

ONLINE BEAMLINE TUNING USING BAYESIAN OPTIMIZATION AT THE HIT MEDICAL SYNCHROTRON

M. Hun^{*1}, C. Steffen^{1,2,3}, E. Feldmeier¹

¹Heidelberg Ion Beam Therapy Centre, Heidelberg, Germany

²Interdisciplinary Center for Scientific Computing, Heidelberg, Germany

³ZUSE School ELIZA, Darmstadt, Germany

Abstract

The Heidelberg Ion Beam Therapy Centre (HIT) plans to implement multi-energy extraction, which allows faster cancer treatment by re-accelerating unextracted beam in the synchrotron. Thus an increase of the number of particles injected into the synchrotron would directly benefit patient treatment time and consequently also the number of patients treated. However, optimization of beam line segments by operators is often time consuming and it is difficult to align with a global optimum for the entire beamline. To address this problem we investigate the implementation of a machine learning algorithm based on Bayesian optimization using gaussian process models for online accelerator tuning. We specifically focus on sample-efficient transmission optimization for larger beamlines in medical accelerators and the effects of input space and objective selection.

INTRODUCTION

Particle therapy uses ions to irradiate a tumor volume by utilizing a Bragg peak, minimizing damage to surrounding healthy tissue. The position of this Bragg peak inside the tissue can be controlled by adjusting the energy of the particles used for irradiation. In conventional synchrotron-based delivery, switching between energies requires dumping of all remaining particles and restarting the full injection and acceleration cycle, introducing significant dead time between energy layers and driving up power costs. Multi-energy extraction reduces this inefficiency by re-accelerating unused particles within the same fill, directly reducing treatment times [1]. A higher available particle count also opens the door to irradiation with a higher dose rate, a regime that has shown promise in improving treatment outcomes.

The beam transport from the ion source through a linear accelerator and to the synchrotron, consists of many individually tunable elements: focusing quadrupole magnets, dipole bending magnets, and beam-steering corrector (steerer) magnets. Optimizing their collective settings to maximize beam transmission is a high-dimensional problem that is difficult to solve by hand or without proper models and simulations. Traditional operator-driven tuning relies on expert knowledge and time consuming iterative manual adjustment. Automated tuning algorithms have therefore attracted significant interest in the accelerator physics community [2]. Bayesian optimization (BO) is a powerful framework for the data-efficient optimization of expensive black-box functions. By

maintaining a probabilistic surrogate model of the objective and using an acquisition function to guide the search, BO can optimize the underlying function within fewer functions evaluations. This property is especially valuable in an online accelerator setting, where each evaluation corresponds to a live beam measurement that consumes scarce beam time and power. In this work we investigate the application of Bayesian optimization with Gaussian process surrogate models to the online transmission optimization of a beamline for a medical synchrotron. We examine the scalability of the approach with respect to the dimensionality of the input space and propose a structured decomposition of the full beamline optimization problem that reduces effective dimensionality and improves sample efficiency.

METHODS

Gaussian Process Surrogate Models

A Gaussian process (GP) is a probabilistic model that defines a distribution over functions [3]. Formally, a GP is fully specified by a mean function $m(x)$ and a covariance (kernel) function $k(x, x')$, such that any finite collection of function values follows a multivariate Gaussian distribution. Given a set of noisy observations $D = (x_i, y_i)$, the GP posterior provides a predictive distribution where the predictive mean acts as a point estimate of the objective, while the predictive variance quantifies model uncertainty. In practice, the radial basis function (RBF) kernel,

$$k(x, x') = e^{(-\|x-x'\|^2/2l^2)}, \quad (1)$$

is widely used, where l is the length scale. Kernel hyperparameters such as l are typically learned by maximizing the marginal log likelihood of the observed data. The marginal log likelihood serves as a measure for the agreement between data and model.

Acquisition Functions

The acquisition function $a(x)$ uses the GP posterior to score candidate points according to predefined algorithm. Optimizing $a(x)$ yields the next query point x^* which in turn is then measured and used to update the GP. The acquisition function used in all described experiments is Upper Confidence Bound (UCB). UCB is derived from the principle of optimism under uncertainty [4]. It scores each candidate point by:

$$a_{UCB}(x) = \mu(x) + \kappa \cdot \sigma(x), \quad (2)$$

* Marcel.Hun@med.uni-heidelberg.de

where $\mu(x)$ and $\sigma(x)$ are the GP posterior mean and standard deviation, respectively, and $\kappa \geq 0$ is a tunable exploration parameter. A larger κ encourages more exploratory behavior by up-weighting uncertain regions, a smaller κ drives exploitation near the current optimum. In this work, UCB is selected as the primary acquisition function due to its simplicity and fast computation.

