



3D convolutional GAN for fast simulation

CHEP 2018

F. Carminati, G. Khattak, S. Vallecorsa

July2018

A DL engine for fast simulation

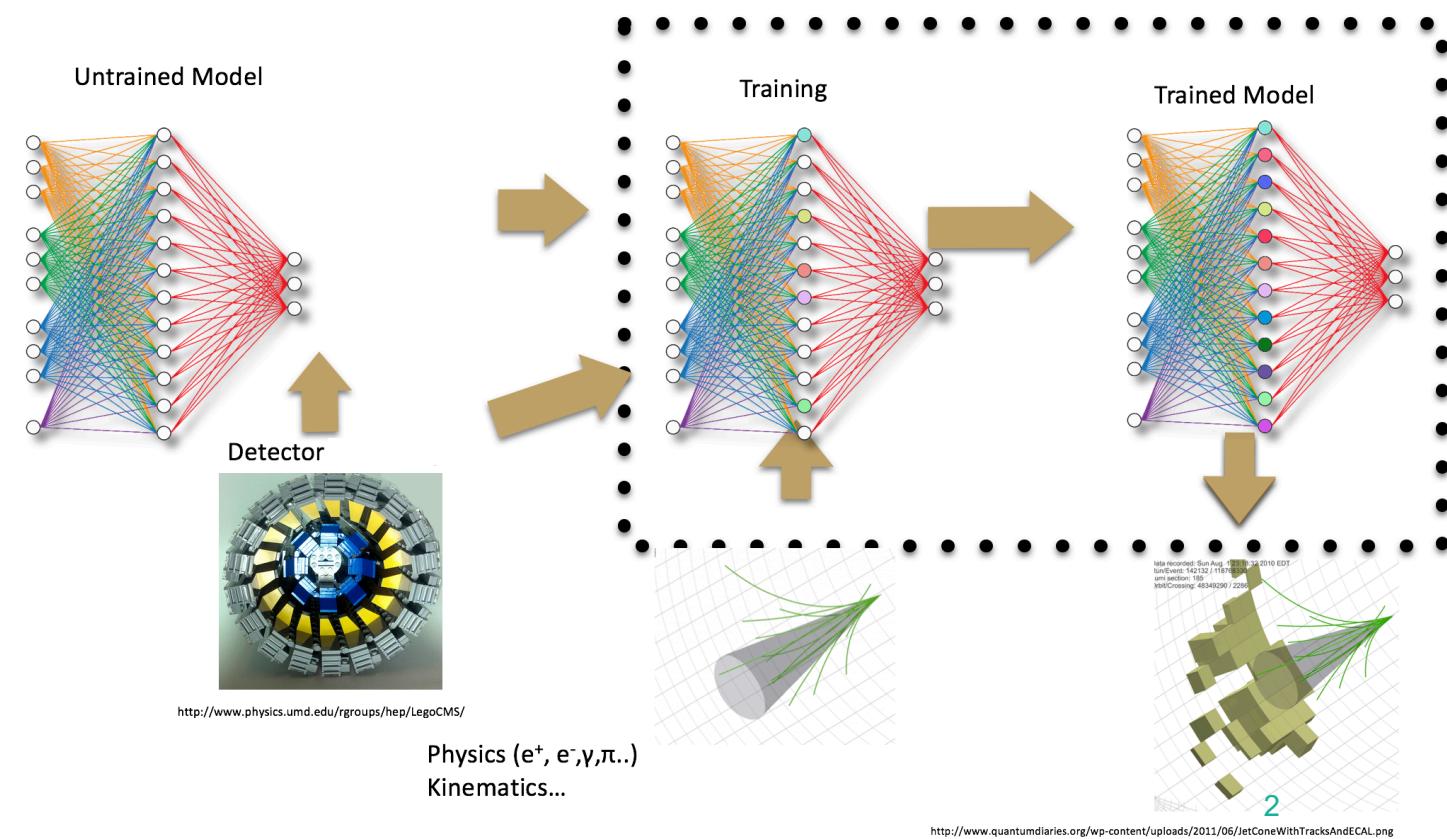
Provide a tool that can be configured and trained for different detectors

Start with time consuming
detectors

Next generation highly
granular calorimeters

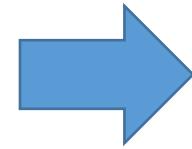
Train on detailed simulation
Test training on real data

Test different models
GAN, RNN



A plan in two steps

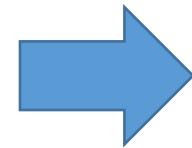
Is generative models output
accurate enough?



Can we sustain the increase in
detector complexity?

- A first proof of concept
- Understand performance and validate accuracy

How generic is this approach?



Can we “adjust” architecture to fit a
larger class of detectors?

- Prove generalisation is possible
- Understand and optimise computing resources

What resources are needed?

