UNDERSTANDING ENERGY-INDUCED OPTICS DISTORTIONS IN THE LHC

J. Gray*, C. Caliari, F. Carlier, J. Dilly, S. Horney, J. Keintzel, E. Maclean, K. Sabin,
K. Skoufaris, F. Soubelet, M. Stefanelli, R. Tomás, W. Van Goethem,
European Organization for Nuclear Research, Geneva, Switzerland
Y. Angelis, Aristotle University of Thessaloniki, Thessaloniki, Greece
A. Oeftiger, John Adams Institute, University of Oxford, Oxford, United Kingdom

Abstract

Small momentum offsets in the LHC can generate significant optics distortions, particularly at low β^*. The beam energy carries a relative uncertainty of approximately 10^{-3}, which is insufficient for precise optics control. To better understand the energy contribution to optics measurements, two beam-based techniques have been explored. The first applies a global linear response matrix between BPM phase advances and relative energy offset, denoted as $\Delta p/p$. While effective in simulation, this method is sensitive to phase errors and tends to underestimate the energy shift. We introduce a new approach based on the principle of Deep Lie Map Networks (DLMN), which fits a differentiable tracking model to measured turn-by-turn BPM trajectories. Using the symplectic single-pass forward differentiation capability of MAD-NG, we can fit our simulation to the optimum $\Delta p/p$ that matches the measurement. The results reveal arc-by-arc variations consistent with varying dipolar errors from arc to arc, providing insight into orbit behaviour around the ring. The measured response also agrees well with intentional energy shifts, demonstrating that the DLMN offers a promising new method for analysing the effect of energy on the optics of the LHC.

INTRODUCTION

Small energy offsets in the LHC have been observed to induce measurable optics distortions that become increasingly important as the machine operates at lower β^*. For example, a relative energy change on the order of 10^{-4} was linked to an observed ~10% increase in beta beating over a two-week period [1]. If the energy shift is left uncorrected, such drifts force repeated optics corrections, unnecessarily increasing time spent commissioning the machine.

The dominant contributors to these energy-dependent optics shifts are the strongest focusing elements in the lattice: the inner triplet quadrupoles surrounding the major interaction points (ATLAS and CMS). With the HL-LHC upgrade and a design β^* of 15 cm, simulations indicate an even greater sensitivity: a relative energy change of 2×10^{-4} can produce up to 50% beta beating [2].

The root cause of energy drifts in the machine has been attributed to changes in the average orbit corrector kicks over time [1]. The RF system maintains path length and so the average corrector strength directly changes the beam energy [1]. Understanding and tracking these energy-induced optics changes is therefore essential to minimise time spent on optics correction and maximise time spent delivering luminosity for physics.

This paper compares two independent beam-based techniques for estimating momentum offsets. The first is a response matrix approach that relates changes in phase advances between BPMs to a change in $\Delta p/p$ [3, 4]. The second is the Deep Lie Map Network (DLMN) method that fits a differentiable symplectic tracking model to turn-by-turn data and extracts arc-level momentum deviations [5, 6].

GLOBAL RESPONSE MATRIX

The first method explored to determine the effect of energy on the optics assumes a linear relationship between changes in momentum offset and changes in observable quantities [3]:

$$\vec{M}(\delta + \Delta p/p) \approx \vec{M}(\delta) + \mathbf{R}(\delta)\Delta p/p \ , \quad (1)$$

where $\vec{M}(\delta)$ represents the model prediction for an observable vector at a given momentum deviation δ, $\Delta p/p$ is the relative momentum offset from this δ, and $\mathbf{R}(\delta)$ is the response matrix at this point [3]. By comparing measurements to model predictions, a correction can be determined [3]:

$$\vec{M}(\delta)+\mathbf{R}(\delta)\Delta p/p \approx \vec{O} \implies \Delta p/p \approx \mathbf{R}^{+}(\delta)\left(\vec{O} - \vec{M}(\delta)\right) \ , \quad (2)$$

where $\vec{O}$ is the stacked measured observable vector at multiple locations along the ring and $\mathbf{R}^{+}$ is the SVD-based pseudo-inverse of the response matrix. For this application, the horizontal and vertical phase advances between BPMs are used as the observable vectors, as they are sensitive to momentum offset, mainly due to the triplets.

