Machine Learning Applications at KARA, KIT

Tobias Boltz on behalf of the KIT team | December 3, 2019

Laboratory for Applications of Synchrotron Radiation (LAS)

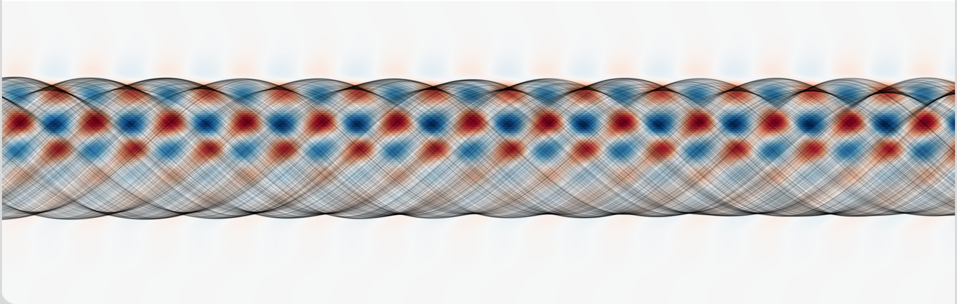

Overview

Ongoing Machine Learning Efforts

Micro-Bunching Control at KARA using Reinforcement Learning

- Feedback Design and Simulations (PhD thesis Physics, Tobias Boltz)
- RL in Simulations (Master's thesis Computer Science, Melvin Klein)
- Hardware Implementation (PhD thesis Electronics, Weijia Wang)

Optimization of Injection Efficiency at KARA

- Bayesian Optimization using Gaussian Processes (Master's thesis Physics, Chenran Xu)

Micro-Bunching Control

Micro-Bunching and CSR Power Fluctuations

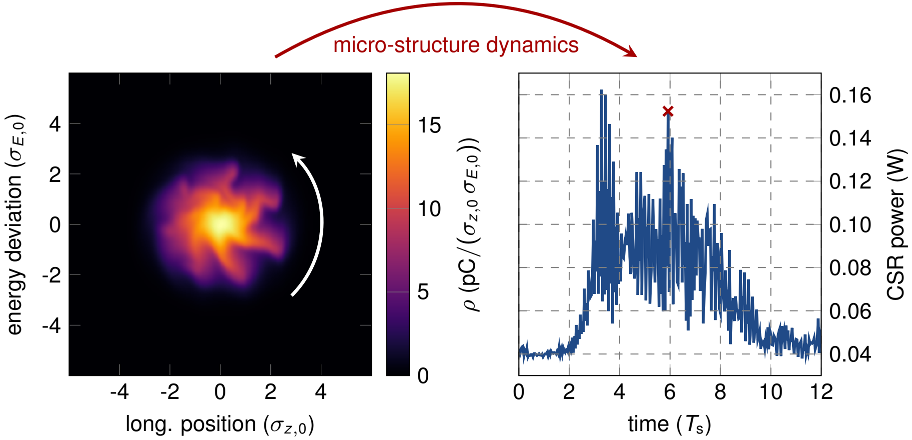

Simulation code: Parallelized VFP solver ***Inovesa*** (`https://github.com/Inovesa/Inovesa`)
Schönfeldt, P. *et al.*, Phys. Rev. Accel. Beams 20 (2017)

