

Efficient dynamic and momentum aperture optimization for lattice design using multipoint Bayesian algorithm execution

Z. Zhang¹, I. Agapov², S. Gasiorowski¹, T. Hellert³, W. Neiswanger⁴,
X. Huang¹ and D. Ratner^{5,6,*}

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

²*Deutsches Elektronen-Synchrotron DESY, 22607, Hamburg, Germany*

³*Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*

⁴*University of Southern California, Los Angeles, California 90089, USA*

⁵*Thomas Jefferson National Accelerator Facility, Newport News, Virginia 23606, USA*

⁶*Old Dominion University, Norfolk, Virginia 23529, USA*



(Received 12 December 2025; accepted 1 September 2026; published 1 October 2026)

We demonstrate that multipoint Bayesian algorithm execution (multipointBAX) can overcome fundamental computational challenges in storage ring design optimization. Dynamic aperture and momentum aperture optimization is a multipoint, multiobjective design task for storage rings, ultimately informing the flux of x-ray sources and luminosity of colliders. Current production black-box optimization methods require extensive particle-tracking simulations for each trial configuration; the high computational cost restricts the extent of the search to $\sim 10^3$ configurations and therefore limits the quality of the final design. We remove this bottleneck using multipointBAX, which selects, simulates, and models each trial configuration at the single-particle level. We demonstrate our approach on a design for a fourth-generation light source, with neural-network-powered multipointBAX achieving equivalent Pareto front results using more than 2 orders of magnitude fewer tracking computations compared to genetic algorithms. The significant reduction in cost positions multipointBAX as a promising alternative to black-box optimization, and we anticipate multipointBAX will be instrumental in the design of future light sources, colliders, and large-scale scientific facilities.

DOI: [10.1103/4t6l-7jmg](https://doi.org/10.1103/4t6l-7jmg)

Designing modern scientific facilities—from synchrotron light sources to particle colliders—requires optimizing hundreds of parameters in complex, nonlinear systems, where a single design evaluation can take hours of computation. In storage rings, this challenge is exemplified by dynamic aperture (DA) and momentum aperture (MA) optimization, where maximizing the regions of particle stability directly determines injection efficiency, beam lifetime, and ultimately the photon flux or luminosity achievable in next-generation facilities.

The computational bottleneck is severe: maximizing DA and MA (DAMA) is a type of multipoint optimization, where evaluating a single lattice design requires tracking tens of thousands of particles for hundreds of thousands of turns, making global optimization prohibitively expensive. Moreover, there is a trade-off between maximizing DA and

MA area, so the standard approach is to find a Pareto front (PF); i.e., DAMA optimization is not just multipoint but also multiobjective. The *de facto* production methods, e.g., multiobjective genetic algorithms (MOGA) and multiobjective particle swarm optimization (MOPSO) [1–3], typically explore $\sim 10^3$ designs, forcing compromises between optimization scope and computational cost that result in suboptimal performance—a limitation that has been identified as a grand challenge for the field [4,5].

We present an approach combining multipoint Bayesian algorithm execution (multipointBAX) [6], an extension of Bayesian algorithm execution (BAX) to multipoint objectives, with deep neural network (NN) surrogates to achieve 2 orders of magnitude speedup over widely used methods. Our approach enables simultaneous optimization of dozens of lattice parameters while tracking particles in up to millions of potential configurations, which previously was computationally intractable.

Applied to the SSRL-X storage ring lattice, our method discovers Pareto-optimal solutions with $< 1\%$ of the function evaluations needed by the *de facto* production methods. This demonstration highlights how modern machine learning (ML) can overcome fundamental computational barriers in large-scale scientific facility design,

*Contact author: ratner@jlab.org

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

with immediate implications for current facility upgrades and the next generation of light sources and colliders.

Background and related work. In classical DAMA optimization, for each proposed design configuration, thousands of particles with different initial transverse positions (for DA) and momenta and longitudinal positions (for MA) are simulated for thousands to hundreds of thousands of turns around the storage ring. For DA, particles are typically chosen along rays of increasing distance from the origin, and the particle with maximum radius that survives the full number of turns defines the outer edge of the DA. The metric is the area enclosed by those points. At the end of the optimization, each potential solution is then reevaluated using a random collection of “error seeds,” i.e., random changes to the design configuration that approximate a realistic ring’s performance. A similar process is applied to the longitudinal positions and momentum offset for MA [7,8]. Ultimately, our goal is to identify a Pareto front of solutions in the two-dimensional DA and MA space.

