Beam-based RF Station Fault Identification at LCLS

R. Humble[1], F. O'Shea[2], W. Colocho[2], M. Gibbs[2],
H. Chaffee[2], E. Darve[1], D. Ratner[2]
[1] Institute for Computational and Mathematical Engineering, Stanford University
[2] SLAC National Accelerator Laboratory

SLAC NATIONAL ACCELERATOR LABORATORY

Stanford University

Outline

1. **Overview**: Anomalies at LCLS
2. **Deep Dive**: RF Station Fault Identification
3. **Extension Work**: CovAD

Anomalies at LCLS

LCLS

- Linac Coherent Light Source (LCLS) is a hard X-ray FEL
- Aim is to deliver an X-ray laser to users around the clock
- User experiments demand stability
- Produces over 200,000 data streams

Anomalies at LCLS

Anomalies vary in degree of impact

1. **Downtime** (beam goes offline)
   - ~3% of availability lost to unplanned downtime caused by faults
   - >180 hours
   - 3+ full user experiments equivalent
2. **FEL performance degradation**
   - Unstable beam causes problems for users
   - LCLS anomalies → noise/anomalies in user data

To guard against anomalies, operators

- Effectively limit the operational range of accelerator
- Hold certain components (e.g., spare RF stations) in reserve

There are lots of LCLS failure modes!

Deep Dive: RF Station Fault Identification

RF Station Faults

- Accelerator is powered by RF stations
  - 82 stations
  - Split across 4 regions (L0-L3)

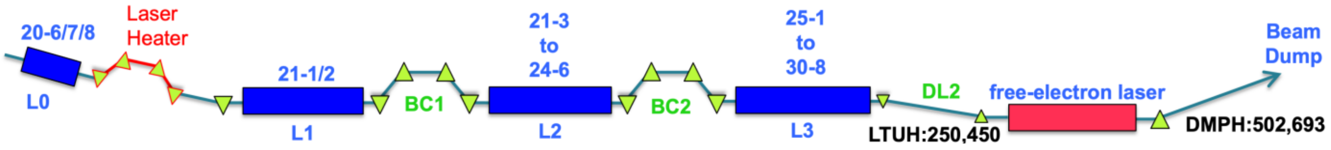

- RF station anomalies are a known, high-priority failure mode
- **Research goals:**
  1. Identify anomalies due to RF stations
  2. Identify the origin of the fault (e.g., which RF station failed and when)

