5th ICFA Beam Dynamics Mini-Workshop on Machine Learning for Particle Accelerators

Tuesday 8 April 2025 - Friday 11 April 2025

CERN

Book of Abstracts

Contents

Poster session / **3**

Exploring interpretable deep learning architectures for anomaly detection in SOLARIS synchrotron

Author: Michał Piekarski[1]

[1] *NSRC SOLARIS, Jagiellonian University*

Corresponding Author: michal.piekarski@uj.edu.pl

The National Synchrotron Radiation Center SOLARIS, third generation light source, is the only synchrotron located in Central-Eastern Europe, in Poland. The SOLARIS Center, with seven fully operational beamlines, serves as a hub for research across a diverse range of disciplines. The most important aspect of such research infrastructure is to provide stable working conditions for the users, operators and the conducted projects. Due to its unique properties, problem complexities, and challenges that require advanced approaches, the problem of anomaly detection and automatic analysis of signals for the beam stability assessment is still a huge challenge that has not been fully developed. To address this problem, an automatic analysis of diagnostic signals on the example of transverse beam profiles has been proposed. Pinhole beamlines are typically installed in the middle and high-energy synchrotrons to thoroughly analyze emitted X-rays and therefore assess electron beam quality. To address the problem, pre-trained convolutional neural network (CNN) models in soft fine-tuning strategy were utilized. Created from scratch database contains over one million transverse beam profile images. The best proposed solution, based on the InceptionV3 architecture, can assess beam quality automatically, based solely on the image itself. It classifies pinhole beamline images into two classes (anomaly/non anomaly) with 94.1% accuracy and 96.6% precision. Finally, interpretability algorithms were employed to perform an analysis of the models and achieved results.

Poster session / **4**

Image segmentation using deep neural networks to eliminate shift in laboratory sample images caused by thermal drift

Author: Mateusz Floras[None]

Corresponding Author: mateusz.floras@uj.edu.pl

The National Synchrotron Radiation Centre SOLARIS, a third-generation light source, is the only synchrotron facility in Central-Eastern Europe, located in Poland. The SOLARIS Centre, equipped with seven fully operational beamlines, serves as a key research hub for a wide array of scientific disciplines. The Centre requires advanced software tools to support the analysis of experimental data and has developed a specialized application to assist scientists in interpreting their results. One significant challenge faced during research is correcting image shifts in laboratory samples due to thermal drift. An additional challenge is the varying contrast, which depends on the amount of energy absorbed by the tested material. While traditional phase correlation algorithms based on Fourier transform have not provided entirely satisfactory outcomes, SOLARIS has integrated deep neural networks to improve image stability and enhance accuracy. In particular, the U-Net model is utilized for segmentation, enabling the detection of elements in often blurry images. Model training is conducted using thousands of images generated from experiments utilizing X-ray radiation, ensuring high precision in data analysis.

MLOps, Infrastructure and Scalability / **5**

Application of boosted decision trees in beam anomaly detection with high throughput data streaming system - 10’+5’

Authors: Jacek Biernat[1]; Maciej Mleczko[1]

[1] *National Synchrotron Radiation Centre*

Corresponding Authors: maciej.mleczko@uj.edu.pl, jacek.biernat@uj.edu.pl

The National Synchrotron Radiation Centre SOLARIS is a third generation light source. SOLARIS, as a big science facility with seven fully operational beamlines, is obligated to provide the best possible conditions for conducting research. One of the ways to create favorable environment is delivering precise tools for teams working across many different fields in SOLARIS. The general problem that still remains alive and it is not fully covered is the detection of beam anomalies. The challenge we faced was to find common points among all of synchrotron's fields and combine them into comprehensive, reliable system which aims to provide continuous operation and improve quality of beam shared for synchrotron users. Solution we would like to present is a system which gathers diagnostics signals (e.g. from BPMs, magnets system, vibration detecting system and other sensors) and ingest them in real-time into machine learning (ML) unit. The current system consists of two parts. The first one is high capable data streaming layer which handles high frequency signal (fast acquisition). The second one is ML unit which uses boosted decision trees (BDT) to perform real-time classification whether beam parameters are good or not and predict future beam dump at the early stage. At this moment the main data source for ML model are signals from Beam Position Monitors. Moreover, algorithms such as Fast Fourier Transform (FFT), Probability Density Function (PDF) or Principal Component Analysis (PCA) are working simultaneously in background and gives diagnostic team factors to make a proper evaluation and take actions before automatic beam dump mechanism will be activated. Additionally, PDF and PCA algorithms give ability to detect exact location of issues.

LLMs and AI Assistants / 9

Towards Agentic AI on Particle Accelerators- 15'+5

Author: Raimund Kammering[None]

Co-authors: Antonin Sulc [1]; Jason St John [2]; Thorsten Hellert

[1] *DESY*

[2] *Fermi National Accelerator Laboratory*

Corresponding Authors: raimund.kammering@desy.de, stjohn@fnal.gov, thellert@lbl.gov, antonin.sulc@desy.de

As particle accelerators grow in complexity, traditional control methods face increasing challenges in achieving optimal performance. This paper envisions a paradigm shift: a decentralized multi-agent framework for accelerator control, powered by Large Language Models (LLMs) and distributed among autonomous agents. We present a proposition of a self-improving decentralized system where intelligent agents handle high-level tasks and communication and each agent is specialized control individual accelerator components.

This approach raises some questions: What are the future applications of AI in particle accelerators? How can we implement an autonomous complex system such as a particle accelerator where agents gradually improve through experience and human feedback? What are the implications of integrating a human-in-the-loop component for labeling operational data and providing expert guidance? We show two examples, where we demonstrate viability of such architecture.

MLOps, Infrastructure and Scalability / 10

Machine learning and computer vision in robotic interventions at CERN boosting maintainability- 10'+5'

Authors: David Guenter Forkel[1]; Eloise Matheson[2]; Luca Rosario Buonocore[2]; Mario Di Castro[2]; Álvaro García González[None]

[1] *University Jaume I (ES)*

[2] *CERN*

Corresponding Authors: eloise.matheson@cern.ch, luca.rosario.buonocore@cern.ch, alvaro.garcia.gonzalez@cern.ch, mario.di.castro@cern.ch, david.guenter.forkel@cern.ch

Robots are used in the CERN accelerator complex for remote inspections, repairs, maintenance, monitoring, autopsy and quality assurance, to both improve safety and machine availability. Past interventions mostly relied on teleoperation of robotic bases and arms, while some current and many future interventions will use autonomous behaviors, largely based on advances in machine learning and computer vision. This transition will reduce the operator workload and allow less expert operators to use robots. Autonomous execution can also ensure consistent control and measurement conditions, which are crucial for maintaining repeatability of results, particularly for radiation surveys which data can be used in historical comparisons, and sensor calibrations. As an example, mobile robots in the SPS currently use an advanced autopiloted measurement procedure with adaptable parameters, a dual solution for automated safety gate crossing and parking based on computer vision and point cloud renderings, and the use of graph-based Simultaneous Localization and Mapping (SLAM) for mapping and navigation in challenging areas of the SPS, such as the six access points and the beam-dump area. Machine learning algorithms such as OCR are used to understand the robot position in the SPS based on plate references on the machine. In the LHC, machine learning and computer vision techniques are combined to detect Beam Loss Monitor (BLM) sensors around the beamline. The pose estimation of these sensors is then used as target locations for a robotic arm integrated in the Train Inspection Monorail (TIM) robot to calibrate each sensor using a radioactive source. Point cloud imagery is also used to buildup environmental awareness of the tunnel, facilitating obstacle avoidance, robust operation and digital twin possibilities for future interventions.

Surrogate Modelling and Digital Twins / 11

Compensating hysteresis for the CERN SPS main magnets with transformers- 15’+5’

Author: Anton Lu[1]

Co-authors: Carlo Petrone [2]; Maurus Otto Taupadel [3]; Michael Schenk [2]; Verena Kain [2]; Vincenzo Di Capua [2]

[1] *Technische Universität Wien (AT)*

[2] *CERN*

[3] *Technische Universitaet Darmstadt (DE)*

Corresponding Authors: carlo.petrone@cern.ch, vincenzo.di.capua@cern.ch, michael.schenk@cern.ch, maurus.taupadel@cern.ch, anton.lu@cern.ch, verena.kain@cern.ch

This work presents a machine learning-based approach for compensating magnetic hysteresis in the main dipole and quadrupole magnets of the multi-cycling CERN SPS, utilizing time series neural architectures like the Temporal Fusion Transformers trained on magnetic field measurements. The predicted magnetic fields enable feed-forward, cycle-by-cycle, compensation through the CERN accelerator control system, thereby improving beam reproducibility and enhancing operational flexibility. Correcting deterministically for hysteresis effects removes the necessity of energy-, and time-consuming magnet-precycling while guaranteeing reproducible beam parameters in the SPS. Achieving a magnetic field prediction accuracy in the range of 5e-5 T, this approach demonstrates significant potential for beam optimization, energy savings and improved sustainability, with potential for implementation for any synchrotron magnet circuit. In addition to addressing the technical implementation and benefits, the challenges associated with real-time autoregressive application, model accuracy, and performance evaluation for cycle-to-cycle operation without feedback will be discussed.

LLMs and AI Assistants / 12

The Journey of Developing Specialized Text Embedding Models-15'+5'

Author: Thorsten Hellert[None]

Co-authors: Andrea Pollastro ; João Montenegro [1]; Marco Venturini [1]

[1] *LBNL*

Corresponding Authors: joaomontenegro@lbl.gov, mventurini@lbl.gov, thellert@lbl.gov, andrea.poll93@gmail.com

The specialized terminology and complex concepts inherent in physics present significant challenges for Natural Language Processing (NLP), particularly when relying on general-purpose models. In this talk, I will discuss the development of physics-specific text embedding models designed to overcome these obstacles, beginning with PhysBERT—the first model pre-trained exclusively on a curated corpus of 1.2 million arXiv physics papers. Building upon this foundation, we turn our attention to accelerator physics, a subfield with even more intricate language and concepts. To effectively capture the nuances of this domain, we developed AccPhysBERT, a sentence embedding model fine-tuned specifically for accelerator physics literature. A key aspect of this development involved leveraging Large Language Models (LLMs) extensively to generate annotated training data, enabling AccPhysBERT to facilitate advanced NLP applications such as semantic paper-reviewer matching and integration into Retrieval-Augmented Generation systems.

Poster session / 13

Batch Spacing Optimization via Reinforcement Learning

Authors: Francesco Maria Velotti[1]; Matthias Remta[2]; Sharwin Rezagholi[3]

[1] *CERN*

[2] *University of Vienna (AT)*

[3] *UAS Technikum Vienna*

Corresponding Authors: matthias.remta@cern.ch, francesco.maria.velotti@cern.ch, sharwin.rezagholi@technikum-wien.at

Beams designated for the LHC are injected in multiple batches into the SPS. With a tight spacing of 200 ns between these batches, the injection-kickers have to be precisely synchronised with the injected beam so that injection oscillations are minimized. Due to machine drifts the optimal settings for the kickers vary regularly. In this paper a Reinforcement Learning agent was developed as an active controller, counteracting the machine drifts by adjusting the settings. The agent was trained entirely on a simulation environment and directly transferred to the accelerator. Slightly higher losses than with the current solution, numerical optimization via the BOBYQA algorithm, were achieved but the agent attained these results much faster. Further research is required to completely replace BOBYQA with an RL-agent.

Keynotes, tutorials and lectures / 14

Large Language Model Chatbots for Enhanced Documentation Access

Author: Florian Rehm[1]

[1] *CERN*

Corresponding Author: florian.matthias.rehm@cern.ch

This tutorial applies Retrieval Augmented Generation (RAG) as a method to improve documentation retrieval in accelerator physics. Participants will learn how combining information retrieval with generative AI models can provide precise, context-aware answers from vast technical resources. The session includes a hands-on demonstration of implementing RAG in combination with Large Language Models (LLMs) and retrieving information from your documentation.

Surrogate Modelling and Digital Twins / 15

Graph Learning for Explainable Operation of Particle Accelerators-15'+5'

Author: Chris Tennant[None]

Co-authors: Daniel Moser [1]; Jundong Li [2]; Song Wang [2]; Theo Larrieu [1]

[1] *Jefferson Laboratory*

[2] *University of Virginia*

Corresponding Author: tennant_cd@yahoo.com

We describe research in deep learning on graph representations of the injector beamline at the Continuous Electron Beam Accelerator Facility (CEBAF) to develop a tool for operations. We leverage operational archived data –both unlabeled and labeled configurations –to train a graph neural network (GNN) via our methods of self-supervised training and supervised fine tuning. We demonstrate the ability of the GNN to distill high-dimensional beamline configurations into low-dimensional embeddings and use them to create an intuitive visualization for operators. By mapping out regions of latent space characterized by good and bad setups, we describe how this could provide operators with more informative, real-time feedback during beam tuning compared to the standard practice of interpreting a set of sparse, distributed diagnostic readings. We further describe the results of a framework that provides users with explanations for why a configuration changes location in the latent space.

Poster session / 16

Simulations on injection optimization with Bayesian algorithm at Korea-4GSR

Author: Hyunchang Jin[1]

Co-authors: Seunghwan Shin [1]; Ki-Moon Nam [1]; Younguk Sohn [1]; Jaeyu Lee [2]; Jaehyun Kim [2]

[1] *Korea Basic Science Institute*

[2] *Pohang Accelerator Laboratory*

Corresponding Author: hcjin98@kbsi.re.kr

The Korea-4GSR project, which was launched in 2021 to accelerate and store electron beams with 4 GeV beam energy, is under construction. The beams extracted from the booster after the acceleration will be injected horizontally to the storage ring with the off-axis injection system, and a sufficient amount of injection space in horizontal direction has been secured through various beam physics simulations. However, the deviation of the injection angle and position caused by various machine errors of the accelerator devices can lower the storage ring injection efficiency and adversely affect the beam operation. Therefore, the study for the optimization of the beam injection from the booster to the storage ring is required, and here we will present the results of the injection optimization

simulations with the Bayesian algorithm in the Korea-4GSR. Also, the simulation results obtained by the Bayesian algorithm will be compared to the results with other algorithms.

Poster session / 17

Modeling of a high-current injector for beam optimization

Author: Evan Scott[None]

Corresponding Author: scotter@nv.doe.gov

End-to-end simulations of intense relativistic electron beams generated by linear induction accelerators (LIA) often involve two-step processes whereby the beam creation is simulated using particle-in-cell (PIC) methods before a handoff to less computationally-expensive methods, e.g. beam envelope solvers, to determine sufficiently robust beam tunes. Because of this hand-off, fields that affect the PIC simulation of the A-K gap region are usually untouched during the tuning process. To allow for magnetic guide field optimization including magnets close to the A-K gap, a machine learning model of an LIA injector system is under development to allow for rapid end-to-end simulations of the electron beam for use in beam optimization problems, e.g. automated magnetic transport field tuning.

This work was done by Mission Support and Test Services, LLC, under Contract No. DE-NA0003624 with the U.S. Department of Energy, the National Nuclear Security Administration's Office of Defense Programs, and supported by the Site-Directed Research and Development Program. DOE/NV/03624–2065.

Surrogate Modelling and Digital Twins / 18

Improving fast beam transport simulations using transfer learning-15'+5'

Author: Paul Stanik III[1]

Co-authors: Brendan Morris [1]; Evan Scott ; Piotr Wiewior [2]; Trevor Burris-Mog [2]

[1] *University of Nevada–Las Vegas*

[2] *Nevada National Security Sites*

Corresponding Authors: wiewiopp@nv.doe.gov, scotter@nv.doe.gov, brendan.morris@unlv.edu, burristj@nv.doe.gov, stanik@unlv.nevada.edu

Fast simulations of intense relativistic electron beams can be sufficiently accurate to allow for tuning of an accelerator's magnetic transport field, but are incapable of capturing all relevant beam physics due to limitations in the model. Because methods that do capture these effects are significantly more computationally-expensive, e.g. particle in-cell simulations, they are fundamentally less useful for optimization problems. Here, transfer learning with high-fidelity, full-physics models is applied to the output of a machine learning model trained on a dataset generated by a fast particle beam simulation to bring the simulation results more in line with experimental data.

This work was done by Mission Support and Test Services, LLC, under Contract No. DE-NA0003624 with the U.S. Department of Energy, and the National Nuclear Security Administration. DOE/NV/03624–2062.

Poster session / 19

Bayesian Optimization for Model-Based Automated Multi-Turn Extraction Tuning at the Proton Synchrotron.

Author: Levente Alex Foldesi[1]

Co-authors: Alexander Huschauer [2]; Massimo Giovannozzi [2]; Michael Schenk [2]; Oleksandr Naumenko [3]; Verena Kain [2]

[1] *University of Zurich (CH)*

[2] *CERN*

[3] *Hamburg University (DE)*

Corresponding Authors: levente.foldesi@cern.ch, oleksandr.naumenko@cern.ch, massimo.giovannozzi@cern.ch, michael.schenk@cern.ch, verena.kain@cern.ch, alexander.huschauer@cern.ch

The beam for CERN's North Area proton physics program is produced through a Multi-Turn Extraction (MTE) scheme at the Proton Synchrotron (PS). Using fourth-order resonant excitation, the beam is split into five beamlets in horizontal phase space, with extraction occurring over five consecutive turns. The quality of the splitting is measured by the uniformity of intensities across the beamlets. The process requires precise tuning of several parameters, including the horizontal tune, transverse feedback gain, and excitation frequency. Given the varying conditions in the PS, optimal parameter settings shift frequently, making manual optimization both challenging and inefficient.

This paper focuses on developing and implementing continuous control for the MTE beam using Bayesian optimization (BO) with Gaussian Processes (GP). We discuss experimental studies on various techniques to improve GP convergence and investigate the use of adaptive BO. The experimental results demonstrate the viability of the proposed solution, eliminating the need for manual tuning for MTE beam production.

Poster session / 20

Enhancing Beamline Simulation in Particle Linear Accelerators through Advanced Deep Learning Techniques

Author: yaxin hu[1]

[1] *Institute of Modern Physics*

Corresponding Author: huyaxin22@mails.ucas.ac.cn

This study focuses on the time-series prediction problem in beamline simulation for particle linear accelerators, aiming to improve simulation accuracy and computational efficiency by introducing advanced deep learning techniques. We compare and evaluate the performance of several deep learning models, including Multi-Layer Perceptron (MLP), Long Short-Term Memory (LSTM) networks, Transformer architectures and their variants (Decoder-Only), as well as hybrid LSTM-Transformer structures, across various beamline tasks.

The experimental accelerator used in this study is the CAFE2 facility located at the Institute of Modern Physics in Lanzhou, and the selected numerical simulation software is TraceWin.

Experimental results demonstrate that the final model achieves an accuracy of 92.1%, and can complete complex beamline simulations in approximately 0.2 seconds, representing an 400-fold speedup compared to traditional numerical computation methods. This study marks the first application of advanced Attention mechanisms in the numerical regression domain of full beamline simulation, and evaluates their effectiveness in accelerator time-series prediction.