Proof of concept, benchmarking and validation

CLIC calorimeter simulation

Data is essentially a 3D image

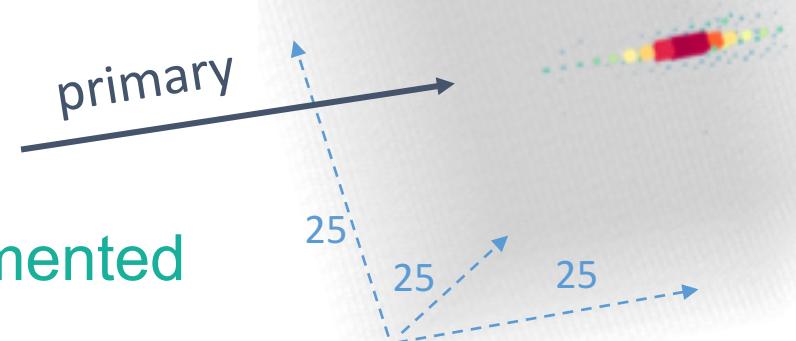
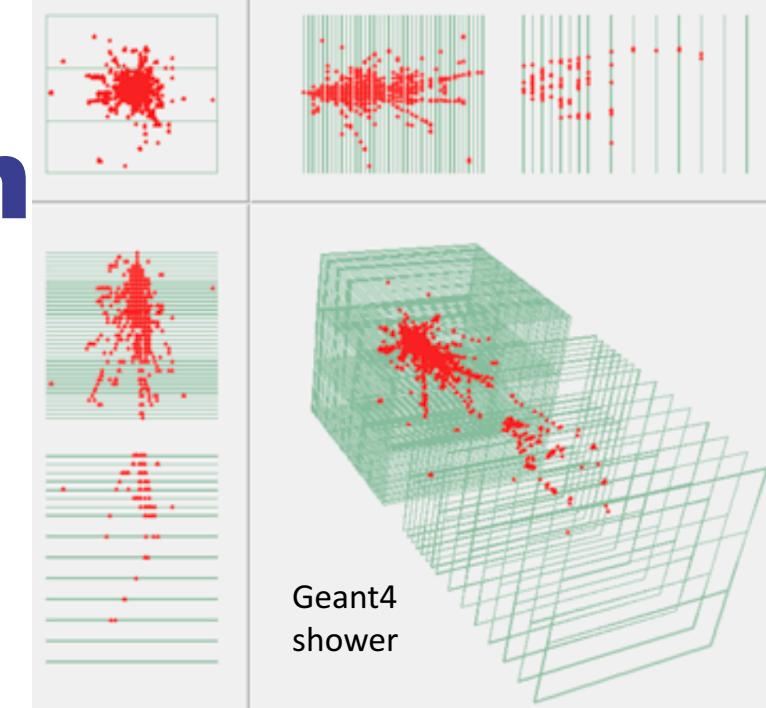
Electromagnetic calorimeter detector design^(*)
(Linear Collider Detector studies)

- 1.5 m inner radius, 5 mm×5 mm segmentation:
25 tungsten absorber layers + silicon sensors

1M single particle samples (e,γ,π)

- Flat energy spectrum (10-500) GeV
- Orthogonal to detector surface
- +/- 10° random incident angle

Highly segmented
Sparse.



(*) <http://cds.cern.ch/record/2254048#>

3D convolutional GAN

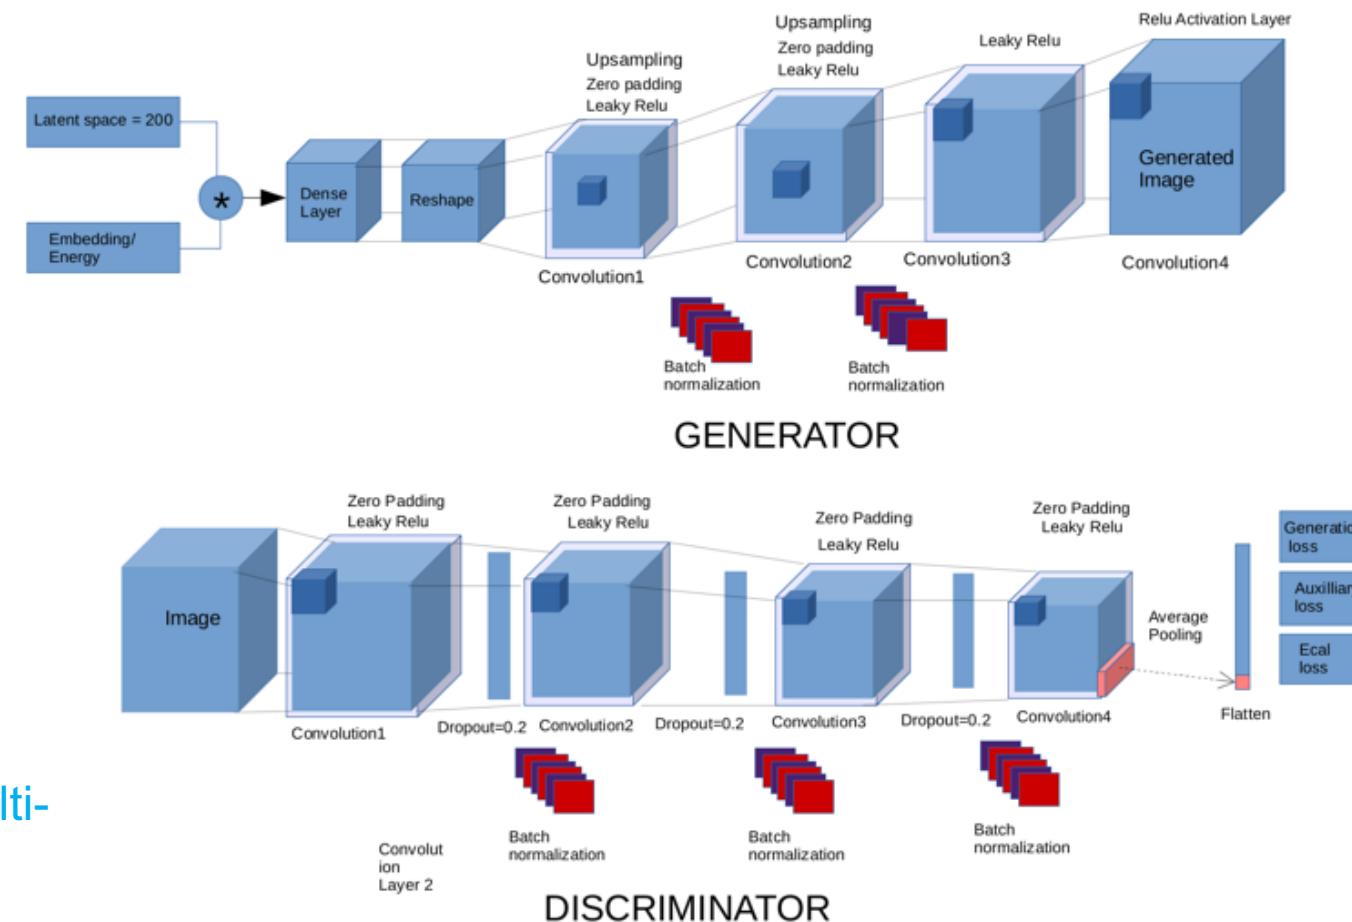
Similar discriminator and generator models

3d convolutions (keep X,Y symmetry)

Condition training on several input variables

Auxiliary regression tasks assigned to the discriminator: cross check

Easily generalisable to multi-class approach (or multi-discriminator approach)



Validation and optimisation

Detailed GAN vs GEANT4 comparison (More than 200 Plots!)

