

AUTOMATION OF GSI KEY BEAM MANIPULATIONS WITH AI METHODS

S. Appel*, M. Bajzek, E. Kazantseva, P. Madysa, S. Pietri, S. Sorge, H. Weick

GSI Helmholtzzentrum für Schwerionenforschung, Darmstadt, Germany

H. Alsmeier, R. Findeisen, S. Hirt, E. Lenz, M. Pfefferkorn

TU Darmstadt, CCPS, Darmstadt, Germany

O. Boine-Frankenheim, D. Kallendorf, O. Kazinova, C. Reinwald

TU Darmstadt, TEMF, Darmstadt, Germany

B. Halilovic, S. Hirlaender

Paris Lodron Universität Salzburg, Salzburg, Austria

Abstract

We present the Geoff framework for automated accelerator tuning, demonstrated in real-world experiments at GSI. Using classical optimizers like BOBYQA, Geoff enables fast deployment, control room integration, and efficient beam optimization, reducing SIS18 injection losses from 45 % to 15 % and speeding up FRS setup. This work also reports the first application of multi-objective and multi-fidelity Bayesian optimization to SIS18 injection tuning. Complementary simulation studies employ model predictive control via model-based reinforcement learning for fast, constraint-aware tuning. These model-based methods outperform classical optimizers by guiding experiments with probabilistic surrogate and dynamic models.

Geoff's modular design supports easy switching between algorithms and integration with modern ML tools, bridging accelerator operations and data-driven optimization.

GEOFF AT GSI/FAIR

The Facility for Antiproton and Ion Research (FAIR) is going to be an international center of heavy-ion accelerators, that will drive the forefront of heavy-ion and antimatter research [1]. The complexity of FAIR requires a high level of automation for future operation [2]. One part of this automation effort is to provide a framework that allows both machine experts and operators to solve concrete optimization problems, and to make these solutions reusable in an operational context. We call this project the “Generic Optimisation Frontend and Framework”, or *Geoff* for short [3].

Geoff is based on the Python programming language, which is widely used in scientific research, has a vibrant ecosystem of machine learning (ML) algorithms, and is perceived as very beginner-friendly. Both language and framework have proven themselves flexible enough to be quickly adapted to new problems.

Geoff is used extensively at CERN—from linacs to SPS and ISOLDE—and is being deployed at GSI/FAIR [4]. It is usually embedded in a GUI application, but can also be used via command-line scripting to shorten the edit-check-run cycle between code changes and test runs. *Geoff* standardizes interfaces for optimization tasks [5] and provides adapters for




Figure 1: Model of *Geoff* and its components. The optimizer interface includes BOBYQA, Bayesian optimization (BO), and reinforcement learning (RL).

various third-party packages, e.g., *SciPy*, *Stable Baselines 3*, and *Scikit-Optimize*. Optimization problems, implemented as plugin packages, can scale to arbitrary complexity and depend on any Python package. Custom figures can be generated and updated continuously to monitor the algorithm's progress. It can use any control system and communicate with external simulation tools, as long as Python bindings exist (see Fig. 1).

MULTI-TURN INJECTION

The SIS18 serves as a booster for SIS100. Multi-turn injection (MTI) into it is a key bottleneck for achieving FAIR intensity goals. Beam loss-induced vacuum degradation limits intermediate-charge-state beam intensity, so injection losses—occurring at the septum or acceptance—must be minimized. Four time-varying bumper magnets create a closed-orbit bump, guiding each incoming beam into the available horizontal phase space near the previous injection [6, 7].

In November 2023 and May 2024, optimization runs at GSI evaluated the use of *Geoff*. The goal was to minimize SIS18 injection losses by adjusting five injection parameters—orbit bump amplitude and reduction, two septum steerers, and a timing offset—and four TK steerers. To avoid bias from favorable initial states, parameters were randomized before each run. Loss was estimated from the difference between ideal and measured SIS18 currents (Fig. 2). BOBYQA was chosen as the optimizer, with an initialization phase of $N = 19$ evaluations for nine parameters, as

* s.appel@gsi.de




Figure 2: Nominal (black) and ideal (red) beam current during injection. The dashed line marks end of incoming beam-lets. The reward is based on their difference.




