A Smart Alarm for the CEBAF Injector

Chris Tennant*, Brian Freeman, Reza, Kazimi, Daniel Moser | Jefferson Lab

Dan Abell, Jonathan Edelen, Joshua Curtis-Einstein | radiasoft

We present initial results from a proof-of-concept "smart alarm" for the CEBAF injector. Because of the injector's large number of parameters and possible fault scenarios, it is highly desirable to have an autonomous alarm system that can quickly identify and diagnose unusual machine states. Our approach leverages a trained neural network to not only identify an anomalous machine state, but also to identify the root-cause by pinpointing the specific element or region responsible. We developed an inverse model trained on data collected during normal operations. Using the inverse model, measurements from the machine are used to compute machine settings, which are then compared to EPICS setpoints. Instances when predictions differ from EPICS setpoints by a user-defined threshold are flagged as anomalies, and the user is alerted to the issue. We present the results of our data collection efforts, model training and performance, and initial performance metrics.

Authored by Jefferson Science Associates, LLC under U.S. DOE Contract No. DE-AC05-06OR23177

U.S. DEPARTMENT OF ENERGY | Office of Science

Introduction and Motivation

A data-driven tool capable of alerting operators

1. when an anomalous condition exists in the beamline
2. identifying the element setting that is the root cause

- The tool is based on an inverse model that maps beamline *readings* (diagnostic readbacks) to *settings* (beamline attributes operators can modify)
- The model leverages machine learning (ML) and is trained on data representing normal conditions
- Model-predicted settings are compared to Experimental Physics and Industrial Control System (EPICS) setpoints → instances where predictions exceed the EPICS setpoints by a user-defined threshold are flagged as anomalous

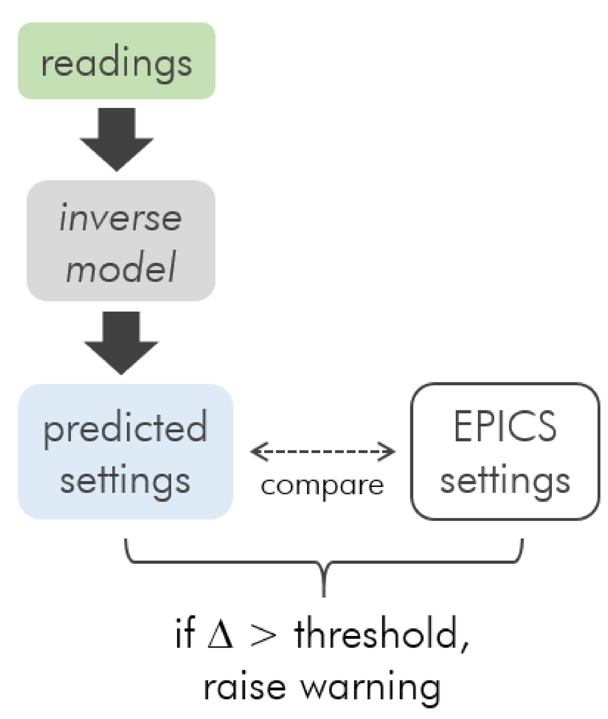

Data Preparation and Model Development

- As a case study we consider a 95 m section of the CEBAF injector
- Data for model training/testing was extracted from the CEBAF archiver
- Data was averaged for 1-minute and collected at 1-minute intervals from May 24, 2021 to January 7, 2022
  - ✓ *329,132 samples*
  - ✓ *for each sample 215 setting PVs and 234 reading PVs*
- Raw data is filtered with an emphasize on data quality rather than quantity
  - ✓ *94,327 samples*

