IAEA Technical Meeting on AI for Nuclear Technology and Applications

Machine Learning Applications at the CERN particle accelerators

P. Arpaia, G. Azzopardi, F. Blanc, X. Buffat, L. Coyle, E. Fol,
F. Giordano, M. Giovannozzi, V. Kain, N. Madysa, T. Pieloni, R. Prevete,
S. Redaelli, B. Salvachua, B. Salvant, M. Schenk, M. Solfaroli Camillocci,
R. Tomas, **G. Valentino**, F. Van der Veken, F. Velotti, J. Wenninger

L-Università ta' Malta

CERN

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Università degli Studi di Napoli Federico II

Outline

- Why is ML useful for particle accelerators?

- Review of successful use cases at CERN
  - Anomaly detection in simulations
  - Collimator beam-based alignment
  - Sample-efficient reinforcement learning
  - Tune estimation

Why is ML useful for accelerators?

- **Anomaly detection and machine protection**
  - High energy / high intensity machines are equipped with machine protection systems (MPS) which should extract the beam from the machine before catastrophic damage can occur.
  - MPS are critical and therefore hardware-based.
  - However, machine learning can be used to capture **operational issues** which impact beam quality, or **precursors of faults** which could lead to downtime (e.g. for machine refilling).

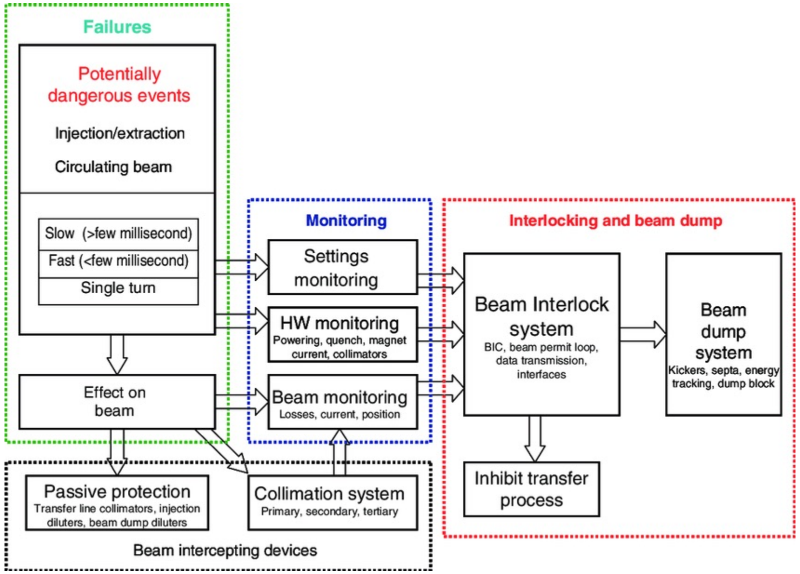

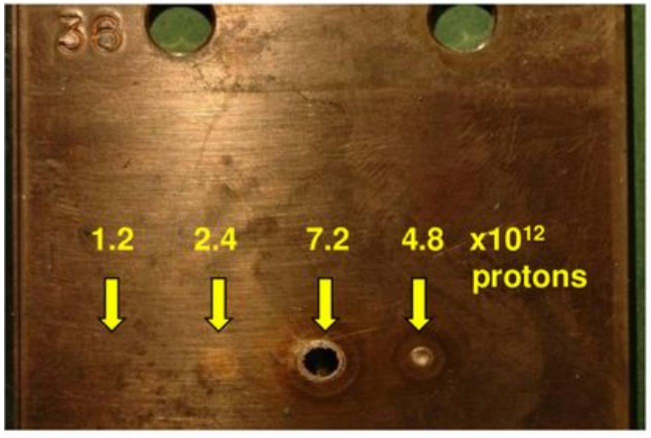

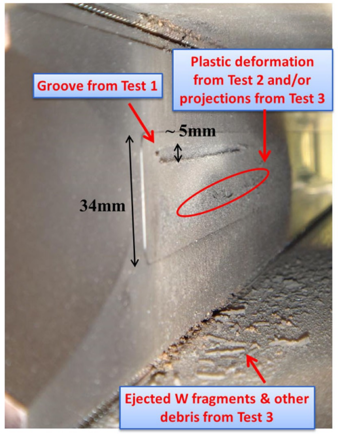

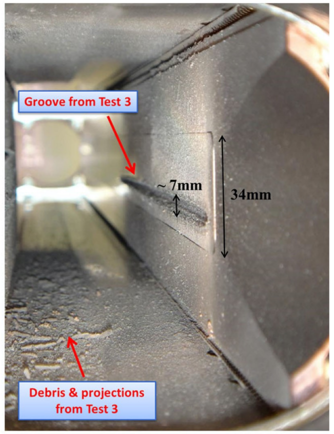

Why is ML useful for accelerators?