The *de facto* production approaches for DAMA optimization, e.g., MOGA [2,9] and MOPSO [10], are slow for two reasons: first, they require many iterations to converge, and second, each iteration involves simulating thousands of particles to calculate the objectives. These are sometimes referred to as “outer” and “inner” loop problems, respectively [11]. ML methods have been explored as solutions to both challenges.

Outer-loop approaches reduce the number of iterations by learning a surrogate model of the objective. For example, Bayesian optimization (BO) uses a surrogate to construct an acquisition function to guide sampling. BO has been widely adopted for accelerator applications [12], including for DAMA optimization [13]. Similarly, Ref. [14] uses K-means clustering (rather than a Gaussian process and BO) to improve population fitness within MOGA. However, there is a ceiling on the efficiency gains of simply reducing the number of iterations—even a small number of iterations is costly due to the expense of each inner-loop acquisition.

Inner-loop approaches reduce the cost of each evaluation. One class of methods uses fast chaos indicators—such as frequency map analysis [15,16], forward-reversal integration [17], and auto-differentiation-based indicators [18]—to replace long-turn tracking with cheaper stability proxies that require far fewer turns per particle [19–21]. However, they remain proxies: like earlier metrics such as tune diffusion, they do not always accurately predict DA performance. Another class uses ML surrogates to replace calls to the physics simulation entirely. Because surrogates can be orders of magnitude faster than physics simulations, calls to the surrogate are effectively free. For example, replacing the physics simulator with a NN reduces computational expense, even when using an inefficient optimizer (e.g., MOGA). These surrogates can either

model the DAMA objective, i.e., directly predict the survival area of the DA and MA maps [22–25], or can model the full maps, i.e., predict the survival time for each individual particle simulation [26–29]. (Relatedly, NNs can also extend DAMA estimates to large numbers of turns beyond the limit of conventional tracking [27,28,30] and identify outliers [27].)

For NN approaches, the assembly of the training set often dominates the computational cost. Training on static, preexisting datasets risks that the surrogate will be inaccurate in the optimal region. Consequently, recent activity has focused on active learning, i.e., continually acquiring training data during the training and/or optimization processes [13,22,23,25,31], typically selecting an entire map at each iteration. Still, after amortizing the cost of the training dataset, the overall speedup compared to these production methods is typically modest.

This paper presents a path to DAMA optimization, using the outer loop to maximize efficiency of the inner-loop training dataset. The work involves the following innovations: (1) We select simulations and train the surrogate at the granularity of single particles, rather than using full maps. (2) We use an algorithmically driven acquisition function (i.e., multipointBAX), which selects individual points that maximally impact the optimization outcome, rather than just reducing model uncertainty. (3) We implement an NN-based BAX framework, avoiding the n^3 scaling of the Gaussian processes used by previous BO and BAX applications in accelerator physics. The NN approach is critical to handle the > 100 k points modeled in DAMA optimization.

Approach. Multipoint optimization refers to tasks in which each iteration of the optimizer requires a scan of variables in a secondary space. The term originated in aeronautic design (e.g., see Refs. [32–35]) but is common in scientific fields in general (e.g., see Refs. [36–38]), with multiple examples in recent years for accelerators [6,39,40]. In DAMA optimization, each acquisition in the configuration space requires a secondary scan of particles in position and momentum space. The expense of the secondary scan—in DAMA, each acquisition requires thousands of particle-tracking simulations to return a scalar objective—is the crux of why traditional multipoint optimization is difficult.

To address the multipoint challenge, multipointBAX trains a surrogate model to construct the full DA and MA stability maps, i.e., the surrogate model will predict the number of turns a particle will survive as a function of both configuration (magnet values) and position (x and y coordinates for DA, and s and δ coordinates for MA). Scanning over the position coordinates then constructs the maps, from which we can calculate the area metrics analytically. We can optimize directly on the cheap surrogate model, calculating the DA and MA area metrics for each configuration, to predict the Pareto front. This

procedure (DAMA calculation plus Pareto front finding) is the “algorithm” in multipointBAX [41]. Finally, multipointBAX uses the output of the surrogate model optimization to construct an acquisition function, which selects new points to acquire from the physics simulator. Specifically, the acquisition function selects new simulations near the predicted stability boundary, but only in regions predicted to have optimal performance. In this work, we use a single deterministic neural-network surrogate together with the greedy, mean-prediction-based MeanBAX acquisition function [42], a published instance of the BAX framework; fully probabilistic, information-based variants (e.g., InfoBAX) are a direction for future work (see Supplemental Material [43]).

