

# Painting pretty pictures of particles: Generative AI for detector simulation and new physics discovery

Gregor Kasieczka

Email: [gregor.kasieczka@uni-hamburg.de](mailto:gregor.kasieczka@uni-hamburg.de)

Twitter: [@GregorKasieczka](https://twitter.com/GregorKasieczka)

LPSC Grenoble Seminar

15.6.2023

**CLUSTER OF EXCELLENCE**  
QUANTUM UNIVERSE

**UH**  
  
Universität Hamburg  
DER FORSCHUNG | DER LEHRE | DER BILDUNG

  
KISS

**CDCS**  
CENTER FOR DATA AND COMPUTING  
IN NATURAL SCIENCES

  
PUNCH  
4 NFDI

**DASHH**

  
PIER  
Partnership of  
Universität Hamburg and DESY

Emmy  
Noether-  
Programm  
Deutsche  
Forschungsgemeinschaft  
**DFG**  


  
Bundesministerium  
für Bildung  
und Forschung





Please list in a concise way the three key developments expected for particle physics in the next decade



As an AI, I don't predict the future, but as of my last update in September 2021, the following are some key developments expected in the field of particle physics over the next decade based on the direction of ongoing research and proposed experiments:

1. **Completion of the High-Luminosity Large Hadron Collider (HL-LHC):** This upgrade to CERN's Large Hadron Collider (LHC) will greatly increase the collider's luminosity, resulting in much more data for physicists to analyze. The HL-LHC could help to answer questions about the Higgs boson, dark matter, and supersymmetry. The project started in 2019 and is expected to be completed by the middle of this decade.
2. **Progress in Neutrino Physics:** Several experiments around the world, like the Deep Underground Neutrino Experiment (DUNE) in the U.S. and the Hyper-Kamiokande in Japan, aim to learn more about the properties of neutrinos, particularly their mass hierarchy and charge conjugation parity (CP) violation. If neutrinos are found to violate CP symmetry, it could provide a clue as to why the universe is made mostly of matter and not antimatter.
3. **Advancements in Dark Matter Detection:** Dark matter is one of the most pressing unsolved mysteries in physics. Experiments like the Xenon1T and its successor the Xenon-nT in Italy, and the Large Underground Xenon (LUX) and its upgrade, LUX-ZEPLIN (LZ) in the U.S., aim to directly detect dark matter particles. The next decade may see a breakthrough in our understanding of this elusive form of matter.

It's important to note that while these are major expected developments, the field of particle physics is vast and inherently unpredictable. Unexpected discoveries could lead to entirely new lines of inquiry.

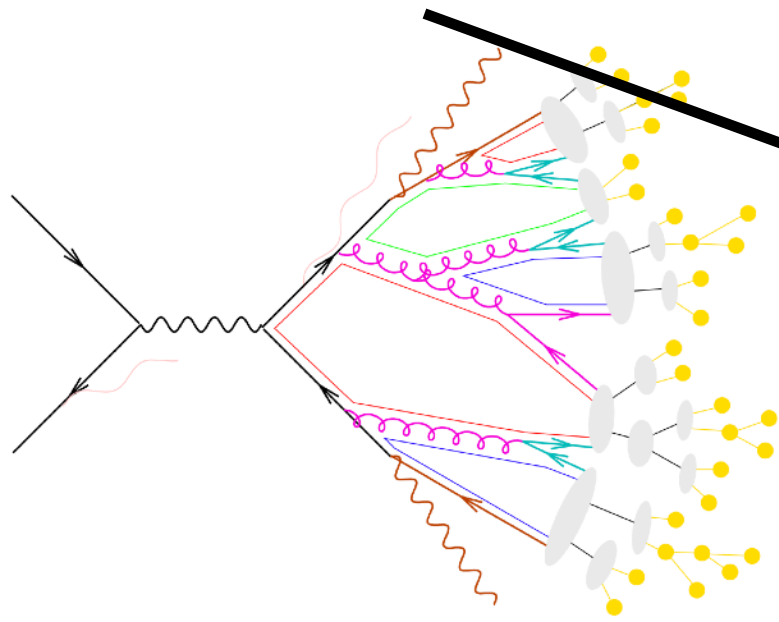




a baroque oil painting of joseph fourier as DJ in a dance club



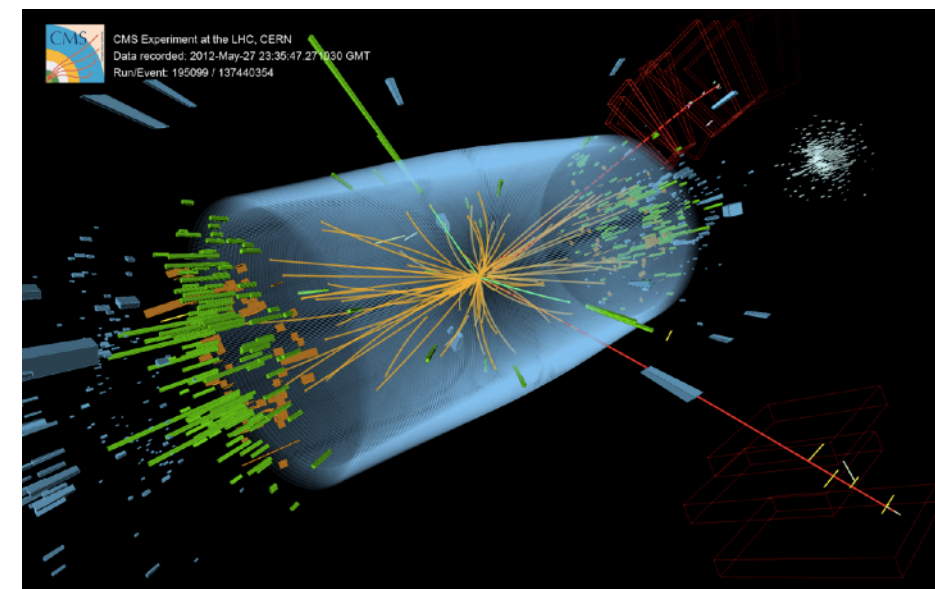
# Experimental particle physics workflow



This is what happens in the experiment



This is what we want to know



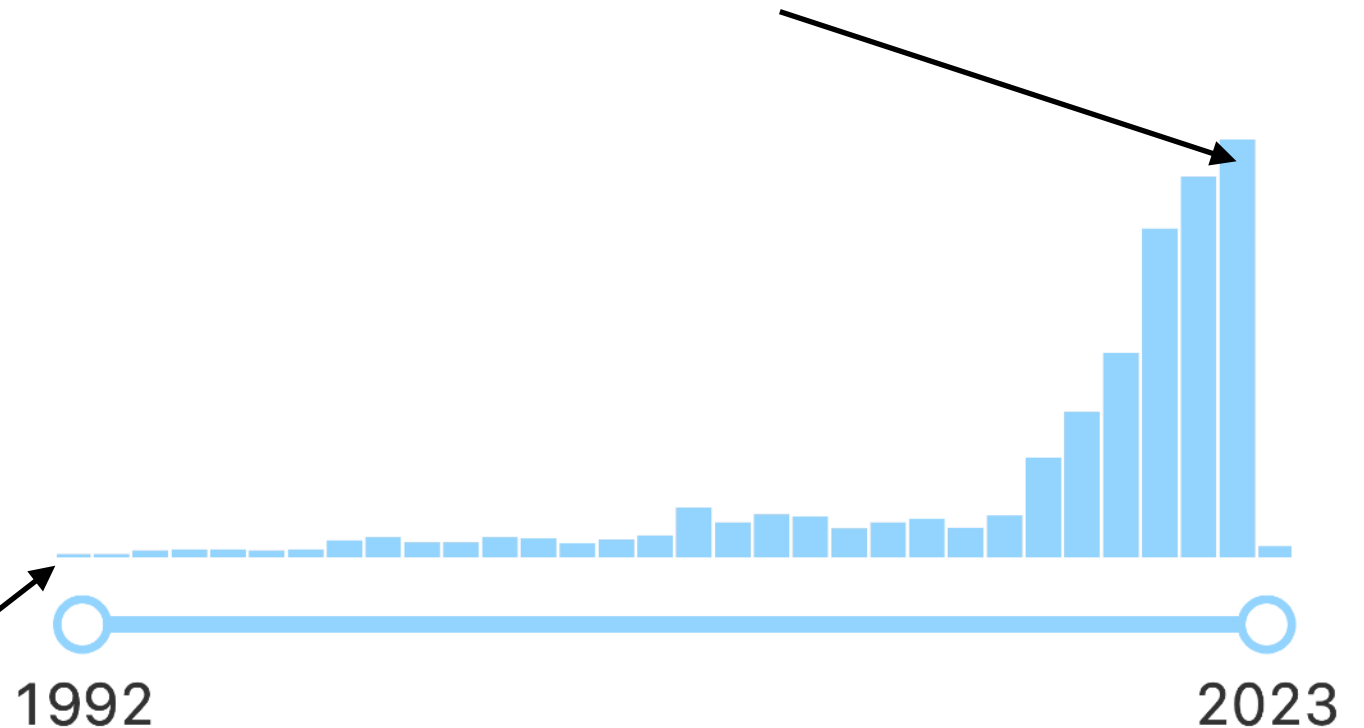
Connect observational data  
with underlying theory:  
Statistics & simulation



# The rise of ML in particle physics

- Machine learning (ML) considered in all aspects of particle physics
- Rapid rise in publications and considered topics

420 papers in 2022



## THE USE OF NEURAL NETWORKS IN HIGH ENERGY PHYSICS\*

BRUCE DENBY  
*Fermi National Accelerator Laboratory*  
M.S. 318  
Batavia, Illinois 60510 U.S.A.  
denby@fnal.bitnet

### ABSTRACT

In the past few years a wide variety of applications of neural networks to pattern recognition in experimental high energy physics has appeared. The neural network solutions are in general of high quality, and, in a number of cases, are superior to those obtained using 'traditional' methods. But neural networks are of particular interest in high energy physics for another reason as well: much of the pattern recognition must be performed online, i.e., in a few microseconds or less. The inherent parallelism of neural network algorithms, and the ability to implement them as very fast hardware devices, may make them an ideal technology for this application.

Inspire: ("machine learning" or "deep learning" or neural) and (hep-ex or hep-ph or hep-th)



# The rise of ML in particle physics

- Machine learning (ML) considered in all aspects of particle physics
- Rapid rise in publications and considered topics
- **To maximise impact:**  
Solid understanding of physics and machine learning techniques required
  - Domain knowledge of the physics data — e.g. symmetries
  - Phrase the actual physics task tractable to direct optimisation

See e.g. for a broader overview:  
<https://arxiv.org/abs/2112.03769>

Focus on LHC-style analysis but underlying ideas applicable in general

The image shows a screenshot of a review article from the journal 'Physics Reports'. The title is 'Machine learning in the search for new fundamental physics'. The authors listed are Georgia Karagiorgi, Gregor Kasieczka, Scott Kravitz, Benjamin Nachman, and David Shih. The abstract discusses the use of machine learning in high-energy particle physics, mentioning the Large Hadron Collider and neutrino experiments. The article is part of a series of reviews in the journal, as indicated by the 'REVIEWS' label in the top right corner. The page number is 399, and the date is June 2022.

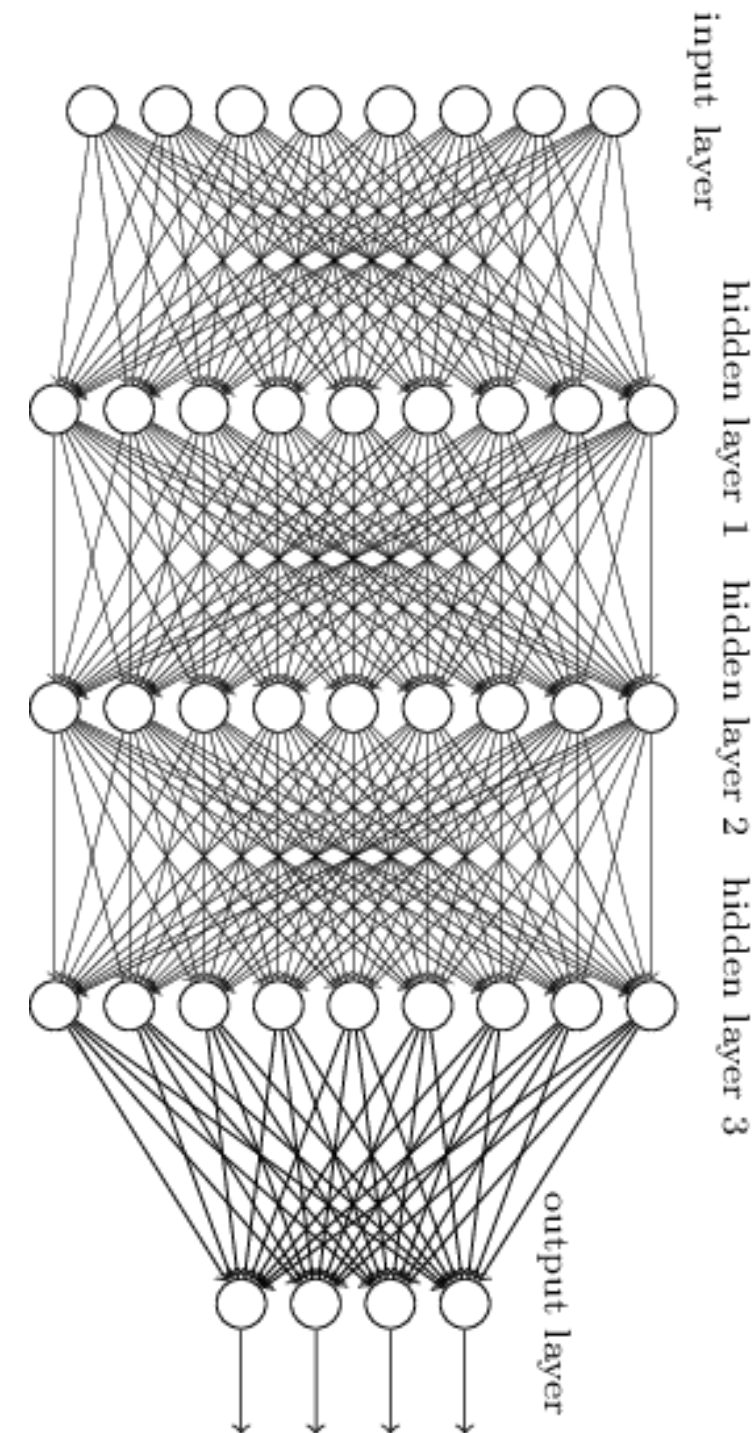


# Micro-Intro: Machine Learning

- Rephrase task as a minimisation problem..
- ..and “simply” solve:

$$\theta^* = \operatorname{argmin}_{\theta} \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [\mathcal{L}(f_{\theta}(\mathbf{x}), \mathbf{x})]$$

- Modern ML: *function  $f$  is a deep neural network & minimisation carried out via gradient descent*
- Devil in the details:
  - How to map physics objective to loss function  $L$
  - How to structure  $f$  to make maximum use of physics knowledge
  - How learn in a robust way from minimum amount of data
  - ...



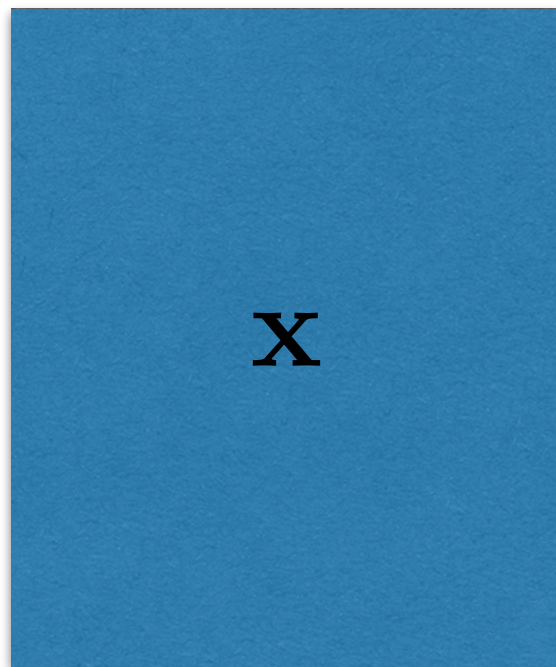
# Loss function: Supervised

## Supervised Learning:

Attempt to infer some target (*truth label*):

classification, regression (often also clustering/inference)

Use training data with known labels  
(often from Monte Carlo simulation)



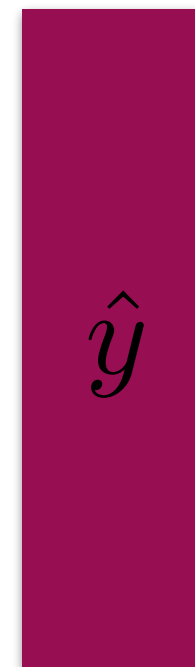
observable features  
such as kinematics,  
tracks,...



truth label  
(e.g. true energy)

Learn to predict:

$$\hat{y} = f_{\theta}(\mathbf{x})$$



predicted energy

**Regression: Minimize mean squared error:**

$$\mathcal{L} = (y - \hat{y})^2$$

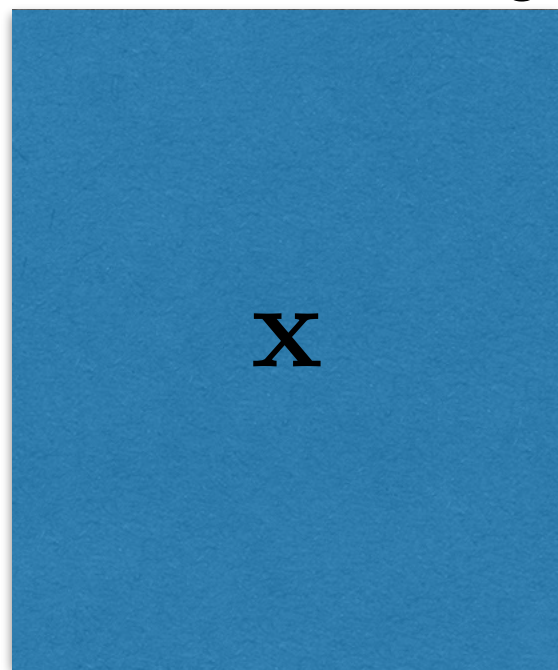


# Loss function: Unsupervised

## Unsupervised Learning:

No target, learn the probability distribution (directly from data)

Can use for sampling, anomaly detection, unfolding, ...



Learn to predict:

$$\hat{p}(\mathbf{x}) = f_{\theta}(\mathbf{x})$$



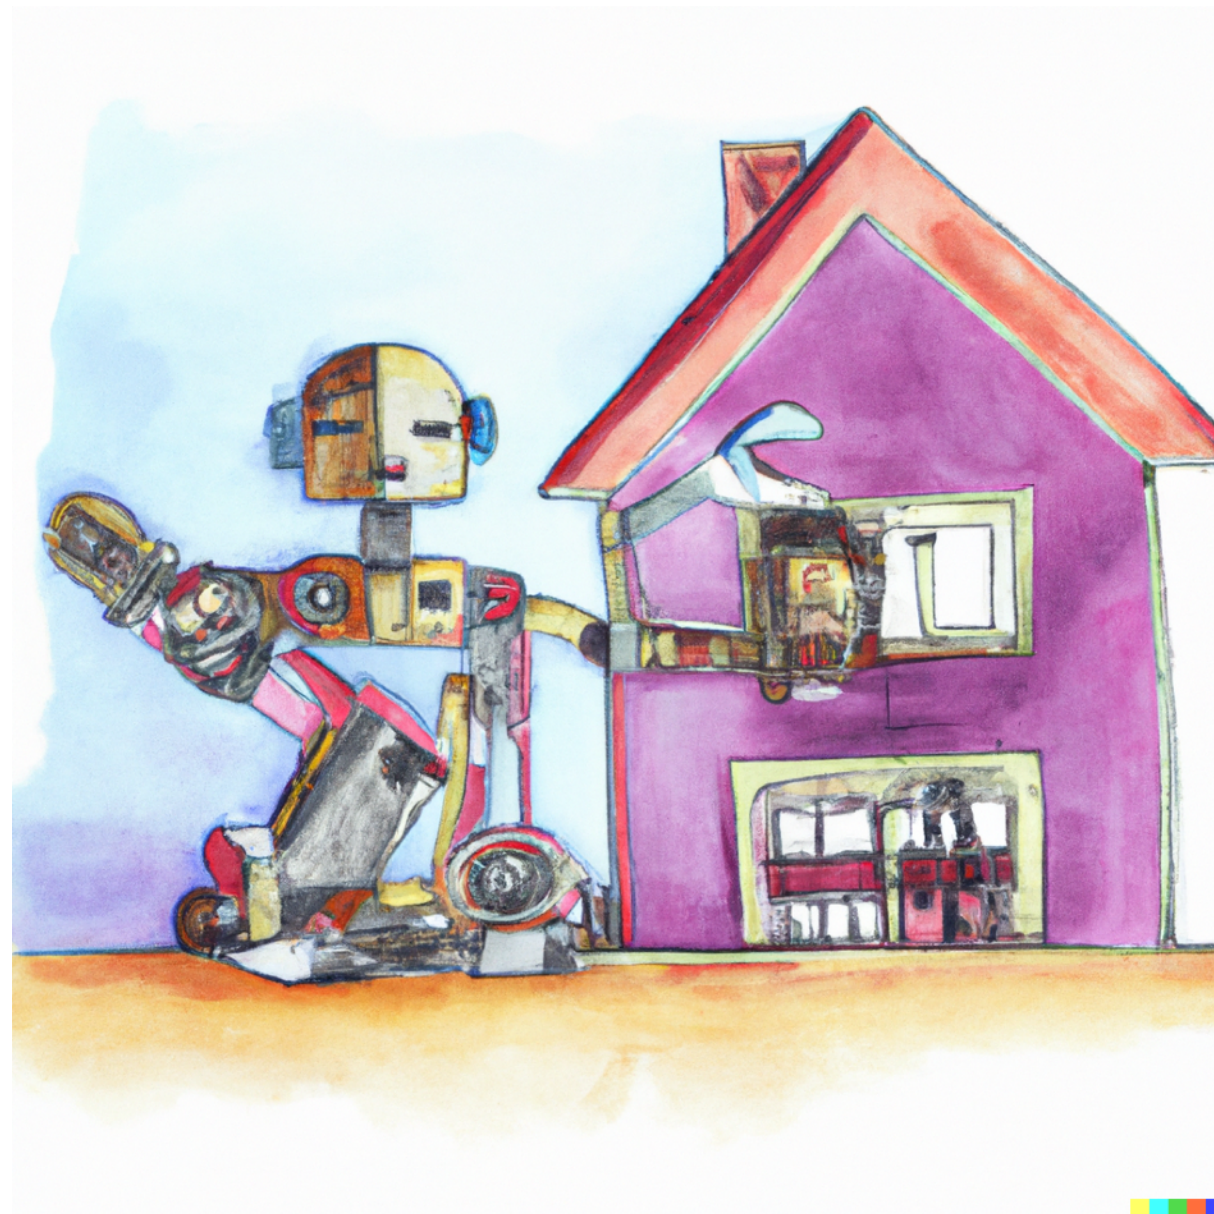
$$p(\mathbf{x})$$

True probability density

**Distribution learning: Maximise likelihood (minimize log-likelihood):**  
(either directly or with approximations)

$$\mathcal{L} = -\log(\hat{p}(\mathbf{x}))$$

# Fast simulation



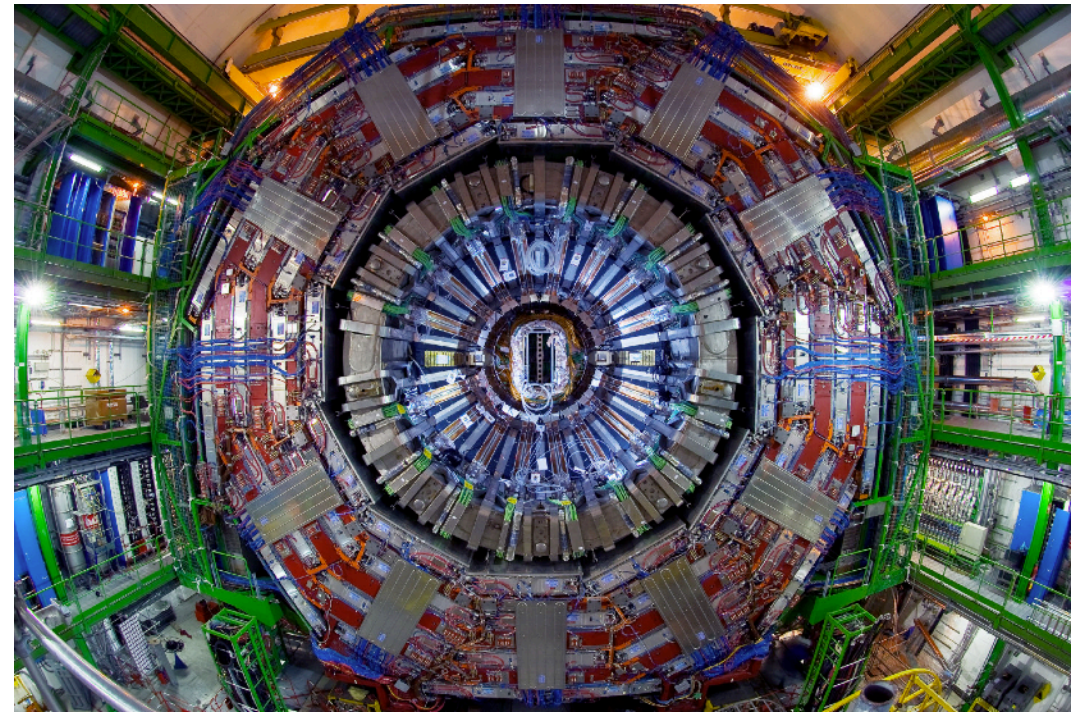


# Generative Models

Simulation is crucial to connect experimental data with theory predictions

$$\begin{aligned}\mathcal{L} = & -\frac{1}{4} F_{\mu\nu} F^{\mu\nu} \\ & + i\bar{\psi} \not{D} \psi + h.c. \\ & + \chi_i Y_{ij} \chi_j \phi + h.c. \\ & + |D_m \phi|^2 - V(\phi)\end{aligned}$$

Complex chain  
of  
simulations



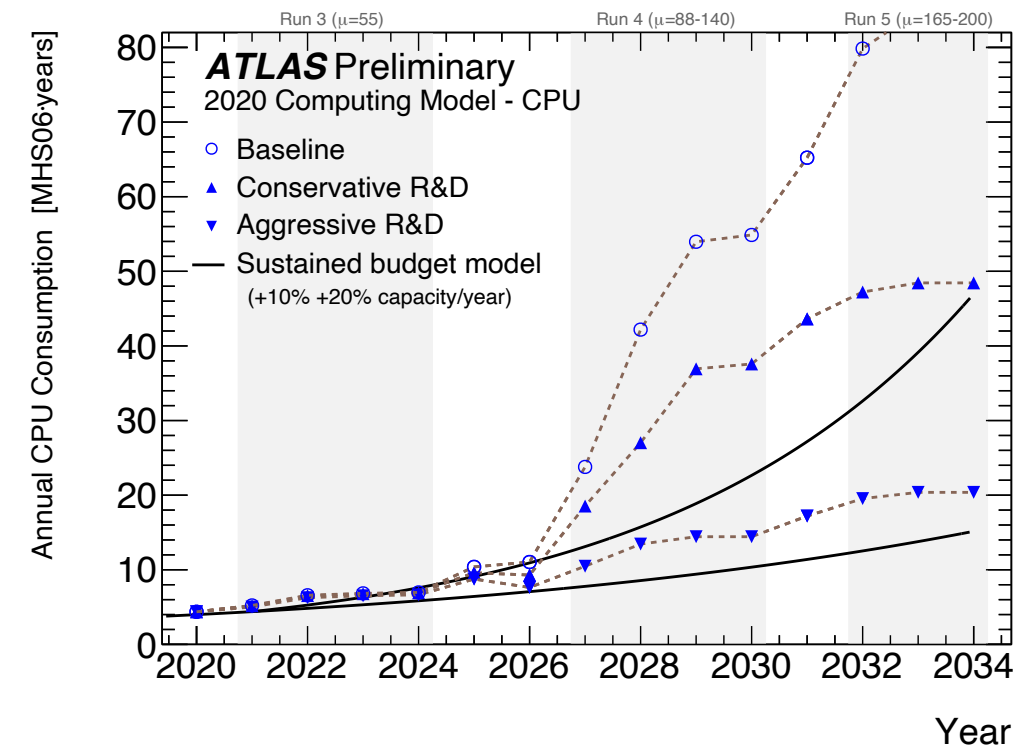
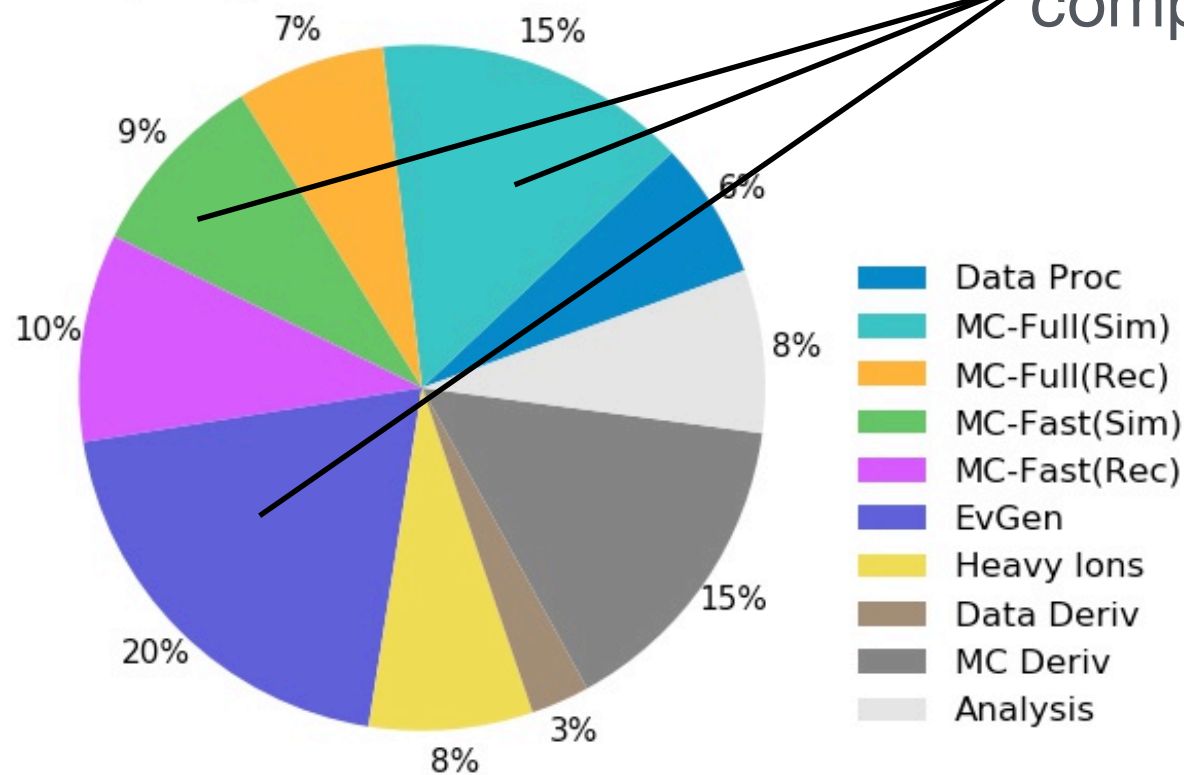
# Generative Models

Simulation is crucial to connect experimental data with theory predictions  
**but computationally very expensive**

Simulation and Generation steps over 40% of ATLAS compute effort..

..and projected to exceed available resources.