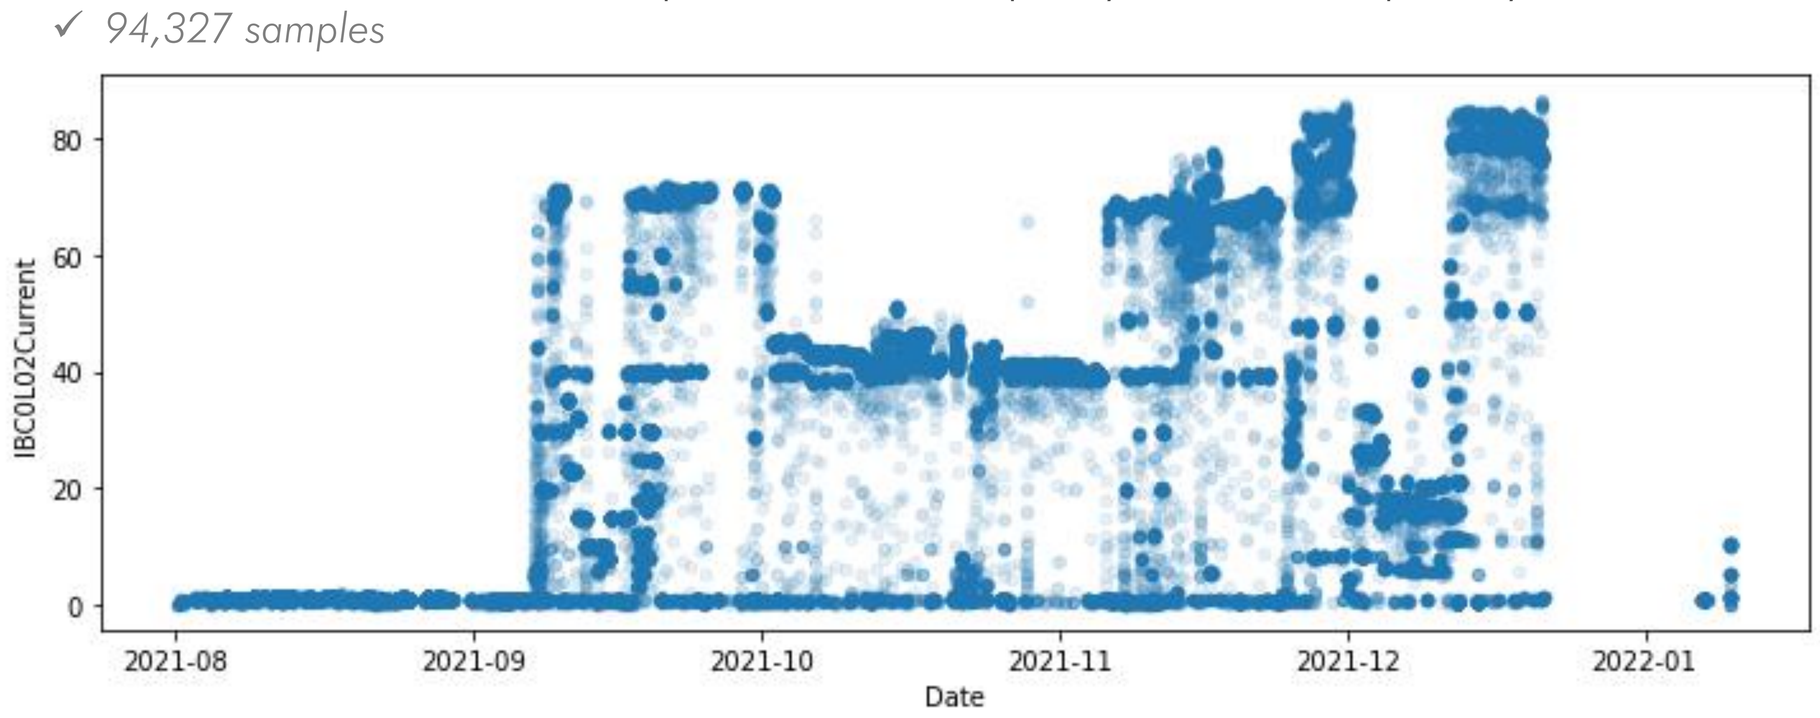

Plot of the current from an injector beam current monitor for the filtered data. The month of August was used for machine setup at low current and operation to user end stations commenced in early September. The gap beginning at the end of December and extending to early January represents the holiday shutdown.

- The model architecture utilizes a fully-connected neural network with three hidden layers of [100, 200, 400] neurons
- A combination of the Adam optimizer for the initial 2,000 epochs, followed by stochastic gradient descent (SGD) for an additional 895 epochs, resulted in the best model performance

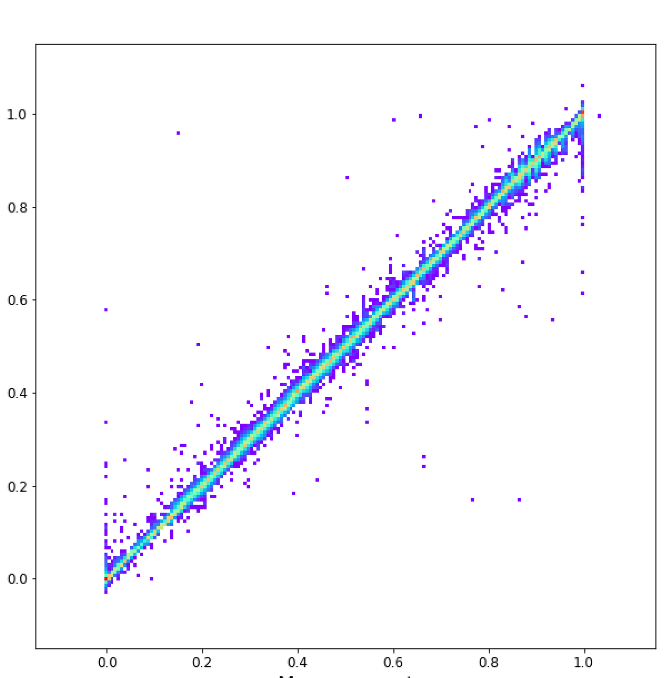

Results

We perform three tests:

