

HOW CAN MACHINE LEARNING HELP FUTURE LIGHT SOURCES?

A. Santamaria Garcia*, C. Xu, L. Scomparin, E. Bründermann, M. Caselle, G. De Carne, A.-S. Müller, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany

Abstract

Machine learning (ML) is one of the key technologies that can considerably extend and advance the capabilities of particle accelerators and needs to be included in their future design. Future light sources aim to reach unprecedented beam brightness and radiation coherence, which require challenging beam sizes and accelerating gradients. The sensitive designs and complex operation modes that arise from such demands will impact the beam availability and flexibility for the users, and can render future accelerators inefficient. ML brings a paradigm shift that can re-define how accelerators are operated. In this contribution we introduce the vision of ML-driven facilities for future accelerators, address some challenges of future light sources, and show an example of how such methods can be used to control beam instabilities.

INTRODUCTION

Frontier Accelerators

Both the photon science and high energy physics research communities generally aim at increasing the performance of accelerators, reduce their cost, and make them more power efficient. These goals are even more relevant for frontier particle accelerators, driven by ambitious research programs that require demanding beam parameters, often outpacing the progress of traditional accelerator technologies [1–3]. The current cost of frontier accelerators is estimated at more than 1 billion dollars, where larger facilities can cost up to 10 times more [4]. This cost is directly related to the technology these accelerators are based on, and can be reduced with advancements in such technology. Given the size, cost, and technological advancements required, frontier accelerators are one of the most challenging scientific endeavors.

The Potential of Machine Learning for Particle Accelerators

The potential of ML methods in accelerators was already identified back in 2018 [5], and their popularity has been rapidly increasing since then, as shown in Fig. 1.

This is due not only to the general rise in popularity of artificial intelligence (AI), but thanks to the continuous development of easily accessible ML software libraries and recent advances of ML methods, applied to a variety of fields. While interesting improvements have already been observed in existing accelerator facilities, particularly in tuning and optimization tasks [6, 7], a promising avenue for ML methods in accelerators is their potential to help overcome the challenges of frontier accelerators [8], which could become a liability in their development if unaddressed (e.g., technical impossibilities, insufficient beam availability, inefficient

* andrea.santamaria@kit.edu




Figure 1: Number of publications with the words "machine learning" or "artificial intelligence" in the abstract, scraped from the JACoW database.

design). Some of the challenges that can be approached with ML are listed in Fig. 2, for different frontier accelerator design trends.




Figure 2: Some trends and related challenges of frontier accelerators.

ML methods can yield fast predictions at a reduced computational cost compared to analytical or classical numerical methods, can take into account the non-linear correlations in the data, and can adapt the predictions to the drifts in the machine state. These capabilities are highly desirable in accelerators since they open the door to a very robust and tailored online detection, prediction, optimization, and control. They can also help design future accelerators by alleviating the computational cost of numerical simulations and guiding the search for optimal parameters in a high-dimensional parameter space. Table 1 summarizes some applications of ML in current accelerators, split by the type of task. There are also numerous applications in particle physics that are not covered in this publication [9, 10].

Looking ahead into the future we can imagine a completely autonomous accelerator [11] where the operation is user-centric and guided by the changing demands of more complex and maybe ML-driven experiments. Such an accelerator would be energy responsible, with an automated

Table 1: Machine learning opportunities in particle accelerators (online and in simulation)

Task	Goal	Methods/Concepts	Examples ¹
Detection	Detect outliers and anomalies in accelerator signals for interlock prediction, data cleaning	- Anomaly detection - Time series forecasting - Clustering	- Collimator alignment - Optics corrections - SRF quench detection
Prediction	Predict the beam properties based on accelerator parameters	- Virtual diagnostics - Surrogate models - Active learning	- Beam energy prediction - Accelerator design - Phase space reconstruction
Optimization	Achieve desired beam properties or states by tuning accelerator parameters	- Numerical optimizers - Bayesian optimization - Genetic algorithm	- Injection efficiency - Radiation intensity
Control	Control the state of the beam in real time in a dynamically changing environment	- Reinforcement learning - Bayesian optimization - Extremum Seeking	- Trajectory steering - Instability control

¹ non-exhaustive

start-up and operation, failure and interlock prediction, virtual diagnostics, and intelligent control of beam dynamics phenomena (see Fig. 3).




Figure 3: A vision for future accelerators, driven by ML methods.

