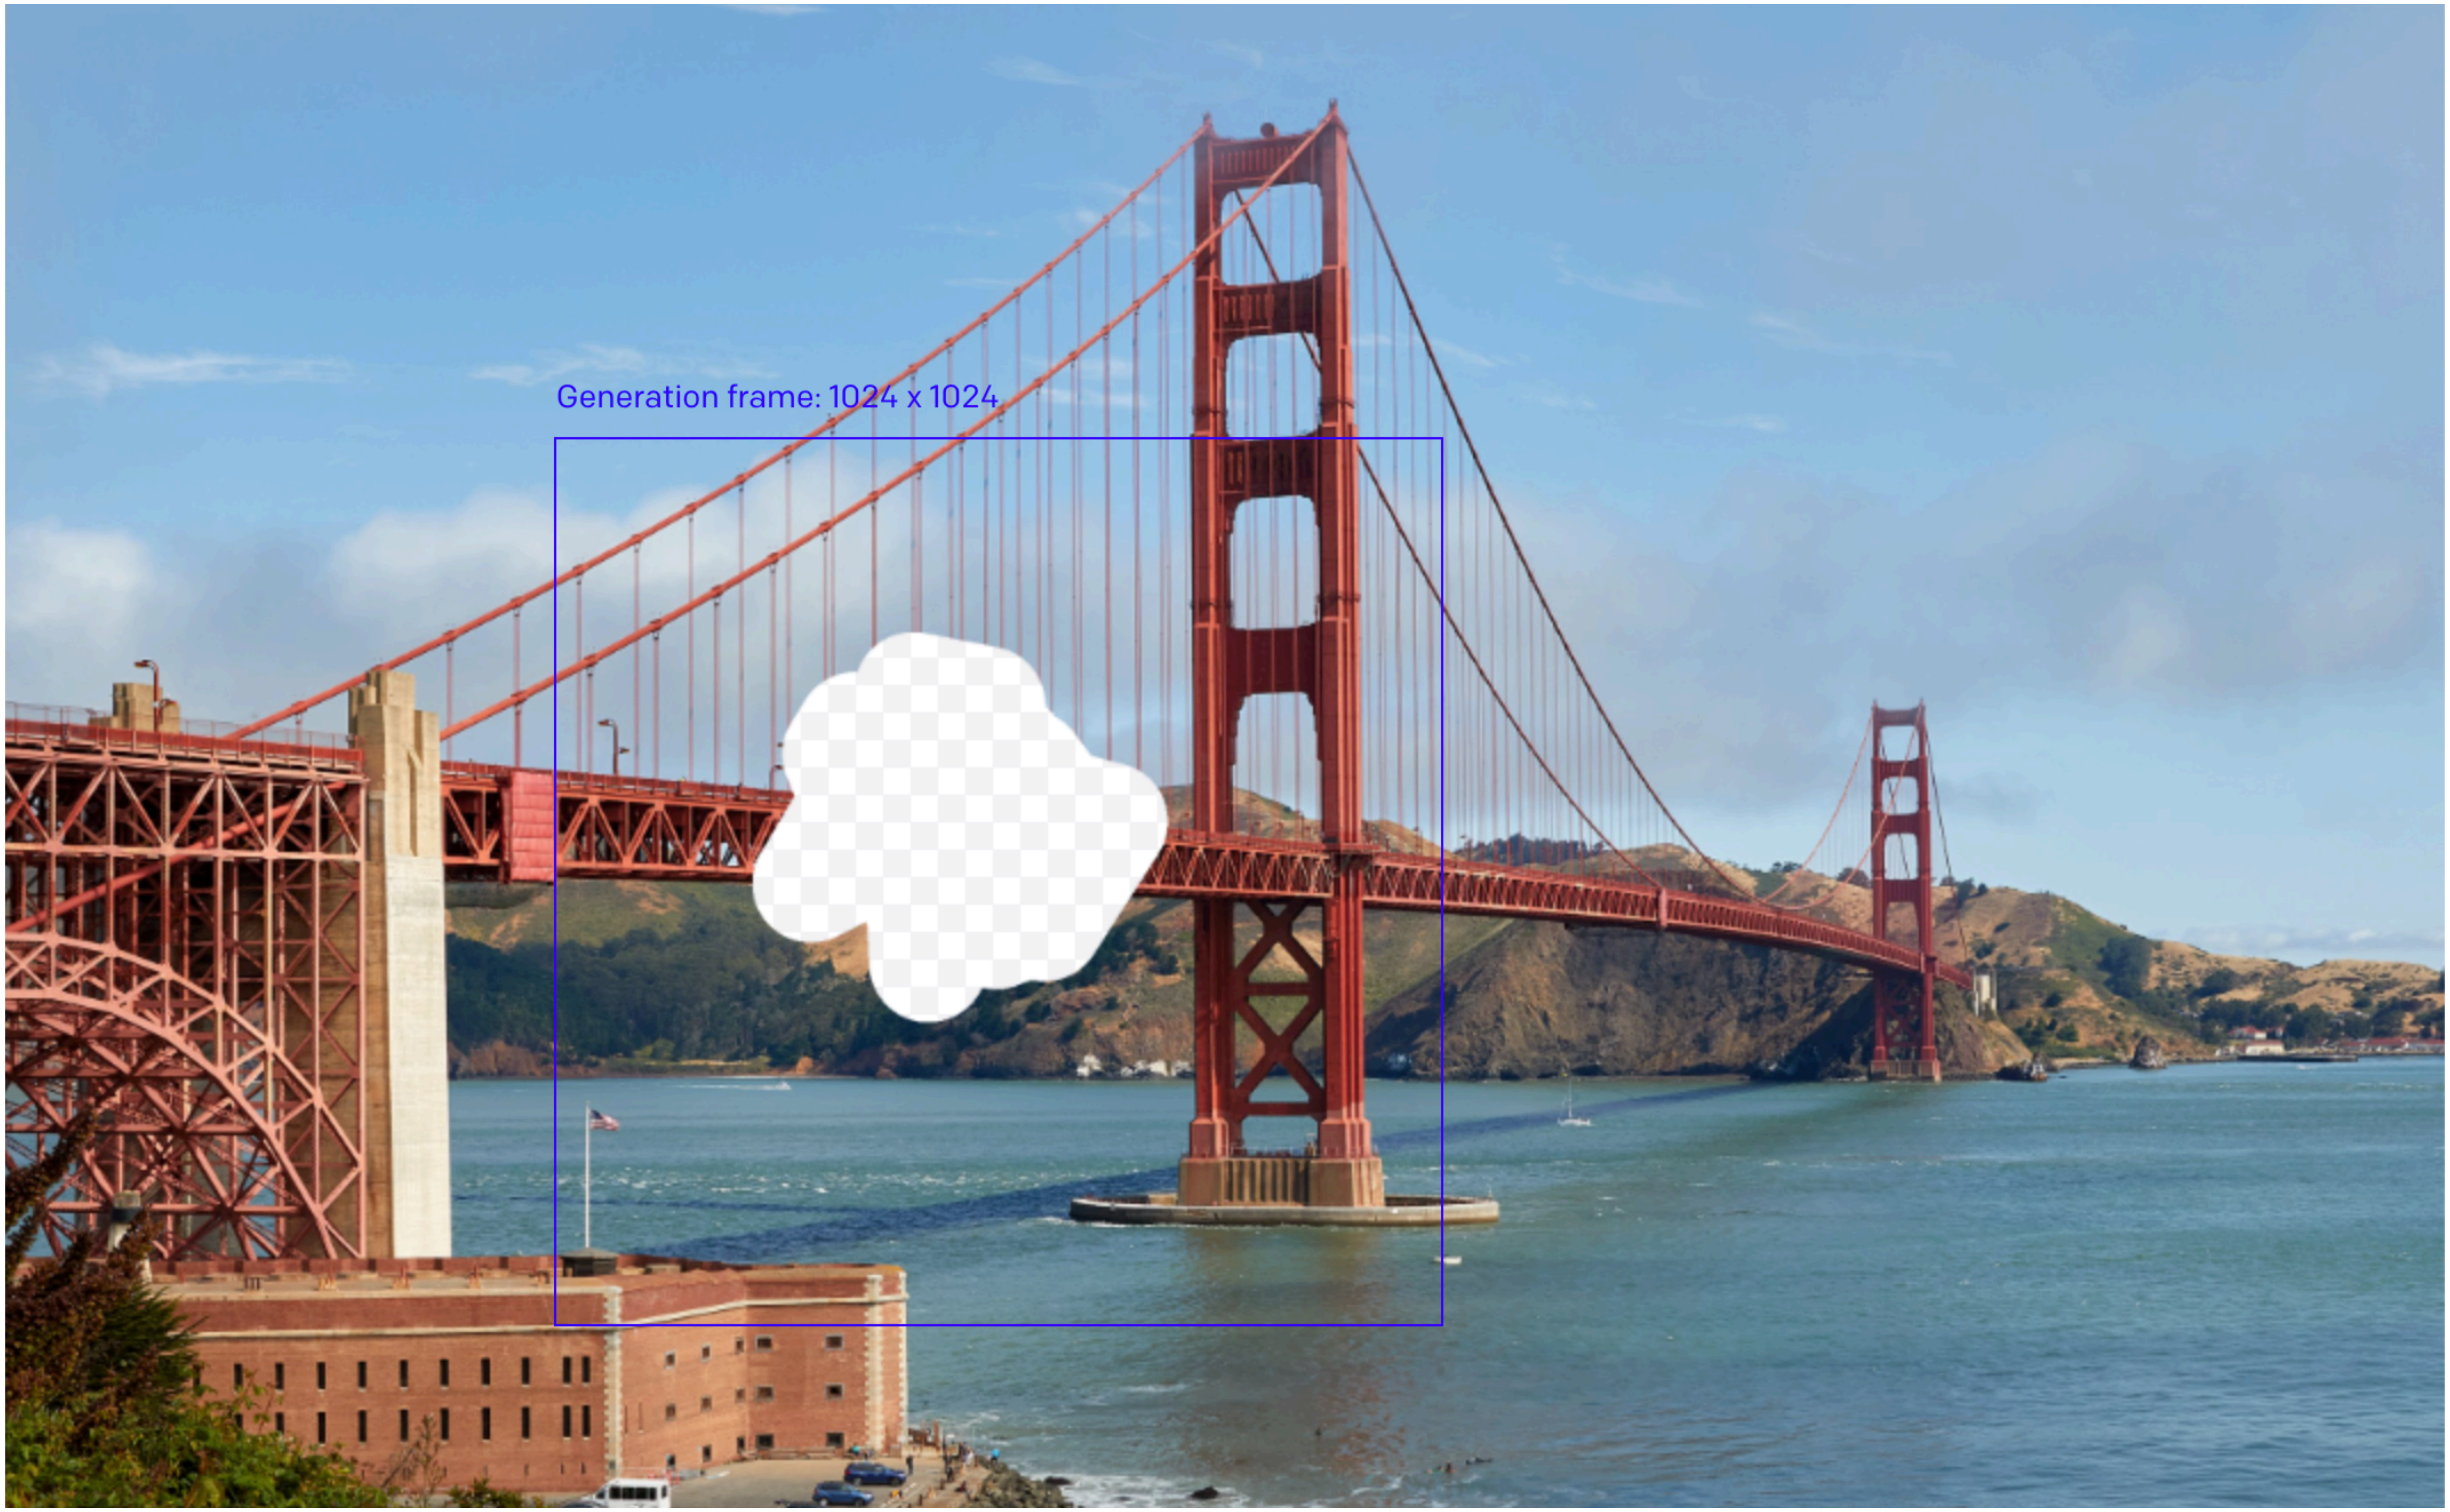

Define a signal cut-out

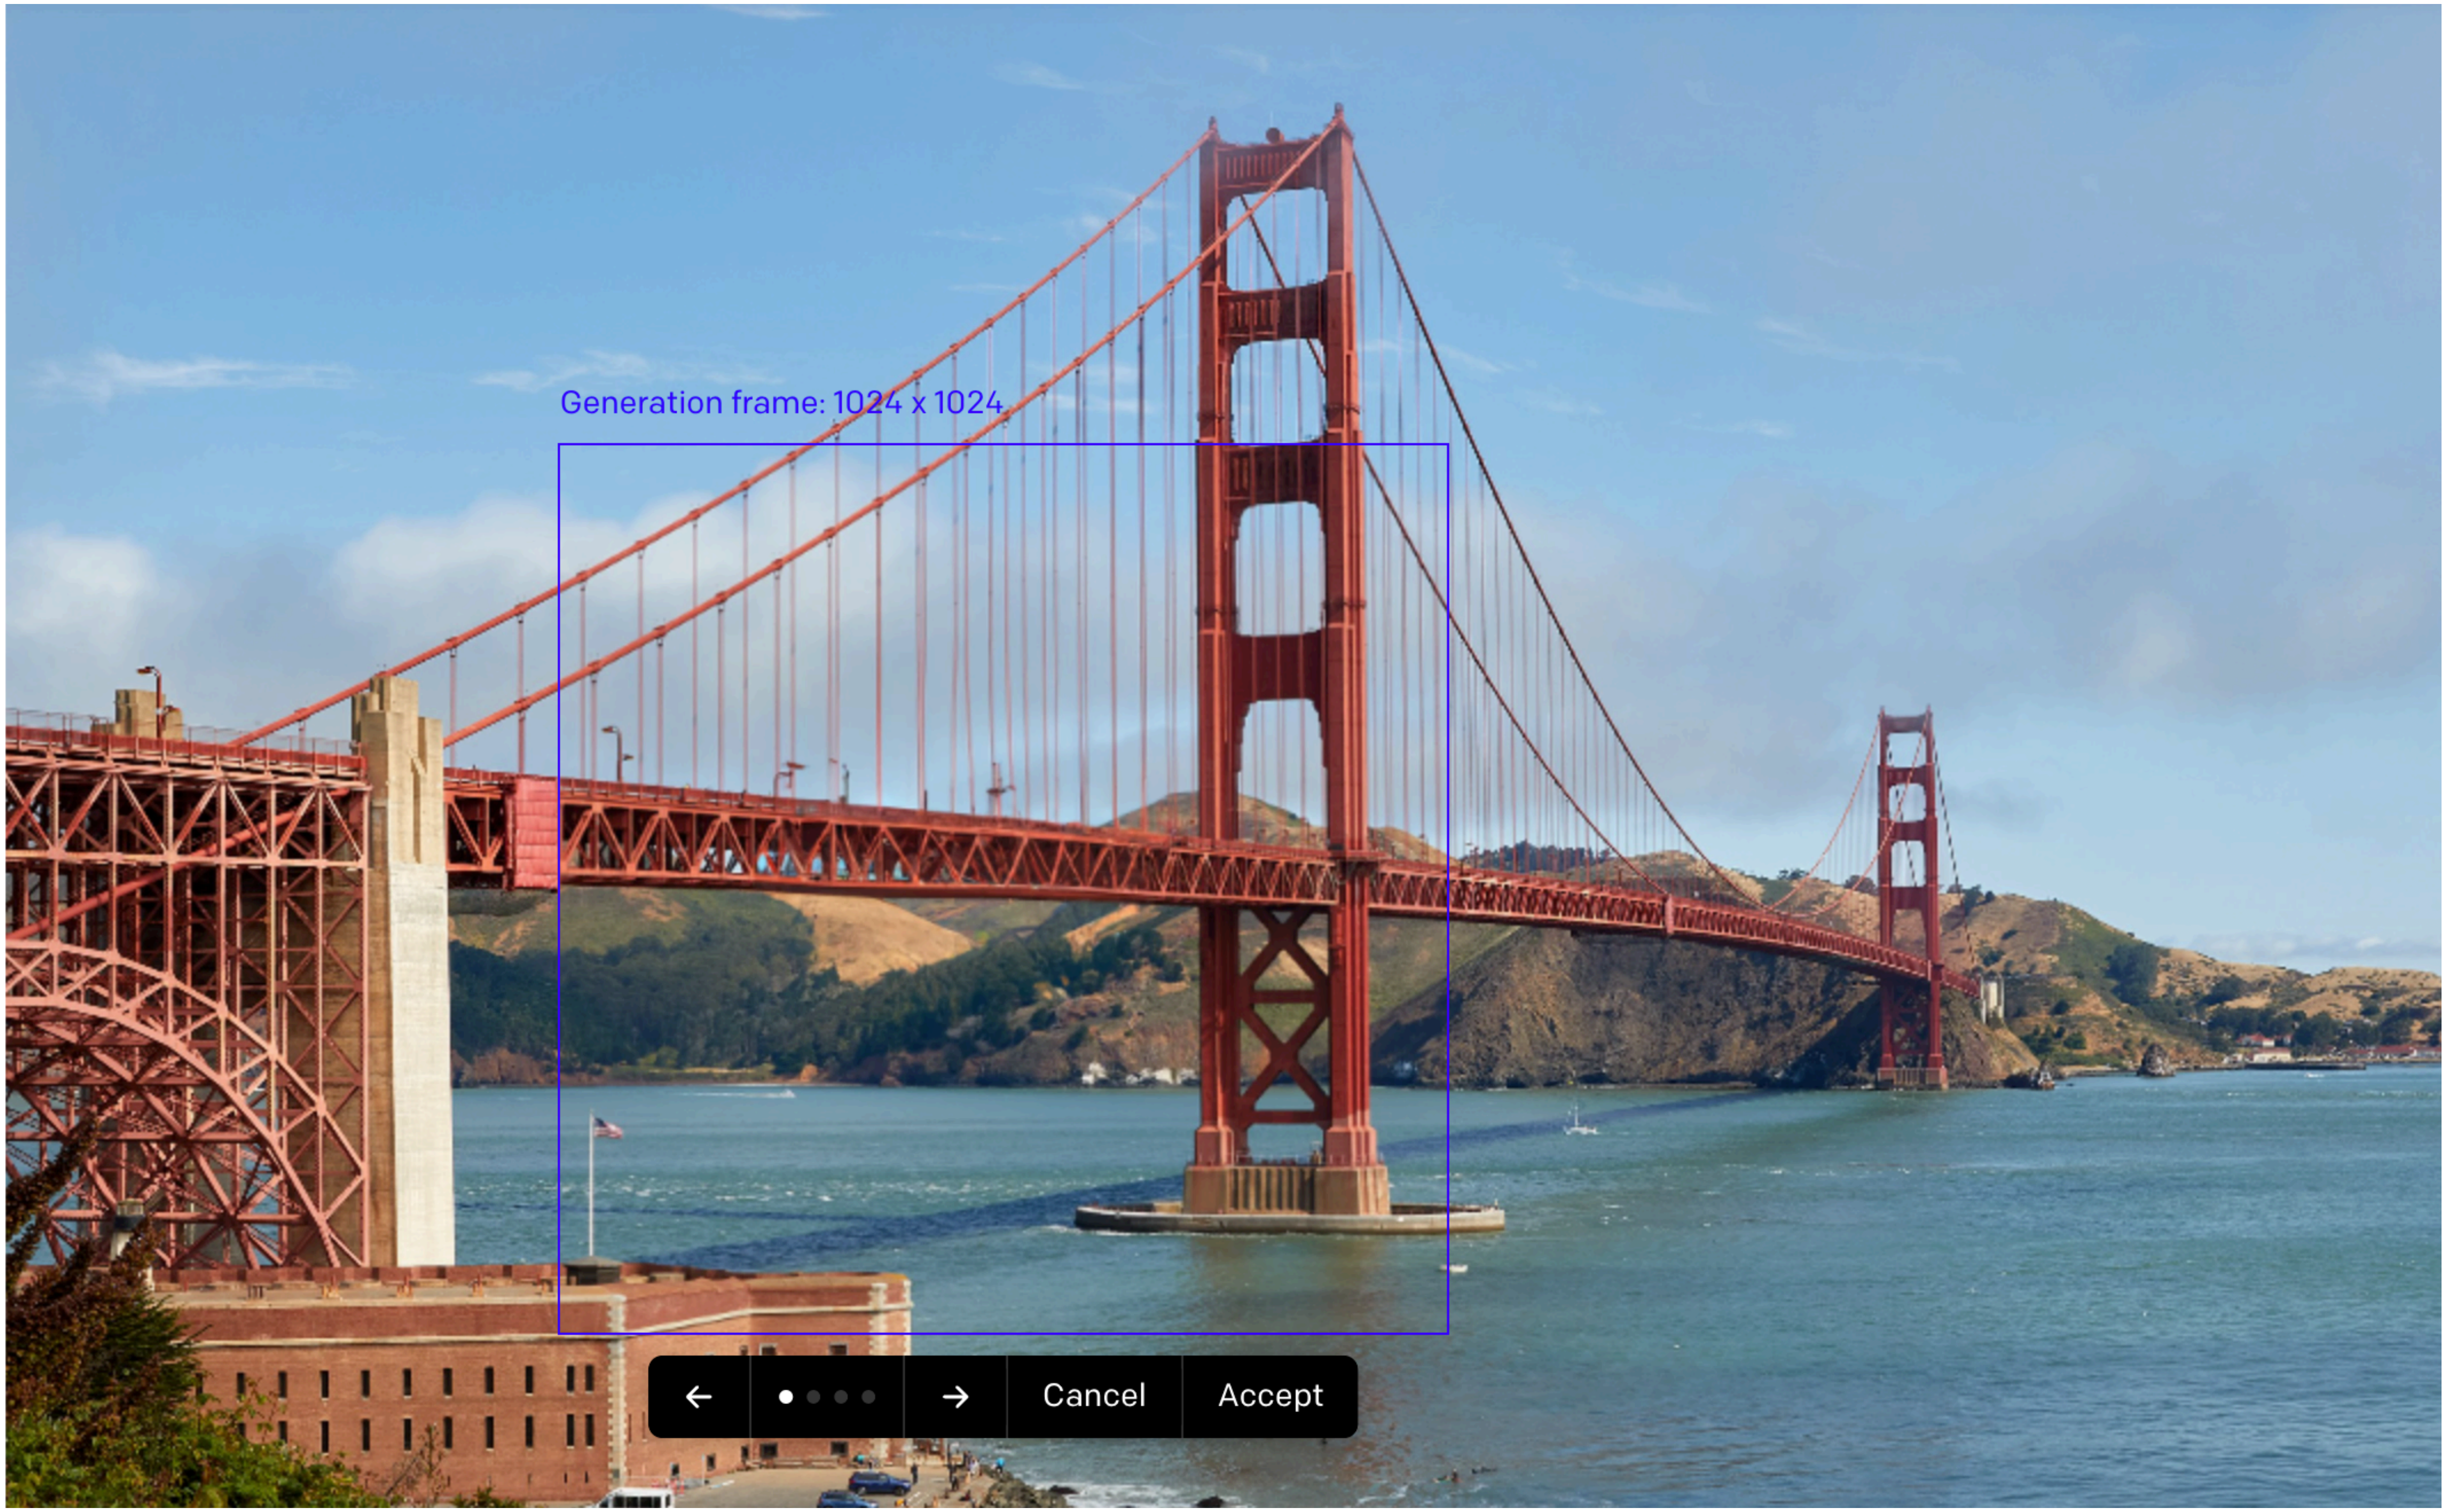

And interpolate the background

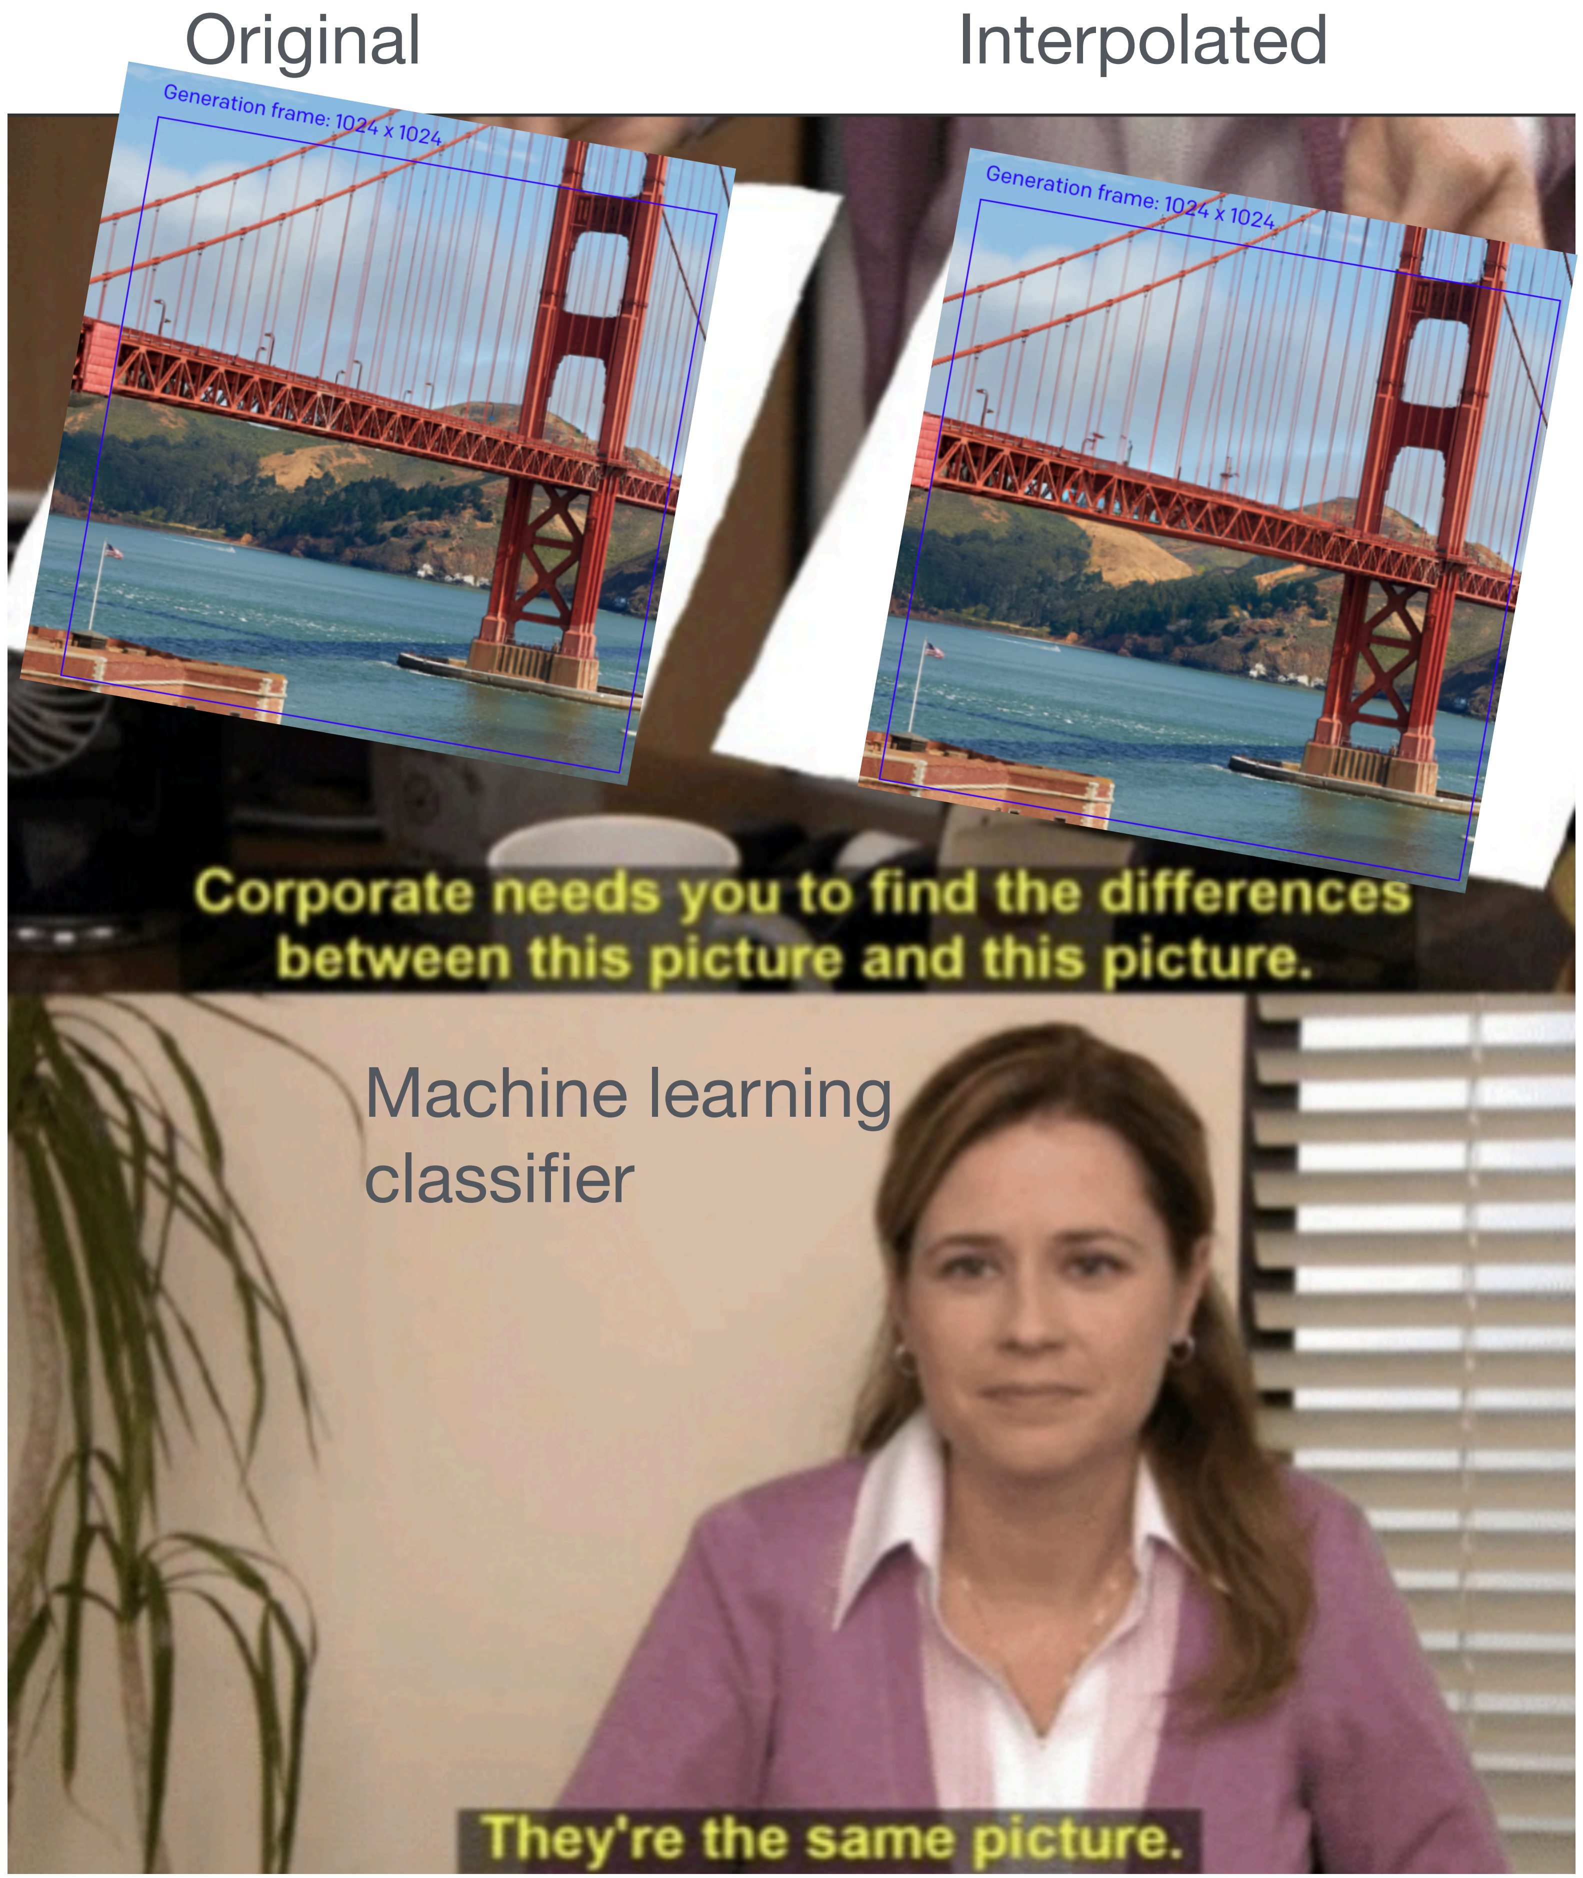
Original
Interpolated
Generation frame: 1024 x 1024
Generation frame: 1024 x 1024
Corporate needs you to find the differences between this picture and this picture.
Machine learning classifier
They're the same picture.