1. evaluate how accurately the model is able to identify the (setting) PV being changed given only information about the readings
2. test how well the model can identify anomalous configurations
3. compare Smart Alarm performance with existing methods of identifying anomalous conditions

Anomalous Data Collection

- A dedicated beam study was used to collect anomalous injector configurations
- Specific beamline elements were varied (solenoid, corrector and quadrupole strengths, RF cavity gradients and phases) one by one in a systematic way
- Changes generated a measurable downstream response, but small enough that beam was still transmitted to an insertable dump
- Data were taken at a variety of current settings: (1, 5, 10) μA → 354 unique injector configurations

1. Comparison with Ground Truth

- Consider the 354 examples in which a setting PV was varied
- For each instance predicted settings are subtracted from the actual settings in the machine at that time and generate reconstruction errors for each PV
- The PVs corresponding to the three largest reconstruction errors are reported and compared to the ground truth

	Top 1	**Top 2**	**Top 3**
Accuracy	77.4%	92.1%	94.6%

- 6 of the instances not correctly identified by the model were for R01XPSETCG, which is a composite signal of the four chopper cavity phases ganged together.
- this PV had been inadvertently left out of the training data → model had no knowledge of the existence of the PV
- However, is that the model's top four reconstruction errors consistently predicted each of the four chopper phases as being the source of the anomalous condition

2. Flagging Anomalous Machine States

- Supplement data from the beam study with data from the period of normal operations
- Thresholds are established for each setting PV and when the reconstruction error exceeds the threshold it is flagged as anomalous
  - ✓ *we establish a threshold for each setting PV individually by taking the maximum reconstruction error from the 18,866 test instances*
- Approximates how a fully deployed version of the model would function
- If we consider if any of the top three reconstructed errors exceed their respective thresholds, 275 of the 708 examples are flagged as anomalous
- To visualize the model's performance we use t-SNE to reduce the dimensionality of the reading PVs for each injector configuration from 234 down to 2 dimensions
- All flagged configurations are from the beam study when PVs were varied