From simulation in a lattice without errors, we find that this method has only a few percent error in the calculated relative momentum offset over the range of -3×10^{-4} to 3×10^{-4}, as can be seen in Fig. 1.

DEEP LIE MAP NETWORKS

The second method adapts the DLMN framework [5], replacing multipole-error optimisation with a single free parameter per arc, $\Delta p/p$, fitting the orbit by minimising position residuals via $\partial\vec{x}/\partial(\Delta p/p)$ from MAD-NG's symplectic forward differentiation [6]. Phase-space coordinates

* joshua.mark.gray@cern.ch

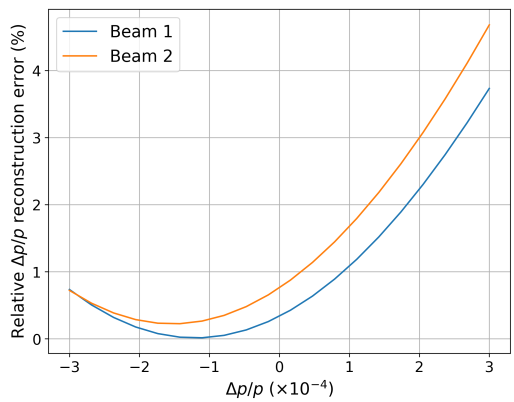

Figure 1: Relative reconstruction error of the linear response matrix method as a function of applied momentum offset for Beam 1 and Beam 2.

(x, p_x, y, p_y) are reconstructed from position-only BPM measurements.

To estimate the transverse momenta, we treat the beam as a single particle executing betatron oscillations on an ellipse in phase space [7]. In a linear machine, BPMs separated by $\pi/2$ in betatron phase correspond to a 90-degree rotation in phase space, which transforms the position coordinate into the momentum coordinate [7]. A position measurement at phase $\phi + \pi/2$ then provides an estimate of the momentum at phase ϕ [7]. The closed orbit (x, p_x, y, p_y) at each BPM is then retrieved by averaging the reconstructed coordinates and momenta across multiple turns.

Since the LHC arcs provide more favourable conditions for applying the DLMN, i.e. BPMs with a higher signal-to-noise ratio and smaller magnetic errors than the rest of the machine, the closed orbit was fitted independently in each arc. This arc-by-arc approach enables us to observe local variations in the energy-optics correlation, which are expected to differ around the ring due to the non-uniform distribution of dipolar field errors and other misalignments.

Before performing the fit, we extract the field strengths in the machine at the measurement time, so that the simulation includes the measured dipole, quadrupole, sextupole currents.

IDENTIFYING ENERGY SHIFTS

To evaluate the performance of the two methods, we took several measurements on both beams in the LHC where we deliberately introduced a known momentum offset by changing the average corrector magnet strengths along with a compensation for the resulting tune shift, using a pre-calculated set of strengths. During measurements, we varied an intentional momentum offset from -2×10^{-4} to 2×10^{-4} in steps of 1×10^{-4}, and then we used both techniques to evaluate the relative momentum offset from the BPM data. The shifts to the energy were then verified using the measured dispersion and measured magnetic currents.

Phase-Based Method

Figure 2 shows the momentum offset calculated using the global response matrix method plotted against the known introduced momentum offset. While the method exhibits an overall linear trend, there is a significant variation between some points for beam 2. Linear fits yield slopes of 0.93 for Beam 1 and 0.89 for Beam 2, indicating a systematic underestimation in the measurement of the change in the $\Delta p/p$.

The error bars represent only the propagation of the phase measurement uncertainty through the response matrix inversion (Eq. 2), and do not account for systematic uncertainties in the response matrix itself. These systematics arise from model imperfections, phase errors, phase reference drifts, and the fundamental limitation of the linear approximation when applied to the non-linear optics response of the LHC. The quality of the fit and the scatter in the data highlight that the phase-based method works but is not reliable.

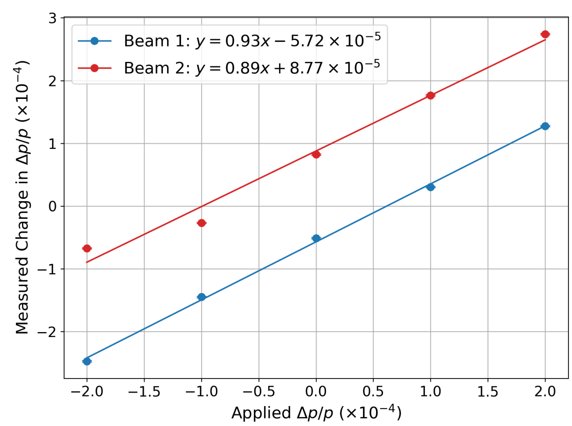

Figure 2: Comparison of the calculated momentum offset using the phase-based method to the known introduced momentum offset.

