Anomaly Detection and Forecasting for the **KFA71/79** Extraction Kicker

SY-ABT-BTP

Algelly Malik

Supervisors: *Francesco Velotti* & *Kostas Papastergiou*

Content

1. Context
EPA WP8 goals,
KFA71/79 importance.

2. KFA71/79 System Overview
Key components,
challenges.

3. Proposed Approach
Anomaly detection
pipeline.

4. Data & Models
Waveform analysis,
VAE/CVAE.

5. Results & Outlook
Insights and future
directions.

6. Continual Learning
Requirements and
benefits.

Context

- **PhD at CERN**: ***Continual Learning*** for anomaly detection and forecasting in accelerator systems.
- Part of the ***Efficient Particle Accelerators*** **(EPA)**[1] project:
  - *Work Package 8* **(WP8) : *Equipment Automation***[2] : reduce downtime through ***predictive maintenance*** and ***automation of critical equipment***.
- **Current focus**: Anomaly forecasting for the ***KFA71/79 extraction kicker magnet*** in the Proton Synchrotron (PS).

System Overview – The KFA71/79 Extraction Kicker

Purpose: Fast-pulsed magnet system to extract particle beams from the Proton Synchrotron (PS).

Components: 12 generator modules operating simultaneously in vacuum tanks.

- **9 modules**: Section 71 (KFA71).
- **3 modules**: Section 79 (KFA79).

Output: High-voltage pulses (~80 kV peak, ~4 μs total duration).

Focus: Main region of interest is ~1 μs within the pulse.

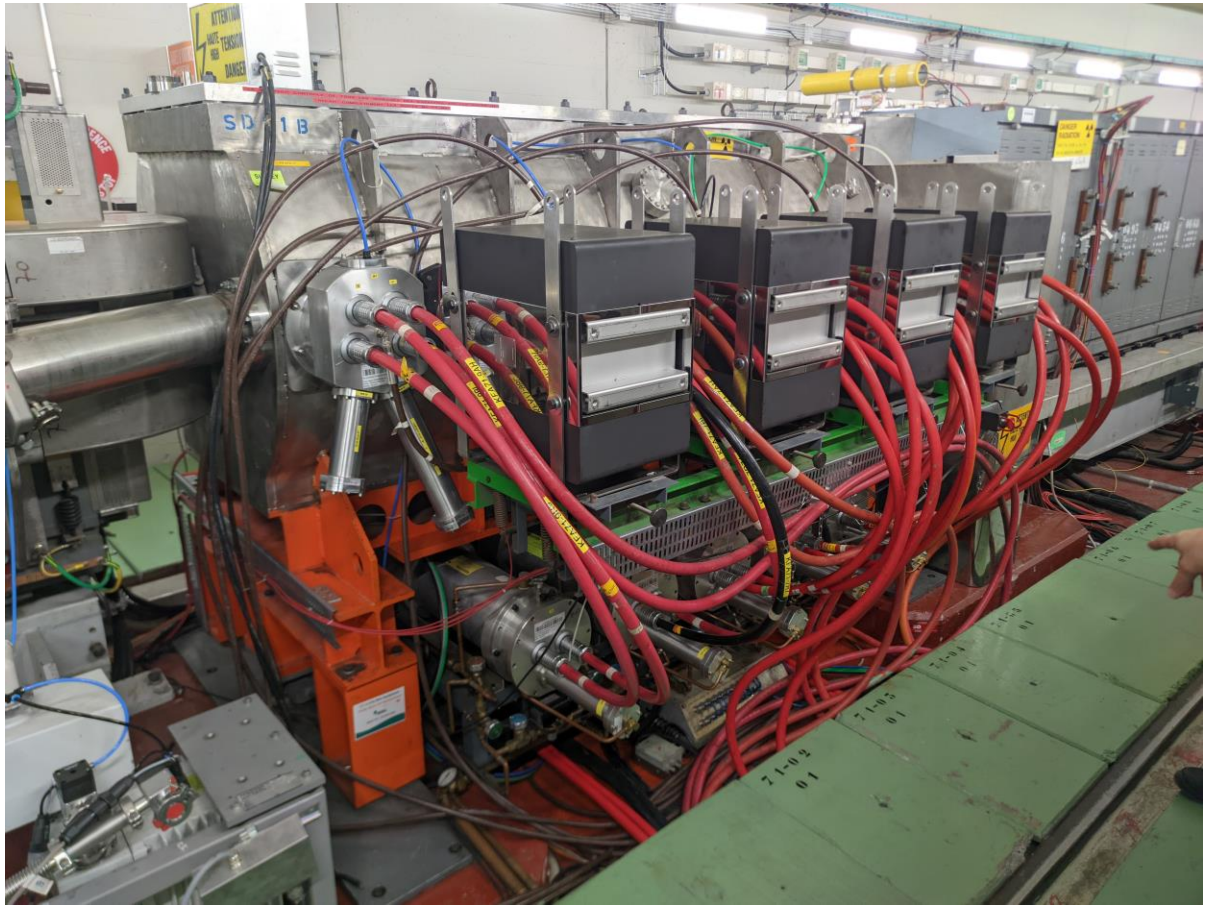

Picture: Modules of KFA71

System Overview – The KFA71/79 Extraction Kicker

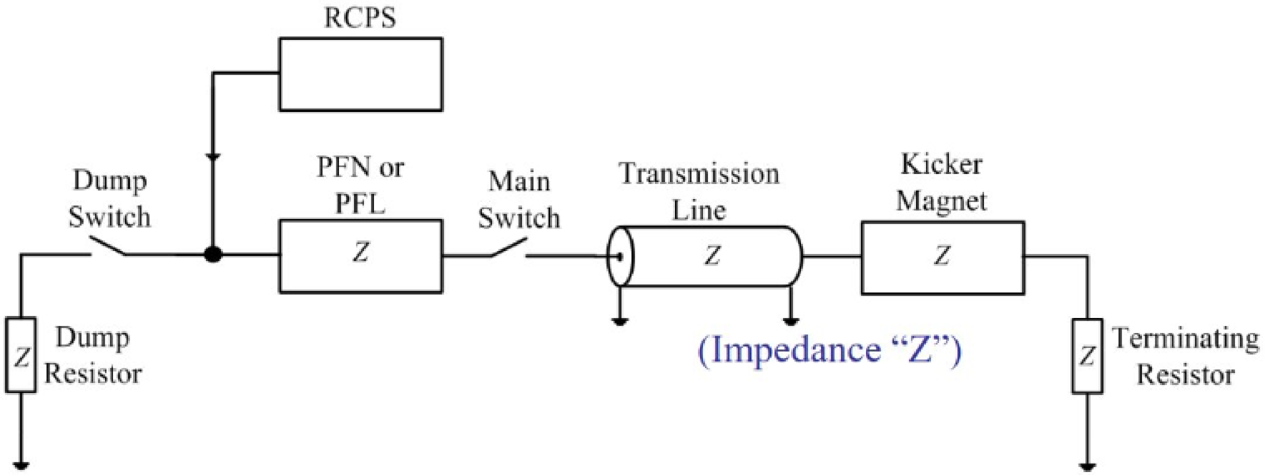

Schema: *Simplified schematic of a kicker module* [4]

PFL: Pulse Forming Line
RCPS: Resonant Charging Power Supply

- **Charging**: RCPS charges the PFL to the required voltage.
- **Switch Timing:** The Main Switch closes at the precise moment for beam extraction.
- **Pulse Generation**: High voltage discharges through the magnet.
- **Energy Dumping:** The Dump Switch redirects remaining energy.
- **Beam Extraction**: The beam is deflected onto the transfer line.

System Overview – The KFA71/79 Extraction Kicker

*Picture: CPS **PFL** DRUM Winding 2005*

System Overview – The KFA71/79 Extraction Kicker

Historical and Future Outlook

- System installed in the 1970s.
- Undergoing a major consolidation project during LS3 [5] to improve:
  - Reliability and availability.
  - Safety (e.g. replacing mineral oil with ester oil, managing SF_6 gas issues).
  - Diagnostics and remote monitoring.
  - Obsolescence and environmental impact.

Why Focus on KFA71/79

- Complex sub-components: HV switches (thyratrons), cables, transmission lines, ferrite magnets.
- High risk of beam losses during module failure.
- Rich waveform datasets enable machine learning studies.
- Aging system = higher anomaly rates.

Current Monitoring Limitations

- Threshold-based alarms on selected signals.
- Detection occurs **after anomalies happen**.
- Reactive maintenance, not proactive or predictive.
- Gaps in addressing long-term reliability.

Data Description

Waveform Characteristics:

- ***Sampling Rate***: 1 sample every 4 ns for 10 μs → 2500 data points
- ***Signal Content***: Short rise and fall times, short plateau region.
- ***12 generators*** → 12 waveforms per cycle
- ***Pulse Settings***: Includes desired pulse **strength**, pulse **length**, enabled generators, etc.

Waveforms have been stored in NXCALS since the end of September 2024

→ Current analysis and training focus on October 2024 data.

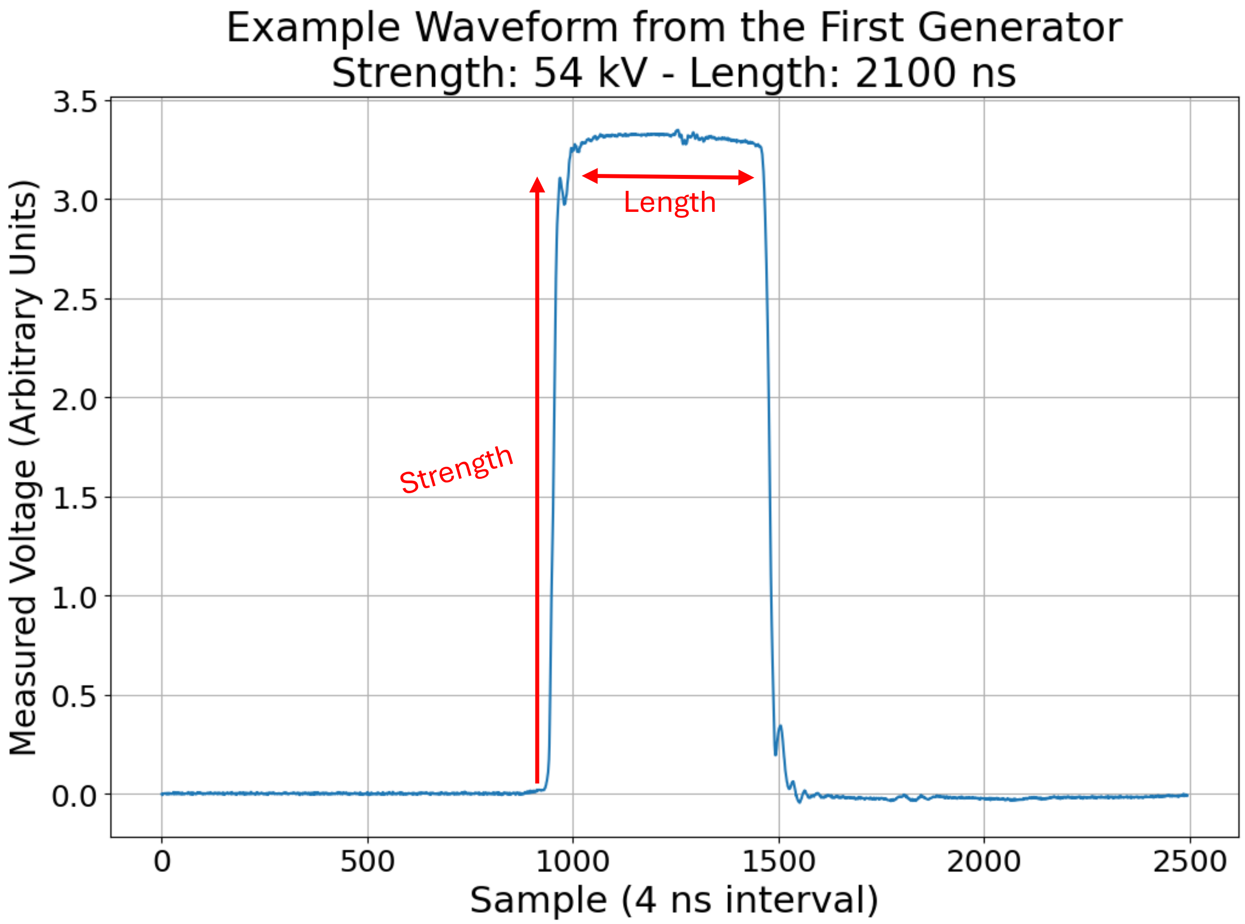

Proposed Approach: Anomaly Detection & Prediction

General Idea:

1. Train a model on nominal waveforms.

2. Detect deviations using performance metrics.

3. Monitor trends to identify drifts or early anomalies.

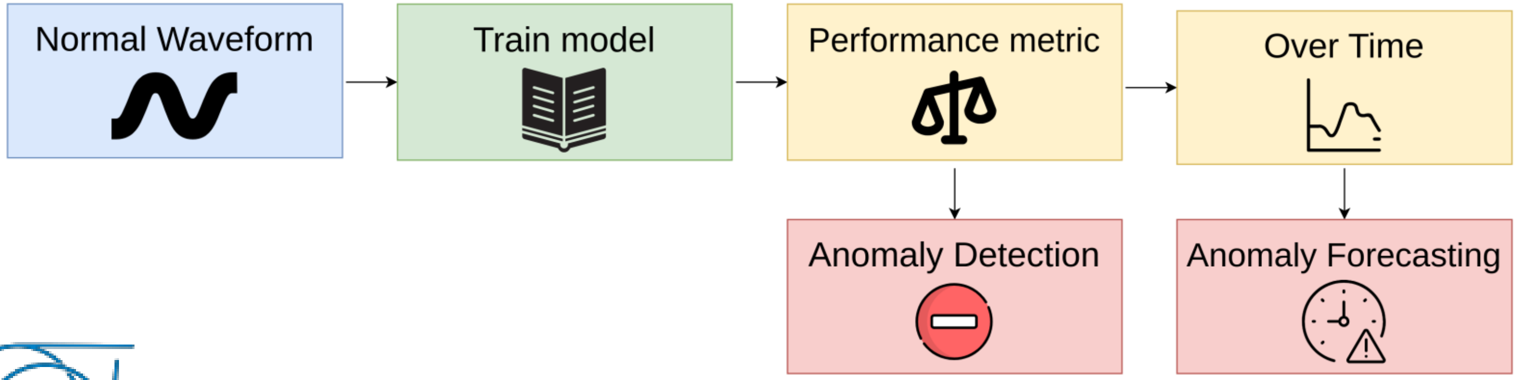

Goals:

- Real-time anomaly detection.
- Minimized reliance on labeled failures.
- Adaptation to variations via **Continual Learning**.
- Automated recovery support.

Data Description

Data & Settings Overview:

- **Enable Setting:** 12-bit integer, indicating which generators are switched on or off (e.g., 4095 if 12 generator enabled).

- **Main Strength Setting:** Divided among enabled generators (e.g., if 9 out of 12 are enabled, total strength is split by 9).

- **Main Length Setting:** Applied both globally and per generator to define the pulse duration.

