TOWARDS FULLY AUTOMATED TRANSFER LINE COMMISSIONING AT THE CERN SUPER PROTON SYNCHROTRON

A. Menor de Oñate*, M. Schenk, V. Kain, B. Rodriguez Mateos, G. Dal Maso, N. Charitonidis
CERN, Geneva, Switzerland

Abstract

Beam commissioning of slow extracted beams from the CERN Super Proton Synchrotron (SPS) to the North Area experimental targets requires trajectory control through multiple transfer lines using corrector magnets—a process that traditionally demands significant expert intervention. Previous work demonstrated the feasibility of applying reinforcement learning (RL) for automated trajectory correction based on secondary emission monitor (SEM) split-foil intensity measurements, successfully centering the beam on target under nominal conditions [1]. However, this approach fails when the beam is lost or its position exceeds the SEM's active surface, and when the corrector magnets' polarities are not known; common sources of uncertainty during commissioning.

We present an extended multi-stage optimization scheme that addresses these critical limitations by automating beam threading when the trajectory exceeds the SEMs' acceptance, systematically identifying corrector magnet polarity configurations, and optimizing the impact angle to maximize beam intensity at the fixed-target stations, measured by scintillators arranged around the target. The threading algorithm employs quasi-random search combined with Bayesian optimization (BO) to center the beam in the SEMs, before handing over to the RL controller. The automated polarity determination uses online system identification to resolve sign ambiguities in the correctors, eliminating a common source of commissioning delays when using RL or other dedicated steering algorithms. Finally, BO is used to optimize the position of the movable SEM monitors at the targets' locations, maximizing target intensity.

INTRODUCTION

The CERN Super Proton Synchrotron (SPS) sends DC proton and heavy ion beams to the North Area (NA) experimental targets through a number of transfer lines. During commissioning, the beams need to be steered through the lines, and the yield of the NA targets needs to be maximized. The steering actions are done using dipole magnets (correctors), while the beam trajectory along the line is monitored using secondary emission monitors (SEMs), as conventional beam position monitors are not available in the line. SEMs have two foils (either aligned horizontally or vertically) that each measure beam intensity [2]. If the beam hits both foils, position can be inferred by using the intensity difference between foils and by fitting a Gaussian beam profile assuming known emittance and nominal optics. However, when the beam falls on one foil only, the position cannot be inferred. In this scenario, classical steering algorithms that rely on beam position measurements fail. Previous work demonstrated that reinforcement learning (RL) controllers can steer the line in such cases [1].

While this is a step forward, if the beam is completely off the foils, or the magnet polarities are changed from the previous commissioning, the RL agent described in [1] can no longer steer the line, as it relies on valid SEM readings and known corrector polarities. Therefore, an additional automated procedure is needed to optimize the beam intensity at the targets.

With this in mind, the following sequential automated procedures have been tested during commissioning in the horizontal plane of the line: first, steering from severe beam trajectory offsets using Bayesian optimization (BO); second, estimating corrector polarities with kick response measurements that are compared against a model of the transfer line; third, steering the beam to a high symmetry S (done with RL [1] or MICADO [3]) in the SEMs (see Equation 1, where I_1 and I_2 are the intensities on foil 1 and 2 in the SEM), and finally optimizing the intensity at the fixed targets.

$$S = \sqrt{1 - \frac{|I_1 - I_2|}{I_1 + I_2}} \tag{1}$$

STEERING FROM SEVERE BEAM TRAJECTORY OFFSETS (THREADING)

At the beginning of commissioning, the beam may be lost along the transfer line, or its position may fall outside the active surface of the SEMs. In order to bring the beam back to an acceptable symmetry (here targeted at 0.6), the transfer line is divided into segments mapping one or more corrector magnets to the downstream SEM(s), located before the next corrector. Each of these segments is sequentially steered by first taking quasi-random actions that explore the action space of the magnet, generating space-filling samples to maximize exploration. Once the beam is detected in any of the SEMs in the segment, BO guides the steering process. This sequential optimization is referred to as "threading".