- **System modeling**
  - A challenge in accelerators is to have machine models which are
    - Sufficiently accurate
    - Can execute quickly enough to be useful during operation, using real input from beam instrumentation (i.e. **online modeling**).
  - Most simulation tools are too slow to be used in control systems or provide guidance to operators during machine operation.
  - System modeling through ML allows to learn representations that combine information from physics-based simulations with measured data
    - aka **surrogate modeling**

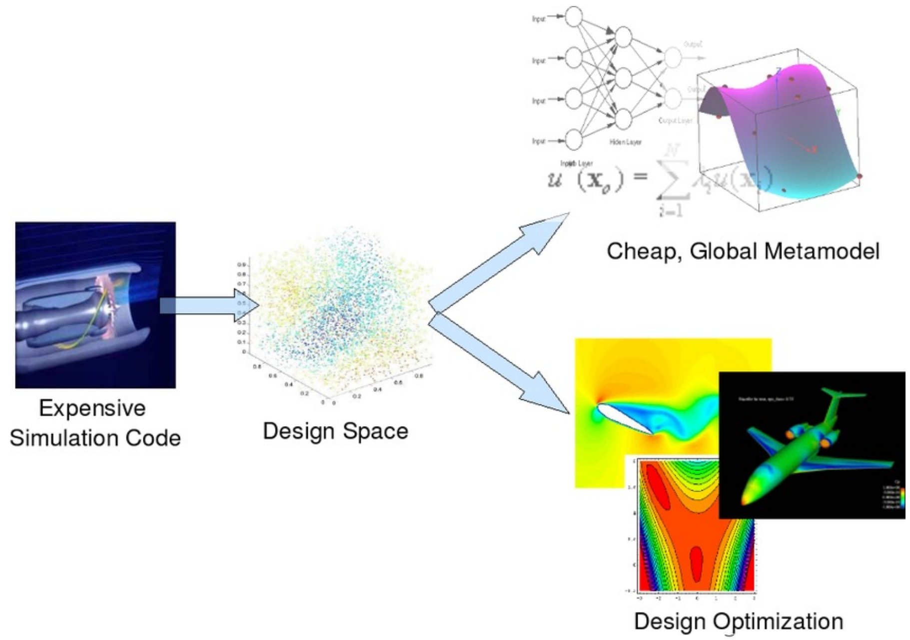

Why is ML useful for accelerators?

- **Tuning and control**
  - Operational settings tend to be established through theory and optimized during beam commissioning.
  - These settings may need to be tuned over time due to effects of ground motion, changes in beam parameters, radiation to electronics etc.
  - Accelerators may have dozens of operational modes and complex operational cycles.
  - A lot of beam time is spent by operators ‘reconfiguring’ the machine to desired parameters.

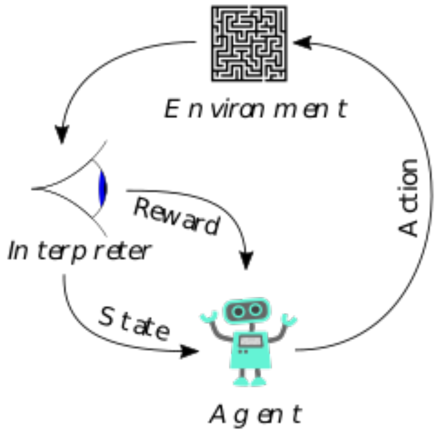

Reinforcement Learning and optimization techniques are well suited for these types of problems, and can be easily re-trained over time

Issues: sample efficiency and transfer learning

Why is ML useful for accelerators?

- **Advanced data analysis**

  - Particle accelerators generate huge volumes of data
    - E.g. LHC: 1.5 million signals are logged from all the accelerator sub-systems = 2 TB/day

  - Data analysis and visualization is a constant aspect of operation: to validate performance, try to diagnose and understand faults etc

  - Unsupervised learning techniques could be very useful to understand hidden structures in data, e.g. understand which machine parameters contribute to beam losses

  - If 'unknown' issues are uncovered, mitigation measures can be taken to improve machine performance and availability.

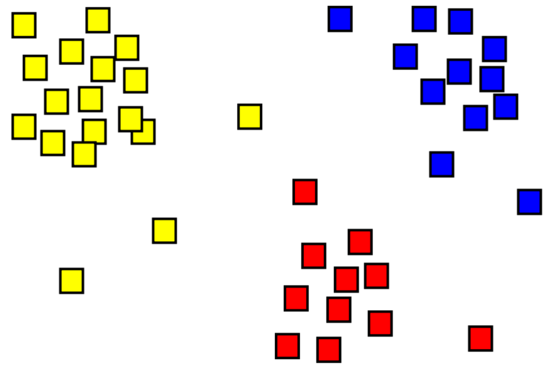

ML in accelerators: review of case studies

1. Anomaly detection (DA)

2. Collimator beam-based alignment

3. Sample efficient particle accelerator control

4. Improved LHC tune estimation

These are only some examples – for several others e.g. optics measurements and corrections, optimisation of beam lifetime and losses, detection of instabilities see P. Arpaia et al., "Machine learning for beam dynamics at the CERN Large Hadron Collider", NIM A vol. 985, 2021.

