TrackML Report

David Rousseau

LAL-Orsay

rousseau@lal.in2p3.fr @dhpmrou

IML workhop, 16th April 2019

Who and How

- Organisation: Jean-Roch Vlimant (Caltech), Vincenzo Innocente, Andreas Salzburger (CERN), Sabrina Amrouche, Tobias Golling, Moritz Kiehn (Geneva University), David Rousseau, Yetkin Yilmaz (LAL-Orsay), Paolo Calafiura, Steven Farrell, Heather Gray (LBNL), Vladimir Vava Gligorov (LPNHE-Paris), Laurent Basara, Cécile Germain, Isabelle Guyon, Victor Estrade (LRI-Orsay), Edward Moyse (University of Massachussets), Mikhail Hushchyn, Andrey Ustyuzhanin (Yandex, HSE)

Platforms:

kaggle Phase 1

CodaLab Phase 2

accuracy

throughput

TrackML sponsors

CERN

kaggle

NVIDIA

UNIVERSITÉ DE GENÈVE

Paris-Saclay Center for Data Science

INSTITUTE DATAIA
Data Science, Intelligence & Society

CHALEARN

iris hep
Institute for Research & Innovation in Software for High Energy Physics

CERN openlab

Inria
INVENTORS FOR THE DIGITAL WORLD

PyTorch

cnrs

erc

Fermilab

COMMON GROUND

inqnet
Alliance for Quantum Technologies

UNIVERSITÉ PARIS SUD

université PARIS-SACLAY

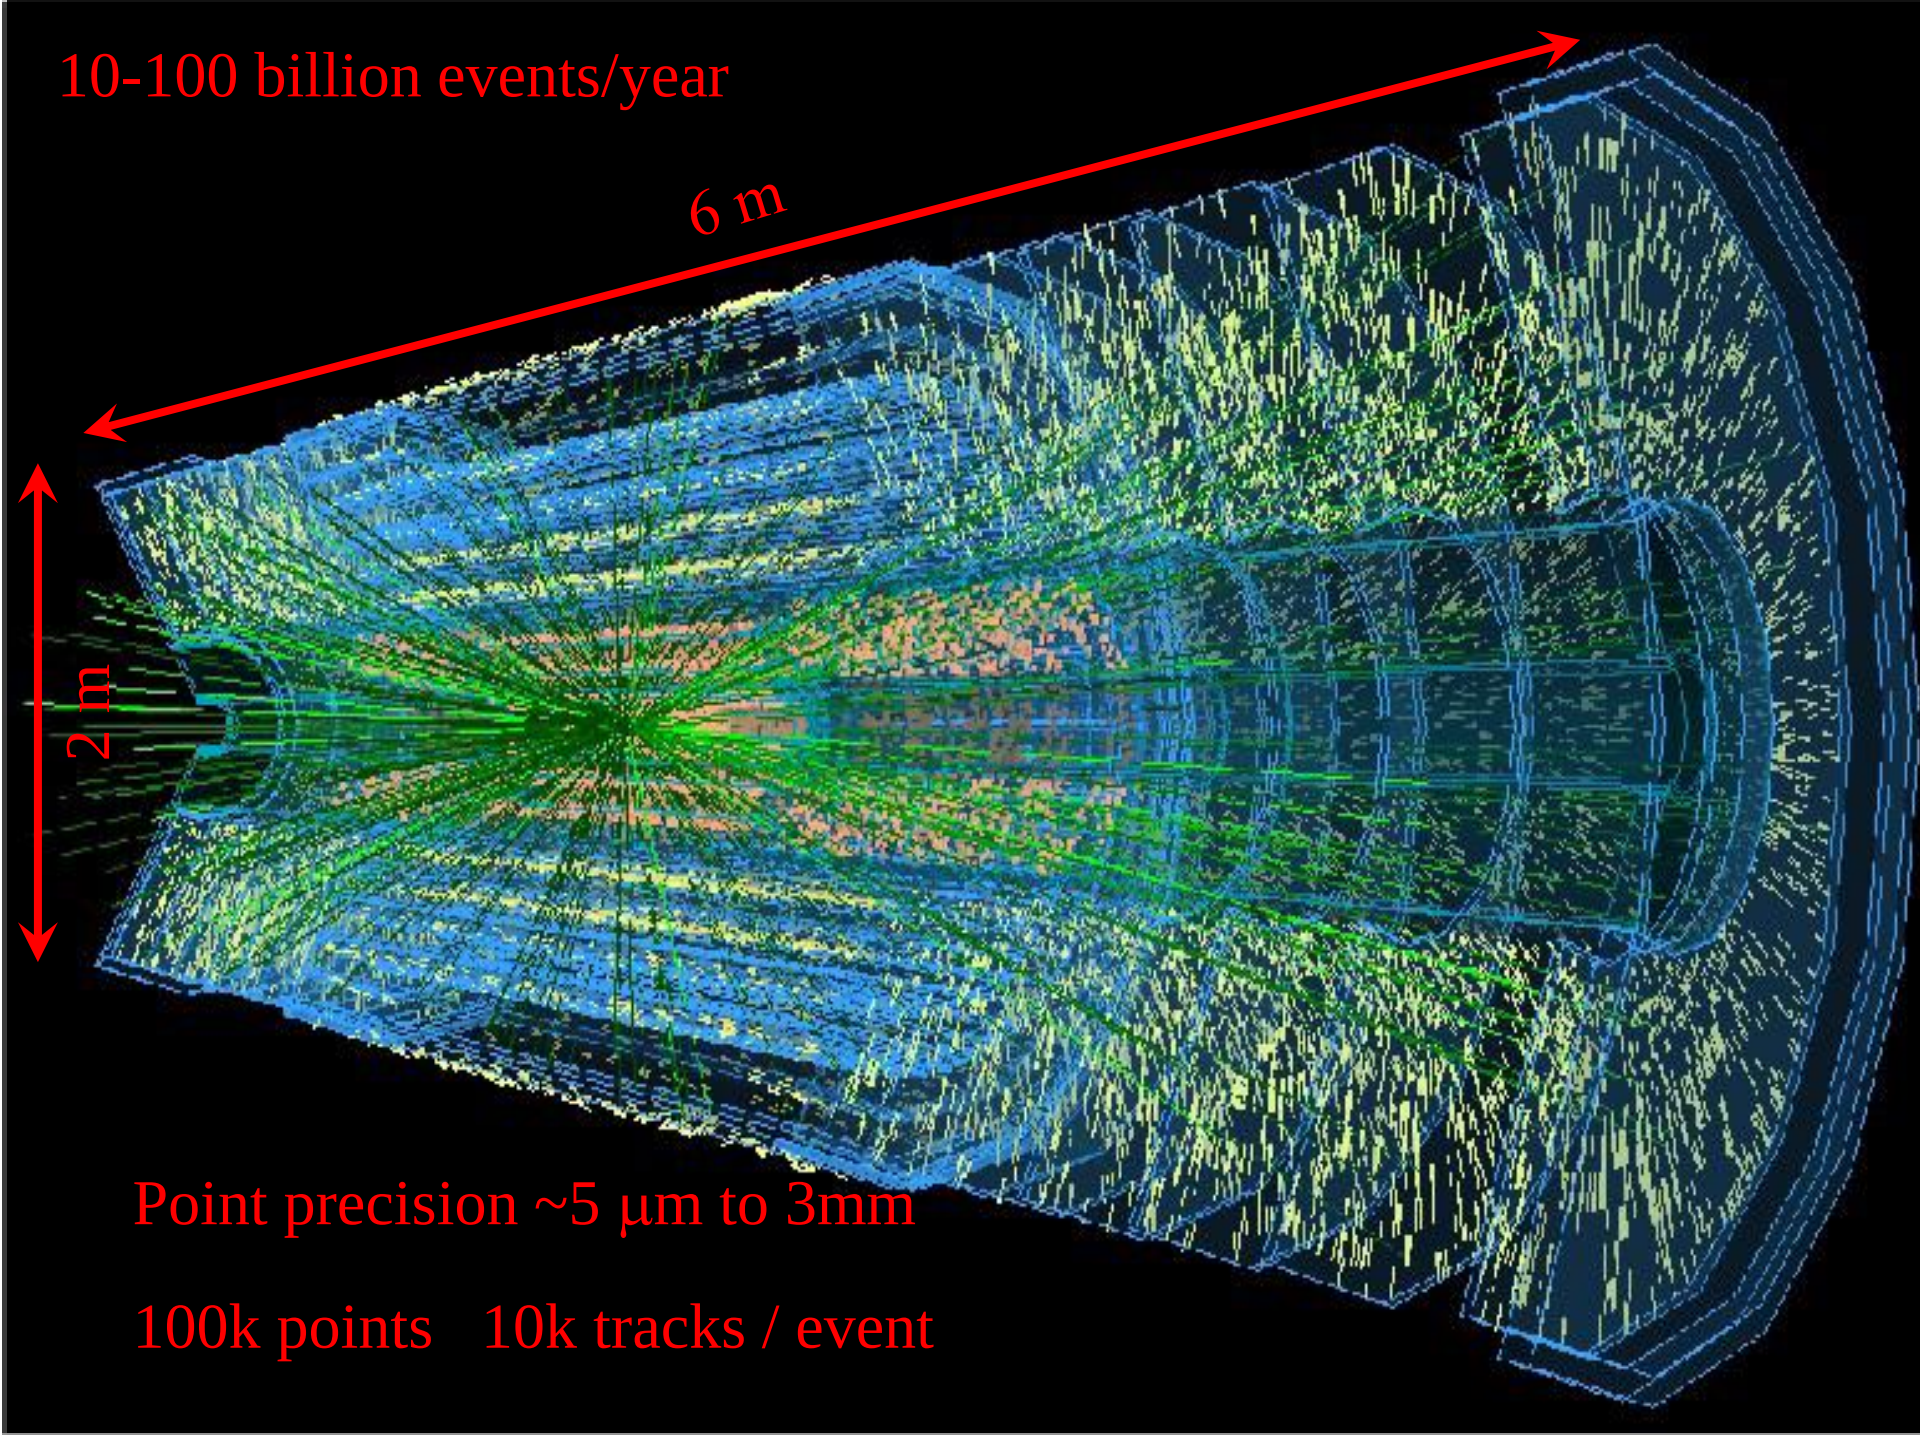
10-100 billion events/year
6 m
2 m
Point precision ~5 μm to 3mm
100k points 10k tracks / event

Tracking crisis

- Tracking (in particular pattern recognition) dominates reconstruction CPU time at LHC
- HL-LHC (phase 2) perspective : increased pileup :Run 1 (2012): <>~20, Run 2 (2015): <>~50,Phase 2 (2025): <>~200
- CPU time quadratic/exponential extrapolation (difficult to quote any number)
- Large effort within HEP to optimise software and tackle micro and macro parallelism. Sufficient gains for Run 2 but still a long way for HL-LHC.
- >20 years of LHC tracking development. Everything has been tried?
  - Maybe yes, but maybe algorithm slower at low lumi but with a better scaling have been dismissed ?
  - Maybe no, brand new ideas from ML
- ➔challenge TrackML !!!