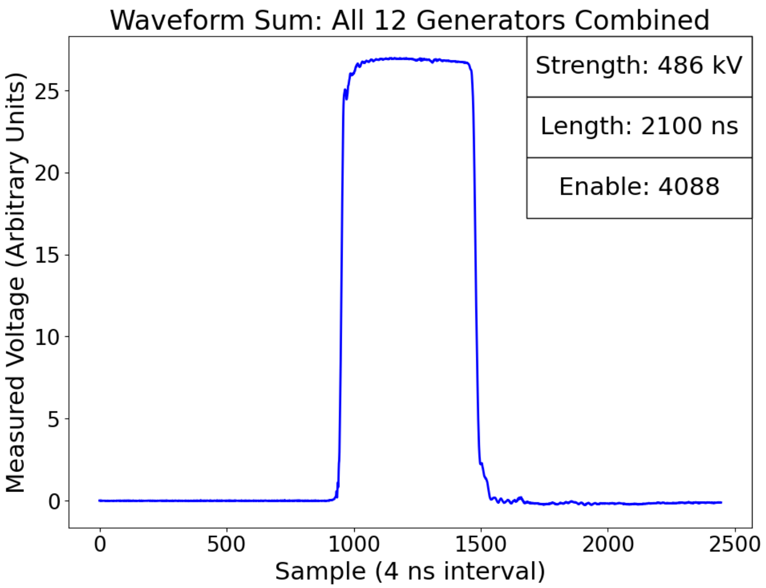

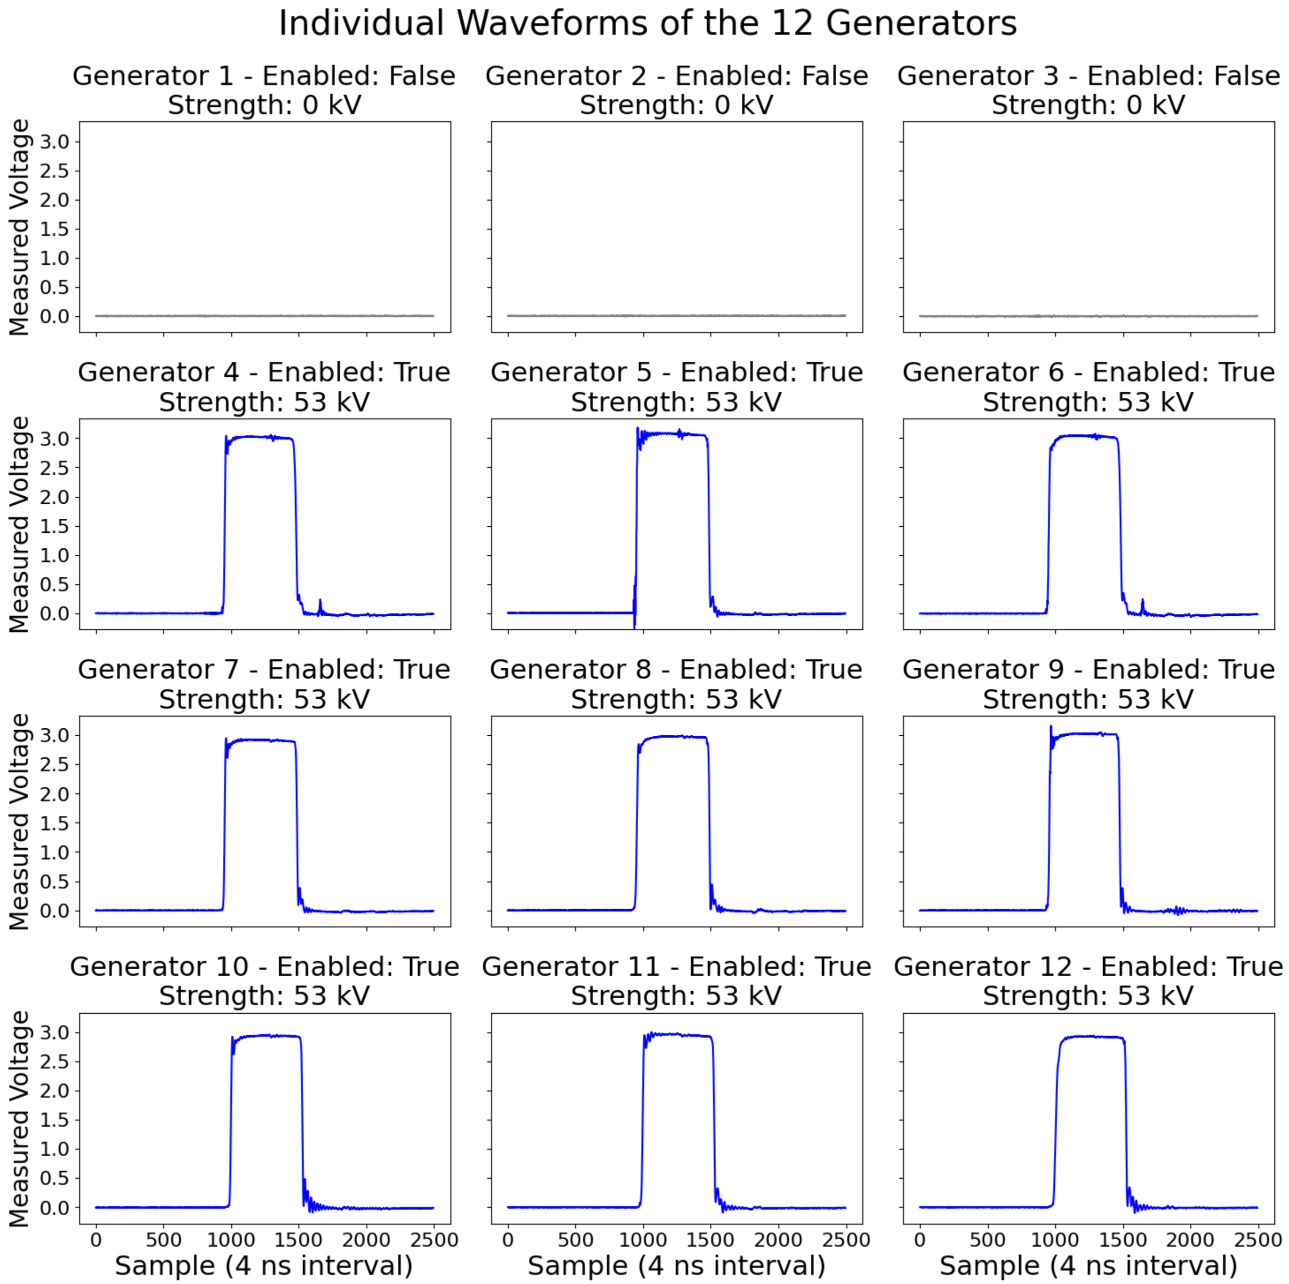

Settings Distribution – **Strength** and **Length**

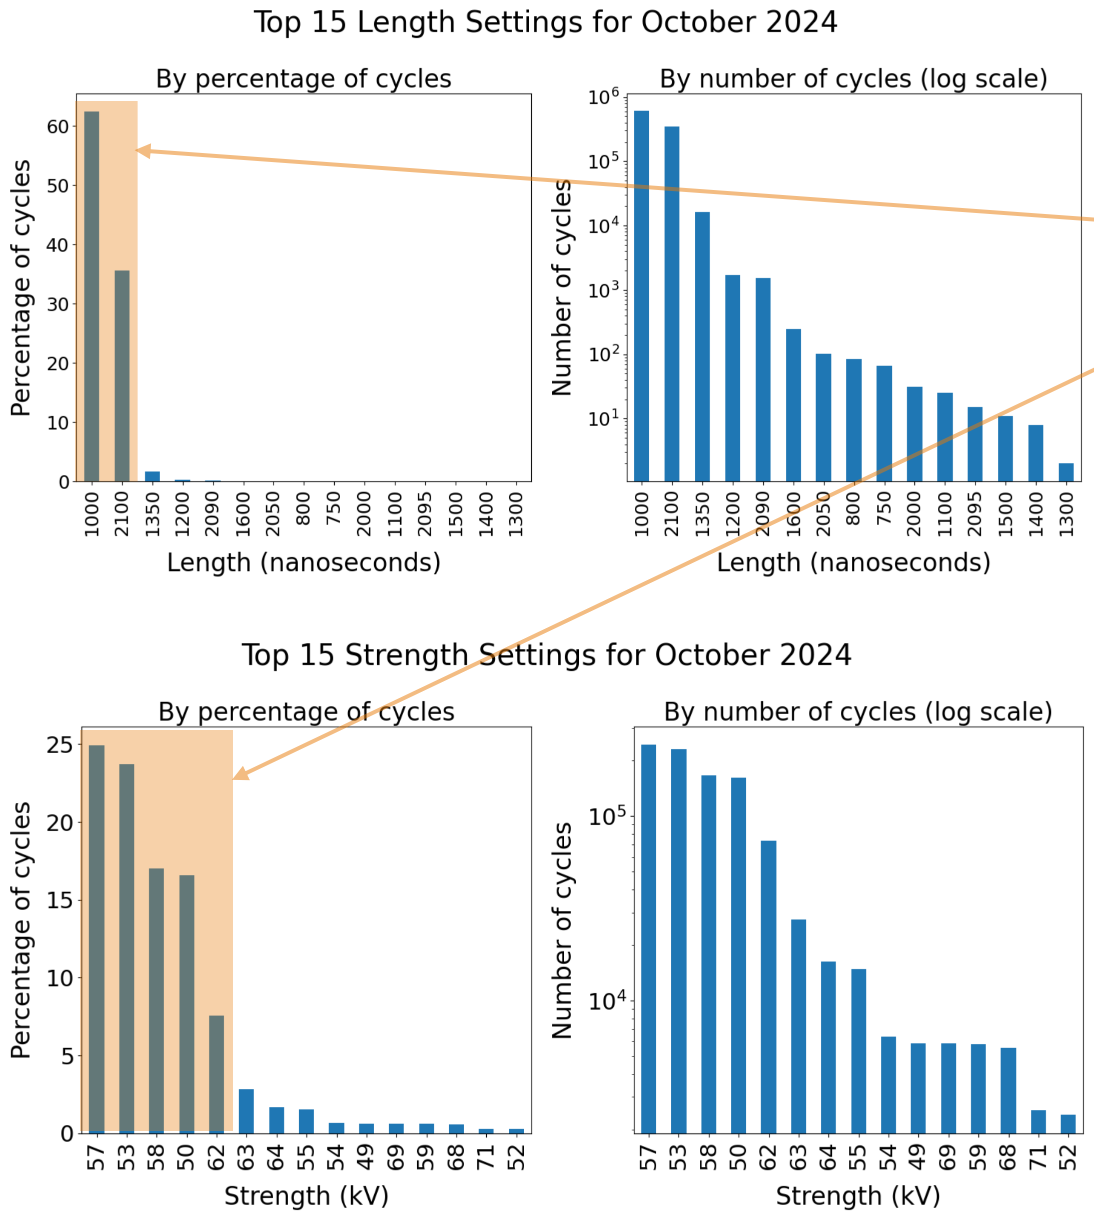

<u>Key Challenges and Implications:</u>

Imbalanced Data:

- **~98%** of waveforms share 2 length values.
- **~80%** of waveform share 6 voltage values.
- Risk: **Overfitting** to dominant settings.

Consequences:

- Rare configurations **misclassified** as anomalies.
- Reduces detection accuracy and increased biases.

Recommendations

- **Dataset Balancing:** Sampling, augmentation, reweighting.
- **Performance Monitoring:** Focus on rare settings.
- **Leverage Diversity:** Use rare configurations to improve robustness.

Labeling Process: Challenges and Key Steps

Context:

- How to label a subset of waveforms from millions of records ?

Approach:

- **Comparing IPOC Data**:
  - Measured pulse properties against expected settings.
  - Detect issues like *missing pulses* or *faulty shot*.
- **Median Waveform Computation**:
  - Group waveforms by *strength* and *length*.
  - Compute median waveforms.
  - Compute deviations (e.g., L2 norm).

Outcome:

- Preliminary set of anomalies for evaluation.
- Manual verification feasible due to reduced candidate anomalies.

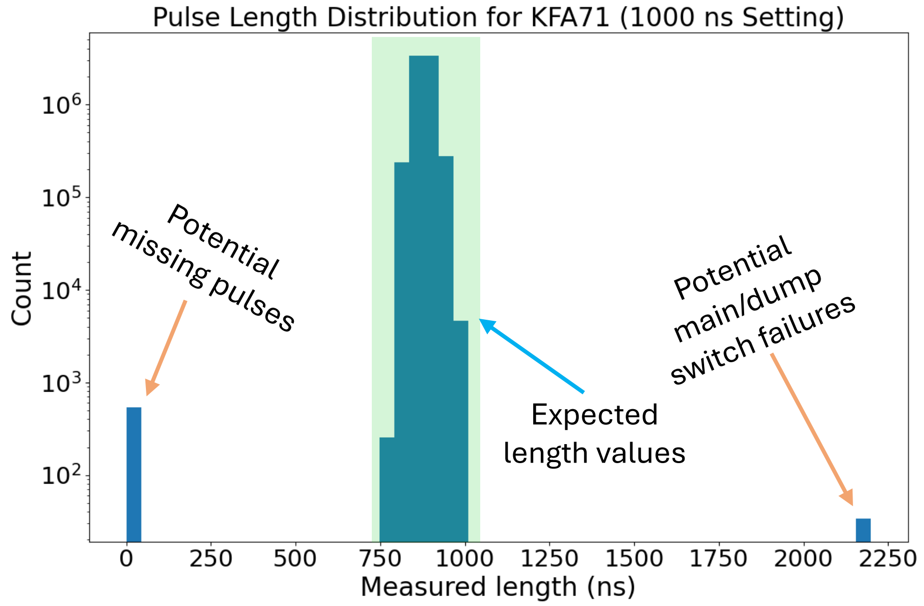

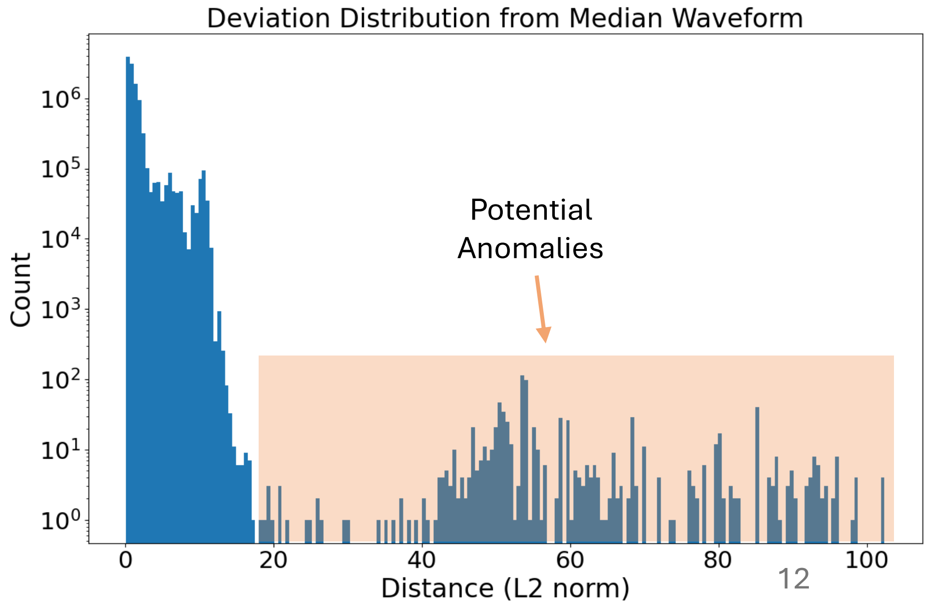

Data Approach

Moving from *Multi-Generator* to *Per-Generator Analysis*

Initial Approach:

- **Data Format**: Treated 12 waveforms as a 1-channel image (12 rows × 2500 columns).
- **Grouping**: Total strength, pulse length, and enable settings combined into a single data point.
- **Challenges:**
  - Complex relationships between waveforms, strength, and enable settings were hard to learn.
  - Increased model complexity reduced interpretability and made training harder.

Per-Generator Approach:

- **Independent Circuits:** Each generator operates with separate, parallel circuits.
- **Simplified Modeling**: Analyze one generator's waveform at a time.
- **Benefits**:
  - **Reduced Complexity**: Easier to train models.
  - **Better Interpretability**: Anomalies traced to specific generators.
  - **Faster Data Handling**: Fewer variables for each analysis.

Data Approach

In the following sections, we focus specifically on the first generator of the KFA71

→ more than **600k** cycle where this generator should have pulsed.

