Accelerating Data Collection and Processing at the Large Hadron Collider

Benjamin Nachman

Lawrence Berkeley National Laboratory and the Simons Institute

BERKELEY LAB

SIMONS INSTITUTE for the Theory of Computing

BERKELEY EXPERIMENTAL PARTICLE PHYSICS

Real Time Decision Making for Applications in the Natural Sciences and Physical Systems, Feb. 28, 2018

Science at the LHC

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

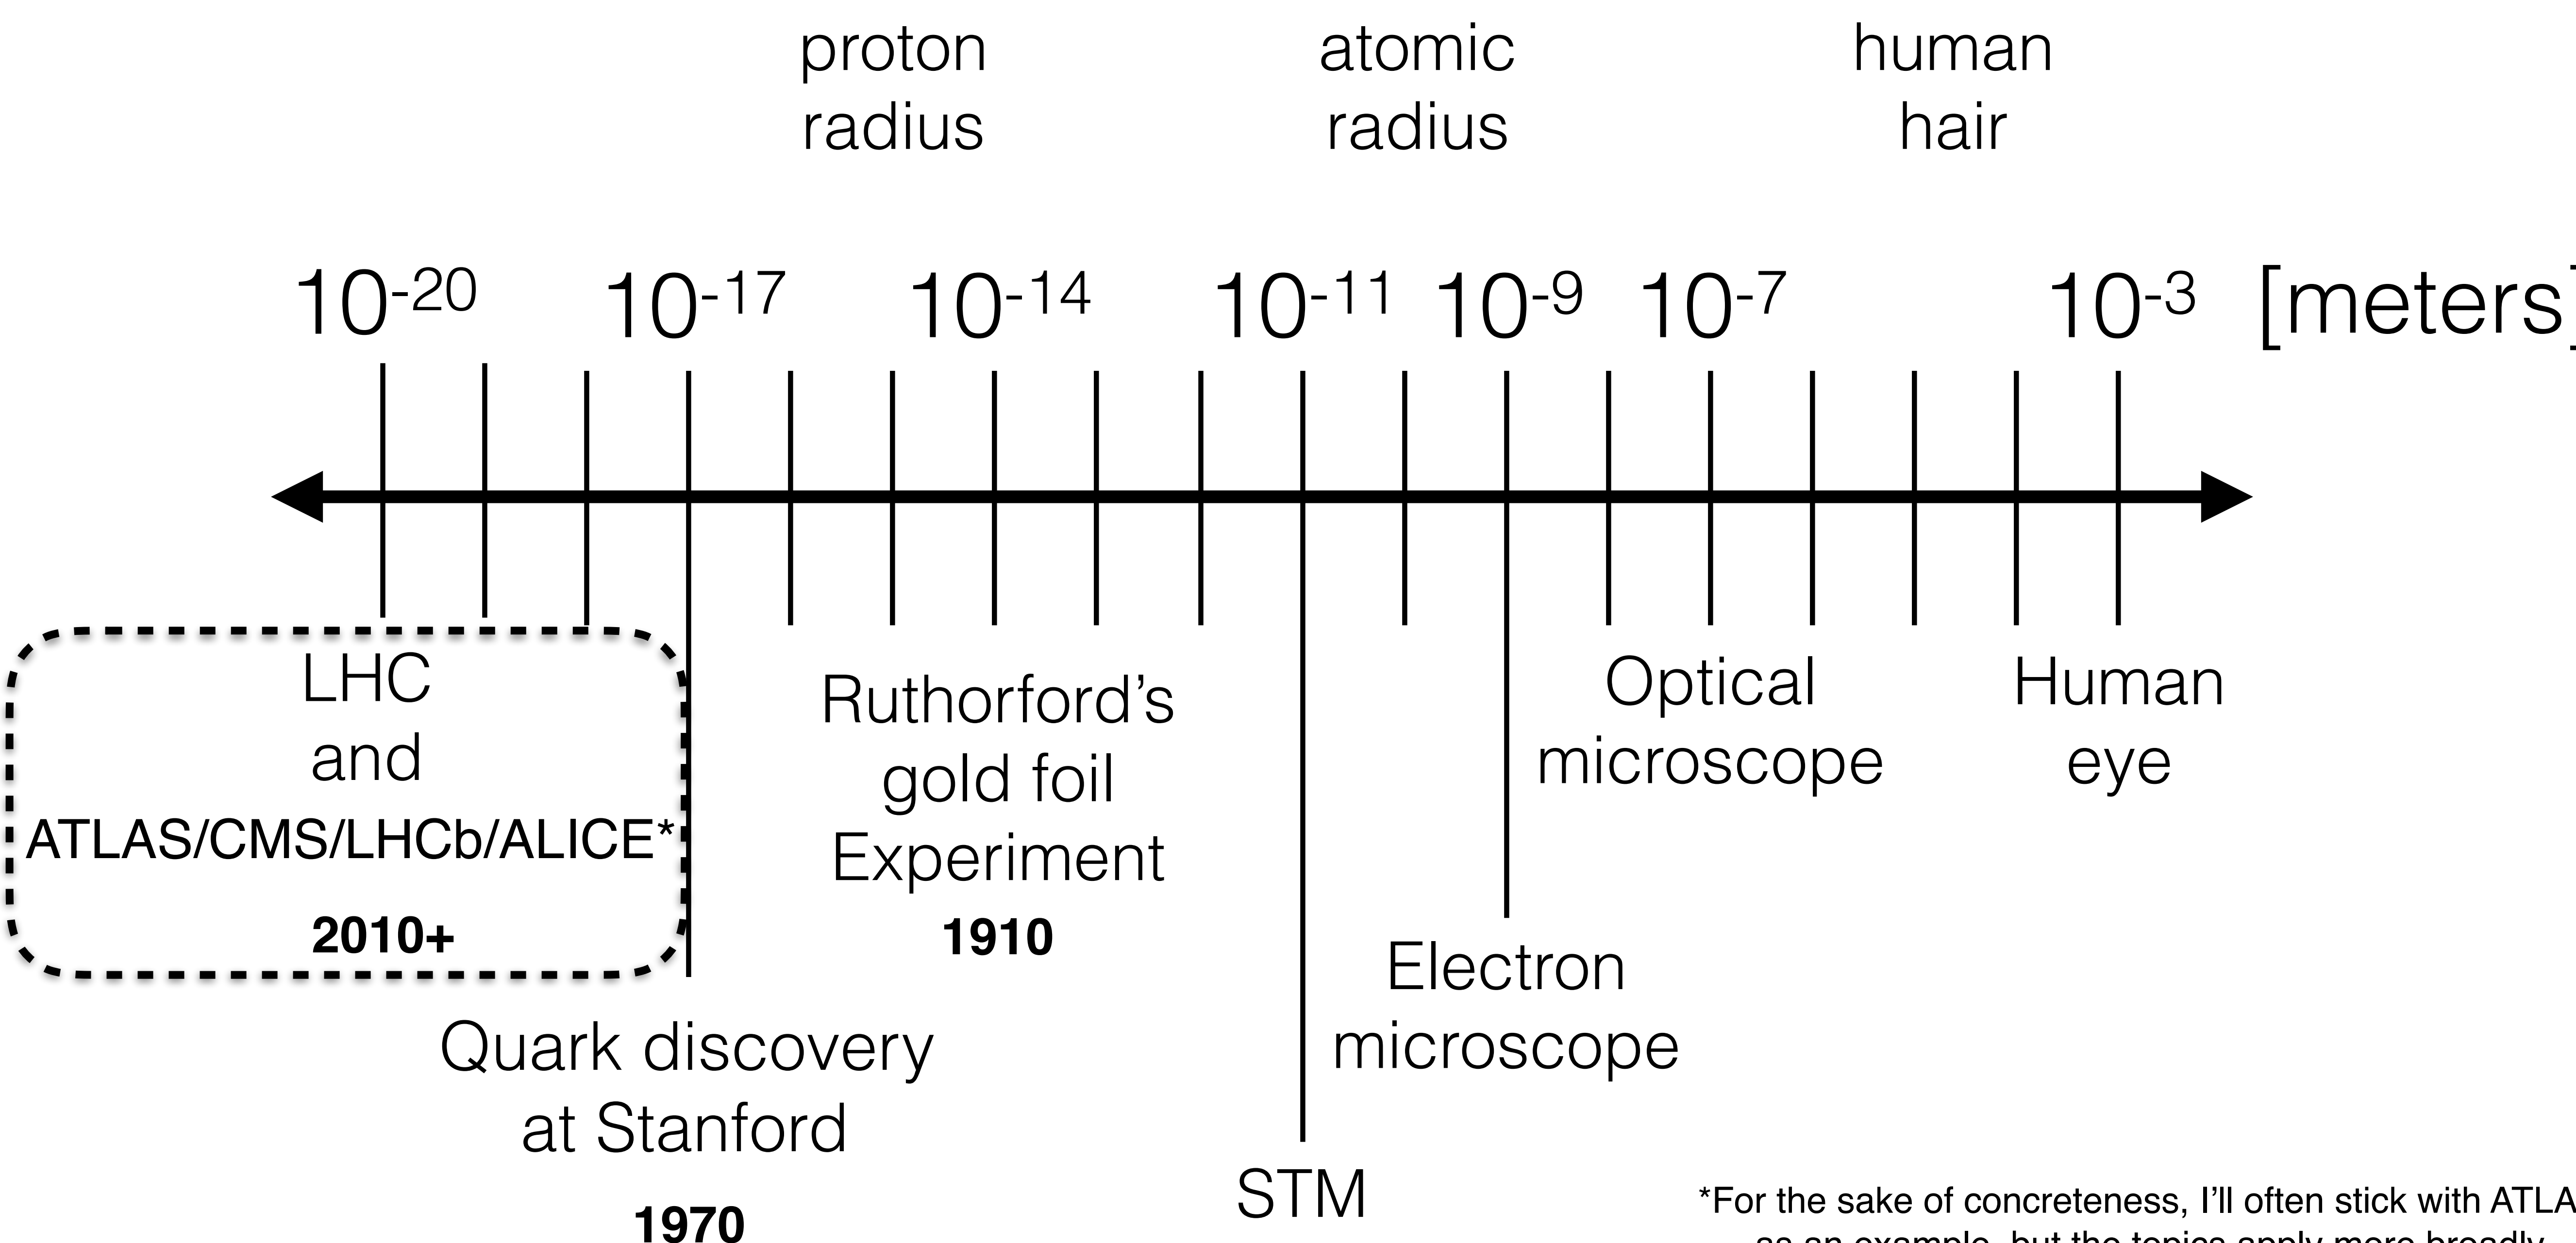

*For the sake of concreteness, I'll often stick with ATLAS as an example, but the topics apply more broadly

Data pipeline at the LHC

data / simulation

25 ns		O(100) *pp* collisions
10^{-19}-10^{-15} s	~ms	(sub-)nuclear physics
0.01-20 ns	~min	out-going particles interact with detector
1-100 ns	~100 ms	detector response (signal formation + digitization)
2.5 µs		hardware-based trigger decision
200 ms		software-based trigger decision
~100 ms		event reconstruction
		event processing (skim, thin, augment)
~months		final data analysis (uses millions of events)

Data pipeline at the LHC

data / simulation

Time (data / simulation)	Stage
25 ns	O(100) *pp* collisions
10^{-19}-10^{-15} s / ~ms	(sub-)nuclear physics
0.01-20 ns / ~min	out-going particles interact with detector
1-100 ns / ~100 ms	**detector response** (signal formation + digitization)
2.5 μs	hardware-based trigger decision
200 ms	software-based trigger decision
~100 ms	event reconstruction
	event processing (skim, thin, augment)
~months	final data analysis (uses millions of events)

Collider-based HEP detectors are like onions

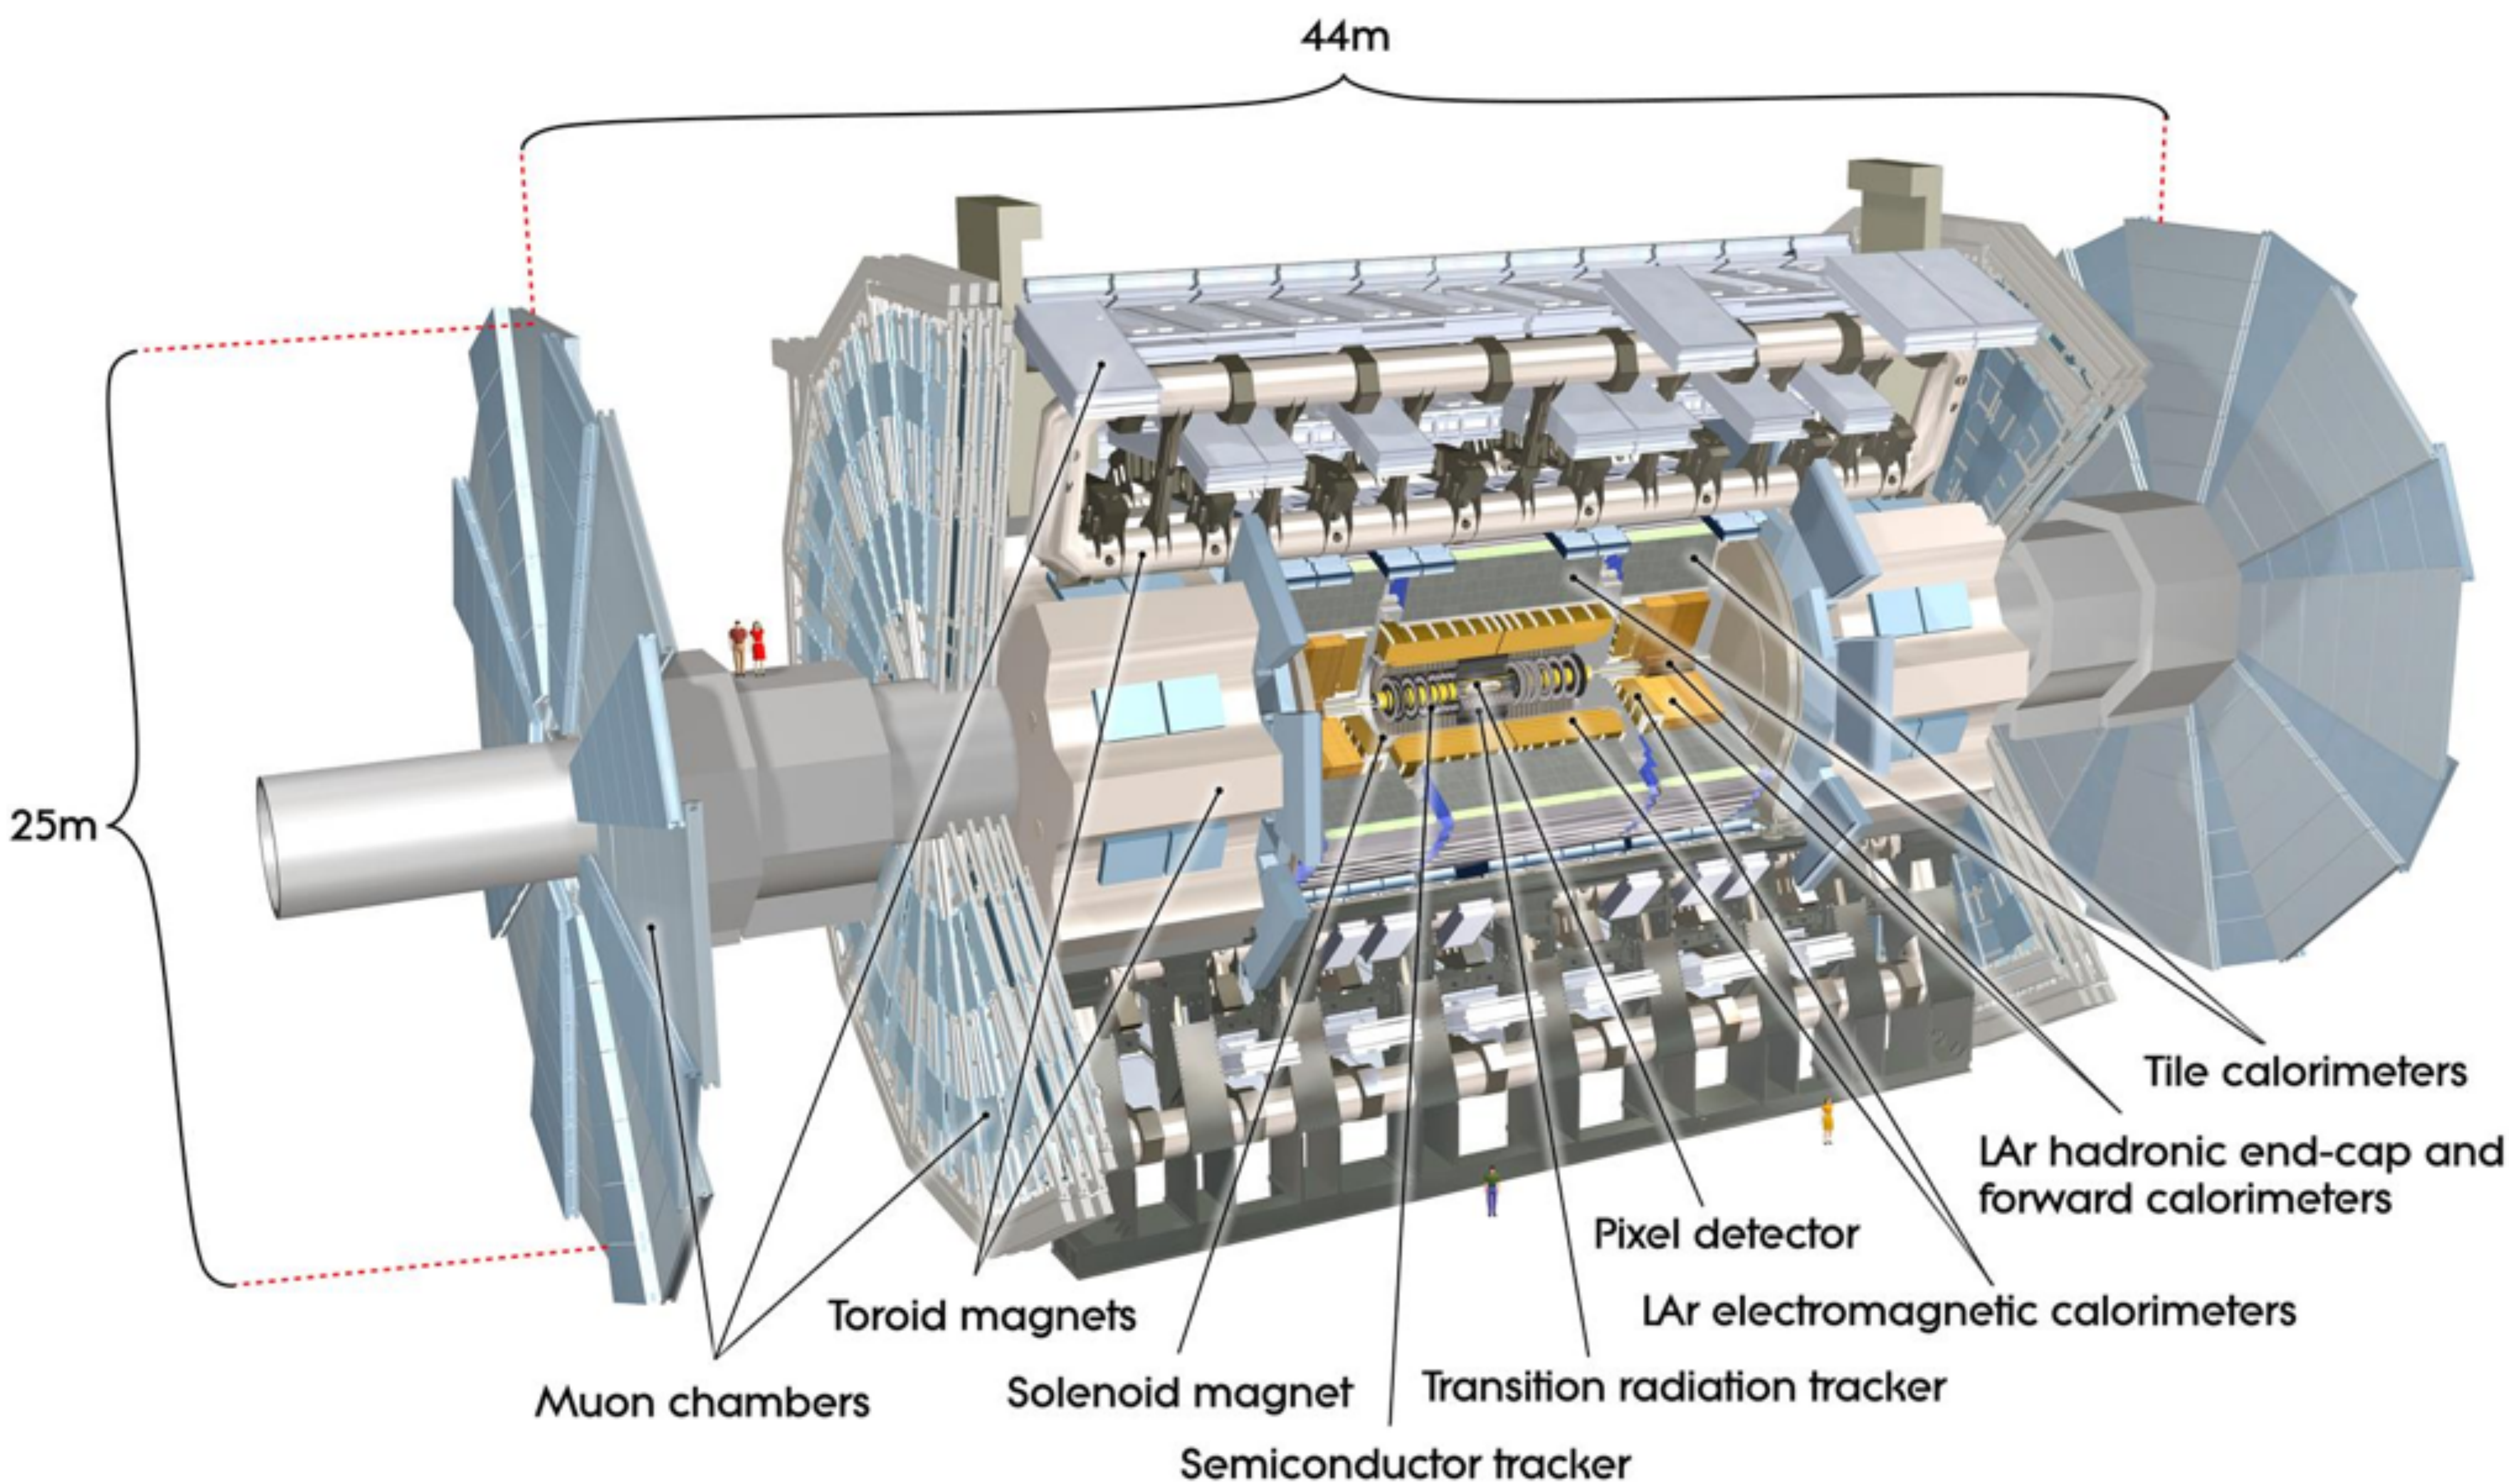

Biggest challenge for data volume is **innermost** layer

Collider-based HEP detectors are like ~~onions~~ leeks

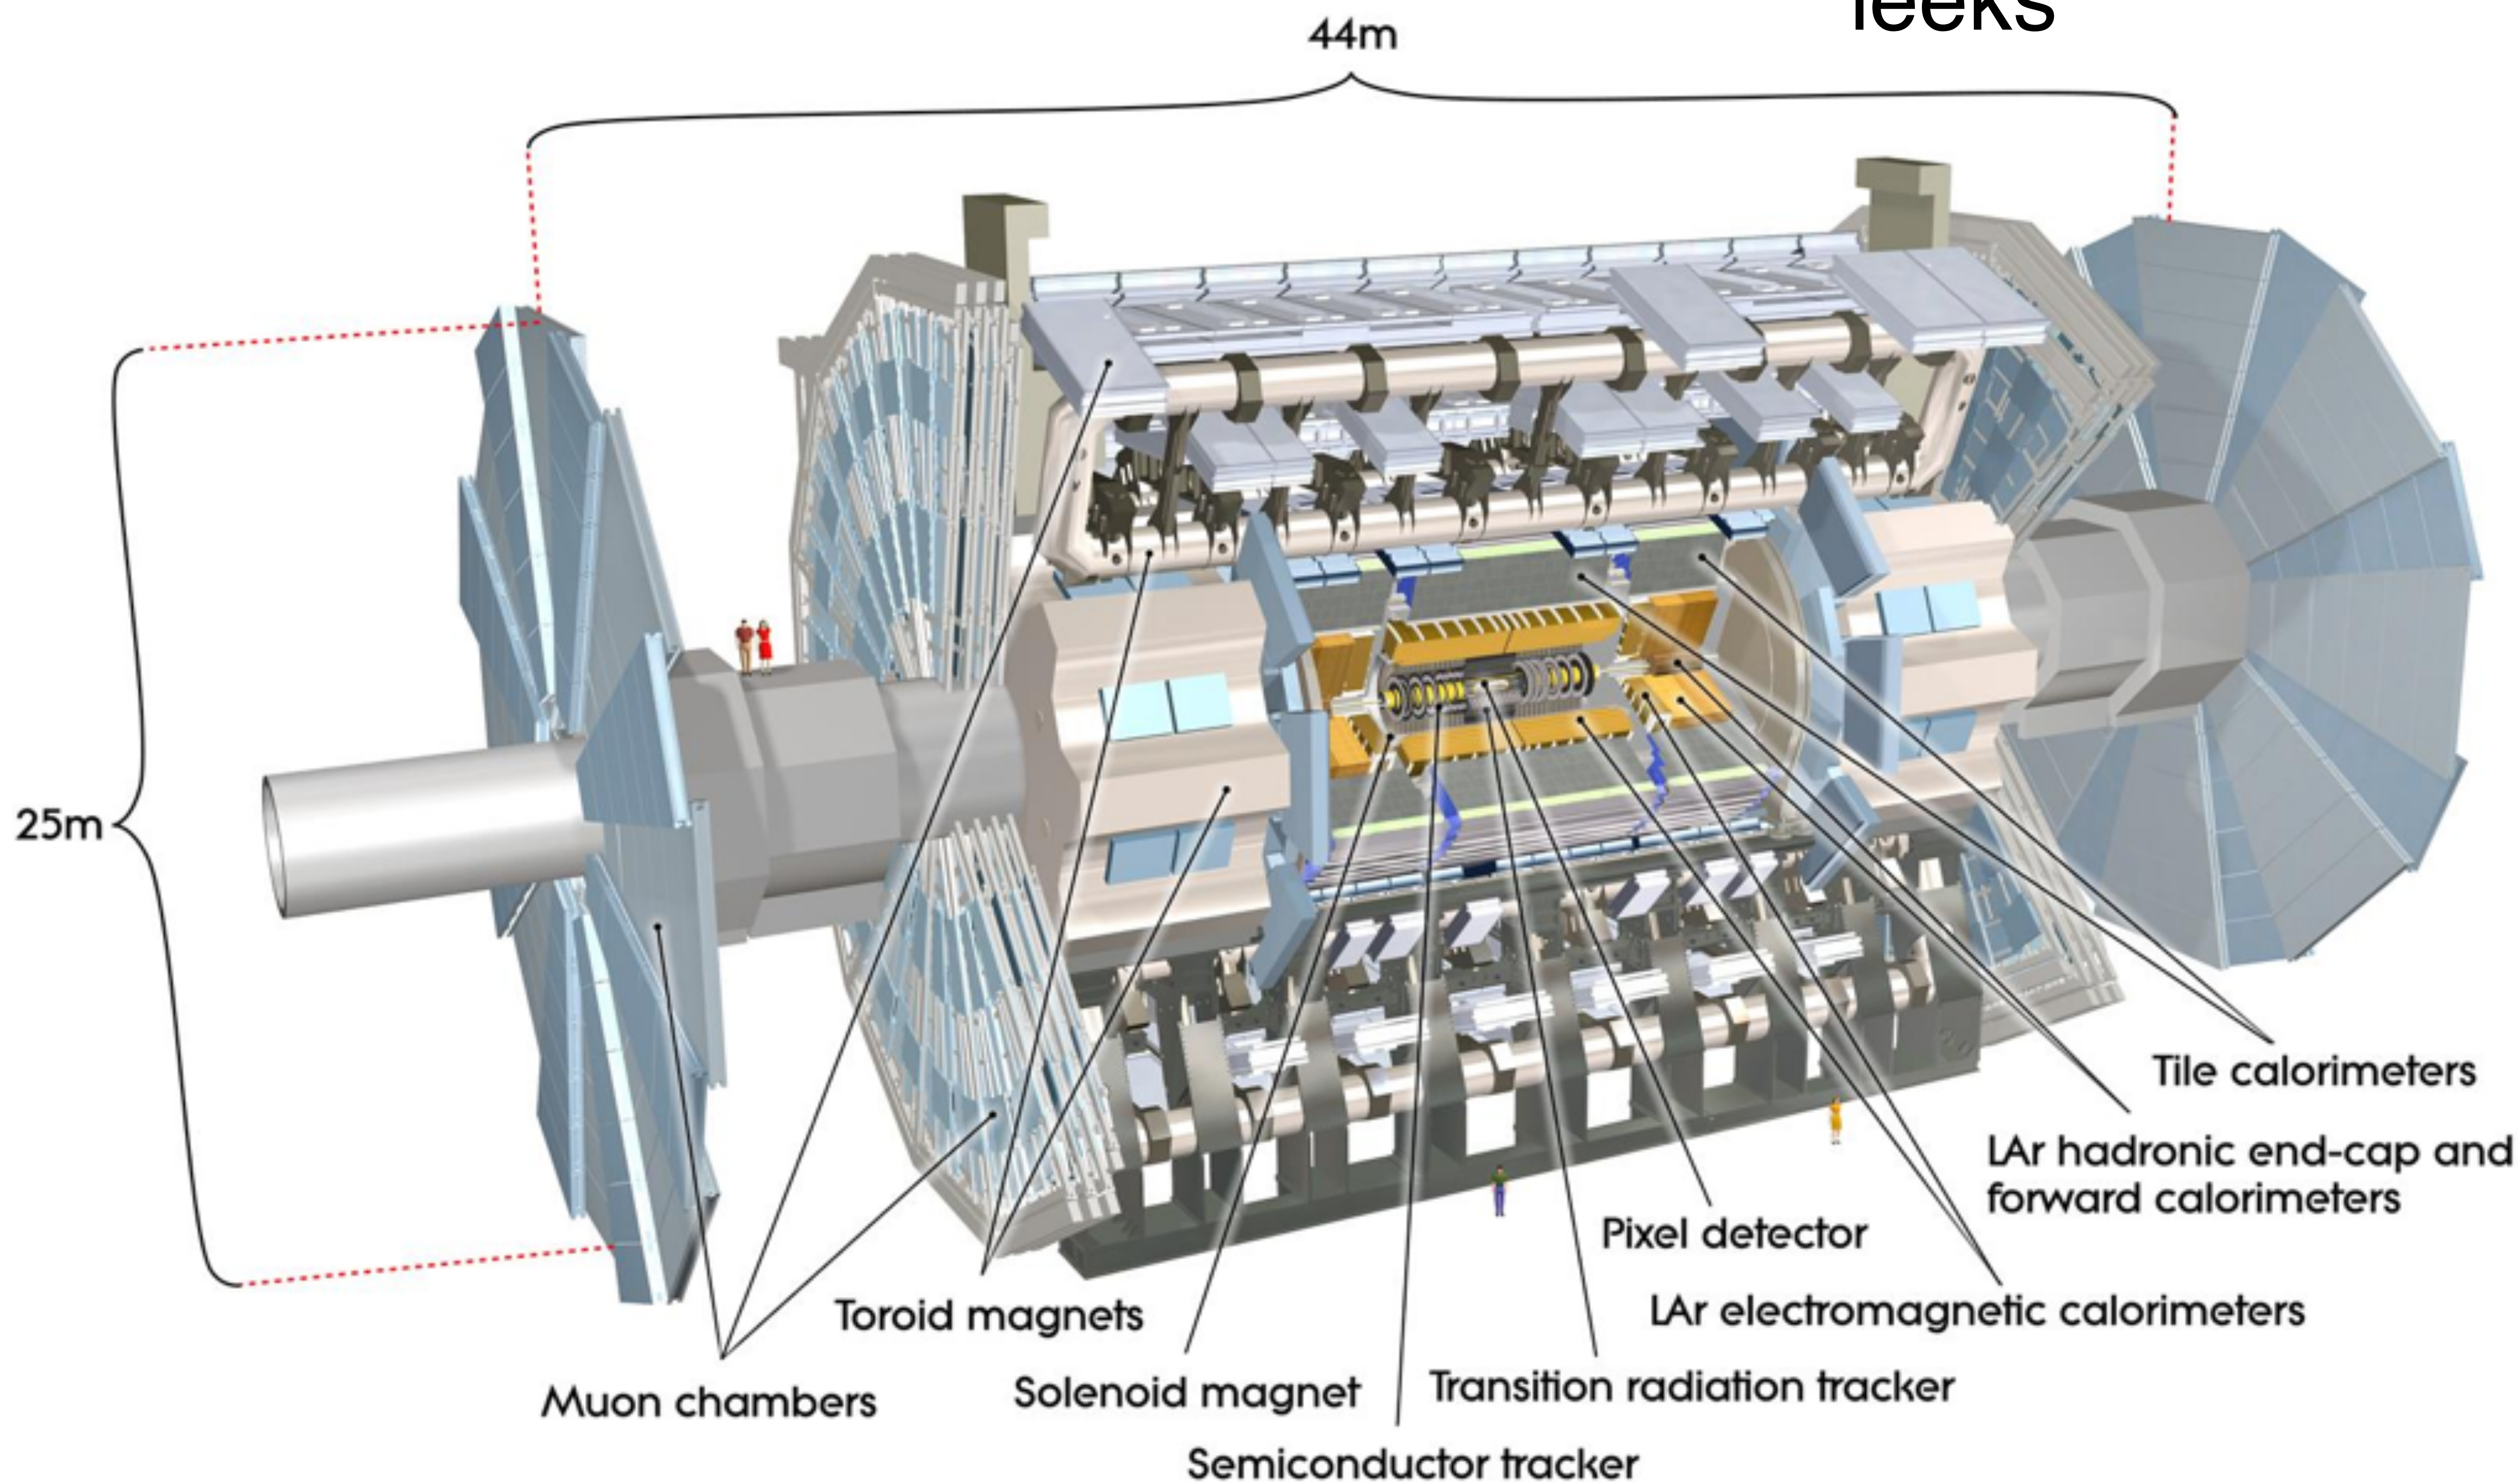

Biggest challenge for data volume is **innermost** layer

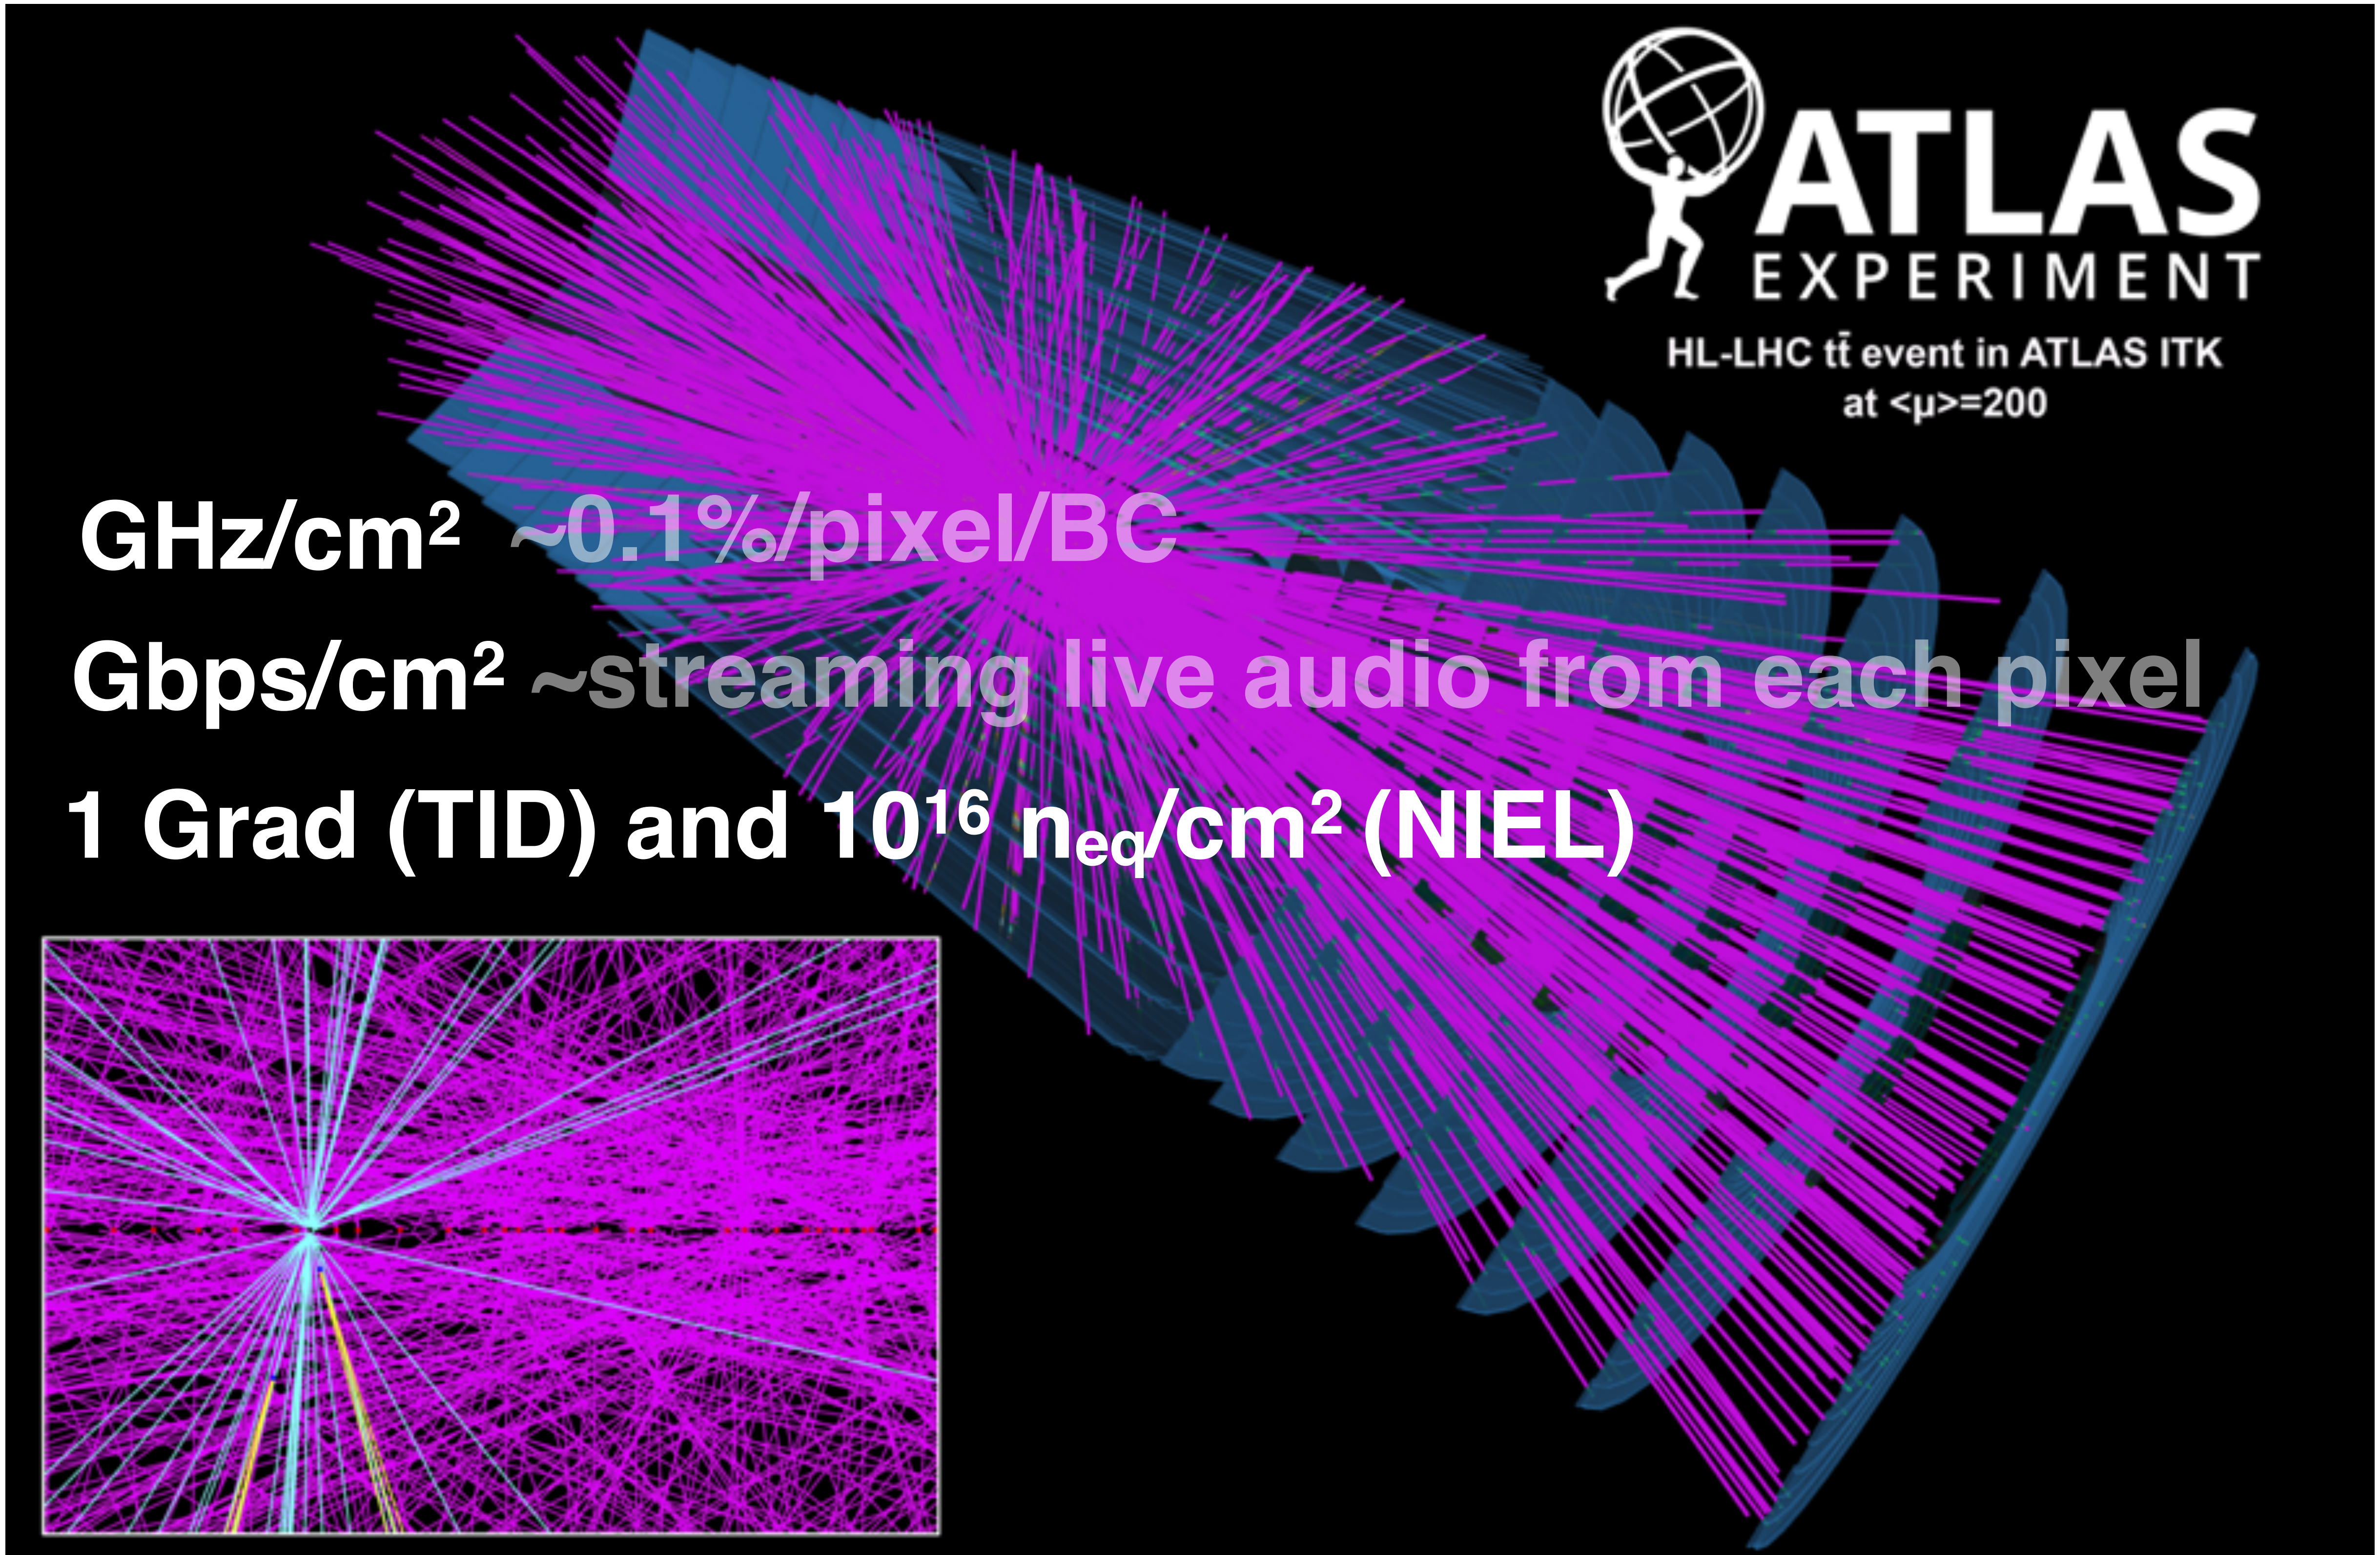

GHz/cm^2 ~0.1%/pixel/BC

Gbps/cm^2 ~streaming live audio from each pixel

1 Grad (TID) and 10^{16} n_{eq}/cm^2 (NIEL)

Challenge of the LHC upgrade (~2025; HL-LHC)

Generation	Run 1 (FEI3, PSI46)	Runs 2+3 (FEI4, PSI46DIG)	Runs 4+5
Chip Size	7.5 x 10.5 mm² 8 x 10 mm²	20 x 20 mm² 8 x 10 mm²	> 20 x 20 mm²
Transistors	3.5 M 1.3 M	87 M	~1 G
Hit Rate	100 MHz/cm²	400 MHz/cm²	~2 GHz/cm²
Hit Memory / Chip	0.1 Mb	1 Mb	~16 Mb
Trigger Rate	100 kHz	100 kHz	200 kHz - 1MHz
Trigger Latency	2.5 μs 3.2 μs	2.5 μs 3.2 μs	6 - 20 μs
Readout rate	40 Mb/s	320 Mb/s	1-4 Gb/s
Radiation	100 Mrad	200 Mrad	1 Grad
Technology	250 nm	130 nm 250 nm	65 nm
Power	~1/4 W/cm²	~1/4 W/cm²	1/2 - 1 W/cm²

Challenge of the LHC upgrade (~2025; HL-LHC)

Generation	Run 1 (FEI3, PSI46)	Runs 2+3 (FEI4, PSI46DIG)	Runs 4+5
Chip Size	7.5 x 10.5 mm^2 8 x 10 mm^2	20 x 20 mm^2 8 x 10 mm^2	> 20 x 20 mm^2
Transistors	3.5 M 1.3 M	87 M	~1 G
Hit Rate	100 MHz/cm^2	400 MHz/cm^2	~2 GHz/cm^2
Hit Memory / Chip	0.1 Mb	1 Mb	~16 Mb
Trigger Rate	100 kHz	100 kHz	200 kHz - 1MHz
Trigger Latency	2.5 µs 3.2 µs	2.5 µs 3.2 µs	6 - 20 µs
Readout rate	40 Mb/s	320 Mb/s	1-4 Gb/s
Radiation	100 Mrad	200 Mrad	1 Grad
Technology	250 nm	130 nm 250 nm	65 nm
Power	~1/4 W/cm^2	~1/4 W/cm^2	1/2 - 1 W/cm^2

e.g. the camera in your phone on steroids, next to a nuclear reactor

(unfortunately, Apple doesn't make one of these)

The big LHC experiments are combing forces

A significant component of the design and testing is happening right here at Berkeley!