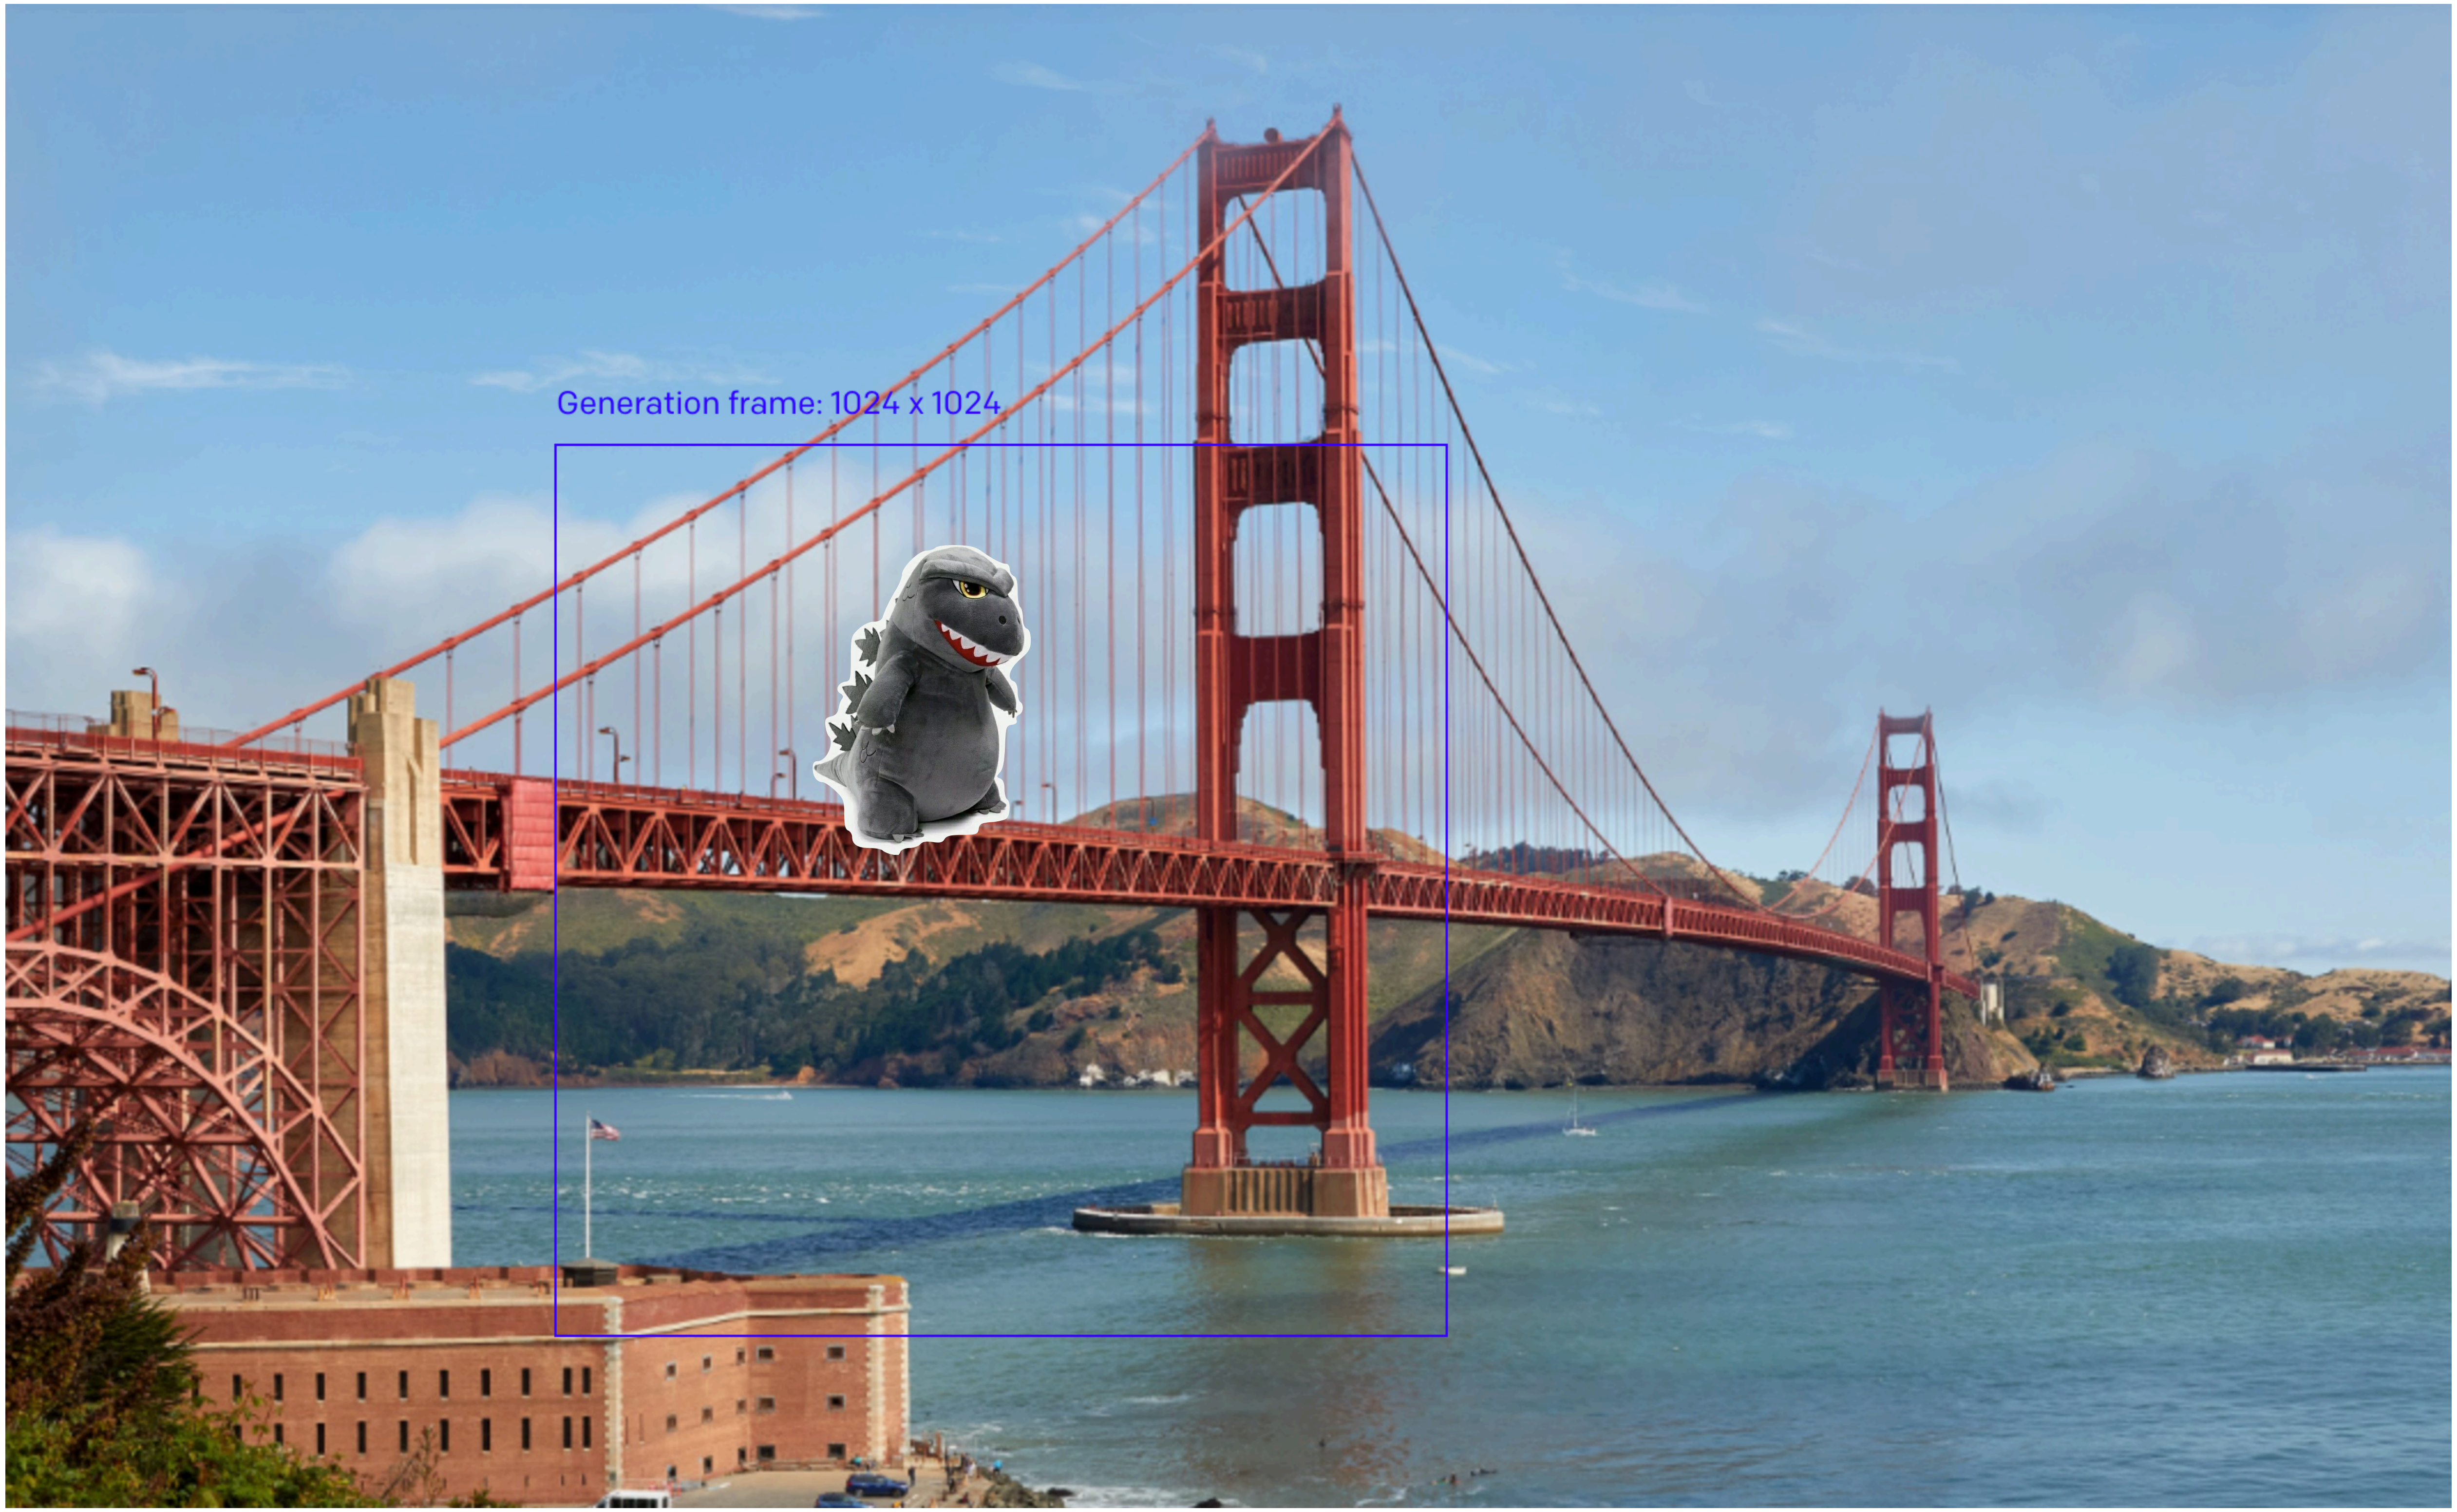
Generation frame: 1024 x 1024

Original Interpolated

CATHODE

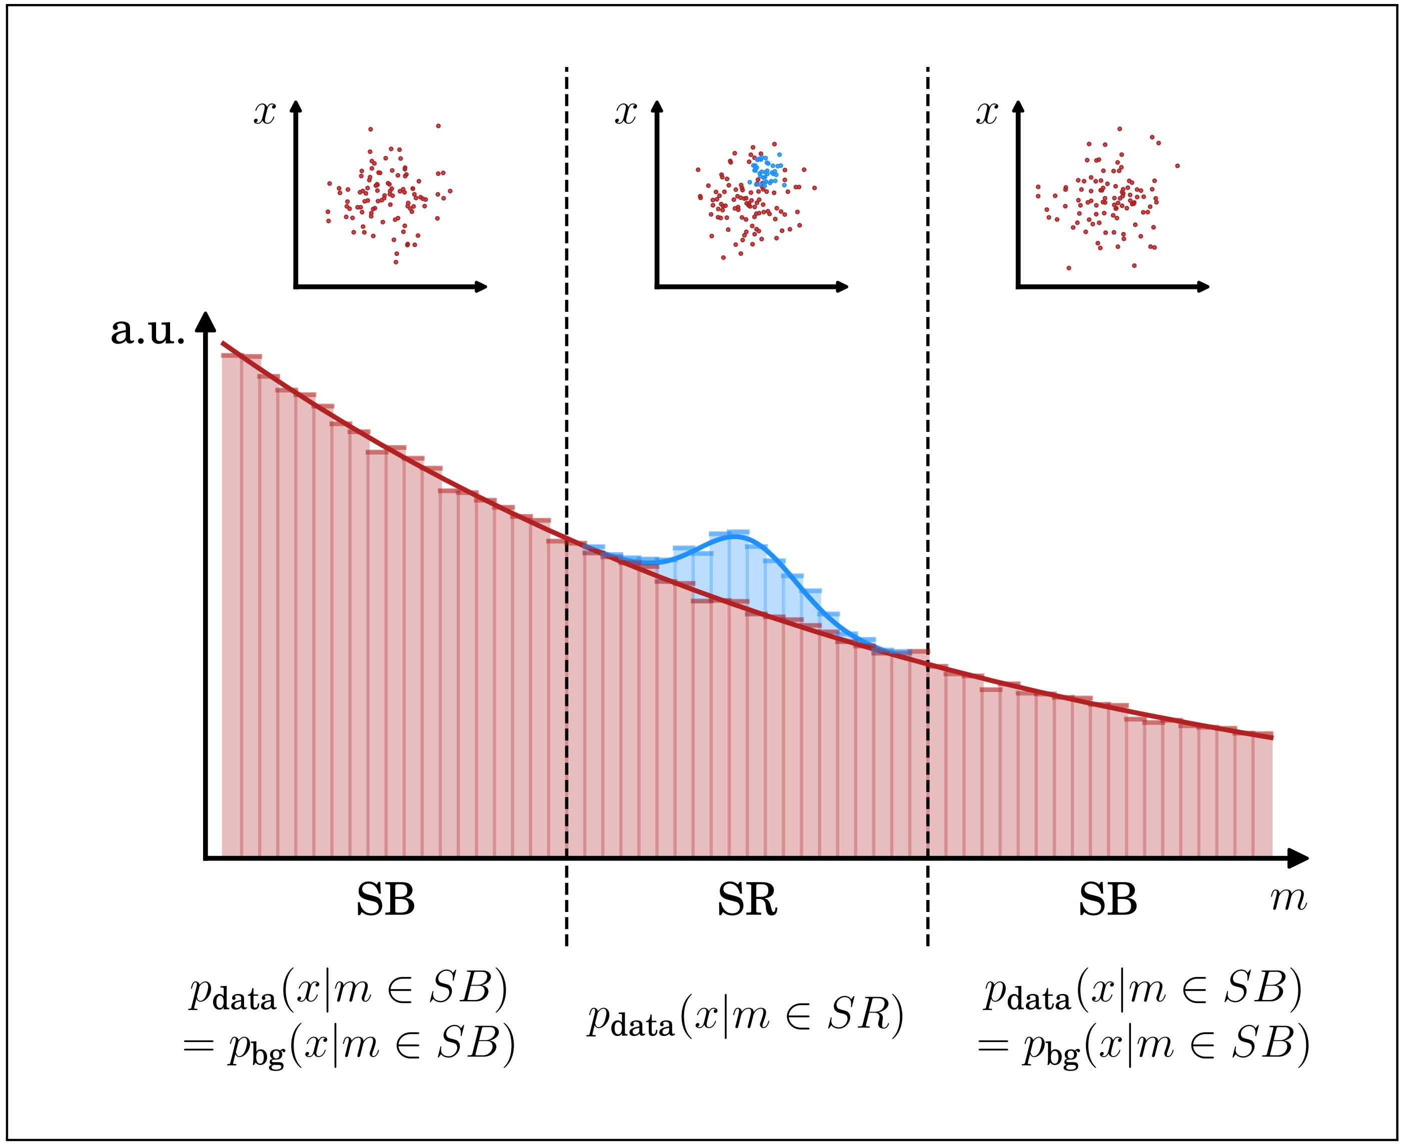

Consider resonant anomalies: fully data-based construction of anomaly detection score

GK, Nachmann, Shih et al 2101.08320;
Hallin, .., **GK** et al 2109.00546;

CATHODE

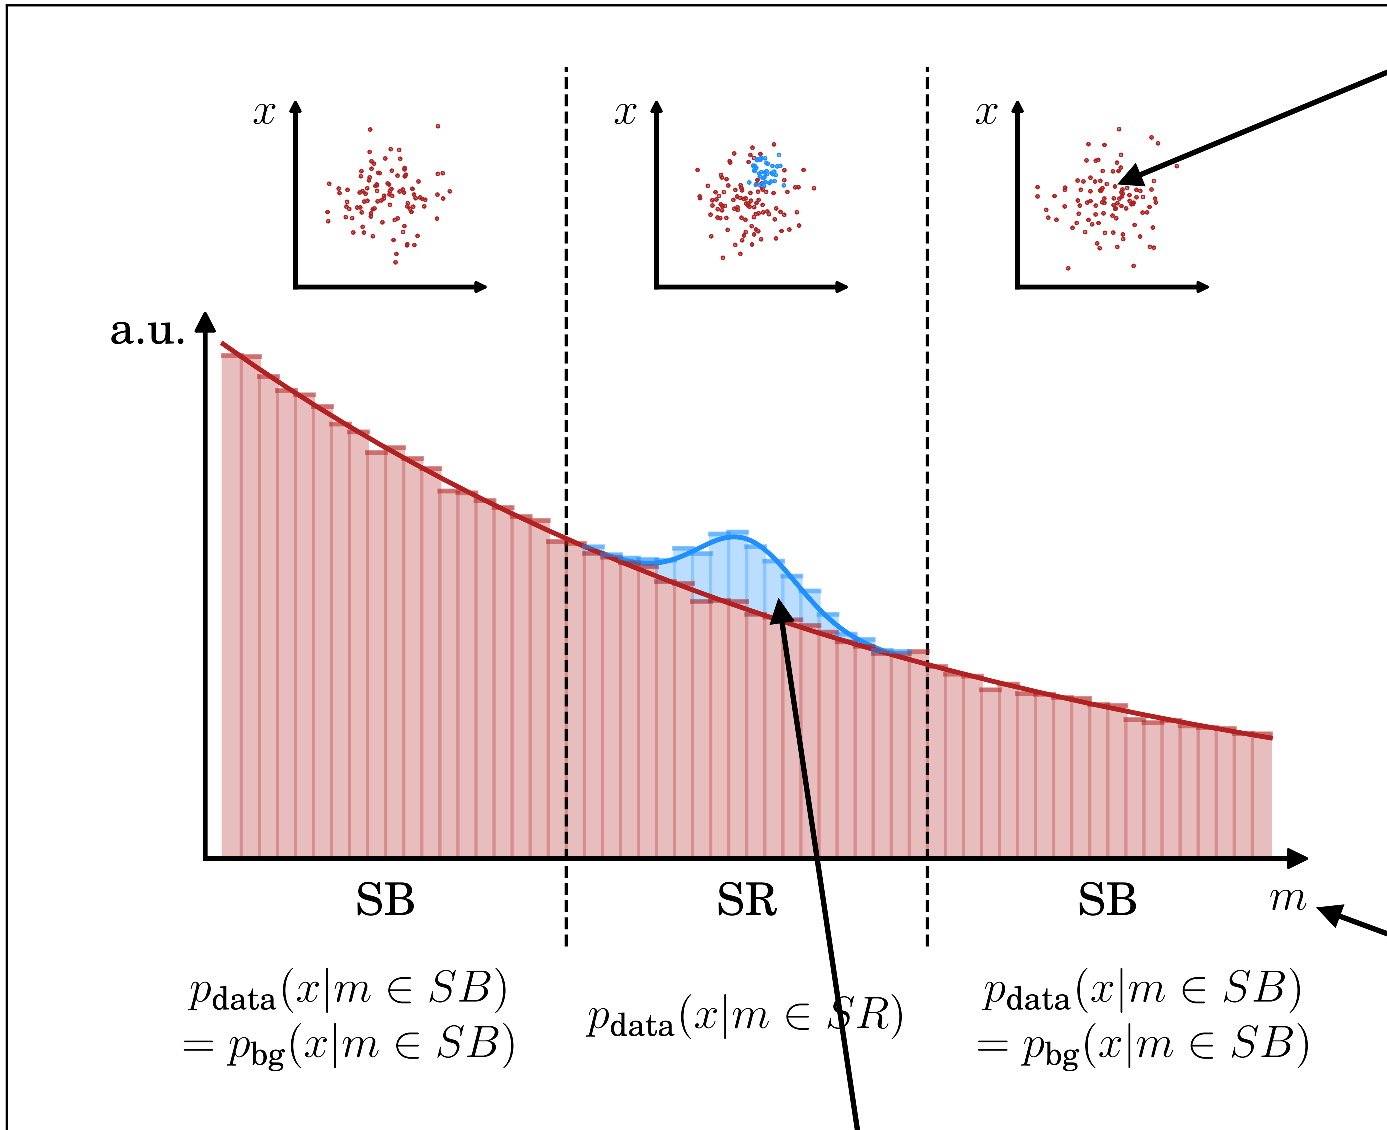

4 additional features

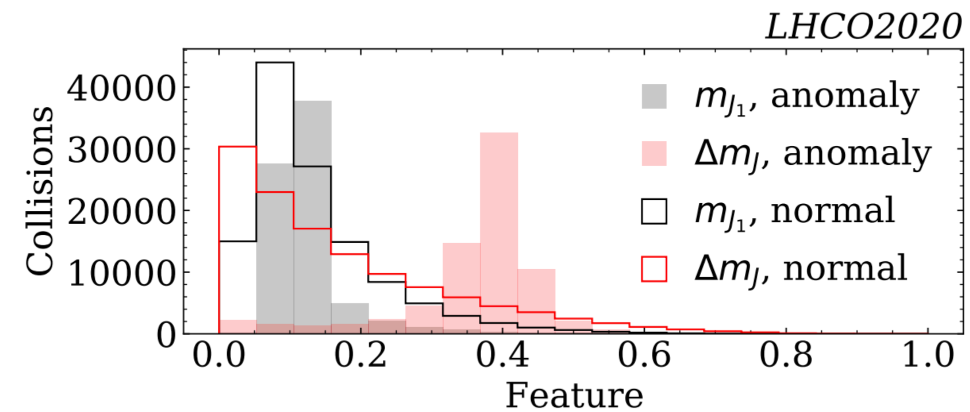

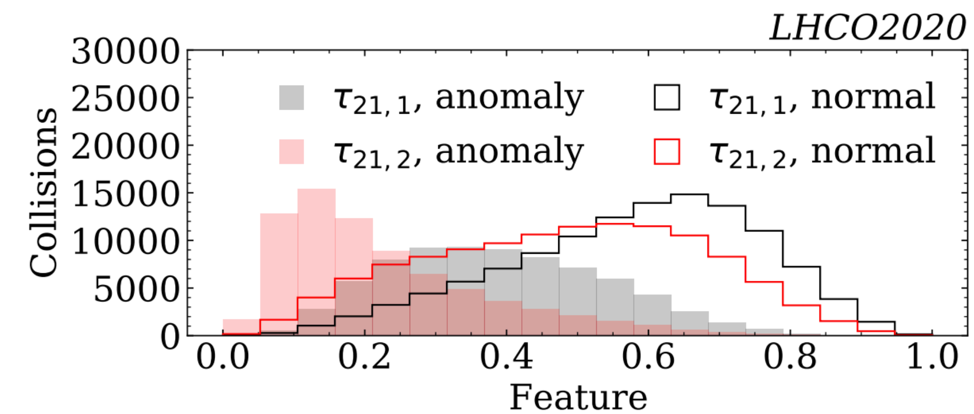

one resonant feature

Signal too small
to be visible in
inclusive distribution

GK, Nachmann, Shih et al 2101.08320;
Hallin, .., **GK** et al 2109.00546;

