Machine Learning Processing and Compression of Signal Shared AC-LGADs

Don C. Wong[1]

on behalf of

Gaetano Barone[1], Diallo Boye[2], Gabriele Giacomini[2], Vagelis Gkougkousis[3], Rohan Kher[1], Ben Kilminster[3], Anna Macchiolo[3], Samantha Sunnarborg[1], Jessica Tang[1], Alessandro Tricoli[2]

Brown University[1], Brookhaven National Laboratory[2] ,University of Zurich[3]

Introduction and Motivation

We use machine learning to extract maximal information from AC-LGAD waveforms to predict particle hit coordinates, with the goal of outperforming matrix inversion baselines.

- Traditional methods (analytic charge-sharing and matrix inversion methods) use reduced summaries of waveforms (e.g. peak amplitudes or relative amplitudes between channels).
- Neural networks can take in full waveforms output by AC-LGADs, and learn nonlinear patterns that capture the effects of small variations in timing and amplitude.
- Project conducted under the proposal *Advancing the Pixelated Resistive Silicon Readout and Charge Collection Techniques*, PI: Gaetano Barone

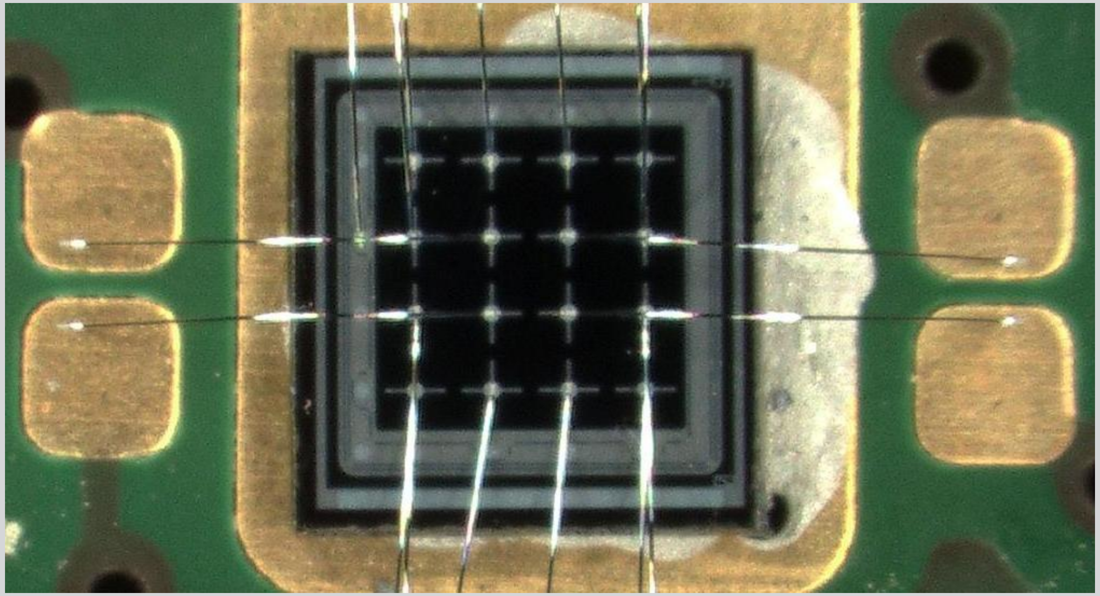

BNL-Manufactured AC-LGAD

Active thickness: 30 μm

Pad size: 200 μm

Pitch: 500 μm

2x2 pads readout

Recurrent Neural Network

RNNs are currently outperforming traditional methods for position reconstruction.

- RNNs model dependencies on past time steps, making them optimal for time-series data, such as waveforms read out by AC-LGADs.
- A **Long Short-Term Memory** (LSTM) layer serves as the core temporal-processing unit. For long sequences and small timesteps, models with LSTM layers are more stable than standard RNNs, as LSTMs more effectively capture the longer time dependencies present in waveforms.

Layer (type)	Output Shape	Param #
input_layer (InputLayer)	(None, 4, 350)	0
lstm (LSTM)	(None, 8)	11,488
dense (Dense)	(None, 16)	144
dense_1 (Dense)	(None, 16)	272
dense_2 (Dense)	(None, 16)	272
dense_3 (Dense)	(None, 2)	34

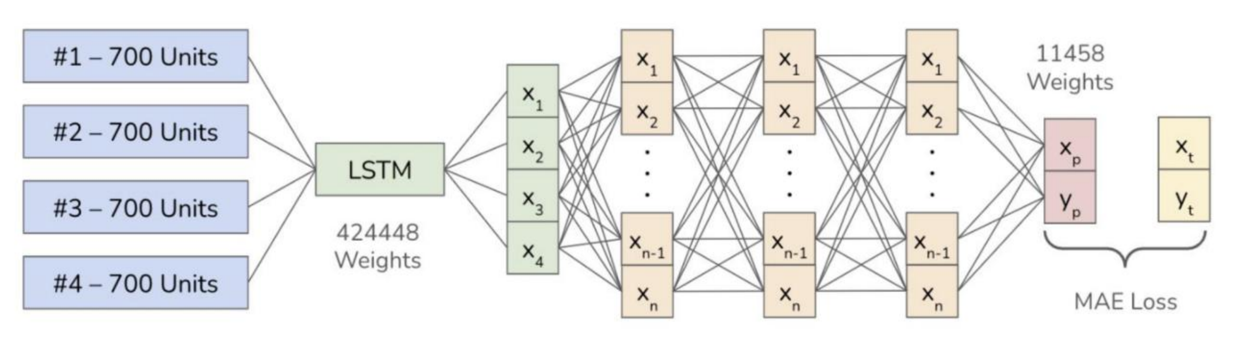

Top: Example of RNN architecture used for our reconstruction analysis.

Bottom: Schematic of RNN architecture.

Transient Current Technique

We use the Transient Current Technique (TCT) lasers to obtain waveforms.

