Time-aware Bayesian optimization for adaptive particle accelerator tuning

Nikita Kuklev
Advanced Photon Source
Argonne National Laboratory
Lemont, IL 60439
nkuklev@anl.gov

Yine Sun
Advanced Photon Source
Argonne National Laboratory
Lemont, IL 60439

Hairong Shang
Advanced Photon Source
Argonne National Laboratory
Lemont, IL 60439

Michael Borland
Advanced Photon Source
Argonne National Laboratory
Lemont, IL 60439

Gregory Fystro
Advanced Photon Source
Argonne National Laboratory
Lemont, IL 60439

Abstract

Particle accelerators require continuous adjustment to maintain beam quality. At the Advanced Photon Source (APS) synchrotron facility this is accomplished using a mix of operator-controlled and automated tools. We have recently implemented Bayesian optimization (BO) as one of automated options, significantly improving sampling efficiency. However, poor BO performance was observed in certain scenarios due to time-dependent device drifts. In this work, we discuss extending BO to an adaptive version (ABO) that can compensate for distribution drifts through explicit time-awareness, enabling long-term online operational use. Our contributions include advanced kernels with physics-informed time dimension structure, age-biased data history subsampling, and auxiliary time-aware safety constraint models. Benchmarks show better ABO performance in several simulated and experimental tests. Our results are an encouraging step for the wider adoption of ML-based optimizers at APS.

1 Introduction

Modern particle accelerators face increasing performance demands, resulting in tighter tolerances on accuracy and stability [1]. Due to cost, physical limits, and external factors, some amount of continuous parameter adjustment is constantly required. Historically, this tuning required expert guidance and intuition, with software tools only allowing for a partial automation. In critical cases, custom automated solutions were implemented, such as for beam orbit, energy, and tune feedback in the Large Hadron Collider [2]. With the rise of machine learning, there is immense interest in making use of the newly available algorithms to implement generic tools to improve reliability, reduce expert workload, and provide higher performance to the users.

A key application of ML for accelerators is in parameter optimization, whereby one or multiple objectives are tuned through an intelligent search of the parameter space. A number of conventional optimization methods are already in use, including simplex[3, 4], RCDS [5], genetic algorithms [6], extremum seeking [7], and several others. New ML methods include Bayesian optimization (BO) [8], reinforcement learning [9], and others. BO is of special interest since it allows efficient black-box function optimization with few samples, taking advantage of any prior physics model knowledge provided to the algorithm. This work first reviews the basic BO process, and then discusses our

Machine Learning and the Physical Sciences workshop, NeurIPS 2022.

contributions - a set of improvements that permits for continuous, robust, and adaptive BO use for optimizing time-varying systems.

2 Time-aware Bayesian Optimization

In standard BO process, system output is described by

$$\mathbf{y} = f(\mathbf{x}) + \varepsilon \tag{1}$$

where $f(\mathbf{x})$ is the black-box function of interest and $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ the added noise. Vector $\mathbf{x}$ (bounded to domain X) has dimension of $n \,\mathrm{x}\, d$ where d is the parameter space size and n the number of measurements. Using Gaussian Process (GP), a surrogate model for f can be parameterized through a multivariate normal distribution with a mean $m(\mathbf{x})$ and covariance kernel $k(\mathbf{x}, \mathbf{x}')$ as

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')) \tag{2}$$

$$k(\mathbf{x}, \mathbf{x}') = \mathrm{Cov}\left[f(\mathbf{x}), f(\mathbf{x}')\right] \tag{3}$$

The kernel is used to evaluate the similarity between values of f at $\mathbf{x}$ and $\mathbf{x}'$, and its' appropriate choice is critical for good GP convergence. Existing knowledge about the system can be encoded through prior distributions on kernel and mean, with the distribution parameters called *hyper-parameters*. During model fitting hyper-parameters are updated using Bayes' rule (conditioned on observed data) and posterior probability distribution $\mathrm{p}(\mathbf{f} \mid \mathbf{y}, \mathbf{x})$ can then be sampled to get model predictions [10]. BO evaluates a special 'acquisition' function over a fitted GP model so as to predict the best next location(s) to sample. A variety of analytic and Monte-Carlo acquisition functions exist, with one of simplest being the upper confidence bound (UCB)

$$\mathrm{UCB}(\mathbf{x}) = \mu(\mathbf{x}) + \sqrt{\beta} * \sigma(\mathbf{x}) \tag{4}$$

where mean μ and variance σ are provided by the GP model. The parameter β allows for trade-off between exploration (taking risks for high reward) and exploitation (use of known good configurations).

