Likelihood free generative modeling for high energy physics

Benjamin Nachman

Lawrence Berkeley National Laboratory

BERKELEY LAB

BERKELEY EXPERIMENTAL PARTICLE PHYSICS

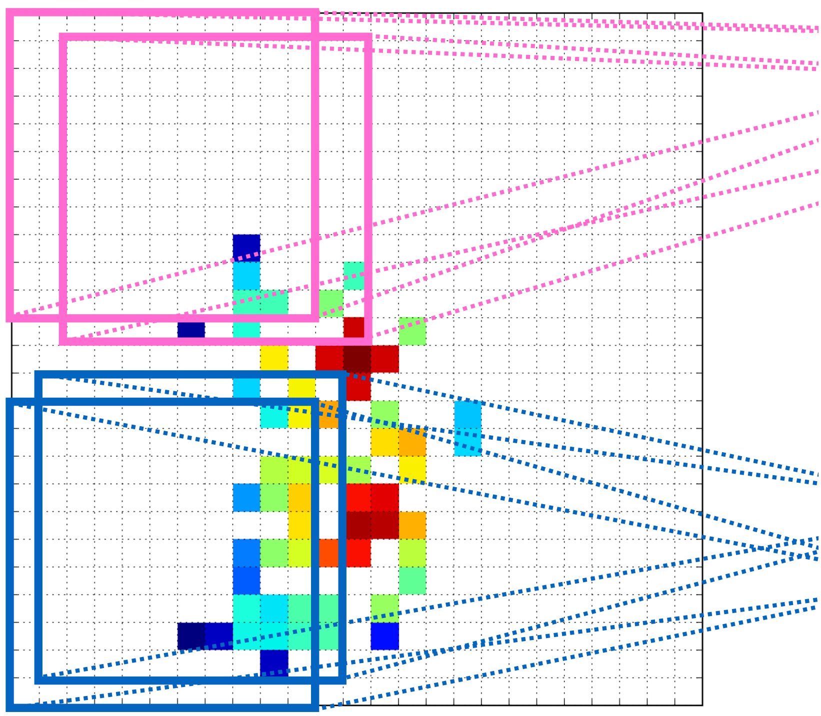

IPAM, UCLA

From Passive to Active: Generative and Reinforcement Learning with Physics

September 27, 2019

Likelihood free generative modeling for high energy physics

In contrast to flows and related methods, I won't ever write down a way to evaluate the PDF.

B

Lawrence Berkeley National Laboratory

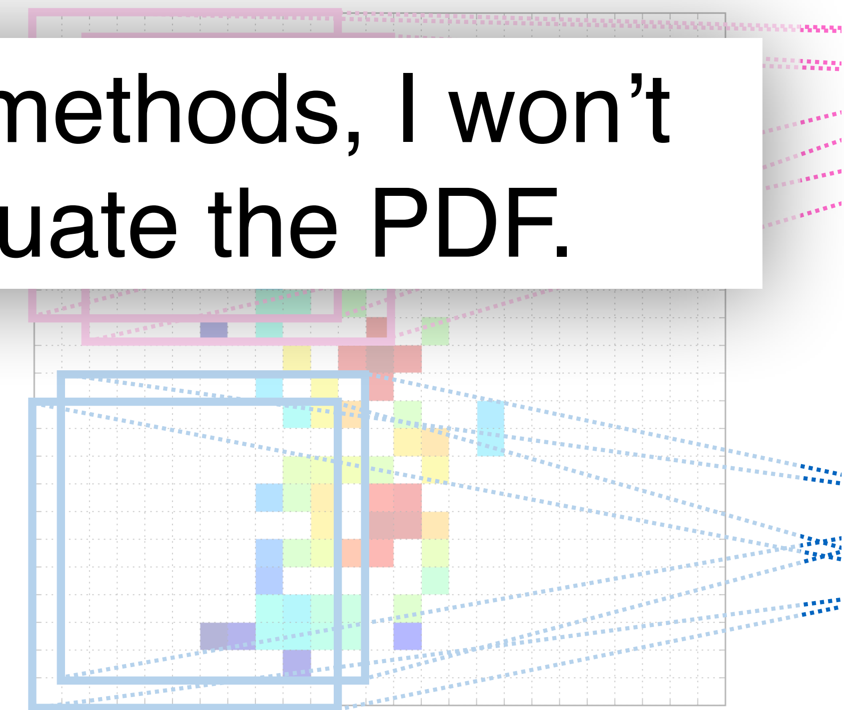

BERKELEY LAB

BERKELEY
EXPERIMENTAL
PARTICLE
PHYSICS

IPAM, UCLA

From Passive to Active: Generative and Reinforcement Learning with Physics

September 27, 2019

First, the science of high energy physics

Goal: **we want to study the structure of the smallest building blocks of matter**. For this, we need the most powerful microscope ever built!

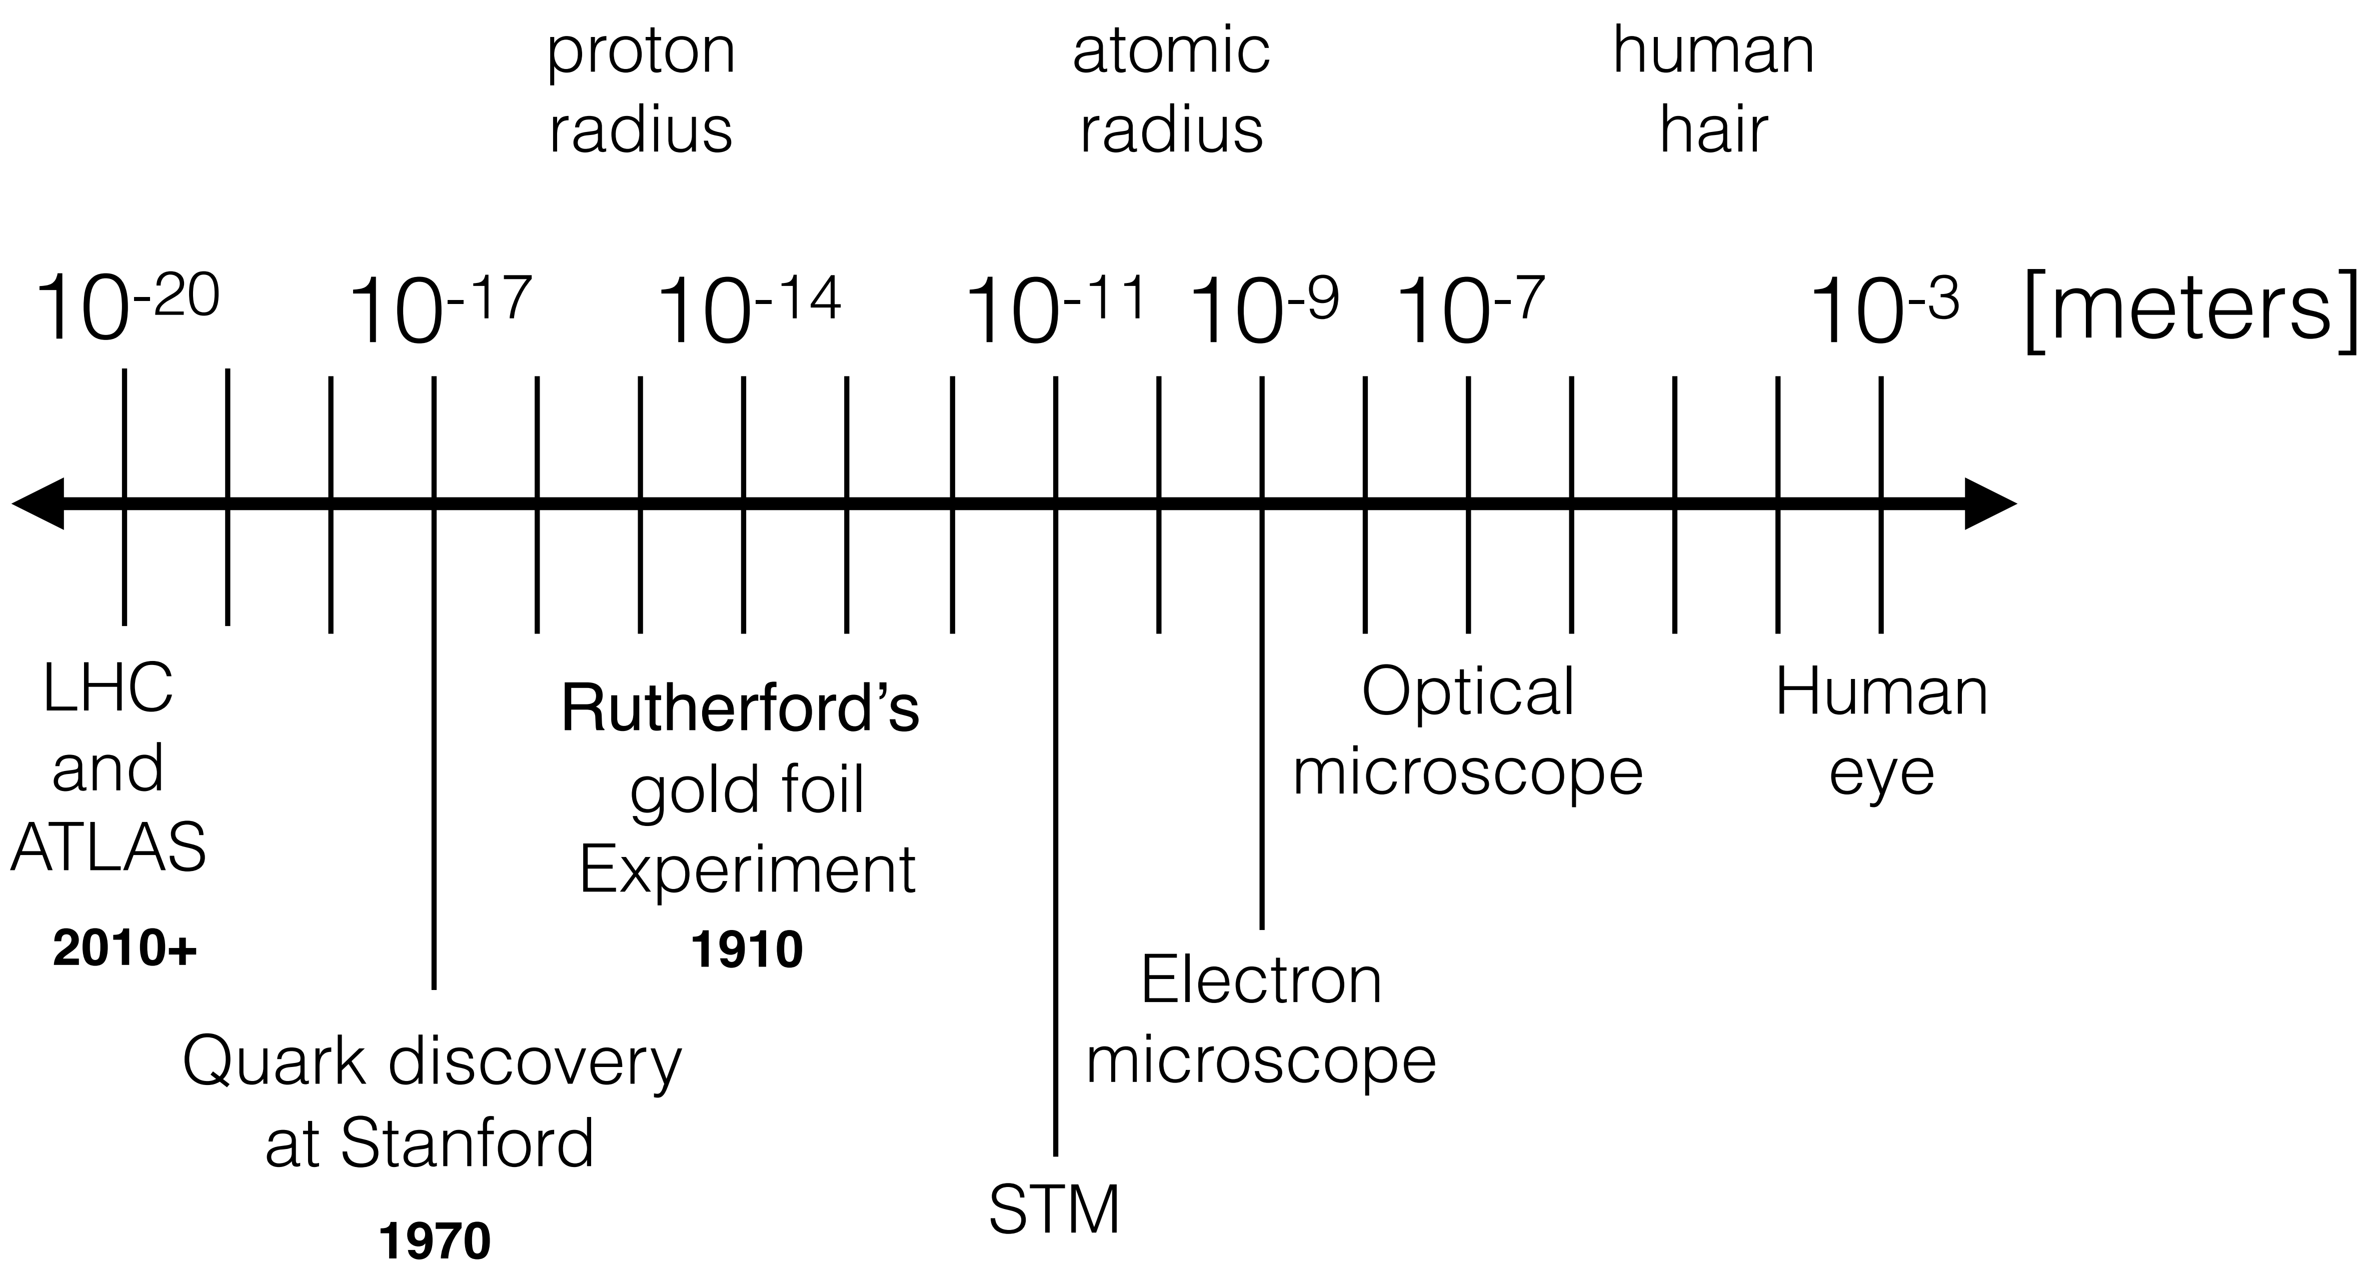

…of high energy physics

4

…**to study the structure of the …g blocks of matter**. For this, we …powerful microscope ever built!

atomic radius

human hair

…4 10^{-11} 10^{-9} 10^{-7} 10^{-3} [meters]

Optical microscope

Human eye

…rd's …il …

2010+

1910

Quark discovery at Stanford

1970

Electron microscope

STM

Optical Microscope

Resolution: wavelength of light (> 100 nm)

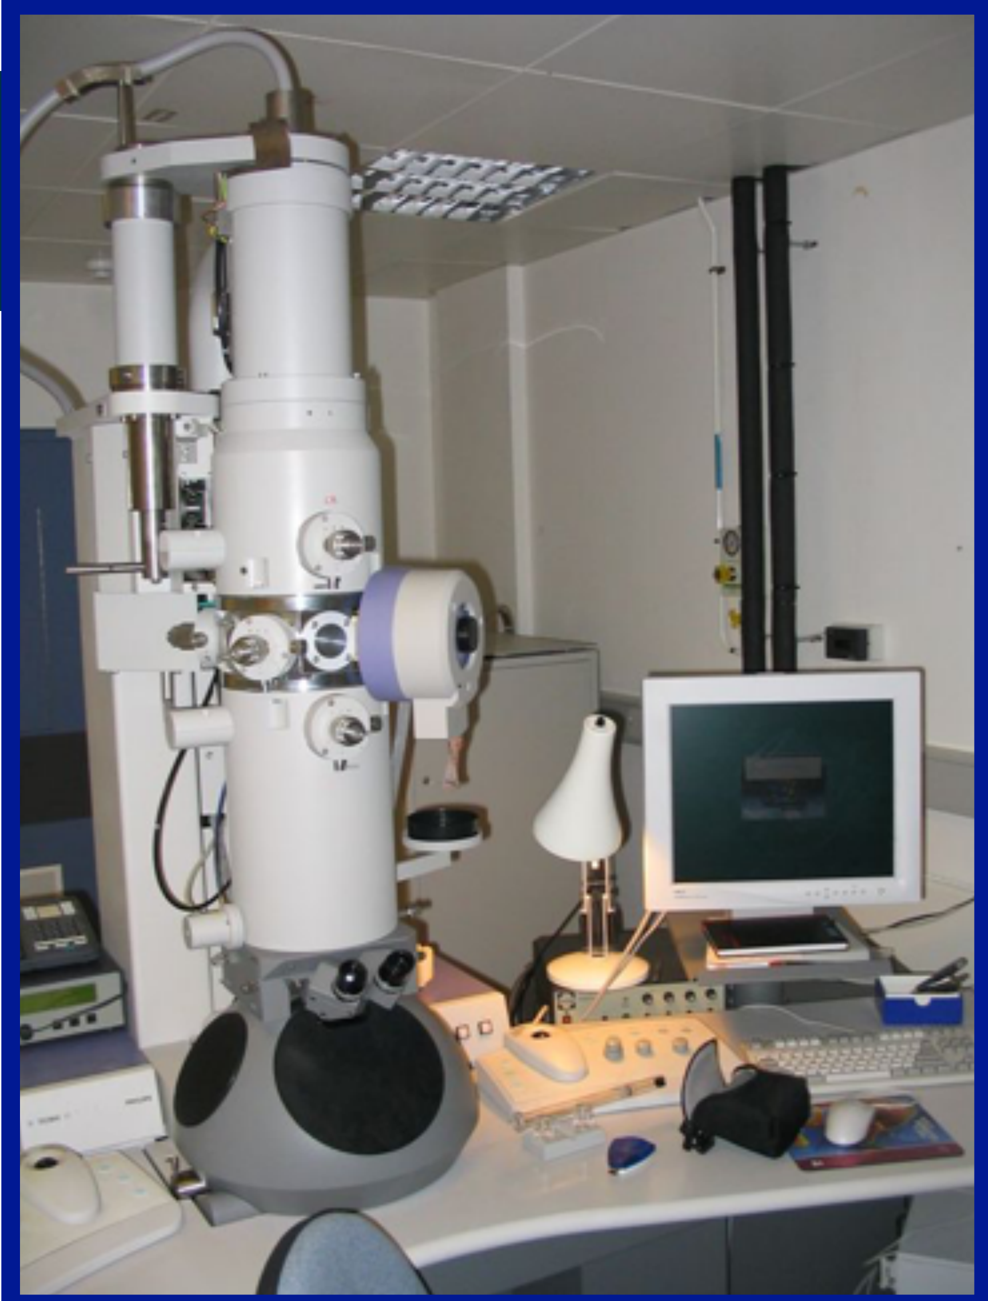

...e of high energy physics

...nt to study the structure of the ...ng blocks of matter. For this, we ...t powerful microscope ever built!

atomic radius

human hair

-14 10^{-11} 10^{-9} 10^{-7} 10^{-3} [meters]

LHC and ATLAS **2010+**

Rutherford's gold foil Experiment **1910**

Quark discovery at Stanford **1970**

STM

Electron microscope

Optical microscope

Remember:

$$E \sim 1/\lambda$$

Electron Microscope

Resolution: ~1 nm

Image: https://upload.wikimedia.org/wikipedia/commons/c/c5/Electron_Microscope.jpg

First, the scienc

Goal: **we wa**
smallest buildi
need the mos

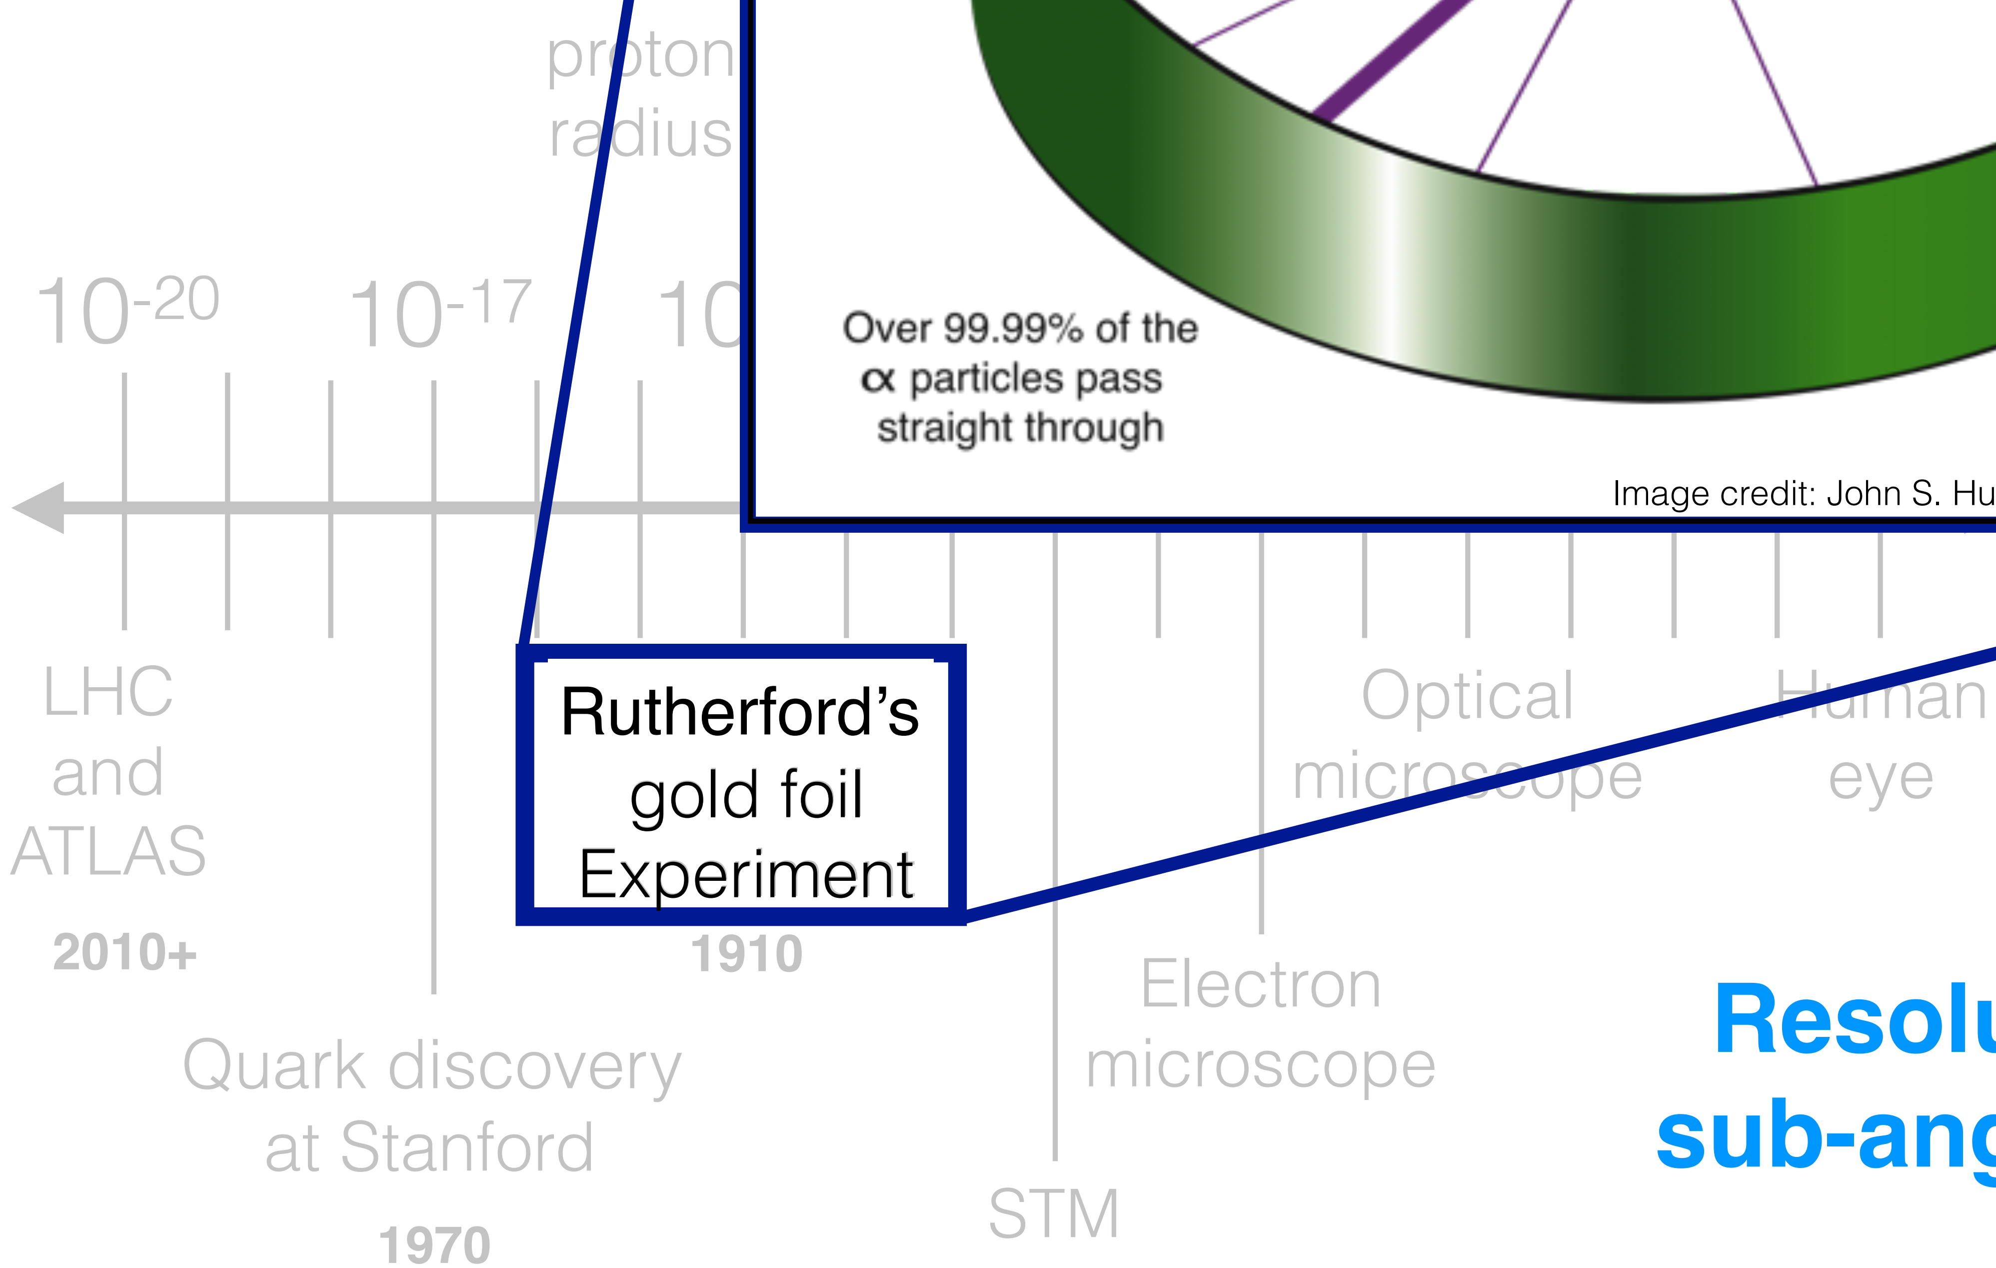

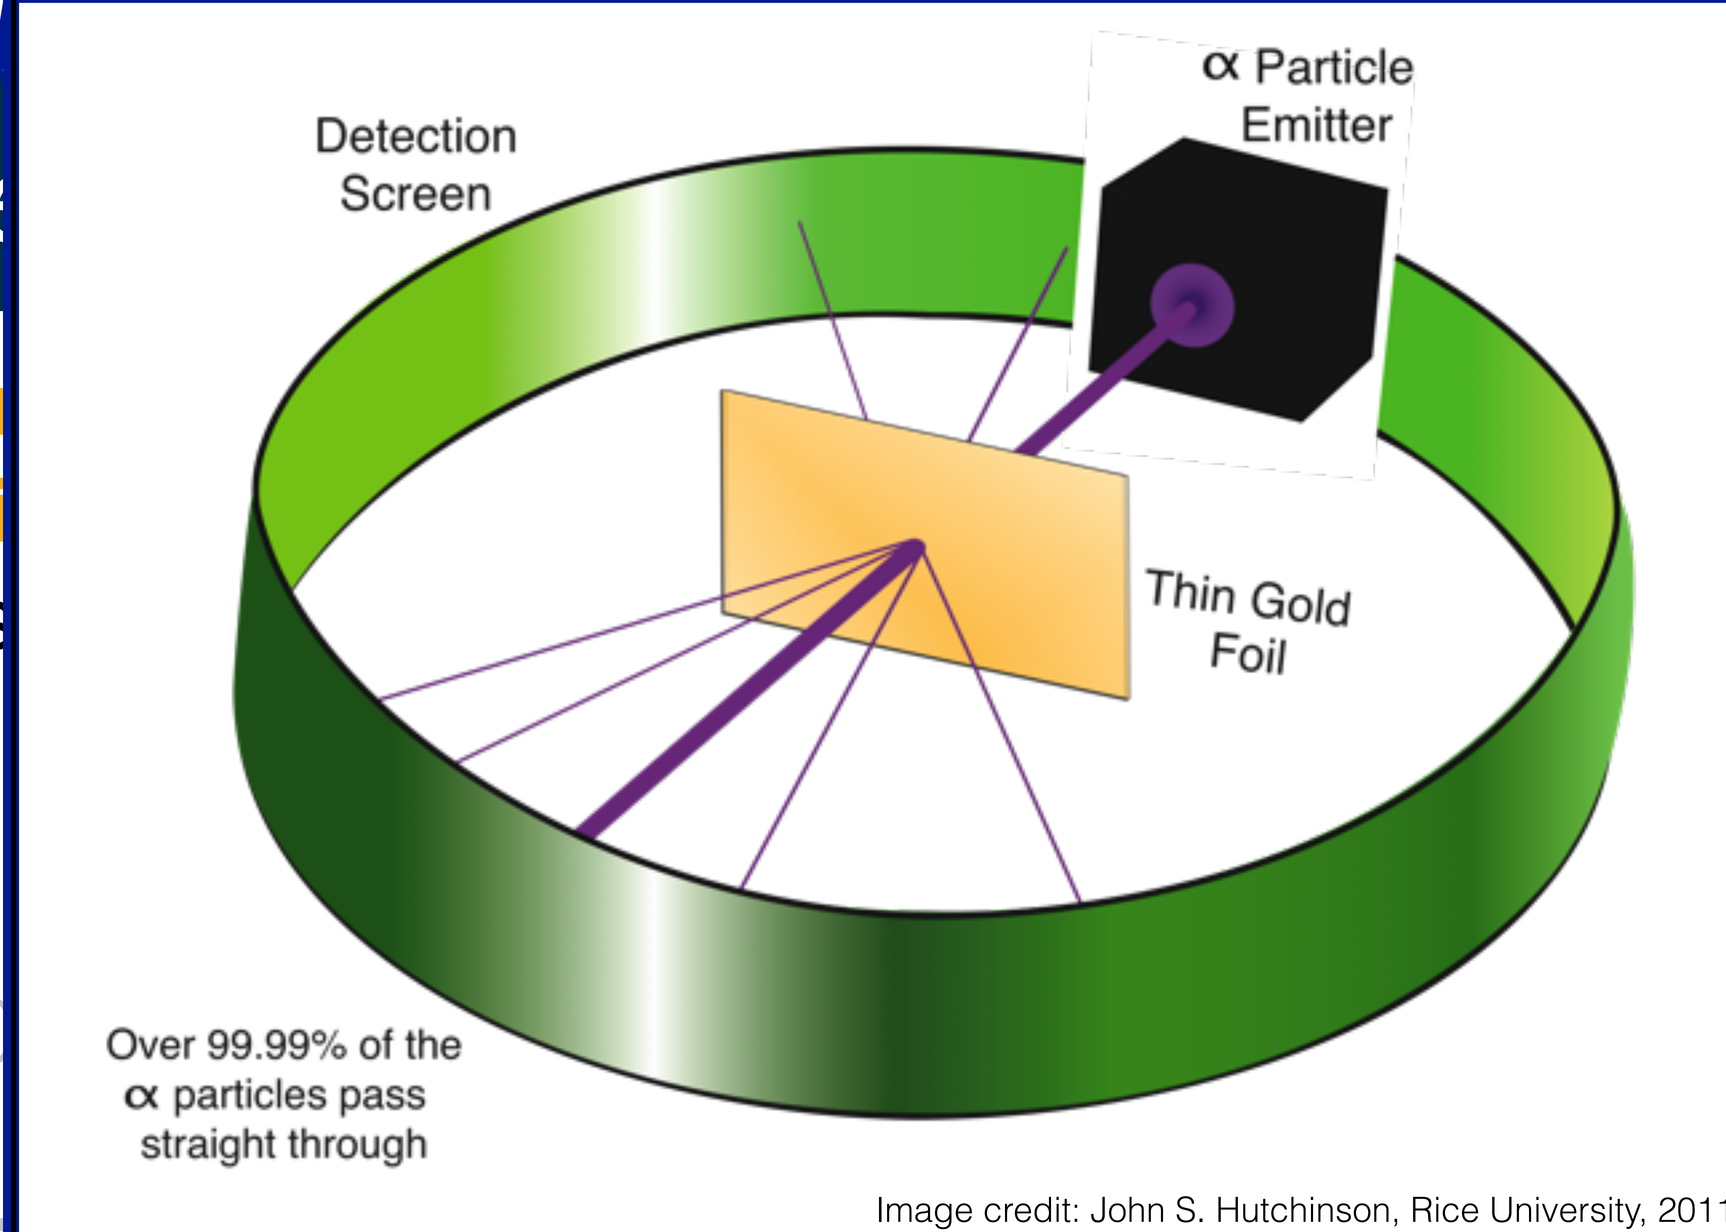

**Resolution:
sub-angstrom**

First, the scie

Goal: **we**
smallest bu
need the r

pr
ra

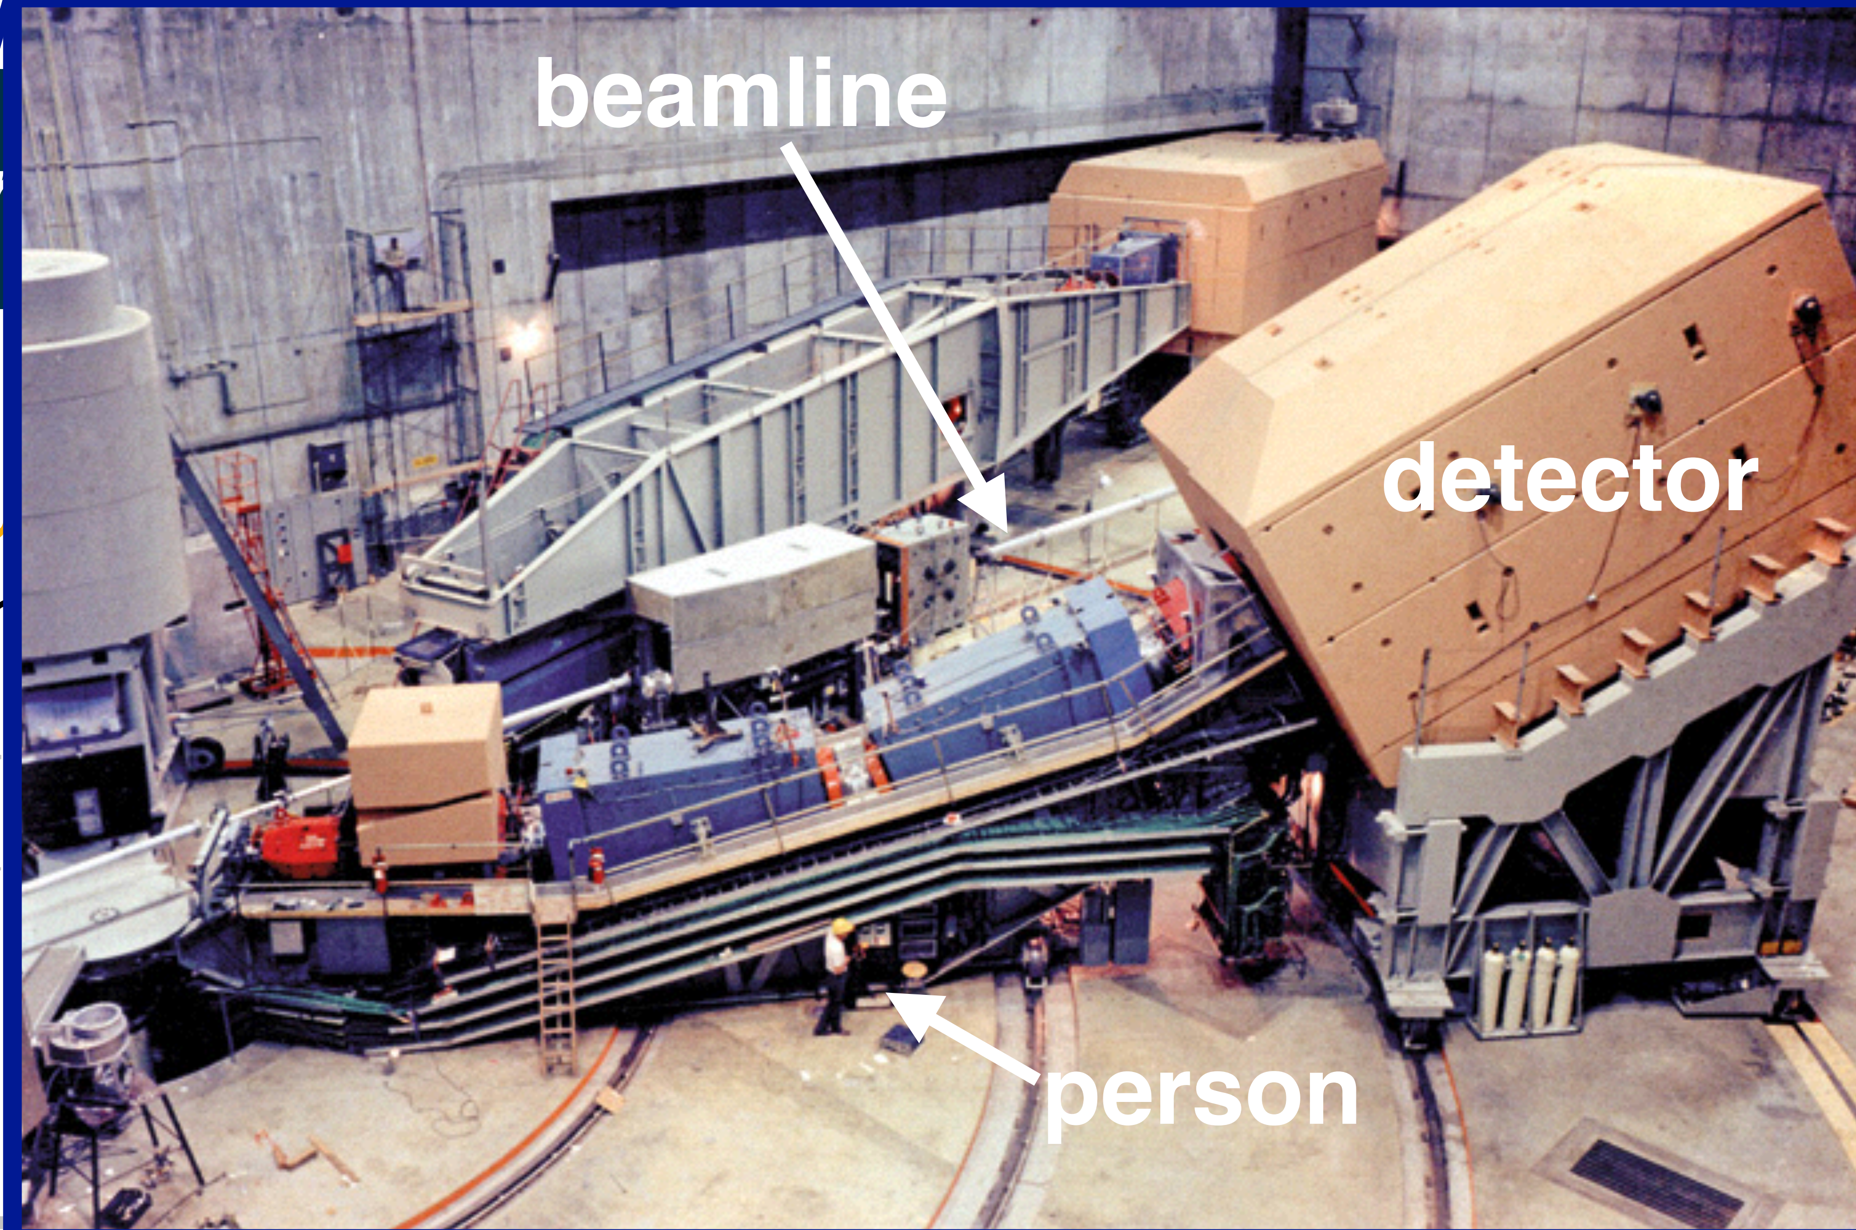

Image credit: SLAC National Accelerator Laboratory.

