DETERMINISTIC METHODS AND BAYESIAN OPTIMIZATION ALGORITHMS APPLIED TO THE UH MĀNOA LINAC

N. Bidault*, S.H.P. Chan, C. Komo, S. Li, E. Valetov†, University of Hawai'i at Mānoa, HI, USA

Abstract

The University of Hawai'I at Mānoa (UHM) linac delivers up to 45 MeV electron beams to a Free-Electron Laser (FEL) oscillator. As the linac is being recommissioned for renewed FEL operation, we are developing simulation and optimization tools to recover operational settings and to explore the landscape of beam-manipulation techniques for future experimental apparatus. This paper benchmarks classical deterministic methods and Bayesian optimization (BO) algorithms on three representative beam-optics tuning scenarios using a beam dynamics simulation model developed in-house. For problems with only a few free parameters, classical methods converge reliably, while finite-difference derivative information improves the performance of constrained gradient-based solvers. For the higher-dimensional case, BO with a Gaussian Process (GP) surrogate and Sobol initialization provides a more robust path toward convergence. The emphasis is on the number of optimization iterations required for each scenario, in order to anticipate the computational cost of applying the same workflow to higher-fidelity models.

INTRODUCTION

The University of Hawai'i Free Electron Laser (UH FEL) facility, previously led by Prof. John Madey, comprises a compact S-band linear accelerator and an infrared FEL system originally developed as the Mark-III FEL undulator. The facility entered recommissioning in 2024 after a nearly decade-long standby period. A description of the recommissioning activities is reported in Ref. [1]. Additional research plans for the UHM accelerator and FEL laboratory were presented in Ref. [2].

The present work aims to benchmark optimization algorithms to determine the optimal beam transport parameters at the entrance to the FEL undulator and at an upstream Interaction Point (IP) used for the Inverse-Compton Scattering (ICS) experiment. This is a typical scenario where the Diagnostic Chicane beamline is used to shape the beam between the linac and the FEL. The Diagnostic Chicane beamline consists of a Chasman-Green Double Bend Achromat (DBA) configuration adapted to linear geometry; it offers several focusing doublet and triplet sections, allowing for the minimization of the transverse envelopes at the IP and a specific focusing of the beam at the FEL undulator entrance.

The optimization algorithms were implemented within a more general beam-dynamics simulation framework, FELsim [3]. The framework includes several widely used beam-dynamics codes (RF-Track [4], XSuite [5], COSY INFINITY [6]) and a fast linear beam-matrix-theory code developed in-house. We run the optimization algorithms on the fast linear beam-matrix-theory code as a proof of concept, before scaling the optimizers to other, more complex codes in future studies. We therefore do not aim to measure the code's ultimate execution time. Instead, we use it to estimate the number of iterations required to solve various beam-matching scenarios.

The simulation model includes quadrupoles with fringe fields, dipoles with wedge angles, and uses quadrupole currents as the optimization variables. The code returns the 6D beam parameters, and, in the context of the studies, the parameters of interest are the transverse Courant-Snyder (β_x, β_y, α_x, α_y), dispersion functions, and beam envelopes at selected locations along the transport line. The practical questions addressed here are which optimization algorithms are reliable for beam matching, how many evaluations they require, and how their performance changes with the number of free quadrupole settings. Three tuning scenarios of increasing complexity are considered. *Scenario A* uses a two-quadrupole doublet and adjusts the currents to meet the conditions $\alpha_x = 0$ and $\alpha_y = 0$ at a downstream observation point, representative of the transfer line between the linac and the first dipole of the DBA lattice. *Scenario B* extends the study to the entire Diagnostic Chicane beamline up to the undulator entrance (UE), covering the interaction point (IP). The problem is decomposed into 12 sequential stages; each stage optimizes the relevant subset of quadrupoles and passes the beam state to the following section. *Scenario C* is tested to optimize a high-dimensional parameter space by controlling 16 quadrupoles simultaneously over the IP–UE region. It combines the IP focusing constraints, an intermediate dispersion target, and the undulator Courant-Snyder matching into a single objective.

For all scenarios, the objective is the weighted mean-squared error between simulated observables and target values:

$$\mathrm{MSE} = \frac{\sum_i w_i \left(\frac{o_i - g_i}{s_i} \right)^2}{\sum_i w_i}, \tag{1}$$

where o_i is the simulated observable, g_i the target observable, s_i a normalization scale, and w_i is a weight for each observable.

