ISBN: 978-3-95450-227-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2022-TUPOST054

EXPERIMENT OF BAYESIAN OPTIMIZATION FOR TRAJECTORY ALIGNMENT AT LOW ENERGY RHIC ELECTRON COOLER*

Y. Gao[1,†], K. A. Brown[1], X. Gu[1], J. Morris[1], S. Seletskiy[1]
W. Lin[2,‡], G. H. Hoffstaetter[1,2], J. A. Crittenden[2]
[1]Collider-Accelerator Department, Brookhaven National Laboratory, Upton, NY, USA
[2]CLASSE, Cornell University, Ithaca, NY, USA

Abstract

As the world's first electron cooler that uses radio frequency (rf) accelerated electron bunches, the low energy RHIC electron cooling (LEReC) system is a nonmagnetized cooler of ion beams in RHIC at Brookhaven National Laboratory. Beam dynamics in LEReC are different from the more conventional electron coolers due to the bunching of the electron beam. To ensure an efficient cooling performance at LEReC, many parameters need to be monitored and fine-tuned. The alignment of the electron and ion trajectories in the LEReC cooling sections is one of the most critical parameters. This work explores using a machine learning (ML) method - Bayesian Optimization (BO) to optimize the trajectories' alignment. Experimental results demonstrate that ML methods such as BO can perform control tasks efficiently in the RHIC controls system.

INTRODUCTION

The Low Energy RHIC electron Cooler (LEReC) is commissioned by the Collider-Accelerator Department (C-AD) at Brookhaven National Laboratory (BNL) to increase the collision rate [1] at the Relativistic Heavy Ion Collider (RHIC). During 2020 and 2021 runs, LEReC has proven to be successful in increasing the ion luminosity at RHIC.

Figure 1 shows the layout of the LEReC system at BNL. Electrons are generated by the gun and accelerated to 1.6 - 2 MeV by the superconducting rf cavity to match the energy of ions in RHIC. The electron bunches have a frequency of 704 MHz, and they are grouped into 30 - 36 macro-bunches with a frequency of 9 MHz. The accelerated electron bunches then travel through the transport line to interact with the ions in two cooling sections (CS) in the "Yellow" and "Blue" RHIC ring, each 20 meters long, connected by a 180-degree bend. Thus, ions in both rings of the collider can be cooled simultaneously. After interacting with the ions, the electrons are extracted from the Blue CS and discarded in the beam dump.

In the two cooling sections, ions experience a cooling force from the co-traveling electrons due to Coulomb interaction. As a result, the energy spread of the ion beam is reduced, and its phase-space density is increased [2, 3]. One of the key factors affecting the magnitude of the cooling force, hence affecting the cooling performance, is the relative angle of the electron beam trajectory with respect to the ion beam.

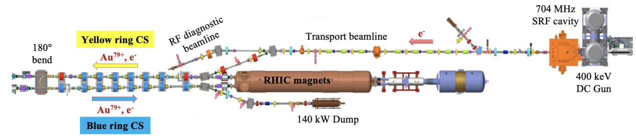

Figure 1: LEReC system layout (not to scale).

In this work, we use a machine learning (ML) method called Bayesian optimization (BO) to optimize the electron-ion trajectory alignment by matching the electron and ion positions observed at all of the beam position monitors (BPMs) in the cooling sections.

Bayesian optimization (BO) is a powerful tool for finding the optimum of an expensive objective function f with as few samples as possible. It contains two important components: a surrogate model, which learns and then predicts the behavior of the objective function, and an acquisition function, which determines quantitatively which inputs are most likely to generate optimal output. Bayesian optimization is applied in various fields, including finance, engineering, environmental science, and robotics. A basic BO background and some of its applications are detailed in [4].

To perform trajectory optimization at LEReC, preliminary simulation studies were conducted using both BO and another ML method called Q-Learning on the LEReC system simulator [5]. After obtaining satisfactory results from the simulation studies, we present in this paper the experimental results from testing the BO method on the real LEReC system. A more detailed description of the experiment, including simulation studies and BO algorithm structures, can be found in [6].

