

Multi-Objective Bayesian Active Learning for MeV-ultrafast electron diffraction

Fuhao Ji^{1*}, Auralee Edelen¹, Ryan Roussel¹, Xiaozhe Shen¹,
Sara Miskovich¹, Stephen Weathersby¹, Duan Luo¹,
Mianzhen Mo¹, Patrick Kramer¹, Christopher Mayes¹,
Mohamed A. K. Othman¹, Emilio Nanni¹, Xijie Wang¹,
Alexander Reid¹, Michael Minitti¹, Robert Joel England^{1*}

¹SLAC National Accelerator Laboratory, Menlo Park, 94025, California,
USA.

*Corresponding author(s). E-mail(s): fuhaoji@slac.stanford.edu;
england@slac.stanford.edu;

Abstract

Ultrafast electron diffraction using MeV energy beams (MeV-UED) has enabled unprecedented scientific opportunities in the study of ultrafast structural dynamics in a variety of gas, liquid and solid state systems. Broad scientific applications usually pose different requirements for electron probe properties. Due to the complex, nonlinear and correlated nature of accelerator systems, electron beam property optimization is a time-taking process and often relies on extensive hand-tuning by experienced human operators. Algorithm based efficient online tuning strategies are highly desired. Here, we demonstrate multi-objective Bayesian active learning for speeding up online beam tuning at the SLAC MeV-UED facility. The multi-objective Bayesian optimization algorithm was used for efficiently searching the parameter space and mapping out the Pareto Fronts which give the trade-offs between key beam properties. Such scheme enables an unprecedented overview of the global behavior of the experimental system and takes a significantly smaller number of measurements compared with traditional methods such as a grid scan. This methodology can be applied in other experimental scenarios that require simultaneously optimizing multiple objectives by explorations in high dimensional, nonlinear and correlated systems.

Introduction

Ultrafast electron diffraction using MeV energy beams (MeV-UED) has led to a new paradigm in ultrafast electron scattering [1, 2]. MeV electron probes lead to atomic spatial resolution together with sub-picosecond temporal resolution[6]. Furthermore, MeV electrons can provide pure kinetic scattering signals for samples with a given thickness, which makes it possible to achieve quantitatively understood observables to validate physical modeling. Recently, MeV-UED techniques have reached a high level of maturity with several beamlines operating as user facilities. Rapid advancements in MeV fs electron probes, sample delivery and manipulation techniques, high efficiency electron detection and novel data analysis techniques [3–7] have enabled unprecedented scientific opportunities in ultrafast structural dynamics, including observations of bond dissociation, coherent ground state wavepacket motion after electrocyclic ring-opening [8], simultaneous observation of electronic and nuclear structure changes in a molecule undergoing non-adiabatic dynamics [9], hydrogen bond strengthening in liquid water [10], and phase switches of materials with exotic functional properties [11, 12].

Broad scientific applications usually pose different requirements for electron beam properties, such as temporal length, spot size and transverse momentum resolution in reciprocal space (q-resolution). Measuring these beam properties as a function of the input parameters using limited and time-consuming diagnostics becomes a necessary part of planning and optimization of experiments. Due to the complex, nonlinear and correlated nature of accelerator systems, electron beam property optimization often relies on extensive and time-taking hand-tuning by experienced human operators. This comes at the cost of reducing the beam time available for scientific experiments. Uniformly spaced grid-like parameter scans are routinely conducted in one or two dimensions for sampling the parameter space and obtaining an overview of the beam property behaviors. However, grid scans suffer from poor scaling as the number of sampling points increases exponentially with input parameter dimensions which maybe impractical during a limited beam time. Algorithm based efficient online tuning strategies are highly desired as they hold the potential to reducing the setup time, allowing human experts to focus on more challenging high-level scientific problems.

Recently, Artificial intelligence and Machine learning (AI/ML) have reached sufficient maturity for applications in real-world tasks such as autonomous driving, natural language processing [13], protein folding prediction [14] and fusion reactor control [15]. Speeding up and aiding online optimizations of complex particle accelerators is one of the key areas where AI/ML can make substantial contributions [16–23]. Bayesian Optimization (BO) with Gaussian processes(GPs) [27], often referred to as Bayesian active learning[22, 24, 25], represents a class of ML algorithms in which a probabilistic model and its uncertainty at each measurement step are used to select the optimal next observation point. Recently, BO based optimization methods have been successfully applied for online tuning at large-scale accelerator systems such as LCLS [19], with a lower number of observations needed than Gradient-based or Nelder-Mead simplex algorithms. These algorithms are designed to optimize a single beam property (an “Objective”). In practice, beam tuning often seeks to simultaneously optimize multiple Objectives. This pose a significant challenge as individual beam properties may not be independent optimizables due to complex inter-correlations between them. For

