Online accelerator optimization with a machine learning-based stochastic algorithm

MACHINE
LEARNING
Science and Technology

OPEN ACCESS

RECEIVED
30 June 2020

REVISED
17 October 2020

ACCEPTED FOR PUBLICATION
5 November 2020

PUBLISHED
24 December 2020

PAPER

Online accelerator optimization with a machine learning-based stochastic algorithm

Zhe Zhang[1,3], Minghao Song[1,2,3] and Xiaobiao Huang[1]

[1] SLAC National Accelerator Laboratory, Menlo Park, CA 94025, United States of America
[2] Department of Physics, Illinois Institute of Technology, Chicago, IL 60616, United States of America
[3] Both authors contributed equally to this work.

E-mail: xiahuang@slac.stanford.edu

Keywords: online optimization, Gaussian process, evolutionary algorithm

Abstract

Online optimization is critical for realizing the design performance of accelerators. Highly efficient stochastic optimization algorithms are needed for many online accelerator optimization problems in order to find the global optimum in the non-linear, coupled parameter space. In this study, we propose to use the multi-generation Gaussian process optimizer for online accelerator optimization and demonstrate that the algorithm is significantly more efficient than other stochastic algorithms that are commonly used in the accelerator community.

1. Introduction

Automated online optimization of accelerators has become increasingly more important in recent years as new machines continue to push the limit in performance and in turn, the operation challenges. For example, the diffraction limited storage rings built on multi-bend achromat lattices reduce the emittance by a factor of 10–100 from the level of third generation storage rings of the same size. In the mean time, the new rings become very challenging in non-linear beam dynamics as more and stronger sextupole magnets are employed in the lattices to correct the chromatic aberrations. Because of the many error sources in the lattice, the non-linear beam dynamics performance of a storage ring, namely, the dynamic aperture and the local momentum aperture, can be significantly different from that of the design. Online optimization is an effective method to restore the performance of the machine [1–4].

In another example, the photon power of a free electron laser depends strongly on the transverse profile of the electron beam in the undulators. Quadrupole magnets are used to adjust the transverse optics in the transport line before the undulator section to match the beam profile to the ideal setting that maximizes the photon beam power. However, since both the initial distribution of the electron beam and the ideal setting are not precisely known, in practice, the quadrupole set points are found by online optimization [5].

In online optimization, control knobs, such as the set points of magnet power supplies, are varied to optimize an objective function, which is a measure of the machine performance. Online optimization for complex problems with multiple variables requires advanced algorithms that are suitable for online applications. Traditional optimization algorithms, could be sensitive to the inevitable measurement noise in the objective function. The robust conjugate direction search (RCDS) method was developed for online optimization of noisy functions and has been successfully applied to many accelerators [1–4, 6–11]. More recently, a robust simplex (RSimplex) algorithm was also proposed for online optimization [12]. However, the RCDS and RSimplex algorithms are local optimizers, which seek local optimum in the vicinity of the starting point and can be attracted to and trapped by one on its convergence path.

In many applications, there can be potentially multiple local minima in the complex parameter space, while it is critical to find the global optimum. Stochastic optimization algorithms are often used to locate the global optimum in such cases as they employ stochastic operations that search the parameter space more completely and enable them to break out from local minima. There are also many cases where it is desirable to optimize multiple objective functions simultaneously. Two multi-objective stochastic optimization algorithms, non-dominated sorting genetic algorithm II (NSGA-II) [13] and multi-objective particle swarm optimization [14], are commonly used in the accelerator community for design optimizations [15–17], including optimization of dynamic aperture and local momentum aperture in simulations. They have also found use in online optimization of single objective problems, for example, in reference [10, 18] for NSGA-II and in reference [1, 19] for particle swarm optimization (PSO), respectively.

However, stochastic optimization algorithms are usually inefficient. It takes many evaluations of the objective functions to converge to the global optimum. This is because the trial solutions proposed by the stochastic operations are not making full use of the information in the data sampled from the parameter space. And, in the case of NSGA-II, often many redundant solutions are tried in the vicinity of existing good solutions, leading to a decreased efficiency and potentially causing premature convergence. Measurement noise for the objective function further reduces the efficiency of NSGA-II as solutions biased by noise tend to be selected to spawn new trial solutions, which defeat the working mechanism of the algorithm [20]. In the application of minimization of the vertical emittance of a storage ring, the NSGA-II algorithm took about 20 000 evaluations to reach the optimum [18], while PSO took about 3000 evaluations [19] and RCDS took only about 300 evaluations [20]. Therefore, efficient stochastic optimization algorithm suitable for online applications are needed for the accelerator community.

In this study, we propose to use the multi-generation Gaussian process optimizer (MG-GPO) for online optimization of problems with complex parameter spaces and demonstrate its application in both simulated and real-life experiments. The method is also naturally applicable to multi-objective problems. MG-GPO was originally proposed for design optimization [21]. It uses posterior Gaussian process (GP) regression models to filter for good trial solutions in an iterative process similar to NSGA-II or PSO. It has been demonstrated in storage ring non-linear lattice design optimization [22].

This paper is organized as follows. Section 2 briefly describes the MG-GPO algorithm. Section 3 shows the applications of MG-GPO to analytic and simulated problems. Section 4 presents experimental results when the method is applied to a real accelerator. Section 5 gives the conclusion.

