



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Bayesian Optimization of a Free-Electron Laser

J. Duris, D. Kennedy, A. Hanuka, J. Shtalenkova, A. Edelen, P. Baxevanis, A. Egger, T. Cope, M. McIntire, S. Ermon, and D. Ratner

Phys. Rev. Lett. **124**, 124801 — Published 25 March 2020

DOI: [10.1103/PhysRevLett.124.124801](https://doi.org/10.1103/PhysRevLett.124.124801)

Bayesian optimization of a free-electron laser

J. Duris,¹ D. Kennedy,^{1,2} A. Hanuka,¹ J. Shtalenkova,¹ A. Edelen,¹
P. Baxevanis,¹ A. Egger,¹ T. Cope,¹ M. McIntire,³ S. Ermon,³ and D. Ratner¹

¹*SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA*

²*Department of Physics, University of California, Santa Cruz, Santa Cruz, CA 95064, USA*

³*Department of Computer Science, Stanford University, Stanford, CA 94305, USA*

(Dated: February 24, 2020)

The Linac Coherent Light Source X-ray free-electron laser is a complex scientific apparatus which changes configurations multiple times per day, necessitating fast tuning strategies to reduce setup time for successive experiments. To this end, we employ a Bayesian approach to maximizing X-ray laser pulse energy by controlling groups of quadrupole magnets. A Gaussian process model provides probabilistic predictions for the machine response with respect to control parameters, enabling a balance of exploration and exploitation in the search for the global optimum. We show that the model parameters can be learned from archived scans, and correlations between devices can be extracted from the beam transport. The result is a sample-efficient optimization routine, combining both historical data and knowledge of accelerator physics to significantly outperform existing optimizers.

Modern large-scale scientific experiments have complicated operational requirements, with performance degraded by errors in controls and dependencies on drifting or random variables. A prime example of this is the Linac Coherent Light Source (LCLS) [1], an x-ray free electron laser (FEL) user facility that supports a wide array of scientific experiments. At LCLS, skilled human operators tune dozens of control parameters on-the-fly to achieve custom photon beam characteristics, and this process cuts into valuable time allocated for each user experiment. Model-independent optimizers can help automate tuning, with successful demonstrations using simplex [2, 3], extremum seeking [4–6], and robust conjugate direction search [7, 8]. However, these methods require a large number of expensive acquisitions and can become stuck in local optima. To improve the efficiency of optimization beyond these methods, a model of the system is necessary [9]; in this work, we use Bayesian optimization with Gaussian process (GP) models during live tuning of LCLS. First, we use archived data to estimate the length scales of tuning parameters in the model. Second, we show that adding physics-inspired correlations between parameters further speeds convergence. Finally, we discuss possible directions for improvement.

Bayesian optimization is a sample efficient and gradient-free approach to global optimization of black-box functions with noisy outputs [10–12]. This efficiency comes from application of Bayes’ theorem to incorporate prior knowledge and previous steps to maximize the value of each new measurement. Numerical optimization of an acquisition function, incorporating the model’s expectation and uncertainty, guides the selection of each new point to sample. This gives Bayesian optimization the capability of balancing exploration with exploitation, which makes it ideal for optimizing noisy or uncertain targets. Prior knowledge such as historical data, simulation, or theory can improve the efficiency of optimization and constrain the search in low signal-to-

noise states. Bayesian optimization has been applied to a quickly growing number of domains; for example, resource prospecting [13, 14], active policy search for reinforcement learning [15], neural network parameter tuning [16], experimental control [17], etc. Large scientific experiments are expensive to run, have noisy and uncertain inputs and outputs, and are supported by detailed physical theory, simulations, and extensive data sets; all of these qualities make complex physics experiments well-suited for control via Bayesian optimization.