EXPERIMENTAL RESULTS

Preliminaries

Due to its high magnetic rigidity, the ion beam has a straight trajectory in the LEReC cooling sections. The lower-rigidity electron trajectory is controlled by 8 pairs of horizontal/vertical correctors and is monitored by 8 BPMs in each cooling section. Currently, there is an orbit correction program in place to manipulate correctors based on the BPM feedback, so electrons always stay at desired positions throughout the cooling sections. After calibration, the straight ion trajectory is kept at the center of the cooling

* Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy, and by the U.S. National Science Foundation under Award PHY-1549132.
† ygao@bnl.gov
‡ wl674@cornell.edu

ISBN: 978-3-95450-227-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2022-TUPOST054

Algorithm 1 Initial Sampling Routine

Require: Observation dataset $\mathscr{D}_{\mathscr{N}}$, random sampling function f_R, random sampling radius r, step size S, statistic period t_s.

Set $\mathscr{D}_{\mathscr{N}} = \emptyset$.
Set an anchor point $x_{ach} = -3$.
Set the initial operation $s_{op} =$ '+'.
for $t = 1, 2, \ldots, 40$ **do**
  if x_{ach} is outside of the range $[-3, 3]$ **then**
    Reverse x_{ach} to the previous value, flip $s_{op} = -s_{op}$.
  end if
  Randomly uniformly sample around the anchor point, $x_{new} = f_R(x_{ach}, r)$.
  Set the BPMs at x_{new}.
  Collect transverse beam size data during t_s, and calculate the cooling rate as $y_{new} = (1/\delta)(d\delta/dt)$.
  Add (x_{new}, y_{new}) to the dataset $\mathscr{D}_{\mathscr{N}}$.
  According to the s_{op}, modify the anchor point by a step size $x_{ach} = x_{ach} + / - S$.
end for

section. Therefore, the electrons should also be tuned to the center position $(x = 0, y = 0)$. Bayesian optimization (BO) offers an independent method to find the optimized electron trajectories. It serves both as an alternative approach and as a validation to the current beam-based alignment routine in case there are still errors in the BPM offsets.

The objective function for this experiment is the transverse cooling rate λ, defined as the decreasing speed of the transverse ion beam size δ. It is calculated as $\lambda = (1/\delta)(d\delta/dt)$. More negative λ means faster cooling, so the goal is to maximize $-\lambda$ by tuning electron positions with the orbit correction program.

This experiment only considers the first 4 BPM measurements due to the limited machine time. The BO algorithm is trained on 40 initial samples, in which the beam moves within a range of $[-3, 3]$ mm for all BPMs. The sample inputs span the entire range with a fixed step size, but a uniformly random noise is added to the designed sample value for each step size. This sampling routine ensures the algorithm learns the objective function behavior in an organized manner while also not getting redundant information with the added randomness. Algorithm 1 outlines the detailed procedure of the sampling process.

Figure 2 shows the initial sample inputs and the corresponding objective values. We can easily see the pattern of the output values in the bottom plot, thanks to the systemic iteration of the inputs. The higher cooling rates are generated by input positions close to 0, which is consistent with the original assumption.

Objective Function Sensitivity

During the experiment, we discovered that the ion beam size data[1] is noisy, as shown[2] in Fig. 3. This means when

[1] The ion beam size is measured by H-jet [7] and the plot shows the rms values

[2] Logged data displayed by system tool "LogView".

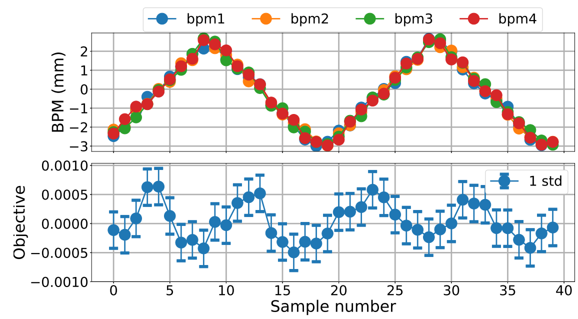

Figure 2: 40 training points sampled throughout the input range of $[-3, 3]$ mm using Algorithm 1.

calculating cooling rate $\lambda = (1/\delta)(d\delta/dt)$, the large noise present in point value δ (ion beam size) causes instability in the objective function. Figure 4 shows this phenomenon. For samples 7 to 11 and 13 to 17, we can see that the objective values change too quickly even when the inputs have converged to the optimal positions. As a result, the algorithm is confused and diverges because it does not learn the correct correlation between inputs and outputs.