example, in UED it is difficult to simultaneously minimize the electron pulse length and spot size/q-resolution due to Space Charge(SC) forces. The goal of multi-objective optimization is to determine the Pareto Front (PF) which represents the system performance limit and gives the trade-offs between key objectives. In this domain, one of the most popular methods are the evolutionary algorithms for their simplicity in implementation and broad applicability [26]. However, the evolutionary algorithms are extremely data inefficient as they require evaluation of large number of candidate solutions to converge to PF. As a result, the evolutionary algorithms are nearly impossible to use for online accelerator tunings [23]. To this end, the Multi-Objective Bayesian Optimization (MOBO) scheme [20, 40] was recently introduced. This method uses a set of GP surrogate models, along with a multi-objective HyperVolume (HV) improvement acquisition function, to reduce the number of observations needed. This method can converge by at least an order of magnitude more quickly than evolutionary algorithms, such as Multi-Object Genetic Optimization (MOGA), and is thus a critical step toward online multi-objective optimization on real accelerator systems.

In this work, we demonstrate Multi-Objective Bayesian Active Learning for speeding up online beam tuning on an MeV-UED facility at SLAC National Accelerator Laboratory. The MOBO algorithm was used for searching the parameter space efficiently in a serialized manner, and actively learn the PF giving trade-offs of key beam properties (electron pulse length, q-resolution and spot size at sample) that have a direct impact on the outcome and quality of scientific user experiments. The achieved performance was comparable with that obtained by experienced human operators. The MOBO scheme enables an unprecedented overview in the multi-dimensional parameter space with a significantly smaller number of measurements required compared with traditional exploration methods such as a Grid Search(GS) scan. While our focus here is the application of MOBO in beam tuning for MeV-UED, the explained methodology can be applied in other experimental scenarios involving high dimensional, nonlinear and correlated systems.

Results

We conducted experiments at the SLAC MeV-UED facility to demonstrate the feasibility of applying MOBO for online exploration and optimizations of electron beam properties in a multi-dimensional parameter space. The experimental setup and MOBO workflow are illustrated in Fig. 1. The MeV-UED beamline generates 4.2 MeV energy electron pulses using an S-band photoinjector triggered by 60 fs FWHM laser pulses. The electron pulse length downstream is sensitive to the launch phase of the electrons relative to the photoinjector RF field(referred to as the "gun phase" and characterized by the parameter ϕ). The electron pulses are then focused using two magnetic solenoids. One is referred to as the "gun solenoid" and characterized by the parameter B_1 , the other is referred to as the "micro-focus solenoid" and characterized by the parameter B_2 . The electrons are delivered to the Interaction Point(IP) for performing pump-probe electron diffraction experiments. The electron diffraction patterns are then recorded and used to extract fs time scale structural dynamics signals. Here, three electron beam properties are optimized. The electron pulse length(σ_t) at IP sets

the smallest temporal resolution that can be achieved for resolving fast dynamical processes. The spot size(σ_x) at IP sets the transverse probing area and the q-resolution (σ_q) on the diffraction detector represents the transverse momentum resolution in the reciprocal space. Due to space charge forces, σ_t , σ_x and σ_q are competing properties that can not be simultaneously minimized. We used MOBO to explore the response of σ_t , σ_x and σ_q to the input beamline parameters $[\phi, B_1, B_2]$ and obtained the PFs giving trade-offs between them. In each iteration, MOBO evaluates a candidate point in the parameter space and updates the GP surrogate model with the newly acquired data. The Expected Hypervolume Improvement (EHVI) [29] acquisition function is then used to calculate the optimal next observation point (see methods). The process is performed iteratively and ultimately providing an approximation of the best achievable PF.

