Machine Learning Activities at KIT for Accelerators

Andrea Santamaría García, Tobias Boltz, Erik Bründermann, Bastian Härer, Akira Mochihashi, Chenran Xu, Anke-Susanne Müller

AMALEA Final Meeting (17-12-2020)

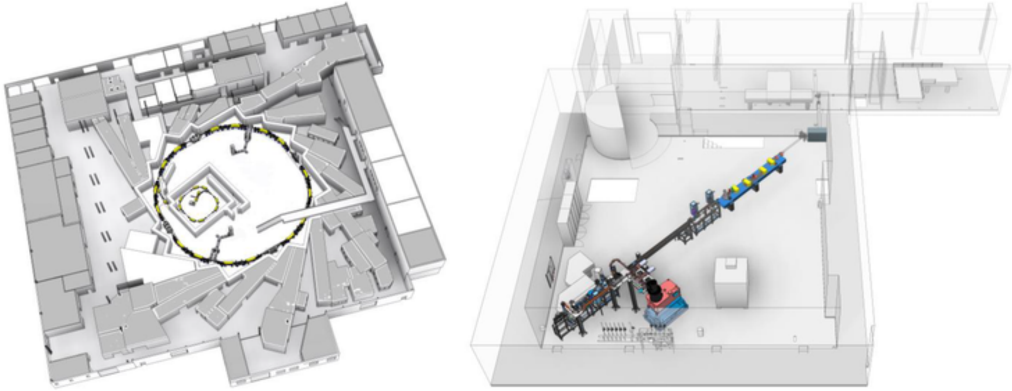

Machine Learning Activities for Accelerators

KARA

Test facility for accelerator physics
Synchrotron light source

- **Timing Modes for Advanced Light Sources: Control of the Micro-Bunching Instability with Reinforcement Learning**
  Tobias Boltz, Bastian Härer, Andrea Santamaría García
- **Bayesian Optimization of the Injection Efficiency**
  Chenran Xu, Akira Mochihashi

FLUTE

Test facility for accelerator physics

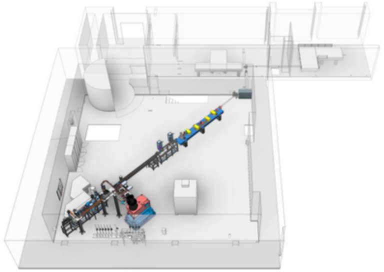

- **Machine Learning Towards Autonomous Accelerators: Control of the longitudinal bunch profile with Reinforcement Learning**
  Andrea Santamaría García, Chenran Xu, Erik Bründermann

Control of the Micro-Bunching Instability with Reinforcement Learning

Laboratory for Applications of Synchrotron Radiation (LAS)

Karlsruhe Institute of Technology

Federal Ministry of Education and Research

Project: **Ti**ming **Mo**des for Advanced Light Sources (TiMo)
BMBF-Verbundforschung

Project partners:
- Helmholtz-Zentrum Berlin
- TU Dortmund

GOAL AT KIT

Development of a longitudinal feedback system to control the micro-bunching instability in short bunch operation mode with Reinforcement Learning (RL) to **tailor Coherent Synchrotron Radiation (CSR) emission**

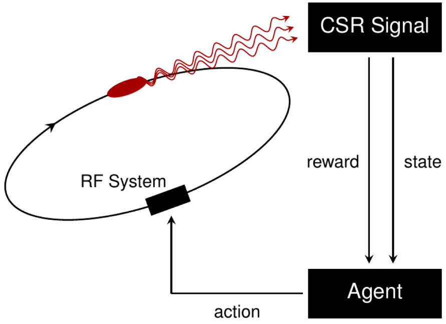

Coherent Synchrotron Radiation (CSR)

There is a strongly increased radiation at wavelengths longer than the bunch length ($\sigma_z < \lambda$) where the synchrotron radiation is emitted coherently (CSR)

CSR = high radiation power → interesting for users! → operation in short bunch mode (low ∝ optics)

microstructures appear making the CSR power fluctuate

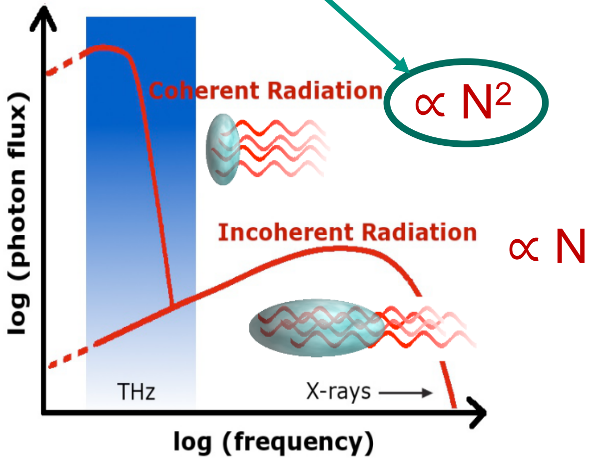

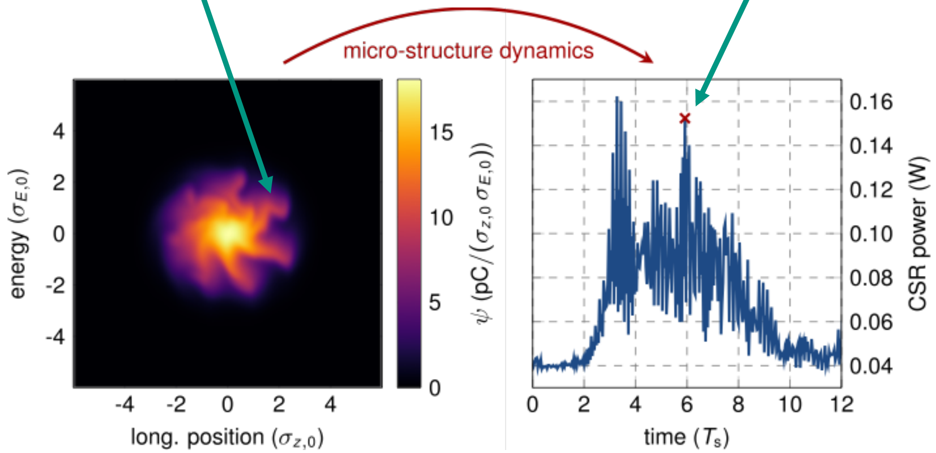

Image courtesy of A.-S. Müller

Simulation code: Parallelized VFP solver ***Inovesa***

T. Boltz et al, MOPGW017, IPAC'19

Coherent Synchrotron Radiation (CSR)

There is a strongly increased radiation at wavelengths longer than the bunch length ($\sigma_z < \lambda$) where the synchrotron radiation is emitted coherently (CSR)

