13th Int. Particle Acc. Conf. IPAC2022, Bangkok, Thailand JACoW Publishing
ISBN: 978-3-95450-227-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2022-TUPOST055

TOWARD MACHINE LEARNING-BASED ADAPTIVE CONTROL AND GLOBAL FEEDBACK FOR COMPACT ACCELERATORS*

F. Cropp†, P. Musumeci
Department of Physics and Astronomy, University of California Los Angeles,
Los Angeles, CA 90095, USA
A. Scheinker
Applied Electrodynamics Group, Los Alamos National Laboratory,
Los Alamos, NM 87545, USA
D. Filippetto, A. Gilardi, S. Paiagua, D. Wang
Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

Abstract

The HiRES beamline at Lawrence Berkeley National Laboratory (USA) is a state-of-the-art compact accelerator providing ultrafast relativistic electron pulses at MHz repetition rates, for applications in ultrafast science and for particle accelerator science and technology R&D. Using HiRES as testbed, we seek to apply recent developments in machine learning and computational techniques for machine learning-based adaptive control, and eventually, a full control system based on global feedback. The ultimate goal is to demonstrate the benefits of such a suite of controls to UED, including increased temporal and spatial resolution. Concrete steps toward these goals are presented, including automatic, model-independent tuning for accelerators, and energy virtual diagnostics with direct application to improving UED temporal resolution.

INTRODUCTION

Because of the complexity of accelerators and their natural parameter movement, both in the short and long term (hereafter, jitter and drift, respectively), accelerators could see major improvement from machine learning (ML). The application of ML has been shown to solve or mitigate a plethora of accelerator control and diagnostic problems, for example, for navigating efficiently the multi-dimensional parameter space to find control set points [1, 2], for inverting a large parameter space to make a parasitic diagnostic [3, 4], or for non-destructive virtual diagnostics [5, 6]. Further, ultrafast electron diffraction (UED) has benefited from ML-based static models and virtual diagnostics [7, 8]. The advantages of static ML models are incontrovertible, but they have limits. For example, it is an open question as to how an optimal ML-based control system should treat a system when parameter drift brings a the state of the machine outside the training set [9]. At HiRES [10], a state-of-the-art MHz-class UED facility with short, high 6-D brightness beams, a model-independent optimization method is detailed for dealing with this case. Further work is shown to apply time of arrival virtual diagnostics to increase the short- and long-term stability of the already-state-of-the-art stability at HiRES to the 10^{-4} level and below. Such work will make up the building blocks of an adaptive control and global feedback system, with application to UED measurements.

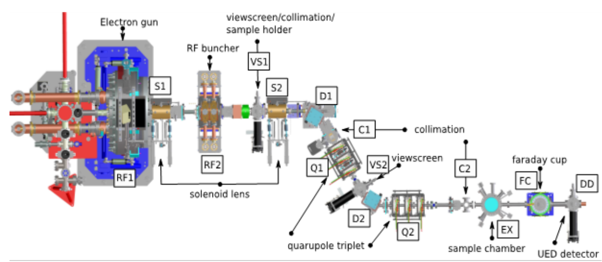

Figure 1: The HiRES beamline, from [10].

INITIAL EXPERIMENTAL TESTS OF ML-ENHANCED OPERATIONS

The ultimate goal is to demonstrate the benefits of such a suite of controls to UED, including increased temporal and spatial resolution. The inital experimental results, showcasing the potential for enhanced stability and autonomous operations, are shown.

Adaptive Control

Extremum seeking (ES) [11–13] is a powerful, model-independent optimization routine that can be applied to optimize quickly and control particle accelerators. It has myriad uses for particle accelerators, including optimizing an electron beam via automatic tuning of accelerator parameters [14] to tuning the latent space of a convolutional neural network-based digital twin of the accelerator to make it more robust to drifting parameters [9]. This work seeks to combine these approaches to build a tool for adaptive, on-line control of accelerators.