1. Anomaly detection for Dynamic Aperture simulations

- DA is region of phase space in which the particle's motion remains bounded over a finite number of turns.
- Outlier identification is important in DA simulations as outliers may have an impact on the DA_{min}.
- They may be due to the excitation of particular resonances as a result of the distribution of nonlinear magnetic errors, which is highly seed-dependent.

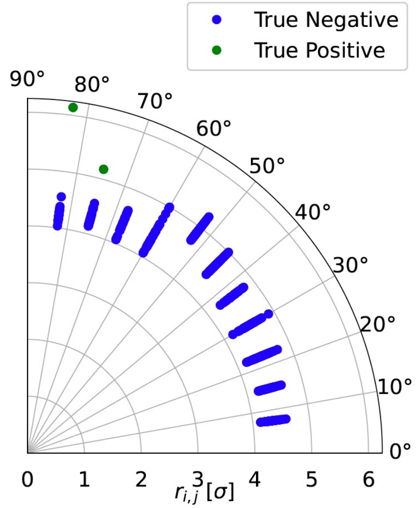

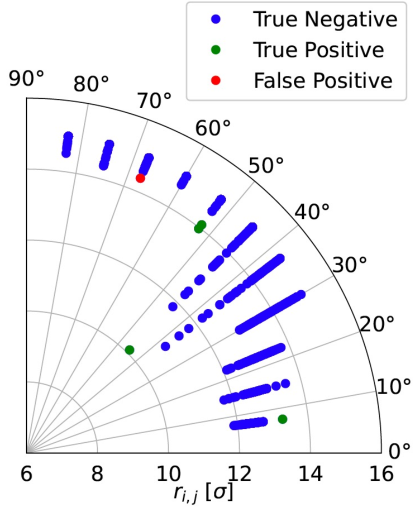

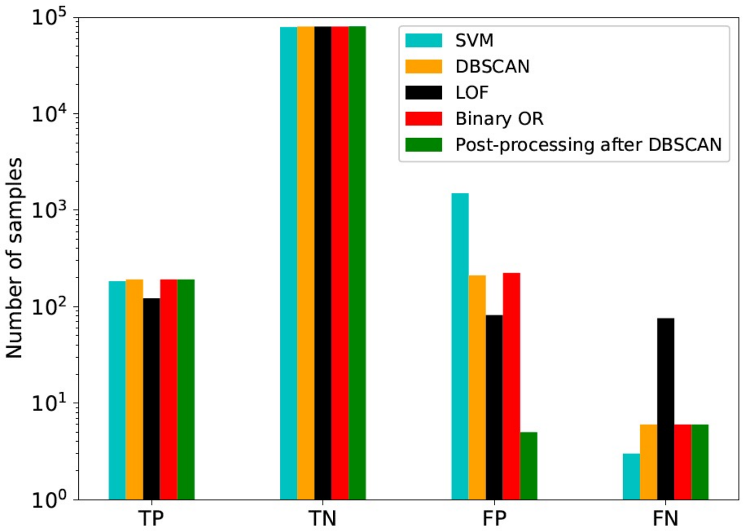

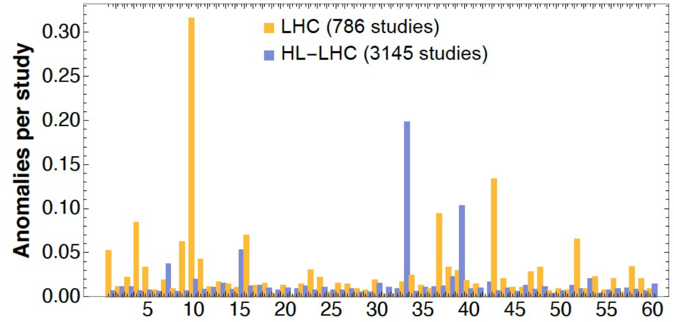

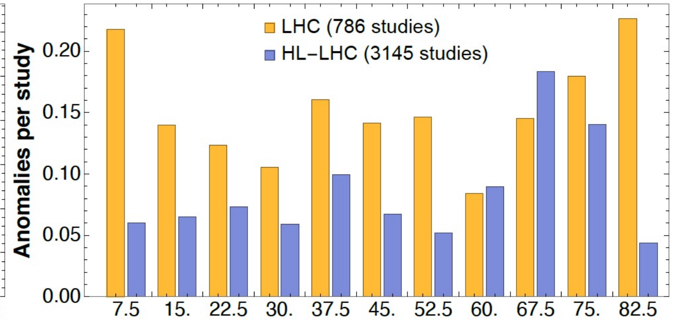

M. Giovannozzi et al., "Machine Learning applied to the analysis of nonlinear beam dynamics simulations for the CERN Large Hadron Collider and its luminosity upgrade", MDPI Information, 2021.

