

PAPER • OPEN ACCESS

A smart alarm for particle accelerator beamline operations

To cite this article: Chris Tennant *et al* 2023 *Mach. Learn.: Sci. Technol.* **4** 015021

View the [article online](#) for updates and enhancements.

You may also like

- [Theoretical characterization of uncertainty in high-dimensional linear classification](#)
Lucas Clarté, Bruno Loureiro, Florent Krzakala et al.
- [Spectrally adapted physics-informed neural networks for solving unbounded domain problems](#)
Mingtao Xia, Lucas Böttcher and Tom Chou
- [Shift-curvature, SGD, and generalization](#)
Arwen V Bradley, Carlos A Gomez-Uribe and Manish Reddy Vuyyuru



PAPER


A smart alarm for particle accelerator beamline operations

OPEN ACCESS

RECEIVED
7 November 2022REVISED
17 January 2023ACCEPTED FOR PUBLICATION
6 February 2023PUBLISHED
16 February 2023

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Chris Tennant^{1,*} , Brian Freeman¹, Reza Kazimi¹, Daniel Moser¹, Dan Abell², Jonathan Edelen² and Joshua Einstein-Curtis²¹ Jefferson Laboratory, 12000 Jefferson Avenue, Newport News, VA 23606, United States of America² RadiaSoft LLC, 1790 38th Street, Suite 306, Boulder, CO 80301, United States of America

* Author to whom any correspondence should be addressed.

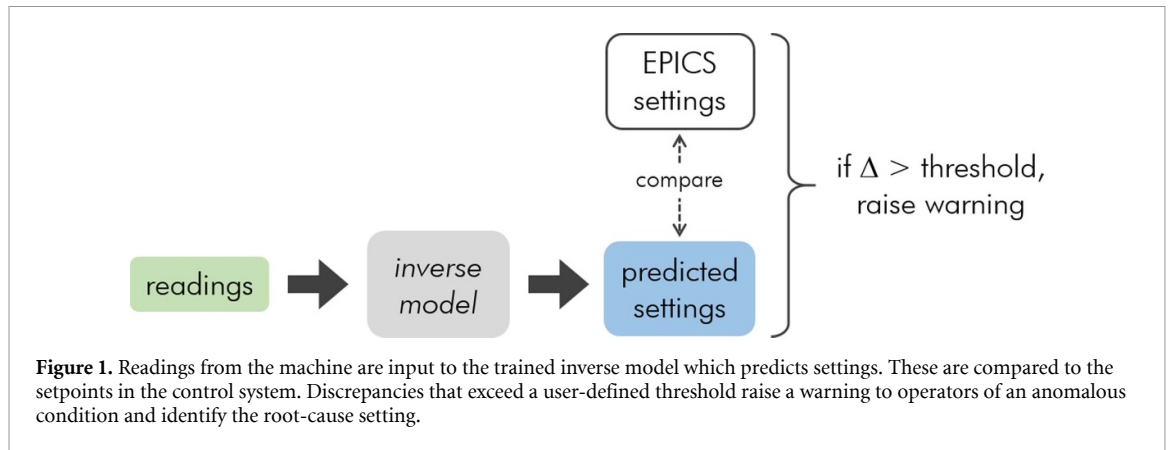
E-mail: tennant@jlab.org**Keywords:** particle accelerator, machine learning, operations, alarm system, anomaly detection**Abstract**

We present the initial results of a proof-of-concept ‘smart alarm’ for the Continuous Electron Beam Accelerator Facility injector beamline at Jefferson Lab. To minimize machine downtime and improve operational efficiency, an autonomous alarm system able to identify and diagnose unusual machine states is needed. Our approach leverages a trained neural network capable of alerting operators (a) when an anomalous condition exists in the beamline and (b) 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 takes as input readings from the machine and computes machine settings which are compared to control setpoints. Instances where predictions differ from setpoints by a user-defined threshold are flagged as anomalous. Given data corresponding to 354 anomalous injector configurations, the model can narrow the root cause of an anomalous condition to three potential candidates with 94.6% accuracy. Furthermore, compared to the current method of identifying anomalous conditions which raises an alarm when machine parameters drift outside their normal tolerances, the data-driven model can identify 83% more anomalous conditions.

1. Introduction and motivation

A non-trivial aspect of accelerator operations is identifying the root cause of a faulty machine state. For example, if the machine protection system indicates a trip due to an excessive reading on a beam loss monitor (BLM), what is the underlying cause? Sometimes the reason is obvious, while other times it is not. Existing alarm systems are commonly used to indicate when specific machine parameters are drifting outside their normal tolerances. However, operators are still required to interpret these alarms in the context of many interacting systems and take appropriate corrective action. Configuring these alarm systems is often based on hard-coded heuristics, making them difficult to easily adapt to intentional changes in beamline setups. It is even more difficult to diagnose situations that arise from changes or drifts that do not exceed established tolerances, nor cause the machine to trip. Evidence of these changes is exhibited in the downstream response of the beam orbit, an aspect of the beam itself (beam size, emittance), or a parameter that the end users are sensitive to (beam position stability, bunch charge). To address this latter issue, we have developed a data-driven method capable of alerting operators (a) when an anomalous condition exists in the beamline, and (b) identifying the element setting that is the likely 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. The model takes as input readings from the machine and computes machine settings which are compared to Experimental Physics and Industrial Control System (EPICS) setpoints [1]. Instances where predictions exceed the EPICS setpoints by a user-defined threshold are flagged as anomalous. This workflow is illustrated in figure 1.

To demonstrate this concept we developed a tool, referred to as a Smart Alarm, for the Continuous Electron Beam Accelerator Facility (CEBAF) injector beamline. In section 2 we provide a brief review of



related work regarding anomaly detection in the context of particle accelerators. In section 3 we describe the data preparation, including how data is mined from an operational archiver, filtered, split, and preprocessed, and in section 4 we provide a brief overview of model training. In section 5 we discuss results of several tests, using a data set of anomalous injector configurations collected during a dedicated beam study, aimed at quantifying the model's performance for detecting anomalous conditions and in its ability to identify the root cause. Section 6 discusses planned future work, including the implementation of scheduled training to maintain model performance over time. And in section 7 we conclude with a summary of the work to date.

