

Accelerator tuning via model-coupled optics and Bayesian steering

O. Hassan^{1,2,*} O. Shelbaya^{1,2,*†} P. M. Jung^{1,2}
O. Kester^{1,2} T. Planche^{1,2} and W. Fedorko¹

¹TRIUMF, 4004 Wesbrook Mall, Vancouver, British Columbia, V6T 2A3, Canada

²Department of Physics and Astronomy, University of Victoria,
Victoria, British Columbia, V8W 2Y2, Canada



(Received 19 February 2026; accepted 1 September 2026; published 7 October 2026)

We present an on-line tuning strategy for the isotope separator and accelerator postaccelerator that presets machine optics with a digital twin and then performs Bayesian optimization for steering during operation with beam. The model computes end-to-end configurations in seconds and interfaces with the control system under device bounds, slew-rate limits, and loss interlocks. We report three experimental case studies demonstrating that decoupling optics from steering yields faster and more reliable convergence than a fully Bayesian optics-plus-steering baseline under identical conditions. Across these cases, the number of iterations required to reach high-transmission solutions was substantially reduced, with final average transmissions in the mid- to high-90% range. By removing optics from the steering optimization, the dimensionality of the parameter space is reduced, convergence becomes more predictable, and operational safeguards are easier to enforce.

DOI: [10.1103/1vy9-r7c6](https://doi.org/10.1103/1vy9-r7c6)

I. INTRODUCTION

The design and operation of particle accelerators are shaped by a wide range of scientific, technological, and societal applications. These include nuclear and particle physics studies of nuclear structure and fundamental interactions [1–4], the production of radioisotopes and beams for medical diagnostics and therapy [5,6], materials science investigations using ion implantation and irradiation, accelerator mass spectrometry, and emerging applications in industry and energy research. Across these domains, experiments impose stringent requirements on beam quality, stability, energy control, and species purity. Meeting these demands motivates the development of increasingly sophisticated accelerator modeling, tuning, and control strategies capable of delivering diverse beam species over broad energy ranges with high reliability and reproducibility.

Isotope-separation on-line (ISOL) facilities produce and purify rare-isotope beams (RIBs) *in situ*, then postaccelerate them with dedicated linear accelerators (linacs) or cyclotrons. Several examples include ISOLDE (CERN) [7], ISAC

(TRIUMF) [8], and SPIRAL 1 (GANIL) [9]. Many heavy-ion linacs employ a low-energy front end consisting of an ion source and a radio-frequency quadrupole (RFQ), followed by structures such as interdigital-H mode (IH) cavities, quarter-wave resonators, or other independently phased accelerating cavities. In a separated-function layout [10], these cavities and bunchers are individually tunable, offering considerable operational flexibility at the cost of a larger parameter space.

With this system's complexity comes a growing challenge for operators: optimizing beam quality and transmission while taking into account operational overhead. At facilities such as TRIUMF's isotope separator and accelerator (ISAC), linacs must be retuned frequently to accommodate different ion species, each with unique charge-to-mass ratios and energy requirements. Manual tuning is highly variable, time-consuming, and places a heavy burden on operators.

The commissioning of the Advanced Rare Isotope Laboratory (ARIEL) [11,12] will further amplify the complexity of ISAC's operational scope by tripling its rare-isotope beam availability. ARIEL's new superconducting electron linac [13] is designed to produce high-intensity, neutron-rich isotopes for both research and applications, spanning from medical isotope production to generation of intense THz radiation to fundamental science related to dark matter. Together with TRIUMF's existing 520 MeV proton cyclotron, the new installation will broaden coverage of the nuclear chart, and once ARIEL is fully online, simultaneous delivery of three RIBs and one stable beam will be possible.

*These authors contributed equally to this work.

†Contact author: oshelb@triumf.ca

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/). Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

To address this challenge, we present an approach to accelerator tuning that decouples the optimization of the beam optical elements from that of the corrective steerers. This enables more intuitive and accessible machine control for operators through rapid, semiautonomous tuning for beam delivery. However, several parameters, such as rf phases and ion source settings, remain under manual control; these are the subject of ongoing studies and are outside the scope of this paper. At its core is a digital twin of the linac and beamline that continuously computes beam envelopes and responds to live system parameters. This allows operators to visualize machine states in real time and supports beam optics setpoint computations from first principles. Combined with a previously reported Bayesian optimizer for beam steering [14], this method facilitates efficient parameter calculations and automated beam threading: the process of guiding the beam through the full accelerator lattice using corrective steerers to avoid losses. This stands in contrast to other methods presently being reported, in which the quadrupoles and steerers are both altered using optimization algorithms [15–19]. By reducing operational complexity and decision load, our approach improves tuning efficiency and reduces complexity for beam delivery. The framework has been successfully implemented at certain sections of ISAC and is now being expanded into several facilities across TRIUMF, including the ARIEL beamline network and the cyclotron injection line. Ultimately, this is an adaptation to the growing demand for more precise accelerator tuning to support TRIUMF’s scientific mission in years to come.

A. Relation to prior work

Prior work [20–22] established fast envelope modeling and sequential optimization within subsections of ISAC (e.g., the DTL) but did not integrate automated beam threading, which was separately reported in [14]. Our contribution is an operational pipeline that (i) computes end-to-end configurations online for the entire ISAC chain with model coupled accelerator tuning (MCAT) and (ii) applies steering-only Bayesian optimization for ion steering (BOIS) on the live machine. Compared to fully Bayesian optics-plus-steering strategies (surveyed in [23] and recent applications [24]), this removal reduces dimensionality, improves reliability, and aligns with operational safety constraints. We provide multispecies experimental validation, quantitative convergence and transmission benchmarks, and deployment outcomes.

Rather than simply combining two existing tools, the present work establishes a framework in which online model computation, sequential optics optimization, steering optimization, control-system integration, and operational safety constraints function as a unified beam delivery workflow.




FIG. 1. Overview of the ISAC-I RIB postaccelerator, showing major components.

II. THE ISAC-I FACILITY

The methodology described in this work builds on extensive operational experience at ISAC, which produces rare isotope beams using a proton beam from TRIUMF’s 500 MeV cyclotron to bombard production targets. These are composed of a variety of materials such as UC_x , Ta, and NiO [25]. The resulting isotopes are ionized, extracted, separated [26], and delivered to experiments. Stable beam is provided by the off-line ion source (OLIS) [27]. The postaccelerator, shown in Fig. 1, is divided into low, medium, and high energy sections: electrostatic ion optics (before the RFQ) define the low energy section, while magnetic optics (after the RFQ) are used in the medium and high energy sections.