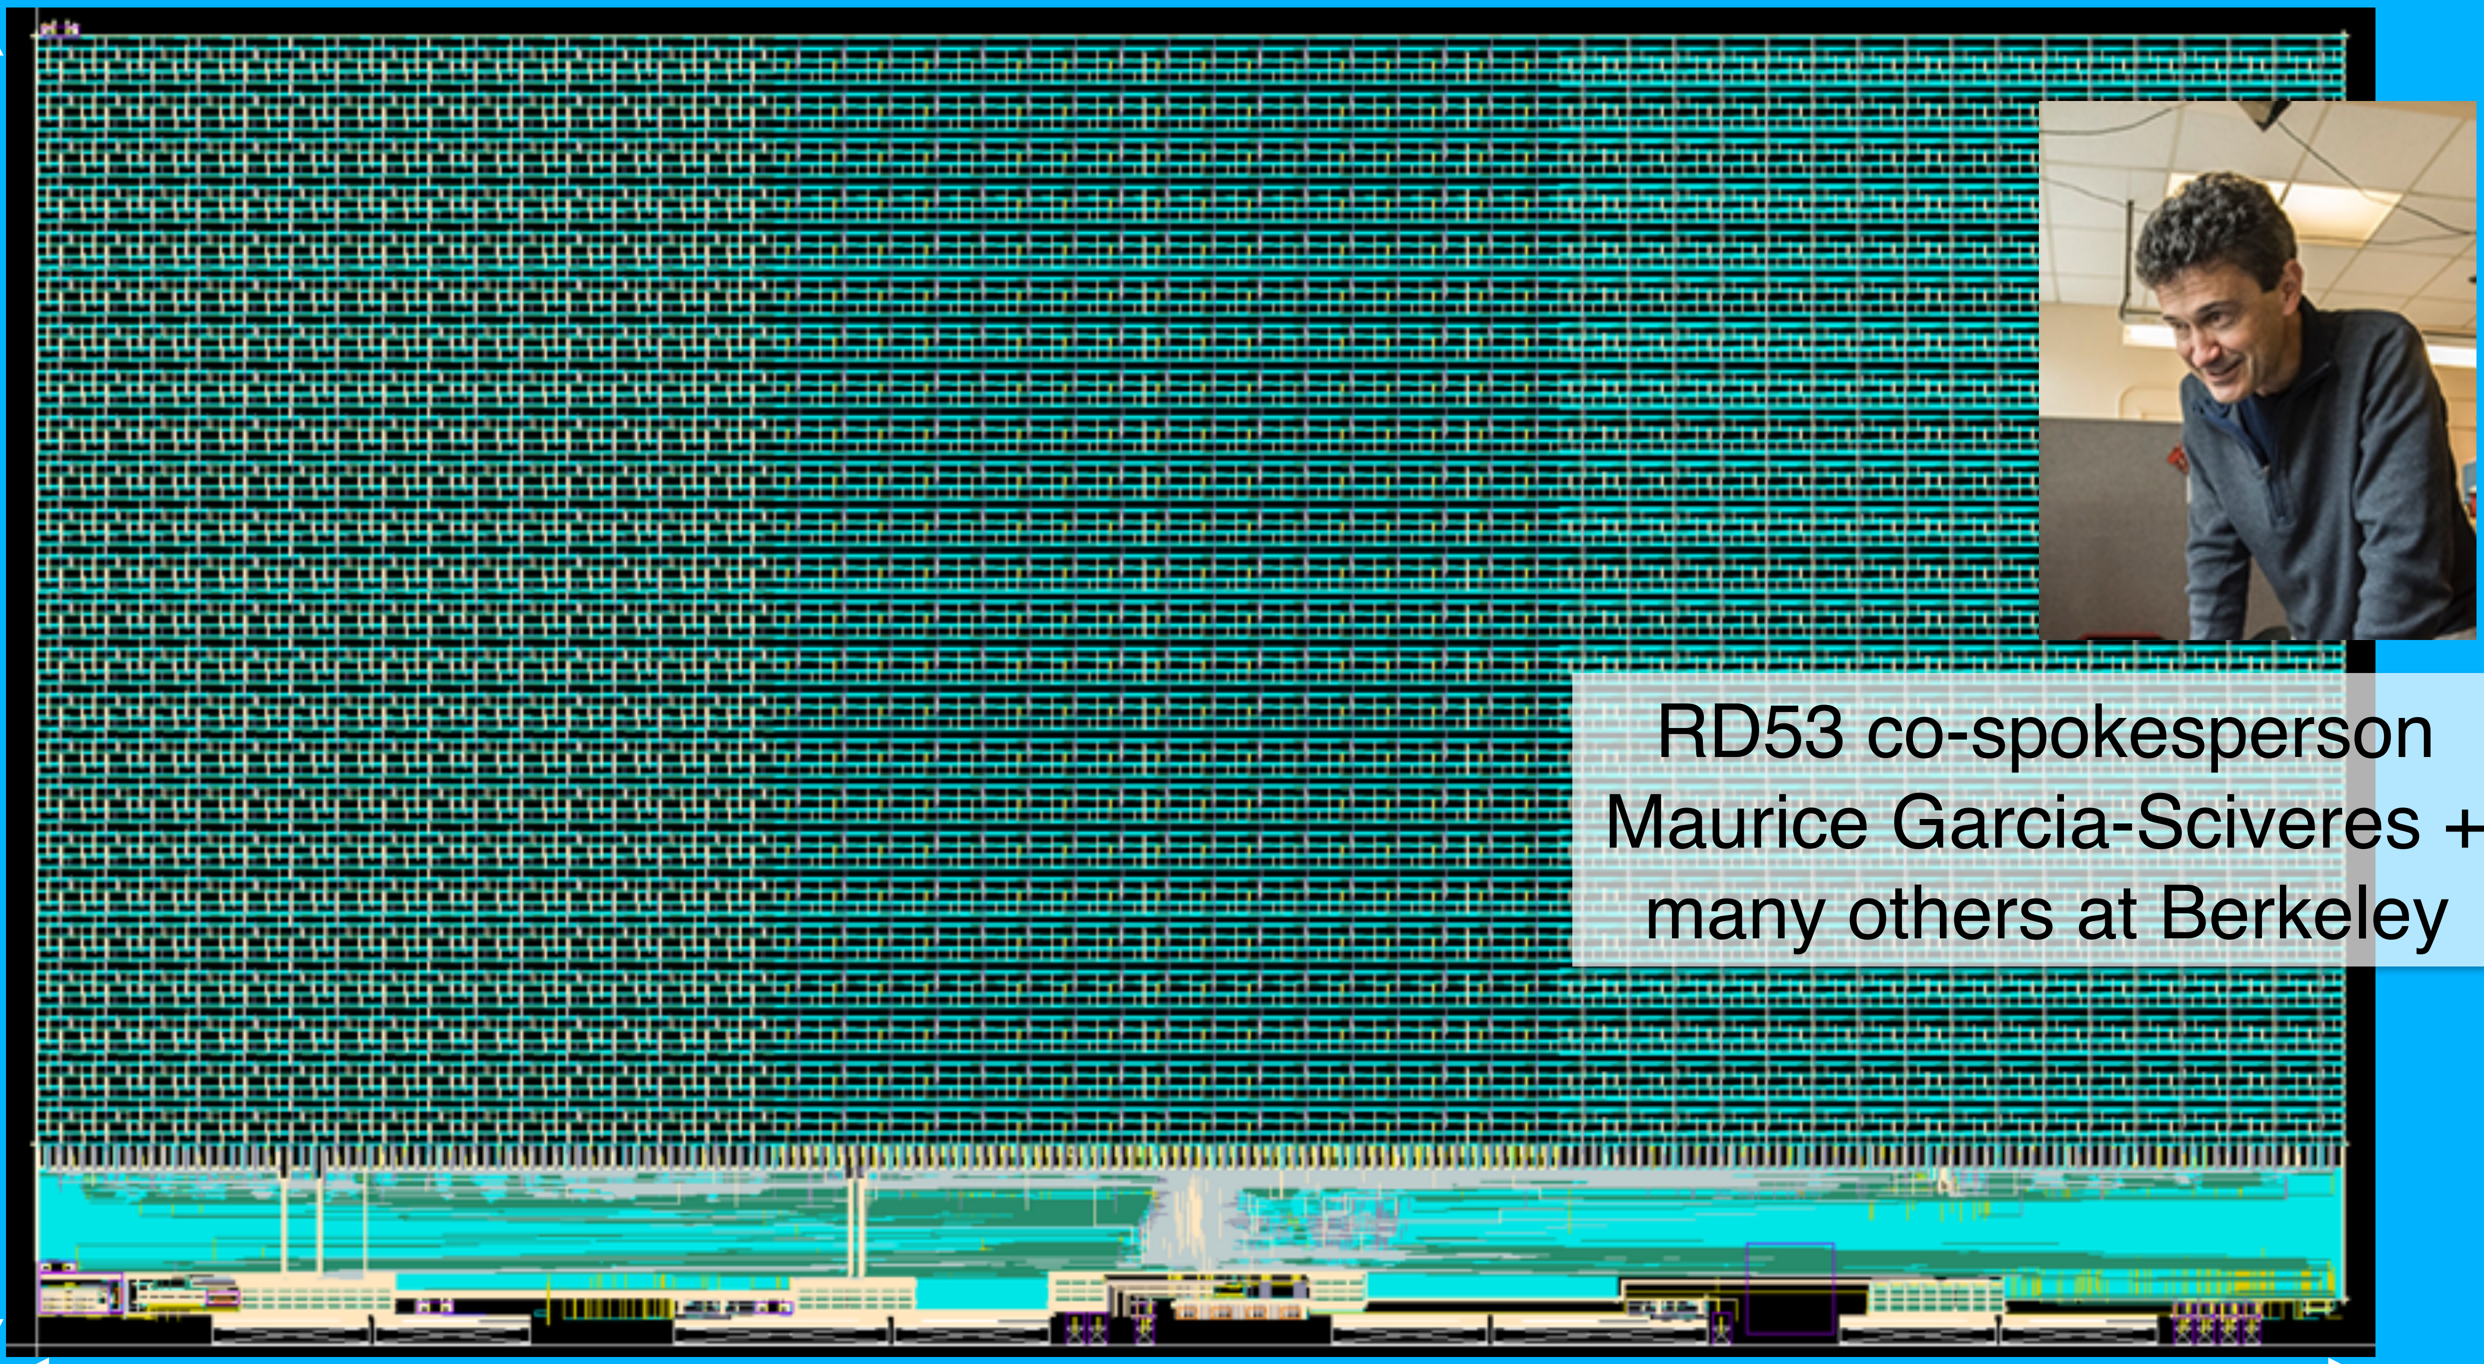

RD53 co-spokesperson Maurice Garcia-Sciveres + many others at Berkeley

11.8 mm ; 192 pixels

50 x 50 μm² pixels

20 mm ; 400 pixels

Real time decision making: what to readout?

Readout regions N x M pixel regions; helps to recover small charge hits. What is optimal?

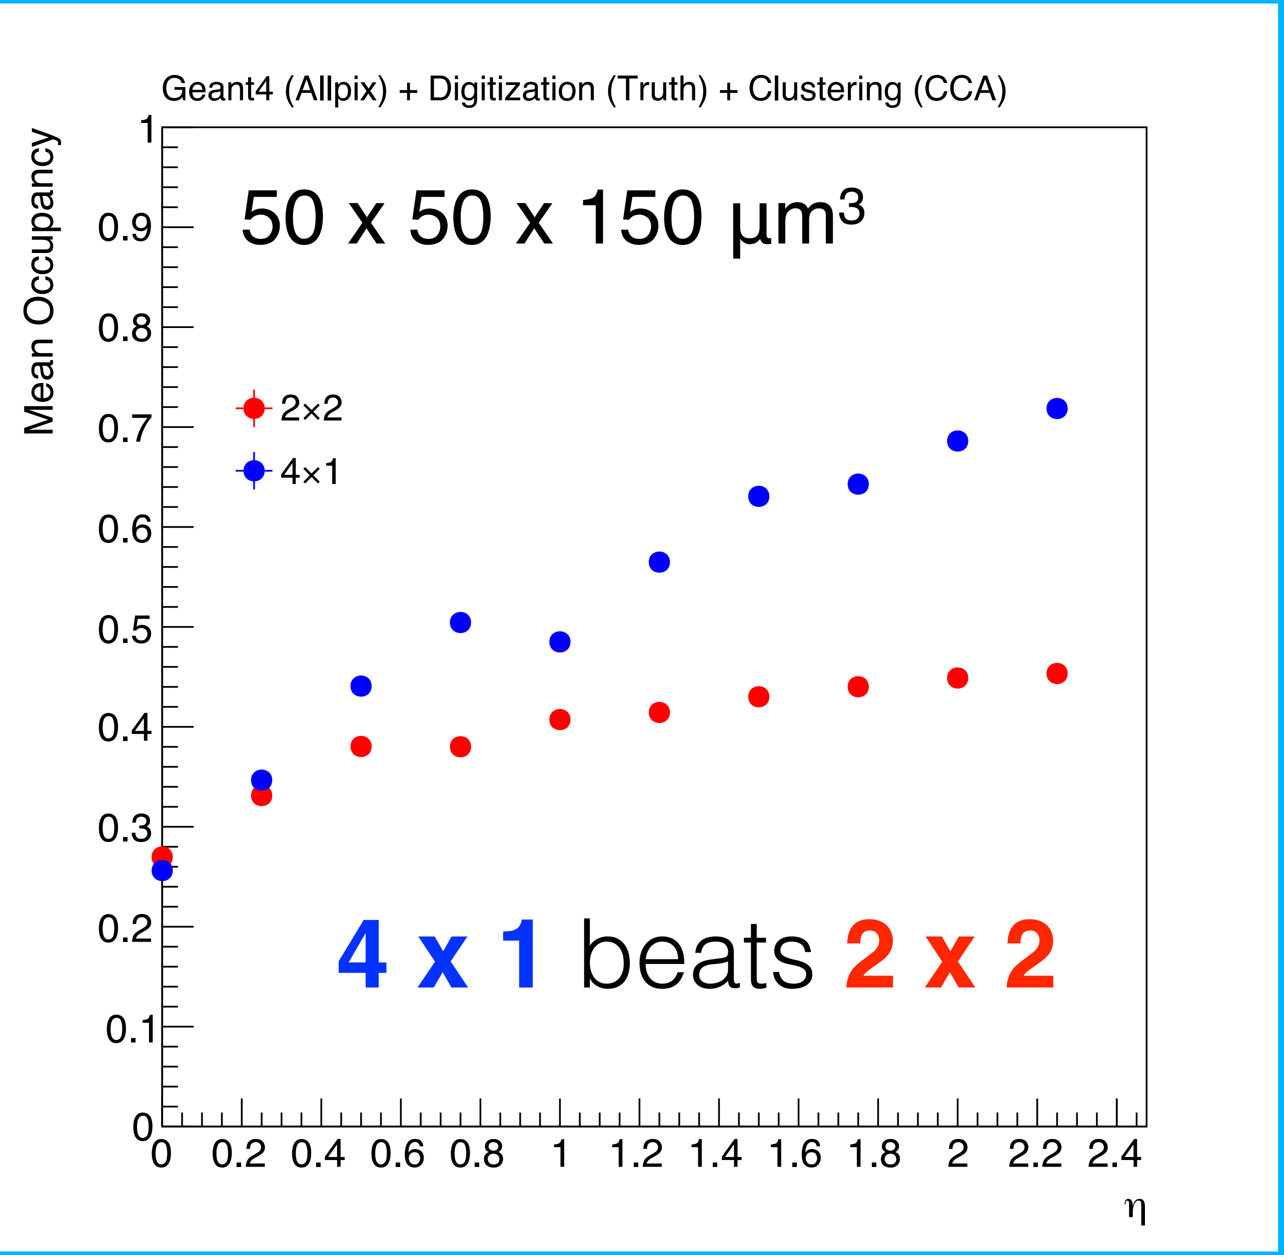

This is just one of many choices we are currently studying!

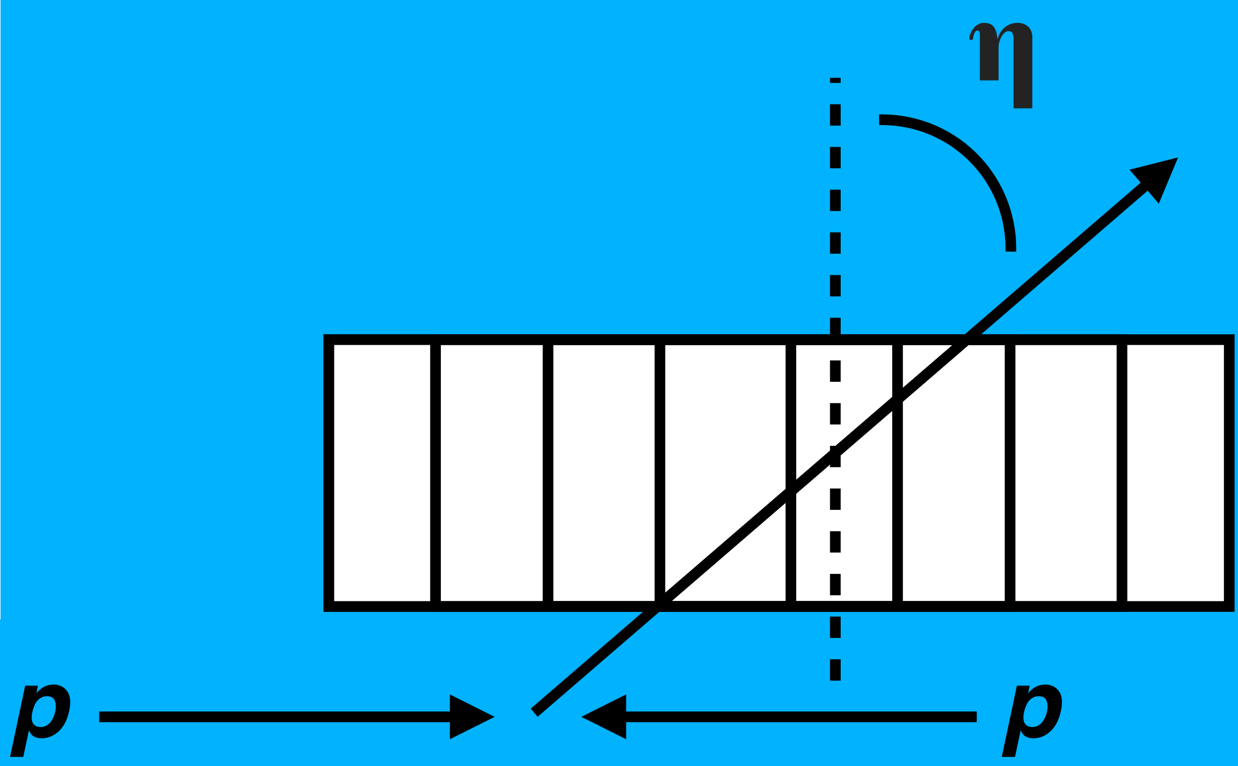

Real time decision making: how to read it out?

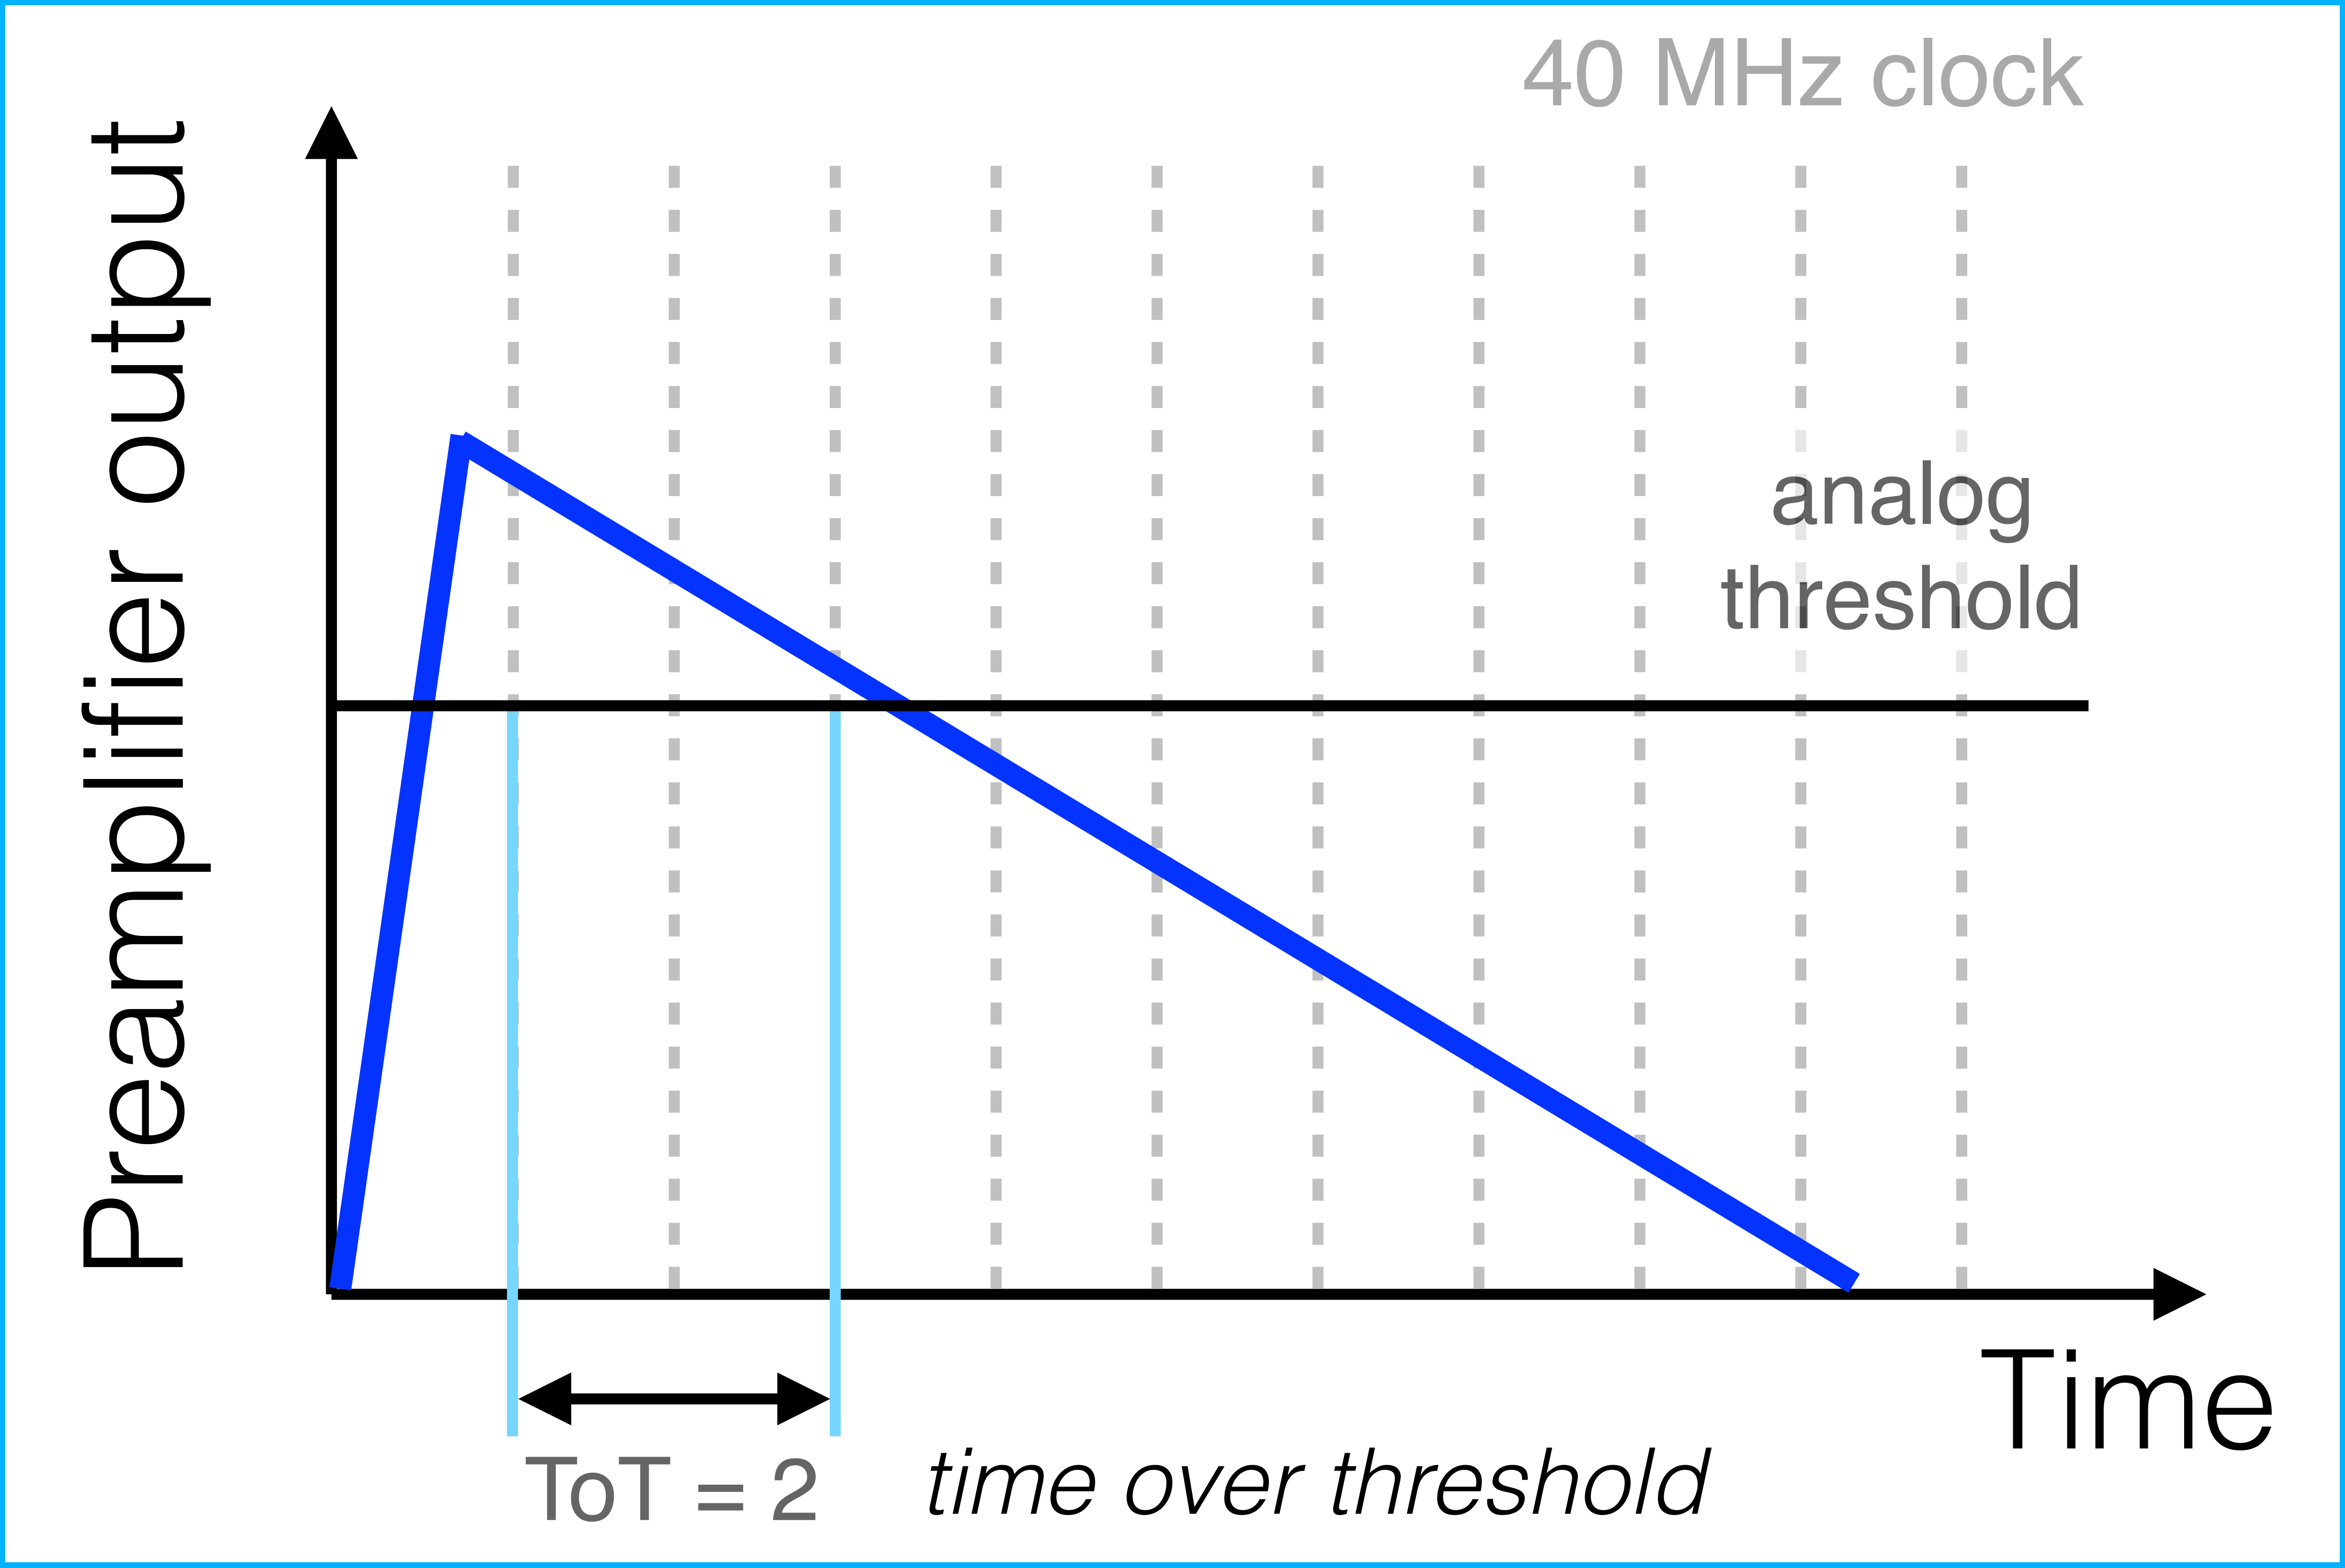

Optimization studies: Use of charge

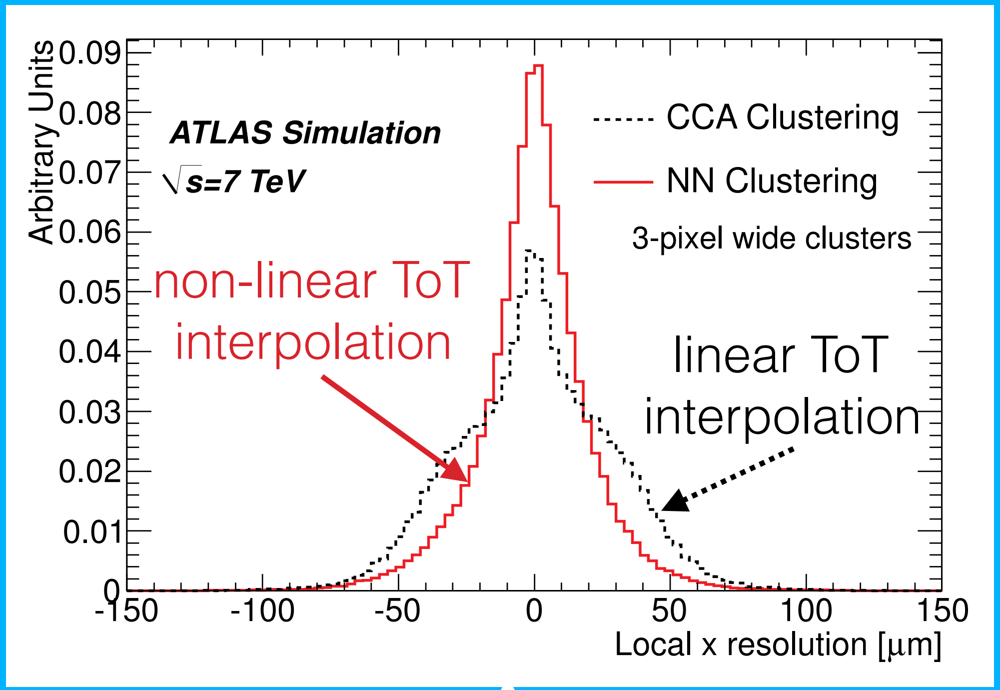

JINST 9 (2014) P09009

Offline, the ToT is used for many purposes; however if not saved properly, performance may suffer!

Position estimation

Particle identification

Particle multiplicity

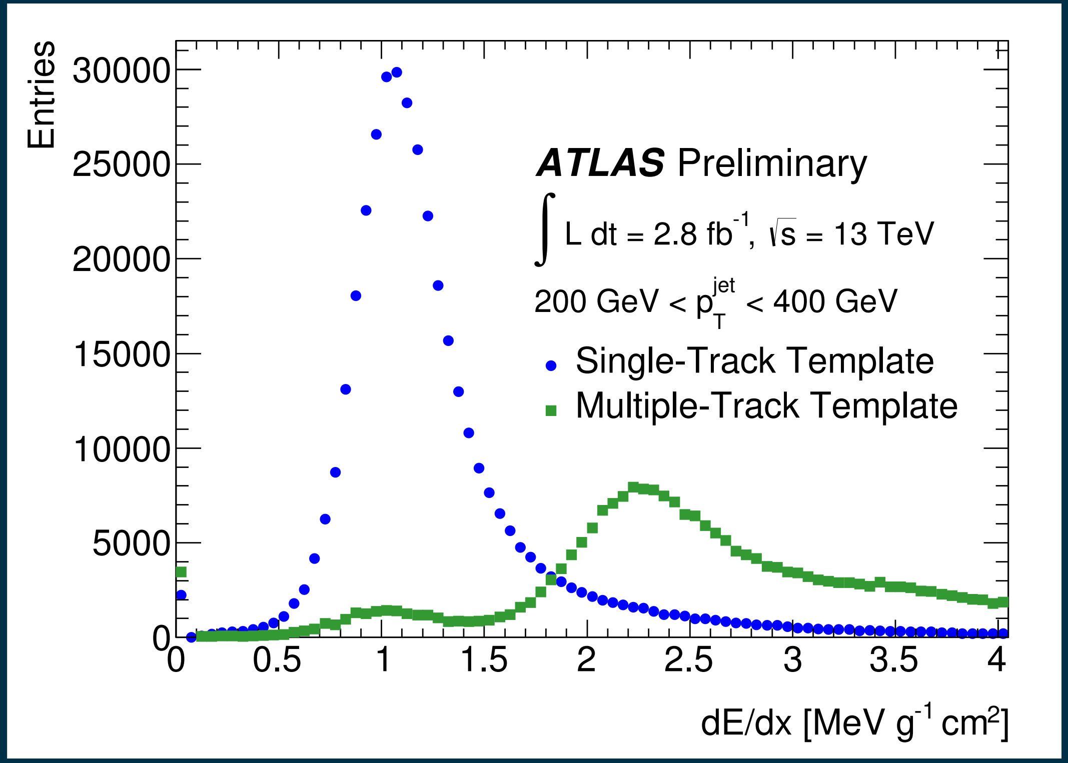

ATL-PHYS-PUB-2016-007

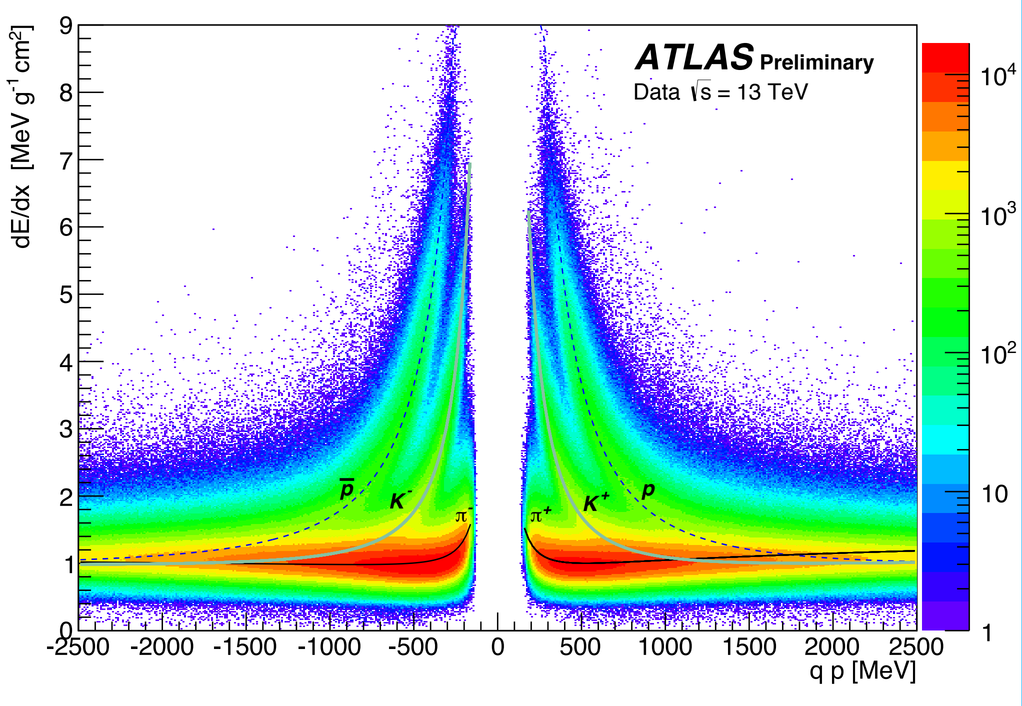

ATLAS-PIX-2015-002

Currently, ATLAS uses 4/8 bit ToT with a linear charge to ToT conversion.

Optimization studies: Use of charge

Offline, the ToT is used for many purposes; however if not saved properly, performance may suffer!

