Journal of
Applied Physics

RESEARCH ARTICLE | OCTOBER 10 2025

Pillar embedding visualization for muon-scattering tomography

Special Collection: Muography: Discoveries, Innovations, and Applications

A. Bueno Rodriguez ; F. Sattler ; B. Carrillo Perez ; M. Perez Prada ; J. M. Alameddine ;
M. Stephan ; S. Barnes

J. Appl. Phys. 138, 144903 (2025)
https://doi.org/10.1063/5.0288257

AIP Publishing

Journal of
Applied Physics

ARTICLE

pubs.aip.org/aip/jap

Pillar embedding visualization for muon-scattering tomography

Cite as: J. Appl. Phys. **138**, 144903 (2025); doi: 10.1063/5.0288257
Submitted: 30 June 2025 · Accepted: 22 September 2025 ·
Published Online: 10 October 2025

View Online Export Citation CrossMark

A. Bueno Rodriguez,[a)] F. Sattler, B. Carrillo Perez, M. Perez Prada, J. M. Alameddine, M. Stephan, and S. Barnes

AFFILIATIONS

German Aerospace Center (DLR), Institute for the Protection of Maritime Infrastructures, Fischkai 1, 27572 Bremerhaven, Germany

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

[a)] **Author to whom correspondence should be addressed:** angel.bueno@dlr.de

ABSTRACT

Muon-scattering tomography (MST) utilizes the deflection of cosmic-ray muons to non-invasively reconstruct the three-dimensional internal structure and material composition of concealed objects, such as those in maritime cargo. Yet, the high dimensionality of reconstructed MST volumes and sparsity of muon hits hinder reliable material discrimination and structural interpretation. We present an unsupervised workflow that visualizes learned data embeddings for material identification. The pipeline couples the Blender-to-Geant4 simulation framework, enabling the rapid prototyping of complex 3D scenes with a standard and widely adopted MST reconstruction algorithm, the Point of Closest Approach (PoCA), to reconstruct the scenes. A structured muon-data sampling grid, termed pillars, feeds an exploratory embedding technique that reveals discriminative material patterns in the reconstructed outputs. Experimental results demonstrate that the proposed approach mitigates key machine-learning challenges in MST; at the same time, they reveal the intrinsic limitations of PoCA estimates for mainstream material classification with machine-learning approaches, and we introduce corrections that enhance visualization and enable data-driven analysis in practical MST deployments.

I. INTRODUCTION

Muon-scattering tomography (MST) images the internal structure and material composition of a concealed Volume of Interest (VOI) by tracking naturally occurring cosmic-ray muons.[1,2] As illustrated in Fig. 1, each muon is measured twice: once in the upper detector plane before it enters the VOI and once in the lower detector plane after it exits. These two position and direction measurements, or muon hits, define the incoming and outgoing trajectory vectors, from which the single-scatter deflection angle θ can be estimated. Collecting many such muon hits over the exposure time yields a hit list that can be transformed into a volumetric image using a tomographic reconstruction method, such as the Point of Closest Approach (PoCA)[3] or Angle Statistics Reconstruction (ASR)[4] algorithms. The common working principle of these algorithms is to back-project the two trajectory vectors into the VOI and assign scattering scores based on the deflection angle θ. In PoCA, the score is assigned to the single voxel where the projected paths come closest.[3] In ASR, the θ-based score is distributed across every voxel intersected by the straight-line track pair that meets the "traversed by both tracks" criterion, without assuming a single scattering point.[4] Repeating this procedure for every muon populates a 3D grid with per-voxel *scattering scores*. Finally, aggregating the scores across all muons produces a three-dimensional scattering density map whose voxel intensities correlate with the local material density and atomic number. These 3D scattering density maps can be used in a wide range of non-destructive MST applications, from contraband detection in cargo and border security,[5] through civil-engineering and infrastructure monitoring,[6] to industrial non-destructive testing[7] and nuclear-fuel cask surveillance.[8] In practice, most MST applications deduce material properties solely from the geometric intersection of reconstructed tracks, leaving the recovered information inherently incomplete.[9] Precise muon-momentum measurements would tighten the link between scattering angle and density. Yet, existing

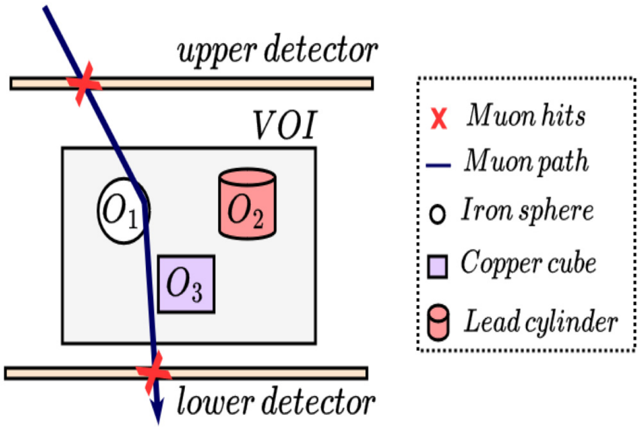

FIG. 1. Visualization of a canonical muon-scattering setup: Muons passing through an object scatter due to Coulomb interactions with atomic nuclei. Muon tomography algorithms use muon hits to estimate muon pathways and image a VOI as a 3D scattering density map.

hardware detector solutions, whether as meter-scale weak magnets or passive scatterers, are too bulky for field measurement setups and introduce large, irreducible uncertainties from single-muon statistics.[1,2]

MST measurement campaigns typically yield small, noisy, and weakly labeled data sets, posing significant challenges for data-driven analyses. This scarcity and noise stem primarily from the limitations of precise momentum measurements, which, although this could in principle be addressed by extending exposure times in MST campaigns to collect more muons, are constrained by the inherent low natural rate of cosmic-ray muon flux, requiring extended exposure times (often hours or days) to accumulate sufficient statistics for reliable imaging, making real-world data collection time-intensive and resource-heavy. Additionally, the relatively short operational history of MST, spanning only a few decades since its inception, has resulted in a limited archive of publicly available data sets. The weak labeling arises because concealed objects in real-world scenarios often lack precise ground truth annotations, relying on expert interpretation or additional invasive verification methods that are rarely available. To overcome these constraints, high-fidelity synthetic data can be simulated to closely mirror real conditions, carry perfect ground truth, and explore corner-case geometries that would be infeasible to measure in the field. Typically, MST setups are simulated using Geant4, a sophisticated particle-physics simulation software.[10] Yet, simulating these scenes at scale requires hand-coding and recompiling the simulation for each design iteration, a bottleneck that sharply limits throughput. The recently released Blender-to-Geant4 (B2G4) workflow[11] bridges this gap: users design scenes interactively in Blender, an open-source 3D computer software program,[12] and export them directly into Geant4, thereby eliminating the need for hand-coded geometry and recompilation cycles.