ISAC was commissioned between the late 1990s and mid-2000s, during which time simulation and tuning relied on contemporary beam dynamics codes. This led to a set of modular, piecewise simulation models tailored to specific sections of the accelerator. The low-energy beam transport and matching sections were primarily designed using the envelope-tracking code TRANSOPTR [28], while high-energy linac and beamline sections were modeled using ray-tracing tools such as PARMTEQ [29,30], LANA [31,32], and TRACEWIN [33]. As a result, parameter computations for the entire accelerator have historically relied on multiple independent simulation models. The integration of these into a unified control or modeling framework has recently been completed in the code TRANSOPTR [34].

Historically, beam delivery configurations at ISAC have been stored in two formats: HTML webpages and spreadsheets. Beam tuning initially required operators to manually input device setpoints into the EPICS control system [35], consulting various information sources for each setting. This process was both time-consuming and error-prone. Beam from one of ISAC’s ion sources would be threaded through the system, from one Faraday cup (FC) to the next, with manual adjustment of transverse steerers and quadrupole lenses to optimize transmission.




FIG. 2. All operator input commands sent into the control system over a 24-h time period, displayed in 5-min timebins and shown over 5 consecutive days. In this instance, operators were manually tuning the ISAC-I linac in preparation for scheduled beam delivery. Power law normalization applied for clarity.

While this made the tuning process robust and self-optimizing, it posed a challenge on operational resources due to its intensive nature for operators.

For example, Fig. 2 illustrates a typical scenario where ISAC-I is prepared for experimental beam delivery. In the figure, all beam optics and rf control system commands issued by operators are accumulated into 5 min bins and plotted over 5 consecutive days; each row corresponds to a day and the horizontal axis shows the time of day. The intensive nature of ISAC postaccelerator tuning can be appreciated by the density of these commands with minimal empty intervals, as operators must continually adjust the machine throughout each day. Automating the tuning of the optics and steerers can therefore offload a substantial fraction of these routine commands and increase operators' capacity to handle more complex tasks or parallel beam tuning.

ISAC supports a broad program of experiments, namely β -NMR [36], DRAGON [37], GRIFFIN [2], TRINAT [38], and TITAN [39]. With the advent of ARIEL and increased user expectations, rapid, accurate, and reliable beam delivery is essential to the facility's future scientific output. The aforementioned fragmentary nature of simulations complicated whole-machine studies and hindered investigations of emergent phenomena, for example, time-of-flight variations arising from phase instabilities and their propagation through the accelerator lattice, or the propagation of couplings in the beam. We therefore adopt the unified theoretical framework implemented in TRANSOPTR, as outlined below.

III. MODEL-BASED OPTICS TUNING

A. Linear optics

We adopt the linear-optics formalism familiar from classical optics, but now applied to a six-dimensional phase space. A particle is specified by its state vector $\mathbf{X} = (x, p_x, y, p_y, z, \delta)$ and the collective beam distribution is described by a 6×6 covariance (or beam) matrix $\sigma(s)$. The beam matrix contains the second moments of the distribution, and its diagonal elements give the squared beam sizes. The beam envelopes, defined as the rms sizes,

are $\sqrt{\sigma_{11}}$ in x and $\sqrt{\sigma_{33}}$ in y . Here, $p_x = P_x/P_0$ and $p_y = P_y/P_0$ are the transverse momenta normalized to the reference momentum P_0 , and $\delta = (P - P_0)/P_0$ is the relative momentum deviation. In the paraxial limit, p_x and p_y are approximately equal to the geometric slopes $x' = dx/ds$ and $y' = dy/ds$, respectively. Instead of time, the longitudinal position s is our independent coordinate as described in a co-moving Frenet-Serret coordinate system [40]. The covariance matrix propagates across the beam-line through the first-order ordinary differential equations:

$$\frac{d\sigma}{ds} = \mathbf{F}(s)\sigma + \sigma\mathbf{F}(s)^T, \quad (1)$$

where the matrix $\mathbf{F}$ describes the Hamiltonian dynamics through the local electromagnetic fields [21,41]. Consequently, any element within the accelerators and beamlines can be modeled in TRANSOPTR using potentials derived analytically, from field maps, or via finite-element simulations. Computing the infinitesimal transfer matrix at each step gives a prediction of how the beam distribution evolves through the accelerator. Over a small step ds , the corresponding transfer matrix is given by $\mathbf{R} = \mathbf{I} + ds \cdot \mathbf{F}$, which will be convenient for formulating constraints in the next subsection. The rms geometrical emittance, $\epsilon_{\text{rms}} = \sqrt{\det \sigma_{2 \times 2}}$, represents the area of the beam's phase-space distribution. The evolution of the beam matrix is described by a symplectic first-order transfer matrix if the EM fields are free from higher-order aberrations. This means that for optimally efficient beam transport and acceleration through the machine, beams should remain small compared to device aperture sizes and remain close to the central axis through the accelerator. Here, $\sigma_{2 \times 2}$ is the covariance submatrix for a single plane extracted from the full 6×6 sigma matrix, either (x, x') , (y, y') , or (z, δ) , which can be considered independently if there are no couplings beyond the canonically conjugate pairs. The Courant-Snyder (Twiss) parameters [42,43], illustrated in Fig. 3, are obtained by factoring out the horizontal emittance ϵ_{rms} from this 2×2 block matrix; for the (x, x') pair they are defined as




FIG. 3. Evolution of a converging beam through a drift. (a) Starting distribution where beam is converging $\alpha_x < 0$. (b) Beam reaches a waist where $\alpha_x = 0$. (c) Beam starts diverging $\alpha_x > 0$.

$$\beta_x = \frac{\sigma_{11}}{\epsilon_{\text{rms}}}, \quad \gamma_x = \frac{\sigma_{22}}{\epsilon_{\text{rms}}}, \quad \alpha_x = -\frac{\sigma_{12}}{\epsilon_{\text{rms}}}, \quad (2)$$

where β_x is related to the width of the phase-space distribution, γ_x is related to the height, and α_x determines the beam divergence: $\alpha_x < 0$ beam is converging, $\alpha_x > 0$ beam is diverging, and $\alpha_x = 0$ corresponds to a waist. These parameters provide a compact description of the beam for transport and acceleration, and serve as ideal constraints in optics optimizations.

B. Implementation

The backbone of our method is a digital twin of the accelerator: a benchmarked model that is continuously synchronized with the real machine's configuration. This enables online simulations and automated optimization of optical elements (quadrupoles, solenoids, and dipoles) and rf parameters. It possesses sufficient predictive power to compute operational configurations during beam delivery. Because the model runs in parallel with the physical machine, any discrepancy between simulated and measured values can be observed by a physicist, signaling miscalibrations in the model or changes in the machine configuration. The resulting feedback loop drives continuous model refinement.

