

AUTOMATIC CLASSIFICATION OF POST MORTEM DATA FOR REDUCED BEAM DOWN TIME

M. C. Chalmers*, Y. E. Tan, ANSTO - Australian Synchrotron, Clayton, Australia

Abstract

Time spent recovering from faults that result in a rapid loss of stored current (total loss of beam) can be costly to the Australian Synchrotron facility and its researchers. The identification of a fault leading to total beam loss is assisted by a large variety of investigative tools for specific tasks, but they do not often give a thorough overview of all systems required to store beam. Post mortem data uniquely provides insight into how the beam was behaving at the specific time the beam loss occurred. With machine learning, we find that we can automatically and rapidly identify many types of total beam loss events by learning about the unique characteristics in the post mortem data.

INTRODUCTION

Each year the Australian Synchrotron loses up to 50hrs of scheduled beamline operations due to accelerator faults [1]. The most common faults which result in total beam loss occur due to magnet power supply failures, RF and cooling water (plant) faults (see Fig. 1). The post mortem data often carries a signature pattern depending on which of these faults caused the beam loss.

The purpose of this investigation is to begin assessing the viability of a diagnostic tool, which will automatically attempt to classify every total beam loss event. It's envisioned that doing so will reduce the amount of time searching through diagnostic tools in order to identify faults, in turn leading to faster recovery times, and reduced down time.

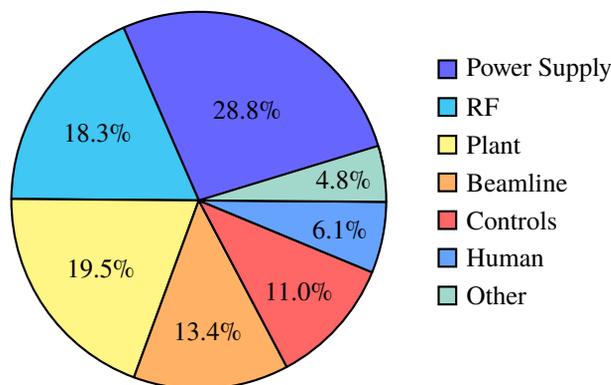


Figure 1: Frequency of beam downtime categorised by the system which caused the outage during beamline operations.

Post Mortem Data

Many facilities capture and characterise post mortem information, it is common to review this data to assess complex to-

* madeleine.chalmers@ansto.gov.au

tal beam loss events [2]. The post mortem information used here is collected from 98 beam position monitors around the storage ring, capturing 16-20,000 samples of turn-by-turn data [3]. The post mortem data cache is frozen and saved when the stored beam position drifts outside a defined limit, and activates the Equipment Protection System (EPS). The most common failures have visually recognisable characteristics. Figures 2 & 3 show examples of magnet power supply failures and Fig. 4 shows when the RF is inhibited by the EPS.

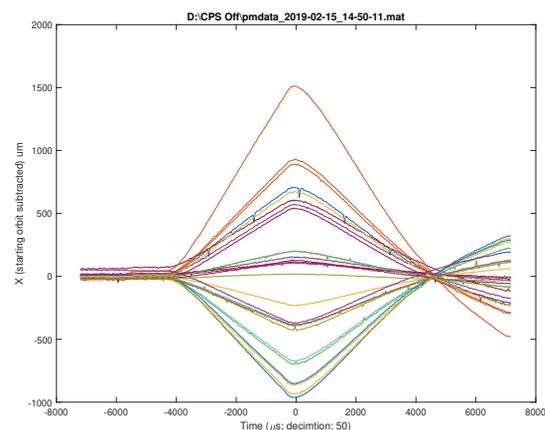


Figure 2: Corrector fault example.

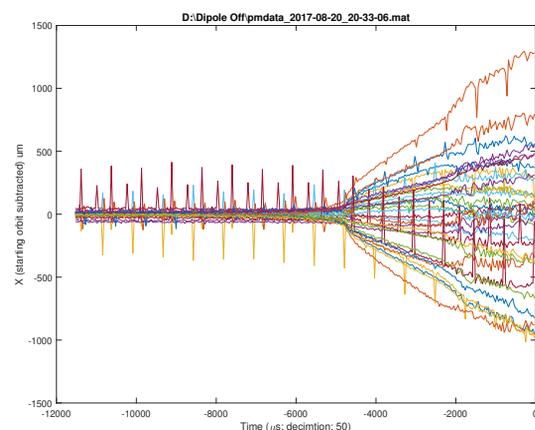


Figure 3: Dipole fault example.

