RL4AA'23: meet & greet session

Dr. Sabrina Appel

2008 starting working with accelerators, PhD 2011

Since 2014: Automated optimization with numerical + nature optimizer

s.appel@gsi.de, http://web-docs.gsi.de/~sappel/index.html

RL4AA'23: meet & greet session

Dr. Sabrina Appel
2008 starting working with accelerators, PhD 2011
Since 2014: Automated optimization with numerical + nature optimizer

s.appel@gsi.de, http://web-docs.gsi.de/~sappel/index.html

- EUROLABS WP5 Task 3:
- scientific staff member for 3 years at GSI

➢ The aim is to find the optimum framework for usable ML algorithms for accelerator problems.

➢ **Open Data**: Share algorithms and application across facilities + Beam optimization algorithms and ML library

➢ Based on GEOFF (CERN gitlab)

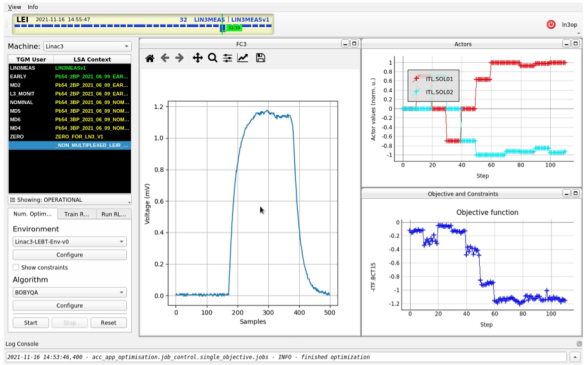

GSI topics

Injection optimization

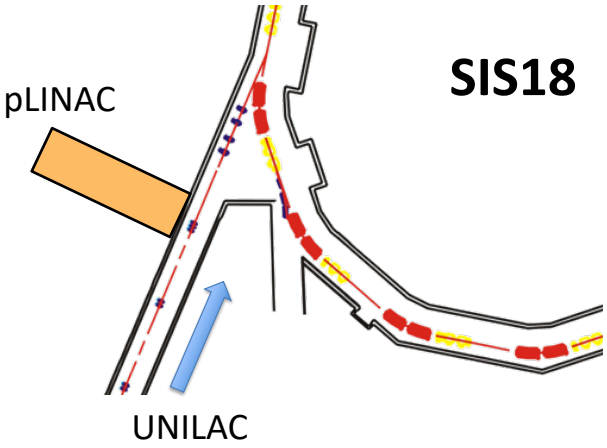

MTI optimization with Bayesian and maybe with RL

Space charge + loss-induced vacuum degradation
Septum protection: Uranium ions can destroy wires

S. Appel et al: International Journal of Modern Physics A Vol. 34, 1942019, (2019),
S. Appel et al: Nucl. Instrum. Methods A 852 (2017), pp. 73-79,
S. Appel et al., J. Phys. Conf. Ser 1350 (2019) 012104

Optimize transport of fragmentation beams

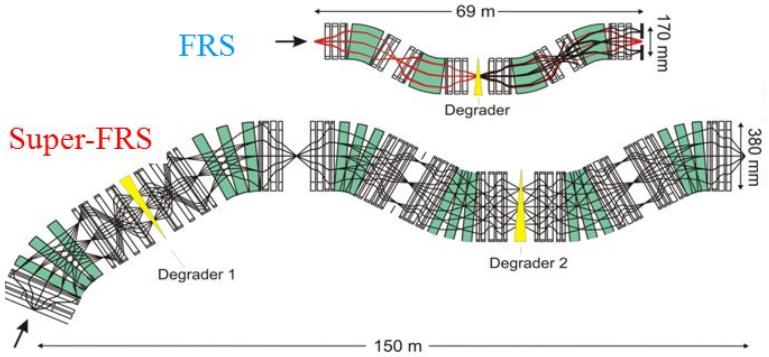

FRS/Super-FRS (GSI/FAIR)

- Optical coefficient tuning
- Detector/degrader calibration
- Optimization of spectrometer parameters

CERN topics

Ion Sources optimization:

FEBIAD

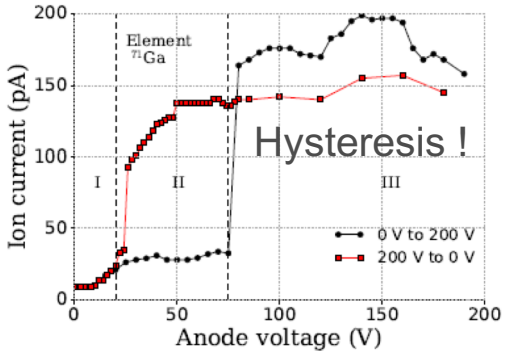

Y. Martinez Palenzuela, Thesis (https://lirias.kuleuven.be/handle/123456789/636675)

and RILIS

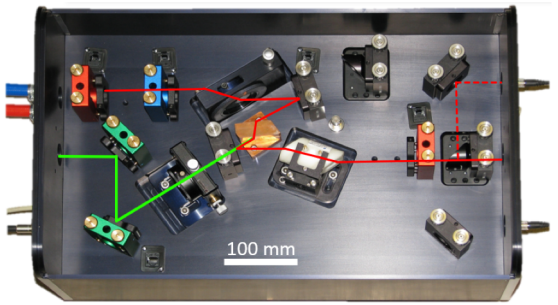

Valentin Fedosseev et al 2017 J. Phys. G: Nucl. Part. Phys. 44 084006, Doi: 10.1088/1361-6471/aa78e0

- Laser cavity has >5 parmeters
- Optomechanics has hysteresis !

ISODLE beam line

ISOLDE offline 2:
Testing optimizer and RL algorithms
+ and AWAKE beam line

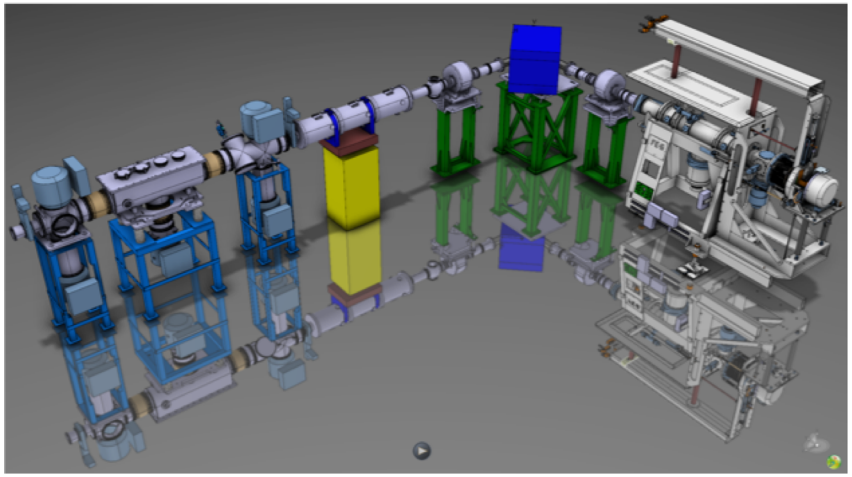

and more …

Awal Awal

RWTH Aachen University
GSI

Phd. Student 2021-present
a.awal@fz-juelich.de
https://github.com/meawal

Involvement in RL:

- Optimization of the Injection Beam Line (IBL)
- Autonomous injection optimization of the beam into the Cooler Synchrotron COSY

Interests in RL:

- Sim2real
- Hierarchical RL
- Derivative-free optimization
- Sample efficiency
- Multi-agent RL

Cooler Synchrotron COSY

- Designed for protons and deuterons
  - Polarized and unpolarized proton beams in energy range upto 2.7 GeV
  - Deuteron beams in energy range upto 2.1 GeV
  - Stochastic cooling
- Injection beam line (IBL) is the transfer line from the cyclotron (JULIC) to COSY
  - Transferring negatively charged protons and deuterons throughout the IBL
  - 45 MeV protons & 76 MeV deuterons
- Electrons are stripped at the injection point through a stripping foil
- IBL length ≈ 94m. Operated manually through 15 quadrupoles and 27 steerers