We emphasize that multipointBAX never acquires a full DA or MA stability map; each acquisition returns the number of turns for a single particle, and the algorithm typically selects only a single-particle simulation for each proposed configuration. Crucially, the single particle is not used as a proxy for nonlinear performance; rather, each single-particle acquisition is the minimal data unit from which the NN surrogate incrementally learns the full stability map. It is the surrogate—not any individual particle—that serves as the model of nonlinear optics performance. Only when we have finished the full optimization will we evaluate the full DA and MA maps using the physics simulation. Figure 1 compares multipointBAX and BO, and Fig. 2 shows the full multipointBAX scheme.

To clarify the interaction between the two loops: in conventional methods, the outer loop proposes a configuration and the inner loop evaluates it by simulating a full particle map ($\sim 10^3$ particles). In multipointBAX, the outer and inner loops are unified through the surrogate model,




FIG. 1. Pedagogical comparison of BO and multipointBAX methodologies. There are two key differences: First, in multipointBAX, we simulate a single particle at each acquisition, whereas BO simulates a full map of more than 1000 particles. Second, the BO surrogate models the objective function (configuration $\rightarrow$ objective), whereas the BAX surrogate models the underlying physics function (i.e., configuration $\rightarrow$ map). Because of the increased complexity and large number of data points to model, BAX also requires the use of an NN model, whereas BO can use a simpler Gaussian process model.




FIG. 2. Schematic of the multipointBAX implementation for DAMA optimization. Particle tracking is performed with the Python build of the Accelerator Toolbox (pyAT) [46].

which spans both configuration space and particle-position space simultaneously. This enables the algorithm to jointly select which configuration *and* which particle position to simulate next—acquiring only the most informative single-particle data points rather than exhaustive maps. It is this joint selection that enables single-particle-level acquisitions and the resulting orders-of-magnitude efficiency gains.

Through the multipoint optimization lens, we can now reinterpret the innovations described previously: selecting individual points—and selecting only those that impact the optimization outcome—improves the sampling efficiency, while modeling individual particle outcomes improves information efficiency. Given the large number of simulations in each BO iteration, from an information and sampling perspective, multipointBAX can be orders of magnitude more efficient.

On the other hand, the need to model millions of points (rather than thousands of DA areas) makes it infeasible to use Gaussian processes as the underlying surrogate due to the poor n^3 scaling in the number of modeled data points. Consequently, we use an NN as the surrogate model. We have implemented a neural field [45] architecture: the network takes as input both the accelerator configuration and a particle position and predicts the survived turns for that one particle. Selecting a single particle at each iteration minimizes the total compute hours used, but given the availability of high-performance computing clusters, as well as the per-iteration overhead of retraining the NN and rerunning the genetic-algorithm search on the surrogate, to minimize the wall-clock time, we launch a batch of 50 physics simulations in parallel in each iteration.

Detailed descriptions of the algorithm, model architecture, additional results, and computational analysis are provided in Supplemental Material [43].

SSRL-X problem description. To demonstrate the method, we apply multipointBAX to the design of the proposed SSRL-X ring. SSRL-X is a low-emittance storage ring lattice that is an option for a future upgrade of the Stanford

Synchrotron Radiation Lightsource (SSRL). The design has a natural emittance of 86 pm at 4 GeV [47], based on the hybrid six-bend achromat (H6BA) lattice structure [48]. The lattice employs several different combinations of H6BA cells to produce long and short straight sections to host insertion devices, superbends, or injection components, so it lacks the multifold periodicity aimed for by many other low-emittance storage ring lattices. With the use of the H6BA lattice structure and by enforcing transparency conditions, the lattice achieves a high nonlinear-dynamics performance. However, numerical optimization of the dynamic aperture and momentum aperture is still necessary.

The MOPSO algorithm has been previously used to optimize the dynamic aperture and the Touschek lifetime with six sextupole families [47]. In Ref. [47], the strengths of the sextupole families are varied freely within the six-dimensional space such that optimization can change—and optimize—the chromaticities [49]. However, additional constraints are necessary to prevent a large positive horizontal chromaticity and a negative vertical chromaticity. In the current study, we use the chromaticity response matrix of the six sextupole families to create four knobs that do not change chromaticities, in the same manner as found in Ref. [49]. This approach eliminates the need to implement additional constraints on chromaticities. Additionally, SSRL-X uses an off-axis injection scheme from the $-x$ side of the horizontal plane. We account for this by applying a weight of 2.0 to the $-x$ side of the stable region in the DA objective calculation.

