

RESEARCH ARTICLE | NOVEMBER 20 2025

Optimization of a cosmic muon tomography scanner for cargo border control inspection

Special Collection: **Muography: Discoveries, Innovations, and Applications**Z. Zaher ; H. Lay ; T. Dorigo ; A. Giammanco ; V. Gulik ; C. Hrytsiuk ; V. A. Kudryavtsev ; M. Lagrange ; T. Metspalu ; G. C. Strong ; C. Turkoglu ; P. Vischia *J. Appl. Phys.* 138, 194903 (2025)<https://doi.org/10.1063/5.0287758>

Articles You May Be Interested In

A comprehensive framework toward the seamless integration of muon reconstruction algorithms with machine learning

J. Appl. Phys. (October 2025)

Pillar embedding visualization for muon-scattering tomography

J. Appl. Phys. (October 2025)

Experimental demonstration of multiple monoenergetic gamma radiography for effective atomic number identification in cargo inspection

J. Appl. Phys. (April 2018)Nanotechnology &
Materials ScienceOptics &
PhotonicsImpedance
AnalysisScanning Probe
Microscopy

Sensors

Failure Analysis &
Semiconductors

Unlock the Full Spectrum.

From DC to 8.5 GHz.

Your Application. Measured.

[Find out more](#)

Optimization of a cosmic muon tomography scanner for cargo border control inspection

Cite as: J. Appl. Phys. 138, 194903 (2025); doi: 10.1063/5.0287758

Submitted: 26 June 2025 · Accepted: 28 October 2025 ·

Published Online: 20 November 2025



Z. Zaher,^{1,a)} H. Lay,² T. Dorigo,^{3,4,5,b)} A. Giammanco,¹ V. Gulik,⁶ C. Hrytsiuk,^{6,7} V. A. Kudryavtsev,² M. Lagrange,¹ T. Metspalu,⁷ G. C. Strong,^{8,c)} C. Turkoglu,² and P. Vischia⁹

AFFILIATIONS

¹Centre for Cosmology, Particle Physics and Phenomenology (CP3), Université catholique de Louvain, B-1348 Louvain la Neuve, Belgium

²Department of Physics and Astronomy, University of Sheffield, Sheffield S3 7RH, United Kingdom

³Department of Computer Science, Electrical and Space Engineering, Luleå University of Technology, 971 87 Luleå, Sweden

⁴INFN, Sezione di Padova, Padova, Italy

⁵Universal Scientific Education and Research Network, Padova, Italy

⁶Institute of Physics, University of Tartu, W. Ostwaldi 1, 50411 Tartu, Estonia

⁷GScan OÜ, Mäealuse 2/1, 12618 Tallinn, Estonia

⁸MODE Collaboration, Oviedo, Spain

⁹Universidad de Oviedo and ICTEA, Oviedo, Spain

Note: This paper is part of the Special Topic on Muography: Discoveries, Innovations, and Applications.

^{a)}Author to whom correspondence should be addressed: zahraa.zaher@uclouvain.be

^{b)}URL: <https://usern.org/>.

^{c)}URL: <https://mode-collaboration.github.io/>.

ABSTRACT

The past several decades have seen significant advancement in applications using cosmic-ray muons for tomography scanning of unknown objects. One of the most promising developments is the application of this technique in border security for the inspection of cargo inside trucks and sea containers in order to search for hazardous and illicit hidden materials. This work focuses on the optimization studies for a muon tomography system similar to that being developed within the framework of the “SilentBorder” project funded by the EU Horizon 2020 scheme. Current studies are directed toward optimizing the detector module design, following two complementary approaches. The first leverages TomOpt, a Python-based end-to-end software that employs differentiable programming to optimize scattering tomography detector configurations. While TomOpt inherently supports gradient-based optimization, a Bayesian Optimization module is introduced to better handle scenarios with noisy objective functions, particularly in image reconstruction-driven optimization tasks. The second optimization strategy relies on detailed GEANT4-based simulations, which, while more computationally intensive, offer higher physical fidelity. These simulations are also employed to study the impact of incorporating secondary particle information alongside cosmic muons for improved material discrimination. This paper highlights the outcomes and key findings from these optimization studies.

© 2025 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0287758>

I. INTRODUCTION

Muon tomography, also known as muography, is a non-invasive imaging technique that leverages naturally occurring cosmic-ray muons to probe the internal structure of large objects.

These muons originate from interactions between high-energy cosmic rays and atomic nuclei in the upper atmosphere, producing showers of secondary particles, primarily pions and kaons. As these mesons decay, they generate muons that travel toward the Earth's

surface with minimal energy loss. At sea level, the flux of cosmic-ray muons is approximately $170 \text{ muons m}^{-2} \text{ s}^{-1}$, with an average energy of vertical muons of the order of 3–4 GeV. Due to their relatively large mass—roughly 200 times that of an electron—these highly energetic muons have a strong penetrating ability, allowing them to traverse several meters of dense material. This makes them particularly valuable for imaging applications where conventional radiation-based methods, such as x-ray imaging, are ineffective due to limited penetration depth.¹

Muon tomography can be classified into two main approaches: muon transmission imaging and muon scattering tomography (MST). The transmission method relies on detecting attenuation in muon flux as particles pass through an object, providing insights into its internal composition. This technique is particularly useful for large-scale applications, such as studying volcanic structures² or uncovering hidden chambers in archaeological sites.³ However, it typically requires prolonged data acquisition periods, which makes it less suitable for applications where rapid imaging is essential. On the other hand, muon scattering tomography (MST) operates by measuring the deflections of muon trajectories caused by multiple Coulomb scattering interactions with the material they pass through. High-atomic-number (Z) materials, such as lead and uranium, tend to cause more significant scattering than lower- Z materials like aluminum or plastic. This property makes MST particularly useful for detecting dense and high- Z substances in diverse applications that span many fields, including nuclear reactor inspection, civil engineering, archaeology, and various industrial and homeland security uses.⁴ In border security applications, MST is a powerful tool for identifying hidden threats and dense contraband in cargo containers, such as radioactive materials or explosives. Furthermore, MST is well-suited for scanning smaller objects and offers the advantage of rapid data acquisition, making it an ideal technology for real-time security applications, where both speed and precision are critical for border security operations.⁵

As muons traverse a material medium, they primarily interact with material nuclei through multiple Coulomb scatterings, where small-angle deviations are most frequent, while occasional large-angle deflections occur as described by Rutherford scattering. After crossing a macroscopic distance, the overall scattering angle distribution consists of a Gaussian core, which accounts for the majority of deflections due to the central limit theorem, along with a longer tail from less frequent large-angle scatterings. For a muon with momentum p (MeV) passing through a material of radiation length X_0 (m) over a distance x (m), the root mean square (RMS) width of the Gaussian component of the scattering angle distribution can be approximated by⁶

$$\theta_{\text{RMS}} = \frac{13.6 \text{ MeV}}{\beta c p} \sqrt{\frac{x}{X_0}}, \quad (1)$$

where βc is the muon's velocity. Hence, by analyzing the scattering angles of muons before and after traversing a volume of interest, valuable insights into the material densities within can be obtained, given the inverse relationship between X_0 and the material density.

Optimizing the design parameters of an MST scanner is crucial to achieving optimal performance, as different configurations impact key factors such as acceptance, detection efficiency, and angular resolution. Design parameters such as the spatial arrangement and size of detectors play a critical role in determining the system's ability to accurately track muon trajectories, maintain a high signal-to-noise ratio, and effectively classify materials. A suboptimal design may lead to reduced tracking precision, lower detection rates, and decreased overall scanning effectiveness. Prior optimization of MST systems has primarily relied on simulations combined with grid-search approaches to evaluate discrete configurations. However, a systematic, end-to-end optimization strategy that directly links detector design to material discrimination or scanning performance has, to our knowledge, not yet been explored. This motivates the approach presented in this work, which seeks to optimize the scanner design in a holistic, performance-driven manner.

In this work, we focus on optimizing the SilentBorder scanner (<https://silentborder.eu/>), a modular MST system designed for cargo screening applications. The scanner consists of multiple hodoscopes, each of which is a self-contained detector module comprising three equally spaced sensitive detector panels inside a protective casing. The panels are made of a mat of plastic scintillator fibers 1 mm in diameter with a pitch of 7.5 mm, achieving a spatial resolution of $120 \mu\text{m}$ and an angular resolution of a few tenths of a milliradian at meter-scale separations.⁵ Compared to drift tubes scanners—for example, those used at LANL⁷—which achieve sub-millimeter resolutions but require complex infrastructure including gas handling, high-voltage systems, and bulky mechanical support, scintillating fibers offer a more compact solution at the cost of a high channel density and demanding optical/electronic readout. In this sense, fibers represent a balance between operational complexity and intrinsic spatial granularity. These hodoscopes are arranged around the scanned cargo, forming a tracking system capable of measuring muon trajectories before and after interaction with the inspected volume. Due to its modular design, the SilentBorder scanner offers flexibility in configuration, allowing different setups to be explored for improved detection performance. To systematically optimize the design parameterization of the hodoscopes, we adopt a dual-method approach combining Monte Carlo simulations and differentiable optimization techniques.^{8,9}

- **GEANT4-based simulation framework:** We employ a GEANT4¹⁰ simulation framework to model different SilentBorder scanner configurations and various cargo materials. GEANT4, a widely used toolkit for simulating particle interactions with matter, provides a realistic modeling environment for evaluating the performance of different scanner layouts.
- **TomOpt-based optimization:** TomOpt¹¹ is a differentiable simulation framework developed for optimizing MST scanner designs through gradient descent.