- 385 grid positions, 25 × 25 μm^2 spacing, 111 triggers per position
- 1064 nm laser, laser spot Gaussian with σ ~ 9 μm, ~1 μm spatial resolution, two pulses separated by 100 ns
- Currently taking more data at the University of Zurich at bias voltages of 60, 70, 80, 90, 100, 110, 120

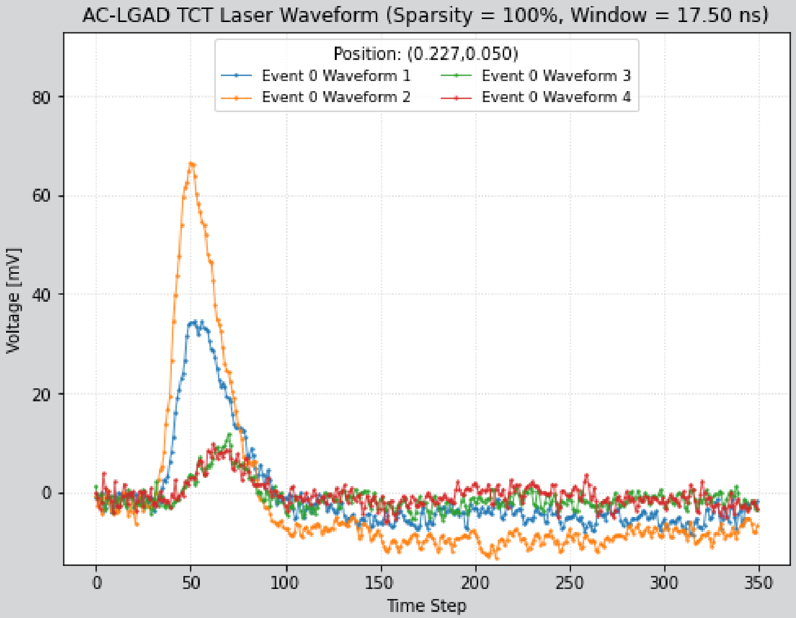

Waveforms from 4-channel AC-LGAD from TCT. Charge sharing exemplified in waveform shape, correlated to hit-to-pad proximity.

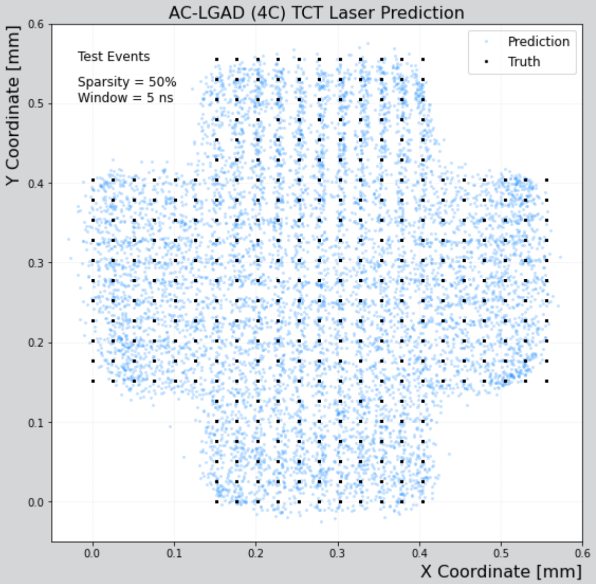

TCT grid of predicted and actual positions.

2025 UZH TCT

Goal: Comprehensive TCT waveform dataset for better neural network training.

Process

- Examined AC-LGAD wire bonds under microscope
- X-Y alignment: LED for general alignment, then stages and motors
- Z alignment: macro script to adjust height of laser while measuring voltage to find optimal distance

Manual IV: Determined rough breakdown voltage at 130V. Taking measurements from 60-120V at +10V increments.