10^{-20} 10^{-17}

LHC and ATLAS
2010+

Quark discovery at Stanford
1970

Rutherford's gold foil Experiment
1910

STM

Electron microscope

Optical microscope

Human eye

Resolution: sub-femtometer

20 mile beam line and a 5 story detector

access to a new frontier!

10^{-20} 10^{-17} 10

LHC and ATLAS

2010+

Rutherford's gold foil Experiment

1910

Quark discovery at Stanford

1970

STM

Ele micr

Images: CERN

High Energy Physics at the LHC

Center-of-mass energy = 13 TeV

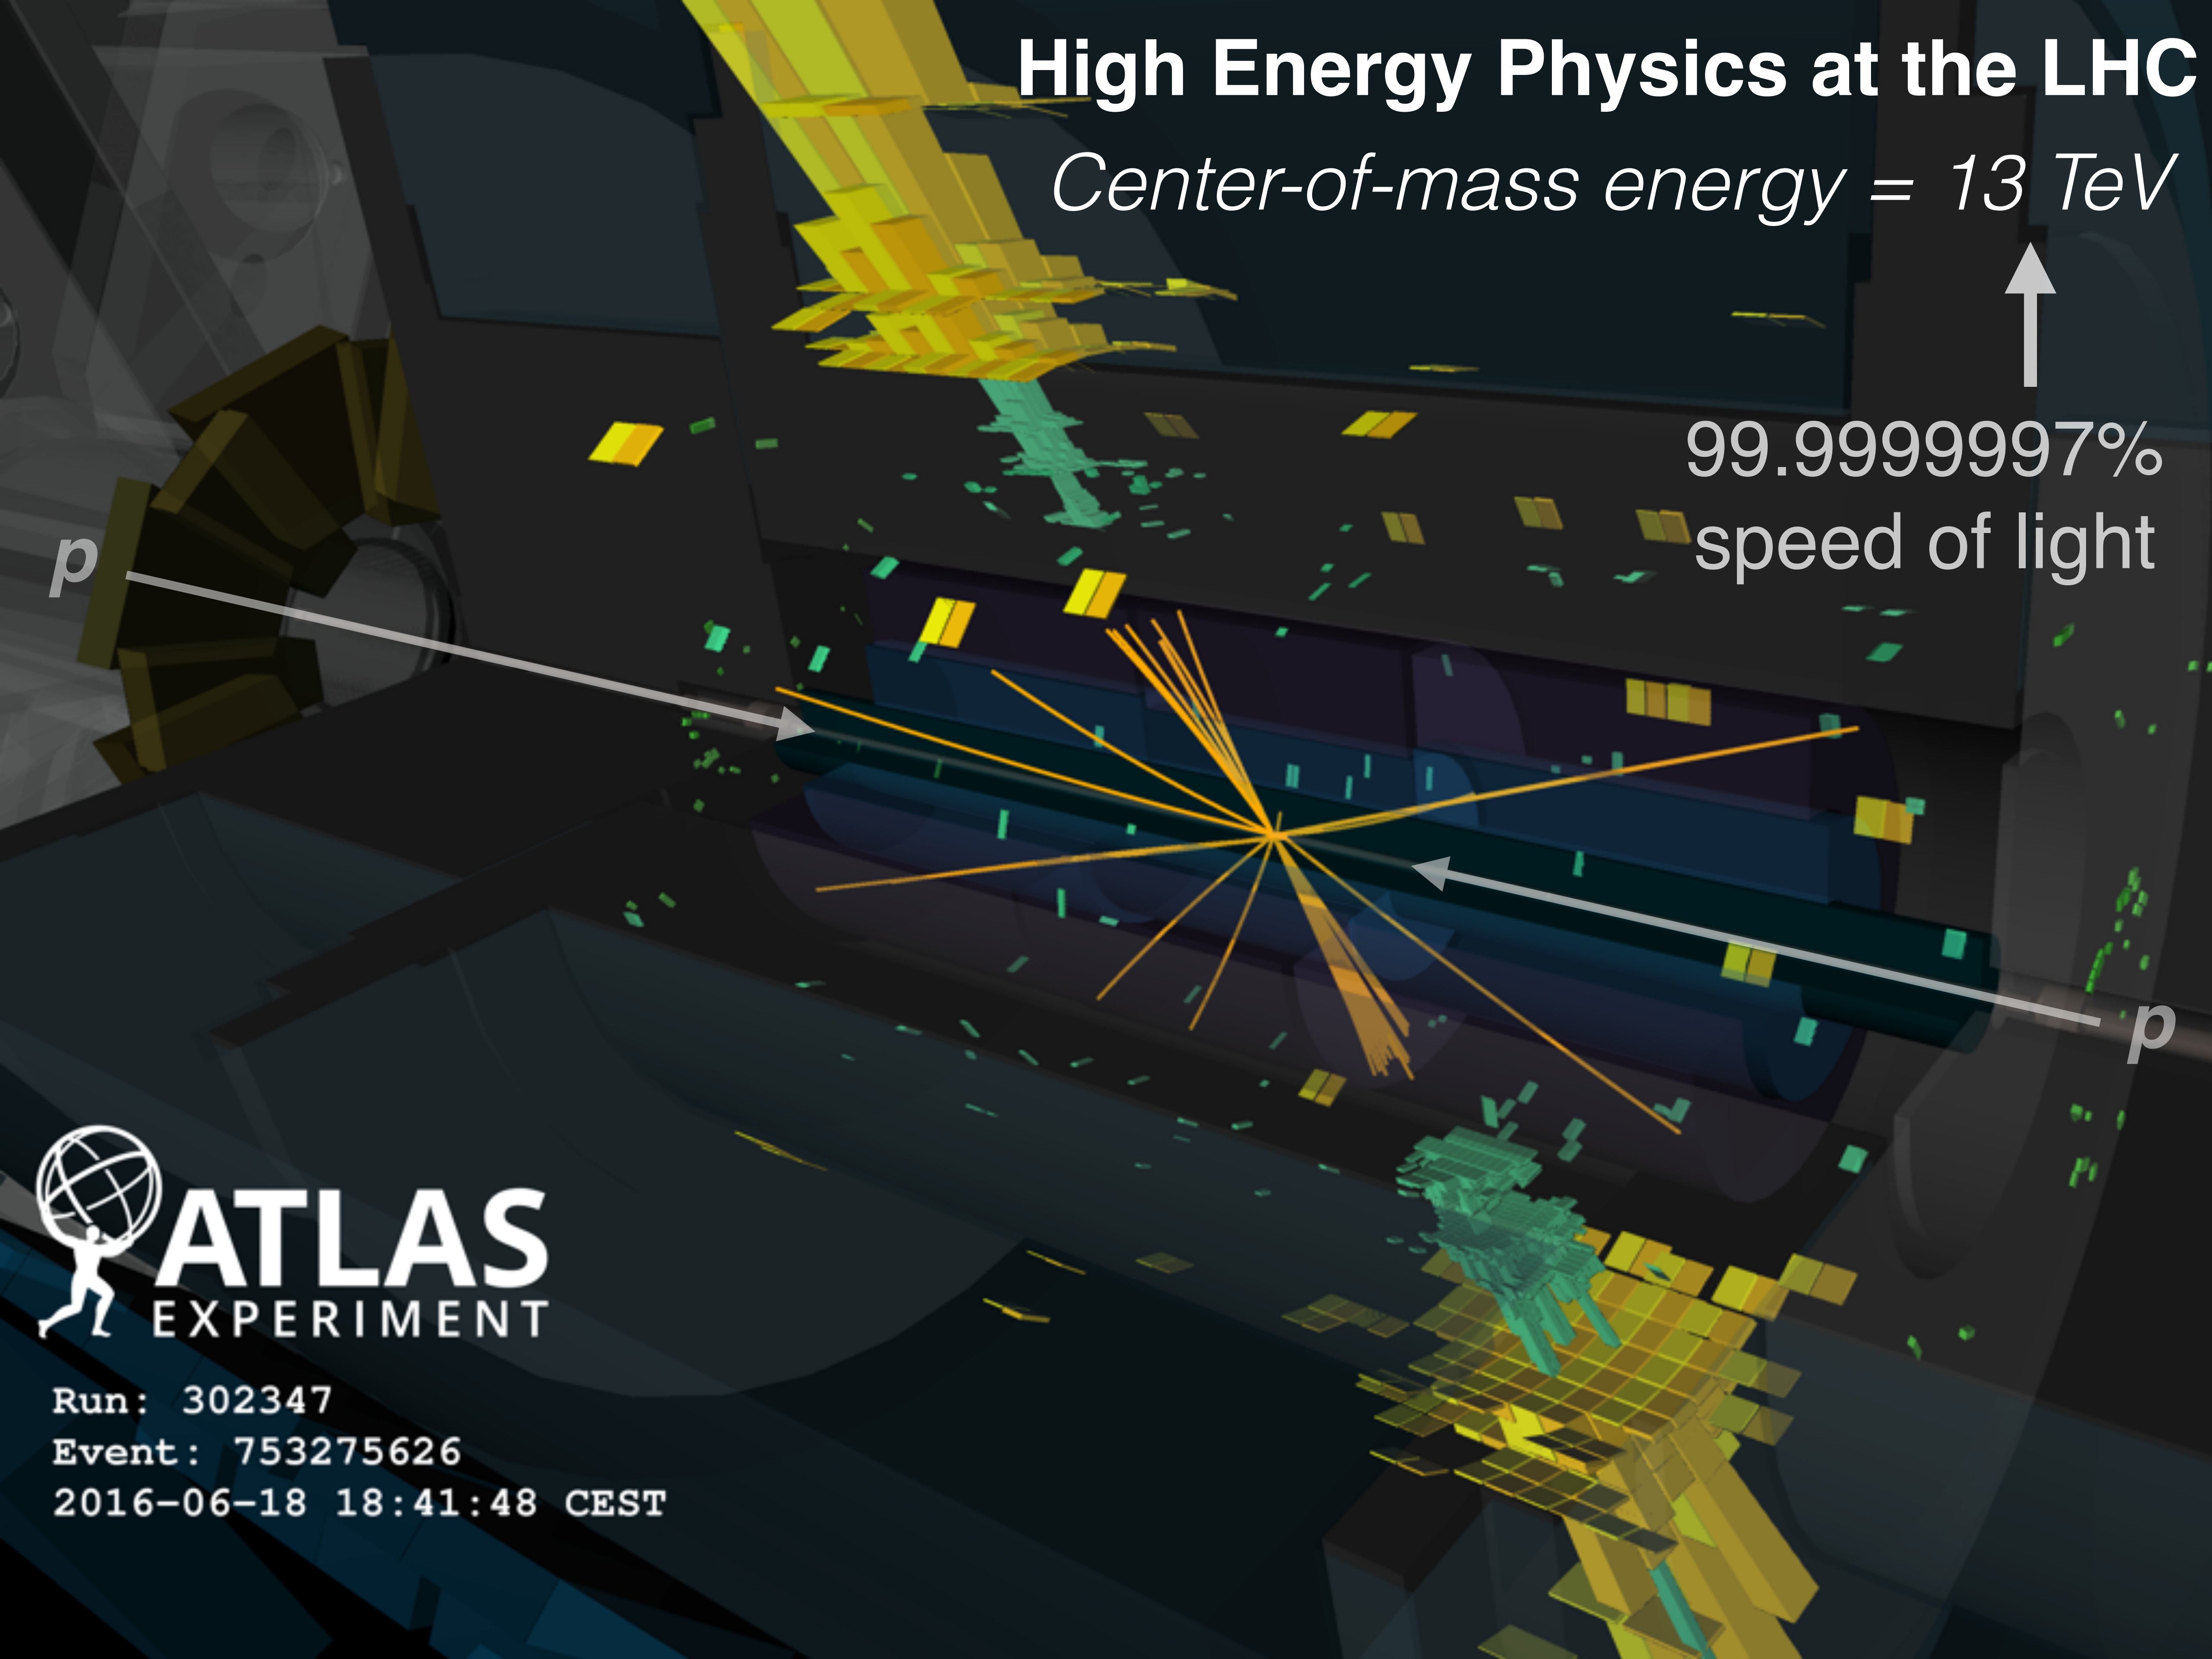

High Energy Physics at the LHC

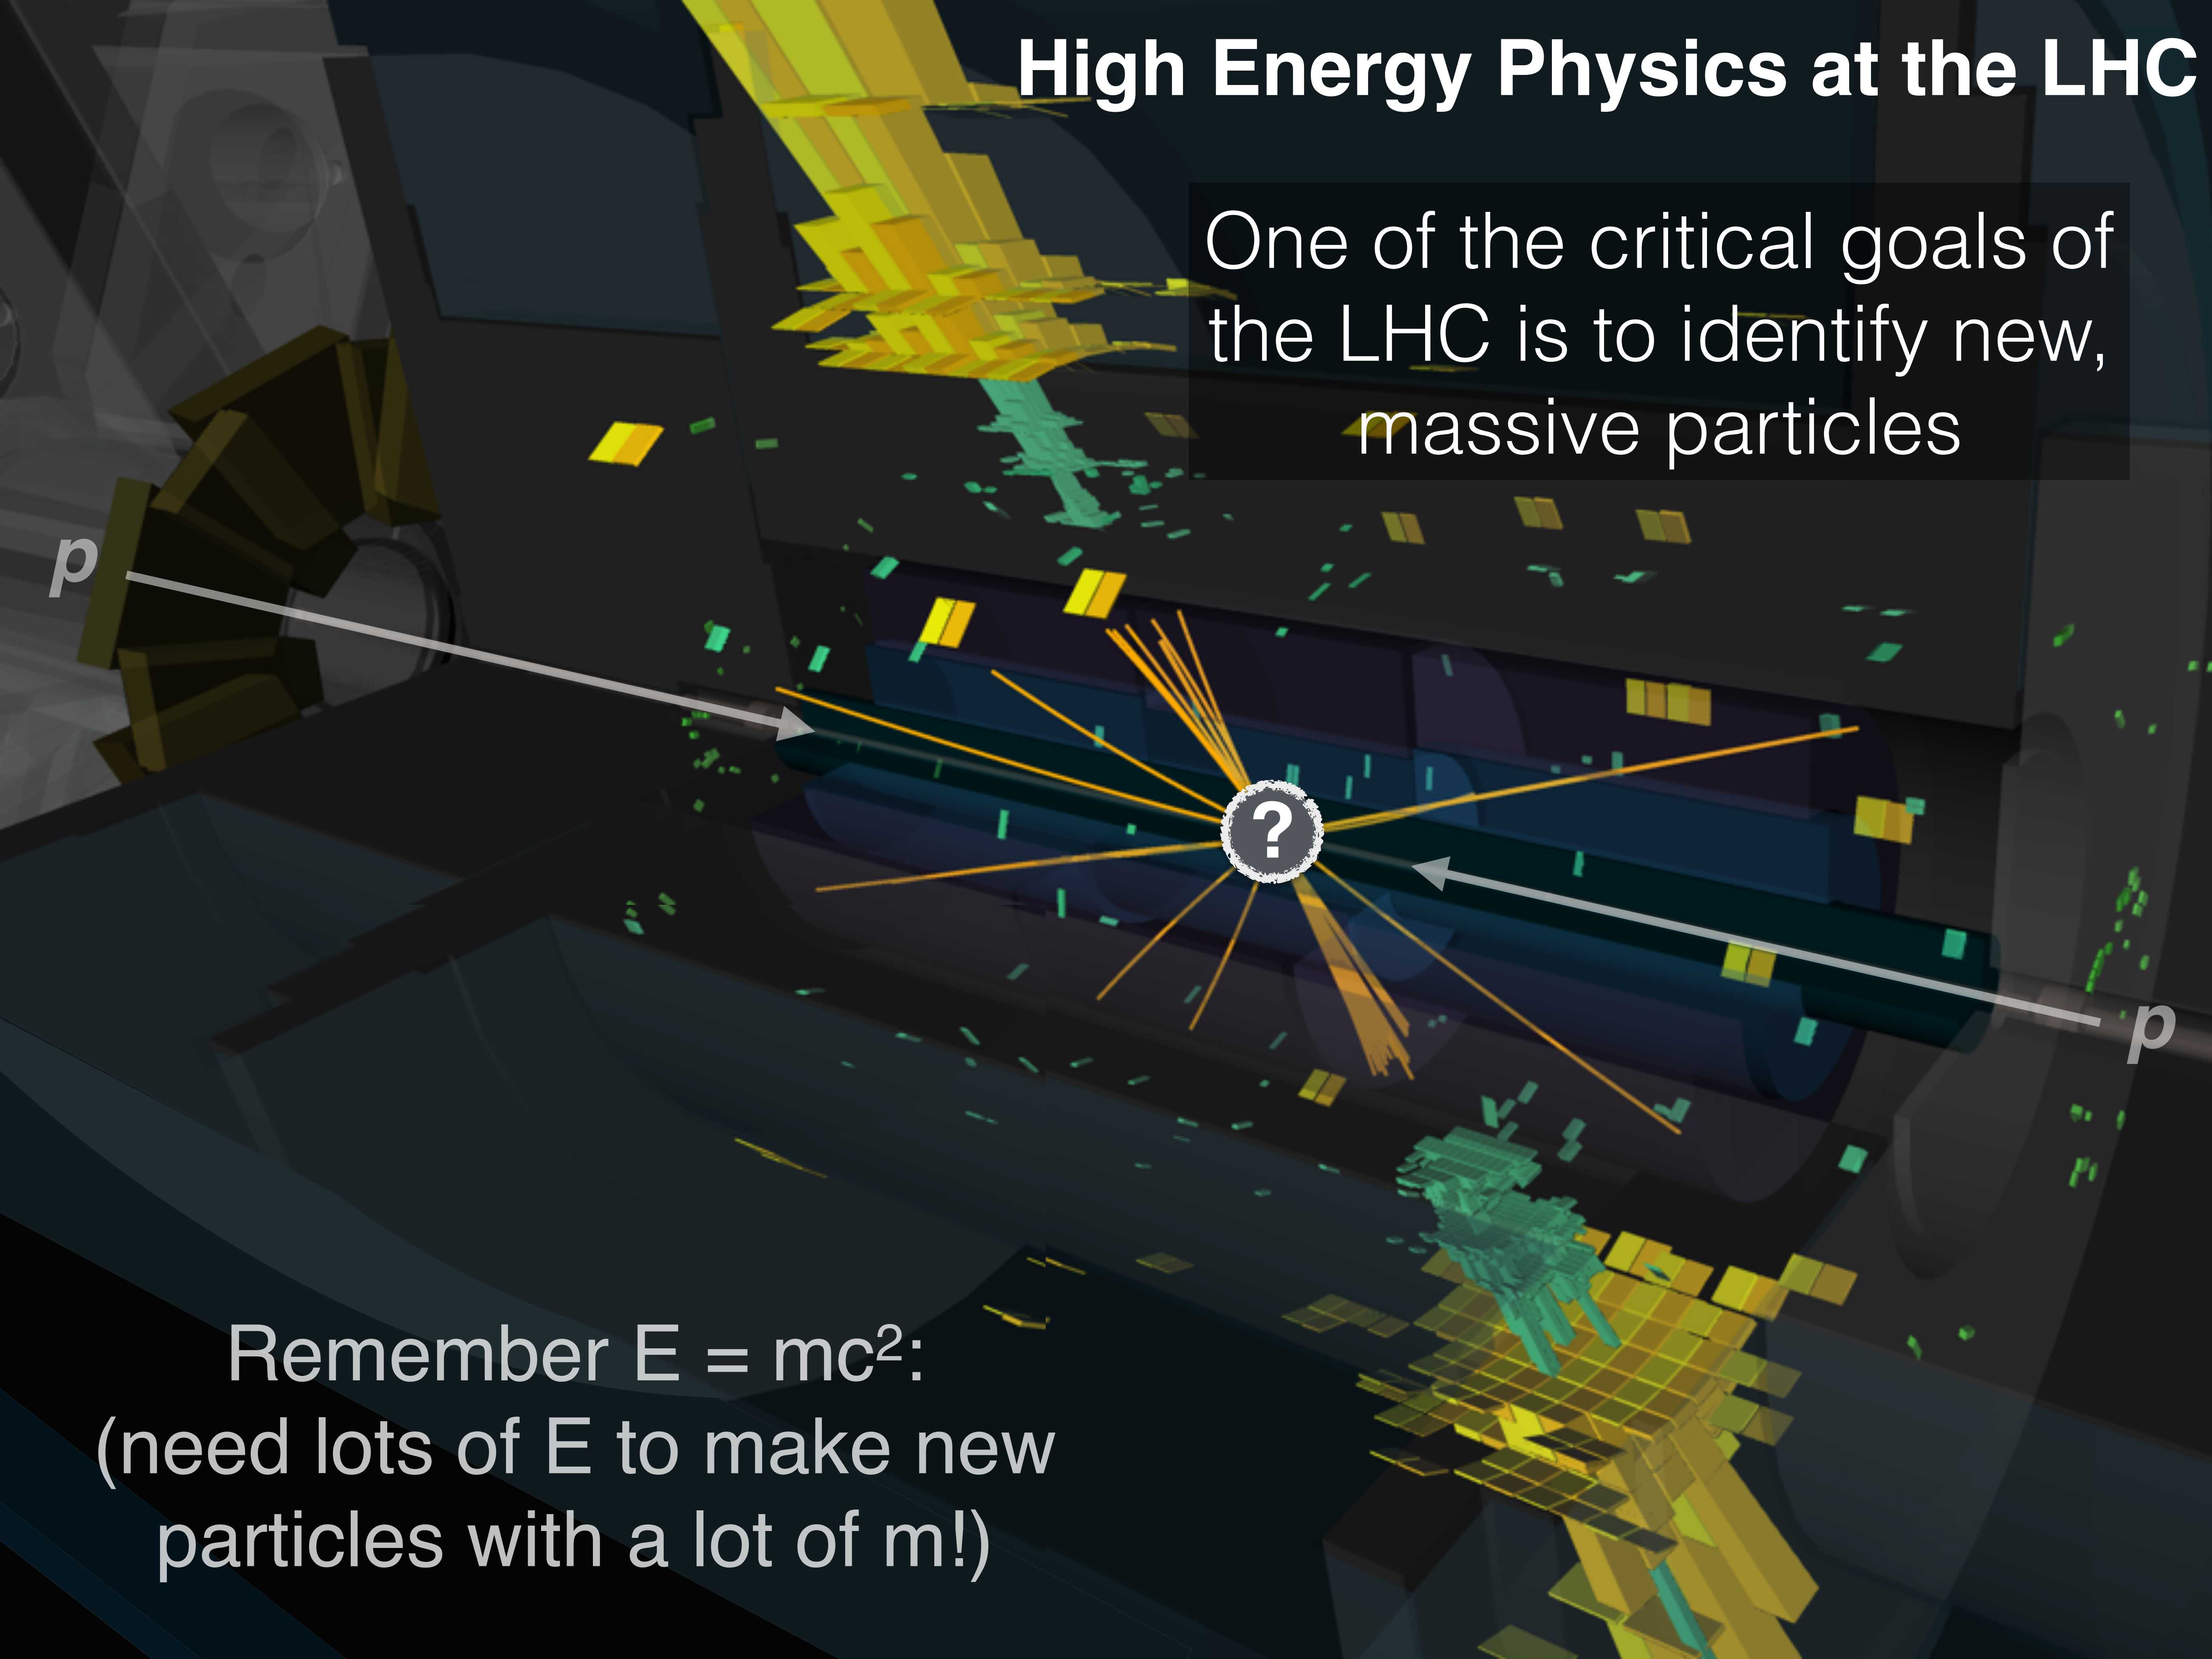

One of the critical goals of the LHC is to identify new, massive particles

Remember $E = mc^2$:
(need lots of E to make new particles with a lot of m!)

Generative models for HEP

…power the inference machine

Theory of everything

↓

Physics simulators

↓

Detector-level features

↓

Pattern recognition ⟷ Pattern recognition

Nature

↓

Experiment

↓

Detector-level features

↓

Pattern recognition

another place where we do a lot of deep learning

…to connect our theories to nature

Part I: *GANs in HEP*

An ab initio generator.

Part II: *High-dimensional reweighting*

Higher fidelity with help from existing simulations.