In Sec. II, we present the GEANT4 simulation framework. Using this framework, Sec. III provides an analysis of the impact of hodoscope positioning on the reconstruction, and Sec. IV presents a study on the impact of secondary particle hits on event

reconstruction. Section V introduces the TomOpt framework, detailing the reconstruction algorithms used. In Sec. VI, we extend TomOpt with Bayesian optimization and compare the optimization results between the reconstruction algorithms introduced in Sec. V. Finally, in Sec. VII, we summarize the main findings of this work and outline future directions for improving the SilentBorder scanner design and reconstruction performance.

II. GEANT4 SIMULATION FRAMEWORK

A GEANT4-based framework was developed by the SilentBorder consortium to provide a simplified model of the SilentBorder scanner concept with the flexibility to simulate a variety of different scanner configurations and cargo contents.

Geometry Description Markup Language (GDML)¹² is used to describe the size, shape, and location of the relevant objects and materials. First, an empty shipping container of external dimensions $6.07 \times 2.44 \times 2.60 \text{ m}^3$ is placed at the center of the geometry with a plywood floor of depth 40 cm. The container is surrounded by a total of 48 hodoscopes, a 3×4 grid on each of the long sides of the container. This arrangement is visualized in Fig. 1. Each hodoscope is itself comprised of three layers (plates) of solid plastic scintillator each of dimensions $2.1 \text{ m} \times 1.1 \text{ m} \times 3.3 \text{ mm}$ and separated by 10 cm. Note the casing and electronics for each hodoscope are not simulated in this simplification, but the 10 cm gaps are introduced to account for the “dead regions” they introduce. The plate is simulated as a solid scintillator rather than as a mat of scintillator fibers. This significantly reduces the complexity of the geometry description and can be considered a good approximation of the fiber mats once the smearing for detector resolution effects is applied later.

A number of different cargo scenarios can be simulated. The setup shown in Fig. 1 involves the inclusion of a $1 \times 1 \times 0.5 \text{ m}^3$ volume in the center of the container. This volume can be simulated using the bulk properties of a number of materials including

- metals, such as iron, steel, or lead
- explosives, such as TNT, RDX, or PETN
- illegal drugs, such as cocaine or cigarettes
- typical legal cargos, such as bananas or toilet roll
- or nuclear materials, such as uranium.

Options are also available to fill the entire container with a single cargo, split the container between a number of cargos, or include a single small element of illicit material and fill the remaining space with a legal material.

Air shower cosmic-ray muons can be simulated in this volume using the Muons for Silent Border (MuSiBo) generator. MuSiBo utilizes a Gaisser parameterization of the muon flux, with modifications to account for muon decay and the Earth's curvature. It is based on work presented in Ref. 13. Muons are sampled on the surface of a cube, a method which better accounts for the high-angle flux when compared to sampling on a single surface above the region of interest. For this simulation, a box of $11 \times 5 \times 4 \text{ m}^3$ was used to fully enclose the volume of the container and all of the hodoscopes.

Finally, the initial muons simulated with MuSiBo are propagated through the simulated geometry via use of the GEANT4

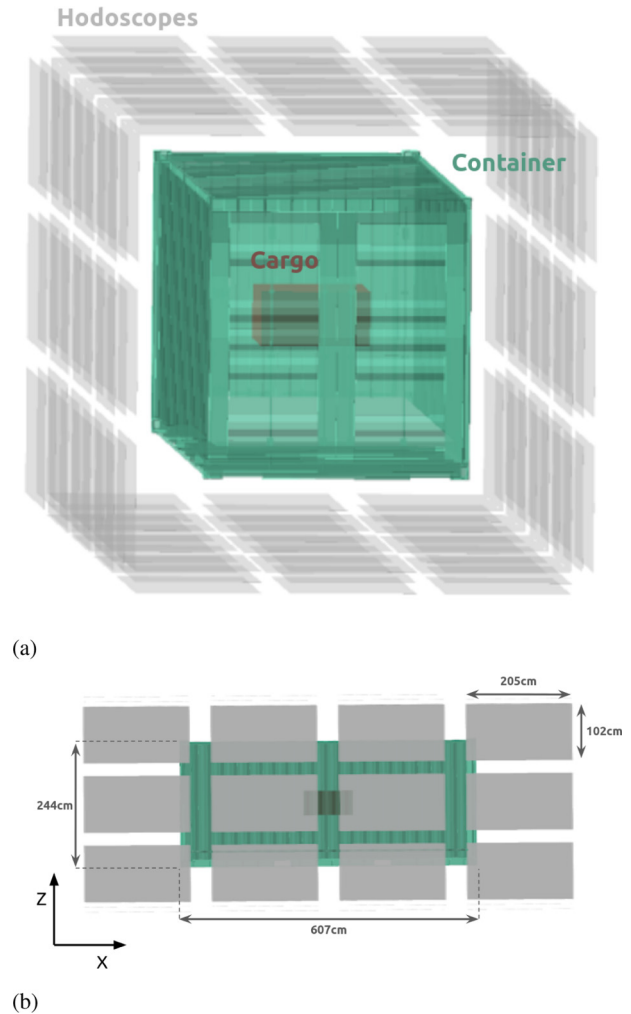


FIG. 1. A visualization of a particular configuration of the detector setup in the GEANT4 simulation. The sensitive detector panels are shown in gray, the shipping container in green, and a small cargo object in red. The front view (a) shows that the coverage extends around all four long sides of the container while the gaps between hodoscope panels are emphasized in the side view (b).

software package.¹⁰ Specifically, the QGSP_BERT physics list is utilized in version 11.0.p03 of the software. Muons, and any secondary particles produced, are stepped through the full geometry with information saved about their trajectory and total energy deposition within each logical volume.

III. GEANT4 STUDY OF HODOSCOPE POSITIONING

The relative positioning of hodoscope plates can impact the performance of the scanner system. This impact can be appreciated in both the efficiency in collecting “interesting” muons and the angular resolution of their trajectories.

A study is performed using the GEANT4 simulation framework described in Sec. II. For the rest of this section, we will use a coordinate system that is relative to each hodoscope group; x will refer to the long in-plane dimension, y to the short in-plane dimension, and z to the out-of-plane dimension. This ensures that moving hodoscope plates in one coordinate is conceptually the same regardless of the rotational orientation of the hodoscope group (top, bottom, or sides).

A baseline configuration was established with the container filled entirely with wood and the hodoscopes as close together in x and y as possible. Note that due to space reserved for electronics and casing this still leaves a ~ 10 cm gap between the active scintillator areas of adjacent hodoscopes. The default spacing in z between the three plates that forms a single hodoscope is 10 cm.

A number of scenarios are created by independently varying the spacing in x , y , and z . Due to physical constraints (such as not overlapping any hodoscope regions), the variations are different for each direction. The cargo is kept identical and for each scenario a 15 min exposure to the cosmic flux is provided by the MuSiBo generator. This 15 min exposure corresponds to a total of 12 683 700 muons in the volume used by MuSiBo, and the GEANT4 simulation required an average of 14.5 CPU h to propagate and save information for each of these muons.

The energy depositions recorded by GEANT4 in the active hodoscope plates pass through a pseudo-detector simulation. At this stage, in order to reduce computational time, we only consider events with two or more depositions before and two or more depositions after the cargo. Without this requirement a muon will not be reconstructed later on, so we can save resources by discarding this event here. This step is later referred to as “filtering.” Any depositions recorded within 1 mm and 20 ns of each other are combined and a Gaussian positional smearing, with width of 1 mm, is then applied to their computed barycenter. This smearing, larger than the real intrinsic per-plate resolution of $120\ \mu\text{m}$, is applied to conservatively approximate the combined effects of clustering, optical cross talk, and readout not explicitly modeled in GEANT4. Finally, a threshold of 0.4 MeV is applied to ensure the deposition would have created a detector “hit.”

These hits form the input to a simple reconstruction chain. Tracks are formed when one hit is found on at least two hodoscope plates. These two plates must be found in different layers of the same hodoscope group (above, below, left, or right). If a hodoscope plate has received multiple hits, it is excluded from track building. This only impacts about 3.5% of track candidates, and, therefore, the small reduction in efficiency is worth accepting for the improved angular resolution. While a more sophisticated algorithm could be designed to make use of this subset, the reduction in operating times for field-based MST scanners would be minimal.

Linear regression is used to find the best fit line between the hits providing the direction of the reconstructed muon track. The final stage of the reconstruction takes events in which the muon had tracks reconstructed in two different hodoscope groups. A simple Point-of-Closest-Approach (PoCA) reconstruction is performed from which the location and angle of the scatter point is determined. An example of the output of the PoCA reconstruction is shown in Fig. 2. The reconstruction chain is significantly less computationally expensive than the simulation, each scenario




FIG. 2. Locations of reconstructed PoCA points within the container in the xy -plane (looking down on the container). The scenario used a $100 \times 100 \times 50$ cm cargo of lead, centered in the container.

requiring an average of 30 CPU minutes to perform the full reconstruction and save the output.

A number of metrics are used to quantify the change in performance when the hodoscope positions are adjusted. They can be split into two categories.

The first three metrics assess the efficiency of the system. As a proportion of the total number of muons simulated by MuSiBo (12 683 700) we can evaluate:

- The number of events with a “Muon of Interest” defined as the number of muons that make any deposition in any of the hodoscopes’ scintillator or the cargo of interest.
- The number of events in which a PoCA point was reconstructed.
- The number of events in which a PoCA point was reconstructed inside the container—thus providing information about the cargo.