CSR = high radiation power → interesting for users! → operation in short bunch mode (low ∝ optics)

microstructures appear making the CSR power fluctuate

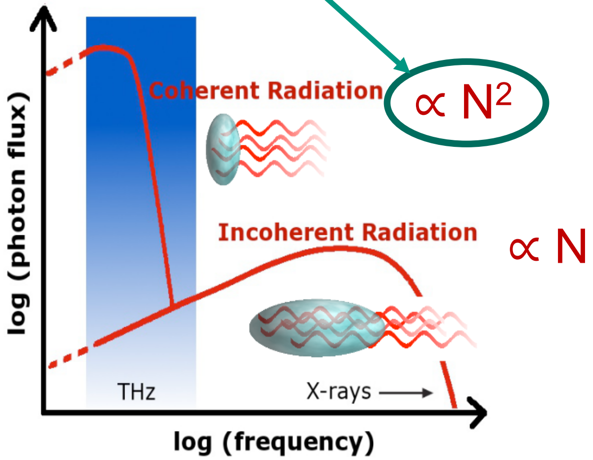

Image courtesy of A.-S. Müller

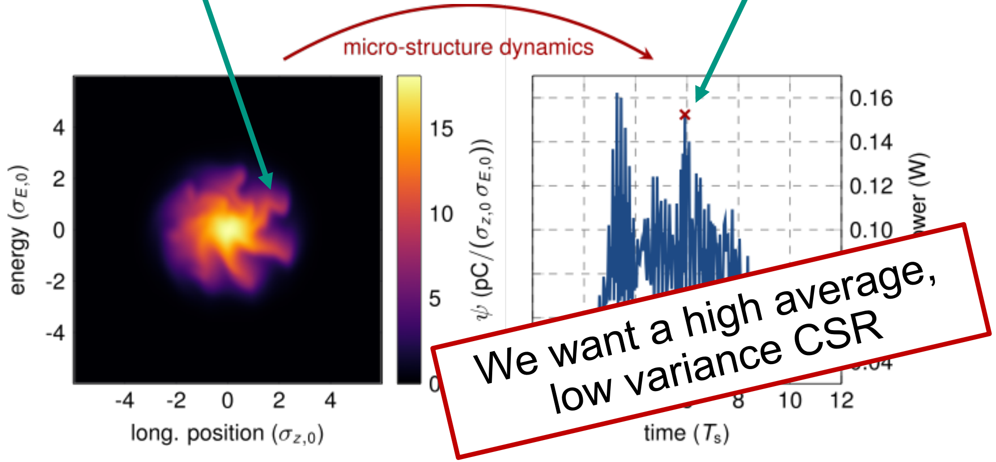

Simulation code: Parallelized VFP solver ***Inovesa***

T. Boltz et al, MOPGW017, IPAC'19

Influencing the Microbunching Instability

The **CSR self-interaction** contributes to the effective potential that the beam is subjected to, and is continuously changing during micro-bunching dynamics

$$V_{eff}(q) = V_{RF}(q) + V_{CSR}(q)$$

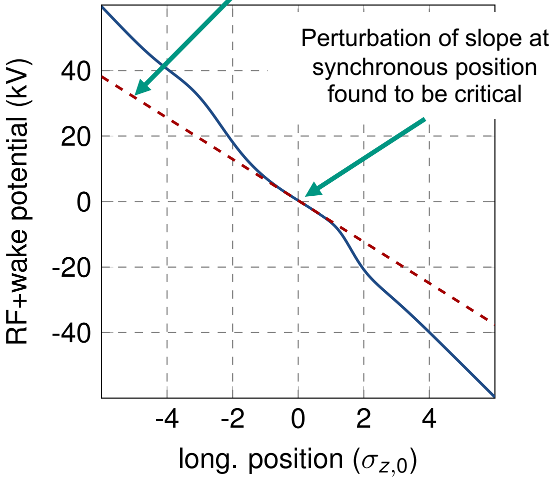

Images courtesy of T. Boltz

Influencing the Microbunching Instability

Karlsruhe Institute of Technology

The **CSR self-interaction** contributes to the effective potential that the beam is subjected to, and is continuously changing during micro-bunching dynamics

$$V_{eff}(q) = V_{RF}(q) + V_{CSR}(q)$$

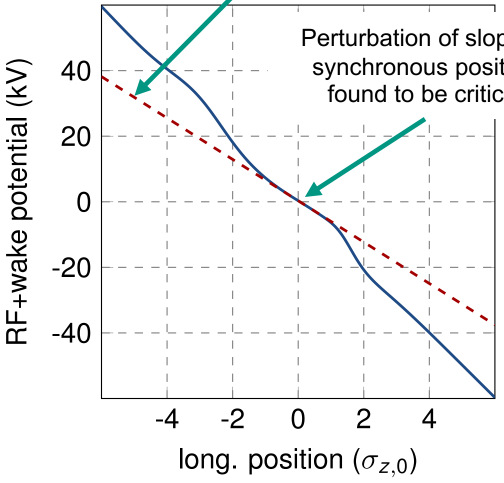

IDEA: compensate the effect of the CSR perturbation and go back to the original restoring force provided by the accelerating voltage by **modulating the RF voltage (amplitude)**

$$V_{RF} = \hat{V}(t)\sin(2\pi f_{RF} t)$$

$$\hat{V}(t) = \hat{V}_0 + A_{mod}\sin(2\pi f_{mod} + \varphi_{mod})$$

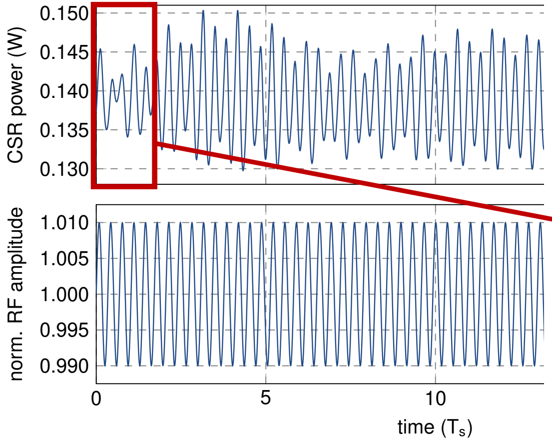

Trying the idea in simulation with a constant modulation:

Initial damping, but quickly out of sync…we need **dynamic control**!