2.1 Time-aware GP model

In accelerator physics, input parameter vector $\mathbf{x}$ represents physical device settings, for example magnet currents, or a physics-informed combination of device inputs called a 'knob'. Standard BO optimization treats all inputs equally (but with potentially different lengthscales) by using multidimensional isotropic kernels, such as Matérn or radial basis function (RBF) [10]. Furthermore, starting samples are randomized and only freshly collected data is used. To improve convergence speed, previous work has successfully used historic data to initialize kernel covariance distributions [11]. Such pre-training works well when experimental conditions are reproducible. However, some accelerators also have undesired and poorly modelled time-dependent drifts. In standard BO, such drifts smear out the outputs, effectively increasing noise levels and producing suboptimal steady-state solutions. Several model-free adaptive approaches have been previously attempted [1], including extremum seeking and adaptive neural networks. Our work seeks instead to use an explicitly time-aware model as part of BO algorithm [12], forming adaptive BO (ABO). Formally, we extend f with explicit time dimension t, such that the system is now described by

$$\mathbf{y} = f(t, \mathbf{x}) + \varepsilon \tag{5}$$

In the simplest implementation, above-mentioned isotropic kernels can be extended to cover time dimension. The only BO algorithm change is in special treatment of acquisition function, where the next sample time t_{next} has to be predicted based how long ABO takes to run and how soon next experimental sample can be measured. Acquisition function is optimized at that fixed future time

$$\mathbf{x}_{next} = \underset{\mathbf{x} \in X}{\arg\max}\, \mathrm{UCB}(t_{next}, \mathbf{x}) \tag{6}$$

However, time effects often have different patterns compared to the other inputs, and it is desirable to encode this physics knowledge into the GP model. Our approach achieves this through composition of specialized sub-kernels. Kernel multiplication and addition can be thought of as logical AND and OR operations along any shared dimensions, and as independent operations on other dimensions. For example, multiplication of periodic and RBF kernels produces a 'locally periodic' correlation - a periodic (not necessarily sinusoidal) function that can slowly change shape:

$$k_{lp} = \sigma^2 \exp\left(\frac{2\sin^2\left(\pi|t-t'|/p\right)}{l^2}\right)\exp\left(\frac{-(t-t')^2}{2l^2}\right) \tag{7}$$

Hyper-parameters of this kernel are output variance σ, period p, and lengthscale l. In experimental applications, some intuition about the underlying drift processes can be gained from historical data but precise values for (hyper)parameters like number of periodic signals are difficult to specify ahead of time. An example of experimental beam position data from the APS linear accelerator (linac) is shown in Fig. 1, demonstrating long-term monotonic drift as well as two distinct periodic signals. For unknown reasons, one of the periodic signal components changes in strength day to day, demonstrating the need for not only time-aware but also time-adaptive methods. Also, note how one of the patterns exhibits a significantly larger noise level, which is yet another time-dependent linac behavior. Based on historical data, noise levels do not appear to be correlated with the oscillation pattern changes, and are fitted through the noise hyperparameter in all models.

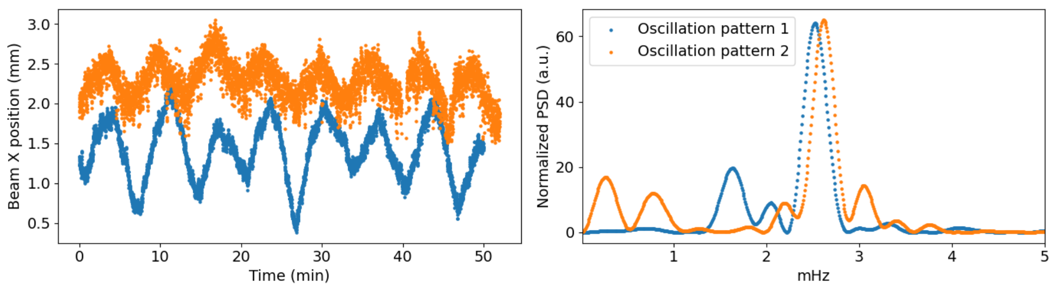

Figure 1: Linac beam position drift and corresponding signal power spectrum density when all feedback is disabled, showing two distinct types of signals observed on different days.