2. The MG-GPO algorithm

The MG-GPO algorithm leverages the observed data points (from the evaluations of the objective function) to guide the evolution direction of the stochastic operations [21]. It utilizes the observed data to build a GP model that approximate the objective function. In each generation, MG-GPO makes use of the stochastic operations such as flocking [14], crossover, and mutation [13] to generate a large number of trial solutions, usually tens of times larger than the original population size. It then uses the GP model to filter out a fixed number of trial solutions with the highest potential, which are to be evaluated on the actual system. MG-GPO is capable of seeking the global optima due to its stochastic nature while achieving a high efficiency as it makes full use of the information contained in the measured data.

The MG-GPO was originally designed for design optimization. Small modifications are made to its implementation in the present study for online applications.

2.1. Algorithm architecture

A schematic of the MG-GPO algorithm is shown in figure 1. The algorithm includes two parts: the initializer and the evolution loop. The initializer generates the initial population by the latin hypercube sampling [23] method or by the simple random sampling. The generated population enter the evolution loop to spawn the trial solutions by performing flocking [14], simulated binary crossover (SBX) [24] and polynomial mutation (PLM) [25] operations. During this process, the population size could be multiplied by 10–40 times.

At the same time, the population of the previous generation are used to build a GP model, which is optimized to give the prediction values and uncertainties of the trial solutions. An acquisition function called lower confidence bound (LCB) is then applied to the trial solutions to assign each one a score per objective. The non-dominated sorting is employed to select the best scored trial solutions. After the selection, the population size is reduced back to the original, and the selected ones will be evaluated by the online system to form the next generation. The loop repeats until meeting the termination condition.

2.2. Evolutionary operations

Several stochastic evolutionary operations are built into the MG-GPO algorithm, which introduce randomness to the behavior of the algorithm and make the algorithm more robust to the noisy and potentially non-convex online optimization problems. The flocking operation is borrowed from the PSO

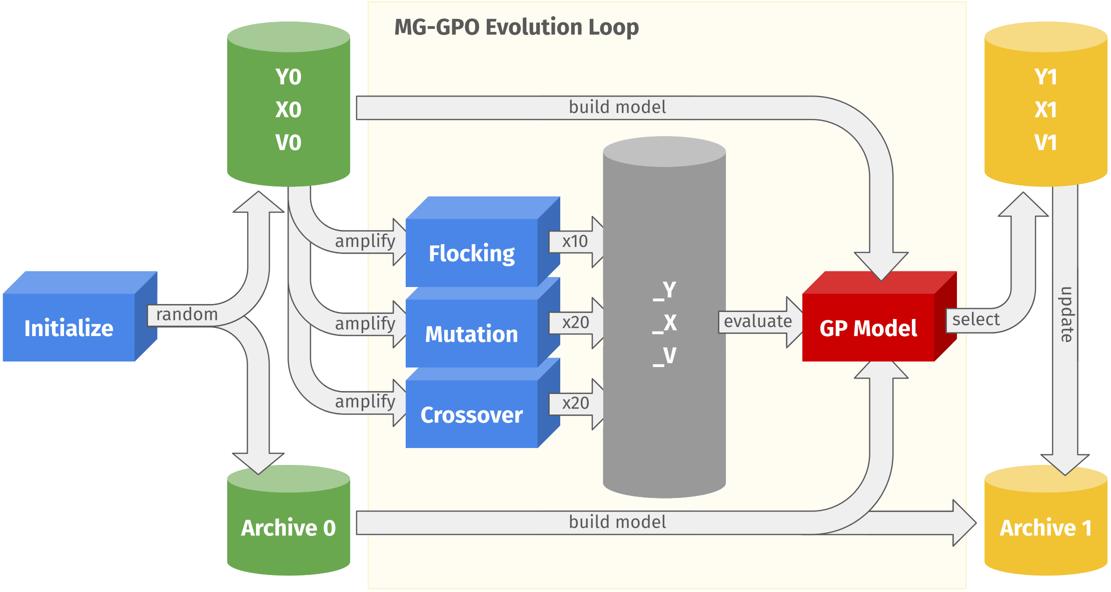

Figure 1. MG-GPO architecture. The green cylinders denote the previous generation, the yellow cylinders denote the next generation, the shaded area highlights the evolution loop, the blue/red cubes denote the operators, and the gray cylinders stand for the intermediate variables/states within one loop.

algorithm. Flocking in PSO has five hyperparameters in total: the three coefficients ω, c_1, c_2 are set to 0.4, 1 and 1 accordingly [17], and the amplitudes r_1, r_2 are random values between 0 and 3 in MG-GPO. The SBX and PLM operations that applied in NSGA-II are also imported in MG-GPO. The control parameter η in SBX and PLM is set to 20.

The stochastic operations are used to amplify the population as visualized in figure 1. A possible amplifying strategy is as follows: for an original population size of N_0 and the amplification factors for flocking, crossover, and mutation set to η_f, η_c, and η_m, respectively, the new solutions generated with the three types of operations will be $\eta_f N_0$, $\eta_c N_0$, and $\eta_m N_0$, respectively, which makes the total amplification factor $\eta_f + \eta_c + \eta_m$. The purpose of the amplifying process is to explore the search space more thoroughly. Comparing to the original population, the amplified population is far more diverse, while still inheriting the genes in the parent generation.

2.3. GP regression

GP regression is a probabilistic method that gives a prediction of function values with a confidence level from data samples. We construct the GP regression model with the history data that are collected during the optimization. In MG-GPO, the history data is composed of the evaluated solutions from the last few generations and the collection of the best solutions. For example, for a case with the population size $N_0 = 30$, data from the last five generations can be used. The number of data points used to build the GP model is limited by the GP computation complexity.

