Warsaw University of Technology

DNN for distortion fluctuation calibration of ALICE TPC

Maja Kabus

ALICE

1st MODE workshop, 07.09.2021

Introduction to TPC space-charge distortions

Particle trajectory reconstruction in ALICE

ALICE is one of the 4 experiments located at the **Large Hadron Collider (LHC)** at the CERN laboratory.

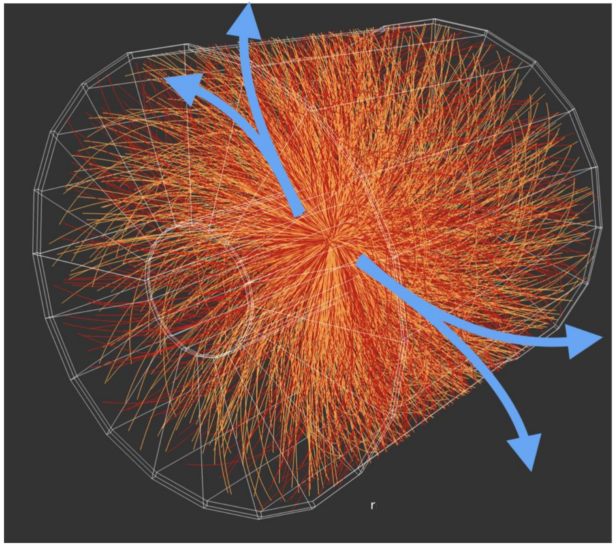

One of **the most challenging** High-Energy Physics (HEP) tasks is the reconstruction of **particle trajectories** or "**tracks."**

In ALICE, tracking is performed by using a gas detector called **Time Projection Chamber (TPC)**.

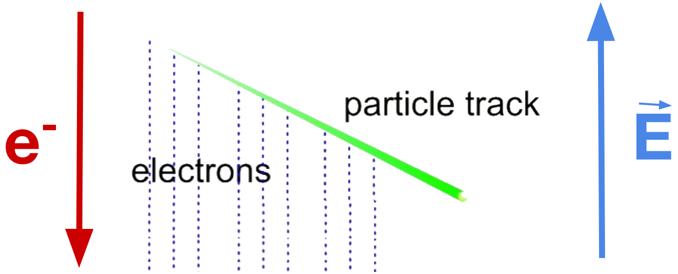

We measure the **electrons ionized** by the passage of charged particles in the gas.

How does the Time Projection Chamber work?

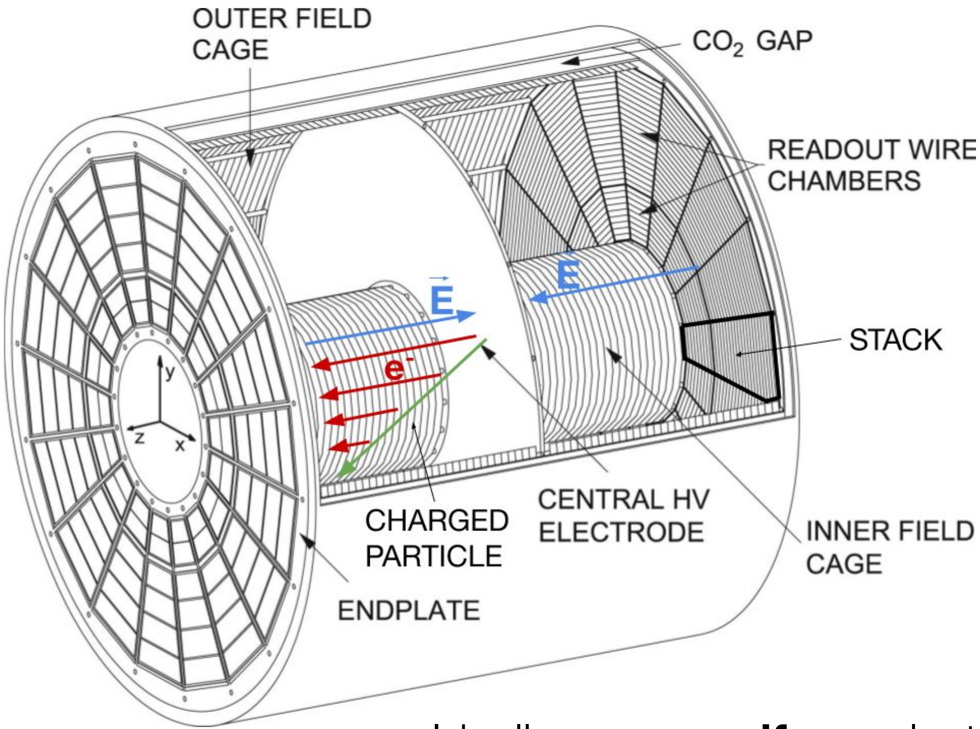

1. Gas ionization by charged particles.
2. The drift of the ionization electrons to the readout chambers.
3. Signal amplified and collected.

3D information for each track point:

- $r\varphi$ via the position in the readout chamber
- z via speed and time of drift, $s = vt$

Ideally: a very **uniform** electric field –> accurate tracks measurements

TPC ongoing upgrade

Run 3 – the next data-taking period

- huge interaction rate: **50 kHz**
  –> much more particles and ions inside the detector
- continuous readout instead of the triggered mode
- no **gates** at the readout chambers
  –> much more amplification ions moving from the readout to the drift volume

