Offline Reinforcement Learning-Based Control of LEIR Injection Efficiency via Data-Driven Surrogate Modeling

Borja Rodriguez Mateos

CERN & UPC

April 3, 2025

Acknowledgements: Michael Schenk, Verena Kain, Penny Madysa, Simon Hirländer, Theodoros Argyropoulos, Felix Carlier, Maciej Slupecki

Outline

Low Energy Ion Ring (LEIR) controls

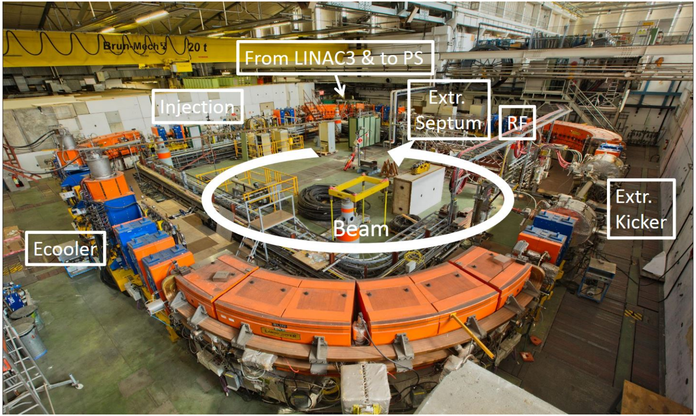

Supervision and operation of the machine:

- Complex system by design
- Many hours of manual maintenance/recovery of performance
- High repetition rate
- Low energy beam

→ Perfect testbed for auto-pilots!

Challenges of LEIR injection

- Injection efficiency into LEIR severely depends on energy distribution of the incoming pulses from LINAC3
  - Accumulation of several pulses with multi turn injection using 6D phase-space painting
  - Cool and drag in momentum space the beam already circulating in LEIR
- Stripping foil has a limited lifespan and degrades over time
  - Impacts the mean ion energy and changes the energy distribution width
  - Decreases the injection efficiency over time and needs to be compensated manually
- We aim to develop stateful optimizers for maintaining or re-establishing desired intensity in LEIR

Stateful optimization

Stateful optimization of injected beam intensity

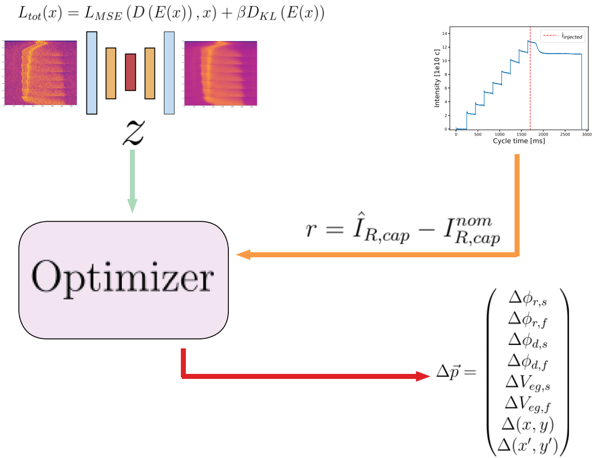

- Training a VAE of longitudinal Schottky spectra
  to obtain encoded latent vector **z**
- Optimizer observes encoding of Schottky spectra **z** and intensity reading from the BCT
- Optimizer modifies relevant parameters for improving the intensity in the ring

Simulating the longitudinal LEIR injection

Injection from Linac3 to LEIR is complex to simulate:

- Coasting beam with multi turn injection $\rightarrow$ no RF focusing

$$\Delta f_{rev} = -f_0 \frac{\Delta P}{P_0} = 0$$

- Interplay of several collective effects: electron cooling, impedance, IBS, and space charge

- Possible to compute the simulated equilibrium longitudinal Schottky spectrum accurately

- Takes several hours to simulate a few 10k turns of multi injection coasting beam with high accuracy

$\rightarrow$ Need to develop a reliable data-driven surrogate model for fast training!

