ISBN: 978-3-95450-221-9 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2021-THAL04

MACHINE LEARNING BASED TUNING AND DIAGNOSTICS FOR THE ATR LINE AT BNL

J. P. Edelen*, K. Bruhwiler, E. Carlin, C. C. Hall, RadiaSoft LLC, Boulder, CO, USA
K. A. Brown, V. Schoefer, Brookhaven National Laboratory, Upton, NY, USA

Abstract

Over the past several years machine learning has increased in popularity for accelerator applications. We have been exploring the use of machine learning as a diagnostic and tuning tool for the transfer line from the AGS to RHIC at Brookhaven National Laboratory. In our work, inverse models are used to either provide feed-forward corrections for beam steering, or as a diagnostic to illuminate quadrupole magnets that have excitation errors. In this paper we present results on using machine learning for beam steering optimization for a range of different operating energies. We also demonstrate the use of inverse models for optical error diagnostics. Our results are from studies that use both simulation and measurement data.

INTRODUCTION

Machine learning (ML) has seen a significant growth in its adoption for widespread applications. In particle accelerators ML has been identified as having the potential for significant impact on modeling, operation, and controls [1,2]. These techniques are attractive due to their ability to model nonlinear behavior, interpolate on complicated surfaces, and adapt to system changes over time. This has led to a number of dedicated efforts to apply ML, and early efforts have shown promise.

For example, neural networks (NNs) have been used as surrogates for traditional accelerator diagnostics to generate non-interceptive predictions of beam parameters [3, 4]. Neural networks have been used for a range of machine tuning problems utilizing inverse models [5, 6]. When used in conjunction with optimization algorithms neural networks have demonstrated improved switching times between operational configurations [7]. Neural network surrogate models have also been demonstrated to significantly speed up multi-objective optimization of accelerators [8]. Additionally, ML has been of interest for anomaly detection, using autoencoders, for root cause analysis [9], and for outlier detection, using large data-sets of known good operational states [10].

In this work we seek to apply ML methods — for both tuning and anomaly detection — on the AGS to RHIC transfer line at Brookhaven National Laboratory. Specifically, we employ the use of inverse models for these applications. The application of inverse models for anomaly detection is a burgeoning area of research in many other fields that has not seen much attention in particle accelerators. Here we present our work towards implementing inverse models to detect errors in quadruples using only beam position monitors and corrector data. We will demonstrate the utility of this approach using a toy model, and then show how it scales to a larger system such as the AGS to RHIC transfer line. We will then show results of training inverse models using data from the machine and discuss future work for this effort.

* jedelen@radiasoft.net

THE ATR LINE

The transfer line between the Alternating Gradient Synchrotron (AGS) and RHIC, or the so-called the ATR line [11, 12], must be retuned for different energies when RHIC changes its operating point. The transfer line controls the orbit matching, optics matching, and dispersion matching of the beam into RHIC. The transfer line is broken down into four sections. The U and W lines are seen by all beams entering RHIC while the X and Y lines are used for injection into the Blue and Yellow rings respectively. In this paper we focus our studies on the UW subset of the ATR line. The length of the transfer line presents challenges for tuning unto itself. The problem is further complicated by a 1.7 m vertical drop in order to get the beam from the AGS to RHIC.

The first part of the ATR (referred to as the U-line) starts with fast extraction from the AGS and stops before the vertical drop from the AGS to RHIC. The U-line consists of two bends. The first bend is 4.25°, consisting of two A-type dipole magnets. The second bend is an 8° bend consisting of four C-type combined-function magnets (placed in a FDDF arrangement), and thirteen quadrupoles. Optics in the U-line are configured to accomplish several goals. The Twiss parameters at the AGS extraction point must be matched, and provide achromatic transport of the beam to the exit of the 8° bend. The beam must be focused at the location of a thin gold foil which is placed just upstream of the quadrupole Q6 of the U-line. The Twiss parameters of the U-line must be matched to the ones at the origin of the W-line. Finally, the beam size should be kept small throughout to minimize losses.

