

# On (track) reconstruction algorithms for the HL-LHC

Jan Stark Laboratoire des 2 Infinis - Toulouse (L2IT)









#### Reconstruction in ATLAS and the Athena software

#### (clickable link)

EUROPEAN ORGANISATION FOR NUCLEAR RESEARCH (CERN)



Eur. Phys. J. C 85 (2025) 234 DOI: 10.1140/epjc/s10052-024-13701-w



#### Software and computing for Run 3 of the ATLAS experiment at the LHC

The ATLAS Collaboration

The ATLAS experiment has developed extensive software and distributed computing systems for Run 3 of the LHC. These systems are described in detail, including software infrastructure and workflows, distributed data and workload management, database infrastructure, and validation. The use of these systems to prepare the data for physics analysis and assess its quality are described, along with the software tools used for data analysis itself. An outlook for the development of these projects towards Run 4 is also provided.

© 2025 CERN for the benefit of the ATLAS Collaboration.

Reproduction of this article or parts of it is allowed as specified in the CC-BY-4.0 license.

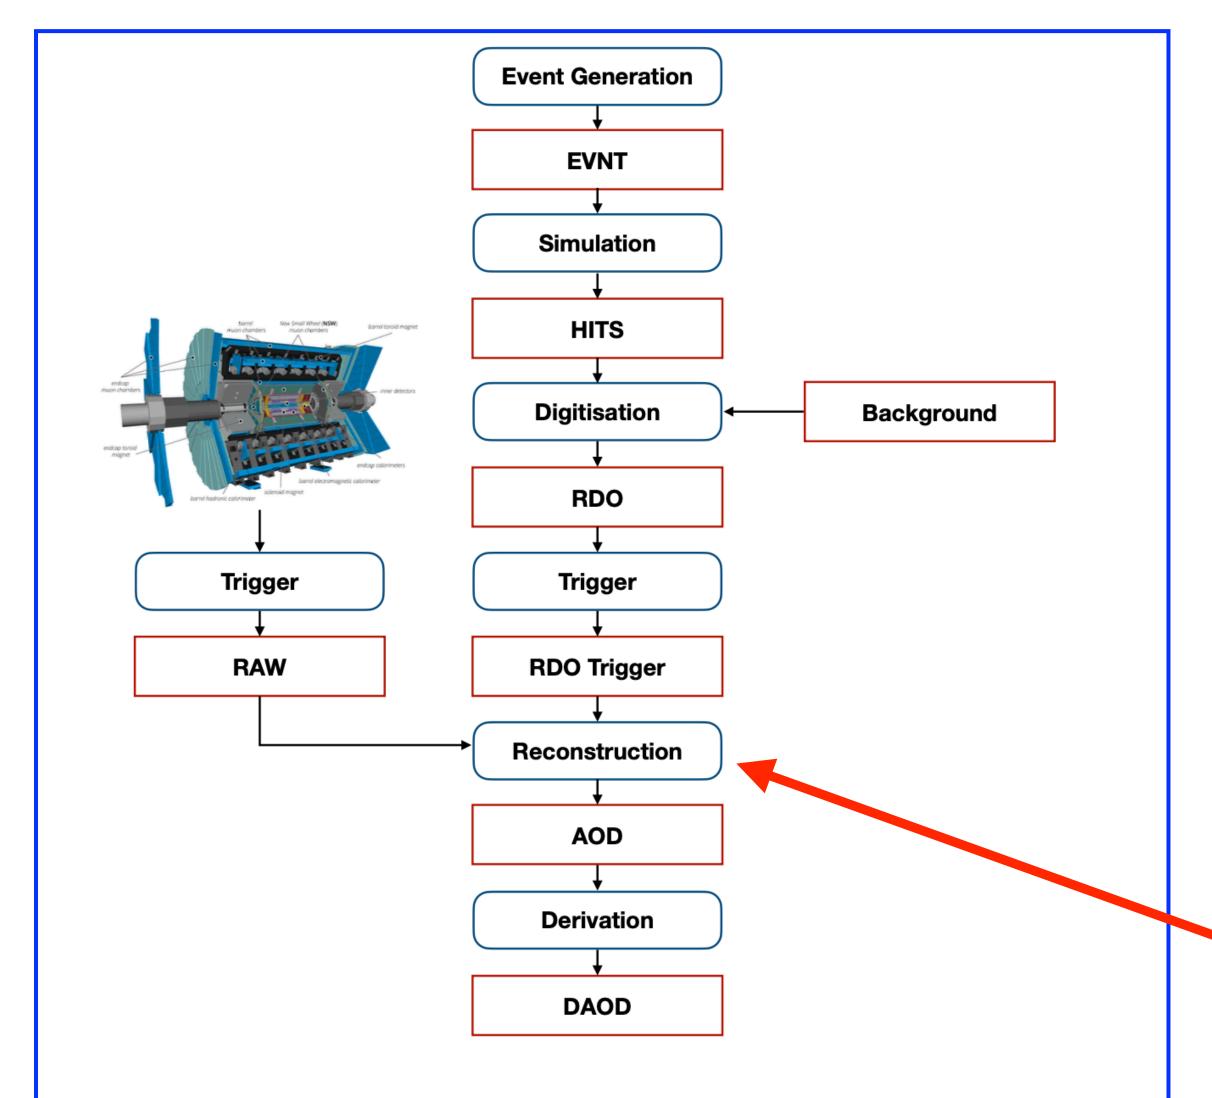


Figure 1: The standard software workflow of the ATLAS experiment. Processing steps are represented by blue ovals, with output formats represented as red boxes. The various steps and data formats are described in the text. The background entering digitisation may be additional simulated HITS files, pre-digitised RDO files, or specially processed RAW detector data.

Language	Files	Comment	Code
C++	17,273	457,373	2,608,231
Python	9,478	211,655	1,009,088
C/C++ Header	20,475	469,490	843,679
<b>Custom Configuration</b>	307	0	368,828
XML	954	12,800	204,169
Shell	1,243	12,283	48,782
CMake / make	2,070	11,021	35,751
Fortran	166	7,674	24,024
Web (HTML, CSS, PHP)	44	289	7,085
CUDA	28	648	5,445
Other	171	3,235	24,027
Total	52,191	1,186,288	5,178,472

5.2 million lines of code,+ 1.2 million lines of comments

Reconstruction

#### Reconstruction in ATLAS — the Athena software

3	Core software components	13			
	3.1 Core software	15			
	3.1.1 Updates for multithreading	19			
	3.2 Athena configuration	20			
	3.3 Conditions data handling	21			
	3.4 Event data model	22			
	3.5 Detector description	26			
	3.5.1 Detector description in Runs 1, 2, and 3	27			
	3.5.2 Detector description in Run 4	29			
	3.6 Machine learning and software infrastructure	30			
4	Data and Monte Carlo production and processing	31			
	4.1 Event generation	31			
	4.2 Detector simulation	35			
	4.2.1 Full simulation	36			

many of these aspects are discussed in the talk tomorrow by Francesco Di Bello (clickable link)

#### 3.6 Machine learning and software infrastructure

Machine learning (ML) is used in a wide range of applications in the ATLAS Collaboration, including physics analysis, simulation and physics object identification in the reconstruction and trigger system. Models trained for classification and regression, as well as generative models, are core components of the software and analysis chain. This section describes the infrastructure and tools for ML used within the collaboration.

Classification models are the most commonly employed. For physics analyses, these models are used to categorise events into signal-like or background-like samples, efficiently distinguishing events of interest. Often multivariate event classifiers based on boosted decision trees (BDTs) are used [72–85]. BDTs have also been used for regression [86]. Other analyses have explored the use of neural networks (NNs) for classification, for example in top-quark physics analyses [87–94] to separate signal and background processes. Classification models are also used extensively to accurately identify physics objects, with applications in the identification of hadronically-decaying  $\tau$ -leptons [95], boosted jets [96] and heavy-flavour jets [32]. More recently, recurrent neural networks (RNNs) [97] are exploited for their ability to harness sequential characteristics to take into account correlations between track impact parameters, thereby improving physics performance for jet flavour tagging [98] and  $\tau$ -lepton reconstruction [99] (see Section 4.4).

Artificial neural networks have also been used for lower-level reconstruction tasks. For example, to identify merged clusters in the pixel detector, which are created by multiple charged particles traversing the silicon, a clustering algorithm is implemented to identify and split the clusters and provide improved position estimates for the individual cluster pieces [100].

(...)

The first general purpose inference framework integrated into the ATLAS software environment LWTNN [111], which uses the Boost library [112] as a JSON parser and EIGEN [113] for tensor operations. Users can convert their trained NN models into JSON format, defining the weights and layer operations. This is then handled by a lightweight class for running inference on single events. For the inference of BDTs, a custom wrapper was written that converts models trained with TMVA, LIGHTGBM or XGBoost into a common ROOT TTree format for inference at runtime. However, as BDTs are mostly used in physics analyses, other libraries are compiled in standalone ntuple processing frameworks to provide inference. Due to the rise of more complex models, which use layers and operations that are not supported by LWTNN, the ONNX Runtime (ORT) library [114] was chosen as a general inference library. Most modern ML libraries can save models in an ORT format, or are supported with conversion tools. These include, but are not limited to, PyTorch, TensorFlow, LightGBM and XGBoost. Using only a single inference library to support many models trained in different frameworks significantly reduces the maintenance load. ORT is observed to provide faster inference than previous implementations and support for batched

### ML model inference in C++ — and in Athena

#### Slides taken from:

Antonio Giannini, « ONNX-based inference in ATLAS », ILC software meeting, 15 May 2024 (clickable link)



#### ML models vs analyses

- ML applications have had a huge boost in Physics analyses at collider experiments
  - during my ATLAS experience, ~O(1) in 2017 up to ~O(10) in 2024 of ML applications in analyses
  - searches better use cases for applications/development than measurements
- Wide variety of python based frameworks are used now
  - TensorFlow, Keras, PyTorch, sci-kit learn, etc etc
  - relatively easy to get started, many tutorials around, many people willing to pass the knowledge
- This implies
  - a large population of "custom" frameworks doing mostly the same things
  - a large population of datasets, intermediate/output formats, ROOT, pandas, h5, etc

ILC software meeting

• All of these ingredients can make really hard the reproducibility of results!!!





- ML models inference in C++ is crucial for two main aspects
- Reproducibility of results outside the collaboration
  - event classifier (NN>x), final discriminant (to complete the analysis ingredients)
  - anomaly detections (when the ML is the main analysis ingredient...)
  - Deployment of a tool that can be used for several analyses
    - jet tagging, low-level variables based classifiers, etc

Inference is not a plus, nor an aesthetic element, it is a significant part of the

R&D of a ML tool for physics!