Existing System

- Current detection process is largely manual for operators

1. Watch for drop in X-ray power

2. After drop, search for RF stations with warning/fail status

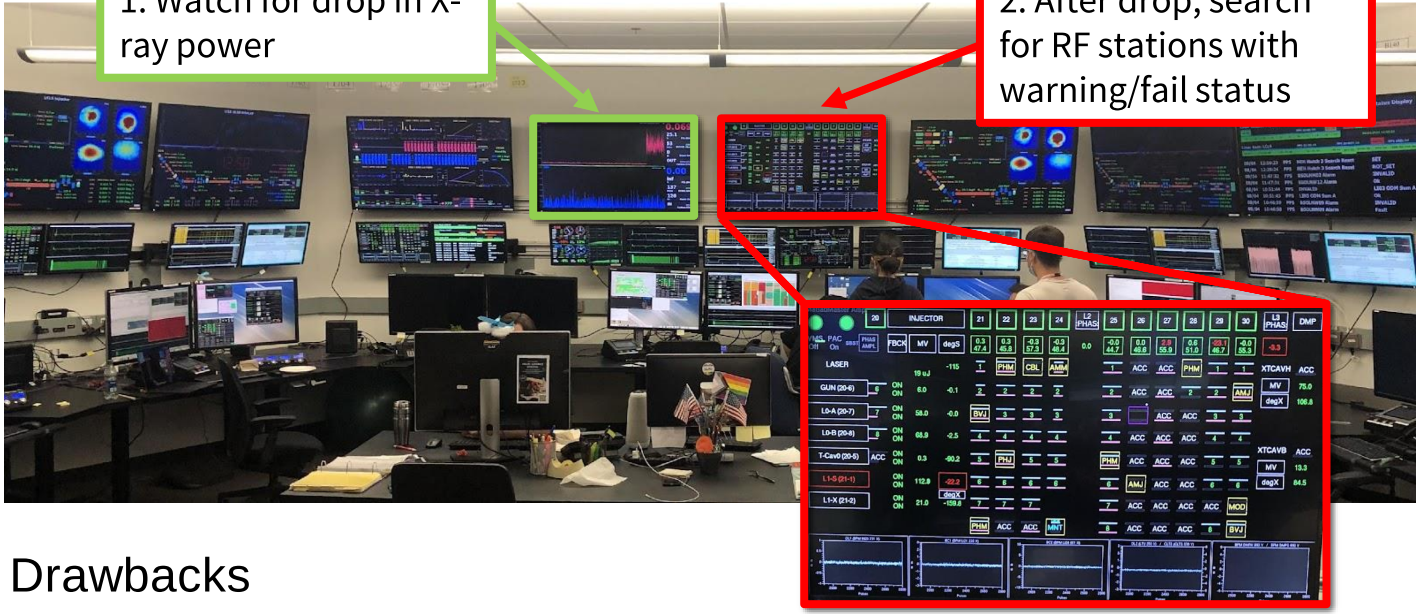

- Drawbacks
  1. Requires human attention for a routine activity
  2. Only addresses significant/prolonged anomalies

RF Station Diagnostics

- Several diagnostics per RF station already exist
- Shortcomings:
  1. Limited bandwidth
  2. Diagnostics are often incorrect

- Diagnostics → fault attribution

Beam Data

- Beam position and charge measurements
- Shortcomings:
  1. No direct attribution info
  2. Sensitive to other failure modes

- Beam data → fault verification

Need both signals to corroborate the anomaly

Example

RF Station Health

Beam Energy

Beam Charge

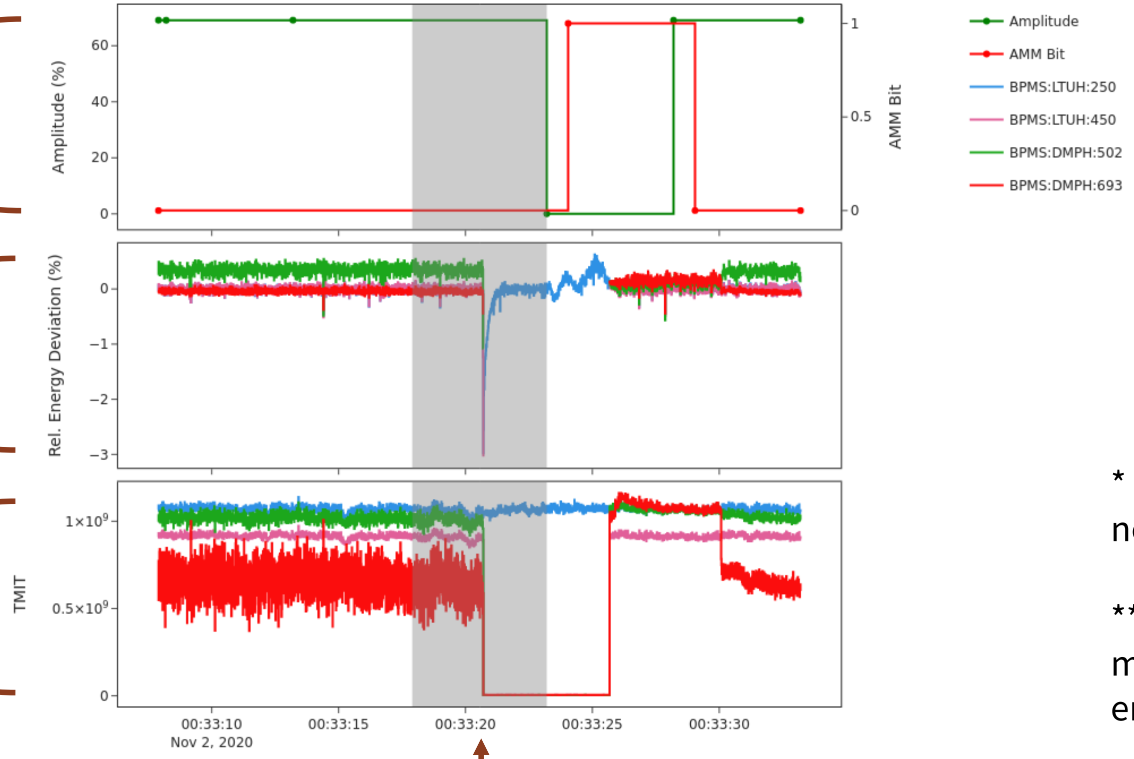

Anomaly

* RF Station health is not synchronous.