In BO, adaptive kernels fall under topic of automatic kernel selection and deep kernel learning, with several promising recent results [13]. In this work we chose standard spectral mixture (SM) kernel [14]. SM can approximate any stationary kernel, including ones with oscillatory and local correlations, and can be physically interpreted via kernel spectral density. It is also the basis for more complex deep kernels, which we plan to study in the future. To account for non-stationary correlations such as monotonic drifts, a linear kernel can be added when necessary. For spatial dimensions, we use a standard Matérn kernel ($\nu = 2.5$) with automatic relevance determination enabled. Output scale parameters are added to all components except SM. Thus, our final ABO-SM kernel is given by

$$k(t, t', \mathbf{x}, \mathbf{x}') = \left(k_{SM}(t, t') + \sigma_l^2 k_{linear}(t, t')\right) \times \sigma^2 k_{Matern}(\mathbf{x}, \mathbf{x}') \tag{8}$$

Our ABO code was implemented with BoTorch/GPyTorch libraries, using Monte Carlo samplers and acquisition functions, as well as standard fitting routines with defaults changed for higher fidelity. Evaluations were performed on NVIDIA A100 GPU, with single sample latency below 10s ($\sim$ 600s per run).

3 Results

3.1 Simulations

We first tested both Matérn (ABO-ISO) and spectral mixture (ABO-SM) time kernel versions on several simulated problems, ranging from synthetic functions to linac particle tracking simulations. Figure 2 shows an example with sinusoidal corrector drift that produces downstream beam oscillations. After 20 initial points observed without any input changes (blue line), methods were allowed to run with identical settings. After approximately a single period, ABO-SM converges on the oscillation frequency and starts following the drift without any lag, demonstrating that it is in fact predicting the future location via the model. ABO-ISO tracks changes but has notable lag, resulting in slight deviations away from the optimum point. Standard BO meanwhile sees the drift as noise and converges to central suboptimal input value, losing significant performance.

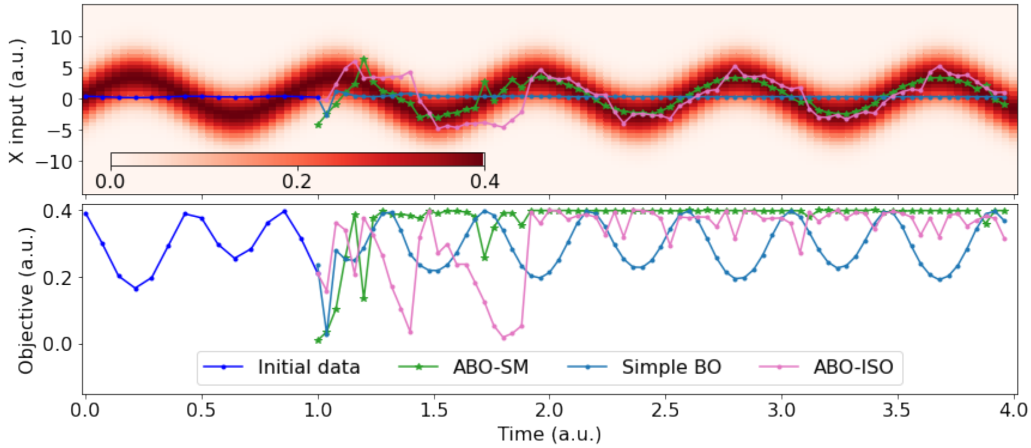

Figure 2: ABO-SM, ABO-ISO, and standard BO results for a drift simulation with a sinusoidal perturbation. Top plot shows input parameter values at each step, with colormap denoting the objective function landscape. Bottom plot shows the objective value observed by each method (i.e. output for the input of the top plot), with objective value of 0.4 being the true maximum at all times.

3.2 Experimental tests

Experimental tests were done at the APS injector. It consists of a linac [15] and two rings [16, 17] that bring the electrons to an energy of $7\,\mathrm{GeV}$, to be injected into the storage ring. To stabilize beam parameters, several proportional feedback controllers, called control laws [18], are used. They operate with pre-computed inverse response matrices derived from experimental data. A recent analysis of beam parameters noted elevated high frequency jitter levels with control laws enabled, potentially caused by beam position monitor (BPM) noise or deviations from expected lattice parameters. With control laws off however, slower but larger amplitude oscillations were observed, as shown in Figure 1. While not large enough to impact overall injector efficiency, eliminating both the slow drifts and the high frequency noise is desirable for downstream experiments. We tested ABO for that purpose by picking a suitable BPM and corrector pair at the end of the linac, and using as objective the absolute value of the trajectory error away from setpoint. Results of this test are given in Figure 3.