We first conducted σ_t vs σ_q optimizations by varying gun phase (ϕ) and solenoid strengths (B_1 and B_2). The initial pulse charge was set to be 10 fC. The gun amplitude was 90 MV/m. A 200 μ m collimator was used as an emittance filter located after the gun solenoid. We initialized the exploration with 10 random sampling points in the chosen subdomain of input parameter space (Table 1). Fig. 2 illustrates the exploration result after 121 measurements. For each measurement, three observations of σ_t and σ_q were recorded for averaging the fluctuations due to system jitters. These measurements are then used to train GP models of σ_t and σ_q responses to the input parameter ϕ , B_1 and B_2 . A Matérn kernel with $\nu = 5/2$ was used and hyperparameter training was done by maximizing the log marginal likelihood of the GP model with the gradient based Adam optimization algorithm. Measurements of σ_t projected to the $\phi - B_2$ subspace, along with the posterior mean predicted by the GP model, are plotted in Fig. 2(a). Estimated length scales from the GP models are shown in Table 1. The $\mathcal{L}(\sigma_t)$ for ϕ , B_1 and B_2 are at the same scale, implying that σ_t depends on all three input parameters. This is consistent with previous observations as σ_t depends on both the temporal-energy chirp at gun exit(changed by ϕ) and SC induced temporal broadening during beam transport(changed by B_1 and B_2). Fig. 2(b) and (c) illustrates the σ_t and σ_q responses projected to the $B_1 - B_2$ subspace. The $\mathcal{L}(\sigma_q)$ for B_1 is significantly shorter than that for B_2 , indicating that σ_q varies faster with B_1 . Again, this is consistent with previous observations that q-resolution mainly depends on the transmitted beam emittance and gun solenoid strength. On the other hand, The $\mathcal{L}(\sigma_q)$ for ϕ is significantly larger than the $\mathcal{L}(\sigma_q)$ for B_1 and B_2 , indicating that the q-resolution is weakly correlated to gun phase. The GP predicted PF curve(denoted with green squares) is plotted in Fig. 2(b) and (c). The GP predicted PF curve connects the GP predicted optimal set points for lowest σ_t at $[B_1 = 0.131, B_2 = 0.240]$ and lowest σ_q at $[B_1 = 0.123, B_2 = 0.217]$. Here, we define the empirical non-dominated data points as the approximate PF. The approximate PF points are plotted with green stars, the difference between approximate PF and GP predicted PF indicates that further exploration is needed to obtain the ultimate PF, and the lighter colored regions (both in σ_t and σ_q plots) are particularly valuable and remains to be explored. Fig. 2(d) and (e) show the calculated EHVI acquisition function projected onto the $\phi - B_2$ and $B_1 - B_2$ space, the most valuable next observation point was found to be at $[\phi = 46.0, B_1 = 0.125, B_2 = 0.237]$. This is consistent with the expectation that the

Table 1 Beamline input parameter ranges and kernel length scales obtained using GP regression, length scales are normalized to unit cube input space

Parameter	Range /Unit	$\mathcal{L}(\sigma_t)$	$\mathcal{L}(\sigma_q)$
Gun phase (ϕ)	[35, 60] degree	1.095	0.822
Gun sol. strength(B_1)	[0.11, 0.17] T	0.765	0.077
Micro-focus sol. strength(B_2)	[0.15, 0.30] T	0.852	0.215

EHVI should suggest exploring points on the predicted PF curve to efficiently improve the hypervolume. Here, the advantage of the MOBO scheme can be clearly seen. It strategically proposes the next observation point which optimally increase the Pareto Frontier HV, and is thus more data efficient than a broad, undirected search for the PF. Fig. 2(f) shows the measured data points in the objective space. The approximate PF shows clear trade-off between σ_t and σ_q . The smallest σ_q was found to be $0.162 \text{ \AA}^{-1}$ at $\sigma_t = 108 \text{ fs}$, while smallest $\sigma_t = 67 \text{ fs}$ can be achieved with $\sigma_q = 0.217 \text{ \AA}^{-1}$.

We then used MOBO to conduct spot size(σ_x) vs q-resolution(σ_q) explorations by varying B_1 and B_2 at 10 fC, 50 fC and 100 fC initial pulse charges, the input parameter ranges used are shown in Supplementary Table 1. For each run, 10 random sampling points were used to initialize the exploration and in total 150 measurements were taken. Fig. 3(a) illustrates the exploration results, all measured data points are plotted in the $\sigma_x - \sigma_q$ space. The approximate PFs clearly shows trade-offs between σ_x and σ_q at different initial pulse charges. We observe that the approximate PF obtained at lower pulse charge (10 fC) shows significant improvement compared with higher pulse charge cases (50 fC and 100 fC). This observation is consistent with previous observations that smaller σ_x or σ_q can be achieved by lowering initial pulse charge, as spot size and q-resolution are limited by SC induced emittance growth at 4.2 MeV beam energy. At 10 fC initial pulse charge, the smallest σ_x was found to be $45 \text{ \mu m}$ at $\sigma_q = 0.351 \text{ \AA}^{-1}$, while smallest $\sigma_q = 0.162 \text{ \AA}^{-1}$ can be achieved with $\sigma_x = 100 \text{ \mu m}$. We also observe that σ_q flattens with σ_x higher than $100 \text{ \mu m}$, implying that other source of broadening on the diffraction detector, such as the detector Point Spread Function(PSF) are the primary limitation for further improvement in the q-resolution. We ran 9 independent optimizations for the 10 fC initial pulse charge case, each with a different set of 10 randomly sampled initial points. After each measurement, the hypervolume of the approximate PF point set was calculated in normalized objective space. Fig. 3(b) shows the convergence plot of approximate PF hypervolume versus number of iterations. One can clearly see improvement of hypervolume across the optimization process. The averaged hypervolume (blue) reaches 95% of its maximum within 30 measurements. We conducted a simulated Grid Search (GS) with step size equal to 10% of the parameter space subdomain using the interpolated dataset. The hypervolume obtained using GS was 62% of that obtained using MOBO after 30 measurements and 91% after 100 measurements. A direct comparison between MOBO and GS scheme is plotted in Fig. 3(b), the advantage of MOBO to improve both search efficiency and maximum achievable hypervolume can be clearly seen.

