ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-THPL005

LONG SHORT-TERM MEMORY NETWORKS FOR ANOMALY DETECTION IN STORAGE RING POWER SUPPLIES*

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

Abstract

We present an approach for detection of anomalous behavior of magnet power supplies (PSs) in storage rings, which may serve as an early indication of an impending PS trip. In this new method, we train a Long Short-Term Memory (LSTM) neural network to predict the temperature of several components of a PS (transistors, capacitors) based on the PS current, PS voltage, room temperature, and cooling water temperature. For training and testing, years of historical data are used from the Advanced Photon Source (APS). The neural network is trained on the data corresponding to the normal operation of the PSs. Anomalous behavior of a PS can be detected when the observed PS temperature starts to deviate significantly from the LSTM prediction. This may allow for preemptive action by the operators or PS group.

INTRODUCTION

There are 1320 power supplies (PSs) for the corrector, quadrupole, and sextupole magnets in the Advanced Photon Source (APS) storage ring [1, 2]. On average, PS trips caused about seven complete beam losses per year of operation [3], which resulted in unexpected interruptions of users' experiments. More than 15 years of historical data for PS parameters in APS can be used to explore various techniques for detection of trip precursors. Previously, we reported [4] on several unsupervised anomaly detection techniques, such as autoencoders [5–7] and saliency detection [8] for PS temperatures, currents, and voltages. These methods were successful in detecting precursors for some PS trips, but produced many false positives. Here we present a new method for anomaly detection in PS temperature data, based on the Long Short-Term Memory (LSTM) [9] neural networks. This new method is more readily interpreted, which allows assessing the severity of the detected anomalies and filtering out clear false positives.

METHOD DESCRIPTION

Historical data for APS power supplies are available at a 64 second interval from May 2008 to April 2023. In this contribution, we focus on PS currents and temperatures. Changes in PS current result in changes in PS temperature—see the blue points in Fig. 1 (observations)—which can be used to detect anomalies in PS temperatures. Namely, one can create a model that predicts the PS temperature based on recent values of the PS current. If the prediction deviates significantly from the observation, an anomaly is declared. Figure 1 and Table 1 compare three possible implementations of such predictive models. In this example, by PS temperature we mean the temperature of one of two transistors of a bipolar horizontal corrector PS. The ten-day training data interval included one intentional full-range PS ramp, followed immediately by a one-month-long test interval.

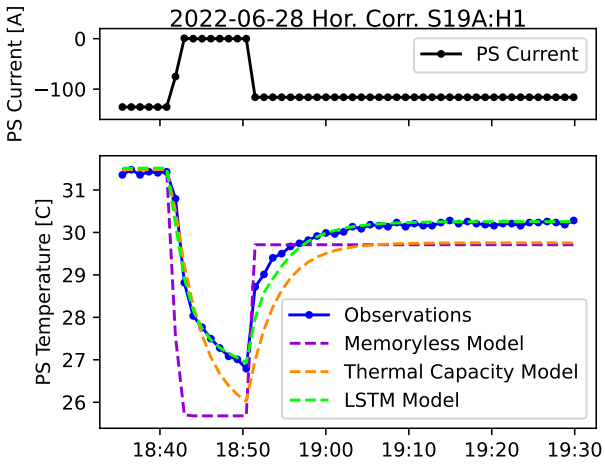

Figure 1: Prediction performance comparison. Short part of a one-month-long test interval.

Table 1: Mean Squared Error (MSE) comparison.

Model	Train MSE [C]	Test MSE [C]
Memoryless	3.12×10^{-2}	2.73×10^{-2}
Thermal Capacity	1.06×10^{-2}	1.18×10^{-2}
LSTM	6.98×10^{-3}	7.31×10^{-3}

First, we consider a polynomial **memoryless model**, which assumes that the PS is always in thermodynamic quasi-equilibrium, i.e., the PS temperature is a polynomial function of PS current at the same moment in time,

$$T(t) = \beta + \sum_{i=1,2} \alpha_1^{(i)} I_i(t) + \alpha_2^{(i)} I_i^2(t) + \alpha_3^{(i)} I_i^3(t), \quad (1)$$

where $i = 1, 2$ corresponds to the positive and negative currents of the bipolar corrector PS.