The basis of the digital twin is a unified model of the entire postaccelerator system, which was developed using the beam envelope code TRANSOPTR. Continuously developed at TRIUMF since 1984 [44], the code has advanced to handle time-dependent elements whose transfer matrices are not analytically known *a priori*. It now supports the inclusion of RFQ [21] and multigap accelerators [20,22,45], allowing for the continuous tracking of the

TABLE I. Initial beam parameters used for the sequential optimization.

Parameter	Symbol	Units	Value
Initial energy per nucleon	E_i	keV/u	2.04
Final energy	E_f	keV/u	153
Mass number	A	u	20
Charge state	q	e	4
Transverse size	x, y	cm	0.7
Transverse divergence	x', y'	mrad	20
Bunch length	z	cm	0.25
Relative momentum spread	δ	...	0.01
Horizontal correlation	r_{12}	...	-0.92
Vertical correlation	r_{34}	...	0.92

beam and transfer matrices along the structure. TRANSOPTR allows parameter optimization through user-defined constraints. Input parameters such as quadrupole strengths, drift lengths, and rf amplitudes can be optimized by imposing constraints on the beam matrix, the transfer matrix, or the reference particle's energy or time. The present version includes two optimizers: Downhill simplex and simulated annealing.

With the MCAT web application [46], TRANSOPTR is deployed for routine control-room use. To minimize the computational latency of online tuning, we employ sequential optimization; the next subsection details this strategy.

C. Sequential optimization

Our digital twin uses sequential optimization, where large optimization problems are broken down into smaller and more manageable sequences. Sequential optimization allows for the single execution of the envelope code, which progressively solves each sequence and moves on to the next, while retaining the optimum values found at each step. This way, the lattice can be adjusted online, with the optimization converging quickly. Although the accelerators and beamlines at TRIUMF-ISAC are shown for the remainder of this paper, the method is general, provided that the system can be represented by the Courant-Snyder Hamiltonian [42], which in turn requires the potential functions to be known. The lattice from OLIS to the medium-energy beam transport (MEBT) was selected to illustrate sequential optimization. This section contains 27 electrostatic and 5 magnetic quadrupoles, since several quadrupoles are wired in parallel, there are a total of 24 independently tunable elements. The optimization was decomposed into nine sequences and achieved a near-identical solution to the design [47]. Table II lists the optimization sequences and constraints, along with the elements varied at each step, using beam parameters listed in Table I. While most elements have their initial value kept at 0, some are given an informed starting value based




FIG. 4. TRANSOPTR simulation from the stable ion source line through the RFQ to the MEBT section. The sequential optimization solution is shown alongside the design solution in solid and dashed lines, respectively. The x and y beam envelopes are 2 rms. The focal strengths are scaled up for the quadrupoles and dipoles and scaled down at the RFQ for visual clarity.

on their typical operating procedure to boost the optimizer. Some constraints are also weighted more heavily as they are more critical to the optimization. Figure 4 compares the design optics with the sequentially optimized solution obtained using simulated annealing. Unlike what was reported in [20], which used a PYTHON wrapper for sequential optimization execution, the method is now implemented natively in TRANSOPTR's FORTRAN codebase. On an i7-3770 S CPU with a single thread and 8 GB memory, the problem from Table II converges in 2 to 3 s. Representative convergence histories for selected optimization sequences are shown in Fig. 5.

The largest deviation from the historical design tune occurs in the horizontal MEBT envelopes. This is intentional. The MEBT optics admit considerable flexibility because the quadrupoles operate at relatively low gradients while transporting highly eccentric phase-space distributions. Rather than reproducing the historical envelope shape, the sequential optimizer searches for an optics

solution that satisfies the required matching conditions while minimizing the onset of optical aberrations. The resulting envelopes therefore differ from the original design, due to TRANSOPTR's aberration reduction routine, which estimates the prevalence of second order effects based on the beam's emittance and distribution [48], resulting in a tune with a smaller overall beam size which in turn reduces nonlinearities from sampling quadrupoles with a beam whose size is large compared to the aperture.

With the development of MCAT and its use of sequential optimization, we can compute end-to-end optics solutions for the entire system in seconds. However, optics alone do not guarantee the beam is transmitted: this requires threading the beam by tuning steerers to compensate for beam alignment errors. To address this, we make use of a machine-learning approach to beam threading, described next (Sec. IV), which computes steerer settings that reliably yield high-transmission solutions.

TABLE II. Sequential optimization steps from OLIS (IOS) to MEBT.

Sequence	Quadrupole IDs	Description	Optimization goal
1	IOS:Q10-13	First achromat	$R_{16} = R_{26} = 0$
2	IOS:Q9, IOS:Q33	Fine-tuning achromat	$\sigma_{12} = \sigma_{34} = 0$
3	ILT:Q34-Q41	Triplet + end quads	$\sigma_{12} = \sigma_{34} = 0$; $\sigma_{11} = \sigma_{33} < 0.5 \text{ cm}^2$; Match Q34 $\rightarrow$ Q41
4	ILT:Q35-Q40	Triplet only	$\sigma_{12} = \sigma_{34} = 0$; $\sigma_{11} = \sigma_{33} < 0.5 \text{ cm}^2$
5	ILT:Q43-Q46	Second achromat	$R_{16} = R_{26} = 0$
6	ILT:Q42, ILT:Q47	Fine-tuning achromat	$\sigma_{12} = \sigma_{34} = 0$
7	ILT:Q47-Q50	Periodic section	Match Twiss parameters for periodicity
8	IRA:Q1-Q4	RFQ injection	Form waist at RFQ entrance
9	MEBT:Q1-Q5	MEBT matching	Form waist at MEBT end

Sequences 1 and 5 tune the achromat sections around each bend. These achromats cancel the dispersion introduced by the dipoles, ensuring that off-momentum particles exit with no position or angular dependence on δ . In practice, the achromatic quadrupoles enforce $\mathbf{R}_{16} = \mathbf{R}_{26} = 0$ at the exit of a horizontal bend, nullifying transverse-longitudinal couplings which can cause unexpected beam size changes leading to transmission loss. If the dispersion is uncorrected, the focal point of quadrupoles will be smeared in space based on δ , which in turn can create a halo around the beam, compromising the ability to focalize through narrow apertures downstream, for example at the experiment's detector. While it is possible to use sequential optimization to optimize the RFQ, this was not done in this work for simplicity; a static solution was used instead.

IV. STEERING CORRECTIONS

The initial centroids of the beam distribution are unknown due to several factors: (1) Lack of diagnostics in the OLIS beam transport section. (2) Mechanical machine misalignments, such as floor-slab settling effects [49]. (3) Stray fields from both the multicharge ion source terminal at OLIS and its steerers, which are affected by steerer lensing [50,51], exhibiting quadrupolelike effects that have been found to be proportional to the sum of the parallel plate voltages [52], which arise due to circular skimmer ground plate apertures together with identically biased opposing steering plates [53]. Given the different energies with each delivered beam, the use of theoretical solutions for setting the corrective steerers is not viable. This motivates the use of a Bayesian method for beam threading.