When inspecting the optimization landscape (in simulation) of a given segment of the transfer line that needs to be optimized (see left plot in Figure 1, showing the action spaces of two corrector magnets, and the resulting objective value), we observe that a ridge-like solution is present when using two magnets or more, corresponding to the magnet combinations that result in a good objective value:

$$\text{Objective} = \sum_{j=1}^{N_{\text{SEM}}} \begin{cases} S_j^2 & \text{if } I_j^{\text{tot}} \geq I_{\text{min}}, \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

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

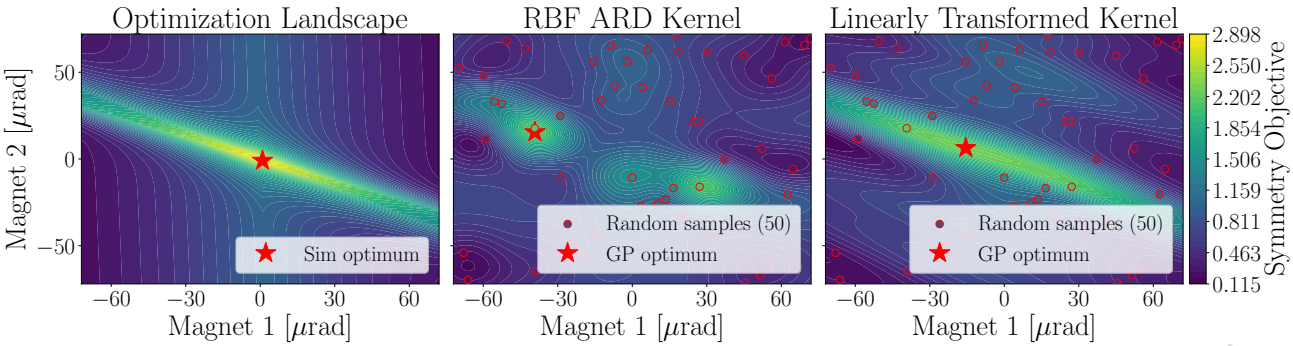

Figure 1: Simulated optimization landscape of a transfer line segment consisting of two corrector magnets (left). The standard RBF ARD kernel (middle) and the linearly transformed input space kernel (right) model this landscape using a GP, with the latter capturing the ridge structure more accurately.

Here, S_j is the symmetry (as defined in Equation 1) , I_j^{tot} the total intensity ($I_{1,j} + I_{2,j}$), and I_{min} the beam-detection threshold in intensity in a given SEM. The width of this ridge is proportional to the beam size, and the sign of the slope also changes with the polarities of the magnets.

Modeling this problem using a Gaussian Process (GP) (to be used as surrogate model for BO) with a standard kernel with axis-aligned lengthscales (ARD) fails to capture the ridge structure, as such kernels can only represent anisotropy along the coordinate axes (see mid plot in Figure 1). To address this, we introduce a lower triangular matrix $\mathbf{L}$ whose entries (a, b, c) are free parameters optimized jointly with the GP hyperparameters via marginal likelihood maximization. This matrix is of size **num magnets** × **num magnets** (here exemplified when optimizing two magnets), and represents a rotation and scaling of the input space, effectively making the ridge axis-aligned in the transformed coordinate system and allowing the isotropic base kernel k_{base} to learn a large lengthscale along the ridge and a short one across it (the lengthscale governs how rapidly the kernel correlation decays with distance in each input direction, so a large lengthscale implies the function varies slowly along that direction, while a short one implies rapid variation). The transformation is applied directly to the kernel inputs as shown in Equation 3.

$$k(\mathbf{x}, \mathbf{x}') = k_{\text{base}}\left(\mathbf{L}^\top \mathbf{x},\ \mathbf{L}^\top \mathbf{x}'\right), \quad \mathbf{L} = \begin{pmatrix} a & 0 \\ b & c \end{pmatrix} \tag{3}$$

Using this kernel, the GP models the optimization landscape, as shown in the right plot of Figure 1. We benchmarked this kernel against a standard RBF (radial basis function) ARD kernel in the BO loop, using Log Expected Improvement [4] as the acquisition function in both cases. As shown in Figure 2, the transformed kernel leads to a more sample-efficient optimization of the corrector magnets' settings, and was used for online optimization in the SPS.