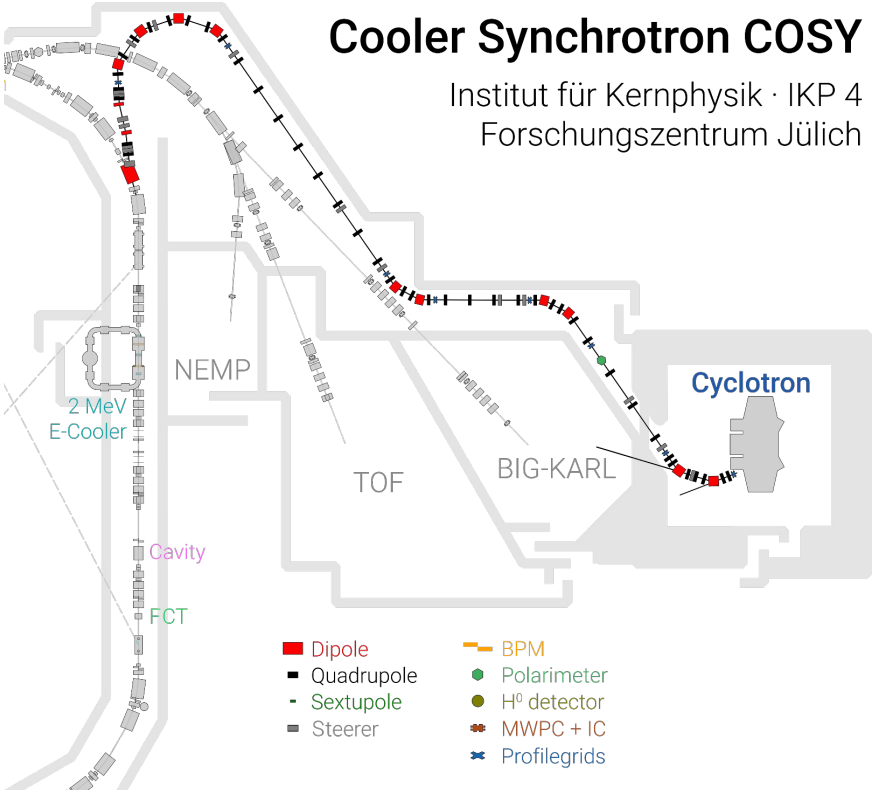

RL in Simulation and Real Environment

Training

Real world

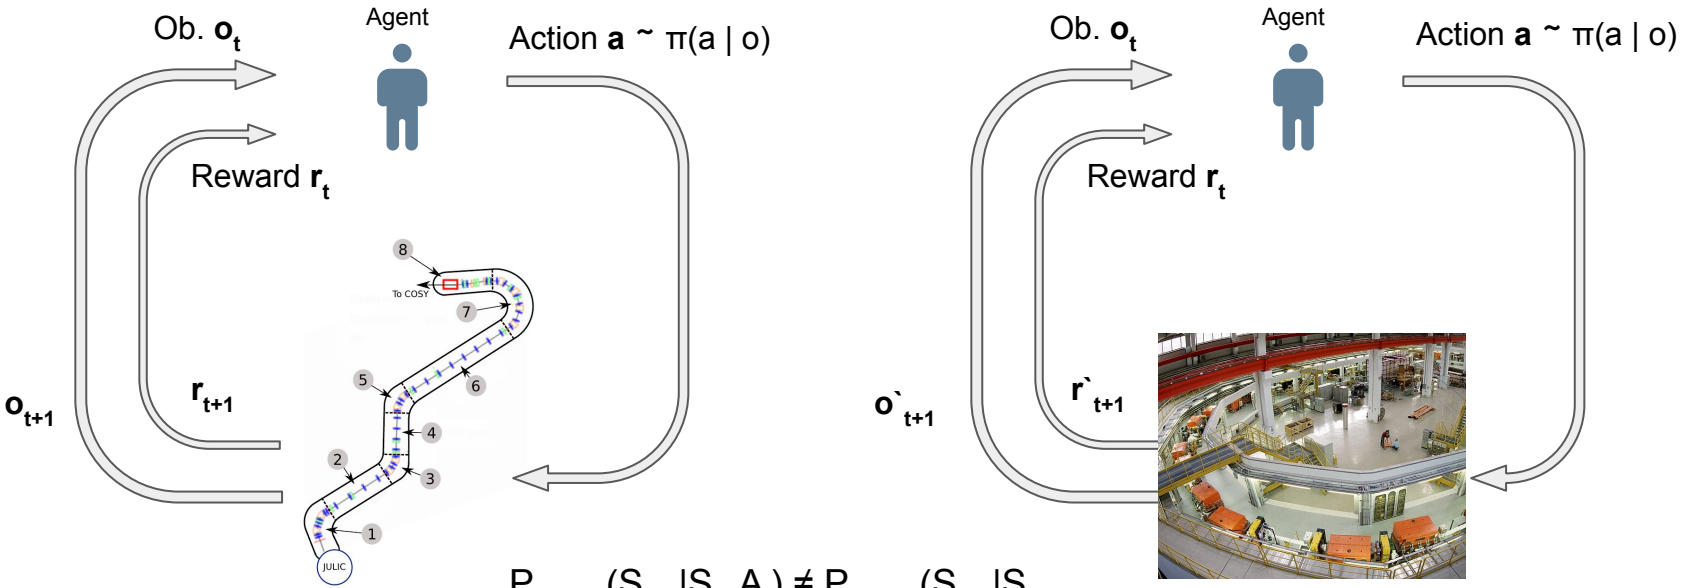

$P_{MADX}(S_{t+1}|S_t, A_t) \neq P_{COSY}(S_{t+1}|S_t, A_t)$

IBL Optimization

5000 particles
Step 0
Efficiency: 0%

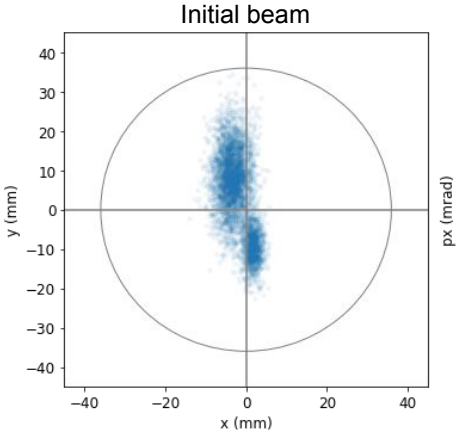

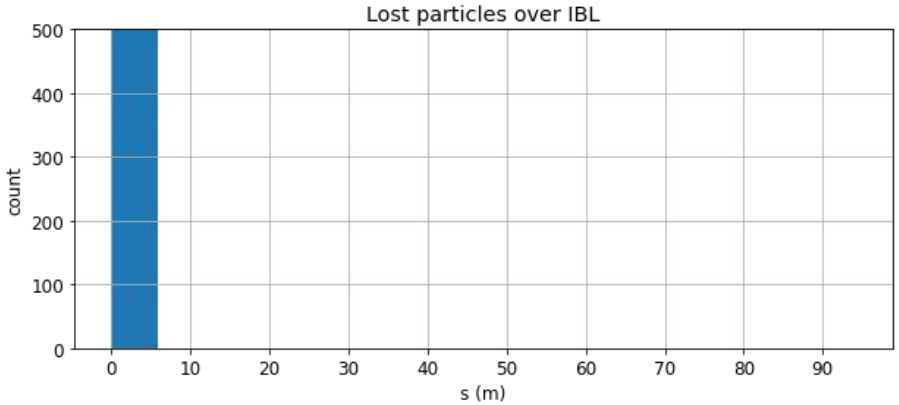

5000 particles
Step 32
Efficiency: 89.9%

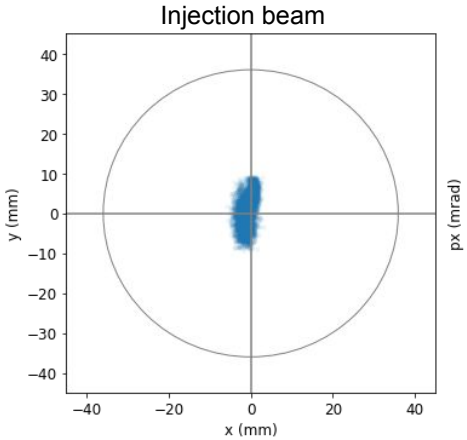

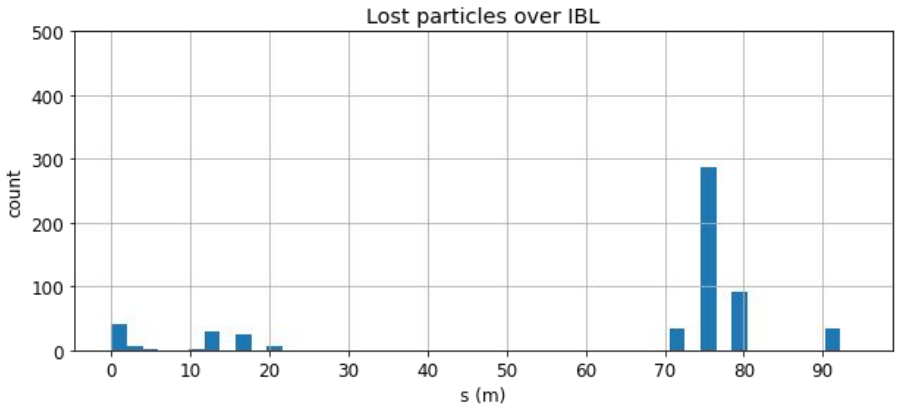

Autonomous Injection

Before

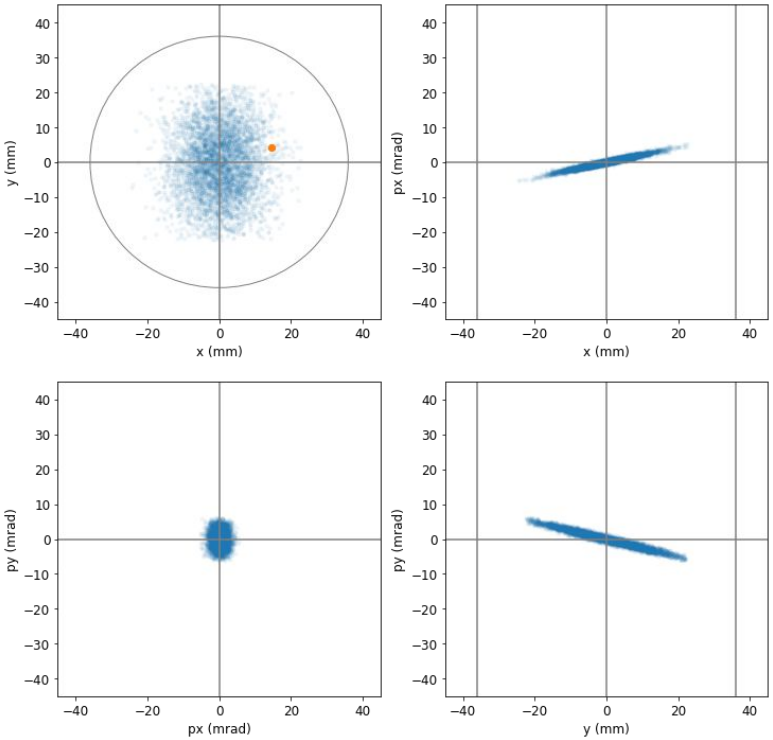

After

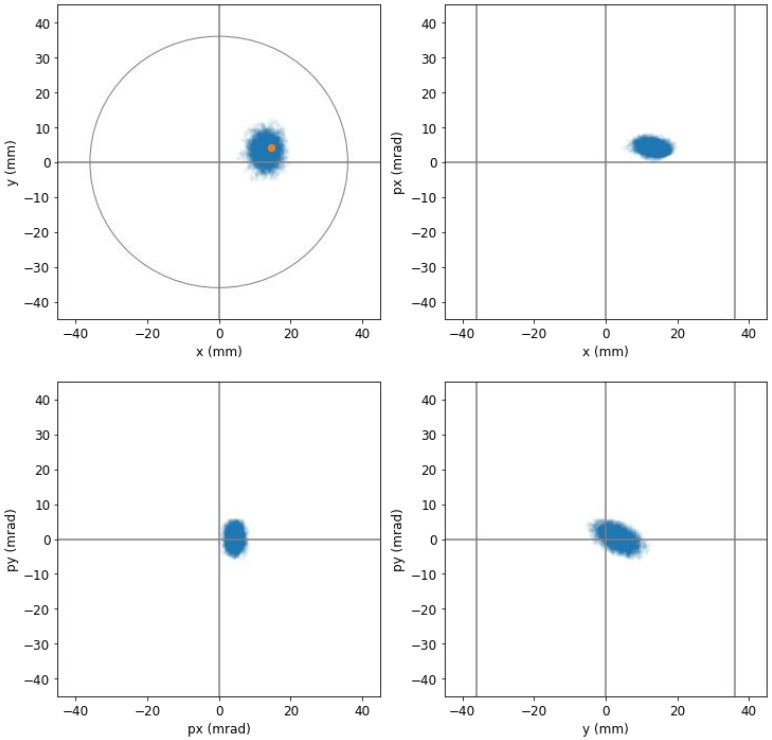

Niky Bruchon

CERN
SY-RF-BR

niky.bruchon@cern.ch
gitlab.cern.ch/nbruchon
researchgate.net

2017 – 2021
Doctoral student
RL methods for FERMI FEL
Performance improvement
Università degli Studi di Trieste
Elettra Sincrotrone Trieste

2021 – 2021
Researcher
BPM anomaly detection
KIT Karlsruher Institut für Technologie

2021 – present
Fellow
Automation of RF procedures in the SPS
CERN

RL activities:

- Investigation of model-free and model-based RL algorithms to improve the FERMI FEL performance
- Custom OpenAI-Gym environment for developing and testing

Interests:

- Autonomous accelerator (one button machine)
- Real-world and online optimization problems
- Reinforcement Learning applications
- Simulation to real word transfer
- Sample efficiency

RL Studies on RL Techniques to Improve FERMI FEL Performance

SY
RF
BR

Problem: transverse overlapping of the electron and laser beams to optimize the light radiation intensity

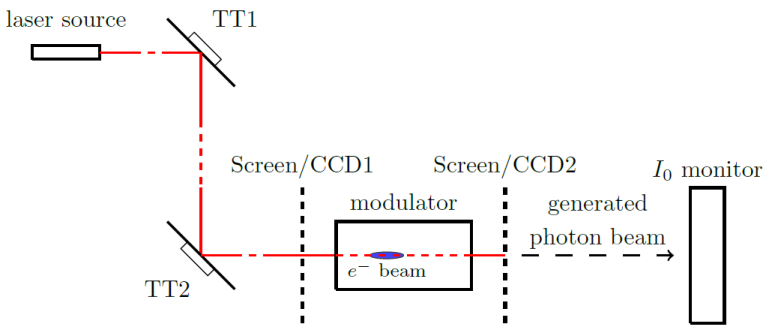

Simple block scheme

Investigated Algorithms

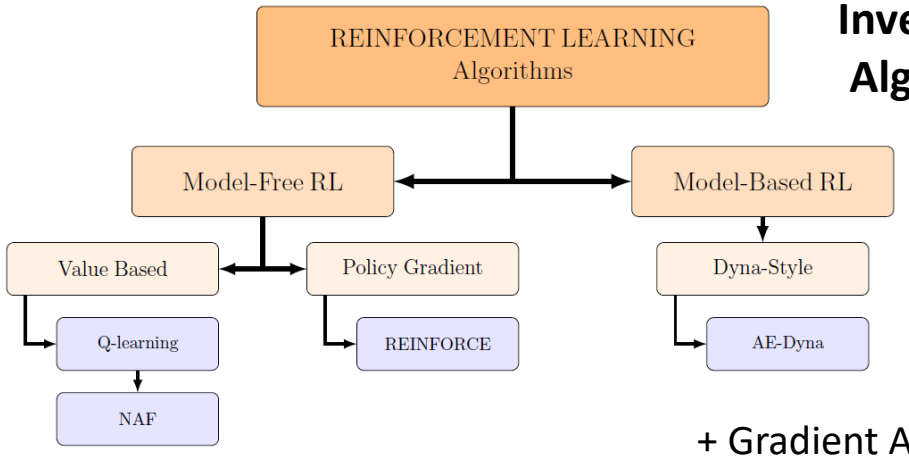

+ Gradient Ascent
+ iLQR

Performance in the training/identification phase

Algorithm	Data points
Q-learning	3128
NAF	1074
NAF2	824
AE-Dyna with TRPO	450
AE-Dyna with SAC	500
Gradient Ascent	1024
iLQR	1024

Performance in the test phase

Algorithm	Episode length (mean)	Final intensity (mean)
Q-learning	11.28	-
NAF	2.56	1.0019
NAF2	2.64	0.9995
AE-Dyna with TRPO	4.46	1.0150
AE-Dyna with SAC	3.28	1.0427
Gradient Ascent	3.82	0.9911
iLQR	2.54	1.0019

Conclusions

Two different tasks have been successfully faced:

I. Attainment of the optimal working point
II. Recovery of the optimal working point after machine drifts

[1] Feasibility investigation on several reinforcement learning techniques to improve the performance of the FERMI free-electron laser

TECHNISCHE UNIVERSITÄT DARMSTADT

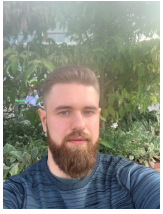

Conrad Caliari, M.Sc.

- graduated in accelerator physics
- doctoral student at TU Darmstadt

GSI FAIR

Involvement in RL

- B.Sc. thesis at S-DALINAC / TU-Darmstadt
  - applied to beam focusing and steering in transfer line

Interested in

- Setup & continuous control with RL
- Multi-turn injection into synchrotron
- Identification of magnetic field errors

S-DALINAC @ TU-Darmstadt

TECHNISCHE UNIVERSITÄT DARMSTADT

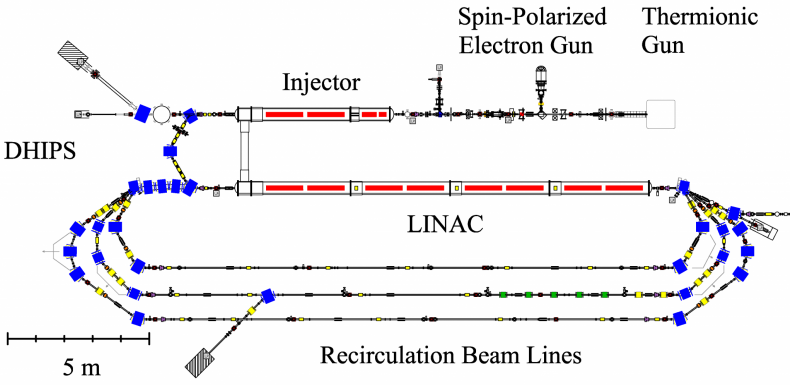

S-DALINAC @ TU-Darmstadt

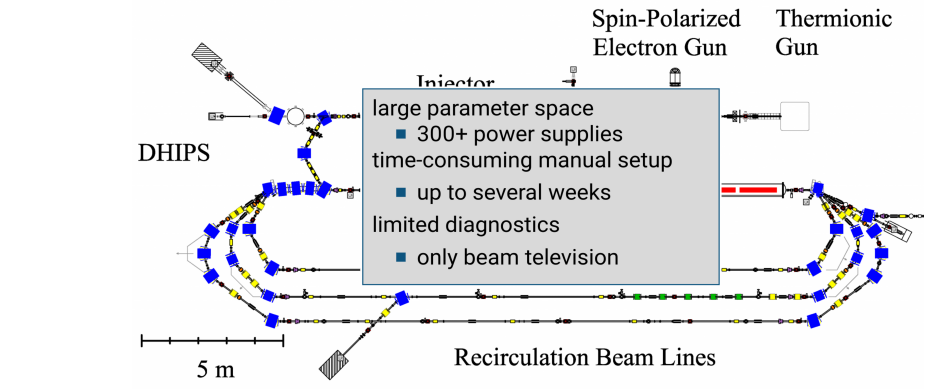

Beam Focusing and Steering with RL

TECHNISCHE UNIVERSITÄT DARMSTADT

Goal

- steer beam to center of target, minimize beam spot

Task

- continuous state space
  - magnet currents
  - location beam spot / rms-size
- continuous action space
  - adjust magnet current

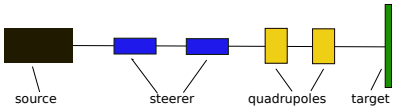

- position and slope of beam randomly initialized
- 6 degrees of freedom
- reward
  - + beam spot moved towards target center
  - + reduction of rms-size
  - — constant penalty per step

Results

TECHNISCHE UNIVERSITÄT DARMSTADT

Implemented algorithms

- Deep-Q-Networks (DQN)
- REINFORCE
- Advantage-Actor-Critic (A2C)

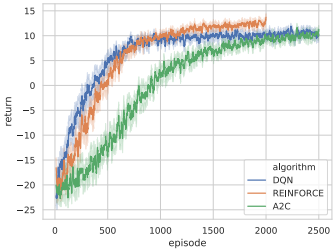

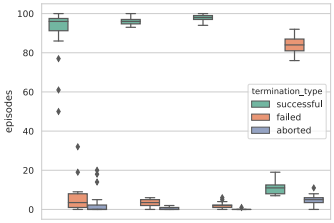

Karlsruhe Institute of Technology

IBPT
Institute for Beam Physics and Technology

PhD Student Felipe Donoso

Background: Electrical Engineering - Automatic Control - Smart Machines

RL4AA'23 workshop | 20th February 2023

Vlasov-Fokker- Planck

$$\frac{\partial f}{\partial t} + \frac{\partial f}{\partial q} \cdot \frac{\partial H}{\partial p} - \frac{\partial f}{\partial p} \cdot \frac{\partial H}{\partial q} = 2\beta \frac{\partial}{\partial p} \cdot (pf) + D \frac{\partial}{\partial p} \cdot \frac{\partial}{\partial p} f$$

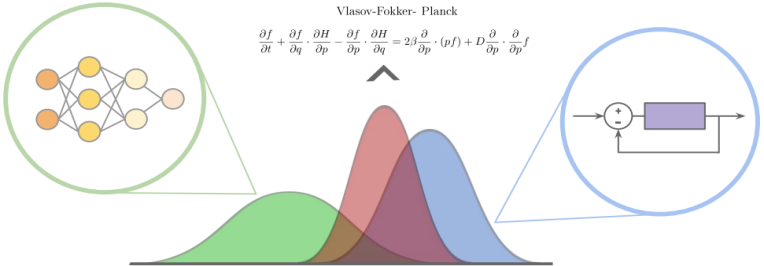

I'm interested in:

My research at IBPT

- The connection of physics models (e.g., partial differential equations) with machine learning (neural networks) for solving inverse problems.
- Tomography methods for the phase-space reconstruction of electron bunches in accelerators.
- Terahertz tomography for medical applications.

- Control of non-linear processes, complex simulations, and algorithm development for science and industrial applications.

AUTOMATIC SETUP OF THE SPS CONTROLLED LONGITUDINAL EMITTANCE BLOW-UP

SY
RF
BR

Without blow-up

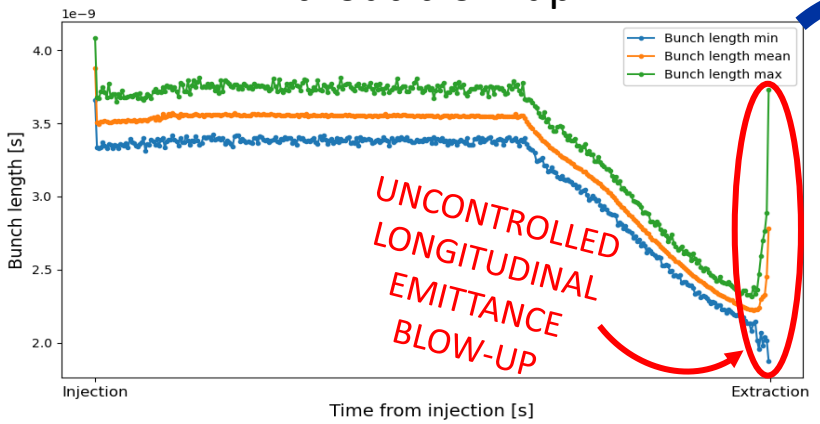

With blow-up

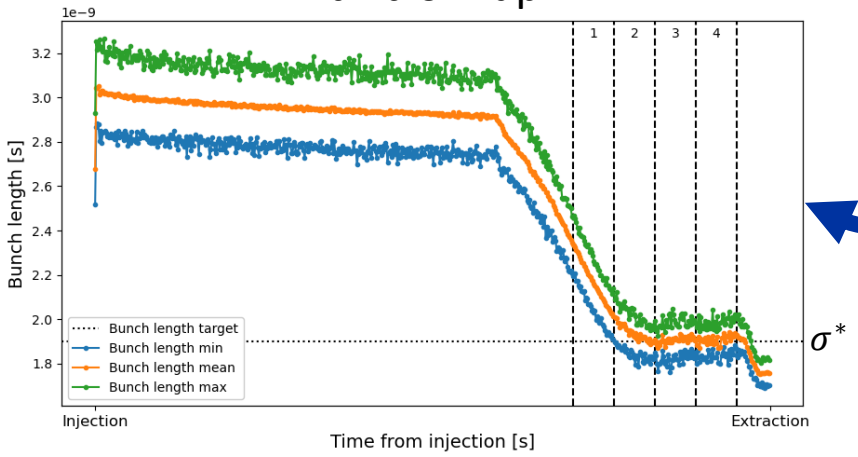

Blow-up needed to gain stability

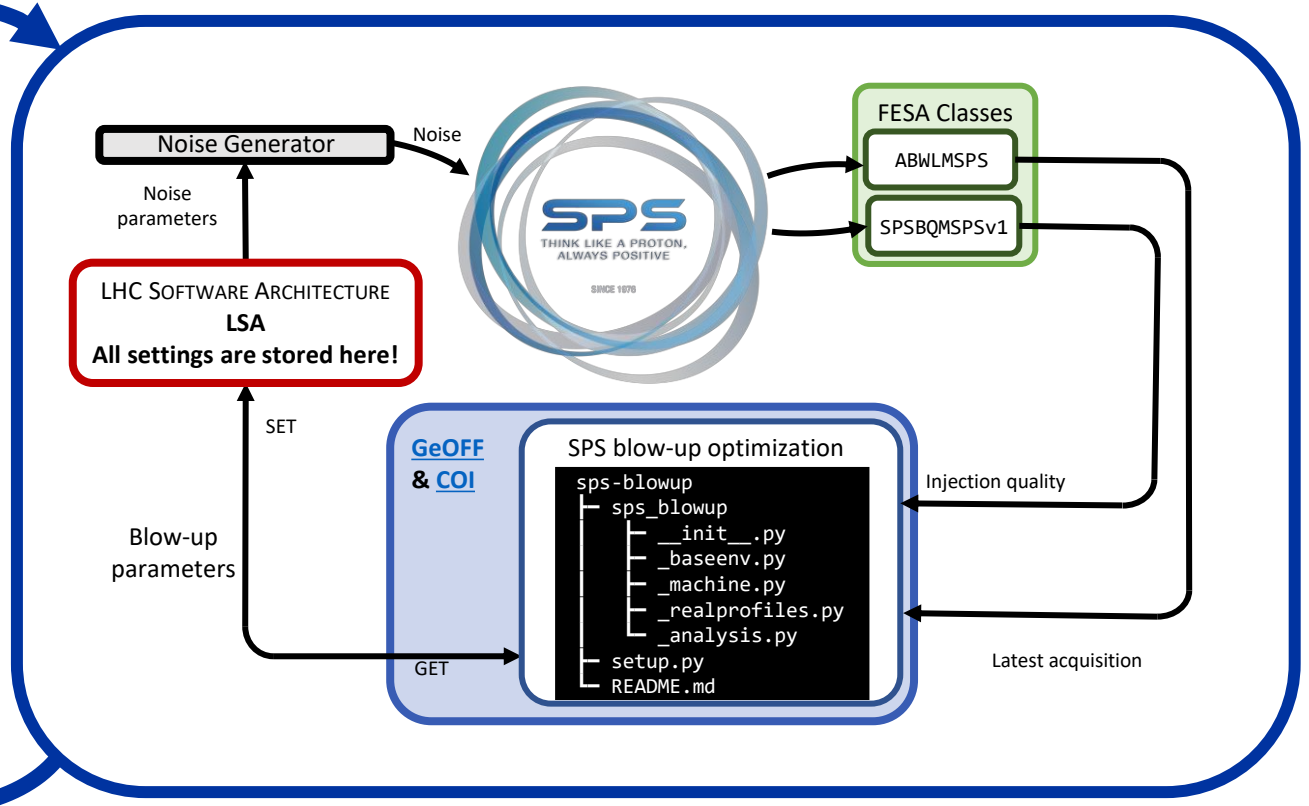

Blow-up setup **successful** for **high intensity** proton LHC beams:

- Beam kept at desired bunch length and stable.