2. LHC Collimator Alignment

- Almost 100 LHC collimators need to be aligned to the beam at various points of the machine cycle to establish their operational settings:

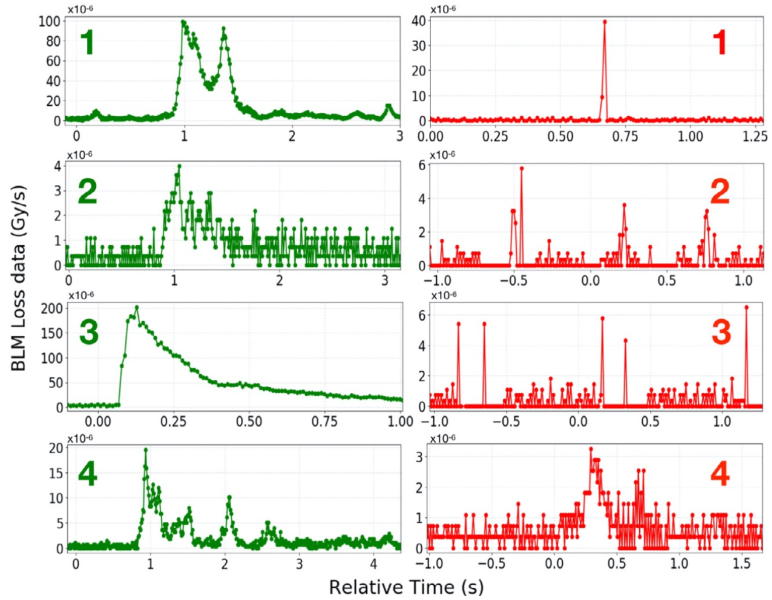

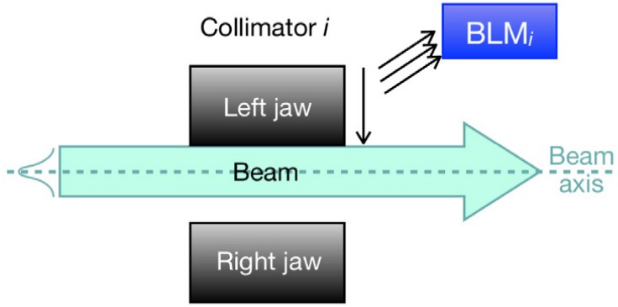

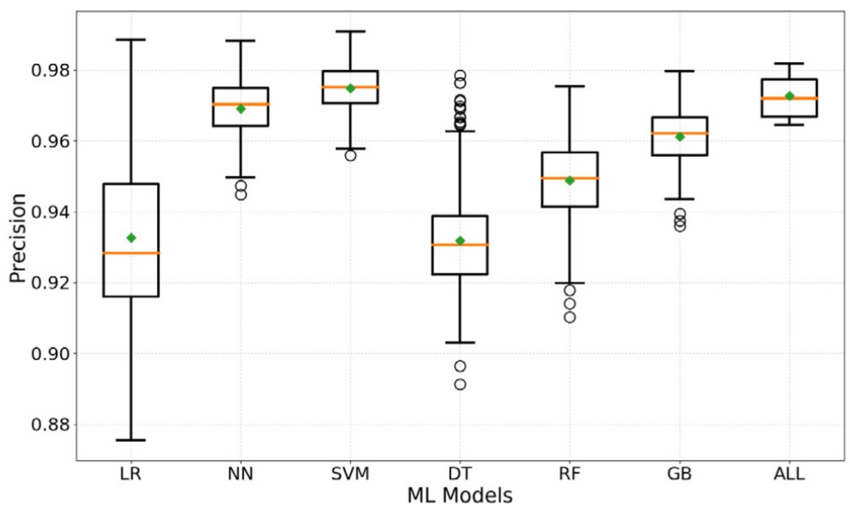

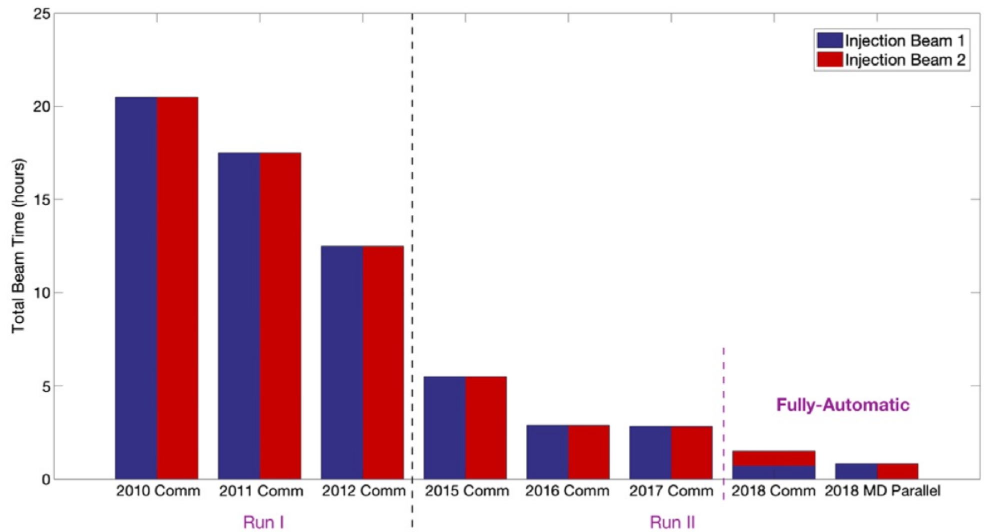

G. Azzopardi et al., “Automatic spike detection in beam loss signals for LHC collimator alignment”, NIM A vol. 934, pp. 10-18, 2019.

3. Sample-efficient reinforcement learning

Advanced Proton Driven Plasma Wakefield Acceleration Experiment (AWAKE)

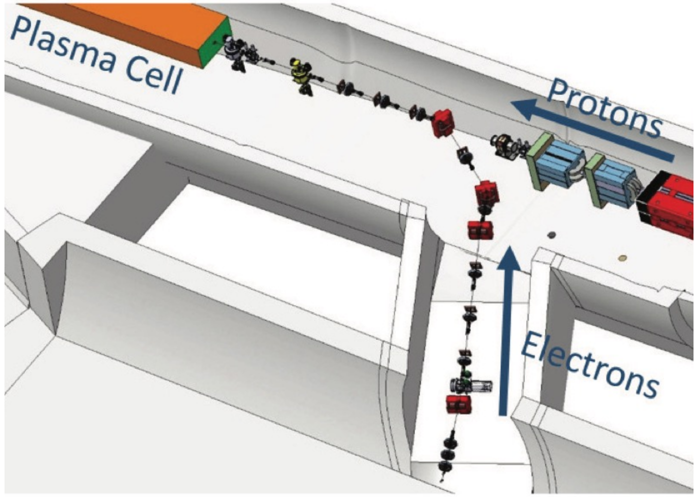

