BAYESIAN OPTIMIZATION OF LONGITUDINAL PHASE SPACE IN THE MAX IV LINAC

J. Lundquist*, J. Björklund Svensson, S. Werin and F. Curbis, Lund University, Lund, Sweden
E. Mansten, MAX IV Laboratory, Lund, Sweden

Abstract

Reaching design performance in modern particle accelerators is a challenge involving many tasks which are time-consuming and difficult to perform. It is always an advantage to be able to simplify high-level operational tasks and measurements through the assistance of optimization techniques. In this work we applied Bayesian optimization via the XOpt framework with the aim to simplify and enhance the operations in the MAX IV linac. The focus of this work has been longitudinal phase-space optimization using signals from a transverse deflector system. Further, a new approach in the optimization of longitudinal phase-space parameters with the use of virtual diagnostics has been developed and implemented.

INTRODUCTION

Within modern accelerators there is a growing need for precise control of both transverse and longitudinal beam parameters. With linear electron accelerators providing beam for light sources requiring femtosecond resolution, the longitudinal phase space (LPS) of the generated beams is of particular importance, as it directly influences the brightness and stability of the resulting photon beams [1]. Achieving and maintaining the key parameters in the LPS of the electron beam is a key issue in such instances.

At MAX IV the linac serves as a full energy injector to two storage rings, at 1.5 and 3 GeV, and also delivers beam for a short pulse facility (SPF) when not injecting to the rings [2]. The SPF houses two separate beamlines, one experimental beamline for light production, and a diagnostic beamline with a transverse deflecting structure (TDS) setup for measuring the beam LPS [3]. The complexity and practicalities of operating the TDS, along with the coupled parameter space and non-linear controls, can make manual tuning of the beam LPS both time-consuming and intrusive to user activities in the experimental beamline. Automated optimization could alleviate many of these issues and enable more efficient, simplified operation of the linac as a whole. In this work we have implemented and tested Bayesian optimization techniques for LPS control in the MAX IV linac, both with integration in the TDS diagnostic system and using a virtual diagnostic as an alternative, non-invasive evaluator of the beam LPS [4–6].

MAX IV Longitudinal Phase Space Control and Diagnostics

The longitudinal phase space of the beam is defined as its extent in the (t, δ) space, with $\delta = \frac{\Delta p}{p}$ being the relative deviation in energy. Direct measurement of the LPS at MAX IV is performed using the TDS system installed at the end of the linac [3]. The key components of the diagnostic system are two 3 m radio-frequency (RF) structures set up to deflect the beam onto a scintillating YAG screen 20 m downstream. This screen is set up in a dispersive region, thus also enabling the extraction of the beam distribution in δ. While this diagnostic system provides detailed LPS information for automatic optimization, it is also invasive and difficult to operate, providing a use case for the virtual diagnostic covered in a later section.

LPS control in the MAX IV linac is achieved with two separate stages of bunch compression, using an arc-like bunch compression scheme [7]. The principal controlling beam parameter of the compression is the energy chirp of the incoming beam, i.e. the linear correlation between the relative energy deviation and the longitudinal position within the electron bunch. The energy chirp results in longitudinal compression when moving through the bunch compressor arc as the transfer elements of R_{56} and T_{566} result in an energy dependent shift in the longitudinal positions. The energy chirp in the incoming beam is controlled using the off-crest phase of the accelerating RF fields experienced by the beam. This phase is in turn controlled through the filltime in the SLED units providing RF to the accelerating structures, as a feedback system acting on the RF phase will maintain the design energy for the bunch compressors [8]. To control beam compression, one provides a new setpoint in SLED filltime, and as the RF phase is changed by the feedback, the beam will enter the compressor with a different energy chirp.

Bayesian Optimization

Bayesian optimization (BO) is an efficient and effective way of optimizing environments which are expensive to query and have noisy or difficult objectives [9–11]. BO relies on a live probabilistic modeling of the environment, allowing it to set up an approximate model of the system's response to parameter changes. To tune the response model an acquisition function decides the selection of new evaluation points and the model is recalculated with these new points added. In order to thoroughly optimize the environment, it is key that the acquisition function achieves a balance between exploration of unexplored areas and exploitation of the model to reach the highest performance. For this project, the upper confidence bound acquisition function was used in all cases shown [12].

