Karlsruhe Institute of Technology

67th ICFA Advanced
Beam Dynamics
Workshop on Future
Light Sources **FLS 2023**

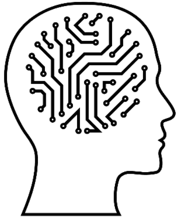

How can machine learning help future light sources?

A. Santamaria Garcia, C. Xu, L. Scomparin, E. Bründermann,
M. Caselle, G. De Carne, A. -S. Müller

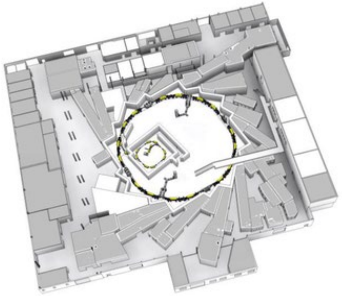

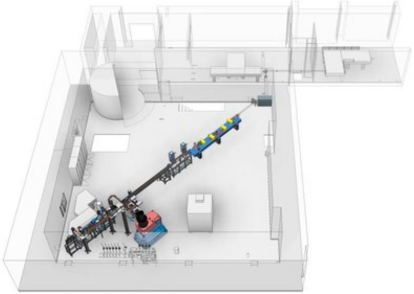

Accelerator roadmap

Energy frontier

Ability to generate new particles via high-energy collisions

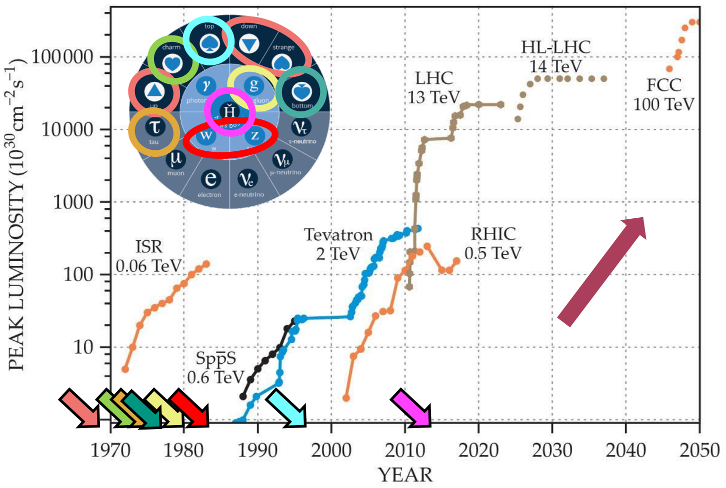

Photon science

Ability to probe atomic structures

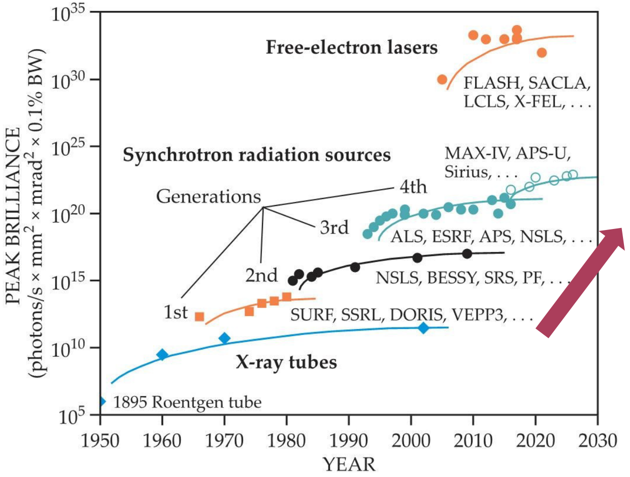

Technological innovation is needed to keep up with the challenging goals

Source: "Particle beams behind physics discoveries" (Physics Today)

Trends and challenges of frontier accelerators

Karlsruhe Institute of Technology

Denser beams for higher luminosity and brilliance

- Complex beam dynamics
- Complex accelerator design and operation

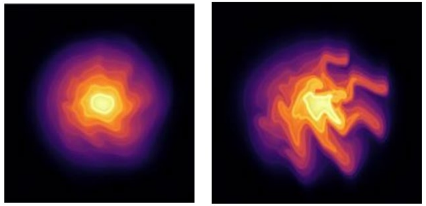

Image: KIT

Larger circular colliders for higher energies

- Orders of magnitude more signals
- Machine protection limits

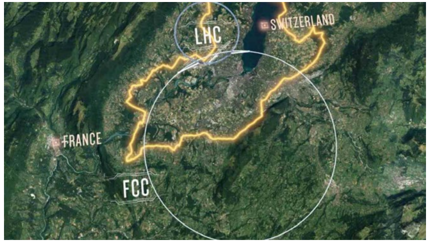

Image: CERN

Compact plasma accelerators with higher gradients

- Tight tolerances
- High-quality beams required

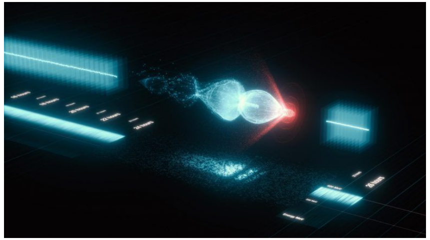

Image: DESY

What is machine learning?

ARTIFICIAL INTELLIGENCE (AI)

Computers mimic human behaviour

- First chatbots
- Robotics
- Expert systems
- Natural language processing
- Fuzzy logic
- Explainable AI

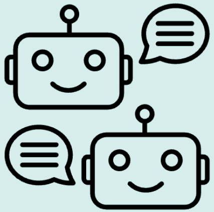

Narrow AI

MACHINE LEARNING (ML)

Computers learn without being explicitly programmed to do so and improve with experience

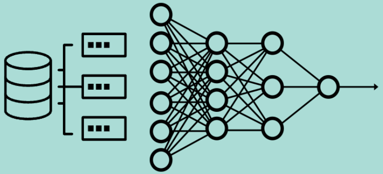

DEEP LEARNING (DL)

Multi-layered neural networks perform certain tasks with high accuracy

- Speech/handwriting recognition
- Language translation
- Recommendation engines
- Computer vision

Machine Learning

Learning style	Task	Popular algorithms
Supervised Learning **We know the input & output (labeled data)**	Classification (discrete variables) Regression (continuous variables)	▪ Neural networks (e.g. stochastic gradient descent, backpropagation) ▪ Support Vector Machine ▪ K-nearest neighbor ▪ Decision Tree algorithms (e.g. Classification and Regression Tree) ▪ Random Forest (ensemble) ▪ Uni or multivariate, linear or logistic
Unsupervised Learning **We only know the input (unlabeled data)**	Clustering	▪ K-means ▪ K-medians ▪ Expectation Maximization (EM) ▪ Hierarchical clustering
	Association	▪ Apriori algorithm ▪ Eclat algorithm

