Machine Learning in High-Energy Physics

Daniele Bonacorsi (University of Bologna / DIFA and INFN)

***AlmaHAI "Hard Sciences" KOW** - 28 Apr 2021*

Setting the scene

Scientific computing has been **<u>enabling the HEP program</u>**

- until today, software and computing have <u>NOT</u> been a limiting factor for Physics

For this to remain true, we need to face the challenges ahead of us:

- The *<u>usual</u>* challenge: computing remains a significant cost driver
  - → measure-optimise-measure cycles, Computing models undergo "adiabatic" (more or less) evolutions
- The *new* challenge(s): ramp-up in global resource needs in the next decade(s)
  - e.g. HL-LHC, theory, astro-particle - in addition, new experiments
  - → further/deeper optimisations, evaluation and adoption of new (even "disruptive") paradigms

Additional complexity from uncertainties in quantitative definition of needs, and specification of computing environments

Scale of the HEP challenges

<u>Disclaimer</u>: not a complete list, and only on experimental physics

HL-LHC hilumilhc.web.cern.ch

- 10x trigger rate, 6x event complexity, plus detector complexity: >60x resources needs (main concern is disk)

SKA (Square Kilometre Array) skatelescope.org

- aims at collecting ~300 PB/yr
  major challenge on software, computing, data movement

LSST (Large Synoptic Survey Telescope) lsst.org

- aims at collecting ~50 PB/yr
  same as for SKA

VIRGO-LIGO virgo-gw.eu, ligo.caltech.edu

- multi-messenger astronomy
  processing velocity as a challenge, more than data volume

NOTE: Not only more events, but also high-granularity detectors. *HEP will pay the price of not having computing costs folded in since the design phase*

E.g. ATLAS and CMS towards HL-LHC

Disk

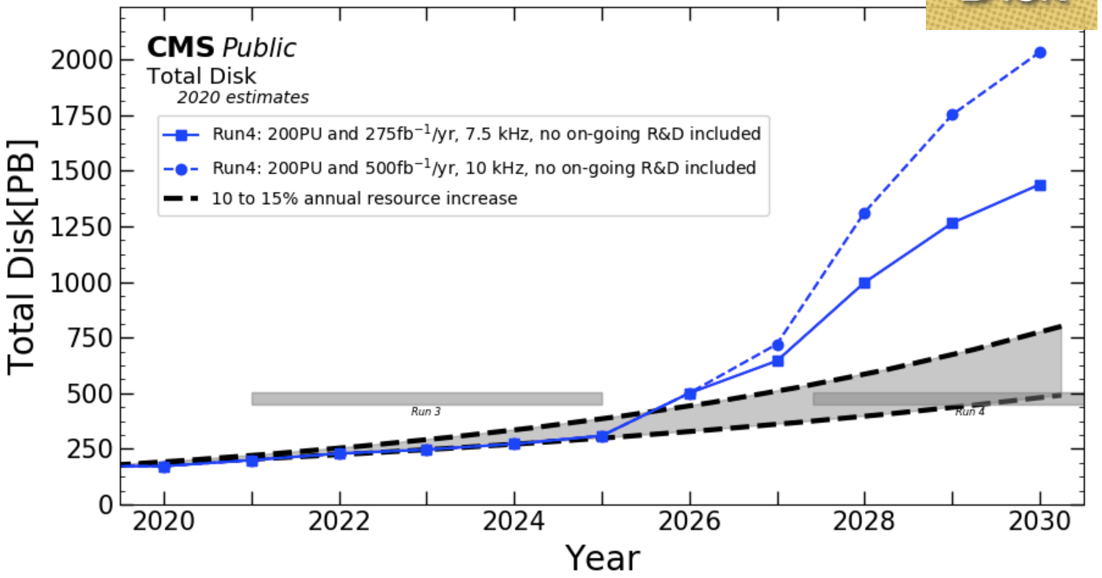

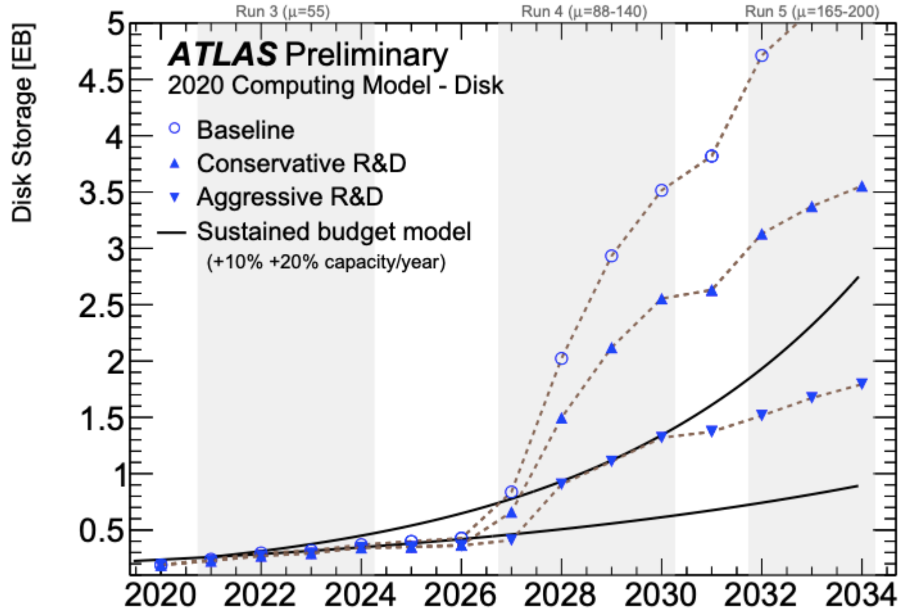

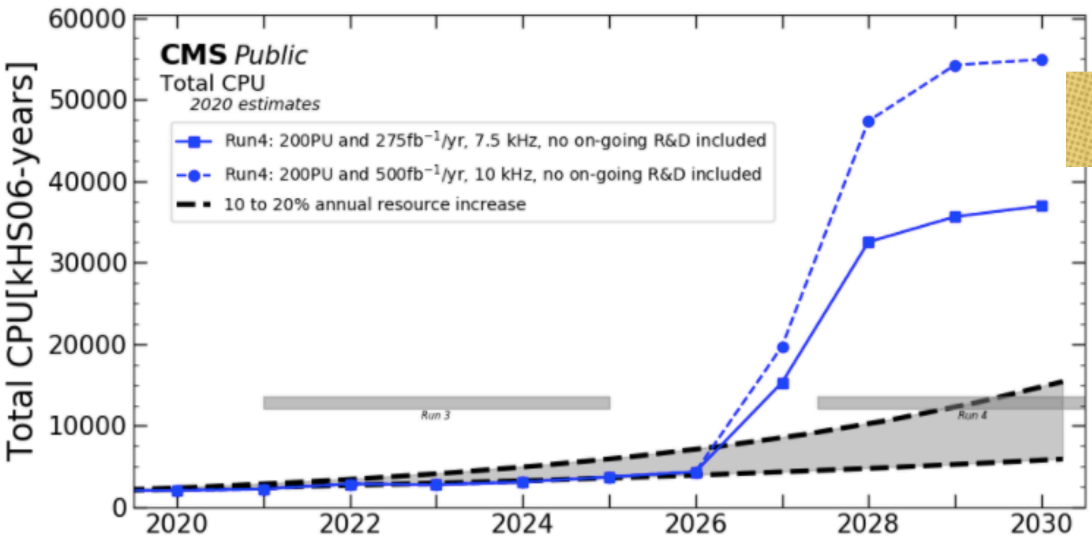

CPU

E.g. needed CPU cycles below **~4x** what is achievable with the technological improvement alone

https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ComputingandSoftwarePublicResults
https://twiki.cern.ch/twiki/bin/view/CMSPublic/CMSOfflineComputingResults

Can ML be “part of the solution”?

Most computationally-intense parts of our workflows are known

- clear target

A set of **modular and challenging goals** that fit

- e.g. (relatively) easy to map our needs to solutions that may come from discriminative or generative models

Data-driven modelling is viable

- plenty of data (also simulated data) we understand well → fuel for training

Access to **high-performances computing infrastructures**

- we designed, deployed and operated the Worldwide LHC Computing Grid
- ML Ops at scale → towards unprecedented levels (hyper-fast, low-power, ..)

Synergy with the non-HEP world is growing *from the inside of HEP*

- formulate a HEP problem in a way CS/ML practitioners can contribute to
- in addition, large traction in HEP software/computing topics from new generations

Discuss **ML**@HEP with a focus on **ML**?

