

ROBUST REAL-TIME OPTIMIZATION OF SIS18 INJECTION USING GAUSSIAN PROCESS MPC*

S. Hirllaender, B. Halilovic, Paris Lodron University, Salzburg, Austria

S. Appel, P. Madysa, GSI Helmholtz Centre for Heavy Ion Research, Darmstadt, Germany

Abstract

We present advancements in the data-driven Model Predictive Control (MPC) framework for optimizing multi-turn injection (MTI) into the SIS18 synchrotron. Building on our prior work on safe, sample-efficient optimization, we systematically investigate the impact of current noise and transverse emittance fluctuations. By incorporating realistic error models derived from dedicated measurements of fluctuations on injected current and emittance into simulations, we demonstrate that the Gaussian Process model effectively filters random uncertainty, maintaining robust operation where standard numerical optimizers degrade. Furthermore, we report on the successful deployment of the framework during live SIS18 tuning.

INTRODUCTION

FAIR, the Facility for Antiproton and Ion Research, will provide antiproton and ion beams of unprecedented intensity and quality [6]. Multi-turn injection (MTI) into SIS18 is a key bottleneck for achieving the FAIR intensity goals. This is particularly important for intermediate charge-state ions, where it is critical to avoid loss-induced vacuum degradation, which is specifically triggered by MTI losses [7, 8].

Data-driven control and reinforcement learning (RL) offer strong potential for accelerator control problems beyond the scope of classical methods [1, 2]. However, RL in accelerator applications is challenged by limited data availability and the trade-off between training stability and data efficiency [3]. Model-based reinforcement learning (MBRL) addresses this issue by using predictive environment models to estimate future states and rewards. This reduces the required interactions with the accelerator [4, 5]. Gaussian processes (GPs) provide probabilistic models of the system dynamics and capture uncertainties essential for safe exploration.

GP-MPC [5] learns a dynamics model using n independent Gaussian processes (one per state dimension s_i) from observed transitions $(s_k, a_k) \rightarrow s_{k+1}$ collected in simulation or on the accelerator, where s_k denotes beam observables (e.g. losses) and a_k machine settings such as injection parameters. The controller operates around a pre-optimized injection point using incremental actions Δa to correct for drifts and further reduce losses. Each GP employs an ARD-RBF kernel with learned lengthscales that infer input relevance. At each step, the model propagates uncertainty over a finite horizon H via moment matching, and actions are selected by minimizing an uncertainty-aware cost using L-BFGS-B. The dataset is updated online after each interaction and GP hyperparameters are re-optimized, forming a closed-loop learning




Figure 1: SIS18 injection measurement: stored current (blue), incoming TK beam current (orange) and ideal current (green) during the optimization. The integral values measure, in addition to the septum loss and reward, the states. The transfer channel (TK) denotes the injection beamline.

scheme. A general introduction to reinforcement learning can be found in [13]. This approach has been successfully applied to the SIS18 injection system.

MULTI-TURN INJECTION

The MTI employs four bumper magnets to create a closed-orbit bump that allows successive beamlets to fill the horizontal phase space in accordance with Liouville's theorem. After injection, the beamlets undergo betatron oscillations with horizontal tune Q_x and pass the injection point again after each turn. Particle losses occur at the septum and at the machine acceptance and are quantified by the loss fraction η , defined as the ratio of lost to injected particles. Lower η indicates higher injection efficiency, with $\eta = 0$ corresponding to loss-free injection. A more detailed description of the MTI can be found in [9–11]. The validated SIS18 MTI model is implemented in Xsuite and benchmarked against experiments [10–12].

The effect of **injection quality** on the accumulated beam current is shown in Fig. 1. Ideal settings yield a monotonic increase, while mismatched injection leads to losses. The reward is the normalized integral deviation between actual and ideal beam currents, scaled by the maximum ideal current to reduce the impact of shot-to-shot fluctuations. In simulation, six injection parameters are varied: orbit bump amplitude x_c and angle x'_c , injected beam position x and angle x' , lattice mismatch M parameter, and per-turn bump reduction red. As shown in Fig. 2, TK pulses exhibit significant shot-to-shot fluctuations. The **ideal current** I_{id} accounts for these variations and is computed as the cumulative sum of measured

* simon.hirllaender@plus.ac.at




Figure 2: TK pulse fluctuations from nearly 150 measurements. Top: individual mean currents compared to the overall mean. Bottom: mean pulse profile μ of I_{TK} with standard deviation band σ . Pulse fluctuations depend on the ion species and production.

TK pulses per turn,

$$I_{id}[n] = \sum_{i=0}^n I_{TK}[i], \quad n \in \text{pulse turns}, \quad (1)$$

instead of relying on the simpler estimate obtained by multiplying the average TK current by the total number of injected turns.