**ATLAS Preliminary**  
 2020 Computing Model -CPU: 2030: Baseline

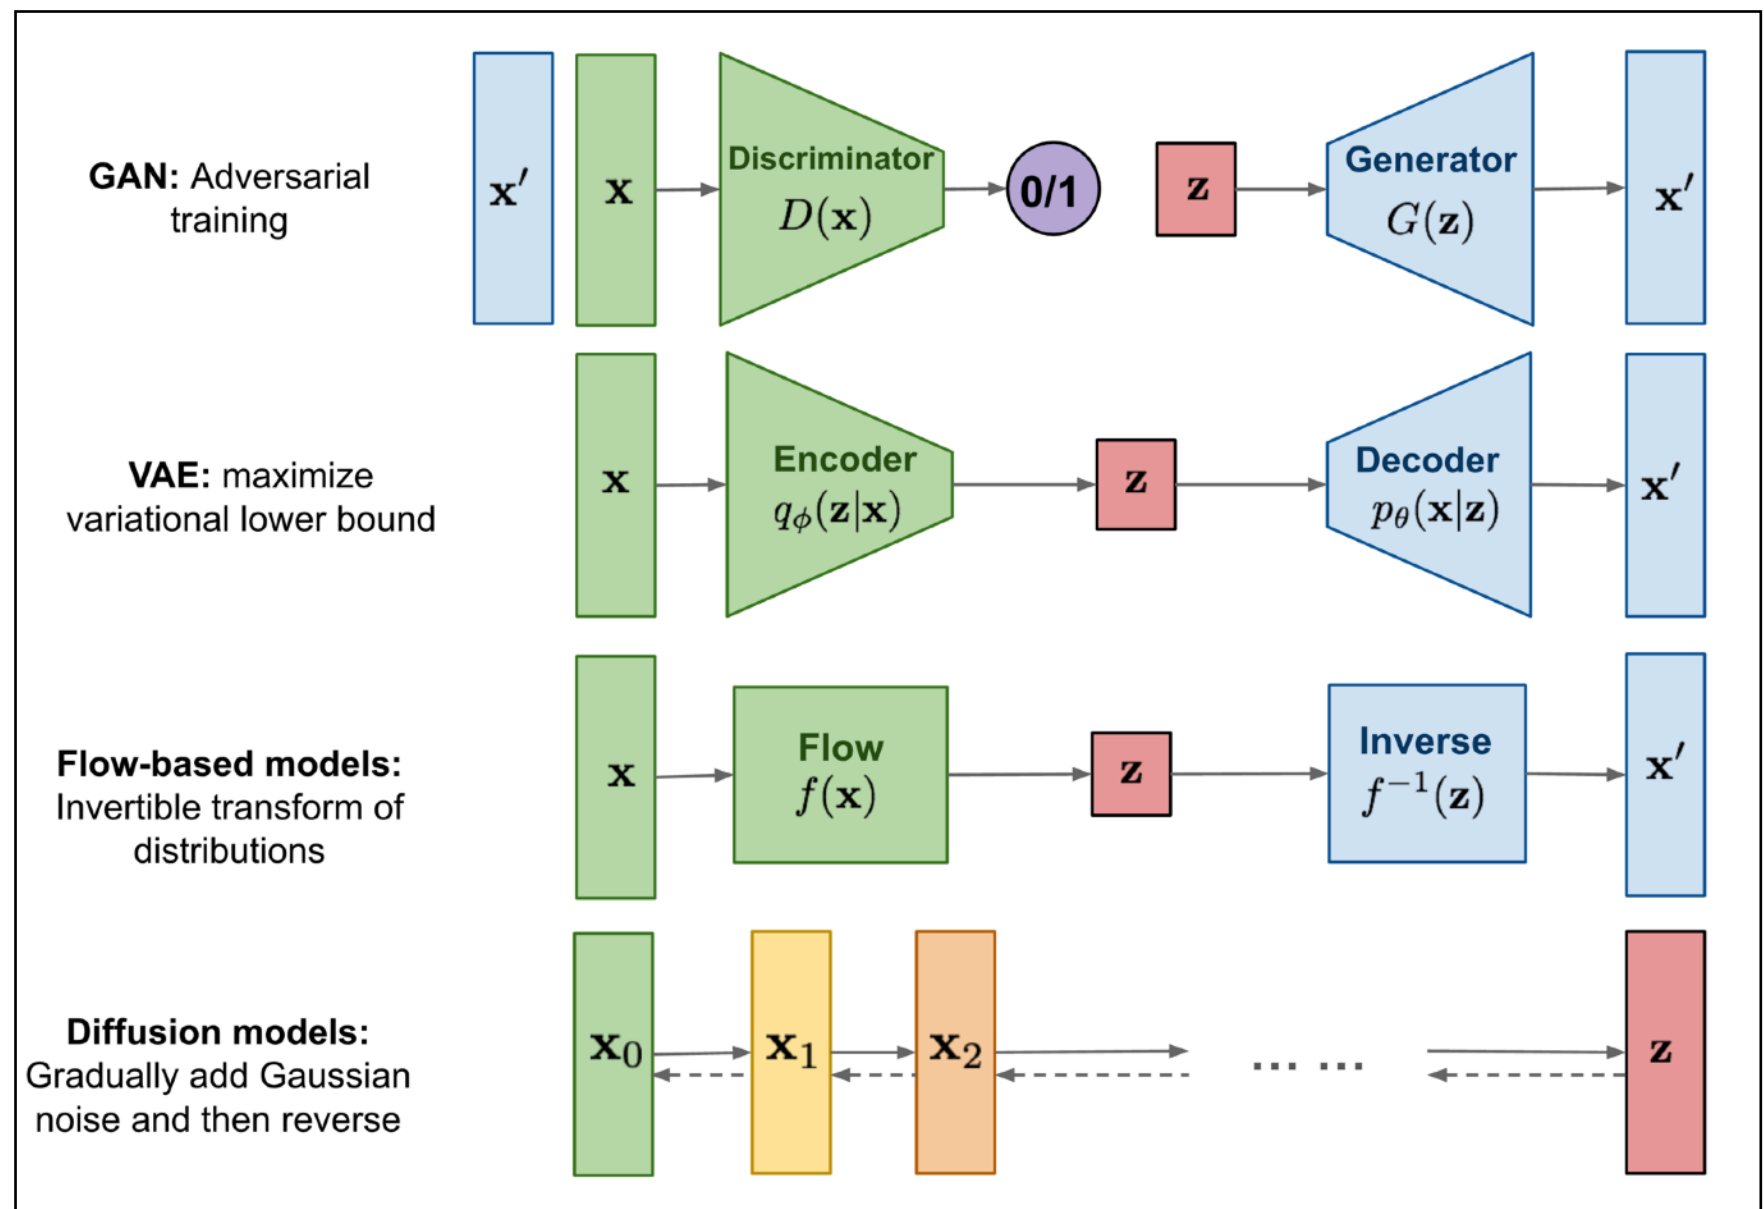
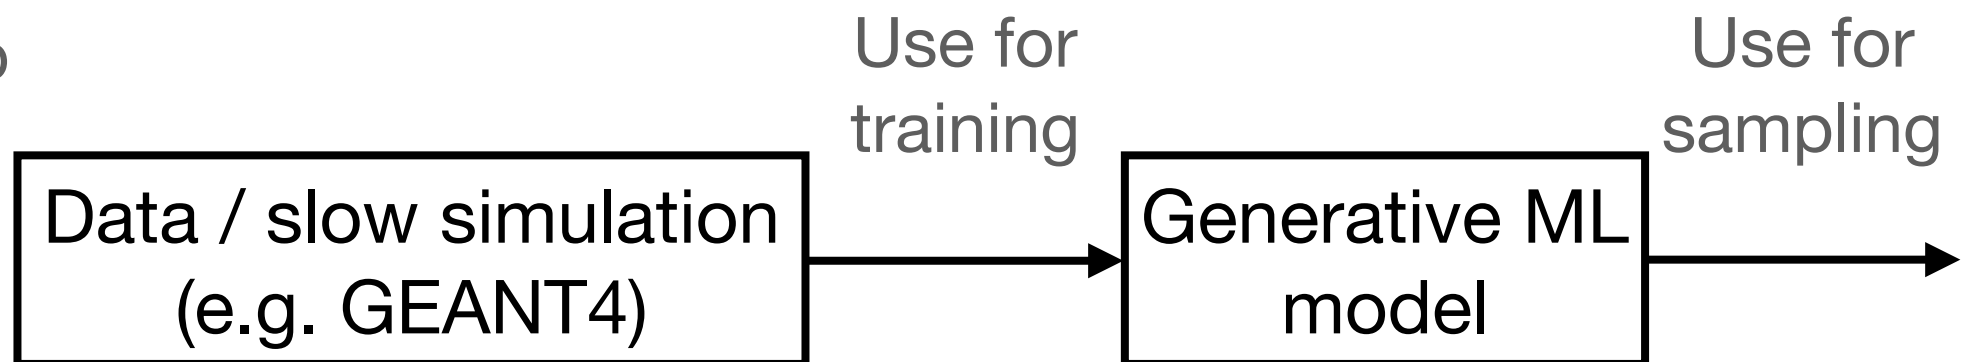




# Generative Models

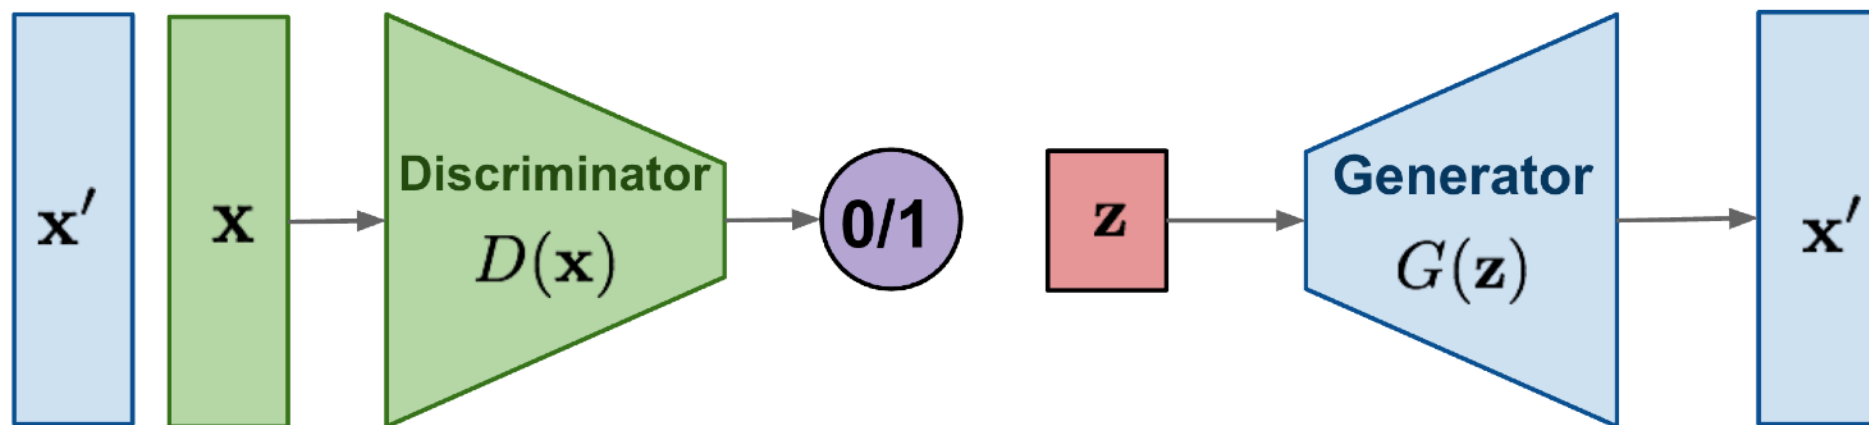
Simulation is crucial to connect experimental data with theory predictions  
**but computationally very expensive**

**Use generative models trained on initial data to augment statistics**

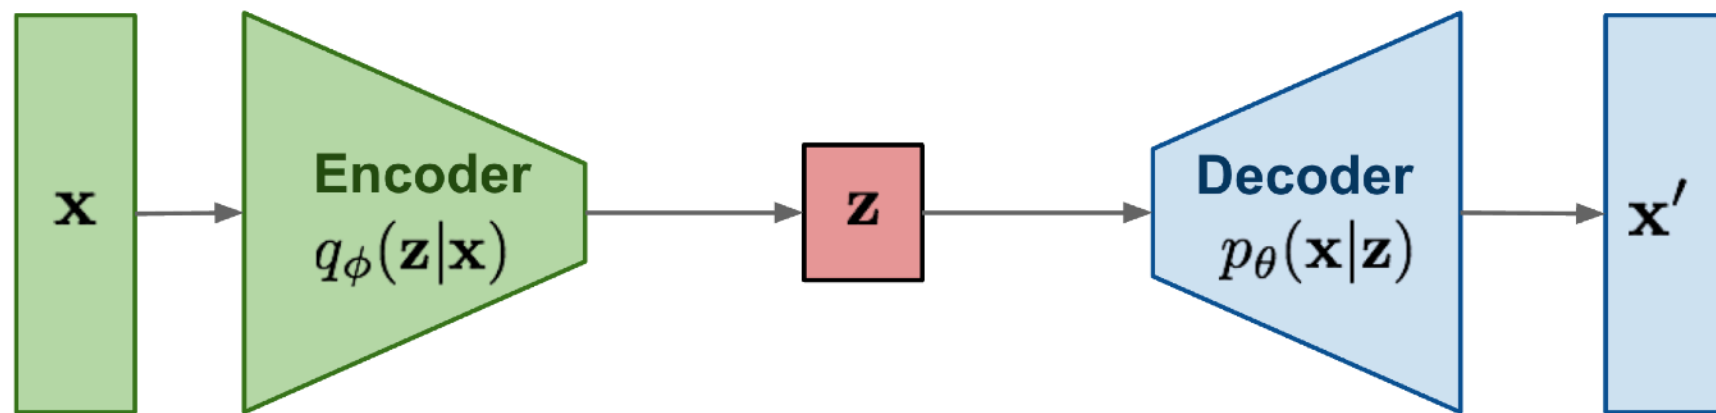




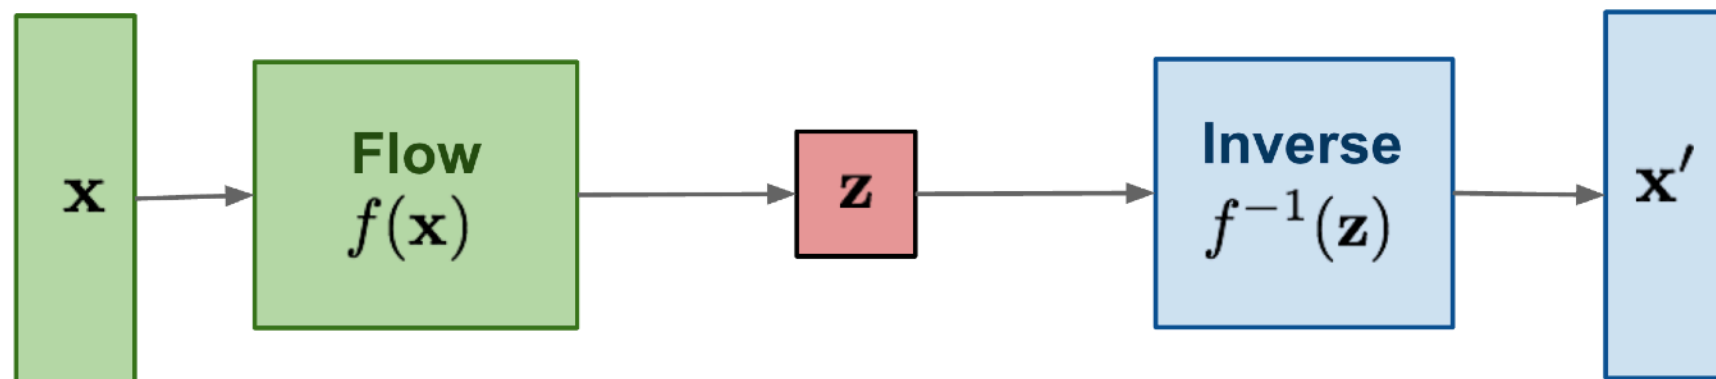
**GAN:** Adversarial training



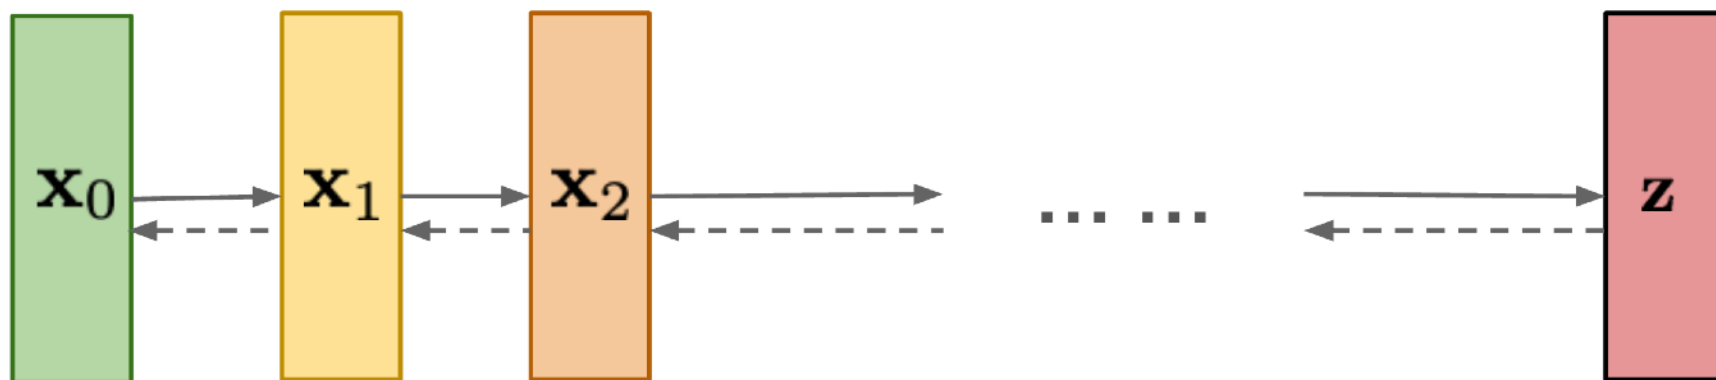
**VAE:** maximize variational lower bound



**Flow-based models:**  
Invertible transform of distributions



**Diffusion models:**  
Gradually add Gaussian noise and then reverse

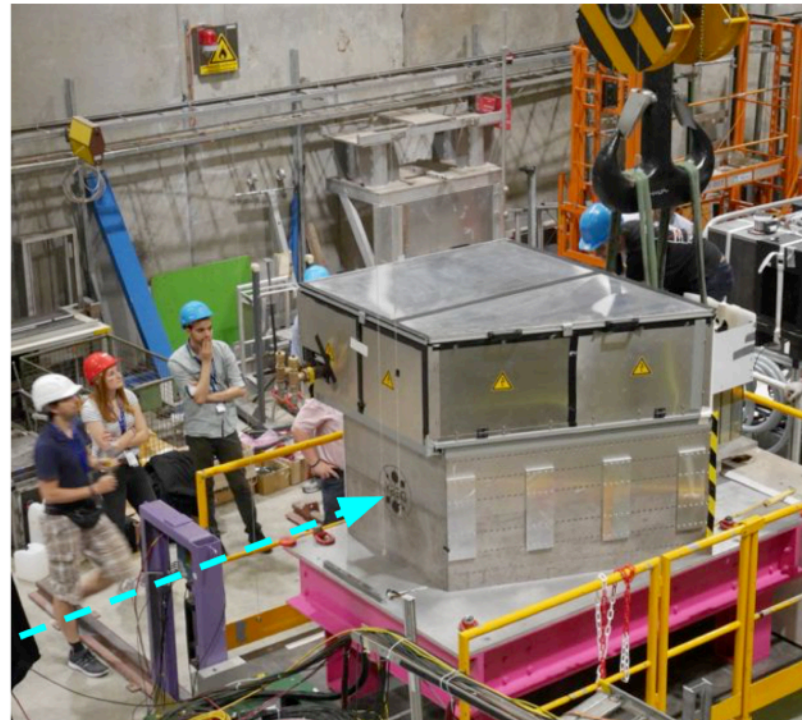


# Generative Models

Simulation is crucial to connect experimental data with theory predictions  
**but computationally very expensive**

Use generative models trained on initial data to augment statistics

**Goal: simulation of particle showers in complex calorimeters like for the ILD or the CMS High Granularity Calorimeter (HGCAL)**



CALICE AHCAL testbeam

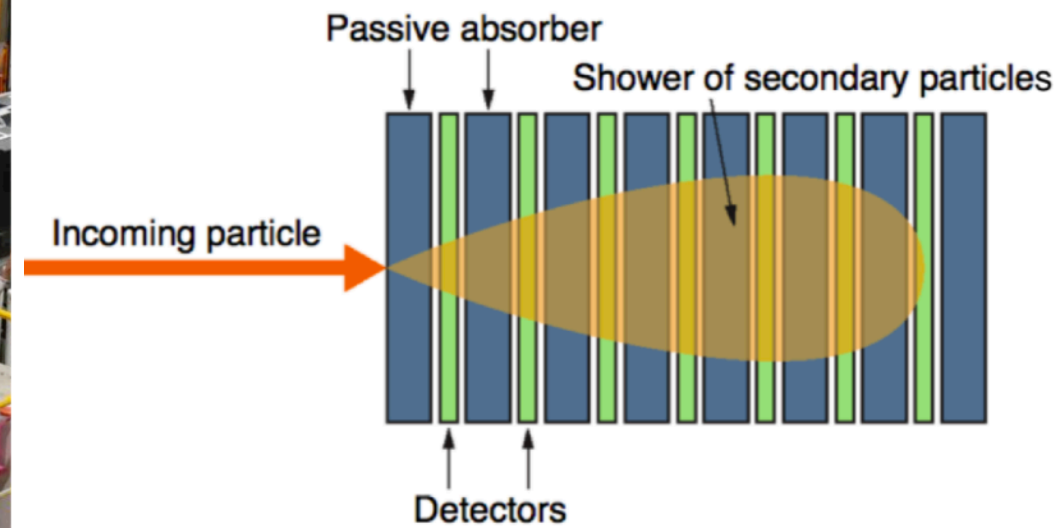
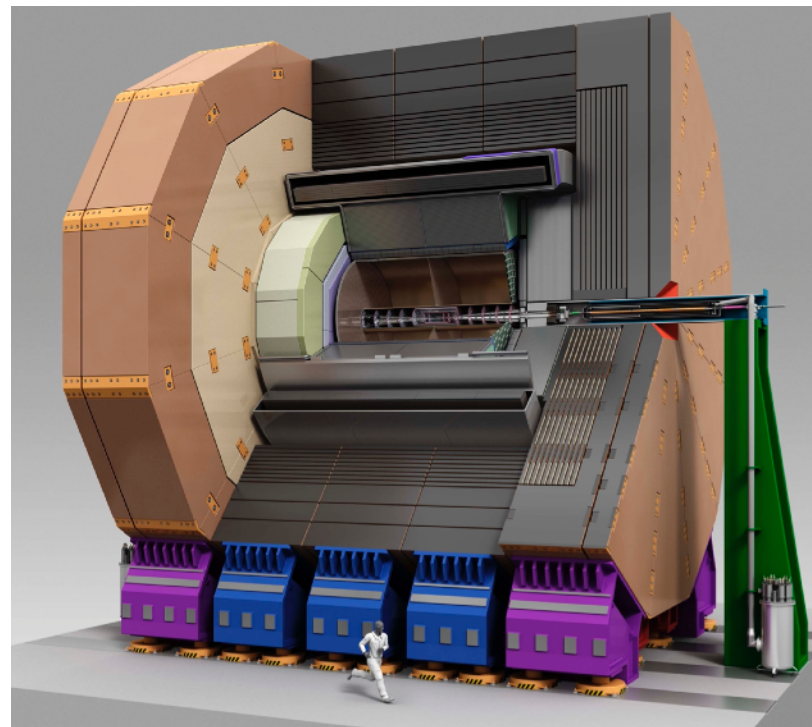
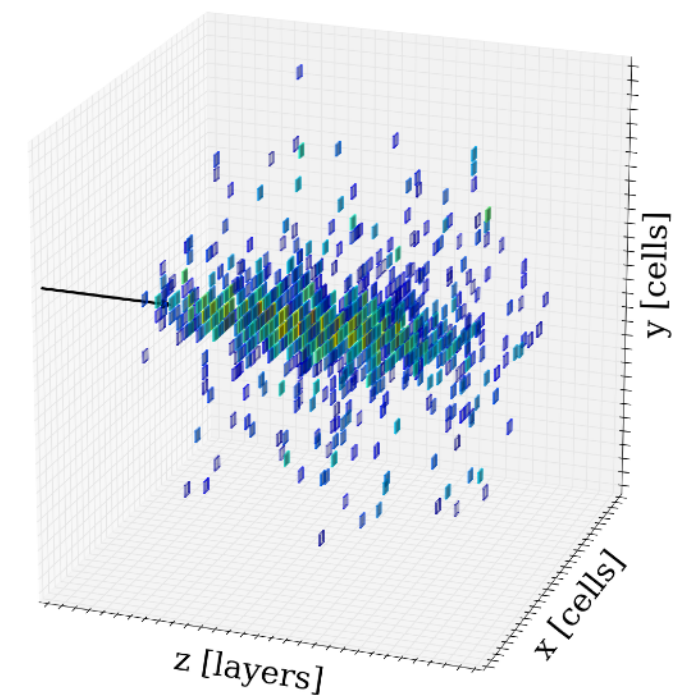


Illustration of particle shower in a sampling calorimeter



ILD Detector



One data example



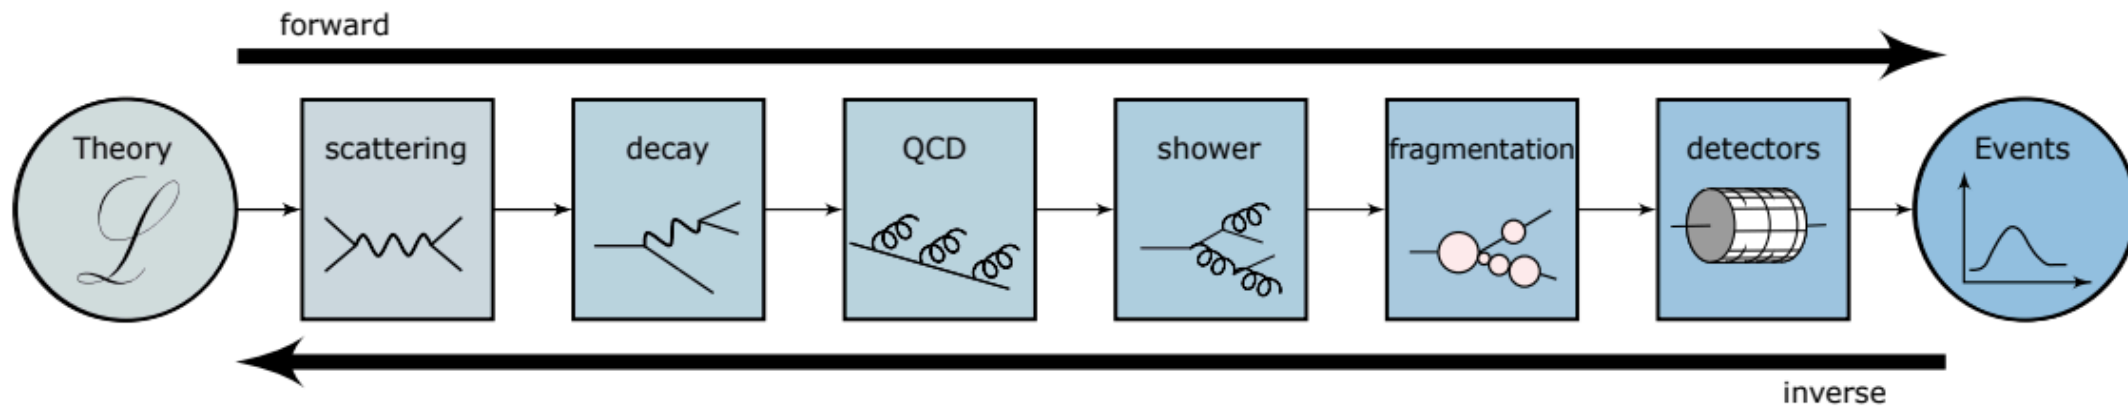
# Aside

Large effort also exists to learn surrogate models of theory simulations // event generation

Similar techniques are used, but different key challenges

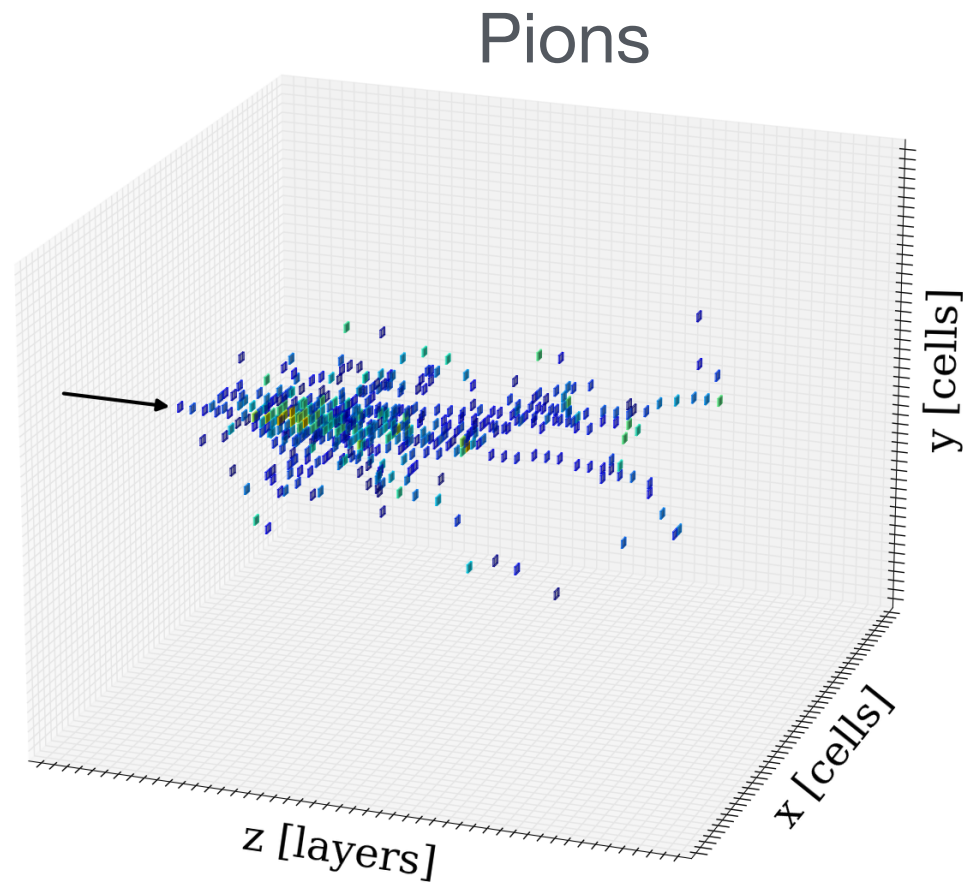
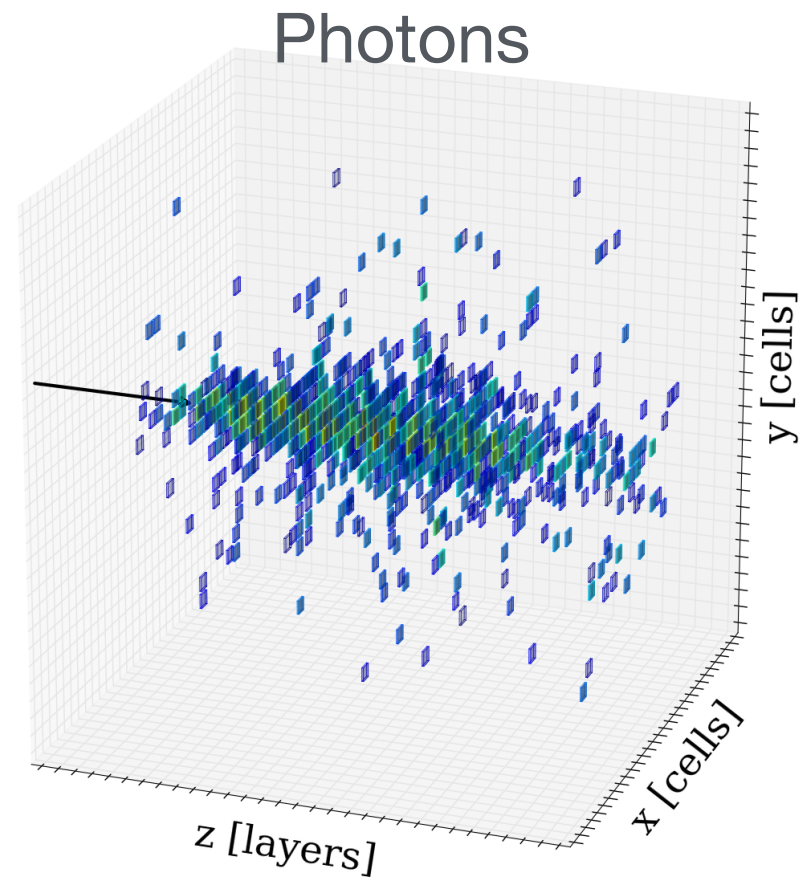
## Machine Learning and LHC Event Generation

