Towards a Real Time ML Solution for Quench Characterization

MDP Collaboration Meeting

2 May 2024

Joseph DiMarco, <u>Maira Khan</u>, Steve Krave, Vittorio Marinozzi, Jennifer Ngadiuba, Stoyan Stoynyev, Nhan Tran

Real Time Characterization of Ramp Events

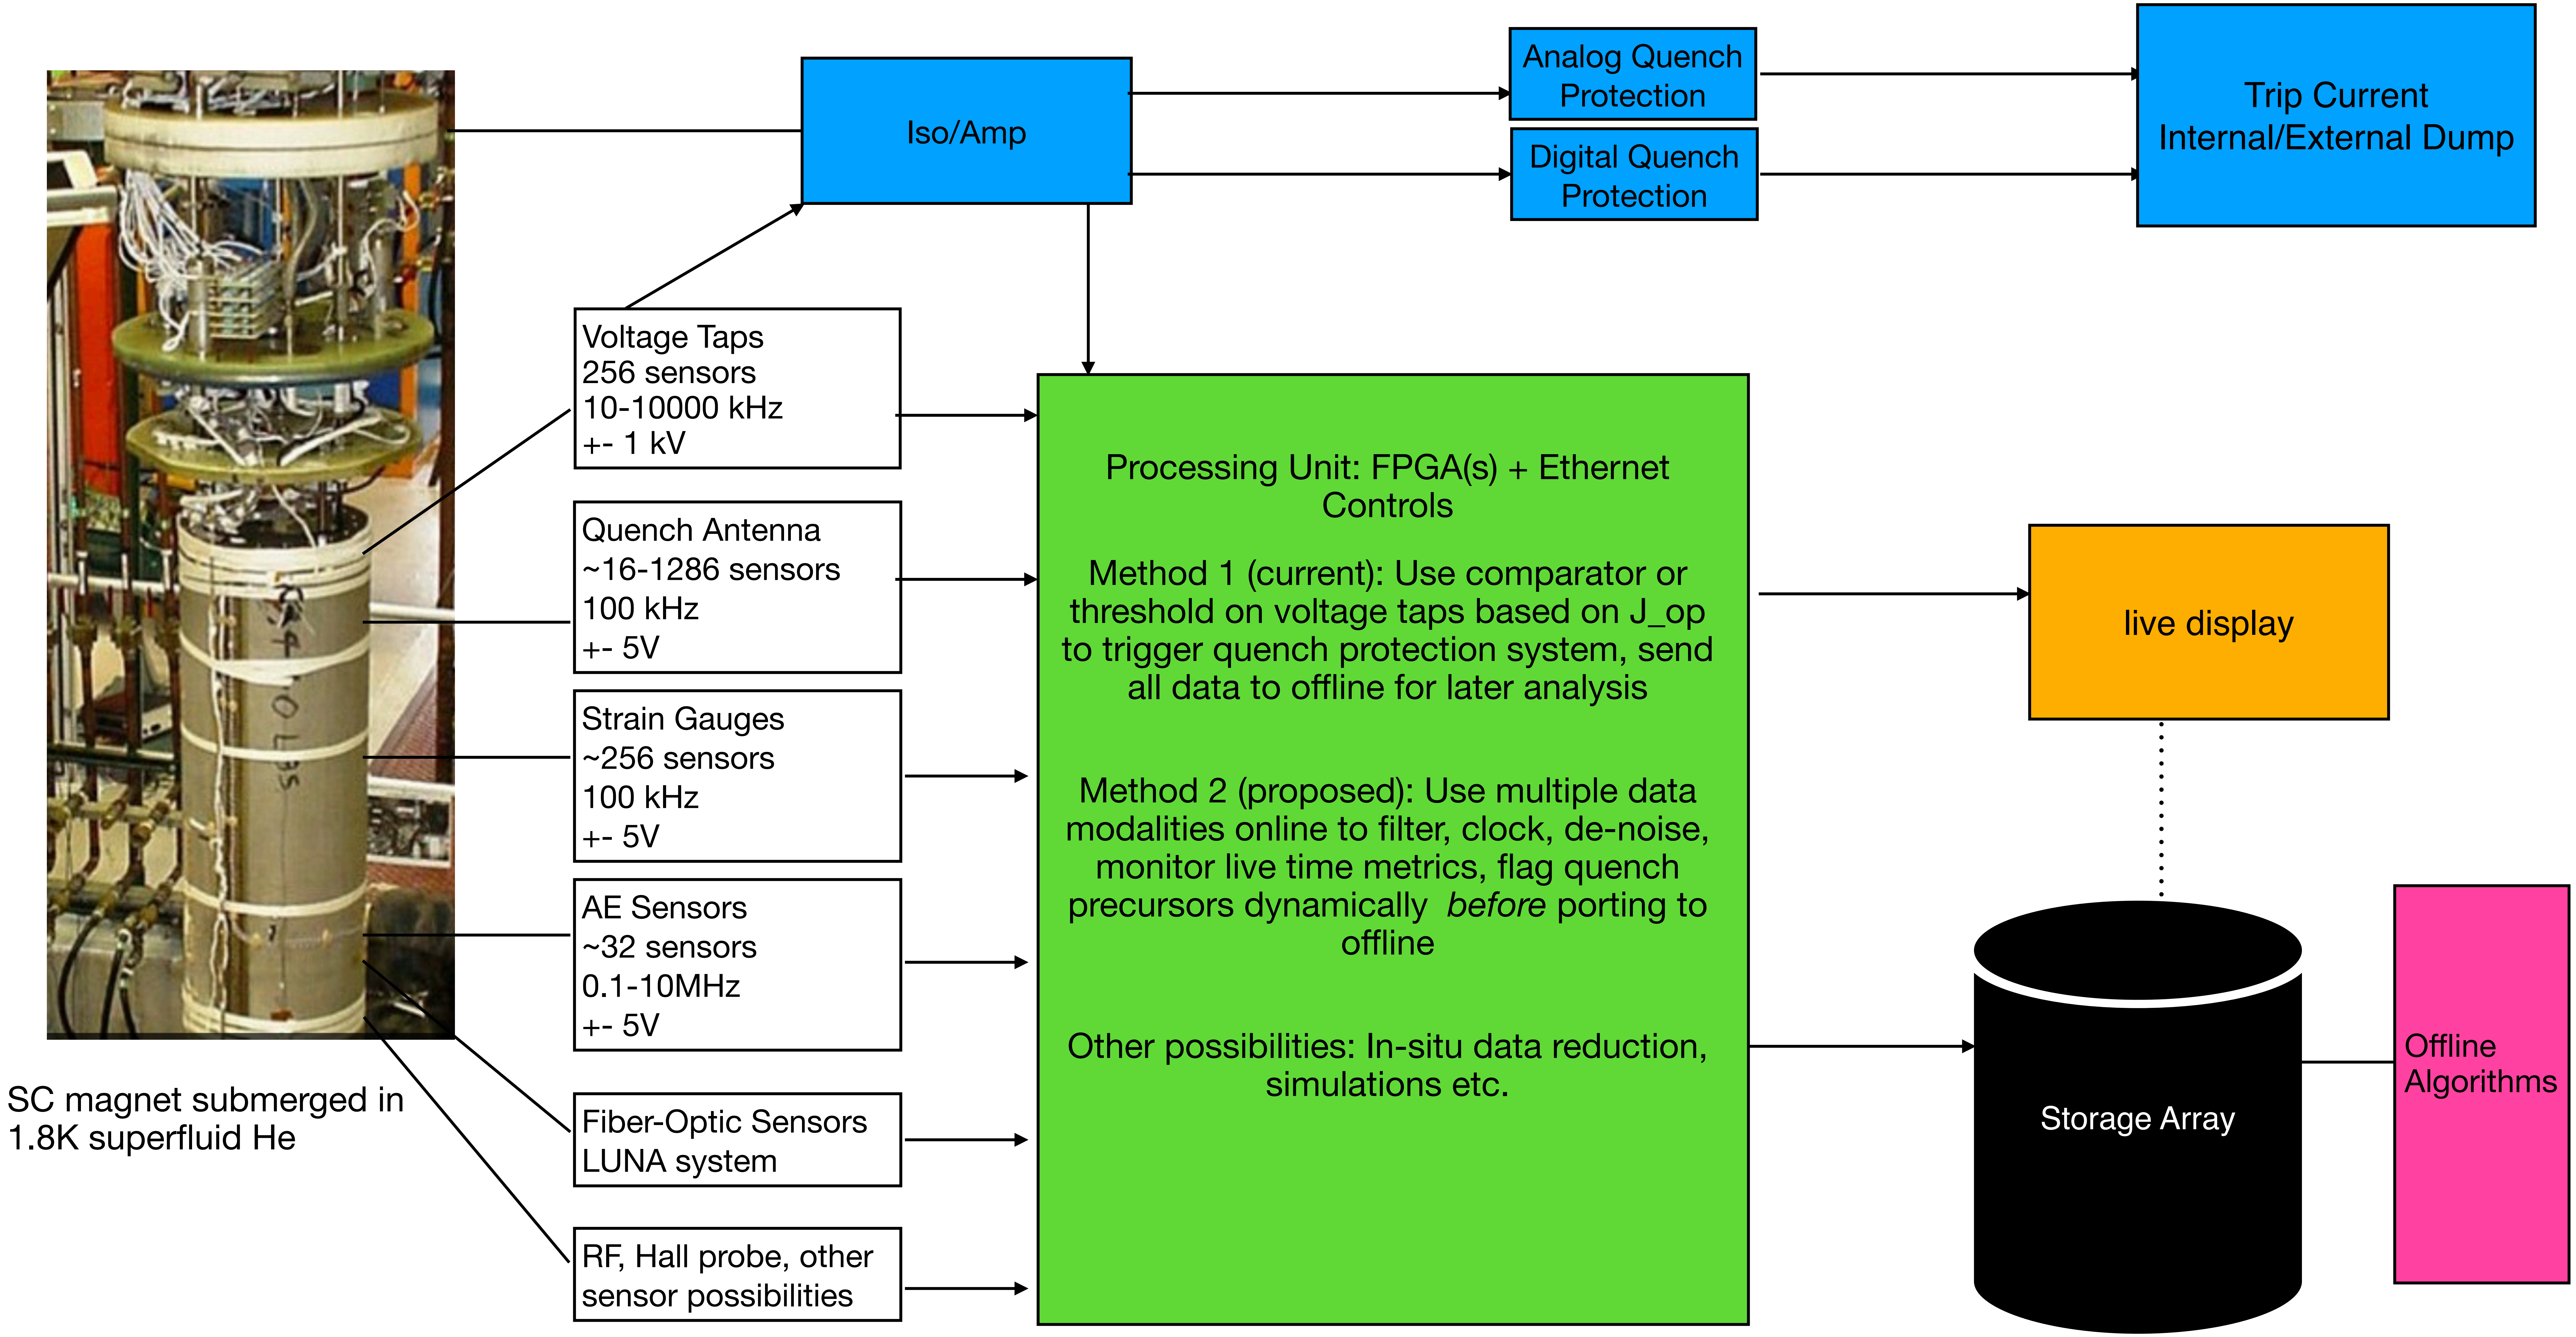

Data Driven Quench Characterization

- **System Objectives:**
  - *Goal:* Quench detection before some *ms* of quench event
  - Modular and scalable with different sensor types, sample rates, number of channels
  - Robust to changes in power supply noise
  - Prevents false triggers during training ramp (high amplitude of signal that does not correlate with quench)
  - Records metrics for interpretability of quench propagation in offline analysis
  - Interfaces well with a database storage system
  - **Software task**: Develop an algorithm that processes signals, and flags quench precursors within time and memory constraints
  - **Physics goal**: Detect quench in a way that captures energy loaded to the magnet, flux jumps, and current re-distribution during the training ramp to determine when a quench will occur

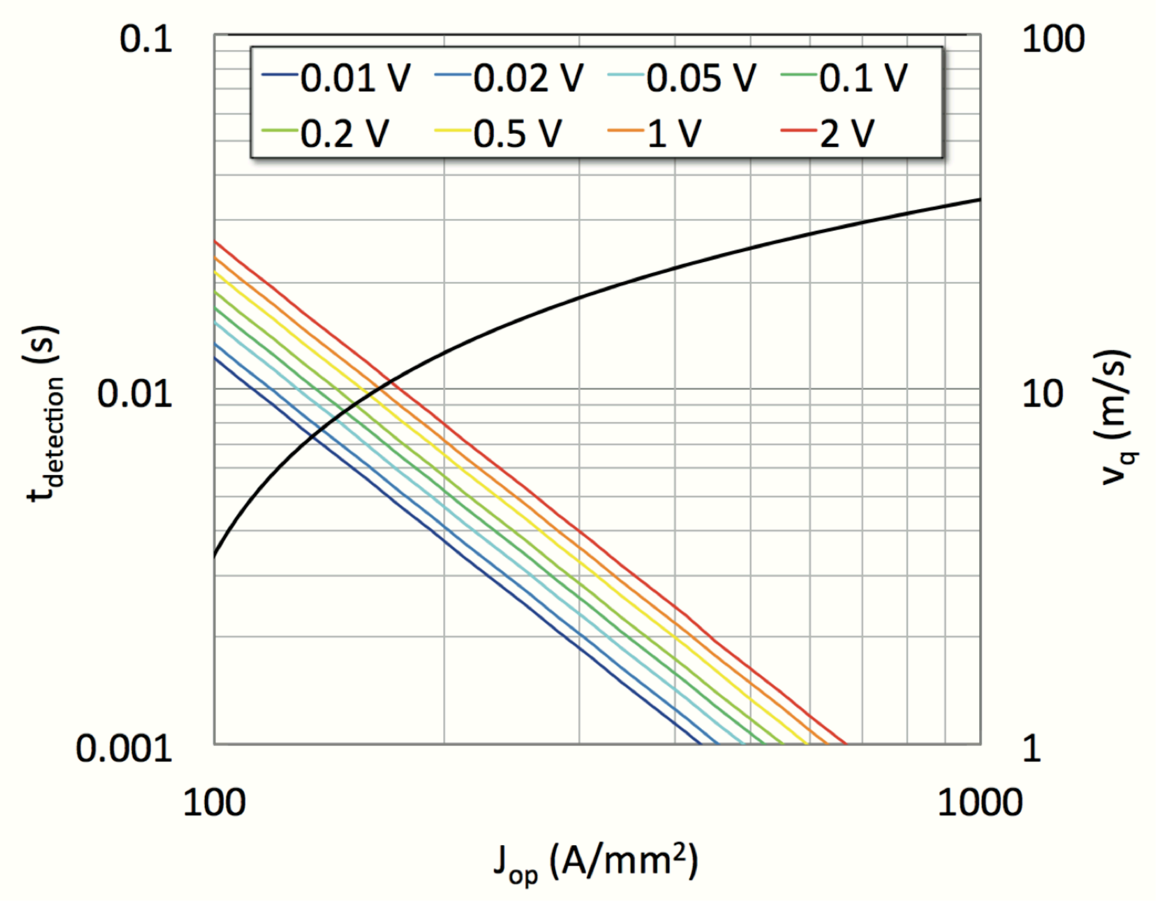

Existing method: Rescale according to current density and choose particular voltage. Source: Magnet Quench 101

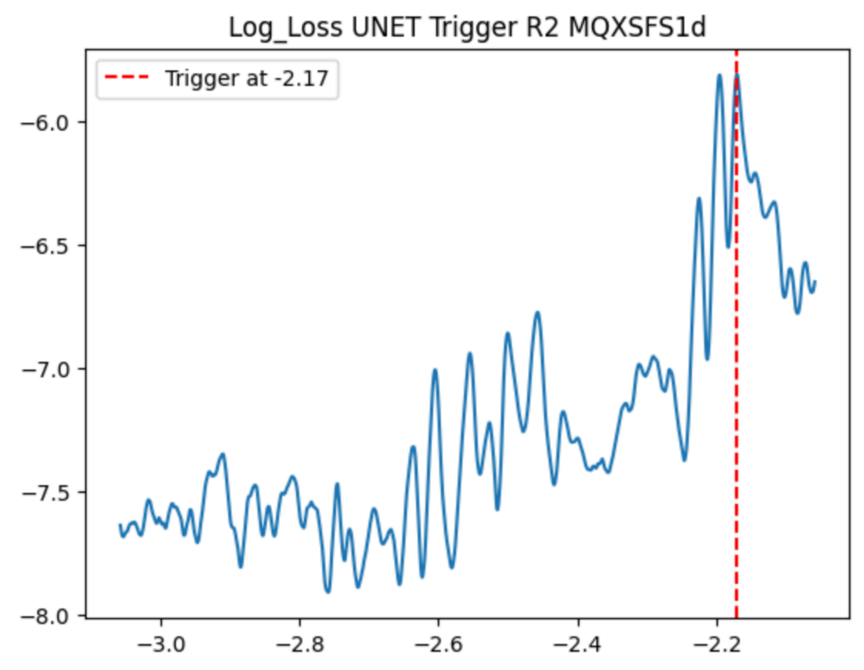

Our sample ML Trigger that uses acoustics + QA data

Previous Work

- FNAL group has studied real time machine learning on acoustics data only where the model was given raw data with statistics with a focus on real time anomaly direction
- Events were developed on two magnets with acoustic sensors
- Focus was on exploring what was possible with viewing data with ML
- Our current work focuses on developing *denoising* algorithms, running on *multiple* data types, and considering how these models may be interpreted with respect to raw data with MBHSM03 and MQXFA series magnets