** Gray region is maximum timing error.

Beam-based RF Station Fault Identification

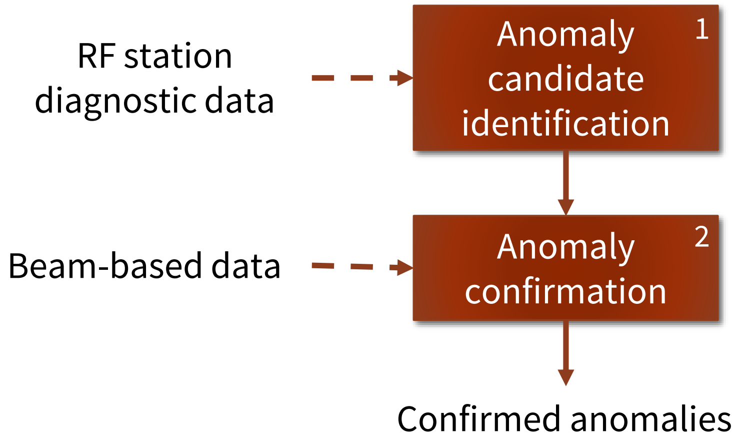

- Fully automated method

1. **Candidate identification**
   - Uses diagnostic data
   - Finds candidate anomalies
   - Custom built rules for each diagnostic
2. **Candidate confirmation**
   - Uses beam data from dispersive regions (i.e., sensitive to energy changes)
   - Unsupervised time-series method
     - Based on robust z-scores
     - Aggregates scores across BPMs to boost signal-to-noise ratio

Results

- Study period: ~2 months in Winter 2020-2021
- Evaluation: Hand-labeled every anomaly candidate

Variant	Precision	Anomalies Found
Only AMM[1] diagnostic	0.31	385
AMM[1] + Beam data	0.91	368
AMPL[2] + Beam data	0.88	504

- Takeaways
  - Reduce false positive rate by 20x by checking AMM with beam data
  - Only reduce AMM-based anomalies found by 4% with beam data
  - Switch to AMPL diagnostic increases anomalies found by 44%

[1] AMM: Amplitude mean out of tolerance (binary-valued) [2] AMPL: Amplitude (real-valued)

RF Station Attribution

- Study period: Nov. 2020 to Feb. 2022 [1]

Variant	Anomalies Found
AMM + Beam data	1666
AMPL + Beam data	3968

- Equivalent to 2-5 RF station anomalies per day
- Top-5 RF stations account for ~36% of identified anomalies

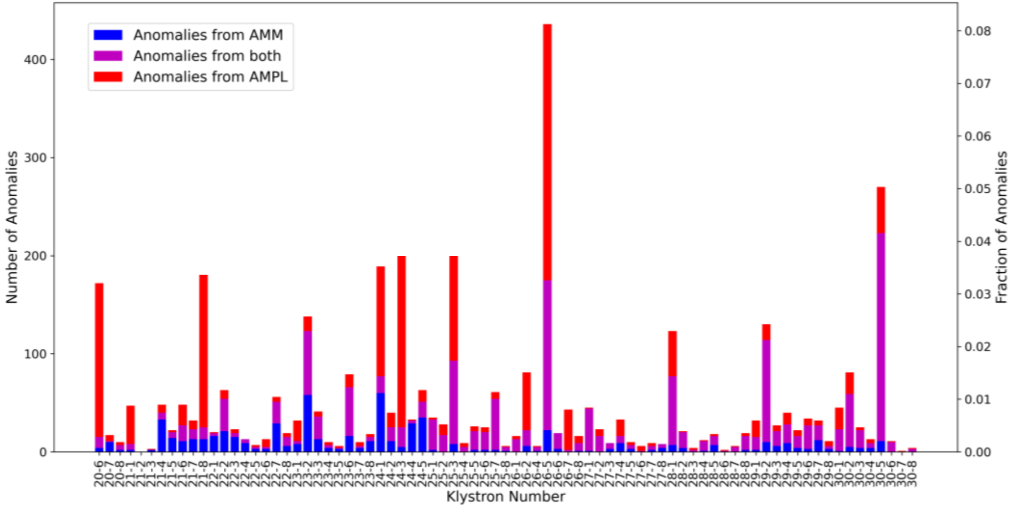

[1] Only contains 80 days worth of runtime due to various filtering criteria

Extension Work: CovAD

CovAD: Covariance Learning for Unsupervised AD

- Our RF station method was fundamentally <u>solving a labeling problem</u>
  - Human labels don't widely exist and are time-consuming
  - Raw diagnostics alone led to, at best, noisy labels
  - Beam measurements were used to corroborate the bad labels
  - Needed a lot of LCLS knowledge and custom anomaly methods