Although every material leaves characteristic scattering *signatures*, classifying them, whether on real or simulated data, remains a challenge. The voxel statistics are inherently noisy and sparsely sampled; moreover, each value represents an integral along the unknown muon path, intertwining object geometry, material composition, and track length in a nonlinear manner. These factors blur class boundaries and feed downstream machine-learning models ambiguous, occasionally contradictory data samples, limiting their reliability. In this work, we present a structured methodology for data sampling and feature extraction explicitly designed for material identification. We first use B2G4 to create two large, fully synthetic MST data sets that span a wide range of shapes and material densities. From the recorded hit pairs, we recover the incident and exiting muon directions and reconstruct a 3D scattering density map of the VOI. We adopt PoCA as the reconstruction algorithm given that it stands as the mainstream method in MST: its closed-form simplicity, linear scaling with track count, and extensive experimental validation make it an ideal baseline.

Next, each 3D scattering density map is decomposed into a grid of vertical pillars, or 3D columns, aligned along the $+Z$ axis in the coordinate space, approximating the predominant muon flight direction. For every pillar, we compute a 12-element feature vector that blends simple statistics with spectral measures to capture local variations in scattering power and, by extension, material density. To suppress PoCA artifacts, we test several pre-filters, including median kernels and lightweight 3D convolutions, that denoise the volume before pillar extraction. The resulting feature cloud is embedded with t-Stochastic Neighbor Embedding (t-SNE),[13] yielding a two-dimensional map in which pillars from different materials cluster according to their scattering signature. The proposed end-to-end workflow, from PoCA reconstruction to t-SNE projection, serves as a diagnostic lens on MST data, revealing where PoCA's single-scattering assumption breaks down, how energy-dependent scattering biases affect cluster structure, and where materials of similar density are prone to confusing machine-learning models. Hence, the workflow offers a robust data pipeline for coupling large, annotated MST data sets with a tool to sample, visualize and anticipate the challenges that machine-learning algorithms may encounter in MST deployments, where measurements are typically conditioned by a broad spectrum of physical phenomena affecting the reconstruction and materials of density are prone to be confused.[14]

II. RELATED WORK

Machine learning has been explored in MST to analyze muon data from Geant4 simulations,[7,15] with feature discriminators to detect waste drum materials,[16,17] classification of measured data,[18] nuclear threat detection[19] or as the core component in muon imaging,[20–23] among others. Despite their state-of-the-art performance, these approaches employ small synthetic data sets that lack detailed 3D geometries. The scalable simulation of detailed synthetic 3D data for machine-learning applications in MST is paramount for developing models that bridge the gap between real and simulated scenarios. However, the simulation of MST data in Geant4 remains a largely hand-crafted process,[24] requiring manual coding of solids, placements, and materials in C++, with geometry errors only detected at compilation time, and no interactive modeler provided. Geometry Description Markup Language

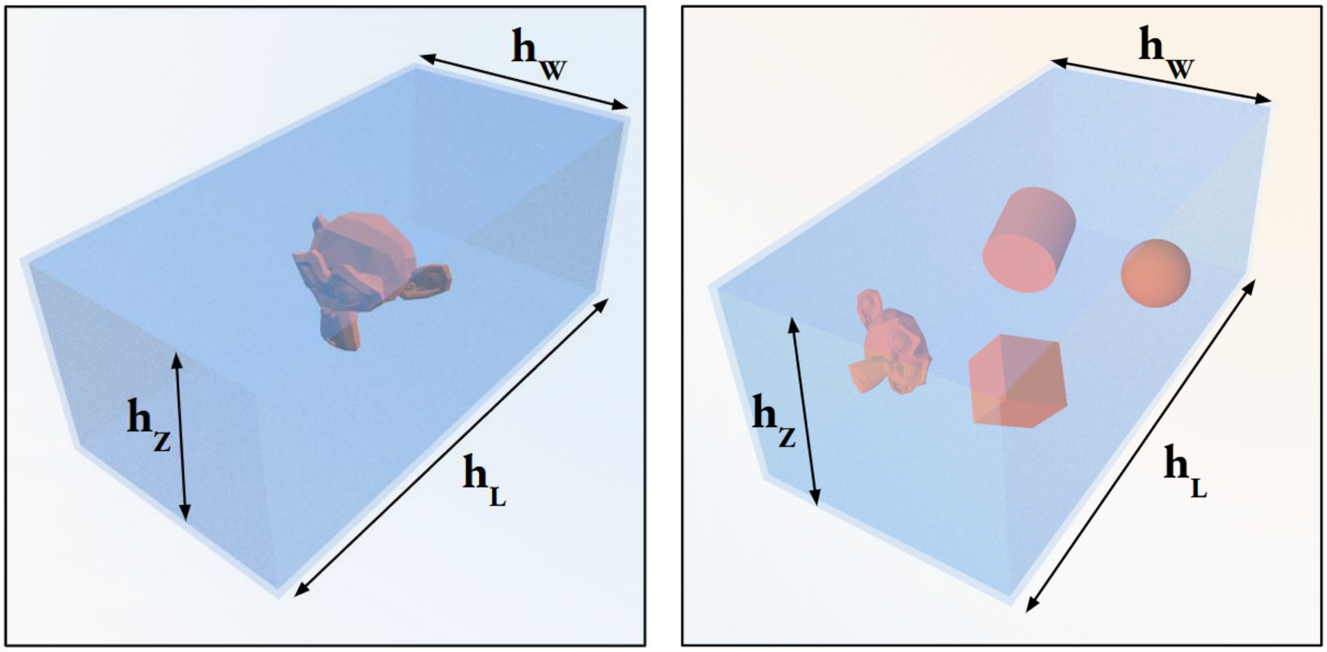

(a)

(b)

FIG. 2. Simulation setup rendered using the B2G4 framework, with the VOI in blue background. (a) MonkeyContainer: a VOI with a single monkey head object inside, made of a single material. (b) 4ShapesContainer: a VOI containing four distinct geometric shapes (cylinder, sphere, monkey head, and a cube). All the shapes in the 4ShapesContainer are made of the same material. The dimensions of the VOI are $h_w:3$ (width), $h_l:6$ (length), and $h_z:2.4$ (height), corresponding to those of a maritime cargo container.