4

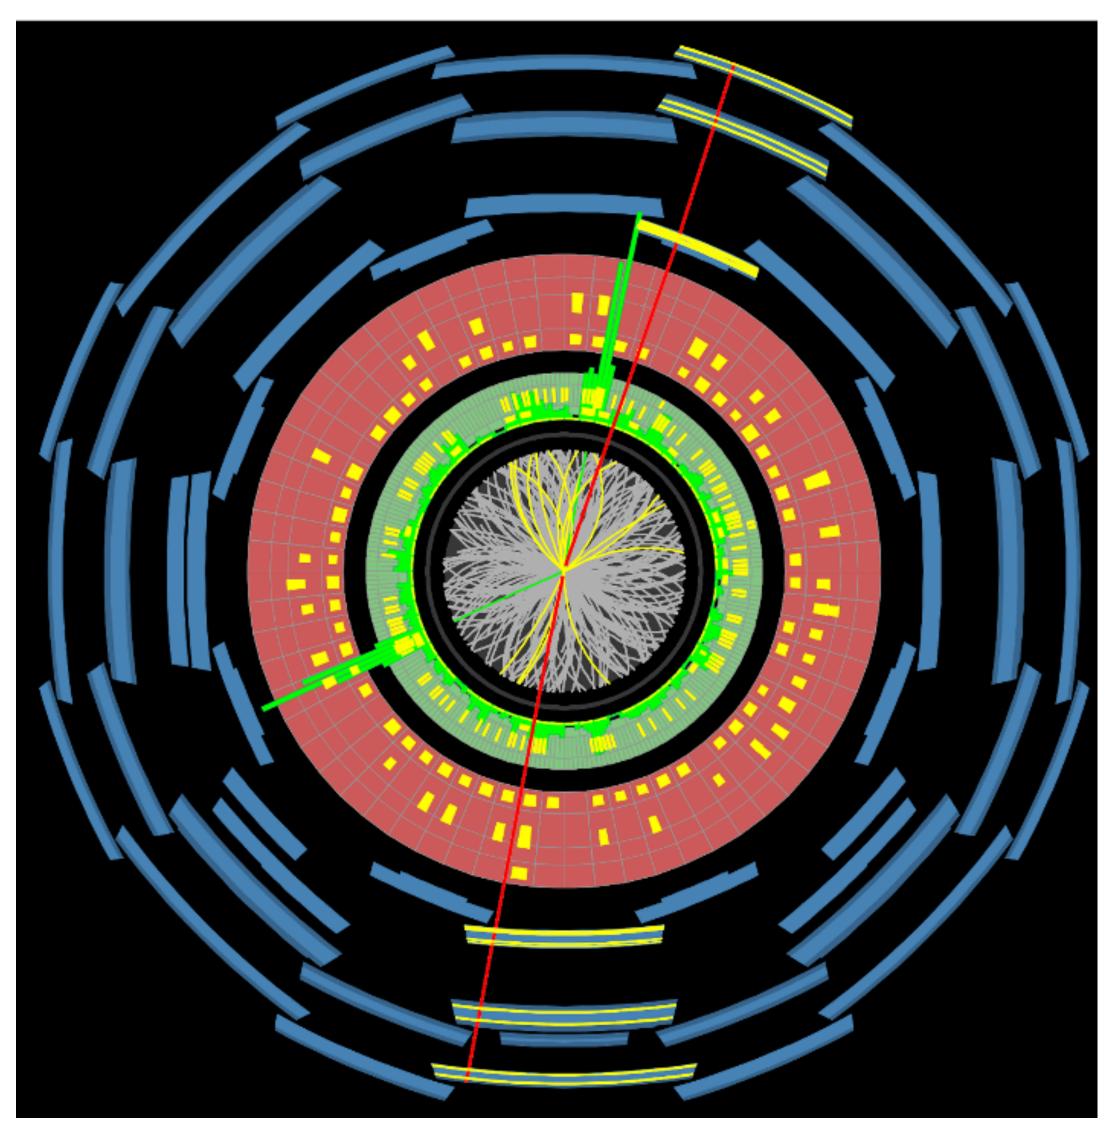
ILC software meeting

7

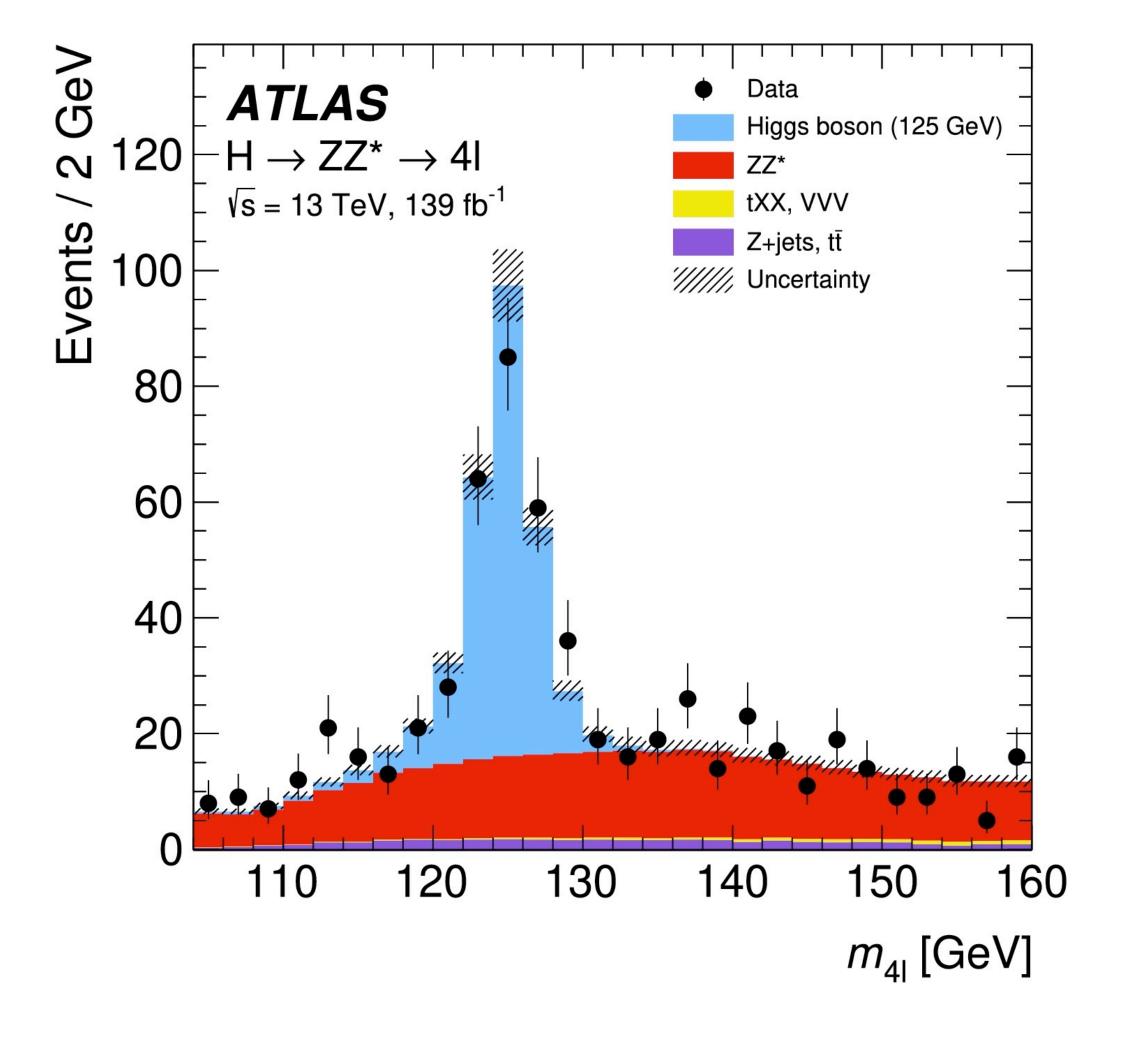
including in Athena, for running on CPU farms

# Higgs boson

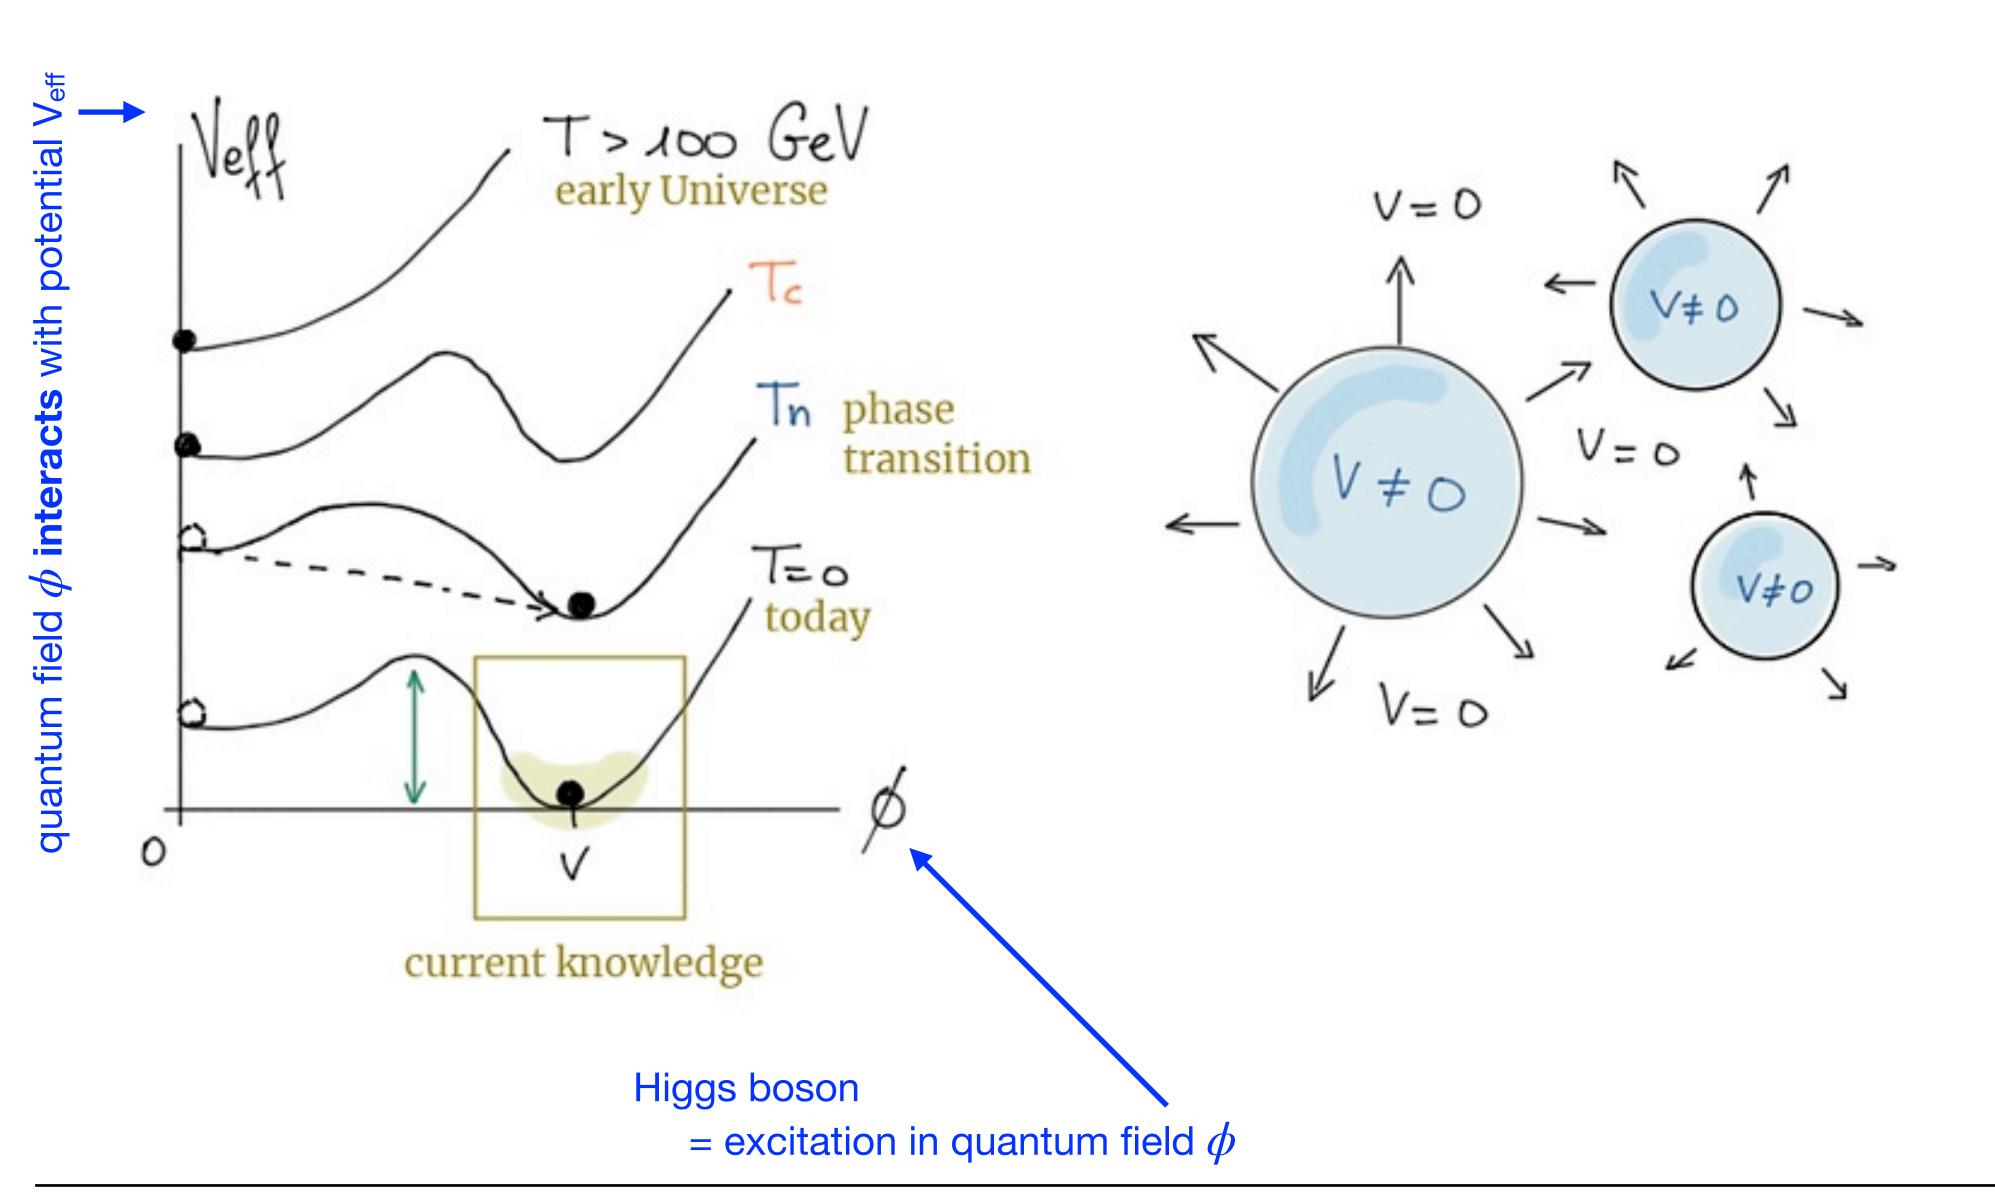
#### One event seen in the ATLAS detector



# Counting events in different intervals of **reconstructed** H -> 4I mass $(m_{4l})$



# Higgs boson: its role in the early Universe



Need much more data to study the shape of V<sub>eff</sub>

(because information is given by rare events)

Figure: Kateryna Radchenko Cluster of Excellence "Quantum Universe" Universität Hamburg

#### Taking data at the LHC is like drinking water out of a firehose



Rheinfall (Rhine Falls): 750 m<sup>3</sup>/second

Data from detector
40 million events / second
60 TB / second



(big) firehose:19 liters/second

Data recorded to **disk** 1000 events /second 1.5 GB / second

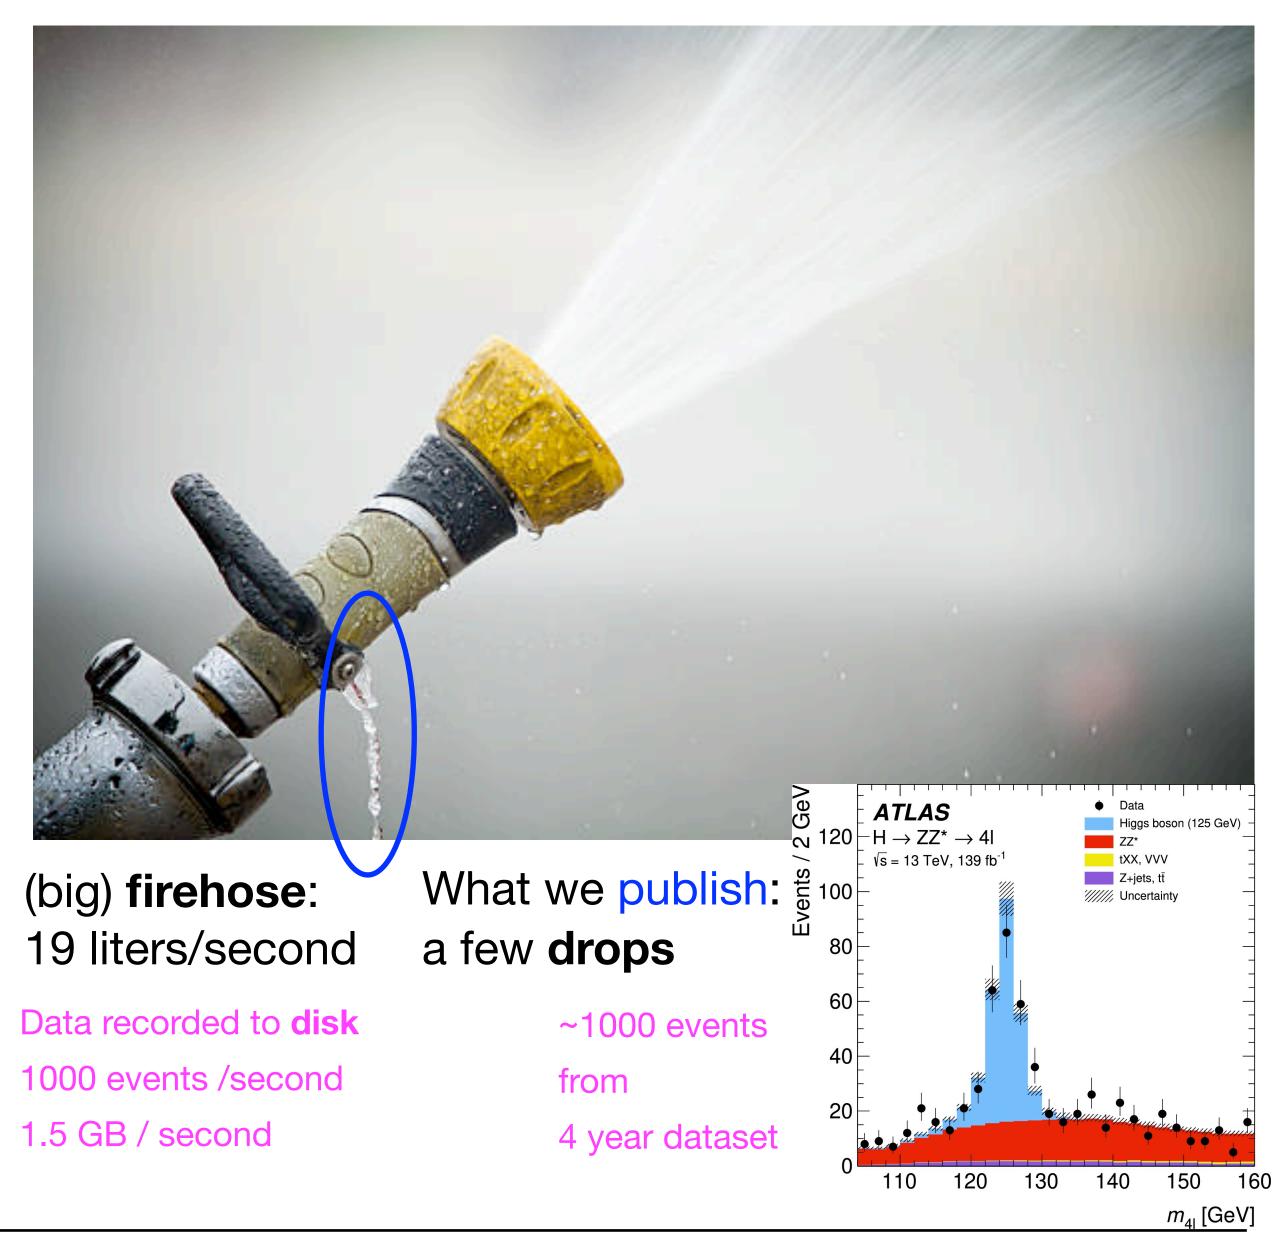
#### Taking data at the LHC is like drinking water out of a firehose