Figure 1 shows that the assumption of quasi-equilibrium does not hold. The memoryless model fails to adequately

* 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: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-THPL005

describe the regions with quick changes in PS current. Therefore, next, we considered a **thermal capacity** model,

$$C\frac{\mathrm{d}T}{\mathrm{d}t} = P_{\mathrm{in}} - P_{\mathrm{diss}}, \tag{2}$$

where C is the effective thermal capacity of the PS, P_{in} is the heating power generated inside the PS, and P_{diss} is the power dissipated to the surrounding environment and the cooling water. As an example, consider the following model for the dissipated power,

$$P_{\mathrm{diss}} = \alpha \times (T - T_{\mathrm{amb}}), \tag{3}$$

where α is a constant and T_{amb} is the ambient temperature. Eq. (2) can then be solved using Laplace transforms, giving

$$T(t) = T(0)e^{-\kappa t} + \int_0^t \kappa e^{-\kappa(t-t')}\left(\frac{1}{\alpha}P_{\mathrm{in}}(t') + T_{\mathrm{amb}}(t')\right)\mathrm{d}t', \tag{4}$$

where $\kappa = \alpha/C$. The heating power term can be approximated by a polynomial, similar to Eq. (1),

$$\frac{1}{\alpha}P_{\mathrm{in}}(t) = \sum_{i=1,2} a_1^{(i)} I_i(t) + a_2^{(i)} I_i^2(t) + a_3^{(i)} I_i^3(t). \tag{5}$$

The optimal value of κ, learned from the training data, was $\kappa = 4.7\times10^{-3}\,\mathrm{s}^{-1}$. One can see in Fig. 1 that the thermal capacity model accounts for the transitory effects and performs significantly better than the memoryless model. However, a PS cannot be accurately represented by a single effective thermal capacity C, since there is heat exchange among different components of the PS. Also, the temperature reading corresponds to only one point on the PS and the exact model for the dissipated power is unknown.

While one could continue to refine the thermal capacity model, we elected to use recurrent neural networks, namely, Long Short-Term Memory (LSTM) networks [9, 10]. Here the only input to the LSTM is the sequence of PS current values. The initial PS temperature is not used, because the system quickly forgets about initial conditions, as can be seen in the first term of Eq. (2). The chosen number of units in the LSTM cell is six and hyperbolic tangent activation is used for the cell output. As can be seen in Fig. 1 and Table 1, the **LSTM model** outperforms the thermal capacity model. The LSTM model is also more flexible, since its complexity can be changed by simply changing the number of parameters in the LSTM cell, as opposed to re-writing the assumptions of the underlying physical processes [Eqs. (2–5)].

Apart from the PS current, the PS temperature is also affected by the cooling water temperature, room temperature, and humidity. We have limited historical data for these parameters. When they are included in the LSTM model input, the prediction performance improves, see Fig. 2(c), as opposed to Fig. 2(b). Figure 2 presents predictions for the temperature of the capacitor of a corrector PS. The transistor temperature is less sensitive to changes in the environment. For this reason, hereafter, we focus on predicting PS transistor temperatures, based solely on PS currents.

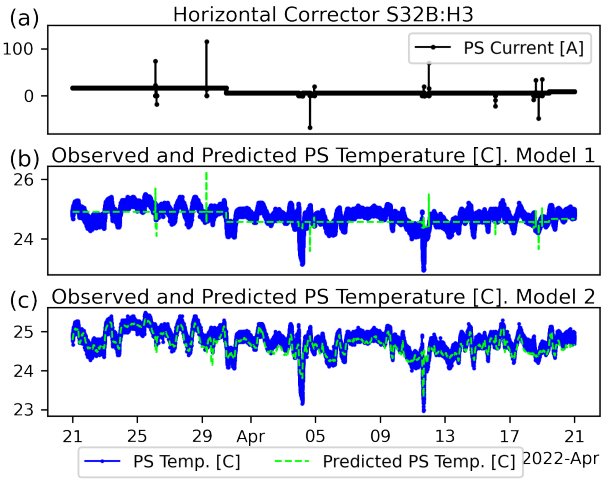

Figure 2: Prediction performance comparison. (a) PS current; (b) Model 1, based on PS current; (c) Model 2, based on PS current, room temperature, and cooling water temperature.