(GDML)[25–27] offers a surrogate through XML-based descriptions, yet it still involves verbose text editing without live visual feedback. Blender-to-Geant4 (B2G4) addresses these limitations by coupling the open-source 3D modeler Blender with Geant4 geometrical hierarchies through a custom data format for 3D geometry exchange[11] and interactive scene design for runtime simulations. Yet, a workflow that couples scalable 3D scene creation (e.g., via B2G4) with machine-learning pipelines, allowing algorithms to be challenged across the full complexity spectrum of modern MST deployments, remains to be explored.

III. FRAMEWORK DESCRIPTION

This section introduces the use of pillars[28] as a sampling technique in muon tomography data to examine material density distribution within a scanned volume. These pillar structures traverse the reconstructed VOI while extracting features, thus allowing the 3D scattering density maps to be divided into manageable segments, namely, feature descriptors, for further analysis. To explore the separability of materials, we embed these feature descriptors using t-SNE. The 2D embedding produced by t-SNE reveals clusters driven purely by muon-scattering behavior, providing an intuitive map of where different densities lie and a diagnostic view of the challenges machine-learning models face in real measurements.

A. Data set generation

The data set generation follows the B2G4 end-to-end data pipeline, ranging from the specification of the scene in B2G4, to the simulation phase, and the subsequent muon tomography reconstruction. Each of these data set generation steps are as follows:

Scene creation: Fig. 2 depicts the two primary data sets generated with the B2G4 framework, the MonkeyContainer, and the 4ShapesContainer. Each data set contains ten scenes, and each scene contains a container with objects made of the specified materials. These materials are categorized into three groups: high-density materials (Au, Pb, W), medium-density materials (Al, Fe, Cu), and low-density materials (concrete, air, Si, water). The dimensions of this container remain the same for all the scenes, in both data sets. In meters, the dimensions are $h_w:3$ (width), $h_l:6$ (length), and $h_z:2.4$ (height), corresponding to those of standard application of MST analysis in cargo containers. Given that our interest lies in determining whether material densities can be identified from the reconstruction of 3D scattering density maps, the composition of the container is initially set as a vacuum. In terms of 3D geometries, the data sets are defined as follows:

1. MonkeyContainer [Fig. 2(a)]: Designed to test if materials can be distinguished with a consistent shape. This data set features

an empty container with a single monkey head (dimensions: 1.92 m width, 1.2 m height, 1.38 m depth) centered inside. In each of the 10 scenes, the monkey head is made of a different material (high, medium, or low density).

2. 4ShapesContainer [Fig. 2(b)]: Aimed at evaluating whether object shape influences material signatures. It includes four diverse shapes in a container, symmetrically placed in quadrants: a sphere (0.41 m radius), cube (0.64 m sides), cylinder (0.93 m height, 0.46 m radius), and scaled monkey head (1.74 m depth, 0.794 m width, 0.918 m height). All shapes in a given scene are made of the same material, with 10 scenes total (varying materials). This setup mimics real-world variability in object size and form to challenge the MST pipeline.

Following the data set specification setup, and once all of the scenes are exported using B2G4,[11] we proceed to the simulation of cosmic-ray muons.

Simulation phase: The simulations were performed using Geant4 version 11.1.2 and the Cosmic-Ray Shower Library (CRY) version 1.7,[29] with standard CRY settings: muon energies ranging from 1 MeV to 100 TeV and zenith angles from 0° to 90°. All the scenes in both data sets were simulated with 5×10^6 muons generated over a 1 m^2 plane, corresponding to an equivalent exposure time of 500 min under typical cosmic-ray muon flux conditions at sea level, that is, $\approx$10 000 muons/m^2/min. Only muons are retained: for every particle that intersects a detector plane, the precise hit positions, directions, and particle energy are recorded. The detector planes themselves are modeled as *ideal* sensors: perfect efficiency, full geometric acceptance, zero edge effects, and infinitesimal spatial resolution to avoid instrumental smearing that could obscure muon-scattering physics and material signatures for subsequent analysis.

Tomographic reconstruction: After generating and processing the muon-hit lists, we reconstruct every scene using the PoCA algorithm with a uniform voxel size of 5 cm. Finer voxels offer minimal added insight, as smaller voxel sizes reduce muon hits, trading resolution for increased noise while quadratically increasing the computational load. The PoCA algorithm determines the minimal separation between the incident and exiting muon trajectories to identify a single scattering point.[3] To do so, the PoCA algorithm uses the positional and directional data of the muon's path before and after deflection to solve the parametric equations of straight lines.[3] Once the scattering point is estimated, the scattering score is assigned to the corresponding voxel. Then, to estimate the scattering angle θ for each voxel, the mean angle of all scatterings that occurred in this voxel is chosen and angles below 0.01° (0.174 mrad) are filtered out to exclude insignificant deflections while ensuring focus on material-relevant events. Once the tomographic reconstructions for each data set are complete, a volumetric array for each scene is produced, containing the estimated PoCA points, that is, the 3D scattering density maps, for the pillar sampling step.