The other two metrics assess the angular reconstruction of the muon scattering point. The difference between the reconstructed and true scattering angles is asymmetric due to the positive-only definition of a scattering angle. For smaller angles, the reconstructed angle is more likely to be larger in magnitude than the true angle. Figure 3 shows the resolution in two different example scenarios. We can calculate the mean and standard deviation of this difference distributions as metrics to use in the variation studies.

Figure 4 shows the variation in the efficiency and angular resolution metrics as a function of the gaps between the hodoscopes in x . Note that 0 cm corresponds to the minimum gap between the hodoscope objects. However, a ~ 10 cm gap between the active scintillator will still be present due to the casing and electronics space reserved on each hodoscope. The behavior is identical for y -variations due to the fact that these variations are both in the

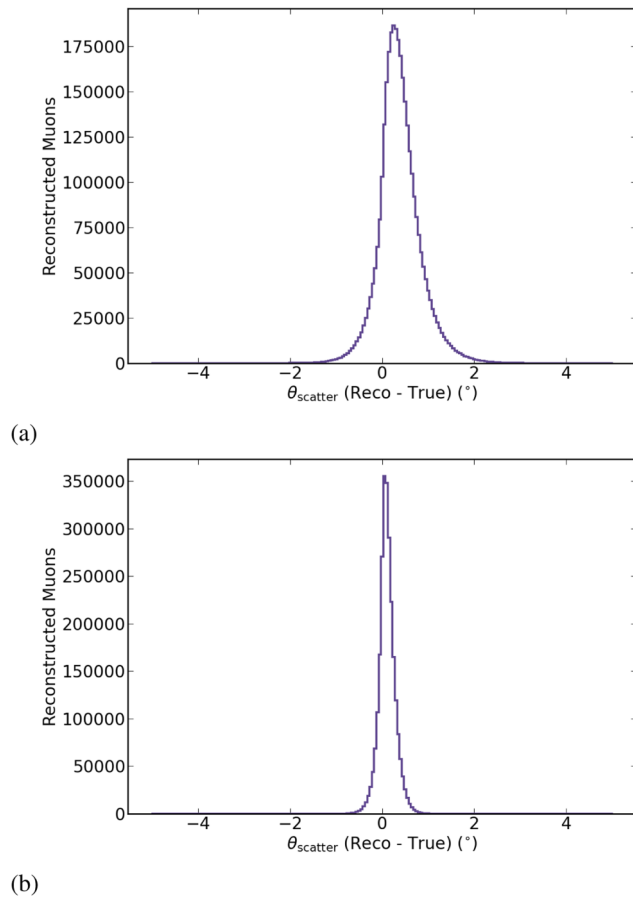


FIG. 3. The resolution of the reconstructed scattering angle in two different example scenarios, (a) the nominal configuration and (b) where the vertical spacing between hodoscope plates has been increased.

plane of the hodoscope plates. Due to physical limitations, however, there is a much larger range of motion in the x -direction (the length of the container) than the y -direction (the width of the container).

In both directions, the number of muons of interest increases as the gaps increase, this is purely because the hodoscopes as an ensemble now cover a larger volume within the total simulation world and hence more of the muons are likely to intercept at least one hodoscope. However, the number of PoCA points, and specifically those in the cargo, decreases due to the gaps that are created in the regions closest to the cargo itself. There is a subsequent reduction in the number of muons that intercept hodoscopes before and after the container and pass through the cargo. The reduction in useful PoCA points per muon will mean that a longer exposure time is required in order to achieve the same level of sensitivity.

There is no visible change in the standard deviation of the angular difference metric, so the angular resolution does not improve. There is a small increase in the mean when going to the




FIG. 4. The change in efficiency (top) and angular resolution (bottom) metrics as the gap between hodoscopes in the x -direction is varied.

largest gaps. This is likely due to the reduced available phase space in the scattering angle created by the gaps.

Figure 5 shows the variation in efficiency and angular resolution metrics as a function of the vertical spacing between the hodoscope plates in z . The nominal gap between the plates is 10 cm. As with the x variations, the number of interesting muons increases as the total scanner volume increases with larger spacing. The number of PoCA points slightly decreases. This occurs because the probability of the muon hitting enough plates to ensure track reconstruction before and after the cargo decreases, but the effect is very minimal. The angular resolution on the other hand is dramatically impacted by the widening of the vertical spacing. By increasing the spacing both the standard deviation and mean parameters reduce significantly before beginning to flatten out. As the distance




FIG. 5. The change in efficiency (top) and angular resolution (bottom) metrics as the gap between hodoscopes in the z -direction is varied.

between the two or three detection points increases, the positional resolution of each plate has a smaller impact on the reconstructed angle, hence improving the angular resolution. The reduction in the mean occurs due to the asymmetry of the distribution. As the resolution improves, asymmetry is naturally reduced in magnitude.

These variation studies can be used to inform decisions about the design for a full-scale SilentBorder-like scanner system. As could be anticipated, minimizing the in-plane gaps between hodoscopes improves the efficiency of the system, thus reducing the necessary exposure time and making the system quicker to use and more viable in border security. Increasing the vertical out-of-plane spacing between hodoscope plates improves the angular resolution due to the lever arm effect, although it does so at the expense of efficiency. For a particular system an appropriate trade-off balance can be found, in the scenario examined in this study Fig. 5 suggests that up to ~ 20 cm the increase in out-of-plane spacing improves the angular resolution enough for the reduction in efficiency to be worthwhile. Beyond this, the resolution improvements plateau and therefore would be unlikely to be worth the efficiency reduction. Note that these studies were completed with a simple PoCA reconstruction and decisions were made to ignore plates with multiple hits or events with three or more tracks. These considerations will have an impact on the magnitude of the effects presented but should not undermine the trajectories of the changes and thus the conclusions drawn.

IV. GEANT4 STUDY OF SECONDARY HITS

As muons pass through the hodoscope, cargo, container, and surrounding air, they can produce secondary particles, primarily delta rays. An assessment is made in this section of the impact of any hodoscope hits induced by such secondary activity. Note that activity due to other particles generated in cosmic air showers is not considered in this study but can also be considered as a source of “secondary” information in MST.

The GEANT4 simulation described in Sec. II shows that about 11 % of the energy depositions recorded in the hodoscope plates resulted from non-muons. However, many of these depositions would likely be irresolvable from the accompanying muon hits with the SilentBorder scanner setup. Using approximate numbers provided by the GScan team, a pseudo-detector simulation was performed (same as Sec. III) in which depositions within the same $1 \times 1 \text{ mm}^2$ area and 20 ns of each other were combined into a single hit. For these amalgamated hits all the relevant energy depositions were summed. Following this amalgamation procedure 8.3 % of the remaining hits resulted purely from non-muon activity, with the remaining 91.7% having at least some contribution from the muon.

Figure 6 shows the energy deposited per hit (in a single hodoscope plate) for these two categories of hits. The energy deposited by a minimally ionizing particle (MIP) crossing the full 3.3 mm plate depth is around 0.6 MeV, and a peak is seen in this region for both muon and secondary hits. The flat shelf at lower energy depositions is created by particles “clipping” a hodoscope plate rather than passing fully through it. The secondary distribution has a third feature, another peak at very low depositions. This results




FIG. 6. The total true energy deposition for each hodoscope plate hit following the application of detector resolution constraints. The distribution is split into hits with some contribution from the primary muon (blue) and hits with no contribution from the primary muon (red).


from the full containment of secondaries produced with very low kinetic energy, which only travel a very short distance in the scintillator before coming to a stop.

The goal of this study was to evaluate what impact deposits from secondaries could have in degrading the performance of MST with the SilentBorder scanners. It is clear from Fig. 6 that above a certain energy threshold there is no method to distinguish between secondary and muon hits based on the scale of the energy deposition.


For the first study a threshold of 0.4 MeV is applied, retaining only those depositions that lie within the MIP peak region. A cargo block of size $1 \times 1 \times 0.5 \text{ m}^3$ is placed at the center of the container with a variety of material types considered. The MuSiBo generator is used to provide muons corresponding to a 15 min exposure time. Figure 7(a) shows the total number of hits recorded across all hodoscopes per event, regardless of whether they were produced via secondaries or muons. In general, the peak structure around 6 hits is produced via the 3 hits left by the muon, before and after the cargo, while the small tail is produced by the presence of any secondary hits. While it appears that there is some variation between the different cargo types this variation stems from the number of events passing the filter, not from a significant difference in the number of secondary hits produced per muon. Higher density and higher-Z materials reduce the number of events that pass the filtering process. Figure 7(b) shows the same distribution but normalized by the number of events that passed the filtering, very limited variation per cargo is now observed.

Figure 8 shows the same analysis performed for a more “realistic” smuggling scenario. For this scenario, the container is filled with bananas and then a $40 \times 40 \times 40 \text{ cm}^3$ box of cocaine is placed in the center of the bananas. This is compared to the scenario in which only the bananas are stored. Again no significant variation is seen in the number of hits recorded indicating the presence of secondary hits cannot, in and of itself, be used to discriminate between different cargoes.

05 December 2025 11:58:21



(a)



(b)

FIG. 7. The total number of hits recorded across all hodoscopes for a variety of cargo types (each simulated as a $1 \times 1 \times 0.5 \text{ m}^3$ box centered in the container) with an otherwise identical configuration. (a) shows this distribution without normalization while (b) shows this distribution after normalizing for the number of events that pass the filtering stage.