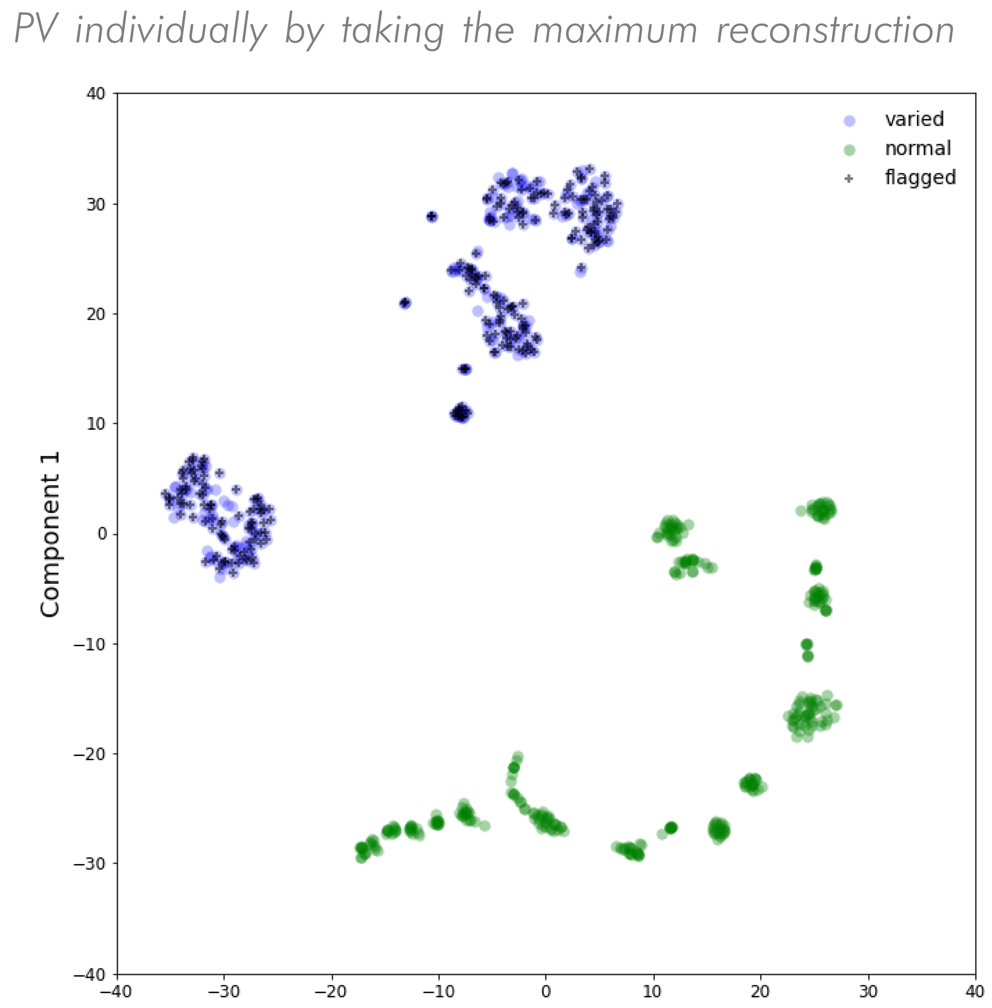

Injector configurations associated with the beam studies are denoted by blue markers and configurations taken from normal operation are denoted by green markers. Machine states flagged by the model as anomalous are represented by a black marker.

3. Comparison with Existing Methods of Anomaly Detection

- Current method is based on a configuration file that lists particular PVs and specifies upper and lower limits to trigger a warning
  - ✓ *hard-coded, heuristic approach that is unable to dynamically adapt to changes*
- Current method and the Smart Alarm agree on 142 instances as being anomalous
- While the current method identifies 8 anomalous instances that the Smart Alarm does not, *the Smart Alarm identifies 133 anomalous instances that the current method does not*

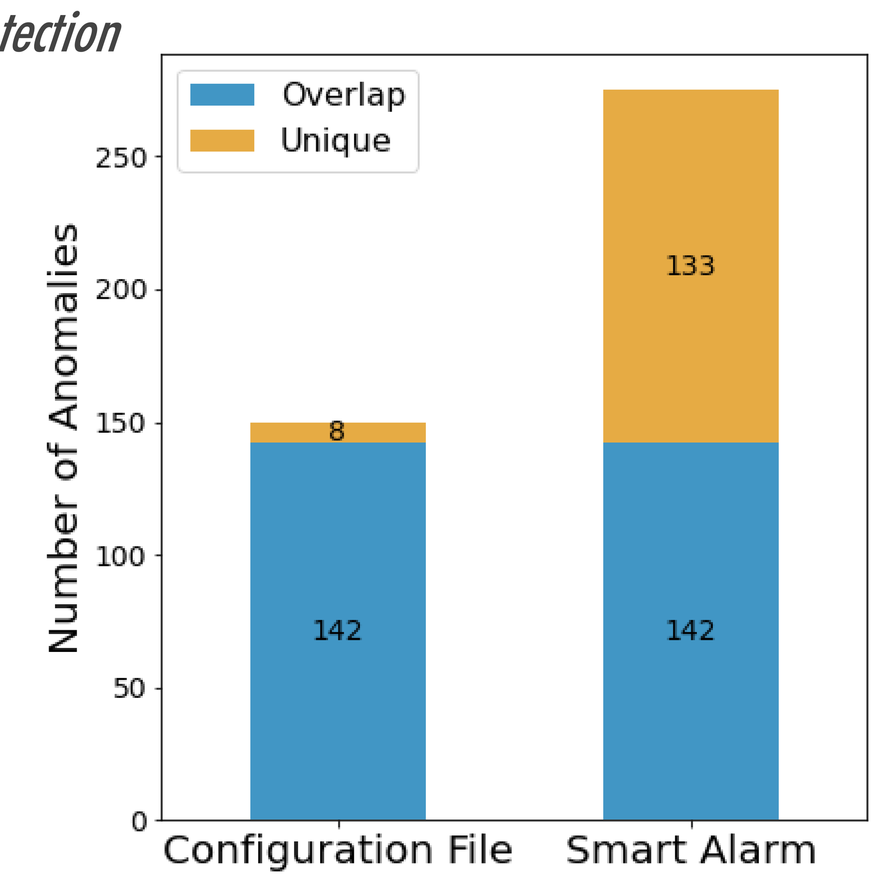

Future Work

- Explore potential of the Smart Alarm to identify the geographic location of the root cause of an anomalous condition – *even if the root cause itself is not associated with a PV*
- Implement scheduled training in order to maintain model performance and guard against
  - ✓ *concept drift*: a change in the relationship between inputs and outputs, or
  - ✓ *data drift*: changes in the underlying distribution of the inputs
- Extend the framework to other regions, and/or larger beamlines, in CEBAF
  - ✓ data for model training is collected passively by mining the operational archiver

*tennant@jlab.org