MACHINE LEARNING TO SUPPORT THE ATLAS LINAC OPERATIONS AT ARGONNE*

B. Mustapha[†], J. Martinez, K. Bunnell, D. Stanton, E. Letcher, B. Blomberg and C. Dickerson
Physics Division, Argonne National Laboratory, Lemont, Illinois, USA

Abstract

The use of artificial intelligence can significantly reduce the time needed to tune the ATLAS heavy ion linac. After establishing automatic data collection procedures and analyzed the data, we have developed, and tested machine learning models to tune and control the machine. Models based on Bayesian Optimization (BO) and Reinforcement Learning (RL) are presented and their performance compared and discussed. RL and BO are well known AI techniques, often used for control problems. Results from beam tuning models are presented for a subsection of the ATLAS linac that contains complex elements such as the radio-frequency quadrupole. Another model was developed to assist in the commissioning of a new beamline. These models will be later generalized to the whole ATLAS linac, and similar models can be developed to control any accelerator with a modern control system.

INTRODUCTION

ATLAS is the DOE/NP User Facility for low-energy Nuclear Physics with heavy ions [1]. It operates ~6000 hours per year. In addition to delivering any stable beam from proton to uranium, the facility also provides radioactive beams from the CARIBU source [2] and via the in-flight radioactive ion separator, RAISOR [3]. The facility uses 3 ion sources and serves 6 target areas at energies from ~1-15 MeV/u. To accommodate the large number and variety of approved experiments, the ATLAS linac is tuned for a different ion species every 3 or 4 days over 40 weeks of operation per year. The start-up time varies from ~12-48 hours depending on the complexity of the tuning, which will increase with the upcoming Multi-User Upgrade to deliver beams to two experimental stations simultaneously [4]. DOE/NP is funding a project to use AI/ML to support ATLAS operations. The project aim is to reduce the accelerator tuning time and improve machine performance by developing and deploying artificial intelligence methods. The project goals are three-fold:

- Establish data collection, organization and classification, towards a fully automatic and electronic data collection for both machine and beam data
- Develop online tuning models to optimize operations, shorten beam tuning time and make more beam time available for the experimental program
- Develop a virtual machine model to enhance our understanding of the machine behavior, improve machine performance, develop and optimize particular and new operating modes.

It is expected that these developments will increase the scientific throughput of the facility and the quality of the data collected.

Significant progress has been made on the data collection and AI/ML modelling of the machine, and some of the results are presented below.

ESTABLISHING AUTOMATED DATA COLLECTION

The importance of data cannot be emphasized enough, it's essential for any AI/ML project, to the point that a data scientist spends 80% or more of his or her time collecting, organizing and labelling data. At ATLAS, we have two kinds of data, first the machine settings or tunes and second the beam data. Up until the start of this project, the machine data in terms of tunes was systematically collected and often used to re-tune the machine by scaling the settings from one ion species to the next. However, the beam data, such as beam currents and profile projections along the linac, were not systematically or automatically collected. The Faraday cup readings (beam currents) were recorded manually onto a sheet of paper during the tuning, while beam profiles were looked at on-the-fly but never recorded. Therefore, the beam data were either missing or not actually correlated to the machine data. As part of this project, we developed a system capable of collecting all the data required and correlating the beam data to the machine data in order to associate each tune with the corresponding beam information such as transmissions and profile projections throughout the different linac sections. Figure 1 shows a schematic of the data collection interface which is a python code connecting a server to the control system and saving the necessary data. The bottom half of the figure shows the data that are now collected automatically during the tuning and also systematically during an experimental run. It is important to note that the Faraday cup readings and beam profiles had to be digitized before saving.

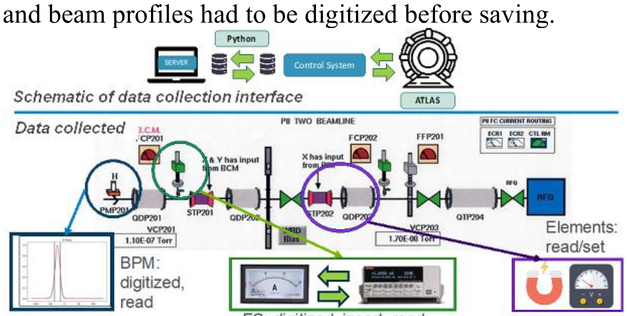

Figure 1: Schematic of the data collection interface (top). The Data being collected; element settings, beam currents and beam profiles at different locations on the linac.

