10th Int. Beam Instrum. Conf. IBIC2021, Pohang, Rep. of Korea JACoW Publishing
ISBN: 978-3-95450-230-1 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2021-TUPP41

MACHINE LEARNING METHODS FOR SINGLE SHOT RF TUNING

J. S. Lundquist*, S. Werin, Lund University, Lund, Sweden
N. Milas, ESS, Lund, Sweden

Abstract

The European Spallation Source, currently under construction in Lund, Sweden, will be the world's most powerful neutron source. It is driven by a proton linac with a current of 62.5 mA, 2.86 ms long pulses at 14 Hz. The final section of its normal-conducting front-end consists of a 39 m long drift tube linac (DTL) divided into five tanks, designed to accelerate the proton beam from 3.6 MeV to 90 MeV. The high beam current and power impose challenges to the design and tuning of the machine and the RF amplitude and phase have to be set within 1% and 1° of the design values. The usual method used to define the RF set-point is signature matching, which can be a challenging process, and new techniques to meet the growing complexity of accelerator facilities are highly desirable. In this paper we study the use of ML to determine the RF optimum amplitude and phase, using a single pass of the beam through the ESS DTL1 tank. This novel method is compared with the more established methods using scans over RF phase, providing similar results in terms of accuracy for simulated data with errors. We also discuss the results and future extension of the method to the whole ESS DTL.

INTRODUCTION

The European Spallation Source (ESS) is a state of the art neutron science facility under construction in Lund, Sweden [1]. The basic process used by the facility is spallation, wherein one impinges a high neutron material, in this case Tungsten, with high energy protons, causing the target to shed excess neutrons. The high energy protons are provided by the ESS linear accelerator (linac), a 600 m long accelerator consisting of many different sections utilizing varied accelerator technologies in order to raise the proton energy from the 75 keV source output to the final 2.0 GeV arriving on the target. A crucial part of this machine is the 39 m long drift tube linac (DTL) divided into five tanks, designed to accelerate the proton beam from 3.6 MeV to 90 MeV. As the machine is expected to deliver beam of high current and power, a primary concern is to avoid slow beam losses, as these lead to radiation activation of surrounding equipment. In order to avoid such losses, proper and careful tuning of the RF fields is crucial. As a result the requirement for accuracy of the RF set point is to be within 1% in RF amplitude and 1° in phase [1]. In order to achieve this type of accuracy, much work has been performed in the last decades to develop new techniques to meet the growing scale and complexity of facilities [2–4]. Within this paper we will investigate how Machine Learning (ML) may serve this purpose. This paper presents our current strides in the development of a tuning technique using ML, with simulated data used in such a way that a single pass through the untuned cavity could be sufficient for setting it up.

* johan.lundquist@ess.eu

RF TUNING

RF Phase Scan

In order to be able to quantify how the beam responds to changes to the RF set-point, a diagnostic sensitive to the beam time of flight through the cavity must be used. For those cases a Beam Position Monitor (BPM) can be used. As the beam passes a BPM both the amplitude and phase of the fields excited on the BPM sensor by the passing beam are recorded. Although this phase alone doesn't hold much usable information for cavity tuning, by comparing two BPM phases we can get a fast measurement which is proportional to the time-of-flight, or looking with respect to acceleration in a RF cavity, the energy gain between the two devices. It is important to stress that this measurement is relative and that extracting the absolute values of the energy is not an easy task. For this technique, using only the relative phase changes has proven to be enough [2–4].

The BPMs are used to measure the energy gain (or time-of flight) as a function of the set points in the accelerating cavity. As the BPM's measured phase is closely dependent on the energy of the beam, scanning RF amplitude and phase in a cavity and plotting out the resulting phase differences will give rise to different curves depending on the proximity to the ideal set point for the cavity. A few of these signature curves can be seen in Fig. 1, where the ideal set point can be found from the signature for the ideal amplitude A_0 = 6.89 kV, the ideal input beam energy E_0 = 3.62 MeV and the -35° phase set point.

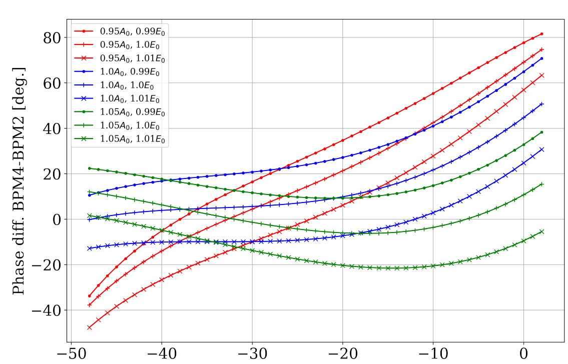

Figure 1: The phase curves for different RF amplitude and input energy set points. BPM phases simulated as comparison between two BPMs in the first DTL tank in the ESS linac.