The next question of interest is whether these hits degrade the usefulness of the muon hits. This can be evaluated by assessing the number of hits within a single hodoscope that result from muons and secondaries, respectively, within each event. The default assumption is that three hits from a muon, one per hodoscope plate, can be used to produce a muon track direction through that hodoscope. As is illustrated in Fig. 9, the presence of more than one secondary hit can limit the success of this approach, as identifying which hits correspond to the muon track becomes more complicated. Note, these schematics make two assumptions that are not always true. First, that just one secondary results in the production of hits in a single hodoscope and secondly that the secondary is produced by the muon while passing through the hodoscope (as oppose to the container or cargo). These are simplifications for the purpose of visualization, all such situations are considered in the analysis.

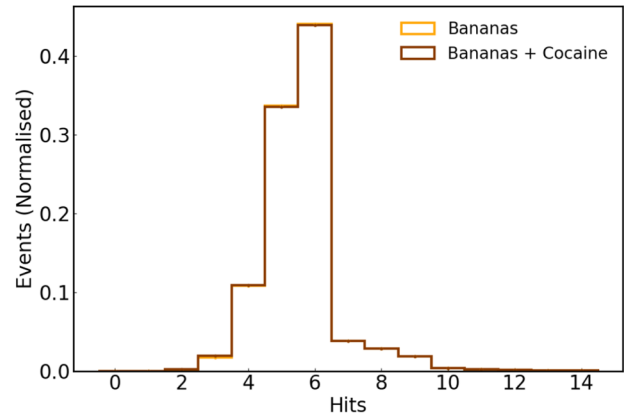


FIG. 8. The total number of hits recorded across all hodoscopes normalized for the number of events that pass a filtering stage. Two cargo scenarios are shown, in the first the container is filled with bananas, in the second a $40 \times 40 \times 40 \text{ cm}^3$ box of cocaine is placed in the center of the bananas.

A confusion matrix visualizes this impact in Fig. 10. The matrix is column normalized such that the values represent the proportion of times in which y secondary hits are recorded given that x muon hits are recorded. In this version of the figure, the container is empty of cargo (the cargo is “Air”). We see that in 95% of occasions in which three muon hits are recorded in a single hodoscope, no secondary hits are recorded. A further 2% of those occasions have just a single secondary hit, which still allows for identification of this hit as the outlier. Therefore, we categorize 3% of the occasions with three muon hits as having potential for muon tracking degradation due to the presence of secondary hits. If this value is kept low, then such contaminated events can just be removed from consideration without significant impact to the scanner’s sensitivity.

We express the contamination proportion (CP) as

$$\text{CP} = \frac{\text{Hodoscopes with 3 muon hits and } \geq 2 \text{ secondary hits}}{\text{Hodoscopes with 3 muon hits}}. \quad (2)$$

We can study how this number varies in a number of different scenarios. Figure 11 shows that the contamination proportion is unaffected when we vary among the different cargo varieties. As with previous studies to this point, a threshold of 0.4 MeV was used, with all lower energy hits assumed to be unobservable. Figure 12 shows how the contamination proportion varies with respect to this threshold. It is clear that as the threshold is reduced, the contamination effect increases, as would be expected given the distributions presented in Fig. 6. However, the effect remains <6% for all scenarios and is therefore unlikely to drastically affect the SilentBorder scanner performance even with a much lower sensitive threshold.

Finally, we consider the scenario in which the hodoscope is produced with just two plates not three. In this scenario, just a single secondary hit will result in a confusion effect and so the

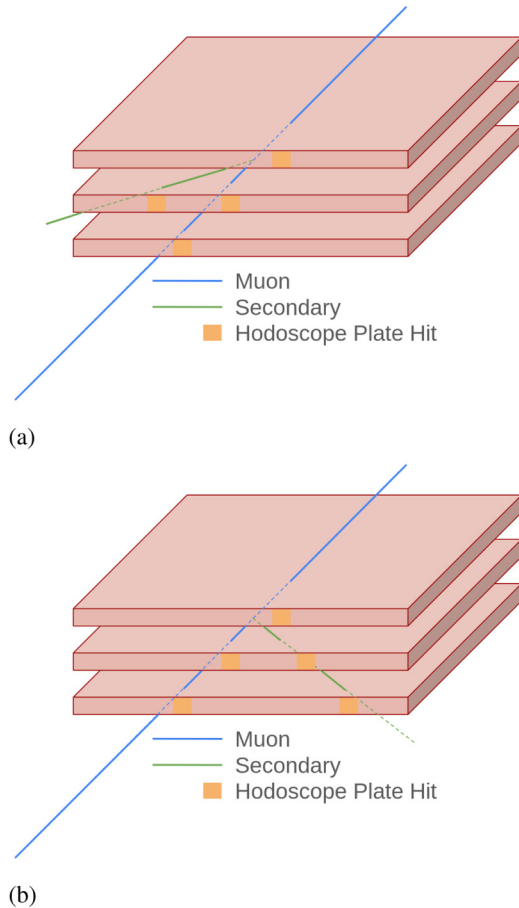


FIG. 9. Two schematic examples of a muon passing through the three plates in a single hodoscope. In (a), the emitted secondary produces just a single hit and the muon trajectory is still clear, in (b) the emitted secondary produces two hits and the muon track is much harder to discern.

metric becomes

$$CP = \frac{\text{Hodoscope with 2 muon hits and } \geq 1 \text{ secondary hits}}{\text{Hodoscope with 2 muon hits}}. \quad (3)$$

Figure 13 shows that while the contamination effect increases as expected, it is still at a very low level.

In conclusion, it is clear that with the current SilentBorder scanner setup the impact of secondary hits is very limited. The variation in the number of hits recorded when different cargoes are stored is too small to be used as a handle to investigate cargo compositions and the degradational impact on the muon tracking for MTS appears to also be very small. These studies are constrained to the impact of secondaries in a simple SilentBorder-like setup, there is significant other work in the realm of using secondary particles to add complementary information to muon tomography




FIG. 10. The number of hits (split into muon hits and secondary hits) recorded within a single hodoscope for a geometry containing an empty container. The matrix is column normalized, such that the values represent the proportion of times y secondary hits are recorded given that x muon hits are recorded.

techniques. Many of these use specialized separate detector elements to better detect secondary particles.^{14–17}

V. TomOpt

A. TomOpt pipeline

TomOpt¹¹ is a software tool developed for task-specific and constraint-aware optimization of the geometries and design specifications of detectors used for muon scattering tomography. Its modular Python¹⁸ architecture features a fully differentiable,




FIG. 11. The contamination proportion for a variety of cargo types. The entries on the left of the line show different cargoes each simulated as a $1 \times 1 \times 0.5 \text{ m}^3$ box centered in the container. The entries on the right of the dashed line show the more “realistic” scenarios where first a container is filled with bananas and then in the second scenario a $40 \times 40 \times 40 \text{ cm}^3$ box of cocaine is placed in the center of the bananas.




FIG. 12. The dependence of the contamination proportion on the energy threshold for a variety of cargo types (each simulated as a $1 \times 1 \times 0.5 \text{ m}^3$ box centered in the container) with an otherwise identical configuration.

end-to-end pipeline that employs differentiable programming through utilizing the automatic differentiation capabilities of the leveraged PyTorch library.¹⁹ The optimization pipeline seamlessly integrates the various stages, including muon generation, muon detection, and their propagation in scanned volumes leading to volume reconstruction, loss evaluation, and gradient backpropagation with respect to the detector design parameters. The detector parameters are then updated through gradient descent in the direction that minimizes the loss. A schematic representation of this pipeline is shown in Fig. 14, illustrating the main stages in the process, which are further detailed below. The modularity of the simulation chain allows for tailored user-specific implementations of muon source and propagation models, detector parameterization, volume reconstruction methods, loss function definitions, etc.




FIG. 13. The contamination proportion for a variety of cargo types (each simulated as a $1 \times 1 \times 0.5 \text{ m}^3$ box centered in the container). The hodoscopes are simulated with either two (hollow bars) or three (solid bars) scintillator plates.




FIG. 14. TomOpt optimization pipeline, as described in Ref. 11. Reproduced with permission from Strong *et al.*, 5(3), 035002 (2024). Copyright 2024 IOP Science; licensed under a Creative Commons Attribution (CCBY 4.0) license.

