ISBN: 978-3-95450-184-7 doi:10.18429/JACoW-IPAC2018-WEPAF022

APPLICATION OF MACHINE LEARNING TO MINIMIZE LONG TERM DRIFTS IN THE NSLS-II LINAC*

R. P. Fliller III†, C. Gardner, P. Marino, R. Rainer, M. Santana, G. Weiner, E. Zeitler, X. Yang, Brookhaven National Laboratory, Upton NY, USA

Abstract

Machine Learning has proven itself as a useful technique in a variety of applications from image recognition to playing Go. Artificial Neural Networks have certain advantages when used as a feedforward system, such as how the predicted correction relies on a model built from data. When applied to a particle accelerator, this allows for the Artificial Neural Network to compensate for effects that are difficult to model such as low level RF adjustments to compensate for long term drifts. The NSLS-II linac suffers from long terms drifts from a number of sources including thermal drifts and klystron gain variations. These drifts have an effect on the injection efficiency into the booster and if left unchecked, portions of the bunch train may not be injected into the booster, and the storage ring bunch pattern will ultimately suffer. This paper covers the application of Artificial Neural Networks to compensate for long term drifts in the NSLS-II linear accelerator. The Artificial Neural Network is implemented in python allowing for rapid development of the network. The design and training of the network, along with results of using the network in operation are also discussed.

INTRODUCTION

The NSLS-II is a 3 GeV electron storage ring with a full energy injector consisting of a 200 MeV linac and a 3 GeV booster. The NSLS-II linac was commissioned in 2011 and has been described in other publications [1]. In short, the linac consists of the electron gun, two standing wave buncher cavities, and 5 travelling wave cavities. The Linac has three klystrons available for power. Two are needed for operation, with one as a ``hot spare". For the purposes of this paper there are two important components to consider. The first is the subharmonic prebuncher (SPB), which provides the initial bunching. The second is the klystron which powers the high energy portion of the Linac. As there are three klystrons available to power the linac, and only two are used at any time, the klystron powering the low energy sections will be referred to as klystron A, and the klystron powering the high energy section will be reference to as klystron B.

Top Off operation at the NSLS-II began in 2015 requiring injections once approximately every two minutes [2]. With the advent of top off, it was noted that the Linac to Booster transport line (LTB) efficiency gradually deteriorates over time. This loss results not only in the loss of charge in the injector, but also a change in the bunch train shape from the injector. This can occur if the energy spread of the train is correlated along the train and bunches at the ends of the train are outside of the energy acceptance to the LTB. Once the top off program senses that not enough charge is injected, it will start requesting more charge from the injector. This can lead to further losses and a runaway situation. At some point the top off program will not be able to maintain the storage ring current and will issue an alarm.

The root cause is drifts in the linac RF system, particularly power drifts in the two linac klystrons. This leads to changes in the beam energy and the energy spread. Power drifts in klystron A affect the beam energy and the energy spread, while power drifts in klystron B affect only the energy. During operations, it was found that adjusting the SPB phase is the easiest way to correct for errors in the bunching and energy spread, and the klystron B power output is adjusted to correct for energy errors.

In order to provide some online means to correct for these drifts, machine learning techniques were used to implement a feedforward correction to the SPB phase and the klystron B power [3]. This approach allows for implementation of a model of the machine without recourse to an offline model. Proper modelling of this requires tracking in the linac to accurately model the low energy section of the linac where the bunching occurs, and accurate knowledge of all of the linac, LTB, and booster apertures. The Artificial Neural Network (NNet) developed for this application tracks the beam motion in the LTB and the first turns in the Booster to optimize the charge transport efficiency of the injector. The NNet is compared to a response matrix and high order polynomial fit.

DATA COLLECTION

The SPB phase and the klystron B amplitude are the parameters of choice to correct for drifts in the linac. The response of the beam to these adjustments must therefore be measured. The beam position on all beam position monitors (BPMs) in the linac to booster transport line and the booster, bunch charges at various diagnostics, linac RF settings, and some magnet settings were recorded for a variety of different bunch charges. Data was collected over a 22 degree range of the SPB phase and a 6 MW change in the klystron power which is beyond the typical drift range on the linac.