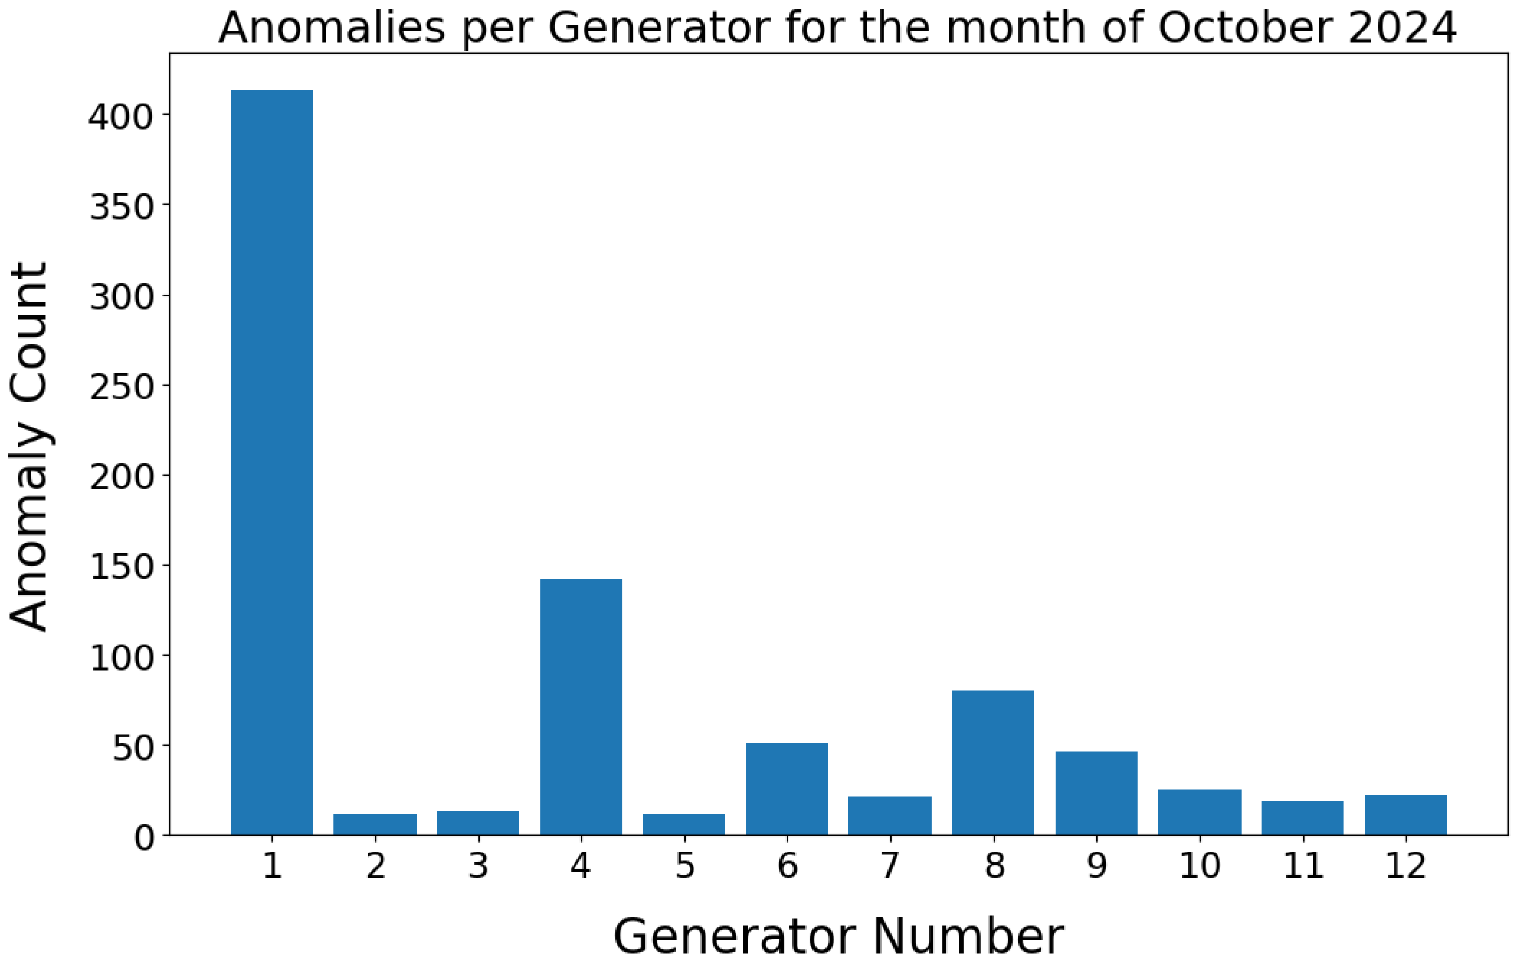

Preprocessing - Creating Training and Validation Dataset

Creating Training Dataset

	Number of Sample
1. Time Period Selection: • October 1st to 17th.	→ ~500k
2. Data Filtering: • Removed inactive generator cycles and known anomalies.	→ ~350k
3. Setting Combination Processing: • Kept combinations with ≥100 cycles (from 89 to 32 combinations).	→ ~349k
4. Setting Balance: • Balanced strength/length settings	→ ~600

Creating Validation Dataset

Three subsets:

1. Settings **present** in training.
2. Settings **absent** in training.
3. Known anomalies.

Key Features

- Excludes inactive generator cycles.
- Removes low-sample combinations (<10).
- Limits to max 50 samples per combination for diversity.
- **October 18th to 24th**, including anomalies for evaluation.

ML Models – Variational Autoencoders (VAE)

Challenge:

- **Unlabeled Data**: Supervised learning not applicable

Goal:

- Model normal waveform distribution
- Minimize reconstruction error.

VAE Components:

- **Encoder**: Maps waveforms to a compressed representation.
- **Latent Space**: Probabilistic representation.
- **Decoder**: Reconstructs waveforms.

Loss Function:

- **Reconstruction Loss (MSE)**: Measures how well the waveform is reconstructed.
- **KL Divergence**: Aligns latent space with a Gaussian distribution.

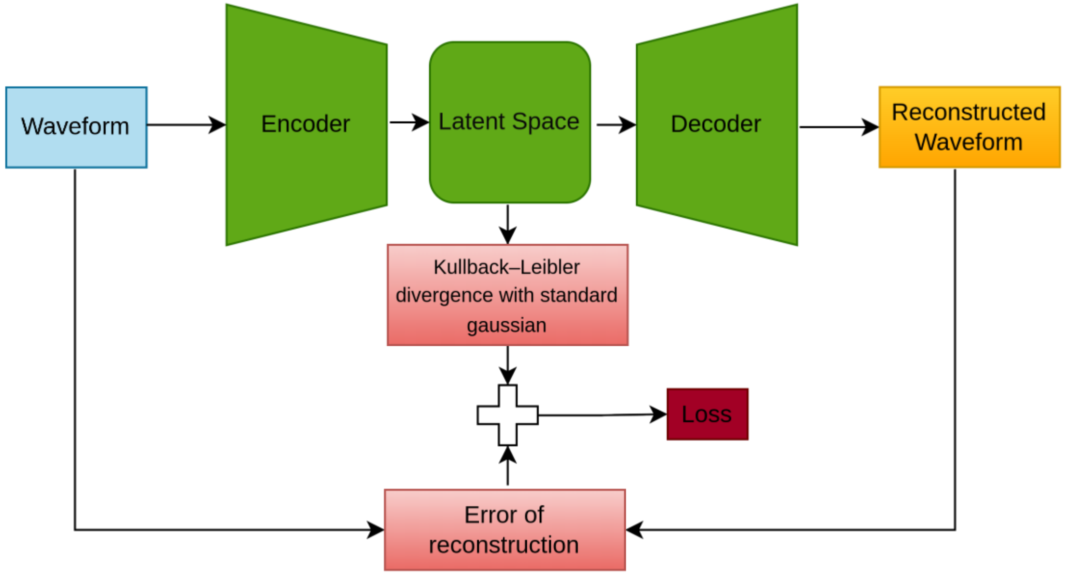

Mathematical Formulation:

$$\mathbb{E}[L] = (1-\epsilon)\cdot\mathbb{E}[l_w(X_{good})] + \epsilon\cdot\mathbb{E}[l_w(X_{anom})]$$

- $l_w(X)$: Loss for waveform X with model's weights w.
- ϵ : Fraction of anomalies.

Anomaly Detection

High Error: Indicates data deviates from normal distribution → **Anomaly**
Low Error: Data aligns with normal waveforms → **Normal**

ML Models – Conditional Variational Autoencoders (CVAE)

Problem with VAEs:

- VAEs often **reconstruct subtle anomalies too well** [6].
- Assume all data shares one global distribution.

Goal:

- Use contextual information (**e.g., strength, length**)

Conditioning Inputs:

- Incorporates external conditions (settings) into the model.

Why CVAEs Solve This:

- Learn conditional distributions, not a single global one.
- Use conditions to adapt reconstructions
- Increasing errors for out-of-context anomalies.
- Generalize better with **conditional distributions**. [14].

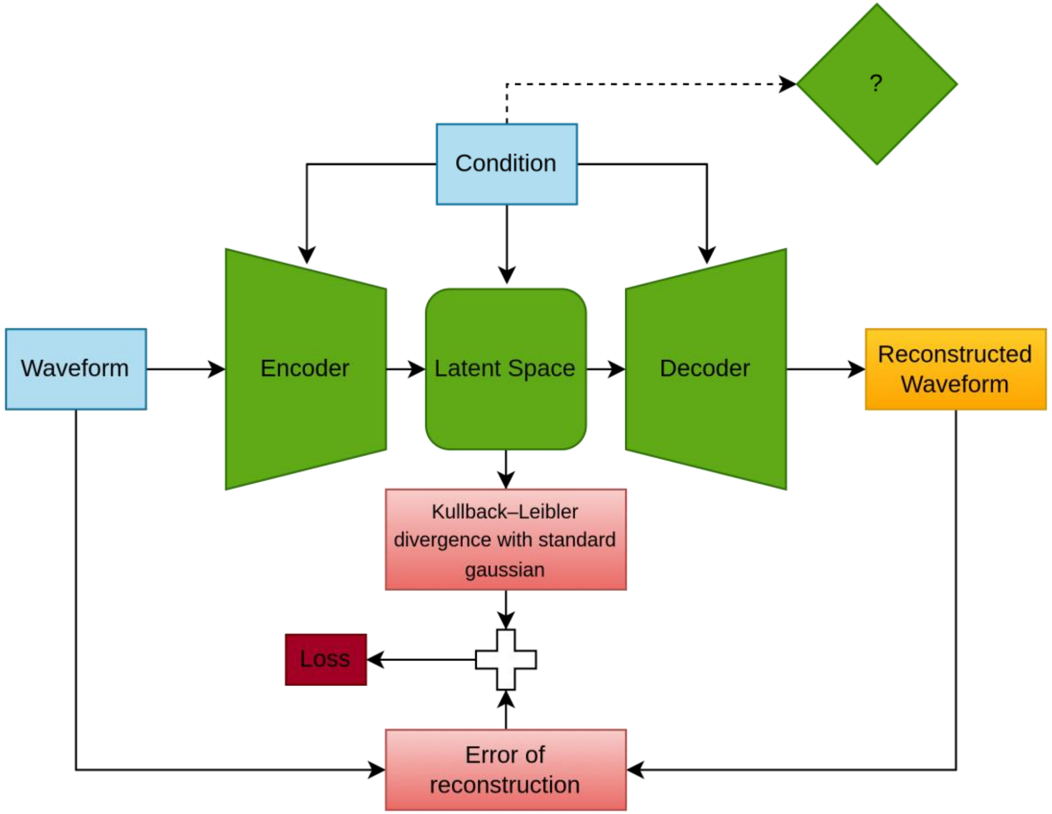

ML Models – CVAE 1

Specificity:

- **Condition Integration:** Conditions are concatenated with the latent space before reconstruction.
- **Contextual Reconstruction:** Adjusts outputs based on waveform settings.

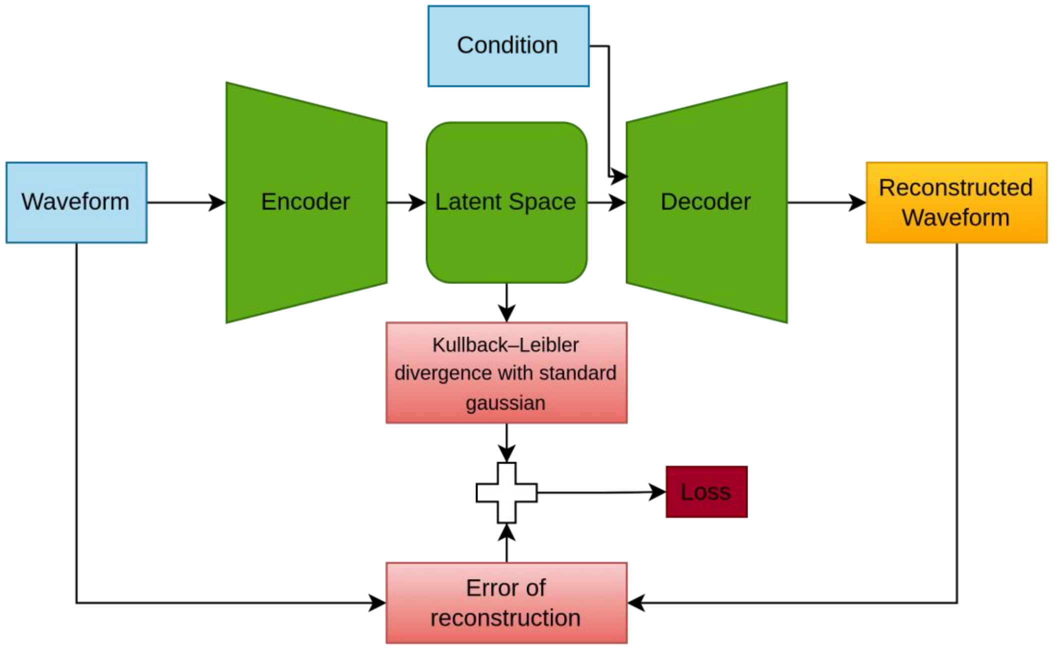

ML Models – CVAE 1

Specificity:

- **Condition Integration:** Conditions are concatenated with the latent space before reconstruction.
- **Contextual Reconstruction:** Adjusts outputs based on waveform settings.

Issues:

- **Over-reliance on Conditions**: Neglects actual waveform data.
- **Anomalies Dilution**: Reconstructs averaged waveforms (e.g., missing pulses ignored).
- **Still in Progress:** Requires tuning to enhance performance.
- **Plot:** Reconstruction matches the waveform based on the conditions, ignoring the input anomaly.

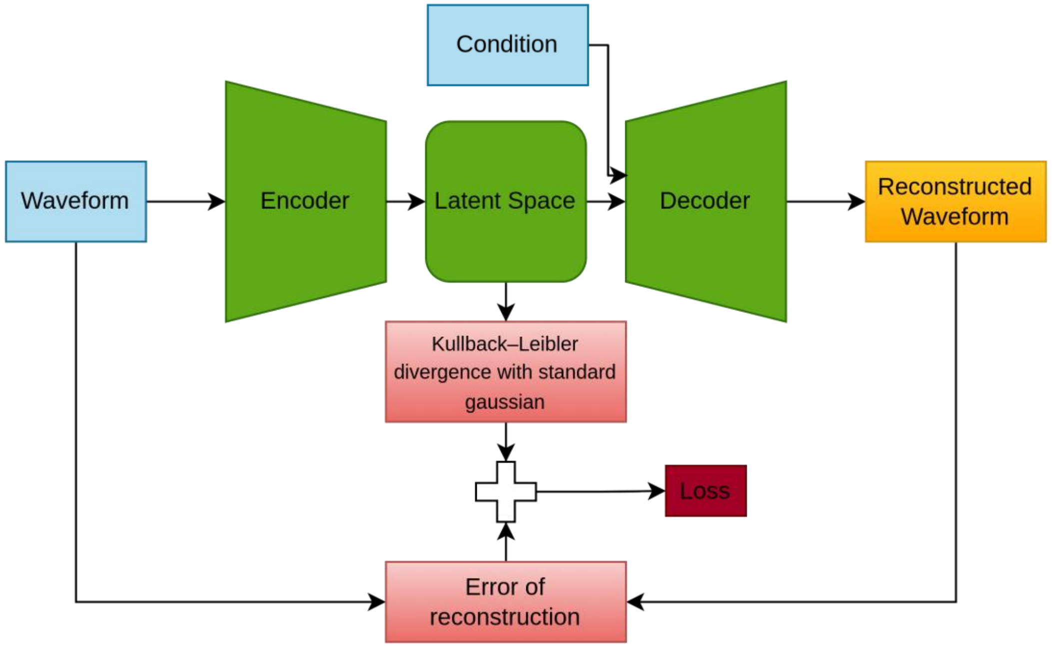

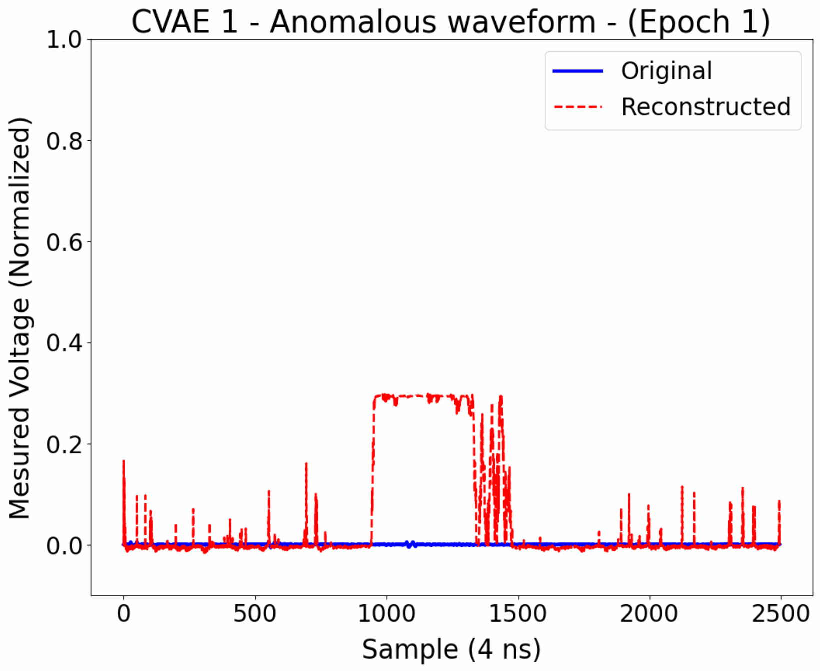

ML Models – CVAE 2

Specificities:

- **Condition Decoder:** Predicts the conditions from the latent space using an additional neural network.
- **Prediction Loss (MAE):** Minimizes error between predicted and true conditions.
- **Indirect Conditioning:** Decoder uses latent space information without direct condition input.
- **Latent Space:** Encodes condition-specific features for better representation.