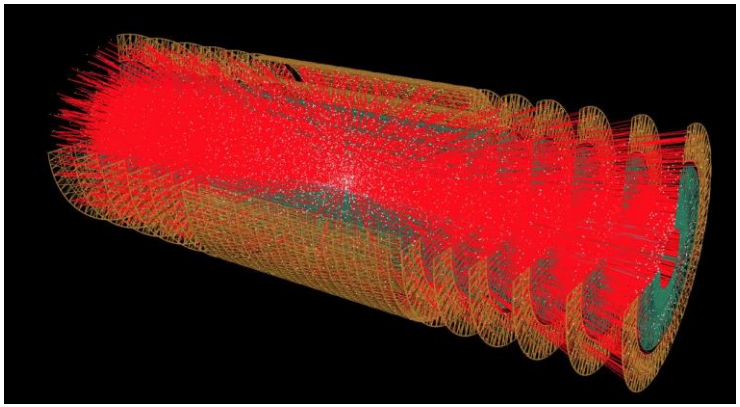

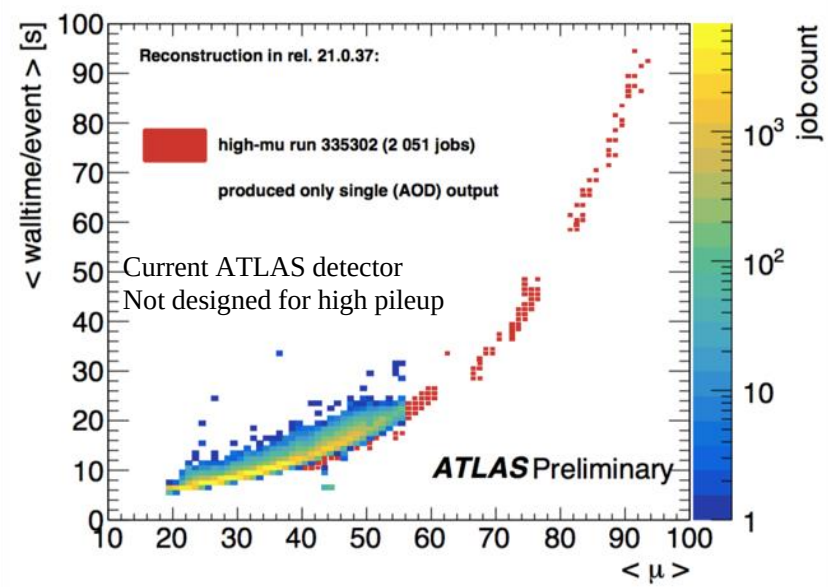

200

There are newer plots, same message

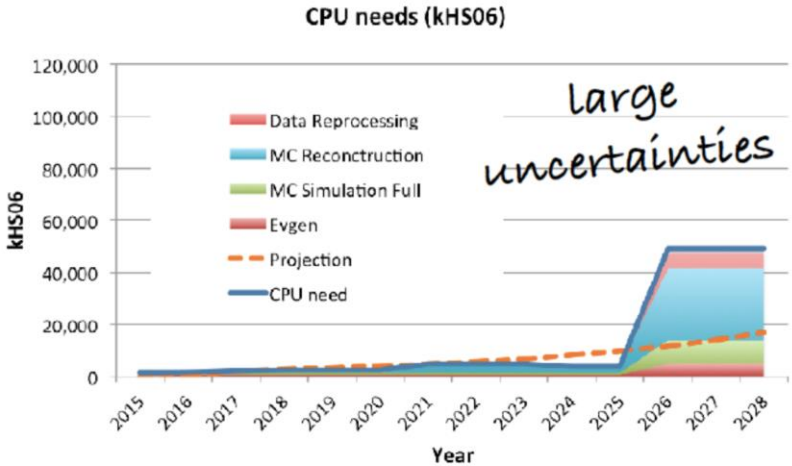

An early attempt

Computer Physics Communications 49 (1988) 429–448
North-Holland, Amsterdam

NEURAL NETWORKS AND CELLULAR AUTOMATA IN EXPERIMENTAL HIGH ENERGY PHYSICS

B. DENBY

Laboratoire de l'Accélérateur Linéaire, Orsay, France

Received 20 September 1987; in revised form 28 December 1987

Bruce Denby

- 1987 Very first Neural Net in HEP paper known
- NN for tracking and calo clustering
- B. Denby then moved from Delphi at LEP to CDF at Tevatron. He still active outside HEP: 2017 analysis of ultrasonic image of the tongue
- 1992 JetNet Carsten Peterson, Thorsteinn Rognvaldsson (Lund U.) , Leif Lonnblad (CERN) (~500 citations) really started NN use in HEP

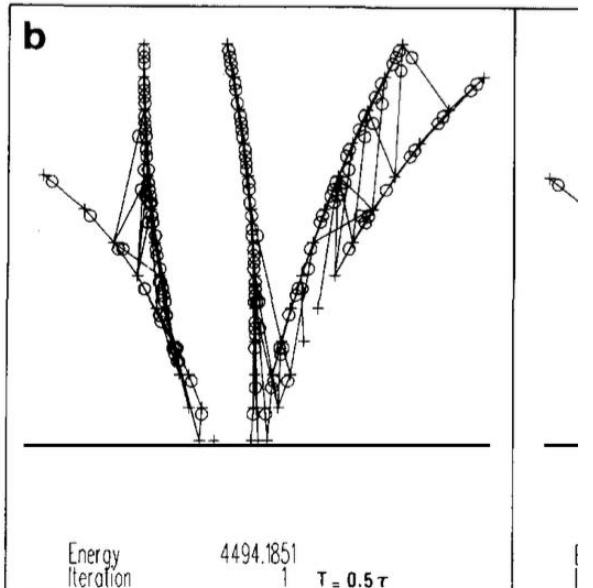

A recent success with ν : NOVA

arXiv 1604.01444 Aurisano et al

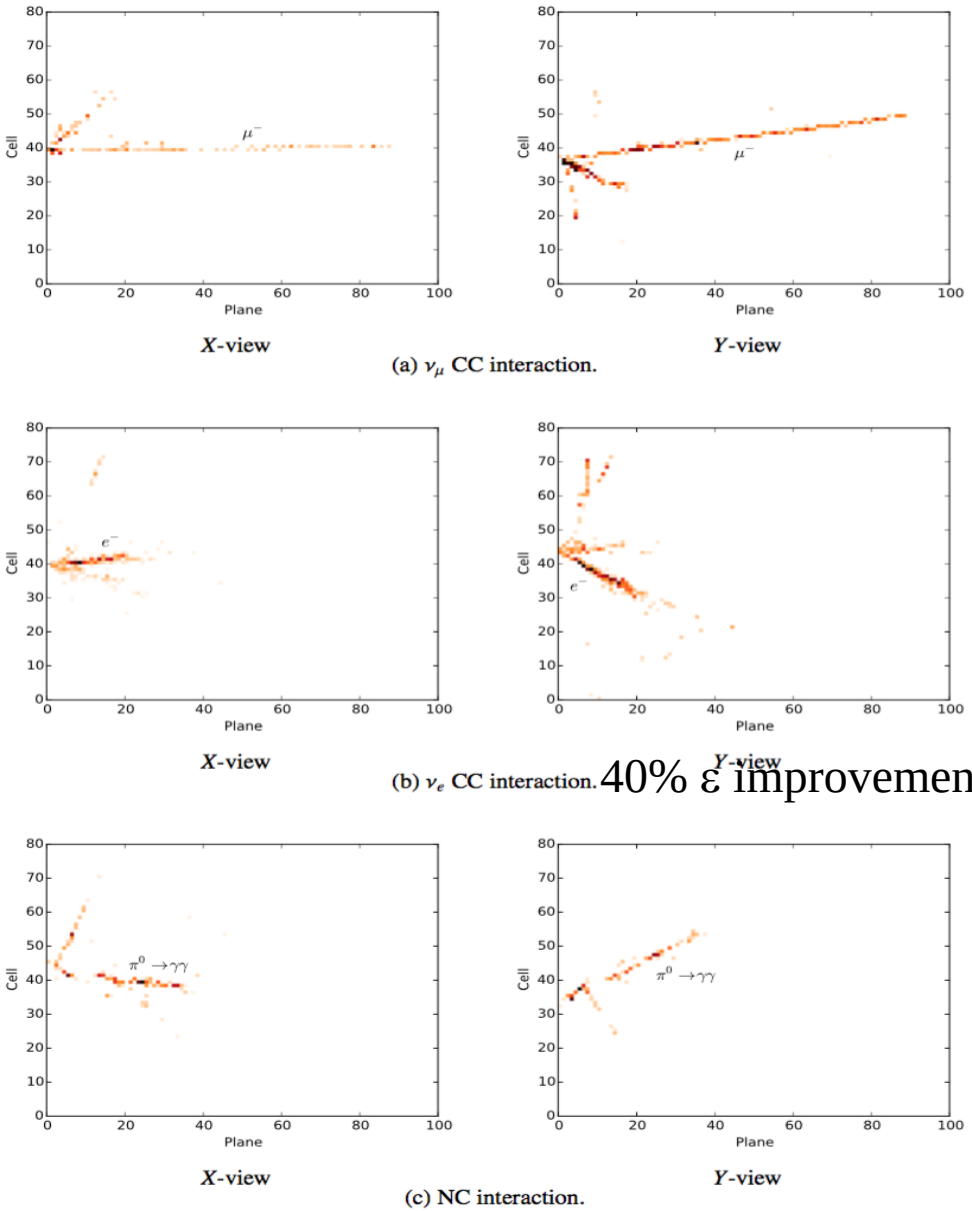

(a) ν_μ CC interaction.

(b) ν_e CC interaction.

(c) NC interaction.

40% ε improvement

Neutrino interaction classification
Using Convolutionnal Neural Network (GoogleNet)
Actually used in physics results 1703.03328 and 1706.04592

Cosmic ray in LArTPC (1)

1903.05663

Simulated data

Simulated signal

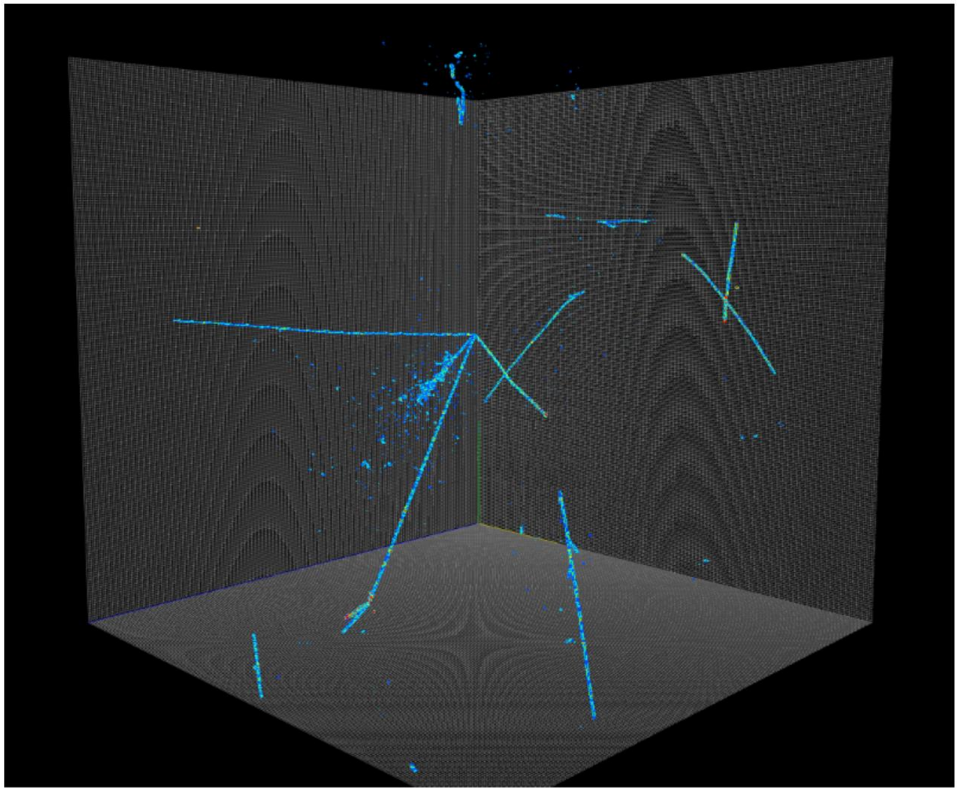

Truth label

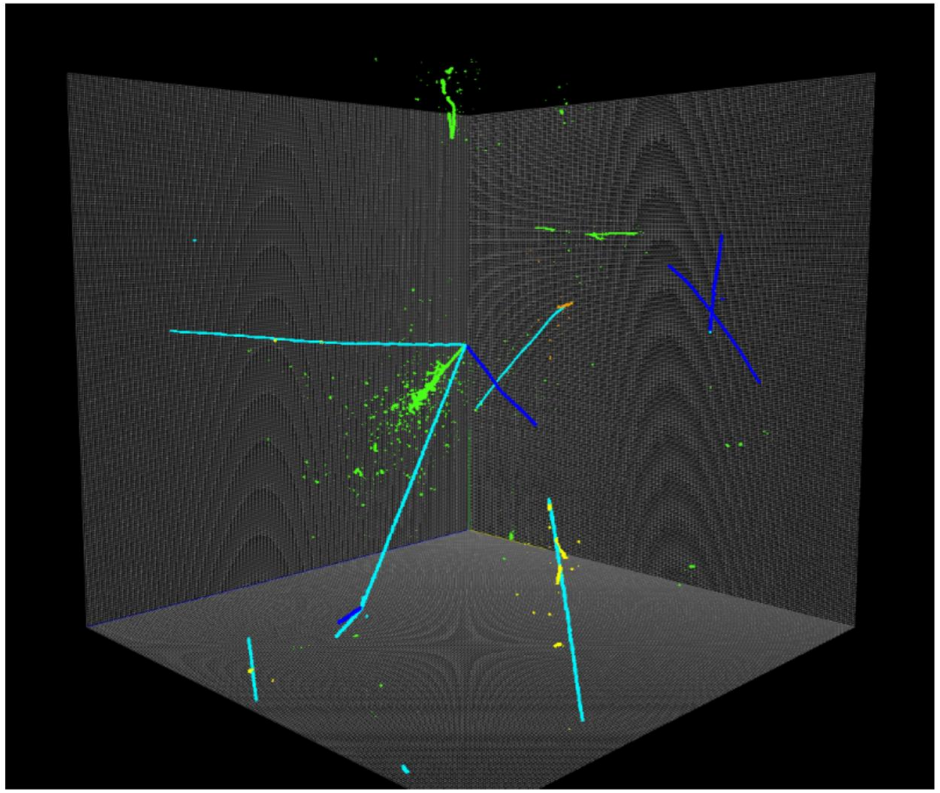

Full 3D information, sparse (non zero voxels <1%)

