MULTI-OBJECTIVE BAYESIAN OPTIMISATION (MOBO) FOR HIGH-QUALITY PHOTOINJECTOR OPTIMISATION

E. J. Brookes*, T. Kamps, Helmholtz-Zentrum Berlin für Materialien und Energie, Berlin, Germany
S. Zeeshan, Deutsches Elektronen-Synchrotron DESY, Zeuthen, Germany

Abstract

Optimising SRF photoinjector parameter space is a challenging task due to the high-dimensional, non-linearly coupled parameters and competing objectives such as transverse emittance and bunch length. Conventional methods such as manual tuning or MOGA require thousands of evaluations and are impractical for routine operation or computationally expensive simulations. This work presents a multi-objective Bayesian optimisation approach that uses Gaussian-process surrogate models and tunable, uncertainty-aware acquisition functions to identify Pareto-optimal solutions in an order of magnitude fewer evaluations. When applied to the 1.4-cell SRF photoinjector at SEALab, and the 1.6-cell SRF gun and 20 m injector beamline for EuXFEL, this optimisation outperforms MOGA in solution-efficiency and provides interpretable sensitivity information for injector tuning. These results demonstrate the potential of MOBO as an efficient, machine-ready strategy for SRF photoinjector optimisation.

INTRODUCTION

High-brightness SRF photoinjectors are required to provide electron beams with stringent longitudinal and transverse properties. Since the quality of the beam delivered to downstream sections is strongly influenced by the beam generated at the source, efficient injector tuning is essential for reliable accelerator operation. The optimisation of photoinjector settings is a coupled and multi-objective problem. High-brightness photoinjector performance is critically reliant on the simultaneous minimisation of bunch length and transverse emittance, two competing objectives. The machine controls for these problems form a multi-dimensional parameter space in which manual tuning can be slow and incomplete. Population-based methods such as Multi-Objective Genetic Algorithms (MOGA) have been successfully used to identify high-quality Pareto fronts in accelerator physics [1], but they often require many thousands of evaluations, and can therefore be inefficient for expensive evaluations. Rather than densely sampling the full parameter space, Multi-Objective Bayesian Optimisation (MOBO) builds probabilistic surrogate models of the accelerator response and uses the uncertainty of these models to select informative new samples.

This paper applies the same MOBO framework to the simultaneous minimisation of bunch length and transverse emittance on two SRF photoinjector optimisation problems: a 1.74 m SEALab beamline [2] and a 20 m EuXFEL injector beamline [3].

* brookesemily17@yahoo.com

MULTI-OBJECTIVE BAYESIAN OPTIMISATION

MOBO aims to solve for the vector, x, which optimises the objectives, $[\sigma_z(\mathbf{x}), \epsilon_n(\mathbf{x})]$, subject to $Q_{\text{transported}}/Q_0 > Q_{\text{min}}$ [4]. Here, σ_z and ϵ_n are the bunch length and normalised transverse emittance measured at the observation point, and $Q_{\text{min}} = 0.95$ characterises the minimum acceptable transport efficiency.

The MOBO workflow is represented by Fig 1. MOBO is an iterative procedure, in which the parameter space is sampled based on the previous iteration results. In this work, for each iteration, five new samples are identified and evaluated in parallel. These measured points are used to train the Gaussian process surrogate models for both objectives, and a third model for the constraint parameter (here, Q_{min}). The models provide both a prediction and uncertainly for each point in the input space which are then used by the acquisition function to determine the next most useful location to sample in by balancing the improvement of the identified optima or by exploring in regions of high uncertainty [4]. This intelligent approach to sampling allows for efficient navigation in high-dimensional spaces and thus reduces the total number of evaluations necessary. Throughout optimisation, the algorithm attempts to increase the hypervolume score (HV). This determines how much of the output space is being dominated by the Pareto front, relative to a user-defined reference point.

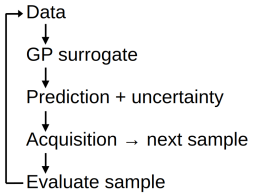

Figure 1: MOBO workflow used for photoinjector tuning.