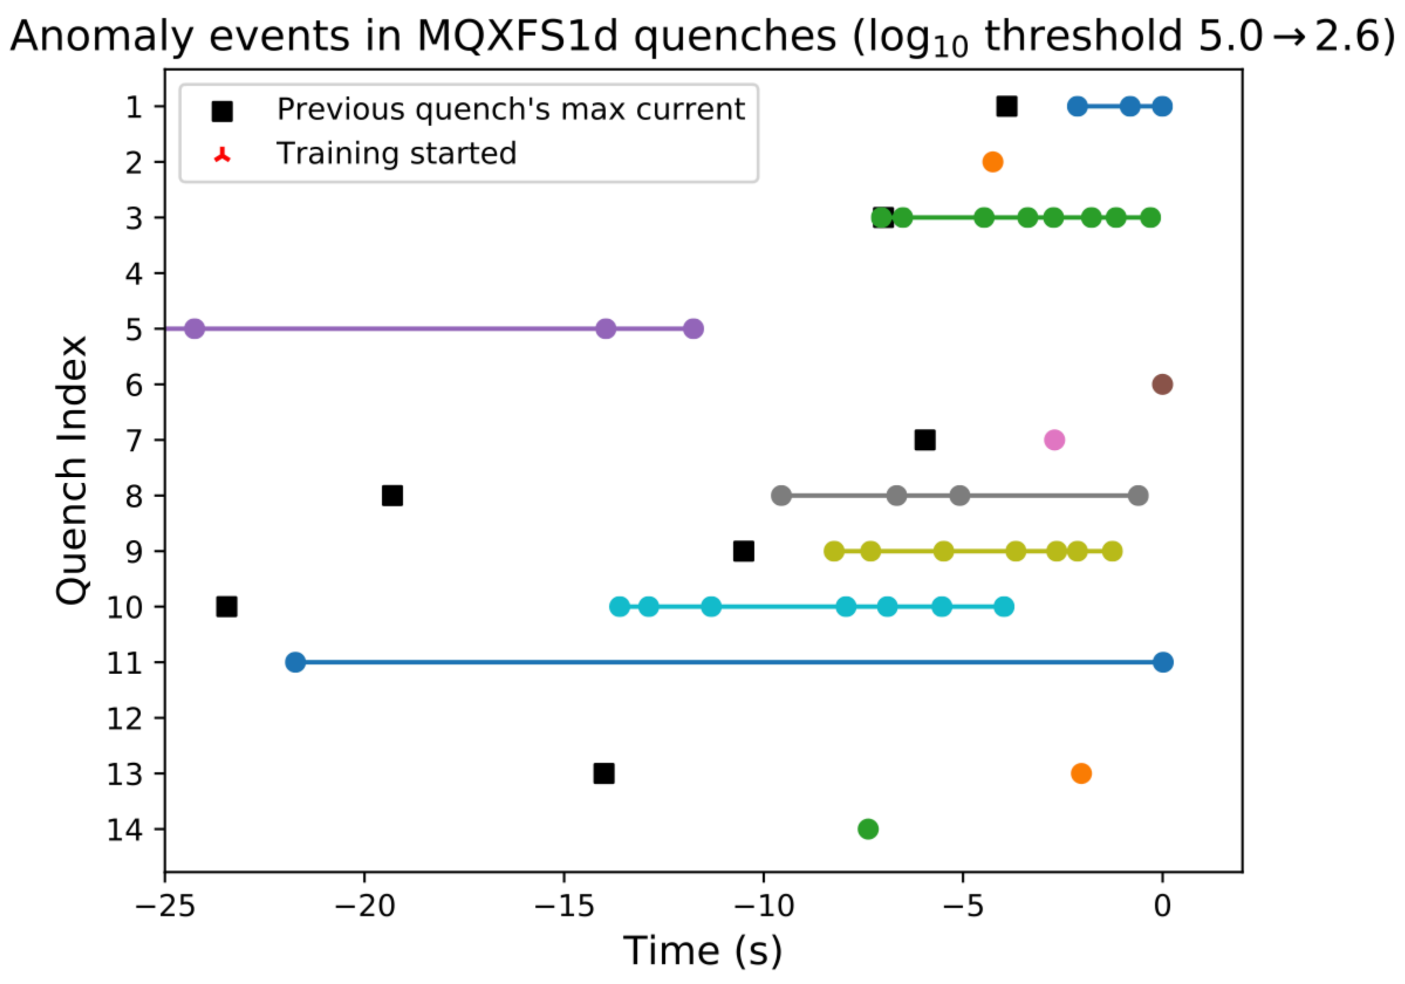

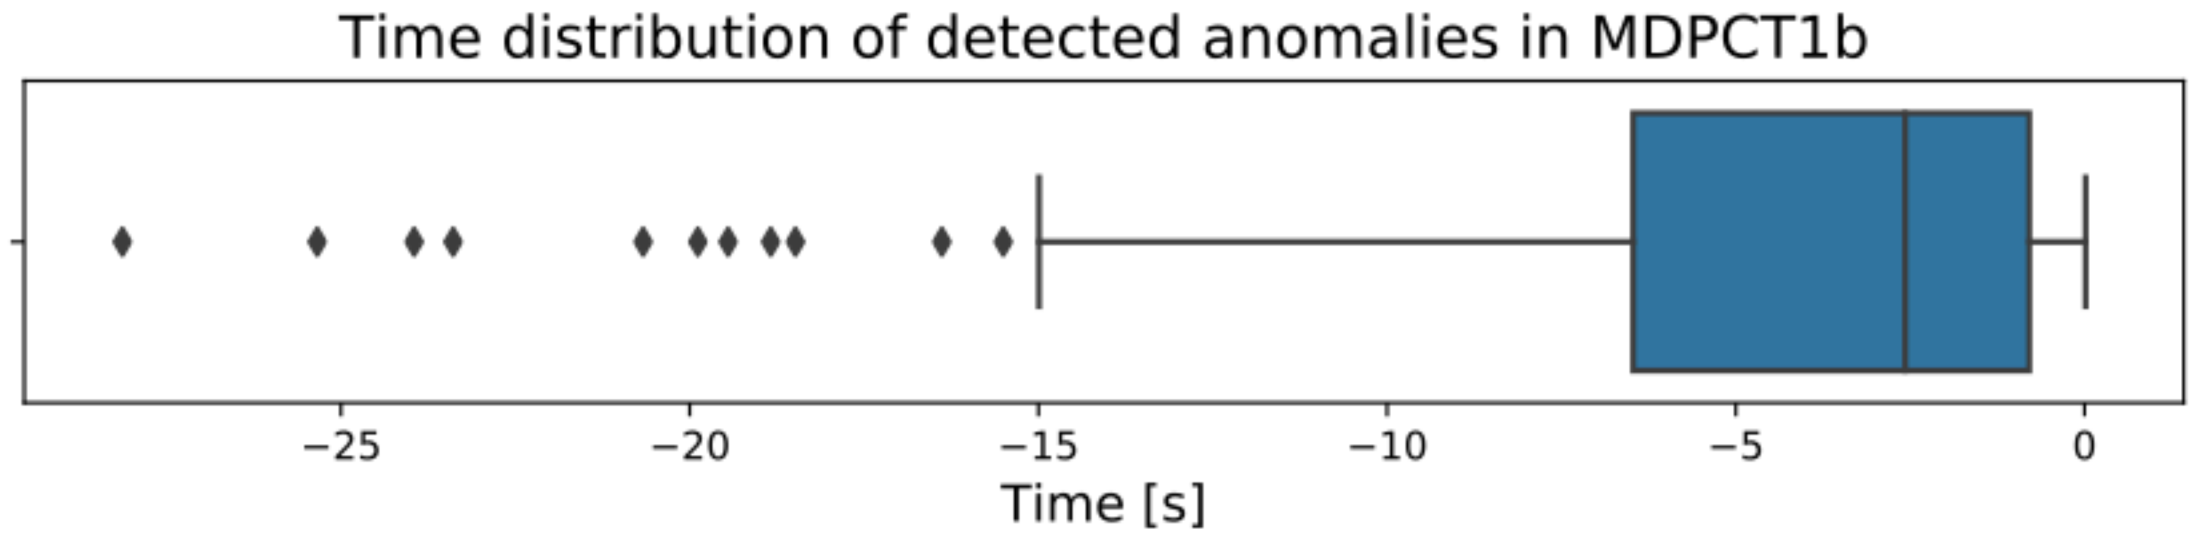

source

Standard Algorithms vs. ML Approaches

	Standard Algorithm	ML Inference Techniques	Evaluation Metric
Signal Denosing	frequency domain technique (bandpass, FFT, etc.)	denoising auto-encoder, blocked weight matrix	SNR of 1D Signals Number of signals above noise floor at quench
Quench Boundary Detection	Weighted threshold by J_op	reconstruction loss on dynamic threshold	rate of false and true triggers ~s of quench, exceeding previous current
Quench Precursor Identification and Recall	limited, possible state machine history	reconstruction loss history n-k ramping history	*dB/dt, dE/dt* Triggers above previous quench current
Cross Sensor Correlation (Sensor and Data Type)	correlation coefficients between single RMS at a given *t*	latent space projections of channels	correlation density of channels and sensors
Spatial Localization	signal amplitude normalized to noise RMS at (*x_i, y,_i, t_j)*	reconstruction loss at (*x_i, y,_i, t_j)*	$\frac{\|s\|_{x,y}}{\|L\|_{x,y}}$
Event Characterization	k-cluster of signal features	latent space separation	Transient events, flux jumps, current redistribution, Kaiser effect (in development)

*The ***latency*** and ***parameter*** footprint (memory resources) are also studied

Filtering and SNR at Quench Time

- **Raw Inputs:**
  - MBHS03 Magnet
  - QA voltages at 100kHz
  - Acoustic voltages 1MHz
  - Current at 1 MHz
  - Rolling window of 20ms, 20us step
- **Filtering and Preprocessing**
  - Removed multiples of 60Hz for QA and Acoustic signals
  - Computed spectral entropy of raw windowed signals
  - Why? Spectral entropy is robust to variations in noise, potentially helping with issues with power supply
  - Work in progress: Build single inference matrix to replace SE/FFT filter