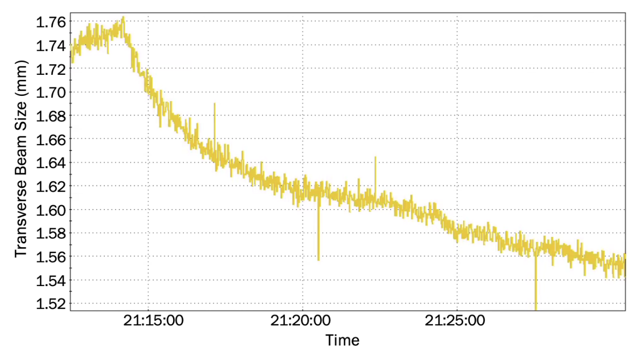

Figure 3: An example of the noise in real-time transverse ion beam size data during the experiment.

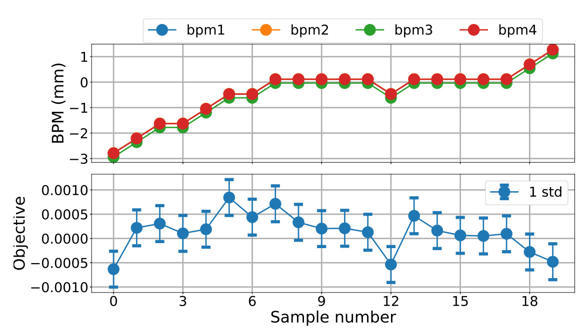

Figure 4: Demonstration of the volatility of the original objective function. The objective changes too fast (bottom) even when the algorithm converges to the optimal locations (from steps 7 to 11 or 13 to 17). This behavior eventually leads to the divergence of the algorithm (after step 17).

To smooth out the noise, we utilize a new parameter, called the number of averaging points, in our formula to calculate

ISBN: 978-3-95450-227-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2022-TUPOST054

the cooling rate. Instead of using point values, the cooling rate is now calculated by $\lambda = (1/avg(\delta))(avg(d\delta)/dt)$. Averaging provides stability to the objective behavior, making it easier for the algorithm to learn and converge. The number of ion beam size points δ to average is defined as the "Number N" of the BO algorithm, a parameter that controls the sensitivity of the objective function. Figure 5 demonstrates the different behaviors in the objective function for different choices of N values. The larger the N value is, the less sensitive the objective is.

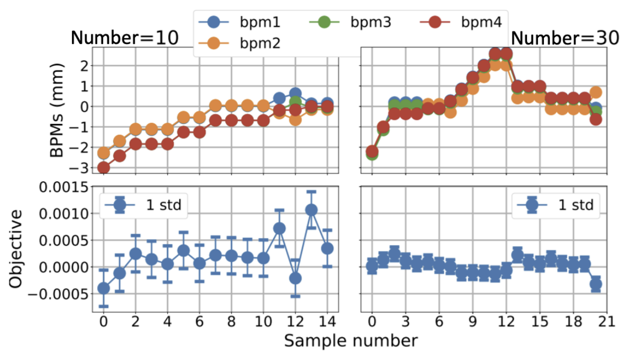

Figure 5: Effects of different numbers of averaging points on the objective's behaviors, which in turn affect the behavior of the BO algorithm.

Results

After experimenting with different N values, we pick a N value of 15, and the optimization results are shown in Fig. 6. The BO algorithm converges quickly, reaching a close neighborhood of the real optimum ($x = 0, y = 0$) in 3 steps. It is worth noting that the objective values (bottom plot) from the Bayesian samples (point 40 to 60) are lower than the ones from the first 40 initial samples due to averaging, but the process becomes more stable, as shown by the error bars. This feature helps the algorithm to converge quickly.

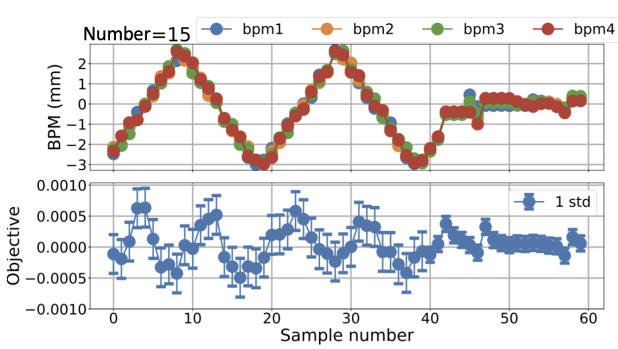

Figure 6: Final experiment results using averaging number of 15.

Figure 7 shows the electron trajectory logged during the optimization process. It indicates that the BO algorithm can tune electrons from the farthest positions (−3 mm) back to the optimal positions and maintain the trajectories. Furthermore, the optimal solution ($x = 0, y = 0$) found by the BO method agrees with and therefore validates the traditional orbit correction program and BPM calibrations.