DLMN Method

In the DLMN method fitting is performed directly on the position measurements rather than phase advances. Due to the chosen method, we obtain arc-by-arc momentum offset values that reflect local optics distortions. As shown in Fig. 3, individual arcs exhibit significant variability, over 4×10^{-4} in a single measurement. However, all arcs respond consistently to changes in the applied momentum offset, with each arc shifting by approximately the same amount as the input $\Delta p/p$ is varied. This consistency across arcs at different applied offsets demonstrates the method's stability and reveals genuine differences between the arcs, most likely caused by magnet misalignments, variations in average bending magnet strength, and BPM alignment.

Figure 4 shows the average of the differences in the relative momentum in each arc, relative to zero offset, for each beam, revealing good agreement with the input. Linear fits yield slopes of (0.97 ± 0.01) for Beam 1 and (0.93 ± 0.03) for Beam 2, with the slopes lying within 3 standard deviations of

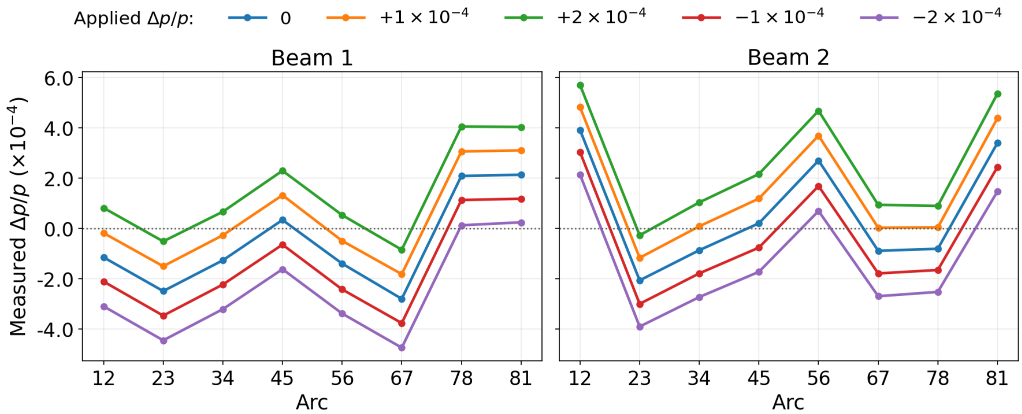

Figure 3: Arc-by-arc DLMN momentum offset for Beam 1 (left) and Beam 2 (right) at each $\Delta p/p$ setting (colour-coded). The blue line marks the no-shift reference. The arc label, *ab*, describes the arc between the interaction points *a* and *b* in the LHC.

unity. This represents a significantly better fit than the phase-based method and demonstrates that the DLMN approach successfully captures the momentum-dependent changes in the beam trajectories. The small error bars indicate good precision in measuring relative momentum changes, enabling the ability to identify energy drifts in the machine with precision and correct for them before they cause significant optics distortions.

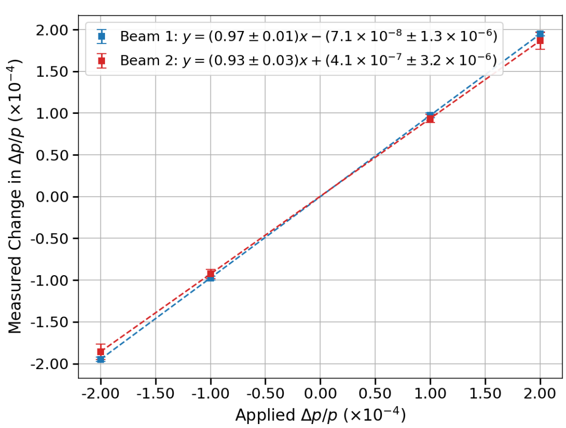

Figure 4: DLMN change in momentum offset relative to the zero-offset measurement reference point. Low error bars in estimated change in relative momentum demonstrate good precision in measuring these changes.