Prior to optimization, detectors positioned above and below a volume of interest—referred to as the passive volume—are instantiated with sub-optimal parameters. These parameters can be the position in xyz , dimension, efficiency, etc. The passive volume is assigned a voxelized ground truth profile of radiation length (X_0). The optimization then runs over a specified number of epochs, where the detector parameters are updated at the end of each iteration. For each epoch, generated muons are sampled from defined flux models taken from the literature^{20,21} and propagated through the detectors and passive volume setup. The scattering model reported by the Particle Data Group⁶ is used for the simulation of the multiple Coulomb scatterings of muons in the passive volume and the calculation of their scattering angles and spatial displacements. The resolution and efficiency of the detectors are modeled by a differentiable surrogate function, represented as a double-sigmoid (Fig. 15). This function achieves maximum resolution and efficiency within the active area of the detector and smoothly decreases while transitioning outside the detector boundaries. Each muon hit in the detectors then is assigned an efficiency and a spatial resolution in xy which is used to obtain the corresponding reconstructed hit with simulated detector response effect. This is done by sampling uncertainties in x and y from Gaussian distributions centered at the true position and of standard deviations equal to the corresponding spatial resolutions. Hits outside the detector region are still recorded but with degraded efficiency and resolution, which is a necessary feature to maintain the differentiability of the subsequent volume reconstruction with respect to the detector parameters. This allows for gradient descent to optimize the configuration which records the hits in an optimal way that yields best performance of the studied task. The hits above and below the passive volume are then fitted into incoming and outgoing tracks, respectively, using an analytical likelihood maximization method

05 December 2025 11:58:21




FIG. 15. Example of the surrogate model used to describe resolution and efficiency variations along the x -direction. For illustrative purposes, it is shown in one dimension (along the x axis), but in practice, it is applied independently along both the x - and y -directions.

accounting for the uncertainties on the x and y positions. The Point-of-Closest-Approach algorithm is a common method used to reconstruct the scattering vertices and obtain a 3D density distribution of these vertices. This approach computes the closest point between the incoming and outgoing muon tracks in 3D space, which corresponds to the midpoint of the common perpendicular to the two tracks. The algorithm assumes that the entirety of the multiple Coulomb scatterings of a muon occurs at this reconstructed point. Subsequent inference methods utilize the spatial and/or angular variables of the reconstructed PoCA points to infer the physical properties of the passive volume under investigation. The discrepancy between the inferred, or predicted, passive volume quantities and their true values is then used to compute a task-specific loss function. Additionally, an optional term reflecting the budget constraints of the detector, calculated based on the unit surface area cost, may be included in the loss function. The loss is then back-propagated through the reconstruction pipeline, and its gradients are calculated with respect to the detector parameters. The detector optimization problem is formulated as a minimization task to be solved through gradient descent. This technique iteratively updates the detector parameters a in the direction that minimizes the loss function,

$$a_{n+1} = a_n - \gamma \cdot \nabla_a \mathcal{L}(a_n), \quad (4)$$

where γ is a learning rate of the parameter and $\nabla_a \mathcal{L}(a_n)$ is the gradient of the loss function with respect to a . This optimization cycle is repeated for multiple epochs until convergence is attained and the detector setup reaches its final optimized configuration.

B. Volume reconstruction

1. Radiation length estimation

The primary inference method in TomOpt estimates the radiation length (X_0) of the voxelized passive volume by leveraging

reconstructed PoCA variables. Each PoCA vertex is modeled by a probability distribution function using three independent Gaussian distributions in the x , y , and z dimensions, which expresses its probability to occur within a voxel. The Gaussian distributions \mathcal{N} are characterized by their mean μ and standard deviation σ , where the mean corresponds to the nominal reconstructed position of the PoCA, and the standard deviation represents the uncertainty on that position within each respective dimension, computed by automatic differentiation,

$$\text{PoCA pdf} = \prod_{i \in \{x,y,z\}} \mathcal{N}(x_i; \mu_i, \sigma_i). \quad (5)$$

The probabilistic modeling of PoCA vertices allows the influence of each reconstructed PoCA point to extend across multiple voxels within the passive volume. For each muon batch, each voxel is assigned a probability for every muon's PoCA point of its being located within that voxel. These probabilities are used as weights to compute a weighted root-mean-square (RMS) of the scattering angle distribution inside each voxel, taking into account all PoCA points in the muon batch. This RMS value is used to estimate the scattering behavior within that voxel. By inverting Eq. (1), the RMS of the scattering angles is employed to infer the radiation length X_0 of the material in each voxel. Although the inferred radiation lengths exhibit a systematic bias toward lower values,²² the resulting 3D image effectively distinguishes between high- Z and low- Z voxels, providing a clear representation of the material properties within the scanned volume (Fig. 16).

2. A clustering density inference

Another volume inference method that we adopt in TomOpt is based on a binned clustering algorithm²³ that groups reconstructed scattering vertices, i.e., PoCA points in this case, into clusters indicating different materials in a voxelized volume. This clustering is done based on the spatial proximity and scattering angles of the scattering vertices. Since materials with different densities produce different scattering distributions, the properties of clustered PoCA points can reveal the density profile of the volume and determine high- and low-density regions.

Within each voxel, PoCA points are ranked in descending order of their scattering angles, and a fixed subset is selected to ensure uniform treatment across voxels. For each pair of selected PoCA points, their distance metric is computed and scaled by the inverse of the product of their scattering angles and normalized momenta. The normalized momentum is defined as

$$\tilde{p} = \frac{p}{p_{\text{norm}}}, \quad (6)$$

where p_{norm} is the nominal average cosmic muon momentum at sea level. Since most muon tomography setups lack direct momentum measurements for practical reasons, we assume a fixed momentum of 5 GeV for all muons, effectively setting $\tilde{p} = 1$. The median of the natural logarithm of the per voxel weighted metric distribution serves as an indicator of material density, with lower




FIG. 16. 3D volume inference using the radiation length estimation method. The scanned volume consists of a $1 \times 1 \times 0.5 \text{ m}^3$ air-filled region, into which four material blocks—iron, uranium, aluminum, and beryllium—each measuring $30 \times 30 \times 20 \text{ cm}^3$, are inserted. These blocks span the second and third slices along the z axis (denoted as z layer 1 and z layer 2). The right column presents cross-sectional heatmaps of the ground truth radiation lengths, while the left column shows the corresponding inferred values. Radiation lengths are displayed in a normalized natural logarithmic scale to enhance contrast between materials with similar densities.

values corresponding to regions where scattering vertices are both larger and closer together. The natural logarithm is applied to improve the separation between closely clustered values. This median value defines the corresponding voxel score, resulting in a per volume distribution of voxel scores. The algorithm subsequently extracts the minimum score value, which serves as a discriminator for identifying the presence of a uranium block in the volume by applying a threshold. While this binary classification

approach has been explored in previous work,²⁴ the present study shifts focus toward the 3D image reconstruction of the passive volume (Fig. 17), utilizing the full voxel score distribution. The previously mentioned extended PoCA pdf in Sec. VB 1 is now used, eliminating the need to select a subset of the largest scattering




FIG. 17. 3D volume inference using the binned clustering algorithm (BCA) method. The scanned volume consists of a $1 \times 1 \times 0.5 \text{ m}^3$ void (air), containing four inserted blocks of different materials: iron, uranium, aluminum, and beryllium. Each block measures $30 \times 30 \times 20 \text{ cm}^3$ and spans the second and third layers along the z axis (labeled z layer 1 and z layer 2). The right column displays cross-sectional heatmaps of the ground truth radiation lengths in the xy -plane at each z layer, visualized in a normalized natural logarithmic scale to enhance contrast between materials of similar density. The left column shows the corresponding inferred BCA voxel scores obtained using the BCA method, represented in a normalized linear scale. Note that the BCA algorithm inherently incorporates a logarithmic transformation, hence no additional log-scaling was applied to the inferred values.

vertices. This allows for a more efficient voxel population, reducing the number of muon events typically required to achieve uniform coverage.

Figure 18(a) illustrates the relationship between the BCA voxel score and the ground truth value of the corresponding radiation length (X_0) for different materials, ranging from low- to high-Z materials. The statistics are derived from 300 scans of a single voxel, distributed across the range of materials shown. The BCA score inference is performed using the extended PoCA version of the algorithm, with 1000 sampled muons. A strong linear correlation is evident when considering only the high-to-mid Z range, while this linearity breaks down for low-Z

materials, such as air. The natural logarithm of the X_0 values is used on the x axis, as the BCA score calculation inherently involves the natural logarithm. In comparison, a similar plot [Fig. 18(b)] showing the relationship between inferred and ground truth x_0 using the previous inference method reveals a significant drop in inferred values for air relative to the true value. This discrepancy underscores the limitations of the previous approach, particularly in accurately capturing low-Z materials. This relationship supports the use of the BCA inference as a reliable method for optimization, by minimizing the mean squared error (MSE) between the BCA predictions and the corresponding X_0 values (in natural logarithmic scale), both of which are normalized, to define the loss function.

VI. BAYESIAN OPTIMIZATION WITH TomOpt

A. The Ax library

While TomOpt's standard optimization approach relies on gradient descent, we investigate Bayesian Optimization (BO)²⁵ as an alternative due to challenges posed by the volume inference methods. In particular, the presence of outliers in the inferred volume properties—especially in the X_0 inference approach—leads to a noisy mean squared error (MSE) loss function. Although gradient-based optimization is, in most cases, capable of improving detector coverage by optimizing the spatial positions of the hodoscopes in simple scenarios, such as configurations with a few hodoscopes, the sharp fluctuations in the MSE loss during optimization prevent the final configuration from being considered a fully optimized solution. This issue becomes more pronounced when extending the parameter space to include non-trivial design factors, such as the relative spacing of panels within the hodoscopes.

While ongoing efforts focus on mitigating inference biases, we present here a BO framework as a robust alternative for cases where the loss function is inherently noisy.