CATHODE

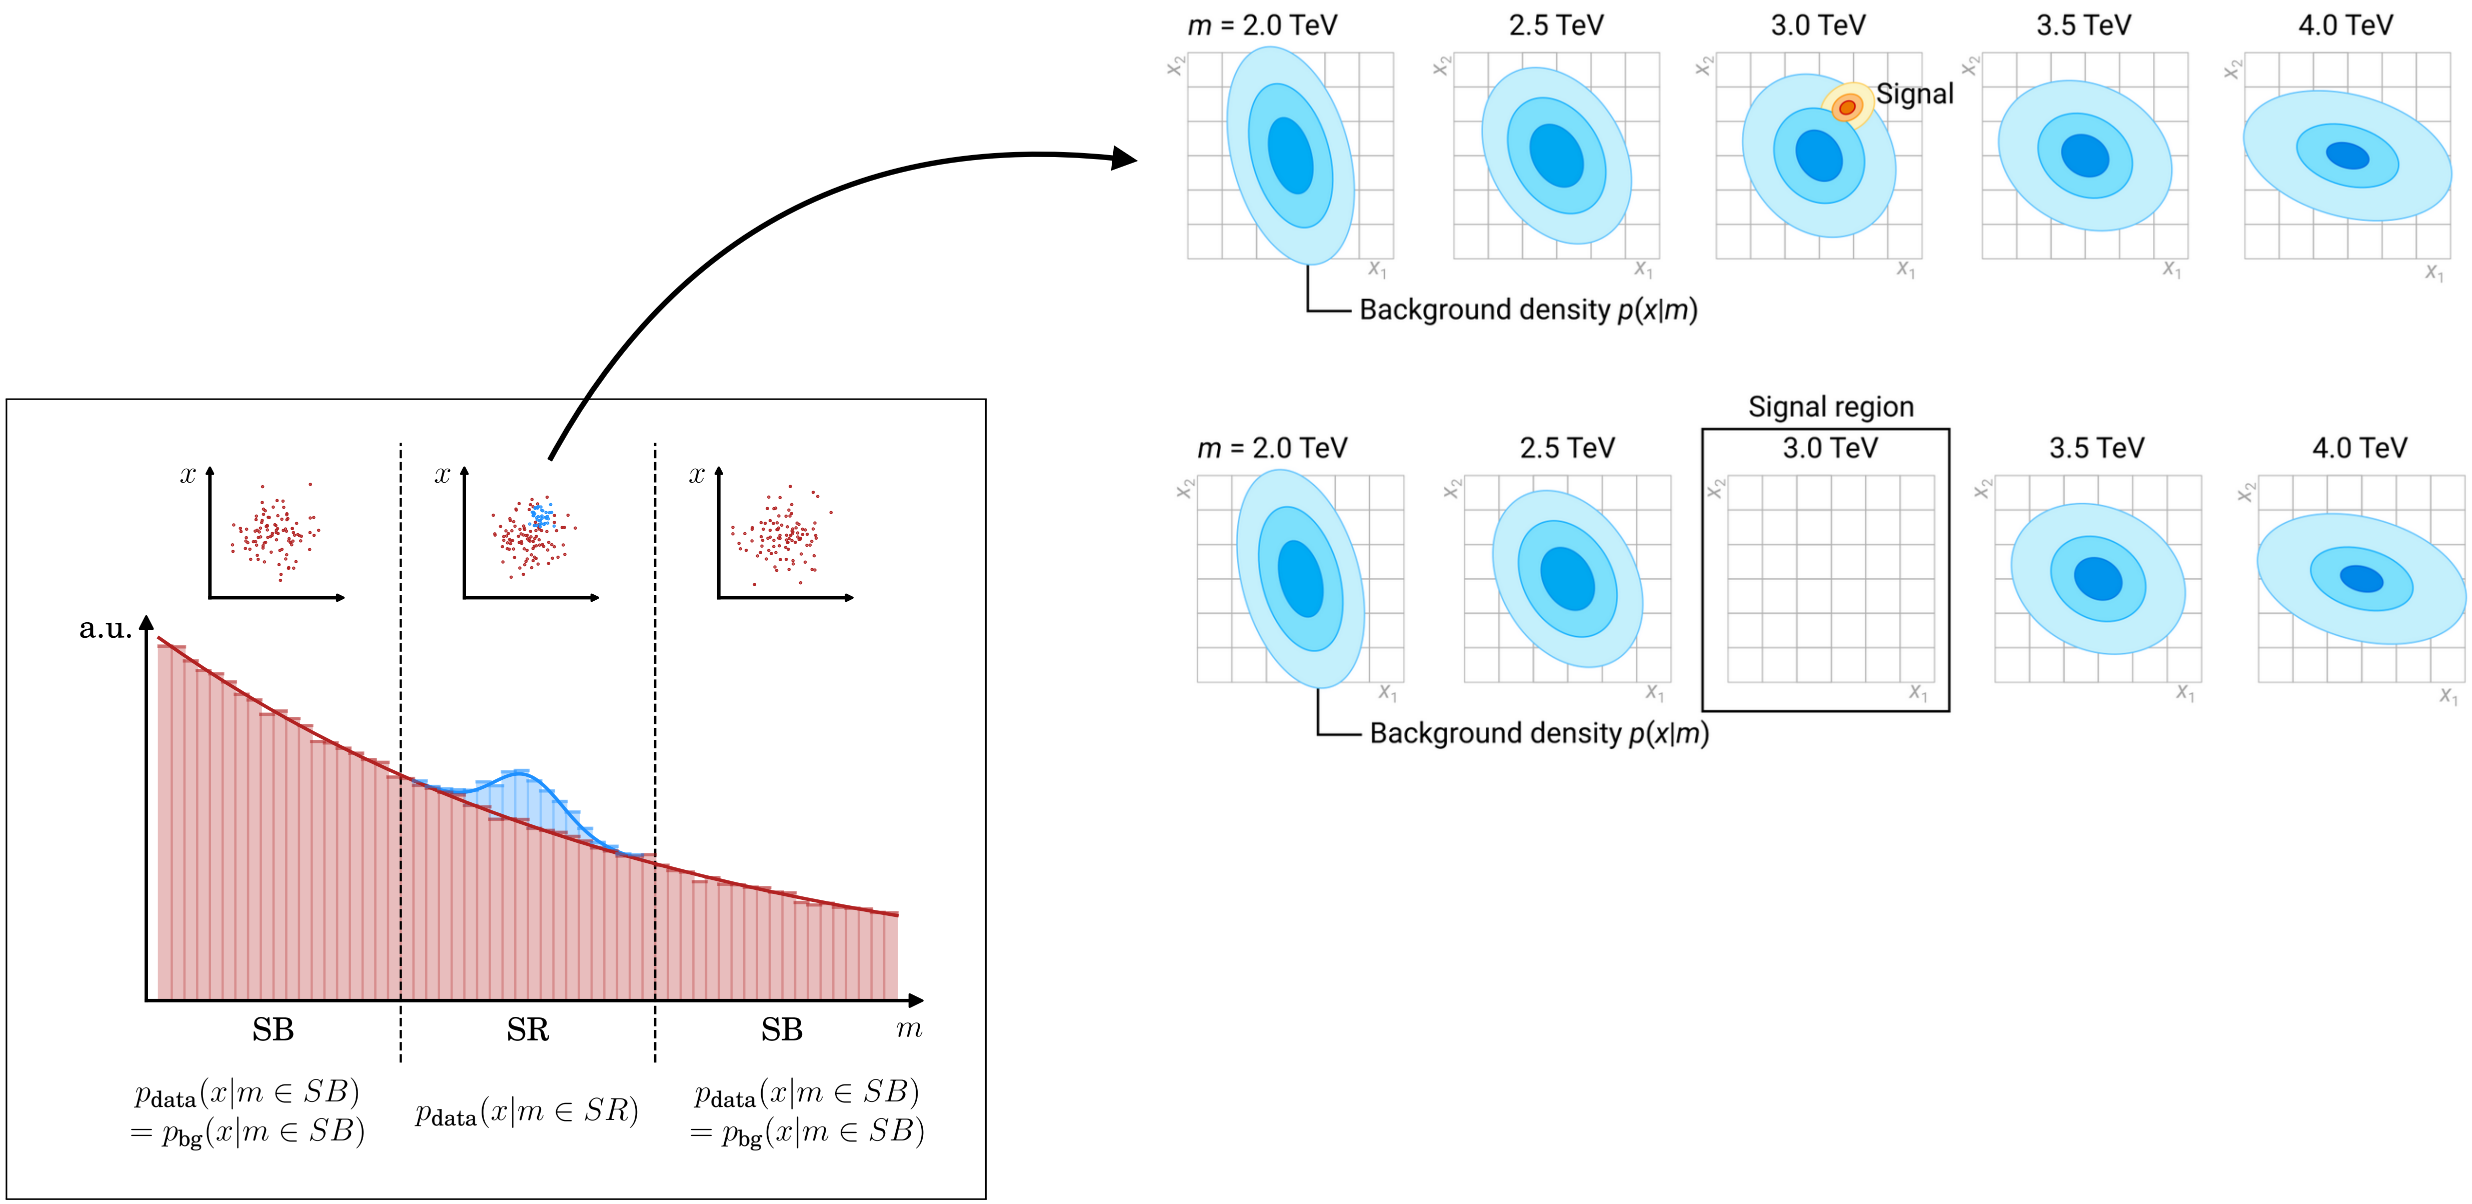

GK, Nachmann, Shih et al 2101.08320; Hallin, .., **GK** et al 2109.00546; Figure by L. Moreaux