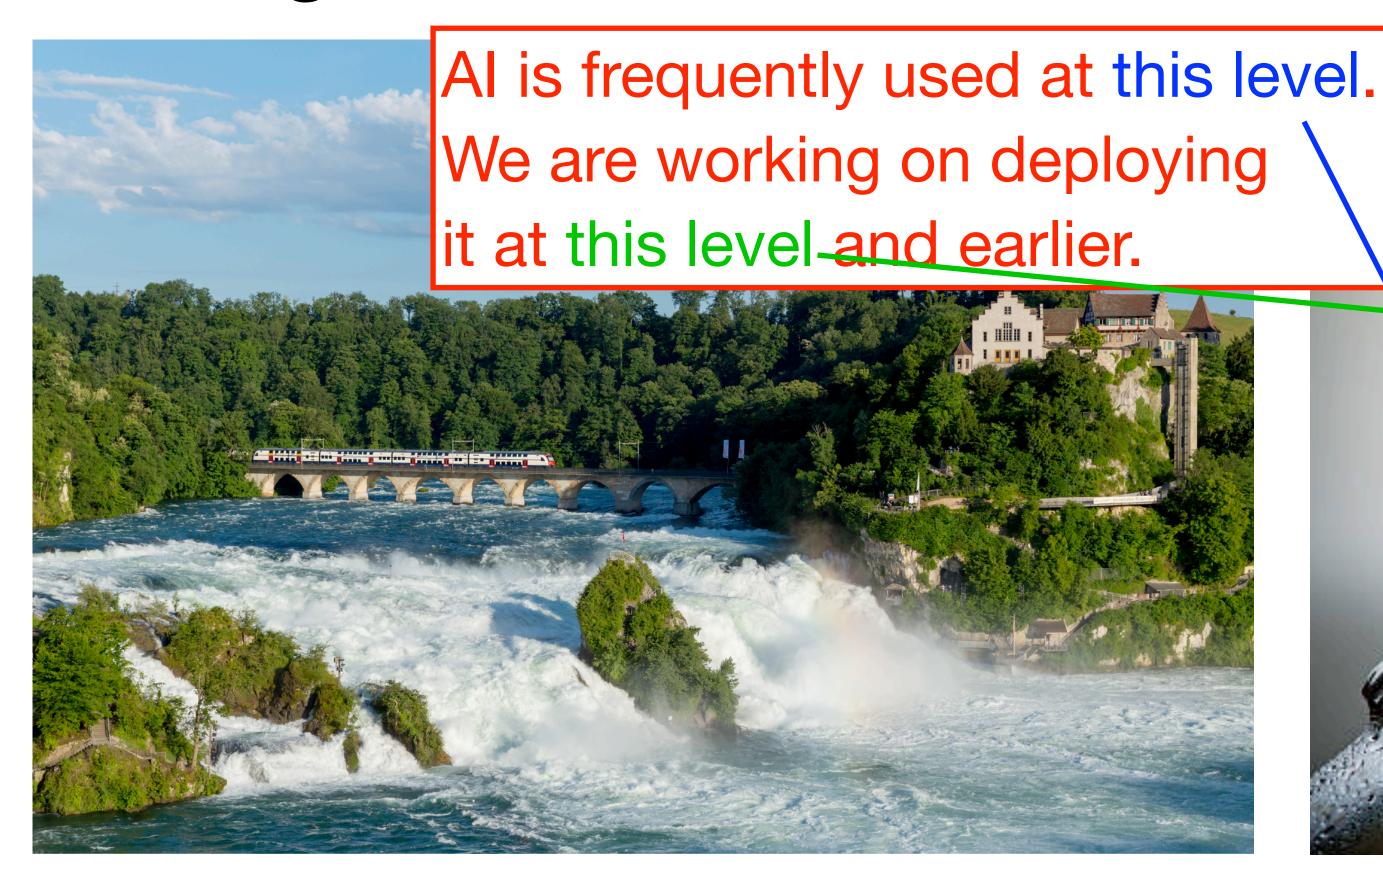
Rheinfall (Rhine Falls): 750 m<sup>3</sup>/second

Data from detector
40 million events / second

60 TB / second



#### Taking data at the LHC is like drinking water out of a firehose



Rheinfall (Rhine Falls): 750 m<sup>3</sup>/second

Data from detector

40 million events / second 60 TB / second

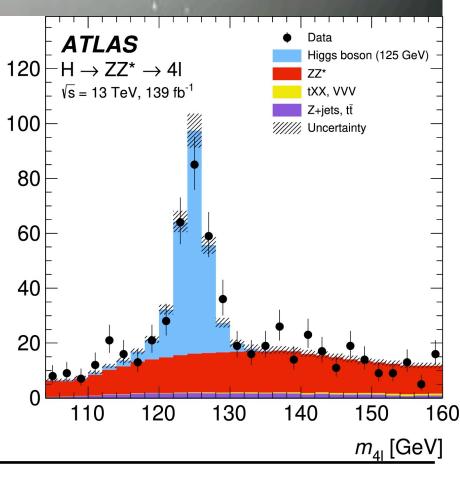
(big) firehose: 19 liters/second Data recorded to **disk** 1000 events /second

1.5 GB / second

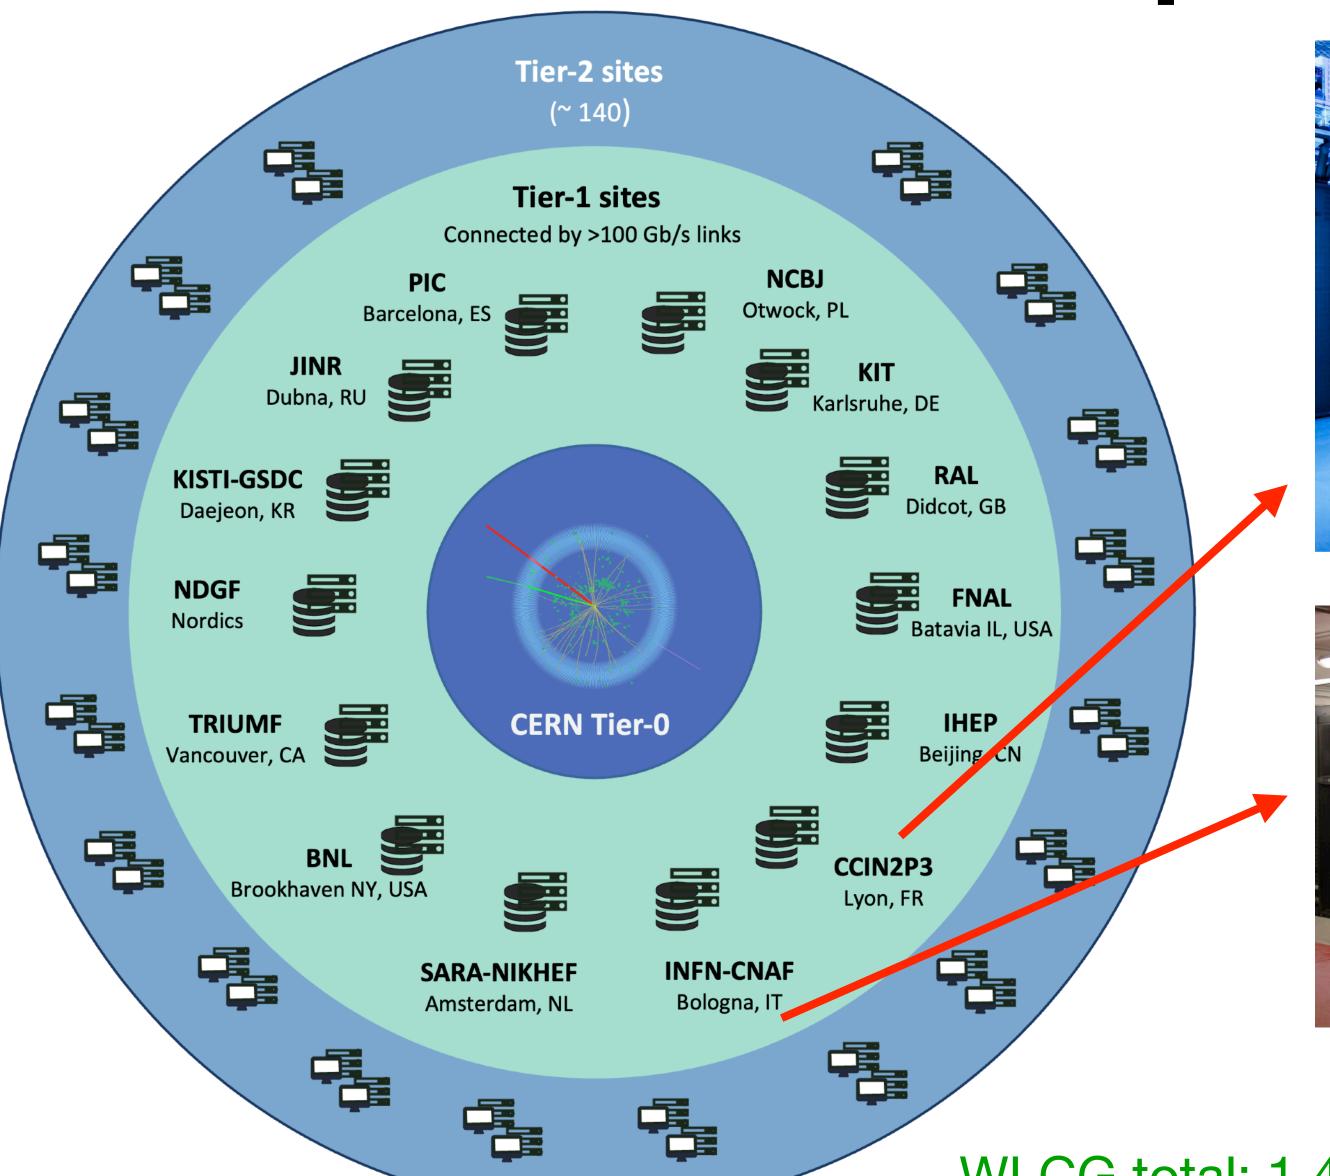
What we publish: \frac{4}{5} 100 \rightarrow a few **drops** ~1000 events

from

4 year dataset



# Worldwide LHC computing grid (WLCG)





Centre de Calcul de l'IN2P3 (Lyon, FR)





WLCG total: 1.4 million CPU cores and 1.5 exabytes of storage

#### "10 years to prepare ourselves" for HL-LHC (statement from 2017)

- Community white paper (2017)
  - Algorithms, infrastructure, data access...

Great overview of ongoing changes in computing industry, current practices in HEP and required R&D activities in key domains:

- Physics generators
- Detector simulation
- Software trigger & event reconstruction
- Data analysis
- Machine learning
- Data management, ...
- Facilities, distributed computing

- ....

2020 update of the European strategy for particle physics



arXiv.org > physics > arXiv:1712.06982 Help | Advanced Search Physics > Computational Physics Download: A Roadmap for HEP Software and Computing R&D for the 2020s PDF Other formats Johannes Albrecht, Antonio Augusto Alves Jr, Guilherme Amadio, Giuseppe Andronico, Nguyen Anh-Ky, Laurent Aphecetche, John Apostolakis, Makoto Asai, Luca Atzori, Marian Babik, Giuseppe Bagliesi, Marilena Bandieramonte Current browse context: Sunanda Banerjee, Martin Barisits, Lothar A.T. Bauerdick, Stefano Belforte, Douglas Benjamin, Catrin Bernius, Wahid physics.comp-ph Bhimji, Riccardo Maria Bianchi, Ian Bird, Catherine Biscarat, Jakob Blomer, Kenneth Bloom, Tommaso Boccali, Brian < prev | next > new | recent | 1712 Bockelman, Tomasz Bold, Daniele Bonacorsi, Antonio Boveia, Concezio Bozzi, Marko Bracko, David Britton, Andy Buckley, Predrag Buncic, Paolo Calafiura, Simone Campana, Philippe Canal, Luca Canali, Gianpaolo Carlino, Nuno Change to browse by: Castro, Marco Cattaneo, Gianluca Cerminara, Javier Cervantes Villanueva, Philip Chang, John Chapman, Gang Chen, hep-ex physics Taylor Childers, Peter Clarke, Marco Clemencic, Eric Cogneras, Jeremy Coles, Ian Collier, David Colling, Gloria Corti, Gabriele Cosmo, Davide Costanzo, Ben Couturier, Kyle Cranmer, Jack Cranshaw, Leonardo Cristella, David Crooks, References & Citations INSPIRE HEP Sabine Crépé-Renaudin, Robert Currie, Sünje Dallmeier-Tiessen, Kaushik De, Michel De Cian, Albert De Roeck, (refers to | cited by )
NASA ADS Antonio Delgado Peris, Frédéric Derue, Alessandro Di Girolamo, Salvatore Di Guida, Gancho Dimitrov, Caterina Doglioni, Andrea Dotti, Dirk Duellmann, Laurent Duflot, Dave Dykstra, Katarzyna Dziedziniewicz-Wojcik, Agnieszka Export citation Google Scholar Dziurda, Ulrik Egede, Peter Elmer, Johannes Elmsheuser, V. Daniel Elvira, Giulio Eulisse, Steven Farrell, Torben Ferber, Andrej Filipcic, Ian Fisk, Conor Fitzpatrick, José Flix, Andrea Formica, Alessandra Forti, Giovanni Franzoni, Bookmark James Frost, Stu Fuess, Frank Gaede, Gerardo Ganis, Robert Gardner, Vincent Garonne, Andreas Gellrich et al. (210 Science Wise additional authors not shown) (Submitted on 18 Dec 2017 (v1), last revised 19 Dec 2018 (this version, v5)) Particle physics has an ambitious and broad experimental programme for the coming decades. This programme requires large (link) recorded. In planning for the HL-LHC in particular, it is critical that all of the collaborating stakeholders agree on the software goals

D. Large-scale data-intensive software and computing infrastructures are an essential ingredient to particle physics research programmes. The community faces major challenges in this area, notably with a view to the HL-LHC. As a result, the software and computing models used in particle physics research must evolve to meet the fature needs of the field. The community must vigorously pursue common, coordinated R&D efforts in collaboration with other fields of science and industry, to develop software and computing infrastructures that exploit recent advances in information technology and data science. Further development of internal policies on open data and data preservation should be encouraged, and an adequate level of resources invested in their implementation.