A Gaussian process (GP) is a popular choice of model for Bayesian optimization [18]. Whereas a Gaussian distribution is characterized by a mean and covariance $y \sim N(\boldsymbol{\mu}, \Sigma)$, a Gaussian *process* is determined by mean and covariance *functions*: $f(\boldsymbol{x}) \sim GP(\boldsymbol{\mu}(\boldsymbol{x}), k(\boldsymbol{x}, \boldsymbol{x}'))$, where \boldsymbol{x} are possible inputs to the objective. The mean $\boldsymbol{\mu}(\boldsymbol{x})$ is a function encoding the expected value of the objective function. Prior to updating the model with observations, $\boldsymbol{\mu}(\boldsymbol{x})$ may encode prior understanding of the objective, such as a fit to data or the most likely value of the objective given random samples on a relevant domain, or may be set to zero without loss of generality. As observations are added to the GP, the mean function begins to interpolate near measured points, and represents the most probable value of the objective given the measured data. The prediction’s uncertainty $\sigma(\boldsymbol{x})$ incorporates the measured data with a covariance function or kernel $k(\boldsymbol{x}, \boldsymbol{x}')$, which describes the similarity between pairs of points \boldsymbol{x} and \boldsymbol{x}' . Parameters of the mean and covariance functions are referred to as hyperparameters and are fit to a training data set by maximizing the probability of the data given the model, which incorporates a measure of the data fit and a natural model complexity penalty which discourages over-fitting (see Section 5.4.1 of [18]). Whereas the mean function may change day to day, the covariance function is fundamental to the underlying physics and thus is relatively constant and more robust for online estimation of the observed data.




FIG. 1: LCLS beamline showing locations of linacs, quadrupole magnets, and undulator magnets. A focusing system comprised of nearly 200 quadrupoles (upper lines denote focusing whereas lower lines denote defocusing for one plane) transports the electron beam downstream (from left to right above) and several groups of matching quadrupole magnets (shown in red) are routinely used to match the beam into the focusing system. An electron beam dump and X-ray pulse energy measurement system (not shown) are positioned downstream of the undulators. Our studies focus on 8 matching quads in the third linac and 4 in the transport line leading to the undulators.

In contrast to parametric models such as neural networks, GPs are non-parametric models, meaning that they are constructed directly from the training instances themselves. This allows for the model’s complexity to grow with the number of observations and predicts a class of functions consistent with a previously learned or physics motivated covariance structure near the current observations. Furthermore, in addition to the mean, GPs also provide predictions’ uncertainties, which are needed to construct the acquisition function. For both of these reasons, GPs are popular models for Bayesian optimization.

In this Letter we apply Gaussian process optimization to the problem of maximizing the LCLS x-ray FEL pulse energy. The FEL instability is a collective effect so performance depends strongly on the current density and therefore beam size along the FEL undulator line. A series of quadrupole magnets (with positions shown in Figure 1) transports and focuses the beam into the undulator line which is composed of 30 undulator modules of length 3.3 m, each followed by 0.6 m of space for diagnostics and corrector magnets to keep the beam on axis and quadrupole magnets to provide strong focusing and maintain a small beam size. The FEL pulse energy is therefore a function of all the quadrupole magnet strengths (field integral in kiloGauss, kG) and the input electron beam parameters. The machine is reconfigured at least twice per day, and hysteresis in magnets and motors throughout the machine necessitates retuning following each change. Consequently, LCLS operators perform tuning scans to optimize a subset (see Fig. 1) of the quadrupoles’ strengths many times per day. In 2015 we found that LCLS was spending more than 500 hours/year on the single task of quadrupole optimization due to the high dimensional parameter space. Reducing this lost experimental time became an important goal for operations.

A central component of Bayesian optimization is a system model. In our case, we model the FEL pulse energy dependence on quadrupole amplitude with a Gaussian

process. Given the FEL pulse energy’s smooth response to quadrupole magnets when all other parameters are fixed, we choose the popular radial basis function (RBF) kernel,

$$k_{\text{RBF}}(\mathbf{x}, \mathbf{x}') = \sigma_f^2 \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{x}')^T \Sigma (\mathbf{x} - \mathbf{x}')\right) \quad (1)$$

Here, \mathbf{x} and \mathbf{x}' are vectors with quadrupole values as coordinates, T denotes a matrix transpose operation, σ_f^2 is the covariance function amplitude, and Σ is a diagonal matrix of inverse square length scales. The length scales set the distance over which the function changes in each dimension. The amplitude captures the variance of the target function values with respect to variations in the inputs and therefore determines the prior prediction uncertainty far from any sampled points. The amplitude and length scale hyperparameters are chosen by fits to archived data; see the Supplemental Material for more details on the choice of the kernel and its hyperparameters [19]. Learning the covariance of the function (e.g. via a GP fit) rather than the function itself (e.g. a neural network fit) provides uncertainty estimates by favoring a subset of possible functions compatible with observations and allows for robust optimization since the peak of the objective may change day-to-day, but the general shape remains the same.

