

---

# Online Bayesian Optimization of Electron Beam Phase Space Shaping

---

**Jingyi Tang**

SLAC National Laboratory  
Stanford University  
Stanford, CA 94305  
jytang@stanford.edu

**Joseph Duris**

SLAC National Laboratory  
Stanford University  
Stanford, CA 94305  
jduris@slac.stanford.edu

**Auralee Linscott Edelen**

SLAC National Laboratory  
Stanford University  
Stanford, CA 94305  
edelen@slac.stanford.edu

**Adi Hanuka**

SLAC National Laboratory  
Stanford University  
Stanford, CA 94305  
adiha@slac.stanford.edu

**Agostino Marinelli**

SLAC National Laboratory  
Stanford University  
Stanford, CA 94305  
marinelli@slac.stanford.edu

## Abstract

Particle accelerators require adjustment of system settings to control the final electron beam characteristics for different applications. At present, this optimization process is often done manually by human operators by visual inspection. Electron beam transport can involve collective effects, making their control not intuitive for human operators. In addition, accelerators are often outfitted with image-based beam measurements that can be challenging to incorporate directly into online optimization algorithms. In this paper, we introduce an online optimization method that directly incorporates images of the electron beam as well as a desired target distribution into a metric in order to automatically shape the beam towards the target distribution. We compare various metrics quantifying the distance between the measured and desired distributions and employ one to guide the beam shaping via Bayesian optimization.

## 1 Introduction

Particle accelerators require adjustment of system settings to control the final electron beam characteristics for different applications. At present, this optimization process is often done manually by human operators by visual inspection. Electron beam transport can involve collective effects, making their control not intuitive for human operators. In addition, accelerators are often outfitted with image-based beam measurements that can be challenging to incorporate directly into online optimization algorithms.

Previous work has used Bayesian optimization with Gaussian Processes to control scalar output variables for linear accelerators [1] and storage rings [2]. In addition, preliminary studies have examined the use of neural networks to provide suggested initial settings for target beam images [3]. In that case, previously-recorded data was used to train the neural network, and so it is limited to setups or parameter spaces which were previously explored. On its own, it is also not suitable for fine-tuning of new setups or compensating for changes in the machine response over time. A different approach is needed to address the problem of algorithmically setting up a new beam shape in a previously unexplored region of the parameter space. This is particularly important for accelerators that regularly must accommodate new beam shape requests (e.g. [4, 5, 6, 7, 8]), where data will only be incrementally available as the search is conducted.

In this paper, we introduce an online optimization method that directly incorporates images of the electron beam as well as a desired target distribution into a metric in order to automatically shape the beam towards the target distribution. We compare various metrics quantifying the distance between the measured and desired distributions and employ one to guide the beam shaping via Bayesian optimization.

## 2 Method

### 2.1 Metrics for comparing the distributions

There are multiple metrics for comparing the similarity of the two distributions. The straightforward one is the root mean square error (RMSE) which calculate the mean error of each point of the distribution.  $f$ -divergences such as Kullback–Leibler divergence (KL) and Jensen–Shannon divergence (JS) are widely used on probability and machine learning study to compare two probability distribution. Another class of metrics calculate the "energy" difference of two distributions, such as the energy distance and earth mover’s distance (EMD, or Wasserstein distance), which are related to the "energy" needed to transform one distribution into the other.

Although all these metrics measure the differences between two distributions, they have different features that are suitable for different situations. As an example, in Fig. 1 we calculated the EMD, RMSE and KL of two normal distributions displaced relative to each other. All of three metrics reach a minimum when the two distributions are overlapped. When the offset increases and the two distributions become disjoint, RMSE and KL become constant while EMD continues to grow. This property has made the EMD metric useful for training generative adversarial networks[9] which necessarily compare low dimensional distributions embedded within high dimensional spaces.

