Optimisation of Muon Portals for Border Controls Using TomOpt

Zahraa Zaher[1], Maxime Lagrange[1], Samuel Alvarez[2], Giles C. Strong[3], Florian Bury[4], Tommaso Dorigo[3,5], Andrea Giammanco[1], Aitor Orio[6], Pietro Vischia[2], Haitham Zaraket[7,8]

TomOpt is a software package designed to optimize the geometric configuration and specifications of detectors intended for tomography using cosmic-ray muon scattering. Differentiable programming is utilized in this software to model muon interactions with detectors and scanned volumes, infer volume properties, and perform loss minimization in an optimization cycle. In this paper, we introduce the implementation in TomOpt of a case study related to cargo scanning at border control points, such as cargo harbors and truck inspection stations.

KEY WORDS
~ Muon tomography
~ Automatic differentiation
~ Detector design
~ Differentiable programming
~ Optimization
~ Muography

[1] Universite catholique de Louvain, Centre for Cosmology, Particle Physics and Phenomenology (CP3), Louvain-la-Neuve, Belgium
[2] Universidad de Oviedo and ICTEA, Oviedo, Spain
[3] Istituto Nazionale di Fisica Nucleare, Padova, Italy
[4] University of Bristol, Bristol, United Kingdom
[5] Lulea University of Technology, Lulea, Sweden
[6] Muon Tomography Systems S.L., Bilbao, Spain
[7] Lebanese University, Faculty of Science, Multi-Disciplinary Physics Laboratory, Optics and Fiber Optics Group, Beirut, Lebanon
[8] Universite Grenoble-Alpes, Laboratoire de Physique Subatomique et de Cosmologie, Grenoble, France
e-mail: zahraa.zaher@uclouvain.be
doi: 10.7225/toms.v14.n03.w01

Received: 21 Apr 2025 / Revised: 4 Jun 2025 / Accepted: 6 Jun 2025 / Published: 25 Jul 2025

This work is licensed under 

1. INTRODUCTION

High-performance computing and neural networks have revolutionized the optimization of complex systems, which can now use automatic differentiation to track the gradient of a task-specific objective function (Dorigo et al., 2023). This function, used to quantify the performance or outcome of a system with respect to a specific objective, maps the system's parameters to a quantity that represents the cost associated with this set of parameters. The optimal system's design is the one corresponding to the supremum of this objective function. This approach eliminates the need for brute-force scans, whose computational cost is very high, especially for systems whose parameter space is of large dimension. For such optimization tasks, it is necessary to construct a fully-differentiable end-to-end modelling of the flow that translates raw data to an inferred quantity that in turn is fed to the objective function. This allows for the backwards pass through the function's chain of gradients during the optimization cycle. This approach is not currently feasible for large multipurpose particle detectors such as the CERN LHC experiments, which involve complex interconnected systems of tens of thousands of detection elements reading out highly interdependent physical parameters. The complexity of this type of detection systems prevents the construction of an end-to-end differential pipeline from raw detector data to a measurement of a physical quantity. While that is seen as a very ambitious future goal (Dorigo et al., 2023), an intermediate but already impactful objective is pursued by the TomOpt software project (Strong et al., 2024a); Strong et al., 2024b) (summarized in the next section) which currently addresses, instead, particle detectors used in cosmic-ray muon tomography applications. Detectors used in this field are of relatively low complexity (Tanaka et al., 2023), such that a fully-fledged end-to-end modelling is already possible. This is owing to the lack of correlation between muons, which makes it possible to reconstruct only one muon at a time, unlike LHC particles that are highly correlated.

Upon their passage through a material volume, cosmic muons get deflected due to Coulomb scattering by nuclei of the medium. Most of these deflections occur with small angles, but occasionally large-angle scatterings may occur, as described by the Rutherford experiment (Rutherford, 1911). After multiple deflections in a macroscopic volume, the scattering angular distribution can be described by a Gaussian core, which represents the 98% of the real distribution (Workman et al., 2022), following the central limit theorem. The RMS width of this Gaussian core for a muon of momentum p [MeV] crossing a length x [cm] is related to the radiation length X_0 [cm] of the traversed material through the following approximation:

$$\theta_{RMS} = \frac{13.6MeV}{\beta cp}\sqrt{\frac{x}{X_0}} \quad \text{..... (1)}$$

where $\beta c$is the muon velocity. The value of X_0 for a material is in a direct relationship with its atomic number Z (Workman et al., 2022). Thus, inference of the radiation length X_0 for a passive volume can be used to identify its composition.

After the first proposal of this technique two decades ago (Borozdin et al., 2003), a large variety of applications of muon scattering tomography has been studied in different sectors related to industry, archaeology, nuclear waste monitoring, etc. (International Atomic Energy Agency, 2022). Searches for illicit and hazardous high-Z materials in cargo containers at border controls remain one of the main drivers for research in this field (Barnes et al., 2023).

Following the first case study of industrial furnace ladle fill-level estimation (Strong et al., 2024a), this paper introduces another TomOpt benchmark study, addressing cargo scanning and material classification, as part of our contribution to the SilentBorder project (Georgadze et al., 2024).

2. TOMOPT

The software package used for simulation and optimization in this study is referred to as TomOpt, and it is available open-source on GitHub (Strong et al., 2024b). TomOpt is a Python-based package that offers a full suite of tools and resources for optimizing a scattering tomography detector. Its modular design, provided by a set of abstract classes, allows the user to easily implement through inheritance the desired detector, inference method, muon source, etc. which fit best the specific task studied. The fully-differentiable nature of the detector modelling and inference pipeline is provided by the PyTorch automatic differentiation framework (Paszke et al., 2019).

