Reinforcement Learning at the Workshop:

- Tutorial
  - Tutorial on Reinforcement Learning
    - Speaker: Kishansingh Rajput (JLAB)
- RL Session Talks:
  - Machine Learning for Slow Spill Regulation in the Fermilab Delivery Ring
    - Mattson Thieme (NorthWestern)
  - Reinforcement Learning Applied to Optimization of LHC Beams in the CERN Proton Synchrotron
    - Joel Axel Wulff (CERN)
  - A Closer Look at Reinforcement Learning for Beam-Based Feedback Systems
    - Leander Grech (Univ. of Malta)
- Other Session Talks:
  - GeoFF/COI and How to get ML into the Control Room - the CERN ML Frameworks
    - Nico Madysa
  - Uncertainty Aware Anomaly Detection to Predict Errant Beam Pulses and GradCAM Analysis
    - Kishansingh Rajput (Jefferson Lab)
  - Machine Learning Tools to support the ATLAS Ion Linac Operations at Argonne
    - Jose L. Martinez-Marin (ANL)
- Posters:
  - Learning-based Optimisation of Particle Accelerators Under Partial Observability Without Real-World Training
    - Jan Kaiser (DESY)
  - Apply Machine Learning in Orbit Control and Accelerator Stabilization
    - Zeyu Dong (Stony Brook University)
  - Machine Learning Tools to support the ATLAS Ion Linac Operations at Argonne
    - B. Mustapha

Reinforcement Learning Applications

- Tune beam magnets (DESY, FNAL, LHC)
- Focus heavy ion linac beam (ATLAS)
- Maximize beam transmission
- Control rate of proton extraction (FNAL)
- Proton synchrotron bunch spacing (CERN)
- Control electron orbit (NSLS-II)
- Commissioning (CERN)
- Replace linear and symmetric control systems

Themes

- Train on simulation, transfer to real world
- Several work-in-progress implementations
- Training and tuning can be lengthy
- Seek sample efficient RL approaches
- RL framework ecosystem for optimization and control (eg. GeOFF)
- Compare to other optimization/control methods
- Use RL in concert with other techniques (eg CNN for feature extraction)
- Divide RL into domain-specific agents
- Desire for increased transferability of models – tend to be brittle and specific to a domain
- Seek better understanding of limits of RL for complexity and generality

Discussion

- Questions/comments/feedback on optimization or a specific subtopic?
- What is missing that the community needs?
- How can the community collaborate, educate, exchange information to address those needs?
- Is this workshop representative of state-of-the-art for optimization techniques for this domain?