The key component of a GP is its kernel, which describes how closely two points in the input space are related. In this study, we choose the most widely used radial basis function (RBF) kernel:

$$k_{\mathrm{RBF}}(\mathbf{x},\mathbf{x}') = \sigma_f^2 \exp\left(-\frac{1}{2}(\mathbf{x}-\mathbf{x}')^T \boldsymbol{\Theta}^{-2}(\mathbf{x}-\mathbf{x}')\right). \tag{1}$$

The RBF kernel has two sets of hyperparameters: the variance σ_f^2 and the length scales $\Theta = \mathrm{diag}(\theta_1, \theta_2, \ldots, \theta_D)$, where D is the dimension of the input space. To simplify the hyperparameter tuning, we assume the same length scale for all the input variables per objective, so that the total tunable hyperparameters of the GP model have been reduced to 3 for a single objective problem: the length scale θ, the variance σ_f^2, and the Gaussian noise variance σ_n^2 which is included by default for a GP model. We employed the log marginal likelihood metric as the criteria to tune the hyperparameters. A python package called GPy [26] is used to auto-tune the hyperparameters, it applies special treatments for highly singular kernel matrices to make the GP modeling process more stable.

After a thorough hyperparameter optimization, the GP model could produce the prediction and the confidence of the given trial solution as shown below [27]:

$$\mu = \mathbf{k}_*^T(\mathbf{K}+\sigma_n^2\mathbf{I})^{-1}\mathbf{y}, \tag{2}$$

$$\sigma^2 = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^T(\mathbf{K}+\sigma_n^2\mathbf{I})^{-1}\mathbf{k}_*, \tag{3}$$

where $\mathbf{k}_*$ is the kernel vector for the trail solution, $\mathbf{K}$ the kernel matrix, $\mathbf{y}$ the objective values for the observed solutions, $k(\mathbf{x}_*, \mathbf{x}_*)$ the kernel value for the trail solution, $\mathbf{I}$ the identity matrix. The μ and σ^2 in the equations above are the predicted objective value and the corresponding variance, respectively, which are then transformed to a score per objective by the adaptive acquisition strategy.

2.4. Adaptive acquisition strategy

The acquisition function we choose is the LCB function [28] given by

$$\mathrm{LCB}(\mathbf{x}) = \mu(\mathbf{x}) - \kappa\sigma(\mathbf{x}), \tag{4}$$

where μ and σ are as in equations (2) and (3). The κ parameter in LCB is a balance factor, which determines the degree of the search space exploration and exploitation. A larger κ encourages the GP to explore further (try out the points that have larger uncertainty), and a smaller κ leads to a more conservative strategy.

Instead of using a constant κ strategy, in MG-GPO, κ is allowed to decay from generation to generation as shown in equation (5) (with $\rho < 1$) [21]. With this strategy, it emphasizes exploration at the beginning to avoid being trapped by local optima, and then gradually shifts toward exploitation in order to look for the best solutions in a relatively small area:

$$\kappa_n = \rho \cdot \kappa_{n-1}. \tag{5}$$

In this study, the initial balance factor is set to $\kappa_0 = 2$ and the decay rate $\rho = 0.9$ for the adaptive acquisition strategy to achieve the averagely best performance. This setup encourages strong exploration at the beginning, and, by attenuating kappa with generation, emphasizes exploitation in the late stage of the optimization.

3. Simulation

Simulation studies were conducted to test the performance of the MG-GPO algorithm in online optimizations. An analytic function, the Rosenbrock function with a dimension of 4, and two real-life accelerator problems were used in the tests. The first accelerator problem is to use skew quadrupoles to minimize the vertical emittance of a storage ring by compensating the linear coupling and vertical dispersion errors [20]. The second one is to use sextupole magnets to maximize the dynamic aperture of a storage ring [1]. In both cases, application on the SPEAR3 storage ring is assumed.

Most hyperparameters of MG-GPO are set to the default values as mentioned in section 2.2. There are a few tweaks to better adapt the algorithm to the test problems: we adjust the amplification factors η_f, η_c and η_m for each test problem. For the Rosenbrock problem and the vertical emittance problem, $(\eta_f, \eta_c, \eta_m) = (20, 10, 10)$, while for the dynamic aperture problem, $(\eta_f, \eta_c, \eta_m) = (2, 4, 4)$.

3.1. Minimization of Rosenbrock function

The Rosenbrock function with dimension N is defined as

$$f(\mathbf{x}) = \sum_{i=1}^{N-1}[100(x_{i+1} - x_i^2)^2 + (1 - x_i)^2]. \tag{6}$$

It is a non-convex function frequently used in testing optimization algorithms. We choose the case with $N = 4$ in this test, with the parameter range of x_i set to $[-2, 2]$ for $i = 1$–4. The global minimum is located at $\mathbf{x}_0 = (1, 1, 1, 1)$, with $f(\mathbf{x}_0) = 0$, while a local minimum is at $\mathbf{x}_1 = (-1, 1, 1, 1)$.

In the test, Gaussian noise with $\sigma = 0.01$ is added to the objective function. The PSO and MG-GPO algorithms were both run for ten times, with 1020 evaluations in each run. The comparison of the median, mean and standard deviation of the best solutions among 10 runs is shown in the left subplot of figure 2. Clearly, MG-GPO performs better than PSO in terms of the convergence speed and the final objective values. The MG-GPO converges to approximately zero with 600 evaluations, while PSO has not converged within 1020 evaluations.