2.1. Optimisation Cycle

A diagram summarizing the optimization process is shown in Figure 1. We start with initial detectors, which are detection panels placed above and below the volume of interest and whose initial designparameters (e.g. position, area of detection panels, etc.) are assigned suboptimal values. Muon generation is performed by sampling from flux models taken from the literature (Guan et al., 2015, Shukla and Sankrith, 2018). Muons are propagated through a passive volume using the scattering model proposed by the Particle Data Group (Workman et al., 2022) to calculate the scattering angles and spatial displacements at each step in the volume. Hits recorded in the upper and lower detectors are fitted into linear incoming and outgoing tracks using an analytical likelihood-maximisation whilst considering hit uncertainties on x and y positions. These tracks are then used for point of closest approach (PoCA) points reconstruction. The PoCA approach (Hoch et al., 2009) assumes that the entirety of the scattering of a muon occurs at a single point, which is called the PoCA vertex. Inference is then based on the calculation of various physical quantities out of the PoCA spatial and/or angular variables. An example is the inference of the voxel-wise X_0 of the passive volume, through inverting Eq. 1, by calculating the RMS of the PoCA scattering angles in each voxel of sidelength x.

The objective function, also referred to as loss function, is tailored by the user to be suitable for the task at hand. It has an inference error term that considers the performance of the detector, represented by the difference between inferred and true values (e.g. mean squared error, cross entropy for classification tasks, etc.). Optionally, the loss function can also include the detector budget, calculated as an approximate cost in currency units depending on the surface area and cost per m^2 of its panels. Two budget modes are implemented in TomOpt: a fixed-budget mode where a maximum budget is defined as a constraint for the optimization not to exceed and budget weights per panel are learnt which are used to compute scaling factors of panel (xy) spans; and a budget-penalization mode where a loose target budget is set and the cost is modelled to rapidly increase if this target is exceeded. For the latter mode, the user defines the parameters of the cost function, as well as a weight coefficient of the cost which is then added to the inference error term in the loss function. In this work, the detectors are defined with fixed dimensions throughout the optimisation, hence no cost component or budget constraint is imposed.

Detector optimization is then reduced to the minimization of the loss function through gradient-descent, arriving at updated detector parameters that guarantee optimal performance with minimal cost.

2.2. Detector design implementation for cargo scanning tasks

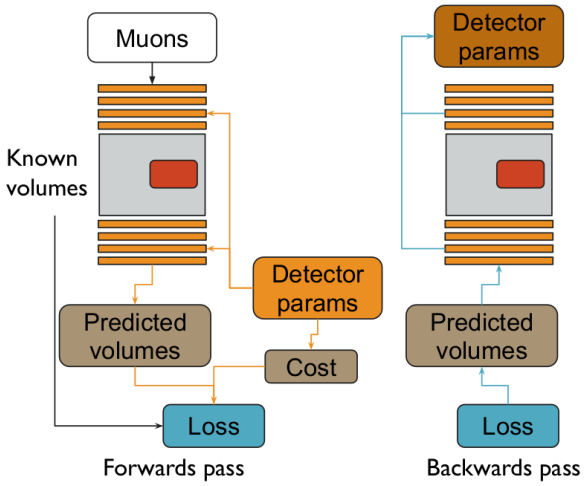


Figure 1. The TomOpt optimisation cycle. Reproduced from (Giles et al., 2024a).

Cycle, as shown in Figure 1, consists of a forwards pass which simulates the muon generation, propagation and detection resulting with a prediction of the scanned volume. This prediction enters the calculation of the loss along with other costs related to detector budget. Then, the backwards pass computes the gradients of the loss with respect to detector parameters in order to arrive at the optimal parameters through gradient descent.

The first step in the TomOpt cargo scanning benchmark study (https://github.com/vischia/TomOptCargo/) is to implement a suitable detector design. Under the current default detector configuration, a single detector is a thin detector panel, whose optimisable parameters are its area and spatial position. Several of these panels can be placed above and below the passive volume.

On the other hand, the SilentBorder detector prototype, illustrated in Figure 2a, is designed as a ''portal'', i.e. the volume of interest to be scanned is additionally surrounded by detectors on its left and right. Trucks are supposed to be driven through this portal to be entirely scanned; the presence of detector modules also on the left and the right, unusual for this type of applications, is meant to increase the statistics accumulated per unit time, while also minimizing the blind zones within the scanned volume. The building blocks of this prototype are referred to as ''hodoscopes'', which are modules consisting of three parallel detector panels, placed within a protective case. For reasons related to hardware and electronics limitations, an empty gap (≈ 10 cm) is present between the panel ends and the outer protective case for each hodoscope, which creates ''dead zones'' between adjacent hodoscopes. The presence of dead zones has a prominent impact on reconstructed images and predictions.

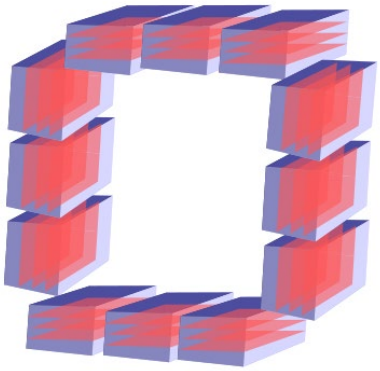


Figure 2a. A diagram of the SilentBorder muon portal prototype. Each block is a hodoscope consisting of 3 active detector panels (red).