Reinforcement Learning	Control
	Classification
		▪ Model-free
		  ▪ Value based (Q-learning)
		  ▪ Policy based
		    ▪ Deterministic or stochastic
		    ▪ On or off policy
		    ▪ Policy gradient or actor-critic
		▪ Model based
		    ▪ Learn the model
		    ▪ Model given

Karlsruhe Institute of Technology

Deep Learning Networks

- Convolutional Neural Networks
- Recurrent Neural Networks
- Long Short-Term Memory Networks
- Autoencoders
- Deep Boltzmann Machine
- Deep Belief Networks

Bayesian Algorithms

- Naive Bayes
- Gaussian Naive Bayes
- Bayesian Network
- Bayesian Belief Network
- Bayesian optimization

Regularization, dimensionality reduction, ensemble, evolutionary algorithms, computer vision, recommender systems, …

Machine learning opportunities in accelerators

Advantages of ML methods:

- Yield fast predictions at a **reduced computational cost**
- Take into account **non-linear** correlations
- Adapt the predictions to the **drifts** in the machine state

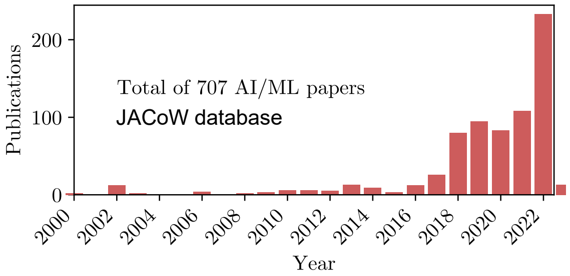

Task	Goal	Methods/Concepts	Examples[1]
Detection	Detect outliers and anomalies in accelerator signals for interlock prediction, data cleaning	• Anomaly detection • Time series forecasting • Clustering	• Collimator alignment • Optics corrections • SRF quench detection
Prediction	Predict the beam properties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

[1] non-exhaustive

A vision for future accelerators, driven by ML

Karlsruhe Institute of Technology

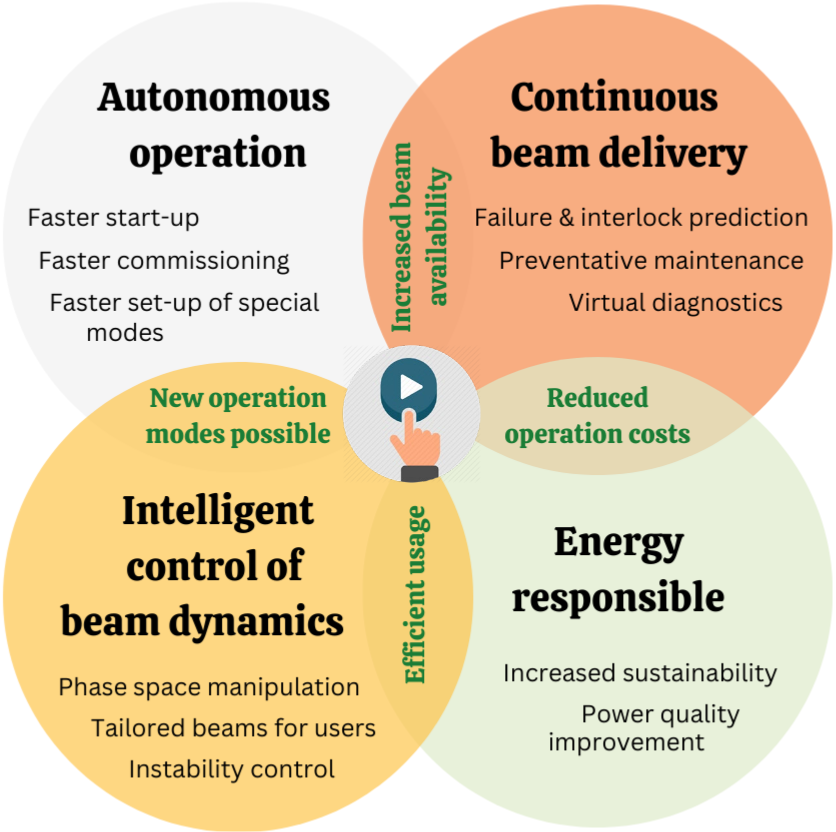

KITTEN

KIT Testfeld für Energieeffizienz und Netzstabilität in großen Forschungsinfrastrukturen

= KARA + Energy Lab 2.0

Energy-responsible infrastructure R&D from accelerator **components to systems**

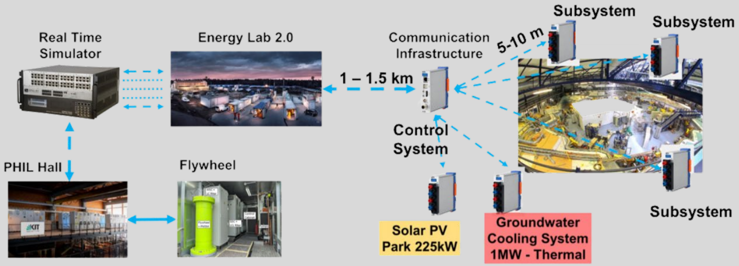

Real-time power consumption data from the accelerator will be fed to a **digital twin** that can emulate with high fidelity the **accelerator power and energy dynamics** during power systems studies.

ML for light sources – some examples

Source stabilization

Instability control

- Low-latency intelligent feedbacks

Assisting operation

- Automated set-up / tuning
  - Injection optimization, beam steering and focusing, PBA tuning
- Special operation modes
  - Negative α_c
- Pulse optimization
  - Energy, E-field, spectrum
- Faster commissioning
- Virtual diagnostics

Faster lattice design

In the future:

- ❑ Uncertainty quantification
- ❑ Explainability/interpretability
- ❑ Robustness
- ❑ Safety
- ❑ …

Accelerator facilities at KIT

Karlsruhe Institute of Technology

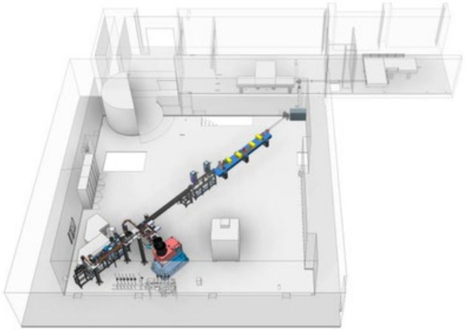

FLUTE

Linac-based THz source

41 MeV top energy

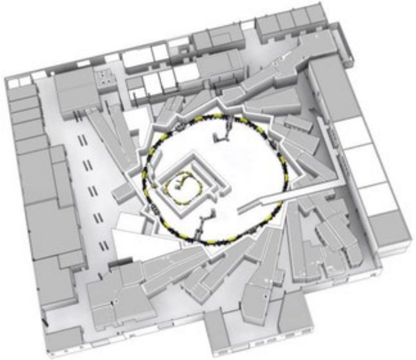

KARA

Synchrotron light source and storage ring

2.5 GeV top energy

Tailoring THz radiation with machine learning

Example: start-to-end pulse optimization in linac (FLUTE)

Karlsruhe Institute of Technology

Surrogate model

input
1. RF gun phase
2. RF gun amplitude
3. Solenoid current
4. Bunch charge

C. Xu et al. TUPOPT070, IPAC22

output
1. Mean energy
2. Energy spread
3. RMS bunch length
4. Beam size
5. Emittance
6. % remaining part.