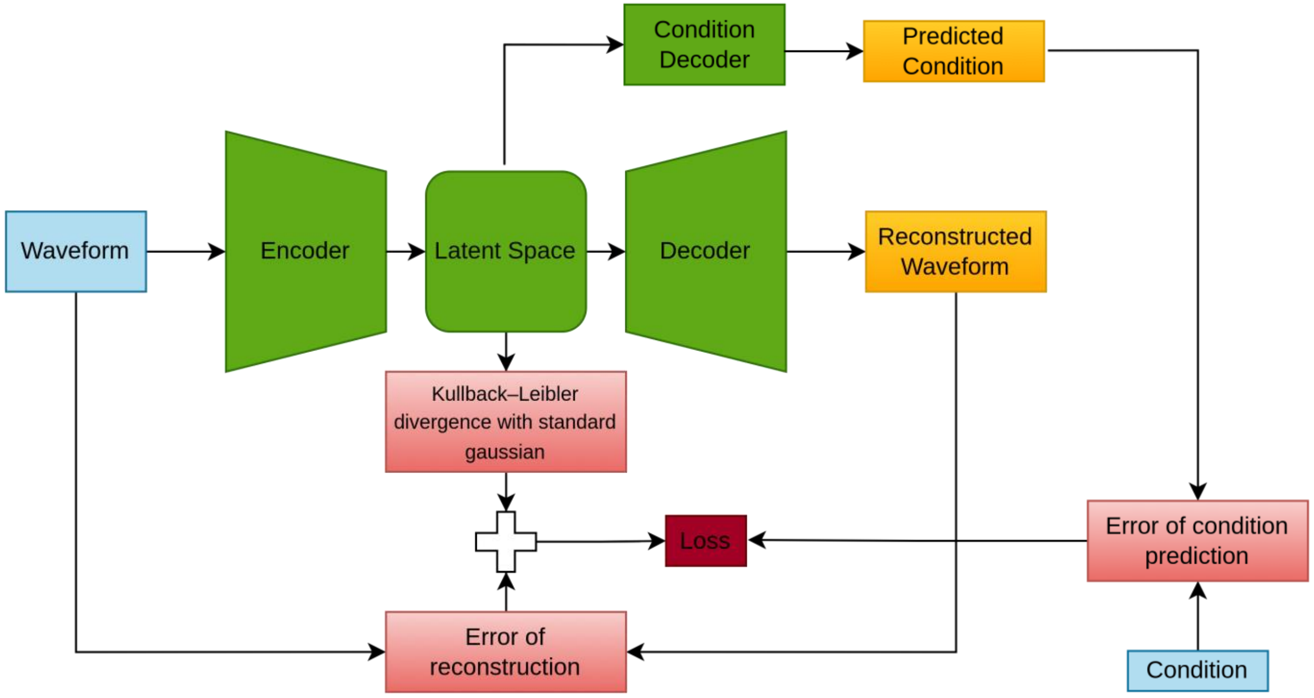

ML Models – CVAE 2

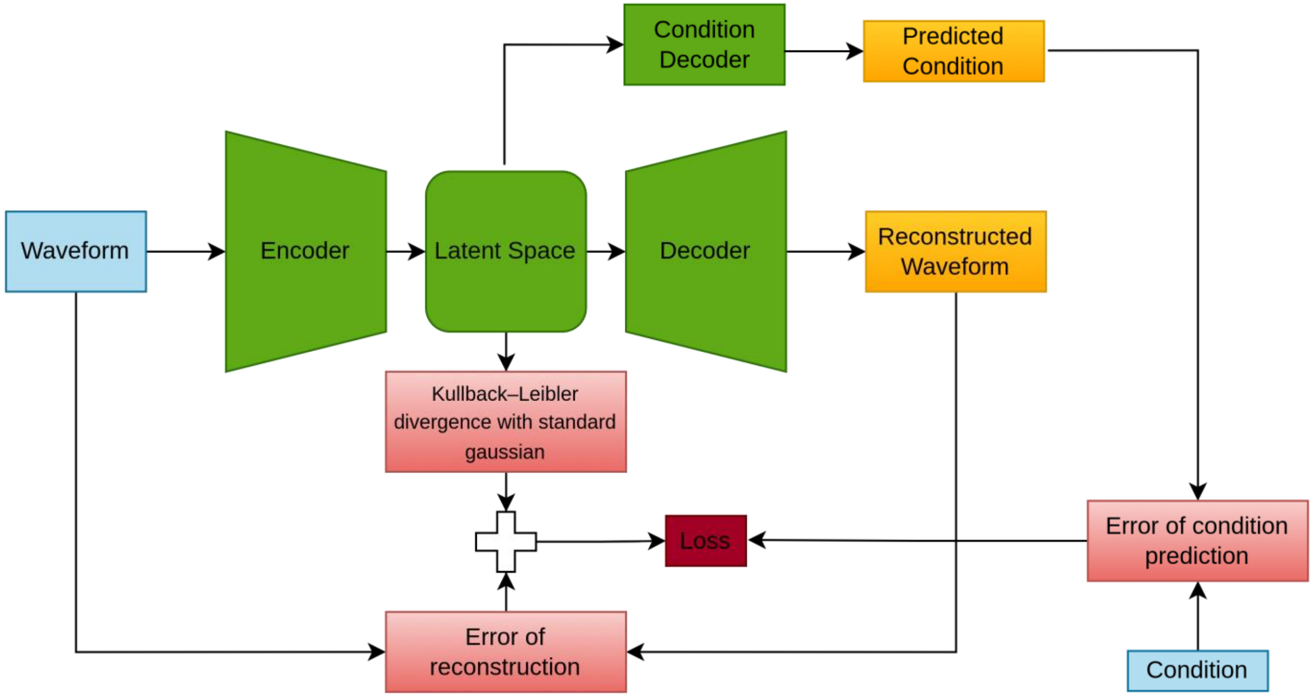

Specificities:

- **Condition Decoder:** Predicts the conditions from the latent space using an additional neural network.
- **Prediction Loss (MAE):** Minimizes error between predicted and true conditions.
- **Indirect Conditioning:** Decoder uses latent space information without direct condition input.
- **Latent Space:** Encodes condition-specific features for better representation.

Advantages:

- **Reduced Over-reliance:** Avoids direct dependence on condition inputs, improving robustness.
- **Improved Performance:** Addresses issues of CVAE 1 by better capturing waveform anomalies.

- **Plot**: No waveform-like reconstruction for large anomalies

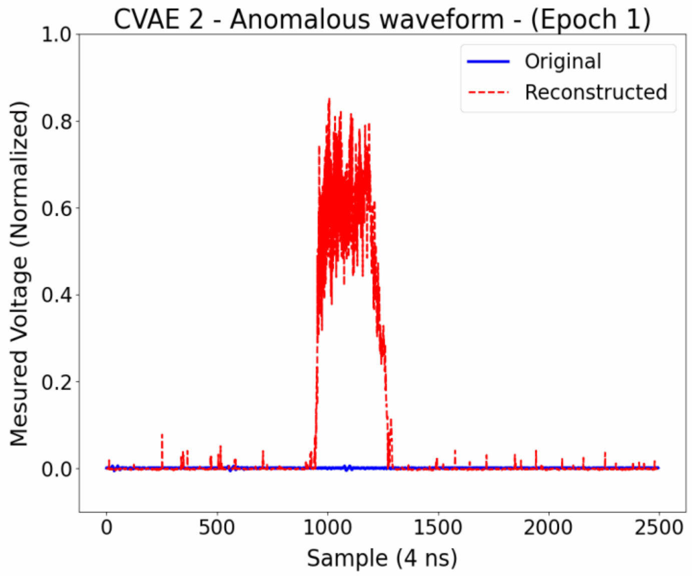

ML Models – CVAE 2: Architecture

```
==========================================================================================
Layer (type:depth-idx)                   Output Shape              Param #
==========================================================================================
CVAE 1                                   [1, 1, 2496]              --
├─WaveformEncoder: 1-1                   [1, 18]                   --
│    └─Conv1d: 2-1                       [1, 8, 1243]              96
│    └─Conv1d: 2-2                       [1, 16, 619]              912
│    └─Conv1d: 2-3                       [1, 32, 309]              1,568
│    └─Linear: 2-4                       [1, 18]                   178,002
│    └─Linear: 2-5                       [1, 18]                   178,002
├─ConditionDecoder: 1-2                  [1, 2]                    --
│    └─Linear: 2-6                       [1, 30]                   570
│    └─Linear: 2-7                       [1, 2]                    62
├─WaveformDecoder: 1-3                   [1, 1, 2496]              --
│    └─Linear: 2-8                       [1, 9888]                 187,872
│    └─ConvTranspose1d: 2-9              [1, 16, 619]              1,552
│    └─ConvTranspose1d: 2-10             [1, 8, 1243]              904
│    └─ConvTranspose1d: 2-11             [1, 1, 2495]              89
==========================================================================================
Total params: 549,629
Trainable params: 549,629
Non-trainable params: 0
Total mult-adds (Units.MEGABYTES): 4.02
==========================================================================================
Input size (MB): 0.01
Forward/backward pass size (MB): 0.50
Params size (MB): 2.20
Estimated Total Size (MB): 2.70
==========================================================================================
```

Results – Validation Normal

Best Model:

CVAE 2 provided the best results.

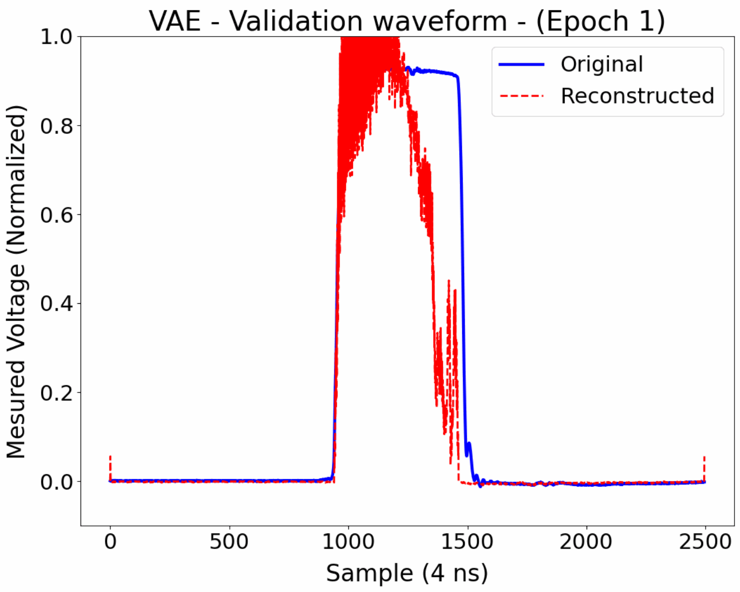

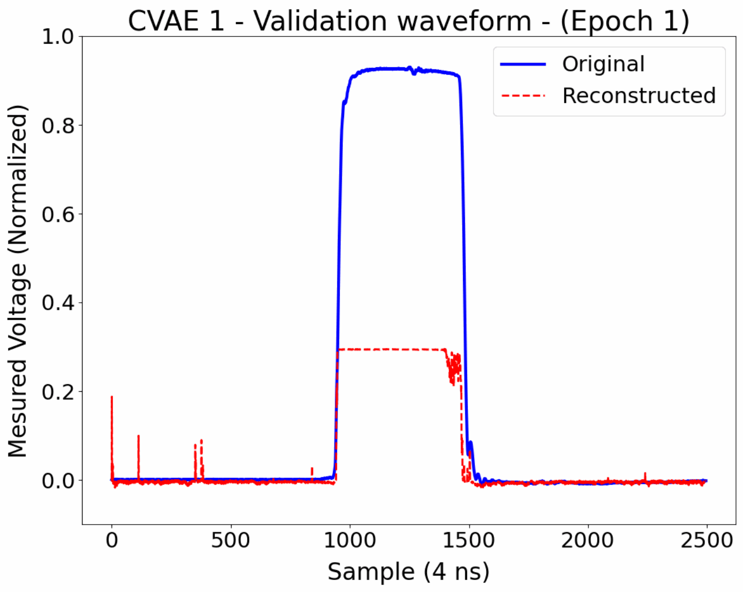

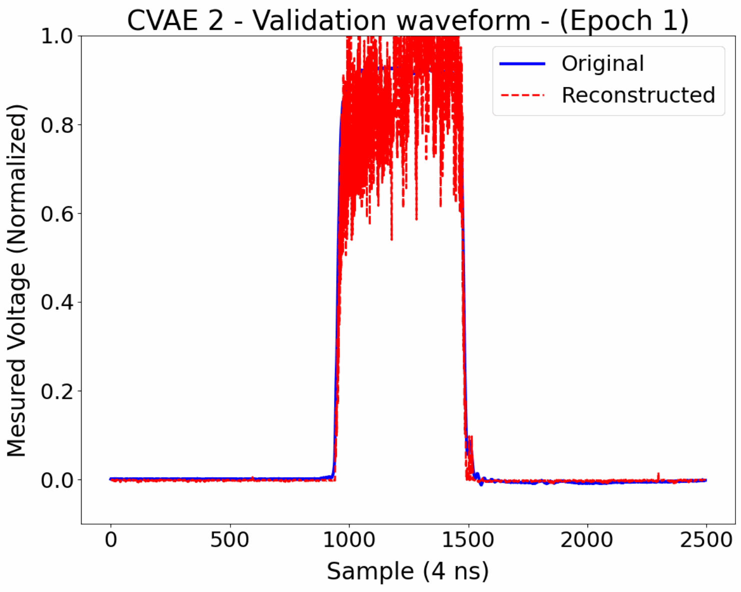

CERN

Results – Validation New Settings

Best Model:

CVAE 2 provided the best results.