+ likelihood free inference

Part I: *GANs in HEP*

Accelerating simulations

replace or augment physics simulator

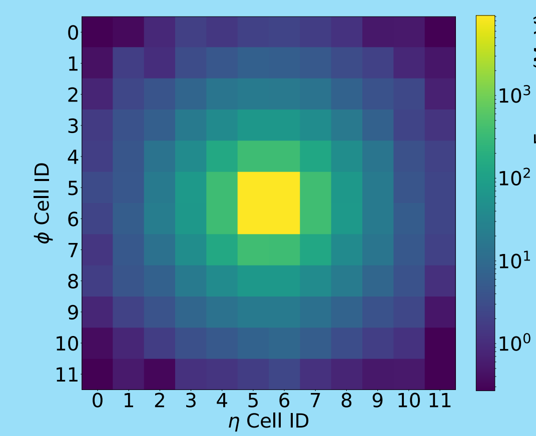

M. Paganini, L. de Oliveira, BPN, PRL 120 (2018) 042003

Saving disk space

replace libraries with on-the-fly generation

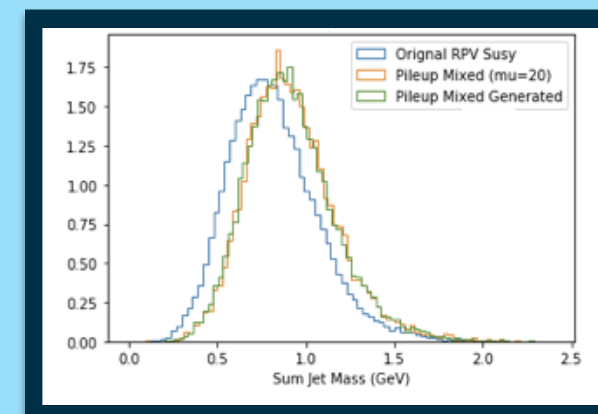

W. Bhimji, W. Blair, S. Farrell, BPN, CHEP 2018

Unbinned, high-dimensional interpolation

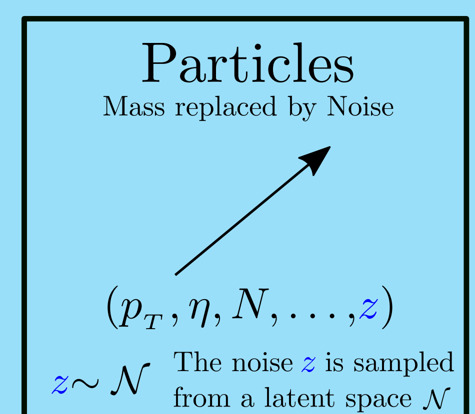

J. Lin, W. Bhimji, BPN, JHEP 05 (2019) 181

Part I: *GANs in HEP*

Focus for today

Accelerating simulations
replace or augment physics simulator

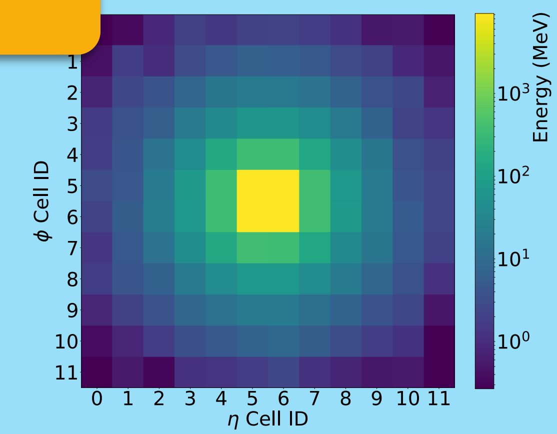

M. Paganini, L. de Oliveira, BPN, PRL 120 (2018) 042003

Saving disk space
replace libraries with on-the-fly generation

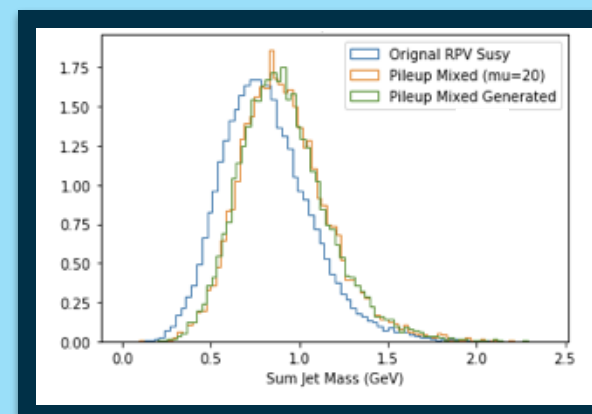

W. Bhimji, W. Blair, S. Farrell, BPN, CHEP 2018

Unbinned, high-dimensional interpolation

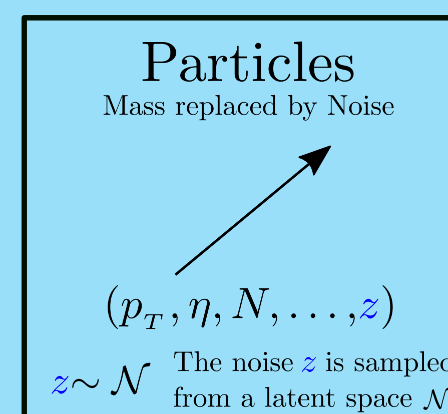

J. Lin, W. Bhimji, BPN, JHEP 05 (2019) 181

Simulations at the LHC

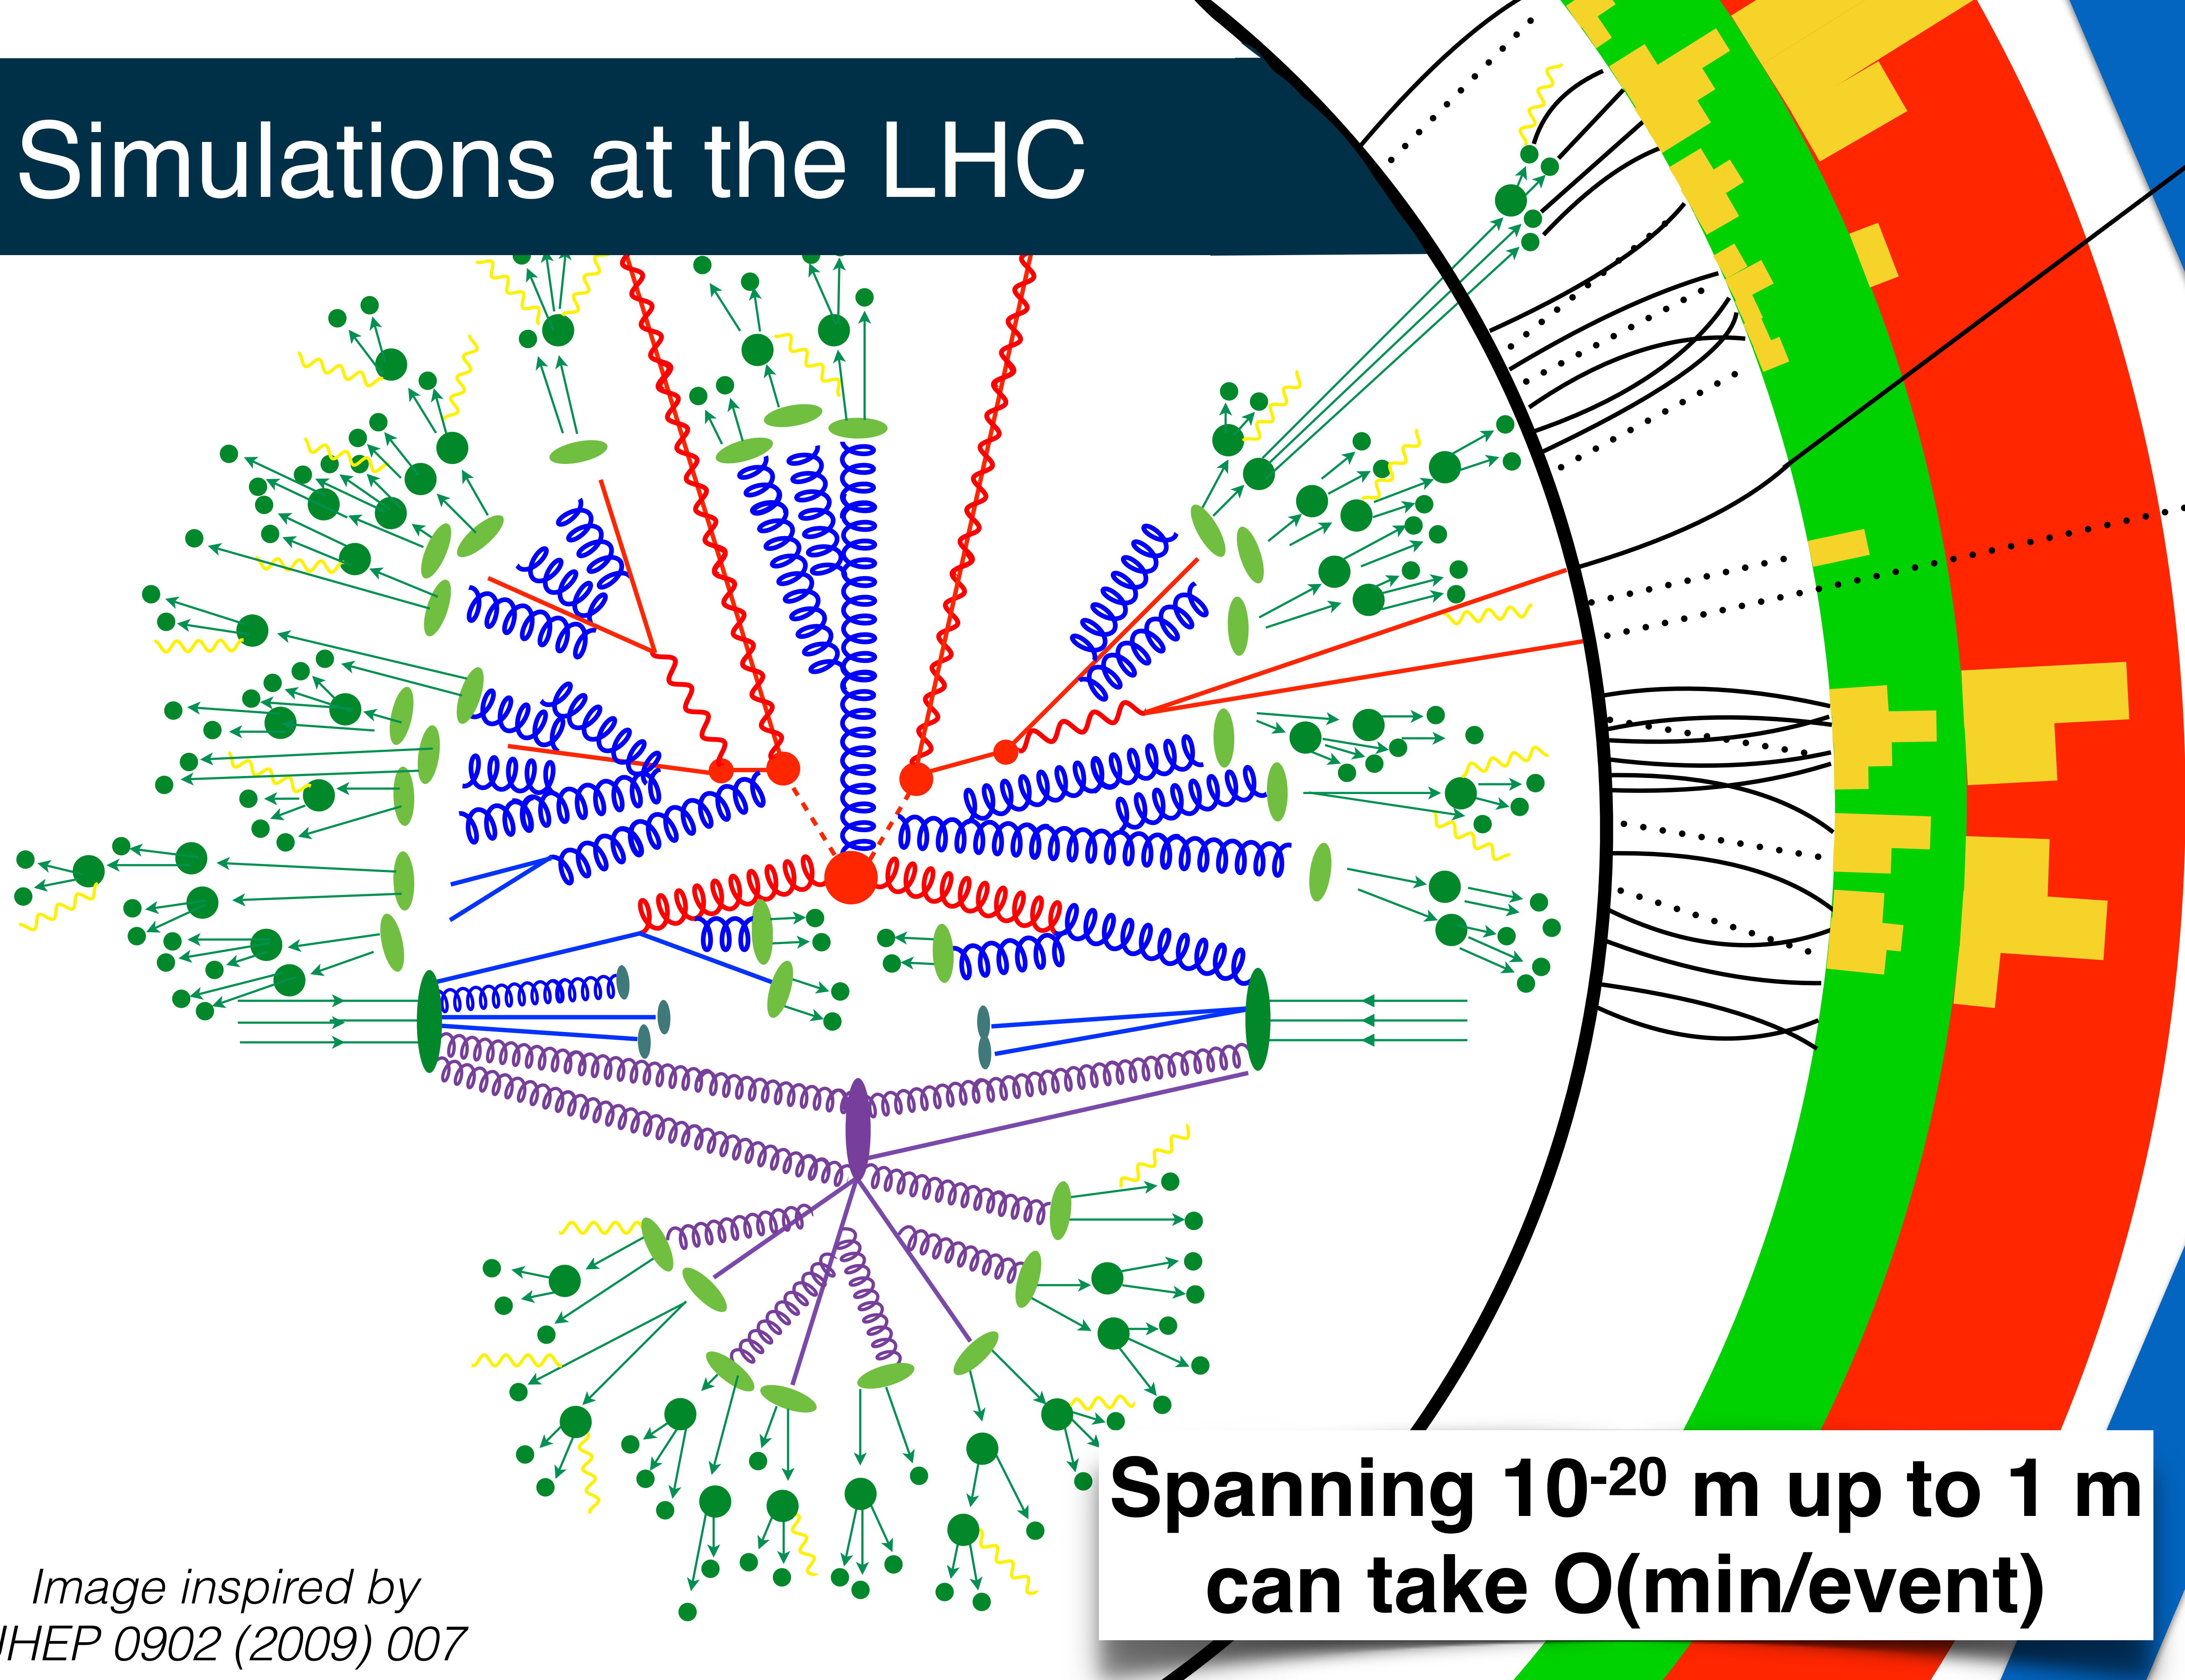

Image inspired by
JHEP 0902 (2009) 007

Spanning 10^{-20} m up to 1 m can take O(min/event)

Simulations at the LHC

State-of-the-art for material interactions is Geant 4.

Includes electromagnetic and hadronic physics with a variety of lists for increasing/decreasing accuracy (at the cost of time)

This accounts for O(1) fraction of all HEP computing resources!

Geant 4

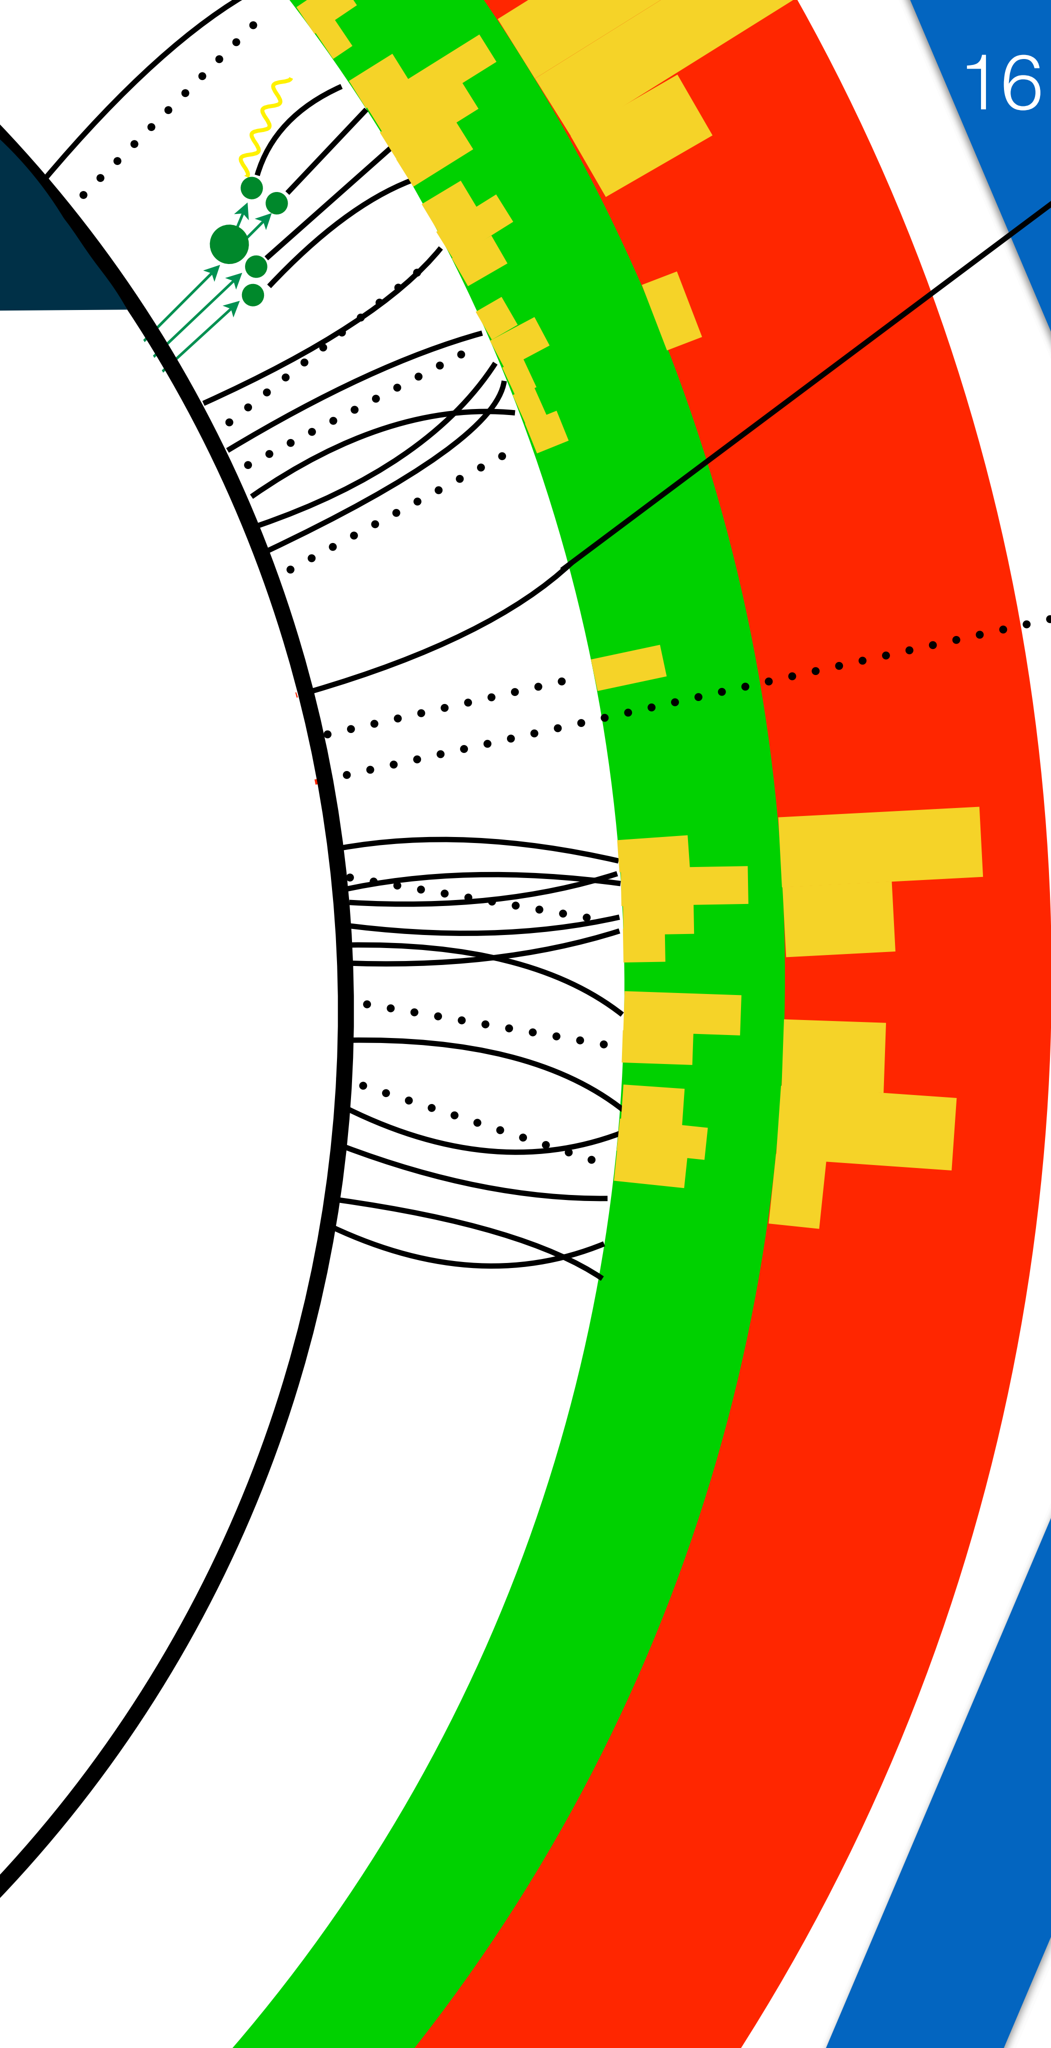

Goal: replace (or augment) simulation steps with a faster, powerful generator based on state-of-the-art machine learning techniques

This work: attack the most important part:
Calorimeter Simulation

Factorization

We are not trying to generate an entire event (O(1000) particles)) all at once - it would be **very had to validate!** Instead, generate a single particle shower (before electronics) and appeal to combinatorics.

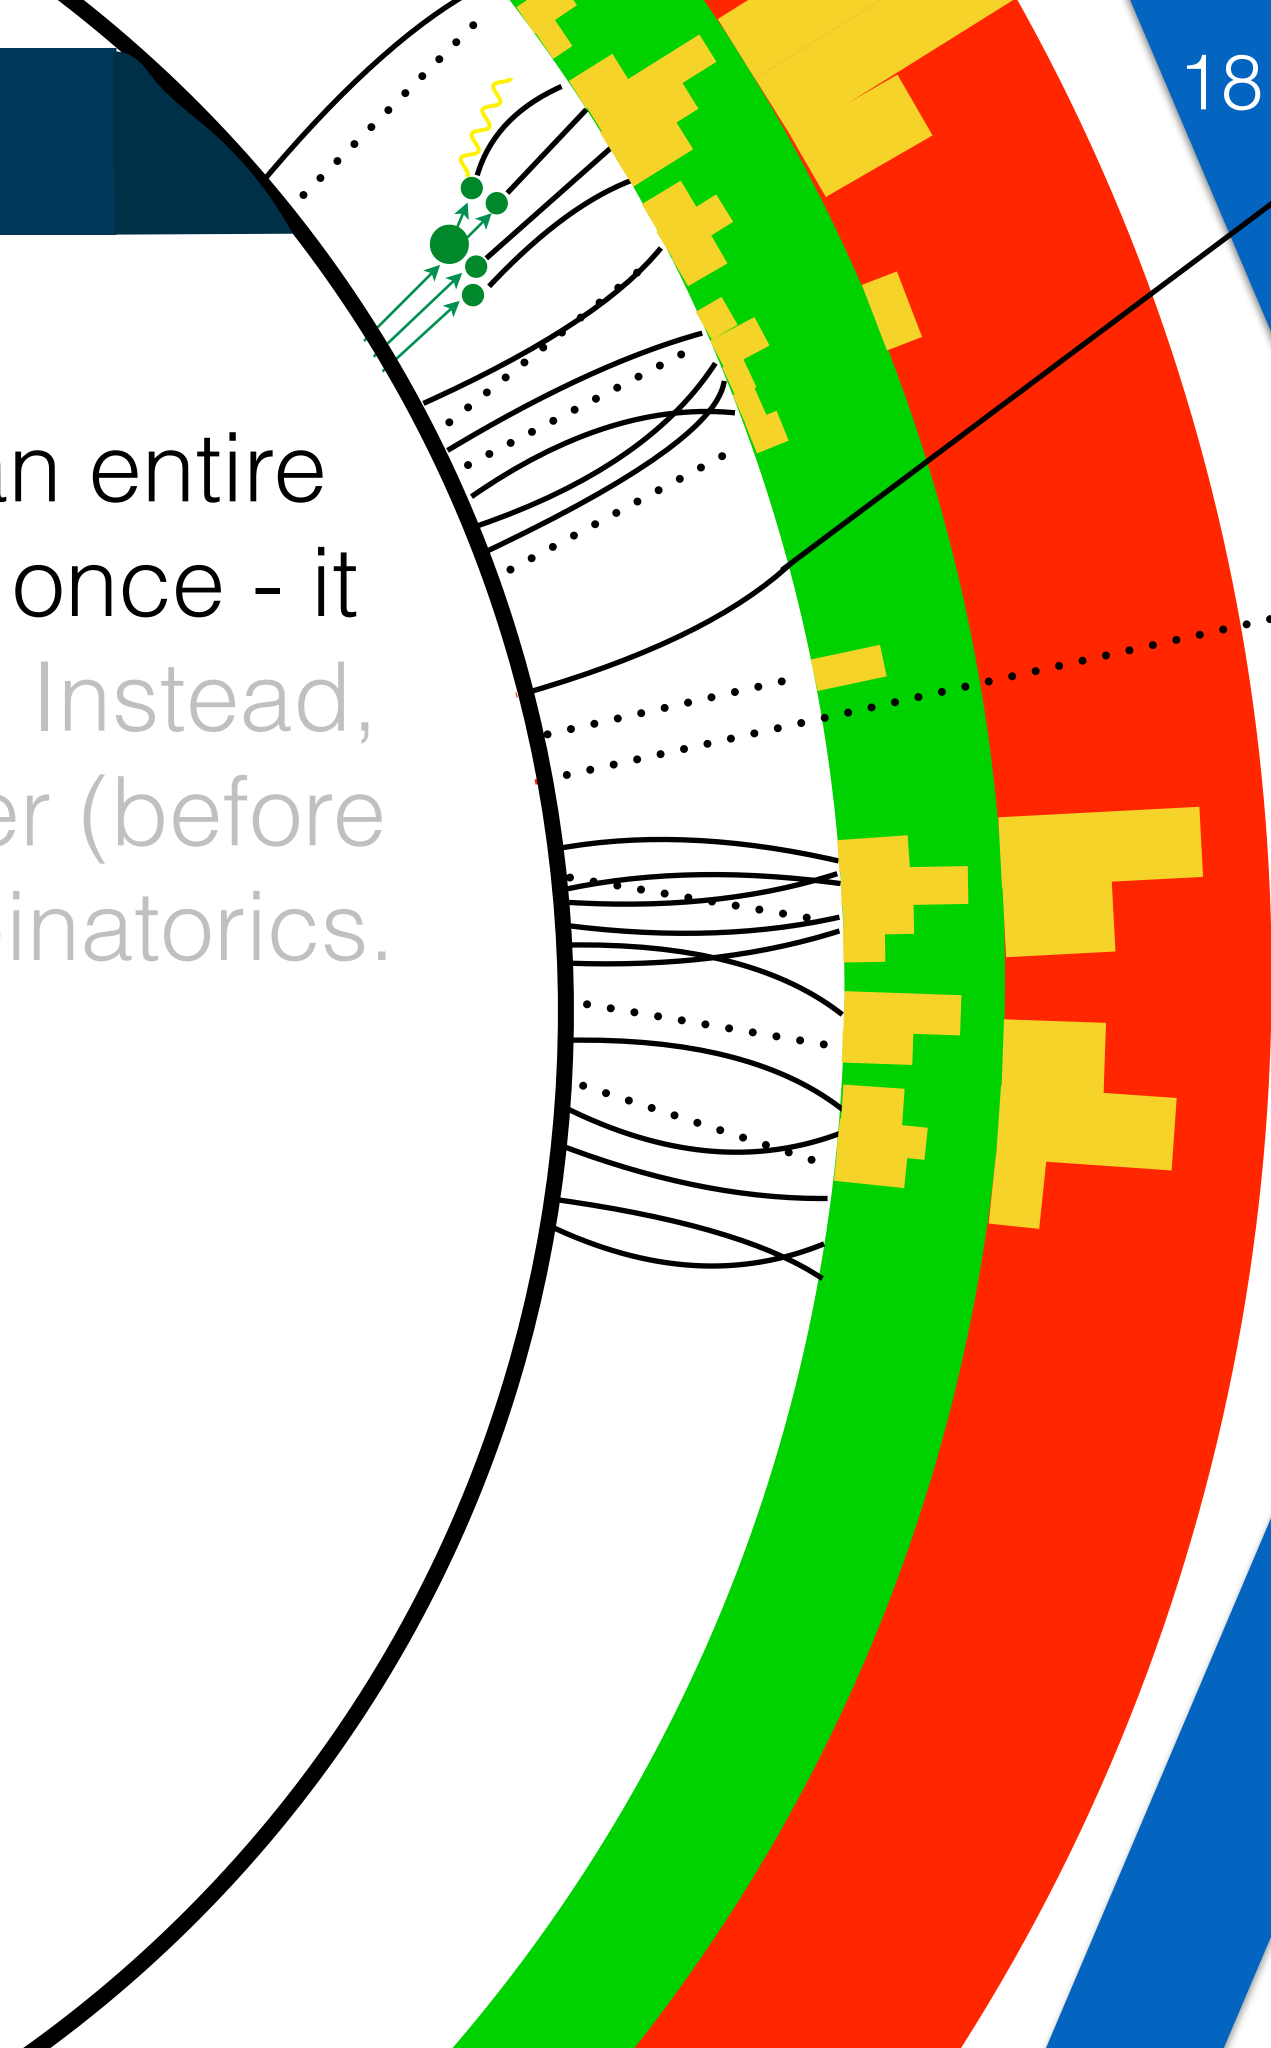

Factorization

We are not trying to generate an entire event (O(1000) particles)) all at once - it would be **very had to validate!** Instead, generate a single particle shower (before electronics) and appeal to combinatorics.

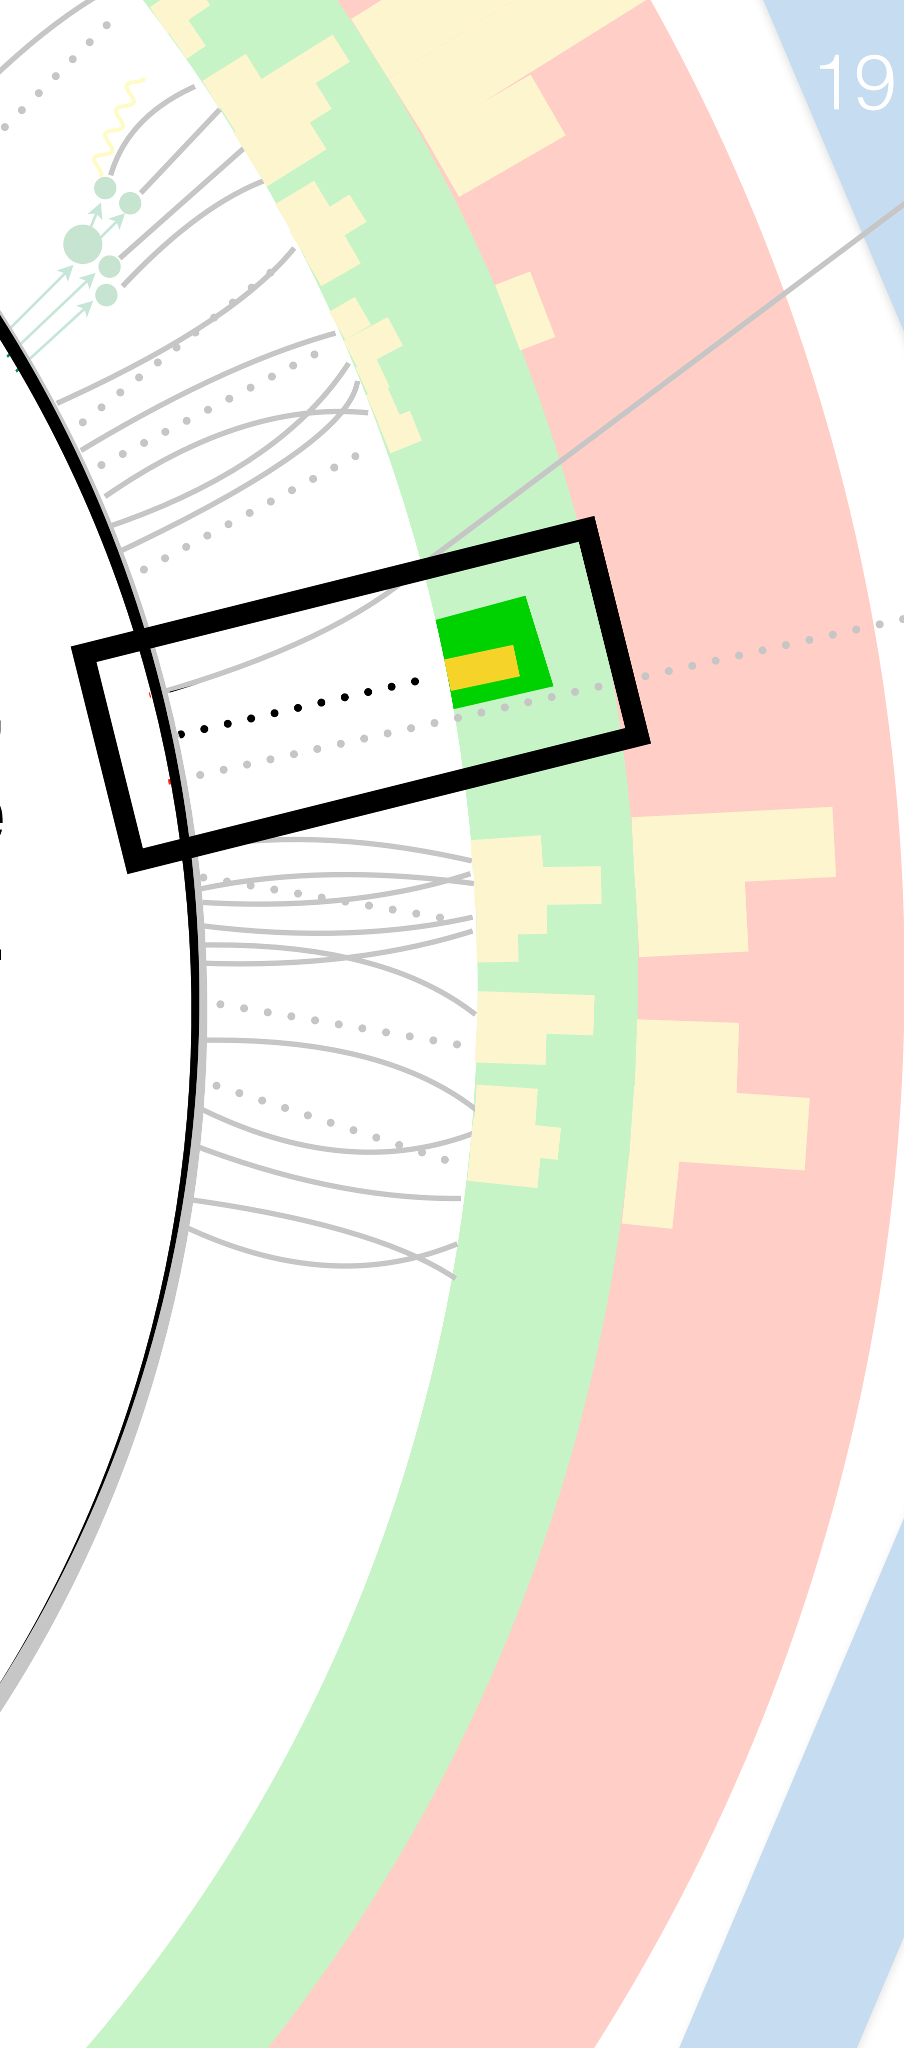

Factorization

We are not tryin[...]
event (O(1000) p[...]
would be **very ha**[...]
generate a single[...]
electronics) and a[...]

N.B. **calorimeter energy deposits** **factorize** (sum of the deposits is the deposit of the sum) but **digitization (w/ noise) does not!**

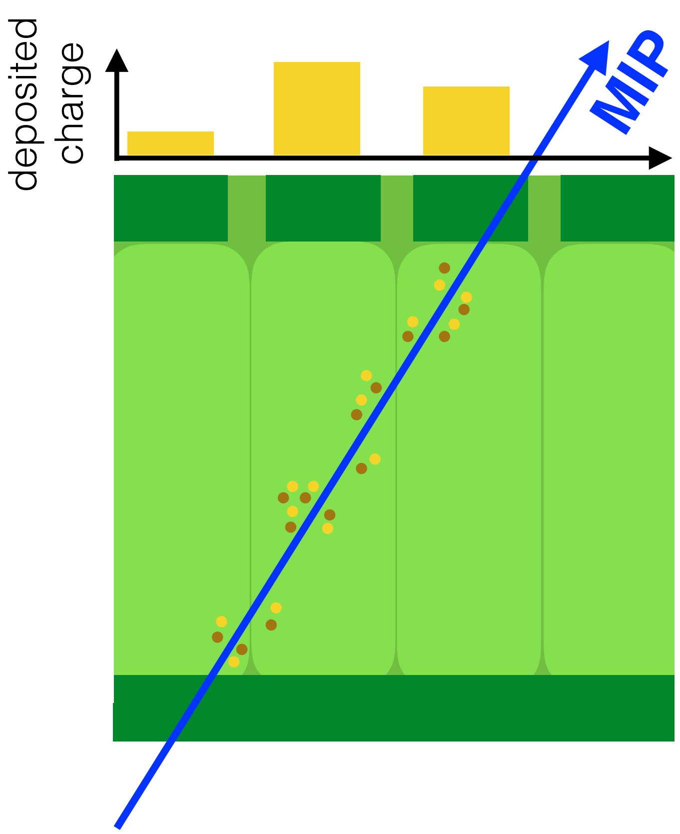

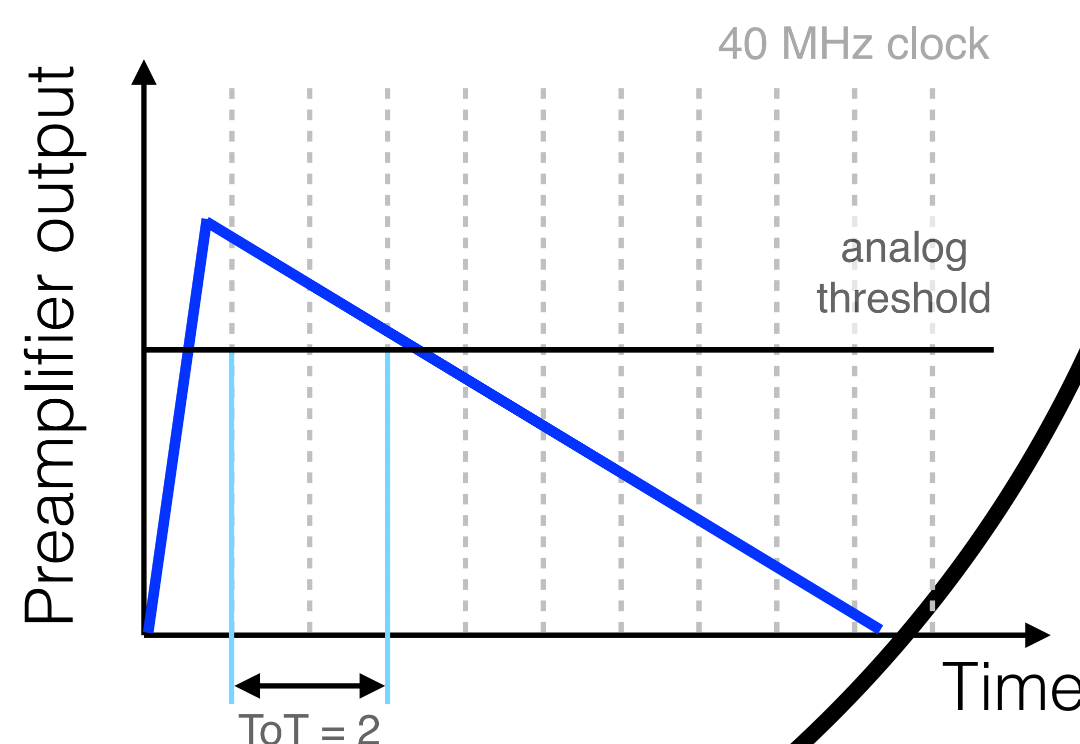

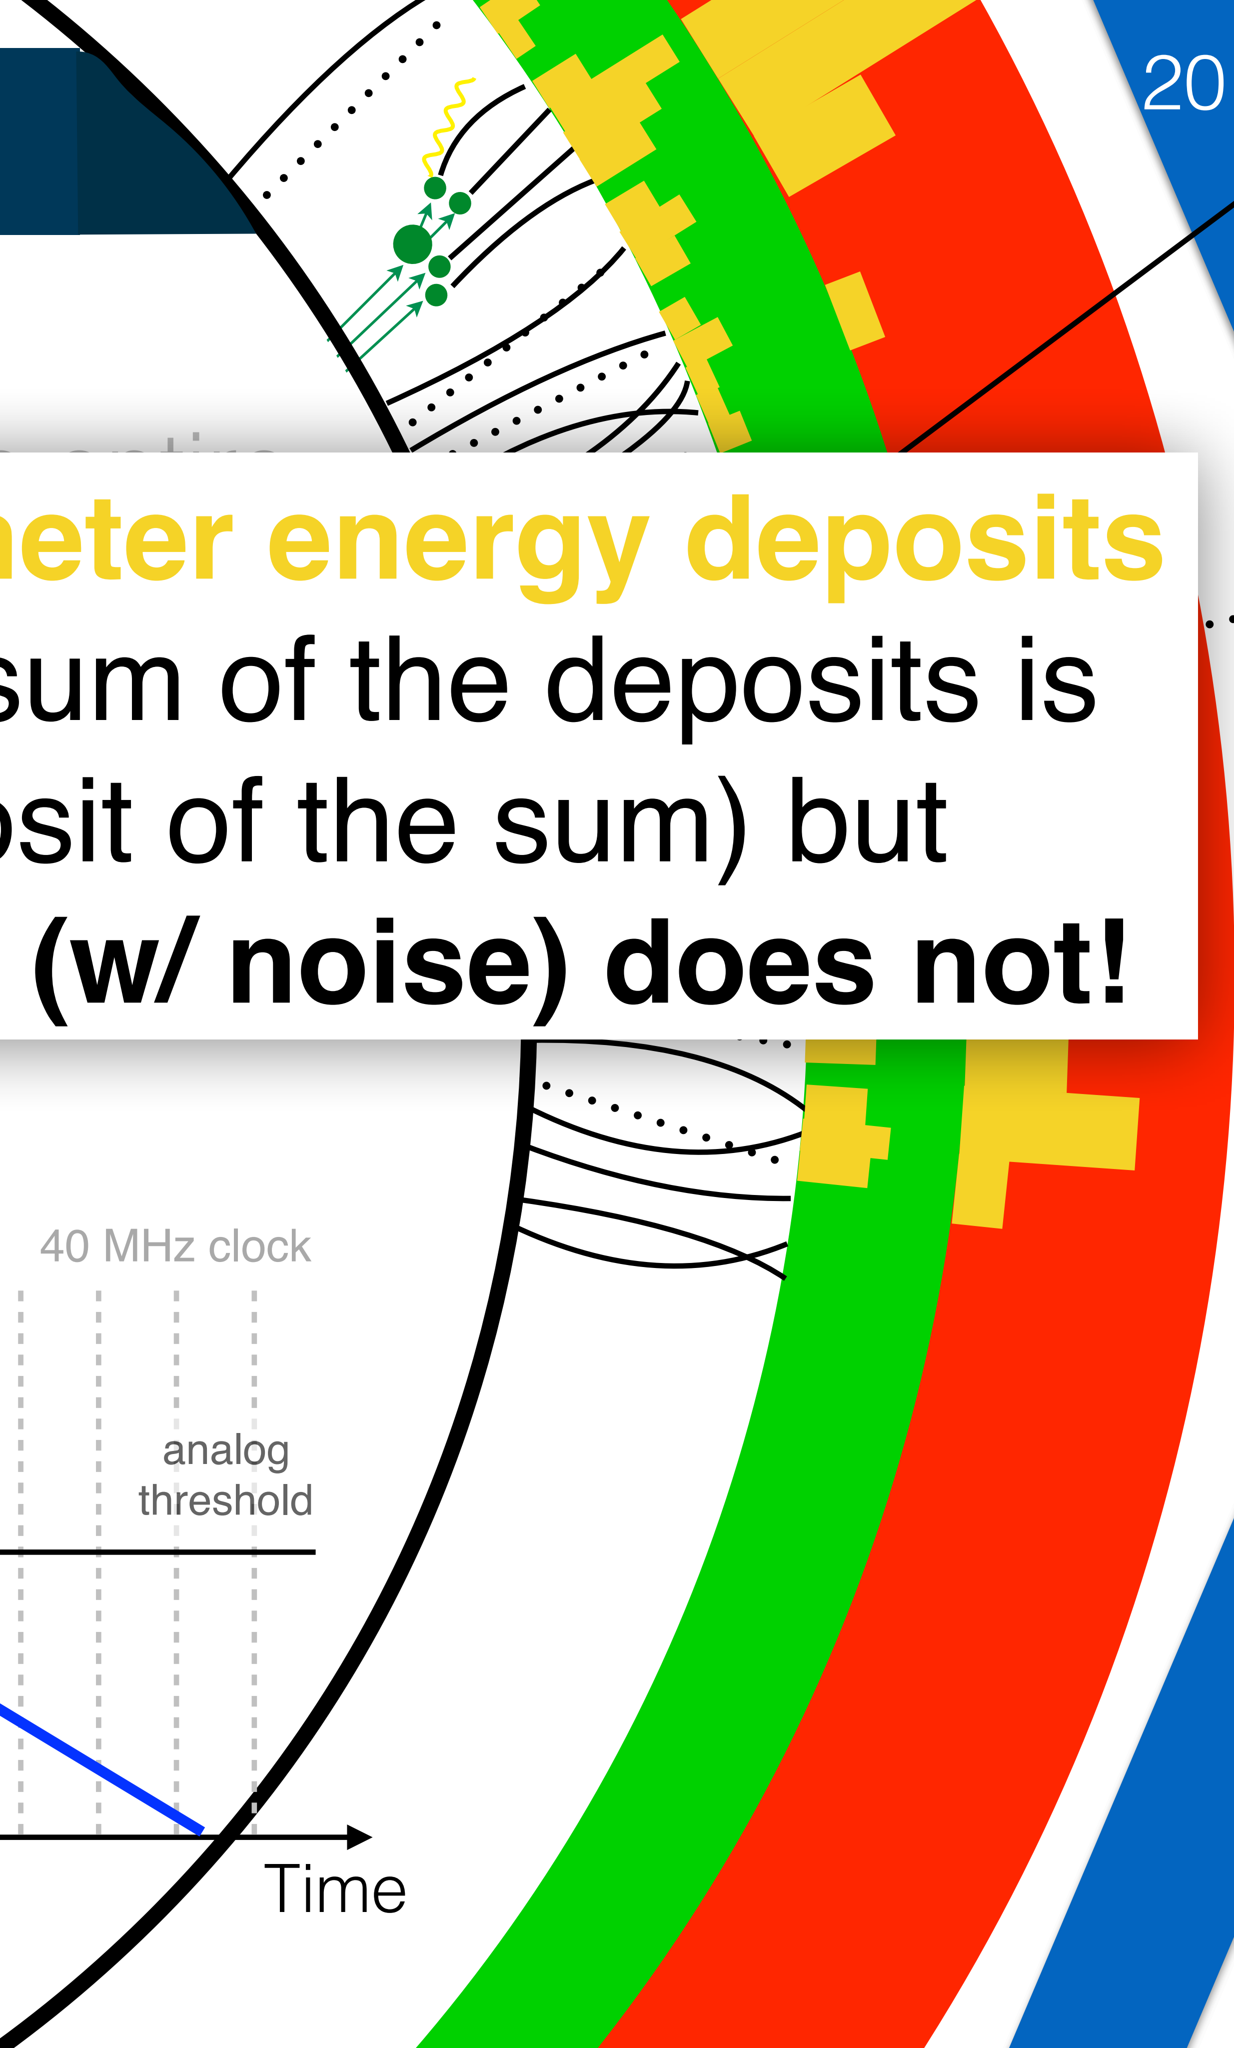

Now to the machine learning

A **generator** is nothing other than a function that maps random numbers to structure.

