DATA-DRIVEN HYSTERESIS-COMPENSATION IN THE CERN SPS MAIN MAGNETS

A. Lu*[1], Technische Universität Wien, Vienna, Austria
V. Kain, C. Petrone, V. di Capua, M. Schenk, CERN, Geneva, Switzerland
M. Taupadel[1], Technische Universität Darmstadt, Darmstadt, Germany
[1]also at CERN, Geneva, Switzerland

Abstract

Magnetic hysteresis and eddy current decay continue to challenge beam quality and operational consistency in multicycling machines like the Super Proton Synchrotron (SPS) at CERN. Building on our previous work, this paper presents improvements in the data-driven approach for magnetic field modelling to enhance the reproducibility of SPS dipole and quadrupole fields and thus maintain stable beam parameters across all operational cycles. The method is based on feedforward correction using magnetic field forecasting with machine learning. It now includes additional operational experience and demonstrates that the field error compensation can reliably be used in operation. This contribution proves that hysteresis compensation can be achieved without a feedback system based on expensive installations with online field measurements in reference magnets. The performance improvements achieved by eliminating the need for manual adjustments and reducing time- and energy-consuming accelerator pre-cycles are presented. The paper also sets the stage for future applications in higher-order magnets, like sextupoles and octupoles, as well as on other CERN synchrotrons.

INTRODUCTION

Building on earlier tests of a machine learning-based feedforward compensation method for hysteresis and eddy current effects in the SPS main dipoles [1], this contribution presents extended studies under realistic operational conditions. The approach was further developed to improve robustness and was evaluated across representative cycle sequences at the CERN Super Proton Synchrotron (SPS).

The focus remains on the SPS Fixed Target Proton (SFTPRO) cycle, which injects at the lowest energy (14 GeV) and is currently the only cycle relying on a dedicated pre-cycle, labeled MD1 in Figure 1, to mitigate hysteresis. Its low beam rigidity compared to LHC-type cycles (26 GeV injection energy) makes it more sensitive to magnetic field variations. In addition, fixed-target beams use chromatic resonant slow extraction [2, 3], where field stability at flat-top is critical. Omitting the pre-cycle introduces hysteresis-induced field deviations in the dipoles corresponding to approximately $\Delta p/p \approx 1 \times 10^{-3}$ at injection. The acceptable tolerance for reproducible beam conditions is $\Delta p/p \approx 1\text{–}2 \times 10^{-4}$ (corresponding to $\approx 2 \times 10^{-5}$ T) at injection, while at extraction it is significantly more relaxed (1×10^{-4} T).

Dynamic effects from eddy current decay, while visible in magnetic field measurements, do not always manifest on the beam as previously assumed. This suggests the need for beam-based modeling methods, which are discussed later.

This paper presents the latest results from offline and online testing of the compensation algorithm, discusses integration aspects, and outlines the path towards future deployment in routine operation, as well as extension to higher-order magnets and other machines in the CERN complex.

IMPROVED HYSTERESIS MODELING AND COMPENSATION

Model Pre-training and Fine-tuning

To improve the hysteresis model, we pre-train a Temporal Fusion Transformer (TFT) [4] on pseudo-random, pseudo-realistic unipolar waveforms generated with a simple Jiles-Atherton model [5], and then fine-tune it on measured SPS data. Simulation parameters are chosen to produce hysteresis curves that resemble measured magnet behavior, allowing the model to learn general hysteresis features. Simple exponential eddy current decays are also included.

Pre-training leads to up to five times faster convergence during fine-tuning and improves validation and test performance. As shown in Table 1, the results on the main dipoles approach the accuracy and resolution limits of the online B-Train field measurement system at 2×10^{-5} T [6].

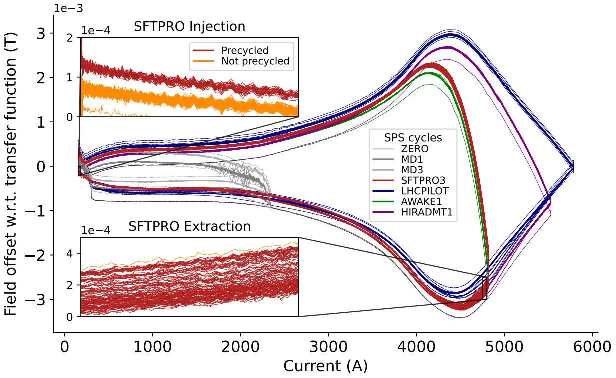

Figure 1: Cycle-to-cycle magnetic hysteresis in the SPS main dipoles under operational conditions. Insets highlight the SFTPRO cycle at injection (14 GeV, 0.063 T) and at the 400 GeV slow extraction plateau (1.80 T). A consistent field deviation of up to 6×10^{-5} T is measured at injection when the SFTPRO cycle is played without a pre-cycle.