The selected optimal model exhibits superior performance in both accuracy and efficiency, providing a robust solution for accelerator design. Additionally, we conduct a comprehensive comparison of time-series prediction models suitable for accelerators, analyzing the strengths and weaknesses of

each model and assessing their applicability to different beamlines. Key factors for selecting the optimal model are summarized, offering valuable insights for future research and development.

This work not only introduces new technical approaches for beamline simulation in particle linear accelerators but also provides a solid theoretical foundation and practical tools for optimizing future accelerator designs. The findings contribute significantly to the field of accelerator physics and pave the way for more efficient and accurate simulation methodologies.

Poster session / 21

Bayesian optimization for the SIS18 injection

Authors: Christoph Reinwald[1]; Nico Madysa[2]; Sabrina Maria Appel[None]

[1] *TU Darmstadt*

[2] *GSI - Helmholtzzentrum fur Schwerionenforschung GmbH (DE)*

Corresponding Authors: nico.madysa@cern.ch, reinwald.christoph2019@outlook.de, s.appel@gsi.de

The complexity of the GSI/FAIR accelerator facility demands a high level of automation to maximize the time for physics experiments. Accelerator laboratories across the globe are investigating numerous techniques to achieve this goal, including classical optimization, Bayesian optimization (BO), and reinforcement learning. This presentation will provide an overview of recent activities in these domains at GSI.
Beginning with conventional optimization, the beam loss during the multi-turn injection into the SIS18 synchrotron was reduced from 40% to 15% in approximately 15 minutes, whereas manual adjustments may takeup to 2 hours.
The implementation of Multi-Objective BO has resulted in the first physical measurement of a Pareto front on the SIS18 injection. Physics-informed BO for automated injection optimization is currently under investigation and has been utilized for a potential two-plane multi-turn injection via simulation.

Poster session / 22

Machine Learning-Driven Optimization of a Silicon Bent Crystal System for Enhanced Slow Extraction Efficiency at CERN-SPS

Author: Francesco Maria Velotti[1]

Co-authors: Francisco Huhn [1]; Luigi Salvatore Esposito [1]; Matthew Alexander Fraser [1]; Michael Schenk [1]; Verena Kain [1]

[1] *CERN*

Corresponding Authors: luigi.salvatore.esposito@cern.ch, francisco.xavier.moreira.huhn@cern.ch, matthew.alexander.fraser@cern.ch, francesco.maria.velotti@cern.ch, verena.kain@cern.ch, michael.schenk@cern.ch

The CERN-SPS slow extraction has been recently equipped with a silicon bent crystal to reduce losses at the electrostatic septum (ES) wires. Such a concept exploits the coherent deflection that a thin crystal can give to part of the separatrix to avoid the ES wires.
In this contribution, we show how machine learning played a fundamental role in the design and operational deployment of the Si bent crystal for slow extraction loss reduction. First, we present how a multi-fidelity Gaussian Process (GP) based optimization was used to optimize the crystal design. Then, we detail how we exploited a neural network based surrogate model to design an efficient live controller. The final result shows a system that should be able to reduce extraction losses by one order of magnitude.

Poster session / 23

Machine Learning-Based Anomaly Diagnostics for Optimizing SRF Operation at CAFE2

Author: Lijuan Yang[1]

Co-authors: Feng Qiu [1]; Yuan He [1]; JingWei Yu [1]; Jiayi Peng [1]; Jinying Ma [1]; Zhenglong Zhu [1]; Tiancai Jiang [1]; Zheng Gao [1]

[1] *Institute of Modern Physics, Chinese Academy of Sciences*

Corresponding Authors: hey@impcas.ac.cn, gaozheng@impcas.ac.cn, zhuzhl@impas.ac.cn, qiufeng@impcas.ac.cn, yanglijuan@impcas.ac.cn, yujingwei@impcas.ac.cn, jiangtiancai@impcas.ac.cn, majingying@impcas.ac.cn, pengjiayi@impcas.ac.cn

In 2021, the Chinese ADS Front-end demo superconducting radio-frequency (SRF) linac, known as CAFe, successfully conducted a commissioning of a 10 mA, 200 kW continuous wave proton beam. During this commissioning, it was observed that the SRF faults are the leading causes of short machine downtime trips, contributing to approximately 70% of total beam trips. Analyzing fault data and identifying fault types is a time-consuming and laborious process, especially for large modern accelerators with hundreds of RF cavities. Here, we propose a machine learning (ML) based model for automating SRF cavity fault recognition. First, a comprehensive study of the cavity fault mechanisms was conducted, leading to the identification of several distinct fault patterns. Next, we converted the "expert reasoning"process into a "model inference"process though feature engineering. We demonstrate the feasibility of this method using the CAFE2 (CAFe had been upgraded to CAFE2 after 2021) facility, achieving an accuracy of over 90%. By combining ML mdoel with big data analysis, we identified three different mechanisms of quench at CAFE2 and developed an automatic detection algorithm to distinguish them, which is crucial for the prevention of quench. Although the specific faults in SRF cavities may vary across different accelerators, similarities exist in the RF signals. Therefore, this study provides valuable guidance for the fault analysis of the whole SRF community.

Anomaly Detection and Diagnostics / 24

Anomaly Forecasting and Adaptive Learning in Fast Kicker Magnet Systems- 15'+5'

Author: Malik Marco Algelly[1]

Co-authors: Alexandros Kalousis [2]; Verena Kain [3]; Kostas Papastergiou [3]; Patrick James Ellison ; Francesco Maria Velotti [3]

[1] *Universite de Geneve (CH)*

[2] *University of Geneva*

[3] *CERN*

Corresponding Authors: francesco.maria.velotti@cern.ch, verena.kain@cern.ch, patrick.ellison@cern.ch, alexandros.kalousis@unige.ch, malik.marco.algelly@cern.ch, k.papastergiou@cern.ch

Kicker magnets are essential for particle beam injection and extraction within CERN's accelerator complex, where high reliability is crucial to maintaining the availability needed for numerous scientific experiments. This study proposes a machine learning approach for forecasting anomalies in these systems, aiming to proactively identify and isolate potential faults before failure occurs. To keep the anomaly detection model accurate over time, continual learning techniques are employed, allowing the model to adapt to evolving system dynamics without frequent retraining. This combination enhances the efficiency and stability of accelerator operations by ensuring the model remains up-to-date in the face of non-static data.

Poster session / 25

High-Speed Differentiable Simulations as an Enabler for Reinforcement Learning-based FEL Tuning Agents

Author: Jan Kaiser[None]

Co-authors: Ryan Roussel ; Zihan Zhu ; Jenny Morgan [1]; Kishansingh Rajput [2]; Malachi Schram ; Auralee Edelen ; Daniel Ratner [3]; Annika Eichler [4]

[1] *University of strathclyde*

[2] *Jefferson Lab*

[3] *SLAC*

[4] *DESY*

Corresponding Authors: jan.kaiser@desy.de, zihan.zhu@desy.de, dratner@slac.stanford.edu, jenny.morgan@strath.ac.uk, annika.eichler@desy.de, edelen@slac.stanford.edu, rroussel@slac.stanford.edu, schram@jlab.org, kishan@jlab.org

The photon pulse intensity is one of the key performance metrics of Free Electron Laser (FEL) facilities and has a direct impact on their experimental yield. To date, FEL intensity tuning is a time-consuming manual task that requires expert human operators to have significant skill and experience. Autonomous tuning methods have been demonstrated to reduce setup times and improve the attained working points on other tuning tasks, but existing numerical optimisation algorithms struggle with the high dimensionality and complexity of FEL intensity maximisation. Reinforcement learning-trained optimisers (RLOs) promise the capabilities to successfully be applied to even such complex tuning tasks, but their requirement for training interactions makes online training infeasible. Offline training in simulation has previously been proposed as an effective solution, but FEL simulations are prohibitively complex and slow for use in RLO training. We propose the coupling of novel high-speed differentiable beam dynamics simulations with modular neural network surrogates to address this issue at the example of tuning the FEL intensity at LCLS. Not only have similar approaches been shown to reduce the wall-time required for training by orders of magnitude, making training for FEL tuning feasible, but the availability of gradients also promises an avenue for profoundly reducing the number of required training samples through gradient-based reinforcement learning.

Poster session / 26

Application of Reinforcement Learning on online Coupling Correction for Storage Ring

Author: Yihao GONG[1]

[1] *Shanghai Advanced Research Institute*

Corresponding Author: gongyh@sari.ac.cn

In modern synchrotron light sources, maintaining beam stability is critical for ensuring high-quality synchrotron radiation performance. Light source stability is governed by stability of current, beam position and beam size. Beam size stability on the order of several microns need to be improved for future experiments. Reinforcement learning (RL) offers a promising approach for real-time beam size feedback system. The RL framework consists of an intelligent agent that interacts with the environment to maximize a cumulative reward, based on state observations and actions. Beam size measurement at one point and vertical dispersion are observations of RL environments, which can present beam size distribution along the storage ring. Through simulation and real experimental setups, we demonstrate the efficacy of the PPO algorithm which adapt to discrete action spaces in controlling beam stability and correcting coupling. Further optimization of hyperparameters during simulation environment is applied in real operation. The approach provides a significant improvement in online, real-time correction of coupling errors, offering a faster and more adaptable solution compared to traditional methods.

Surrogate Modelling and Digital Twins / 27

Optimizing High-Energy Beam Transport with AI: Advances in IFMIF-DONES Design- 15'+5'

Authors: Galo Gallardo Romero[1]; Guillermo Rodríguez Llorente[2]

Co-authors: Lucas Magariños Rodríguez [1]; Nikita Khvatkin Petrovsky [1]; Roberto Gómez-Espinosa Martín [1]; Rodrigo Morant Navascués [1]; Ruben Lorenzo Ortega [3]

[1] *HI Iberia*

[2] *Artificial Intelligence Engineer at HI Iberia, Department of Mathematics at UC3M, Gregorio Millán Barbany Institute*

[3] *IFMIF-DONES Spain*

Corresponding Authors: nkhvatkin@hi-iberia.es, grodriguez@hi-iberia.es, ruben.lorenzo@ifmif-dones.es, ggallardo@hi-iberia.es, robertogemartin@hi-iberia.es, rmorant@hi-iberia.es, lmagarinos@hi-iberia.es

A major challenge in constructing future nuclear fusion power plants is understanding how reactor materials are damaged by the neutron flux generated during the fusion process. In order to address this challenge, the IFMIF-DONES neutron source is being built for material irradiation, generating the necessary neutron flux through a stripping reaction between accelerated deuterons and a lithium circuit. In this work, within the DONES-FLUX project, Fourier Neural Operators are employed as Deep Learning Surrogate Models for optimizing the design of the High-Energy Beam Transport Line of the IFMIF-DONES accelerator. The trained models, which predict the deuteron beam statistical functions and beam profile distributions, are roughly 3 orders of magnitude faster than traditional simulations while keeping mean absolute percentage errors below 5%. This significant reduction in inference time, along with the models' differentiability, enables the use of optimization algorithms like online Reinforcement Learning, Bayesian Optimization, and Gradient Descent. Additionally, this last method was implemented and tested for finding optimal quadrupole values for different beam configurations, where solutions are reached within minutes. These positive results highlight the synergy between different Deep Learning architectures and offer a promising collaboration between the field of Artificial Intelligence and accelerator facilities.

Poster session / 29

Automation of Device Server creation in TANGO Controls using Large Language Models

Author: Krzysztof Klimczyk[None]

Corresponding Author: krzysztof.klimczyk@s2innovation.com

This project explores methods for automating the creation of device servers in the TANGO Controls environment using Large Language Models (LLMs). The primary goal is to streamline and accelerate the coding process for device servers, reducing the time and effort required by developers. The application features a web-based user interface where users specify device attributes, commands, and properties, while the backend generates functional server code using LLMs such as GPT, Claude, or Gemini.

A key aspect of this project is using technical documentation as a contextual knowledge source for the LLMs. By employing the Retrieval-Augmented Generation (RAG) technique, the application automatically selects relevant documentation fragments to enhance code accuracy, particularly for implementing device-specific features.

This poster presents the system architecture, underlying technologies, and real-world testing outcomes, illustrating the potential of LLMs to improve device server development in distributed control environments.

The project concept was initially presented at the last TANGO Community Meeting during the "lightning talks"and is part of a master's thesis conducted at Jagiellonian University in Poland, in cooperation with the commercial company S2Innovation, which specializes in TANGO Controls software solutions.

Poster session / 30

Time-varying Bayesian optimisation of injection in the Proton Synchrotron Booster

Author: Francisco Huhn[1]

Co-authors: Chiara Bracco [1]; Francesco Maria Velotti [1]

[1] *CERN*

Corresponding Authors: francisco.xavier.moreira.huhn@cern.ch, chiara.bracco@cern.ch, francesco.maria.velotti@cern.ch

The Proton Synchrotron Booster (PSB) receives 160 MeV H^- ions, which are converted to protons at injection via a charge exchange mechanism, an upgrade that allows the production of high-intensity beams ($> 10^{13}$ per ring). Nevertheless, with the increase in intensity, space-charge losses arise. To mitigate these effects, horizontal phase-space painting is performed with a system of fours kickers whose pulse is customisable via time and amplitude parameters.

Recent work has shown that classical optimisation algorithms can find the optimal parameter values on both a digital twin and the real machine. However, these techniques: do not handle system-state time variations, do not continually update the parameters during operation, require non-negligible dedicated beam time and are usually not robust to observation noise.

We suggest time-varying Bayesian optimisation and show that it performs well, addressing each of the previous issues. Moreover, it is simple to develop and deploy (i.e. cheap in human capital), and runs continually, with minimal human interaction. We also discuss context-based variants of this algorithm. This work improves the operation of the PSB and contributes towards the goal of automating the operation of particle accelerators.

Poster session / 31

EPU Coupling Correction by Bayesian Optimization

Author: Mau-Sen Chiu[1]

[1] *National Synchrotron Radiation Research Center*

Corresponding Author: chiu.ms@nsrrc.org.tw

The APPLE-II types of elliptically polarized undulators (EPU) are extensively utilized in synchrotron light sources. Manufacturing imperfections in the EPU inevitably lead to the creation of a residual skew quadrupole component, which couples horizontal betatron oscillation and dispersion to the vertical plane, consequently altering the vertical beam size. To regulate the vertical beam size, EPU is equipped with a pair of skew quadrupole magnets to compensate for the residual skew quadrupole compo-nent of the EPU. The paper focuses on building a 2D coupling feed-forward table by using Bayesian optimiza-tion at the Taiwan Photon Source (TPS)

Poster session / 32

Anomaly Detection for fault analysis in the Proton Synchrotron using Machine Learning

Author: Joel Axel Wulff[None]

Co-author: Alexandre Lasheen [1]

[1] *CERN*

Corresponding Authors: joel.wulff@cern.ch, alexandre.lasheen@cern.ch

The CERN Proton Synchrotron (PS) is equipped with several RF systems covering a wide range of revolution harmonics (7 to 21, 42, 84) with heterogeneous hardware of different age. Despite a good track record of high availability (over 99%), a degradation observed from 2021 to 2023 highlighted a need for enhanced fault diagnosis of these RF systems. Accurate identification of failure sources is crucial to maintaining optimal performance and high availability, particularly under more demanding beam conditions such as increased intensities. While some faults are easily detected and understood through self-reporting devices, others may cause subtle deterioration that goes unnoticed. To enable more accurate root-cause analysis of faults, we propose machine learning-based anomaly detection models trained on data from both beam and hardware signals. These models could be deployed online to provide early alerts for changing machine states, or they could facilitate post-mortem analyses by identifying the last known normal state prior to a fault.

Poster session / **33**

Operational Results of Data-Driven Automated Intensity Optimization at CERN's LEIR.

Author: Borja Rodriguez Mateos[1]

Co-authors: Maciej Slupecki [2]; Michael Schenk [2]; Theodoros Argyropoulos [2]; Verena Kain [2]

[1] *Universitat Politecnica Catalunya (ES)*

[2] *CERN*

Corresponding Authors: michael.schenk@cern.ch, theodoros.argyropoulos@cern.ch, verena.kain@cern.ch, borja.rodriguez.mateos@cern.ch, maciej.slupecki@cern.ch

Aging of the stripper foil and unexpected machine shutdowns are the primary causes for reduction of the injected intensity from CERN's linac3 into the Low Energy Ion Ring (LEIR). As a result, the set of optimal control parameters that maximizes beam intensity in the ring tends to drift, requiring daily adjustments to the machine control settings. In this paper, several data-driven methods such as Bayesian Optimization (BO) and Reinforcement Learning (RL) are compared for the design of an autonomous controller of linac3 and LEIR parameters to maximize beam intensity before RF capture. Through the evaluation of both black-box and stateful approaches to solving this high-dimensional maximization problem, we aim to design an optimal controller that meets both performance and sample efficiency constraints.

Optimisation and Control / **34**

Machine Learning for Online Control of Particle Accelerators-15'+5'

Authors: Xiaolong Chen[1]; Zhijun Wang[1]; Yuan He[1]; Hong Zhao[2]; Chunguang Su[1]; Shuhui Liu[1]; Weilong Chen[1]

[1] *Institute of modern physics*

[2] *XIAMEN University*

Corresponding Authors: chenweilong@impcas.ac.cn, wangzj@impcas.ac.cn, chenxiaolong@impcas.ac.cn, suchunguang@impcas.ac.cn, hey@impcas.ac.cn, zhaoh@xmu.edu.cn, liush@impcas.ac.cn

Particle accelerators play a critical role in modern scientific research. However, existing manual beam control methods heavily rely on experienced operators, leading to significant time consumption and potential challenges in managing next-generation accelerators characterized by higher beam current and stronger nonlinear properties. In this paper, we establish a dynamical foundation for designing the online adaptive controller of accelerators using machine learning. This provides a guarantee for dynamic controllability for a class of scientific instruments whose dynamics are described by spatial-temporal equations of motion but only part variables along the instruments under steady states are available. The necessity of using historical time series of beam diagnostic data is emphasised. Key strategies involve also employing a well-established virtual beamline of accelerators, by which various beam calibration scenarios that actual accelerators may encounter are produced. Then the reinforcement learning algorithm is adopted to train the controller with the interaction to the virtual beamline. Finally, the controller is seamlessly transition to real ion accelerators, enabling efficient online adaptive control and maintenance. Notably, the controller demonstrates significant robustness, effectively managing beams with diverse charge mass ratios without requiring retraining.This controller enables global control, achieving up to 42-dimensional synchronous regulation across the entire superconducting section of the China Accelerator Facility for Superheavy Elements.

