



UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

TrackML throughput challenge on CodaLab

Marcel Kunze, Heidelberg University

PROCESSORS

MACHINE LEARNING



Introduction

TrackML was a data science competition organized in 2018 on Kaggle and CodaLab platforms.

The aim of the challenge was to

- stimulate development of new particle tracking algorithms for the HEP community
- Get the best ideas and techniques from the Machine Learning community



UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386



sponsors



NVIDIA.



UNIVERSITÉ
DE GENÈVE



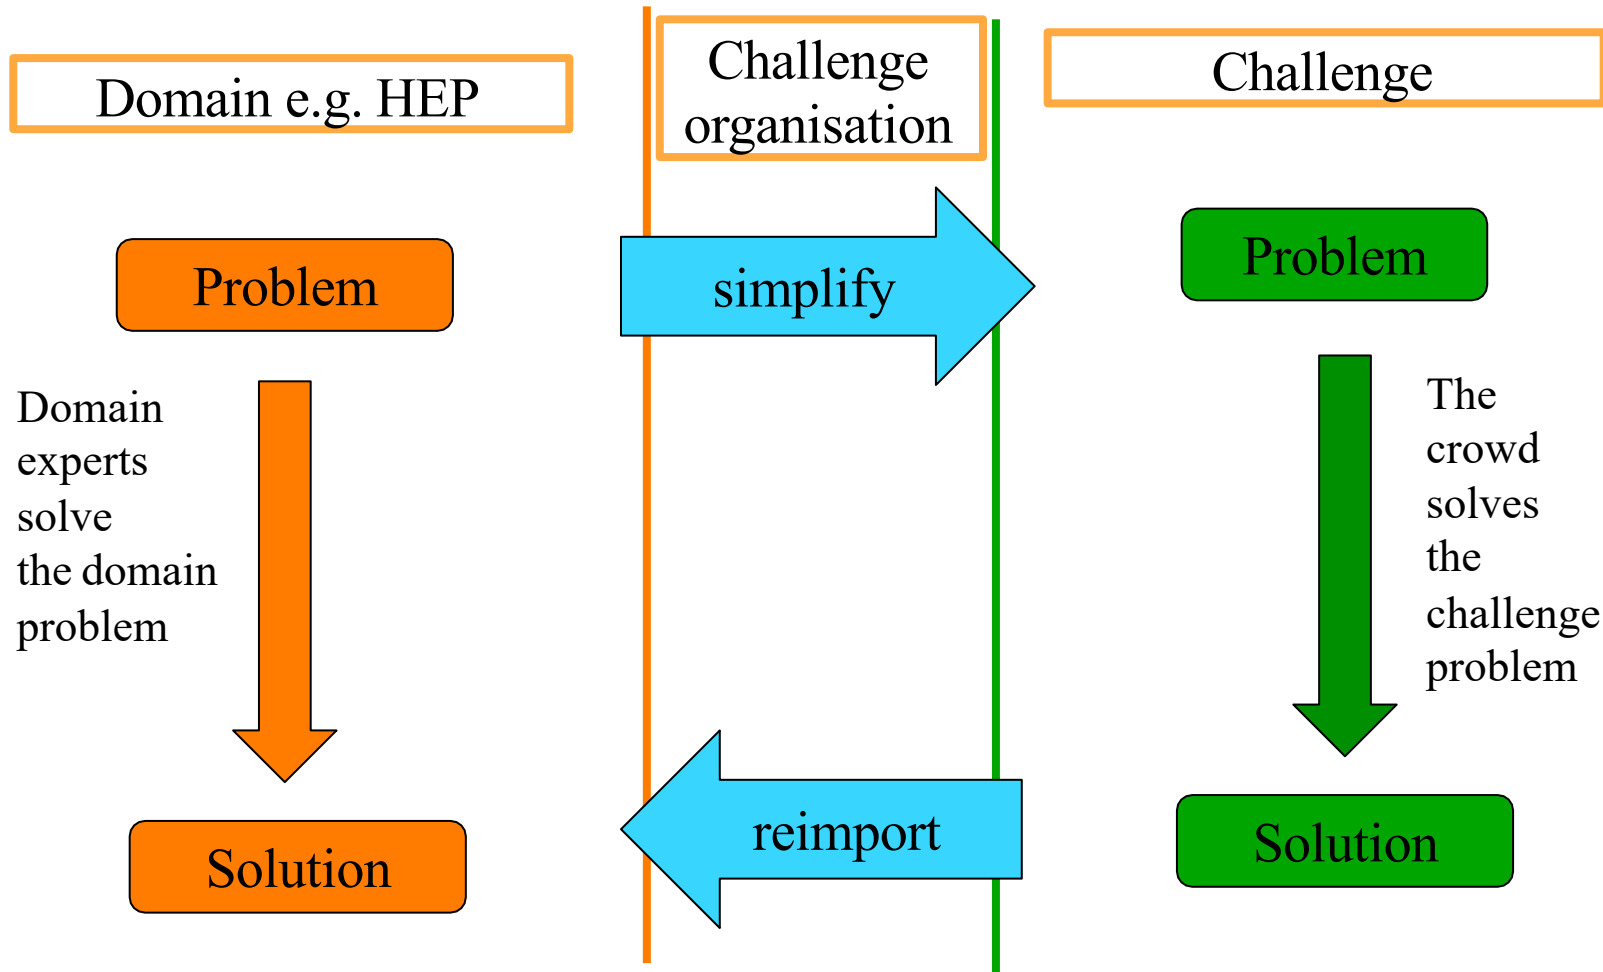
Organisation team

Jean-Roch Vlimant (Caltech),
Vincenzo Innocente, Andreas Salzburger (CERN),
Isabelle Guyon (ChaLearn),
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, Victor Estrade (LRI-Orsay),
Edward Moyse (University of Massachusetts),
Mikhail Hushchyn, Andrey Ustyuzhanin (Yandex, HSE)