* This work was supported by the U.S. Department of Energy, under Contract No. DE-AC02-06CH11357. This research used the ATLAS facility, which is a DOE Office of Nuclear Physics User Facility. Funded through a "Data, Artificial Intelligence, and Machine Learning at DOE Scientific User Facilities" Grant from the DOE's Office of Nuclear Physics.
† brahim@anl.gov

BAYESIAN OPTIMIZATION FOR INITIAL ONLINE TUNING

Due to the limited amount of beam data correlated with machine tunes we collected so far (~ 30 tunes), we have chosen Bayesian Optimization (BO) for our initial tuning model. BO is widely used [5] because it combines the complementary strengths of human and numerical optimization: life-long learning, learning by experience, not requiring large amounts of data, estimating its own uncertainty, and reaching global optimum in a minimum number of steps. This method starts with a prior belief regarding the objective function and then updates it based on samples drawn from the system for a better approximation, a posterior belief. It uses a probabilistic surrogate model for approximating the objective function and an acquisition function that instructs the model where to query the system next for more likely improvement.

A BO model was developed online for the tuning of two different beams; $^{14}N^{3+}$ and $^{40}Ar^{9+}$, through the ATLAS RFQ beamline to maximize beam transmission. The line consists of a triplet and two 2D steerers, varying a total of 7 parameters. For the case of $^{14}N^{3+}$, the starting data set included 29 historical tunes and 33 randomly generated tunes. Although, due to different initial beam distribution, the historical tunes didn't produce the expected high transmission, they're still more useful than the random tunes. For the case of $^{40}Ar^{9+}$, only the 29 historical tunes were used to start the model. Figure 2 shows the beamline and the results of the tuning model compared to the operators' manual tunes. We can clearly see that the model converges to a beam transmission similar to the operator's (dashed green line) in ~ 30 iterations. Note that the typical RFQ transmission is ~ 80%. On the right plots, we see how much the settings have changed from the operator's values. We notice that in some cases the change is minimal which means that the operator's tune was already good and there's not much room for improvement. In this case, the focus will be on how fast the model could tune the beamline compared to the manual tuning by the operators, future work.

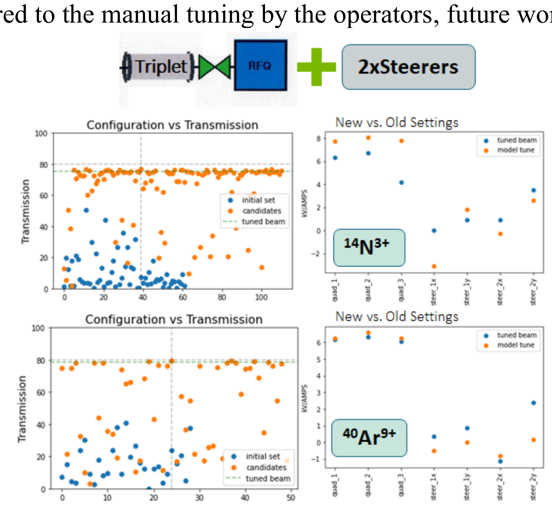

Figure 2: Top – Beamline, Middle – Results for $^{14}N^{3+}$ and Bottom – Results for $^{40}Ar^{9+}$. Right – Initial data set in blue and model iterations in orange. Left – Model settings for best beam transmission vs. original operators' settings.

REINFORCEMENT LEARNING FOR FINE TUNING / CONTINEOUS CONTROL

Reinforcement Learning (RL) is one of the three basic machine learning paradigms, alongside supervised and unsupervised learning. RL does not require labelled data because it learns from interactions between an AI agent and its environment. The idea behind using RL to tune/control a particle accelerator arises from the complexity of the system. In classic control, creating a large single function is more difficult than building a control system with piecewise subcomponents; however, this is where RL can help. In essence, RL maps situations to actions to maximize a numerical reward. There are different kinds of algorithms that can be applied, and the one selected here is Deep Deterministic Policy Gradient (DDPG) [6], which is an "actor-critic" approach that mixes policy optimization and Q-learning. Policy optimization methods tend to be more stable and reliable, and Q-learning is substantially more sample efficient. Although Q-learning is not consistent with continuous action spaces, which is the case of accelerators, DDPG supports continuous action spaces, because the critic only looks at the latest single action taken by the actor and does not try to find the best action by evaluating all of them. The actor is a neural network that takes what it thinks is the best action given the current state, as seen within the policy function method. The critic is a second neural network that estimates the value of the state and the action that the actor took, as seen within the value function method.