2. Related work

Particle accelerators represent some of the most complex scientific instruments ever designed, built, and operated. Seeking ways to improve operational efficiency so as to maximize scientific output remains a high priority across all facilities. ML has provided new tools to address this issue, with application to anomaly detection, classification, and prognostics in particular. These algorithms work by identifying unpermitted deviations from acceptable, usual or standard conditions in an automated way [2]. Given that radio-frequency (RF) cavities are the fundamental building blocks of particle accelerators, and given that these devices generate information-rich data, a lot of research has been directed toward detection, isolation, classification, and prediction of anomalies in RF systems [3–6]. Recent work also applies anomaly detection methods to superconducting magnets [7], to identify and remove malfunctioning beam position monitors (BPMs) [8], and classify or predict errant signals [9, 10], among many other applications [11–15].

3. Data preparation

We consider the CEBAF injector as a test bed for our technique. The injector beamline is well suited for ML development due to its manageable size, independent beam dump, and abundance of tuning data reflected in the operational archiver. More specifically, because the formation and evolution of the beam at low energy is critical to performance, the injector represents a region where there is a lot of beam tuning. This translates to a wealth of historical data that can be used for training models. Additionally, the injector is ideal for beam studies as it can be operated independently from the rest of the machine. As a result there are more opportunities for dedicated beam studies to collect data and/or test our algorithm. Specifically we model the beamline starting from the electron gun and extend to an insertable dump 102 m downstream.

3.1. Data mining

The CEBAF archiver provides a record of historical operational information from which we collected data to train the model [16]. An effective model learns how a well-defined set of inputs map to a well-defined set of corresponding outputs. Therefore, a necessary first step is to examine the injector beamline and identify process variables (PVs) of interest as either a setting or a reading. A setting is defined as any PV that an operator can adjust during routine beam tuning. These include solenoid, dipole, quadrupole, and corrector strengths, and the phase and gradient setpoints of RF cavities. Readings, on the other hand, are characterized by readbacks of various diagnostic systems. These include readings from BLMs, BPMs (the horizontal and vertical position as well as the wire sum), vacuum signals, beam current monitors (BCMs), and statistical descriptions of the beam extracted from a synchrotron light monitor image. Using these categories we

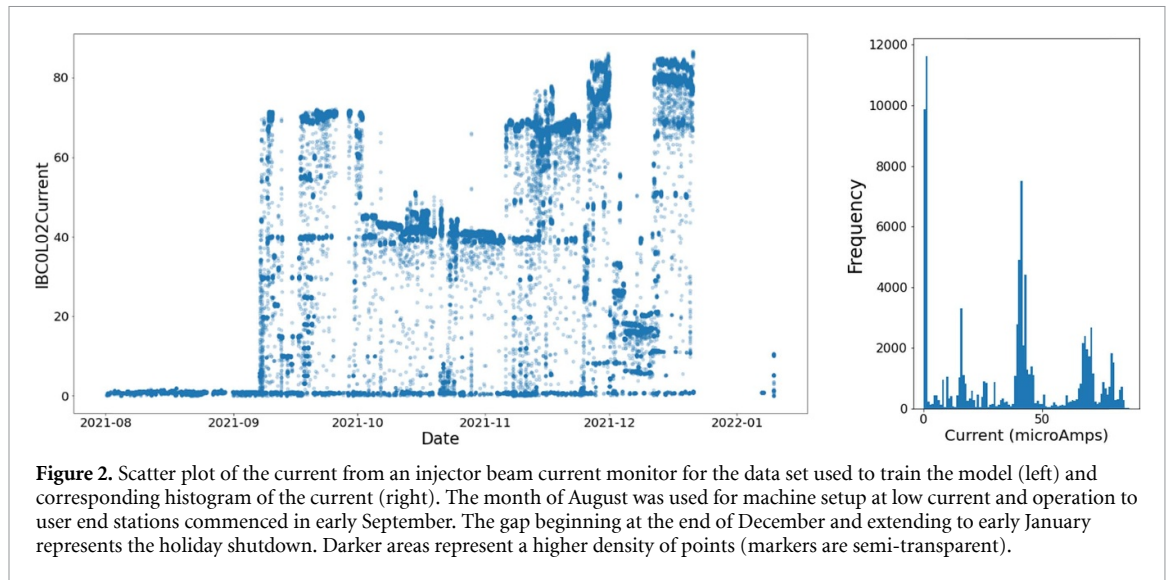


Figure 2. Scatter plot of the current from an injector beam current monitor for the data set used to train the model (left) and corresponding histogram of the current (right). The month of August was used for machine setup at low current and operation to user end stations commenced in early September. The gap beginning at the end of December and extending to early January represents the holiday shutdown. Darker areas represent a higher density of points (markers are semi-transparent).

identified 215 settings and 234 readings in the injector. Data from the archiver was averaged for 1 min and collected at 1 min intervals from 24 May 2021 to 7 January 2022. This results in 329 132 samples. In addition to the settings and readings, several ancillary signals were also collected to aid in filtering the data, including PVs that describe allowable beam modes.

3.2. Data filtering

Prior to training we performed initial data cleaning to aid in our model development. Our aim is to emphasize quality rather than quantity, but note that there is a lot of flexibility in how that is achieved. The data is filtered using the following criteria:

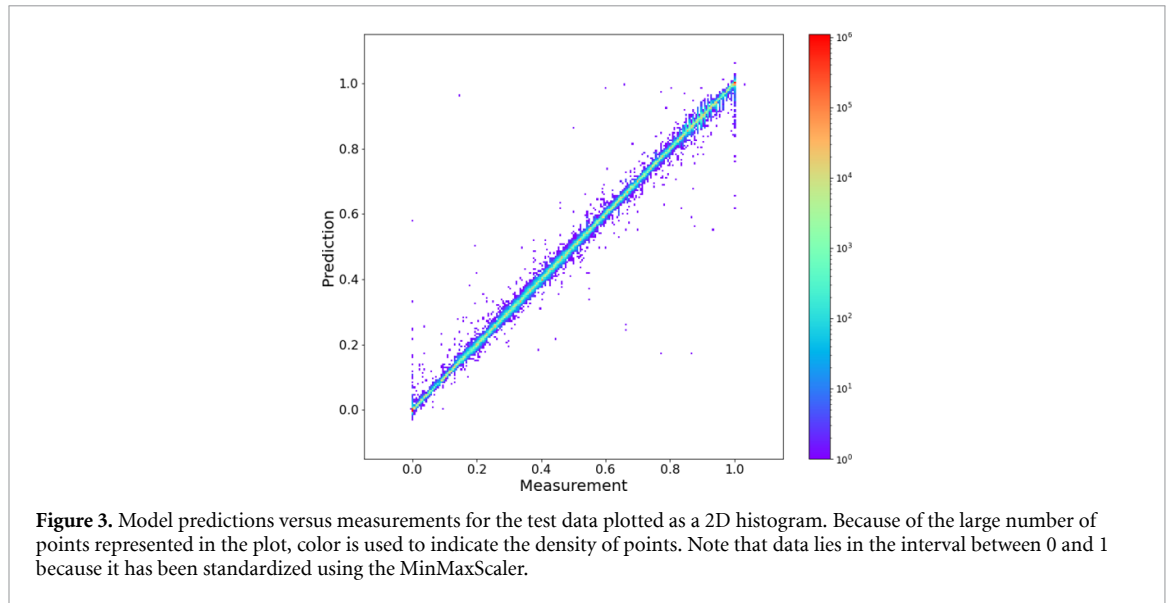
- The beam mode must be in tune, continuous wave (CW), or user mode, and the current from an injector BCM (PV name IBC0R08CRCUR1) must exceed $0.1 \mu\text{A}$. This step removes samples through the end of July—when the operational run started in earnest—and also removes samples during machine trips and planned maintenance.
- Remove samples corresponding to a machine state transition (i.e. from tune-up to CW beam, or vice-versa).
- Remove duplicate settings. Apart from small fluctuations due to inherent system noise, the reading PVs can change value either from (a) a change to setting PVs, or (b) a change in the beam current. For instance, in an uncoupled, idealized system, a change in a horizontal corrector will result in a downstream orbit displacement that will cause a change in the horizontal BPM readings. On the other hand, simply increasing the beam current does not (ideally) change the horizontal position, but the BPM wire sum—which serves as a proxy for beam current—will change. It follows that if the setting PVs are all the same and the beam current is the same, then the reading PVs will be the same and these duplicate samples are removed (because the BCM PV is recorded to the nearest hundredth of a nA, we first round the current to the nearest tenth of a μA for comparison across samples).

Removal of duplicate data has important implications for ML performance as well. When randomly splitting the data for training, duplicate data points could be in both the training and validation sets, and result in overfitting. By removing duplicates we ensure a higher degree of robustness in our training process.

Following the application of these criteria the data size is reduced from 329 132 samples to 94 327 samples. The BCM current from the final data set is plotted in figure 2 and clearly indicates periods of commissioning, running for users, and planned down time.

4. Model development

Prior to training the model the data were randomly split into training (60%), validation (20%), and testing (20%) sets and the MinMaxScaler used to pre-process the data. See appendix for examples of visualizations used to inspect splits. Following the data collection, filtering, splitting, and pre-processing, the final number of training examples is 56 595 while validation and testing sets each have 18 866 examples.



As the name suggests, an inverse model takes as input the reading PVs and predicts the setting PVs. Adjusting the weights and biases in the model to minimize a loss function is what constitutes ‘learning’. For this work we use the mean-squared error (MSE)

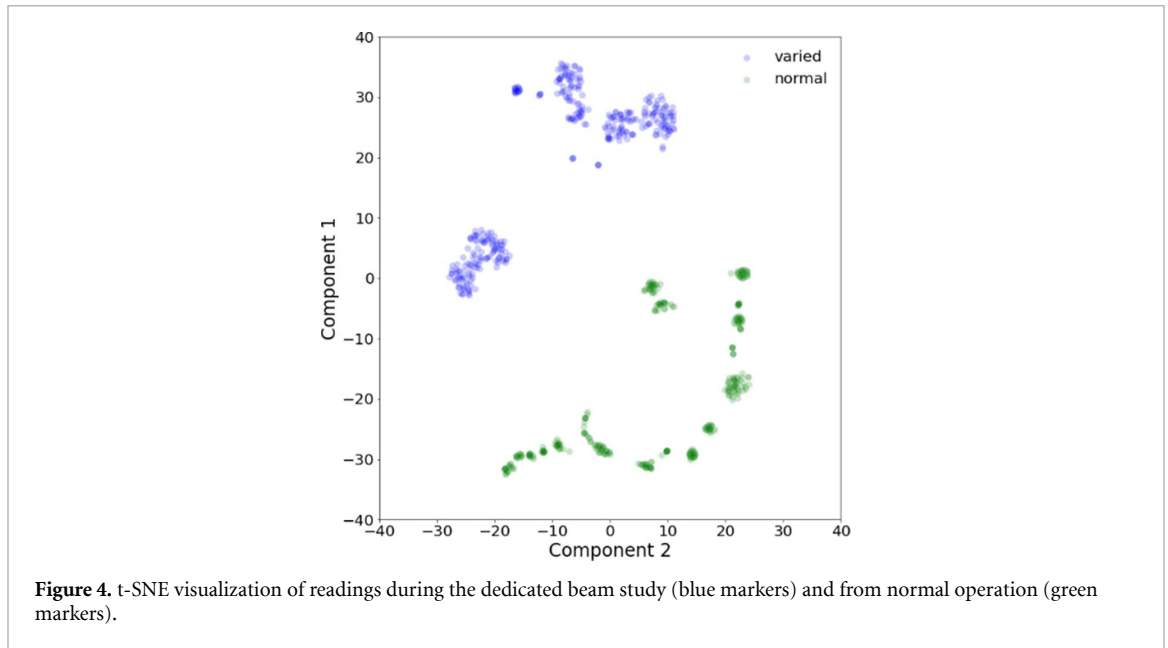
$$\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2$$

where N is the number of samples, y_i is the ground truth value and \hat{y}_i is the predicted value. The MSE loss provides a metric for evaluating model performance. For simplicity, the model architecture utilizes a fully-connected neural network and, as result of exploring the parameter space, three hidden layers of [100, 200, 400] neurons. As part of the standard exercise of hyperparameter optimization, we explored different optimizers, a variety of learning rates (including learning rate schedules), batch sizes, and the use of dropout and batch normalization layers. The best performance was achieved using a batch size of 512 with a combination of the Adam optimizer (learning rate of 0.75×10^{-3}) for the initial 2000 epochs, followed by stochastic gradient descent for an additional 895 epochs [17]. Drop and batch normalization layers provided no additional benefit and were not used. Overfitting was not observed during training. For each of the 18 866 test examples, the model predicts the value for 215 setting PVs. Plotting each of the 4056 190 predicted values ($18\,866 \times 215$) against their actual values provides a way to qualitatively evaluate model performance [18]. This is shown in figure 3. A perfect model would be evident by all points on the line $y = x$.