Anja Butter<sup>1,2</sup>, Tilman Plehn<sup>1</sup>, Steffen Schumann<sup>3</sup>, Simon Badger<sup>4</sup>, Sascha Caron<sup>5, 6</sup>, Kyle Cranmer<sup>7,8</sup>, Francesco Armando Di Bello<sup>9</sup>, Etienne Dreyer<sup>10</sup>, Stefano Forte<sup>11</sup>, Sanmay Ganguly<sup>12</sup>, Dorival Gonçalves<sup>13</sup>, Eilam Gross<sup>10</sup>, Theo Heimel<sup>1</sup>, Gudrun Heinrich<sup>14</sup>, Lukas Heinrich<sup>15</sup>, Alexander Held<sup>16</sup>, Stefan Höche<sup>17</sup>, Jessica N. Howard<sup>18</sup>, Philip Ilten<sup>19</sup>, Joshua Isaacson<sup>17</sup>, Timo Janßen<sup>3</sup>, Stephen Jones<sup>20</sup>, Marumi Kado<sup>9,21</sup>, Michael Kagan<sup>22</sup>, Gregor Kasieczka<sup>23</sup>, Felix Kling<sup>24</sup>, Sabine Kraml<sup>25</sup>, Claudius Krause<sup>26</sup>, Frank Krauss<sup>20</sup>, Kevin Kröniger<sup>27</sup>, Rahool Kumar Barman<sup>13</sup>, Michel Luchmann<sup>1</sup>, Vitaly Magerya<sup>14</sup>, Daniel Maitre<sup>20</sup>, Bogdan Malaescu<sup>2</sup>, Fabio Maltoni<sup>28,29</sup>, Till Martini<sup>30</sup>, Olivier Mattelaer<sup>28</sup>, Benjamin Nachman<sup>31,32</sup>, Sebastian Pitz<sup>1</sup>, Juan Rojo<sup>33,34</sup>, Matthew Schwartz<sup>35</sup>, David Shih<sup>25</sup>, Frank Siegert<sup>36</sup>, Roy Stegeman<sup>11</sup>, Bob Stienen<sup>5</sup>, Jesse Thaler<sup>37</sup>, Rob Verheyen<sup>38</sup>, Daniel Whiteson<sup>18</sup>, Ramon Winterhalder<sup>28</sup>, and Jure Zupan<sup>19</sup>



# Example

Simulate showers initiated by photons and charged pions



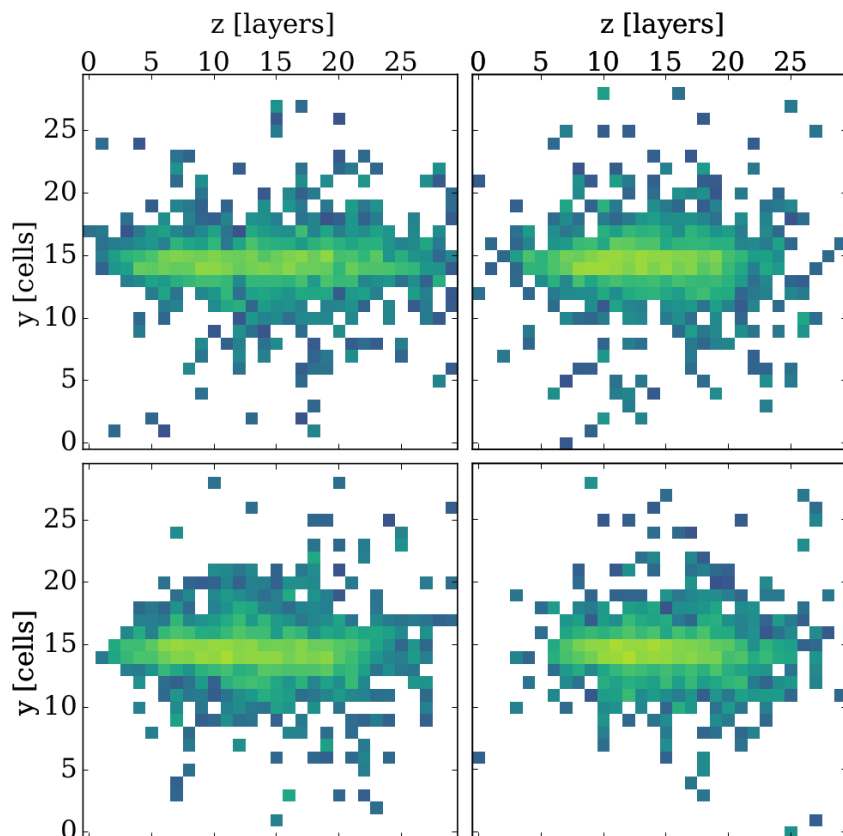


# Example

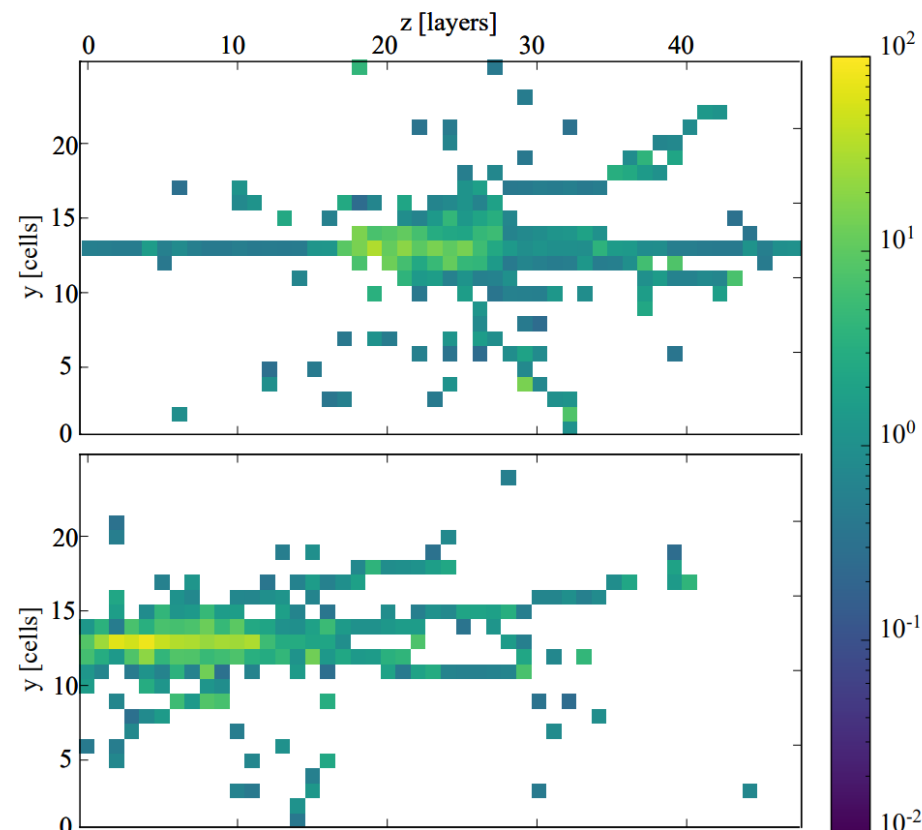
Simulate showers initiated by photons and charged pions

Photons

Pions



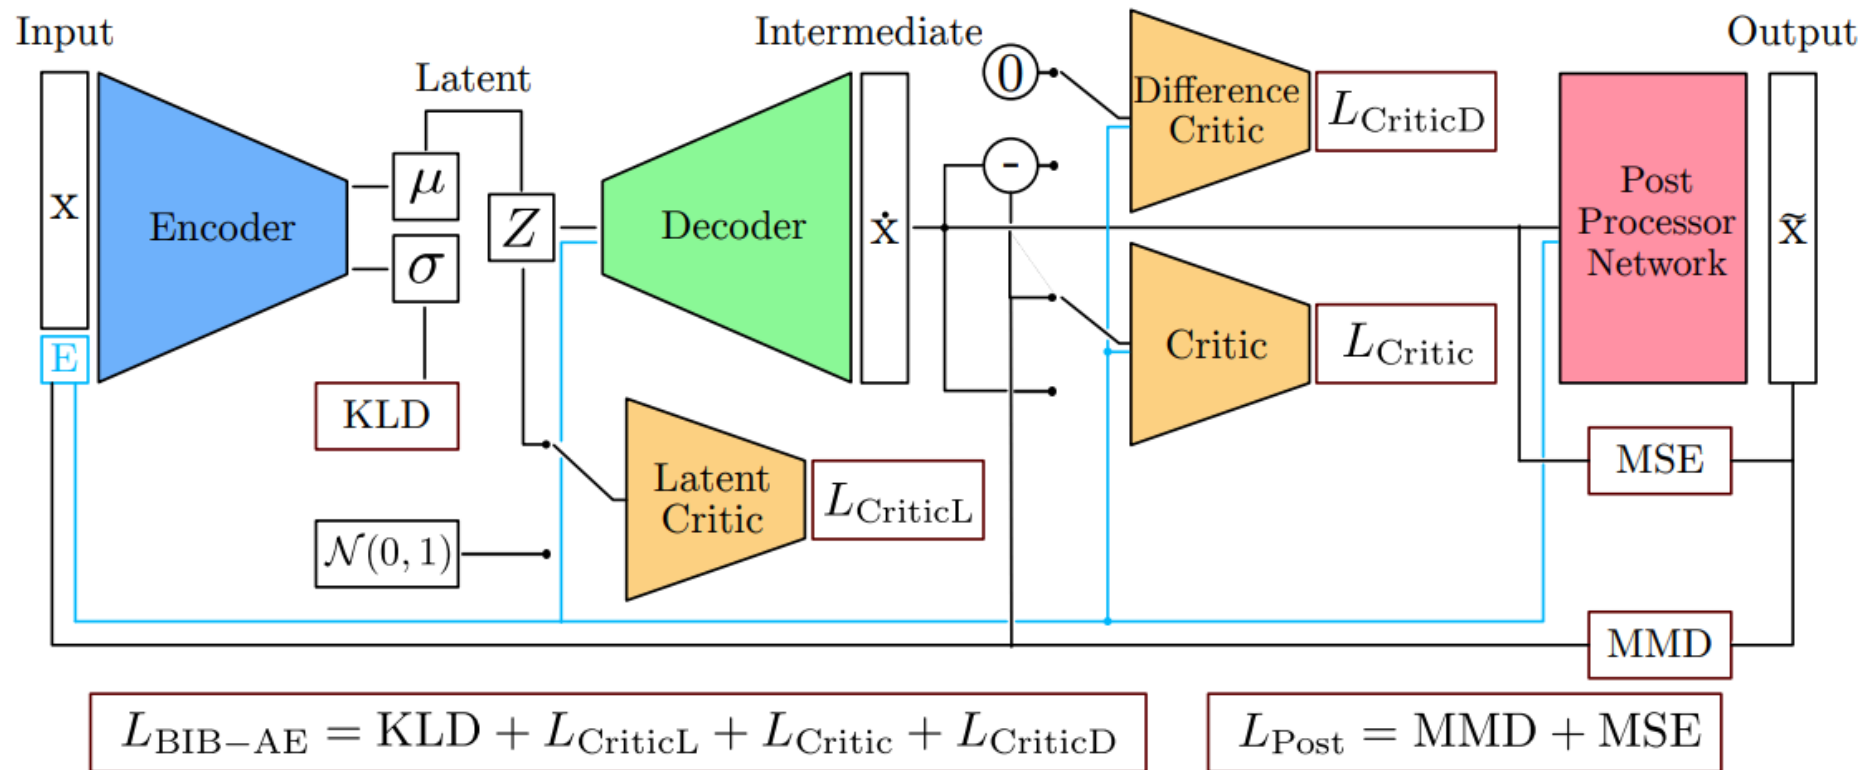
ILD ECal  
Silicon/Tungstate



ILD AHCal  
stainless steel absorber  
plates and scintillator  
tiles with SiPMs

Pion showers are significantly more complex

# Example



Simulate showers initiated by photons and charged pions

Pion showers are significantly more complex

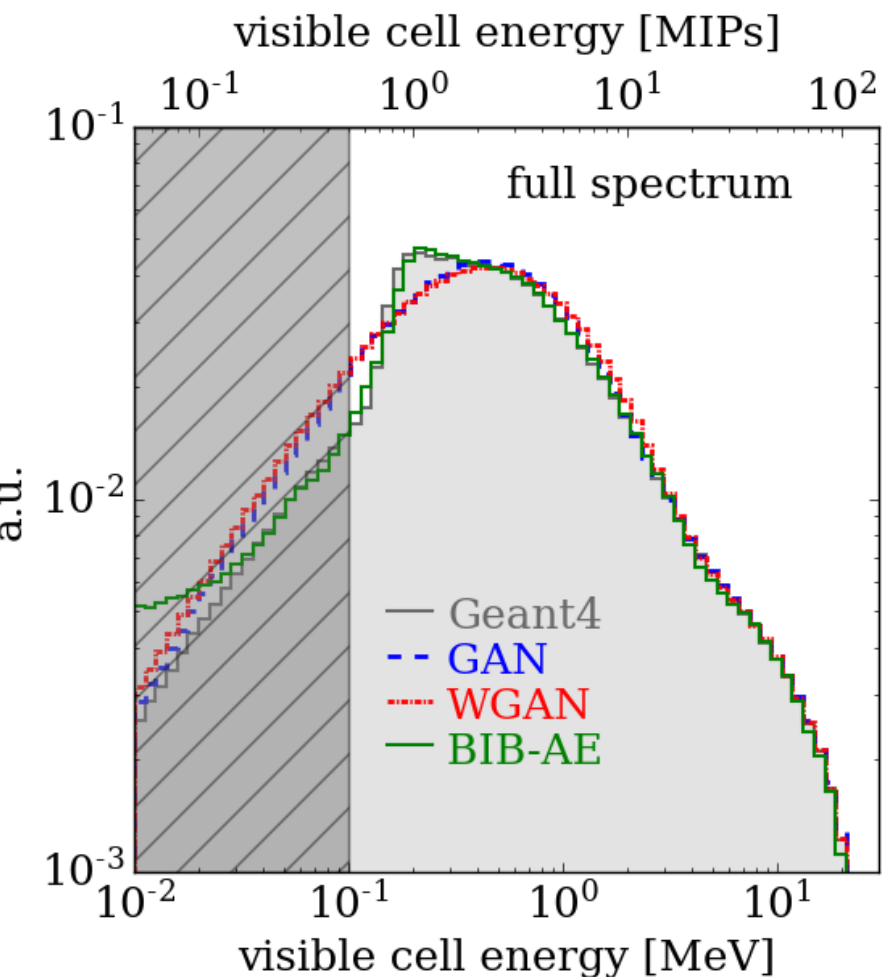
Use modified BIB-AE architecture

$$\begin{aligned}
 L_{\text{BIB-AE}} = & -\beta_{C_L} \cdot \mathbb{E}_{x \sim p_{\text{data}(x)}} [C_L(E(x))] \\
 & -\beta_C \cdot \mathbb{E}_{x \sim p_{\text{data}(x)}} [C(D(E(x)))] \\
 & + \beta_{\text{KLD}} \cdot \text{KLD}(E(x)) \\
 & + \beta_{\text{MMD}} \cdot \text{MMD}(E(x), \mathcal{N}(0, 1)).
 \end{aligned}$$

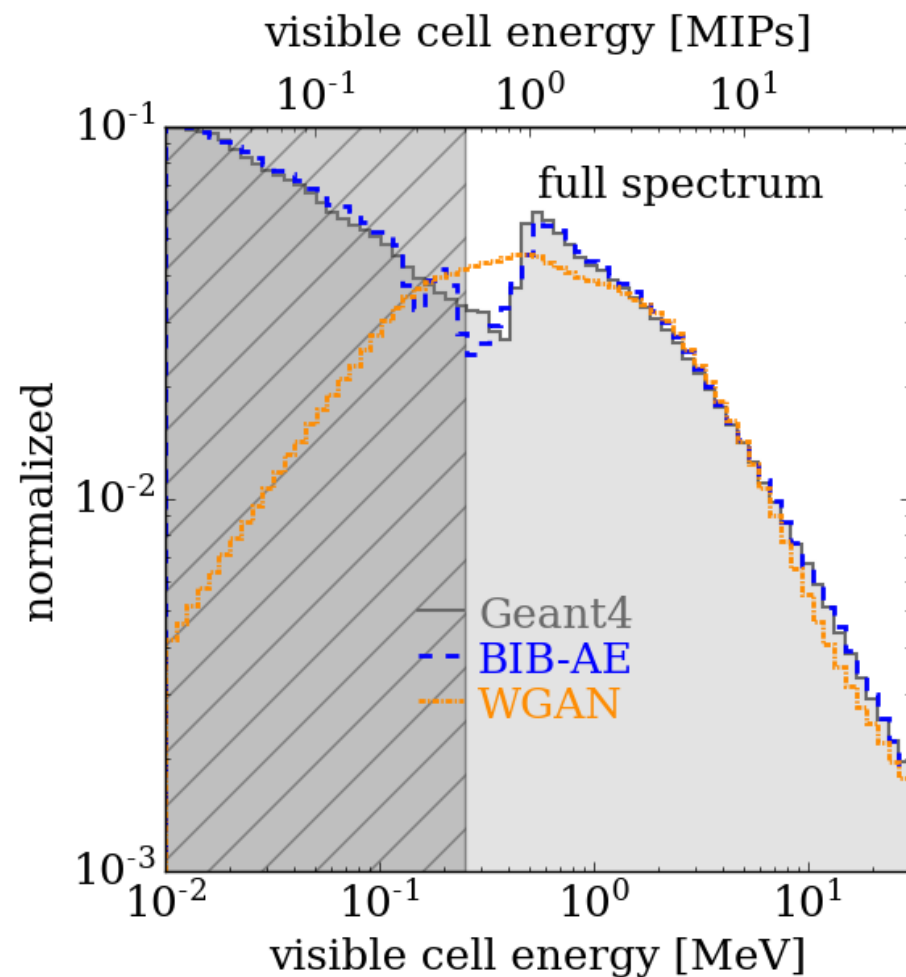


# Example

## Photons



## Pions



Simulate showers initiated by photons and charged pions

Pion showers are significantly more complex

Use modified BIB-AE architecture

Achieves excellent modelling of differential distributions — first for high-granularity calorimeter

# Timing

Simulator	Hardware	Batch Size	15 GeV	Speed-up	10-100 GeV Flat	Speed-up
Geant4	CPU	N/A	1445.05 ± 19.34 ms	-	4081.53 ± 169.92 ms	-
WGAN	CPU	1	64.34 ± 0.58 ms	<b>x23</b>	63.14 ± 0.34 ms	<b>x65</b>
		10	59.53 ± 0.45 ms	<b>x24</b>	56.65 ± 0.33 ms	<b>x72</b>
		100	58.31 ± 0.93 ms	<b>x25</b>	58.11 ± 0.13 ms	<b>x70</b>
		1000	57.99 ± 0.97 ms	<b>x25</b>	57.99 ± 0.18 ms	<b>x70</b>
BIB-AE	CPU	1	426.60 ± 3.27 ms	<b>x3</b>	426.32 ± 3.62 ms	<b>x10</b>
		10	422.60 ± 0.26 ms	<b>x3</b>	424.71 ± 3.53 ms	<b>x10</b>
		100	419.64 ± 0.07 ms	<b>x3</b>	418.04 ± 0.20 ms	<b>x10</b>
WGAN	GPU	1	3.24 ± 0.01 ms	<b>x446</b>	3.25 ± 0.01 ms	<b>x1256</b>
		10	6.13 ± 0.02 ms	<b>x236</b>	6.13 ± 0.02 ms	<b>x666</b>
		100	5.43 ± 0.01 ms	<b>x266</b>	5.43 ± 0.01 ms	<b>x752</b>
		1000	5.43 ± 0.01 ms	<b>x266</b>	5.43 ± 0.01 ms	<b>x752</b>
BIB-AE	GPU	1	3.14 ± 0.01 ms	<b>x460</b>	3.19 ± 0.01 ms	<b>x1279</b>
		10	1.56 ± 0.01 ms	<b>x926</b>	1.57 ± 0.01 ms	<b>x2600</b>
		100	1.42 ± 0.01 ms	<b>x1017</b>	1.42 ± 0.01 ms	<b>x2874</b>

Photons

Hardware	Simulator	Time / Shower [ms]	Speed-up
CPU	GEANT4	2684 ± 125	×1
	WGAN	47.923 ± 0.089	×56
	BIB-AE	350.824 ± 0.574	×8
GPU	WGAN	0.264 ± 0.002	×10167
	BIB-AE	2.051 ± 0.005	×1309

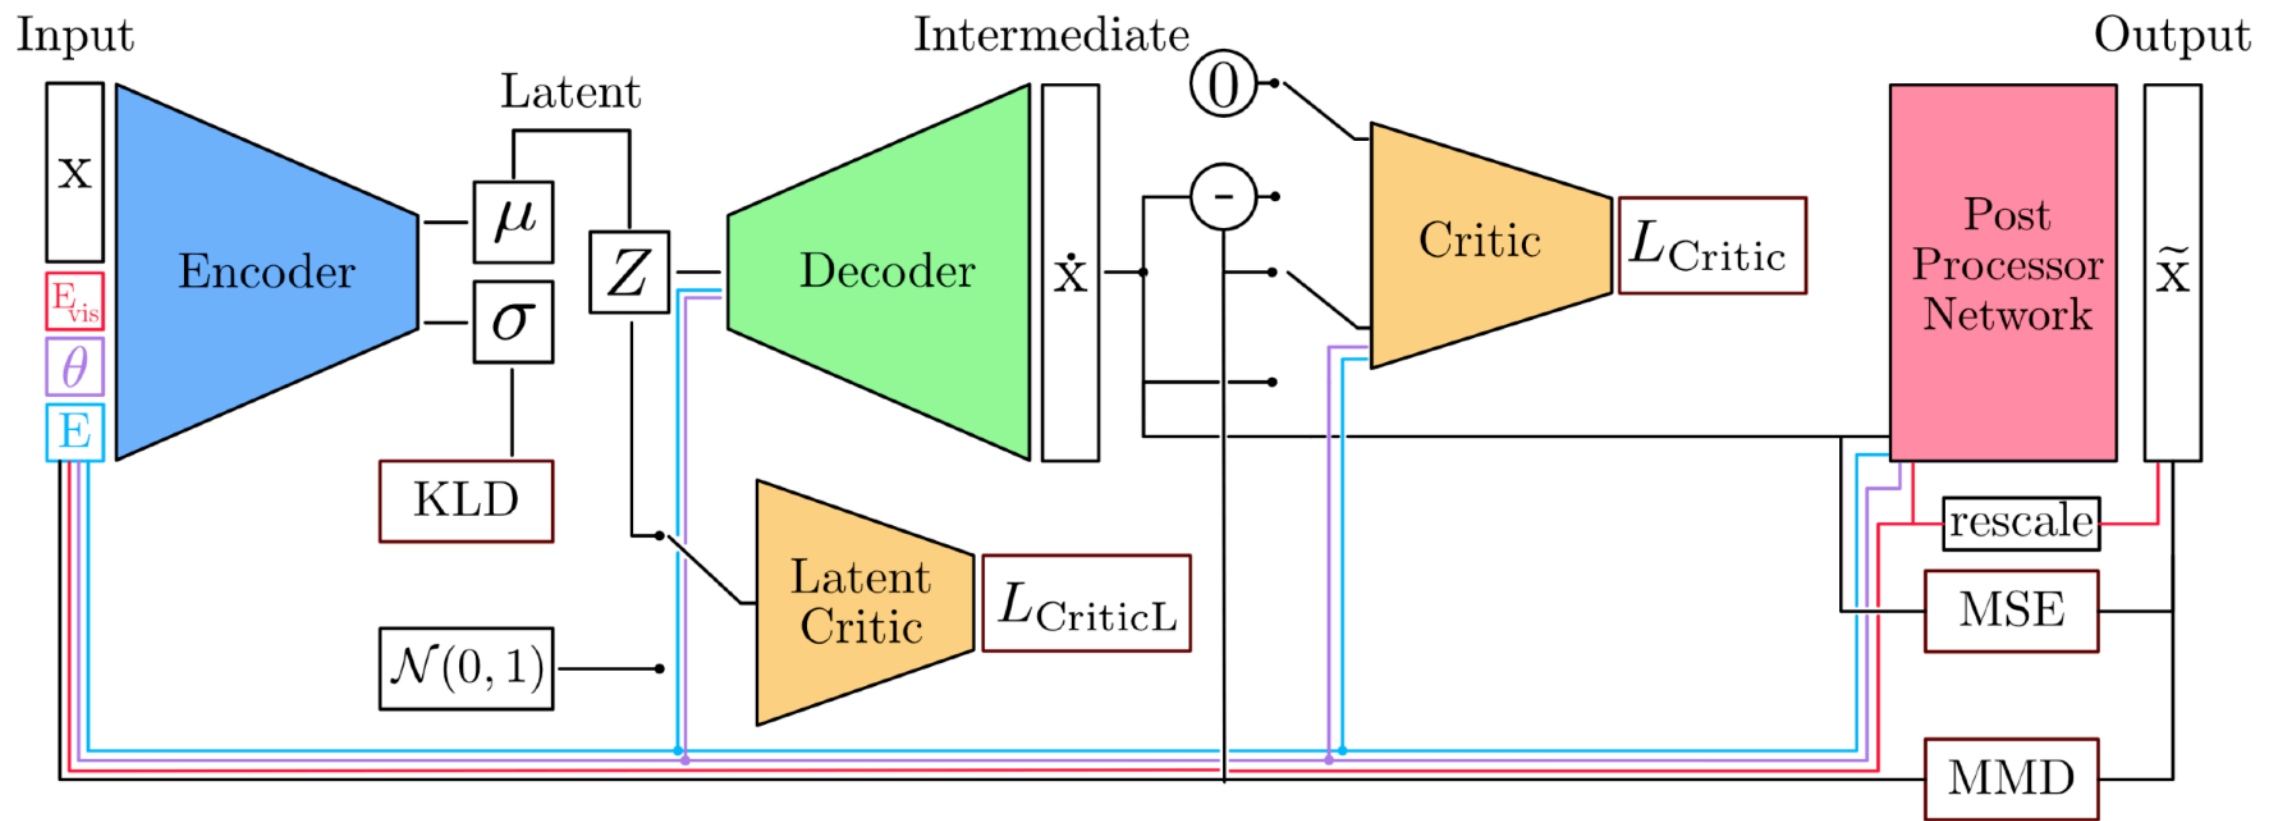
Pions

Buhmann, .., GK et al 2005.05334;

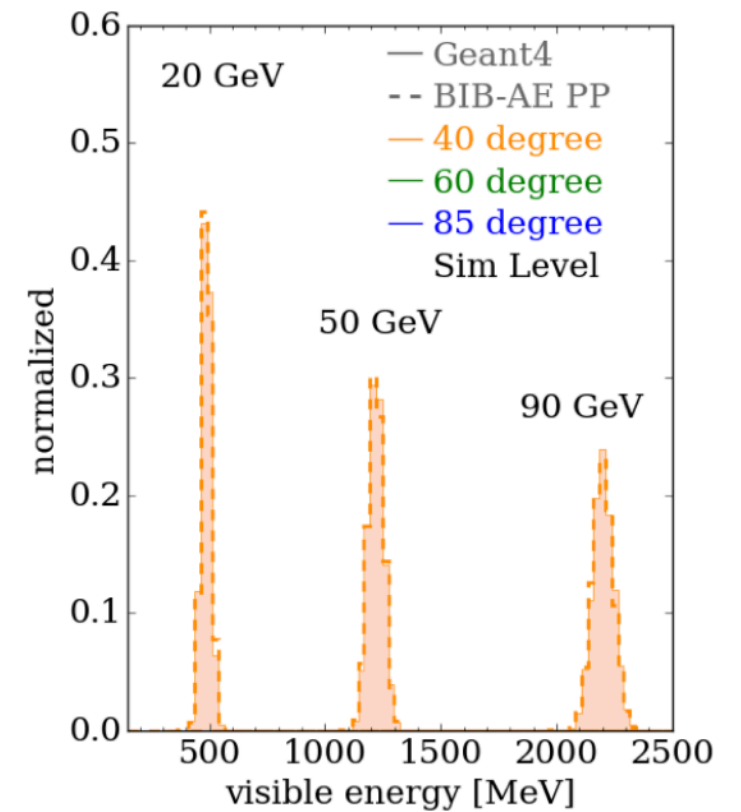
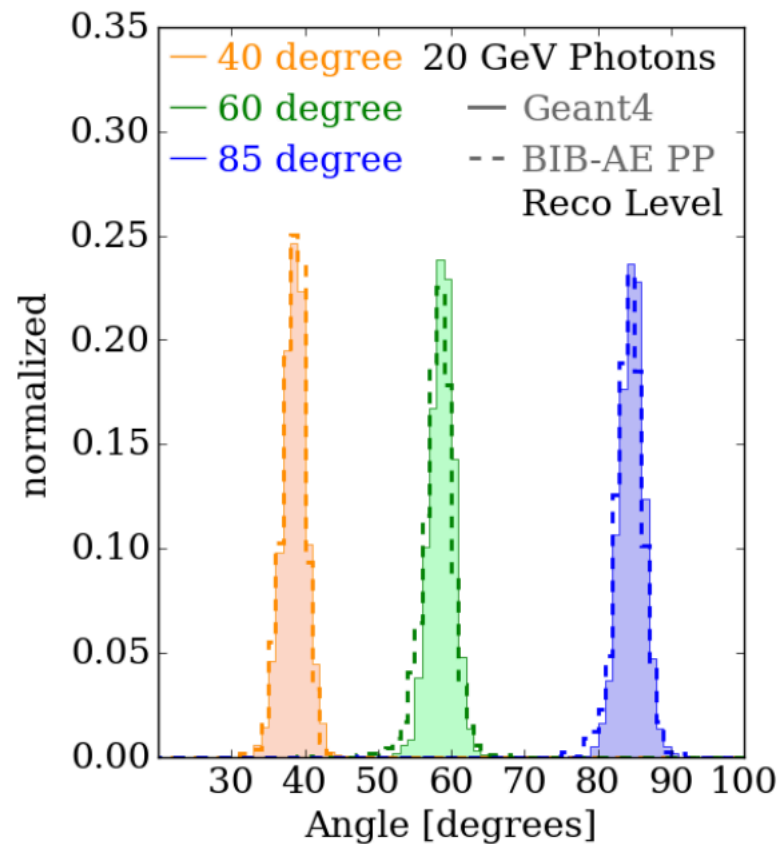
Buhmann, .., GK et al 2112.09709;



# Dual conditioning

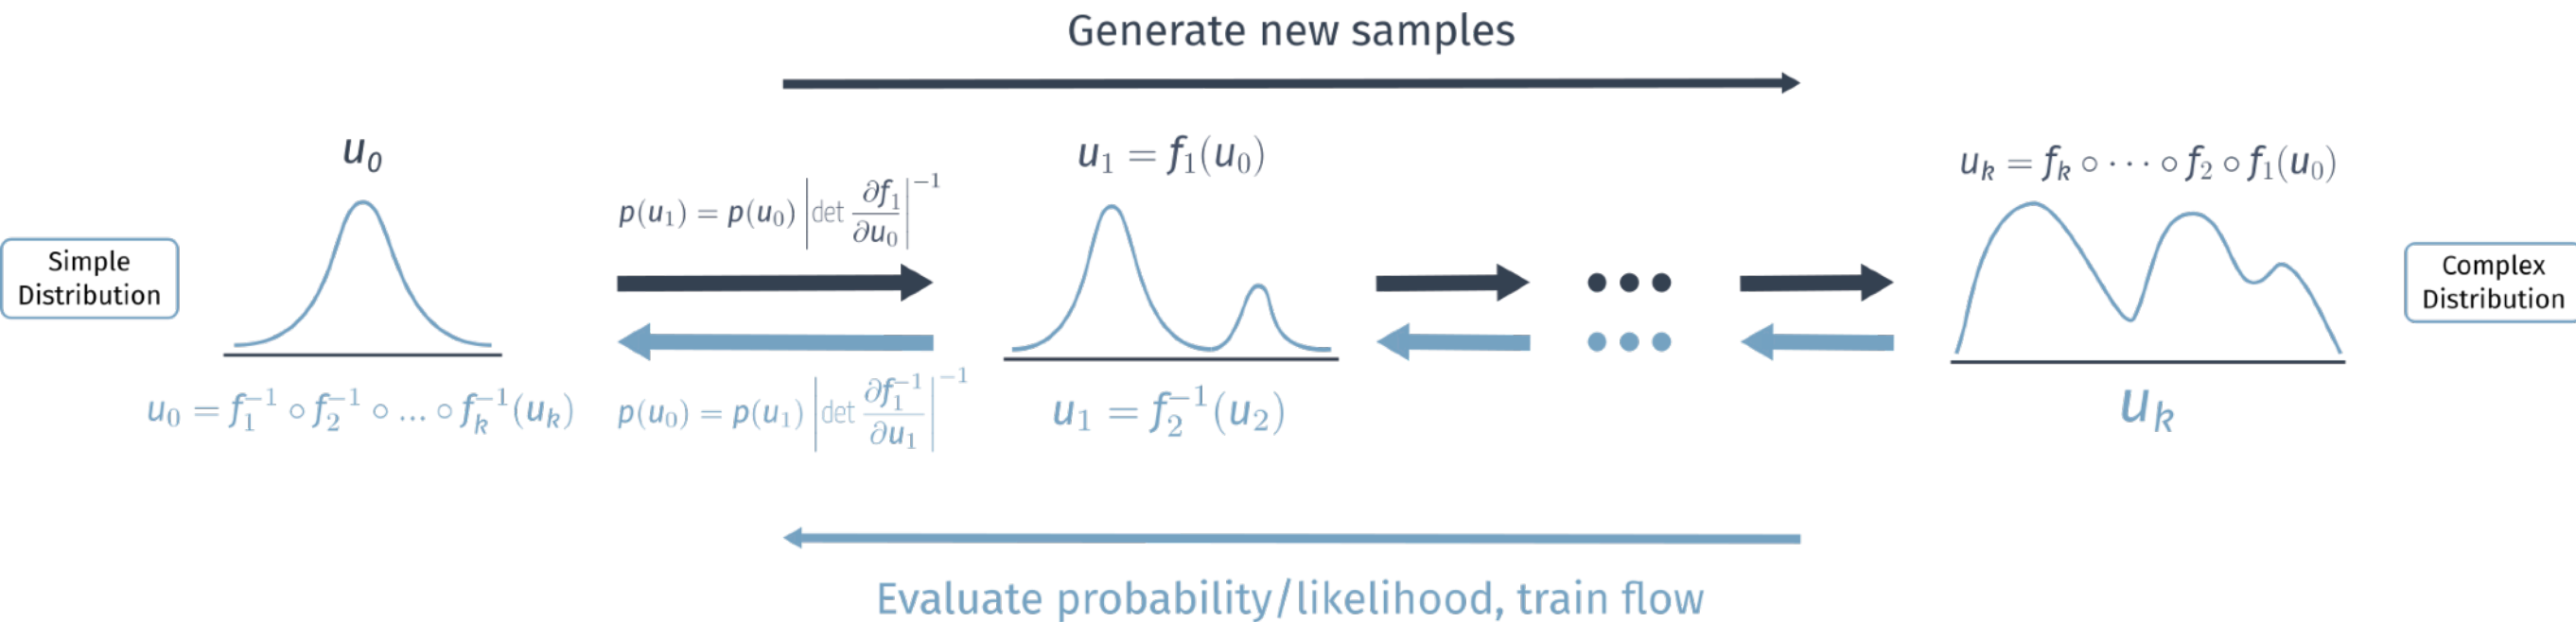


Additional control over one incoming particle angle



# Flow based generation

How can we further improve the quality of generated distributions?