Poster session / 35

Surrogate Model Training and Applications at the S-DALINAC based on FCNNs

Author: Dominic Schneider[1]

Co-authors: Felix Schliessmann ; Jonny Birkhan ; Michaela Arnold ; Norbert Pietralla

[1] *Institut für Kernphysik, TU Darmstadt*

Corresponding Authors: fschliessmann@ikp.tu-darmstadt.de, jbirkhan@ikp.tu-darmstadt.de, dschneider@ikp.tu-darmstadt.de, marnold@ikp.tu-darmstadt.de, pietralla@ikp.tu-darmstadt.de

Machine learning methods provide a significant potential for the optimized operation of complex facilities, such as particle accelerators. In this contribution, the first training and application of surrogate models to the electron accelerator S-DALINAC based on Fully-Connected Neural Networks (FCNN) will be presented.

An exhaustive data-mining algorithm has been developed to generate the training data using the live accelerator. The architecture and training of the surrogate model, as well as its introspection using both Sobol-based sensitivity analysis and Shap impact analysis will be presented. Additionally, the test of optimization algorithms on the surrogate model prior to the application on the live system are presented in this contribution. Future prospects of transforming this surrogate model to an online digital twin application, as well as drawbacks in the applicability of the model, will be discussed.

Poster session / 36

Research on AI-Assisted Operation System for Particle Accelerators

Authors: Zhijun wang[1]; Xiaolong Chen[2]; Yuan He[2]; Weilong Chen[2]; Chunguang Su[2]

[1] *IMP,CAS*

[2] *Institute of modern physics*

Corresponding Authors: wangzj@impcas.ac.cn, hey@impcas.ac.cn, chenxiaolong@impcas.ac.cn, suchunguang@impcas.ac.cn, chenweilong@impcas.ac.cn

The long-term stable operation of high-power accelerators demands high availability, reliability, and safety, requirements that traditional labor-intensive maintenance methods struggle to meet. In recent years, with the rapid advancement of artificial intelligence and machine learning technologies, various efficient algorithms have demonstrated great potential in the field of particle accelerators. However, a unified and efficient framework to integrate these tools for intelligent tuning and operation is still lacking. This contribution proposes a three-layer system architecture based on predictive control, comprising a data layer, a modeling layer, and a decision layer. The data layer is responsible for collecting and processing accelerator operation data, establishing a solid foundation for model construction and system decision-making. The modeling layer leverages machine learning algorithms to build predictive models for key accelerator components and subsystems. The decision layer plays a critical role in two areas: optimizing the beam tuning process through intelligent decision-making and providing dynamic forecasts of the accelerator's operational state and early fault detection based on predictive modeling, enabling real-time intervention and intelligent control of the system. This architecture is expected to significantly enhance the tuning efficiency and operational performance of accelerators, offering new solutions and technical support for the development of future intelligent accelerator systems.

LLMs and AI Assistants / 37

Multi-modal LLMs at EuXFEL: Knowledge Assistants, Data Exploration, and Policy- 15'+5'

Authors: Danilo Enoque Ferreira De Lima[1]; Illia Derevianko[1]; Krzysztof Wrona[1]; Lisa Brodersen[1]; Luca Gelisio[1]; Mahmoud Adjami[1]; Maurizio Manetti[1]; Robert Rosca[1]; Steffen Hauf[None]

[1] *European XFEL GmbH*

Corresponding Authors: illia.derevianko@xfel.eu, danilo.enoque.ferreira.de.lima@xfel.eu, lisa.brodersen@xfel.eu, robert.rosca@xfel.eu, steffen.hauf@xfel.eu, luca.gelisio@xfel.eu, krzysztof.wrona@xfel.eu, maurizio.manetti@xfel.eu, mahmoud.ajami@xfel.eu

Multimodal Large Language models extend LLMs'capabilities to input beyond text, often images. At the European XFEL, these models are used as Retrieval-Augmented Generative (RAG) Knowledge assistants in technical and administrative domains. We present a selection of current applications and prototypes: chatbot assistants for data service support, business travel aid, vision-based document exploration systems, coding assistants, and an application rating control system panels for accessibility and design. The aforementioned examples were studied as part of an effort to create a policy for the usage of LLMs by facility staff. We highlight how this policy is influenced by the aforementioned examples and associated legal and data privacy considerations.

Anomaly Detection and Diagnostics / 38

A path to efficient machine learning-based beam diagnostics: complete six-dimensional generative phase space reconstruction without RF deflecting cavity- 15'+5'

Authors: Seongyeol Kim[None]; Haeryong Yang[1]; Gyujin Kim[1]; Ryan Roussel[None]; Juan Pablo Gonzalez Aguilera[2]; Auralee Edelen[None]

[1] *Pohang Accelerator Laboratory*

[2] *University of Chicago*

Corresponding Authors: jpga@uchicago.edu, edelen@slac.stanford.edu, sykim12@postech.ac.kr, ilyoukim@postech.ac.kr, rroussel@slac.stanford.edu, highlong@postech.ac.kr

Generative phase space reconstruction method based on neural networks and differentiable simulations has become a novel beam diagnostic technique to obtain the beam phase space information. Recent studies show that four-dimensional phase space can be successfully obtained by using only YAG images with different quadrupole magnet strength, allowing us to understand both uncoupled and coupled phase spaces. Furthermore, it has been experimentally demonstrated that the complete six-dimensional phase space can be reconstructed by additionally utilizing a spectrometer dipole magnet and RF transverse deflecting cavity. In addition to the previous research activities, we are currently investigating the complete six-dimensional phase space reconstruction method that does not require the RF transverse deflecting cavity. We demonstrate in simulation that our proposed method can also provide complete six-dimensional phase spaces including all the transverse-longitudinal couplings, which successfully represent the ground truth distributions. In this study, we present how to perform the reconstruction without such an advanced diagnostic instrument. In addition, we show the reconstruction results with synthetic examples and actual experimental data obtained at the Pohang Accelerator Laboratory X-ray Free Electron Laser (PAL-XFEL) facility.

Poster session / 39

Using Hyper-GAN to improve Uncertainty Quantification for Surrogate Models

Author: Kathryn Baker[None]

Corresponding Author: k.baker@stfc.ac.uk

Gaussian Processes are the most popular and accurate modelling technique for estimating uncertainty in a system. However, they struggle to scale with high dimensional data and large datasets that are often required to train surrogate models of accelerator components. Typical alternatives include deep ensembles, Monte-Carlo Dropout and quantile regression. However, each of these methods has shortfalls, resulting in either a highly overconfident or underconfident estimate of the uncertainty in the model. This work will explore whether Hyper-GANs, generative models that learn distributions of neural network parameters, can be used to address these problems and provide accurate estimates of uncertainty for high and multi-dimensional outputs for surrogate models of accelerators. We will use a model of the ISIS MEBT as a case study and compare the quality of the uncertainty predictions against predictions using quantile regression.

Poster session / 40

Exploring the Adaptation of LLM and the Formation of AI Governance at European XFEL

Author: Mahmoud Ajami[1]

Co-authors: Luca Gelisio [1]; Paul Drews

[1] *European XFEL GmbH*

Corresponding Authors: luca.gelisio@xfel.eu, paul.drews@leuphana.de, mahmoud.ajami@xfel.eu

At the European XFEL, Large Language Models provide opportunities to facilitate technical and administrative workflows, with applications such as chatbots for technical support, travel assistance and coding assistance. A task force was formed to evaluate these applications and to develop a European XFEL AI policy, focusing on technical assessment, framework creation, and ensuring ethical and regulatory compliance to guarantee safe adaptation and effective governance. Following a

socio-technical perspective, this work utilizes grounded theory and the Gioia method to explore and analyze decision-making, learning outcomes, storytelling, and the unique considerations of implementing such technology in a research setting like European XFEL. This contribution also elaborates on task force members' future speculations regarding this technology at the European XFEL, and provides input for other organizations considering similar technologies.

Poster session / 41

First Attempts at Tuning the ISIS LEBT with Xopt and Bayesian Optimization

Authors: Kathryn Baker[None]; Raunakk Banerjee[None]

Corresponding Authors: k.baker@stfc.ac.uk, raunakk.banerjee@stfc.ac.uk

Until now, the majority, if not all, accelerator tuning at the ISIS Accelerators has been done manually. With the migration of the control system to EPICS as well as the development of accessible optimization frameworks such as Xopt, we have made the first attempts to automatically tune the Low Energy Beam Transport (LEBT) at ISIS. In this presentation we will discuss the specifics of the optimization problem being addressed including the variables, constraints and objectives. We will then share results of the first optimization runs using Xopt, discussing some of the setbacks and issues we faced during the process and the future work required to address each of these.

Poster session / 42

Machine Learning Applied to BGI Profile Measurements

Author: Javier Martinez Samblas[None]

Co-author: Manuel Gonzalez Berges [1]

[1] *CERN*

Corresponding Authors: manuel.gonzalez@cern.ch, javier.martinez.samblas@cern.ch

A key measurement in CERN accelerators for beam diagnostics is transverse size. The Beam Gas Ionization (BGI) instrument enables non-destructive observation of transverse beam profiles by detecting free electrons produced through beam-gas ionization using a Timepix-family detector. However, BGI profiles often suffer from artifacts, such as beam losses, which degrade profile quality and significantly hinder analysis. This contribution addresses the challenge of background removal through machine learning techniques, employing both supervised and unsupervised approaches to enhance the accuracy and reliability of BGI profile measurements. Additionally, aspects such as the performance and time complexity of these methods are analyzed to ensure the instrument remains fully operational.

Poster session / 43

Simulation package for the Beam Synchrotron Radiation Longitudinal density monitor (BSRL)

Author: Alexander Nicholas Jury[1]

Co-authors: Stefano Mazzoni [2]; Alexander Hill [3]; Hao Zhang [4]; Manuel Gonzalez Berges [2]

[1] *University of Liverpool (GB)*

[2] *CERN*

[3] *Liverpool*

[4] *University of Liverpool/Cockcroft Institute*

Corresponding Authors: hao.zhang@cockcroft.ac.uk, manuel.gonzalez@cern.ch, alexander.nicholas.jury@cern.ch, ahill94@liverpool.ac.uk, stefano.mazzoni@cern.ch

The Beam Synchrotron Radiation Longitudinal density monitor (BSRL) at the LHC leverages time-correlated single-photon counting to provide high-dynamic-range measurements of particle populations within each bunch in the LHC including monitoring of "ghost"and "satellite"bunches, which represent charge captured in nominally empty buckets, thereby enhancing the accuracy of luminosity calibration. To further improve BSRL performance and assess the impact of hardware modifications and data processing adjustments on measurement accuracy, we are developing a simulation package of the BSRL. This will allow for in-depth analysis and optimization of this complex device.

Optimisation and Control / 44

Progress on Automating Experiments at the Argonne Wakefield Accelerator Using Advanced Bayesian Optimization Algorithms - 15'+5'

Author: Ryan Roussel[None]

Co-authors: Auralee Edelen ; John Power ; Dylan Kennedy [1]; Eric Wisniewski [2]

[1] *SLAC National Laboratory*

[2] *Argonne National Laboratory*

Corresponding Authors: ewisniew@anl.gov, edelen@slac.stanford.edu, kennedy1@slac.stanford.edu, rroussel@slac.stanford.edu, jp@anl.gov

The goal of machine learning for accelerator control is to automate the start-up, optimization, and execution of experiments at accelerator facilities with limited-to-no human operator input. To address this challenge, we have been pursuing a research program to completely automate sequential accelerator beamline configuration tasks at the Argonne Wakefield Accelerator (AWA). These tasks include beam alignment with magnetic elements, RF cavity phasing synchronized to the photoinjector laser system, and transverse phase space control, all in tightly constrained parameter spaces. In this work, we describe progress towards automating the AWA beamline and the development of novel Bayesian optimization algorithms needed to efficiently address each of these tuning tasks. We highlight the development of Bayesian Algorithm Execution (BAX) acquisition functions to generalize Bayesian optimization algorithms to optimize so-called "virtual"objectives that are inferred from Gaussian process models of observable quantities. Finally, we describe the path forward for achieving full autonomous execution of beamline configuration and experiment execution in the near future.

Poster session / 45

Model Predictive Control for energy efficient HVAC systems for CERN accelerators

Authors: Diogo Monteiro[1]; Nikolina Bunijevac[None]

[1] *CERN*

Corresponding Authors: nikolina.bunijevac@cern.ch, diogo.monteiro@cern.ch

In times of concern over the environmental impact of high-energy physics organizations, our research in CERN's Cooling and Ventilation group (EN/CV) investigates energy-saving strategies for heating, ventilation, and air conditioning (HVAC) systems. Widely used in both residential and industrial settings, HVAC systems contribute up to 40% of residential and 70% of industrial consumption, making their optimization a global concern. At CERN, cooling and ventilation systems are used to ensure appropriate temperature and humidity conditions in the accelerator complex. Together with general water-cooling systems, these systems account for up to 15% of total electricity consumption of CERN's flagship accelerator.
Despite their energy intensity, these systems are typically managed by classical controllers, which are reliable but not optimal in terms of energy efficiency.
This study aims to quantify the potential energy savings in HVAC systems using model predictive control (MPC), a modern advanced control strategy that incorporates behaviour prediction and external data, such as weather forecasts. The methodology involves coding both classical and advanced controllers in a virtual environment, developing a digital twin model for a selected plant, and running simulations to confirm the improved thermal performance and electricity reduction with the MPC approach. In a quest for a reproducible solution that can be easily adapted to different HVAC plants, the digital twin is built using neural networks, following recent advances in academic research.

Poster session / **46**

Updates to the Xopt/Badger Ecosystem for Advanced Optimization and Online Control

Author: Ryan Roussel[None]

Co-authors: Auralee Edelen ; Dylan Kennedy [1]; Jeff Larson [2]; John-Luke Navarro [2]; Remi Lehe ; Sara Miskovich [1]; Stephen Hudson [2]; Zhe Zhang [1]; Yekta Yazar [1]

[1] *SLAC National Accelerator Laboratory*

[2] *Argonne National Laboratory*

Corresponding Authors: yazar@slac.stanford.edu, zhezhang@slac.stanford.edu, smiskov@slac.stanford.edu, rlehe@lbl.gov, shudson@anl.gov, jmlarson@anl.gov, kennedy1@slac.stanford.edu, rroussel@slac.stanford.edu, edelen@slac.stanford.edu, jnavarro@anl.gov

The Xopt/Badger ecosystem offers a versatile suite of tools designed to address the growing needs of advanced optimization and online control in scientific applications. The goal of these tools is to standardize the implementation and use of advanced optimization algorithms at arbitrary scientific facilities for the benefit of the wider accelerator community. In this work, we provide a summary of updates to Xopt and Badger that enable new capabilities and improve ease of use. This includes new developments in trust-region approaches to Bayesian optimization and GUI-based online visualization of surrogate models. Finally, we discuss the implementation of generator standards established between Xopt, libensemble, and optimus packages to allow for future interoperability, enabling robust usage of advanced optimization algorithms in both experiment and at high performance computing clusters.

Poster session / **48**

Surrogate Model-Driven Reinforcement Learning for Efficient Beam Tuning in Accelerators

Author: Chunguang Su[1]

Co-authors: Xiaolong Chen [2]; Zhijun wang [3]

[1] *Institute of modern physics, Chinese Academy of Sciences*

[2] *Institute of modern physics*

[3] *IMP,CAS*

Corresponding Authors: suchunguang@impcas.ac.cn, wangzj@impcas.ac.cn, chenxiaolong@impcas.ac.cn

Beam tuning in particle accelerators is a complex task, especially when physical modeling is impractical due to the lack of complete beam diagnostics. Traditional methods often rely on iterative manual tuning by operators, which can be inefficient. Reinforcement learning (RL) algorithms offer a promising alternative for automating this process. In this work, we demonstrate the successful application of RL-based policies to beam tuning for the High-Intensity Proton Injector (HIPI), where traditional physical modeling was not feasible.
The policy is trained using a surrogate model, which was built from online data collection. Our results show that the surrogate model significantly enhances the training efficiency, reducing the time required for the RL agent to learn an effective control policy. Moreover, the policy demonstrated robust performance in real-world testing, achieving approximately 90% beam transmission within minutes. This approach provides a practical solution for environments where traditional physical models are unavailable, showcasing the potential of RL in optimizing accelerator operations.

Poster session / 49

Testing for Beam Controls Applications Development

Author: Gopika Bhardwaj[1]

[1] *Fermi National Accelerator Laboratory*

Corresponding Author: gopikab@fnal.gov

In the process of converting a tuning script written in Fermilab's in-house accelerator scripting language (ACL) to Python, we have implemented new, innovative testing strategies.
This was done because we do not want to test in production with live variables and Github Actions test suite doesn't have access to real-time data.

For code testing, we wrote tests before writing functions with 3 different strategies -
1. We constructed mock variables in our Accelerator Controls Network which contain historical values of the live variables so we can perform end-to-end testing.
2. We have a pytest suite integrated with Github Actions to test functions that perform mathematical computations like checking beam loss ratio and whether beam change is extreme or beam is oscillating.
3. Finally for functions that require input from the Accelerator Controls Network, we have interactive tests to check if get and set variables are working as expected.

Poster session / 50

Fluid Refill Prediction

Author: Gopika Bhardwaj[1]

[1] *Fermi National Accelerator Laboratory*

Corresponding Author: gopikab@fnal.gov

Cooling water systems are common to almost every accelerator and necessary for their smooth operations. Currently, fluid refill time is computed manually from level sensors. The task of ensuring that the fluid level remain sufficient for operations can be automated in a way that supports expert human effort.

As an initial approach, simple linear regression method was explored. The challenges involve automatically finding refills and limited device resolution. A website has been created which is in use by support staff to ensure water cooling refills happen in time and where possible, during normal work hours ensuring that beam is available around the clock.

A machine learning approach is being explored and compared to the simple linear regression technique for better accuracy and robustness.

Optimisation and Control / 51

Efficient Dynamic and Momentum Aperture Optimization for Lattice Design Using Multipoint Bayesian Algorithmic Execution-15'+5'

Authors: Daniel Ratner[1]; Xiaobiao Huang[None]; Zhe Zhang[None]

[1] *SLAC*

Corresponding Authors: zhezhang@slac.stanford.edu, dratner@slac.stanford.edu, xiahuang@slac.stanford.edu

This paper presents a novel application of Multipoint Bayesian Algorithmic Execution (multipoint-BAX) to optimize dynamic aperture (DA) and momentum aperture (MA) in lattice design. DAMA optimization is a critical design task for storage rings, ultimately determining the flux of x-ray sources and luminosity of colliders. Traditionally, solving this multi-objective optimization problem has relied on genetic algorithms (GA) and/or Bayesian optimization (BO), requiring extensive particle tracking simulations for each trial configuration, which in turn limits the quality of the final design. Here we instead use multipointBAX to select, simulate, and model a single particle for each trial configuration, resulting in two orders of magnitude higher end-to-end efficiency, while also integrating random magnet error seeds to increase robustness. The multipointBAX implementation involved several advances, including a neural-network surrogate model, a batch acquisition strategy, and the concept of a "Pareto front region" to improve stability. A proof-of-principle demonstration on the SPEAR3 storage ring design validates our approach, with multipointBAX achieving equivalent Pareto front results with only 1% of the tracking computations required by GA. We are now applying the method to the design of future light sources and colliders.

