Machine Learning Methods for Single Shot RF Tuning

J. S. Lundquist[1,2], S. Werin[2] and N. Milas[1], [1]ESS and [2]Lund University, Lund, Sweden

IBIC'21 September 13 - 17, 2021

IBIC 2021

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].
- 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.
- A simple model of the ESS linac can be seen in Fig. 1.
- **The requirement for accuracy of the RF set point in the DTL is to be within 1% in RF amplitude and 1° in RF phase [1].**

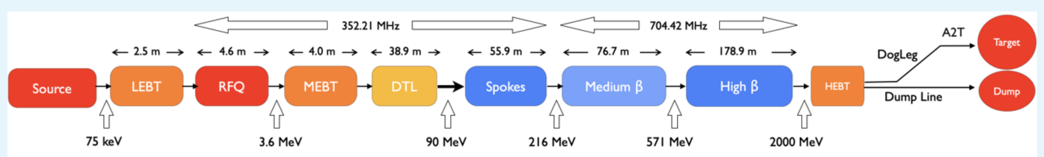

Figure 1: Simplified figure of the ESS linac.

Simulations

- Simulations of DTL Tank 1 done in OpenXAL [5].
- Envelope simulations were used, rather than trajectory tracking, which has proved to be sufficient information for RF tuning.
- ML Networks trained with perfect data sets, with no errors deviating from the lattice design files.
- Networks were then tested with new datasets with the errors shown in Table 2.

Table 1: The different types of errors used in simulations and their corresponding magnitude.

Error Type	Magnitude
BPM Δs	$\pm 100\,\mu$m
BPM $\Delta\phi$	$\pm 1^\circ$
RF Amplitude	$\pm 2\%$
RF Phase	$\pm 0.5^\circ$

RF Phase Scan

- By comparing two BPM phases we can get a fast measurement which is proportional to the time-of-flight.
- 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, 3, 4].
- 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.
- RF phase was scanned to produce signature curves shown in Fig. 2.

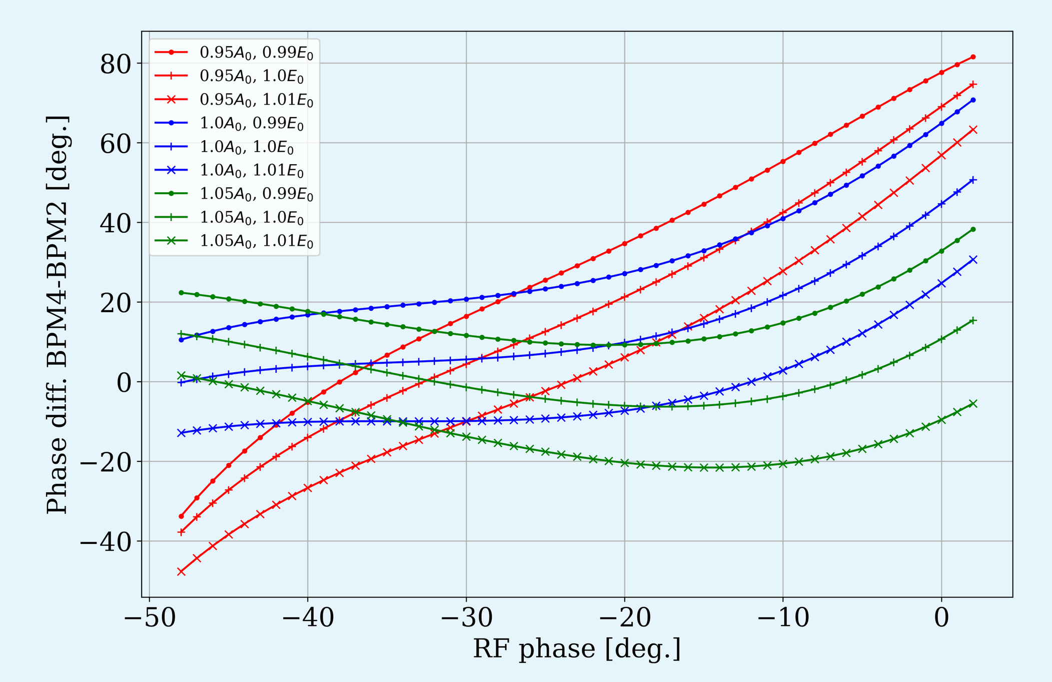

Figure 2: The phase curves for different RF amplitude and input energy set points. BPM phases simulated as comparison between two BPMs in the first DTL tank in the ESS linac.

- Identifying these types of signatures is the basis of most established techniques for cavity tuning [2, 3, 4].

Single Shot Measurement

- 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 3 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.