CATHODE

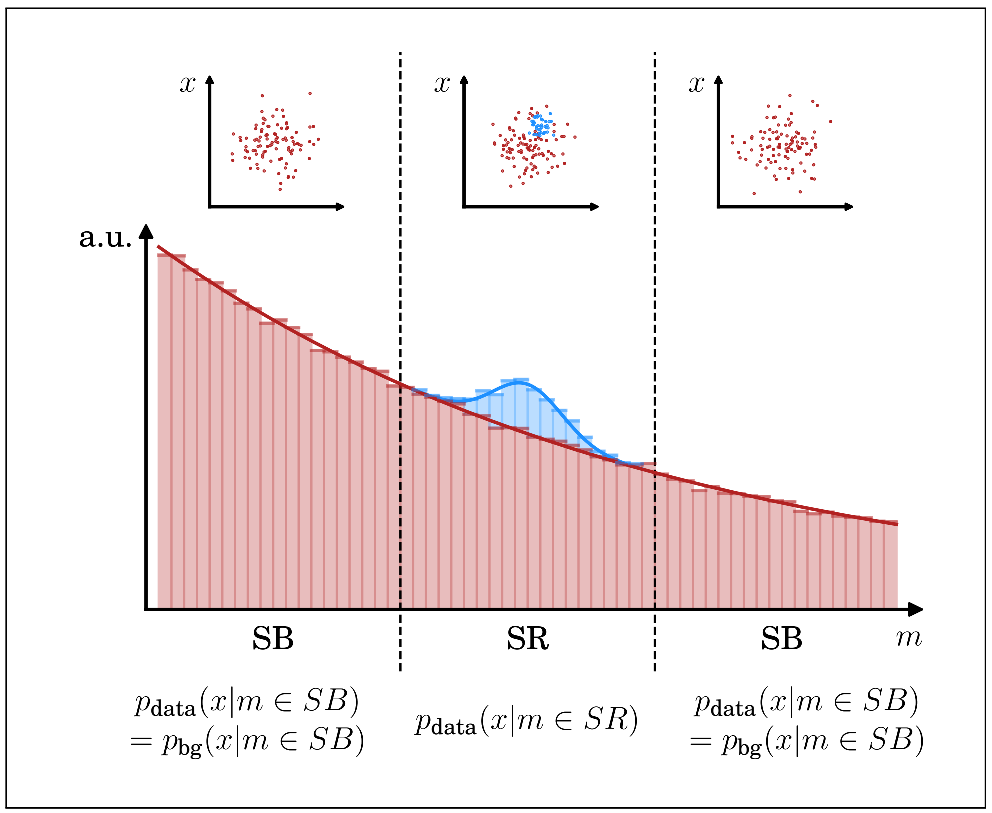

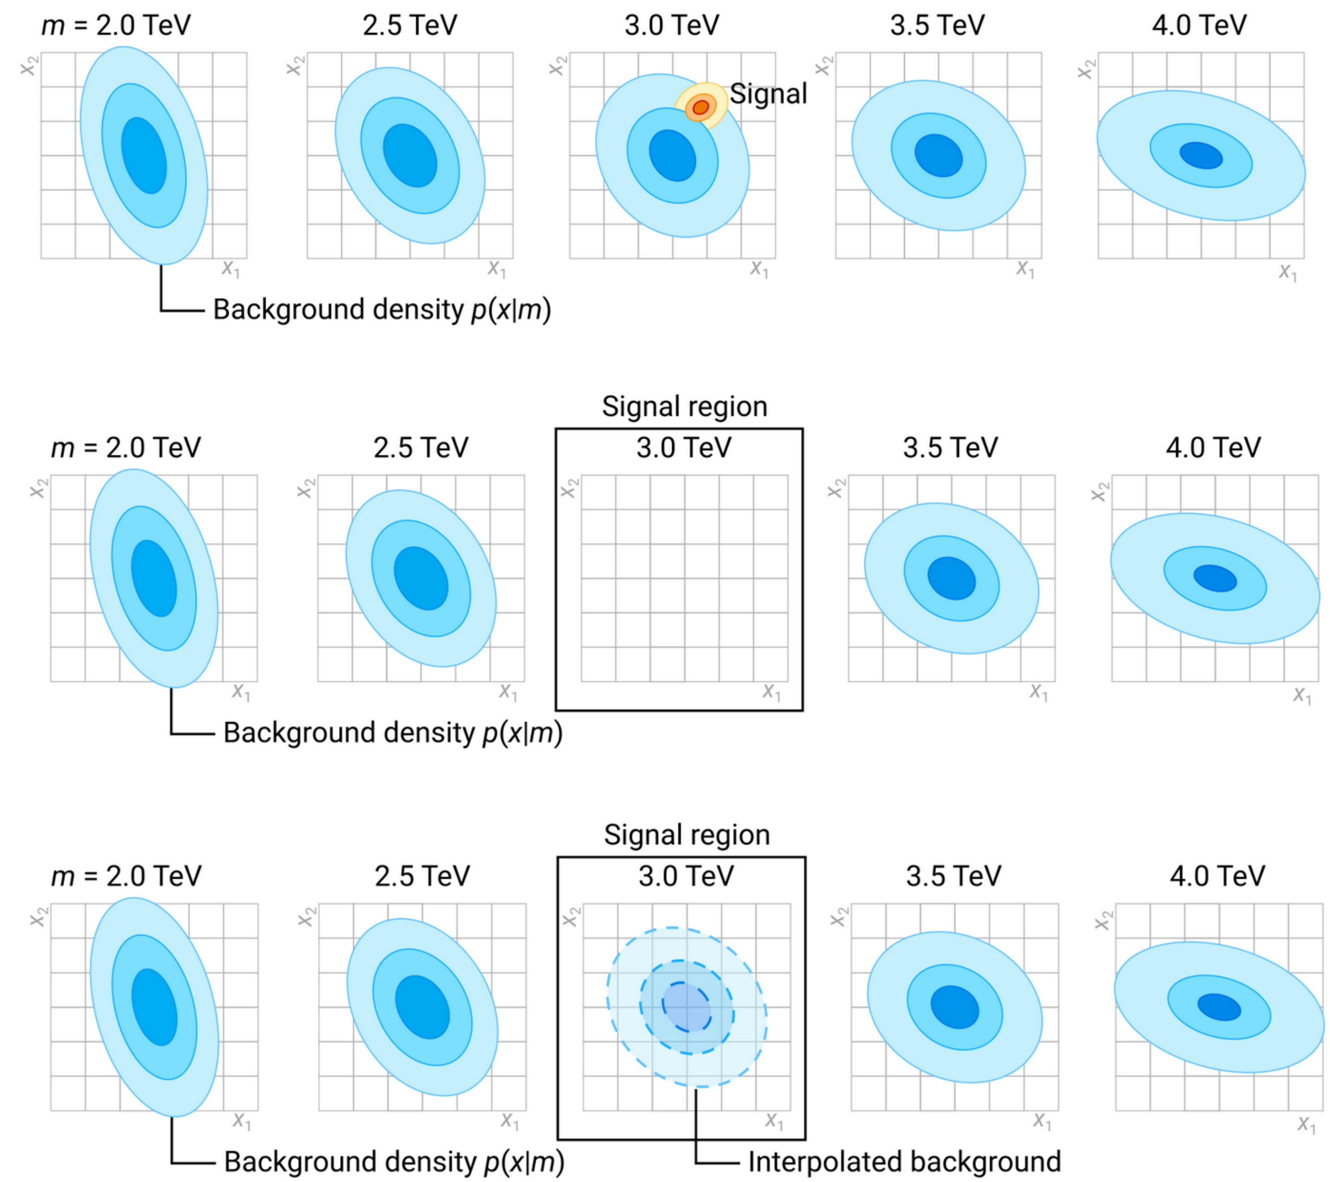

Train conditional generative*
model and interpolate

*e.g. normalising flow /
diffusion models

CATHODE

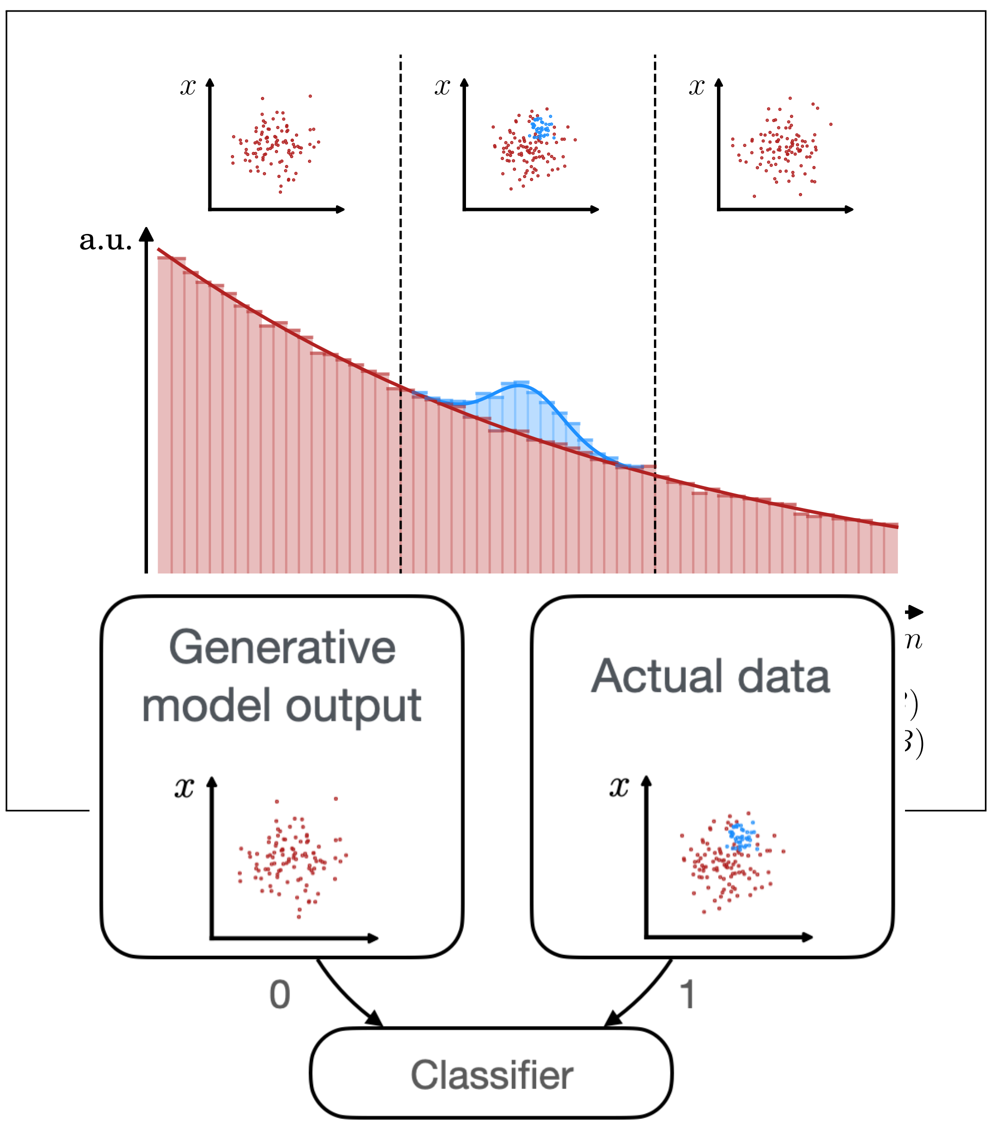

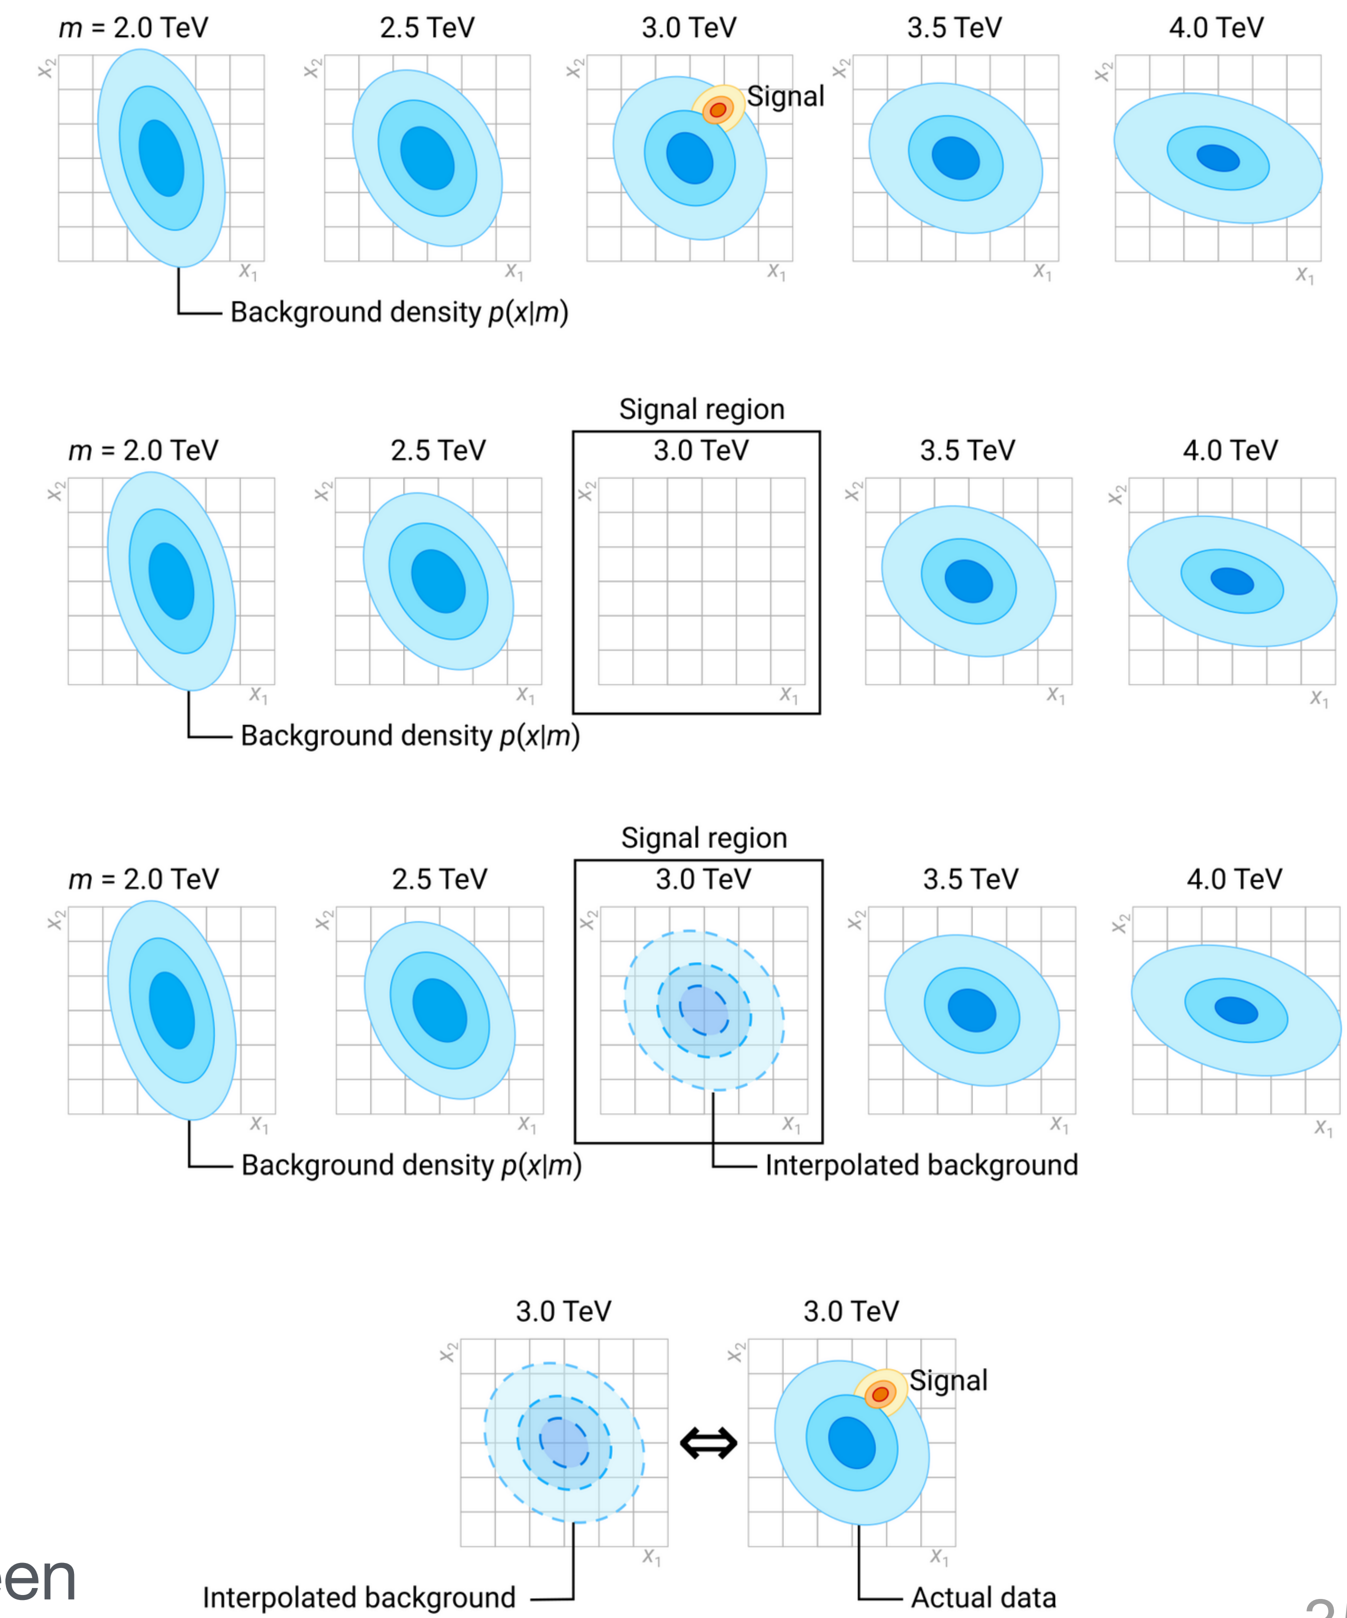

Train a classifier between
prediction vs data

CATHODE

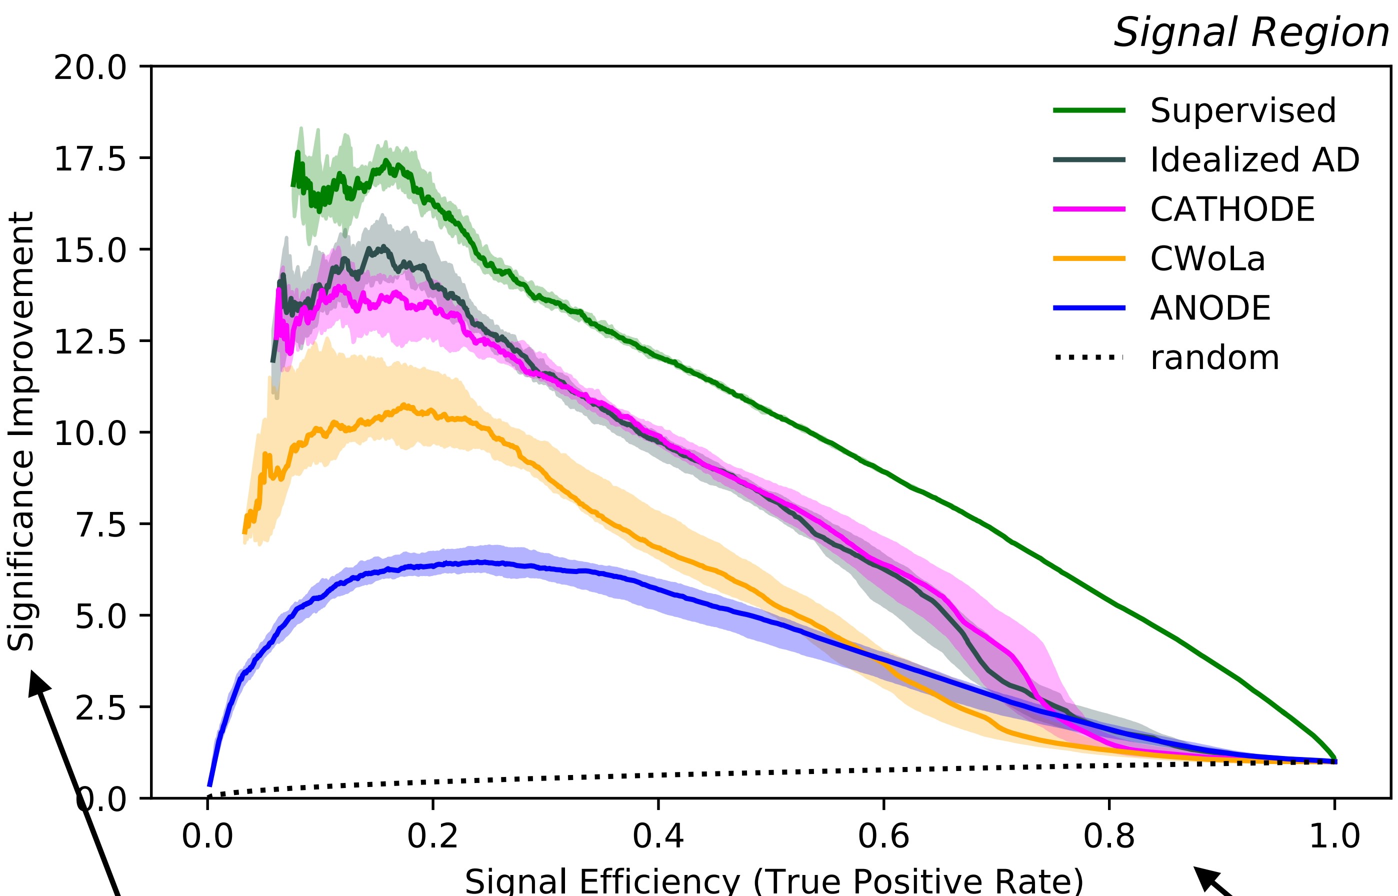

Use classifier to identify anomalies

$$\mathrm{SIC} = \frac{\epsilon_S}{\sqrt{\epsilon_B}}$$

Promising, but does this work on data?

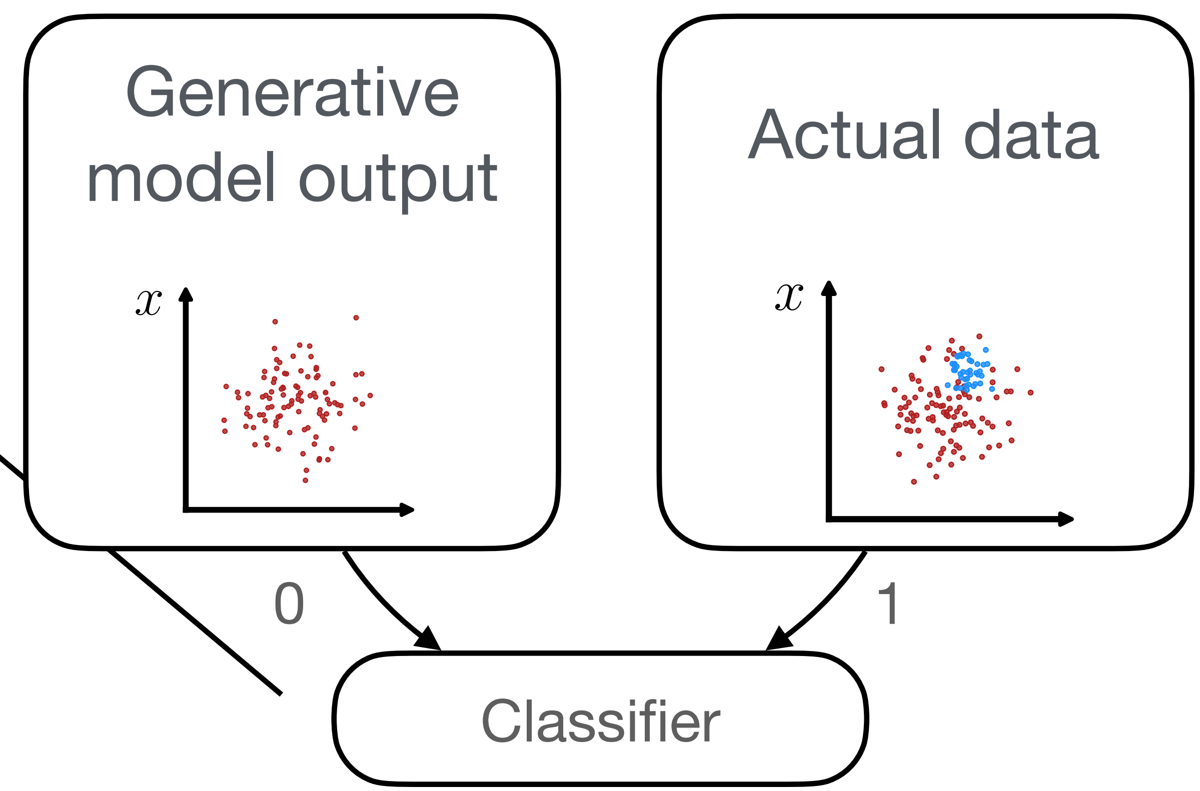

CASE

- Result by the CMS collaboration
- Full Run 2 dataset
- 6 anomaly detectors in parallel