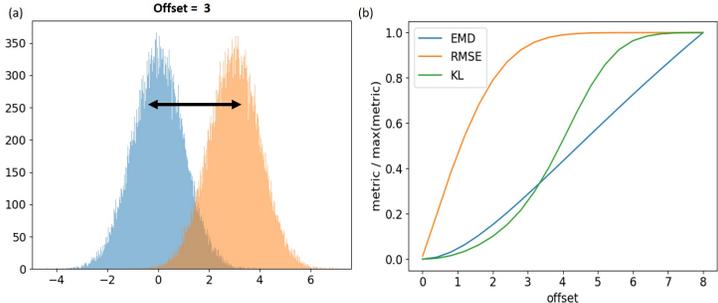


Figure 1: Comparing metrics for two normal distributions as one is displaced relative to the other. Here, EMD is the earth mover’s distance, RMSE is the root mean square error, and KL is the Kullback–Leibler Divergence.

### 2.2 Bayesian optimization

We use Bayesian optimization to control the parameters of the physical system in order to minimize the difference between the target and measured distributions. In Bayesian optimization, at each iteration a probabilistic surrogate model is used in conjunction with an acquisition function to select the next point to sample [10, 11]. The model is updated with new data as the space is searched. The acquisition function enables the trade-off between exploration of the space and exploitation of learned information to be balanced during the search. The Gaussian Process (GP) is an appealing choice for a surrogate model because it produces uncertainty estimates and can provide reliable predictions in the low-data regime (which is ideal for online training). The metrics quantifying the difference between distributions described below are roughly parabolic in shape for deviations from their optima. We negate and exponentiate these metrics to create targets which are roughly Gaussian shaped. In this work, we use a radial basis function (RBF) with automatic relevance determination plus white noise as our kernel to model the resultant objective. In the absence of prior data, we optimize the selection of kernel hyperparameters at each step by maximizing the GP marginal likelihood.

### 3 Physical system

In this work we use a simplified particle simulation of optical compression for enhanced self amplified spontaneous emission FEL(Free Electron Lasers) scheme [12] as an example, in which case the 2D longitudinal phase space (time vs. energy) of electron beam can be manipulated by controlling the collective force on the beam via laser modulation, compression, and drift. Electrons uniformly distributed in the time (Fig. 2(a)) first go through an energy modulator (Fig. 2(b)), which introduces a periodic energy modulation and varies the mean energy along the electron bunch. A compression section (Fig. 2(c)) turns the energy modulation into a density or current modulation and changes the time duration of the electron bunch (FWHM current). Finally, a space charge drift (Fig. 2(d)) allows the collective effects within the beam to further modify the 2D phase space distribution. The longitudinal space charge force is proportional to the derivative of the current distribution and increases the energy spread of the beam as it drifts in free space.

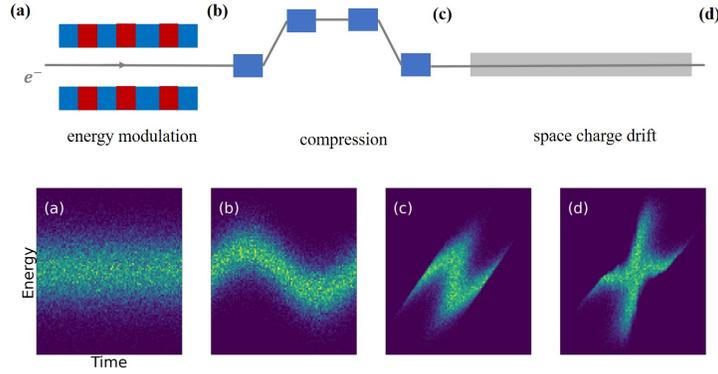


Figure 2: Physics system. The upper figure shows the physics model we used for the optical compression simulation [12]. The longitudinal phase space of initial electron beam (a) is manipulated via energy modulation (b), compression (c), and space charge drift (d).

The amplitude, period and phase of the energy modulation, the mean and rms beam energy, the compression factor and the length of the space charge drift serve as the input features for the Bayesian optimizer. The goal is to minimize the difference between the target 2D distribution and the simulated final phase space. The target distribution can be abstract and not existent in the solution set of the physics model. It can be in an ideal shape that reflects the main features people desire. In this work we focus on creating a "spike" in the phase space, which is represented by a beam distributed uniformly in a rectangular area, with constrained length in time and extended in energy (blue areas in Fig. 3(b)).