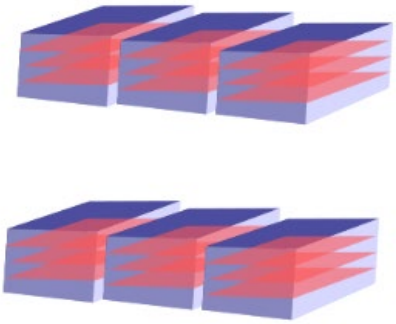


Figure 2b. An example of a hodoscope configuration that we focus on optimizing with TomOpt. Similar configurations can have several hodoscope layers (in the z direction), and different number of hodoscopes in each layer in (xy) plane.

For example, Figure 3 illustrates the PoCA point density obtained from GEANT4 simulated hits. The detectors used in this simulation are 4 hodoscope modules (1.2 m × 1.2 m × 0.4 m) of 3 thin sensitive detector panels (of active area 1 m × 1 m and a separation distance of 10 cm) each, 2 of them placed above and the other 2 below a centered iron block of size 1 m × 0.5 m × 0.5 m (Figure 4). The hodoscopes in each pair are placed adjacent to each other, but dead zones are nonetheless present due to the $10cm$ gap between the inner sensitive panel and the outer protective envelope of each hodoscope. This impacts the PoCA distribution, since the muon tracks passing through this gap, which is at the center of the object in this case, fail to be reconstructed into PoCA points. In effect, this region will acquire a lower density of PoCA points.

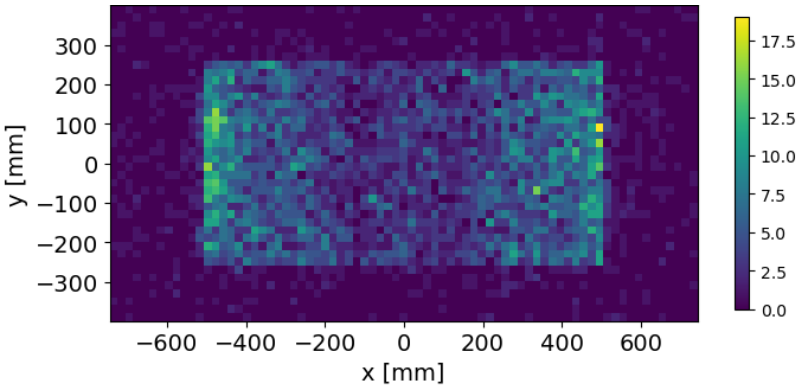


Figure 3. XY projection of the PoCA density cloud reconstructed from GEANT4 simulated data. The effect of the dead zones is visible as a localized deficit in PoCA points in the center of the block.

To study and optimise this industrial scenario, it has been necessary to implement a new detector design in TomOpt: the hodoscope, illustrated in Figure 4. This hodoscope, which is in itself made up of three detector panels, will be attributed with the optimisable parameters, whereas the inner detector panels' parameters are updated based on their parent hodoscopes' parameters. The optimization of such detector design focuses on hodoscope positioning in order to arrive at a configuration with optimal performance which minimizes the negative effect of the inter-hodoscope dead zones. Given the

current limitation of TomOpt to horizontal panels in the (xy) plane, we will only study configurations with hodoscopes placed above and below the passive volume (Figure 2b). The implementation of vertical panels is planned for future work, after which a full portal-like design will be studied. The ultimate goal after fully integrating this new detector design and performing development tests is to define the inferred physical quantity that signifies the presence of a high-Z anomaly in the container, facilitating the construction of a reliable loss function that is suitable for this task. This is currently a work in progress.

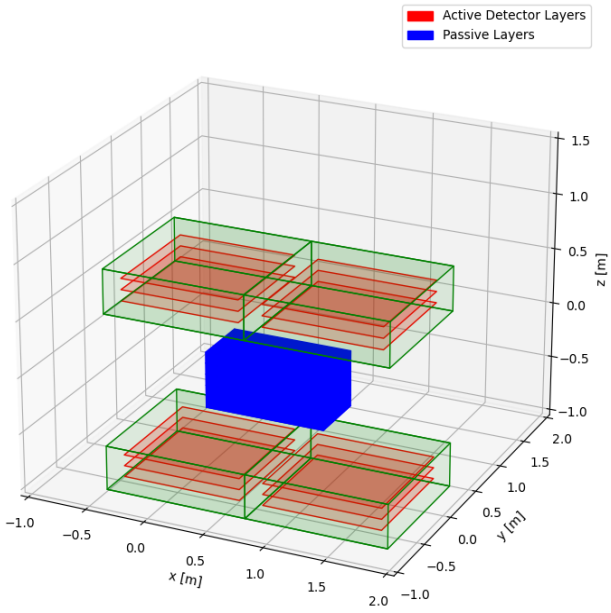


Figure 4. The hodoscope design implemented in TomOpt for this study. Each hodoscope has an outer protective case (green) encompassing 3 detector panels (red). In this diagram, a pair of such hodoscopes are placed above and below a passive volume (blue).