Simulation Results

Using simulation data for the case of focusing an ion beam on target using a quadrupole triplet (3 parameters), an RL model was developed. Figure 3 shows the results from the training and testing of the model. In this case the quads are electrostatic and were varied between 2 and 10 kV with alternating polarity and an action step of 0.25 kV. In the training results, we clearly see that the model learns the quads limits first, then starts learning how to focus the beam on target. In the testing results, the model quickly finds the quad setting to focus the beam on target from five different random starting points.

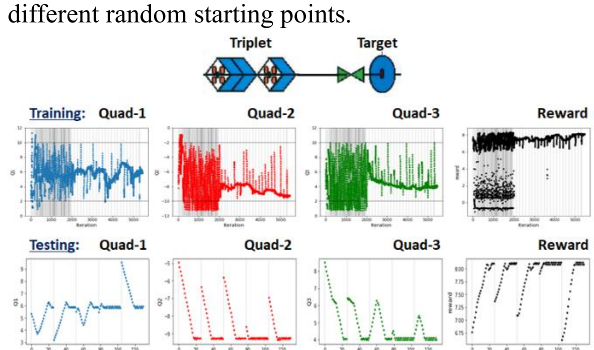

Figure 3: RL model results for the triplet beamline (top). The training results (middle) show the model learning the quads limit first then starts the optimization by maximizing the reward. The testing results (bottom) show the model quickly converging on the same solution from five different random starting points.

Experimental Results

For the experimental problem, we used the RFQ beamline, similar to the one used for the BO model, but varying two quad doublets and two steerer magnets, a total of 8 parameters, to maximize beam transmission. The electrostatic quads voltages were varied between 2 and 10 kV with action steps of 0.25 kV while the magnetic steerer current were varied between -1 A and +1 A with a step of 0.25 A. The model training results are shown in figure 4. We can see that the model started learning the limits for some of the parameters but didn't start the optimization process. This is mainly due to the limited beam time for this experiment. Although promising, this RL model will require a significantly longer time to converge than the BO model. A better approach would be to train the RL model offline first using simulation data, then re-train it and apply it experimentally online. Due to the direct control of the action step size, an RL model could be more useful for fine tuning and continuous control by watching the objective function and adjusting the beamline settings as they drift in time.

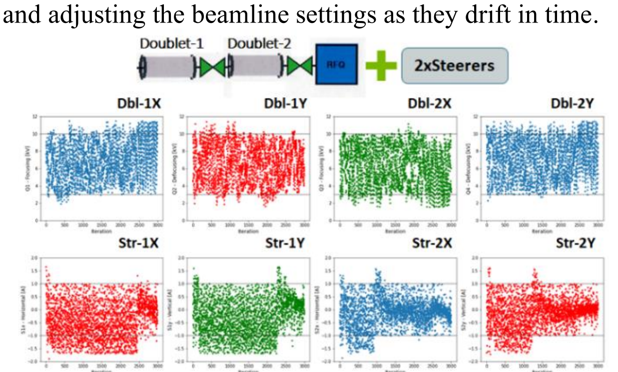

Figure 4: Results of RL model training online for the RFQ beamline with two quad doublets and two steerer magnets to maximize the beam transmission. The results show the model started learning some parameters limits but will require more time to start the optimization process.

AI/ML SUPPORTING THE COMMISSIONING OF A NEW BEAMLINE

As mentioned above, a BO model could be developed using a minimal amount of data, which could be very useful for the commissioning of a completely new beamline. At ATLAS, we have recently built a new beamline for material irradiation, AMIS (ATLAS Material Irradiation Station). Figure 5 shows the new beamline consisting of two dipoles and a singlet. Not shown are a triplet and two steerers on the main line and 10-deg line that are used to match the beam and direct it towards the AMIS target.

Optimizing beam transmission

A BO model was developed to maximize beam transmission using an initial data set of 20 randomly generated settings for a quad triplet, two dipoles and two steerers. Figure 6 shows the results for two attempts, both were successful improving the beam transmission from ~ 40 to 70%. 40% was the transmission obtained by a first tune by the operators. In the first attempt, the model found a glitch in the problem setup and took advantage of it. The first steerer happened to be before the first Faraday cup, and the model learned to optimize the transmission ratio by steering away the beam, thus minimizing the input current. In the second attempt, the first steerer was not included and the higher transmission was obtained for a higher beam current.