B. Pillar sampling in MST

To facilitate material identification, we employ a strategy to extract sparse feature vectors from the 3D scattering density maps through a structured sampling process we term *pillarization*. This method, inspired by 3D pillar-based discretization techniques developed originally for aggregating multiscale features in real-time object detection from point clouds in autonomous driving,[28] computer vision applications,[30] and volumetric characterization,[31] divides the reconstructed volume into vertical pillars, that is, columns of voxels sharing the same (x, y) muon-scattering footprint. By aggregating statistics along the z axis, pillarization collapses these vertical voxel stacks into compact feature vectors, respecting the predominantly downward paths of cosmic-ray muons while capturing essential characteristics of the 3D imaged volume. Figure 3 provides a schematic illustration of the proposed pillarization strategy. We define a sampling grid spanning X, Y, and Z dimensions to fully enclose the container boundaries, aligned with the VOI. The grid resolution corresponds to one pillar per position cell in the ideal detector, computed for each $(+X, +Y)$ location in the positive directions of width and height, extending along the positive z-depth $(+Z)$. Note that while pillars are aligned vertically along the $+Z$ axis to approximate the predominant downward muon flux, simulations include muons from all incident angles (0°–90° zenith) using standard CRY settings; oblique muons contribute to the pillars they traverse, ensuring the approach remains consistent with realistic MST measurements. After the pillarization step, we obtain P_{jk} pillars. Each pillar can then be treated as a discrete 1D spatial signal $s[i]$ along the z axis, thus enabling the extraction of 12 feature descriptors. These features provide insights into the pillar density distribution and materials within the 3D reconstructed scattering density

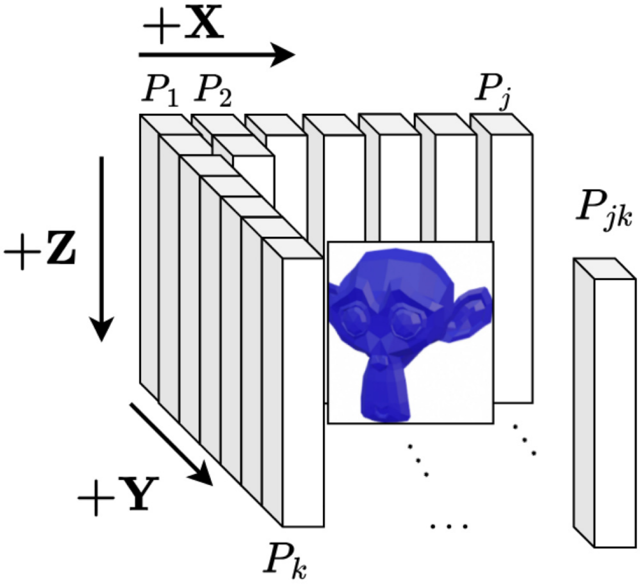

FIG. 3. Schematic representation of the proposed pillar sampling (pillarization): The sampling process yields multiple pillars ($P_{11}, P_{12}, \ldots, P_{21}, P_{22}, \ldots, P_{jk}$) aligned vertically in the X, Y, and Z axes. Here, j and k represent the total number of grid divisions along the X and Y axes, respectively. We conduct grid sampling across the VOI in the positive X and Y directions, and from each pillar P_{nm}, we extract a 12-dimensional feature vector along the positive $+Z$ direction for analysis.

TABLE I. Features extracted from each pillar, treated as a discrete 1D spatial signal $s[i]$ along the z axis.

Statistical features

Feature	Definition	Interpretation (per pillar)
Mean	$\mu_s = \frac{1}{n}\sum_i \lvert s[i]\rvert$	Average signal value
Median	Midpoint of sorted $s[i]$	Signal midpoint; robustness against noise spikes
Standard deviation	$\sigma_s = \sqrt{\frac{1}{n-1}\sum_i (s[i]-\mu_s)^2}$	Signal variability; measures dispersion
Interquartile range	Q3 – Q1	Central spread; ignores outliers
Skewness	$\frac{1}{n}\sum_i \left(\frac{s[i]-\mu_s}{\sigma_s}\right)^3$	Pillar distribution asymmetry
Kurtosis	$\frac{1}{n}\sum_i \left(\frac{s[i]-\mu_s}{\sigma_s}\right)^4$	Pillar distribution peakedness
Maximum	$\max(s[i])$	Peak pillar value; identifies regions of high reconstructed scattering density
Minimum (non-zero)	$\min\{s[i] \mid s[i] \neq 0\}$	Lowest non-zero pillar value; captures baseline reconstructed scattering density

Complexity and spectral features

Feature	Definition	Interpretation (per pillar)
Energy	$\sum_i s[i]^2$	Total signal power; quantifies intensity
Normalized entropy	$-\frac{1}{\log_2(n)}\sum_s p(s)\log_2(p(s))$	Signal disorder; higher for varied profiles
Total intensity	$\sum_i s[i]$	Cumulative signal strength
Maximum Fourier amplitude	$\max_k \lvert\mathscr{F}(s)[k]\rvert$	Strongest Fourier component; reveals periodic patterns

maps. Let n be the pillar length, μ_s the mean, σ_s the standard deviation, and $p(s)$ the probability of amplitude level s, we group 12 feature descriptors as a combination of statistical and complexity/spectral features, thus creating a feature vector. Table I lists all features together with their mathematical definitions and a brief per-pillar interpretation, which we define as follows:

1. Statistical features: The selection of features here provides insight into the muon-scattering distribution along the pillar profile: mean, median, standard deviation, interquartile range, skewness, kurtosis, the maximum and minimum (non-zero) pillar values sampled from the reconstructed scattering densities.
2. Complexity and spectral features: These features offer insights into the signal energy distribution of the 3D reconstructed scattering density maps: energy, normalized entropy, total intensity, and maximum Fourier amplitude.[31]

Note that all features are computed independently per pillar from the discrete 1D spatial signal $s[i]$ (length n). The pillarization strategy constructs vertical pillars aligned with the prevailing muon flight direction. Once all reconstructed tracks are accumulated into the 3D scattering-density map, the typical transverse displacement of a muon is less than 5 cm.[3,5] Consequently, each P_{jk} column collects the scattering contribution of essentially all muons traversing it, so a pillar encodes the aggregated statistics of many paths rather than a single track. This path integration yields a robust, path-integrated estimate of local scattering power and makes pillars a natural unit for reasoning about material density. Pillars corresponding to empty regions exhibit (near-)zero deflection and are excluded from the analysis to focus on meaningful data.