Karlsruhe Institute of Technology

IBPT
Institute for Beam Physics and Technology

Dr. Julian Gethmann

Accelerator Physicist
Controls Group (KIT)

julian.gethmann@kit.edu
https://chaos.social/smartsammler
https://www.linkedin.com/in/ansantam/
https://github.com/smartsammler/

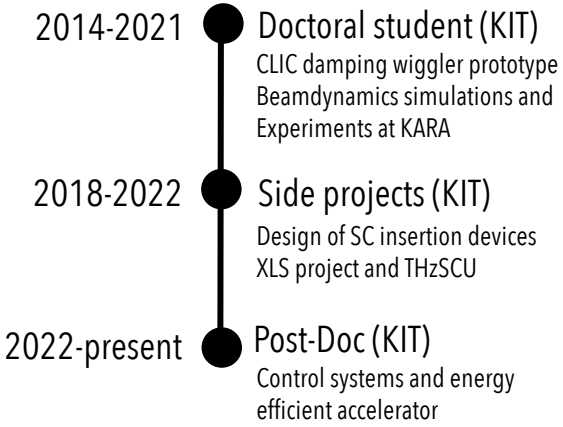

I'm interested in:

- Integration of ML into accelerator related topics
- Data engineering
- Power consumption optimisation
- Utilising RL for cooling plant optimisation
- (Ethics of ML/AI)

Possible future projects with RL advantages:

- Integrate the thermal wells into our cooling system
- Adjust the cooling system to external factors (weather, beam time / operation mode, …)
- Support my colleagues with their RL projects

Simple Python wrapper

IT-Infra.

Accelerators

Approach: Wrap accelerator specific interfaces with coherent, accessible and maintainable Python libraries.

```python
from ibpt import accelerator
from ibpt.accelerator import kara
from ibpt.epics import get_pv
from ibpt.pvs import get_pv_string
from ibpt.utils.network import is_internal_network

if is_internal_network():
    energy = epics.get_pv("beam_energy")
    print(f"{accelerator.get} runs with {energy} GeV")
    # -> KARA runs with 2.5 GeV
    energy_pv = get_pv_string("beam_energy")
    print(f"You can cross check it with 'caget {energy_pv}'")
    # -> You can cross check it with 'caget A:SR:BeamInfo:01:Energy'
```

```python
from ibpt import accelerator
from ibpt.accelerator import flute
from ibpt.epics import get_pv
from ibpt.pvs import get_pv_string
from ibpt.utils.network import is_internal_network

if is_internal_network():
    energy = epics.get_pv("beam_energy")
    print(f"{accelerator.get} runs with {energy} GeV")
    # -> FLUTE runs with 0.05 GeV
    energy_pv = get_pv_string("beam_energy")
    print(f"You can cross check it with 'caget {energy_pv}'")
    # -> You can cross check it with 'caget F:LIN-1:BeamInfo:01:Energy'
```

```python
from ibpt import accelerator
from ibpt.accelerator import kara
from ibpt.epics import get_pv
from ibpt.pvs import get_pv_string

if is_internal_network():
    energy = epics.get_pv("beam_energy")
    print(f"{accelerator.get} runs with {energy} GeV")
    # -> KARA runs with 2.5 GeV
    energy_pv = get_pv_string("beam_energy")
    print(f"You can cross check it with 'caget {energy_pv}'")
    # -> You can cross check it with 'caget A:SR:BeamInfo:01:Energy'
```

Sets defaults for KARA

Only control group/IT knows how to check this

Uses the correct PV for KARA

Coherent naming

EPICS

cassandra

ELOG

PostgreSQL

MML

python

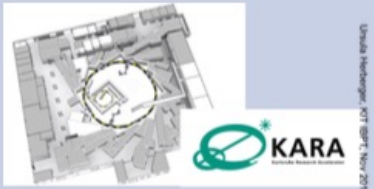

Ursula Hechtenger, KIT IBPT, Nov 2017

Steffen Schott, KIT IBPT, 2021

...

Interplay of the sub-systems

Long-term (>12 hours) storage solutions

Geothermal

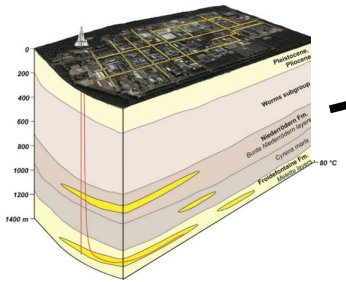

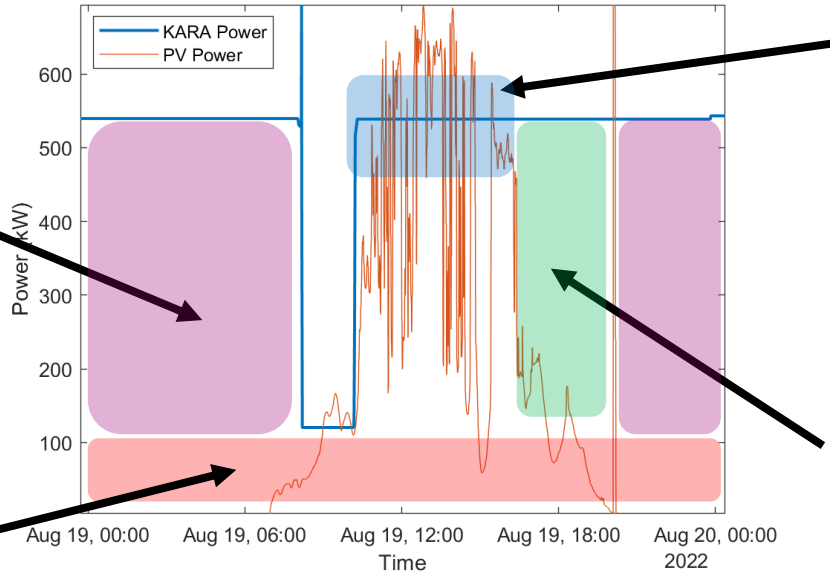

Fast dynamics solutions

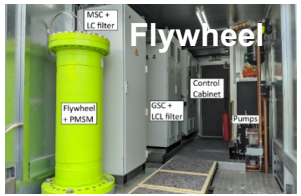

Medium-term solutions

Thermal wells

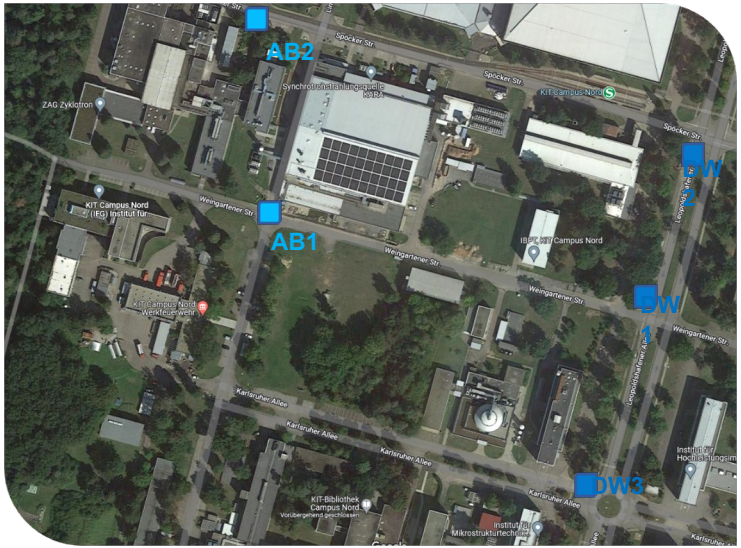

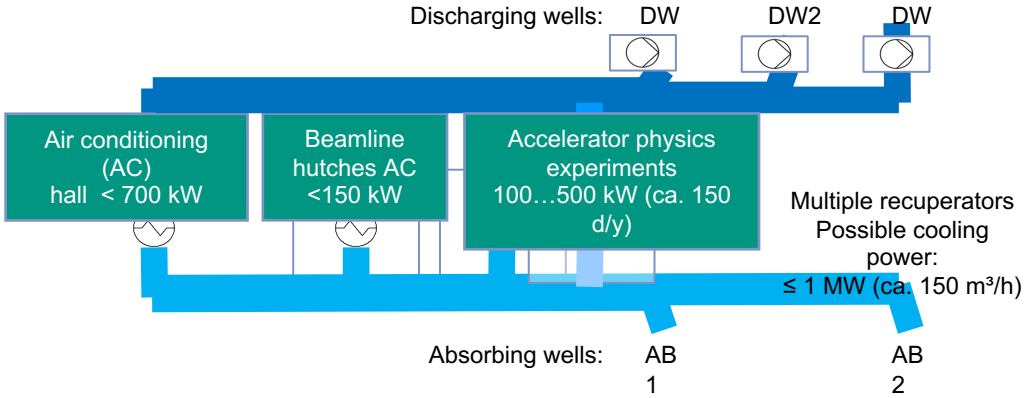

Jan Henry Hetzel

- Accelerator Physicist @ COSY in Jülich since 2014 (PhD in 2018), now employee @ GSI
- Background in Magnetic Measurements and Beam Dynamics for HESR @ FAIR
- Since 2018 part of Supervisor Crew @ COSY
- Strong Interest in Application of Modern Methods and Concepts to Support Operation of Accelerators

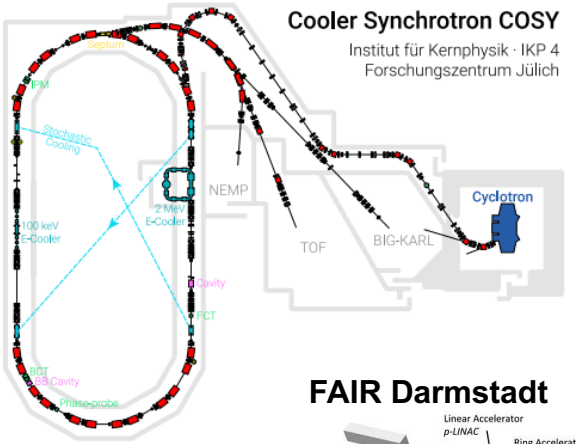

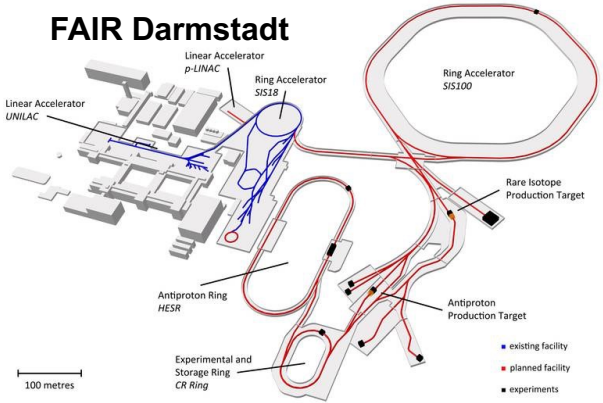

Recent Topics (extract)

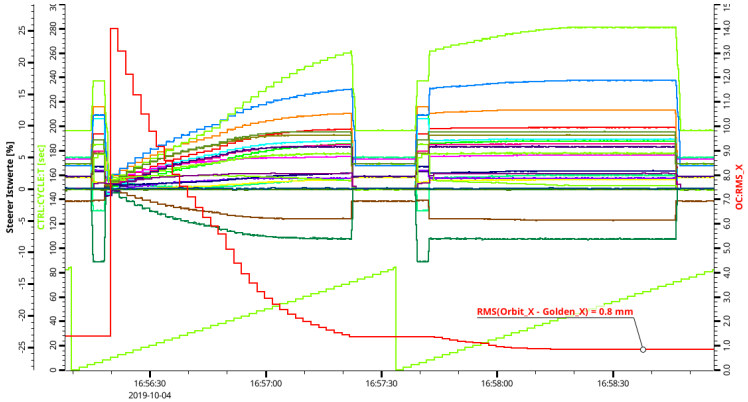

Automated Orbit Correction

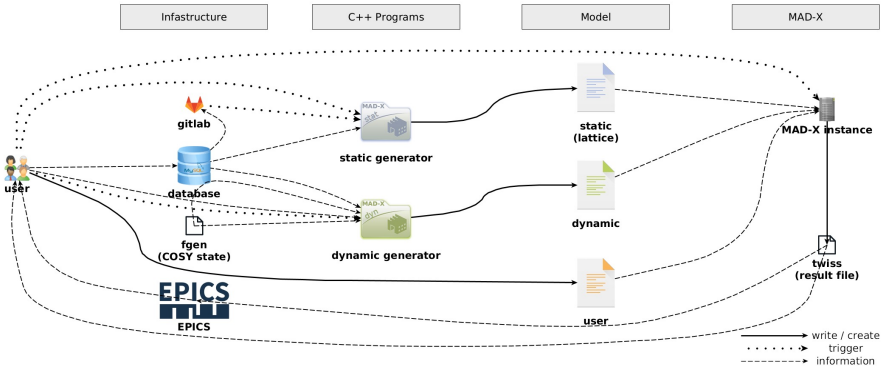

Generation of Online Model from Component Database and Display in Control System

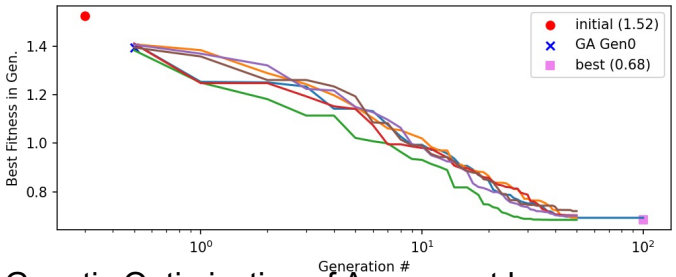

Genetic Optimization of Agreement beween Model and Machine

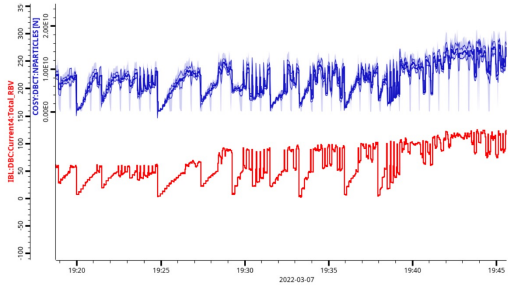

Automated Optimization of Particle Transmission through Injection Beam Line

Simon Hirlaender

Team lead:
Smart analytics and RL
IDA LAB
Artificial intelligence and human interfaces
Digital and analytical Sciences
University of Salzburg

Working in:
Artificial intelligence and machine learning, especially deep learning and reinforcement learning,
data science, statistics.
Research in industry and academia

Interested in sample-efficient, safe and robust reinforcement algorithms.