- High level quantities (shower shapes)

- Calorimeter response (single cell response)

- Particle properties (primary particle energy)

Optimisation on

- Network Architecture (Layers, filters, kernels, initialisation)

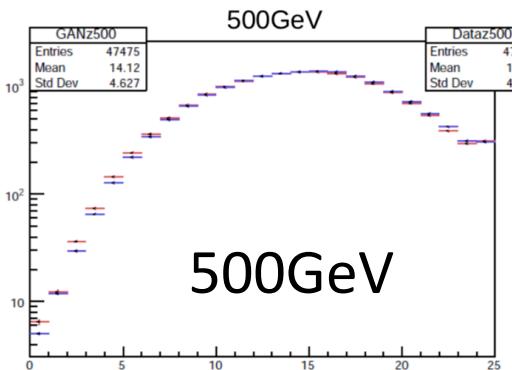
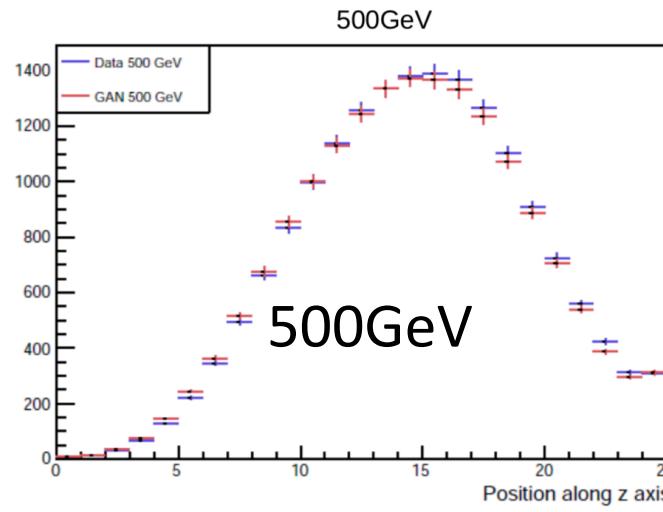
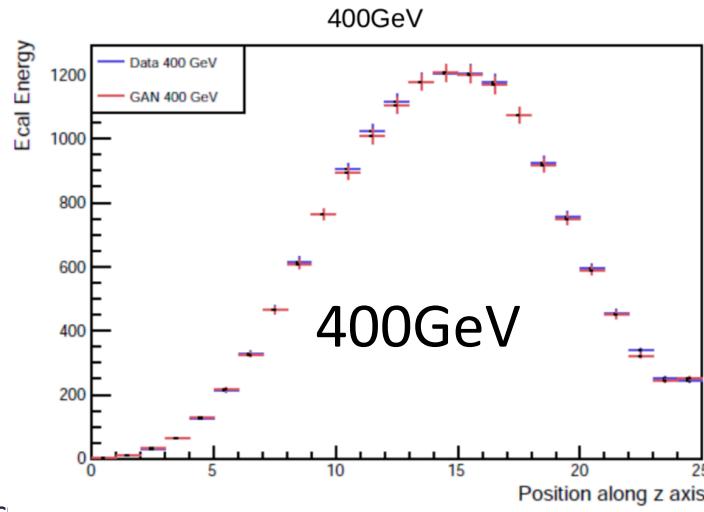
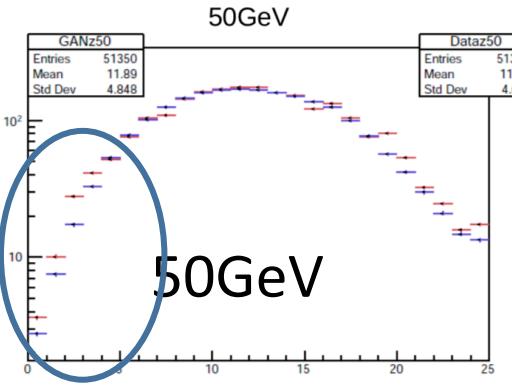
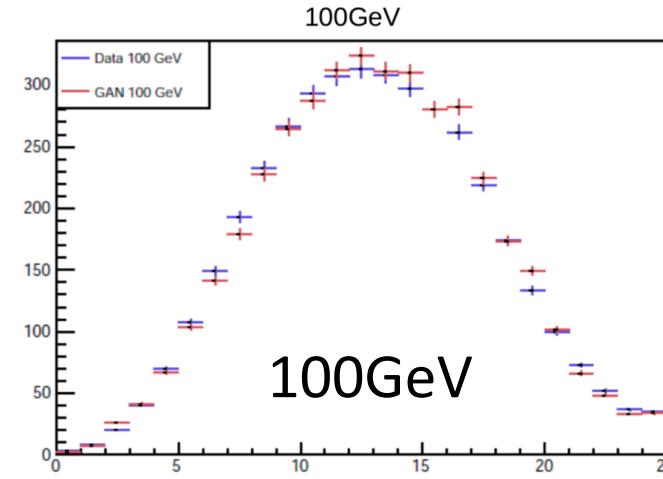
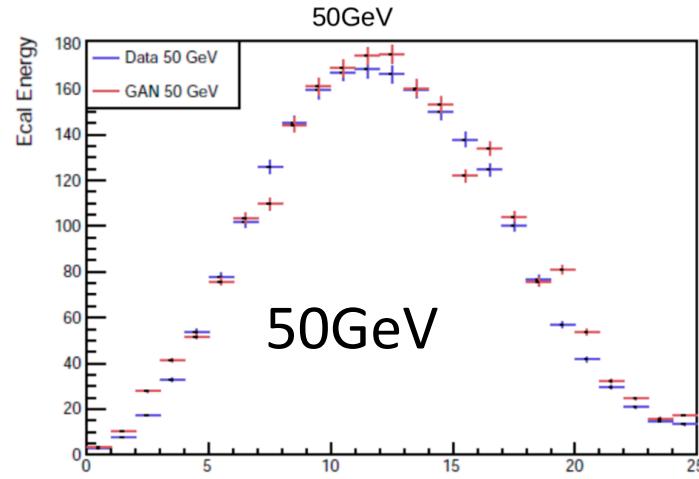
- Losses definition