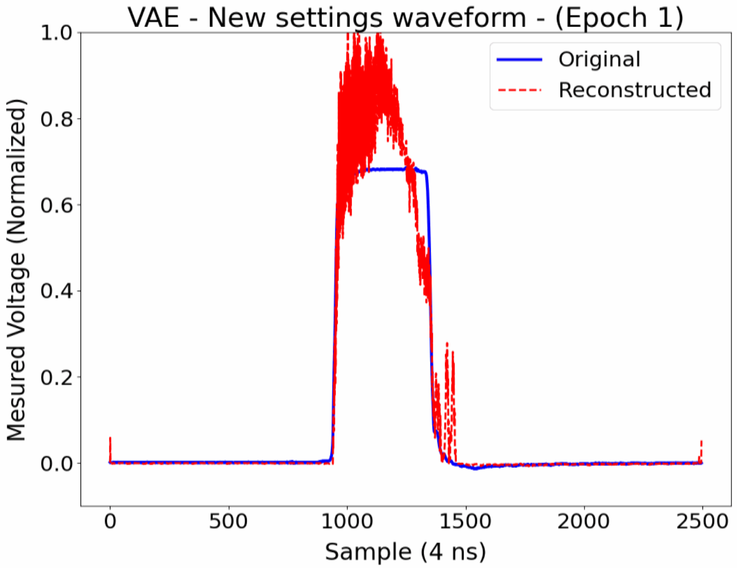

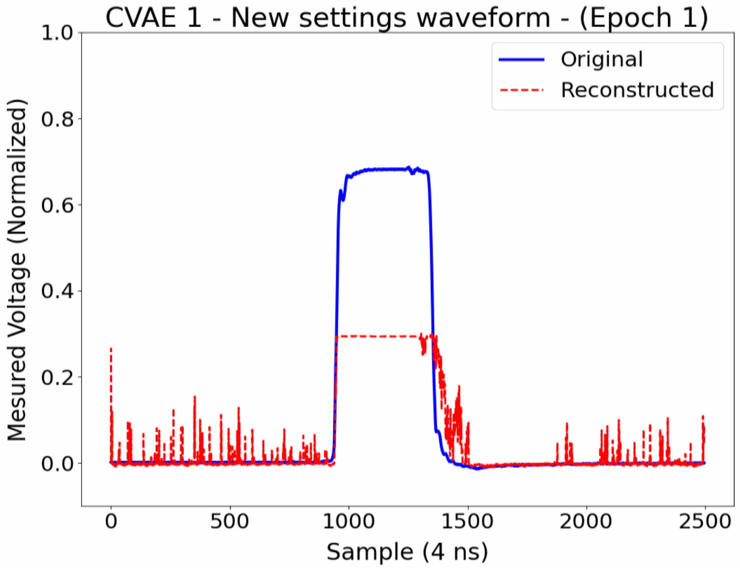

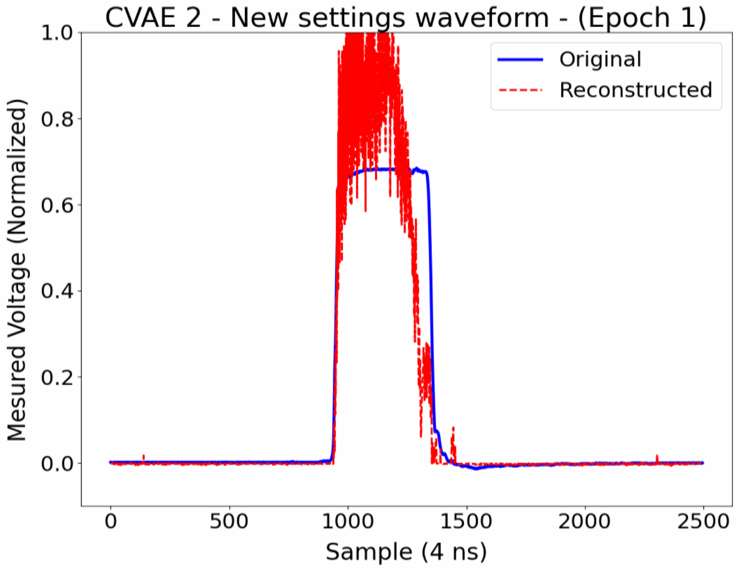

Results – Validation Anomalies

Best Model:

CVAE 2 provided the best results.

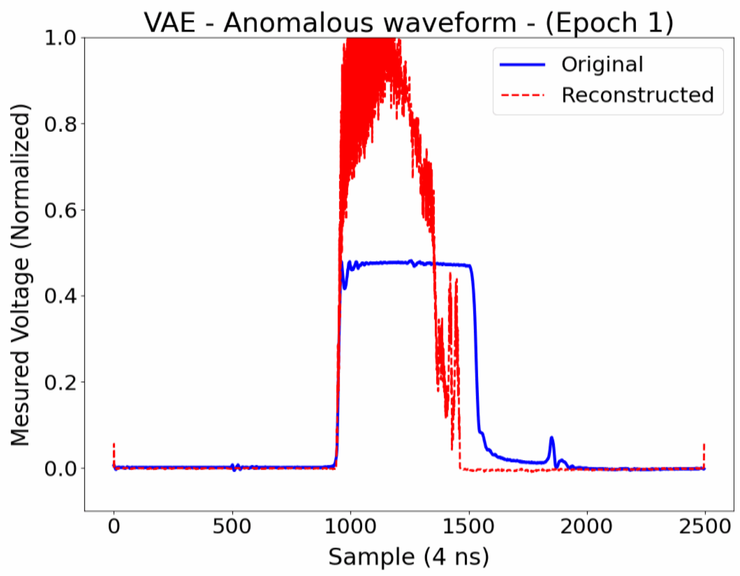

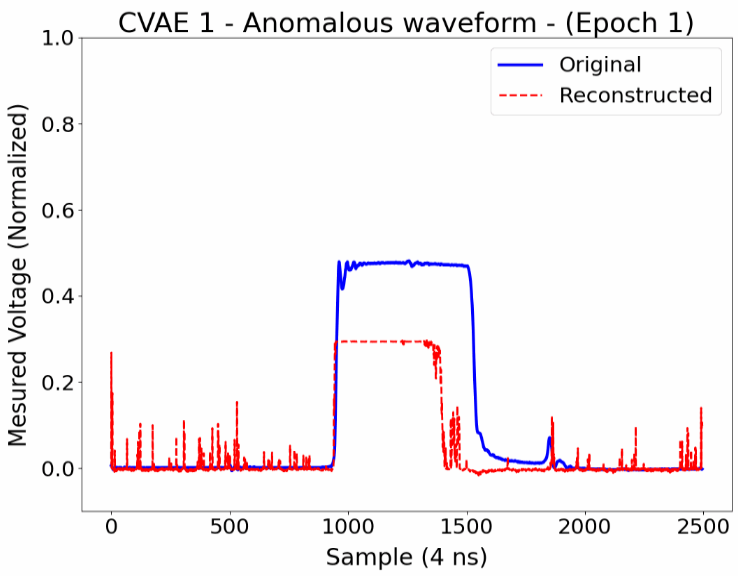

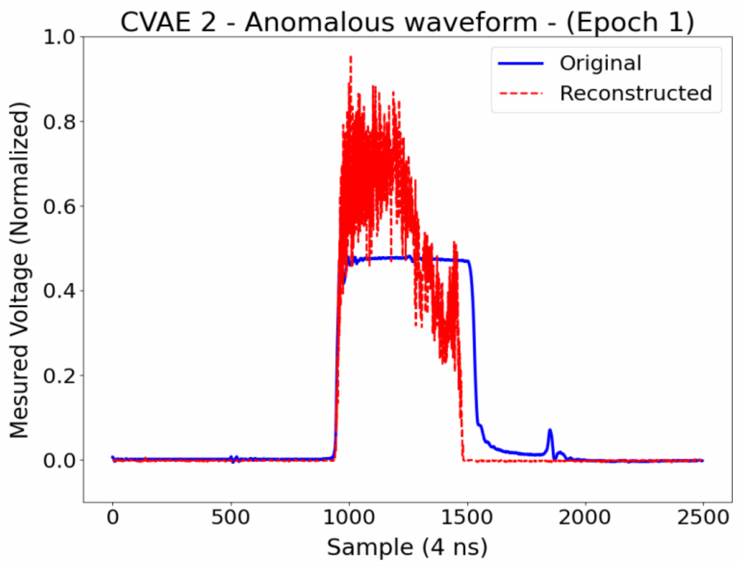

Results – Reconstruction errors

- **High Reconstruction Errors**: Observed in some data.

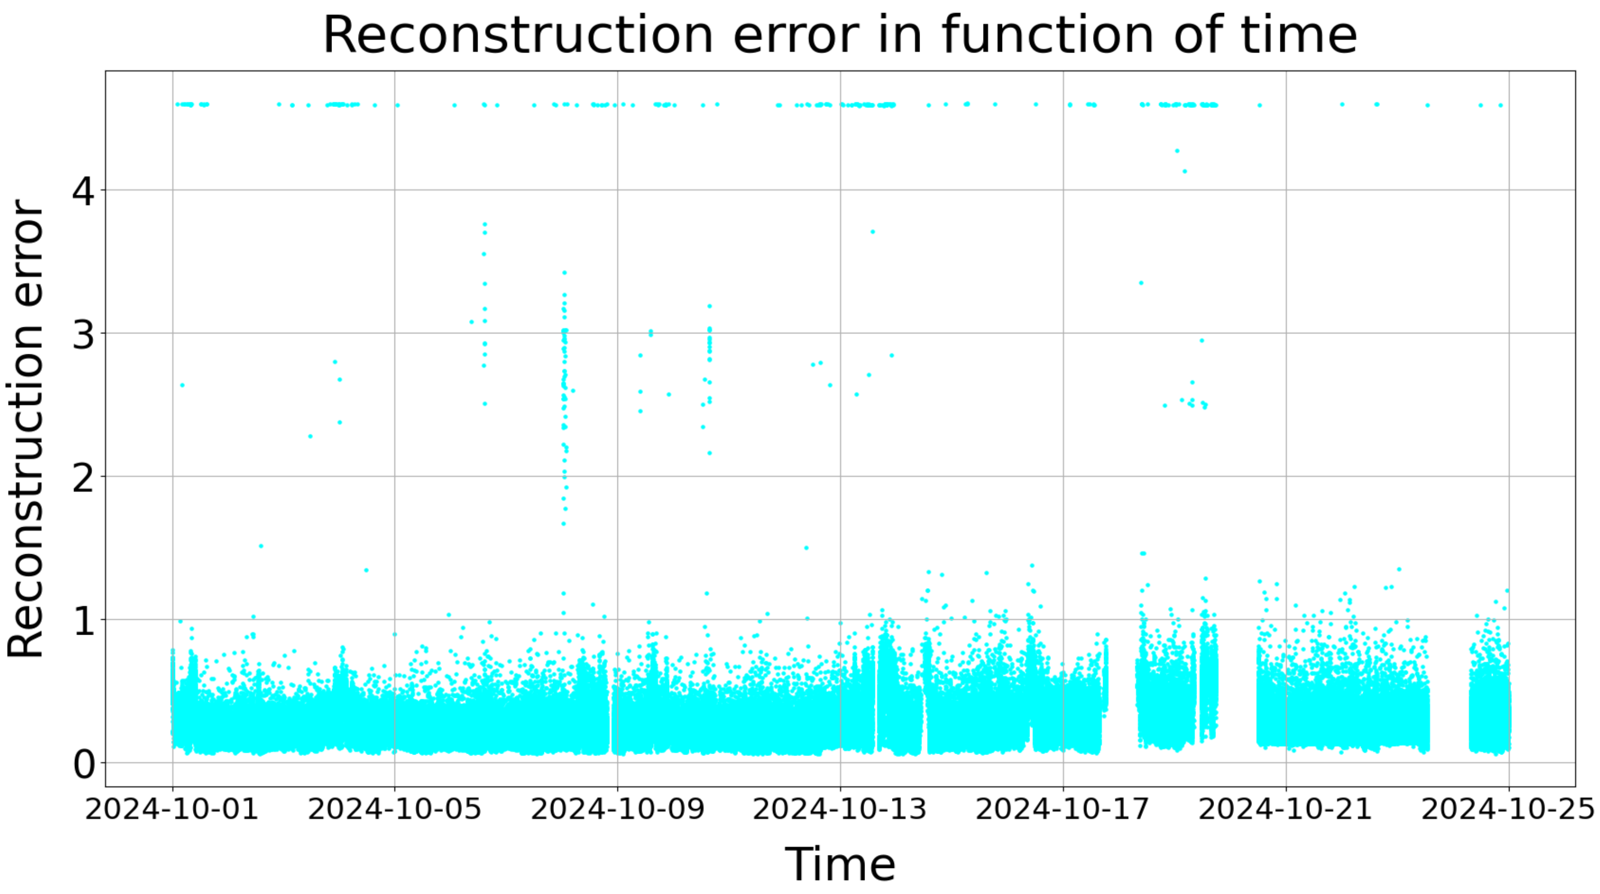

Results – Reconstruction errors

- **High Reconstruction Errors**: Observed in some data.
- **Known Anomalies**: Associated with high errors.

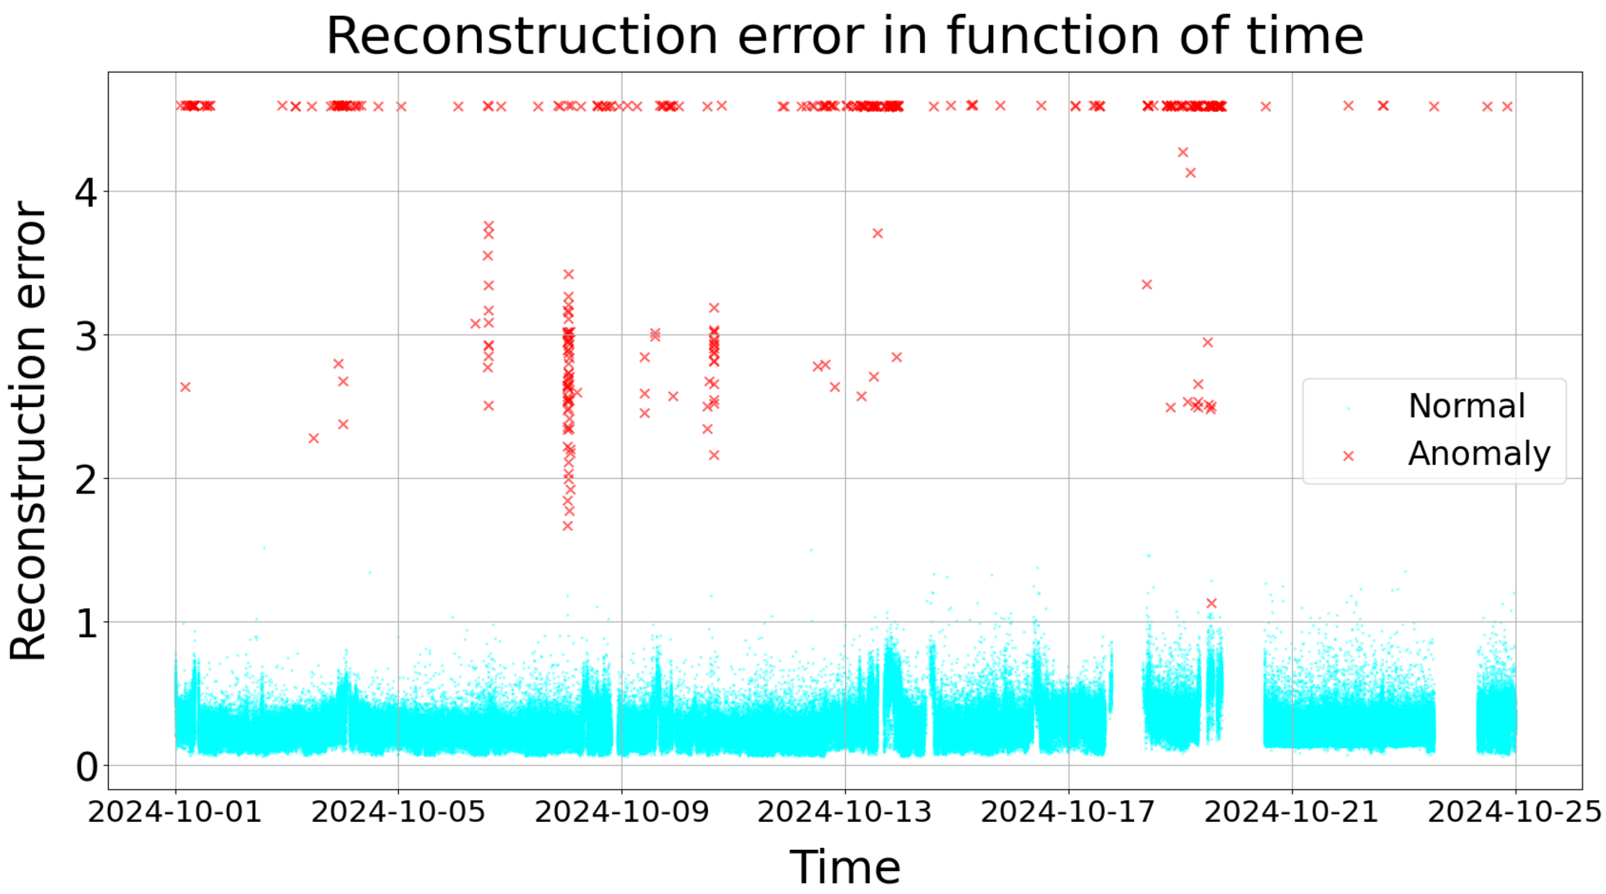

Results – Reconstruction errors

- **High Reconstruction Errors**: Observed in some data.
- **Known Anomalies**: Associated with high errors.