Cosmic ray (2)

U-Resnet for semantic segmentation using Submanifold Sparse Convolutional Network

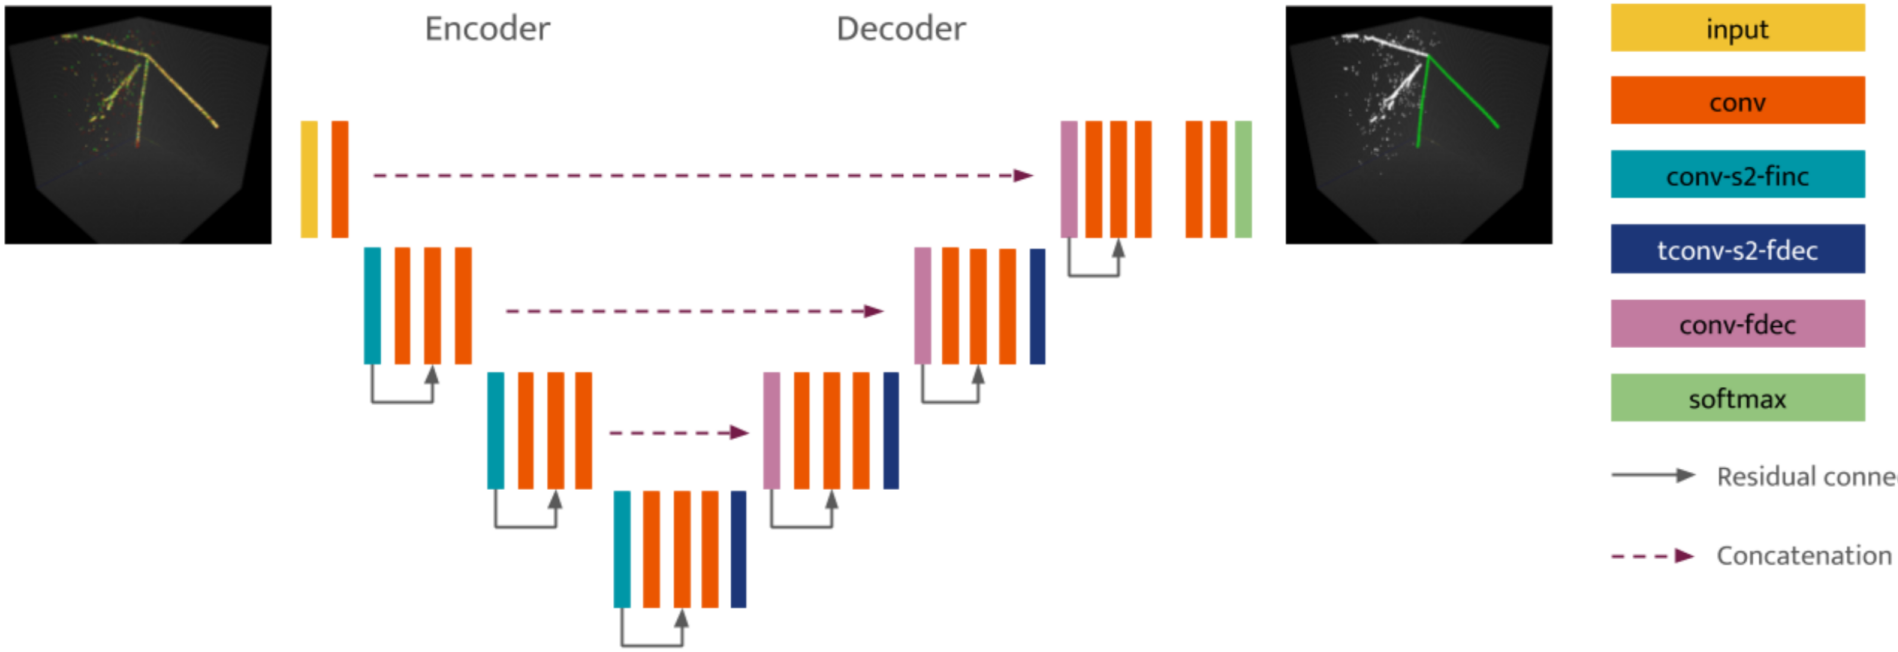

100 times less memory, 10 times faster than normal CNN, high accuracy

From domain to challenge and back

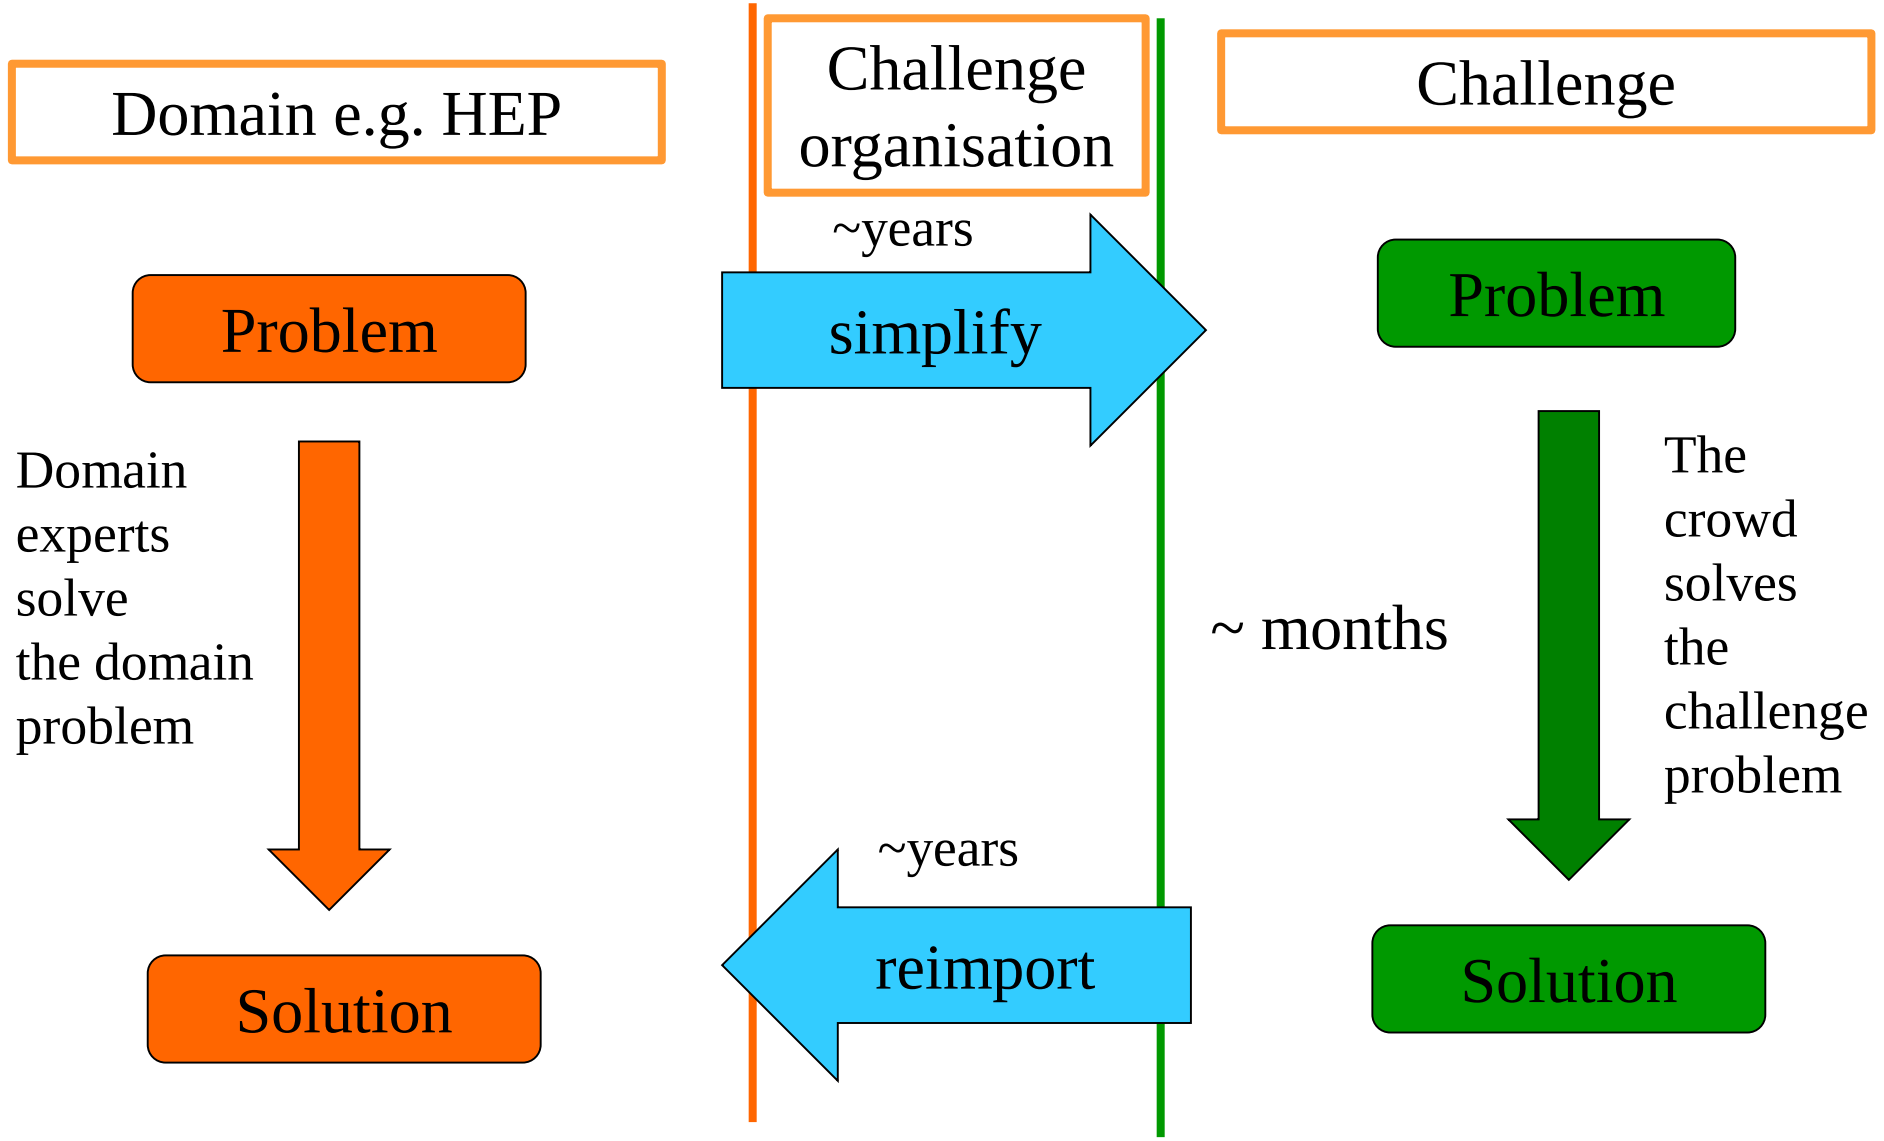

TrackML in a nutshell

- Accurate simulation engine (ACTS https://gitlab.cern.ch/acts/acts-core) to produce realistic events
  - One file with list of 3D points
  - Ground truth : one file with point to particle association
  - Ground truth auxiliary : true particle parameter (origin, direction, curvature)
  - Typical events with ~200 parasitic collisions (~10.000 tracks/event)
- Large training sample 10k events, 0.1 billion tracks, 1 billion points, ~100GByte
- Accuracy phase (May to August 2018) on Kaggle
  - Participants are given the test sample (with usual split for public and private leaderboard) and run the evaluation to find the tracks
  - They should upload the tracks they have found
  - A track is a list of 3D points
  - Score : fraction of points correctly grouped together
  - Evaluation on test sample with per-mille precision on 100 event
- Throughput phase Sep to Mar 2019 on Codalab
  - Strong CPU incentive