* nbidault@hawaii.edu
† Department of Physics and Astronomy, Michigan State University, East Lansing, MI 48824, USA

DETERMINISTIC METHODS

Algorithms

We first compare five classical local optimizers available through `scipy.optimize` [7]: Nelder-Mead, COBYLA, L-BFGS-B, SLSQP, and trust-constr. L-BFGS-B and SLSQP use a two-point finite-difference Jacobian, while trust-constr combines the same Jacobian estimate with a BFGS Hessian approximation. Nelder-Mead and COBYLA are derivative-free, providing a useful comparison when gradient information is not available (black-box optimization).

Results

Scenario A is a two-parameter matching problem used as a controlled case for comparing convergence on a map with several local minima. The currents are varied over 0–1.5 A. For each method, 100 random restarts are performed within these bounds. The classical methods are likely to start near a local minimum (MSE = 4.2×10^{-3}) on the 2-parameter map we use in *Scenario A*. We adopt a subjective rule for analyzing the convergence of classical algorithms, discarding optimization runs that end with an MSE above 10^{-3}, while acknowledging the resulting bottleneck in these methods. Figure 1 shows the geometric mean convergence and one-standard-deviation bands in logarithmic space.

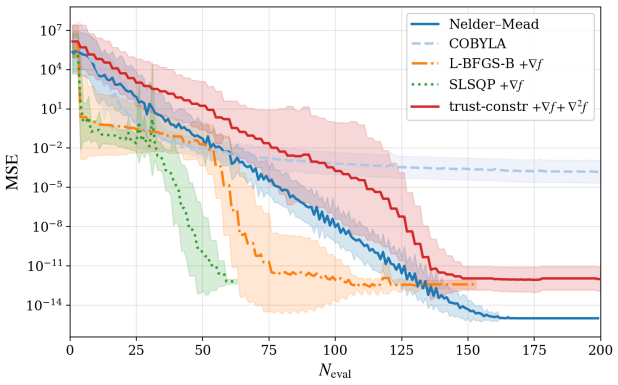

Figure 1: Convergence of the local optimization algorithms for Scenario A. The curves show the geometric mean over the random restarts, with bands corresponding to one standard deviation.

Scenario B extends the same approach to the entire Diagnostic Chicane beamline by splitting the matching problem into 12 sequential stages. Figure 2 compares Nelder-Mead, SLSQP, and trust-constr for this staged procedure. The average cumulative number of evaluations over the course of the entire Diagnostic Chicane optimization for the trust-constr method with the Jacobian estimate and a BFGS Hessian approximation is relatively high (5800) compared to the Nelder-Mead method (2600) and the SLSQP with Jacobian (1800). We also calculated the cumulative MSE for the different methods across the different stages of *Scenario B*. The SLSQP and Nelder-Mead methods yield similar values; however, the trust-constr method yields a cumulative MSE twice that of the other two methods. In general, for solving *Scenario B*, all methods shown in Fig. 2 are viable in terms of final MSE, and with faster optimum search for the SLSQP method with Jacobian.

In *Scenario C*, where the IP and undulator matching are optimized simultaneously, the local methods become less reliable. The greater number of coupled quadrupole settings increases the likelihood of finding local minima and makes the number of required iterations less predictable. After several repeated runs with 1000 evaluations for each method, the classical algorithms all yielded MSE $\geq 10^5$. This motivates the use of Bayesian optimization for the high-dimensional case.

BAYESIAN OPTIMIZATION

Framework and Algorithms

Bayesian optimization (BO) builds a probabilistic surrogate of the objective function and uses an acquisition function to select the next point to evaluate [8]. This is attractive for beamline tuning because the number of model evaluations is the relevant cost when moving from a linear-optics model to higher-fidelity tracking tools. In this work, the surrogate is a Gaussian Process implemented through Xopt [9] and BoTorch [10]. We compare Upper Confidence Bound (UCB), Expected Improvement (EI), and UCB-TurBO [11], the latter using a trust region to improve scaling in larger parameter spaces.

Initialization with Sobol Sequences

The GP is initialized with quasi-random Sobol samples [12], which cover the parameter space more uniformly than purely random points. Three history levels are compared: a cold start, a small history generated from 10^2 Sobol points, and a medium history generated from 10^4 points. For the warm-start cases, a diverse subset of the best history points is selected before the BO loop. This gives the surrogate more information at the start of the scan, at the price of additional up-front evaluations.

Scenario A: Parameter Space Map

Figure 3 shows the MSE landscape for Scenario A over the same two current variables as Fig. 1. The overlaid BO points compare UCB, EI, and UCB-TurBO. Each BO run uses 15 randomly selected initial points, followed by 200 BO steps. The BO methods converge to an MSE similar to that of the classical algorithms in this simple 2-parameter optimization case.