Evaluation of DA and MA in the previous study [47,48] is done by tracking thousands of particles launched with initial coordinates spanning the phase space (for DA) and from various locations with different energy errors (for MA) for thousands of turns. The optimization is computationally costly and time-consuming, and typically only the performance for one error seed (for linear optics and coupling errors) is evaluated for each sextupole solution. Using multipointBAX can shorten the optimization cycle, allowing more robust optimization across multiple error seeds, and search over larger parameter spaces.

Comparison of bax and nsga-ii performance. The full multipointBAX algorithm implemented for this study is shown schematically in Fig. 2. We highlight that multipointBAX focuses only on promising regions of parameter space: once the model has determined a region is not promising, multipointBAX does not resample those configurations, improving the efficiency of training data acquisition. On the other hand, multipointBAX continually resamples in regions with high DA and MA, increasing prediction confidence in promising regions (Fig. 3).

To assess performance against the *de facto* production baseline, we require both a benchmark algorithm and a metric. We choose the NSGA-II (a standard genetic algorithm with similar performance to the MOPSO used in Ref. [47]) as both the production baseline and the internal optimizer within multipointBAX. For a metric, we select the hypervolume enclosed by the Pareto front and a




FIG. 3. Model learning process: (a) shows the configuration input distribution projected on two dimensions (u_0 and u_1) using a uniform manifold approximation and projection [50] transform fit on initial samples of BAX. Gray points show all sampled configurations during the BAX run, blue triangles show the final PF from nondominated sorting genetic algorithm II (NSGA-II), red squares show the final PF from BAX, and green circles show the initial PF from BAX. (b),(c) show DA maps predicted by the NN, and (d),(e) show MA ones. The MA panels are plotted versus the index of the longitudinal position “spos” along the ring (horizontal) and the index of the momentum offset (vertical). Panels (b),(d) show a “good” configuration predicted to lie on the Pareto front [black dot labeled (b,d) in (a)], while panels (c),(e) show a random “bad” configuration [black dot labeled (c,e) in (a)] far from the predicted Pareto front. For both DA and MA, from left to right in each subplot: prediction after initial training (Initial DAMA), prediction after final training (Final DAMA), and ground truth (GT DAMA). The color scale in all panels (b)–(e) represents normalized survived turns (white = lost on first turn, blue = survived the full number of tracked turns), and the contour marks the DA or MA boundary.




FIG. 4. BAX vs NSGA-II performance as a function of the number of particle-tracking simulations using hypervolume. We run both BAX and NSGA-II twice and show both individual run curves (thin, light lines), as well as the mean of the two (thick lines). For BAX, solid lines show optimization cost only, and dashed lines include the final evaluation (20 configurations); for NSGA-II, only the optimization cost is shown, as the evaluation cost (~ 720 k simulations) is negligible relative to the ~ 22 M optimization simulations. Black stars mark BAX's early-stopping point on the optimization-only (25k simulations) and total-cost (265k simulations) curves, and the point at which NSGA-II reaches comparable performance (11M simulations); the dotted line indicates that hypervolume level.

reference point, with the comparison as a function of simulation queries shown in Fig. 4.

For both multipointBAX and NSGA-II, we randomly select 1 of 10 error seeds for each call to the physics simulator, which we observe improves the performance of both approaches. After the optimization concludes, we evaluate every point in each final selected configuration using 10 new error seeds, none of which were used during the optimization process. For BAX, the evaluation cost is significant relative to the small optimization cost, so Fig. 4 shows both optimization-only and total curves. For NSGA-II, only the optimization cost is shown, as the evaluation cost is negligible compared to the ~ 22 M optimization simulations.

We find that multipointBAX converges with 40 times fewer acquisitions, all-inclusive, compared to the number of acquisitions required by NSGA-II to reach approximately the same performance (Fig. 4). However, this undersells the efficiency of multipointBAX: only 10% of the simulations occur during the BAX routine, with the vast majority occurring in the final evaluation after the optimization is complete. If we only consider the BAX acquisitions, multipointBAX converges to an optimum with more than 400 times fewer simulation calls compared to NSGA-II. The high fraction of simulations in the final evaluation is in part due to the relative simplicity of the four-dimensional search. For harder tasks (e.g., higher input dimensionality, more complex behavior), the final

evaluation will be a smaller fraction of the computational cost, and the all-inclusive gain will increase.