On the one hand, this automation would increase the beam availability for users due to faster commissioning times, faster set-up of operation modes, interlock prediction, less destructive measurements thanks to virtual diagnostics, and the reduction in the mean time between failures thanks to preventative maintenance. On the other hand, an online targeted phase space manipulation could not only deliver tailored beams and radiation to the users, but also actively mitigate beam instabilities. Paired with the increased speed in accelerator operations, new special operation modes would also be accessible. Swift changes between these operation modes would be possible, with direct feedback from the sample position at a beamline. Finally, the reduction of the

mean time between failures and energy responsibility efforts would reduce the costs of operating an accelerator, making it more reliable and sustainable. More details about energy responsible accelerators are given in the following section.

Energy Responsible Accelerators

From a grid perspective, particle accelerators are electrical loads in the order of hundreds of MWs with stringent requirements on power quality and low flexibility in power demand. The electrical grid is experiencing an increased number of disturbances (faster, more frequent, and more severe) caused by a decreased system inertia and increased variability in the power production (i.e., renewables). In addition, the energy cost steeply increased in many countries due to ongoing geopolitical conflicts. In view of these trends, future accelerators can profit from being more energy-efficient and resilient to external disturbances. Some steps that can be taken in that direction are:

- Inclusion of renewable energy sources combined optimally with energy storage systems to maximize the local energy production over the absorption from the main grid.
- Use of alternative cooling sources, such as geothermal energy, to increase the cooling efficiency and reduce the dependence from the electrical grid.
- Implement flexibility options for accelerator operation, to enable fast load variations in case of emergency: data centers, cooling systems, experiment scheduling based on green power production have great potential to improve flexibility.
- Develop novel devices and systems that are able to decrease the energy usage during accelerator operations: solid state amplifiers, permanent magnets, efficient cryogenic systems, superconducting power supply and distribution.

A joint test field was created at KIT to address the energy efficiency challenge, merging the Karlsruhe Research Accelerator (KARA) and the Energy Lab 2.0, Europe's largest research infrastructure for renewable energy. This joint-lab is called KITTEN (KIT Testfeld für Energieeffizienz und Netzstabilität in großen Forschungsinfrastrukturen) and aims to research energy sustainable and resilient solutions for accelerators to reduce their carbon footprint. In the BMBF project "ACcelerator Energy System Stability - ACCESS" the real-time power consumption data from the accelerator will be fed to a digital twin that can emulate with high fidelity the accelerator power and energy dynamics during power systems studies. This will allow to test energy storage technologies or power electronics solutions in varying testing conditions without affecting the real hardware. In this context, the fast inference of ML methods can be used to provide important real-time insights, identify patterns, analyze trends, and predict control actions to make the accelerators more energy sustainable, efficient, and stable.

MACHINE LEARNING FOR LIGHT SOURCES

Storage Rings

The common desiderata for photon beams produced in 4th generation light sources (4GLS) are high brilliance and flux, coherent radiation, and tunability of wavelength, beam size, polarization, and time structure. Additionally, the beam has to be as stable in energy, intensity, position, and size as possible.

Operating with ultra-low emittances comes with various challenges [12], like for example a reduced lifetime due to an increased Touschek scattering effect. The Touschek scattering can be decreased using round beams, but flat beams have a higher coherence and brightness and are therefore preferred by the users. A multi-objective optimization algorithm can find a compromise between both, optimizing for both lifetime and brightness.

A stable beam size (source size) is also desired, as some experiments are highly sensitive to intensity fluctuations. Insertion device (ID) gap variations induce coherent tune shifts that translate into orbit distortions, caused by their integral field errors. This is usually compensated with local and/or global orbit and linear optics corrections schemes, but non-linear residual field errors persist and can significantly affect the quality of the beam. For example, residual skew quadrupole errors induce vertical beam size variations from undesired coupling, and residual higher-order components reduce the beam lifetime when the tune is near their resonance [13]. In both cases, the quadrupole strengths (skew or straight) can be used to change the optics and correct the betatron coupling or move the betatron tune away from resonances, respectively. ML methods have already been applied to this problem, where for example neural networks (NNs) were trained to predict the orbit distortion at 239 beam position monitors (BPMs) induced by 18 different ID gaps [14] and extremum seeking (ES) was used to tune skew

quadrupoles to minimize the vertical emittance [15]. A more detailed study can be found in Ref. [16], where beam size predictions were done with NNs and used as a feed-forward to stabilize the beam size. It was found that the levels of stability achieved were roughly one order of magnitude better than previously observed using model-based schemes, fulfilling the requirements for future light sources.