ANALYSIS OF ALL HISTORICAL DATA

We applied the LSTM-based approach to all the historical data from 2008 to 2023. The models for all 1320 PSs were retrained every other year, using a ten-day interval, which included intentional full-range PS ramps (performed regularly by the PS group). A part of one such interval is shown in Fig. 3(a) along with the LSTM-based temperature prediction (dashed lime line). During training, we used a weighted cost function, which emphasised regions with more activity in PS currents and temperatures. MinMax scaler [11] was chosen. The optimal parameters of the weight function and the optimal number of units in LSTM cells (six) were determined by running a hyperparameter optimization with RayTune [12]. For unipolar (quadrupole and sextupole) PSs, the input and output sequences were one-dimensional, while for bipolar corrector PSs, they were two-dimensional—positive and negative polarity currents in the input, temperatures of positive and negative polarity transistors in the output.

There were some outliers (defined as temperatures below 15 C or above 80 C) in the training data, as well as some missing data. If the fraction of outliers and missing data exceeded 5 %, the PS was deemed unsupported. If the fraction was less than 5 %, the outliers and missing data points were filled with linearly interpolated values. More than 90 % of the PSs were supported at any point in time. To review the predictions and the detected anomalies for all the 15 years of historical data, a graphical user interface application was created. Some examples of the detected anomalies are shown in Fig. 3(b)–(f) as the red shaded regions, where observations deviate from predictions by more than 3 C. In panels (a),(b), and (f) of Fig. 3, corresponding to bipolar PSs, only the temperatures of the positive polarity transistors are shown.

ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-THPL005

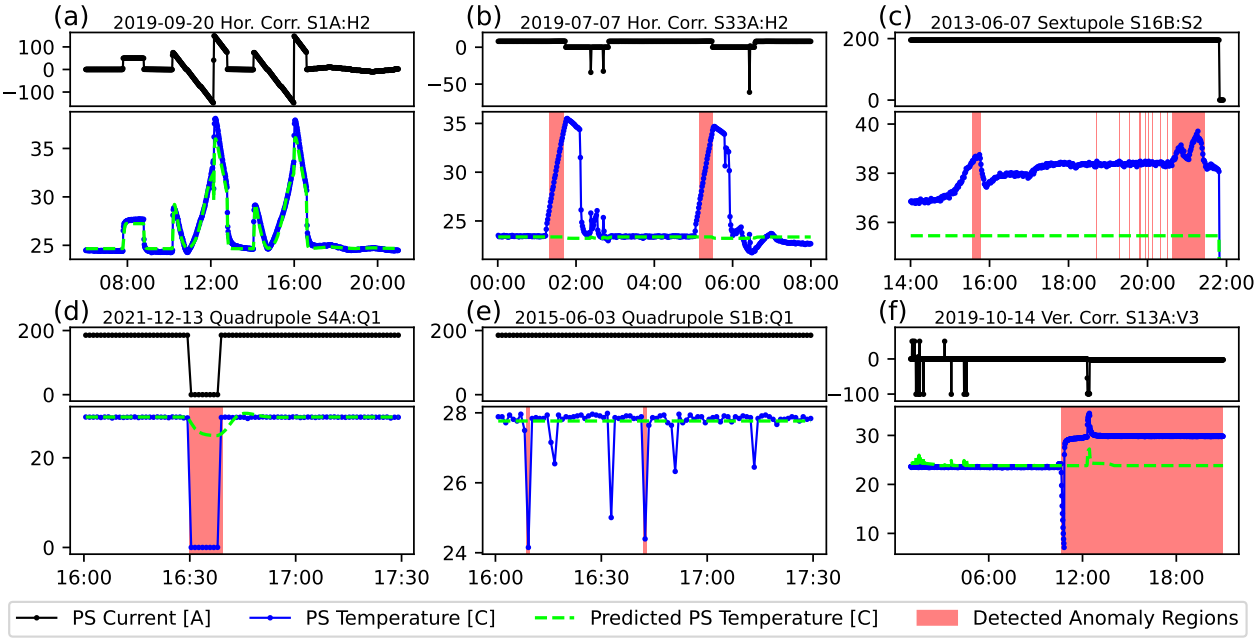

Figure 3: Examples of anomalies detected with the LSTM-based method in the historical APS data from 2008 to 2023.

DISCUSSION AND CONCLUSIONS

The example in Fig. 3(b) corresponds to a stuck mixing valve for the cooling water (two incidents). When cooling water stopped circulating, the temperatures of multiple PSs increased quickly. When one of them reached 50 C, it tripped (by design). Of course, other less sophisticated methods could detect this too, such as an alarm on a fixed temperature threshold. However, the LSTM-based method can provide a much earlier warning.