Trial 2D scan

- 20mm increments
- 300 waveforms per points
- 60 V bias

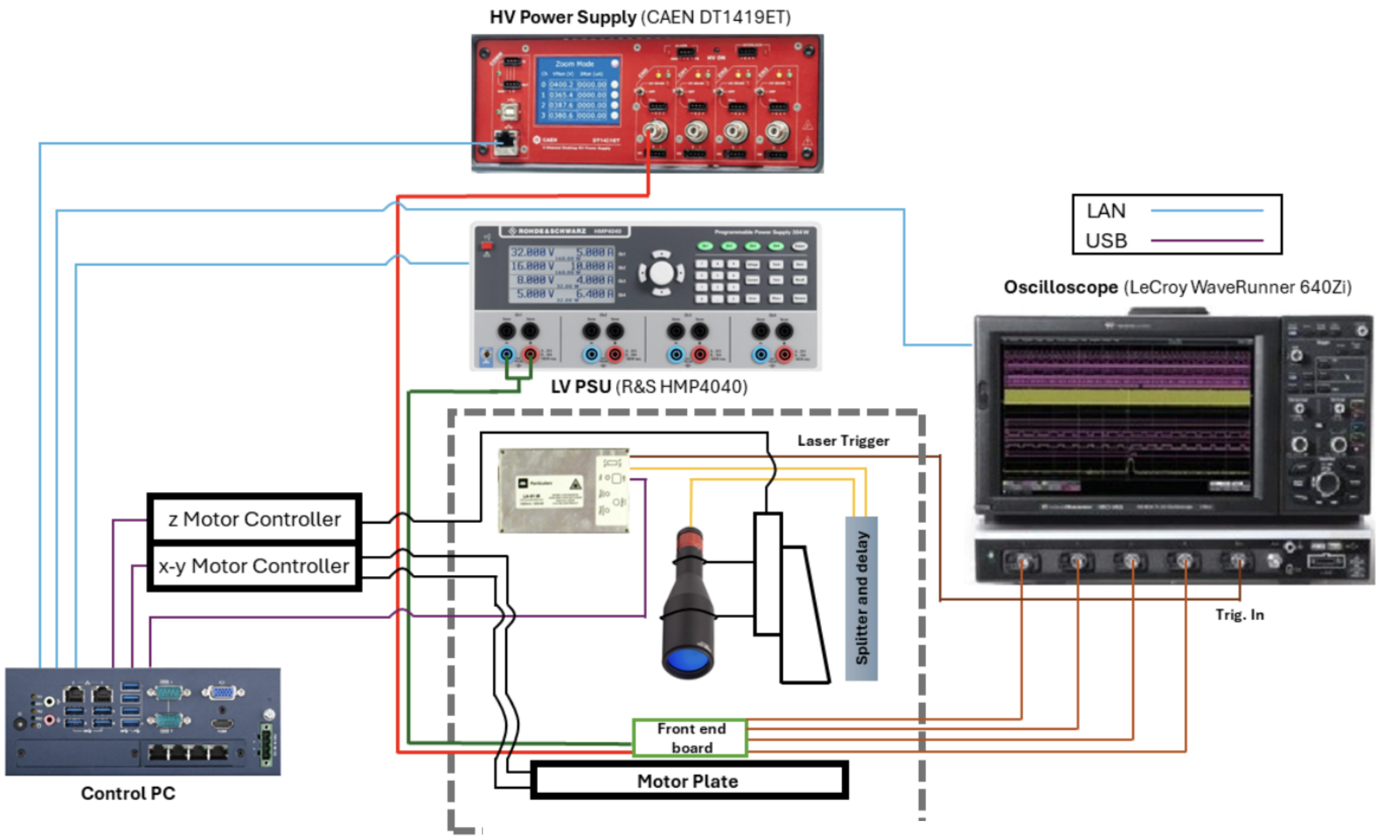

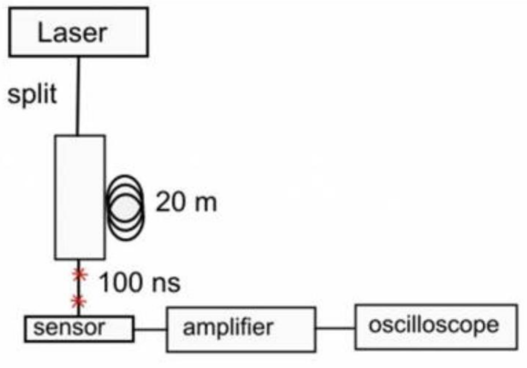

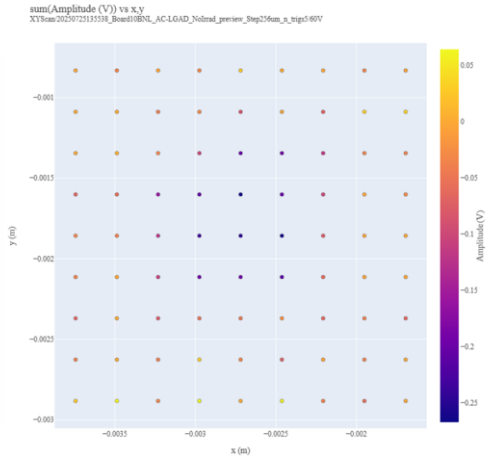

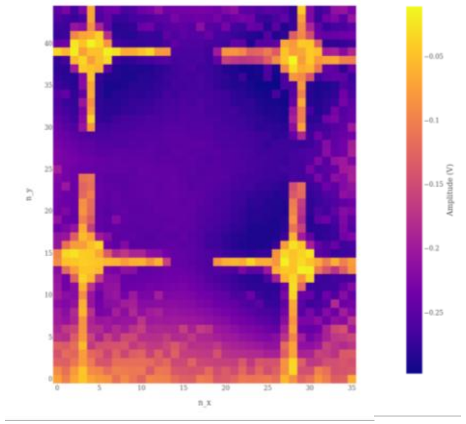

RNN on TCT Results

With full waveform processing, we achieve a 50% improvement over current ML methods and a 7-fold improvement over binary readout.

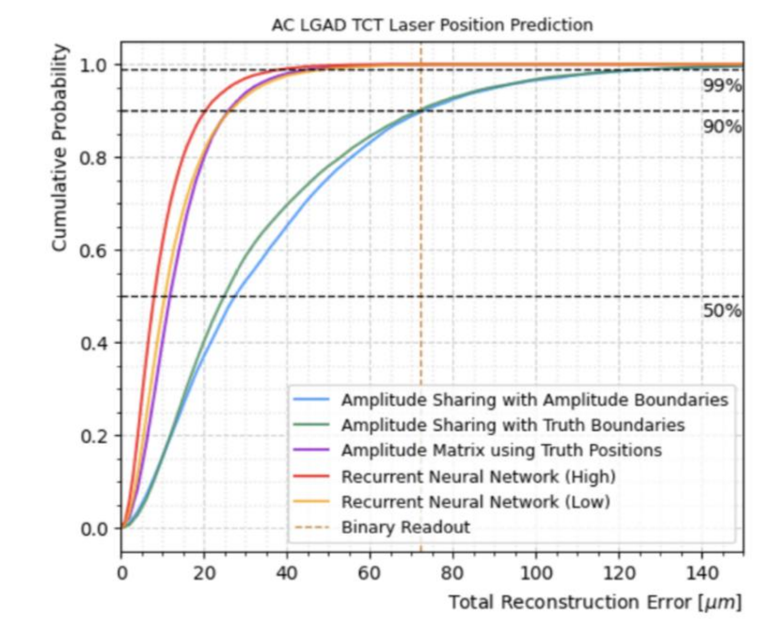

Indicates a potential resolution of ~10μm from pixels with 500 μm x 500 μm.

