ISBN: 978-3-95450-232-5 ISSN: 2673-7000 doi:10.18429/JACoW-NAPAC2022-TUPA29

MACHINE LEARNING FOR PREDICTING POWER SUPPLY TRIPS IN STORAGE RINGS*

I. Lobach†, M. Borland, G. Fystro, A. Sannibale, Y. Sun
Argonne National Laboratory, Advanced Photon Source, Lemont, IL, USA
A. Diaw, J. Edelen, RadiaSoft, Boulder, CO, USA

Abstract

In the Advanced Photon Source (APS) storage ring at Argonne National Lab, trips in the magnet power supplies (PSs) lead to a complete electron beam loss a few times a year. This results in unexpected interruptions of the users' experiments. In this contribution, we investigate the historical data for the last two decades to find precursors for the PS trips that could provide an advance notice for future trips and allow some preventive action by the ring operator or by the PS maintenance team. Various unsupervised anomaly detection models can be trained on the vast amounts of available reference data from the beamtime periods that ended with an intentional beam dump. We find that such models can sometimes detect trip precursors in PS currents, voltages, and in the temperatures of magnets, capacitors and transistors (components of PSs).

INTRODUCTION

In the APS storage ring, trips in the magnet power supplies (PSs) lead to a complete electron beam loss about 7 times a year, on average. This results in unexpected interruptions of the users' experiments. This paper explores the possibility of using Machine Learning (ML) to give an advance warning about an impending PS trip, so that a preventive action can be taken by the operator or by the PS maintenance group. We analyzed two decades of historical data. Since the total number of recorded PS trips is relatively low (149), supervised ML methods cannot be used as they would be prone to overfitting. Instead, we focus on unsupervised anomaly detection methods, such as spectral residual saliency detection [1] and neural network autoencoders, which can be trained and tested on the vast amounts of available historical data.

HISTORICAL DATA

There is a spreadsheet with detailed records for each APS run dating back to 1997 [2]. The records include each beam fill time, each beam dump time, and the reason for the beam dump. We are interested in the records with an intentional beam dump (end of period) and in the records with a beam loss due to a magnet PS trip. More specifically, we focus on trips in the PSs for quadrupoles, sextupoles, horizontal and vertical correctors. The APS storage ring consists of 40 sectors. The total number of considered PSs is 1320 [3, 4]. Several parameters of each PS and each magnet are constantly monitored and logged. We analyzed magnet temperatures, PS currents, voltages, capacitor temperatures, transistor temperatures. The capacitors are used to smoothen the input voltage. The transistors are used to regulate the currents in the magnets. Quadrupole and sextupole PSs have one transistor each. PSs for horizontal and vertical correctors have two transistors each, for positive and negative current. These parameters are logged at 1 point per 64 seconds, and the historical data go back to 2008. Another source of relevant data is the independent noise monitor [5] for the PS currents. Its raw-data acquisition rate is 100 Hz. However, only the processed data are logged. First, the mean and the mean absolute deviation (MAD) are calculated for every 500-sample window. Then, for the two obtained quantities, the minimum, the mean, and the maximum values are calculated in a 64-second long window. The 6 obtained process variables (MeanMin, MeanMean, MeanMax, MADMin, MADMean, MADMax) are logged at 1 point per 64 seconds for the most recent year. The older data are down-sampled to 1 point per 10 minutes and go back to 2001. PS data are labeled by the names of the corresponding magnets. Hence, the data may not always refer to the same physical PS, because of occasional PS replacements.

We analyzed the run history spreadsheets from 2001 to 2022 and found 629 beamtime periods that ended with an intentional dump and 149 beamtime periods that ended with a trip of one of the PSs (the records indicate which one). For further analysis, we pulled the data from the data archive for these beamtime periods and obtained 629 "reference" data files and 149 "trip" data files. These files can be as short as a half hour. However, most of them are several days long.