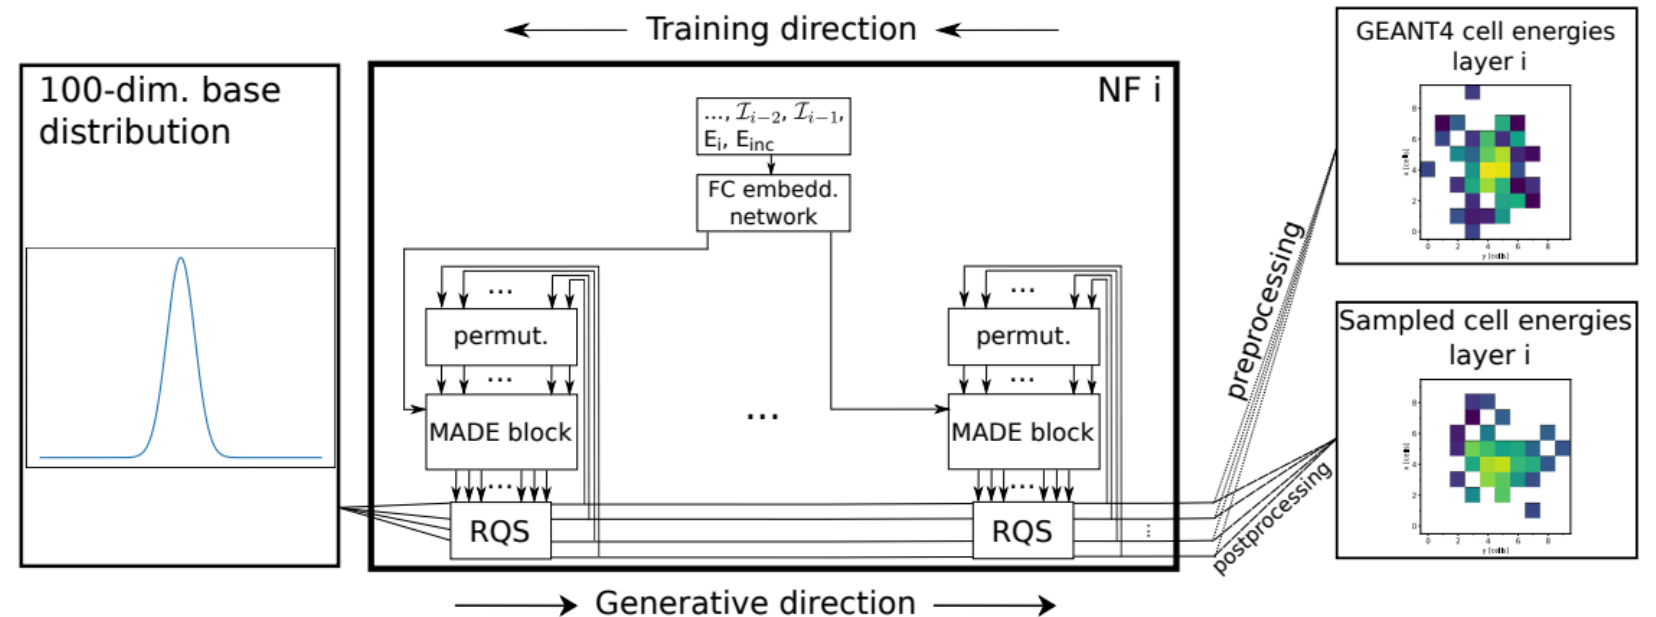
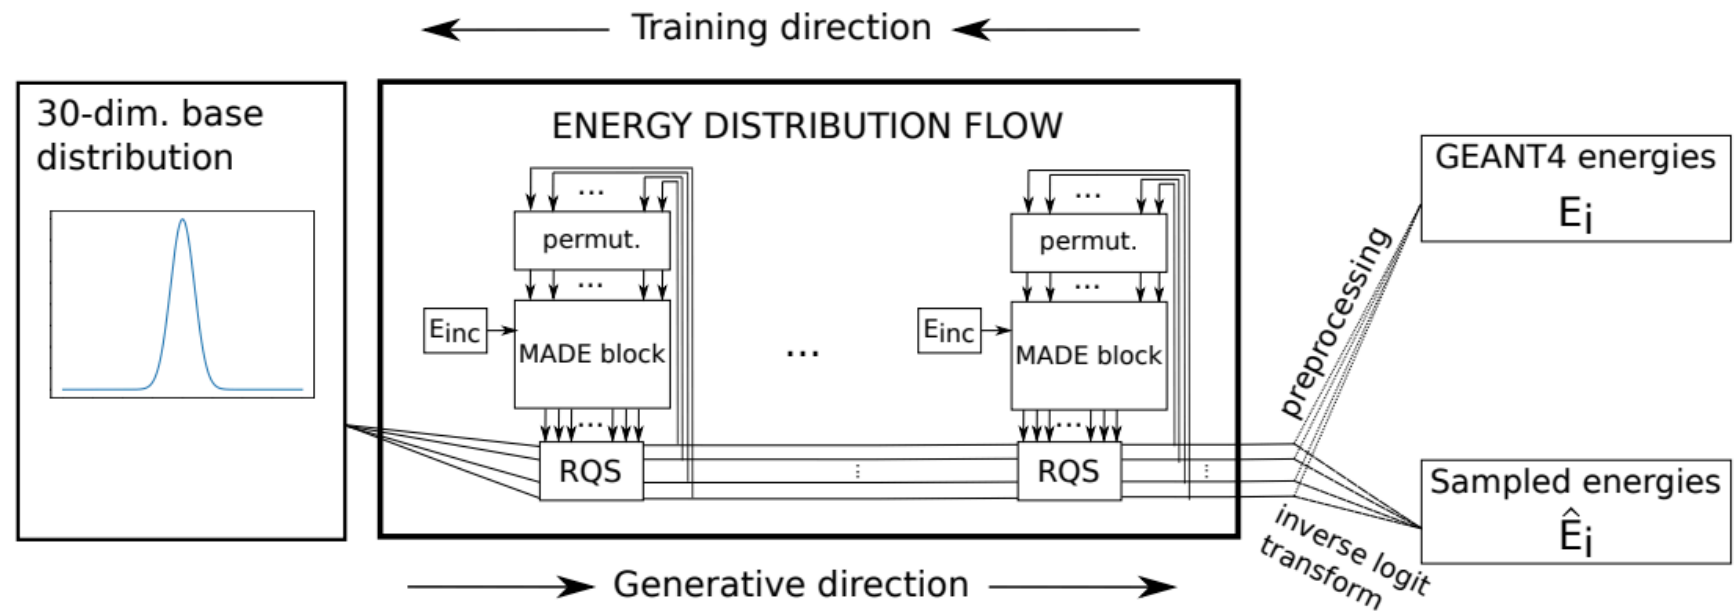
Flow based models learn an exact mapping to a Gaussian latent space, can be trained by direct likelihood minimisation (no adversarial terms, more stable!)



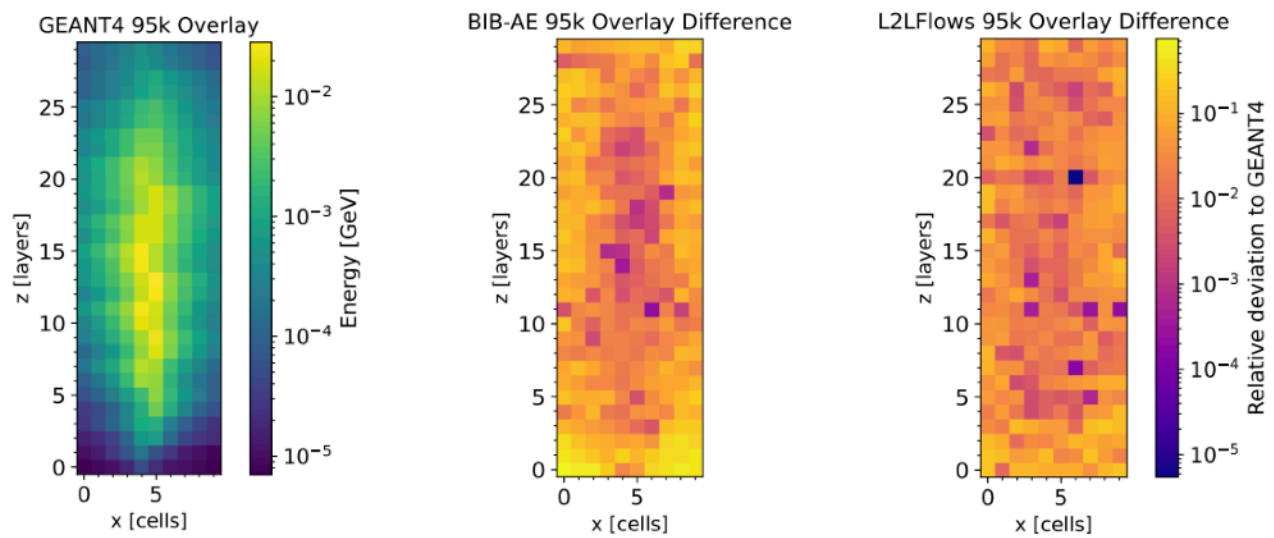
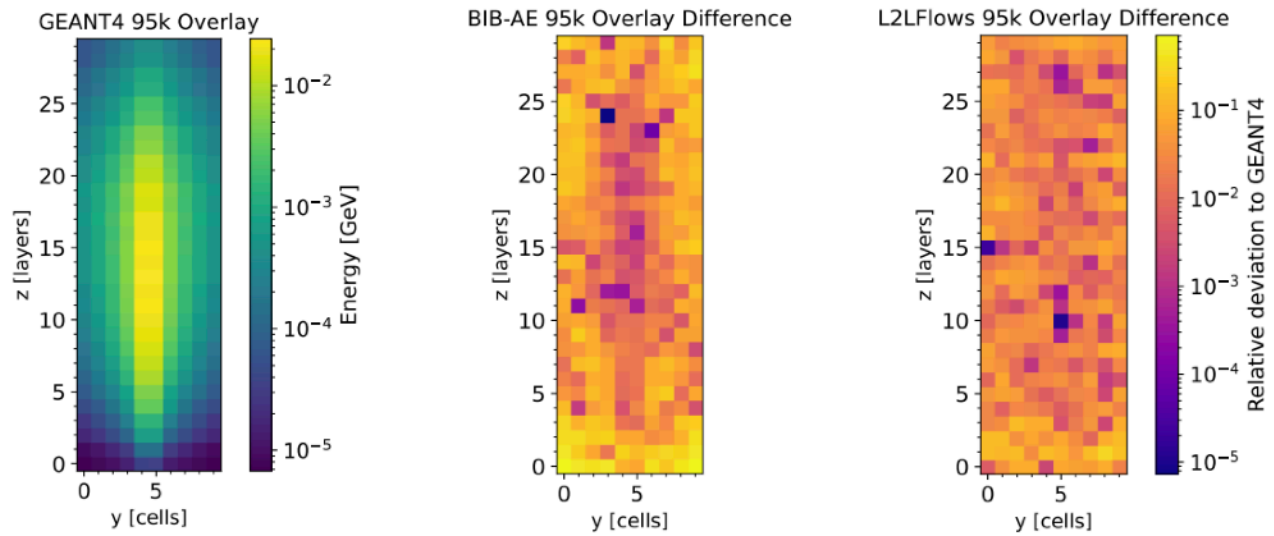
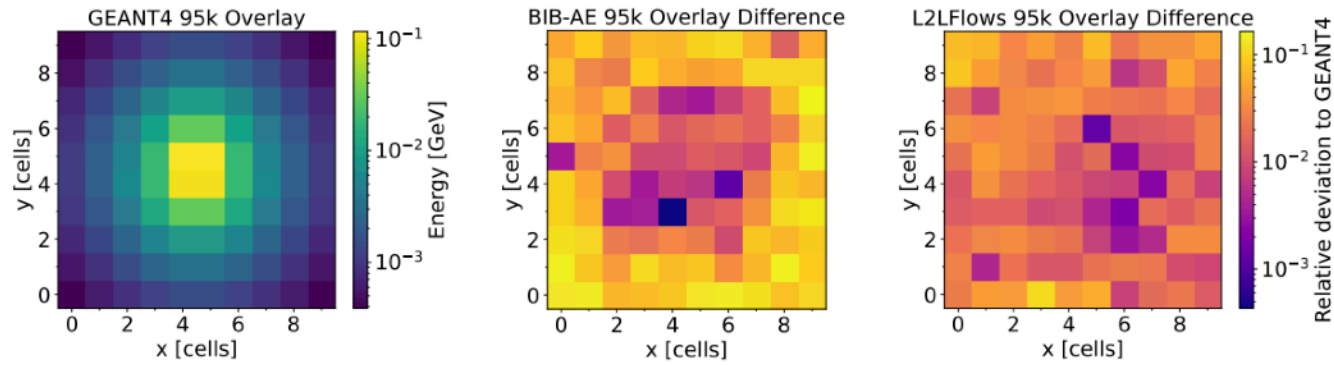
# Flow based generation

Disadvantage: As invertible model, latent space dimension is same as data space dimension, bad scaling with large inputs

Solution:  
Two step generation,  
Sequential conditioning



# L2L Flows



# Showers per simulator	AUC GEANT4 vs L2LFlows	AUC GEANT4 vs BIB-AE
95k	$0.8518 \pm 0.0042$	$0.9947 \pm 0.0025$
190k	$0.8768 \pm 0.0029$	—
380k	$0.8962 \pm 0.0024$	—
760k	$0.9402 \pm 0.0011$	—

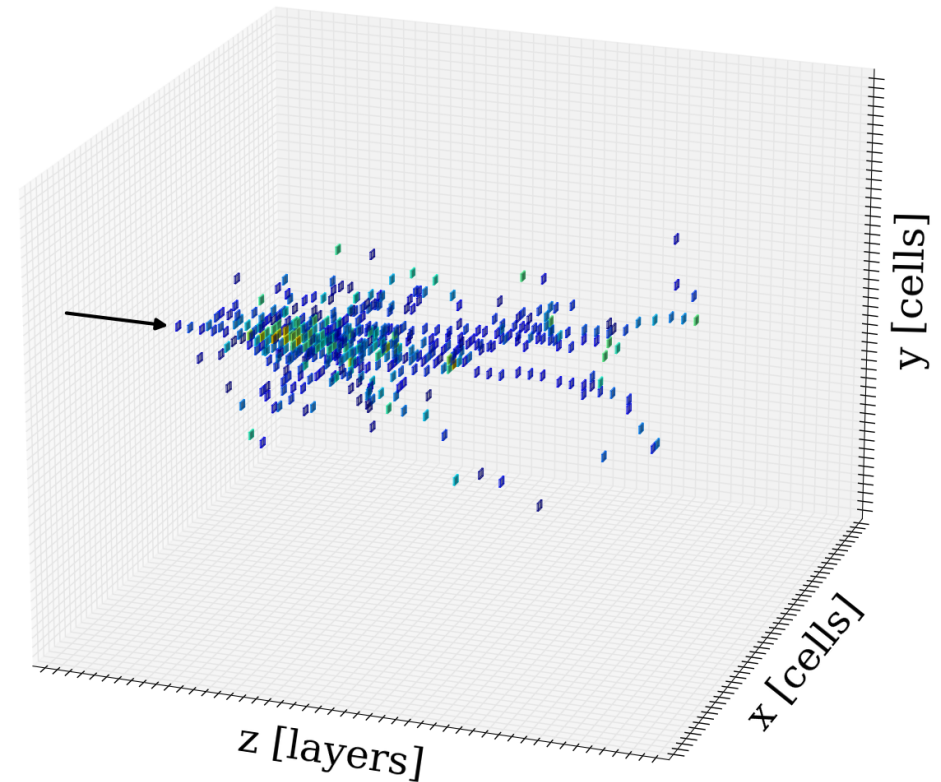
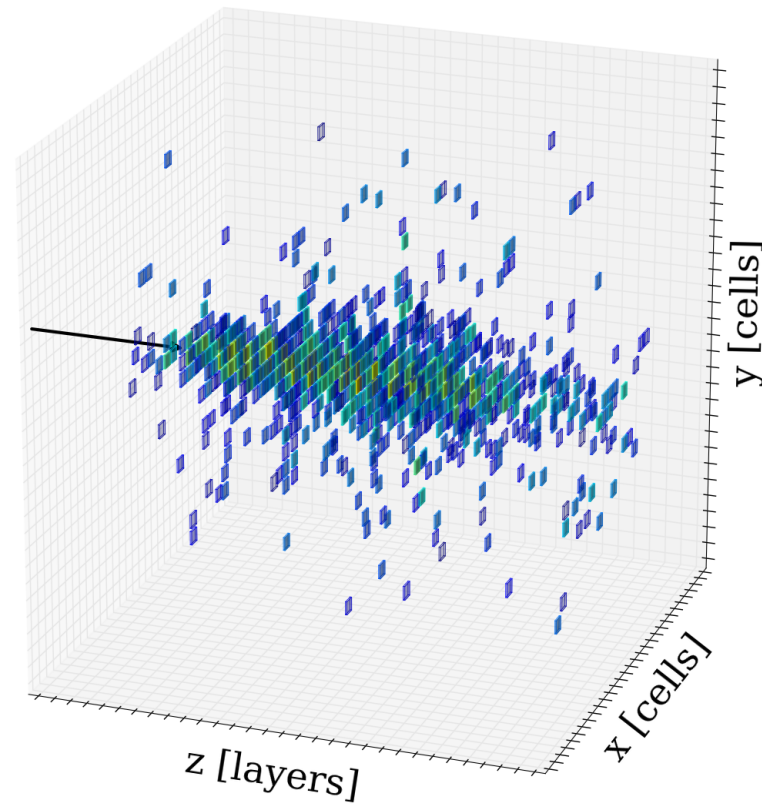
Better generative fidelity than BIB-AE,  
working on scaling up!



# Point Cloud Generation

For high-resolution detectors, images will be very **sparse** and geometry **impossible** to map to fixed grid. Instead treat showers as **graphs or point clouds**

The physics of particle shower is **simple**: Point cloud (instead of graph) more efficient and **much faster**

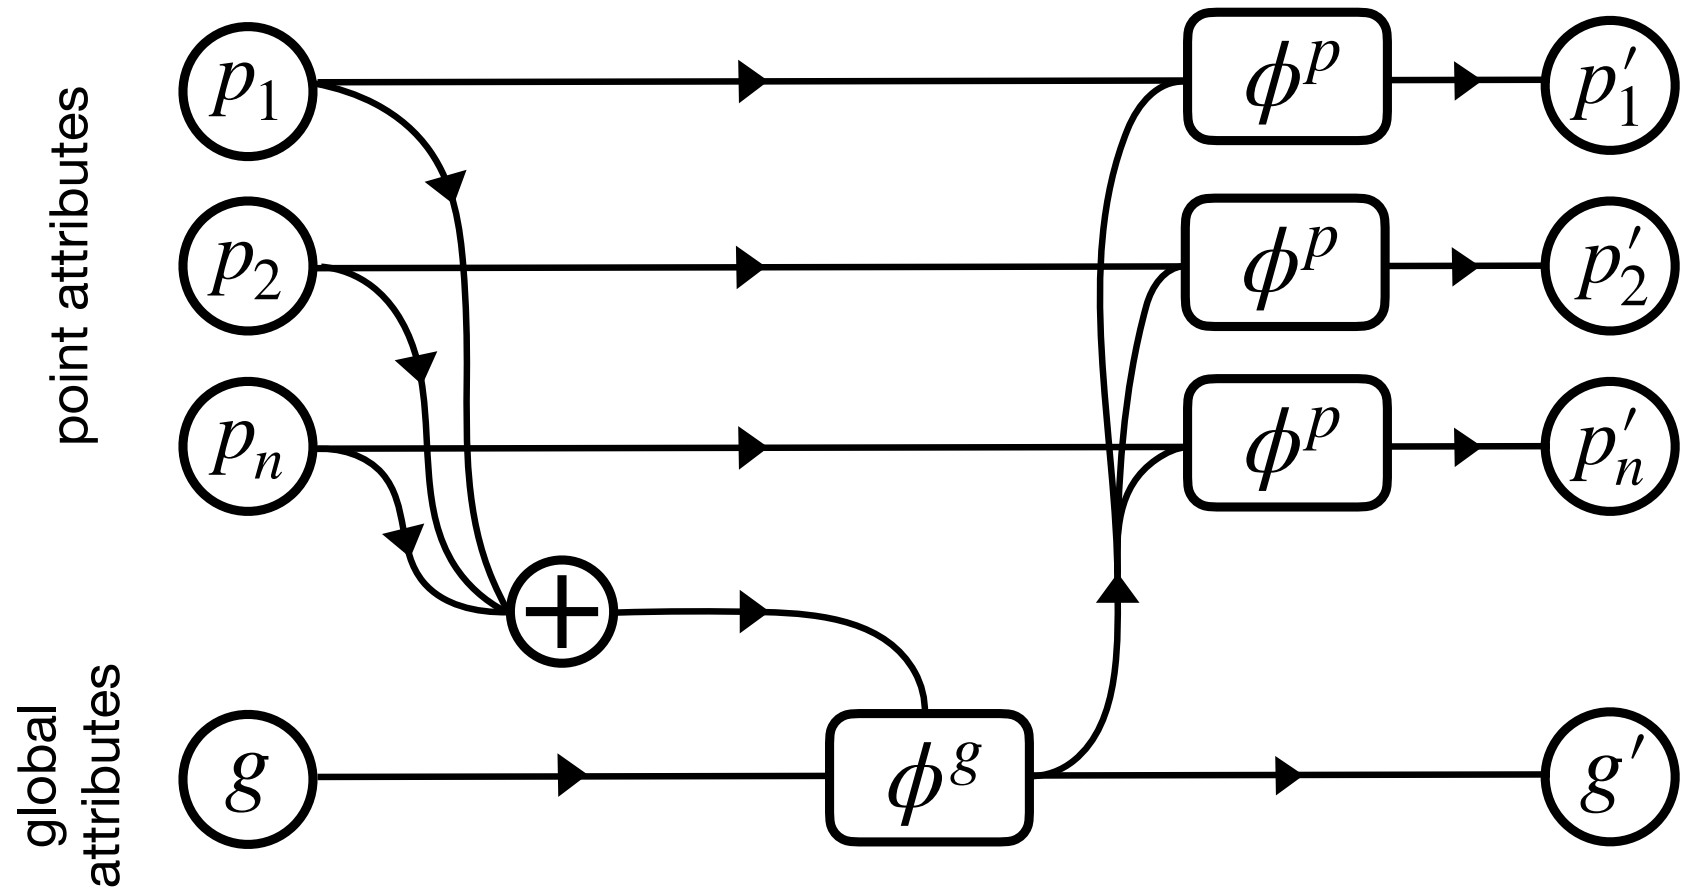


Buhmann, GK, Thaler 2301.08128;  
Kansal et al 2106.11535; Käch et al  
2211.13630; Buhmann, ... GK, et al  
2305.04847

# Point Cloud Generation

For high-resolution detectors, images will be very **sparse** and geometry **impossible** to map to fixed grid. Instead treat showers as **graphs or point clouds**

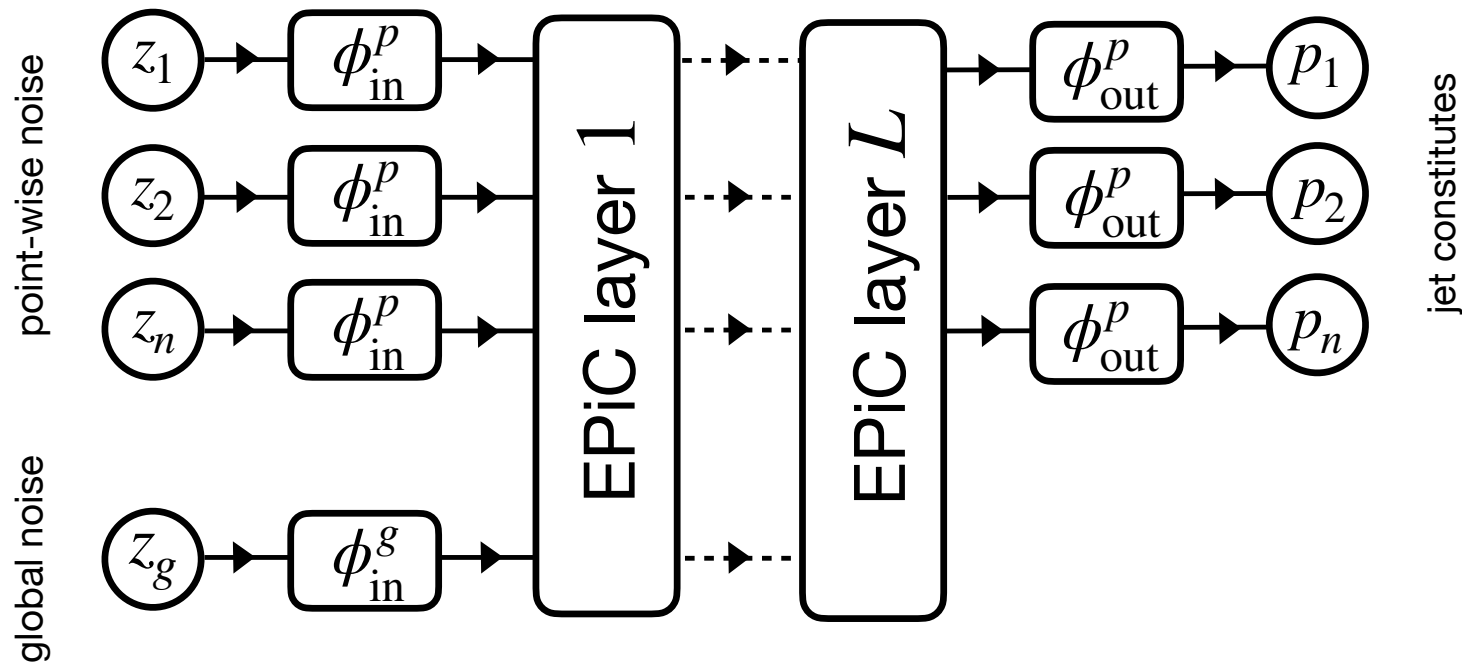
The physics of particle shower is **simple**: Point cloud (instead of graph) more efficient and **much faster**



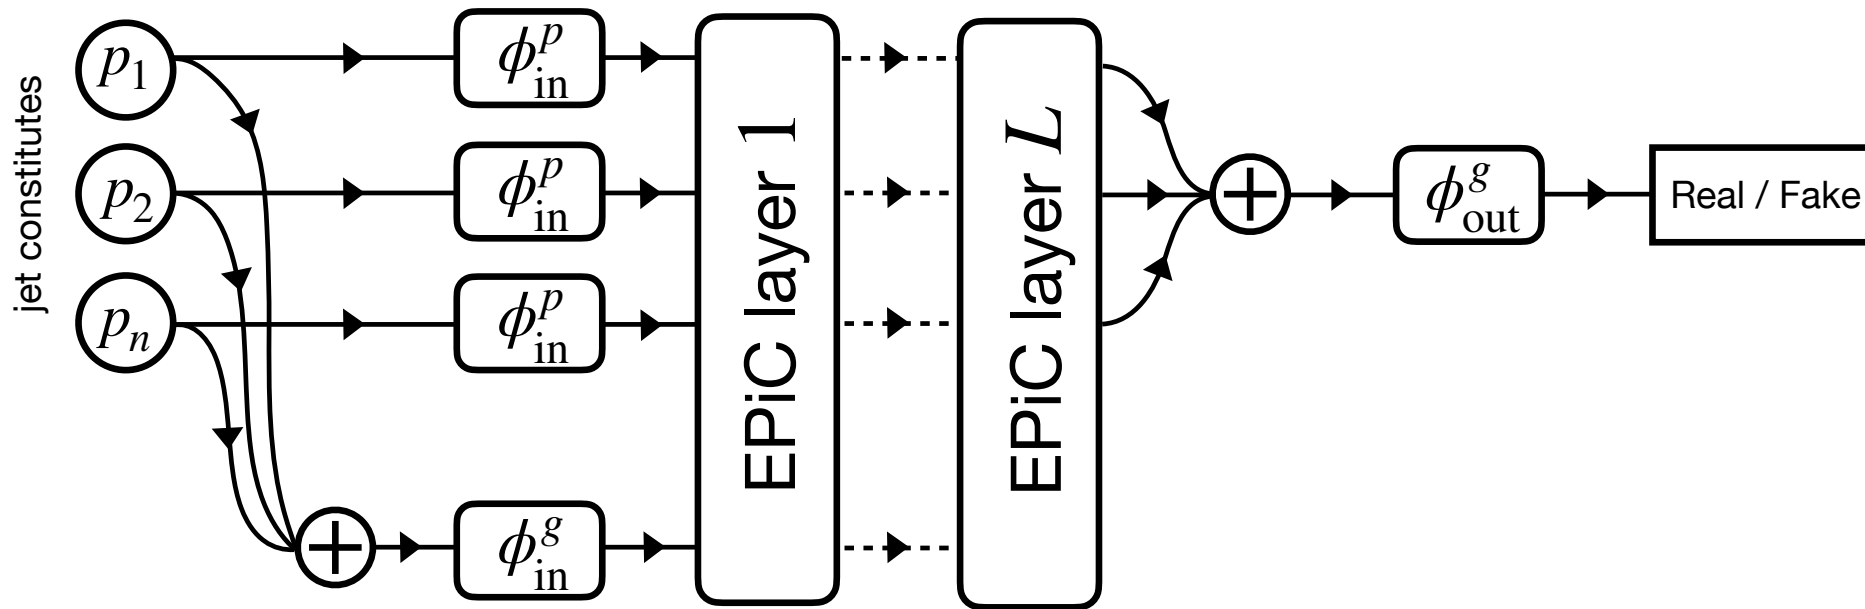
Equivariant Point Cloud interaction (EPIc) block: Similar to deep sets, but with additional global information exchange.