Often, when operating an accelerator, it is necessary to change modes of operation. At HiRES, ES was demonstrated to be able to follow a moving cost function. In this example,

* This work was supported by the DOE Office of Science Graduate Student Research (SCGSR) program, by the DOE Office of Basic Energy Sciences under Contract No. DE-AC02-05CH11231, by the DOE Office of Science, Office of High Energy Physics under contract number 89233218CNA000001 and DE-AC02-05CH11231 and by the NSF under Grant No. PHY-1549132.

† ericcropp@physics.ucla.edu

ISBN: 978-3-95450-227-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2022-TUPOST055

data was taken on VS2 (as shown in Fig. 1). The quadrupoles in Q1 were the parameters on which ES was optimizing. To find a starting point, a coarse grid quadrupole scan was performed and interpolated using a four-layer feed-forward neural network. Inputs were the three quadrupole settings. Outputs were the root-mean-square transverse beam sizes. Minimization of the NN of the following cost function was performed:

$$f(q_1, q_2, q_3) = |x_{rms}| + |y_{rms}| + |x_{rms} - y_{rms}| \qquad (1)$$

to find the optimal quadrupole settings to make a small, round beam. Using the resultant point as a starting point, a sinusoidally varying cost function was introduced, as shown in Fig. 2. As can be seen, ES keeps the cost function low by varying the beam size as necessary.

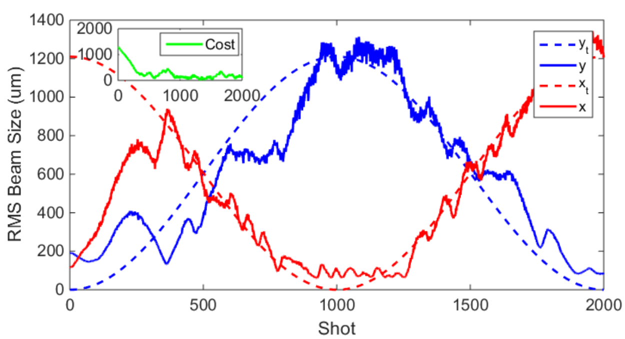

Figure 2: ES minimizing a variable cost function: RMS beam sizes and corresponding targets. Inset: cost function minimization.

Another situation that often requires operator expertise is the compensation of large parameter drifts. In fact, severe enough drifts for a static, ML-based system will require operator intervention, if the system drifts outside of the range of the training set. As a trial, at HiRES, such a severe drift was induced by moving upstream parameters and was compensated by means of the ES algorithm. The data for this experiment were acquired using a PI-MAX 4 intensified camera at approximately 2 Hz on DD (see Fig. 1). The second quadrupole of the first quadrupole triplet in the dogleg (Q1 in Fig. 1) was varied in a sinusoidal pattern. As can be seen from Fig. 3, ES was able to keep the cost function (Eq. 1) minimized by only varying the three quadrupoles in the second triplet. The nominal case of keeping the three quadrupoles constant is also shown. It is worth noting that beam sizes in the both dimensions remained small, so long as the rate of change of the drift was sufficiently small.

ES is shown to be a powerful tool for automatic, model-independent optimization of a changing system, in two cases: 1) if the system needs to change, and 2) if the system changes, but the objective stays the same. As such, this will be a powerful building block for a ML-based control system if the system changes outside the training of the model.