---

* anton.lu@cern.ch

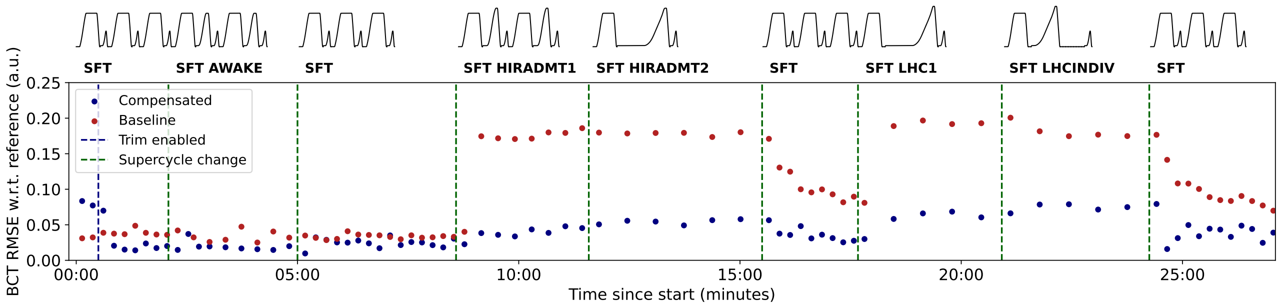

Figure 2: Operational supercycle changes, with and without hysteresis compensation, on the SFT flat-top, showing spill quality via the Root Mean Square Error (RMSE) between normalized measured and reference Beam Current Transformer (BCT) intensity values. Lower RMSE indicates better reproducibility. Overall spill improvement can be seen in Figure 3. The magnetic sequences (top) include 400 GeV SFTPRO, 450 GeV LHC injection, 440 GeV HiRadMat, and 400 GeV AWAKE cycles. The 200 GeV MD1 pre-cycle precedes every SFTPRO cycle.

Table 1: RMSE calculated between the measured magnetic field from SPS main dipoles and the hysteresis model prediction, across the entire dataset. Local errors may be higher.

	W/o pre-training (T)	Fine-tuned (T)
Validation	6.57×10^{-5}	3.38×10^{-5}
Test	9.39×10^{-5}	6.50×10^{-5}

Field Predictions without Feedback

For higher-order magnets in the SPS, and many other magnetic circuits in CERN synchrotrons, autoregressive prediction is required. Operational tests in 2024 demonstrated that the deployed online model can stably predict and compensate the field for over an hour, even with unfamiliar supercycles, improving spill quality compared to uncompensated operation. Robust performance depends on the model being trained on sequences similar to those encountered in operation, which is feasible given the typically predictable SPS magnetic sequences.

Autoregressive prediction was also found advantageous for SPS main dipoles, even when direct field measurements are available. This compensates non-reproducibility attributed to drift in fluxmeter integration in the online field measurements, which can vary between cycles and affect prediction accuracy. As the model cannot distinguish between drift and true hysteresis effects, using autoregression may yield more robust cycle-to-cycle predictions.

Eddy Current Compensation

Eddy current decay introduces time-dependent field decays after fast ramp-downs, visible in field measurements. However, operational studies in 2024 revealed a significant mismatch between field-based correction and actual beam response. While our previous approach suggested successful correction through field forecasting, new evidence indicates that while the hysteresis model compensates injection dipole fields with respect to the B-Train measurements, it fails to fully correct radial beam drift. Furthermore, some longer measured field drifts, previously attributed to dynamic decay, may actually result from field measurement integration drift not experienced by the beam.

To address this complexity, we now model eddy current decay as a sum of exponentials and train the hysteresis model on measurements with eddy current effects removed, avoiding compensating for the same phenomenon twice. This approach not only prevents overcompensation but also simplifies hysteresis modeling by reducing the importance of prediction time resolution. This complementary approach allows for analytical modeling of dynamic effects directly on the beam, and thus more robust feed-forward corrections.