Figure 3 shows the threading results using this algorithmic logic in the real machine, in a large section of the transfer line referred to as TT20, consisting of 8 SEM devices and 7 corrector magnets. The transfer line was decomposed into different segments that were steered from a no-beam-reading

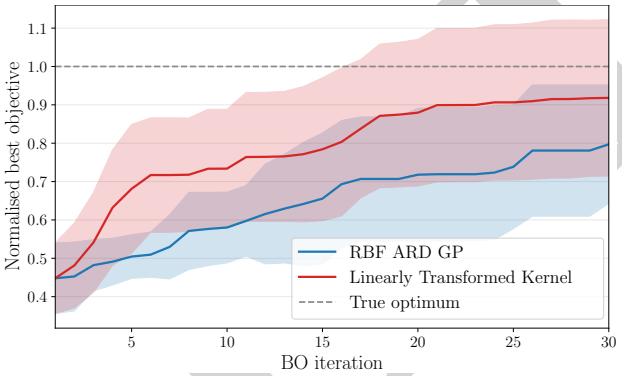

Figure 2: Bayesian optimization of a set of correctors-SEMs in simulation. The linearly transformed kernel in Equation 3 leads to a more sample-efficient optimization. Statistical results are obtained from an ensemble of 10 runs.

status in the SEMs to a symmetry ≥ 0.6, effectively paving the way for high-symmetry steering using dedicated algorithms. The beam intensity was 3×10^{11} p; a low intensity value that is safe for the machine, but which results in noisy readings at the SEMs, representing the most challenging optimization scenario, and at the same time the most useful one from an operation's perspective.

POLARITY ESTIMATION

Once the beam is detected in all SEMs using the logic described above, a high-symmetry steering of the line can be performed using algorithms such as RL steering [1], or MICADO [3] as implemented in e.g. YASP [5]. However, all these algorithms rely on the assumption that the corrector polarities are known, which may not always hold in practice, as incorrect polarities can occasionally occur due to human error during hardware modifications. With this in mind, a system identification procedure identifies the polarities before triggering the aforementioned algorithms.

Specifically, the procedure determines the sign of the corrector magnet's polarity $p_j \in \{+1, -1\}$ relating each corrector j to the accelerator optics model by sequential perturbation testing. For a given corrector, perturbation actions of alternating sign $\delta k_j^{(n)}$ (where n is the iteration index) are applied, and the expected response $e_{n,i} = R_{ij}\, \delta k_j^{(n)}$

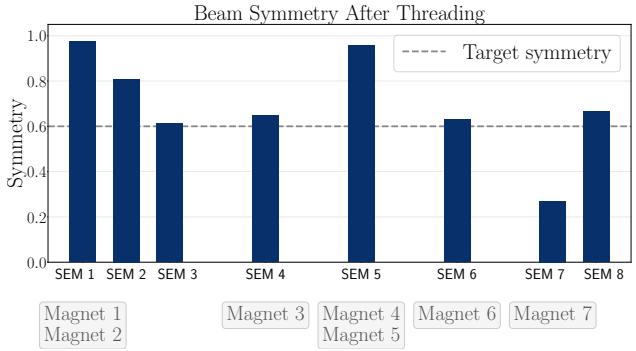

Figure 3: Beam threading of the TT20 line segment connecting the SPS to the North Area experimental hall. Corrector magnets used for each SEM segment are shown in gray boxes. The beam was initially undetected in any SEM; optimization brought it into all SEMs and steered it to a symmetry above 0.6. The below-threshold symmetry of SEM 7 is due to the optics characteristics in that line segment.

predicted by the optics response matrix R (linking corrector j to SEM i) is compared against the measured response $o_{n,i} = \Delta I^{(i)}_{\text{perturbed},n} - \Delta I^{(i)}_{\text{base}}$. All $(e_{n,i}, o_{n,i})$ pairs across SEMs and iterations are used to update the running Pearson correlation coefficient r, which is then converted to the test statistic $t = r\sqrt{(N-2)/(1-r^2)}$ and evaluated against the Student's t-distribution to yield a confidence c. The iteration terminates when $c \geq 0.95$, and the polarity is accepted as $p_j = \text{sign}(r)$, encoding whether the physical corrector response agrees with $(+1)$ or is inverted (-1) relative to the model prediction.

This procedure was tested over two trials in the TT20 segment to identify the horizontal correctors, correctly identifying all 7 out of 7 magnets and 6 out of 7 magnets, respectively.