Optimisation and Control / 52

Serval Applications of Machine Learning at HEPS- 15'+5'

Authors: Xiaohan Lu[None]; Yi Jiao[None]; wei bao[None]

Corresponding Authors: baowei@ihep.ac.cn, luxh@ihep.ac.cn, jiaoyi1981@gmail.com

The High Energy Photon Source (HEPS) is a fourth-generation synchrotron radiation facility under construction in Beijing, China. Since the beam commissioning of the storage ring commenced in July 2024, progress has proceeded smoothly, and the first light was achieved in October. During the construction and beam commissioning of HEPS, we explored machine learning to address critical technical challenges. In the offline optimization phase, we applied serval machine learning methods, focusing on data-driven evaluation and optimization of Touschek lifetime. With approximately 95% evaluation accuracy, we reduced computation time by about 90% compared to traditional methods, thereby significantly improving computational efficiency. In the beam commissioning phase, we employed unsupervised learning methods to identify abnormal states in the power supplies, successfully detecting anomalies. Furthermore, we are optimizing beam lifetime by applying machine learning algorithms to accelerator parameters and Beam Loss Monitor (BLM) data in an ongoing study. In recognition of the importance of high-quality data for machine learning applications, we developed the FAIR-Compliant AI-Ready Accelerator Data Platform (FARAD), aimed at generating AI-ready datasets to advance research in the accelerator field. All the beam experiments described

above were conducted on the FARAD platform, greatly enhancing research efficiency. This paper will provide a detailed description of the progress in the research outlined above, along with future plans.

Poster session / 53

Efficient data-driven model predictive control for online accelerator tuning

Authors: Andrea Santamaria Garcia[1]; Annika Eichler[2]; Borja Rodriguez Mateos[3]; Chenran Xu[4]; Christian Hespe[2]; Jan Kaiser[2]; Simon Hirlaender[5]

[1] *University of Liverpool*

[2] *DESY*

[3] *Universitat Politecnica Catalunya (ES)*

[4] *KIT*

[5] *University of Salzburg*

Corresponding Authors: chenran.xu@kit.edu, annika.eichler@desy.de, simon.hirlaender@plus.ac.at, ansantam@liverpool.ac.uk, borja.rodriguez.mateos@cern.ch, jan.kaiser@desy.de, christian.hespe@desy.de

Reinforcement learning (RL) is a promising approach for the online control of complex, real-world systems, with recent success demonstrated in applications such as particle accelerator control. However, model-free RL algorithms often suffer from sample inefficiency, making training infeasible without access to high-fidelity simulations or extensive measurement data. This limitation poses a significant challenge for efficient real-world deployment. In this work, we explore data-driven model-predictive control (MPC) as a solution. Specifically, we employ Gaussian processes (GPs) to model the unknown transition functions in the real-world system, enabling safe exploration in the training process. We apply the GP-MPC framework to the transverse beam tuning task at the ARES accelerator, demonstrating its potential for efficient online training. This study showcases the feasibility of data-driven control strategies for accelerator applications, paving the way for more efficient and effective solutions in real-world scenarios.

Poster session / 54

Challanges in designing an anomaly detection at ISIS

Authors: Kathryn Baker[None]; Mateusz Filip Leputa[None]; Mihnea Romanovschi[1]

Co-author: Pinilla Sanchez Samuel [2]

[1] *UKRI/STFC/ISIS*

[2] *UKRI/STFC/SCD*

Corresponding Authors: k.baker@stfc.ac.uk, mihnea.romanovschi@stfc.ac.uk, samuel.pinilla@stfc.ac.uk, mateusz.leputa@stfc.ac.uk

The controls group at ISIS has been exploring anomaly detection and its associated challenges. This overview highlights the challenges faced, unsuccessful attempts, and lessons learned. Initially, the group implemented a machine learning anomaly detection system on the methane moderator for Target Station 1. The anomaly detection work began before the system upgrade, rendering previous challenges and solutions obsolete. Initially, the focus was on a single temperature channel, with the main challenge being data labelling of normal and anomalous. Solutions included dimensionality reduction, initial classification in lower dimensions, and convolutional neural networks.

Post-upgrade, the machine's behavior changed, introducing new challenges such as analyzing differences between moderator versions, feature engineering for time series, data relabeling, and developing new evaluation methods. The project's scope expanded to detecting anomalies across multiple channels and understanding their interactions, aiding users in identifying root causes rather than just symptoms.

Poster session / 55

Preliminary Findings on the Application of Reinforcement Learning for Real-Time Optimization of Crystal Channelling in the LHC

Author: Andrea Vella[1]

Co-authors: Daniele Mirarchi [2]; Gianluca Valentino [1]; Rongrong Cai ; Stefano Redaelli [2]

[1] *University of Malta (MT)*

[2] *CERN*

Corresponding Authors: rongrong.cai@cern.ch, stefano.redaelli@cern.ch, daniele.mirarchi@cern.ch, andrea.vella@cern.ch, gianluca.valentino@cern.ch

The Large Hadron Collider (LHC) operates with high intensity proton and heavy ion beams that necessitate a robust collimation system to prevent damage to sensitive equipment along the ring. However, the efficiency of cleaning ion beams is approximately 100 times less efficient than with protons. To address this, bent silicon crystals were implemented to enhance collimation efficiency. The first operational use of crystal-assisted collimation took place during the 2023 lead ion run, achieving the required performance improvements to safely manage high-intensity beams. Despite this success, unwanted crystal rotations were observed, leading to suboptimal performance. Current understanding attributes these rotations to mechanical deformations caused by energy deposited due to impedance effects, a problem that cannot be addressed until the next long shutdown. In response, a conventional numerical optimiser was employed to monitor and counteract mechanical deformations using data from a series of beam-loss monitors. This challenge lends itself to the application of reinforcement learning (RL) techniques, which can maintain continuous optimal channelling, reduce convergence time, and eventually support the optimisation of crystals across multiple planes simultaneously.

Optimisation and Control / 56

Optimising Injection Efficiency at Diamond Light Source using Gaussian Processes with Non-Gaussian Likelihoods - 15'+5'

Author: Shaun Preston[1]

Co-authors: Ian Martin [2]; Neven Blaskovic-Kraljevic [2]

[1] *University of Oxford*

[2] *Diamond Light Source*

Corresponding Authors: shaun.preston@physics.ox.ac.uk, neven.blaskovic-kraljevic@diamond.ac.uk, ian.martin@diamond.ac.uk

A general Bayesian optimisation tool is being developed at Diamond Light Source to improve machine performance by constructing surrogates from Gaussian Process (GP) models. Priors are placed on covariance kernel hyperparameters to guide an optimiser and prevent overfitting. The model has been integrated with the machine control system. During an experiment aimed at improving injection efficiency by changing magnet currents, hysteresis was identified and compensated. Multiple measurements were taken at each current to determine the measurement error. However, large

variations in the error with respect to current were observed; this heteroscedastic behaviour was handled robustly with non-Gaussian likelihoods incorporated into the inference step. The injection efficiency was increased running the model in less time than a manual scan from operators.

Poster session / 57

Using Surrogate Models to make BO safer for particle accelerators

Authors: Kathryn Baker[None]; Raunakk Banerjee[None]

Corresponding Authors: raunakk.banerjee@stfc.ac.uk, k.baker@stfc.ac.uk

Tuning of the accelerator at the ISIS Neutron and Muon Source has traditionally been done manually, relying on experts to vary parameters in the control system to achieve optimum beam efficiency and intensity. Bayesian Optimisation (BO) is an effective way to automatically tune the Low Energy Beam Transport (LEBT), a section within the linear accelerator. However, due to losses and radiation caused in the accelerator by the initial exploratory nature of BO, the machine can be put into unsafe operating states. To mitigate these losses, we explore how surrogate models can be used for pre-training, exploring their function as low fidelity information sources to achieve high beam intensity without the noisy initial exploration. In addition, such models can be used to test BO algorithms specifically modified to be safe, before running them on the accelerator.

Poster session / 58

The Progress of Developing a FAIR-Compliant AI-Ready Accelerator Data Platform

Authors: Xiaohan Lu[None]; Yi Jiao[None]; Yuliang Zhang[1]; wei bao[None]

[1] *IHEP*

Corresponding Authors: baowei@ihep.ac.cn, zhangyl@ihep.ac.cn, jiaoyi1981@gmail.com, luxh@ihep.ac.cn

To advance the development of machine learning applications in particle accelerators, a FAIR-compliant AI-Ready Accelerator Data Platform (FARAD) has been proposed. By the end of 2024, our project has made several significant advancements. First, we successfully completed a new data acquisition system that enables higher data sampling rates while ensuring data synchronization, providing more accurate real-time data support for subsequent machine learning models. Second, we developed and refined a parameter scanning system that enables the platform to generate a rich, diverse data set through multidimensional parameter scanning, providing valuable resources for improving the generalization and accuracy of machine learning models. In addition, we established a unified standard for data annotation and publication to ensure consistency and standardization across data collection, processing, annotation, and sharing. These enhancements not only further strengthen data integrity and reliability, but also increase data richness and practical value to better support the intelligent operation, maintenance, and optimization of particle accelerators. These enhancements provide a stronger foundation for machine learning applications in accelerator engineering.

Poster session / 59

Surrogate Models for Longitudinal Beam Manipulations in the CERN Proton Synchrotron

Authors: Alexandre Lasheen[1]; Stamatina Detsi[2]

[1] *CERN*

[2] *National Technical University of Athens, School of Electrical and Computer Engineering*

Corresponding Authors: alexandre.lasheen@cern.ch, stamatina.detsi@cern.ch

The Proton Synchrotron (PS) at CERN is equipped with numerous RF systems allowing for evolved longitudinal beam manipulations to adapt the number of bunches and their spacing. The beam produced for the LHC undergoes several bunch splittings, merging and batch compression. Each manipulation must be carefully adjusted to minimize the spread in bunch parameters at PS extraction. The design of these complex RF processes relies on the non-linearities of the longitudinal equations of motion. This step is usually performed applying numerical optimization with dedicated simulation codes and requires manual processing for deployment in operation. To improve the design of existing manipulations and explore new schemes, surrogate models are proposed in view of improved integration in the accelerator control system. In this work, surrogate models in form of artificial neural networks are built to control the parameters of the RF bucket based on the desired bucket areas and shapes, and generating the expected RF parameters (voltage and phase) to the control system.

Optimisation and Control / 60

Eliminating mains noise with Machine Learning- 15'+5'

Author: Verena Kain[1]

Co-authors: Elias Walter Waagaard [2]; Fabio Follin [1]; Hannes Bartosik [1]; Konstantinos Paraschou [1]; Michael Schenk [1]; Olivier Michels [1]; Pablo Arrutia [1]; Quentin King [1]

[1] *CERN*

[2] *EPFL - Ecole Polytechnique Federale Lausanne (CH)*

Corresponding Authors: hannes.bartosik@cern.ch, michael.schenk@cern.ch, elias.walter.waagaard@cern.ch, verena.kain@cern.ch, fabio.follin@cern.ch, konstantinos.paraschou@cern.ch, quentin.king@cern.ch, pablo.andreas.arrutia.sota@cern.ch, olivier.michels@cern.ch

Power supply ripples at various frequencies - characteristic to the magnet circuits or from the electrical network - have always been an issue in accelerator operations, with several mitigation measures put in place over the years. This contribution summarises recent efforts in the CERN SPS to compensate the ripple at 50 Hz and its harmonics in the main quadrupole circuits with ML methods. It will start with introducing the detrimental effects of the ripple at low energy for LHC-type beams and at top energy for slow extracted beams. For optimal conditions of slow extracted beams, a continuous control algorithm had to be conceived. The implementation required hardware modifications on the power converter electronics side, additional new controls infrastructure and the development of adaptive algorithms that can deal with the changes of the electrical distribution network throughout the day. Adaptive continuous control with adaptive Bayesian Optimisation has been in place for slow extracted spill control throughout 2024. The obtained improved spill quality over the year will be discussed. First impressive results with 50 Hz compensation for the LHC ion cycles in the SPS during the ion run at the end of 2024 will also be presented. Finally, ideas for further R&D towards one-shot corrections for beams that are only played on-demand (i.e. LHC beams) will be proposed.

MLOps, Infrastructure and Scalability / 61

Lume Deployment - standardising model deployment- 15'+5'

Author: Mateusz Leputa[None]

Co-authors: Kathryn Baker ; Mihnea Romanovschi [1]; Raunakk Banerjee

[1] *UKRI/STFC/ISIS*

Corresponding Authors: mateusz.leputa@stfc.ac.uk, mihnea.romanovschi@stfc.ac.uk, raunakk.banerjee@stfc.ac.uk, k.baker@stfc.ac.uk

Ensuring efficient use of resources and longevity of machine learning projects requires careful consideration of the full machine learning lifecycle especially when models are deployed to interact with live control systems or end users. We present *Lume Deployment* a framework of standardised modules built for rapid development and deployment of machine learning models and their integration to the control system. The framework is built around a modular approach which separates the system, the data and the model from each other via well-defined interfaces, which also allows easy expandability of the framework. We showcase several case studies and use-cases that demonstrate the framework's flexibility in various real-world scenarios. Additionally, we outline current developments and possible future developments focusing on the framework's stability and automation model evaluation tasks.

Surrogate Modelling and Digital Twins / 62

Predicting Space-Charge Potentials and E-Fields Using CNN and KAN- 15'+5'

Authors: Emanuele Laface[1]; Isabella Vojskovic[None]

[1] *European Spallation Source ERIC*

Corresponding Authors: emanuele.laface@ess.eu, isabella.vojskovic@fysik.lu.se

This study explores various neural network approaches to simulate beam dynamics, specifically addressing non-linear space charge effects. We introduce a convolutional encoder-decoder architecture with skip connections, achieving a relative error of 0.5% in predicting both transversal and coupled 3D electric self-fields. Additionally, to enhance interpretability and robustness, we investigate an auxiliary Kolmogorov-Arnold Network (KAN) designed to replicate the solution flow of a Finite Element Method (FEM). Our findings suggest that these advancements offer a potentially more efficient alternative to traditional numerical methods for non-linear space-charge calculations in beam dynamics simulations, delivering substantial speed-up.

Poster session / 63

Digital Twin Implementation for FCC-ee Arc Half-Cell: Advancing Mechanical Stability and Predictive Analytics

Author: Lakshit Bahl[1]

Co-author: Oscar Sacristan De Frutos [2]

[1] *Sapienza Universita e INFN, Roma I (IT)*

[2] *CERN*

Corresponding Authors: lakshit.bahl@cern.ch, oscar.sacristan@cern.ch

The Future Circular Collider (FCC) project demands unprecedented mechanical stability, requiring nanometre-scale (10^{-9} meters) precision for beam optics and operation. The Arc Half-Cell, a critical and frequently repeated mechanical unit in the FCC-ee design, must maintain sub-micron alignment despite ground motion and vibrations. To address these challenges, we are in the early phase of developing a Digital Twin (DT) of the Short-Straight Section (SSS) of FCC-ee Arc Half-Cell, enabling real-time simulation, prediction, and optimization of its behaviour within the accelerator proposed infrastructure.

Our approach combines high-fidelity physical modelling, validated through experimental measurements and finite element analysis, with machine learning techniques. By incorporating live data from seismic sensors and vibration analysis, the Digital Twin improves the accuracy of simulations and predictive models, enabling precise monitoring of key mechanical components such as girders and magnets. This infrastructure will facilitate real-time monitoring and prediction of the magnetic center motion in the assembly due to the external excitation sources, a critical factor influencing beam dynamics.
This initiative builds upon previous work on Digital Twins within CERN's Engineering Department's Mechanical and Materials Engineering group (EN-MME), including a proof of principle developed for the crab cavity supporting system in the High-Luminosity Large Hadron Collider (HL-LHC) project. The crab cavities'performance is highly sensitive to mechanical stability; minimal vibrations can induce phase errors and beam instabilities. The experience gained from this project informs our current efforts.
The development of the FCC-ee Arc Half-Cell Digital Twin follows a structured roadmap, starting with a robust data architecture for real-time data acquisition, management, and reliable storage. Ongoing work focuses on integrating real-time data from remote sensors to synchronize the virtual model with the physical system, enhancing the Digital Twin's ability to predict mechanical behaviour. This includes validating FEA simulations by analysing the effects of ground motion and vibrations on the mechanical stability of the arc half-cell. Machine learning techniques and reduced-order modelling are employed to optimize simulation efficiency and improve the predictive accuracy of the Digital Twin.
This endeavour exemplifies a shift towards a comprehensive, data-driven approach to accelerator design and monitoring. By leveraging real-time data integration, predictive analytics, and AI, we address the stringent mechanical stability requirements of future collider projects.
In this presentation, we will discuss the architecture and methodologies used for real-time data acquisition, IoT-based data streaming with Kafka, and the integration of AI for anomaly detection. It will also be addressed the implementation of scalable computational methods that address the FCC-ee's mechanical stability requirements. Our findings demonstrate that the Digital Twin framework can assist in performance monitoring and provide early warnings for potential mechanical instabilities, supporting predictive maintenance. This work lays the groundwork for further integration of Digital Twins into the analysis and validation of mechanical systems, offering insights to improve the operational efficiency of high-energy particle accelerators like the FCC.

Poster session / 64

Enhancements of beam diagnostics and simulation environments toward ML-based XFEL improvements at SACLA

Author: Hirokazu Maesaka[1]

Co-authors: Chikara Kondo [2]; Eito Iwai [2]; Hitoshi Tanaka [1]; Kenji Yasutome [1]; Shinichi Matsubara [2]; Takahiro Inagaki [1]; Takashi Ohshima [2]; Takashi Tanaka [1]; Takato Tomai [2]; Toru Hara [1]

[1] *RIKEN SPring-8 Center*

[2] *Japan Synchrotron Radiation Research Institute*

Corresponding Authors: yasutome@spring8.or.jp, ohshima@spring8.or.jp, tanaka@spring8.or.jp, iwai@spring8.or.jp, toru@spring8.or.jp, tomai@spring8.or.jp, ztanaka@spring8.or.jp, ckondo@spring8.or.jp, matsubara@spring8.or.jp, inagaki@spring8.or.jp, maesaka@spring8.or.jp

An ML-based optimizer has been working on maximizing XFEL performance at SACLA [1]. The spectral brightness of XFEL was successfully optimized by using a new high-resolution inline spectrometer [2]. To improve the XFEL performance further, we plan to enhance beam diagnostics and simulation environments for the ML-based optimizer. As for beam diagnostics, we are developing an X-band RF deflecting structure to measure the longitudinal phase-space profile of an electron beam after the undulator beamline. The temporal resolution is expected to be better than 1 fs. We will be able to observe the lasing part of the electron beam and the precise longitudinal phase-space information will contribute to improving the XFEL performance by using an ML-based optimizer. The simulations of an electron beam and XFEL are also important for the temporal phase-space analysis after lasing. Therefore, we are developing an ecosystem for a smooth connection between the

simulation tools, such as the XFEL simulation code SIMPLEX [3]. In this presentation, we will show the development status and plans for these enhancements.