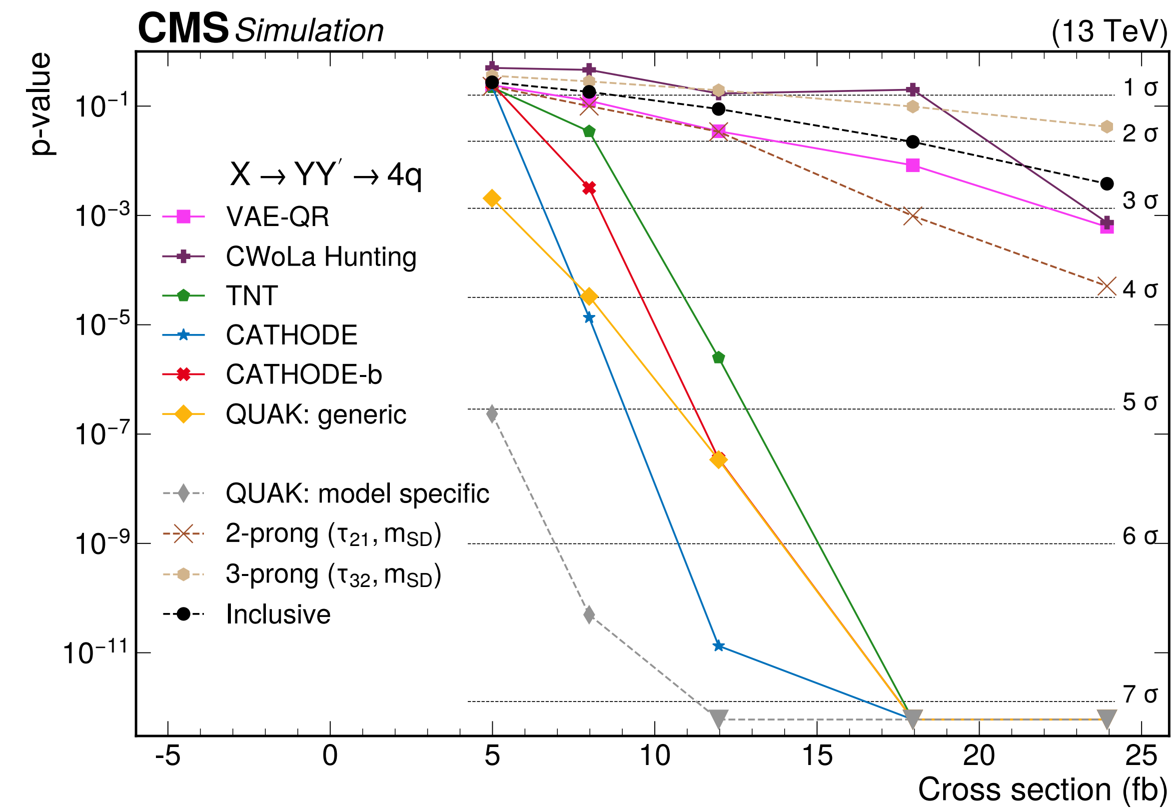

Test sensitivity gain via injected signals in simulation

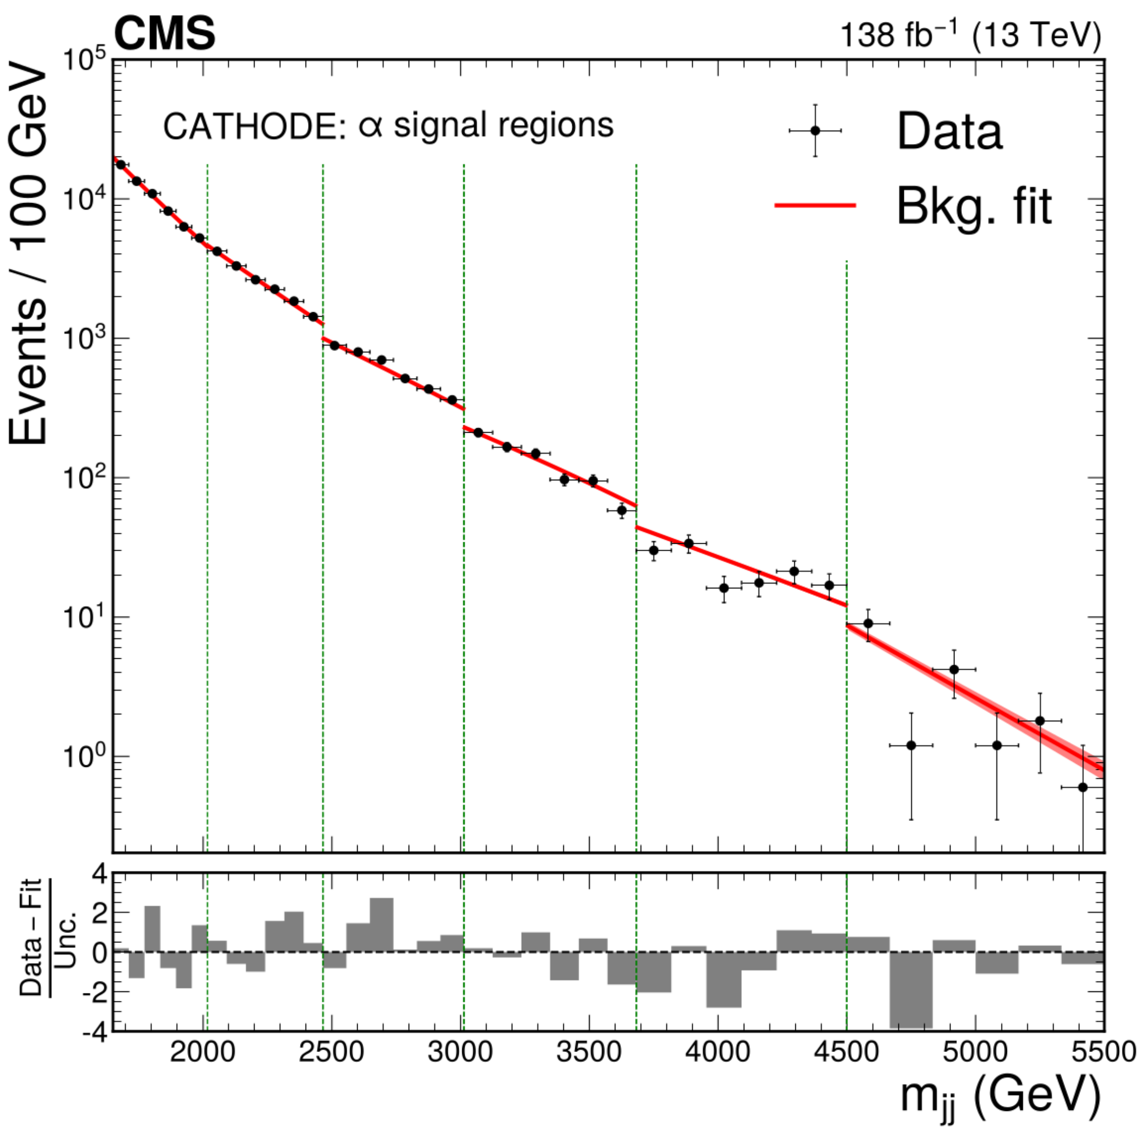

Fully train CATHODE on data

Select top 1% most anomalous events, perform bump-hunt

No signal-like outlier: set limits

Improve breadth and sensitivity

CMS 2412.03747

Outline

Tools for Discovery

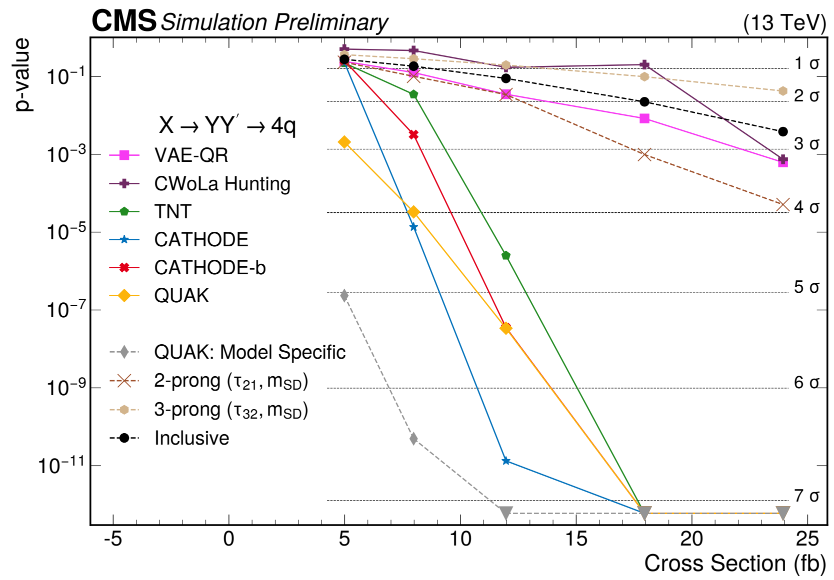

Discovery Strategies

Autonomous Discovery

Qualitatively new approaches enabled by AI

Outline

Tools for Discovery

Discovery Strategies

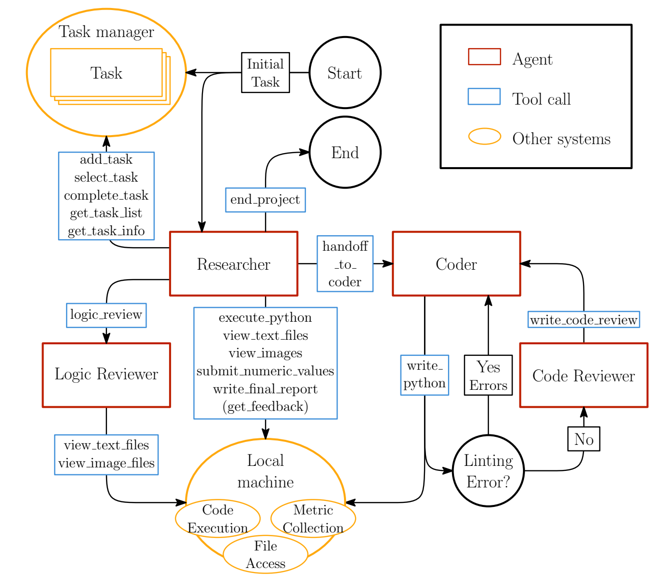

Autonomous Discovery

Increasing autonomy of AI systems

Large Language Models

Most impressive growth:
Large language models

Impact for physics?

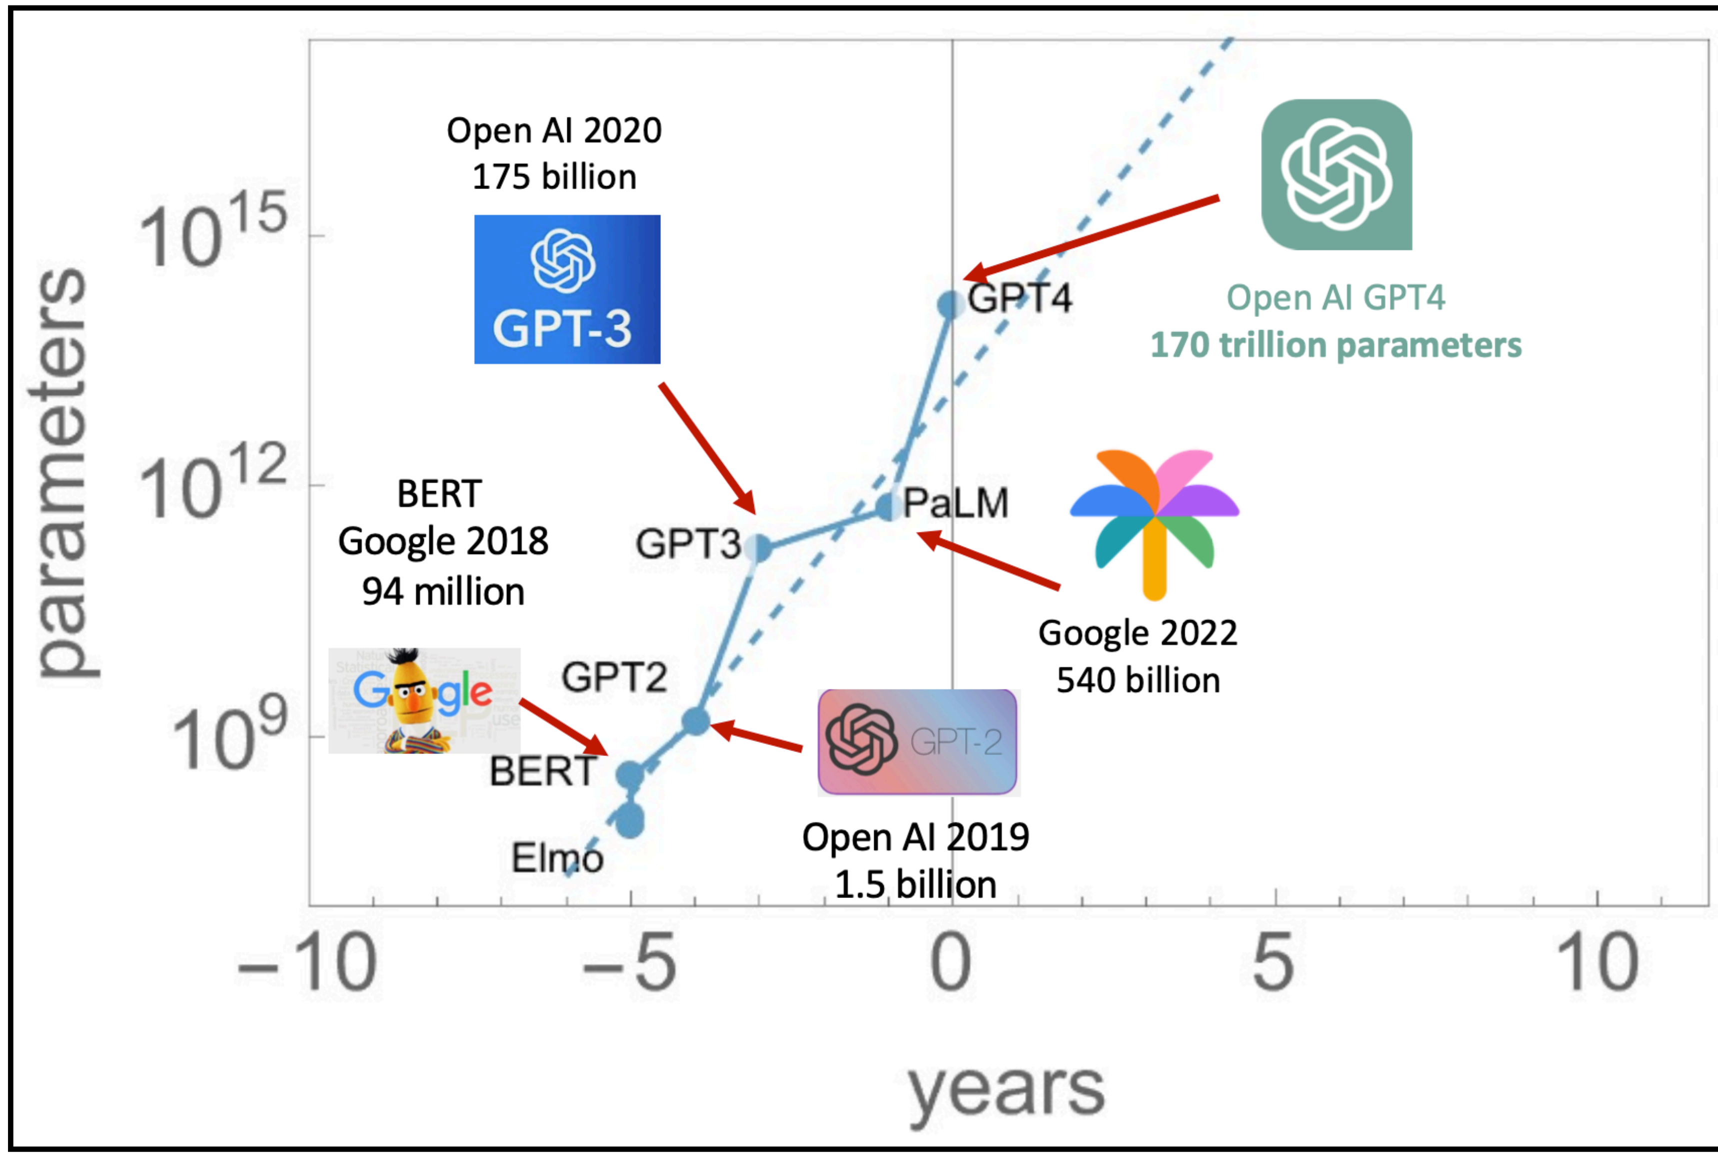

Large Language Models

Most impressive growth:
Large language models

Impact for physics:
Numerics
Symbolic

What else?

Virtual lab powered by 'AI scientists' super-charges biomedical research

Could human–AI collaborations be the future of interdisciplinary studies?

By Helena Kudiabor

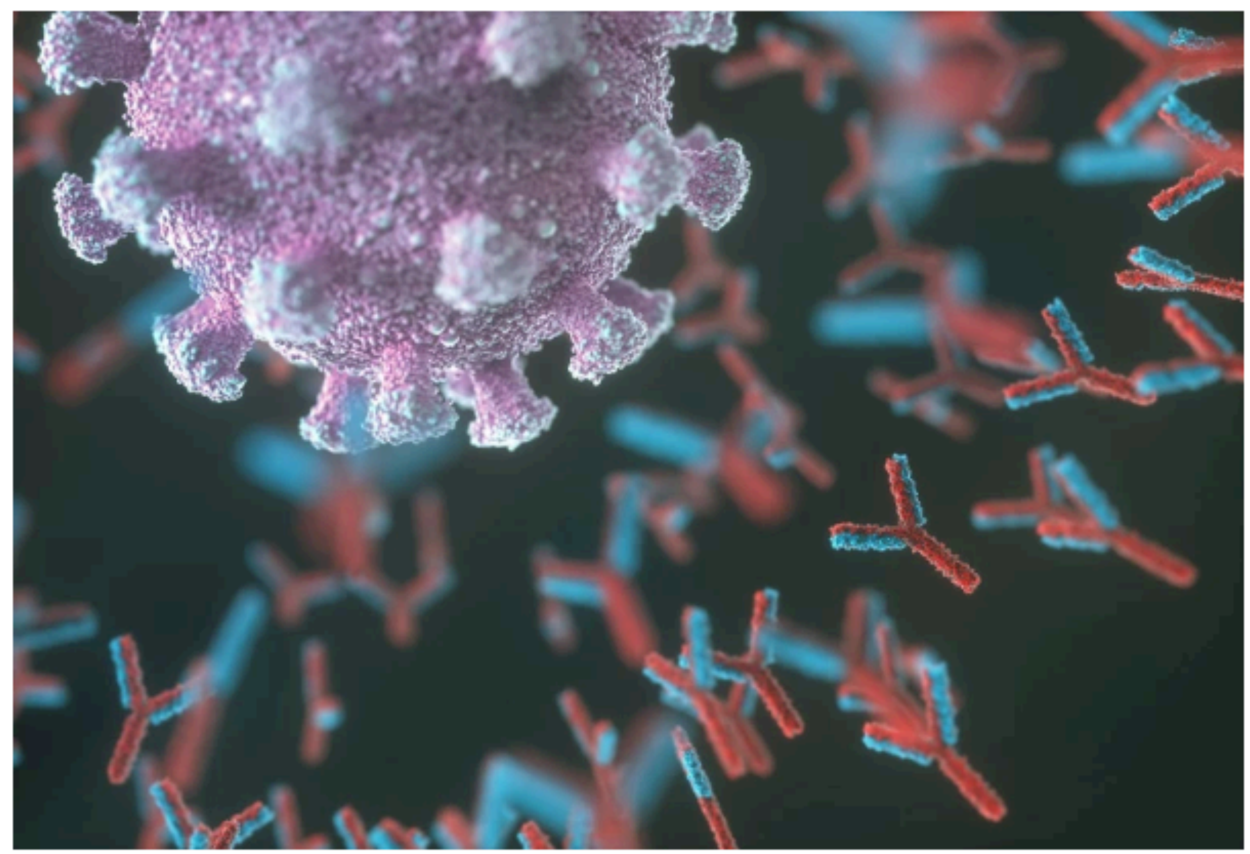

Large Language Models

Most impressive growth:
Large language models

Impact for physics:
Numerics
Symbolic

Include agent-based models as collaborators?