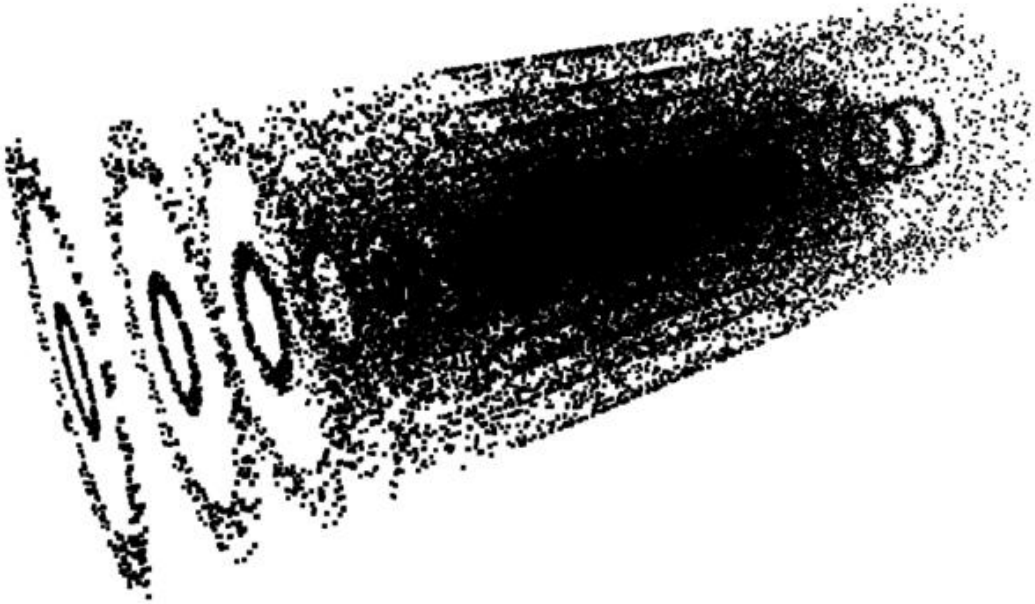

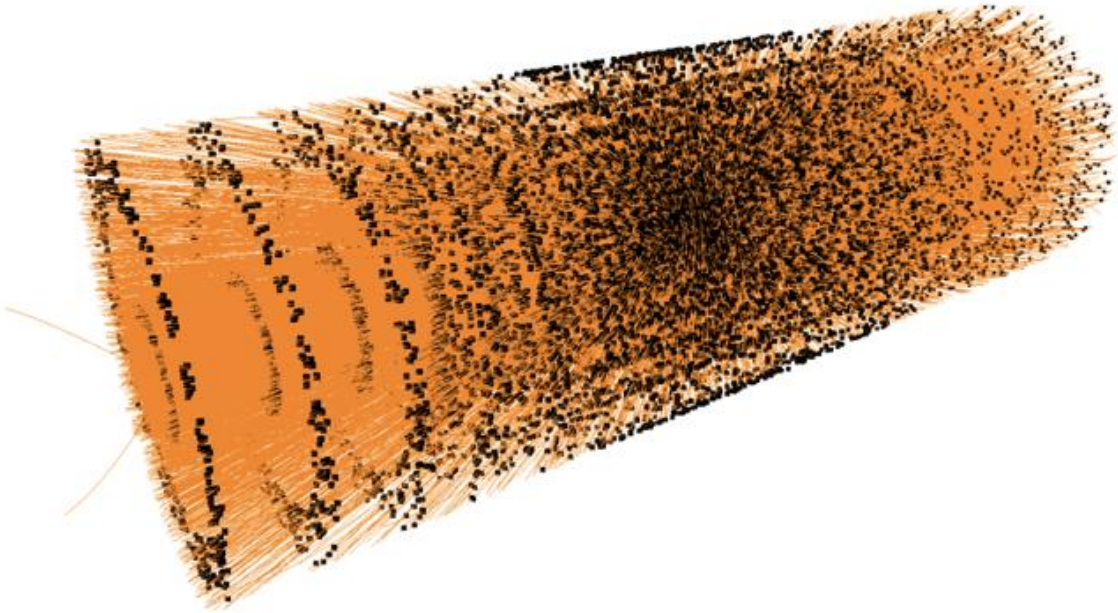

Score

- 2017 CMS tracker Technical Design Report : Chapter 6 expected performance 31 pages 58 figures
- ATLAS Si strip Technical Design Report Chapter 4 ITk Performance and Physics Benchmark Studies 54 pages 80 figures

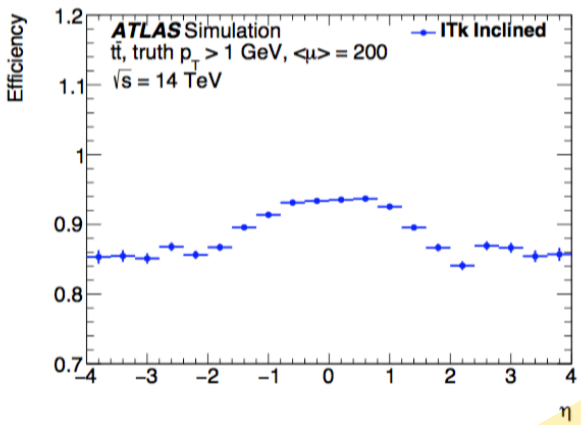

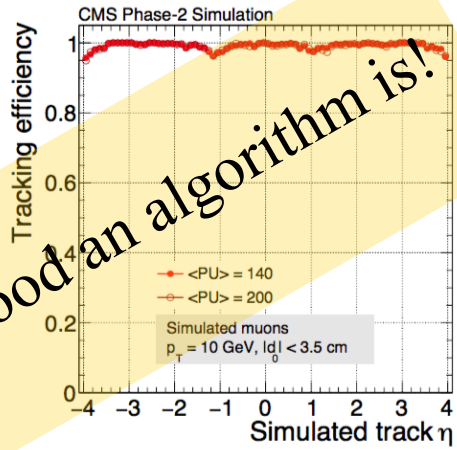

We need 1 number to specify how good an algorithm is!
plus CPU time

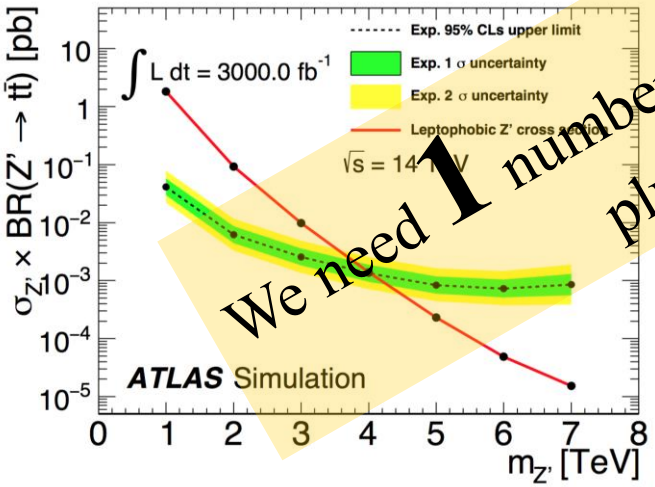

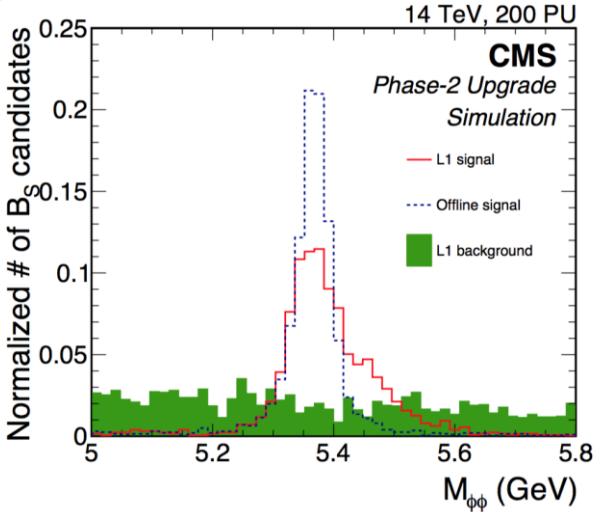

Track evaluation

good track	not so good track
many compatible hits	short tracks
completeness	holes
uniqueness	shared hits
low χ^2/ndf	bad fit quality, outliers
small impact parameter (for primaries)	
clusters are compatible	

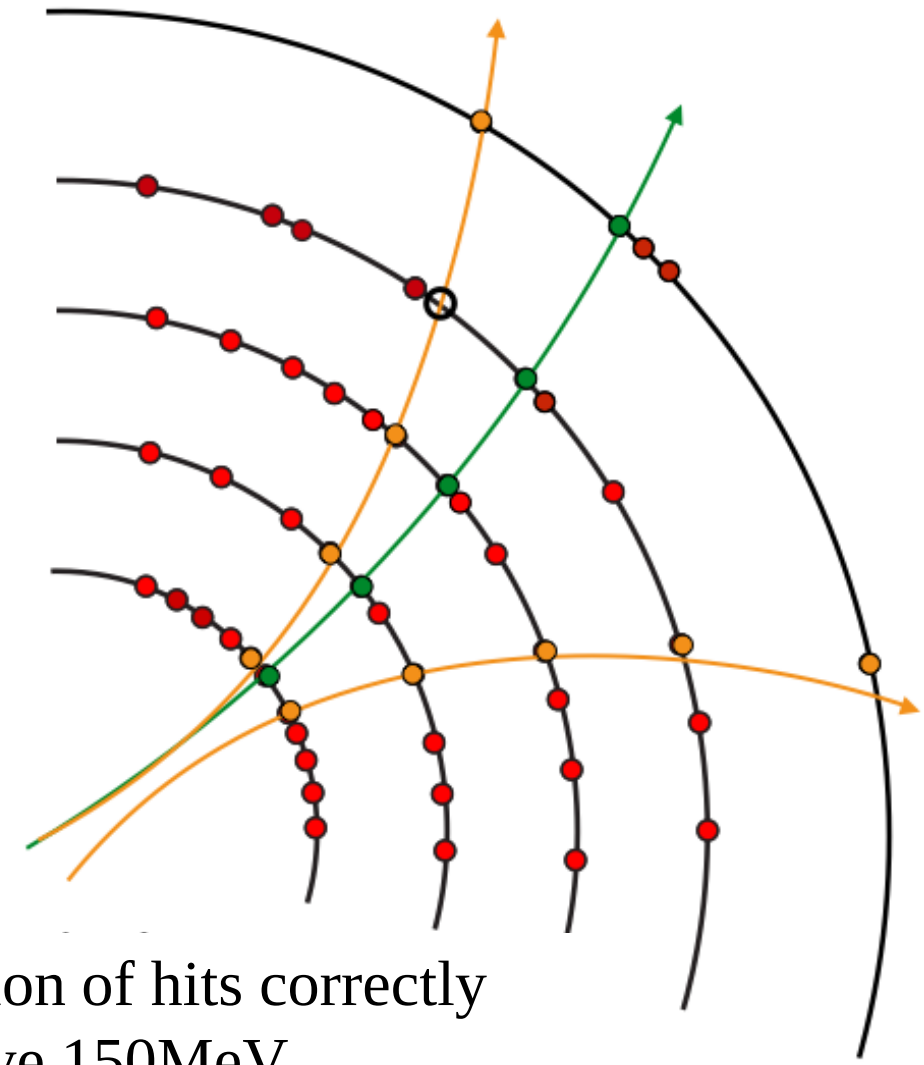

Big decision : score is ~ « the fraction of hits correctly associated ». Include all tracks above 150MeV

Real life vs challenge

1. Wide type of physics events
2. Full detailed Geant 4 / data
3. Detailed dead matter description
4. Complex geometry (tilted modules, double layers, misalignments…)
5. Hit merging
6. Allow shared hits
7. Output is hit clustering, track parameter and covariance matrix
8. Multiple metrics (see TDR's)

1. One event type (ttbar)
2. ACTS (MS, energy loss, hadronic interaction, solenoidal magnetic field, inefficiency)
3. Cylinders and slabs
4. Simple, ideal, geometry (cylinders and disks)
5. No hit merging
6. Disallow shared hits
7. Output is hit clustering
8. Single number metrics

Simpler, but not too simple!