Poster session / 65

Machine Learning driven beam emittance optimization at EuXFEL

Authors: Alexander Klemps[1]; Denis Ilia[2]; Pradeep Kr. Banerjee[1]

Co-authors: Henrik Tünnermann [2]; Nihat Ay [1]; Ye Chen [2]

[1] *Hamburg University of Technology (TUHH)*

[2] *DESY*

Corresponding Authors: henrik.tuennermann@desy.de, pradeep.banerjee@tuhh.de, alexander.klemps@tuhh.de, denis.ilia@desy.de, nihat.ay@tuhh.de, ye.lining.chen@desy.de

Planned upgrades of the European X-Ray Free Electron Laser (EuXFEL) target higher photon energy and a high duty-cycle operation up to CW-operation using a superconducting RF gun with lower gradient. An operation in this regime though critically depends on improvements of the beam slice emittance of the electron gun. Within the OPAL-FEL project, we are addressing this challenge by developing a data-driven optimization framework for longitudinal drive laser shapes to minimize beam emittance, thereby ensuring the delivery of high-quality electron beams.

Our approach centers on the application of deep learning techniques to create an inverse model that predicts optimal parameter configurations for the photoinjector, enabling targeted control of beam emittance. This methodology involves generating synthetic training data through comprehensive beam dynamics simulations and introduces a machine learning-based strategy for temporal pulse shaping, accommodating a broad family of pulse distributions beyond flattop and Gaussian shapes.
We present results from trained neural networks with various architectures and establish a theoretical foundation for the invertibility of the forward model by connecting our approach to the theory of inverse problems. In particular, we draw on Whitney's embedding theorem within the framework of attractor reconstruction to validate model invertibility.

Leveraging extensive simulations, data-driven modeling and theoretical insights, our approach offers a robust pathway for improved optimization and control of photoinjector parameters in CW mode, with potential applications for further advancements in FEL performance.

Optimisation and Control / 66

Beam halo losses reduction with simulation constrained Bayesian Optimization- 15'+5'

Author: Andrea De Franco[1]

Co-authors: Akihiko Mizuno [1]; Dominique Gex [2]; Fabio Cismondi [2]; Francesco Scantamburlo [2]; Herve Dzitko [2]; Jibong Hyun [1]; Kai Masuda [1]; Keitaro Kondo [1]; Kouki Hirosawa [1]; Tomonobu Itagaki [1]; Tomoya Akagi [1]

[1] *National Institutes for Quantum Science and Technology (QST)*

[2] *F4E*

Corresponding Author: defranco.andrea@qst.go.jp

The Linear IFMIF Prototype Accelerator (LIPAc) is designed to accelerate 125 mA of D+ to 9 MeV in CW. The very high power stored in the beam (˜1.1 MW) and the use of superconductive RF cavities requires precise control of beam losses (target <1e-6). On the other hand the intense beam is affected by strong space charge forces that easily results in significant halo formation. This contribute is difficult to simulate because requires large number of particle tracking and precise knowledge of input phase space. In this work we present the use of Bayesian optimization of transport optic (4 quadrupoles + 4 steerers) to minimize the halo losses by observing vacuum response in the sector. As a single beam pulse with inappropriate optics could lead to permanent damage of the machine, we include constraint of negligible losses from beam core in simulations. The algorithm successfully and safely reduced vacuum in sector by a factor of ˜3 within two hours of operation. After analysis of newly set-point we find that the proposed optics results both in a compromise matching of core and halo distribution and a reduction of particles ejection from the former to the latter.

Anomaly Detection and Diagnostics / 67

Measurement of CSR effects using generative phase space reconstruction-15'+5'

Author: Juan Pablo Gonzalez-Aguilera[1]

Co-authors: Ryan Roussel [2]; Auralee Edelen [2]; Young-Kee Kim [1]

[1] *University of Chicago*

[2] *SLAC National Accelerator Laboratory*

Corresponding Authors: jpga@uchicago.edu, ykkim@hep.uchicago.edu, edelen@slac.stanford.edu, rroussel@slac.stanford.edu

Coherent synchrotron radiation (CSR) is a limiting effect in linear accelerators with dispersive elements due to its contribution to projected transverse emittance growth. This effect becomes a limitation for highly compressed beams. Even though CSR-induced projected emittance growth has been widely studied, conventional measurement techniques are not detailed enough to resolve the multi-dimensional structure of the beam, namely the different translations and rotations of transverse phase space slices throughout the longitudinal coordinate. In this work, we use a state-of-the-art method to reconstruct the phase space of a beam affected by CSR at the Argonne Wakefield Accelerator Facility. This detailed, efficient and multi-dimensional phase space reconstruction method enables better understanding of the CSR effects in a double dogleg where shielding is limited.

Poster session / 68

Machine Learning based Optical Distortion measurement by Phase Advance in SSRF

Author: liu xinzhong[None]

Co-authors: Liyuan Tan ; Yihao GONG [1]

[1] *Shanghai Advanced Research Institute*

Corresponding Authors: liuxinzhong@sari.ac.cn, tanliyuan@sinap.ac.cn, gongyh@sari.ac.cn

Optical distortion measurement and correction are pivotal for the stable operation of accelerators. This study introduces a machine learning-based approach to optical distortion measurement and correction implemented on the Shanghai Synchrotron Radiation Facility (SSRF). We trained models

from modulated orbits to phase advance and from phase advance to quadrupole models, establishing a data-driven method for optical distortion correction. A key advantage of this method is the elimination of the need to update the Jacobian matrix, thereby reducing computational iteration time. Our machine learning-driven optical distortion correction method completes the process in 2 minutes, significantly faster than traditional Local Closed Orbit Correction (LOCO) methods. Furthermore, our method, based on steady-state orbit data measurement with orbit modulation, achieves higher accuracy and requires lower measurement conditions compared to turn by turn (TBT) data measurement-based optical distortion correction. This research underscores the potential of machine learning in the field of accelerator optical distortion correction and introduces an efficient strategy for future accelerator operations.

Poster session / 69

Beam based alignment using a neural network

Author: Wei Xu[1]

[1] *University of Science and Technology of China*

Corresponding Author: wxu@ustc.edu.cn

Beams typically do not travel through the magnet centers because of errors in storage rings. The beam deviating from the quadrupole centers is affected by additional dipole fields due to magnetic field feed-down. Beam-based alignment (BBA) is often performed to determine a golden orbit where the beam circulates around the quadrupole center axes. For storage
rings with many quadrupoles, the conventional BBA procedure is timeconsuming, particularly in the commissioning phase, because of the necessary iterative process. In addition, the conventional BBA method can be affected by strong coupling and the nonlinearity of the storage ring optics. In this study, a novel method based on a neural network was proposed to determine the golden orbit in a much shorter time with reasonable accuracy. This golden orbit can be used directly for operation or adopted as a starting point for conventional BBA. The method was demonstrated in the HLS-II storage ring for the first time through simulations and online experiments. The results of the experiments showed that the golden orbit obtained using this new method was consistent with that obtained using the conventional BBA. The development of this new method and the corresponding experiments are reported in this poster.

Poster session / 70

Symplectic machine learning model for space charge effect simulation

Author: Jinyu Wan[None]

Co-authors: Yue Hao ; Ji Qiang

Corresponding Authors: jqiang@lbl.gov, haoyue@msu.edu, wanjinyu@msu.edu

We propose to use a generative adversarial network (GAN) based on a U-Net architecture to learn the space-charge Hamiltonian. By differentiating the predicted Hamiltonian, we obtain the one-step symplectic transverse transfer map of the particles. This model can preserve the symplecticity in the simulation and reduce non-physical errors for long-term particle tracking.

Optimisation and Control / 72

Data-Driven Feedback Optimization for Particle Accelerator Control-15'+5'

Authors: Christian Hespe[1]; Jan Kaiser[None]; Jannis Lübsen[2]; Annika Eichler[1]

[1] *DESY*

[2] *TUHH*

Corresponding Authors: annika.eichler@desy.de, jan.kaiser@desy.de, christian.hespe@desy.de, jannis.luebsen@tuhh.de

Feedback control is an essential component for the successful operation of particle accelerators. However, achieving the desired closed-loop performance requires precise model knowledge, which is difficult to obtain in complex accelerator systems. For this reason, we present an application of a combined optimization approach that estimates the response matrix online while optimizing the chosen performance measure, eliminating the need for first-principles modelling or a priori identification experiments. Specifically, a Kalman filter is employed to construct a linearization of the system response around its operating point from noisy input-output measurements, iteratively improving the available knowledge about the system. In parallel, this knowledge is exploited by a feedback optimizer, which is incrementally driving the system to its optimal operating point while maintaining safety of operation as formulated by constraints on inputs and outputs. As a consequence of the continuous online response estimation and in contrast to other modelling approaches, the scheme is able to instantaneously react to changes and drifts in the system behaviour. This is demonstrated on the orbit feedback of the main electron beam dump line of the European XFEL.

Poster session / 73

Exploring Reinforcement Learning for Optimal Bunch Merge in the AGS

Author: Yuan Gao[None]

Co-authors: Armen Kasparian [1]; Daria Kuzovkova [2]; Eiad Hamwi ; Freddy F Severino [3]; Georg Hoffstaetter ; John Morris [3]; Jonathan Unger [2]; Keith Zeno [3]; Kevin Brown [4]; Levente Hajdu [3]; Malachi Schram ; Matthew Signorelli [2]; Vincent Schoefer ; Weijian Lin [4]; Xiaofeng Gu [3]; Yinan Wang [5]

[1] *JLab*

[2] *Cornell Univ.*

[3] *Brookhaven National Lab*

[4] *Brookhaven National Laboratory*

[5] *RPI*

Corresponding Authors: mgs255@cornell.edu, brownk@bnl.gov, eh652@cornell.edu, jeu8@cornell.edu, jtm@bnl.gov, xgu@bnl.gov, wangy88@rpi.edu, armenk@jlab.org, dk638@cornell.edu, severino@bnl.gov, wlin1@bnl.gov, georg.hoffstaetter@cornell.edu, lbhajdu@bnl.gov, zeno@bnl.gov, ygao@bnl.gov, schram@jlab.org, vincent.schoefer@gmail.com

In BNL's Booster, the beam bunches can be split into two or three smaller bunches to reduce their space-charge forces. They are then merged back after acceleration in the Alternating Gradient Synchrotron (AGS). This acceleration with decreased space-charge forces can reduce the final emittance, increasing the luminosity in RHIC and improving proton polarization. Parts of this procedure have already been tested and are proposed for the Electron-Ion Collider (EIC). The success of this procedure relies on a series of RF gymnastics to merge individual source pulses into bunches of suitable intensity. In this work, we explore an RF control scheme using reinforcement learning (RL) to merge bunches, aiming to dynamically adjust RF parameters to achieve minimal longitudinal emittance growth and stable bunch profiles. Initial experimental results and ongoing system developments are presented and discussed.

Poster session / 74

Improve Beam Brightness with Bayesian Optimization at the AGS Booster Injection at BNL

Author: Weijian Lin[1]

Co-authors: Eiad Hamwi ; Georg Hoffstaetter ; Kevin Brown [1]; Levente Hajdu [1]; Petra Adams [1]; Vincent Schoefer [1]; Xiaofeng Gu [2]

[1] *Brookhaven National Laboratory*

[2] *Brookhaven National Lab*

Corresponding Authors: padams@bnl.gov, eh652@cornell.edu, brownk@bnl.gov, vincent.schoefer@gmail.com, wlin1@bnl.gov, xgu@bnl.gov, lbhajdu@bnl.gov, georg.hoffstaetter@cornell.edu

Alternating Gradient Synchrotron (AGS) and its Booster serve as part of the injector compound for RHIC and the future EIC at Brookhaven National Laboratory. Injection and early acceleration processes set maximum beam brightness for the collider rings. Such processes have many control parameters and are traditionally optimized empirically by operators. In an effort to streamline the injection processes with machine learning (ML) techniques, we develop and test a Bayesian Optimization (BO) algorithm to automatically tune the Linac to Booster (LtB) transfer line magnets to maximize beam brightness after injection into the Booster. We present experimental results that demonstrate BO can be applied to optimize Booster injection efficiency. Beam studies also indicate transverse coupling in LtB, which has been difficult to quantify due to instrument limitations. We plan to develop further studies to investigate the coupling effect.

Optimisation and Control / 75

Explainable physics-based constraints on reinforcement learning for accelerator controls- 15'+5'

Authors: Jonathan Colen[1]; Malachi Schram[None]; Kishansingh Rajput[2]; Armen Kasparian[3]

[1] *Old Dominion University*

[2] *Jefferson Lab*

[3] *JLab*

Corresponding Authors: armenk@jlab.org, schram@jlab.org, jcolen@odu.edu, kishan@jlab.org

We present a reinforcement learning (RL) framework for controlling particle accelerated experiments that builds explainable physics-based constraints on agent behavior. The goal is to increase transparency and trust by letting users verify that the agent's decision-making process incorporates suitable physics. Our algorithm uses a learnable surrogate function for physical observables, such as energy, and uses them to fine-tune how actions are chosen. This surrogate can be represented by a neural network or by a sparse dictionary model. We test our algorithm on a range of particle accelerator controls environments designed to emulate the Continuous Electron Beam Accelerator Facility (CEBAF) at Jefferson Lab. By examining the mathematical form of the learned constraint function, we are able to confirm the agent has learned to use the established physics of each environment. In addition, we find that the introduction of a physics-based surrogate enables our reinforcement learning algorithms to reliably converge for difficult high-dimensional accelerator controls environments.

Optimisation and Control / 76

Harnessing the Power of Gradient-Based Simulations for Multi-Objective Optimization in Particle Accelerators- 15'+5'

Author: Kishansingh Rajput[1]

Co-authors: Adam Carpenter [1]; Armen Kasparian [2]; Auralee Edelen ; He Zhang [1]; Jay Benesch [3]; Jonathan Colen [4]; Malachi Schram ; Ryan Roussel

[1] *Jefferson Lab*

[2] *JLab*

[3] *Thomas Jefferson National Accelerator Facility*

[4] *Old Dominion University*

Corresponding Authors: edelen@slac.stanford.edu, benesch@jlab.org, schram@jlab.org, rroussel@slac.stanford.edu, armenk@jlab.org, kishan@jlab.org, jcolen@odu.edu

Particle accelerator operation requires simultaneous optimization of multiple objectives. Multi-Objective Optimization (MOO) is particularly challenging due to trade-offs between the objectives. Evolutionary algorithms, such as genetic algorithm (GA), have been leveraged for many optimization problems, however, they do not apply to complex control problems by design. This paper demonstrates the power of differentiability for solving MOO problems using a Deep Differentiable Reinforcement Learning (DDRL) algorithm in particle accelerators. We compare DDRL algorithm with Model Free Reinforcement Learning (MFRL), GA and Bayesian Optimization (BO) for simultaneous optimization of heat load and trip rates in the Continuous Electron Beam Accelerator Facility (CEBAF). The underlying problem enforces strict constraints on both individual states and actions as well as cumulative (global) constraint for energy requirements of the beam. A physics-based surrogate model based on real data is developed. This surrogate model is differentiable and allows back-propagation of gradients. The results are evaluated in the form of a Pareto-front for two objectives. We show that the DDRL outperforms MFRL, BO, and GA on high dimensional problems.

Poster session / **77**

Integrating Machine Learning in the Brookhaven Control System with Badger

Author: Levente Hajdu[1]

Co-authors: Eiad Hamwi [2]; Georg Hoffstaetter [2]; Kevin Brown [3]; Ryan Roussel [4]; Weijian Lin [3]; Zhe Zhang

[1] *Brookhaven National Laboratory - C-AD Controls*

[2] *Cornell University*

[3] *Brookhaven National Laboratory*

[4] *SLAC*

Corresponding Authors: lbhajdu@bnl.gov, zhezhang@slac.stanford.edu, eh652@cornell.edu, brownk@bnl.gov, wlin1@bnl.gov, rroussel@slac.stanford.edu, georg.hoffstaetter@cornell.edu

The Brookhaven Pre-injector Accelerator Facility, which serves RHIC, NSRL, BLIP, and the future EIC, requires occasional tuning of its transfer beam line optics by control room operators to optimize parameters like beam current and emittance. Machine learning (ML) can significantly speed up this tuning process by helping operators quickly identify optimal settings. To facilitate this, ML algorithms must be seamlessly integrated into the control system and accessed through a consistent, familiar interface. We explore this approach using the Badger software stack, which provides a user interface built atop Xopt, a comprehensive package of advanced ML optimization algorithms. This paper presents our experience in developing Badger plug-ins for transfer beam line optimization and

interfacing with multinet, a non-EPICS control system, to streamline tuning and enhance operational efficiency.

Optimisation and Control / 78

Virtual to Physical: Reinforcement Learning to Optimize SNS Particle Accelerator Controls- 15'+5'

Author: Armen Kasparian[1]

Co-authors: Alexander P. Zhukov [2]; Brandon Cathey [3]; Carrie Elliott [2]; Jonathan Colen [4]; Kishansingh Rajput [1]; Malachi Schram ; Trent Thompson [2]; Willem Blokland [2]

[1] *Jefferson Lab*

[2] *Oak Ridge National Lab*

[3] *Spallation Neutron Source*

[4] *Old Dominion University*

Corresponding Authors: schram@jlab.org, kishan@jlab.org, armenk@jlab.org, thompsonte@ornl.gov, catheybl@ornl.gov, blokland@ornl.gov, elliottcl@ornl.gov, jcolen@odu.edu, zhukovap@ornl.gov

Complex accelerators must have control systems that can handle dynamic nonlinear environments. This makes traditional control methods unsuitable as they can struggle to adapt to these uncertainties. This provides an ideal environment for reinforcement learning algorithms as they are adaptable and generalizable. We present a reinforcement learning pipeline that can effectively handle the dynamics of a complex accelerator. We test and prove our pipelines capabilities on multiple environments including the Spallation Neutron Source (SNS) and the Beam Test Facility (BTF) at Oakridge National Lab (ORNL). Due to the limited time available to train an online algorithm like reinforcement learning on a real accelerator, we utilize a virtual twin accelerator (VIRAC) developed by ORNL to pretrain the policy and show its ability to converge in the virtual environment. We then test the adaptability of the pretrained RL model by applying it on the real accelerator and comparing the results. Utilizing our Scientific Optimization and Controls Toolkit (SOCT) and open-source standards such as Gymnasium we create and solve for a MEBT orbit correction problem in the SNS and an emittance maximization problem in the BTF. We show how Twin Delayed Deep Deterministic Policy Gradient (TD3) can solve this optimization environment in the virtual accelerator and transfer this policy onto the real accelerator for inference and model retraining. We show how reinforcement learning can be utilized as a control system for complex accelerators and provide a model pipeline for how an implementation performs and can be adapted to new accelerator control problems.