Bayesian Optimization (BO) is a powerful approach for optimizing black-box functions that are expensive to evaluate, lack a known analytical form or exhibit noise and uncertainty. Unlike gradient-based optimization methods, which require the objective function to be differentiable, BO constructs a probabilistic surrogate model of the objective function that approximates the function and guides the search for the optimal solution. At the core of BO is Bayes' theorem which states that—given prior knowledge about a function f and newly observed data \mathcal{D} , the posterior distribution of f is proportional to the likelihood of \mathcal{D} given f multiplied by the prior distribution of f ,²⁶

$$p(f | \mathcal{D}) \propto p(\mathcal{D} | f)p(f). \quad (7)$$

A common choice for the surrogate model in BO is a Gaussian Process (GP), a flexible, non-parametric model that provides a probabilistic estimate of the objective function. A GP assumes that the function values at different points are correlated and defines a prior distribution over possible functions. Given a set of observed data points, the GP updates its belief about the function by computing a posterior distribution, which provides both a mean prediction (expected function value) and a variance

05 December 2025 11:58:21

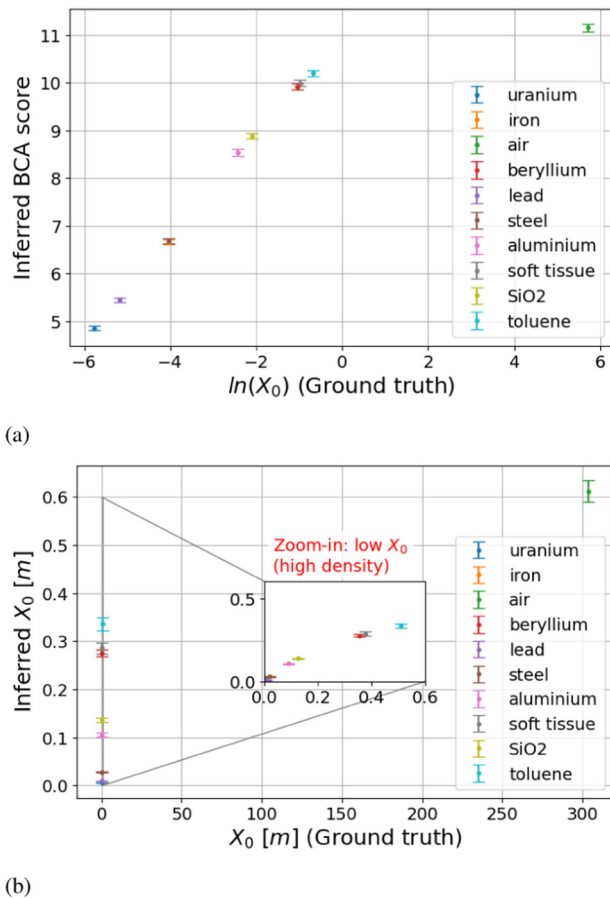


FIG. 18. Relationship between inferred quantities and ground truth X_0 for different materials. (a) shows the correlation between BCA voxel scores and the natural logarithm of the corresponding radiation length. A strong linear trend is observed for mid-to-high-Z materials, while the correlation weakens for low-Z materials such as air. The natural logarithm of X_0 is taken to ensure a valid comparison with the BCA score, as the latter inherently includes a logarithmic transformation in its metric. (b) compares inferred and true X_0 values, demonstrating a significant drop in inferred values observed for air, while higher-Z materials exhibit a near-linear relationship. A zoomed-in view reveals the enhanced linearity at higher Z.

estimate (uncertainty in the prediction). This uncertainty measure is particularly advantageous in BO, as it enables informed decisions about where to evaluate the function next. The selection of evaluation points is determined by an acquisition function, which defines a utility function from the model posterior that determines the next point to evaluate.²⁶ The selection is performed in a way that strategically balances exploration (sampling uncertain regions) and exploitation (sampling regions with promising performance). One widely used acquisition function is Expected Improvement (EI), which quantifies how much improvement we expect by sampling at a particular point. In the case of minimizing the objective function, the improvement at a new point x is defined as the positive difference between the current best observed minimum f^* and the predicted function value $f(x)$, if the new point offers an improvement,

$$I(x) = \max(f^* - f(x), 0), \quad (8)$$

where $f^* = \min_{i=1,\dots,n} f(x_i)$ is the lowest observed function value so far and $f(x)$ is the function value at the candidate point x . The EI is then computed by taking the expectation of $I(x)$ over the posterior distribution of $f(x)$, modeled by the GP,²⁷

$$\mathbb{E}[I(x)] = \mathbb{E}[\max(f^* - f(x), 0)]. \quad (9)$$

A higher EI indicates that sampling at that point is expected to yield a larger improvement over the current best, making it a good candidate for the next evaluation. After evaluating a new point, the surrogate model is updated iteratively, refining the search for the global optimum while minimizing the number of function evaluations.

Ax²⁸ is an open-source platform for adaptive experimentation and Bayesian Optimization, developed by Meta. It provides a high-level interface for optimizing parameters without requiring explicit definition of the acquisition function. Ax internally leverages BoTorch,²⁹ a library for Bayesian Optimization, and automatically selects an acquisition function based on the optimization scenario.

In this work, BO is implemented using the Ax Loop API, which provides an automated framework for sequential BO. The optimization pipeline consists of defining the search space, constructing and updating the surrogate model, and iteratively selecting and evaluating parameter configurations. The search space defines the range of possible values for the parameters being optimized, in our case the xyz positions of hodoscopes. This parameter space is passed to the experiment object, which acts as the central controller for the optimization process. The experiment object in Ax manages the entire optimization process, including the collection of observations, selection of new trials, and updating of the surrogate model. It stores all configurations tested so far along with their corresponding objective function values. Evaluation of each candidate set of parameters is done through a user-defined evaluation function that quantified the performance of the parameters. We take the MSE of the volume inference as our evaluation function, which quantifies the accuracy of the voxelized reconstructed material distribution and directly relates to the ability to discriminate high-density or suspicious objects in

border security applications. BO with Ax's Loop API is conducted using a hybrid sampling strategy that consists of two sequential phases:

- Sobol generation strategy:

The first 12 trials are evaluated using a low-discrepancy quasi-random Sobol sequence that ensures uniform coverage of the search space. These initial evaluations serve to bootstrap the Gaussian Process model by providing an initial data set that prevents premature convergence to a local optimum. Sobol sampling is preferred over purely random initialization as it minimizes clustering and ensures better space-filling properties.

- Gaussian Process Regression (GPEI):

After the initial Sobol trials, Ax switches to Bayesian Optimization generation strategy using a Gaussian Process surrogate model prior of the objective function coupled with an Expected Improvement acquisition function (GPEI).

This hybrid strategy ensures efficient exploration of the parameter space while minimizing the number of function evaluations required for convergence. At the end of the optimization process, the best parameter configuration is extracted based on the lowest observed objective function value.

For our purpose, we integrate Ax's BO framework with TomOpt, coupling the BO optimization pipeline to the TomOpt simulation pipeline while bypassing the gradient-based optimization module. Instead of using gradient descent to optimize detector parameters, we employ Ax's BO strategy, where each evaluation trial consists of updating the current detector parameter set, running a full muon tomography simulation with the updated detector configuration, and performing volume inference. These inferred values are then used to evaluate the surrogate objective function through a MSE evaluation function, providing feedback for the next trial.

Since the MSE loss function is sensitive to small errors and exhibits noise, BO is particularly useful for optimizing our detector parameters. The surrogate model in BO allows us to smooth out fluctuations in the loss landscape, leading to more reliable parameter tuning. Additionally, using BO enables a controlled comparison of our two inference methods: radiation length inference and the binned clustering algorithm. By setting aside gradient-based updates, we can assess how each inference method performs under the same optimization framework, providing insights into their stability and effectiveness in reconstructing material properties from muon scattering data.

We use a passive volume with dimensions $1 \times 1 \times 0.5 \text{ m}^3$, simulated as a lorry with an iron base and filled with a random arrangement of $10 \times 10 \times 10 \text{ cm}^3$ iron, beryllium, and air blocks to increase scenario variability and improve the generalizability of the results. Additionally, a 2000 cm^3 uranium block is randomly placed inside the lorry with a probability of 0.5 of being present. The initial detector setup consists of a hodoscope pair, each measuring $1.5 \times 1.5 \times 0.4 \text{ m}^3$, placed above and below the passive volume. The maximum of the surrogate for spatial resolution is 0.1 mm, corresponding to the uncertainty on a single muon hit position, as in TomOpt each muon is generated and processed individually.

With current computational resources, TomOpt comfortably handles objects of this scale. For significantly larger volumes, longer exposure times are required, translating into a larger number of simulated muons and correspondingly bigger PyTorch tensors storing both muon and voxel information. This can become a bottleneck if computations are run solely on CPU. While GPU acceleration can alleviate part of this burden, scaling to much larger objects (e.g., tens of meters) remains a practical challenge rather than a conceptual one. Initially, hodoscopes are offset by 0.5 m in both x - and y -directions, leaving part of the passive volume uncovered (Fig. 19).

The reason for placing the hodoscopes at different z positions is to avoid overlap along the z axis. This is due to the current implementation of hodoscopes in TomOpt, where each hodoscope is defined as a single unit consisting of individual detector plates. The fitting of hits into tracks accounts for hits in all detector plates, where hits outside an active plate surface are down-weighted through lower resolution and efficiency. Hence, placing two plates at the same z level would cause ambiguities in track reconstruction. While this constraint does not prevent the optimisation study focussing on inference methods, it does limit the ability to model perfectly co-planar panels or more complex detector layouts, where modules intentionally share the same z -position. This panel overlap limitation is a technical constraint of the current version and is being addressed in an updated implementation that will be presented in a forthcoming paper.

Starting from this suboptimal configuration, optimization is expected to converge toward an optimal final setup where the hodoscopes perfectly cover the passive volume. BO is then executed for 100 evaluation trials. Before each evaluation, the detector parameters are set to the values sampled for evaluation.

B. Optimization with the BCA inference method

In this section, BO is conducted using the BCA inference method for volume reconstruction. The surrogate objective function is evaluated at each trial by computing the MSE between inferred BCA voxel scores and the logarithmic ground truth X_0 , both normalized to unity. This formulation captures the




FIG. 19. The initial configuration of the hodoscope setup in XZ (left) and YZ (right) views.




FIG. 20. Convergence of the best minimum objective (ideal optimum = 0) across 100 BO trials using BCA inference.

underlying correlation between the two quantities, as described in Fig. 18(a).