5. Results

To evaluate the model’s ability to identify anomalous conditions in the injector beamline, we collected data during a dedicated beam study where incremental changes were made to a variety of beamline elements and the downstream response recorded. Using this data we perform two tests. First, we measure how accurately the model is able to identify the (setting) PV being changed given only information about the readings. For the second test we supplement the data from the beam study with data from the period of normal operations (see figure 2). Thresholds are established for each setting PV based on the model performance on the data during normal operation. When the reconstruction error exceeds one of these thresholds it is flagged as anomalous. This ability to discern anomalous configurations from normal operational conditions tests how a fully deployed version of the model would function.

Several hours were spent collecting data of anomalous machine states on 9 January 2022 as part of a dedicated beam study. Specific beamline elements were varied one by one in a systematic way. For each PV change, the system was allowed to settle for 5 s and the downstream response was written to a file. Beamline components varied include solenoid, corrector and quadrupole strengths, as well as RF cavity gradients and phases. The changes to magnet strength, gradient or phase were such that a measurable downstream response was generated, but small enough that beam was still transmitted to an insertable dump at the end of the injector. Furthermore, data were taken at a variety of current settings: (1, 5, 10) μA . In total, 354 unique



injector configurations were collected. The same pre-processing steps were applied as for the training/validation/testing data described in section 3.

It is important to verify that the data collected during the study represent a parameter space that is outside of normal operations. To examine this we used t-distributed stochastic neighbor embedding (t-SNE) to reduce the dimensionality of the readings from 234 down to 2. This relatively new, non-linear, dimensionality reduction technique is especially well-suited for visualizing high-dimensional datasets [19]. Its purpose is to take high-dimensional data and faithfully map it to a low-dimensional space where it can easily be visualized. The results are shown in figure 4. We see a clear delineation between the normal operation data (blue markers) and the beam study data (green markers), which confirms that our test set represents anomalous machine states.

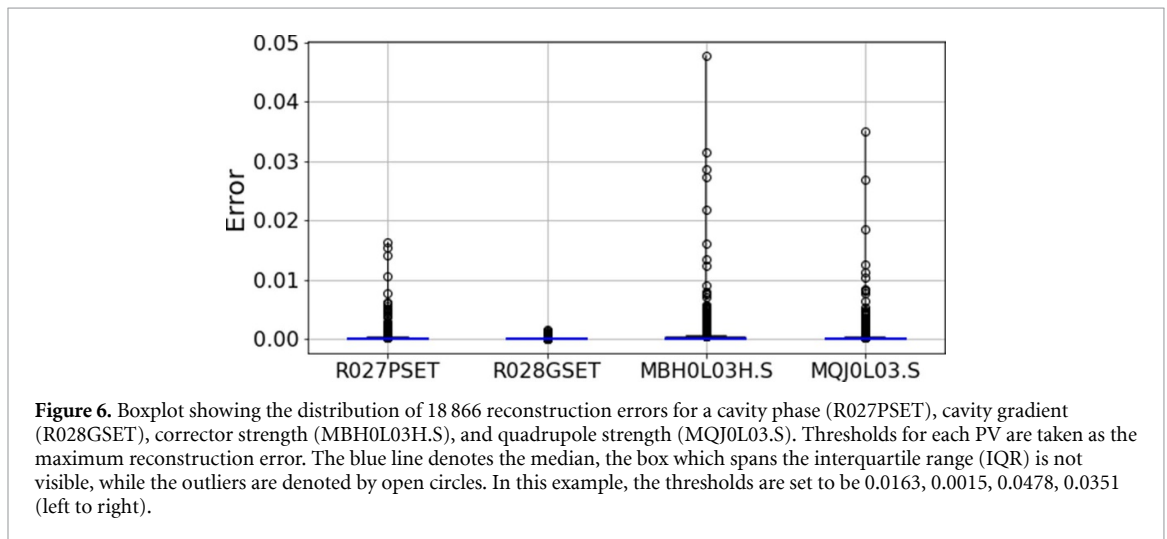
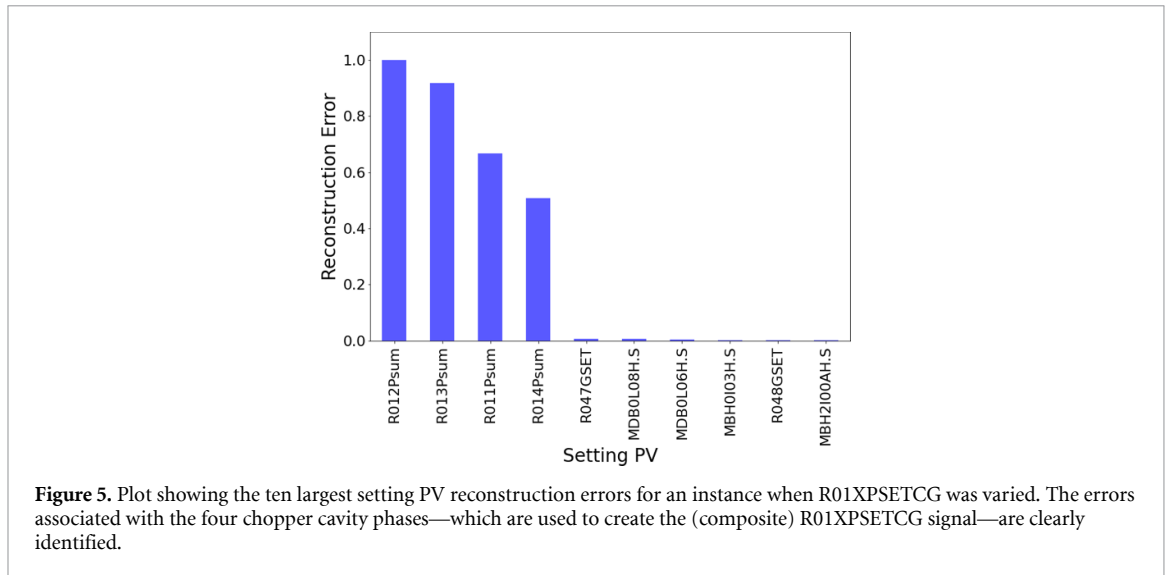
5.1. Comparison with ground truth