- We can rephrase our problem in a more general framework

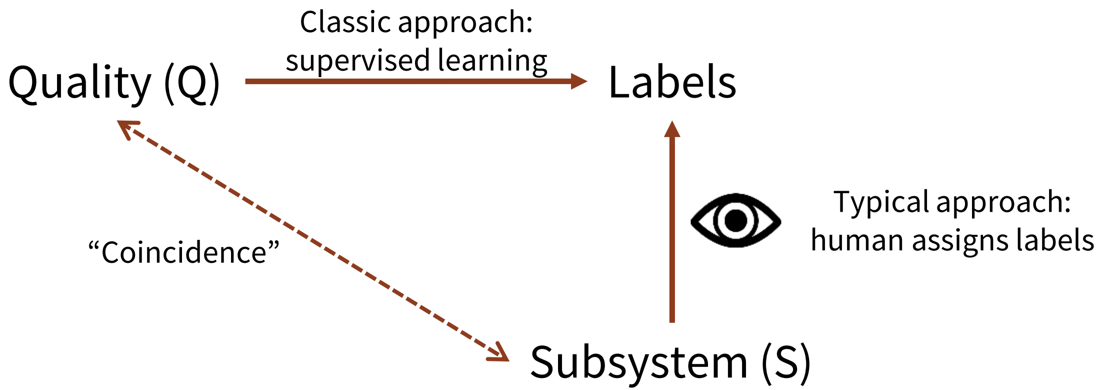

CovAD

- Assume we have candidate algorithms $A(\theta_s)$ and $A(\theta_q)$ to identify anomalies in each stream
- Only joint events are deemed anomalies
- Better coincidence rate → higher precision
- More coincidences → higher recall

