

DATA-DRIVEN OPTIMIZATION OF OPEN-LOOP CONTROL FUNCTIONS AT THE CERN SUPER PROTON SYNCHROTRON

A. Menor de Oñate*, M. Schenk, V. Kain, T. Prebibaj, I. Mases Sole, B. Rodriguez Mateos
CERN, Geneva, Switzerland

Abstract

To minimize beam intensity loss during a cycle in the CERN Super Proton Synchrotron (SPS), several machine parameters must be adjusted as functions of cycle time, spanning injection, injection plateau, acceleration, and extraction plateau. Today, these functions are typically tuned manually – a cumbersome procedure that can require hours of operator effort. This paper presents the progress towards automatically tuning time-dependent parameter functions. Using Bayesian optimization (BO), we aim to minimize intensity loss throughout the cycle with intensity measurements as the primary feedback signal. We report results from applying this method to an intentionally detuned machine development beam in the SPS, as a step towards deployment on the operational fixed-target beams. The approach is generic and applicable to time-dependent parameter optimization problems in other machines.

INTRODUCTION

The Super Proton Synchrotron (SPS) provides beams to a range of users, including the LHC and several fixed-target and beam-physics facilities. Each beam is delivered through a “cycle”, in which the beam is first injected from the Proton Synchrotron (PS) – possibly over several injections – onto the SPS injection plateau (“flat bottom”). The beam is then accelerated to its target energy and extracted to its destination. An example cycle for a slow-extracted beam is shown in Figure 1, displaying the total beam intensity and momentum alongside three parameter functions targeted by the optimization: the horizontal (Q_H) and vertical (Q_V) betatron tunes and the total RF voltage (V_{RF}).

Throughout the cycle, the beam is subject to a variety of processes and beam-dynamics phenomena: injection, RF capture, intensity-dependent tune shifts that increase stepwise with each injection, decaying eddy currents compensated by adjustments of the momentum, tune, and chromaticity functions on the flat bottom, and transition crossing during acceleration. This requires multiple machine parameters to be adjusted as functions of cycle time to minimize intensity loss. These parameter functions must be defined before the cycle is executed and are applied open-loop: they cannot be modified during a cycle, only between cycles.

Moreover, the optimal settings depend on the requested beam intensity, which is adjusted several times per year as experimental needs vary, making manual retuning a recurring operational burden. This is particularly relevant for the future Search for Hidden Particles (SHiP) [1] experiment at the SPS, which will ultimately be a high-intensity operational




Figure 1: Example SPS cycle for a slow-extracted beam, showing two injections on the injection plateau, acceleration from 14 GeV/c to 400 GeV/c and extraction (top panel). Horizontal (Q_H) and vertical (Q_V) betatron tunes and RF voltage (V_{RF}) functions are also shown (bottom panels).

beam delivered at high repetition rate, making the minimization of transmission losses crucial to keep activation of SPS beam equipment at acceptable levels.

Among the parameter functions, the betatron tunes and RF voltage have the largest impact on intensity loss (or equivalently, beam transmission) for the fixed-target beams under study here and are the primary focus of this work. Other parameter functions, such as the chromaticities, will be addressed in future studies. From an operations perspective, optimizing these parameters manually is time-consuming and becomes increasingly difficult at high intensities, where key observables such as tune measurements are unreliable or unavailable, leaving intensity measurements as the primary feedback. This further motivates the use of automated, data-driven optimization, an approach recently applied to related problems at CERN, including Multi-Turn Extraction tuning at the PS [2] and hysteresis compensation in the SPS

* adrian.menor.de.onate@cern.ch

main magnets [3], as well as at other particle accelerator laboratories [4]. This work focuses on the data-driven optimization of the open-loop parameter functions governing beam transmission within an SPS cycle.