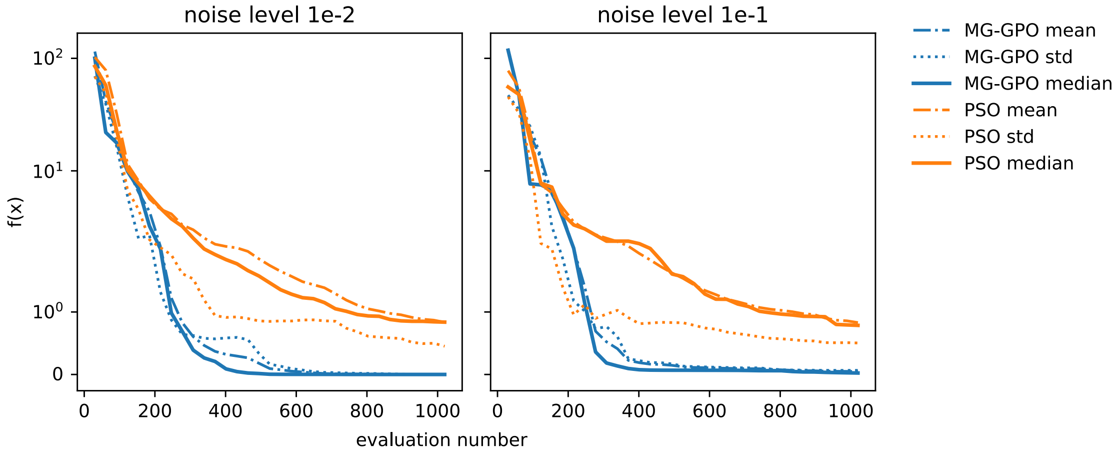

Figure 2. The history of the median, mean and standard deviation of the best solutions among 10 runs for the Rosenbrock four-dimensional minimization test for MG-GPO and PSO. Left plot: Gaussian noise $\sigma = 0.01$; right plot: $\sigma = 0.1$.

Tests with a larger Gaussian noise, $\sigma = 0.1$, have also been conducted. The comparison results for two different noise levels on the Rosenbrock test function are shown in figure 2 (right plot). The results indicate that MG-GPO is more noise-resilient than PSO for this test case. There is no clear performance decrease for either algorithms when the noise level is increased from 0.01 to 0.1.

3.2. Minimization of vertical emittance

In a storage ring, the vertical emittance is usually dominated by the residual vertical dispersion and linear coupling between the horizontal and vertical planes. Both of those can be compensated by skew quadrupoles magnets, which can be used as tuning knobs. Since the vertical projected emittance varies with location in the ring and the SPEAR3 storage ring has diagnostics to measure the beam size at only one location, in experiments Touschek beam lifetime is used as an indirect, global measure of the vertical emittance. The Touschek lifetime is proportional to the vertical beam size. For a high current beam in a third generation storage ring such as SPEAR3, Touschek scattering is the dominant beam loss mechanism. Therefore, minimizing the vertical emittance can be achieved by maximizing the beam loss rate for a high current beam [20].

In simulation, the vertical emittance can be calculated with the lattice model for the storage ring. It is then used to calculate the beam loss over a period of 6 s for a beam current of 500 mA in SPEAR3, which is converted to beam loss rate. Gaussian random noise are added to the beam currents at the beginning and the end of the 6 s period to simulate measurement errors. It introduces noise to the beam loss rate, i.e. the objective function. For beam current noise sigma of $\sigma_I = 0.003$ mA, the objective function noise sigma is about 0.04 mA min^{-1}.

The SPEAR3 has 13 powered skew quadrupoles that are used to correct linear coupling and vertical dispersion. The strengths of these magnets are the decision variables of the optimization problems. In the simulation, the strengths are specified in the normalized gradients of the magnets, each with a range $[-0.1, 0.1]$ m^{-2}. Initial coupling errors are introduced to the lattice model by changing other skew quadrupoles to reproduce the level of coupling in the machine when all 13 skew quadrupoles are turned off. The initial coupling ratio (i.e. the ratio of vertical emittance to horizontal emittance) is 1.08%.

MG-GPO and PSO are both applied for the optimization problem. Due to noise in the objective function and the stochastic nature of the algorithms, the convergence path is different every time. Each algorithm is tested ten times, with 1020 evaluations in each run. Figure 3 shows the comparison of the median, mean and standard deviation among the ten best solutions as they evolve during the 1020 evaluations. MG-GPO clearly outperforms PSO in the convergence speed and the best objective value achieved in 1020 evaluations. The performance variation within the 10 runs for MG-GPO is slightly larger than that of PSO in this test. The coupling ratio of the run for which the performance ranked the fifth among all 10 MG-GPO runs is 0.039%, while the corresponding value for PSO is 0.050%.

