Autonomous beam flattening using reinforcement learning at the CLEAR facility at CERN

Giacomo Tangari[1,2], Wilfrid Farabolini[1], Ondrej Franek[1,3], Davide Gamba[1], Antonio Gilardi[1], Johan Leygonie[1,4,6], Alexander Malyzhenkov[1,5,6], Ciara Moore[1,9], Sara Simonsson[1,7], Sabrina Wang[1,8], Andrea Mostacci[2] and Roberto Corsini[1]

[1] European Organization for Nuclear Research (CERN), Geneva, Switzerland
[2] Sapienza University of Rome, Rome, Italy
[3] Czech Technical University in Prague, Prague, Czech Republic
[4] Université Paris-Saclay, Gif-sur-Yvette, France
[5] Swiss International Institute Lausanne (SIIL), Prilly, Switzerland
[6] Theryq, Rousset, France
[7] Lund University, Lund, Sweden
[8] University of Oxford, Oxford, United Kingdom
[9] University of Exeter, Exeter, United Kingdom

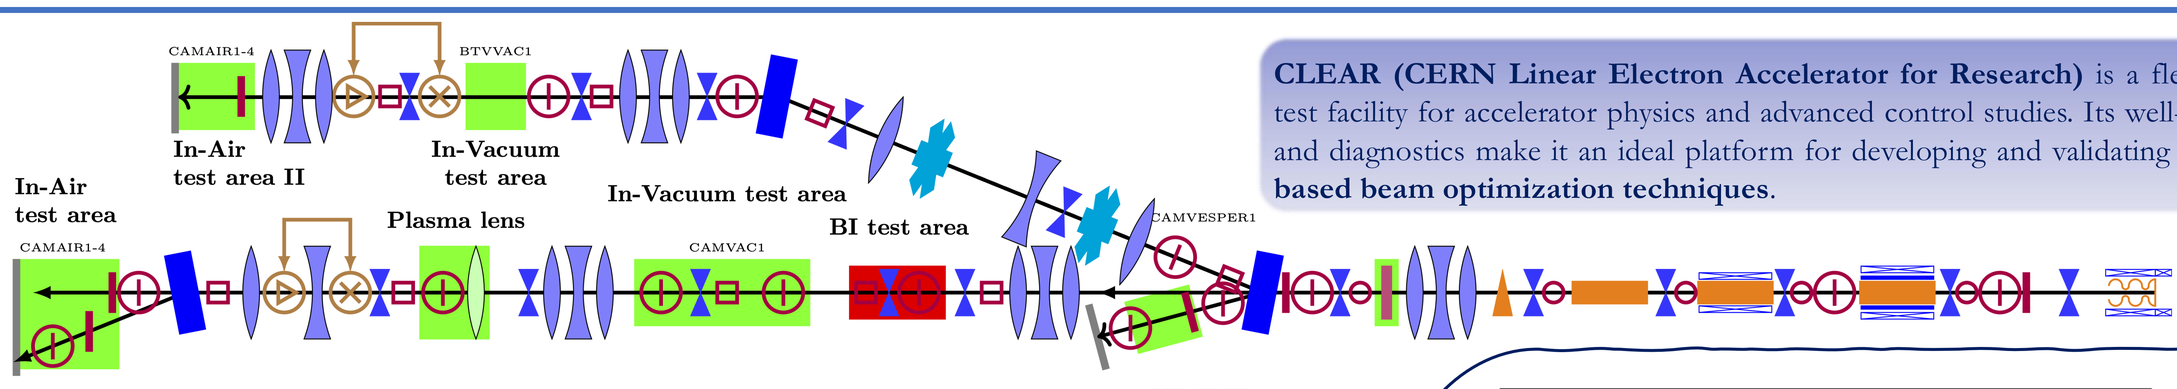

CLEAR (CERN Linear Electron Accelerator for Research) is a flexible electron-beam test facility for accelerator physics and advanced control studies. Its well-characterized optics and diagnostics make it an ideal platform for developing and validating **machine learning-based beam optimization techniques**.

Producing and maintaining a **uniform transverse beam distribution** is essential in many applications. This must be achieved under varying beam conditions (e.g. charge, optics), making the problem:

- **Nonlinear**
- **High-dimensional**
- **Time-dependent**

As a result, manual tuning of quadrupoles and correctors is **challenging and not scalable**.