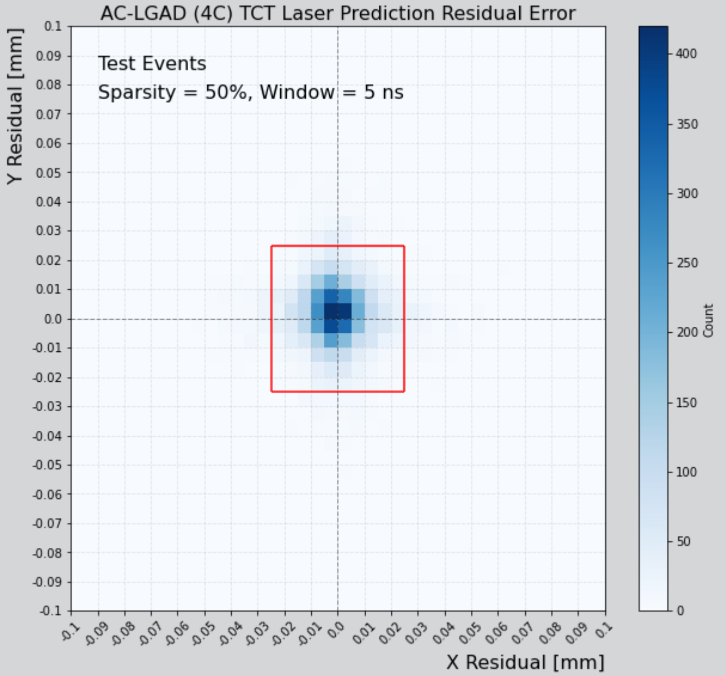

Accurate and precise enough to distinguish two grid points.

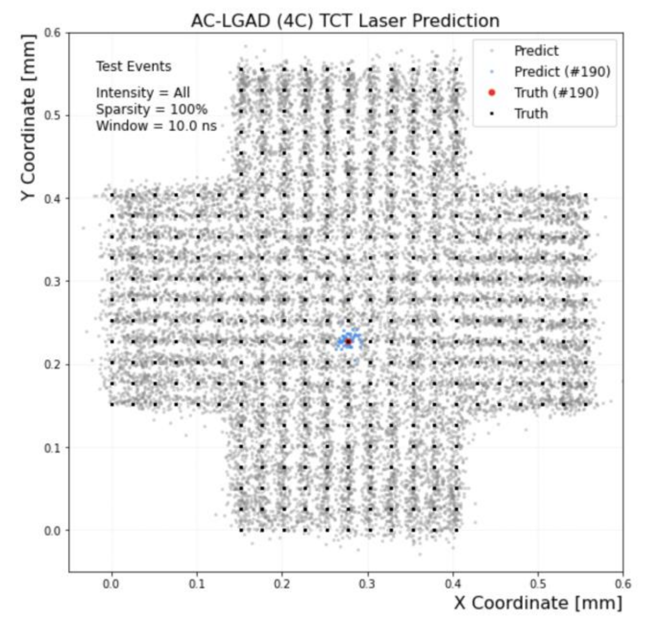

Central regions perform better than the edges.

CERN SPS H6 Test Beam (June/October Campaign)

Goal: Comprehensive test beam dataset for non-irradiated AC-LGAD for neural network training.

- CERN SPS H6 beam, 120 GeV pions
- Chubut 2, 4 channels readout board
- Mimosa telescope
- 100V Bias (June) and 120V Bias (October)
- CAEN DT5742 digitizer, 500 MHz @ 5 GS/s
- Cold box for irradiated DUTs @ -12°C (AC-LGAD was non-irradiated)

Special thanks to the DRD3 community for continued collaboration in test beam campaigns!