Data-driven surrogate model

Building the data-driven surrogate model

Due to high dimensionality of action space and limitation of resources, data collection needs to be done efficiently:

Static sampling methods:

- Latin Hypercube sampling (ensures stratification in each dimension)
- Sobol sampling (low discrepancy sequence)

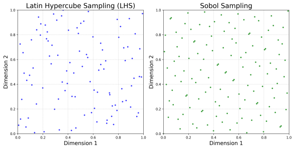

Dynamic sampling methods:

- Bayesian Optimization on the injection efficiency with high β UCB acquisition function $a_{\mathrm{UCB}}(x;\beta) = \mu(x) + \sqrt{\beta}\sigma(x)$

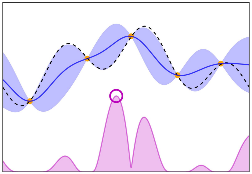

Observation space: Encoding the longitudinal Schottky spectra

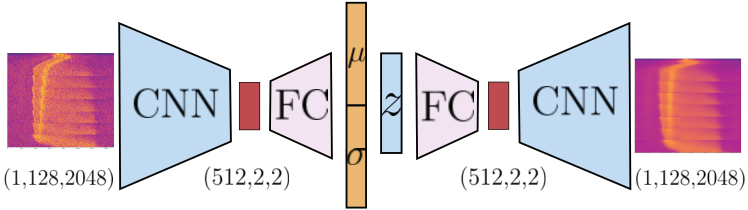

Loss used for training :

$$L_{tot}(x) = L_{MSE}(D(E(x)), x) + \beta |D_{KL}(E(x)) - C|$$

where C is increased throughout the training to control the encoding capacity of the VAE's bottleneck

- Longitudinal Schottky spectra has a dimension of $\approx (1, 90, 1514)$
- Need to reduce it for using it as observation for RL agent
- Found optimal latent dimension for reconstruction error $\mathbf{z} \in \mathbb{R}^{20}$
- Right balance for reconstruction vs disentanglement with $\beta = 0.05$

Action space: Spanning the whole parameter space

List of actionable parameters:

- Debunching and ramping cavity phases (3 parameters)
- Horizontal and vertical cooler bumps (4 parameters)
- Electron gun voltage (2 parameters)
- Horizontal and vertical injection bumps (4 parameters)
- Injection corrector magnet intensity per pulse (n_{pulse} parameters)

- The whole possible parameter space is explored
- Two sets of problems with different action space dimensions:
  - $d_a = 9$
  - $d_a = 21$ (for Pb ions)
- Will only focus here on purely longitudinal problem with $d_a = 9$

Offline training and evaluation of RL agent

Offline training of an RL agent on the surrogate model

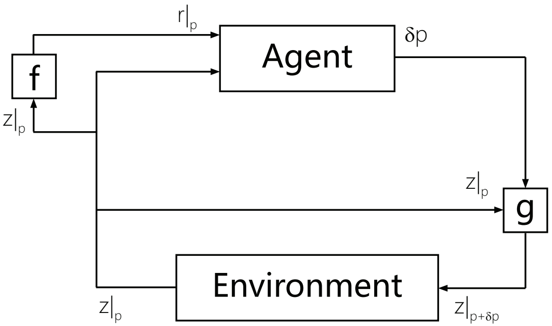

Two models are used for the training :

- Reward model : $f(\mathbf{z}|_{\mathbf{p}}) = r|_{\mathbf{p}}$
- Transition model : $g(\mathbf{z}|_{\mathbf{p}}, \delta\mathbf{p}) = \mathbf{z}|_{\mathbf{p}+\delta\mathbf{p}}$