BEAMLINE SETUPS

First, the SEALab case is discussed which represents a transverse Gaussian electron bunch through 1.4-cell SRF gun cavity followed by a 1.74 m beamline including an emittance compensation solenoid, two steerer magnets and a compact set of diagnostics. More details on this setup can be found in [2]. This study is designed to simulate the commissioning stage of operation, with bunch charges from 5-9 pC, and beam energies limited to ~2 MeV from the gun. The low charges limit the effect of space charge, but low energies decrease the adiabatic damping. The SEALab gun aims to produce beams with transverse emittance <1mm mrad, with

flexible bunch lengths for a variety of experimental setups whilst maintaining high beam transport.

The EuXFEL case consists of a ellipsoidal bunch through a 1.6-cell SRF L-band gun, and a 20 m beamline which includes an emittance compensation solenoid followed by an accelerating linac formed of eight 9-cell 1.3GHz TESLA cavities. Further details on this beamline can be found in [3]. This beamline is designed to feed a free electron laser, requiring higher bunch charges and energies than the SEALab commissioning setup, at 1 nC and an energy of 145 MeV, thus balancing space-charge effects with stronger adiabatic damping and compression at the linac. The aim of this injector is to maximise brightness for further use downstream.

The optimisation parameters for both setups are given in Tables 1 and 2.

Table 1: Optimisation Parameters for SEALab

Parameter	Value
RMS pulse duration [ns]	[1e-3,10e-3]
RMS transverse beam size [mm]	[0.5,10]
Solenoid strength [T]	[0,0.2]
Gun phase [deg]	[-10,60]
Gun amplitude [MV/m]	[7,20]
Bunch charge [pC]	[5,9]

Table 2: Optimisation Parameters for EuXFEL

Parameter	Value
RMS pulse duration [ns]	[3.1e-3,9.48e-3]
RMS transverse beam size [mm]	[0.4,1.4]
Solenoid strength [T]	[0.184,0.2]
Gun phase [deg]	[-6,1]
Gun field amplitude [MV/m]	55
A1 phase (1st half) [deg]	[-9,1]
A1 amplitude (1st half) [MV/m]	32
A1 phase (2nd half) [deg]	[-5,3]
A1 amplitude (2nd half) [MV/m]	32

SEALab RESULTS

Figure 2 demonstrates that MOBO can rapidly locate the bunch length/emittance trade-off for a compact SRF photoinjector beamline. The hypervolume increases rapidly during the early optimisation iterations, indicating that the algorithm quickly discovers feasible operating regions. After this initial rise, the hypervolume increases step-wise as points within these feasible operating regions are improved. A plateau begins to form after 42 iterations, suggesting that additional samples produce only marginal improvement in the Pareto set. Since 5 samples are evaluated in each iteration, this amounts to 210 samples at the time of plateau.

The final Pareto set defines the accessible trade-off between short bunch length and low transverse emittance. The smallest bunch length found in the SEALab case is 0.058 mm, while the smallest emittance is 0.35 mm mrad. A representative balanced Pareto point gives a bunch length of

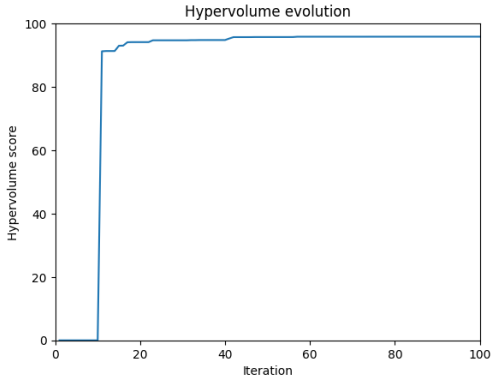

Figure 2: The hypervoume score for the SEALab optimisation, reaching a plateau after approximately 210 samples.

