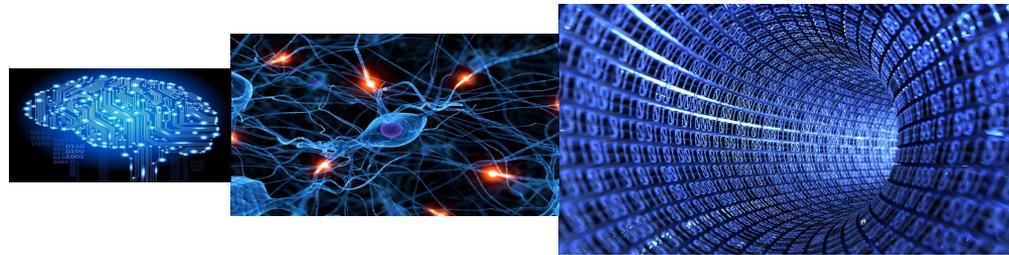


Machine Learning at the LHC: Status, prospects and challenges

Sergei Gleyzer

University of Florida



September 12, 2018

Today's Outline



- **Machine Learning in Practice**
- **Challenges**
- **Applications in HEP**



What is Machine Learning?

- Study of algorithms that improve their performance **P** for a given task **T** with more experience **E**

Sample tasks: identifying faces, Higgs bosons



Preferred approach to:

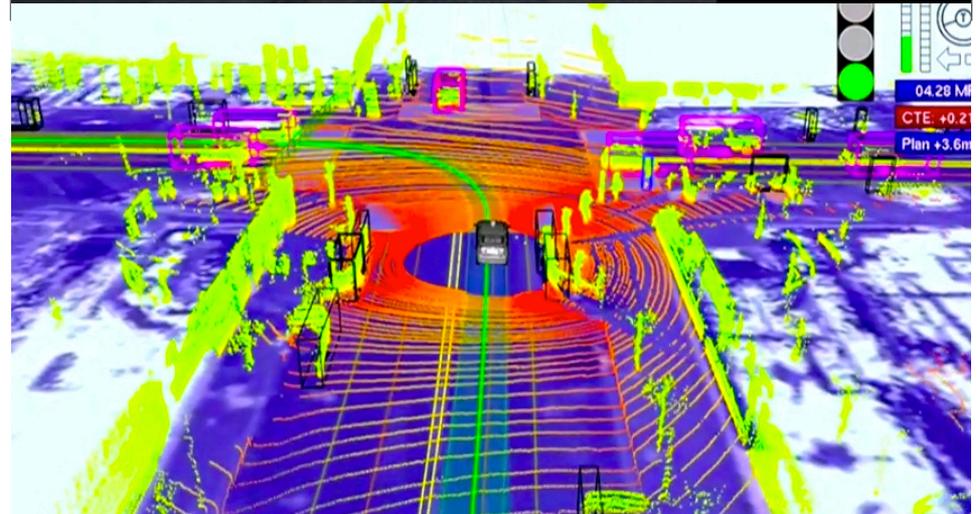
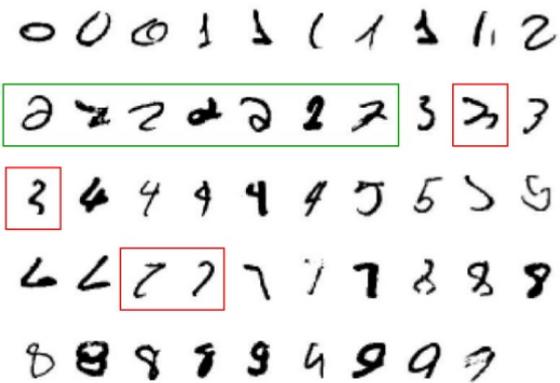
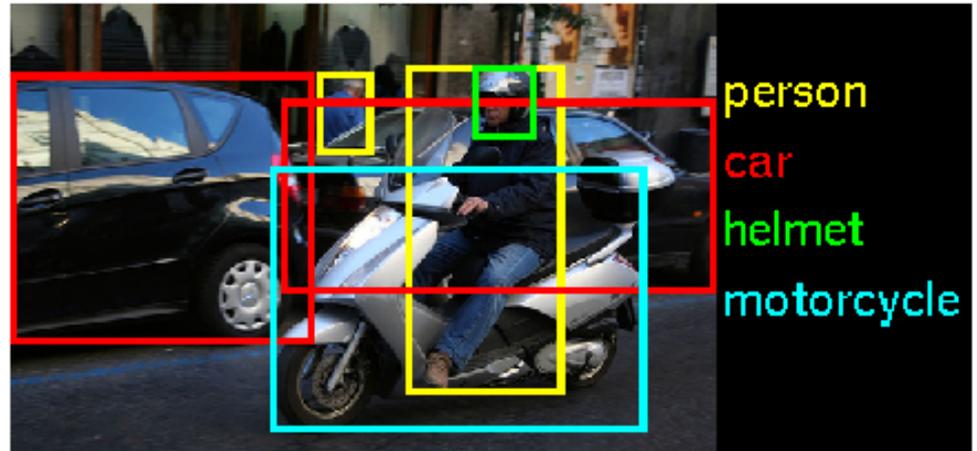
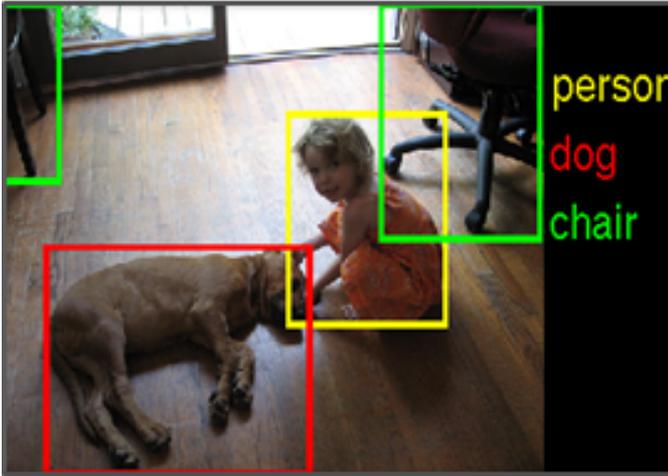
- Speech recognition, natural language processing
- Computer vision, Robot control
- Medical outcomes analysis



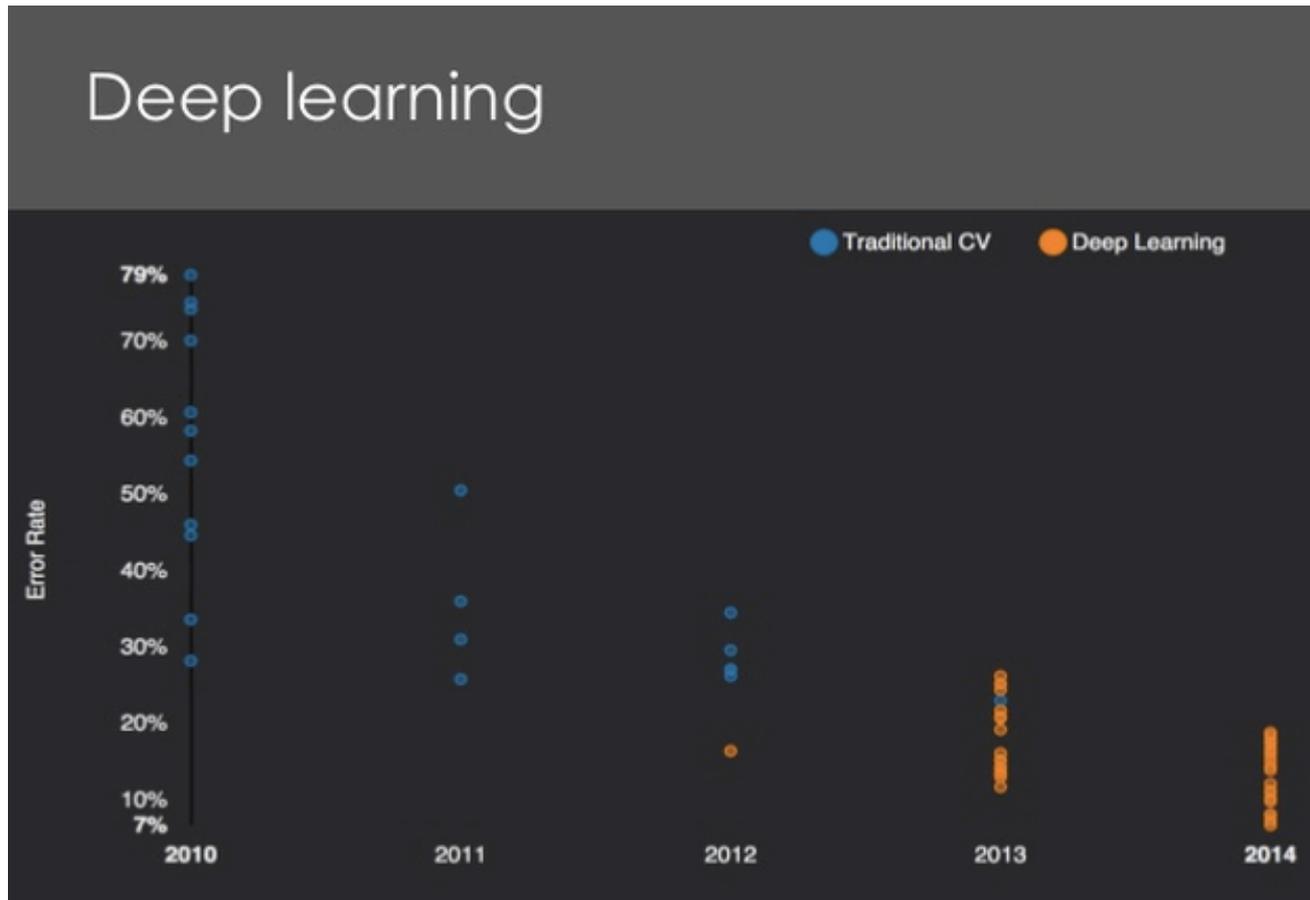
Growing fast:

- Improved algorithms
- Increased data capture
- Software too complex to write by hand

Examples



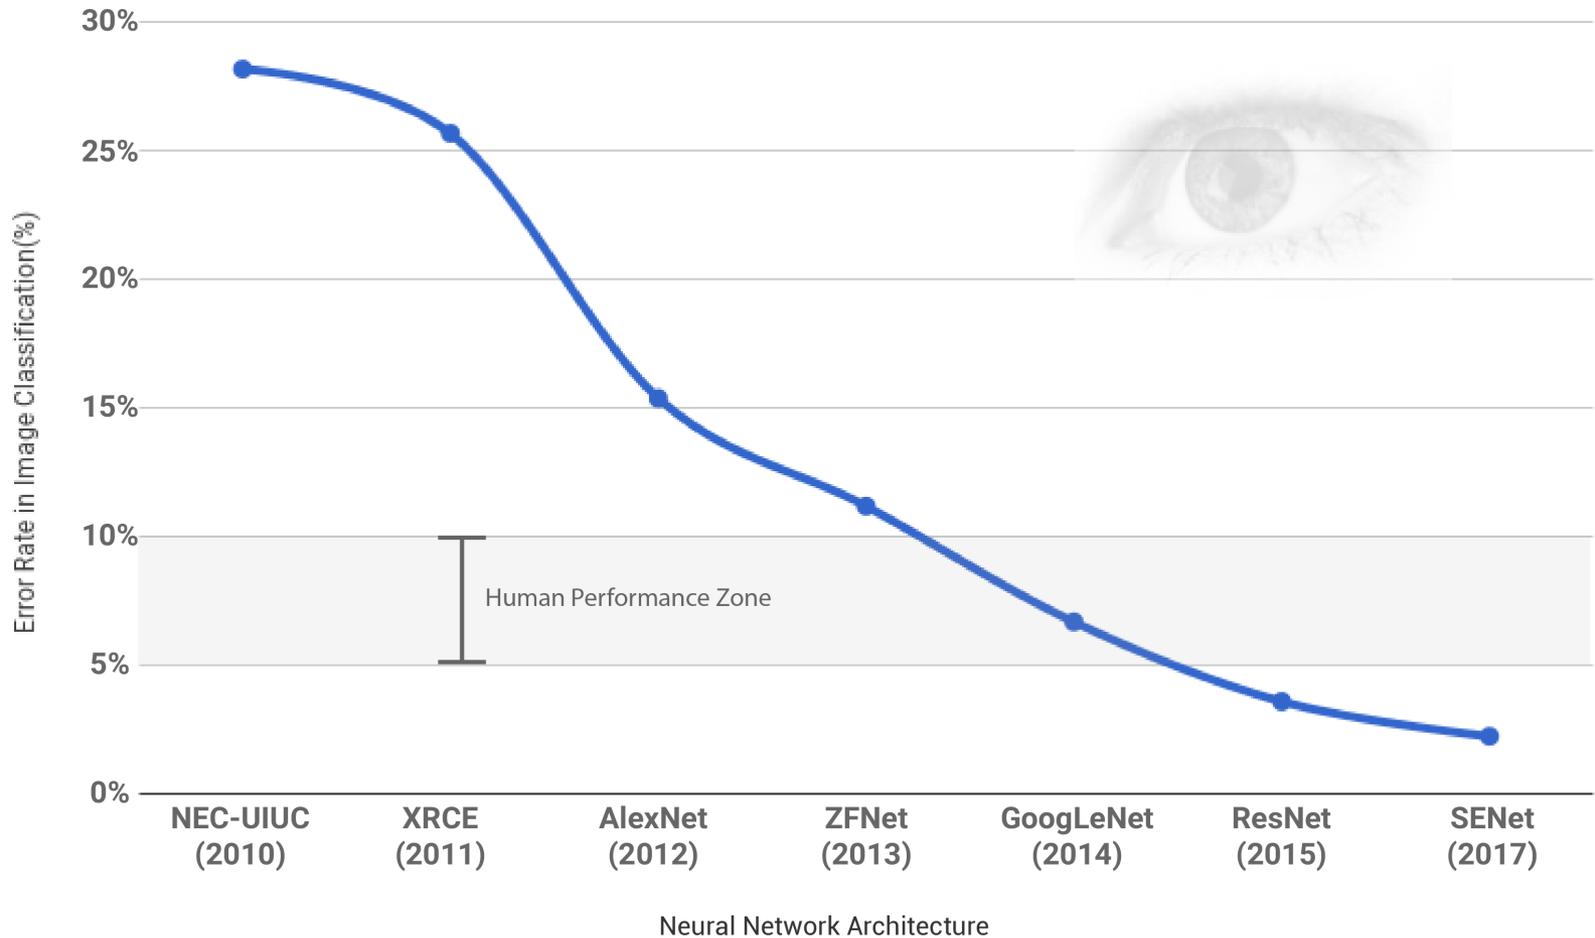
Diving Deeper



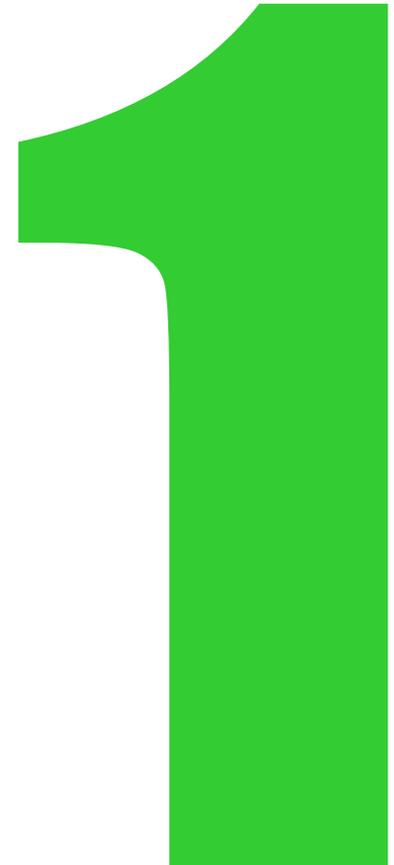
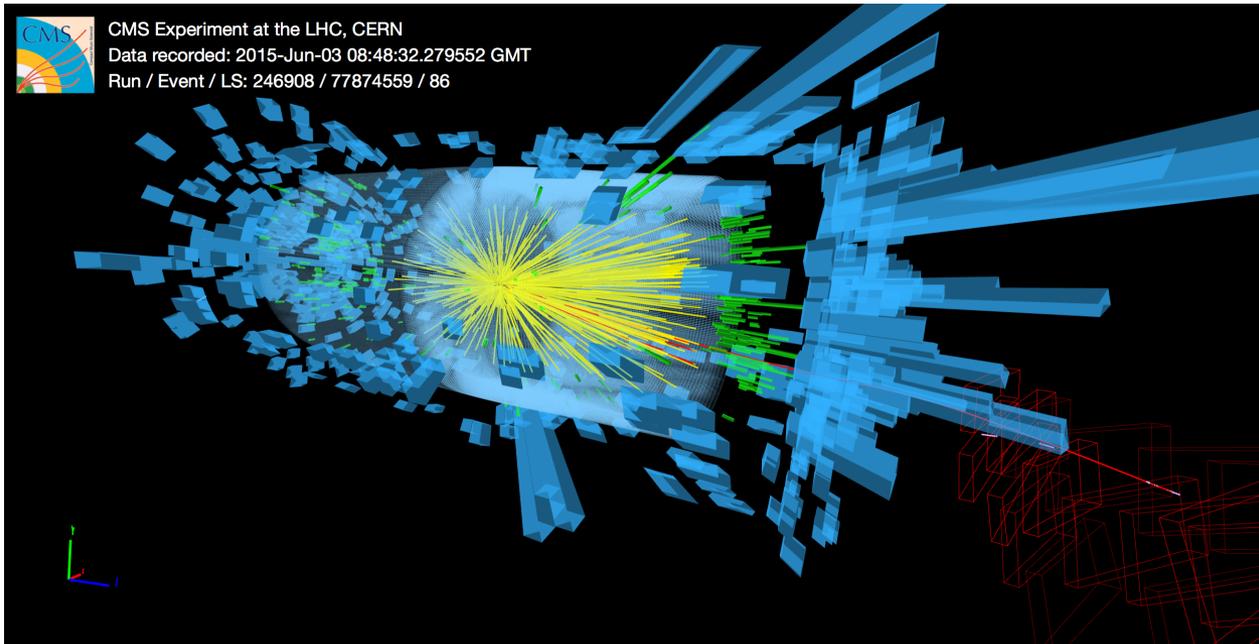
**Huge
Progress**



Diving Deeper



ess



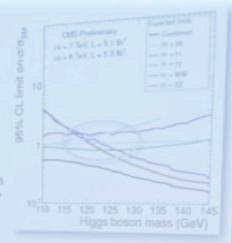
In Particle Physics

Higgs Boson Discovery

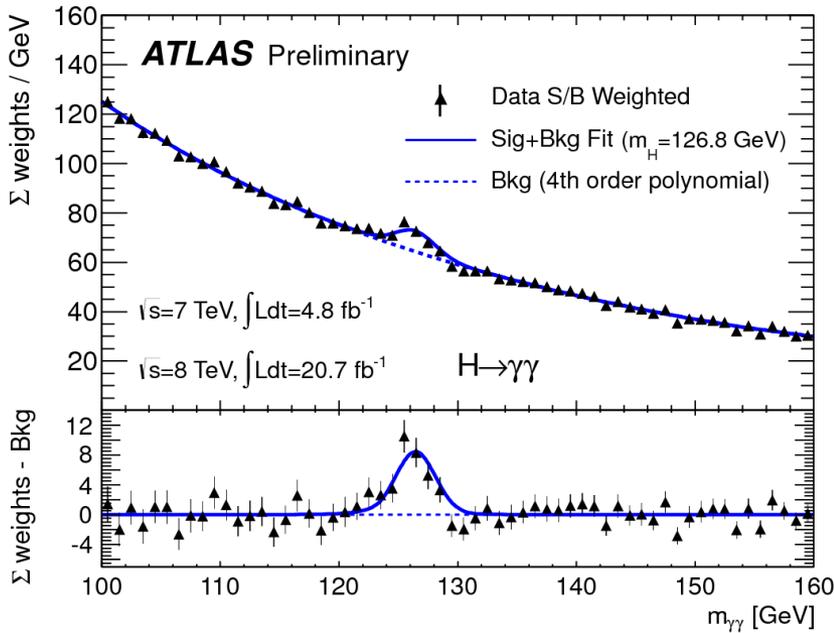


- Not-yet-excluded region: ~ 133 $\gamma\gamma$ GeV
- The five decay modes discussed today have comparable sensitivities for exclusion.
- Most analyses used in this combination have been re-optimized. In order to avoid the possibility of an unintended bias, all selection criteria in the analyses of the 2011 and 2012 data were fixed before looking at the result in the signal region.