Technique Comparison

Figure 5 shows the mean absolute momentum offset from DLMN, obtained by averaging arc-by-arc fits, alongside the phase-based response from Figure 2. The DLMN absolute error bars (shown on the plot) are substantially larger than the relative error bars in Figure 4, a consequence of averaging the varying arc-by-arc values. Therefore, these large error bars reflect uncertainties on the intercept of the fit, not the gradient.

Within the uncertainties, both methods give consistent estimates. The phase-based method assumes phase beating and the tune correction originate solely from energy, which need not hold if other sources contribute. DLMN instead attributes the orbit shift between settings to energy; constant dipolar and quadrupolar contributions cancel in the difference, and the measured corrector strengths are embedded in the simulation, making this a more robust assumption.

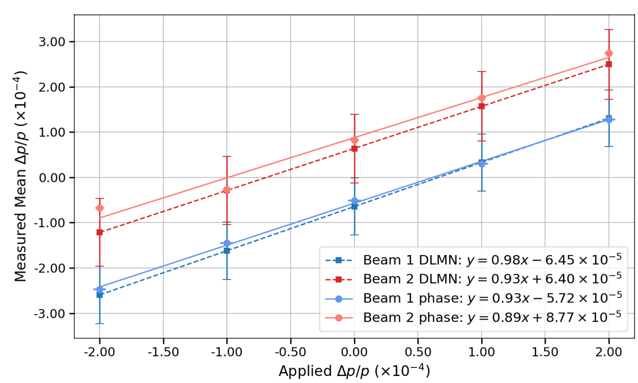

Figure 5: DLMN measurements averaged over all arcs versus known introduced momentum offset. Large error bars reflect arc-by-arc variations in the reconstruction, and poor knowledge of the absolute relative momentum offset.

CONCLUSION

We compared two beam-based techniques for measuring momentum offsets in the LHC, based on experimental data. The phase-based linear response matrix method shows significant scatter and 7–11% underestimation of the change, while the DLMN approach achieves slopes 3–7% underestimation but with superior precision. Despite different causes of systematic errors, both methods produce consistent results, validating the physics and demonstrating that DLMN enables reliable tracking of energy drifts. This will allow early detection and correction of momentum-induced optics distortions, minimising repeated optics corrections and improving optics commissioning for the LHC and HL-LHC, as demonstrated during the 2026 LHC commissioning where it proved to be instrumental.

REFERENCES

[1] F. Carlier *et al.*, "LHC run 3 optics corrections", in *Proc. IPAC'23*, Venice, Italy, pp. 535–538, May 2023. doi:10.18429/JACoW-IPAC2023-MOPL015

[2] S. Horney *et al.*, "Energy sensitivity of the high luminosity LHC optics at the end of the beta* squeeze", in *Proc. IPAC'25*, Taipei, Taiwan, Jun. 2025, pp. 318–321. doi:10.18429/JACoW-IPAC2025-MOPM018

[3] J. W. Dilly, L. Malina, and R. Tomas Garcia, "An Updated Global Optics Correction Scheme", *CERN Internal Note*, 2018. https://cds.cern.ch/record/2632945

[4] E. Fol, R. Tomás, and G. Franchetti, "Supervised learning-based reconstruction of magnet errors in circular accelerators", *The European Physical Journal Plus*, vol. 136, no. 4, p. 365, Apr. 2021. doi:10.1140/epjp/s13360-021-01348-5

[5] C. Caliari, A. Oeftiger, and O. Boine-Frankenheim, "Beam-based identification of magnetic field errors in a synchrotron using deep Lie map networks", *Phys. Rev. Accel. Beams*, vol. 28, no. 2, p. 024601, Feb. 2025. doi:10.1103/PhysRevAccelBeams.28.024601

[6] J. Gray *et al.*, "Deep Lie map networks from single–pass forward differentiation", in *Proceedings of the 17th International Particle Accelerator Conference (IPAC'26)*, May 2026.

[7] R. Tomás, M. Bai, R. Calaga, W. Fischer, A. Franchi, and G. Rumolo, "Measurement of global and local resonance terms", *Phys. Rev. ST Accel. Beams*, vol. 8, no. 2, p. 024001, Feb. 2005. doi:10.1103/PhysRevSTAB.8.024001