JINST 9 (2014) P09009

Position estimation

Particle identification

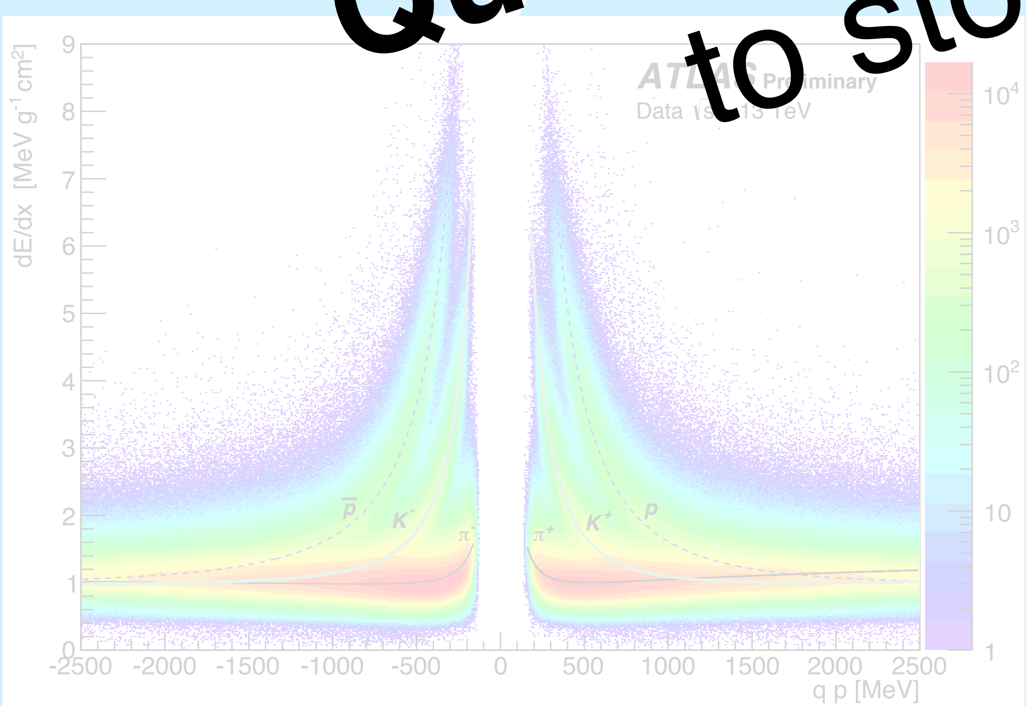

ATLAS-PIX-2015-002

Particle multiplicity

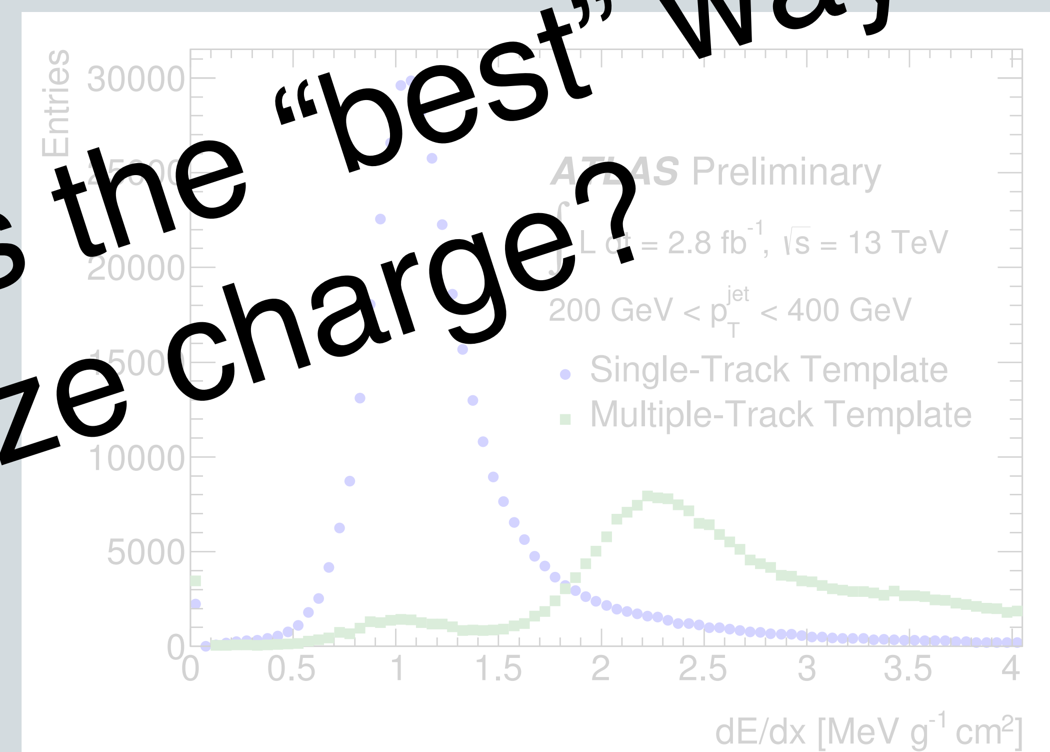

ATL-PHYS-PUB-2016-007

Currently, ATLAS uses 4/8 bit ToT with a linear charge to ToT conversion.

Question: what is the "best" way to store/utilize charge?

Y. Chen, E. Frangipane, M. Garcia-Sciveres, L. Jeanty, **BPN**, S. Pagen Griso, F. Wang, NIMA (2017); arXiv:1710.02592

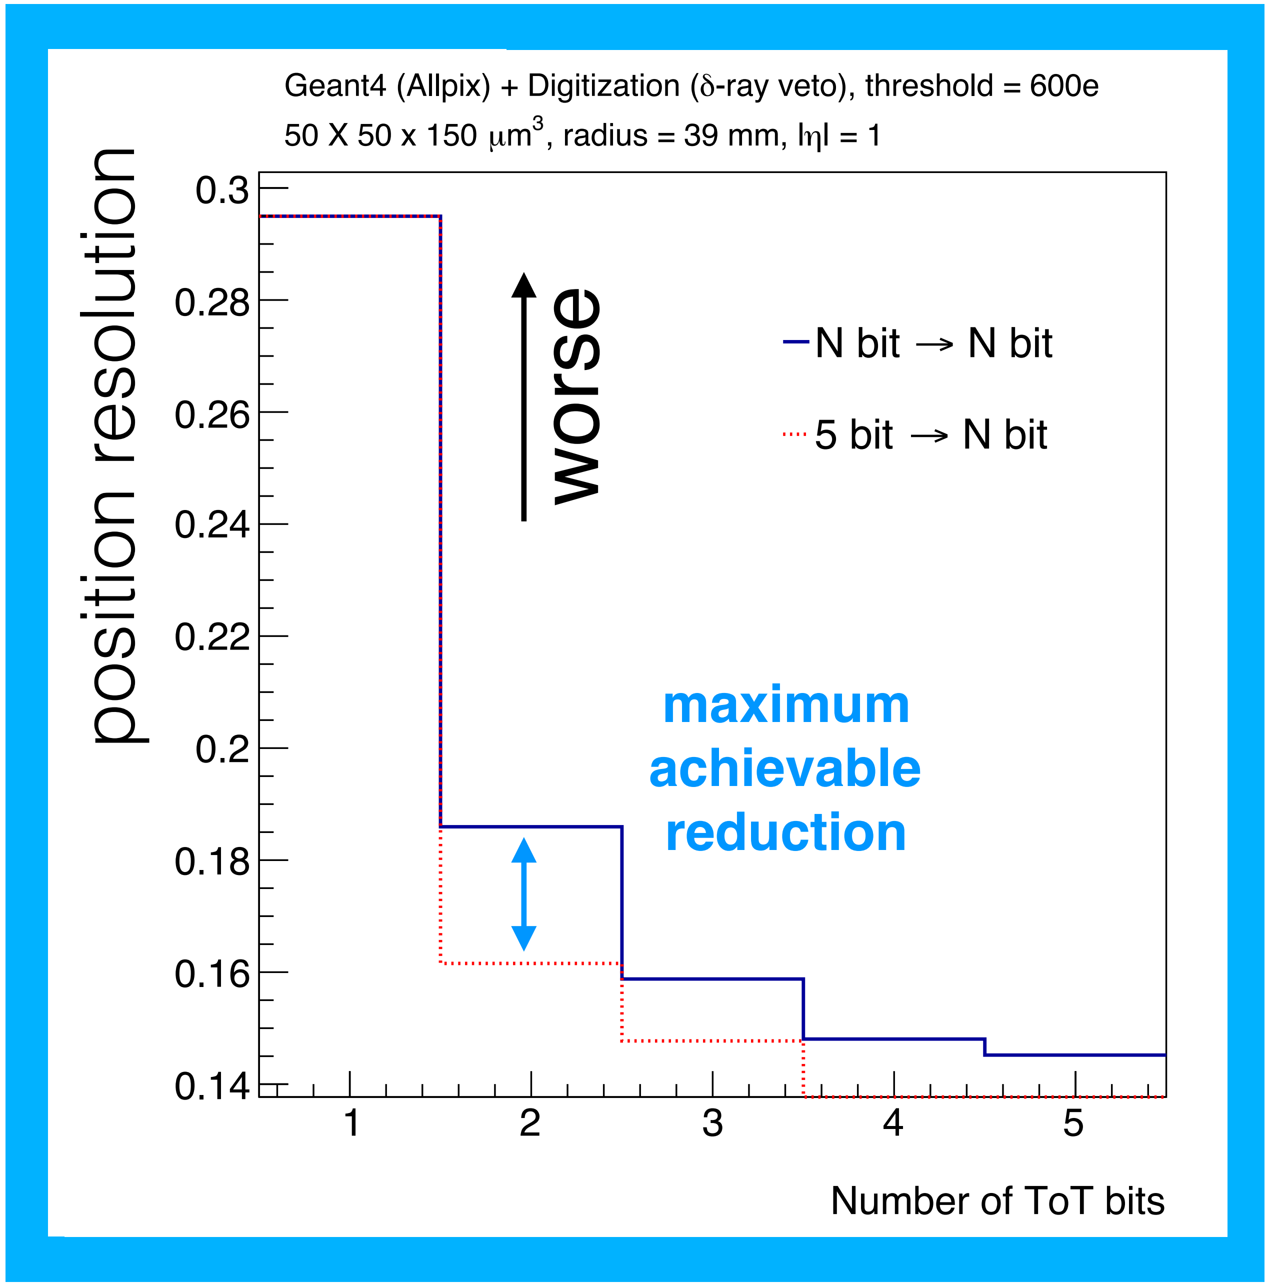

Can add digital logic so that N digitized bits are stored as M ≤ N bits.

There are $\binom{2^N - 2}{2^M - 2}$ possible functions mapping N to M bits.

We could gain space without much loss in performance

+ many more considerations for readout design!

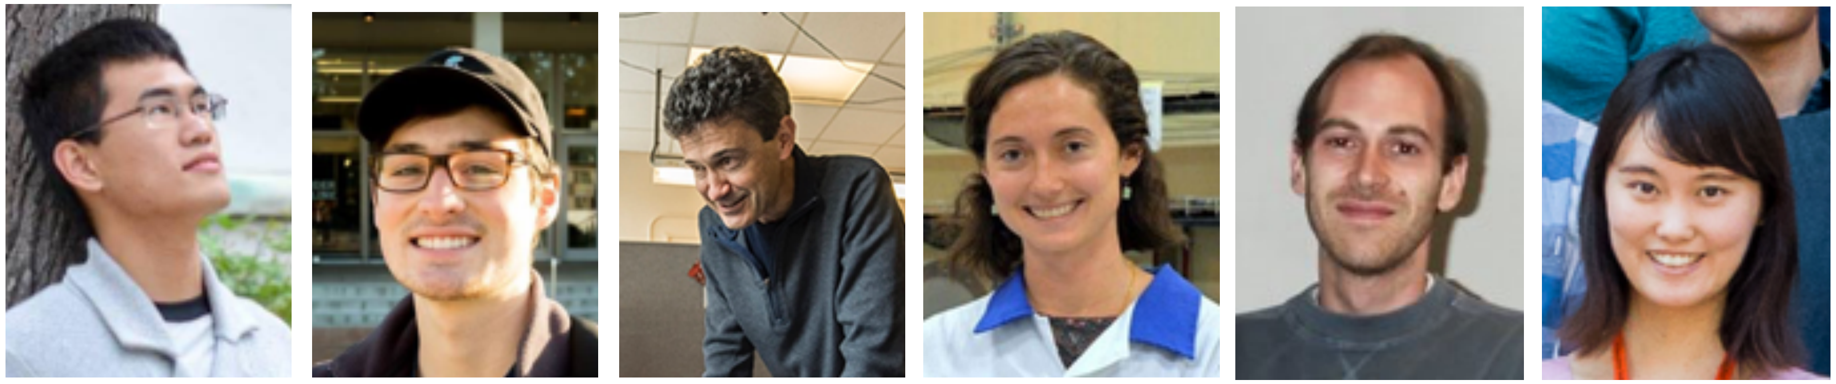

Data pipeline at the LHC

data / simulation

Time (data / simulation)	Stage
25 ns	O(100) *pp* collisions
10^{-19}-10^{-15} s / ~ms	(sub-)nuclear physics
0.01-20 ns / ~min	out-going particles interact with detector
1-100 ns / ~100 ms	detector response (signal formation + digitization)
2.5 µs	**hardware-based trigger decision**
200 ms	software-based trigger decision
~100 ms	event reconstruction
	event processing (skim, thin, augment)
~months	final data analysis (uses millions of events)

Cross-section is dominated by final states with no electrons, no muons, and no neutrinos

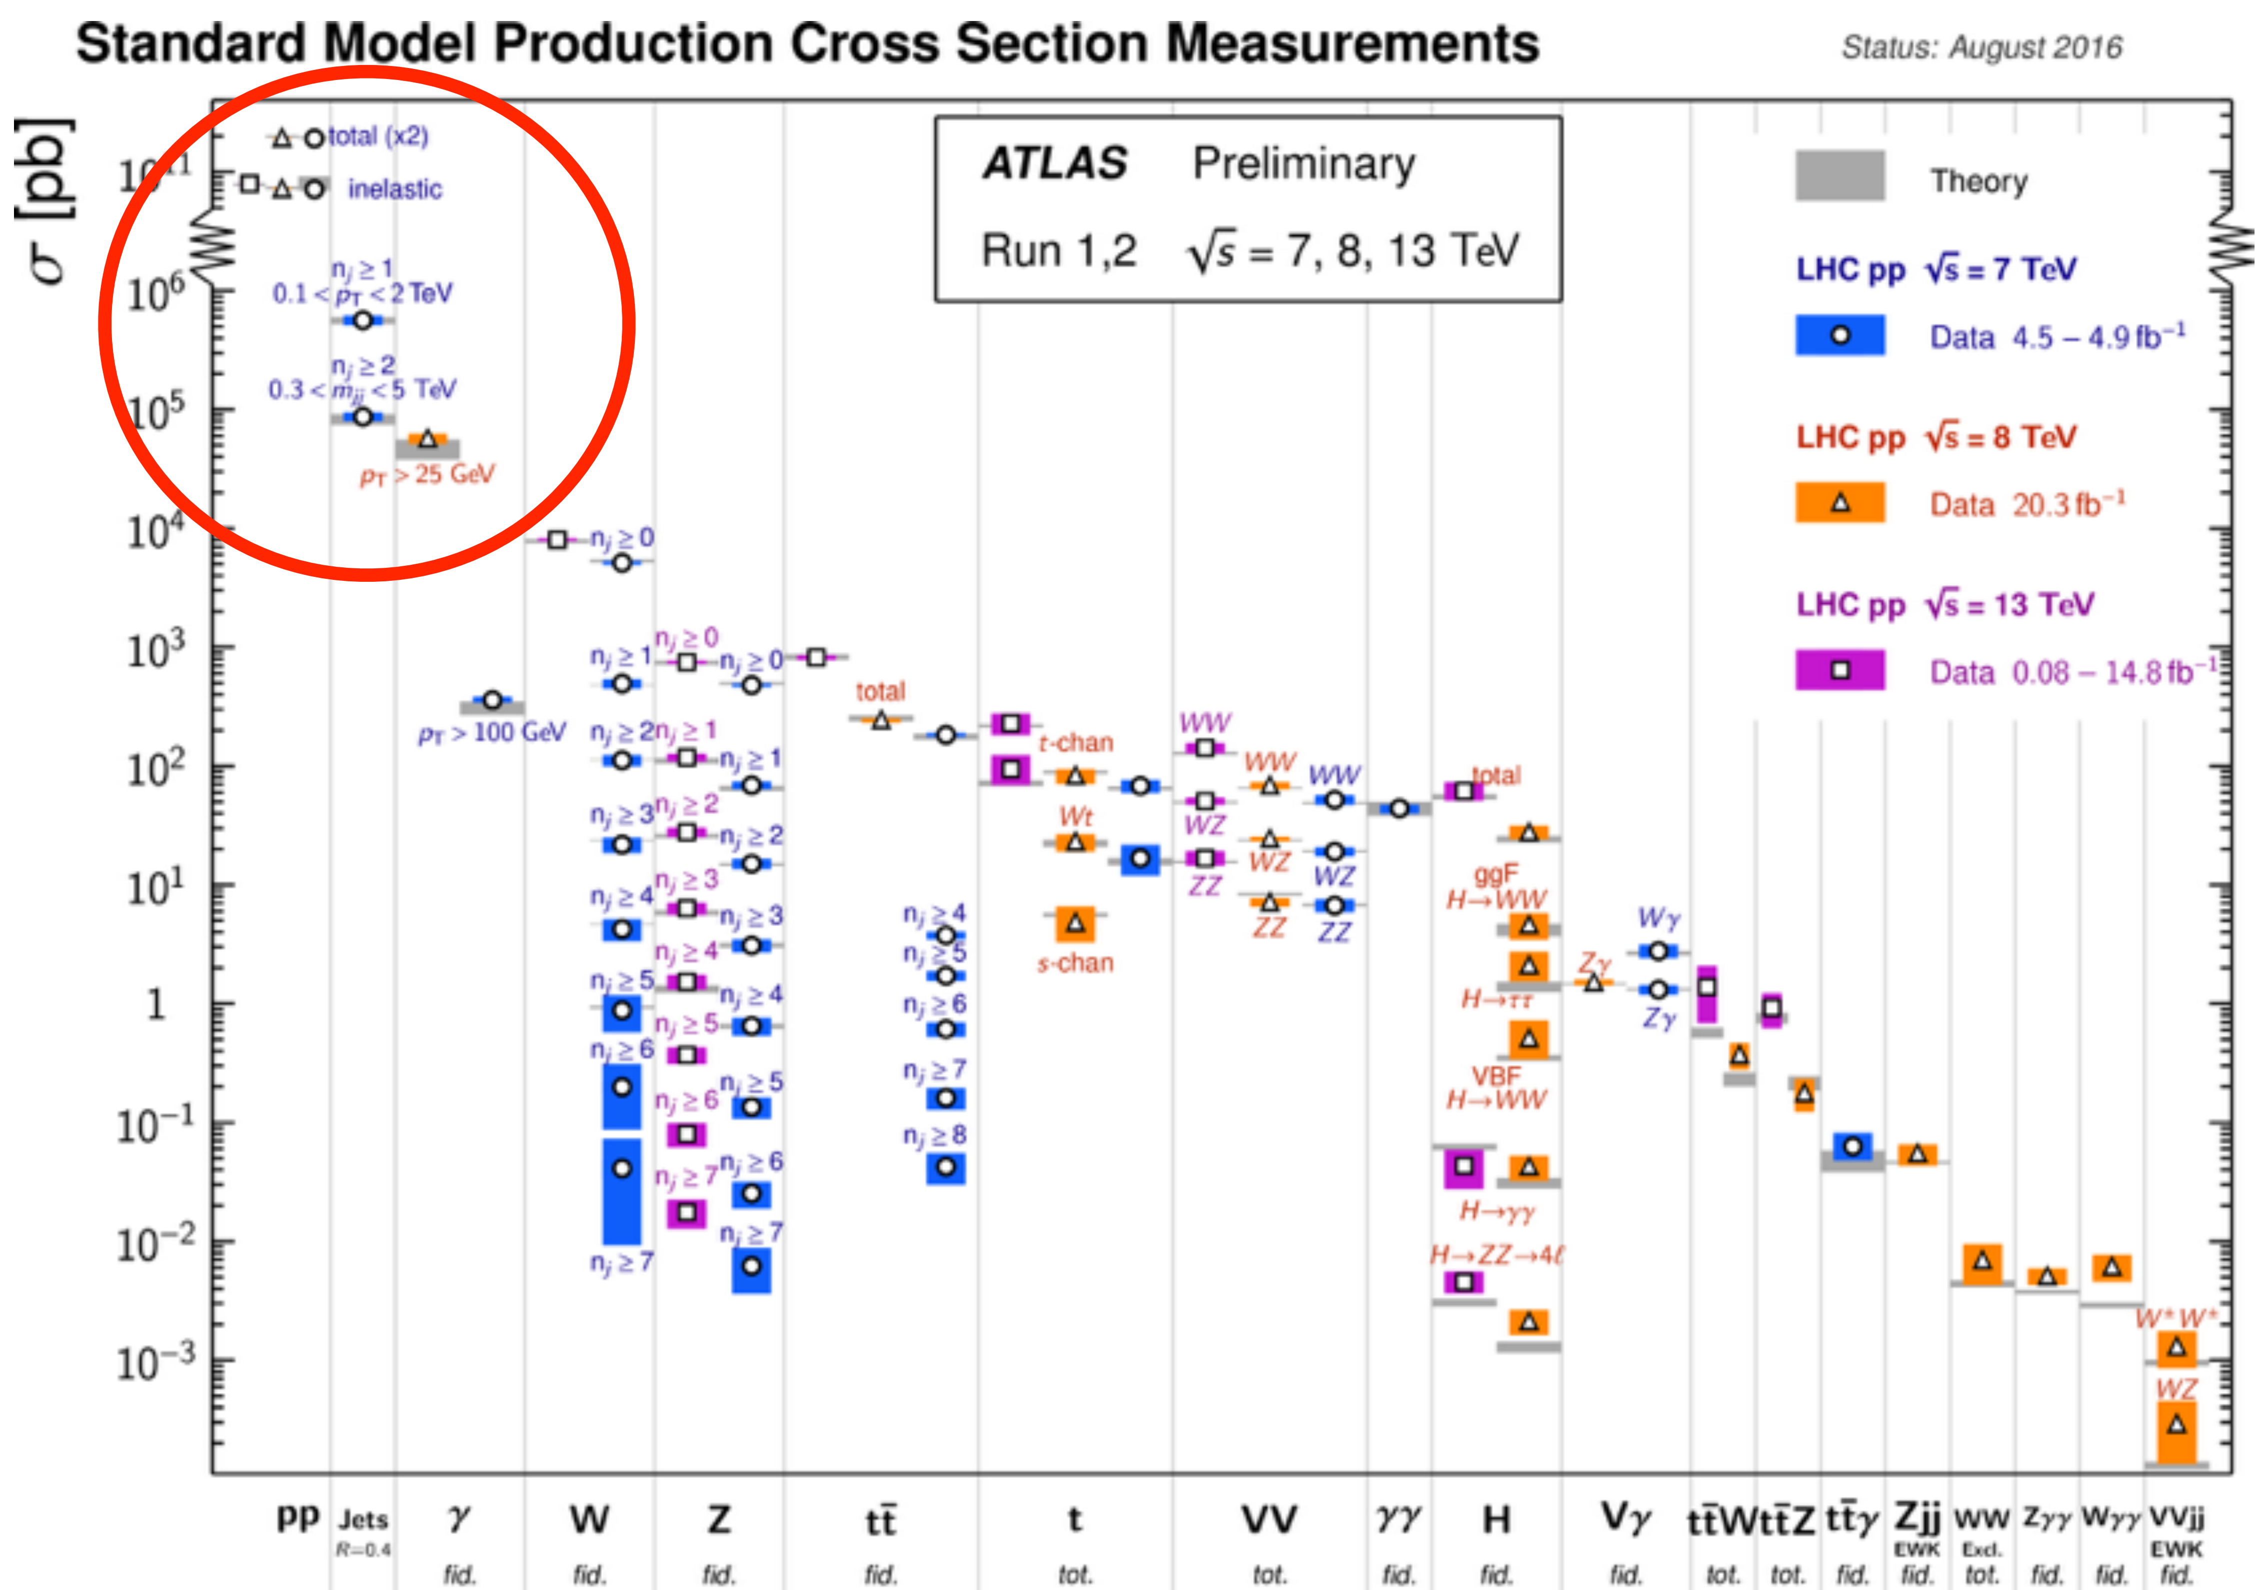

Recorded events are dominated by final states with electrons, muons, neutrinos

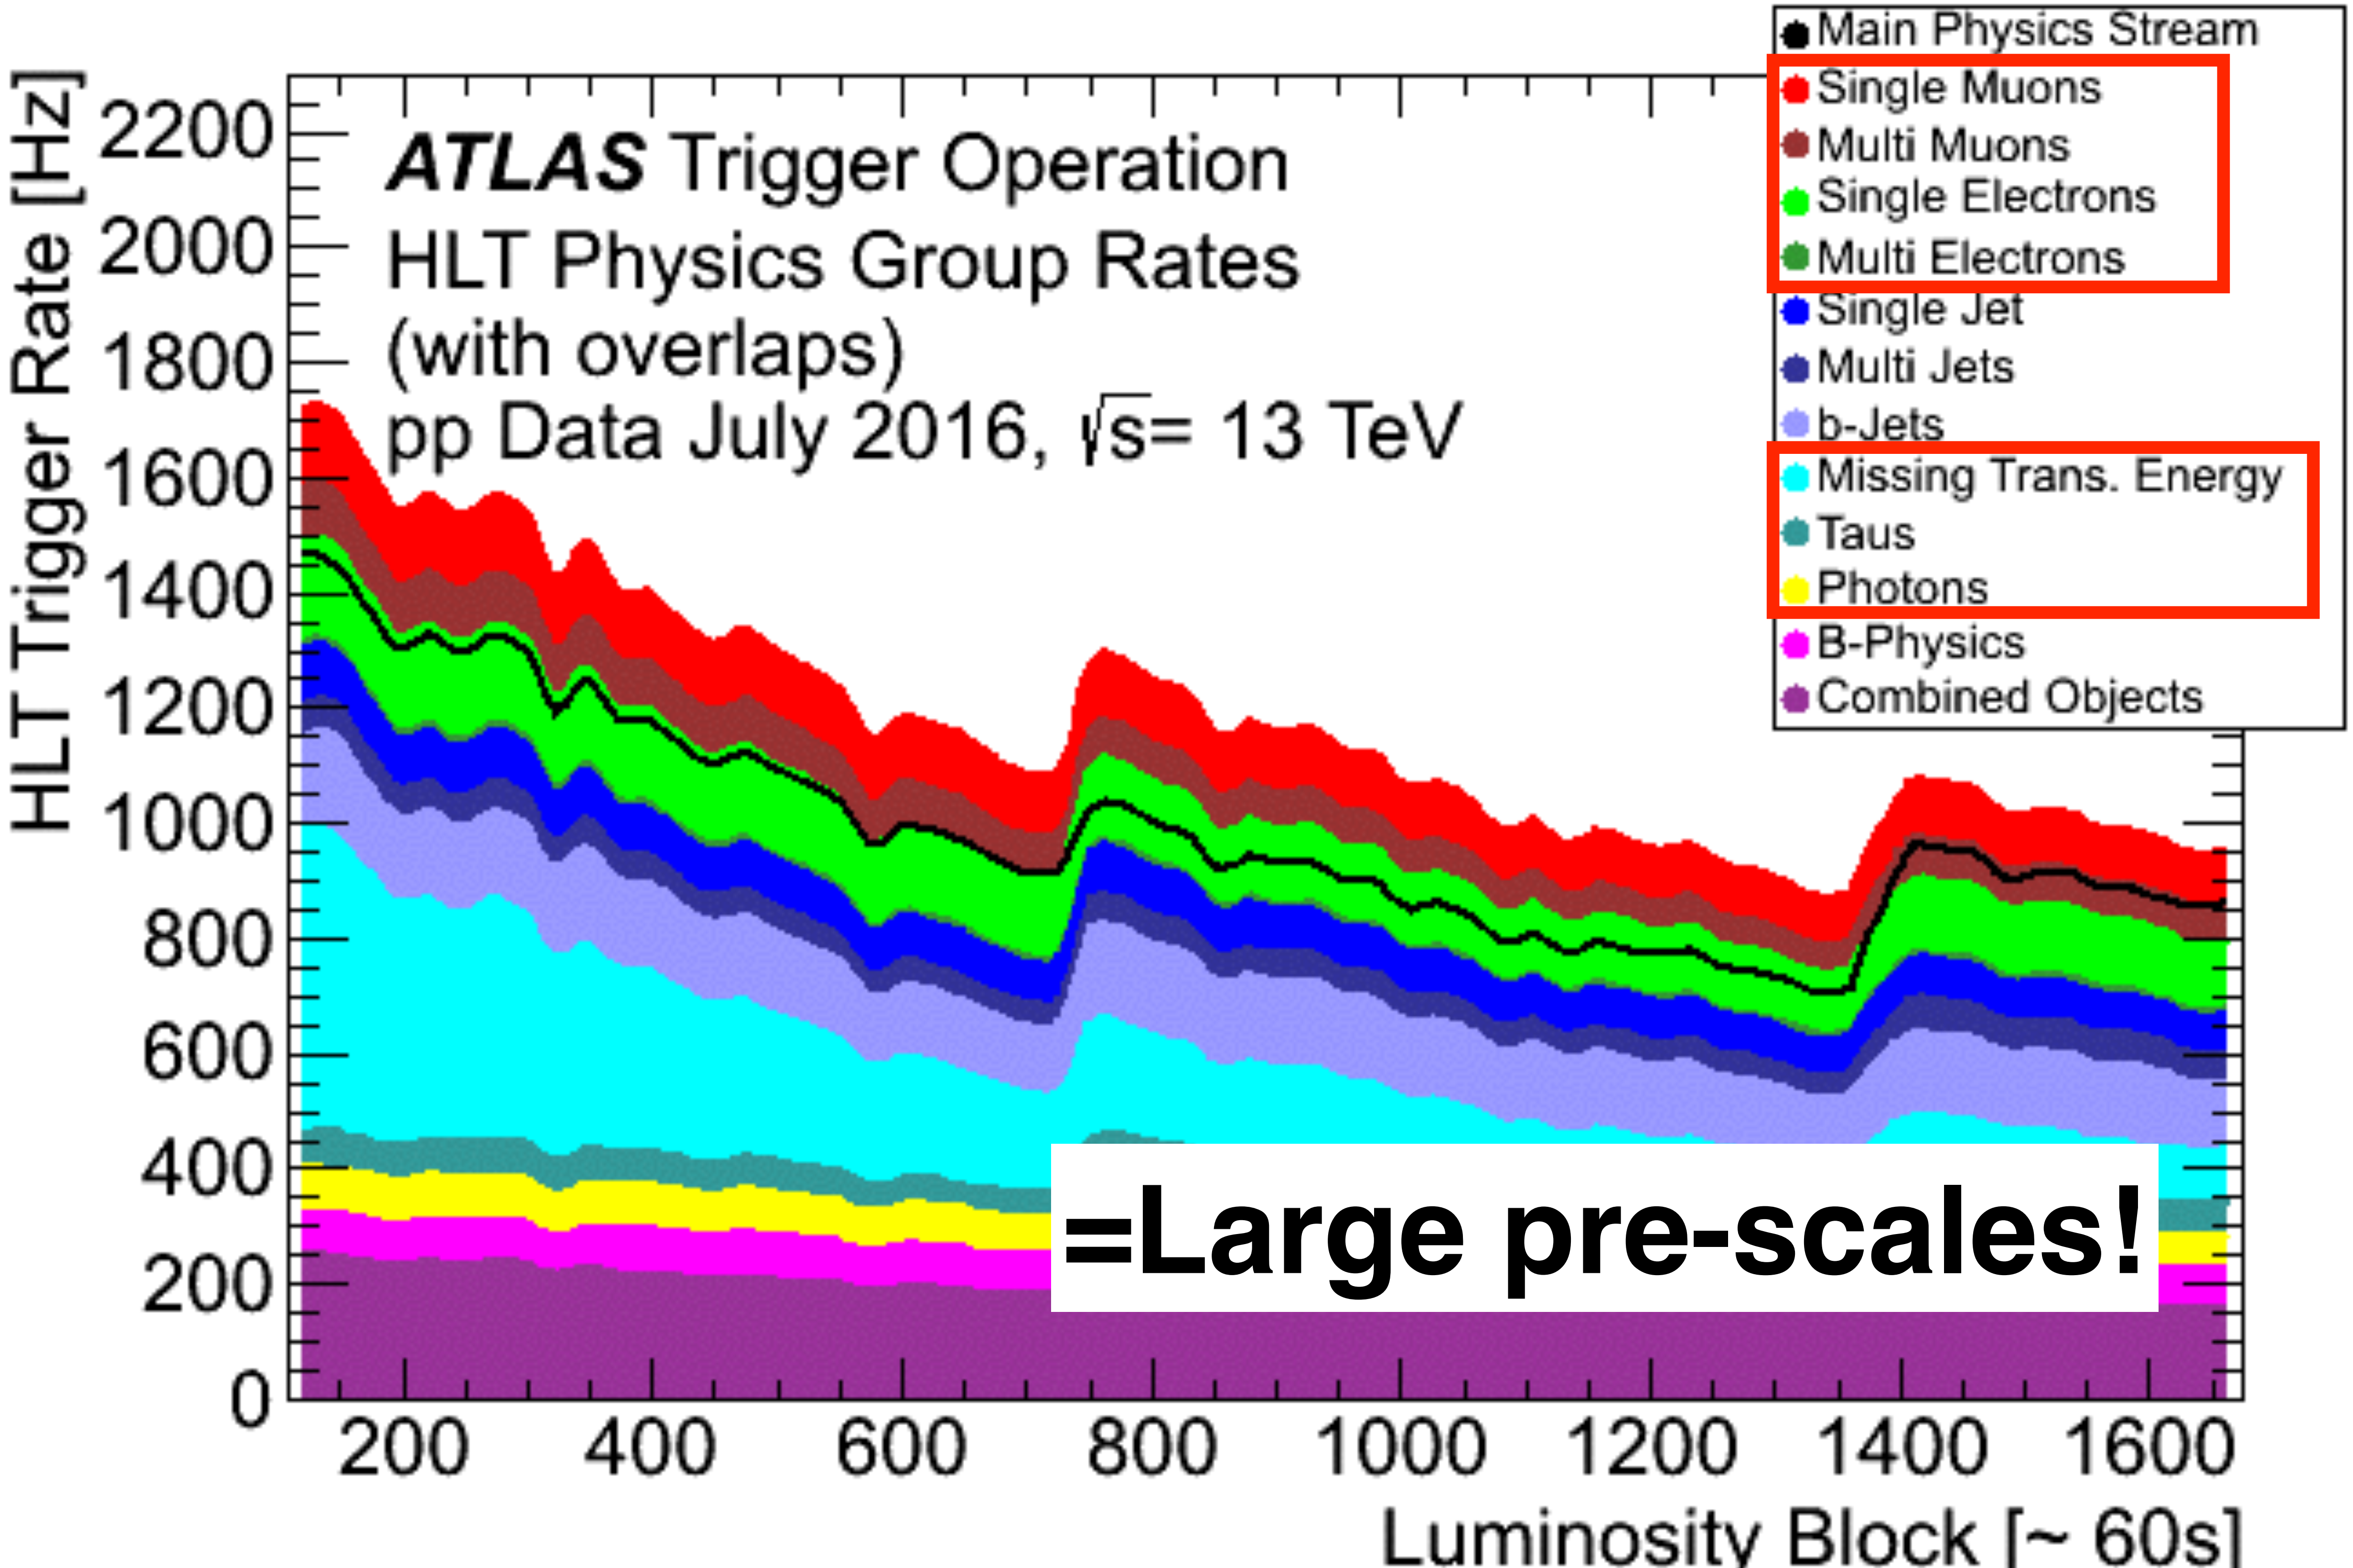

HLT = software trigger (L1 = hardware trigger)

usual paradigm

If your favorite process cannot be triggered on, then it will ***not be recorded*** *and there is* ***no way to analyze it****.*

new paradigms

"If your favorite process cannot be triggered on inclusively, look for ***associated production*** *with an object you can trigger."*

-I.S. Radiation, 2015

"If your favorite process cannot be triggered on at HLT, make your analysis ***faster and simpler*** *and do it after L1."*

-T.L. Analysis, 2012

(see Caterina's talk)

usual paradigm

If your favorite process cannot be triggered on, then it will ***not be recorded*** *and there is* ***no way to analyze it****.*

new paradigms

“If your favorite process cann… …ely, look for associat… …you can trigger.”

Caveat: **large effective prescale** from the low cross-section from the associated production

-I.S. Radiation, 2015

“If your favorite process cannot be trigger… …make your analysis faster an… …do it after L1.”

Caveat: limited by the L1 rate (usually **harsher than HLT**)

-T.L. Analysis, 2012

ATLAS EXPERIMENT

Run Number: 266904, Event Number: 25884352

Date: 2015-06-03 13:41:54 CEST

trigger on pink

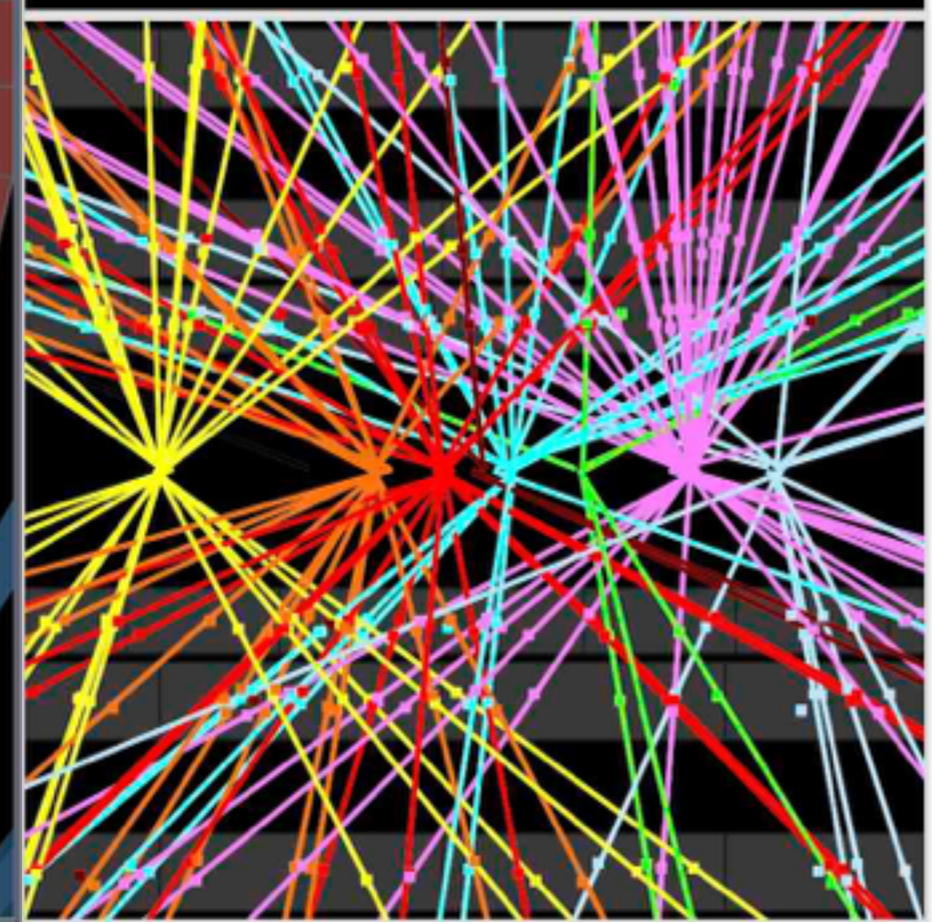

your favorite process could be in any of these ones

L1 (TLA) or HLT (offline) rate

$$\left(\begin{array}{c}\text{“Zero bias} \\ \text{trigger” data}\end{array}\right) = \frac{h}{H} \times \left(\begin{array}{c}\text{Triggered} \\ \text{data amount}\end{array}\right)$$

40 MHz

# *pp* interactions	L1 Rate	HLT Rate	ZBT	ZBT @HLT
20	100 kHz	100 Hz	4×10^5	400
80 (~now)	100 kHz	1 kHz	4×10^4	400
200	400 kHz	10 kHz	4×10^3	100

How many events?

If you run a zero-background search and can't beat a trigger efficiency of ~0.02%, then you should be using the ZBT!

trigger" data H data amount

…and if you can do TLA, that number goes up to ~1%!

# *pp* interactions	L1 Rate	HLT Rate	ZBT	ZBT @HLT
20	100 kHz	100 Hz	4 x 10^5	400
80 (~now)	100 kHz	1 kHz	4 x 10^4	400
200	400 kHz	10 kHz	4 x 10^3	100

Associated photon production

Associated quark/gluon production

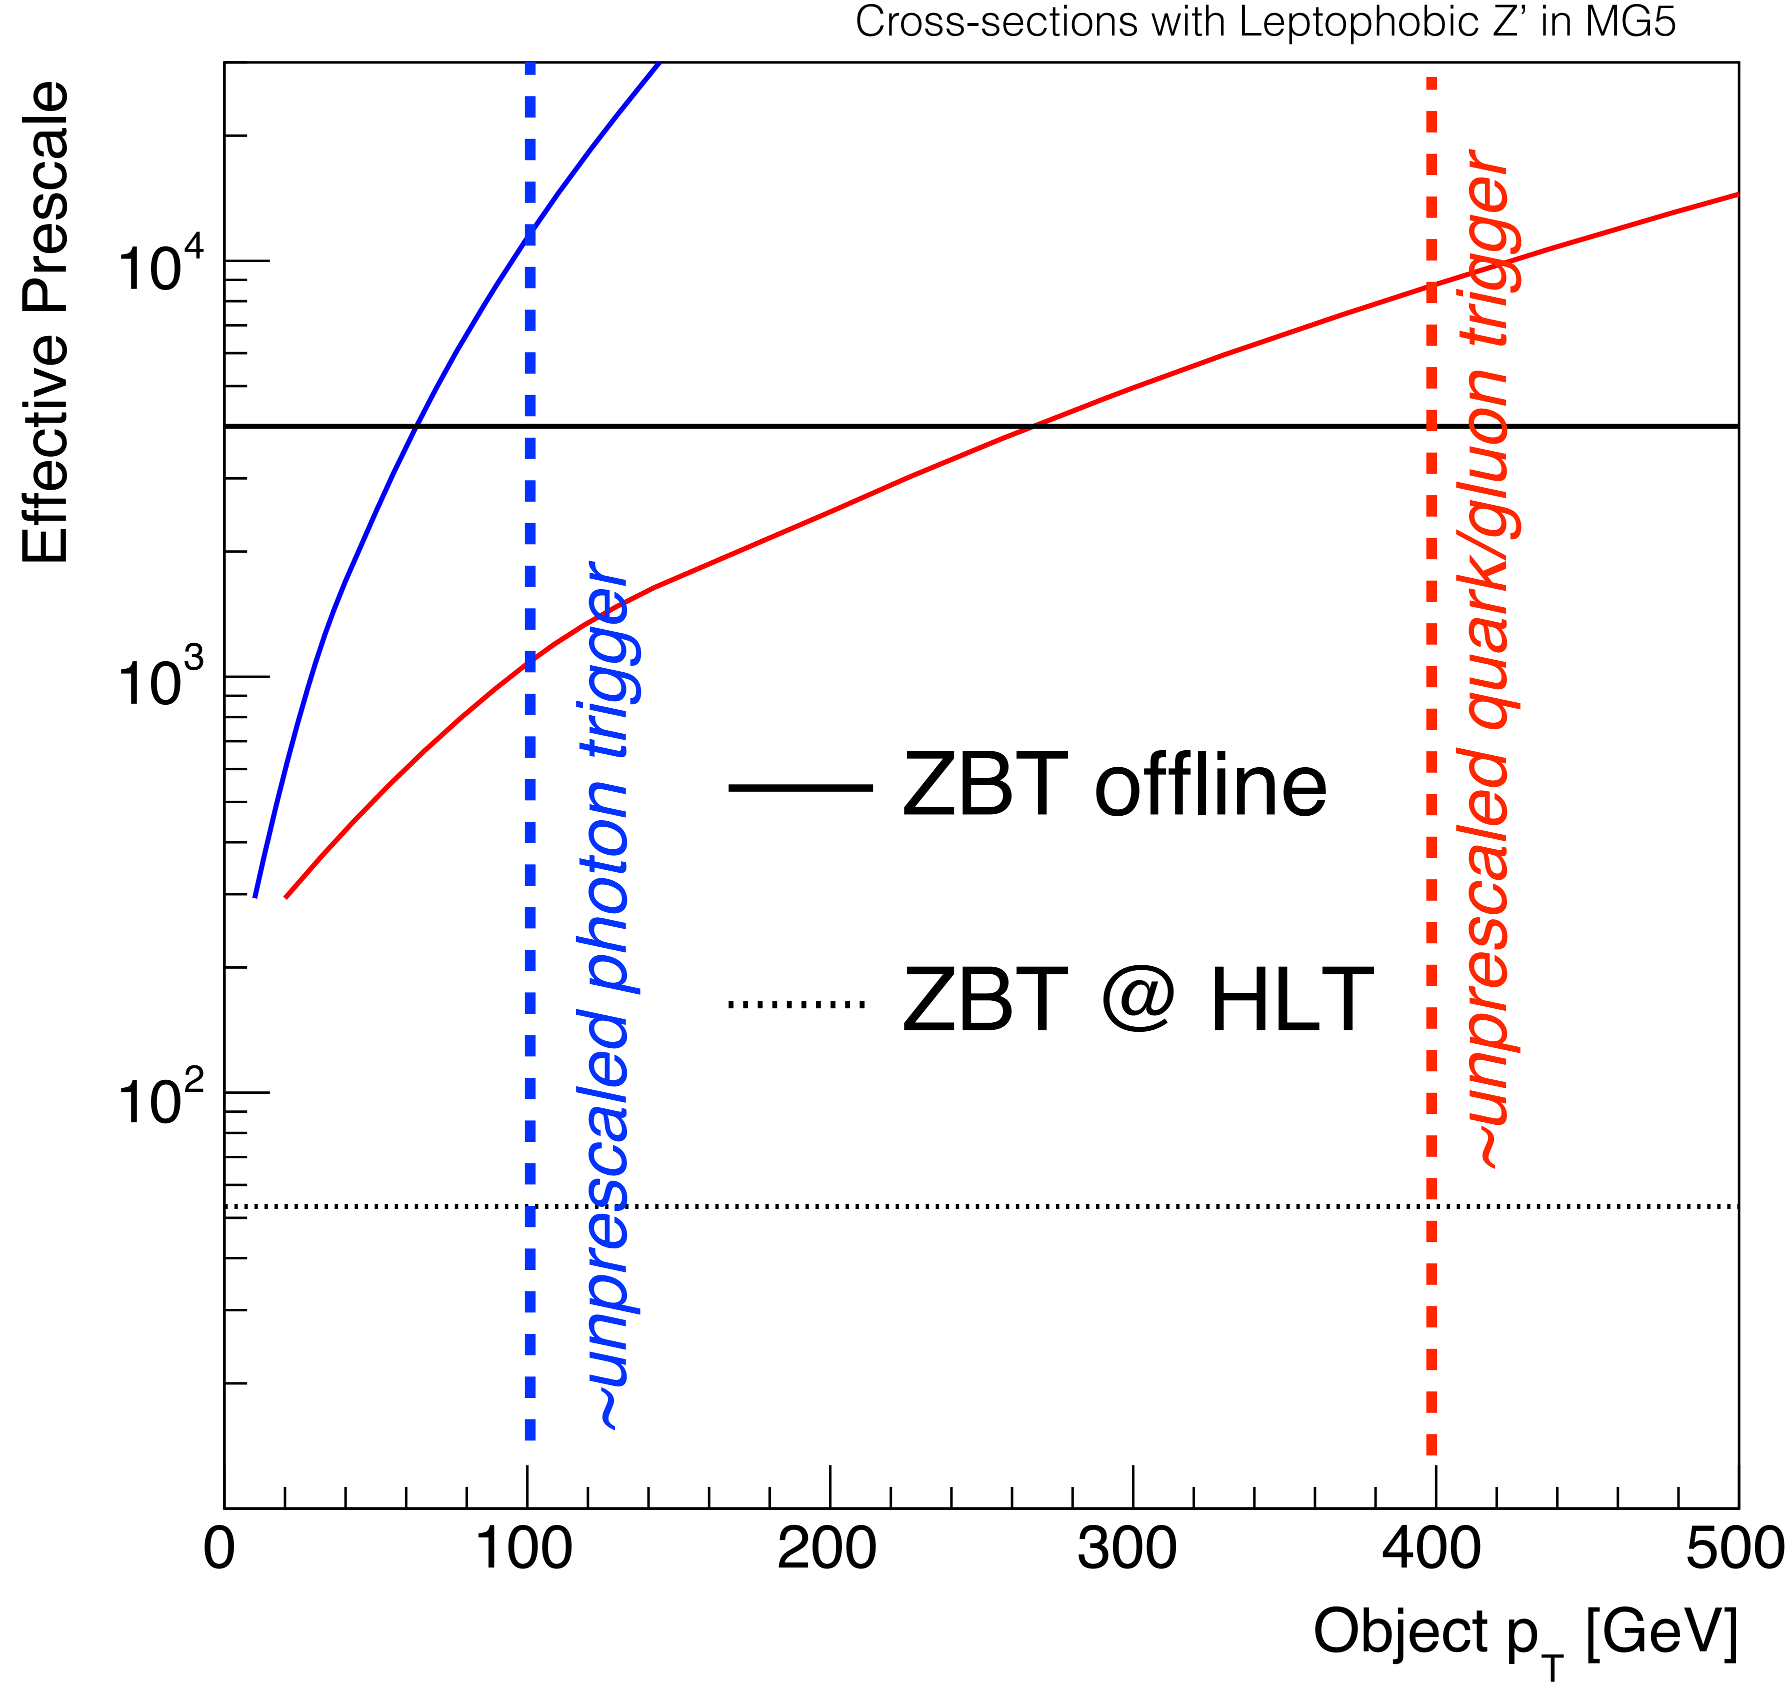

There is a huge dataset that we are currently ignoring.