From Domain to Challenge and back



UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386



CodaLab Schematic

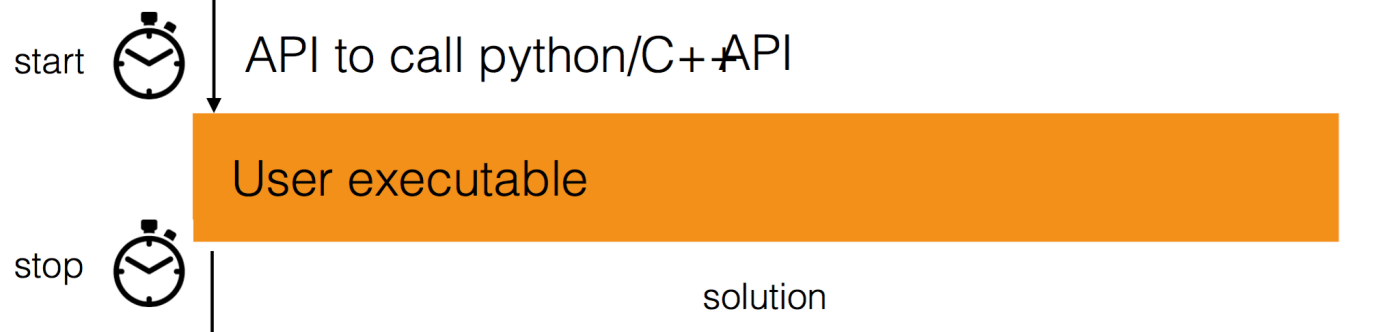


UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

CodaLab

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

event(s) are loaded in memory



VM 2 cores, 4 Gb memory



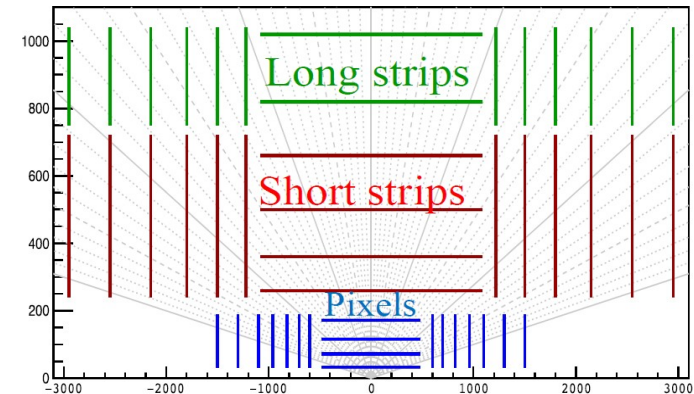
TrackML challenge in a nutshell



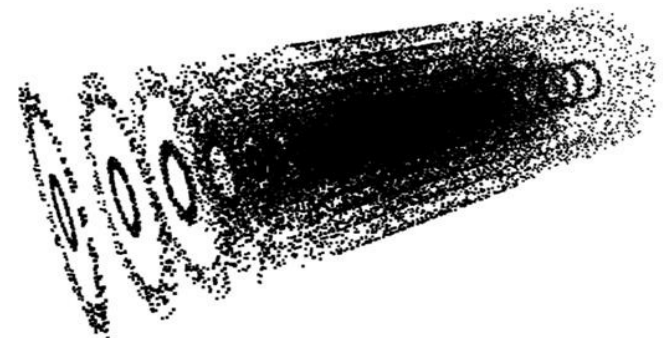
UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

- **Based on a simplified, yet realistic detector model**
 - non-uniform magnetic field similar to ATLAS solenoid
 - detailed simulation of particle interactions with detector material
 - three types of Si-detectors: pixel, shortstrips, long strips
- **The goal is reconstruct all tracks in the detector**
 - 10K tracks/event, min pT = 120 MeV, min number of hits = 4
- Test data: 50 events, each event consists of
 - a list of particle position measurements (hits) in 3D space(x,y,z)
 - a list of individual silicon detector cells associated with each hit
- Training data (10K events) : the above + ground truth
 - 0.1 billion truth tracks, 1 billion hits, size O(100 Gb)
- **Solution**
 - unique hit-to-track associations for test events