#### "10 years to prepare ourselves" for HL-LHC (statement from 2017)

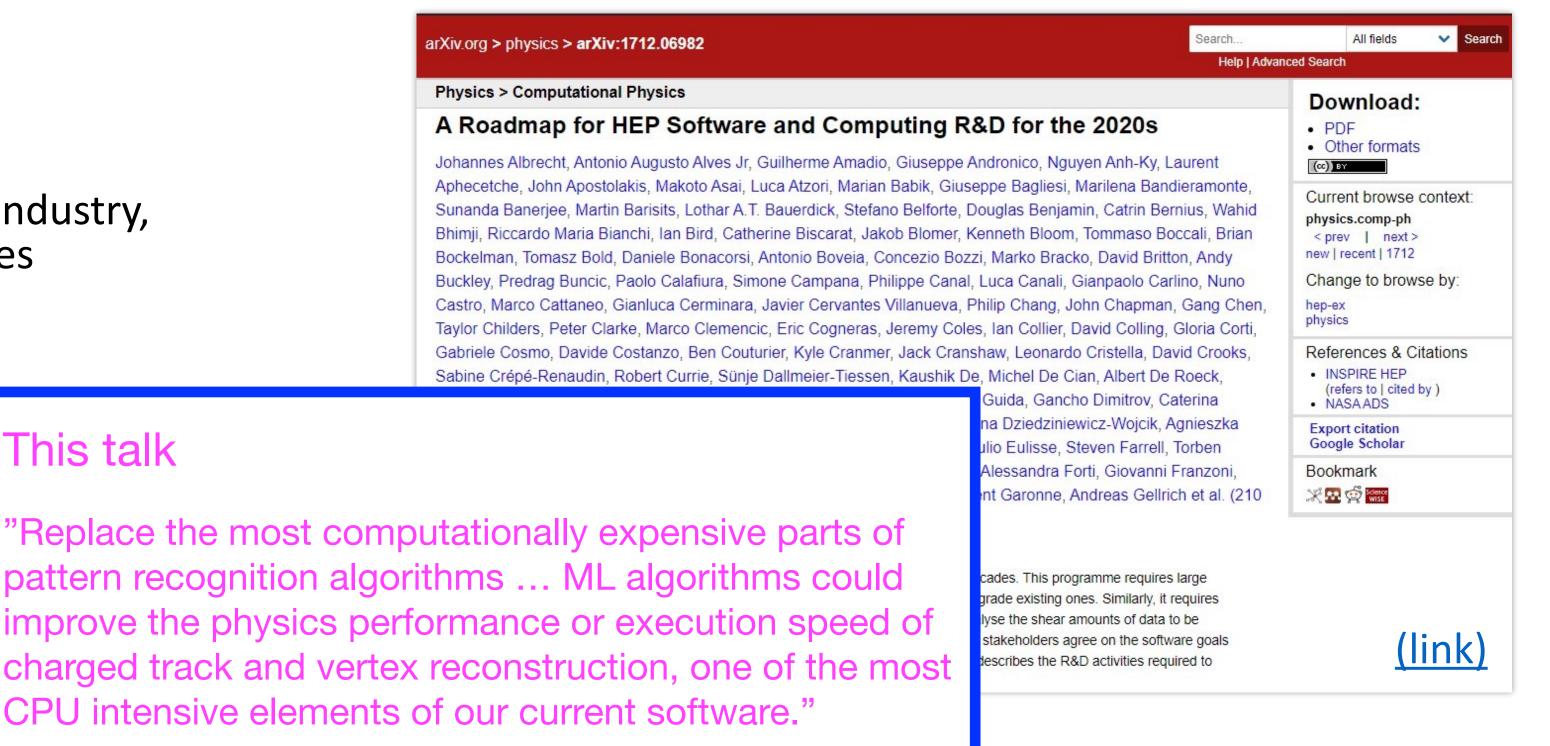
- Community white paper (2017)
  - Algorithms, infrastructure, data access...

Great overview of ongoing changes in computing industry, current practices in HEP and required R&D activities in key domains:

- Physics generators
- Detector simulation
- Software trigger & event reconstruction -
- Data analysis
- Machine learning
- Data management, ...
- Facilities, distributed computing

J. Stark

2020 update of the European strategy for particle physics





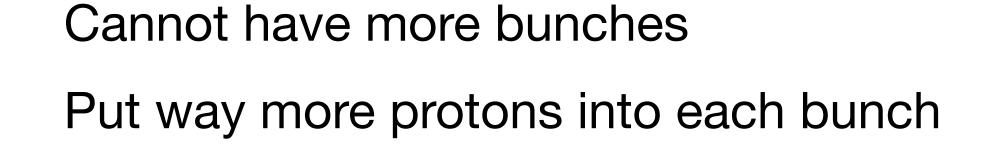
This talk

Large-scale data-intensive software and computing infrastructures are an essential ingredient to particle physics research programmes. The community faces major challenges in this area, notably with a view to the HL-LHC. As a result, the software and computing models used in particle physics research must evolve to meet the future needs of the field. The community must vigorously pursue common, coordinated R&D efforts in collaboration with other fields of science and industry, to develop software and computing infrastructures that exploit recent advances in information technology and data science. Further development of internal policies on open data and data preservation should be encouraged, and an adequate level of resources invested in their implementation. (link)

## LHC

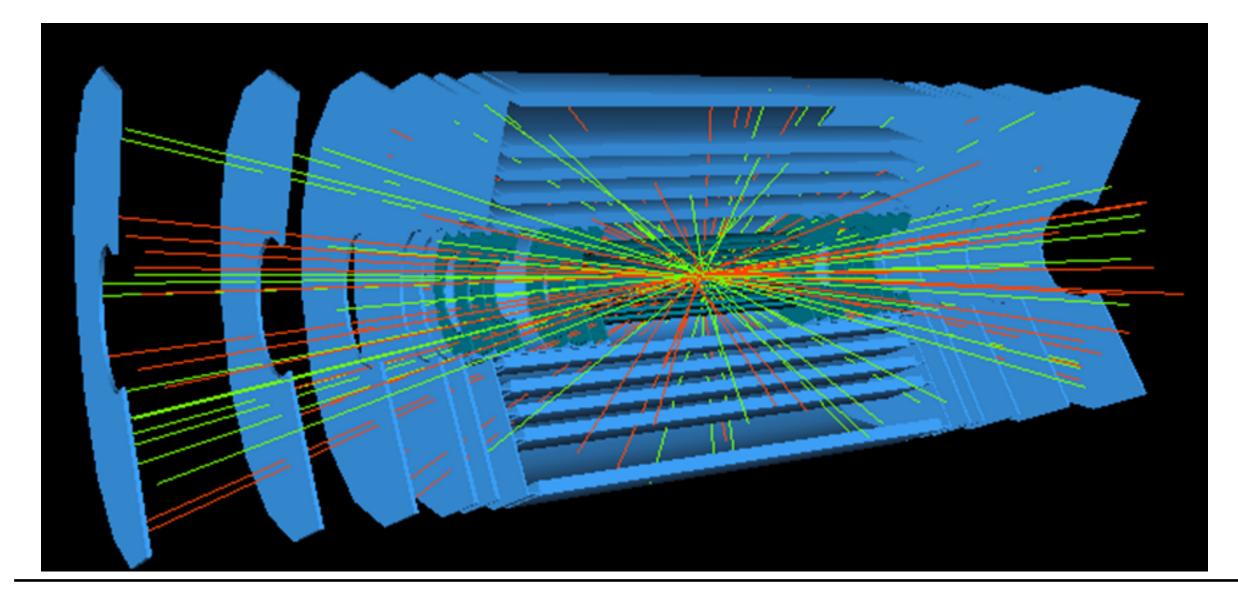
## HL-LHC

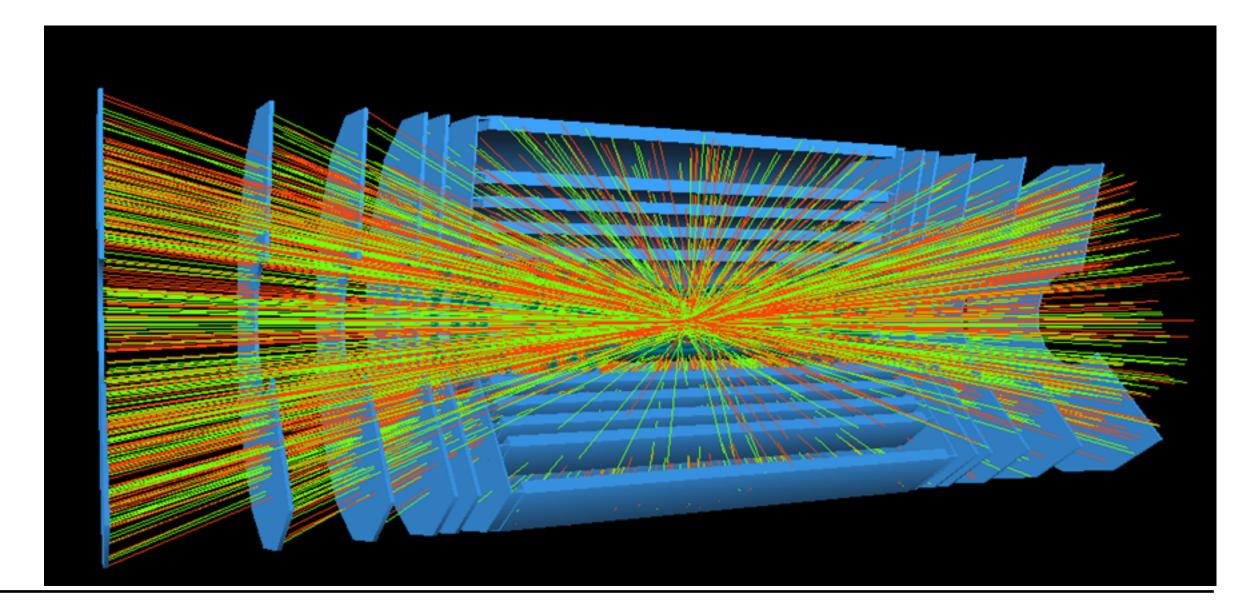
Beams in the LHC: bunches of protons 40 million bunch crossings per second











## Machine Learning for track pattern recognition?



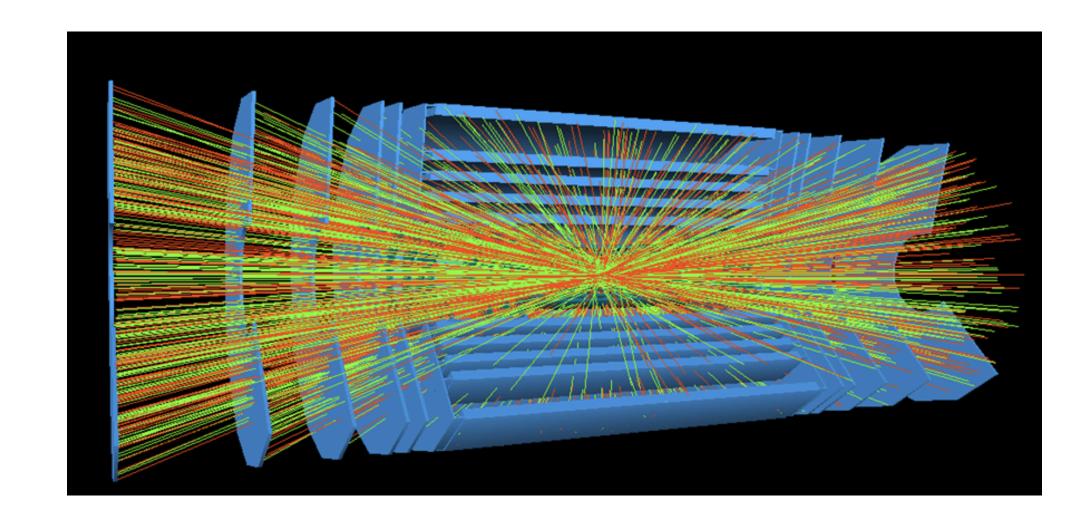
622 \* 415 pixels

a large fraction carries information about the person



Can't use the same tools

How to present tracking data to a neural network?



ATLAS tracker for HL-LHC:

5 \* 10<sup>9</sup> readout channels

~3 \* 10<sup>5</sup> 3D space-points per event

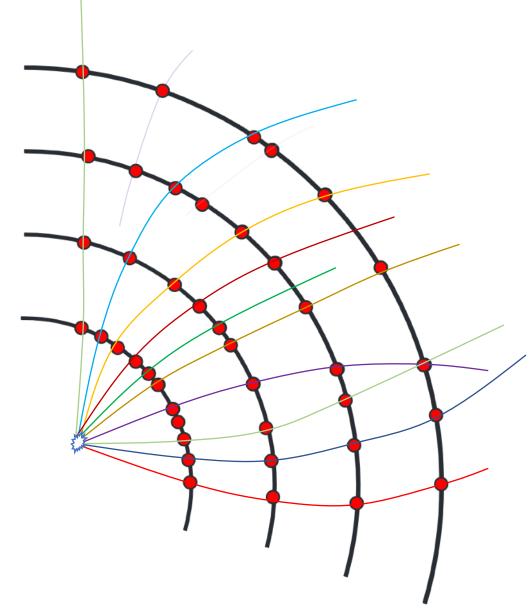
=> data are *sparse* 

# Representing tracking data using graphs

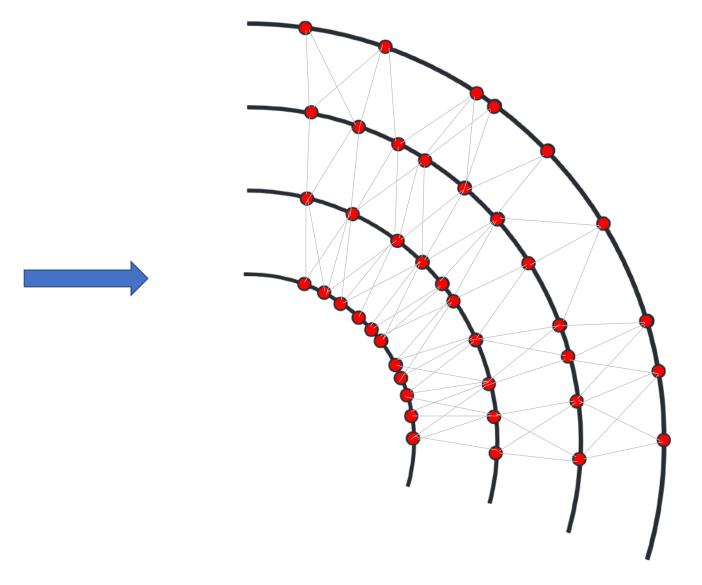
F. Siklér, "Combination of various data analysis techniques for efficient track reconstruction in very high multiplicity events", Connecting the Dots conference 2017 (link)

S. Farrell *et al.*, "Novel deep learning methods for track reconstruction", proceedings of *Connecting the Dots* conference 2018 (link)