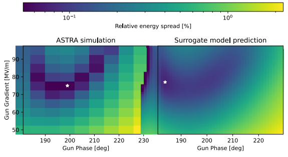

1 min 1 ms

Deep neural networks

Control over laser pulse length, pulse shape, spot size, spot position with spatial light modulators (SLMs)

C. Xu et al. WEPAB289, IPAC21

Can inform/guide the optimization with smart initial guesses

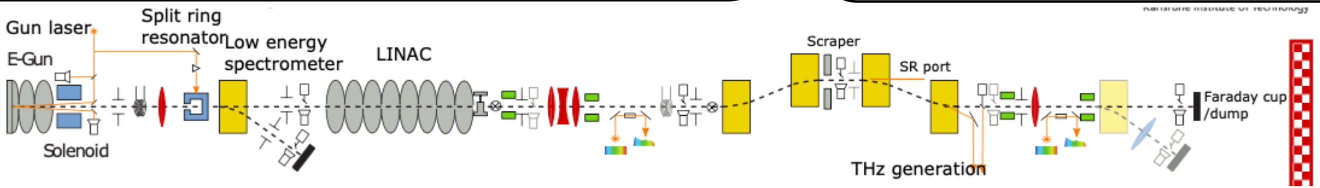

Parallel Bayesian optimization

C. Xu et al. WEPOMS023, IPAC22

input
1. RF gun phase
2. RF gun amplitude
3. Solenoid current
4. Linac phase
5. Linac amplitude
6. Chicane bending radius

objective
- Min. RMS bunch length after chicane
- Max. peak E-field of CSR pulse

observation
- Long. phase space
- Spectral intensity
- Form factor
- Bunch current profile
- THz pulse E-field

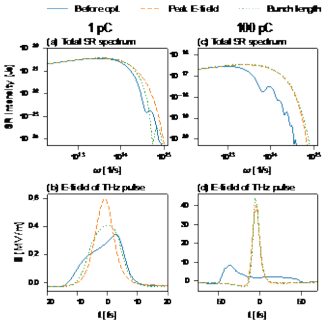

Bayesian optimization algorithm transferred to EuXFEL

- Time to inject to KARA cut in half with automated tuning by BO algorithm
  [C. Xu et al., PhysRevAccelBeams.26.034601]
- Emitted THz radiation at FLUTE optimized with parallel BO in simulation
  [C. Xu et al, IPAC'22-WEPOMS023]
- Transfer of algorithm to EuXFEL to tune SASE emission [C. Xu et al, IPAC'23-THPL028]

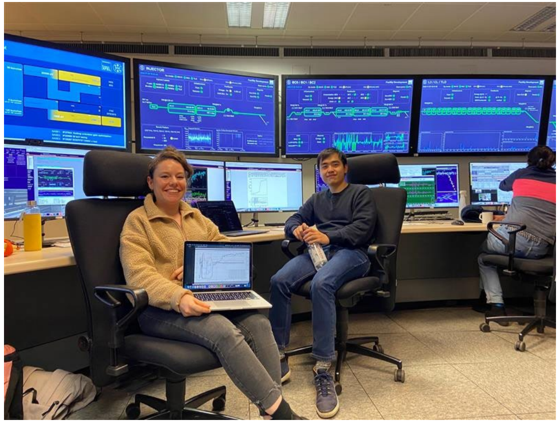

SASE1/ 250pC / 9.3keV

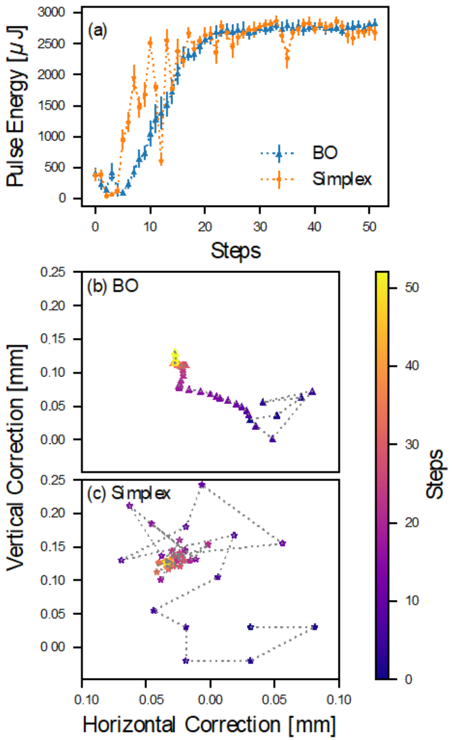

Code available at: https://github.com/cr-xu/bo-4-euxfel

Bayesian Optimization

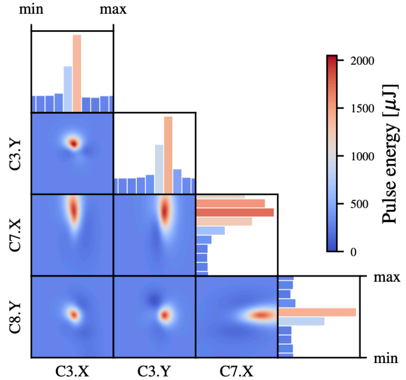

GP model can be used to visualize the sensitivity of actuators with respect to an objective and assist operators

First detailed comparison of BO and RL in a real accelerator

[J. Kaiser, C. Xu et al, arxiv: 2306.03739]

Reinforcement Learning

Bayesian Optimization

- **Task**: focus and position the electron beam
- **Actuators**: 3 quadrupole magnets + 2 corrector magnets
- **Observation**: beam image on the diagnostic screen