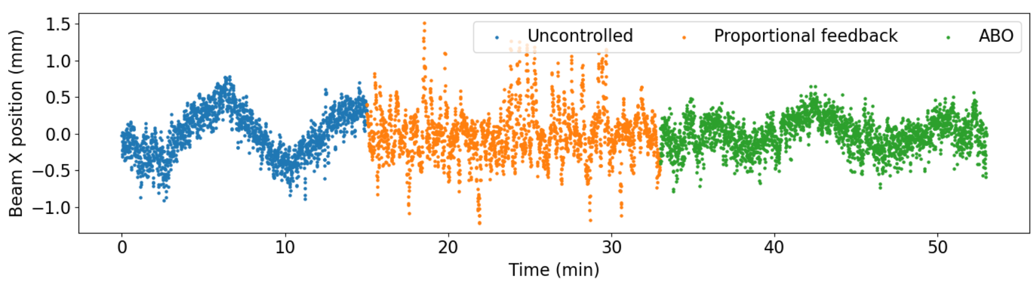

Figure 3: No input, proportional feedback, and ABO-SM trajectories in the experimental study.

For the particular dataset shown, beam position exhibited the ‘two frequency’ regime described previously and ABO-SM with 3 mixture components was used (additional components required notably more points for a stable fit). Previous 15 minutes of history were used as the training set for each run, subsampled to 100 points to reduce noise. Results show that ABO was able to remove the main oscillatory signal at $2.5\,\mathrm{mHz}$, with corresponding RMS jitter lowered to $0.23\,\mathrm{mm}$ as compared to $0.35\,\mathrm{mm}$ for control law and $0.33\,\mathrm{mm}$ for uncontrolled cases. Simulations indicate that the remaining oscillatory component could have be modelled robustly with longer data collection time, but convergence speed is strongly dependent on noise levels.

3.3 Limitations

Above results demonstrate promising ABO performance, and we aim for it to eventually become a production-ready robust stabilization method for APS accelerators. To that end, we incorporated and are testing additional features such as time-aware safety constraint modelling (extending [19]) and a digital twin test harness. Main fundamental limitation of all BO methods is poor scaling with number of dimensions and number of points ($\mathcal{O}(n^3)$ for exact GP, although better scaling algorithms [20] are faster for large datasets). Given real-time nature of accelerator optimization, this limits tractable problems to those with fewer than ~ 20 parameters. We presently mitigate dataset growth by intelligently subsampling based on expected data and drift bandwidths, and are exploring recent work on high-dimensional spaces [21] to improve evaluation performance.

4 Conclusion

Improving robustness and adaptability of ML-based optimizers is a crucial step in making these tools useful in day-to-day accelerator operation. We have demonstrated that BO can be usefully modified with the explicit addition of the time dimension so as to fit a wide variety of experimentally-relevant distribution drifts. With a special spectral mixture kernel, ABO can applied to problems with unknown or varying number of periodic and local correlations. Results on both simulated and experimental tasks, while not perfect, significantly improve on standard BO performance. Future work will focus on improving GP hot-start with pre-training on more varied historical data, and exploring DKL neural networks for further kernel model improvements.

5 Broader Impact

Particle accelerators have often been called the engines of discovery, since their advances underpin much of the progress in other fields like particle physics, material science, structural biology, and beyond. Synchrotrons in particular each serve dozens of experimental beamlines, with user communities comprising hundreds of scientists. Our work will contribute to the improved performance and reliability of the APS accelerator systems, allowing for more beam time to be delivered to users and thus improving their scientific output. It is hard to speculate as to the effects of our work beyond accelerator community. ABO is an in-place upgrade for any process where standard Bayesian optimization can be used, and expands BO applicability to strongly drifting systems. It will thus increase the impact (positive or negative) of the underlying process, but on its own, does create any biases or other issues.

Acknowledgments and Disclosure of Funding

Authors acknowledge ANL Laboratory Computing Resource Center (LCRC) for providing the computational resources. We also thank APS operators for their help in late-night debugging and data collection.

This work is supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No. DE-AC02-06CH11357.