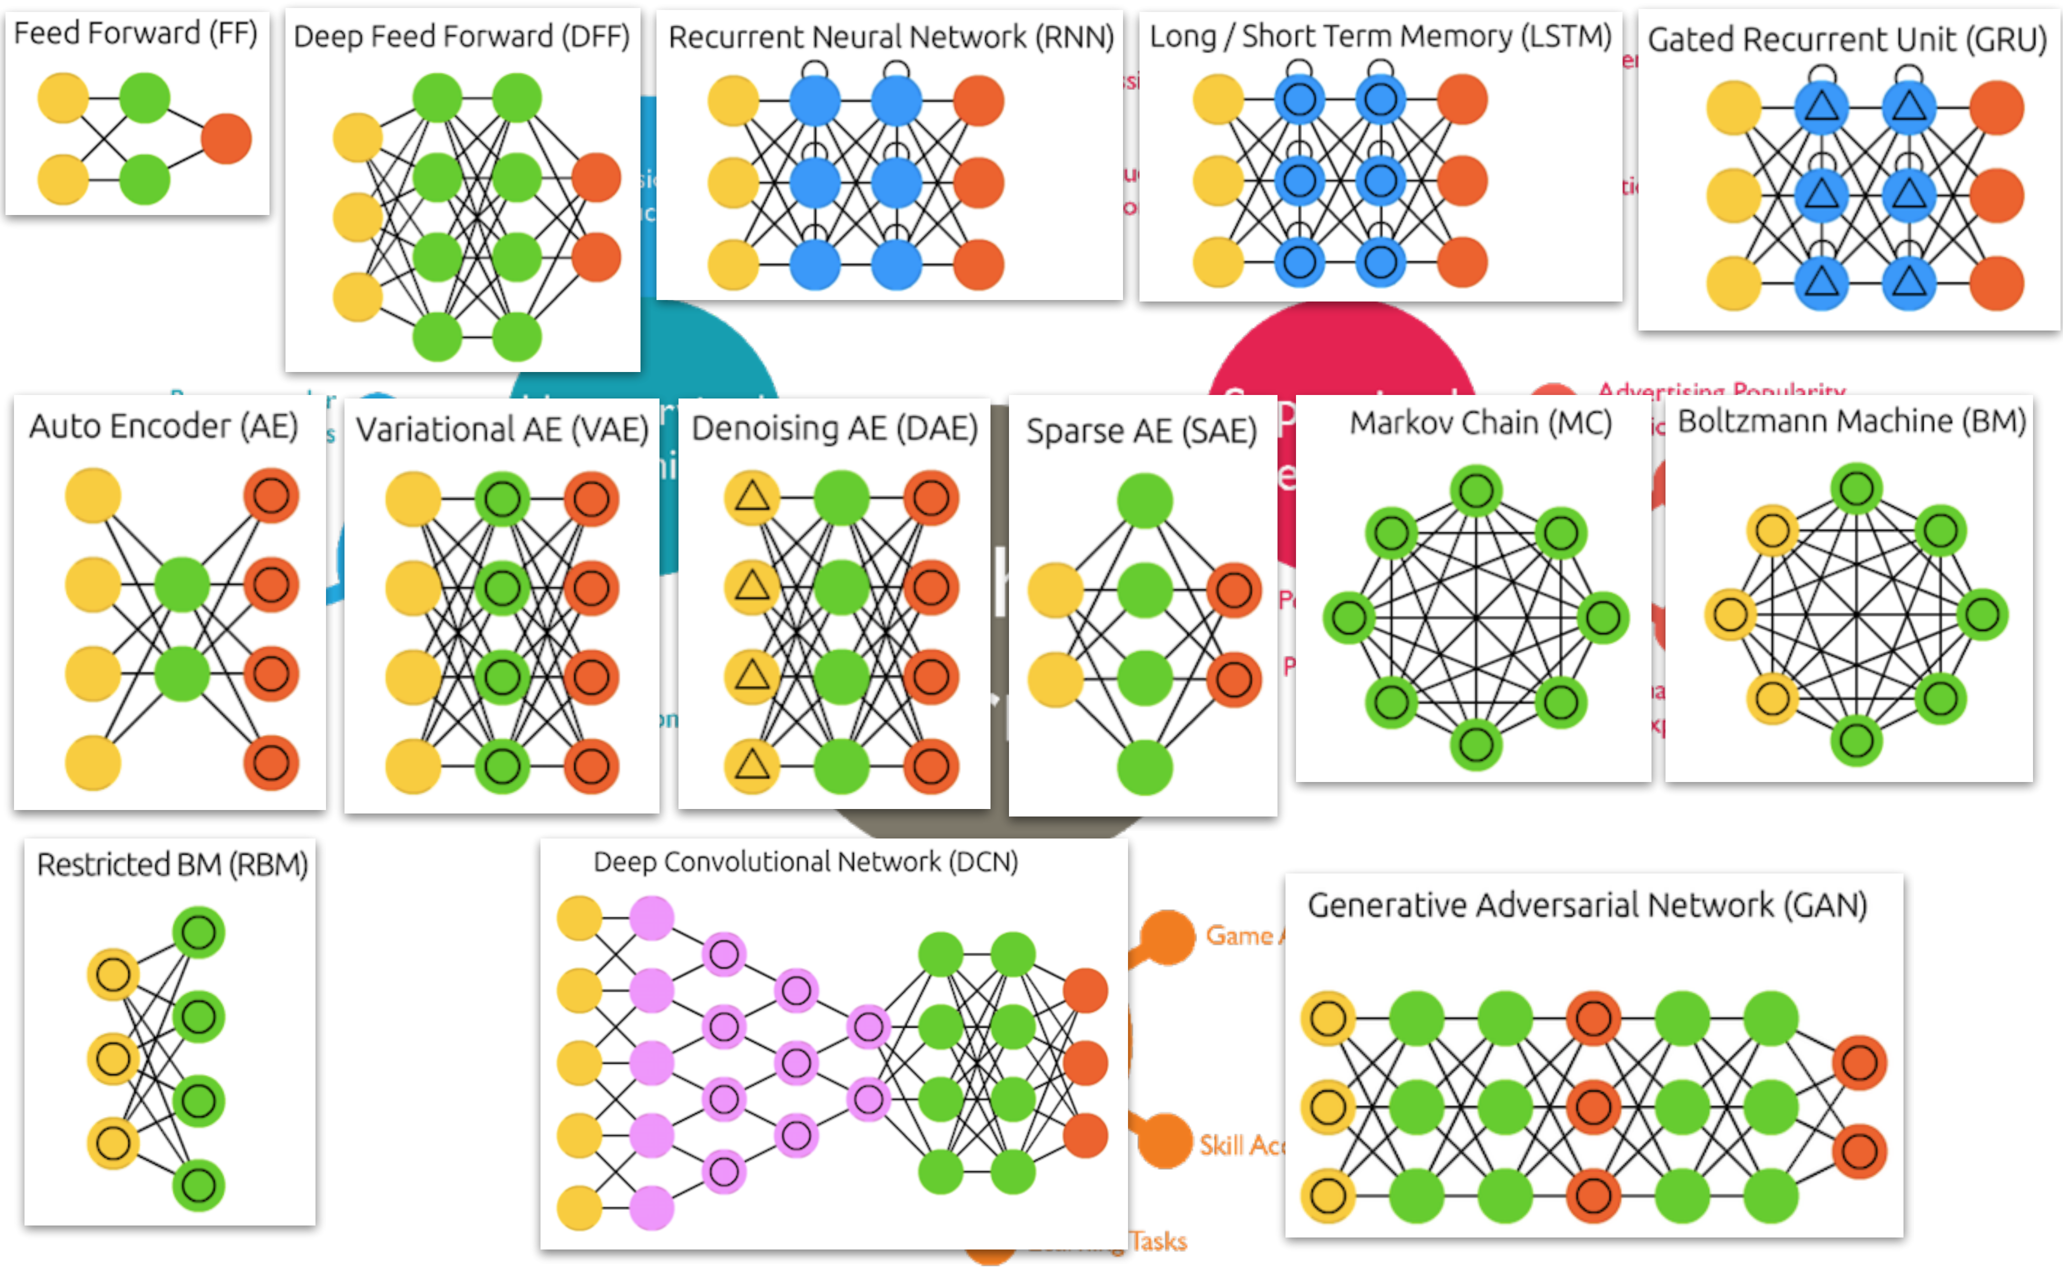

Discuss ML@HEP with a focus on HEP activities?

ML in **data acquisition** and **triggering**

- Bkg and trigger rate reduction
- Signal specific trigger paths
- Anomaly detection in data taking
- Unsupervised new physics mining

Trigger for **all HEP experiments** is a driver of ***high-performance* ML applications in HEP**

- design of next-generation triggering processes as key enabler of real-time reconstruction (that in turn enables real-time analysis)

One of the challenges is the **trade-off** in **algorithm complexity** and performance under **strict inference time constraints**

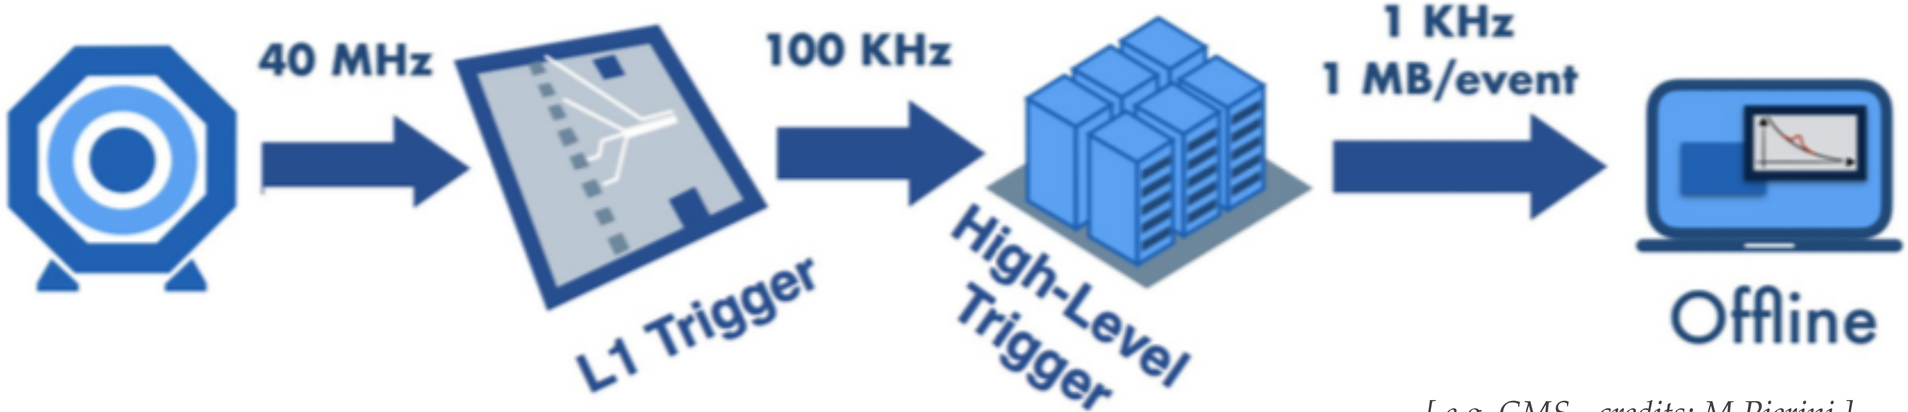

[e.g. CMS - credits: M.Pierini]

Discuss ML@**HEP** with a focus on **HEP** activities?

ML in Event **Simulation**

Event production through full/fast simulation today is extremely computing-intensive (up to potentially impacting the Physics reach of experiments). ML may help reducing such load

- Calorimeter shower surrogate simulator
- Analysis level simulator
- Pile-up overlay generator
- Monte Carlo integration
- ML-enabled fast-simulation
- Invertible full-simulation (probabilistic programming, …)
- …

ML in Event **Reconstruction**

Online/offline reconstruction may be in part replaced by surrogate models (approximate → faster) or by novel algorithms (yielding unprecedented performance)

- Charged particle tracking (GraphNN, vertexing, …)
- Calorimeter reconstruction (local, clustering, …)
- Particle flow (GraphNN, …)
- Particle identification (boosted jets, isolation, …)
- Pileup mitigation
- Energy regression (end-2-end, …)
- …

Discuss ML@**HEP** with a focus on **HEP** activities?

ML for **Operations**

Application of ML on non-collision (meta-)data may help to increase efficiency and reduce manpower burden in Ops, by automating selected tasks, creating intelligent/adaptive systems, ultimately expedite the entire chain from data collection to final analysis

- Detector control
- Data quality monitoring
- Operation intelligence
- Predictive maintenance
- …

Methodology in ML

"Ask not what ML can do for HEP, ask what physicists/scientists can do for ML" (semi-cit)

Plenty of work that Hard Sciences can do to contribute to build a solid theory behind ML empirical success

- Continual learning
- Interpretability and explainability
- Uncertainty quantification
- Incorporating domain knowledge
- Cast into optimisation problems
- …

Discuss ML@HEP with a focus on **HEP** activities?

REVIEW

nature
https://doi.org/10.1038/s41586-018-0361-2

Machine learning at the energy and intensity frontiers of particle physics

Alexander Radovic[1]*, Mike Williams[2]*, David Rousseau[3], Michael Kagan[4], Daniele Bonacorsi[5,6], Alexander Himmel[7], Adam Aurisano[8], Kazuhiro Terao[4] & Taritree Wongjirad[9]

Our knowledge of the fundamental particles of nature and their interactions is summarized by the standard model of particle physics. Advancing our understanding in this field has required experiments that operate at ever higher energies and intensities, which produce extremely large and information-rich data samples. The use of machine-learning techniques is revolutionizing how we interpret these data samples, greatly increasing the discovery potential of present and future experiments. Here we summarize the challenges and opportunities that come with the use of machine learning at the frontiers of particle physics.