Images courtesy of T. Boltz

Testing the idea with manual control

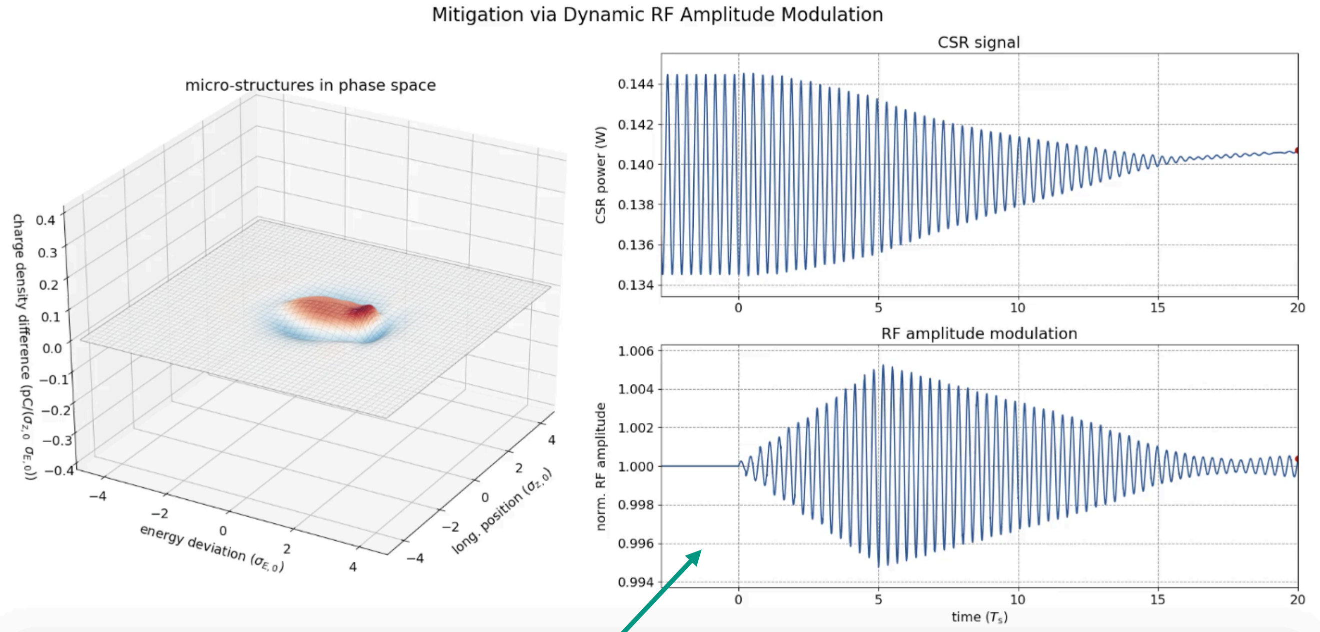

Courtesy of T. Boltz

$$V_{RF} = \hat{V}(t)\sin(2\pi f_{RF} t) \qquad \hat{V}(t) = \hat{V}_0 + A_{mod}\sin(2\pi f_{mod} + \varphi_{mod})$$

Testing the idea with manual control

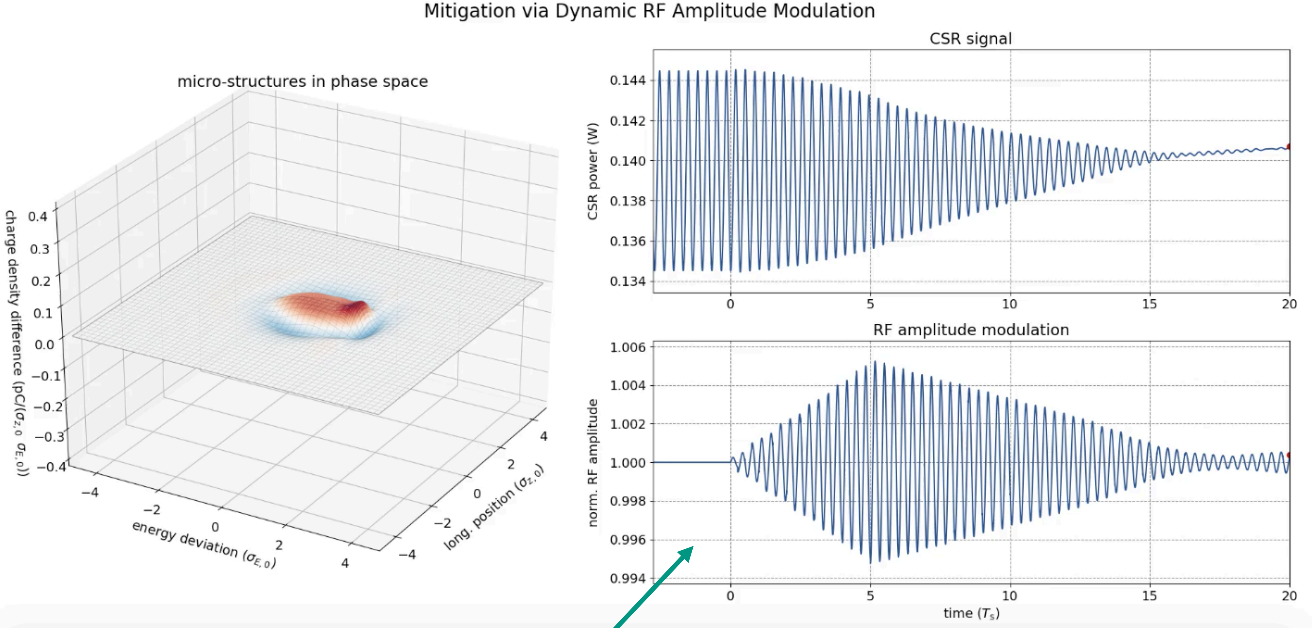

Courtesy of T. Boltz

$$V_{RF} = \hat{V}(t)\sin(2\pi f_{RF} t) \qquad \hat{V}(t) = \hat{V}_0 + A_{mod}\sin(2\pi f_{mod} + \varphi_{mod})$$

High average, low variance CSR. So far so good…now to Reinforcement Learning!

Applying Reinforcement Learning to our case

Action
Modulation of the RF amplitude

$\hat{V}(t) = \hat{V}_0 + A_{mod}\ \sin(2\pi f_{mod} + \varphi_{mod})$

Simulation observable (state definition):
Charge distribution

Input: (256x256) matrix + (5x1) feature vector

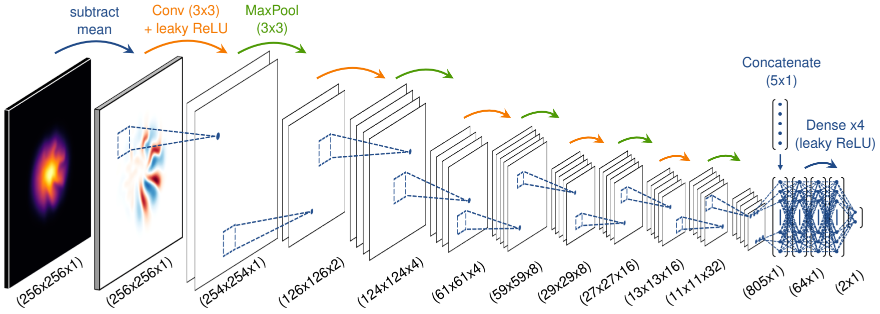

Images courtesy of T. Boltz

Reward
R = $\mu_{CSR} - w\ \sigma_{CSR}$ where w is a weight

Could we improve the reward definition?

