DEVELOPING A PHYSICS-INFORMED GAUSSIAN PROCESS MODEL TO CONSTRUCT AN UNCERTAINTY-QUANTIFIED MACHINE MODEL USING BAYESIAN OPTIMIZATION

V. Isensee*, O. Boine-Frankenheim[1], TEMF, Technical University Darmstadt, Darmstadt, Germany
A. Oeftiger, John Adams Institute for Accelerator Science, University of Oxford, Oxford, UK
E. Kazantseva, GSI Helmholtz Centre for Heavy Ion Research, Darmstadt, Germany
[1]also at GSI Helmholtz Centre for Heavy Ion Research

Abstract

For maintaining the designed optical performance of machines like the GSI heavy-ion synchrotron SIS18 or the FAIR fragment separator SFRS, accurate knowledge of magnetic field and element alignment is crucial. A Gaussian Process model with a physics-informed kernel based on a stochastic ensemble of lattices is proposed. Key advantages compared to Linear Optics from Closed Orbits (LOCO) include (1.) fitting of Gaussian probability distributions for parameters, which inherently model uncertainty, (2.) incorporation of measurement uncertainty from BPM noise, (3.) uncertainty-enabled orbit prediction between BPMs, and (4.) an active-learning strategy for greater sample efficiency than measuring an orbit response matrix. Applied in a simulation of SIS18, the method constructs an effective machine model with minimal orbit uncertainty around the ring, enabling orbit correction with uncertainty-quantified minimal deviation at any location. The SFRS is identified as a compelling target for future application, where sparse instrumentation and complex optics make uncertainty-aware modeling particularly valuable. This physics-informed GP framework provides uncertainty-enabled modeling for rings and beamlines, with potential applications extending to broader optics-correction and beam-tuning tasks.

INTRODUCTION

Accurate knowledge of element alignment and magnetic field errors is a prerequisite for maintaining the designed optical performance of modern accelerator facilities. Conventional methods such as Linear Optics from Closed Orbits (LOCO) [1] reconstruct optics from orbit response matrix (ORM) measurements but return a single best-fit solution without any quantification of residual uncertainty. For facilities operating at the edge of their dynamic range – such as during slow extraction in SIS18, or when commissioning the complex optics of the SFRS [2] – knowing not just the most probable orbit but also the confidence in that prediction is essential.

SIS18 is the heavy-ion synchrotron at GSI in Darmstadt, Germany: a 216 m-circumference ring of 12 identical sections, each with two dipoles and three quadrupoles, equipped with 12 BPMs [3]. The SFRS is the superconducting fragment separator under construction at FAIR, designed to separate exotic rare isotopes produced by projectile fragmentation or fission [4].

Building on the physics-informed Gaussian Process (GP) framework introduced in [5], we present an extended framework with adaptive steerer selection. The core idea is to represent machine variability through a stochastic ensemble of accelerator lattices (generated with MAD-X [6]; alternative simulation codes are equally applicable). The physical correlations are encoded into the GP kernel empirically, following the approach of [7]. The resulting model provides probabilistic orbit predictions at arbitrary positions along the ring or beamline and supports parameter inference via Bayesian optimization [8]. The active-learning component is extended with a physics-informed steerer selection strategy: a phase advance prefilter eliminates geometrically blind steerers, and a submodular batch-selection criterion greedily assembles complementary measurements to reduce posterior uncertainty efficiently.

EXTENDED FRAMEWORK

Adaptive Steerer Selection

In [5], the active learning loop selected steerer settings based on a scalar acquisition function applied to the full GP posterior. A limitation of that approach is that it does not account for the betatron phase advance between a steerer and the quadrupole whose misalignment is to be inferred. The phase advance prefilter is based on the closed-orbit response of a steerer s at quadrupole location p. For a ring, the differential closed orbit at position p due to a kick θ at steerer s is

$$\Delta x_{sp} = \frac{\sqrt{\beta_s \beta_p}}{2\sin(\pi q)} \cos\left(\Delta\phi_{sp} - \pi q\right) \cdot \theta, \qquad (1)$$

where β_s and β_p are the Courant–Snyder β-functions at the steerer and quadrupole locations, respectively. $\Delta\phi_{sp}$ is the betatron phase advance from steerer to quadrupole, and q is the fractional tune [9]. The sensitivity of this response to the phase advance is entirely encoded in the cosine factor, and a steerer is therefore discarded if

$$\left|\cos\left(\Delta\phi_{sp} - \pi q\right)\right| < \varepsilon. \qquad (2)$$

We choose $\varepsilon = 0.1$. This criterion is dimensionless and lattice-independent: a steerer is retained only if its betatron phase advance places it in a regime where it can produce

* victoria.isensee@tu-darmstadt.de

at least 10 % of the maximum attainable orbit response at the target quadrupole. Such phase-blind steerers convey no information about the target quadrupole regardless of kick amplitude and are therefore eliminated by this cheap analytic prefilter before any simulation is invoked.