Acquisition Function Optimization

Finding the optimum of $\alpha(x)$ is itself an optimization problem that must be solved at each BO iteration. Because $\alpha(x)$ is cheap to evaluate, requiring only GP posterior inference, a naïve Monte Carlo (MC) sampling approach was chosen as the primary method of this work. For MC sampling a large number of candidate points is drawn uniformly from the input space, the acquisition function is evaluated at each, and the point with the highest score is returned. While conceptually simple, MC sampling suffers from poor scaling with input dimensionality. The number of samples required to maintain adequate coverage of the space grows exponentially. In low-dimensional problems (two to four parameters) MC sampling is effective, but as the number of dimensions increases, the approach becomes too expensive for online optimization. To address this limitation, a gradient-based optimization using the Adam (Adaptive Moment Estimation) optimizer was implemented [5]. Because the GP posterior mean and variance are differentiable with respect to the input x (given a differentiable kernel), the gradient of the acquisition function $\Delta_x \alpha(x)$ can be computed analytically and via automatic differentiation. Testing of the Adam-based acquisition optimizer in the online beamline has not yet been possible at the time of writing.

Application to Online Beamline Optimization

In the context of online beamline optimization the input of the GP is some form of setting(s) which controls beamline elements while the observations are measurements taken with beam instrumentation devices. The experiments performed in this work used only the bending strength of quadrupoles as inputs. Measurements of the beam current taken by a transformer coil as the observation of transmission. As shown in Fig. 1 there are two optimization loops, the outer loop which takes the costly beam measurements and the inner loop which uses the surrogate model. After some initial measurements the GP model is used to predict the next point by MC sampling the inner loop and selecting the best result as the next point. The measurement of this next point is then taken by the accelerator and used to update the model. Criteria for the completion of the outer loop could be some fixed measurement budget or stagnation of the transmission improvement for several cycles.

Proposed Method: Structured Decomposition of the Optimization Problem for Large Beamlines

Even with a gradient-based inner optimizer, the dimensionality of the full beamline optimization problem, which




Figure 1: Flow chart of the online optimization process. Green fields are measurements taken with the accelerator, blue ones are computations.

may involve well over 20 independent magnetic elements, remains a challenge for online GP-based surrogate modeling. To address this, two methods are proposed. One is a simple segmentation along the length of the beamline. Wherever there is sufficient beam instrumentation to measure the transmission and position the beamline is divided and optimized separately with only a small overlap between segments to allow for adjustment between segments. The other is a physically motivated decomposition of the optimization problem that separates the roles of different magnet families. The injection beamline contains two functionally distinct device classes: (i) quadrupole magnets, which primarily control beam focusing; (ii) steerer magnets and dipole bending magnets, which adjust the transverse beam position and angle; In an ideal setting where the beam is kept in the center of the beamline both classes should be optimizable independently of each other. Since this will likely not be achievable in a real world setting a multiobjective approach is necessary to find a trade-off between transmission and centration.

To do so we propose training two separate Gaussian process surrogate models with distinct objectives, each operating over a reduced input space corresponding to a subset of device families. The first GP is trained on beam transmission as the objective and takes the settings of quadrupole magnets as inputs. The second GP is trained on a beam centration objective quantifying how well the beam travels along the nominal axis of the beamline and operates over the settings of steerer and dipole magnets.

This decomposition reduces the number of devices that must be co-optimized in a single GP from the full comple-

ment of up to 20 elements to a subset of approximately 6–10 quadrupoles and 6–10 steerers and dipoles.

RESULTS

Experiments were conducted on isolated short segments of the injection beamline, in which only a small number of magnetic elements (two to four quadrupole magnets) were optimized simultaneously. Beam transmission was measured using a current transformer at the end of the optimized segment. The GP surrogate model was initialized with as little as 5 random observations, after which BO iterations were performed for a fixed observation budget.

For two dimensions the Monte Carlo acquisition optimizer with 500,000 randomly sampled candidates performed reliably in two dimensions. As can be seen in Fig. 2 the GP surrogate accurately captures the landscape within few evaluations, and UCB-guided optimization converged to current values comparable to or exceeding the best operator-tuned settings. There is however room for improvement as the GP does not accurately model the height of the optimum. The number of required machine interactions can be kept below fifty even in the four dimensional case which is particularly significant given the limited beam time available outside of clinical operations.