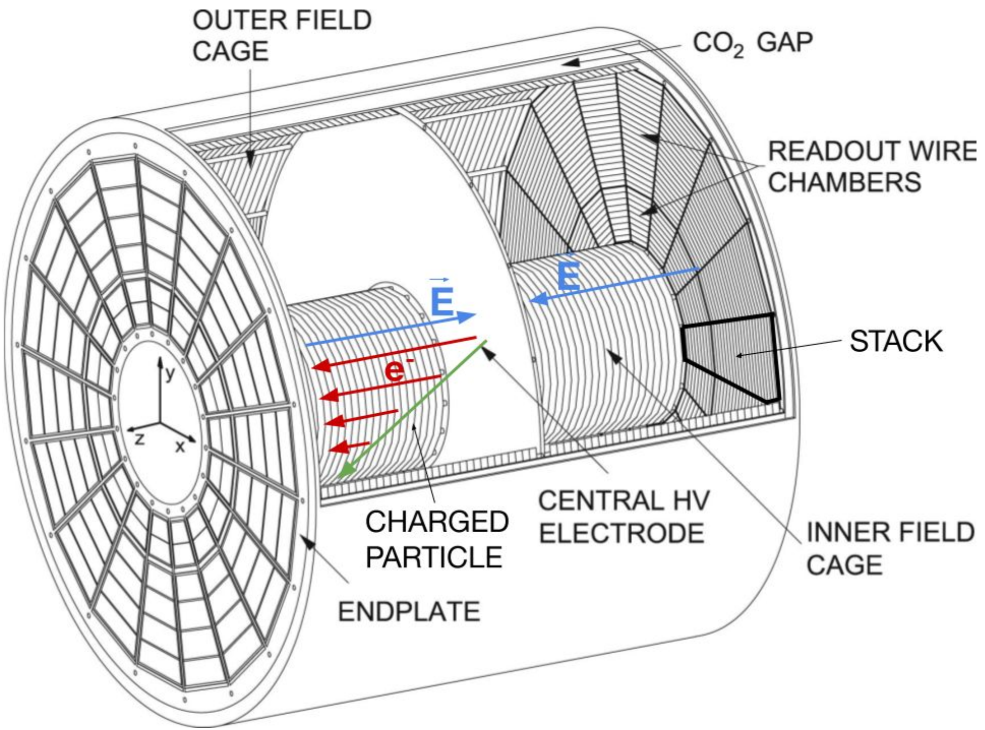

The research problem: TPC distortions

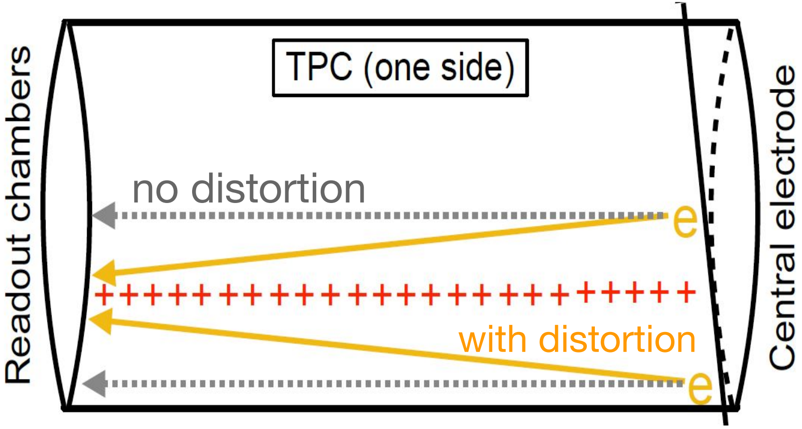

The positive ions are very slow and distort the electric field lines.
–> The electric field is **not constant** and **not uniform!**

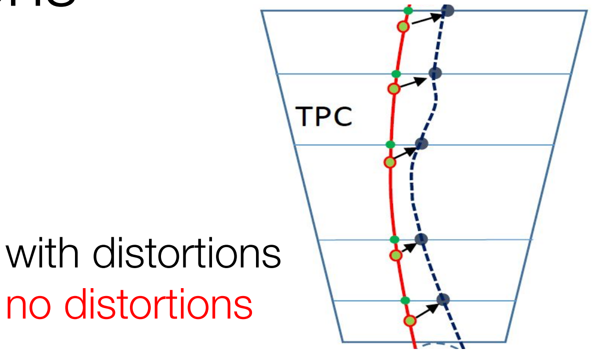

Shifted reconstructed point positions
–> worse reconstruction accuracy.

- distortions ($\vec{E}$ not uniform)
- distortion fluctuation ($\vec{E}$ not constant over time) – **most difficult** to correct

CRITICAL: No proper correction –> no precise reconstruction of particle trajectories
–> cannot perform almost **any physics analysis!**