The second part of the ATR (referred to as the W-line) introduces the vertical drop for injection into RHIC, and the matching sections for the injection lines. It contains eight C-type combined-function magnets that each make a of 2.5° bend, followed by six quadrupoles. The eight combined function magnets form a 20° achromatic horizontal bend placed in a (F-D) configuration. The W-Line is also responsible for lowering the beam elevation by 1.7 m. This is accomplished using two dipoles in an achromatic dogleg configuration. Along the line there are also a number of BPMs and correctors that are required to match the orbit of the beam into RHIC.

ISBN: 978-3-95450-221-9 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2021-THAL04

FODO TEST PROBLEM

Before applying inverse models for anomaly detection on the full UW line, we first demonstrate the efficacy of this technique using a toy problem. Our toy problem is a FODO lattice, containing BPMs and correctors, in addition to the quadrupoles. We built a model that predicts the corrector settings for a given sequence of BPM settings. Because the quadrupoles supply a dipole kick when the beam does not traverse them on axis, any error in the predicted corrector setting should be strongly correlated to an error in the quadrupole strength. Here we utilized a neural network based inverse model training on simulation data collected using MAD-X. A schematic of the beam-line used for this study is shown in Fig. 1.

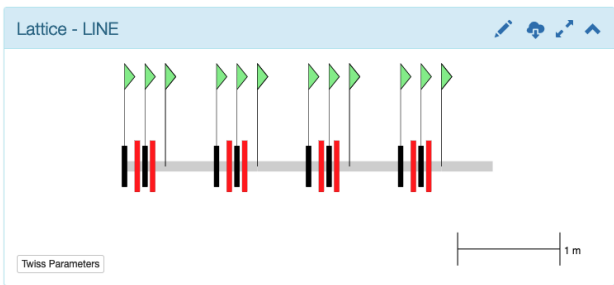

Figure 1: Screenshot of Sirepo visualization of toy lattice used for studying inverse models for anomaly detection. Quadrupoles are shown in red, correctors are black, and BPMs are indicated by green flags.

The beam line is composed of four identical FODO sequences. Each segment has two BPM/corrector pairs followed by a BPM after the second quadrupole. Having a large number of BPMs available, ensured that the inverse model was easy to train, which will allow us to use it to detect errors in the quadrupoles.

The training data consisted of 5000 examples simulated by randomly changing corrector strength and the initial beam position. The NN architecture was optimized as a function of the number of layers as well as the nodes per layer to improve training loss without over-fitting. Gaussian noise was used as a regularizer which leads to an increased noise on the validation loss as the epoch increases. Figure 2 shows the loss as a function of training epoch for the training and validation data.

Here we see that The training loss is a bit above the validation loss which indicates we could train longer. However, Fig. 3 shows the model prediction compared to the ground truth for each of the correctors (kickers). The relationship is almost perfectly linear in all cases, indicating the model is well trained.

Next we test the model using data with systematically introduced errors in individual quadrupoles. We ran the MAD-X simulations with random corrector strengths and initial beam positions. Then, the BPM outputs were used as inputs to the NN model — trained without quadrupole errors — to predict the corrector settings. These predicted corrector settings were then compared to the actual corrector settings.

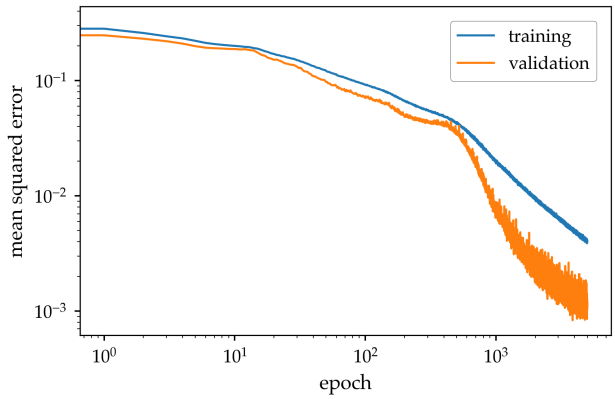

Figure 2: Mean squared error for the training and validation data as a function of the training epoch.