Discussion

Here we have experimentally demonstrated MOBO scheme to autonomously and efficiently optimize a subset of critical electron beam properties for an MeV-UED facility. The MOBO algorithm was used for sampling the multi-dimensional parameter space efficiently with little prior knowledge of the experimental system, which makes it practical for use in online beam tuning. By proposing solutions which optimally increase the hypervolume at each measurement step, the MOBO algorithm was capable to active-learn the PFs which give the trade-offs between key beam properties of interest for scientific experiments. The PF offers an unprecedented overview of the machine's performance limitations and can greatly assist human scientists in rapid decision-making, thereby maximizing the scientific output within the very limited beam time. The MOBO scheme was capable to achieve system performance comparable with that obtained by experienced operators and requires a significantly fewer measurements compared with traditional exploration methods such as a GS scan. It is anticipated that this advancement will have a direct and substantial impact on the scientific experiments conducted at the SLAC MeV-UED facility, particularly in the studies which require simultaneously optimizing electron pulse length and q-resolution/spot size, such as the studies of charge density wave orders [33], ultrafast energy transfer in small-scale, low-dimensional systems [34], and other complex phenomena such as unconventional superconductivity in twisted-angle graphene super lattices [35]. While this work focuses on the application of the MOBO scheme in electron beam tuning for MeV-UED, this algorithm is flexible, efficient and can be used to replace GS in any experimental scenarios that require online explorations in a multi-dimensional, nonlinear parameter space while simultaneously optimizing multiple objectives, and we expect it to have broad impact across researches in different disciplines (physics, chemistry and biology) at large scale, complex Scientific User Facilities, such as Free Electron Lasers, Synchrotron Radiation Light Sources and Neutron Sources.

While our demonstration of MOBO for online tuning was conducted in a relatively low-dimensional input parameter space, it is expected that MOBO can be scaled up to higher dimensional parameter space [30]. During our experiment, the extra computation time associated with GP fitting and EHVI acquisition function optimization is small relative to the reduction in electron beam property evaluation time associated with faster convergence of HV (as shown in Supplementary Fig. 2). One major challenge is the rapidly increasing computational cost when the number of data points becomes large, as is often the case when conducting explorations in high dimensional parameter spaces. One solution is to use Neural Network (NN) surrogate models as the prior mean to improve the BO efficiency, even with erroneous predictions [36]. Work in the multi-objective Bayesian optimization field has produced efficient methods for HV and acquisition function calculations [31]. Finally, if the evaluations can be conducted in parallel, as in simulations or under the multi-beamline scenario, recently developed parallelized acquisition functions, such as q-Expected Hypervolume Improvement (q-EHVI) [32] can be used to significantly reduce overall optimization time while maintaining the sampling efficiency advantages of EHVI.

Methods

Gaussian Process Modeling

Nonparametric GP surrogate models were used to predict the value of an objective function $f(\mathbf{x})$ using Bayesian statistics [27]. The GP model includes a distribution of possible functions $y(\mathbf{x}) \sim \mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$. With $\mu(\mathbf{x})$ presenting the predicted mean of objective function $y(\mathbf{x})$, $k(\mathbf{x}, \mathbf{x}')$ is the kernel function based on the objective function behavior [20]. Given N previous measurements $\mathcal{D} = \{(x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)\}$, the predictive probability distribution of the function value is given by a normal distribution:

$$p(y|\mathcal{D}, \mathbf{x}) = \mathcal{N}(\mu(\mathbf{x}), \sigma^2(\mathbf{x})) \quad (1)$$

where

$$\mu(\mathbf{x}) = \mathbf{k}^T [\mathbf{K} + \sigma_n^2 \mathbf{I}]^{-1} \mathbf{y} \quad (2)$$

$$\sigma(\mathbf{x}) = k(\mathbf{x}, \mathbf{x}) - \mathbf{k}^T [\mathbf{K} + \sigma_n^2 \mathbf{I}]^{-1} \mathbf{k} \quad (3)$$

$$\mathbf{k} = [k(\mathbf{x}, \mathbf{x}_1), \dots, k(\mathbf{x}, \mathbf{x}_N)]^T \quad (4)$$