OPTIMIZING BEAM INTENSITY ON THE TARGETS

After steering, the beam's impact angle on the 500 mm Beryllium target in the transfer line must be optimized, as it determines the intensity delivered to the two downstream beam lines it services, H2 and H4 [6], as measured by scintillators arranged in these beam lines. These scintillators were specifically installed for the purpose of target alignment, and are not a permanent fixture of the beam line. Since the previous commissioning, the targets' position may have drifted or been modified compared to the previous installation, resulting in misalignment and hence a low intensity reading at the scintillators. To adjust the impact angle, movable SEM monitors at the targets are displaced transversely to candidate positions. For each position, an inner proportional controller (integrated in YASP) steers the upstream transfer line until the SEM symmetries are ≥ 0.9, ensuring the beam passes through the displaced monitor centrally and thus hits the target at the corresponding angle. Once the inner loop converges, the scintillator intensities are read as a measure of the resulting beam-target interaction quality. An outer-loop optimizer then proposes the next SEM position based on these scintillator readings, and the process repeats.

We optimize the outer loop using BO (where conventionally, a grid scan method is used, significantly hindering the sample efficiency of the optimization) to select the new SEM positions. Figure 4 shows the optimization results for the H2 and H4 beam lines, improving the intensity in both tests.

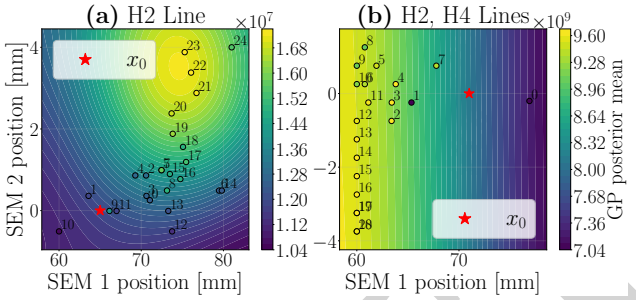

Figure 4: Beam scan measurement for scintillator optimization at the H2 line only (a) or both H2 and H4 lines simultaneously (b). Axes show the movable SEM positions (mm), while the objective value (proportional to scintillator intensity counts) is color-coded. Data are fitted using a GP with RBF kernel and ARD. x_0 is the initial objective value.

CONCLUSION

When commissioning the SPS-North Area transfer lines, the beam needs to be steered from an initial status where it might even be lost along the line, to a steered configuration that maximizes the intensity at the targets. This is achieved by first recovering beam signal on the SEMs that are distributed along the transfer line, then steering the line to achieve a high mean symmetry, and finally finding the optimum movable SEM positions at the targets.

This paper shows preliminary results to automate all these commissioning steps. A BO-based steering algorithm, combined with quasi-random search, sequentially recovers the beam from a no-signal state to a symmetry of ≥ 0.6 across all SEMs. Then, the line is steered to high symmetry values using specialized steering algorithms that rely on the knowledge of the corrector magnets' polarities, which are estimated by performing a kick response analysis comparing the machine readings against the expectations from optics models. Finally, the intensity at the targets is maximized by using a nested control logic that optimizes the movable SEM positions and steers the line accordingly.

All these automation routines have been experimentally validated in the SPS, representing a first proof of concept. Future efforts will focus on simultaneous commissioning in both planes, more accurate polarity estimation, and robust criteria to discriminate beam presence from SEM noise.

REFERENCES

[1] A. Menor De Oñate *et al.*, "Trajectory steering for DC beams at the CERN SPS using reinforcement learning based on intensity measurements", in *Proc. IPAC'25*, Taipei, Taiwan, Jun. 2025,

pp. 2928–2930.
doi:10.18429/JACoW-IPAC2025-THPM113

[2] K. Bernier *et al.*, "Calibration of secondary emission monitors of absolute proton beam intensity in the CERN SPS North Area", CERN, Geneva, Switzerland, CERN-97-07, 1997.
doi:10.5170/CERN-1997-007

[3] B. Autin and Y. Marti, "Closed Orbit Correction of A.G. Machines Using a Small Number of Magnets," CERN, Geneva, Switzerland, Rep. CERN-ISR-MA/73-17, 1973.

[4] S. Ament *et al.*, "Unexpected Improvements to Expected Improvement for Bayesian Optimization", Jan. 2025.
doi:10.48550/arXiv.2310.20708

[5] J. Wenninger, "YASP: Yet Another Steering Program User Guide", CERN, Geneva, Switzerland, 2005.

[6] L. Gatignon, "Design and Tuning of Secondary Beamlines in the CERN North and East Areas," CERN, Geneva, Switzerland, Rep. CERN-ACC-NOTE-2020-0043, 2020.