Storage rings can provide photons at MHz repetition rates but at a relatively low power. The intensity of the photons can be amplified by increasing the spatial coherence between the emitted waves, in which case the intensity will scale quadratically with the number of electrons at full coherence instead of linearly. This can be achieved by reducing the bunch length to the scale of the emitted wavelength or with the presence of substructures in a longer bunch, like the ones created by the microbunching instability (MBI). The MBI is a longitudinal collective instability that happens above a certain current threshold and is driven by the self-interaction of the bunch with its own wakefield. It results in the formation of substructures that emit bursts of coherent synchrotron radiation (CSR) [17]. These bursts happen at a particular frequency corresponding to the rate of growing and damping of the charge substructures, a periodic phenomenon driven by the non-equilibrium between the driving wake potential and radiation damping, diffusion, and filamentation mechanisms. This is observed as a partially periodic fluctuation of the CSR amplitude in time, and a growing and shrinking effect in longitudinal phase space. This bursting makes the CSR power considerably fluctuate on timescales that are difficult for users to average. Radio-frequency (RF) modulations can be used to stabilize the CSR power [18], but requires a more intelligent control than a simple feedback due to the continuously-evolving charge densities in longitudinal phase space. Studies have been carried out towards the control of the MBI with reinforcement learning (RL), and a summary is presented in the last section of this paper. Achieving control over the longitudinal phase space is particularly relevant for 4GLS, as they are mainly designed for beam brilliance optimization and lack pulse flexibility and variability for spectroscopy and timing experiments. Unprecedented control of pulse length and pulse repetition rate can make very flexible and tailored modes of operation possible. Additionally, such an intelligent feedback system can open the door to the control of other types of instabilities that deteriorate the beam quality and limit the bunch current range for stable operation, which is especially relevant in 4GLS where instabilities are more significant.

Reaching ultra-low emittances require stronger sextupole strengths for chromatic correction, which in turn reduce the dynamic aperture and therefore the transverse acceptance. For accelerators targeting two orders of magnitude reduction in emittance the injection needs to be done on-axis to avoid injection oscillations. Off-axis injection into storage rings with methods such as Bayesian optimization (BO) have already been successfully implemented in light sources [19], and similar or more advanced methods could assist novel injection schemes.

Special modes that are challenging to operate in existing machines can also be assisted by ML, like for example negative momentum compaction factor α_c operation. It is designed to increase the dynamic aperture by reducing the strength of the sextupoles and requires operation at negative chromaticities and α_c to avoid head-tail instabilities [20]. While the lifetime of this operation mode is lower than for positive α_c , it has a considerable bunch shortening effect that could be useful for particular experiments [21].

ML methods can certainly speed-up commissioning campaigns in any type of accelerator [19, 22], and more effort could be invested in developing an all-encompassing commissioning tool that leverages the advantages of ML.

Finally, ML can help in simulation by greatly reducing accelerator design stages and developing new optics. 4GLS have strongly nonlinear lattices where the optimization of the dynamic aperture and momentum acceptance is complex due to the dimensionality of the problem, the sensitive correlation among those parameters, and stringent constraints. Deep learning techniques have been developed to accelerate lattice evaluation for 4GLS [23], which allows a faster convergence to an optimal design. Accelerator design can be further sped-up by active learning, where the surrogate model is built iteratively with simulation points selected by a model that guides the parameter space exploration based on the uncertainty of the model (e.g., [24]).

Linear Accelerators

The existing and future free electron lasers (FELs) aim to reach a higher repetition rate and provide unprecedented precise control over the light pulses, such as sub-femtosecond pulses, higher peak power, flexible spectrum, and tunable polarization [25, 26]. As opposed to a storage ring, which can serve dozens of beamlines simultaneously, an FEL only have a handful of beamlines and requires special modes like multi-beam operation with a switchyard to provide the light to the beamlines simultaneously. Therefore, it is common that an FEL needs to switch operation modes multiple times a day to provide the light tailored to the requirement of user experiments. This necessitates methods and routines that can automate and speed up the process of setting up the accelerator for different operation modes. Various methods have been designed and applied to aid such tuning tasks, including the Nelder-Mead simplex algorithm, ES [27, 28], and robust conjugate direction search (RCDS) [29, 30]. An alternative approach, BO, is able to perform global optimization efficiently. It has been successfully tested at multiple FEL facilities for tasks like optics matching to SASE pulse energy optimization [31–33]. The BO algorithm can also be modified to incorporate correlations of the tuning parameters, safety constraints, and drifting systems, making it applicable to a wide range of daily accelerator tasks. Several software frameworks are developed collaboratively to provide a standardized implementation of the advanced optimization algorithms mentioned above and aid general accelerator tuning tasks [34–36]. Such frameworks simplify the sharing of models and algorithms across different facilities,

helping the transition of state-of-the-art ML methods from research projects to operational tools.