Figure 3: SIS18 MTI online steering with BOBYQA. Gray area: initialization; final point: extra evaluation at optimum.

shown in Fig. 3. Each iteration used the median of three measurements to reduce fluctuations. In November 2023, the beam loss decreased from 45 % after manual tuning to 15 % after full optimization. In May 2024, the loss dropped from 30 % manually to 15 % after optimization. Each run took 15–20 minutes, depending on external factors [8].

FRAGMENT SEPARATOR OPTIMIZATION

At GSI, setting up the Fragment Separator (FRS) for experiments currently requires manual adjustments. This takes up to one third of experimental time. Optical mismatches between FRS and downstream ESR can make some experiments infeasible, as manual tuning is insufficient to reach the required beam intensity due to limited diagnostics and many influencing parameters. The FAIR Super-FRS, with four times more magnets than the FRS, will increase complexity and setup time, making automated tuning essential. Initial automation uses *Geoff* for numerical optimization.

The first application aligned the beam at a target using four correctors, measured via profile grids and Gaussian fits. In a June 2024 run [9], convergence occurred in 20 iterations




Figure 4: Online focusing with BOBYQA: Converges after ~15 iterations, aligning U^{91+} beam vertically via quadrupoles.

(about 15 minutes), though one corrector was incidentally near the optimum. Measurement speed is limited by the slow-extraction cycle and grid readout [8].

Beam focusing at the main dispersive focal plane (S2) was then automated next (Fig. 4). A tantalum target produced multiple uranium charge states; the U^{91+} beam was focused by adjusting quadrupoles. Particle trajectories were reconstructed using a TPC detector and classified via *scikit-learn*, then converted to a single figure of merit describing beam spot focusing. The algorithm found the optimum after 15 of 25 iterations [8, 9].

Ongoing developments include computing transfer maps from particle-by-particle phase-space data using ML and symplectic map theory, along with improved goal functions for sequential quadrupole and sextupole tuning [10].

ADVANCED TECHNIQUES

As accelerator facilities grow in complexity and demand higher beam quality and efficiency, more advanced optimization methods are needed. Classical algorithms like BOBYQA are robust and efficient for many tasks but can struggle with high-dimensional or highly non-linear parameter spaces. Techniques such as Bayesian optimization (BO) and data-driven Model Predictive Control (MPC) provide frameworks that model uncertainties, learn from past data, and adapt dynamically.

Bayesian Optimization

Bayesian optimization (BO) is a sample-efficient method for optimizing expensive black-box functions [4, 11]. It builds a probabilistic surrogate (typically a Gaussian process) and uses an acquisition function to balance exploration and exploitation. This makes it ideal for accelerator tuning, where evaluations are costly or beam losses must be minimized [6, 7, 12].

For SIS18 MTI, BO was applied using the XSuite simulation model and the input parameters in Table 1. The model was carefully validated against experiments [6, 7, 12]. Both single-objective (minimizing injection loss) and multi-objective optimization (trading off injection loss and multi-

Table 1: Injection Simulation Input Parameters

Input parameter	Description
x_{orb}	Orbit bump amplitude at IP ^a
x'_{orb}	Orbit bump angle at IP ^a
x_{dist}	Mean injection position at IP ^a
x'_{dist}	Mean injection angle at IP ^a
M	Injected distribution mismatch [13]
Δ	Orbit bump decay per turn
Q_x	Horizontal tune


^a Injection point

Figure 5: MOBO results: blue dots for all measurements and orange circles for Pareto front. BOBYQA results: red for SIS18 parameters, green for full optimization.

plication factor) were considered. Multi-Objective Bayesian optimization (MOBO) approximates the Pareto front by selecting candidates that improve the expected hypervolume, enabling better trade-offs than single-objective methods [4, 14]. Figure 5 shows that MOBO identifies more solution sets and achieves better trade-offs than BOBYQA. The expected increase in losses with higher stored turns (multiplication factor) is visible, consistent with simulations using genetic algorithms [14].