Figure 2: Online Transmission Optimization for a quadrupole triplet. The red dots are the actual accelerator observations taken in the outer loop. The surface plot shows the beam current prediction of the GP surrogate model.

When the optimization was extended to larger beamline sections involving six and eight simultaneously varied devices, the MC-based acquisition optimizer exhibited clear degradation in performance. Even at 5×10^6 candidates, the effective resolution of the acquisition function maximum deteriorated and the optimizer consistently returned suboptimal next query points causing the improvement in beam current to plateau early into the optimization. This was confirmed by later analyzing the six dimensional data using the Adam optimizer which found much better settings closer to the operator tuned values compared to the MC sampling.

Further it was found that the Adam optimizer is also prone to become stuck in local optima and a number of randomly initialized restarts were needed to consistently reach the perceived global optimum. The required number of restarts scales with the dimension of the problem just as the MC sampling, but not exponentially.

Including steerers and dipoles into the transmission optimization proved to be difficult because of the sensitivity even to small deviations, causing the beam current to drop down to the noise level of the measuring device. This led to the proposed method of splitting quadrupoles, and steerers and dipoles into different groups and using a multiobjective optimization approach which could not yet be tested.

CONCLUSION

The results demonstrate that Bayesian optimization with Gaussian process surrogates is a viable and sample-efficient approach to online transmission optimization for short beamline segments. The UCB acquisition function provides an intuitive and principled trade-off between exploration and exploitation, and low-dimensional experiments confirm fast convergence with a limited number of beam measurements. Monte Carlo acquisition optimization, while adequate in low dimensions, fails to scale due to the exponentially growing volume of the input space. The proposed transition to Adam-based gradient optimization of the acquisition function shows better behaviour in an offline setting, is practically straightforward to implement within the existing BO framework, but requires experimental validation in the online setting.

The primary technical challenge remains scalability to the full injection line. The dual-GP decomposition might be able to address the complementary challenge scalability in high dimensions. By separating the roles of focusing elements (quadrupoles) and trajectory elements (steerers and dipoles), and by assigning each family to a distinct GP with a physically meaningful objective, the effective dimensionality of each sub-problem is substantially reduced. Whether this method holds up still needs to be proven in future experiments. Difficulties might include the steering effects of quadrupoles on off-center beams and sensitivity of steerers and dipoles prompting a more iterative approach negatively impacting sample efficiency.

Recent related work has also provided improvements upon the methods used in this work, such as improved acquisition functions or much desired incorporation of prior knowledge and historical data [6].

ACKNOWLEDGEMENTS

Corinna Steffen is supported by the Konrad Zuse School of Excellence in Learning and Intelligent Systems (<https://eliza.school1>) through the DAAD programme Konrad Zuse School of Excellence in Artificial Intelligence, sponsored by the Federal Ministry of Education and Research.

REFERENCES

- [1] M. Galonska, E. Feldmeier, T. Haberer, A. Peters, and C. Schoemers, “Multi-Energy Trial Operation of the HIT Medical Synchrotron: Accelerator Model and Data Supply”, in *Proc. IPAC'17*, Copenhagen, Denmark, May 2017, pp. 4644–4646. [doi:10.18429/JACoW-IPAC2017-THPVA082](https://doi.org/10.18429/JACoW-IPAC2017-THPVA082)
- [2] R. Roussel, A. Hanuka, and A. Edelen, “Multi-objective bayesian optimization for accelerator tuning”, *Phys. Rev. Accel. Beams*, vol. 24, p. 062801, 2021. [doi:10.1103/PhysRevAccelBeams.24.062801](https://doi.org/10.1103/PhysRevAccelBeams.24.062801)
- [3] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006. [http : / / www . GaussianProcess.org/gpml](http://www.GaussianProcess.org/gpml)
- [4] P. Auer, N. Cesa-Bianchi, and P. Fischer, “Finite-time analysis of the multiarmed bandit problem”, *Mach. Learn.*, vol. 47, no. 2–3, pp. 235–256, 2002. [doi:10.1023/A:1013689704352](https://doi.org/10.1023/A:1013689704352)
- [5] D. P. Kingma and J. Ba, “Adam: a method for stochastic optimization”, 2017. [doi:10.48550/arXiv.1412.6980](https://doi.org/10.48550/arXiv.1412.6980)
- [6] T. Boltz *et al.*, “Leveraging prior mean models for faster bayesian optimization of particle accelerators”, 2025. [doi:10.48550/arXiv.2403.03225](https://doi.org/10.48550/arXiv.2403.03225)