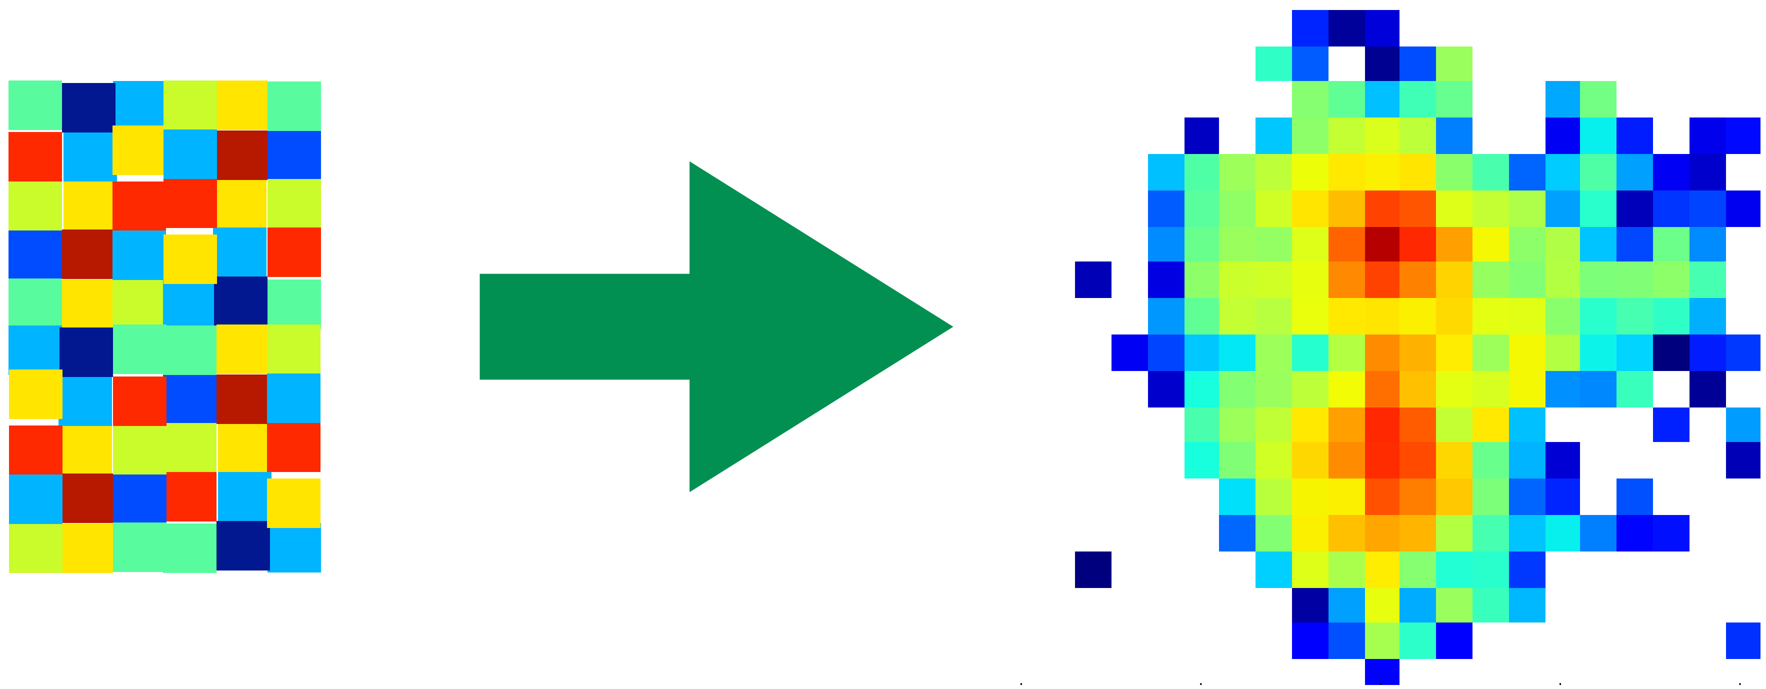

Our structure: **calorimeter images**

Calorimeter images

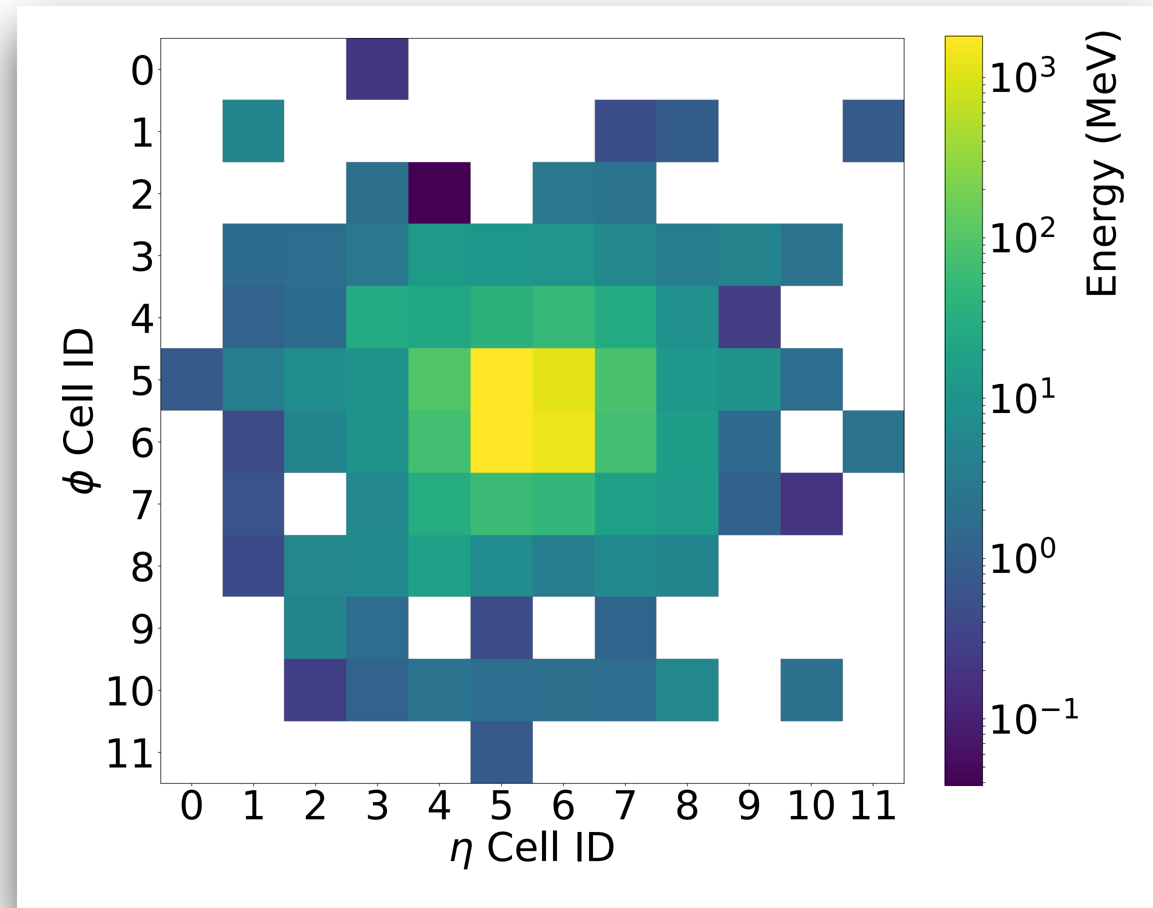

Grayscale images:
Pixel intensity =
energy deposited

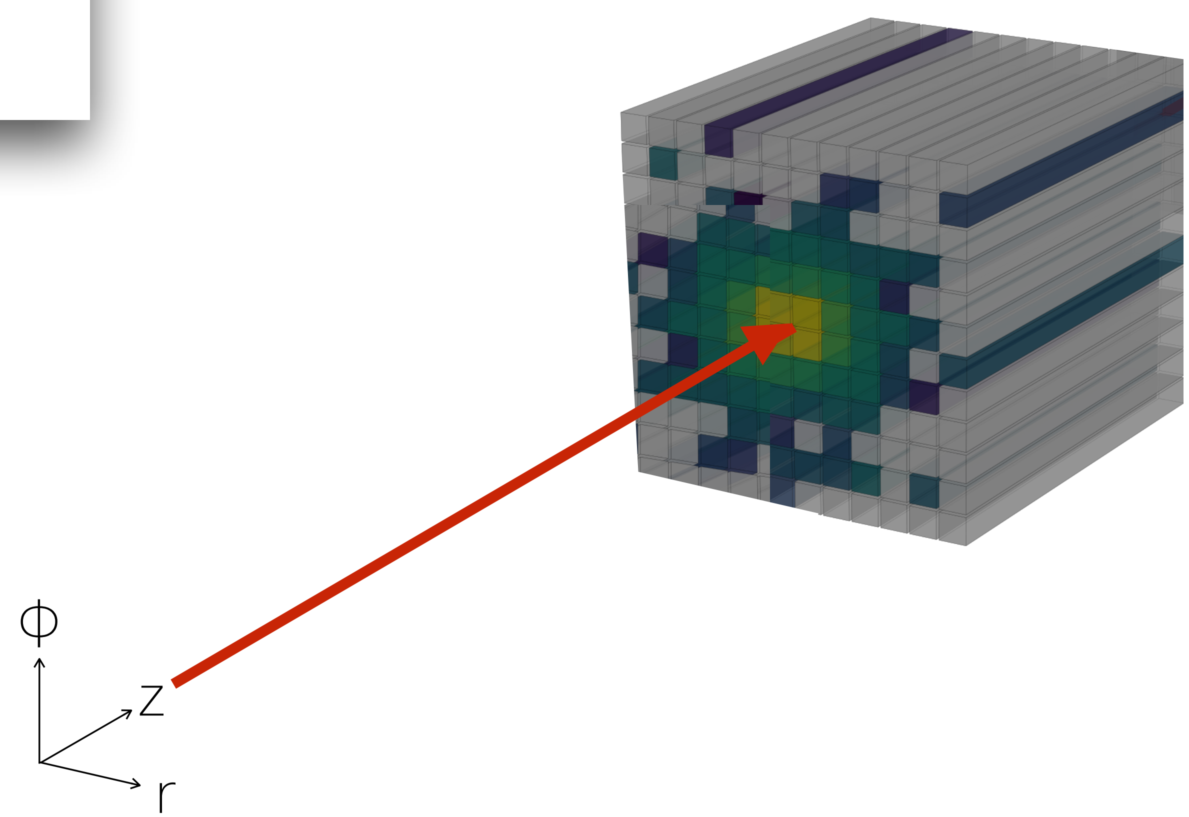

Calorimeter images

Challenge: **multiple layers** with **non-uniform granularity** and a **causal relationship**?

N.B. images are O(1000) dimensional

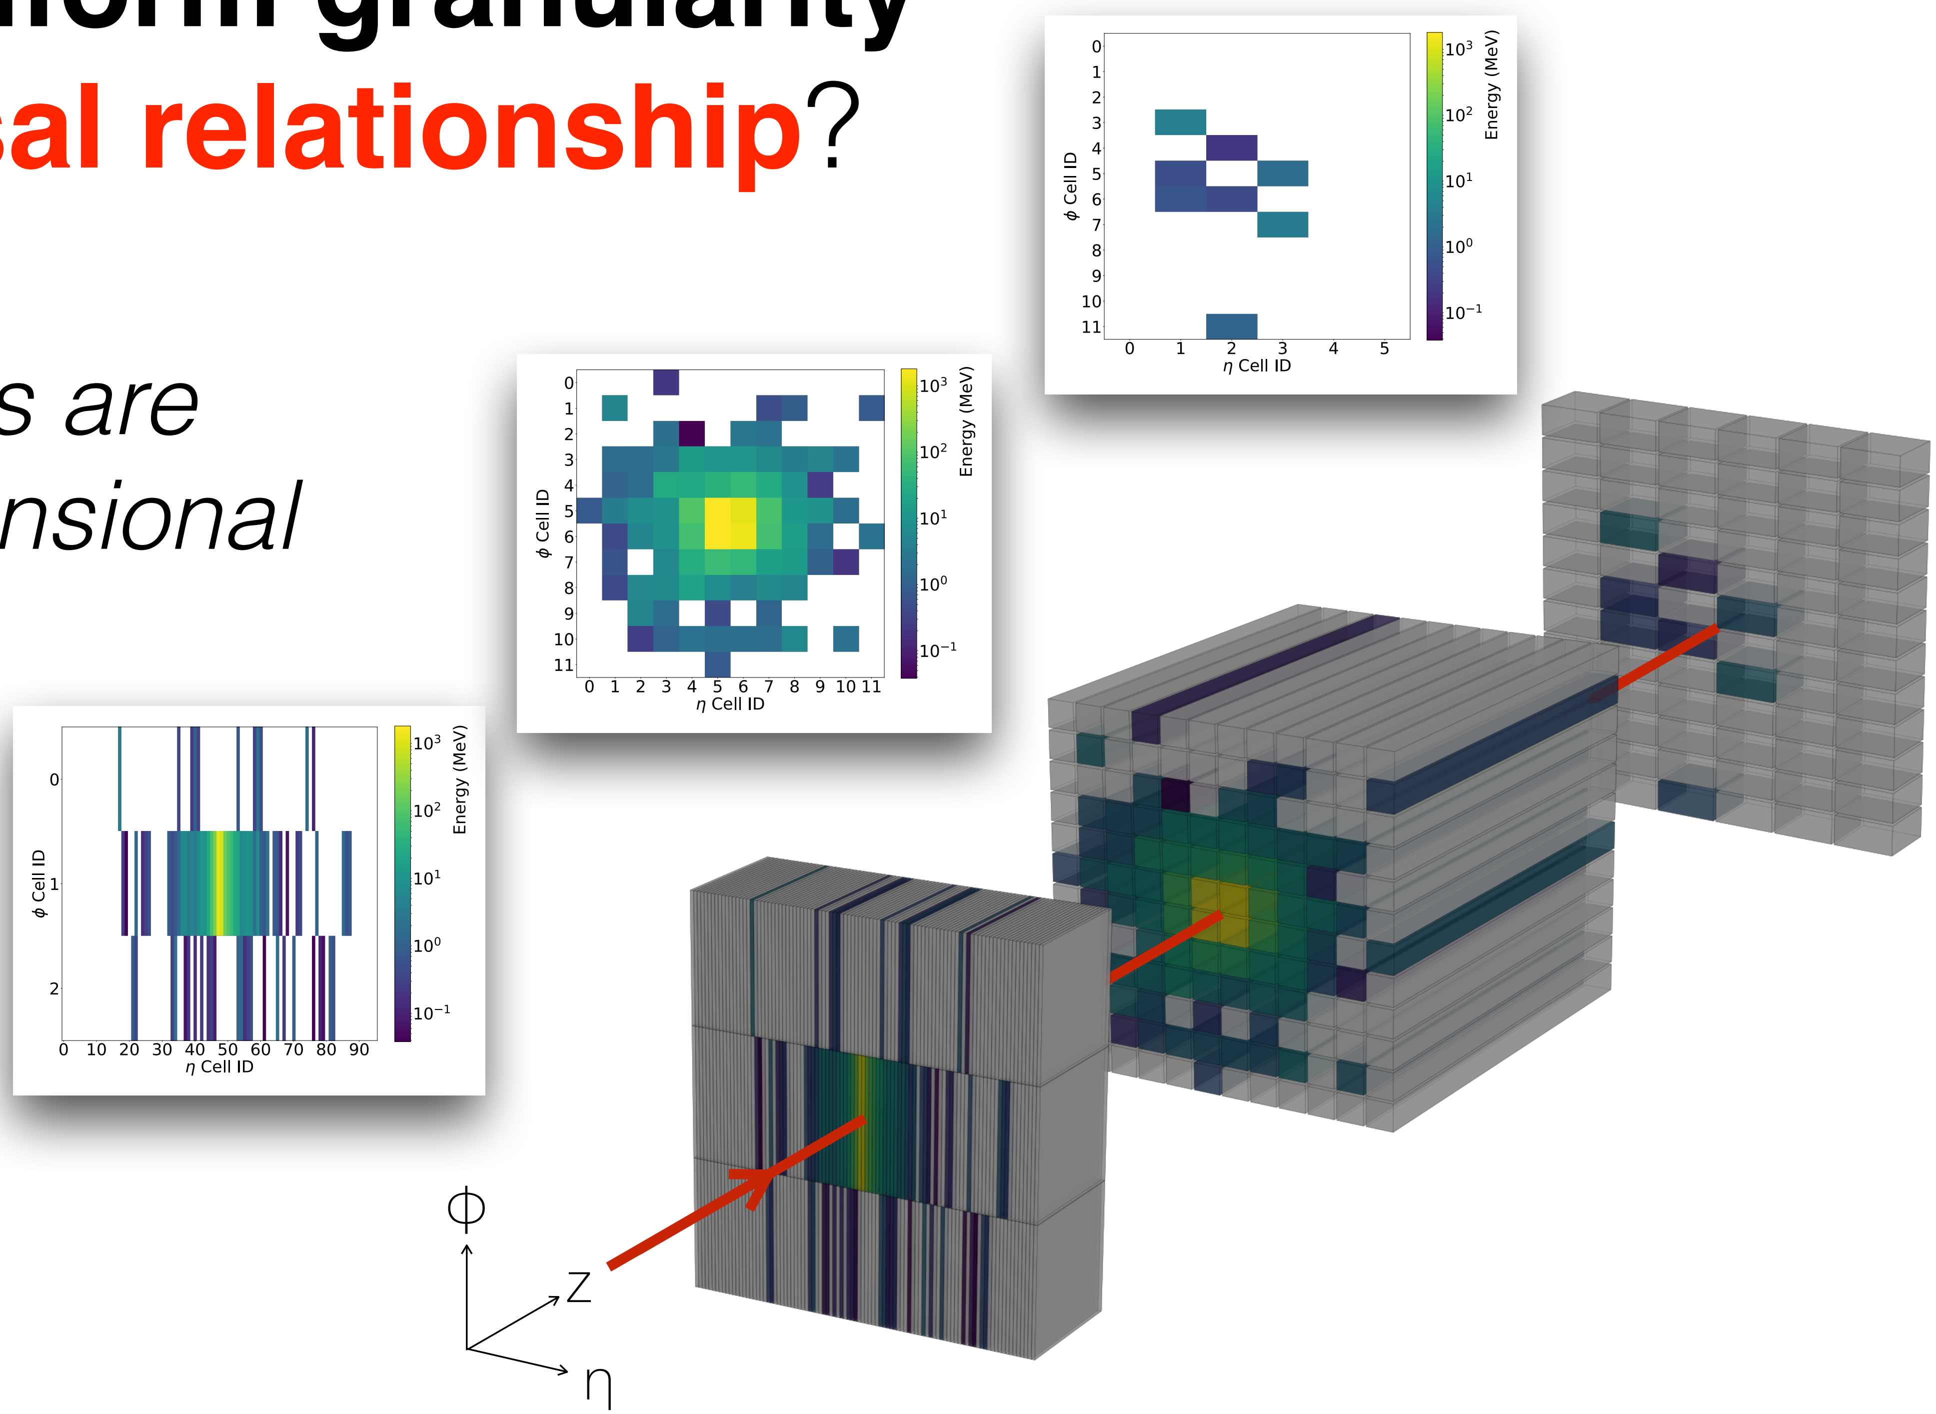

Reminder: GANs

Generative Adversarial Networks (GAN):
A two-network game where one ***maps noise to images*** *and one* ***classifies images as fake or real****.*

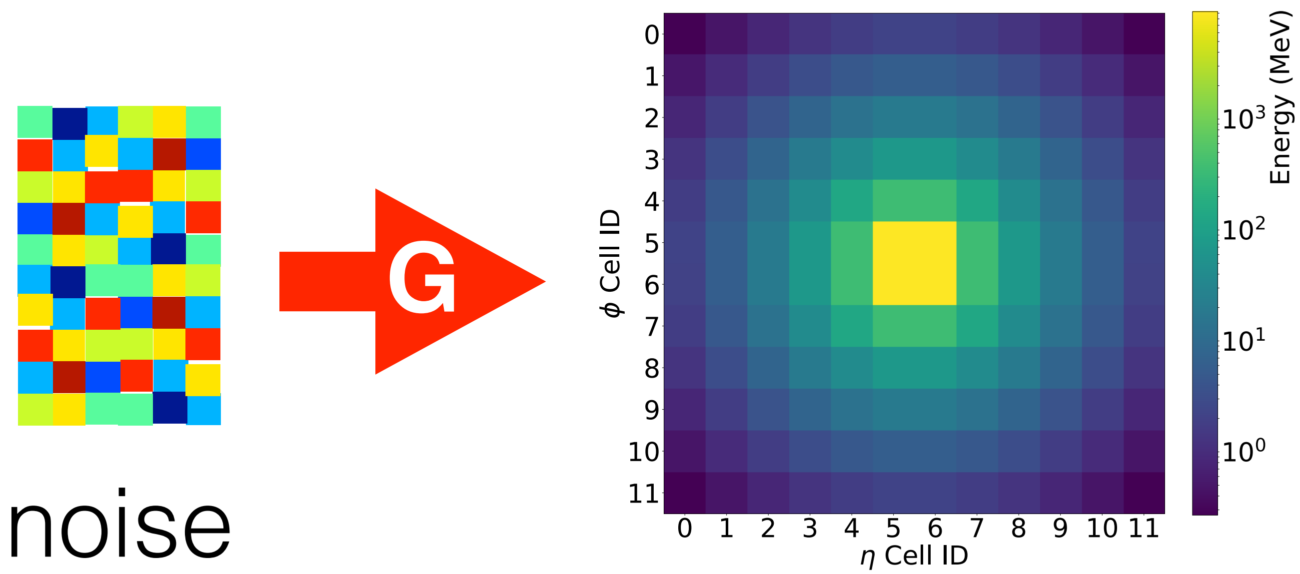

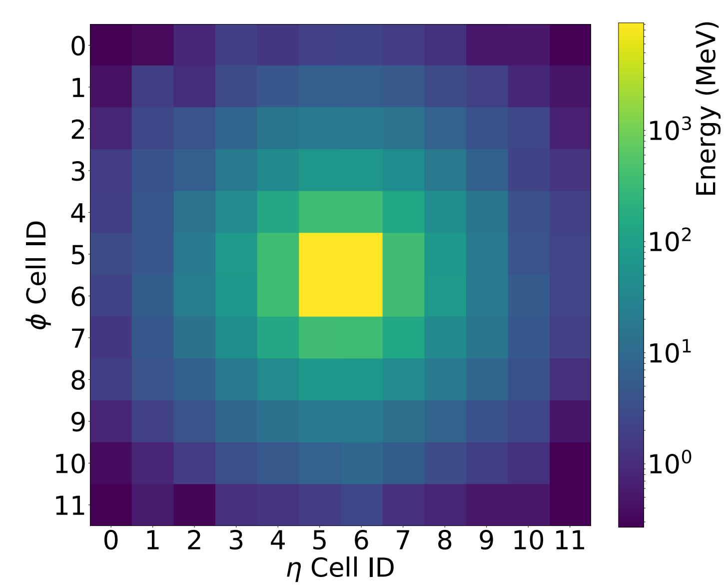

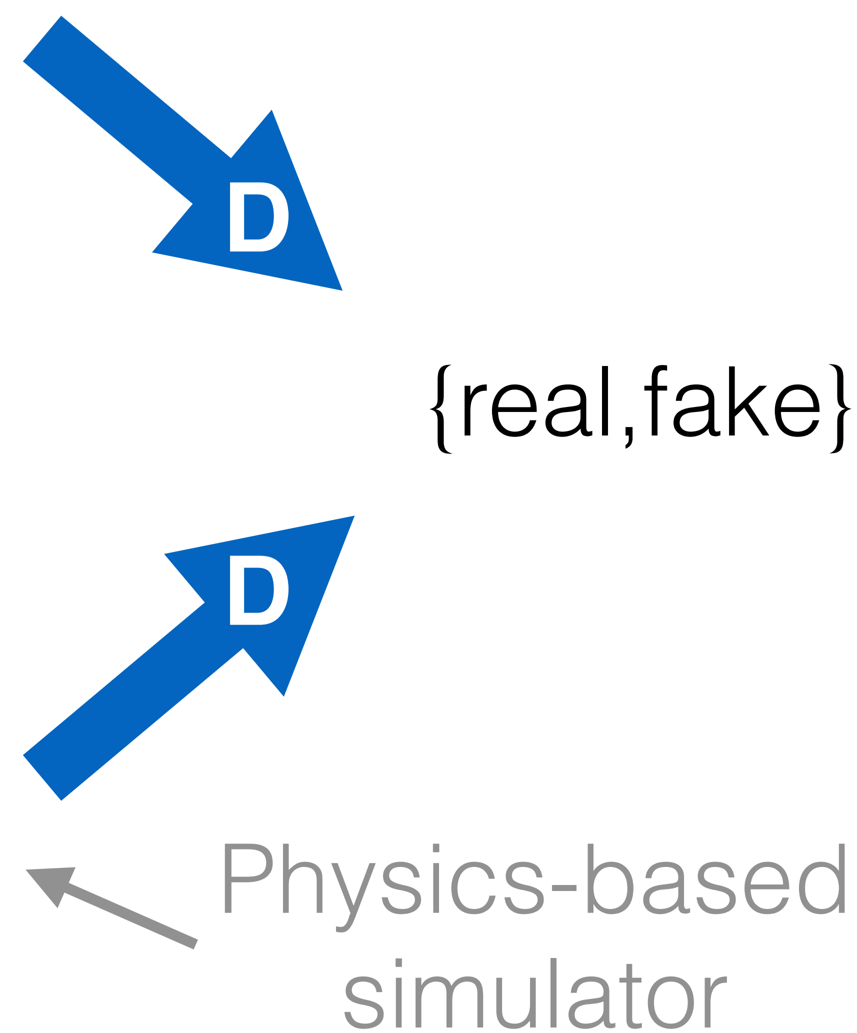

When **D** is maximally confused, **G** will be a good generator

Introducing CaloGAN

One image per calo layer

One network per particle type; input particle energy

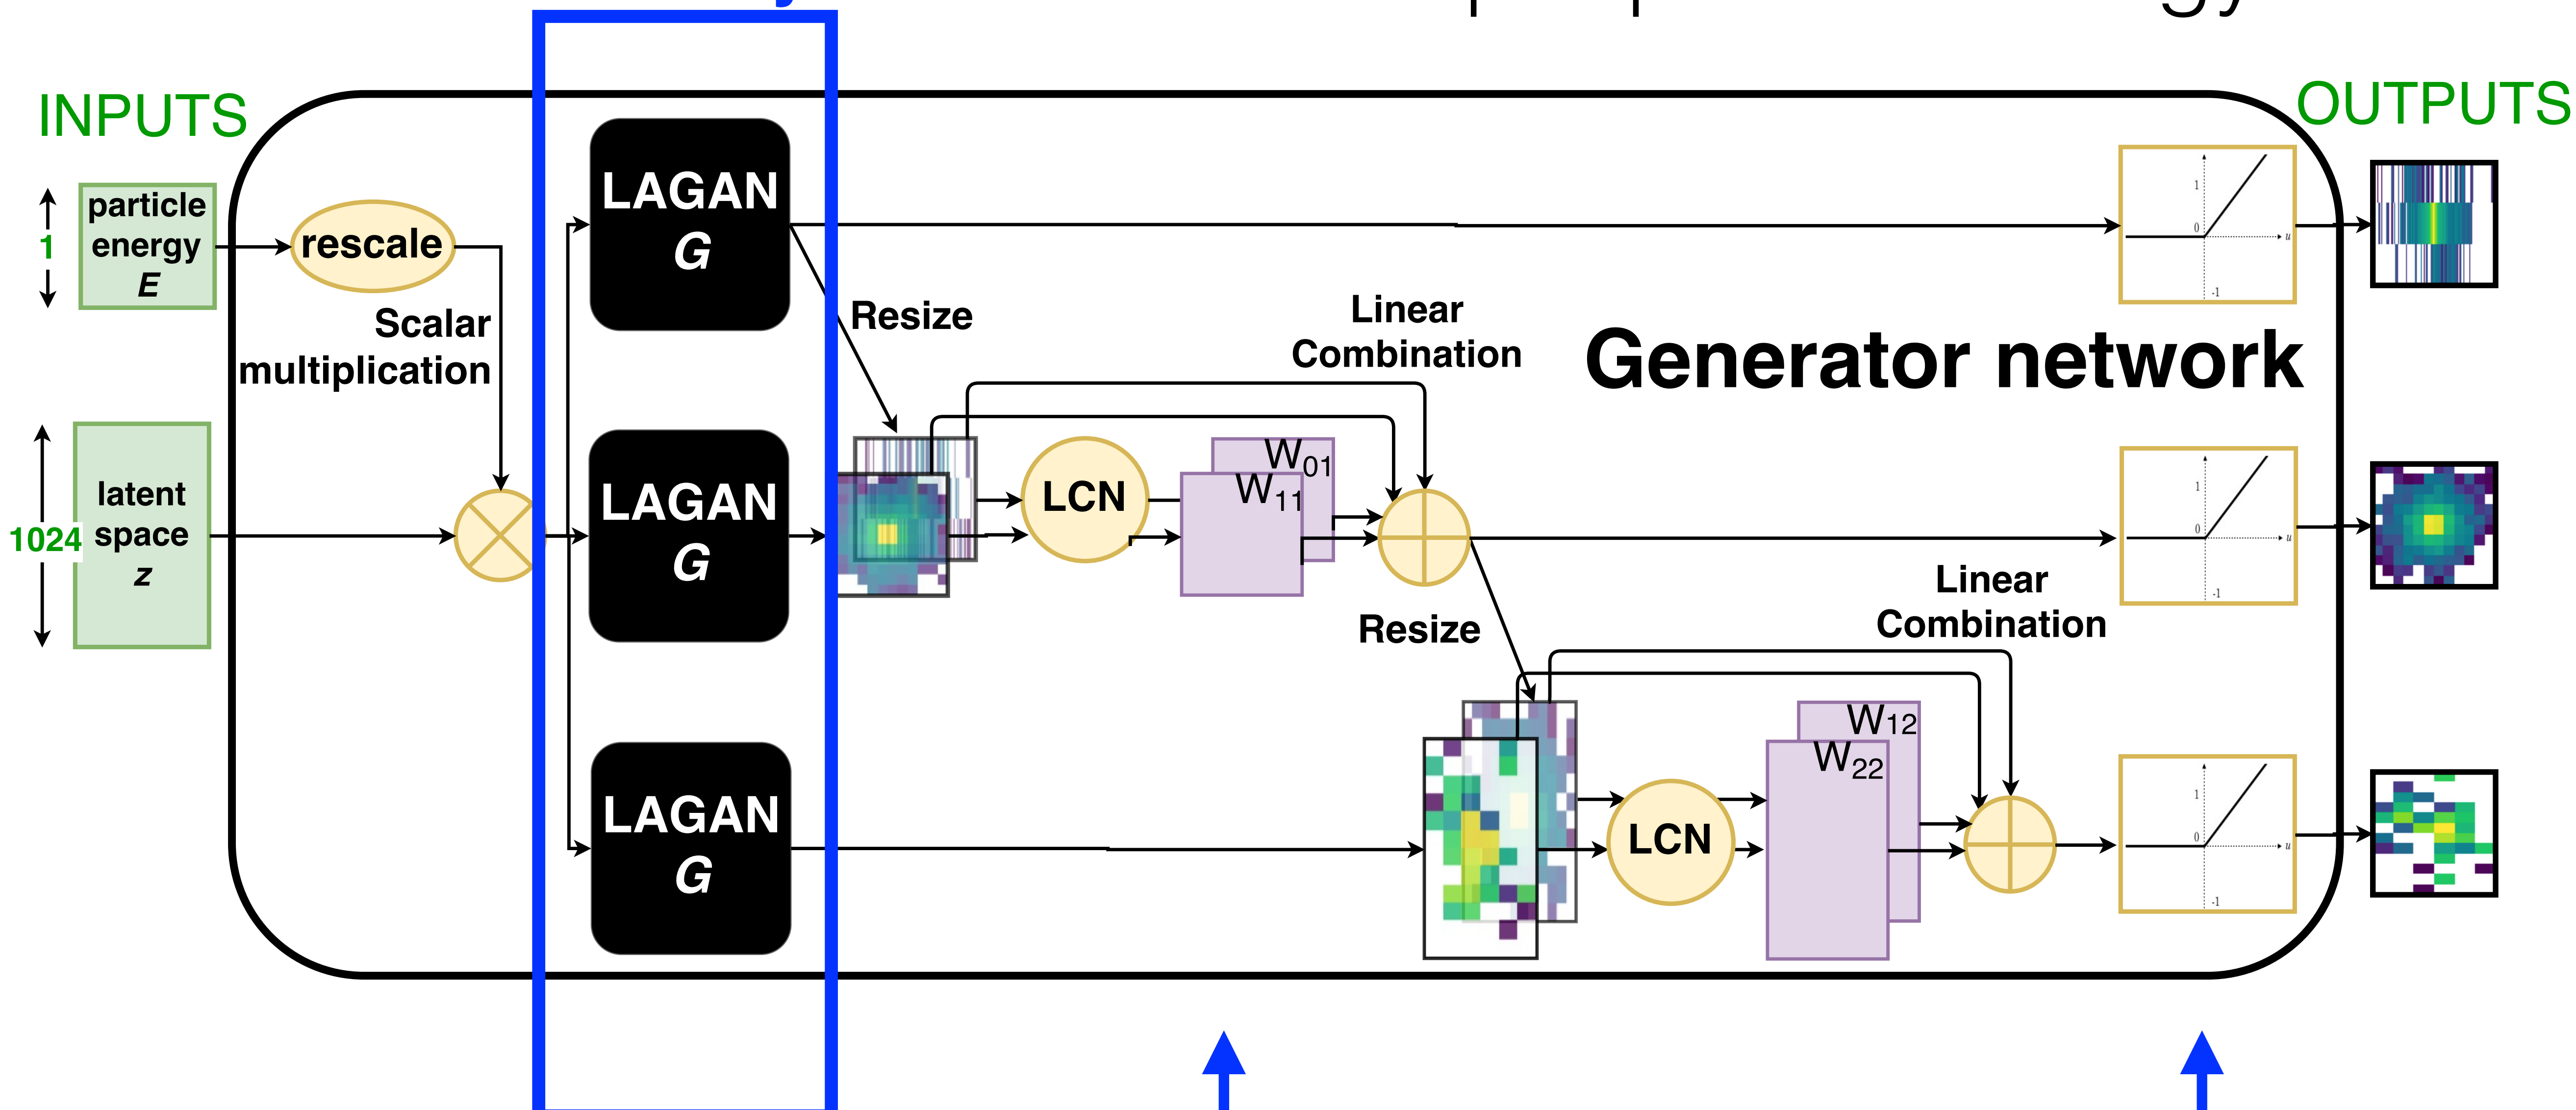

use layer i as input to layer i+1

ReLU to encourage sparsity

Introducing CaloGAN

Mode collapse: learns to generate one part of the distribution well, but leaves out other parts.