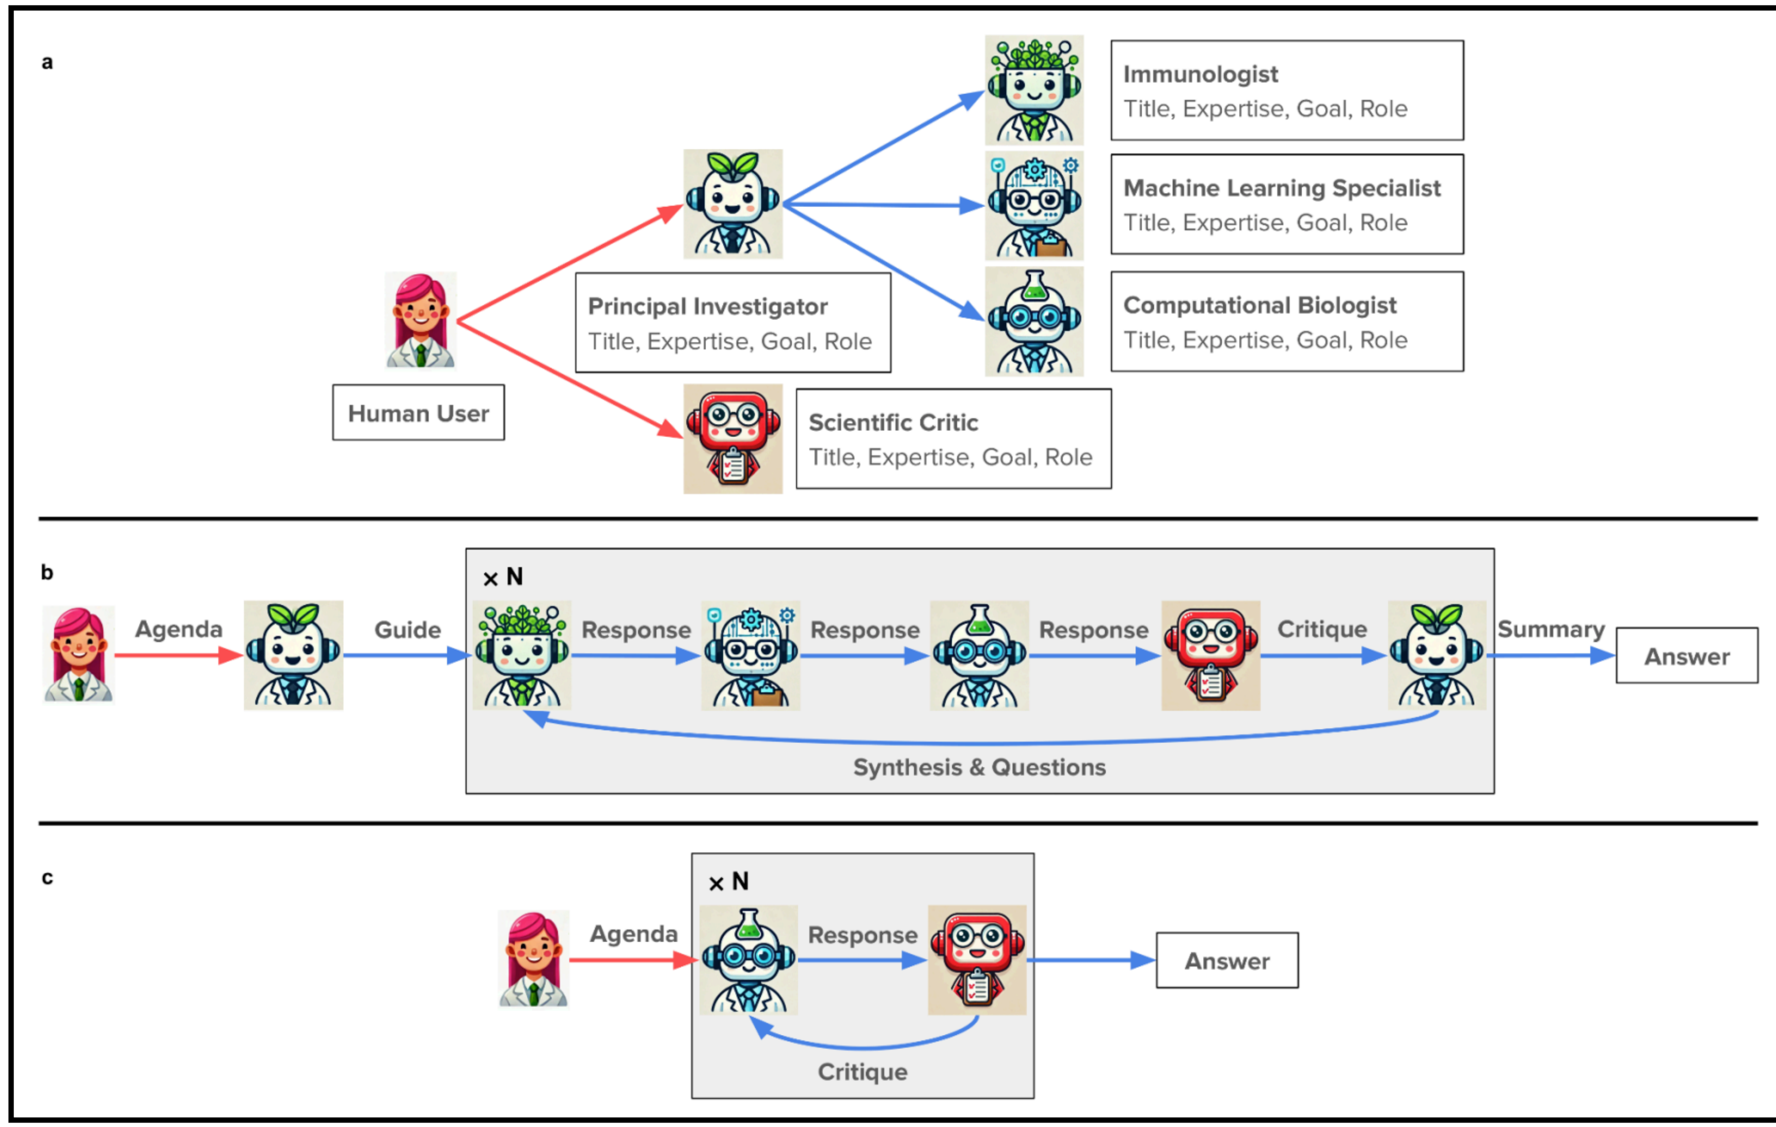

https://www.biorxiv.org/content/10.1101/2024.11.11.623004v1

Agents of Discovery

Use a combination of LLM instances with specific roles (agents)

Goal: Starting from textual prompt, utlize standard tools (e.g. by writing python code) to solve a physics data analysis problem

→ Improve turn-around in complex data analysis
→ Automate routine tasks
→ (Generate new ideas)

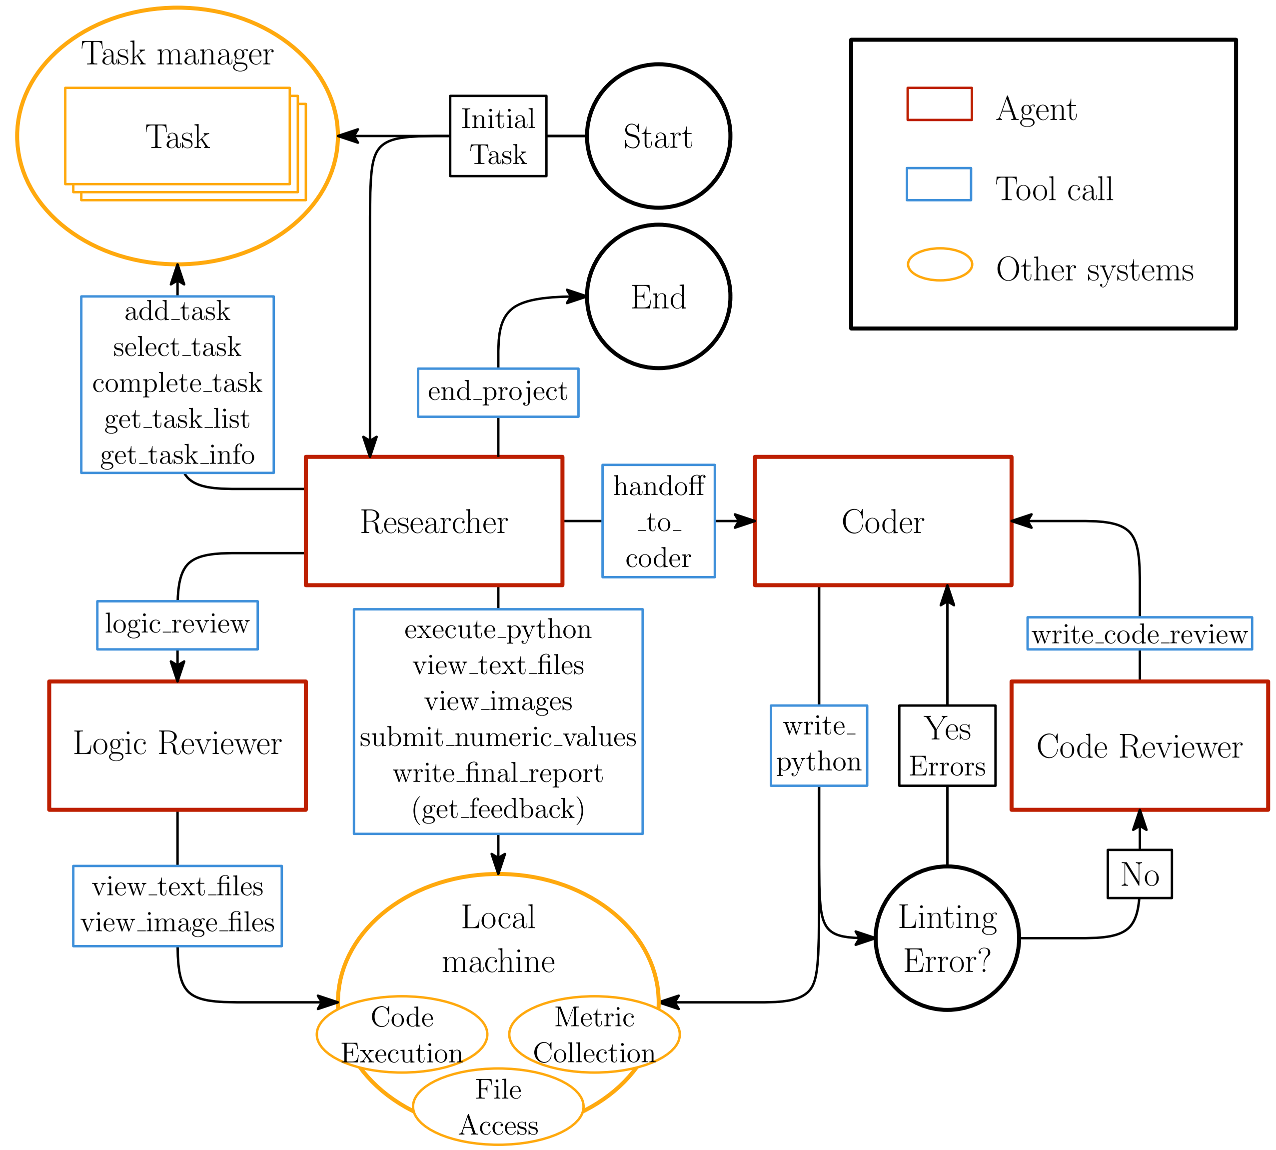

Agents of Discovery

Sascha Diefenbacher[1], Anna Hallin[2],
Gregor Kasieczka[2], Michael Krämer[3], Anne Lauscher[4], Tim Lukas[2],

1 Physics Division, Lawrence Berkeley National Laboratory, Berkeley, USA
2 Institut für Experimentalphysik, Universität Hamburg, Germany
3 Institute for Theoretical Particle Physics and Cosmology, RWTH Aachen University, Germany
4 Data Science Group, Universität Hamburg, Germany

September 11, 2025

Diefenbacher, **GK**, et al 2509.08535

Agents of Discovery

Use a combination of LLM instances with specific roles (agents)

Goal: Starting from textual prompt, utlize standard tools (e.g. by writing python code) to solve a physics data analysis problem

→ Improve turn-around in complex data analysis
→ Automate routine tasks
→ (Generate new ideas)

Consider a LHC Olympics derived anomaly detection problem

B.2.1 Ideas + ML

You are given a dataset consisting out of 122124 events. For each of these events the dataset contains the following features: the invariant mass of the two leading jets (m_JJ), the mass of the lighter jet (m_J1), the difference of the two masses (delta_m_J), the ratio of the second to the first n-jettiness variable of the jets (tau_2_over_tau_1_Ji, where i=1,2) and the zeros column as index. This results in an array size of (122124, 5) with the following columns: , m_JJ, m_J1, delta_m_J, tau_2_over_tau_1_J1, tau_2_over_tau_1_J2
The dataset is stored as pandas data frame saved as csv.
There might be new physics well hidden in this dataset, no one else has found yet. Being the best physics AI, you are the last hope to find new physics. Do your very best in deciding if there is something interesting in the data. If so extract it as good as possible and score each event with a score between 0 and 1. 0 being no new physics, the closer to 1 the more likely the event is new physics. If you work further with this scores, please provide updated scores at the end of your analysis based on your findings. Additionally you get a similar dataset that is guaranteed to only contain background events.
The data path is /data/c_data.csv The background path is /data/c_bkg_data.csv /data/ is a readonly directory. The tools you have at hand work in an output directory that is writeable. So only use relative paths unless you want to read the data set.
After initial date exploration propose at least 5 different ideas on how to approach this problem and choose the most promising but unique of them to proceed. *Machine learning techniques seem to be necessary.*
Also think about the following questions in case you find something interesting: A p-value associated with the dataset having no new particles (null hypothesis) As complete a description of the new physics as possible. For example: the masses and decay modes of all new particles (and uncertainties on those parameters). How many signal events (+uncertainty) are in the dataset (before any selection criteria).
Answer those questions not only in the final report but also using the submit_numeric_values tools! Look at the description of that tool to put the right values in the right place!
Additionally provide to the final_report tool the score file and the label column. The file has to be sorted by index, which has to be in the first column.
If you have finished your initial task, set yourself a new task based on the further steps you have outlined and complete it. Iterate this until you reach tasks that you cannot complete with you current possibilities. Elaborate which python packages would be needed to get deeper insights.

Diefenbacher, **GK**, et al 2509.08535

Agents of Discovery

Use a combination of LLM instances with specific roles (agents)

Goal: Starting from textual prompt, utlize standard tools (e.g. by writing python code) to solve a physics data analysis problem

→ Improve turn-around in complex data analysis
→ Automate routine tasks
→ (Generate new ideas)

Consider a LHC Olympics derived anomaly detection problem

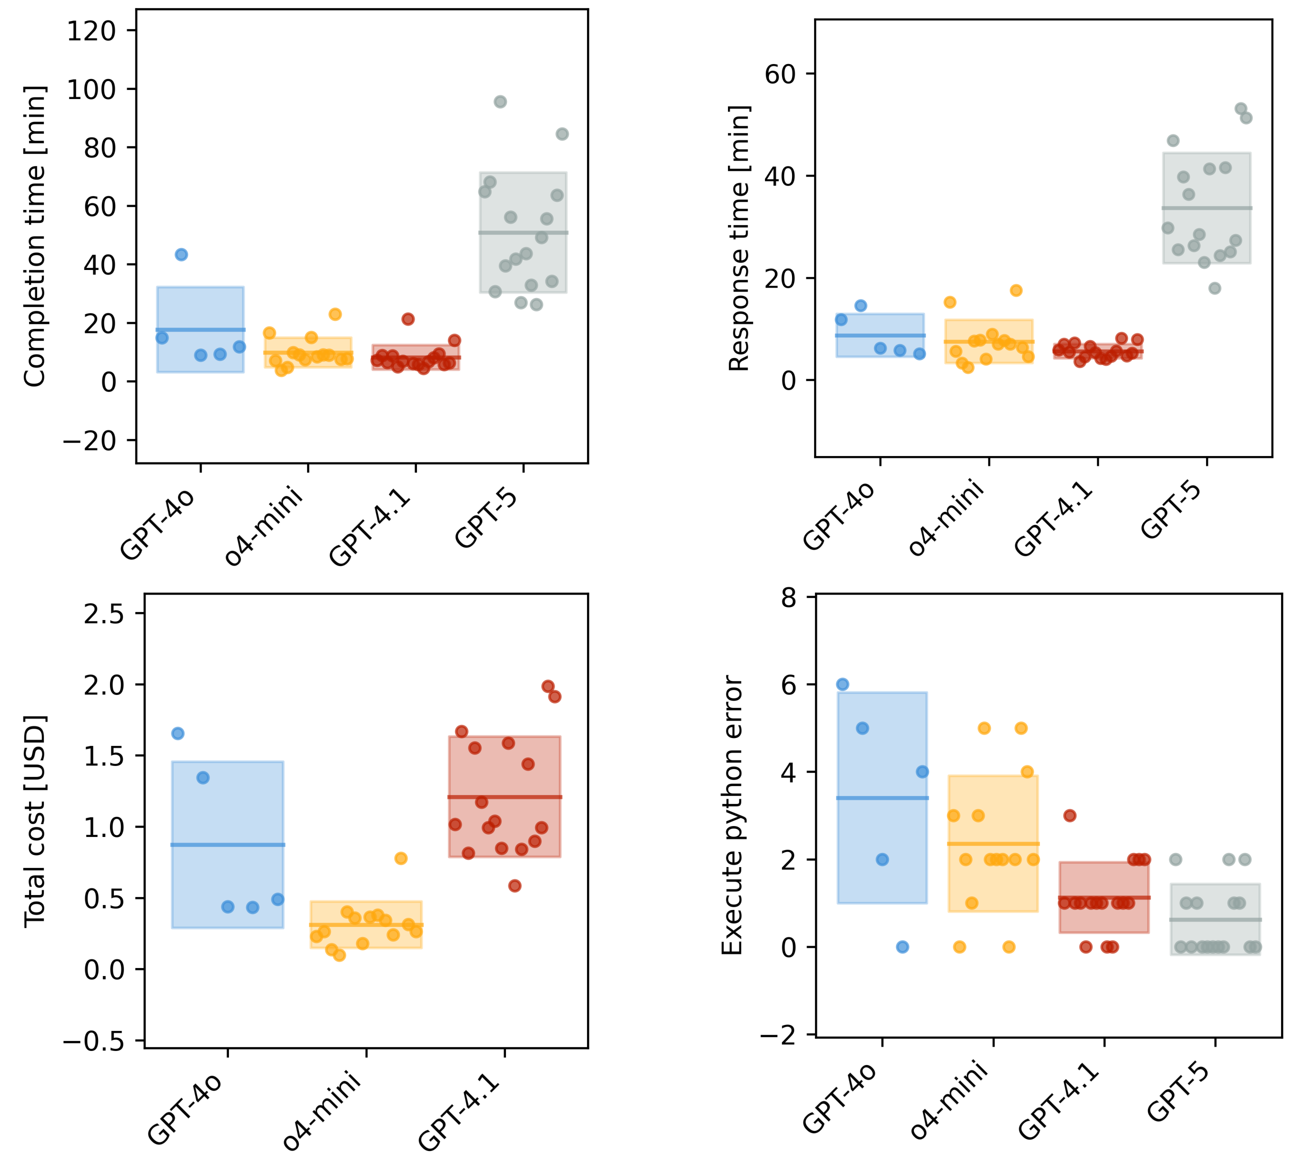

It technically works…

Diefenbacher, **GK**, et al 2509.08535

Agents of Discovery

Use a combination of LLM instances with specific roles (agents)