Online optimization proceeds by first measuring the initial state of the machine and then initializing the Gaussian process model with the appropriate kernel and first measured point. The GP provides a probabilistic surrogate model for the machine, and an upper confidence bound (UCB) acquisition function [20] is constructed from the GP prediction mean and standard deviation: $UCB(x) = \mu(x) + \kappa\sigma(x)$, where κ is of order unity. The UCB acquisition function employs a strategy of optimism in the face of uncertainty, balancing the trade-off between exploiting known promising regions (via the prediction mean term $\mu(x)$) and exploring uncertain areas (via the prediction uncertainty $\sigma(x)$). The point maximizing the UCB function is chosen for the next measurement, which

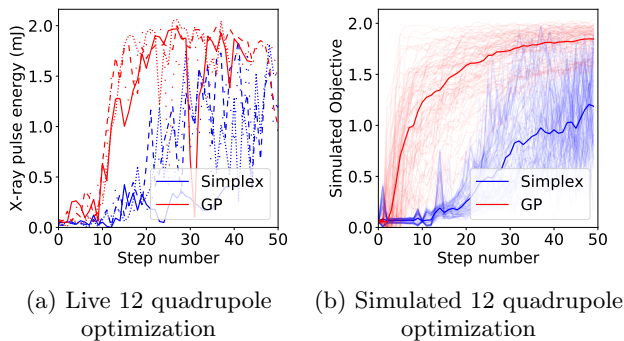


FIG. 2: (a) Comparison of optimization of FEL pulse energy over 12 quadrupole magnets for Bayesian optimization (red curves) vs Nelder-Mead simplex optimization (blue). Each scan was performed 4 times with identical starting conditions, shown with different dashing. Each step corresponds to approximately 3 to 4 seconds. (b) Simulations using the conditions of (a).

100 individual scans for each method, with means shown by thick lines, are consistent with measurements. Note occasional drops in FEL pulse energy, e.g. at step 30 in the solid red line, are caused by momentary machine glitches. It is encouraging that the optimizer is not thrown off course by such brief glitches.

is then acquired and added to the GP, finishing one step through the optimization process (see the Supplemental Material for details of our specific implementation [19]). The optimization continues in this way until reaching a time limit or a target performance.

There are many Gaussian process codes available [21–25], and many of these packages employ techniques to limit computation times for large data sets. In our case, we use an online GP model [26] interfaced to LCLS via the Ocelot optimization framework [3]. The online model saves frequent computations to speed subsequent predictions. In practice at LCLS, computation time per step is shorter than the acquisition time (1 to 2 seconds to set magnets and allow feedback systems to correct the trajectory and 1 second to measure FEL pulse energies for 120 shots).

Figure 2a shows results from live optimization of the FEL pulse energy simultaneously on 12 quadrupoles. In this example, Bayesian optimization is approximately 4 times faster than LCLS’s standard Nelder-Mead simplex algorithm, and reaches a higher optimum. The different lines for each algorithm correspond to scans with identical starting conditions.

We also compare simplex and Bayesian optimization in a simulation environment. Ideally we would use physics codes such as elegant [27] and Genesis [28] to model the transport and FEL behavior, however due to mismatch between the codes and measured performance, as well as computational expense for each simulation, we instead fit a beam transport model as described below. Though

the beam transport model does not capture the full complexity of the real machine, it allows us to compare the relative performance of simplex and Bayesian optimization with a simulated objective function, which we find consistent with live scans (Fig. 2b).

We can further improve the optimizer by leveraging our knowledge of accelerator physics to introduce correlations to the model. Correlated kernels become advantageous when a system’s response to one input dimension depends strongly on one or more of the other input dimensions as is exemplified with a correlated binormal test function in Figure 3a. Figure 3b shows a GP regression on noisy samples (RMS noise is 10% of the signal peak) from a correlated ground truth (Fig. 3a) with an isotropic kernel, while Figure 3c shows a GP regression on the *same* samples but with a kernel sharing the same correlation as the ground truth ($\rho = 0.8$). The latter model is more representative of the system. To demonstrate the effectiveness of a correlated model for regulating the system, we perform Bayesian optimization with and without kernel correlations for various dimensional spaces with nearest-neighbor correlation coefficients of $\rho_{i,i+1} = 0.5$. Each point in Figure 3d shows the average number of steps to achieve $> 90\%$ of the ground truth peak amplitude for 100 runs starting at a random position such that the starting signal-to-noise ratio is unity. The relative efficiency of the correlated kernel grows exponentially with the number of dimensions, making it attractive for high-dimensional optimization at accelerators.