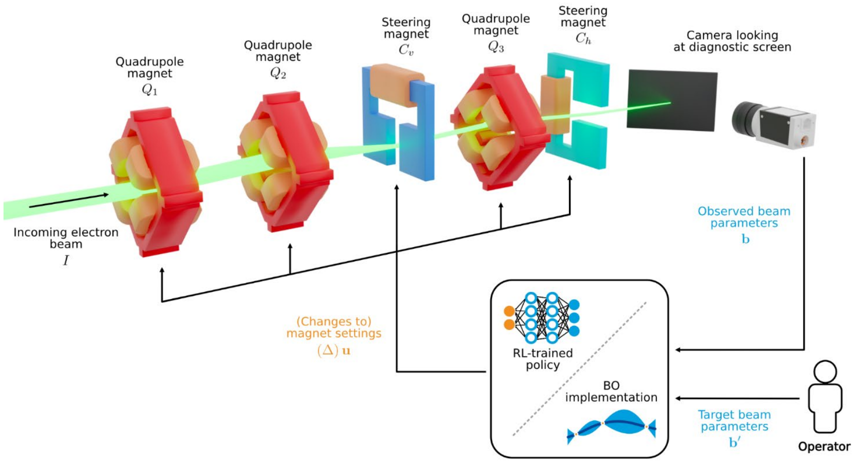

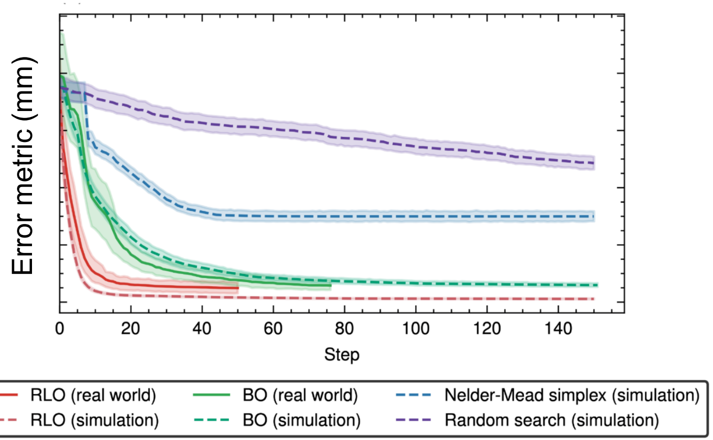

RL optimization outperforms BO

First detailed comparison of BO and RL in a real accelerator

Reinforcement Learning

Bayesian Optimization

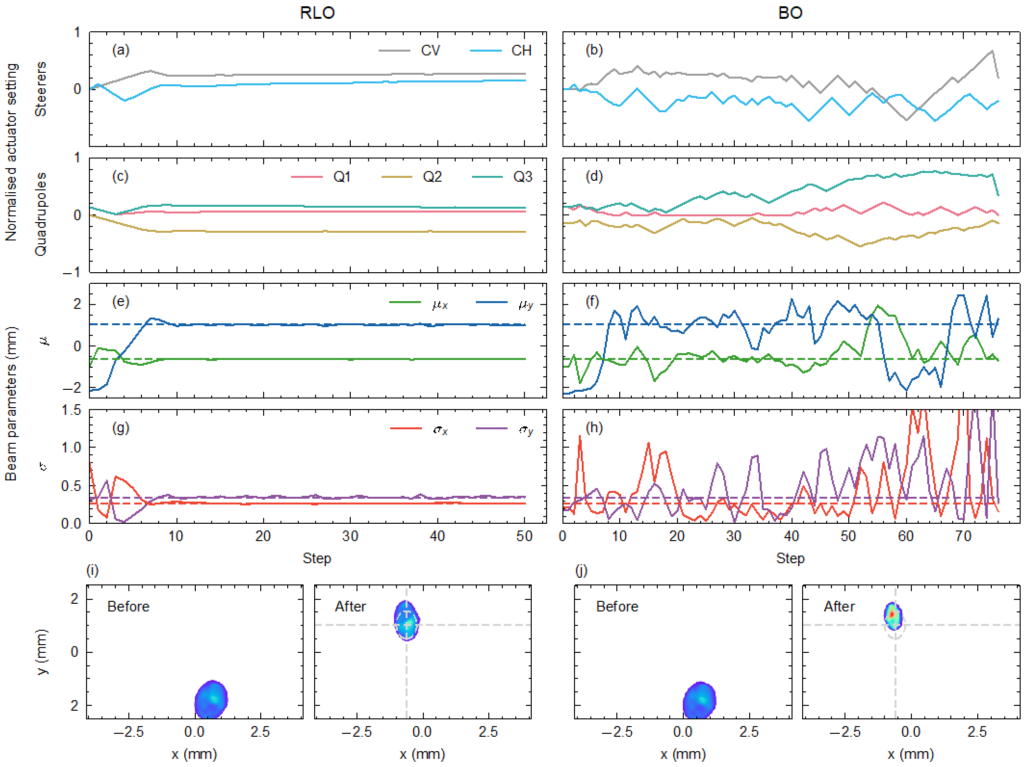

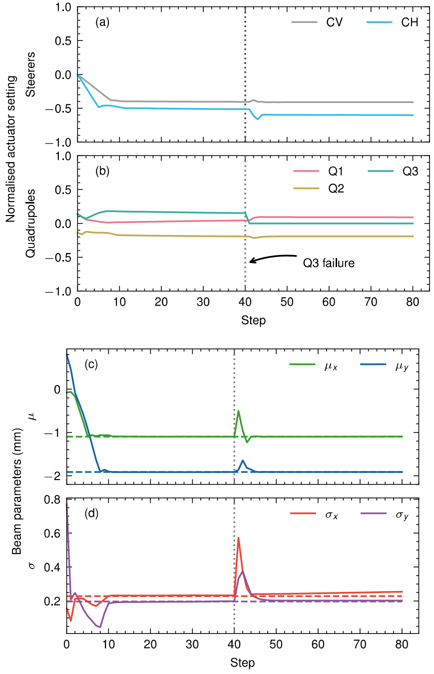

RL optimization can be used as feedback (e.g., magnet failure)

Lattice agnostic RL → Generalizable RL

- **Goal**: train a generalizable RL agent for transverse control
- **Method (domain randomization DR)**:
  - randomize magnet positions in training
  - keep QQQCC order