- An orbit feedback system is necessary to ensure that the particles are on the correct trajectory (orbit):

$$\Delta x = W \Delta\theta$$

where the response matrix $R = U.W.V^T$

- Can RL be used to learn a more efficient policy?

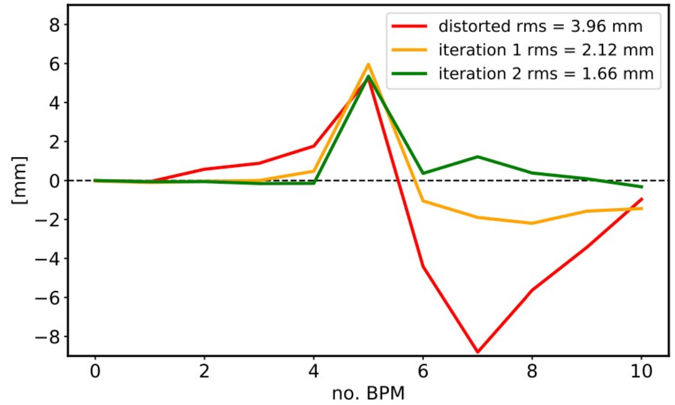

- **Reward** = Beam Position Monitor (BPM) RMS
- Trained continuous model-free Normalized Advantage Function (NAF) agent modified to have Prioritized Experience Replay
  - **PER-NAF:** experience sampled from replay buffer
- Training performed on both real machine and in an OpenAI simulation environment using a response matrix generated through MAD-X.