Optimization of the electron beam LPS is a clear use-case for BO. Querying the environment, i.e. performing a full TDS measurement, is expensive and time consuming, and

* johan.lundquist@maxiv.lu.se

the response is often non-linear and complex, depending on which aspect of the LPS one seeks to optimize. The BO integration into accelerator systems is an ongoing field of development. For the applications in this study, the *XOpt* [13] framework was used to set up the basic BO components (i.e., the surrogate model of the response and the acquisition function), while the evaluation of new TDS measurements (taking a set of TDS images and extracting FWHM in each projection) and execution of new acquisition steps (writing new SLED filltimes to the relevant accelerating sections and allowing feedbacks to adjust the machine) was integrated with the accelerator using *PyTango*. This full system was integrated and tested on the MAX IV linac with optimizations of FWHM bunch length, energy spread, and the sum of the two. A novel approach to the evaluation was also added, leveraging a virtual diagnostic to non-destructively evaluate the relevant LPS parameters.

Virtual Diagnostics

To achieve fully non-invasive, online optimization of the LPS of the beam, a virtual diagnostic (VD) system was set up for the optimizing agent to evaluate instead of the real TDS measurement. This was done by training an artificial neural network to learn the mapping between 65 different machine parameters and the resulting LPS. The machine parameters and non-destructive diagnostic readings used as input for the VD included the SLED filltimes used in the optimizer, as well as RF phases and amplitudes, BPM readings, current transformer signals and an amplitude reading from a horn antenna. Parameter scans of the optimizer parameters were done, scanning filltimes in the accelerating sections before each of the two bunch compressors. In total, a dataset of 1400 images was collected, along with the relevant input vectors to train the network on. The network was set up using *TensorFlow* and consisted of two 64-node dense layers followed by three layers of convolutional neural networks [14, 15]. Training was done using the ADAM optimizer, and leveraged a scheduled learning rate and early stopping of training to avoid overfitting [16]. As a final metric of performance, predicted and measured images were compared in terms of coefficient of determination, R^2.

With the VD trained and validated, it could be implemented in the optimization loop in place of the TDS measurements. Following the same optimization procedures as for measurements, the VD was used in optimizing FWHM bunch length and energy spread, as well as the sum of the FWHM in both projections. The evaluation procedure now consisted of collecting the relevant non-destructive diagnostics used in setting up the network, normalizing their values and feeding into to the pre-trained network. With all of this in place, live optimization on a VD proved efficient and reliable.

IMPLEMENTATION AND RESULTS

The results of this study are largely divided into three subsections: The LPS optimization directly using the TDS measurement as evaluator, the results of the VD training, and finally the results from LPS optimization using the VD as evaluator.

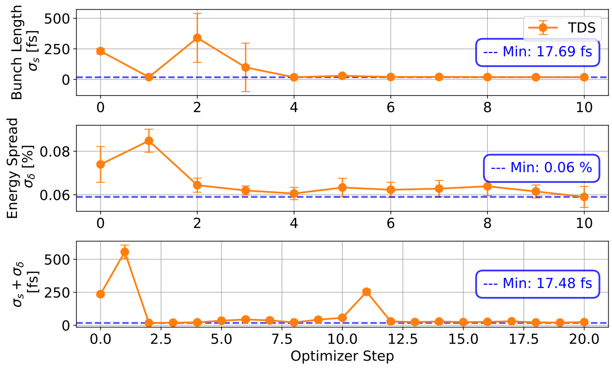

Figure 1: Optimization runs on LPS parameters using the TDS measurements, optimizing on bunch length (σ_s), energy spread (σ_δ) and the sum of the two. The dashed line indicates the minimum parameter value encountered during the optimization.

TDS Optimization

Three parameters had been selected for a test case of LPS optimization, the longitudinal FWHM of the bunch σ_s, the FWHM energy spread σ_δ and the sum of these two parameters. This was done in three consecutive optimization runs, the results of which are presented in Fig. 1, with error bars representing $\pm 1\sigma$ of 10 acquisitions performed in each optimization step. We can see how the optimizer improves each parameter in a few steps, and the minimum encountered during the optimization run is shown in each plot. It should be noted that this minimum is not necessarily the optimum according to the BO procedure, but it always found to be close to the minimum during data collection.

VD Training

With the initial success in optimizing on LPS with the TDS as evaluator, a VD was set up to predict the TDS images. The result of this procedure is summarized in Fig. 2, here showing the performance of the VD on 10% of the collected data isolated before training the model, i.e. on a so called test dataset. In Fig. 2a we can see a few examples of predictions of the LPS alongside corresponding measurements, chosen based on extreme points in FWHM bunch length. In Fig. 2b a histogram of the R^2 performance across the dataset is shown, with lines demarcating the cases shown in Fig. 2a. We can see the model achieves a median R^2 of nearly 90%. For a more interpretable numeric measure of performance, Fig. 2c shows a correlation plot of the predicted and measured FWHM bunch length σ_s, with the cases shown in Fig. 2a marked in the corresponding colors. There is a clear correlation between measured and predicted bunch lengths, and we can extract a Pearson correlation coefficient of 93.8%.

