DEEP LIE MAP NETWORKS FROM SINGLE-PASS FORWARD DIFFERENTIATION

J. Gray*, C. Caliari, F. Carlier, L. Deniau, K. Skoufaris
European Organization for Nuclear Research, Geneva, Switzerland
A. Oeftiger, John Adams Institute, University of Oxford, Oxford, United Kingdom

Abstract

A new implementation of Deep Lie Map Networks (DLMNs) based on the single-pass forward-differentiation capability of MAD-NG is presented. DLMNs were introduced in earlier work as an optics-modelling framework in which lattice parameters are fitted directly to measured BPM trajectories. In the present implementation, exact derivatives of particle coordinates with respect to lattice parameters are computed during tracking. Rather than relying on backpropagation or finite differences, gradients are evaluated on the fly during tracking, so intermediate states do not need to be stored. This reduces memory use and computation time while preserving precision. The result is a scalable tool for data-driven optics reconstruction that is well suited to larger multi-turn data sets and, ultimately, to digital-twin applications in which model parameters are inferred continuously from live measurements.

INTRODUCTION

Achieving and maintaining good performance in synchrotrons relies on accurate knowledge of the underlying lattice optics [1, 2]. Magnet strength errors, misalignments, and higher-order field components all lead to optics distortions that can degrade beam quality and reduce dynamic aperture, beam lifetime, and machine performance [1, 2]. Identifying and correcting these errors is therefore critical.

Awidely used method for linear optics correction is Linear Optics from Closed Orbits (LOCO), originally introduced in [3]. LOCO determines lattice errors by fitting a model orbit response matrix to measured data, enabling the extraction of quadrupole strength errors and beam position monitor (BPM) calibration factors [3, 4]. This method can be slow and cumbersome for large machines, may fail to converge, and remains sensitive to noise and systematic errors [1, 5]; moreover, typically it is limited to only linear optics, although there exists extensions to nonlinear optics from off-energy closed orbits (NOECO) [6, 7]. On the other hand, methods widely used for example at the Large Hadron Collider (LHC), rely on optics measurements of AC-dipole-driven turn-by-turn data; however, these approaches typically target a specific machine state and do not inherently provide a simulation model applicable to all or even any machine state.

Deep Lie Map Networks (DLMNs), introduced in [8], offer a framework in which an accelerator lattice is treated as a parametrised model whose parameters are optimised directly against measurements. The first implementation of this approach was developed in Julia, using backpropagation techniques from the machine-learning community. In that work, a custom tracking code was implemented to optimise quadrupole and sextupole strengths against measured data [6].

In this work, we present a new DLMN implementation based on the forward-differentiation capabilities of Methodical Accelerator Design – Next Generation (MAD-NG) [9]. By exploiting its truncated power series algebra (TPSA) framework, derivatives of particle coordinates with respect to arbitrary lattice parameters can be computed exactly during tracking. This enables step-by-step gradient evaluation without finite differences or backpropagation, improving efficiency, numerical stability, and memory usage. In contrast to LOCO, which relies on dedicated and expensive beam-time-intensive measurements and structured orbit-response data, DLMN can make use of more routine machine data. The approach has already been applied to LHC measurements to identify beam-energy offsets [10].

DLMN DESIGN

The parametrised lattice model can be interpreted as a machine-learning network in which magnet parameters, such as strengths or misalignments, are optimised to reproduce measured beam positions [6]. The optimisation aims to reduce the loss function using stochastic gradient descent, $\mathscr{L}$, defined as

$$\mathscr{L} = \frac{1}{2} \sum_i \left[(x_{s,i} - x_{m,i})^2 + (y_{s,i} - y_{m,i})^2 \right], \qquad (1)$$

which quantifies the discrepancy between the measured and simulated beam position monitor (BPM) coordinates, $(x_{m,i}, y_{m,i})$ and $(x_{s,i}, y_{s,i})$, respectively. The index i spans all BPMs, turns, and initial conditions. To optimise a parameter k, the derivative of the loss with respect to k is required. Applying the chain rule gives

$$\frac{\mathrm{d}\mathscr{L}}{\mathrm{d}k} = \sum_i \left(\frac{\partial \mathscr{L}}{\partial x_{s,i}} \frac{\mathrm{d}x_{s,i}}{\mathrm{d}k} + \frac{\partial \mathscr{L}}{\partial y_{s,i}} \frac{\mathrm{d}y_{s,i}}{\mathrm{d}k} \right). \qquad (2)$$