We implement an early-stopping procedure based on convergence of the model's predictions [6], resulting in the black stars on the BAX curves in Fig. 4. We also note that Fig. 4 compares the computational cost based on the number of simulations acquired, ignoring the varying runtime of individual particles. Because multipointBAX preferentially samples near the stability boundary where particles survive longer, its average per-simulation cost is higher, but the difference is relatively modest (~ 2.6 s for multipointBAX vs ~ 1.9 s for NSGA-II) and does not affect the conclusions. See Supplemental Material [43] for details. In total, the BAX optimization required ~ 100 CPU core-hours (~ 1 k including final evaluation), compared to ~ 38 k core-hours for NSGA-II. We further assess the robustness of these results to the choice of DAMA objective weighting and to run-to-run variation in Supplemental Material [43].

Beyond the efficiency gains, multipointBAX has an additional advantage of greater interpretability compared to the production baseline methods. MOGA and MOPSO provide no information about configurations beyond those sampled during the optimization, and each new query requires an expensive physics simulation. Classical BO algorithms produce a surrogate model of the objective function only. By contrast, the multipointBAX routine produces an NN surrogate of the full DA and MA maps, enabling subsequent exploration of new sextupole settings by hand. The NN surrogate provides physicists with intuition about the optimization, for example, revealing how the predicted DAMA maps change as a function of sextupole knobs or suggesting how to promote the growth of DA in desirable areas.

Conclusion. We have presented an application of multipointBAX to DAMA optimization, highlighting the benefit of acquiring and modeling simulations at the individual particle level. We observe $40\text{--}90\times$ speedups for the proposed SSRL-X lattice while achieving equivalent performance. We anticipate gains could be even higher when applied to more complex optimization tasks, e.g., for the Future Circular Collider [51], where the optimization is expected to take longer, and the final evaluation will be a negligible fraction of compute cost: during the optimization procedure, we observe a speedup factor of more than 400 compared to NSGA-II.

The implications extend far beyond DAMA. Any multipoint optimization—from beam diagnostics and machine protection at accelerators, to scientific facilities broadly—can benefit from this approach. Several advances could amplify these gains further: probabilistic models to quantify uncertainty, improved batch sampling strategies, pretraining on existing simulation databases, and combining multipointBAX with fast chaos indicators [17–19] that

reduce the per-simulation cost. Most significantly, the ability to simultaneously optimize performance and robustness opens new design paradigms where reliability is built into the optimization process rather than verified *post hoc*.

Finally, the ability to model a physics function, rather than an objective function, provides multiple benefits that future work could explore. For example, the trained surrogate could support optimization with new objective functions without the need for additional simulations. The physical surrogate could even support joint optimization of entirely different design tasks simultaneously. For example, diagnostics and machine protection systems also require detailed tracking and could build off the same set of simulations and even the same surrogate model developed during a multipointBAX DAMA optimization. Joint optimization across an entire facility design would provide further improvements in efficiency.

Acknowledgments. This work was supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No. DE-AC02-76SF00515.

Data availability. The data and code that support the findings of this study will be made publicly available in a repository upon publication of this letter.

