

Reinforcement Learning for Intensity Tuning at Large FEL Facilities

4th ICFA Machine Learning Workshop




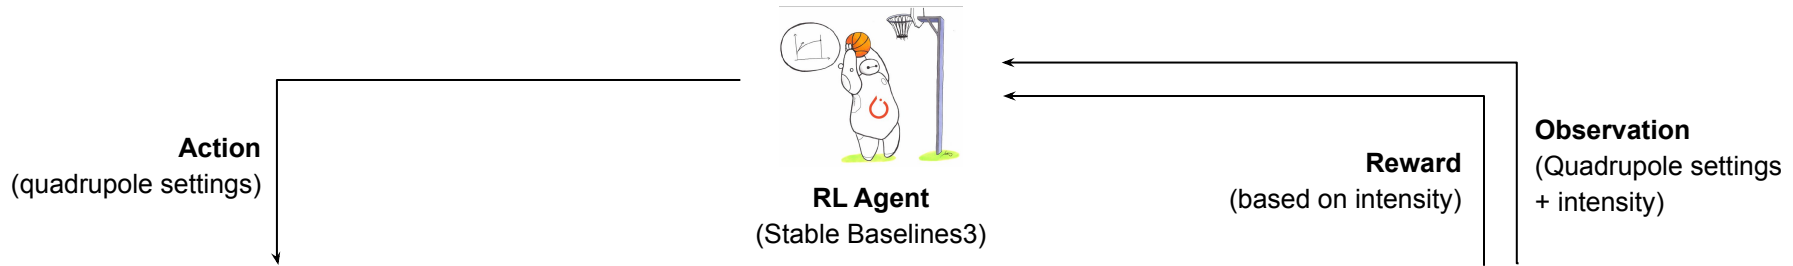
Jan Kaiser, Annika Eichler, Auralee Edelen, Daniel Ratner, Malachi Schram and Kishansingh Rajput
Gyeongju, 7 March 2024

FEL Intensity Tuning at LCLS

Reinforcement learning-trained optimisation (RLO)

- **Maximise FEL intensity** using 14 quadrupole magnets
- Challenges:
 - Slow simulation -> **Cheetah** and **neural network surrogate modelling**
 - High dimensionality and easy failure -> **Curriculum learning**

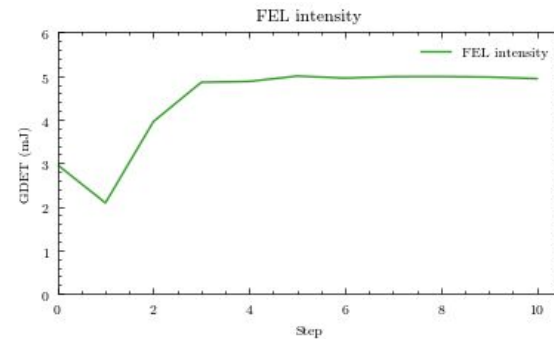
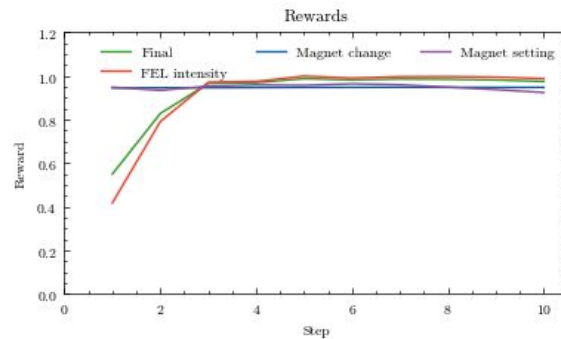
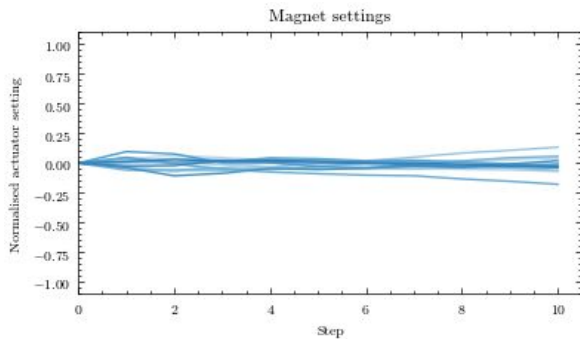
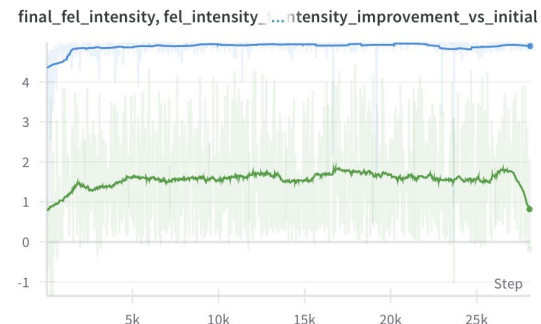
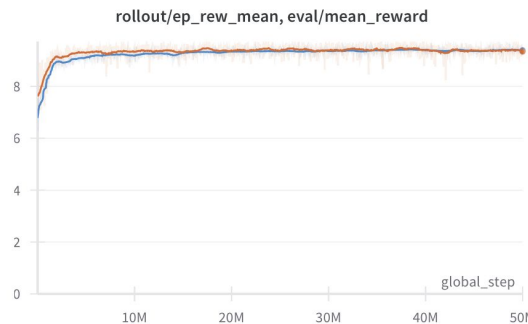
Ultimately also
transfer to



Results with PPO


It works!

- **Proximal Policy Optimisation (PPO)** algorithm from Stable Baselines3
- Training for **50 Million** environment interactions
- **1 day 16 hours** on a HPC cluster node



Gradient-based Reinforcement Learning

Significantly improving sample-efficiency and reducing training times



We are using **Cheetah!**


Cheetah supports **automatic differentiation.**



Gradient-based RL with true policy gradient.



Achieve same performance in **45x fewer samples.**



Contact

DESY. Deutsches
Elektronen-Synchrotron

www.desy.de

Jan Kaiser
Machine Beam Controls (MSK)
jan.kaiser@desy.de