→ Agent can't memorize magnet settings

Reinforcement Learning

C. Xu et al, IPAC23-THPL029

Universal agent that can be deployed at similar but different accelerators

ARES

FLUTE

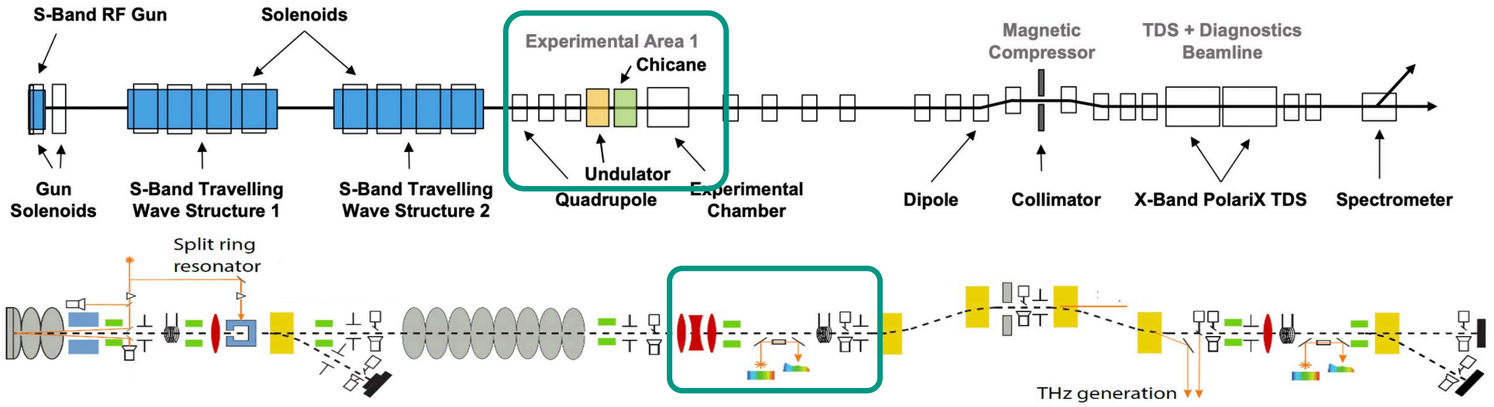

Lattice agnostic RL → Generalizable RL

Reinforcement Learning

C. Xu et al, IPAC23-THPL029

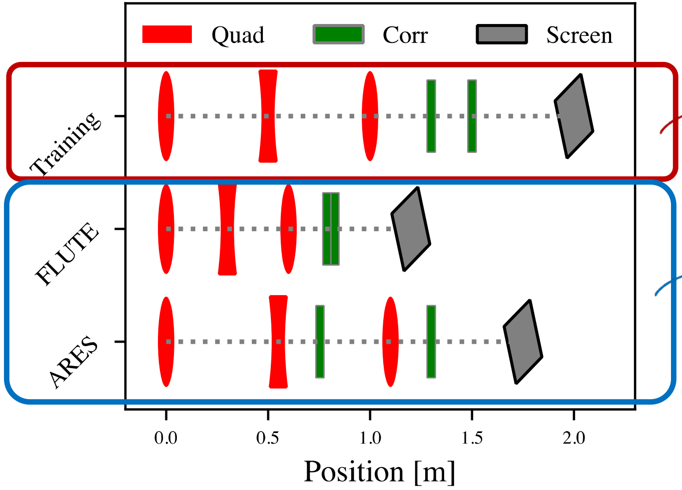

Used during training, with randomized positions but following order (=DR)

Test lattices

Reduction of performance when used in different lattices

Fine tuning (FT) → re-training with new lattice only 2% of the original training samples

Average of 100 tasks, max. 50 steps

Algorithm	Averaged best beam MAE [µm]			
	Tr-Ch	FL-Ch	FL-Oc	AR-Ch
Random	958	621	548	930
Vanilla-SAC	43	154	150	92
SAC-DR	86	63	60	67
SAC-DR+FT	52	31	31	36

Ch = Cheetah, tensorized optics simulation
Oc = OCELOT (with space charge)

Control of the microbunching instablity with RL

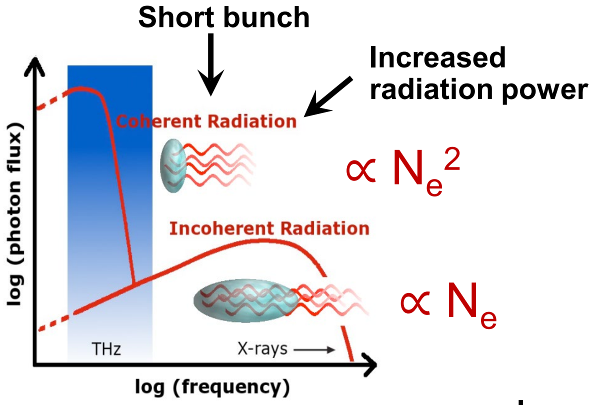

Low-α_c optics → MBI

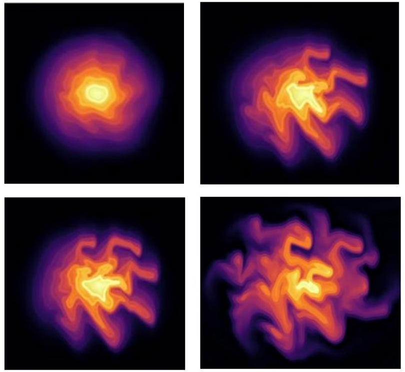

Bursting can be controlled with RF modulations

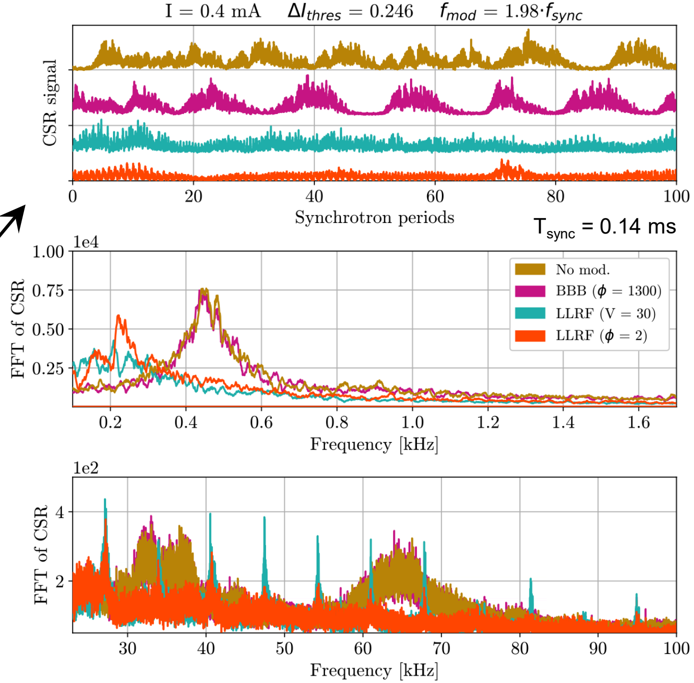

A. Santamaria Garcia et al, IPAC23-WEPA018

Control of the microbunching instablity with RL

Proposed control loop

[L. Scomparin et al, IBIC'22-MOP42]
[W. Wang et al, doi: 10.1109/TNS.2021.3084515]
[T. Boltz, doctoral thesis]