PhD @ CERN
Fellowship in SPS-OP
ML Coffee Cofounder (with Verena Kain)

Lecturing in
Reinforcement Learning
Precision medicine

IDA LAB
INTELLIGENT DATA ANALYTICS SALZBURG

PARIS LODRON UNIVERSITÄT SALZBURG

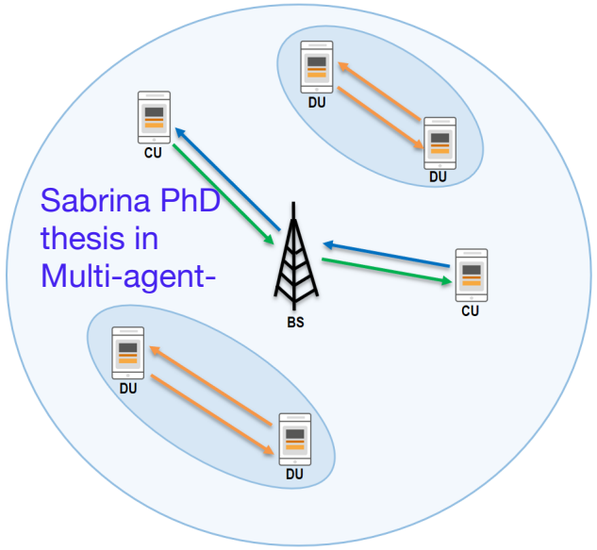

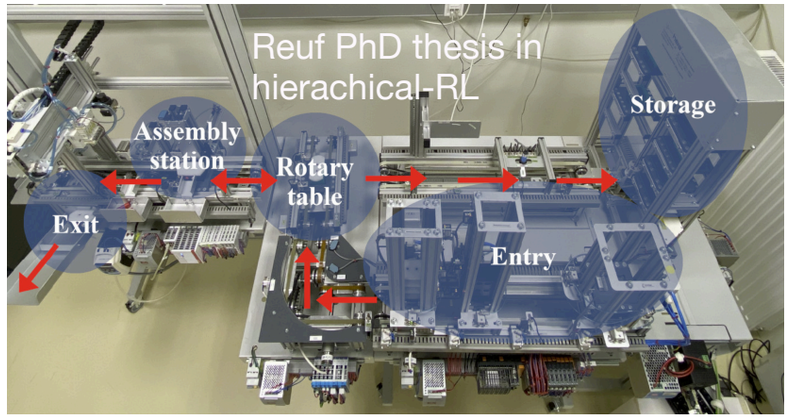

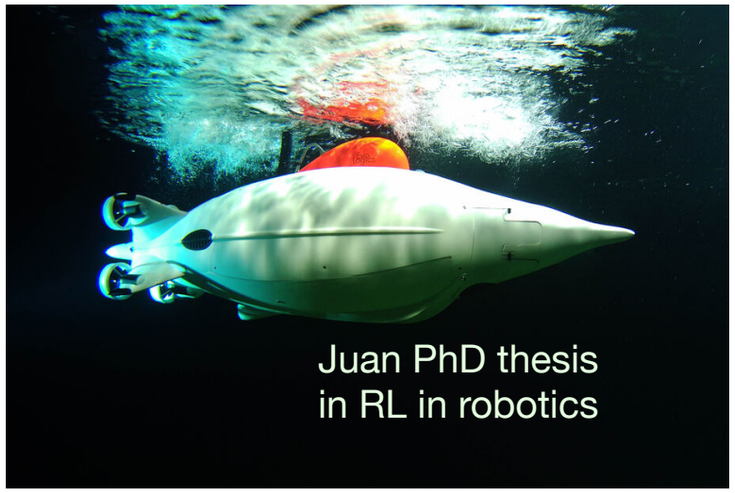

Sascha master thesis in interpretable RL in medicine

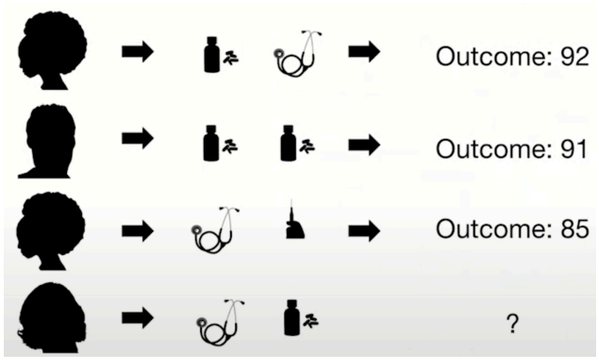

Lukas master thesis in Meta RL

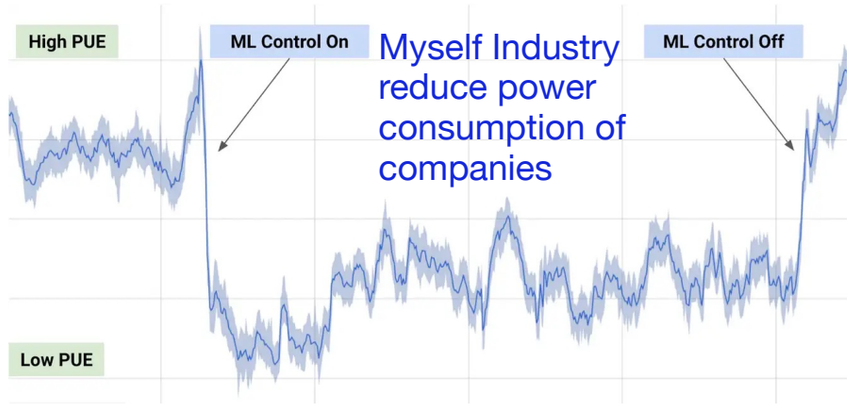

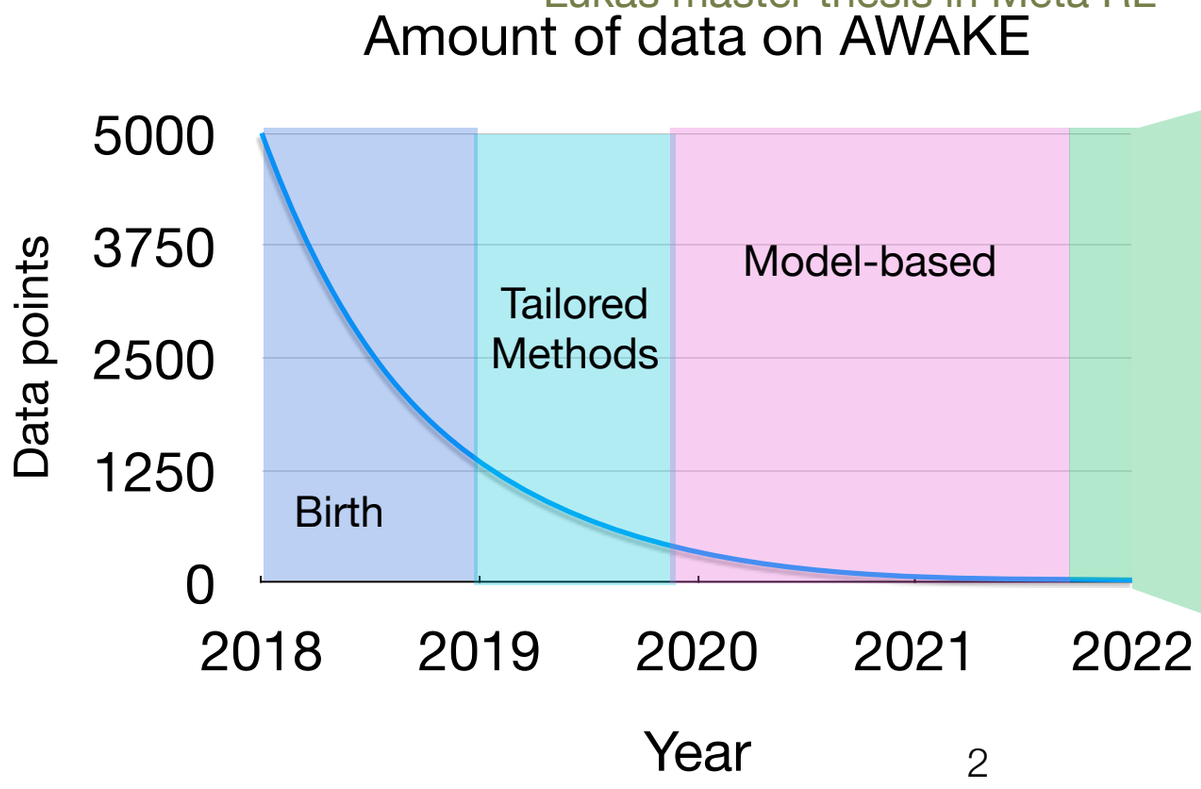

Shown at IPAC 2023

Jan Kaiser

Computer Scientist

DESY MSK IPC, Hamburg

jan.kaiser@desy.de
https://www.linkedin.com/in/jank324/
https://github.com/jank324

2015 - 2018 — **Hamburg University of Technology**
Computer Science (B.Sc.)

2018 - 2020 — **Hamburg University of Technology**
Computer Science (M.Sc.)
ANN-based Data Augmentation
Anomaly Detection for UAVs

2020 - present — **DESY**
Doctoral Researcher
Reinforcement Learning for Accelerator Control
ANN-based Virtual Diagnostics

Involvement in RL

- PhD project developing RL solutions to various accelerators
  - Transverse beam parameter tuning at ARES
  - Multi-agent RL-based beam threading at ARES (ongoing)
  - Beam dump loss and temperature feedback at European XFEL (ongoing)
  - FEL intensity tuning and LCLS (ongoing)

Interested in …

- RL for tuning and as a feedback on complex real-world systems
- Sim2real transfer
- Reward design
- Observation design (including feature engineering)

Transverse Beam Tuning at ARES

Successfully deploying RL to an accelerator with zero-shot learning

- Deploy a RL-trained optimisation algorithm to the **real-world** ARES accelerator with **zero-shot learning**
- Equivalent of **3 years of experience** tuning the transverse beam parameters
- **Faster optimisation** that alternative optimisation algorithms
- **Better final beam** than alternative optimisation algorithms

Algorithm	**MAE** Median (mm)	**Convergence** Median (Steps)
Do Nothing	1.122	0
Zero	0.588	1
FDF	0.699	1
Random	0.267	101
Powell	0.259	119
COBYLA	0.105	34
Nelder-Mead	0.007	112
Bayesian	0.081	101
Ours	0.008	7
Ours (Machine)	0.036	12

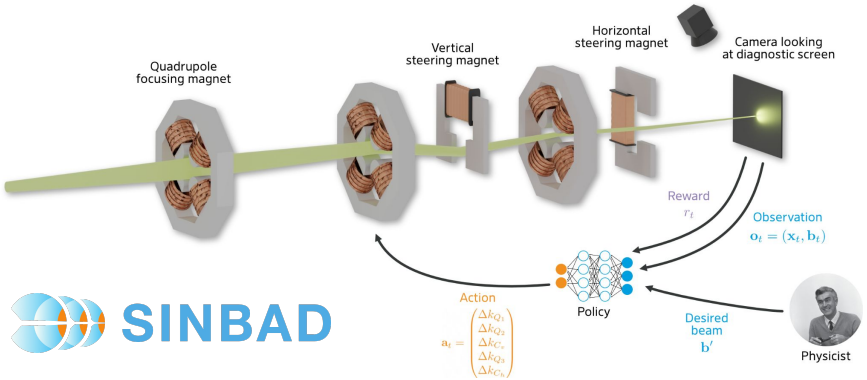

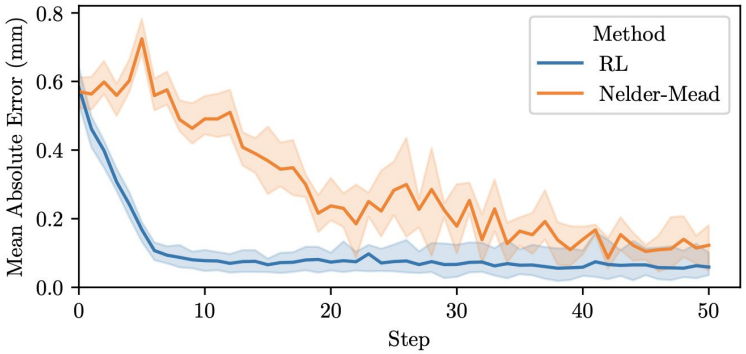

Transverse Beam Tuning at ARES

Successfully deploying RL to an accelerator with zero-shot learning

- **Autonomously** achieve tune in less than **5 minutes** what takes human operators over **20 minutes**
- Deployed application for using RL agent **in production**

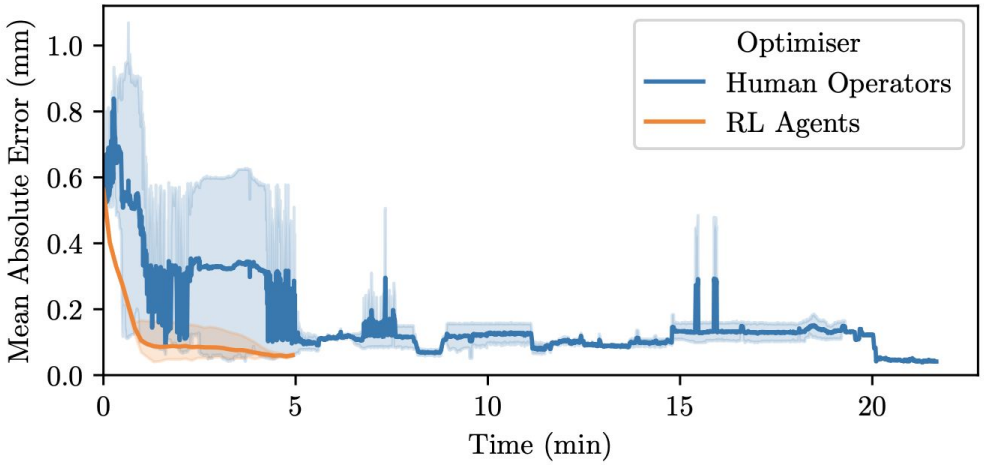

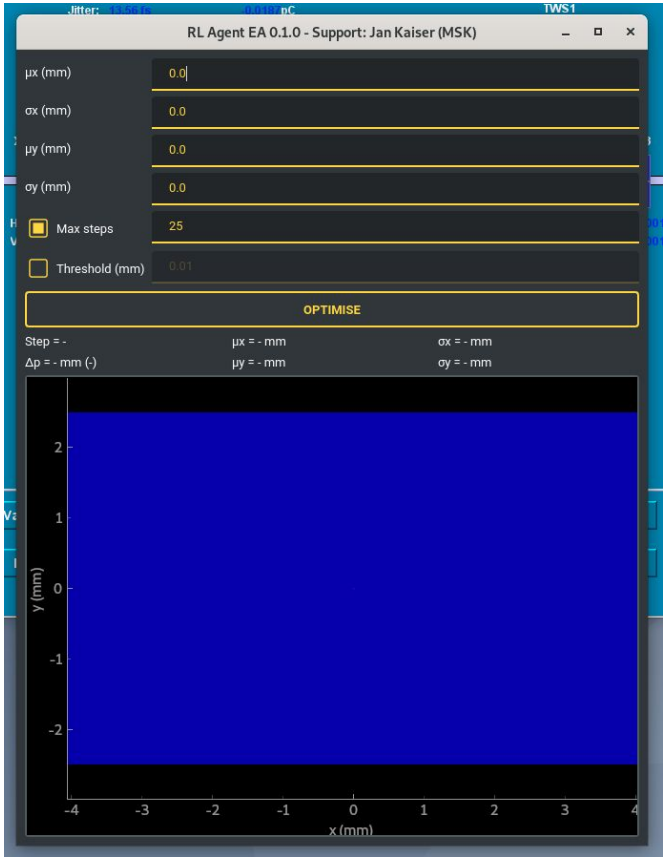

Karlsruhe Institute of Technology

IBPT
Institute for Beam Physics and Technology

Dr. Stephan-Robert Kötter
Accelerator Physicist
Member of AI4Accelerators team at IBPT (KIT)

stephan-robert.koetter@kit.edu

2016 - 2022 **Doctoral student (DELTA)**
- Orbit correction
- Beam dynamics
- Optimization (cx & non-lin)

2022 - present **Postdoc (KIT)**
- Orbit correction
- Anomaly detection
- Laser-pulse shaping

Involvement in RL:

- None so far

I'm interested in:

- Learning about RL
  - Methods
  - Applications
- Identifying usecases for my research

Research interest I: Orbit correction/stability

- Orbit steering (possible candidate for RL)
- Identifying faulty hardware (looked into BPMs)
- Clustering historic machine states (nothing done yet)

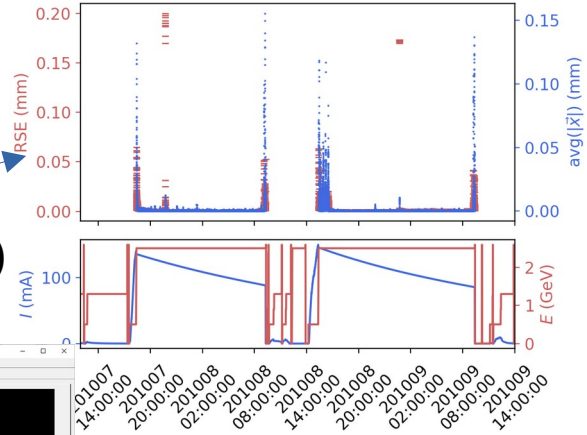

Research interest II: Laser-pulse shaping

- Goal: control transverse wavefront
- Actuator: spatial light modulator (SLM)
- CMOS camera & wavefront sensor
- Method: neural network

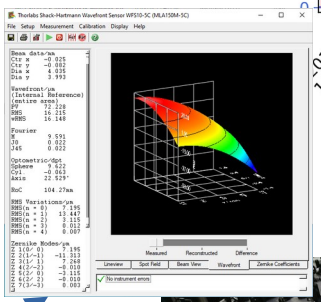

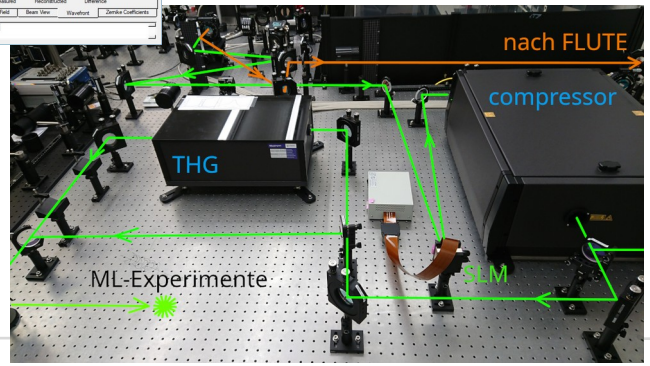

Originally planned setup

FH Salzburg

Reuf Kozlica
Junior Researcher (FH Salzburg)

reuf.kozlica@fh-salzburg.ac.at
https://github.com/reufko
https://www.linkedin.com/in/reuf-kozlica/

2015-2018 **Bachelor's Degree**
(FH Salzburg)
Media Informatics

2018-2020 **Master's Degree**
(FH Salzburg)
Data Science & Adaptive Software Systems

2020-present **Junior Researcher** (FH Salzburg)
Doctoral Student (University of Salzburg)

Involvement in RL:

- Reinforcement Learning in Production Plants
  - ➢ Automatic assembly and sorting of products
- Teaching-Assistant

I'm interested in:

- Advanced RL algorithms with focus on hierarchical methods
- Robustness and scalability of RL algorithms
- Deployment of RL algorithms to real world tasks and architectures

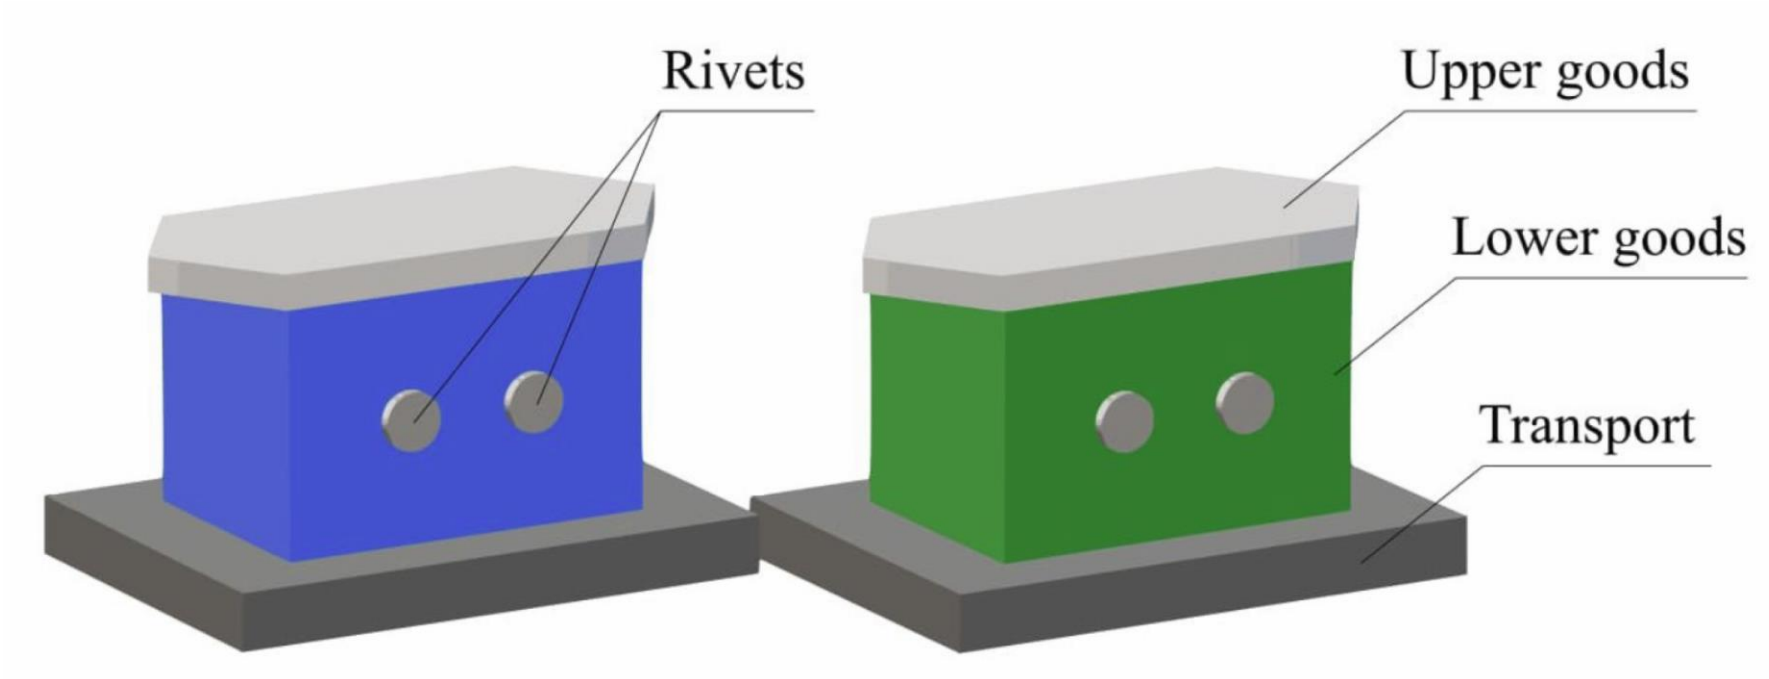
Rivets
Upper goods
Lower goods
Transport

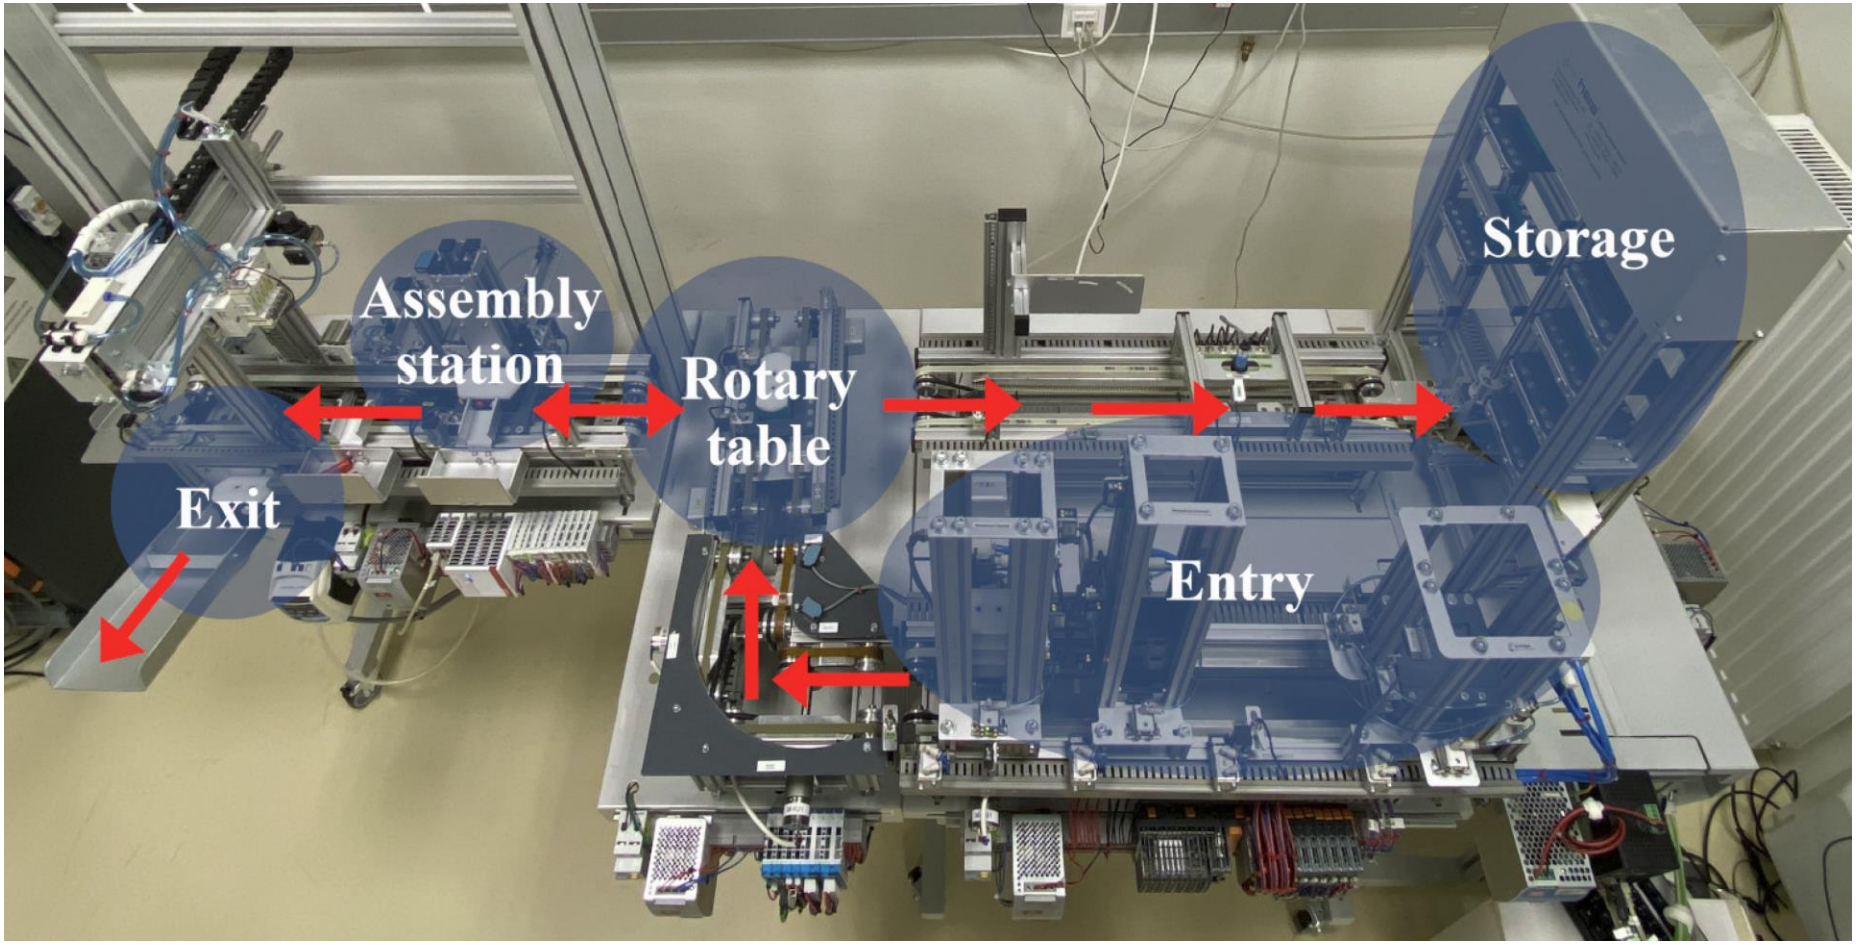
Assembly station
Rotary table
Storage
Exit
Entry

Petri Nets as a Simulation

- Mathematical modeling language
- Directed bipartite graph:
  - Places
  - Transitions
  - Arcs
  - Tokens
- Graphical notation for stepwise processes
- Application areas:
  - Process modeling
  - Control engineering
  - Simulation etc.

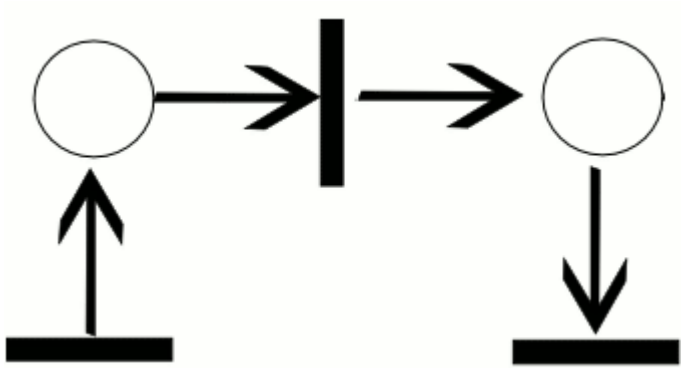

Petri Nets as a Simulation

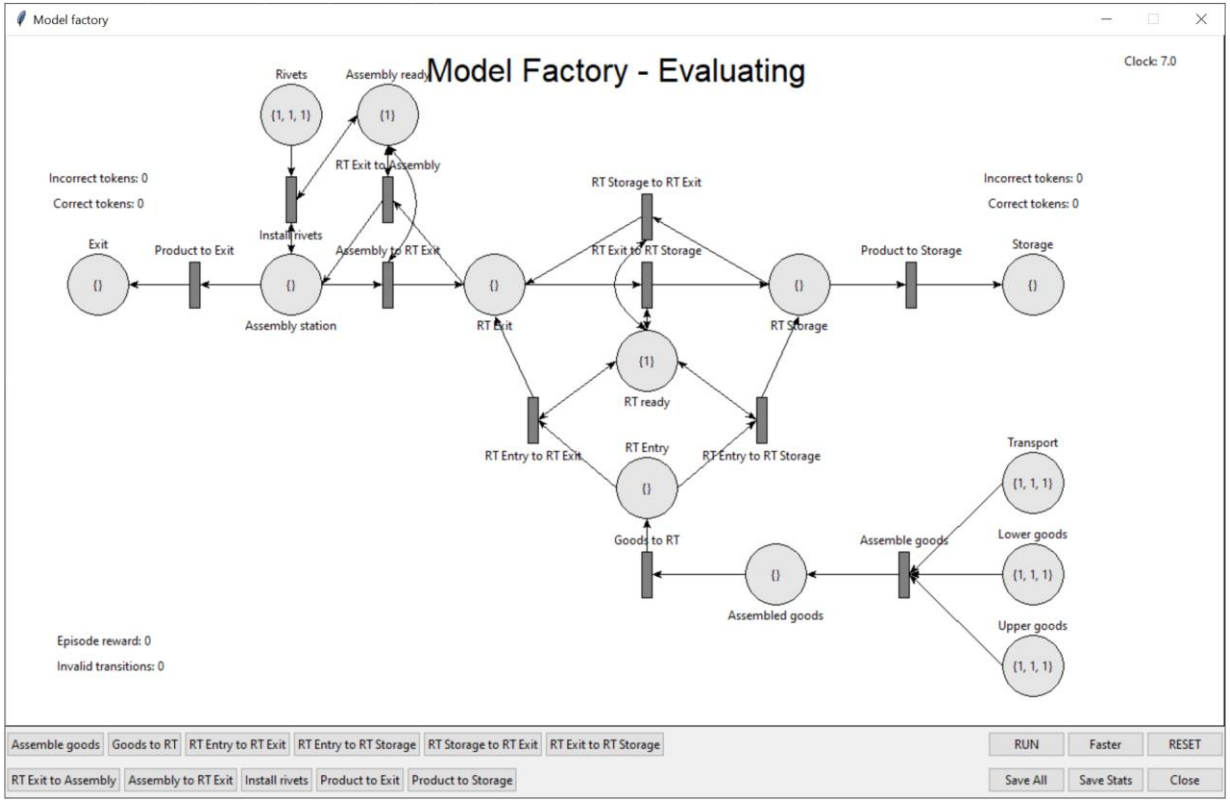

Deep Q-Learning and PPO results

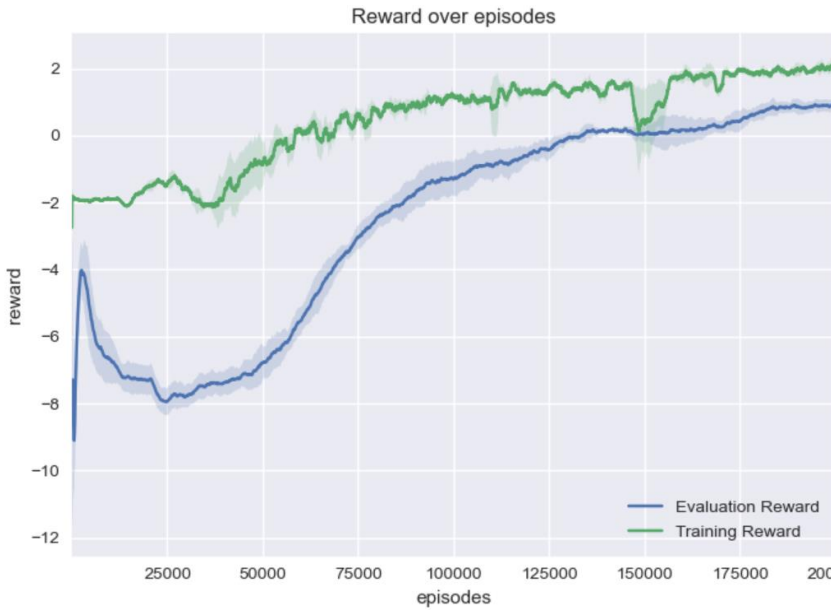

DQN

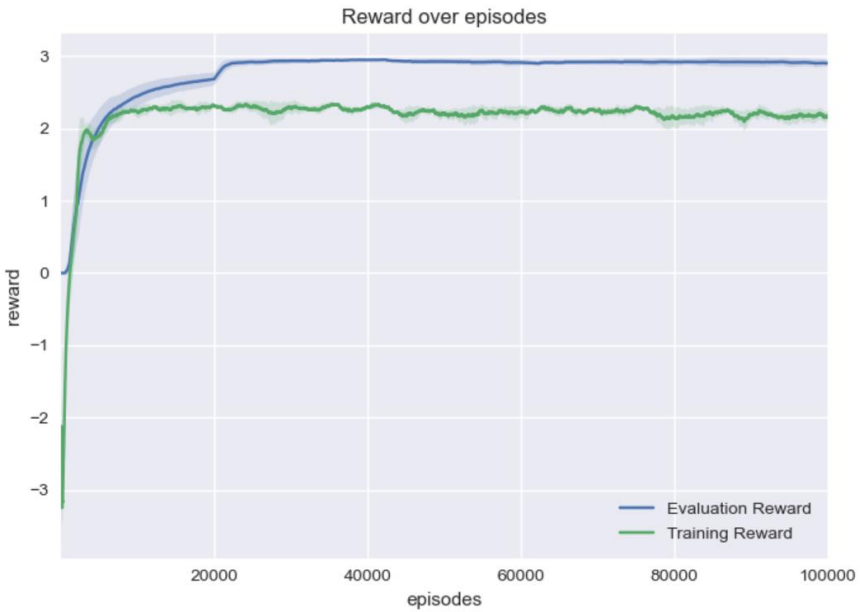

PPO

Sabrina Pochaba

Data Scientist & PhD student

Salzburg Research Forschungs GmbH

sabrina.pochaba@salzburgresearch.at

2013-2018 **Student (University Ulm)**
Mathematics & Biology
State examination

2018-2021 **Student (University Heidelberg)**
Mathematics
Master of Science

2021-present **Data Scientist (Salzburg Research)**
PhD Student (University Salzburg)
Machine Learning in Network communication

Involvement in RL:

- RL for network communications
- Multi-Agent RL
  - Game Theoreticel approaches for Multi-Agent Systems
  - Multi-Agent RL applied to wireless communication

I'm interested in:

- Deep learning algorithms
- Real life applications of RL
- Advanced RL algorithms (Hierarchical RL, Multi-Agent RL, Meta RL, Inverse RL, mix with other algorithms)
- Exciting ML stuff

RL for Network Communications

Problem:
Reliable Communication without regulation of BS

Solution:
Multi-Agent RL

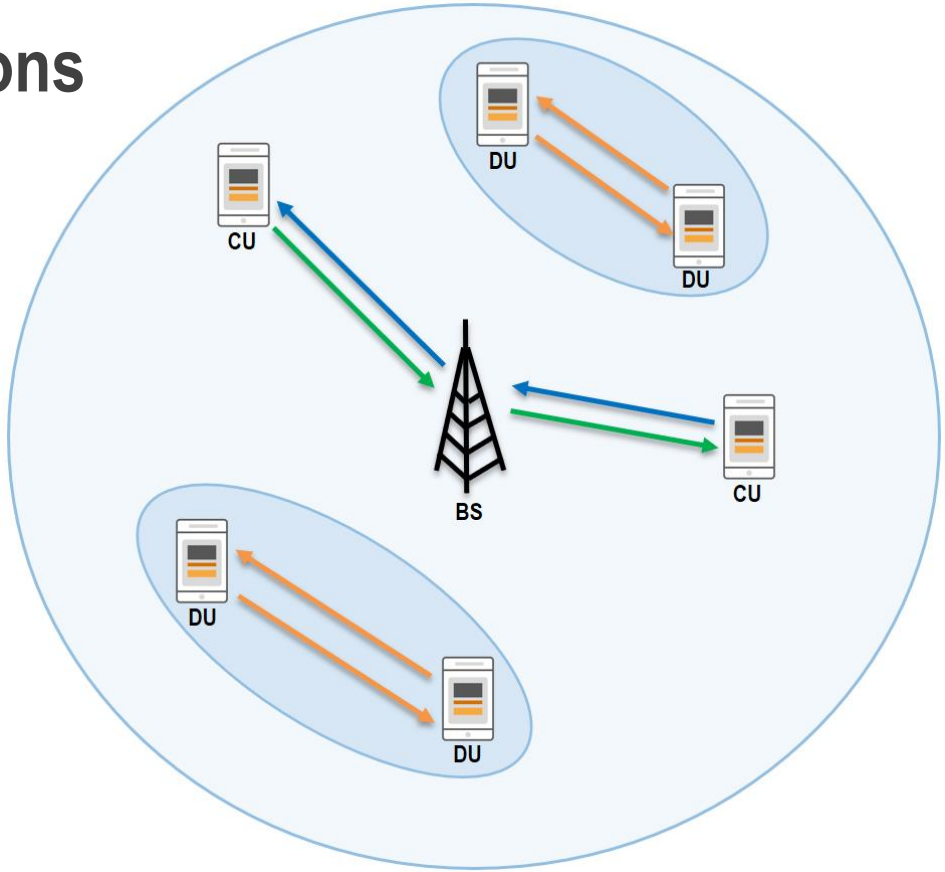

- **Environment:** Communcation cell
- **Agents:** Devices

- **Action:** Choice of Frequencyband
- **State:**
  - Own Frequencyband
  - Satisfaction (QoS)
  - Neighbors
  - Frequencyband of Neighbors
- **Reward:** Satisfaction (QoS) of all devices

Multi-Agent RL

Challenges

- Non-Unique Learning Goals
- Non-Stationarity
- Scalability Issue
- Various Information Structures

Game Theory for MARL

- Interaction of many player
- Searching for Equilibria
- NashQ

Karlsruhe Institute of Technology

IBPT
Institute for Beam Physics and Technology

Dr. Andrea Santamaria Garcia

Accelerator Physicist
Leading AI4Accelerators team at IBPT (KIT)

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

2014-2017 Doctoral student (CERN, EPFL)
- High luminosity LHC project
- Crab cavity failures
- Tracking simulations

2017-2019 Fellow (CERN)
- LHC Injectors Upgrade project
- PSB operation

2020-present Researcher (KIT)
- Machine learning activities

Involvement in RL:

- Autonomous accelerator project with DESY (more info during tutorial)
  - ➢ Automatic steering and focusing of beam
- Control of the microbunching instability
  - ➢ Enhancement of coherent synchrotron radiation

I'm interested in:

- Advanced RL algorithms (safety, robustness, hierarchical, multi-agent, meta RL, mix with other algorithms)
- Simulation to real world
- Feature engineering / dimensionality reduction
- Non deep learning algorithms
- Continuous fast feedback with RL

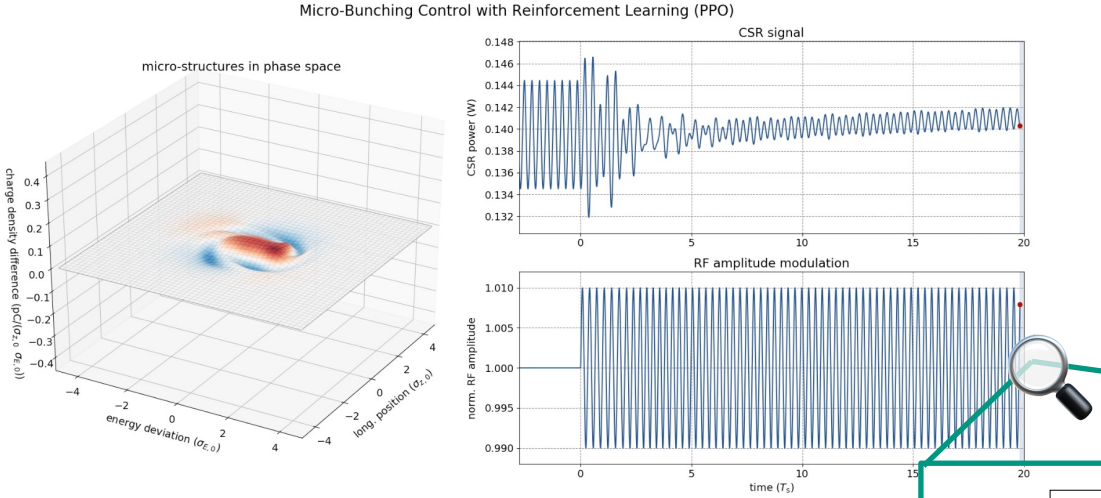

$$V_{RF} = \hat{V}(t)\sin(2\pi f_{RF} t)$$

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

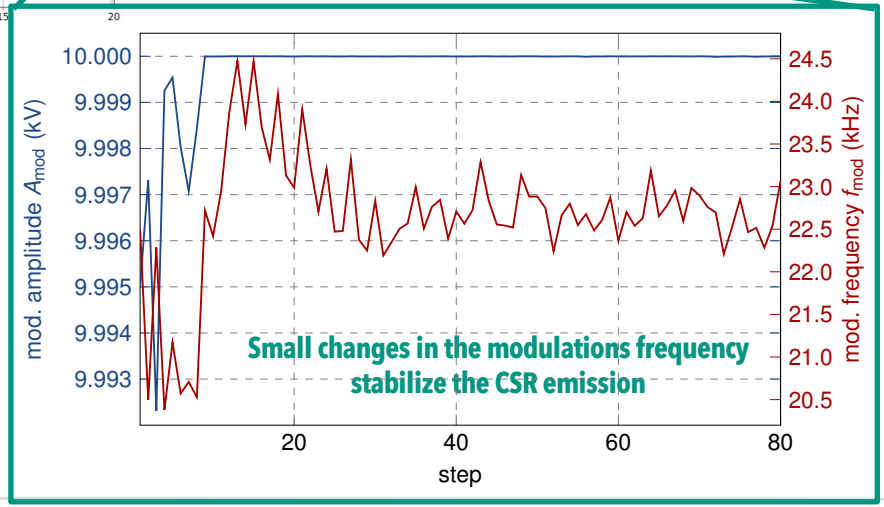

Courtesy of T. Boltz

IN PRACTICE: WE NEED HARDWARE!

Fast feedback for real-time optimization

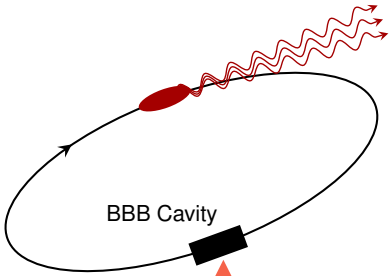

Detector

Schottky diode
50 GHz – 2 THz

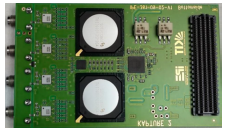

KAPTURE
Fast digitizer

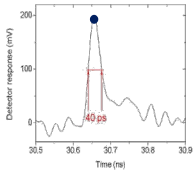

HighFlex 2 DAQ card
Xilinx ZYNQ

Custom modular readout DAQ optimized for AI applications

Ethernet 1GbE

Δt < 1 ms

Feedback

Bunch-by-bunch / LLRF

RF modulation computed by Versal

Versal board
Xilinx Versal ACAP

Scalar processor (ARM-PS) for the implementation of the "critic" neural network

Large **FPGA**/DSP for data management

Dedicated programmable **AI-engines** for the deployment of fast inference "actor" neural network + feature extraction

All linked by high bandwidth Network on a Chip (NoC)
full customization of dataflow

400 tiles x 8 FP multiplications x 0.5 GHz
→ up to 1.6 Tera FP operations per second !!!

Latency measured with Versal: 4.5 µs

"Micro-Bunching Control at Electron Storage Rings with Reinforcement Learning", T.Boltz

"KINGFISHER: a framework for fast machine learning inference for autonomous accelerator systems", L. Scomparin

"Accelerated deep reinforcement learning for fast feedback of beam dynamics at KARA," W. Wang

Michael Schenk
Accelerator Physicist
Fellow in Data Science for Beam Operations
michael.schenk@cern.ch
https://www.linkedin.com/in/m-schenk/

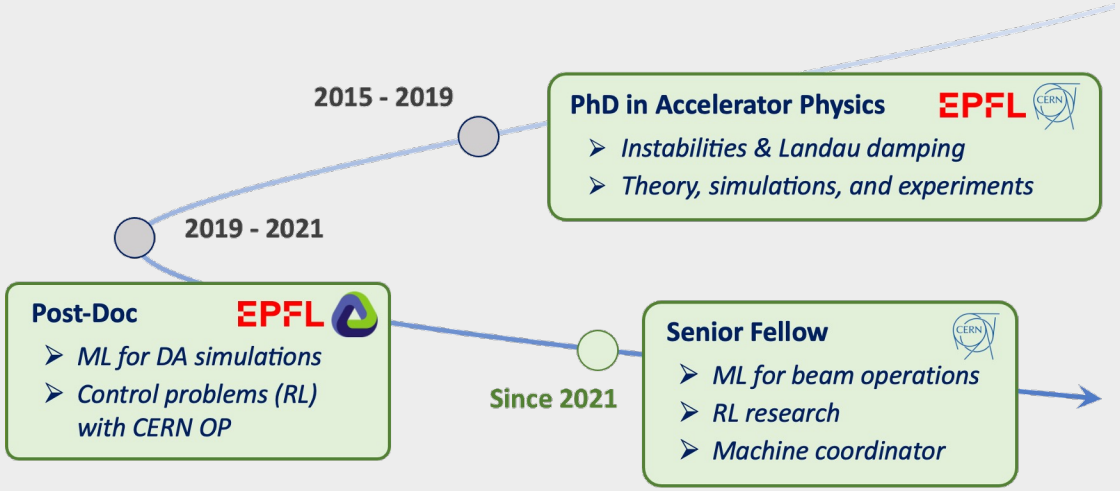

RL projects & interests

- **2019:** first RL steps on LEIR together with Simon
- **Now**
  Sample-efficient RL with quantum Boltzmann machines trained on D-Wave hardware
- **Future**
  - Continue N. Madysa's LEIR Schottky RL project
  - Hierarchical RL: student's project

Other ML / control experience

- Adaptive sampling / Bayesian exploration
- Conditional GANs
- Clustering
- Automatic tune adjustment using numerical optimizers

Free energy based RL (FERL)

RL with quantum Boltzmann machines

Q-learning

Various function approximators for $Q(s, a)$

- **Traditionally:** look-up table
- **DQN:** feed-forward neural net
- **FERL:** quantum Boltzmann machine (QBM)

Why FERL

- **Better learning efficiency?**
  ⇨ more **cost-effective** for beam operations
- **No FERL** algorithm exists for **continuous control**

1D beam steering

SPS North Area transfer line

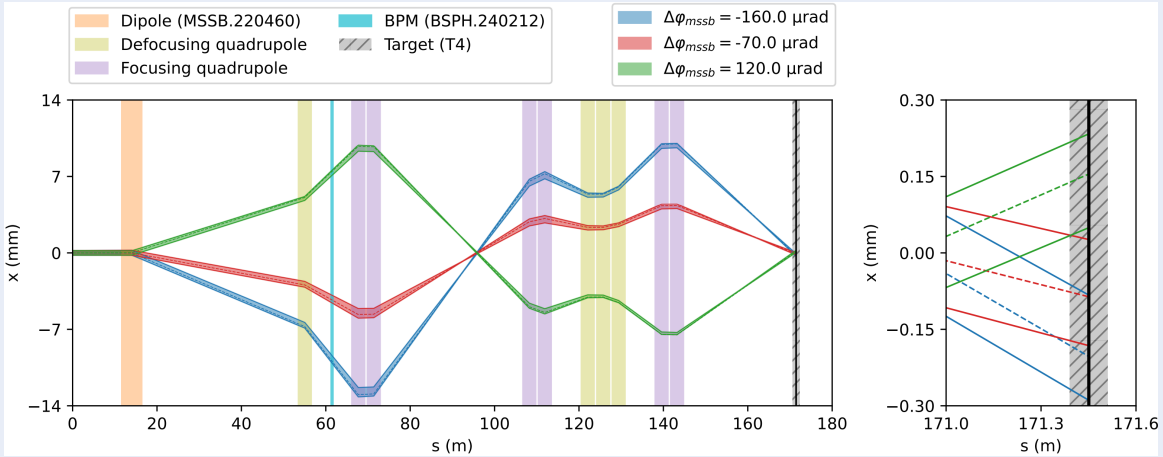

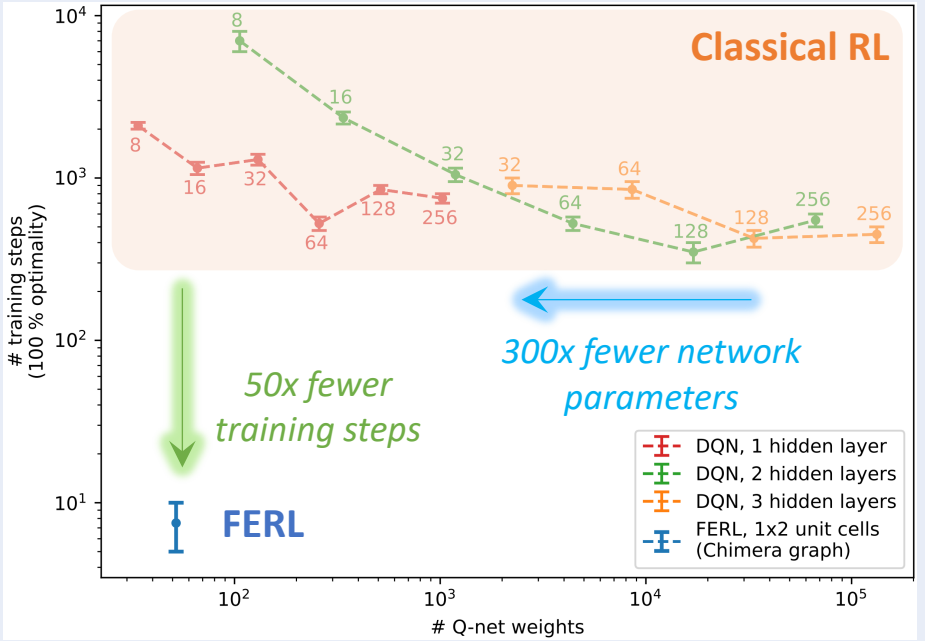

Free energy based RL (FERL)

RL with quantum Boltzmann machines

- **Inspired** by **classical actor-critic** scheme, developed a **hybrid algorithm for continuous state-action space**

  Intuition: *if critic learns more efficiently, can provide more valuable feedback early on during training*

- Trained on **D-Wave quantum annealer** and tested on **CERN's AWAKE** beam line