LEARNING METHOD

Data used in this study was manually classified by reviewing operations logs in order to attribute a fault to each beam loss event [4]. Logistic regression was applied to the data

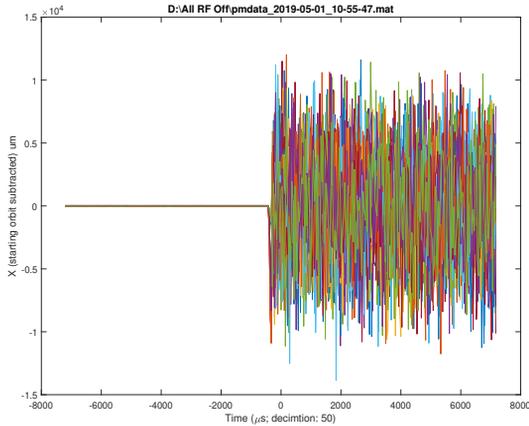


Figure 4: RF inhibit example.

using gradient descent to minimise the cost function and optimise the model [5]. The logistic function is:

$$J(\theta) = \frac{1}{m} \text{Cost}(h_{\theta}(x^{(i)}), y^{(i)}) \quad (1)$$

$$= \frac{1}{m} \left[\sum_{i=1}^m y^{(i)} \log h_{\theta}(x^{(i)}) + (1 - y^{(i)}) \log(1 - h_{\theta}(x^{(i)})) \right]$$

$$h_{\theta}(x) = \frac{1}{1 + e^{-\theta^T x}} \quad (2)$$

The function for gradient descent with regularisation is:

$$\theta_0 := \theta_0 - \alpha \sum_{i=1}^m (h_{\theta}(x^{(i)}) - y^{(i)}) x_0^{(i)} \quad (3)$$

$$\theta_j := \theta_j - \alpha \left[\sum_{i=1}^m (h_{\theta}(x^{(i)}) - y^{(i)}) x_j^{(i)} + \frac{\lambda}{m} \theta_j \right] \quad (4)$$

Generating Learning Metrics

In all models, the x and y beam position as well as sum current signals were used to characterise the beam at regular intervals in time, to generate features for the algorithm to learn from. Files were divided into 16 equal time segments to characterise the beam loss evolution, and the noise levels were quantified by calculating the root mean square for all signals in all segments.

Other features were generated by sorting the data in each segment by amplitude. The average of the first 10 minimum data points and 10 maximum data points in each segment were calculated. This effectively places an upper and lower threshold on the beam position during each time segment.

RESULTS

Three groups of fault comparisons were analysed in order to test the classifiability of post mortem data, as follows:

RF Inhibit vs Noise

The RF inhibit examples were easily distinguishable from the noise produced in the beams absence. The model achieved 100% accuracy after training on a single example of each fault and being tested on 20 independent events.

RF Inhibit vs All Faults

This model was trained on up to eight examples of each fault (16 training examples in total), and tested on four examples of each fault (eight examples in total), to determine the optimal number of examples to train on for maximum accuracy. When comparing all faults versus RF Inhibit events, the model performed less accurately than before (Table 1). This is likely due to the fact that the training data included examples with beam instabilities which resulted in rapid beam loss which was similar in appearance to an RF inhibit. Secondly, the entire collection of post mortem files is remarkably diverse and the model likely had difficulty characterising the non RF inhibit data, especially with such a small training dataset. This should be addressed by increasing the number of training and test datasets available, and also increasing the number fault categories along with better features to differentiate between them with.

Table 1: Models were created for each run to determine the optimal number of examples to train on. Accuracy could not be improved past 87.5% (Run 2) due to limited testing and training examples.

Run	Test Accuracy	Training Examples
Run 1	37.50%	8
Run 2	87.50%	12
Run 3	87.50%	16

RF Inhibit vs Magnet (CPS & QPS) Faults vs Noise

In this section, up to six training examples of each fault type were used (24 training examples in total) and four test examples of each (16 test examples in total), to determine the optimal number of examples to train on for maximum accuracy. When restricting the model to include only four common event types, the algorithm again demonstrated some success (Table 2). During each run the number of examples per fault category were increased one at a time, which improved the models accuracy at first. Improvements were not able to be made past Run 3, again due to the small number of training and test examples available.