$$\mathbf{K} = \begin{bmatrix} k(\mathbf{x}_1, \mathbf{x}_1) & \cdots & k(\mathbf{x}_1, \mathbf{x}_N) \\ \vdots & \ddots & \vdots \\ k(\mathbf{x}_N, \mathbf{x}_1) & \cdots & k(\mathbf{x}_N, \mathbf{x}_N) \end{bmatrix} \quad (5)$$

where $\mathbf{y} = [y_1, y_2, \dots, y_N]^T$ and σ_n represents the noise hyperparameter. In this work, a Matérn kernel function was used for GP modeling:

$$k_{\text{Matérn}}(\mathbf{x}, \mathbf{x}') = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu} \|\mathbf{x} - \mathbf{x}'\|}{\mathcal{L}} \right)^\nu K_\nu \left(\frac{\sqrt{2\nu} \|\mathbf{x} - \mathbf{x}'\|}{\mathcal{L}} \right), \quad (6)$$

With $\|\mathbf{x} - \mathbf{x}'\|$ represents the Euclidean distance between inputs, Γ is the gamma function, and K_ν is the modified Bessel function. $\nu = 2.5$ was used in this work for modeling functions in the absence of prior information. The hyperparameter $\mathcal{L}$ determines the characteristic length-scale of the kernel, where small and large values correspond to rapidly and slowly varying functional behavior respectively.

Multi-Objective Bayesian Optimization

The goal of multi-objective optimization is to search the input parameter space of the experimental system and to optimize the vector objective function $\mathbf{y}(\mathbf{x}) = \{y_1(\mathbf{x}), y_2(\mathbf{x}), \dots, y_M(\mathbf{x})\}$. Usually, there is no best single solution $\mathbf{x}^*$ that can simultaneously optimize all objectives. Instead, multi-objective optimization algorithms attempt to find a set of points, known as the PF that optimally balances the trade-offs between the competing objectives. The Pareto Frontier set $\mathcal{P}$ is defined as the set of non-dominated points in objective space. For the minimization case,

an objective vector $\mathbf{y}(\mathbf{x})$ dominates another vector $\mathbf{y}(\mathbf{x}')$ if $y_m(\mathbf{x}) \leq y_m(\mathbf{x}')$ for all $m = 1, \dots, M$ and there exists at least one n such that $y_n(\mathbf{x}) < y_n(\mathbf{x}')$. HV is an often used metric to evaluate the quality of a PF. The HV is defined as the M-dimensional Lebesgue measure of the subspace dominated by $\mathcal{P}$ and bounded below by a reference point $\mathbf{r}$ [29]:

$$HV(\mathcal{P}) = \lambda_M(\cup_{\mathbf{y} \in \mathcal{P}} [\mathbf{y}, \mathbf{r}]) \quad (7)$$

With λ_M being the Lebesgue measure on $\mathbb{R}^M$. A larger HV corresponds to a better solution set $\mathcal{P}$. The MOBO technique aims to maximize the HV through the following process. Given a set of N observations: $D_N = \{(\mathbf{x}_1, \mathbf{y}_1), (\mathbf{x}_2, \mathbf{y}_2), \dots, (\mathbf{x}_N, \mathbf{y}_N)\}$, each objective is modeled as an independent GP surrogate model:

$$y_m(\mathbf{x}) \sim \mathcal{GP}_m[\mu_m(\mathbf{x}), k_m(\mathbf{x}, \mathbf{x}')] \quad (8)$$

The second part of MOBO is to construct an acquisition function to propose points which are likely to maximally increase the Pareto Frontier HV. A commonly used acquisition function is the Expected Hypervolume Improvement (EHVI) [29], which calculates the average increase in hypervolume using the posterior probability distribution of each objective function from the surrogate model. The EHVI acquisition function is defined as:

$$\alpha_{EHVI}(\mu, \sigma, \mathcal{P}, \mathbf{r}) := \int_{\mathbb{R}^M} H_I(\mathcal{P}, \mathbf{y}, \mathbf{r}) \cdot \mathcal{N}_{\mu, \sigma}(\mathbf{y}) d\mathbf{y} \quad (9)$$

With $H_I(\mathcal{P}, \mathbf{y}, \mathbf{r})$ being the hypervolume improvement from an observed point $\mathbf{y}$ in objective space, $\mathcal{N}_{\mu, \sigma}$ is the multivariate normal distribution function defined by the GP predicted mean μ and uncertainty σ . The hypervolume improvement is defined by the exclusive HV contribution to the current PF by adding $\mathbf{y}$ to the Pareto set. Combining the GP surrogate model and acquisition function, MOBO optimization can be performed. In each iteration, MOBO evaluates a candidate point in the parameter space and updates the GP surrogate model with the newly acquired data. The EHVI acquisition function is then used to calculate the next observation point in the parameters space. The process is performed iteratively and ultimately providing an approximation of the best achievable PF.