#### 3.1 Comparing the metrics

The metric that compares the differences of the simulated distribution from the ideal target can be essential for the success of the optimization. As discussed in Section 2 it should have a minimum where the two distributions are the same, and should continuously increase for smooth transforms of the distribution from the goal. Second it should be able to provide effective information even when the two distributions are totally disjoint. For these reasons we choose EMD as our metric.

Another complication for this particular case is that there is a local minimum when the simulation produces multiple spikes while targeting a single spike. As shown in Fig. 3(a), when the energy modulation period approaches zero, both EMD (yellow line) and KL divergence (green line) show a deep local minimum. This can be problematic because it can takes numerous random steps during the optimization to get out of the local minimum and thus slow down the optimization process. In order to improve the performance, we calculate the KL divergence in  $k$  space (e.g. between the power spectral density or square magnitude of the 2D Fourier transform of the two distributions) and multiply the KL divergence with EMD in real space as a combined metric, as shown by the blue line in Fig. 3(a). We use KL-divergence in  $k$  space because it is more computationally efficient compared with EMD and displacement is no longer a problem in  $k$  space since the spectral density is always centered at the origin point. The product allows us to avoid scaling each metric so that one

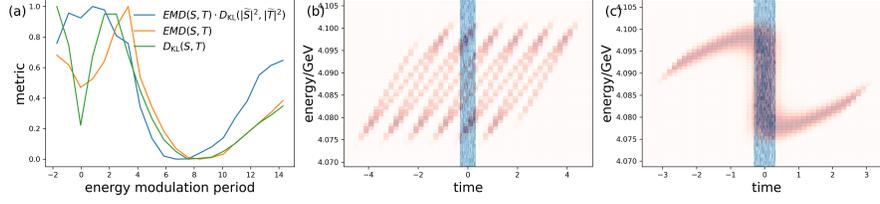


Figure 3: Comparison between metrics. (a) KL divergence of distributions in phase space (green), EMD (orange) and combined metric (blue) as a function of energy modulation period. Here the metrics are scaled such that all of them lies in  $(0,1)$ .  $S$  and  $T$  are the simulated and target phase space respectively and  $\tilde{S}, \tilde{T}$  are their Fourier transform. (b) The simulated phase space (red) and the target (blue) when energy modulation period is small (period=1). (c) The simulated phase space (red) and the target (blue) when energy modulation period is large (period=7).

doesn't dominate the other and we can be assured that neither will be zero since the target is not a valid solution in this case. The high frequency components of multiple spikes leads to a large KL divergence and therefore the effect of the local minimum is mitigated.

### 3.2 Automatic beam shaping tests

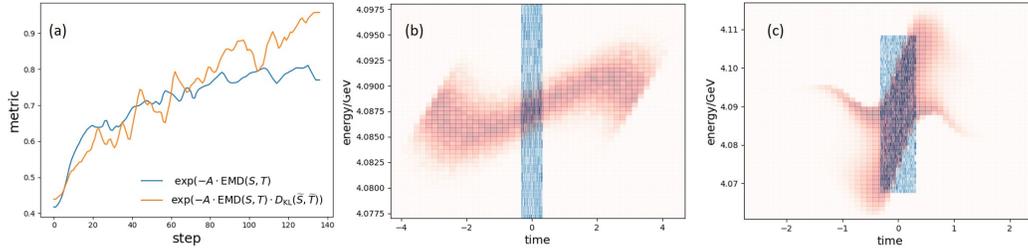


Figure 4: (a) Comparison of optimization with EMD (blue) and combined metric (orange).  $S$  and  $T$  are the simulated and target phase space respectively and  $\tilde{S}, \tilde{T}$  are their Fourier transform.  $A$  is a scaling factor such that those two metrics are comparable (b) One example of simulated phase space (red) and target (blue) at the beginning of the optimization. (c) One example of simulated phase space (red) and target (blue) at the end of the optimization.