C. Embedding visualization

In machine learning, an embedding is a compact, low-dimensional representation of high-dimensional data that preserves key relationships. We employ embeddings as an essential tool for interpreting reconstructed 3D scattering density maps in MST. Specifically, after pillar sampling, we embed the 12-dimensional feature vectors using t-distributed Stochastic Neighbor Embedding (t-SNE).[13] The t-SNE algorithm[13] converts Euclidean distances between pillar pairs into conditional probabilities representing pairwise affinities in the high-dimensional space, with the user-defined perplexity parameter (per) controlling the effective neighborhood size. A randomly initialized two-dimensional map is then iteratively optimized to ensure its affinities, q_{ij} (computed using a heavy-tailed Student's t distribution), closely match the high-dimensional affinities p_{ij}. Formally, the algorithm minimizes the Kullback–Leibler divergence as

$$KL[P\|Q] = \sum_i \sum_j p_{ij} \log \frac{p_{ij}}{q_{ij}}, \tag{1}$$

with P and Q representing the probability distributions in both the high- and low-dimensional spaces, and p_{ij} and q_{ij} denoting the pairwise similarities between data points. This asymmetric divergence prioritizes the preservation of local structure and topology, making t-SNE particularly effective for revealing cluster structure

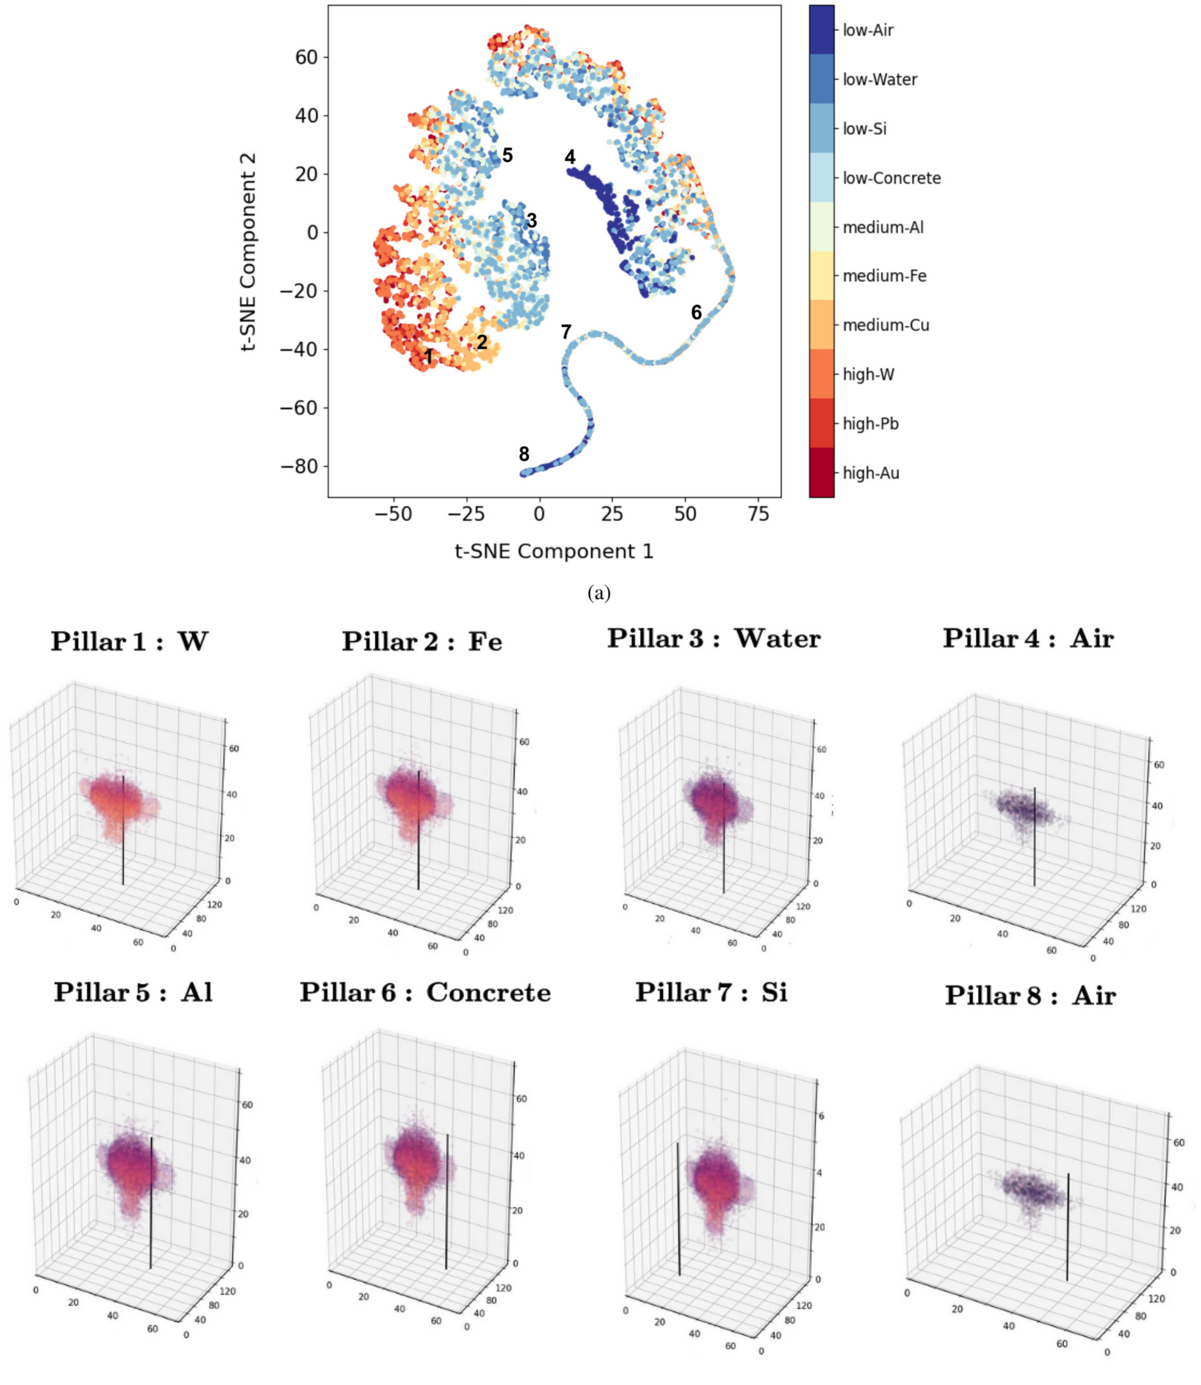

(a)