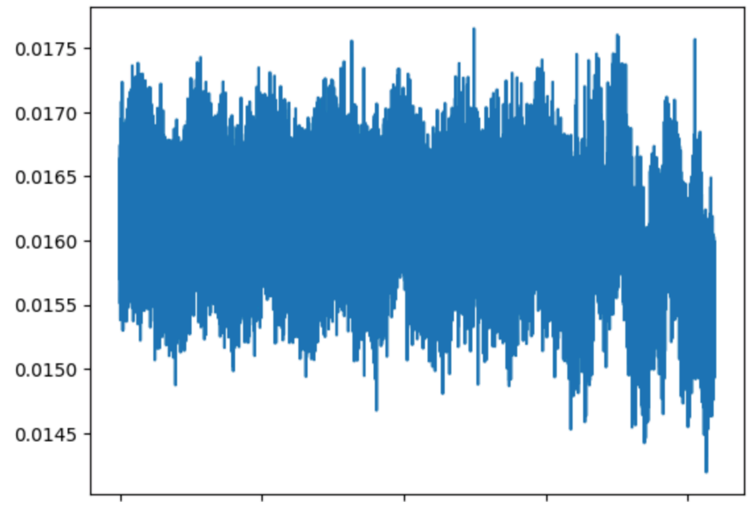

Raw waveform

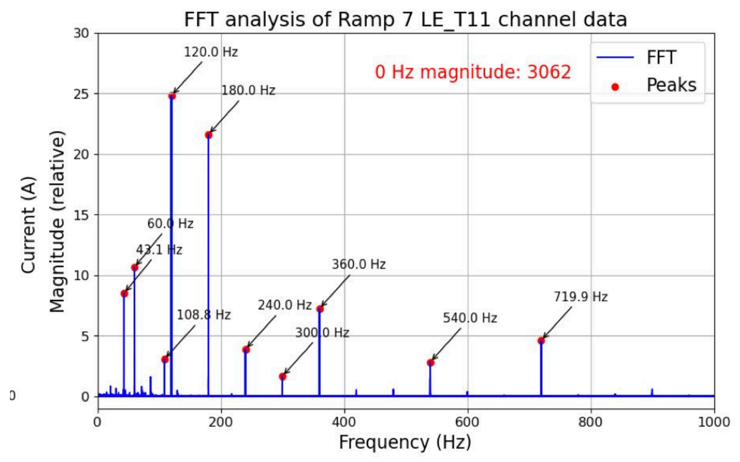

STFFT to remove 60Hz Harmonics

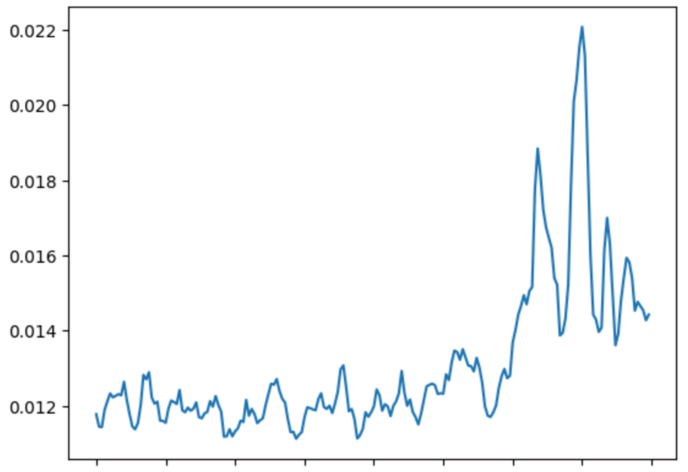

Spectral entropy over window

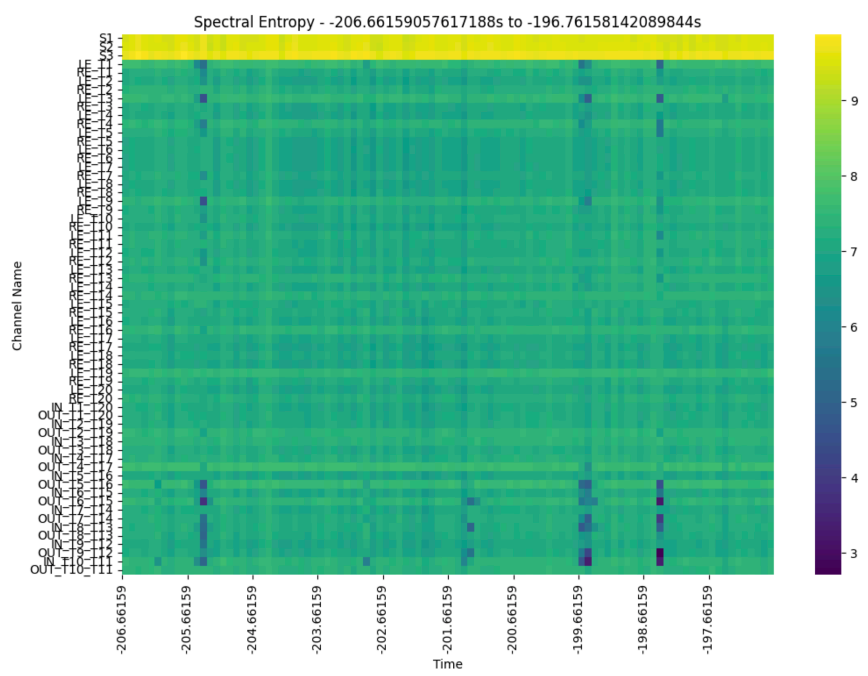

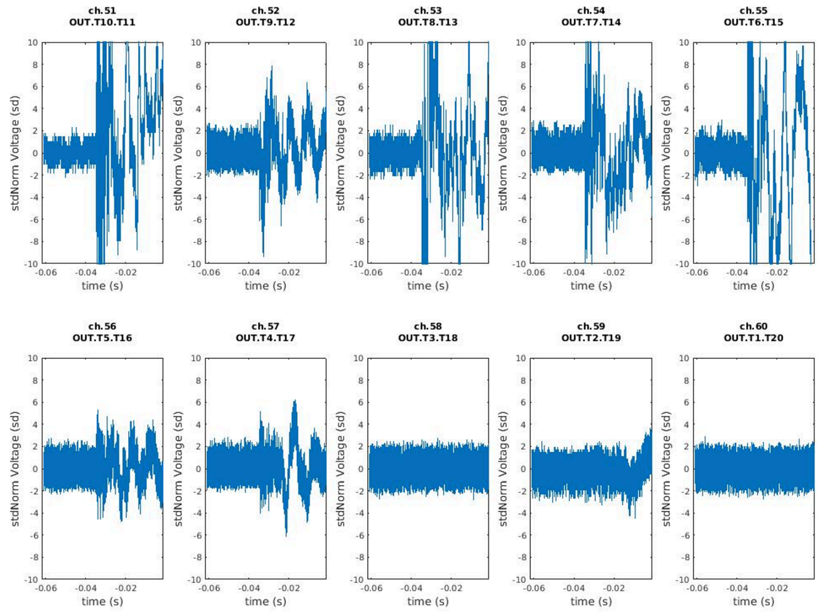

Noise floor remains consistent in time

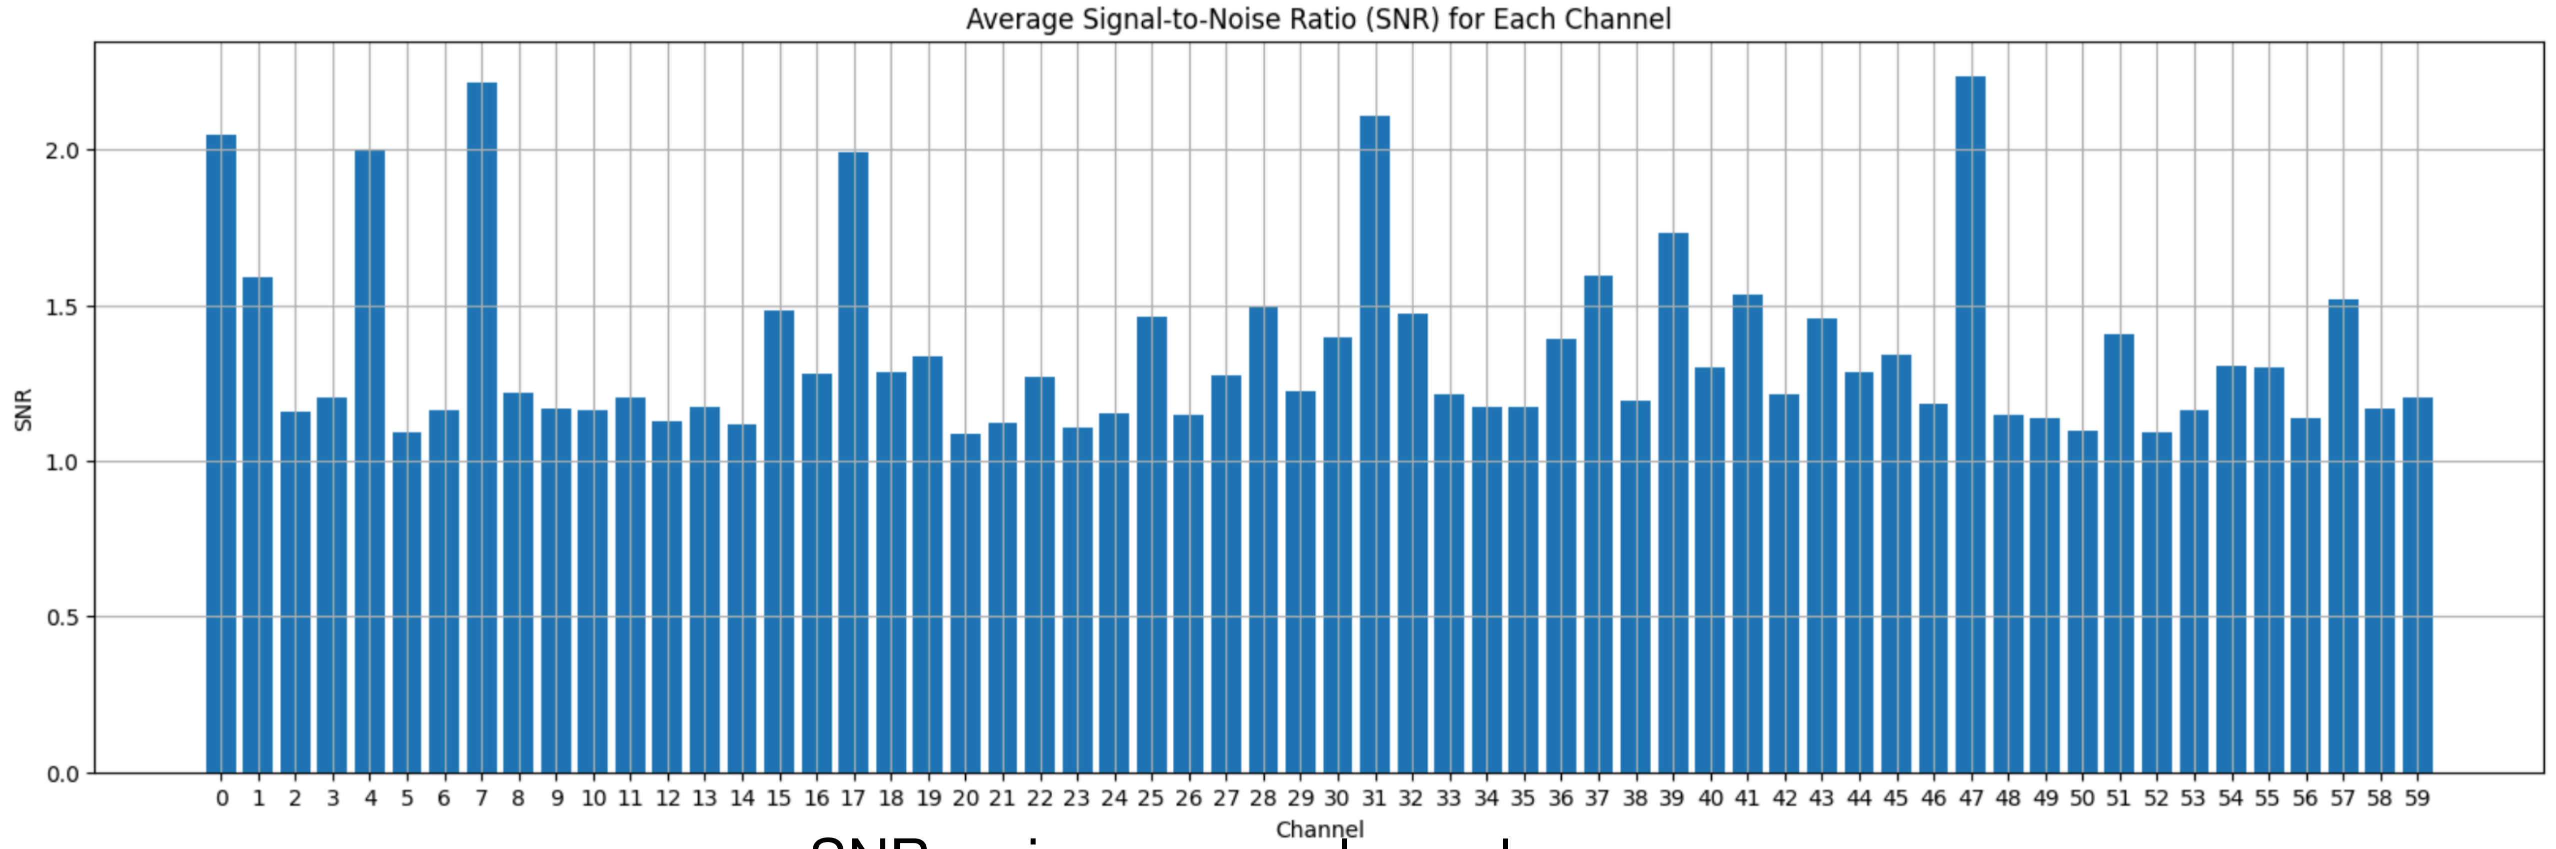

SNR varies across channels, some more sensitive than others

Algorithm Logic

Input: QA channels

Raw waveform of 20ms with 20us step

Filter 60Hz

Spectral entropy filter

Output: float value of spectral entropy in the window

Input: AE channels

Raw waveform of 20ms with 20us step

Filter 60Hz

Spectral entropy filter

Output: float value of spectral entropy in the window

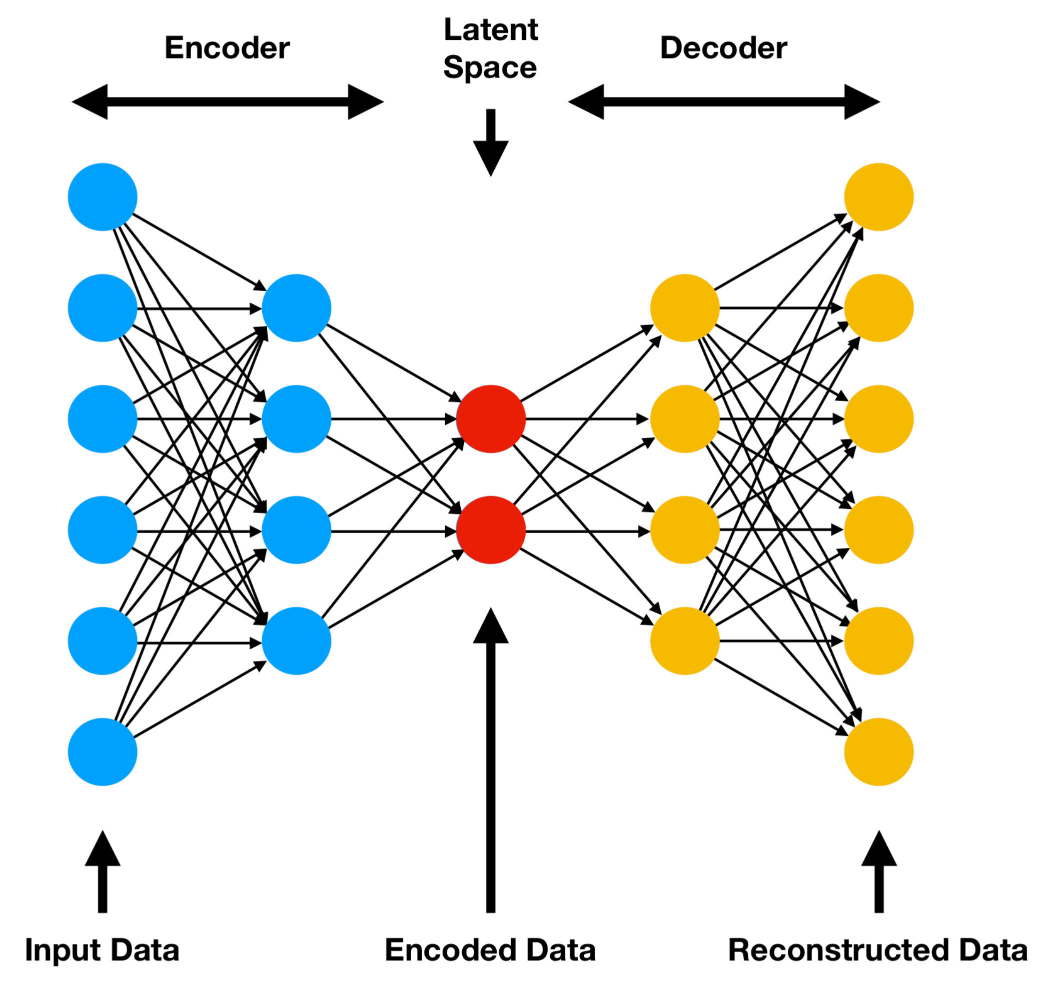

Is L1 difference between actual and reconstructed data > n sigma of RMS floor of previous

~10s of acoustic data?

~3s on QA?

If yes, trigger

If not trigger, set n to n-1

Neural network updates gradient values every 7s

We are taking into account the following

1. Acoustic events tend to have much higher relative amplitude

2. QA events tend to occur more densely, and transients may occur in a short time span

3. Investigations into adaptive weighting and normalization of two model trigger streams are underway

Quench “Anomaly” Trigger

- To detect the quench boundary, we weight a dynamic threshold based on the RMS of the reconstruction loss
- We evaluated for 29 ramps, ramps n >1. The QA trigger missed 1 ramp. The acoustic trigger missed 6 ramps