Mean space-charge density

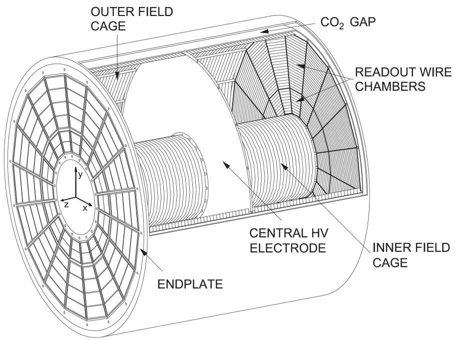

Small variations along φ

Dead zones between readout chambers

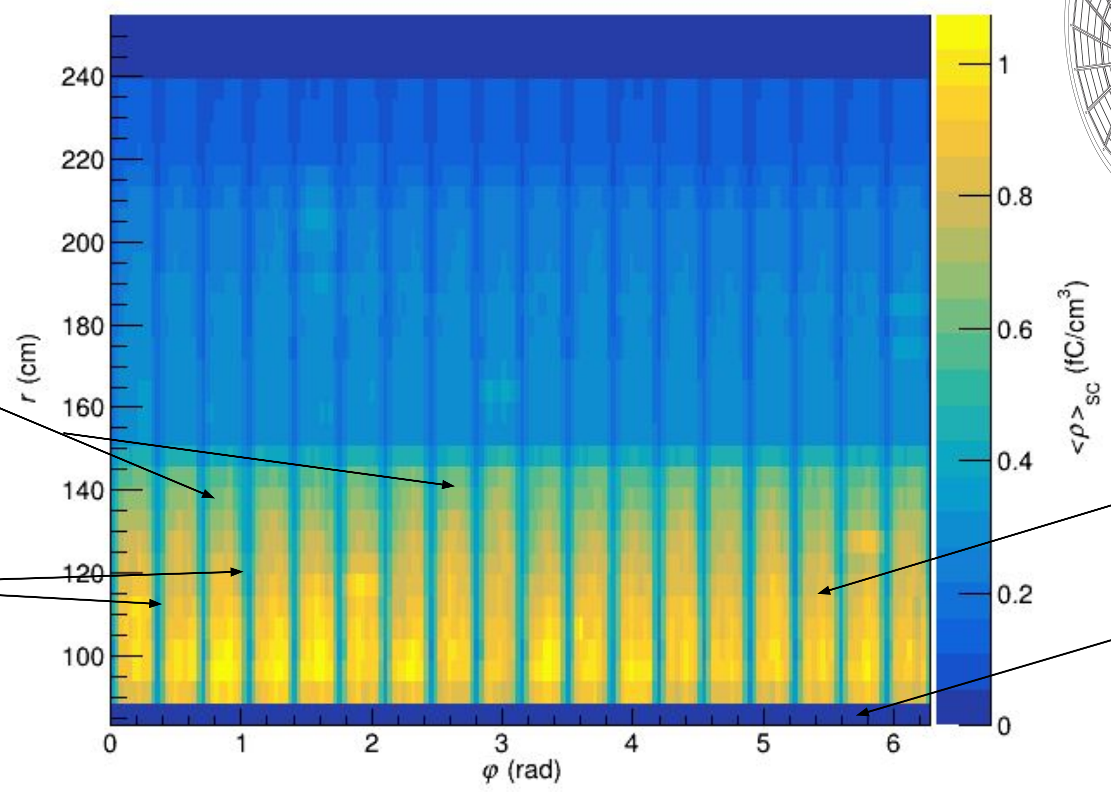

Biggest densities

Inner field cage boundary

Random space-charge density (1 scenario)

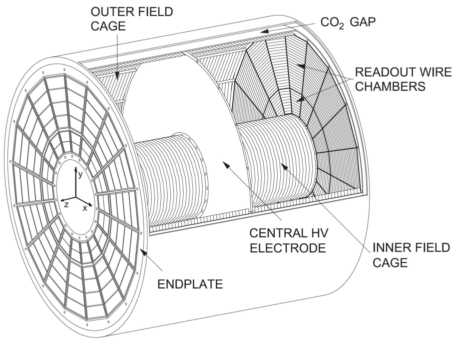

Pile-up of 8000 events

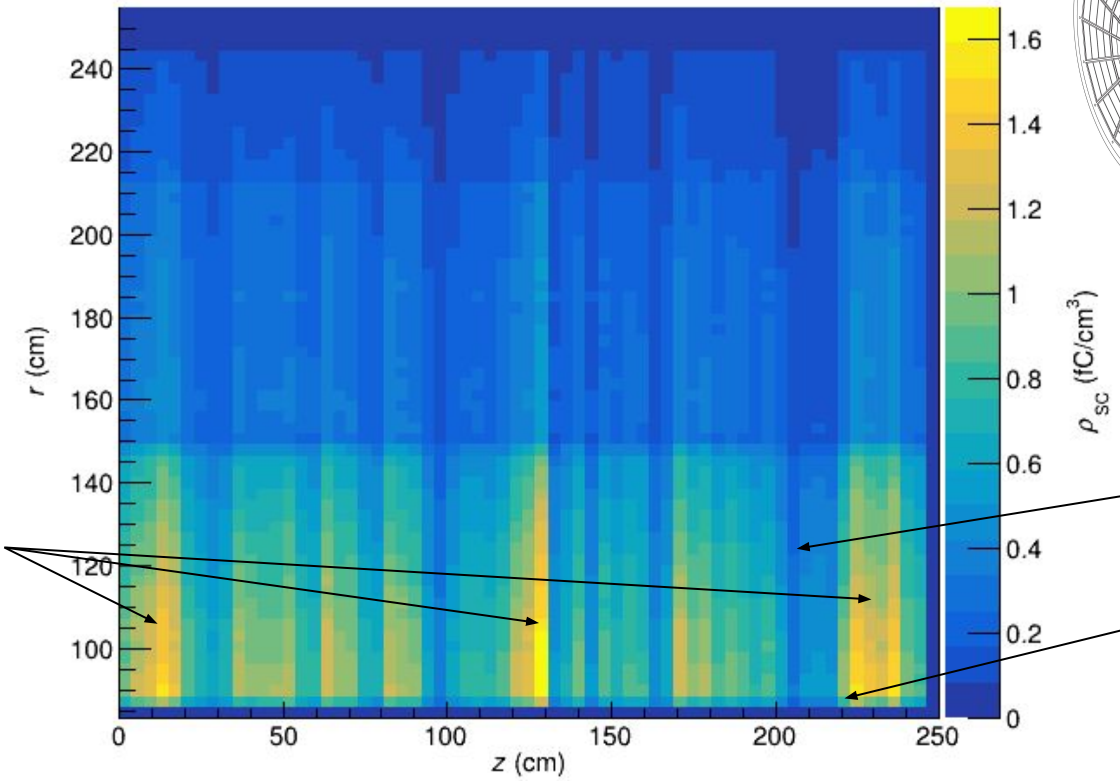

Fluctuations of the event multiplicities and of the number of pile-up events

Biggest densities

Inner field cage boundary

Space-charge distortion in r direction

biggest magnitude of distortions and distortion fluctuations (on average)