New physics may be hiding in these data
and we are collecting them anyway

Most powerful when combined with trigger-level analysis (**so need to design ASAP**!)

Takeaway message: the baseline is the ZBT >> 0!

Think creatively about new possibilities…
the sky, and not the trigger, is the limit!

(also, remember ZBT offline has ~infinite time for processing)

Data pipeline at the LHC

data / simulation

25 ns	O(100) *pp* collisions
10^{-19}-10^{-15} s	(sub-)nuclear physics
~ms	
0.01-20 ns	out-going particles interact with detector
~min	
1-100 ns	detector response (signal formation + digitization)
~100 ms	
2.5 µs	hardware-based trigger decision
200 ms	software-based trigger decision
~100 ms	event reconstruction
	event processing (skim, thin, augment)
~months	final data analysis (uses millions of events)

The extra pileup collisions add unwanted soft radiation on top of the event

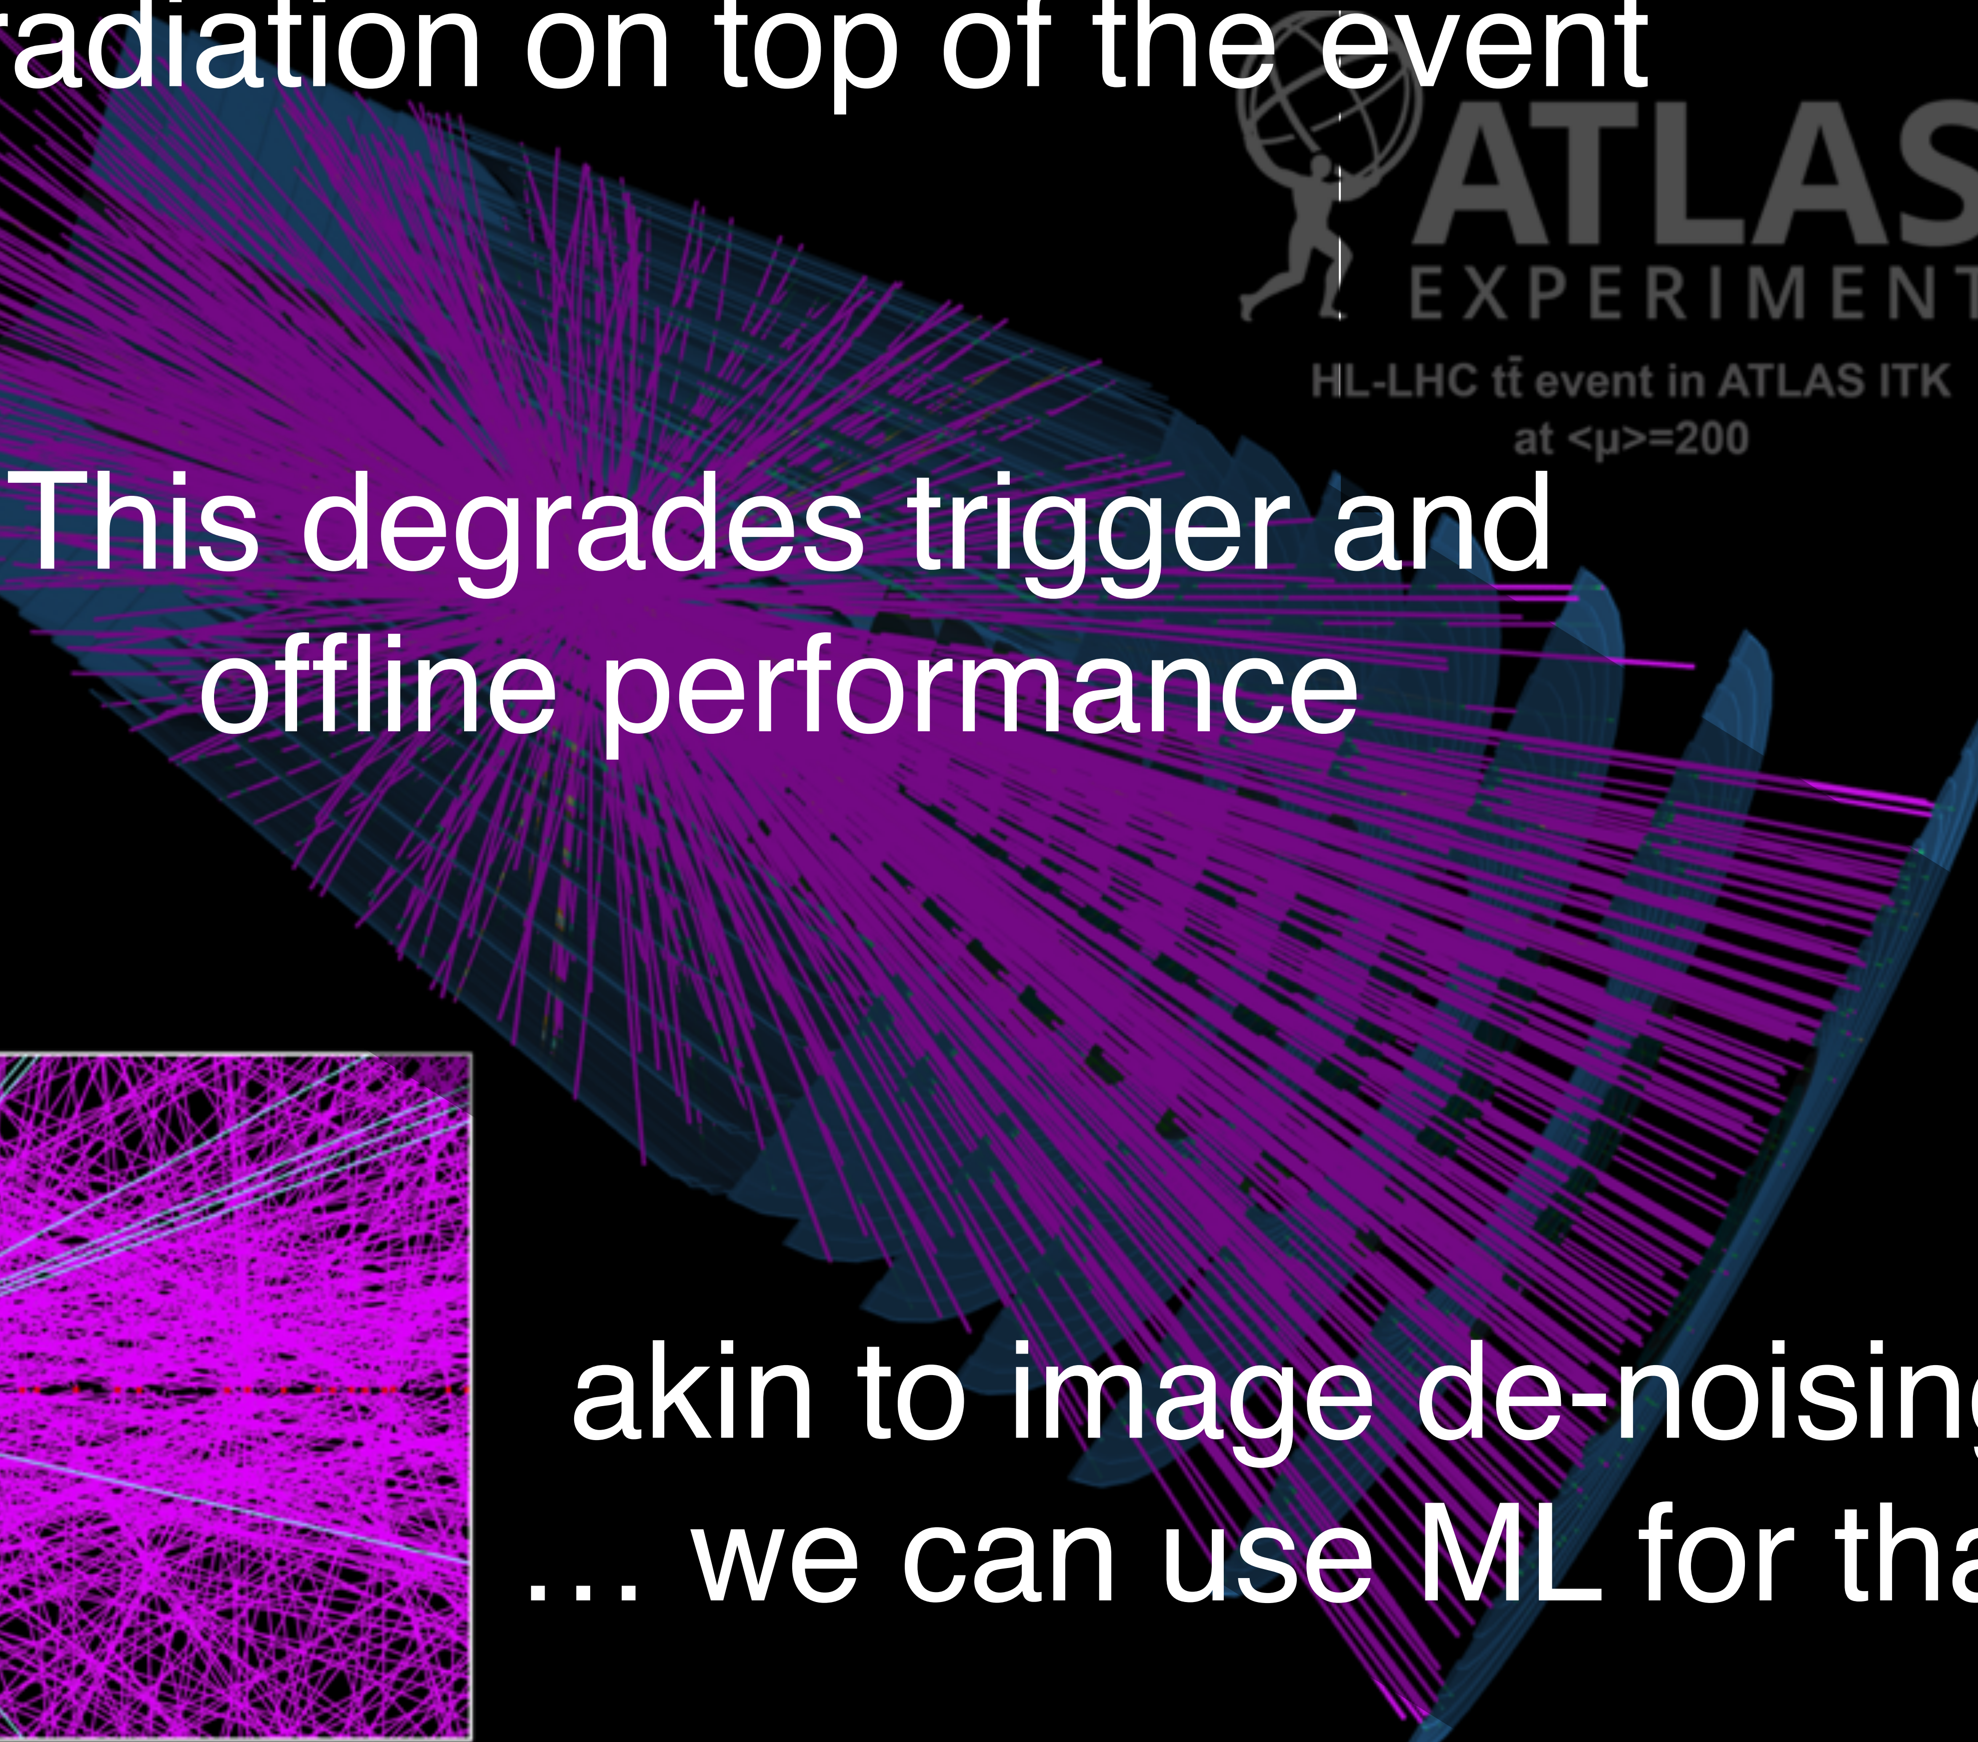

This degrades trigger and offline performance

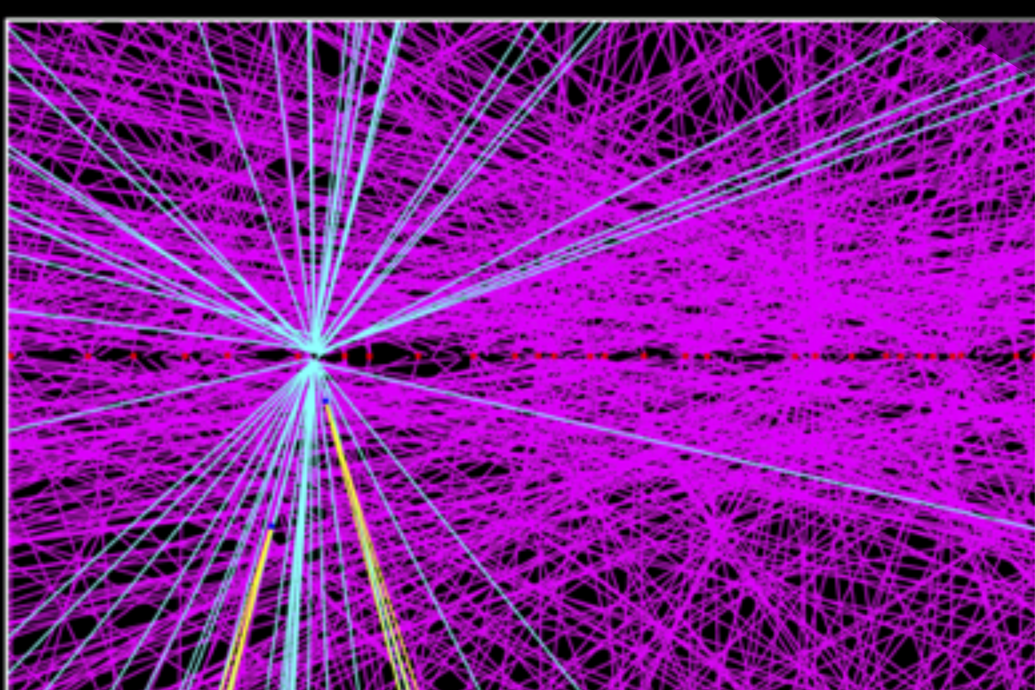

akin to image de-noising

… we can use ML for that!

Pileup mitigation with machine learning

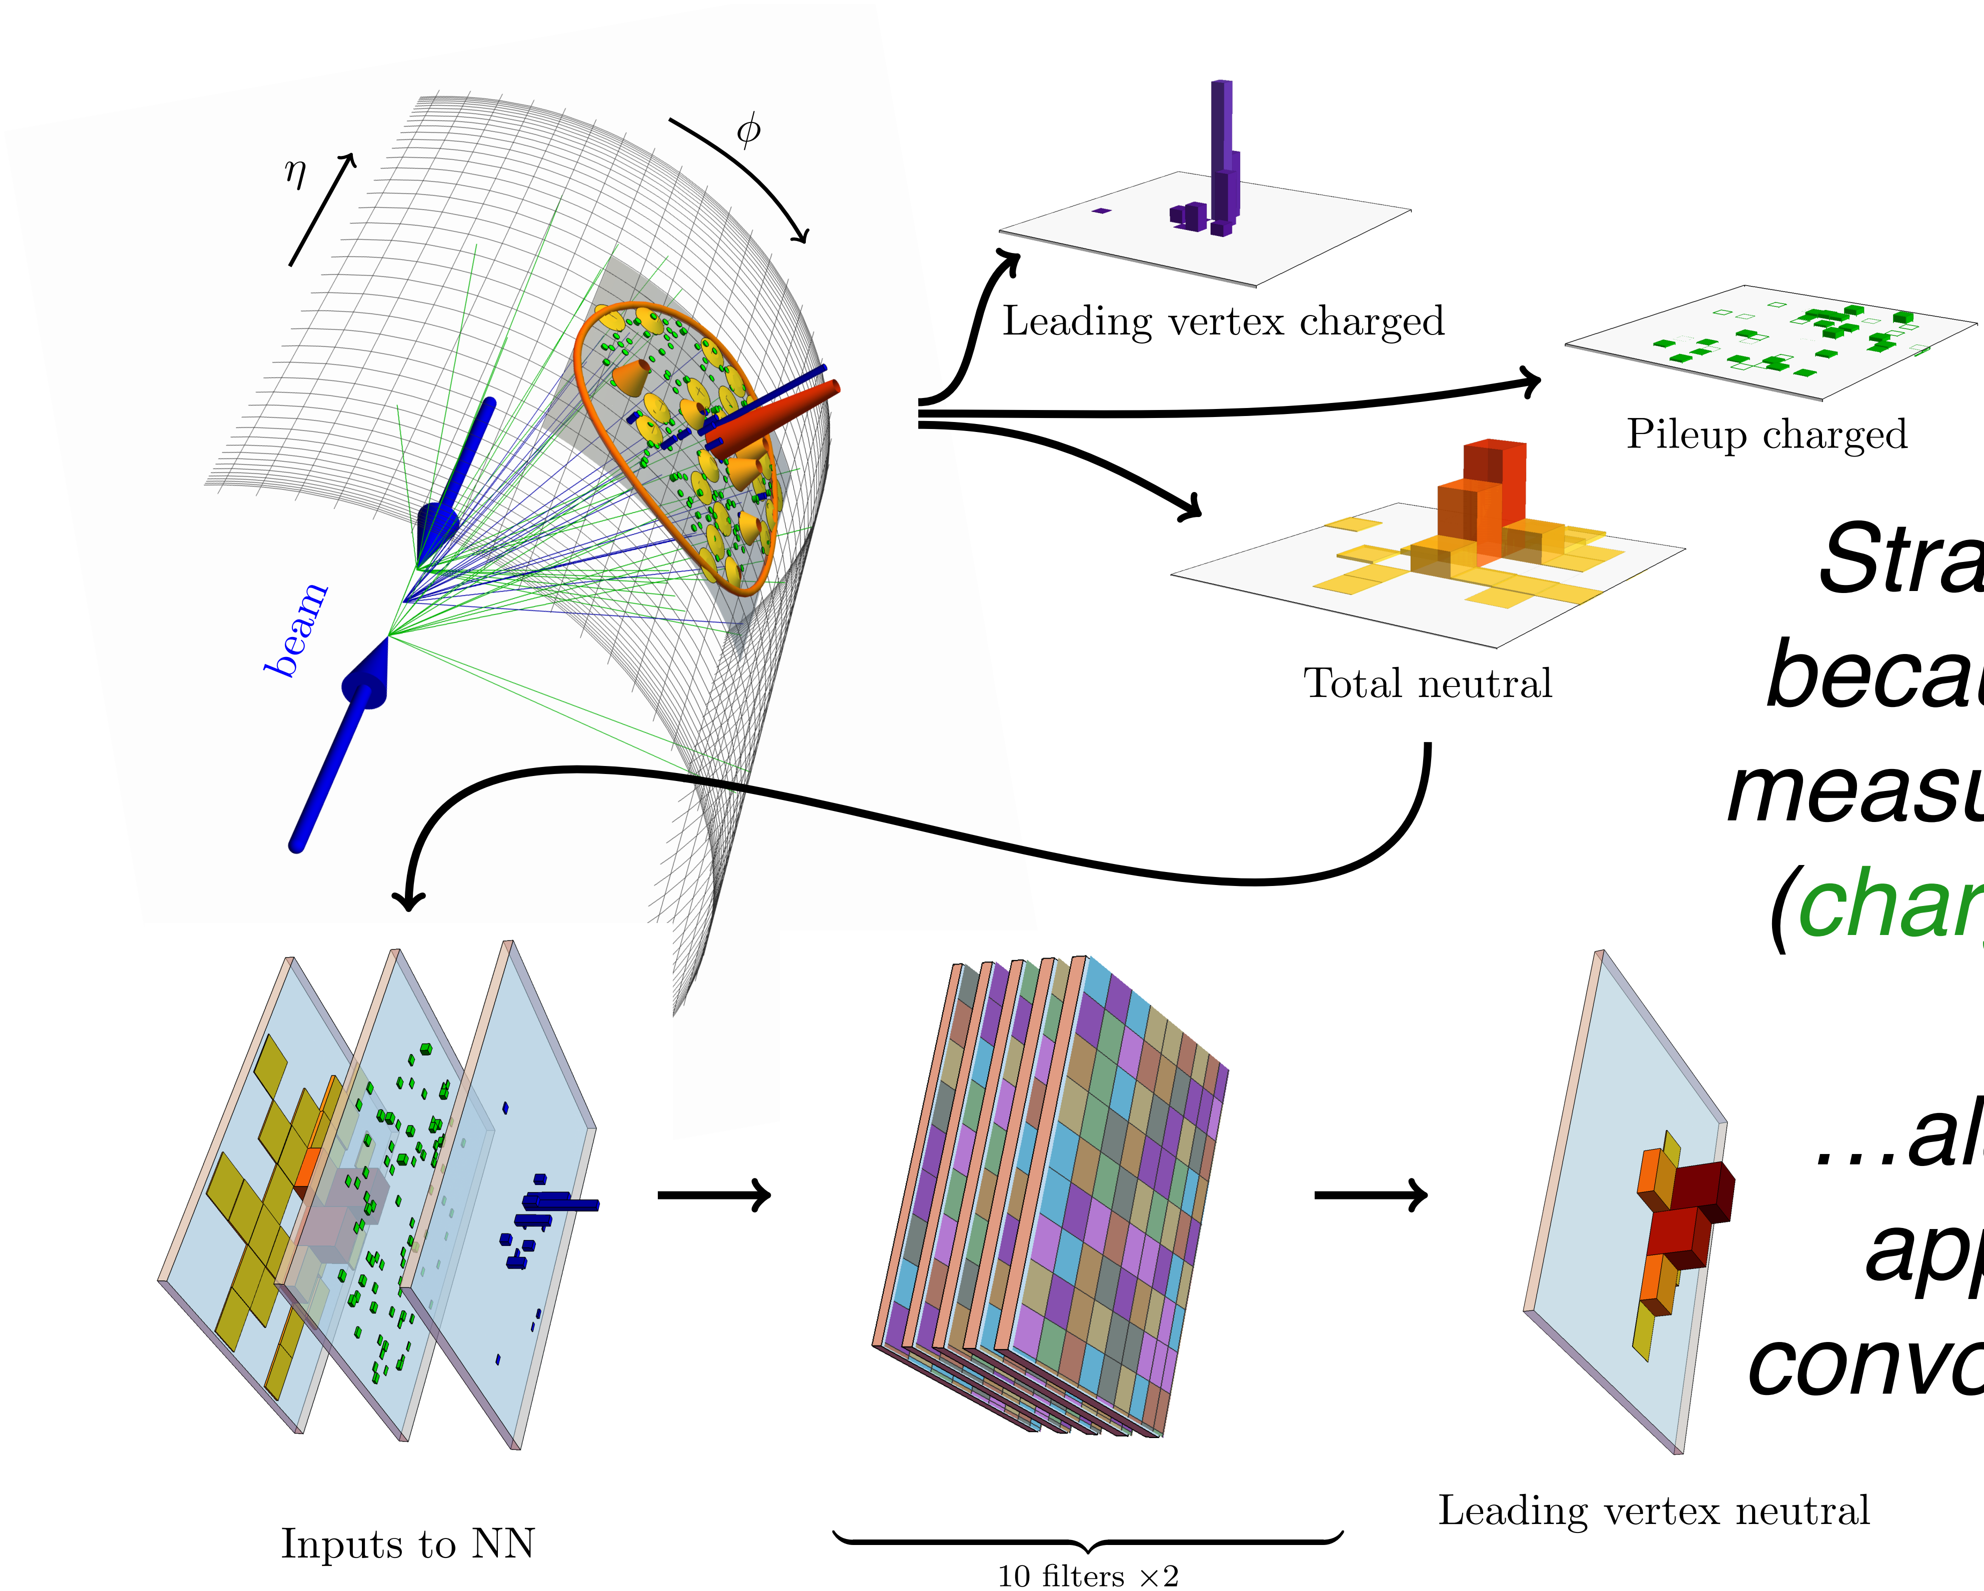

Strange noise because we can measure ~2/3 of it (charged pileup)

…also a natural application of convolutional NNs!

Pileup mitigation with machine learning

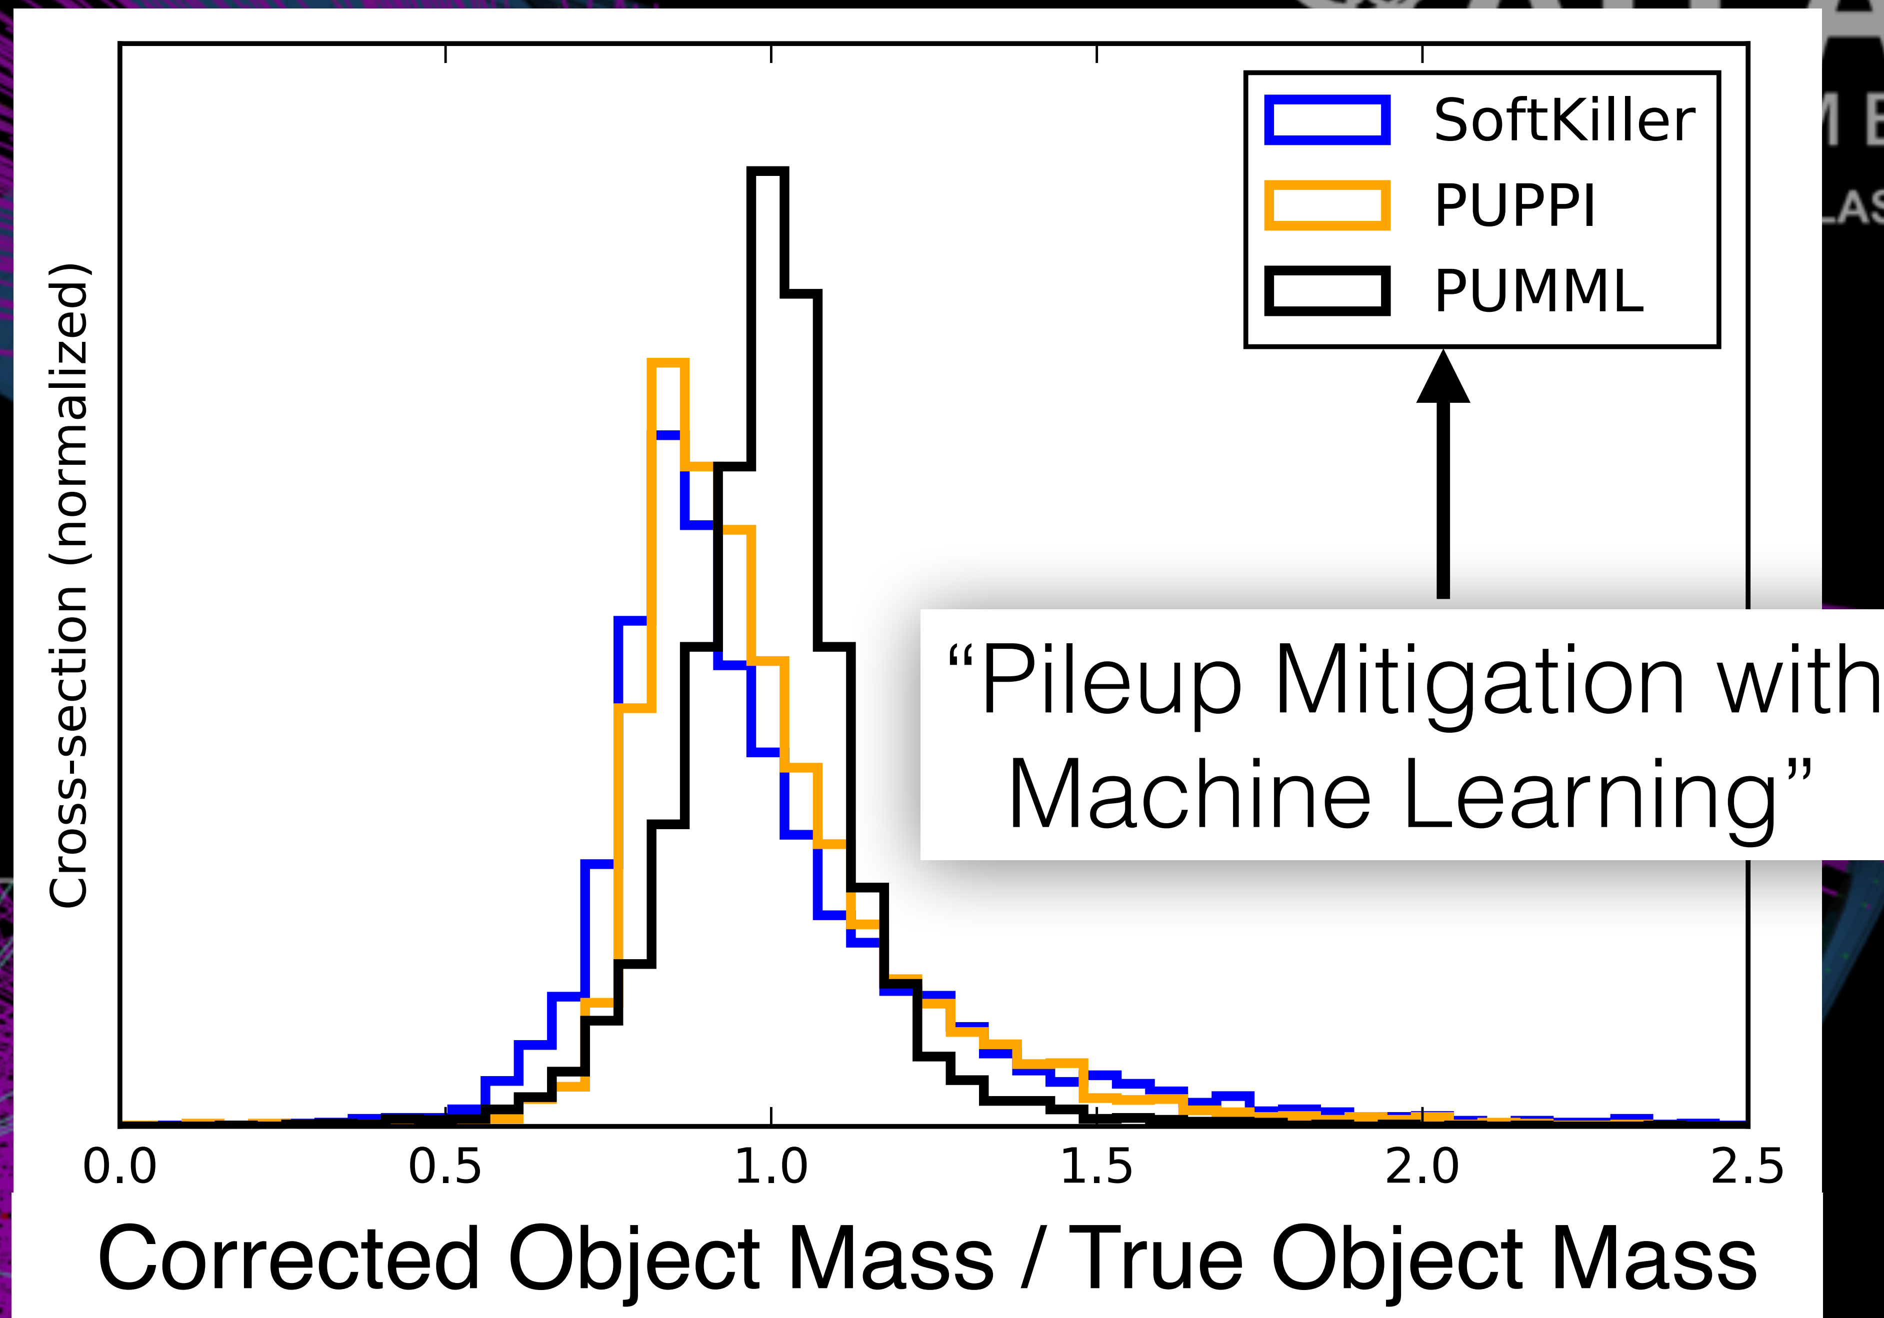

P. Komiske, E. Metodiev, **BPN**, and M. Schwartz, JHEP 12 (2017) 051

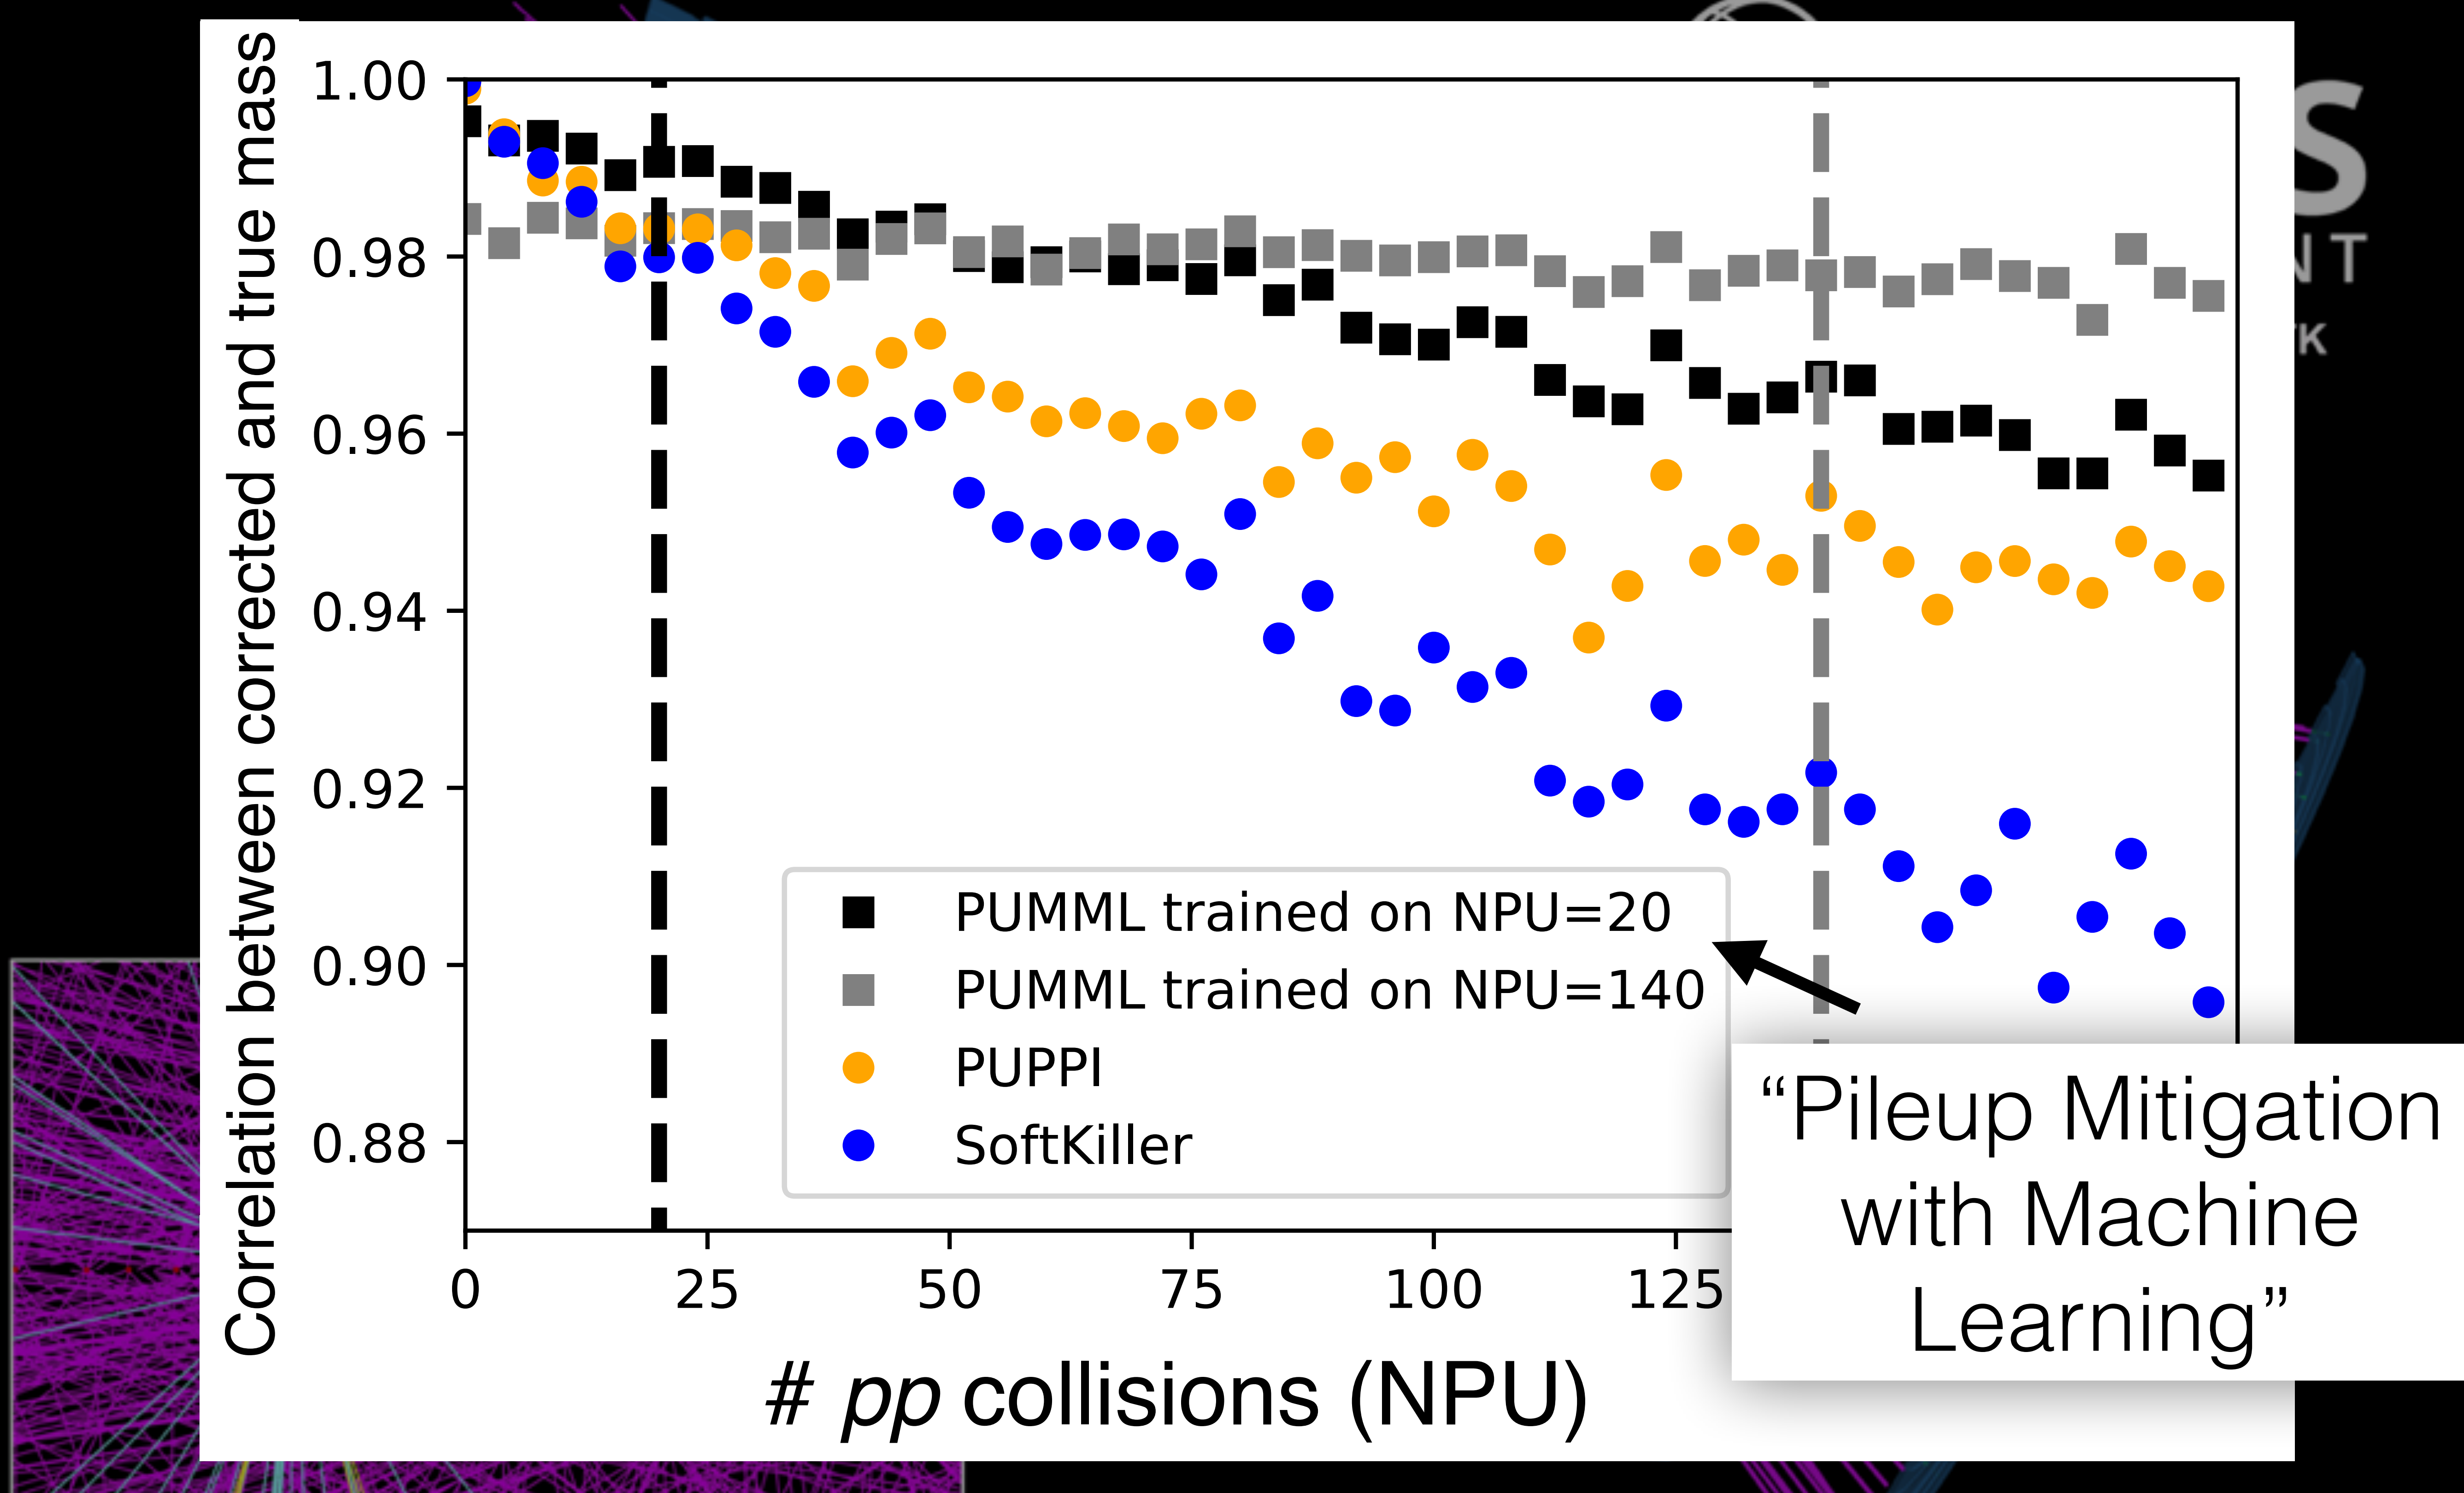

P. Komiske, E. Metodiev, **BPN**, and M. Schwartz, JHEP 12 (2017) 051

Data pipeline at the LHC

data / simulation

data time	simulation time	step
25 ns		O(100) *pp* collisions
10^{-19}-10^{-15} s	~ms	(sub-)nuclear physics
0.01-20 ns	~min	out-going particles interact with detector
1-100 ns	~100 ms	detector response (signal formation + digitization)
2.5 µs		hardware-based trigger decision
200 ms		software-based trigger decision
~100 ms		event reconstruction
		event processing (skim, thin, augment)
~months		final data analysis (uses millions of events)