3.3. Optimization of dynamic aperture

A sufficiently large dynamic aperture is needed for injecting beam into a storage ring. A large dynamic aperture is preferred in order to achieve high injection efficiency. For a low emittance storage ring, strong sextupole magnets are used to correct chromaticities. The non-linear beam motion in the magnetic fields of

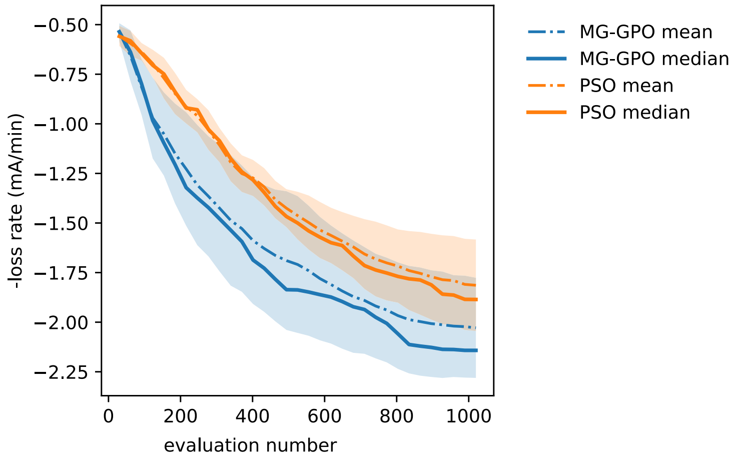

Figure 3. The history of the median and mean of the best solutions among 10 runs for the vertical emittance minimization test for MG-GPO and PSO. The shaded areas denote the corresponding one-sigma intervals.

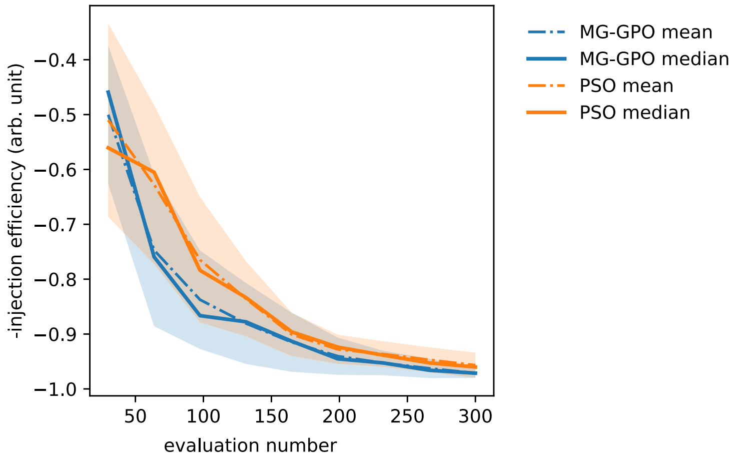

Figure 4. The history of the median and mean performance of 10 runs of injection efficiency during optimization in simulation for the MG-GPO and PSO algorithms. The shaded areas denote the corresponding one-sigma intervals.

the sextupoles limits the dynamic aperture. Magnetic field errors in the real machine can significantly decrease the dynamic aperture from the design performance. Optimizing dynamic aperture with sextupole magnets has been proven to be an effective remedy [1, 3, 4].

SPEAR3 storage ring has a total number of 72 sextupoles, which are powered in groups (i.e. families) according to the symmetry configuration of the ring sextupoles. There are ten sextupole families in total. Because it is necessary to keep the chromaticities fixed during optimization, we construct eight free knobs using the null space of the 2-by-10 chromaticity response matrix, as was done in reference [1]. The eight knobs are obtained with the singular value decomposition method. Each knob is a combination of the ten sextupole families that corresponds to a pattern with a zero singular values. The knobs do not change chromaticity to the linear order.

In simulation, the injection efficiency is evaluated by placing the injected beam represented by 1000 randomly generated particles according to the actual distribution into the horizontal phase space and counting the particles that fall between the dynamic aperture and the septum wall. Illustrations of injection efficiency calculation in the simulation can be found in figure 5. The dynamic aperture is obtained by particle tracking simulation, with particles launched from 19 rays on the upper x–y plane that extend from the origin. The average x-value of the last surviving particles on the first three rays on the negative side is used to define the phase space ellipse that represents the acceptance. The SPEAR3 upgrade lattice with an emittance of 6 nm is used. The separation between the stored beam and the injected beam is intentionally increased to make injection more challenging and hence create room for improvement for the dynamic aperture.

Figure 4 shows a comparison of the negated injected efficiency for the mean, median and standard deviation values of 10 runs for the MG-GPO and PSO algorithms. The MG-GPO algorithm clearly has faster convergence. The PSO algorithm eventually catches up as the objective function has an upper limit near

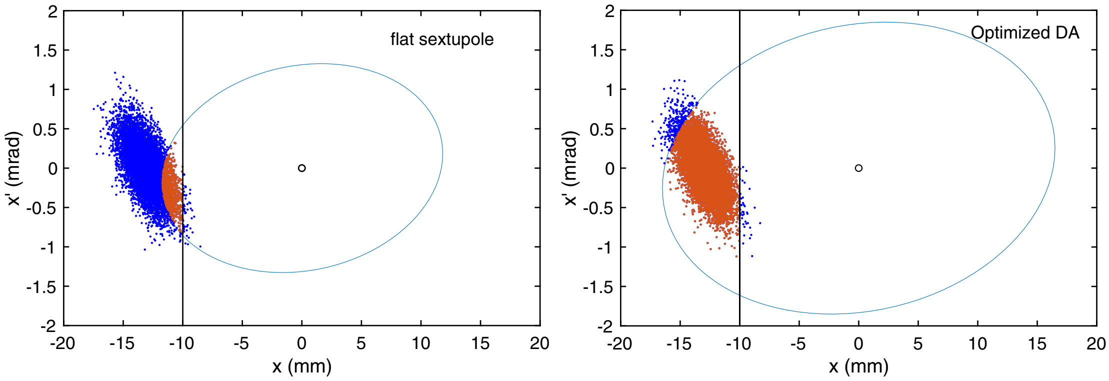

Figure 5. Injection diagram before (left plot) and after (right plot) DA optimization. The dark vertical line represents the septum wall. The orange dots represents the beam that will be captured. The dark circle indicates the location of the stored beam.