Still permutation equivariant

# Point Cloud Generation



(a) Generator

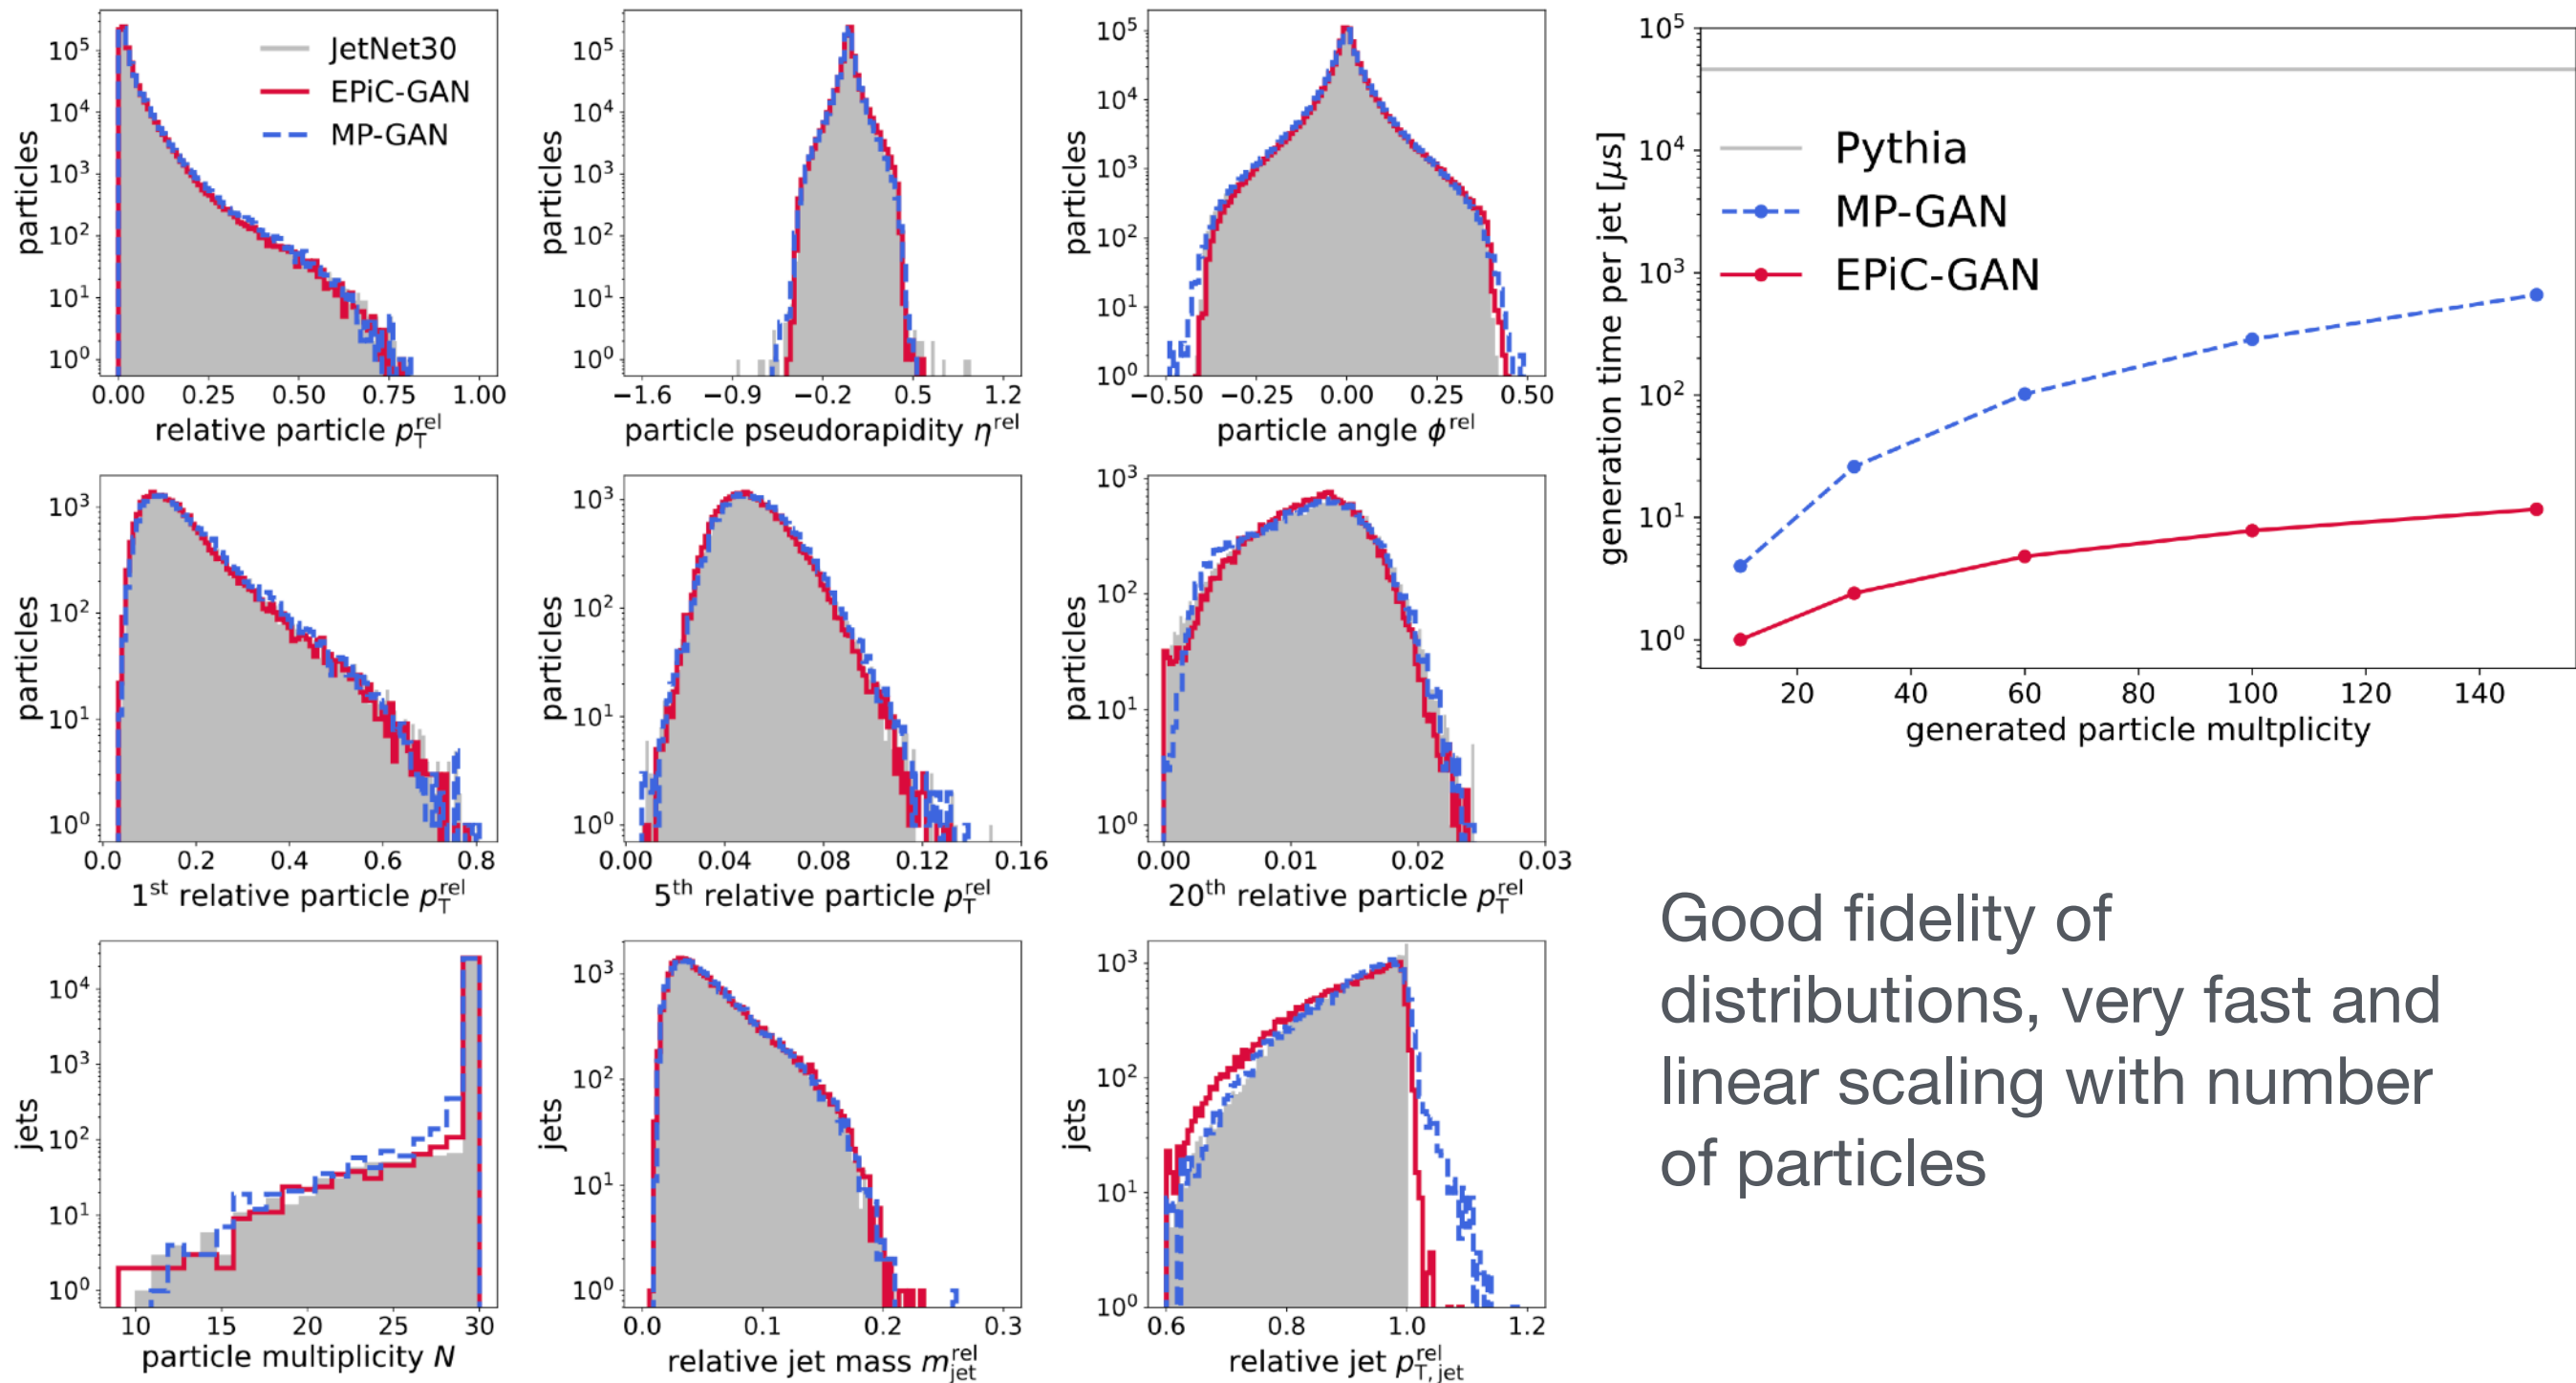


(b) Discriminator

Can use to build generator and discriminator blocks for classical GAN architecture, but fully respecting permutation symmetry.



# Point Cloud Generation

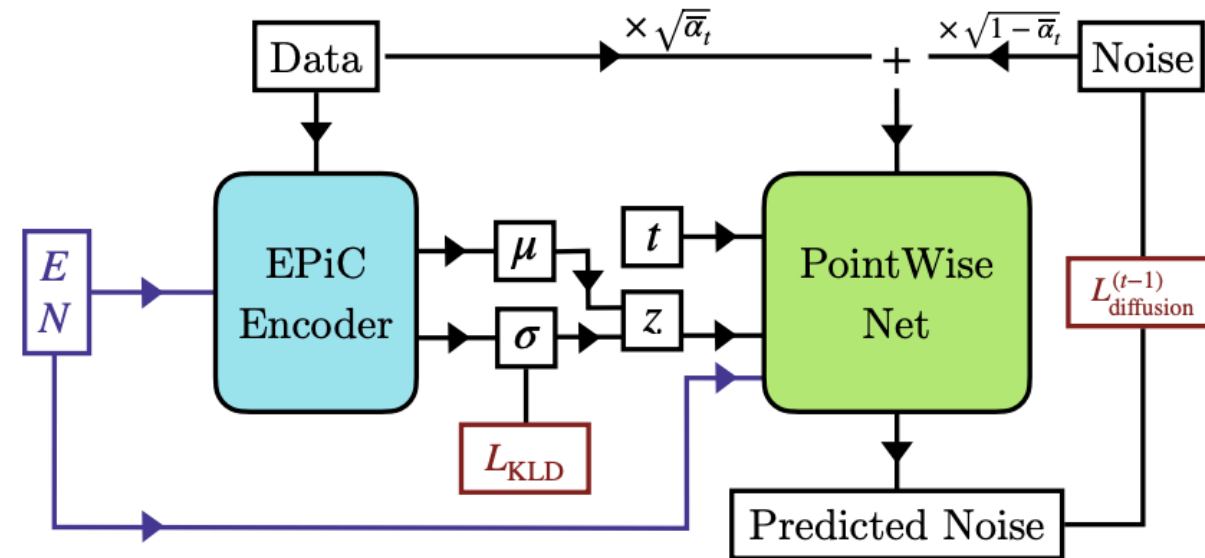


Good fidelity of distributions, very fast and linear scaling with number of particles

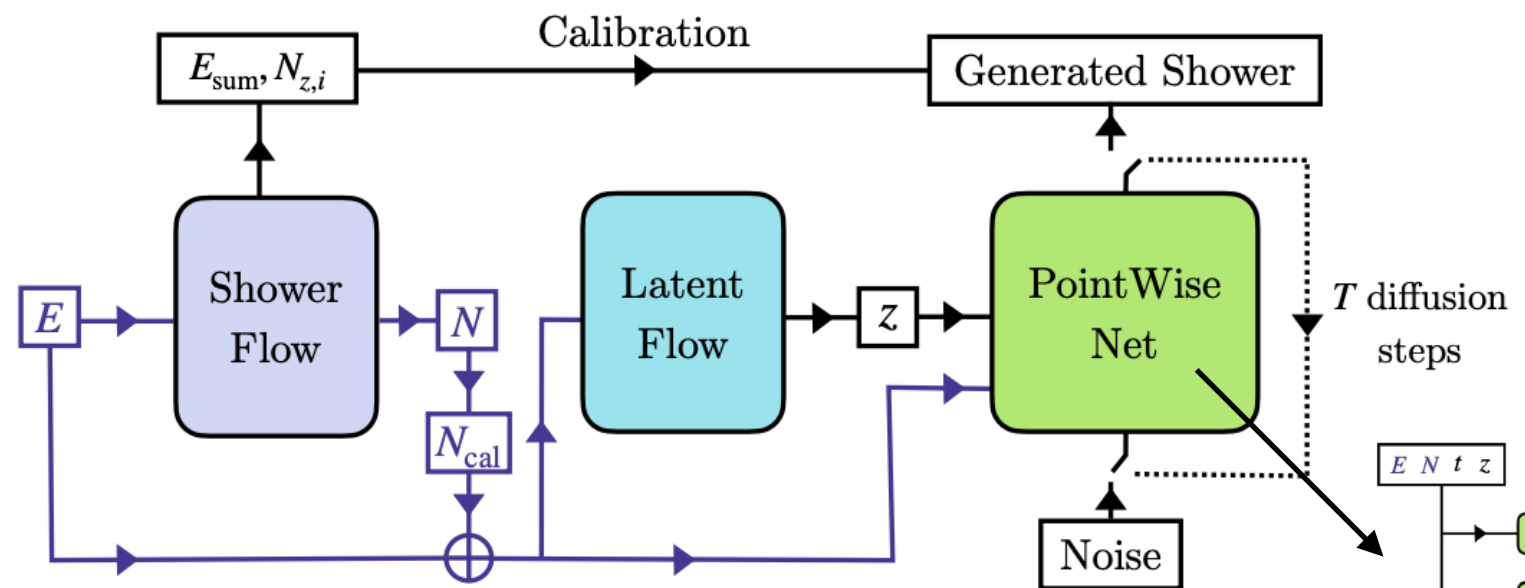
# Point Cloud Generation

To improve the generative fidelity, move from GAN to diffusion model:

Learn step-wise demonising



(a) Training at random time step  $t$



(b) Sampling with reverse diffusion through all time steps  $T$

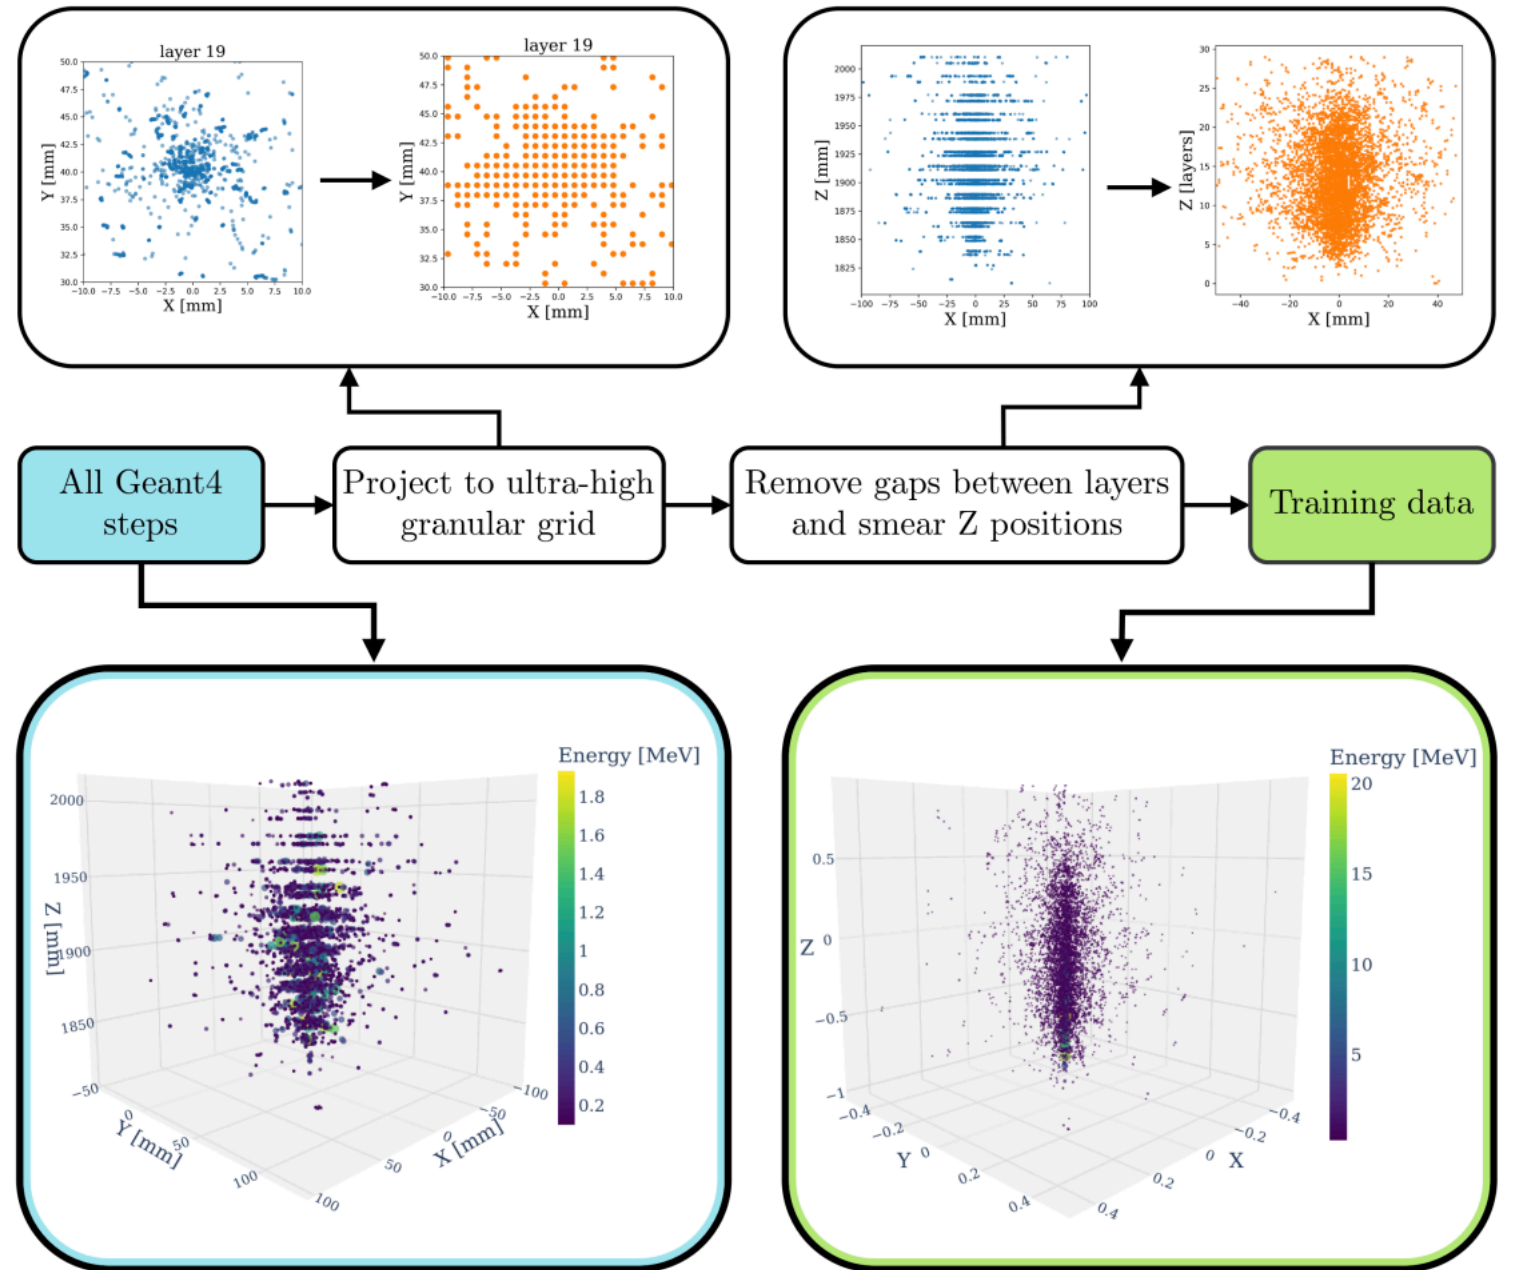
Buhmann, GK, Thaler 2301.08128;  
 Kansal et al 2106.11535; Käch et al  
 2211.13630; Buhmann, ... GK, et al  
 2305.04847

# Point Cloud Generation

To improve the generative fidelity, move from GAN to diffusion model:

Learn step-wise demonising

Some additional pre-processing needed

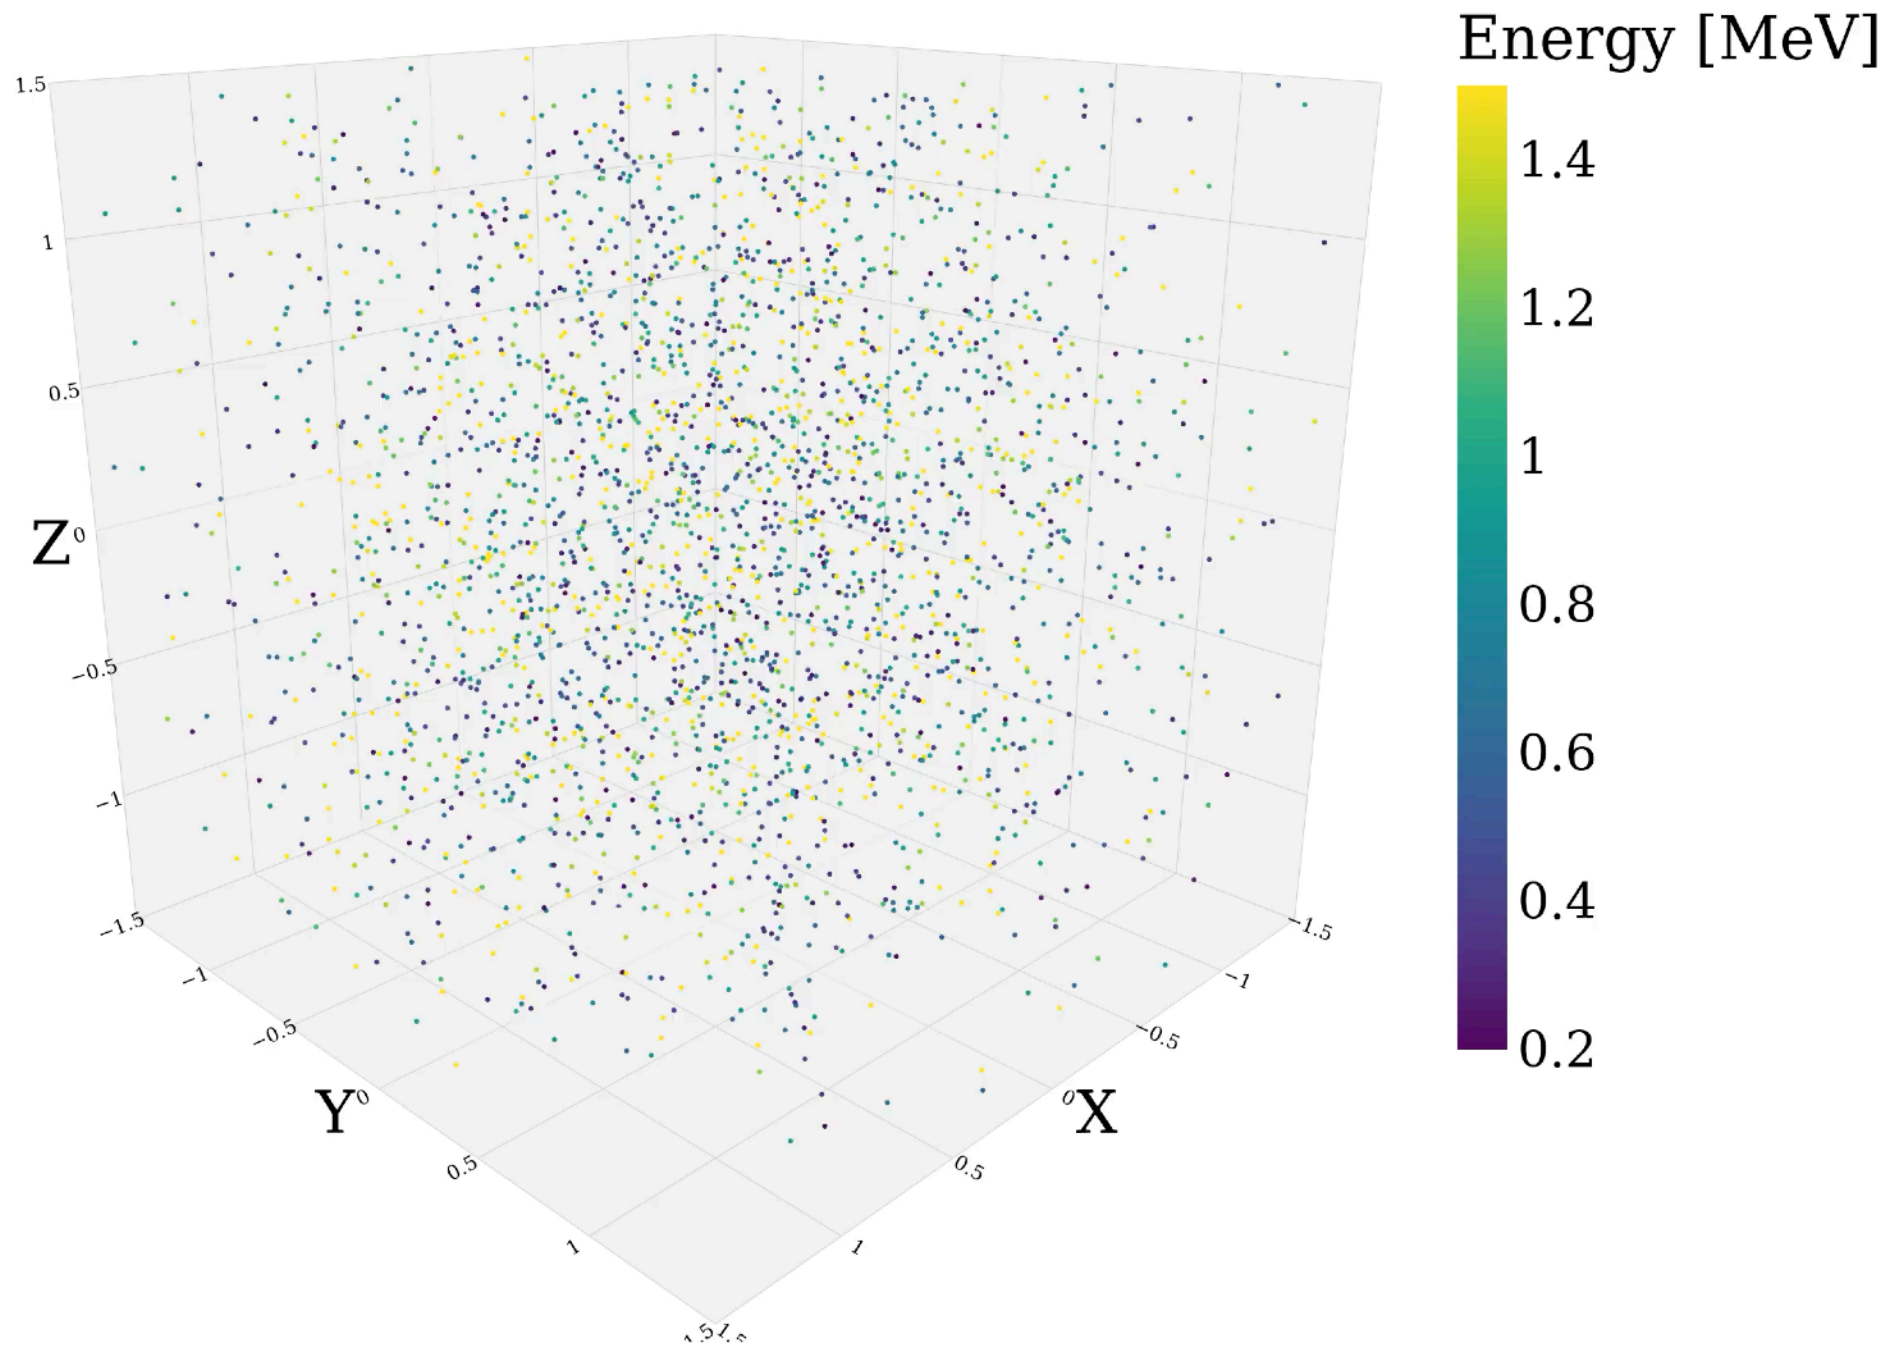


Buhmann, GK, Thaler 2301.08128;  
Kansal et al 2106.11535; Käch et al  
2211.13630; Buhmann, ... GK, et al  
2305.04847