Experimental observable (state definition):
CSR signal

We can consistently and reliably measure the THz emission thanks to KAPTURE II and provides information about the micro-bunching dynamics. Input: (8x1) feature vector

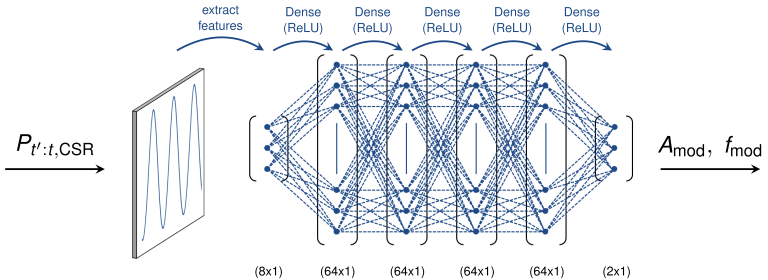

Training the agents (CSR signal)

Algorithm: PPO

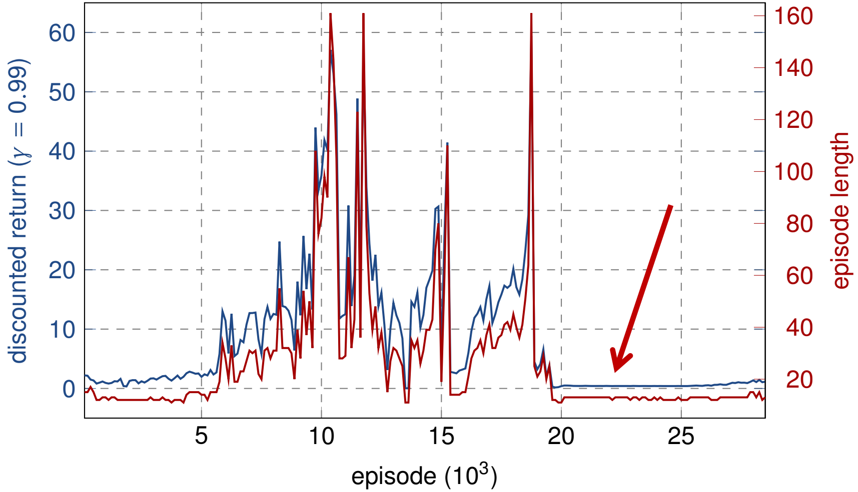

Algorithm: DDPG

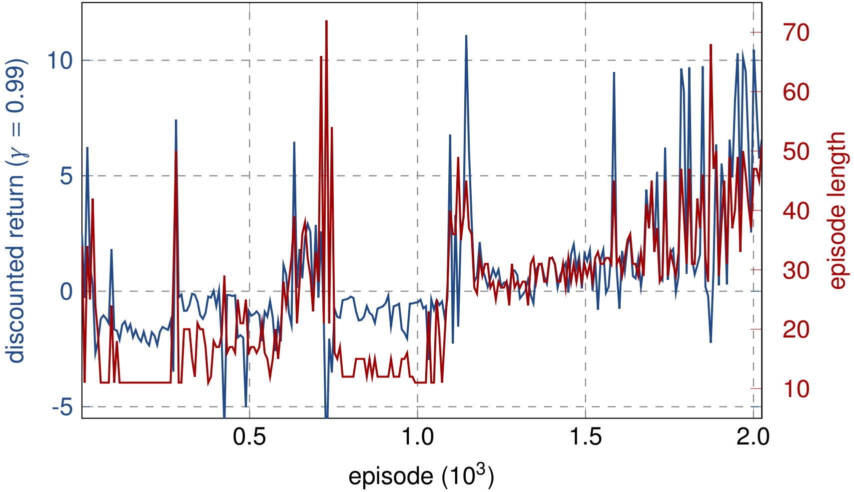

- The NNs are updated at every episode
- Performance (reward) drops to zero after a certain number of episodes with the PPO agent
- The results on the next slide show peak performance / the best deployable agent