Scenario B: Sequential Matching

For *Scenario B*, BO was applied to the same staged matching procedure used for the local optimizers. Without history, the cumulative number of evaluations required by the BO algorithms is typically between 2700–3100, which is comparable to the Nelder-Mead result in Fig. 2. The main difference appears when prior samples are used to initialize the surrogate model. With a history generated from 10^2 Sobol points at each stage, the cumulative number of additional

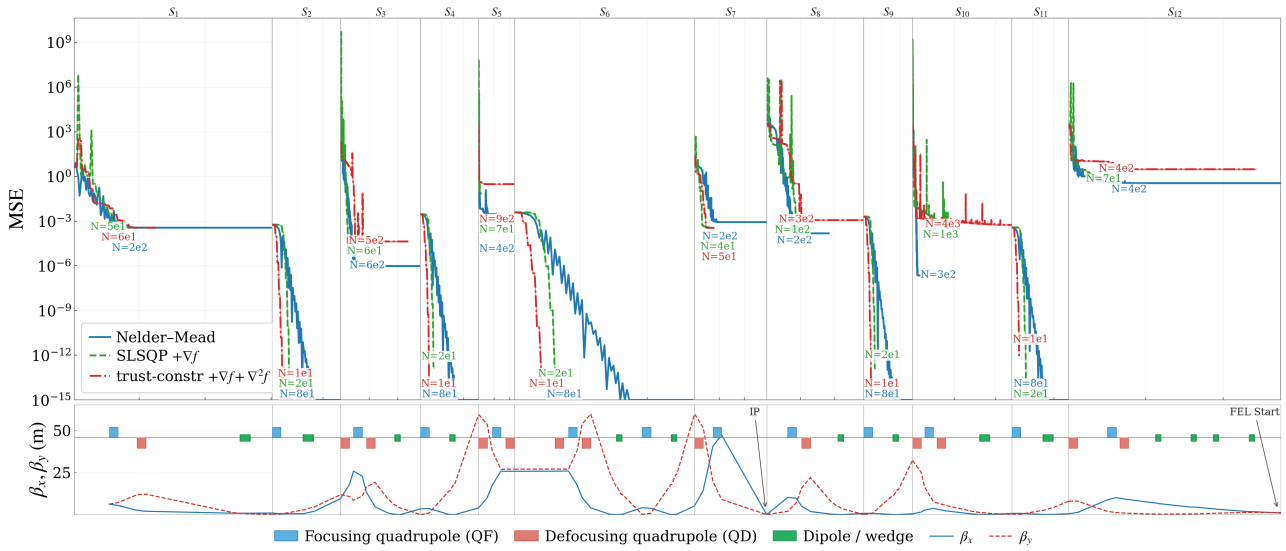

Figure 2: The upper panels show the convergence of the local optimization algorithms over the 12-stage Scenario B optimization. The lower panel shows the beamline schematic and the corresponding β-function propagation.

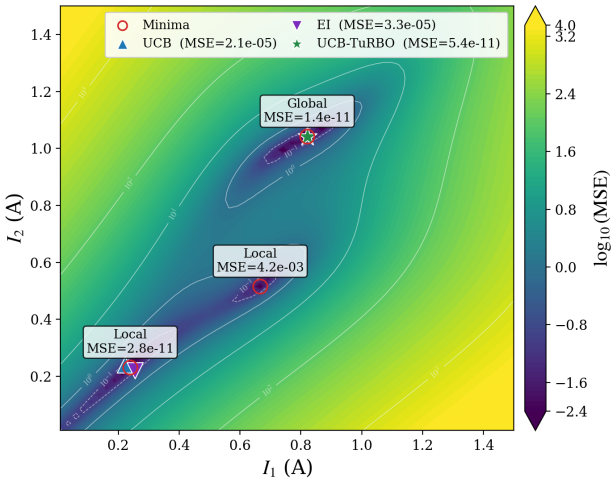

Figure 3: Scenario A MSE landscape with Grid local minima are marked together with the best BO solutions obtained with UCB, EI, and UCB-TurBO for 200 evaluations.

BO evaluations decreases to about 300–600, depending on the acquisition function. This shows that even a modest pre-sampled history can substantially accelerate the staged optimization.

The benefit of larger histories is more limited in this case, since *Scenario B* is decomposed into a sequence of low-dimensional problems, mostly involving two or three variables and only one four-variable stage. Larger histories are more relevant for the higher-dimensional case discussed below.