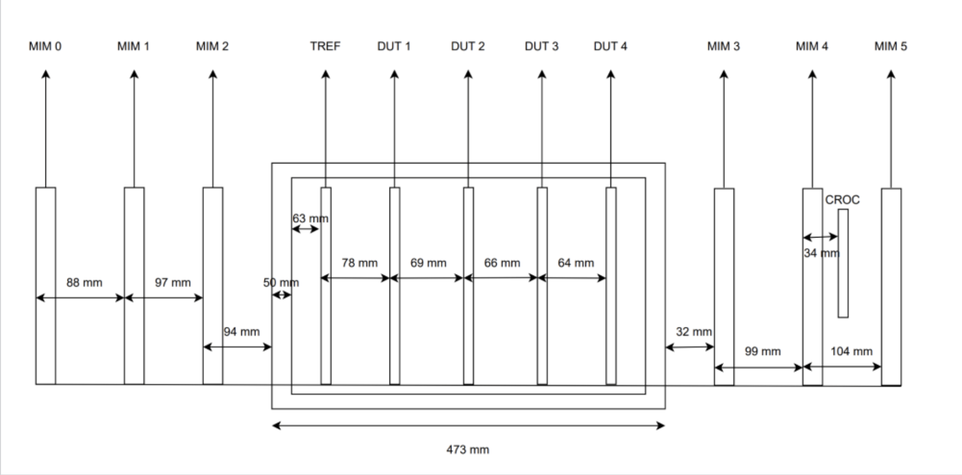

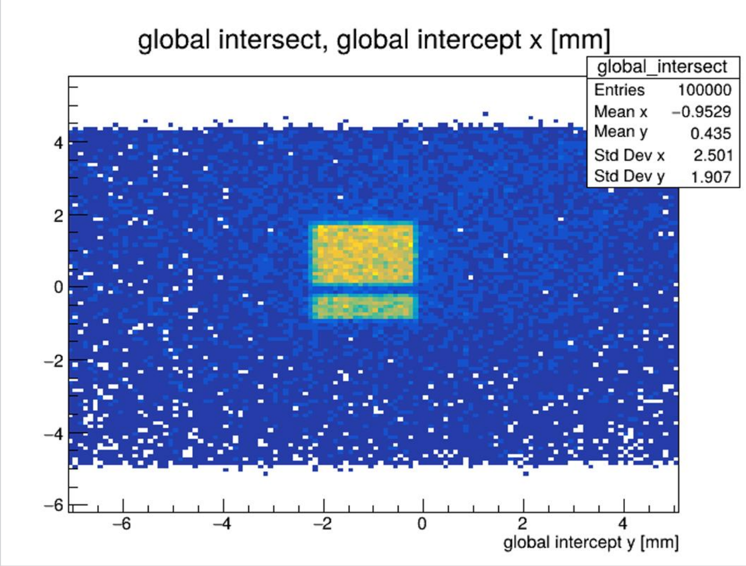

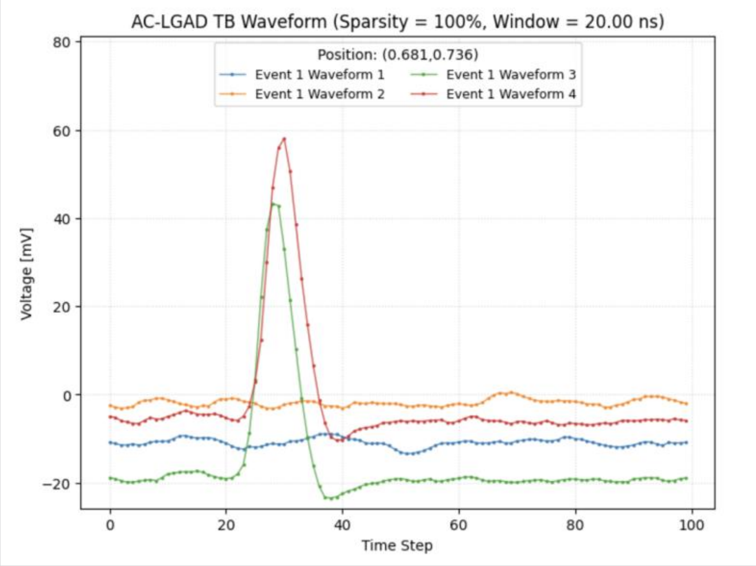

Corryvreckan Modifications for Test Beam Data

We use a custom pipeline to process and reconstruct events. Position data is obtained with [TreeWriterDUT].

Pipeline:

1. Masking run using [MaskCreator] module to mask bad pixels.
2. Prealignment run to align telescope.
3. Final alignment run to align DUTs.
4. Custom [TreeWriterDUT] module for position data.

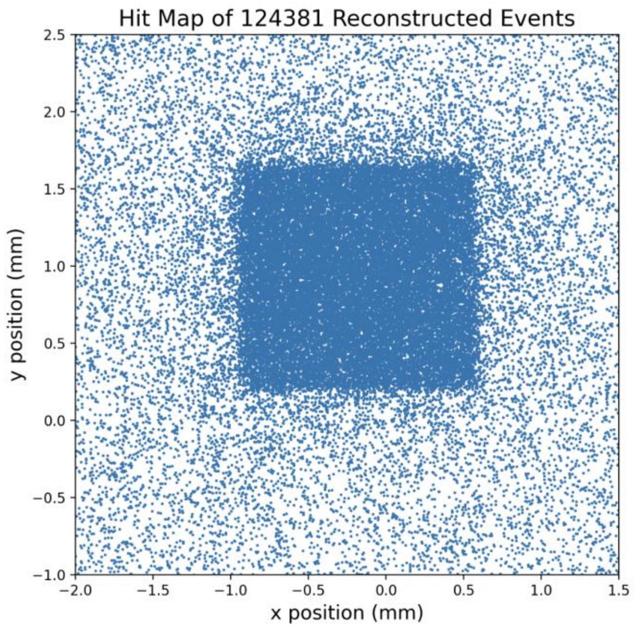

[TreeWriterDUT] selects clustered hits from the DUT plane to deliver positions instead of using a generalized reconstructed track as used in [TreeWriterTracks].