The standard model of particle physics is supported by an abundance of experimental evidence, yet we know that it cannot be a complete theory of nature because, for example, it cannot incorporate gravity or explain dark matter. Furthermore, many properties of known particles, including neutrinos and the Higgs boson, have not yet been determined experimentally, and the way in which the emergent properties of complex systems of fundamental particles arise from the

Big data at the LHC

The sensor arrays of the LHC experiments produce data at a rate of about one petabyte per second. Even after drastic data reduction by the custom-built electronics used to readout the sensor arrays, which involves zero suppression of the sparse data streams and the use of various custom compression algorithms, the data rates are still too large to store the data indefinitely—as much as 50 terabytes per second,

nature.com/articles/s41586-018-0361-2

The evolution of ML adoption in HEP

“Traditional” ML

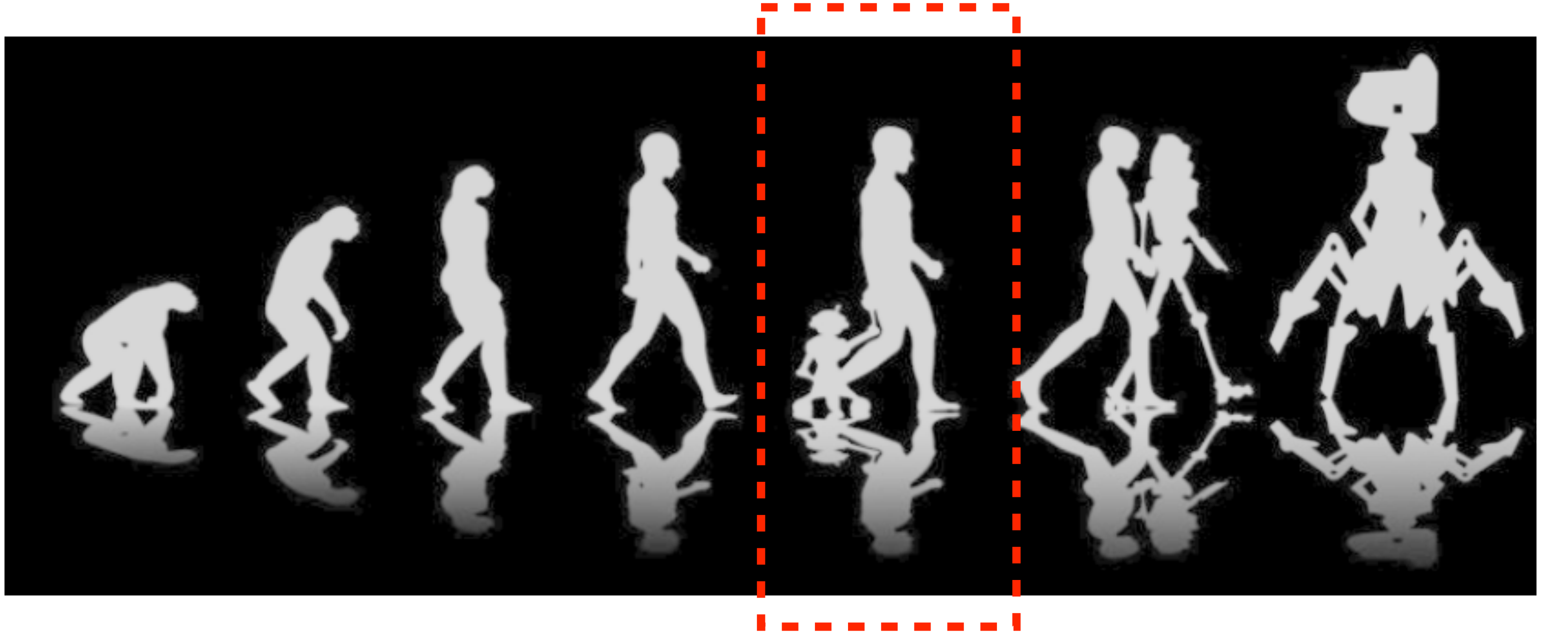

Until few years ago, the overall ML@HEP scenario was based on exploiting field-specific knowledge for **feature engineering**

The approach:

- **use physicist-designed high-level features as input to traditional shallow ML algorithms**

Particle properties: energy resolution

Using ML to improve the determination of particle properties is now commonplace in **all LHC experiments**

E.g. energy deposited in calorimeters is recorded by many sensors, which are clustered to reconstruct the original particle energy

- **CMS** is training **BDT**s to learn corrections using all information available in the various calorimeter sensors - thus resulting in a **sizeable improvement in resolution**

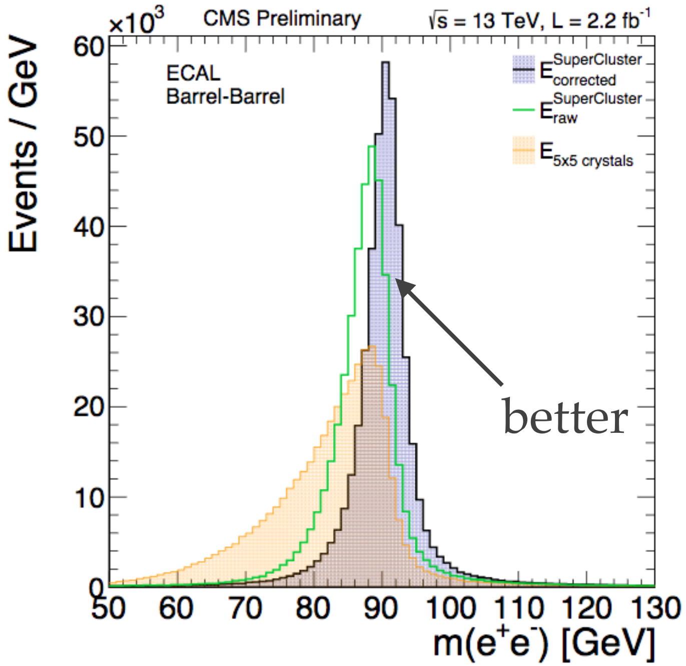

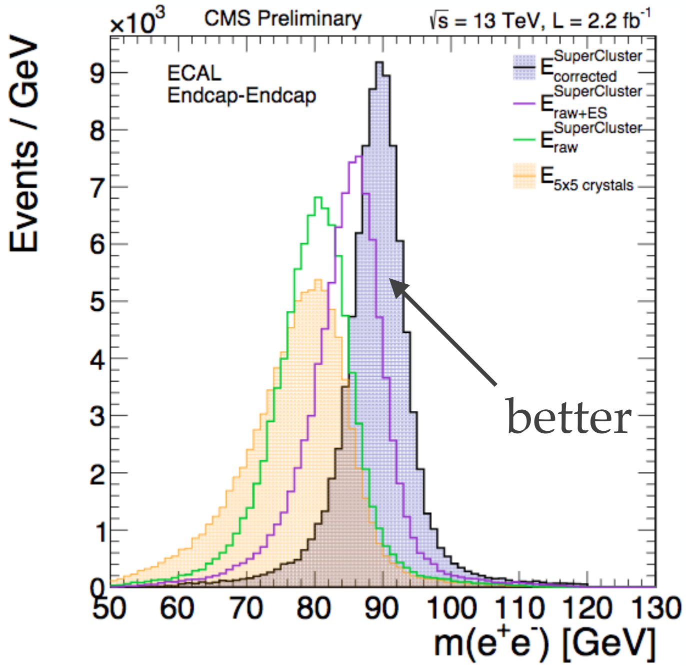

Improvements to the Z→e+e- energy scale and resolution from the incorporation of more sophisticated clustering and cluster correction algorithms (energy sum over the seed 5x5 crystal matrix, bremsstrahlung recovery using supercluster, inclusion of pre-shower energy, **energy correction using a multivariate algorithm**)

[2015 ECAL detector performance plots, CMS-DP-2015-057. Copyright CERN, reused with permission]

Particle identification

Similarly, ML is commonly used to identify particle types

- e.g. **LHCb** uses **NN**s trained on O(30) features from all its subsystems, each of which is trained to identify a specific particle type

- **~3x less mis-ID bkg /particle**. Further estimates indicate that **more advanced algorithms may reduce bkg by another ~50%**

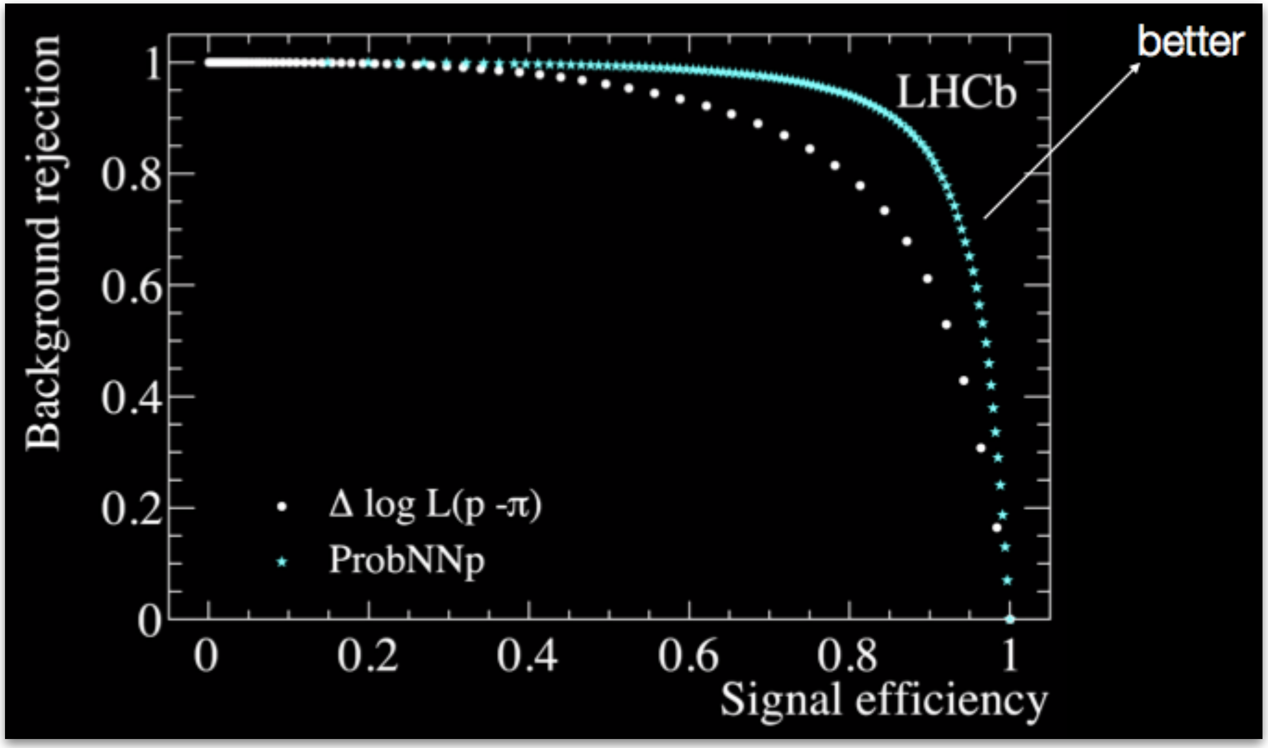

[courtesy: M.Williams]

Discovery of the Higgs boson

ML played a key role in the discovery of the Higgs boson

- especially in the diphoton analysis by **CMS**, where **BDTs** (used to improve the resolution and to select/categorise events) **increased the sensitivity by roughly the equivalent of collecting ~50% more data**

[courtesy M.Pierini]