We tested the described shaping method by calculating the EMD with POT [13], the KL divergence with SciPy [14], and performing the Gaussian process regression with scikit-learn [15]. Fig. 4 shows the results of Bayesian optimization with EMD and the combined metric. Figure 4a shows the presented metrics are averaged over 10 runs. With the combined metric, the algorithm generally produces a single spike after 60 steps. As comparison, the EMD metric converges more slowly due the effects of the local minimum. Additionally for the runs without the spectral comparison, we occasionally observe multiple spikes. This is not observed when the spectra of the distributions are compared as well.

## 4 Conclusion

We frame the problem of beam shaping as an optimization problem, and use Bayesian optimization with a distance metric between measured and target distributions as an objective. We find that the earth mover distance in conjunction with the KL divergence of the spectra of the measured and target distributions to be most robust for this beam shaping task. This proof of principle demonstration here controls a simplified beam model neglecting transverse effects. Initial studies using this technique with higher fidelity simulations and different shaping mechanisms are promising but require more work. For example, constraints of the system may make it impractical to place a current spike anywhere on the beam, yet by comparing the power spectral densities of the observed and target

current distributions or their autocorrelations would allow solutions with a spike imprinted anywhere on the beam.

The lasing medium of a X-ray free-electron laser (XFEL) is an electron beam allowing control of the produced X-ray pulse via control of the energy and current of the electron beam. Other applications may be targeting the measured properties of the produced X-ray pulses such as the measured spectrum. As the number of XFEL schemes based on beam shaping proliferate [4, 5, 6, 7, 8] and the number of concurrent operating beamlines multiplies at LCLS-II [16], the need increases for more automatic tools to assist experimental setups. Whereas existing methods for setup require visual inspection [17] or previously measured setups [3], the approach in this study may enable faster initial control and discovery of new setups.

## Broader Impact

Modern scientific facilities often involve complicated operations and the optimization of such facilities requires experienced operators and can take up a large amount of research time. Moreover, the high-dimensional parameter space and complicated collective effects can make the control not intuitive for human operators. In this paper, we introduce an online optimization method that directly incorporates images-based target into a metric which can enable automatic phase space shaping of electron beam in particle accelerators. This approach can be applicable in a wide range of research areas to improve efficiency of research activities and boost the discovery of new experimental setups. Possible applications exist in a number of beam shaping based XFEL experiments [4, 5, 6, 7, 8, 18, 19] and the multiple concurrent operating beamlines at LCLS-II [16], where it can work as a automatic tools to assist setups. The algorithmic beam shaping can also play an important role in the study of extreme compact high energy accelerators for which a ramped drive beam can increase transformer ratio in plasma acceleration[20]. More potential applications may be explored shaping the spectrum for XFEL and HHG[21]. Failure of the system can possibly lead to slower tuning and failure to find a desired phase space. In this paper the data are based on simulation, while real measurements of the phase space may be resolution limited to few femtoseconds due to finite beam size. Negative ethical or societal impacts are not applicable for this paper.

## Acknowledgments

This work was supported by U.S. Department of Energy (DOE) contract No. DE-AC02-76SF00515

## References

- [1] Joseph Duris, Dylan Kennedy, Adi Hanuka, Jane Shtalenkova, Auralee Edelen, P Baxevanis, A Egger, T Cope, M McIntire, S Ermon, et al. Bayesian optimization of a free-electron laser. *Physical Review Letters*, 124(12):124801, 2020.
- [2] Adi Hanuka, X Huang, J Shtalenkova, D Kennedy, A Edelen, VR Lalchand, D Ratner, and J Duris. Physics-informed gaussian process for online optimization of particle accelerators. *arXiv preprint arXiv:2009.03566*, 2020.
- [3] Alexander Scheinker, Auralee Edelen, Dorian Bohler, Claudio Emma, and Alberto Lutman. 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.*, 121:044801, Jul 2018.
- [4] Y. Ding, K. L. F. Bane, W. Colocho, F.-J. Decker, P. Emma, J. Frisch, M. W. Guetg, Z. Huang, R. Iverson, J. Krzywinski, H. Loos, A. Lutman, T. J. Maxwell, H.-D. Nuhn, D. Ratner, J. Turner, J. Welch, and F. Zhou. Beam shaping to improve the free-electron laser performance at the linac coherent light source. *Phys. Rev. Accel. Beams*, 19:100703, Oct 2016.
- [5] A. Marinelli, R. Coffee, S. Vetter, P. Hering, G. N. West, S. Gilovich, A. A. Lutman, S. Li, T. Maxwell, J. Galayda, A. Fry, and Z. Huang. Optical shaping of x-ray free-electron lasers. *Phys. Rev. Lett.*, 116:254801, Jun 2016.