- **Does NOT require labels!**

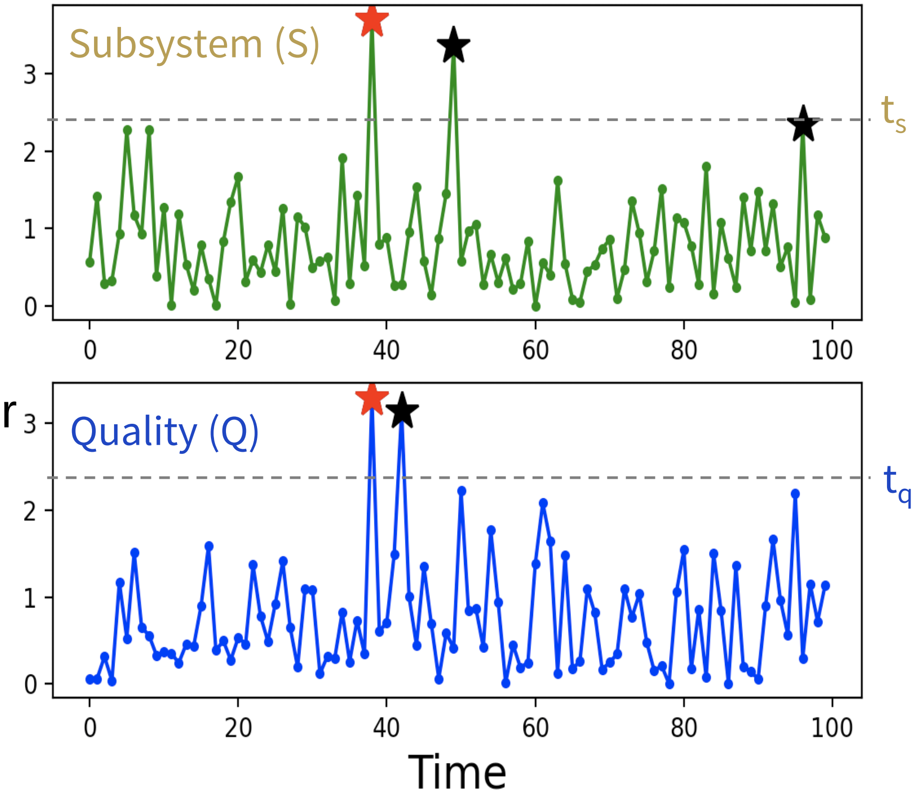

CovAD

- Suppose null hypothesis that s and q are independent and anomalies are rare
- We can then define an *unsupervised* analogue to recall
  - $J(\theta_s, \theta_q)$: number of joint events
  - $N(\theta_s, \theta_q)$: number of joint events under a null hypothesis
  - $R(\theta_s, \theta_q)$: estimate of true number of anomalies detected (analogous to recall)
- Which gives

$$R(\theta_s, \theta_q) = cov\left(A_{\theta_s}, A_{\theta_q}\right) = \sum_i A_{\theta_s}(s_i)\, A_{\theta_q}(q_i) - \left(\frac{1}{n}\sum_i A_\theta(s_i)\right)\left(\sum_i A_{\theta_q}(q_i)\right)$$

CovAD

- **Method**: Maximize the covariance $cov\left(A_{\theta_s}, A_{\theta_q}\right)$
- Upon convergence, both algorithms $A(\theta_s)$ and $A\left(\theta_q\right)$ identify features that predict anomalies
  - This is *despite* being fed different inputs
- Equivalent to a clustering method that uses covariance to identify clusters

CovAD: Running on RF Station Faults

- We ran CovAD over the raw RF diagnostics and beam data
- Used two neural networks as the scoring algorithms
- Result: CovAD only labels coincident events as truly anomalous

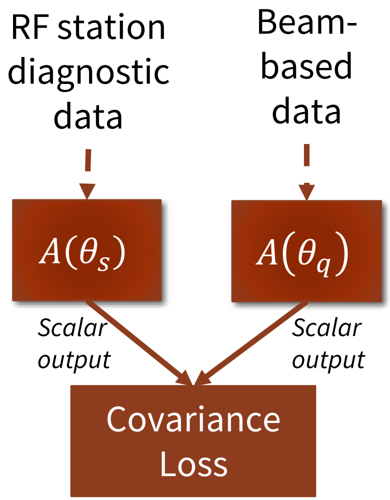

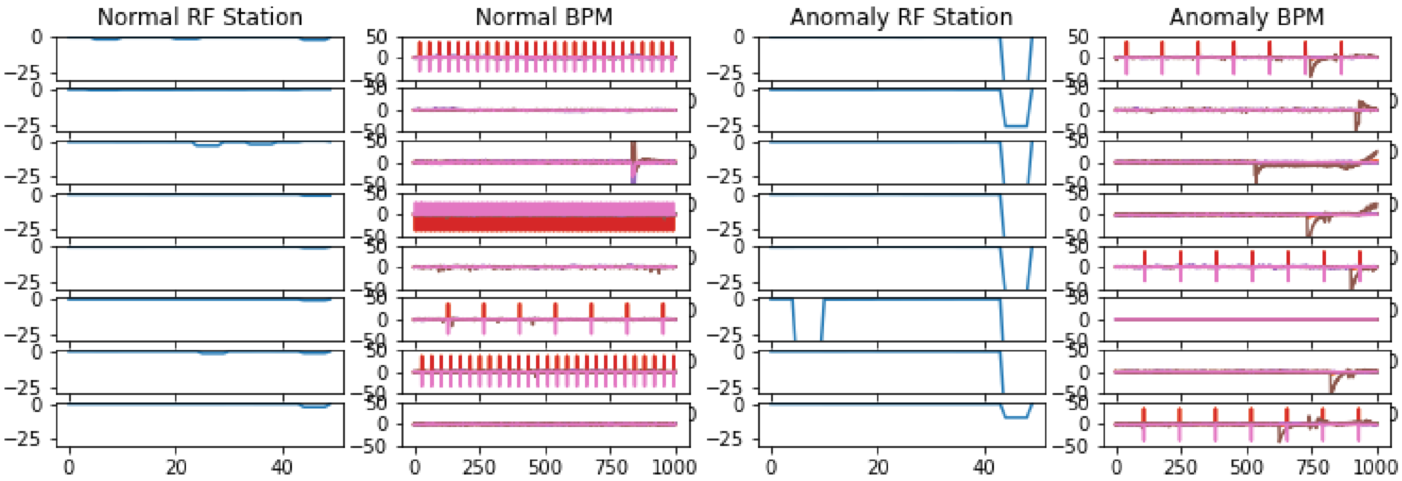

Thank you!