-
- [1] M. Borland, V. Sajaev, L. Emery, and A. Xiao, Direct methods of optimization of storage ring dynamic and momentum aperture, in *Proceedings of PAC09, Vancouver, BC, Canada* (JACoW Publishing, Geneva, Switzerland, 2009), pp. 3850–3852, paper TH6PFP062.
 - [2] L. Yang, D. Robin, F. Sannibale, C. Steier, and W. Wan, Global optimization of an accelerator lattice using multi-objective genetic algorithms, *Nucl. Instrum. Methods Phys. Res., Sect. A* **609**, 50 (2009).
 - [3] X. Huang and J. Safranek, Nonlinear dynamics optimization with particle swarm and genetic algorithms for SPEAR3 emittance upgrade, *Nucl. Instrum. Methods Phys. Res., Sect. A* **757**, 48 (2014).
 - [4] S. Nagaitsev, Z. Huang, J. Power, J.-L. Vay, P. Piot, L. Spentzouris, and J. Rosenzweig, Accelerator and beam physics research goals and opportunities, Fermi National Accelerator Laboratory, Tech. Rep. FERMILAB-FN-1123, 2021.
 - [5] S. Nagaitsev, V. Shiltsev, A. Valishev, T. Zolkin, J.-L. Vay, M. Bai, Y. Cai, M. J. Hogan, Z. Huang, J. Seeman *et al.*, Accelerator and beam physics: Grand challenges and research opportunities, [arXiv:2203.06816](https://arxiv.org/abs/2203.06816).
 - [6] S. A. Miskovich, W. Neiswanger, W. Colocho, C. Emma, J. Garrahan, T. Maxwell, C. Mayes, S. Ermon, A. Edelen, and D. Ratner, Multipoint-BAX: A new approach for efficiently tuning particle accelerator emittance via virtual objectives, *Mach. Learn.: Sci. Technol.* **5**, 015004 (2024).
 - [7] C. Steier, D. Robin, L. Nadolski, W. Decking, Y. Wu, and J. Laskar, Measuring and optimizing the momentum aperture in a particle accelerator, *Phys. Rev. E* **65**, 056506 (2002).
 - [8] J. Kim, X. Huang, K. Tian, M. Song, J. Safranek, and P. Raimondi, A hybrid multi-bend achromat lattice design for SSRL-X, in *Proceedings of IPAC2022* (JACoW Publishing, Bangkok, Thailand, 2022), pp. 207–209, paper MOPOST054.
 - [9] T. Murata and H. Ishibuchi, MOGA: Multi-objective genetic algorithms, in *Proceedings of 1995 IEEE International Conference on Evolutionary Computation* (IEEE, Perth, WA, Australia, 1995), Vol. 1, pp. 289–294.
 - [10] C. A. Coello Coello and M. S. Lechuga, MOPSO: A proposal for multiple objective particle swarm optimization, in *Proceedings of the 2002 Congress on Evolutionary Computation (CEC'02)* (IEEE, Honolulu, HI, USA, 2002), Vol. 2, pp. 1051–1056.
 - [11] R. Stevens, V. Taylor, J. Nichols, A. B. Maccabe, K. Yelick, and D. Brown, AI for science: Report on the department of energy (DOE) town halls on artificial intelligence (AI) for science, Argonne National Laboratory, Argonne, IL, Tech. Rep. ANL-20/17, 2020.
 - [12] R. Roussel, A. L. Edelen, T. Boltz, D. Kennedy, Z. Zhang, F. Ji, X. Huang, D. Ratner, A. Santamaria Garcia, C. Xu *et al.*, Bayesian optimization algorithms for accelerator physics, *Phys. Rev. Accel. Beams* **27**, 084801 (2024).
 - [13] X. Huang, Z. Zhang, and M. Song, Multi-objective multi-generation Gaussian process optimizer, in *Proceedings of IPAC2021* (JACoW Publishing, Campinas, Brazil, 2021), pp. 3383–3386, paper WEPAB304.
 - [14] Y. Li, W. Cheng, L. H. Yu, and R. Rainer, Genetic algorithm enhanced by machine learning in dynamic aperture optimization, *Phys. Rev. Accel. Beams* **21**, 054601 (2018).
 - [15] J. Laskar, Frequency analysis for multi-dimensional systems. Global dynamics and diffusion, *Physica D* **67**, 257 (1993).
 - [16] Y. Papaphilippou, Detecting chaos in particle accelerators through the frequency map analysis method, *Chaos* **24**, 024412 (2014).
 - [17] Y. Li, Y. Hao, K. Hwang, R. Rainer, A. He, and A. Liu, Fast dynamic aperture optimization with forward-reversal integration, *Nucl. Instrum. Methods Phys. Res., Sect. A* **988**, 164936 (2021).
 - [18] J. Qiang, J. Wan, A. Qiang, and Y. Hao, Fast chaos indicator from auto-differentiation for dynamic aperture optimization, *Nucl. Instrum. Methods Phys. Res., Sect. A* **1087**, 171427 (2026).
 - [19] K. Hwang, C. Mitchell, and R. Ryne, Rapidly converging chaos indicator for studying dynamic aperture in a storage ring with space charge, *Phys. Rev. Accel. Beams* **23**, 084601 (2020).
 - [20] C. E. Montanari, R. B. Appleby, A. Bazzani, A. Fornara, M. Giovannozzi, S. Redaelli, G. Sterbini, and G. Turchetti, Chaos indicators for nonlinear dynamics in circular particle accelerators, *Eur. Phys. J. Plus* **140**, 603 (2025).
 - [21] E. Andrés Sánchez, A. Flores, J. Hernández-Cobos, and A. Antillón, Multiple quasi-invariant surfaces for dynamic aperture optimization in low-emittance synchrotrons, *Commun. Nonlinear Sci. Numer. Simul.* **154**, 109533 (2026).