TrackML detector geometry : r-z view



TrackML event : 100K points, 10K tracks



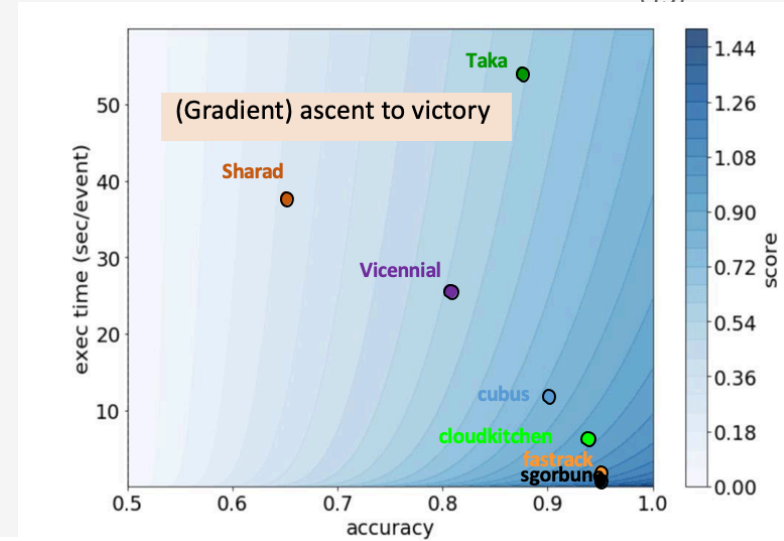
Throughput phase Leader Board



UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

RESULTS

#	User	Entries	Date of Last Entry	score ▲	accuracy_mean ▲	accuracy_std ▲	computation time (sec) ▲	computation speed (sec/event) ▲	Duration ▲
1	sgorbuno	9	03/12/19	1.1727 (1)	0.944 (2)	0.00 (14)	28.06 (1)	0.56 (1)	64.00 (1)
2	fastrack	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)			
7	Sharad	57	03/10/19	0.2918 (7)	0.674 (7)	0.02 (4)			
8	WeizmannAI	5	03/12/19	0.0000 (8)	0.133 (11)	0.01 (11)			
9	harshakoundinya	2	03/12/19	0.0000 (8)	0.085 (13)	0.01 (6)			
10	iWit	6	03/10/19	0.0000 (8)	0.082 (15)	0.01 (8)			



Phase 2 Mikado



Author: Sergey Gorbunov

third in Phase-1

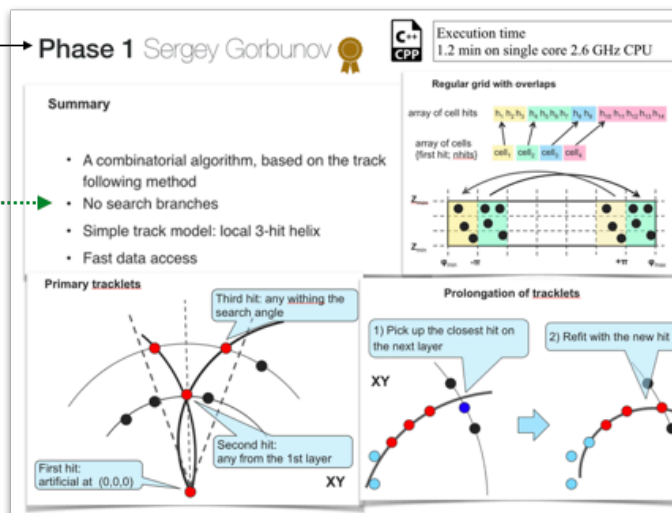


Accuracy: 0.944
Time/event: 0.56 sec
Memory: 0.1/0.178 Gb (1core/2 cores)

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

no machine learning used



Phase 2 FASTrack

Author: Dmitry Emelianov




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
---	---	----------	---	---------	----	-----

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 order RK** & **simplified field**
 - includes **clone identification** & **track merging**

3 passes (hit removal):

- high momentum
- low momentum
- rest

Throughput phase 3rd place



UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

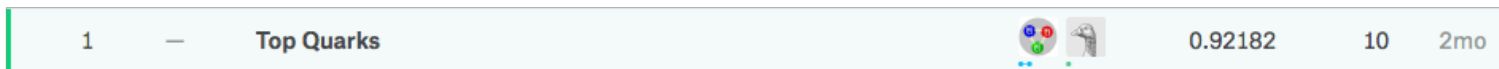
Phase 2 cloudkitchen

Author: Marcel Kunze

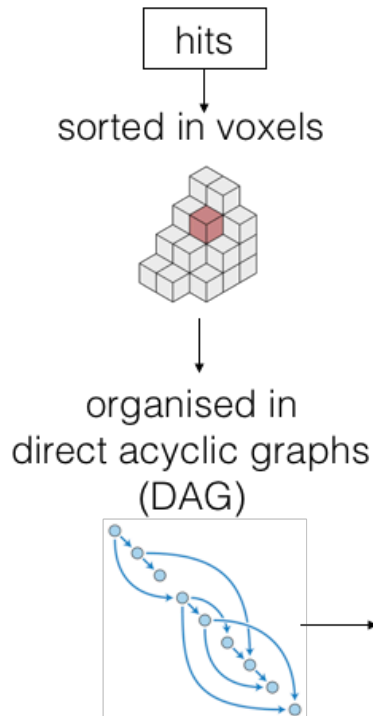


Accuracy: 0.93
Time/event: ~7 sec
Memory: 0.7 Gb

partly based on top quarks Phase 1 solution



Algorithm outline



Main steps

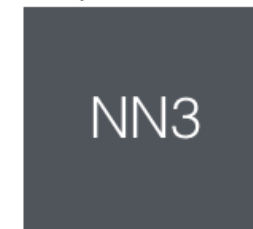
- Select promising pairs
 - 7 million / 0.99
- Extend pairs to triples
 - 12 million / 0.97
- Extend triples to tracks
 - 12 million / 0.95
- Add duplicate hits to tracks
 - 12 million / 0.96
- Assign hits to tracks
 - 90% of hits / 0.92

DAGs are pre-trained on ~25 events ground truth

DAGs are used to fast navigate through voxel space

Disc section
Tube section

Triplet finder



ca. 300k
97.2%

doublet finder



ca. 500k
99.4%

Threaded

ca. 2 Mio.