Schottky diode
analog pulse signal
50 GHz - 2 THz

→ coaxial →

KAPTURE-2
signal digitization

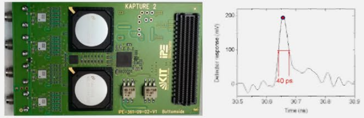

Low-latency high-throughput sampling
500 MS/s, 8 channels

→ FMC →

HighFlex 2
bunch labeling

Custom modular readout card

↓ fiber, aurora protocol 64b/66b

Xilinx Versal VCK190
decide action

Low-latency RL inference platform
1.6 Tera FP operations/s

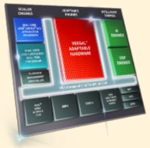

AI engines: feature extraction and agent inference
ARM processor: slow-control
FPGA: dataflow management

← serial

Feedback system
execute action
Low-level RF amplitude and phase modulation control
every 6 revolutions

Measured latency without re-training 2.5 us

⇄ 1 Gb ethernet

CPU/GPU
re-train agent

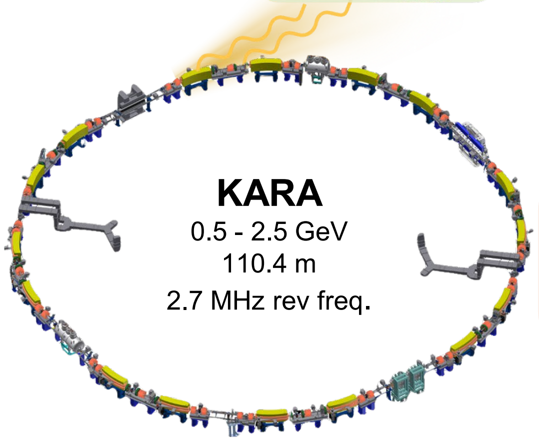

KARA
0.5 - 2.5 GeV
110.4 m
2.7 MHz rev freq.