Data pipeline at the LHC

data / simulation

25 ns — O(100) *pp* collisions

10^{-19}-10^{-15} s — (sub-)nuclear physics

~ms

0.01-20 ns — out-going particles interact with detector

~min

1-100 ns — detector response (signal formation + digitization)

~100 ms

2.5 μs — hardware-based trigger decision

200 ms — software-based tr…

~100 ms — e…

e… processing (skim, thin, augment)

~months — final data analysis (uses millions of events)

Not needed for “real time” decisions, but we need **real time speed** to make this work

Simulation at the LHC

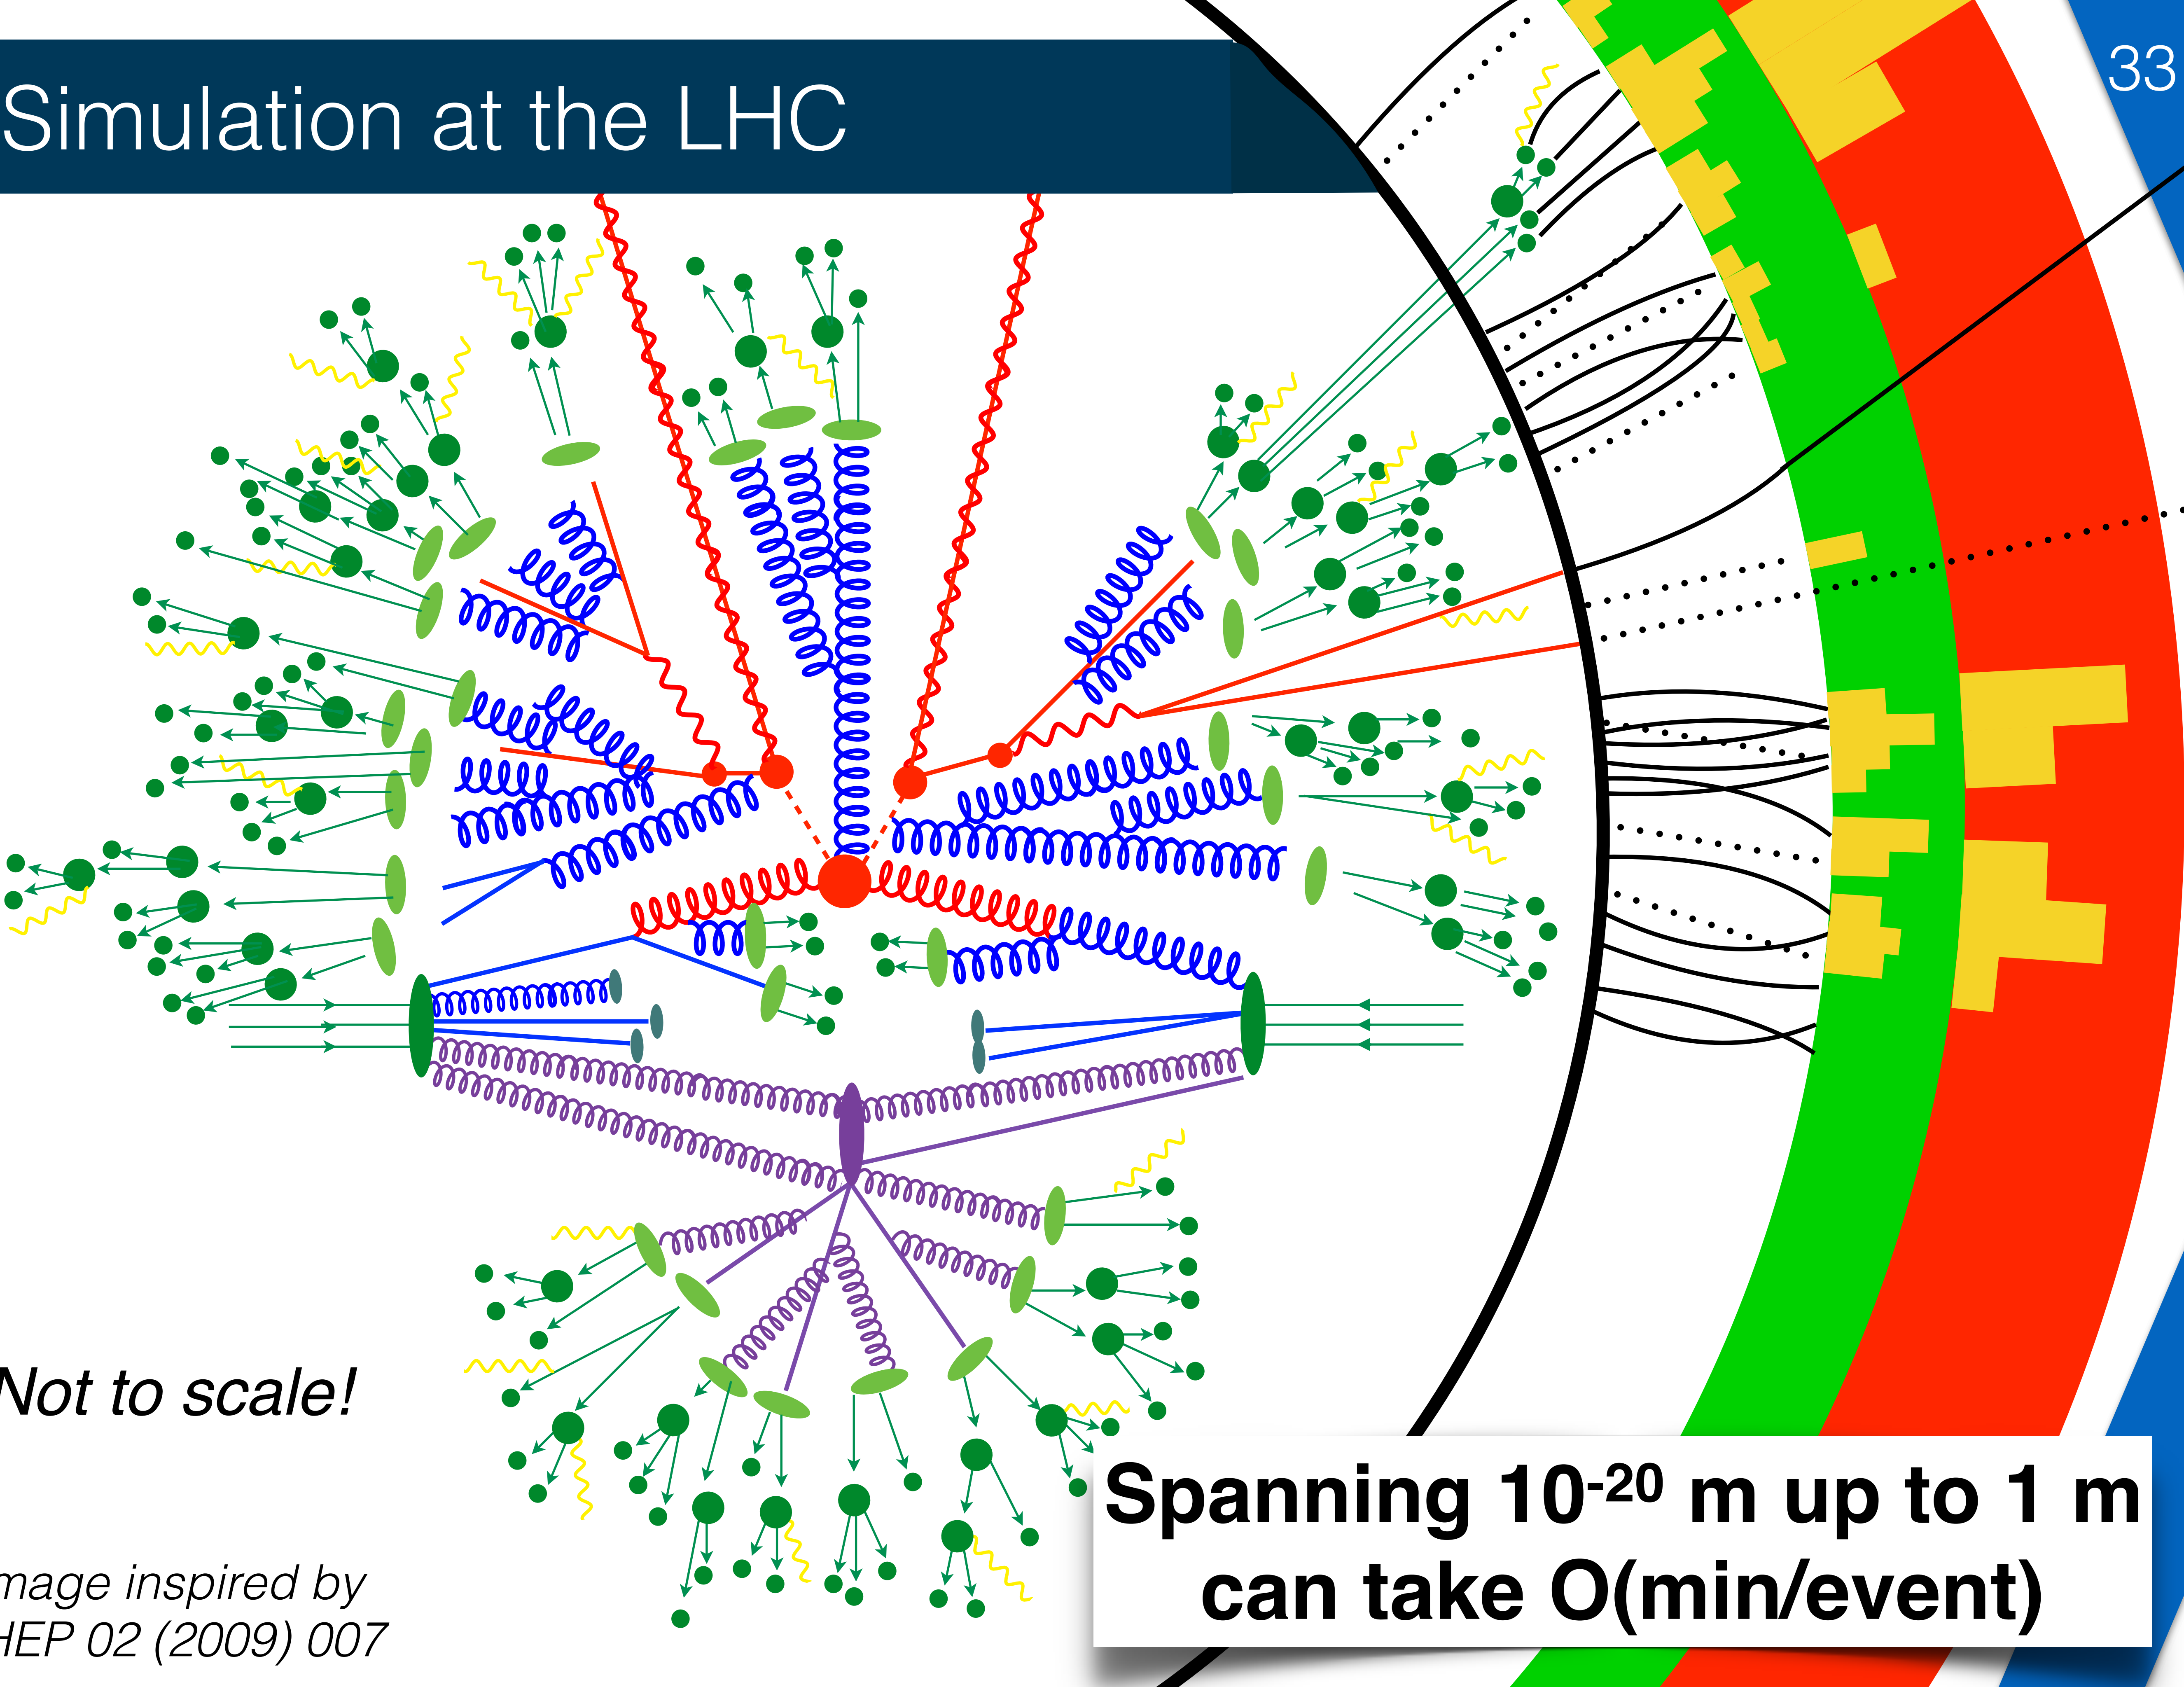

Not to scale!

Image inspired by JHEP 02 (2009) 007

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

Simulation at the LHC

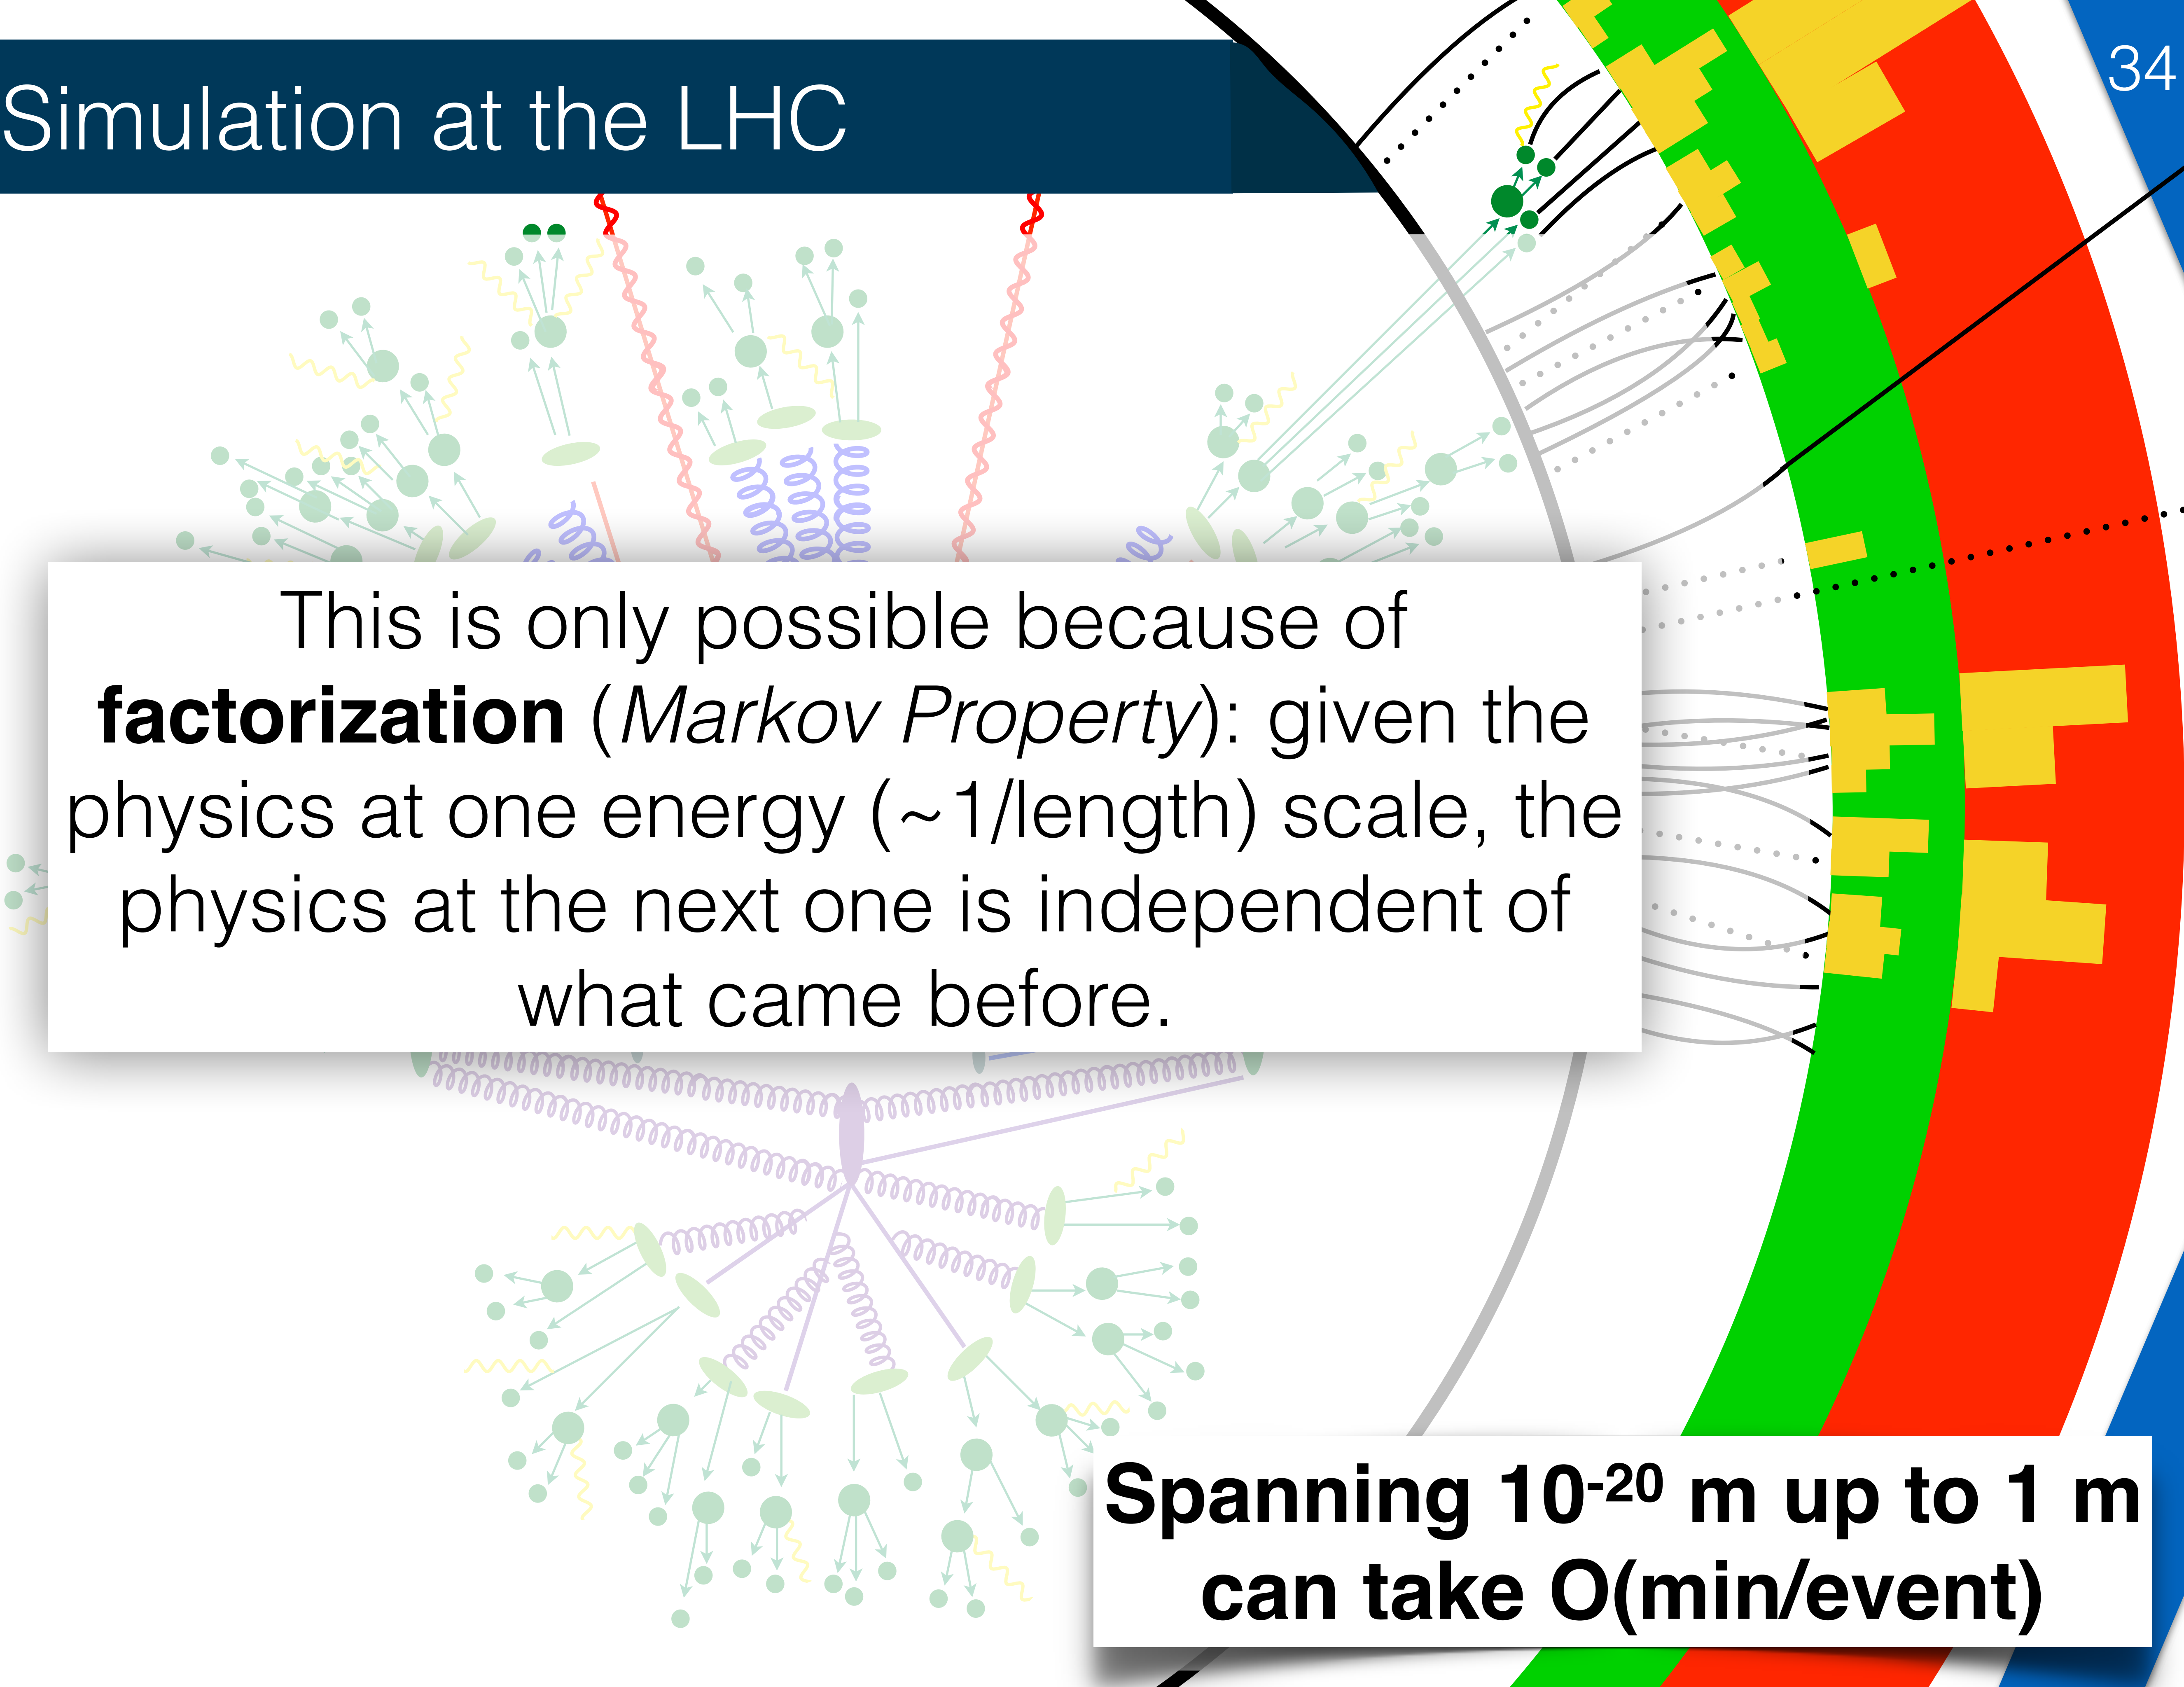

This is only possible because of **factorization** (*Markov Property*): given the physics at one energy (~1/length) scale, the physics at the next one is independent of what came before.

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

Part I: “Hard-scatter”

We begin with equations of motion →

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + i\bar{\psi}\not{D}\psi + \psi_i y_{ij}\psi_j\phi + \text{h.c.} + |D_\mu\phi|^2 - V(\phi) + ???$$

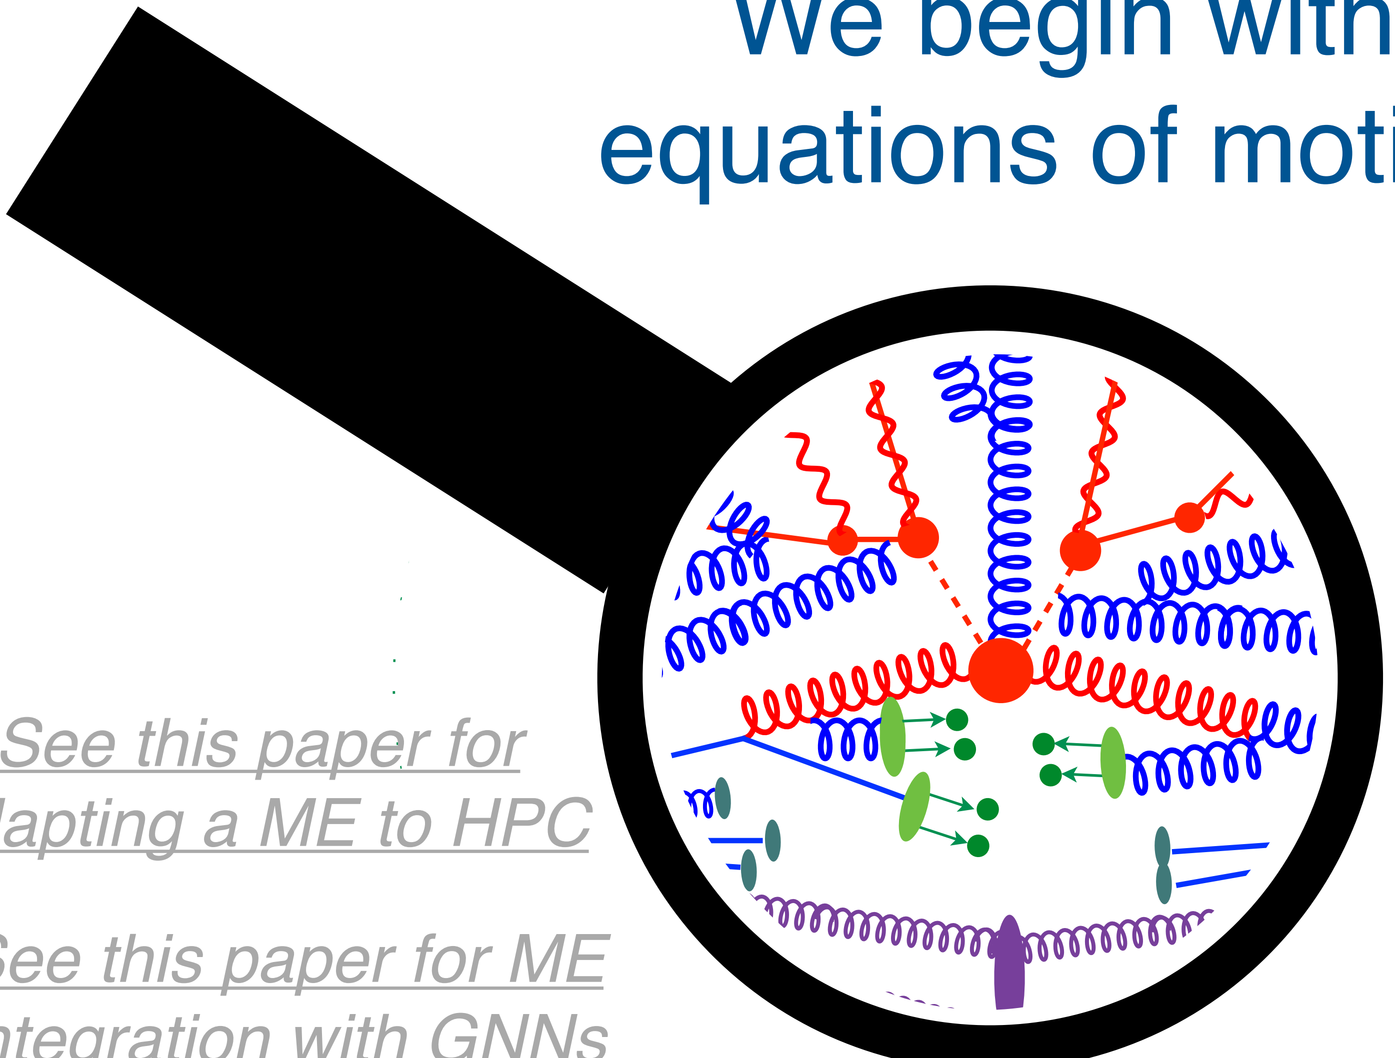

See this paper for adapting a ME to HPC

See this paper for ME integration with GNNs

```
*********************************************************
*                                                       *
*                  W E L C O M E  to                    *
*           M A D G R A P H 5 _ a M C @ N L O           *
*                                                       *
*                                                       *
*              *                         *              *
*                *       * *           *                *
*                  * * * * 5 * * * *                    *
*                *       * *           *                *
*              *                         *              *
*                                                       *
*********************************************************
```

Many tools exist for automating this highest energy step

For many cases, this is slow but not limiting (yet)

Part II: Quarks → protons (+ friends)

Makes heavy use of MCMC

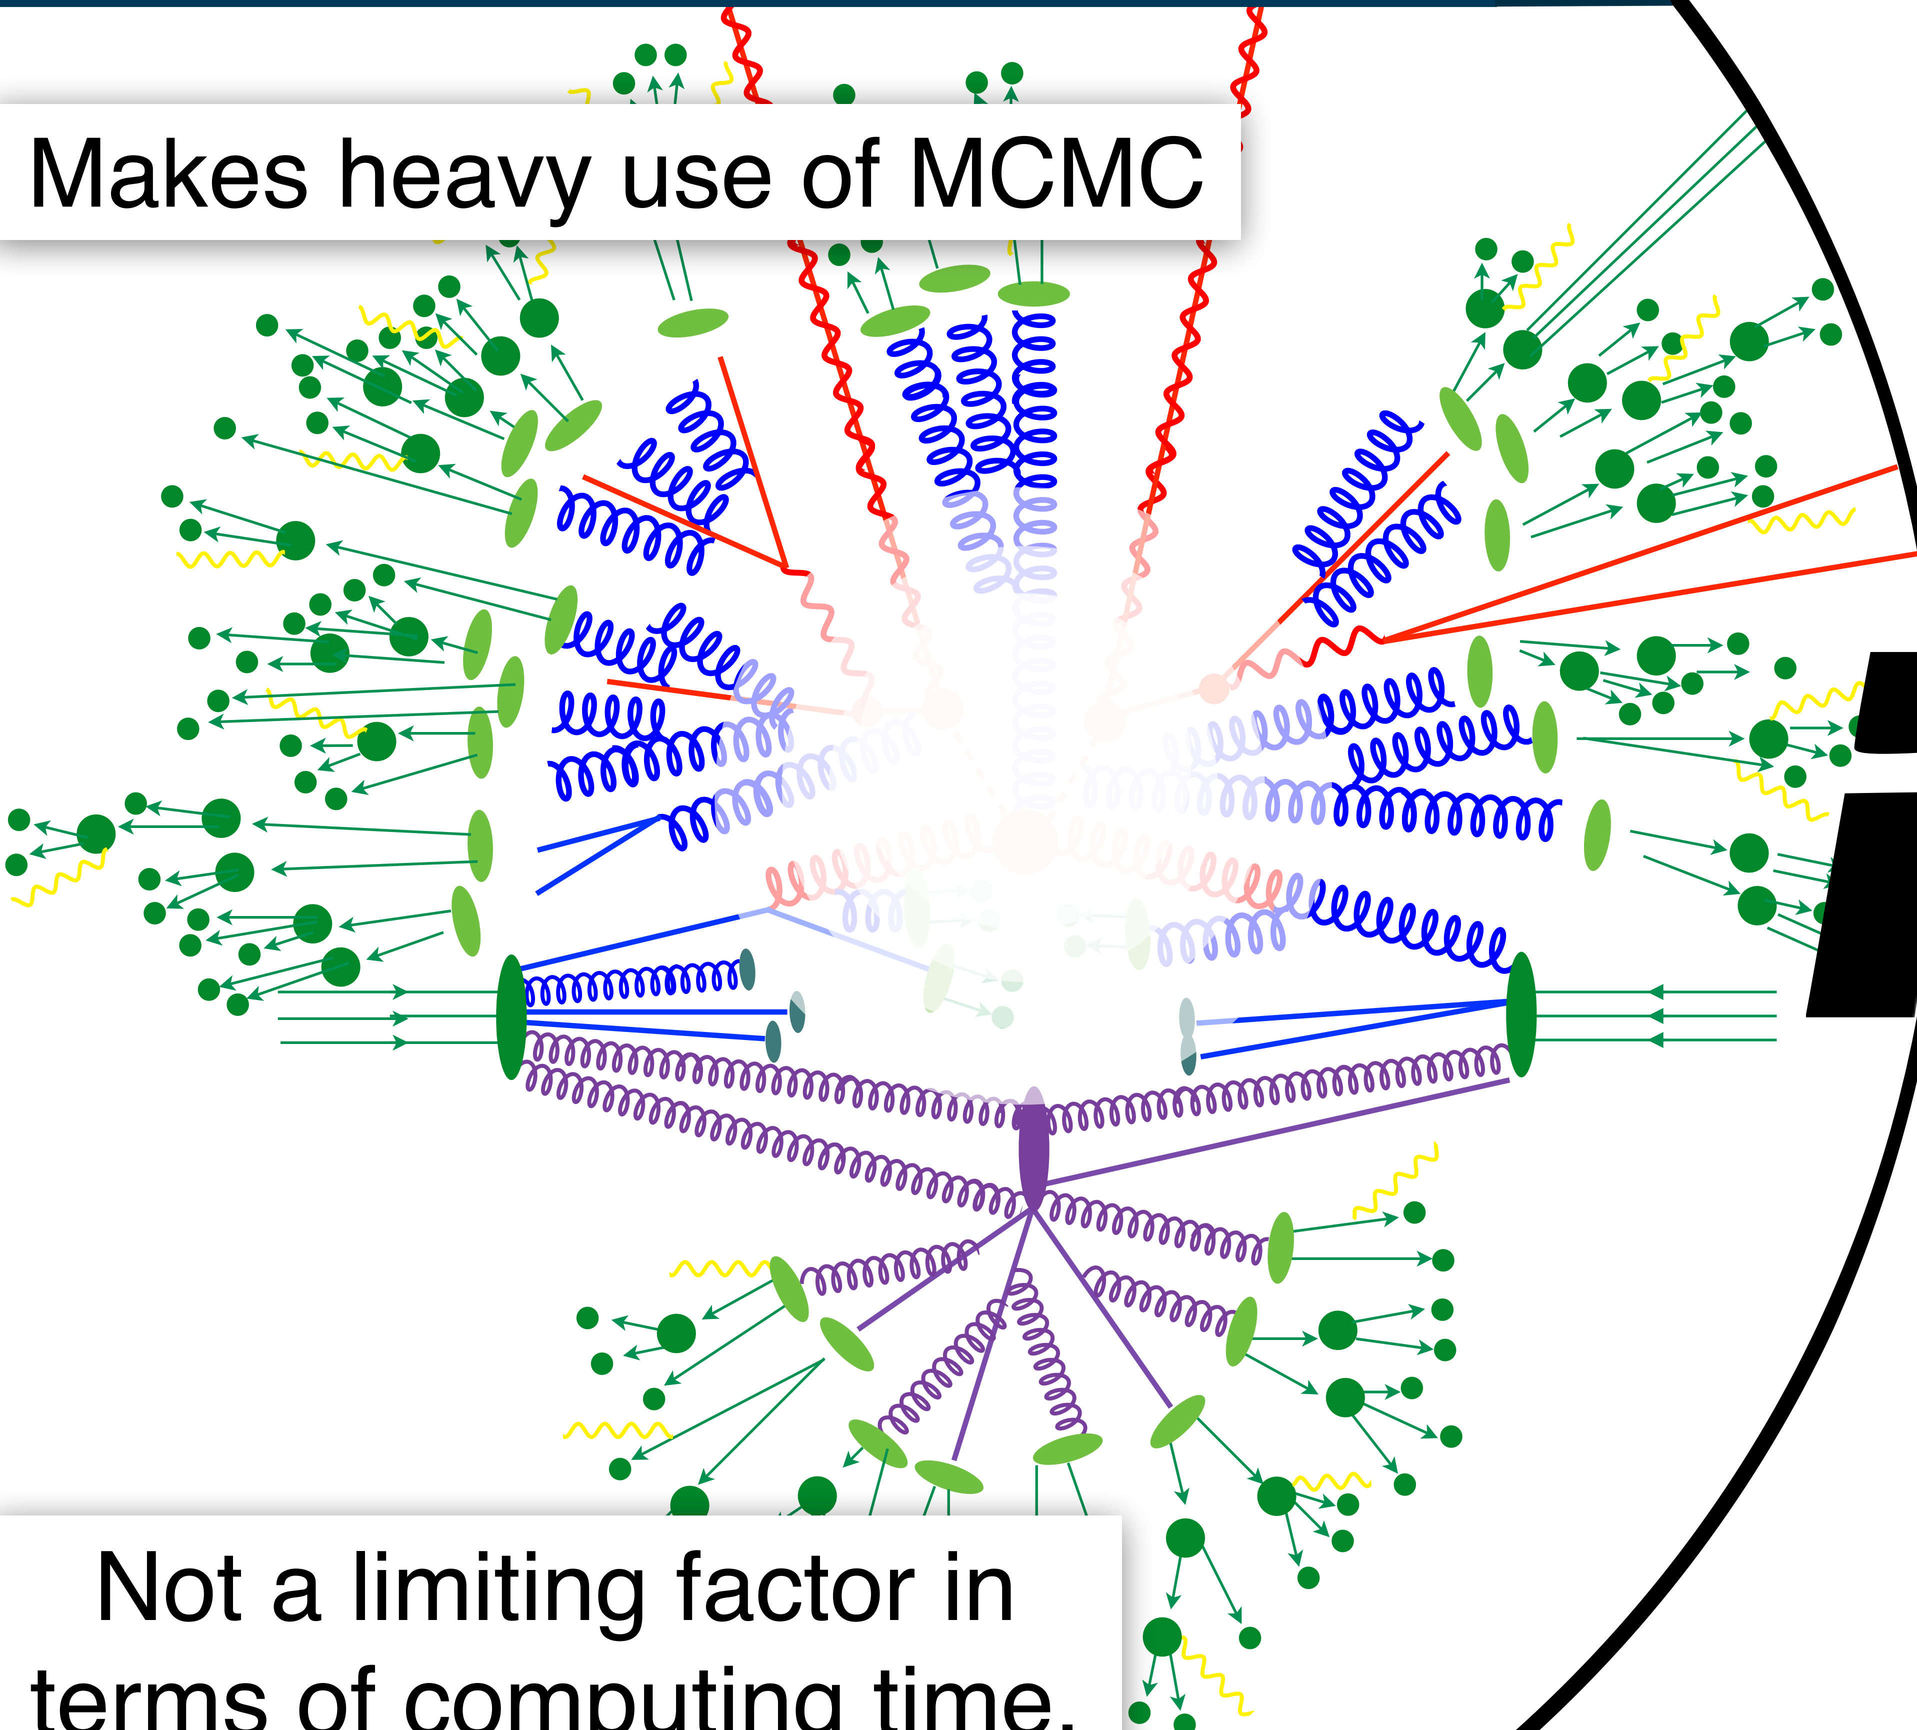

Not a limiting factor in terms of computing time.

Part III: Material Interactions

State-of-the-art for material interactions is Geant4.

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 competing resources!