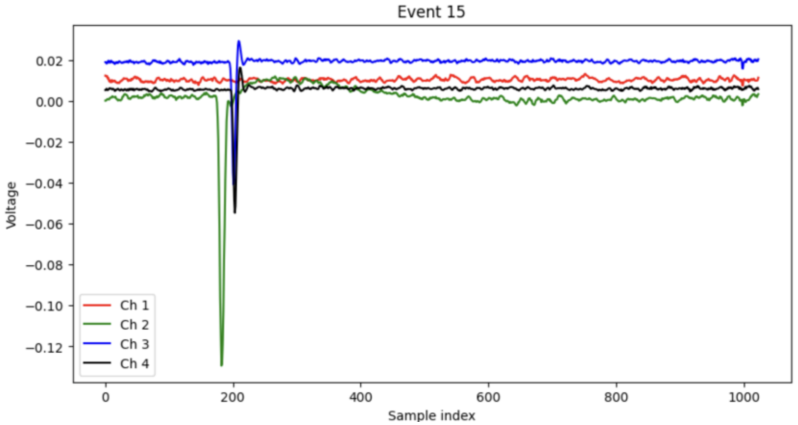

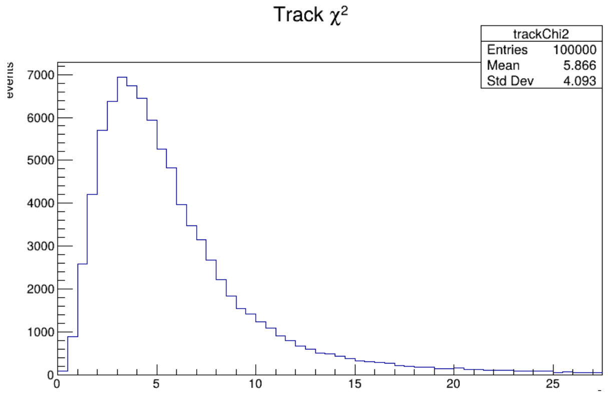

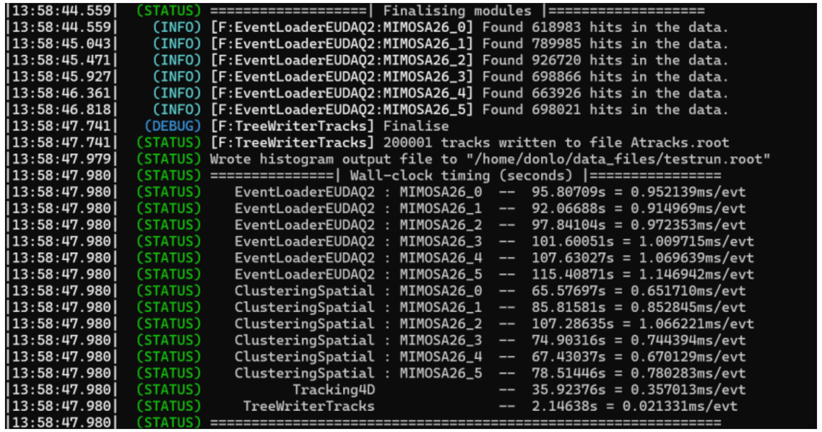

```
|13:58:44.559|  (STATUS) ===================| Finalising modules |===================
|13:58:44.559|    (INFO) [F:EventLoaderEUDAQ2:MIMOSA26_0] Found 618983 hits in the data.
|13:58:45.043|    (INFO) [F:EventLoaderEUDAQ2:MIMOSA26_1] Found 789985 hits in the data.
|13:58:45.471|    (INFO) [F:EventLoaderEUDAQ2:MIMOSA26_2] Found 926720 hits in the data.
|13:58:45.927|    (INFO) [F:EventLoaderEUDAQ2:MIMOSA26_3] Found 698866 hits in the data.
|13:58:46.361|    (INFO) [F:EventLoaderEUDAQ2:MIMOSA26_4] Found 663926 hits in the data.
|13:58:46.818|    (INFO) [F:EventLoaderEUDAQ2:MIMOSA26_5] Found 698021 hits in the data.
|13:58:47.741|   (DEBUG) [F:TreeWriterTracks] Finalise
|13:58:47.741|  (STATUS) [F:TreeWriterTracks] 200001 tracks written to file Atracks.root
|13:58:47.979|  (STATUS) Wrote histogram output file to "/home/donlo/data_files/testrun.root"
|13:58:47.980|  (STATUS) ===============| Wall-clock timing (seconds) |================
|13:58:47.980|  (STATUS)    EventLoaderEUDAQ2 : MIMOSA26_0  --  95.80709s = 0.952139ms/evt
|13:58:47.980|  (STATUS)    EventLoaderEUDAQ2 : MIMOSA26_1  --  92.06688s = 0.914969ms/evt
|13:58:47.980|  (STATUS)    EventLoaderEUDAQ2 : MIMOSA26_2  --  97.84104s = 0.972353ms/evt
|13:58:47.980|  (STATUS)    EventLoaderEUDAQ2 : MIMOSA26_3  --  101.60051s = 1.009715ms/evt
|13:58:47.980|  (STATUS)    EventLoaderEUDAQ2 : MIMOSA26_4  --  107.63027s = 1.069639ms/evt
|13:58:47.980|  (STATUS)    EventLoaderEUDAQ2 : MIMOSA26_5  --  115.40871s = 1.146942ms/evt
|13:58:47.980|  (STATUS)    ClusteringSpatial : MIMOSA26_0  --  65.57697s = 0.651710ms/evt
|13:58:47.980|  (STATUS)    ClusteringSpatial : MIMOSA26_1  --  85.81581s = 0.852845ms/evt
|13:58:47.980|  (STATUS)    ClusteringSpatial : MIMOSA26_2  --  107.28635s = 1.066221ms/evt
|13:58:47.980|  (STATUS)    ClusteringSpatial : MIMOSA26_3  --  74.90316s = 0.744394ms/evt
|13:58:47.980|  (STATUS)    ClusteringSpatial : MIMOSA26_4  --  67.43037s = 0.670129ms/evt
|13:58:47.980|  (STATUS)    ClusteringSpatial : MIMOSA26_5  --  78.51446s = 0.780283ms/evt
|13:58:47.980|  (STATUS)           Tracking4D               --  35.92376s = 0.357013ms/evt
|13:58:47.980|  (STATUS)     TreeWriterTracks               --  2.14638s = 0.021331ms/evt
|13:58:47.980|  (STATUS) ================================================================
```

RNN on SPS H6 Test Beam (2025) Results

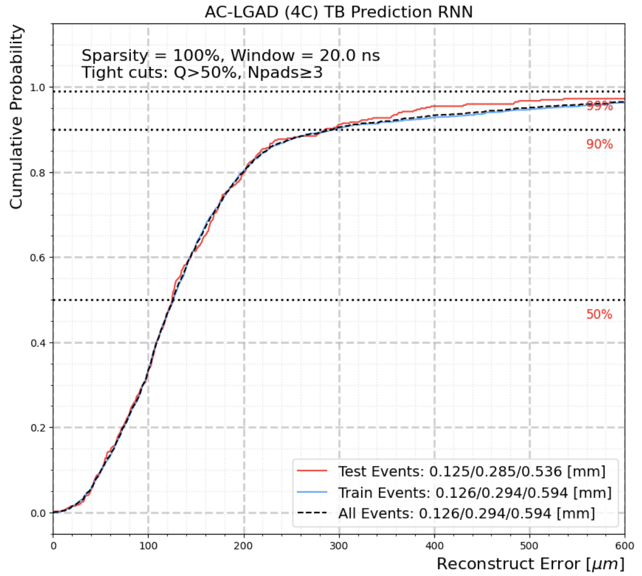

We ran the identical pipeline used for TCT on our test beam data.

October TB – fresh results, not fully digested yet (!)