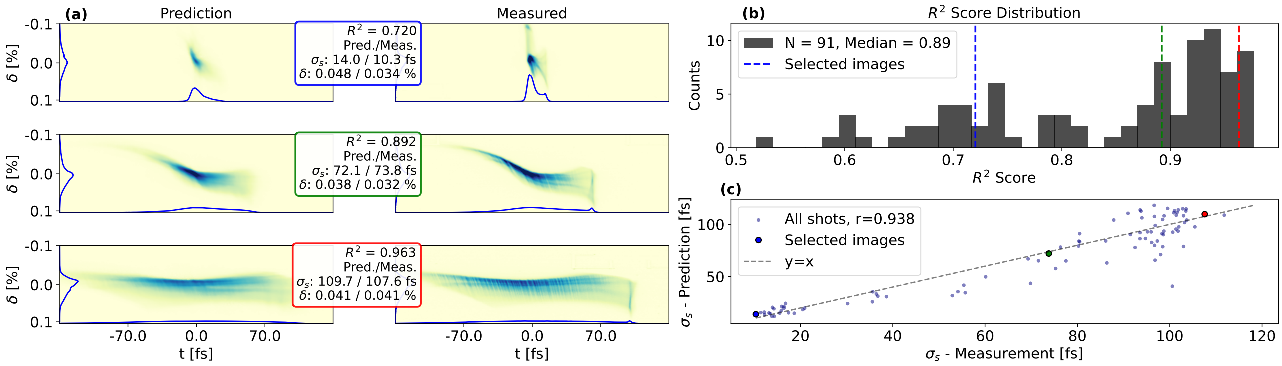

Figure 2: Results from the VD LPS predictions. Figure **(a)** shows examples of the VD predictions along with the corresponding measured images, sorted from top to bottom by FWHM bunch length. Figure **(b)** shows a histogram of the R^2 performance, with the performance in the shown cases highlighted. Figure **(c)** shows a correlation plot of FWHM bunch length from the VD predictions and the measurements, with the shown cases highlighted.

VD Optimization

With a VD set up and performing to a satisfactory degree, the final implementation into the BO loop, with the VD acting as evaluator, could be performed. The results from three such VD optimization campaigns are shown in Fig. 3, optimizing on the longitudinal FWHM of the bunch σ_s, the FWHM energy spread σ_δ and the sum of these two parameters. Error bars in the results represent $\pm 1\sigma$ of 5 acquisitions done in each step. Along with the parameters as extracted from the VD LPS prediction, in each case the parameters from the live TDS measurement are also shown. For this comparison, the parameters are also normalized to unity, as some offset was present from differences in the image pipeline. One can see how the VD and TDS cases in each parameter overlap, suggesting the VD performs well using online data for the predictions. One can also see the VD optimizations reach comparable minimum parameters as shown in the real TDS optimizations in Fig. 1.

It should be noted that the optimization campaigns were performed approximately 24 h after the VD data collection and training. Although the machine was returned to similar operating conditions, the satisfactory performance achieved using the day-old VD further demonstrates the method's robustness and resilience to drifts in the machine state.

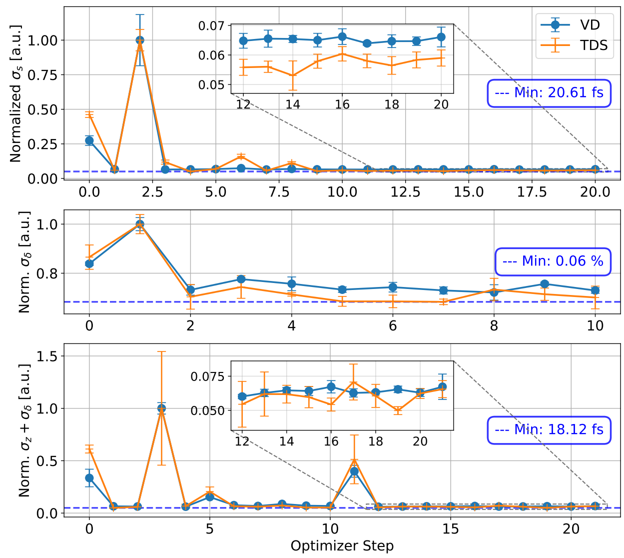

Figure 3: Optimization runs on the virtual diagnostic signal, optimizing on bunch length (σ_s), energy spread (σ_δ) and the sum of the two. Also shown is the relevant parameters as measured simultaneously using the TDS, along with a dashed line indicating the minimum value encountered during the optimization.