FEL quadrupole optimization is an example of a highly correlated system. Strong focusing in charged particle transport relies on a series of oppositely polarized quadrupole fields [29]. Each quadrupole focuses in one transverse plane while defocusing in the orthogonal plane, and repeated application of alternate focusing/defocusing results in net focusing in both planes. Increasing the strength of one quadrupole field necessitates increasing the strength of the next quadrupole (with opposite sign) to achieve net focusing, resulting in negative correlations between nearby focusing elements. Figure 4a shows the average FEL pulse energy response to variation of two adjacent quadrupoles in a matching section just upstream of the FEL undulator magnets (near position 1400 m in Figure 1). Similarly, correlations exist between all pairs of quadrupoles.

To find the correlation hyperparameters, we exploit our knowledge of strong-focusing and FEL physics to calculate correlations from a beam physics transport model. The FEL pulse energy (denoted as U) is a correlation-preserving function of the transverse area of the beam, σ^2 , averaged along the interaction, with $\log U \propto \langle \sigma^2 \rangle^{-1/3}$ [30]. As a consequence, the beam size and pulse energy have similar correlations with respect to variations in the quadrupole magnets.

We model the average beam size in the undulator line from a linear transport model [29], with quadrupole values based on the archived settings. The result shown

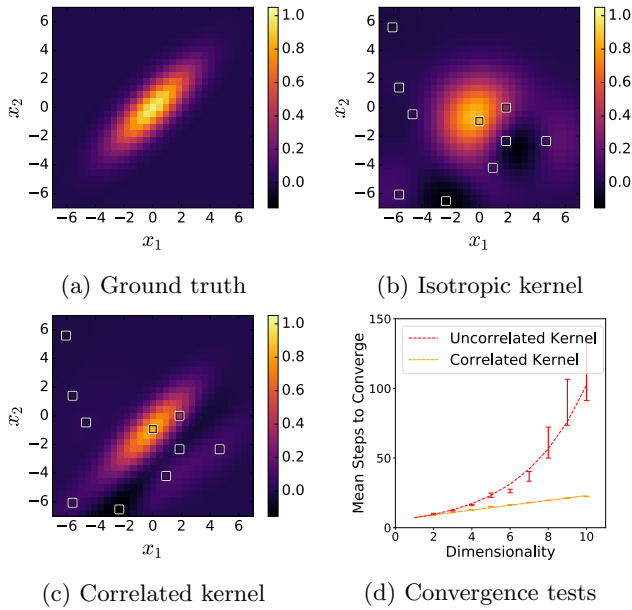


FIG. 3: (a) Ground truth of a 2D multinormal test function with unit slice widths (evaluated holding all but one coordinate fixed at a time) and correlation coefficient of 0.8. (b) GP regression with isotropic kernel on noisy samples from the ground truth. (c) GP regression with correlated kernel on identical samples as b. (d) Bayesian optimization convergence tests on a correlated ground truth with and without kernel correlations. Each bar shows the standard error about the mean for 100 trials. The correlated GP kernel (blue linear fit) performs as well as optimization of an isotropic ground truth with an isotropic GP kernel, growing linearly with the number of dimensions. Steps to convergence with mismatched kernel grows exponentially (red exponential fit).

in Figure 4b is a modeled beam size with correlations that match those of the measured FEL pulse energy (Fig. 4a). Finally, we combine the correlations from the linear model with the length scales fit from archived data to give the correlated RBF kernel. In principle, we would calculate correlation hyperparameters from fits to the archive data, as was done for the length scales. However, due to sparsity in the sampled data, we find that our modeled correlations are more accurate. Conversely, in principle we could follow the same modeling procedure to calculate the length scale hyperparameters, but because our simple model does not handle details such as FEL saturation, we found empirically that the length scales from archived data are more accurate.