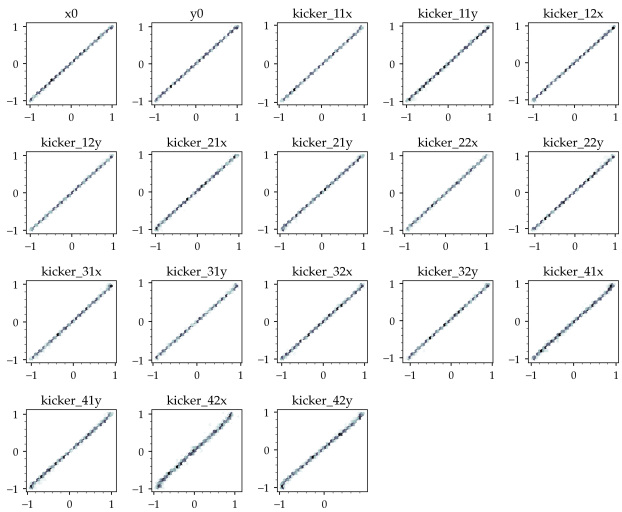

Figure 3: Predicted output vs. ground truth for the initial beam position and the kicker settings. The relationships are almost perfectly linear in all cases.

Figure 4 shows the predicted corrector setting versus the ground truth when testing the model using data generated with a strength error in one of the quadrupoles. Here it is very clear that the correctors in the first section of the beamline (kicker_11 and kicker_12) are not well reconstructed while all the other correctors are well reconstructed. This points to quadrupole strength errors in the first part of the beam-line.

We then studied the prediction error for a wide range of quadrupole errors and compared our ability to reconstruct the corrector settings. Figure 5 shows the normalized prediction error vs corrector index for data with different quadrupole errors. For each test case we introduced a single quadrupole error at different points along the beamline and we observed prediction errors in both the horizontal and vertical corrector settings.

Here we see that for a given quadupole there is a unique error signature for the correctors. Additionally, the corrector indices are arranged by their position in the beamline. As the quadropoles that are varied move down the beamline so do the errors in the corrector predictions. This is likely because

ISBN: 978-3-95450-221-9 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2021-THAL04

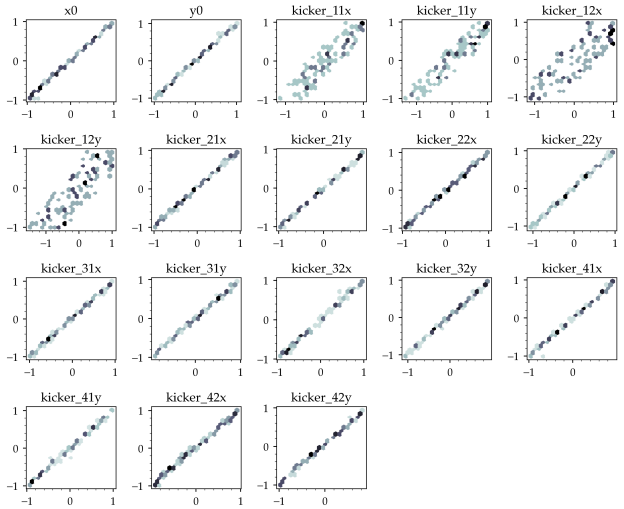

Figure 4: Predicted output vs. ground truth for data with a single quadrupole strength error.

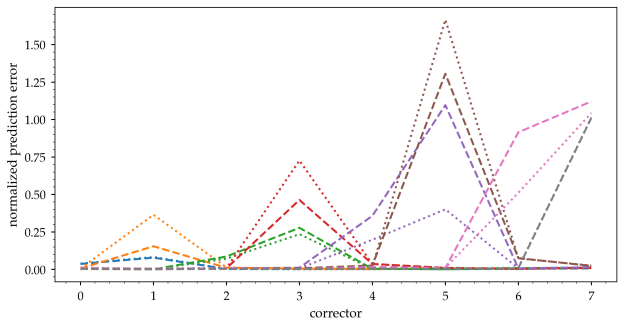

Figure 5: Normalized prediction error of corrector strengths given a range of Quadrupole errors. The dashed line is the horizontal corrector and the dotted line is the vertical corrector. The color indicates which Quadrupole was varied.

the BPM nearest the quadrupole with the error will be the most sensitive to this perturbation. Additionally due to the degenerate solutions in this type of model, it can be difficult to differentiate upstream errors from downstream errors. If we want to use this diagnostic for tuning, it is important to also examine the error in the corrector prediction as a function of the change in quadruple strength. Figure 6 shows the normalized prediction error as a function of quadrupole strength error for a single quadrupole.