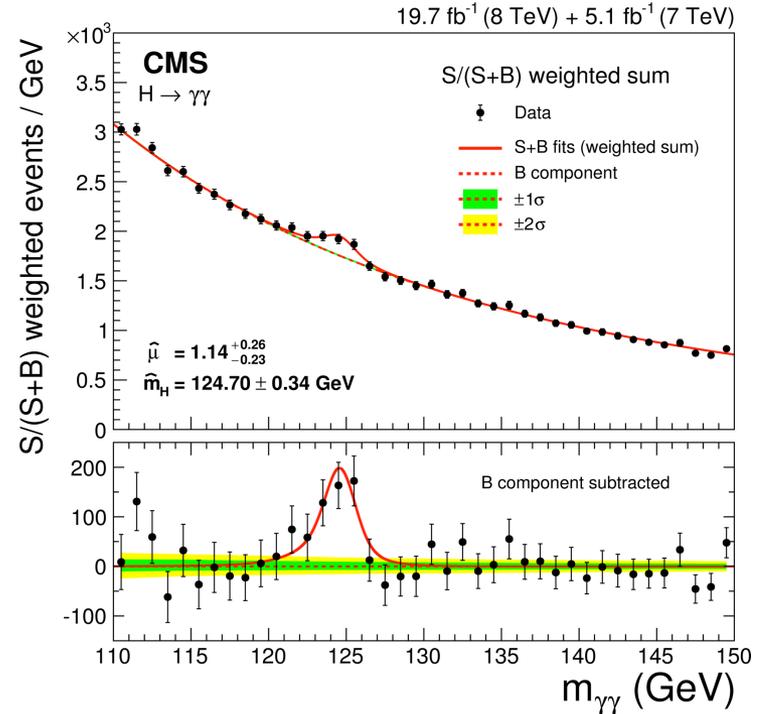
July 4, 2012



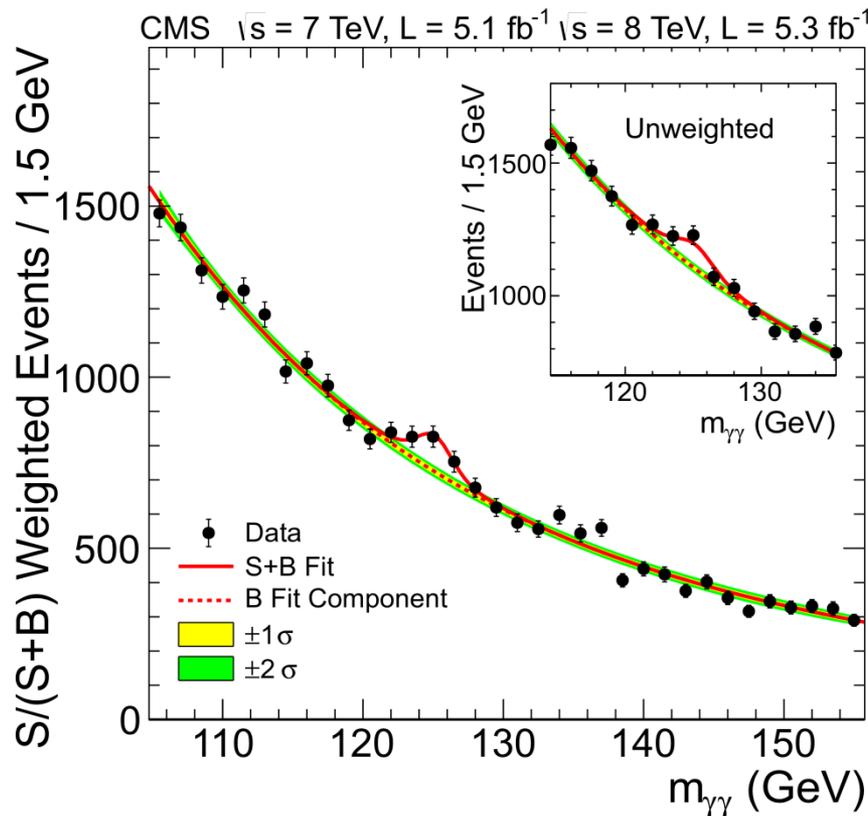
Higgs to di-photons



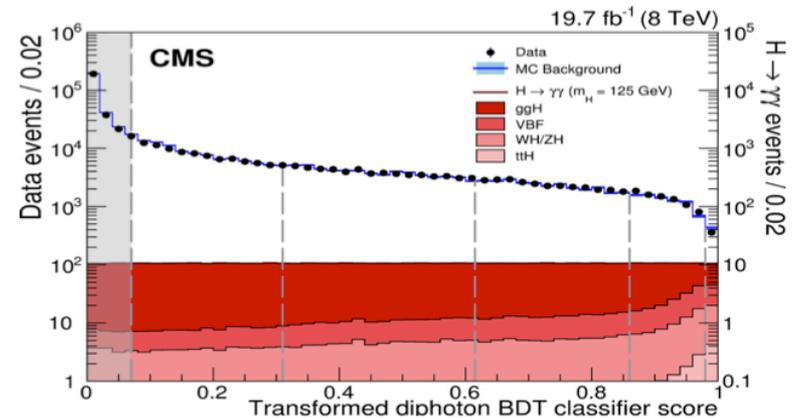
ATLAS



CMS



- Identification of particles
- Identification of interactions
- Energy regression
- Event selection



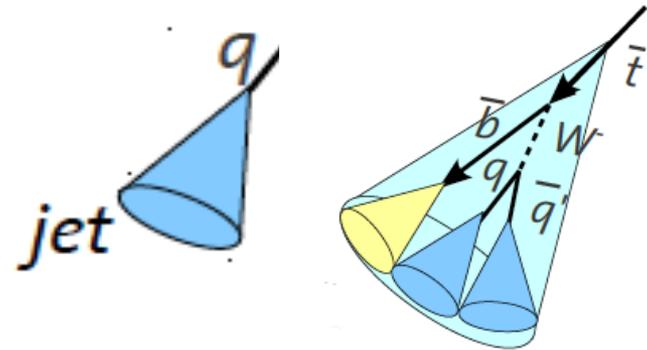
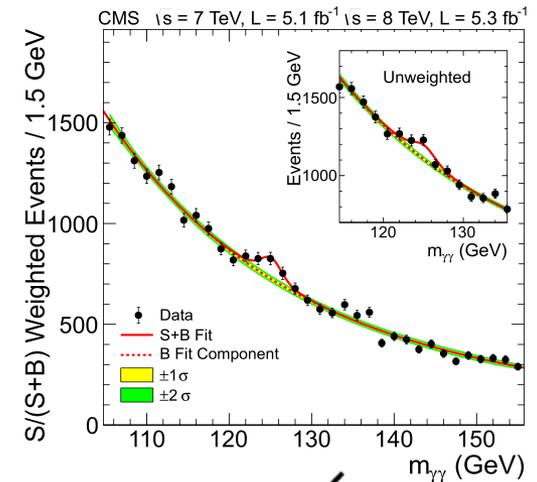
Improvement in analysis from all four areas



Machine learning already at forefront of what we do:

- Physics object **identification**
- Event type **classification**
- Object properties **regression**

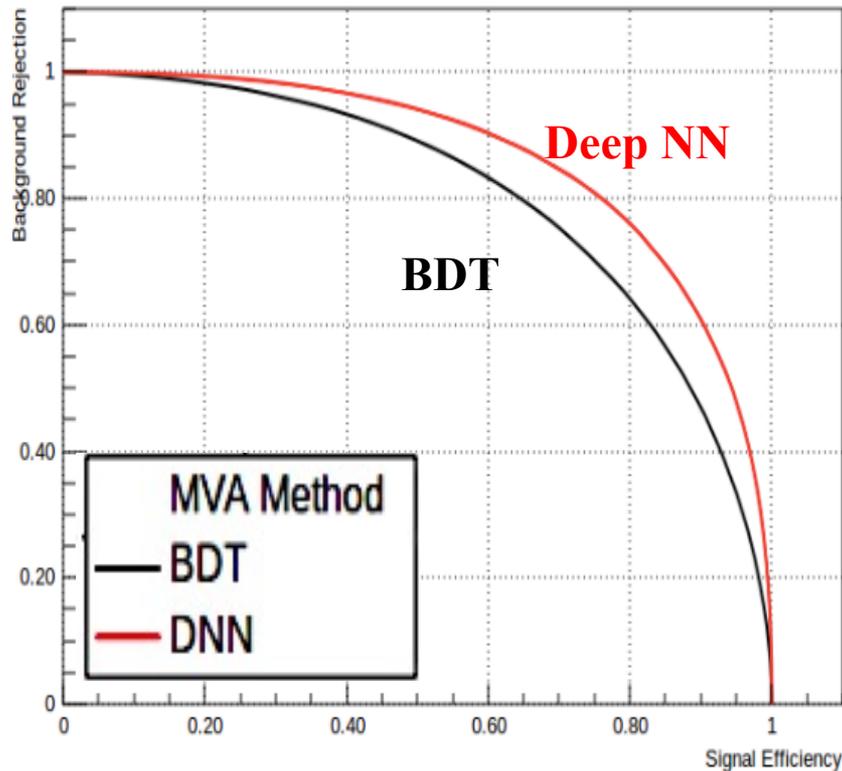
Expanding quickly



Deep Learning



Background Rejection vs. Signal Efficiency



Higher performance compared to previous ML methods

Deep Learning



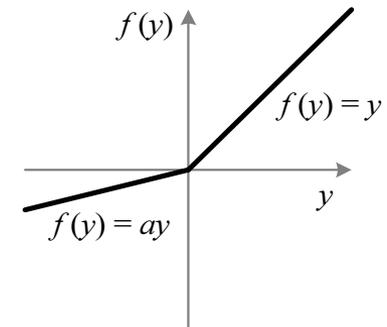
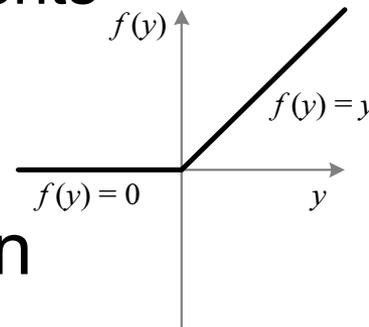
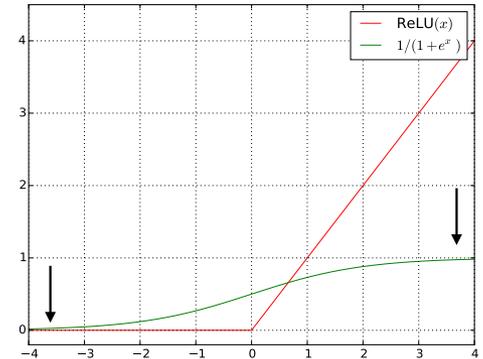
- **Training more complex models**
 - Increased Depth
 - Enlarged Width
 - Specialized Architectures: Convolutional, Recurrent, Graphs
 - Novel activation functions: ReLU
- **Effective strategies avoiding overfitting**
 - Regularization: L1/2, Data Augmentation, DropOut

ReLU



Rectified Linear Unit (ReLU)

- Rectified neuron
- Faster training convergence
 - Better solutions than sigmoids
 - Vanishing gradients
 - Trained by back-propagation

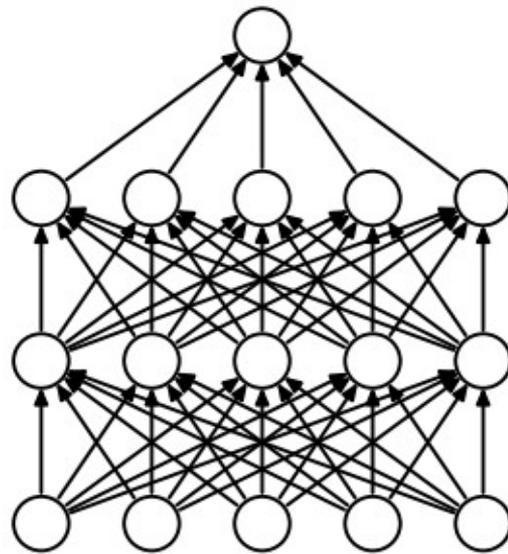


ReLU and Parametric PReLU

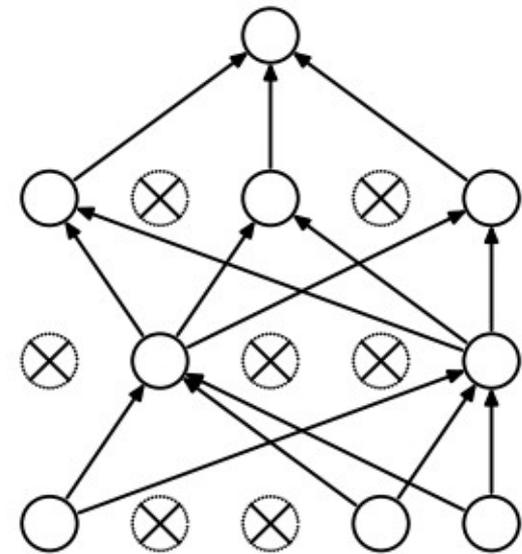
Regularization



- i.e. Drop-Out

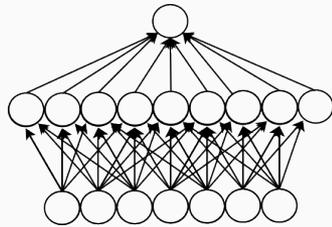


(a) Standard Neural Net

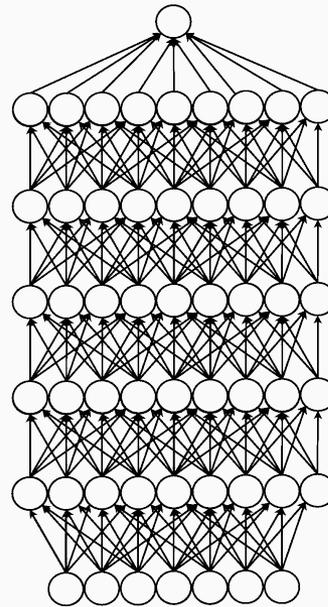


(b) After applying dropout.

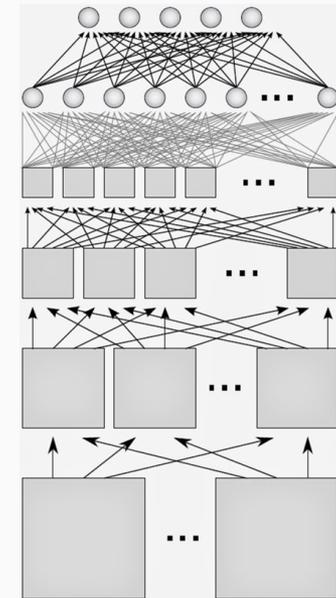
Convolutional



Neural Network (NN)



Deep NN

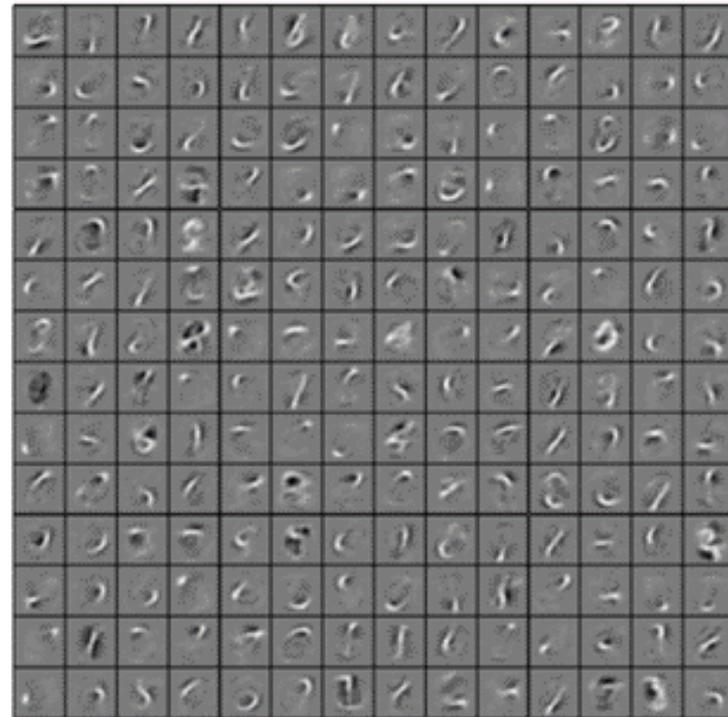


Convolutional NN

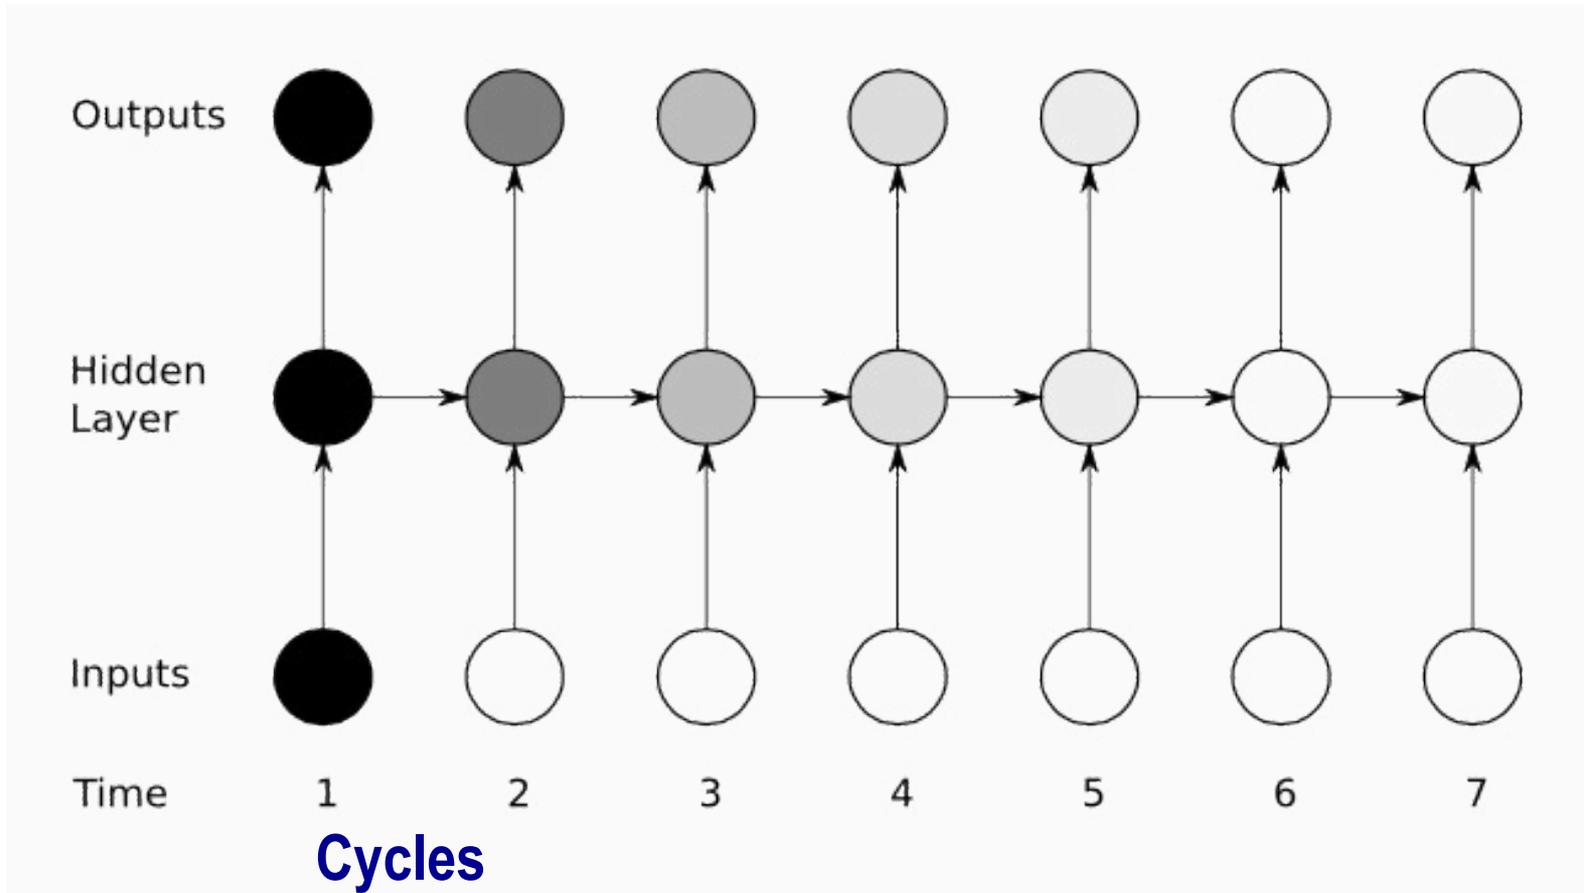
Convolutional Neural Networks:

Unsupervised Feature Learning

5 0 4 1 9 2 1 3 1 4 3 5
 3 6 1 7 2 8 6 9 4 0 9 1
 1 2 4 3 2 7 3 8 6 9 0 5
 6 0 7 6 1 8 7 9 3 9 8 5
 9 3 3 0 7 4 9 8 0 9 4 1
 4 4 6 0 4 5 6 7 0 0 1 7
 1 6 3 0 2 1 1 7 9 0 2 6
 7 8 3 9 0 4 6 7 4 6 8 0
 7 8 3 1 5 7 1 7 1 1 6 3
 0 2 9 3 1 1 0 4 9 2 0 0
 2 0 2 7 1 8 6 4 1 6 3 4
 5 9 1 3 3 8 5 4 7 7 4 2



Recurrent NN

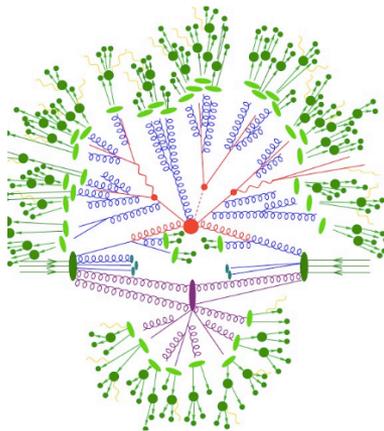


Questions

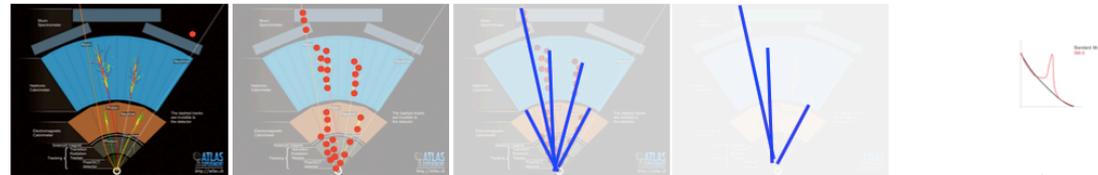


Can we fully exploit the detectors?

- Raw data, low-level variables



Raw	Sparsified	Reco	Select	Ana
1e7	1e3	100	50	1



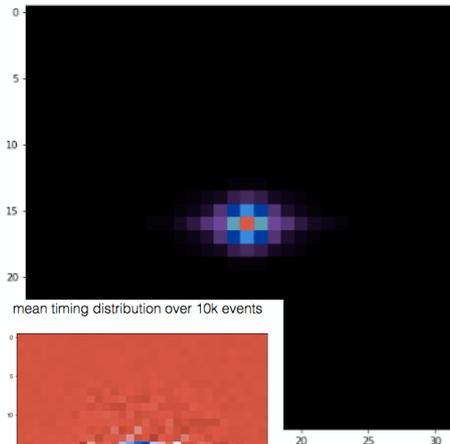
Example



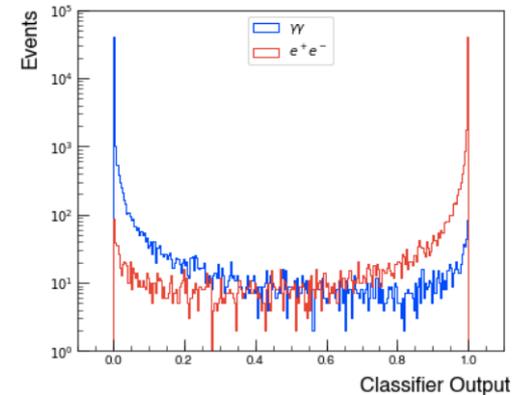
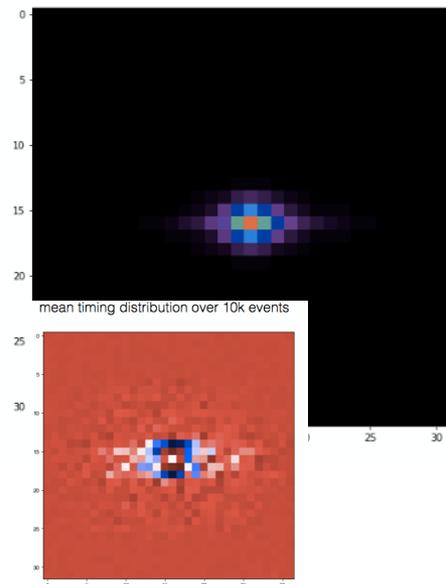
“End-to-end learning” Andrews et al., 2018

- By-passing traditional reconstruction

Photon-Induced EM Shower
mean energy distribution over 10k events

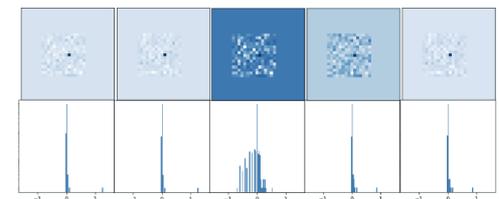


Electron-Induced EM Shower
mean energy distribution over 10k events



ResNet-23

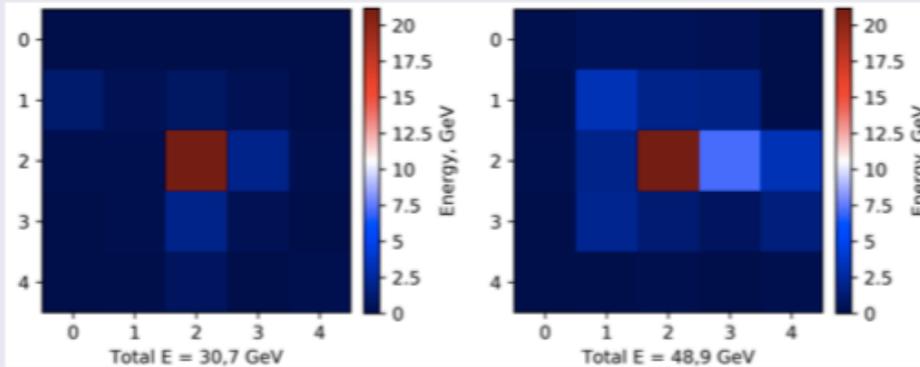
Test Set ROC AUC 0.997



Particle ID



Clusters for photon and π^0 photon

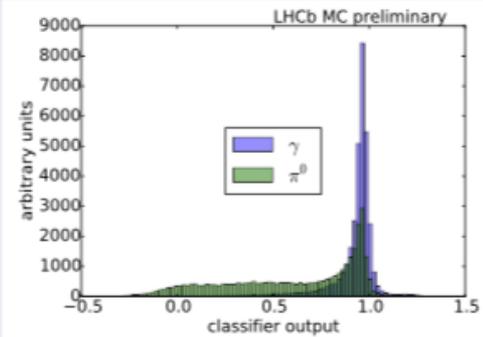


Responses from single photon(left) and merged π^0 (right)