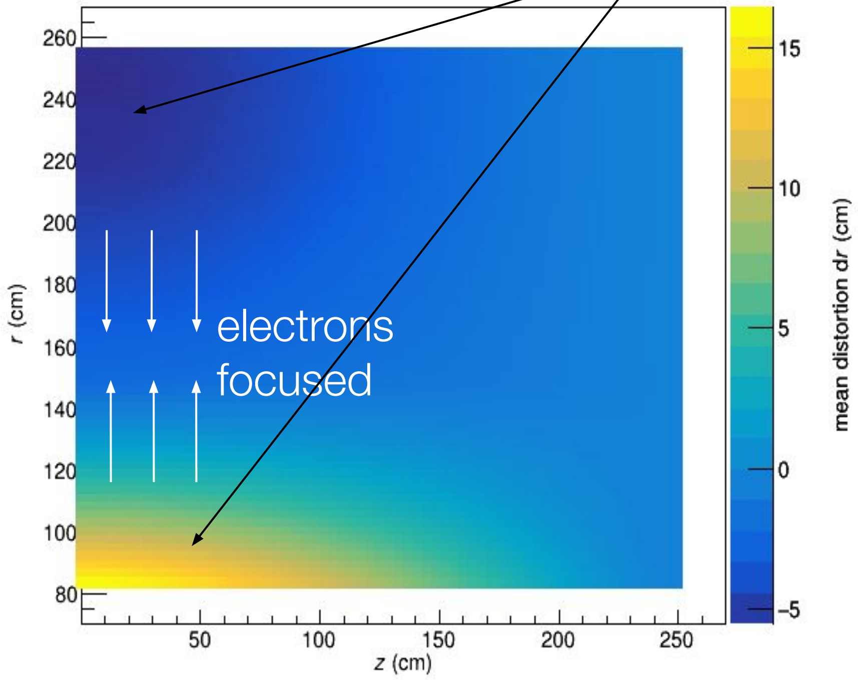

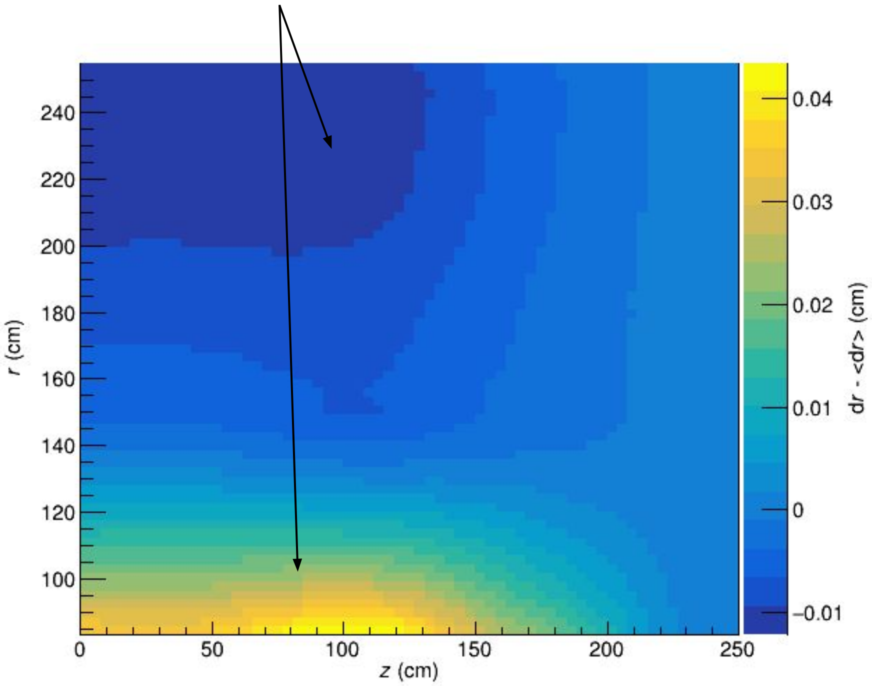

TPC distortion corrections

Problem: calculate the correction of the **distortions** and **distortion fluctuations** from the electric field modifications

Requirements:

- precision **~200 µm** (TPC resolution)
- new distortion correction for each **~5 ms** data interval

Analytic calculations: too slow and potentially less accurate

ML and DNN: effective and fast methods for correcting the **fluctuations**

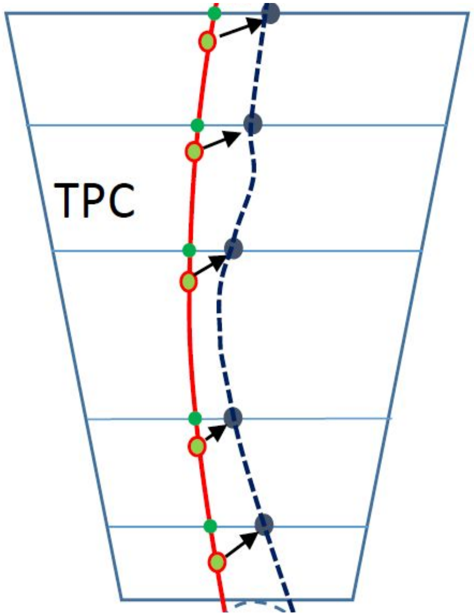

with distortions
no distortions

Machine learning solution

Global correction with Boosted Decision Trees

Global = time-dependent variation of the total ion charge.