The first dataset consisted of 2700 points collected over two shifts on two separate days using the third klystron as klystron B. A neural network and the response matrix were developed from this dataset. During the course of these studies, the conditions necessitated using the second

* This manuscript has been authored by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy
† rfliller@bnl.gov

ISBN: 978-3-95450-184-7 doi:10.18429/JACoW-IPAC2018-WEPAF022

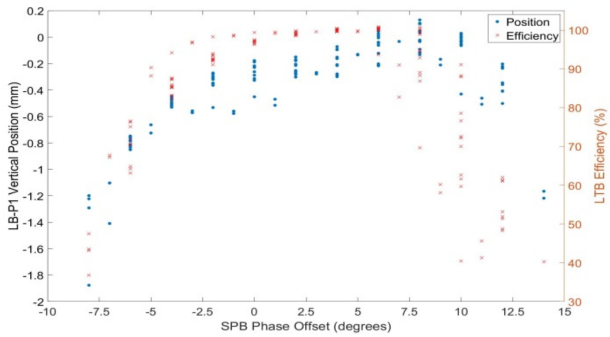

Figure 1: LB-P1 position and LTB efficiency vs the SPB phase for the klystron power unchanged.

klystron as klystron B. Therefore a second dataset consisting of 5400 points collected on two consecutive shifts. A second neural network and the polynomial functions were developed from this dataset. As the two datasets used different klystrons it is not possible to merge them into a larger dataset.

Based on operational experience and the data, it was determined that the input parameters would be the LTB transport efficiency, the mean difference in horizontal position between turns 2 and 20 in the booster, and the vertical position at the first BPM in the LTB, named LB-P1. These data can be used to tune the SPB phase and the klystron amplitude. Figure 1 shows the variation of the LB-P1 vertical position and LTB efficiency vs. the SPB phase for constant klystron power.

MODEL DEVELOPMENT

Two neural networks were developed, one for each dataset. Each neural network consisted of three layers, an input layer of three nodes, a hidden layer, and a two node output layer. The activation function was the rectified linear unit for the input and hidden layer, and linear for the output layer. The first neural network, based on the 2700 point dataset had 19 nodes in the hidden layer. During training of this network, dropout was used, cutting a random 20% of the connections each training epoch. The second neural network utilized 18 nodes in the hidden layer and no dropout during training. For purposes of this paper the two neural nets were named NNet2700 and NNet5400 respectively, to account for the number of points in the applicable dataset. In each case 75% of the dataset was used to train the neural network with the remaining 25 % to validate the fit to ensure there was no overfitting.

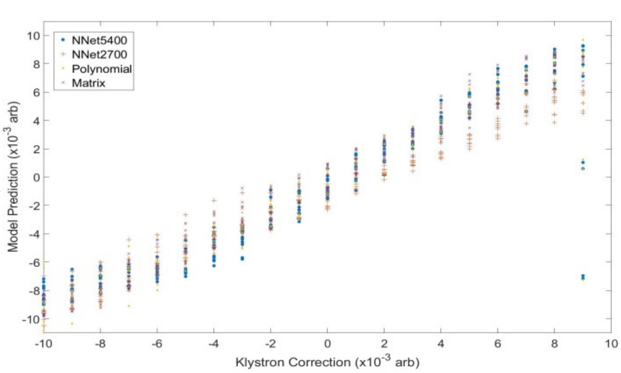

Figure 2: Klystron responses vs. offsets for each of the models.

Figure 2 shows a fit of the model outputs vs required klystron correction from the training data for the two neural networks. The required SPB phase correction is zero in the graph.

Figure 3 shows the SPB phase correction predictions from the models. It shows that the polynomial predicts the correction well, as the R matrix fails outside of the range to which it was fit.