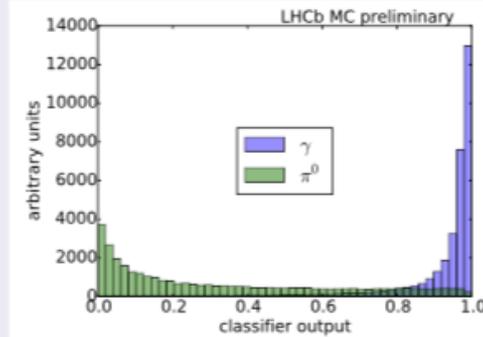
LHCb, CHEP2018

New approach

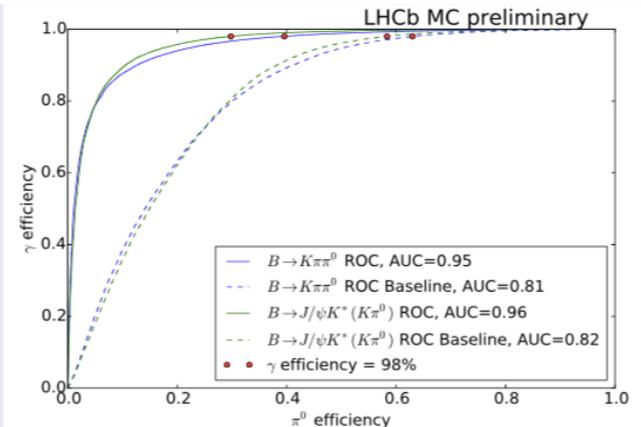
- Consider 5*5 cluster
- Use energy in each cell as a feature
- Use several models and look for the best one



Baseline output



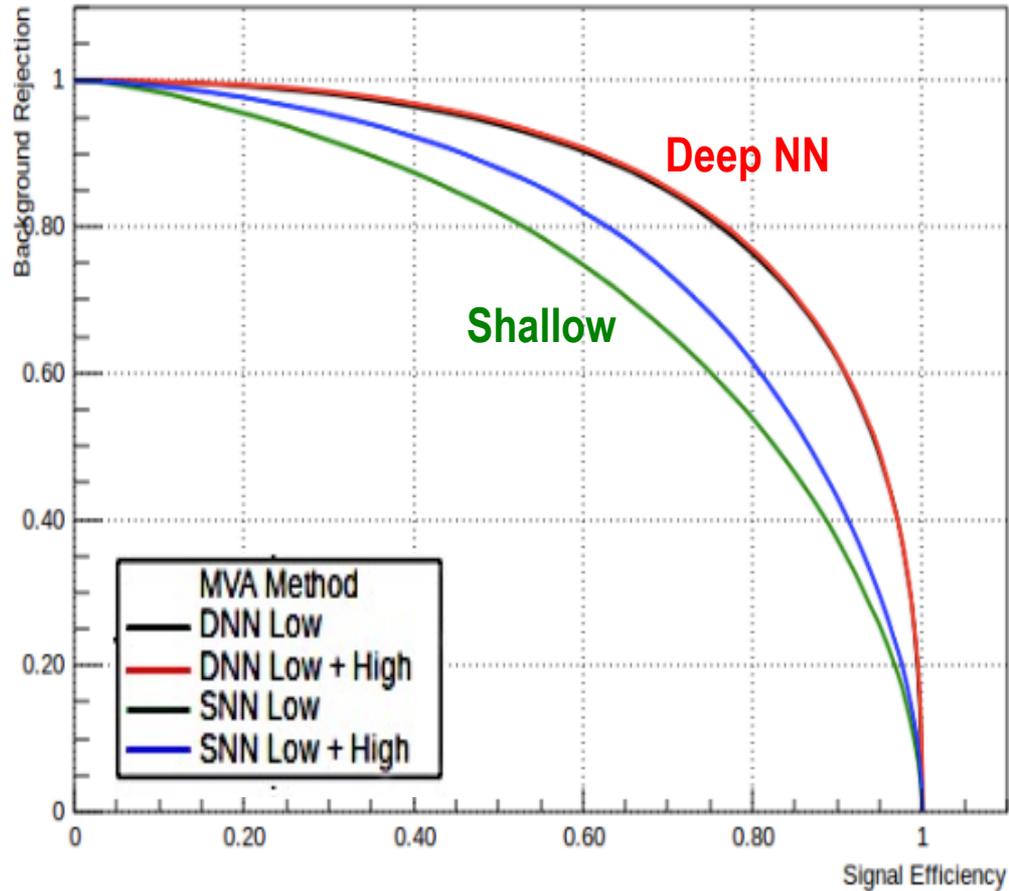
XGBoost approach output

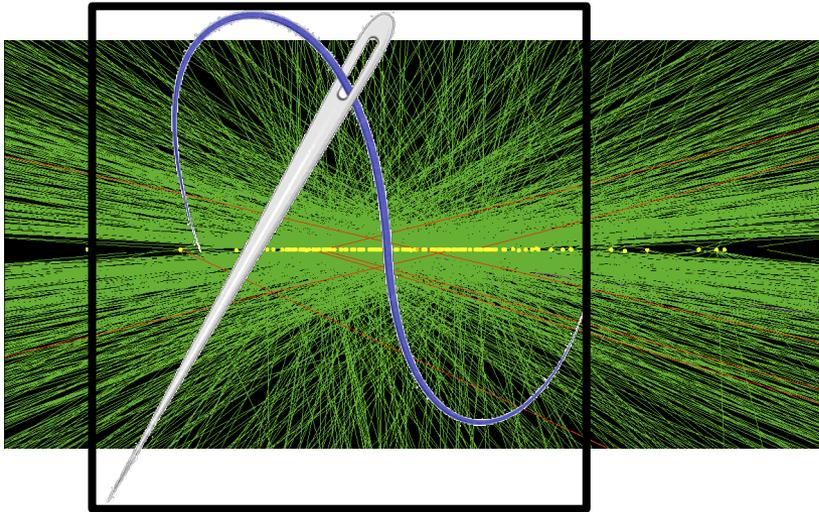


Feature Extraction



Background Rejection vs. Signal Efficiency



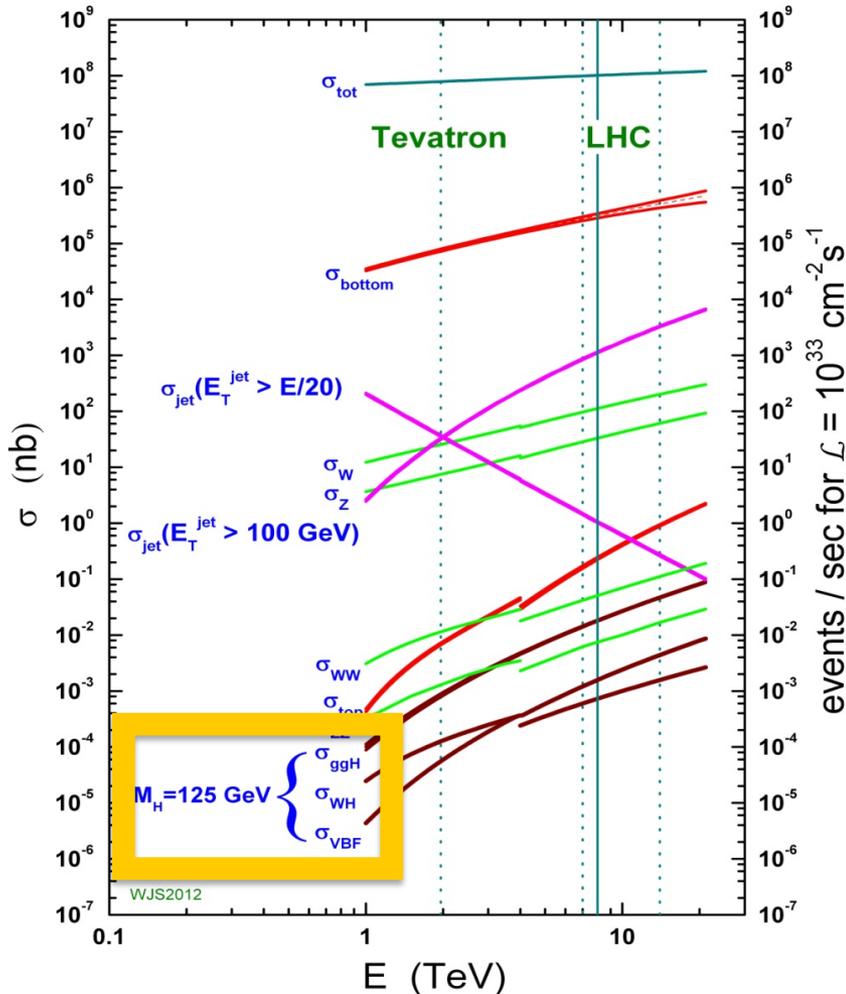


Present and Future Challenges

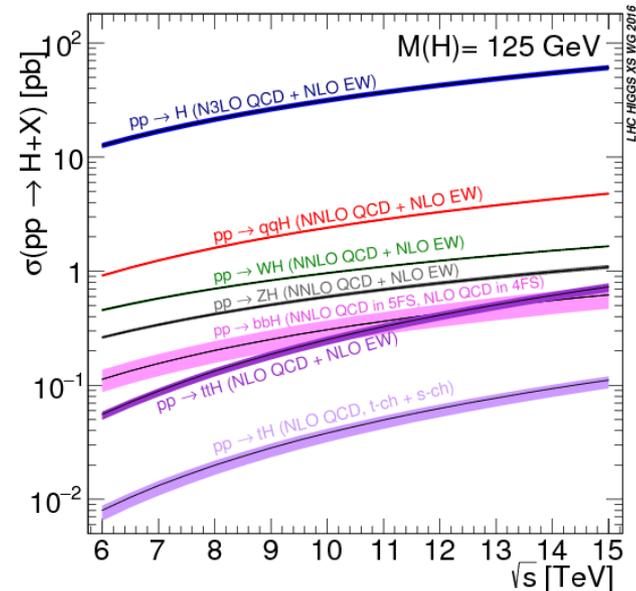




proton - (anti)proton cross sections



Orders of magnitude between signals and backgrounds



*Observations*

New Higgs Boson Observations Reveal Clues on the Nature of Mass

For the first time the scientists have observed the famous Higgs boson, responsible for imparting mass, interacting with the heaviest particle in the universe

By Don Lincoln on June 6, 2018

Purch

PHYSICS

Physicists Observe the Higgs Boson's Elusive Decay

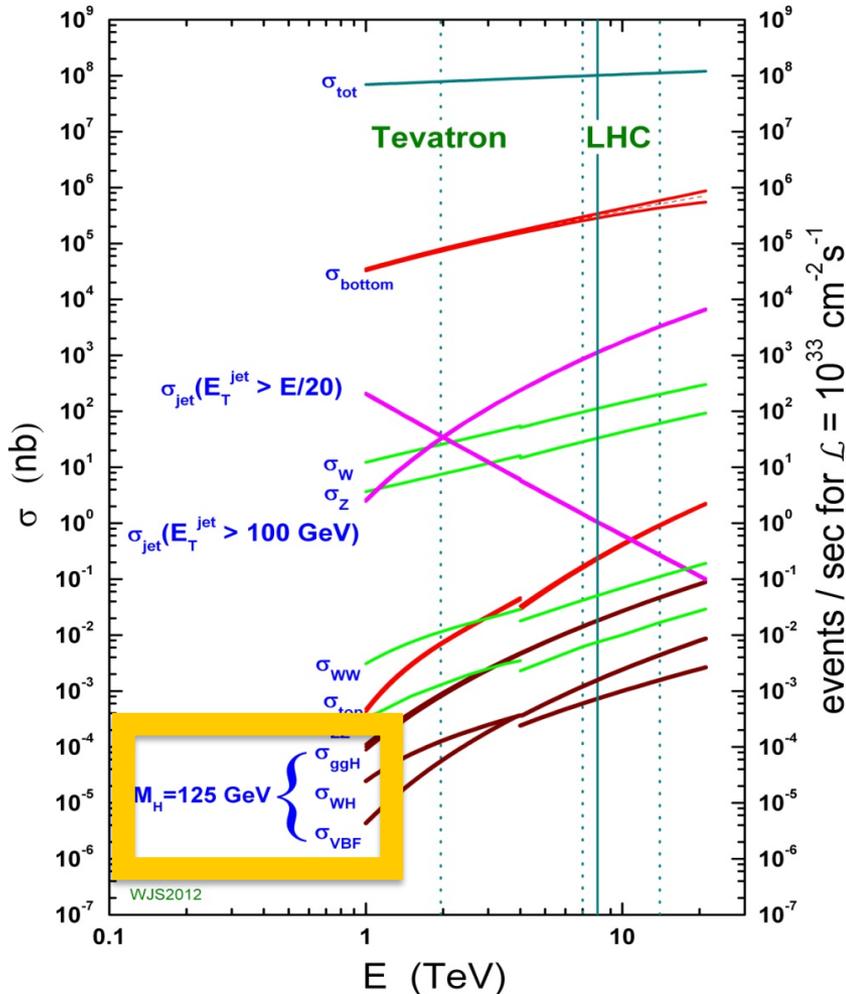
Directly detecting this long-predicted phenomenon further validates the Standard Model of particle physics

By Chelsea Gohd, SPACE.com on August 29, 2018

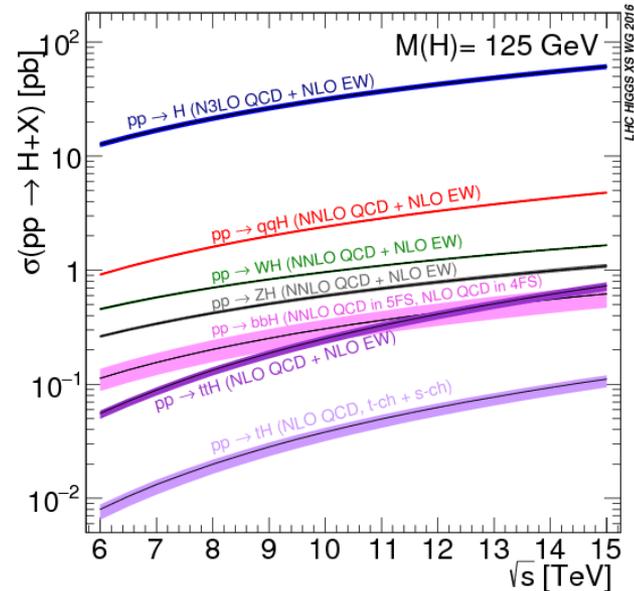
E (TeV)