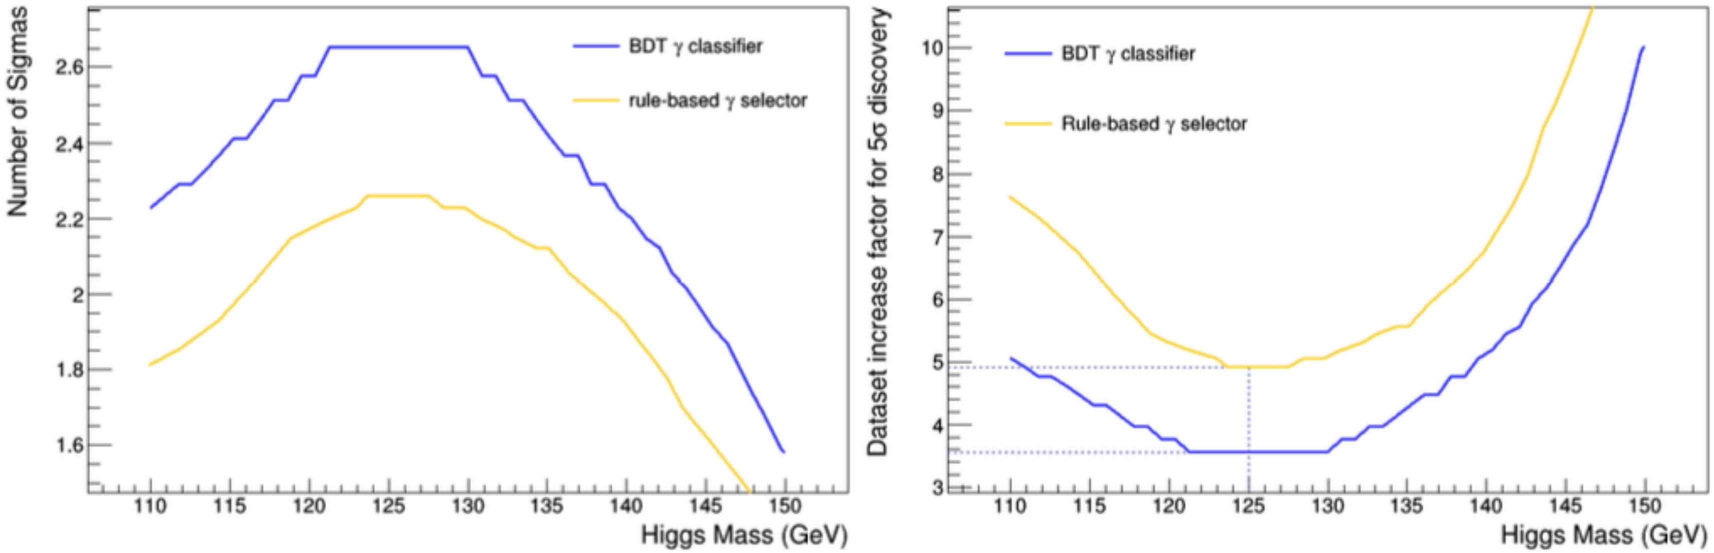

We were not supposed to discover the Higgs boson **as early as 2012**

- Given how machine progressed, we expected discovery by end 2015 / mid 2016

We made it earlier thanks (also) to ML

Study of Higgs properties

[1] JHEP 04 (2015) 117

E.g. analysis of τ leptons at LHC complex, as they decay before detection + loss of subsequently produced neutrinos + bkg from Z decays

- e.g. **ATLAS** divided the data sample into 6 distinct kinematic regions, and in each a **BDT** was trained using 12 weakly discriminating features [1] → **improved sensitivity by ~40% vs a non-ML approach**

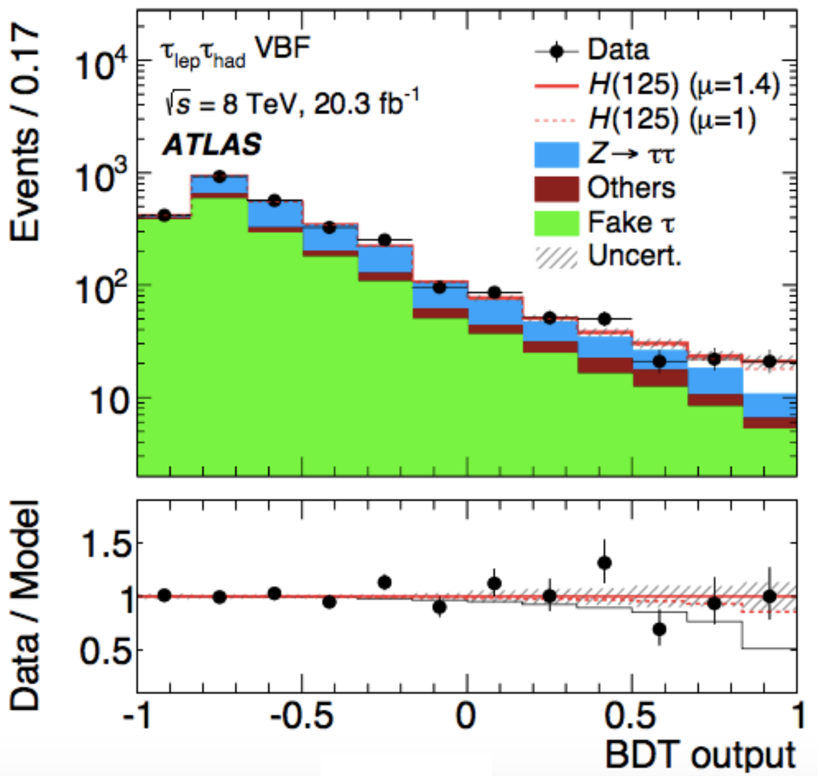

High-precision tests of the SM

[1] Nature 522 68–72 (2015)

[2] *Phys.Rev.Lett.* 118 (2017) 19, 191801

CMS and **LHCb** were the first to find evidence for the $B^0_s \rightarrow \mu^+\mu^-$ decay with a combined analysis [1] (as rare as ~ 1 / 300 billion pp collisions..)

- **BDT**s used to reduce the dimensionality of the feature space - excluding the mass - to 1 dimension, then an analysis was performed of the mass spectra across bins of BDT response

- decay rate observed is consistent with SM prediction with a precision of ~25%, placing stringent constraints on many proposed extensions to the SM

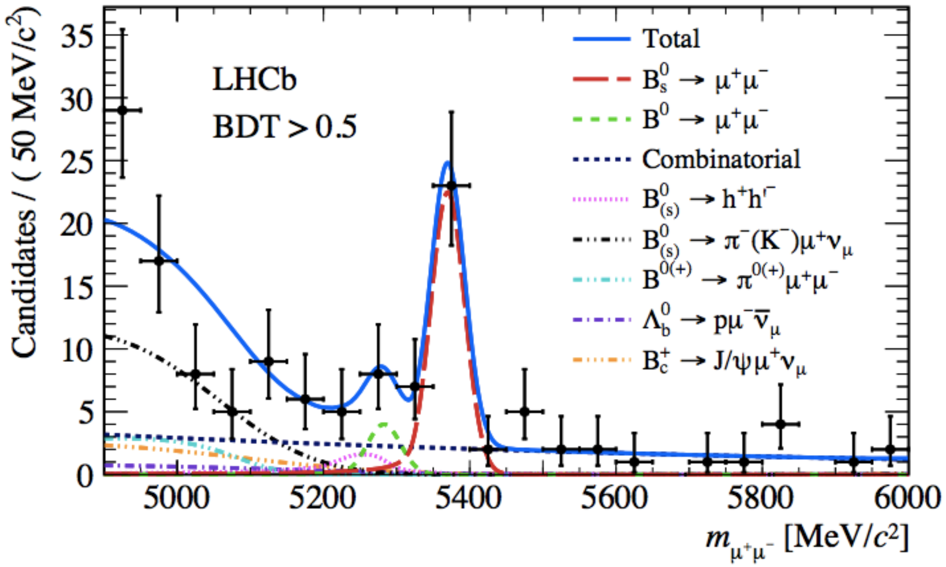

Mass distribution of the selected $B^0 \rightarrow \mu^+\mu^-$ candidates with BDT > 0.5 [2]

To obtain the same sensitivity without ML by LHCb as a single experiment would have required ~4x more data

The evolution of ML adoption in HEP

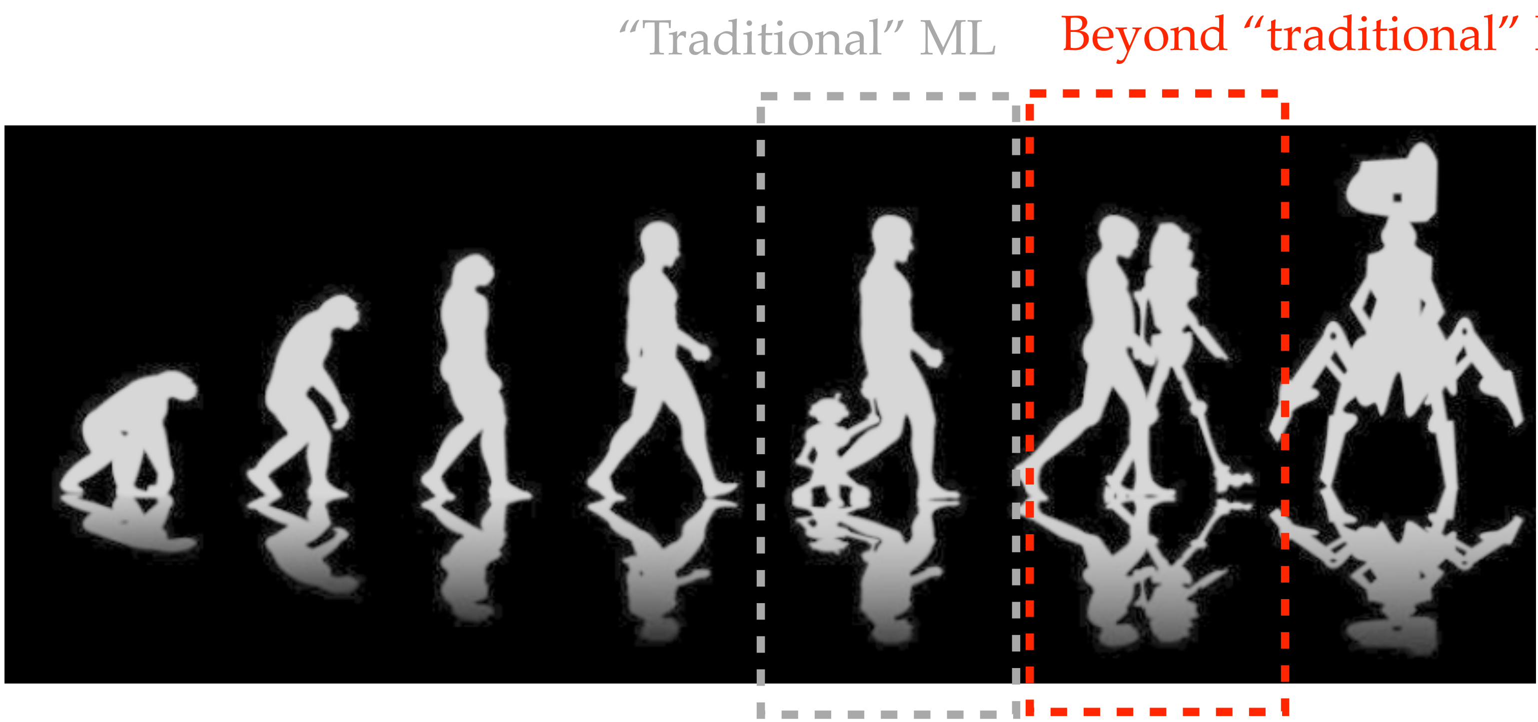

Since few years, ML@HEP exploits cutting-edge ML algorithms

- **multiple architectures of Neural Networks** (NNs), depending on specific use-cases

The approach:

- use of full high-dimensional feature space to train **Deep NNs**; growing effort in HEP to <u>skip the feature-engineering step</u>
    - in analogy with progresses in Computer Vision and Natural Language Processing

Deep Learning for jet flavour identification

[1] JINST 13 (2018) P05011
[2] JINST 15 (2020) 12, P12012

Why Deep NN?