3. METHODOLOGY

3.1. Hodoscope module

The TomOpt detector is a thin panel, defined by 5 learnable parameters a_n, namely its position in (x, y, z) and its span in (xy). The end-to-end differentiable simulation pipeline tracks the analytical effects of these parameters on the predicted output in order to back-propagate the loss gradient through gradient-descent:

$$a_{n+1} = a_n - \gamma . \nabla L(a_n),$$

where γ is a learning rate of the parameter and L is the loss function.

In order to meet the needs of the novel cargo scan application, a new detector design of the hodoscope module has been implemented. The hodoscope is intended to be the building block of the cargo scan detector setup, hence the learnable parameters are assigned to the hodoscopes. These optimizable parameters are reduced to three, i.e. the hodoscope's position in (x, y, z), in order to maintain a fixed detector size across all hodoscope modules. Muon hits are still recorded in thin detector panels fixed inside each hodoscope, so that the fitting of panel hits into tracks is not perturbed. When a hodoscope with given initial parameters is created, inner panels are simultaneously initialized. The default number of panels is three, and their separation distance is 10 cm. The horizontal gap between panel ends and the outer hodoscope case, as well as the vertical gap between upper/lower panels and the case, are of 10 cm. During optimisation and following every update to the hodoscope parameters, the inner panel parameters are automatically adapted, keeping their relative positions with respect to their parent hodoscopes intact.

3.2. Detector modeling

The differentiability of the reconstructed hits in the detector with respect to the detector parameters is preserved through a differentiable surrogate model of the detector, which is used during optimisation. With this modelling, muon hits on the detector are assigned, in both x and y directions, a resolution (σ_{xy}) and efficiency (ε_{xy}) based on a double-sigmoid model that maps the true hit positions to the corresponding per-hit resolution/efficiency, so that the resolution and efficiency

are maximum in the center of the detector and smoothly decrease outward. This allows for the recording of hits outside the detector, but at low resolution and efficiency. The reconstructed hit position $hit_{reco,xy}$ is then sampled from a Gaussian distribution $G\left(\mu = hit_{true,xy}, \sigma = \sigma_{xy}\right)$, and the per-hit efficiency is the product of the efficiencies in x and y. Refer to (Strong et al., 2024a) for a more detailed description.

In default TomOpt, the individual detector panel parameters, i.e. (x, y, z) positions and xy span, being the learnable parameters to be optimized, are explicitly related to the reconstructed hit positions through the surrogate model. Switching to the modular hodoscope detector design, the (x, y, z) positions and a corresponding surrogate model of the hodoscope are instead passed during recording of hits in the panels inside a given hodoscope.

3.3. Detector layers

The detector setup is composed of several detector layers in the z direction, each consisting of several hodoscopes in (xy) plane. This configuration allows for a better detector coverage which is diminished due to the inter-hodoscope dead zones. Any dead zones in a layer would be conveniently covered by hodoscopes in another layer. Another benefit of this layer design is the prevention of overlap among hodoscopes following parameter updates during optimization. Each layer is given a certain range in the z direction, so this means that hodoscopes in a layer do not overlap with hodoscopes in a different layer in the z direction. Within the same layer and depending on the number of hodoscopes initialized, each hodoscope would be assigned a free region in x and y which it cannot trespass. This eliminates any overlap among hodoscopes in x and y directions in the same layer.

4. OPTIMISATION STUDY

For preliminary studies of the performance of this new detector modelling, the figure of merit is the voxel-wise X_0 predictions of the passive volume. The MSE between predictions and true values is used as the loss function that measures the performance.

We study, in this development stage, two basic scenarios of detector configuration, both detailed below. In both cases, the volume of interest is a block of iron of radiation length $X_0 = 0.01757$ m. Such high-Z material is chosen in order to obtain a better prediction of radiation length with minimal bias compared to the one obtained with low-Z materials, which substantially affects the performance of the mean squared error (MSE) loss function.

4.1. Hodoscope centering check

In this configuration, the detector setup is composed of two hodoscopes, each having a size of 1.5 m × 1.5 m × 0.4 m. One of the hodoscopes is placed above the iron block 1 m × 1 m × 0.5 m, and the other is placed below. The bottom hodoscope is centered with respect to the passive volume, while the top hodoscope is shifted by -0.5m in x and y directions. This initial configuration is shown in Figure 5. The aim here is to study the evolution of the positioning of the shifted hodoscope and of the loss over several training epochs. In each epoch, muon batches are passed through the passive volume, the loss is evaluated and hodoscope parameters are updated. The number of muons generated for each training fit is 1000, which is a moderate number for our volume dimensions. The first five epochs constitute a warm-up cycle in which suitable learning parameters are estimated without updating the hodoscope parameters.

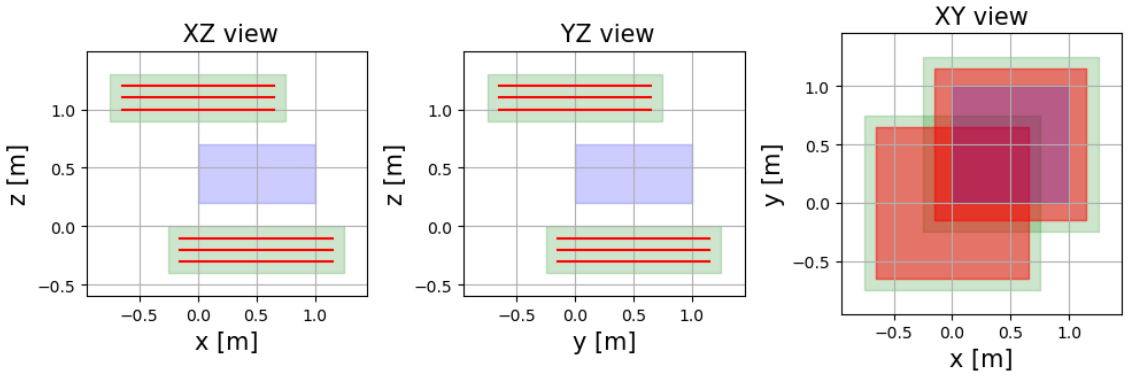


Figure 5. The initial configuration of detector setup for the optimisation test of hodoscope centering, illustrated for projections in the (x, z) (left), (y, z) (center) and (x, y) (right) planes.