Designed as a complement to MCAT, BOIS provides steering corrections for beam threading. Developed in BoTORCH [54], BOIS optimizes the corrective steerers while reading back the current from the Faraday cup. The BOIS algorithm has been previously reported in [14] and is used in this work to locally optimize groups of steerers while maximizing beam on a Faraday cup in the line. While this work focuses on one section, BOIS is typically applied in a manner analogous to sequential optimization: the problem can be broken into sequences, allowing rapid and efficient local optimizations [24], from low- to high-energy segments in the machine.

Bayesian optimization has become a popular tool for many problems in accelerator physics [23]. Its main advantage is sample efficiency; it can converge on optimal solutions after only a handful of evaluations, making it an ideal choice for online optimization tasks, such as accelerator tuning and real-time beam delivery. This is a probabilistic technique based on Bayesian inference where a surrogate model, typically a Gaussian process (GP) [55], is defined on prior information. The objective is to find the

input values that maximize an unknown noisy objective function $\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$.

The selection of new evaluation points is determined by the acquisition function (AF), which quantifies the utility of sampling at a given location in the input space. Acquisition functions balance exploration (sampling in regions of high uncertainty) against exploitation (sampling near regions with high predicted objective values). Two of the most commonly used AFs are the upper confidence bound (UCB) and the expected improvement (EI) [56,57], defined as

$$\text{UCB}(\mathbf{x}) = \mu(\mathbf{x}) + \sqrt{\beta}\sigma(\mathbf{x}), \quad (3)$$

$$\text{EI}(\mathbf{x}) = \mathbb{E}[\max(f(\mathbf{x}) - f^* - \xi, 0)], \quad (4)$$

where $\mu(\mathbf{x})$ and $\sigma(\mathbf{x})$ are the GP posterior mean and standard deviation, respectively, and f^* is the best observed objective value. UCB favors points with either high posterior mean or uncertainty, whereas EI targets points with the largest expected gain over f^* , which combines both the probability and magnitude of improvement. The parameters β (for UCB) and ξ (for EI) control the trade-off between exploration and exploitation. Larger β or ξ increase the incentive to explore uncertain regions of the input space, while smaller values bias the search near the current optimum. Striking the right balance is important for reliably locating a global maximum. In this work, BOIS was used in a standard Bayesian optimization configuration: (i) Objective and measurement. The objective is to maximize the FC current. Performance is quantified by the transmission, defined as the FC current normalized to the upstream FC current. Typical measurement noise is on the order of $\sigma \approx 1\% - 3\%$ of full scale. This noise is treated as a hyperparameter, which is learned by the GP. (ii) Bounds and safeguards. Each element is subject to predefined limits; for example, magnetostatic steerers are restricted to $[-100, 100]$ A. Power supplies can span these ranges without limitation. (iii) Zero-current crossing. Magnetic steerers are powered by unipolar supplies coupled with polarity switches, which introduce a finite delay for crossing zero current. For this reason, when BOIS causes a steerer to cross zero, a 5 s delay is applied to prevent loss of communication with the device [58], which is unresponsive during a polarity switch. (iv) BOIS settings. The GP surrogate uses a Matérn kernel ($\nu = 5/2$) with a Gamma probability distribution as the prior [14]. Inputs were normalized and outputs standardized. The random sampling stage was initialized at the midpoints of the parameter bounds [24]. BOIS allows for the usage of UCB ($\beta = 3$) or EI ($\xi = 0.1$) acquisition functions.




FIG. 5. Optimization history for select sequences. $\|\chi\|$ shown in a red solid line alongside the tunable elements in dotted lines. $\|\chi\|$ was normalized to $[0,1]$. Here, $\|\chi\|$ denotes the scalar objective function minimized by TRANSOPTR’s optimizer, formed from the weighted residuals associated with the user-defined beam optics constraints. The tunable elements were normalized using the maximum and minimum of all the parameters.

V. RESULTS AND DISCUSSION

We evaluated two tuning strategies on the MEBT to DTL-injection segment, which contains eight quadrupoles and six steerers. In the decoupled method, MCAT’s digital twin with sequential optimization sets the quadrupoles, and BOIS then tunes only the six corrective steerers. In the fully Bayesian method, the optimizer is allowed to vary all 14 parameters (both quadrupoles and steerers), consistent with practice at other facilities [15–19]. A stable OLIS beam was delivered in three configurations, and for each method we ran six trials using two acquisition functions: three with EI and three with UCB.

The quadrupole optimization performed with MCAT and BOIS addresses related but distinct problems. MCAT seeks to find the quadrupole settings that best satisfy a set of optics constraints, for example, realizing an achromatic bend or a periodic section. By matching these manually defined constraints, MCAT yields beam envelopes that correspond to high-transmission solutions through the transport line. In contrast, BOIS directly searches for quadrupole settings that maximize the measured transmission. Although they take different approaches, both methods ultimately aim to maximize beam transmission.

Although TRANSOPTR computes beam centroids throughout the lattice, centroid information is not used by the optimizer in this work. The Bayesian optimizer is deliberately applied in a minimal configuration, using only measured transmission as the objective. The objective is to demonstrate that a simple transmission-based optimization strategy is sufficient once the optics are constrained by the physics model. This isolates the contribution of the model-coupled framework and avoids conflating improvements in beam delivery with increasingly elaborate objective functions or trajectory constraints.

A naive strategy was chosen for the initial sampling stage, where sampling occurs around the midpoints of the

parameter bounds for both the steerers and the quadrupoles. In the case of steering, sampling occurs around zero steering, which ensures no biasing toward any direction (e.g., no steering beam only to the right). For quadrupoles, the midpoints follow the typical linear optics design. A variable number of sampling points is typically used, ranging from a minimum of $2d + 1$ to a maximum of $10d$ [59], where d is the number of elements used in the optimization. However, in this paper, a fixed number of sampling points based on the minimum $2d + 1$ was chosen due to limited beam time.

A. Tuning through the DTL

Before considering automated steering corrections, it is first necessary to establish that the underlying optics model provides a sufficiently accurate representation of the machine. Sequential optimization, implemented as reported in [20], enables the ISAC-DTL optics to be computed in real time for any attainable output beam energy. This capability removes the need for manual tuning of the DTL and HEBT quadrupoles, historically one of the most time-consuming steps in beam delivery. Figure 6 compares the resulting model predictions with measured beam profiles for a $^{16}\text{O}^{3+}$ beam, demonstrating good agreement and achieving 91% transmission through the DTL following only manual steering corrections at its injection. Transmission losses through the DTL are due to the relatively small apertures in the accelerator lattice [60]. The MEBT quadrupole model required for accurate DTL injection, including fringe-field effects, is presented in [61]. Having established the operational predictive accuracy of the optics model, the remaining task is to compensate for beam trajectory errors through automated steering corrections.