Images courtesy of T. Boltz

Algorithm: PPO

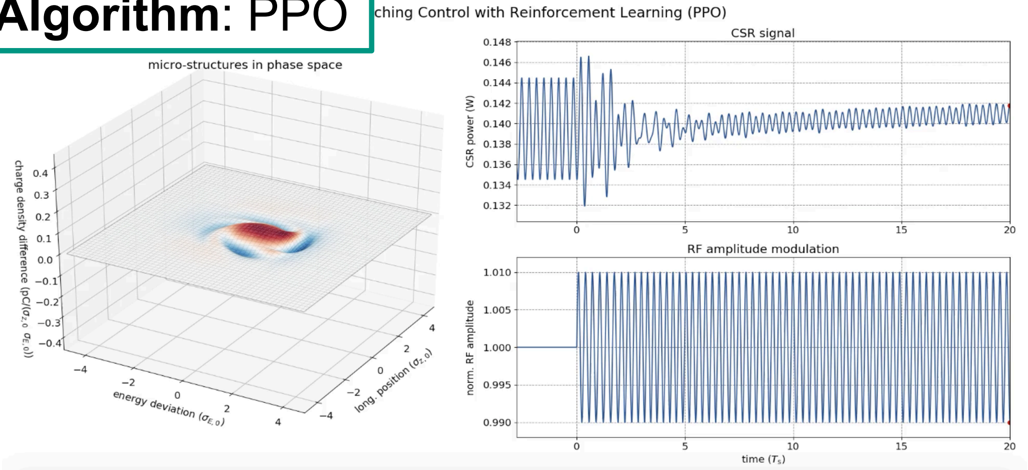

Algorithm: DDPG

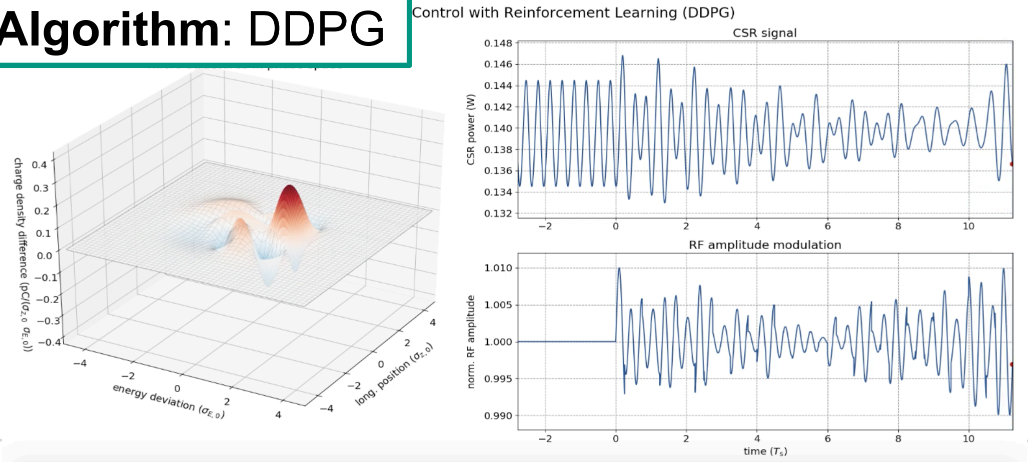

Courtesy of T. Boltz

Algorithm: PPO

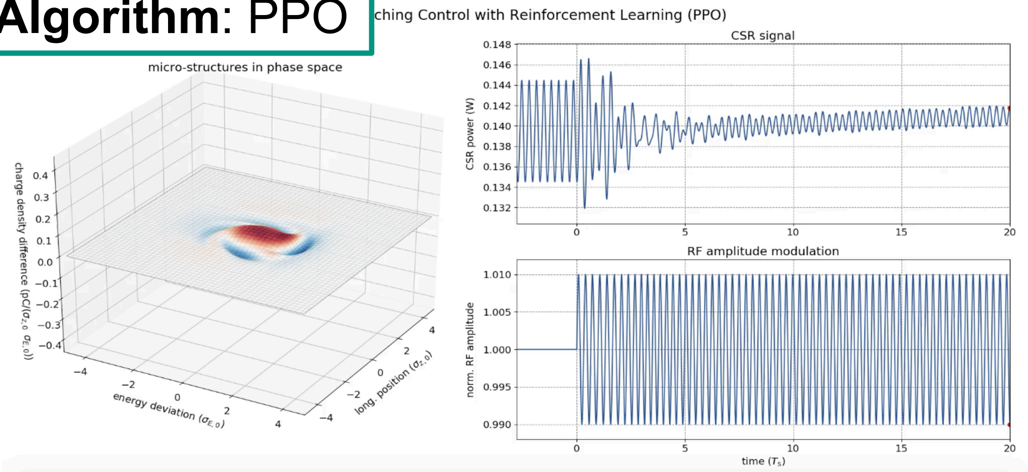

Quickly finds a stable regime and the CSR emission is considerably improved

→ The exploration noise is built into the algorithm and it's naturally reduced at low rewards

Algorithm: DDPG

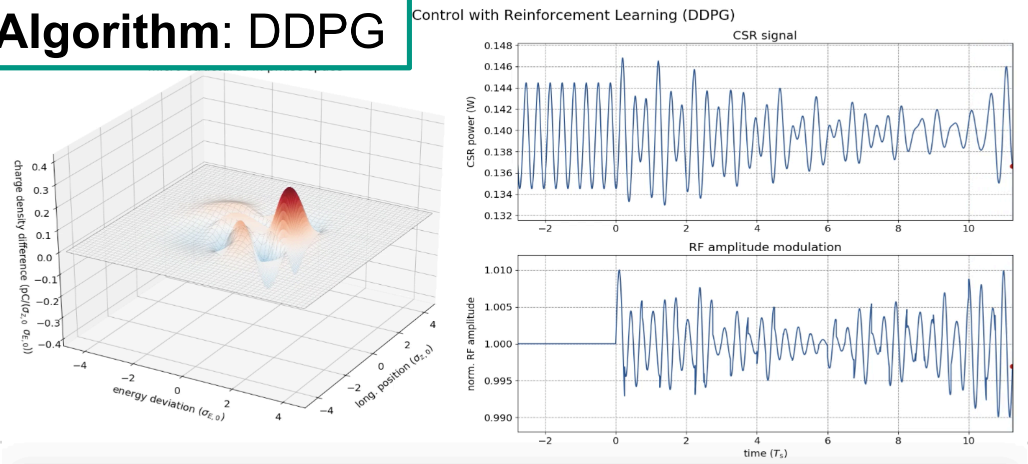

Episode is terminated because it does not improve

→ Exploration noise = Ornstein-Uhlenbeck

Can we improve performance by adjusting the exploration noise?

Courtesy of T. Boltz

Evolution of the actions with time (PPO)

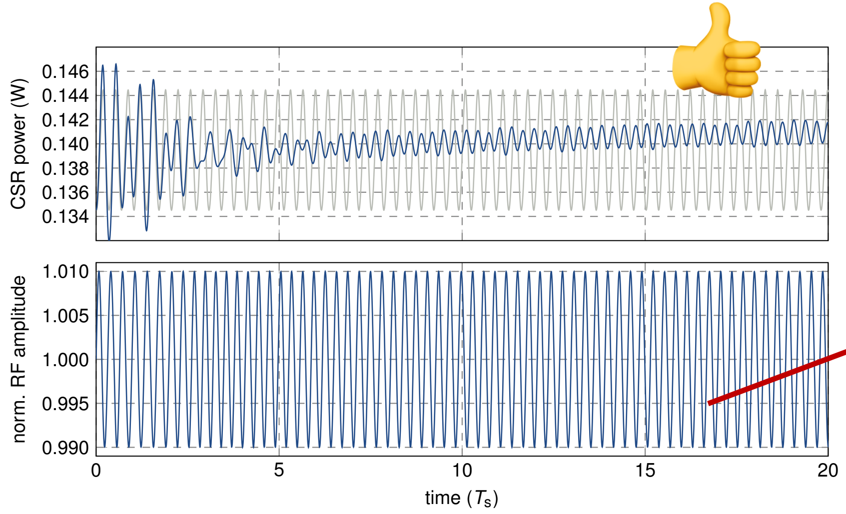

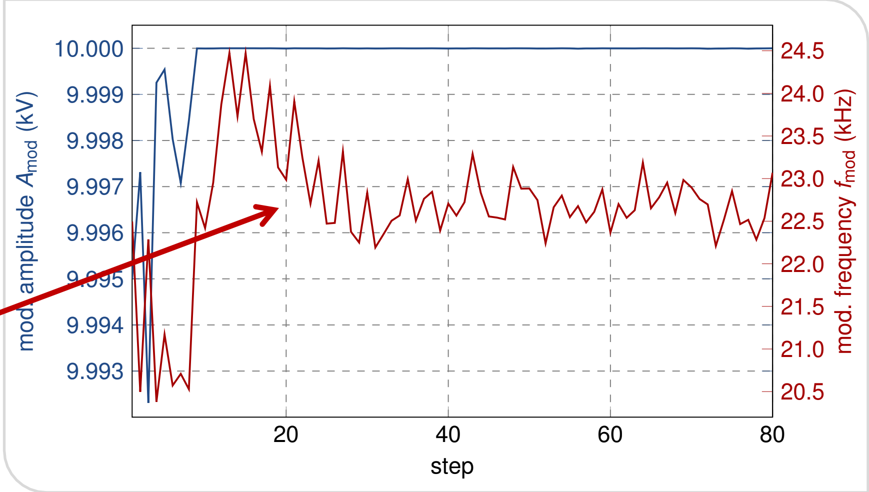

It cannot be appreciated in the RF voltage modulation by the naked eye but…

Small changes in the modulations frequency stabilize the CSR emission

Images courtesy of T. Boltz

1 step = 0.25 synchrotron periods (chosen small enough for the agent to be able to react to the changing micro-structure dynamics)

In practice: we need hardware!

