

ANOMALY DETECTION IN ACCELERATOR FACILITIES USING MACHINE LEARNING

A. Das*, D. Ratner

SLAC National Accelerator Laboratory, California, USA

M. Borland, L. Emery, X. Huang, H. Shang, G. Shen

Argonne National Laboratory, Illinois, USA

R. Smith, G. Wang†

Brookhaven National Laboratory, New York, USA

Abstract

Light source facilities usually operate about 5000 hours per year to support multiple beamline operations. Reliability is a key parameter in such user facilities to evaluate machine performance. Some facilities have achieved more than 95% beam reliability. However, there are still many hours of unplanned beam downtime and every hour lost is a waste of operational costs. Beam downtime also interrupts the completion of scheduled user experiments. Preventive maintenance of subsystems and quick recovery from downtimes are the basic strategies to improve reliability. Current recovery incorporates significant human diagnosis efforts. To circumvent this problem of unprecedented downtimes requiring recuperation, we take steps to build solutions that can detect anomalous conditions caused by faulty subsystems. In this paper, we share our findings from an initial assessment of production logs and provide an overview of some potential future directions.

BACKGROUND

Accelerator facilities experience beam degradation or complete beam dumps due to subsystem faults such as RF (Radio-Frequency) system trips, power supply faults etc. Degradation implies poor beam performance not leading to downtime, while beam dump indicates machine downtime. Besides wastage of operational costs associated with downtimes, beam dumps also disrupt scheduled user experiments. The Mean Time To Recover (MTTR) from a failure usually ranges between one to a few hours. For example, in 2018, the MTTR for NSLS-II was 1.5 hours. The MTTR includes the time needed in diagnosing failure causes, repairing the faulty devices, and setting the system back to the standard operational state. Efficient recovery requires better understanding of the system health based on the archived and real-time measurements, in a timely manner.

A light source facility may archive more than 100,000 measured signals (referred to as process variables, or PVs). Currently, fault diagnosis is done manually by the system managers or operation specialists as they analyze the operation history data prior to the observed faults. This can be time-consuming, as the hints to the root cause can be buried amongst hundreds of PVs related to the affected subsystem

and it demands the subsystem specialists to be available for diagnosis any time such faults appear. Efficient automation of failure detection can greatly reduce the recovery time and ease the burden on professional staff. We seek to leverage this high dimensional parameter space of irregular time-series data for anomaly detection using principles of statistics and machine learning (ML). Data labeling is cumbersome as found in recent studies [1]. Moreover, supervised methods are becoming less applicable in facilities with limited or no ground truth [2]. Keeping in mind the need for both efficiency and correctness, we analyze PVs associated with a few subsystems and describe our observations. The derived insights can be helpful in designing unsupervised methods for fault detection.

FAULT ANALYSIS

We identify days with reported subsystem or device faults and downtimes. Days with no reported faults or downtimes are considered healthy time-frames. In the discussed plots, x-axis corresponds to the timesteps and y-axis to the signal observations, unless otherwise mentioned.

For a specific power supply, Figs. 1 and 2 show the trends of two PVs namely, magnet temperature and output voltage. As seen in Fig. 1, before the power supply tripped there is a drop in magnet temperature. The variation in voltage is comparatively less distinguishable between anomalous and healthy times. We observe similar difference in PV variations for other subsystems as well. Several PVs of a specific device may not exhibit helpful trends to aid anomaly detection. It is imperative to choose a suitable set of PVs indicative of anomaly, and ensure that helpful signal patterns do not get obscured during any multivariate analysis performed to represent device state.





Figure 1: P-Supply (fault).



We analyze water system cracks that do not immediately affect beam performance, but eventually cause beam dump

* anwesha@slac.stanford.edu

† gwang@bnl.gov







Figure 3: W-Supply (FM1). Figure 4: W-Supply (FM1). Figure 5: W-Supply (FM4). Figure 6: W-Supply (FM4).




Figure 7: Current Dip.




Figure 8: Downtime.




Figure 9: Steady Current.

due to monitoring by the Equipment Protection System (EPS). Pump leaks are eventually fixed by operators for healthy machine performance. These water cells with multiple flow-meters are part of the EPS, all of which do not show similar variations during healthy and faulty times.