Micro-Bunching Control

CSR Power Spectrogram: Dependency on Bunch Current

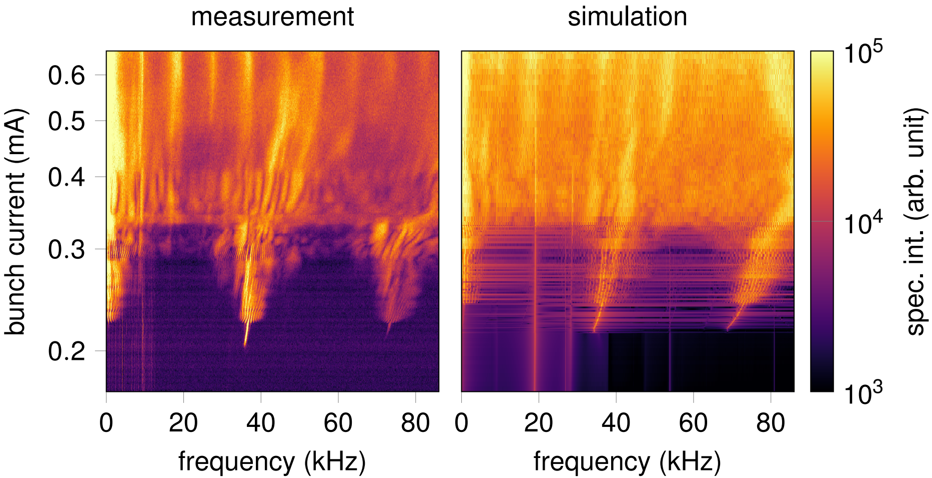

Micro-Bunching Control

CSR self-interaction

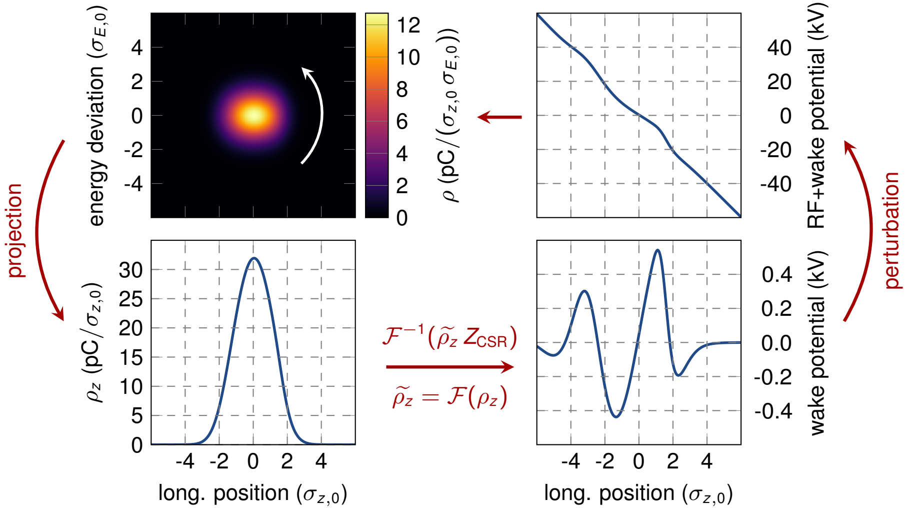

Micro-Bunching Control

…via Dynamic RF Amplitude Modulation

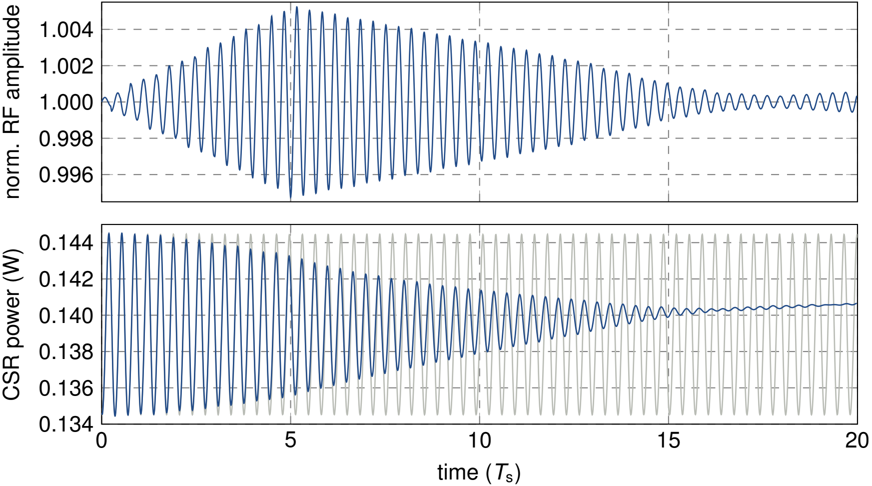

Micro-Bunching Control

RL in a Nutshell: Learning from Interaction

- mathematical foundation: Markov decision process (MDP)

  "The future is independent of the past given the present."

policy: defines agent's behavior

goal: maximize cumulative reward

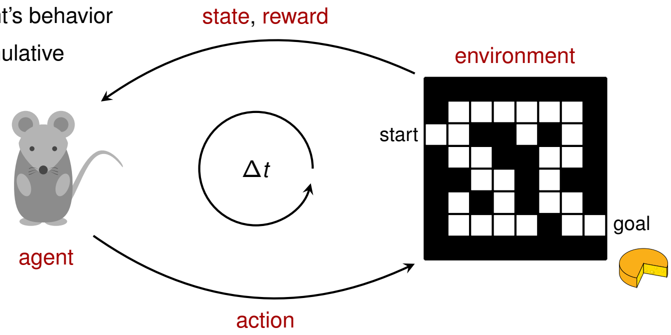

Micro-Bunching Control

RL – developing some Intuition . . .