Evolution of leaderboard

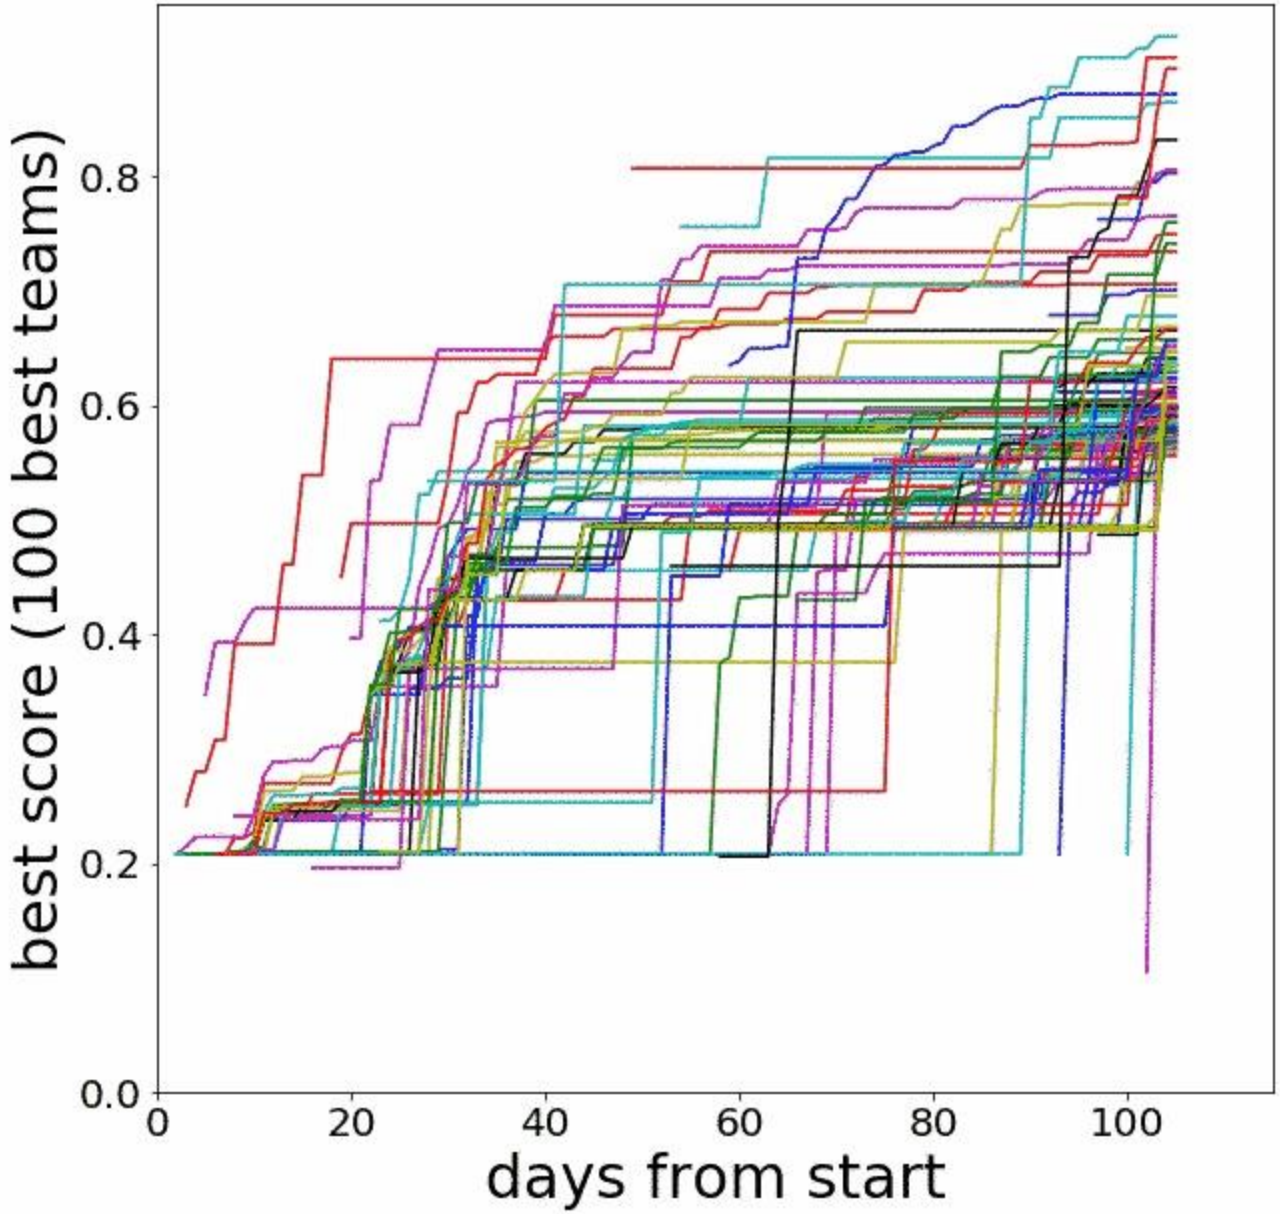

Final Leaderboard

1	—	Top Quarks	0.92182	10	19d
2	—	outrunner	0.90302	9	18d
3	HEP	Sergey Gorbunov	0.89353	6	18d
4	HEP	demelian	0.87079	35	1mo
5	—	Edwin Steiner	0.86395	5	18d
6	—	Komaki	0.83127	22	18d
7	—	Yuval & Trian	0.80414	56	18d
8	—	bestfitting	0.80341	6	18d
9	—	DBSCAN forever	0.80114	23	18d
10	—	Zidmie & KhaVo	0.76320	26	18d
11	—	Andrea Lonza	0.75845	15	18d
12	—	Finnies	0.74827	56	18d
13	—	Rei Matsuzaki	0.74035	12	18d
14	—	Mickey	0.73217	10	2mo
15	—	Vicens Gaitan	0.70429	19	1mo
16	—	Robert	0.69955	3	21d
17	—	Yuval-CPMP tribute band	0.69364	20	20d
18	—	N. Hi. Bouzu	0.67573	9	22d
19	—	Steins;Gate	0.66763	12	19d
20	▲1	Victor Nedel'ko	0.66723	4	2mo