We tested FEL optimization using GPs built with and without correlations for four adjacent matching quadrupole magnets located in the transport between the accelerator and the FEL undulators. The results, presented in Figure 4c, show that the correlated GP model outperforms the uncorrelated GP. Furthermore, we find

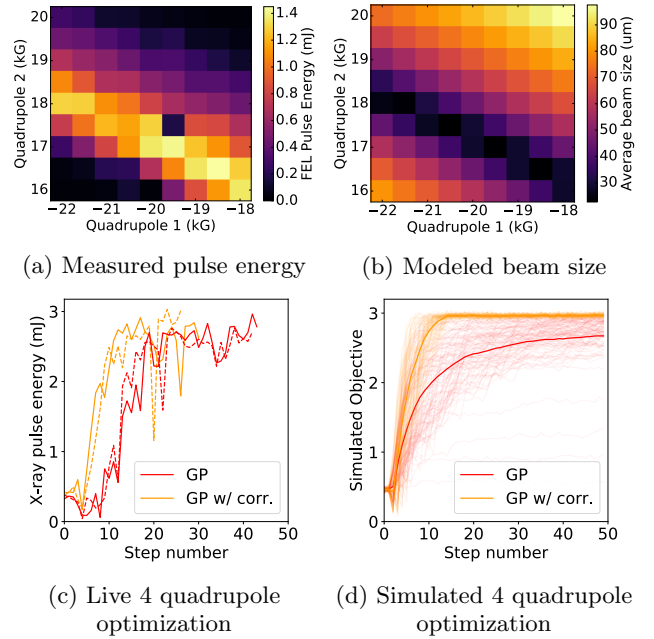


FIG. 4: (a) Average FEL pulse energy vs two adjacent quadrupoles. The dark spot near the center of the plot was due to a momentary machine malfunction. (b) The modeled average electron beam size in the undulator with the same quadrupoles shows the same correlations. (c) Optimization test for 4 quadrupoles: GP (red) vs GP with correlations (yellow). Each scan was performed twice with identical starting conditions, shown with different dashing. Each step takes approximately 3 to 4 seconds. (d) Simulations using a beam matrix model for the conditions of c. 100 individual scans for each method, with means shown by thick lines, are consistent with measurements.

consistency of these results with simulations based on correlations from the beam transport model as shown in Figure 4d. It is interesting to note that the Bayesian optimizer maximized the FEL with a similar number of steps for the 12 quadrupole case (Fig. 2a), seemingly in contradiction to Figure 3d. The similarity is due to the fact that to leading order (assuming a monoenergetic beam and linear optics), only four quadrupole magnets are needed to match the four Twiss parameters into the undulator line. However, optimizing more quadrupoles further increases the FEL pulse energy by reducing chromatic effects which increase the electron beam emittance and suppress FEL gain. While four quadrupoles can recover a significant fraction of peak performance, in practice operators typically cycle through subsets of all of the controllable quadrupoles. As the number of control parameters increases, the GP with correlations is expected to provide an exponentially growing benefit over the standard GP.

In this Letter, we showed that online Bayesian optimization with length scales estimated from archived

historical data can tune the LCLS FEL pulse energy more efficiently than the current standard, Nelder-Mead simplex algorithm. Moreover, we showed that adding physics-informed correlations, obtained from beam optics models, further improves tuning efficiency. This latter effect was demonstrated with four control parameters, and the improvement is expected to grow dramatically as the number of dimensions increases. The flexibility of the GP enables training from archived data, simulations, and physical models, and incorporating all this information into the model allows it to improve in efficiency and adapt to new tasks—all without the need for specialized hardware such as graphics processing units (GPUs) or high performance computing clusters.

We expect Bayesian optimization will become a standard tuning method for LCLS (and later for LCLS-II) when operations resume, improving efficiency and operational ability at SLAC as well as other accelerators worldwide. Whereas the present study focused on the most time-consuming task of tuning quadrupoles, we have plans for additional applications to other accelerator and beamline tasks, e.g. controlling FEL bandwidth, optimizing undulator field strengths [31], controlling electron beam collimation and peak current, FEL focusing and alignment, etc. More expressive GP models, such as deep kernel learning [32–34] or deep GPs [35], may help represent more complicated relationships between variables and extract additional value from historical data. Furthermore, adding ‘safety’ constraints to the acquisition function can guide exploration while ensuring operational requirements are met (for instance, avoiding transient drops in FEL pulse energy or keeping losses low) [36].