There is a clear quadratic relationship between the quadrupole strength error and the normalized prediction error. This means that we can clearly detect anomalies in quadrupoles using our inverse model trained on the BPMs and correctors. Moreover, this means that we can vary the quads to minimize the model error based on measured BPMs and correctors, allowing us to tune the lattice without using destructive beam optics measurements.

ATR INVERSE MODEL

Given the success of this effort on the toy model we next applied this concept to the UW line in the ATR. The problem

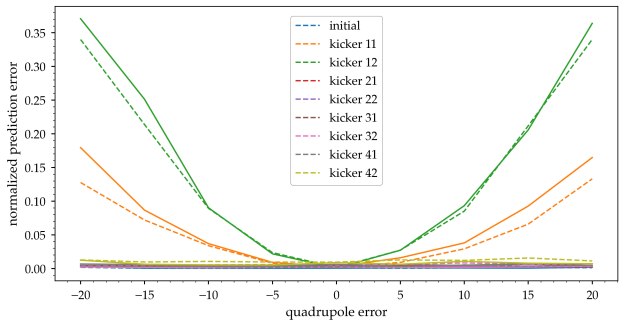

Figure 6: Normalized prediction error as a function of quadrupole strength error for each kicker (m^{-2}). The solid line is the horizontal kicker and the dashed line is the vertical kicker.

is a bit more complex as there are now not just quadrupoles, correctors, and BPMs, but also combined function bends and vertical bends.That said there are still comparatively few BPMs and correctors — only 26 and 14 respectively — making the model of similar scale to the toy problem. However there are 19 quadrupoles which significantly increases the complexity.

We trained the inverse model using 5000 samples, randomly varying the corrector strengths and beam initial positions. During our initial training of the inverse model four correctors (utv4, uth6, utv7, and wth1) were not well fit. This is likely due to the degenerate solutions that arise from the length of the transfer line. We attempted a wide range of solutions including sampling more data and working with image representations of the data 1-D convolutional layers. None of these methods were successful. When training inverse models there is always the possibility that the problem will not be fully invertable. In future work we will address this issue, however, for studies presented here we simply removed those four correctors from the prediction. Figure 7 shows the training and validation loss for the inverse model trained on the UW line.

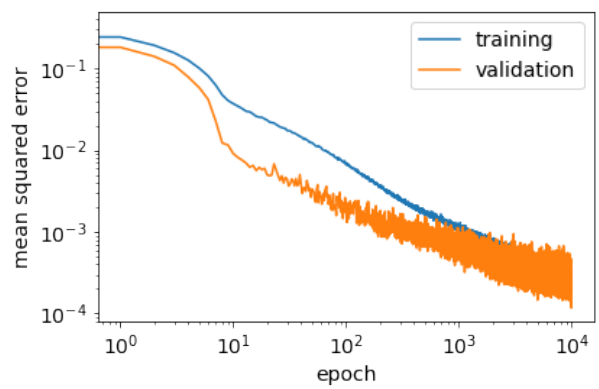

Figure 7: Training and validation loss as a function of training epoch for the UW line.

Here we see very good agreement between the training and validation set. Additionally the loss continues to decrease, showing that we are not over-fitting. For this study the data

ISBN: 978-3-95450-221-9 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2021-THAL04

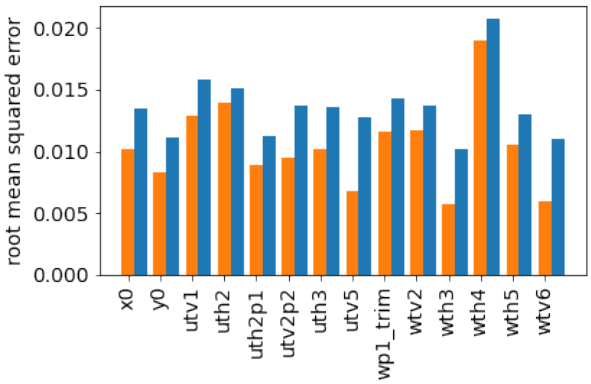

Figure 8: Root mean squared error on the training (blue) and validation (orange) datasets for the individual correctors.

were split into 80 % training and 20 % validation. Figure 8 shows the root mean squared error for each of the parameters individually.