CONCLUSION AND OUTLOOK

A successful integration of Bayesian optimization for LPS control has been implemented in the MAX IV linac. Using the *XOpt* framework to provide the BO algorithms and *PyTango* for control system integration, an automated optimization procedure for key LPS parameters has been implemented. The optimization was tested on live TDS measurements, minimizing the parameters of FWHM bunch length and energy spread, as well as the sum of the two, showing that the BO implementation worked fully in the control room to provide new working points. Going further, to enable non-invasive operation of the optimizer, a virtual diagnostic was setup and trained to predict the TDS images based on non-destructive beam measurements and machine parameters. This VD achieved high performance in terms of R^2 and correlation with measurements in terms of FWHM bunch length. When deployed inside the optimizer, the VD reproduced the optima observed with real TDS measurements, demonstrating that reliable online optimization could be achieved without interrupting user operation. While future work should be done to implement control of higher-order parameters such as energy chirp and slice emittance, and further testing of the reliability of VDs is required, these early results indicate that BO combined with virtual diagnostics provides a novel and powerful tool for automated beam tuning and control.

REFERENCES

[1] Z. Huang and K.-J. Kim, "Review of x-ray free-electron laser theory," *Phys. Rev. ST Accel. Beams*, vol. 10, p. 034 801, 3 2007, doi:10.1103/PhysRevSTAB.10.034801

[2] MAX IV Laboratory, *Detailed design report*, https://www.maxiv.lu.se/beamlines-accelerators/accelerators/accelerator-documentation-2/, 2010.

[3] N. Blaskovic Kraljevic, L. Isaksson, E. Mansten, and S. Thorin, "Longitudinal Electron Beam Characterisation at the MAX IV Linac," *JACoW*, vol. IPAC2023, THPL136, 2023, doi:10.18429/JACoW-IPAC2023-THPL136

[4] C. Emma, A. Edelen, M. J. Hogan, B. O'Shea, G. White, and V. Yakimenko, "Machine learning-based longitudinal phase space prediction of particle accelerators," *Phys. Rev. Accel. Beams*, vol. 21, p. 112 802, 11 2018, doi:10.1103/PhysRevAccelBeams.21.112802

[5] S. Bettoni et al., "Machine learning based longitudinal virtual diagnostics at swissfel," *Review of Scientific Instruments*, vol. 95, no. 1, p. 015 110, 2024, doi:10.1063/5.0179712

[6] J. Lundquist, S. Werin, and F. Curbis, "Virtual diagnostics for longitudinal phase space imaging," English, in *Proc. IPAC'23*, Venice, Italy, 2023, pp. 4471–4474, doi:10.18429/JACoW-IPAC2023-THPL022

[7] P. H. Williams, G. Pérez-Segurana, I. R. Bailey, S. Thorin, B. Kyle, and J. B. Svensson, "Arclike variable bunch compressors," *Phys. Rev. Accel. Beams*, vol. 23, p. 100 701, 10 2020, doi:10.1103/PhysRevAccelBeams.23.100701

[8] Z. D. Farkas, H. A. Hoag, G. A. Loew, and P. B. Wilson, "SLED: A Method of Doubling SLAC's Energy.," in *9th International Conference on High-Energy Accelerators*, 1974, pp. 576–583.

[9] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas, "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

[10] M. McIntire, T. Cope, S. Ermon, and D. Ratner, "Bayesian Optimization of FEL Performance at LCLS," in *7th International Particle Accelerator Conference*, 2016, WEPOW055, doi:10.18429/JACoW-IPAC2016-WEPOW055

[11] J. Duris et al., "Bayesian optimization of a free-electron laser," *Phys. Rev. Lett.*, vol. 124, p. 124 801, 12 2020, doi:10.1103/PhysRevLett.124.124801

[12] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine Learning*, vol. 47, no. 2, pp. 235–256, 2002, doi:10.1023/A:1013689704352

[13] R. Roussel, A. Edelen, A. Bartnik, and C. Mayes, "Xopt: A simplified framework for optimization of accelerator problems using advanced algorithms," in *Proc. IPAC'23*, Venezia, 2023, pp. 4796–4799, doi:doi:10.18429/jacow-ipac2023-thpl164

[14] Martín Abadi et al., *TensorFlow: Large-scale machine learning on heterogeneous systems*, Software available from tensorflow.org, 2015, https://www.tensorflow.org/

[15] K. Fukushima, "Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position," *Biological Cybernetics*, vol. 36, no. 4, pp. 193–202, 1980, doi:10.1007/BF00344251

[16] D. P. Kingma and J. Ba, *Adam: A method for stochastic optimization*, 2017.