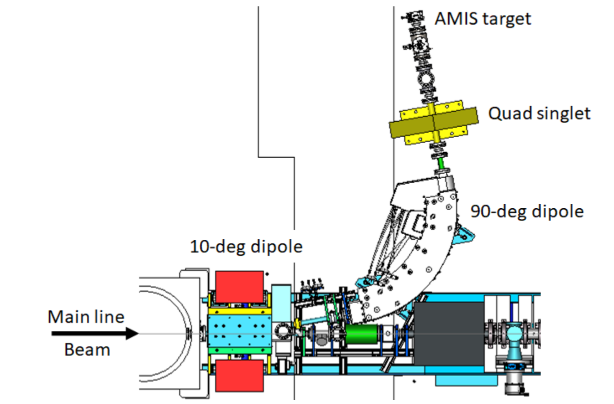

Figure 5: AMIS line, recently built, being commissioned.

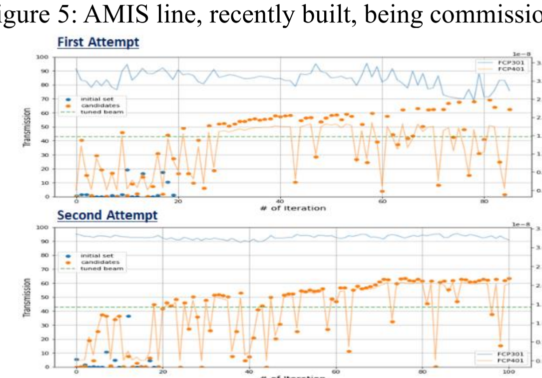

Figure 6: BO model results to maximize beam transmission through the AMIS line leading to 30% improvement.

Optimizing beam profiles

Following the optimization of beam transmission, we added the task of optimizing the beam profiles at a given location by tuning only a triplet and a steerer. The results are shown in Figure 7. The online training clearly shows a slow but steady convergence, top plot, with clear improvement in the beam profiles, lower plot.

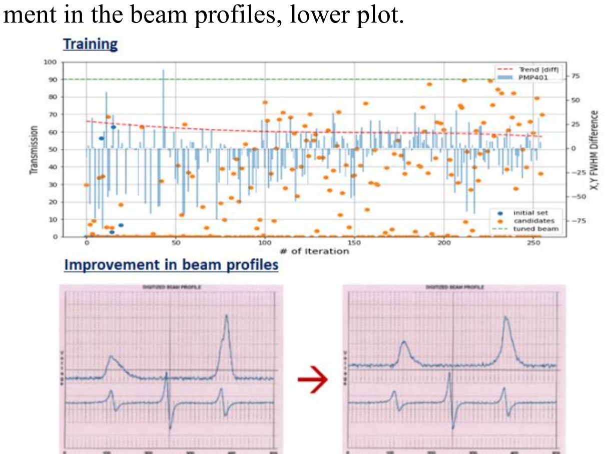

Figure 7: BO model results in optimizing the beam profiles. Model training showing slow but steady convergence– top. Improvements in beam profiles – bottom.

REFERENCES

[1] P. N. Ostroumov et al., “Completion of Efficiency and In-tensity Upgrade of the ATLAS Facility”, in Proc. LIN-AC’14, Geneva, Switzerland, Aug-Sep. 2014, paper TUPP005, pp. 449-451.

[2] G. Savard et al., “Radioactive beams from gas catchers: The CARIBU facility”, Nucl. Instrum. Meth. B, vol. 266, pp. 4086-4091, Oct. 2008. doi:10.1016/j.nimb.2008.05.091

[3] B. Mustapha et al., “An In-Flight Radioactive Ion Separator Design for the ATLAS Facility”, in Proc. LINAC’14, Ge-neva, Switzerland, Aug-Sep. 2014, paper TUPP004, pp. 446-448.

[4] B. Mustapha et al., “The ATLAS multi-user upgrade and po-tential applications”, J. Instrum., vol. 12, p. T12002, Dec. 2017. doi:10.1088/1748-0221/12/12/t12002

[5] M. W. McIntire, T. M. Cope, S. Ermon, and D. F. Rat- ner, “Bayesian Optimization of FEL Performance at LCLS”, in Proc. 7th Int. Particle Accelerator Conf. (IPAC’16), Busan, Korea, doi:10.18429/JACoW-IPAC2016-WEPOW055, 2016.

[6] V. Kain, S. Hirlander, B. Goddard, F. M. Velotti, G. Z. Del-la Porta, N. Bruchon, and G. Valentino, “Sample- efficient rein-forcement learning for cern accelerator control,” Phys. Rev. Accel. Beams, vol. 23, p. 124801, Dec 2020. [Online]. Avail-able: https://link.aps.org/doi/10.1103/PhysRevAccel-Beams.23.124801