The largest discrepancy occurs in the vertical profile at the MEBT section. This region has historically been the




FIG. 6. Model-computed x and y 2 rms beam envelopes through the ISAC drift-tube linac (solid lines) compared with rms beam sizes from profile monitors (markers) for a $^{16}\text{O}^{3+}$ beam. Optics were generated in real time via sequential optimization [20] using an MEBT quadrupole model with fringe-field effects from [61]. The quantity *foc. str.* denotes the quadrupole focal strength. Dataset also used in [61] for model validation, but shown here to illustrate operational use of the parallel-modeling framework. The plotted markers represent fitted 2 rms beam sizes. Statistical fit uncertainties are smaller than the marker size; the dominant uncertainties are systematic and arise from profile-monitor calibration, background subtraction, and the imperfectly known incoming beam distribution.

most challenging to model accurately because of the sensitivity of the MEBT optics to quadrupole transformation errors, fringe-field effects, and highly eccentric beam distributions, as discussed in Refs. [61,62]. These effects motivated the subsequent refinement of the MEBT quadrupole model. Consequently, the model is not intended to reproduce every individual profile measurement exactly; rather, the available profile monitors constrain the overall envelope evolution and provide an operational consistency check. Despite the localized discrepancy, the measured profiles follow the predicted trends sufficiently well to support the computed optics and the routinely achieved high-transmission operation.

B. Decoupling advantage

Figure 7 compares BOIS optimization through the DTL for two sets of tuning cases. In the first case, all quadrupoles are preset using MCAT, and BOIS is applied only to the steerers. Here, convergence to high-transmission solutions is both rapid and reproducible, as indicated by the tight cluster of solutions on the left side of the plot. In the second case, BOIS also varies the MEBT quadrupoles. This step was historically necessary prior to refining the quadrupole model [24,61], as fringe-field and effective-length errors prevented the decoupled approach. Allowing BOIS to control the MEBT quadrupoles approximately doubles the dimensionality of the search space, leading to




FIG. 7. $^7\text{Li}^+$ BOIS optimization performance through the DTL for two optimization configurations: steering only (left cluster of points) and steering plus MEBT quadrupoles (right cluster). One step/iteration corresponds to one write-measure cycle for BOIS. Different hues of blue/red correspond to different runs, totaling six runs each. Some runs were terminated earlier as they reached a high transmission and were thus considered successful due to limited beam time.

slower average convergence and reduced reliability. Because BOIS begins with a random sampling phase, an unfavorable initial sample can leave the algorithm without any high-transmission starting points, forcing it to spend considerably more iterations “hunting” before effective optimization can occur.

These results highlight the advantage of our decoupled strategy: by using MCAT for quadrupoles, with constraints derived from expert knowledge of the facility, we reduce the dimensionality of the problem that BOIS must solve. This leads to faster and more reliable convergence, enabling high-performance automated tuning for beam delivery.

Table III compares the decoupled and fully Bayesian methods. The decoupled method consistently outperforms the fully Bayesian approach: it converges in fewer iterations, achieves higher average transmission, and is more reliable, with reduced variability between runs. Figure 8 shows the mean and standard deviation from averaged runs. The $^{16}\text{O}^{3+}$ test had three different convergence regions for the fully Bayesian approach, with the last region only including one trial, as seen in Fig. 8. For the $^7\text{Li}^+$ test, there was a stripping foil inserted downstream of DTL, leading to an amplified current reading due to the multiple resulting charge states. In this case, transmission was calculated by normalizing to the highest found current across all runs. Together, these results show that the decoupled method is more reliable and efficient for ISAC-linac tuning.

This improvement arises from several factors. First, as noted above, the decoupled approach significantly reduces the dimensionality of the problem, so the Bayesian optimizer searches a much smaller configuration space

TABLE III. Performance at the 95% transmission target over six independent trials. Cycle counts are median [IQR] among successful trials. An em dash indicates that no trial attained the target.

Species	Strategy	Success	T_{95}	Final transmission (%)
$^{16}\text{O}^{3+}$	Fully Bayesian	1/6	167 ^a	80.3 ± 25.2
	Decoupled	4/6	16[10–21]	94.6 ± 3.5
$^4\text{He}^+$	Fully Bayesian	0/6	...	4.7 ± 8.8
	Decoupled	6/6	24[22–25]	98.4 ± 0.4
$^7\text{Li}^+$	Fully Bayesian	1/6	68 ^a	90.7 ± 5.6
	Decoupled	4/6	25[22–29]	94.7 ± 2.2

^aDenotes a target attained by only one trial.

for high-transmission solutions. Second, continually varying the quadrupoles with BOIS introduces complications due to magnetic hysteresis. In contrast, MCAT mitigates this by using unipolar degaussing [63], which minimizes quadrupole gradient errors. While the same procedure could in principle be incorporated for the fully Bayesian approach, it is impractical due to the time required for each unipolar degaussing cycle (~ 2 min).

Adding quadrupoles raises the dimension from d_s (steerers) to $d_s + d_q$. For a fixed initialization size k and evaluation budget T , the space-filling resolution of the initial sample scales roughly like $k^{-1/d}$, so increasing d dilutes coverage per knob; BOIS therefore expends more iterations on exploration (“hunting”) before stabilization, reducing reliability at fixed T .

Another substantial advantage of our method is low diagnostic requirements. No new diagnostics are required beyond those used for standard tuning. The only diagnostics required are a set of beam profile monitors and Faraday cups, which are standard and found in most facilities. The server requirements are also minimal: TRANSOPTR is CPU-based and Bayesian optimization is more efficient with a CPU as long as the number of training points remains low [23]. This sidesteps the need for greater compute power, which is the most expensive component of some machine learning approaches.

C. Statistical summary (small sample)

All tests done at ISAC occurred during scheduled machine development periods. Beam availability is limited, and the accelerator configuration varies with the preceding experiment. The three case studies were therefore acquired on separate development shifts. Given these constraints, we have prioritized comparisons with identical beam compositions but different optimization methodology.

Because we have $N = 3$ such cases (distinct species/energies) with six runs per method per case, we report per-case effects (convergence time and final transmission) and avoid more involved hypothesis tests. We have compared




FIG. 8. Optimization performance through the DTL for three beam configurations, and two optimization methods: decoupled and fully Bayesian. Each method is averaged over six trials. Shaded band shows median and 2σ interval across runs.

the proposed method (decoupled) to a baseline (fully Bayesian optics+steering) within each case (species/energy). One *step* is a single write-measure cycle (apply new steerer settings, read the Faraday cup).