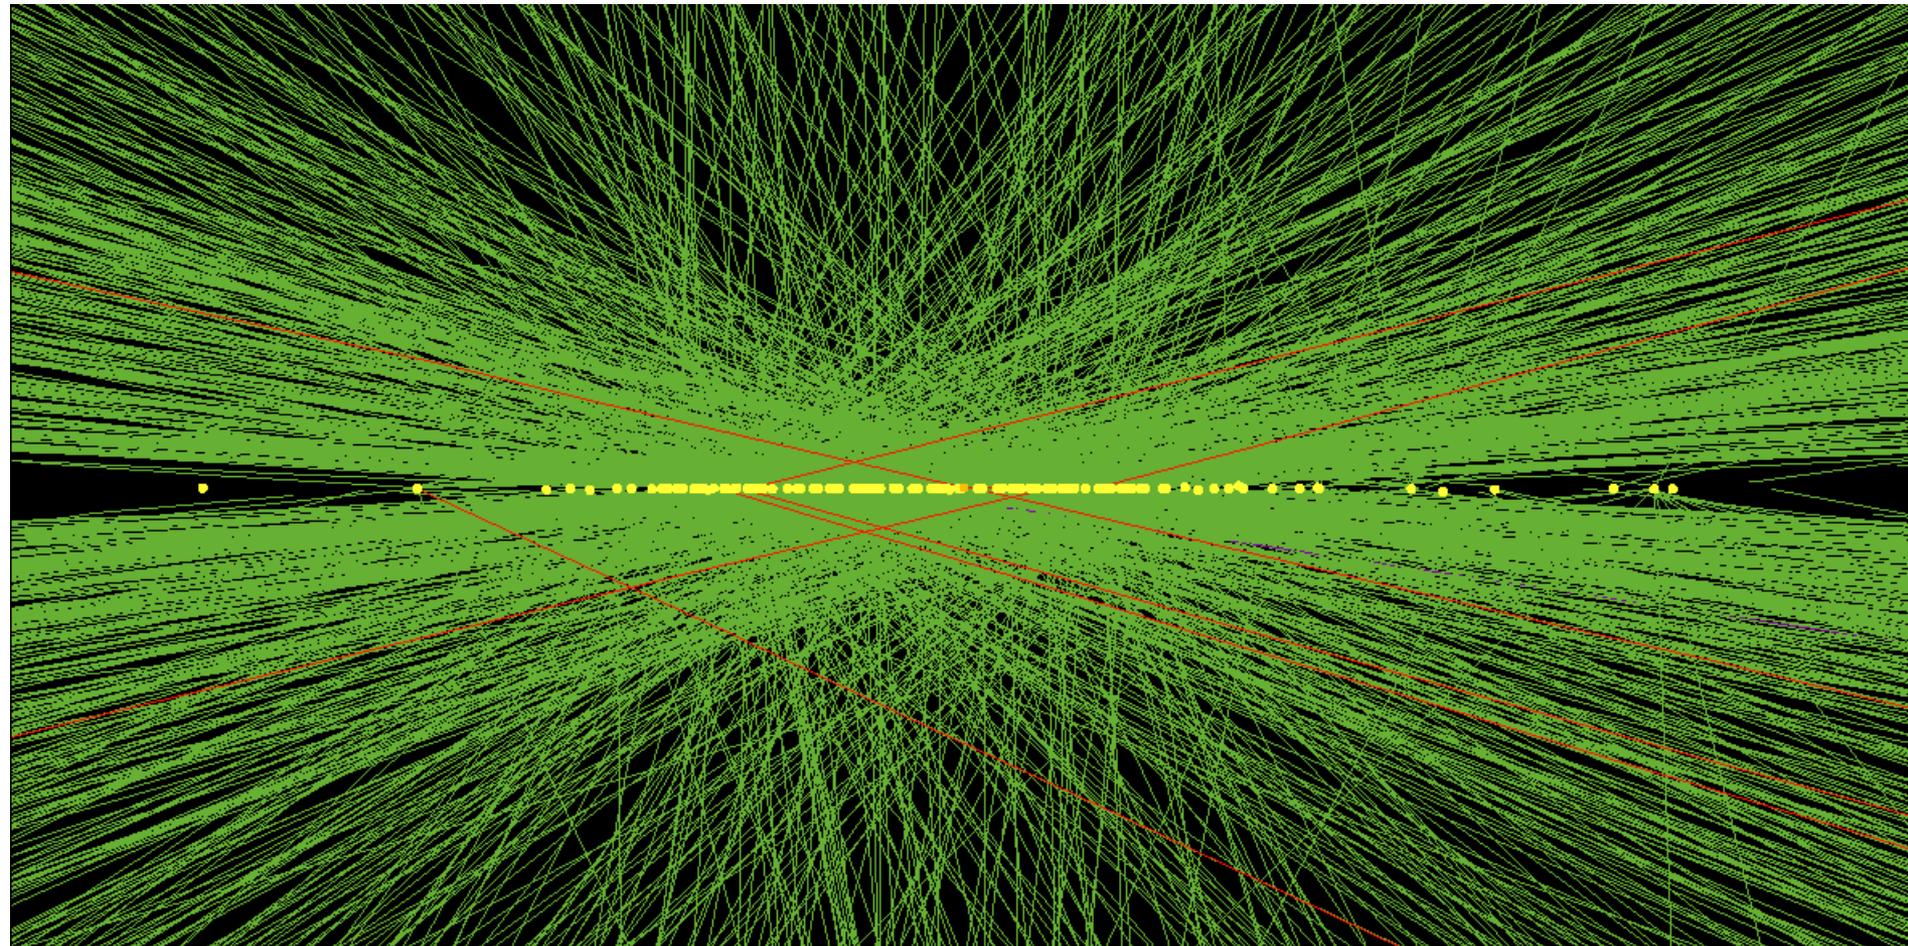
proton - (anti)proton cross sections



Orders of magnitude between signals and backgrounds



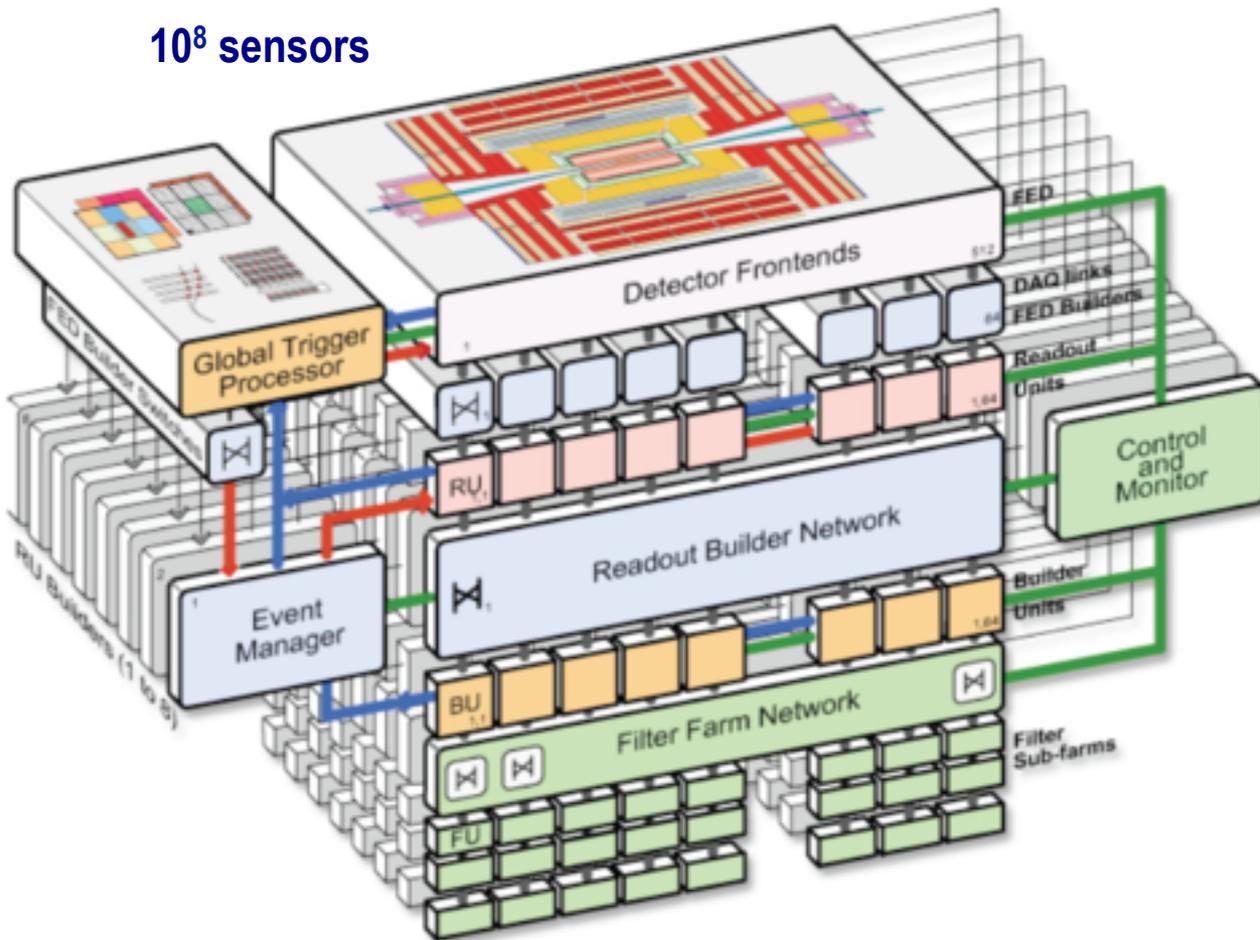
Event Complexity



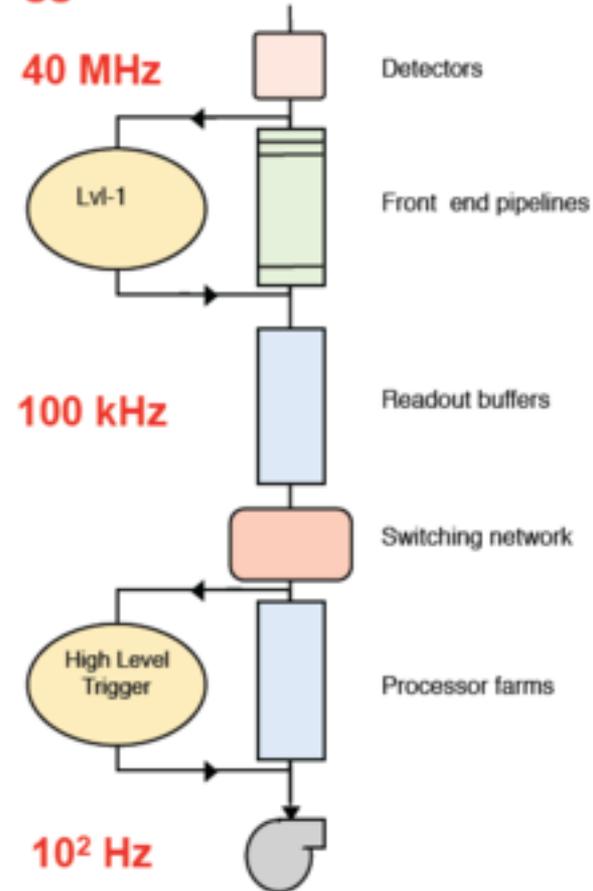
Trigger

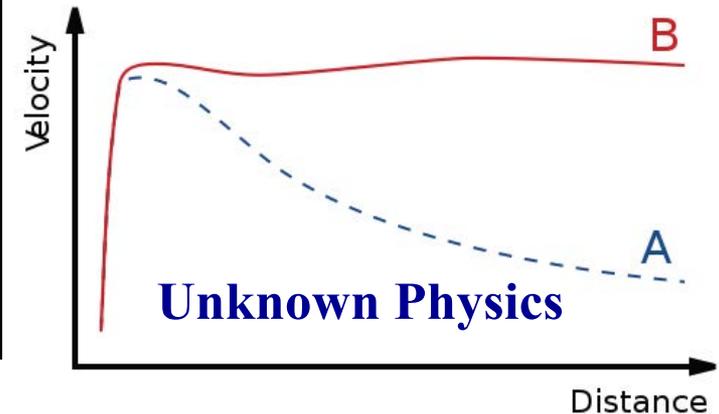
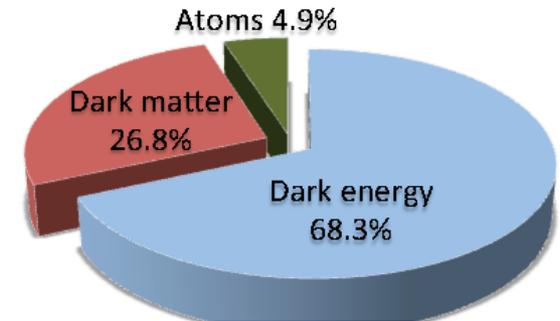
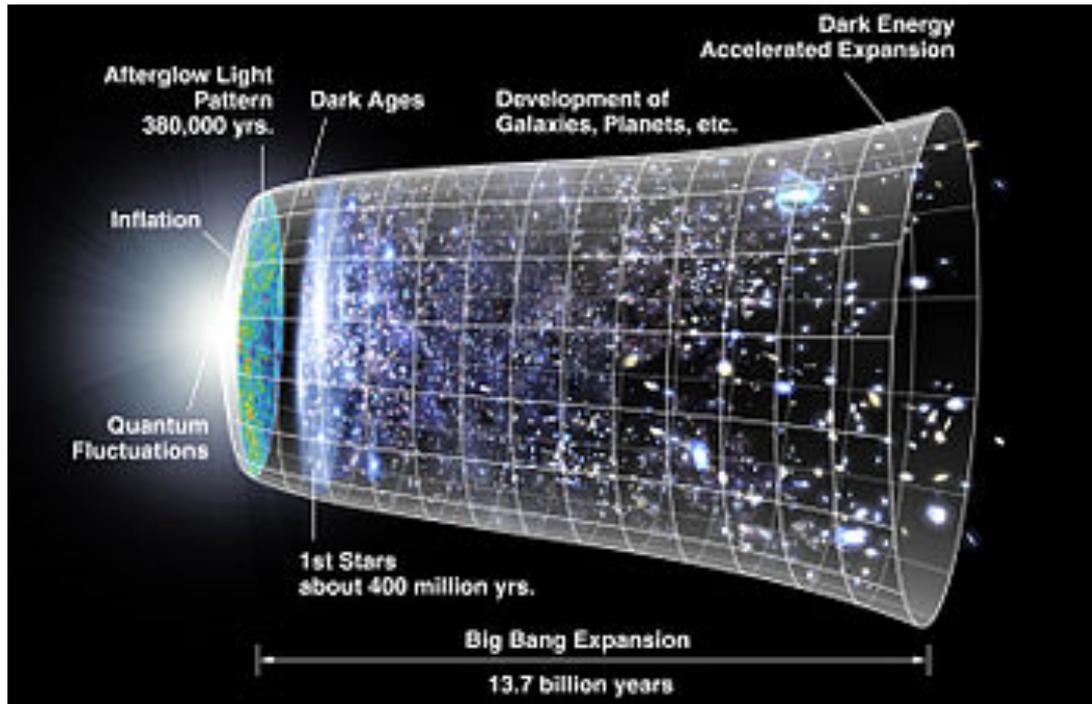


10^8 sensors



Trigger Rate

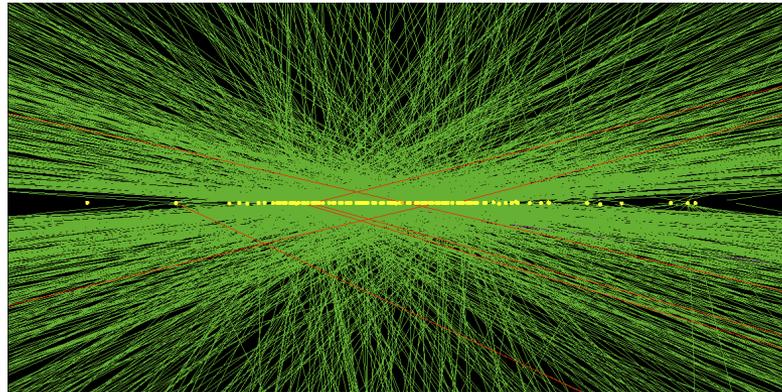




Data size:

- LHC **15,000,000 Tb** 2010 – 2035
- Resources not up as fast as data volume

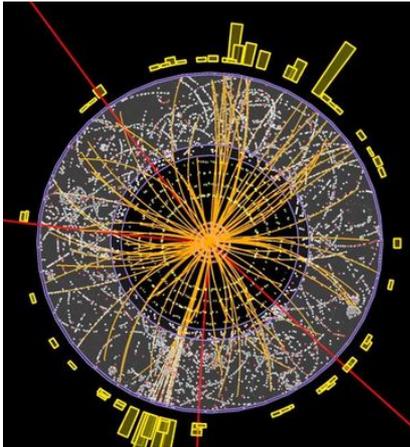




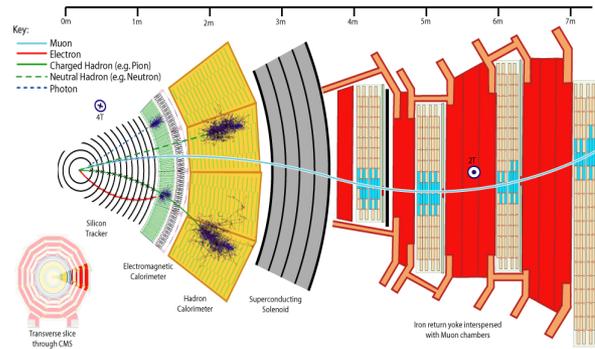
HEP Applications



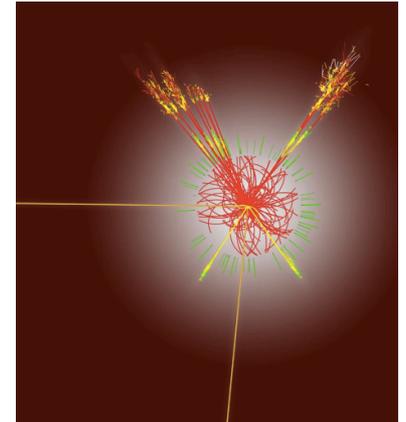
Interesting areas



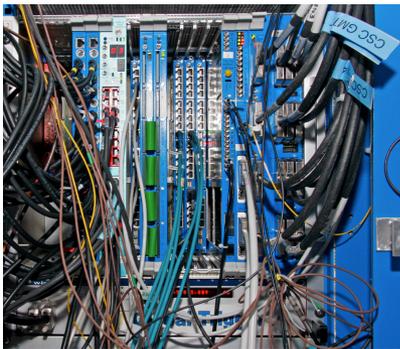
Tracking



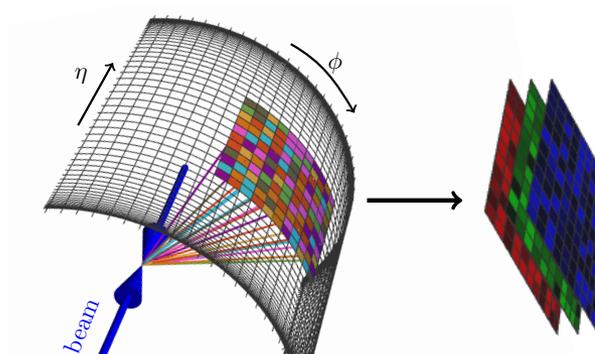
Fast Event Simulation



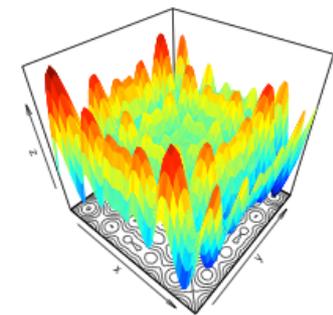
Object Identification



Trigger

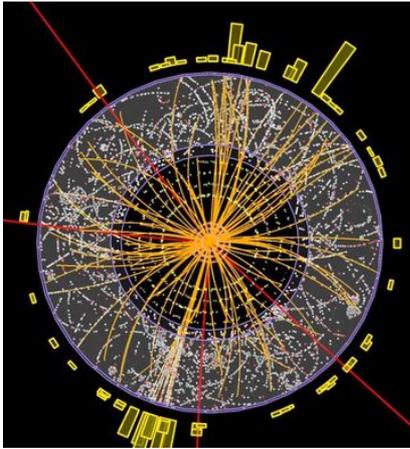


Imaging Techniques

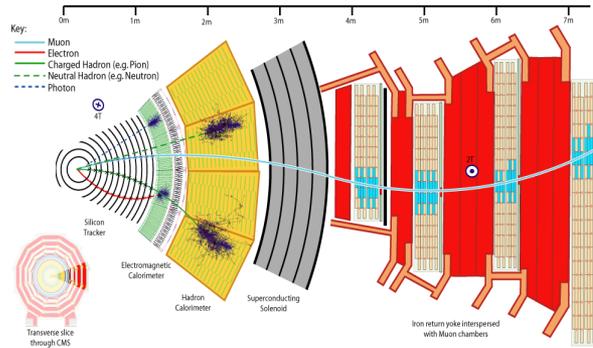


Simulation

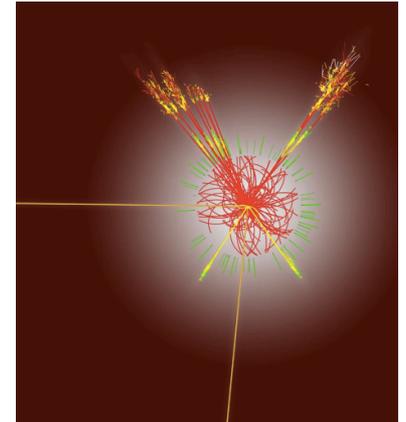
Interesting areas



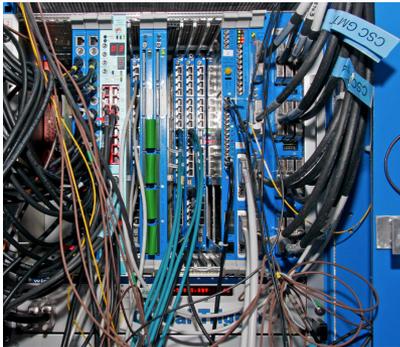
**Deep Kalman,
LSTMs, GNN**



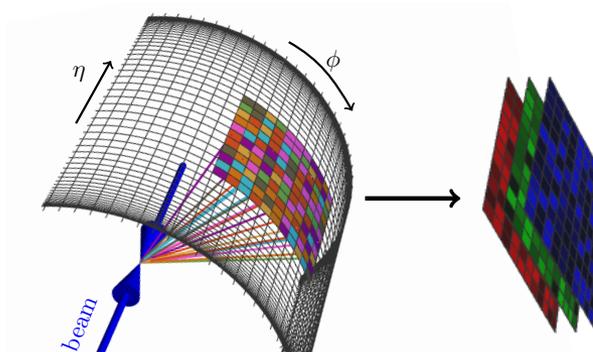
**Generative Models,
Adversarial Networks**



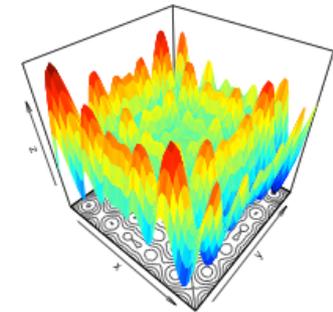
**FCN, Recurrent,
LSTMs**



Deep ML +FPGA



Convolutional NN



Multiobjective Regression



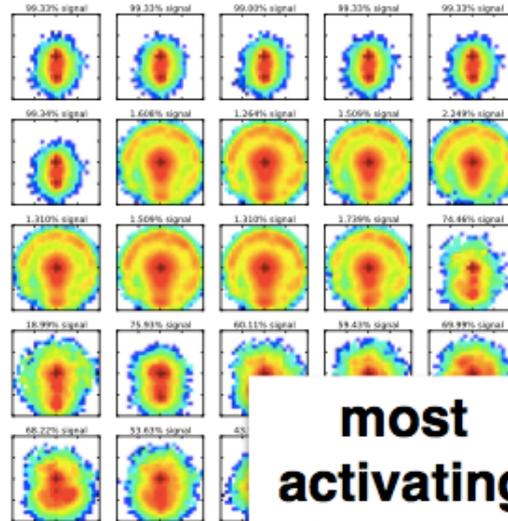
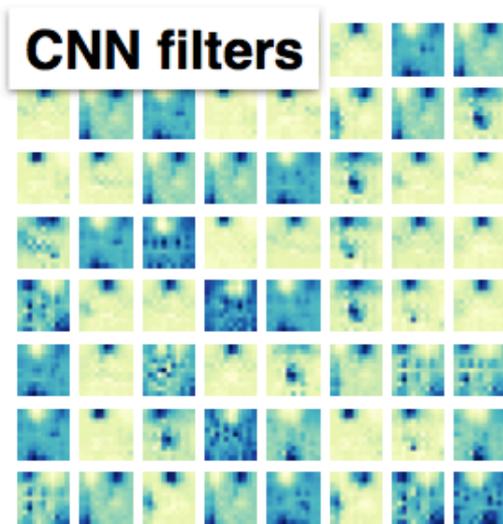
If a problem can be expressed as a known problem

- Apply **existing algorithms**
 - **Example:** convolutional neural networks from computer vision

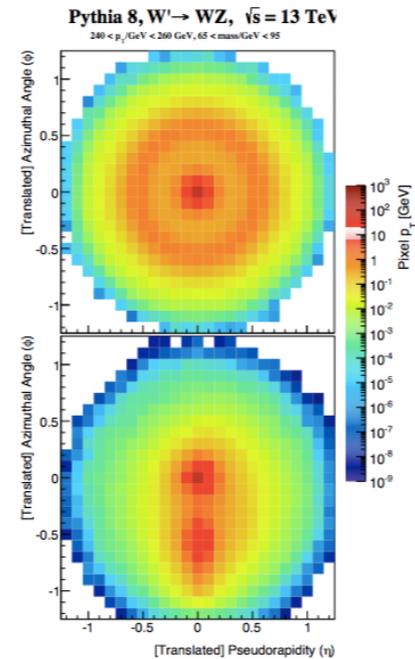
If a problem has not been solved

- Push the knowledge boundary forward

Jet images with convolutional nets



**most
activating
images**

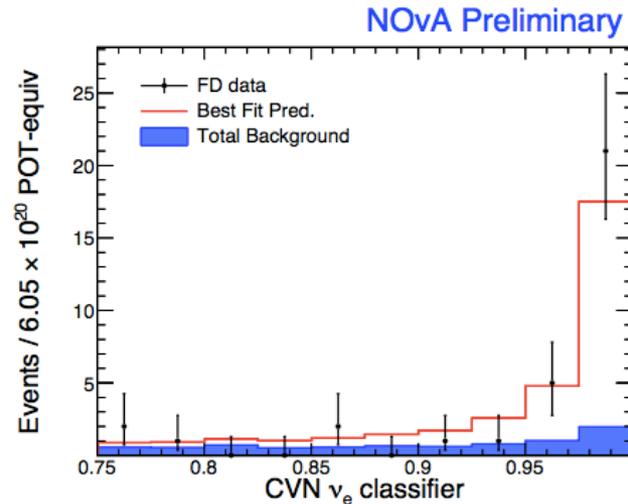
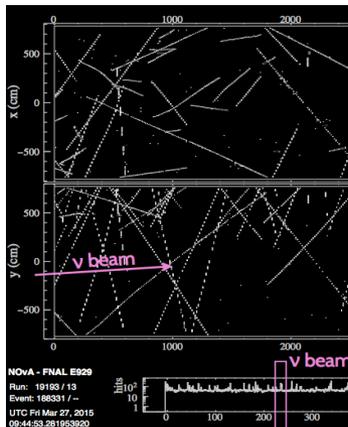


L. de Oliveira et al., 2015

Examples



Neutrinos with convolutional nets

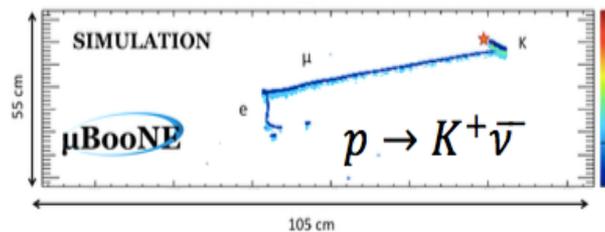
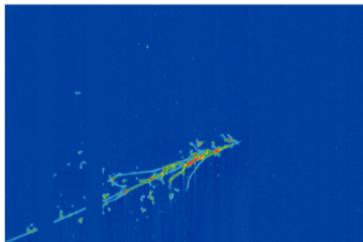


76% Purity
73% Efficiency

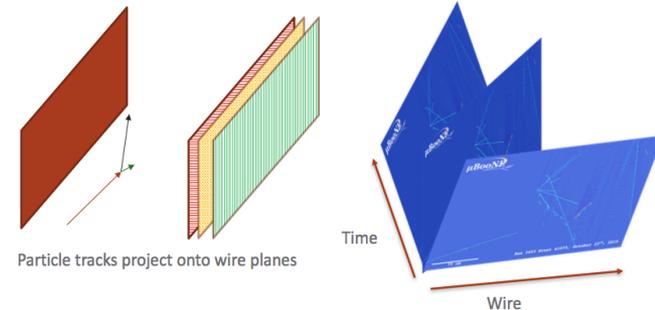


An equivalent increased exposure of 30%

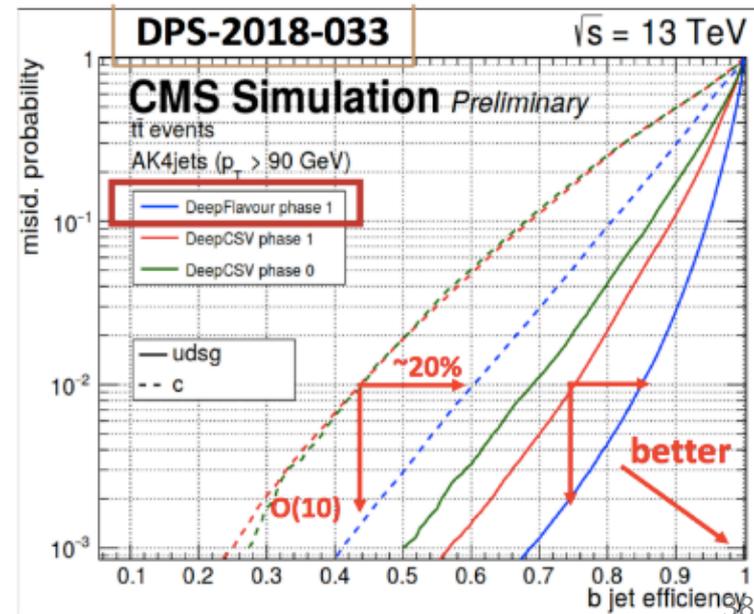
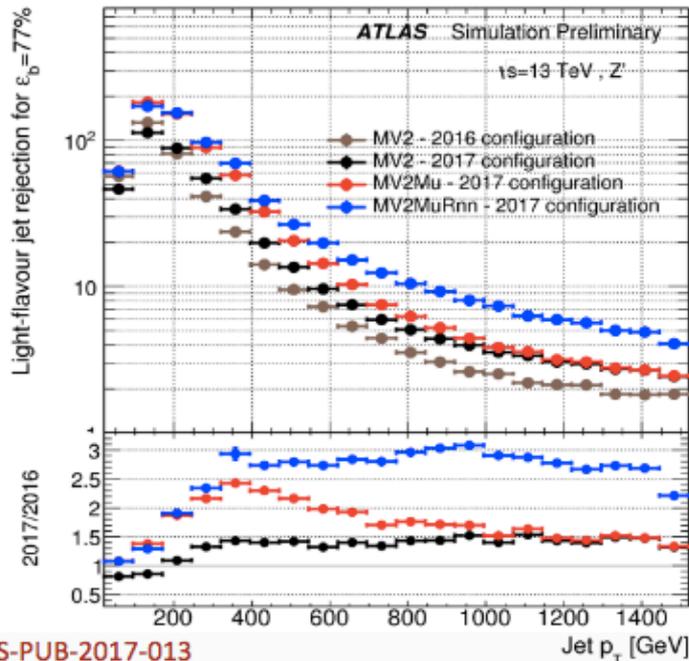
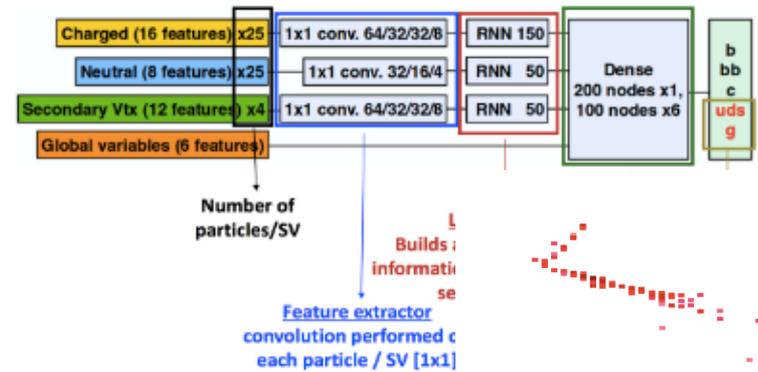
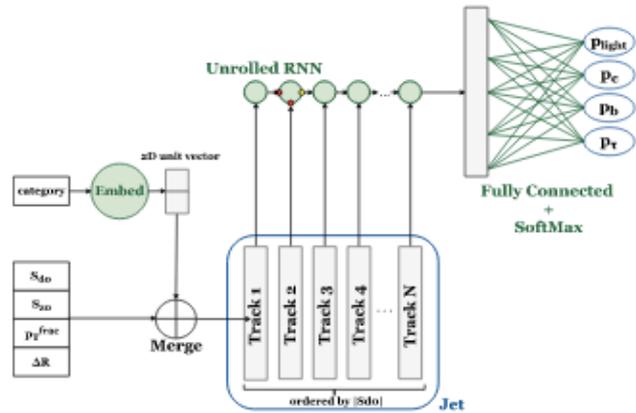
Aurisiano et al. 2016