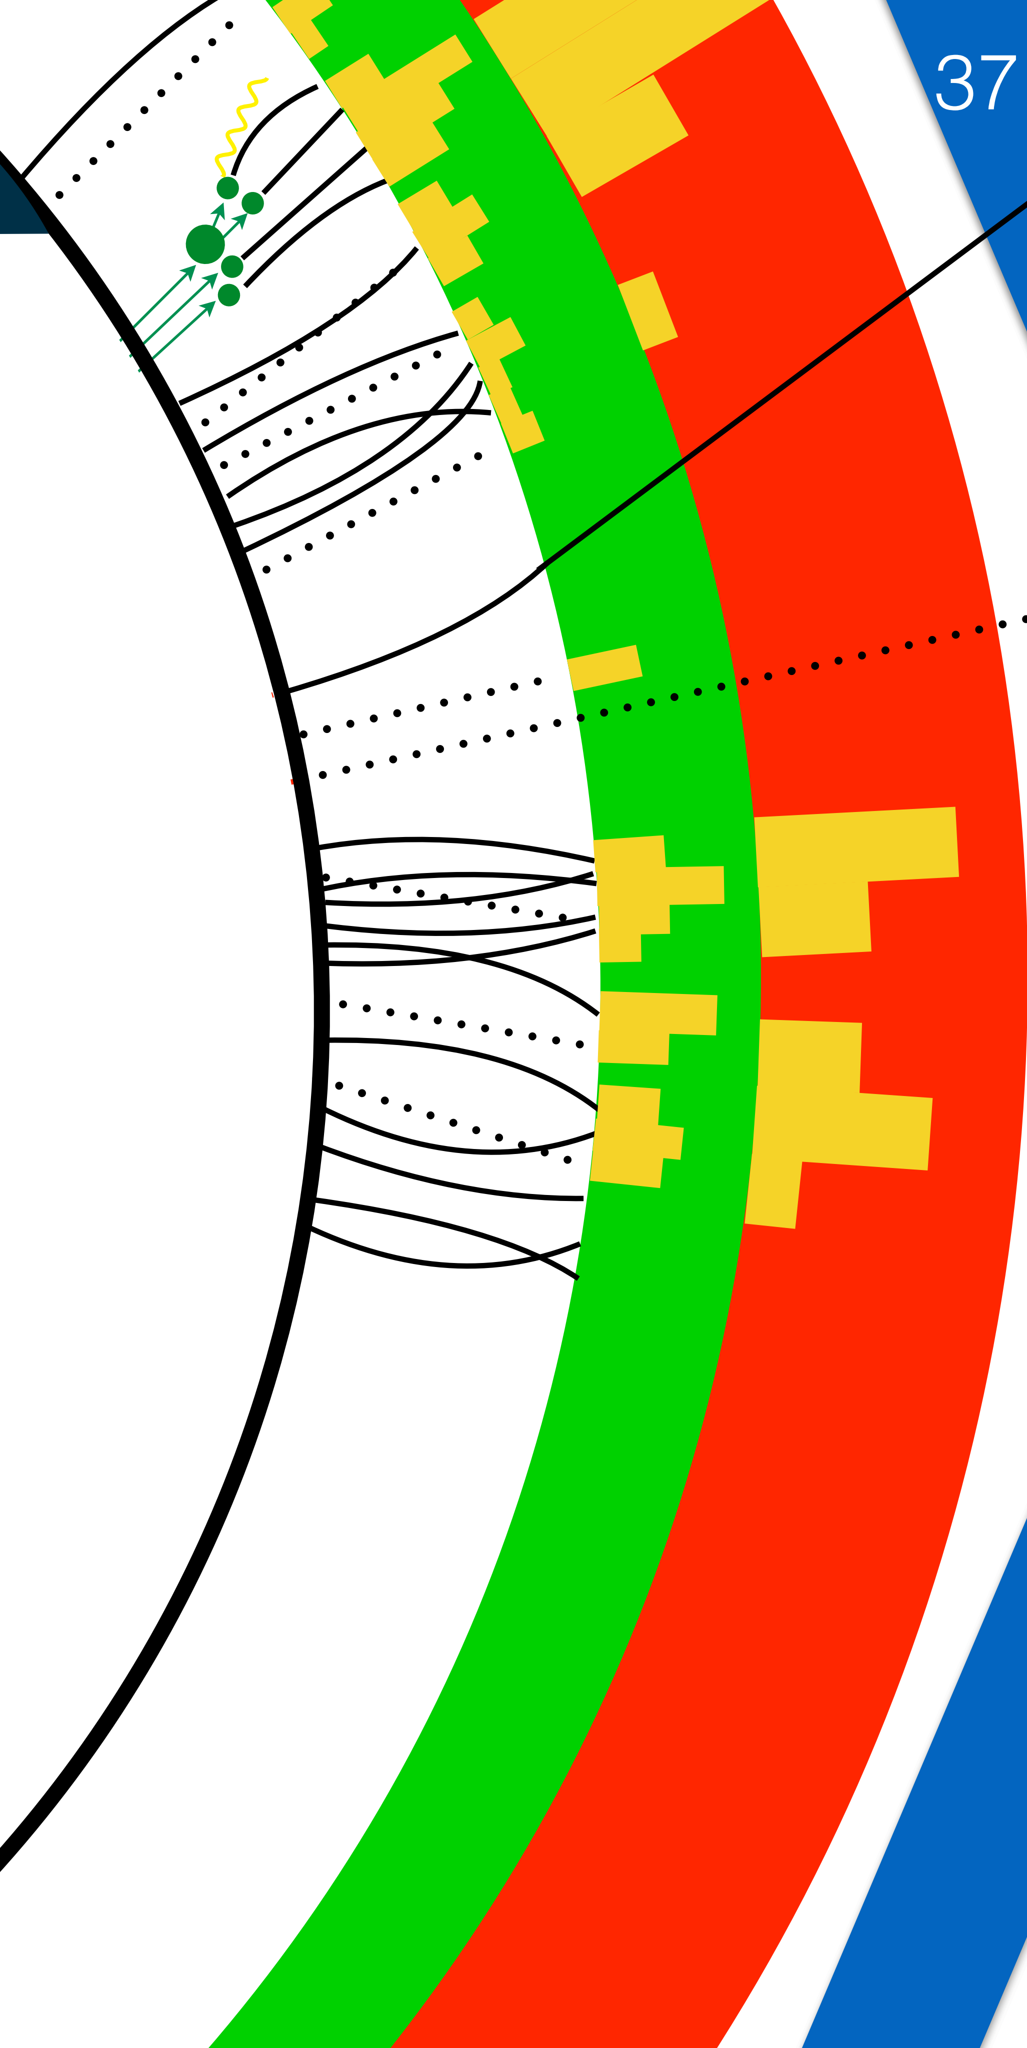

Part IV: Digitization

It is important to mention that **after** Geant4, each experiment has custom code for ***digitization***

this can also be slow; but is usually faster than Geant4 and reconstruction

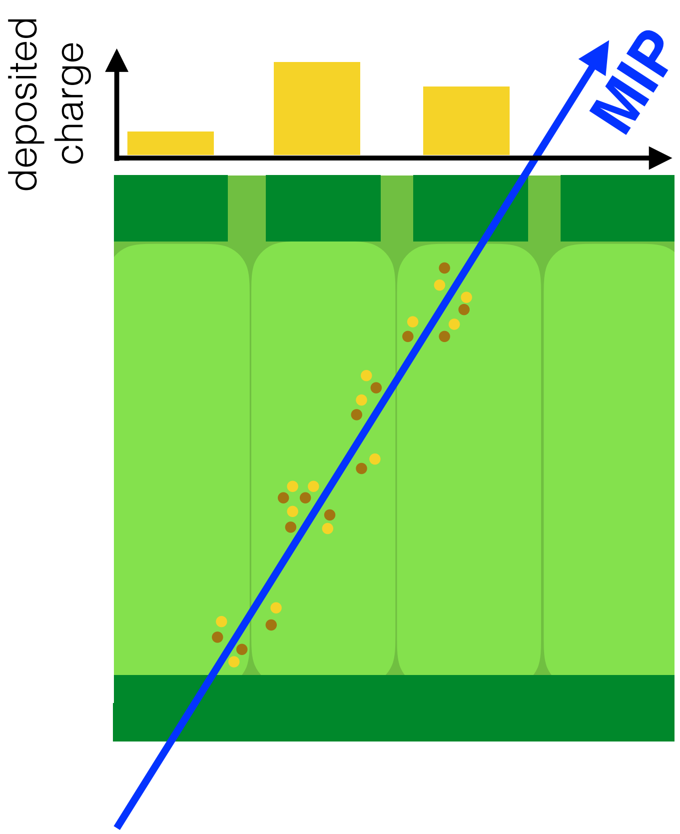

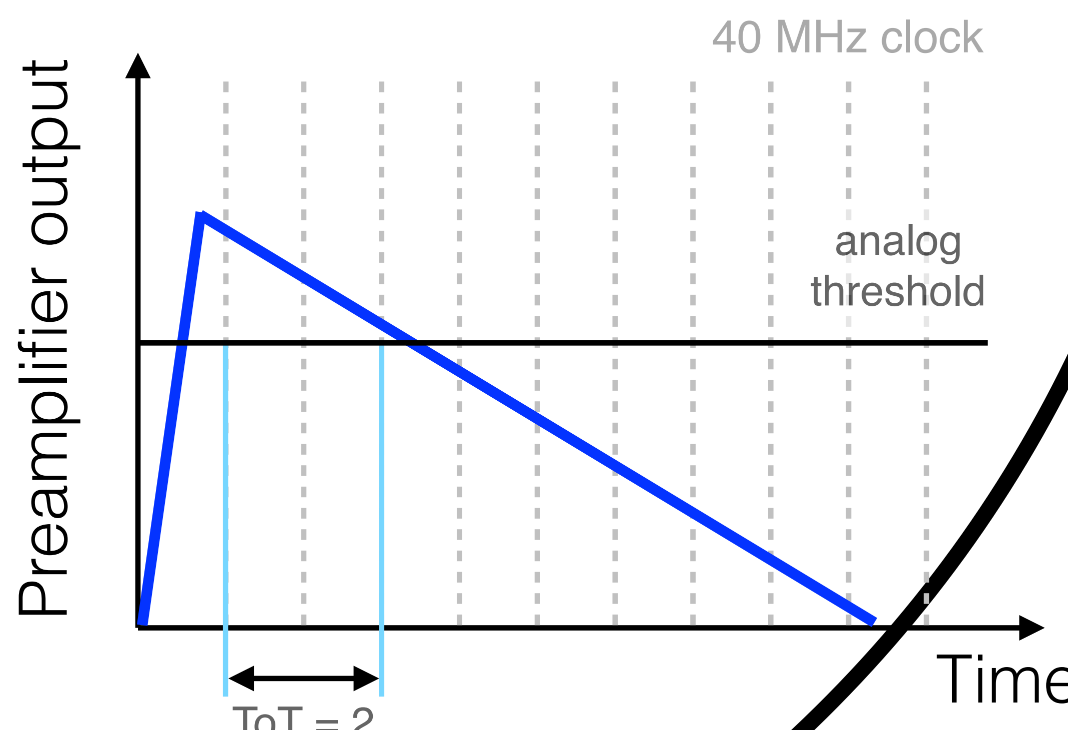

Part IV: Digitization

It is important to mention that **after** Geant4, each experiment has custom code for *digitization*

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

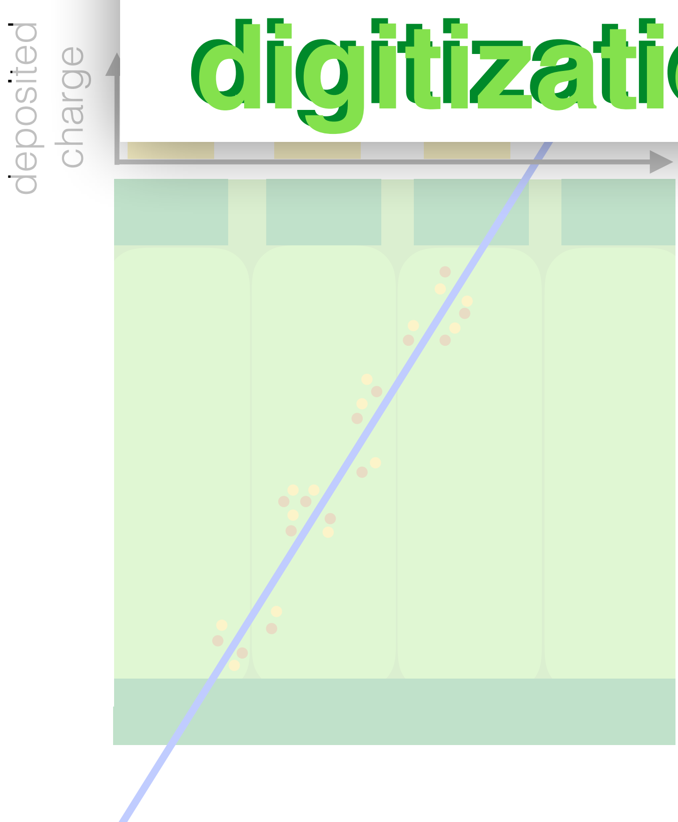

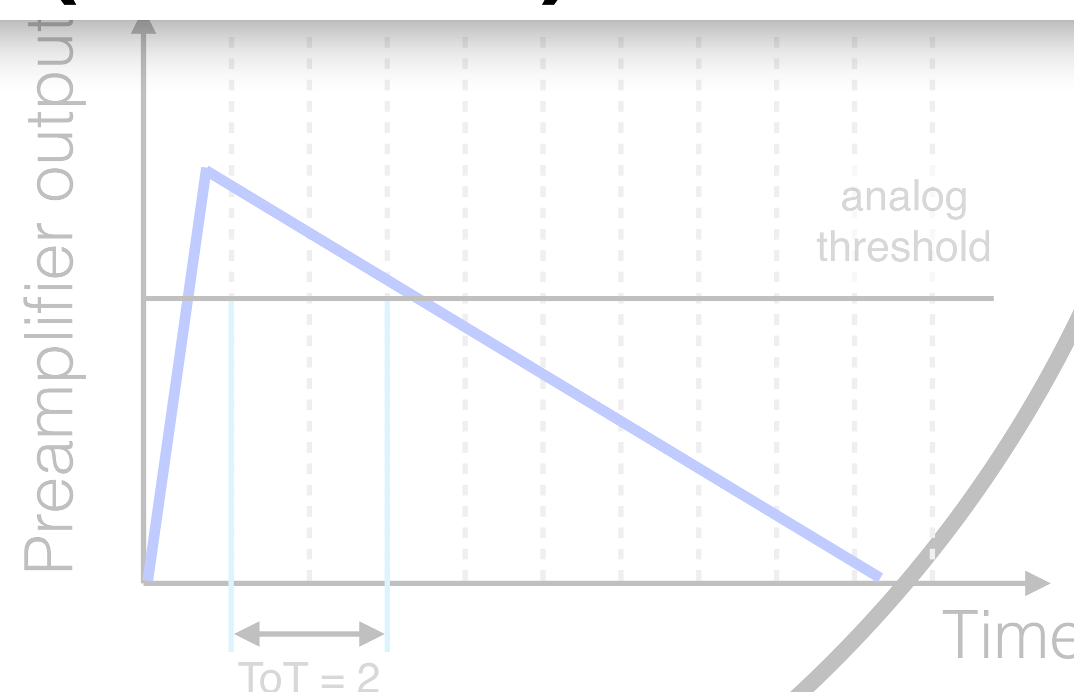

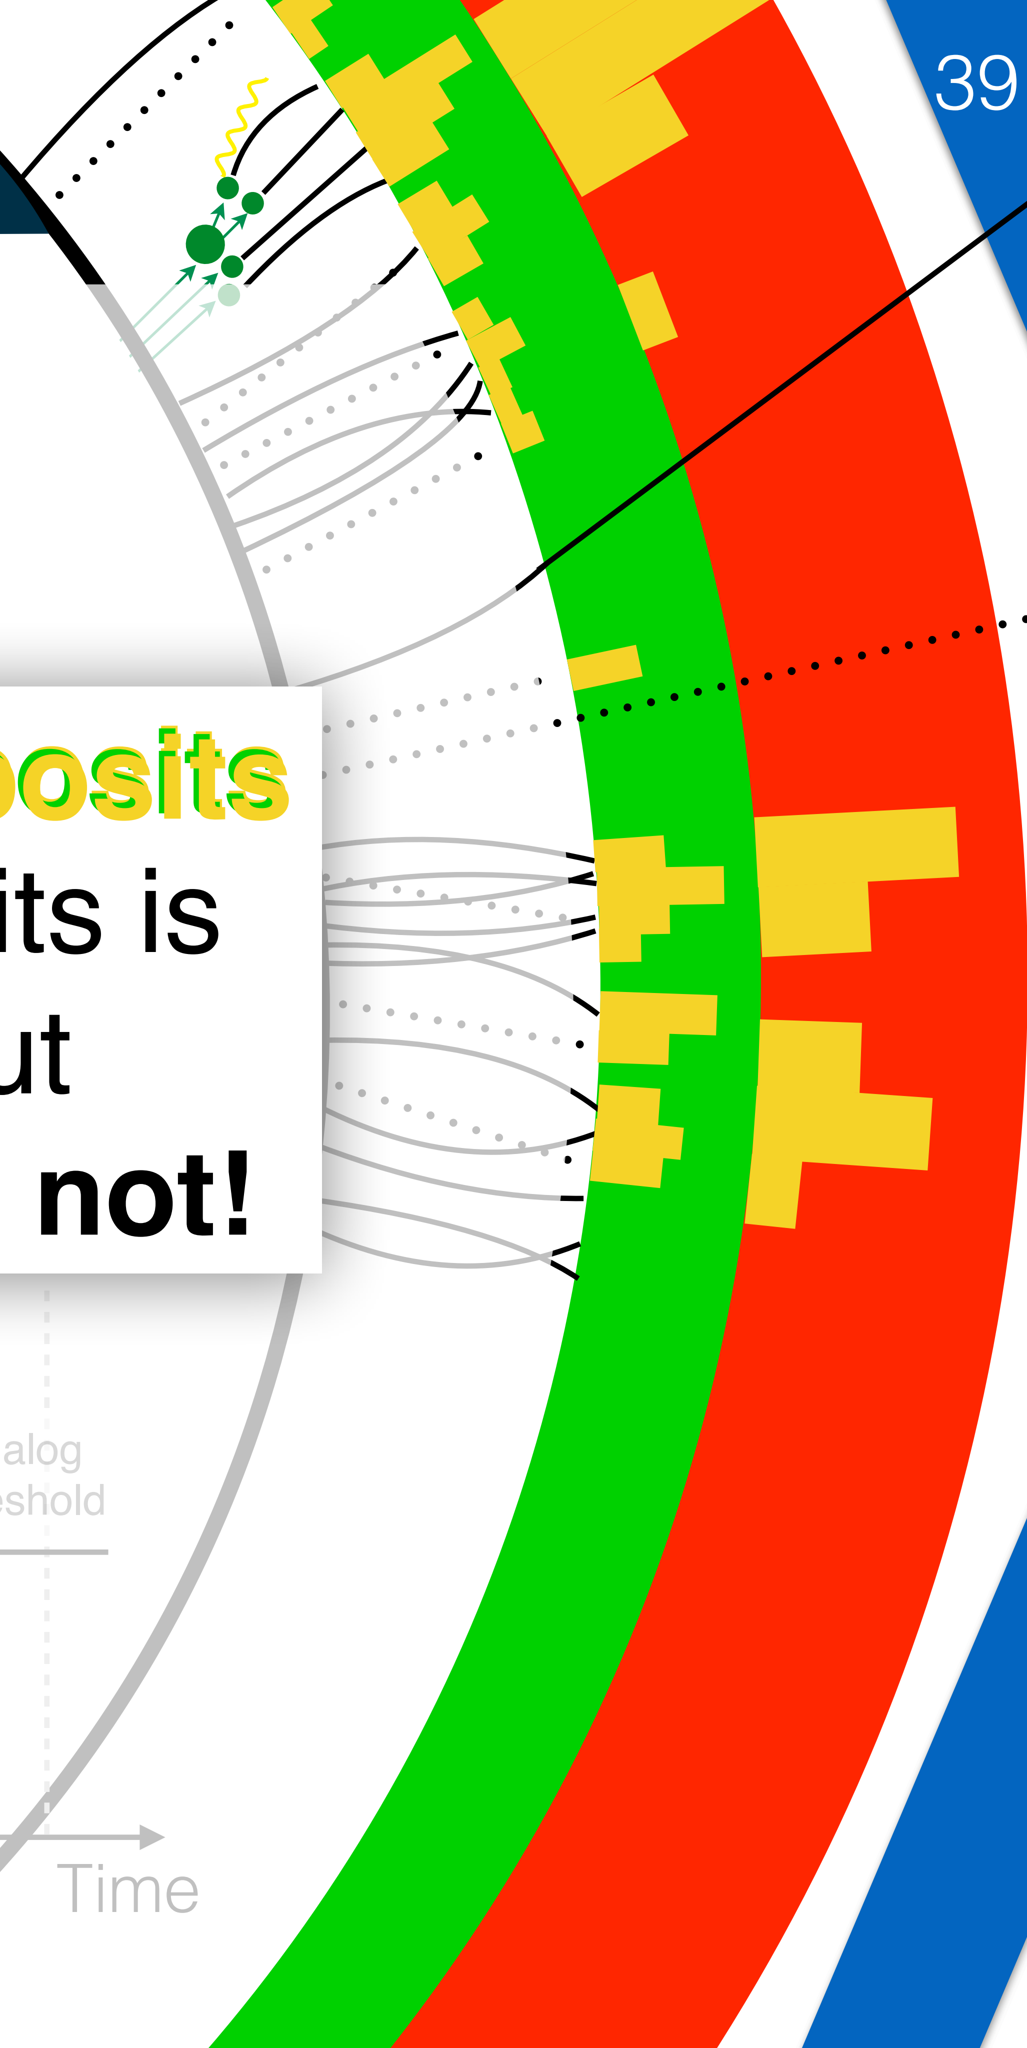

Goal: replace (or augment) simulation steps with a faster, powerful generator based on state-of-the-art machine learning techniques

To start: attack the most important part:
Calorimeter Simulation

Why should **you** care?
Many analyses are forced to use a Geant4-based simulation as current fast sim. is not good enough.

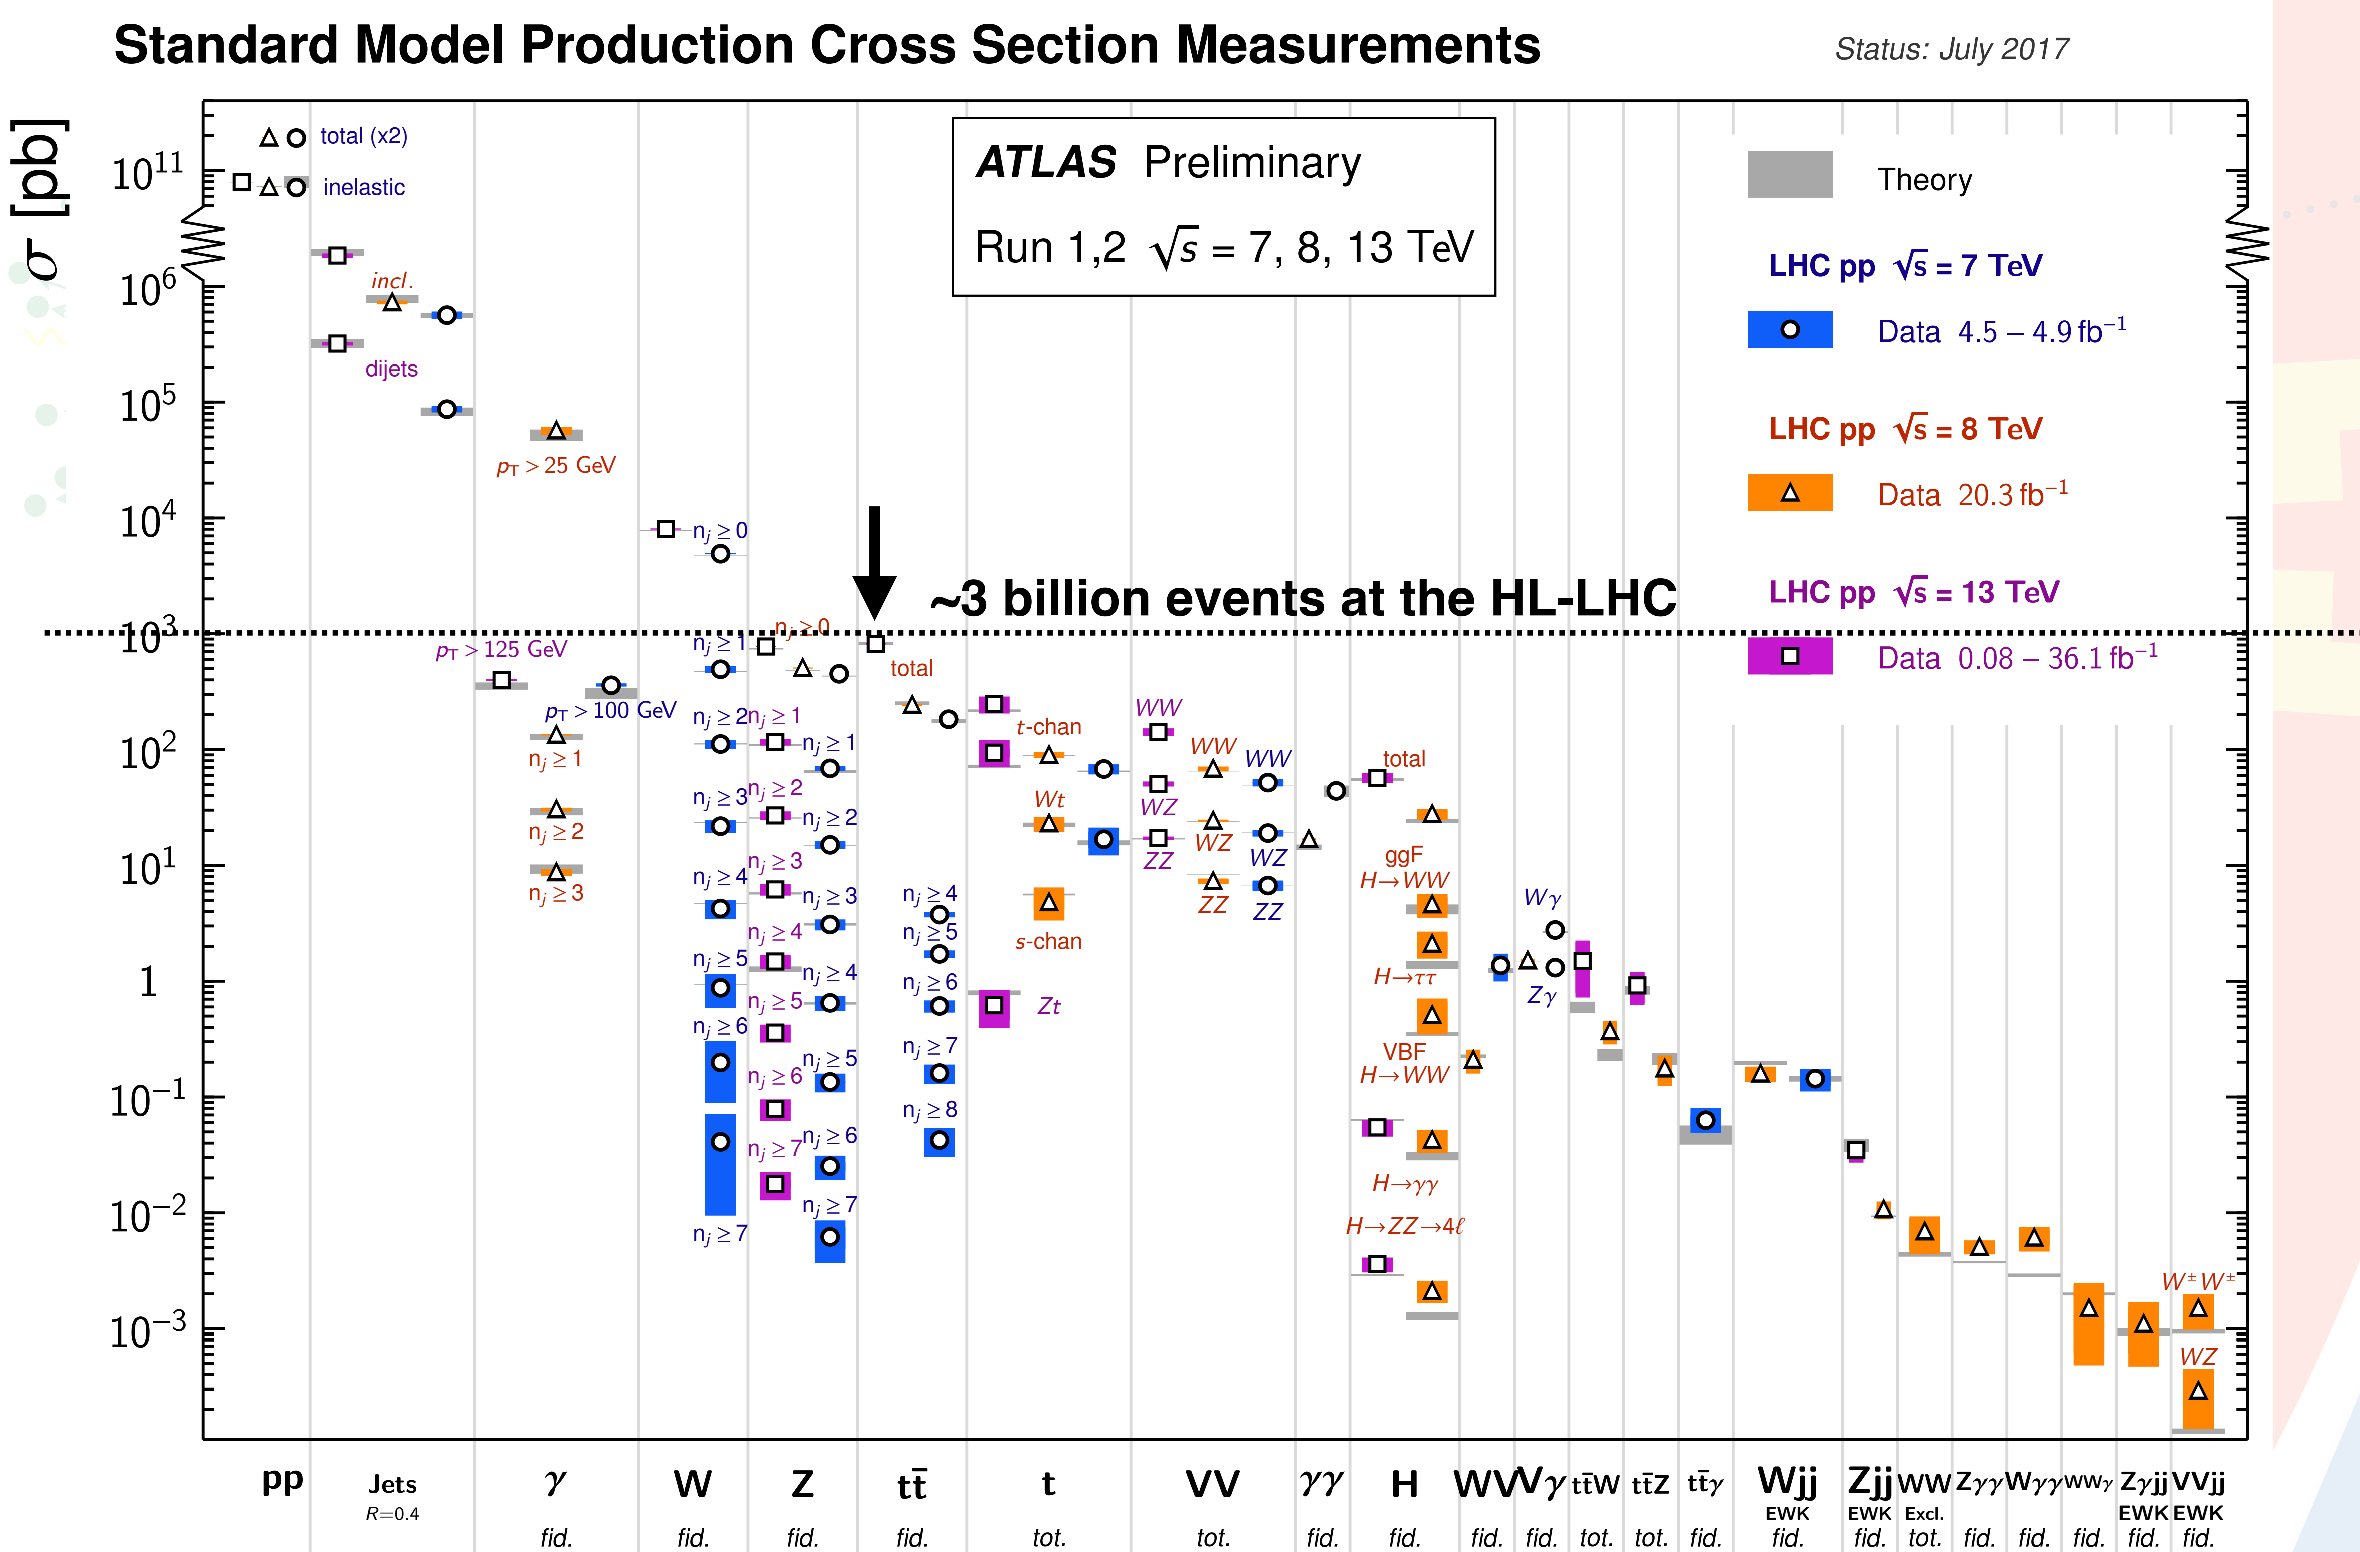

Why should **you** care?

Many analyses are forced to use a Geant4-based simulation as current fast sim. is not good enough.

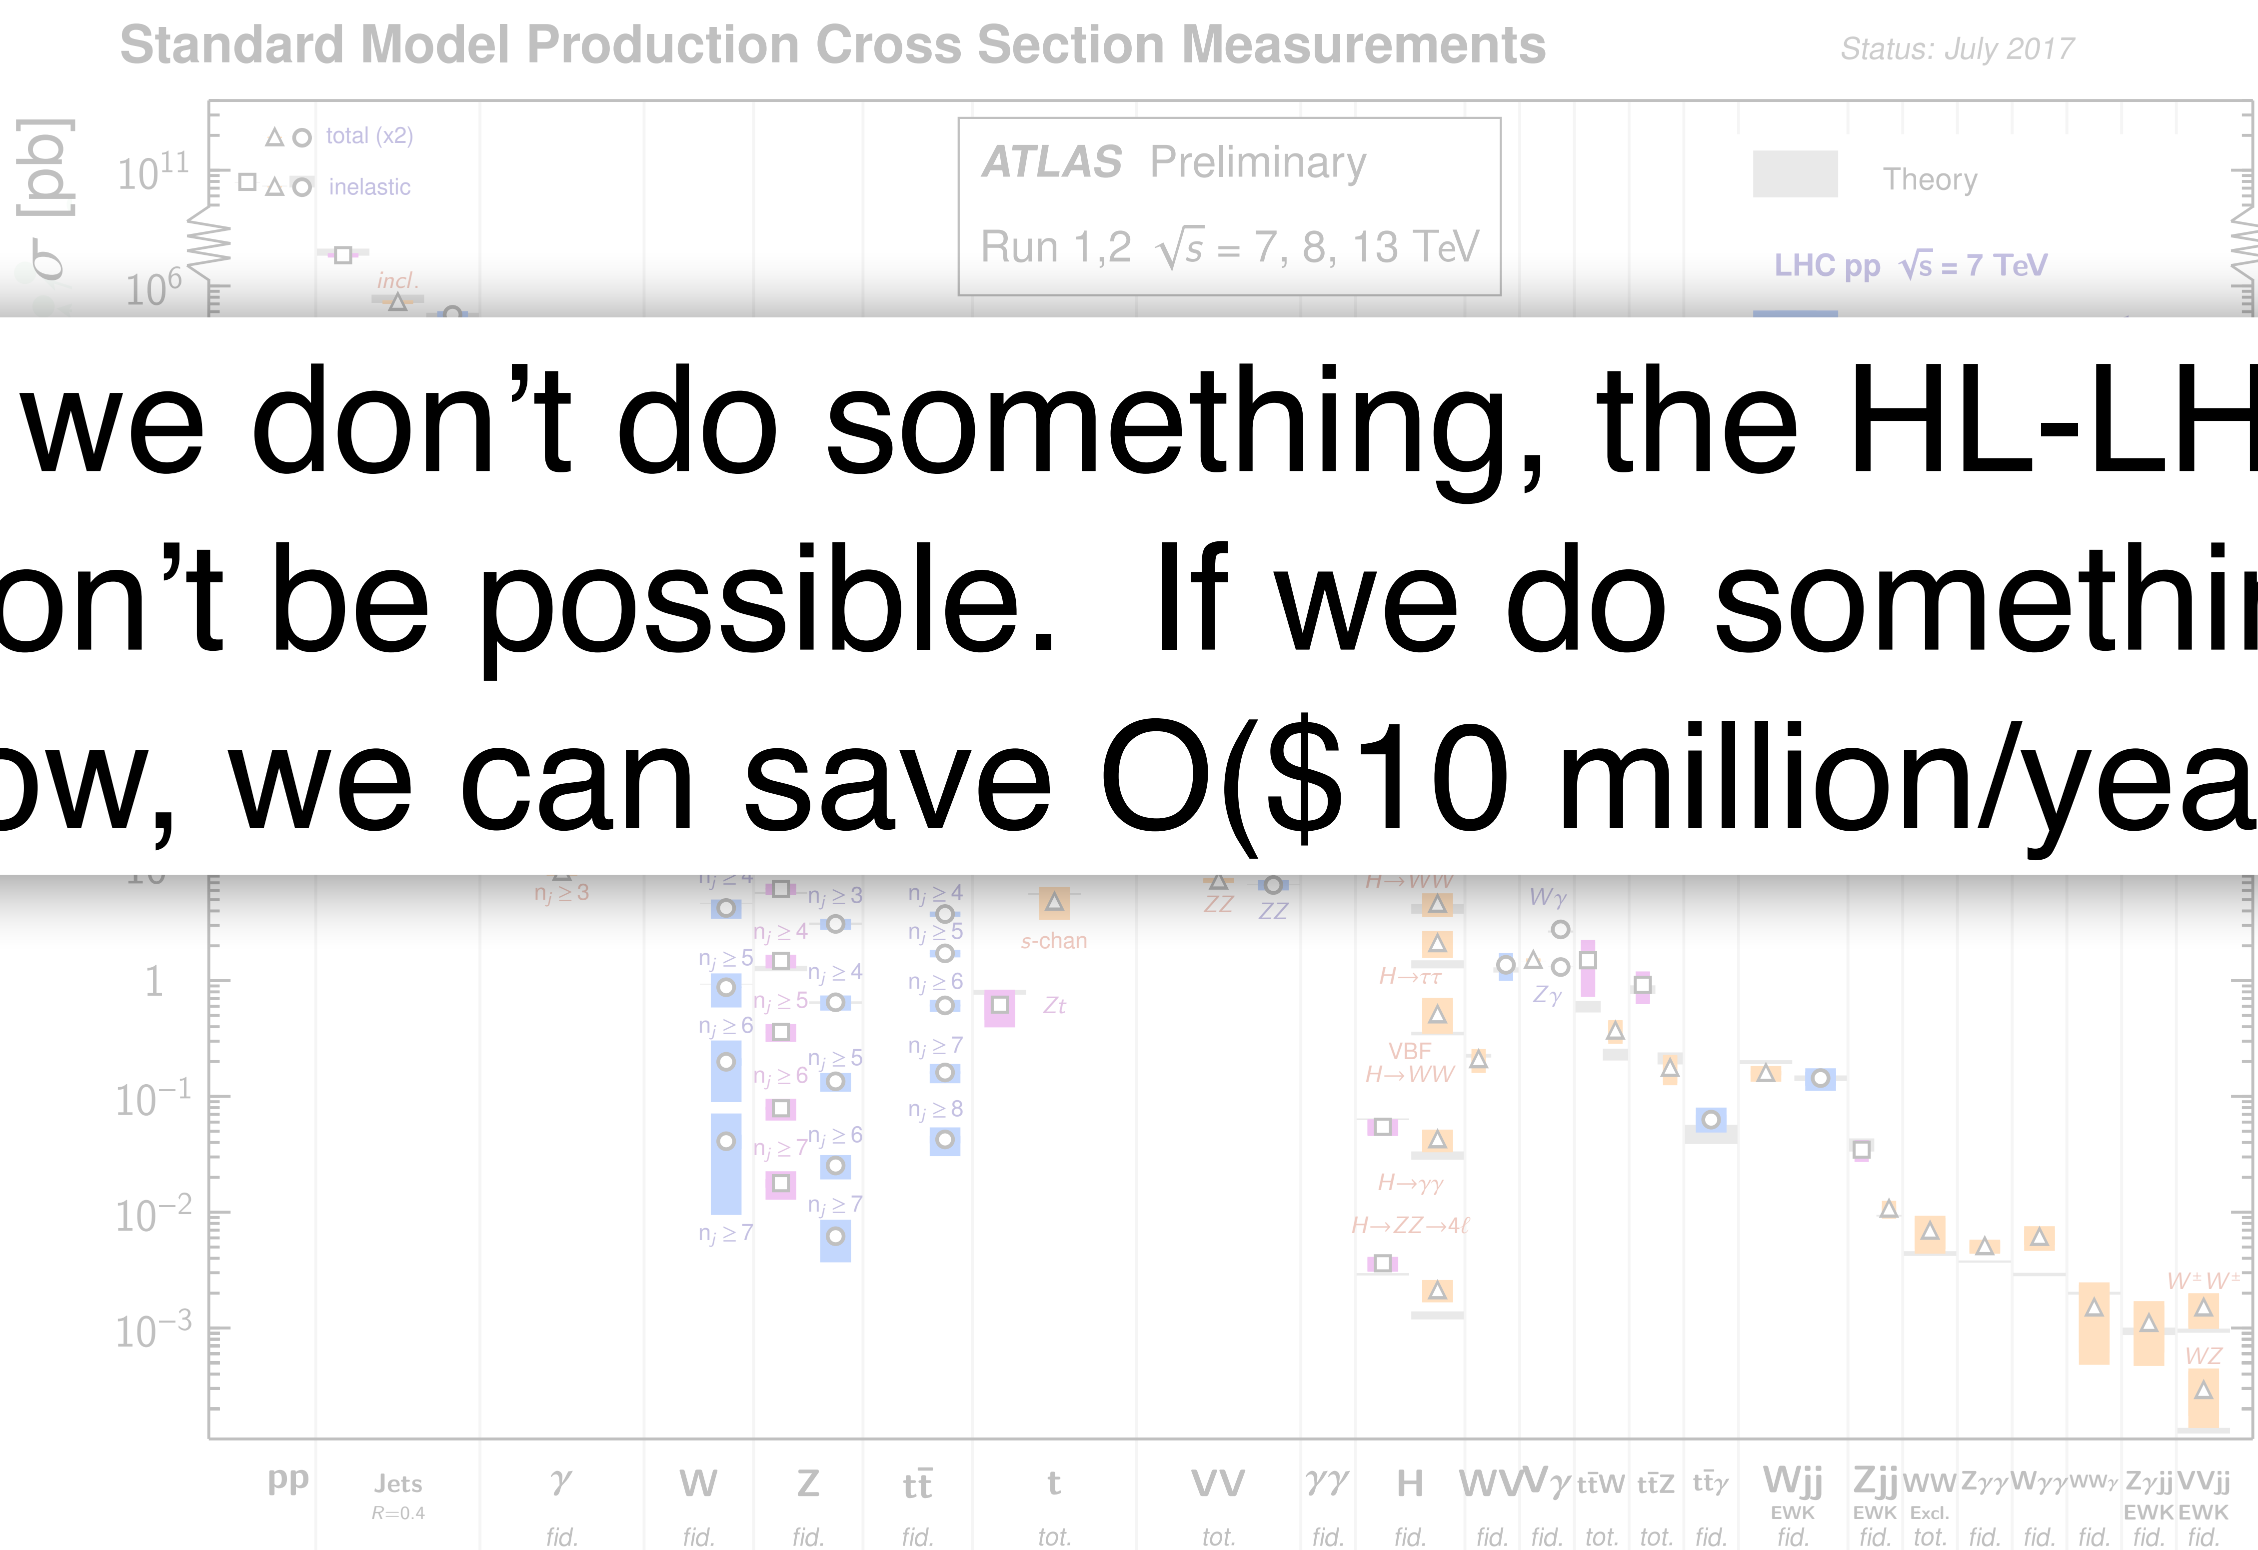

If we don't do something, the HL-LHC won't be possible. If we do something now, we can save O($10 million/year).