In MTI optimization, a parameter adjustment (action) changes the measured current and affects subsequent optimization steps. Since decisions depend on the current beam conditions (states), MTI optimization can be formulated as a Markov Decision Process (MDP) as follows:

- **State:** $[\eta, \eta_S, \bar{I}_1, \bar{I}_2]$
- **Reward:** $-\frac{\int (I_{id}(t) - I_{SIS18}(t)) dt}{I_{id}^{max}}$
- **Action:** $[\Delta x_c, \Delta x'_c, \Delta x, \Delta x', \Delta M, \Delta red]$

The observation vector includes total loss η after 50 turns, septum loss η_S , and normalized time-integrated beam currents $\bar{I}_1, \bar{I}_2$. Although this 4D observation does not capture the full physical state (the particle distribution in phase space), the GP transition model conditions on the joint (s_k, a_k) pair, providing sufficient statistics for local prediction despite the partial observability. Episodes start from random perturbations and terminate upon success ($R > -R_{best}$), failure ($R < -R_{worse}$), or after 25 steps.

SIMULATION RESULTS

Extending our previous research on safe and sample-efficient optimization [14], we study how current noise and transverse emittance variations affect performance. Our goal




Figure 3: Combined 5 % particle and emittance noise with baseline reference. The GP-MPC agent remains robust under realistic compound perturbations. Training across 30 seeds with 95 % confidence intervals.

is to demonstrate robustness against stochastic uncertainties and thus bridge the sim-to-real gap.

Note that the initial state was intentionally set close to the global optimum, so the results characterize *local convergence* rather than global exploration. This design choice reflects the operational scenario where the machine is already coarsely tuned and the controller performs fine adjustment.

Baseline Training: Under ideal conditions (no noise), with constant emittance and an ideal rectangular TK pulse, the GP-MPC agent converges within approx. 300 interactions across 30 seeds with 95 % confidence intervals, confirming reliable convergence.

Robustness to Noise: The GP-MPC agent was evaluated across particle noise ($\sigma_p \in \{0, 1, 3, 5, 8\} \%$) and emittance noise ($\sigma_e \in \{0, 1, 3, 5, 8, 10\} \%$), totaling 180 runs across 6 seeds per condition. Figure 3 shows the training trajectory under compound 5 % perturbations, demonstrating that the GP’s inherent uncertainty quantification filters aleatoric noise without degrading convergence.

The full noise grid (Fig. 4) confirms robust optimization across all 30 conditions, with mean episode lengths remaining below 10 steps.

Measurement-Based Injection: To bridge the sim-to-real gap, the agent was tested using measured TK beam currents [15] instead of idealized rectangular pulses. The GP-MPC agent handles these realistic injection profiles without retraining (Fig. 5), confirming generalization beyond idealized conditions.

Observation Space Reduction: Reducing the observation vector from the 4D $[\eta, \eta_S, \bar{I}_1, \bar{I}_2]$ to the 2D $[\eta, \eta_S]$ increases the mean steps to converge from 11.4 ± 2.5 to 17.5 ± 3.2 per episode. With fewer observations, the GP transition model faces greater state aliasing—distinct physical configurations become indistinguishable in the reduced representation—requiring more interactions to resolve the input-output mapping. The ARD kernel automatically weights each observation dimension by its predictive relevance, making the richer 4D setting more data-efficient. In




Figure 4: Mean episode length across the 5×6 noise grid ($\sigma_p \in \{0, 1, 3, 5, 8\} \%$, $\sigma_e \in \{0, 1, 3, 5, 8, 10\} \%$, 180 runs total). Lower values indicate faster convergence. Cell annotations show average steps per episode; the GP model maintains convergence quality across all tested conditions.




Figure 5: Measurement-based TK injection with baseline reference. The agent generalizes to realistic, fluctuating beam profiles without retraining, confirming sim-to-real transfer capability.

practice, only the 2D observation ($\mathcal{I}_1$ and $\mathcal{I}_2$) is currently available from standard SIS18 beam instrumentation.

FIRST REAL-MACHINE RESULTS AT GSI

The GP-MPC method was implemented at the SIS18 synchrotron in July 2025 through the Geoff framework [16]. Previous successful optimization of SIS18 injection with Bayesian optimization and BOBYQA is reported in [15].

At the machine, only a 5D subset of injection parameters was accessible (bumper knob, fall time, chopper offset, two corrector kicks), compared to the 6D simulation parameter space. The mismatch parameter M has been removed, as it cannot be controlled directly with the available high-level knobs.

Figure 6 provides an overview of the **measured optimization**, showing normalized system states (top), actor outputs (middle), and RMS reward (bottom) over 100 steps. The controller systematically explored the five-dimensional parameter space, improving the RMS reward from -0.32 to -0.28 . The GP-MPC controller required fewer than 20




Figure 6: GP-MPC optimization trajectory at GSI (measurement): states (top), actor outputs (middle), and RMS reward (bottom) over 100 optimization steps.

optimization steps to achieve a setting comparable to the manually tuned operational state of 50% loss and fewer than 60 steps to surpass it, reducing to 38% losses after 90 steps, confirming the viability of the approach for online deployment.