Over 20 epochs, the evolution of the loss for 10 optimisation runs is presented in Figure 6. An average reduction of 45.37% of the loss was noted between the fifth epoch (end of warm-up cycle) and last epoch, but fluctuations are prominent. The final configuration of one of the runs is shown in Figure 7. We notice that the off-centered hodoscope is updated towards a slightly better centered position, and the loss tends to decrease with epochs. One of the main causes of the fluctuations observed in the loss throughout the optimization is the effect of the under-biased predictions of voxel-wise X_0, evidenced in Figure 8, which originates from our inference method based on the PoCA reconstruction method. PoCA assumes that the scattering occurs entirely at the point of closest approach between the incoming and outgoing muon track, whereas in reality, multiple small-angle scatterings occur along the muons' path through the material. This simplification, though allows for an easy implementation of the algorithm, renders the PoCA vertex an inaccurate estimation of the muon path through the material. The inferred PoCA variables are then used in a second line of inference to compute the voxel-wise X_0 predictions, which adds to the significance of the uncertainties on PoCA variables.

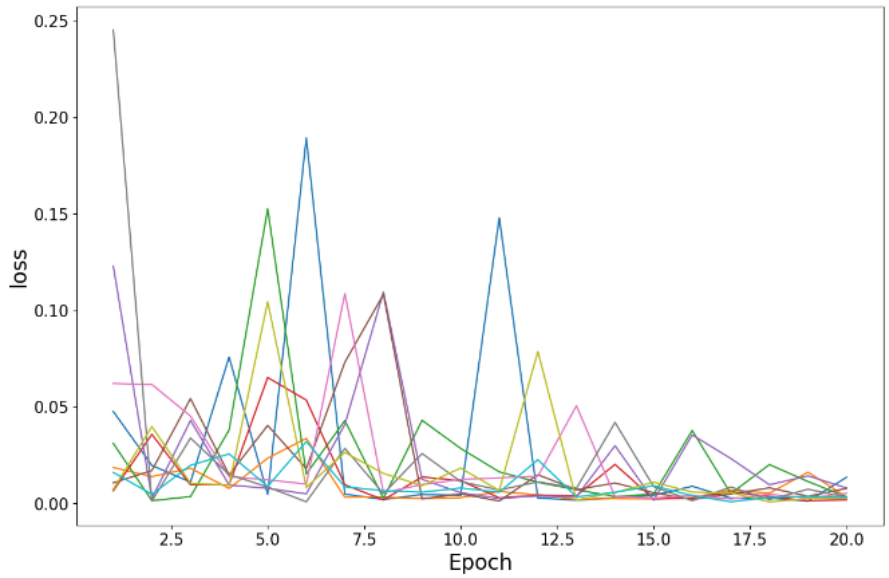


Figure 6. The evolution of the loss over 20 epochs, repeated for 10 hodoscope centering optimization runs with the same initial configuration. An overall decrease in loss is noticed, though sharp fluctuations are predominant throughout the optimisation.

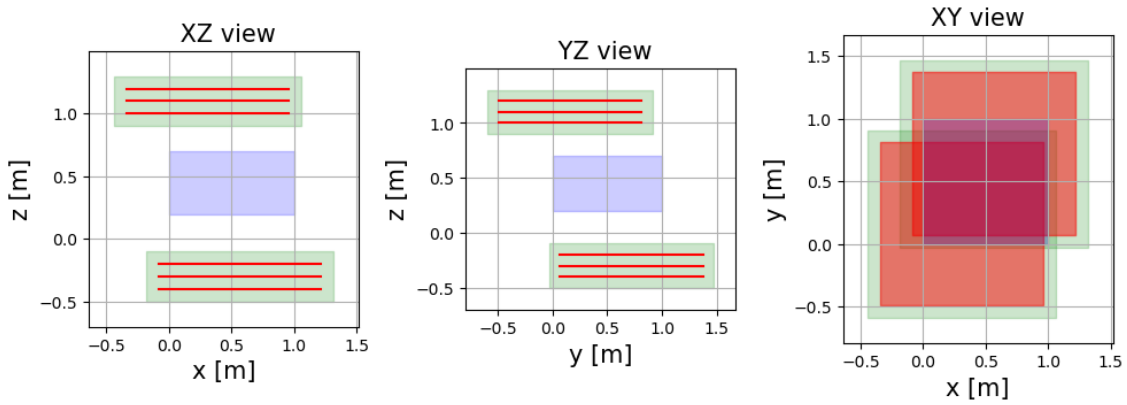


Figure 7. The final configuration of detector setup for the optimisation test of hodoscope centering, shown for projections in the (x, z) (left), (y, z) (center) and (x, y) (right) planes.