Training our algorithm on a quantum computer

Trajectory steering on simulated AWAKE e^- beam line

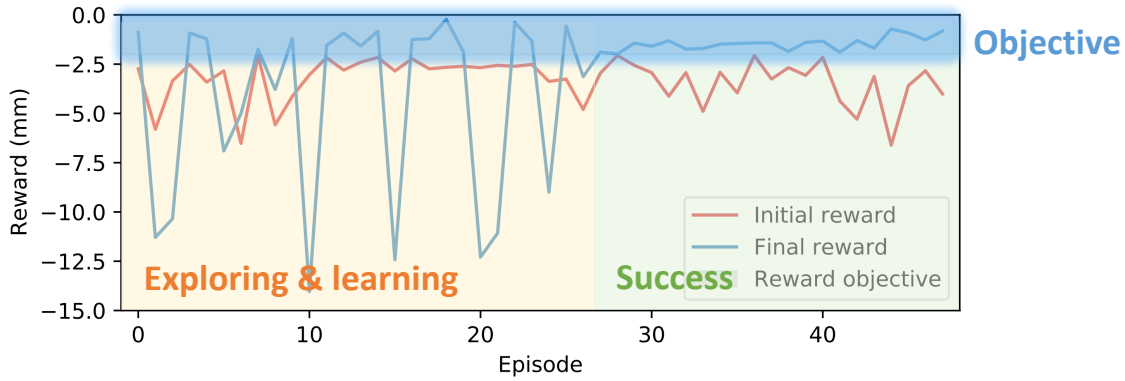

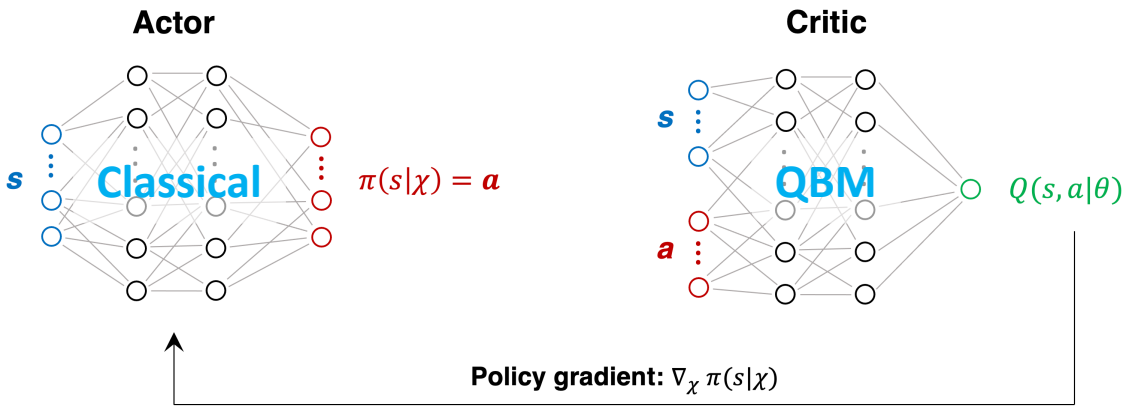

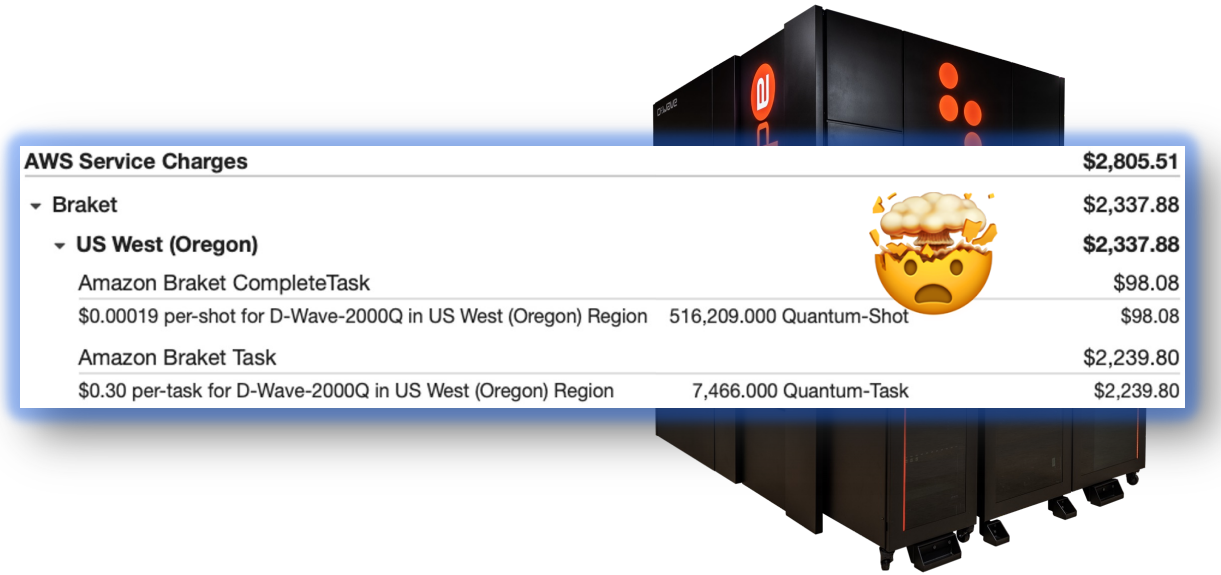

Meet & greet talk

Alexander Schütt

Helmholtz-Zentrum Berlin

RL4AA'23

Bachelor thesis (Without RL)

- Task: Predict the vertical beam size of the electron beam
- Why?
  - Derive vertical beam size from device settings
  - Long term: Adjust noise generator for constant beamsize
    $\Rightarrow$ minimize impact from one user to another
  - Preparation for BESSY III: round beam

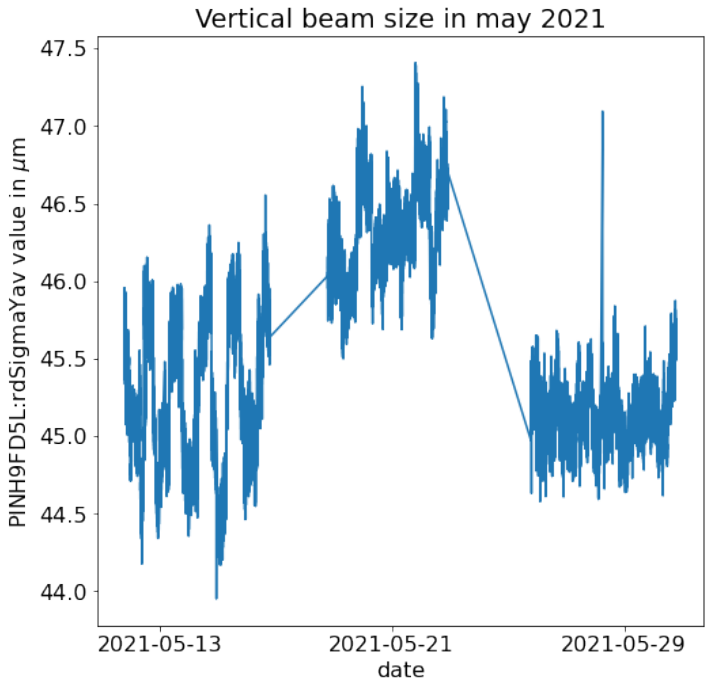

Figure 1: Vertical beam size of electron beam, measured using a pinhole camera.

Recurrent neural network regression

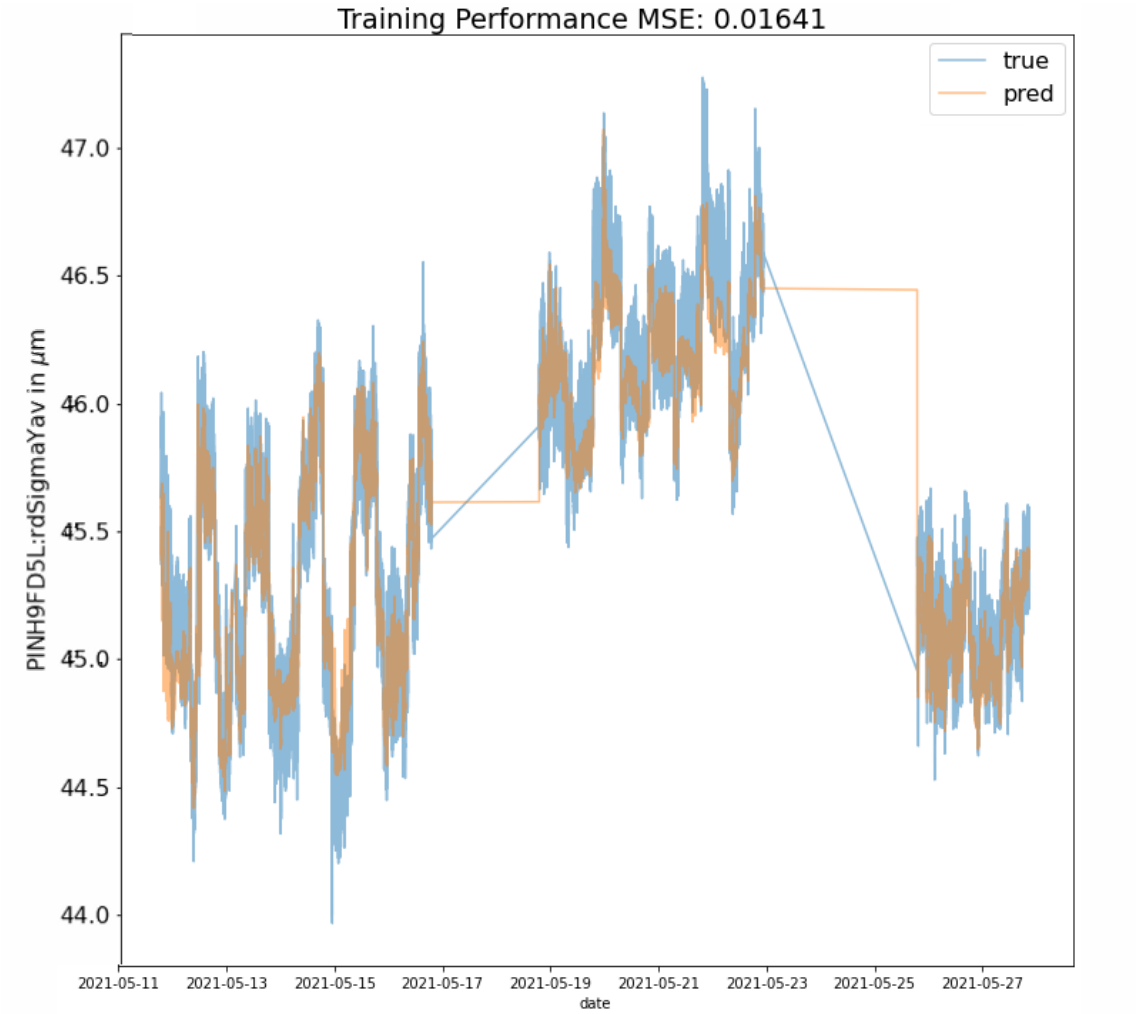

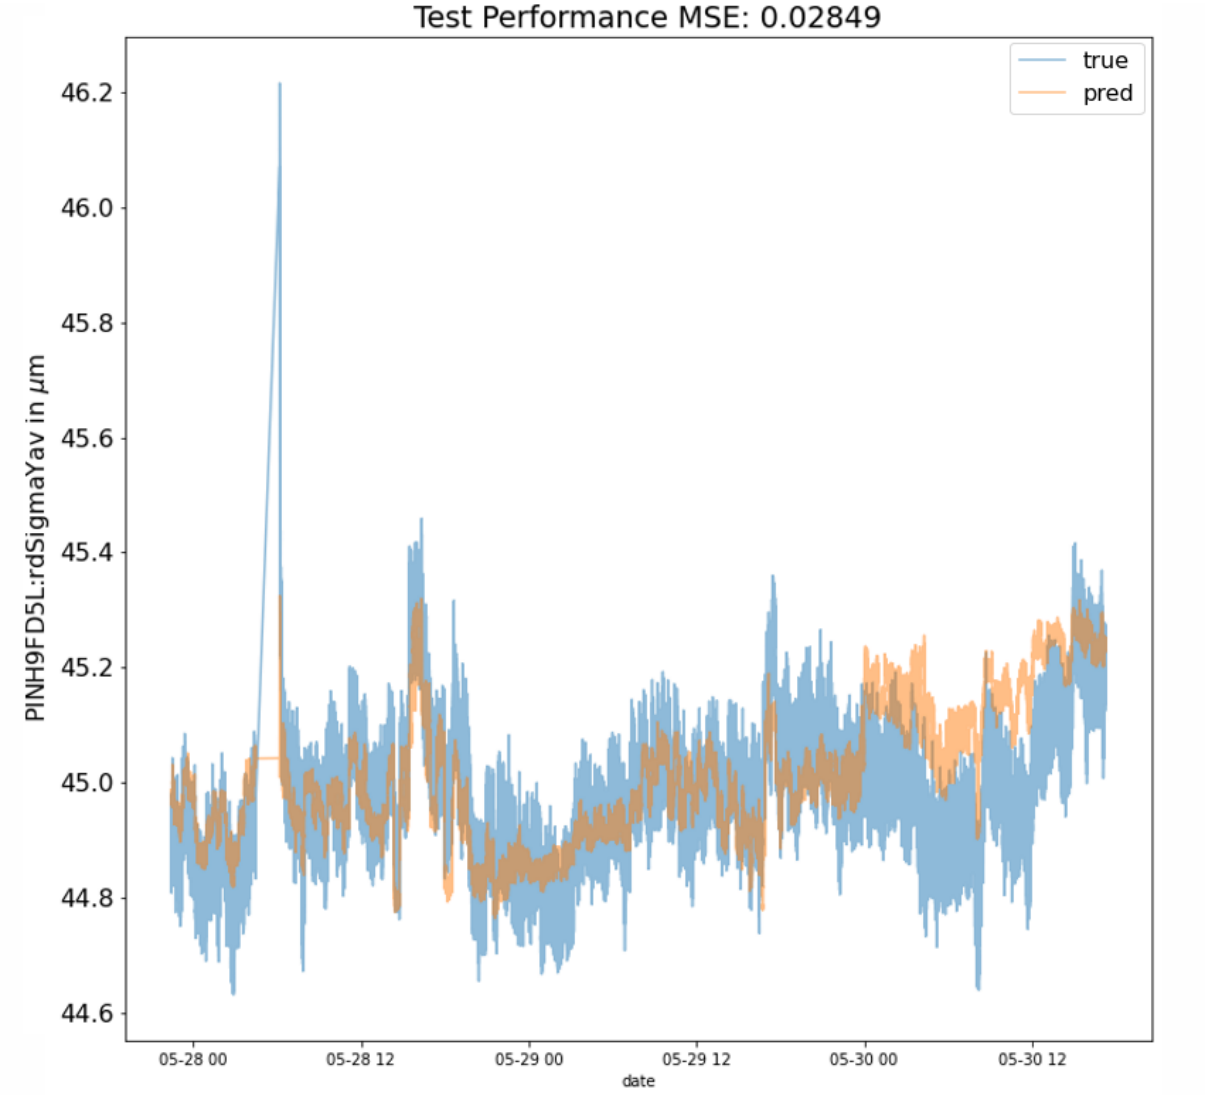
Test Performance MSE: 0.02849
true
pred
PINH9FD5L:rdSigmaYav in µm
46.2
46.0
45.8
45.6
45.4
45.2
45.0
44.8
44.6
05-28 00
05-28 12
05-29 00
05-29 12
05-30 00
05-30 12
date

Performance Analysis

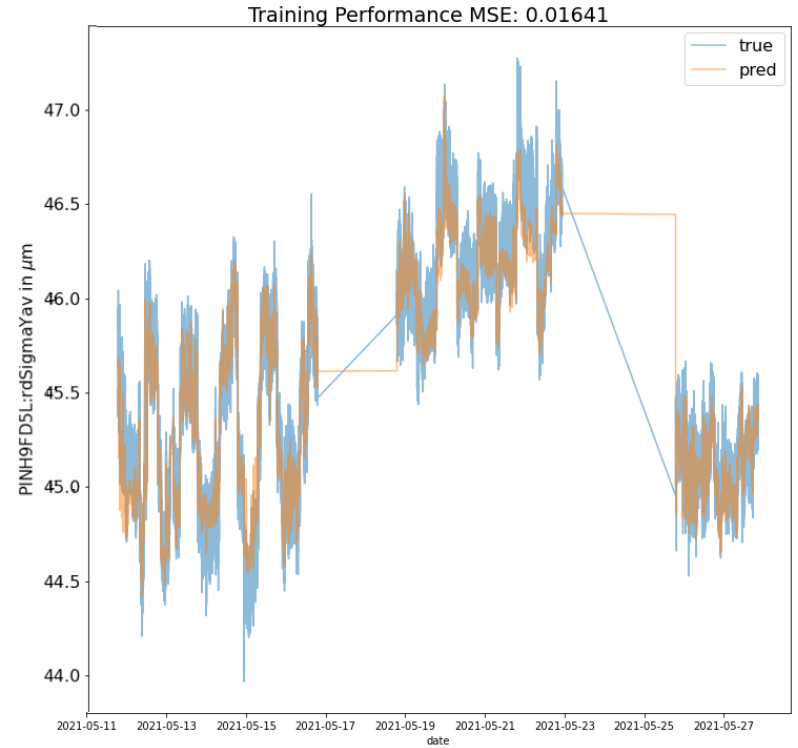

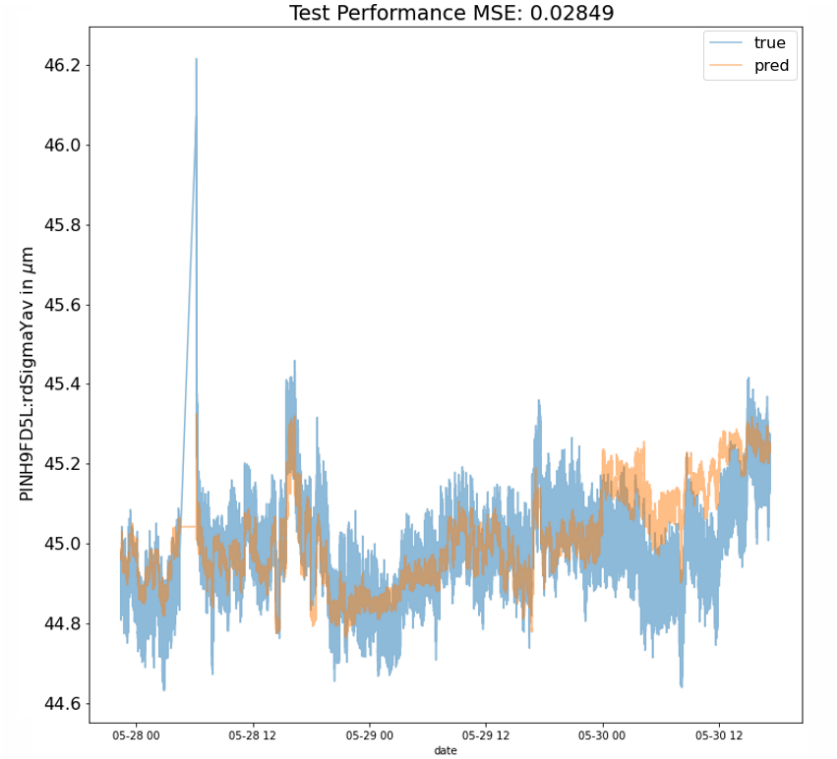

Idea for Master thesis (With RL)

- Task: Optimize steerer control
- Currently: Beam controlled via Least squared method

$$\Delta s_{t_i} = (R^T R + \lambda \mathrm{Id}_{64})^{-1} R^T (-b_{t_i}),$$

- **Problem:** Response matrix R might change in time
- Easy solution: Update response matrix via

$$R_{\mathrm{new}} = R_{\mathrm{old}} - 2\alpha (R_{\mathrm{old}} s_{t_i} - b_{t_{i+1}}) s_{t_i}^T$$

- Better solution: With RL?