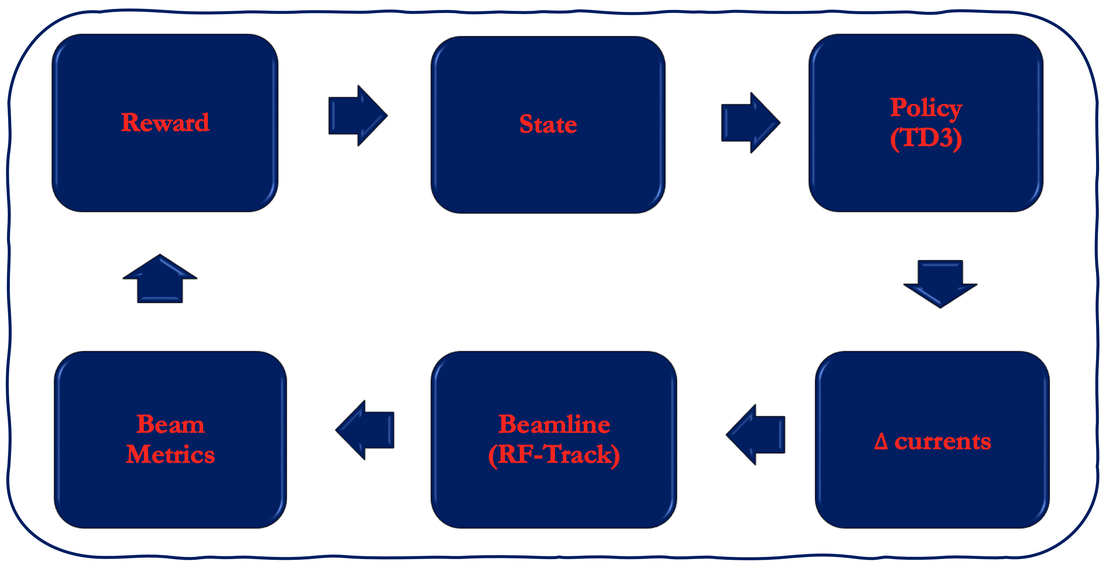

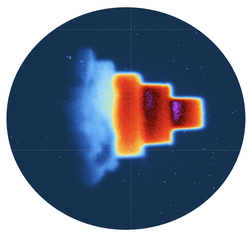

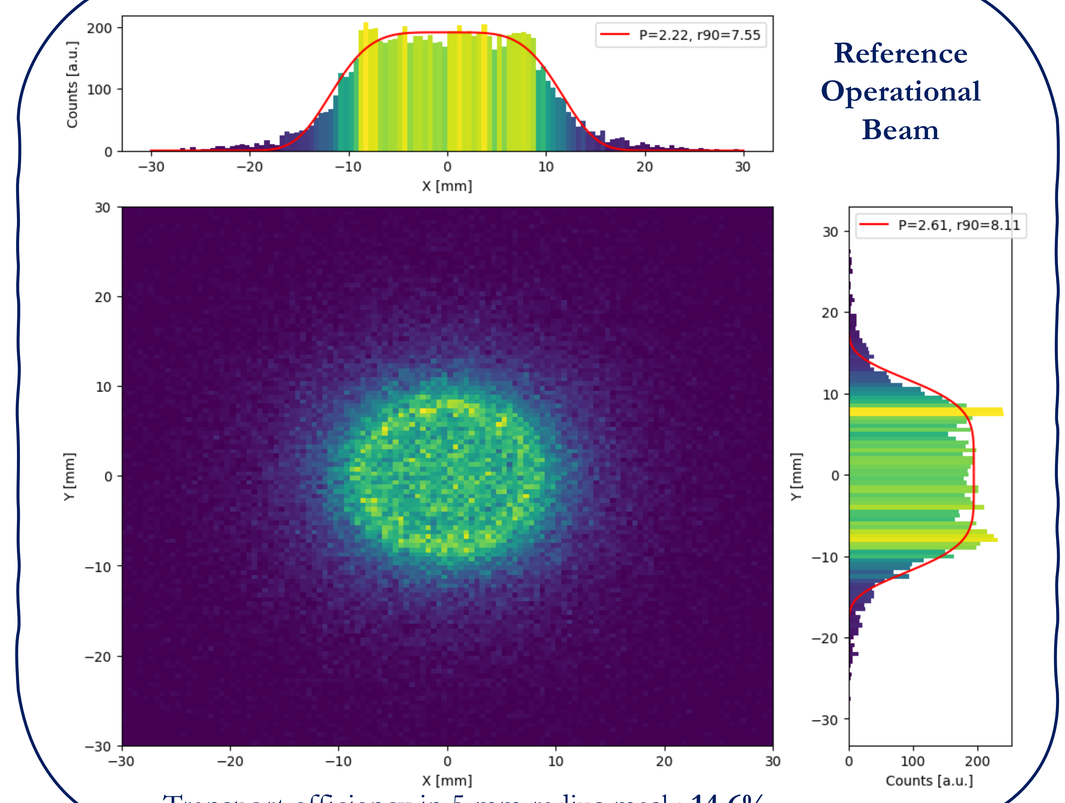

Reference Operational Beam

Transport efficiency in 5 mm radius mask: **14.6%**

Reinforcement Learning (RL) control for automated beam tuning:

- **TD3** (actor-critic, continuous control)
- **Incremental magnet tuning** (Δ currents)
- **Reward:** beam quality + constraints

Goal: uniform and stable beam profile

Workflow:

- Train and evaluate on RF-Track model
- Transfer policy to CLEAR accelerator

Key questions:

- Sim-to-real transfer