- Data pre-processing

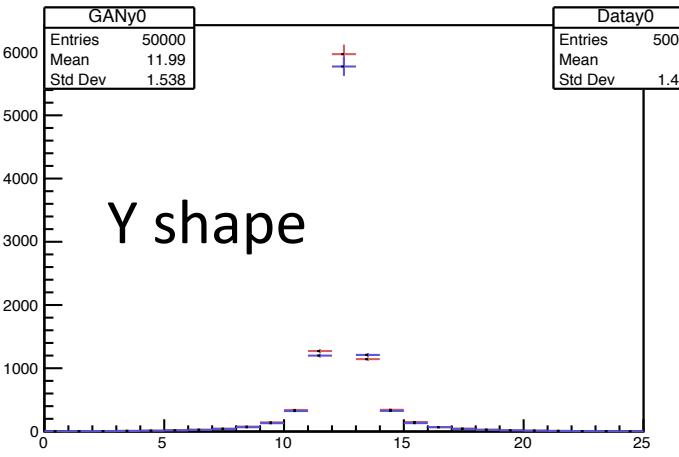
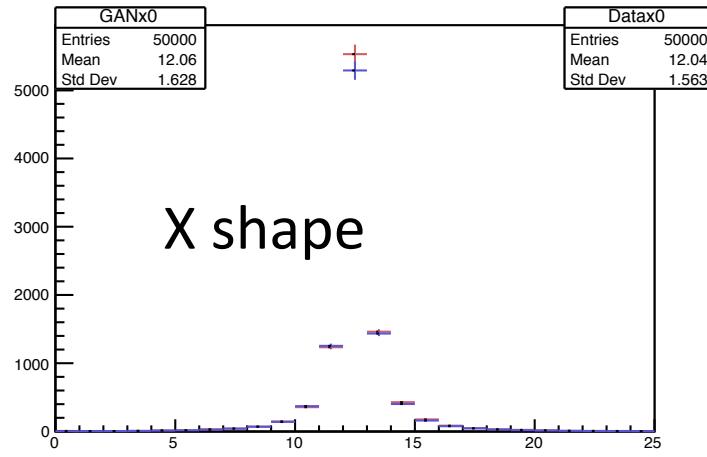
- Rely on GAN losses only !! No physics variable explicitly constrained!

Results agree within a few % to Geant4 (**labelled “DATA” in next slides ☺**)

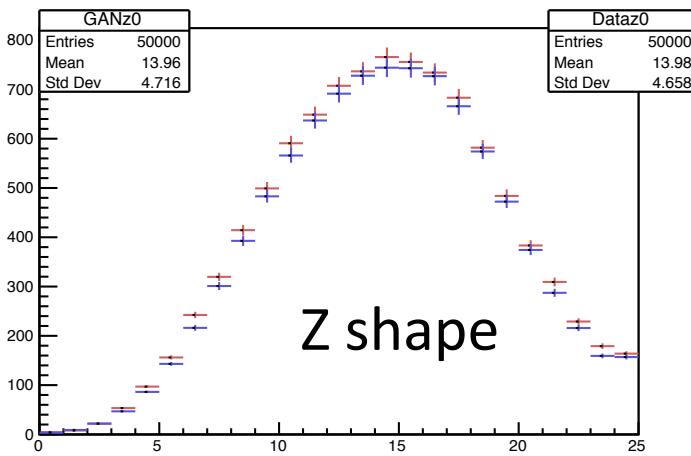
Electrons shower shapes



Neutral Pions

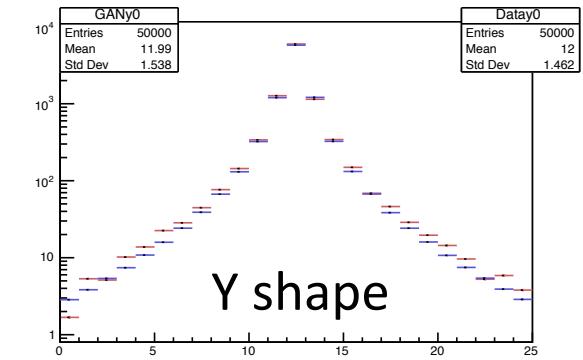
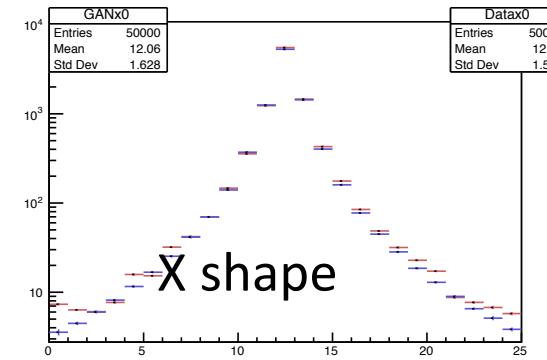