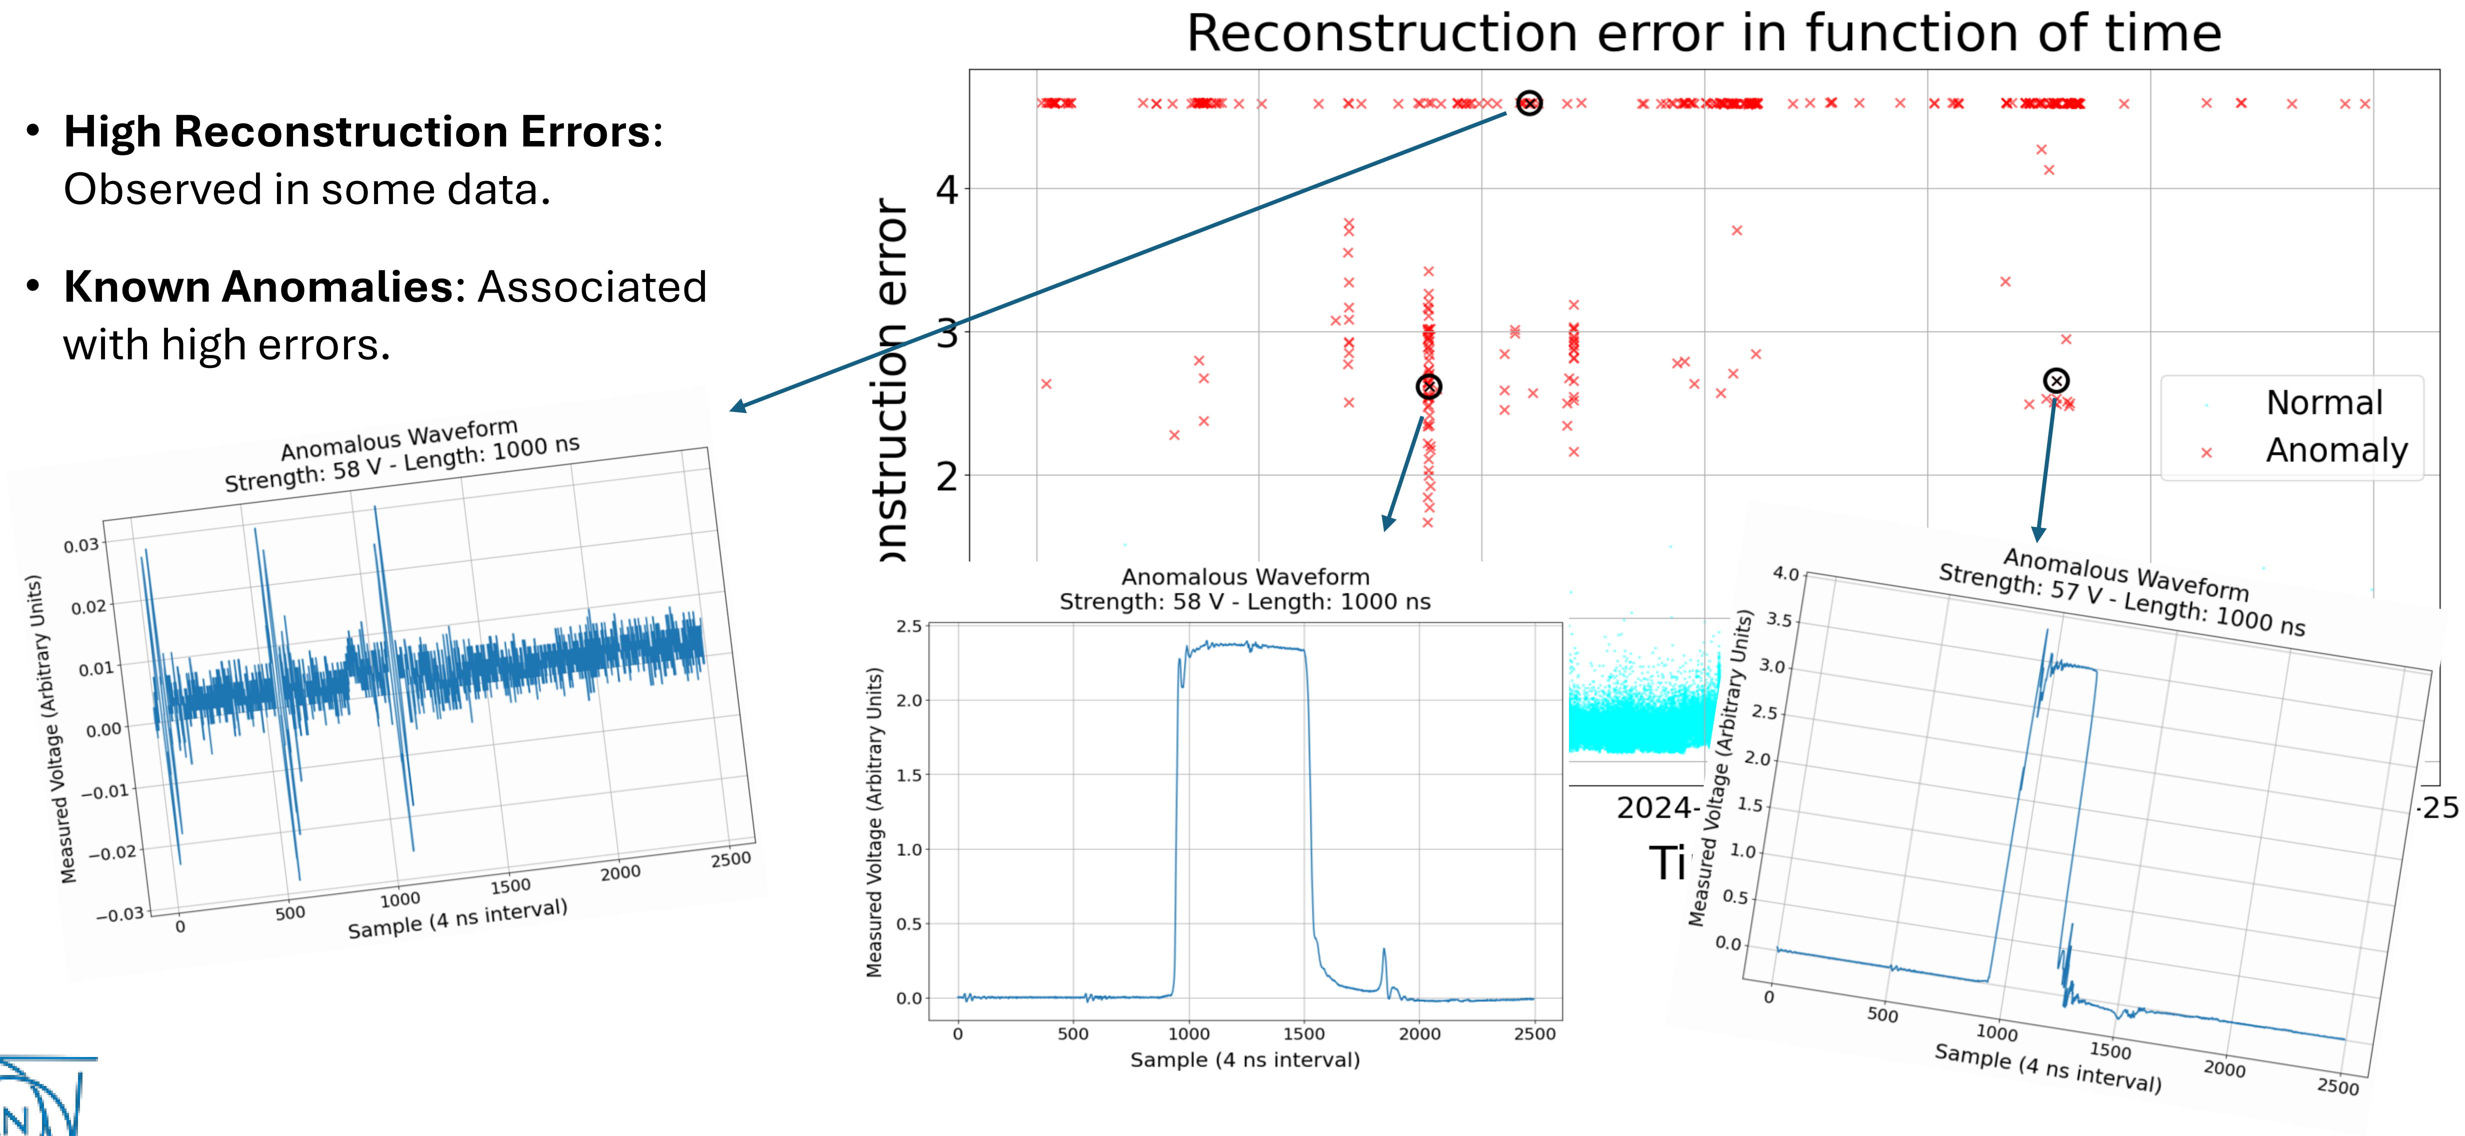

Results – Reconstruction errors

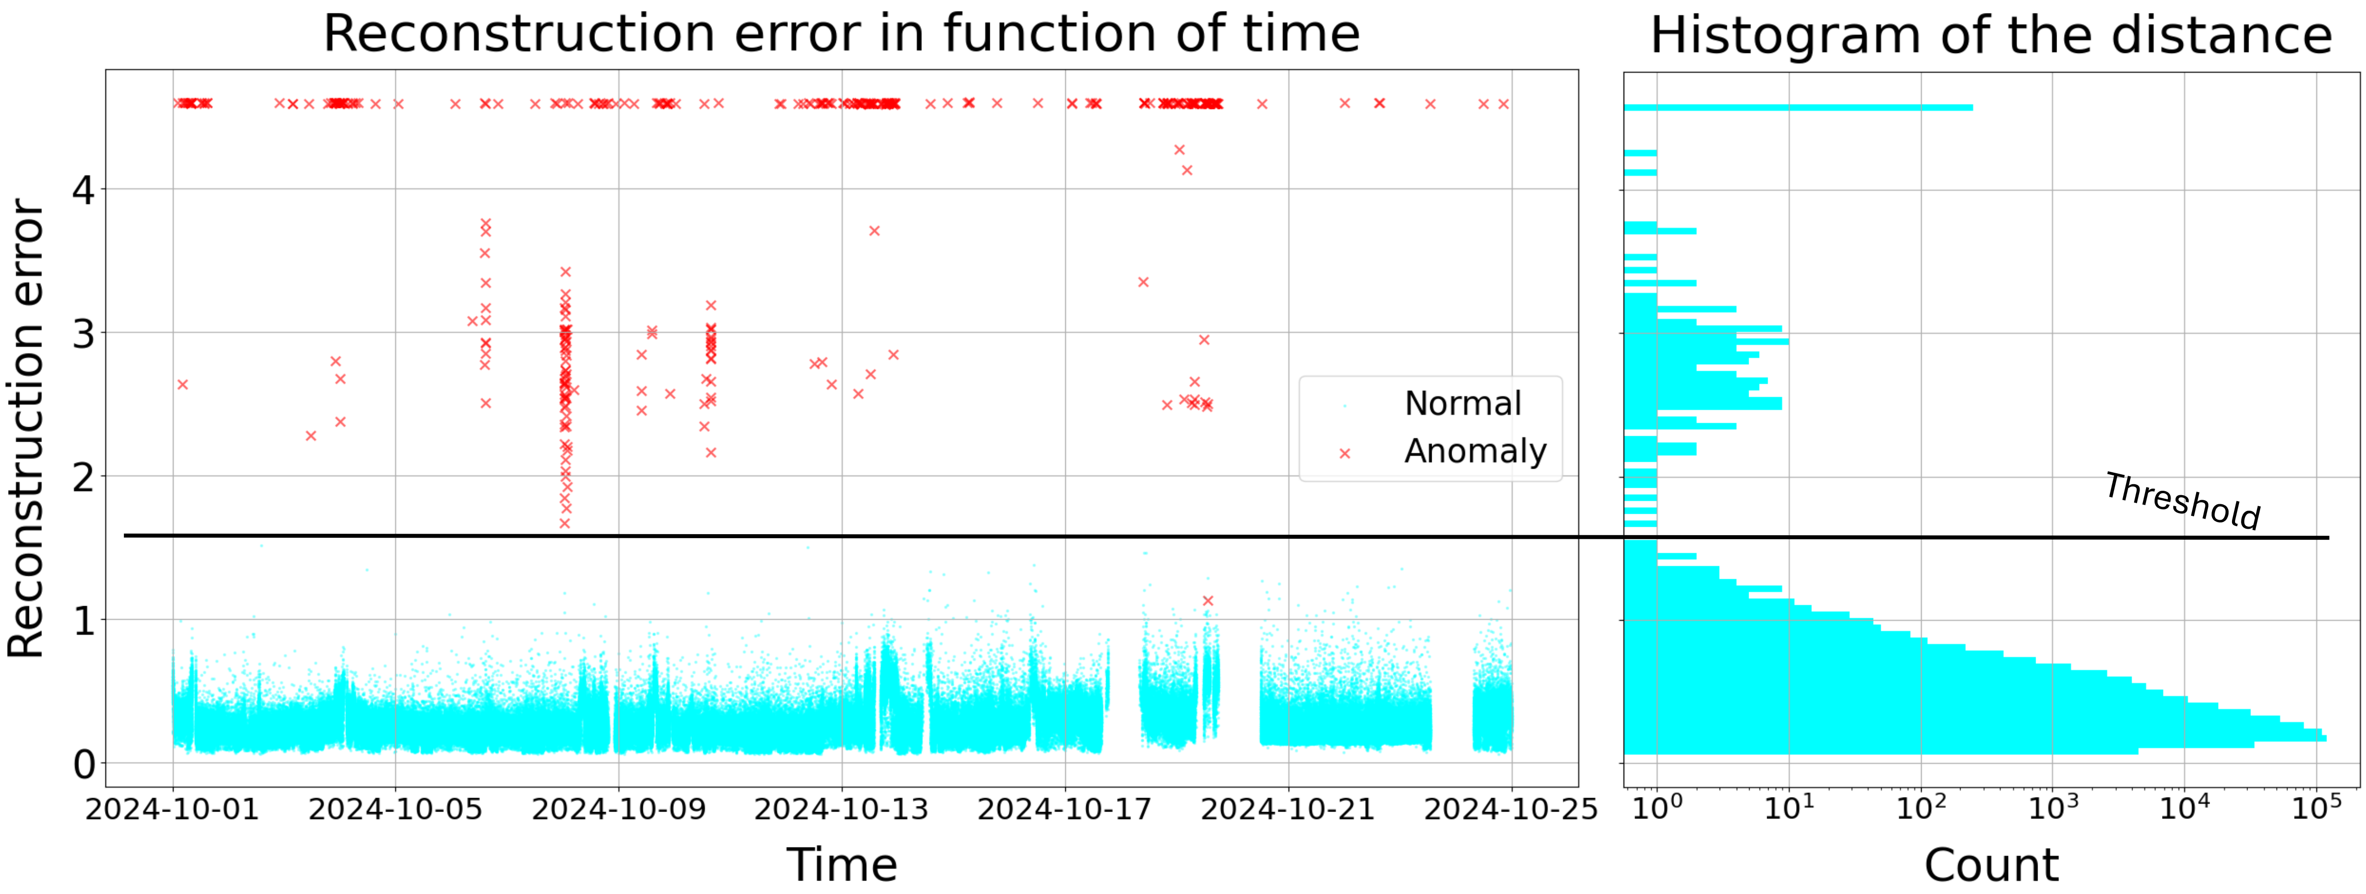

Results – Reconstruction errors

- **High Reconstruction Errors**: Observed in some data.
- **Known Anomalies**: Associated with high errors.
- **Visualization**: Logarithmic y-axis for clarity.