- goal-directed learning from interaction (trial-and-error search)
- finding a sequence of actions (e.g. moves in a game)
- taken actions affect the following states (e.g. board positions)

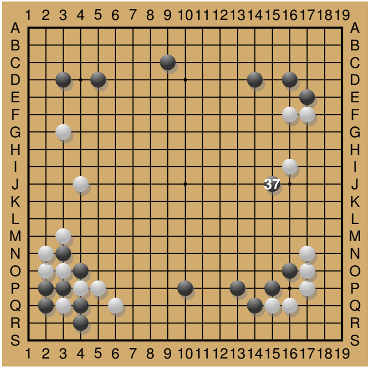

cartpole balancing

(a textbook RL problem)

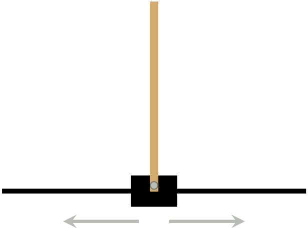

Micro-Bunching Control

Feedback Scheme

goal: control micro-bunching (longitudinal beam dynamics) to optimize emitted CSR

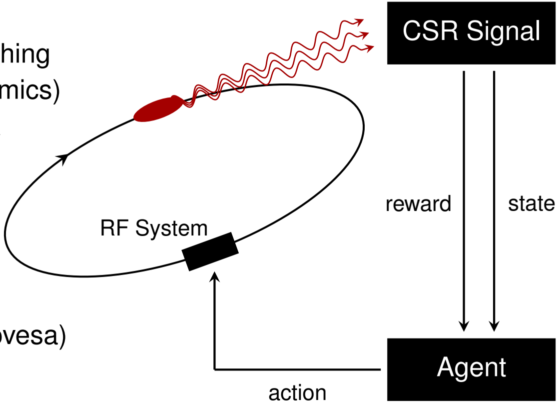

proof of principle:
control in simulation (Inovesa)

implementation:
THz diagnostics (KAPTURE) and RF system at KARA

Micro-Bunching Control

Observation, Reward and Action

- observation: state of electron bunch / micro-bunching dynamics

  1) in theory/simulations: longitudinal charge distribution

  2) measurable: CSR power signal $\Rightarrow$ feature vector encoding the state

- reward function: optimization of emitted CSR signal

  $R = w_1 \mu_{\text{CSR}} - w_2 \sigma_{\text{CSR}} \quad \text{with} \quad w_1, w_2 > 0$

  $\Rightarrow$ maximize μ_{CSR}, minimize σ_{CSR}!

- action: RF amplitude modulation

  dynamically adjust amplitude & frequency

  $\Rightarrow$ counteract perturbation by CSR

  $\Rightarrow$ include RF phase?

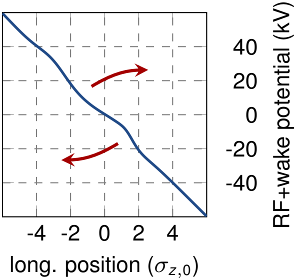

Micro-Bunching Control

…behind the curtain: Actor-Critic System using NNs

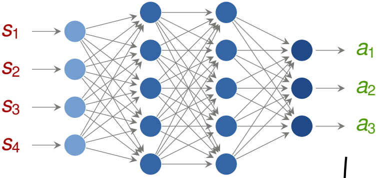

actor network

chosen action to be evaluated

evaluation and update based on estimation of expected reward $\hat{q}(s, a)$

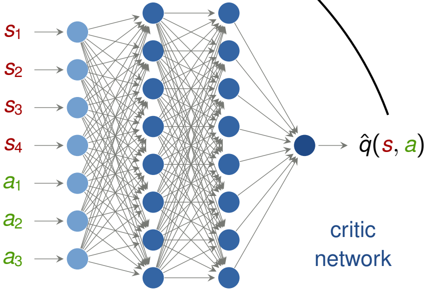

critic network

state $s = (s_1, s_2, s_3, s_4)^T$
action $a = (a_1, a_2, a_3)^T$

e.g. DDPG algorithm*
(*Deep Deterministic Policy Gradient*)

***Continuous control with deep reinforcement learning*, Lillicrap, T.P. *et al.* (2015), `https://arxiv.org/abs/1509.02971`