OPTIMIZATION PROBLEM AND METHOD

Automatic tuning of parameter functions presents a functional optimization problem: the decision variables are functions of cycle time, and the goal is to find the function shapes that minimize intensity loss throughout. Formally, we define the transmission T as the objective functional of the problem, given by the ratio of intensities I measured at two times $t_0 < t_f$ in the cycle (Equation 1). The dependence of T on the parameter functions enters implicitly through the beam dynamics. The pair $[t_0, t_f]$ specifies an optimization window which can be any segment of the cycle and the parameter functions are correspondingly varied only within this window. The optimization problem is then given by $\max_{Q_H(t), Q_V(t), V_{RF}(t)} T[Q_H, Q_V, V_{RF}]$, with

$$T[Q_H(t), Q_V(t), V_{RF}(t)] = \frac{I(t_f)}{I(t_0)} \quad (1)$$

Sample efficiency is critical, since machine time is scarce. Two strategies reduce the high-dimensional problem to a more tractable one. First, $Q_H(t)$, $Q_V(t)$, and $V_{RF}(t)$ are parametrized by a smaller number of values: the tune functions are defined by skeleton points interpolated linearly or by splines. Splines provide two distinct benefits: they are C^2 -smooth, reducing stress on the main quadrupole power converters, and offer greater representational power than piecewise-linear interpolation for the same number of knots. The RF voltage function is divided into domains to which constant shifts are applied. These choices were guided by input from machine operators and specialists. Second, since the full cycle is too high-dimensional to optimize at once, it is segmented into windows that are optimized sequentially. Figure 2 illustrates the actions on each of the parameters.

All parametrized values are concatenated into a single action vector, which is modified with Bayesian optimization (BO), and then mapped back to the parameter functions before applying them to the machine. A key advantage of this joint formulation is that BO can act on all parameter functions and skeleton points simultaneously across the optimization window. Manual tuning, by contrast, typically adjusts one parameter and one skeleton point at a time, and such coordinate-wise changes can miss optima where parameters interact.

For the machine tests carried out in this work, the tune skeleton points are constrained to $\Delta Q_{H,V}^{\max} = \pm 0.01$ with respect to the values present in the machine at the start of the optimization, and the RF voltage is bounded between 0 and 7.5 MV. Depending on the parametrization, the action vector has up to 28 dimensions in the tests reported here. The optimization uses BO with a Gaussian process (GP) surrogate model, a radial basis function (RBF) kernel with




Figure 2: Example actions on the tune (Q_H , Q_V) and RF voltage (V_{RF}) functions over the course of an optimization. The tune functions are parametrized by splines, while the RF voltage is divided into domains to which constant shifts are applied.

automatic relevance determination (ARD) lengthscales, and Log Expected Improvement as the acquisition function [5]. The GP hyperparameters are re-optimized at each iteration. The optimization is warm-started with 5 to 7 random actions before BO begins proposing candidate actions.

RESULTS

To validate the optimization logic, several tests were carried out using the cycle foreseen for the SHiP experiment at the SPS. This cycle was chosen due to its similar acceleration profile to existing fixed-target cycles in the SPS (which require frequent retuning to accommodate changes in beam intensity), but also because of the aforementioned relevance of minimizing transmission losses in SHiP beams; ultimately designed to be a high-intensity operational beam delivered at high repetition rate. The tests were limited to a single injection and a total intensity of 3.3×10^{12} p, safe for the machine in case of severe intensity losses. Note that the cycles already had acceptable transmissions $T \geq 0.97$ and were hence purposely detuned by perturbing the tune and RF voltage functions away from their nominal values to create degraded starting conditions. This both provided the optimizer with measurable headroom and mimicked the state of an uncommissioned cycle, where automatic tuning would be most useful in operation.