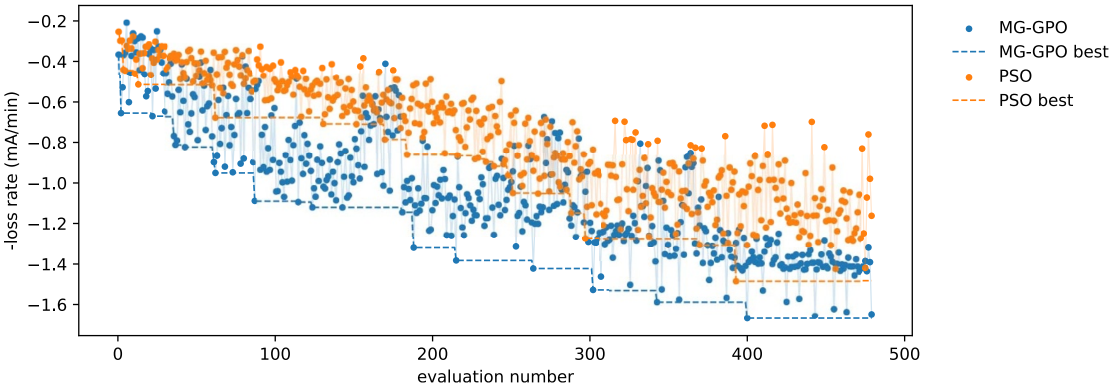

Figure 6. Comparison of the history of all evaluated solutions and the best-to-date solution during the beam loss rate optimization experiment with the MG-GPO and PSO algorithms.

(but below, as some particles are lost to the septum wall) 100% and, as the algorithms approach the limit, the convergence slows down. MG-GPO is also slightly more stable compared to PSO. Figure 5 plots show the injected beam in the horizontal phase space with the initial dynamic aperture or the optimized dynamic aperture, respectively. The initial dynamic aperture is for the sextupole setting with all SF and SD sextupoles in the standard cells set to equal values. The dynamic aperture is improved from 11.8 to 16.5 mm.

4. Experimental application of the MG-GPO algorithm

The MG-GPO method has been applied experimentally to the SPEAR3 storage ring to demonstrate its online optimization capability. In the following we show its application to the vertical emittance minimization problem and the dynamic aperture optimization problem. The problems are the same as was discussed in the previous section in simulation.

4.1. Vertical emittance minimization

In the experimental study of vertical emittance minimization, the strengths of the same 13 skew quadrupoles are used as tuning knobs. In this case, the strengths are changed by modifying the power supply set points of the magnets. During the experiment, the beam current is maintained at nearly 500 mA with frequent re-fills at a 5 min interval. Since the loss rate depends on the total current, the decrease of beam current in the 5 min period has an impact, which needs to be accounted for by normalizing the loss rate with the beam current squared [29]. The objective function is the beam loss rate normalized to a 500 mA total beam current.

Figure 6 shows the evolution of the negated loss rate in the optimization runs for MG-GPO and PSO under the same conditions. Within about 480 evaluations, the loss rate reached -1.66 mA min^{-1} for MG-GPO, nearly the maximum loss rate achieved in the SPEAR3 ring in recent studies. For PSO, the convergence speed is slower, and it only reached a loss rate of -1.48 mA min^{-1}.

While MG-GPO requires additional computation time for the algorithm itself, such an overhead is usually negligible in online optimization scenarios. For example, in the injection efficiency optimization

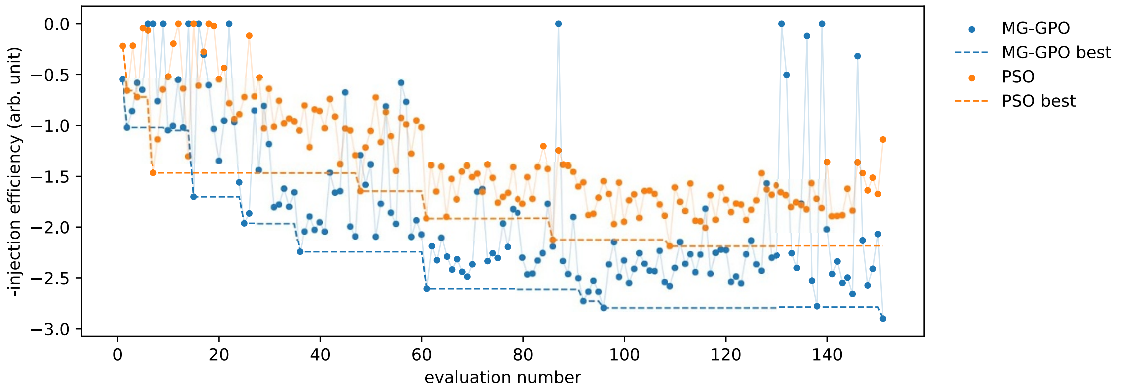

Figure 7. Comparison of the history of all evaluated solutions and the best-to-date solution during the injection efficiency optimization experiments with the MG-GPO and PSO algorithms, starting from the flat sextupole solution.