OPERATIONAL RESULTS

Figure 2 shows hysteresis compensation in use in currently operational magnetic sequences, including beams to the HiRadMat [7] and AWAKE [8] experiments. Using B-Train measurements for predictions (rather than autoregression), we compensate the dipole field on the SFTPRO flat-top, where the most significant beam degradations due to hysteresis are shown. The reduction in spill variation between different supercycles (Figure 3, main) shows the importance of using active field compensation, as well as the benefit of combining this approach with analytic modeling of dynamic effects on the beam (Figure 3, insets). However, the spill is not perfectly reproduced, likely due to still uncompensated hysteresis in the main quadrupole circuits.

Operational flat-top field predictions consistently achieved 3×10^{-5} T accuracy relative to measurements, a big improvement over previous results at 5×10^{-5} T. Rare mispredictions up to 2×10^{-4} T during extended tests degraded beam quality transiently, but were automatically recovered within two cycles.

Offline validation of SFTPRO cycles without magnetic pre-cycle (Figure 4) demonstrates below 5×10^{-5} T accuracy across cycles when the beam is present. While we can often demonstrate $\approx 3 \times 10^{-5}$ T accuracy at SFTPRO injection, further work is required to prove robustness for full operational deployment without beam loss. Still, these

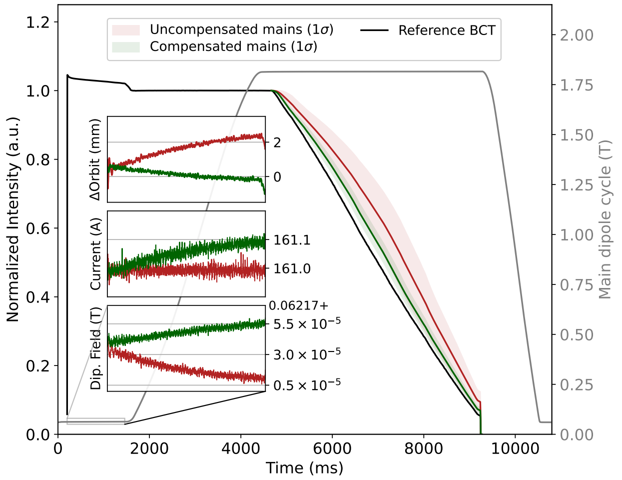

Figure 3: SFTPRO cycle with and without hysteresis compensation at extraction. The reference BCT shows an ideal, linear, spill to the SPS North Experimental Area. The shaded areas show the measured spills of the SFTPRO cycles in Figure 2 with and without compensation. The insets show analytical feed-forward eddy current correction for the same cycle at injection.

results represent important progress toward energy-efficient, flexible beam scheduling.

Potential Operational Benefits

Hysteresis compensation primarily reduces the need for manual adjustments when switching magnetic sequences during supercycle changes or dynamic operation. Since quadrupoles are not yet compensated, small tune corrections by operators are still required to reach an optimal beam quality. Full automation will become more important in future, when dynamic scheduling and automated LHC filling will increase the frequency of sequence changes.

The pre-cycling using the MD1 cycle, significantly reduces beam availability and increases energy consumption. During the 2024 physics run, the 3.6 s MD1 cycle was played over one million times, consuming 5 GWh of power and occupying 18 % of operational time. Successful hysteresis compensation therefore has the potential to significantly improve both machine availability and energy efficiency. However, once the pre-cycles are removed, I_{RMS} restrictions for the magnetic circuits will limit the number of additional high-energy cycles that can be played, in relation to lower energy machine development cycles.

LIMITATIONS

Although the hysteresis model predicts the field more accurately than 5×10^{-5} T at injection, this is not yet sufficient to eliminate the SFTPRO pre-cycle. The main challenges are limited training data for cycles without a pre-cycle (most operational data include the MD1 pre-cycle), and measurement instabilities, such as integration drift in the B-Train system, which can introduce errors exceeding 5×10^{-4} T. This drift is much greater than the required accuracy and critically limits its compensation performance. Laboratory measurements of the main dipoles [9] are also constrained by hardware limitations, preventing replication of operational cycles and reaching adequate absolute field measurement precision.

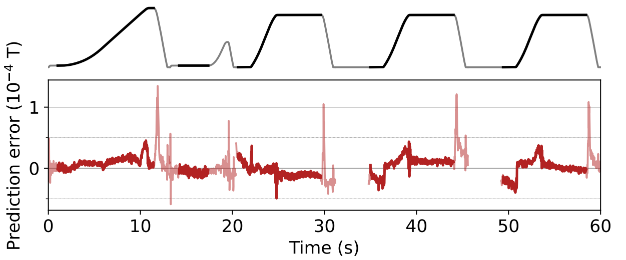

Figure 4: Field predictions on sequences without precycle, only for cycles with beam. The top plot shows the magnetic cycles, where the second cycle has a similar effect to the MD1 pre-cycle. Darker lines indicate beam presence.