Additionally, since FELs are single-pass accelerators and one bunch only radiates at a single beamline, they possess more freedom to modulate and tailor the pulses to the user requirements [37, 38]. With more ML-based tuning tools available for the operation of FLS, it can be expected that the machine operation modes will be dynamically changed, providing pulses according to the user's needs throughout the experiment.

In recent studies, RL proved to be able to solve various beam tuning tasks at simpler setups [39–41], outperforming existing numerical optimization methods. Once trained, the RL policy can also handle system drifts and be used as a continuous controller. With enough redundancy in the system, it can even deal with unexpected scenarios such as magnet power supply failures [42]. For future accelerators, it is foreseeable that RL methods will be deployed as robust controllers for complicated system dynamics, which will be otherwise challenging or not possible using conventional feedback controls.

Another challenge that future FELs face is diagnostic devices. Especially for the case of longitudinal phase space (LPS), existing diagnostics like transverse deflecting cavities are often destructive, lack resolution for ultra-short bunches, and need multi-shot measurements to reconstruct the full phase space information. ML methods, such as NNs, can be trained to provide rapid non-destructive predictions of the LPS [43, 44], which can have higher resolutions compared to the single shot measurements [45, 46]. The virtual diagnostics driven by the ML models can be used in combination with real diagnostics to provide higher-fidelity information on the electron bunch [47]. Future applications of ML-based virtual diagnostics are expected to drastically increase the information that can be obtained during operation to the full 6D phase space of the electron bunch and simultaneous prediction at various points along the accelerator.

The current X-ray free electron lasers (XFELs) are predominantly driven by RF-based accelerating structures, which places them among the largest facilities worldwide. There is active research to construct more compact XFELs by accompanying advanced accelerating schemes. The most promising one is the plasma-based accelerator (PBA), including the laser wakefield accelerator (LWFA) and plasma wakefield accelerator (PWFA) [48, 49]. The FELs impose stringent constraints on the upstream beam parameters like the energy spread and divergence, which remains an issue of PBAs. The acceleration process in plasma depends nonlinearly on a large number of parameters, which are evolving due to the dynamic nature of the plasma. An analytic solution is often not possible or not accurate enough, so online tuning of the input parameters is constantly required in operation. Methods such as genetic algorithms and BO have been successfully applied at LWFAs to improve various aspects of the electron bunch, increasing the bunch charge and the bunch energy, minimizing the energy spread, and reaching a more stable state [50–53]. In addition to online tuning,

Bayesian models and active learning can be used to build an accurate data-driven model using either the particle-in-cell (PIC) simulations or the experimental data [54–56]. This allows to identify the underlying correlation between the parameters, extract new knowledge from the physics process, and obtain an optimized set of design parameters trading off different objectives for future operation.

TOWARDS THE CONTROL OF THE MICROBUNCHING INSTABILITY AT KIT

The real-time control of the MBI with RL is actively being researched at the storage ring KARA. Due to the dynamically changing nature of the instability, classical feedback schemes are not able to provide the required level of control. RL is a powerful learning paradigm that is particularly well-suited to tackle control problems in large environments, can learn from experience without the need of a model of the dynamics, and can deal with delayed consequences. RL applications are very promising, but their deployment in accelerators is challenging and has been done only a handful of times. One of the difficulties of training RL agents is that they need numerous interactions with the environment to train, and this is too time-consuming to be done in low repetition rate accelerators. This is why the RL agents are usually pre-trained in simulation, although the transfer to the real accelerator can be problematic when the gap between the simulation and experiment is too large. In our case, the data is generated in the accelerator at a faster rate than in simulation, overcoming one of the main limitations of this method. The goal of this project is to control the radiation emitted by the MBI with RF modulations in order to stabilize and maximize the radiation power. This will be achieved with a control feedback loop, composed of the following elements:

- **CSR detection:** broadband Schottky diode.
- **Pulse digitization:** KAPTURE-2 board, a low-latency and high-throughput sampling system for continuous sampling of ultra-short pulses developed at KIT [57].
- **Data readout:** HighFlex 2, a custom modular readout card (Xilinx ZYNQ family).
- **Low-latency RL inference platform:** Xilinx Versal VCK190 evaluation board, where the KIT-developed KINGFISHER [58] platform allows to more easily train agents on the accelerator.
- **Feedback system:** low-level RF (LLRF) amplitude and phase modulation control, possible every six revolutions.