Determination of temporal length

The measurements of temporal length of electron pulses at IP was accomplished through interaction with quasi-single-cycle THz pulses generated via optical rectification of 800 nm laser pulses in a LiNbO3 crystal [37, 38]. An Off-Axis Parabolic (OAP) mirror was used to focus the THz pulses into a high field enhancement parallel-plate waveguide (PPWG) [39] streaking structure placed at IP. The THz pulses are then focused at the 70 μm slit in the PPWG and impart a transverse momentum kick to the electron beam which streaks the longitudinal profile onto the y axis on the downstream detector. The single-shot temporal resolution of THz streaking diagnostics was better than 1.5 fs with 3 μJ THz pulse energy.

Determination of q-resolution

The q-resolution quantifies the extent to which tiny features can be resolved within a given diffraction pattern. The instrumental q-resolution can be expressed as[41]:

$$\sigma_q = \sqrt{\sigma_{ND}^2 + \sigma_E^2} \quad (10)$$

σ_{ND} is the non-dispersive broadening which depends on transverse beam emittance and focusing solenoid settings. σ_E is the dispersive component which depends on beam energy spread ΔE and scattering wave vector q :

$$\sigma_E = \frac{\Delta E}{E_k} \left(\frac{\gamma}{\gamma + 1} \right) q \quad (11)$$

E_k and γ are the beam kinetic energy and Lorentz factor. During experiment, σ_{ND} was measured using the electron diffraction detector (phosphor screen coupled with an Electron-Multiplying CCD) located 3.2 m downstream the IP. Reciprocal space calibration was performed using a single crystal gold sample. Within the experimental parameter range, σ_E is calculated to be smaller than $0.011 \text{ \AA}^{-1}$, thus the q-resolution is dominated by σ_{ND} .

Data availability

The raw data generated in this study have been deposited in the Zenodo repository under accession code: <https://doi.org/10.5281/zenodo.11095450>[44]

Code availability

This research used open source Python software libraries, including BoTorch[42] and Xopt[43]. The data analysis codes are available from the corresponding author upon request.

References

- [1] D. Filippetto, P. Musumeci, R. K. Li, B. J. Siwick, M. R. Otto, M. Centurion and J. P. F. Nunes, "Ultrafast electron diffraction: visualizing dynamic states of matter," Review of Modern Physics, vol. 94, p. 045004 (2022)
- [2] SLAC Mega-electronvolt Ultrafast Electron Diffraction Instrument <https://lcls.slac.stanford.edu/instruments/mev-ued>
- [3] Kim, H.W., Vinokurov, N.A., Baek, I.H. et al. "Towards jitter-free ultrafast electron diffraction technology", Nat. Photonics 14, 245–249 (2020)
- [4] F. Qi et al., "Breaking 50 Femtosecond Resolution Barrier in MeV Ultrafast Electron Diffraction with a Double Bend Achromat Compressor", Phys. Rev. Lett. 124, 134803 (2020)

- [5] F. Ji, D. B. Durham, A. M. Minor, P. Musumeci, J. G. Navarro and D. Filippetto, "Ultrafast Relativistic Electron Nanoprobes", *Comm. Phys* 2, 54 (2019)
- [6] X. Shen et al., "Femtosecond gas-phase mega-electron-volt ultrafast electron diffraction", *Struct Dyn* 6, 054305 (2019)
- [7] J. P. F. Nunes et al., "Liquid-phase mega-electron-volt ultrafast electron diffraction", *Struct Dyn* 7, 024301 (2020)
- [8] T. J. A. Wolf et al., "The photochemical ring-opening of 1,3-cyclohexadiene imaged by ultrafast electron diffraction", *Nat. Chem.* 11, 504 (2019)
- [9] J. Yang et al. , "Simultaneous observation of nuclear and electronic dynamics by ultrafast electron diffraction", *Science* 368, 885-889(2020)
- [10] Yang, J., Dettori, R., Nunes, J.P.F. et al. "Direct observation of ultrafast hydrogen bond strengthening in liquid water", *Nature* 596, 531–535 (2021)
- [11] Sie, E.J., Nyby, C.M., Pemmaraju, C.D. et al. "An ultrafast symmetry switch in a Weyl semimetal". *Nature* 565, 61–66 (2019)
- [12] Aditya Sood et al. ,"Universal phase dynamics in VO2 switches revealed by ultrafast operando diffraction". *Science* 373,352-355(2021)
- [13] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser and I. Polosukhin, "Attention is All you Need", in *Advances in Neural Information Processing Systems(NIPS)* 2017
- [14] J. Jumper et al, "Highly accurate protein structure prediction with AlphaFold," *Nature* 596, 583-589 (2021)
- [15] J. Degraeve et al, "Magnetic control of tokamak plasmas through deep reinforcement learning," *Nature* 602, 414-419 (2022)
- [16] A. Edelen, N. Neveu, M. Frey, Y. Huber, C. Mayes, and A. Adelman, "Machine learning for orders of magnitude speedup in multiobjective optimization of particle accelerator systems," *Phys. Rev. Accel. Beams*, vol. 23, p. 044601 (2020)
- [17] A. Sheinker et al, "Demonstration of Model-Independent Control of the Longitudinal Phase Space of Electron Beams in the Linac-Coherent Light Source with Femtosecond Resolution," *Phys. Rev. Lett*, vol. 121, p. 044801 (2018)
- [18] S. C. Leemann et al, "Demonstration of Machine Learning-Based Model-Independent Stabilization of Source Properties in Synchrotron Light Sources," *Phys. Rev. Lett*, vol. 123, p. 194801 (2019)
- [19] J. Duris et al, "Bayesian Optimization of a Free-Electron Laser," *Phys. Rev. Lett*, vol. 124, p. 124801 (2020)