For a single-pass beamline such as the SFRS, Eq. (1) does not apply. The orbit response at position p due to an upstream kick at s is instead given by the transfer matrix element

$$\Delta x_{sp} = \sqrt{\beta_s \beta_p}\ \sin\big(\Delta\phi_{sp}\big) \cdot \theta, \tag{3}$$

where $\Delta\phi_{sp} > 0$ requires the steerer to be located upstream of the target quadrupole [9]; downstream steerers produce no response and are excluded trivially. The corresponding prefilter condition becomes

$$\big|\sin\big(\Delta\phi_{sp}\big)\big| < \varepsilon. \tag{4}$$

Among the remaining candidates, a batch of k complementary steerers is assembled greedily by maximizing the mutual information criterion between the new observations and the unknown misalignment vector [10, 11]. The full adaptive loop proceeds as follows:

1. **Observe** the orbit under the current θ-matrix (one row per (s, θ) pair).
2. **Fit GP**: Nelder–Mead maximizes the log marginal likelihood to obtain posterior estimates $\{\mu_j, \sigma_j\}$.
3. **Rank quads** by uncertainty σ_j.
4. **Phase prefilter**: drop steerers with $|\cos(\Delta\phi - \pi q)| < \varepsilon$ using posterior-mean Twiss parameters.
5. **Submodular batch**: greedily pick steerers maximizing mutual information.
6. **Theta selection**: GP surrogate finds kick amplitude minimizing $\sigma_{\mathrm{orbit}}^{\mathrm{rms}}$.

We apply the framework, extended with adaptive steerer selection, to the full horizontal closed orbit of SIS18.

SIS18 Application

We apply the extended framework to the full horizontal closed orbit of SIS18, modeled via an ensemble of MAD-X lattices with Gaussian-distributed horizontal misalignments in six selected quadrupoles. The framework constructs a physics-informed GP machine model with uncertainty quantification: given a limited number of steerer–kick observations, the model infers the posterior distribution over quadrupole misalignments and provides uncertainty-quantified orbit predictions at any longitudinal position s, including locations without BPMs. Two operational objectives are considered: (a) minimizing the epistemic uncertainty of the orbit prediction at the extraction septum and (b) minimizing the integrated orbit uncertainty across the full ring. The evolution of the predicted orbit uncertainty σ_{orbit} at the extraction septum over iterations is shown in Fig. 1. Starting

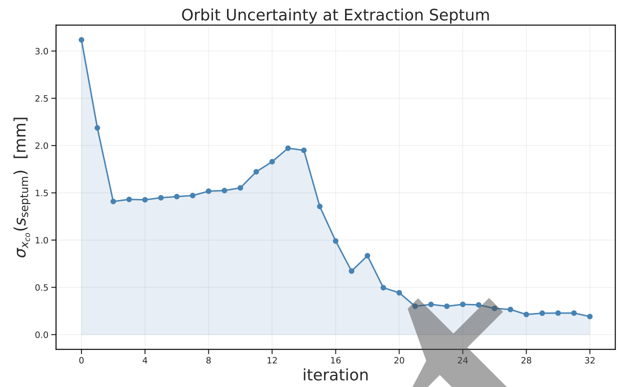

Figure 1: Evolution of the closed-orbit uncertainty σ_{orbit} at the extraction septum. At each iteration, a new batch of steerer–kick observations is added to the training set. A total reduction of 93.9 % is achieved within 32 iterations.

