

Reinforcement Learning Beyond Greedy Optimisation for Delayed-Consequence Accelerator Control

Tuesday, 31 March 2026 12:00 (2 hours)

Most accelerator control systems assume that the effect of an action can be evaluated locally and immediately. While greedy approaches work in near-linear regimes and Bayesian Optimisation (BO) is now standard for black-box tuning, both are essentially static optimisers and struggle in dynamic tasks with delayed consequences, where even adaptive BO remains time-myopic and lacks explicit temporal credit assignment for system memory and long-range machine evolution.

We investigate three relevant forms of delayed consequences: explicit action latency (field settling delays response), magnetic hysteresis (output depends on change history), and ballistic amplification (small upstream kicks grow through nonlinear optics and apertures, causing downstream loss).

Using a high-fidelity XSuite model of the AWAKE electron line, we benchmark a reinforcement learning controller against an inverse-response greedy optimiser and BO. The learning-based method anticipates delayed effects and avoids failure regions where both baselines fail, indicating that delayed-consequence regimes are a key class of accelerator control problems where horizon-aware model-based or learning-based methods clearly outperform current practice.

Student

Yes

Primary authors: BJÖRKBOM, Kajsa Miho (PLUS University Salzburg); HIRLAENDER, Simon (PLUS University Salzburg)

Co-authors: TRAUSNER, Sarah (University of Salzburg); MIRONOVA, Olga; FISCHL, Lorenz (MedAustron GmbH, Wiener Neustadt, Austria); KAIN, Verena (CERN)

Presenters: BJÖRKBOM, Kajsa Miho (PLUS University Salzburg); HIRLAENDER, Simon (PLUS University Salzburg)

Session Classification: Poster session