help avoid 'mode collapse'

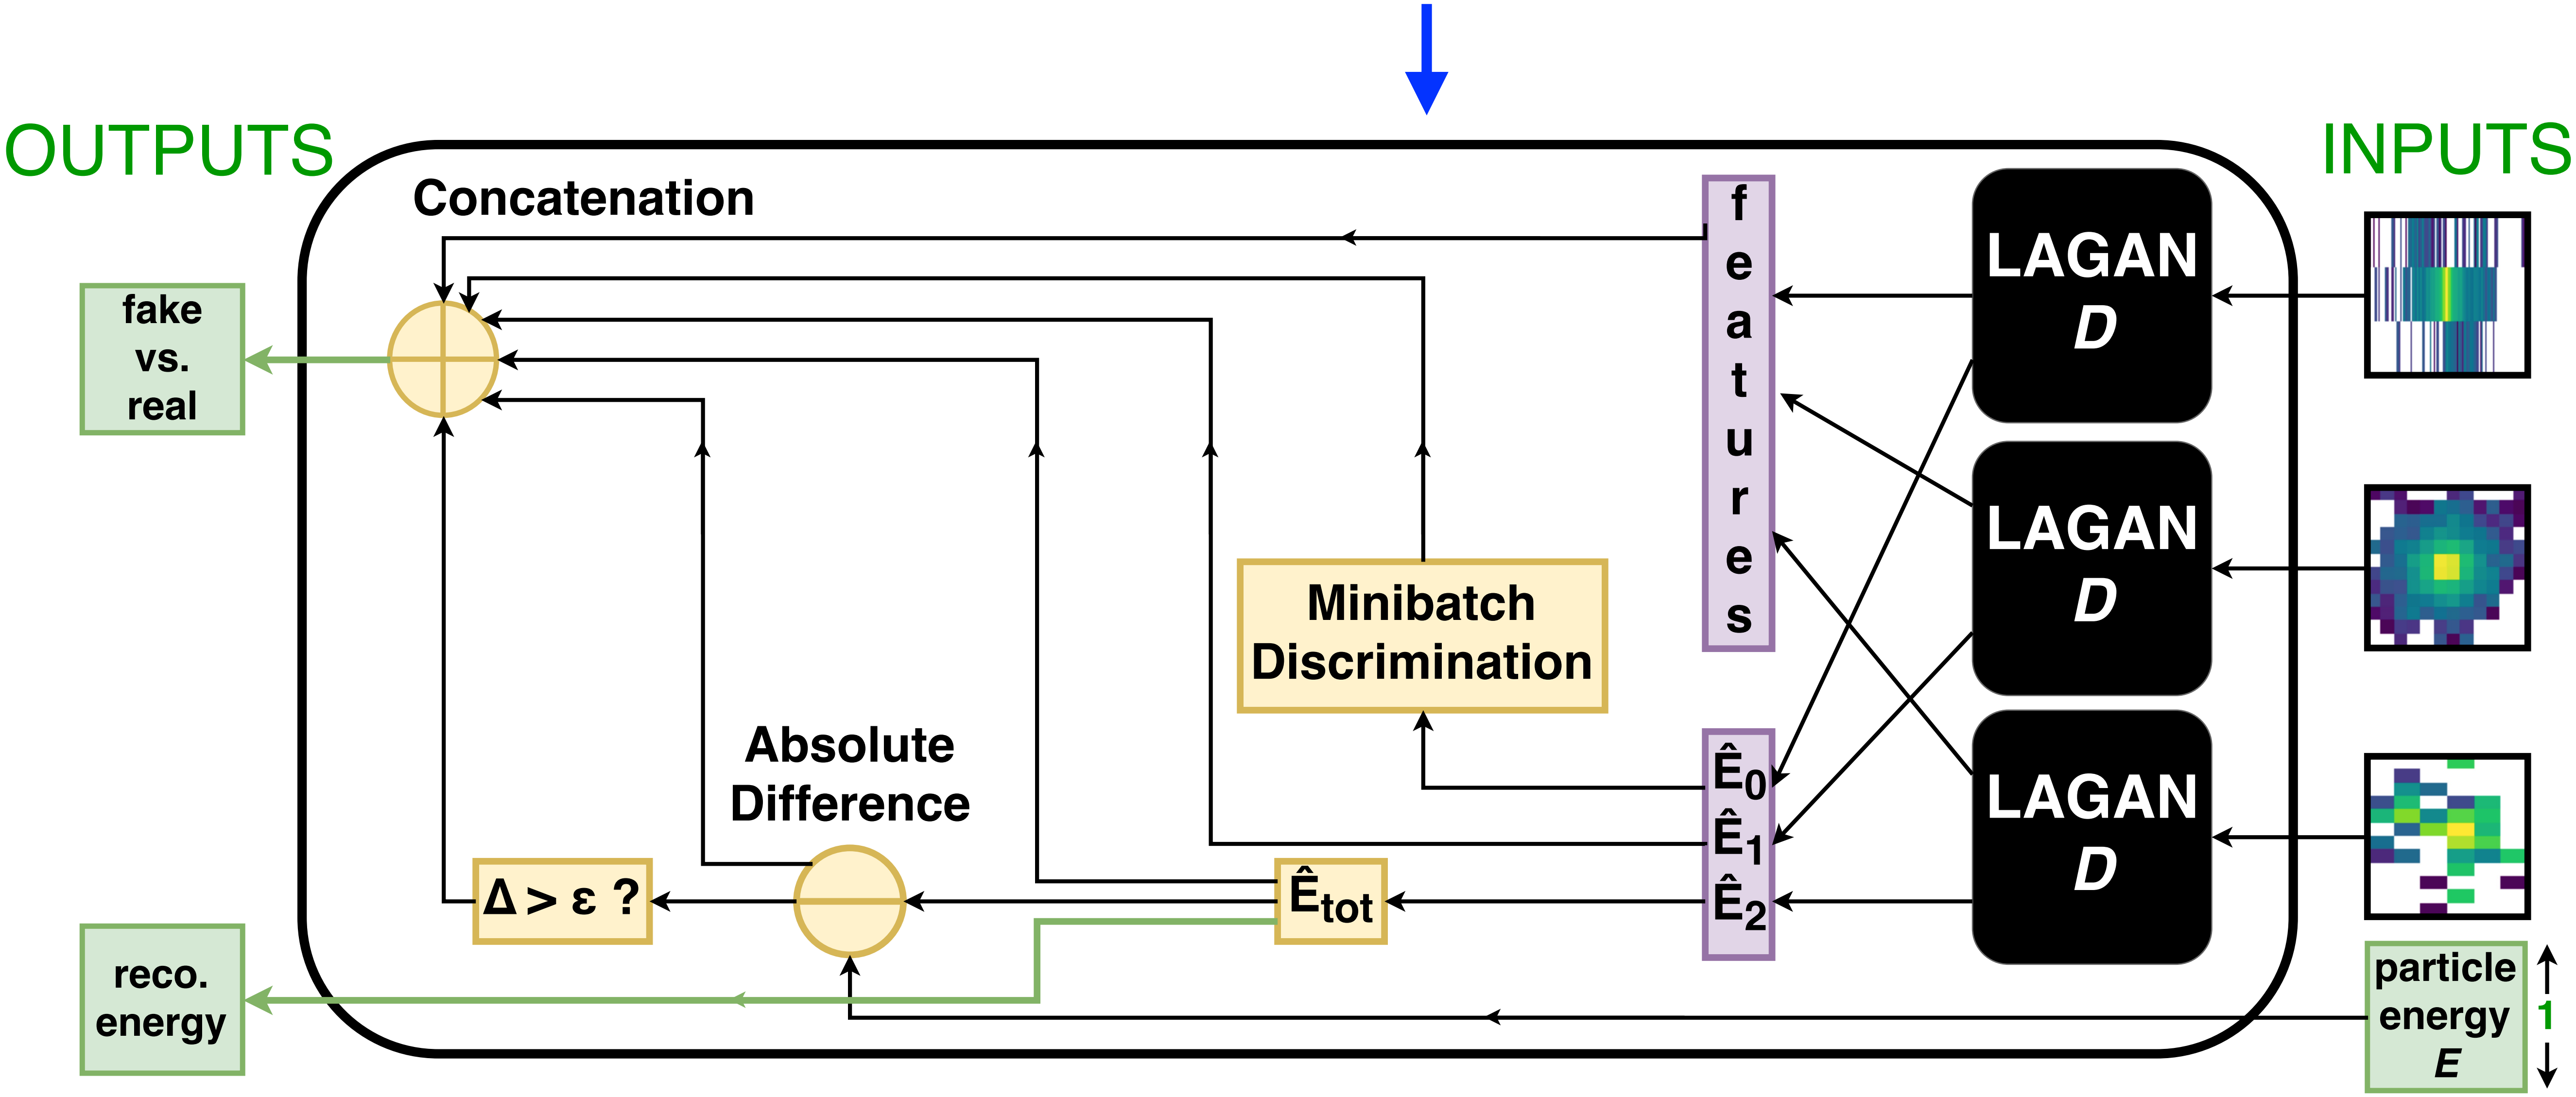

Discriminator network

Locally connected layers

Due to the structure of the problem, we do not have translation invariance.

Classification studies found fully connected networks outperformed CNNs

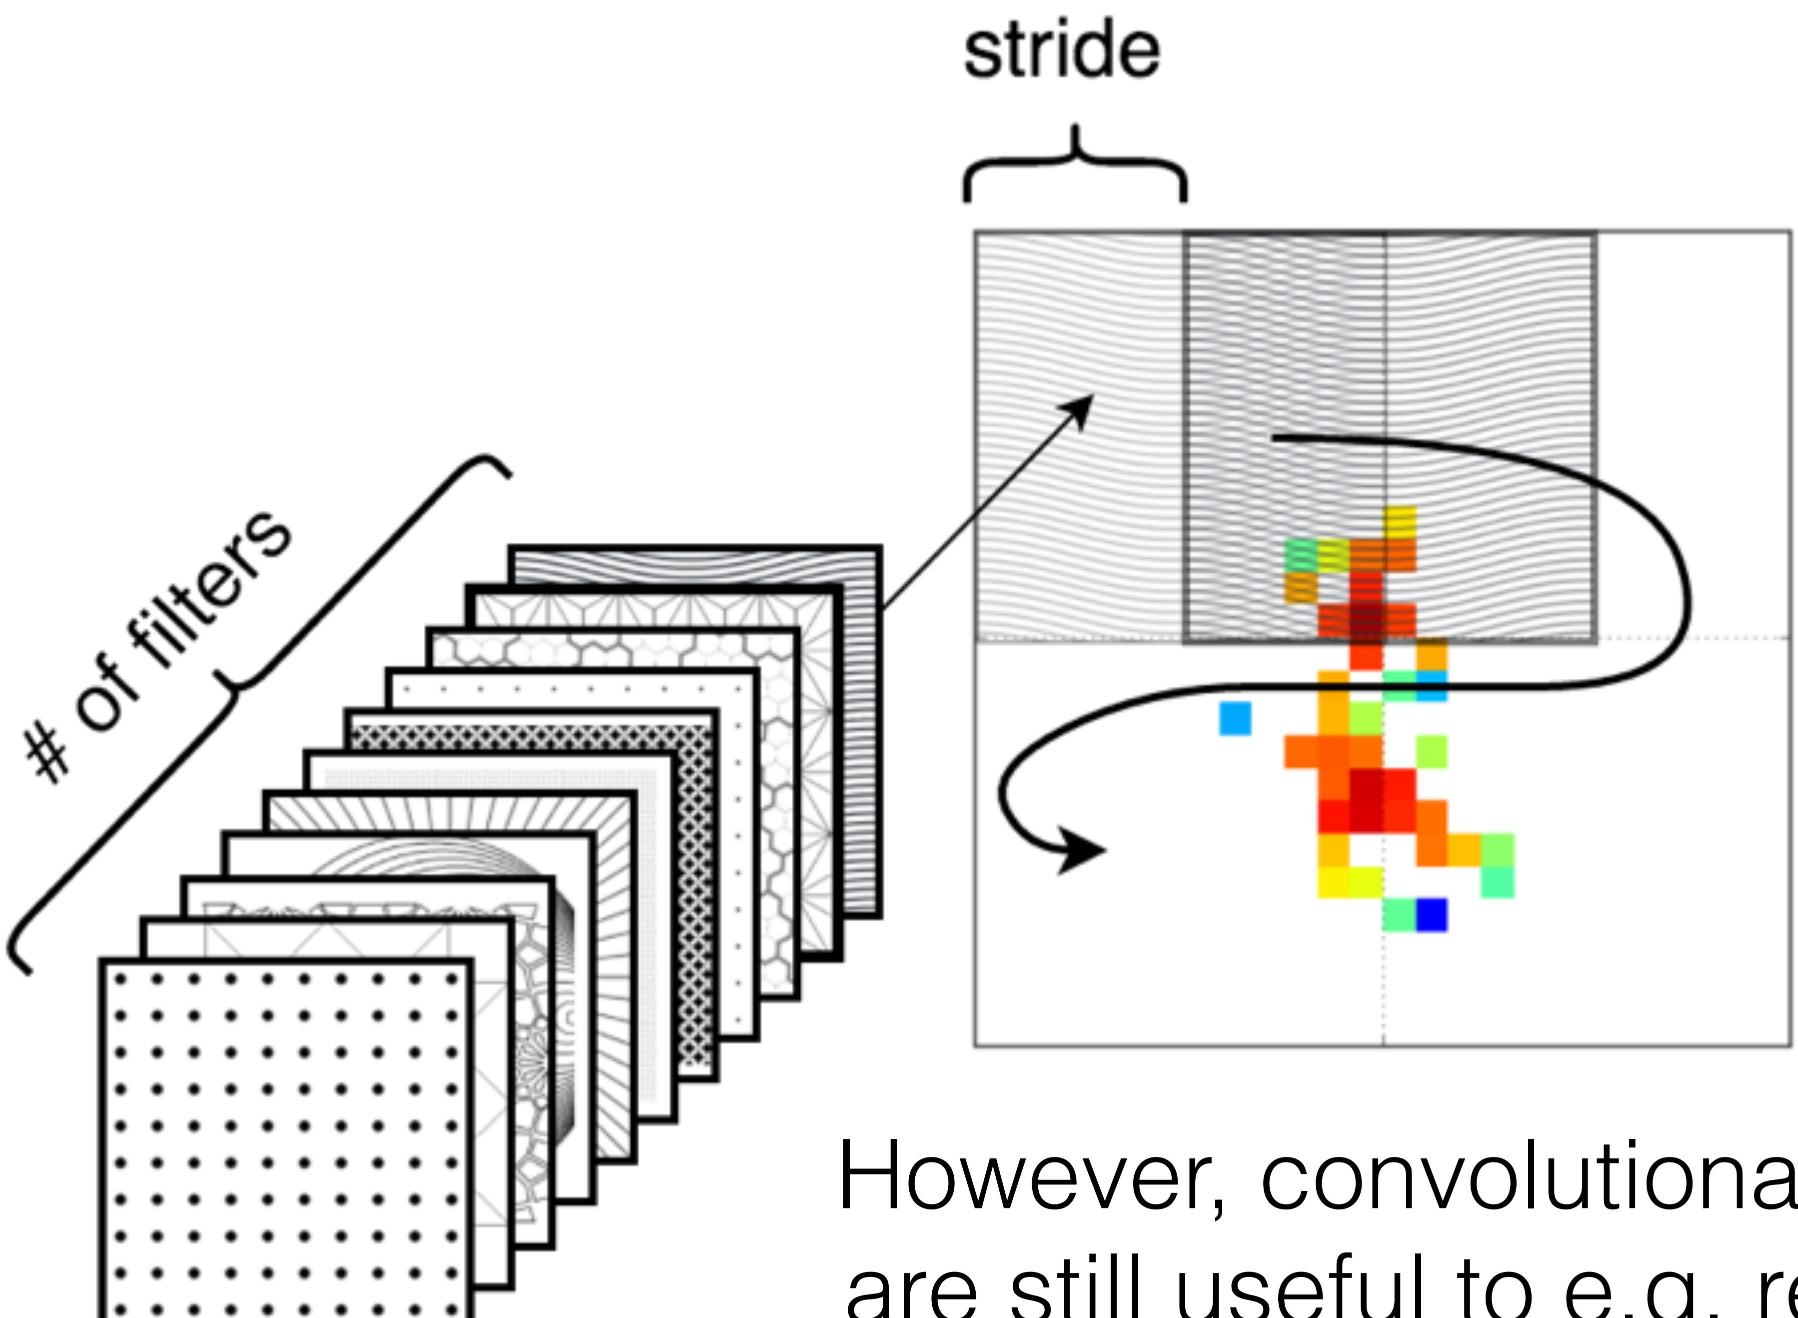

However, convolutional-like architectures are still useful to e.g. reduce parameters

Locally connected layers

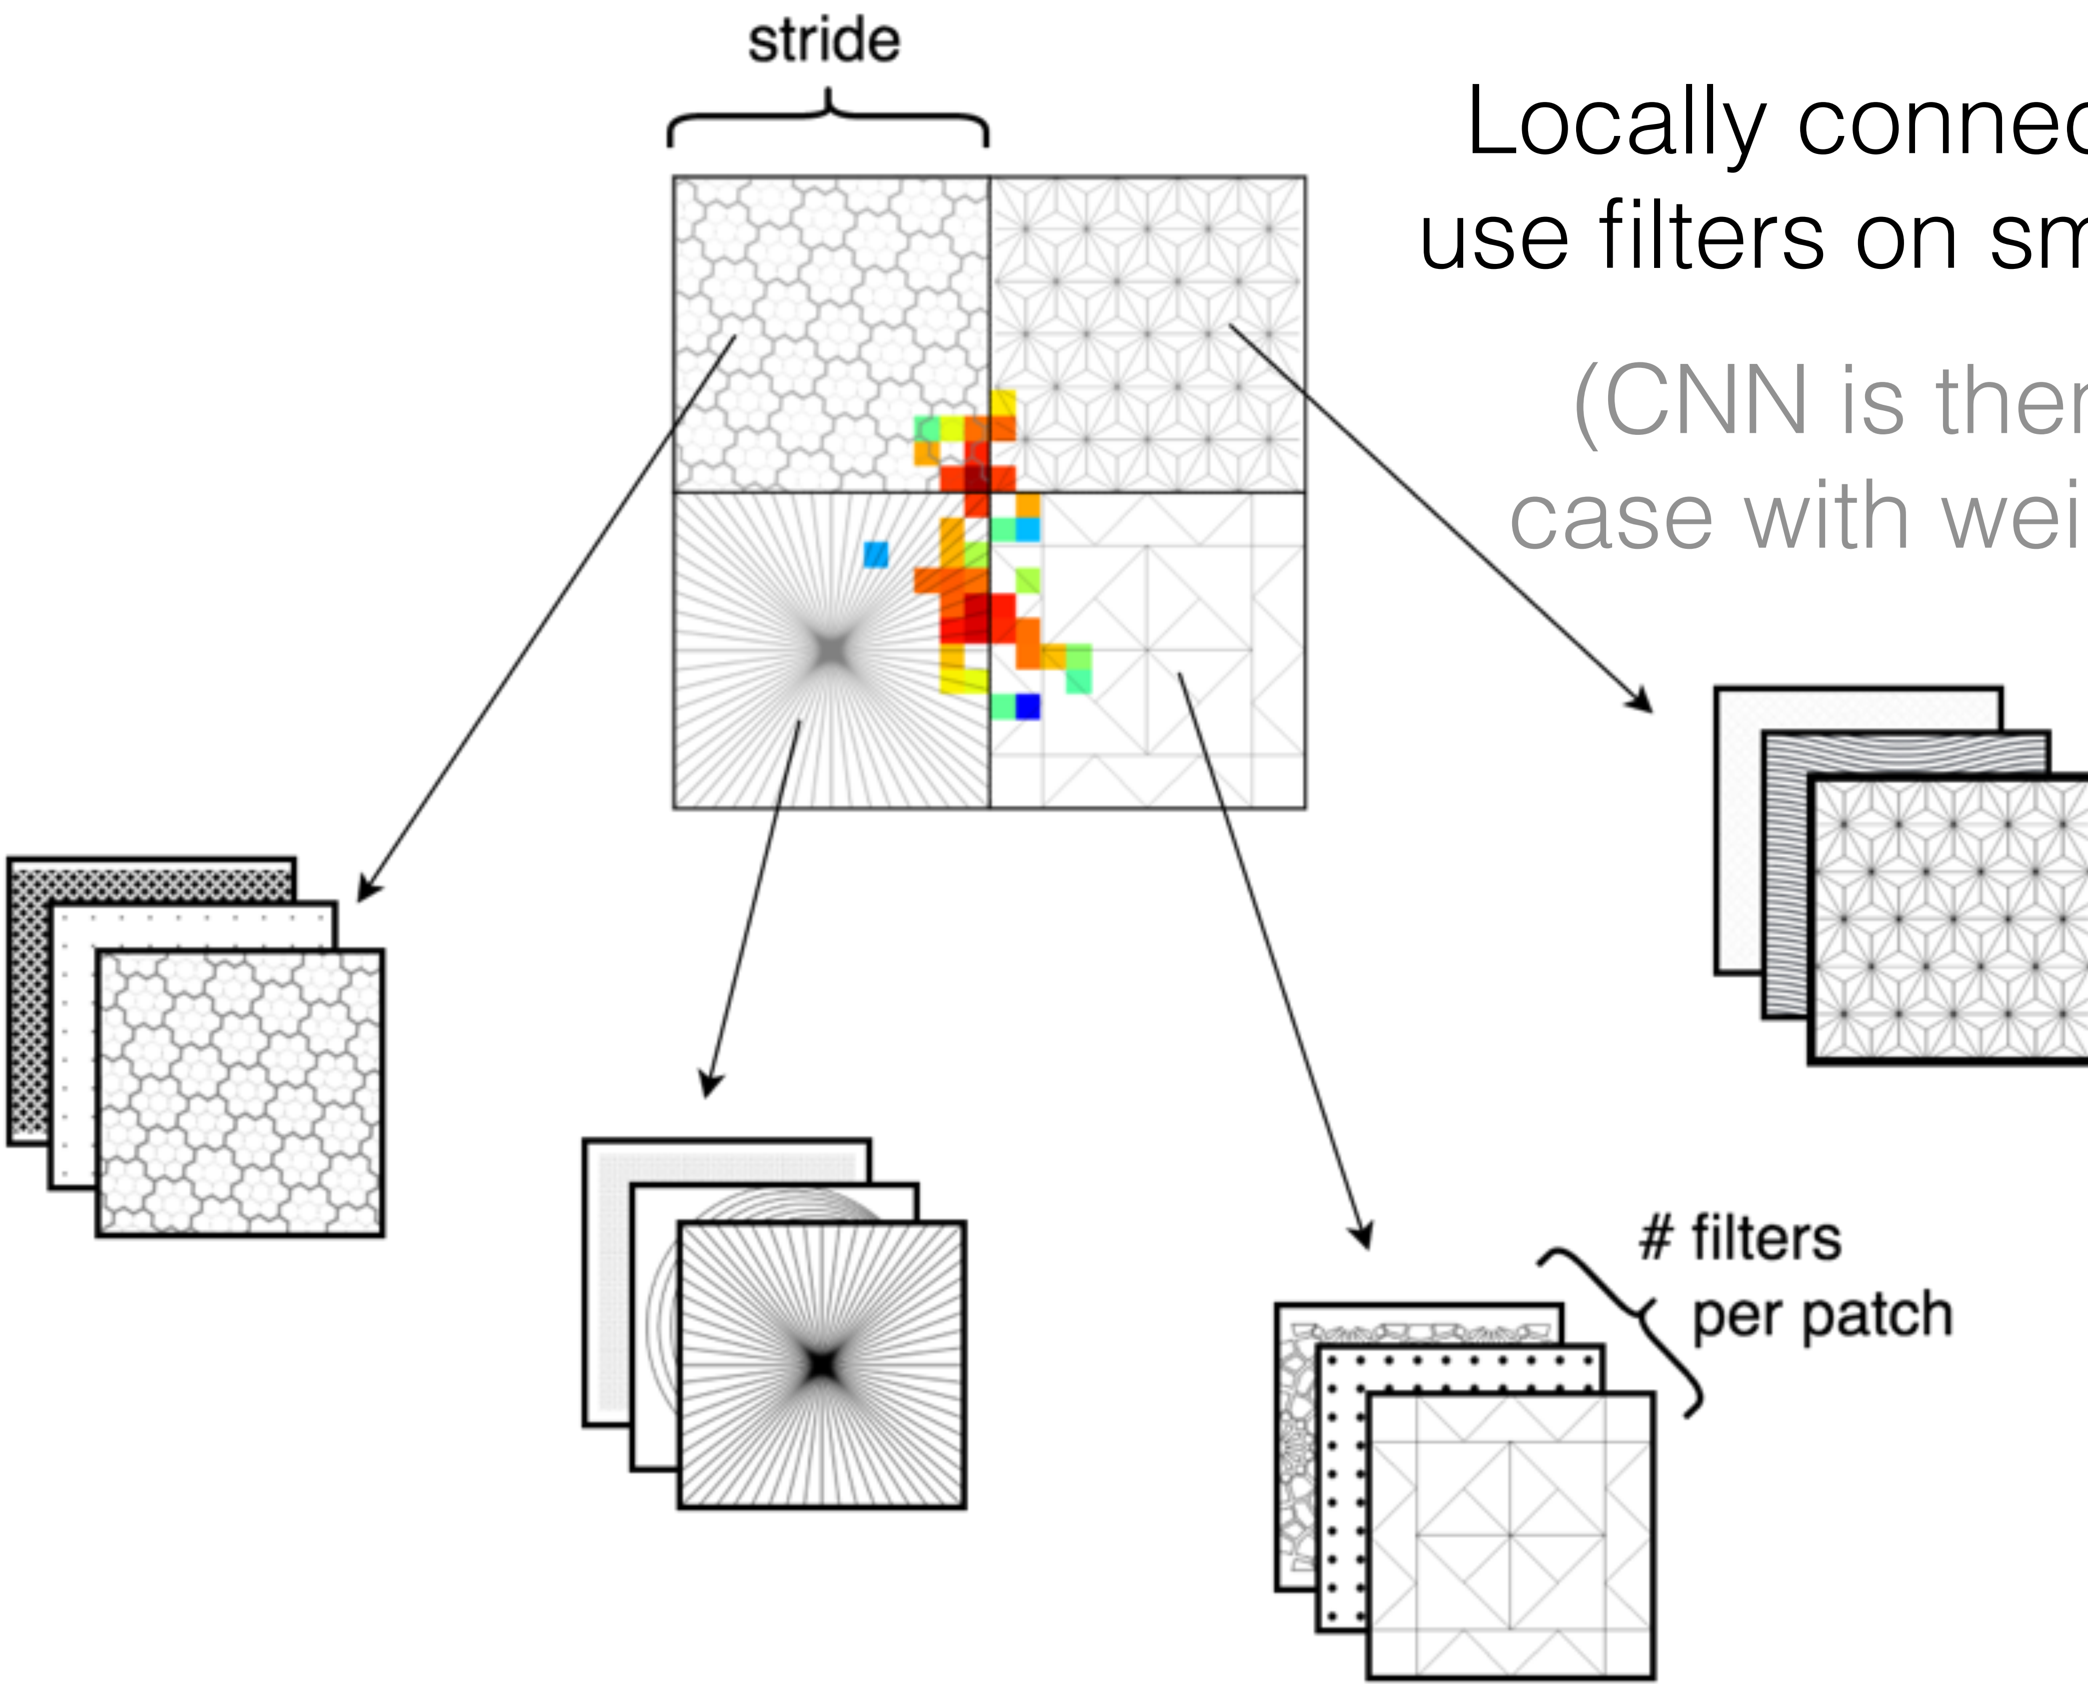

Locally connected layers use filters on small patches

(CNN is then a special case with weight sharing)