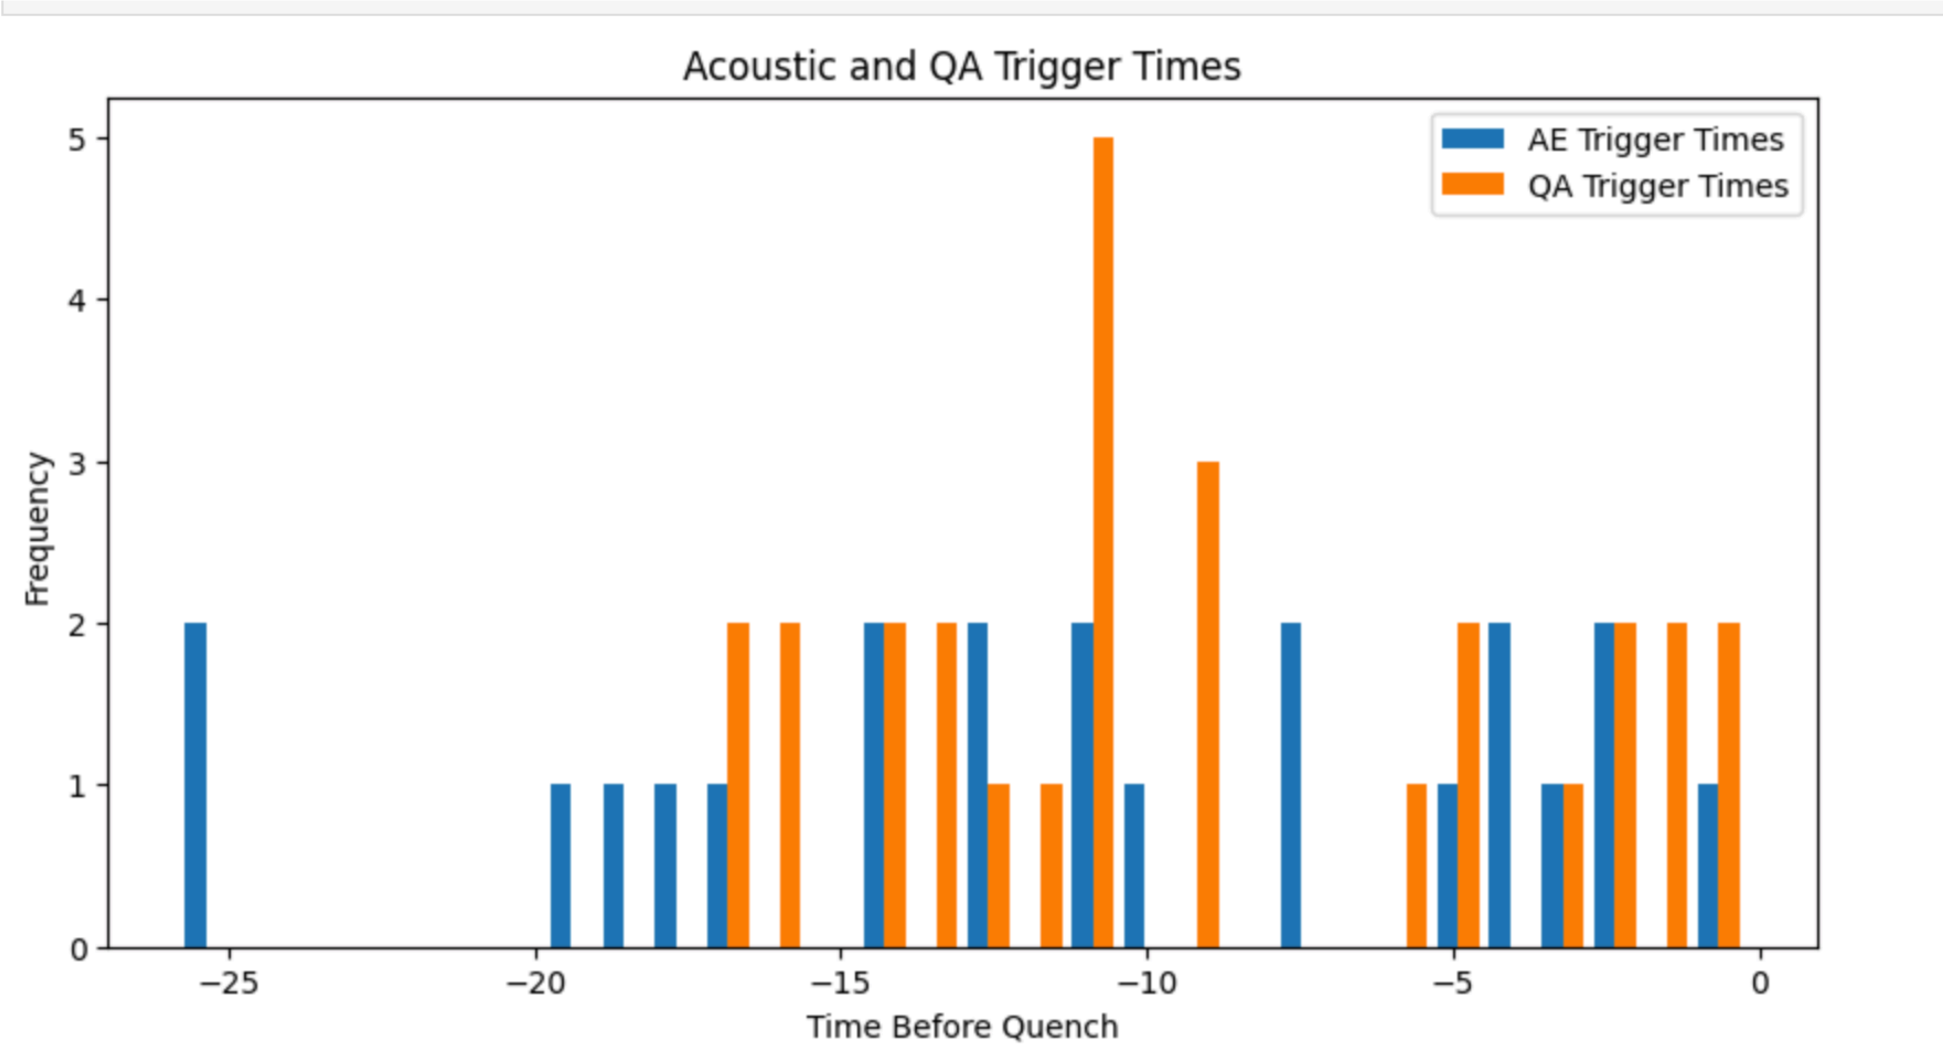

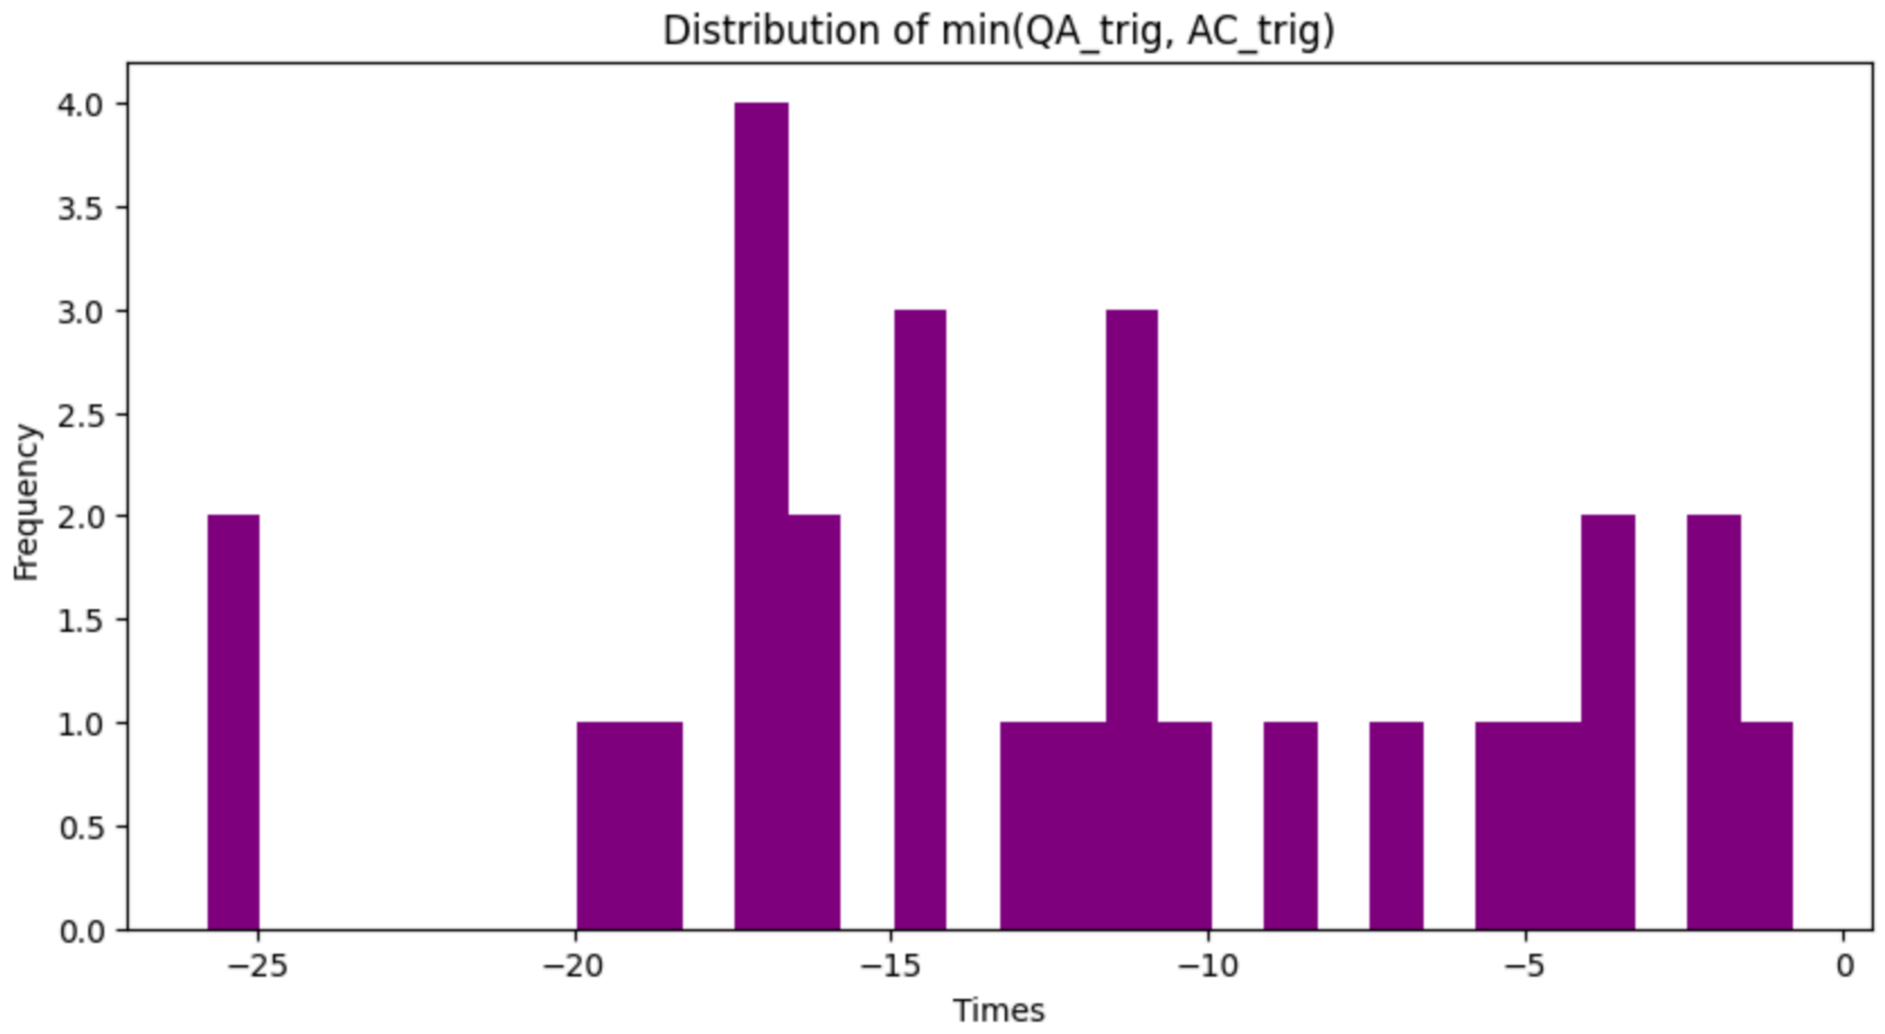

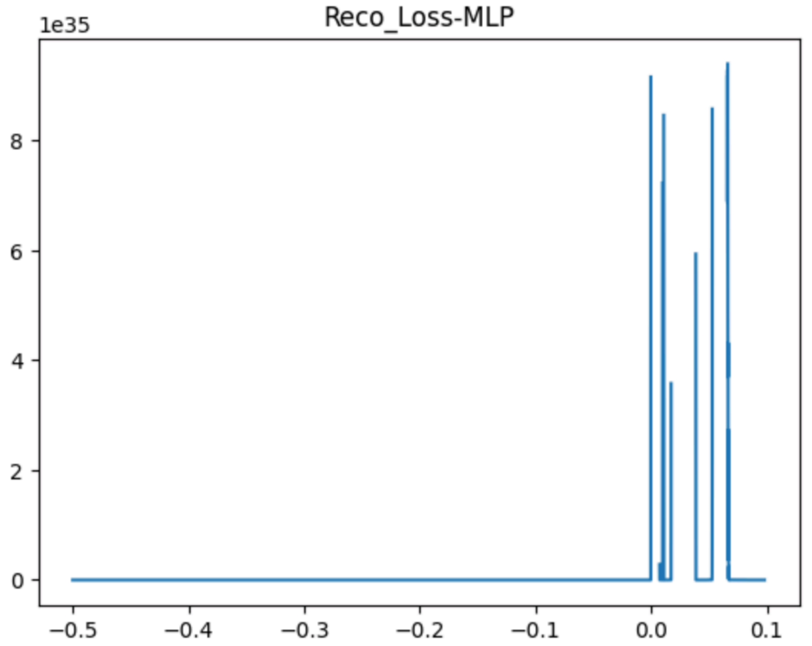

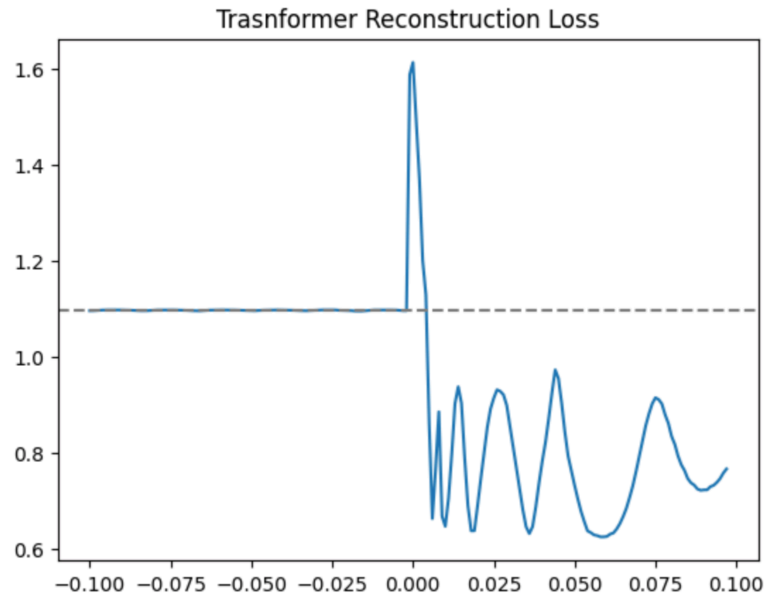

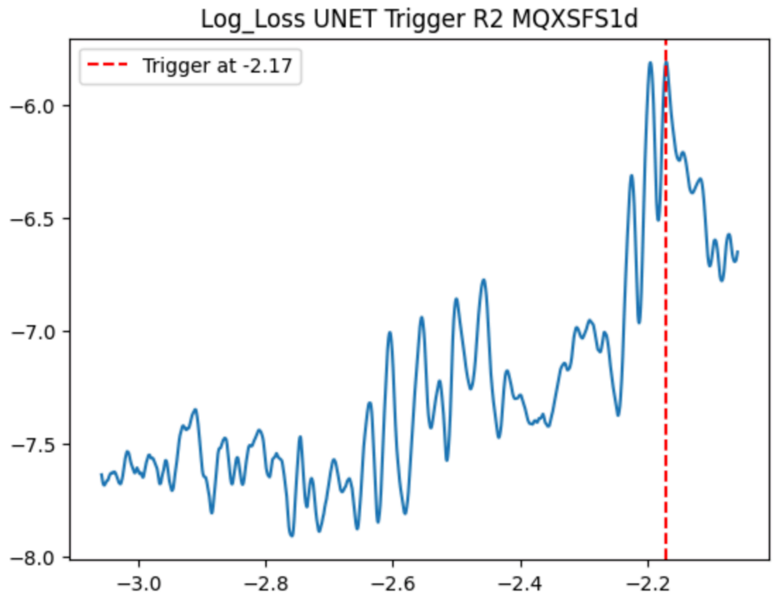

The model seems to trigger a “precursor or anomalous event earlier with acoustics than quench antenna

Possible causes: Power supply changes, event density by current boosting

Sample loss triggers by different machine learning architectures

Quench Precursor Identification

- We can compare the changes in reconstruction loss of the auto-encoder with changes in energy and field of the magnet. At higher currents, the loss is higher, and detects hotspots close to the quench