Charged particles leave hits in the detector



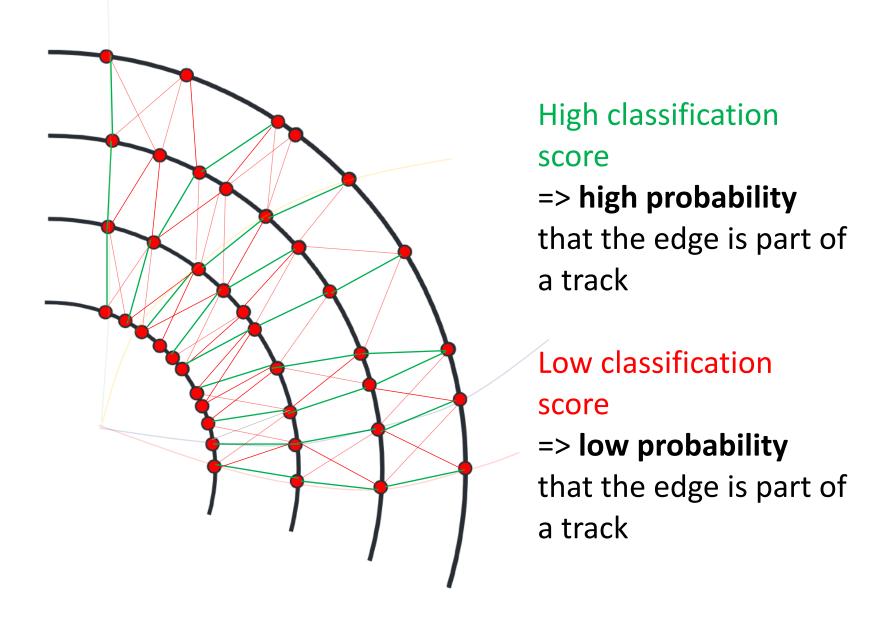
Represent the data using a graph



One node of the graph = one hit in the detector

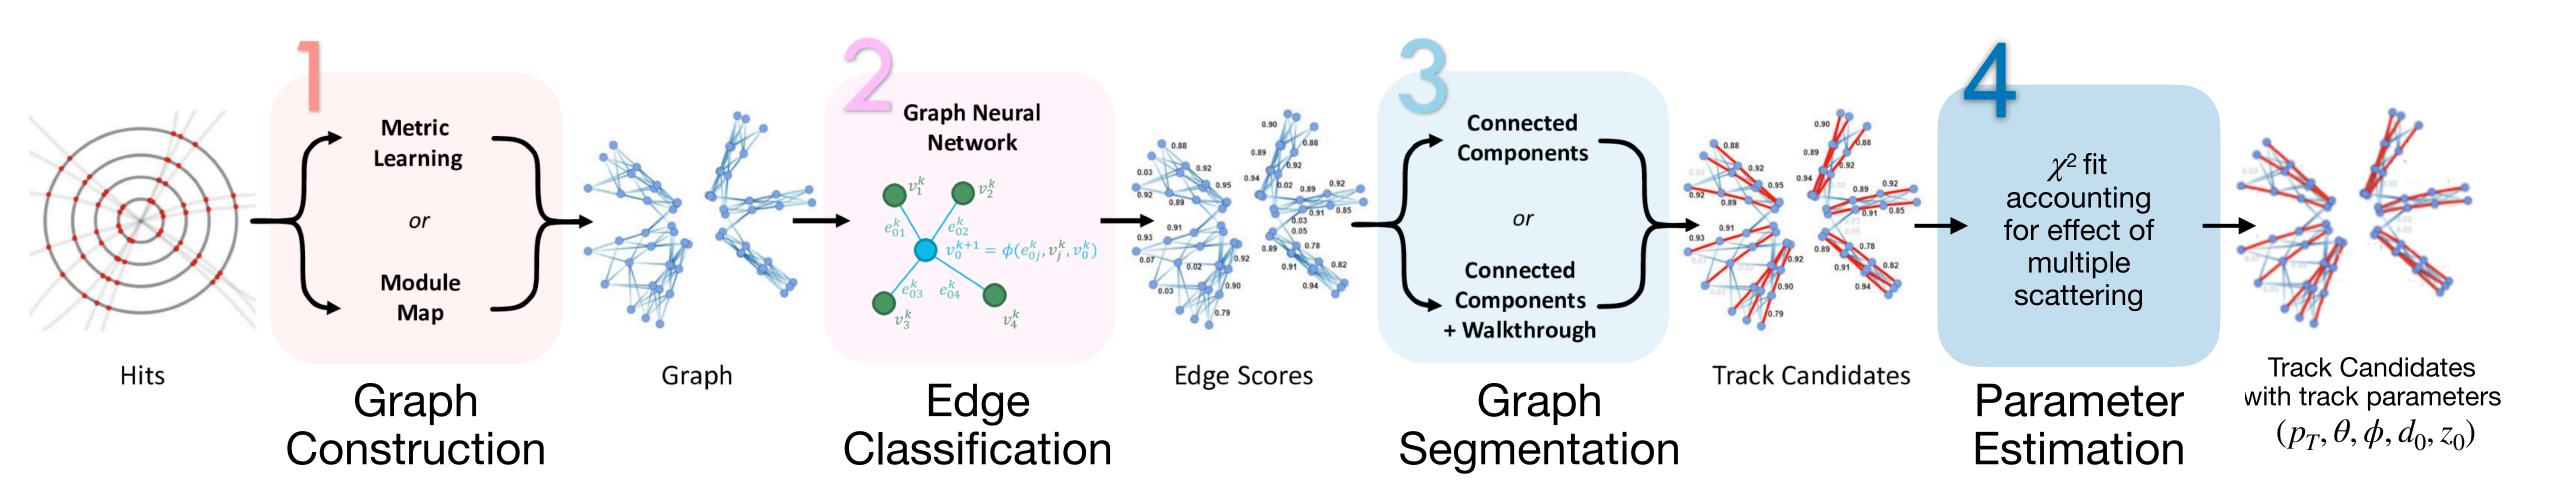
Connect two nodes using an edge if "it seems possible" that the two hits are two (consecutive) hits on a track

## Goal: classify the edges of the graph

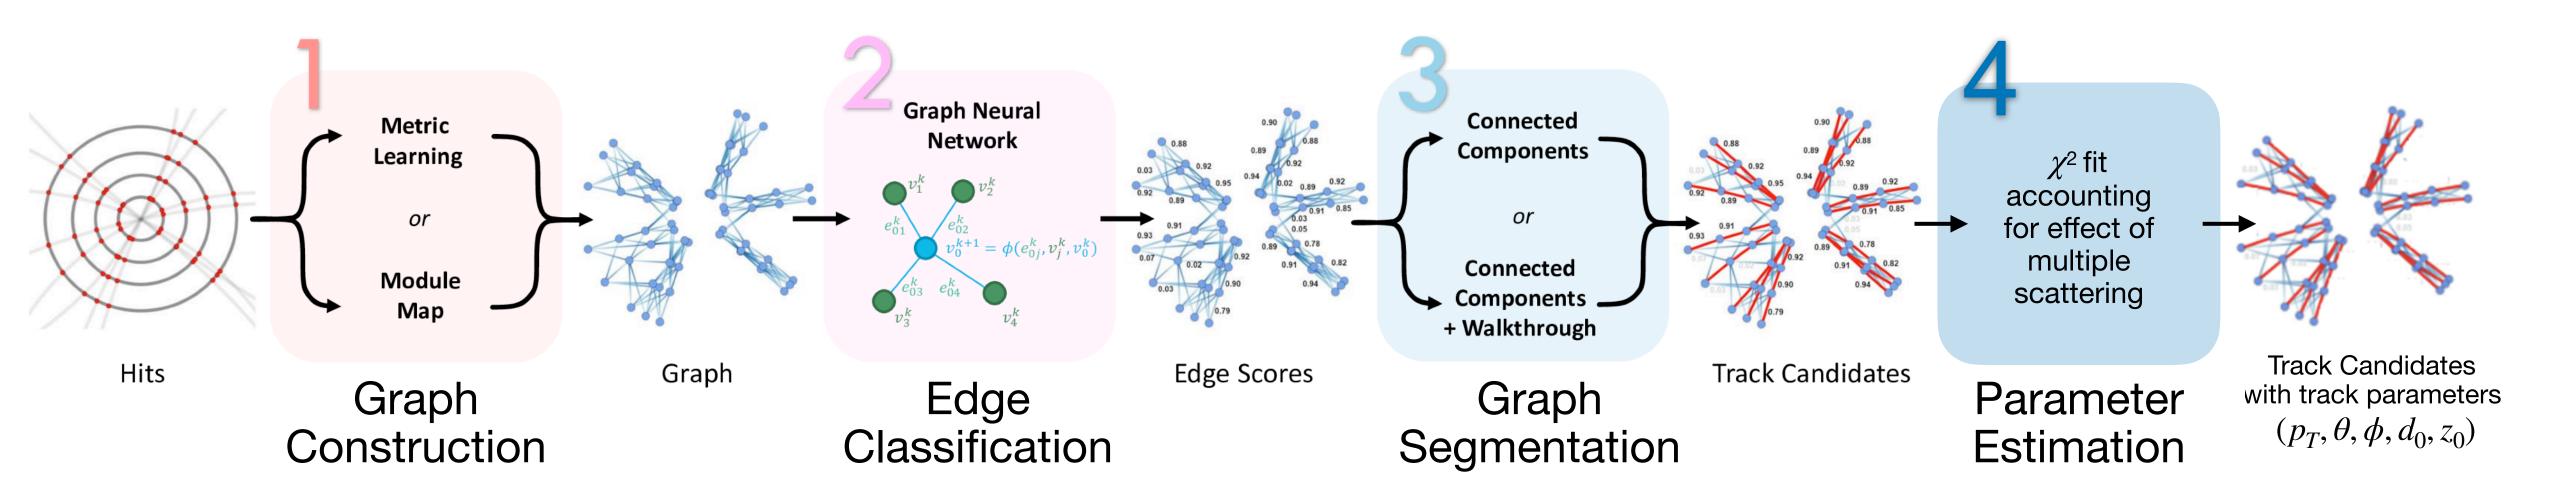


More general review article: "GNNs at the LHC" (link)

# Track pattern recognition using GNNs



# Track pattern recognition using GNNs



F. Siklér, talk at Connecting the Dots 2017

C. Biscarat, S. Caillou, C. Rougier, J. Stark and J. Zahreddine, EPJ Web of Conferences 251, 03047 (2021)

X. Ju et al., Eur. Phys. J. C 81, 876 (2021)

## ML versus classical algorithms

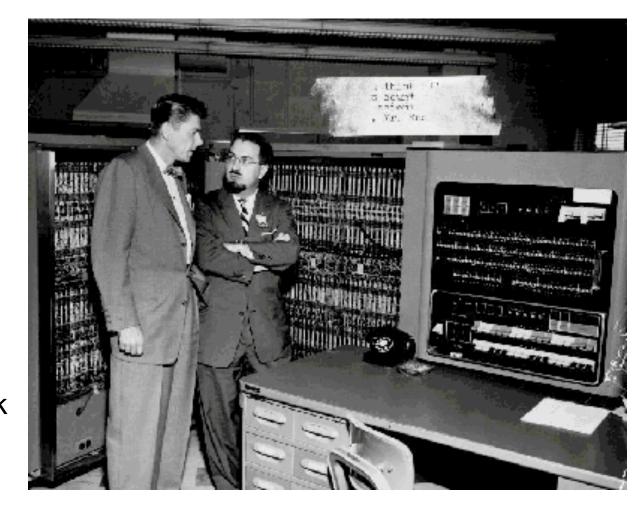
#### Automatic translation of text

IBM Press release, January 8, 1954

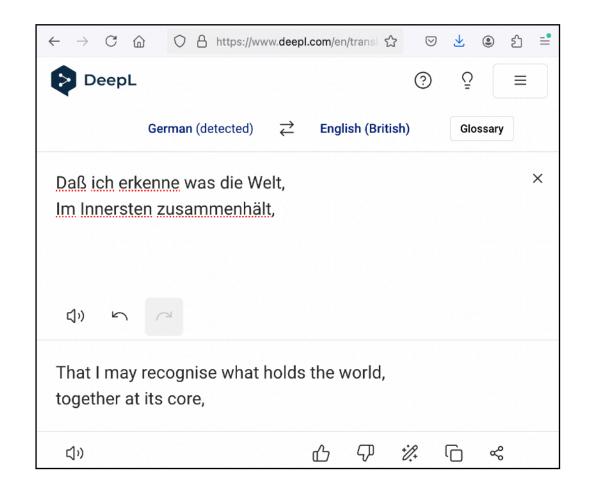
New York, January 7..... Russian was translated into English by an electronic "brain" today for the first time.

[...]

A girl who didn't understand a word of the language of the Soviets punched out the Russian messages on IBM cards. The "brain" dashed off its English translations on an automatic printer at the breakneck speed of two and a half lines per second.

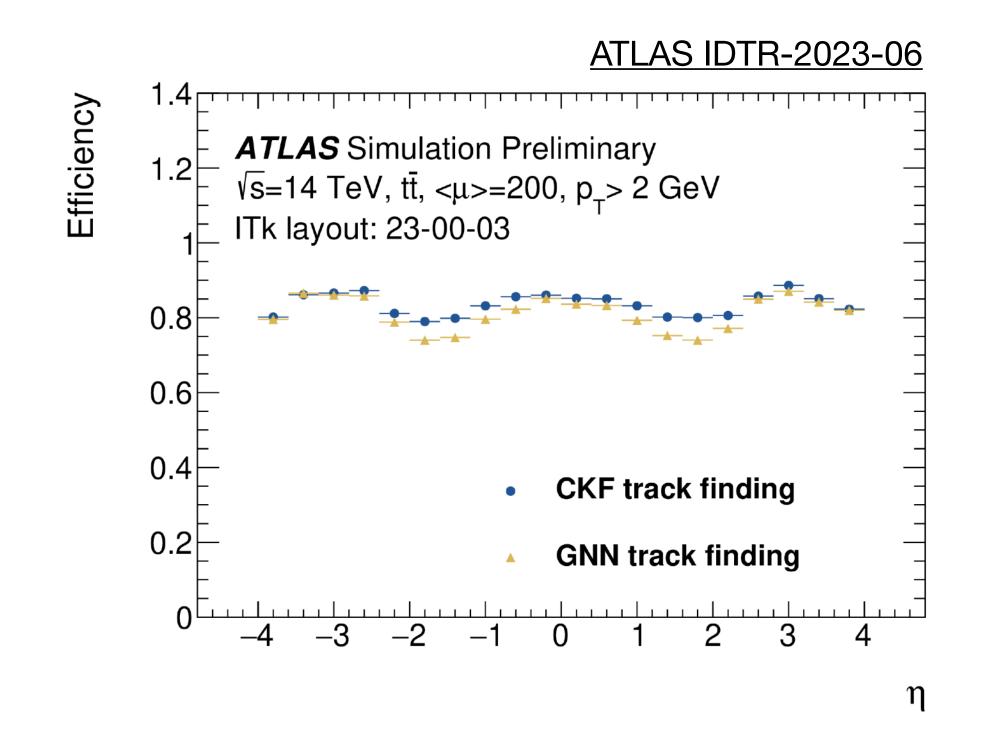


However, the triumphant headlines hid one little detail. No one mentioned ... (link)