How can ML help?

Training NN's is slow, but evaluation is **fast**

Physics-based simulations are **slow**

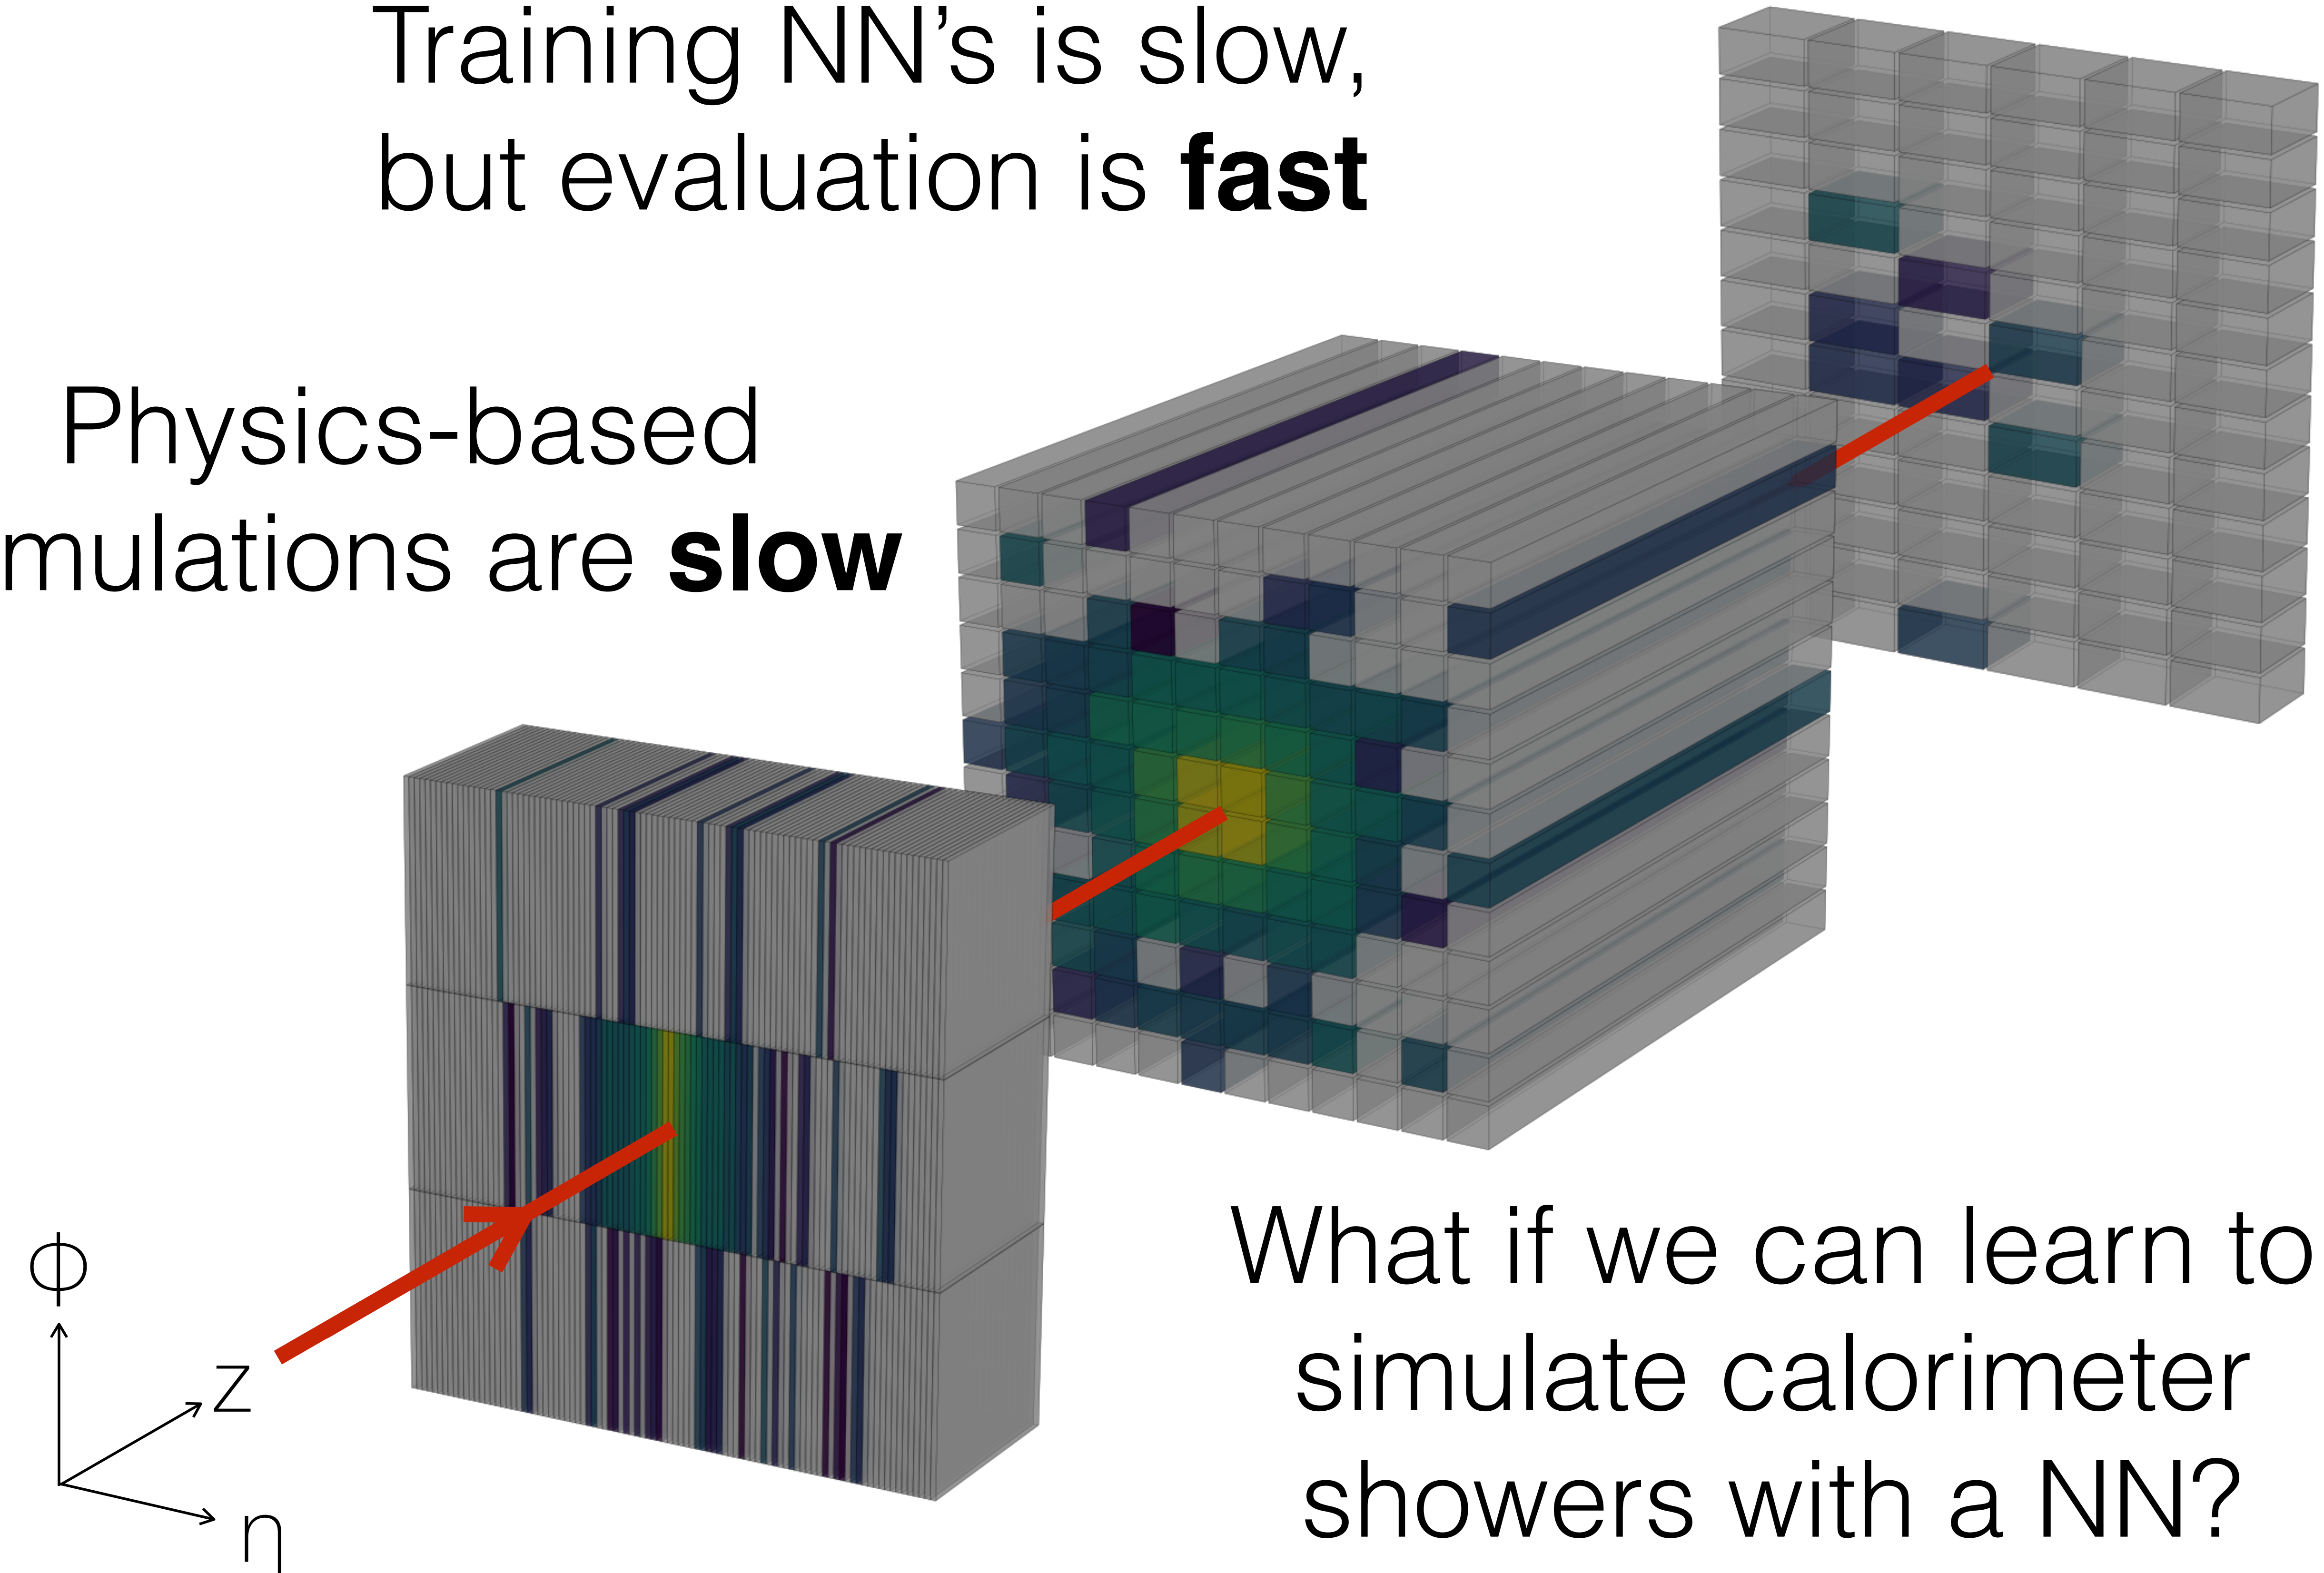

What if we can learn to simulate calorimeter showers with a NN?

Generative Adversarial Networks (GAN):

A two-network game where one ***maps noise to images*** *and one* ***classifies images as fake or real****.*

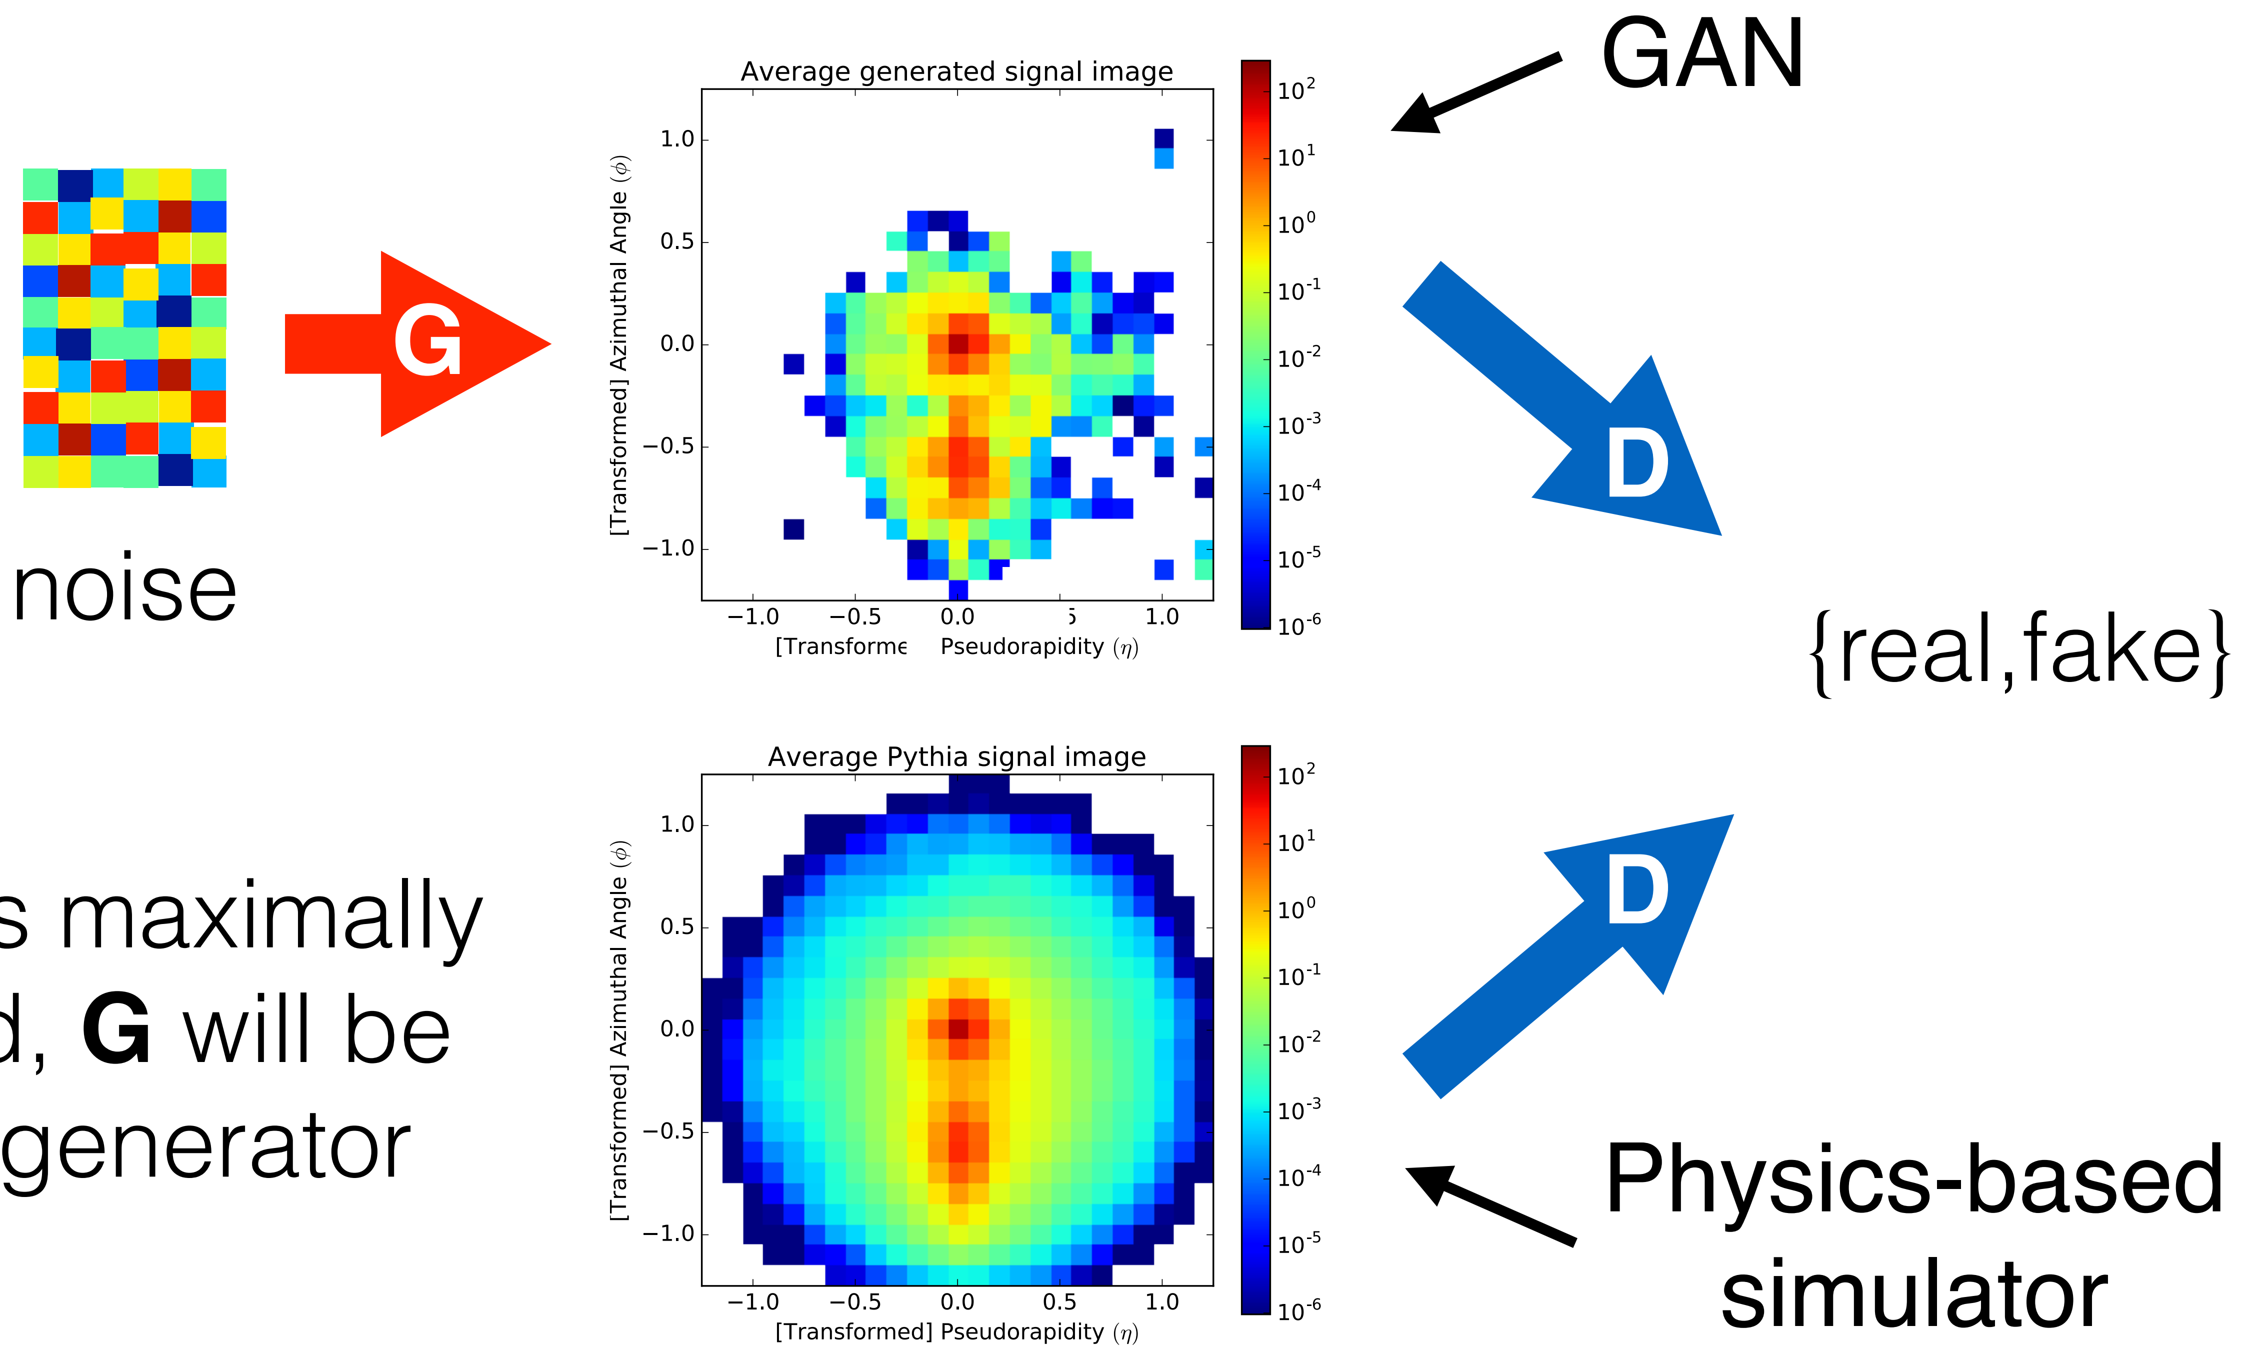

When **D** is maximally confused, **G** will be a good generator

+ More Layers for Generation

What about **multiple layers** with **non-uniform granularity** and a **causal relationship**?