RF STATION FAULT IDENTIFICATION: HTTPS://ARXIV.ORG/ABS/2206.04626

Appendix

Limiting to “good” beam conditions

- We limit our method to time periods when machine is running at “good” conditions
  1. Beam stopper is inactive (STPR:BSYH:2:STD2_IN_A == 0)
  2. Beam rate is 120Hz (IOC:BSY0:MP01:PC_RATE == 8)
  3. Beam split is 120 Hz HXR/0 Hz SXR (IOC:IN20:EV01:RG02_ACTRATE == 10)
  4. Beam is actively logged (BPMS:IN20:221:TMITCUH1H logged at 1Hz)
  5. Beam has “good” charge (BPMS:IN20:221:TMITCUH1H > 0.5e9)
  6. Beam satisfies the above for at least 5 minutes

- Allow temporary (1 minute) violation of conditions 2-5 (machine will automatically respond to “catastrophic” errors; without this condition, we filter out many of the worst faults)

Types of Anomalies

- During hand labeling, we found a natural anomaly categorization:
  1. *Faults*: beam is lost
  2. *Sustained anomalies*: initial deviation followed by a recovery period

- *Sustained anomalies* accounted for ~60% of all anomalies during the study period
- Most *sustained anomalies* were likely missed by operators in real-time
  - Short-lived: typically last <10 seconds
  - FEL degradation: causes beam degradation, not cause beam loss

Fault

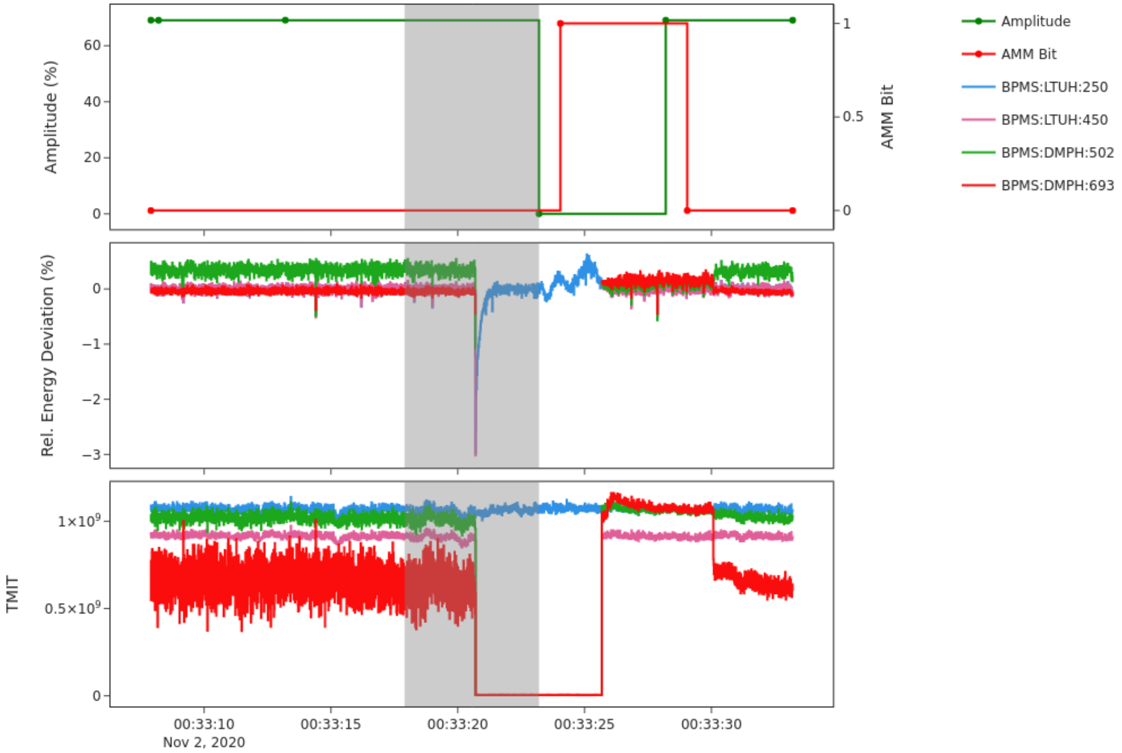

Amplitude (%)
60
40
20
0
1
0.5
0
AMM Bit
Amplitude
AMM Bit
BPMS:LTUH:250
BPMS:LTUH:450
BPMS:DMPH:502
BPMS:DMPH:693
Rel. Energy Deviation (%)
0
−1
−2
−3
TMIT
1×10^9
0.5×10^9
0
00:33:10
Nov 2, 2020
00:33:15
00:33:20
00:33:25
00:33:30