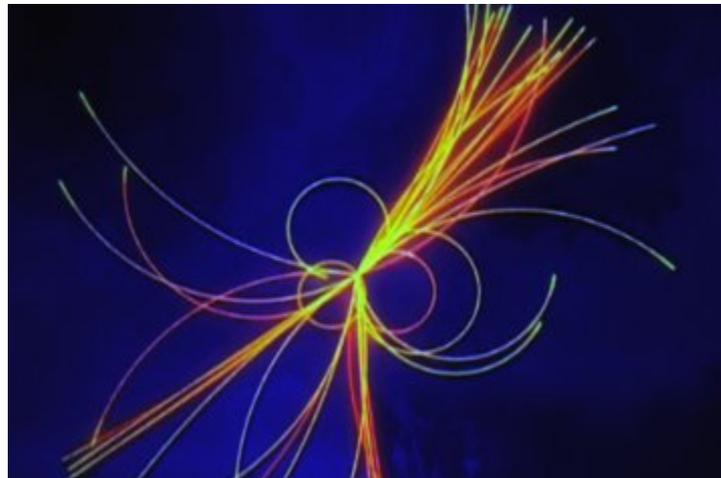
μBooNE



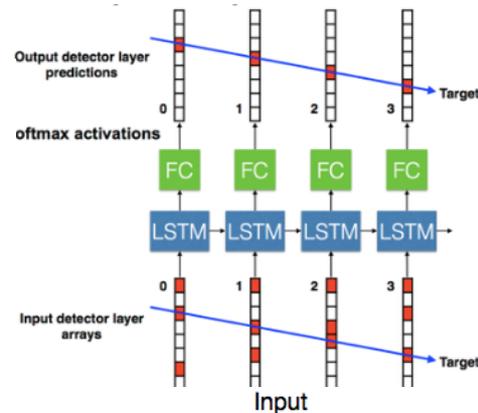
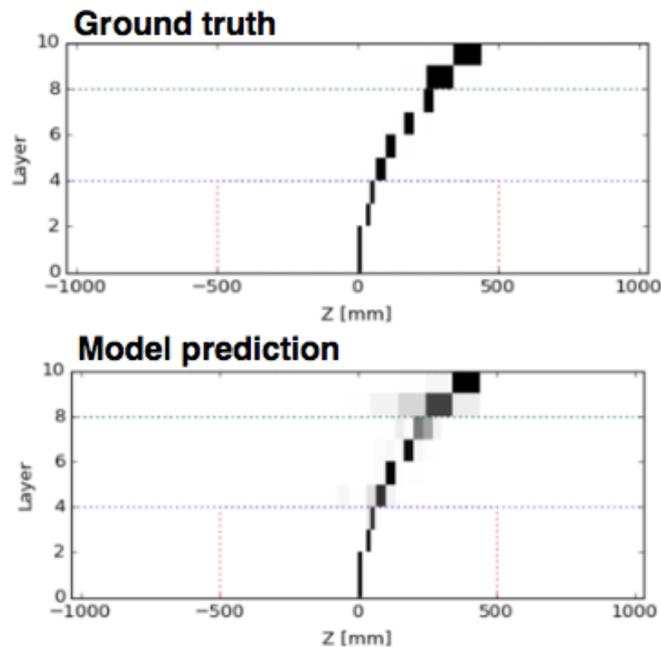
Flavour Tagging



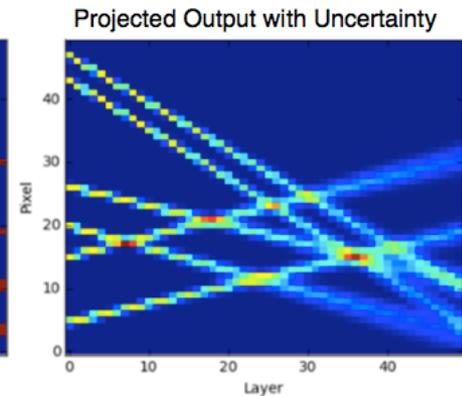
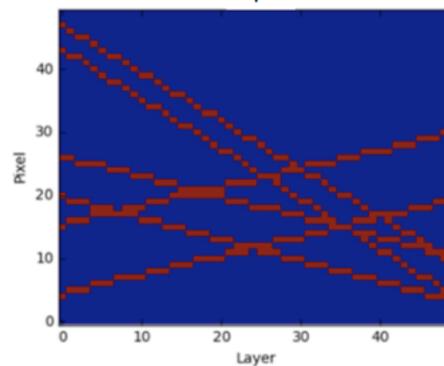
Tracking



Tracking with recurrent nets (LSTM)

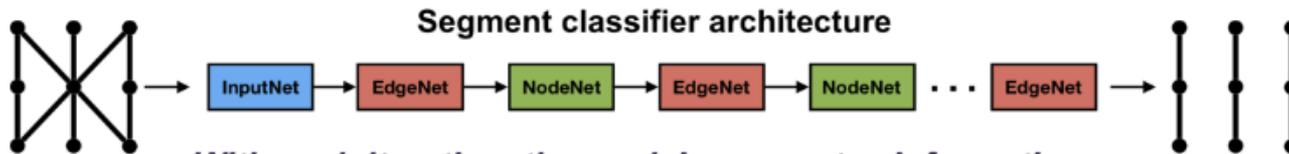


Time dimension
(state memory)



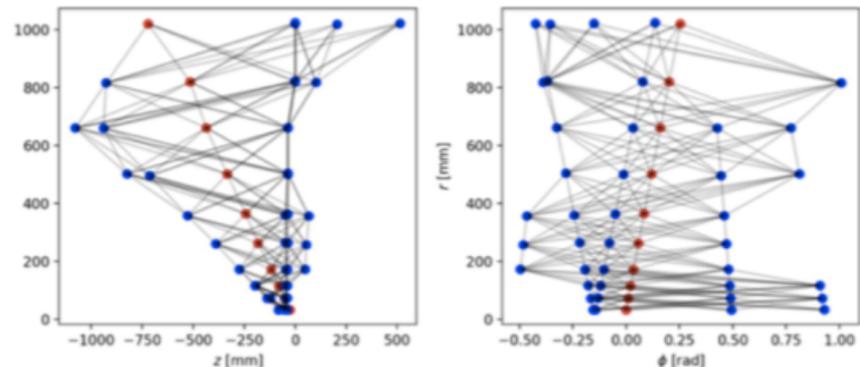
HEP.TrkX, CHEP 2018

Tracking with Graph Neural Network



With each iteration, the model propagates information through the graph, strengthens important connections, and weakens useless ones.

- Unseeded hit-pair classification
- Model predicts the probability that a hit-pair is valid

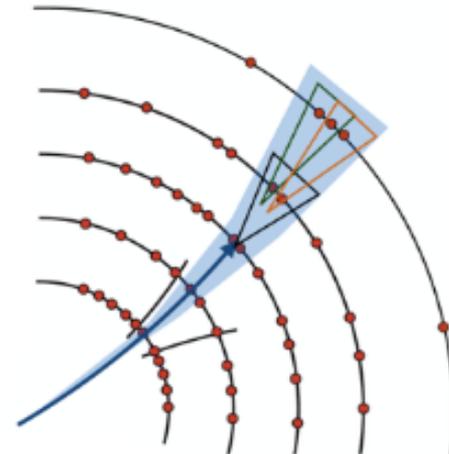
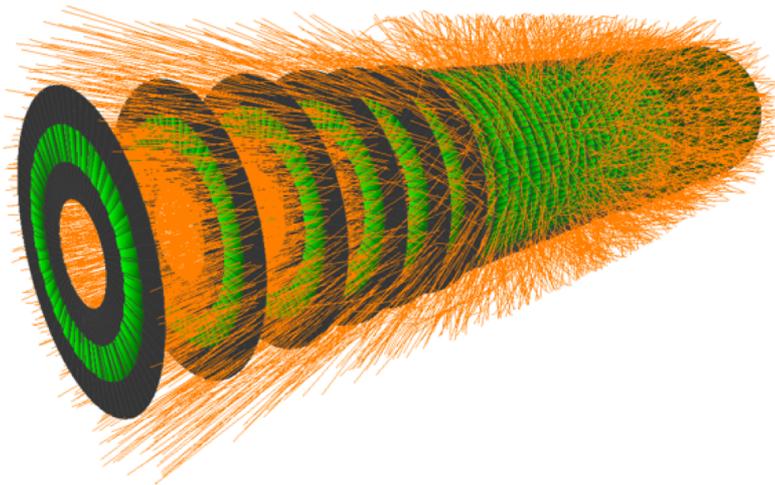


HEP.TrkX, CHEP 2018

TrackML Challenge



<https://www.kaggle.com/c/trackml-particle-identification>



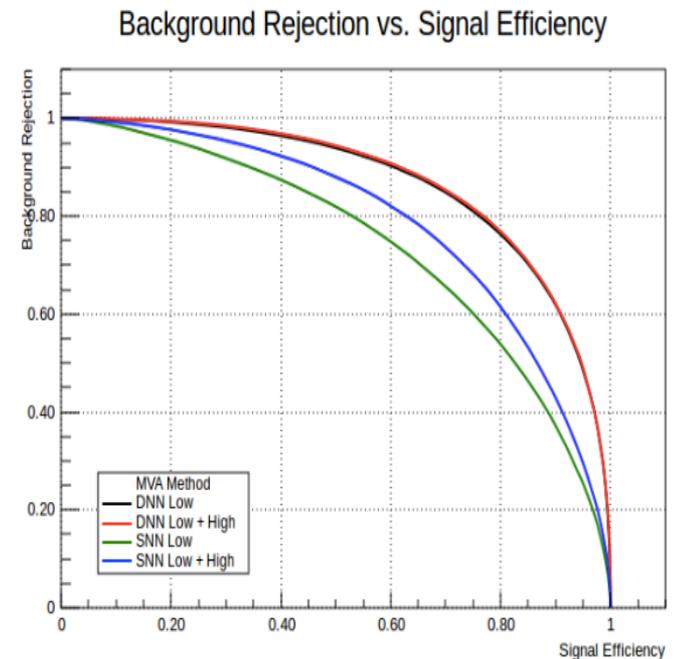


Can we extract features with meaningful physics?

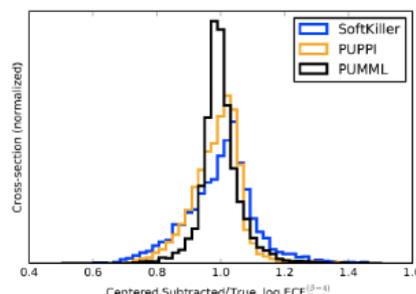
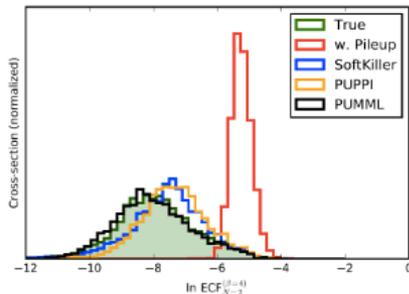
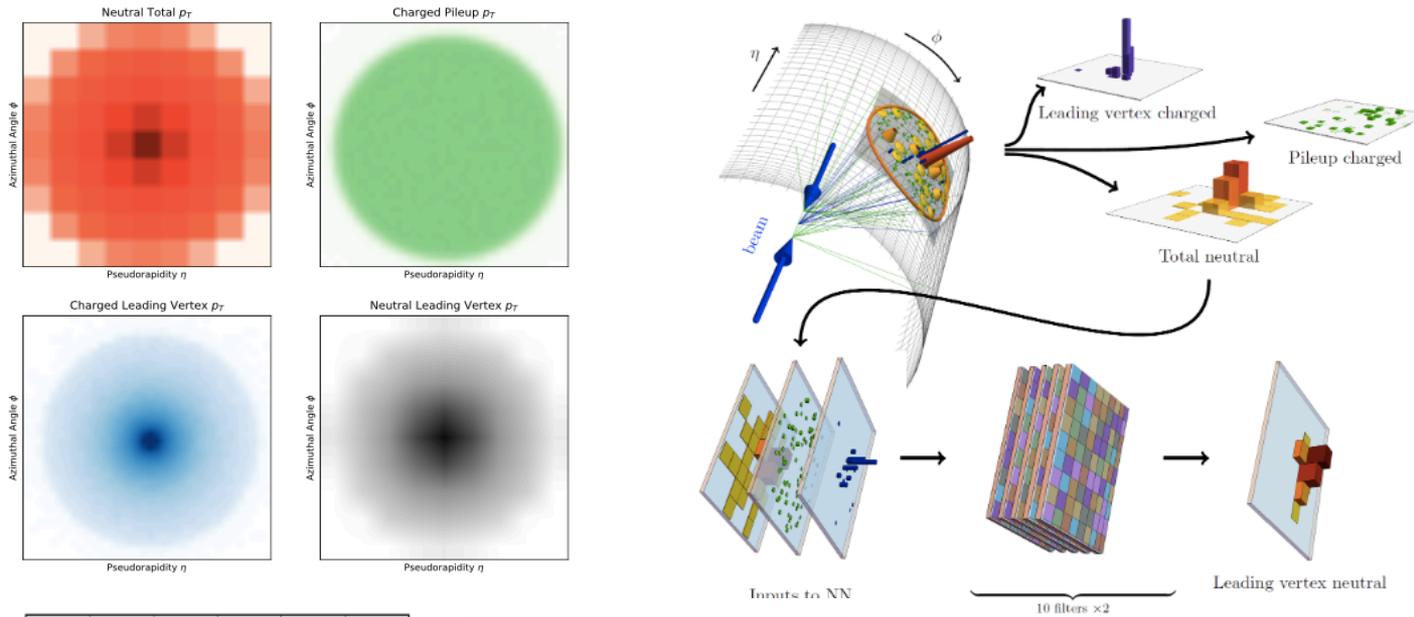
- from low-level variables

Are we able to understand ML models

- physics interpretations

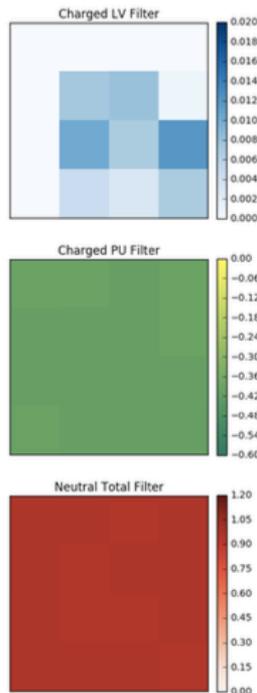


Pile-up removal with CNN



Komiske et al., 2017

Pile-up removal with CNN



What is learned?

- Train a single 4×4 filter and inspect it.
- Pixel-wise: $p_T^{N,LV} \approx p_T^{N,tot} - \frac{1}{2}p_T^{C,PU}$
- This is linear cleansing with $\bar{\gamma}_0 = 2/3!$

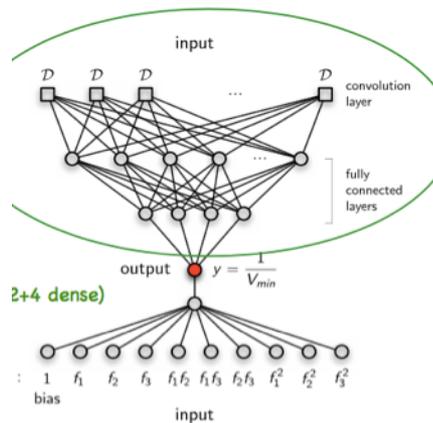
$$p_T^{N,LV} = p_T^{N,tot} + \left(1 - \frac{1}{\bar{\gamma}_0}\right)p_T^{C,PU}$$

Komiske et al., 2017

How to best use domain knowledge we have accumulated?

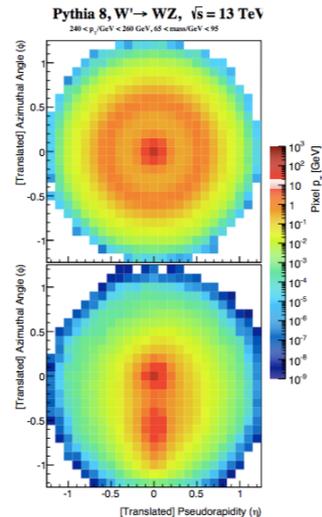
- in designing the algorithms

Strings



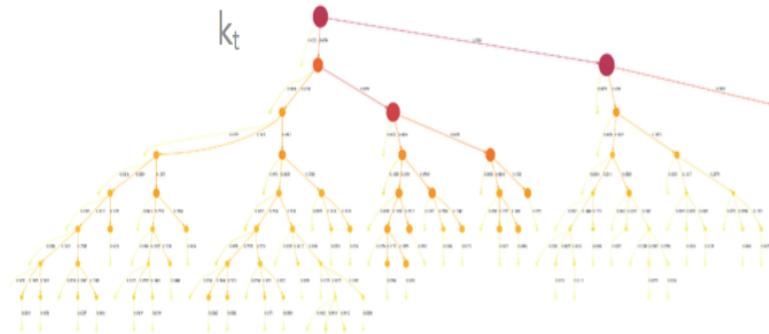
Krefl, 2017

Jet Images



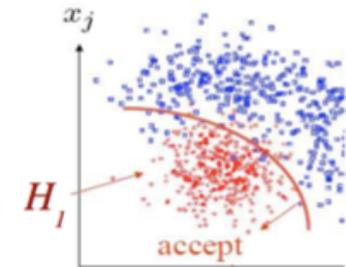
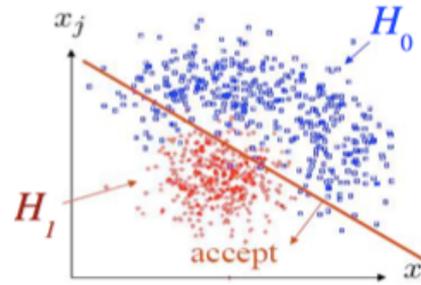
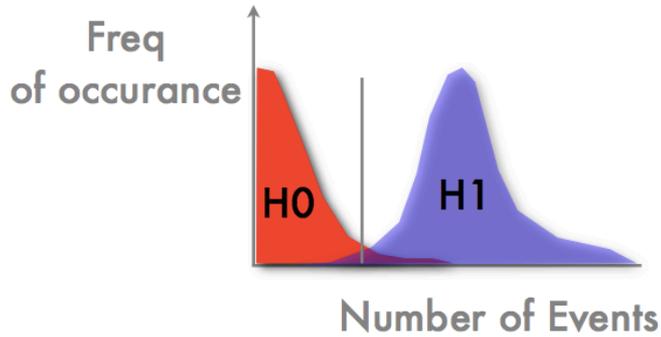
de Oliveira et al., 2015

Jet Clustering

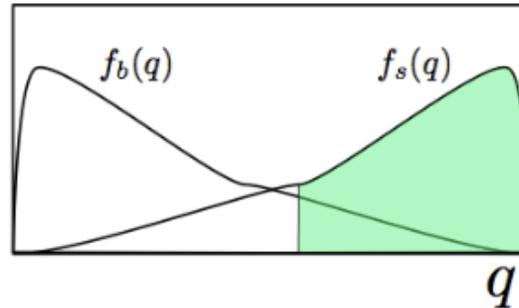


Loupe et al., 2017

Uncertainties

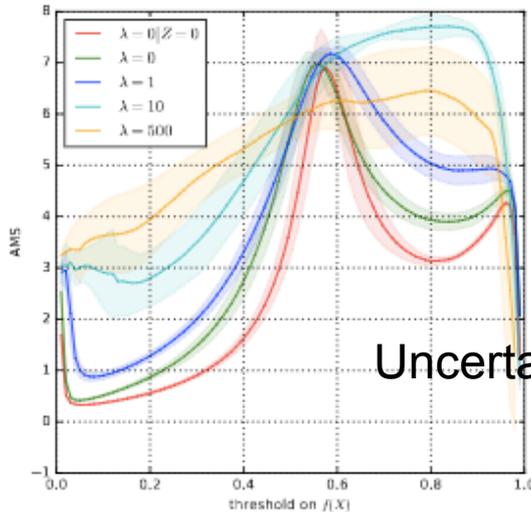


A threshold makes sense.