Surrogate Modelling and Digital Twins / 79

ML-Based Multi-Fidelity Model Calibration Toward Precision Control of Electron Beams- 15'+5'

Author: Eric Cropp[None]

Co-authors: Ryan Roussel ; Remi Lehe ; Auralee Edelen

Corresponding Authors: rlehe@lbl.gov, rroussel@slac.stanford.edu, edelen@slac.stanford.edu, cropp@slac.stanford.edu

Precision control of electron beams is one of the main charges of beam physics, as producing high-brightness beams is critical to numerous accelerator deliverables, including high-quality x-rays from XFELs and high-quality ultrafast probes for UED/UEM. Critical to this effort is a set of accurate system models that can inform control policies. To be useful, these models must accurately reflect the behavior of the accelerator. In this work, a systematic, ML-based approach toward this model calibration problem is outlined. We use ML-based, time-efficient approaches, such as multi-fidelity

Bayesian optimization, to balance the flow of information from high- and low-fidelity models. Additionally, the application of this work to online digital twins toward higher brightness beams will be discussed.

MLOps, Infrastructure and Scalability / 80

Software Infrastructure Plans for End-to-End and Bottom-Up AI/ML Capabilities at the Electron-Ion Collider (EIC)- 15'+5'

Author: Linh Nguyen[None]

Co-authors: Georg Hoffstaetter ; James Jamilkowski [1]; Kevin Brown ; Kyle Kulmatycski [1]; Seth Nemesure ; Weijian Lin [1]; Yuan Gao

[1] *Brookhaven National Laboratory*

Corresponding Authors: ygao@bnl.gov, lnguyen@bnl.gov, jpjamilk@bnl.gov, wlin1@bnl.gov, seth@bnl.gov, kbrown@bnl.gov, georg.hoffstaetter@cornell.edu, kkulmatycski@bnl.gov

Plans for the Electron-Ion Collider (EIC), to be built at Brookhaven National Laboratory, include end-to-end and bottom-up capabilities in artificial intelligence (AI) and machine learning (ML). Enabling these capabilities, especially for EIC Operations, will require the large-scale integration of software platforms and tools for the reliable and efficient management of AI/ML-related data, workflows, and processes. We present on our software infrastructure plans.

Anomaly Detection and Diagnostics / 81

Enhancing Quench Detection in SRF Cavities at the European XFEL: Machine Learning Approaches and Practical Challenges- 15'+5'

Authors: Lynda Boukela[1]; Annika Eichler[1]; Julien Branlard[2]; Burak Dursun[1]; Nadeem Shehzad[1]

[1] *DESY*

[2] *Deutsches Elektronen-Synchrotron DESY*

Corresponding Authors: annika.eichler@desy.de, nadeem.shehzad@desy.de, julien.branlard@desy.de, burak.dursun@desy.de, lynda.boukela@desy.de

At the European XFEL, detecting anomalies in superconducting cavities is essential for reliable accelerator performance. We began with a model-based fault detection approach focused on residual analysis to identify anomalies. To improve fault discrimination, particularly for quench events, we augmented this system with machine learning (ML) models. Key challenges included the scarcity of labeled data, which we addressed by integrating expert feedback through an optimized labeling process, and the transition to real-time operation, requiring computational and integration adjustments. For the online application, we deployed two servers in the tunnel at one of the 25 stations to detect failures in real-time with a software-based solution. In parallel we pushed the development of an FPGA-based solution, that will allow to counteract on real-time in the future. The resulting detection system delivers reports across various timescales, supporting both immediate responses and long-term maintenance. It will provide new insights to the online data, which was never explored in the past.

Poster session / 82

Machine Learning for Quench Detection in SRF Cavities at the European XFEL

Authors: Lynda Boukela[None]; Annika Eichler[1]; Julien Branlard[2]

[1] *DESY*

[2] *Deutsches Elektronen-Synchrotron DESY*

Corresponding Authors: lynda.boukela@desy.de, annika.eichler@desy.de, julien.branlard@desy.de

Superconducting Radio Frequency (SRF) cavities are essential components in particle accelerators, where quenches, which can cause an abrupt loss of superconductivity, remain a significant challenge. While quench detection has traditionally relied on single-parameter analysis, these methods are often limited in terms of robustness and scalability. In this work, we explore the use of machine learning (ML) to enhance the detection of quenching events at the European XFEL. We employ a variety of ML techniques, including different neural network architectures. This analysis augments the physical model explored to obtain representative calculations from the linac operation. Initial experiments show strong performance in terms of detection accuracy and robustness. We also discuss the challenges inherent in applying ML to SRF data, including the need for expert feedback, which requires a streamlined process. Achieving consistency between offline implementations and various online calculations could also be challenging, primarily due to numerical differences.

Poster session / 83

Machine Learning Applications for CLARA

Author: Amelia Pollard[None]

Co-author: David Dunning

Corresponding Authors: amelia.pollard@stfc.ac.uk, david.dunning@stfc.ac.uk

CLARA is an electron accelerator test facility at STFC's Daresbury Laboratory in the UK. The first phase has been operated since 2018 for a wide range of accelerator applications. A major upgrade is presently being commissioned that increases the electron beam energy from ~35 MeV to 250 MeV, and adds an experimental area featuring a new high-power laser. Machine learning will play an important role in the future development of the facility, with aims to deliver bespoke beam properties rapidly, to detect and diagnose anomalies, and to provide virtual diagnostics. This paper summarises machine learning applications under development, including photo-injector laser pulse shaping and a model to learn magnet hysteresis conditions to reduce the need for degaussing.

Poster session / 84

Large language Models for Human-Machine Collaborative Particle Accelerator Tuning through Natural Language

Author: Jan Kaiser[None]

Co-authors: Anne Lasucher [1]; Annika Eichler [2]

[1] *Universität Hamburg*

[2] *DESY*

Corresponding Authors: anne.lauscher@uni-hamburg.de, jan.kaiser@desy.de, annika.eichler@desy.de

Autonomous tuning of particle accelerators is an active and challenging research field with the goal of enabling advanced accelerator technologies and cutting-edge high-impact applications, such as physics discovery, cancer research and material sciences. A key challenge with autonomous accelerator tuning remains that the most capable algorithms require experts in optimisation and machine learning to implement them for every new tuning task. Here, we propose the use of large language models (LLMs) to tune particle accelerators. We demonstrate on a proof-of-principle example the ability of LLMs to tune an accelerator subsystem based on only a natural language prompt from the operator, and compare their performance to state-of-the-art optimisation algorithms, such as Bayesian optimisation and reinforcement learning-trained optimisation. As part of our study, we investigate how to correctly prompt LLMs, evaluating different prompts, where the task is phrased as both an accelerator tuning task and an application-agnostic optimisation task. In doing so, we also show how LLMs can perform numerical optimisation of a non-linear real-world objective. Considering the high computational costs incurred by LLMs, we further evaluate the environment and monetary impact that using them for accelerator tuning would have. Ultimately, this work represents another complex task that LLMs can solve and promises to help accelerate the deployment of autonomous tuning algorithms to day-to-day particle accelerator operations.

Poster session / 85

Towards Lattice-Agnostic Reinforcement Learning Agents for Transverse Beam Tuning

Authors: Chenran Xu[None]; Jan Kaiser[None]; Annika Eichler[1]; Andrea Santamaria Garcia[2]

[1] *DESY*

[2] *University of Liverpool*

Corresponding Authors: annika.eichler@desy.de, jan.kaiser@desy.de, chenran.xu@kit.edu, ansantam@liverpool.ac.uk

Reinforcement learning (RL) has been successfully applied to various online tuning tasks, often outperforming traditional optimization methods. However, model-free RL algorithms typically require a high number of samples, with training processes often involving millions of interactions. As this time-consuming process needs to be repeated to train RL-based controllers for each new task, it poses a significant barrier to their broader application in online tuning tasks. In this work, we address this challenge by extending domain randomization to train general lattice-agnostic policies. We focus on a common task in linear accelerators: tuning the transverse positions and sizes of electron bunches by controlling the strengths of quadrupole and corrector magnets. During training, the agent interacts with environments where the magnet positions are randomized, enhancing the robustness of the trained policy. Preliminary results demonstrate that this approach enables policies to generalize and solve the task on different lattice sections without the need for additional training, indicating the potential for developing transferrable RL agents. This study represents an initial step toward rapid RL deployment and the creation of lattice-agnostic RL controllers for accelerator systems.

Keynotes, tutorials and lectures / 86

Differentiable Simulations

Authors: Jan Kaiser[None]; Chenran Xu[None]; Andrea Santamaria Garcia[1]; Juan Pablo Gonzalez Aguilera[2]

[1] *University of Liverpool*

[2] *University of Chicago*

Corresponding Authors: chenran.xu@kit.edu, ansantam@liverpool.ac.uk, jpga@uchicago.edu, jan.kaiser@desy.de

Many accelerator physics problems, such as beamline design, beam dynamics model calibration, online tuning and phase space measurements rely on solving high-dimensional optimisation problems over beam dynamics simulations. Numerical optimisers have successfully been applied to such tasks,

but they struggle as the dimensionality and complexity of the objective function increase. In machine learning, gradient-based optimisation algorithms are successfully used to optimise billions of model parameters over complex loss functions when training large neural network models. This is made possible by reverse-mode automatic differentiation, which enables the fast computation of gradients of complex functions. In this tutorial, you will learn to use novel beam dynamics simulations with support for automatic differentiation to your advantage and harness the power of gradient-based optimisation in accelerator physics. Multiple hands-on examples using the Cheetah beam dynamics code will allow you to try these methods for yourself. While we will present multiple example applications of gradient-based optimisation on differentiable beam dynamics simulators, the space of potential applications here is vast, and we believe that participants will go on to discover numerous novel applications for differentiable beam dynamics simulations that were intractable to solve with existing methods.

Keynotes, tutorials and lectures / 87

Introduction to Control Theory

Author: Annika Eichler[1]

[1] *DESY*

Corresponding Author: annika.eichler@desy.de

Control theory is a pivotal field of study that focuses on the behavior of dynamical systems and the development of strategies to influence these systems towards desired outcomes. The principle of control theory find its application in plenty of disciplines including engineering, economics, biology and beyond. It were control concepts like the Kalman filter that has flew the Apollo to the moon. Also for the operation of particle accelerator there are thousands of control loops running, ranging from low-level hardware-focused systems, as magnet powering, synchronization or RF control, over fast transverse feedback systems to beam steering.
In this tutorial, we begin by exploring the basic concepts of control, emphasizing the importance of stability and robustness. The tutorial then introduces the mathematical modeling of systems using differential equations and the representation of these systems in state-space form. Starting from classical control strategies, we will go towards modern optimal control techniques and will finally see the link to optimization and ML approaches like reinforcement learning.

Poster session / 88

Analyzing the Impact of Seismic Activities on EuXFEL's Optical Synchronization System Using Data Mining

Author: Arne Grünhagen[None]

Co-author: Annika Eichler [1]

[1] *DESY*

Corresponding Authors: arne.gruenhagen@desy.de, annika.eichler@desy.de

In pump-probe experiments with free-electron lasers, the arrival time stability between the FEL pump pulse and the probe pulse is of utmost importance. An optical synchronization system is used to synchronize several components of the accelerator and the pump-probe laser. Different seismic activities cause the tunnel length and thus the length of the installed optical fibers to change. In this study, we investigate the impact of such seismic activities, more precisely earthquakes, ocean generated microseism, and civilization noise, on the optical synchronization system. To this end, we analyze the controller Input/Output of phase locked loops in length-stabilized fibers links. By comparing the controller data with external data, we were able to identify specific disturbances and their

effects on the control signals. Our results show that even earthquakes that are approximately 5000 km away cause remarkable fluctuations in the in-loop control signals. Ocean-generated microseism in particular has an enormous influence on the in-loop control signals due to its constant presence. The optical synchronization system is so highly sensitive that it can even identify vibrations caused by civilization, such as road traffic or major events like concerts or sport events. The phase-locked loops manage to eliminate more than 99 of the existing interference.

Poster session / 89

Toward Autonomous Control: Reinforcement Learning for Improving accelerator performance in CLEAR

Author: Antonio Gilardi[None]

Co-authors: Alfred Olof Petersson [1]; Andrea Pollastro [2]; Borja Rodriguez Mateos [3]; Levente Alex Foldesi [4]; Michael Schenk [1]; Roberto Corsini [1]; Verena Kain [1]

[1] *CERN*

[2] *University of Naples Federico II, Naples, Italy*

[3] *Universitat Politecnica Catalunya (ES)*

[4] *University of Zurich (CH)*

Corresponding Authors: alfred.olof.petersson@cern.ch, verena.kain@cern.ch, roberto.corsini@cern.ch, antonio.gilardi@cern.ch, levente.foldesi@cern.ch, borja.rodriguez.mateos@cern.ch, andrea.polastro@unina.it, michael.schenk@cern.ch

Particle accelerators, such as the CERN Linear Electron Accelerator for Research (CLEAR), play a critical role in various scientific fields.
Ensuring their operation is automatic, stable, and reproducible is vital for the scalability of future large-scale accelerator projects.
This paper presents an initial step toward autonomous control of the CLEAR beamline, beginning with a basic beam steering challenge and progressing to more complex issues, such as absolute alignment within quadrupoles, which are critical for CLEAR's operational stability.
The proposed solution leverages Reinforcement Learning (RL) agents that learn in real-time using beam screens.
These strategy was used to optimize the sampling efficiency taking into consideration the highly invasive and expensive nature of data collection in particle accelerator environments.
The goal is to achieve single-shot optimization that can be directly applied in real operational scenarios, potentially eliminating the need for further manual adjustments.

The results are highly promising, demonstrating that, with only a few hours of training, it is possible to achieve single-step corrections of the CLEAR beamline's experimental section.
This success has motivated the operational team to further explore and develop this approach.

Poster session / 90

Machine Learning Assisted Bayesian Calibration of Accelerator Digital Twin from Orbit Response Data

Author: Weijian Lin[1]

Co-authors: Christopher Kelly [1]; Georg Hoffstaetter ; Kevin Brown [1]; Nathan Urban [1]

[1] *Brookhaven National Laboratory*

Corresponding Authors: georg.hoffstaetter@cornell.edu, brownk@bnl.gov, wlin1@bnl.gov, ckelly@bnl.gov, nurban@bnl.gov

Digital twins of particle accelerators are used to plan and control operations and design data collection campaigns. However, a digital twin relies on parameters that are hard to measure directly, e.g., magnet alignments, power supply transfer functions, magnet nonlinearities, and stray fields. These parameters can be constrained by beam position and profile measurements. We use Bayesian statistical inference to estimate the parameters, and their uncertainties, probabilistically by calibrating the Bmad digital twin to beam measurements. The inference is computationally accelerated using a machine learning emulator of the physical accelerator digital twin trained to a perturbed-parameter ensemble of Bmad simulations. The result is a joint posterior distribution over parameters (control currents, individual magnet transfer function coefficients, and beam monitor errors) which is propagated to uncertainties in predicted beam positions and profiles, which we validate against beam responses measured at the AGS booster at Brookhaven National Laboratory.

Poster session / 91

eLog Analysis for Accelerators: Status and Future Outlook

Author: Antonin Sulc[1]

Co-authors: Raimund Kammering ; Thorsten Hellert

[1] *DESY*

Corresponding Authors: raimund.kammering@desy.de, thellert@lbl.gov, antonin.sulc@desy.de

This work presents a systematic analysis of electronic logbook (eLog) systems and their analytical capabilities at the accelerator facilities of DESY and Lawrence Berkeley National Laboratory (LBNL). We evaluate contemporary tools and methodologies for enhanced information retrieval, focusing on extracting operational insights from eLog entries through state-of-the-art natural language processing approaches and the integration of structured data with existing accelerator control systems.

The study examines current challenges in eLog analysis through practical implementations at both institutes, demonstrating applications and limitations while proposing architectural modifications to prepare facilities for seamless integration of eLog systems into modern AI-driven operational pipelines. We investigate recent advancements in data accessibility and knowledge extraction methodologies at these facilities, systematically identifying areas that require further enhancement.

Based on our findings, we outline potential developments in analytical capabilities, providing a technical framework for future eLog system improvements and integration strategies.

Poster session / 92

Automated Visual Recognition of Accelerator Components

Author: Thorsten Hellert[None]

Co-author: Antonin Sulc [1]

[1] *DESY*

Corresponding Authors: antonin.sulc@desy.de, thellert@lbl.gov

This study presents a novel application of computer vision techniques for identifying magnets and magnet components in crowded accelerator tunnels at ALS. Utilizing the SAM2/YOLO tracking, we trained a system on a combination of CAD renderings and real photographs of components from

the Advanced Light Source Upgrade (ALSU) project. Our methodology involved creating a comprehensive dataset of a few manually labeled images, derived from CAD models and on-site manually labeled photographs. We show potential applications in streamlining maintenance procedures, enhancing safety protocols, and improving overall accelerator operations towards automated visual inspection and inventory management in complex accelerator environments.

Poster session / **93**

Enhancing ALS Injector Performance Through Data Analysis

Author: Thorsten Hellert[None]

Co-author: Antonin Sulc [1]

[1] *DESY*

Corresponding Authors: thellert@lbl.gov, antonin.sulc@desy.de

This study presents a data-driven methodology aimed at enhancing the performance and reliability of the injector at ALS.
We show a data acquisition system for capturing and analyzing the parameters affecting the injection process to find patterns and improve reliability. We analyze the recorded injection parameters to find key correlations and patterns within the multidimensional parameter space, gaining insights into injector dynamics and potential areas for optimizing the injection process.

Furthermore, we develop a parametric digital twin of the ALS injector based on the recorded data to enable more precise predictions of injector behavior, facilitate rapid troubleshooting, and support the development of advanced control strategies.

Poster session / **95**

Lattice optimization using ML-assisted MOGA for ultralow- emittance storage ring

Author: Saroj Jena[None]

Co-authors: Marco Apollonio [1]; Pedro Fernandes Tavares [2]; Åke Andersson [2]

[1] *MAX IV Laboratory, Lund University, Sweden*

[2] *Lund University*

Corresponding Authors: saroj.jena@maxiv.lu.se, pedro.fernandes_tavares@maxiv.lu.se, ake.andersson@maxlab.lu.se, marco.apollonio@maxiv.lu.se

The Multi-Objective Genetic Algorithm (MOGA) is a powerful and increasingly adopted method for optimizing both linear and nonlinear beam dynamics in accelerator lattices, particularly for ultralow-emittance storage rings. Key objectives in this optimization include minimizing beam emittance for high brightness, maximizing dynamic aperture to ensure efficient particle injection, and expanding momentum aperture to enhance beam lifetime. However, these objectives are constrained by various strict lattice parameters, and magnet strengths which must be satisfied for a physically viable machine design. Integrating Machine Learning (ML) techniques with MOGA offers a promising solution to accelerate the optimization. By training ML models on the objective values that meet predefined constraints throughout MOGA's evolution, we can predict optimal variable sets more efficiently. This ML-assisted MOGA approach not only speeds up the optimization process but also improves convergence towards feasible solutions. Our current work applies this combined ML and MOGA optimization approach to MAX 4U, an upgrade of the MAX IV 3 GeV ring, a state-of-the-art ultralow-emittance storage ring, with the goal of achieving sub-100 pm·rad emittance.