Sustained Anomaly

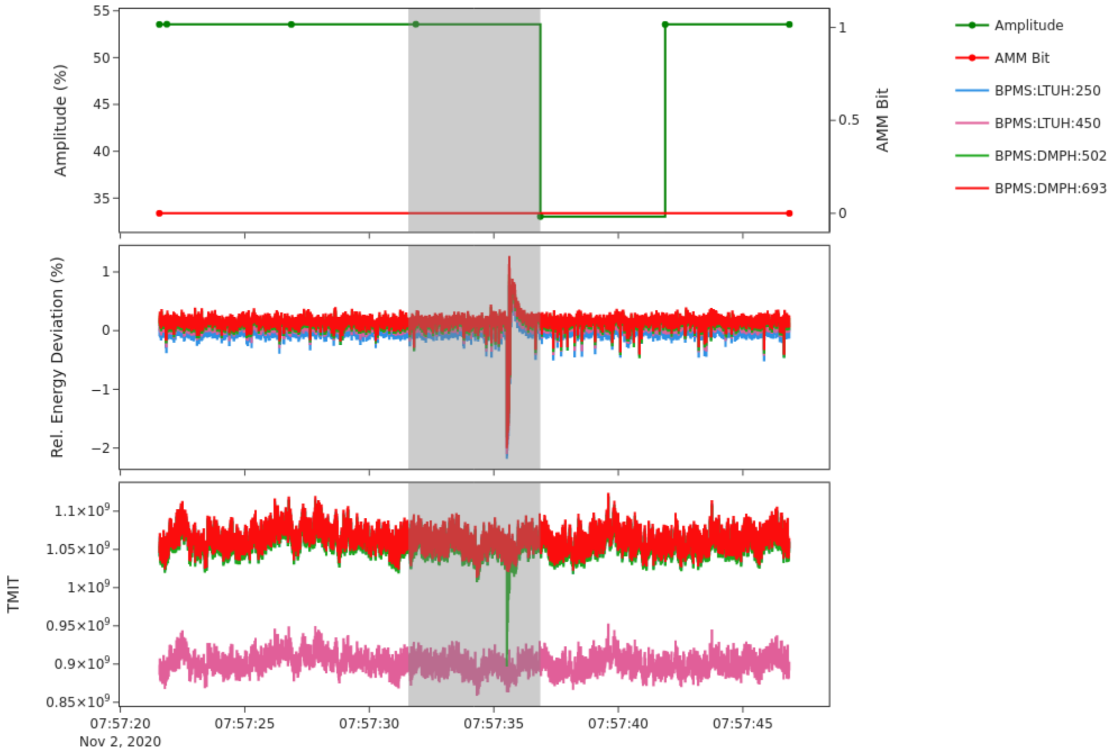

ResilientVAE: Unsupervised AD with contaminated data

- We have focused on finding anomalies in a particular subsystem(s)
- But LCLS has a lot of failure modes!
- Can we find general anomalies in the LCLS beam data?
- Can we also attribute the anomaly back to particular beam readings?

- The challenge is that the beam data is
  - **Unlabeled** → no supervised methods
  - **High (>500) dimensional** → curse of dimensionality for many traditional (distance-based/vector space partitioning) anomaly detection algorithms
  - **Contaminated** → no “normal” training set exists for one-class classifiers

ResilientVAE

- We propose a variant of a Variational AutoEncoder (VAE)
  - v: cleanliness of full input
  - z: latent code
  - w_d: cleanliness of d^th input feature
  - x_d: d^th input feature