Directed Graphs

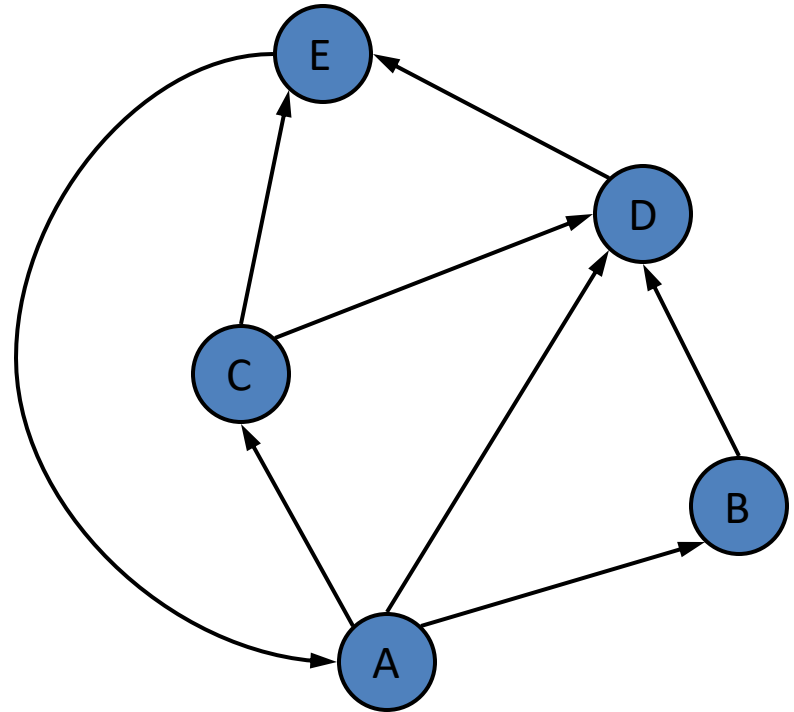


UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

A **directed Graph** is a graph whose edges are all directed

Applications

- one-way streets
- flights
- task scheduling
- ...

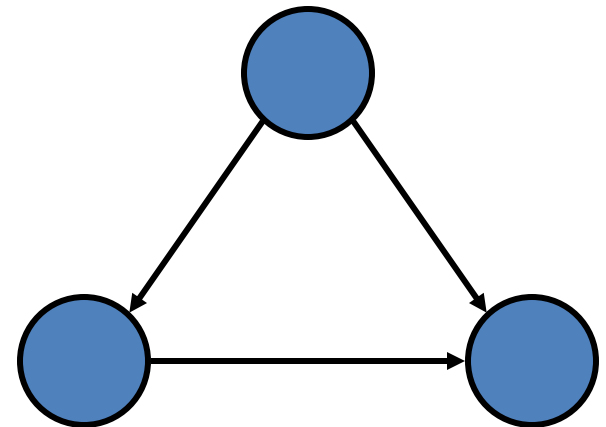
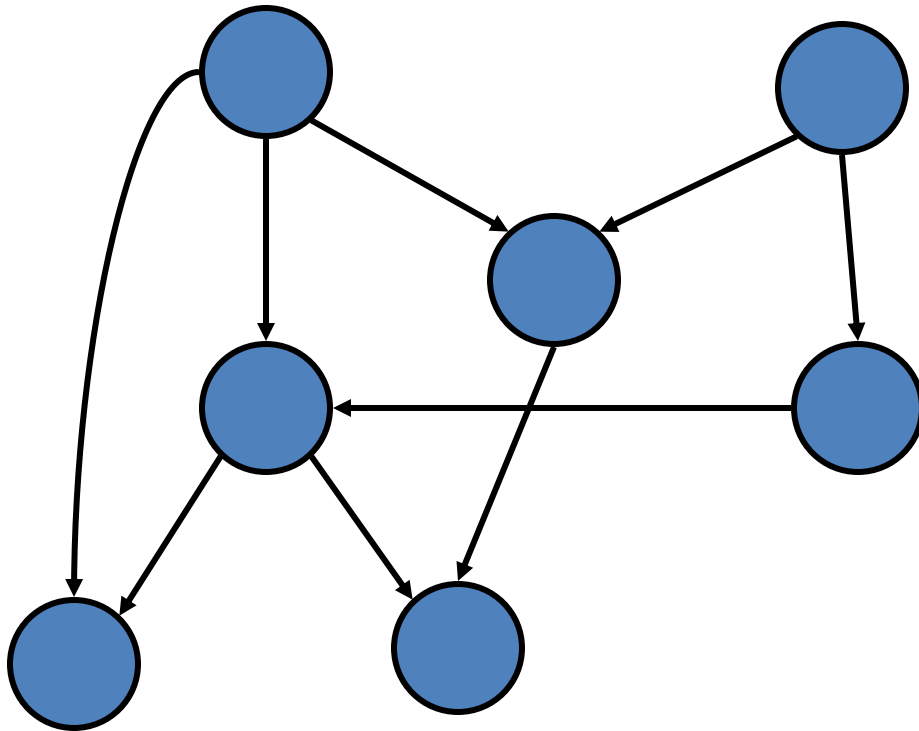