Convergence studies tested various kernels and acquisition strategies via PyTorch, GPyTorch, and BoTorch [15–17]. Portfolio strategies GP-Hedge and GP-Exp3 [18, 19] outperformed single acquisition functions, with GP-Exp3 preferred for computational efficiency. Analysis of 180 optimized inputs (Fig. 6) revealed that some parameters (e.g. x'_{dist}) have little effect, allowing a reduced input space Ω_s with fixed x'_{orb} and x'_{dist} to further improve optimization performance [20].

Multi-Fidelity Bayesian Optimization

Beam injection into SIS18 is currently tuned with black-box optimization, where high losses and many iterations occur, especially early in tuning. This is dangerous in case of uranium beams, where minimizing loss-induced radiation is essential. Effective strategies must deliver high performance with few costly machine interactions. Multi-fidelity BO ad-





Figure 6: Distribution of 180 optimized input values of larger than 0.9 injection efficiency.

Figure 7: Schematic of multi-fidelity BO: low-fidelity simulations ($s < 1$) and high-fidelity experiments ($s = 1$) combined to efficiently optimize expensive functions.

dresses this by treating *fidelity* as an input to the surrogate model [21]. Low-fidelity simulations provide approximate but cheap information, while high-fidelity machine experiments are accurate but costly. The algorithm adaptively balances both, accelerating tuning and reducing machine wear without compromising quality (Fig. 7).

Remote Geoff extends *Geoff* by offloading computationally demanding algorithms—such as multi-fidelity BO—to external resources like HPC clusters, while securely interacting with the accelerator [22].

Using the MTI simulation framework, initial experiments showed substantial reductions in required machine interactions. Figure 8 shows fourteen low-fidelity simulations in blue, eleven measurements in red, and the optimized result in orange. Due to time constraints, iterations were capped at 25 and optimization limited to the five SIS parameters. Despite this, beam losses were reduced from about 45 % (manual tuning) to 35 %, enabled by sufficiently accurate low-fidelity guidance.

Data-Driven Model Predictive Control

Data-driven MPC combines model-based reinforcement learning (MBRL) with control theory. MBRL learns a probabilistic surrogate, while MPC rolls it forward over a finite horizon to minimize beam-loss cost under limits. This yields a controller that converges rapidly, requires far fewer machine interactions than brute-force tuning, and preserves safety [23–27]. Unlike greedy optimization, which treats each step independently, MPC builds a reusable dynamics




Figure 8: Multi-fidelity MTI optimization: interactions reduced from 25 to 11 using high-fidelity (red) and low-cost simulations (blue).




Figure 9: Episode resets appear as jumps, and shaded regions show prediction uncertainty, which decreases with more data.

model, plans ahead to respect constraints, and reaches target performance with fewer trials.

MPC was applied in simulation to improve MTI [23]. Results are shown in Fig. 9. The top panel depicts the system state—beam loss, septum loss, and two current integrals (cf. Fig. 2); the middle shows the evolution of control actions from Table 1 (excluding tune); and the bottom tracks the reward, defined as the RMS deviation between ideal and actual beam current. As training progresses, rewards improve and model uncertainty decreases. Episodes are reinitialized with randomized starting conditions and end early if rewards exceed a threshold or action limits are violated; otherwise, after 25 steps.

OUTLOOK

This work demonstrates the flexibility and practical value of the *Geoff* framework for automated accelerator tuning.

Starting with classical algorithms like BOBYQA, *Geoff* enables rapid deployment, integration into control environments, and efficient optimization of beam parameters.

Experimental results at GSI show substantial gains: SIS18 injection losses dropped from 45 % to 15 %, and FRS setup times were significantly reduced. Losses during slow extraction were also mitigated through automated optimization, supporting more robust and efficient operation. This work also introduces multi-objective and multi-fidelity Bayesian optimization for SIS18, alongside model predictive control and constraint-aware tuning. These model-based approaches outperform classical methods by learning probabilistic or dynamic models that guide experiments and reconcile competing objectives.

Geoff's modular design supports easy switching between optimization strategies and integration with external analysis tools and modern ML libraries, bridging the gap between accelerator operations with data-driven methods.