The prediction error on the training and validation datasets are quite good. For this model we used 5 dense layers with 45 nodes each and Gaussian noise for regularization. The model used rectified linear units for the activation functions.

INVERSE MODEL ERROR STUDIES

With a trained inverse model we moved on to the error studies. Quadrupoles were individually varied over a range of ±20 % of their design excitation. We then simulated the UW line with random initial beam positions and random corrector settings. The BPM output from these simulations was used to predict the corrector settings using the model trained on data with the design quadrupole strengths. These simulations were run in two configurations, one where the initial positions were also varied randomly and one where the initial positions were not varied. The goal here was to understand if the model can differentiate between these two cases. During operations it is likely that the initial position will be relatively static.

Figure 9 shows the predicted corrector settings vs the ground truth for the validation set without quadrupole errors (black), the test set with a single quadrupole error and random initial position errors (red), and the test set with a single quadrupole error without initial position errors (blue).

Here we can clearly see that some correctors have a significant error in the prediction compared to others. For example Utv2p2 has a nice linear relationship but with a clear jump near zero. While UTH3, on the other hand, has a general increase in the error, as compared with WTV6, where the errors are relatively low. The model also shows a lack of ability to distinguish between cases where the initial position is varying and is not varying. The spread in the errors appear consistent between the two data sets. Indeed, Figs. 10 and 11 — which show the sensitivity of each corrector prediction to a particular quadrupole error — demonstrate that there is negligible difference between the cases with and without initial position errors. Here sensitivity is defined as the model

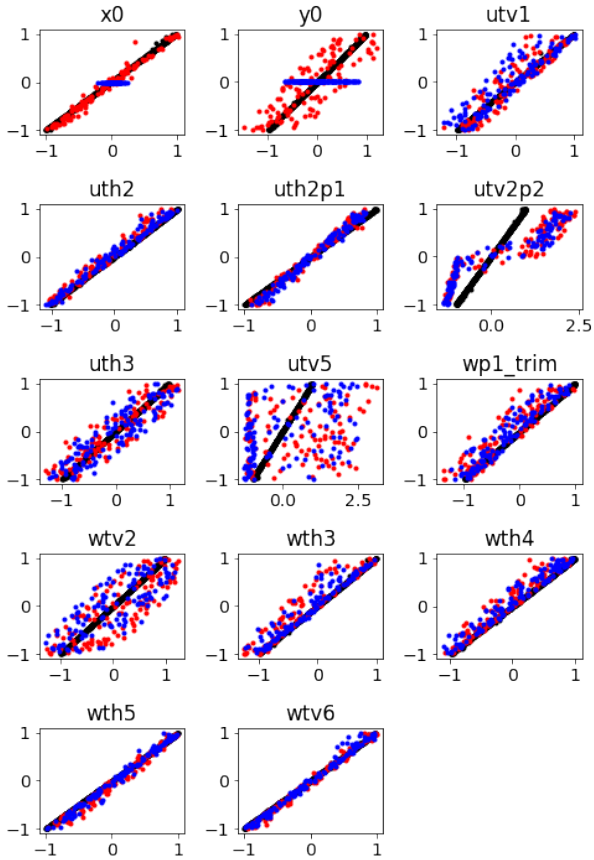

Figure 9: Predicted corrector settings as a function of the ground truth for, the three test cases. Without quadrupole errors: black. A single quadrupole error and random initial position errors: red. A single quadrupole error without initial position errors: blue.

error divided by the fractional change in quadrupole strength. These figures also show that there is a unique signature for each quadrupole and that the model clearly identifies errors in these magnets without any explicit knowledge of their existence.

These results are promising for the adaption of this method on the machine. Especially because the UW line is designed to be linwar, we expect the methods developed in simulation to behave similarly when transferred to the machine.

MACHINE STUDIES AND BPM INVERSE MODELING

We are in the early stages of testing our method on the machine, and have collected BPM and corrector data for the nominal machine configuration. Working with the UW line allows us to take data between injections to RHIC, parasitic to operations. These studies took approximately four hours of beam time without any interruption to operations.