The KAPTURE-2 board has 8 parallel sampling channels with a sampling rate of 500 MS/s, designed for bunch-by-bunch diagnostics at synchrotron light sources. This high data throughput is handled by the Highflex 2 board, where the analog-to-digital converters (ADCs) samples are labeled

with the bunch number information and optional metadata. These data are then sent to the Versal board through a high-speed fiber optic link, where the RL algorithm is implemented. This board combines an FPGA, an ARM processor, and programmable AI-engines, which are interconnected by a high-bandwidth Network-on-a-Chip (NoC) and allow full customization of the data flow. In this feedback system the more computationally heavy operations, namely feature extraction and agent inference, are carried out by the AI Engines, while the ARM processor runs the slow-control and the training algorithms as Petalinux applications. The system is designed as an experience accumulator, where the interaction of the RL agent with the accelerator is stored in the DDR memory. This data can then be used to train the model on several different platforms, ranging from the conventional CPU of a control room computer to a GPU equipped server, or to the ARM processor on a Versal board for cases with a particularly simple model, which would reduce the time needed for data access.

In order to influence the longitudinal beam dynamics the control feedback loop needs to act within a few synchrotron periods, imposing a latency constraint of tens of microseconds. This system was tested at KARA, showing a latency of 2.5 μ s for inference [59].

Results

The feasibility of the MBI control with RL was first tested in simulation, with encouraging results [60].

Due to needed hardware and firmware modifications to the LLRF to accept the continuous signal generated by the RL agent, the testing of the feedback loop for MBI control was postponed. However, in order to already test the concept, the feedback loop was adapted to sample a BPM signal and the RL agent was designed to damp transverse oscillations with a stripline kicker, replacing the bunch-by-bunch feedback system. The RL agent was deployed in the Versal board and tested during beamtime without being previously trained. The agent learned purely through interaction with the machine in several episodes of 2048 turns, and was re-trained after every episode on the control room computers. The RL agent performed equally or better than the conventional feedback system, validating the design of the control feedback loop. The results will be summarized for publication.

SUMMARY AND OUTLOOK

ML methods are powerful tools that can improve the performance of existing accelerators and create a new generation of autonomous ones, helping future accelerators become viable and sustainable and changing the way they are operated. As shown in this contribution, ML is already being used by the accelerator community to solve a variety of problems, where an untapped wealth of applications remains to be discovered. More advanced concepts like explainability, robustness, safety, and uncertainty quantification will need to be considered, as well as hardware infrastructure upgrades, to make ML methods become regular tools in accelerators.

ACKNOWLEDGEMENTS

A. Santamaria Garcia acknowledges funding by the BMBF ErUM-Pro project TiMo (FKZ 05K19VKC). G. De Carne research has been supported by the BMBF "Verbundprojekt 05K2022 - ACCESS: ACcelerator Energy System Stability".