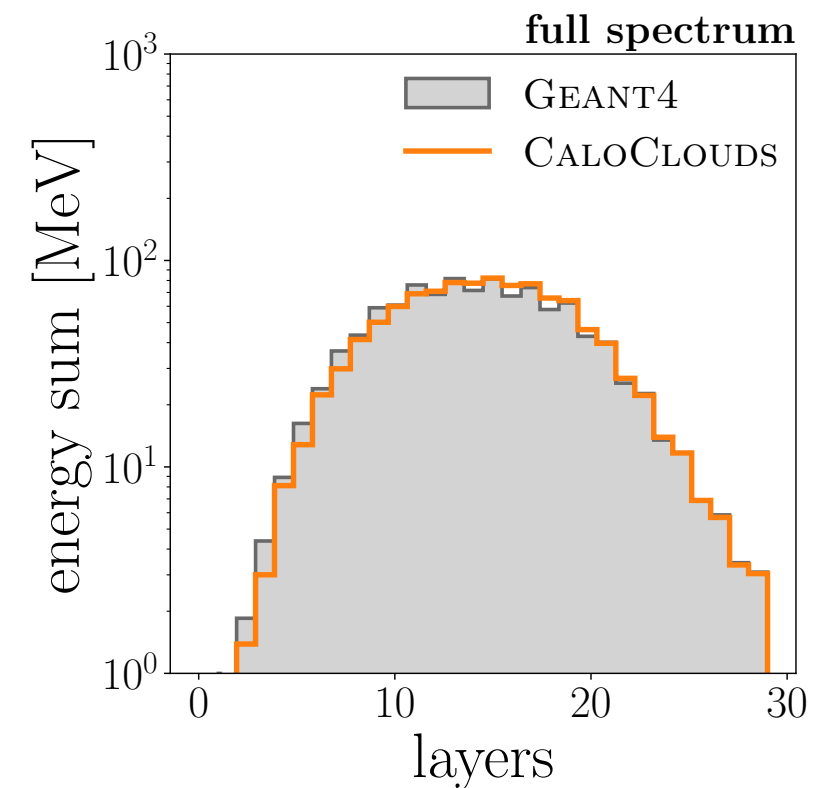
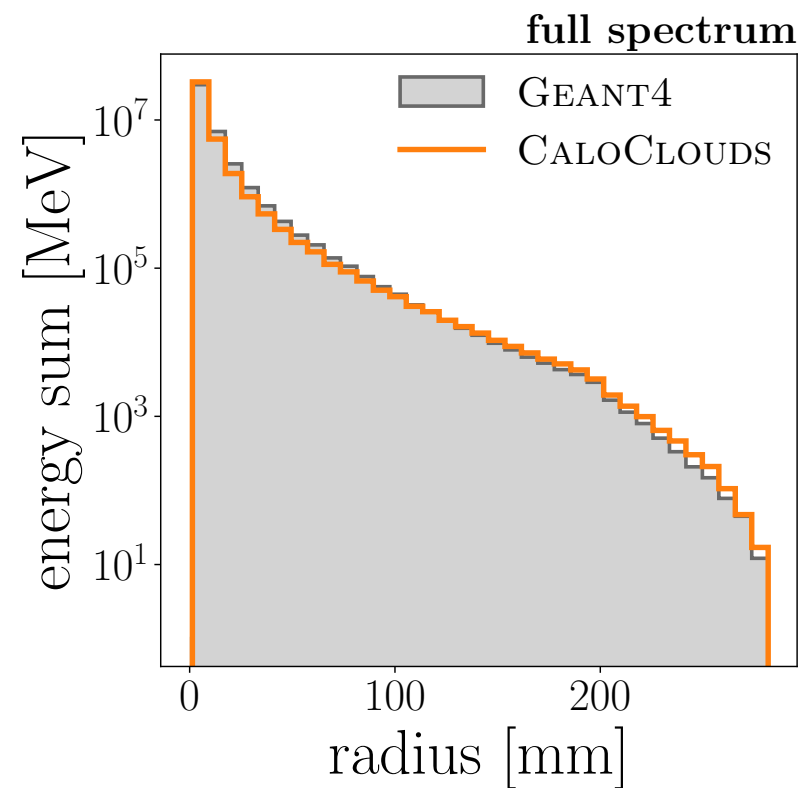
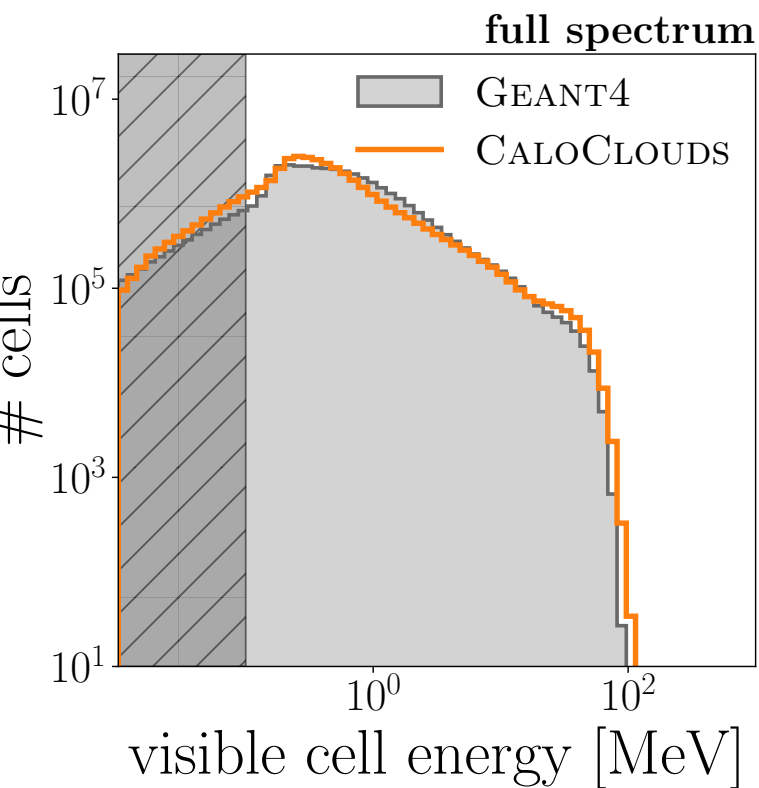


# Diffusion

CaloCloud, time stamp:  $t_{99}$



# Results



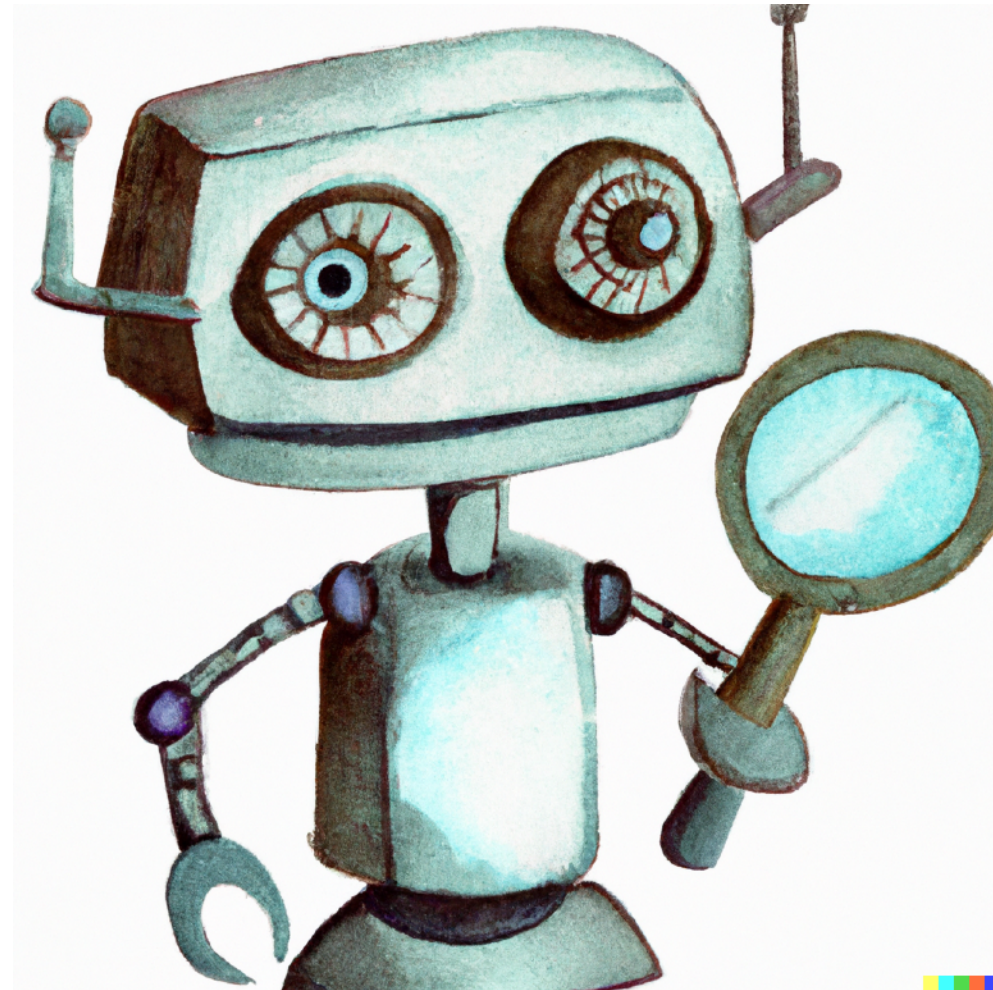
First successful point-cloud model for 6k points

Fidelity and speed can still be improved

For much more: See recent **CaloChallenge**

Hardware	Simulator	Time / Shower [ms]	Speed-up
CPU	GEANT4	4082 ± 170	×1
	CALOCLOUDS	3509 ± 220	×1.2
GPU	CALOCLOUDS	38 ± 3	×107

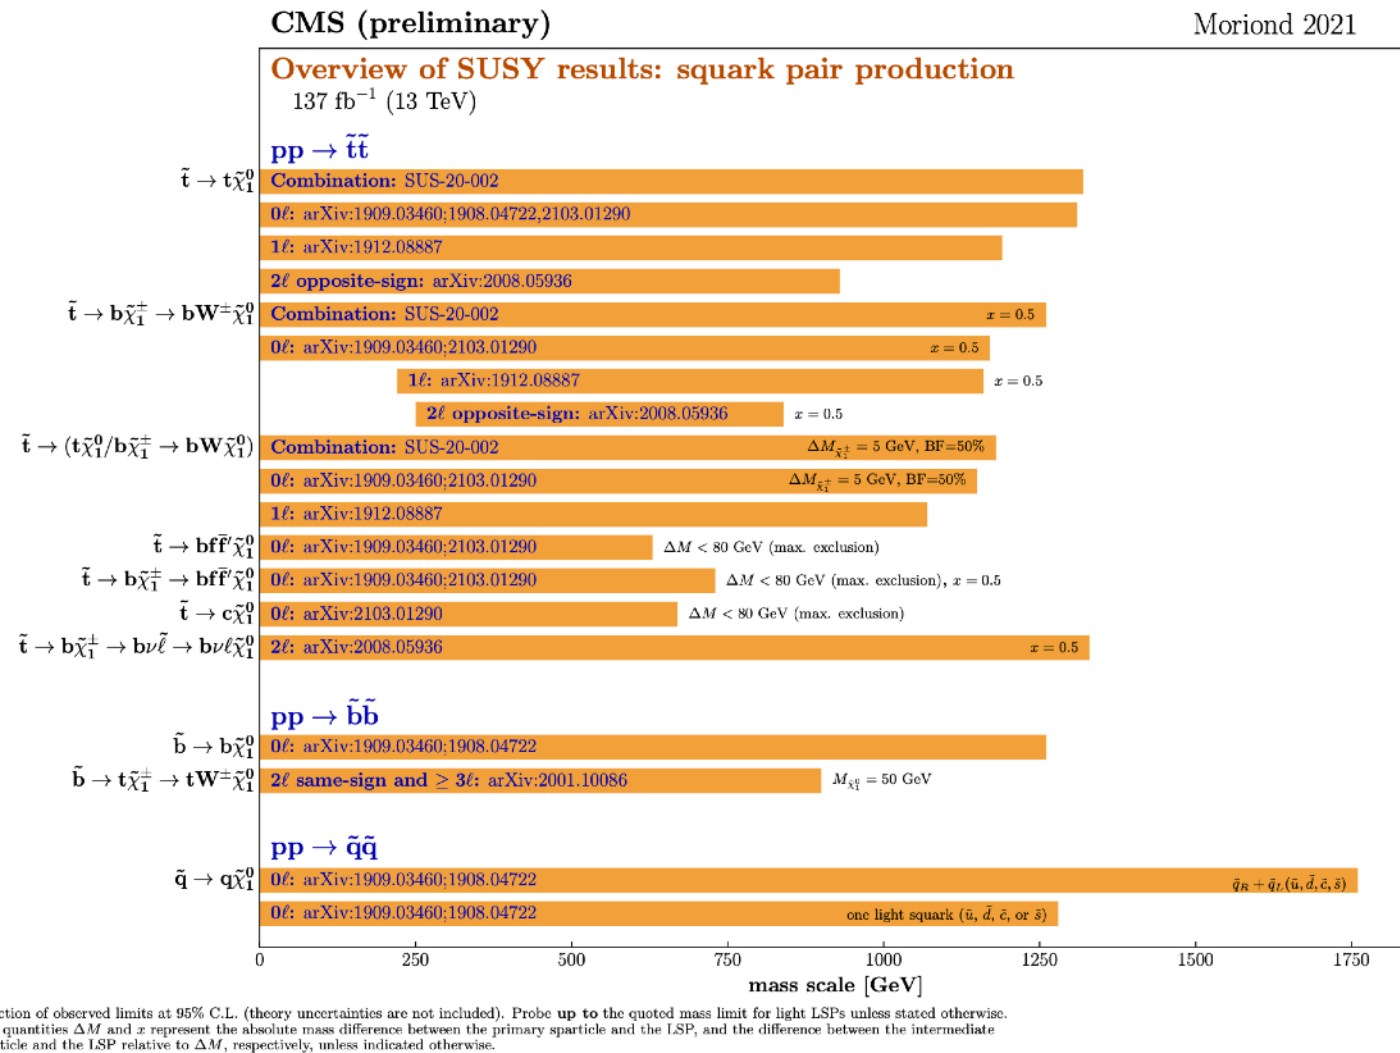
# Discovering new Physics





# Searches for new physics

- Theoretical and experimental reasons to expect new physics beyond the Standard Model
- However, so far only negative results in direct (model driven) searches
- Make sure that we do not miss potential discoveries at the LHC  
→ **Anomaly detection**

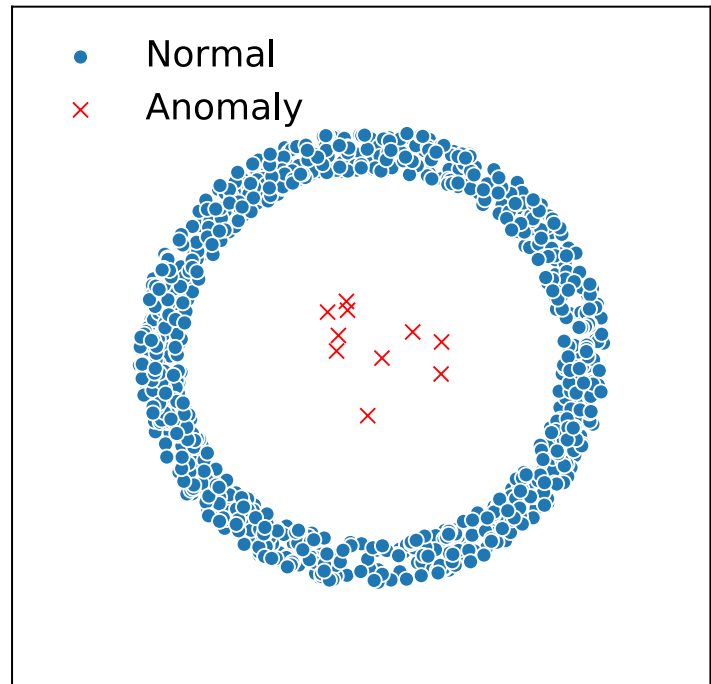
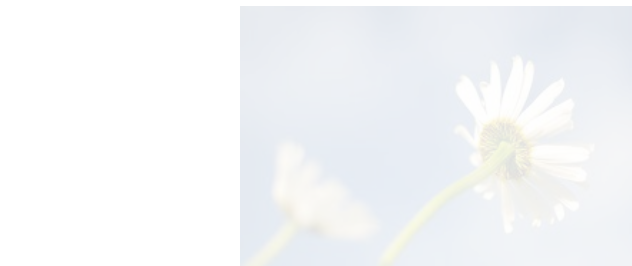
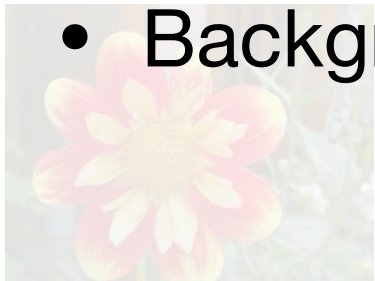


# What is an anomaly?



# Point anomaly

- **Outliers:** Datapoints far away from regular distribution
- **Examples:**
  - Detector malfunctions
  - Background-free search

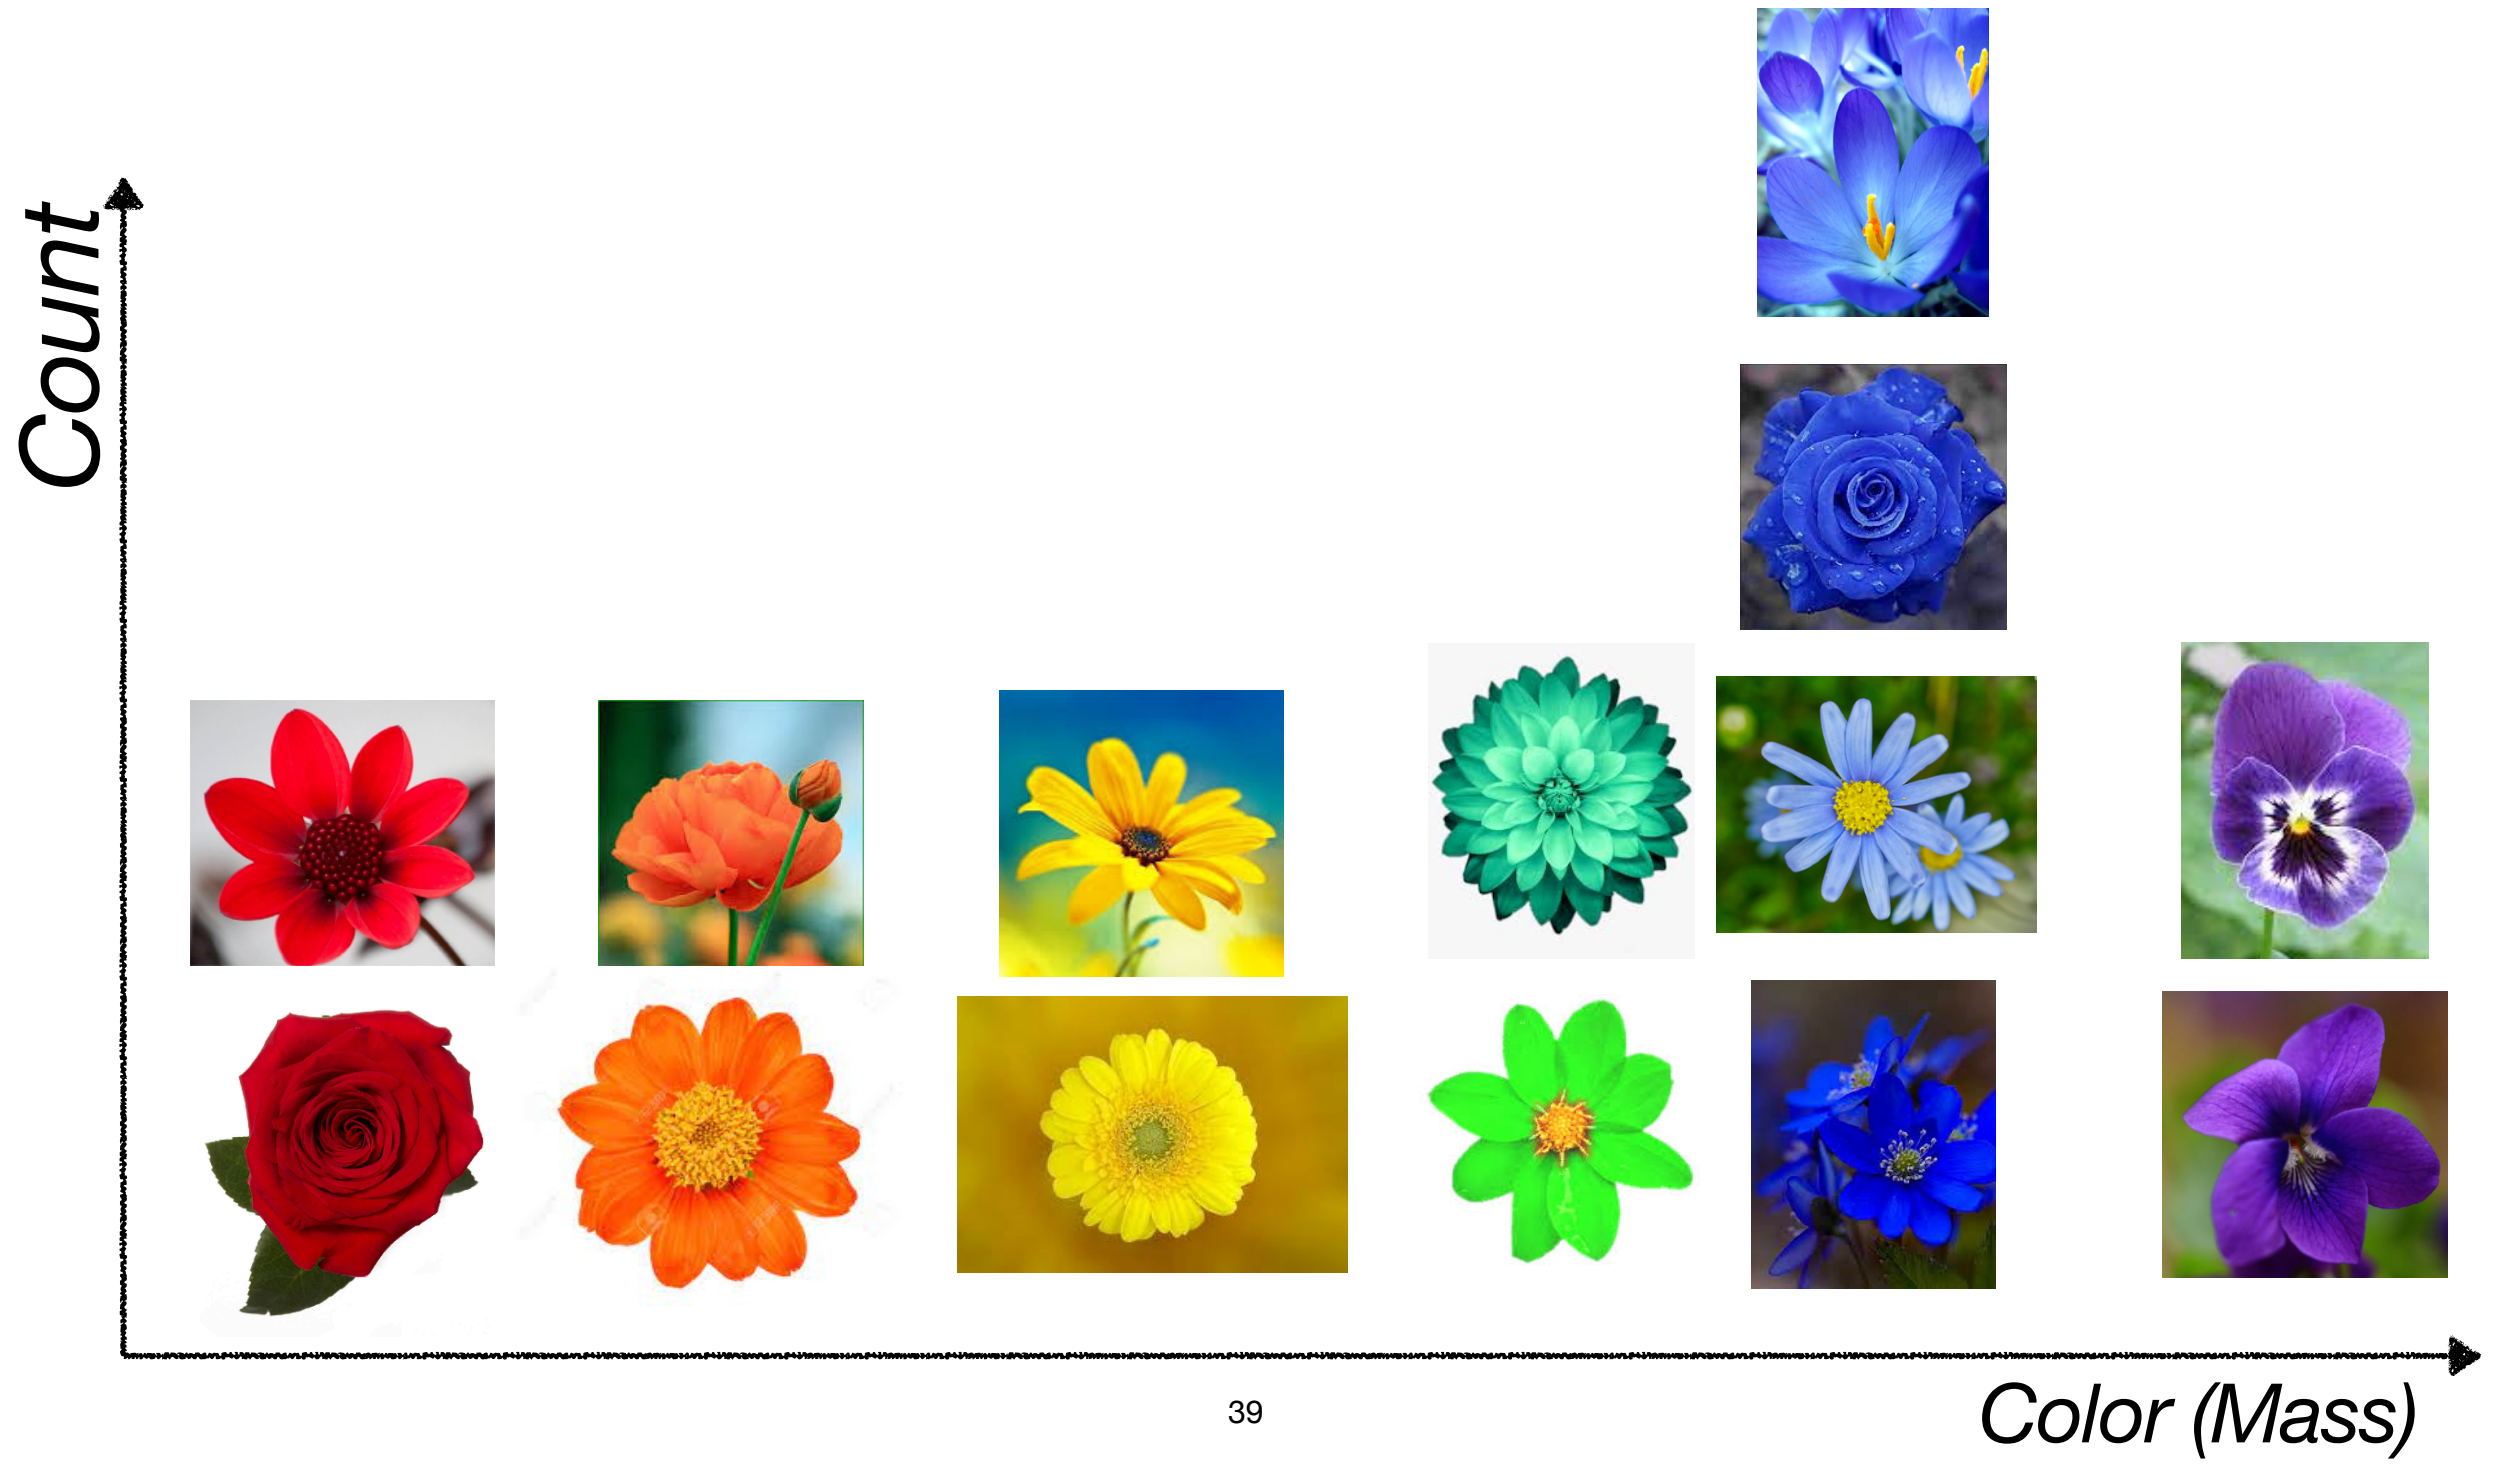




# And now?



# Group anomaly

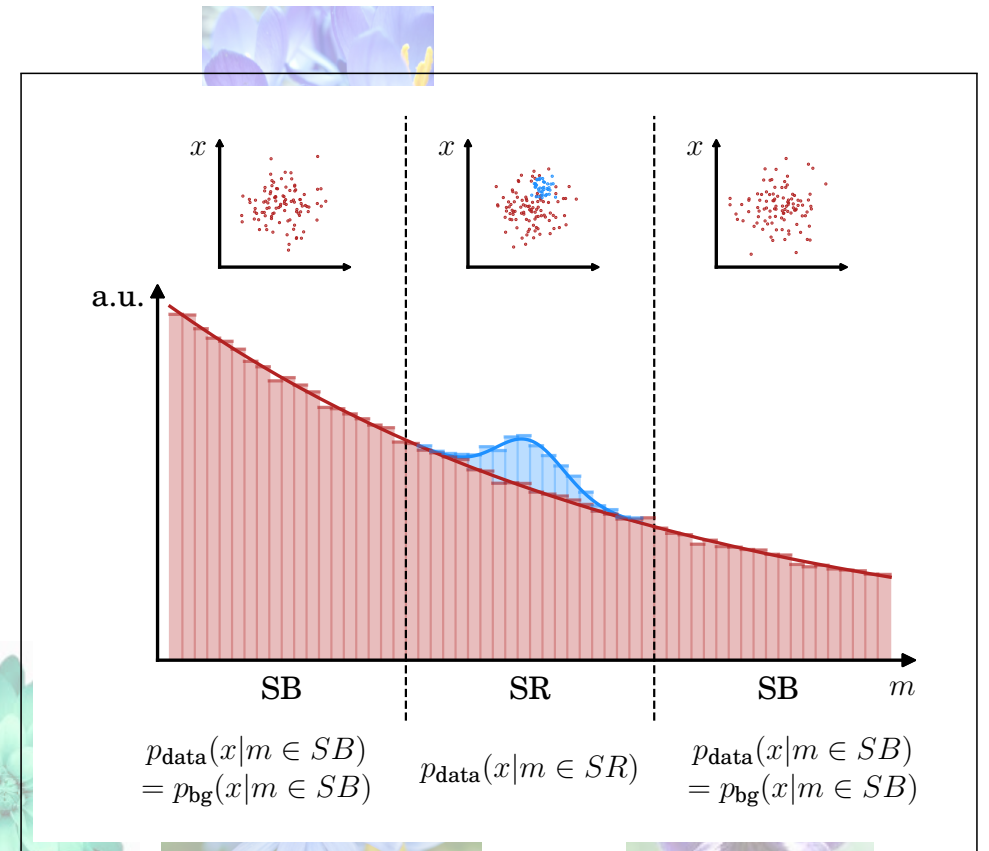
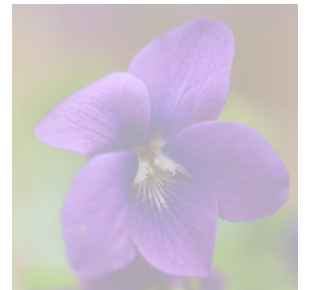
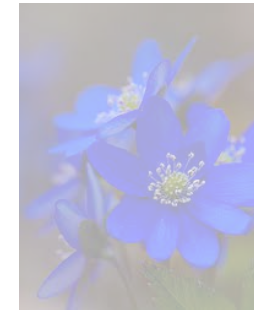
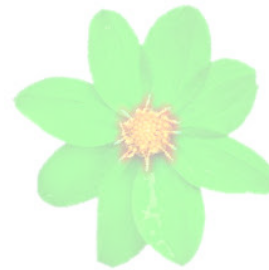
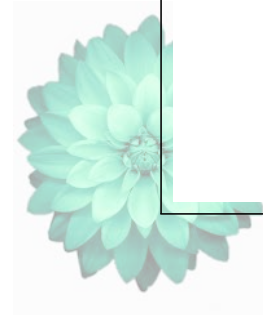
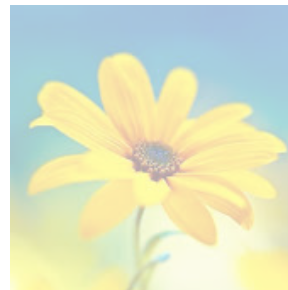
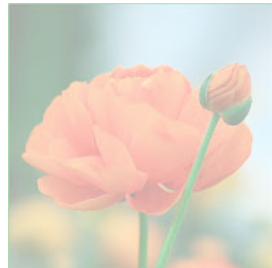




# Group anomaly

- Individual examples not anomalous, but interesting collective behaviour
- Examples:
  - New physics searches, e.g. resonances
  - Excess in time series