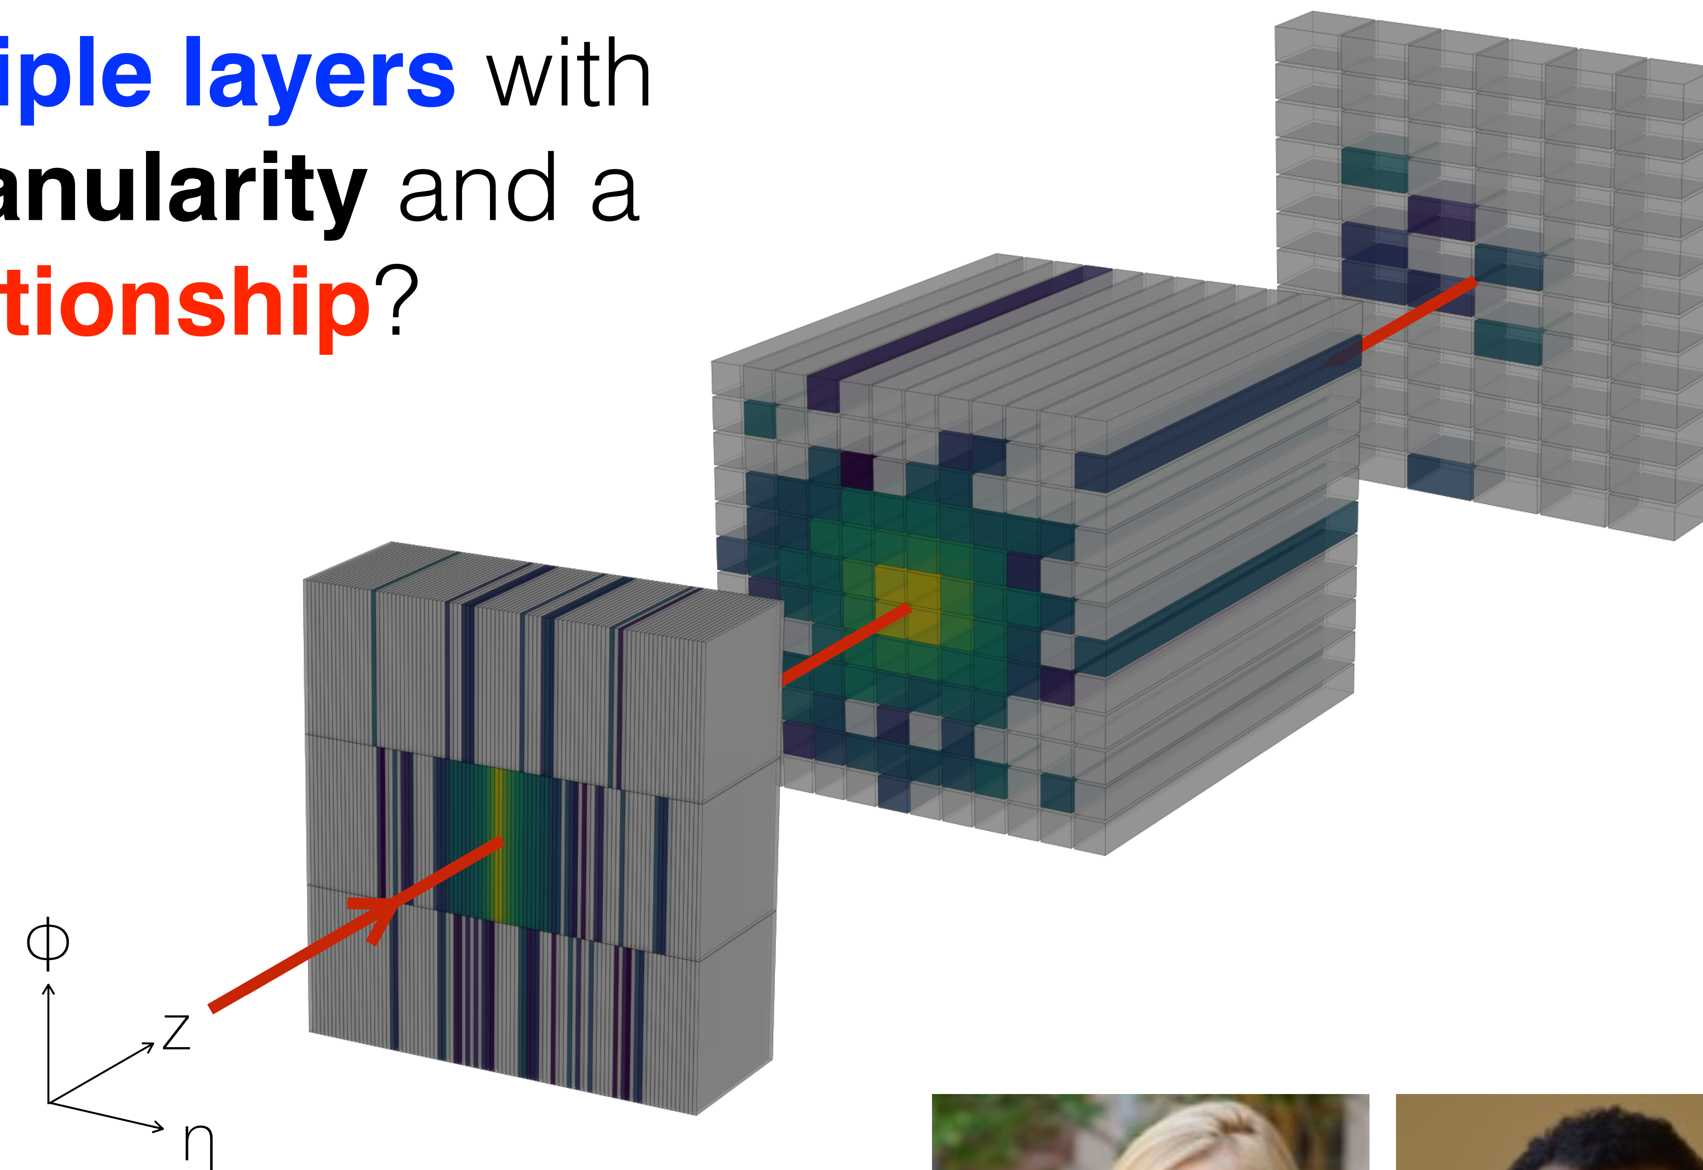

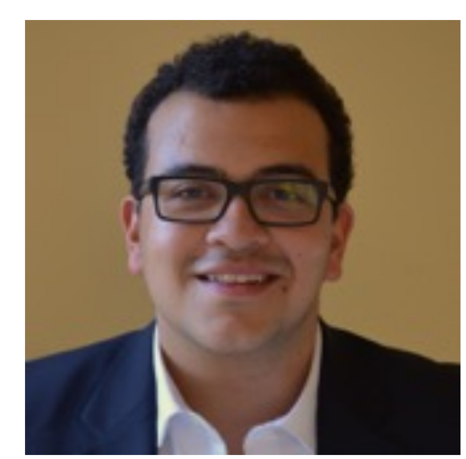

M. Paganini, L. de Oliveira,
***BPN**, PRL 120 (2018) 042003*

Generator Network

One image per calo layer

One network per particle type; input particle energy

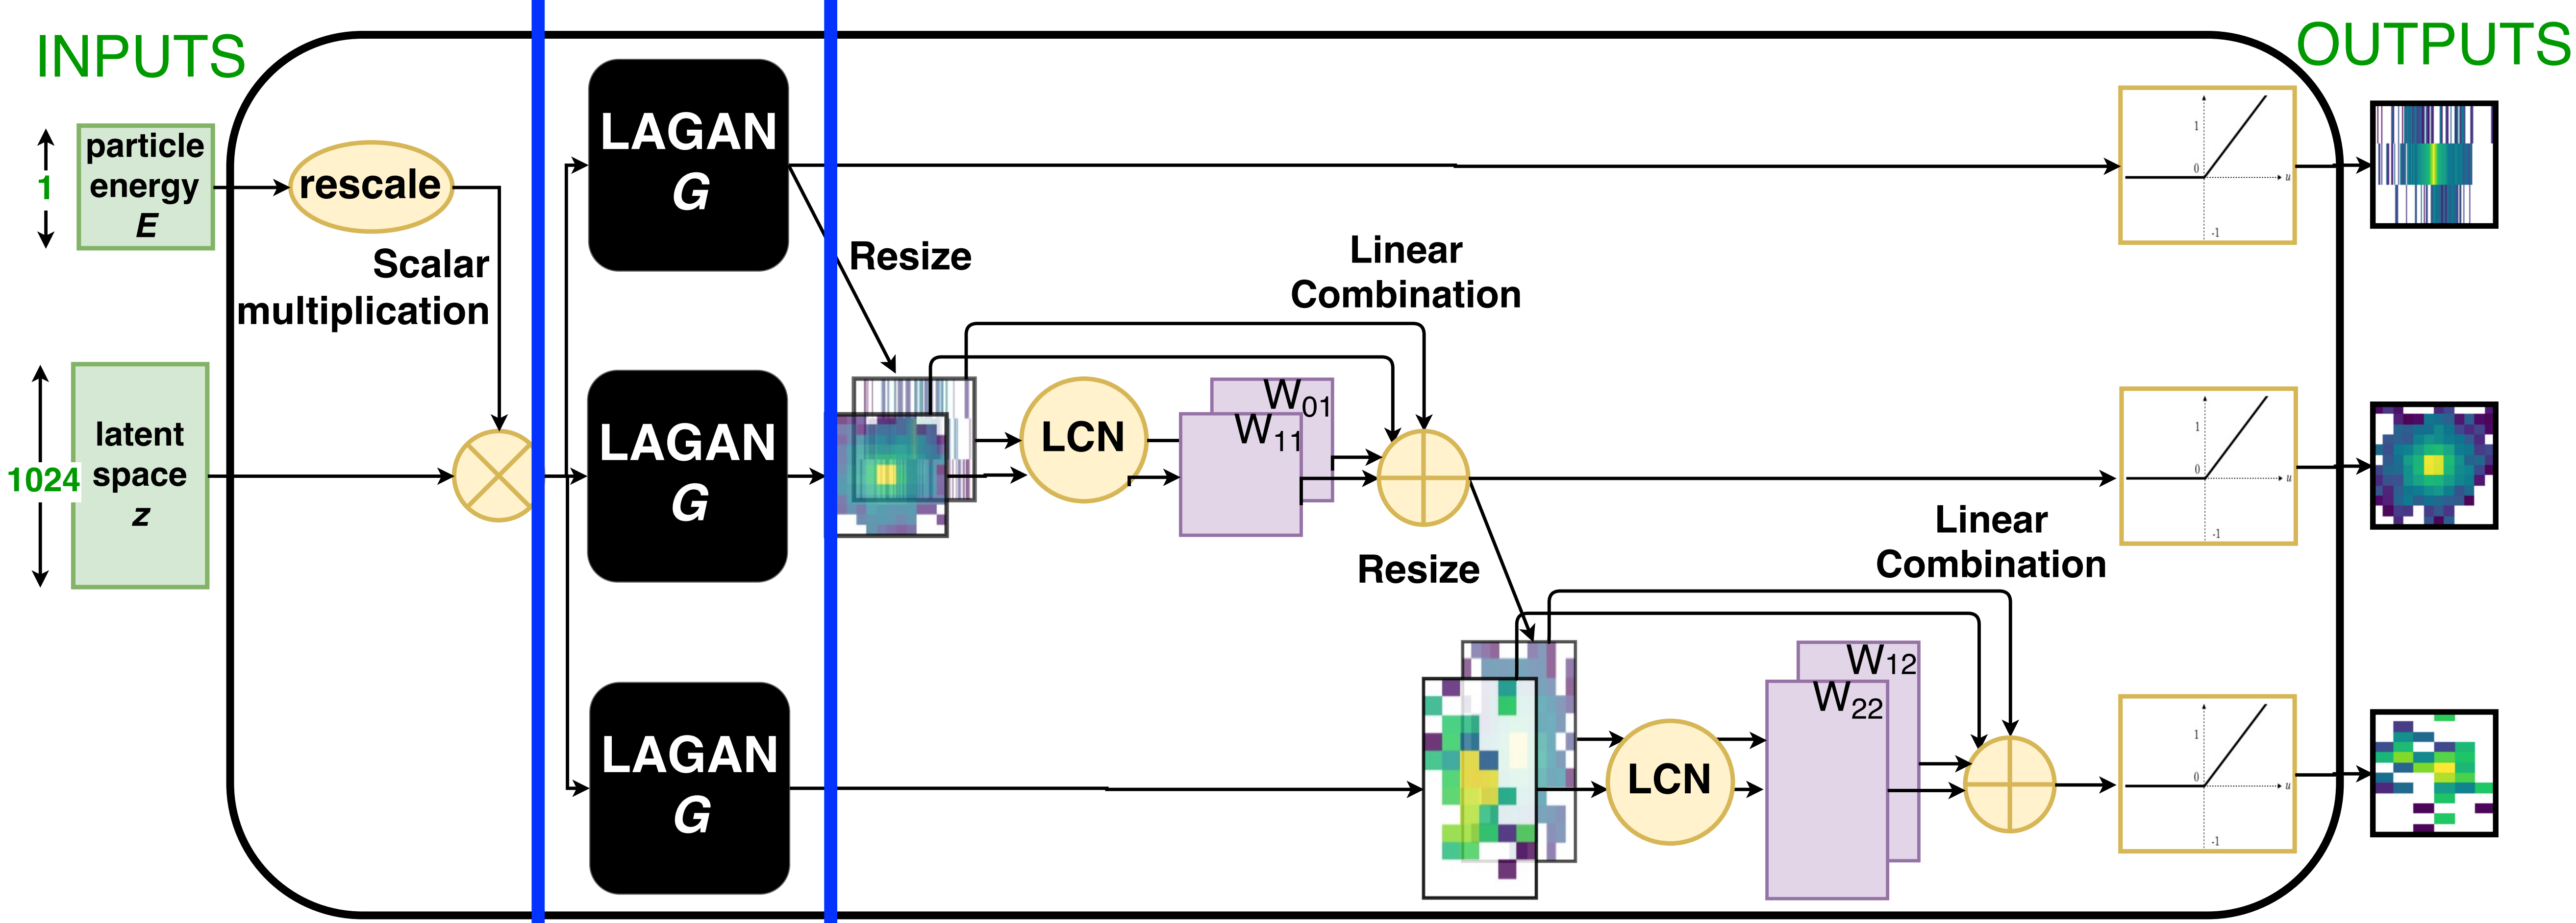

use layer i as input to layer i+1

ReLU to encourage sparsity

M. Paganini, et al., PRL 120 (2018) 042003

Average Images

Geant4

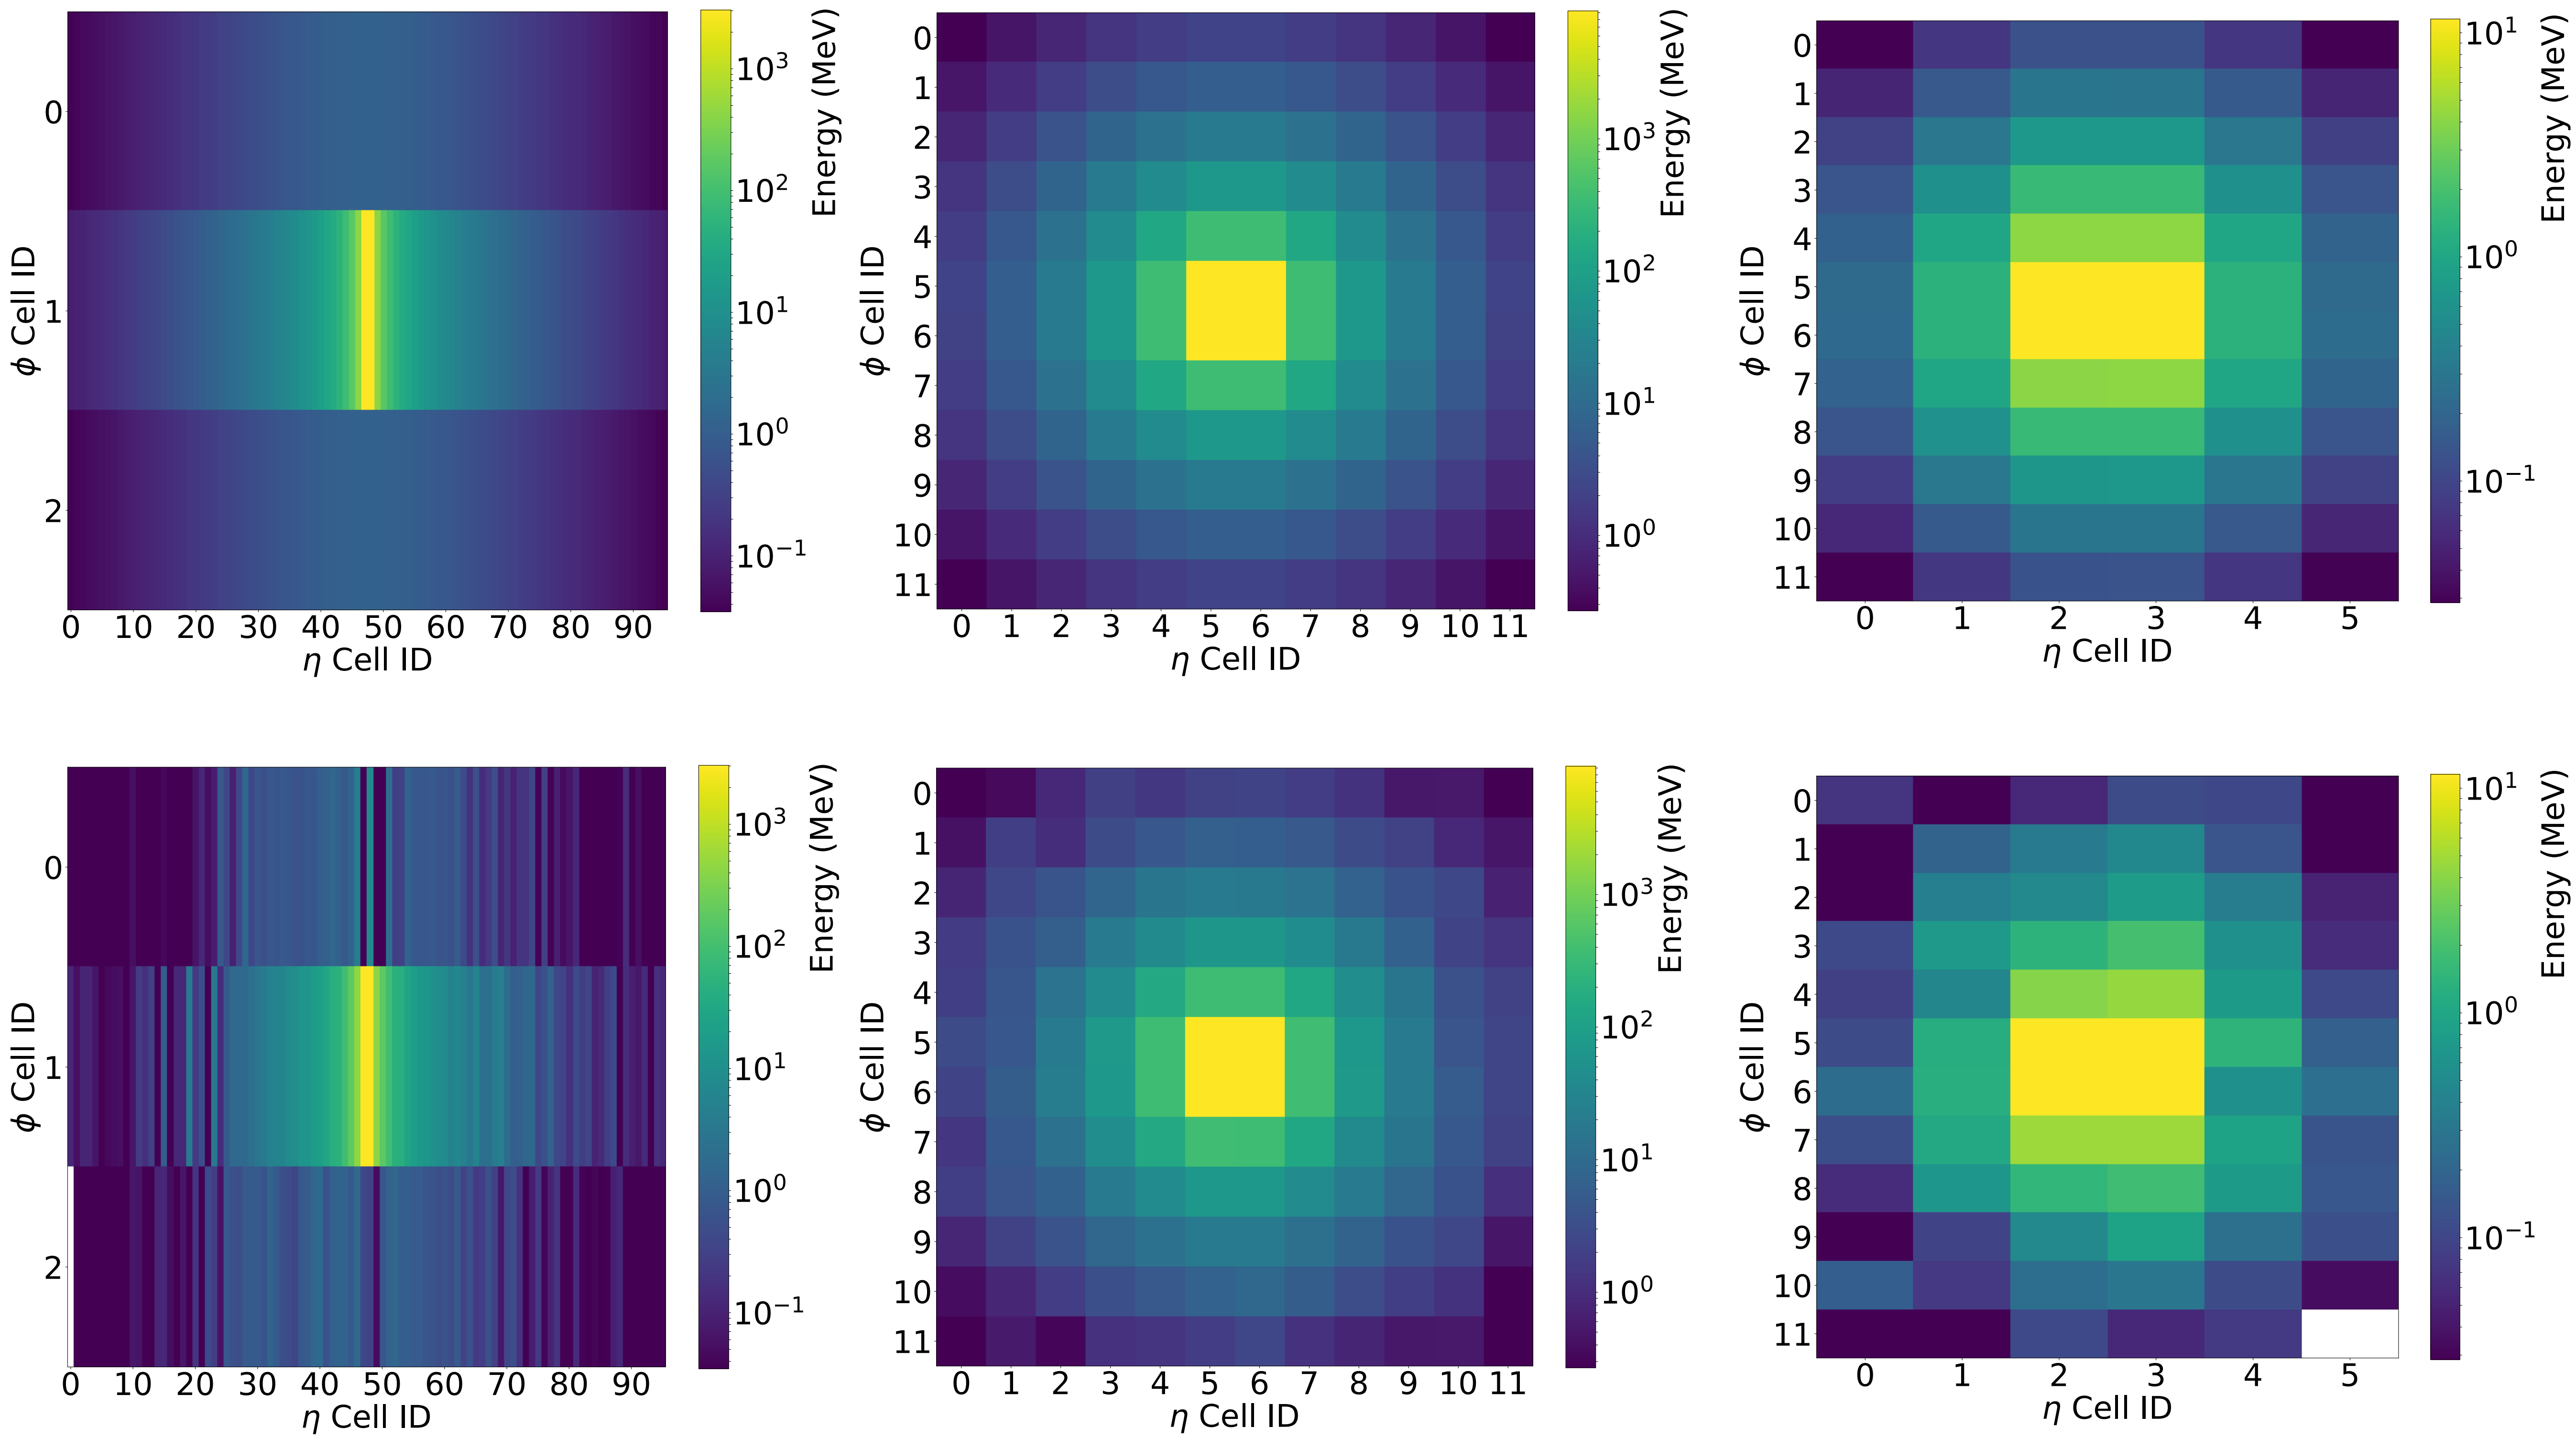

M. Paganini, et al., PRL 120 (2018) 042003

Shower shapes

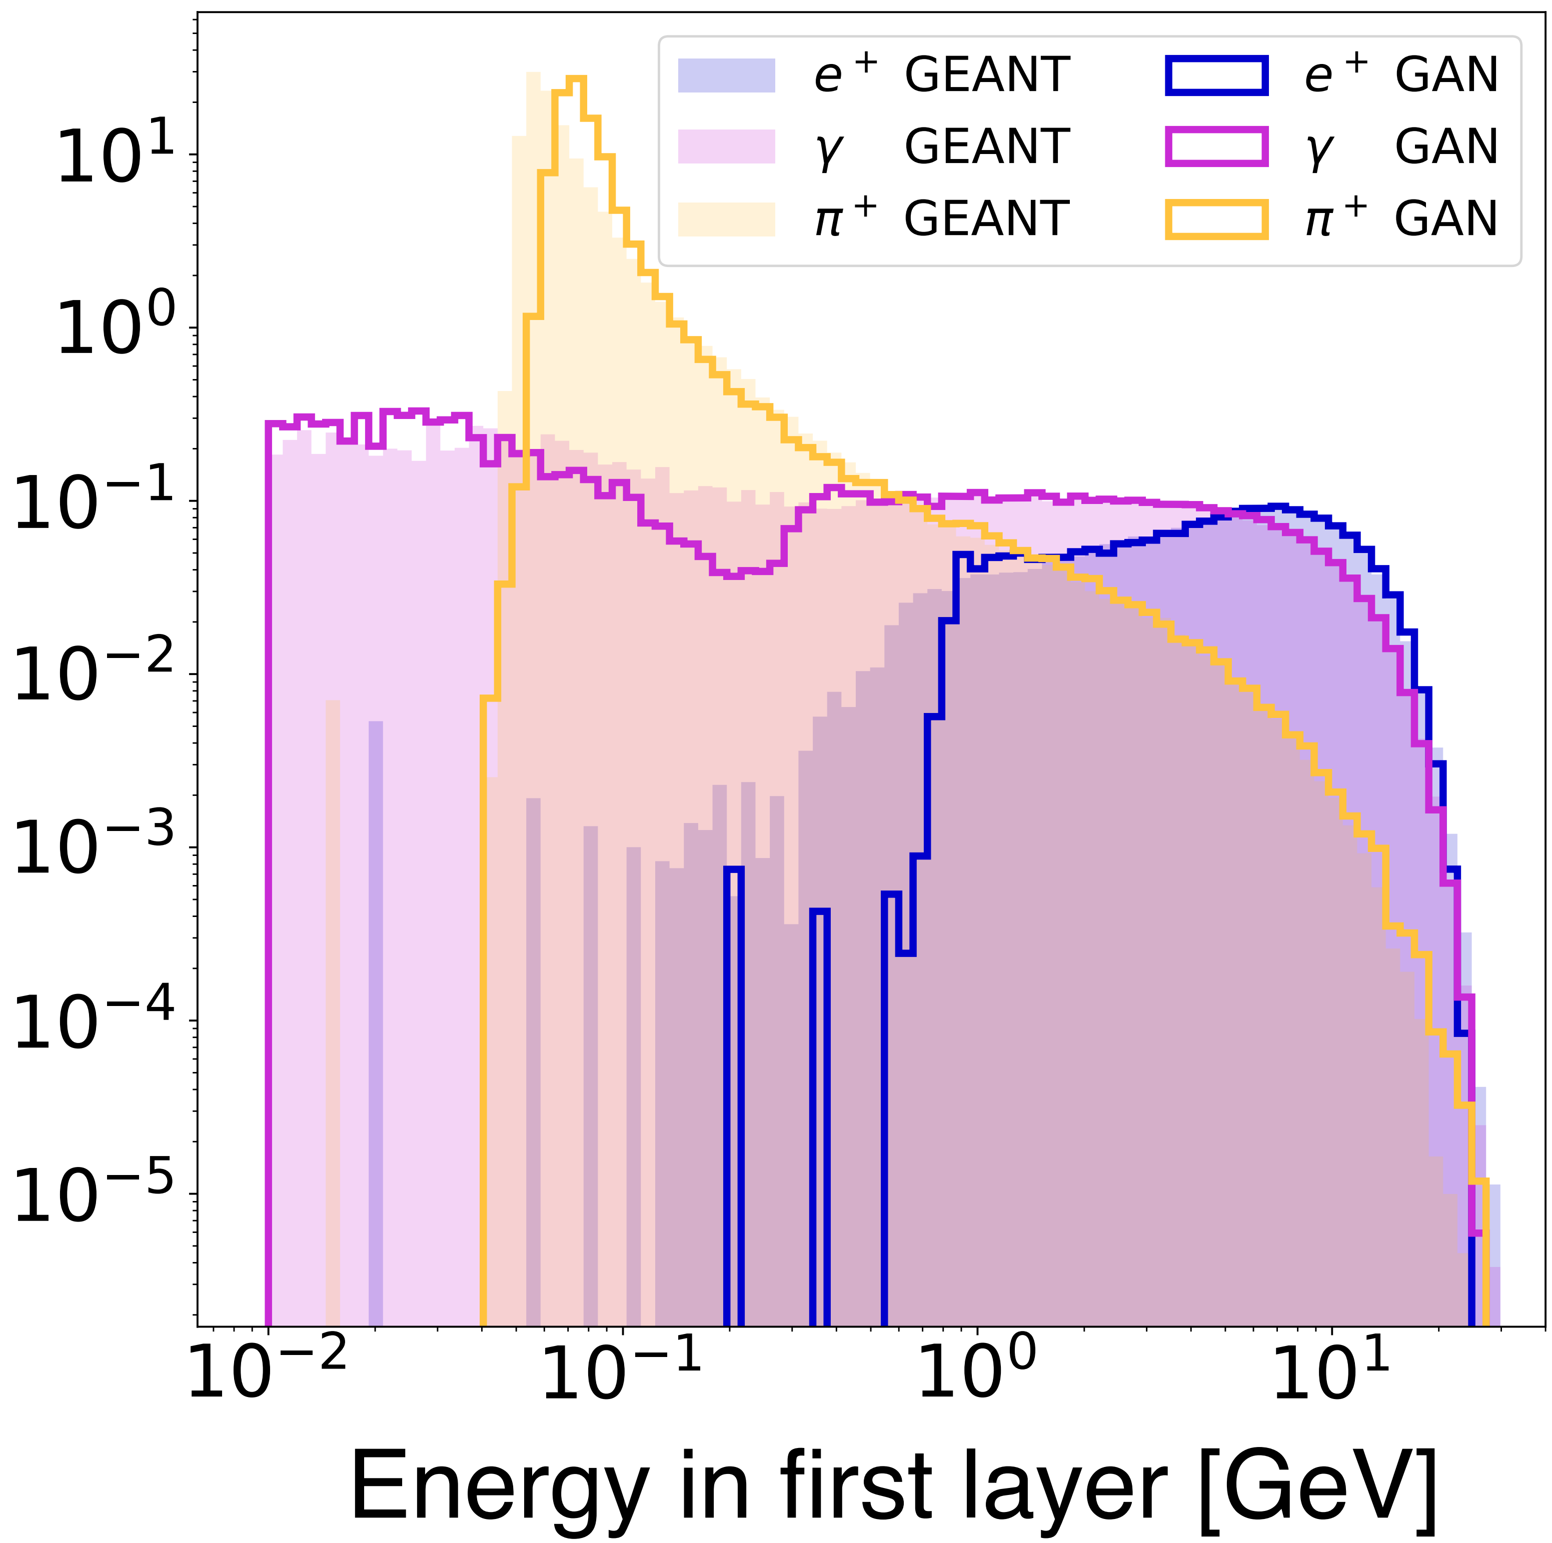

Qualitative agreement; clearly also room for improvement.

In fact, one could add this into the training; for now held out for validation.

Timing

Generation Method	Hardware	Batch Size	milliseconds/shower
GEANT4	CPU	N/A	1772 ←
CALOGAN	CPU	1	13.1
		10	5.11
		128	2.19
		1024	2.03
	GPU	1	14.5
		4	3.68
		128	0.021
		512	0.014
		1024	0.012 ←

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

Conclusions and outlook

The LHC is a unique science tool with extreme challenges related to the data rate: real time / ultra fast algorithms are required.

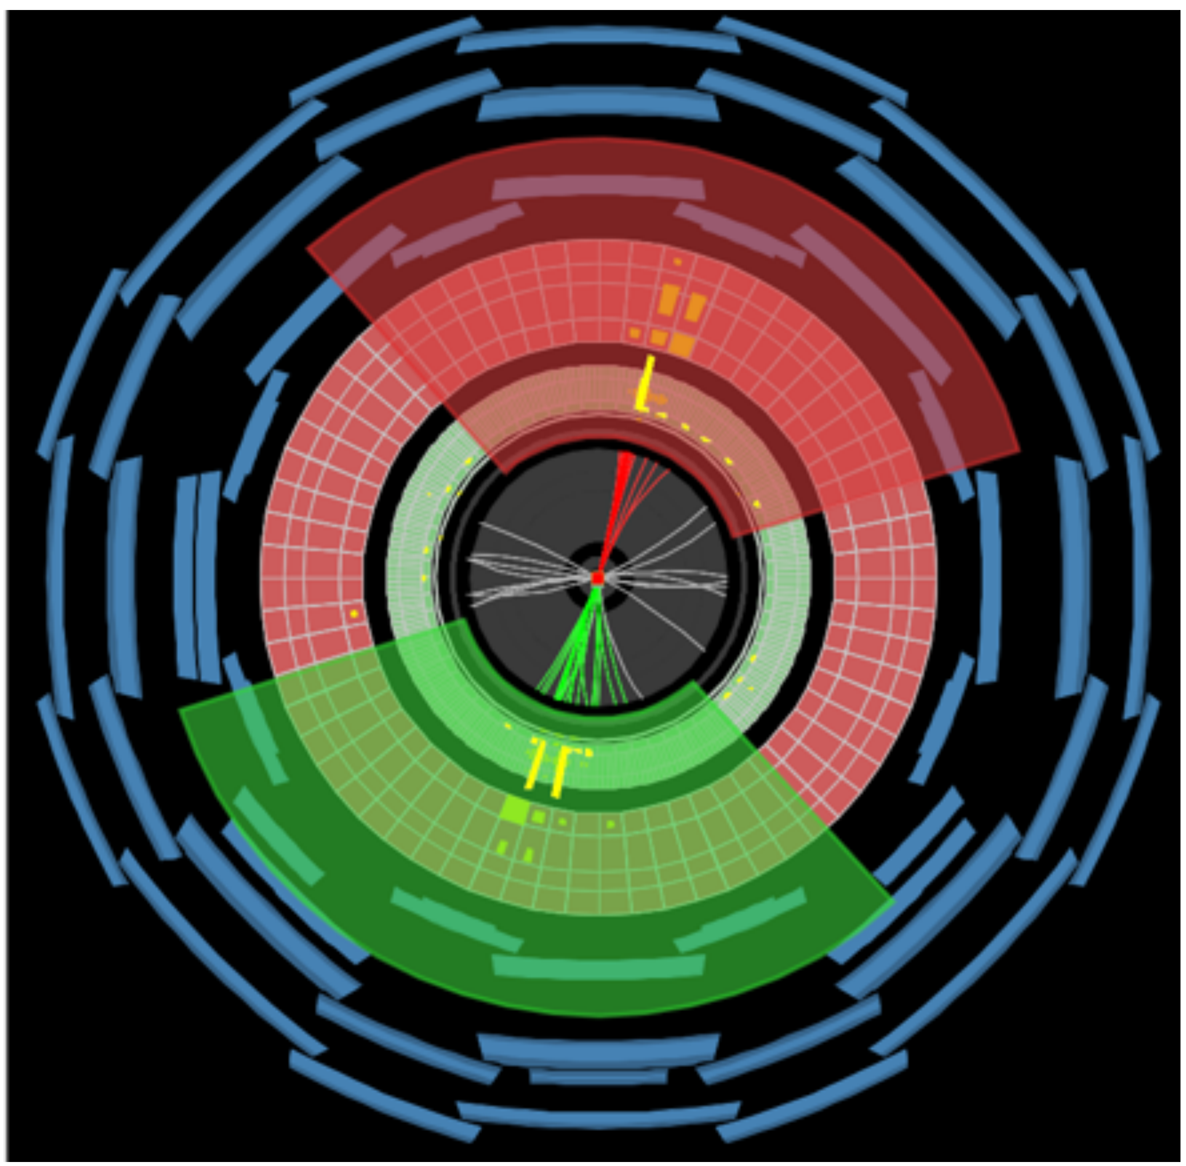

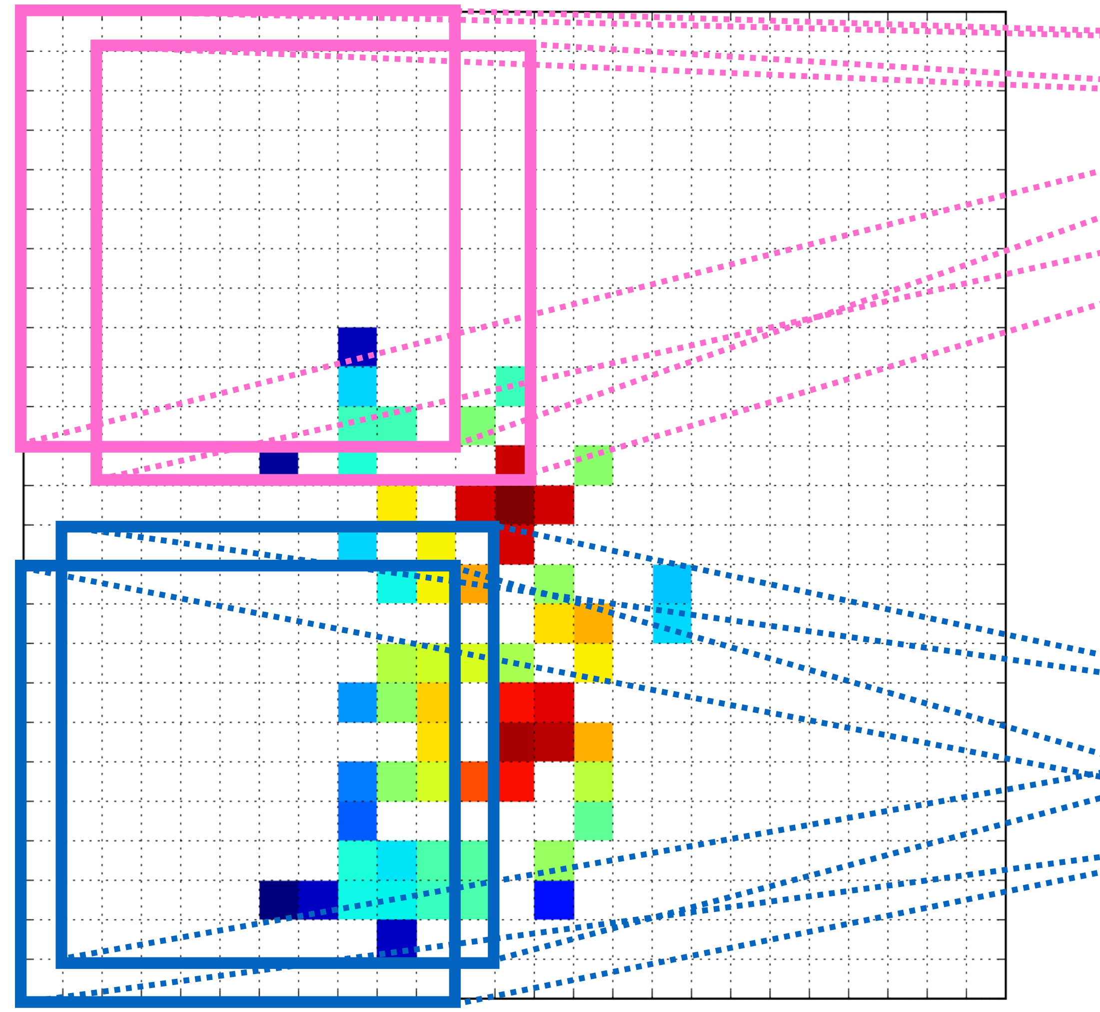

There are many exciting opportunities and ideas for fully exploiting our data
we must make sure no stone is left unturned !

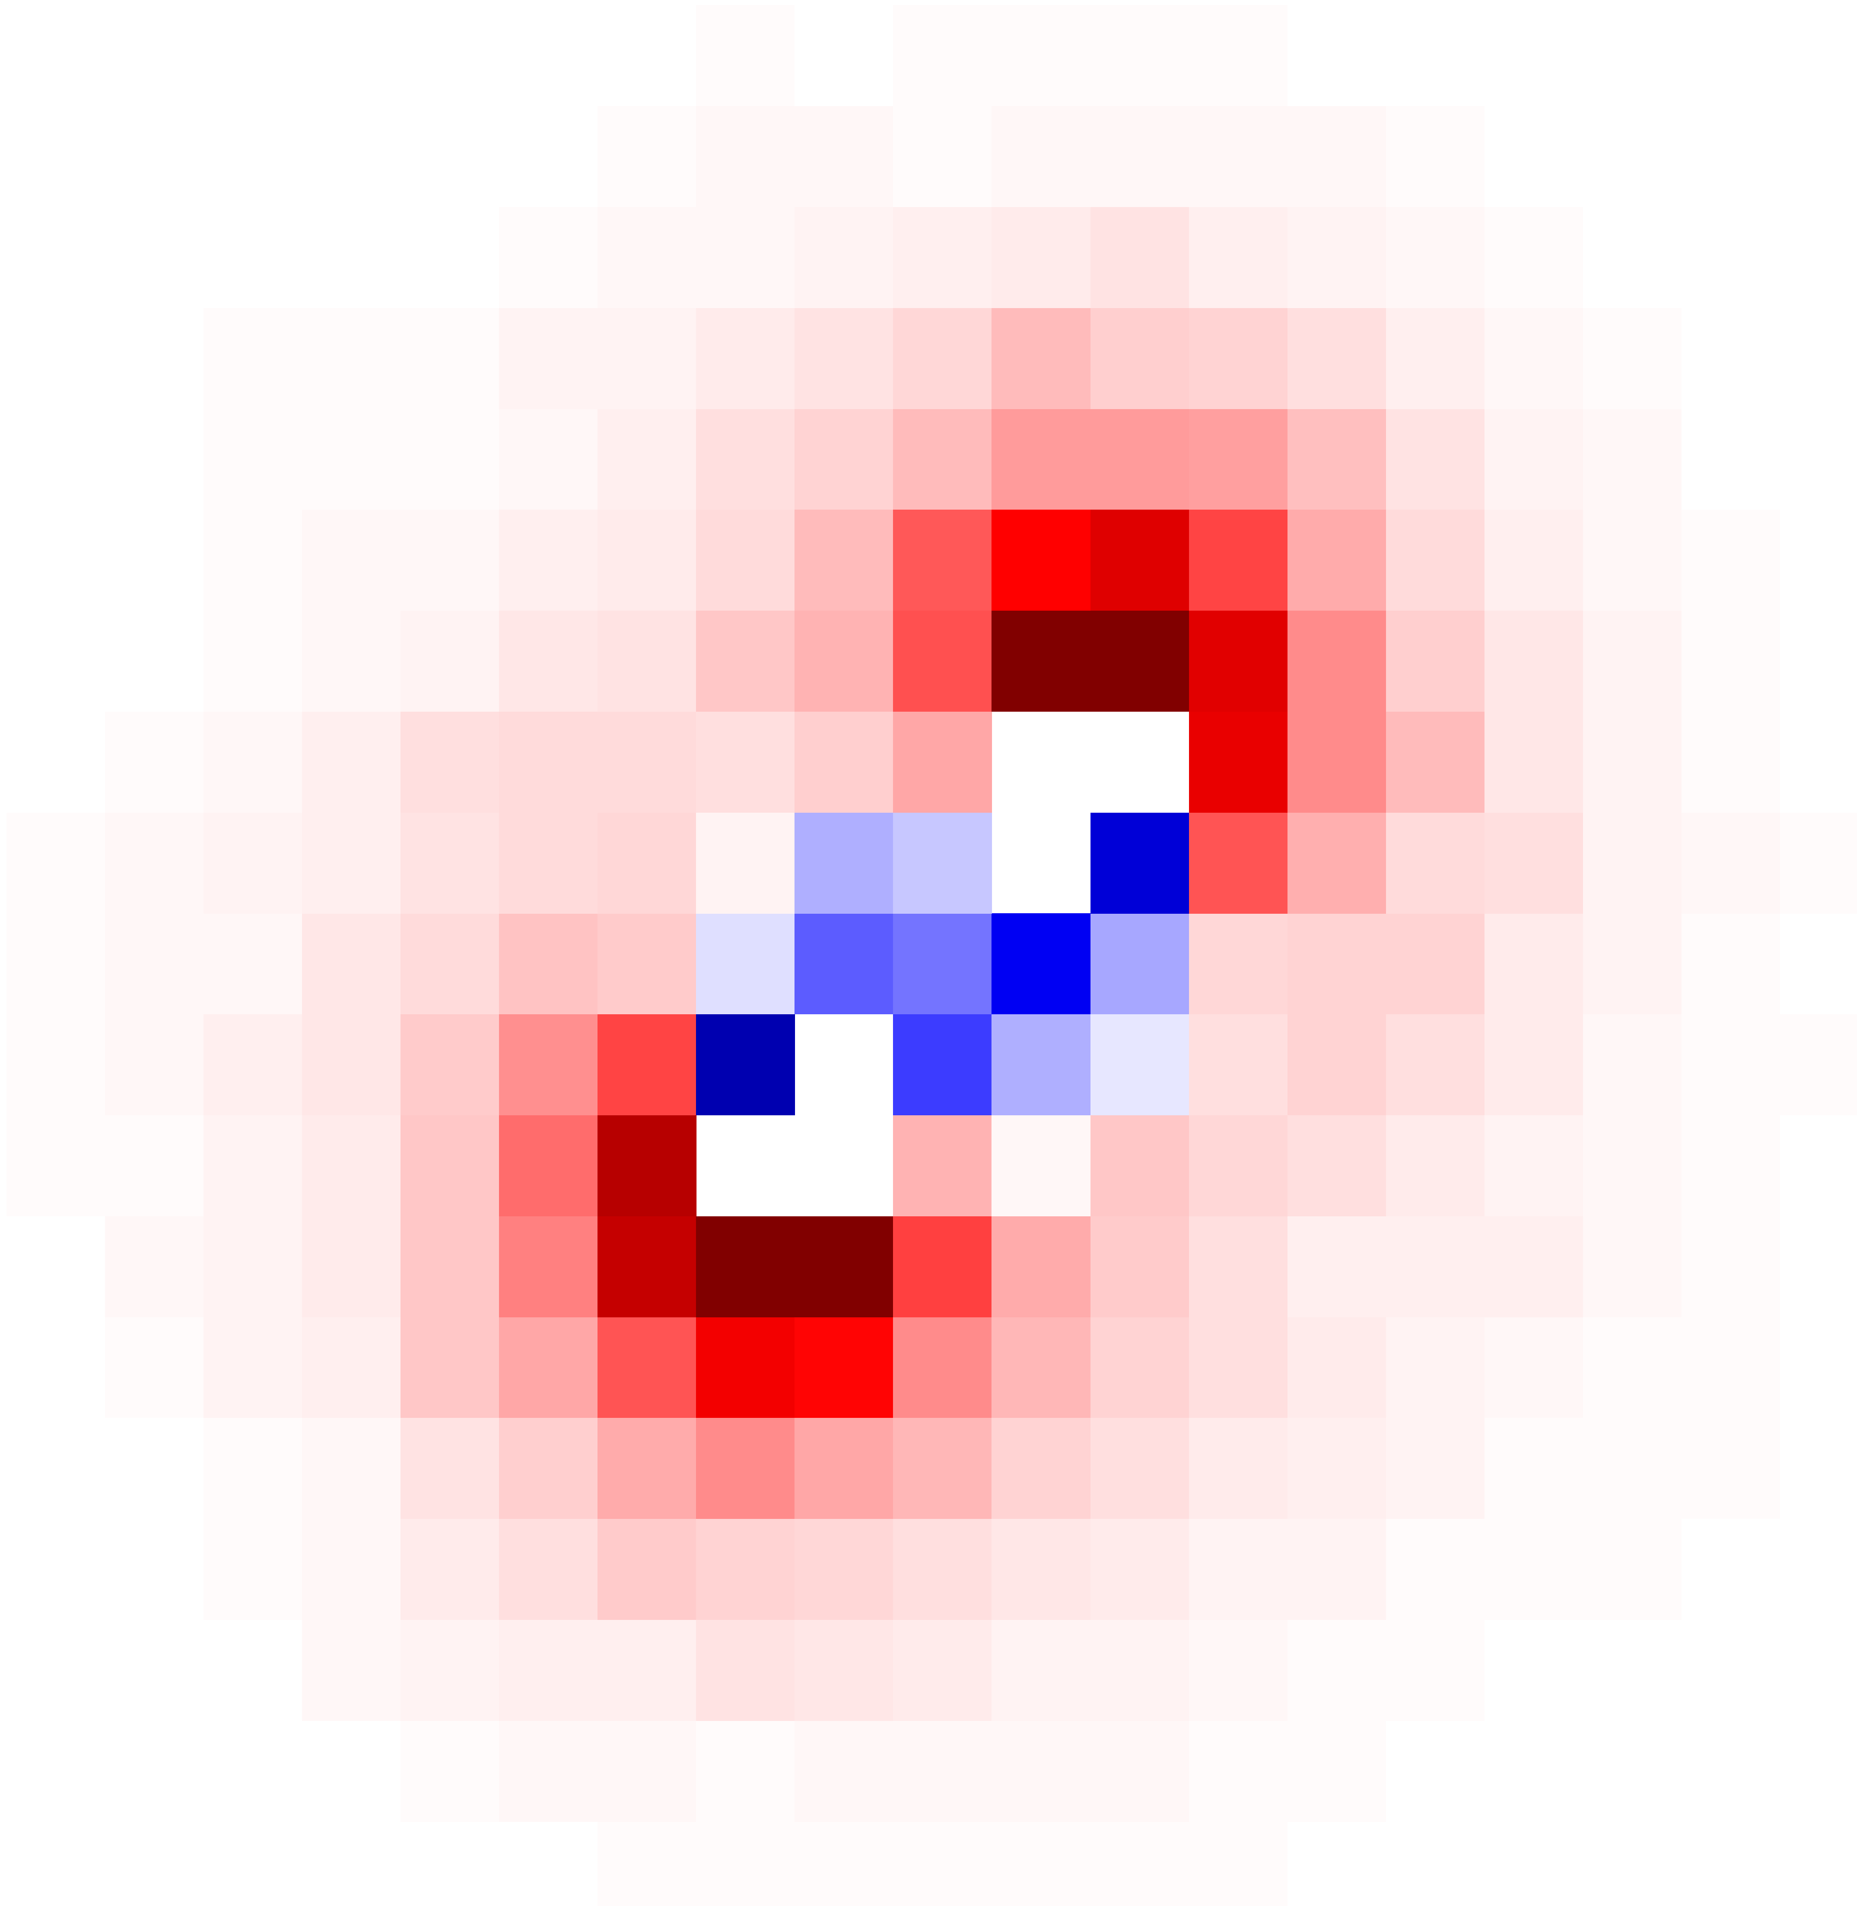

Fin.

Discriminator Network for CaloGAN

help avoid
'mode collapse'

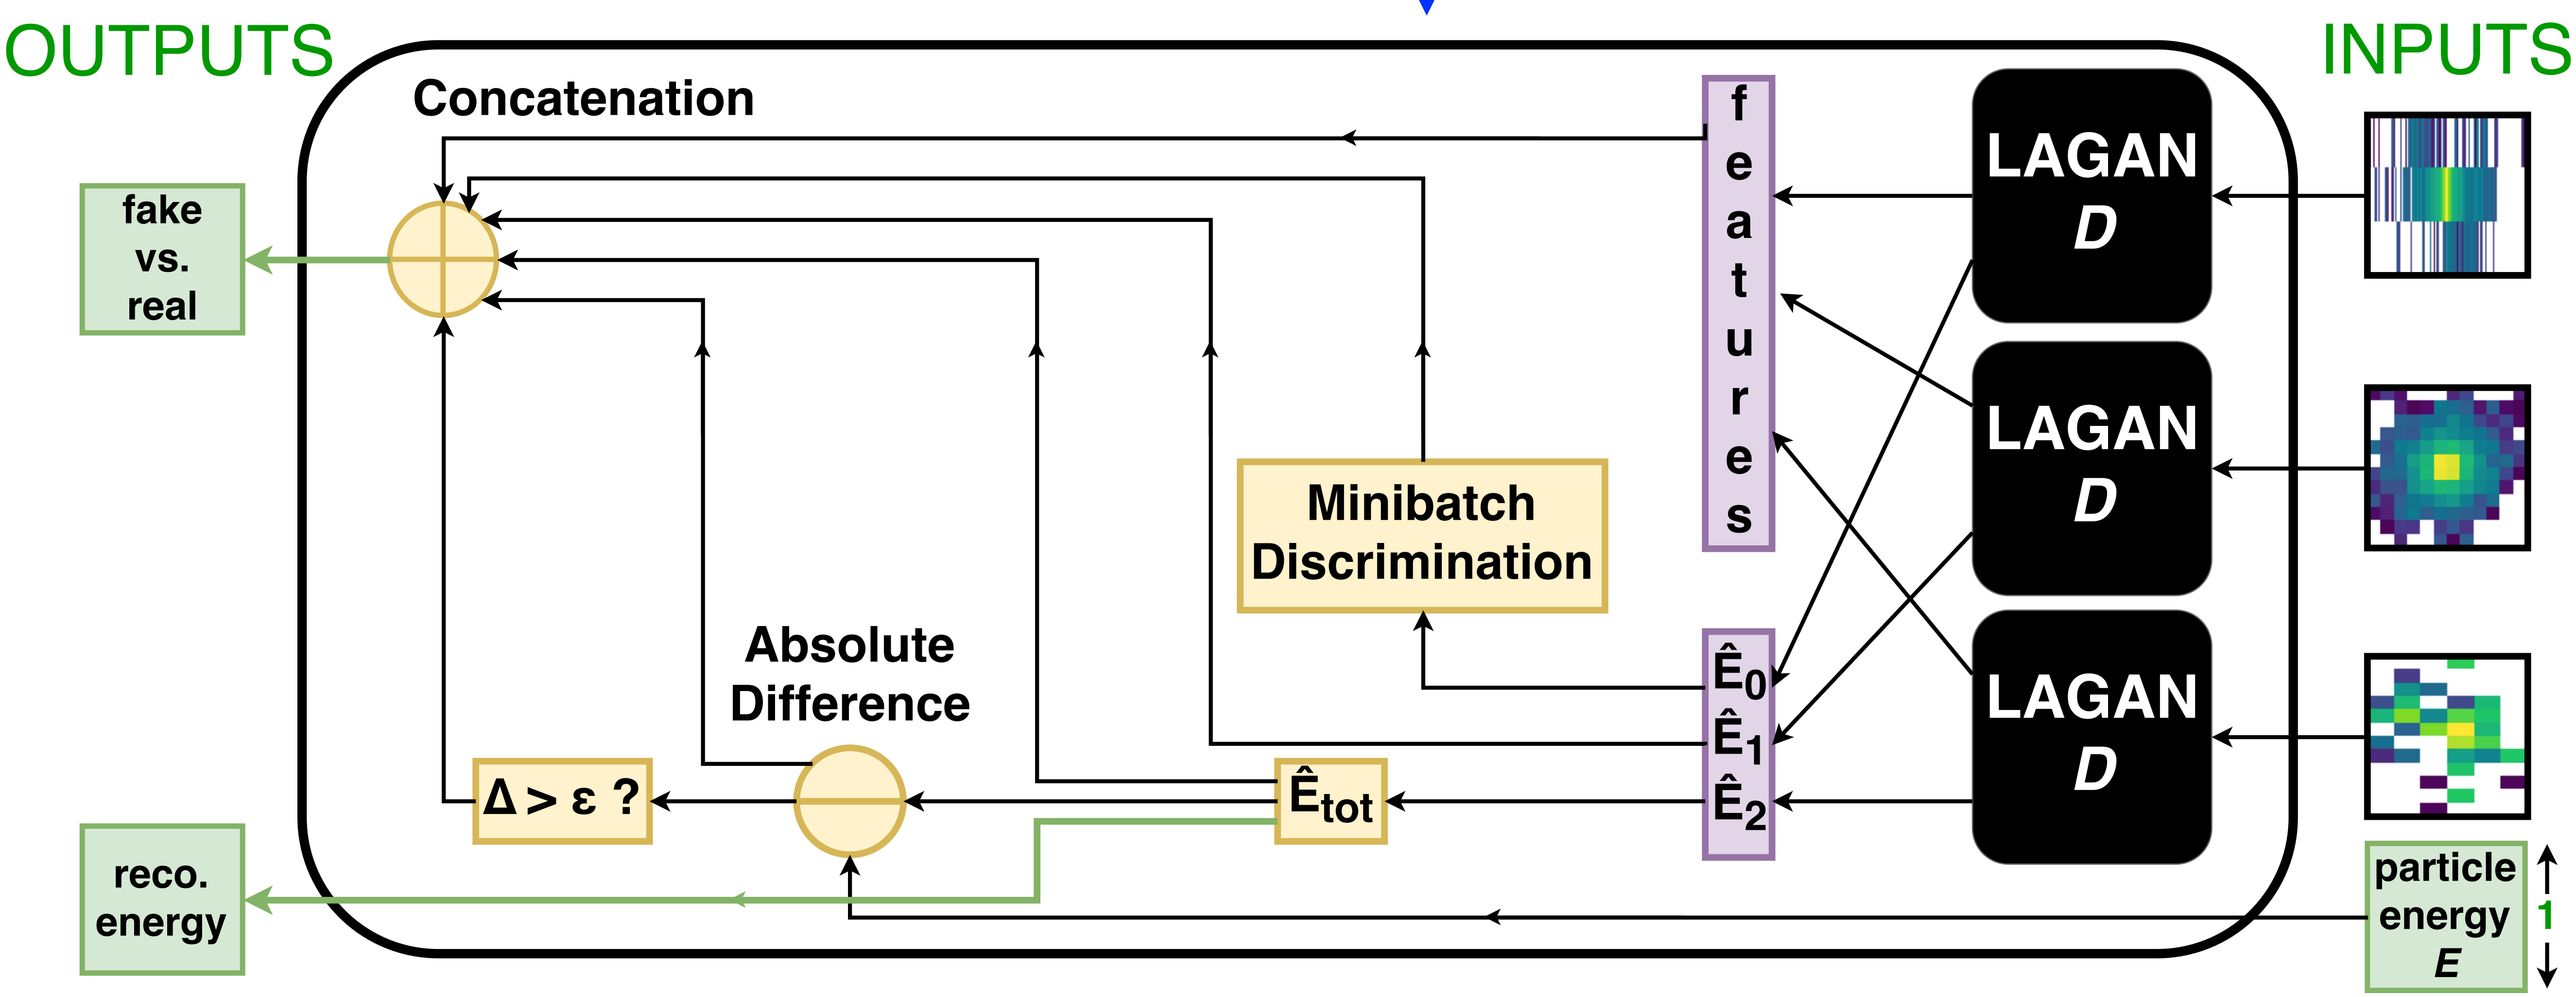

M. Paganini, et al., PRL 120 (2018) 042003

“Overtraining”

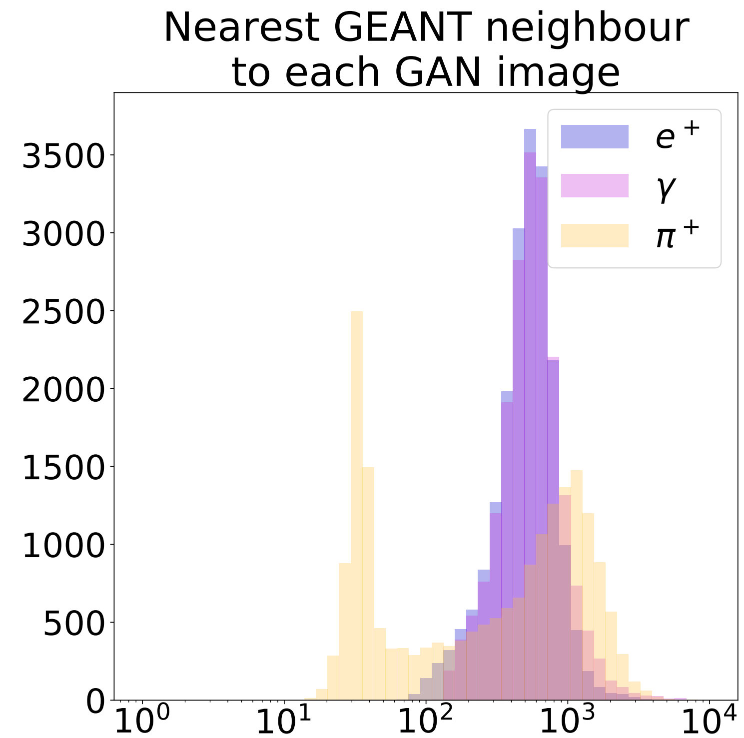

not memorizing

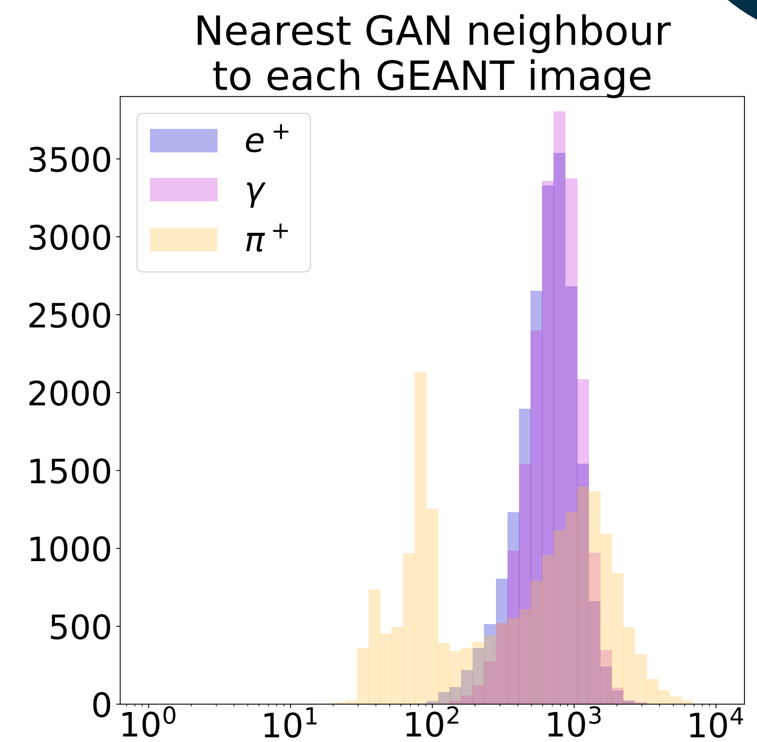

A key challenge for GANs is the diversity of generated images.
This does not seem to be a problem for CaloGAN.

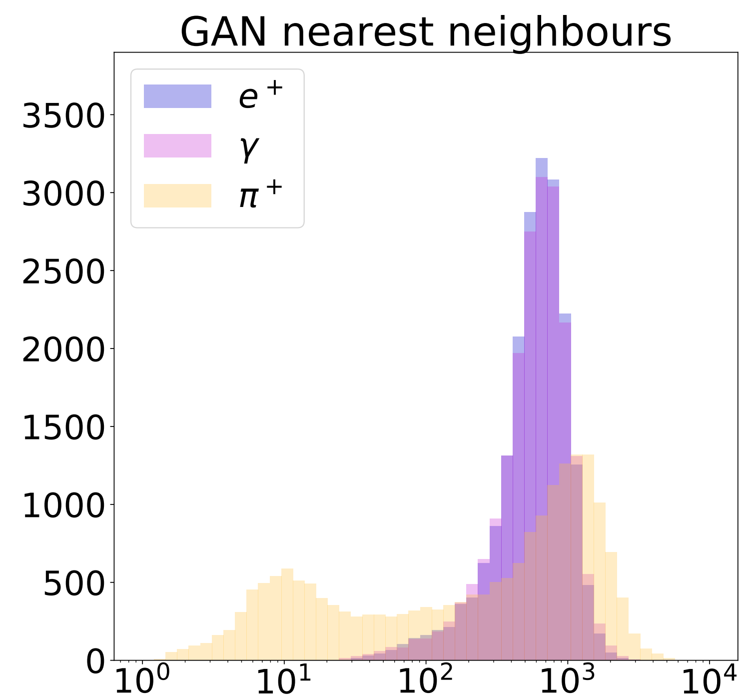

no mode collapse