ANOMALIES IN TEMPERATURE DATA

The exact mechanism of most of the observed PS trips is not well understood. However, there are scenarios, in which a PS is programmed to trip, e.g., when the transistor temperature reading reaches 50 °C. Further, there are thermal switches on the magnets, which open and cause a trip at 71 ± 3 °C (water-cooled magnets) and at 95 ± 3 °C (air-cooled magnets). However, a thermal switch may open before the temperature reading reaches the manufacturer specified threshold, because the temperature is measured independently. Clearly, one could create an algorithm that warns the operator when the magnet temperature or the PS temperature crosses a certain threshold. However, we wanted to see whether there are more general anomaly detection methods that would not only be sensitive to higher-than-

* The work is supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No. DE-AC02-06CH11357.
† ilobach@anl.gov

ISBN: 978-3-95450-232-5 ISSN: 2673-7000 doi:10.18429/JACoW-NAPAC2022-TUPA29

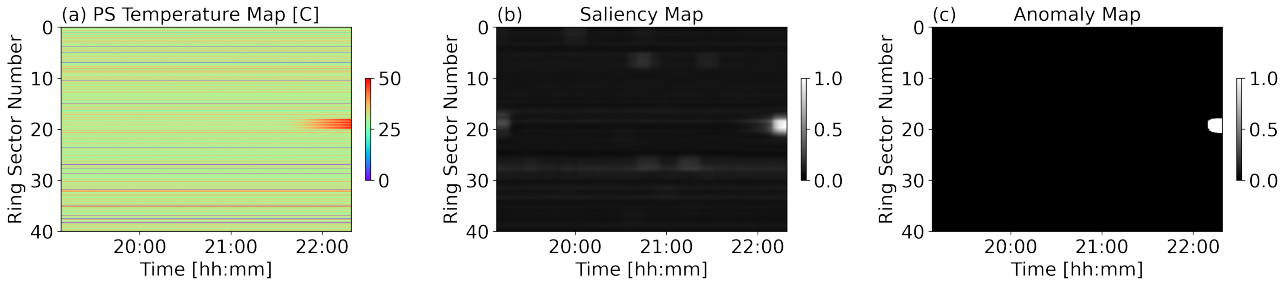

Figure 1: (a) A map of quadrupole and sextupole temperatures (680 levels) for the last 3 hours of Fill #15 of Run 2019-3. This beamtime period ended with a trip of one of the PSs in sectors 19 and 20. Color represents temperature in centigrade. Vertical axis represents the position around the storage ring. (b) Spectral residual saliency map, obtained from the temperature map. (c) Anomaly map, obtained from the saliency map using a 0.5 anomaly threshold.

normal temperatures, but also to other less trivial anomalies, that could be earlier signs of an impending trip.

One mechanism of PS trips is the following. Some magnets and PSs are water-cooled. Each pair of adjacent ring sectors has a shared cooling system with a single mixing valve controlling the inflow of cold water. If a mixing valve gets stuck and the cold water stops flowing into the system, the temperature will start rising in the two neighboring sectors. This signature may be detected before a fixed temperature threshold is reached to warn the operator. There have been at least 5 mixing valve faults that resulted in a PS trip. One example is illustrated in Fig. 1(a). Henceforth, we will focus on the transistor temperature only and we will refer to it as the "power supply temperature".

The detection of the temperature gradient in sectors 19 and 20 in Fig. 1(a) is analogous to the detection of objects in photographs. Hence, we employed the spectral residual saliency detection method [1, 6]. The map of the saliency score (from 0 to 1) is shown in Fig. 1(b). The object map (anomaly map), presented in Fig. 1(c), successfully spotlights the stuck-mixing-valve event. Depending on the parameters, the advance warning can be issued up to 30 minutes before the trip. In real time, the saliency detection can be applied to a sliding-window temperature map. The optimal window length was about 3 hours. The vertical size of the temperature map was 680, the number of considered PSs. The method is sensitive to the temperature map resolution, false positives may occur. We are working on using the saliency maps as input for a convolutional neural network for further classification of the detected anomaly candidates. This may reduce the number of false positives [7].

Now, let us consider anomaly detection with an autoencoder—a neural network with equal dimensions of input and output layers. It contains one or more hidden layers, and there is always a bottleneck layer in the middle, with a lower size than input and output layers, see Fig. 2. Autoencoders learn a compressed representation of the training data by minimizing the difference between the input and the output. An autoencoder can be used for anomaly detection in the following way. First, it is trained on the reference data, so that it can learn various patterns, typical for the reference data only. Then, when it encounters an anomalous data sample, it is unlikely to reconstruct it well. Hence, the reconstruction error constitutes an anomaly score. The threshold can be chosen based on the reconstruction error in the training data.

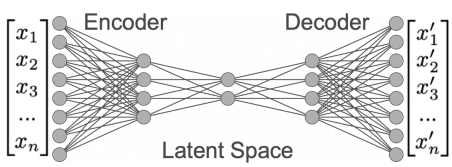

Figure 2: Example of an autoencoder [8].