Directed Acyclic Graphs (DAG)



UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

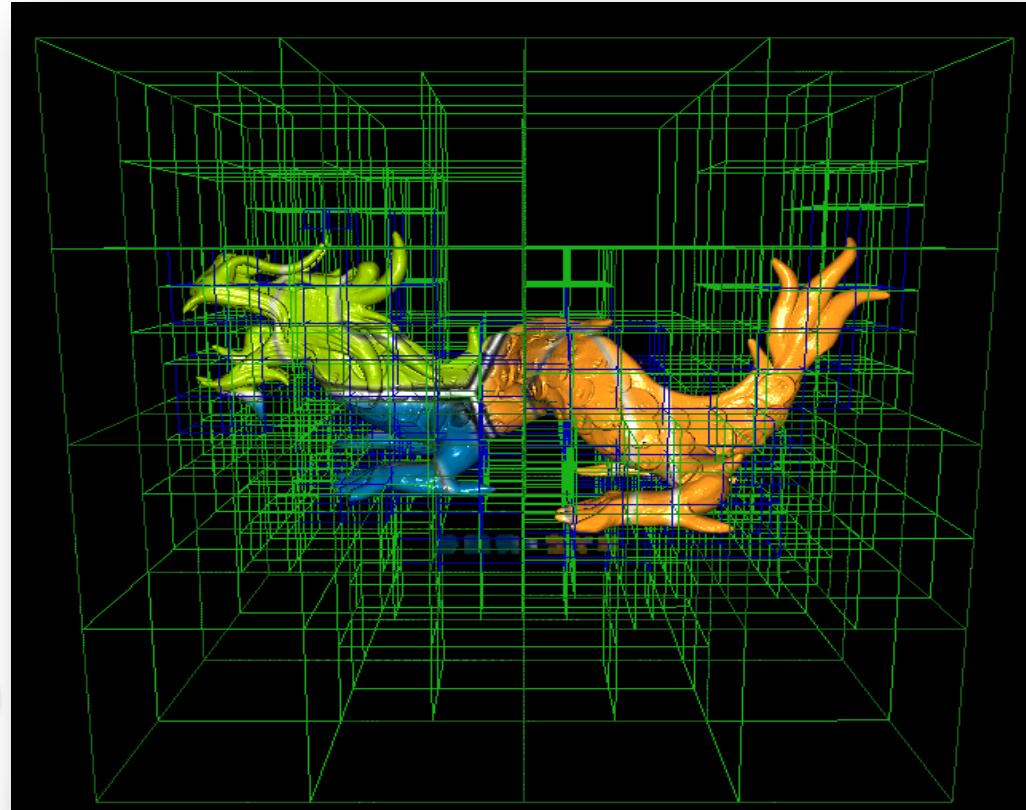
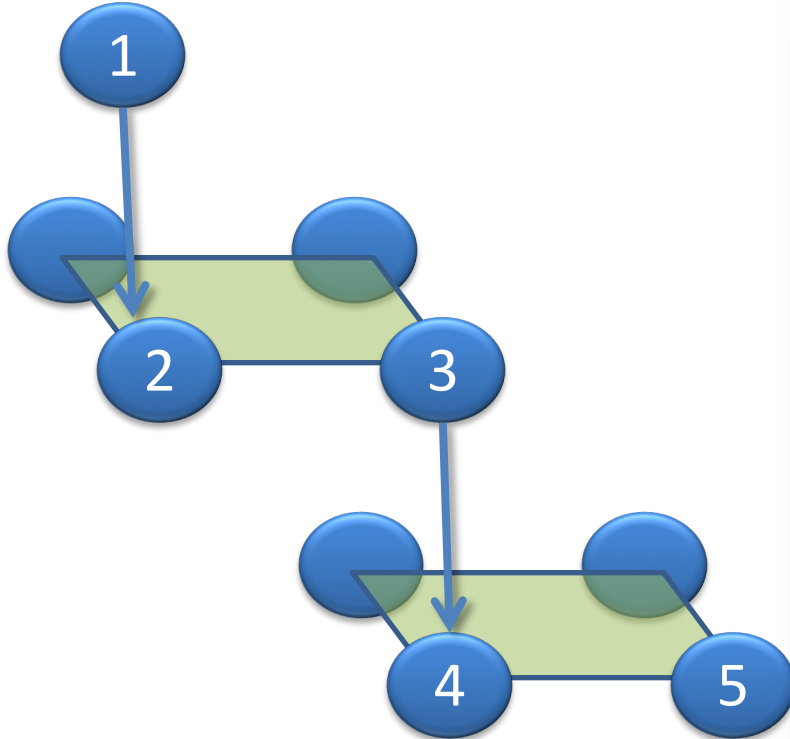
A *directed acyclic graph* or *DAG* is a directed graph with no directed cycles:



Gaming: Sparse Voxel Octrees (SVO)



UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386



- Raytracing
- Compression of data
- Multi-scale resolution

Voxel (Volume Pixel)



Define spatial elements in $\phi*\theta$ (voxel)