10-500 GeV

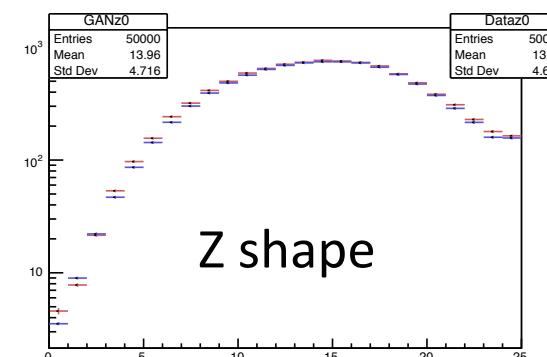


— Data
— GAN

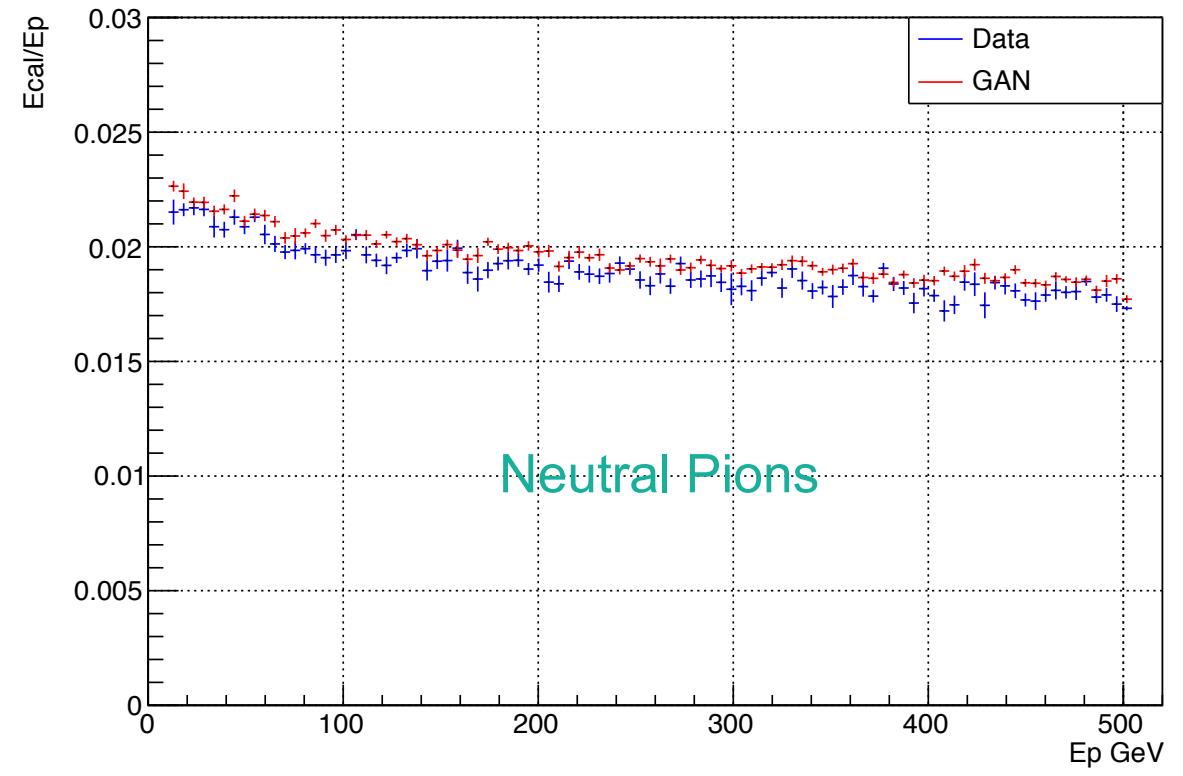
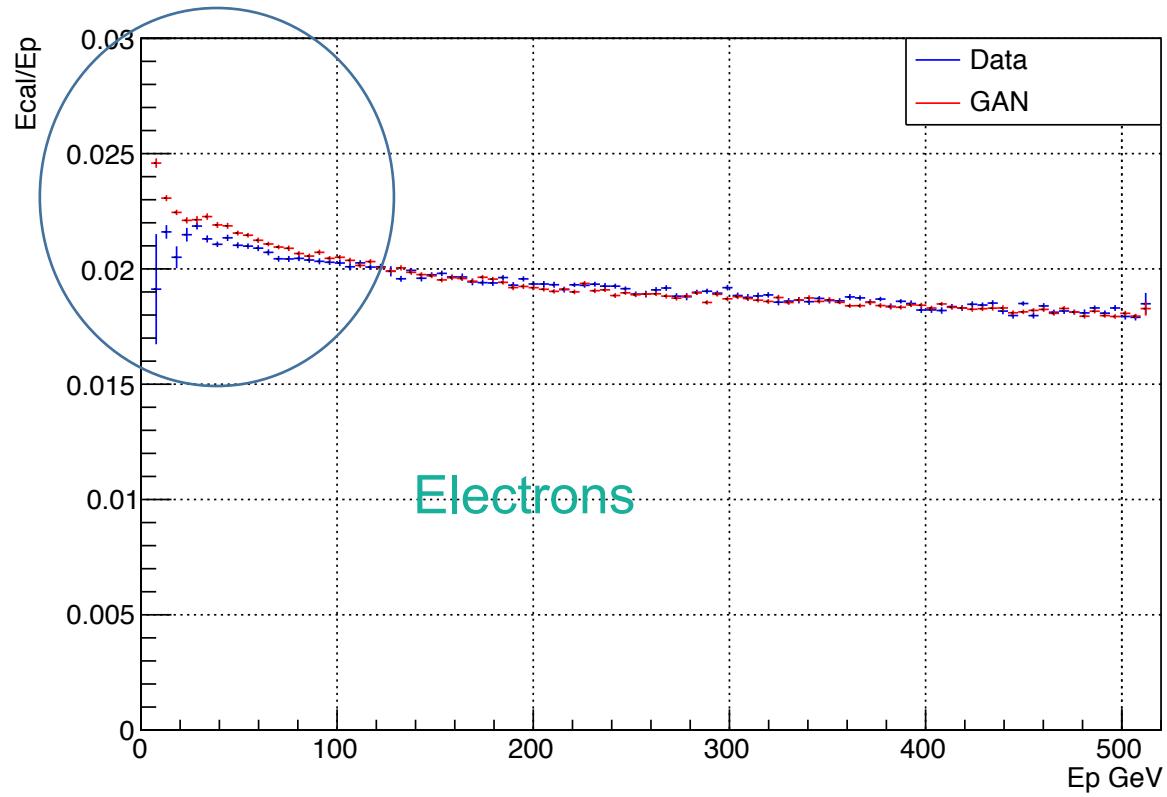
Log scale