Input: **~TPC "currents" ~ instantaneous charge densities** along *z*-axis (1D)
Output: corrections of global **distortion fluctuations** in each direction in 3D space

Note: This is a **recent addition** to the correction strategy.
Most of the results are presented for the old strategy (no BDT)

dmlc XGBoost + scikit learn

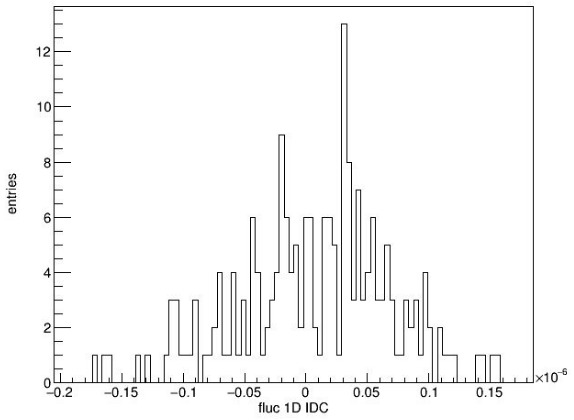

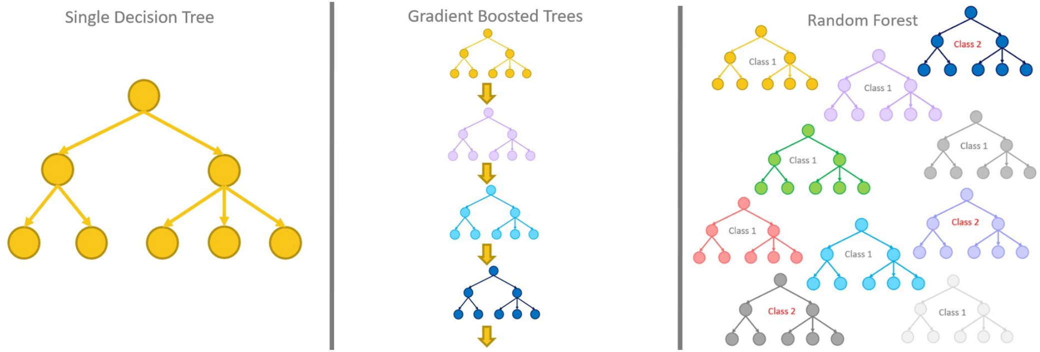

Local corrections with deep learning

Local fluctuations of the ion charge due to **ionization**.

Input: **~ion "current" in each TPC 3D point**
Output: correction of local **distortion fluctuations** in each direction in 3D space

U-Net: a convolutional neural network for biomedical image segmentation.

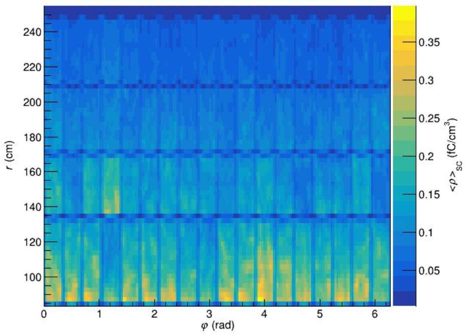

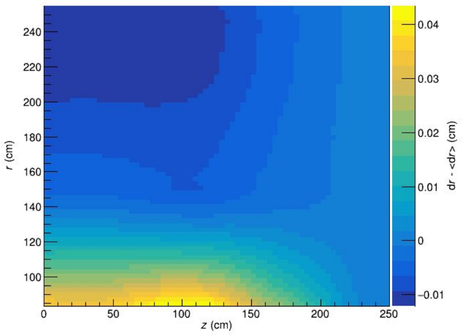

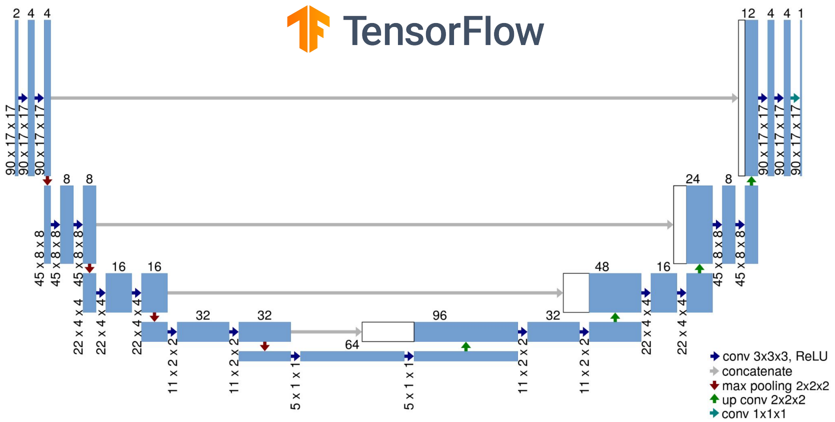

Why U-Net?

We want to predict a fluctuation **value** for **each voxel**. **Output size = input size.**
Image segmentation: assign a class to each voxel.
Difference: **continuous** output values –> another loss function

Fully Convolutional Network (FCN):

- convolutions with no dense layers
- contraction + an upsampling layer
- skip connections

U-Net:

- an upsampling path mirroring the contracting path
- learn effectively from few samples
- easily adapted to more than 2 dimensions

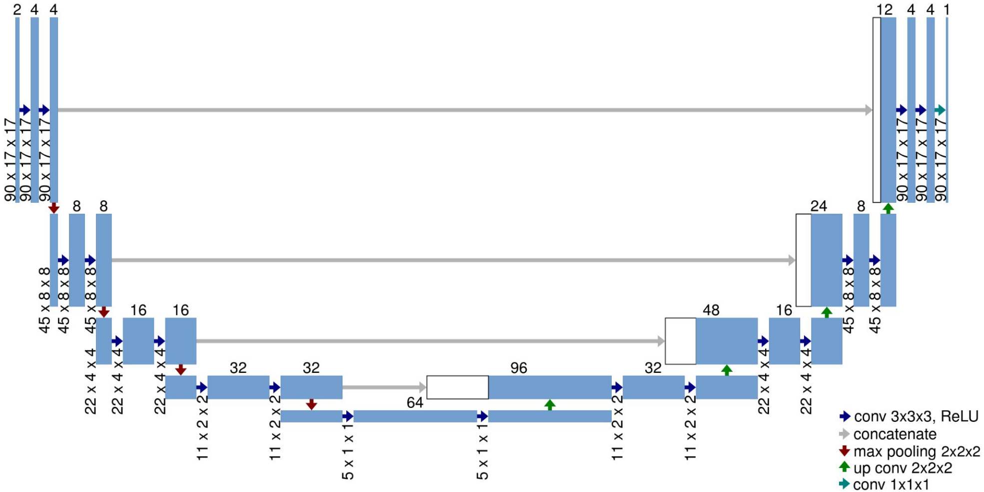

Bayesian optimization

Both algorithms have **many hyperparameters**
–> very hard to find the optimal settings