Figure 3(c) shows a peculiar anomaly in one PS, which tripped shortly after. It is noteworthy that this unusual behavior of the PS temperature would not be detected with a simple threshold method, for example, since the PS temperature was not particularly high. However, it did deviate from what it is expected to be at this PS current, and the LSTM-based method immediately detected it.

The anomalies detected in Fig. 3(d),(e), and (f) do not immediately lead to a PS trip. In Fig. 3(d), the temperature and current readings briefly went to zero, which could be a network problem. In Fig. 3(e), the anomalies flag a very noisy signal from the temperature sensor. Perhaps the contact between the sensor and the surface was poor. The anomaly type shown in Fig. 3(f), may be observed when a water pump is reset and the water flow encounters a step change, or when a PS is replaced by a new one during the run. It is possible to retrain the LSTM model whenever PSs are swapped out. However, the PS maintenance records in APS do not appear to be reliable enough, and we cannot try this approach on the historical data. Still, it is possible to implement this retraining strategy in the future in the APS Upgrade [13].

Even though some kinds of detected anomalies [Fig. 3(b),(c)] led to PS trips within several minutes or hours, the majority of the anomalies [Fig. 3(d)–(f)] can be on-going for months without causing a trip. However, they still remain true anomalies. According to the PS group experts, any unexpected deviation beyond 3–4 C must be reported. This makes it difficult to quantitatively summarize the method's performance on the entire historical data set. In the future, it may be possible to train a classifier to automatically assign a priority to the detected anomalies, so that Fig. 3(b),(c) would receive higher priority than Fig. 3(d),(e),(f). Perhaps a more important metric is by how much the proposed anomaly detection method could improve the cost efficiency. In this regard, it may be acceptable to use a method with low precision of trip prediction but high recall [14], because the cost of checking a PS is about 1/300 the cost of an hour of beam time.

The historical data collected during the operation of APS provided a unique opportunity to test anomaly detection techniques for the storage ring PSs. The lessons learned will help shape the methods used for the APS Upgrade, which promise to be more accurate thanks to higher data rates and better monitoring of the parameters of the PSs, cooling systems, and the tunnel.

ACKNOWLEDGMENTS

The authors thank L. Emery, J. Wang, G. Fystro, 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: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-THPL005

REFERENCES

[1] *APS storage ring magnets*, https://www3.aps.anl.gov/APS_Engineering_Support_Division/Mechanical_Operations_and_Maintenance/Subsystems/Magnets/submagnet.html, Accessed: 2022-07-30.

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

[3] *APS run history*, https://ops.aps.anl.gov/statistics/, Accessed: 2022-07-30.

[4] I. Lobach *et al.*, "Machine Learning for Predicting Power Supply Trips in Storage Rings," in *Proc. 5th Int. Particle Accel. Conf. (NAPAC'22)*, Albuquerque, NM, USA, 2022, paper TUPA29, pp. 413–416, doi:10.18429/JACoW-NAPAC2022-TUPA29

[5] J. Edelen and N. Cook, "Anomaly detection in particle accelerator using autoencoders," in *Proceedings of the 2021 Improving Scientific Software Conference*, 2021, doi:10.26024/p6mv-en77

[6] H. Xu *et al.*, "Unsupervised anomaly detection via variational auto-encoder for seasonal kpis in web applications," in *Proceedings of the 2018 World Wide Web Conference*, 2018, pp. 187–196, doi:10.1145/3178876.3185996

[7] I. Higgins *et al.*, "Beta-VAE: Learning basic visual concepts with a constrained variational framework," in *International Conference on Learning Representations*, 2017, https://openreview.net/forum?id=Sy2fzU9gl

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

[9] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997, doi:10.1162/neco.1997.9.8.1735

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

[11] F. Pedregosa *et al.*, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[12] R. Liaw, E. Liang, R. Nishihara, P. Moritz, J. E. Gonzalez, and I. Stoica, "Tune: A research platform for distributed model selection and training," *arXiv preprint arXiv:1807.05118*, 2018.

[13] M. Borland *et al.*, "The Upgrade of the Advanced Photon Source," in *Proc. IPAC'18*, Vancouver, Canada, Apr.-May 2018, pp. 2872–2877, doi:10.18429/JACoW-IPAC2018-THXGBD1

[14] K. M. Ting, "Precision and recall," in *Encyclopedia of Machine Learning*. 2010, pp. 781–781, doi:10.1007/978-0-387-30164-8_652