Results: average images

Geant4

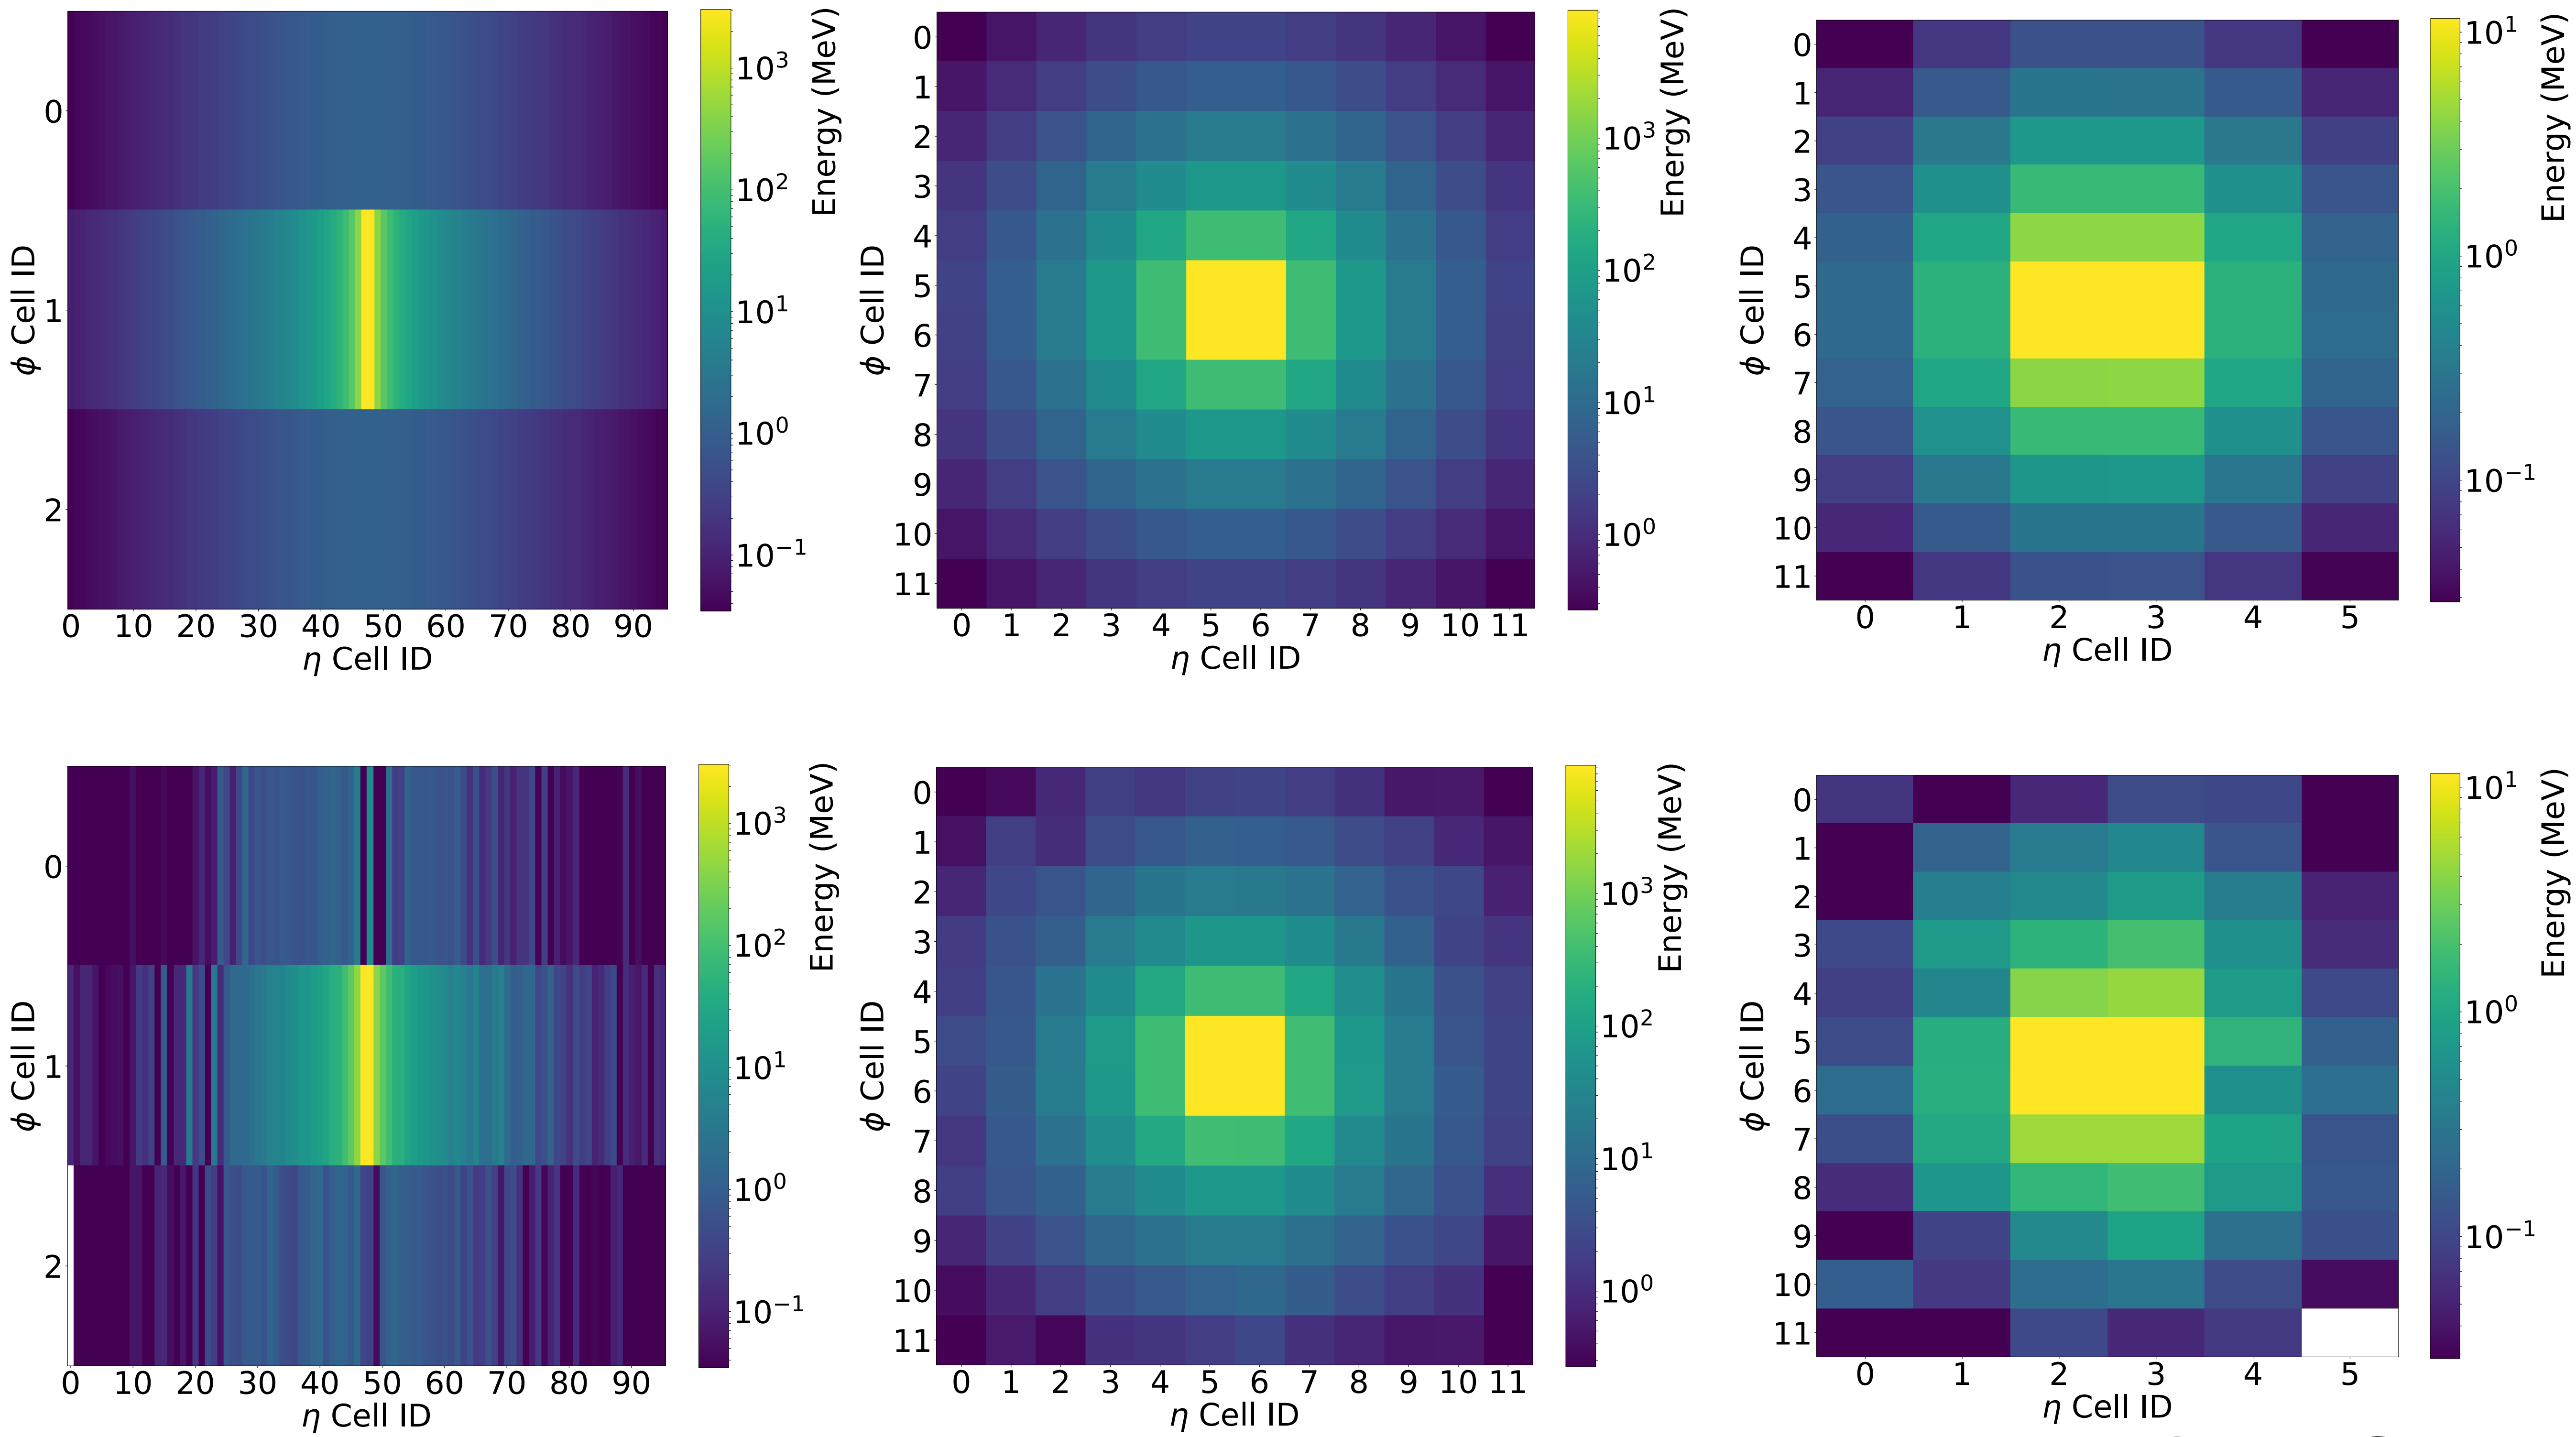

CaloGAN

Energy per layer

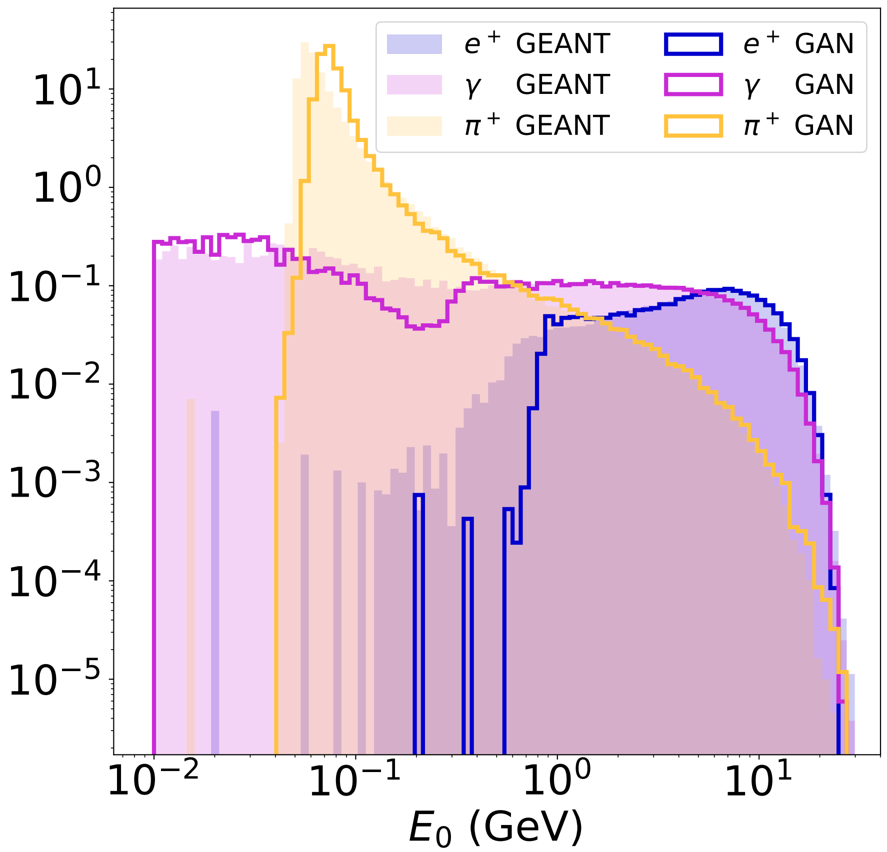

Pions deposit much less energy in the first layers; leave the calorimeter with significant energy

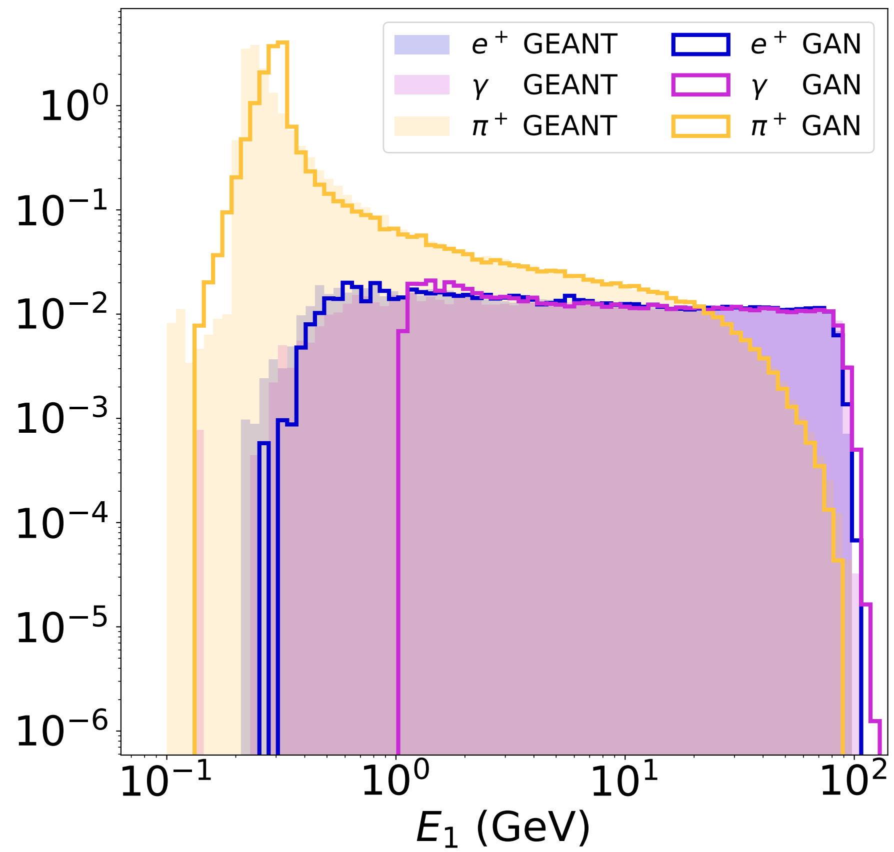

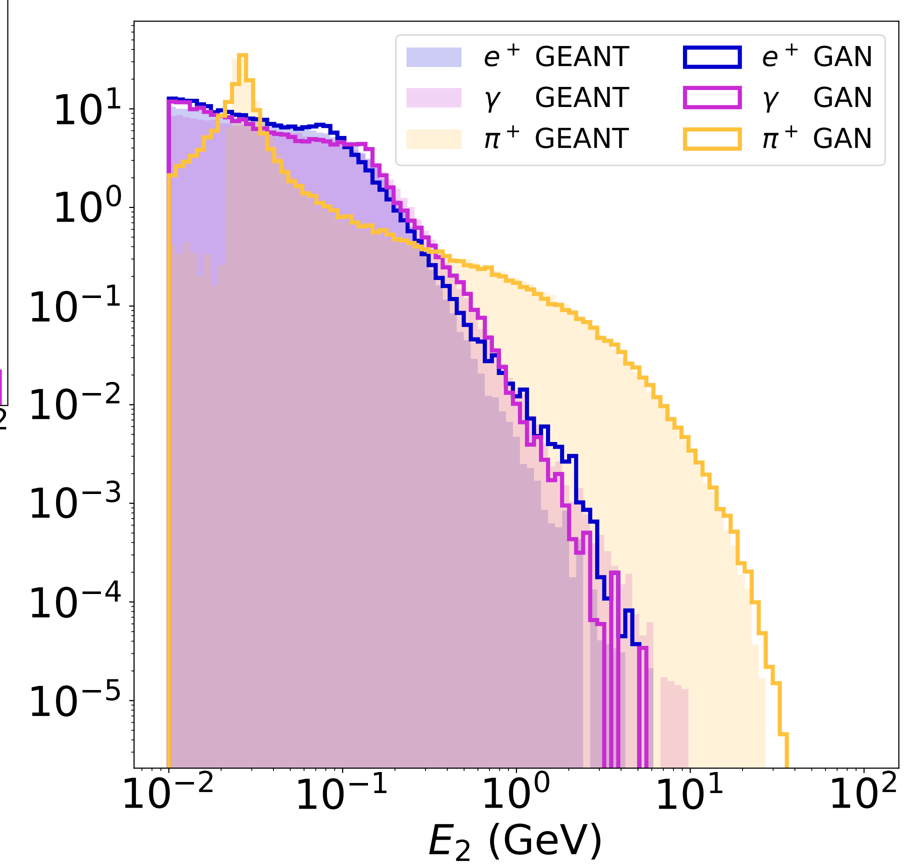

N.B. can always add these (and others) explicitly to the training

Warning: challenge with GANs

Unlike for classifiers, it is not easy to figure out which GAN is a good GAN - trying to learn a O(1000) generative model and not a single likelihood ratio!

…this is a place where science applications can make a big impact on ML.

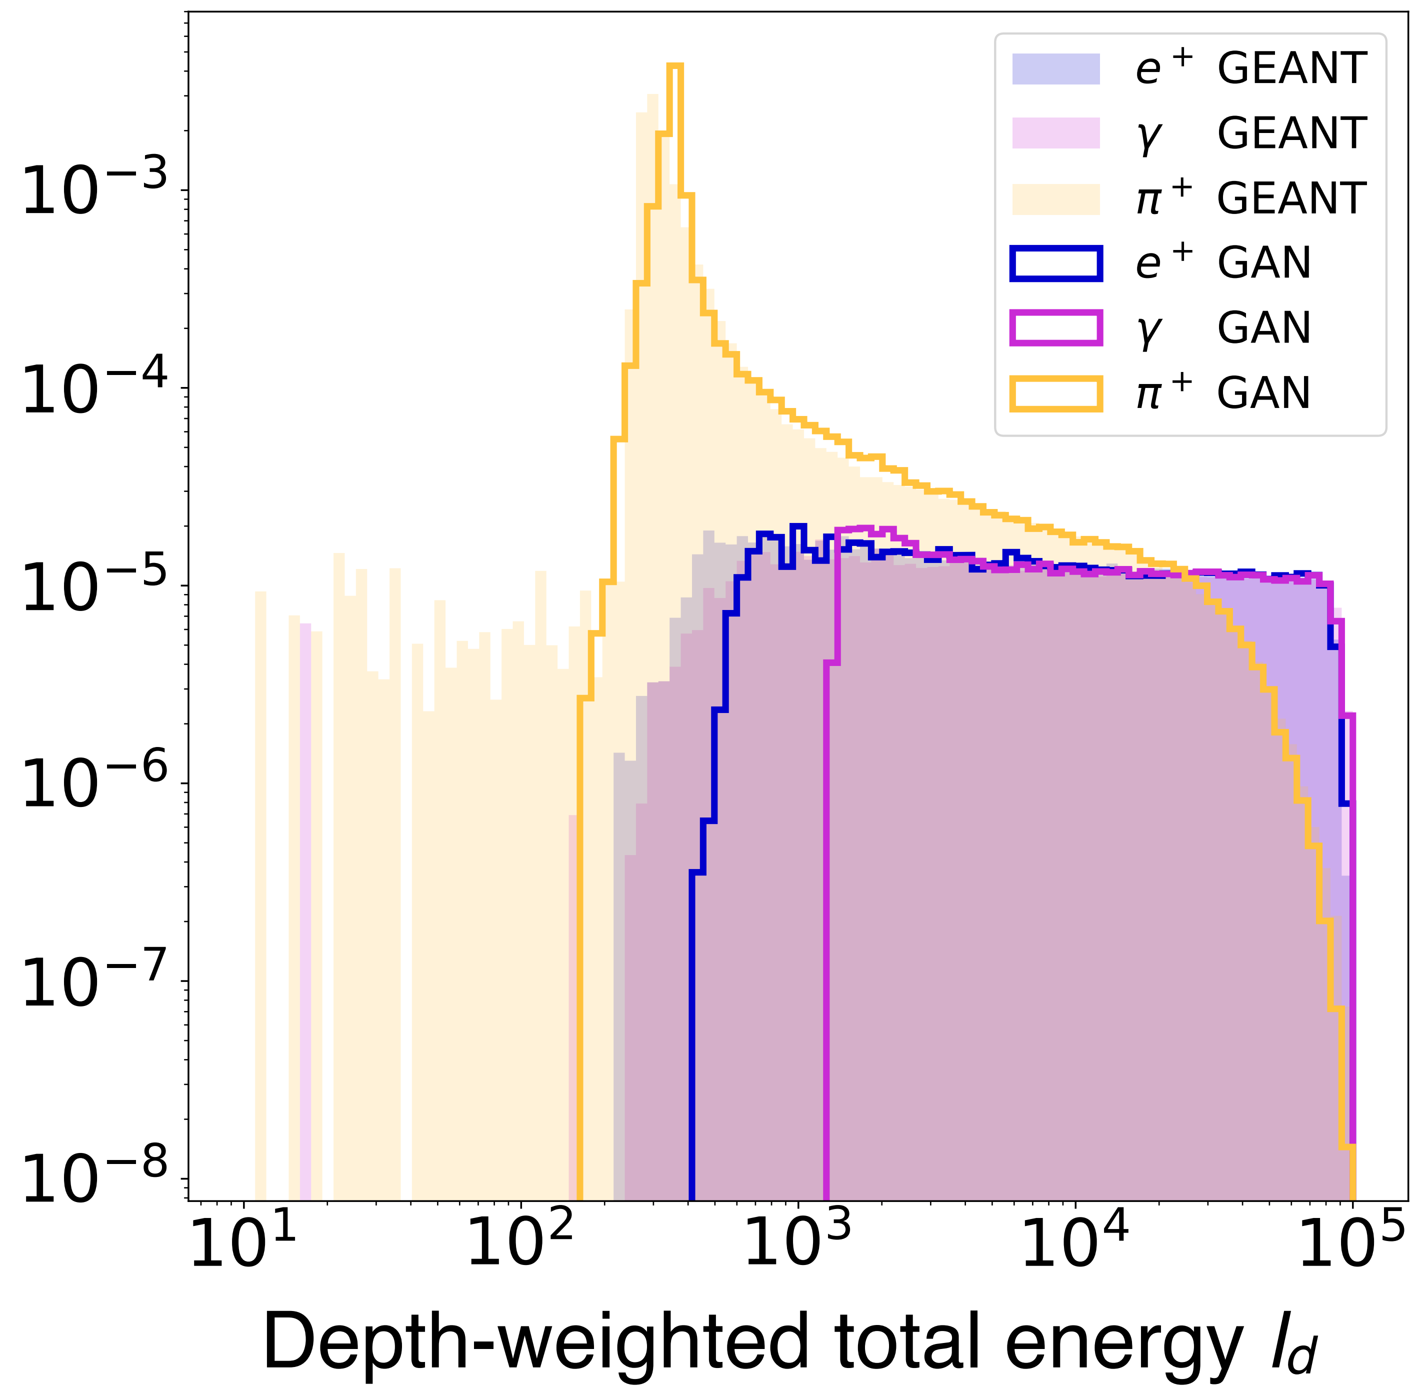

“Overtraining”

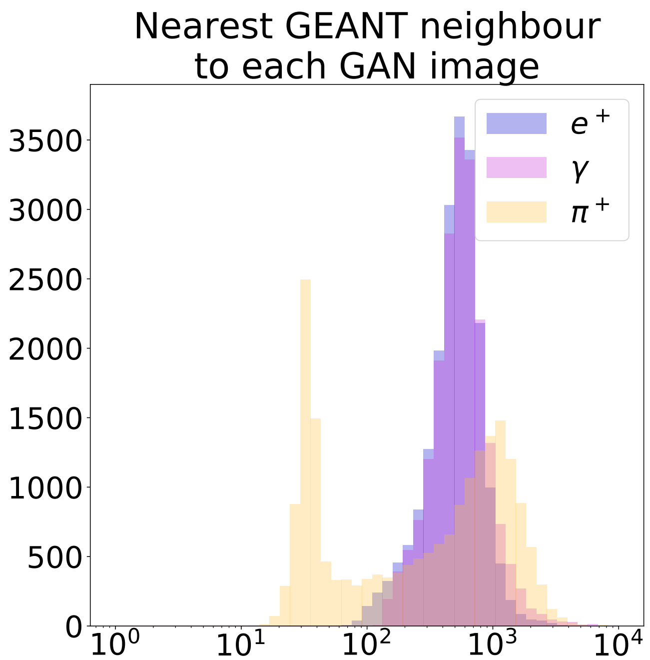

~not
memorizing

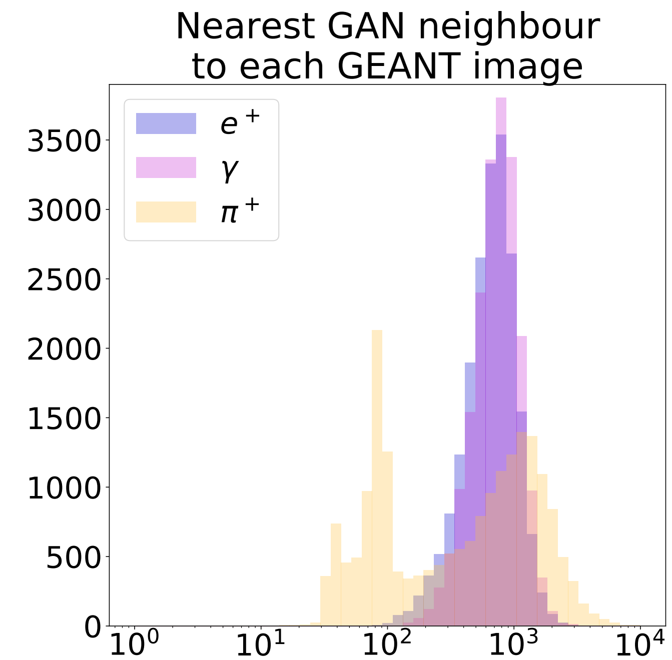

A key challenge in training GANs is the diversity of generated images. This does not seem to be a (big) problem for CaloGAN.

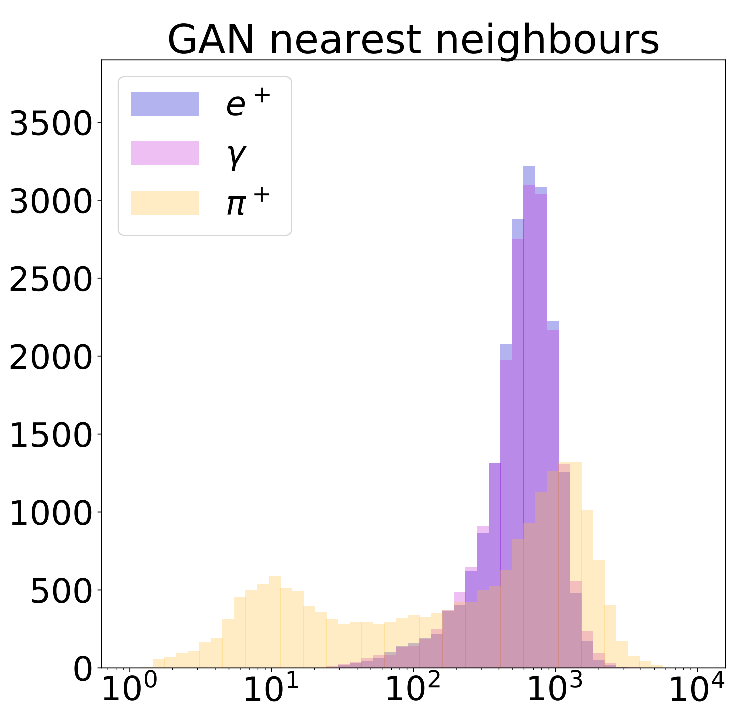

~no mode
collapse

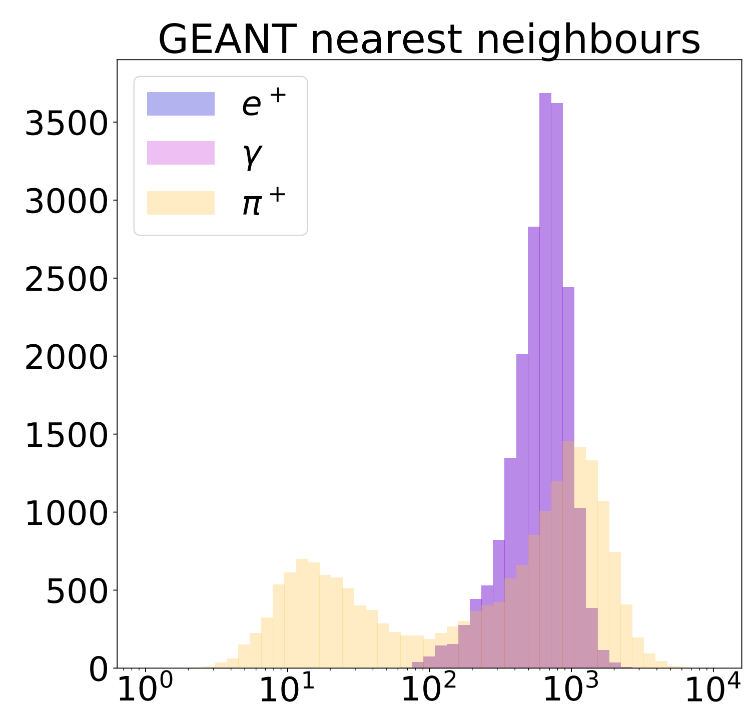

Extrapolating

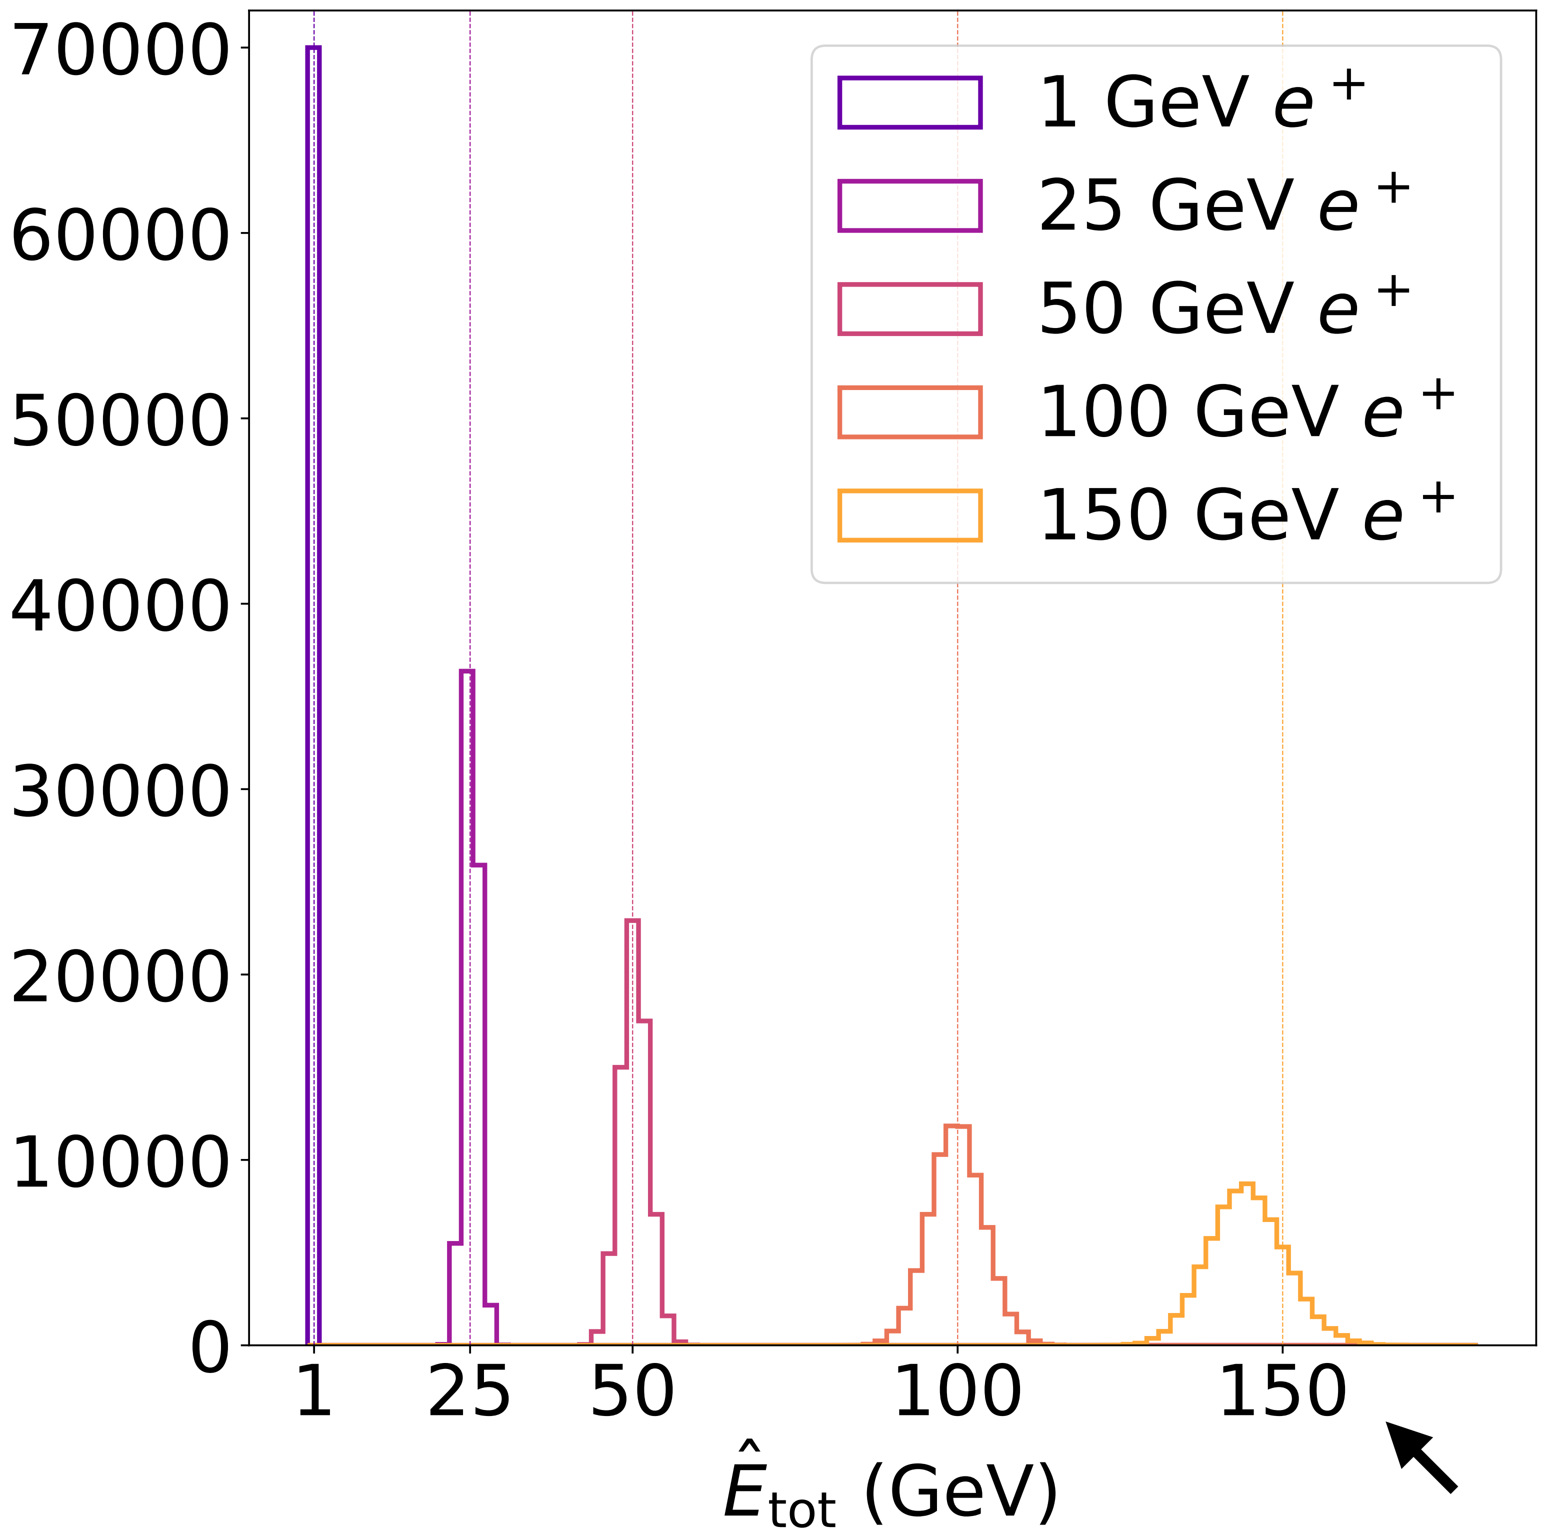

GANs are not designed to extrapolate, but in some cases, they can smoothly go on!

works here until there is no new physical principles which turn on at some energy

Beyond our training sample!