REFERENCES

- [1] S. Gourlay *et al.*, *Snowmass'21 accelerator frontier report*, 2022. doi:10.48550/arXiv.2209.14136
- [2] CERN Yellow Reports: Monographs, *Cern yellow reports: Monographs*, vol. 1 (2022): *European strategy for particle physics - accelerator R&D roadmap*, en, 2022. doi:10.23731/CYRM-2022-001
- [3] R. W. Assmann *et al.*, *Eupraxia conceptual design report*, 2020. doi:10.1140/epjst/e2020-000127-8
- [4] V. Shiltsev, "Particle beams behind physics discoveries," *Physics Today*, vol. 73, no. 4, 2020. doi:10.1063/pt.3.4452
- [5] A. Edelen *et al.*, *Opportunities in machine learning for particle accelerators*, 2018. doi:10.48550/arXiv.1811.03172
- [6] X. Huang, *Beam-based Correction and Optimization for Accelerators*. doi:10.1201/9780429434358
- [7] Z. Geng and S. Simrock, *Intelligent Beam Control in Accelerators*. Springer Nature, 2023.
- [8] C. Allaire *et al.*, "Artificial Intelligence for the Electron Ion Collider (AI4EIC)," in *Artificial Intelligence for the Electron Ion Collider*, 2023. doi:10.48550/arXiv.2307.08593
- [9] A. Radovic, M. Williams, D. Rousseau, *et al.*, "Machine learning at the energy and intensity frontiers of particle physics," *Nature*, vol. 560, 2018. doi:10.1038/s41586-018-0361-2
- [10] M. D. Schwartz, "Modern Machine Learning and Particle Physics," *Harvard Data Science Review*, vol. 3, no. 2, 2021. doi:10.1162/99608f92.beeb1183
- [11] A. Eichler *et al.*, "First Steps Toward an Autonomous Accelerator, a Common Project Between DESY and KIT," in *Proc. IPAC'21*, Campinas, Brazil. doi:10.18429/JACoW-IPAC2021-TUPAB298
- [12] R. Walker *et al.*, "Summary of Working Group B: Ring-Based Light Sources," in *Proc. FLS'18*, Shanghai, China. doi:10.18429/JACoW-FLS2018-FRA1PL02
- [13] A. Papash *et al.*, "High order magnetic field components and non-linear optics at the ANKA storage ring," in *Proc. IPAC'17*, Copenhagen, Denmark, 2017. doi:10.18429/JACoW-IPAC2017-WEPAB011
- [14] B. Veglia *et al.*, "Neural networks for ID gap orbit distortion compensation in PETRA III," presented at IPAC'23, Venice, Italy, May 2023, paper WEP095, to appear in the proceedings.
- [15] B. Veglia *et al.*, "Extremum seeking for accelerator optimisation," presented at IPAC'23, Venice, Italy, May 2023, paper MOPA137, to appear in the proceedings.
- [16] S. C. Leemann *et al.*, "Demonstration of machine learning-based model-independent stabilization of source properties in synchrotron light sources," *Phys. Rev. Lett.*, vol. 123, p. 194801, 19 2019. doi:10.1103/PhysRevLett.123.194801
- [17] G. Stupakov and S. Heifets, "Beam instability and microbunching due to coherent synchrotron radiation," *Phys. Rev. ST Accel. Beams*, vol. 5, p. 054402, 5 2002. doi:10.1103/PhysRevSTAB.5.054402
- [18] A. Santamaria Garcia *et al.*, "Systematic study of longitudinal excitations to influence the microbunching instability at KARA," presented at IPAC'23, Venice, Italy, May 2023, paper WEP018, to appear in the proceedings.
- [19] C. Xu, T. Boltz, A. Mochihashi, A. Santamaria Garcia, M. Schuh, and A.-S. Müller, "Bayesian optimization of the beam injection process into a storage ring," *Phys. Rev. Accel. Beams*, vol. 26, no. 3, p. 034601, 2023. doi:10.1103/PhysRevAccelBeams.26.034601
- [20] P. Schreiber *et al.*, "Beam dynamics observations at negative momentum compaction factors at KARA," in *Proc. IPAC'22*, Bangkok, Thailand, 2017. doi:10.18429/JACoW-IPAC2022-THPOPT006
- [21] P. Schreiber *et al.*, "Effect of negative momentum compaction operation on the current-dependent bunch length," in *Proc. IPAC'21*, Campinas, Brazil. doi:10.18429/JACoW-IPAC2021-WEPAB083
- [22] J. Li *et al.*, "Particle swarm optimization algorithm applied in online commissioning at the MLS and BESSY II," in *Proc. IPAC'17*, Copenhagen, Denmark, 2017. doi:10.18429/JACoW-IPAC2017-THPAB008
- [23] Y. Lu *et al.*, "Demonstration of machine learning-enhanced multi-objective optimization of ultrahigh-brightness lattices for 4th-generation synchrotron light sources," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 1050, p. 168192, 2023. doi:https://doi.org/10.1016/j.nima.2023.168192
- [24] A. Oeftiger, A. Santamaria Garcia, S. Hirländer, and J. B. Lagrange, "Active Deep Learning for Nonlinear Optics Design of a Vertical FFA Accelerator," presented at IPAC'23, Venice, Italy, May 2023, paper MOPA137, to appear in the proceedings.
- [25] T. Raubenheimer, W. Decking, L. Giannessi, and H.-S. Kang, "Summary Report: Linac-Based Light Sources," in *Proc. 60th ICFA Advanced Beam Dynamics Workshop (FLS'18)*, Shanghai, China, 5-9 March 2018, Shanghai, China, 2018, pp. 199-201. doi:10.18429/JACoW-FLS2018-FRA1PL01
- [26] N. Huang, H. Deng, B. Liu, D. Wang, and Z. Zhao, "Features and futures of X-ray free-electron lasers," *The Innovation*, vol. 2, no. 2, p. 100097, 2021. doi:https://doi.org/10.1016/j.xinn.2021.100097
- [27] A. Scheinker *et al.*, "Model-independent tuning for maximizing free electron laser pulse energy," *Physical Review Accelerators and Beams*, vol. 22, 8 2019. doi:10.1103/PhysRevAccelBeams.22.082802