The most significant limitation for our work is therefore the stability and accuracy of field measurements. When provided with sufficient high-quality data, the model can achieve prediction consistency on a par with, or better than, direct measurements, particularly when integration drift is accounted for. Routine compensation without a pre-cycle across all cycles, especially at injection energy, requires a field accuracy better than 2×10^{-5} T. This is beyond current pulsed measurement capabilities for the SPS main magnets at CERN. Eliminating the pre-cycle also depends on extending compensation to quadrupoles and sextupoles, which are expected to present similar challenges with their respective offline field measurements.

CONCLUSIONS AND FUTURE WORK

This contribution demonstrates substantial progress in data-driven magnetic hysteresis compensation for particle accelerator magnets. The field prediction accuracy is improved to below 3×10^{-5} T at flat-top and below 5×10^{-5} T at injection energy with respect to measurement, and beam reproducibility is consistently improved within the errors of higher-order magnetic hysteresis. While operational benefits are currently limited by rigid beam scheduling, hysteresis and eddy current compensation prepare the SPS for further automation and improved beam parameter reproducibility. While the removal of the pre-cycle is conditioned on successful compensation of the main quadrupole and sextupole circuits, the proposed transfer learning method significantly simplifies learning a new accurate model.

The success of this solution, without feedback systems or reference magnets, suggests this as a practical, deployable solution for magnets where online field measurements are missing. As the online compensation algorithms are now fully integrated in the CERN control system, flexible model selection allows for continuous refinement and rapid deployment to other magnets and other multi-cycling synchrotrons at CERN.

REFERENCES

[1] A. Lu, V. Kain, C. Petrone, V. D. Capua, M. Schenk, and C. Zannini, “First operational experience with data-driven hysteresis compensation for the main dipole magnets of the CERN SPS”, in *Proc. IPAC’24*, Nashville, TN, USA, pp. 874–877, 2024. `doi:10.18429/JACoW-IPAC2024-MOPS66`

[2] V. Kain, K. Cornelis, and E. Effinger, “New Spill Control for the Slow Extraction in the Multi-Cycling SPS”, in *Proc. IPAC’16*, Busan, Korea, pp. 1371–1373, 2016. `doi:10.18429/JACoW-IPAC2016-TUPMR051`

[3] V. Kain *et al.*, “Resonant slow extraction with constant optics for improved separatrix control at the extraction septum”, *Phys. Rev. Accel. Beams*, vol. 22, no. 10, p. 101 001, 2019. `doi:10.1103/PhysRevAccelBeams.22.101001`

[4] B. Lim, S. O. Arik, N. Loeff, and T. Pfister. “Temporal Fusion Transformers for Interpretable Multi-horizon Time Series Forecasting”, `http://arxiv.org/abs/1912.09363` pre-published.

[5] D. C. Jiles and D. L. Atherton, “Theory of ferromagnetic hysteresis (invited)”, *Journal of Applied Physics*, vol. 55, no. 6, pp. 2115–2120, 1984. `doi:10.1063/1.333582`

[6] D. Giloteaux *et al.*, “First operational results of new real-time magnetic measurement systems for accelerator control”, in *Proc. IPAC’23*, Venice, Italy, pp. 3996–3999, 2023. `doi:10.18429/JACoW-IPAC2023-THPA024`

[7] F. J. Harden, A. Bouvard, N. Charitonidis, and Y. Kadi, “HiRadMat: A Facility Beyond the Realms of Materials Testing”, in *Proc. IPAC’19*, Melbourne, Australia, May 2019, pp. 4016–4019. `doi:10.18429/JACoW-IPAC2019-THPRB085`

[8] E. Gschwendtner *et al.*, “AWAKE, The Advanced Proton Driven Plasma Wakefield Acceleration Experiment at CERN”, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 829, pp. 76–82, 2016. `doi:10.1016/j.nima.2016.02.026`

[9] M. Taupadel *et al.*, “Linear weight optimization of local magnetic field sensors for the integral field measurement in accelerator magnets”, presented at IPAC’25, Taipei, Taiwan, Jun. 2025, paper WEPB024, this conference.