Conditioning

Fix noise, scan latent variable corresponding to energy

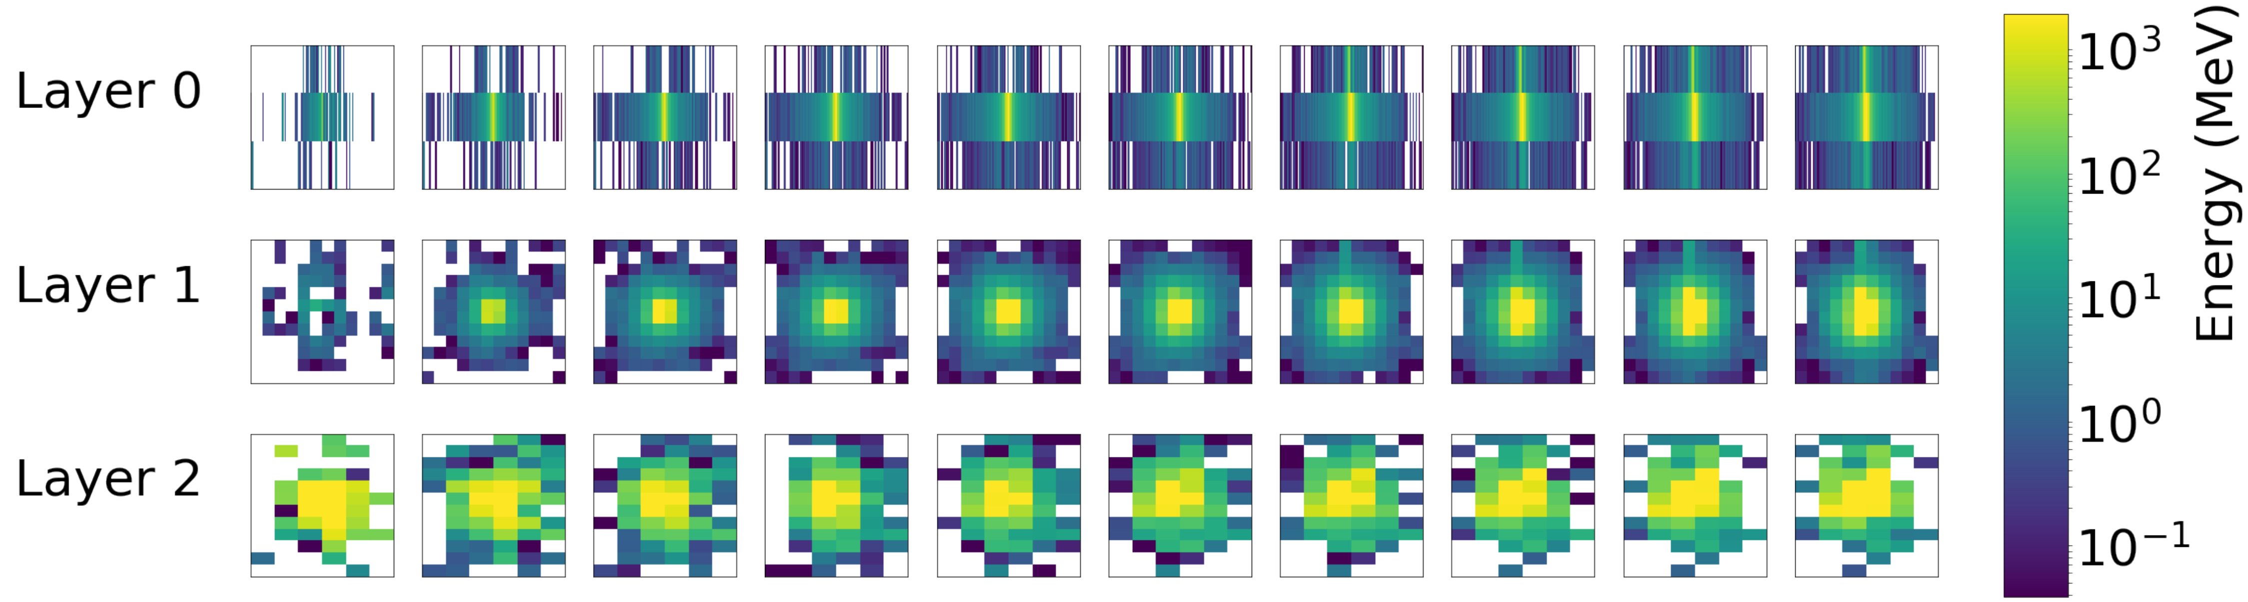

Fix noise, scan latent variable corresponding to x-position

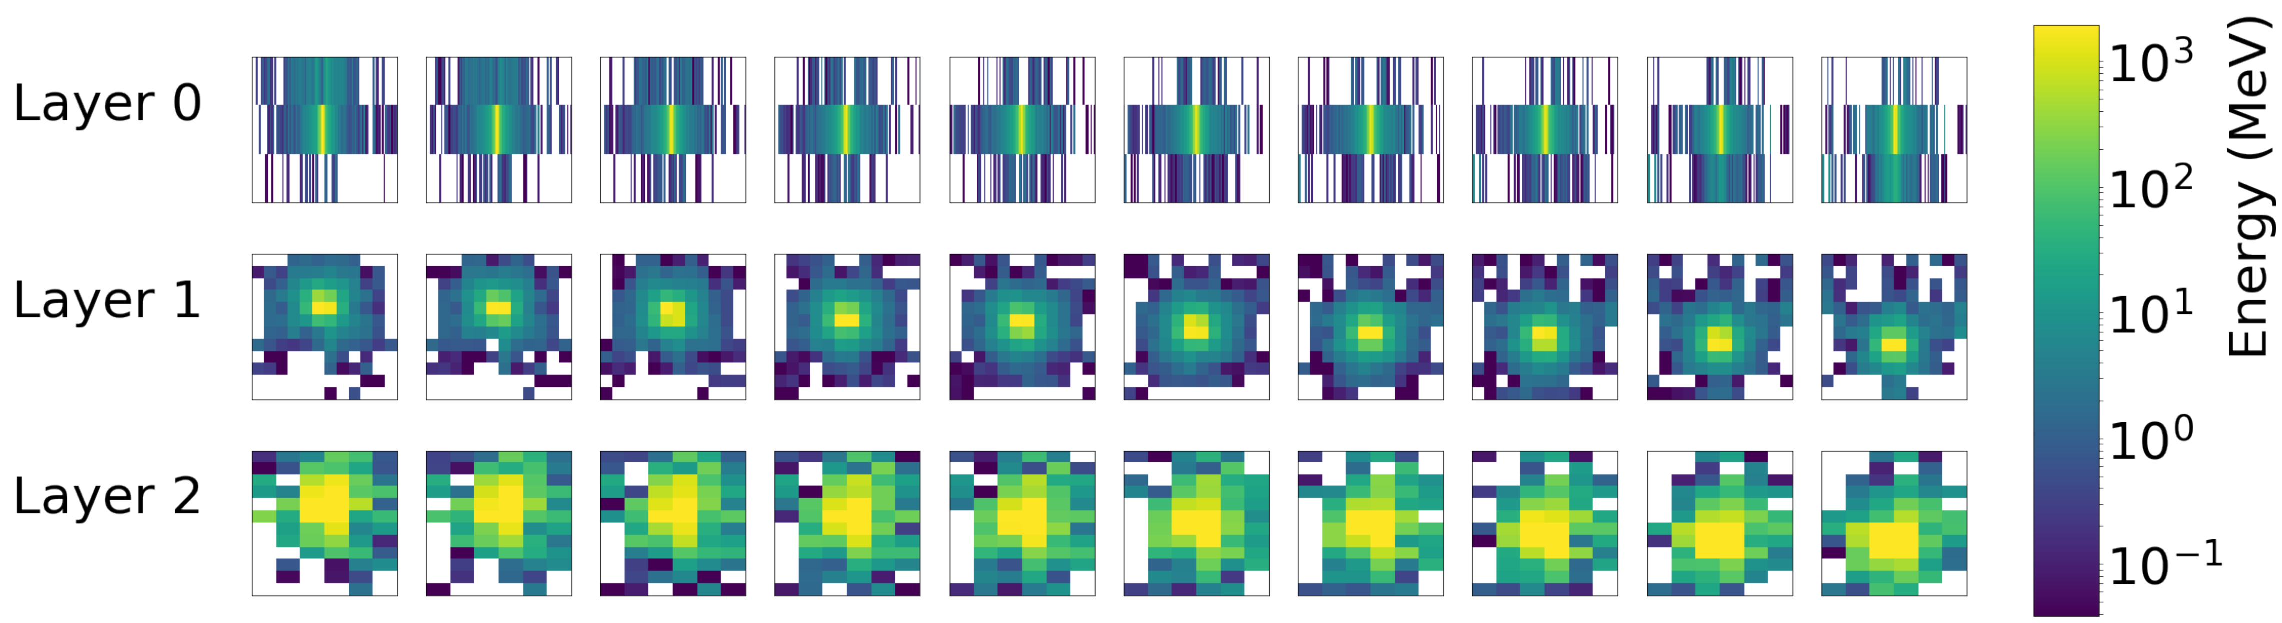

Timing

Generation Method	Hardware	Batch Size	milliseconds/shower
GEANT4	CPU	N/A	1772
CALOGAN	CPU *Intel Xeon E5-2670*	1	13.1
		10	5.11
		128	2.19
		1024	2.03
	GPU *NVIDIA K80*	1	14.5
		4	3.68
		128	0.021
		512	0.014
		1024	0.012

(clearly these numbers will change as both technologies improve - this is simply meant to be qualitative and motivating!)

Collaboration workflow

Integrating these techniques into a full detector simulation is another layer of complication, but is possible and hopefully worth the effort!

ATL-SOFT-PUB-2018-001

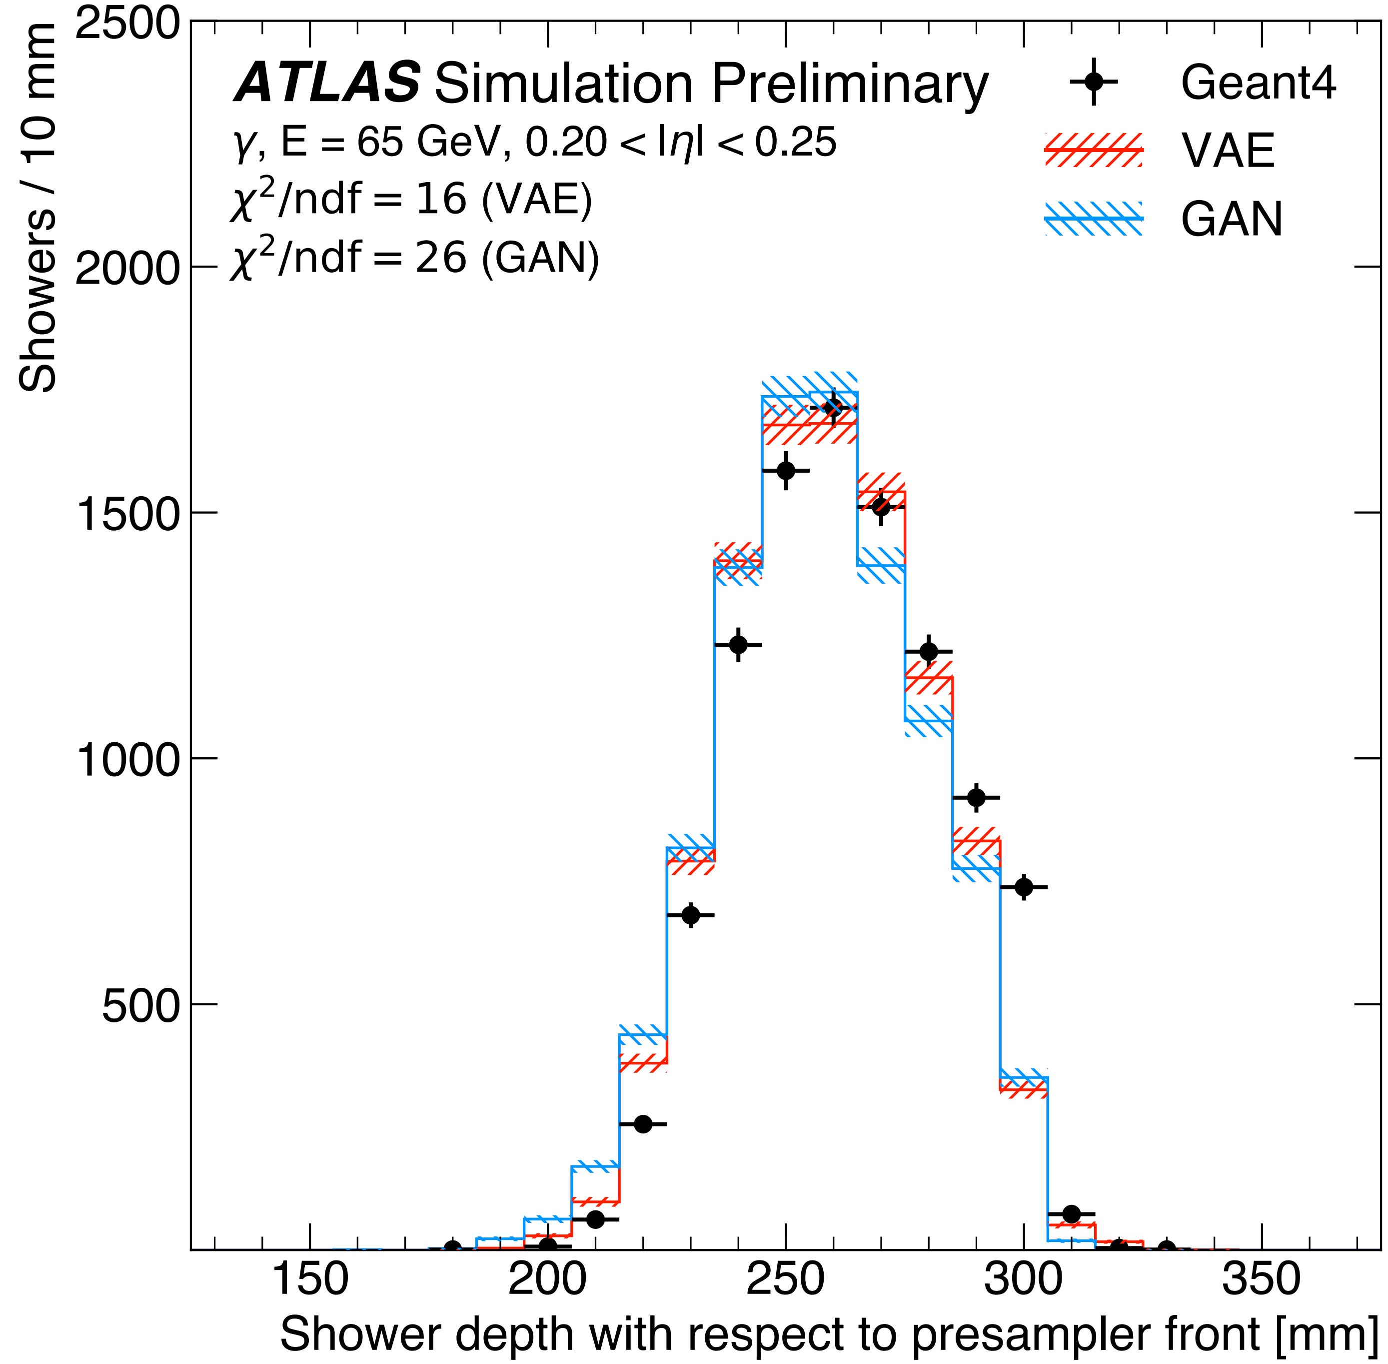

GAN studies by Aishik Ghosh and others


ATLAS ~ 3000 people
(~10% here celebrating the Higgs boson discovery)

Any paper by the ATLAS collaboration

G. Aad[48], T. Abajyan[21], B. Abbott[111], J. Abdallah[12], S. Abdel Khalek[115], A.A. Abdelalim[49], O. Abdinov[11], R. Aben[105], B. Abi[112], M. Abolins[88], O.S. AbouZeid[158], H. Abramowicz[153], H. Abreu[136], B.S. Acharya[164a,164b], L. Adamczyk[38], D.L. Adams[25], T.N. Addy[56], J. Adelman[176], S. Adomeit[98], P. Adragna[75], T. Adye[129], S. Aefsky[23], J.A. Aguilar-Saavedra[124b,a], M. Agustoni[17], M. Aharrouche[81], S.P. Ahlen[22], F. Ahles[48], A. Ahmad[148], M. Ahsan[41], G. Aielli[133a,133b], T. Akdogan[19a], T.P.A. Åkesson[79], G. Akimoto[155], A.V. Akimov[94], M.S. Alam[2], M.A. Alam[76], J. Albert[169], S. Albrand[55], M. Aleksa[30], I.N. Aleksandrov[64], F. Alessandria[89a], C. Alexa[26a], G. Alexander[153], G. Alexandre[49], T. Alexopoulos[10], M. Alhroob[164a,164c], M. Aliev[16], G. Alimonti[89a], J. Alison[120], B.M.M. Allbrooke[18], P.P. Allport[73], S.E. Allwood-Spiers[53], J. Almond[82], A. Aloisio[102a,102b], R. Alon[172], A. Alonso[79], F. Alonso[70], A. Altheimer[35], B. Alvarez Gonzalez[88], M.G. Alviggi[102a,102b], K. Amako[65], C. Amelung[23], V.V. Ammosov[128,*], S.P. Amor Dos Santos[124a], A. Amorim[124a,b], N. Amram[153], C. Anastopoulos[30], L.S. Ancu[17], N. Andari[115], T. Andeen[35], C.F. Anders[58b], G. Anders[58a], K.J. Anderson[31], A. Andreazza[89a,89b], V. Andrei[58a], M-L. Andrieux[55], X.S. Anduaga[70], S. Angelidakis[9], P. Anger[44], A. Angerami[35], F. Anghinolfi[30], A. Anisenkov[107], N. Anjos[124a], A. Annovi[47], A. Antonaki[9], M. Antonelli[47], A. Antonov[96], J. Antos[144b], F. Anulli[132a], M. Aoki[101], S. Aoun[83], L. Aperio Bella[5], R. Apolle[118,c], G. Arabidze[88], I. Aracena[143], Y. Arai[65], A.T.H. Arce[45], S. Arfaoui[148], J-F. Arguin[93], E. Arik[19a,*], M. Arik[19a], A.J. Armbruster[87], O. Arnaez[81], V. Arnal[80], C. Arnault[115], A. Artamonov[95], G. Artoni[132a,132b], D. Arutinov[21], S. Asai[155], S. Ask[28], B. Åsman[146a,146b], L. Asquith[6], K. Assamagan[25], A. Astbury[169], M. Atkinson[165], B. Aubert[5], E. Auge[115], K. Augsten[127], M. Aurousseau[145a], G. Avolio[163], R. Avramidou[10], D. Axen[168], G. Azuelos[93,d], Y. Azuma[155], M.A. Baak[30], G. Baccaglioni[89a], C. Bacci[134a,134b], A.M. Bach[15], H. Bachacou[136], K. Bachas[30], M. Backes[49], M. Backhaus[21], J. Backus Mayes[143], E. Badescu[26a], P. Bagnaia[132a,132b], S. Bahinipati[3], Y. Bai[33a], D.C. Bailey[158], T. Bain[158], J.T. Baines[129], O.K. Baker[176], M.D. Baker[25], S. Baker[77], P. Balek[126], E. Banas[39], P. Banerjee[93], Sw. Banerjee[173], D. Banfi[30], A. Bangert[150], V. Bansal[169], H.S. Bansil[18], L. Barak[172], S.P. Baranov[94], A. Barbaro Galtieri[15], T. Barber[48], E.L. Barberio[86], D. Barberis[50a,50b], M. Barbero[21], D.Y. Bardin[64], T. Barillari[99], M. Barisonzi[175], T. Barklow[143], N. Barlow[28], B.M. Barnett[129], R.M. Barnett[15], A. Baroncelli[134a], G. Barone[49], A.J. Barr[118], F. Barreiro[80], J. Barreiro Guimarães da Costa[57], P. Barrillon[115], R. Bartoldus[143], A.E. Barton[71], V. Bartsch[149], A. Basye[165], R.L. Bates[53], L. Batkova[144a], J.R. Batley[28], A. Battaglia[17], M. Battistin[30], F. Bauer[136], H.S. Bawa[143,e], S. Beale[98], T. Beau[78], P.H. Beauchemin[161], R. Beccherle[50a], P. Bechtle[21], H.P. Beck[17], A.K. Becker[175], S. Becker[98], M. Beckingham[138], K.H. Becks[175], A.J. Beddall[19c], A. Beddall[19c], S. Bedikian[176], V.A. Bednyakov[64], C.P. Bee[83], L.J. Beemster[105], M. Begel[25]

T. Mkrtchyan[180], M. Mlynarikova[131], T. Moa[148a,148b], K. Mochizuki[97], P. Mogg[51], S. Mohapatra[38], S. Molander[148a,148b], R. Moles-Valls[23], R. Monden[71], M.C. Mondragon[93], K. Mönig[45], J. Monk[39], E. Monnier[88], A. Montalbano[150], J. Montejo Berlingen[32], F. Monticelli[74], S. Monzani[94a,94b], R.W. Moore[3], N. Morange[119], D. Moreno[21], M. Moreno Llácer[32], P. Morettini[53a], S. Morgenstern[32], D. Mori[144], T. Mori[157], M. Morii[59], M. Morinaga[157], V. Morisbak[121], A.K. Morley[152], G. Mornacchi[32], J.D. Morris[79], L. Morvaj[150], P. Moschovakos[10], M. Mosidze[54b], H.J. Moss[141], J. Moss[145,ak], K. Motohashi[159], R. Mount[145], E. Mountricha[27], E.J.W. Moyse[89], S. Muanza[88], R.D. Mudd[19], F. Mueller[103], J. Mueller[127], R.S.P. Mueller[102], D. Muenstermann[75], P. Mullen[56], G.A. Mullier[18], F.J. Munoz Sanchez[87], W.J. Murray[173,133], H. Musheghyan[181], M. Muškinja[78], A.G. Myagkov[132,al], M. Myska[130], B.P. Nachman[16], O. Nackenhorst[52], K. Nagai[122], R. Nagai[69,ad], K. Nagano[69], Y. Nagasaka[61], K. Nagata[164], M. Nagel[51], E. Nagy[88], A.M. Nairz[32], Y. Nakahama[105], K. Nakamura[69], T. Nakamura[157], I. Nakano[114], R.F. Naranjo Garcia[45], R. Narayan[11], D.I. Narrias Villar[60a], I. Naryshkin[125], T. Naumann[45], G. Navarro[21], R. Nayyar[7], H.A. Neal[92], P.Yu. Nechaeva[98], T.J. Neep[138], A. Negri[123a,123b], M. Negrini[22a], S. Nektarijevic[108], C. Nellist[119], A. Nelson[166], M.E. Nelson[122], S. Nemecek[129], P. Nemethy[112], M. Nessi[32,am], M.S. Neubauer[169], M. Neumann[178], P.R. Newman[19], T.Y. Ng[62c], T. Nguyen Manh[97], R.B. Nickerson[122], R. Nicolaidou[138], J. Nielsen[139], V. Nikolaenko[132,al], I. Nikolic-Audit[83], K. Nikolopoulos[19], J.K. Nilsen[121], P. Nilsson[27], Y. Ninomiya[157], A. Nisati[134a], N. Nishu[35c], R. Nisius[103], I. Nitsche[46], T. Nobe[157], Y. Noguchi[71], M. Nomachi[120], I. Nomidis[31], M.A. Nomura[27], T. Nooney[79], M. Nordberg[32], N. Norjoharuddeen[122], O. Novgorodova[47], S. Nowak[103], M. Nozaki[69], L. Nozka[117], K. Ntekas[166], E. Nurse[81], F. Nuti[91], K. O'connor[25], D.C. O'Neil[144], A.A. O'Rourke[45], V. O'Shea[56], F.G. Oakham[31,d], H. Oberlack[103], T. Obermann[23], J. Ocariz[83], A. Ochi[70], I. Ochoa[38], J.P. Ochoa-Ricoux[34a], S. Oda[73], S. Odaka[69], H. Ogren[64], A. Oh[87], S.H. Oh[48], C.C. Ohm[16], H. Ohman[168], H. Oide[53a,53b], H. Okawa[164], Y. Okumura[157], T. Okuyama[69], A. Olariu[28b], L.F. Oleiro Seabra[128a], S.A. Olivares Pino[49], D. Oliveira Damazio[27], A. Olszewski[42], J. Olszowska[42], A. Onofre[128a,128e], K. Onogi[105], P.U.E. Onyisi[11,z], M.J. Oreglia[33], Y. Oren[155], D. Orestano[136a,136b],

Accelerating simulations: the future

GANs (and friends) are a promising solution to the collider HEP computing challenge.

(see the growing list of variations that cite the CaloGAN paper)

The key challenge now is achieving precision.

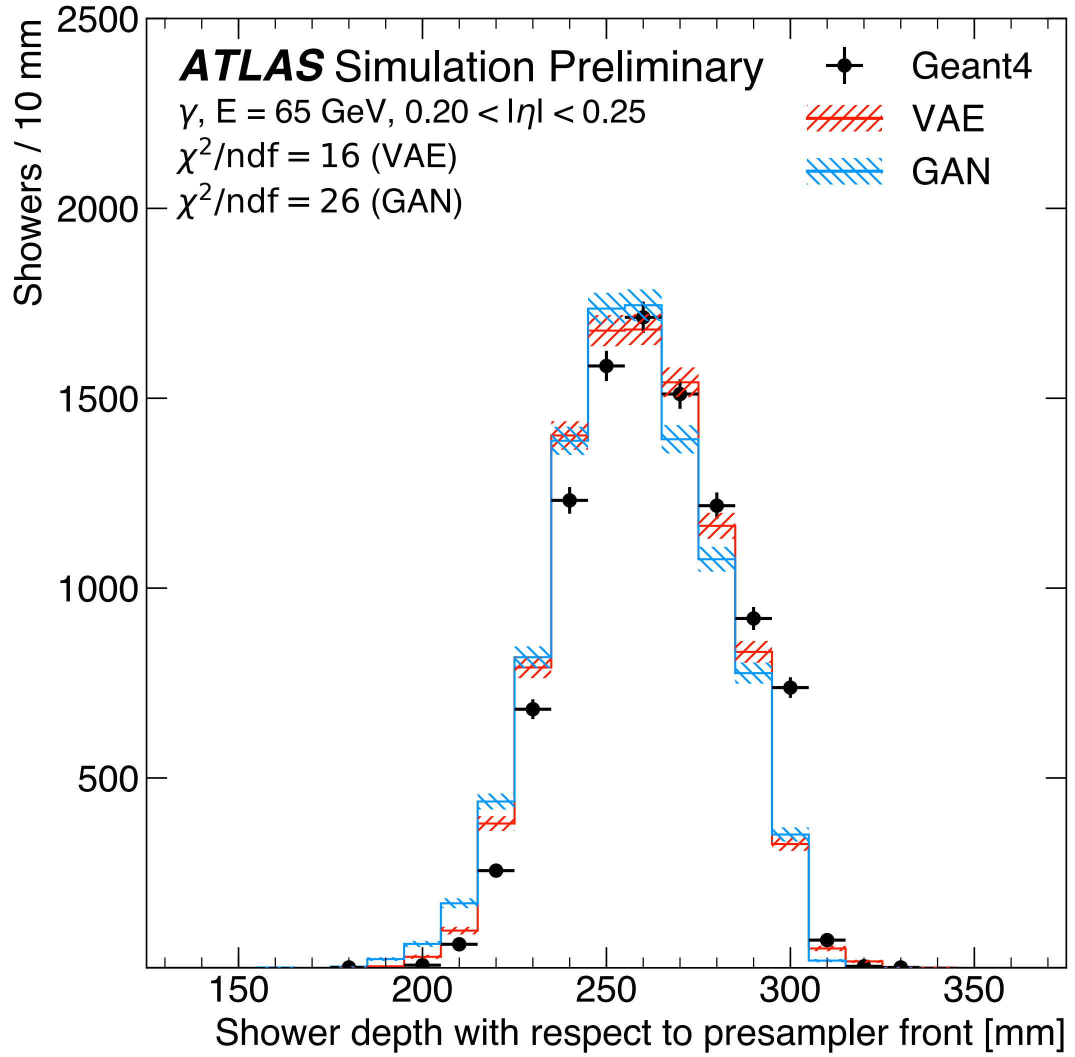

It is difficult hard to know which GAN is a good GAN …
for high-precision tasks, will need to rely on other techniques.

Part II: *High-dimensional reweighting*

New simulations

morph one simulation into another

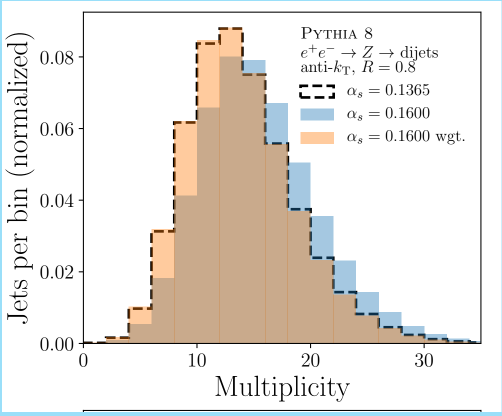

Continuous variations

learn the dependence on parameters

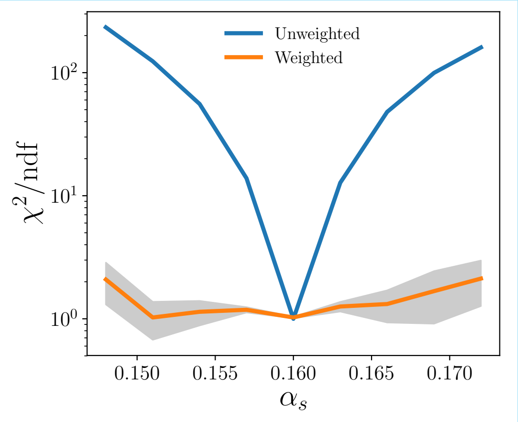

Parameter estimation

use classification loss to fit parameters

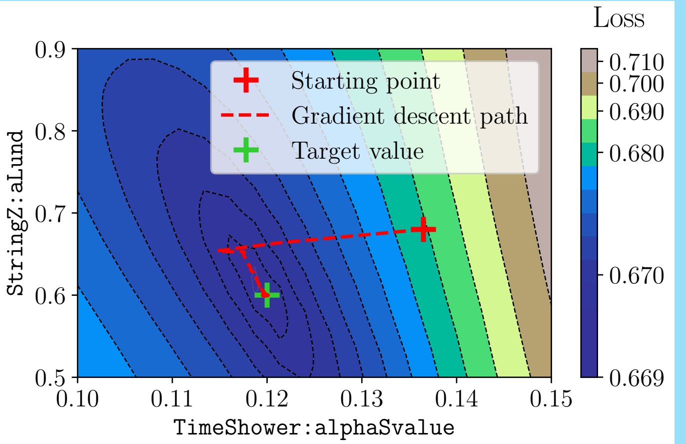

A. Andreassen and BPN, 1909.03081

Part II: *High-dimensional reweighting*

Facts: detector-level simulation is **expensive**. (e.g. **Geant4**)
pre-detector particle simulations are **cheap**.

Imagine we have one high-statistics expensive simulation.

Suppose there is another simulation of the pre-detector dynamics. Can we use the pre-detector parts to achieve a detector version of the new simulation?

Answer: Yes! Full phase space reweighing with neural networks.

Likelihood free reweighting

Let x be a simulated event. It will be composed of many hundreds of particles.

Suppose that **p(x)** and **q(x)** are the densities for the two simulations.

We can reweight the first simulation into the second by assigning per-event weights of **q(x)/p(x)**.

…what if we don't (and can't easily) know **q** and **p**?

Likelihood free reweighting

Solution: train a neural network to distinguish the two simulations. Call this **f**.

It is not hard to show that if **f** is optimal and you train with cross-entropy, then

$$\frac{f(x)}{1-f(x)} \propto \frac{q(x)}{p(x)}$$

(for weighting, we don't care about overall constants - in this case, it is the class imbalance during training)

Likelihood free reweighting

Solution: train a neural ne ... to
distinguish the two s ... this **f**.

It i ... that if **f** is optimal and
... ain with cross-entropy, then

This is great because **classification is easy** while **generation is hard**.

$$\frac{f(x)}{1-f(x)} \propto \frac{q(x)}{p(x)}$$

(for weighting, we don't care about overall constants - in this case, it is the class imbalance during training)

Example: electron-positron collisions

Learn a classifier on the full observable phase space (momenta + particle flavor) and then check with some standard observables.

Our events have a variable number of particles & due to quantum mechanics, are permutation invariant. Thus, we use a deep-sets variant called **particle flow networks**.

PFNs: Komiske, Metodiev, Thaler, JHEP 01 (2019) 121

Deep sets: Zaheer et al., NIPS 2017

Example: electron-positron collisions

Learn a classifier on the full observable phase space (momenta + particle flavor) and then check with some standard 1D observables.

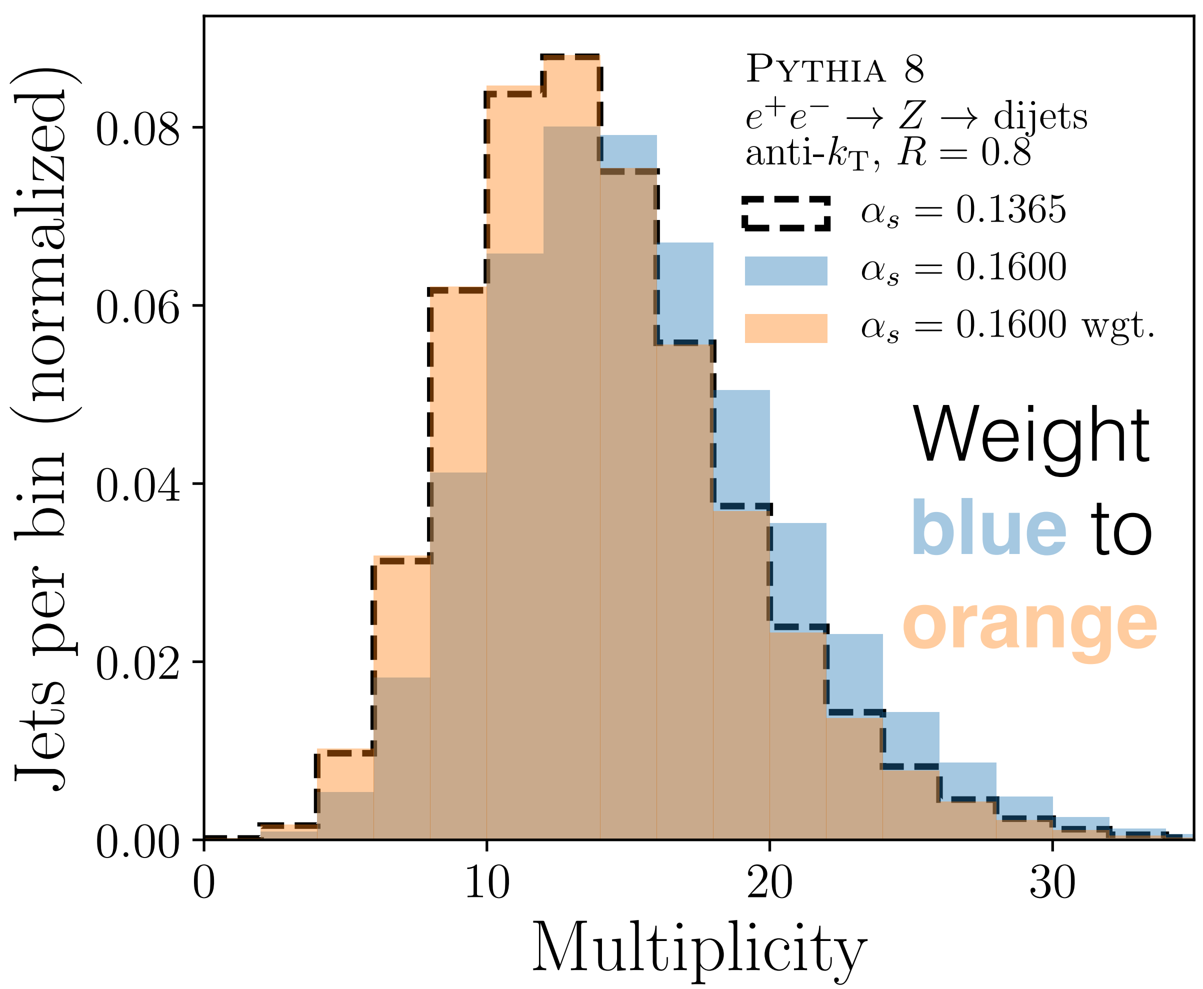

(# of particles)

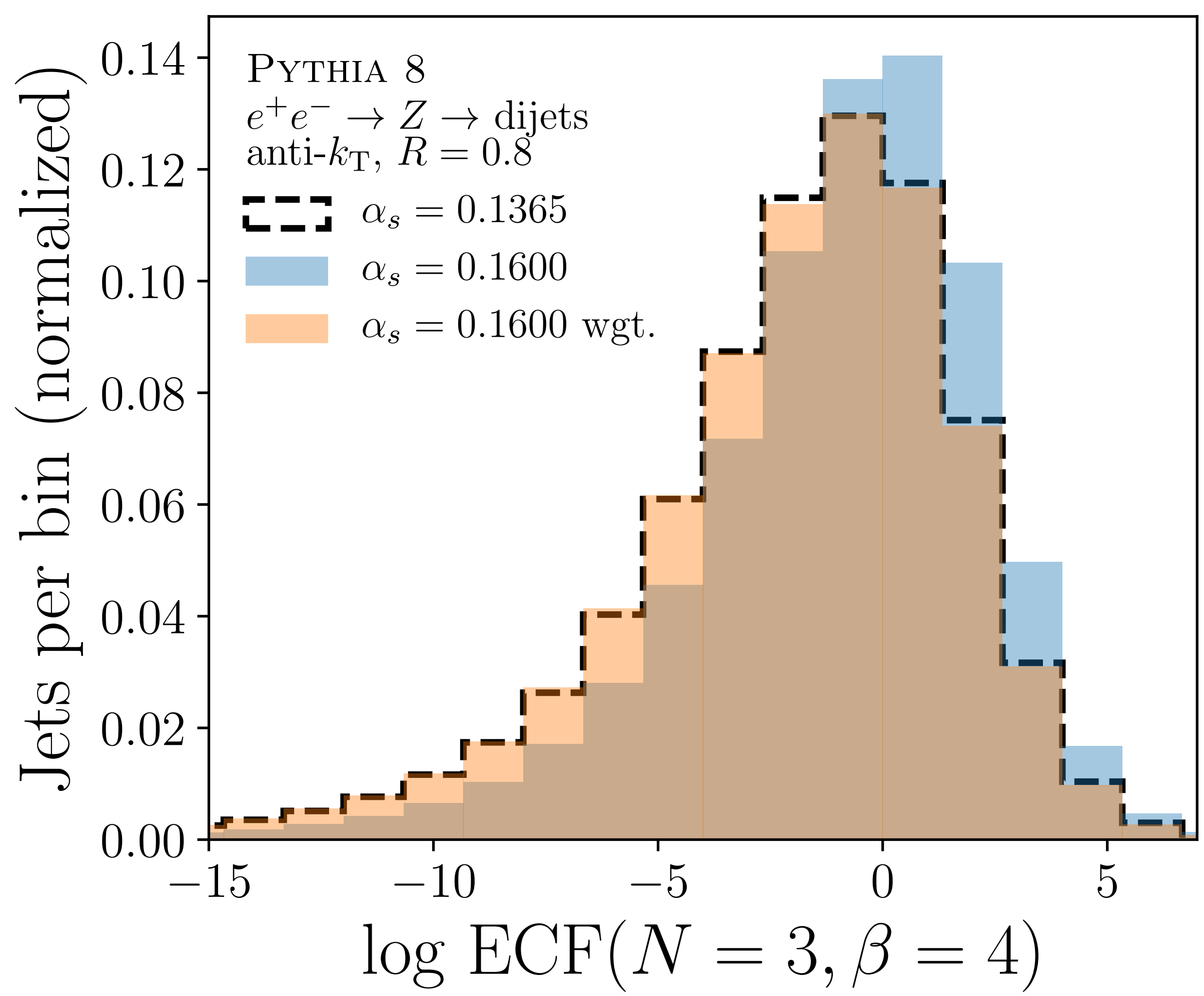

(3-particle correlation function)