Only public LB to private LB rank change

Efficiency #20 Nedel'ko

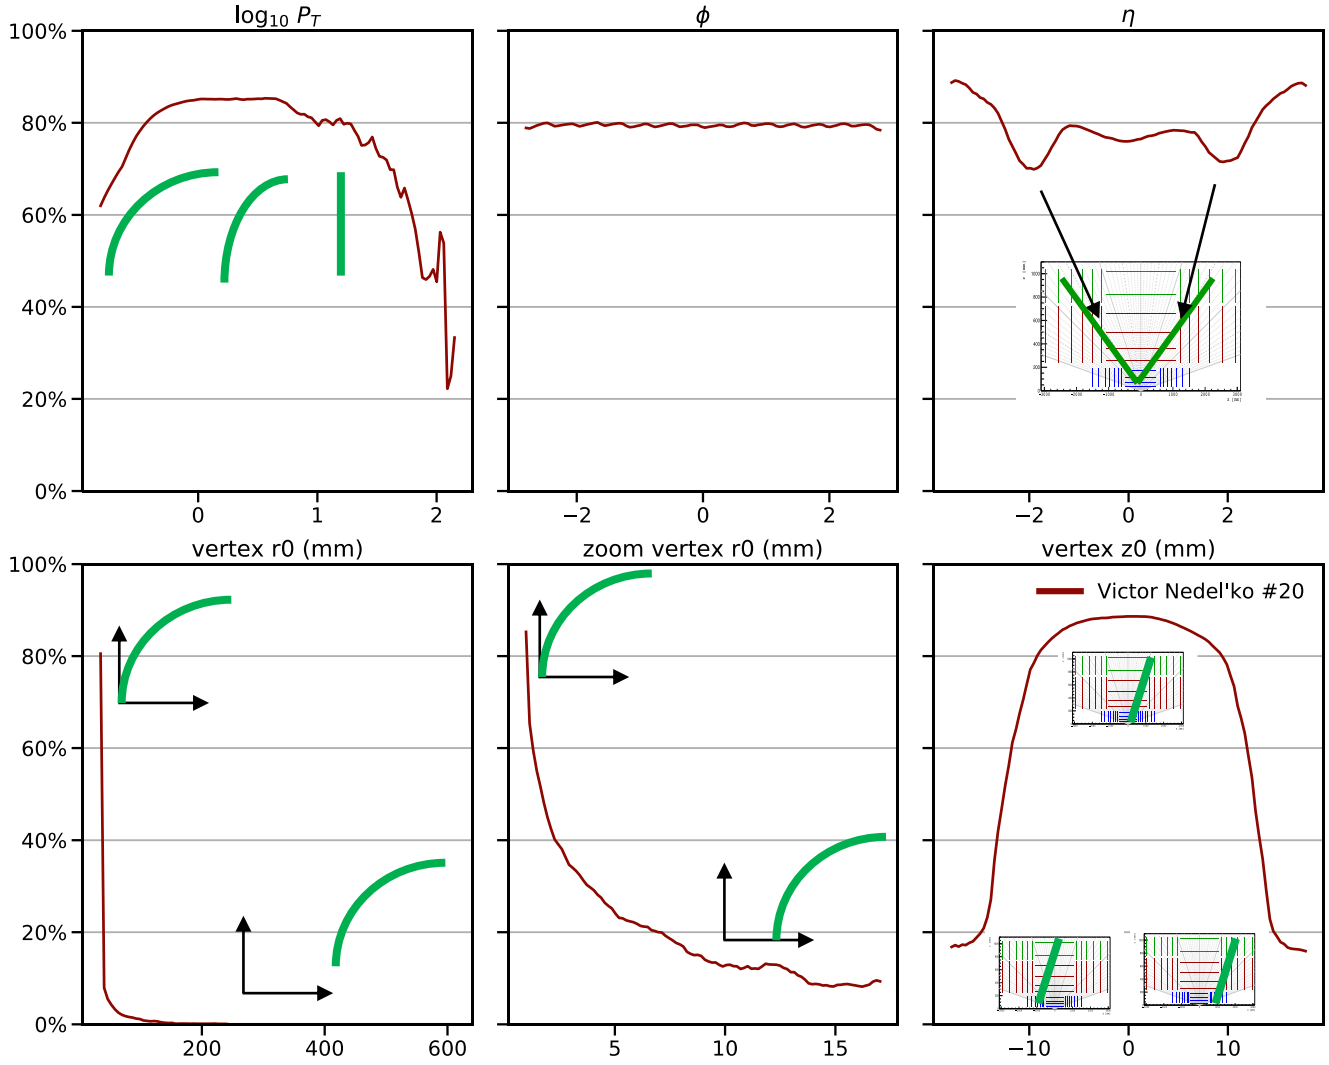

Efficiency all

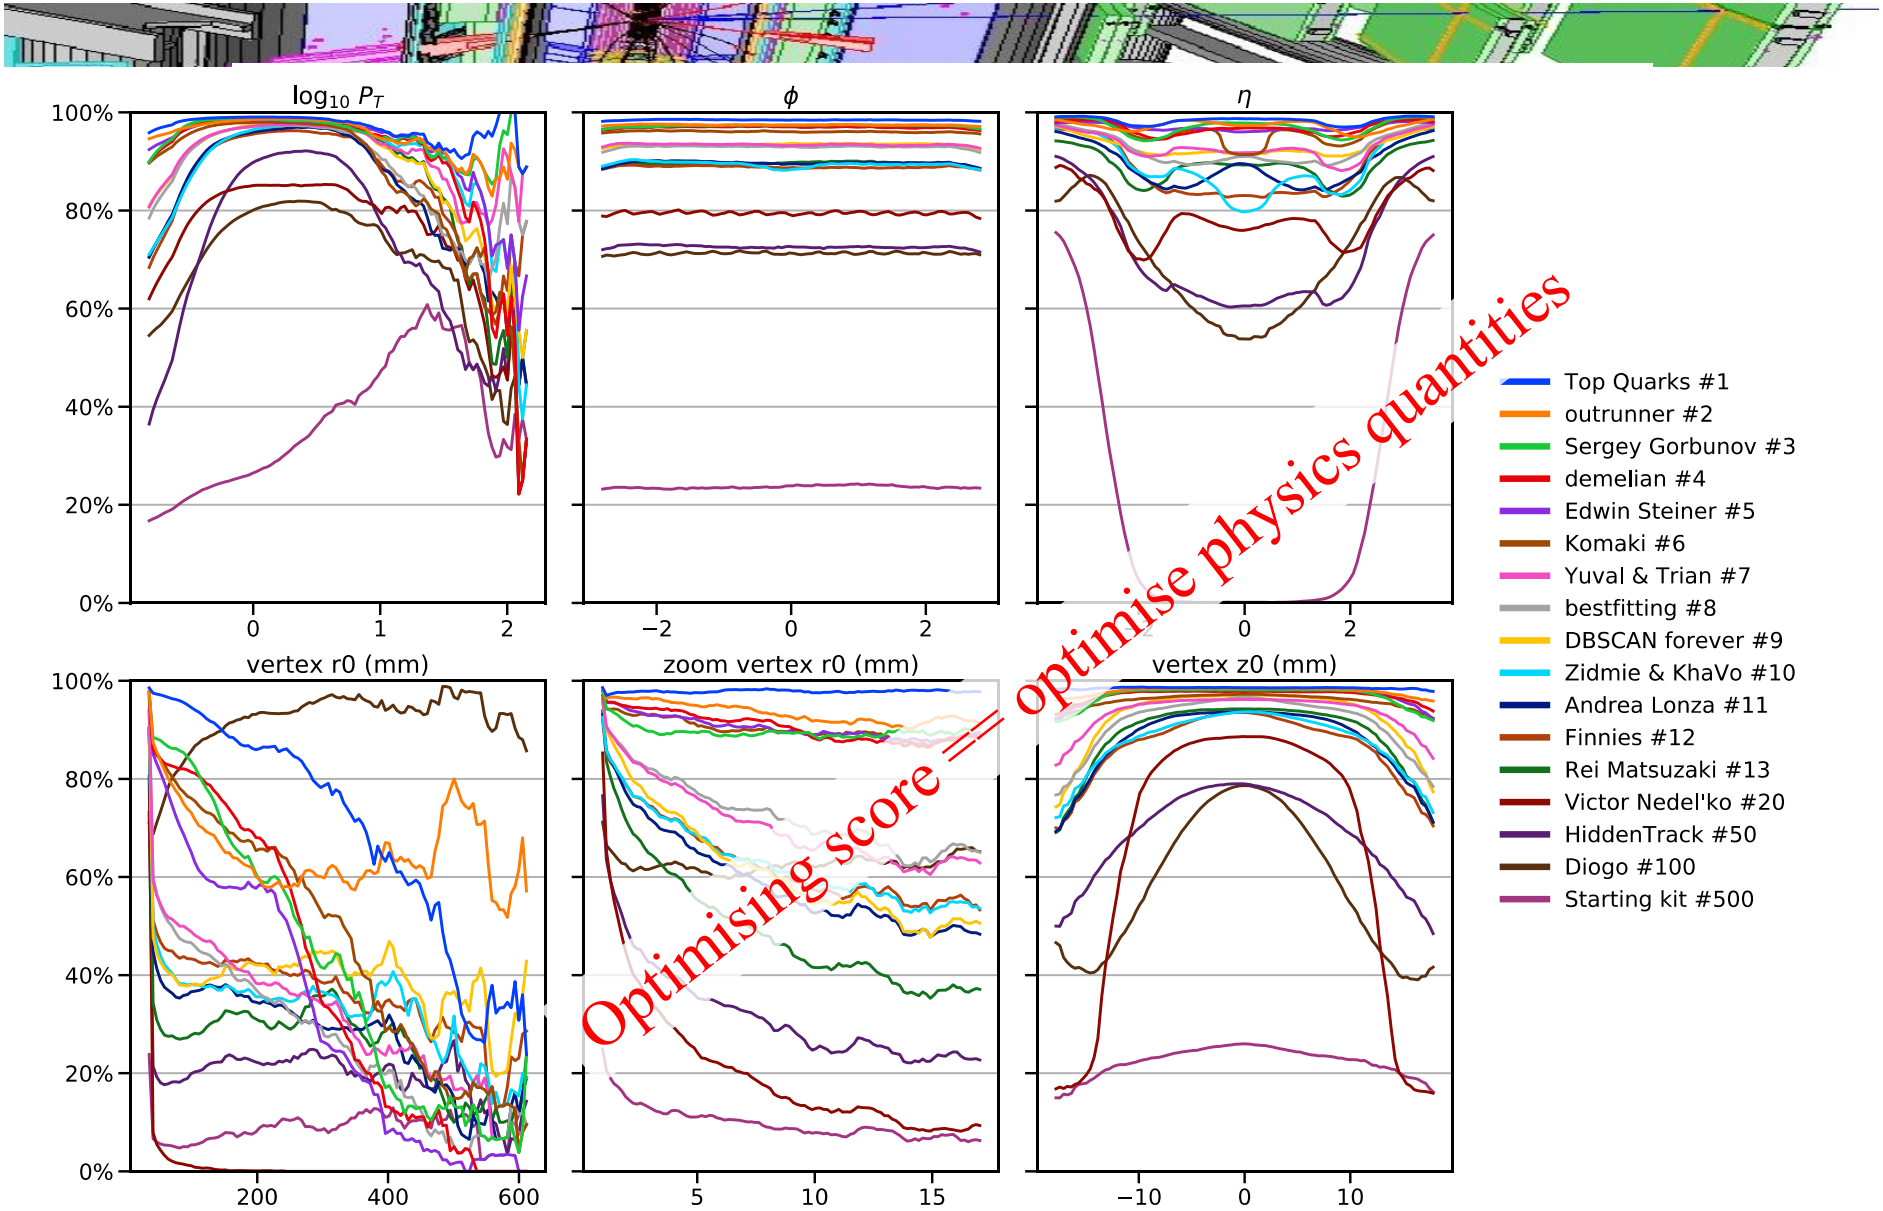

A few competitors

- icecube (Oslo CS master student) #1 92.2 % : combinatorial approach+ML
- outrunner (Taïwan sw engineer) #2 90.3% Deep Learning approach
  - Evaluates probability of any two points to be on same track
  - Very innovative!
  - However takes one full day per event !
- Sergey Gorbunov #3 89.4% demelian #4 87.1% : HEP tracking trigger experts
- Yuval (israëli entrepreneur) & Trian (greek sw engineer) #7 80.4% : unsupervised clustering with randomized projections
- CPMP (french IBM engineer) #9 80.1% : DBSCAN unsupervised clustering algorithm
  - we gave DBSCAN in starting kit, with a 20% score, because in only required a few lines

Participants dendrogram

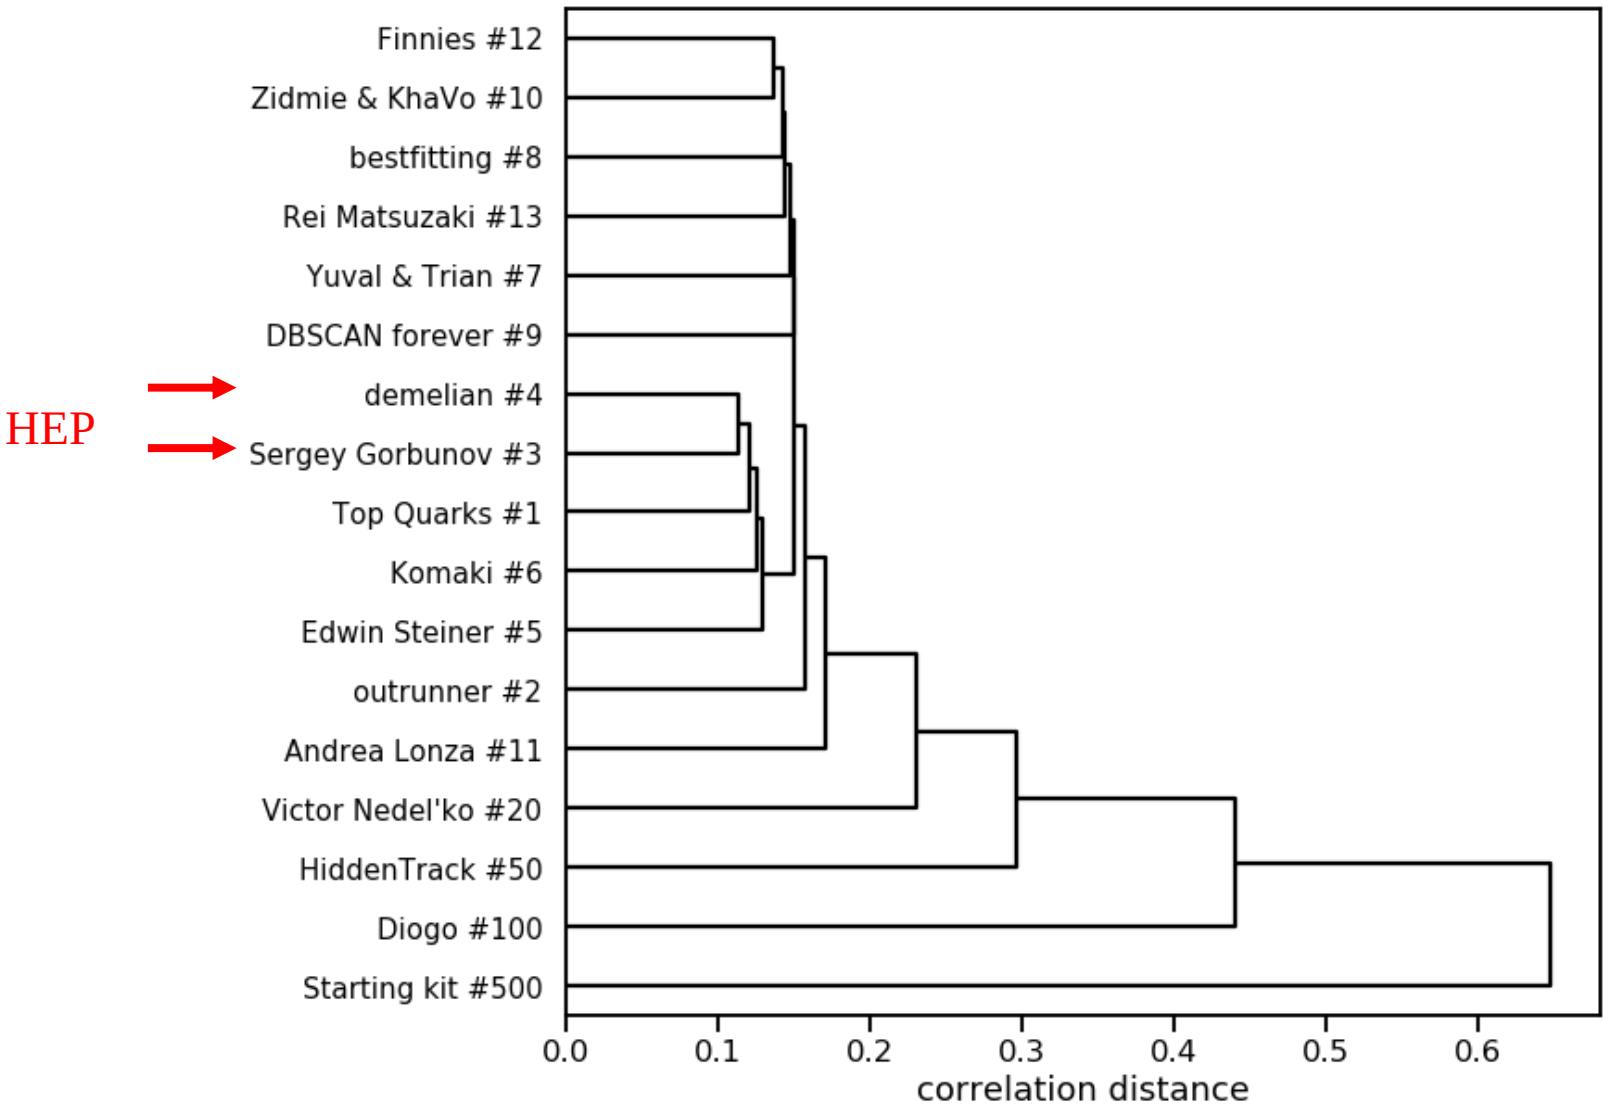

Throughput Phase

Now participants submit their software…
… and are evaluated on accuracy AND speed !

Launched 6th Sep 2018 until 12th March 2019 on Codalab

Codalab Schematic

CodaLab

	hit_id	x	y	z	volume_id	layer_id	module
0	1	-64.409897	-7.163700	-1502.5	7	2	1
1	2	-55.336102	0.635342	-1502.5	7	2	
2	3	-83.830498	-1.143010	-1502.5	7	2	1
3	4	-96.109100	-8.241030	-1502.5	7		1

event(s) are loaded in memory

start

API to call python/C++ API

User executable

stop

solution

Measure both accuracy score and evaluation time

VM 2 cores, 4 Gb memory

Throughput results

- *Ranking score* :
  - 0 if time >600 s or accuracy <50%
  - $\sqrt{log(1 + 600/time)} * (accuracy - 0.5)^2$