- [28] A. Scheinker, A. Edelen, D. Bohler, C. Emma, and A. Lutman, "Demonstration of model-independent control of the longitudinal phase space of electron beams in the linac-coherent light source with femtosecond resolution," *Physical Review Letters*, vol. 121, p. 044 801, 4 2018. doi:10.1103/PhysRevLett.121.044801
- [29] X. Huang, J. Corbett, J. Safranek, and J. Wu, "An algorithm for online optimization of accelerators," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 726, pp. 77–83, 2013. doi:10.1016/j.nima.2013.05.046
- [30] Z. Zhang, M. Song, and X. Huang, "Optimization method to compensate accelerator performance drifts," *Physical Review Accelerators and Beams*, vol. 25, p. 122 801, 12 2022. doi:10.1103/PhysRevAccelBeams.25.122801
- [31] J. Duris *et al.*, "Bayesian optimization of a free-electron laser," *Physical Review Letters*, vol. 124, 12 2020. doi:10.1103/PhysRevLett.124.124801
- [32] J. Kirschner, M. Mutný, A. Krause, J. C. de Portugal, N. Hiller, and J. Snuerink, "Tuning particle accelerators with safety constraints using Bayesian optimization," *Physical Review Accelerators and Beams*, vol. 25, p. 062 802, 6 2022. doi:10.1103/PhysRevAccelBeams.25.062802
- [33] C. Xu *et al.*, "Bayesian Optimization for SASE Tuning at the European XFEL," presented at IPAC'23, Venice, Italy, May 2023, paper THPL028, to appear in the proceedings.
- [34] I. Agapov, G. Geloni, S. Tomin, and I. Zagorodnov, "Ocelot: A software framework for synchrotron light source and fel studies," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 768, pp. 151–156, 2014. doi:10.1016/j.nima.2014.09.057
- [35] Z. Zhang *et al.*, "Badger: The Missing Optimizer in ACR," in *Proc. IPAC'22*, Bangkok, Thailand, 2022, paper TU-POST058, pp. 999–1002. doi:10.18429/JACoW-IPAC2022-TUPOST058
- [36] R. Roussel *et al.*, "Xopt: A simplified framework for optimization of accelerator problems using advanced algorithms," presented at IPAC'23, Venice, Italy, May 2023, paper THPL164, to appear in the proceedings.
- [37] A. Marinelli *et al.*, "Multicolor operation and spectral control in a gain-modulated x-ray free-electron laser," *Physical Review Letters*, vol. 111, p. 134 801, 13 2013. doi:10.1103/PHYSREVLETT.111.134801
- [38] J. Duris *et al.*, "Tunable isolated attosecond X-ray pulses with gigawatt peak power from a free-electron laser," *Nature Photonics* 2019 14:1, vol. 14, pp. 30–36, 1 2019. doi:10.1038/s41566-019-0549-5
- [39] V. Kain *et al.*, "Sample-efficient reinforcement learning for CERN accelerator control," *Physical Review Accelerators and Beams*, vol. 23, p. 124 801, 12 2020. doi:10.1103/PHYSREVACCELBEAMS.23.124801
- [40] J. Kaiser, O. Stein, and A. Eichler, "Learning-based optimisation of particle accelerators under partial observability without real-world training," in *Proceedings of the 39th International Conference on Machine Learning*, vol. 162, 2022, pp. 10 575–10 585. <https://proceedings.mlr.press/v162/kaiser22a.html>
- [41] C. Xu *et al.*, "Beam trajectory cocontrol with lattice-agnostic reinforcement learning," in *Proc. IPAC'23*, Venice, Italy, 2023. doi:10.18429/JACoW-IPAC2023-THPL029
- [42] J. Kaiser *et al.*, *Learning to do or learning while doing: Reinforcement learning and bayesian optimisation for online continuous tuning*, 2023. doi:10.48550/arXiv.2306.03739
- [43] C. Emma, A. Edelen, M. J. Hogan, B. O'Shea, G. White, and V. Yakimenko, "Machine learning-based longitudinal phase space prediction of particle accelerators," *Physical Review Accelerators and Beams*, vol. 21, 11 2018. doi:10.1103/PhysRevAccelBeams.21.112802
- [44] A. Hanuka *et al.*, "Accurate and confident prediction of electron beam longitudinal properties using spectral virtual diagnostics," *Scientific Reports*, vol. 11, p. 2945, 1 2021. doi:10.1038/s41598-021-82473-0
- [45] J. Zhu, Y. Chen, F. Brinker, W. Decking, S. Tomin, and H. Schlarb, "High-fidelity prediction of megapixel longitudinal phase-space images of electron beams using encoder-decoder neural networks," *Physical Review Applied*, vol. 16, p. 024 005, 2 2021. doi:10.1103/PhysRevApplied.16.024005
- [46] J. Kaiser, "Machine learning for combined scalar and spectral longitudinal phase space reconstruction," in *Proc. IPAC'23*, Venezia, 2023, pp. 4413–4416. doi:10.18429/jacow-ipac2023-thpl019
- [47] J. Zhu, N. M. Lockmann, M. K. Czwalianna, and H. Schlarb, "Mixed diagnostics for longitudinal properties of electron bunches in a free-electron laser," *Frontiers in Physics*, vol. 10, p. 903 559, 2022. doi:10.3389/FPHY.2022.903559
- [48] S. Corde *et al.*, "Femtosecond X-rays from laser-plasma accelerators," *Rev. Mod. Phys.*, vol. 85, pp. 1–48, 1 2013. doi:10.1103/RevModPhys.85.1
- [49] W. Wang *et al.*, "Free-electron lasing at 27 nanometres based on a laser wakefield accelerator," *Nature*, vol. 595, pp. 516–520, 7868 2021. doi:10.1038/s41586-021-03678-x
- [50] Z.-H. He, B. Hou, V. Lebailly, J. Nees, K. Krushelnick, and A. Thomas, "Coherent control of plasma dynamics," *Nature Communications*, vol. 6, p. 7156, 1 2015. doi:10.1038/ncomms8156
- [51] S. J. D. Dann *et al.*, "Laser wakefield acceleration with active feedback at 5 Hz," *Phys. Rev. Accel. Beams*, vol. 22, p. 041 303, 4 2019. doi:10.1103/PhysRevAccelBeams.22.041303
- [52] R. J. Shalloo *et al.*, "Automation and control of laser wakefield accelerators using bayesian optimization," *Nature Communications*, vol. 11, pp. 1–8, 1 2020. doi:10.1038/s41467-020-20245-6
- [53] S. J. D. Dann *et al.*, "Bayesian optimization of a laser-plasma accelerator," *Physical Review Letters*, vol. 126, 10 2021. doi:10.1103/PhysRevLett.126.104801
- [54] F. Irshad, S. Karsch, and A. Döpp, "Multi-objective and multi-fidelity Bayesian optimization of laser-plasma acceleration," *Physical Review Research*, vol. 5, p. 013 063, 1 2023. doi:10.1103/PHYSREVRESEARCH.5.013063