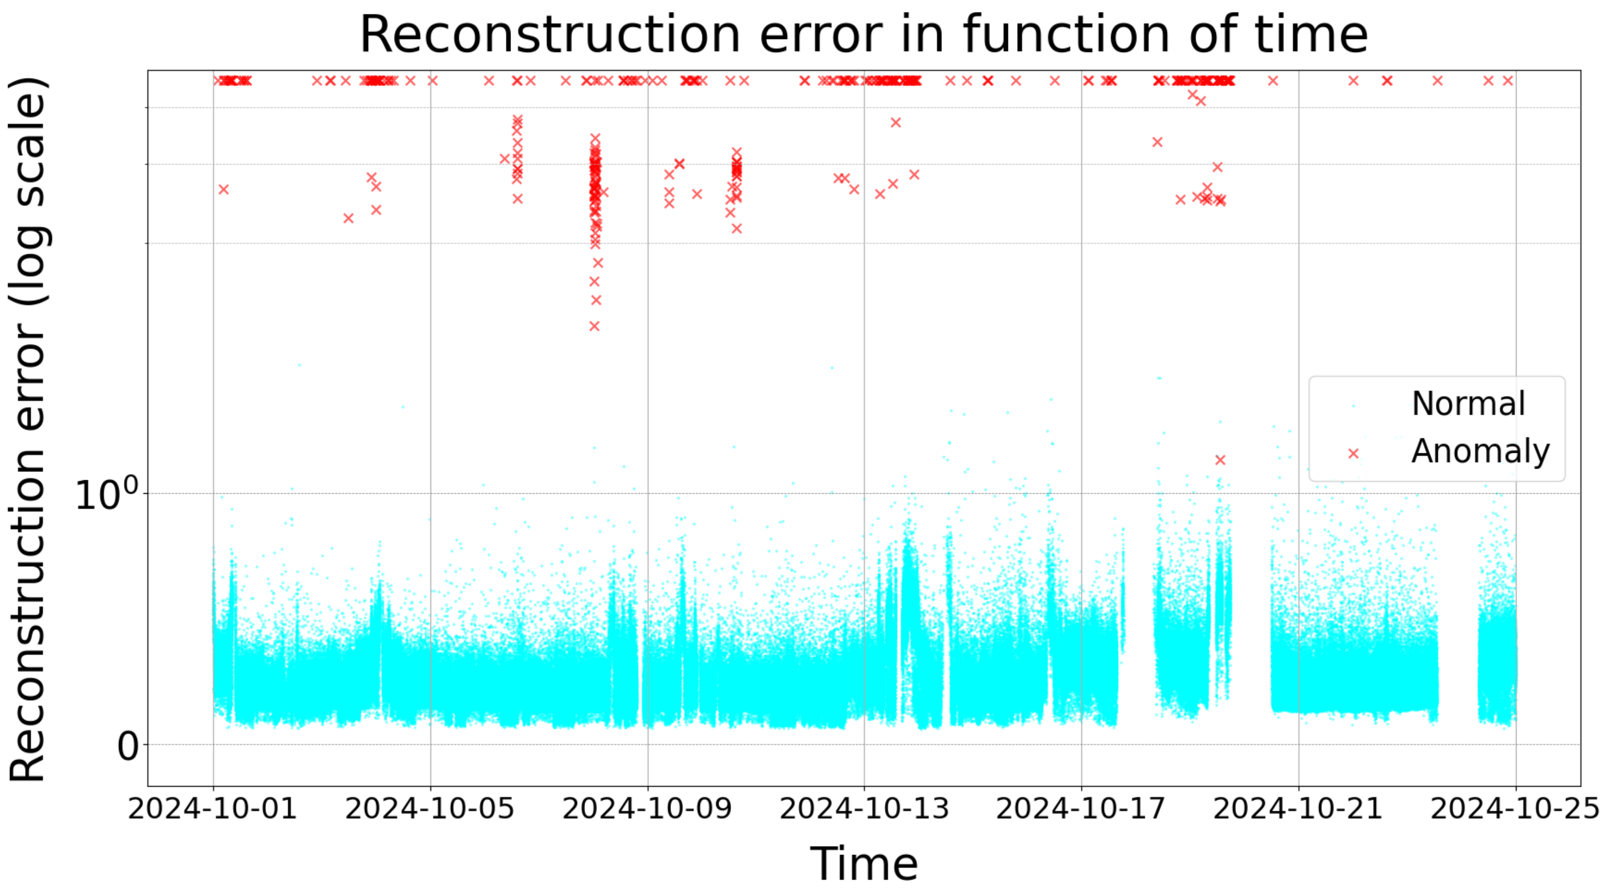

Results – Reconstruction errors

- **High Reconstruction Errors**: Observed in some data.
- **Known Anomalies**: Associated with high errors.
- **Visualization**: Logarithmic y-axis for clarity.
- **Rolling Median**: Provides insights into anomaly precursors.

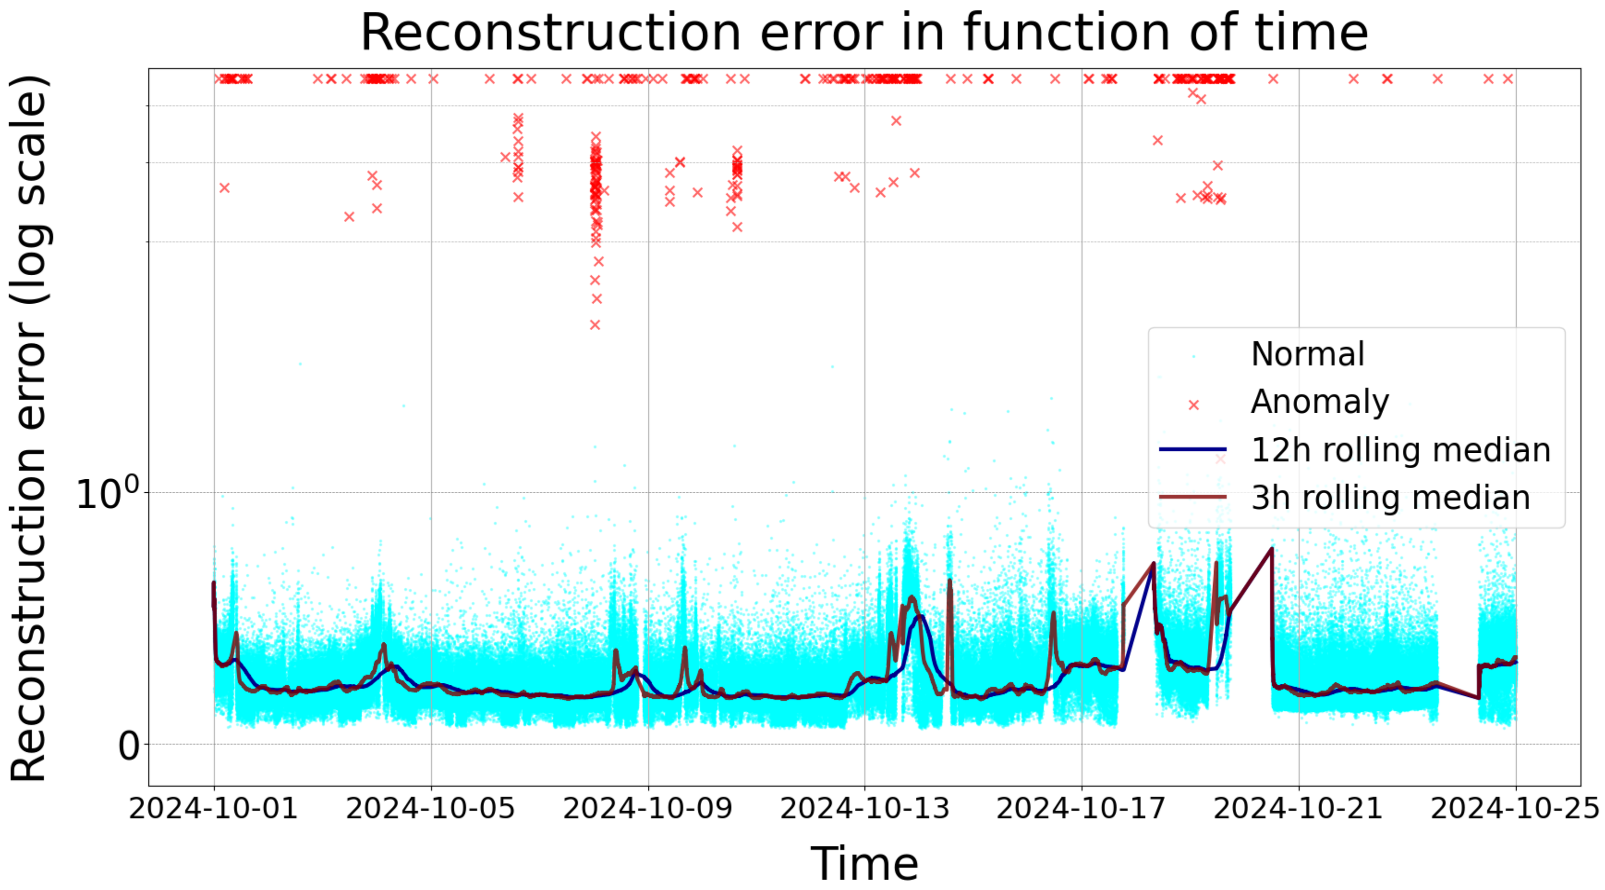

Results – Reconstruction errors

- **High Reconstruction Errors**: Observed in some data.
- **Known Anomalies**: Associated with high errors.
- **Visualization**: Logarithmic y-axis for clarity.
- **Rolling Median**: Provides insights into anomaly precursors.
- **Rolling Threshold:** Dynamically adjusts for trends and variability compute as:

  → ***Rolling Median + 5 * Rolling Std***

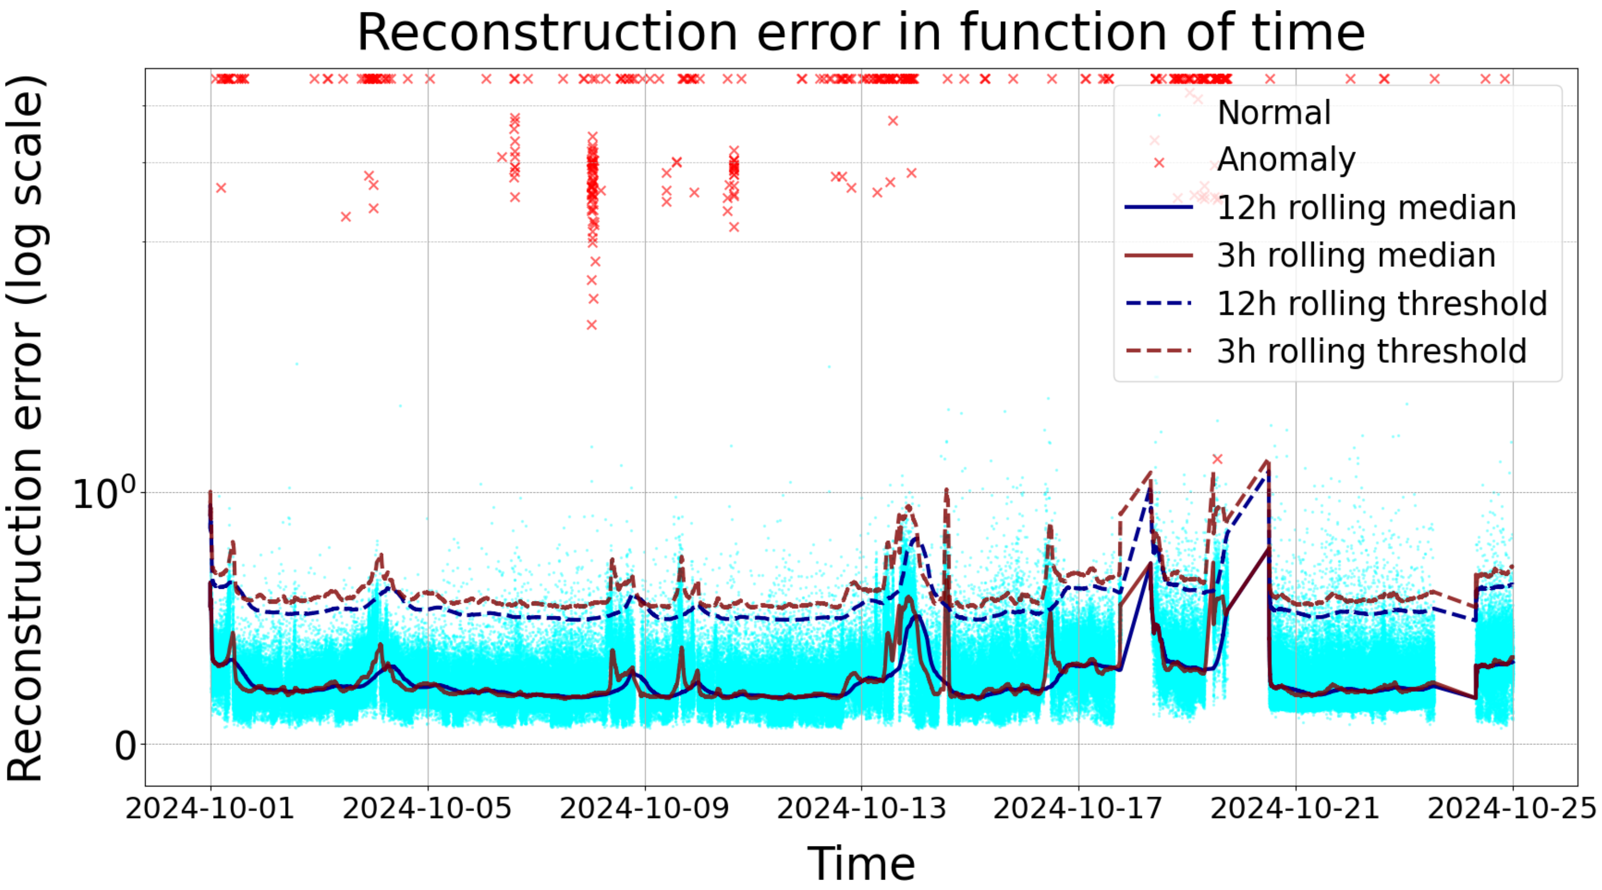

Preliminary Results – Anomaly Forecasting

Observation: Gradual rise in reconstruction error before failures.

Technique: Analyze error trends (e.g., rolling medians) for patterns.

Future Potential: Combine threshold-based and advanced models (e.g., LSTM).

Data Quality: Limited labels and root cause information.

Imbalance: Sparse anomalies, limited diversity in settings.

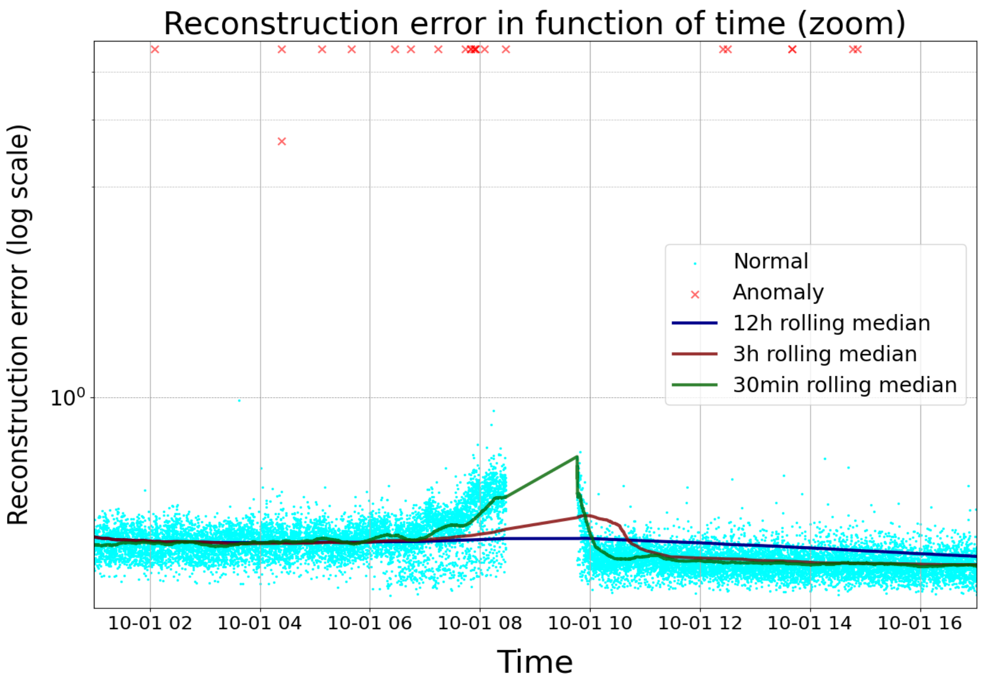

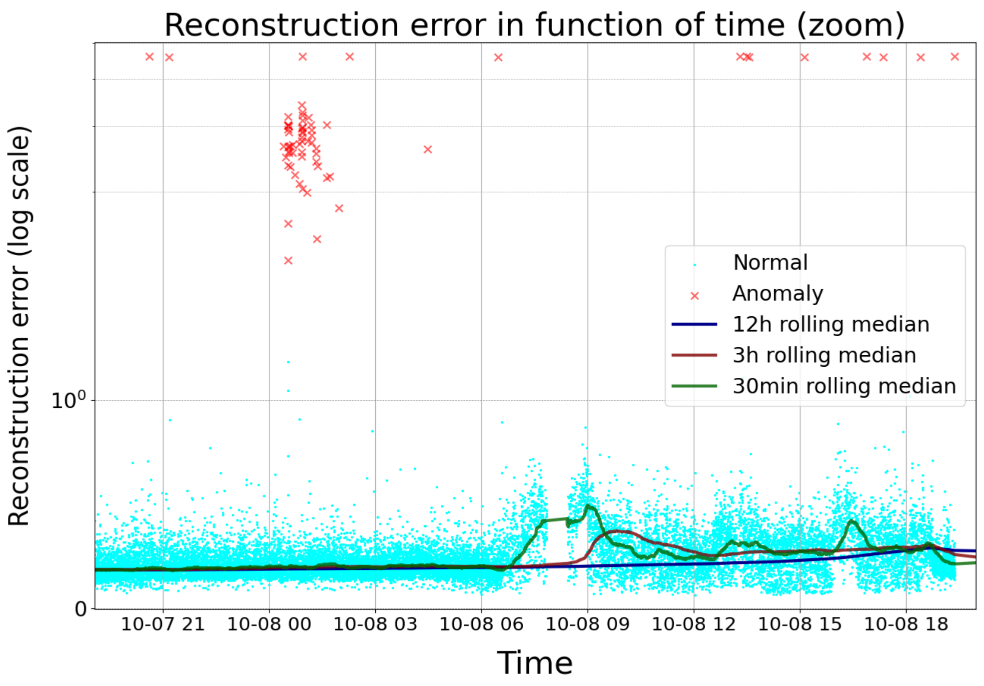

Preliminary Results – Anomaly Forecasting

Threshold-based methods:

- **Risk**: 30 min rolling mean of reconstruction error
  → *Mean gives more weight to anomalies compared to median for risk calculation.*
- **Threshold**: 12h Rolling median + Rolling median absolute value (MAD).
  → *Median and MAD is chosen for a more stable threshold.*
- **Condition**: If *Risk* > *Threshold* for 5 min continuously → **Warning**.
- **Visualization**: Vertical dash lines show the time warnings occur.