- Large volume of simulated events for training (fight overfitting, better generalisation)
- High-level representations are tough → DeepNNs excel when fed with large volumes of low-level features

E.g. CMS reviewed algos for the identification of heavy-flavour (b and c quarks) jets in pp at 13 TeV

- **DeepCSV algorithm** [1]: NN-version of the combined secondary vertex (CSV) taggers used in LHC Run-1
→ **15% improvement** in relative efficiency w.r.t likelihood-based methods

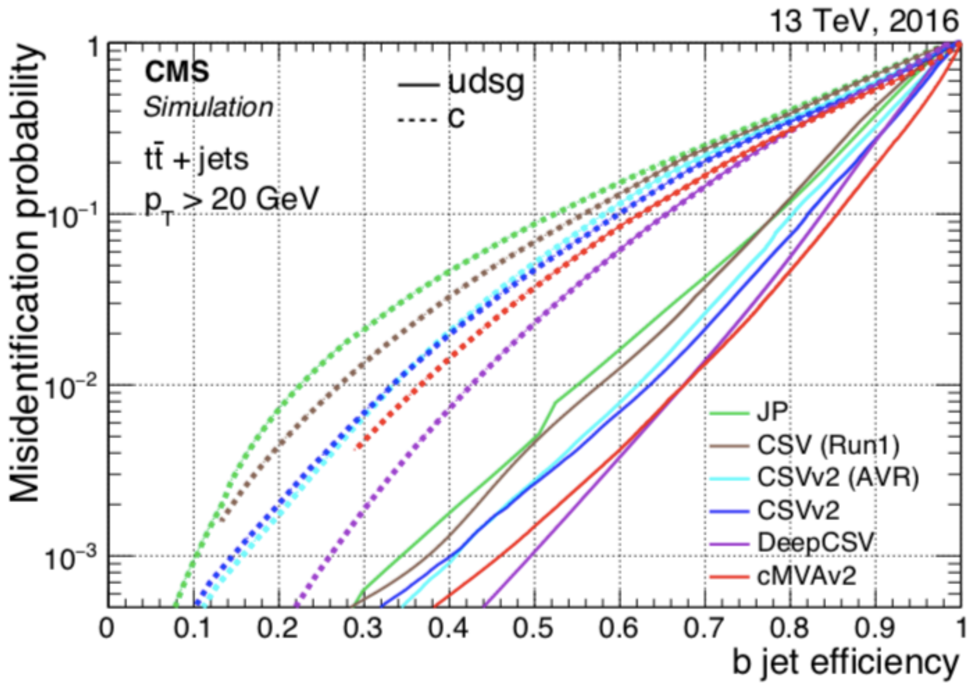

Further work → **DeepJet** [2]

- use full lists of particle flow candidates, secondary vertices, ..
- **performance improvements**, plus extension to quark-gluon tagging

Convolutional Neural Networks (**CNN**)

CNNs are based on strategies that decrease their sensitivity to the absolute position of elements in an image, making them more robust to noise

- Deep CNNs capable to extract complex features from images
  - e.g. use in self-driving cars, owing to translation-invariant feature learning
- particularly suited for HEP neutrino experiments
  - but also in simplified settings in collider experiments

Industry:
large adoption in computer vision tasks

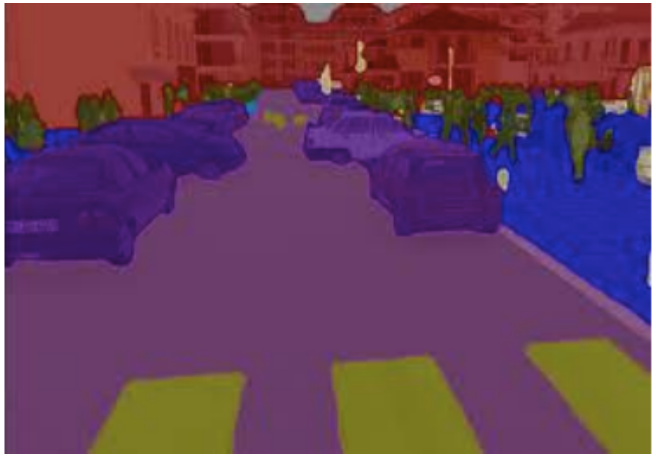

HEP:
3D imaging in detectors, event classification, ..

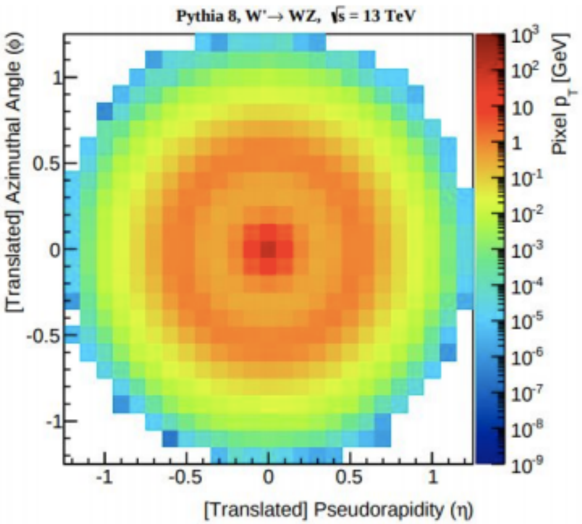

CNNs in NOvA

[1] J. Instrum. 11, P09001 (2016)
[2] Phys. Rev. Lett. 118, 231801 (2017)
[3] J. Instrum. 12, P02017 (2017)

[novaexperiment.fnal.gov]

NuMI Off-axis ν_e Appearance (NOvA): study neutrinos through precision measurements of their oscillation properties

NOvA detector and data collection strategy

- filled with mineral oil, which emits light when a charged particle transverse it
- a NOvA event consists of 2 images, taken from the top and from the side

Novel ML algo [1] **for NOvA: 2 parallel NNs inspired by GoogleNet**

- the NOvA CNN extracts features from both views simultaneously and combines them to categorise neutrino interactions in the detector

→ **improvement of 40% (with no loss in purity) in the efficiency** of selecting ν_e

- used as the event classifier in searches for appearance of ν_e [2] but also for a new type of particle called a sterile neutrino [3]

CNNs in NOvA

Softmax Output
Avg Pooling 6x5
Inception Module
Inception Module
Max Pooling 3x3, stride 2
Inception Module
Inception Module
Max Pooling 3x3, stride 2
LRN
Convolution 3x3
Convolution 1x1
LRN
Max Pooling 3x3, stride 2
Convolution 7x7, stride 2
X View
Y View

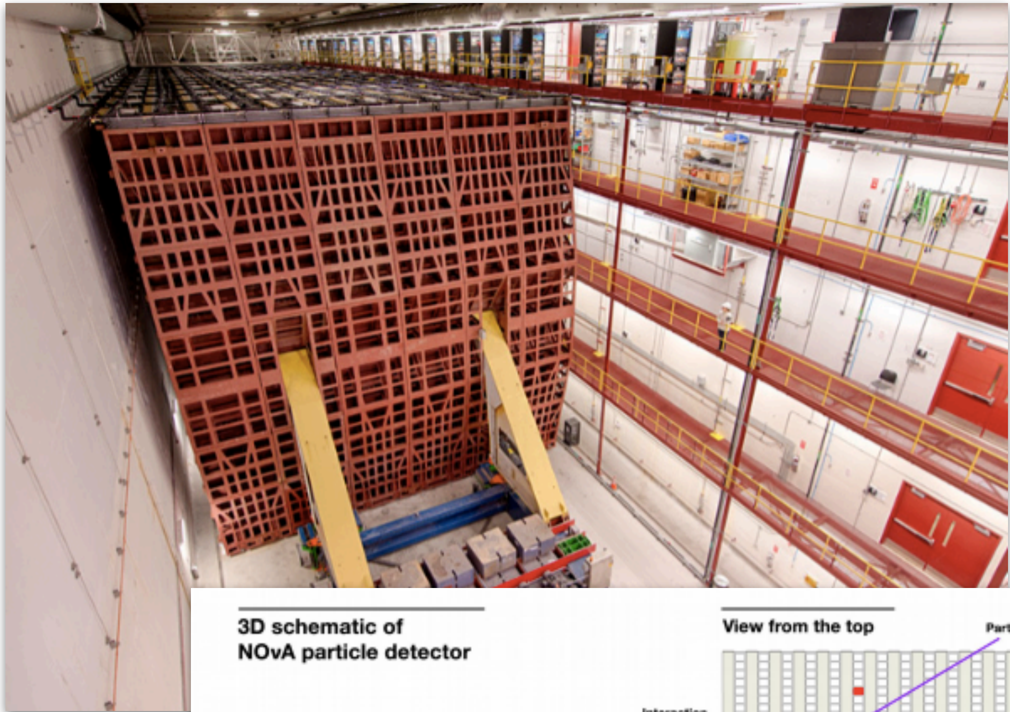

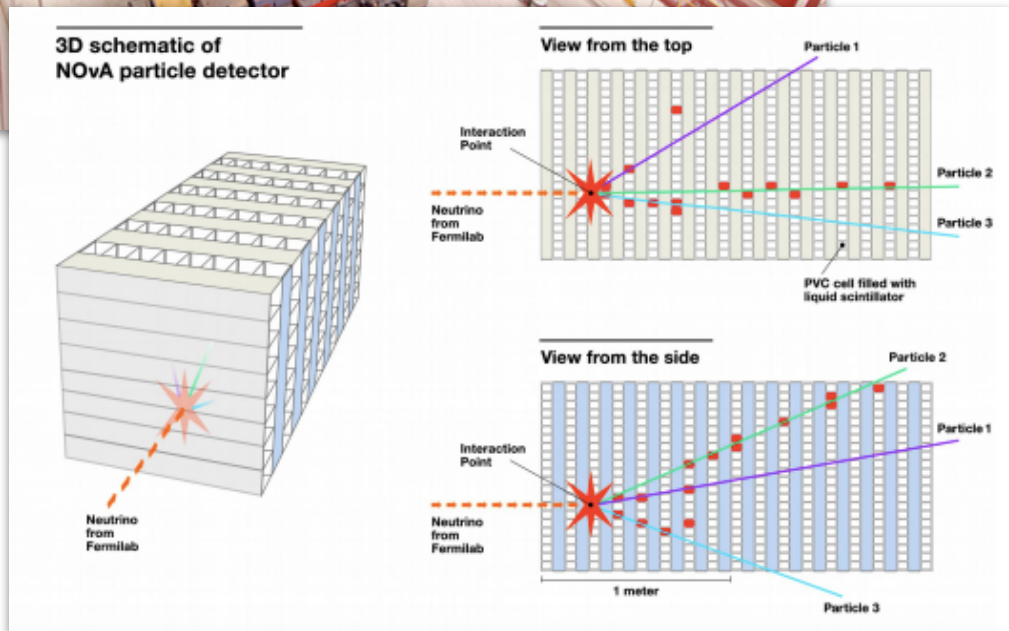

3D schematic of NOvA particle detector
View from the top
Particle 1
Interaction Point
Particle 2
Neutrino from Fermilab
Particle 3
PVC cell filled with liquid scintillator
View from the side
1 meter

CNNs in MicroBooNE

[microboone.fnal.gov]

[1] J. Instrum. 12, P03011 (2017)
[2] IEEE Trans. Pattern Anal. Mach. Intell. **29**, 061137 (2017)

MicroBooNE uses a large LArTPC to measure a suite of low energy ν cross sections, and investigate astro-particle physics

MicroBooNE detector and data collection strategy

- 170-ton of liquid argon to detect ν sent from the booster ν beam-line at FNAL
- a MircoBooNE event corresponds to a 33 Mpixel image that may contain signal of ν interaction and bkg tracks by cosmic rays (size vary from cm's to m's)

Faster-RCNN: a novel CNN method to detect ν interactions [1][2]

- it uses spatially sensitive information from intermediate Conv layers to predict a bounding box that contains the secondary particles produced in a ν interaction

→ **GPU-accelerated CNNs show better performance than any conventional algos** used by previous ν experiments

- ideally suited to any task of real-time image classification and object detection