Substituting Eq. (1) into Eq. (2) gives

$$\frac{\mathrm{d}\mathscr{L}}{\mathrm{d}k} = \sum_i \left[(x_{s,i} - x_{m,i}) \frac{\mathrm{d}x_{s,i}}{\mathrm{d}k} + (y_{s,i} - y_{m,i}) \frac{\mathrm{d}y_{s,i}}{\mathrm{d}k} \right]. \qquad (3)$$

The remaining task is therefore to obtain $\mathrm{d}x_{s,i}/\mathrm{d}k$ and $\mathrm{d}y_{s,i}/\mathrm{d}k$, either by automatic differentiation or by finite differences.

* joshua.mark.gray@cern.ch

IMPLEMENTATION FRAMEWORK

In MAD-NG, the required derivatives are calculated directly during tracking using TPSA objects. MAD-NG combines a LuaJIT front end with computationally intensive C and C++ components, including the TPSA objects and the physics maps [9,11,12]. This structure provides direct access to tracking internals and allows coordinates and derivatives to be extracted within the same tracking pass.

Truncated Power Series Algebra

TPSA is the bedrock of MAD-NG [9] and has been successfully used in the Polymorphic Tracking Code (PTC) [13]. Instead of tracking a particle only through its phase-space coordinates, each coordinate is promoted to a TPSA object representing a truncated Taylor polynomial that is differentiated automatically at every step.

Each object is truncated at a specified order and represents the particle as a Taylor expansion about its phase-space point, exact to that order. Propagation through the lattice uses exact maps where available, thick–kick–thick maps for quadrupoles, and drift–kick–drift (DKD) maps for sextupoles, all integrated using a symplectic scheme, such as the fourth-order Yoshida integrator [14].

MAD-NG also uses a generalised TPSA formalism in which user-defined parameters can be treated as differentiation variables [15], referred to here as parametric tracking. This allows direct extraction of $\mathrm{d}x_{s,i}/\mathrm{d}k$ and $\mathrm{d}y_{s,i}/\mathrm{d}k$ for any fitted parameter during tracking, avoiding backpropagation while retaining full generality.

Python

PyMAD-NG provides access to MAD-NG commands and internal data structures from Python [16]. Its ability to run multiple MAD-NG instances simultaneously allows trajectories with different initial conditions to be executed in parallel, which is useful for large data sets on multithreaded machines. Data extraction and optimisation are handled in Python, so the implementation can use mature optimisation and multiprocessing tools.

PROVIDING INITIAL CONDITIONS

Accurate reconstruction of the initial conditions is essential for reliable DLMN fitting. When these are noisy or poorly constrained, the loss function becomes ill-conditioned, leading to unreliable optimisation and noise-dominated parameter estimates. Earlier studies addressed this by acquiring hundreds of measurements with nearly identical initial conditions using a single-turn kicker [6]. As the kicker excites the beam only once per machine cycle, each cycle yields a single trajectory, making the procedure time-consuming.

In contrast, optics measurements in machines such as the LHC are typically performed with an AC dipole, which produces adiabatic oscillations over thousands of turns [17, 18]. This extended excitation allows multiple short trajectories to be extracted within a single cycle, substantially reducing the required measurement time. However, because the excitation ramps over roughly 1000 turns, reproducing identical oscillations between kicks remains challenging, motivating the need for an alternative reconstruction strategy.

To address this, a more general reconstruction method can be employed in which the beam is treated as a single particle undergoing betatron motion on an ellipse in phase space [19]. The idea is to use neighbouring BPM measurements together with the known optics to infer the unknown transverse momenta. In a linear lattice, BPMs separated by a phase advance of $\pi/2$ correspond to a 90-degree rotation in phase space, effectively mapping position into momentum [19].