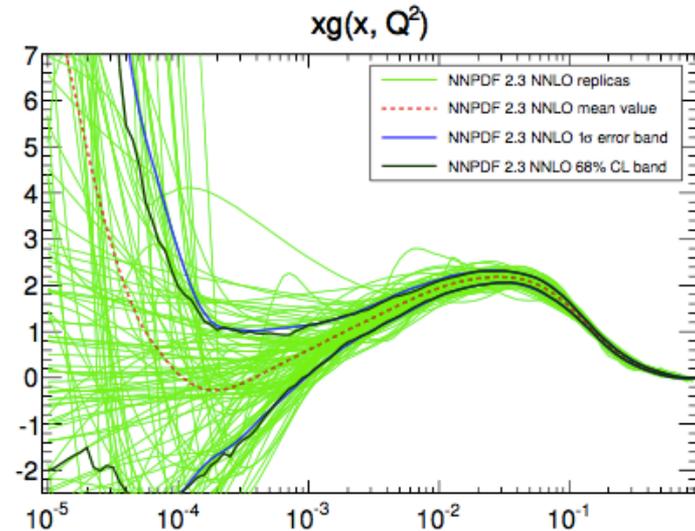


- Decision making**

G. Louppe et al., 2016



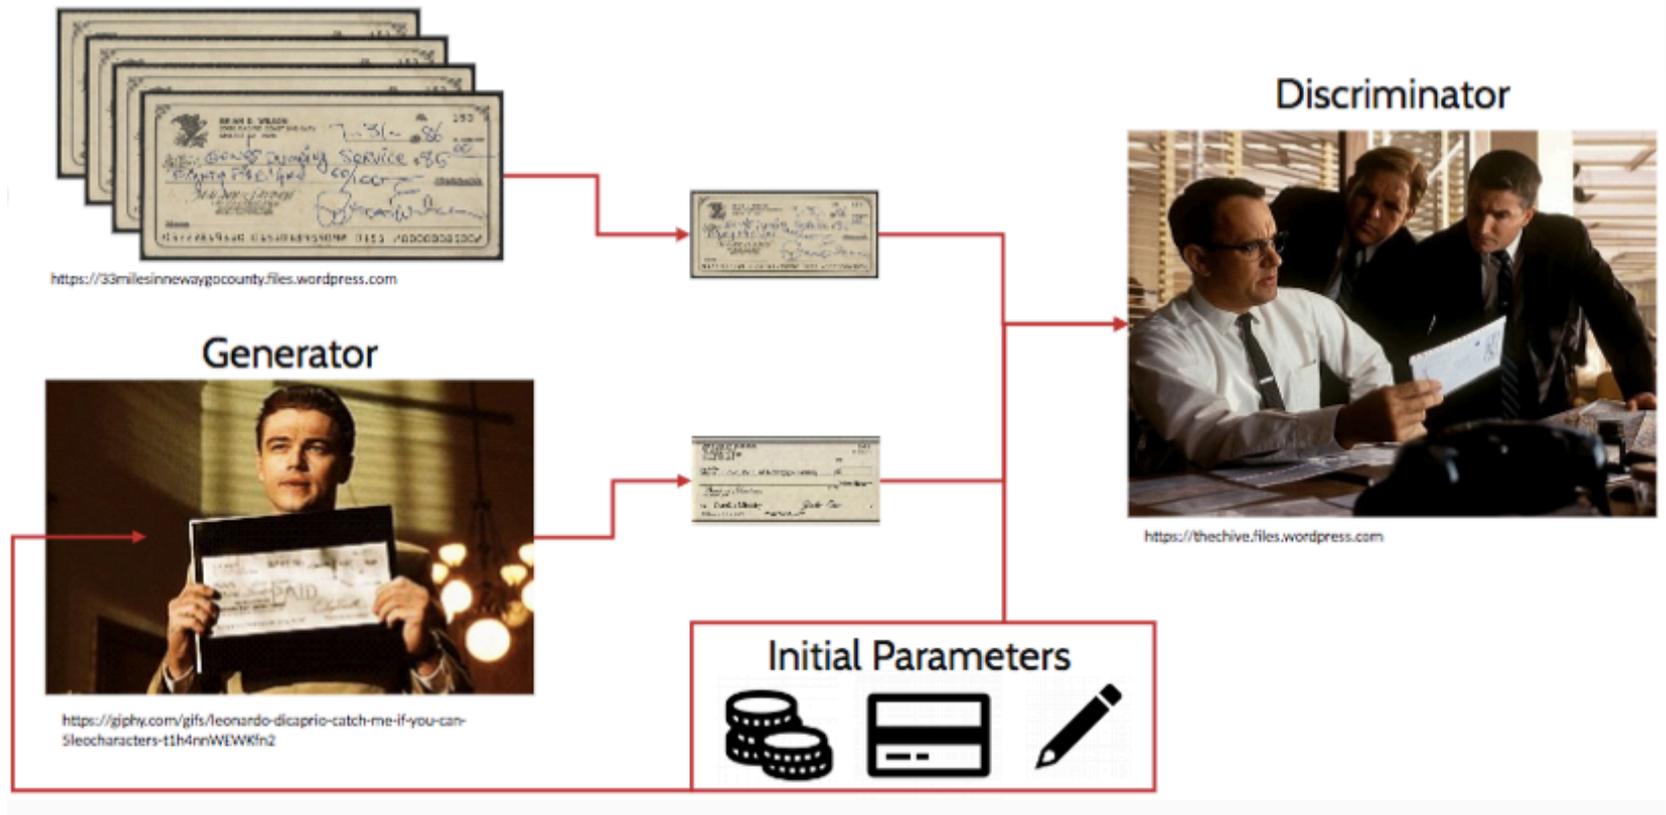
NNPDF Collaboration



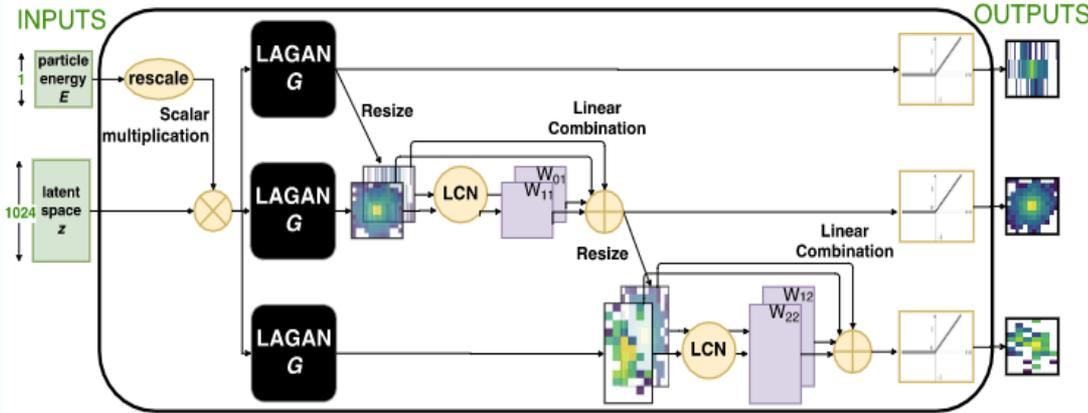
Bayesian connection: Deep neural networks with drop-out approximate variational inference of Bayesian NNs: *Gal and Ghahramani, 2016*



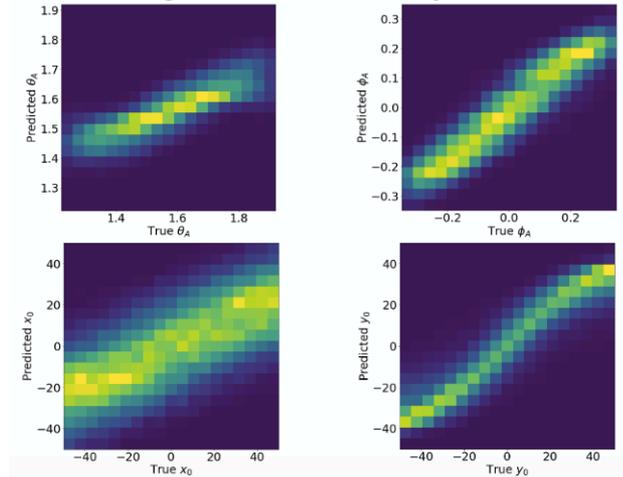
Generative Adversarial Networks



Simulation GANs

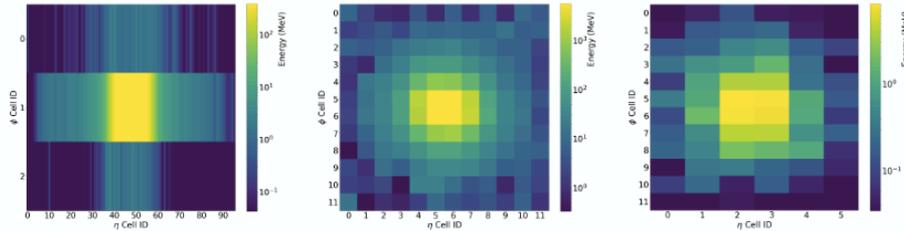


• Regression on $\theta, \varphi, x_0, y_0$

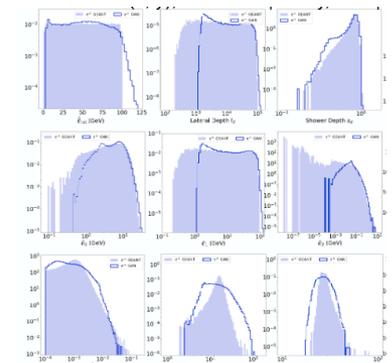
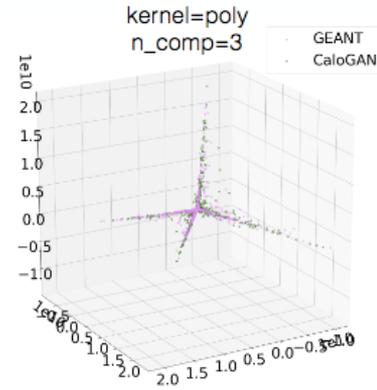
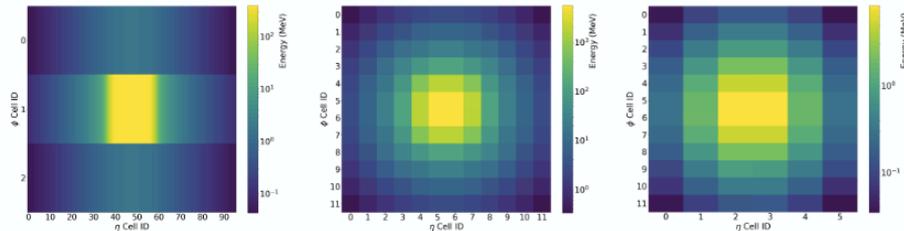


Dataset: 5°; Net: soft sparsity, multiplied E, Conv. attn. and layers

• CaloGAN



• GEANT



L. de Oliveira et al., 2017

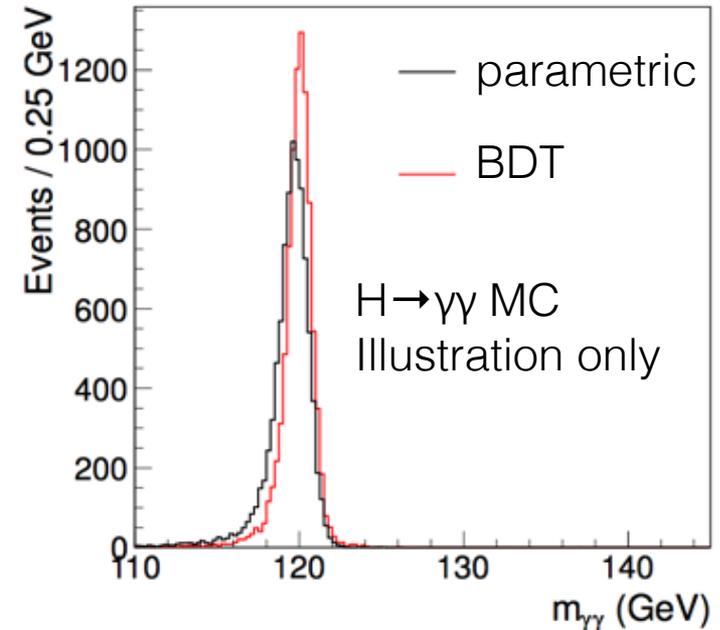
Inputs:

photon coordinates
photon shower information
median event energy

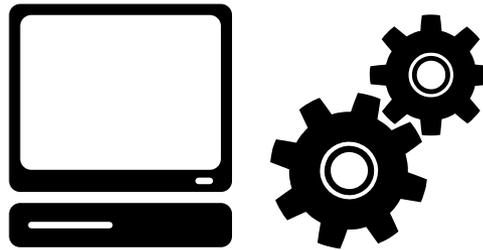
Target Output:

$$E_{\text{MEASURED}}/E_{\text{TRUE}}$$

10-30% improvement with shallow ML



Modify evaluation in induction algorithm

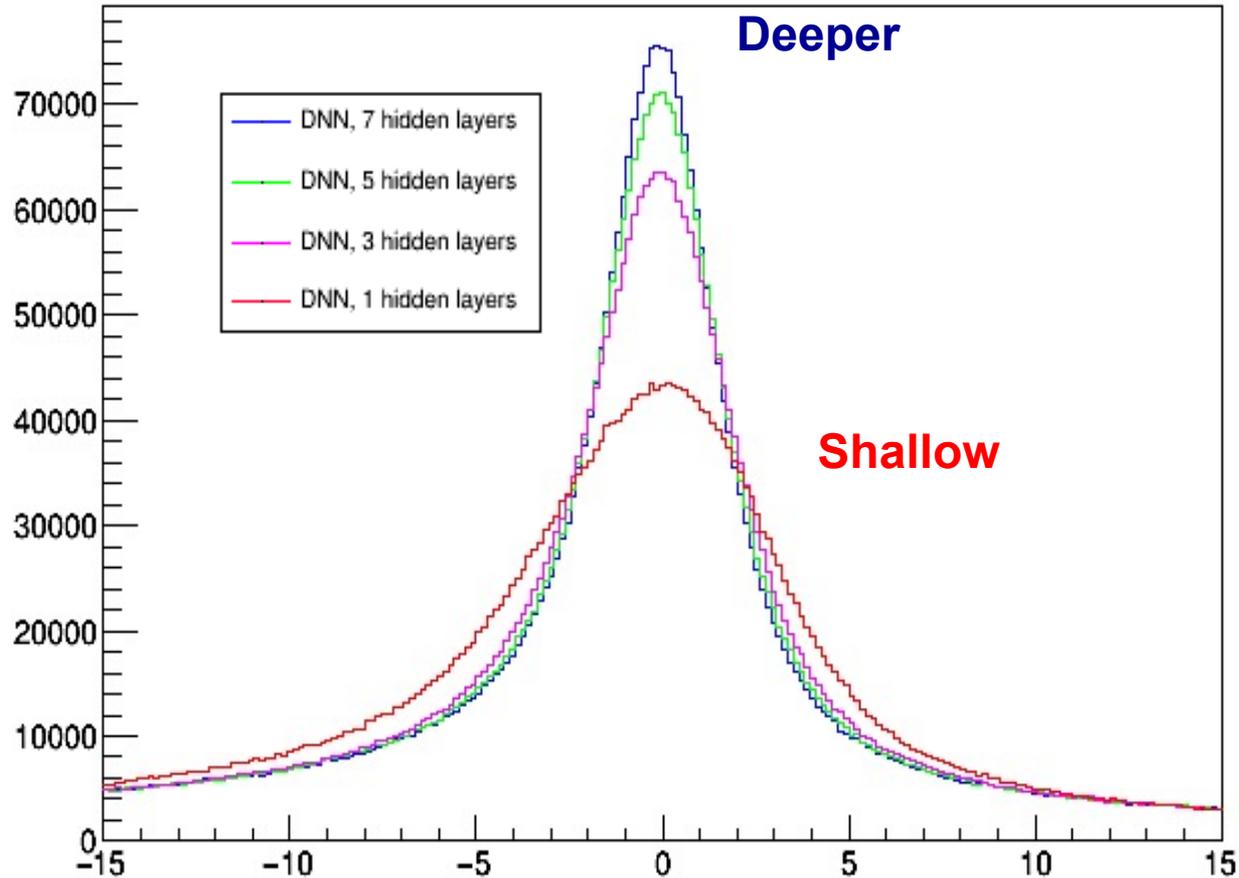


Maximum separation



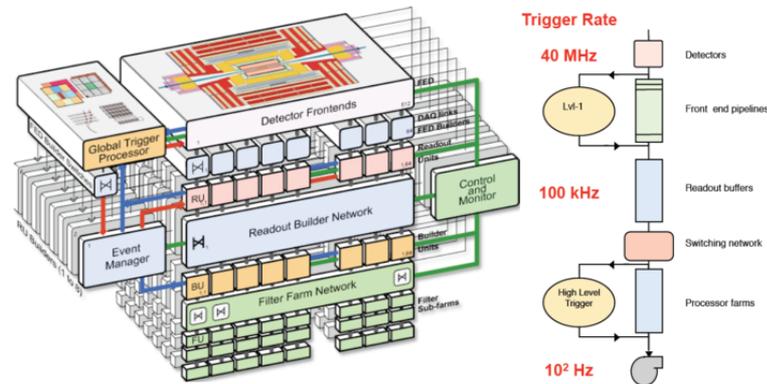
Minimal variance

Prediction Error

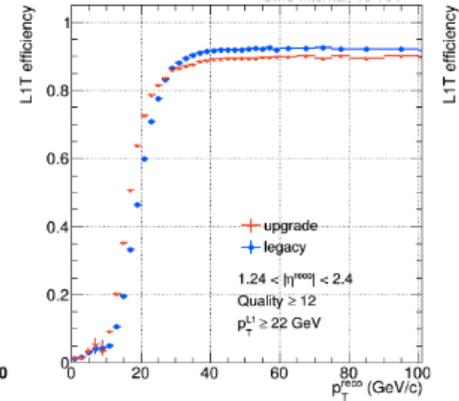
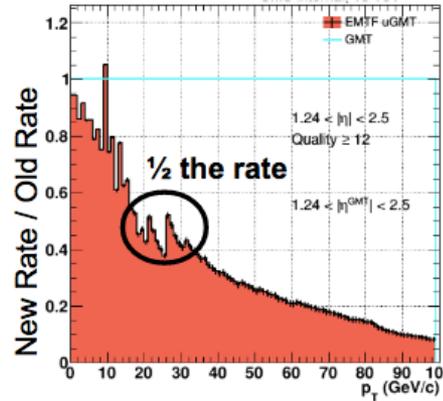
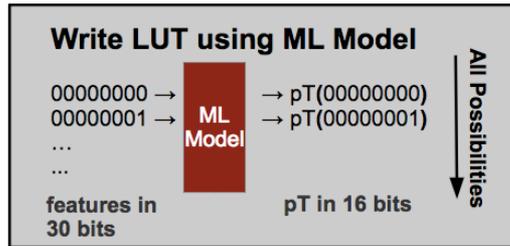


Can we do ML in **real-time?**

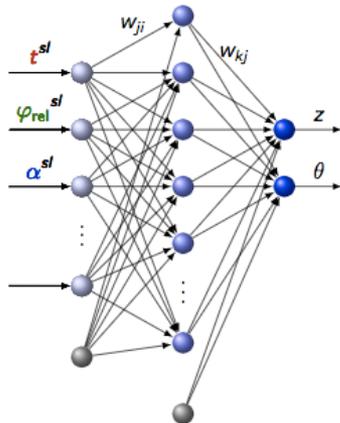
- ML: live video analysis, medical, self-driving cars
- HEP **Trigger Systems** (software and hardware)



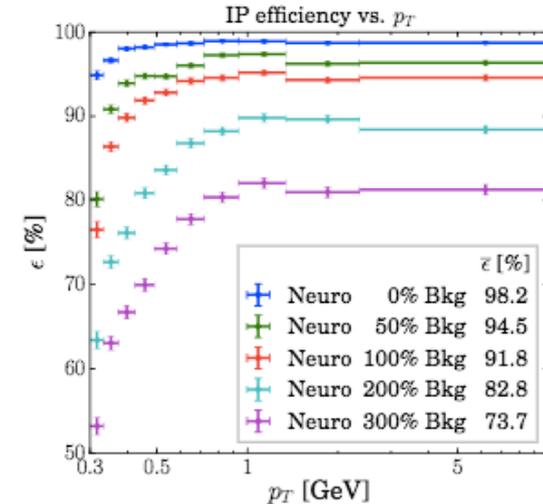
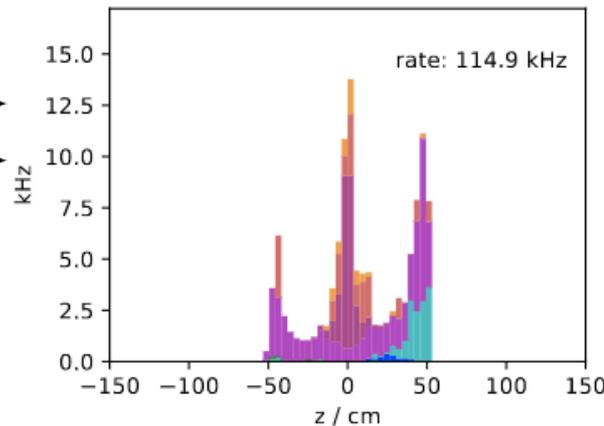
CMS L1



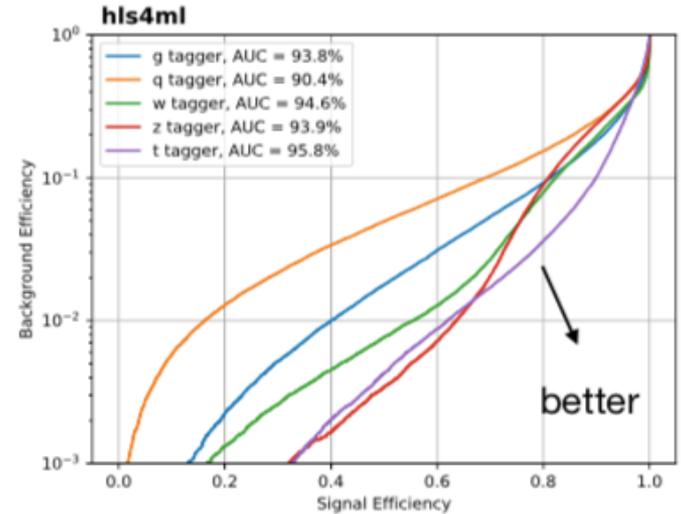
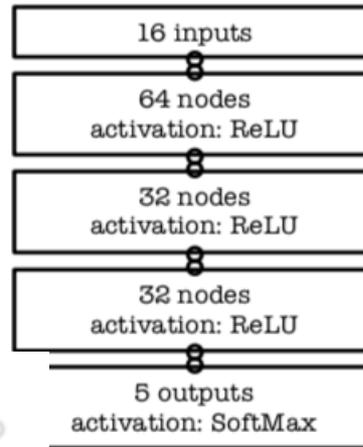
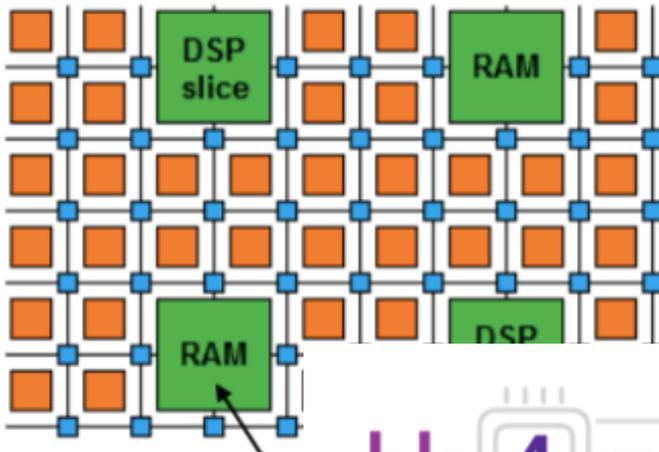
Belle II