(b)

FIG. 4. t-SNE results with selected pillar samples for interpretation. (a) t-SNE visualization of the MonkeyContainer data set, where each point represents a pillar's 12D feature vector; proximity indicates similar pillar features; clusters reflect discriminative material patterns, while tails indicate low-density or noisy regions. The numbers correspond to selected pillar samples for visualization, and the colorbar represents material density, from blue (low density) to red (high density). (b) Numbered t-SNE samples for tungsten (W), iron (Fe), aluminum (Al), concrete, silicon (Si), water, and air, with black lines indicating the position of sampled pillars in the 3D scattering density map. Axes are labeled as X, Y, Z (dimensions: $h_w:3\,\mathrm{m}$, $h_l:6\,\mathrm{m}$, $h_z:2.4\,\mathrm{m}$).

in complex data sets. By projecting these embeddings into a low-dimensional space, t-SNE captures hidden similarities among pillars, facilitating visual inspection and material identification in the reconstructed volume.

IV. RESULTS AND DISCUSSION

A. Single shape, multiple materials

Here, we analyze the MonkeyContainer data set, which features a monkey head placed in the center of the container made of one of the following materials: Au, Pb, W, concrete, Cu, Fe, Air, Al, Si, or water. For each reconstructed scene, we extract the pillars and run the t-SNE algorithm with the selected perplexity value of $per = 42$ to balance the preservation of global and local data structures without excessive oversimplification.[13]

Material identification: Fig. 4 summarizes the t-SNE analysis for the *MonkeyContainer* data set. The black lines in Fig. 4(b) mark the regions in the 3D scattering-density maps from which pillars are sampled; the numeric labels in Fig. 4(a) correspond to those sampled pillars. Cluster density reflects the consistency of sampling across similar pillars, separation indicates the discriminability of materials, and outliers often correspond to artifacts or low-occupancy regions, such as those affected by high-energy muons or sparse sampling. We observe that material densities are clearly separable in the embedding, thus indicating that PoCA-based 3D scattering density maps can effectively discriminate materials when their scattering density (and underlying mass density) are sufficiently different [e.g., W/Fe vs water/air in Fig. 4(b)]. For instance, we observe a separation across high-, medium-, and low-density materials, as depicted in Fig. 4(b) with pillar 1 (W), pillar 2 (Fe), pillar 3 (water), and pillar 4 (air). Discrimination across materials of similar densities proves more challenging: t-SNE points overlap near materials of similar densities, such as Si ($\approx 2.33\ \mathrm{g/cm^3}$), Al ($\approx 2.7\ \mathrm{g/cm^3}$), and concrete ($\approx 2.4\ \mathrm{g/cm^3}$). At cluster boundaries, the t-SNE analysis shows that the effective material thickness along the muon trajectory could significantly impact machine-learning applications. For instance, we can observe that pillar sampling regions with higher path-integrated density yield much more tightly grouped points in the t-SNE analysis, as seen in pillar 1 (W) and pillar 5 (Al) samples [Fig. 4(b)]. Furthermore, note that pillars sampling from regions with reduced effective material thickness along the path, such as the ears of the monkey in pillar 6 (concrete), yield an embedding vector similar to pillars sampling from regions with greater path-integrated density in less dense materials. This effect is inherent to MST since the physical scattering caused by the muon is also dependent on the path length through the material along with its density,[1] as seen in Fig. 4(b) for pillar 5 (Al), pillar 6 (concrete), and pillar 7 (Si). Partial-volume effects occurring at boundaries and thin structures, where a pillar typically lies along smooth transitions between high- and low-scattering areas instead of forming isolated clusters, suggest that geometry and thickness influence the pillar descriptors.

Reconstruction biases: The PoCA single-interaction assumption and the broadband energy spectrum of incoming muons introduce characteristic artifacts that bias machine learning applications. Despite filtering small angles during reconstruction, a diffuse *halo* appears around objects and in low-scattering regions due to aggregated contributions from many paths and energies. An example of this phenomenon is shown in Fig. 5, where the standard PoCA 3D reconstruction shows a diffuse *halo* effect with a large number of outliers and border effects that impact pillar sampling, resulting in noisy embeddings in the t-SNE map. High-energy muons penetrate deeper and scatter less than low-energy muons; PoCA considers all energies, producing residual background in the 3D scattering-density maps. When pillars sample these zones, their descriptors inherit that background, which manifests in the t-SNE map as elongated tails and neighborhoods that are not cleanly separable, most evident for low-scattering materials such as pillar 7 (Si) and pillar 8 (Air) in Fig. 4(b). Consequently, tails in the t-SNE plot are influenced by the pillar encoding of aggregated scattering statistics from multiple muon paths in the PoCA algorithm, compounded by the geometry of the imaged object.

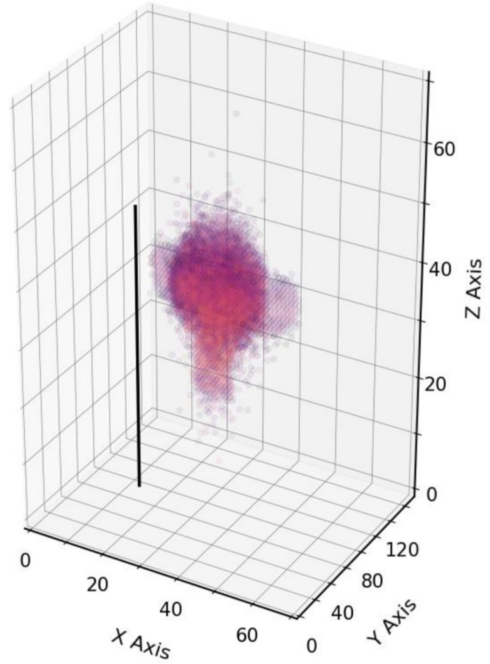

FIG. 5. Halo effect: The 3D scattering density map of a Si monkey head shows a diffuse halo, with voxel outliers and border effects that impact pillar sampling, potentially leading to an incorrect embedding. Here, axes are in voxels (1 voxel = 5 cm), and the black line is a pillar sample.