ACKNOWLEDGEMENTS

The EURO-LABS project has received funding from the European Union's Horizon Europe Research and Innovation programme under Grant Agreement no. 101 057 511.

REFERENCES

- [1] H. Gutbrod *et al.*, “Executive summary”, in *FAIR baseline technical report*. Darmstadt: GSI, 2006. <https://repository.gsi.de/record/54062>
- [2] S. Reimann, M. Sapinski, P. Schütt, and M. Vossberg, “Building an operation team for FAIR nearly from scratch”, in *Proc. WAO'16*, Sep. 2016. <https://www.researchgate.net/publication/316324398>
- [3] P. Madysa, S. Appel, V. Kain, and M. Schenk, “Geoff: the generic optimization framework & frontend for particle accelerator controls”, *SoftwareX*, vol. 32, p. 102335, 2025. [doi:10.1016/j.softx.2025.102335](https://doi.org/10.1016/j.softx.2025.102335)
- [4] R. Roussel *et al.*, “Bayesian optimization algorithms for accelerator physics”, *Phys. Rev. Accel. Beams*, vol. 27, no. 8, p. 084801, 2024. [doi:10.1103/PhysRevAccelBeams.27.084801](https://doi.org/10.1103/PhysRevAccelBeams.27.084801)
- [5] P. Madysa, COI — common optimization interfaces, <https://cernml-coi.docs.cern.ch/>
- [6] S. Appel and O. Boine-Frankenheim, “Multi-turn injection into a heavy-ion synchrotron in the presence of space charge”, *arXiv*, 2014. [doi:10.48550/arXiv.1403.5972](https://doi.org/10.48550/arXiv.1403.5972)
- [7] S. Appel *et al.*, “Injection optimization through generation of flat ion beams”, *Nucl. Instrum. Methods Phys. Res. A*, vol. 866, pp. 36–39, 2017. [doi:10.1016/j.nima.2017.05.041](https://doi.org/10.1016/j.nima.2017.05.041)
- [8] S. Appel and et al, “Automated optimization of accelerator settings at GSI”, in *Proc. IPAC'24*, Nashville, TN, USA, May 2024, pp. 882–885. [doi:10.18429/JACoW-IPAC2024-MOPS68](https://doi.org/10.18429/JACoW-IPAC2024-MOPS68)
- [9] S. Appel, N. Madysa, and S. Dobosz, “D5.4 EURO-LABS Toolkit for optimizing beam delivery (deployed in two facilities)”, *Zenodo*, Aug. 2024. [doi:10.5281/zenodo.13809056](https://doi.org/10.5281/zenodo.13809056)