Our main goals with this first study is to answer two key questions 1) how much data do we need to train an inverse model for the transfer line and 2) establish the feasibility of a neural network based inverse model for detecting quadrupole

ISBN: 978-3-95450-221-9 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2021-THAL04

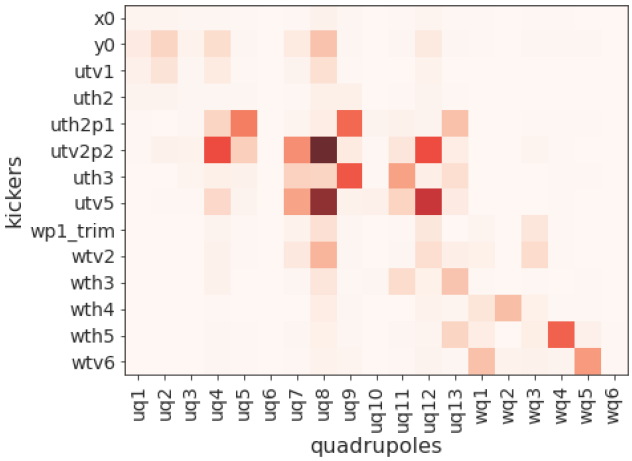

Figure 10: Sensitivity of each corrector prediction to individual quadruple errors. In this case the initial position was randomly varied along with the correctors.

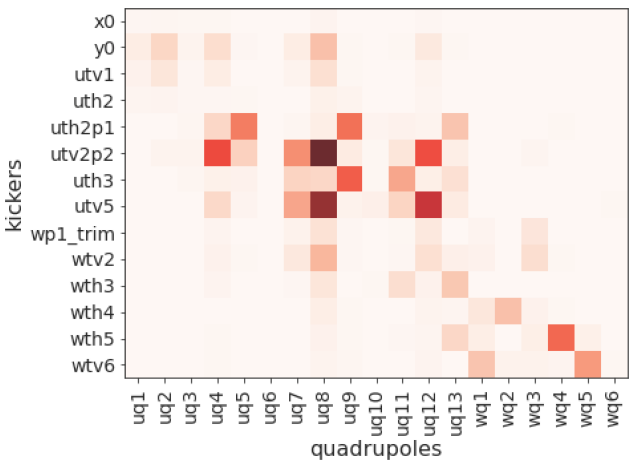

Figure 11: Sensitivity of each corrector prediction to individual quadruple errors. In this case the initial position not varied.

errors in the ATR line. With these data we can evaluate the accuracy of the ATR line MAD-X model, study the ability to transfer our NN inverse model from simulation data to measured data, and finally test our ML models on real data from the machine.

Figure 12 shows histograms of the beam position monitor data collected on the ATR line. Note that our initial dataset is relatively small, and has regions where there are quite a few outliers. The outliers in the dataset make it difficult to effectively sample the space for training neural networks. This will present a challenge when training on machine data alone. We plan to overcome this by combining the measured data with simulation data prior to testing on the machine.

With this dataset we performed some initial studies in order to understand the limitations of the measurements, and to help identify areas for additional data collection, or where simulations would be particularly useful. We are focused on using inverse models that predict the corrector settings from the BPM readings.

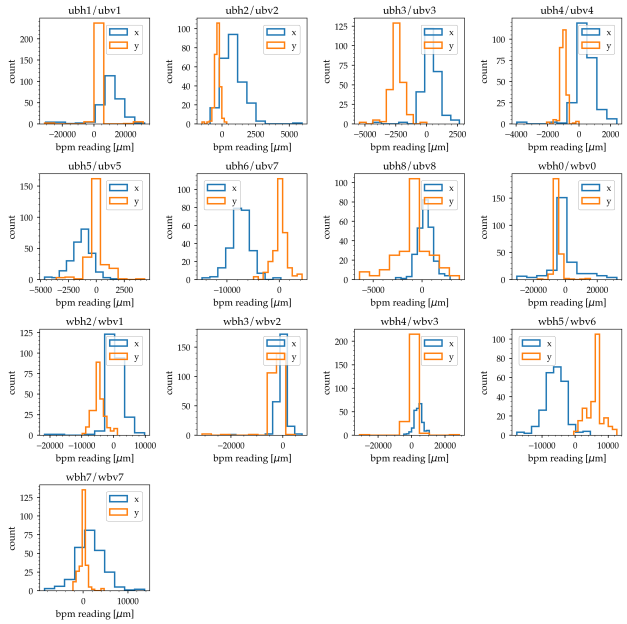

Figure 12: Histogram of BPM data collected from the ATR line. Horizontal (x) and vertical measurements are shown in blue and orange respectively.