For the first test we consider the 354 examples in which a setting PV was varied. For each instance, the model takes as input the 234 readings and outputs predictions for the 215 settings. The predicted settings are subtracted from the actual settings in the machine at that time which generates reconstruction errors for each PV. The PVs corresponding with the three largest reconstruction errors are reported and compared to the ground truth. If we consider only the largest reconstructed error, the agreement of the model's prediction with the ground truth is 77.4%. In other words, in 274 out of the 354 cases, the model correctly identifies the PV that was varied by analyzing the injector's response (i.e. the reading PVs). The accuracy increases to 92.1% if we consider whether the first or second highest reconstructed error matches the ground truth, and reaches 94.6% if we consider the three largest reconstructed errors.

In taking a closer look at the 19 instances where the model incorrectly identified the PV that was varied, 6 involved a variable named R01XPSETCG, which is a composite signal of the four chopper cavity phases ganged together. That is, during the beam study R01XPSETCG was changed six times but the model was unable to identify the correct PV in each instance. Further investigation revealed that R01XPSETCG had been inadvertently left out of the training data. The model did not correctly identify R01XPSETCG as having changed because it had no knowledge of the existence of the PV. What is noteworthy, however, is that the model's largest four reconstruction errors consistently predicted each of the four chopper phases as being the source of the anomalous condition (see figure 5). While the exact PV was not predicted (it was absent from the training data), the model was able to correctly identify the constituent components of the PV as being the root cause.

5.2. Flagging anomalous machine states

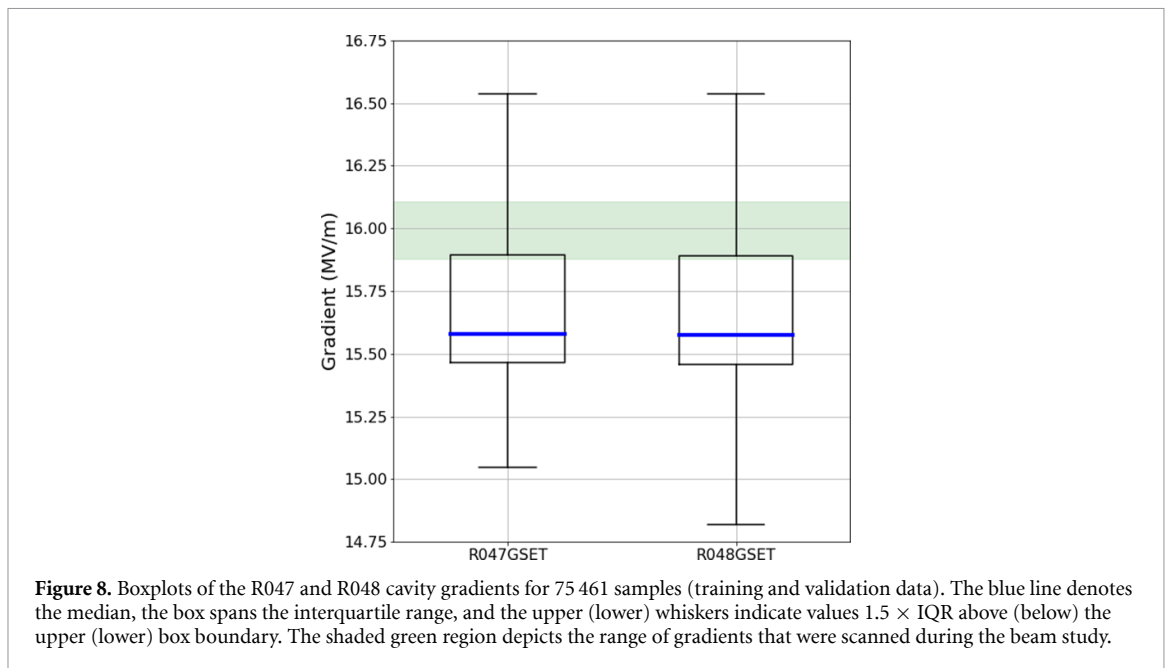
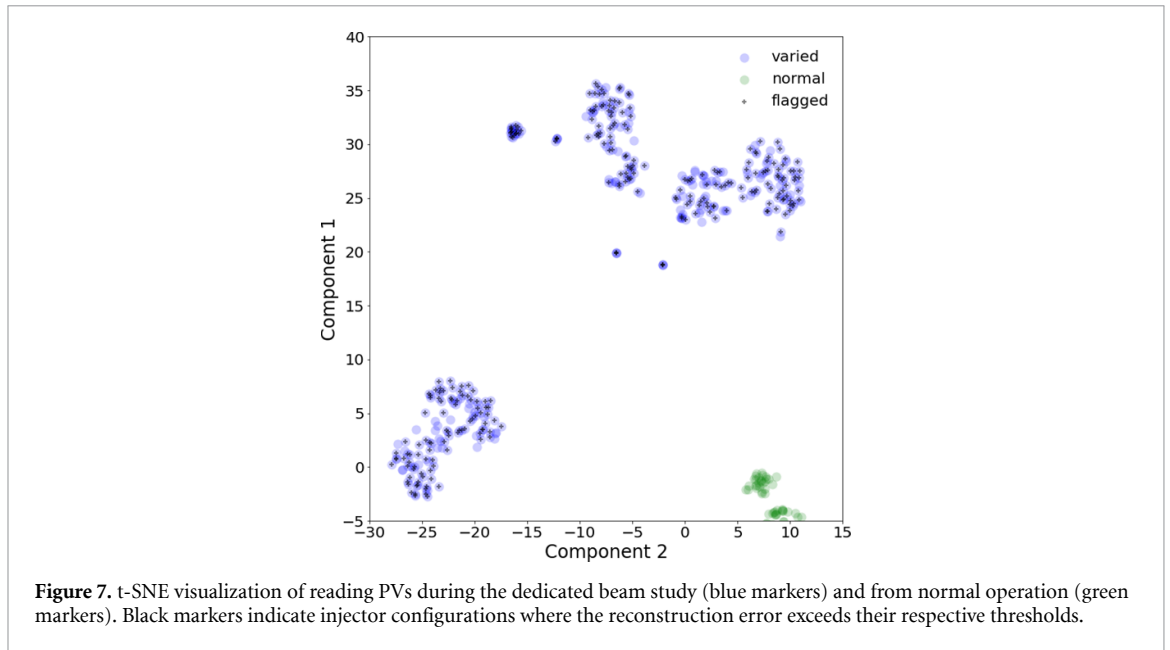
For the second test, and as a way to assess the model performance as if it were deployed, we create a new dataset by setting aside 354 samples collected from the same time period as the original training data (but not from the training data) to supplement the 354 examples of an element being varied during the beam study, resulting in 708 total samples. Unlike the previous test, simply reporting the largest n reconstruction errors is not an appropriate measure of performance as the model will always report a largest reconstructed