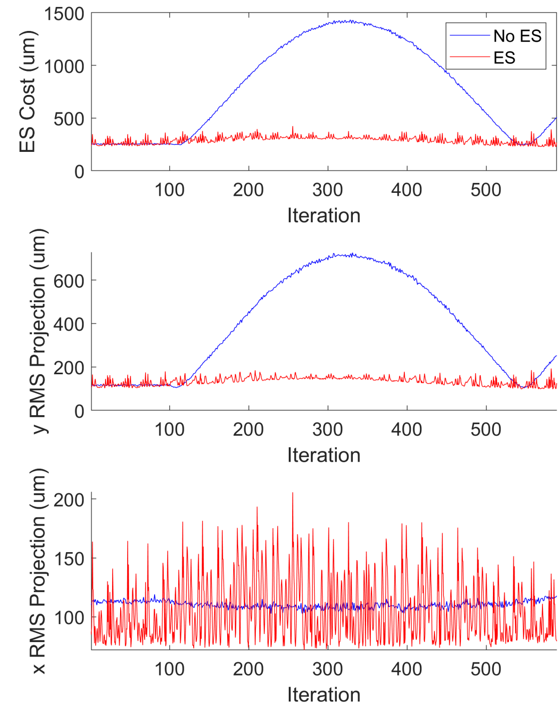

Figure 3: ES minimizing a static cost function (Eq. 1) in variable, drifting conditions. Top: Cost function with and without ES feedback. Middle: RMS beam size in y-dimension with and without ES. Bottom: RMS beam size in x-dimension with and without ES. Note that the feedback does not sacrifice beam size in x to make y smaller.

Increasing Stability Beyond State-of-the-Art

Beam stability is critical to consistent and ideal beam operation. For example, for UED, the temporal resolution of the pump-probe experiment is

$$\tau = \sqrt{\Delta t_{e^-}^2 + \Delta t_{laser}^2 + \Delta t_{jitter}^2 + \Delta t_{VM}^2}, \qquad (2)$$

where Δt_{e^-} is the electron bunch length, and Δt_{laser} is the laser pulse length. These quantities can be reduced using a bunching cavity and laser compressor, respectively. Δt_{VM} is the velocity mismatch term, which can be neglected for thin samples. Δt_{jitter} is the time-of-arrival jitter between the laser pulse and electron bunch, which can be reduced using feedback methods and what this work seeks to reduce further.

HiRES already has near state-of-the-art stability, which approaches $\Delta E/E$ stability of $5x10^{-5}$ at short time scales, when feedback is turned on, thanks to the 102 MSPS FPGA-enabled feedback. Drifts on the order of minutes to hours, however, inflate the overall relative energy variation to approach $4x10^{-4}$. In general, it is desired to remove both jitter and drift, but different models for different timescales may be required.

The unknown aspects of parameter drift make for an ideal case for a non-static model; sparse models, where time information is not encoded in the model, may not perform

13th Int. Particle Acc. Conf. IPAC2022, Bangkok, Thailand JACoW Publishing
ISBN: 978-3-95450-227-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2022-TUPOST055

well in a case that is dominated by such a drift. Time series analysis is a field that has applications to changing systems, from scientific (e.g. [15, 16]) to stock markets (e.g. [17, 18]). In time-series analysis, future targets are predicted by present and/or past predictors and the order matters. The decision as to whether to include predictors from the current shot is dependent on multiple factors (see Fig. 4), including prediction speed, and prediction accuracy without including the current shot.

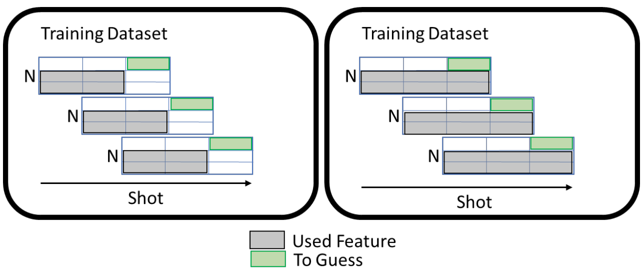

Figure 4: Left: only past predictors (gray) are used to predict the present target (green). Right: Present predictors are included in the prediction of the present target.

Another consideration is whether to use past targets in the prediction of the next target. This can be particularly helpful in the case where unknown parameter drift is driving a change that is not visible in the predictors. However, for a virtual diagnostic, in the testing set, previous measurements of the target are not available. One technique is to include previous predictions from the virtual diagnostic in the next prediction (See Fig. 5).