B. Multiple shapes, single materials

This section focuses on the *4ShapesContainer* data set: four geometric shapes inside the container, all made of the same material.

Proposed corrections: We correct the PoCA output in two steps. First, we pre-select simulated muons with momenta between 2 and $10\ \mathrm{GeV}/c$ to retain particles that carry most of the density-contrast signal while excluding very high energy tracks that scatter minimally. The volumes are then reconstructed with PoCA as in the previous experiments. Then, we introduce a 3D convolutional

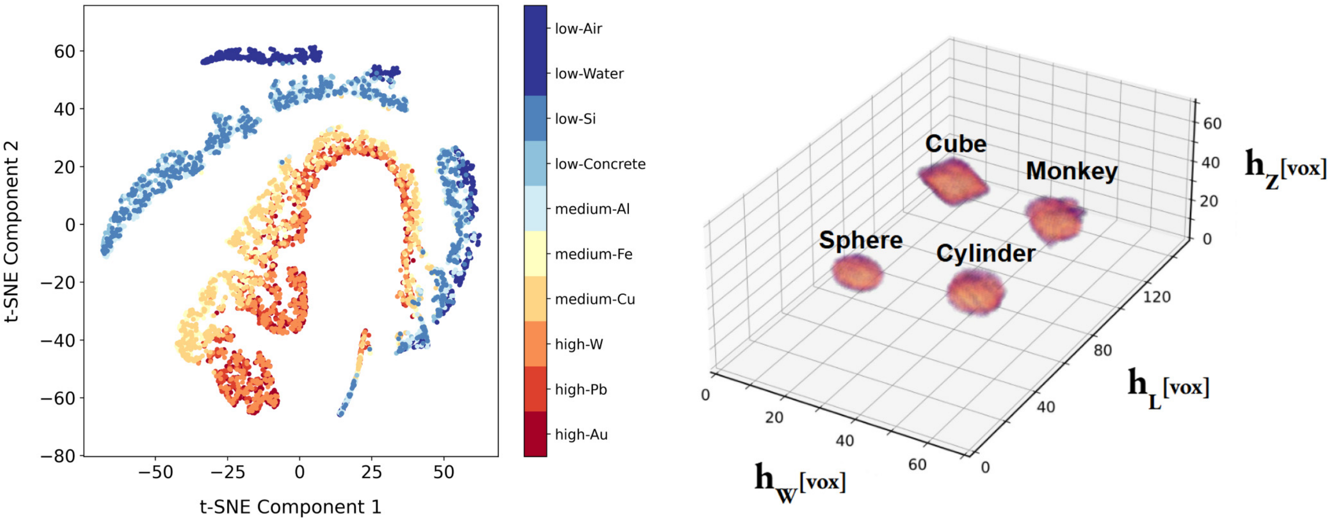

(a) (b)

FIG. 6. Effect of the proposed corrections: (a) t-SNE map of pillar embeddings for the 4ShapesContainer after applying the correction and 3D feature aggregation; colors denote material classes (legend ordered by increasing effective density). The plot is used for visualization only and shows a smoother, density-driven continuum with the previous "halo" removed due to convolutions. (b) 3D reconstructed scattering map of the water scene after the 3D convolutional kernel. Axes are in voxels [vox] with isotropic voxel size of 5 cm (1 voxel = 5 cm). The four objects (sphere, cylinder, cube, monkey) are spatially resolved.

layer with smoothing operation that ingests the 3D scattering density maps and produces a denoised, compact scattering density map representation for pillar sampling. Hence, we define as the lower and upper bounds for the offsets from the center of the kernel k_s. The 3D convolutional layer is, thus, defined as

$$M(x, y, z) = V(x + i, y + j, z + k) \mid i, j, k \in \left[-\left\lfloor\frac{k_s}{2}\right\rfloor, \left\lfloor\frac{k_s}{2}\right\rfloor\right], \qquad (2)$$

with $V(x + i, y + j, z + k)$ being the values in the filtered input 3D scattering density map, with the kernel range (i, j, k) as the displacements within the kernel size, k_s centered at the point (x, y, z). We use a $3 \times 3 \times 3$ kernel (i.e., $k_s = 3$) to exploit spatial relationships between neighboring voxels while preserving edges. For each of the denoised representations, we extract the pillars following the sampling procedure detailed in Sec. III B, and the t-SNE algorithm with a perplexity value of 42 as detailed in Sec. III C. All four geometric shapes are made from the same material, thus yielding a higher number of pillars in the entire data set compared to the previous data set.

Role of 3D convolutions and pillar sampling: Fig. 6(a) highlights the t-SNE results for the newly extracted pillars. Figure 6(b) illustrates the $3 \times 3 \times 3$ median filtering on a Water scene. The neighborhood-aware filter integrates local context from adjacent voxels, producing smoother scattering-density maps and mitigating the single-voxel concentration artifacts typical of unfiltered PoCA. After filtering, the 3D map is sampled using the proposed pillars. The filtered, more contextual pillar descriptors yield tighter, more coherent neighborhoods in the embedding when sampling all materials within the VOI. However, when material densities are close to the background, these can be filtered excessively by the convolution kernels; for instance, the number of points representing air in the t-SNE map has greatly diminished, yet the remaining ones appear grouped in the same region. Compared with Fig. 4, the median kernel preserves interfaces across all three dimensions and averages out uncorrelated noise, improving separability without excessive boundary blurring.