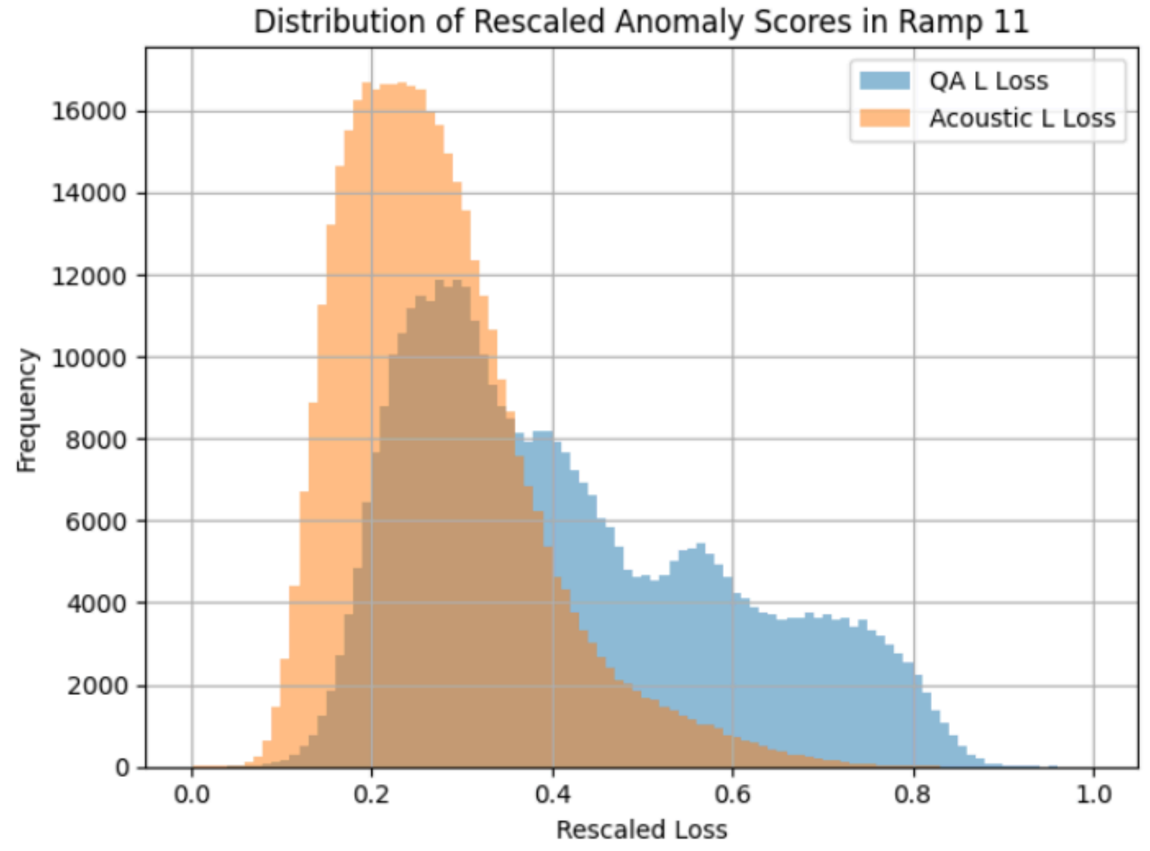

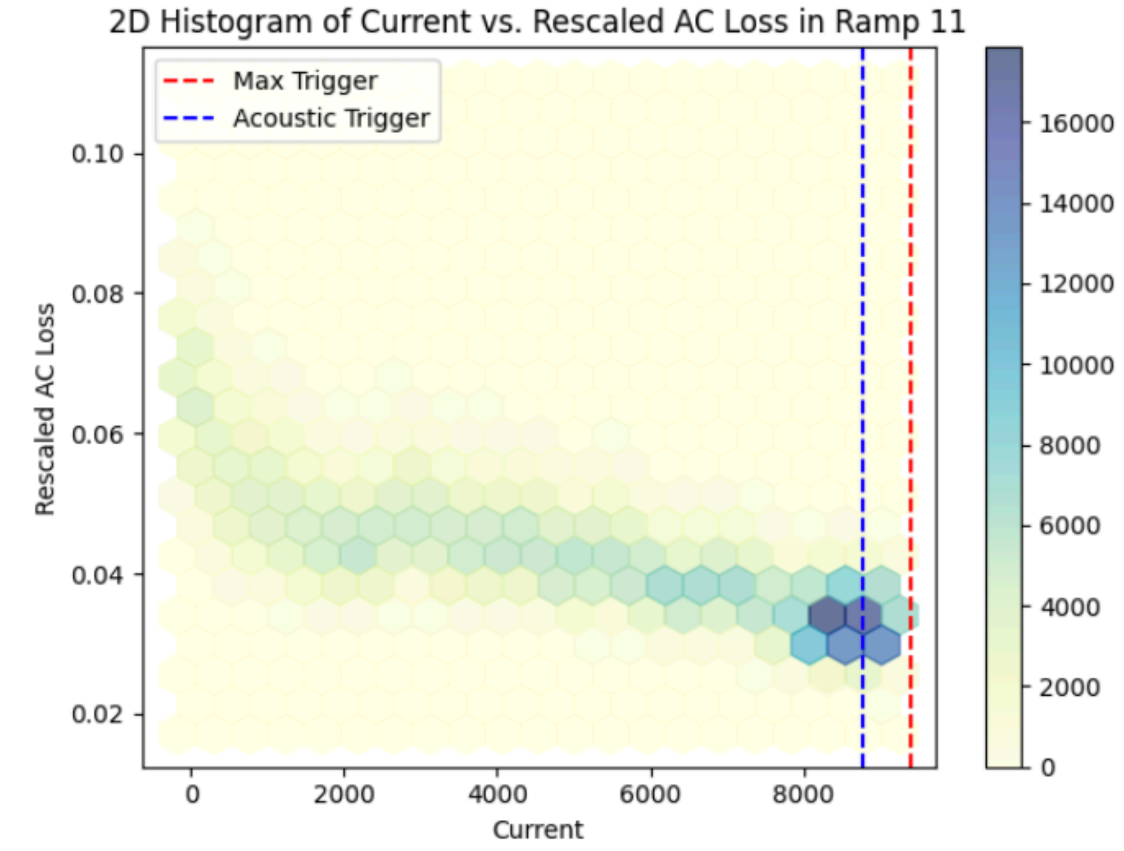

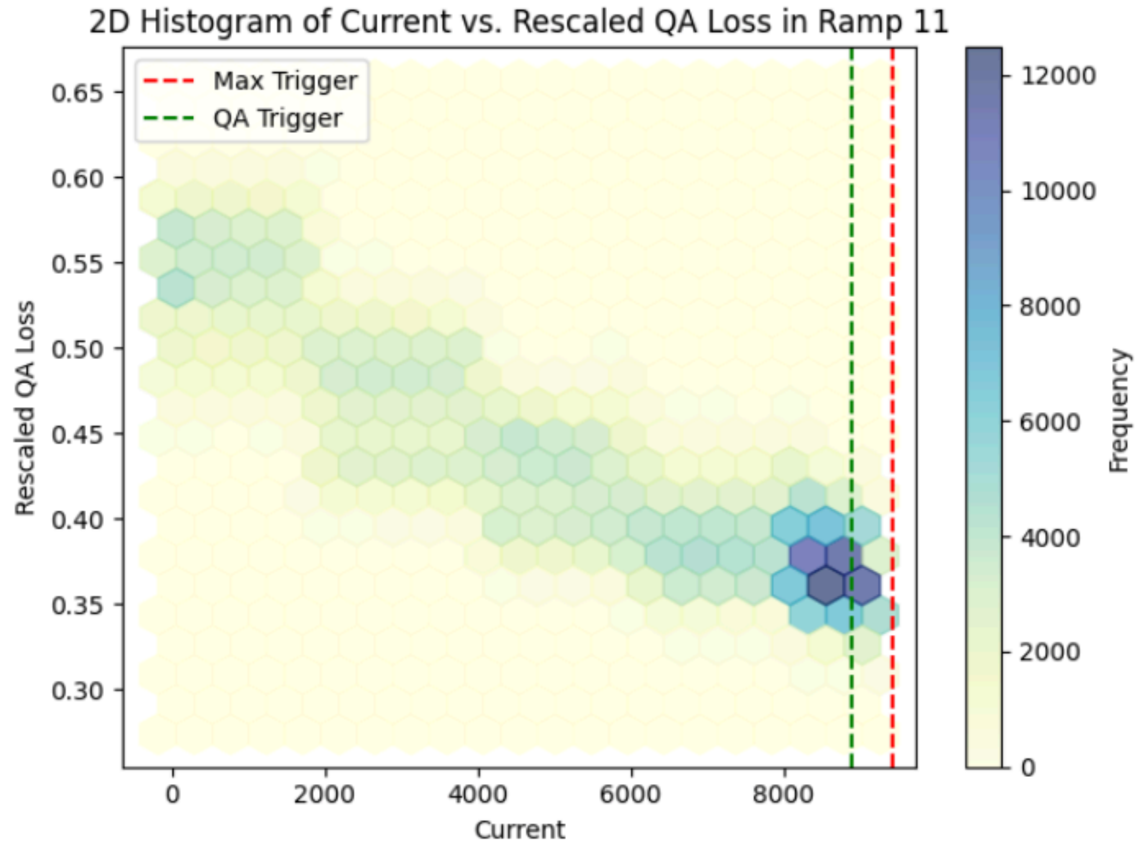

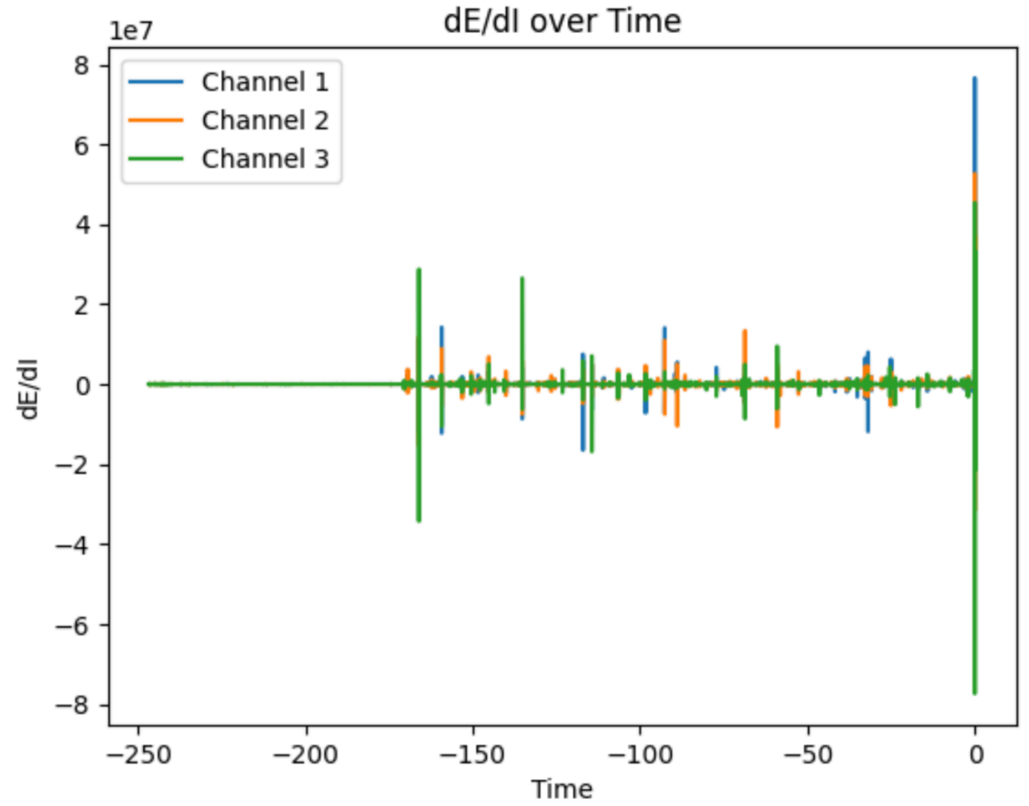

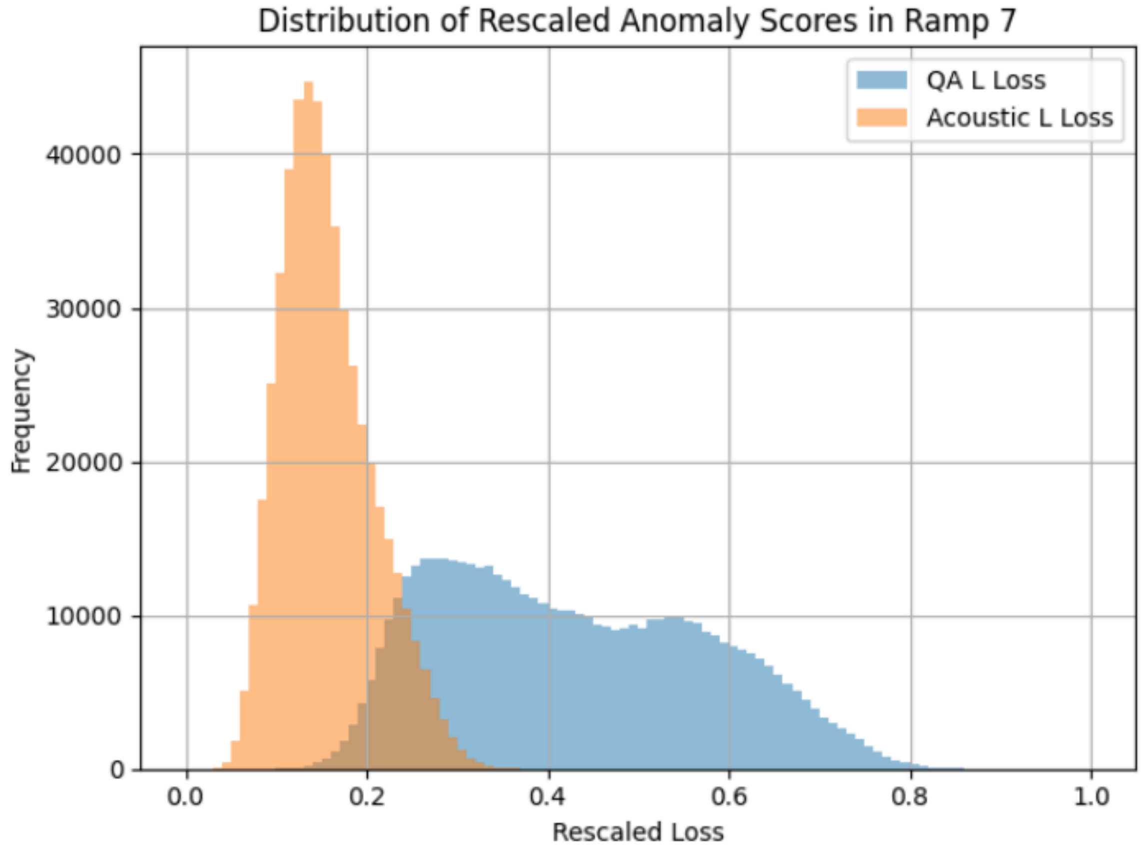

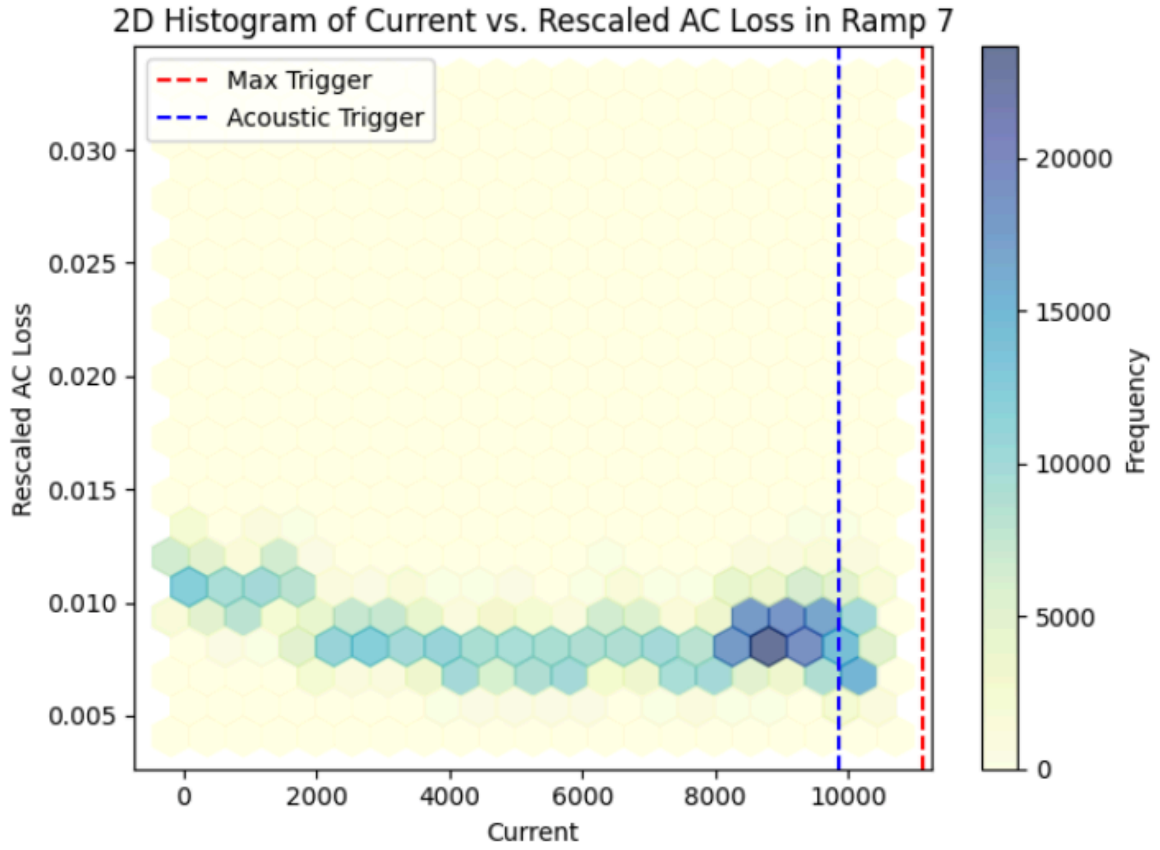

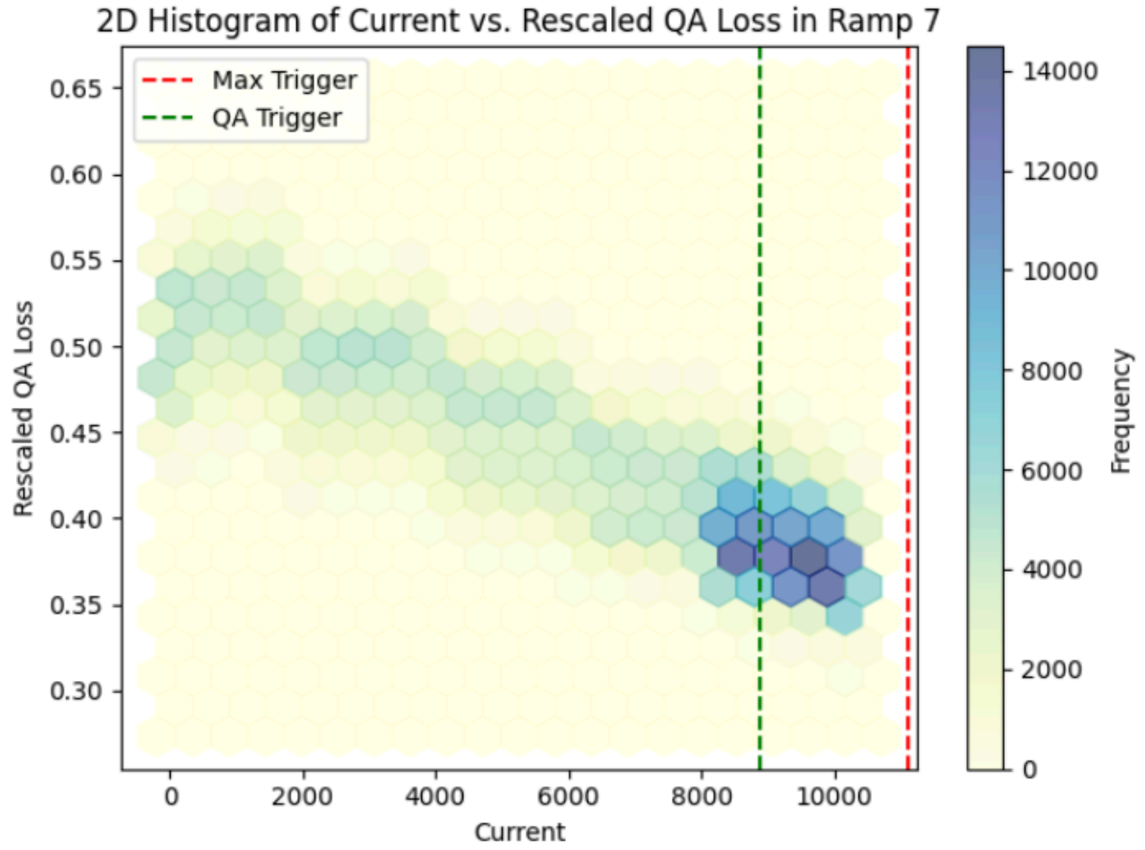