Figure 3 demonstrates how the reconstruction error of an autoencoder for the PS temperatures increases in the anomalous region leading to a trip. This is the same mixing-valve incident, illustrated in Fig. 1. The first signs of the anomaly are detected about an hour before the trip. The inputs of this autoencoder are the PS temperatures, averaged by sector (40 values). Only quadrupoles and sextupoles are considered. The autoencoder layer sizes are the following, $40 \to 10 \to 5 \to 10 \to 40$. Mean squared error (MSE) is used as the loss function. Rectified Linear Units (ReLUs) are used for activation. The autoencoder was trained on 10 reference data files preceding the trip file shown in Fig. 3.

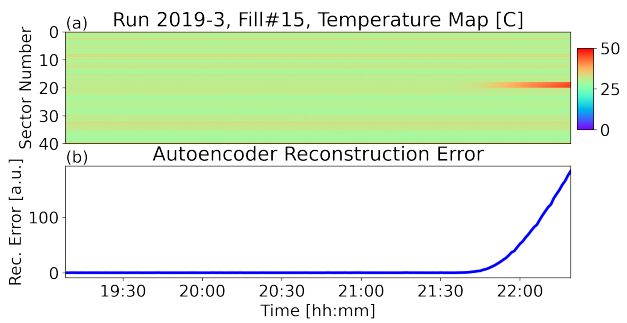

Figure 3: (a) A temperature map (averaged by sector) for the last 3 hours of a beamtime period before a trip. (b) The reconstruction error of the autoencoder as a function of time.

ISBN: 978-3-95450-232-5 ISSN: 2673-7000 doi:10.18429/JACoW-NAPAC2022-TUPA29

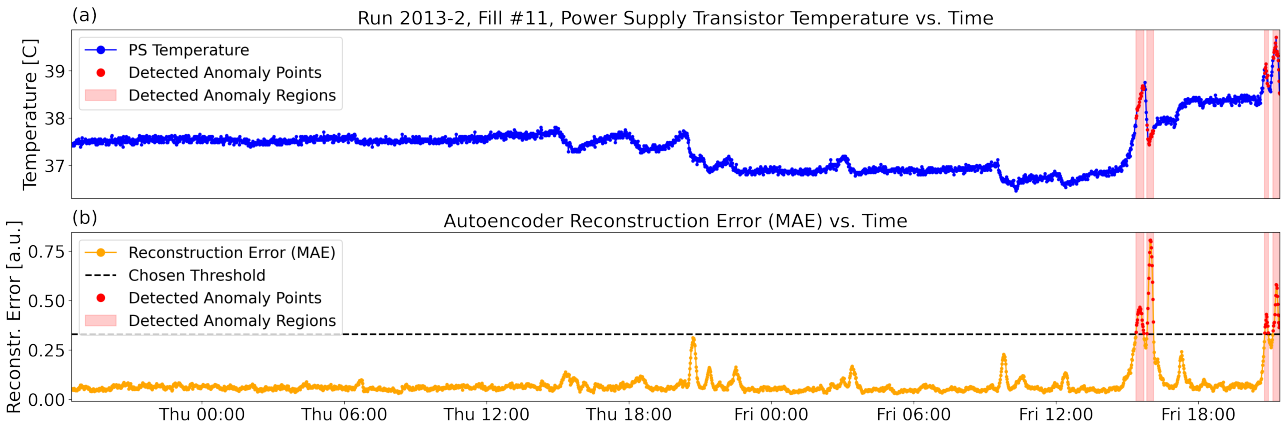

Figure 4: (a) PS temperature vs. time during Fill #11 of Run 2013-2, which ended with a glitch in this PS. (b) Reconstruction error of the autoencoder vs. time. An anomaly is declared when the reconstruction error exceeds the chosen threshold.

One can also use several contiguous temperature values of a single PS as input for an autoencoder. Figure 4 illustrates that such autoencoder was able to detect unusual behavior of a PS temperature, about 6 hours before it tripped. The size of the sliding window was 20 points (about 20 minutes). The architecture of the autoencoder was $20 \rightarrow 10 \rightarrow 5 \rightarrow 10 \rightarrow 20$. Mean absolute error (MAE) was used as the loss function. ReLU was used for activation. The anomaly threshold for the reconstruction error was chosen as 1.5 times the maximum reconstruction error in the training data. The autoencoder was trained on 10 preceding reference data files. Further, this autoencoder was tested on all available data files for the PS of the considered magnet (S16B:S2). It was always trained on 10 preceding reference files. This magnet's PS tripped only once, at the end of Fill #11 of Run 2013-2. The autoencoder detected a precursor for this trip (see Fig. 4) without producing any false positives in the entire observation period from 2008 to 2022.