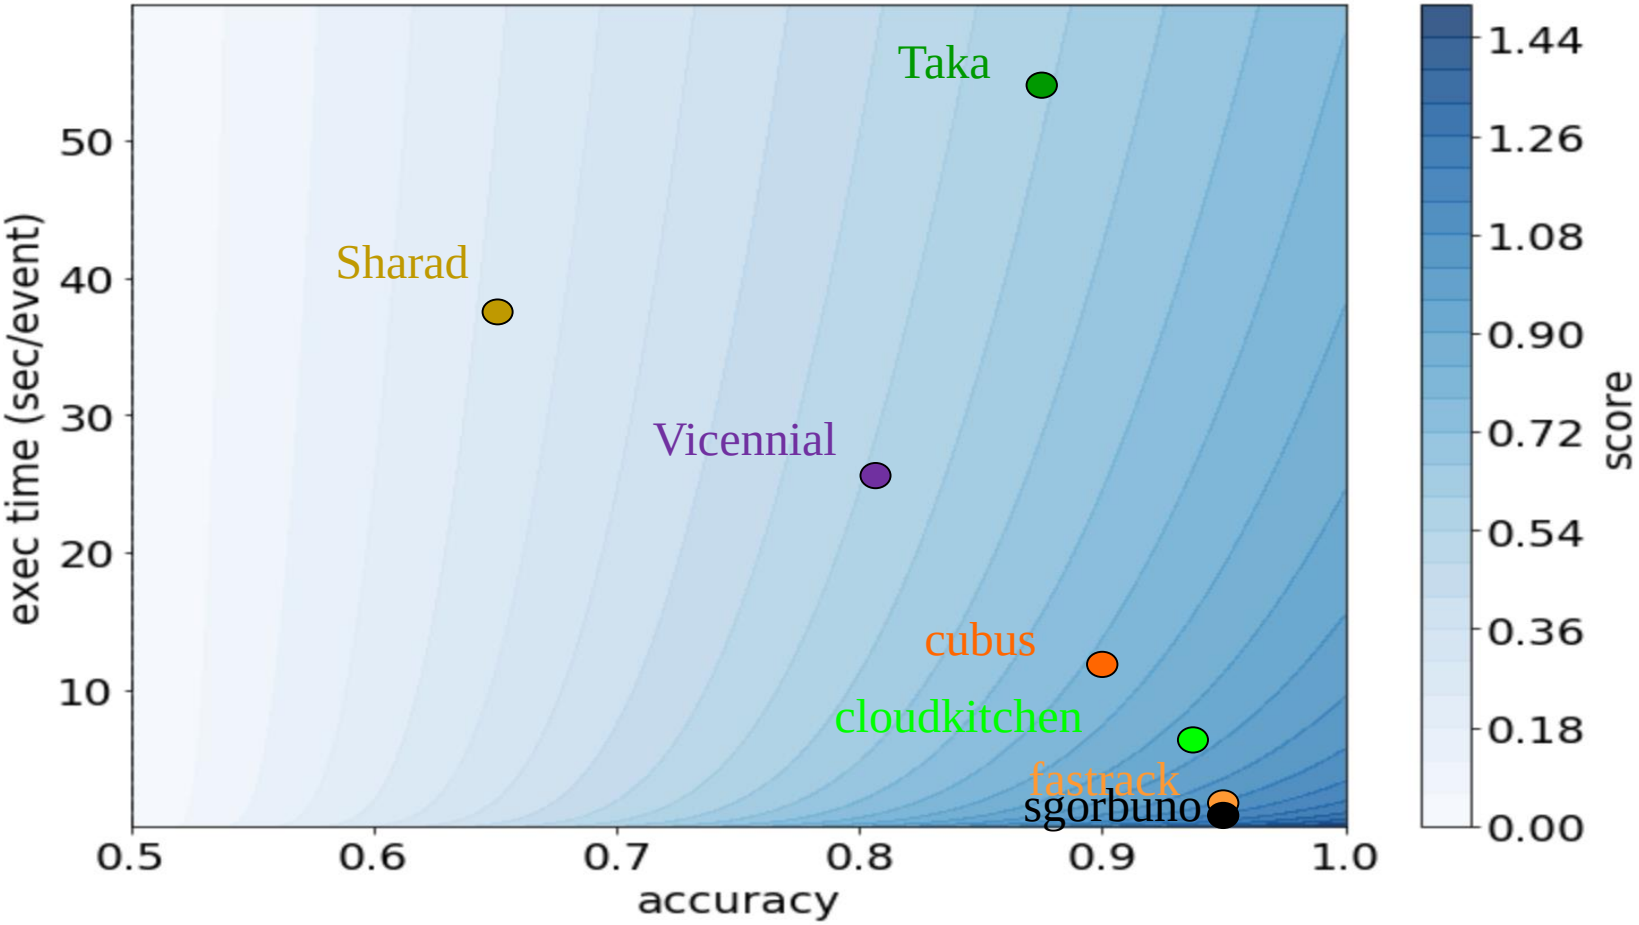

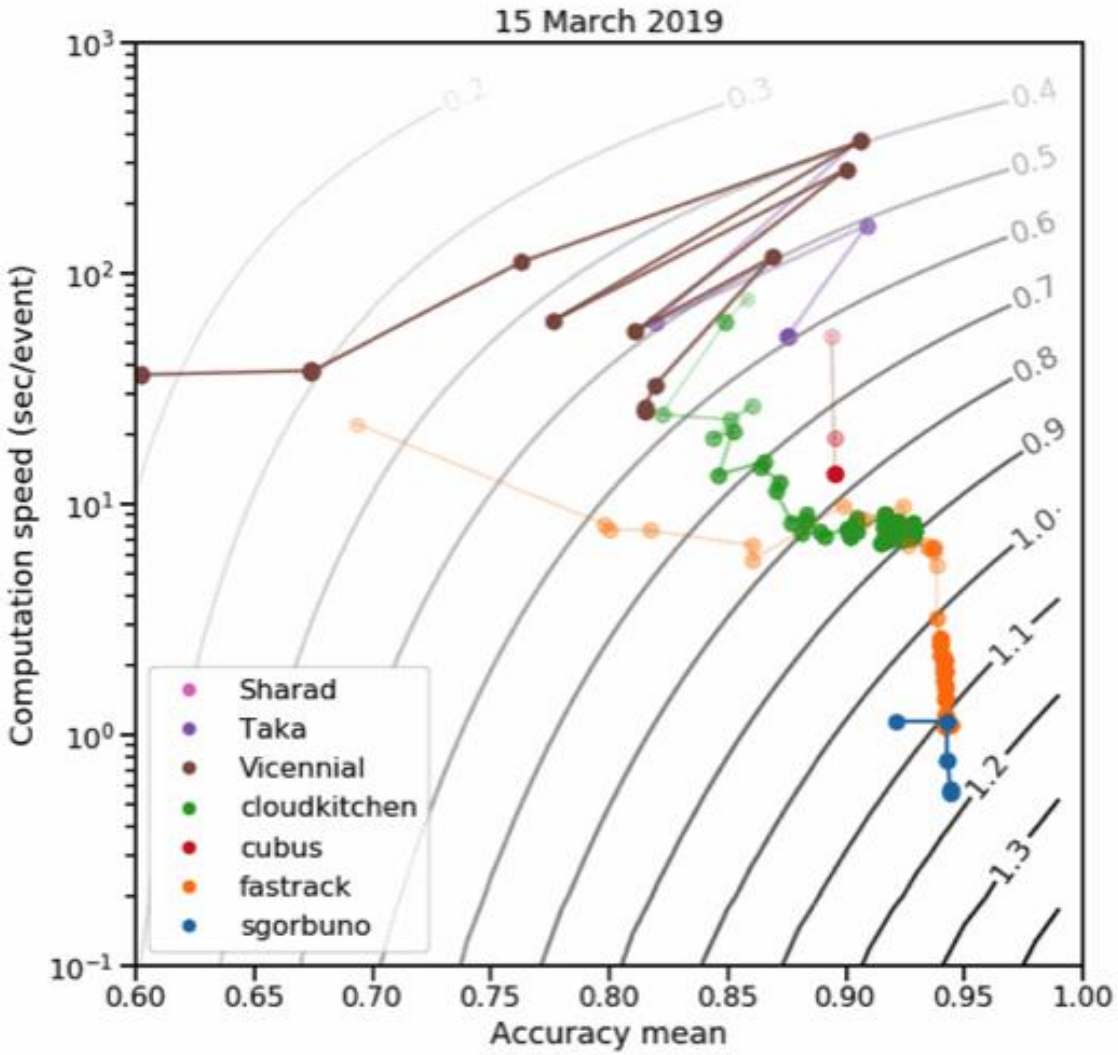

Accuracy mean

Throughput phase LB

RESULTS

#	User	Entries	Date of Last Entry	score ▲	accuracy_mean ▲	accuracy_std ▲	computation time (sec) ▲	computation speed (sec/event) ▲	Duration ▲
1	sgorbuno HEP	9	03/12/19	1.1727 (1)	0.944 (2)	0.00 (14)	28.06 (1)	0.56 (1)	64.00 (1)
2	fastrack HEP	53	03/12/19	1.1145 (2)	0.944 (1)	0.00 (15)	55.51 (16)	1.11 (16)	91.00 (6)
3	cloudkitchen	73	03/12/19	0.9007 (3)	0.928 (3)	0.00 (13)	364.00 (18)	7.28 (18)	407.00 (8)
4	cubus	8	09/13/18	0.7719 (4)	0.895 (4)	0.01 (9)	675.35 (19)	13.51 (19)	724.00 (9)
5	Taka	11	01/13/19	0.5930 (5)	0.875 (5)	0.01 (12)	2668.50 (23)	53.37 (23)	2758.00 (13)
6	Vicennial	27	02/24/19	0.5634 (6)	0.815 (6)	0.01 (10)	1270.73 (20)	25.41 (20)	1339.00 (10)
7	Sharad	57	03/10/19	0.2918 (7)	0.674 (7)	0.02 (4)	1902.20 (22)	38.04 (22)	1986.00 (12)
8	WeizmannAI	5	03/12/19	0.0000 (8)	0.133 (11)	0.01 (11)	88.08 (17)	1.76 (17)	124.00 (7)
9	harshakoundinya	2	03/12/19	0.0000 (8)	0.085 (13)	0.01 (6)	49.22 (8)	0.98 (8)	86.00 (3)
10	iWit	6	03/10/19	0.0000 (8)	0.082 (15)	0.01 (8)	48.23 (3)	0.96 (3)	85.00 (2)
				0.0000					

Throughput phase winner

Phase 2 *Mikado*

Author: Sergey Gorbunov

Accuracy: 0.944
Time/event: 0.56 sec
Memory: 0.1/0.178 Gb (1core/2 cores)

third in Phase-1

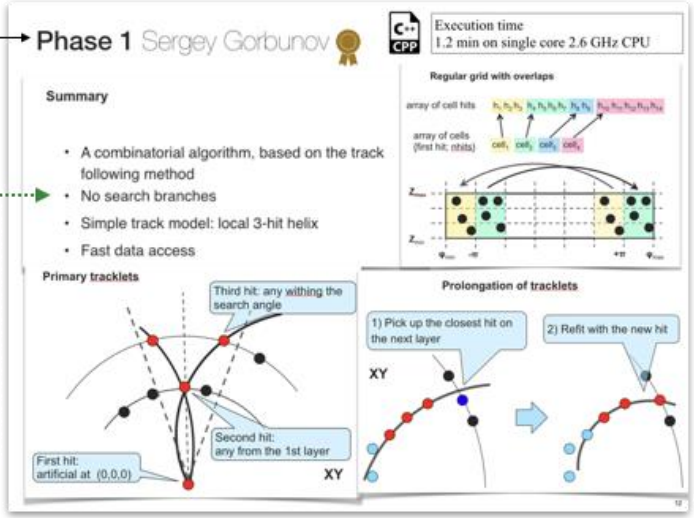

Based on Phase-1 algorithm

- runs iteratively in 80 passes & hit removal from high to low pT
- modifications with respect to Phase 1
  search branches enabled
- every pass has optimised parameters
  results in O(10^4) parameters to be tuned,
  tuning done semi-automated

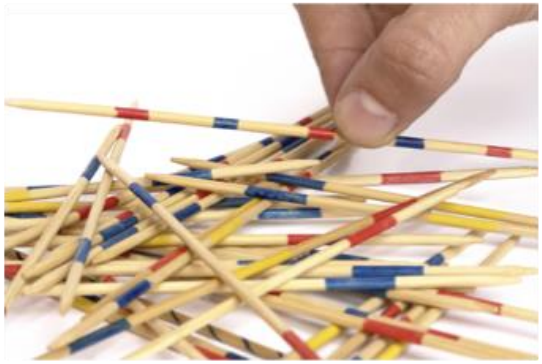

Throughput phase runner up

Phase 2 *FASTrack*

Author: Dmitry Emeliyanov

C++ CPP + OpenMP

Accuracy: 0.944
Time/event: 1.11 sec → 0.8 sec
Memory: 0.6 Gb recently down to

first runner-up to podium in Phase-1

4	—	demelian	0.87079	35	2mo

Phase-1 w/o measurement shapes

Algorithm outline

- using measurement shapes to predict intervals of track inclination
- segment based track following network with embedded Kalman Filter
  - **connection graph** pre-build (&compiled) from `Detector.csv` file
  - run with a **Cellular Automaton (CA), parallelised** with **OpenMP**
  - **candidate building:** graph traversal with applied simplified KF
- combinatorial track following for track completion
  - fast **combinatorial** Kalman Filter using **3rd oder RK** & **simplified field** includes **clone identification** & **track merging**

3 passes (hit removal):
- high momentum
- low momentum
- rest

Throughput phase 3rd place

Phase 2 cloudkitchen

Author: Marcel Kunze

Accuracy: 0.93
Time/event: ~7 sec
Memory: 0.7 Gb

partly based on top quarks Phase 1 solution

1	—	**Top Quarks**		0.92182	10	2mo

Algorithm outline

hits

sorted in voxels

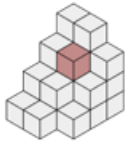

organised in direct acyclic graphs (DAG)

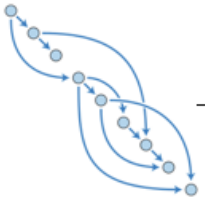

DAGs are pre-trained on ~25 events ground truth

DAGs are used to fast navigate through voxel space
- $\pm z$ graph set
- $\eta - \phi$ graph set

Main steps

- Select promising pairs
  - 7 million / 0.99
- Extend pairs to triples
  - 12 million / 0.97
- Extend triples to tracks
  - 12 million / 0.95
- Add duplicate hits to tracks
  - 12 million / 0.96
- Assign hits to tracks
  - 90% of hits / 0.92