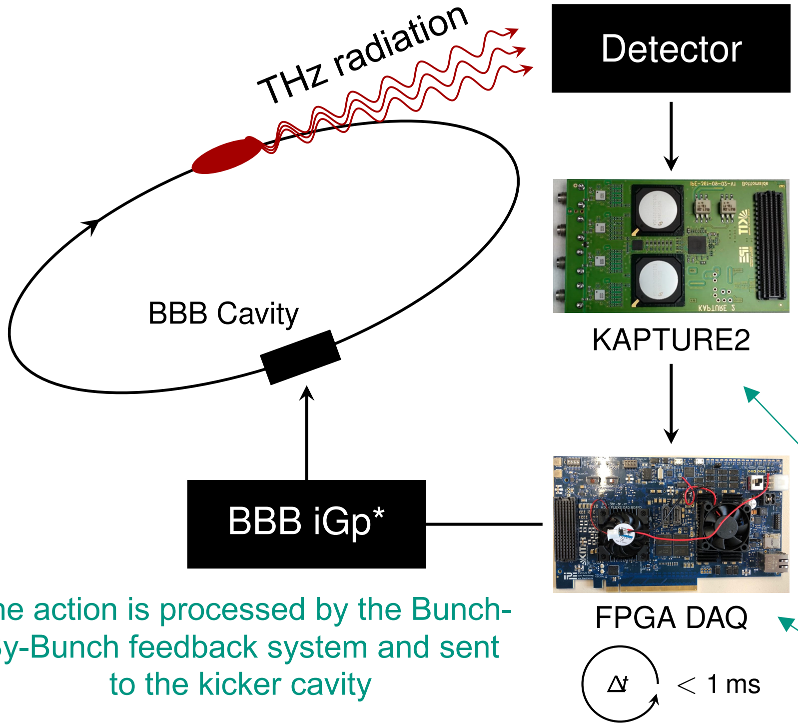

The action is processed by the Bunch-By-Bunch feedback system and sent to the kicker cavity

Cooperation between:

MT ARD + MT DTS

MT-DTS@KIT: Michele Caselle, Andreas Kopmann

Collaboration at KIT with:

- Institute of Data Processing and Electronics (IPE)
- Institute for Anthropomatics and Robotics (IAR)

Edmund Blomley, Miriam Brosi, Michele Caselle, Timo Dritschler, Melvin Klein, Christoph Pohl, Weijia Wang

Extraction of the frequency of the instability

RL algorithm running on the FPGA receives the observation and decides on an action (which RF modulation)

[T. Boltz, W. Wang et al, TUCPL06, ICALEPCS 2019]

Control of the Micro-Bunching Instability with Reinforcement Learning

ACHIEVEMENTS

- Manual control of the microbunching instability in simulation with the Vlasov-Fokker- Planck solver Inovesa through RF amplitude modulation
- Research of different RL algorithms: DDPG, TD3, PPO and SAC
- First experimental studies: capability to influence the microbunching with RF amplitude modulations in a closed loop demonstrated at KARA

Control of the Micro-Bunching Instability with Reinforcement Learning

OPEN QUESTIONS

Benchmark Problem (Learning Process)

- Reproducibility (fully deterministic computation)
- Choosing suitable exploration noise
- Instability of the agent’s learning process
- Reward function and termination condition
- Choosing the RL algorithm (e.g., on-/off-policy learning)

Extension of Control

- Feasibility of control based on solely the CSR
- Signal generalization to different bunch currents
- Generalization to different machine settings

Bayesian Optimization of the Injection Efficiency

- We want to improve the injection rate from the booster to the storage ring
- Manual trial-and-error tuning is time consuming, depends on the operator's experience, and can easily get stuck in local optima
- **Bayesian optimization has been successfully implemented in other facilities and can converge to the global optimum**

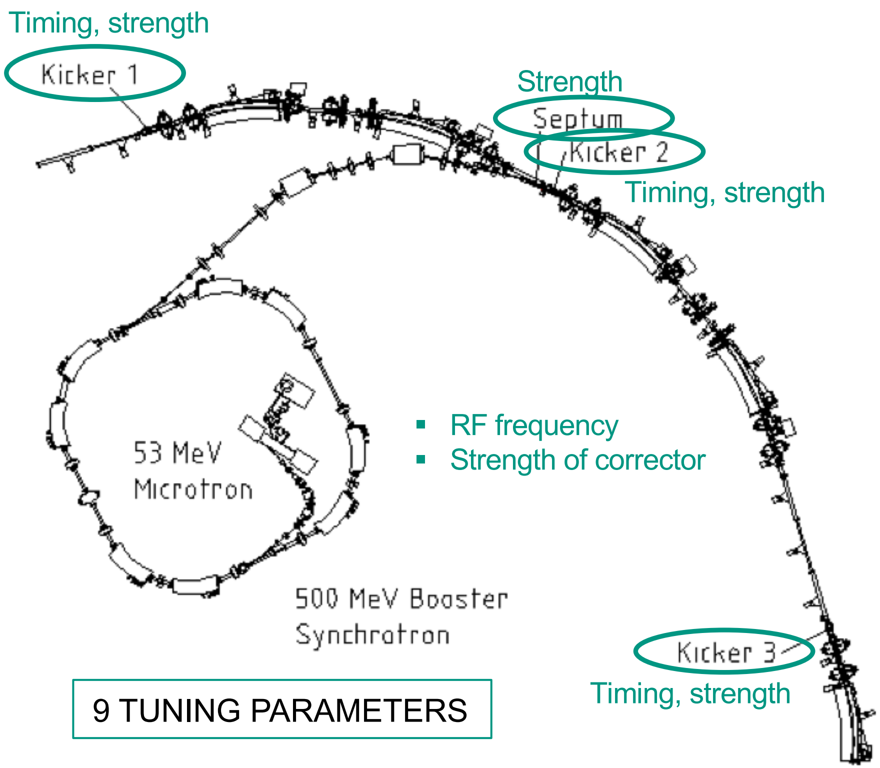

9 TUNING PARAMETERS

C. Xu, Master thesis, KIT, to be published

Preparing and testing the model

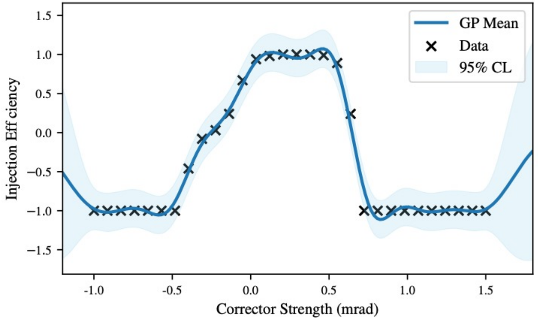

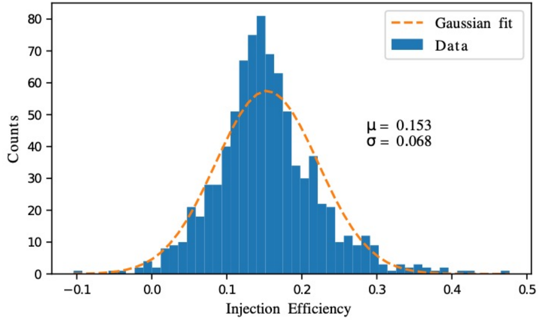

Determination of the Gaussian Process (GP) hyperparameters (signal variance, characteristic length-scales, and noise variance) through measurements. **A GP model is trained to recreate the observed structure.**