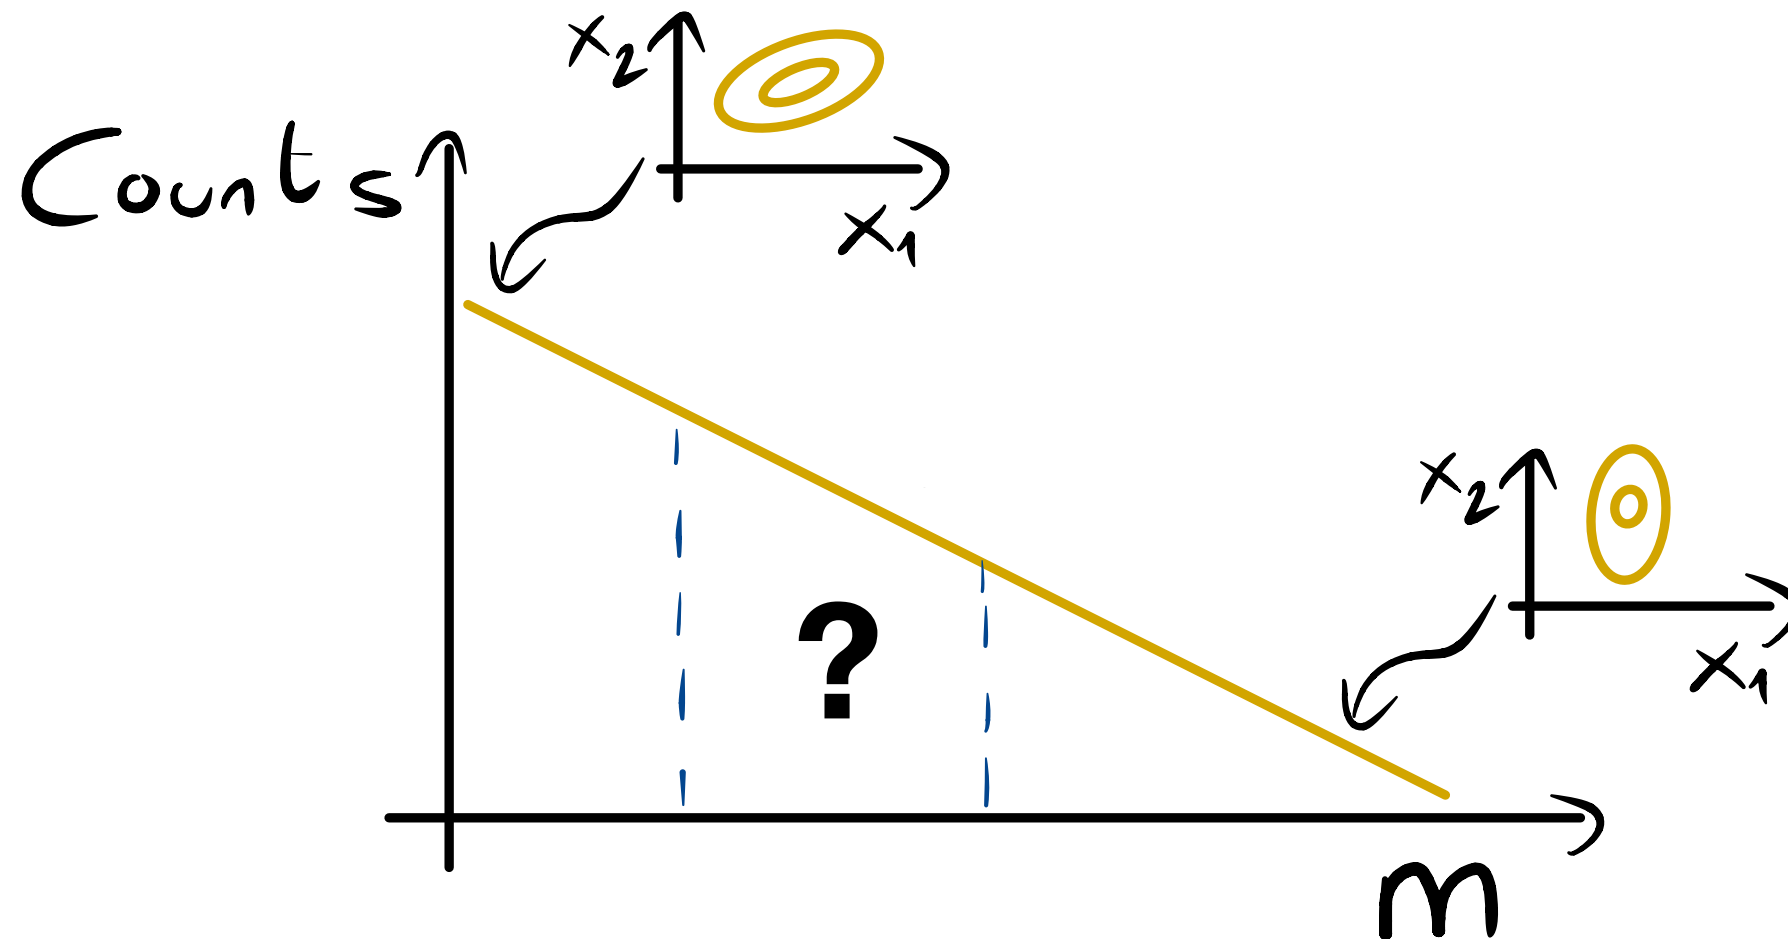
Count



Color (Mass)

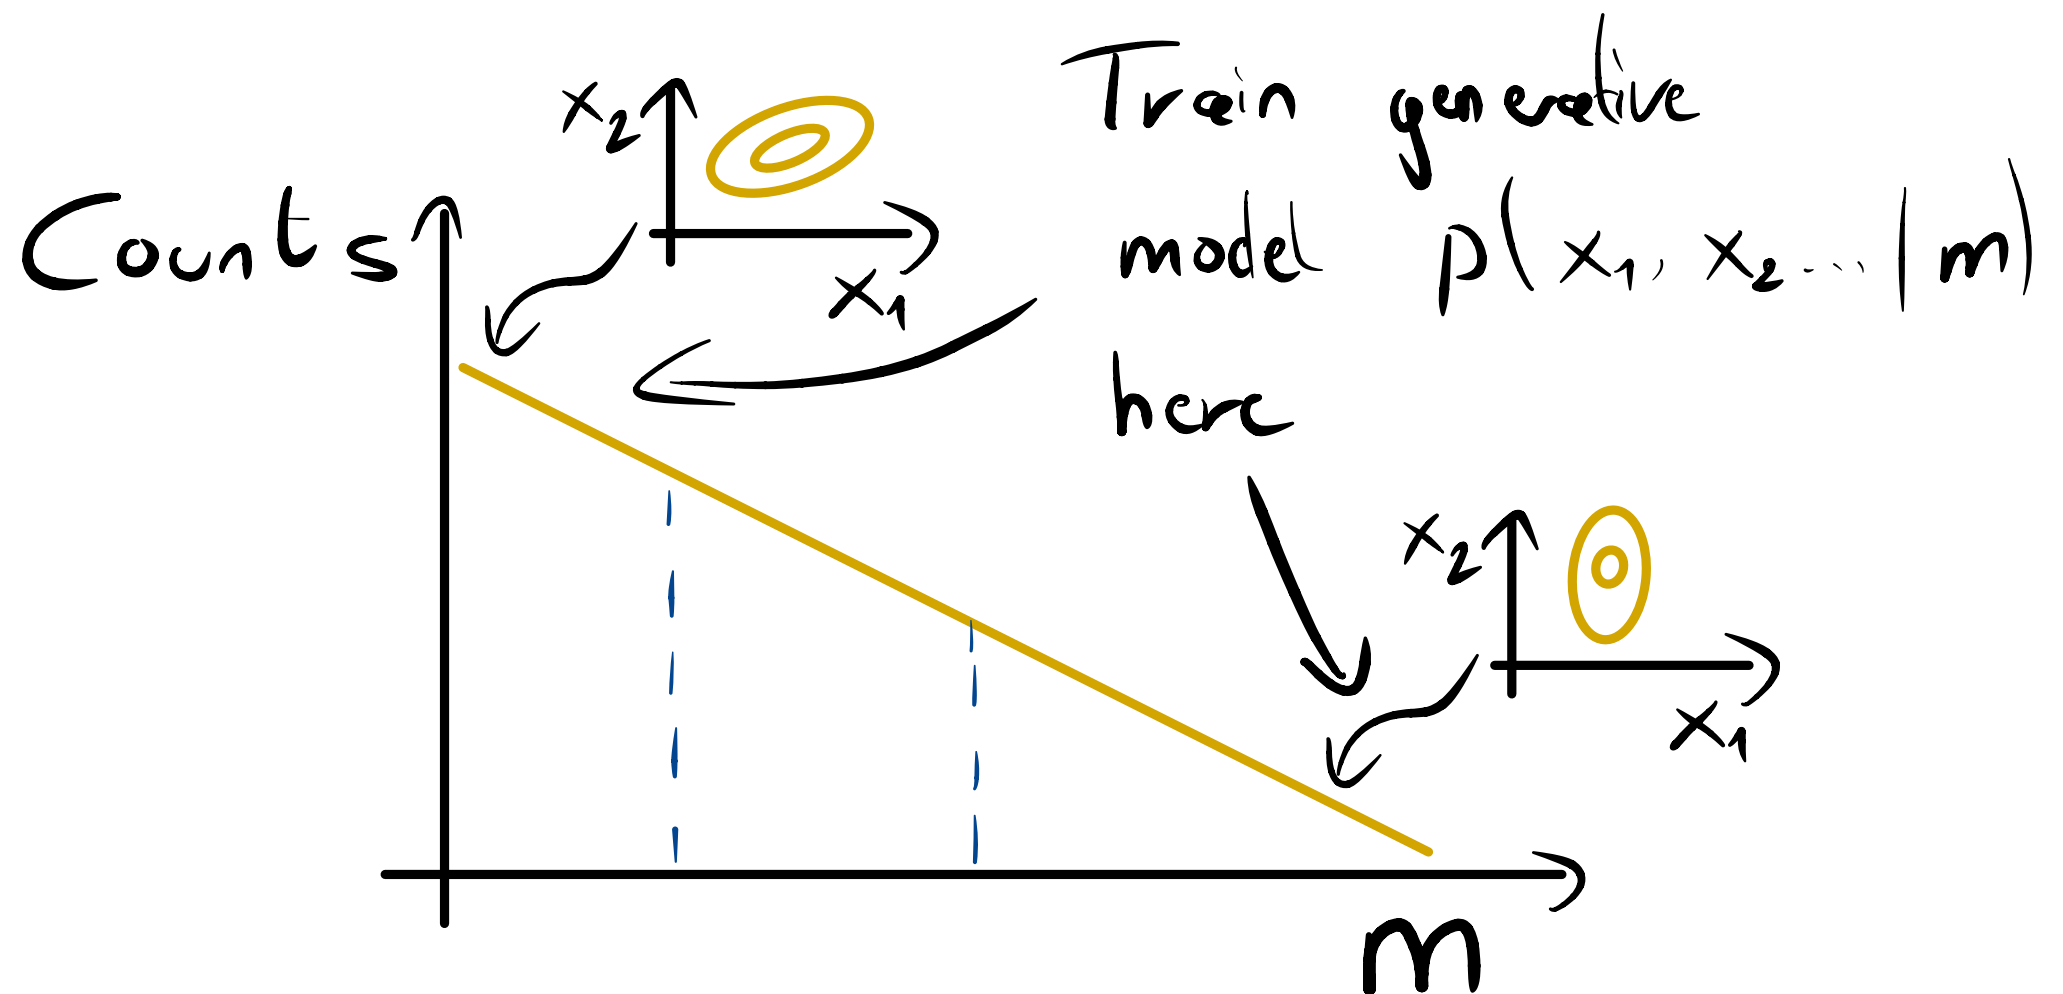


# Shape Interpolation



Can learn distributions in arbitrary dimensions (shape) as long as interpolation in one variable

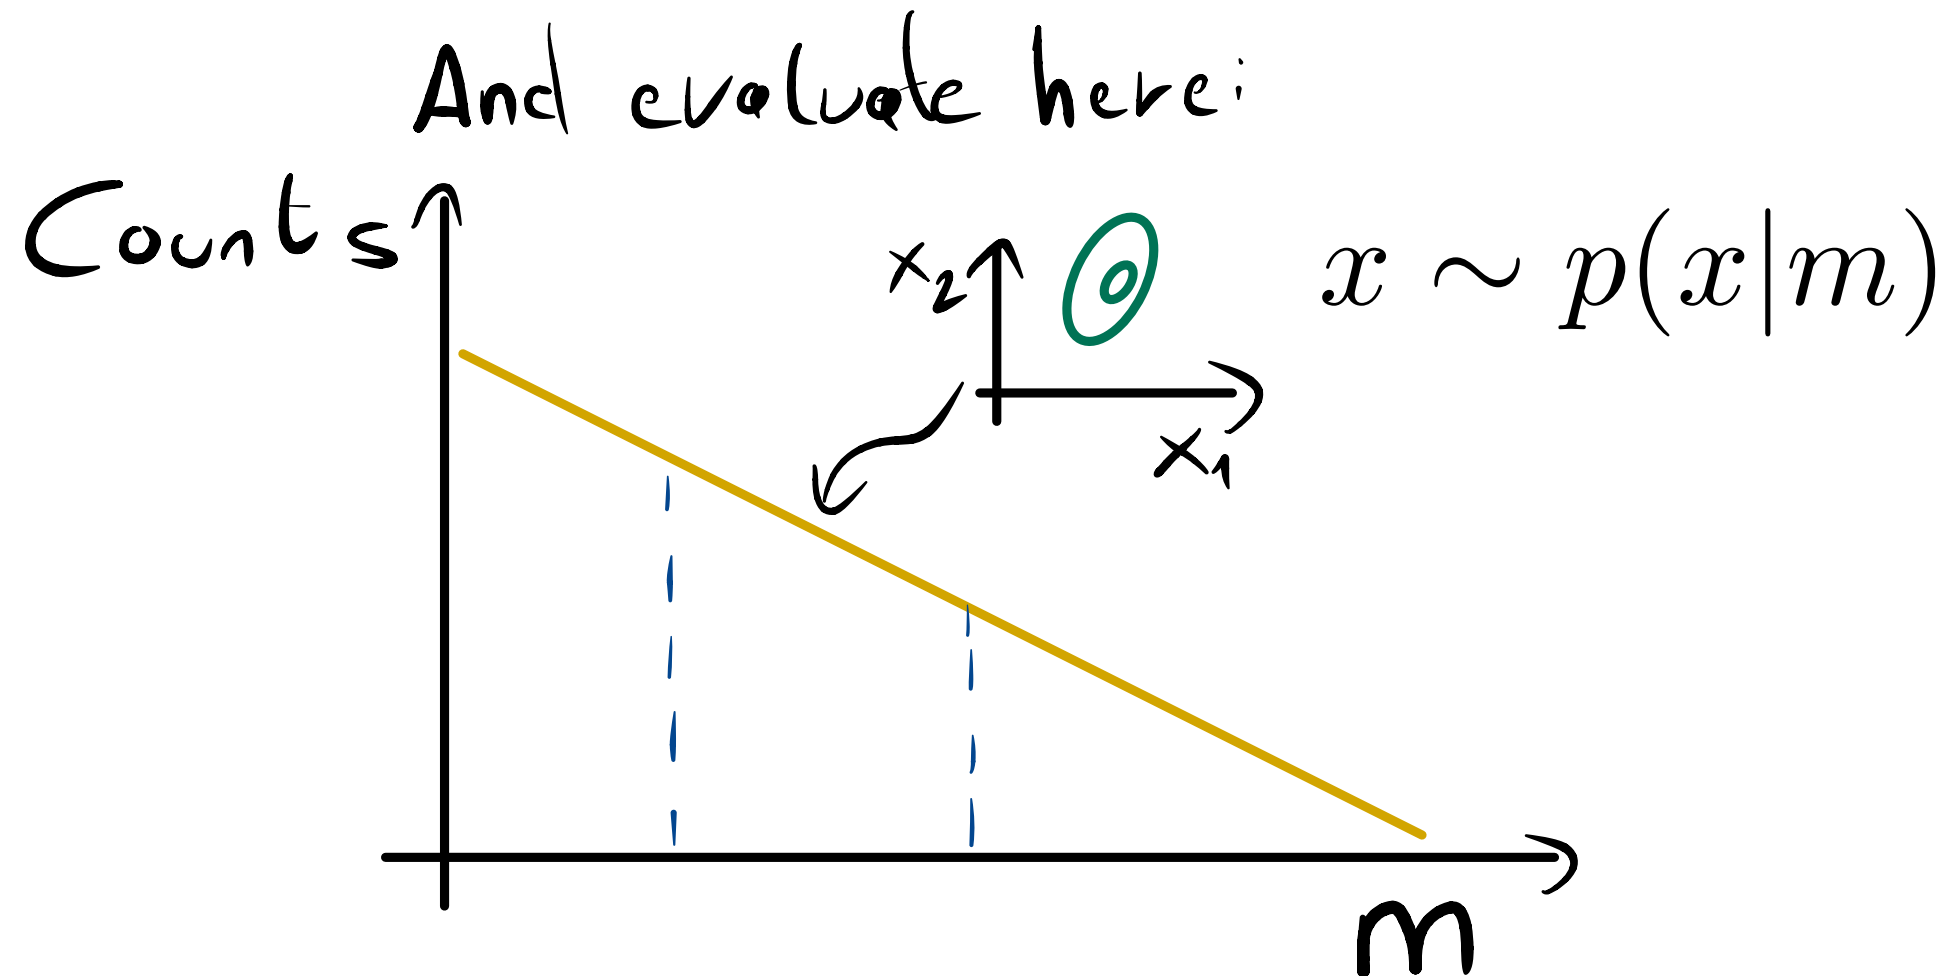
# Shape Interpolation



Can be any conditional generative model  
(GAN, VAE, flow, diffusion)

In practice we use a normalising flow

# Shape Interpolation

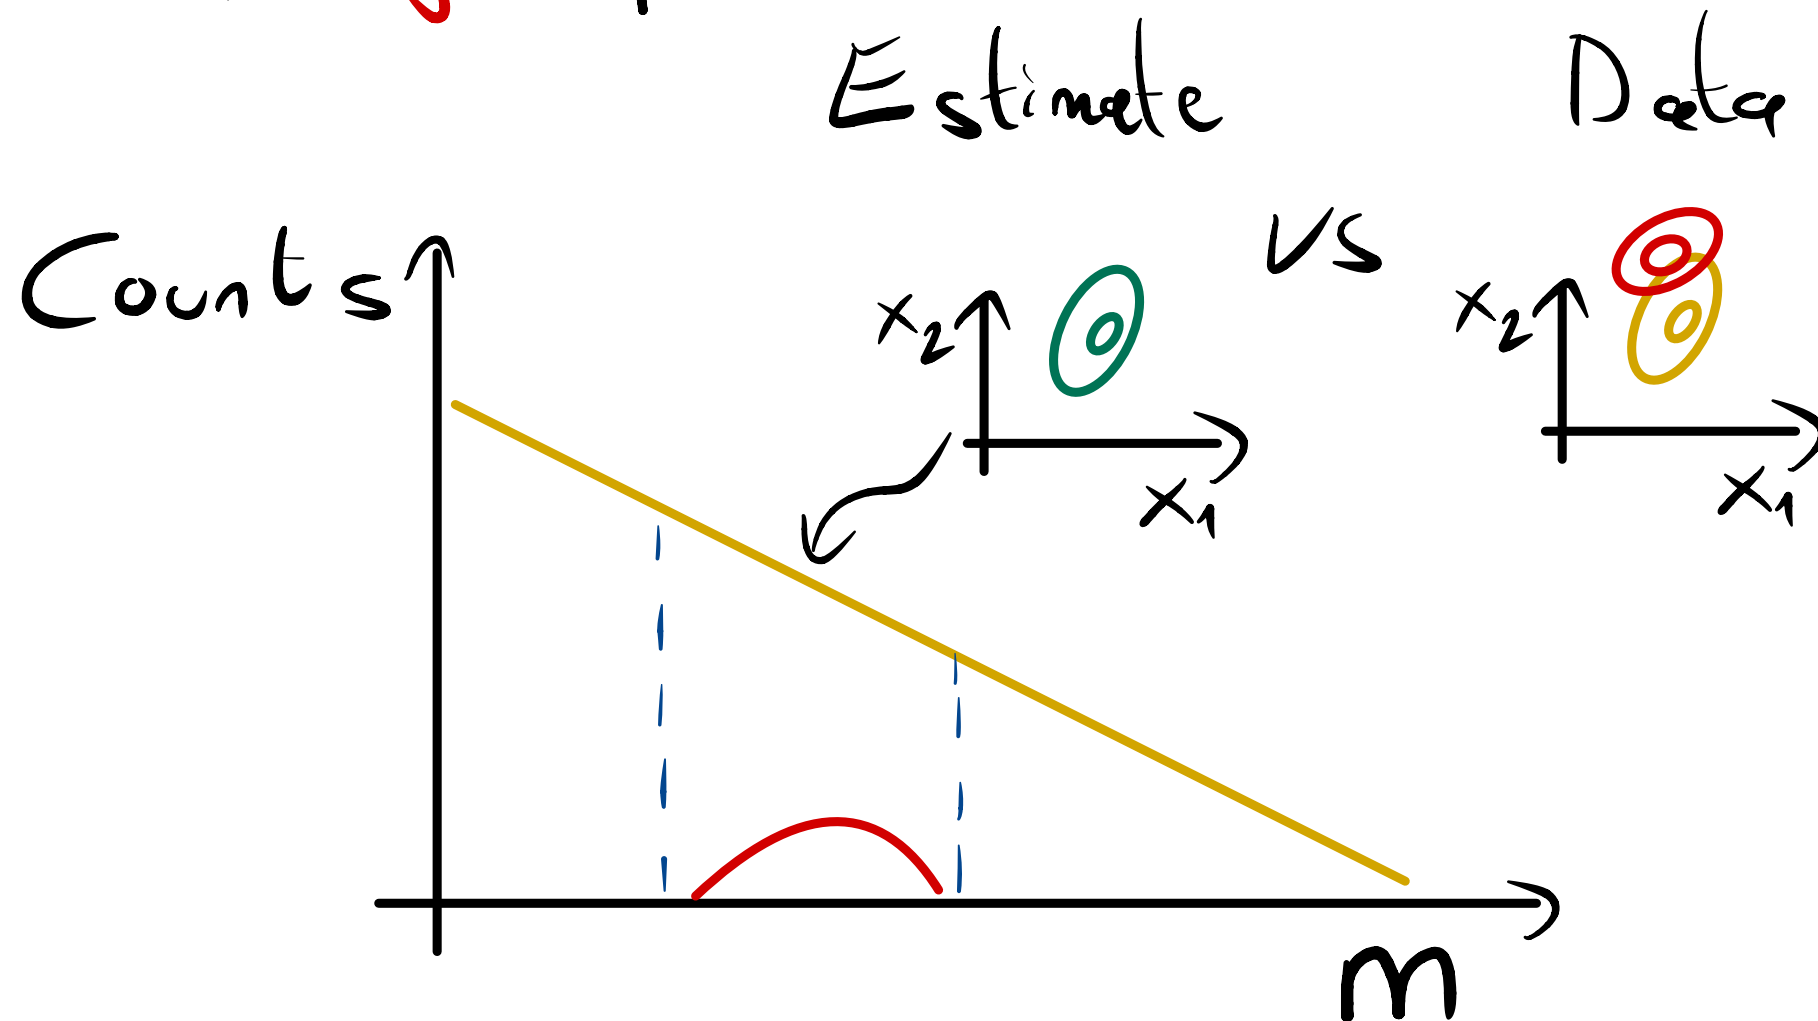


Sampling yields an **in-situ background estimate trained on data**



# Anomalies!

If *signal* present:



Training a binary classifier between estimate and data can detect localised signals (e.g. resonant new physics)

# Aside: Justification

Per Neyman-Pearson: Likelihood-ratio  
is optimal test statistic  
*Unfortunately,  $p(x|\text{anomaly})$  is not  
available*

$$L_{S/B} = \frac{p(x|\text{anomaly})}{p(x|\text{normal})}$$

Build data/background ratio:

$$L_{D/B} = \frac{p(x)}{p(x|\text{normal})}$$

Approximate background density using  
control measurement (e.g. sideband)

$$L_{D/B} \approx \frac{p(x)}{\tilde{p}(x|\text{normal})}$$

Expand

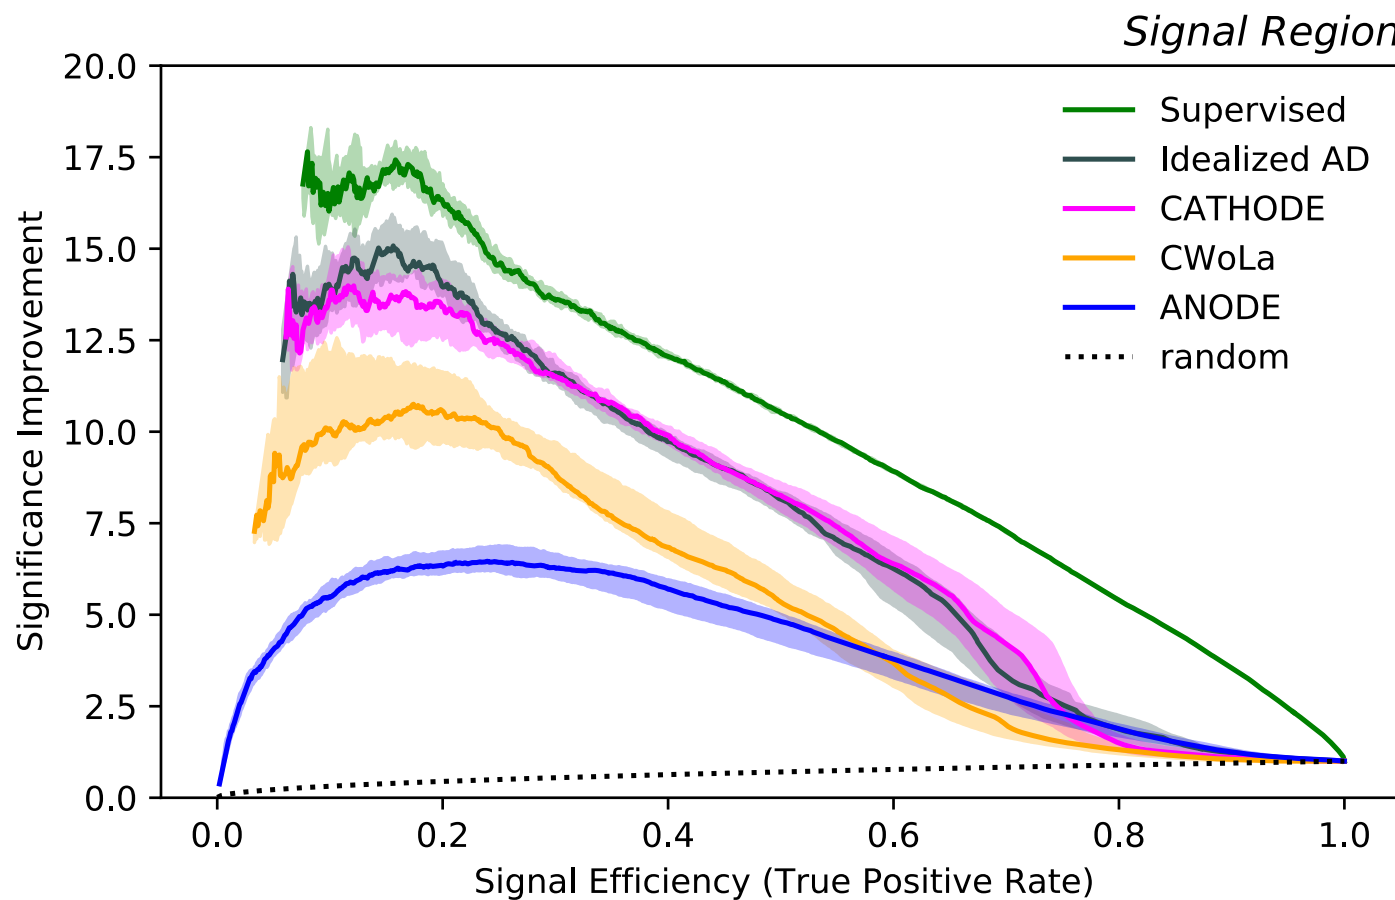
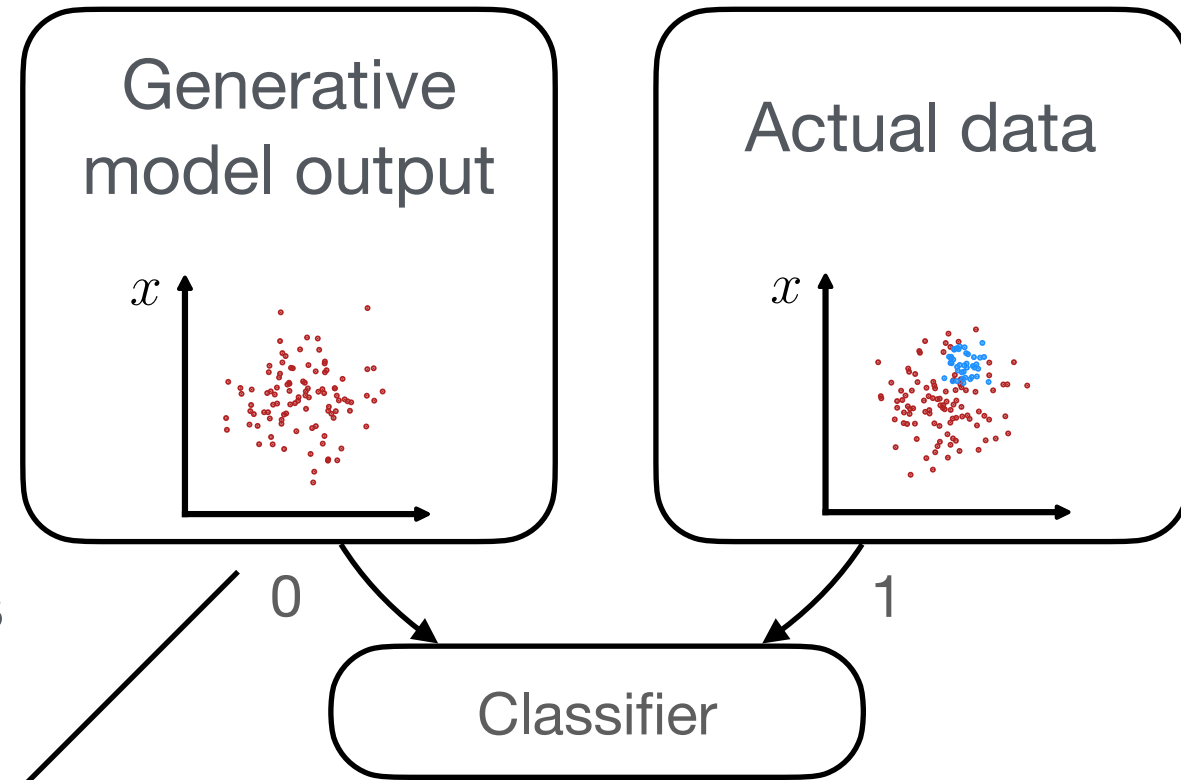
$$p(x) = f_{\text{normal}} p(x|\text{normal}) + f_{\text{anomaly}} p(x|\text{anomaly})$$

And insert:

$$L_{D/B} \approx f_{\text{normal}} + f_{\text{anomaly}} \frac{p(x|\text{anomaly})}{\tilde{p}(x|\text{normal})}$$

# CATHODE & Friends

Cut on classifier output to identify possible anomalies

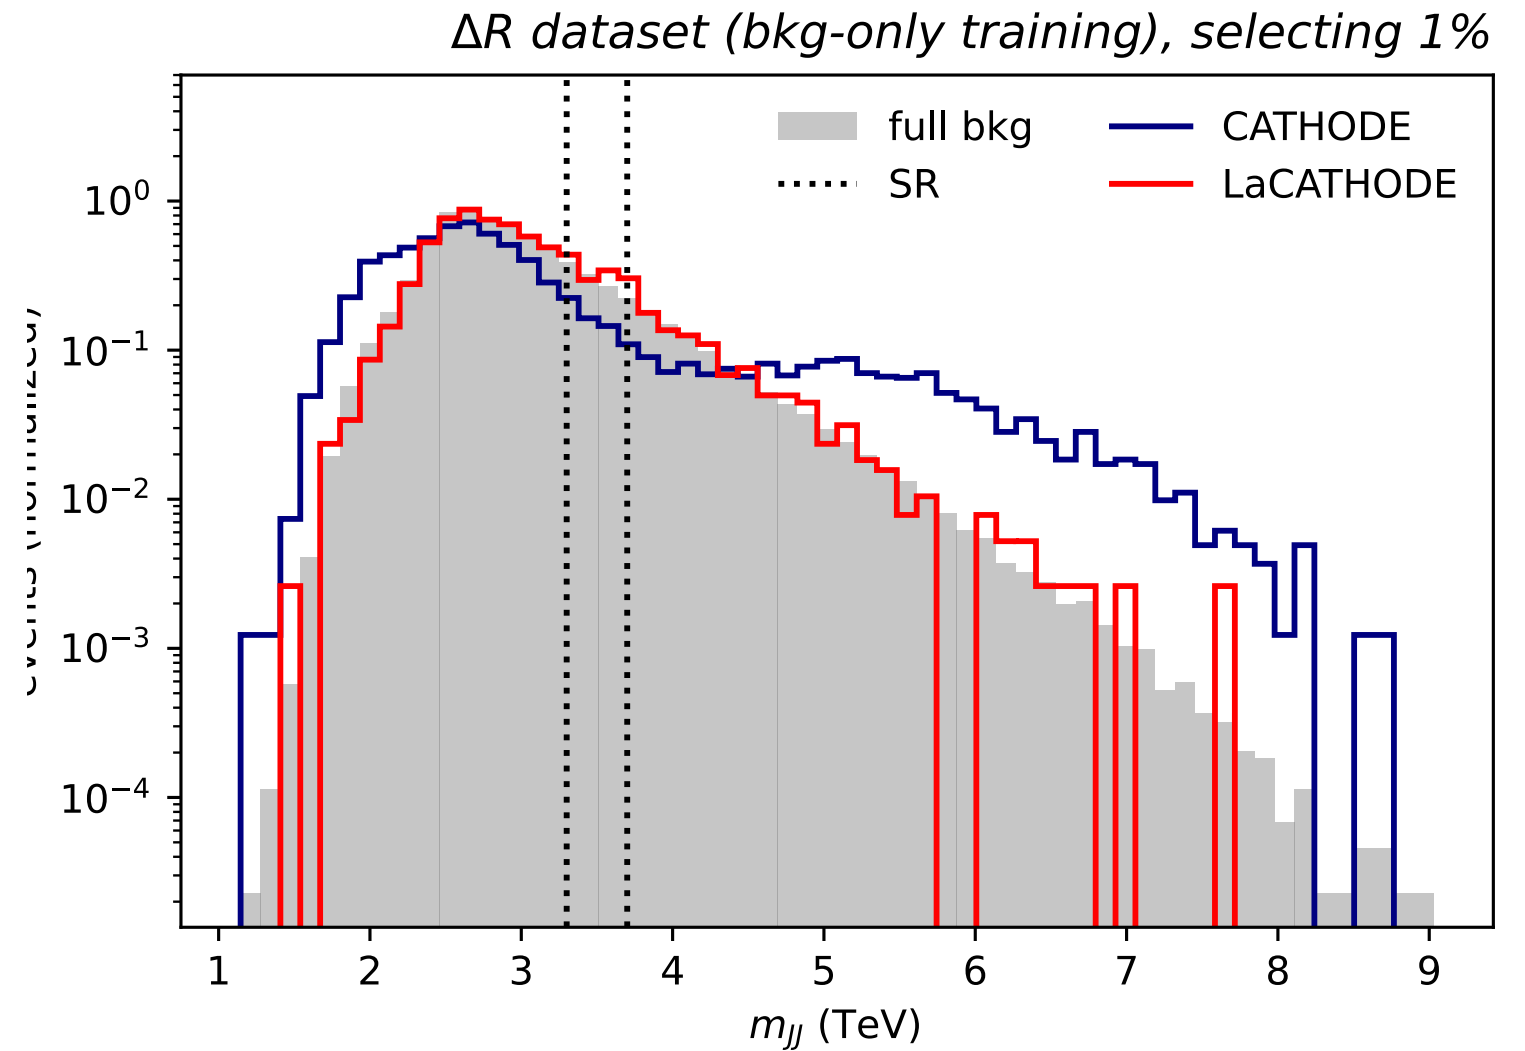


- Most sensitive current resonant anomaly detection technique: CATHODE
- Stability improvement (LaCATHODE)
- Experimental application: XXXXXXXXXX