Nevertheless, the contrast in Table III is clear: Across all presented cases, there is a clear advantage to decoupling, with an average of 27 iterations taken to reach well above 90% transmission through the DTL and into the HEBT line. For the fully Bayesian method, despite significantly more iterations, final transmissions only exceeded 90% in the $^7\text{Li}^+$ case, with both other species underperforming. Within the presented data, a clear trend is observed: Decoupling both reduces the configuration space’s volume and speeds up convergence to a high transmission final result, owing to the predictive power of the envelope model.

D. Limitations and safety

Single-region Bayesian optimization was used in this work, which can become inefficient as the dimensionality of the problem increases. More specialized techniques, such as TuRBO [64], address this limitation by optimizing within adaptively configured local trust regions. These methods can improve scalability in higher-dimensional problems and may lead to faster convergence.

The decoupled strategy assumes correctly calibrated quadrupole strengths. Large calibration errors or strong skew terms can degrade the model optics' predictive power and reduce overall performance. Fast rfphase or energy drifts or high space charge can couple optics and steering; in such cases, local reoptimization may be required. Ion source fluctuations are also currently not monitored by BOIS.

All algorithmic setting changes are constrained by per-device bounds and slew-rate limits enforced at the control system layer. Before each evaluation, we verify interlocks and optics device statuses. If a trip occurs, the system reverts to a safe setting by automatically inserting the source Faraday cup. BOIS never writes to radiofrequency cavities, and all optics and steerer commands are constrained by the control system and cannot bypass the safety system. A vacuum-loss interlock is continuously monitored throughout the optimization. A kill switch allows the user to terminate the algorithm immediately.

VI. CONCLUSION

We have demonstrated a hybrid and physics-first approach to accelerator tuning that uses a validated beam-physics model where predictive power is high and reserves Bayesian optimization for steering corrections, where machine alignment errors and other difficult-to-model effects dominate. Under realistic beam-time and operational constraints, this strategy provides a practical compromise between model-based control and purely data-driven optimization.

Bayesian optimization is applied only to the steering corrections, addressing the beam threading problem in the presence of alignment errors, fringe fields, and limited diagnostics. Removing the optics from the search space avoids the slow convergence and variability seen in fully

Bayesian optimization. Tests on the ISAC postaccelerator show that the decoupled method converges rapidly compared to a fully Bayesian approach, achieves higher average transmission, and is more robust across species and energies.

In operation, the framework turns what was a manual, operator-intensive task into a rapid and repeatable procedure. It requires no additional diagnostics and only modest computing resources, allowing it to be applied during beam delivery without interrupting experiments. The same principles, *model where possible and optimize where necessary*, are portable to other accelerators facing high-dimensional tuning problems.

Work is underway to extend automation to other parameters at TRIUMF, in particular rf phase control, which drifts with temperature [65] and impacts the longitudinal configuration. Integrating these into the same framework will move operations toward a closed-loop, self-optimizing mode suited to high-throughput RIB delivery in the ARIEL era, supported by autonomous tuning methods that the operators can wield at a high level.

ACKNOWLEDGMENTS

We acknowledge the help and support provided by RIB operators during data acquisition, and thank Rick Baartman, Stephanie Rädcl, and Spencer Kiy for their advice and useful discussions. TRIUMF is funded under a contribution agreement with the National Research Council of Canada (NRC). O. H. acknowledges support from the Natural Sciences and Engineering Research Council of Canada (NSERC) under Grant No. SAPPJ-2023-00038, which supported his work and some of the computing resources used in this study. TRIUMF is located on the traditional, ancestral, and uncaded territory of the Musqueam People, who for millennia have passed on their culture, history, and traditions from one generation to the next on this site.

DATA AVAILABILITY

The data that support the findings of this article are not publicly available. The data are available from the authors upon reasonable request.

APPENDIX: FULL RUN PLOTS (${}^7\text{Li}^+$)

The individual ${}^7\text{Li}^+$ optimization trials for the decoupled and fully Bayesian methods are shown in Figs. 9(a)–9(l).




FIG. 9. Individual trials for the decoupled method.




FIG. 10. Individual trials for the fully Bayesian method.

-
- [1] R. Neugart, J. Billowes, M. L. Bissell, K. Blaum, B. Cheal, K. T. Flanagan, G. Neyens, W. Nörtershäuser, and D. T. Yordanov, Collinear laser spectroscopy at ISOLDE: New methods and highlights, *J. Phys. G:Nucl. Part. Phys.* **44**, 064002 (2017).
 - [2] C. E. Svensson and A. B. Garnsworthy, The GRIFFIN spectrometer, *Hyperfine Interact.* **225**, 127 (2014).
 - [3] J. Nielsen, Fundamentals of LHC experiments, in *String Theory and Its Applications* (World Scientific, Singapore, 2011), pp. 127–152.
 - [4] P. Jenni and T. S. Virdee, *The Discovery of the Higgs Boson at the LHC* (Springer International Publishing, Cham, 2020), pp. 263–309.
 - [5] A. K. H. Robertson, C. F. Ramogida, P. Schaffer, and V. Radchenko, Development of ^{225}Ac radiopharmaceuticals:

- TRIUMF perspectives and experiences, *Curr. Radiopharm.* **11**, 156 (2018).
- [6] Z. Liu, A. Trudel, R. S. Augusto, and K. Buckley, Shielding assessment of the IAMI facility, *Radiat. Phys. Chem.* **177**, 109154 (2020).
- [7] R. Catherall, W. Andreazza, M. Breitenfeldt, A. Dorsival, G. J. Focker, T. P. Gharsa, T. J. Giles, J.-L. Grenard, F. Locci, P. Martins, S. Marzari, J. Schipper, A. Shornikov, and T. Stora, The ISOLDE facility, *J. Phys. G* **44**, 094002 (2017).
- [8] G. C. Ball, G. Hackman, and R. Krücken, The TRIUMF-ISAC facility: Two decades of discovery with rare isotope beams, *Phys. Scr.* **91**, 093002 (2016).
- [9] P. Chauveau, V. Bosquet, S. Damoy, P. Delahaye, M. Dubois, P. Jardin, M. Lalande, L. Maunoury, and J. C. Thomas, Latest improvements of the SPIRAL1 facility at GANIL, *Nucl. Instrum. Methods Phys. Res., Sect. B* **541**, 61 (2023).
- [10] R. E. Laxdal, P. G. Bricault, T. Reis, and D. V. Gorelov, A separated function drift-tube linac for the ISAC project at TRIUMF, in *Conf. Proc. C* (IEEE, Vancouver, BC, Canada, 1997), Vol. 970512, pp. 1194–1196.
- [11] J. Dilling, R. Krücken, and L. Merminga, *ARIEL overview* (Springer Netherlands, Dordrecht, 2014), pp. 253–262.
- [12] G. Ball, I. Dillmann, A. Garnsworthy, G. Gwinner, R. Kanungo, G. Morris, and C. Ruiz, The TRIUMF-ISAC facility: Recent highlights in RIB science and future prospects with ARIEL, *Nucl. Phys. News* **30**, 27 (2020).
- [13] R. E. Laxdal and V. Zvyagintsev, ARIEL e-LINAC: Commissioning and development, *J. Phys.: Conf. Ser.* **747**, 012089 (2016).
- [14] E. Ghelfi, A. Katrusiak, R. Baartman, W. Fedorko, O. Kester, G. Kogler Anele, O. Shelbaya, and D. Tanyer, Bayesian optimization for ion beam centroid correction, *Rev. Sci. Instrum.* **96**, 023304 (2025).
- [15] Y. Ong *et al.*, ALPI-PIAVE performance at INFN-LNL with advanced optimization algorithms, in *Proc. HIAT2025, number 16 in HIAT* (JACoW Publishing, Geneva, Switzerland, 2025), pp. 220–223.
- [16] A. Awal, J. Hetzel, R. Gebel, V. Kamerzhiev, and J. Pretz, Optimization of the injection beam line at the cooler synchrotron COSY using Bayesian optimization, *J. Instrum.* **18**, P04010 (2023).
- [17] S. Lopez-Caceres and D. Santiago-Gonzalez, AI-assisted transport of radioactive ion beams, *Phys. Rev. Accel. Beams* **28**, 072802 (2025).
- [18] T. Nishi *et al.*, Development of automatic beam tuning system using Bayesian optimization for high intensity heavy ion beams at RIBF, in *Proc. HIAT2025, number 16 in HIAT* (JACoW Publishing, Geneva, Switzerland, 2025), pp. 140–143.
- [19] S. R. Garces, L. Le, M. Au, A. Schmidt, J. P. Ramos, M. Dierckx, D. Atanasov, I. De Boi, S. Rothe, L. Popescu, and S. Derammelaere, Isotope separator on-line system tuning: Bayesian optimization applied to the transport beamline case, *Nucl. Instrum. Methods Phys. Res., Sect. B* **568**, 165859 (2025).
- [20] O. Shelbaya, T. Angus, R. Baartman, P. M. Jung, O. Kester, S. Kiy, T. Planche, and S. D. Rädcl, Autofocusing drift tube linac envelopes, *Phys. Rev. Accel. Beams* **24**, 124602 (2021).
- [21] O. Shelbaya, R. Baartman, and O. Kester, Fast radio frequency quadrupole envelope computation for model-based beam tuning, *Phys. Rev. Accel. Beams* **22**, 114602 (2019).
- [22] O. Shelbaya, R. Baartman, P. Braun, P. M. Jung, O. Kester, T. Planche, H. Podlech, and S. D. Rädcl, Tuning methods for multigap drift tube linacs, *Rev. Sci. Instrum.* **95**, 033302 (2024).
- [23] R. Roussel *et al.*, Bayesian optimization algorithms for accelerator physics, *Phys. Rev. Accel. Beams* **27**, 084801 (2024).
- [24] O. Hassan, O. Shelbaya, W. Fedorko, T. Planche, and O. Kester, Strategy for Bayesian optimised beam steering at TRIUMF-ISAC's MEBT and HEBT beamlines, *J. Instrum.* **20**, T07005 (2025).
- [25] J. Dilling, R. Krücken, and G. Ball, ISAC overview, *Hyperfine Interact.* **225**, 1 (2014).
- [26] P. Bricault, R. Baartman, M. Dombisky, A. Hurst, C. Mark, G. Stanford, and P. Schmor, TRIUMF-ISAC target station and mass separator commissioning, *Nucl. Phys. A* **701**, 49 (2002).
- [27] K. Jayamanna, Off line ion source terminal, in *ISAC and ARIEL: The TRIUMF Radioactive Beam Facilities and the Scientific Program*, edited by J. Dilling, R. Krücken, and L. Merminga (Springer Netherlands, Dordrecht, 2014), pp. 51–62.
- [28] R. Baartman, ISAC LEBT, in *ISAC and ARIEL: The TRIUMF Radioactive Beam Facilities and the Scientific Program* (Springer, Dordrecht, Netherlands, 2014), pp. 69–77.
- [29] K. R. Crandall and T. P. Wangler, PARMTEQ: A beam-dynamics code fo the RFQ linear accelerator, *AIP Conf. Proc.* **177**, 22 (1988).
- [30] S. Koscelniak, Design history of the ISAC RFQ. Technical Report TRI-DN-95-04, TRIUMF, 1995, <https://lin12.triumf.ca/design-notes/1995/TRI-DN-95-04.PDF>.
- [31] K. L. Junc, LANA, a computer code for beam dynamics simulation in multi-cavity linacs and its application to the Fermilab Linac upgrade 6 1993.
- [32] P. N. Ostroumov, D. V. Gorelov, and R. E. Laxdal, Use of the LANA code for the design of a heavy ion linac, in *Conf. Proc. C* (Fermilab, Batavia, IL, USA, 1997), Vol. 970512, p. 2621.
- [33] R. Laxdal, Design specification for ISAC HEBT. Technical Report TRI-DN-99-23, TRIUMF, 1999, <https://lin12.triumf.ca/design-notes/1999/TRI-DN-99-23.PDF>.
- [34] O. Shelbaya, R. A. Baartman, and O. K. Kester, End-to-end RMS envelope model of the ISAC-I linac, in *12th International Particle Accelerator Conference (IPAC'21)*, THPAB204 (JACoW Publishing, Campinas, SP, Brazil, 2021), p. 4183, .
- [35] L. R. Dalesio, A. J. Kozubal, and M. R. Kraimer, EPICS architecture EPICS Architecture as a contribution to the 1991 ICALEPCS proceedings. Technical report, Los Alamos National Lab., NM (United States), 1991, pp. 278–282.
- [36] G. D. Morris, β -NMR (Springer Netherlands, Dordrecht, 2014), pp. 173–182.