We used **Bayesian optimization** to find the best models:

- find point *x* that maximizes the current acquisition function
- evaluate objective function value *y* at *x*
- update the statistical model

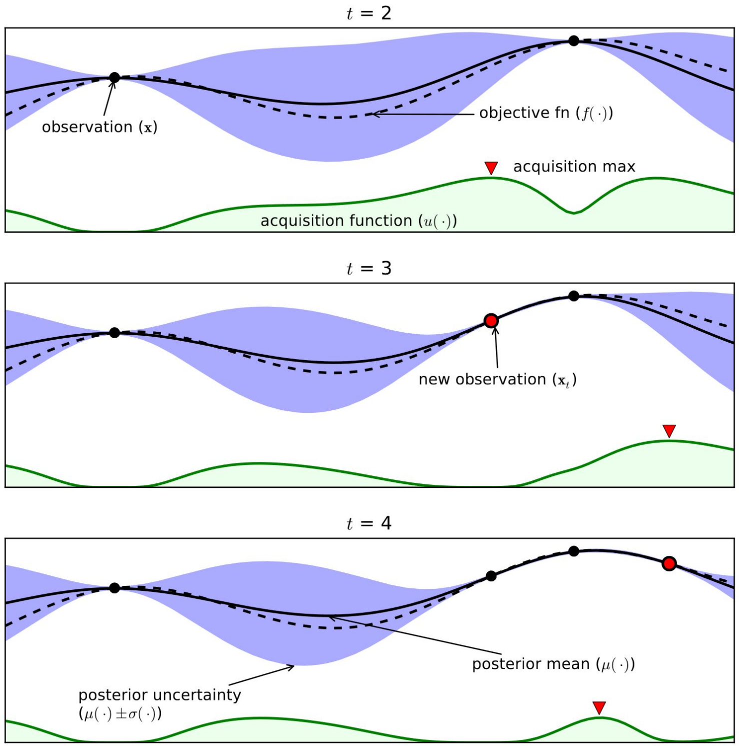

Best U-Net configuration

- dropout: none
- batch normalization: none
- depth: 4
- initial feature channels: 4
- optimizer: Adam
- learning rate: 0.001
- loss function: MSE
- epochs: 20
- batch size: 27