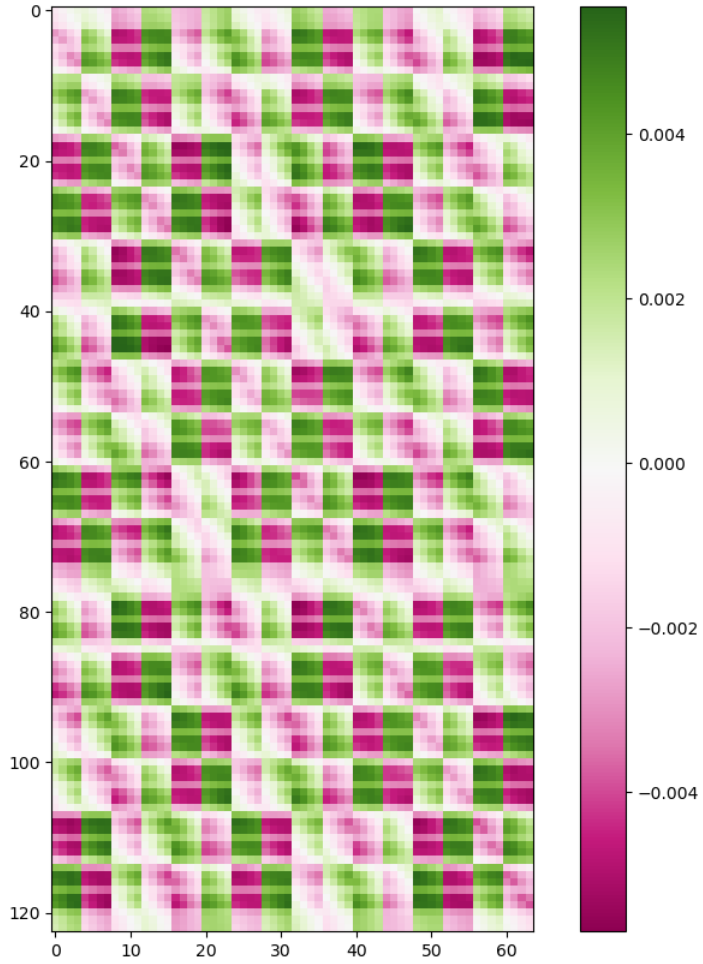

Figure 2: BESSY II response matrix

Better Solution with RL?

- Car driving analogy: „If I can drive 1000 cars, then small changes in the cars will not bother me.“
- Idea: Train the RL-agent to pick optimal steerer settings, after observing the BPMs, on multiple response matrices

$$R_{\text{new}} = R + \varepsilon,$$

  where $\varepsilon \in \mathbb{R}^{123,64}$ random
- To debate: Can it perform better than the easy solution?

Karlsruhe Institute of Technology

HELMHOLTZ
RESEARCH FOR GRAND CHALLENGES

KINGFISHER: Fast Machine Learning Inference for Autonomous Accelerator Systems

Luca Scomparin

Control of Longitudinal Beam dynamics at KARA

Interaction of beam with emitted radiation creates **instabilities** making **power fluctuate** (micro-bunching instability) → **limits user operation**

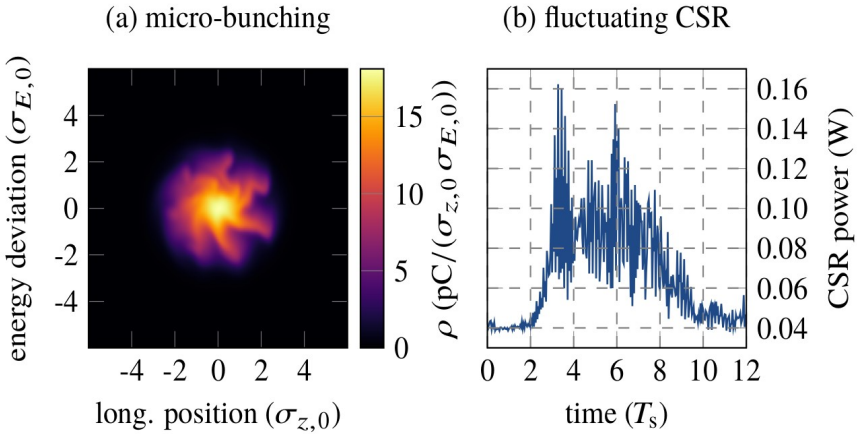

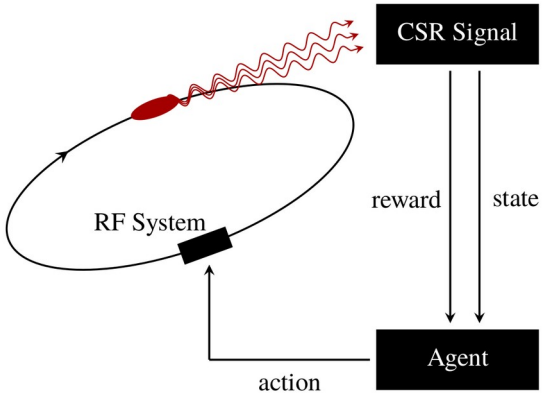

Control loop to limit this effect, **<u>but</u>** control problem not solved → can we used ML methods?
Yes, but inference must be at dynamics timescale O(tens of µs)

Specialized hardware is needed

Xilinx Versal ACAPs

Adaptive Compute Acceleration Platform (ACAP)

Which combines:

- ARM processors (Scalar Engines)
- FPGA (Adaptable Engines)
- AI Engines
- DSP Engines
- Advance interfaces

All linked by high bandwidth Network on a Chip (NoC)

Allows full customization of the dataflow depending on the application

VCK190 Evaluation Kit

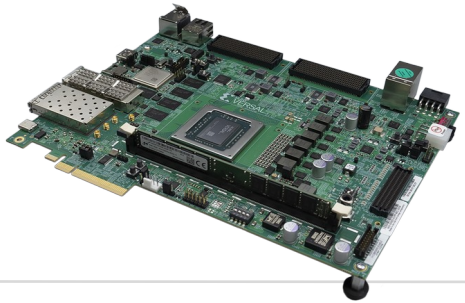

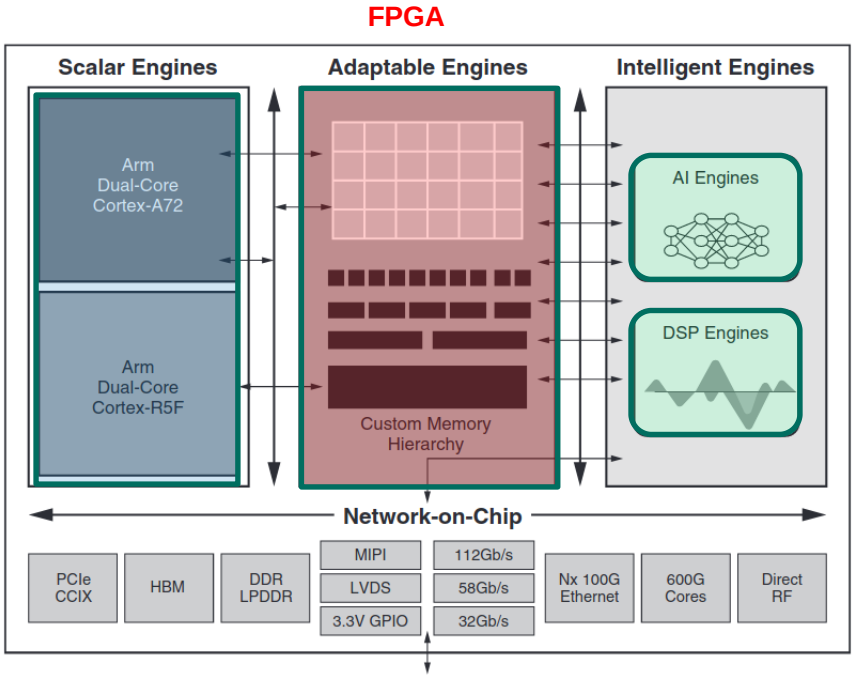

KINGFISHER: the structure

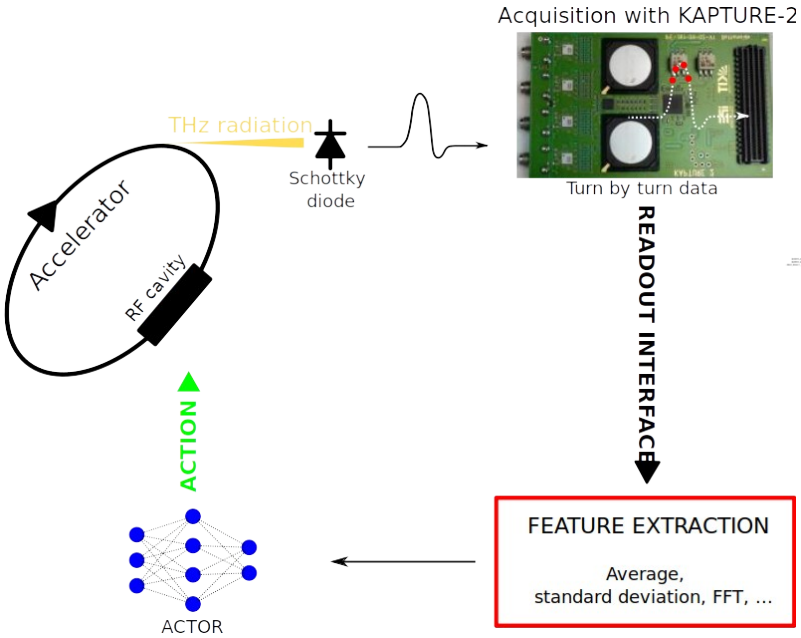

KIT Versal implementation

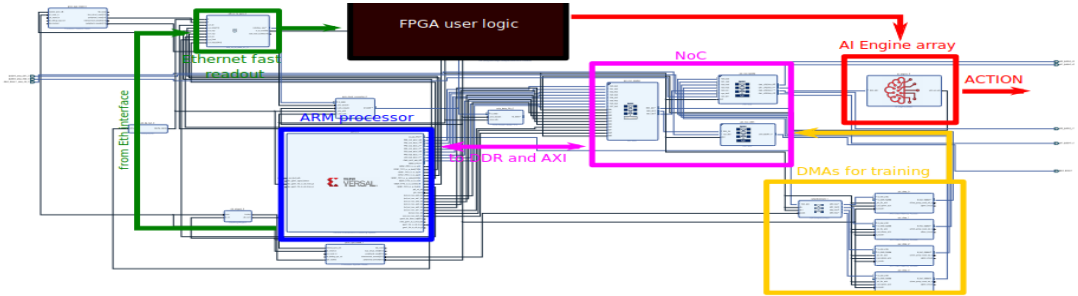

Waheedullah Sulaiman Khail

PhD in Software Engineering from Slovak University of Technology.

Research focus: Agile and lean organization of Patterns.

5+ years of experienced in Object Oriented Programming.

Teaching: Software Modeling, Software Architecture and Object Oriented Programming.

Working on developing digital twin since May 2022.

Karlsruhe Institute of Technology

IBPT
Institute for Beam Physics and Technology

Chenran Xu

Doctoral Researcher

chenran.xu@kit.edu
https://github.com/cr-xu

2018-2020 — **KIT**
Physics (M.Sc)
Bayesian optimization of injection efficiency at KARA storage ring

2020-present — **KIT**
Doctoral Researcher
Accelerator control with machine learning methods at FLUTE

Involvement in RL:

- Autonomous accelerator project with DESY
  - ➢ Automatic steering and focusing of beam
- Automatic beam control for FLUTE (ongoing)
  - ➢ THz radiation optimization

I'm interested in:

- Real world RL applications
- Generalizable RL agent
- Synergy between optimization and tuning
- Advanced RL approaches (hierachical RL, multi-agent RL...)

Previous experience with machine learning

Bayesian Optimization (BO) for Injection optimization at KARA

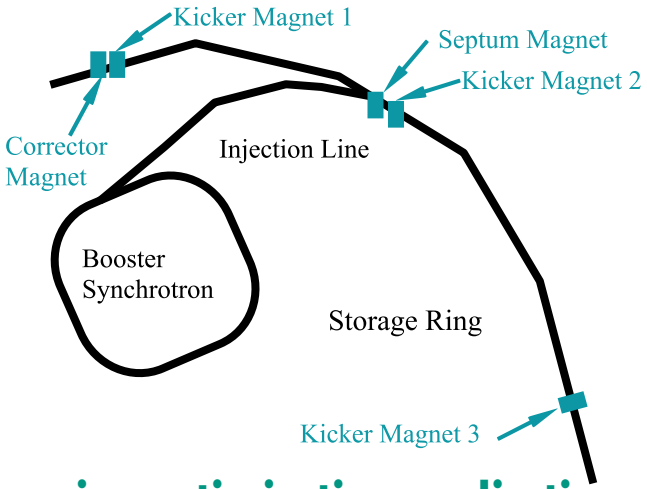

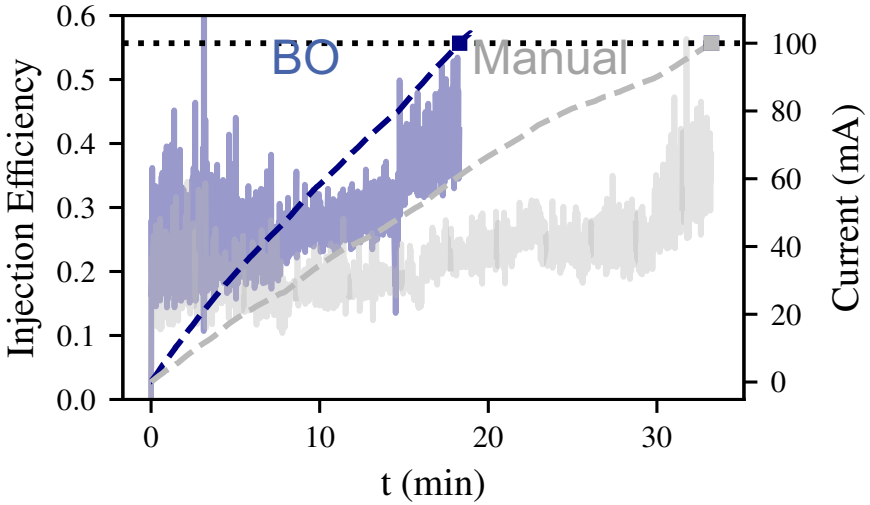

Other Bayesian optimization applications:

- Transverse beam tuning at ARES (benchmark to RL)
- SASE tuning at EuXFEL
- Simulation parameter optimization for THz generation

What I'm working on

Autonomous Accelerartor Project

FLUTE

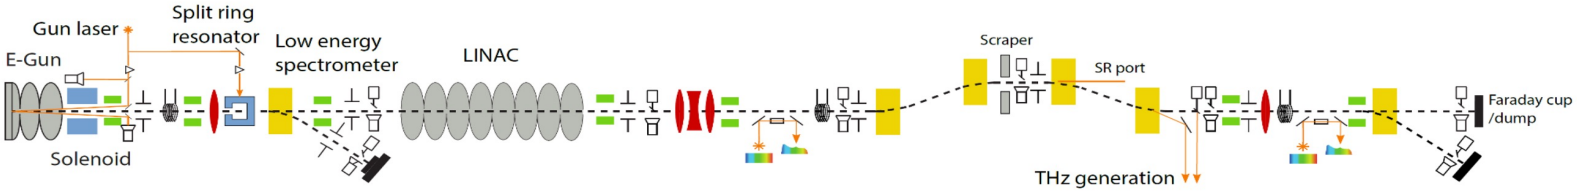

ARES

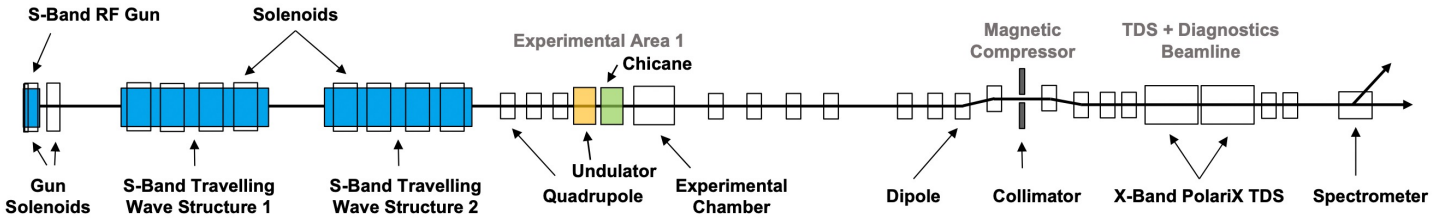

Transfer RL agents to similar tasks (same goal, different lattice)

- Domain Randomization
- (Future) Meta-RL