Testing the model in simulation:

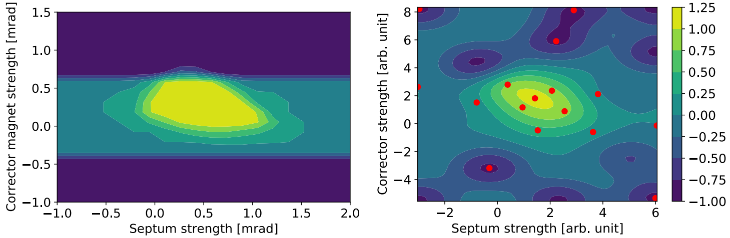

- Left: grid scan of parameter space (~1h)
- Right: Bayesian Optimization (3 min)

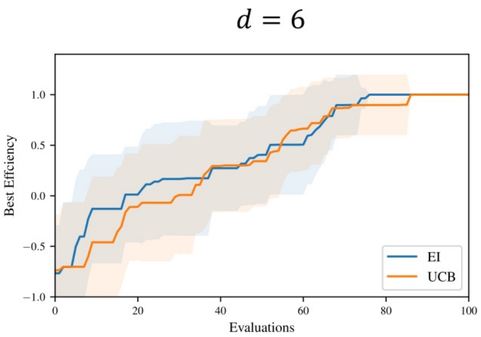

Injection efficiency averaged over 10 runs for two different acquisition functions

Plots courtesy of C. Xu

Correction for beam lifetime dependency with stored current

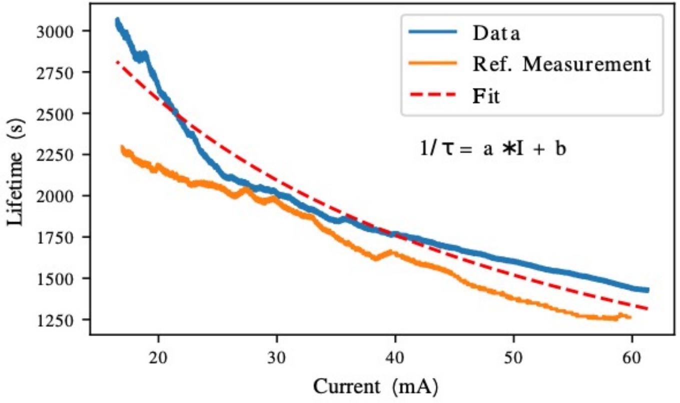

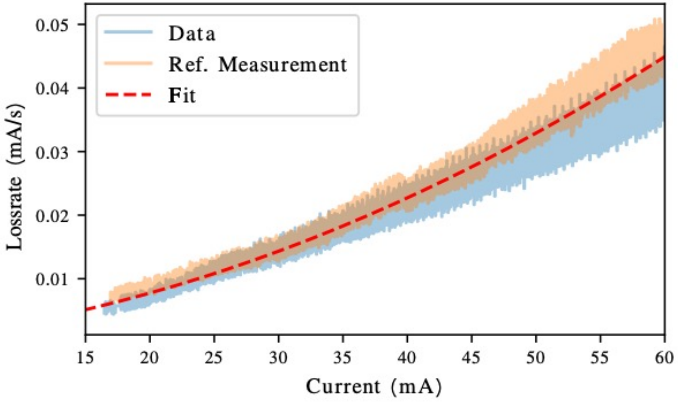

- The beam lifetime depends on the storage ring current (Touschek scattering)
- If left unaccounted for, the injection efficiency values will decrease independently of the optimization algorithm

Plots courtesy of C. Xu

Bayesian Optimization of the Injection Efficiency

ACHIEVEMENTS

- Development of a Bayesian Optimization algorithm with Gaussian Processes with three different acquisition functions and up to 9 input parameters

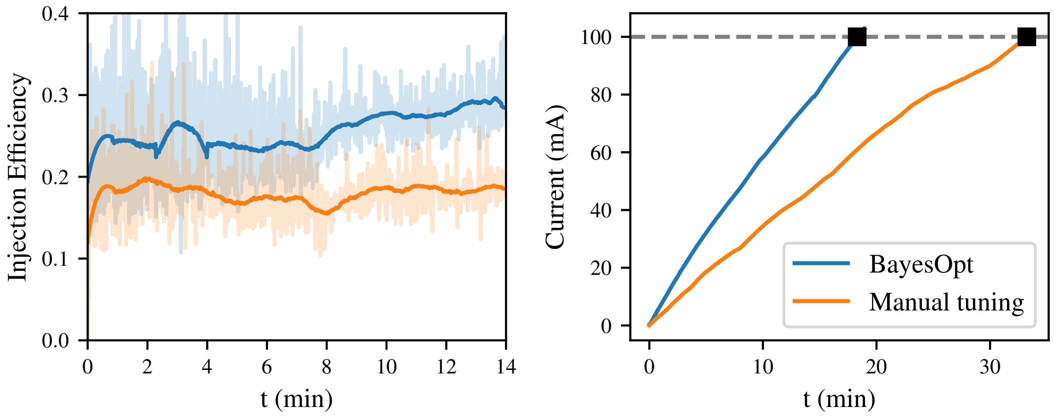

Optimized injection is roughly two times faster!