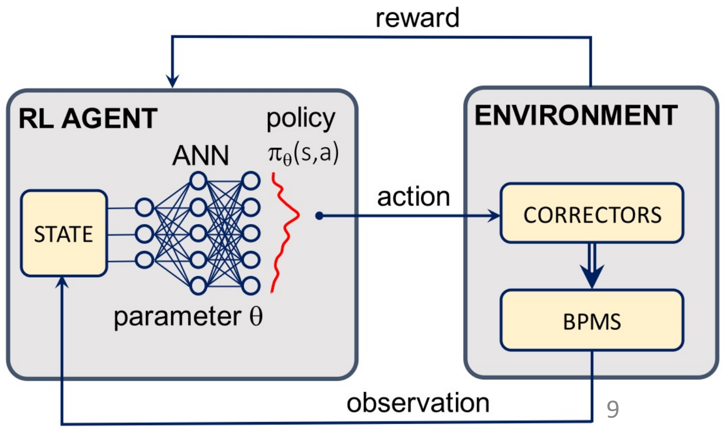

3. Sample-efficient reinforcement learning

Training of NAF agent on the machine

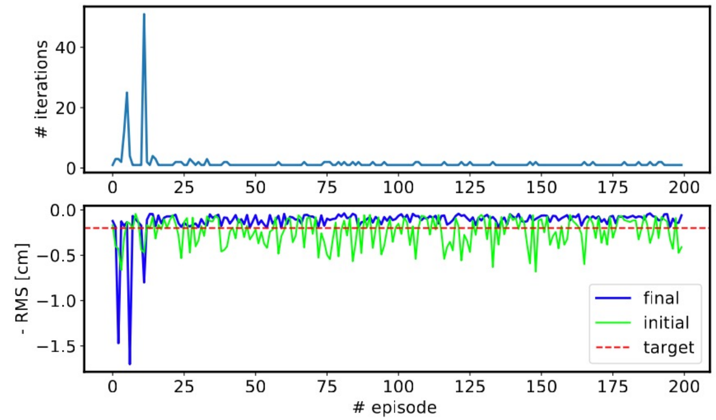

DOF	11
Reward target [cm]	–0.2
Max episode length	50
Max Δ_{corr} [μrad]	300
Min allowed reward [cm]	–1.2

Validation results on the machine following training on simulation

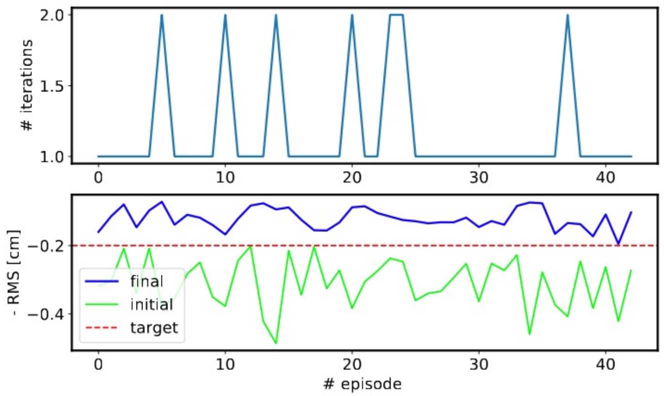

Validation results on the machine following training on the machine

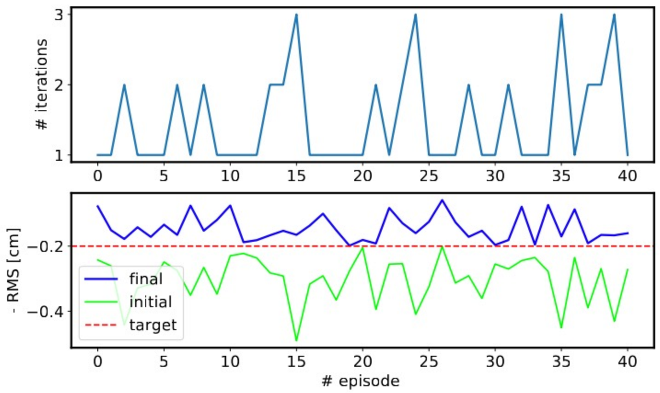

V. Kain et al., “Sample-efficient reinforcement learning for CERN accelerator control”, Phys Rev Accel Beams 23, 124801 2020

4. Tune estimation in the LHC

- The LHC tunes in H and V, B1 and B2 are measured by observing turn-by-turn betatron oscillations using a beam position monitor.