— Data
— GAN



Calorimeter sampling fraction



GAN seems to slightly overestimate
slightly neutral pions energy deposits

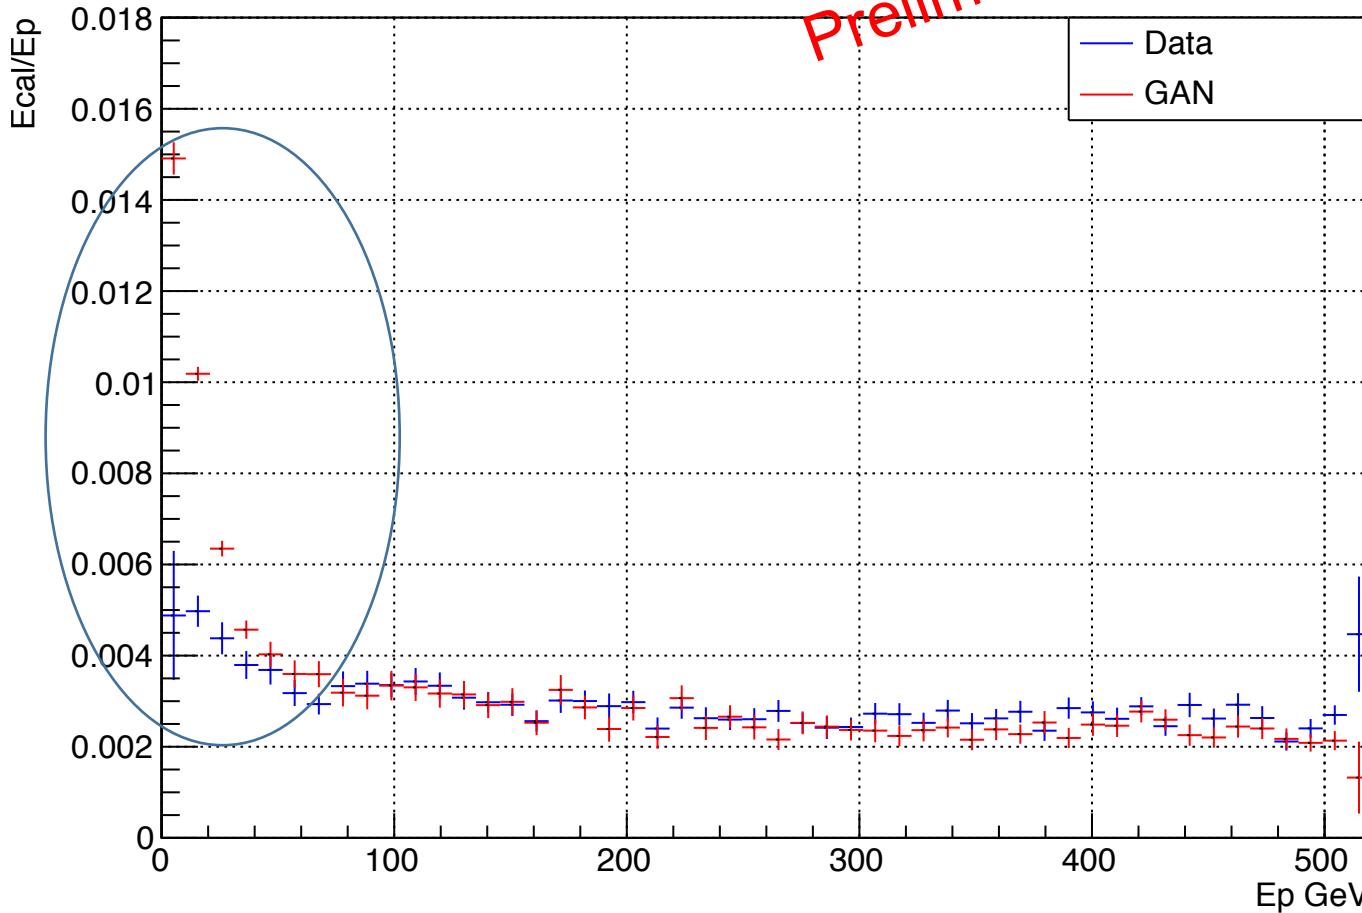
Generalisation & Computing resources

Distributed training

Hyper-parameter scans

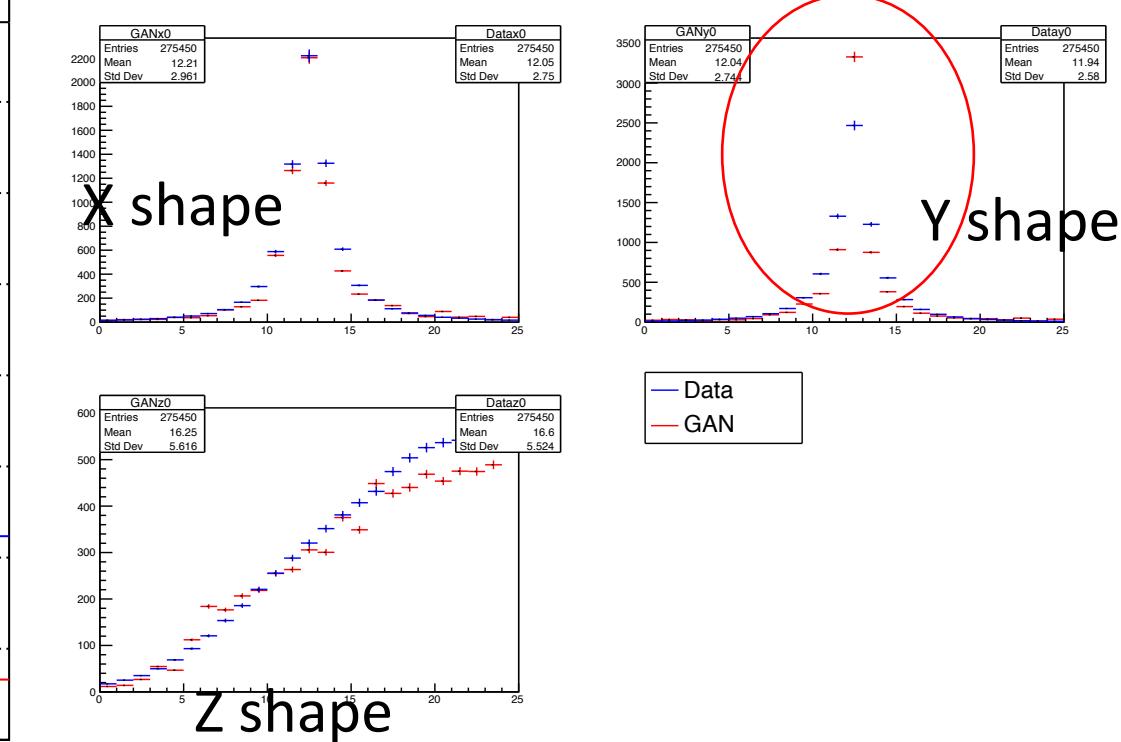
Generalisation

Charged Pions



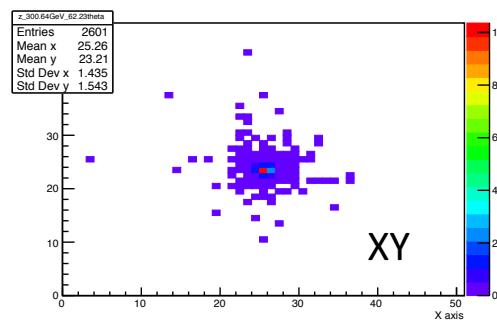
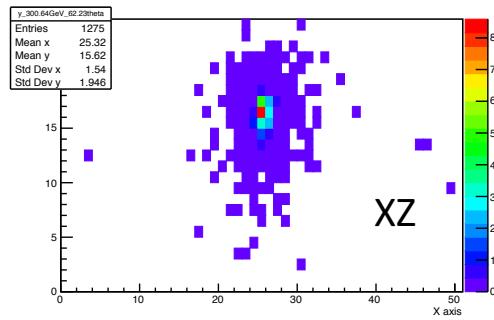
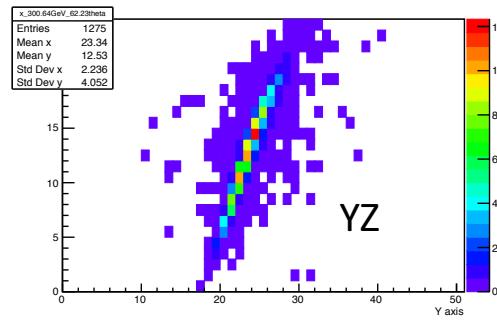
Charged pions have small energy deposits

Energy showers are delayed along Z



Generalisation

Variable angle sample



Preliminary