- [6] S Huang, Y Ding, Y Feng, E Hemsing, Z Huang, J Krzywinski, AA Lutman, A Marinelli, TJ Maxwell, and D Zhu. Generating single-spike hard x-ray pulses with nonlinear bunch compression in free-electron lasers. *Physical review letters*, 119(15):154801, 2017.
- [7] Alberto A. Lutman, Marc W. Guetg, Timothy J. Maxwell, James P. MacArthur, Yuantao Ding, Claudio Emma, Jacek Krzywinski, Agostino Marinelli, and Zhirong Huang. High-power femtosecond soft x rays from fresh-slice multistage free-electron lasers. *Phys. Rev. Lett.*, 120:264801, Jun 2018.
- [8] Joseph Duris, Siqi Li, Taran Driver, Elio G Champenois, James P MacArthur, Alberto A Lutman, Zhen Zhang, Philipp Rosenberger, Jeff W Aldrich, Ryan Coffee, et al. Tunable isolated attosecond x-ray pulses with gigawatt peak power from a free-electron laser. *Nature Photonics*, 14(1):30–36, 2020.
- [9] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.
- [10] Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.
- [11] Eric Brochu, Vlad M Cora, and Nando De Freitas. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*, 2010.
- [12] Alexander A. Zholents. Method of an enhanced self-amplified spontaneous emission for x-ray free electron lasers. *Phys. Rev. ST Accel. Beams*, 8:040701, Apr 2005.
- [13] R’emi Flamary and Nicolas Courty. Pot python optimal transport library, 2017.
- [14] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020.
- [15] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12(85):2825–2830, 2011.
- [16] J Stohr. Linac coherent light source ii (lcls-ii) conceptual design report. 11 2011.
- [17] Y. Ding, C. Behrens, P. Emma, J. Frisch, Z. Huang, H. Loos, P. Krejcik, and M-H. Wang. Femtosecond x-ray pulse temporal characterization in free-electron lasers using a transverse deflector. *Phys. Rev. ST Accel. Beams*, 14:120701, Dec 2011.
- [18] Zhirong Huang, M Borland, P Emma, J Wu, C Limborg, G Stupakov, and J Welch. Suppression of microbunching instability in the linac coherent light source. *Physical Review Special Topics-Accelerators and Beams*, 7(7):074401, 2004.
- [19] Jingyi Tang, Randy Lemons, Wei Liu, Sharon Vetter, Timothy Maxwell, Franz-Josef Decker, Alberto Lutman, Jacek Krzywinski, Gabriel Marcus, Stefan Moeller, et al. Laguerre-gaussian mode laser heater for microbunching instability suppression in free-electron lasers. *Physical Review Letters*, 124(13):134801, 2020.
- [20] Ryan Roussel, Gerard Andonian, Walter Lynn, Kunal Sanwalka, River Robles, Claire Hansel, Aihua Deng, Gerard Lawler, JB Rosenzweig, G Ha, et al. Single shot characterization of high transformer ratio wakefields in nonlinear plasma acceleration. *Physical Review Letters*, 124(4):044802, 2020.

- [21] Randy Bartels, S Backus, E Zeek, L Misoguti, G Vdovin, IP Christov, MM Murnane, and HC Kapteyn. Shaped-pulse optimization of coherent emission of high-harmonic soft x-rays. *Nature*, 406(6792):164–166, 2000.