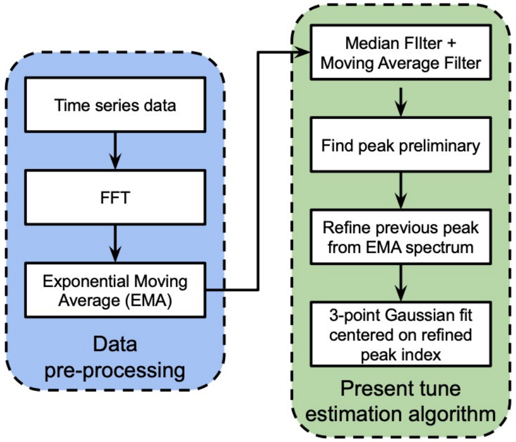

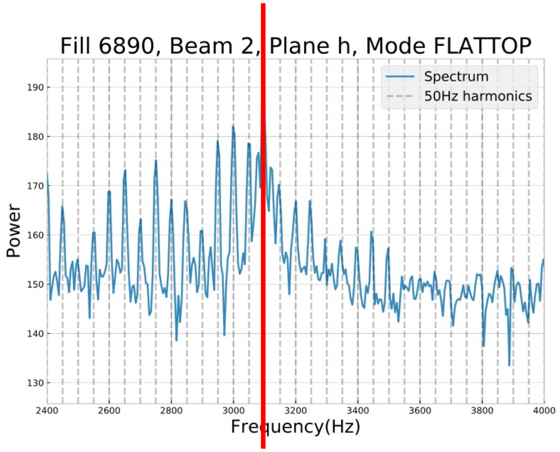

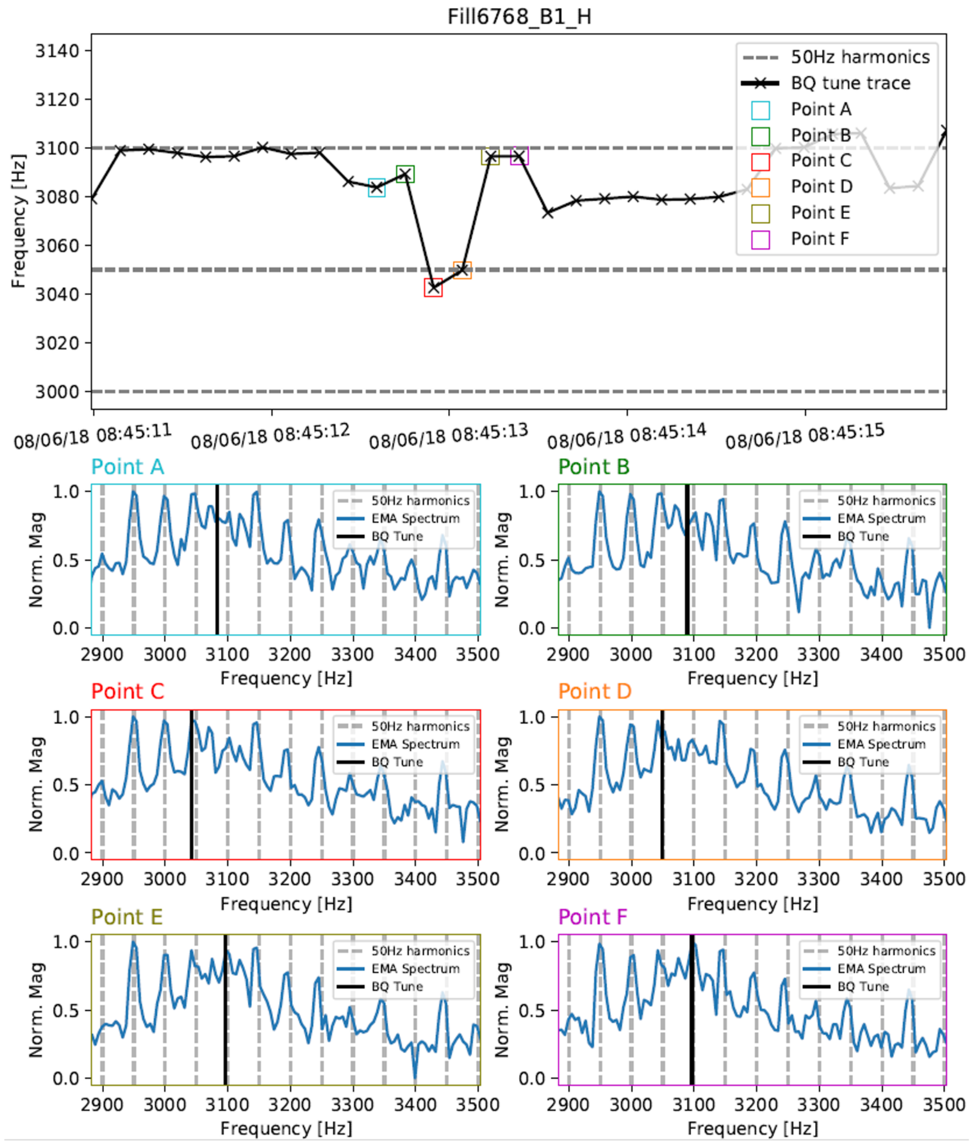

- 50 Hz harmonics
  - Present since start of LHC – due to main dipole magnets.
  - Harmonics perturb the spectrum, which affects reliability of the tune estimates.
  - Unstable tunes cause the Tune Feedback (QFB) system to switch itself off as a preventive measure.

L. Grech et al., MDPI Information, 2021.

4. Tune estimation in the LHC

Second order system simulation of real BBQ spectra:

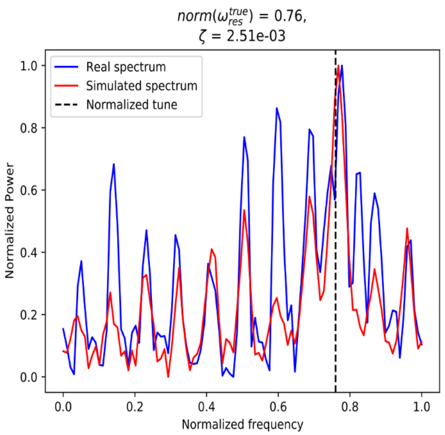

Variant of GAN called SimGAN used to improve simulated spectra:

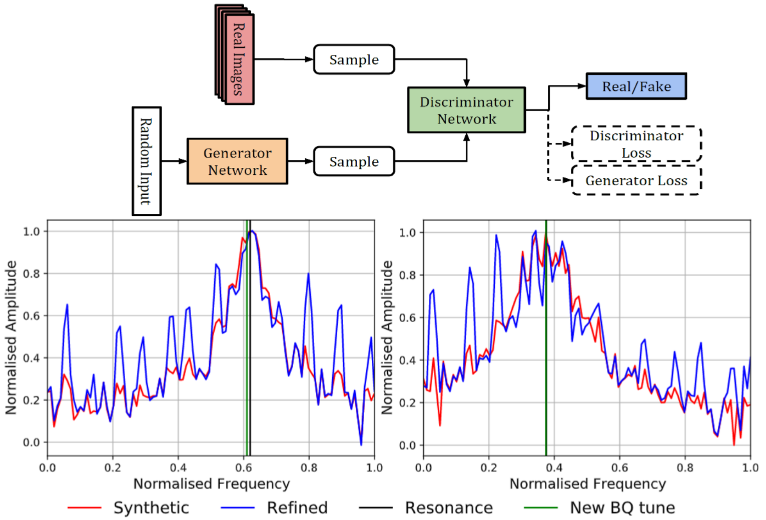

First spectra were simulated, passed through trained SimGAN, and a dataset was created to train ML models.

ML-Refined has same architecture as best simple model (Model #1) but trained over improved SimGAN dataset.

ML-Refined shows better stability than the current tune estimation algorithm, meaning better tune control in LHC.

$$Stability \;=\; q_{t+1} * (1-\alpha) + q_t * \alpha$$

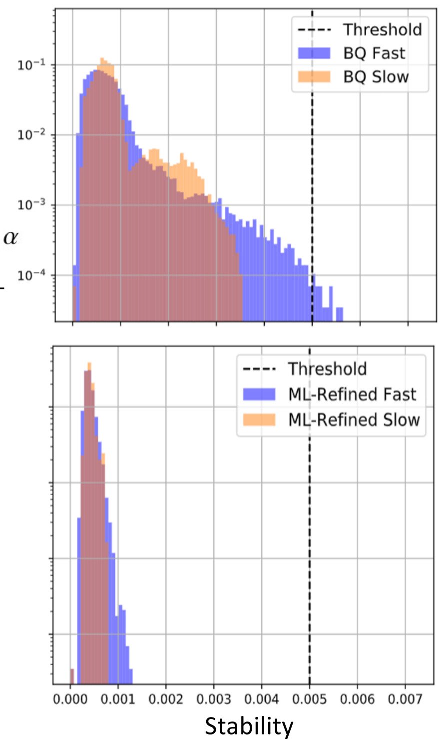

Simple approach
Simulated data trains DNN and CNN

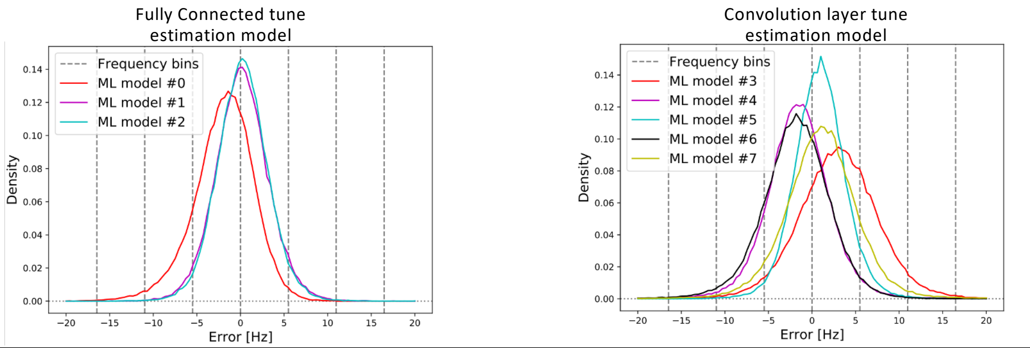

Some limitations, over-fitting to simulated training data:

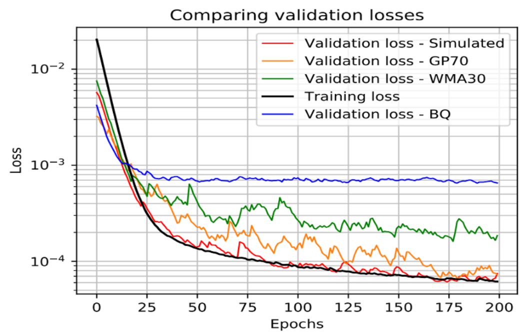

Conclusions

- The past few years have seen an high growth in the take-up of ML by the accelerator community, driven by:
  - Deep learning developments
  - Increase in scale and complexity of machines
  - Availability of data
- See “opportunities in ML for particle accelerators”: https://arxiv.org/abs/1811.03172
- Some of the latest activities involving anomaly detection, pattern recognition and particle accelerator control problems were reviewed.
- ML will be a key tool to help meet demands for higher beam brightness, energy and intensity.