- [10] D. Kallendorf, “Data driven optimization of the ion-optical settings of the FRS”, MA thesis, TU Darmstadt, Darmstadt, Germany, 2025.
- [11] E. Brochu, V. M. Cora, and N. de Freitas, “A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning”, *arXiv*, 2010.
[doi:10.48550/arXiv.1012.2599](https://doi.org/10.48550/arXiv.1012.2599)
- [12] G. Iadarola *et al.*, “Xsuite: An Integrated Beam Physics Simulation Framework”, in *Proc. HB’23*, CERN, Geneva, Switzerland, Oct. 2023, pp. 73–80.
[doi:10.18429/JACoW-HB2023-TUA2I1](https://doi.org/10.18429/JACoW-HB2023-TUA2I1)
- [13] C. R. Prior and G. H. Rees, “Multiturn injection and lattice design for HIDIF”, *Nucl. Instrum. Methods Phys. Res. A*, vol. 415, no. 1, pp. 357–362, 1998.
[doi:10.1016/S0168-9002\(98\)00406-9](https://doi.org/10.1016/S0168-9002(98)00406-9)
- [14] S. Appel, O. Boine-Frankenheim, and F. Petrov, “Injection optimization in a heavy-ion synchrotron using genetic algorithms”, *Nucl. Instrum. Methods Phys. Res. A*, vol. 852, pp. 73–79, 2017. [doi:10.1016/j.nima.2016.11.069](https://doi.org/10.1016/j.nima.2016.11.069)
- [15] A. Paszke *et al.*, “Pytorch: an imperative style, high-performance deep learning library”, in *Proc. NeurIPS’19*, Vancouver, Canada, Dec. 2019. https://proceedings.neurips.cc/paper_files/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf
- [16] J. Gardner, G. Pleiss, K. Q. Weinberger, D. Bindel, and A. G. Wilson, “GPpyTorch: blackbox matrix-matrix Gaussian process inference with GPU acceleration”, in *Proc. NeurIPS’18*, Montréal, Canada, Dec. 2018.
- [17] M. Balandat and et al., “BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization”, in *Proc. NeurIPS’20*, Virtual, Dec. 2020–Dec. 2010.
[doi:10.48550/arXiv.1910.06403](https://doi.org/10.48550/arXiv.1910.06403)
- [18] M. Hoffman, E. Brochu, N. De Freitas, *et al.*, “Portfolio allocation for Bayesian optimization”, in *Proc. UAI’11*, Barcelona Spain, Jul. 2011, pp. 327–336.
- [19] P. Auer, N. Cesa-Bianchi, Y. Freund, and RE. Schapire, “Gambling in a rigged casino: the adversarial multi-armed bandit problem”, in *Proceedings of IEEE 36th Annual Foundations of Computer Science*, Milwaukee, Wisconsin, USA, Oct. 1995, pp. 322–331. [doi:10.1109/SFCS.1995.492488](https://doi.org/10.1109/SFCS.1995.492488)
- [20] C. Reinwald, “Automatization of the multi-turn injection into the SIS18 synchrotron using Bayesian optimization”, MA thesis, Technische Universität Darmstadt, Department of Electrical Engineering and Information Technology, Darmstadt, Germany, Oct. 2024.
- [21] J. Wu, S. Toscano-Palmerin, P. I. Frazier, and A. G. Wilson, “Practical Multi-fidelity Bayesian Optimization for Hyperparameter Tuning”, *arXiv*, 2019.
[doi:10.48550/arXiv.1903.04703](https://doi.org/10.48550/arXiv.1903.04703)
- [22] P. Madysa, S. Appel, L. Dingeldein, J. Fitzek, V. Kain, and M. Schenk, “Geoff developments in 2025”, presented at ICALEPCS’25, Chicago, IL, USA, Sep. 2025, paper WEPD076, unpublished.
- [23] S. Hirilaender, S. Appel, and N. Madysa, “Data-driven model predictive control for automated optimization of injection into the SIS18 synchrotron”, in *Proc. IPAC’24*, Nashville, TN, USA, May 2024, pp. 1800–1803.
[doi:10.18429/JACoW-IPAC2024-TUPS59](https://doi.org/10.18429/JACoW-IPAC2024-TUPS59)
- [24] S. Hirilaender, L. Lamminger, G. Zevi Della Porta, and V. Kain, “Ultra fast reinforcement learning demonstrated at CERN AWAKE”, in *Proc. IPAC’23*, Venice, Italy, May 2023, pp. 4510–4513.
[doi:10.18429/JACoW-IPAC2023-THPL038](https://doi.org/10.18429/JACoW-IPAC2023-THPL038)
- [25] S. Hirilaender and N. Bruchon, “Model-free and Bayesian ensembling model-based deep reinforcement learning for particle accelerator control demonstrated on the FERMI FEL”, *arXiv*, 2022. [doi:10.48550/arXiv.2012.09737](https://doi.org/10.48550/arXiv.2012.09737)
- [26] S. Hirilaender *et al.*, “Towards few-shot reinforcement learning in particle accelerator control”, in *Proc. IPAC’24*, Nashville, TN, USA, May 2024, pp. 1804–1807.
[doi:10.18429/JACoW-IPAC2024-TUPS60](https://doi.org/10.18429/JACoW-IPAC2024-TUPS60)
- [27] V. Kain *et al.*, “Test of machine learning at the CERN LINAC4”, in *Proc. HB’21*, Batavia, IL, USA, Oct. 2021, pp. 181–185. [doi:10.18429/JACoW-HB2021-TUEC4](https://doi.org/10.18429/JACoW-HB2021-TUEC4)