Later trials involved training on data from on all sectors instead of just the first, this was done with one fault category at a time so that the effect of each one could be quantified. The same total beam loss events were used for training and testing as before so that the effects could be measured relative to the previous results (Table 2). Applying this method to the corrector power supplies (CPS) was highly successful

Table 2: Models were created for each run to determine the optimal number of examples to train on. Accuracy was limited to 68.75% after training on three examples.

Run	Test Accuracy	Training Examples
Run 1	62.50%	4
Run 2	62.50%	8
Run 3	68.75%	16
Run 4	68.75%	24

at increasing the prediction accuracy (Table 3). The opposite occurred when the model was trained on all sectors of the quadrupole (QPS) failure data, this appears to be because the data for quadrupole failures is highly inconsistent. Splitting the Noise and RF Inhibit data into sectors provided no accuracy increases.

Table 3: Models were created to determine which faults can be modeled more accurately by using data from all sectors.

Run	Test Accuracy	Improvement
RF	68.75%	0%
QPS	62.50%	-6.25%
Noise	68.75%	0%
CPS	87.50%	+18.75%

LIMITATIONS

This study was limited by the fact that the data was exhumed from the archive, making it susceptible to documentation errors that would result in incorrect classification during training. An upgrade in BPM hardware in 2017 has also meant that without extensive pre-processing, the data is not always comparable for the same event type before and after the upgrade, this resulted in less data being available to use in the study.

FUTURE IMPROVEMENTS

The method of dividing each post mortem example into 16 time segments is crude and imprecise, future studies will investigate improved statistical measurement for better features. A sliding window technique could be used to identify the specific timing of thresholds being reached. For example the length of the decay period could be identified by searching the post mortem data for when the beam first becomes unstable and when the beam current first reaches zero. It's

likely that using the beam dispersion and offset to estimate the beam energy could lead to accuracy improvements as well.

Data produced specifically for this purpose could improve the precision of the models created and remove any errors introduced through the documentation errors already mentioned. Introducing new realtime variables could also extend the functionality of this tool. With the right variables it would be possible to assess the beam quality and provide pre-emptive warning of a potential loss of beam. It will enable early fault intervention and possibly even prevention of total beam loss events. Notable issues where real-time intervention may be possible include those which evolve over a timeframe in the order of minutes or greater, this could typically include feedback system malfunctions and climate control faults.

CONCLUSION

When given specific and well-defined events to classify, the models developed, while limited in scope were effective at classifying post mortem data. With the right training data and distinct faults to categorise, it was possible to predict the cause of a beam loss event with 87% accuracy. Metrics used to generate features could be improved for situations where different faults produce very similar data sets and when identical faults produce highly variable data sets. More work is needed to improve the amount data available for training and testing models.

REFERENCES

- [1] J. Trehwella, D. Morris, G. S. LeBlanc, D. McGilvery "Maximising Beam Availability At The Australian Synchrotron", in *Proc. 25th Particle Accelerator Conf. (PAC'10)*, Kyoto, Japan, May 2010, paper WEPEA002, pp. 2469–2471.
- [2] G. M. Wang *et al.*, "NLS-II Post Mortem Function Development and Data Analysis of Beam Dump", in *Proc. North American Particle Accelerator Conf. (NAPAC'16)*, Chicago, IL, USA, Oct. 2016, pp. 1039–1041. doi:10.18429/JACoW-NAPAC2016-WEPOB66
- [3] Y. E. Tan and R. B. Hogan, "Australian Synchrotron BPM Electronics Upgrade", in *Proc. 7th International Beam Instrumentation Conference (IBIC'18)*, Shanghai, China, Sep. 2018, paper TUPC01, pp. 297.
- [4] Don McGilvery, Greg LeBlanc "Operator Roles At The Australian Synchrotron", in *Proc. 6th International Particle Accelerator Conf. (IPAC'15)*, Richmond, VA, USA, May 2015, paper TUPWA004, pp. 1397–1398.
- [5] A. Ng "CS229 Lecture notes", <http://cs229.stanford.edu/notes/cs229-notes1.pdf>