- [20] R. Roussel et al, "Multiobjective Bayesian optimization for online accelerator tuning," *Phys. Rev. Accel. Beams*, vol. 24, p. 062801 (2021)
- [21] A. Hanuka et al, "Physics model-informed Gaussian process for online optimization of particle accelerators," *Phys. Rev. Accel. Beams*, vol. 24, p. 072802 (2021)
- [22] R. Roussel et al, "Turn-key constrained parameter space exploration for particle accelerators using Bayesian active learning," *Nature Communications* 12, 5621 (2021)
- [23] R. Roussel et al, "Bayesian Optimization Algorithms for Accelerator Physics", *arXiv:2312.05667* (2023)
- [24] Kusne, A.G., Yu, H., Wu, C. et al. On-the-fly closed-loop materials discovery via Bayesian active learning. *Nat Commun* 11, 5966 (2020)
- [25] Abolhasani, M., Kumacheva, E. The rise of self-driving labs in chemical and materials sciences. *Nat. Synth* 2, 483–492 (2023)
- [26] H. R. Maier, S. Razavi, Z. Kapelan, L. S. Matott, J. Kasprzyk, and B. A. Tolson, Introductory overview: Optimization using evolutionary algorithms and other metaheuristics, *Environmental modeling & software* 114, 195 (2019)
- [27] C. E. Rasmussen and K. I. Williams, "Gaussian Processes for Machine Learning," *Adaptive Computation and Machine Learning*, MIT Press, Cambridge, MA, 2006
- [28] M. Emmerich, K. Yang, A. Deutz, H. Wang and C. M. Fonseca, "A multi-criteria generalization of Bayesian global optimization," *Advances in Stochastic and Deterministic Global Optimization*, Springer Optimization and Its Applications, pp. 229-242
- [29] Kaifeng Yang, Michael Emmerich, André Deutz, Thomas Bäck, "Multi-Objective Bayesian Global Optimization using expected hypervolume improvement gradient", *Swarm and Evolutionary Computation* 44, 945 (2019)
- [30] S. Daulton, D. Eriksson, M. Balandat and E. Bakshy, "Multi-Objective Bayesian Optimization over High-Dimensional Search Spaces", *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence*, 2022
- [31] Yang, K., Emmerich, M., Deutz, A. et al. Efficient computation of expected hypervolume improvement using box decomposition algorithms. *J Glob Optim* 75, 3–34 (2019)
- [32] S. Daulton, M. Balandat, and E. Bakshy, "Differentiable expected hypervolume improvement for parallel multi-objective Bayesian optimization," *Advances in*