Micro-Bunching Control

Summary and Status Quo

- goal: control micro-bunching $\Rightarrow$ optimize CSR emission
- CSR self-interaction: perturbation is dependent on the state of the micro-bunching dynamics
  $\Rightarrow$ countermeasure should be state-dependent as well!
- CSR wake potential causes perturbation of effective restoring force
  $\Rightarrow$ compensate via dynamic RF amplitude modulation scheme
- interaction with the bunch changes the micro-bunching dynamics
  $\Rightarrow$ sequence of actions is required *(deal with consequences)*

Status Quo

- ongoing efforts to train an RL agent on simulation data
- FPGA development to meet required time constraints
- connection of THz detectors with BBB system at KARA

Optimization of Injection

Automation of Parameter Tuning

- injection rate from booster at KARA is rather low
- manual trial-and-error tuning of injection is time consuming and might not result in the optimal condition

⇒ application of Bayesian optimization

Relation to Reinforcement Learning

- trial-and-error learning paradigm
- stationary optimization (not a sequence of actions)
- elementary sub-problem in RL (*multi armed bandit problem*)

⇒ similar approach to quadrupole tuning at LCLS II[1] and SwissFEL[2]

[1] McIntire, M.W. *et al.* (2016), IPAC16-WEPOW055 [2] Kirschner, J. *et al.* (2019), `https://arxiv.org/abs/1902.03229`

Optimization of Injection

Bayesian Optimization using Gaussian Processes (UCB)

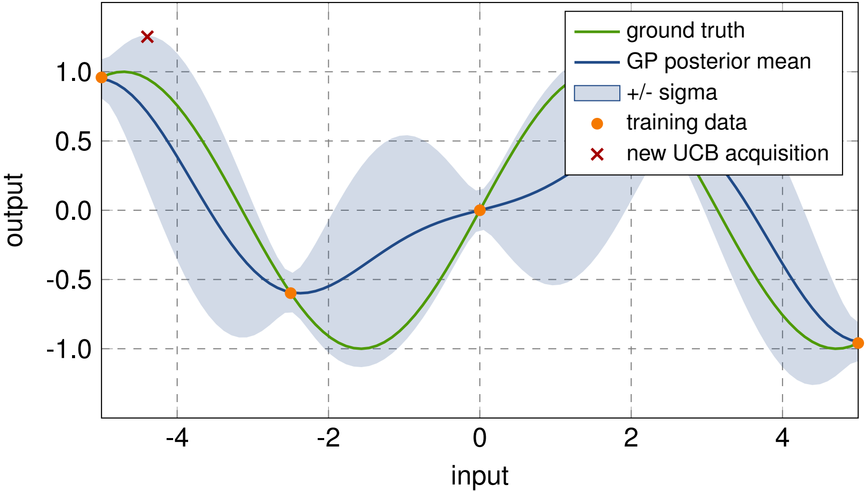

Optimization of Injection

First Results on Simulations

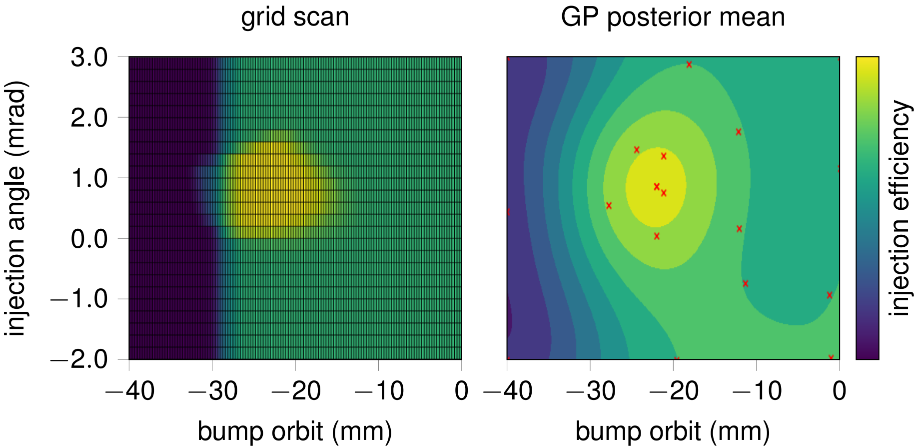

2 parameters: bump orbit (combined kicker strength), injection angle (septum strength)

Optimization of Injection

Summary and Outlook

- goal: automatic tuning of injection parameters
- Bayesian optimization using Gaussian Processes performs reasonably well on simulations (with 2-3 parameters)

Outlook

- adding more parameters, e.g. corrector magnets
- apply algorithm at the real machine, integrate intro control system
- benchmark with other algorithms: random search, genetic algorithm, gradient based, . . .