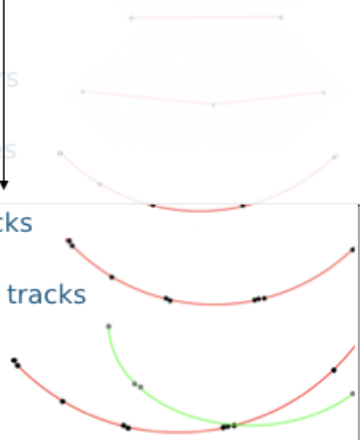

Triplet finder

NN3

doublet finder

NN1 NN2

Threaded

Efficiency

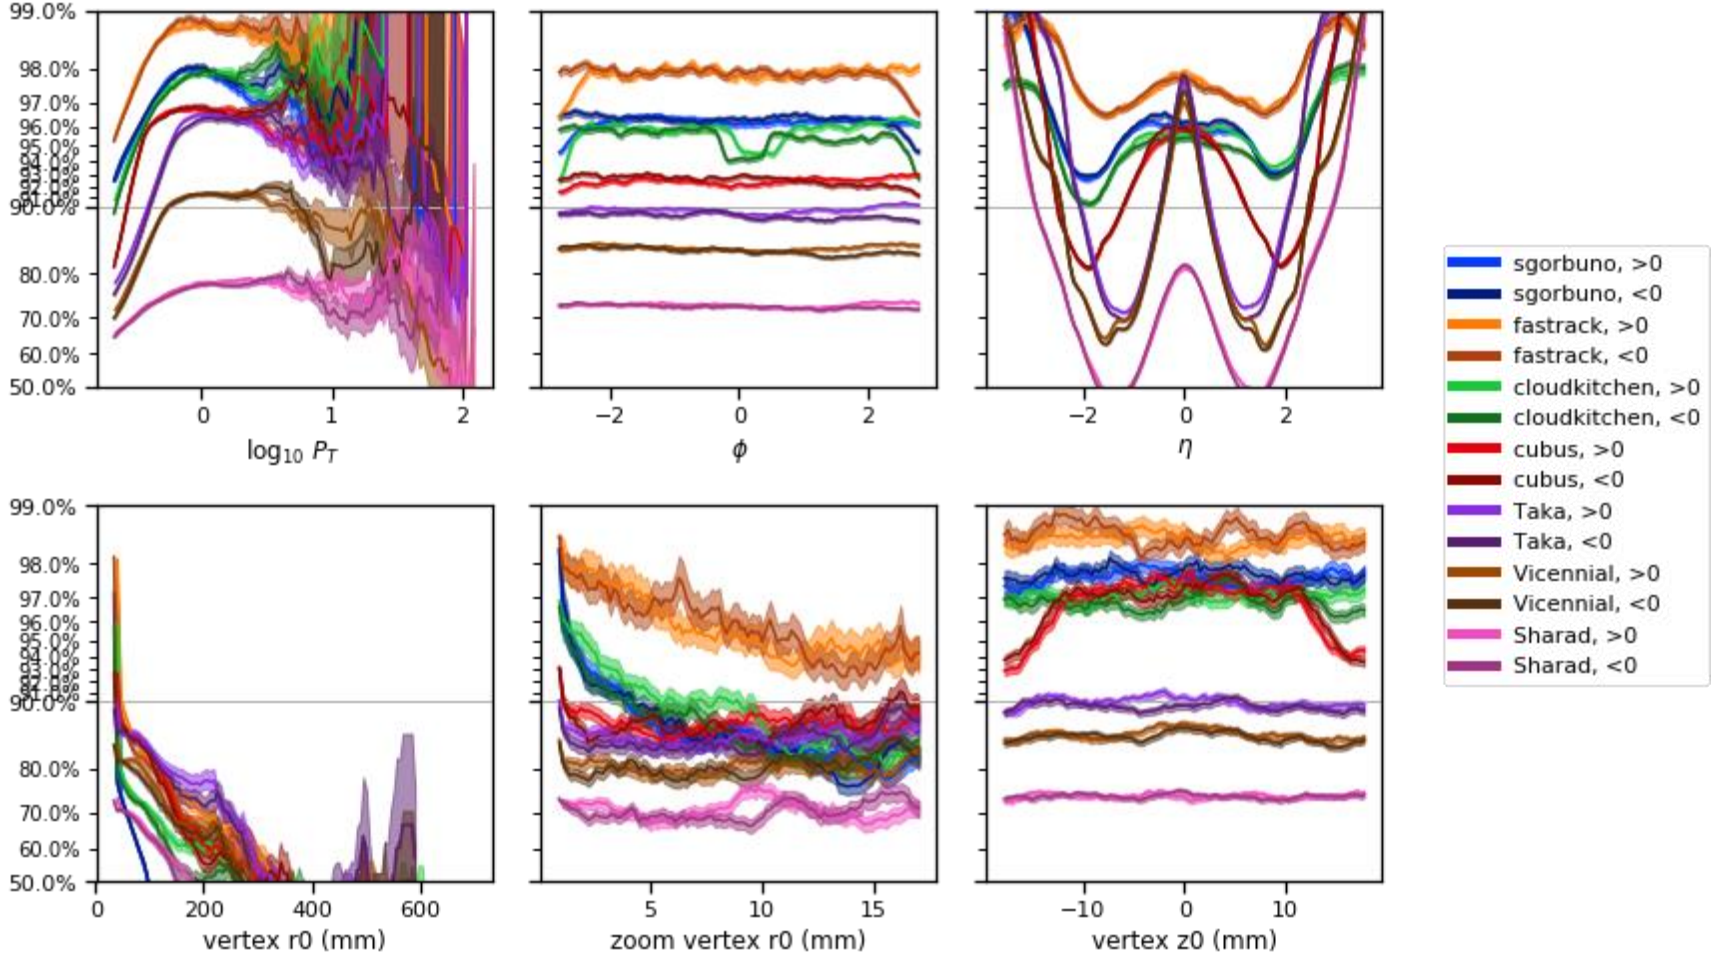

Not quite as good as for accuracy phase: do not spend time when not worth it

Track separation

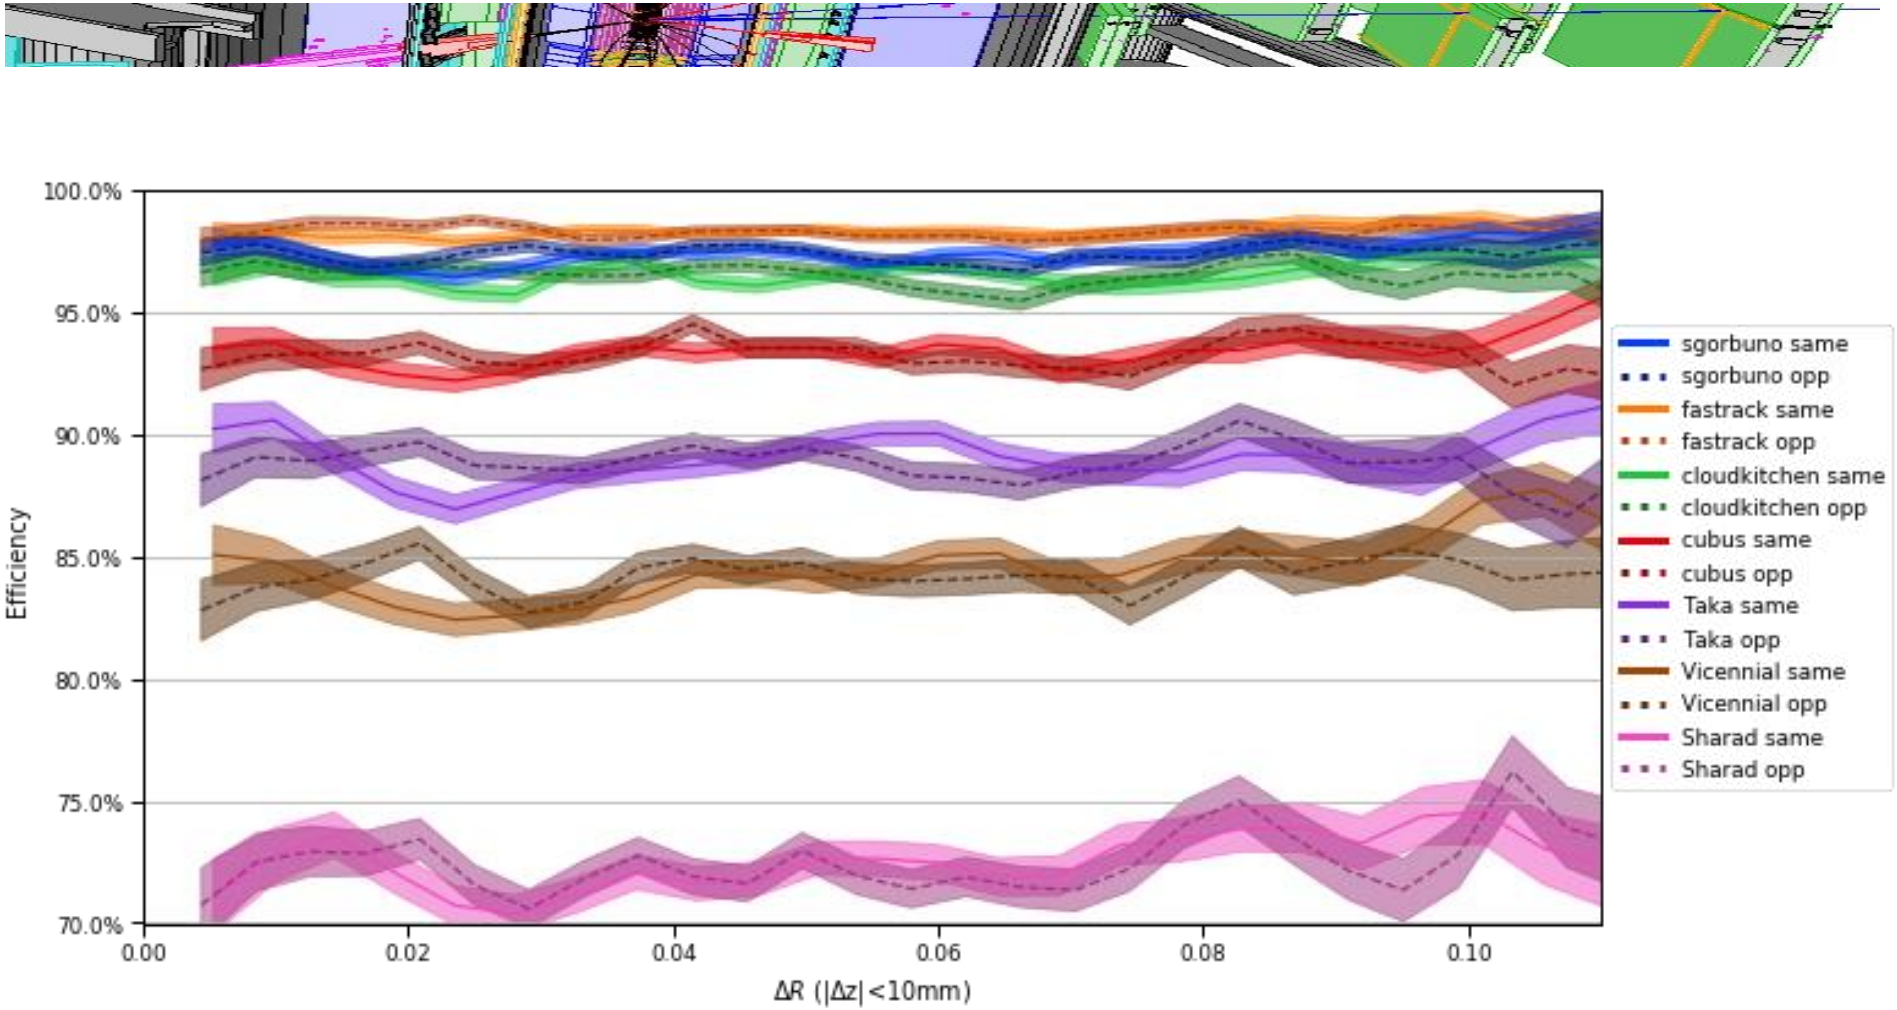

No sign of track confusion (higher stat of high pT jet needed)

Pre-Conclusion

- So we've organised a tracking competion to reach out to Computer Science and winner and runner-up are our own HEP tracking experts…
- ML –enhanced algorithms are on the table available for further study
- Competition on a common dataset with well define figure-of-merit appealed to our experts
- Dataset will be released permanently to serve as a benchmark
- »final» solution will be blended from HEP expertise and new ideas

Conclusion

- Contact : trackml.contact@gmail.com https://sites.google.com/site/trackmlparticle Twitter : @trackmllhc
- Accuracy phase @ Kaggle : https://www.kaggle.com/c/trackml-particle-identification
  - Different approaches identified, sometimes new in the field
  - Now working on decyphering/combining them
  - ➔chapter in the NeurIPS 2018 Competition book arXiv:1904.06778 (just out today!)
- Throughput phase @ Codalab : https://competitions.codalab.org/competitions/20112
  - Superfast (0.5s, 1s, compared to state of the art 10-50s) and accurate solutions submitted
- Port of best submissions to ACTS on-going
- TrackML grand finale workshop at CERN 1st 2nd July 2019 indico place holder
  - Open to all.
  - Best participants invited
- What then (discussion at CTDWIT 2019 Valencia)?
  - "final" document to be delivered this summer
  - Dataset will be released permanently on CERN Open Data portal in the coming weeks
  - Study the possibility to reopen an open-ended "competition" on codalab, possibly with GPU