error—even when no anomalous condition exists. Therefore we must assign a threshold to each predicted PV which, if exceeded by its reconstruction error, flags the condition as anomalous. To do this in a data-driven way the threshold for each setting PV is established by taking its corresponding maximum reconstruction error from the 18 866 test instances. The intuition being that the test data defines stable, ideal operation and configurations outside of that distribution should be flagged. An example of reconstruction error distributions for several PVs is shown in figure 6. Establishing thresholds in a data-driven way represents a departure from the current hard-coded approach (discussed more in the following section). As with the previous test, we report the PVs with the largest three reconstructed errors for each of the 708 examples in the dataset. The machine configuration is flagged as anomalous if one of these reconstruction errors exceeds its associated threshold. If we consider only the largest reconstructed error, 270 of the 708 instances are flagged as anomalous. This number increases to 275 if we consider if any of the top three reconstructed errors exceed their respective thresholds. In practice, a global scaling factor may need to be introduced to better tune model performance based on operational experience. If the model is too sensitive, for instance, scale all PV thresholds by 1.25. Alternatively, thresholds can be modified for individual PVs.

The results of the model's performance can be visualized as shown in figure 7, where 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. Note that all the flagged configurations are from the beam study when PVs were being varied.

Of the 79 varied configurations not flagged by the model, 18 correspond to changes in gradients from cavities R047 and R048. These two cavities are used for the injector energy lock and are allowed to vary in order to maintain constant energy. Therefore the training data will reflect a much larger range of gradient

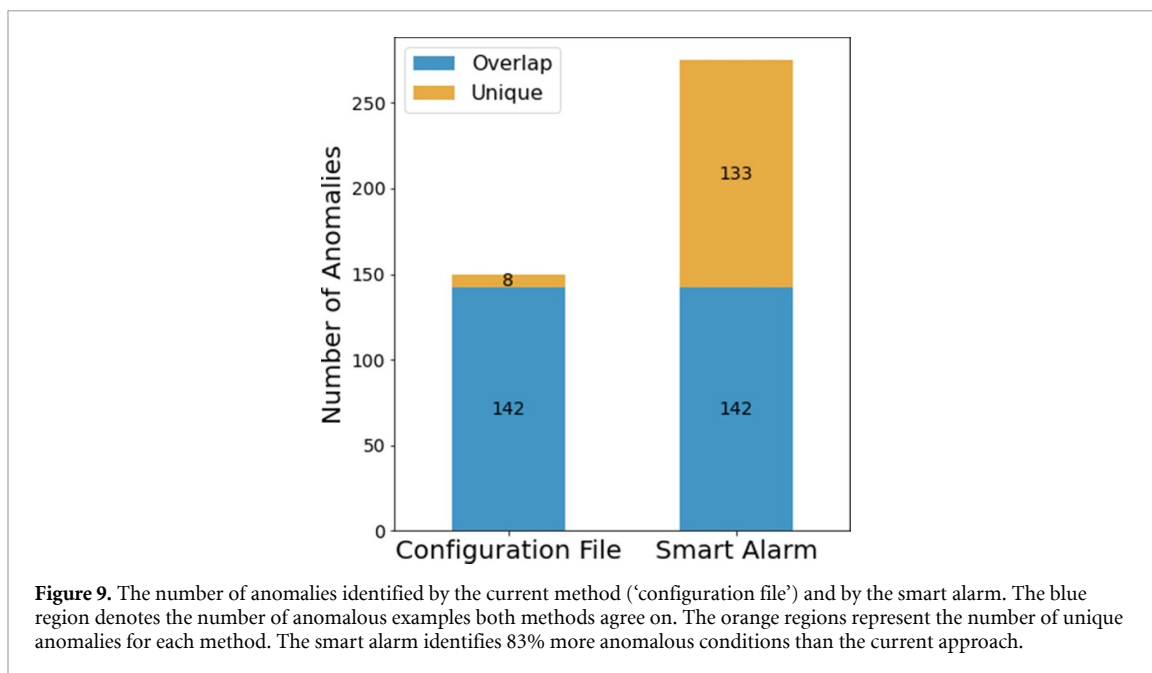


settings compared to other cavities. This is illustrated in figure 8. The model does not view the changes as anomalous since the changes are well within the distribution of data seen during training.

5.3. Comparison to existing methods of anomaly detection

The current method of alerting operators to an anomalous injector condition is based on a configuration file that lists particular PVs of interest and specifies upper and lower limits to trigger a warning, and upper and lower limits to trigger an error. This is a hard-coded approach that is unable to dynamically adapt as injector configurations change from run to run. Not all setting PVs are included and the limits themselves are set heuristically.

The result from the previous section 5.2 showed that of the 354 instances where a PV was varied in the injector, the model predicted that 275 represented anomalous conditions. Subjecting the same 354 examples to the criterion in the configuration file provides a head-to-head comparison of these two methods. Both the current method and the Smart Alarm agree on 142 instances as being anomalous. While the configuration file identifies 8 anomalous instances that the Smart Alarm does not, the Smart Alarm identifies 133 anomalous instances that the current method does not. These results are summarized in figure 9. That the Smart Alarm is able to identify 83% more anomalous conditions than the current method (275 versus 150)



demonstrates how effective a data-driven approach is. It is important to keep in mind that the anomalous conditions that are detected represent configurations where the beam is being transported through the injector. Conditions have not deteriorated to the point where excess beam loss causes the machine protection system to turn the beam off. Therefore, the Smart Alarm represents a valuable tool to alert operators to make preemptive changes to avoid what might otherwise lead to loss of beam and machine downtime.