For each transverse plane $u \in \{x, y\}$, define

$$\tilde{u} = \frac{u - \delta D_u}{\sqrt{\beta_u}}, \qquad \tilde{u}_n^{(\pm)} = \frac{u_n^{(\pm)} - \delta D_{u,n}^{(\pm)}}{\sqrt{\beta_{u,n}^{(\pm)}}}, \tag{4}$$

where $\delta \equiv \Delta p/p$, and $(\pm)$ denote the previous (upstream) and next (downstream) $\pi/2$-neighbour BPMs. Let

$$\phi_u^{(\pm)} = 2\pi\Delta_u^{(\pm)}, \qquad A_u^{(\pm)} = \tan\phi_u^{(\pm)} \mp \alpha_u, \tag{5}$$

where $\Delta_u^{(\pm)}$ is the signed measured phase error relative to $\pi/2$, and α_u, β_u, D_u, and D_{pu} are evaluated at the reference BPM, with model or measurement values, depending on accuracy. The two momentum estimates are

$$p_u^{(\pm)} = \pm\frac{\tilde{u}_n^{(\pm)}\sec\phi_u^{(\pm)} + \tilde{u}A_u^{(\pm)}}{\sqrt{\beta_u}} + D_{pu}\delta. \tag{6}$$

The momentum deviation is calculated using $\delta = \langle D_{x,i}\, CO_{x,i}\rangle / \langle D_{x,i}\rangle$ [20]. They are combined by inverse-variance weighting,

$$\hat{p}_u = \frac{\sum_{\eta=\pm} p_u^{(\eta)}/\sigma_{\eta,u}^2}{\sum_{\eta=\pm} 1/\sigma_{\eta,u}^2}, \tag{7}$$

where $\sigma_{\eta,u}$ denotes the standard deviation associated with each estimate, obtained from the optics parameters and BPM coordinate uncertainties, which are assumed to be uncorrelated. This approach allows the initial momenta to be inferred directly from measured orbit data and known machine optics, without requiring repeated identical excitations. Its main limitation is that it assumes no magnetic errors between the BPMs used for the reconstruction, so the inferred initial conditions can still be biased if the optics in that region are already incorrect. Extending the method to use N BPMs [21] should reduce this sensitivity and improve the momentum estimates.

During optimisation, the updated optics parameters can be used to repeat the momentum reconstruction and improve the accuracy of the inferred initial conditions [6]. Another approach for noisier conditions is to use a much larger number of initial conditions, so that the loss and its gradients are evaluated over a larger sample and the inferred coordinates become less sensitive to noise. This latter approach is now practical thanks to the faster MAD-NG implementation.

CODE COMPARISON

To compare the original Julia implementation of DLMN with the MAD-NG version, both were evaluated on the same SIS18 ring [22] using identical quadrupole perturbations drawn as Gaussian errors with 1% rms. To avoid favouring either code, reference free-oscillation tracking data was generated independently in XSuite [23], without noise and with exact initial momenta.

The comparison was constructed so that the underlying machine model remained identical, while the tracking settings were chosen to make the implementations as consistent as possible. In XSuite, this meant fourth-order Yoshida tracking with the full bend-fringe model, whereas Julia used sixth-order Yoshida DKD maps with two slices per element. This keeps the comparison focused on differences between the implementations rather than tracking inaccuracies. The data was split into four training samples and one validation sample, and each code ran for 3000 iterations.

Accuracy

Figure 1 compares the relative β-beating in the horizontal and vertical planes for the Julia and MAD-NG implementations, before and after fitting 36 quadrupole strengths, with respect to the perturbed XSuite reference model. Before optimisation, both implementations show almost identical baseline discrepancies along the ring. After fitting, both reduce the relative β-beating by roughly one to two orders of magnitude, indicating that the two approaches achieve comparable optics reconstruction quality in this benchmark.

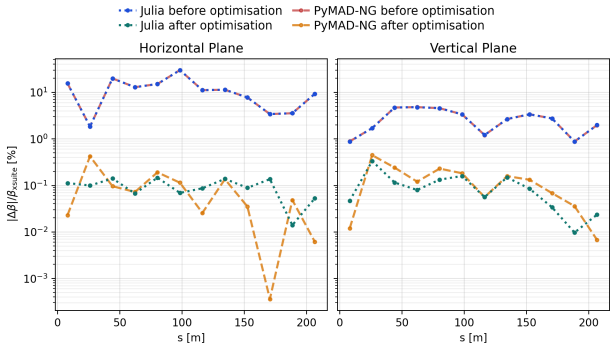

Figure 1: Relative β-beating along the ring in the horizontal and vertical planes, before and after fitting 36 quadrupole strengths for the MAD-NG and Julia implementations. The perturbed XSuite model is used as the reference. Without errors, the maximum relative β-beating between XSuite and MAD-NG is 3×10^{-10}.