- [22] M. Kranjčević, B. Riemann, A. Adelmann, and A. Streun, Multiobjective optimization of the dynamic aperture using surrogate models based on artificial neural networks, *Phys. Rev. Accel. Beams* **24**, 014601 (2021).
- [23] J. Wan, P. Chu, and Y. Jiao, Neural network-based multi-objective optimization algorithm for nonlinear beam dynamics, *Phys. Rev. Accel. Beams* **23**, 081601 (2020).
- [24] Y. Lu, S. C. Leemann, C. Sun, M. P. Ehrlichman, T. Hellert, H. Nishimura, and M. Venturini, Enhancing the MOGA optimization process at ALS-U with machine learning, in *Proceedings of IPAC2021* (JACoW Publishing, Campinas, Brazil, 2021), pp. 387–390.
- [25] Y. Lu, S. C. Leemann, C. Sun, M. P. Ehrlichman, H. Nishimura, M. Venturini, and T. Hellert, Demonstration of machine learning-enhanced multi-objective optimization of ultrahigh-brightness lattices for 4th-generation synchrotron light sources, *Nucl. Instrum. Methods Phys. Res., Sect. A* **1050**, 168192 (2023).
- [26] M. Schenk, L. Coyle, M. Giovannozzi, A. Mereghetti, T. Pieloni, E. Krymova, and G. Obozinski, Modeling particle stability plots for accelerator optimization using adaptive sampling, in *Proceedings of IPAC2021* (JACoW Publishing, Campinas, Brazil, 2021), pp. 1923–1926.
- [27] F. F. Van der Veken, M. Giovannozzi, E. H. Maclean, C. E. Montanari, and G. Valentino, Using machine learning to improve dynamic aperture estimates, in *Proceedings of IPAC2021* (JACoW Publishing, Campinas, Brazil, 2021), pp. 134–137.
- [28] J. Wan and Y. Jiao, Machine learning enabled fast evaluation of dynamic aperture for storage ring accelerators, *New J. Phys.* **24**, 063030 (2022).
- [29] D. D. Croce, M. Giovannozzi, T. Pieloni, M. Seidel, and F. F. Van der Veken, Accelerating dynamic aperture evaluation using deep neural networks, *J. Phys. Conf. Ser.* **2687**, 062032 (2024).
- [30] M. Casanova, B. Dalena, L. Bonaventura, and M. Giovannozzi, Ensemble reservoir computing for dynamical systems: Prediction of phase-space stable region for hadron storage rings, *Eur. Phys. J. Plus* **138**, 559 (2023).
- [31] D. Di Croce, M. Giovannozzi, E. Krymova, T. Pieloni, S. Redaelli, M. Seidel, R. Tomás, and F. F. Van der Veken, Optimizing dynamic aperture studies with active learning, *J. Instrum.* **19**, P04004 (2024).
- [32] M. Nemec, D. W. Zingg, and T. H. Pulliam, Multipoint and multi-objective aerodynamic shape optimization, *AIAA J.* **42**, 1057 (2004).
- [33] G. K. W. Kenway and J. R. R. A. Martins, Multipoint high-fidelity aerostructural optimization of a transport aircraft configuration, *J. Aircr.* **51**, 144 (2014).
- [34] R. P. Liem, G. K. W. Kenway, and J. R. R. A. Martins, Multimission aircraft fuel-burn minimization via multipoint aerostructural optimization, *AIAA J.* **53**, 104 (2015).
- [35] X. Chai, X. Yu, and Y. Wang, Multipoint optimization on fuel efficiency in conceptual design of wide-body aircraft, *Chin. J. Aeronaut.* **31**, 99 (2018).
- [36] T. Yamashita, N. Sato, H. Kino, T. Miyake, K. Tsuda, and T. Oguchi, Crystal structure prediction accelerated by Bayesian optimization, *Phys. Rev. Mater.* **2**, 013803 (2018).
- [37] K. Terayama, H. Iwata, M. Araki, Y. Okuno, K. Tsuda, and B. Rost, Machine learning accelerates MD-based binding pose prediction between ligands and proteins, *Bioinformatics* **34**, 770 (2018).