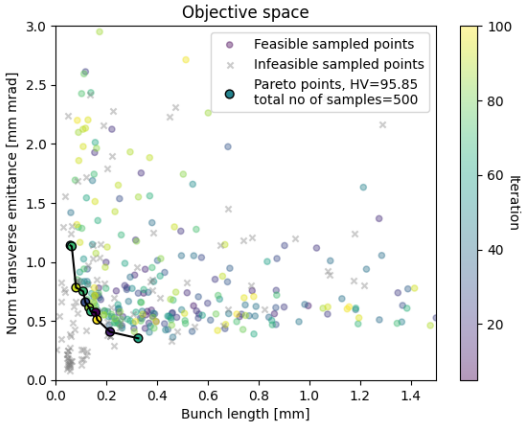

Figure 3: Representation of all samples evaluated as feasible or infeasible samples, and the final located Pareto front.

0.158 mm and a transverse emittance of 0.57 mm mrad. Several points were evaluated which did not meet the constraint threshold (95% charge transported), as shown by the grey crosses. These number less than a third of the total samples evaluated, indicating that whilst exploring the input space, infeasible operating regions were located and the algorithm penalised sampling in these areas. These infeasible regions are concentrated in setups in which the gun RF amplitude is low and RF phase is non-optimal, and thus particles are not accelerated sufficiently to be focused and transported.

EuXFEL RESULTS

The EuXFEL optimisation shows the same MOBO optimisation on an extended beamline, allowing the introduction of booster linacs further downstream, which are able to act as bunch compressors and adiabatic dampers for the bunch length and emittance objectives respectively. In addition, this study provides a direct comparison between MOBO and a MOGA reference front. The MOGA data represent a densely sampled reference optimisation, with approximately 40,000 samples used to identify the high-quality bunch-length–emittance region, as generated in [1]. The output space from this study is shown in Fig. 4.

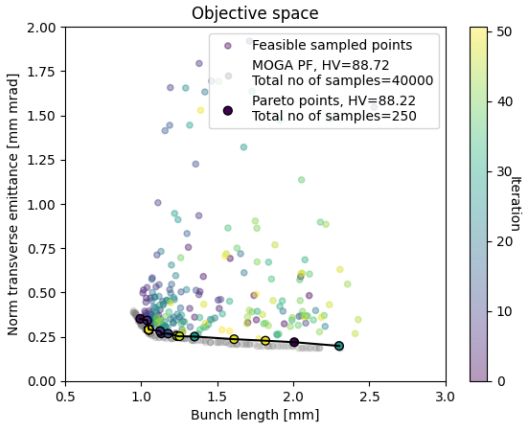

Figure 4: Comparison of MOBO and MOGA optimisation for the EuXFEL injector case, showing the MOGA- and MOBO-identified Pareto fronts and all MOBO samples.

Following the results from the SEALab study, the results for the optimisation for EuXFEL were limited to samples evaluated in fewer than 50 iterations, but the optimisation was allowed to continue to 100 iterations to verify the identification of the plateau. The plateau was found by iteration 47 (235 samples). The located MOBO Pareto front overlaps the MOGA-identified Pareto front with only a few hundred samples, corresponding to a reduction in the number of required evaluations by more than two orders of magnitude, showing that the Bayesian method is able to recover the relevant trade-off without dense sampling of the full parameter space. This is a central result of the EuXFEL study: MOBO does not merely find isolated good points, but reconstructs the physically useful Pareto region with far fewer evaluations.

The smallest bunch length found is 0.99 mm, and the smallest emittance is 0.198 mm mrad. A representative balanced Pareto point has a bunch length of 1.24 mm and a transverse emittance of 0.25 mm mrad. These values differ from those found for the SEALab case, owing mostly to the difference in bunch charge, and the introduction of the booster linacs which allow for increased energies and emittance compensation. The Pareto front points found by MOBO span the range of the MOGA Pareto front, at spaced intervals, indicating that MOBO provides a sufficiently populated Pareto set for practical operating-point selection while maintaining a low evaluation cost.