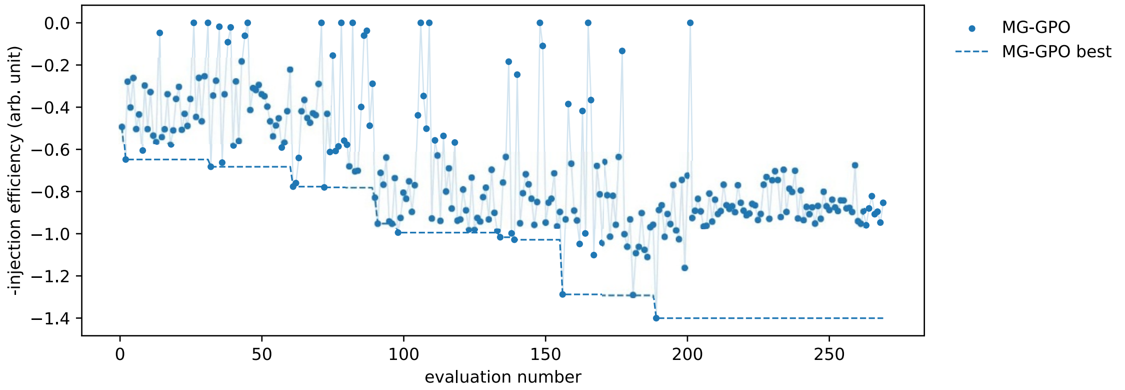

Figure 8. History of all evaluated solutions and the best-to-date solution during the injection efficiency second optimization experiment with the MG-GPO algorithm, starting from the best solution of the first run.

experiment, the time to measure one data point is around 17 s. Since the measurements are sequential, it takes about 8.5 min to evaluate for one generation (30 individuals). The MG-GPO calculation in one generation usually costs less than 8 s (with most of the time spent on GPy hyper-parameter tuning).

4.2. Dynamic aperture optimization

The optimization of dynamic aperture is done by optimizing the injection efficiency with sextupole knobs, in a manner similar to what was done in reference [1]. Injection into the 6 nm upgrade lattice for SPEAR3 was used in the experiments. The injection efficiency is measured by the ratio of beam current change in 10 s and the average beam intensity of the injected beam in the same period. A long duration is needed for the injection efficiency measurement to reduce the noise of the objective function. The power supply set points of the same ten sextupole families are varied through eight combined knobs that do not change the chromaticities as done in the simulated case. The range of the knobs are set such that the maximum variation of the sextupole current is within ±20 A from the initial setting when all SFs and SDs are equal, respectively (this setting is referred to as the flat sextupole solution).

The optimization run starts with the flat sextupole solution. A matched kicker bump with the standard bump size was used initially because the injection efficiency is poor for this solution. The first run took five generations, at which point a substantially improvement of injection efficiency was made. A second run was launched after the kicker bump size was reduced to 80% of the standard kicker bump and was centered around the best solution of the first run. The PSO algorithm was run starting from the same initial condition as the first MG-GPO run.

Figure 7 shows the history of the objective function values for the evaluated solutions and best-to-date solutions during the first run, in comparison to the PSO run. The injection efficiency as shown can be interpreted as the fill rate (in unit of mA min^{-1}) normalized to a certain injector beam intensity when the injector beam is delivered at a 2 Hz repetition rate. In the first run, the injection efficiency reached approximately -2.9 mA min^{-1} with 151 evaluation numbers for MG-GPO, while PSO only reached -2.2 mA min^{-1}. Figure 8 shows the results for the second MG-GPO run. The injection efficiency dropped to

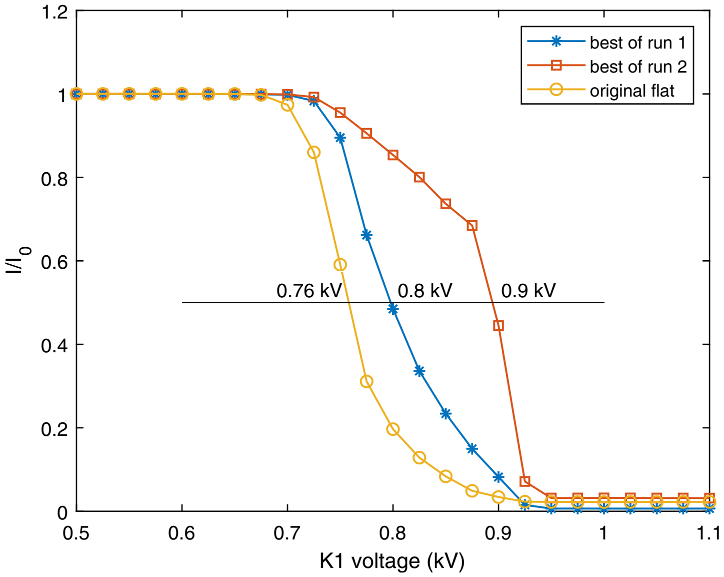

Figure 9. Comparison of measured dynamic aperture for the three solutions: the initial solution, or the best solution of the first and second runs.

a lower value initially because of the kicker bump decrease. It was then increased by the algorithm from -0.5 mA min^{-1} to about -1.4 mA min^{-1}.

Direct dynamic aperture measurement was done to characterize the performance of the optimized solutions. In the measurement a small stored beam current is kicked out with an increasingly larger kick strength. The beam currents (normalized by the initial value) versus kicker voltage are shown in figure 9 for three cases: the flat sextupole solution and the best solutions from the first run and the second run, respectively. The kick strength when the beam current drops to 50% may be used as a measure of the dynamic aperture. The corresponding kick voltage is increased from 0.76 kV (flat sextupole) to 0.9 kV (best solution).