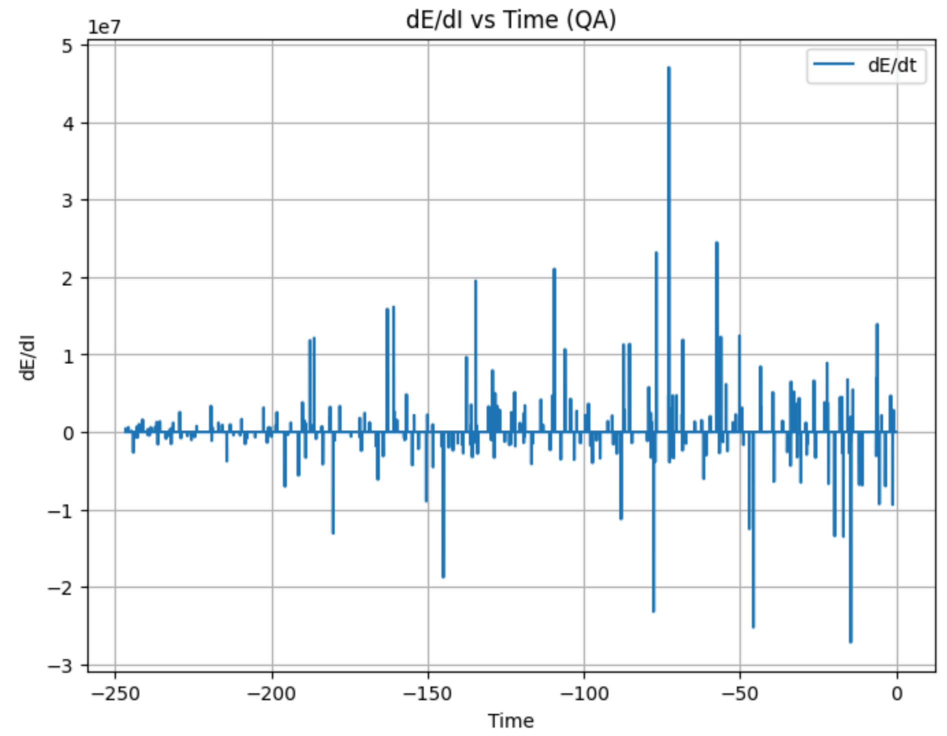

Precursors: Quieter Ramps

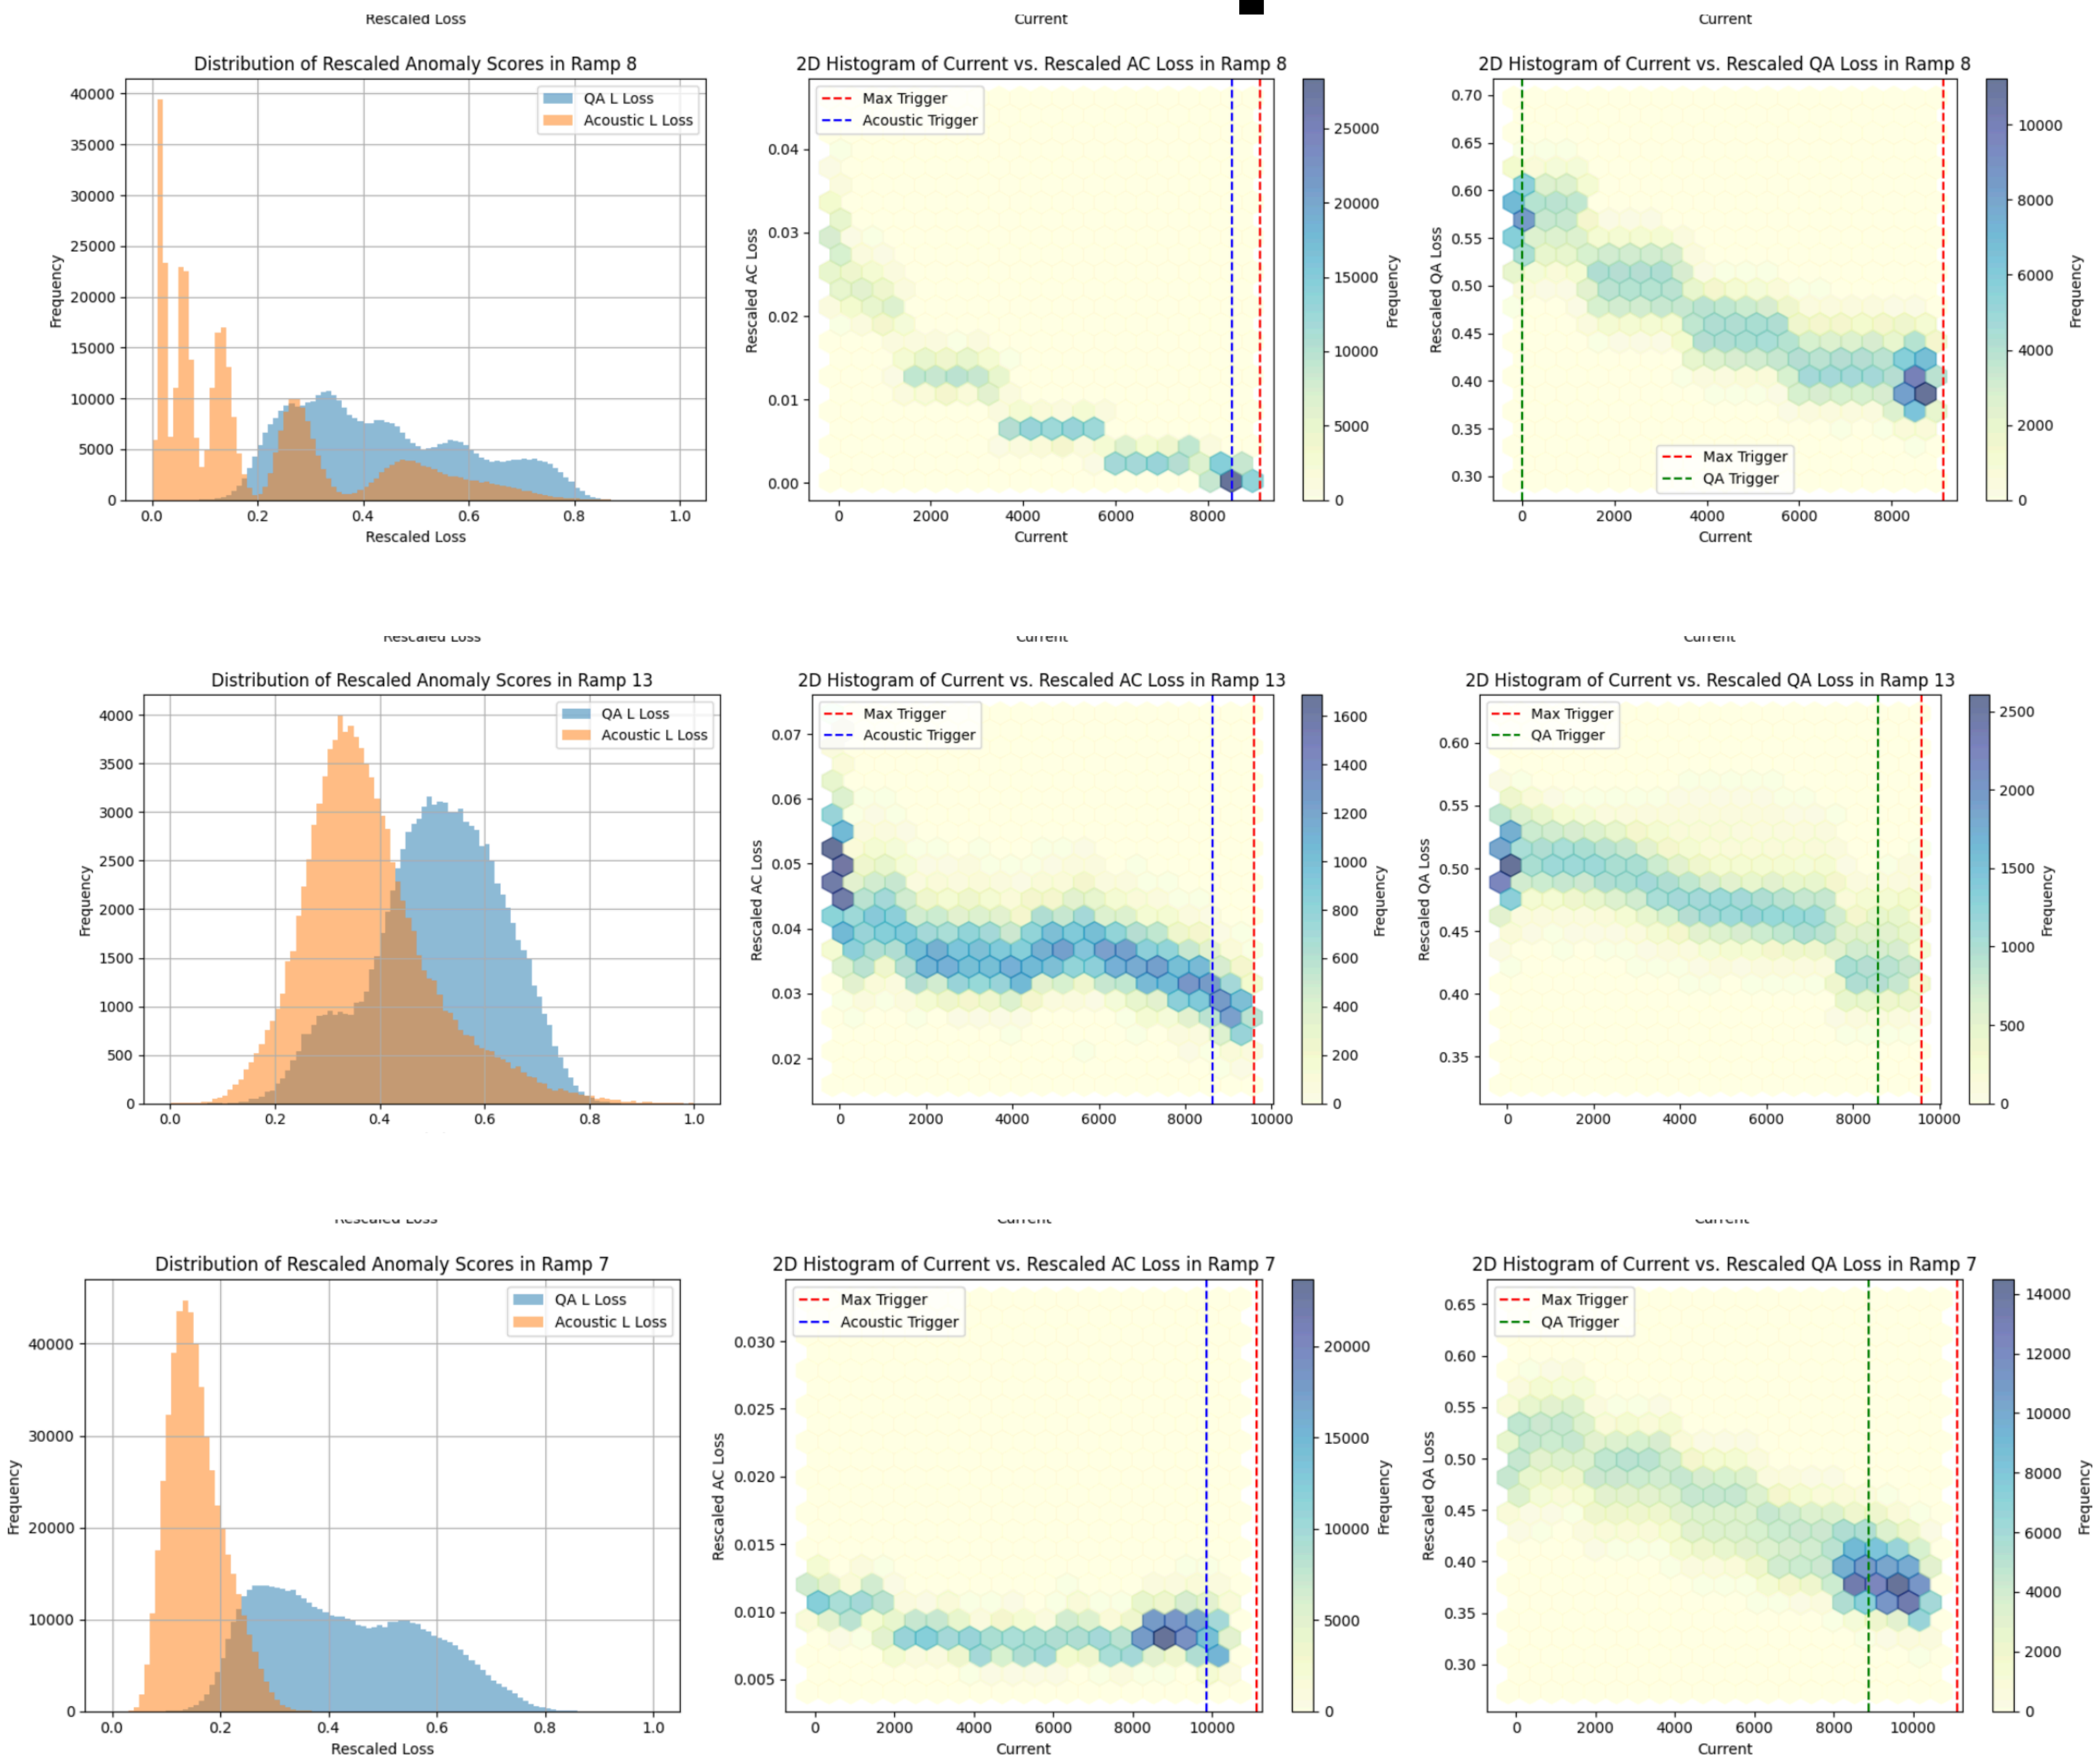

Acoustic triggers seem to line up right before the quench on quieter ramps, or due to small perturbations
Edge effects due to window size over RMS may create a trigger as a false artifact of an anomaly