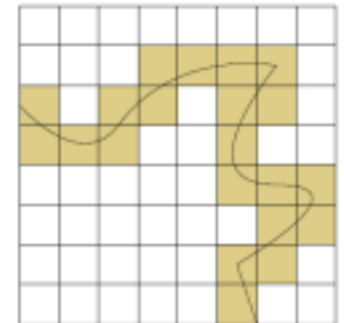
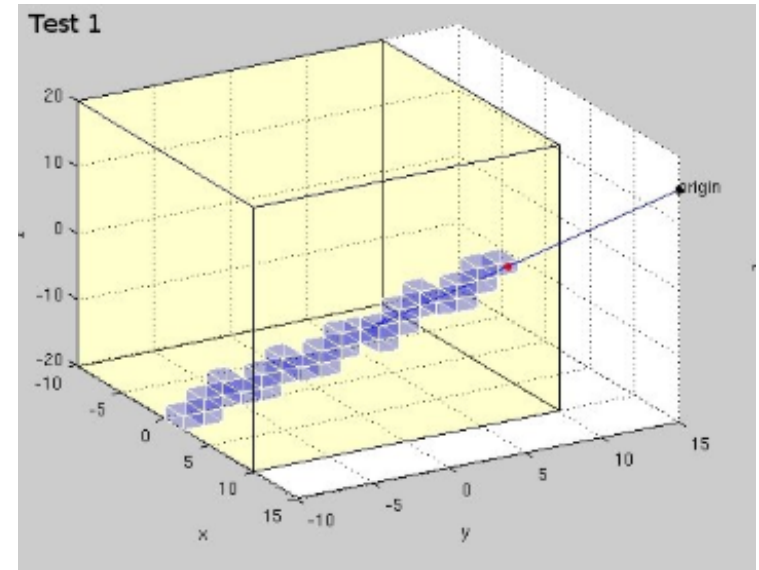
- Organize the voxels in DAGs according to track evolution in radial direction
 $\text{index} = (\text{phi} \ll 32) | (\text{theta} \ll 24) | (\text{layer} \ll 16) | \text{module};$
- Flexible to model even arbitrary paths (kinks, missing hits, outliers, random walk, ..)
- Training is done with MC tracks of typically 15-25 events

Multiscale resolution (Better use SVOs?)

- 2*1 DAGs for pair finding (slices)
- 12*14 DAGs for triple finding (tiles)

Path finding

- Sort event hits into the trained DAGs
- Seed and follow the path strategy





Intuition

- Model free estimator
- Start with basic quantities
- Coordinates, simple derived values
- Only very basic detector specific information

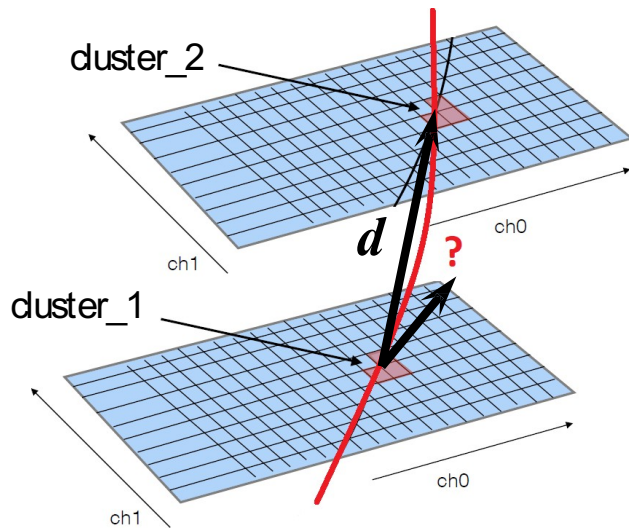
Input parameter space

- Polar coordinates (R_t , ϕ , z)
 - Directional cosines
 - Simple helix calculation (score)
- } In principal not needed, but speeds up the thing !

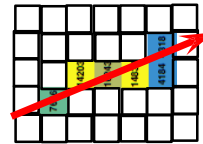
Training

- Supervised: presenting MC ground truth
- Unsupervised: presenting probability density function

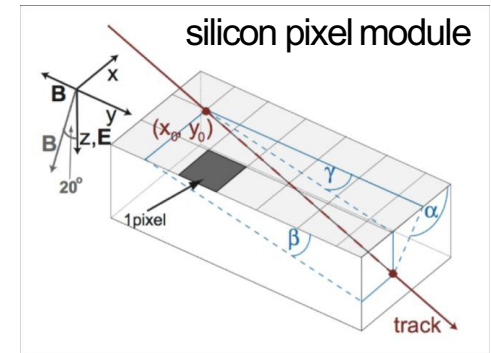
Given two hits (clusters of silicon cells): predict if they belong to the same track



- Estimate track direction from the cluster shape:



eigenvector of covariance matrix of the silicon cells



Features for the training

- Polar coordinates of the hit doublet: (r_1, ϕ_1, z_1) , (r_2, ϕ_2, z_2)
- Triplet finder works the same with a hit triplet
- Simple helix score
- Angle/length deviations of the vector d projection from the values predicted by the shape of cluster 1
- Angle/length deviations of the vector d projection from the values predicted by the shape of cluster 2

Input Parameter Folding



The tracking problem is symmetric wrt. polar coordinates

- Fold the input parameter space into an octagon slice using “abs” function
- Considerable improvement of the separation strength of the parameters
- Need less statistics / yield better results