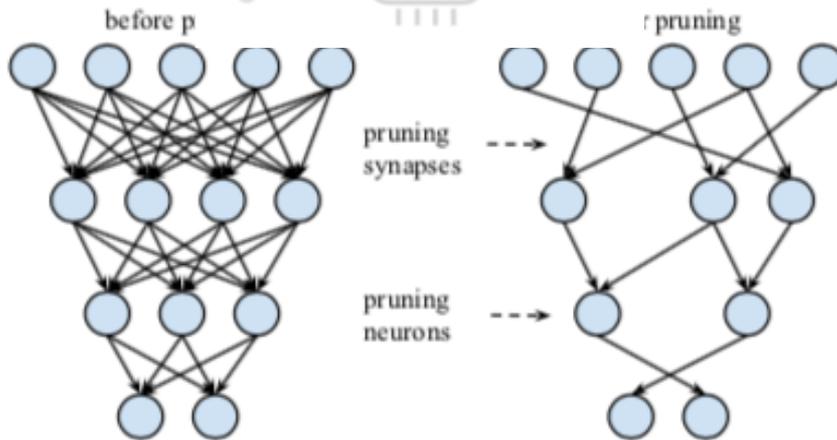
Neural Network Track Estimates



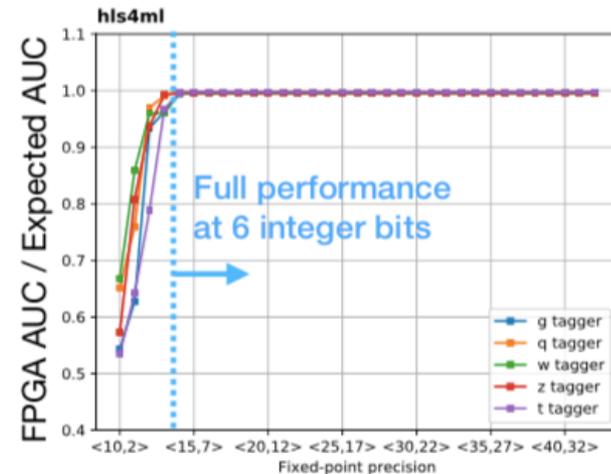
ML on FPGA



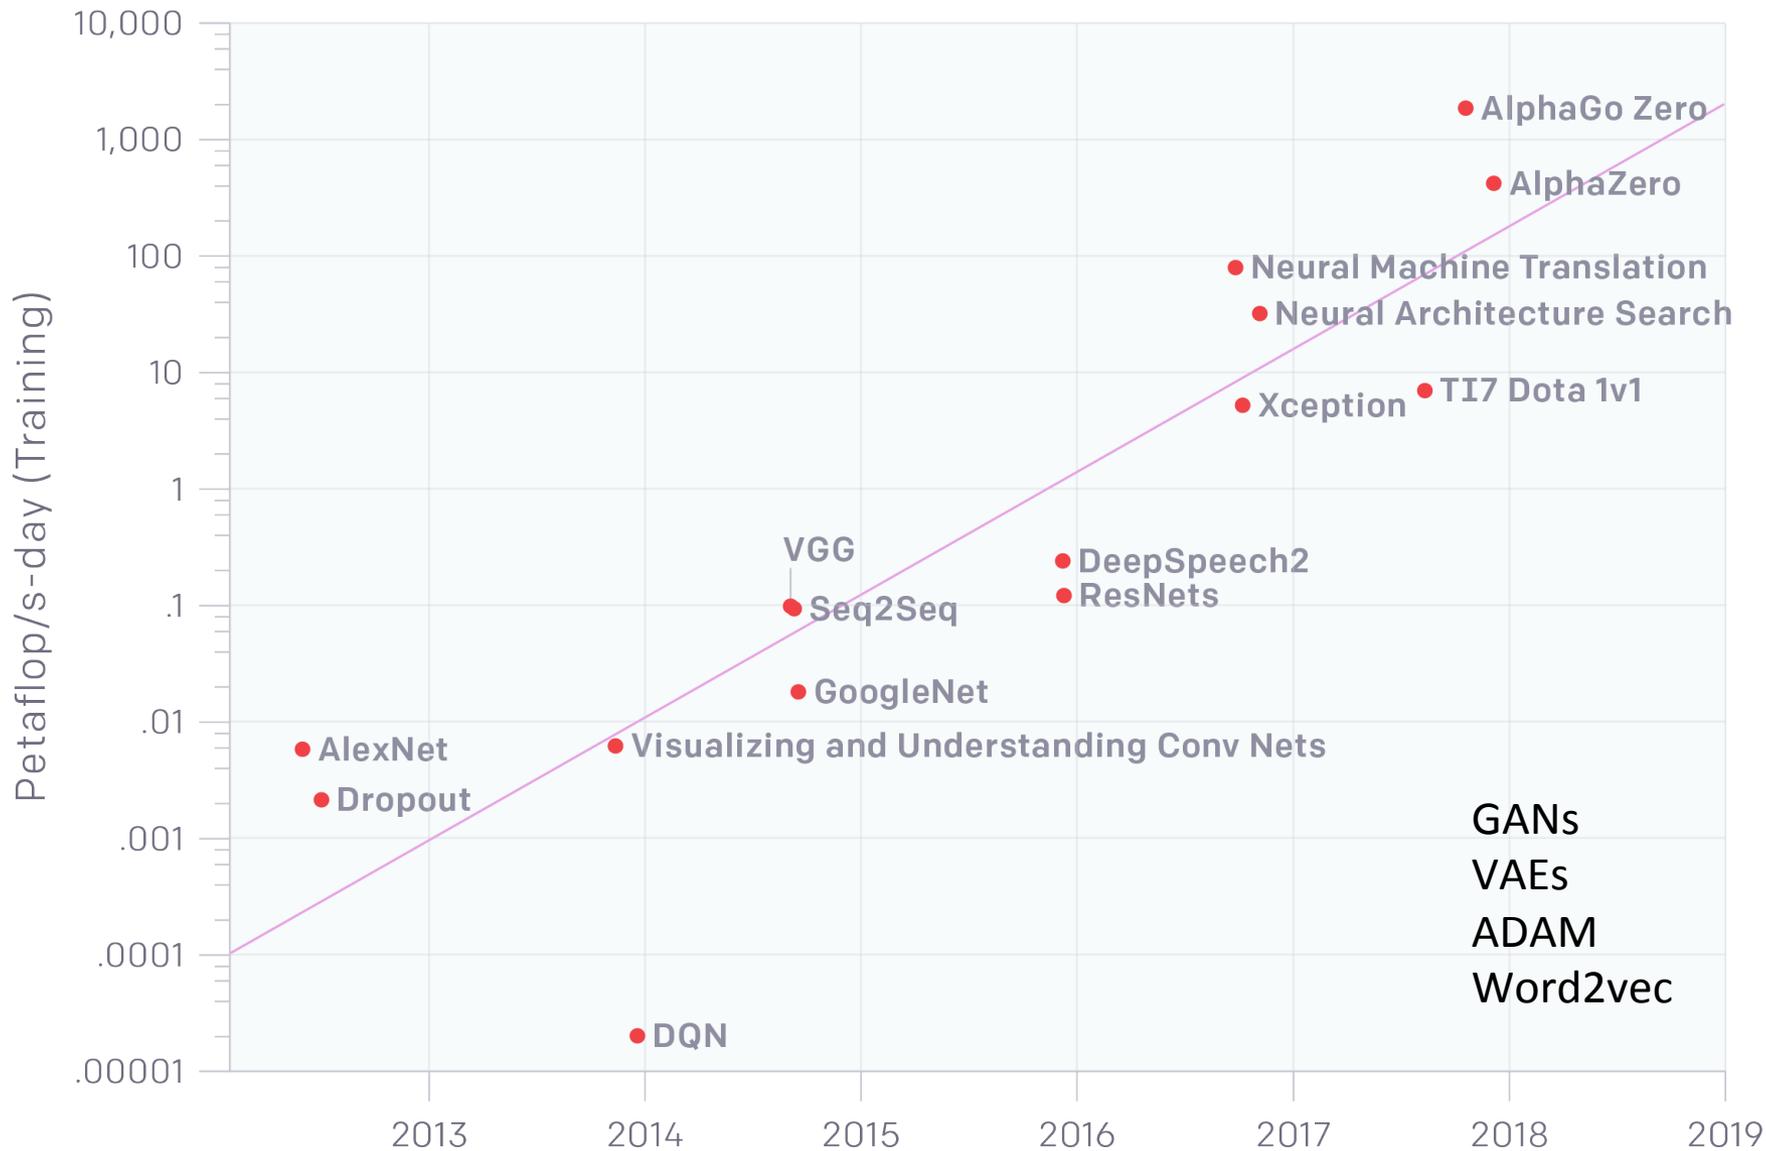
hls4ml



Fractional bits fixed to 8



AlexNet to AlphaGo Zero: A 300,000x Increase in Compute



Machine Learning in High Energy Physics Community White Paper

July 10, 2018

Abstract: Machine learning is an important applied research area in particle physics, beginning with applications to high-level physics analysis in the 1990s and 2000s, followed by an explosion of applications in particle and event identification and reconstruction in the 2010s. In this document we discuss promising future research and development areas in machine learning in particle physics with a roadmap for their implementation, software and hardware resource requirements, collaborative initiatives with the data science community, academia and industry, and training the particle physics community in data science. The main objective of the document is to connect and motivate these areas of research and development with the physics drivers of the High-Luminosity Large Hadron Collider and future neutrino experiments and identify the resource needs for their implementation. Additionally we identify areas where collaboration with external communities will be of great benefit.

Editors: Sergei Gleyzer²⁶, Paul Seyfert¹¹, Steven Schramm²⁸

[arXiv:1807.02876](https://arxiv.org/abs/1807.02876)

Contributors: Kim Albertsson¹, Piero Altoc², Dustin Anderson³, Michael Andrews⁴, Juan Pedro Araque Espinosa⁵, Adam Aurisano⁶, Laurent Basara⁷, Adrian Bevan⁸, Wahid Bhimji⁹, Daniele Bonacorsi¹⁰, Paolo Calafiura⁹, Mario Campanelli⁸, Louis Capps², Federico Carminati¹¹, Stefano Carrazza¹¹, Taylor Childers¹², Elias Coniavitis¹³, Kyle Cranmer¹⁴, Claire David¹⁵, Douglas Davis¹⁶, Javier Duarte¹⁷, Martin Erdmann¹⁸, Jonas Eschle¹⁹, Amir Farbin²⁰, Matthew Feickert²¹, Nuno Filipe Castro⁵, Conor Fitzpatrick²², Michele Floris¹¹, Alessandra Forti²³, Jordi Garra-Tico²⁴, Jochen Gemmler²⁵, Maria Girone¹¹, Paul Glaysher¹⁵, Sergei Gleyzer²⁶, Vladimir Gligorov²⁷, Tobias Golling²⁸, Jonas Graw², Lindsey Gray¹⁷, Dick Greenwood²⁹, Thomas Hacker³⁰, John Harvey¹¹, Benedikt Hegner¹¹, Lukas Heinrich¹⁴, Ben Hooberman³¹, Johannes Junggeburth³², Michael Kagan³³, Meghan Kane³⁴, Konstantin Kanishchev⁷, Przemyslaw Karpiński¹¹, Zahari Kassabov³⁵, Gaurav Kaul³⁶, Dorian Keira³, Thomas Keck²⁵, Alexei Klimentov³⁷, Jim Kowalkowski¹⁷, Luke Kreczko³⁸, Alexander Kurepin³⁹, Rob Kutschke¹⁷, Valentin Kuznetsov⁴⁰, Nicolas Köhler³², Igor Lakomov¹¹, Kevin Lannon⁴¹, Mario Lassnig¹¹, Antonio Limosani⁴², Gilles Louppe¹⁴, Aashrita Mangu⁴³, Pere Mato¹¹, Helge Meinhart¹¹, Dario Menasce⁴⁴, Lorenzo Moneta¹¹, Seth Moortgat⁴⁵, Meenakshi Narain⁴⁶, Mark Neubauer³¹, Harvey Newman³, Hans Pabst³⁶, Michela Paganini⁴⁷, Manfred Paulini⁴, Gabriel Perdue¹⁷, Uzziel Perez⁴⁸, Attilio Picazio⁴⁹, Jim Pivarski⁵⁰, Harrison Prosper⁵¹, Fernanda Psihas⁵², Alexander Radovic⁵³, Ryan Reece⁵⁴, Aurelius Rinkevicius⁴⁰, Eduardo Rodrigues⁶, Jamal Rorie⁵⁵, David Rousseau⁵⁶, Aaron Sauers¹⁷, Steven Schramm²⁸, Ariel Schwartzman³³, Horst Severini⁵⁷, Paul Seyfert¹¹, Filip Siroky⁵⁸, Konstantin Skazytkin³⁹, Mike Sokoloff⁶, Graeme Stewart⁵⁹, Bob Stienen⁶⁰, Ian Stockdale⁶¹, Giles Strong⁵, Savannah Thais¹⁷, Karen Tomko⁶², Eli Upfal⁴⁶, Emanuele Usai⁴⁶, Andrey Ustyuzhanin⁶³, Martin Vala⁶⁴, Sofia Vallecorsa⁶⁵, Mauro Verzetti⁶⁶, Xavier Vilasis-Cardona⁶⁷, Jean-Roch Vlimant³, Ilija Vukotic⁶⁸, Sean-Jiun Wang²⁶, Gordon Watts⁶⁹, Michael Williams⁷⁰, Wenjing Wu⁷¹, Stefan

Summary



- **Machine learning already at the forefront of high-energy physics applications**
 - Improvements in **particle and event classification**, and **property measurements**
 - **Feature extraction**
 - Estimating and incorporating **uncertainty**
 - Emerging machine learning applications for **fast inference** in hardware



Thank You



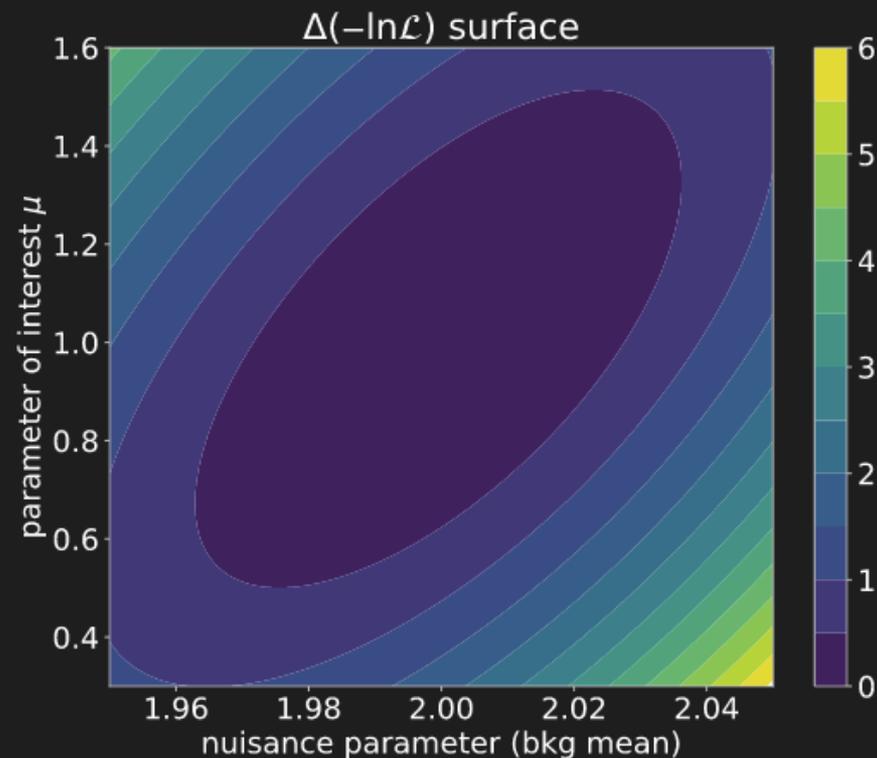
INFERENCE-MOTIVATED LOSS FUNCTION

If we expand the negative log-likelihood around minimum (e.g. Asimov $n_i = \alpha_s \cdot s_i + \alpha_b \cdot b_i$), due to Cramér-Rao bound:

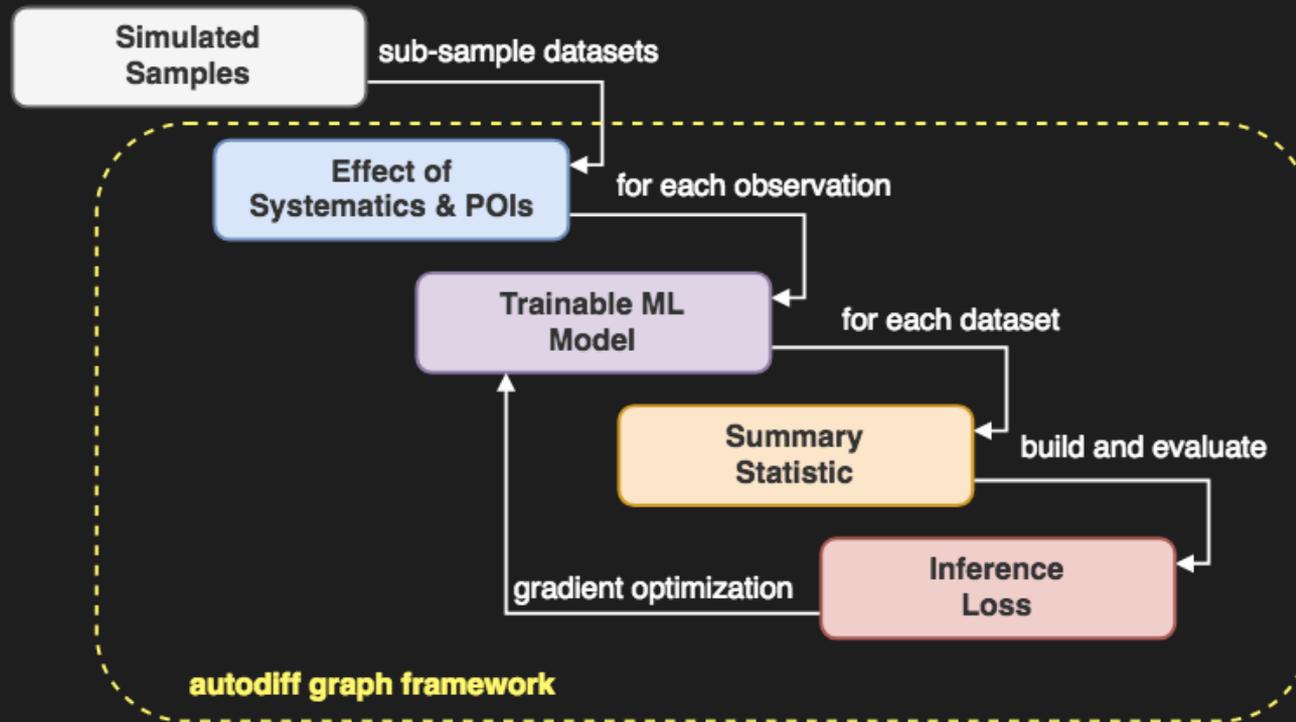
$$\text{covariance} \geq \mathbf{H}^{-1}(-\ln \mathcal{L})$$

which can be computed via autodiff. Can use as loss function directly the variance bound on the parameters of interest

$$\text{loss} \approx \text{Var}(\mu) \quad (\text{expected})$$



END-TO-END DIFFERENTIABILITY FOR LHC ANALYSES



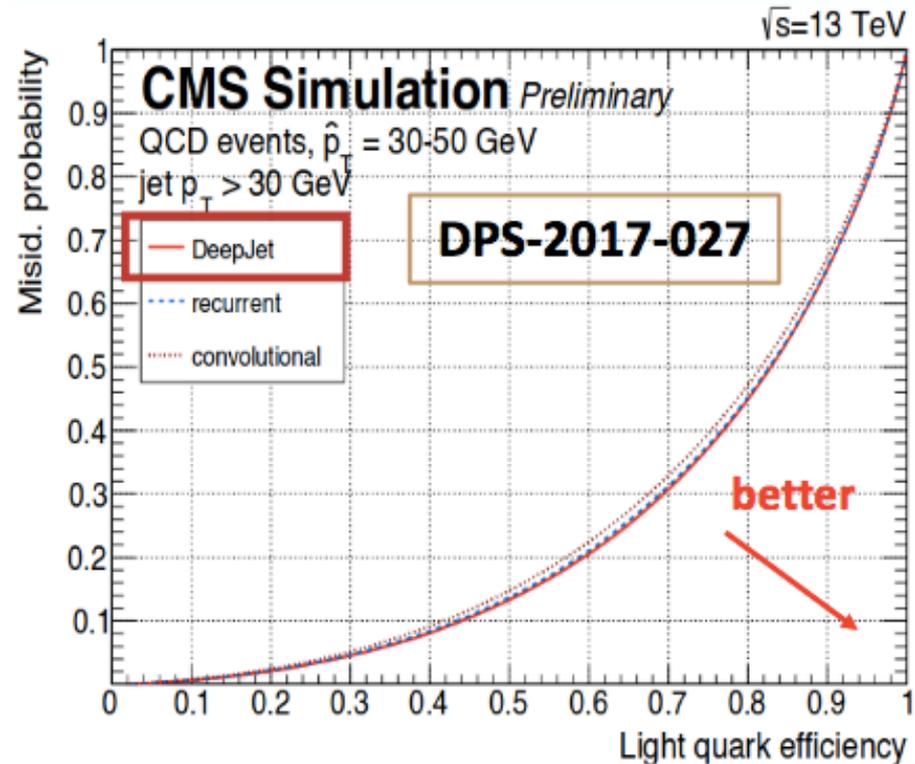
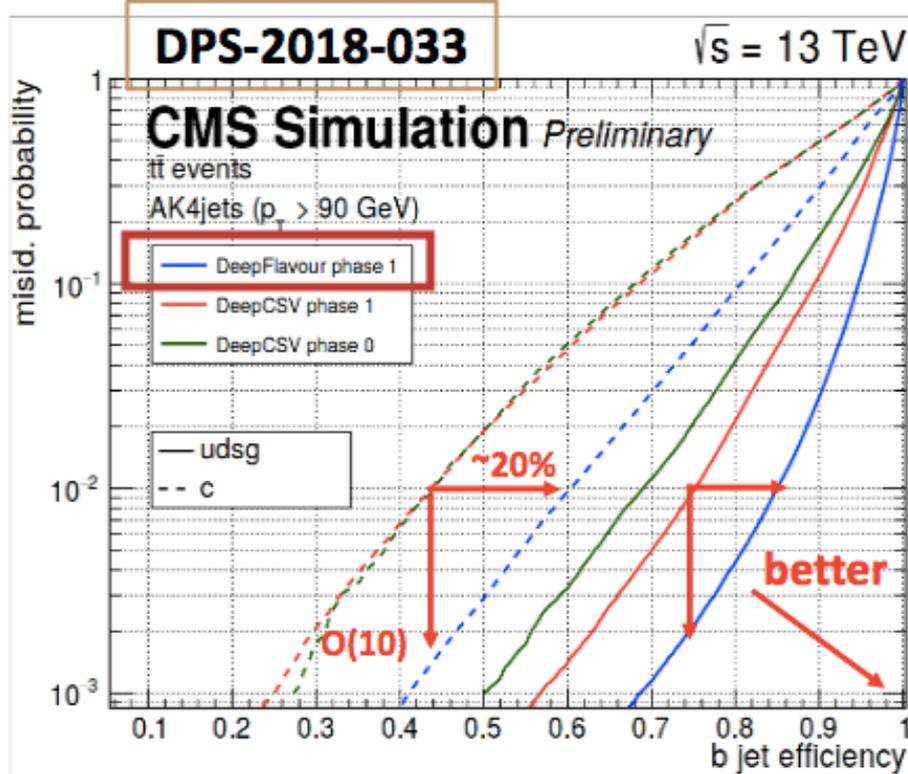
Within this general framework, several approaches are possible, focus here is

DIRECT LEARNING OF SYSTEMATICS-AWARE SUMMARY STATISTICS

Flavor Tagging

b vs. udsg / b vs. c

Quark – gluon separation



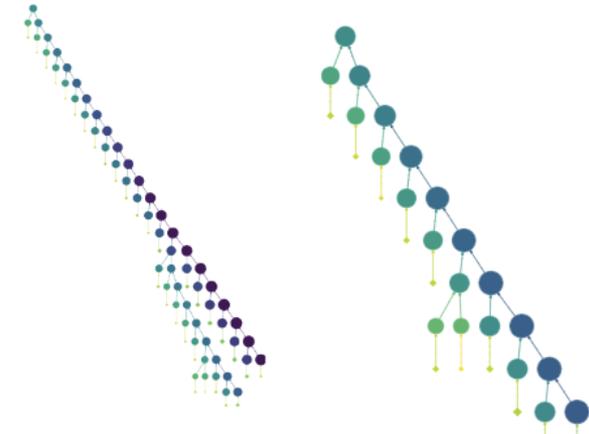
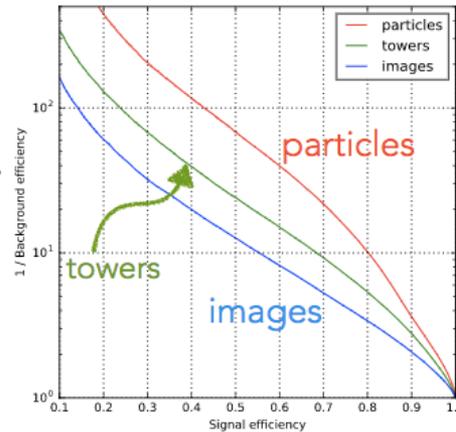
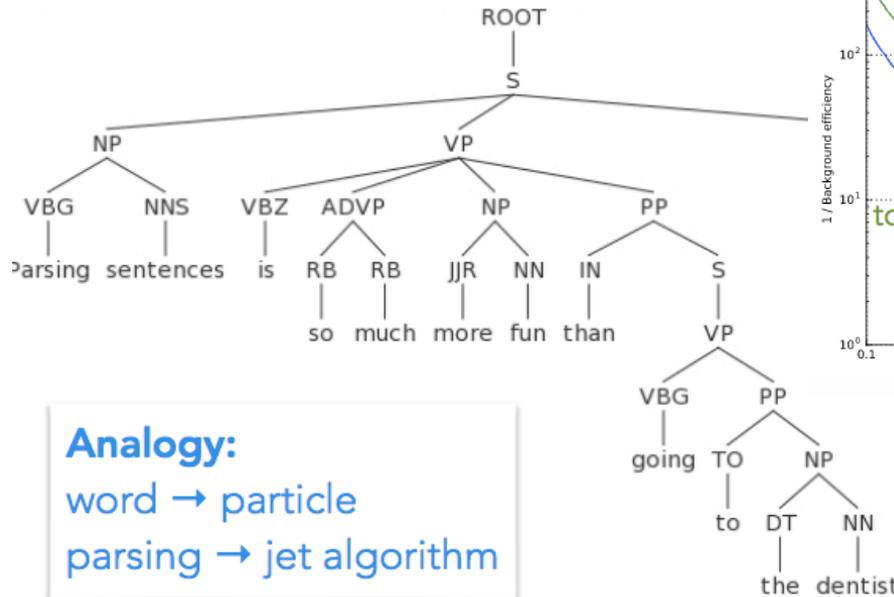
→ Significant gain in performance
 even more significant at higher p_T
 → Large part of the performance loss of
 previous [non particle-based] taggers
 was due to track preselection

→ Generator level light quarks/gluons
 that did not split to heavy flavour
 → Similar performance to simpler &
 dedicated architectures

Jet Algorithm



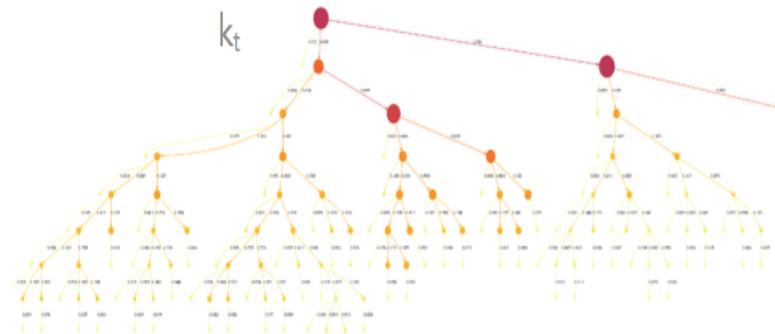
Natural Language Processing, RNNs and Jets



T. Cheng, 2017

Analogy:
word → particle
parsing → jet algorithm

G. Louppe et al., 2017





Butter, Kasieczka, Plehn, MR '17

Beyond images: LoLa

Why not use the jet constituent 4-vectors directly?