Tests end of September

Damping of transverse oscillations

Proof-of-principle

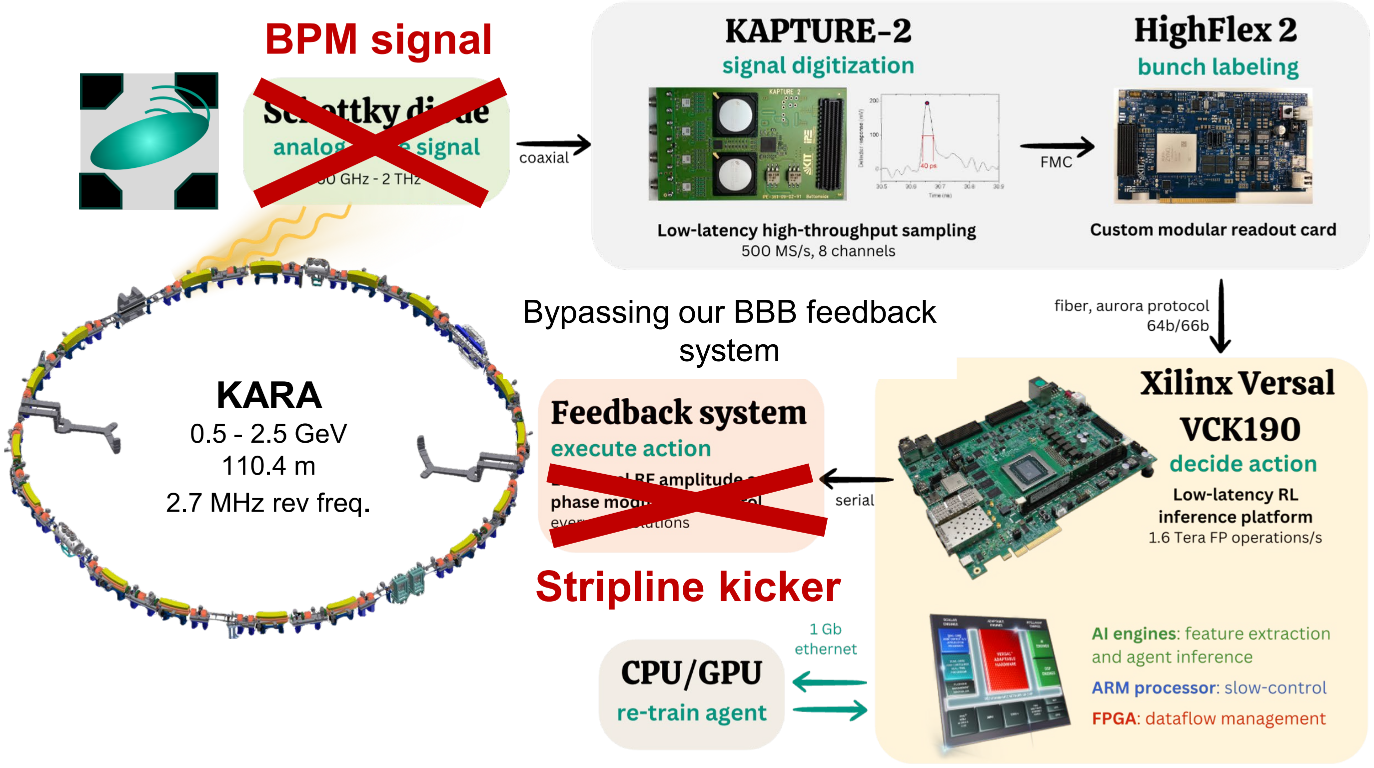

First RL algorithm online training and running on hardware in accelerators

Reinforcement Learning

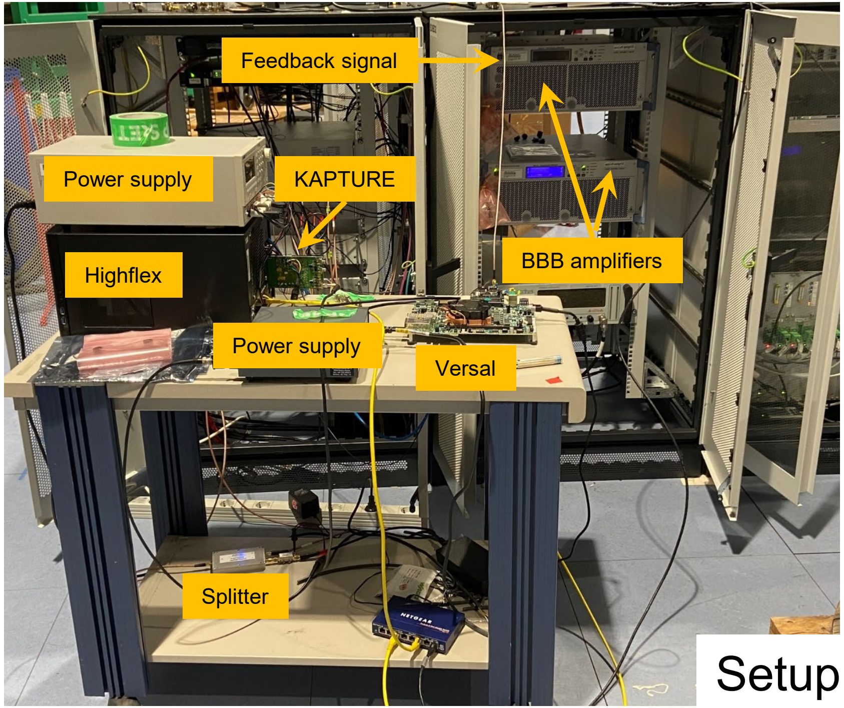

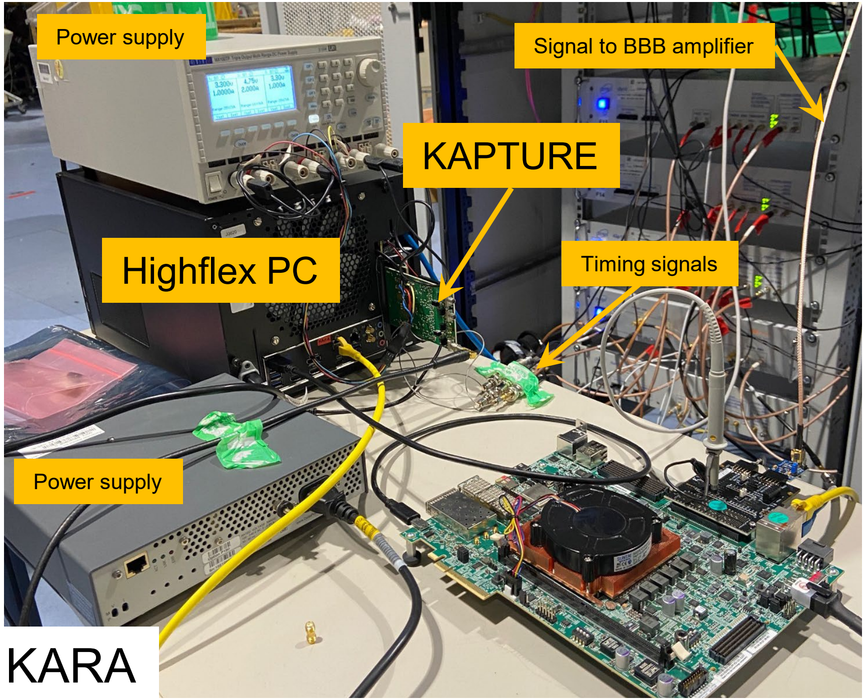

Setup at KARA

First RL algorithm online training and running on hardware in accelerators

Reinforcement Learning

- **Agent**: Vanilla PPO from Stable Baselines 3
- **Actor & critic architecture**: 8-16-1
- **Reward**: metric of the beam position (low as possible)
- **Observation**: last 8 BPM samples
- **Strategy:**

    1. Agent acts during 2048 turns (0.74 ms)
    2. Agent stops and is re-trained in a CPU (~2.6 s)
    3. New weights are sent to Versal board and agent starts again

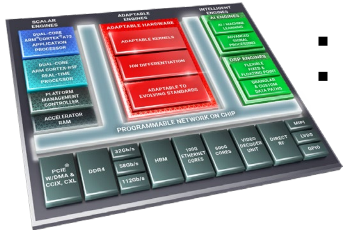

- NNs coded in Versal AIE
- Only forwad pass

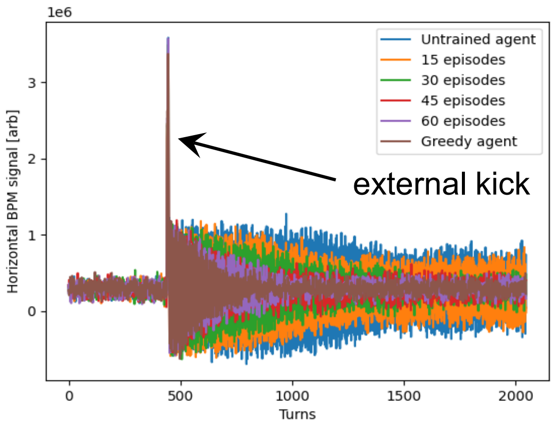

Damping improves with experience: the system is learning!

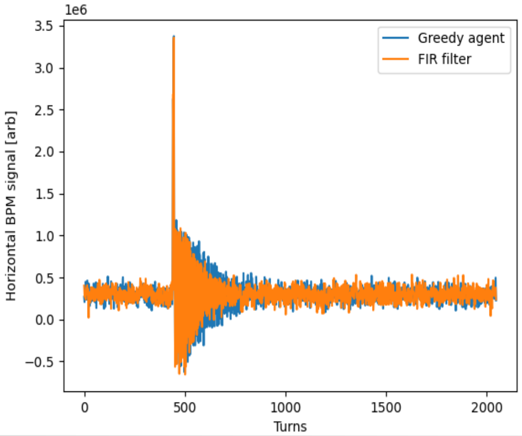

Achieves (sometimes surpassing) performance of FIR filter control (commercial solution)