- [55] A. F. Pousa *et al.*, “Bayesian optimization of laser-plasma accelerators assisted by reduced physical models,” *Physical Review Accelerators and Beams*, vol. 26, p. 084601, 8 2023. doi:10.1103/PhysRevAccelBeams.26.084601
- [56] M. Labat *et al.*, “Seeded free-electron laser driven by a compact laser plasma accelerator,” *Nature Photonics*, vol. 17, pp. 150–156, 2 2023. doi:10.1038/s41566-022-01104-w
- [57] M. Caselle *et al.*, “KAPTURE-2. A picosecond sampling system for individual THz pulses with high repetition rate,” *Journal of Instrumentation*, vol. 12, no. 01, p. C01040, 2017. doi:10.1088/1748-0221/12/01/C01040
- [58] L. Scomparin *et al.*, “KINGFISHER: A Framework for Fast Machine Learning Inference for Autonomous Accelerator Systems,” in *Proc. IBIC2022*, Kraków, Poland. doi:10.18429/JACoW-IBIC2022-MOP42
- [59] L. Scomparin *et al.*, “A Low-Latency Feedback System for Electron Beam Control with Reinforcement Learning,” presented at IPAC’23, Venice, Italy, May 2023, paper THPL027, to appear in the proceedings.
- [60] T. Boltz, “Micro-Bunching Control at Electron Storage Rings with Reinforcement Learning,” Ph.D. dissertation, 2021. doi:10.5445/IR/1000140271