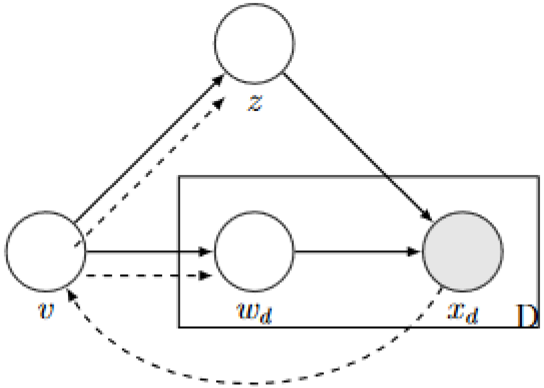

- Neural network is composed of:
  - $p_\theta(x|z)$: encoder
  - $q_\phi(z|x)$: decoder
  - $\pi_d(x)$: feature cleanliness
  - $\gamma(x)$: input cleanliness

$$\mathcal{L}(x) = \sum \gamma(x)\pi_d(x)\mathbb{E}_{q_\phi(z|x)}\left[\log p_\theta(x_d|z)\right] - \lambda_1\gamma(x)D_{KL}\left(q_\phi(z|x)||p(z)\right)$$

ResilientVAE: Example Anomaly

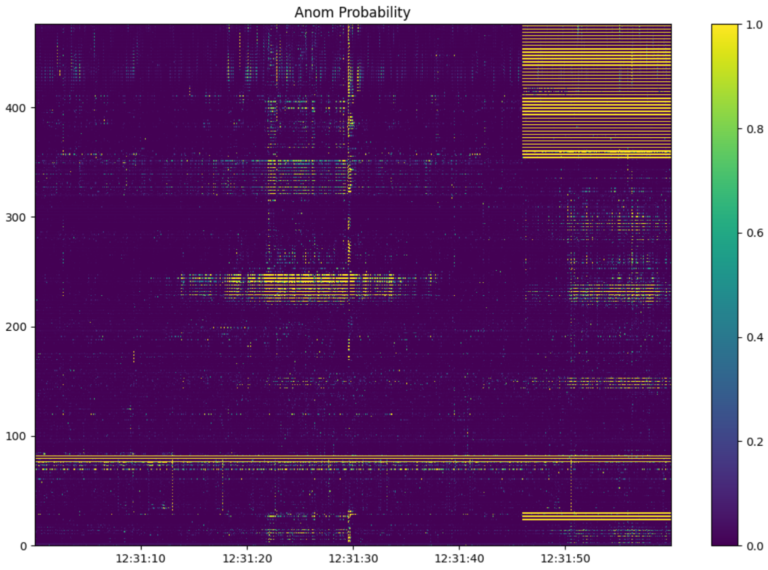

RF station fault

Minor beam loss

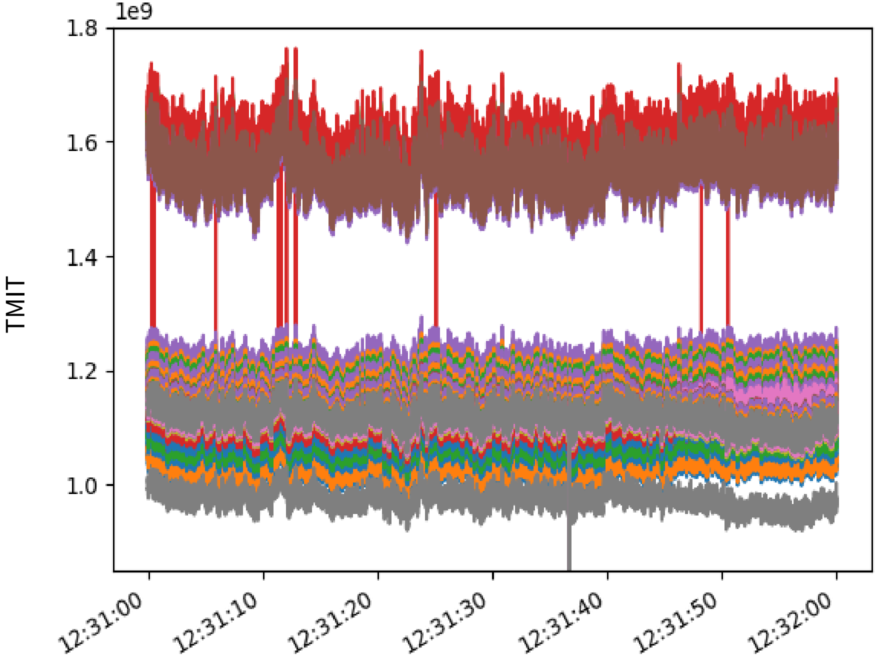

Minor beam loss