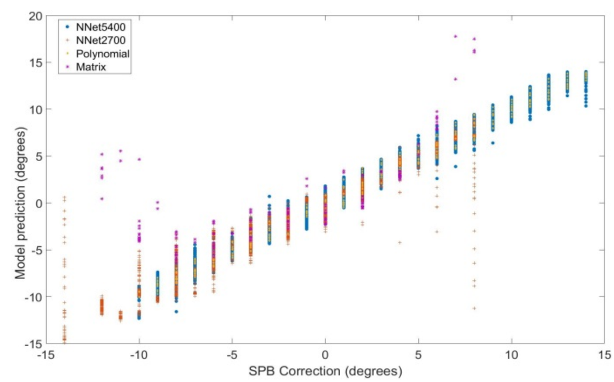

Figure 3: Matrix and Polynomial SPB responses vs. offsets. The matrix only used data between -5 and 5 degrees.

The 2x3 response matrix, R, was fit to a subset of the first dataset. Figure 1 shows the response on LB-P1 is nonlinear, and therefore only the range between ±5° of SPB phase is used in the fitting. Again, 75% of the dataset was used to fit the response matrix.

A polynomial was also used to determine each of the corrections for the second set. The polynomial for the SPB phase was fifth order in LB-P1 position, fifth order in LTB efficiency, and second order in mean difference in booster turns 2 and 20. The polynomial for the klystron correction was seventh order in the LB-P1 position and LTB efficiency, and 5^{th} order in the mean difference between turns 2 and 20 in the booster. The order of these polynominals was determined empirically by increasing the order until the accuracy of the fit was no longer improved.

MODEL OPERATION

NNet2700 and the R matrix were used during the second half of 2017. Figure 4 shows an example of NNet2700 at a time when there was testing of the linac's spare klystron during top off operations in the ring. Note there are two dips in the SPB phase in the upper graph. These dips occurred during the klystron testing at which time some unknown parameter changed the linac. NNet2700 was able to measure the beam motion and correct for the change. This can be seen on the lower graph where the beam position and efficiency remains relatively unchanged during this time.

The R matrix was used on several occasions to correct the linac during this time period as well. There were

ISBN: 978-3-95450-184-7 doi:10.18429/JACoW-IPAC2018-WEPAF022

several issues with the R matrix. The first is that the R matrix would consistently shift the beam 100 μm, with a corresponding drop of the overall injector efficiency from 75% to 72%. This was commensurate with a shift of +2° to the SPB phase. Further analysis showed that the R matrix was not responsive when the LTB efficiency dropped, and could drive the linac into a worse state. NNet2700 under the same conditions showed no such behaviour. Therefore it was deemed that the R matrix was not sufficient to correct for long term linac drifts.

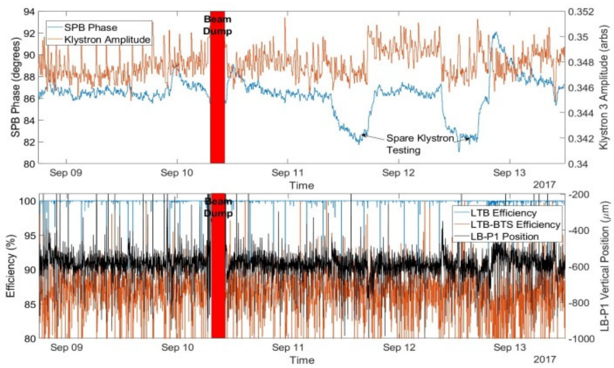

Figure 4: Behavior of NNet2700 during four days of operation. Note the action of NNet during klystron testing. The LTB efficiency remains constant.

NNet5400 and the polynomials were used from the beginning of 2018. Figure 5 shows the behaviour of NNet5400 during a four day period in February. The effect on the beam position is apparent when the NNet is started. A two degree shift in the SPB phase is apparent during the span on the graph. The large upward jump to 50 degrees is after the first klystron in the linac tripped. Upon recovery, the power level is not exactly the same, and NNet5400 corrected the SPB phase to maintain the beam position and efficiency.