Figure 20 shows the convergence of the best minimum objective over 100 trials on CPU, showing a 13.3% reduction throughout the optimization. For simplicity, only a single run is shown, although different random seeds may yield different results. The full optimization takes ~ 26 min, while using gradient descent would result in $\sim 38\%$ increase in runtime. This difference arises because the BCA reconstruction is computationally heavier than the formula-based X_0 inference method, involving PoCA pairwise voxel clustering. The gradient descent optimization requires evaluating many gradients per step, resulting in longer runtime. With BO, the GP model reduces the number of expensive evaluations, so it is faster overall despite the overhead of building the GP model.

Figure 21 presents the response surface for the x positions of the two upper hodoscopes, illustrating both the mean and standard error of the objective function's predictive model. The response surface provides insight into how the objective function varies across the parameter space and highlights regions of lower expected error.


Figure 22 displays the evolution of the gap size between the upper hodoscopes, as determined by the evaluated x coordinates throughout the optimization trials. The marker color scale in this figure represents the observed (evaluated) mean of the objective function. The results demonstrate a clear trend in the optimization process, with BO guiding the hodoscope placement toward regions associated with lower objective values. The response surface and gap evolution plots collectively validate the surrogate model's effectiveness in navigating the noisy objective landscape. Figure 23 presents the final optimized configuration of the detector setup, corresponding to the optimal parameter values identified by the BO process.

C. Optimization with the X_0 inference method

In this approach, BO was performed using the X_0 inference method, where the evaluation function was defined as the mean squared error (MSE) between normalized inferred and ground truth X_0 values. Normalization was necessary due to



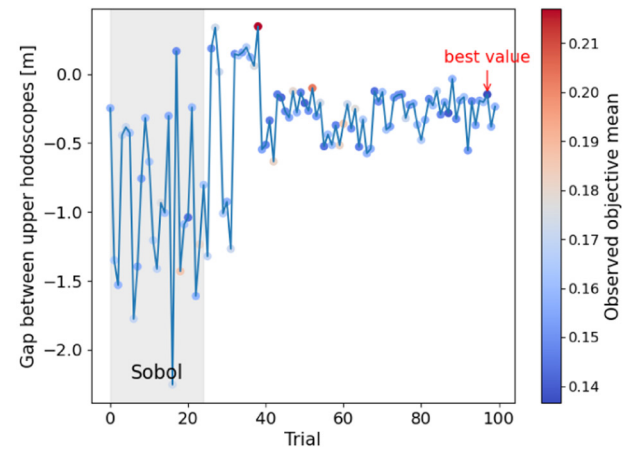
(a)




(b)

FIG. 21. Response surface of the objective mean (a) and standard error (b) for the x positions of the upper two hodoscopes, labeled x_1 and x_2 . The objective function being evaluated is a MSE that captures the correlation between inferred BCA voxel score and true voxel X_0 of the scanned volume.

the significant underbias observed for low- Z materials, particularly air [Fig. 18(b)]. Since air is present in the passive volume, the raw MSE values would otherwise be orders of magnitude larger than what is expected from a typical MSE evaluation, potentially distorting the optimization process. Figure 24 shows the convergence of the best minimum objective, with a total reduction of 22.8%. The complete optimization runs ~ 11 min on CPU with BO, comparable to ~ 9 min with gradient descent. Here, the formula-based X_0 inference method is relatively computationally inexpensive, so the additional runtime in BO stems primarily from the overhead of updating the GP surrogate model. Figure 25 illustrates the response surface for the mean



(a)



(b)

FIG. 22. The evolution of the gap size between upper hodoscopes (a) and between lower hodoscopes (b) across 100 trials during the BO process utilizing the BCA inference method. The color scale of the markers corresponds to the observed (evaluated) mean of the objective. For the first 24 trials, sampled using the Sobol strategy, a broader range of gap sizes is explored. Negative values indicate an overlap or swapping of the hodoscope pair along the x -direction. A trend is observed for the upper gap, where the gap size tends toward zero, suggesting minimal overlap. The lower gap shows a preference for overlap roughly equivalent to half the size of a hodoscope. The best gap values, corresponding to the optimal parameter configurations, are annotated in the plot. Notably, these values align with the minimum observed mean of the objective function.

and standard error of the surrogate objective. Unlike the case with BCA inference, the gap sizes between the upper and lower hodoscope pairs (Fig. 26) do not exhibit a clear optimization trend, with the best-evaluated parameters appearing mid-range in the trials and the observed mean of the objective function remaining relatively stable throughout the GPEI stage following Sobol sampling. This behavior can be attributed to the

05 December 2025 11:58:21




FIG. 23. The BO optimized configuration of the hodoscope setup in XZ (left) and YZ (right) views. The placement of the hodoscopes aligns with the observations in Fig. 22, where the upper hodoscopes provide fuller coverage, while the lower hodoscopes exhibit overlap while maintaining sufficient coverage of the passive volume.

noisiness of the X_0 MSE, which arises from the inherent bias in the X_0 inference method.

Although the X_0 inference approach still led to an optimal final detector configuration (Fig. 27), achieving full coverage placement, its noisier optimization landscape suggests that it may not be as reliable for extending BO to more complex optimization problems. If the goal is to optimize additional detector parameters beyond basic placement, using a more stable inference method, such as BCA, may be necessary for efficient and interpretable optimization.

To statistically compare the inference performance of the optimized detector configurations, we evaluate the Mean Squared Error (MSE) and the Structural Similarity Index Measure (SSIM) for 100 simulated passive volumes, with and without a high-density uranium block (Fig. 28). The performance is compared with a baseline corresponding to a detector design based on human intuition for maximizing acceptance (Fig. 29). While the MSE directly quantifies voxel-wise differences, the SSIM is a perceptual metric that evaluates the structural similarity between predicted and true






FIG. 24. Convergence of the best minimum objective across 100 BO trials using X_0 inference.



(a)



(b)

FIG. 25. Response surface of the objective mean (a) and standard error (b) for the x positions of the upper two hodoscopes, labeled x_1 and x_2 . The objective function being evaluated is the MSE between the normalized inferred and true voxel X_0 values of the scanned volume.

volumes. It is computed over local regions using a sliding window of size 4×4 , and the final SSIM score is obtained by averaging over all such windows. The SSIM is defined as

$$\text{SSIM}(\text{infer}, \text{true}) = \frac{(2\mu_{\text{infer}}\mu_{\text{true}} + C_1)(2\sigma_{\text{infer},\text{true}} + C_2)}{(\mu_{\text{infer}}^2 + \mu_{\text{true}}^2 + C_1)(\sigma_{\text{infer}}^2 + \sigma_{\text{true}}^2 + C_2)}, \quad (10)$$

where μ_{infer} and μ_{true} are the local means, σ_{infer}^2 and σ_{true}^2 are the local variances, and $\sigma_{\text{infer},\text{true}}$ is the local covariance of the inferred and ground truth volumes, where the inferred quantities may





FIG. 26. The evolution of the gap size between upper hodoscopes (a) and between lower hodoscopes (b) across 100 trials during the BO process utilizing the X_0 inference method. The color scale of the markers corresponds to the observed (evaluated) mean of the objective. For the first 24 trials, sampled using the Sobol strategy, a broader scope of gap sizes is explored. Negative values indicate an overlap or swapping of the hodoscope pair along the x-direction. Unlike the results obtained using BCA inference, the optimal parameters are not located near the final trials but instead appear in the mid-trial range. Additionally, the various trial evaluations yield comparable values of the observed objective mean, as indicated by the predominance of blue-colored markers throughout the GPEI stage following Sobol sampling.

represent radiation lengths or BCA scores depending on the inference method, and the ground truth is the true X_0 layout of the volume. The constants C_1 and C_2 are used to stabilize the division when the denominators are small. Higher SSIM values (closer to 1) indicate stronger structural agreement, while lower MSE values indicate higher accuracy. Overall, the BCA method exhibits




FIG. 27. The BO optimized configuration of the hodoscope setup in XZ (left) and YZ (right) views.

superior inference performance compared to the radiation length estimation (X_0) method, both in terms of SSIM and MSE. This suggests that the BCA method is more robust to noise in complex or high-contrast configurations.

Some of the remaining limitations of TomOpt arise from intrinsic challenges in muography: the inverse problem is inherently ill-posed and relies on approximations, with the PoCA method reducing multiple scattering to a single point and




FIG. 28. Statistical distribution of volume inference performance using MSE (left axis) and SSIM (right axis), across the initial, optimized, and baseline detector configurations with the BCA and X_0 inference methods. Each boxplot summarizes the distribution of inference quality over 100 simulated passive volumes: 50 volumes containing blocks of iron, beryllium, and air, and 50 volumes with an additional randomly placed uranium block. In each boxplot, the central line indicates the median; the box edges represent the first and third quartiles (Q1 and Q3); and the whiskers extend to the most extreme data points within 1.5 times the inter-quartile range (IQR). For SSIM, higher values (closer to 1) indicate better structural similarity, whereas for MSE, lower values indicate better inference accuracy. Using the radiation length estimation method, a slight improvement is observed in MSE with the optimized configuration, while SSIM remains similar across all configurations. In contrast, the BCA method shows consistent improvement in both MSE and SSIM when using the optimized detector design, comparable to the baseline detector.

05 December 2025 11:58:21




FIG. 29. The configuration of the baseline detector setup in XZ (left) and YZ (right) views.

introducing systematic biases. Importantly, nearly all current reconstruction methods build on PoCA, meaning that the sensitivity of detector design to volume reconstruction is closely tied to the accuracy of this approximation, highlighting the importance of developing approaches that go beyond PoCA in future work.