For this study we varied numerous hyper parameters for the model, including: the number of nodes per layer, the number of layers, and the batch size. We also varied the regularization parameters to include both L2 regularization and Gaussian Noise. L2 regularization seeks to optimize the neural network architecture while simultaneously eliminating weights that are very small. The result is a more sparse network, but fewer nodes that are tracking noise in the data. Gaussian noise can assist the optimizer in getting out of local minima leading to a more robust solution. In the end, we trained on a batch size of 5 and aggressive Gaussian noise which helped the model fit such a small dataset. Figure 13 shows the training and validation loss as a function of training epoch for our neural network.

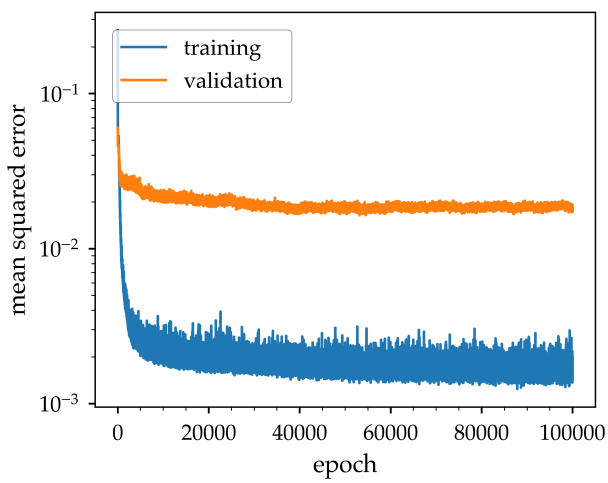

Figure 13: Loss curve for training and validation sets for the BPM data

ISBN: 978-3-95450-221-9 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2021-THAL04

The training and validation loss both decrease quite rapidly and then flatten out. The validation loss continues to improve slightly but does not change significantly after about 40000 epochs. With small datasets it is important to train for a long time in addition to restarting with different batch sizes to introduce perturbations to the network that allow it to generalize better. The large offset between the training and validation loss indicates that the network would benefit from more examples. The fact that the validation loss is not increasing indicates that we are not overfitting though. In spite of the relatively large discrepancy between the training and validation loss, the performance on the validation set is quite reasonable. Figure 14 shows the predicted corrector setting compared to the ground truth for the validation set.

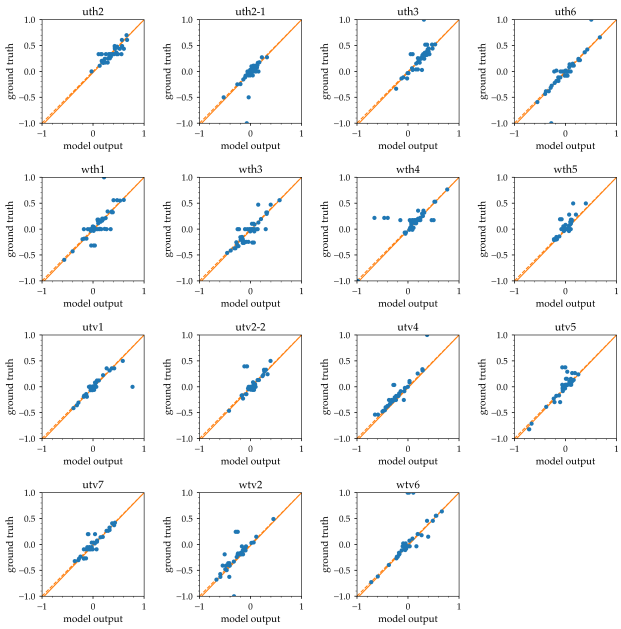

Figure 14: Fit results of the validation data for neural network trained on bpm on our study data

The solid orange line is the linear fit between the ground truth and the model output and the dashed line is the ideal fit should the model accurately reconstruct the corrector settings from the BPMs. In general the model performs relatively well, however as noted earlier it is clearly biased by small samples at extremes.