from an initial value, the active learning loop reduces the septum prediction uncertainty by 93.9 % within 32 iterations. Each iteration adds a batch of three steerer–kick pairs; the algorithm initially draws exclusively from three specific kickers and introduces an additional one from iteration 9 onward as the phase-informed selection adapts to the evolving posterior. The inferred misalignment estimates are compared against ground truth (GT) values in Fig. 2. The posterior means converge toward the GT values and the $\pm 1\sigma$ uncertainty bands narrow progressively across all six quadrupoles, confirming that the Bayesian Optimization loop successfully identifies the machine state within the modeled parameter space. The resulting uncertainty-quantified machine model enables confident orbit prediction at the extraction septum, without requiring a full orbit response matrix measurement. Switching to objective (b), minimizing the integrated orbit uncertainty across the full ring, requires only a modification of the acquisition objective in the adaptive steerer selection step: instead of evaluating the predicted variance at the septum location, the objective is replaced by the ring-averaged posterior variance. No other changes to the framework are necessary. With this substitution, the active learning loop yields analogous convergence behavior, achieving a comparable reduction in orbit uncertainty and a similar progressive narrowing of the misalignment posterior bands across all six quadrupoles, demonstrating the flexibility and modularity of the proposed approach.

SFRS: Future Work

The SFRS at FAIR presents a qualitatively different operating environment compared to SIS18. As a large-acceptance fragment separator, it comprises multiple independently powered dipole and quadrupole sections with relatively sparse BPM coverage and a demanding requirement on optical reproducibility across different fragment settings. The combination of sparse instrumentation and complex, significantly varying optics makes conventional ORM-based methods difficult to apply reliably. The physics-informed GP framework is well suited to address these challenges. The ensemble-

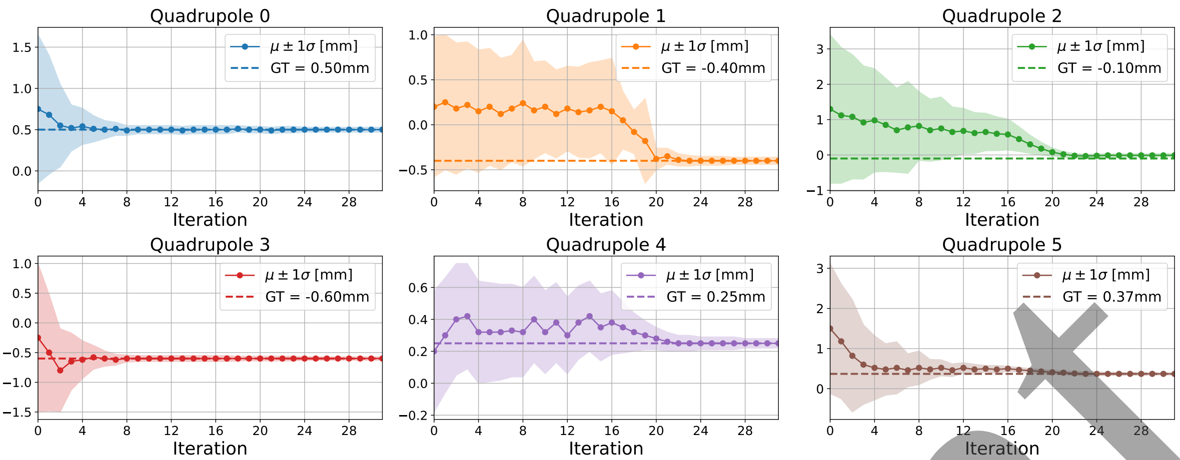

Figure 2: Inferred misalignment mean μ_j and $\pm 1\sigma_j$ band per quadrupole over 32 active learning iterations. Dashed horizontal lines indicate the ground truth (GT) values.

based kernel construction accommodates arbitrary lattice topologies, including beamlines without periodicity, provided that a sufficiently diverse ensemble of simulated lattice realizations can be generated. The uncertainty-enabled prediction between BPMs is especially valuable where detector placement is constrained by hardware or radiation considerations, and the active learning strategy reduces the number of required beam pulses – an important advantage during commissioning when beam time is scarce.

As future work, we will construct a GP model for a representative section of the SFRS using the same MAD-X ensemble approach applied to SIS18. Key open questions include the appropriate choice of free parameters (e.g. dipole field errors and quadrupole misalignments), the scaling of the submodular batch selection to larger numbers of beamline elements, and the impact of beam optics that vary significantly with the magnetic rigidity setting. A detailed simulation study with quantitative uncertainty estimates is planned for the near term, building directly on the validated methodology presented here.