VII. CONCLUSIONS

This work outlines a clear pathway toward the optimization of a future SilentBorder scanner for muon scattering tomography, with current results demonstrating meaningful advances in the studied approaches. Future studies will aim to extend these developments toward more realistic operational scenarios and validation with experimental data. Studies using the GEANT4 framework have shown that minimizing the in-plane gaps between active hodoscope areas, by keeping any casing and electronics sizes as small as possible, will ensure optimal efficiency of the system. However, increasing the vertical spacing between hodoscope plates does show improvement in the angular resolution with little reduction in efficiency and is worth pursuing. A complementary GEANT4 study on the role of secondary particle hits in event reconstruction was also conducted. Our findings suggest that secondary particles do not provide any extra material discrimination power in a “SilentBorder” type scanner but also cause very limited degradation of the reconstruction quality and, therefore, do not need separate measures to exclude them.

On the machine learning side, a Bayesian optimization (BO) module was developed within the TomOpt framework to enable robust design optimization for reconstruction methods affected by noise. This module complements the existing gradient-based approach and is particularly well-suited for scenarios, where the loss landscape is poorly defined or convergence is difficult to achieve. The BO strategy was applied to two distinct reconstruction algorithms: a root-mean-square (RMS)-based X_0 inferrer and a binned clustering-based density inferrer. Results indicate that the RMS-based method is more sensitive to noise, whereas the binned clustering approach exhibits greater robustness. Although BO-optimized configurations demonstrated performance improvements, they did not significantly outperform a baseline detector

configuration derived from human intuition. This motivates further investigation of deep learning-based approaches for volume inference to enhance sensitivity to subtle updates of detector parameters. This could be pursued either by coupling existing reconstruction outputs with post-processing neural networks or by developing fully end-to-end learning frameworks that directly infer volumes from minimally processed detector-level scattering information.

ACKNOWLEDGMENTS

The authors acknowledge funding from the EU Horizon 2020 Research and Innovation Programme under Grant Agreement No. 101021812 (“SilentBorder”) and by the Fonds de la Recherche Scientifique—FNRS under Grant Nos. T.0099.19 and J.0070.21. Pietro Vischia’s work was supported by the “Ramón y Cajal” program under Project No. RYC2021-033305-I funded by the MCIN MCIN/AEI/10.13039/501100011033 and by the European Union NextGenerationEU/PRTR. Finally, the authors gratefully acknowledge the computer resources at Artemisa, funded by the European Union ERDF and Comunitat Valenciana as well as the technical support provided by the Instituto de Física Corpuscular, IFIC (CSICUV).

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Z. Zaher: Investigation (equal); Methodology (equal); Software (equal); Writing – original draft (equal). **H. Lay:** Investigation (equal); Methodology (equal); Writing – original draft (equal). **T. Dorigo:** Supervision (equal); Validation (equal). **A. Giammanco:** Conceptualization (equal); Funding acquisition (lead); Methodology (equal); Project administration (equal); Supervision (lead); Validation (equal). **V. Gulik:** Formal analysis (equal). **C. Hrytsiuk:** Software (equal). **V. A. Kudryavtsev:** Conceptualization (equal); Funding acquisition (equal); Methodology (equal); Project administration (equal); Resources (equal); Supervision (lead). **M. Lagrange:** Investigation (equal); Validation (equal). **T. Metspalu:** Resources (equal). **G. C. Strong:** Conceptualization (equal); Methodology (equal); Software (lead). **C. Turkoglu:** Visualization (equal); Writing – original draft (equal). **P. Vischia:** Supervision (equal); Validation (equal).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

- ¹K. N. Borozdin *et al.*, “Radiographic imaging with cosmic-ray muons,” *Nature* **422**, 277 (2003).
- ²R. D’Alessandro, F. Ambrosino, G. Baccani, L. Bonechi, M. Bonghi, A. Caputo, R. Ciaranfi, L. Cimmino, V. Ciulli, M. D’Errico, F. Giudicepietro, S. Gonzi, G. Macedonio, V. Masone, B. Melon, N. Mori, P. Noli, M. Orazi, P. Passeggio,

- and L. Viliani, "Volcanoes in Italy and the role of muon radiography," *Philos. Trans. R. Soc. Lond. A* **377**, 20180050 (2018).
- ³L. W. Alvarez *et al.*, "Search for hidden chambers in the pyramids," *Science* **167**, 832 (1970).
- ⁴International Atomic Energy Agency, "Muon imaging: Present status and emerging applications," IAEA TECDOC 1212 (IAEA, Vienna, 2022).
- ⁵S. Barnes, A. Georgadze, A. Giammanco, M. Kiisk, V. A. Kudryavtsev, M. Lagrange, and O. L. Pinto, "Cosmic-ray tomography for border security," *Instruments* **7**, 13 (2023).
- ⁶S. Navas *et al.* (Particle Data Group), "Review of particle physics," *Phys. Rev. D* **110**, 030001 (2024).
- ⁷W. C. Priedhorsky, K. N. Borozdin, G. E. Hogan, C. Morris, A. Saunders, L. J. Schultz, and M. E. Teasdale, "Detection of high-Z objects using multiple scattering of cosmic ray muons," *Rev. Sci. Instrum.* **74**, 4294–4297 (2003).
- ⁸T. Dorigo *et al.* (MODE), "Toward the end-to-end optimization of particle physics instruments with differentiable programming," *Rev. Phys.* **10**, 100085 (2023).
- ⁹M. Aehle *et al.* (MODE), "Progress in end-to-end optimization of detectors for fundamental physics with differentiable programming," [arXiv:2310.05673](https://arxiv.org/abs/2310.05673) [physics.ins-det] (2023).
- ¹⁰S. Agostinelli *et al.*, "GEANT4—A simulation toolkit," *Nucl. Instrum. Methods Phys. Res. A* **506**, 250 (2003).
- ¹¹G. C. Strong *et al.*, "TomOpt: Differential optimisation for task- and constraint-aware design of particle detectors in the context of muon tomography," *Mach. Learn. Sci. Tech.* **5**, 035002 (2024).
- ¹²R. Chytrcek, J. McCormick, W. Pokorski, and G. Santin, "Geometry description markup language for physics simulation and analysis applications," *IEEE Trans. Nucl. Sci.* **53**, 2892–2896 (2006).
- ¹³T. Blackwell and V. Kudryavtsev, "Simulation study into the identification of nuclear materials in cargo containers using cosmic rays," *J. Instrum.* **10**, T04003 (2015).
- ¹⁴M. Pérez Prada, S. Barnes, and M. Stephan, "Analysis of secondary particles as a complement to muon scattering measurements," *Instruments* **6**, 66 (2022).
- ¹⁵G. Galgóczi, D. Mrdja, I. Bikit, K. Bikit, J. Slivka, J. Hansman, L. Oláh, G. Hamar, and D. Varga, "Imaging by muons and their induced secondary particles—A novel technique," *J. Instrum.* **15**, C06014 (2020).
- ¹⁶J. D. Bacon, K. N. Borozdin, J. M. Fabritius II, C. Morris, and J. O. Perry, "Muon induced fission neutrons in coincidence with muon tomography," Technical Report [Los Alamos National Laboratory (LANL), Los Alamos, NM, 2013].
- ¹⁷I. Bikit, D. Mrdja, K. Bikit, J. Slivka, N. Jovancevic, L. Oláh, G. Hamar, and D. Varga, "Novel approach to imaging by cosmic-ray muons," *Europhys. Lett.* **113**, 58001 (2016).
- ¹⁸Python Software Foundation, see <https://www.python.org/> for "Python programming language."
- ¹⁹A. Paszke *et al.*, "PyTorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems* 32 (Curran Associates, Inc., 2019), p. 8024.
- ²⁰M. Guan, M.-C. Chu, J. Cao, K.-B. Luk, and C. Yang, "A parametrization of the cosmic-ray muon flux at sea-level," [arXiv:1509.06176](https://arxiv.org/abs/1509.06176) [hep-ex] (2015).
- ²¹P. Shukla and S. Sankrith, "Energy and angular distributions of atmospheric muons at the earth," [arXiv:1606.06907](https://arxiv.org/abs/1606.06907) [hep-ph] (2018).
- ²²Z. Zaher *et al.*, "Optimisation of muon portals for border controls using TomOpt," in *MARESEC 2024* (Zenodo, 2024).
- ²³C. Thomay *et al.*, "A binned clustering algorithm to detect high-Z material using cosmic muons," *J. Instrum.* **8**, P10013 (2013).
- ²⁴Z. Zaher, S. Alvarez, T. Dorigo, A. Giammanco, M. Lagrange, G. C. Strong, P. Vischia, and H. Zaraket, "Optimisation of muon tomography scanners for border control using TomOpt," *Particles* **8**, 53 (2025).
- ²⁵J. Mockus, V. Tiesis, and A. Zilinskas, *The Application of Bayesian Methods for Seeking the Extremum* (Springer Nature, 1978), pp. 117–128.
- ²⁶E. Brochu, V. Cora, and N. Freitas, "A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning," *CoRR* abs/1012.2599 (2010).
- ²⁷D. R. Jones, M. Schonlau, and W. J. Welch, "Efficient global optimization of expensive black-box functions," *J. Global Optim.* **13**, 455–492 (1998).
- ²⁸Meta Open Source, see <https://ax.dev/> for information on Ax Adaptive Experimentation Platform.
- ²⁹M. 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," in *Advances in Neural Information Processing Systems* 33 (Curran Associates, Inc., 2020).