To illustrate the optimizer's performance, Figure 3 shows an example run on an initially degraded SHiP cycle. The cycle was divided into two sequential optimization windows: Q_H , Q_V , and V_{RF} are first optimized simultaneously over the first window, after which the same procedure was applied to the second, with the optimizer constrained to the maximum iteration budget allocated for the tests. The two optimiza-




Figure 3: Example optimization of a purposely detuned SHiP cycle. The cycle was divided into two windows that were sequentially optimized in terms of Q_H , Q_V , and V_{RF} .

tion windows overlapped by 200 ms to capture time-delayed effects: parameter configurations applied in the first window may cause beam losses that only materialize later in the cycle. The overlap also prevents abrupt transitions at the interface between adjacent windows. We chose this overlap empirically; its validity needs to be further investigated for different configurations of parameter functions and cycle regions. In this test, BO sequentially improved the transmission to ≥ 0.98 with iteration budgets of 20 and 40 cycles for the first and second windows, respectively.

Similar machine tests were carried out across the 2025 and 2026 runs, with results summarized in Figure 4. The transmission of the SHiP cycle was improved in 45 of 54 cases, where failures to improve primarily occurred in cycles already having a high initial transmission ≥ 0.97 . All the tests were limited to the flat-bottom and early acceleration phases with a single injection, in line with the safety constraints. These results motivate scaling to higher intensities and to optimization across the full cycle.

CONCLUSIONS

This paper presents progress towards automating the optimization of open-loop machine parameter functions in the




Figure 4: Comparison of the initial transmission with the best transmission found, over 54 tests on initially detuned SHiP cycles. The tests span a range of parameter detunings and cycle time windows, and therefore have different head-room for improvement.

CERN SPS, targeting the horizontal and vertical tunes as well as the RF voltage as a first step. The approach relies on BO with a GP surrogate model to maximize beam transmission using only intensity measurements as feedback, and acting in the parameter space defined by the chosen function parametrizations (linear or spline interpolations, or constant shifts).

The method was validated on purposely degraded SHiP cycles at intensities of 3.3×10^{12} p, where the optimizer recovered transmission by acting on sequential windows along the cycle, varying the tunes and the RF voltage.

Several limitations remain. The current tests are restricted to a single injection at moderate intensity, and the artificial degradation may not capture the full complexity of an actual commissioning scenario, where parameter functions may be poorly known across many degrees of freedom simultaneously. The dimensionality of the optimization is managed through windowed sequential optimization and function parametrization, whose scalability to cycles with multiple injections, higher intensities, and larger parameter families has yet to be demonstrated.

Next steps will focus on incorporating additional domain knowledge, improving beam-state observations, extending the optimization to larger families of parameter functions (including chromaticity), and validating the system at higher beam intensities.

REFERENCES

- [1] S. Alekhin *et al.*, “A facility to search for hidden particles at the CERN SPS: the SHiP physics case”, *Rep. Prog. Phys.*, vol. 79, p. 124201, 2016.
[doi:10.1088/0034-4885/79/12/124201](https://doi.org/10.1088/0034-4885/79/12/124201)
- [2] L. Foldesi *et al.*, “Model-based optimisation for automated multi-turn extraction tuning at the CERN Proton Synchrotron”, in *Proc. IPAC’25*, Taipei, Taiwan, Jun. 2025, pp. 2703–2706.
[doi:10.18429/JACoW-IPAC2025-THPM009](https://doi.org/10.18429/JACoW-IPAC2025-THPM009)
- [3] A. Lu *et al.*, “Data-driven hysteresis compensation in the CERN SPS main magnets”, in *Proc. IPAC’25*, Taipei, Taiwan, Jun. 2025, pp. 1674–1677.
[doi:10.18429/JACoW-IPAC2025-WEAN2](https://doi.org/10.18429/JACoW-IPAC2025-WEAN2)
- [4] R. Roussel *et al.*, “Bayesian optimization algorithms for accelerator physics”, *Phys. Rev. Accel. Beams*, vol. 27, p. 084801, Aug. 2024.
[doi:10.1103/PhysRevAccelBeams.27.084801](https://doi.org/10.1103/PhysRevAccelBeams.27.084801)
- [5] S. Ament *et al.*, “Unexpected Improvements to Expected Improvement for Bayesian Optimization”, arXiv:2310.20708, 2025. [doi:10.48550/arXiv.2310.20708](https://doi.org/10.48550/arXiv.2310.20708)