CNNs in MicroBooNE

E.g. neutrino selection and isolation in MicroBooNE from the output of the CNN

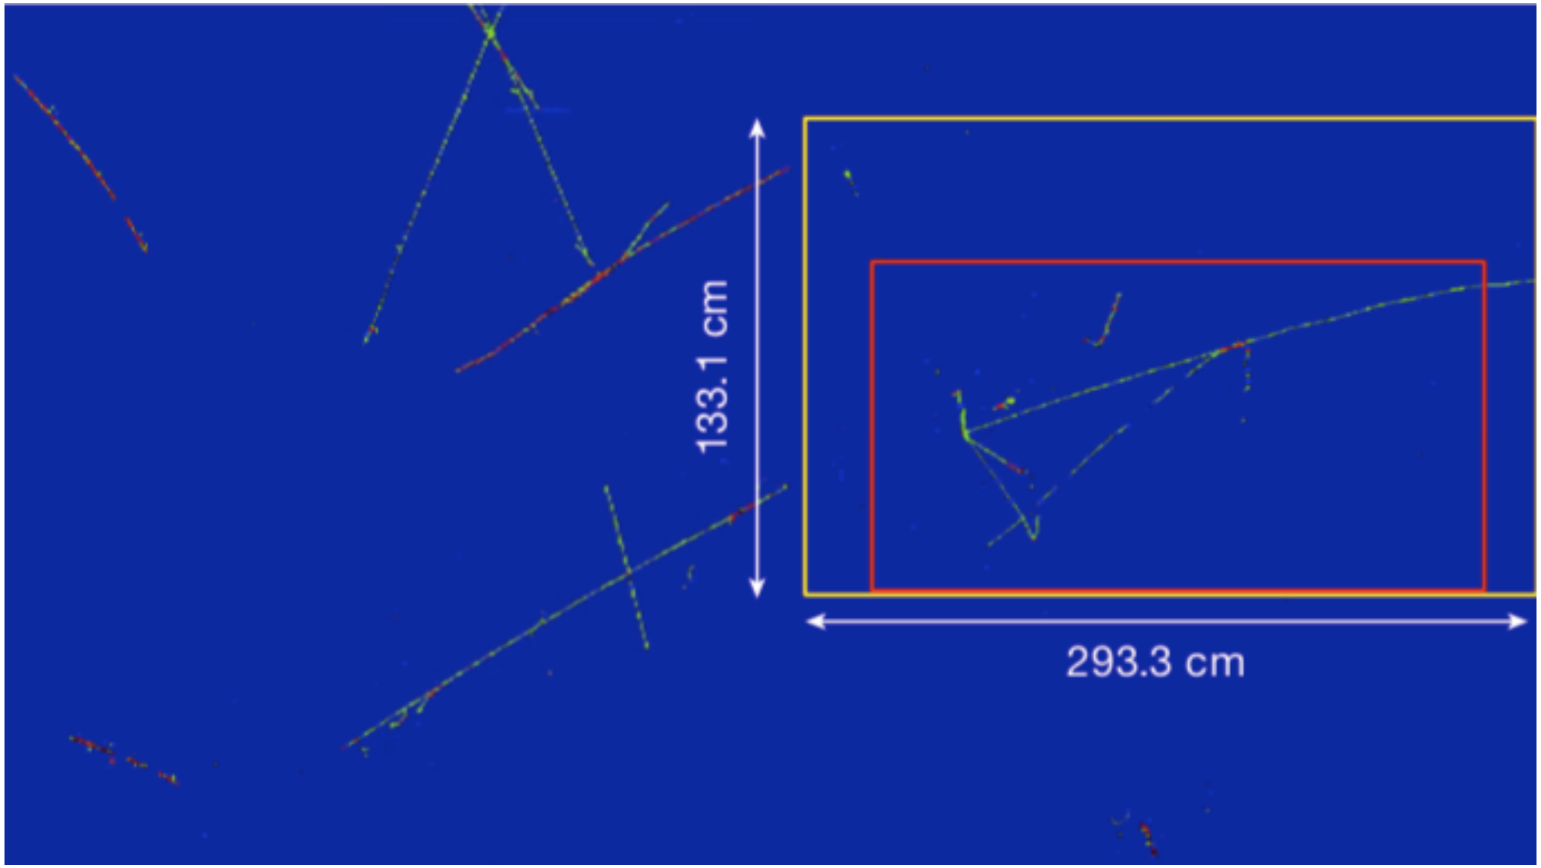

The CNN does not get in input the info as from the **yellow box** …

.. but it is able to draw the **red box**, i.e. it successfully localises the neutrino interaction with high confidence.

".. but HEP is different .."

Bonus feature!

Is it, really, to all extents?

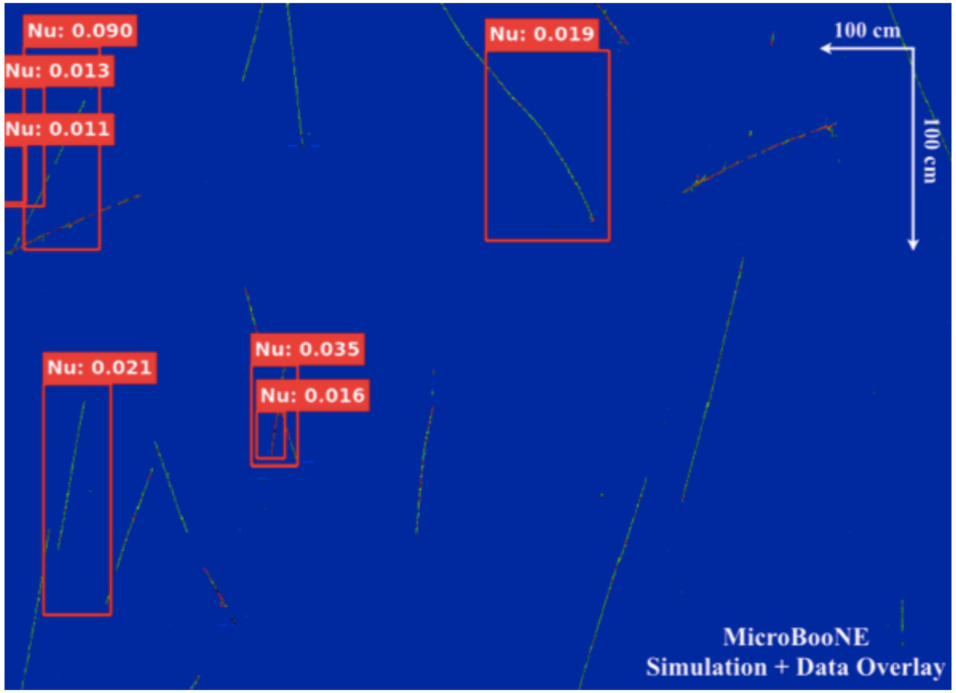

Detection of **neutrinos** on cosmic background event (method: **CNN**)

Detection of **airports** from satellite images (method: **CNN**)

DeepNNs for particle ID and particle properties

A general tactics (TPCs, CALOs..): represent the data as a 2D or 3D image → the problem can be cast as a computer vision task

- even 4D, including timing information..

Deep Learning techniques based on **DeepNNs** to reconstruct images from pixel intensities are good candidates to be used to identify particles and extract parameters **in many experiments**

- promising DL architectures not limited to **CNN**, but also **Recurrent NN (RNN)**

Examples:

- whenever e.g. LArTPCs is the chosen detection technology, see previous examples, see also **DUNE**, ..
- More: **b-tagging in collider experiments** can also exploit RNNs

Recurrent Neural Networks (**RNN**)

RNNs successful at processing long sequences of data

- based on recurrent neurons (with connections pointing backwards)
- able to treat variable-length input and to process time series by accumulating and using all the info across a sequence
  - e.g. current Google translation service

Industry:
managing "time series"
(audio, video, natural language processing)

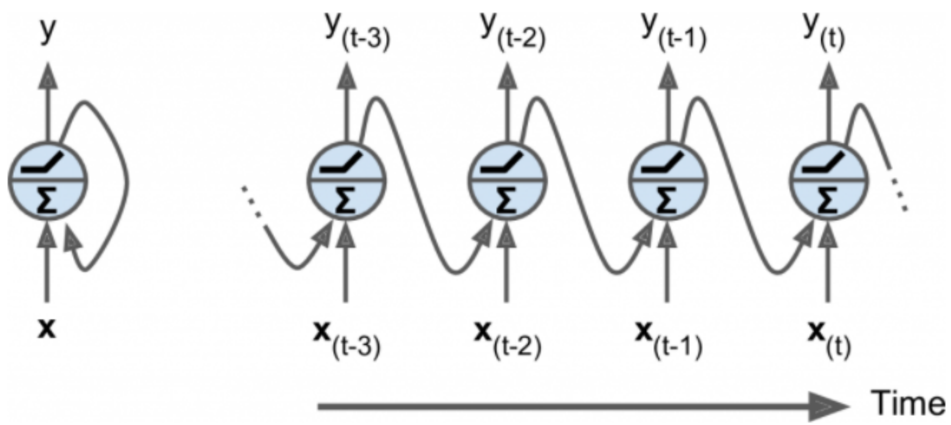

HEP:
classifiers capable to process complex signals, or variable-lenght inputs (tracks, particles in jets, etc)

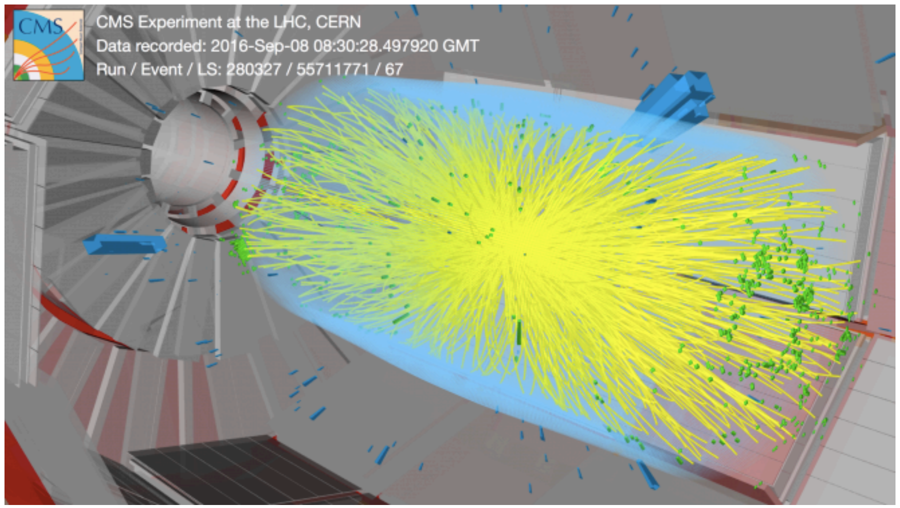

RNNs on b-quark identification

[1] ATL-PHYS- PUB-2017-003 and 013
[2] CMS-DP-2017-005
[3] Phys. Rev. D 94, 112002 (2016)

Beauty (b) quarks identification is important at LHC

- frequently produced in Higgs/top decays, predicted in SUSY decays, and more
- identification of jets from HE b-quarks (time scale is picosecs, flight range is <cm) requires finding a displaced vertex, typically contain 10-50 particles, # of potential discriminating features varies on a per-jet basis

RNNs methods can use low-level particle features within a jet

- order the jet particles into a sequence (e.g. ranking by incompatibility with originating from pp collision point), feed a set of features for each particle to the RNN, and train to discriminate between b-quark jets and other jets

→ **ATLAS**: **mis-identification rate reduced by 4x** w.r.t non-ML algo, with an **additional 3x reduction** when RNN itself used as an input feature in the subsequent training of a BDT or NN [1]

→ **CMS**: **similar approaches** [2], + promising results with more sophisticated RNN structures in a simplified setting [3]

Autoencoders (**AE**) and Variational-AE (**VAE**)

AEs are feed-forward NN (unsupervised) able to compress input into a lower-dimensional representation ("latent-space", a data-specific compression) and to reconstruct the output

VAEs are generative models instead → learn the parameters of a probability distribution representing the data → can generate new input data samples