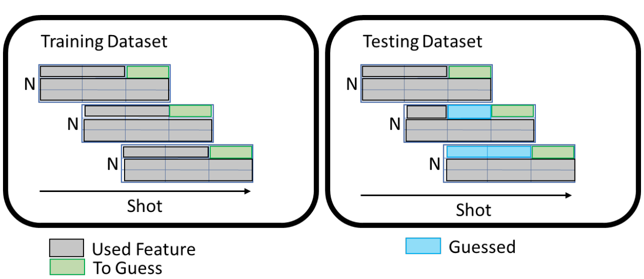

Figure 5: Left: training dataset where past predictors and measurements (gray) are used to predict the present target (green). Right: Testing dataset where past predictions replace measurements where necessary (blue).

As a first test, sparse prediction of beam x-position on VS2 in the dogleg (see Fig. 1) was attempted using multiple linear regression. The final 10% of shots were held in reserve for a test dataset. See the inset of Fig. 6 for the results of the regression. Note that on an hours-long timescale, the RMSE in the test dataset is more than a factor of two smaller than the fluctuations with the hardware feedback turned on.

The virtual diagnostic uses two cavity probes (amplitude and phase), the forward (amplitude and phase) and reverse (amplitude and phase) RF power, the laser phase and the laser position on a virtual cathode camera to predict the position on VS2.

Due to data save rates and camera acquisition times, RF waveforms and beam images are acquired at 1 Hz, downsampled from the high repetition rates of the HiRES facility. Extra care must be taken to ensure that the downsampling results in synchronized data (i.e. the RF and camera shots must align). Further, alignment on the 1 Hz level must also be checked; save delays can exceed 1 second.

Because of the stable nature of a CW-type gun, multiple beams can be produced from the gun for each RF pulse, at the laser repetition rate of 250 kHz. Considering the signal-to-noise ratio of the images, up to 10 beams were averaged for each image. These beams are each only 4us apart, and even short-term jitter is expected to be on on a timescale longer than 100us, given the quality factor of the gun.

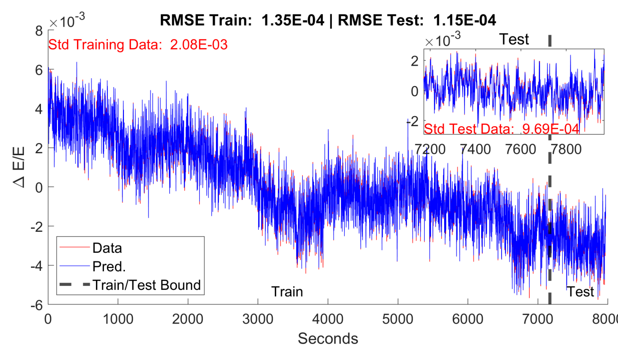

Figure 6: Linear regression prediction with hardware feedback off.

The 10^{-4} RMSE shows much potential for improving stability; one could reduce the overall temporal jitter using this method: 1) by using the virtual diagnostic for feedback or 2) for UED, by using the virtual diagnostic to order shots by time of arrival. Improved modeling, time-series prediction and reducing noise in measurements is planned and promises even better results.

CONCLUSION

In this work, concrete steps toward a novel adaptive control and global feedback system with direct application to UED are presented. First, such a control system will need to have an automated, model-independent feedback system robust to changes that bring the system outside of the training set. ES is shown to perform adequately for that purpose. Second, the already state-of-the-art energy stability is shown to be able to be improved using simple, linear regression-based virtual diagnostics. These virtual diagnostics will be incorporated into such a novel control system for feedback.

ISBN: 978-3-95450-227-1 ISSN: 2673-5490 doi:10.18429/JACoW-IPAC2022-TUPOST055

REFERENCES

[1] A. Edelen, N. Neveu, M. Frey, Y. Huber, C. Mayes, and A. Adelmann, "Machine learning for orders of magnitude speedup in multiobjective optimization of particle accelerator systems," *Physical Review Accelerators and Beams*, vol. 23, no. 4, p. 044 601, 2020.