These results show that the algorithm does not need to densely sample the full six-dimensional parameter space in order to identify useful operating regions. Instead, as the algorithm progresses, the MOBO samples concentrate increasingly near the high-quality region of objective space.

GP SENSITIVITY ANALYSIS

The GP surrogate models were used to investigate the sensitivity of the objectives to the input variables via Shapley conditional sensitivities [5]. This analysis provides an interpretable diagnostic of the optimisation problem by identifying which variables most strongly influence the predicted objectives. The three highest sensitivities for both objectives on each beamline are given in Table 3.

Table 3: Sensitivity Analysis Results

Objective	Input	Sensitivity
SEALab σ_z	**RMS pulse duration**	**0.718**
SEALab σ_z	Gun phase	0.146
SEALab σ_z	RMS trans. beam size	0.138
SEALab ϵ_n	**Solenoid strength**	**0.317**
SEALab ϵ_n	RMS trans. beam size	0.252
SEALab ϵ_n	Gun phase	0.178
EuXFEL σ_z	**RMS pulse duration**	**0.726**
EuXFEL σ_z	Gun phase	0.135
EuXFEL σ_z	RMS trans. beam size	0.123
EuXFEL ϵ_n	**RMS trans. beam size**	**0.386**
EuXFEL ϵ_n	Solenoid strength	0.363
EuXFEL ϵ_n	Gun phase	0.144

In both SEALab and EuXFEL, the bunch length is critically sensitive to the bunch duration which sets the initial bunch length, and secondarily to the gun emission settings which control the extraction and initial bunch compression in the beamline. Conversely, the transverse emittance in both cases are sensitive to the same three variables (solenoid strength, transverse beam size and gun phase) with the solenoid strength and transverse beam size more equally sensitive. This demonstrates the emittance compensation setup of the solenoid, and the dependence on the initial beam size sampling the RF field as a function of the radial distance. These sensitivities help explain why the Pareto fronts occupy particular regions of objective space and provide guidance for future reduced-dimensional or operator-assisted tuning.

CONCLUSIONS

MOBO efficiently identified Pareto-optimal operating regions using 210 or 235 evaluations, a reduction of around two orders of magnitude to the tens of thousands of evaluations used during a MOGA optimisation, while reaching the same high-quality region as MOGA, with interpretable sensitivity and prediciton information from the trained GP models. This combination of efficient optimisation and diagnostic insight makes MOBO a promising approach for simulation-based studies and online accelerator operation.

ACKNOWLEDGEMENTS

The authors thank the EuXFEL injector upgrade studies team at DESY for providing the MOGA optimisation results and the simulation data used in the MOBO optimisations for the EuXFEL study presented in this work.

REFERENCES

[1] S. Zeeshan, D. Bazyl, M. Krasilnikov, X. Li, and I. Zagorodnov, "Beam dynamics optimization in high-brightness Photo Injector with various photocathode laser pulse shapes", in *Proc.*

IPAC'25, Taipei, Taiwan, Jun. 2025, pp. 1454–1457. doi:10.18429/JACoW-IPAC2025-TUPS019

[2] T. Kamps *et al.*, "First beam commissioning of the HZB SRF photoelectron gun", in *Proc. IPAC'25*, Taipei, Taiwan, Jun. 2025, pp. 1718–1721. doi:10.18429/JACoW-IPAC2025-WECN2

[3] M. Altarelli *et al.*, "XFEL The European X-Ray FreeElectron Laser Technical Design Report", DESY, Hamburg, Germany, Jul. 2006. https://bib-pubdb1.desy.de/record/77248/files/european-xfel-tdr.pdf

[4] P. Frazier, "A Tutorial on Bayesian Optimization", *arXiv*. doi:10.48550/arXiv.1807.02811

[5] E. Song, *et al.*, "Shapley Effects for Global Sensitivity Analysis: Theory and Computation", *SIAM/ASA J. Uncertainty Quantification*, vol. 4, num. 1, pp. 106-1083, Jan. 2016, doi:10.1137/15M1048070