- Robustness to machine variations
- Stability across different beam conditions

Start in a **reduced 3-quadrupole space**
Scale to the full **9-quadrupole system**

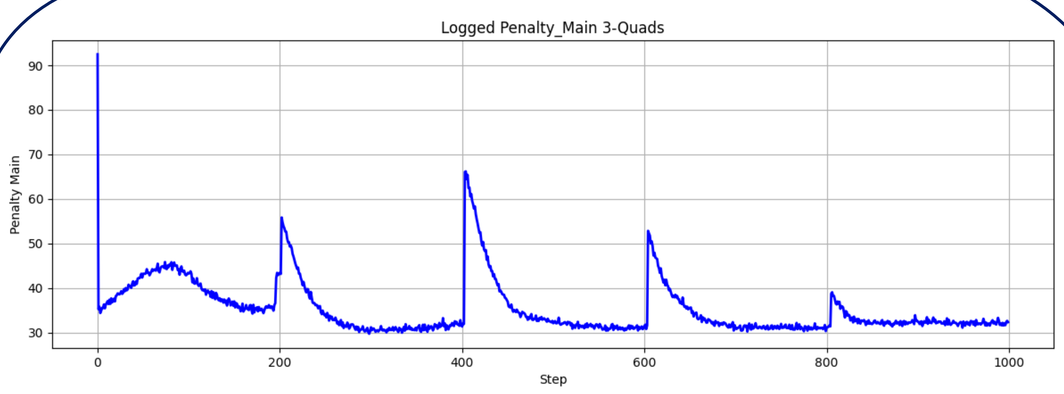

3-Quad RL-Optimized Beam

Transport efficiency in 5 mm radius mask: **25.6%**

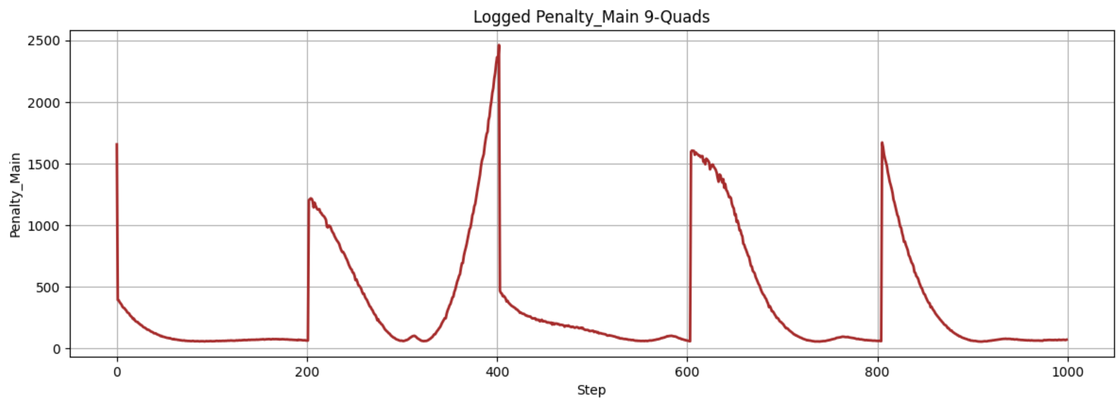

Simulation results:

- **3 quads:** stable learning, interpretable behavior, improvement in beam uniformity and transport
- **9 quads:** stable learning, higher complexity, in development

giacomo.tangari@cern.ch – giacomo.tangari@uniroma1.it