AEs in Industry:
dimensionality reduction, denoising, …

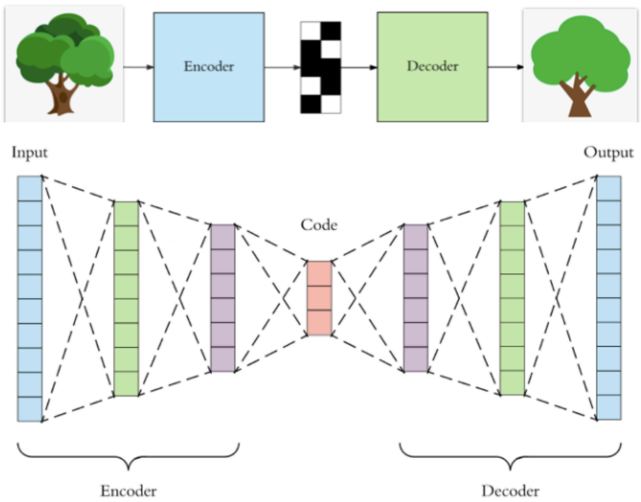

VAEs in HEP:
VAEs could isolate new physics as outliers of known distributions

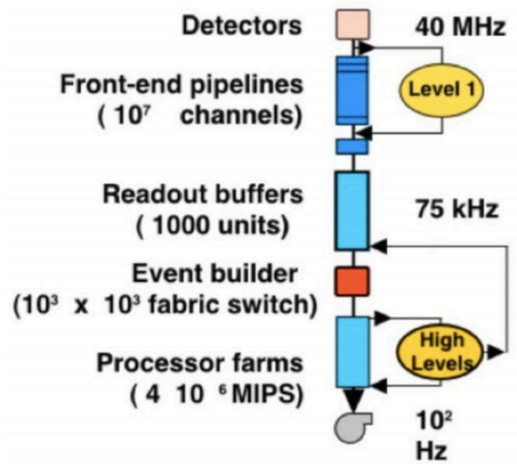

VAEs for new physics mining at LHC

[1] JHEP05 (2019) 036

For all **LHC experiments**, VAEs are proposed to be trained on known physics processes to be used to build "thresholds" to isolate previously unseen physics events as outliers

- training does not depend on any specific new physics signature → assumptions-free, and complementary to classic LHC searches for new physics (typically based on model-dependent hypothesis testing)

Outcome: a catalogue of anomalous events to be further scrutinised

- recurrent event topologies in the catalogue may inspire focussed searches and model building

- plugging this technique in the trigger of LHC exps could avoid discard of potentially interesting events

A promising approach to **extend the physics reach of LHC**

Generative Adversarial Networks (**GAN**)

GANs as generative ML models

- designed as 2-NN game where one (generator NN) maps noise to images, and the other (discriminator NN) classifies the images as real vs fake (the best generator being the one that maximally confuses its adversary)

Industry:
image editing, data generation, security, ..

HEP:
Simulate the detector response (promising alternative to traditional simulation solutions)

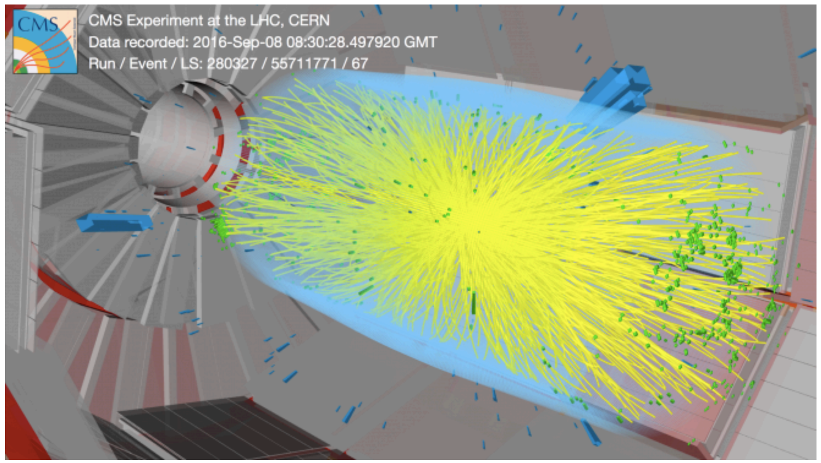

Simulation with VAEs and GANs

Promising alternatives for Fast Simulation may be built on recent progress in high fidelity fast generative models

- e.g. **GAN**s and **VAE**s → ability to sample high dimensional feature distributions by learning from existing data samples

Some simplified first attempts at using such techniques for simulation saw **orders of magnitude increase in speed over existing Fast Simulation techniques**, which **all HEP experiments** would largely benefit from

- not yet reached the required accuracy, though (inherent shortcomings of the methods, instabilities in training such NNs, ..)

E.g. Simulation with GANs

[1] *Phys.Rev.D* 97 (2018) 1, 014021

E.g. **CaloGAN**, a new FastSimulation technique, to simulate 3D HEP showers in multi-layer ECAL systems with GANs

- basically, CaloGAN can generate the reconstructed Calo image using random noise, skipping the GEANT and RECO steps - thus making it **10k faster than GEANT**

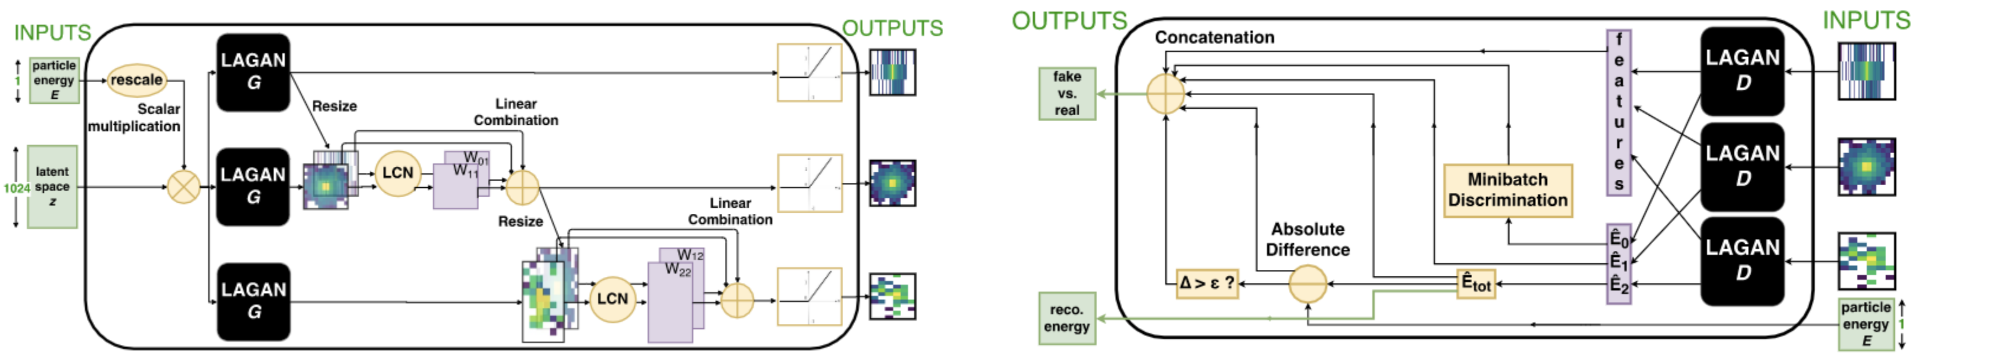

CaloGAN composite generator (left) and discriminator (right)

Towards **Quantum ML** for HEP

"How QC can be used for ML (in HEP)?"

- namely: *"how quantum computers can learn from (HEP) data?"*

Possible approaches to this question:

- foundational approach that reformulates learning theory in a quantum setting
- efforts to find quantum algos that speed up ML with regards to computational complexity measures
- a near-term perspective: develop new ML applications tailored for NISQ devices

The near-term perspective: start from quantum devices available today and investigate how they can be used to solve a ML problem

- **circuit-based** quantum computers → outsource the **prediction** part of ML to QC (i.e. predominantly used to compute the prediction of a QML model that can be trained classically)
- **quantum annealers** → outsource the **training** part of ML to QC (namely, proposed to optimise classical models, i.e. map an optimization problem to a QUBO instance)

Exploring a Quantum-GNN for tracking [1/2]

Charged particle tracking (aka “**tracking**”) is the cornerstone of event reco in particle physics

- in a nutshell, the task of associating sparse detector measurements (aka ”hits”) to the trajectory of a given particle that caused them

Reconstructing particle trajectories with high accuracy will be one of the major challenges e.g. in the HL-LHC experiments

- increase in the expected # of simultaneous collisions + high detector occupancy → tracking extremely demanding in terms of computing resources

Today’s state-of-the-art algos rely on a Kalman filter-based approach

- robust and provide good physics performance, but they are expected to scale worse than quadratically with the increasing # of simultaneous collisions

What’s next? Investigating several possibilities

- deep learning, i.e. introduce an image-based interpretation of the detector data and use **CNNs**
- representation based on space-points arranged in connected graphs could have an advantage given high dimensionality and sparsity of the tracking data → HEPtrkX project developed a set of **GNNs** to perform hits and segments classification

Or.. go quantum! (next)

Exploring a Quantum-GNN for tracking [2/2]

[1] EPJ Web of Conferences 245, 09013 (2020)

Towards **Q-GNN**: explore a quantum perspective of this GNN architecture, i.e. **re-implements the set of GNNs as quantum circuits**.

$$R_y(\theta)\,|0\rangle = \cos(\theta/2)\,|0\rangle + \sin(\theta/2)\,|1\rangle$$

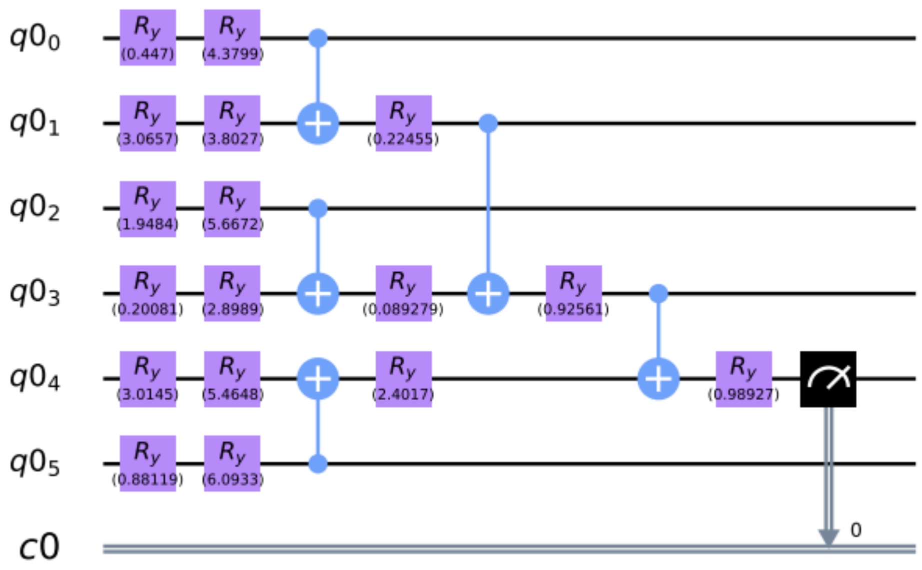

The Quantum Edge Network implemented as a Tree Tensor Network, a hierarchical quantum classifier. The architecture uses Ry rotation gates and CNOT gates. A single output qubit is measured.