In our study, we exploited the physics of strong focusing to learn correlations between parameters while empirically learning relevant length scales. This was necessary since simulation codes that model the FEL process in the postsaturation regime are computationally expensive, limiting their usefulness in online tuning. However in other systems, simulation alone may suffice to calculate length scales and correlations [37]. Looking forward, these codes may be rewritten in a framework that supports automatic differentiation [38] to simplify and accelerate Hessian calculations. Alternatively, it may be possible to replace some simulation codes entirely with fast-to-evaluate surrogate models [39] which enable quick approximations of the system’s covariance. Physics abounds with well verified mathematical models, and incorporating additional physics knowledge into the model—either explicitly in the formulation of the kernel function as shown here or via additional training with simulation—can have a significant effect on optimization efficiency.

The authors are grateful to the SLAC Accelerator Operations group for help with live tests on LCLS, to Lauren Alsberg for a Python script to query the LCLS archive for data, and to Greg White for updated transport matrices.

We would also like to thank Hugo Slepicka and Sergey Tomin for help with the details of the Ocelot-optimizer code and to Tim Maxwell, Jeremy Mock, Ahmed Osman, and Hugo Slepicka for adding helpful LCLS-specific functions to Ocelot-optimizer. We are also appreciative to Gabriel Marcus and Brendan O’Shea for helpful discussions.

This work was supported by the Department of Energy, Laboratory Directed Research and Development program at SLAC National Accelerator Laboratory, under contract DE-AC02-76SF00515.

-
- [1] P. Emma, R. Akre, J. Arthur, R. Bionta, C. Bostedt, J. Bozek, A. Brachmann, P. Bucksbaum, R. Coffee, F.-J. Decker, Y. Ding, D. Dowell, S. Edstrom, A. Fisher, J. Frisch, S. Gilevich, J. Hastings, G. Hays, P. Hering, Z. Huang, R. Iverson, H. Loos, M. Messerschmidt, A. Miahnahri, S. Moeller, H.-D. Nuhn, G. Pile, D. Ratner, J. Rzeplia, D. Schultz, T. Smith, P. Stefan, H. Tompkins, J. Turner, J. Welch, W. White, J. Wu, G. Yocky, and J. Galayda, *Nature Photonics* **4**, 641 (2010).
 - [2] J. A. Nelder and R. Mead, *The Computer Journal* **7**, 308 (1965).
 - [3] S. Tomin, G. Geloni, I. Agapov, I. Zagorodnov, Y. Fomin, Y. Krylov, A. Valintinov, W. Colocho, T. Cope, A. Egger, and D. Ratner, *Proceedings of the 7th International Particle Accelerator Conference* (2016).
 - [4] A. Scheinker, X. Pang, and L. Rybarczyk, *Phys. Rev. ST Accel. Beams* **16**, 102803 (2013).
 - [5] A. Scheinker, D. Bohler, S. Tomin, R. Kammering, I. Zagorodnov, H. Schlarb, M. Scholz, B. Beutner, and W. Decking, *Phys. Rev. Accel. Beams* **22**, 082802 (2019).
 - [6] A. Scheinker, X. Huang, and J. Wu, *IEEE Transactions on Control Systems Technology* **26**, 336 (2018).
 - [7] X. Huang, *Phys. Rev. Accel. Beams* **21**, 104601 (2018).
 - [8] X. Huang and J. Safranek, *Phys. Rev. ST Accel. Beams* **18**, 084001 (2015).
 - [9] R. C. Conant and W. Ross Ashby, *International Journal of Systems Science* **1**, 89 (1970).
 - [10] J. Moćkus, *Optimization Techniques IFIP Technical Conference, Novosibirsk, July 1–7*, 400 (1975).
 - [11] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas, *Proceedings of the IEEE* **104**, 148 (2016).
 - [12] E. Brochu, V. M. Cora, and N. de Freitas, *Univ. of British Columbia Tech. Rep. UBC TR-2009-023 and arXiv:1012.2599* (2010).
 - [13] N. Cressie, *Mathematical Geology* **22**, 239 (1990).
 - [14] J. P. Delhomme, *Advances in Water Resources* **1**, 251 (1978).
 - [15] R. Martinez-Cantin, N. de Freitas, E. Brochu, J. Castellanos, and A. Doucet, *Autonomous Robots* **27**, 93 (2009).
 - [16] J. Snoek, H. Larochelle, and R. P. Adams, *Advances in Neural Information Processing Systems 25 (NIPS 2012)*, 2951 (2012).
 - [17] M. M. Noack, K. G. Yager, M. Fukuto, G. S. Doerk, R. Li, and J. A. Sethian, *Scientific Reports* **9**, 11809 (2019).
 - [18] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (MIT Press,