References

[1] A. Scheinker, "Adaptive machine learning for time-varying systems: Low dimensional latent space tuning," *Journal of Instrumentation*, vol. 16, no. 10, P10008, 1, 2021, `10.1088/1748-0221/16/10/P10008`

[2] L. Grech, G. Valentino, D. Alves, and S. Hirlaender, "Application of reinforcement learning in the lhc tune feedback," *Frontiers in Physics*, vol. 10, 2022, `10.3389/fphy.2022.929064`

[3] L. Emery, M. Borland, and H. Shang, "Use of a General-Purpose Optimization Module in Accelerator Control," in *Proc. PAC'03*, Portland, OR, USA, May 2003, `https://jacow.org/p03/papers/WPPE002.pdf`

[4] X. Huang, "Robust simplex algorithm for online optimization," *Phys. Rev. Accel. Beams*, vol. 21, p. 104 601, 10 2018, `10.1103/PhysRevAccelBeams.21.104601`

[5] X. Huang, J. Corbett, J. Safranek, and J. Wu, "An algorithm for online optimization of accelerators," *Nuclear Instruments and Methods in Physics Research. Section A, Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 726, 2013, `10.1016/j.nima.2013.05.046`

[6] M. Borland, V. Sajaev, L. Emery, and A. Xiao, "Multi-objective direct optimization of dynamic acceptance and lifetime for potential upgrades of the Advanced Photon Source.," ANL/APS/LS-319, 986303, 24, 2010, ANL/APS/LS–319, 986 303, `10.2172/986303`

[7] A. Scheinker, "Adaptive Machine Learning for Robust Diagnostics and Control of Time-Varying Particle Accelerator Components and Beams," *Information*, vol. 12, no. 4, p. 161, 4 2021, `10.3390/info12040161`

[8] M. Balandat *et al.*, "BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization," in *Advances in Neural Information Processing Systems*, vol. 33, 2020, pp. 21 524–21 538, `https://proceedings.neurips.cc/paper/2020/hash/f5b1b89d98b7286673128a5fb112cb9a-Abstract.html`

[9] V. Francois-Lavet, P. Henderson, R. Islam, M. G. Bellemare, and J. Pineau, "An Introduction to Deep Reinforcement Learning," *Foundations and Trends® in Machine Learning*, vol. 11, no. 3-4, pp. 219–354, 2018, `10.1561/2200000071`

[10] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. 23, 2005, `10.7551/mitpress/3206.001.0001`

[11] J. Duris *et al.*, "Bayesian Optimization of a Free-Electron Laser," *Physical Review Letters*, vol. 124, no. 12, p. 124 801, 25, 2020, `10.1103/PhysRevLett.124.124801`

[12] F. M. Nyikosa, M. A. Osborne, and S. J. Roberts, *Bayesian optimization for dynamic problems*, 2018, `10.48550/ARXIV.1803.03432`

[13] A. G. Wilson, Z. Hu, R. Salakhutdinov, and E. P. Xing, *Deep Kernel Learning*, 6, 2015, `10.48550/arXiv.1511.02222`

[14] A. G. Wilson and R. P. Adams, *Gaussian Process Kernels for Pattern Discovery and Extrapolation*, 31, 2013, `http://arxiv.org/abs/1302.4245`

[15] A. Nassiri, G. Mavrogenes, and W. Wesolowski, "The Linac Injector for the ANL 7 GeV Advanced Photon Source," in *Proc. LINAC'90*, Albuquerque, NM, USA, Sep. 1990, pp. 611–613, `https://jacow.org/l90/papers/TH411.pdf`

[16] M. Borland, "Commissioning of the Argonne Positron Accumulator Ring," in *Proc. PAC'95*, Dallas, TX, USA, May 1995.

[17] S. V. Milton, "The APS Booster Synchrotron: Commissioning and Operational Experience," in *Proc. PAC'95*, Dallas, TX, USA, May 1995.

[18] L. Emery, "Commissioning Software Tools at the Advanced Photon Source," in *Proc. PAC'95*, Dallas, TX, USA, May 1995.

[19] D. Eriksson and M. Poloczek, *Scalable Constrained Bayesian Optimization*, 28, 2021, `10.48550/arXiv.2002.08526`

[20] G. Pleiss, J. Gardner, K. Weinberger, and A. G. Wilson, "Constant-time predictive distributions for Gaussian processes," in *Proceedings of the 35th International Conference on Machine Learning*, vol. 80, 2018, pp. 4114–4123, `https://proceedings.mlr.press/v80/pleiss18a.html`

[21] S. Daulton, D. Eriksson, M. Balandat, and E. Bakshy, "Multi-objective bayesian optimization over high-dimensional search spaces," in *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence*, 2022.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [Yes]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No] We ran experiments in an actual particle accelerator facility, which is unique and difficult to reproduce only with simulations. Our code is currently restricted to internal use, but we plan to open-source final version.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] Partial, rest can be obtained from library defaults
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] Partial
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]