General Behavior

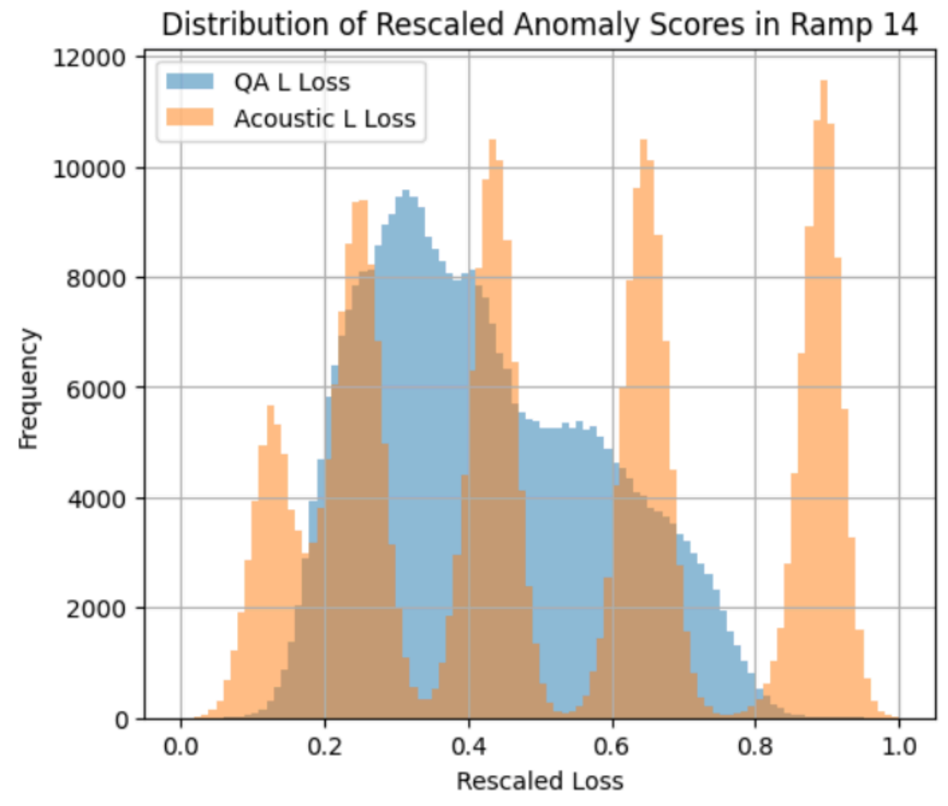

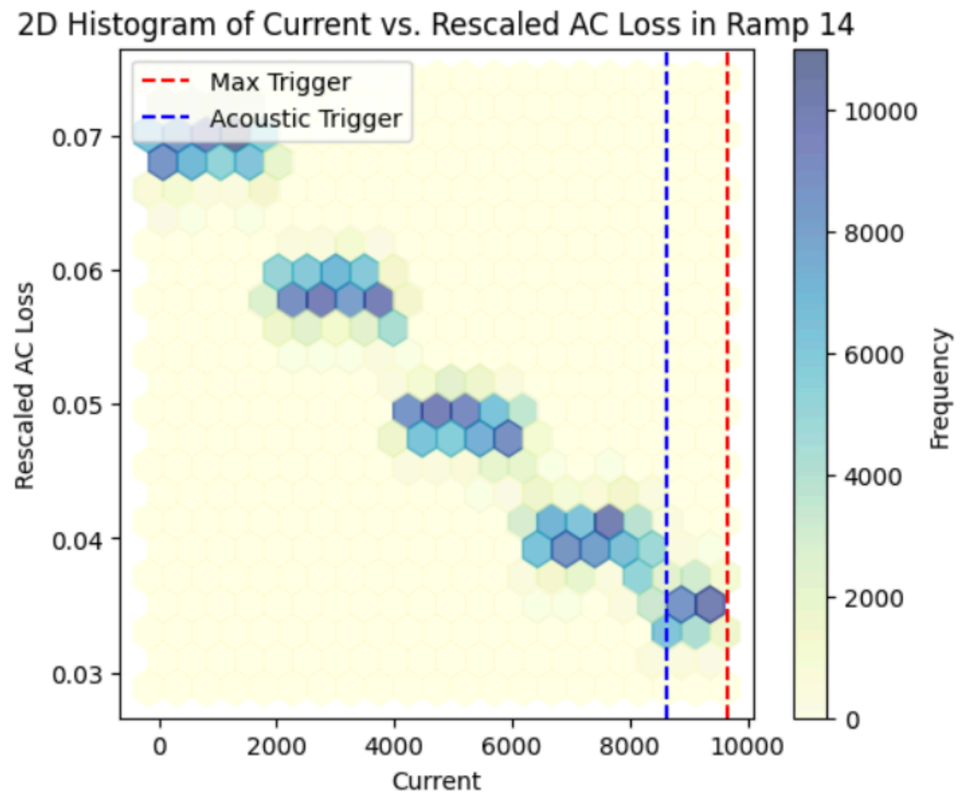

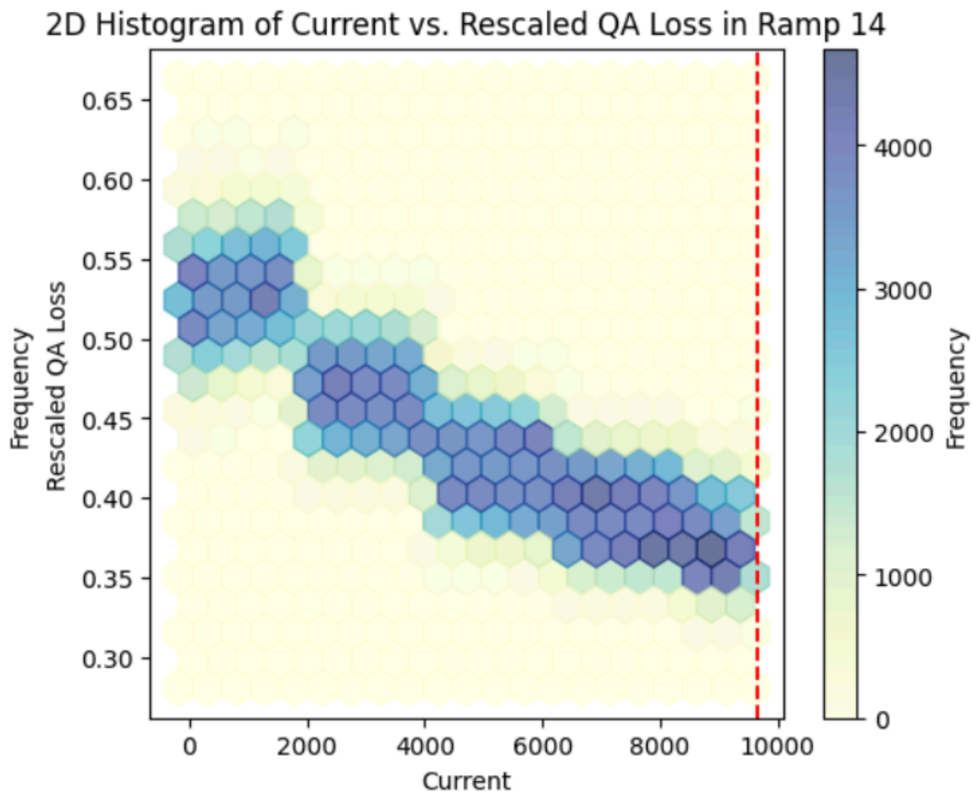

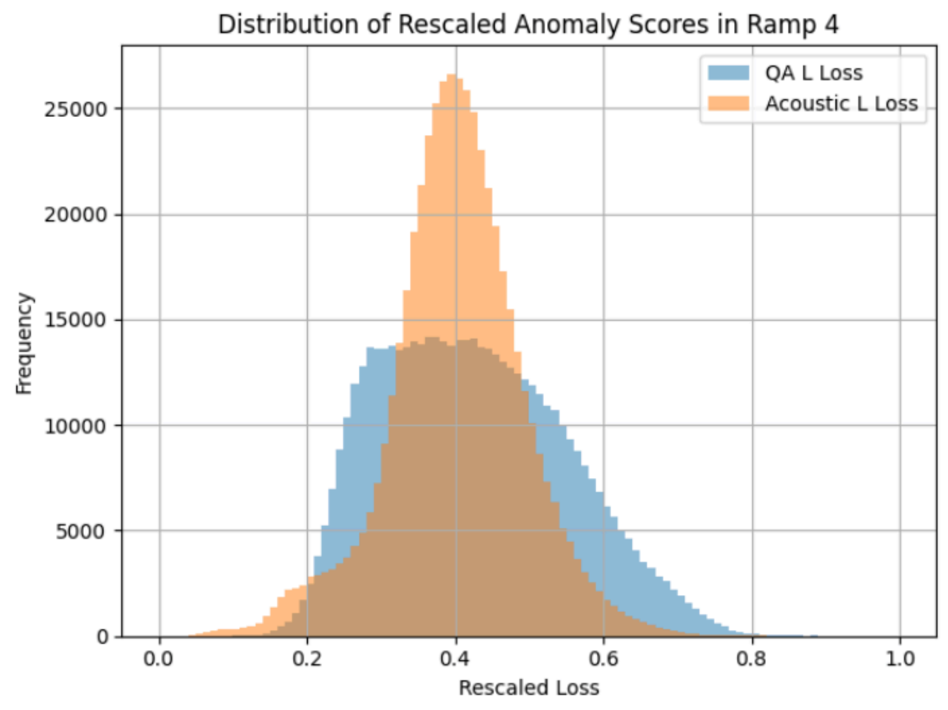

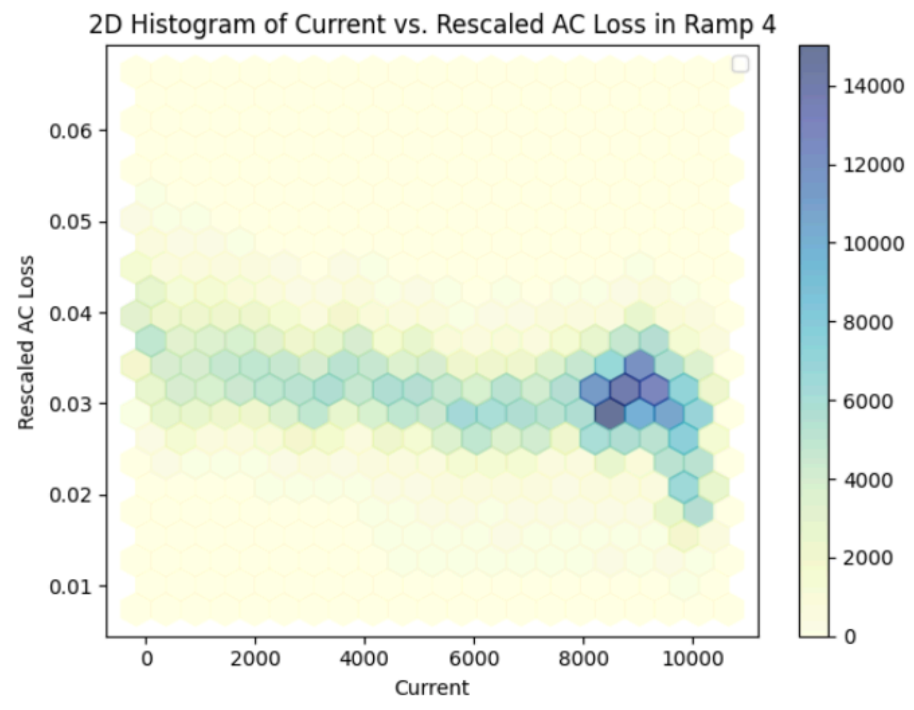

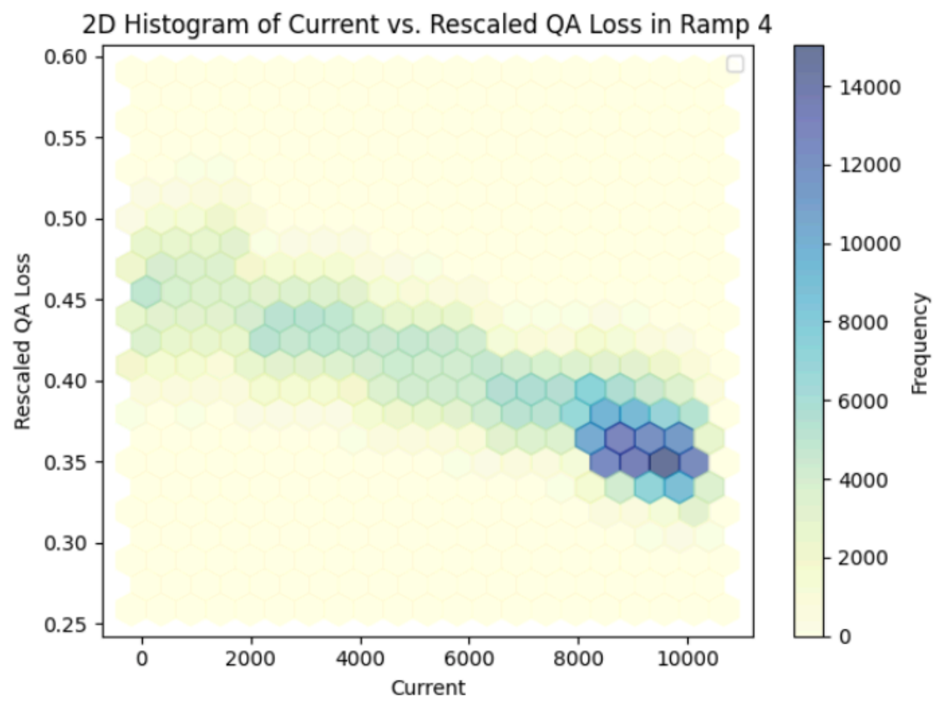

- Across all ramps spread of loss seems to be higher for quench antenna or acoustic. This is probably due to the greater numbers of channels and noise in QA.
- Acoustic events seem to line up with QA events, but smaller more subtle acoustic events many not be captured in our anomaly model due to window sizing, RMS fits of loss, and boundary effects
- We do however see consistent correlations in acoustic and quench antenna using the ML model
- Must now compare with standard statistics
- Event density varies by current ramp rate and previous training