Efficiency

Figure 2 shows that the PyMAD-NG implementation is overall more than 15 times faster than the Julia implementation. Part of this gain comes from parallel execution across multiple MAD-NG instances, while another important contribution arises from the tracking model and TPSA implementation. This computational advantage not only reduces optimisation time but also makes it practical to evaluate the loss and gradients over much larger multi-turn data sets in noisy conditions than was feasible in the earlier Julia implementation. Independent memory measurements, not shown in the figure, give a peak RAM usage of 722 ± 4 MiB for MAD-NG and 805 ± 25 MiB for the Julia code, indicating that the MAD-NG implementation is also slightly more memory efficient despite the use of four parallel MAD-NG instances.

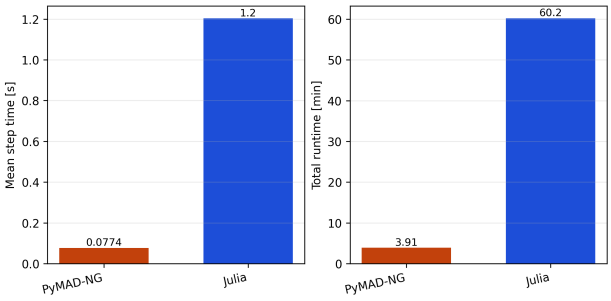

Figure 2: Comparison of mean step time and total optimisation runtime for the PyMAD-NG and Julia implementations.

Generality

The implementation was designed to be largely machine independent and has been used in simulation studies of the LHC, Super Proton Synchrotron, SIS18, and Proton Synchrotron Booster (PSB). The method has been validated on real LHC measurements [10], with further measurements in the PSB currently being analysed.

DLMNs can in principle be used to establish effective accelerator models in a data-driven way by fitting beam energy, quadrupole displacements, fringe parameters, and other quantities of interest; together with its ability to use routine measurement data, this makes it a promising framework for accelerator digital twins. MAD-NG makes this straightforward by enabling differentiation with respect to any parameter entering the tracking [9].

CONCLUSION

A new DLMN implementation based on MAD-NG forward differentiation has been presented. Parametric TPSA tracking computes exact coordinate derivatives with respect to lattice parameters in a single pass, removing the need for custom backpropagation while preserving precision. The approach has been successfully demonstrated on real LHC measurements, validating its applicability to operational accelerators. Compared with the earlier Julia implementation, the MAD-NG version achieves similar optics-reconstruction quality at significantly lower computational cost.

The proposed momentum-reconstruction strategy enables the extraction of multiple short trajectories from a single excitation cycle, reducing the number of synchrotron cycles required to collect trajectory data and saving measurement time. Together, these results make DLMN in MAD-NG a practical basis for measurement-driven optics modelling and future accelerator digital twins.

REFERENCES

[1] R. Tomás, M. Aiba, A. Franchi, and U. Iriso, "Review of linear optics measurement and correction for charged particle accelerators", *Phys. Rev. Accel. Beams*, vol. 20, no. 5, p. 054801, May 2017. doi:10.1103/PhysRevAccelBeams.20.054801

[2] T. H. B. Persson, Y. Inntjore Levinsen, R. Tomás, and E. H. Maclean, "Chromatic coupling correction in the Large Hadron Collider", *Phys. Rev. ST Accel. Beams*, vol. 16, no. 8, p. 081003, Aug. 2013. doi:10.1103/PhysRevSTAB.16.081003

[3] J. Safranek, "Experimental determination of storage ring optics using orbit response measurements", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 388, no. 1, pp. 27–36, 1997. doi:10.1016/S0168-9002(97)00309-4

[4] B. H. Oh, J. Lee, M. Yoon, and S. H. Shin, "Beam optics correction in the PLS-II storage ring", *Journal of the Korean Physical Society*, vol. 63, no. 7, pp. 1268–1273, 2013. doi:10.3938/jkps.63.1268