CONCLUSION

In this paper we explore the use of inverse models to detect errors in quadrupole strengths using BPM and corrector data. We have demonstrated an initial test using a toy model consisting of four FODO cells. We then scaled this to the Uw line on the ATR beamline at RHIC. Our results show that inverse models can identify quadrupole errors by comparing the predicted corrector setting to actual corrector settings. We also show that the each quadruple strength error yields a unique signature. We have begun transferring our work for testing on the machine and have shown a neural network inverse model can be trained on real BPM data from the UW line.

ACKNOWLEDGMENTS

This work was supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, Award Number DE-SC0019682.

REFERENCES

[1] A.L. Edelen, S.G. Biedron, B.E. Chase, D. Edstrom, S.V. Milton, and P. Stabile, "Neural Networks for Modeling and Control of Particle Accelerators", in *IEEE Transactions on Nuclear Science*, vol. 63, no. 2, pp. 878–897, April 2016. doi:10.1109/TNS.2016.2543203

[2] A.L. Edelen *et al.*, "Opportunities in Machine Learning for Particle Accelerators", 2018. arXiv:1811.03172[physics.acc-ph]

[3] C. Emma *et al.*, "Machine learning-based longitudinal phase space prediction of particle accelerators", *Phys. Rev. Accel. Beams*, vol. 21, p. 112802, 2018. doi:10.1103/PhysRevAccelBeams.21.112802

[4] A.L. Edelen, S. Biedron, J.P. Edelen, and S.V. Milton, "First Steps Toward Incorporating Image Based Diagnostics into Particle Accelerator Control Systems Using Convolutional Neural Networks", in *Proc. NAPAC'16*, Chicago, IL, USA, Oct. 2016, pp. 390–393. doi:10.18429/JACoW-NAPAC2016-TUPOA51

[5] A.L. Edelen *et al.*, "Using Neural Network Control Policies For Rapid Switching Between Beam Parameters in a Free Electron Laser", in *Proceedings of the 2017 Deep Learning for Physical Sciences Workshop at the 31st NeurIPS*. https://ml4physicalsciences.github.io/2017/files/nips_dlps_2017_16.pdf

[6] J.P. Edelen, K.A. Brown, N.M. Cook, and P.S. Dyer, "Optimal Control for Rapid Switching of Beam Energies for the ATR Line at BNL", in *Proc. ICALEPCS'19*, New York, NY, USA, Oct. 2019, pp. 789–794. doi:10.18429/JACoW-ICALEPCS2019-TUCPL07

[7] A. Scheinker *et al.*, "Demonstration of Model-Independent Control of the Longitudinal Phase Space of Electron Beams in the Linac-Coherent Light Source with Femtosecond Resolution", *Phys. Rev. Lett.*, vol. 121, p. 044801, 2018. doi:10.1103/PhysRevLett.121.044801

[8] A.L. Edelen *et al.*, "Machine learning for orders of magnitude speedup in multiobjective optimization of particle accelerator systems", *Phys. Rev. Accel. Beams*, vol. 23, p. 044601, 2020. doi:10.1103/PhysRevAccelBeams.23.044601

[9] J.P. Edelen and C.C. Hall, "Autoencoder Based Analysis of RF Parameters in the Fermilab Low Energy Linac", *Information*, vol. 12, no. 6, p. 238, 2021. doi:10.3390/info12060238

[10] *Proceedings of the 2021 Improving Scientific Software Conference*, W. Hu, D. Del Vento, and S. Su, Eds. Boulder, Co, USA: NCAR, 2021. doi:10.26024/p6mv-en77

[11] W.W. MacKay *et al.*, "AGS to RHIC transfer line: Design and commissioning", *Proc. EPAC'96*, Sitges, Spain, June 1996, paper MOP005G, pp. 2376–2378, 1996. https://jacow.org/e96/PAPERS/MOPG/MOP005G.PDF

[12] T. Satogata *et al.*, "Physics of the AGS-to-RHIC transfer line commissioning", in *Proc. EPAC'96*, Sitges, Spain, June 1996, paper MOP006G, pp. 2379–2381, 1996. https://jacow.org/e96/PAPERS/MOPG/MOP006G.PDF