ANOMALIES IN PS CURRENT DATA

Inspired by the autoencoder performance in Fig. 4, we also used this approach for the PS current noise monitor data. We trained an individual autoencoder for each PS. Each autoencoder's input size was 36, i.e., 6 process variables (MeanMin, MeanMean, MeanMax, MADMin, MADMean, MADMax) times 6 contiguous timesteps. For a more efficient use of resources, all 1320 autoencoders could be joined into one neural network [9] and trained on the GPU nodes of the LCRC cluster at ANL. Three years of preceding reference data files were used for training. The model can be re-trained before the start of every APS run. The training takes about an hour and a half. The detected anomalies are saved into a single database, and the plots with the anomalous regions can be quickly accessed in a special program.

The process variable, most relevant to anomaly detection, is the maximum value of the MAD of current in a time window (MADMax). Simpler models can be considered using this variable only. Overall, the models, based on the noise monitor data, can detect precursors in up to 20 % of trips from the historical records. There is also a sizeable number of false positives. Still, the models are up to 500 times more likely to declare an anomaly in the historical data leading to a trip, than in the reference data.

DISCUSSION AND CONCLUSIONS

We have made progress in anomaly detection in the PS current noise monitor data. At the moment, the false positive rate is not insignificant, and it is not feasible to request a response from the operator or from the PS maintenance team after every declared anomaly, in real time. However, we created an application, where an expert from the PS group can review all detected anomaly candidates at the end of each APS run, and decide whether any maintenance needs to be performed during the shutdown. The model accuracy may be improved by properly accounting for occasional PS swap-outs, usage of correctors for photon beam steering, possible dependence of PS current noise on current setpoints.

Anomalies in the PS temperature maps (see Figs. 1 and 3) can be detected very effectively by the spectral residual saliency detection method and by the neural network autoencoder. All 5 incidents with a stuck mixing valve would have an advance warning of 30 minutes to an hour before the trip. The number of false positives can be made very low. It dependents on the parameters of the models, and on their ability to generalize to other kinds of anomalies.

Autoencoders may be able to detect anomalies invisible to less sophisticated algorithms. The temperature anomaly in Fig. 4 was only detected due to its unusual behavior in time. The value of the temperature was not higher than usual.

ACKNOWLEDGMENTS

The authors are thankful to L. Emery, J. Wang, R. Wright, and A. Puttkammer for discussions about PS trips in the APS storage ring. We gratefully acknowledge the computing resources provided on Bebop and Swing, high-performance computing clusters operated by the Laboratory Computing Resource Center at Argonne National Laboratory.

ISBN: 978-3-95450-232-5 ISSN: 2673-7000 doi:10.18429/JACoW-NAPAC2022-TUPA29

REFERENCES

[1] X. Hou and L. Zhang, "Saliency detection: A spectral residual approach," in *Proc. 2007 IEEE Conference on Computer Vision and Pattern Recognition*, 2007, pp. 1–8. doi:10.1109/CVPR.2007.383267

[2] APS run history accessed on 2022-07-30. https://ops.aps.anl.gov/statistics/

[3] G. Decker, "APS Storage Ring Commissioning and Early Operational Experience," in *Proc. PAC'95*, Dallas, TX, USA, May 1995, pp. 290–292.

[4] APS storage ring magnets accessed on 2022-07-30. https://www3.aps.anl.gov/APS_Engineering_Support_Division/Mechanical_Operations_and_Maintenance/Subsystems/Magnets/submagnet.html

[5] T. Fors, J. Carwardine, A. Hillman, J.F. Maclean, and S. Sarkar, "A Method of Measuring Noise and Detecting Glitches in Magnet Power Supplies at the APS Storage Ring," in *Proc. PAC'01*, Chicago, IL, USA, Jun. 2001. https://jacow.org/p01/papers/FPAH001.pdf

[6] G. Bradski, "The OpenCV Library," *Dr. Dobb's Journal of Software Tools*, 2000.

[7] H. Ren *et al.*, "Time-series anomaly detection service at microsoft," in *Proc. 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019, pp. 3009–3017. doi:10.1145/3292500.3330680

[8] Publication-ready NN-architecture schematics. https://alexlenail.me/NN-SVG/

[9] Martín Abadi *et al.*, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015. https://www.tensorflow.org/