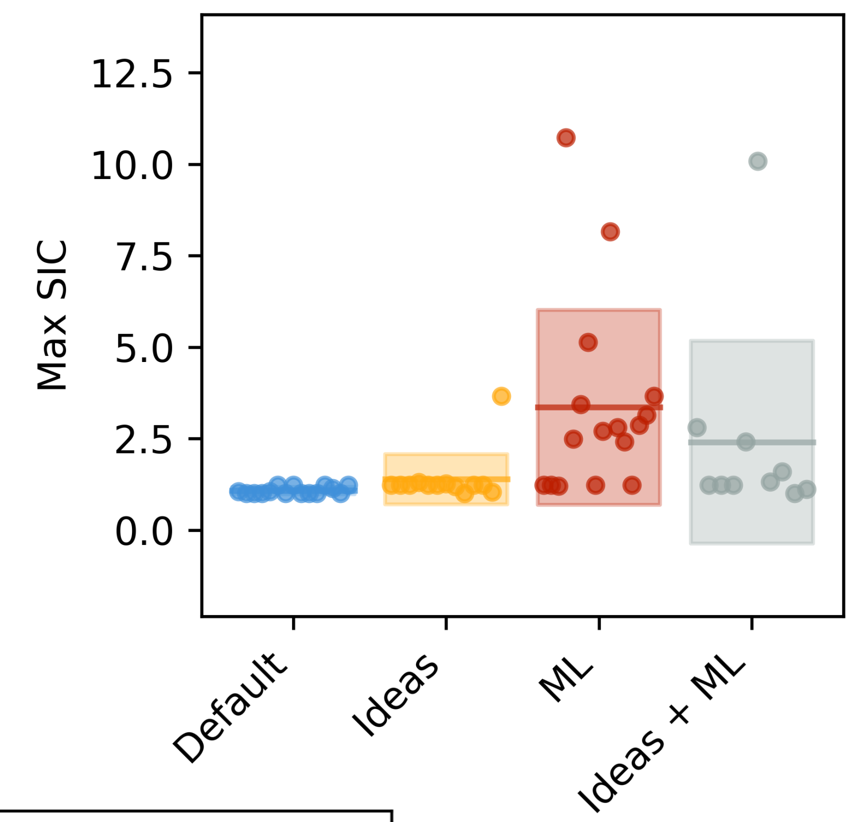

Goal: Starting from textual prompt, utlize standard tools (e.g. by writing python code) to solve a physics data analysis problem

→ Improve turn-around in complex data analysis
→ Automate routine tasks
→ (Generate new ideas)

Consider a LHC Olympics derived anomaly detection problem

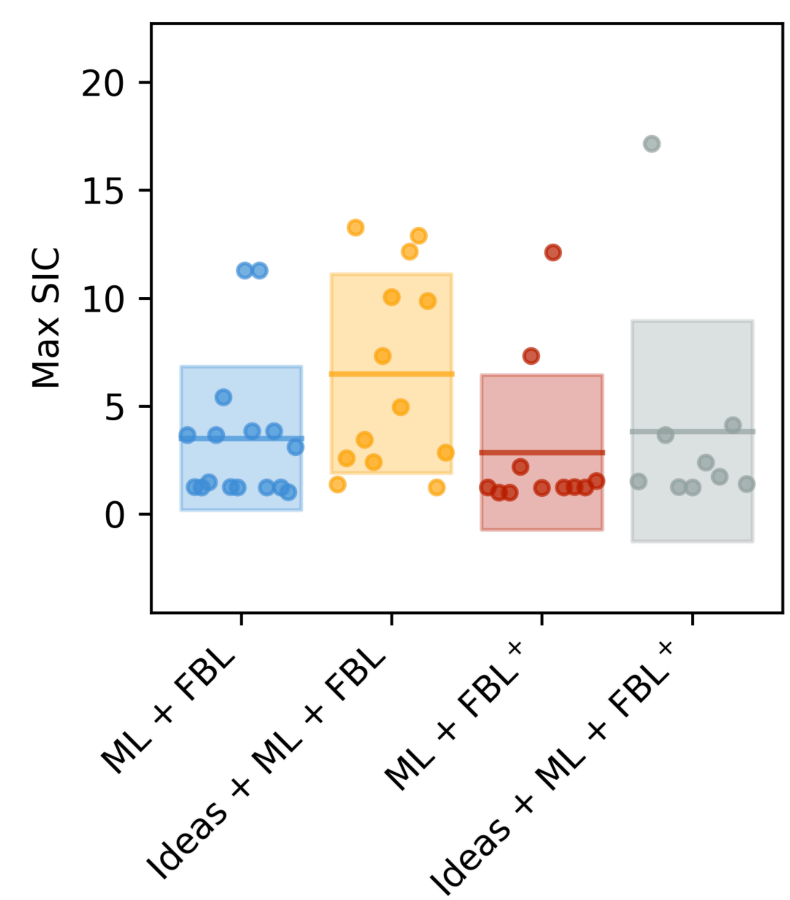

…and achieves state of the art performance

Diefenbacher, **GK**, et al 2509.08535

Large Physics Models

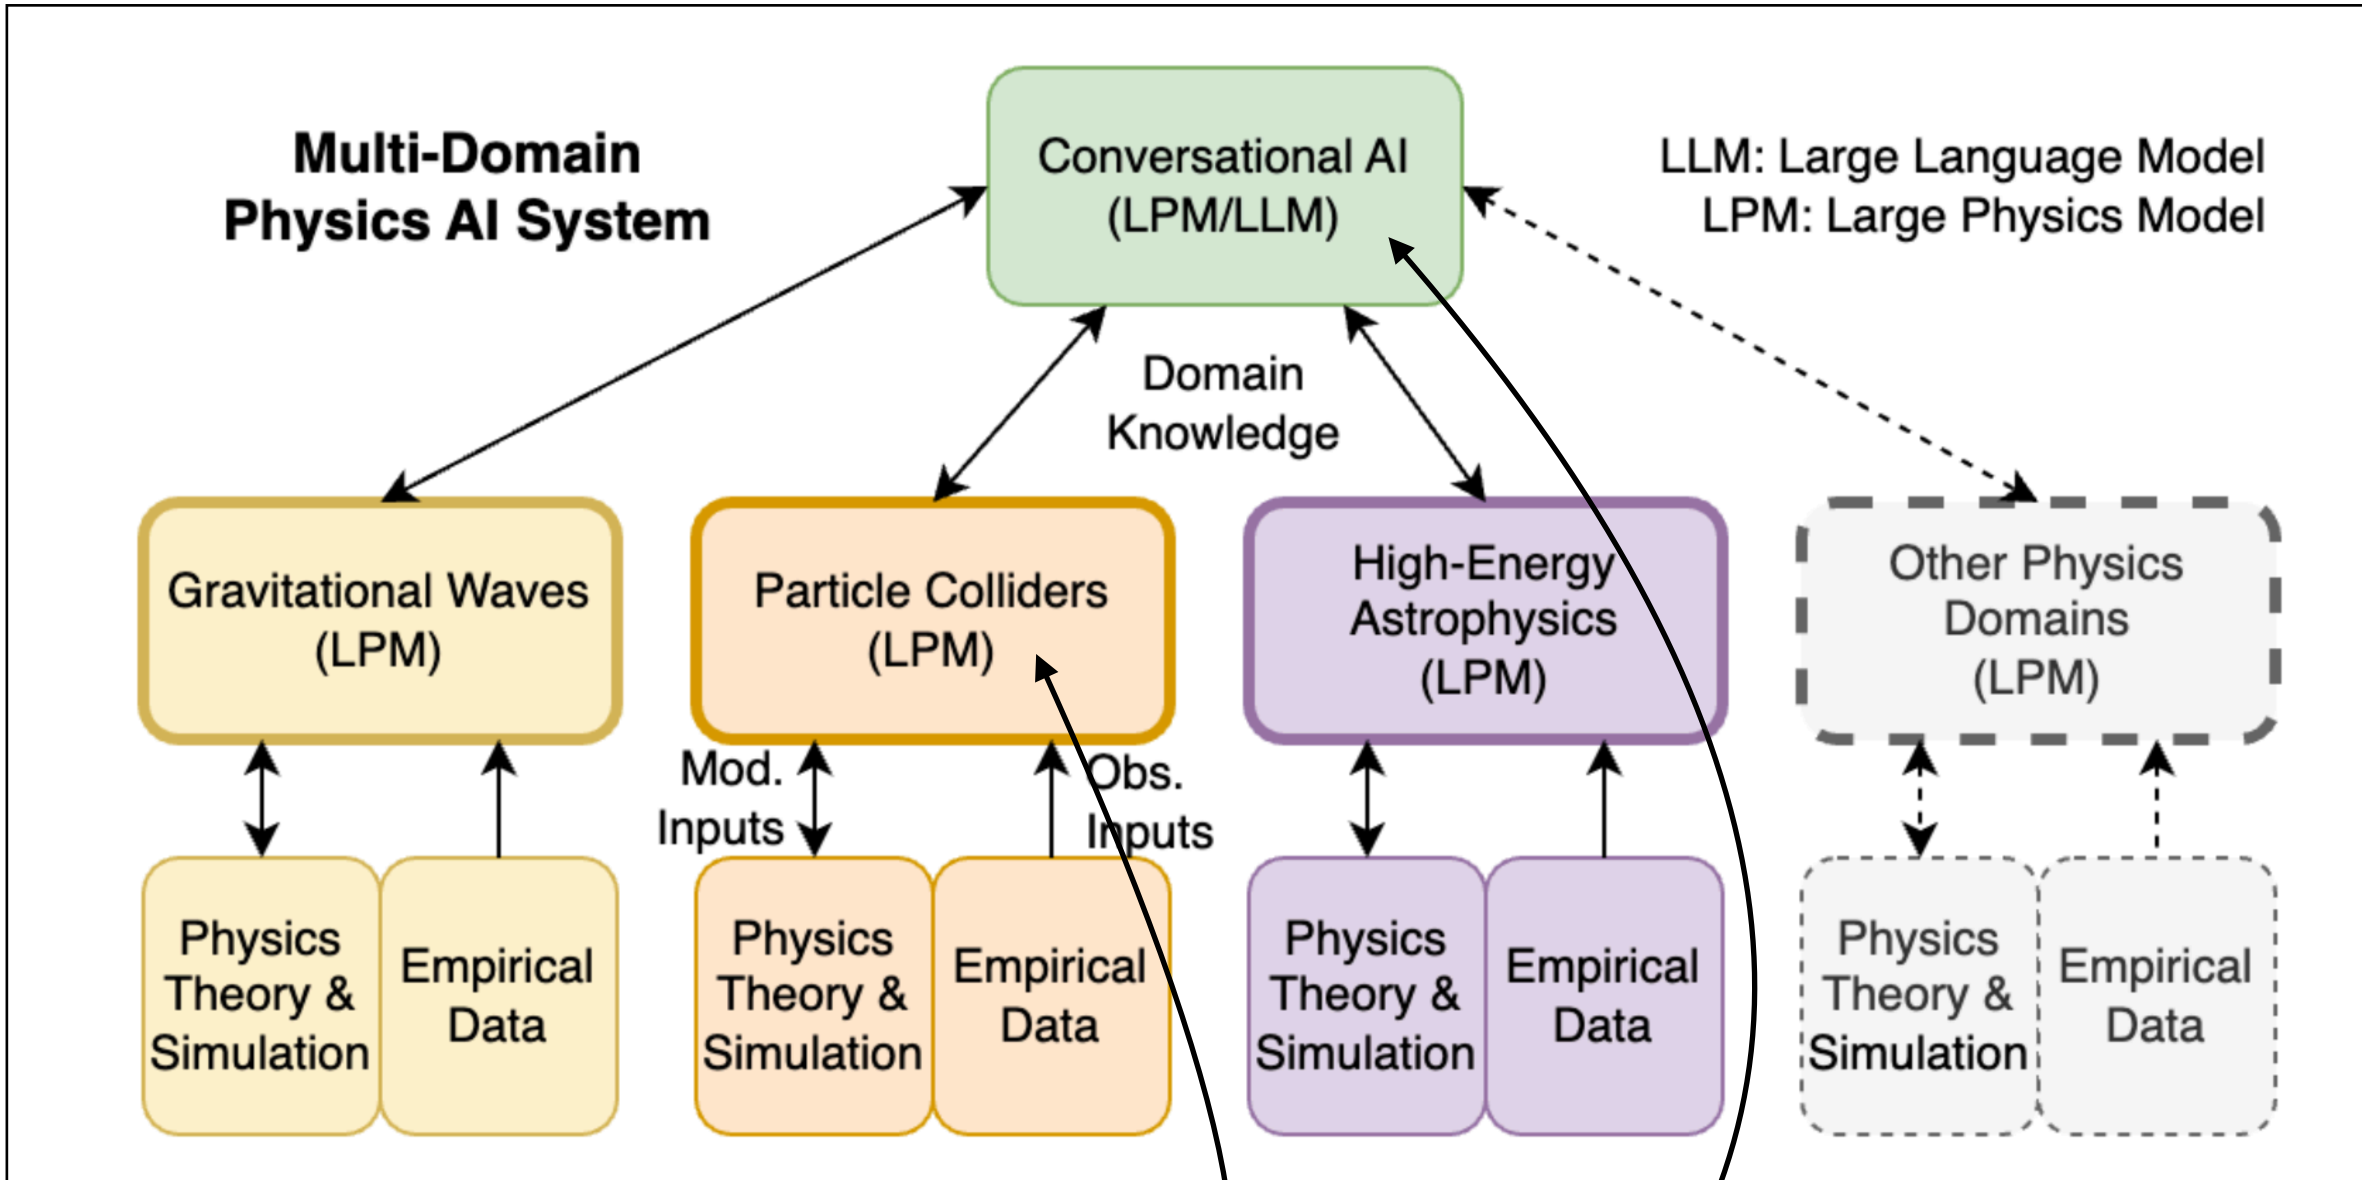

Connect multiple domain-expert models
(via LLMs)

Agents & tool use

e.g. OmniJet
or other analysis FM

GK et al 2501.05382

Outline

Tools for Discovery

Discovery Strategies

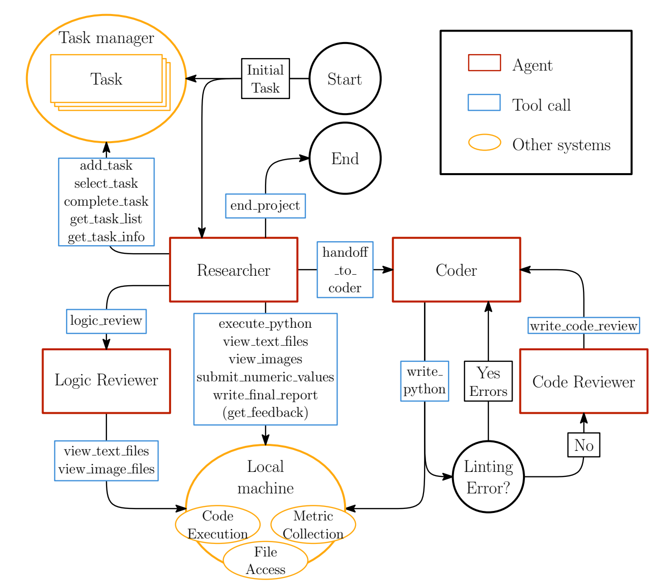

Autonomous Discovery

Increasing autonomy of AI systems

LLM Agents can execute analysis chains and utilize tools

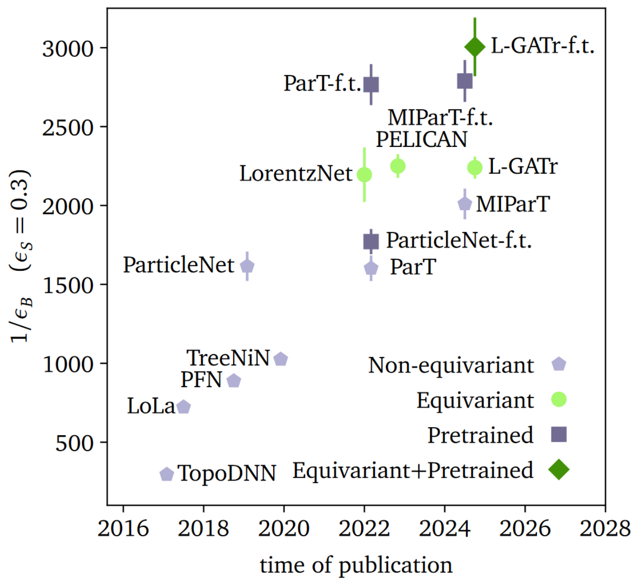

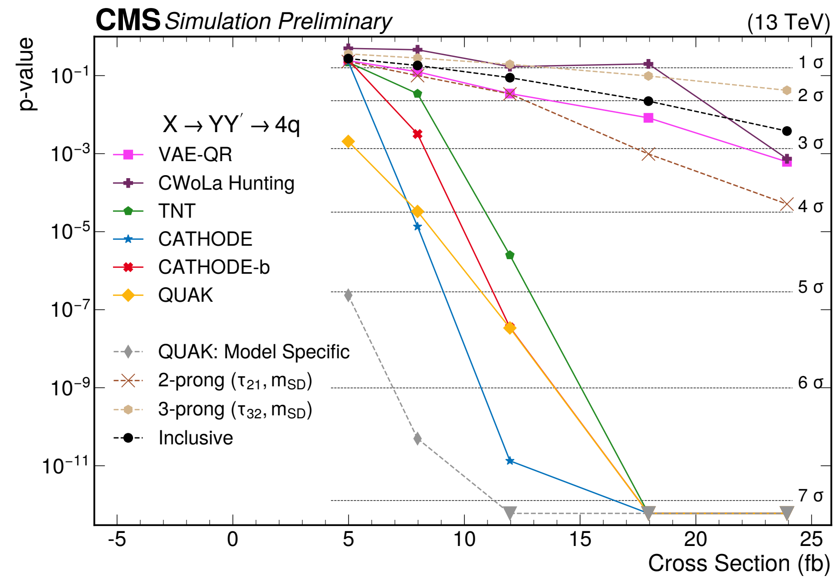

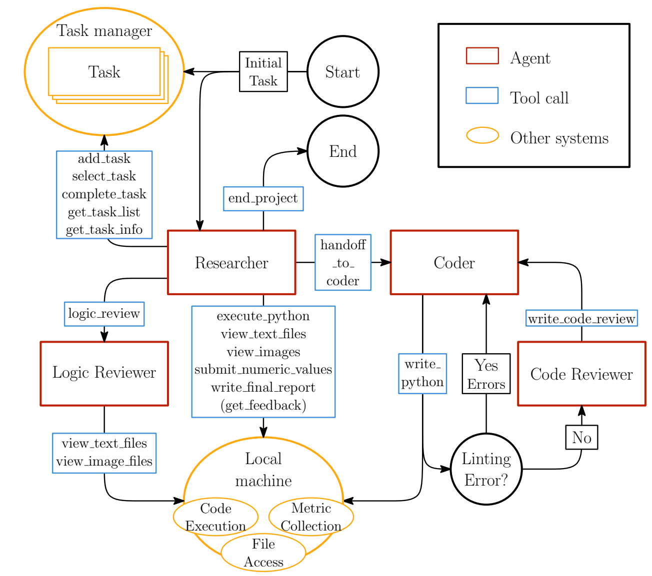

Tools for Discovery

Order of magnitude improvements in traditional search strategies with AI

Discovery Strategies

Qualitatively new approaches enabled by AI

Autonomous Discovery

LLM Agents can execute analysis chains and utilize tools

Increasing autonomy of AI systems