Rank	Variable	Separation
1	log(score)	5.039e-01
2	rz3	5.491e-04
3	phi3	7.552e-05
4	z3	4.986e-05
5	rz2	1.519e-05
6	rz1	9.568e-06
7	phi2	4.101e-06
8	z1	1.967e-06
9	z2	1.965e-06
10	phi1	1.503e-06



Rank	Variable	Separation
1	log(score)	5.978e-01
2	rz3	6.329e-04
3	abs(abs(phi3)-1.57079632679)	1.317e-04
4	abs(z3)	5.522e-05
5	rz2	2.067e-05
6	rz1	1.675e-05
7	abs(abs(phi2)-1.57079632679)	4.335e-06
8	abs(z1)	3.592e-06
9	abs(abs(phi1)-1.57079632679)	3.038e-06
10	abs(z2)	2.963e-06

Hit Doublet / Triplet Classification: MLP

“Shallow learning” ;)

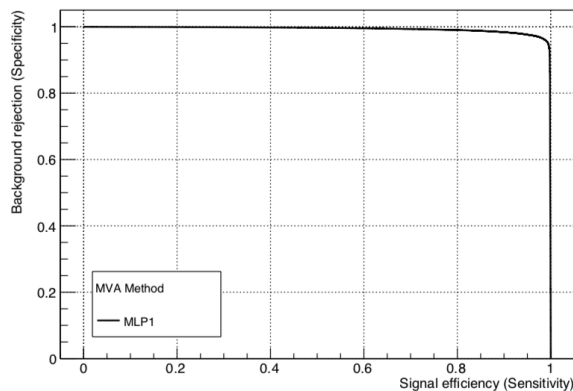


UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

- **Classify the doublets and triplets with neural networks**
 - Multi Layer Perceptron: MLP1 8-15-5-1 / MLP2 9-15-5-1 / MLP3 10-15-5-1
 - Input: hit coordinates, directional cosines towards the clusters, helicity score wrt. origin
 - Output: doublet/triplet quality, supervised training with Monte-Carlo ground truth
 - Training: Typically 10 events, O(Mio) patterns, 500 epochs, one hour on standard PC
 - “*Receiver Operation Characteristics*” (ROC) curves indicate good quality

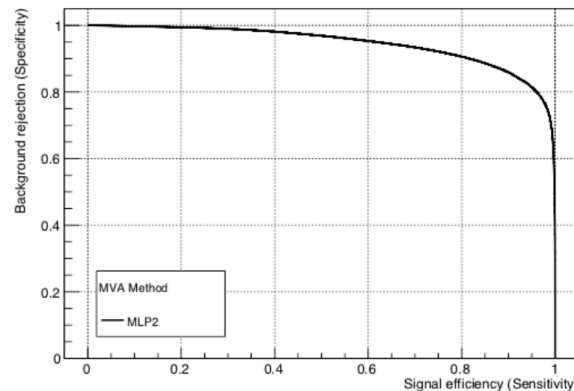
Doublet finder (disc)

Signal efficiency vs. Background rejection



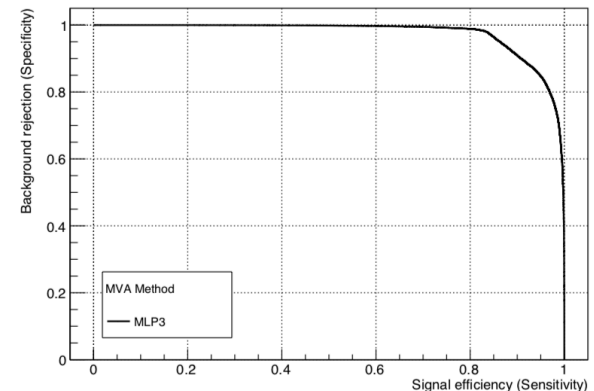
Doublet finder (tube)

Signal efficiency vs. Background rejection



Triplet finder

Signal efficiency vs. Background rejection



Worse due to vertex shift !

Hyperparameter Tuning



Automated tests with docker / singularity to maximize CodaLab score

Test set of 50 events not used by training. Optimize:

- Spatial resolution / training of DAGs
- Network topology and cuts on output wrt. event size
- Run time / accuracy trade-offs

```

Timer0 initHits 12.000000 ms
Timer1 initCells 0.000000 ms
Timer2 initGraphData 27.000000 ms
Timer3 initHitDir 11.000000 ms
Timer4 initPolarModule 202.000000 ms
Timer5 initRecoObjects 0.000000 ms
Timer6 initTasks 0.000000 ms
Timer7 findCandidatesGraph 1136.000000 ms
Timer8 findTriplesGraph 1967.000000 ms
Timer9 findPaths 816.000000 ms
Timer10 addDuplicates 762.000000 ms
Timer11 findAssignment1 774.000000 ms
Timer12 findAssignment2 106.000000 ms
Timer13 mapAssignment 29.000000 ms
Timer14 writeSubmission 10.000000 ms
Processing time per event 5852.000000 ms
    
```

Files 20, Phi / Theta	8	10	12	14	16
8	0.883360		0.878024		
10		0.880177	0.884479	0.887572	0.885034
12	0.878399	0.883600	0.887683	0.889858	0.881736
14		0.880297	0.877356	0.884148	0.878094
16		0.882559	0.885102	0.876590	0.871375