- 2006).
- [19] See Supplemental Material at [URL will be inserted by publisher] for details of the kernel hyperparameter calculation and the acquisition function optimization.
- [20] N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger, *IEEE Transactions on Information Theory* **58**, 3250 (2012).
- [21] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, *Journal of Machine Learning Research* **12**, 2825 (2011).
- [22] “GPpy: A Gaussian process framework in python,” <http://github.com/SheffieldML/GPy> (since 2012).
- [23] A. G. d. G. Matthews, M. van der Wilk, T. Nickson, K. Fujii, A. Boukouvalas, P. León-Villagrà, Z. Ghahramani, and J. Hensman, *Journal of Machine Learning Research* **18**, 1 (2017).
- [24] C. E. Rasmussen and H. Nickisch, *Journal of Machine Learning Research* **11**, 3011 (2010).
- [25] J. Salvatier, T. V. Wiecki, and C. Fonnesbeck, *PeerJ Computer Science* **2**, e55 (2016).
- [26] M. McIntire, D. Ratner, and S. Ermon, *Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence UAI16*, 517 (2016).
- [27] M. Borland, Argonne National Lab. Tech. Rep. No. LS-287 (2000), 10.2172/761286.
- [28] S. Reiche, *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **429**, 243 (1999).
- [29] H. Wiedemann, *Particle Accelerator Physics* (Springer Berlin Heidelberg, Berlin, Heidelberg, 2007).
- [30] K. Kim, Z. Huang, and R. Lindberg, *Synchrotron Radiation and Free-Electron Lasers* (Cambridge University Press, 2017).
- [31] J. Wu, X. Huang, T. Raubenheimer, and A. Scheinker, in *38th International Free Electron Laser Conference* (JACOW, Geneva, Switzerland, 2018) pp. 229–234.
- [32] R. Calandra, J. Peters, C. E. Rasmussen, and M. P. Deisenroth, *Proceedings of the International Joint Conference on Neural Networks*, 3338 (2016).
- [33] A. G. Wilson, Z. Hu, R. R. Salakhutdinov, and E. P. Xing, *Advances in Neural Information Processing Systems* **29**, 2586 (2016).
- [34] A. Jacot, F. Gabriel, and C. Hongler, *Advances in Neural Information Processing Systems* **31**, 8571 (2018).
- [35] A. C. Damianou and N. D. Lawrence, *Proceedings of the 16th International Conference on Artificial Intelligence and Statistics (AISTATS)* **31**, 207 (2013).
- [36] J. Kirschner, M. Mutný, N. Hiller, R. Ischebeck, and A. Krause, *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 3429 (2019).
- [37] A. Hanuka, J. Duris, J. Shtalenkova, D. Kennedy, A. Edelen, D. Ratner, and X. Huang, *Proceedings of the Machine Learning for the Physical Sciences Workshop, NeurIPS* (2019).
- [38] J. Degraeve, M. Hermans, J. Dambre, and F. Wyffels, *Frontiers in Neurobotics* **13**, 6 (2019).
- [39] A. Edelen, N. Neveu, C. Emma, D. Ratner, and C. Mayes, *Proceedings of the Machine Learning for the Physical Sciences Workshop, NeurIPS* (2019).
- [40] D. Duvenaud, J. Lloyd, R. Grosse, J. Tenenbaum, and G. Zoubin, *Proceedings of the 30th International Conference on Machine Learning* **28**, 1166 (2013).
- [41] S. Sun, G. Zhang, C. Wang, W. Zeng, J. Li, and R. Grosse, *Proceedings of the 35th International Conference on Machine Learning*, 4828 (2018).
- [42] R. Akre, D. Dowell, P. Emma, J. Frisch, S. Gilevich, G. Hays, P. Hering, R. Iverson, C. Limborg-Deprey, H. Loos, A. Miahnahri, J. Schmerge, J. Turner, J. Welch, W. White, and J. Wu, *Phys. Rev. ST Accel. Beams* **11**, 030703 (2008).