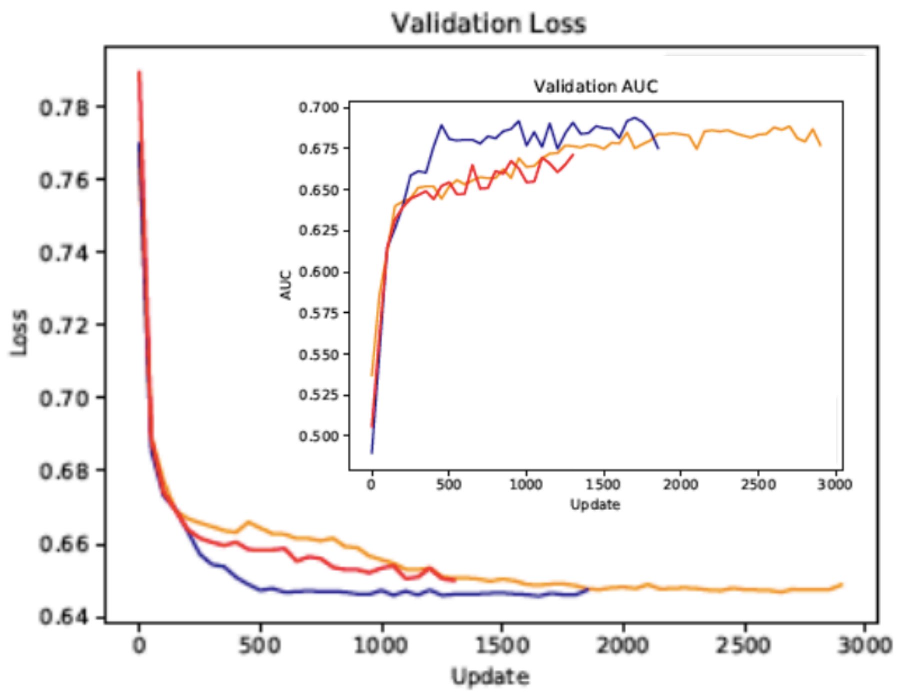

Q-GNN performance. Validation loss and AUC displayed for various # iterations. Results for two epochs, corresponding to 2900 steps (1 epoch = 1450 updates)

QAML for di-γ event classification [1/2]

The use of **quantum adiabatic ML** is proposed to classify events between the $H \rightarrow \gamma\gamma$ signal (S) and irreducible bkg (B) events with 2 uncorrelated γ

- 8 high-level features are measured from the di-γ system
- using such 8 features and their products as input, n=36 weak classifiers $c_i(x_\tau)$ are computed, assuming values in the range [−1, 1], S being represented by positive values
- a strong classifier is then constructed from a binary linear combination of the weak classifiers (with parameter $w_i \in \{0, 1\}$ for each weak classifier index *i*)
- The parameters w_i are then determined by the optimisation of a carefully crafted QUBO

$$E(\boldsymbol{w}) = \sum_{i,j=1}^{n=36} C_{ij} w_i w_j + \sum_{i=1}^{n=36} 2\left(\lambda - C_i\right) w_i$$

$$C_{ij} = \sum_{\tau} c_i\left(x_\tau\right) c_j\left(x_\tau\right)$$

$$C_i = \sum_{\tau} c_i\left(x_\tau\right) y_\tau$$

QAML for di-γ event classification [2/2]

Nature 550, 375–379 (2017)

The optimization is both run on the D-Wave 2X quantum annealer (QA) system and performed with simulated annealing (SA) using variable fractions of the training dataset

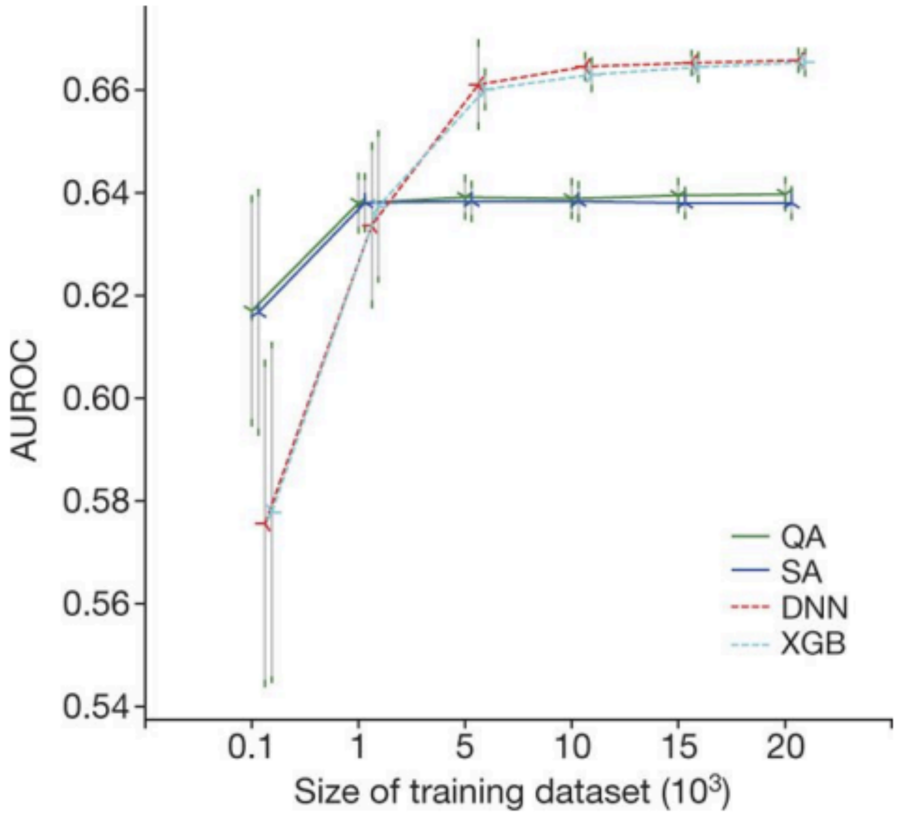

SA and QA are typically on par, and not providing obvious classification advantage over BDT and DNN, although a slight advantage with a small training dataset is noted

Computing infrastructures/tools

[1] arxiv.org/abs/2007.14781 (subm. to Comp.Sw.Big.Sci.)

[2] fastmachinelearning.org

(some of these aspects covered also at the "AI and HPC" kick-off workshop)

A variety of technology research aspects, and related assets.

Osmosis with "Istituto Nazionale di Fisica Nucleare" (INFN)

- one example: INFN Cloud → cloud.infn.it (contact: D.Salomoni)

Plus (not an exhaustive list):

- **(Big) Data Lake for HEP**, **AI-orchestrated data management**, **distributed caches**, etc
- prototyping a "cloudified" **MLaaS** systems [1] for HEP use-cases (and beyond)
- connection with the **hls4ml** [2] project: ultra-fast ML inference on FPGAs
- connection to LHC efforts on **heterogeneous computing**
- …

→ topical items for follow-up (post kick-off) meetings.

The future (?)

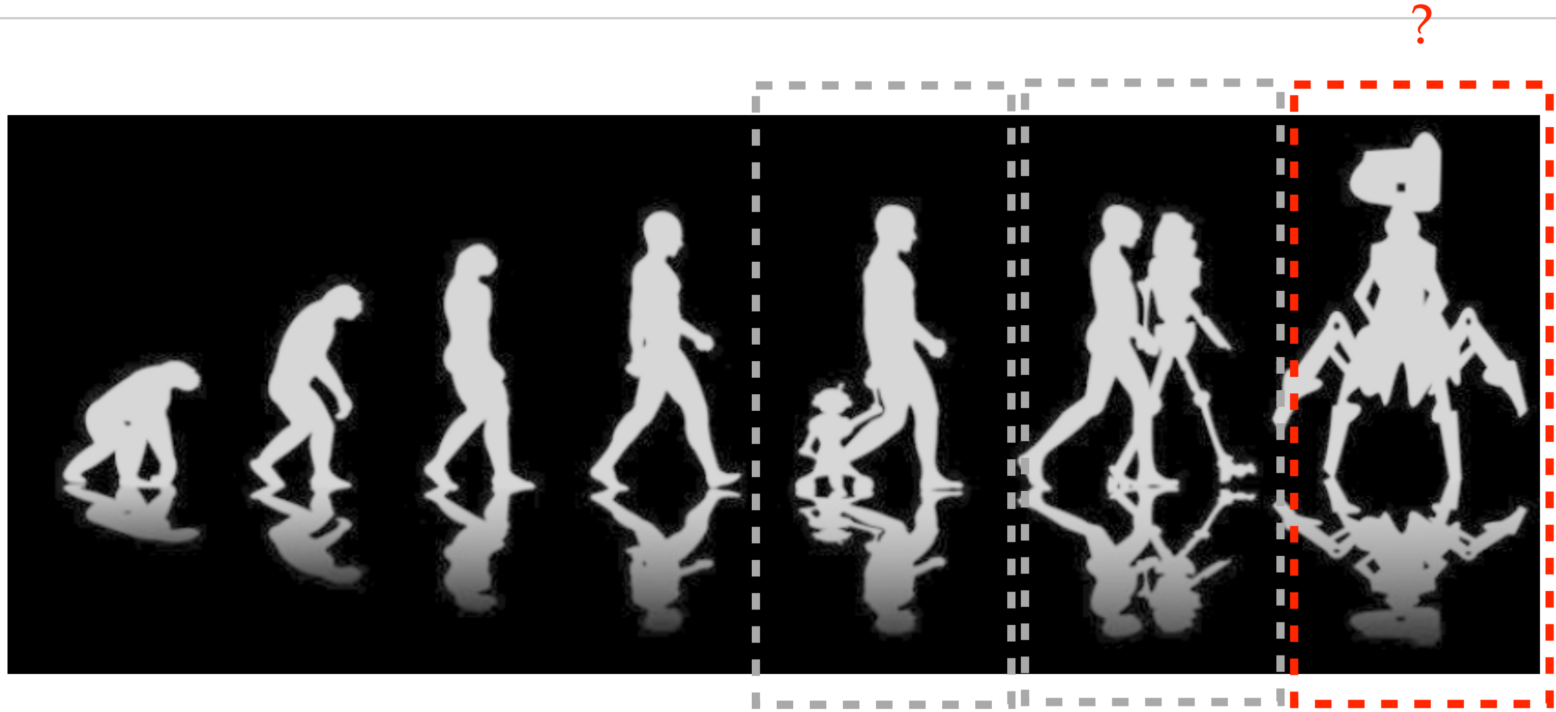

Enhancement of currently-successful NN architectures? Quantum ML? Neuromorphic computing? …?

The HEP community has always been able to contribute to and profit from technology progress. My personal opinion → HEP will do good if it speeds up in: 1) **embracing contamination** with neighbouring disciplines; 2) supporting **career recognition of young HEP physicists** passionate and committed to research in software/computing for HEP, including ML/DL/QML..

Conclusion

ML is ubiquitous in HEP, and continues to grow

- chosen techniques and adoption levels vary a lot across experiments
- techniques embraced over the years changed, and are still evolving

Cooperation among HEP and CS/ML researchers is crucial

- ability to engage ML experts on HEP problems
- ability to open up to non-HEP tools and frameworks, invest in education of youngest HEP physicists on advanced ML skills

If HEP will be able to drive this process, ML in HEP will play a crucial role towards meeting future challenges of data-intensive science at the energy and intensity frontier

Thanks for the attention

Credits

[1] Nature 560, 41–48 (2018)
[2] J. Phys. Conf. Series 1085 (2018) 022008
[3] Comp. Softw. Big Sci. 3, 7 (2019)

Thanks for material, discussion, inspirations to:

- M. Pierini, J.R. Vlimant, P. Vischia, M. Williams, S. May, S. Gleyzer, D. Bourilkov, co-authors of "*ML at the energy and intensity frontiers of particle physics*" [1], co-authors of the "ML in HEP - Community White Paper" [2], co-authors of "A Roadmap for HEP Software and Computing R&D for the 2020s" [3], coordinators of and contributor to the CMS ML forum, coordinators of and contributors to the IML team. With apologies to everyone I might have not mentioned.