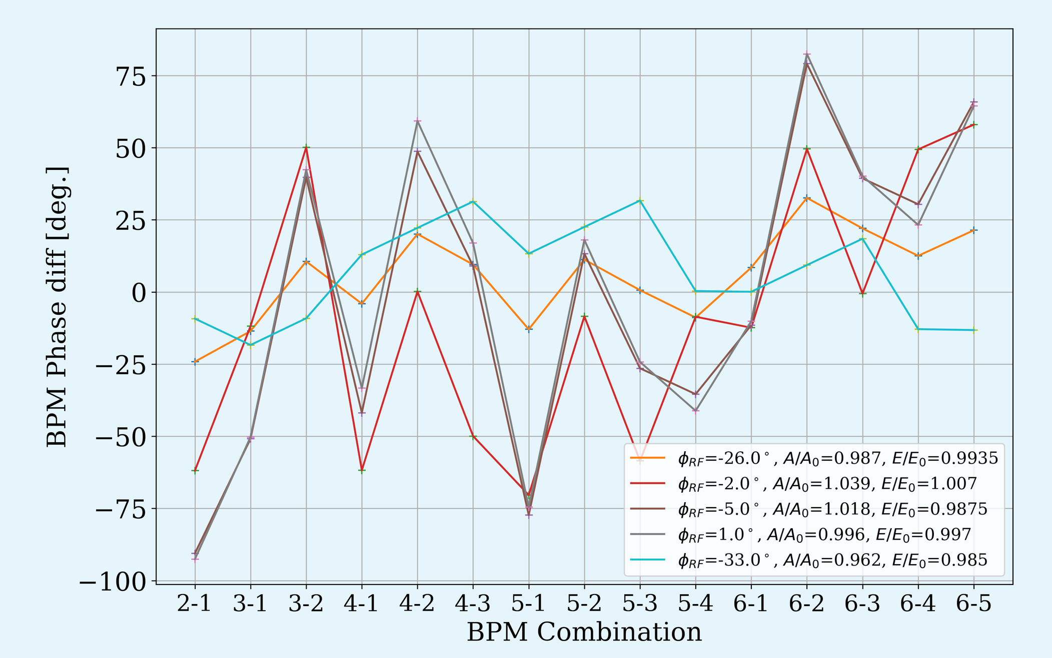

Figure 3: BPM phase differences for each possible BPM coupling, with the different plots each corresponding to a single cavity set point.

- 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.
- This new data format was used to produce the results presented here.

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 [6].
- We compare two types of network, a traditional linear regression structure, and a newer decision tree boosting model called XGBoost.

Linear Regression

- This network was defined using the python library Keras [7].
- Our loss function was mean squared error and our optimization algorithm was ADAM [8].
- Optimization of the network structure and training parameters was done iteratively, looking at generalized performance as the figure of merit.
- 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, which was trained for 20 000 epochs with a learning rate of 0.00001.

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 [9].
- 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.
- Figure 4 shows a simplified model of the principle of decision tree ensembles.

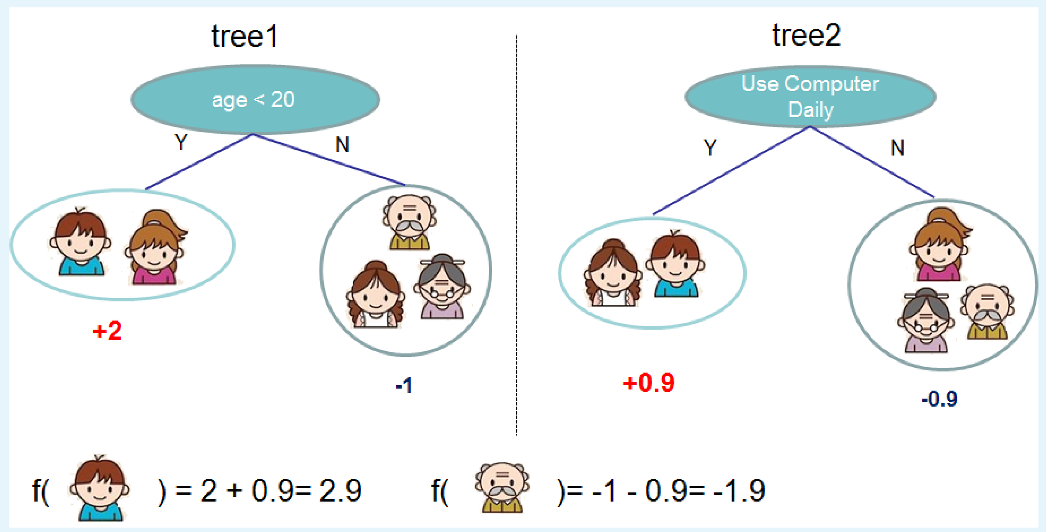

Figure 4: Simplified figure of the working principle of a tree ensemble network. In this case, the final prediction of the ensemble is a summation of the predictions of individual trees [9].

- 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).
- We see good performance in the energy predictions, but both methods fail to produce the sought results in both phase and amplitude.
- The variation in the single shot signatures as a function of phase is quite small, so networks struggling to distinguish between these is understandable.
- XGBoost only slightly outperforms the more traditional linear regressor, but remains far outside the limit in the phase prediction.

Conclusions

- 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 and further optimization of the meta parameters for the training of the networks could reveal better results in the future.
- 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.
- Further applications of this online tuning information could be developed in the future, for use in the ESS control room or elsewhere.

References

[1] R. Garoby *et. al.*,"The European Spallation Source Design", Physica Scripta, vol 93, 014001 (2018).

[2] G. Shen and M. Ikegami, Chin. Phys C, **33**, 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] OpenXAL,`https://openxal.github.io/`.

[6] A. Edelen et. al., "Opportunities in Machine Learning for Particle Accelerators", **1811.03172**, arXiv (2018).

[7] Keras, `https://keras.io`.

[8] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization", **1412.6980**, arXiv (2017).

[9] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System", **1603.02754**, arXiv (2016).

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		$3\sigma_A$ [%]	$3\sigma_\phi$ [$^\circ$]	$3\sigma_E$ [%]	μ_A [%]	μ_ϕ [$^\circ$]	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