Poster session / 96

Machine Learning for Calibration Drift Forecasting in Superconducting RF Cavities

Author: Yue Sun[1]

Co-authors: Andrea Bellandi [1]; Christian Schmidt [1]; Leon Luk Bozo Richter [1]; Marco Diomede [2]; Holger Schlarb [3]; Annika Eichler [3]; Julien Branlard [1]

[1] *Deutsches Elektronen-Synchrotron DESY*

[2] *DESY - MSK*

[3] *DESY*

Corresponding Authors: holger.schlarb@desy.de, christian.schmidt@desy.de, marco.diomede@desy.de, julien.branlard@desy.de, annika.eichler@desy.de, andrea.bellandi@desy.de, bozo.richter@desy.de, yue.sun@desy.de

Superconducting radio frequency (SRF) cavities are critical components in particle accelerators, where accurately calibrated RF signals are essential for assessing cavity bandwidth and detuning, providing key insights into cavity performance and facilitating optimal accelerator operation. In practice, however, calibration drift due to humidity and temperature fluctuations over time poses a significant challenge, potentially resulting in suboptimal operation and reduced efficiency. This study examines long-term calibration drifts and explores how environmental variables such as humidity and temperature affect this phenomenon. Relative humidity, in particular, is difficult to control and has been shown to have a strong impact on calibration drift. Building on these insights, we introduce machine learning-based approaches to forecast both relative humidity and calibration drift in SRF cavities. By leveraging advanced algorithms and historical data on cavity operation and performance, we develop predictive models that identify patterns and trends indicative of relative humidity and calibration drift. Two approaches are presented in this work, including a polynomial NARMAX model and an attention-based deep neural network. These models are trained and validated with real-world environmental and operational data, demonstrating their effectiveness in predicting relative humidity and calibration drift. This forecasting methodology enables real-time drift compensation and timely, automated recalibrations of RF signals, ultimately enhancing overall system stability and performance.

MLOps, Infrastructure and Scalability / 97

An Integrated Research Infrastructure framework for digital twins of laser-plasma acceleration experiments- 10'+5'

Author: Remi Lehe[1]

Co-authors: Johannes Blaschke [1]; Auralee Edelen [2]; Marco Garten [1]; Axel Huebl [1]; Revathi Jambunathan [1]; Aodhan McIlvenny [1]; Lieselotte Obst-Huebl [1]; Ryan Roussel [2]; Olga Shapoval [1]; Jean-Luc Vay [1]; Edoardo Zoni [1]

[1] *LBNL*

[2] *SLAC*

Corresponding Authors: edelen@slac.stanford.edu, amcilvenny@lbl.gov, axelhuebl@lbl.gov, oshapoval@lbl.gov, ezoni@lbl.gov, rjambunathan@lbl.gov, lobsthuebl@lbl.gov, jlvay@lbl.gov, jpblaschke@lbl.gov, rlehe@lbl.gov, rroussel@slac.stanford.edu, mgarten@lbl.gov

Laser-plasma acceleration is a promising acceleration technology for a number of applications due to the large accelerating gradient and unique beam properties that it produces. This technology is in active development, and experimental campaigns typically dedicate significant time to exploring the parameter space in real time, adjusting laser properties, target configuration, and other factors to find the optimal setup. Therefore, having AI/ML-driven digital twins capable of providing real-time optimization guidance, by learning from both experimental and simulation data, would be highly beneficial.

We present progress towards an Integrated Research Infrastructure framework that can run multi-GPU simulations during experimental campaigns, and combine the results with ongoing measurements to provide real-time guidance. The framework is primarily intended for use at LBNL's BELLA Center and leverages several open-source tools, including WarpX and LASY for simulations, lume-model for the AI/ML digital twin, NERSC's superfacility API for submitting multi-GPU simulations in real time and the Prefect platform for overall workflow orchestration. Additionally, we will discuss challenges in building ML models that must learn from distinct and potentially mismatched data sources.

Anomaly Detection and Diagnostics / 98

Improving Coincident Learning for Beam-based RF Station Fault Identification Using Phase Information at the SLAC Linac Coherent Light Source- 15'+5'

Author: Jason Liang[1]

Co-authors: Daniel Ratner [2]; Eric Darve [1]

[1] *Stanford University*

[2] *SLAC*

Corresponding Authors: darve@stanford.edu, dratner@slac.stanford.edu, jialiang@stanford.edu

The vast amount of data generated by accelerators makes manual monitoring impractical due to its labor-intensive nature. Existing machine learning solutions often rely on labeled data, manual inspection, and hyperparameter tuning, which limits their scalability. To address these challenges, we leverage coincidence learning—an unsupervised technique designed for multi-modal tasks—to automatically detect anomalies by identifying coincident patterns of behavior across two distinct segments of the feature space. Specifically, we focus on anomaly detection for radio-frequency (RF) stations at the SLAC Linac Coherent Light Source (LCLS). By analyzing shot-to-shot data from the beam position monitoring system alongside data from RF stations, we can identify the source of changes in the accelerator's status. Previous studies on RF stations produced reasonable results using time-asynchronous amplitude data, but ignored the richer information from time-synchronous phase data due to its complexity. We find that using neural networks to analyze the phase data enables the detection of anomalies that amplitude-based detection missed. Additionally, the time-synchronous phase data provides critical insights, allowing us to distinguish whether an RF station change occurs simultaneously with changes in the accelerator status or in response to them. Additionally, the rich information contained in the phase data facilitates clustering of anomalies into distinct categories, each with unique signatures. This categorization brings us closer to identifying the root causes of issues within the RF stations.

Poster session / 99

Insertion Device Correction Based on Machine Learning Models at MAX IV 3GeV Ring

Author: Carla Caldeira Takahashi[None]

Co-authors: Adnan Al Medawer [1]; Hamed Tarawneh [2]; Michael Holz [2]

[1] *Jönköping University*

[2] *MAX IV*

Corresponding Authors: hamed.tarawneh@maxiv.lu.se, michael.holz@maxiv.lu.se, almedawer.adnan@gmail.com, carla.takahashi@maxiv.lu.se

Insertion devices in particle accelerators introduce orbit distortions that must be compensated to maintain beam stability. At MAX IV, this compensation traditionally relies on feedforward tables, which are time-consuming to generate and sensitive to changes in accelerator optics. This study explores the use of machine learning (ML) to automate the generation of feedforward tables without requiring extensive measurements. Using archived data from ID gaps, beam position monitors (BPMs), and corrector magnets, a neural network-based model was developed to replicate the current ID compensation system. Preliminary results show that the model effectively reproduces existing compensation behaviour and suggests potential for adaptive feedforward tables that refine themselves with online data. In parallel, alternative ML approaches focused on minimising beam position error are being investigated. These efforts aim to improve the maintainability of ID compensation and prepare orbit control for future optical changes and new operational scenarios.

Optimisation and Control / 100

Photon systems automation activities at EuXFEL 15'+5'

Authors: Richard Bean[1]; Johan Bielecki[1]; Sarlota Birnsteinova[1]; David Doblas-Jimenez[1]; Danilo Enoque Ferreira de Lima[1]; Ana Garcia-Tabares[1]; Luca Gelisio[1]; Steffen Hauf[1]; Ivars Karpics[1]; Jayanath Koliyadu[1]; Romain Letrun[1]; Thomas Michelat[1]; Egor Sobolev[1]; Florian Sohn[1]; James Wrigley[1]; Peter Zalden[1]

[1] *European XFEL GmbH*

Corresponding Author: sarlota.birnsteinova@xfel.eu

Large-scale facilities like European XFEL consist of a multitude of subsystems, which often require frequent calibration. Additionally, accurate real-time tuning of many of these subsystems is critical to maintain stable and optimal performance. Automation techniques can be leveraged to reduce operators' time investment and potentially increase the exploitation of allotted beamtime, both in quantity and quality. One example is the automation of photon beam alignment through an instrument, achieved by adjusting multiple optical components, such as mirrors and lenses, using precise and constrained actuation. This process uses Bayesian optimization to iteratively determine the optimal configuration by evaluating system performance metrics, such as beam intensity, shape and position, which requires reliable image processing. In this talk, I will introduce ongoing activities aimed at automating selected components of the photon system at the European XFEL.

Surrogate Modelling and Digital Twins / 101

Utilizing Optical Fibers and Machine Learning for Radiation-Tolerant Beam Imaging- 15'+5'

Author: Qiyuan Xu[None]

Co-authors: Alexander Hill [1]; Carsten Peter Welsch [2]; Federico Roncarolo [3]; Georges Trad [3]; Hao Zhang [4]

[1] *Liverpool*

[2] *Cockcroft Institute / University of Liverpool*

[3] *CERN*

[4] *University of Liverpool/Cockcroft Institute*

Corresponding Authors: georges.trad@cern.ch, hao.zhang@cockcroft.ac.uk, qiyuanxu95@gmail.com, c.p.welsch@liverpool.ac.uk, federico.roncarolo@cern.ch, ahill94@liverpool.ac.uk

Beam imaging presents significant challenges due to the necessity of positioning imaging devices near the beam pipe, an area subjected to high levels of radiation that can damage cameras and their peripheral electronics, reducing their lifespan and reliability. With the global discontinuation of radiation-hardened tube cameras previously used for this purpose, a robust and durable replacement imaging solution is needed. Multimode optical fibers have emerged as viable alternatives, capable of relaying the image signal to a standard CMOS camera location in a radiation-safe environment. A challenge within this approach is mode coupling and scattering within the fiber, which increases the difficulty in accurately reconstructing beam information.

This contribution showcases a method of reconstructing transverse beam distribution parameters from a distorted fiber output. This is achieved with an experimental setup that makes use of a synthetic input dataset generated from multiple high-variance 2D Gaussian fields, multimode optical fibers to propagate and distort these images, and a 2D convolutional autoencoder to reconstruct the inputs. This setup is used as a training dataset, with the input dataset chosen to support generalizability. Our machine learning model is tested on a real dataset of transverse beam distributions collected at CERN's CLEAR facility. We achieve an average RMSE of 2.44% over four key transverse beam parameters after reconstruction on the testset. Our model further demonstrates mitigated bias in beam parameter estimation and strong generalization capability, reconstructing well radically different parameter distributions.

MLOps, Infrastructure and Scalability / 102

Machine Learning for the SNS Accelerator and Target- 15'+5'

Author: Willem Blokland[1]

Co-author: Malachi Schram

[1] *ORNL*

Corresponding Authors: blokland@ornl.gov, schram@jlab.org

We apply Machine Learning techniques at the Spallation Neutron Source (SNS) to improve operations, specifically to deter and prevent errant beam pulses, to speed up minimization of halo beam losses, and to alert operators to anomalies in the target cooling system. We give an overview of the work done and discuss the infrastructure implemented and under development to support the data acquisition, pre-processing, training, and continuous learning.

Poster session / 103

Bayesian optimization in the Karabo SCADA system at the European XFEL

Authors: Danilo Enoque Ferreira de Lima[1]; David Doblas-Jimenez[1]; Florian Sohn[1]; Luca Gelisio[1]; Peter Zalden[1]; Sarlota Birnsteinova[1]; Steffen Hauf[1]

[1] *European XFEL*

Corresponding Authors: florian.sohn@xfel.eu, sarlota.birnsteinova@xfel.eu

The automation of repetitive tasks at FEL light-sources like the European XFEL allows to allocate staff more efficiently and increases reliability and safety. Recurrent procedures often require adjustments of various control parameters to maximize a measured quantity, for example, repositioning of optical components during beam-alignment to maximize intensity. In many cases an analytic form of the underlying objective function is not known. Moreover, movement speeds of the hardware components can be slow, rendering function sampling time-consuming.

In this contribution we present a flexible software approach using Bayesian optimization as an efficient sampling strategy for this type of optimization problems. At the European XFEL, the Karabo SCADA system is used to control the majority of instrumentation in the facility and handles the acquisition of hardware control parameters as well as scientific data. Hence, Karabo is well-suited as

a basis for the implementation and application of automation routines. Our optimization algorithm is implemented in Karabo.

Poster session / 104

Accelerator Surrogate Model Based on Diffusion Models: A Case Study of the MEBT Section in CAFe2

Author: Kunxiang Sun[1]

[1] 近代物理研究所

Corresponding Author: sunkunxiang@impcas.ac.cn

This paper introduces an accelerator surrogate model utilizing Diffusion Models. The proposed model leverages the parameters of the incoming beam bunch as conditional inputs to accurately and efficiently predict the phase space distribution of the beam at the accelerator's exit. By focusing on the Medium Energy Beam Transport (MEBT) section of the China Accelerator Facility for Superheavy Elements (CAFe2), this study constructs an overall surrogate model by progressively substituting each accelerator component with a diffusion model and connecting them in series. This method not only significantly reduces computational complexity but also ensures high prediction accuracy, providing substantial practical significance for the design, optimization, and operational maintenance of accelerators.

Poster session / 105

Machine Learning-Driven Calibration of the Beam Position Monitors in the SPS

Author: Vasiliki Stergiou[1]

[1] *University of Oxford (GB)*

Corresponding Author: vasiliki.stergiou@cern.ch

In the SPS, a flexible machine serving the LHC and a multitude of fixed-target experiments and fast-extraction facilities, reliable monitoring of the transverse beam position across a wide range of different beam structures and intensities is essential for stable and efficient operation. Today, the calibration procedure and signal processing of the beam position monitors (BPMs) of the SPS –and, essentially, the mapping of the digital ADC output to the analog beam signal –rely on statistical interpolation with a third-order polynomial model. This approach introduces position uncertainties reaching up to a few ˜100 um. Depending on the beam intensity, structure and displacement, this compromises precision in challenging experiments, like the ones taking place in the HiRadMat facility where the beam position precision is key for the physics experiments there.
In this work, we present a machine learning-based approach for the BPMs calibration using neural networks, which aims to a more accurate mapping of the BPMs response and thus the effective reduction of these systematic errors. Using moderately deep neural networks, first results already suggest that a machine learning-based approach can outperform the previous calibration method and lead to smaller systematic errors in position measurement readings across the beam intensity and beam displacement phase space.

Poster session / 106

Application of Bayesian Optimization on Booster to AGS Injection at BNL through Xopt –Experiences and Challenges

Author: Eiad Hamwi[1]

Co-authors: Alex Burkhart [2]; Georg Hoffstaetter [1]; Kevin Brown [2]; Rachel Terheide [2]; Vincent Schoefer [2]; Weijian Lin [2]

[1] *Cornell University*

[2] *Brookhaven National Laboratory*

Corresponding Authors: georg.hoffstaetter@cornell.edu, wlin1@bnl.gov, terheide@bnl.gov, aburkhart@bnl.gov, brownk@bnl.gov, schoefer@bnl.gov, eh652@cornell.edu

Accelerator control systems consist of very large numbers of parameters, many of which are continually re-tuned to account for different working conditions and drifting optimal points. In the case of the BNL RHIC Injector Complex, many different ion species are accelerated, and while there is a scarcity of diagnostics, there is a surplus of control knobs for optimizing the injection process. In this report, we investigate the use of Bayesian optimization (BO) to optimize the injection for the highest beam brightness in the AGS. We used up to 4 steering magnets in the BtA and up to 4 quadrupoles. The most suitable magnets were chosen by an investigation of the betatron phase advance to facilitate an efficient BO process. An integrated current transformer captured the injected beam intensity while the emittance was estimated via an Ion Profile Monitor. It was demonstrated that the chosen magnets effectively recovered a high intensity beam from a poorly-tuned configuration, using an Xopt implementation of BO, without increasing the beam profile. Systematic effects and noise from the IPM signal made clear readings challenging, but a new electron IPM is being configured for better profile measurements.

Poster session / 107

Surrogate Models for X-Ray Pulse Diagnostics at the European XFEL

Authors: Danilo Enoque Ferreira De Lima[1]; Fady Bishara[1]; Luca Gelisio[1]

[1] *European XFEL GmbH*

Corresponding Authors: fady.bishara@xfel.eu, danilo.enoque.ferreira.de.lima@xfel.eu, luca.gelisio@xfel.eu

The European XFEL is a scientific research facility that produces ultra-short and ultra-brilliant x-ray pulses. The facility is 3.4 kilometers long and comprises, very simplistically, three main sections: a linear electron accelerator, undulators, and photon beamlines. The entire facility is densely instrumented with various diagnostic devices that produce a vast amount of diverse data.

This talk focuses on an on-going effort that aims to understand the interplay between the various components of the whole facility on the properties of the produced x-ray pulses via a surrogate model. The sheer amount of available diagnostic data makes such a holistic approach possible and potentially invaluable. While still at an early stage, data analysis and preliminary results already provide useful insights into the correlation between electron bunches and x-ray spectral properties at MHz repetition rates. Furthermore, the goal of the program is not only to provide a surrogate model of the machine, but also to allow for its inversion; i.e. providing a systematic method to obtain self-consistent machine setting ranges that reliably produce the desired photon beam properties.

Optimisation and Control / 108

ML optimization methods for APS-U commissioning- 15’+5’

Author: Nikita Kuklev[1]

Co-authors: Louis Emery [2]; Michael Borland [2]; Hairong Shang [2]; Yine Sun [2]; Vadim Sajaev [2]

[1] *Fermilab*

[2] *Argonne National Laboratory*

Corresponding Author: nikita.kuklev@gmail.com

The Advanced Photon Source (APS) facility has just completed an upgrade to become one of the world's brightest storage-ring light sources. For the first time, machine learning (ML) methods have been extensive used as part of the baseline commissioning plan. Most popular such method was Bayesian optimization (BO) –a tool for efficient online high-dimensional single and multi-objective tuning. In this paper we will present our BO development work on experimentally motivated augmentations - uncertainty-aware simulation priors, parameter space and acquisition function refinement for multi-objective optimization, and online execution time improvements. These improvements were integrated into the APSopt optimizer, which was then successfully used for various commissioning tasks. We will show results of tuning linac and booster transmission efficiency, injection trajectory stabilization, and of extensive multi-objective storage ring dynamic/momentum aperture studies. Given the success of BO methods, work is proceeding on tighter integration into the control room.

Surrogate Modelling and Digital Twins / 109

End-to-end differentiable digital twin for the IOTA/FAST facility-15'+5'

Author: Nikita Kuklev[1]

Co-authors: Alexander Romanov [1]; Dean Edstrom [1]; John Wieland [2]; Jonathan Jarvis [1]; Michael Wallbank [1]; Nilanjan Banerjee [1]

[1] *Fermilab*

[2] *MSU*

Corresponding Author: nikita.kuklev@gmail.com

As the design complexity of modern accelerators grows, there is more interest in using controllable-fidelity simulations that have fast execution time or yield additional insights as compared to standard codes. One notable example of additional information are gradients of physical observables with respect to design parameters produced by differentiable simulations. The IOTA/FAST facility has recently begun a program to implement and experimentally validate an end-to-end digital twin to serve as a virtual accelerator test stand, allowing for rapid prototyping of new software and experiments with minimal beam time costs. In this contribution we will discuss our plans and progress. Specifically, we will cover the selection and benchmarking of both physics and ML codes for linac and ring simulation, the development of generic interfaces between surrogate and physics-based sections, and presenting the control interface as either a deterministic event loop or a fully asynchronous EPICS soft IOC. We will also discuss challenges in model calibration and uncertainty quantification, as well as future plans to extend modelling to larger machines like PIP-II and Booster.