4.2. Hodoscope coverage check

In this configuration, the detector setup is composed of four hodoscopes, each having a size of 1 m × 1 m × 0.4 m. Two of the hodoscopes are placed above the iron block (1 m × 0.5 m × 0.5 m), and the other two are placed below. An empty gap of 0.5 m is imposed between each hodoscope pair above and below the volume. This initial configuration is shown in Figure 9. The aim here is to study the coverage of the hodoscopes in terms of reducing any detrimental effect the dead zone between them may have on the X_0 inference. The number of muons generated for each training fit is 1000, same as for the previous case.

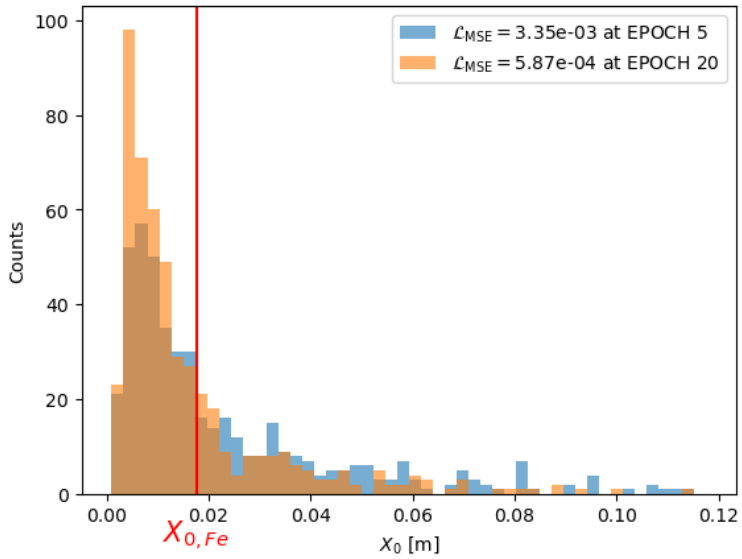


Figure 8. Distributions of the voxel X0 predictions at the end of the warm-up cycle (epoch 5) and at the end of optimisation (epoch 20). Both distributions peak at values lower than the true radiation length of iron (red), although a decrease in L_MSE is observed.

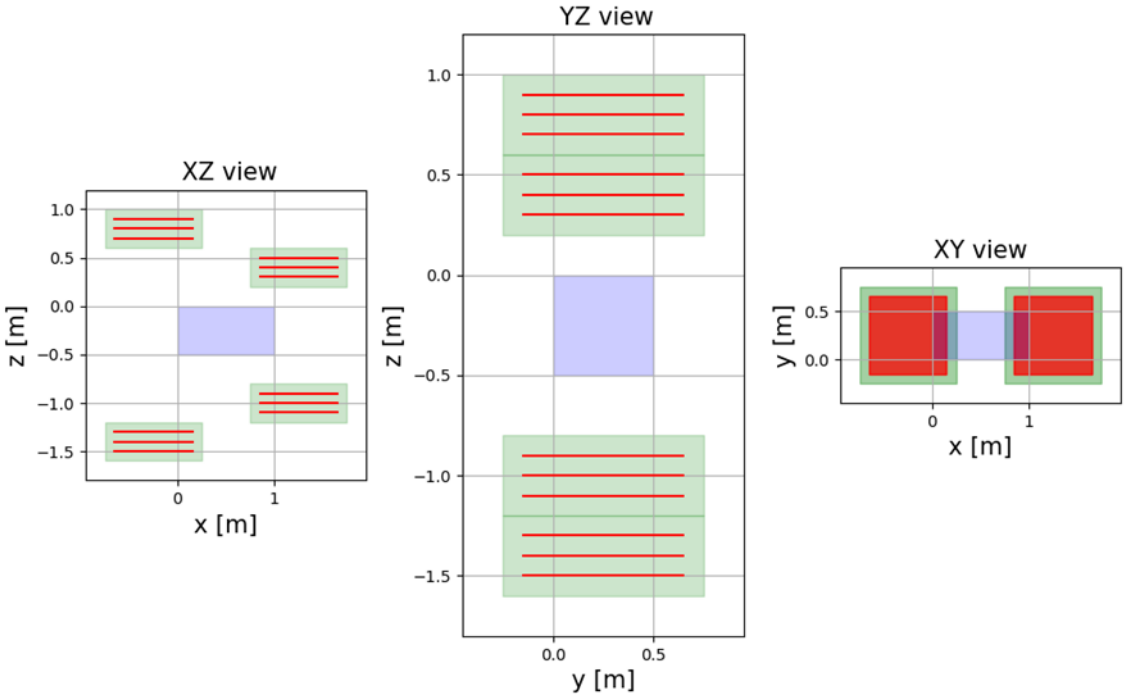


Figure 9. The initial configuration of detector setup for the optimisation test of hodoscope coverage, presented in (x, z) (left), (y, z) (center) and (x, y) (right) views.