[2] J. Duris *et al.*, "Bayesian optimization of a free-electron laser," *Physical Review Letters*, vol. 124, no. 12, p. 124 801, 2020.

[3] S. Li *et al.*, "Electron ghost imaging," *Physical Review Letters*, vol. 121, no. 11, p. 114 801, 2018.

[4] K. Kabra, S. Li, F. Cropp, T. J. Lane, P. Musumeci, and D. Ratner, "Mapping photocathode quantum efficiency with ghost imaging," *Physical Review Accelerators and Beams*, vol. 23, no. 2, p. 022 803, 2020.

[5] A. Scheinker, S. Gessner, C. Emma, and A. L. Edelen, "Adaptive model tuning studies for non-invasive diagnostics and feedback control of plasma wakefield acceleration at facet-ii," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 967, p. 163 902, 2020, doi:10.1016/j.nima.2020.163902

[6] C. Emma, A. Edelen, M. Hogan, B. O'Shea, G. White, and V. Yakimenko, "Machine learning-based longitudinal phase space prediction of particle accelerators," *Physical Review Accelerators and Beams*, vol. 21, no. 11, p. 112 802, 2018.

[7] Z. Zhang *et al.*, "Accurate prediction of mega-electron-volt electron beam properties from ued using machine learning," *Scientific Reports*, vol. 11, no. 1, pp. 1–13, 2021.

[8] Z. Zhang *et al.*, "Toward fully automated ued operation using two-stage machine learning model," *Scientific Reports*, vol. 12, no. 1, pp. 1–12, 2022.

[9] A. Scheinker, F. Cropp, S. Paiagua, and D. Filippetto, "An adaptive approach to machine learning for compact particle accelerators," *Scientific Reports*, vol. 11, no. 1, pp. 1–11, 2021.

[10] D. Filippetto and H. Qian, "Design of a high-flux instrument for ultrafast electron diffraction and microscopy," *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 49, no. 10, p. 104 003, 2016.

[11] A. Scheinker and M. Krstić, "Minimum-seeking for clfs: Universal semiglobally stabilizing feedback under unknown control directions," *IEEE Transactions on Automatic Control*, vol. 58, no. 5, pp. 1107–1122, 2012.

[12] A. Scheinker, "Model Independent Beam Tuning," in *Proc. IPAC'13*, Shanghai, China, May 2013, pp. 1862–1864, https://jacow.org/IPAC2013/papers/TUPWA068.pdf

[13] A. Scheinker and D. Scheinker, "Constrained extremum seeking stabilization of systems not affine in control," *International Journal of Robust and Nonlinear Control*, vol. 28, no. 2, pp. 568–581, 2018, doi:10.1002/rnc.3886

[14] A. Scheinker, A. Edelen, D. Bohler, C. Emma, and A. Lutman, "Demonstration of model-independent control of the longitudinal phase space of electron beams in the linac-coherent light source with femtosecond resolution," *Physical Review Letters*, vol. 121, no. 4, p. 044 801, 2018, doi:10.1103/PhysRevLett.121.044801

[15] A. Prokoph, G. Shields, and J. Veizer, "Compilation and time-series analysis of a marine carbonate δ18o, δ13c, 87sr/86sr and δ34s database through earth history," *Earth-Science Reviews*, vol. 87, no. 3-4, pp. 113–133, 2008.

[16] O. N. Bjørnstad and B. T. Grenfell, "Noisy clockwork: Time series analysis of population fluctuations in animals," *Science*, vol. 293, no. 5530, pp. 638–643, 2001.

[17] C. Q. Cao and R. S. Tsay, "Nonlinear time-series analysis of stock volatilities," *Journal of Applied Econometrics*, vol. 7, no. S1, S165–S185, 1992.

[18] J. Kraft and A. Kraft, "Determinants of common stock prices: A time series analysis," *The Journal of Finance*, vol. 32, no. 2, pp. 417–425, 1977.