5. Conclusions

In this study, we proposed to use the MG-GPO for online optimization of complex parameter spaces. The application of the method on accelerator tuning is demonstrated with two important problems, storage ring vertical emittance minimization with skew quadrupoles and dynamic aperture with sextupoles. Simulation and experiments both show that the algorithm can effectively improve the performance of the machine.

Acknowledgment

This work was supported by the U.S. Department of Energy, under Contract No. DE-AC02-76SF00515, and Office of Science, Office of Basic Energy Sciences, FWP 2018-SLAC-100469 and Office of Advanced Scientific Computing Research under FWP 2018-SLAC-100469ASCR.

Data availability statement

The data that support the findings of this study are available upon reasonable request from the authors.

ORCID iDs

Zhe Zhang https://orcid.org/0000-0002-8143-0381
Minghao Song https://orcid.org/0000-0002-7055-0660
Xiaobiao Huang https://orcid.org/0000-0002-8195-9277

References

[1] Huang X and Safranek J 2015 *Phys. Rev. ST Accel. Beams* **18** 084001

[2] Liuzzo S, Carmignani N, Farvacque L, Nash B P T, Raimondi P, Versteegen R and White S M 2016 RCDS optimizations for the ESRF storage ring *Proc. IPAC2016* (*Busan, Korea*) pp 3420–2

[3] Olsson D K 2018 Online optimisation of the MAX-IV 3 GeV ring dynamic aperture *Proc. IPAC2018* (*Vancouver, BC, Canada*) pp 2281–3

[4] Yang X, Ganetis G, Hidaka Y, Shaftan T, Smaluk V, Wang G, Yu L H and Zuhoski P 2019 Online optimization of Nsls-II dynamic aperture and injection transient *Proc. IPAC2019* (*Melbourne, Australia*)

[5] Duris J *et al* 2020 *Phys. Rev. Lett.* **124** 124801

[6] Martin I, Apollonio M and Bartolini R 2016 Online suppression of the sextupole resonance driving terms in the diamond storage ring *Proc. IPAC2016* (*Busan, Korea*) pp 3381–3

[7] Pulampong T, Klysubun P, Kongtawong S, Krainara S and Sudmuang S 2017 Online optimization applications at SPS *Proc. IPAC2017* (*Copenhagen, Denmark*) pp 4086–8

[8] Zhou G, Jiao Y, Xu G, Wu J and Zhang T 2016 Power improvement of free-electron laser using transverse-gradient undulator with external focusing *Proc. IPAC2016* (*Busan, Korea*) pp 761–2

[9] Wu J *et al* 2017 Recent on-line taper optimization on LCLS *Proc. FEL'2017* (*Santa Fe, NM, US*) pp 229–34

[10] Bergan W F, Bazarov I V, Duncan C J R, Liarte D B, Rubin D L and Sethna J P 2019 *Phys. Rev. Accel. Beams* **22** 054601

[11] Yang X *et al* 2019 *Sci. Rep.* **9** 5115

[12] Huang X 2018 *Phys. Rev. Accel. Beams* **21** 104601

[13] Deb K, Pratap A, Agarwal S and Meyarivan T 2002 *Trans. Evol. Comp* **6** 182–97

[14] Kennedy J and Eberhart R 1995 Particle swarm optimization *Proc. ICNN'95 - Int. Conf. on Neural Networks* vol 4 pp 1942–8

[15] Bazarov I V and Sinclair C K 2005 *Phys. Rev. ST Accel. Beams* **8** 034202

[16] Borland M, Sajaev V, Emery L and Xiao A 2009 Direct methods of optimization of storage ring dynamic and momentum aperture *Proc. PAC'09* pp 3850–2

[17] Huang X and Safranek J 2014 *Nucl. Instrum. Methods Phys. Res.* A **757** 48–53

[18] Tian K, Safranek J and Yan Y 2014 *Phys. Rev. ST Accel. Beams* **17** 020703

[19] Huang X 2016 Development and application of online optimization algorithms *Proc. NAPAC2016* (*Chicago, IL*) pp 1287–91

[20] Huang X, Corbett J, Safranek J and Wu J 2013 *Nucl. Instrum. Methods Phys. Res.* A **726** 77–83

[21] Huang X 2019 Multi-objective multi-generation Gaussian process optimizer for design optimization (arXiv: 1907.00250)

[22] Song M, Huang X, Spentzouris L and Zhang Z 2020 *Nucl. Instrum. Methods Phys. Res.* A **976** 164273

[23] McKay M D, Beckman R J and Conover W J 1979 *Technometrics* **21** 239–45

[24] Deb K and Agrawal R B 1995 *Complex Syst.* **9** 115–48

[25] Liagkouras K and Metaxiotis K 2013 An elitist polynomial mutation operator for improved performance of MOEAS in computer networks *2013 22nd Int. Conf. on Computer Communication and Networks (ICCCN)* (IEEE) pp 1–5

[26] GPy since 2012 GPy: a Gaussian process framework in python (available at: http://github.com/SheffieldML/GPy)

[27] Rasmussen C E and Williams C K I 2006 *Gaussian Processes for Machine Learning* (Cambridge, MA: The MIT Press)

[28] Auer P 2002 *J. Mach. Learn. Res.* **3** 397–422

[29] Huang X 2019 *Beam-Based Correction and Optimization for Accelerators* (Boca Raton, FL: CRC Press)