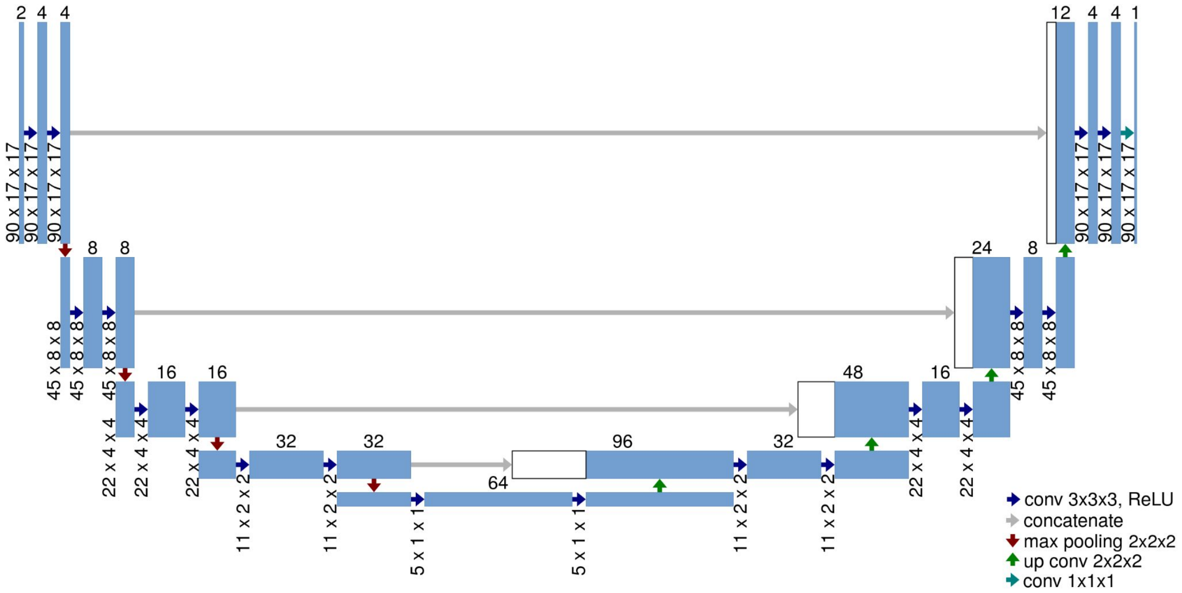

Results

Standalone DNN predictions

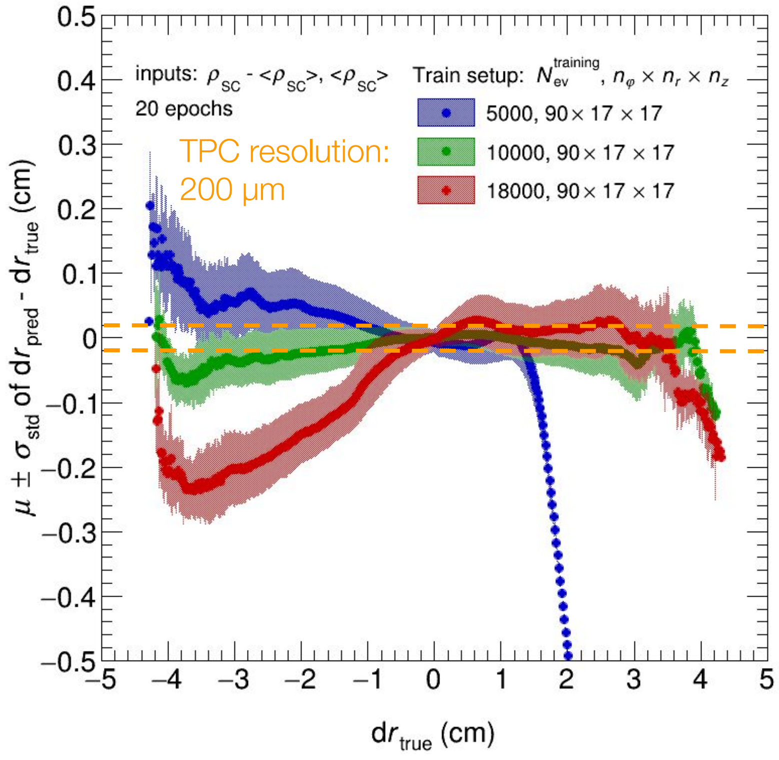

- **mean error:** predicted – expected distortion fluctuation in the r direction
- granularity = number of bins/voxels in (r, φ, z)
- orange lines – TPC resolution = level of precision required
- **best model:** 10k training scenarios, 90x17x17 – **very close** to the TPC resolution
- big errors at **largest** and **smallest** distortion fluctuations
- 5k scenarios: **undertrained**, lack of samples for big distortions
- 18k scenarios: **overtrained**

Standalone DNN predictions

- **RMSE and mean** of the error: predicted – expected distortion fluctuation in r direction
- the U-Net performs **worst** at the **smallest *r*** (biggest distortions)
- **best model:** 10k scenarios, 90x17x17
- 18k scenarios, 180x33x33 – worse performance at the smallest *r*

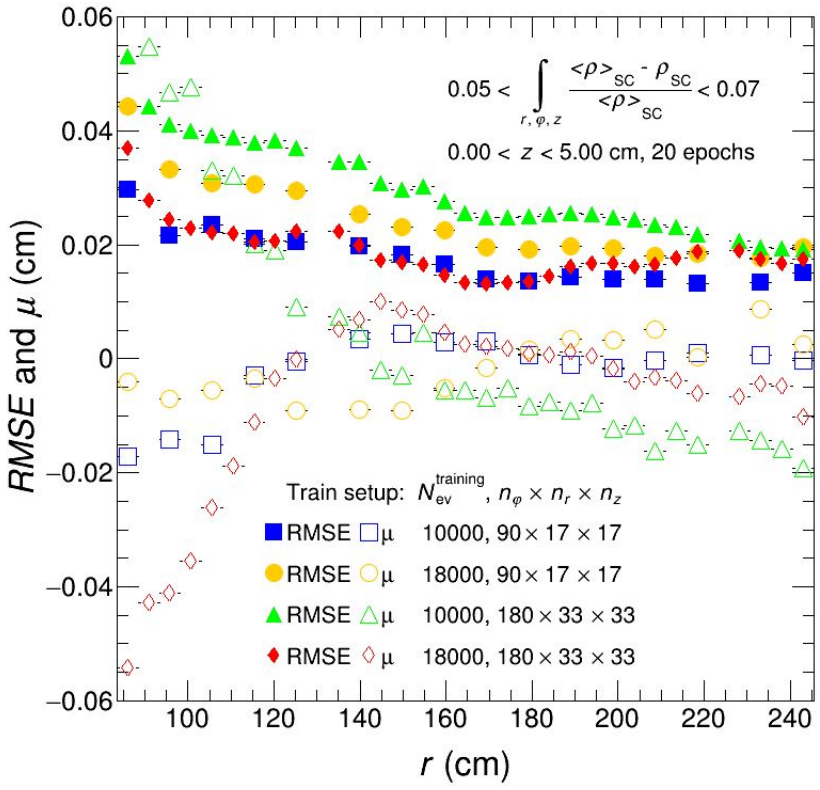

How can we improve

Top plot: numerical calculation of the distortion fluctuations caused by a single line charge.
Bottom plot: CNN predictions.

The problem:

The UNet tries to learn the effects of **non-symmetric boundary conditions** (field cage, readout chambers, central electrode).

Reduce the boundary effects for the neural network:

- precede CNN 3D–>3D correction with data-driven 1D–>3D BDT distortion fluctuations correction
- use data-driven input (numerical derivative of the corrections)