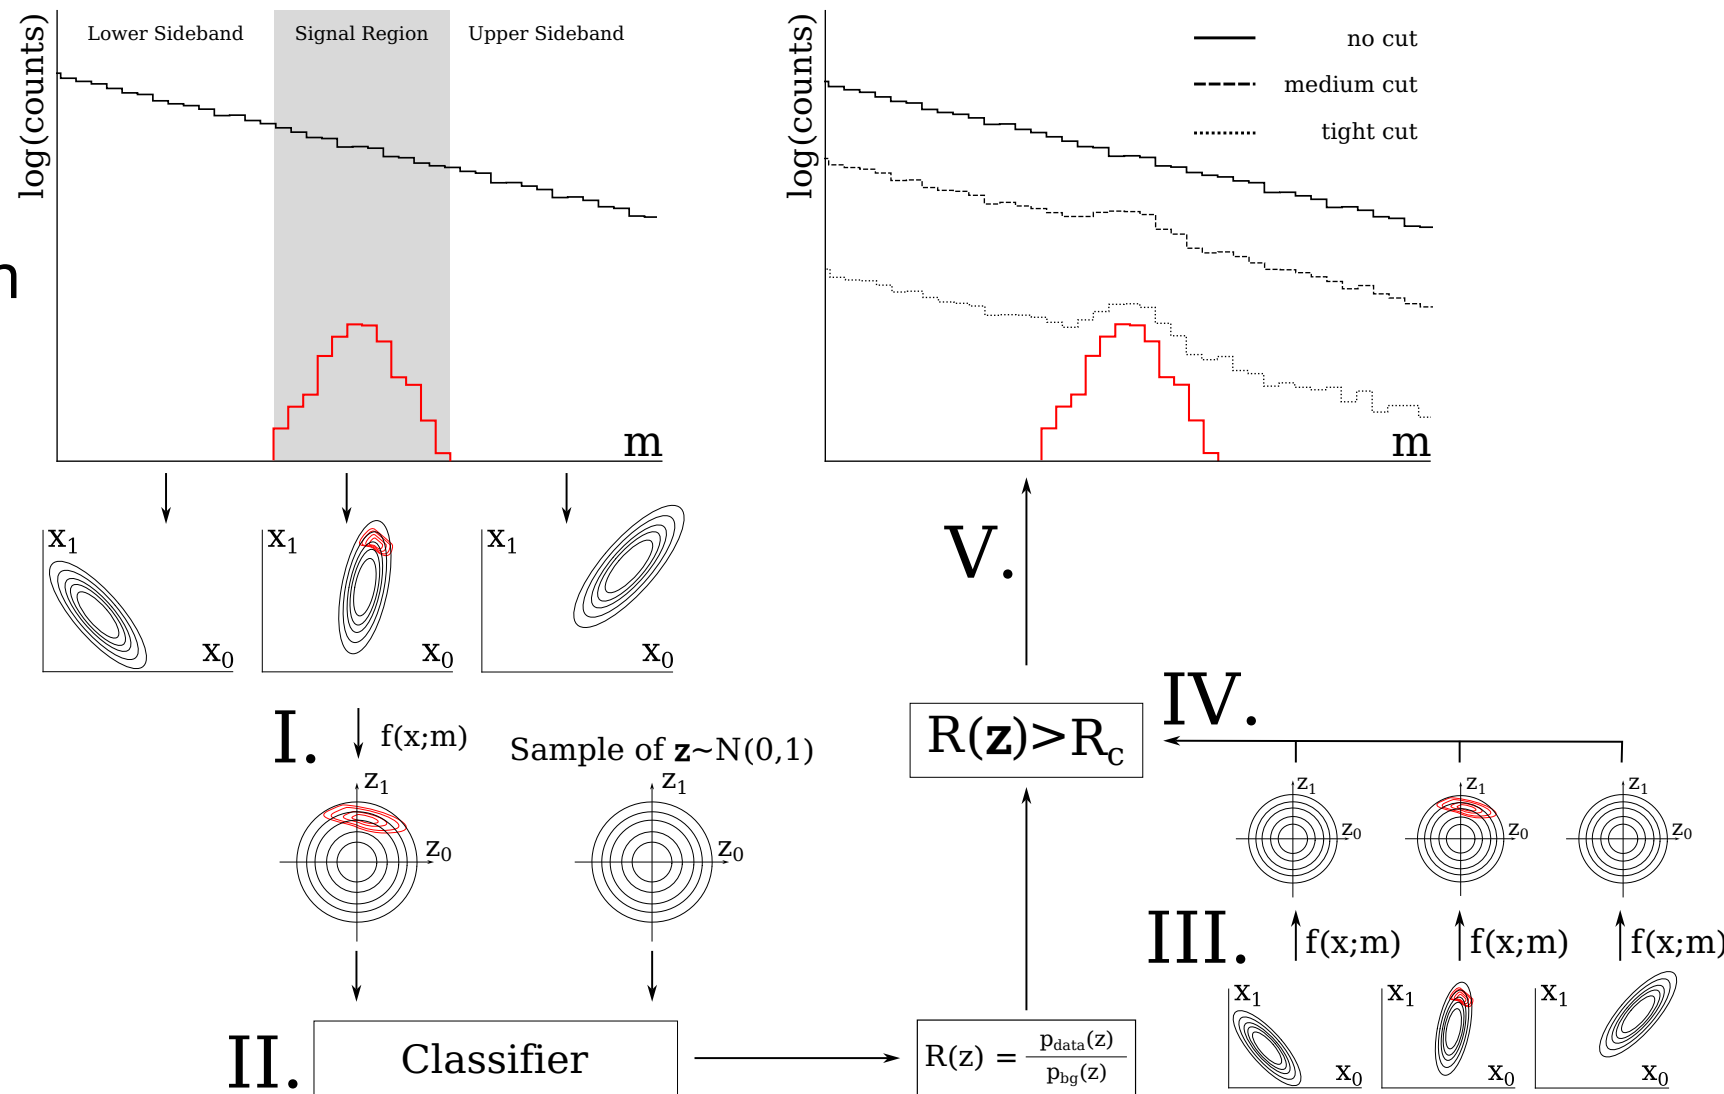
# LaCATHODE

- If  $R(x)$  is only calculated in signal region, its extrapolation is not well-defined
- Potential problem for bump-hunt if it shapes distributions



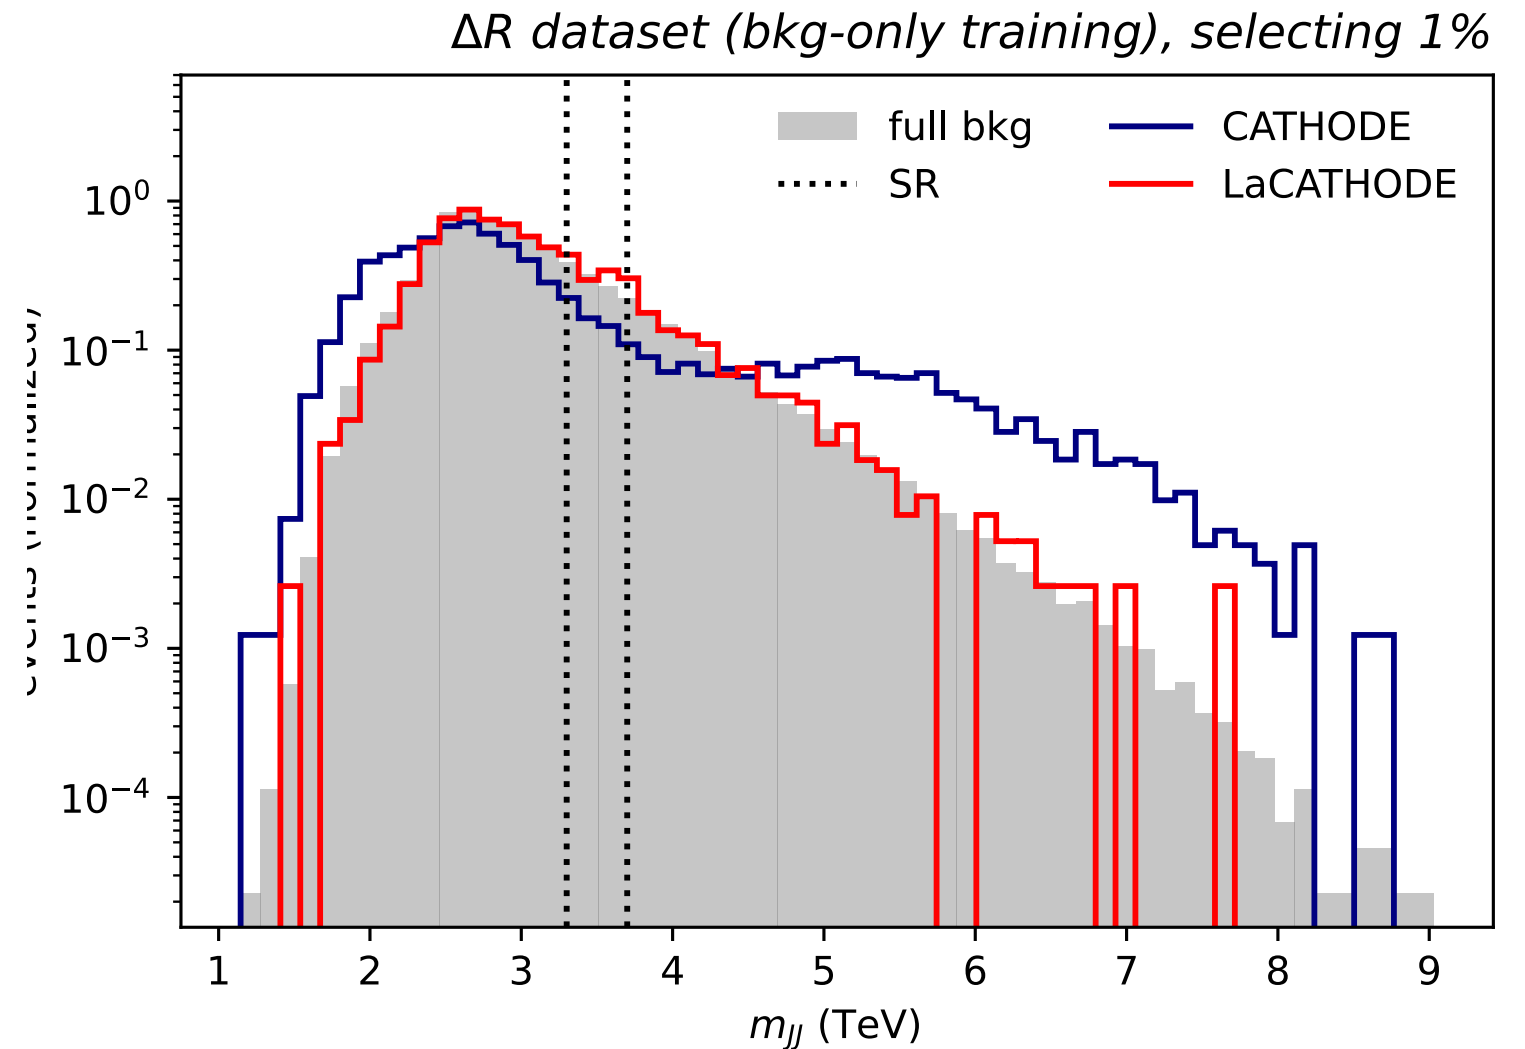
# LaCATHODE

- If  $R(x)$  is only calculated in signal region, it's extrapolation is not well-defined
- Potential problem for bump-hunt if it shapes distributions
- Instead, train classifier in latent space



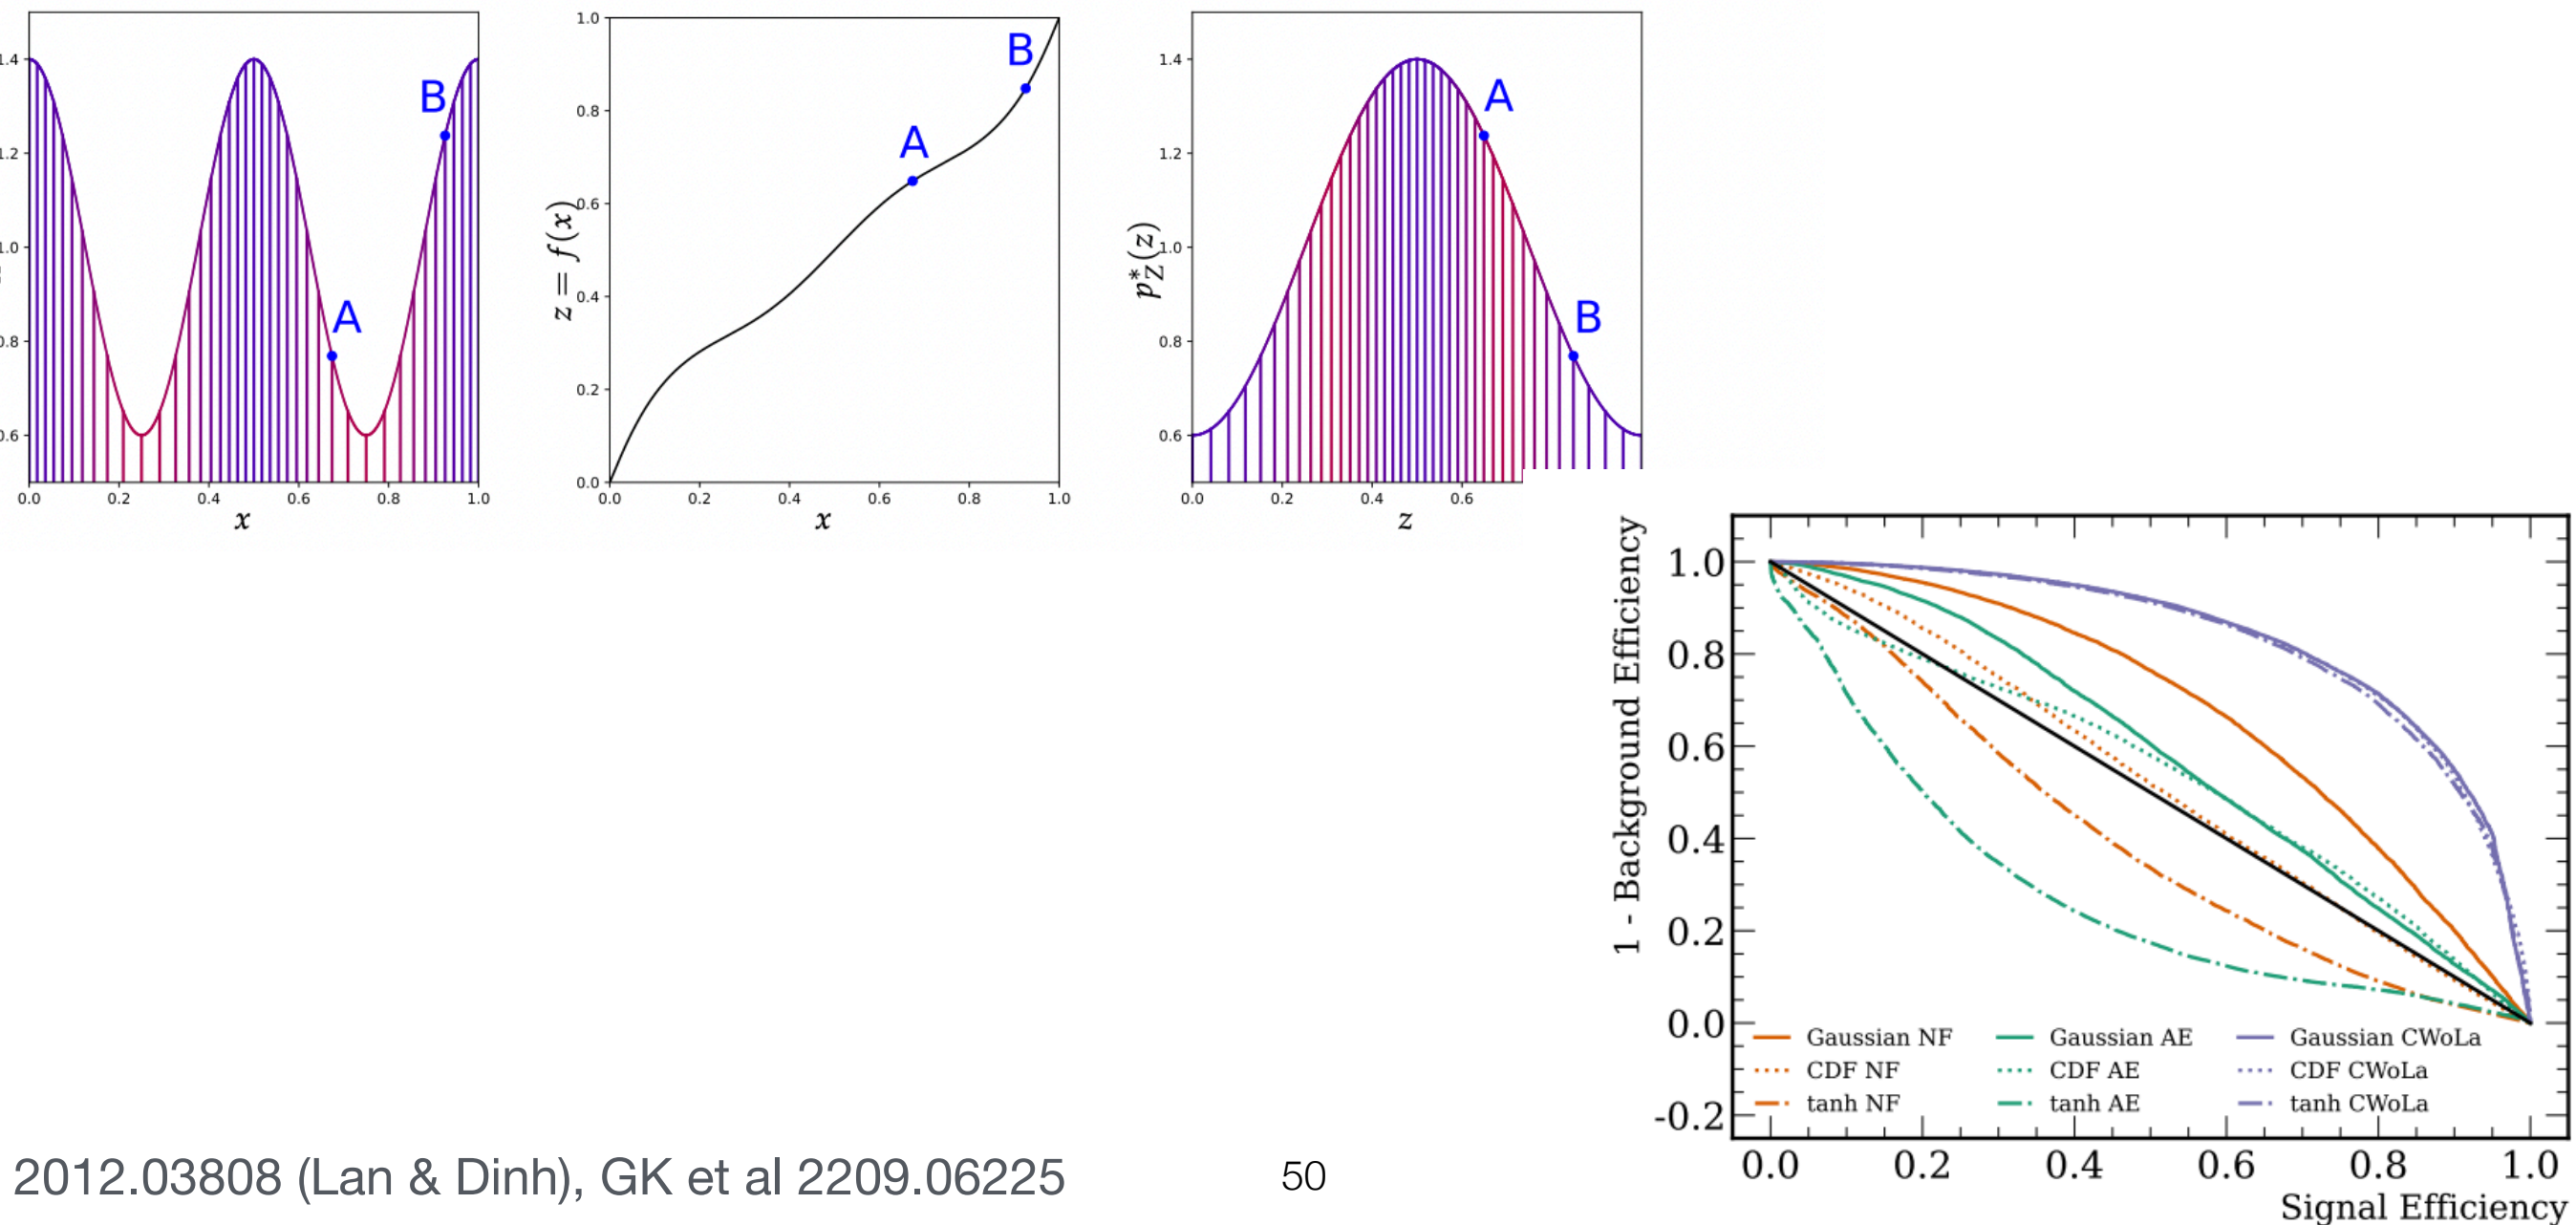
# LaCATHODE

- If  $R(x)$  is only calculated in signal region, its extrapolation is not well-defined
- Potential problem for bump-hunt if it shapes distributions
- Instead, train classifier in latent space to achieve flat distributions



# Comments on anomaly detection

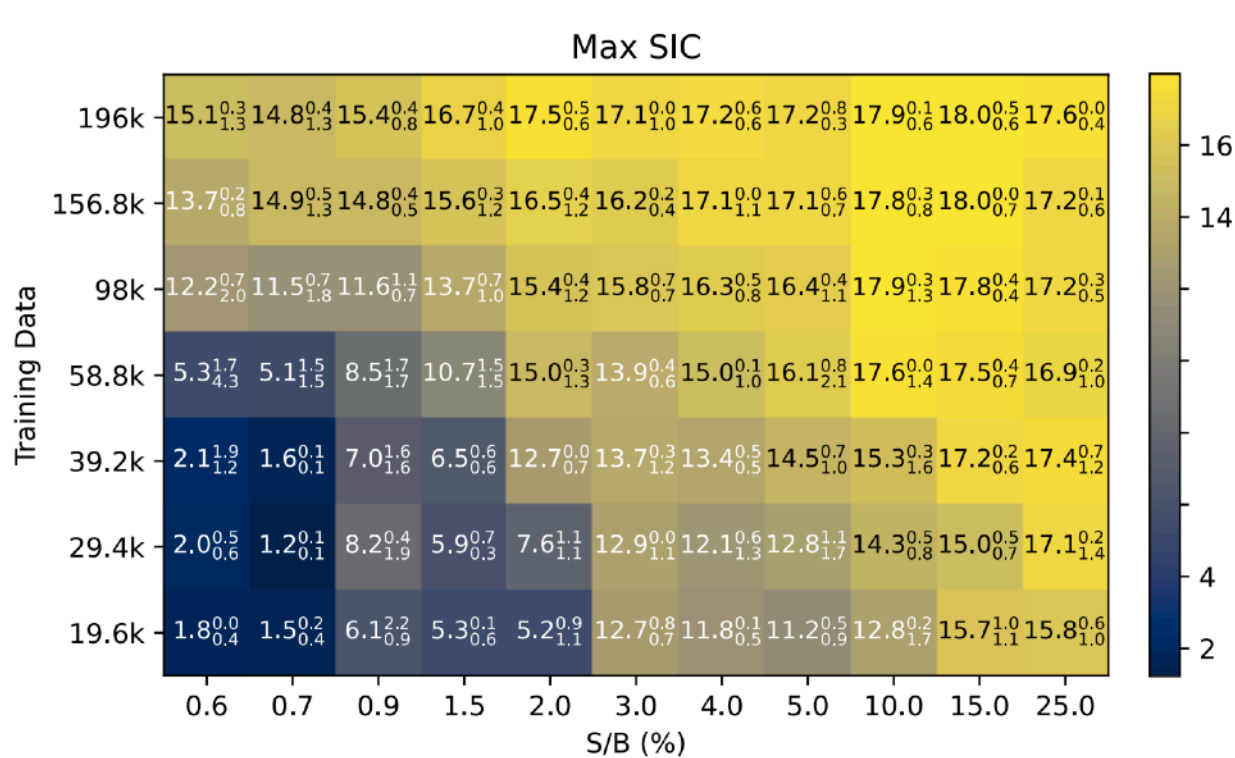
- As CATHODE approximates a likelihood ratio, it should be robust compared to methods that only use  $p_{\text{Background}}$  (e.g. autoencoders)



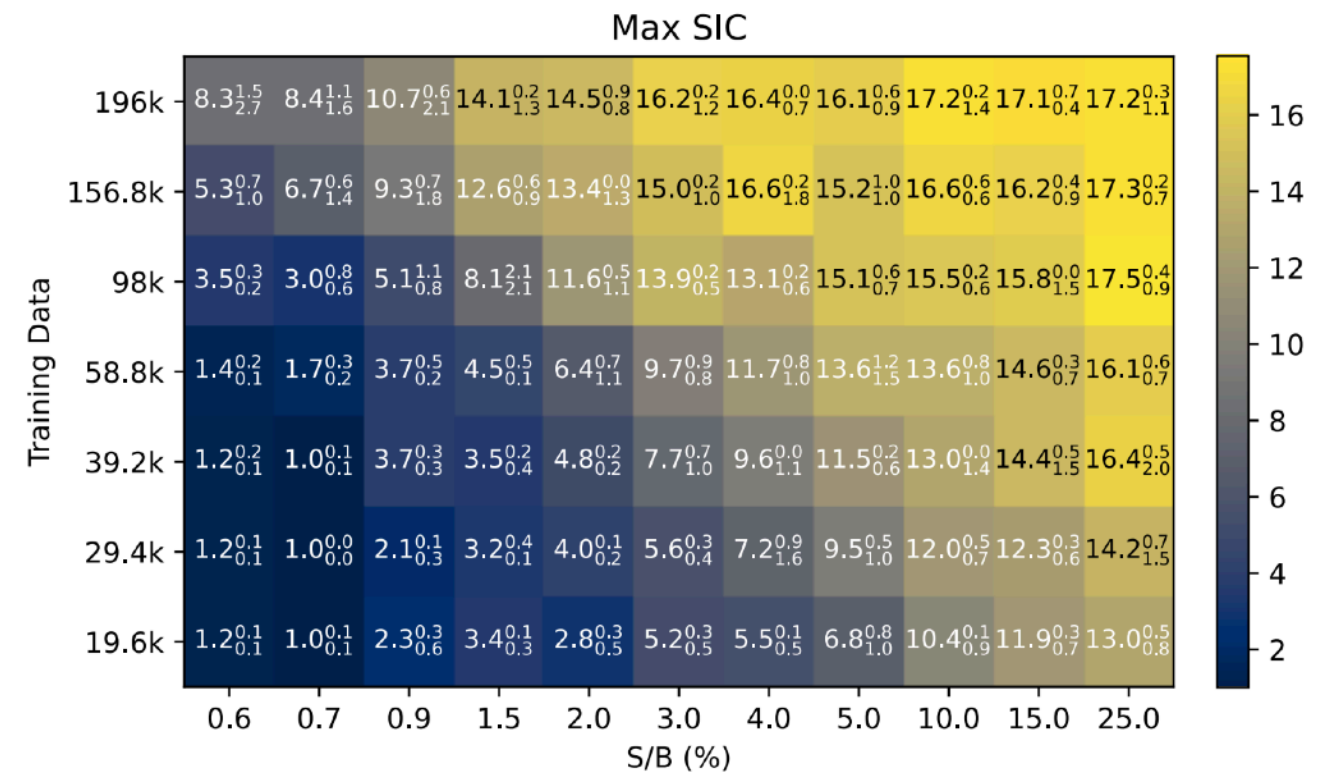


# Comments on anomaly detection

- As CATHODE approximates a likelihood ratio, it should be robust compared to methods that only use  $p_{\text{Background}}$  (e.g. autoencoders)
- However, still can be sensitive to choice of input features
  - Here shown: idealised anomaly detector (perfect DE)



No noisy features



Two noisy features

# Closing



# Advertisement

Consider coming to Hamburg!

ML4Jets from November 6-10

Registration and Abstract submission are **open** under

<https://indico.cern.ch/e/ml4jets>

## ML4Jets2023

6–10 Nov 2023  
DESY  
Europe/Zurich timezone

- Overview
- Call for Abstracts
- Timetable
- Registration
- Participant List
- Code of Conduct

**Contact**  
✉ [ml4jets2023-info@desy.de](mailto:ml4jets2023-info@desy.de)

Machine learning has become a hot topic in particle physics over the past several years. In particular, there has been a lot of progress in the areas of particle and event identification, reconstruction, generative models, anomaly detection and more. In this conference, we will discuss current progress in these areas, focusing on new breakthrough ideas and existing challenges. The ML4Jets workshop will be open to the full community and will include LHC experiments as well as theorists and phenomenologists interested in this topic. Contributions from method scientists as well as adjacent scientific fields facing similar challenges are welcome as well.

This year's conference is organised jointly by [DESY](#) and [Universität Hamburg](#) and hosted at the DESY campus. It follows conferences in [2017](#), [2018](#), [2020](#), [2021](#), and [2022](#).

**Registration and abstract submission are open now and close on October 6th (registration for in-person participation) and September 10 (abstracts).**

The workshop will be organised in a hybrid format (with a Zoom connection option). **We expect speakers to attend in-person.**

Registration for both in-person and Zoom-participation will be free of charge and (at the minimum) include coffee-breaks for in-person participants. We are looking into an opt-in dinner and announce details and potential extra costs closer to the event.

**Local Organizing Committee:**  
Freya Blekman (DESY & Universität Hamburg)  
Andrea Bremer (Universität Hamburg)  
Frank Gaede (DESY)  
Gregor Kasieczka (Universität Hamburg, chair)  
Andreas Hinzmann (DESY)  
Matthias Schröder (Universität Hamburg)

**International Advisory Committee:**  
Florencia Canelli (University of Zurich)  
Kyle Cranmer (NYU)  
Vava Gligorov (LPNHE)  
Gian Michele Innocenti (CERN)  
Ben Nachman (LBNL)  
Mihoko Nojiri (KEK)  
Maurizio Pierini (CERN)  
Tilman Plehn (Heidelberg)  
David Shih (Rutgers)  
Jesse Thaler (MIT)  
Sofia Vallescorsa (CERN)



# Closing

- Deep Learning for particle physics is rapidly developing solutions to a wide range of problems
  - Classification
  - Anomaly detection
  - Robustness and uncertainties
  - Efficient generation
  - Fast processing
- Physics encounters challenges of complex data and large volumes with potential relevance to other domains
- **Contact:**
  - Email: [gregor.kasieczka@uni-hamburg.de](mailto:gregor.kasieczka@uni-hamburg.de)
  - Twitter: [@GregorKasieczka](https://twitter.com/GregorKasieczka)
  - Webpage: <https://www.physik.uni-hamburg.de/en/iexp/gruppe-kasieczka.html>



Thank you