- [33] Kogar, A., Zong, A., Dolgirev, P.E. et al. "Light-induced charge density wave in LaTe3", *Nat. Phys.* 16, 159–163 (2020)
- [34] Sood, A., Haber, J.B., Carlström, J. et al. "Bidirectional phonon emission in two-dimensional heterostructures triggered by ultrafast charge transfer", *Nat. Nanotechnol.* 18, 29–35 (2023)
- [35] Cao, Y., Fatemi, V., Fang, S. et al. "Unconventional superconductivity in magic-angle graphene superlattices", *Nature* 556, 43–50 (2018)
- [36] C. Xu, R. Roussel and A. Edelen, "Neural Network Prior Mean for Particle Accelerator Injector Tuning", in *Machine Learning and the Physical Sciences workshop, NeurIPS 2022*, New Orleans
- [37] E. C. Snively, M. A. K Othman et al, "Femtosecond compression dynamics and timing jitter suppression in a THz-driven electron bunch compressor," *Phys. Rev. Lett.*, vol. 124, p. 054801 (2020)
- [38] R. Li et al, "Terahertz-based subfemtosecond metrology of relativistic electron beams," *Phys. Rev. Accel. Beams*, vol. 22, p. 012803 (2019)
- [39] Mohamed A. K. Othman, Annika E. Gabriel, Emma C. Snively, Michael E. Kozina, Xiaozhe Shen, Fuhao Ji, Samantha Lewis, Stephen Weathersby, Praful Vasireddy, Duan Luo, Xijie Wang, Matthias C. Hoffmann, Emilio A. Nanni; Measurement of femtosecond dynamics of ultrafast electron beams through terahertz compression and time-stamping. *Appl. Phys. Lett.* 3 April 2023; 122 (14): 141107
- [40] S. Jalas et al, "Tuning curves for a laser-plasma accelerator," *Phys. Rev. Accel. Beams*, vol. 26, p. 071302 (2023)
- [41] J. Williams, F. Zhou, T. Sun, Z. Tao, K. Chang, K. Makino, M. Berz, P. M. Duxbury, C.-Y. Ruan; Active control of bright electron beams with RF optics for femtosecond microscopy. *Struct Dyn* 1 July 2017; 4 (4): 044035
- [42] M. Balandat, B. Karrer, D. R. Jiang, S. Daulton, B. Letham, A. G. Wilson, and E. Bakshy. BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. *Advances in Neural Information Processing Systems* 33, 2020.
- [43] R. Roussel., et al., "Xopt: A simplified framework for optimization of accelerator problems using advanced algorithms", in *Proc. IPAC'23, Venezia*.doi:<https://doi.org/10.18429/JACoW-14th International Particle Accelerator Conference-THPL164>
- [44] F. Ji et al, Raw dataset for 'Multi-Objective Bayesian Active Learning for MeV-ultrafast electron diffraction'. <https://doi.org/10.5281/zenodo.11095450> (2024)

Acknowledgments

This work was supported by U.S. Department of Energy Office of Science, Office of Basic Energy Sciences (Contract No. DE-AC02-76SF00515 and DE-SC0021984). M. M. acknowledges the support from DOE Fusion Energy Science under FWP No. 100182. The SLAC MeV-UED program is supported by the U.S. Department of Energy Office of Science, Office of Basic Energy Science under FWP 10075.

Author contributions

F. J., A. E. and R. J. E. designed the experiment. A. E., R. R., F. J. and C. M. wrote the software required to perform data collection and online multi-objective bayesian optimizations. F. J., A. E., X. S., S. M., S. W., D. L., M. M., A. R., M. A. K. O., P. K., X. W. and E. N. conducted the experiment. R. J. E. and M. M. provided guidance and oversights. F. J., R. J. E., X. S. and R. R. performed data analysis and wrote the manuscript. All authors contributed to the completion of the manuscript.

Competing interests

The authors declare no competing interests.




Fig. 1 Cartoon depicting the Multi-Objective Bayesian active Learning experiment at SLAC MeV-UED. Ultrafast electron pulses are generated using an S-band photoinjector and delivered to interaction point for conducting fs electron diffraction experiments. Electron beam properties are extracted from detector readouts and used for constructing GP surrogate models giving overview of machine response in the parameter space. A hypervolume improvement acquisition function is constructed for predicting the optimal next observation point and efficiently search the parameter space for obtaining PFs.




Fig. 2 Temporal length vs q-resolution optimization results. (a), (b) Overview of temporal length response projected to the ϕ - B_2 and B_1 - B_2 space. Gaussian Process surrogate model was constructed using measured data points(gray dots), color map(blue) represents the predicted mean of σ_t . (c) Overview of q-resolution response projected to the B_1 - B_2 space. The color map represents the predicted mean of σ_q . (d), (e) On-the-fly acquisition function projected to the ϕ - B_2 and B_1 - B_2 space, colormap shows the optimal next observation regions predicted by the acquisition function. (f) Measured data points in the objective(σ_t - σ_q) space, approximate Pareto Front points are highlighted with green stars. The corresponding hypervolume convergence plot is shown in Supplementary Fig. 1.




Fig. 3 Beam spot size vs q-resolution optimizations (a) Optimization results with 10 fC, 50 fC and 100 fC initial pulse charges. Measured data points are plotted in the σ_x - σ_q space, approximate Pareto Front points are highlighted. (b) Convergence plot with averaging Pareto Frontier hypervolume (blue) for 9 independent optimizations (gray) for the 10 fC initial pulse charge case, shading area denotes 1 sigma variance. Interpolated grid search result (red) is plotted for comparison.