ISBN: 978-3-95450-230-1 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2021-TUPP41

Identifying these types of signatures is the basis of most established techniques for cavity tuning [2–4]. We are suggesting a new type of identifiable signature for this type of tuning, which could be measured in a single pass through the cavity, rather than requiring scanning over a parameter, described in more detail in the following section.

Single Shot Measurement

RF phase scans are an established and reliable method for extracting the information needed to achieve good tuning, and with a more limited diagnostic output it is the only option available for now. However, with the large number of BPMs within the ESS DTL section, a restructuring of the data can be done such that we can see distinct signatures for each cavity setpoint in amplitude, phase and beam input energy. We look at BPM phase differences, not against RF phase, but against each diagnostic output, the pairing of BPMs. Figure 2 shows an example of this type of plot, where each line represents a cavity set point and is measured in a single pass through the machine, without scanning any parameter. From here we encounter the same problem to be solved as with the phase scan data, needing to accurately identify these new signatures. The nature of the signatures in this data format leaves ML uniquely equipped for the task.

If accurate predictions can be made with this data format, a few new advantages manifest. Being able to tune the machine acceptably with a single shot could cut down set up times. One would also not require to determine a range for the scanned parameters as in more traditional RF tuning [2, 3] but would in principle reload the last machine state with good settings and run a single verification pulse.

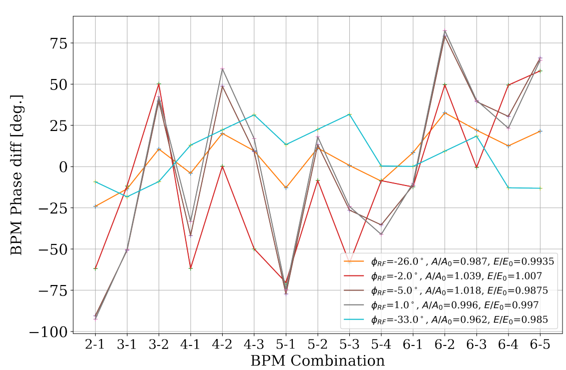

Figure 2: BPM phase differences for each possible BPM coupling, with the different plots each corresponding to a single cavity set point.

Simulations

OpenXAL was used to simulate the first tank of the ESS DTL during acceleration and to reproduce the signals from the six BPMs inside the DTL tank 1 [5, 6]. As phase difference is the data of interest, this results in 15 different BPM combinations, each combination producing one data point for each cavity set point, in RF amplitude and phase and input beam energy. ML requires large amounts of data for training networks and for this purpose an error free dataset was used. This consisted of 110 different amplitude set points, with a variation of ±5.5% around the design RF cavity amplitude A_0, 60 different input beam energy set points, with a variation of ±1.5% around the design input energy E_0, and 55 different phase set points, spread evenly around the -35° design set point.

To this perfect machine four different types of errors were then applied. BPM longitudinal position within the machine was adjusted, potentially caused by installation and construction, as well as the phase readout from these BPMs, produced by electronic limitations. There are also errors arising from production limitations when constructing the cavities. Such limitations could give rise to errors in both RF amplitude and phase gap-to-gap. The different types of errors and their magnitudes are summarized in Table 1.

Table 1: The Different Types of Errors Used in Simulations and Their Corresponding Magnitude

Error Type	**Magnitude**
BPM Δs	±100 µm
BPM $\Delta\phi$	±1°
RF Amplitude	±2%
RF Phase	±0.5°

MACHINE LEARNING

Machine learning algorithms come in many forms and can solve many distinct problems using varying network structures, definitions of loss and optimization algorithms [7]. The problem we are looking at in this project involves reducing larger scans of data down to three dependent variables, RF amplitude, RF phase and input beam energy. We compare two types of network, a traditional linear regression structure, and a newer decision tree boosting model called XGBoost.

Linear Regression Network

This network was defined using the python library Keras [8]. The library comes with predefined versions of our loss function, mean squared error, and our optimization algorithm, ADAM. ADAM has different coefficients which may be tuned to improve the networks performance, although the learning rate is most relevant [9]. Optimization of the network structure and training parameters was done iteratively, looking at generalized performance as the figure of merit. This was quantified as the loss on a subset of data separated during training. Through this process we arrived at a 10-layer structure with an 160-160-80-80-40 symmetrical neural layout, and a final output layer of three neurons. This was trained for 20 000 epochs with a learning rate of 0.00001. This network was used to produce the results presented in the following section.

ISBN: 978-3-95450-230-1 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2021-TUPP41

Table 2: Three standard deviations in difference between predicted and correct values for the RF Amplitude and phase and the input energy. Results shown for both linear regression (LR) and XGBoost (XGB) network structures.