SUMMARY AND OUTLOOK

This paper demonstrates GP-MPC for multi-turn injection optimization at SIS18. The approach converges within ~ 300 interactions, remains robust across 30 noise conditions (particle and emittance noise up to 8% and 10%, respectively), and generalizes to realistic measurement-based TK beam currents. The richer 4D observation yields faster convergence (11.4 ± 2.5 steps/episode) than the reduced 2D case (17.5 ± 3.2 steps/episode), attributed to reduced state aliasing in the GP transition model. The first real-machine deployment at GSI confirms viability for online operation. Future work includes systematic comparison with BOBYQA, sim-to-real transfer via the Xsuite model, integration of simulation-based GP priors, and meta-RL for rapid adaptation.

ACKNOWLEDGMENTS

S. Hirlander gratefully acknowledges the support of the WISS 2025 project ‘IDA-lab Salzburg’ (20204-WISS/225/197-2019 and 20102-F1901166-KZP).

S. Appel and P. Madysa acknowledge the EURO-LABS and iRIS projects funded by the European Union’s Horizon Europe Research and Innovation program under Grant Agreement no. 101057511 and 101275935.

REFERENCES

- [1] S. Hirlander, L. Lamminger, G. Zevi Della Porta, and V. Kain, “Ultra fast reinforcement learning demonstrated at CERN AWAKE”, in *Proc. IPAC’23*, Venice, Italy, May 2023, pp. 4510–4513. doi:10.18429/JACoW-IPAC2023-THPL038
- [2] S. Hirlander and N. Bruchon, “Model-free and Bayesian ensembling model-based deep reinforcement learning for particle accelerator control demonstrated on the FERMI FEL”, Jan. 2022. doi:10.48550/arXiv.2012.09737
- [3] V. Kain *et al.*, “Test of machine learning at the CERN LINAC4”, in *Proc. HB’21*, Batavia, IL, USA, Oct. 2021, pp. 181–185. doi:10.18429/JACoW-HB2021-TUEC4
- [4] S. Hirlander *et al.*, “Towards few-shot reinforcement learning in particle accelerator control”, in *Proc. IPAC’24*, Nashville, TN, USA, May 2024, pp. 1804–1807. doi:10.18429/JACoW-IPAC2024-TUPS60
- [5] S. Kamthe and M. P. Deisenroth, “Data-efficient reinforcement learning with probabilistic model predictive control”, Feb. 2018. doi:10.48550/arXiv.1706.06491
- [6] P. Spiller *et al.*, “FAIR technical design report”, 2008.
- [7] E. Mustafin *et al.*, “A theory of the beam loss-induced vacuum instability applied to the heavy-ion synchrotron SIS18”, *Nucl. Instrum. Methods Phys. Res. A*, vol. 510, no. 3, pp. 199–205, Sep. 2003. doi:10.1016/S0168-9002(03)01811-4
- [8] C. Omet, P. Spiller, J. Stadlmann, and D. H. H. Hoffmann, “Charge change-induced beam losses under dynamic vacuum conditions in ring accelerators”, *New J. Phys.*, vol. 8, no. 11, p. 284, Nov. 2006. doi:10.1088/1367-2630/8/11/284
- [9] S. Appel, O. Boine-Frankenheim, and F. Petrov, “Injection optimization in a heavy-ion synchrotron using genetic algorithms”, *Nucl. Instrum. Methods Phys. Res. A*, vol. 852, pp. 73–79, Apr. 2017. doi:10.1016/j.nima.2016.11.069
- [10] S. Appel and O. Boine-Frankenheim, “Multi-turn injection into a heavy-ion synchrotron in the presence of space charge”, Mar. 2014. doi:10.48550/arXiv.1403.5972
- [11] S. Appel *et al.*, “Injection optimization through generation of flat ion beams”, *Nucl. Instrum. Methods Phys. Res. A*, vol. 866, pp. 36–39, Sep. 2017. doi:10.1016/j.nima.2017.05.041
- [12] G. Iadarola *et al.*, “Xsuite: An integrated beam physics simulation framework”, in *Proc. HB’23*, Geneva, Switzerland, Oct. 2023, pp. 73–80. doi:10.18429/JACoW-HB2023-TUA2I1
- [13] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, MIT Press, Cambridge, MA, USA, 1998.
- [14] S. Appel, S. Hirlander, and N. Madysa, “Data-driven model predictive control for automated optimization of injection into the SIS18 synchrotron”, in *Proc. IPAC’24*, Nashville, TN, USA, May 2024, pp. 1800–1803. doi:10.18429/JACoW-IPAC2024-TUPS59
- [15] S. Appel *et al.*, “Automated optimization of accelerator settings at GSI”, in *Proc. IPAC’24*, Nashville, TN, USA, May 2024, pp. 882–885. doi:10.18429/JACoW-IPAC2024-MOPS68
- [16] P. Madysa, S. Appel, V. Kain, and M. Schenk, “Geoff: The generic optimization framework & frontend for particle accelerator controls”, *SoftwareX*, vol. 32, p. 102335, Dec. 2025. doi:10.1016/j.softx.2025.102335