- [38] S. A. Miskovich, F. Montes, G. P. A. Berg, J. Blackmon, K. A. Chipps, M. Couder, C. M. Deibel, K. Hermansen, A. A. Hood, R. Jain, T. Ruland, H. Schatz, M. S. Smith, P. Tsintari, and L. Wagner, Online Bayesian optimization for a recoil mass separator, *Phys. Rev. Accel. Beams* **25**, 044601 (2022).
- [39] R. Roussel, D. Kennedy, A. Edelen, A. Ody, and E. Wisniewski, Autonomous beam alignment through quadrupole triplets using Bayesian algorithm execution, in *Proceedings of the 32nd Linear Accelerator Conference (LINAC2024)*, Chicago, IL, USA (JACoW Publishing, Geneva, Switzerland, 2024), pp. 765–767, paper THPB066.
- [40] E. Matzoukas, C. Xu, E. Blomley, E. Bründermann, J. Gethmann, G. De Carne, and A.-S. Müller, Efficient accelerator operation with artificial intelligence based optimization methods, in *Proceedings of the 16th International Particle Accelerator Conference (IPAC'25)*, Taipei, Taiwan (JACoW Publishing, Geneva, Switzerland, 2025), pp. 2722–2725, paper THPM018.
- [41] W. Neiswanger, K. A. Wang, and S. Ermon, Bayesian algorithm execution: Estimating computable properties of black-box functions using mutual information, in *Proceedings of the 38th International Conference on Machine Learning, Proceedings of Machine Learning Research*, edited by M. Meila and T. Zhang (PMLR, 2021). Vol. 139, pp. 8005–8015.
- [42] S. R. Chitturi, A. Ramdas, Y. Wu, B. Rohr, S. Ermon, J. Dionne, F. H. da Jornada, M. Dunne, C. Tassone, W. Neiswanger, and D. Ratner, Targeted materials discovery using Bayesian algorithm execution, *npj Comput. Mater.* **10**, 156 (2024).
- [43] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/4t6l-7jmg> for the full multipointBAX algorithm and hyperparameters, the acquisition function, the neural-network architecture, a comparison on an alternative single-configuration metric, the early-stopping indicator, and the average time cost per particle-tracking simulation for BAX and NSGA-II, which includes Ref. [44].
- [44] D. Ginsbourger, R. Le Riche, and L. Carraro, Kriging is well-suited to parallelize optimization, *Computational Intelligence in Expensive Optimization Problems, Adaptation, Learning, and Optimization*, edited by Y. Tenne and C.-K. Goh (Springer, Berlin, 2010), Vol. 2, pp. 131–162.
- [45] Y. Xie, T. Takikawa, S. Saito, O. Litany, S. Yan, N. Khan, F. Tombari, J. Tompkin, V. Sitzmann, and S. Sridhar, Neural fields in visual computing and beyond, *Comput. Graphics Forum* **41**, 641 (2022).
- [46] W. Rogers, N. Carmignani, L. Farvacque, and B. Nash, pyAT: A python build of accelerator toolbox, in *Proceedings of IPAC2017* (JACoW Publishing, Copenhagen, Denmark, 2017), pp. 3855–3857, paper THPAB060.
- [47] P. Raimondi, X. Huang, J. Kim, J. Safranek, and T. Rabedeau, Advanced storage ring lattice options based on hybrid six-bend achromat for stanford synchrotron radiation lightsource upgrade, *Nucl. Instrum. Methods Phys. Res., Sect. A* **1061**, 169137 (2024).

-
- [48] P. Raimondi and S.M. Liuzzo, Toward a diffraction limited light source, [Phys. Rev. Accel. Beams](#) **26**, 021601 (2023).
- [49] X. Huang and J. Safranek, Online optimization of storage ring nonlinear beam dynamics, [Phys. Rev. ST Accel. Beams](#) **18**, 084001 (2015).
- [50] L. McInnes, J. Healy, and J. Melville, UMAP: Uniform manifold approximation and projection for dimension reduction, [arXiv:1802.03426](#).
- [51] M. Benedikt, A. Blondel, P. Janot, M. Mangano, and F. Zimmermann, Future circular colliders succeeding the LHC, [Nat. Phys.](#) **16**, 402 (2020).