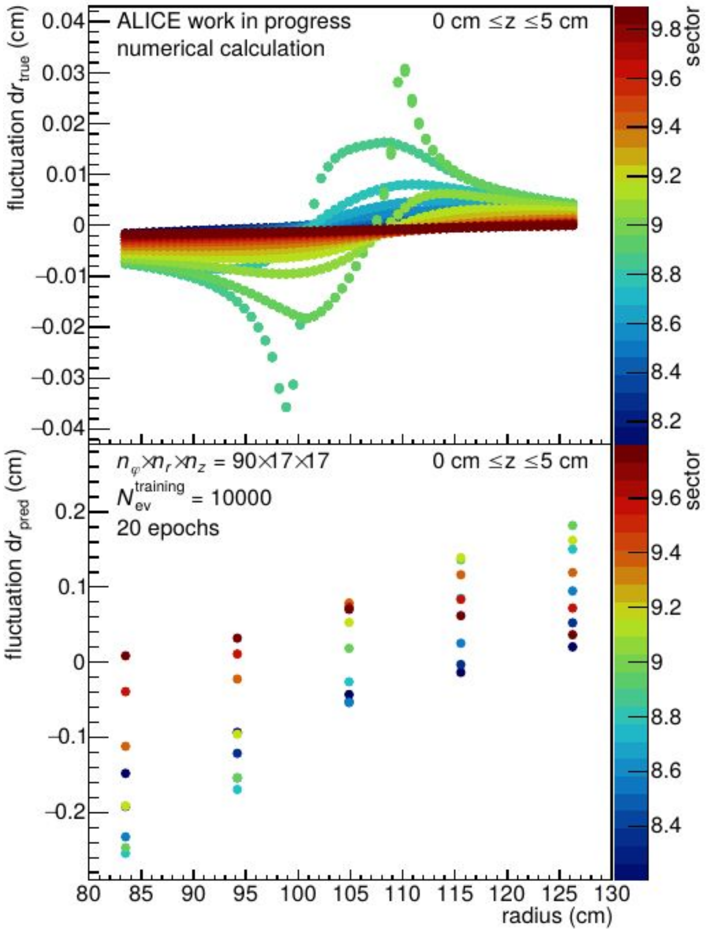

Recent developments: BDTs for global correction

- **Bayesian optimization** is used to find the best model configuration.
- best train performance after ~100 events **but still quite big validation loss:** 0.04 cm (expected: around 0.002 cm)
- **very fast: < 1 min** training time

More analysis ongoing:

- refinement of the choice of the training features

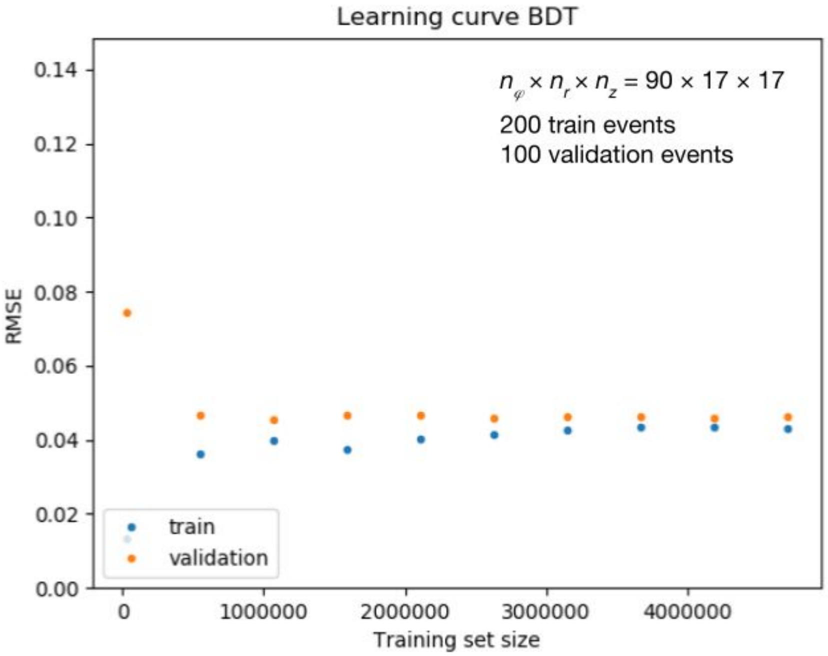

Performance benchmarks

- **AMD Vega20, 16 GB**, 60 cores, HDD disk
- **NVIDIA Tesla V100, 32 GB**, SSD and HDD disks
- measured time spent on **training** alone
- **NVIDIA is 4-5 times faster**, even with slow HDD disk

CPU Time (min)	Configuration				
	AMD 16 GB HDD	NVIDIA 32 GB SSD	NVIDIA 16 GB SSD	NVIDIA 32 GB HDD	NVIDIA 16 GB HDD
user	25:08	7:26	7:23	9:11	9:15
system	24:31	1:41	1:35	3:16	3:26

Conclusions

Summary

What?

- correction of the distortion fluctuations with ML and DNN
- required precision **~200 µm**, **5-ms** data intervals

Why?

- **almost no physics analysis** can be done without the correction!

How?

- global correction: Boosted Decision Trees
- local correction: U-Net – a deep convolutional neural network

Results

- **BDTs** can be trained and applied **very quickly** (work ongoing)
- standalone **DNN** can achieve **the required precision** except for the boundaries
- DNN benchmarks performed with different GPUs

Further work

- train and test on full BDT + U-Net pipeline
- train and test in other directions (z, $r\varphi$) as well
- parallel training of U-Net on multiple GPU cores on Google Cloud (ongoing):
  - part of **the first pilot project** for cloud LHC computing
- develop a strategy to re-train existing DNN models

March 2022: launch of LHC Run 3 – the correction algorithm used right from the start!

© Klaus Barth

Thank you for your attention!

Backup

Strategy for Run 3 corrections

Synchronous

- average distortions (~1 min):
  - pre-calculated correction maps
- **global** distortion fluctuations:
  - Boosted Decision Trees
  - input:
    - 1D IDCs fluctuations
    - numerical derivative of average corrections
- precision: ~**mm**

Analytical corrections

- **too much time** (O(1s - 1min))
- **assume ideal TPC**

Asynchronous

- average distortions (~1 min):
  - correction maps from the data
  - ITS-TRD interpolation
- **global** distortion fluctuations:
  - same as at synchronous stage
- **local** distortion fluctuations:
  - 3D neural network
  - network input:
    - 3D IDCs fluctuations
    - numerical derivative of average corrections
- new distortion correction ~**5ms**
- precision: ~**200µm** (TPC resolution)