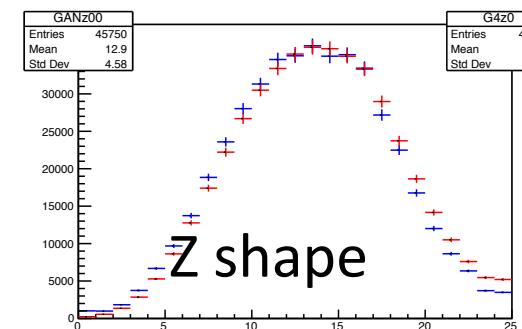
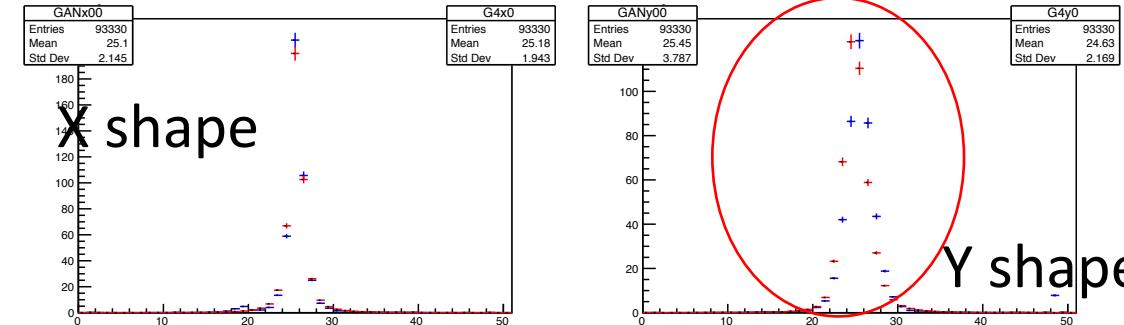
Adjust convolution parameters to improve energy description vs angle

Minimal architecture changes



Electrons enter the calorimeter with a 60° - 120° angle range

Wider/asymmetric image size (51x51x25):



Energy inclusive shower shapes (60° angle)

Computing performance

Distributed training is needed

Inference:

Geant4: 17 s/particle **vs** 3DGAN: 7 ms/particle

→ speedup factor > 2500!!

Training:

45 min /epoch on Tesla P100

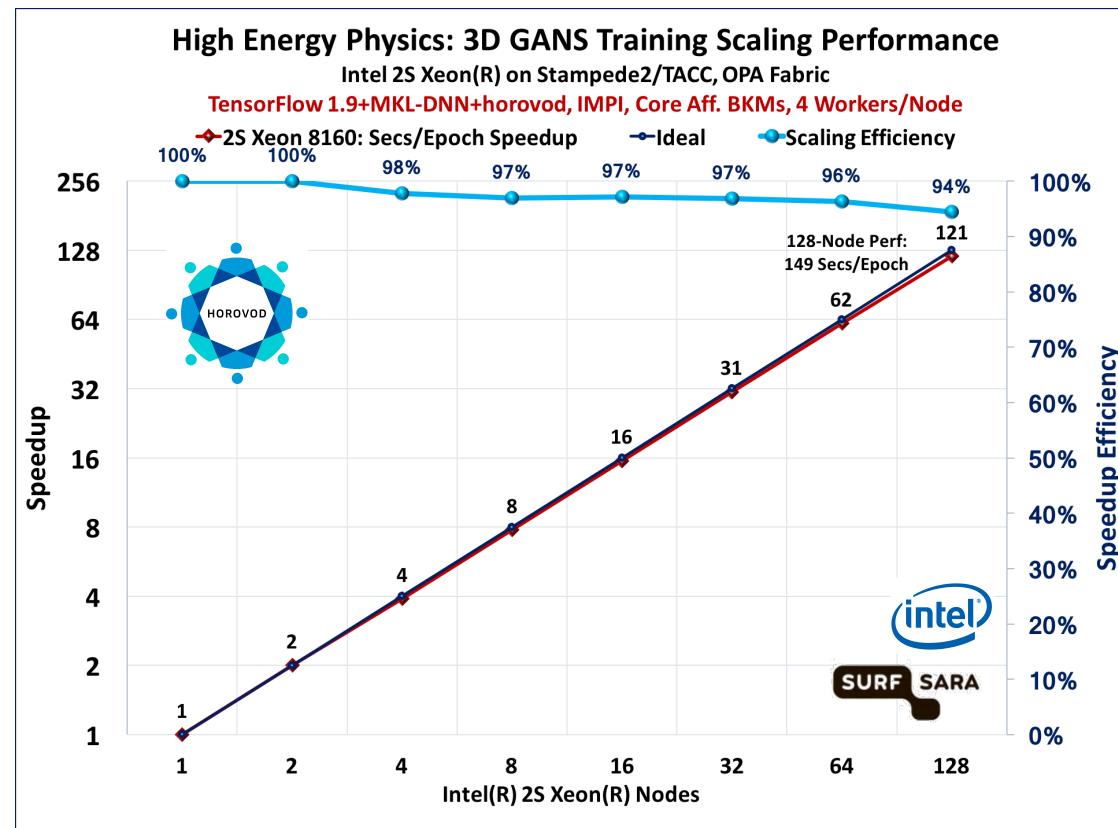
Introduce **data parallel** training based on MPI