Plots courtesy of C. Xu

Machine Learning Toward Autonomous Accelerators

HELMHOLTZAI | ARTIFICIAL INTELLIGENCE COOPERATION UNIT

The "Autonomous Accelerator" is a two-year project funded by Helmholtz AI, one of the five platforms initiated by the Helmholtz Information and Data Science Incubator

Start: September 2020

[Press release]

- **DESY PI (coordinating):** Annika Eichler
- **KIT PI:** Erik Bründermann
- **Research associate:** Andrea Santamaría García
- **Doctoral researcher:** Chenran Xu
- **ARES contact:** Florian Burkart

Goals & Challenges

Goals

- Control of the longitudinal bunch profile (to transversal bunch profile control, to automatic start-up)
- Transfer Learning: apply algorithms to two similar facilities (ARES and FLUTE)

Challenges

- High-dimensional, continuous state and action spaces
- Low repetition rate
- Nonlinear (collective) effects
- (Sub)-femtosecond requirements on bunch duration

Photoinjector
- Laser pulse length
- Laser pulse shape
- Laser spot size
- Laser spot position
- Magnet current

RF gun
- RF amplitude
- RF phase
- Beam arrival
- Time

Traveling wave structure
- Temperature profile
- Amplitude
- Phase

Magnetic chicane
- R56 travel difference of average beam energy

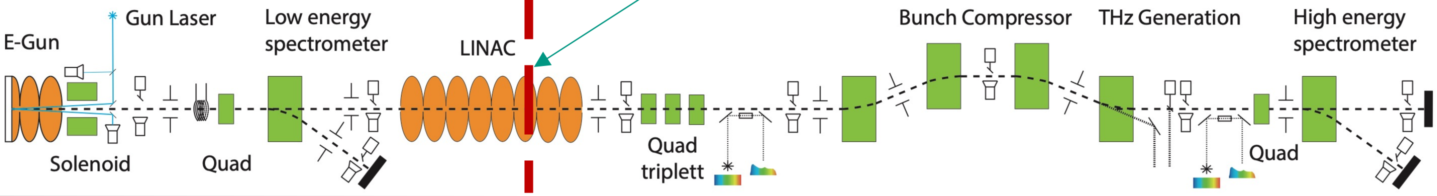

End of 2020 Outlook

Karlsruhe Institute of Technology

FINISHED

Bayesian Optimization of the Injection Efficiency

ONGOING

Timing Modes for Advanced Light Sources: Control of the Micro-Bunching Instability with Reinforcement Learning

STARTING

Machine Learning Towards Autonomous Accelerators: Control of the longitudinal bunch profile with Reinforcement Learning

Let's keep in touch in the international machine learning for accelerators community!

Let me know if you would like to show your work:
andrea.santamaria@kit.edu

The **O**ne **W**orld charged partic**L**e acc**E**lerator (**OWLE**) Colloquium & Seminar Series

Given the impossibility of travel during the COVID-19 crisis the (OWLE) seminar series was established as an inter-institutional global online colloquium and seminar(s).

The **OWLE-Colloquium** is aimed at giving researchers a platform to share research and development results of very broad interest.

The **OWLE-ML seminar series** has a topical focus on machine learning and experimental demonstration of AI-ML.

Colloquium talks are held via Zoom once a month on the first Tuesday at 1:30 PM EDT (19:30 CEST, 10:30 AM PST).

Seminars are held every second and last Tuesdays at 2:30 PM EDT (20:30 CEST, 11:30 AM PST).

https://sites.google.com/view/owle/

Observation vector based on the CSR signal

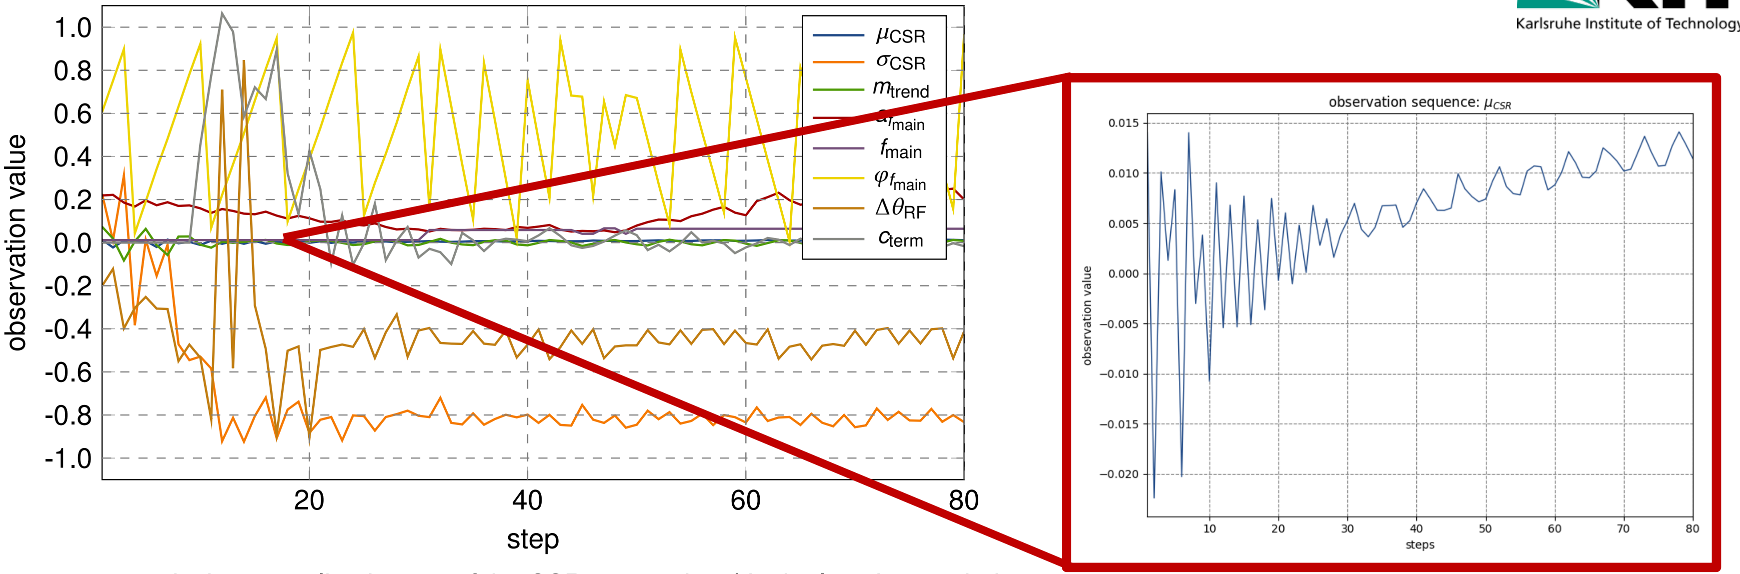

- μ_{CSR} is the normalized mean of the CSR power signal in the last time period.
- σ_{CSR} is the normalized standard deviation of the CSR power signal in the last time period.
- m_{trend} is a slow trend of the CSR power signal
- $a_{f_{main}}$ is the amplitude of the main frequency in the Fourier transformed CSR signal.
- f_{main} is the main frequency in the Fourier transformed CSR signal.
- $\varphi_{f_{main}}$ is the phase of the main frequency in the Fourier transformed CSR signal.
- $\Delta\theta_{RF}$ is the relative phase between the CSR signal and the applied RF signal (amplitude modulation).
- c_{term} models the termination condition (difference between the last reward and the one 10 steps prior).