Figures 3 to 6 show two signals of a specific water system that faulted leading to a beam dump. Figures 3 and 4 relate to flow-meter FM1 that did not deviate much around the faulty time compared to its healthy counterpart. However, flow-meter FM4 reveals dips and spikes (Fig. 5) around the faulty time-window relative to its normal signal behaviour (Fig. 6) indicating anomalous state that could eventually hurt the EPS causing downtime. The overall range of FM1 PV values (≈ 1) is generally lower than FM4 (≈ 3) for different subsystems. We notice that, besides identifying the anomaly indicative PVs, the statistical metric of variation can be different in diverse subsystems or its related signals. This means, while higher temporal variance can be anomaly indicative in a specific PV, a few drops to zeroes or spikes to higher values without significant difference in variance can relate to fault in another PV. These fault-based signatures are inferred based on the observed healthy trends of the corresponding PVs. ML-based model require data that do not result in too many false positives or false negatives. In our examples, signals such as output voltage of power supply or FM1 of water system may lead to detection errors, for which subsystem-level PV analysis can be helpful before building a generic fault detector.

Beam Performance

The signals used for beam quality assessment are broadly similar but not entirely the same in different facilities. Some subsystem faults show tangible impacts on the beam, others do not. Beam current is often monitored in storage rings to examine anomalous beam dumps. We study the correlation of beam current with a few studied subsystem faults, not analyzed in some prior studies [3].

Figure 7 shows a dip in beam current. This dip correlates with a downtime that lasted over 3 hours, related to a faulty water system. In this case, the faulty subsystem lead to a complete beam dump, however there is no early indication of anomaly in the beam current before downtime, to aid preventive maintenance. Figure 8 shows another case of downtime that lasted for several hours comprising of downtime and recovery related to a pump leak. The beam current sharply drops to zero and stays there for a long time. *Current* usually remains steady or drops to zero, unlike energy or intensity-related signals that show gradual variations indicating degraded beam quality. Hence, in the absence of signals having temporal variations (e.g., energy, frequency) subsystem-related signals showing prior trends before fault manifestation is necessary for prediction. Faults not leading to downtimes is expected to show weak correlation with beam quality. In the absence of sustained temporal fluctuations (e.g., temperature in Fig. 1) forecasting is difficult, as opposed to classification [4] or post-mortem analysis. Beam current may not be very helpful in that case. Figure 9 shows beam current values during normal machine operation with no reported faults or downtimes. In this case, we expect current to be mostly steady with minor deviations.

Both signal correlations within a subsystem and their relationship with the beam quality can be considered for fault detection. A signal when examined with other signals can become a differentiator at times, in contrast to independent consideration. For example, we notice that the co-existence of a jittery signal related to current frequency with sporadic degradation of beam intensity can indicate a faulty station. Predicting such faulty stations not causing downtimes yet, can be helpful for preventive maintenance. However, if such jitters do not correlate with beam intensity most of the times or lead to faults, the related short-lived jitters can be considered non-critical.

Signal Combination

We identify certain RF system faults that show variation in measurements prior to trips. Each subsystem has a different set of signals associated with it. In order to form a single representative signal for a specific device that can help in distinguishing anomalous and healthy times, we apply the copula-based approach [5] from statistics. The idea is to first learn the most suitable univariate distribution of each considered signal, from which its marginal distribution is then estimated. A copula function then learns the multivariate joint distribution (i.e., CDF; cumulative distribution function) based on correlations between individual marginal distributions of each signal. From the obtained joint distribution, new data can be estimated. This many to one mapped time-series can then be used for subsystem fault modeling.

Let us consider an RF system A , with just 5 PVs from its set of signals. For a 5-dimensional joint CDF F , we have:

$$F(PV_1, \dots, PV_5) = C[F_1(PV_1), \dots, F_5(PV_5)]. \quad (1)$$

A copula function C exists [Eq.(1)] for univariate marginals F_1 to F_5 related to 5 PVs. If F is a multivariate normal distribution, it is called a Gaussian copula.