Achieving precision

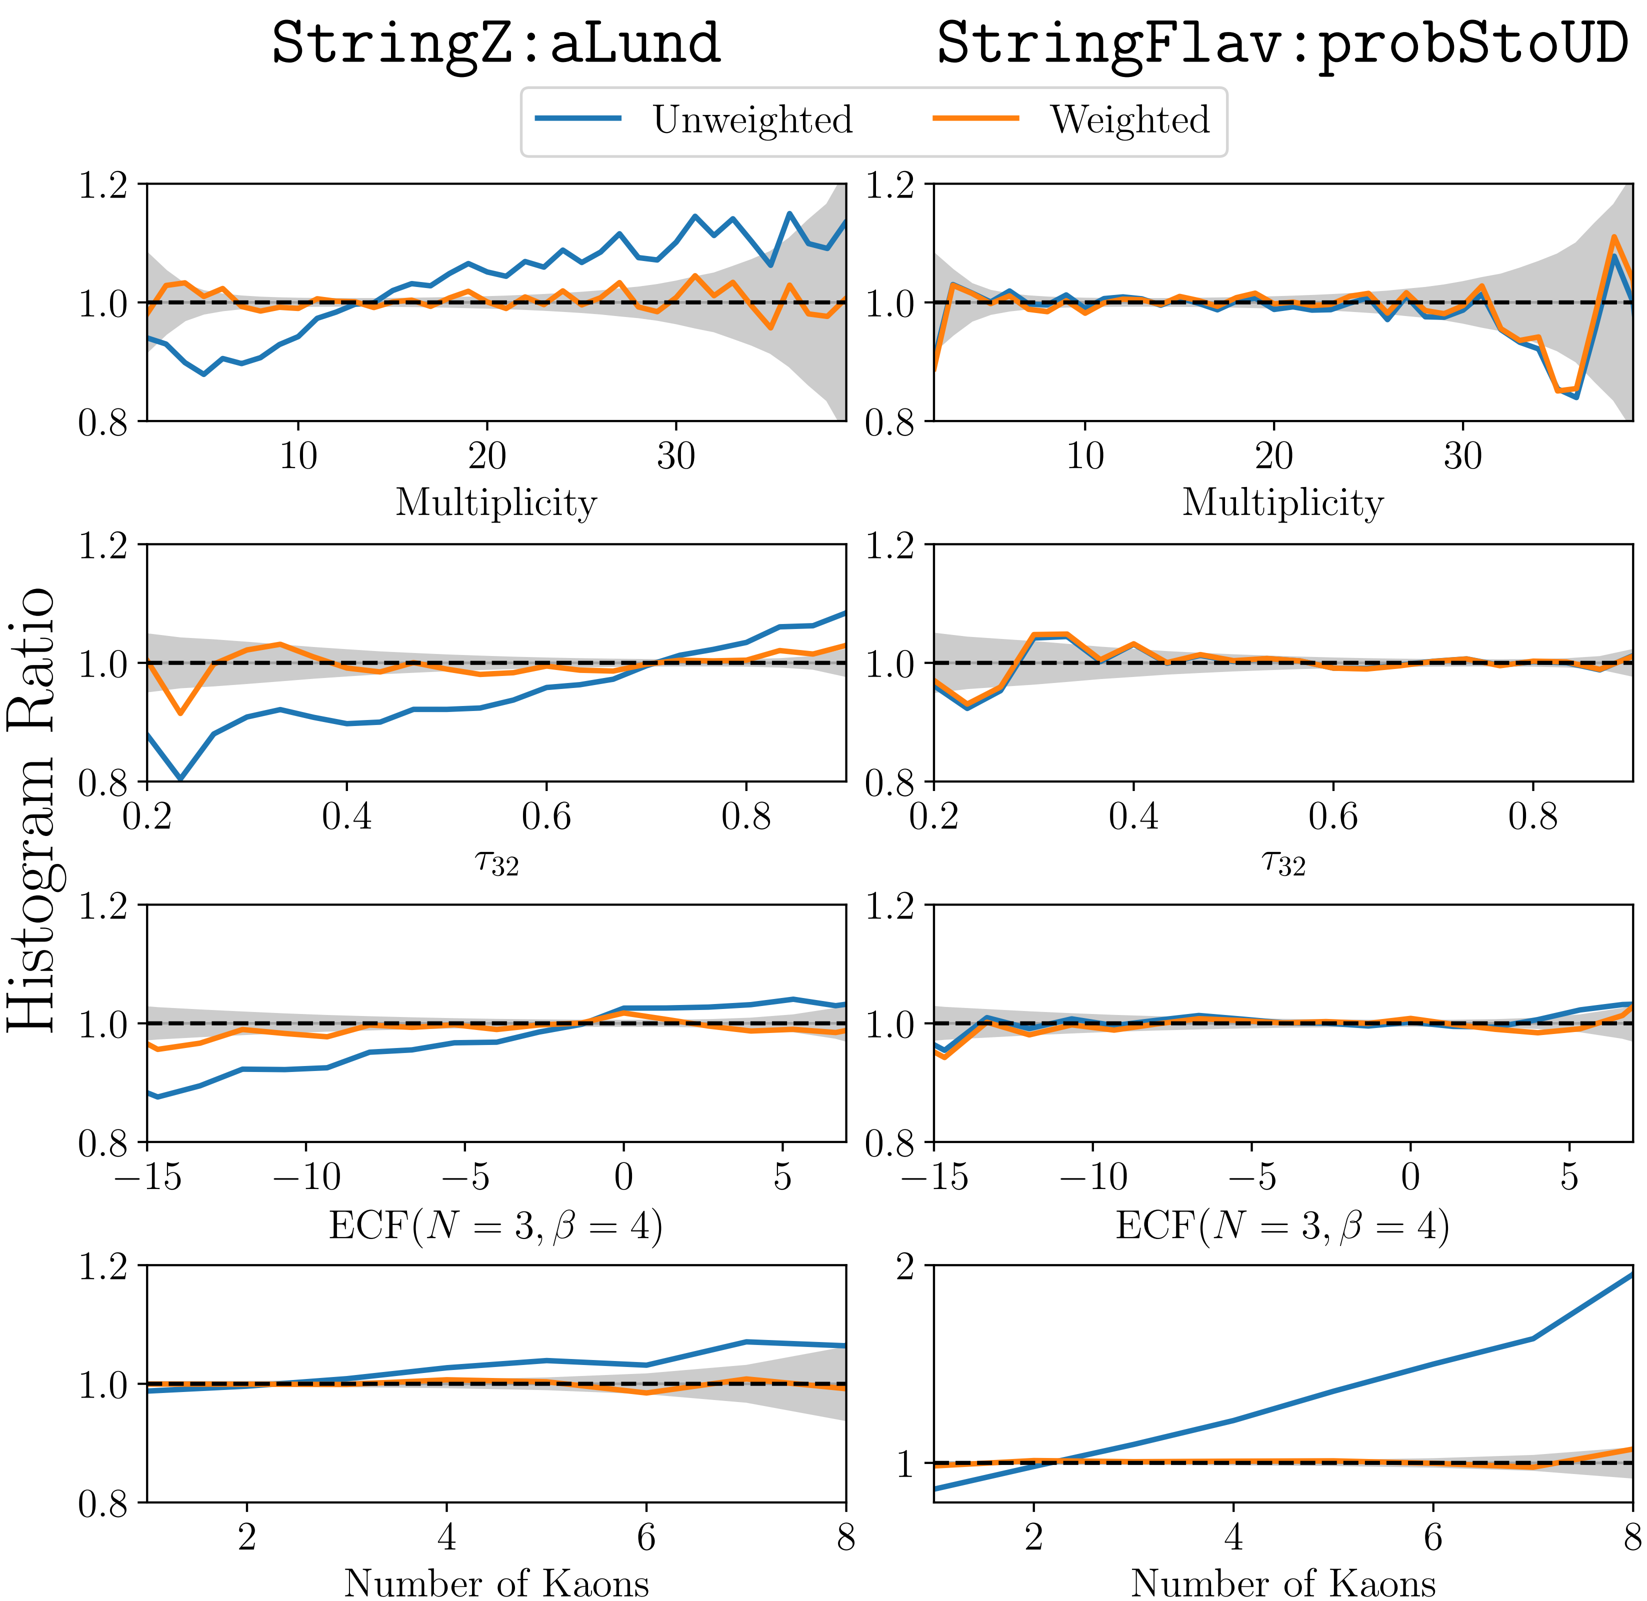

Works also when the differences between the two simulations are small (left) or localized (right).

These are histogram ratios for a series of one-dimensional observables

Parameterized reweighting

What if we have a new simulation parameterized by some parameters θ?

Easy - simply learn a parameterized classifier* !

…simply add the parameter as a feature to the network during training and let it learn to interpolate.

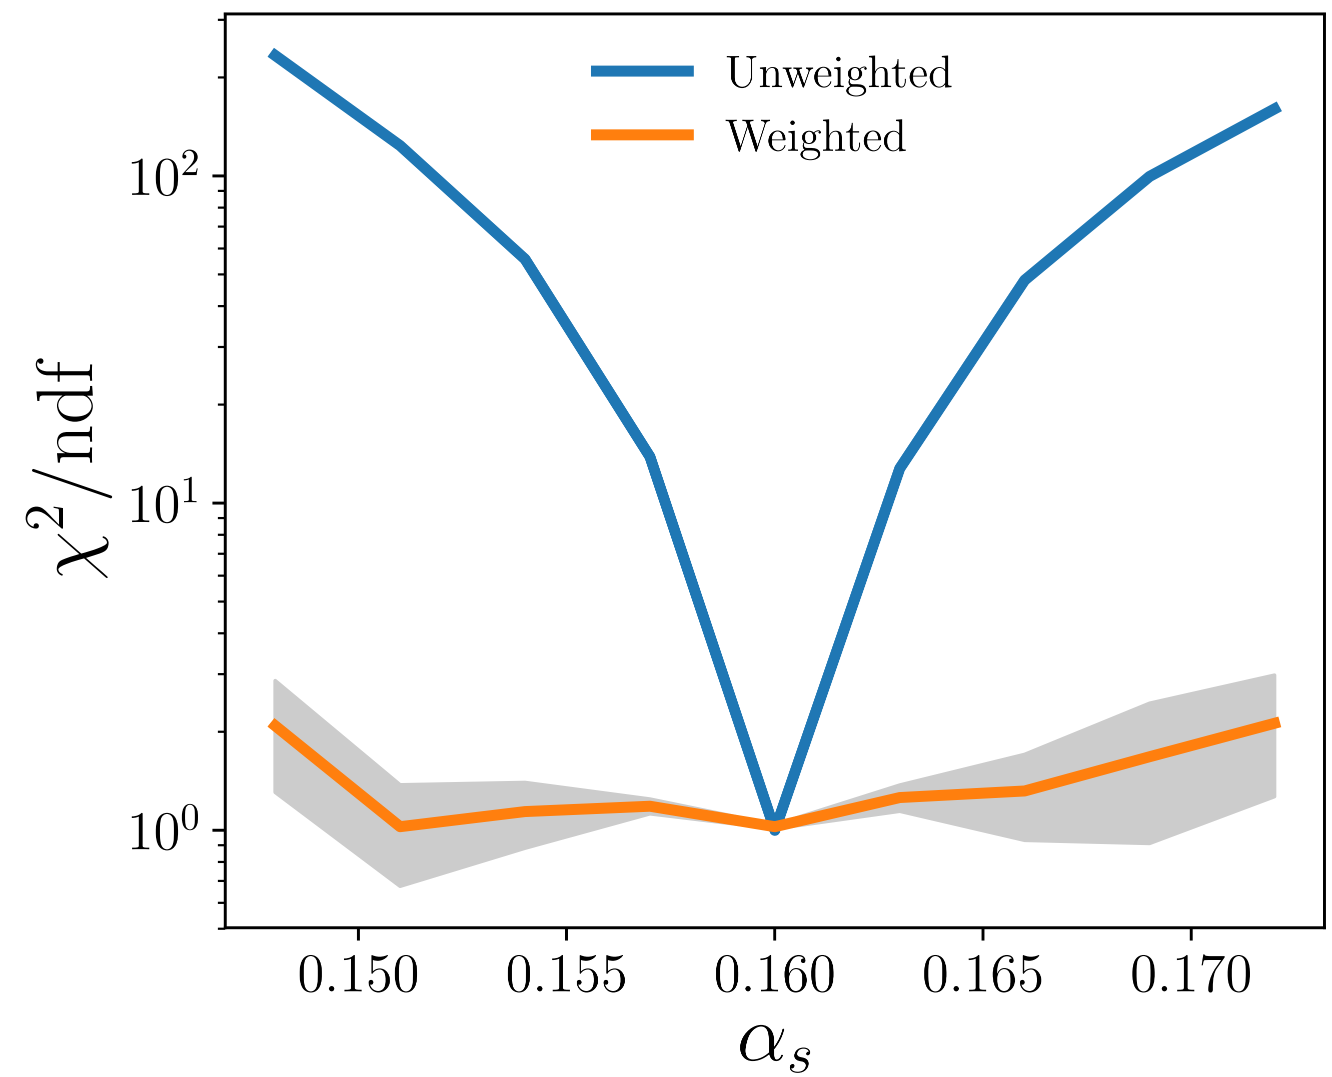

"fine structure constant" of the strong force

*see Cranmer, Pavez, Louppe, 1506.02169

Parameter estimation

One can combine a parameterized reweighting function with a classifier to fit model parameters.

Parameterized classifier:

$$f(x,\theta) = \underset{f'}{\operatorname{argmax}} \sum_{i\in\boldsymbol{\theta_0}} \log f'(x_i,\theta) + \sum_{i\in\boldsymbol{\theta}} \log(1 - f'(x_i,\theta))$$

$$\theta^* = \underset{\theta'}{\operatorname{argmax}} \sum_{i\in\boldsymbol{\theta_0}} \log f(x_i,\theta') + \sum_{i\in\boldsymbol{\theta_1}} \log(1 - f(x_i,\theta'))$$

Classifier loss to fit the parameters

Parameter estimation

Fit 3 (2 shown) parameters using the full phase space!

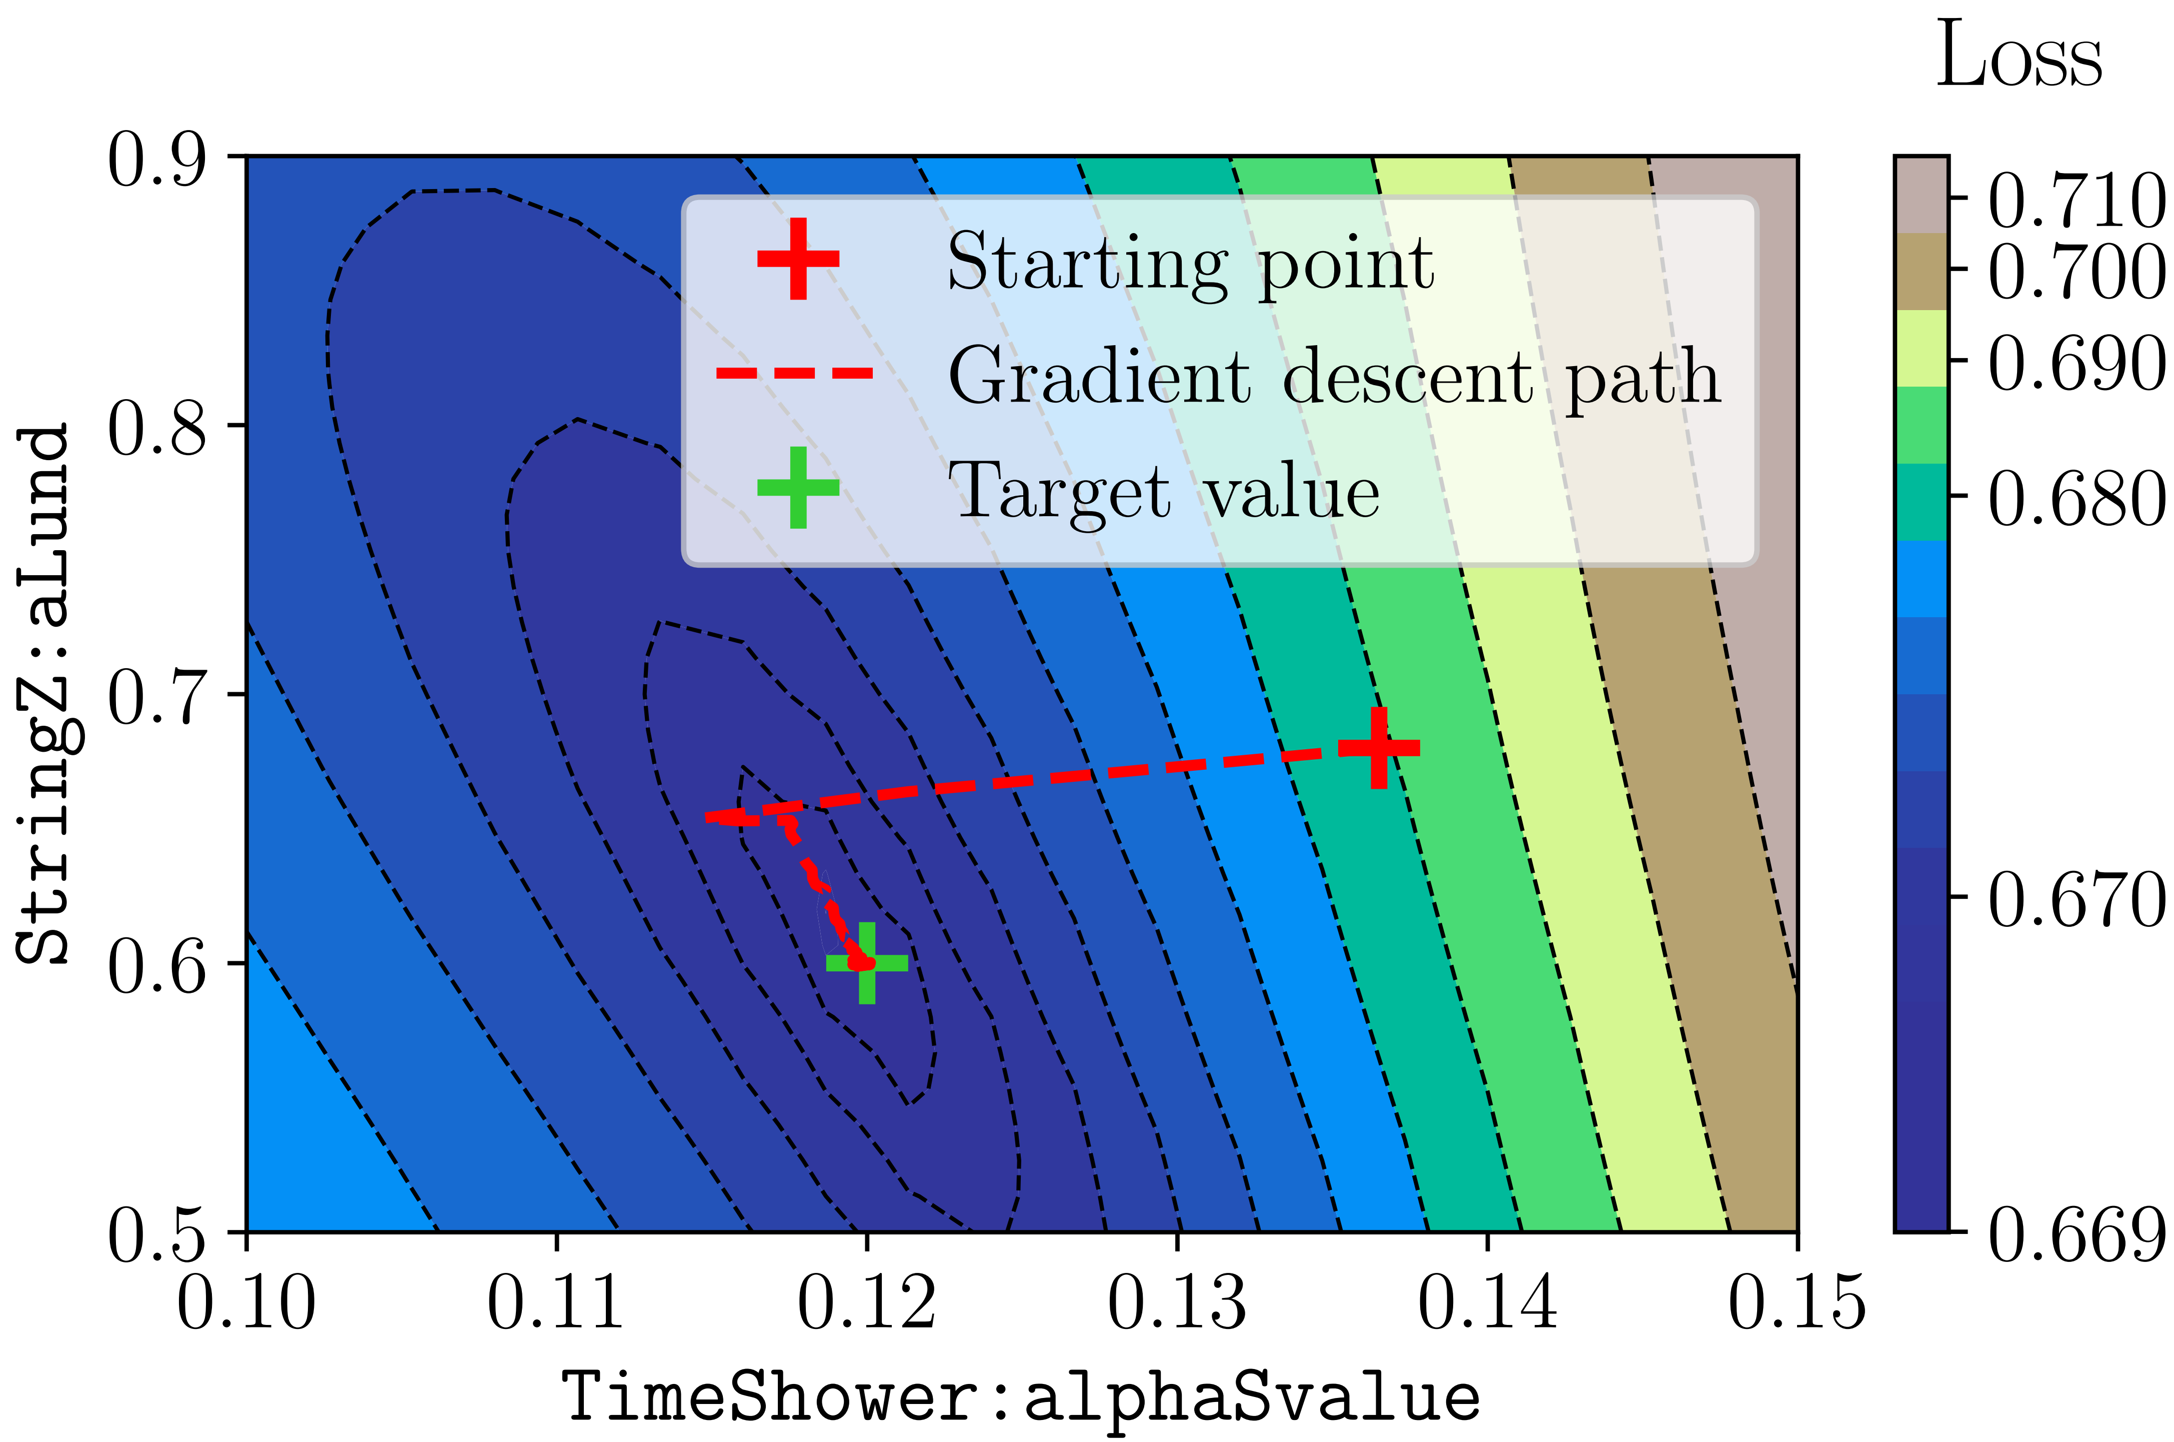

Parameter estimation: v2

What if we want to reweight with **pre-detector particles**, but fit to **detector-level objects**?

$$\theta^* = \underset{\theta'}{\mathrm{argmax}} \min_g \sum_{i \in \theta_0} \log g(x_{D,i}) \quad \text{[data]}$$

[reweighted simulation]

$$+ \sum_{i \in \theta} w(x_{T,i}, \theta) \log(1 - g(x_{D,i}))$$

$$\frac{f(x_T, \theta)}{1 - f(x_T, \theta)}$$

Intuition: reweight until you can't distinguish the data from the (reweighted) simulation!

Part I: *GANs in HEP*

Multiple applications under study; one of the most advanced is for accelerating expensive simulations. Challenge: fidelity.

Part II: *High-dimensional reweighting*

Can use classification to do reweighting and thus recycle simulations. This can be parameterized and used for fitting. High fidelity, but cannot be used to sample new examples.

Conclusions and outlook

Generative models: essential to connect our data to fundamental properties of nature.

Theory of everything

↓

Physics simulators

↓

Detector-level features

↓

Pattern recognition ⟷ Pattern recognition

Nature

↓

Experiment

↓

Detector-level features

↓

Pattern recognition

Deep learning can accelerate and enhance this work!

References

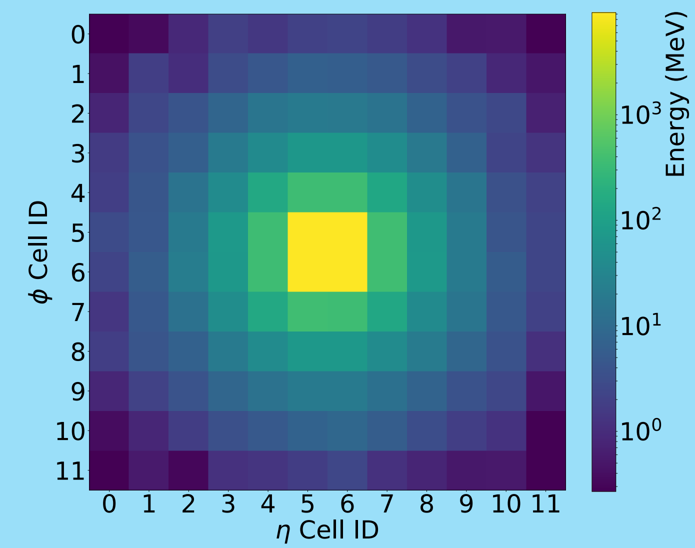

Phys. Rev. Lett. 120, 042003 (2018), 1705.02355

Phys. Rev. D 97, 014021 (2018), 1712.10321

Comput Softw. Big Sci. (2017) 1: 4, 1701.05927

ACAT 2017 Proceedings: https://arxiv.org/abs/1711.08813

NeurIPS 2017 Deep Learning for Physical Sciences

Reweighting: https://arxiv.org/abs/1909.03081

https://github.com/hep-lbdl/CaloGAN

Collaborators


Luke de Oliveira (Twillo)

Michela Paganini (Facebook)

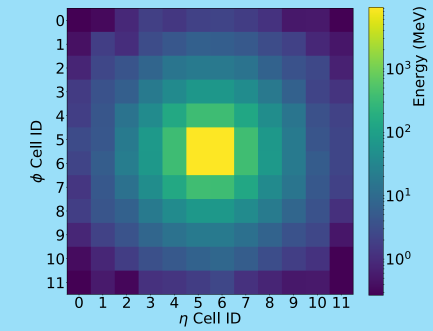


Steve Farrell (NERSC)

Wahid Bhimji (NERSC)

Wyatt Blair (Berkeley)

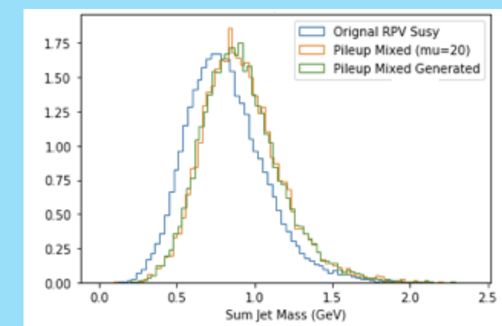

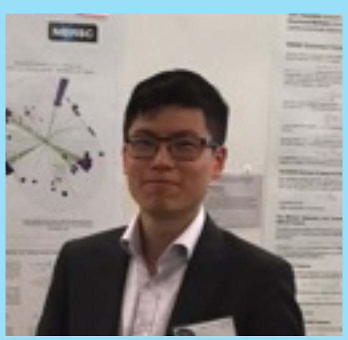


Josh Lin (Berkeley)

Wahid Bhimji (NERSC)

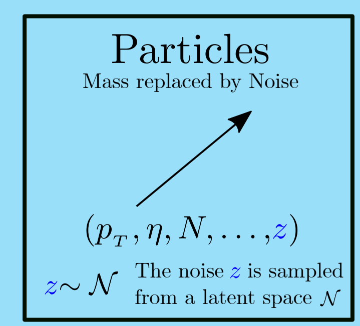

Anders Andreassen (Google)

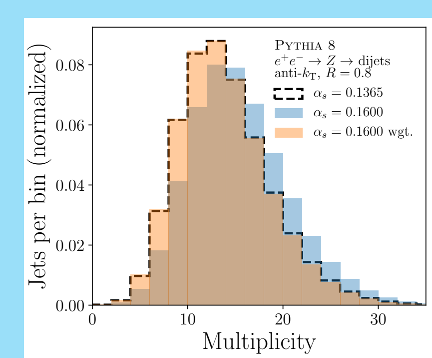

Questions?

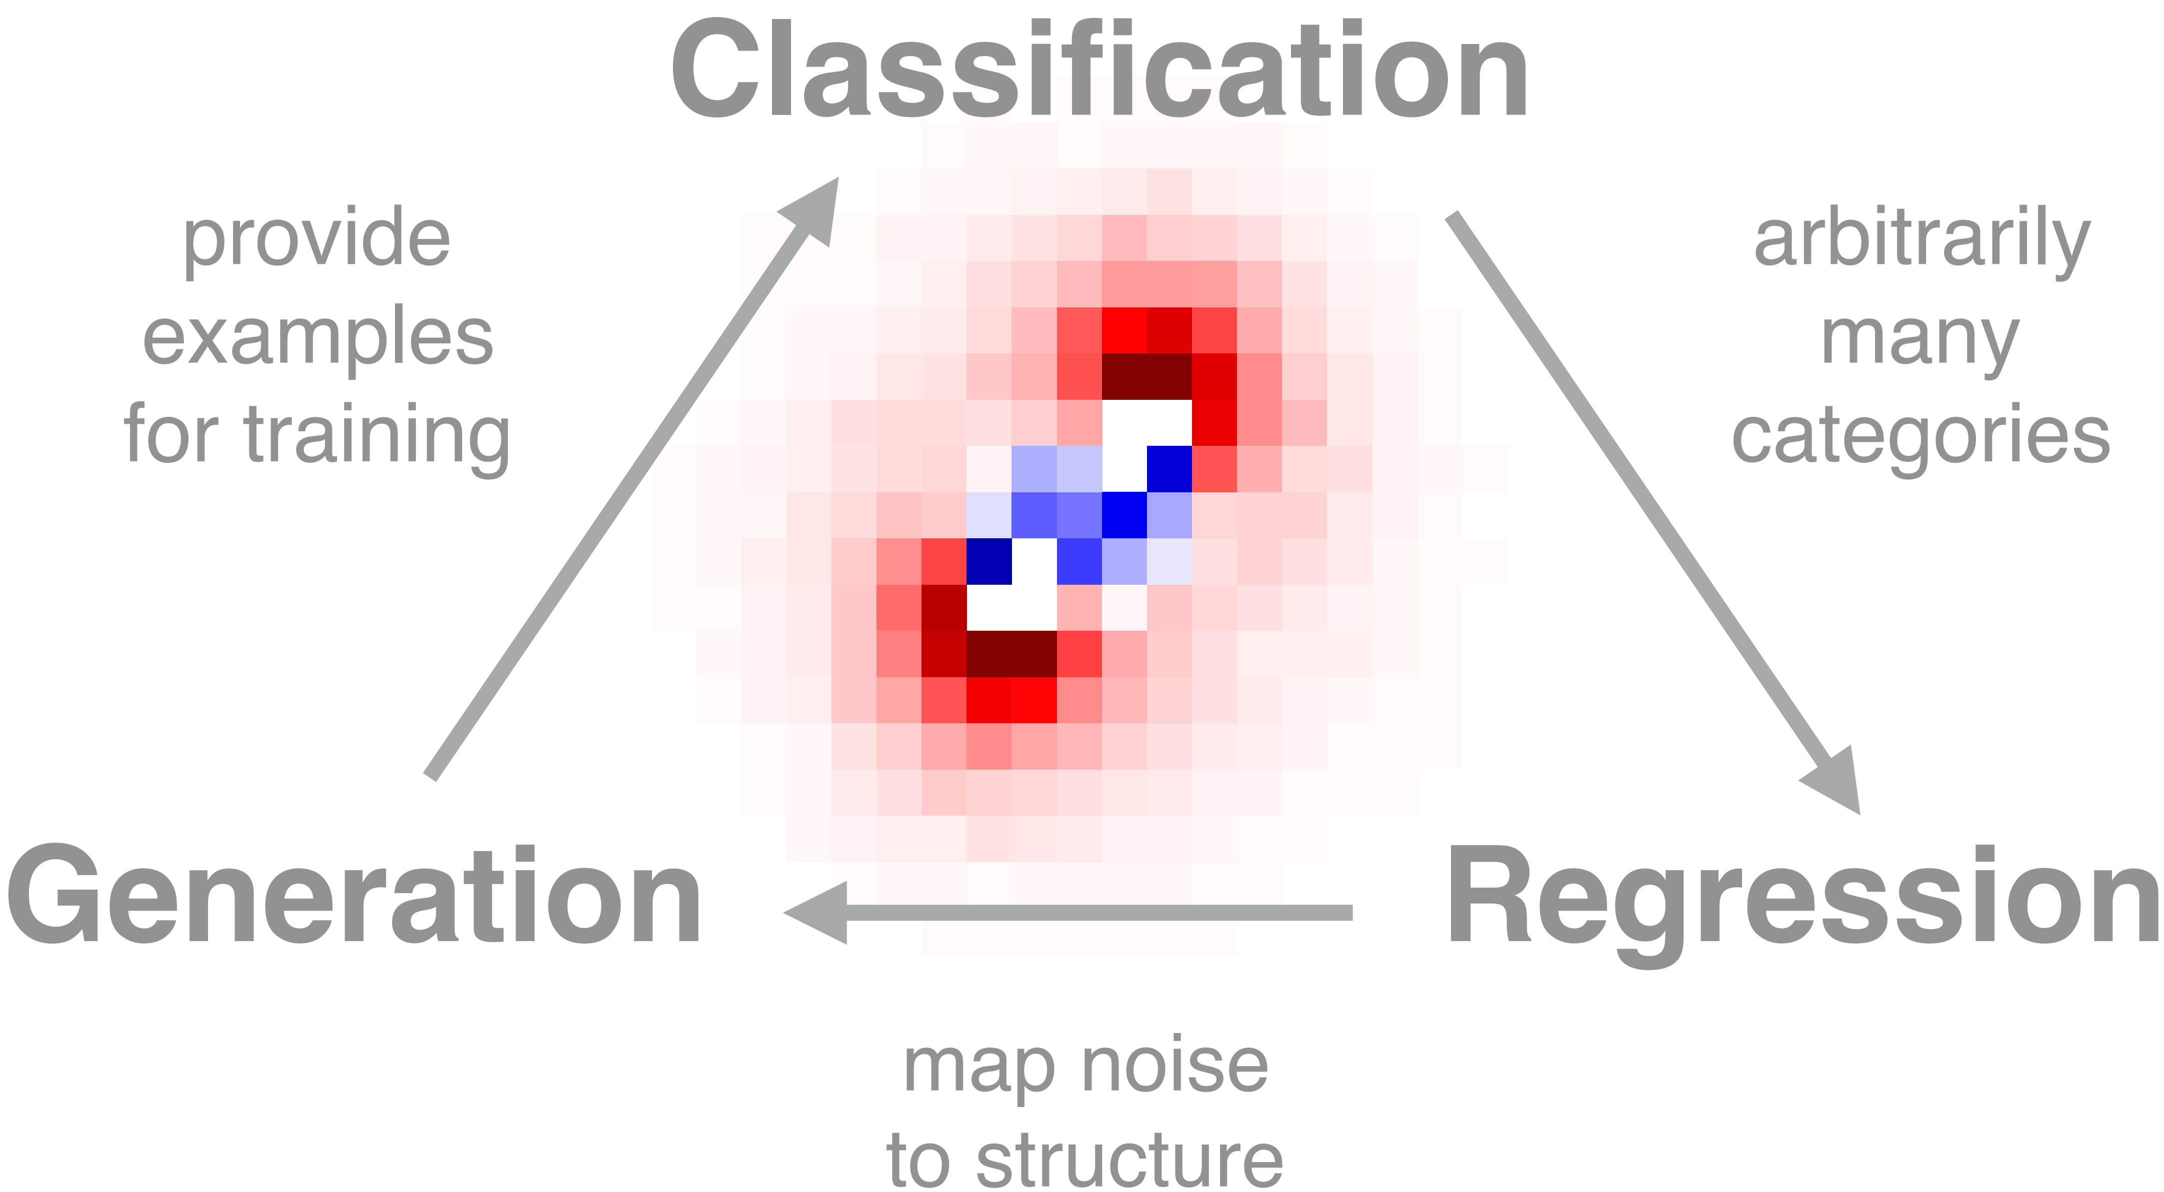
Classification
provide
examples
for training
arbitrarily
many
categories
Generation
Regression
map noise
to structure