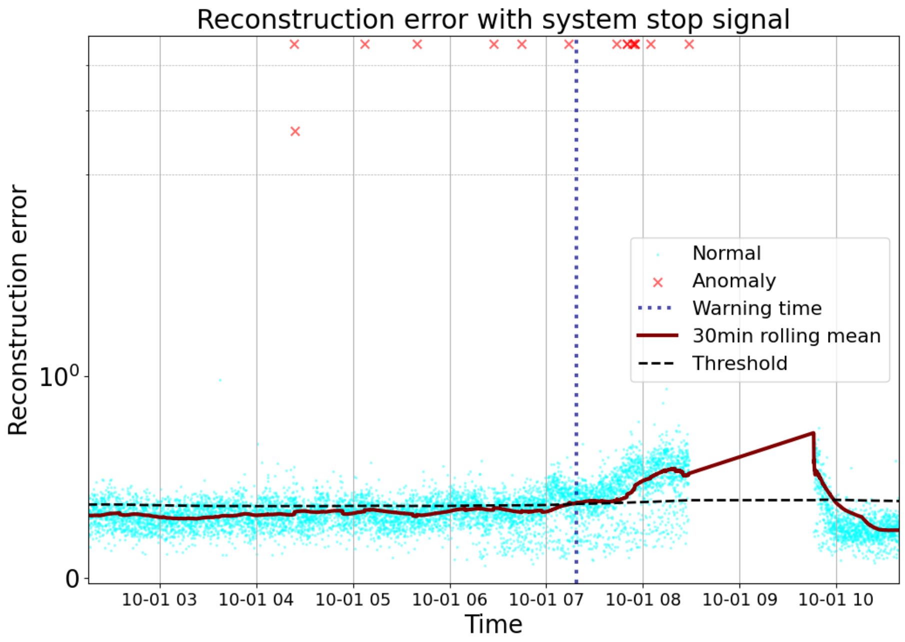

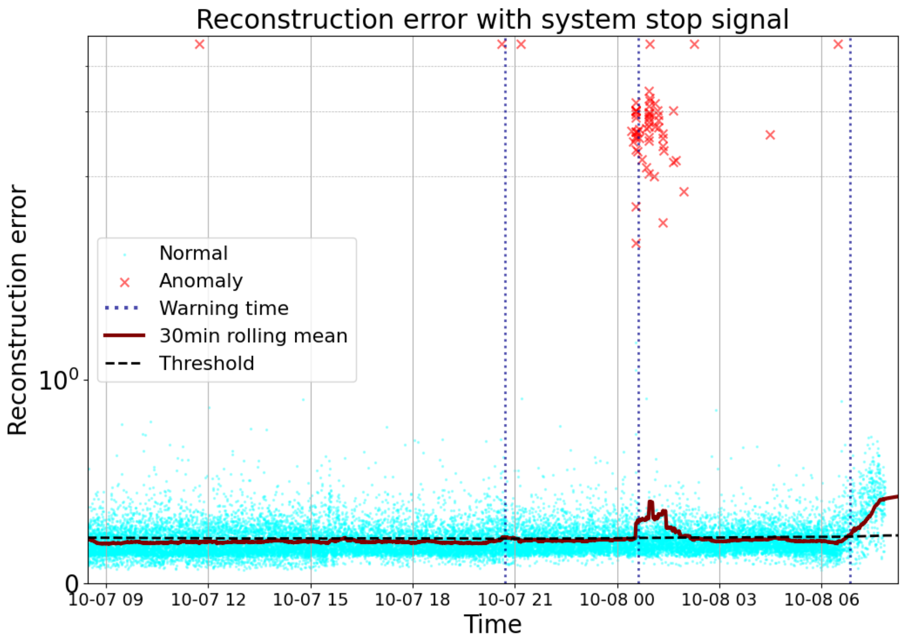

Preliminary Results – Anomaly Forecasting

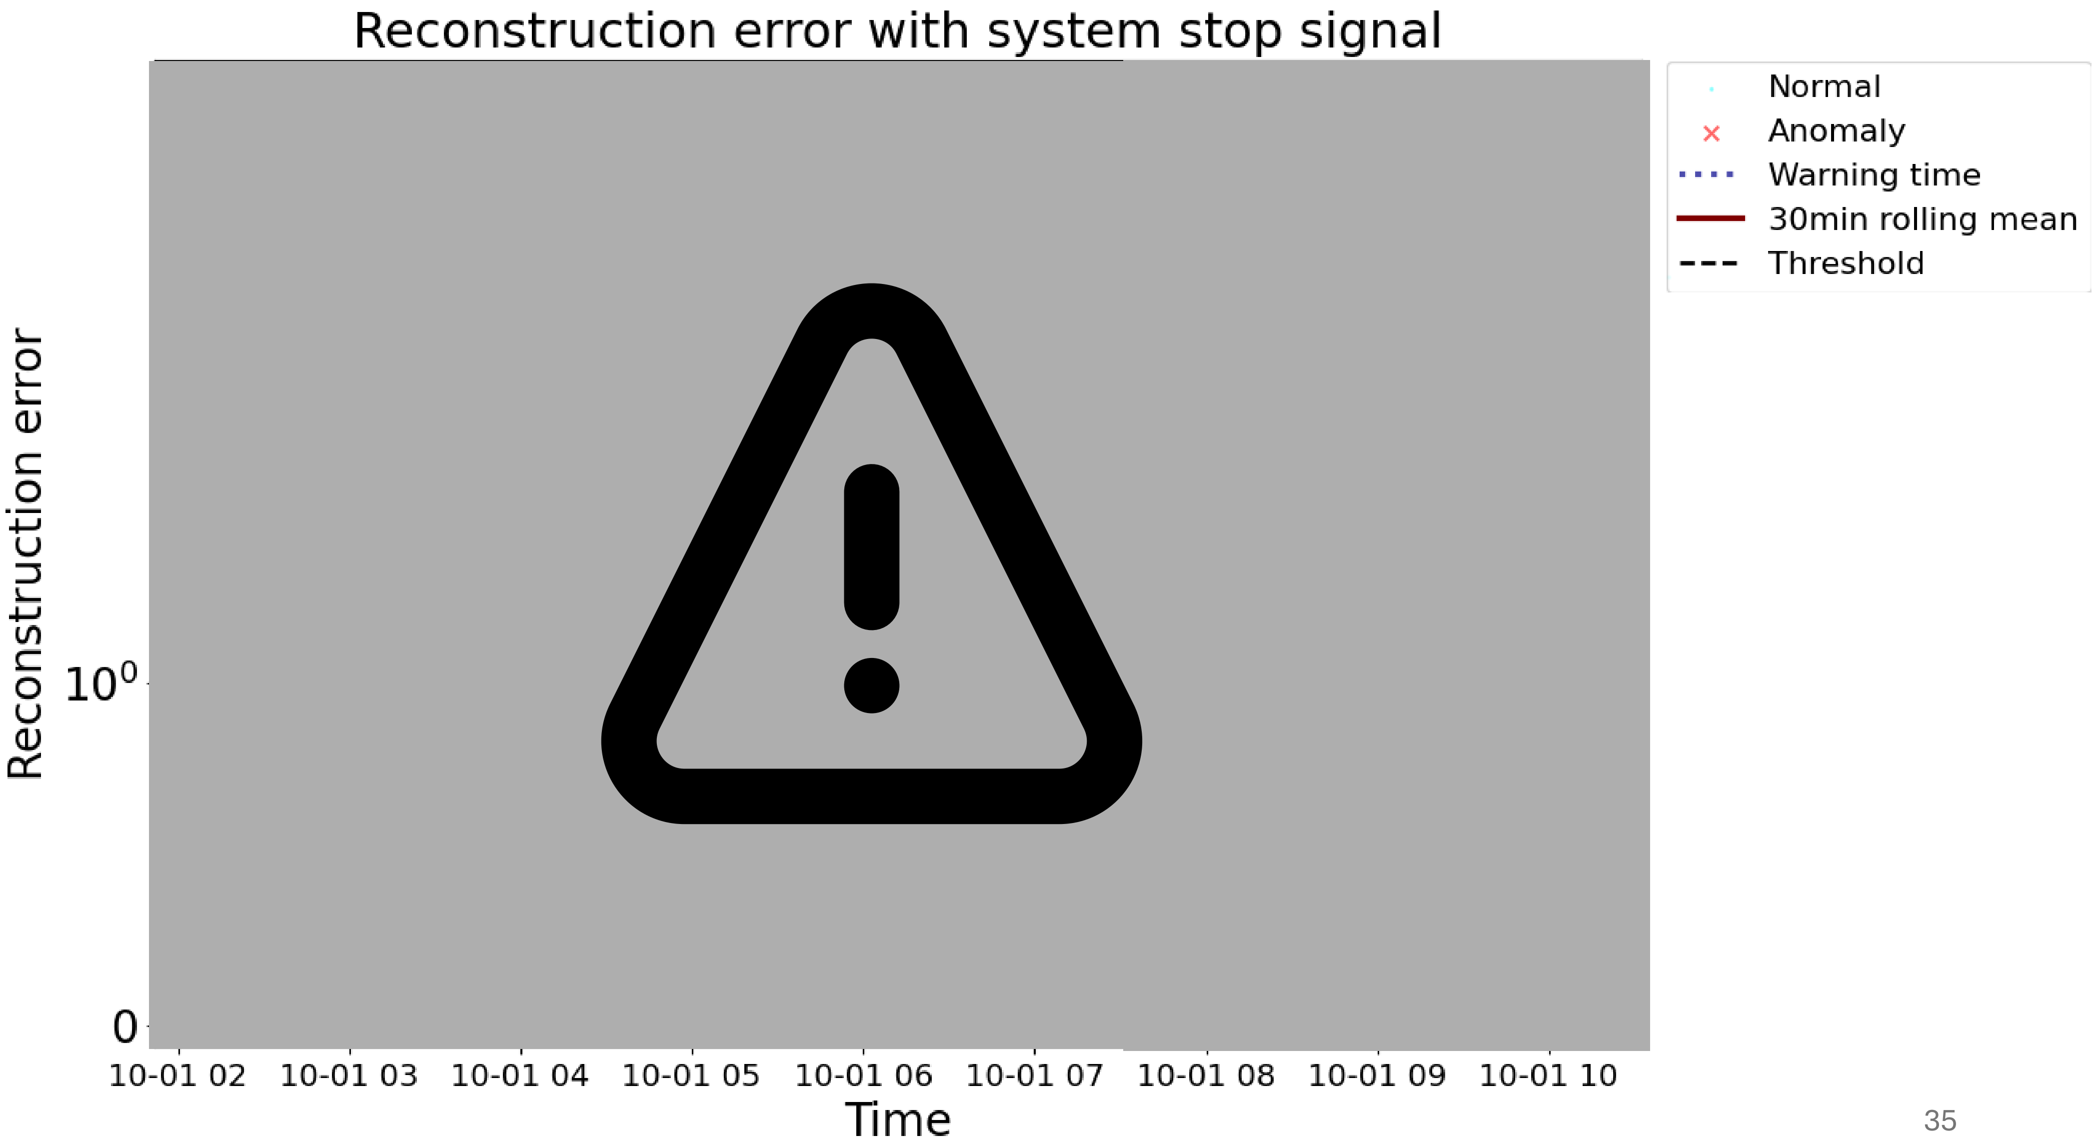

Operational Deployment for Model Validation:

- **UCAP Deployment**: Launch before machine restart.
- **Detection**: Flag anomalies and send warnings.
- **Forecasting**: Two warning levels:
  - **Rising Risk**: Signals moderate issues for monitoring.
  - **Critical Risk**: Alerts for significant risks, advising a machine stop.
- **Feedback**: System stops generate reports with reasons, anomalies, and system trends, logged in the logbook.
- **Improvement**: Refine detection using expert insights.

Continual Learning

Why Continual Learning in This Context:

- **Operational Shifts**: Upgrades, different beam types, new hardware components, varying high-voltage settings across years.
- Traditional static models may *forget* old knowledge or misinterpret new normal conditions as anomalies.

Key Requirements:

- **Adaptation**: Ongoing training with new waveforms while preserving knowledge of previously seen modes.
- **Stability vs. Plasticity**: Avoid *catastrophic forgetting* while still learning new patterns and normal states.
- **Drift Detection**: Distinguish *natural drift* in data (e.g., new standard operation) from true anomalies.

Implementation Approaches:

- **Rehearsal-based**: Keep a small representative buffer of past pulses to retrain or regularize the model.
- **Regularization-based**: Techniques like EWC (Elastic Weight Consolidation) or MAS to preserve crucial model parameters [10, 11].
- **Dynamic Architectures**: Expandable network components to handle truly novel operational regimes [12, 13].
- **Generative replay:** Generate past data with the model and combine them with new data for training [10].
- **And More to Explore**: ...

Conclusions & Outlook

Conclusions:

- CVAE models show potential for anomaly detection.
- Forecasting is promising but needs further validation.
- Challenges: Data imbalance and limited operational diversity.

Outlook:

- **Testing**: UCAP deployment this year with A/B testing of thresholds and forecasting methods.
- **Improvement**: Refine models based on feedback and validated thresholds.
- **Future**: Scale to other systems and enhance adaptability with continual learning.

References

1. EPA core team, *Efficient Particle Accelerators What is on the Horizon,* Joint Accelerator Performance Workshop 2023 , Montreux, 7 December 2023
2. F. Velotti, K. Papastergiou, Alex and V. Kain, *WP8 Automate equipments* (https://docs.google.com/presentation/d/130UolLf9o8UhCt4jOGE1xjmaZxSLKH2Tx90BCM7fwwc/edit#slide=id.p)
3. S. Gilardoni and D. Manglunki, *Fifty years of the CERN Proton Synchrotron - Volume I*, CERN-2011-004.
4. M.J. Barnes*, CAS: Beam Injection, Extraction & Transfer 2 Kicker Systems - Part 1 - Introduction and Hardware,* CERN TE/ABT
5. Luis Miguel Coralejo Feliciano, Matthew Fraser, Thomas Kramer, Christophe Boucly, *Functional Specification for PS Extraction Kickers PE.KFA71 and PE.KFA79,* CERN SY-ABT
6. Bercea et al., *What Do We Learn? Debunking the Myth of Unsupervised Outlier Detection* (2023).
7. Ran et al., *Detecting Out-of-Distribution Samples via Variational Auto-Encoder with Reliable Uncertainty Estimation*, (2021)
8. Bhosale et al., *Anomaly Detection using Diffusion-Based Methods,* (2024)
9. Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. *Optuna: A Next-generation Hyperparameter Optimization Framework. In KDD* (2019).
10. Faber et al., *Lifelong Learning for Anomaly Detection: New Challenges, Perspectives, and Insights* (2023).
11. Dohare et al., *Loss of Plasticity in Deep Continual Learning* (2024).
12. Ross et al., *Incremental Learning for Robust Visual Tracking* (2008).
13. Shenggang Li, *Dynamic Time Series Model Updating* (2024).
14. Adrian Alan Pol and Victor Berger and Gianluca Cerminara and Cécile Germain and Maurizio Pierini, *Anomaly Detection With Conditional Variational Autoencoders* (2020).