- [37] J. Fallis, *Nuclear Astrophysics Measurements with DRAGON* (Springer Netherlands, Dordrecht, 2014), pp. 219–224.
- [38] B. Fenker, A. Gorelov, D. Melconian, J. A. Behr, M. Anholm, D. Ashery, R. S. Behling, I. Cohen, I. Craiciu, G. Gwinner, J. McNeil, M. Mehlman, K. Olchanski, P. D. Shidling, S. Smale, and C. L. Warner, Precision Measurement of the β Asymmetry in Spin-Polarized ^{37}K Decay, *Phys. Rev. Lett.* **120**, 062502 (2018).
- [39] P. Delheij, L. Blomeley, M. Froese, G. Gwinner, V. Ryjkov, M. Smith, and J. Dilling, The TITAN mass measurement facility at TRIUMF-ISAC, *Hyperfine Interact.* **173**, 123 (2007).
- [40] H. Wiedemann, *Elements of Classical Mechanics* (Springer International Publishing, Cham, 2015), pp. 83–98.
- [41] R. Baartman, What the F. Technical Report TRI-BN-24-38, TRIUMF, 2024, https://lin12.triumf.ca/text/public/text/design_notes/TRI-BN-24-38/TRI-BN-24-38.pdf.
- [42] E. D. Courant and H. S. Snyder, Theory of the alternating-gradient synchrotron, *Ann. Phys. (Leipzig)* **3**, 1 (1958).
- [43] R. Q. Twiss and N. H. Frank, Orbital stability in a proton synchrotron, *Rev. Sci. Instrum.* **20**, 1 (1949).
- [44] R. Baartman, TRANSOPTR: Changes since 1984. Technical Report TRI-BN-16-06, TRIUMF, 2016, https://lin12.triumf.ca/text/public/text/design_notes/b2016_06/TRI-BN-16-06_TRANSOPTR.pdf.
- [45] C. M. Lund, P. M. Jung, M. J. Maher, J. Bancheri, T. Planche, R. Baartman, and J. Seuntjens, Development and application of a linear optics model for dielectric wall accelerators, *Phys. Rev. Accel. Beams* **28**, 010401 (2025).
- [46] O. Shalbaya, Model coupled accelerator tuning. Ph. D. thesis, 2023.
- [47] R. Baartman, ISAC LEBT. Technical Report TRI-BN-12-10, TRIUMF, 2012, https://lin12.triumf.ca/text/public/text/design_notes/b2012_10/LEBT.pdf.
- [48] R. Baartman, Low-energy beam transport design, in *Proceedings, CAS—CERN Accelerator School: Beam Injection, Extraction and Transfer, volume 5 of CERN Yellow Reports: School Proceedings* (CERN, Geneva, Switzerland, 2018), Vol. 3, pp. 491–506. Presented at the CAS. 2017 course in Erice.
- [49] O. Shalbaya and R. Baartman, Langevin-like DTL triplet BI fits and analysis of transverse DTL tuning difficulties. Technical Report TRI-BN-19-18R, TRIUMF, 2019, <https://beamphys.triumf.ca/~oshelb/phynotes/dtlInvestigation/dtlBI.pdf>.
- [50] O. Shalbaya and J. Adegun, A record of OLIS steerer lensing. Technical Report TRI-BN-24-05, TRIUMF, 2024, <https://beamphys.triumf.ca/~oshelb/phynotes/OSL/OSL.pdf>.
- [51] O. Shalbaya, Mitigation of steerer lensing effects for radioactive beam transport. Technical Report TRI-BN-24-25, TRIUMF, 2024, <https://beamphys.triumf.ca/~oshelb/phynotes/IMSLensing/IMSLensing.pdf>.
- [52] R. Baartman and T. Planche, Electrostatic steerer lensing effect. Technical Report TRI-BN-24-22, TRIUMF, 2024, https://lin12.triumf.ca/text/public/text/design_notes/2024steerers/2024steerers.pdf.
- [53] I. V. Bylinskii, R. A. Baartman, K. Jayamanna, T. Planche, and Y.-N. Rao, Recent Improvements in Beam Delivery with the TRIUMF's 500 MeV Cyclotron, in *Proc. of International Conference on Cyclotrons and Their Applications (Cyclotrons'16), Zurich, Switzerland, September 11-16, 2016, number 21 in International Conference on Cyclotrons and Their Applications (JACoW Publishing, Geneva, Switzerland, 2017)*, pp. 133–136, [10.18429/JACoW-Cyclotrons2016-TUA04](https://doi.org/10.18429/JACoW-Cyclotrons2016-TUA04).
- [54] M. Balandat, B. Karrer, D. R. Jiang, S. Daulton, B. Letham, A. G. Wilson, and E. Bakshy, BoTorch: A framework for efficient Monte-Carlo Bayesian optimization, *Adv. Neural Inf. Process. Syst.* **33** (2020).
- [55] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*, second ed., (The MIT Press, Cambridge, MA, USA, 2005).
- [56] N. Srinivas, A. Krause, S. Kakade, and M. Seeger, Gaussian process optimization in the bandit setting: No regret and experimental design, in *ICML 2010 Proceedings Contribution* (2010), pp. 1015–1022.
- [57] J. Mockus, V. Tiesis, and A. Zilinskas, *The Application of Bayesian Methods for Seeking the Extremum* (North-Holland, 1978), Vol. 2, pp. 117–129.
- [58] O. Hassan, Polarity switching at the HEBT steerers. Technical Report TRI-BN-24-12, TRIUMF, 2024, https://lin12.triumf.ca/text/public/text/design_notes/TRI-BN-24-12_PolaritySwitching.pdf.
- [59] J. L. Loeppky, J. Sacks, and W. J. Welch, Choosing the sample size of a computer experiment: A practical guide, *Technometrics* **51**, 366 (2009).
- [60] O. Shalbaya, The TRANSOPTR model of the ISAC drift tube linear accelerator—Part I: Longitudinal verification. Technical Report TRI-BN-20-08, TRIUMF, 2020, <https://beamphys.triumf.ca/~oshelb/phynotes/optrDTL/optrDTL.pdf>.
- [61] O. Shalbaya, O. Hassan, R. Baartman, O. Kester, C. Pearce, T. Planche, and L. Zhang, Matching optimization for TRIUMF's rare isotope linac, *Nucl. Instrum. Methods Phys. Res., Sect. B* **573**, 166051 (2026).
- [62] O. Shalbaya, Beam dynamics study of ISAC-MEBT. Technical Report TRI-BN-22-29, TRIUMF, 2022, <https://beamphys.triumf.ca/~oshelb/phynotes/optrMEBT2/optrMEBT2.pdf>.
- [63] J. Nasser, R. Baartman, O. Kester, S. Kiy, T. Planche, S. Radel, and O. Shalbaya, Algorithm to mitigate magnetic hysteresis in magnets with unipolar power supplies, in *IPAC2022 (JACoW, Geneva, Switzerland, 2022)*, pp. 156–159 MOPOST039.
- [64] D. Eriksson, M. Pearce, J. R. Gardner, R. Turner, and M. Poloczek, Scalable global optimization via local Bayesian optimization (2020).
- [65] O. Shalbaya and O. K. Kester, Toward an end-to-end model for ISAC-I accelerators, in *Journal of Physics: Conference Series* (IOP Publishing, Bristol, UK, 2018), Vol. 1067, p. 062028.