Two ingredients:

1. CoLa* - learns the jet clustering history

$$k_{\mu,i} \xrightarrow{\text{CoLa}} \tilde{k}_{\mu,j} = k_{\mu,i} C_{ij}$$

- Test on-shell conditions

$$\tilde{k}_{\mu,1}^2 = (k_{\mu,1} + k_{\mu,2} + k_{\mu,3})^2 = m_t^2$$

$$\tilde{k}_{\mu,2}^2 = (k_{\mu,1} + k_{\mu,2})^2 = m_W^2.$$

$$C = \begin{pmatrix} 1 & 0 & \cdots & 0 & C_{1,N+2} & \cdots & C_{1,M} \\ 0 & 1 & & \vdots & C_{2,N+2} & \cdots & C_{2,M} \\ \vdots & \vdots & \ddots & 0 & \vdots & & \vdots \\ 0 & 0 & \cdots & 1 & C_{N,N+2} & \cdots & C_{N,M} \end{pmatrix}$$

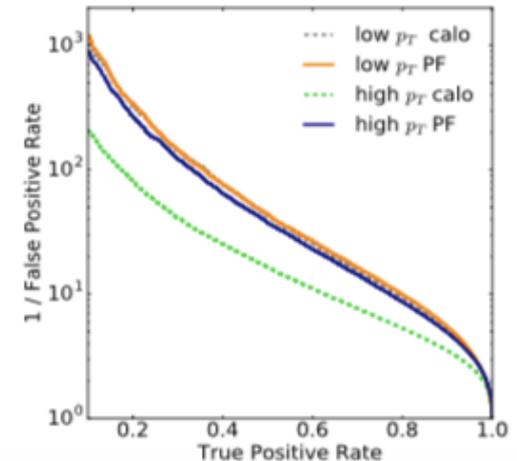
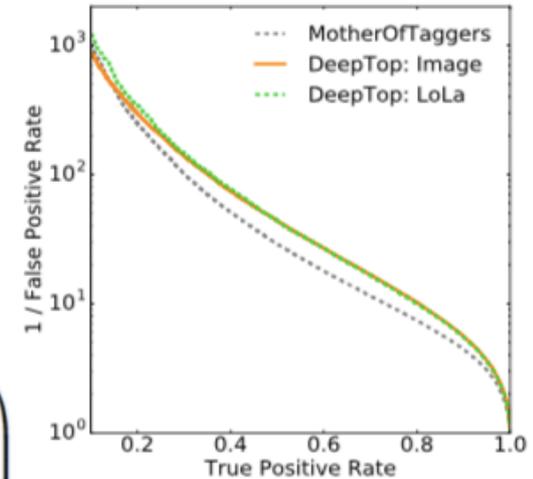
2. LoLa** - learns the kinematics

$$\tilde{k}_j \xrightarrow{\text{LoLa}} \hat{k}_j = \begin{pmatrix} m^2(\tilde{k}_j) \\ p_T(\tilde{k}_j) \\ w_{jm}^{(E)} E(\tilde{k}_m) \\ w_{jm}^{(d)} d_{jm}^2 \end{pmatrix}$$

transform 4-vectors into: invariant mass, p_T , energy and Minkowski distance
effectively a rotation in observable space

* CoLa = Combination Layer

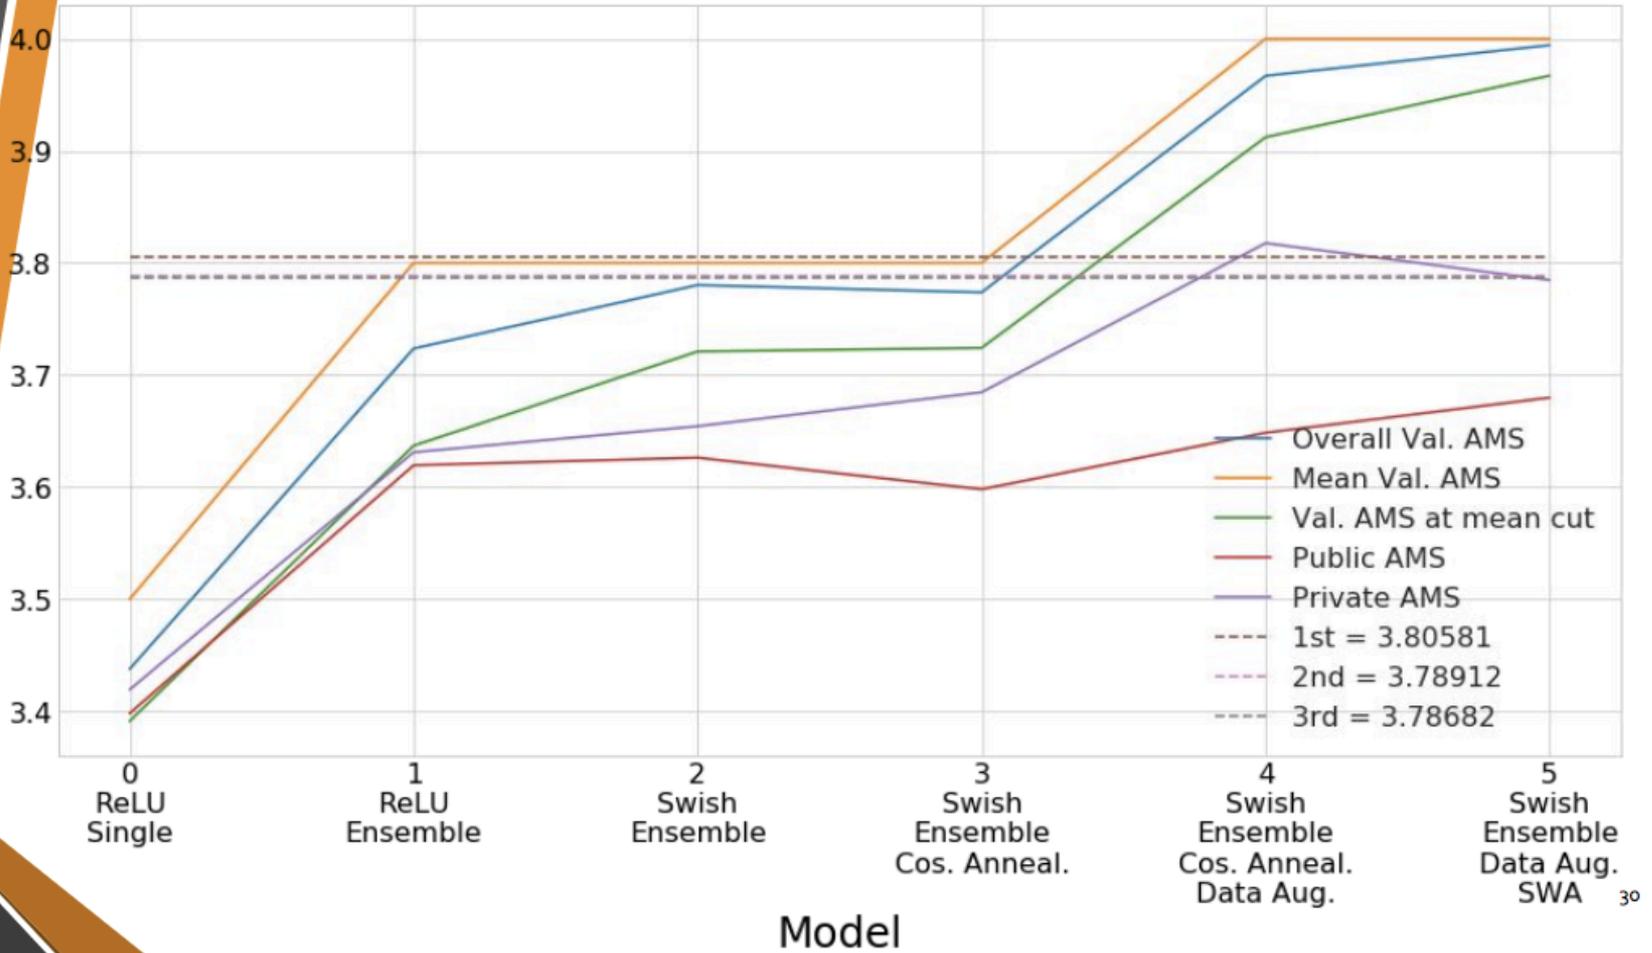
** LoLa = Lorentz Layer



HiggsML Optimization



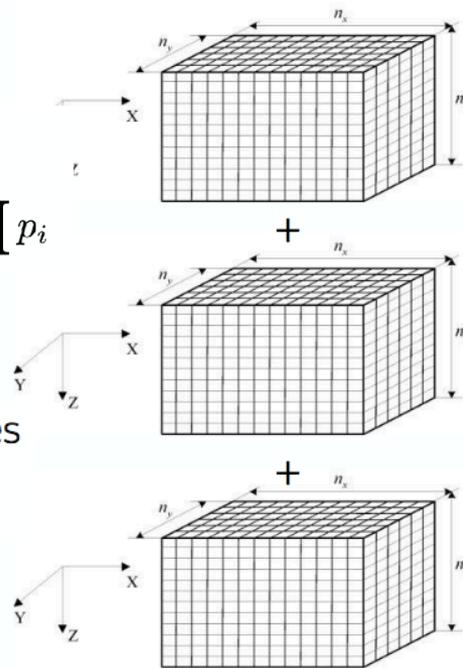
AMS



Trigger Applications

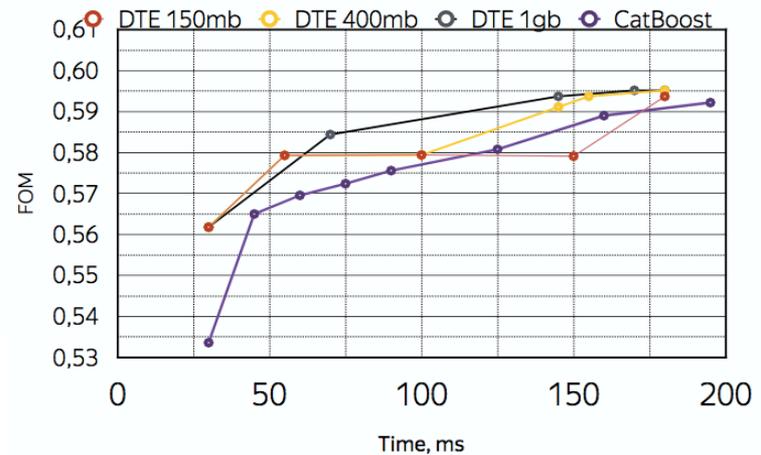
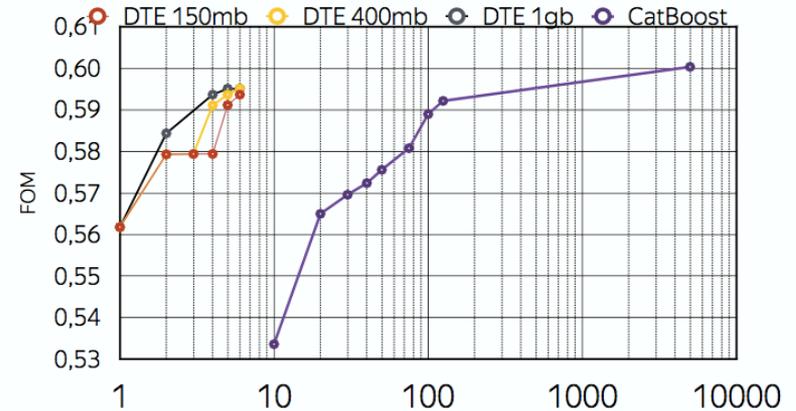


$\rangle D = (f_1, \dots, f_n)$
 $f_i = \{b_{i1}, \dots, b_{ip_i}\}$
 $S(D) \propto \prod |f_i| = \prod p_i$



- > CatBoost
- > Uses oblivious trees
- > Discretize features

LHCb, ACAT 2017





ALICE

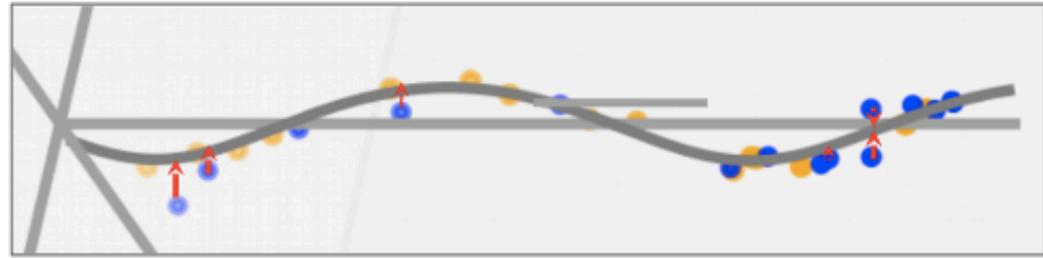
Results



- Mean Squared Error (MSE) from the original helix as a quality measure
- Evaluation conducted on the separate test-set with ~15000 tracks

MSE visualisation:

- Red - error
- Grey - ideal helix
- Orange - original clusters
- Blue - generated clusters

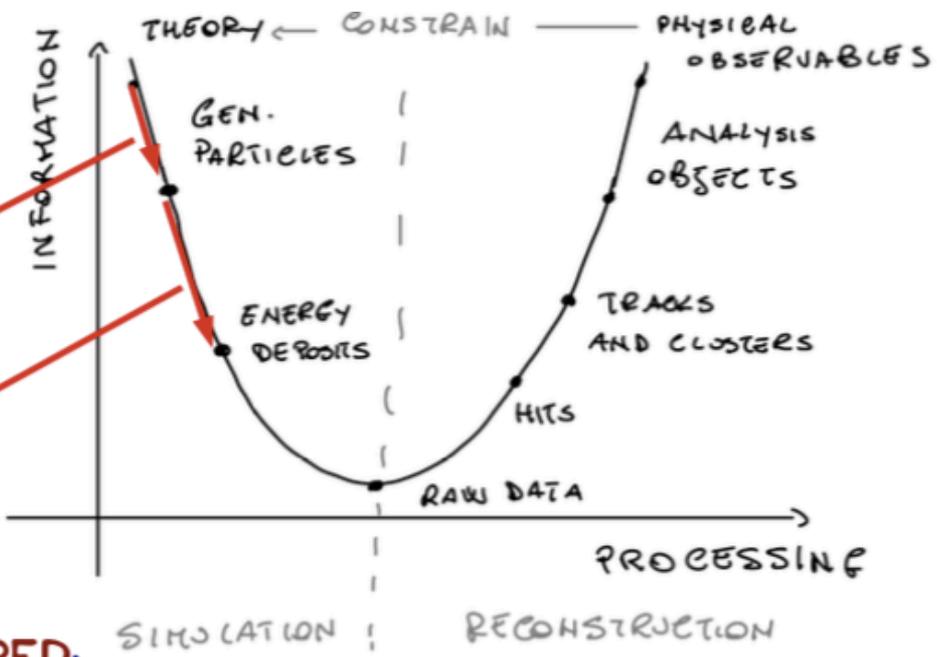


Method	Mean MSE (mm)	Median MSE (mm)	Speed-up
GEANT3	1.20	1.12	1
Random (estimated)	2500	2500	N/A
condLSTM GAN	2093.69	2070.32	100
condLSTM GAN+	221.78	190.17	
condDCGAN	795.08	738.71	25
condDCGAN+	136.84	82.72	

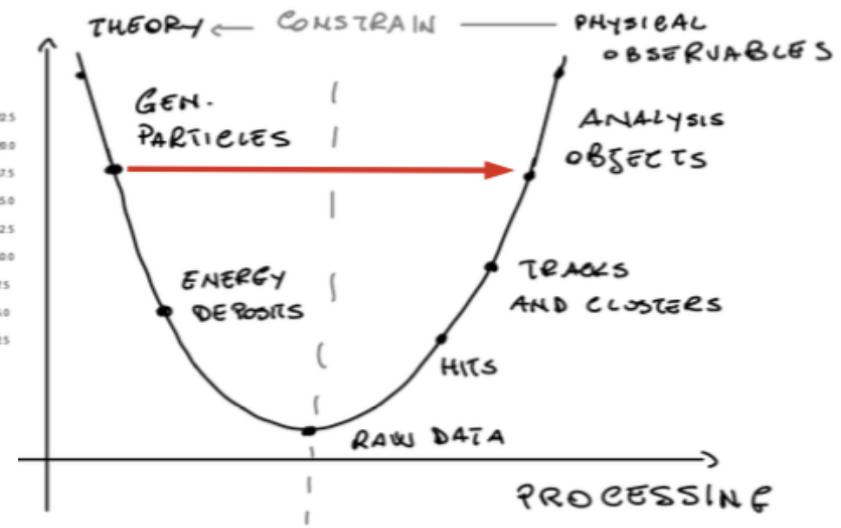
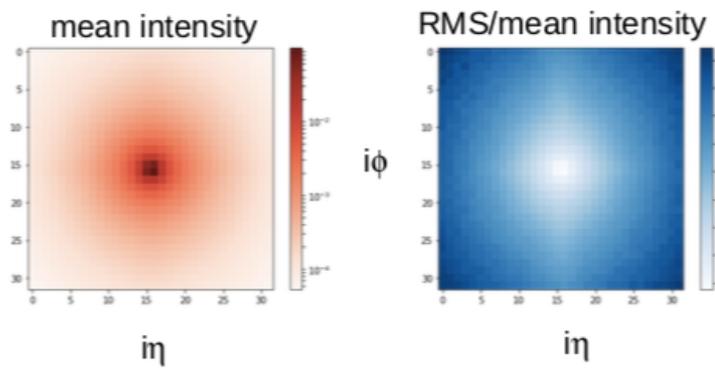


FOCUSES ON LOW LEVEL PART OF THE CHAIN:

- LA-GANS FOR EVENT SIMULATION.
- CALO-GAN AND GEANT-V FOR DETECTOR SIMULATION.

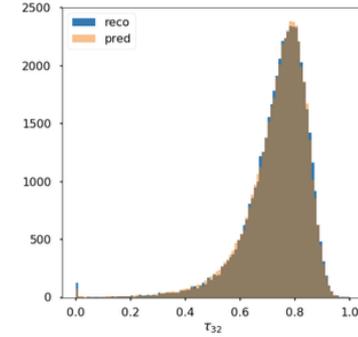
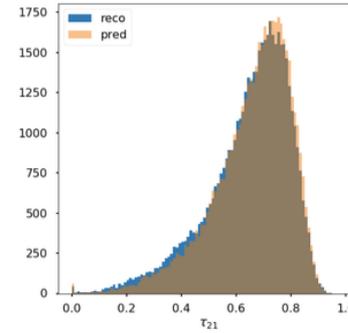
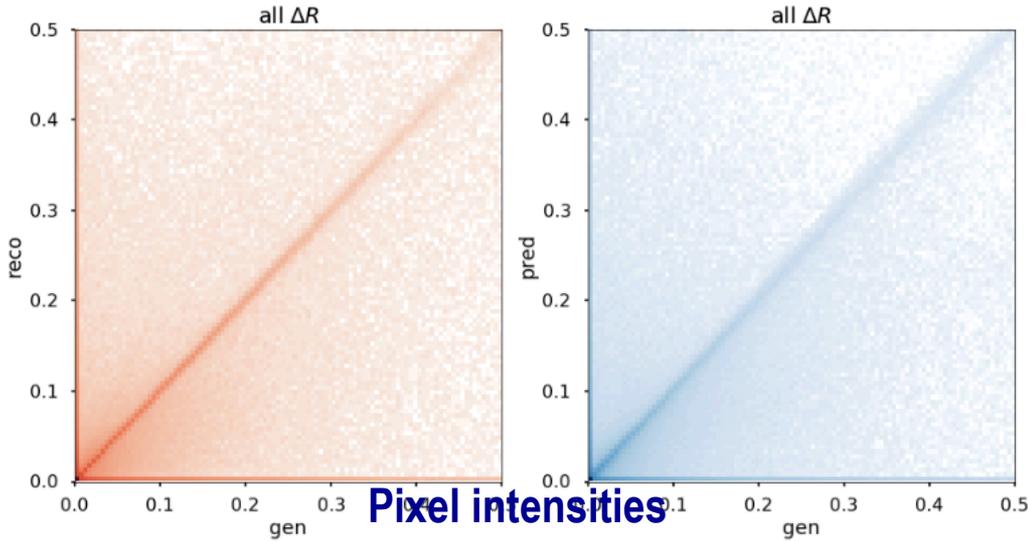


VERY HIGH ACCURACY REQUIRED:



Musella and Pandolfi, 2018

FastSim with GANs



Jet substructure