L. Scomparin

First RL algorithm online training and running on hardware in accelerators

Reinforcement Learning

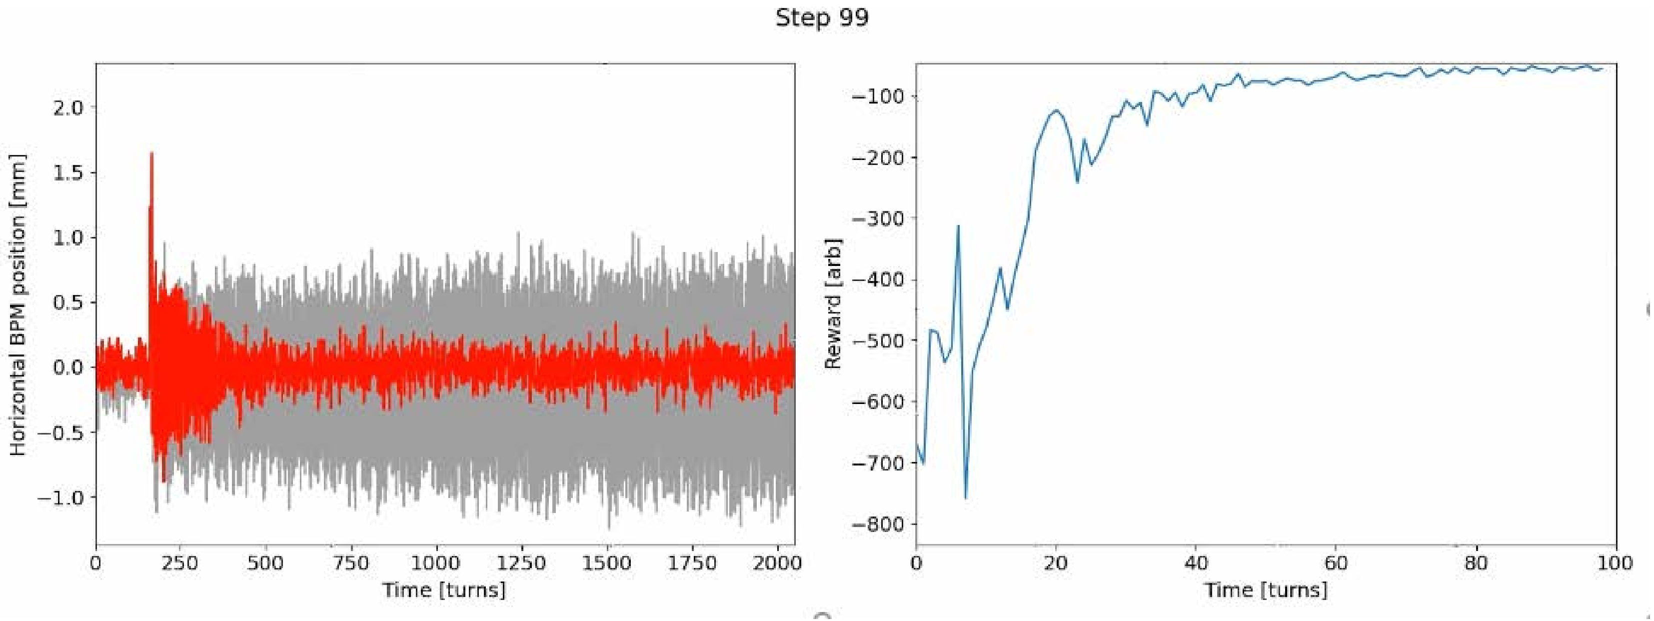

L. Scomparin

First RL algorithm online training and running on hardware in accelerators

Reinforcement Learning

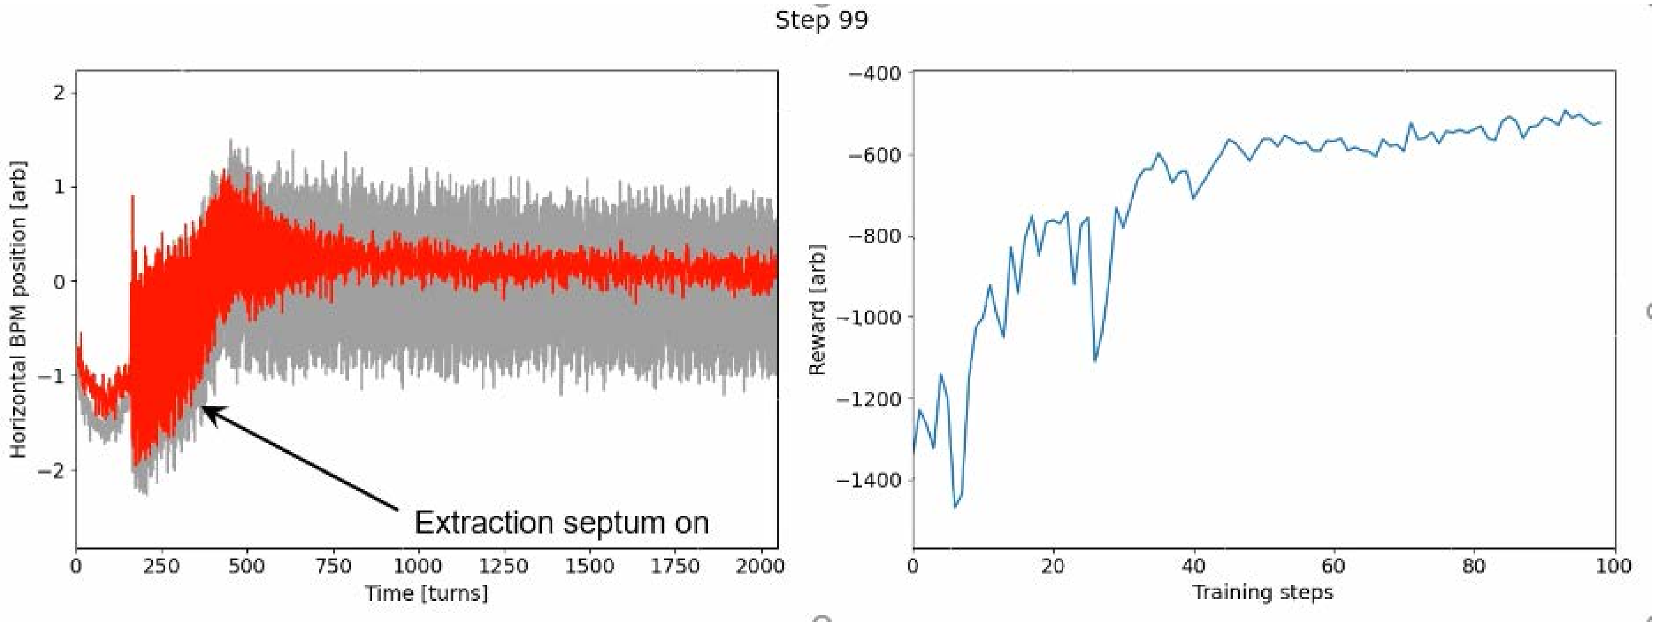

L. Scomparin

Outreach efforts

Creation of the Collaboration on Reinforcement Learning for Autonomous Accelerators (RL4AA)

https://rl4aa.github.io/

- Kick-off with workshop organized at KIT Feb. 2022
  - https://indico.scc.kit.edu/event/3280/overview
  - Expert lectures on reinforcement learning
  - Real application to accelerator tutorials
  - Advanced discussion sessions
- **Registration for Feb. 2023 open!**
  - https://indico.scc.kit.edu/event/3746/

ML tutorials

- https://github.com/RL4AA/RL4AA23
- https://github.com/ansantam/2022-MT-ARD-ST3-ML-workshop
- https://github.com/aoeftiger/TUDa-NMAP-14

Thank you for your attention!

What questions do you have for me?

Karlsruhe Institute of Technology
Institute for Beam Physics and Technology

Dr. Andrea Santamaria Garcia
AI4Accelerators team leader

andrea.santamaria@kit.edu
https://twitter.com/ansantam
https://www.linkedin.com/in/ansantam/
https://github.com/ansantam

paper

(we are hiring!)