6. Discussion and future work

The result of section 5.1 shows the 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. There is strong motivation to test this capability more fully as the potential benefits to operations is significant. By way of example, during the most recent CEBAF operational run a viewer in the front end of the injector would charge up and cause the beam to be missteered. When the viewer discharged, beam delivery to two of the experimental halls was interrupted. In this instance the root cause was not tied to a system PV (the model is unaware of beamline viewers). If the model identified a corrector in the geographic vicinity of the viewer as being the root cause of the viewer-induced steering, even though incorrect, merely localizing the problem region would result in a much faster resolution of the issue. We are unable to test this hypothesis, however, the potential to isolate a root cause issue to a specific geographic region makes this a valuable operational tool that will be explored further.

Another focus of future work will be on continual, or scheduled, training in order to maintain model performance. For any trained model deployed in a real world environment, performance degradation is inevitable. Whether the model drift is due to a change in the relationship between inputs and outputs (concept drift) or due to changes in the underlying distribution of the inputs (data drift), addressing the issue is critical to maintain optimal performance [20, 21].

During the initial training of our model we experienced the consequences of data drift and the impact it can have on performance. We trained the model on nearly seven months of data and tested it on data collected during the beam study. Model performance was very poor and further investigation revealed that in the day between the end of the training data and the beam study, a single horizontal corrector (MBH1I02H) had been changed to a value far outside its range in the training data. Even though it was a single corrector, because it is located in the very front end of the injector beamline, it created a substantial impact downstream. We added 202 examples with the new corrector setting to the training data, retrained the model, and achieved the excellent results reported in section 5.

Therefore, we cannot assume that future data will be similar to the past data used to train the model. One simple approach is to generate training data at the end of each 24 h day. Collecting data every 5 min would create a data set of 288 samples (subject to the constraint that beam is cleanly transported through the

injector). If we consider the base, deployed model as having been trained on dataset A collected during time interval t_1 , our goal is to have good performance on a second dataset B collected during a time interval t_2 . There are several options on how training can proceed. The first option is to maintain the same architecture but randomly initialize model weights and train from scratch on dataset B. The second option is to start with weights of the base model and train with dataset B. A third option is to freeze the weights of some layers and train weights of others using dataset B. In general, if dataset B is substantially different than dataset A then the first two options are preferable (transfer learning), whereas if datasets A and B are similar, the last option is preferred (fine-tuning).

Finally, while the injector beamline provided an ideal testbed to demonstrate the efficacy of a ML-based Smart Alarm, there is nothing to prevent this framework from being extended to other regions, or even to the entirety, of CEBAF. It is worth emphasizing that the data used for training the model was collected passively, by mining the operational archiver, and can likewise be utilized for training models for other parts of the accelerator.

7. Conclusion

In this work we describe the development of a novel data-driven alarm system and report results of several tests applied to the CEBAF injector beamline at Jefferson Lab. The so-called Smart Alarm is a neural network-based inverse model that maps beamline readings to settings. Instances where predictions differ from machine setpoints by a user-defined threshold are flagged as anomalous. Data from an operational archiver is used to train the model, and benefits from the fact that data collection is ongoing and passive.

A data set comprised of 354 anomalous injector configurations was collected for the purpose of testing system performance. In the first test, the model demonstrated the ability to narrow the root cause of an anomalous condition to three potential candidates with 94.6% accuracy. For a 95 m long beamline comprised of several hundred PVs that the operators must monitor, this represents a very encouraging result. In the second test we compared the model's ability to detect anomalous machine conditions with the current method, which raises an alarm when machine parameters drift outside preset, hard-coded tolerances. The Smart Alarm was able to identify 83% more anomalous conditions than the current method, demonstrating the benefits of a data-driven approach.

The development of an autonomous alarm system that is capable of alerting operators when an anomalous condition arises in the beamline, while simultaneously identifying the element setting at the root cause, is a much needed tool for operators at accelerator facilities. The encouraging results reported in this paper motivate continued efforts to expand the scope of the tool at Jefferson Lab. Future work will also address challenges with maintaining model performance in a changing environment through continual, or scheduled, training.

Data availability statement

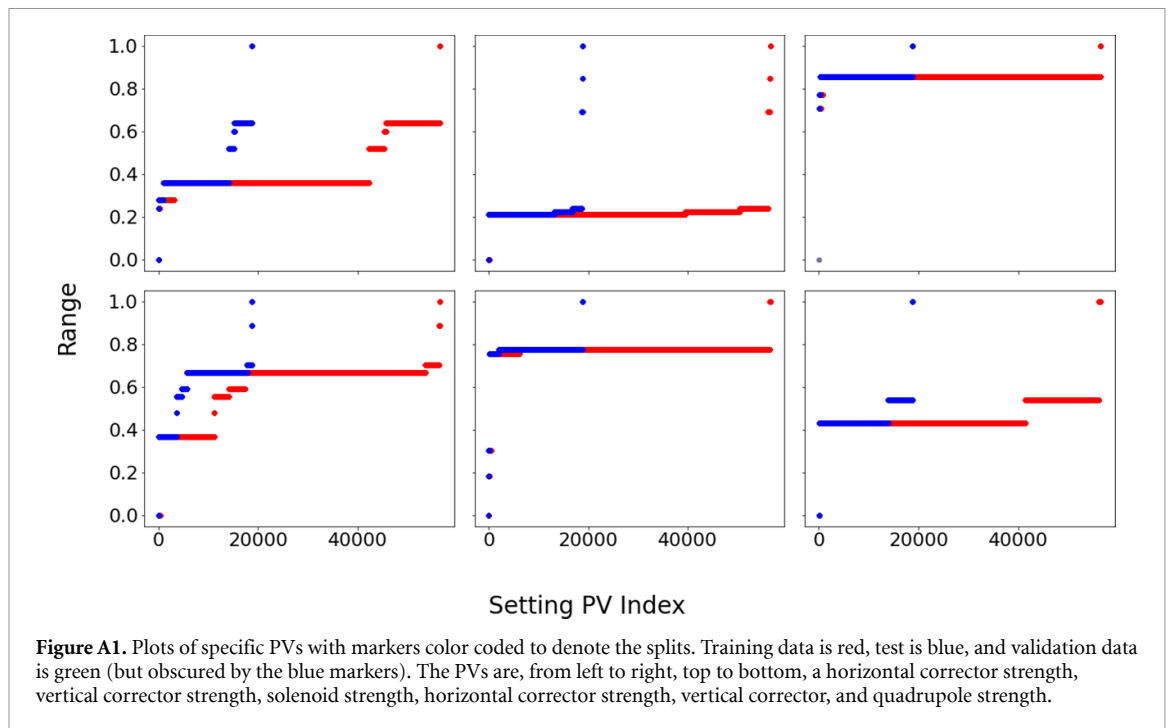
The data that support the findings of this study are openly available at the following URL: <https://github.com/JeffersonLab/smart-alarm>.