DISCUSSION AND CONCLUSION

The active learning loop reduces the orbit prediction uncertainty at the extraction septum from an initial value of approximately 3.1 mm to 0.19 mm within 32 iterations, corresponding to a total reduction of 93.9 %, as seen in Fig. 1. The convergence is not monotonic: a secondary rise in uncertainty around iterations 12–14 reflects the algorithm selecting steerer-kick pairs that reduce ring-wide posterior uncertainty but are less constraining at the septum specifically, consistent with the scalar acquisition target being the septum uncertainty rather than a global objective. The subsequent drop after iteration 14 coincides with the introduction of an additional kicker from iteration 9 onward, whose phase advance provides stronger sensitivity at the septum location. The misalignment inference converges to the ground truth with qualitatively different rates across the six quadrupoles. Quadrupoles 0 and 3 reach tight agreement early, while quadrupoles 1 and 2 exhibit slower convergence and retain wider residual σ bands at iteration 32. Quadrupoles 4 and 5 converge to posterior means that remain slightly offset from the ground truth, which may reflect a partial degeneracy in the orbit response between quadrupoles at similar betatron phase positions; distinguishing such contributions might be possible with more diverse steerer angles. The phase prefilter (Eqs. (2) and (4)) is dimensionless and lattice-independent by construction. This makes it directly portable to other machines without machine-specific tuning. The submodular batch selection adds negligible computational overhead relative to the GP fitting step, so runtime is dominated by the MAD-X ensemble evaluations. The SFRS is a natural next application: its single-pass topology is directly handled by the beamline prefilter, and the uncertainty-enabled orbit prediction between BPMs addresses a concrete operational need given its sparse instrumentation. A dedicated simulation study using the same MAD-X ensemble methodology is planned as the immediate next step, with open questions including the appropriate free-parameter set and the scaling of the batch selection to a larger number of beamline elements. Physics-informed Gaussian Processes provide a statistical framework to align accelerator models with real machine behavior. Uncertainty quantification thereby enables systematic identification of free model parameters with full probabilistic confidence estimates. It opens a path toward a digital twin of the SFRS: a model that continuously updates as new measurements arrive, bridging simulation and reality.

ACKNOWLEDGMENTS

V. Isensee thanks S. Appel and D. Kallendorf.

REFERENCES

[1] J. Safranek, “Experimental determination of storage ring optics using orbit response measurements”, *Nucl.Instrum. Methods Phys. Res. A*, vol. 388, no. 1–2, pp. 27–36, Mar. 1997. doi:10.1016/S0168-9002(97)00309-4

[2] S. Reimann *et al.*, “FAIR commissioning - Towards first science”, in *Proc. IPAC’25*, Taipei, Taiwan, Jun. 2025, pp. 2495–2498. doi:10.18429/JACoW-IPAC2025-THBN2

[3] S. Mirza, R. Singh, P. Forck, and H. Klingbeil, “Closed orbit correction at synchrotrons for symmetric and near-symmetric lattices”, *Phys. Rev. Accel. Beams*, vol. 22, no. 7, Jul. 2019. doi:10.1103/PhysRevAccelBeams.22.072804

[4] H. Geissel *et al.*, “The Super-FRS project at GSI”, *Nucl. Instrum. Methods Phys. Res. B*, vol. 204, 2003. doi:10.1016/S0168-583X(02)01893-1

[5] V. Isensee, A. Oeftiger, and O. Boine-Frankenheim, “Uncertainty-Quantified Machine Model Construction Using Physics-Informed Gaussian Processes and Bayesian Optimization”, in *Proc. IPAC’25*, Taipei, Taiwan, Jun. 2025, pp. 2718–2721. doi:10.18429/JACoW-IPAC2025-THPM017

[6] MAD-X project, https://cern.ch/madx

[7] X. Yang, G. D. Tartakovsky, and A. M. Tartakovsky, “Physics-Information-Aided Kriging: Constructing Covariance Functions using Stochastic Simulation Models”, Sep. 2018. doi:10.48550/arXiv.1809.03461

[8] C. K. I. Williams and C. E. Rasmussen, *Gaussian Processes for Machine Learning*, Cambridge, MA, USA: MIT Press, 2006.

[9] H. Wiedemann, *Particle Accelerator Physics*, 4th ed., Cham, Switzerland: Springer, 2015. doi:10.1007/978-3-319-18317-6

[10] A. Krause, A. Singh, and C. Guestrin, “Near-optimal sensor placements in Gaussian processes: Theory, efficient algorithms and empirical studies”, *J. Mach. Learn. Res.*, vol. 9, pp. 235–284, Feb. 2008. https://jmlr.org/papers/v9/krause08a.html

[11] D. J. C. MacKay, “Information-based objective functions for active data selection”, *Neural Comput.*, vol. 4, no. 4, pp. 590–604, Jul. 1992. doi:10.1162/neco.1992.4.4.590