Contribution ID: **7** Type: **Poster**

Model Predictive Control with Gaussian Processes for Safe and Efficient RL: A Case Study at the CERN SPS

Classical, model-free Reinforcement Learning (RL) has achieved impressive results in areas where interactions with the environment are inexpensive, such as computer games or simulations. However, in many real-world applications, such as robotics or autonomous particle accelerators, interactions with the system are costly, which creates a need for sample-efficient RL algorithms. In addition, safety constraints must be respected even during the exploratory phases of the RL agent. To address both concerns, Reinforcement Learning based on Model Predictive Control using Gaussian Processes (GP-MPC) was proposed by Kamthe and Deisenroth in 2018. We present an implementation of this approach that is compatible with the Gymnasium API and test it on the problem of correcting power supply ripples (multiples of 50 Hz), which negatively affect the slow extraction spill from the CERN Super Proton Synchrotron (SPS) to the fixed-target physics experiments hosted in the North Experimental Area (NA) at CERN.

Author: FERBER, Ferdinand (CERN)

Presenter: FERBER, Ferdinand (CERN)

Session Classification: Poster session