Acknowledgments

This work is supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics under Contract No. DE-AC05-06OR23177.

Appendix . Data splits

Careful attention was given to the training, validation, and testing data splits. In particular, we wanted to ensure that the validation and testing data was similar to the training data. One of the ways we verified this is using visualizations like those depicted in figure A1. Normalized data from each PV is sorted by value (from smallest to largest) and plotted. The train, validation, and test sets are denoted by red, green, and blue markers, respectively. For a given PV, the vertical extent of each dataset should be comparable. Issues would arise, for instance, if a model was trained on data with a particular distribution and then tested on data that had a different distribution.



ORCID iD

Chris Tennant  <https://orcid.org/0000-0003-3814-8417>

References

- [1] Epics Development Team 2013 EPICS: experimental physics and industrial control system (ascl:1302.005) (Astrophysics Source Code Library)
- [2] Isermann R 2006 *Fault-Diagnosis Systems: An Introduction from Fault Detection to Fault Tolerance* (Heidelberg: Springer) [10.1007/3-540-30368-5](https://doi.org/10.1007/3-540-30368-5)
- [3] Tennant C, Carpenter A, Powers T, Shabalina Solopova A, Vidyaratne L and Iftekharuddin K 2020 Superconducting radio-frequency cavity fault classification using machine learning at Jefferson Laboratory *Phys. Rev. Accel. Beams* **23** 114601
- [4] Obermair C et al 2022 Explainable machine learning for breakdown prediction in high gradient rf cavities *Phys. Rev. Accel. Beams* **25** 104601
- [5] Humble R, O'Shea F H, Colocho W, Gibbs M, Chaffee H, Darve E and Ratner D 2022 Beam-based rf station fault identification at the SLAC Linac Coherent Light Source *Phys. Rev. Accel. Beams* **25** 122804
- [6] Nawaz A, Lichtenberg G, Pfeiffer S and Rostalski P 2018 Anomaly detection for cavity signals—results from the European XFEL Proc. in 9th Int. Particle Accelerator Conf. IPAC'18 pp 2502–4
- [7] Wielgosz M, Skoczea A and Mertik M 2017 Using LSTM recurrent neural networks for monitoring the LHC superconducting magnets *Nucl. Instrum. Methods Phys. Res. A* **867** 40–50
- [8] Fol E, Tomás R, De Portugal J C and Franchetti G 2020 Detection of faulty beam position monitors using unsupervised learning *Phys. Rev. Accel. Beams* **23** 102805
- [9] Blokland W, Rajput K, Schram M, Jeske T, Ramuhalli P, Peters C, Yucesan Y and Zhukov A 2022 Uncertainty aware anomaly detection to predict errant beam pulses in the Oak Ridge Spallation Neutron Source accelerator *Phys. Rev. Accel. Beams* **25** 122802
- [10] Radaideh M I, Pappas C and Cousineau S 2022 Real electronic signal data from particle accelerator power systems for machine learning anomaly detection *Data Brief* **43** 108473
- [11] Grünhagen A et al 2021 Fault analysis of the beam acceleration control system at the European XFEL using data mining 2021 *IEEE 30th Asian Test Symp. (ATS)* pp 61–66
- [12] Valentino G et al 2017 Anomaly detection for beam loss maps in the Large Hadron Collider *J. Phys.: Conf. Ser.* **874** 012002
- [13] Piekarski M, Kitka W, Solaris N and Jaworek-Korjakowska J 2019 Deep neural network for anomaly detection in accelerators Proc. 17th Int. Conf. on Accelerator and Large Experimental Physics Control Systems, ICALEPCS'19 p 1379
- [14] Coyle L et al 2021 Detection and classification of collective beam behaviour in the LHC Proc. 12th Int. Particle Accelerator Conf., IPAC'21 pp 4318–21
- [15] Tilaro F et al 2018 Model learning algorithms for anomaly detection in CERN control systems 16th Int. Conf. on Accelerator and Large Experimental Physics Control Systems, ICALEPCS'18 [10.18429/JACoW-ICALEPCS2017-TUCPA04](https://doi.org/10.18429/JACoW-ICALEPCS2017-TUCPA04)
- [16] Slominski C 2009 A MySQL based epics archiver Proc. 12th Int. Conf. on Accelerator and Large Experimental Physics Control Systems ICALEPCS'09 WEP021 pp 447–9
- [17] Keskar N S and Socher R 2017 Improving generalization performance by switching from Adam to SGD (arXiv: [abs/1712.07628](https://arxiv.org/abs/1712.07628))
- [18] Tennant C et al 2023 The saved inverse model, data to evaluate model performance, and data to simulate deployed use is (available at: <https://github.com/JeffersonLab/smart-alarm>)
- [19] van der Maaten L and Hinton G 2008 Visualizing high-dimensional data using t-SNE *J. Mach. Learn. Res.* **9** 2579–605
- [20] Webb G I, Lee L K, Petitjean F and Goethals B 2017 Understanding concept drift (arXiv: [1704.00362](https://arxiv.org/abs/1704.00362))
- [21] Webb G I, Hyde R, Cao H, Nguyen H L and Petitjean F 2016 Characterizing concept drift *Data Min. Knowl. Discov.* **30** 964–94