In the Works

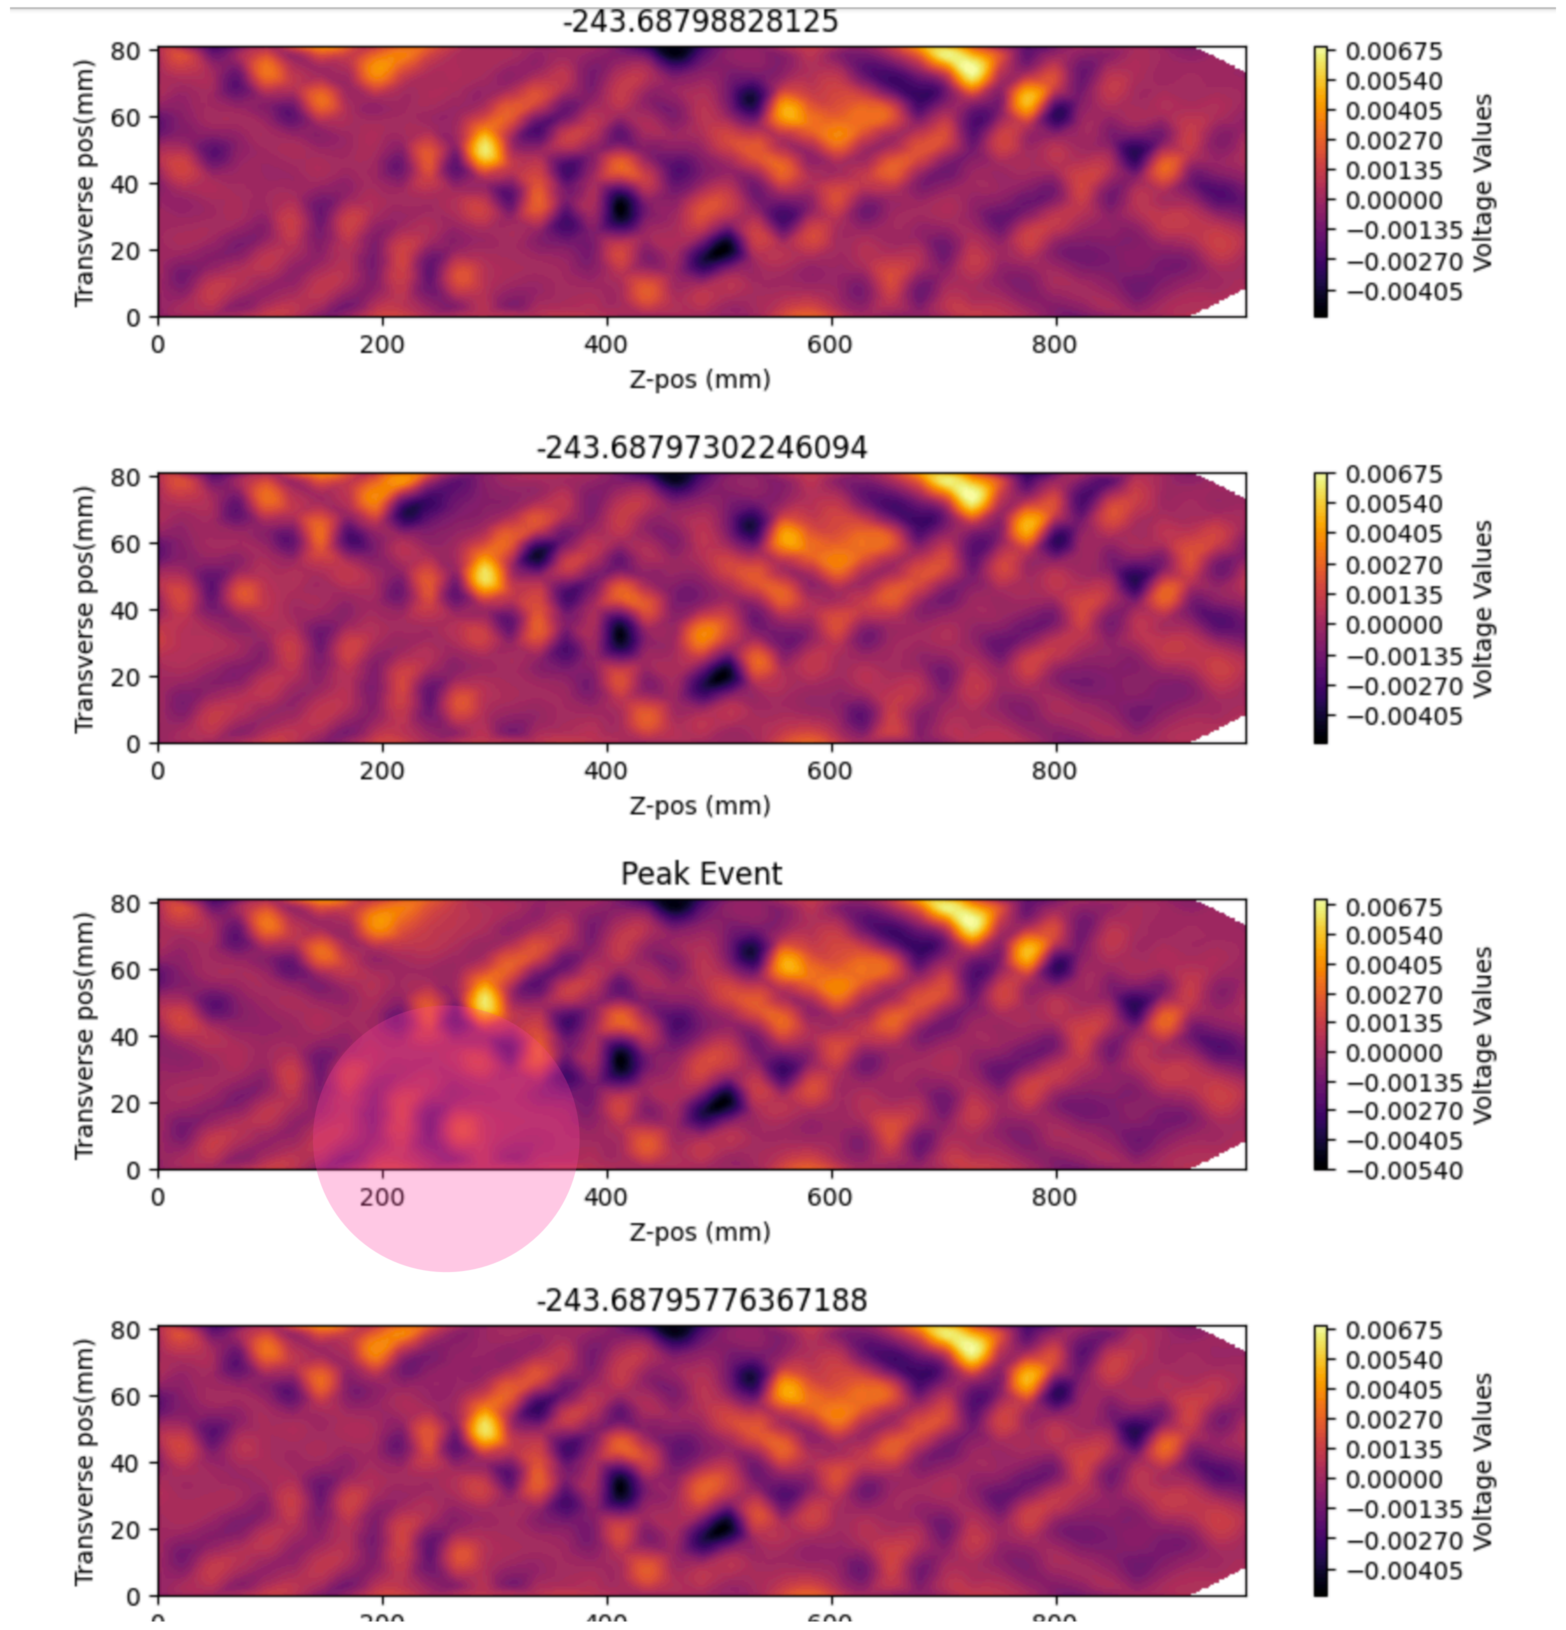

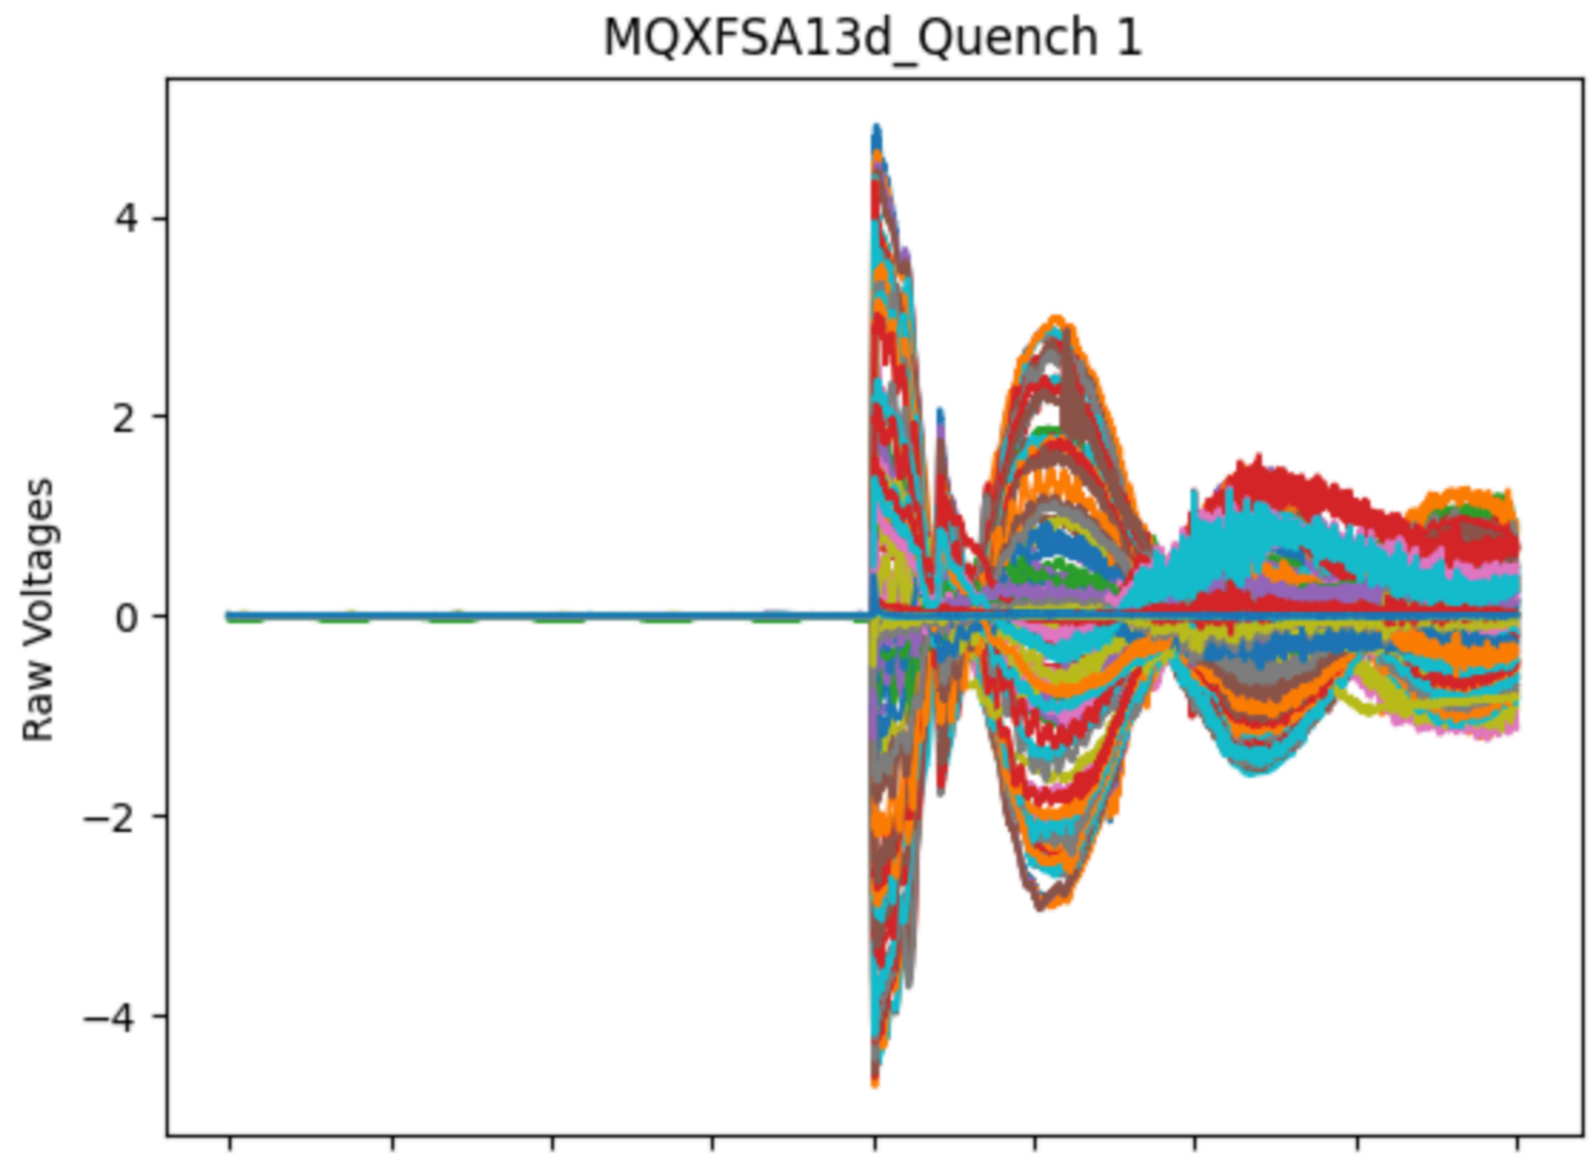

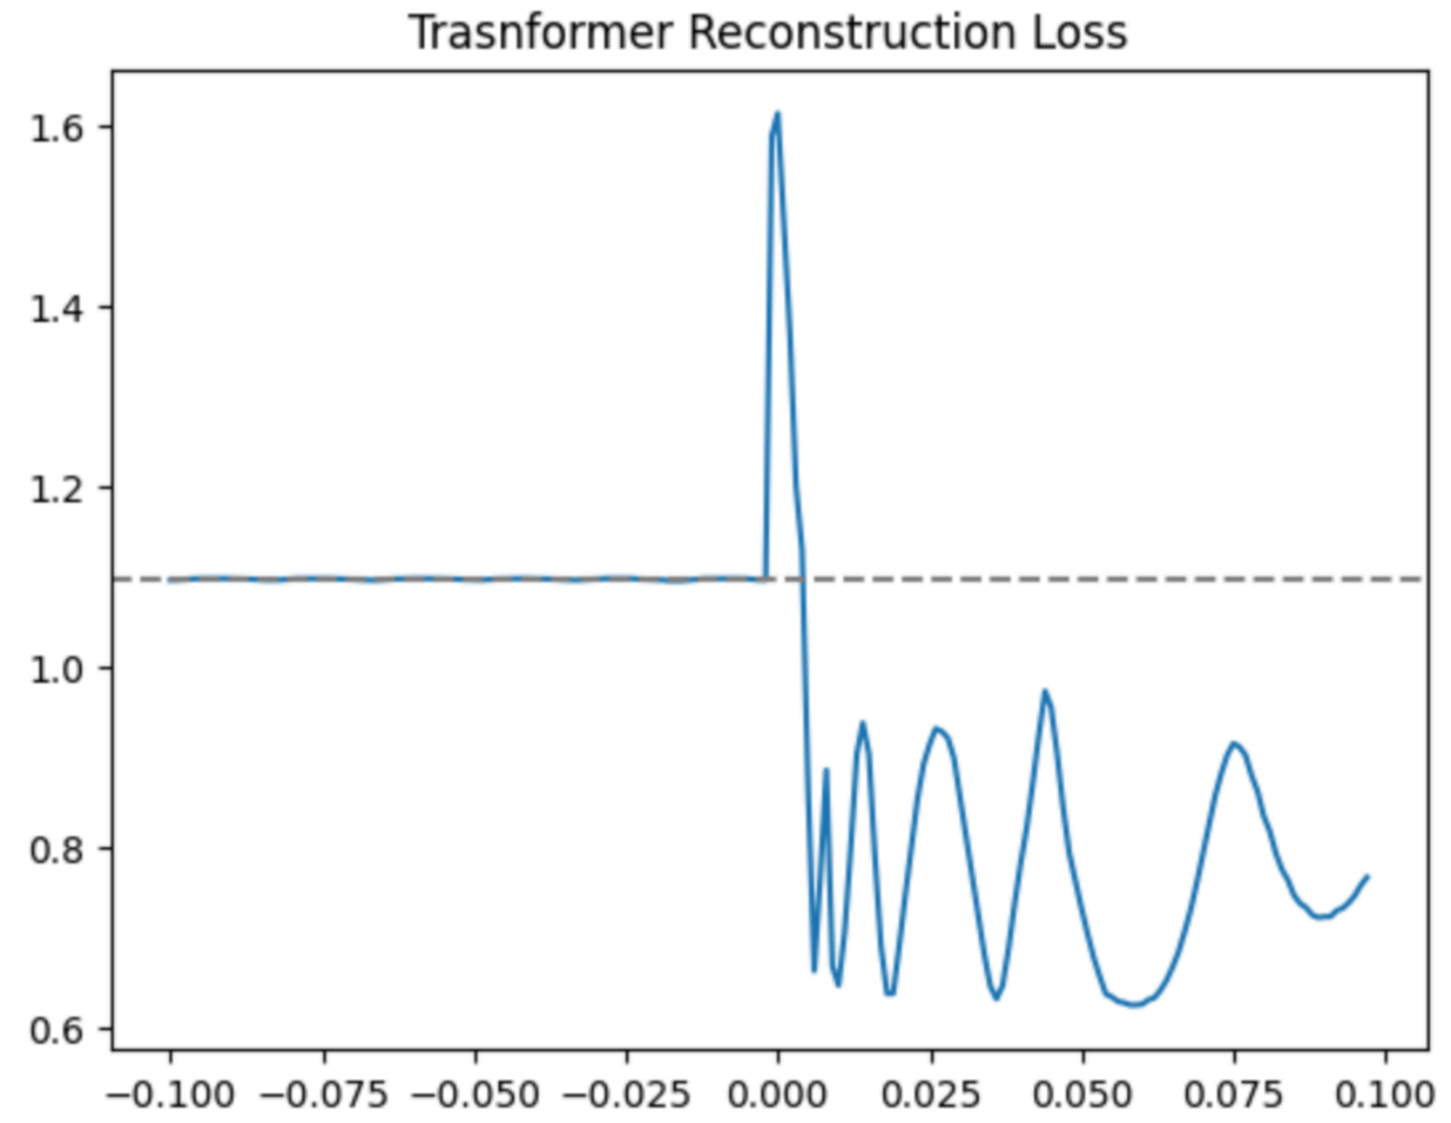

We have overlayed machine learning models with our loss on the actual write geometry

Aim: To see model is able to characterize geometry without explicit direction

We are investigating the use of models using a more consistent dataset: MQXA series magnet to see if our models may better capture the quench boundary or characterize the individual channel

Also, we might return to other feature representations now that r understand the data better

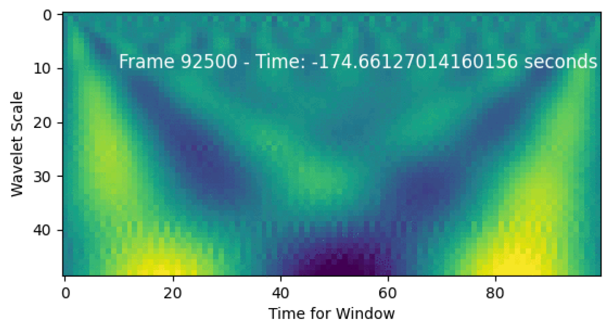

Future: HTS magnet analysis at boundary?

Real Time and Edge Application of Algorithms

	Spectral Entropy Reconstruction	Statistics Computation	Denoising Weight Matrix
GPU processing time of 20ms at 20us step per channel at 16bit precision	1 MHz: ~22us 100KHz:~17us	1 MHz: ~2us 100KHz: ~1us	To be studied
	Dense Autoencoder	**U-Net Autoencoder**	**Transformer Autoencoder**
GPU Inference time of 20ms window inputs at 16-bit precision	Acoustic Trigger: ~54us QA Trigger: ~79us	Acoustic Trigger: ~2ms QA Trigger: ~3ms	QA Trigger: ~12ms
Best Total Time Trigger Efficiency within ~25s	False Triggers: 1/29 Missed Triggers: 0/29 True Triggers: 28/29	False Triggers: 2/29 Missed Triggers: 3/29 True Triggers: 24/29	In progress

	Spectral Entropy Reconstruction	Raw Statistics	Denoising Weight Matrix
Memory Footprint per 20ms window (Bytes)	1 MHz: 16000b 100kHz: 1600 b	1 MHz: 32000B 100kHz: 32000B	To be studied
	Dense Autoencoder	**U-Net Autoencoder**	**Transformer Autoencoder**
Model Parameter Count (SE) **Parameter Count (stats)**	Acoustics: 67,184 QA: 96,956	Acoustics: 463,52 QA: 658,692	Acoustics: 49,280 QA: 196,864

*All timing benchmarks are preliminary, seem stochastic

Conclusions and Outlook

- We have developed a series of tools and metrics for real time denoising and quench characterization
- Our models seem to recover events that occur across the training ramp and establish correlations between acoustic and quench antenna sensor data
- We are still investigating various adaptive weighting algorithms to deal with the multiple types of data and varying window size
- The lack of consistent data in training makes consistent training challenging. Data regularity study to come
- We should eventually evaluate our algorithms on steady state data for actual real time control
- For now, the model serves as an interesting tool for flagging potential precursors and understanding them
- Lots of room for ML in other parts of this problem: objective function optimization, RF type analysis
- Temperature sensors would contextualize many of these readings