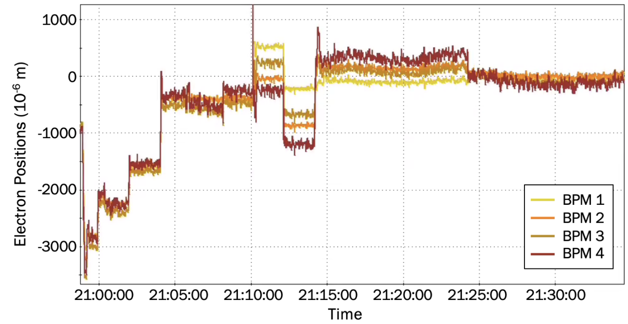

Figure 7: Electron beam trajectory data during optimization process.

FUTURE WORK

There are many areas of accelerator operation where machine learning techniques such as Bayesian optimization can be applied. One possible future project is applying Bayesian optimization to the Coherent electron Cooling (CeC) experiment at RHIC, which is a different electron cooler design being tested at BNL. Similar to LEReC, the CeC experiment also requires high-quality electron beams to interact with and cool ions in RHIC. Currently, the CeC system still requires time-consuming manual tuning by experienced operators to obtain desirable electron beams with efficient qualities. Bayesian optimization can make this tuning process automatic and much faster. There are many promising variants of Bayesian optimization that can be experimented with. Two such variants, physics model-informed (PMI) BO and Contextual Gaussian Process (CGP), are discussed in [6]. PMI BO is faster than the general data-informed BO by incorporating a physics data model, and BO with GCP can optimize dynamic systems with varying environmental factors. Another BO variant worth exploring for future projects is multi-objective BO (MBO). Since MBO can optimize multiple objective functions simultaneously, this would be extremely useful for beam optimization in particle accelerators as multiple beam parameters (i.e., peak current, slice emittance, and slice energy spread) usually need to be optimized together. Studies of MBO applications are also ongoing at other accelerator facilities. One recent approach is presented in [8].

CONCLUSIONS

In this work, we apply a machine learning technique called Bayesian optimization to maximize the cooling rate in the LEReC system. Experimental results show that the BO algorithm effectively finds and maintains electron positions for optimal cooling performance. This success opens up possibilities of trying various machine learning methods to improve operations in the RHIC complex, as well as in the future Electron-Ion Collider (EIC).

ISBN: 978-3-95450-227-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2022-TUPOST054

REFERENCES

[1] K. M. Walsh, "Electron Bunches Keep Ions Cool at RHIC," https://www.bnl.gov/newsroom/news.php?a=215585/

[2] G. I. Budker, "An effective method of damping particle oscillations in proton and anti-proton storage rings," *Sov. Atom. Energ.*, vol. 22, pp. 438–440, 1967. doi:10.1007/BF01175204

[3] Y. S. Derbenev, "Theory of electron cooling," Mar. 2017. doi:10.48550/arXiv.1703.09735

[4] B. Shahriari, K. Swersky, Z. Wang *et al.*, "Taking the human out of the loop: A review of bayesian optimization," *Proceedings of the IEEE*, vol. 104, no. 1, pp. 148–175, 2016. doi:10.1109/JPROC.2015.2494218

[5] Y. Gao, K. A. Brown, P. S. Dyer, S. Seletskiy, and H. Zhao, "Applying Machine Learning to Optimization of Cooling Rate at Low Energy RHIC Electron Cooler", in *Proc. IPAC'21*, Campinas, Brazil, May 2021, pp. 3391–3394. doi:10.18429/JACoW-IPAC2021-WEPAB306

[6] Y. Gao, W. Lin *et al.*, "Bayesian optimization experiment for trajectory alignment at the low energy rhic electron cooling system," *Phys. Rev. Accel. Beams*, vol. 25, no. 1, p. 014 601, 2022. doi:10.1103/PhysRevAccelBeams.25.014601

[7] A. Zelenski, A. Bravar, D. Graham *et al.*, "Absolute polarized h-jet polarimeter development, for RHIC," *Nucl. Instrum. Methods Phys. Res., Sect. A*, vol. 536, no. 3, pp. 248–254, 2005. doi:10.1016/j.nima.2004.08.080

[8] R. Roussel, A. Hanuka, and A. Edelen, "Multiobjective bayesian optimization for online accelerator tuning," *Phys. Rev. Accel. Beams*, vol. 24, no. 6, p. 062801, 2021. doi:10.1103/PhysRevAccelBeams.24.062801