Test several libraries

Run on HPC clusters and Cloud
(HNSciCloud providers)

More info in track T6

Time to create an electron shower		
Method	Machine	Time/Shower (msec)
Full Simulation (geant4)	Intel Xeon Platinum 8180	17000
3d GAN (batch size 128)	Intel Xeon Platinum 8180	7
3d GAN (batchsize 128)	GeForce GTX 1080	0.04



Summary & Plans

3D GAN: first step towards customizable simulation tool

Initial results are very promising

Agreement to Monte Carlo within few percent

Meta-optimization and hyper-parameters scans are key

Tested several distributed training approaches on different platforms
(GPUs/CPUs)

Understand / optimize physics performance at scale

More on optimization and distributed training in J. Vlimant talk
on Thursday (Track 6)

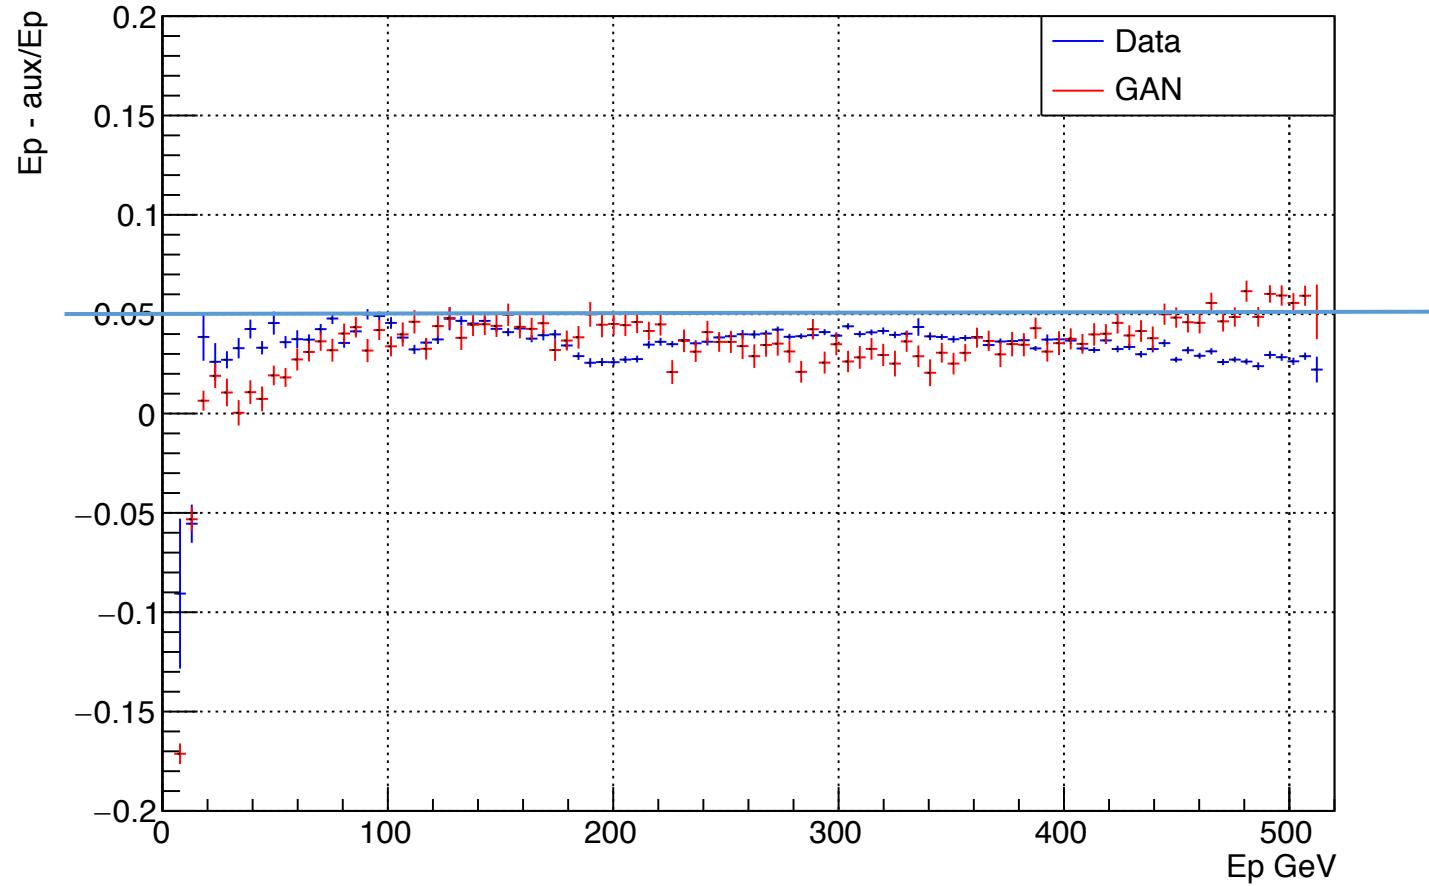


Questions?

Sofia.Vallecorsa@cern.ch

Thanks!

Discriminator regression on input energy



5% error on
auxiliary energy
regression

Conditioning and auxiliary tasks

Loss is linear combination of 3 terms:

Combined cross entropy (real/fake)

Mean absolute percentage error for regression tasks