Over 20 epochs, the evolution of the loss for 10 optimisation runs is presented in Figure 10. An average loss reduction of 71.89% is noted between the fifth epoch (end of warm-up cycle) and last epoch. However, as in the previous case, fluctuations are persistent. The final configuration of one of the runs is shown in Figure 11. We notice that both hodoscope pairs have shifted towards each other in order to cover the dead zone separating them, thus visibly reducing the loss. The fluctuation observed in the loss at the end of the optimization are, again, mainly due to the under-biased predictions of voxel-wise X_0. Depending on the passive object material and dimensions, detector dimensions, number of hodoscopes and hodoscope layers, etc., it is possible that in some cases the presence of gaps in certain configurations results in a better performance. In such cases, this can be justified by an increase in detector acceptance which accounts for muons incident with large angles. In our scenario, adjacent hodoscopes still have a good acceptance, and knowing that most of the incident muons are vertical, we would expect that the gaps are covered during optimization.

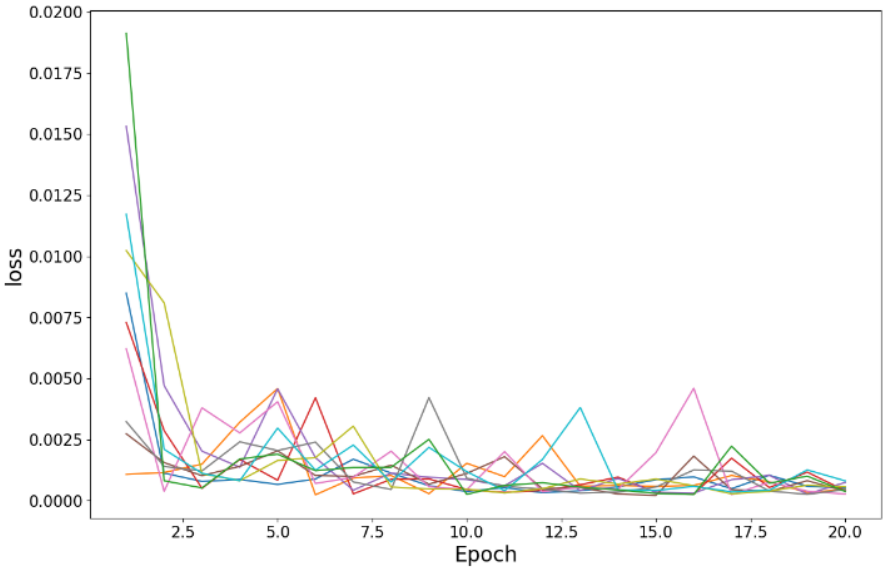


Figure 10. The evolution of the loss over 20 epochs, repeated for 10 hodoscope coverage optimization runs with the same initial configuration.

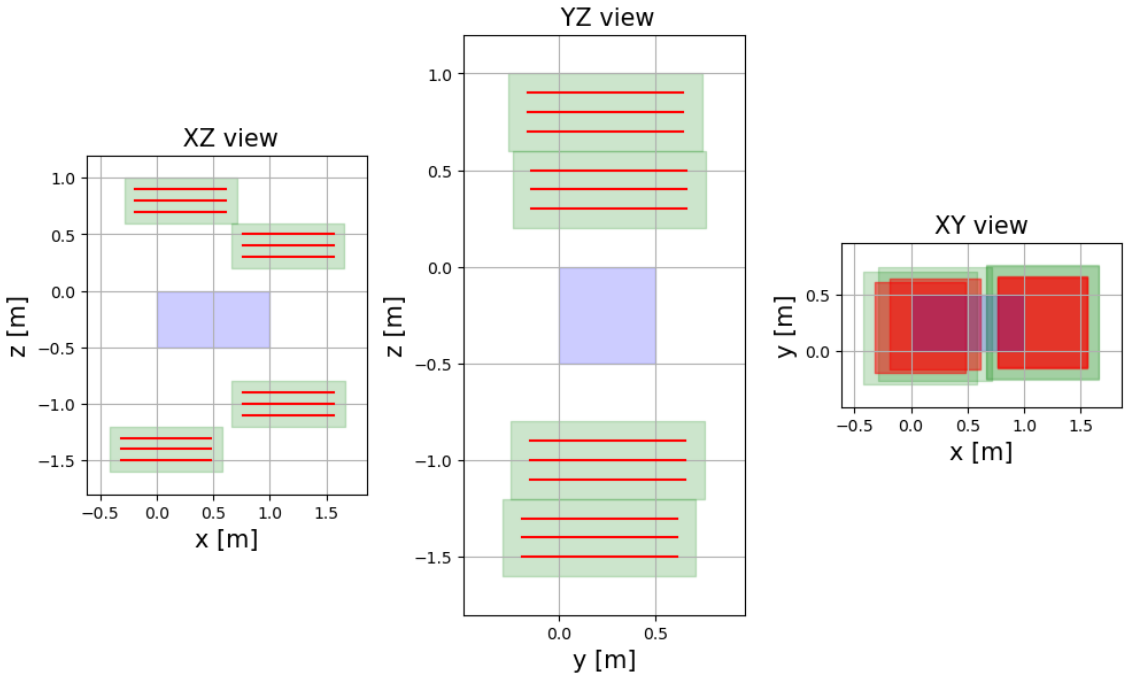


Figure 11. The final configuration of detector setup for the optimisation test of hodoscope coverage, shown for projections in the (x, z) (left), (y, z) (center) and (x, y) (right) planes.

5. ONGOING/FUTURE DEVELOPMENT

After the implementation of a different detector design dedicated for cargo scan applications, the differential backward pass in the TomOpt optimisation cycle has shown its capability to improve hodoscope positions in simple scenarios. A limitation concerns the unreliability of the PoCA algorithm as a reconstruction method of scattering locations, which negatively affects the material inference. Any subsequent anomaly detection algorithm is based on these predictions, and is therefore affected by their accuracy levels. A possible solution is to replace the PoCA method with an algorithm that more accurately estimates the muon path inside the material. One of the algorithms (Barnes et al., 2023) that we deem promising to investigate is the Expectation Maximisation algorithm (Schultz et al., 2007).