Today: ML achieves what classical algorithms do not

#### Reconstruction of charged particle tracks

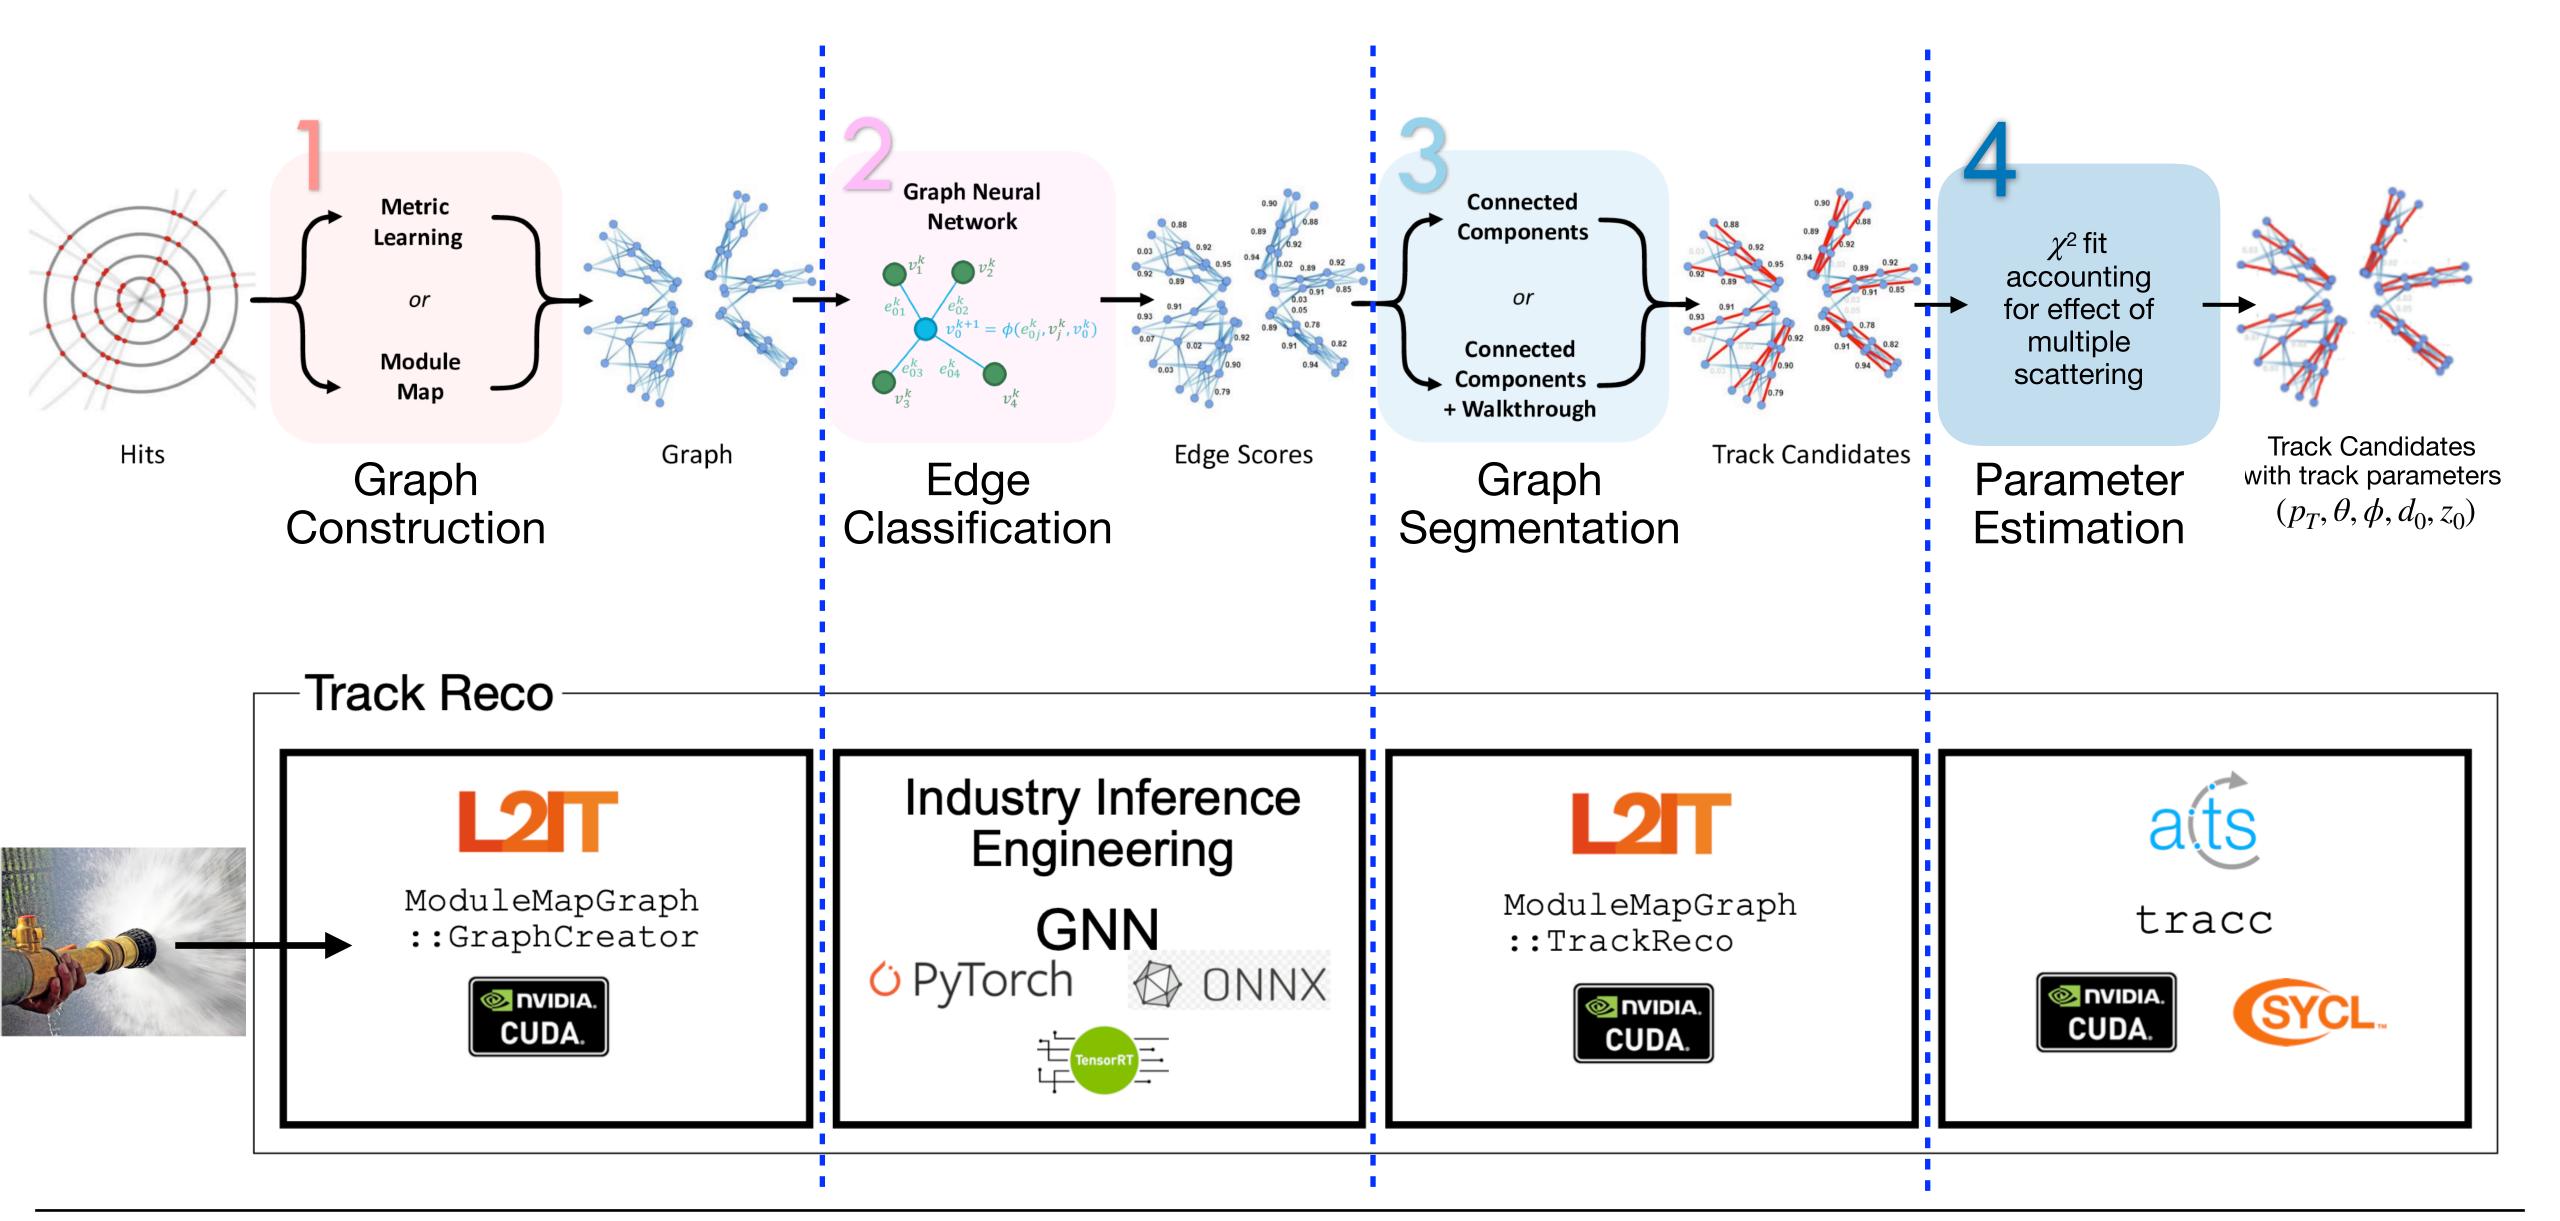


ATLAS input to 2025 update of European strategy for particle physics: <u>ATL-SOFT-PUB-2025-002</u>

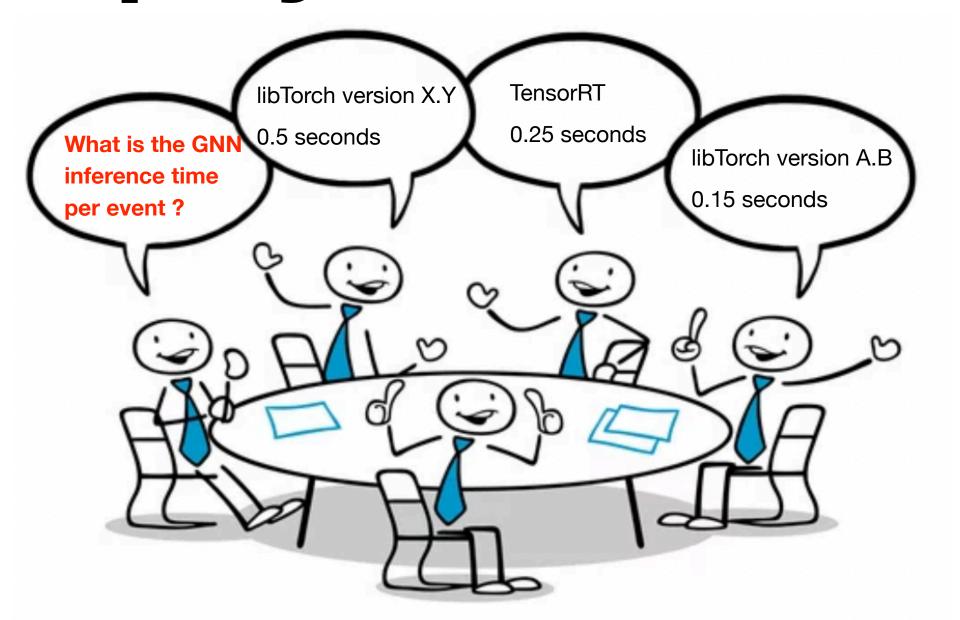
"Now, machine learning is being developed for lower-level tasks [...]

These developments have a very high threshold to adoption: they must improve upon the well-understood and extremely precise methods that have been used in ATLAS searches and measurements to date, not introduce any features [...]"

# Towards deployment



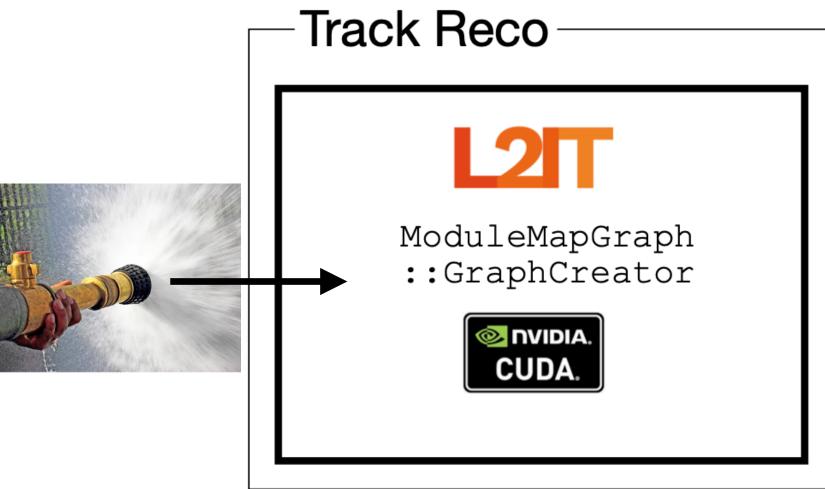
## Towards deployment

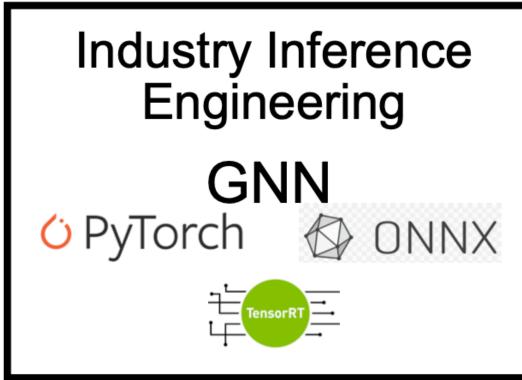


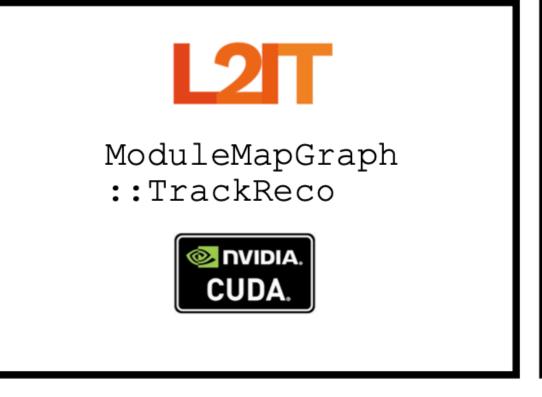
In contrast:
have well-understood
benchmarks for our
CPU-based workloads