- Only 2 or 3 channels are considered here.
- Partial statistics

Known issues we are working on:

- Low efficiency: sensors were underbiased
- Alignment: No default channel-to-event number alignment in the waveform, as opposed to our TCT setup.
- RNN is too small: too few LSTM layers.
- Technicalities in extracting the waveforms from the framework

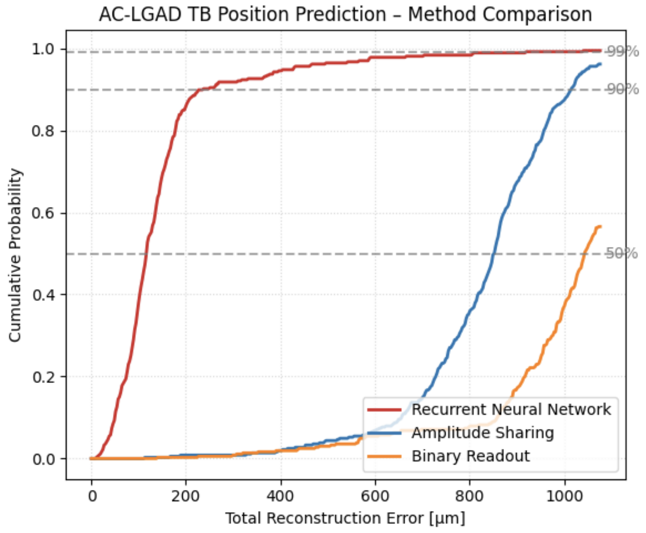

Current work:

- Improved alignment of waveforms and position
- Expand LSTM layer to 64 hidden units and more dense layers
- Laser-aided training: adaptation domain of laser data with MIPs
- Comparing with CNNs.
- Comparing higher bias voltages from October's DR3D campaign.

Summary and Outlook

Proof of concept: RNN improves hit reconstruction with potential resolution of ~10µm.

- TCT resolution improves from **~25µm** (matrix inversion) **to ~10µm** (RNN)
- First attempt at MIP-based training with full waveform regression/compression.
- Reorganizing RNN pipeline to process test beam data.

Ongoing and future work:

- Finish running RNN pipeline on 2025 test beam data and 2025 TCT data
- Using feed-forward DNN to reproduce previous results for 4 channels and compare with the performance for 16 channels
- Reconstruction with triangular different pad patterns

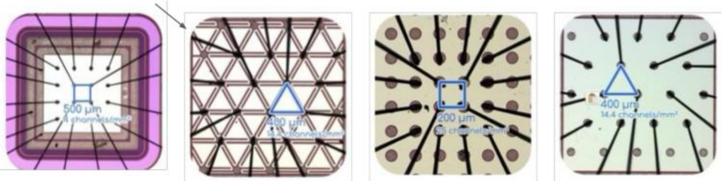

Also see
Studies on ML processing and compression of signal shared AC-LGADs (Tang et al. 2025) 3rd DRD Week
Machine Learning Processing of BNL AC-LGAD Sensors Readout with Signal Sharing (Li et al. 2025) 1st DRD Week

Additional Material

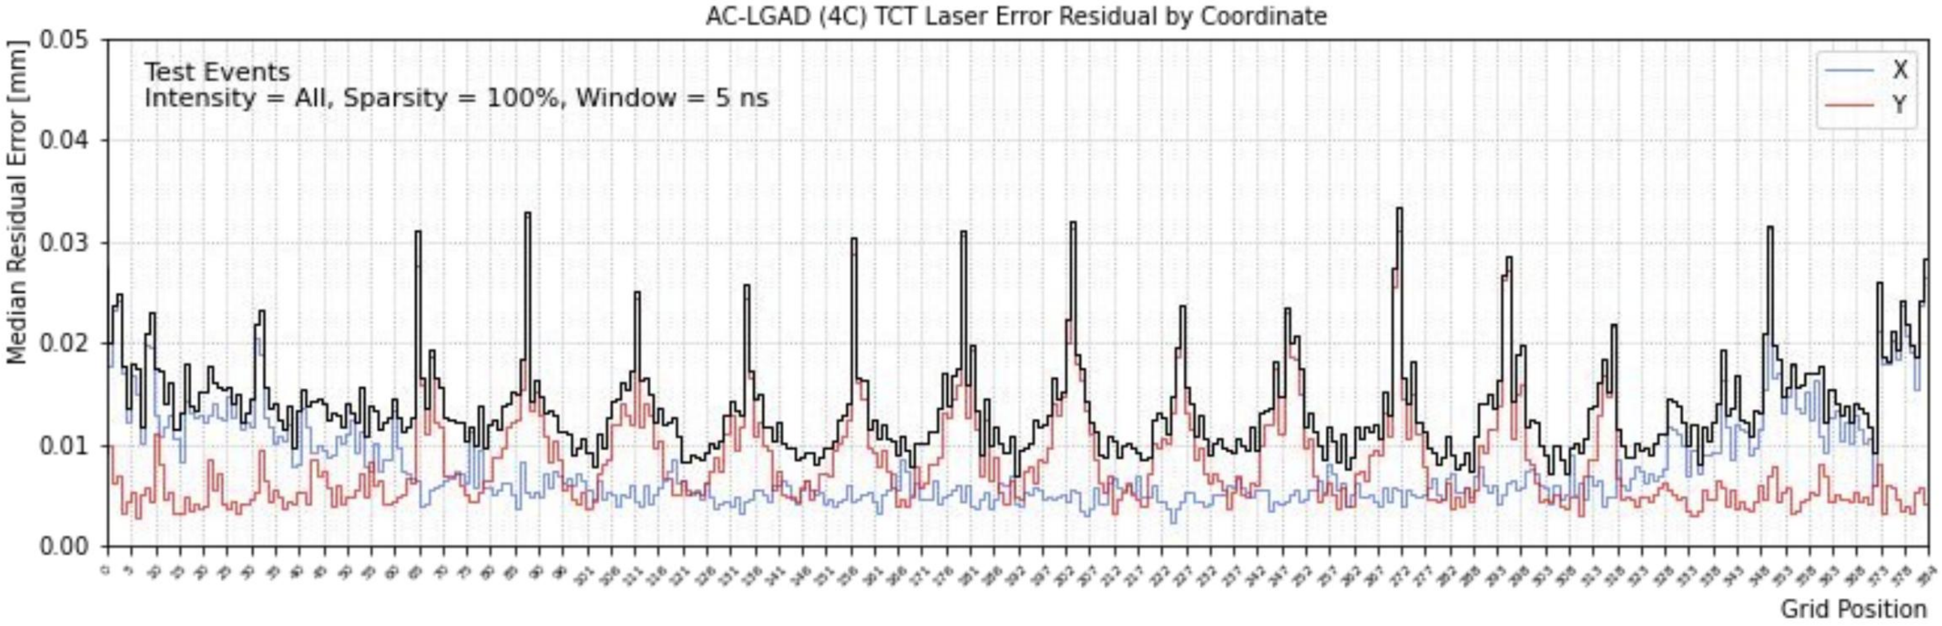

AC-LGAD (4C) TCT Laser Error Residual by Coordinate
Test Events
Intensity = All, Sparsity = 100%, Window = 5 ns
X
Y
Median Residual Error [mm]
0.05
0.04
0.03
0.02
0.01
0.00
Grid Position