6. CONCLUSION

After the successful demonstration of the capabilities of TomOpt in a previous industrial case study (Strong et al., 2024a), providing an end-to-end differentiable and inference-aware optimisation of particle physics detectors, we introduce in this manuscript a second benchmark study devoted to cargo scanning applications such as sea or truck containers. This benchmark requires a novel modelling of the detector design, which has been recently implemented in TomOpt. The cargo scanning case study is an ongoing work in the current development stage, where preliminary results indicate the capability of the software to optimize simple configurations of the implemented modular design of the detector. The results of this work will be used for improvements of possible future upgrades of the SilentBorder detector.

ACKNOWLEDGEMENTS

This is extended version of the paper presented at MARESEC 2024.

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 Grants No. T.0099.19 and J.0070.21.

Pietro Vischia’s work was supported by the "Ramón y Cajal" program under the Project No. RYC2021-033305-I funded by 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 Fisica Corpuscular, IFIC (CSIC-UV).

CONFLICT OF INTEREST

Authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

REFERENCES

Barnes, S., Bonechi, L., D'Alessandro, R., Giammanco, A., and Kudryavtsev, V. (2023) 'Cosmic-ray tomography for border security', Instruments, 7(1), p. 13. Available at: https://doi.org/10.3390/instruments7010013.

Bonechi, L., D'Alessandro, R., and Giammanco, A. (2020) 'Atmospheric muons as an imaging tool', Reviews in Physics, 5, p. 100038. Available at: https://doi.org/10.1016/j.revip.2020.100038.

Borozdin, K. N., Greene, S. R., Milner, E. C., Morris, C. L., Saunders, A., Schultz, L. J., and Teasdale, M. E. (2003) 'Radiographic imaging with cosmic-ray muons', Nature, 422, p. 277. Available at: https://doi.org/10.1038/422277a.

Dorigo, T., Giammanco, A., Vischia, P., Calafiura, P., and Cranmer, K. (2023) 'Toward the end-to-end optimization of particle physics instruments with differentiable programming', Reviews in Physics, 10, p. 100085. Available at: https://doi.org/10.1016/j.revip.2023.100085.

Georgadze, A., Giammanco, A., Kudryavtsev, V., Lagrange, M., and Turkoglu, C. (2024) 'A simulation of a cosmic ray tomography scanner for trucks and shipping containers', Journal of Advanced Instrumentation in Science, 1(482). Available at: https://doi.org/10.31526/jais.2024.482.

Giles, G. C., Smillie, L., Wang, J., Steffens, M., and Hanlet, P. (2024) 'TomOpt: Differential optimisation for task- and constraint-aware design of particle detectors in the context of muon tomography', Machine Learning: Science and Technology, 5(3). Available at: https://doi.org/10.48550/arXiv.2309.14027.

Giles, G. C., Steffens, M., Wang, J., Hanlet, P., and Smillie, L. (2024) 'TomOpt: Differential muon tomography optimisation. Version v0.1.0', GitHub repository. Available at: https://doi.org/10.5281/zenodo.10673885.

Guan, M., Chu, M. C., Cao, J., Luk, K. B., and Yang, C. (2015) 'A parametrization of the cosmic-ray muon flux at sea-level'. Available at: https://arxiv.org/abs/1509.06176.

Hoch, R., Mitra, D., Gnanvo, K., and Hohlmann, M. (2009) 'Muon tomography algorithms for nuclear threat detection', Studies in Computational Intelligence, 214. Available at: https://doi.org/10.1007/978-3-540-92814-0_35.

International Atomic Energy Agency (2022) Muon imaging: Present status and emerging applications. IAEA-TECDOC-2012. Available at: https://www.iaea.org/publications/15182/muon-imaging.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., et al. (2019) 'PyTorch: An imperative style, high-performance deep learning library', Advances in Neural Information Processing Systems. Available at: http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf.

Rutherford, E. (1911) 'The scattering of α and β particles by matter and the structure of the atom', The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 21(125), pp. 669–688. Available at: https://doi.org/10.1080/14786440508637080.

Schultz, L. J., Blanpied, G. S., Borozdin, K. N., Fraser, A. M., Hengartner, N. W., Hogan, G. E., et al. (2007) 'Statistical reconstruction for cosmic ray muon tomography', IEEE Transactions on Image Processing, 16(8), pp. 1985–1993. Available at: https://doi.org/10.1109/tip.2007.901239.

Shukla, P., and Sankrith, S. (2018) 'Energy and angular distributions of atmospheric muons at the earth'. Available at: https://arxiv.org/abs/1606.06907.

Tanaka, H. K., Eino, M., Kuno, M., and Nishiyama, R. (2023) 'Muography', Nature Reviews Methods Primers, 3(1), p. 88. Available at: https://doi.org/10.1038/s43586-023-00270-7.

Workman, R. L., Amsler, C., Bauer, T., Beringer, J., and Dahmes, B. (2022) 'Review of particle physics', Progress of Theoretical and Experimental Physics, 2022(8), p. 083C01. Available at: https://doi.org/10.1093/ptep/ptac097.