Data Set		$3\sigma_A$ [%]	$3\sigma_\phi$ [°]	$3\sigma_E$ [%]	μ_A [%]	μ_ϕ [°]	μ_E [%]
No Errors	LR	0.075	0.051	0.045	0.002	0.000	0.002
	XGB	0.891	1.755	0.153	0.025	-0.013	0.002
All Errors	LR	4.002	5.568	0.804	-0.013	-0.068	-0.009
	XGB	3.171	5.217	0.750	0.022	0.038	0.016

XGBoost

The modern ML system of XGBoost (eXtreme Gradient Boosting) is an open source gradient boosting model, which has proven extremely powerful for solving varied, nonlinear problems [10]. Gradient boosting tree models such as this are based on the decision tree model of network structures, with a regularized objective function. In a decision tree ML system, the parameters adjusted in training are not the weighted connections within a network of nodes, but rather the branching criteria in a large decision tree. A gradient boosting ML system uses an ensemble of many decision trees in order to improve the final predictions, and commonly a regularized loss function which penalises increasing complexity of the model as well as the usual error of predictions. This regularized loss is then applied to the ensemble of trees iteratively to improve the output by training the branching criteria.

For the results produced here an ensemble of 10000 trees was used, each with a max allowed depth, the amount of branching criteria, of 20. A learning rate factor of 0.0001 was applied and an early stopping system was also used during training, forcing the training to halt if the generalized performance of the ensemble network did not improve over a period of 500 iterations.

RESULTS

Table 2 shows three standard deviations (3σ) and the mean (μ) of the difference between the predicted and expected value for the RF Amplitude (A) and phase (ϕ), as well as for the input energy (E). The low mean in all rows shows there is little to no systematic offset to the predictions. We see both types of network performing within the given limits on the training data set, although we do see higher standard deviation from the XGBoost training. However, it is generalized performance on the realistic error data set which presents the more relevant figure of merit. Here we see good performance in the energy predictions, but both methods fail to produce the sought results in both phase and amplitude. XGBoost only slightly outperforms the more traditional linear regressor, but remains far outside the limit in the phase prediction. The variation in the single shot signatures as a function of phase is quite small, so networks struggling to distinguish between these is understandable.

OUTLOOK

While both methods may fail the limits for operation at this stage of investigation, there are still many factors which could prove this method more reliable than suggested by these results. The error data set produced may have been pessimistic in the predictions of one or many of the factors included. Further optimization of the meta parameters for the training of the networks could reveal better results in the future. New data sets could be produced with more distinguished patterns arising from difference in RF phase.

Furthermore, even if this technique would prove unable to produce the sought results for initial cavity tuning, the 1° and 1% error in RF phase and amplitude predictions, it does not render it useless. The single shot nature of the data could allow for updated tuning information during operation of the machine, as well as long term tracking of drifts on the RF parameters. Also, the data available for each set point increases rapidly with added diagnostics. Applying this single shot data to the full ESS DTL section would include more BPM combinations, and the increase in available data could perhaps be sufficient to improve the method to within the restrictions. Further applications of this online tuning information could be developed in the future, for use in the ESS control room or elsewhere.

ACKNOWLEDGEMENTS

This study has only been possible with the help of the entire Beam Physics team at ESS, as well as that of other teams at the facility.

REFERENCES

[1] R. Garoby *et. al.*,“The European Spallation Source Design”, *Physica Scripta*, vol 93, p. 014001, 2018.

[2] G. Shen and M. Ikegami, “Tuning of RF amplitude and phase for the drift tube linac in J-PARC”, *Chin. Phys C*, vol. 33, p. 577, 2009.

[3] T.L. Owens, M. B. Popovic, E. S. McCrory, C. W. Schmidt and L.J. Allen, “Phase Scan Signature Matching for Linac Tuning”, *Particle Accelerators*, vol. 8, no 1, pp. 169-179, 1994.

[4] G. B. Shen, H. Sako, and S. Sato, “RF Amplitude and Phase Tuning of J-PARC SDTL”, in *Proc. 22nd Particle Accelerator Conf. (PAC'07)*, Albuquerque, NM, USA, Jun. 2007, paper TUPAN062, pp. 1529–1531.

[5] J. Galambos, A. V. Aleksandrov, C. Deibele, and S. Henderson, “PASTA An RF Phase and Amplitude Scan and Tun-

ISBN: 978-3-95450-230-1 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2021-TUPP41

ing Application”, in *Proc. 21st Particle Accelerator Conf. (PAC’05)*, Knoxville, TN, USA, May 2005, paper FPAT016, pp. 1529-1531.

[6] OpenXAL, https://openxal.github.io/.

[7] A. Edelen *et al.*, “Opportunities in Machine Learning for Particle Accelerators”, **1811.03172**, arXiv, 2018.

[8] Keras, https://keras.io/.

[9] D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization”, **1412.6980**, arXiv, 2017.

[10] T. Chen and C. Guestrin, “XGBoost: A Scalable Tree Boosting System”, **1603.02754**, arXiv, 2016.