Optimisation and Control / 110

Physics-informed Bayesian inference and optimization of the closed orbit in synchrotrons -15'+5'

Author: Victoria Isensee[None]

Co-authors: Adrian Oeftiger [1]; Oliver Boine-Frankenheim

[1] *University of Oxford*

Corresponding Authors: boine-frankenheim@temf.tu-darmstadt.de, victoria.isensee@tu-darmstadt.de, adrian.oeftiger@physics.ox.ac.uk

Heavy ion synchrotrons, like the SIS18 at GSI, rely on the proven numerical approaches to correct the closed orbit. The SIS18 has a relative moderate amount of BPMs (one per cell) and requires a well corrected and known orbit, especially near the injection/extraction systems. Fluctuations of the BPM signal arise from the electronics. In addition there are systematic errors due to the relative positioning of the BPMs. At specific locations, like the the beam position at the extraction septum, it is desirable to have a prediction, including an uncertainty estimate. An adapted approach towards closed orbit correction is proposed that integrates probabilistic modeling with beam dynamics to infer a closed orbit including uncertainty quantification. Methods, such as LOCO (Linear Optics from Closed Orbits) and NOECO (Nonlinear Optics from Off-Energy Closed Orbits), are limited by the need for extensive orbit response matrix (ORM) measurements and lack uncertainty quantification. The proposed method leverages physics-informed Bayesian regression to develop a surrogate model that not only quantifies uncertainties at beam position monitors (BPMs) but also in between them, reducing the required data. A Gaussian Process (GP) model is used to incorporate beam dynamics by estimating the kernel (and mean function) through the evaluation of simulated realizations, with simulations based on a MAD-X model of the SIS18 lattice. The learned distribution of multipole misalignments enables a model of the closed orbit with integrated uncertainty and noise handling. This model is then used in a Bayesian optimization framework to correct the closed orbit and achieve minimal deviation at specific locations, such as at the septum. The approach has also broader applications towards more general optics corrections.

Surrogate Modelling and Digital Twins / 111

An Online Virtual Model for the ATLAS Ion Linac at Argonne-15'+5'

Author: Adwaith Ravichandran[1]

Co-author: Brahim Mustapha [1]

[1] *Argonne National Laboratory*

Corresponding Authors: aravichandran@anl.gov, brahim@anl.gov

A multifaceted virtual accelerator model that seamlessly integrates with the online experimental system would highly benefit the operators to test and evaluate beam tuning scenarios and apply them online. As part of this effort, the beam dynamics code TRACK is wrapped with control system architecture and the graphic user interface BADGER developed by SLAC. Customizability and task visualization are prioritized based on end user feedback. More important than the user experience is how well the simulation model agrees with the experimental measurements. A recent effort to account for beam steering and misalignment effects has reduced the error between simulation and measurements from ˜ 40% to ˜ 10%. More work to pare down this error and close the gap even further is currently underway. The concurrent execution of simulation with experiment for troubleshooting and parallel simulation-based optimizations to quickly determine ideal element settings that map to desired beam parameters are some of the features to be highlighted. Previously published AI/ML capabilities are incorporated to this extent. Surrogate models based on collected data to reproduce a given set of operating conditions (beam energy, charge state etc.) and use of an LLM to monitor simulation progress and make informed decisions at a higher level will be included as add-on features.

Poster session / 112

Development of AI-ML and Generative AI Tools for Accelerator Design and Optimization

Author: Brahim Mustapha[None]

Co-authors: Adwaith Ravichandran [1]; Onur Gilanliogullari [2]

[1] *Argonne National Laboratory*

[2] *Illinois Institute of Technology*

Corresponding Authors: ogilanli@hawk.iit.edu, aravichandran@anl.gov, brahim@anl.gov

Accelerator design, optimization and simulation, are important steps in the construction of future accelerator-based facilities. Typically, multiple alternatives are investigated, involving multiple design iterations, before selecting a final design. Once a design is selected, the focus shifts to the detailed design of individual accelerator components and end-to-end lattice design and optimization. This phase requires extensive, large-scale and time-consuming electromagnetics and beam dynamics simulations. AI-ML can significantly accelerate and integrate this design process. Additionally, there is significant potential for new discoveries in unexplored areas of the design space, which is currently constrained by the inability to visualize hidden correlations and the limited resources available for exploration. Generative AI can be instrumental in addressing these challenges by uncovering patterns and optimizing designs efficiently. These developments aim to (i) speed up the accelerator design process by developing and applying AI-ML techniques for the design, optimization and simulation of individual accelerator components and full lattices, and (ii) harness the discovery potential of AI-ML by deploying existing Generative AI tools for accelerator design, evaluating their strengths and weaknesses, developing and implementing improvements to enhance their effectiveness for these design tasks.

Poster session / 113

Multi-Objective Bayesian Optimization of Beam Emittances and Energy Spread in Linear Accelerators

Author: Chong Shik Park[1]

[1] *Korea University*

Corresponding Author: kuphy@korea.ac.kr

This study explores a multi-objective optimization and modeling approach for enhancing the performance of a linear accelerator (linac) used in Fourth Generation Synchrotron Radiation (4GSR) facilities, focusing on the minimization of horizontal and vertical emittances as well as energy spread at the linac end. Efficient control of these parameters is critical to achieving high beam quality, essential for advanced synchrotron applications. Traditional optimization methods often face challenges in balancing these competing objectives. Therefore, we employ a Bayesian optimization technique that combines probabilistic modeling with sequential sampling to efficiently navigate the high-dimensional parameter space of the linac configuration. Bayesian optimization allows for the systematic trade-off between exploration and exploitation, leveraging prior knowledge and continuously updating with new observations to optimize performance with fewer evaluations. The approach models the complex, nonlinear relationships between operational parameters and beam quality, providing an adaptive framework that adjusts for uncertainties inherent in linac operations. Our results demonstrate that Bayesian optimization not only reduces horizontal and vertical emittances and energy spread effectively but also achieves this in a computationally efficient manner. The findings suggest that this framework could serve as a robust optimization tool for other accelerator systems, offering a generalizable methodology for improving performance in multi-objective settings.

Poster session / 114

Phase Space Reconstruction of Heavy Ion Linac Beams at RAON Facility Using Machine Learning Techniques

Author: Chong Shik Park[1]

[1] *Korea University*

Corresponding Author: kuphy@korea.ac.kr

This study investigates the application of machine learning techniques for the phase space reconstruction of heavy ion linac beams at the Rare isotope Accelerator complex for ON-line experiments (RAON) facility in Korea. Phase space analysis is a critical component in understanding and optimizing beam dynamics, enabling precise control of beam quality for advanced nuclear physics experiments. Leveraging modern machine learning methods, including neural networks and differentiable simulations, the proposed approach seeks to reconstruct the multidimensional phase space distribution from limited and noisy measurement data. These methods utilize their ability to model nonlinear relationships and infer missing information, overcoming traditional challenges associated with high-dimensional data processing in heavy ion accelerators. The framework incorporates beam diagnostics data, such as beam profiles and time-of-flight measurements, as input to train predictive models capable of reconstructing spatial, angular, and energy distributions with high fidelity. Preliminary results suggest significant improvements in reconstruction accuracy compared to conventional techniques, along with potential for real-time implementation. This work highlights the feasibility and effectiveness of machine learning for beam diagnostics and optimization in state-of-the-art heavy ion linacs, paving the way for enhanced performance and operational efficiency at facilities like RAON.

Poster session / 115

IMPROVEMENTS IN KOOPMAN OPERATOR METHOD FOR NONLINEAR DYNAMICS ANALYSIS USING SYMPLECTIC NEURAL NETWORKS

Author: Kelly Anderson[None]

Co-author: Yue Hao

Corresponding Authors: haoyue@msu.edu, anderske@frib.msu.edu

Analyzing Hamiltonian systems has been aided by advancements in data drive methods such as using these methods to calculate Koopman operators. The Koopman operator is an approach to simplify nonlinear dynamic systems. The symplectic condition is a strong constraint on Hamiltonian systems that when met preserves phase space area and certain constants of motion. By including the symplectic constraint in the calculation of Koopman operators we can ensure these properties are met and be able to more accurately make long term predictions about a system. The development of symplectic neural network layers called SympNets*ensure that the symplectic condition is met without the use of symplectic integrators, extra gradient calculations. The paper will present on improvements in the previous approach** of using SympNets to calculate Koopman operators and approximate constants of motion in two dimensional systems, with some preliminary efforts in four dimensional systems.

*P. Jin, A. Zhu, G.E. Karniadakis, and Y. Tang, "Symplectic networks: Intrinsic structure-preserving networks for identifying Hamiltonian systems,"CoRR abs/2001.03750, (2020). doi:10.48550/arXiv.2001.03750

**Anderson, K., & Hao, Y. (05 2024). Koopman operator method for nonlinear dynamics analysis using symplectic neural networks. Proc. IPAC'24, 713–716. Presented at the Nashville, TN. doi:10.18429/JACoW-IPAC2024-MOPS10

Poster session / 116

Digital Assistants using LLMs for the Electron Ion Collider Project

Authors: Jennefer Maldonado[1]; Kyle Kulmatycski[1]

[1] *Brookhaven National Laboratory*

Corresponding Authors: kkulmatycski@bnl.gov, jmaldonad@bnl.gov, lnguyen@bnl.gov

Brookhaven National Laboratory is home to future Electron Ion Collider (EIC). The EIC will collide electrons with protons and nuclei to produce snapshots of particles'internal structure. This will allow us to study the role of gluons in the matter within and all around us. As the project expands, the growing collection of technical documentation across multiple sub-systems creates challenges in information access and knowledge sharing. Leveraging local LLMs offer greater data privacy, security, and compliance control since possible sensitive information never leaves internal systems, unlike with commercial public cloud-based systems. We propose a digital assistant framework utilizing the Ollama API with locally-deployed LLMs, ensuring data privacy while maintaining full control over the information processing pipeline. Our approach combines open-source models with natural language processing techniques to a domain-specific assistant that processes EIC technical repositories, accelerating project development by providing rapid access to critical information. We will discuss techniques, applications, and use cases of this developed digital assistant.

Poster session / 117

Machine learning activities in support of the Australian Synchrotron 2.0 project

Author: Xuanhao Zhang[1]

Co-authors: Rohan Dowd [2]; Tessa Charles [3]

[1] *ANSTO - Australian Synchrotron*

[2] *Australian Synchrotron - ANSTO*

[3] *Australian Synchrotron (AU)*

Corresponding Authors: fzhang@ansto.gov.au, rohan.dowd@synchrotron.org.au, tessa.charles@cern.ch

A new project is underway to develop the successor to the current Australian Synchrotron facility. The new facility will feature ultra-low emittance levels in-line with upcoming 4th generation light sources. A series of machine learning and optimisation activities are in progress to address challenges associated with the realisation and operation of ultra-low emittance machines. In particular, this paper will discuss three areas of activity. First, optimisation of non-linear optics using a multi-objective genetic algorithm. Second, study on optimal beam position monitor placement using differentiable lattice simulation code and Fisher information. Third, feasibility study on improving closed-orbit correction efficiency by training neural networks on orbit response data.

Anomaly Detection and Diagnostics / 118

Semi-supervised detection of optics errors in beamlines- 15'+5'

Author: Jonathan Edelen[None]

Co-authors: Christopher Hall [1]; Kevin Brown [2]; Nathan Cook [1]; Vincent Schoefer [2]; Yuan Gao

[1] *RadiaSoft LLC*

[2] *Brookhaven National Laboratory*

Corresponding Authors: brownk@bnl.gov, ncook@radiasoft.net, jedelen@radiasoft.net, schoefer@bnl.gov, chall@radiasoft.net, ygao@bnl.gov

Optics tuning in transfer lines and LINACs can be challenging due to the fact that multiple combinations of machine settings can lead to the same diagnostic output. Moreover, the lack of a periodic solution can limit the ability to infer optics in the same way as rings from BPM signals. Model based approaches are often used to assist with the optics tuning in combination with optimization or parameter estimation. Here we have developed a novel approach using machine learning inverse models trained on a known configuration to detect variations in quadrupole settings without explicitly including them in the model. This paper shows a comparison of neural network models and linear models on both a simulation-based study and experimental studies conducted at the AGS to RHIC transfer line at Brookhaven National Lab.

Poster session / 119

Integrated Denoising for Improved Stabilization of RF Cavities

Authors: Jonathan Edelen[None]; Morgan Henderson[None]

Co-authors: Auralee Edelen ; Christopher Hall [1]; Jorge Diaz-Cruz [2]; Joshua Einstein-Curtis [1]

[1] *RadiaSoft LLC*

[2] *SLAC*

Corresponding Authors: jedelen@radiasoft.net, edelen@slac.stanford.edu, mhenderson@radiasoft.net, jeinstei@gmail.com, chall@radiasoft.net

Typical operational environments for industrial particle accelerators are less controlled than those of research accelerators. This leads to increased levels of noise in electronic systems, including radio frequency (RF) systems, which make control and optimization more difficult. This is compounded by the fact that industrial accelerators are mass-produced with less attention paid to performance optimization. However, growing demand for accelerator-based cancer treatments, imaging, and sterilization in medical and agricultural settings requires improved signal processing to take full advantage of available hardware and increase the margin of deployment for industrial systems. In order to improve the utility of RF accelerators for industrial applications we have developed methods for removing noise from RF signals and characterized these methods in a variety of contexts. Here we expand on this work by integrating denoising with pulse-to-pulse stabilization algorithms. In this poster we provide an overview of our noise reduction results and the performance of pulse-to-pulse feedback with integrated ML based denoising.

Poster session / 120

Machine Learning Based Sample Alignment at ORNL

Authors: Jonathan Edelen[None]; Morgan Henderson[None]

Co-authors: Bhargavi Krishna [1]; Christina Hoffmann [1]; Gary Taufer [1]; Joshua Einstein-Curtis [2]; Stuart Calder [1]

[1] *ORNL*

[2] *RadiaSoft LLC*

Corresponding Authors: jedelen@radiasoft.net, mhenderson@radiasoft.net, choffmann@ornl.gov, caldersa@ornl.gov, jeinstei@gmail.com, krishnab@ornl.gov, tauferga@ornl.gov

Neutron scattering experiments are a critical tool for the exploration of molecular structure in compounds. The TOPAZ single crystal diffractometer at the Spallation Neutron Source and the Powder Diffractometer at the High Flux Isotope Reactor study these samples by illuminating them with different energy neutron beams and recording the scattered neutrons. Aligning and maintaining the alignment of the sample during an experiment is key to ensuring high quality data are collected. At present this process is performed manually by beamline scientists. RadiaSoft in collaboration with the beamline scientists and engineers at ORNL has developed a machine learning based alignment software automating this process. We utilize a fully-connected convolutional neural network configured in a U-net architecture to identify the sample center of mass. We then move the sample using a custom python-based EPICS IOC interfaced with the motors. In this poster we provide an overview of our machine learning tools and show our results aligning samples at ORNL.

122

Introduction and welcome

Corresponding Authors: miceli@fnal.gov, verena.kain@cern.ch

123

Introduction to Control Theory

Corresponding Author: annika.eichler@desy.de

Control theory is a pivotal field of study that focuses on the behavior of dynamical systems and the development of strategies to influence these systems towards desired outcomes. The principle of control theory find its application in plenty of disciplines including engineering, economics, biology and beyond. It were control concepts like the Kalman filter that has flew the Apollo to the moon. Also for the operation of particle accelerator there are thousands of control loops running, ranging from low-level hardware-focused systems, as magnet powering, synchronization or RF control, over fast transverse feedback systems to beam steering.
In this tutorial, we begin by exploring the basic concepts of control, emphasizing the importance of stability and robustness. The tutorial then introduces the mathematical modeling of systems using differential equations and the representation of these systems in state-space form. Starting from classical control strategies, we will go towards modern optimal control techniques and will finally see the link to optimization and ML approaches like reinforcement learning.

124

AI for particle accelerators - a roadmap, speaker to be confirmed

125

Multi-objective RL

Corresponding Author: kishan@jlab.org

Multi-objective reinforcement learning (MORL) extends traditional reinforcement learning (RL) by addressing environments where multiple conflicting objectives must be optimized simultaneously. In real-world applications, such as autonomous systems, particle accelerator optimization and control, agents often face trade-offs between competing goals. This lecture provides an overview of the key concepts, techniques, and challenges in MORL. We will explore different approaches to handling multiple objectives, including scalarization methods, Pareto dominance, and reward decomposition. Additionally, we will discuss the concept of the Pareto front and how it can guide decision-making in complex environments. The lecture will also highlight the role of exploration-exploitation trade-offs in multi-objective settings and the development of algorithms that can balance these objectives effectively. By the end of the session, attendees will gain a deeper understanding of how to design and implement multi-objective RL systems capable of navigating the complexities of real-world decision-making tasks.

126

Differentiable simulations

Corresponding Authors: jpga@uchicago.edu, jan.kaiser@desy.de, chenran.xu@kit.edu, ansantam@liverpool.ac.uk

Many accelerator physics problems, such as beamline design, beam dynamics model calibration, online tuning and phase space measurements rely on solving high-dimensional optimisation problems over beam dynamics simulations. Numerical optimisers have successfully been applied to such tasks, but they struggle as the dimensionality and complexity of the objective function increase. In machine learning, gradient-based optimisation algorithms are successfully used to optimise billions of model parameters over complex loss functions when training large neural network models. This is made possible by reverse-mode automatic differentiation, which enables the fast computation of gradients of complex functions. In this tutorial, you will learn to use novel beam dynamics simulations with support for automatic differentiation to your advantage and harness the power of gradient-based optimisation in accelerator physics. Multiple hands-on examples using the Cheetah beam dynamics code will allow you to try these methods for yourself. While we will present multiple example applications of gradient-based optimisation on differentiable beam dynamics simulators, the space of potential applications here is vast, and we believe that participants will go on to discover numerous novel applications for differentiable beam dynamics simulations that were intractable to solve with existing methods.

127

Highlights from RL4AA'25

Corresponding Authors: ansantam@liverpool.ac.uk, annika.eichler@desy.de

128

Poster session

129

Panel Discussion

130

summary, preview, final words

131

CERN visits

132

Registration

133

Large Language Model Chatbots for Enhanced Documentation Access

Corresponding Author: florian.matthias.rehm@cern.ch

This tutorial applies Retrieval Augmented Generation (RAG) as a method to improve documentation retrieval in accelerator physics. Participants will learn how combining information retrieval with generative AI models can provide precise, context-aware answers from vast technical resources. The session includes a hands-on demonstration of implementing RAG in combination with Large Language Models (LLMs) and retrieving information from your documentation.