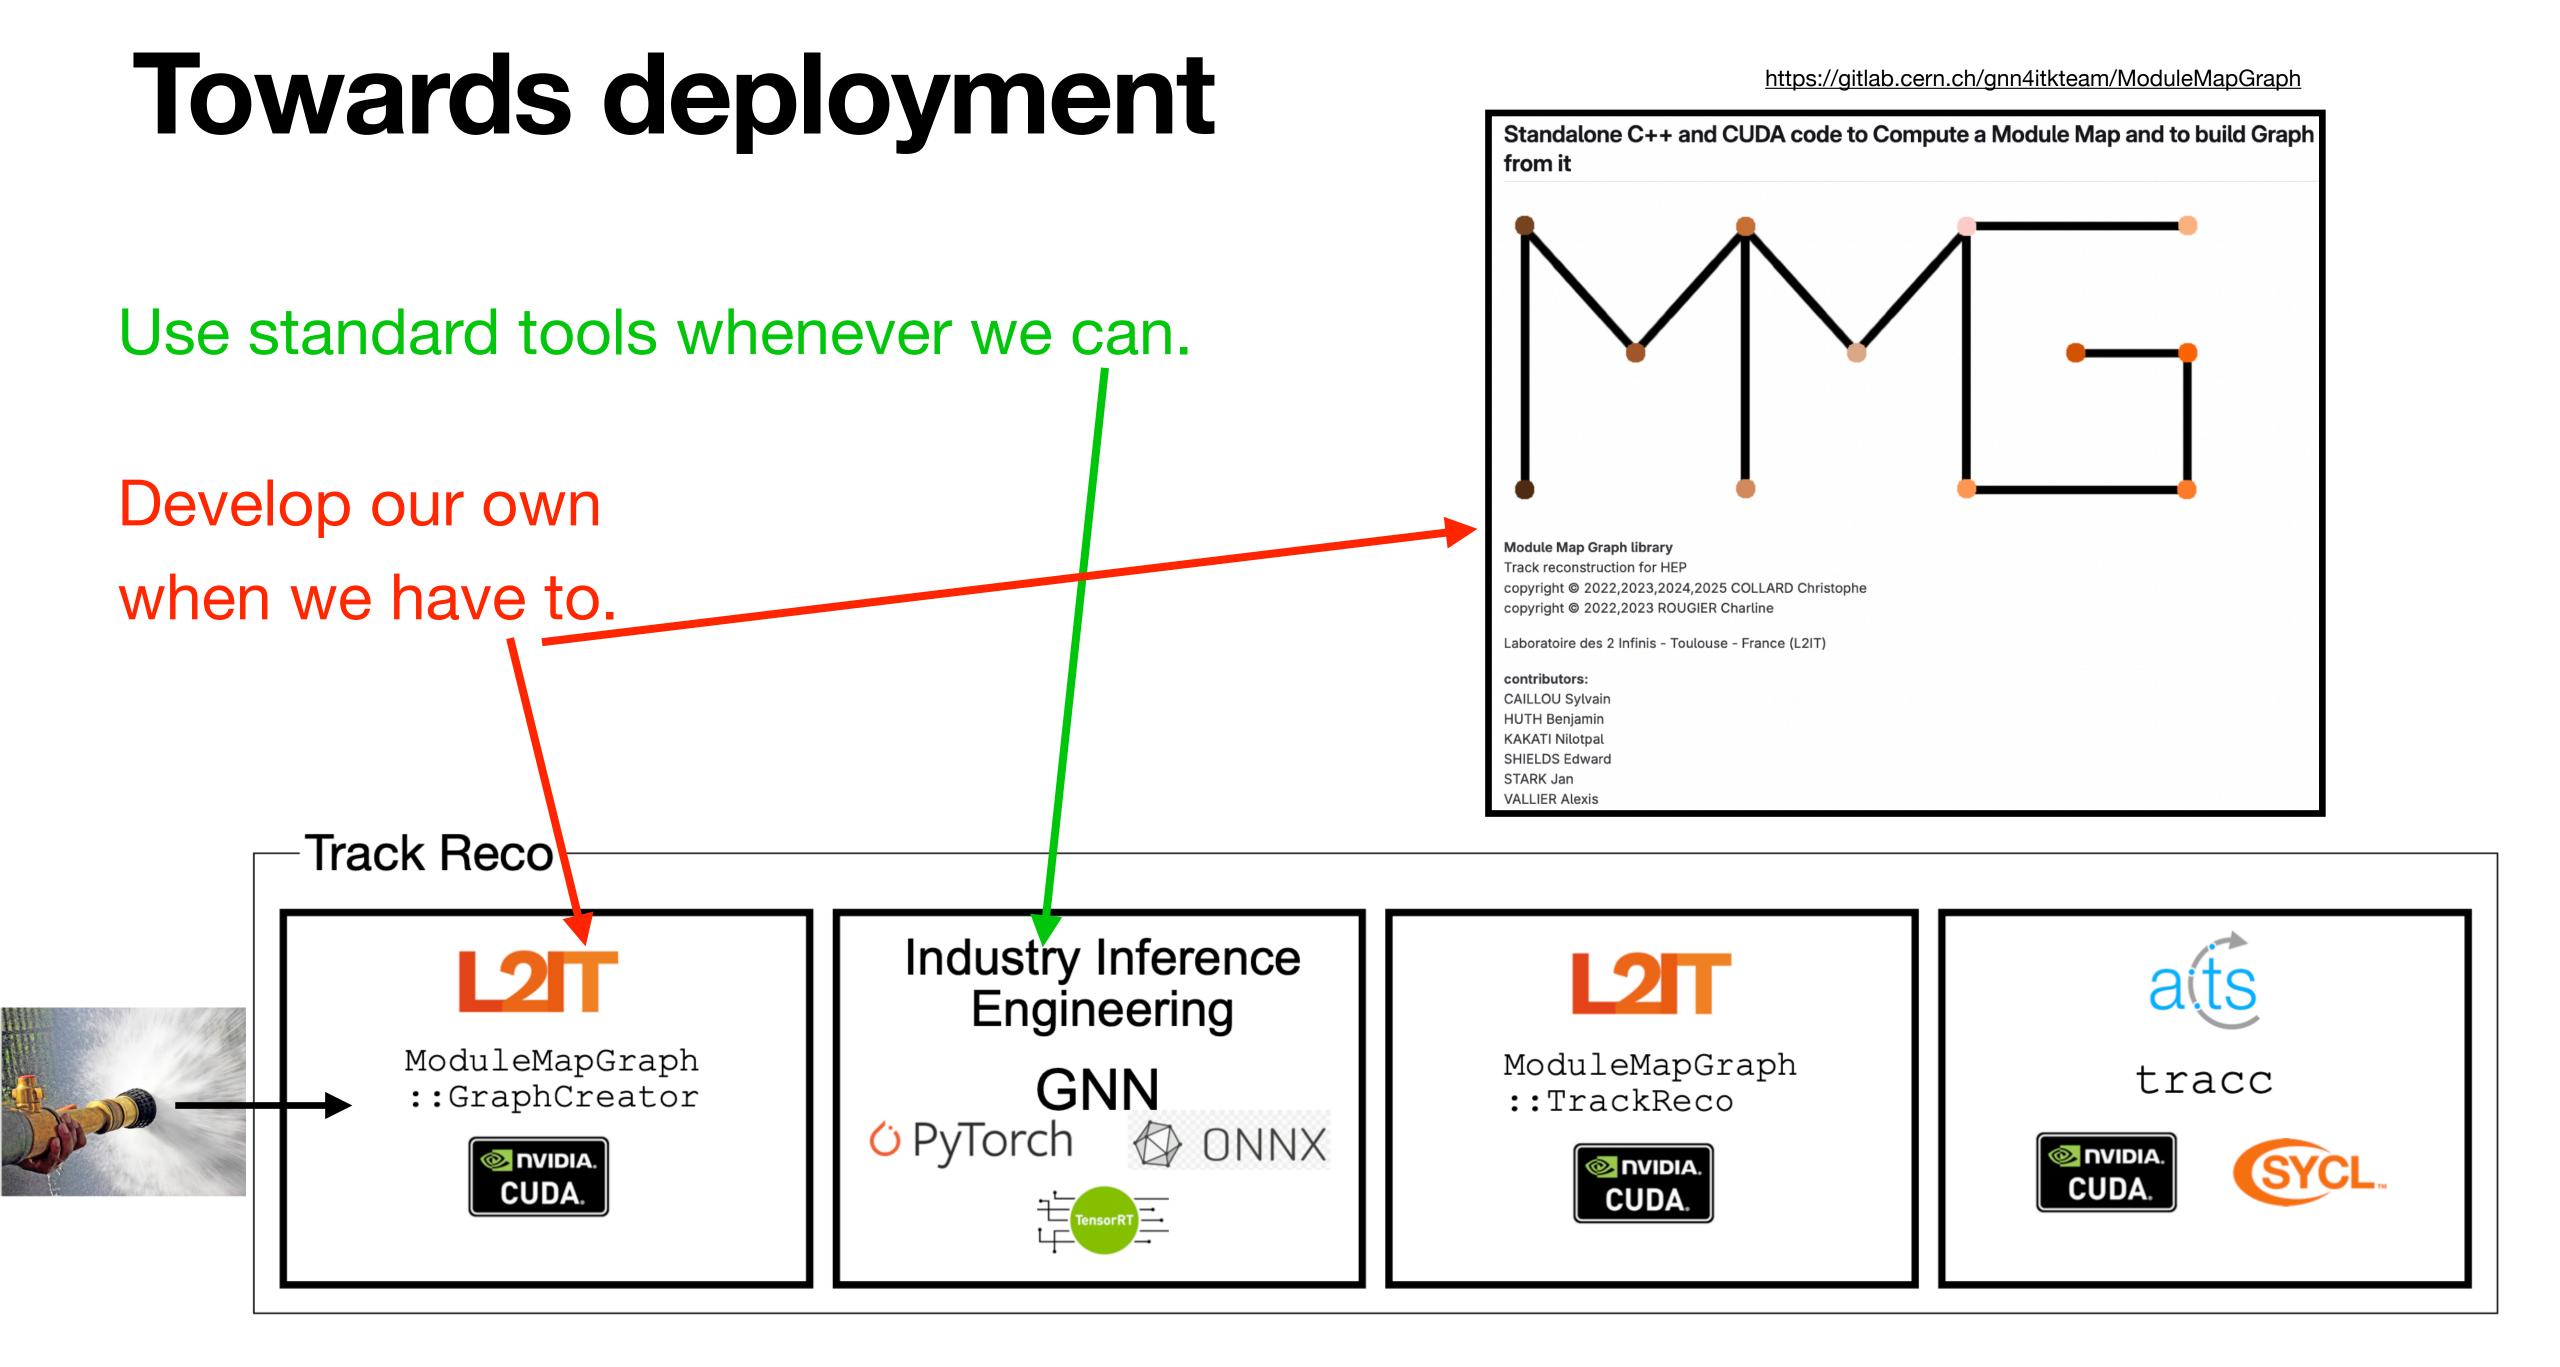
Benchmark Results							
CPU	HS06	Clock speed (MHz)	L2+L3 cache size (grand total, KB)	Cores (runs)			
Intel Xeon E5-2660v3	488	2600	5120+51200	40 HT on			
Intel Xeon E5-4669v4	1836	2200	22528+225280	176 (HT on)			
Intel Xeon E5-2699v4	987	2200	11264+112640	88 (HT on)			
Intel Xeon E5-2620v4	305	2100	4096+40960	32			
Intel Xeon Gold 6130	577	2100	32768+45056	32 (HT off)			











### Conclusions/observations

Reviewed progress in one of the ML-based charged particle tracking algorithms from its start in 2017 until now. Now moving to deployment and production.

Must not forget: we are drinking data out of a firehose! Implementation must run fast.

Fast inference is also important for industry and other other fields.

Benefit a lot from fast inference engines from industry and academia that do the heavy lifting for us.

- would be nice to have standard benchmarks (that test both hardware and inference software),
   à la HS06 or HS23
- ideally these also need to run affordably on CPUs

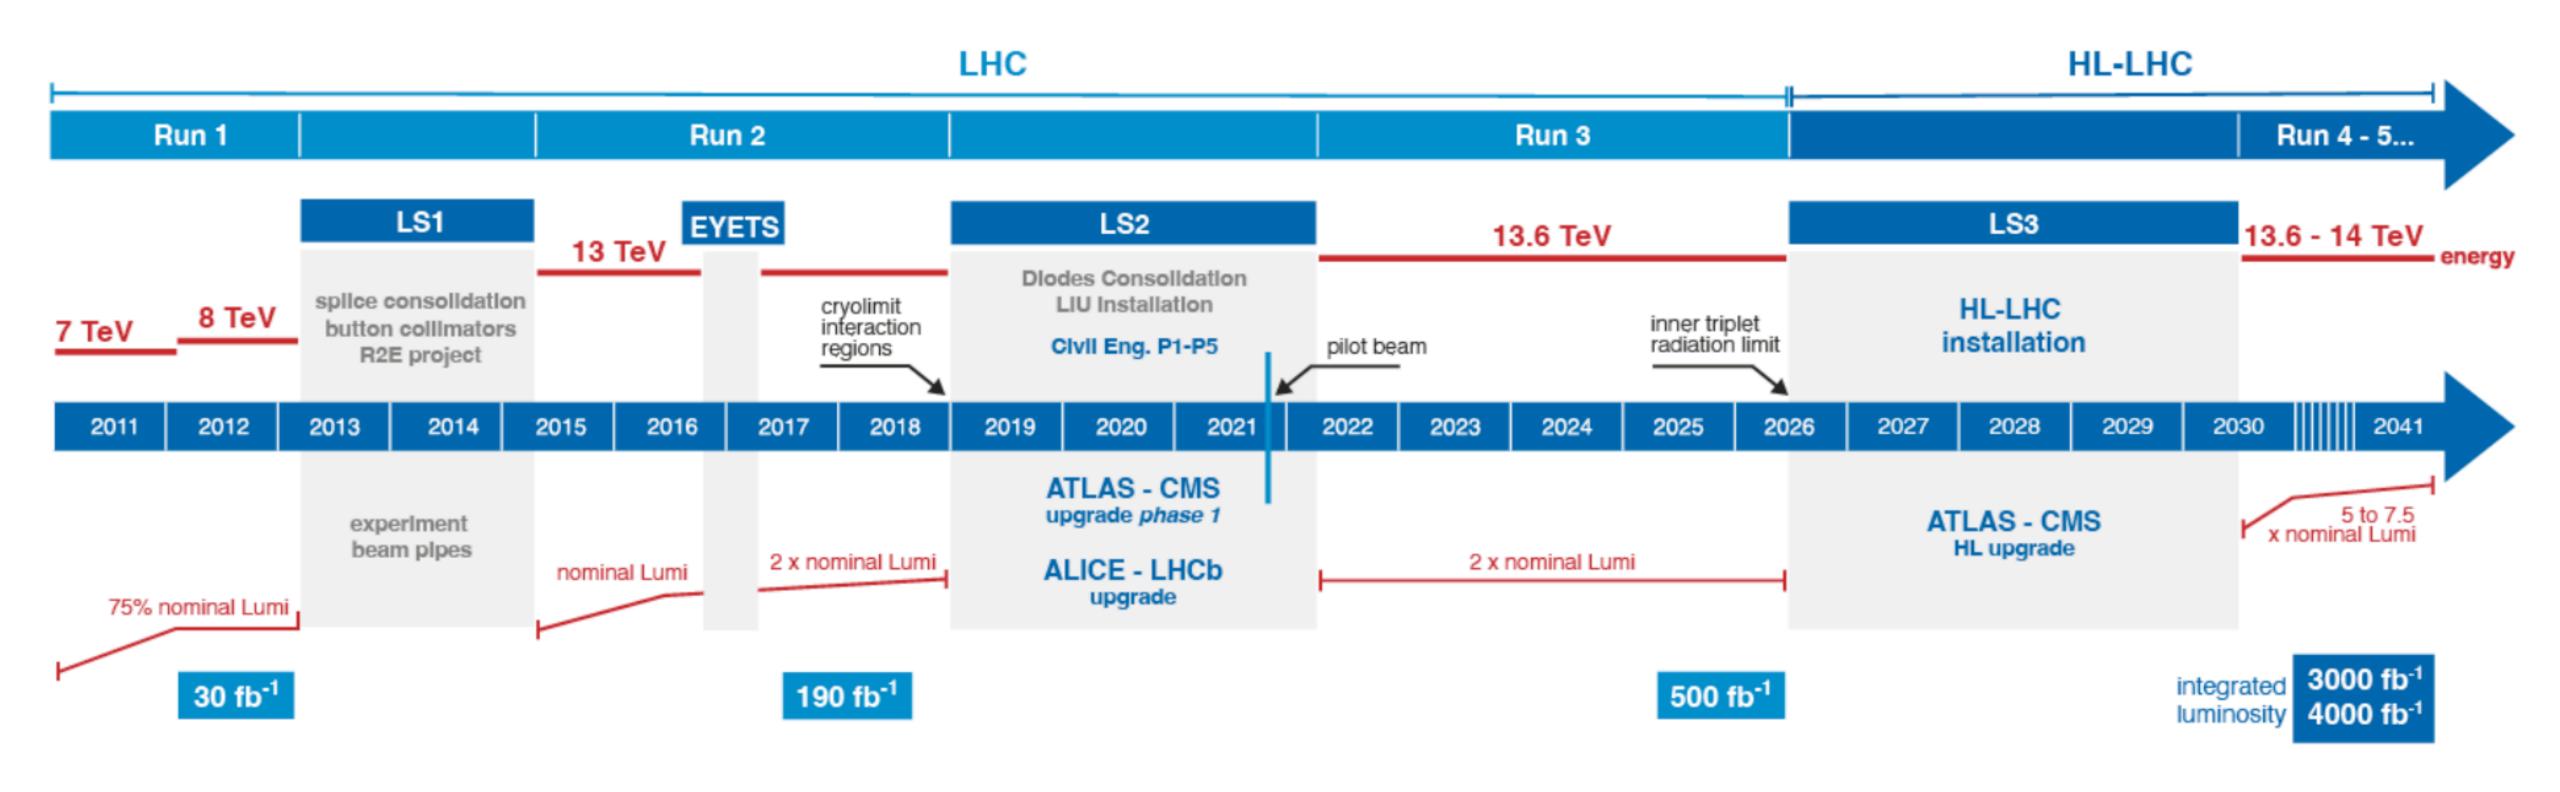
Need very little domain-specific, dedicated GPU code, but it needs to be efficient.