Loss used for training the transition model:

$$\begin{aligned} L =& ||g(\mathbf{z}|_{\mathbf{p}}, \delta p) - \mathbf{z}|_{\mathbf{p}+\delta\mathbf{p}})|| \\ &+ \gamma_1 ||g(g(\mathbf{z}|_{\mathbf{p}}, \delta p), -\delta p) - \mathbf{z}|_{\mathbf{p}}|| \\ &+ \gamma_2 ||g(\mathbf{z}|_{\mathbf{p}+\delta\mathbf{p}}, -\delta p) - \mathbf{z}|_{\mathbf{p}}|| \\ &+ \gamma_3 ||g(\mathbf{z}|_{\mathbf{p}}, 0) - \mathbf{z}|_{\mathbf{p}}|| \\ &+ \gamma_4 ||g(\mathbf{z}|_{\mathbf{p}+\delta\mathbf{p}}, 0) - \mathbf{z}|_{\mathbf{p}+\delta\mathbf{p}}|| \end{aligned}$$

Loss used for training the reward model:

$$L_{\text{reward}} = ||f(\mathbf{z}|_{\mathbf{p}}) - r|_{\mathbf{p}}||$$

→ Still needs to be tested online and evaluate quality of simulation-to-reality transfer

Application of common RL algorithms to the offline LEIR injection task

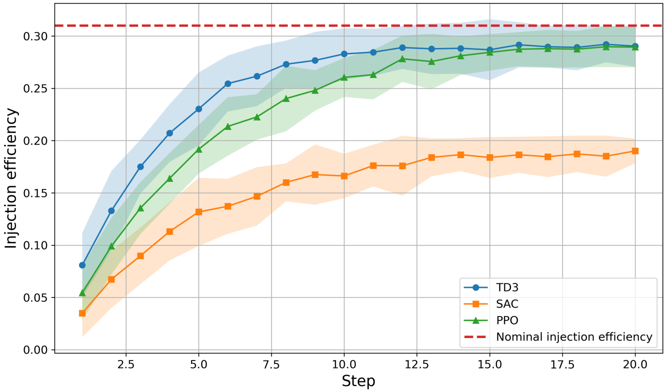

- Reduced action step size for testing quality of trained transition model, possible to solve the task in 1 step.
- Evaluated trained agents over 50 iterations for common on and off-policy DRL algorithms.
- Reward corresponds to the injection efficiency between Linac3 and LEIR

$$r(i_{injected}, i_{L3}) = \frac{i_{injected}}{i_{L3}} = \frac{i_{injected}}{\sum_{j=1}^{8} i_j}$$

Outcomes and future plans

Current outcomes of RL on LEIR injection surrogate model

- Data-driven surrogate models extrapolate poorly: during training, need to clip observations and rewards to values within the distribution of training set

- Small errors accumulate over several steps making sim2real transfer less accurate for optimizers performing a large number of steps

- No accurate indication of stripping foil age in observation $\rightarrow$ plan to add time of flight measurements

Future updates

- Plan to apply it to the Oxygen run of 2025 starting mid-May and the Lead run early September
- Work around clipping surrogate model by using normalizing flows for transfer model
- Use fancier models once transfer model and classic DRL agents reliably obtain nominal injection efficiency

Thank you for your attention

References

Biancacci. "Advancements in injection efficiency modelling for the Low Energy Ion Ring (LEIR) at CERN". In: *JACoW IPAC* 2023 (2023), MOPL016. DOI: 10.18429/JACoW-IPAC2023-MOPL016.

Biancacci, N. and M. Zampetakis. "Simulation and measurement of the interplay between electron cooling, impedance, intrabeam scattering, and space charge for coasting beams". In: *Phys. Rev. Accel. Beams* 27 (1 2024), p. 014201. DOI: 10.1103/PhysRevAccelBeams.27.014201. URL: https://link.aps.org/doi/10.1103/PhysRevAccelBeams.27.014201.

Madysa. "Automated Intensity Optimisation Using Reinforcement Learning at LEIR". In: *JACoW IPAC* 2022 (2022), pp. 941–944. DOI: 10.18429/JACoW-IPAC2022-TUPOST040.