Figure 10: Fault (5 PVs).




Figure 11: No Fault (5 PVs).

Figures 10 and 11 show the difference in CDF estimates during faulty and healthy times. In the plots, x-axis indicates the obtained values after joint modeling of univariates, and y-axis represents the frequency or number of observations corresponding to the values on x-axis, for device A . *Real* implies original data, and *Synthetic* implies the sampled data estimated from the derived joint distribution. The difference in distribution is discernible with lower range of values

(≈ 0.16 vs. ≈ 0.8) during healthy times (Fig. 11) and higher peaks observed during faulty times (Fig. 10). For some cases, the vice-versa is observed as well, i.e., healthy times with higher peaks over faulty times. Both cases can aid anomaly detection as long as healthy and faulty times are recognizable. This is one possible way of handling subsystem signals prior to training a suitable ML-model.

Well known approaches like PCA (principal component analysis) [6] also approximate the behaviour using a multivariate joint distribution for linear correlations. Copulas have been shown to work for certain non-linear correlations as well [5]. However, it is possible to lose signal information post joint modeling based on the scale and irregularity of time-series. We are currently examining univariate models and the potential limitations of joint modeling by analyzing more subsystems from different facilities. Future work comprises of devising ML-models for training and inference based on the transformed subsystem signals.

CONCLUSION

We present some preliminary findings from our studies on subsystem faults. Univariate trends between anomalous and healthy times can be similar, with no early fault indications. A suitable subset of signals pertaining to a specific subsystem needs to be identified for low detection errors. The distinct pattern or statistical metric indicating degrading health can vary across subsystems of same type (e.g., power supply Q vs. S) or different types (e.g., water system vs. power supply). A generic fault detector needs to account for this diversity. Multivariate correlations with their relationship with beam quality has the potential to reveal faulty subsystems. Copula-based approach works to a certain extent in providing a unified view of subsystem state in identifying anomalous conditions. Further studies are necessary for determining suitable methods that can work robustly for ML-based detection. Power supply trips causing downtimes may not have temporal patterns ahead of time to aid prediction. RF system drifts exhibiting beam instabilities has the potential to provide early indicators for fault prediction. The feasibility of prediction depends on the nature of subsystems and fault manifestation, subject to future work.

ACKNOWLEDGEMENTS

This work was supported by the Department of Energy, Laboratory Directed Research and Development program at SLAC National Accelerator Laboratory, under contract DE-AC02-76SF00515; Project FWP#: BES-ML 34573.

REFERENCES

- [1] C. Tennant *et al.*, “Superconducting radio-frequency cavity fault classification using machine learning at jefferson laboratory”, *Phys. Rev. Accel. Beams*, vol. 23, no. 11, p. 114601, Nov. 2020. doi:10.1103/physrevaccelbeams.23.114601
- [2] E. Fol, R. Tomás, J. Coello de Portugal, and G. Franchetti, “Detection of faulty beam position monitors using unsupervised learning”, *Phys. Rev. Accel. Beams*, vol. 23, no. 10,

- p. 102805, Oct. 2020. doi:10.1103/physrevaccelbeams.23.102805
- [3] B. Nachman and D. Shih, “Anomaly detection with density estimation”, *Phys. Rev. D*, vol. 101, no. 7, p. 075042, Apr. 2020. doi:10.1103/physrevd.101.075042
- [4] G. Valentino and B. Salvachua, “Machine learning applied at the LHC for beam loss pattern classification”, *J. Phys. Conf. Ser.*, vol. 1067, no. 7, p. 072036, Sep. 2018. doi:10.1088/1742-6596/1067/7/072036
- [5] D. Salinas, M. Bohlke-Schneider, L. Callot, R. Medico, and J. Gasthaus, “High-dimensional multivariate forecasting with low-rank gaussian copula processes”, *NeurIPS*, pp. 6824–6834, 2019.
- [6] J. Lever, M. Krzywinski, and N. Altman, “Points of significance: Principal component analysis”, *Nat. Methods*, vol. 14, pp. 641–642, Jun. 2017. doi:10.1038/nmeth.4346