- E.g. for "data preparation": graph creation for use with GNNs. Presented MMG package.
- Train more people that can design these codes?

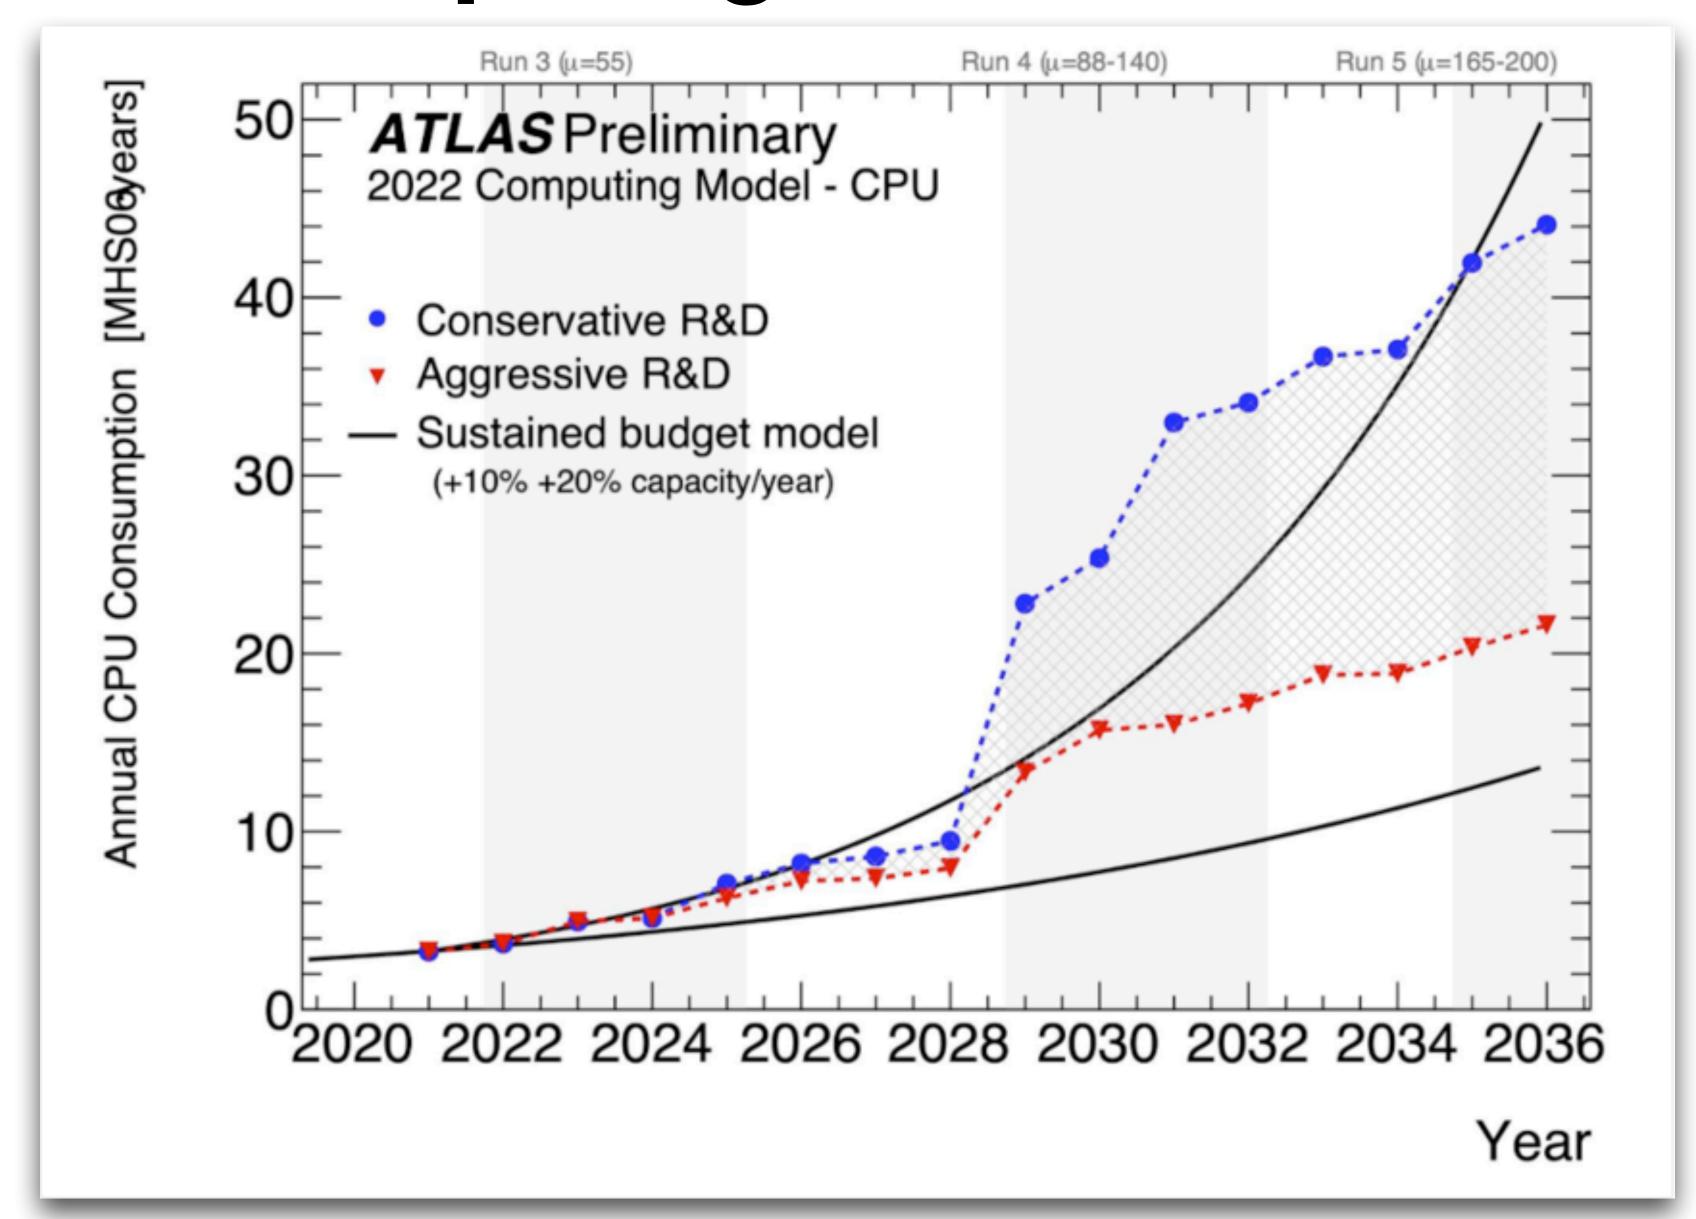
Do not hesitate to get in touch with us; eager to hear experience/needs from others.

# Backup material

## LHC and HL-LHC schedule



# Projected computing needs



#### "10 years to prepare ourselves" for HL-LHC (statement from 2017)

arXiv.org > physics > arXiv:1712.06982

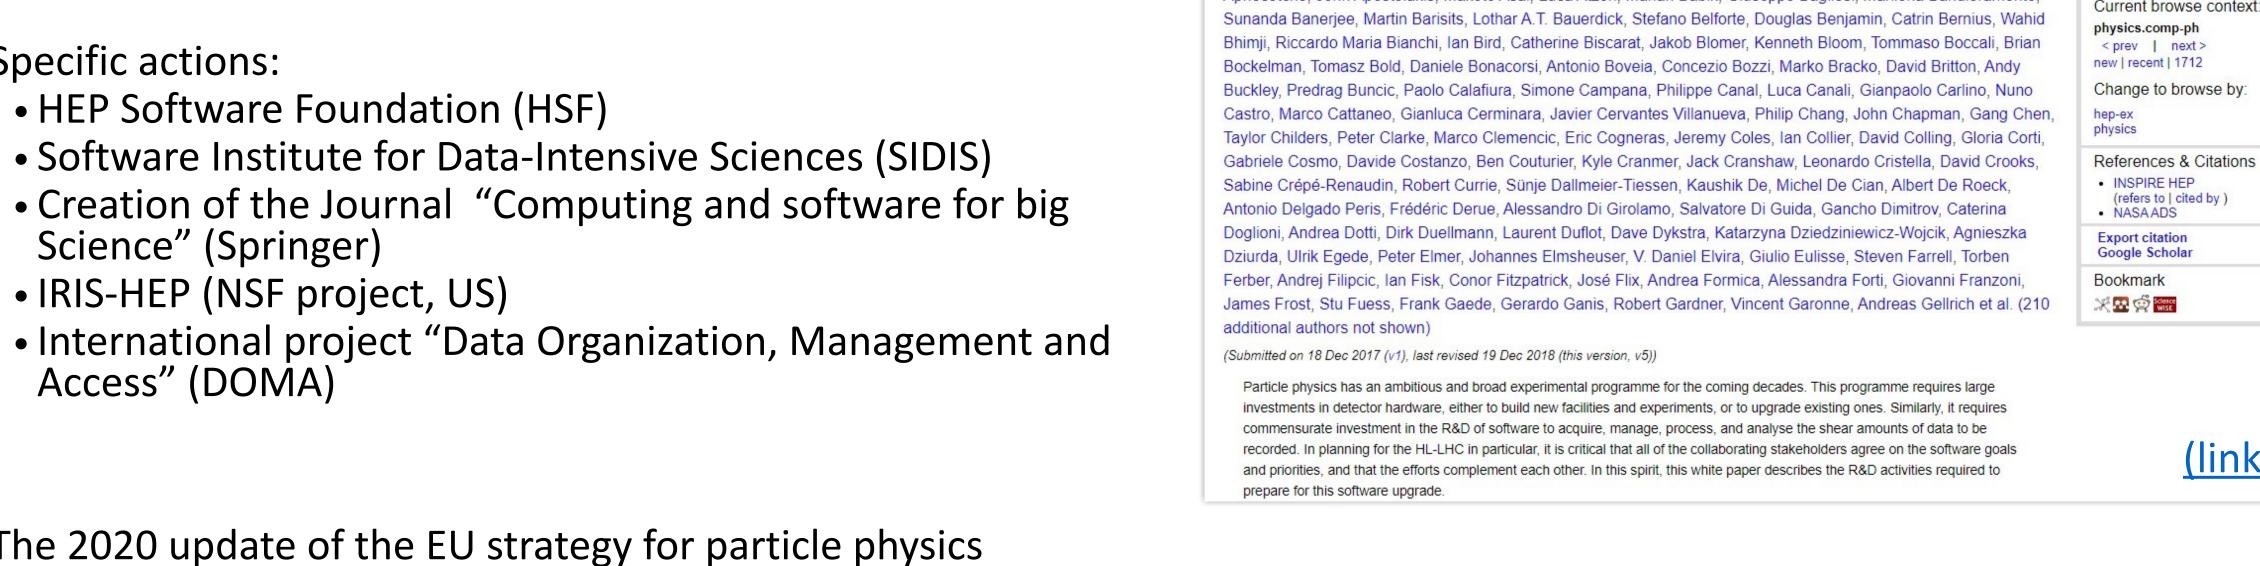
Physics > Computational Physics

A Roadmap for HEP Software and Computing R&D for the 2020s

Johannes Albrecht, Antonio Augusto Alves Jr, Guilherme Amadio, Giuseppe Andronico, Nguyen Anh-Ky, Laurent Aphecetche, John Apostolakis, Makoto Asai, Luca Atzori, Marian Babik, Giuseppe Bagliesi, Marilena Bandieramonte,

- Community white paper (2017)
  - Algorithms, infrastructure, data access...
- Specific actions:

The 2020 update of the EU strategy for particle physics





Large-scale data-intensive software and computing infrastructures are an essential ingredient to particle physics research programmes. The community faces major challenges in this area, notably with a view to the HL-LHC. As a result, the software and computing models used in particle physics research must evolve to meet the future needs of the field. The community must vigorously pursue common, coordinated R&D efforts in collaboration with other fields of science and industry, to develop software and computing infrastructures that exploit recent advances in information technology and data science. Further development of internal policies on open data and data preservation should be encouraged, and an adequate level of resources invested in their implementation. (link) Help I Advanced Search

Download:

Other formats

(link)

PDF

## Machine Learning for track pattern recognition?



Challenge on Kaggle platform (in 2018): (link)

Article in proceedings of CHEP 2018: (link)



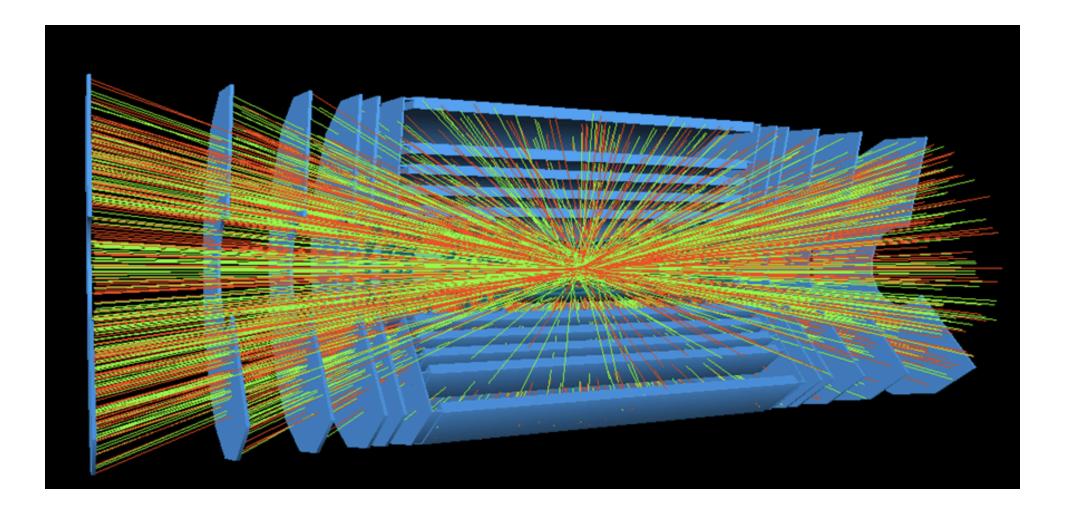
622 \* 415 pixels

a large fraction carries information about the person



Can't use the same tools

How to present tracking data to a neural network?



ATLAS tracker for HL-LHC:

5 \* 109 readout channels

~3 \* 10<sup>5</sup> 3D space-points per event

=> data are *sparse* 

# ACORN tracking software