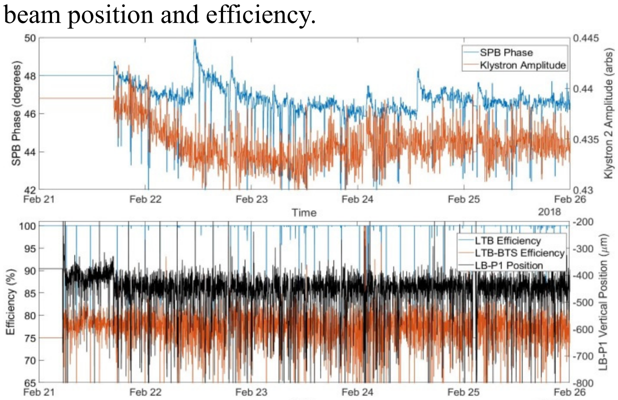

Figure 5: Behaviour of NNet 5400 during five days of operation. The LTB efficiency remains constant as NNet5400 makes changes during the first day.

Figure 6 shows the behaviour of the polynomial pair during operation. After the polynomials were started, there is a 1 degree shift on the SPB phase, with no effect on the efficiency, but a slight beam shift. The SPB controller needed to be reset late on March 8 for unrelated reasons. The polynomials are able to maintain the efficiency after the reset, and hold it through the remainder of the run.

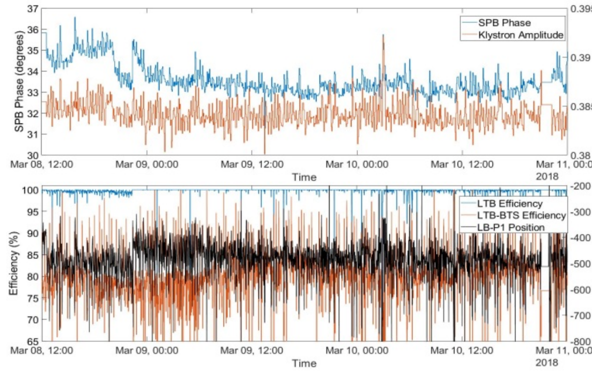

Figure 6: Action of the polynomials during two days of operation.

Three of the four methods used were successful in maintaining the injector efficiency and were roughly equivalent in performance. Table 1 lists the rms corrections to the SPB phase and the klystron amplitude along with the efficiencies through the injector. Since this system is meant to correct for long term drifts with a timescale of hours, and makes corrections every 4-5 minutes, the corrections are expected to be small. Table 1 also shows that each model was roughly equivalent in maintaining the efficiency of the LTB and the whole injector.

Table 1: RMS Corrections From Each Method with Efficiencies

Model	SPB Phase	Klystron Amplitude (x10-3)	LTB efficiency (%)	Injector efficiency (%)
NNet2700	0.8	1.5	97.9±7.5	83±10
NNet5400	2.5	3.7	99.6±3.5	77±4
Polynomials	1.3	2.6	99.3±3.9	80±8

CONCLUSION

Several methods of correcting the NSLS-II linac drifts have been presented. Two of these methods rely on machine learning techniques, while the others are least square fits to data. While the linear approximation is not a good choice, all of the other methods worked with success. Some differences were noted in the performance of the methods, but all worked equally well in operation.

REFERENCES

[1] R. Fliller *et al.,* "Results of NSLS-II LINAC Commissioning," in: *Proc. 4th Int. Particle Accelerator Conf.,* Shanghai, China, May 2013, pp. 2301-2303.

[2] R. Fliller, A. Derbenev, T. Shaftan, G-M. Wang, "Top Off Algorithm Development and Commissioning at NSLS-II,“ in: *Proc. 7th Int. Particle Accelerator Conf,* Busan, Korea, 2016, pp. 2988-2990.

[3] A. L. Edelsen *et al*, "Neural Networks for Modeling and Control of Particles Acclerators," Fermilab, Batavia, IL, FERMILAB-PUB-16-121-AD, 2016.