Scenario C: High-Dimensional Viability

For *Scenario C*, the contrast between local optimization and BO is more pronounced. Repeated local-optimizer runs limited to 1000 evaluations remained trapped at high objective values, with MSE $\geq 10^5$. Cold-start BO with UCB and EI also remained insufficient over the same evaluation budget, converging to values above MSE $= 10^3$. In contrast, cold-start UCB-TurBO already reached MSE $< 10^{-1}$ within 1000 evaluations, indicating that the trust-region strategy is better suited to the coupled high-dimensional search. The use of history is particularly beneficial in this scenario. Initializing UCB-TurBO with a history generated from 10^4 Sobol points further improved the convergence, reaching MSE $< 10^{-2}$ after about 400 additional BO evaluations.

CONCLUSION

We have presented a first benchmark of optimization algorithms applied to the UHM FEL beamlines. Local optimizers are effective for low-dimensional matching and remain useful when the beamline is decomposed into sequential sections. Bayesian optimization with a Gaussian Process surrogate and Sobol initialization provides a more robust path through the higher-dimensional parameter space.

Ongoing work includes neural-network surrogates built with PyTorch and BoTorch, the consolidation of the GP-based BO workflow with Xopt presented here, and a Gymnasium-based reinforcement-learning environment for quadrupole tuning with virtual beam monitors. Together, these developments support the recommissioning of the UH FEL beamline to target higher beam quality and lay the foundation for raising the linac's beam capabilities.

ACKNOWLEDGEMENTS

This work is supported by start-up funds provided by UH Mānoa to N. Bidault and S. Li, and by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No. DE-SC0025583.

REFERENCES

[1] N. Bidault, S. Li, H. Puwar, and A. Weinberg, "Recommissioning of the University of Hawai'i LINAC and Free Electron Laser", in *Proc. IPAC'25*, Taipei, Taiwan, Jun. 2025, pp. 1175–1178. doi:10.18429/JACoW-IPAC2025-TUPM005

[2] A. Weinberg, N. Bidault, and S. Li, "Research plans for the University of Hawai'i Accelerator and Free-Electron Laser Lab", in *Proc. IPAC'25*, Taipei, Taiwan, Jun. 2025, pp. 257–260. doi:10.18429/JACoW-IPAC2025-MOPB110

[3] C. Komo and N. Bidault, "FELsim: simulation and optimization tools for the University of Hawai'i linac and FEL", https://github.com/komochristian/FELsim

[4] A. Latina, "RF-Track Reference Manual", CERN, Geneva, Switzerland: Zenodo, Jun. 2020. doi:10.5281/zenodo.3887085

[5] G. Iadarola *et al.*, "Xsuite: An Integrated Beam Physics Simulation Framework", in *Proc. HB'23*, Geneva, Switzerland, Oct. 2023, pp. 73–80. doi:10.18429/JACoW-HB2023-TUA2I1

[6] M. Berz and K. Makino, "COSY INFINITY and Its Use for Singlepass and Multipass Systems", in *J. Microsc.*, Mar. 2026. doi:10.1093/jmicro/dfag016

[7] P. Virtanen *et al.*, "ciPy 1.0: fundamental algorithms for scientific computing in Python", *Nat. Methods*, vol. 17, pp. 261–272, 2020. doi:10.1038/s41592-019-0686-2

[8] E. Brochu, V. M. Cora, and N. de Freitas, "A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning", Dec. 2010, arXiv:1012.2599 [cs.LG]. doi:10.48550/arXiv.1012.2599

[9] R. Roussel, C. Mayes, A. Edelen, and A. Bartnik, "Xopt: A simplified framework for optimization of accelerator problems using advanced algorithms", in *Proc. IPAC'23*, Venice, Italy, May 2023, pp. 4847–4850. doi:10.18429/JACoW-IPAC2023-THPL164

[10] M. Balandat *et al.*, "BoTorch: A framework for efficient Monte-Carlo Bayesian optimization", in *Adv. Neural Inf. Process. Syst.*, vol. 33, Dec. 2020. doi:10.5555/3495724.3497531

[11] D. Eriksson *et al.*, "Scalable global optimization via local Bayesian optimization", in *Adv. Neural Inf. Process. Syst.*, vol. 32, 2019.

[12] I. M. Sobol, "On the distribution of points in a cube and the approximate evaluation of integrals," *USSR Comput. Math. Math. Phys.*, vol. 7, no. 4, pp. 86–112, 1967. doi:10.1016/0041-5553(67)90144-9