[5] M. Aiba, M. Böge, J. Chrin, N. Milas, T. Schilcher, and A. Streun, "Comparison of linear optics measurement and correction methods at the Swiss Light Source", *Phys. Rev. ST Accel. Beams*, vol. 16, no. 1, p. 012802, Jan. 2013. doi:10.1103/PhysRevSTAB.16.012802

[6] 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

[7] D. K. Olsson, Å. Andersson, and M. Sjöström, "Nonlinear optics from off-energy closed orbits", *Phys. Rev. Accel. Beams*, vol. 23, no. 10, p. 102803, Oct. 2020. doi:10.1103/PhysRevAccelBeams.23.102803

[8] C. Caliari, A. Oeftiger, and O. Boine-Frankenheim, "Identification of magnetic field errors in synchrotrons based on deep Lie map networks", *Phys. Rev. Accel. Beams*, vol. 26, no. 6, p. 064601, Jun. 2023. doi:10.1103/PhysRevAccelBeams.26.064601

[9] L. Deniau, "MAD-NG, a standalone multiplatform tool for linear and non-linear optics design and optimisation", *arXiv*, Jan. 2025. doi:10.48550/arXiv.2412.16006

[10] J. M. Gray *et al.*, "Understanding energy-induced optics distortions in the LHC", in *Proceedings of the 17th International Particle Accelerator Conference (IPAC'26)*, May 2026.

[11] R. Ierusalimschy, "Lua 5.2 reference manual", [Accessed 27-03-2026], Lua.org, 2015, https://www.lua.org/manual/5.2/,

[12] M. Pall, "The LuaJIT Project — luajit.org", https://luajit.org/, accessed 27-03-2026,

[13] F. Schmidt, E. Forest, and E. McIntosh, "Introduction to the polymorphic tracking code", CERN, Geneva, Rep., 2002. https://cds.cern.ch/record/573082

[14] H. Yoshida, "Construction of higher order symplectic integrators", *Physics Letters A*, vol. 150, no. 5, pp. 262–268, 1990. doi:10.1016/0375-9601(90)90092-3

[15] L. Deniau and CI. Tomoiagă, "Generalised truncated power series algebra for fast particle accelerator transport maps", in *Proc. 6th International Particle Accelerator Conference (IPAC'15), Richmond, VA, USA, May 3-8, 2015*, Richmond, VA, USA, pp. 374–377, Jun. 2015. doi:10.18429/JACoW-IPAC2015-MOPJE039

[16] J. Gray and L. Deniau, "PyMAD-NG, methodicalacceleratordesign/MAD-NG.py: Release v0.9.1", version v0.9.1, Mar. 2026, doi:10.5281/zenodo.19351899,

[17] R. Tomás, "Adiabaticity of the ramping process of an ac dipole", *Phys. Rev. ST Accel. Beams*, vol. 8, no. 2, p. 024401, Feb. 2005. doi:10.1103/PhysRevSTAB.8.024401

[18] E. H. Maclean *et al.*, "New approach to LHC optics commissioning for the nonlinear era", *Phys. Rev. Accel. Beams*, vol. 22, no. 6, p. 061004, Jun. 2019. doi:10.1103/PhysRevAccelBeams.22.061004

[19] 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

[20] J. Keintzel *et al.*, "New findings on BPM calibration in the LHC", presented at IPAC'26, Deauville, France, May 2026, paper WEP6080, this conference.

[21] A. Wegscheider and R. T. García, "An N-BPM momentum reconstruction for linear transverse coupling measurements in LHC and HL-LHC", in *Proc. IPAC'22*, Bangkok, Thailand, pp. 519–522, Jul. 2022. doi:10.18429/JACoW-IPAC2022-MOPOTK032

[22] K. Blasche, B. Franczak, B. Langenbeck, G. Moritz, and C. Riedel, "The heavy ion synchrotron SIS-a progress report", in *Proceedings of International Conference on Particle Accelerators*, 3736–3738 vol.5, 1993. doi:10.1109/PAC.1993.309772

[23] G. Iadarola *et al.*, "Xsuite: An integrated beam physics simulation framework", in *Proc. HB'23*, Geneva, Switzerland, pp. 73–80, Mar. 2024. doi:10.18429/JACoW-HB2023-TUA2I1