T3 / hits	90000	100000	110000	120000	130000	140000
0.2	0.896278					
0.3	0.896026					
0.4	0.896748	0.871153	0.847126			
0.5	0.895815	0.871703	0.847288	0.825986	0.806712	0.779419
0.6	0.893367	0.871531	0.847247	0.826128	0.806855	0.780699
0.7			0.846020	0.825648	0.806230	0.780974
0.8						0.779338

- **Well defined algorithmic steps for pattern recognition**
- **Efficient parallelism on the basis of DAGs**
 - Form doublets from seeding hits in a DAG (MLP1, MLP2)
 - Extend the doublets to triplets (MLP3)
 - Extend the triplets to path segments
 - The path segments are merged into tracklets
 - Remove duplicate solutions

The tracklets are merged into a common tracking solution by **serial tasks**

```
Timer0 initHits 12.000000 ms
Timer1 initCells 0.000000 ms
Timer2 initGraphData 27.000000 ms
Timer3 initHitDir 11.000000 ms
Timer4 initPolarModule 202.000000 ms
Timer5 initRecoObjects 0.000000 ms
Timer6 initTasks 0.000000 ms
Timer7 findCandidatesGraph 1136.000000 ms
Timer8 findTriplesGraph 1967.000000 ms
Timer9 findPaths 816.000000 ms
Timer10 addDuplicates 762.000000 ms
Timer11 findAssignment1 774.000000 ms
Timer12 findAssignment2 106.000000 ms
Timer13 mapAssignment 29.000000 ms
Timer14 writeSubmission 10.000000 ms
Processing time per event 5852.000000 ms
```

Serial tasks: ca. 0.3 seconds

Parallel tasks: ca. 4 seconds

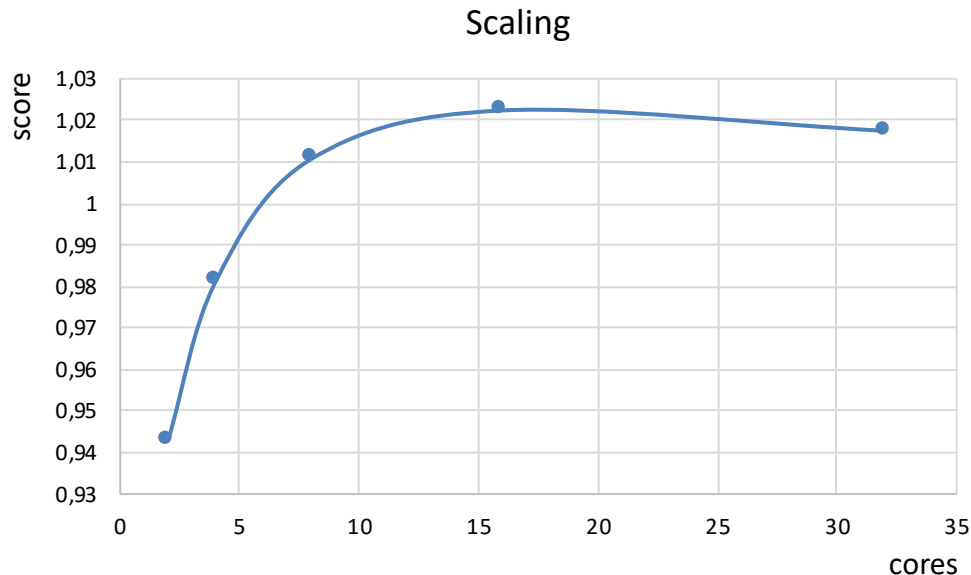
Serial tasks: ca. 0.8 seconds

Scaling Behavior



Scaling tests have been performed with Amazon EC2

- Instance type c5n.9xlarge (36 cores)
- Core power comparable to CodaLab cores
- Code scales up to 16 cores (Score: 1.022, accuracy 92.3%, 1.7s)
- Limited by serial code: Sorting tracklets into tracks (improve by use of OpenMP ?)



Amdahls Law: Speedup is the fraction of code P that can be parallelized:

$$speedup = \frac{1}{1 - P}$$

Model free estimator

- ## Graceful degradation in presence of changes

- ## The DAGs may represent arbitrary tracking paths

-
- A 6x6 grid with a yellow shaded path and a black line graph. The path starts at (1,1), goes to (1,2), (2,2), (2,3), (3,3), (3,4), (4,4), (4,5), (5,5), (5,6), and ends at (6,6). The black line graph starts at (1,1), goes to (1,2), (2,2), (2,3), (3,3), (3,4), (4,4), (4,5), (5,5), (5,6), and ends at (6,6).



Machine Learning Software: Neural Network Objects

Neural Network Objects (NNO) is a C++ class library for Machine Learning based on the ROOT framework

Supervised models

- Multi-Layer Perceptron (TMLP, TXMLP)
- Fisher Discriminant (TFD)
- **Supervised Growing Cell Structure (TSGCS)**
- **Supervised Growing Neural Gas (TSGNG)**
- Neural Network Kernel (TNNK)

Unsupervised models

- Learning Vector Quantization (TLVQ)
- **Growing Cell Structure (TGCS)**
- **Growing Neural Gas (TGNG)**

Published on <https://github.com/marcelkunze/rhonno>

The solution has also been trained with ROOT/TMVA, yields comparable results.

Dr. Marcel Kunze

marcel.kunze@uni-heidelberg.de

Im Neuenheimer Feld 293 / 106

D-69120 Heidelberg



**UNIVERSITÄT
HEIDELBERG**
ZUKUNFT
SEIT 1386