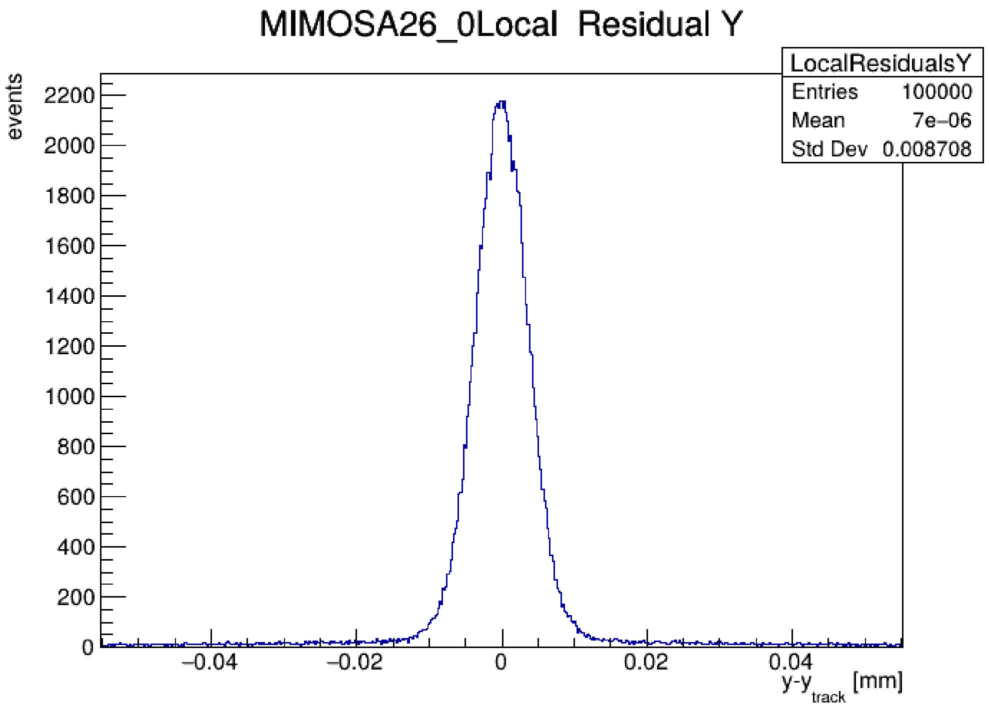
MIMOSA26_0Local Residual Y
LocalResidualsY
Entries 100000
Mean 7e−06
Std Dev 0.008708
events
y-y$_{track}$ [mm]

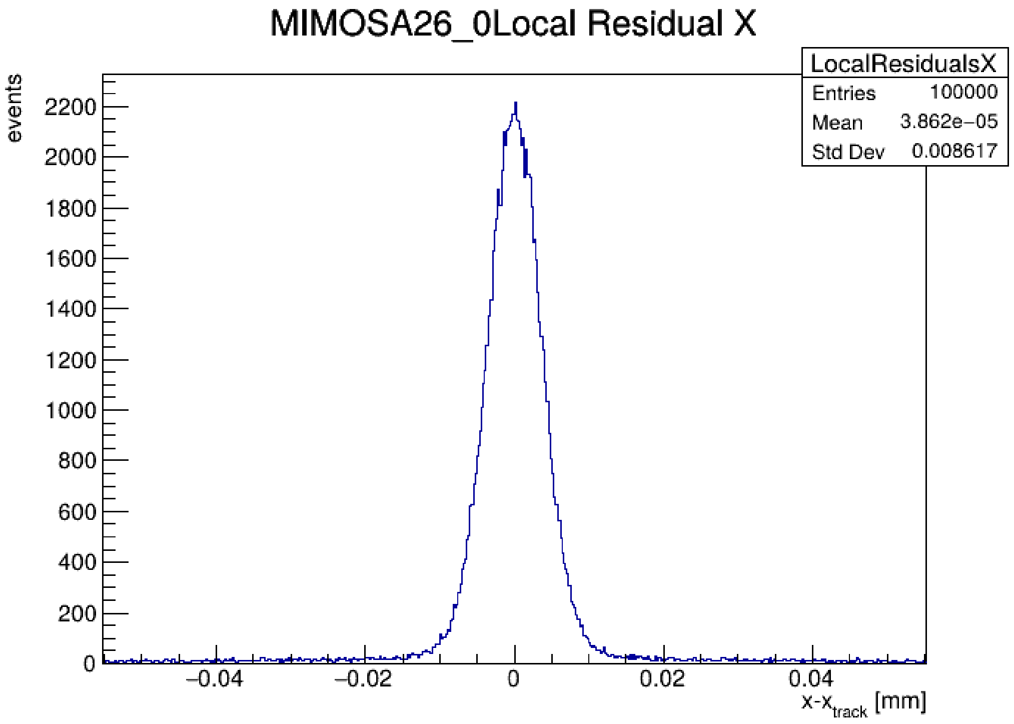
MIMOSA26_0Local Residual X
LocalResidualsX
Entries 100000
Mean 3.862e−05
Std Dev 0.008617
events
x-x$_{track}$ [mm]