V. CONCLUSION

Data scarcity and reconstruction methods remain a critical bottleneck in machine-learning applications for 3D muon-scattering tomography. To address this, we presented a workflow that couples structured, pillar-based sampling with unsupervised embedding (t-SNE) to expose material-specific scattering signatures in fully synthetic B2G4 scenes. Because MST has often relied on geometry-based reconstruction, with PoCA as the mainstream algorithm, we adopt PoCA as our baseline; its single-scatter simplification epitomizes the challenges that learning methods must overcome. The results confirm that objects spanning a wide density range can indeed be separated; however, they also reveal how reconstruction bias and the broad cosmic muon energy spectrum can obscure those differences. To mitigate these effects, we (i) prefilter the muon spectrum before PoCA reconstruction to reduce noise from high-energy muons and (ii) apply 3D convolutions that fuse neighbor context to smooth out artifacts, thereby sharpening the pillar descriptors. Geometry and sampling choices still influence the final embedding, but the proposed convolutional preprocessing operator improves embedding separations. However, materials with near-identical density remain difficult to categorize, as seen in the t-SNE embeddings, suggesting the need for additional physics-informed features, momentum estimates in the reconstruction, and representation-learning models capable of mapping subtle scattering patterns into more discriminative latent spaces than is currently possible.

ACKNOWLEDGMENTS

This research was funded by the SilentBorder project under the Grant Agreement ID No. 101021812 of the European Union's Horizon 2020 research and innovation program.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

A. Bueno Rodriguez: Conceptualization (equal); Data curation (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal). **F. Sattler:** Data curation (equal); Investigation (equal); Methodology (equal); Writing – review & editing (equal). **B. Carrillo Perez:** Data curation (equal); Investigation (equal); Methodology (equal); Software (equal); Writing – review & editing (equal). **M. Perez Prada:** Data curation (equal); Investigation (equal); Validation (equal); Writing – review & editing (equal). **J. M. Alameddine:** Data curation (equal); Validation (equal); Writing – review & editing (equal). **M. Stephan:** Supervision (equal); Validation (equal); Visualization (equal); Writing – original draft (equal); Writing – review & editing (equal). **S. Barnes:** Funding acquisition (equal); Supervision (equal); Validation (equal); Writing – original draft (equal); Writing – review & editing (equal).

DATA AVAILABILITY

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

REFERENCES

[1] K. Borozdin, G. Hogan, C. Morris, W. Priedhorsky, A. Saunders, L. Schultz, and M. Teasdale, "Radiographic imaging with cosmic-ray muons," Nature **422**, 277 (2003).

[2] L. Bonechi, R. D'Alessandro, and A. Giammanco, "Atmospheric muons as an imaging tool," Rev. Phys. **5**, 100038 (2020).

[3] W. Zeng, M. Zeng, X. Pan, Z. Zeng, H. Ma, and J. Cheng, "Principle study of image reconstruction algorithms in muon tomography," J. Instrum. **15**, T02005 (2020).

[4] M. Stapleton, J. Burns, S. Quillin, and C. Steer, "Angle statistics reconstruction: A robust reconstruction algorithm for muon scattering tomography," J. Instrum. **9**, P11019 (2014).

[5] S. Barnes, A. Georgadze, A. Giammanco, M. Kiisk, V. Kudryavtsev, M. Lagrange, and O. Pinto, "Cosmic-ray tomography for border security," Instruments **7**, 13 (2023).

[6] L. Thompson, J. Stowell, S. Fargher, C. Steer, K. Loughney, E. O'Sullivan, J. Gluyas, S. Blaney, and R. Pidcock, "Muon tomography for railway tunnel imaging," Phys. Rev. Res. **2**, 023017 (2020).

[7] P. Arbol, P. Garcia, C. Gonzalez, and A. Alonso, "Non-destructive testing of industrial equipment using muon radiography," Phil. Trans. R. Soc. A **377**, 20180054 (2019).

[8] D. Poulson, J. Bacon, M. Durham, E. Guardincerri, C. Morris, and H. Trellue, "Application of muon tomography to fuel cask monitoring," Phil. Trans. R. Soc. A **377**, 20180052 (2019).

[9] M. P. Prada, S. Barnes, and M. Stephan, "Analysis of secondary particles as a complement to muon scattering measurements," Instruments **6**, 66 (2022).

[10] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrand *et al.*, "GEANT4—A simulation toolkit," Nucl. Instrum. Methods Phys. Res. Sect. A: Accel. Spectrom. Detect. Assoc. Equip. **506**, 250–303 (2003).

[11] A. B. Rodriguez, F. Sattler, M. P. Prada, M. Stephan, and S. Barnes, "B2G4: A synthetic data pipeline for the integration of Blender models in Geant4 simulation toolkit," arXiv:2311.06327 (2023).

[12] Stichting Blender Foundation, "Blender—A 3D modelling and rendering package," http://www.blender.org (2018).

[13] L. van der Maaten and G. Hinton, "Visualizing data using t-SNE," J. Mach. Learn. Res. **9**(86), 2579–2605 (2008).

[14] F. Riggi, M. Bandieramonte, U. Becciani, D. L. Bonanno, G. Bonanno, P. G. Fallica, G. Gallo, A. Grillo, P. La Rocca, D. Lo Presti, C. Petta, C. Pinto, S. Riggi, G. Romeo, G. V. Russo, G. Santagati, and G. Valvo, "Multiparametric approach to the assessment of muon tomographic results for the inspection of a full-scale container," Eur. Phys. J. Plus **136**, 139 (2021).

[15] R. Sehgal, A. Shinde, L. Pant, and B. Nayak, "Application of machine learning in muon tomography," in *Proceedings of the DAE Symposium on Nuclear Physics* (Department of Atomic Energy, 2019), Vol. 64, pp. 848.

[16] T. Stocki, C. Warren, M. Magill, B. Morgan, J. Smith, D. Ong, V. Anghel, J. Armitage, J. Botte, K. Boudjemline *et al.*, "Machine learning for the cosmic ray inspection and passive tomography project (CRIPT)," in *2012 IEEE Nuclear Science Symposium and Medical Imaging Conference Record (NSS/MIC)* (IEEE, 2012), pp. 91–94.

[17] M. Weekes, A. Alrheli, D. Barker, D. Kikoła, A. Kopp, M. Mhaidra, J. Stowell, L. Thompson, and J. Velthuis, "Material identification in nuclear waste drums using muon scattering tomography and multivariate analysis," J. Instrum. **16**, P05007 (2021).

[18] X.-Y. Pan, Y.-F. Zheng, Z. Zeng, X.-W. Wang, and J.-P. Cheng, "Experimental validation of material discrimination ability of muon scattering tomography at the TUMUTY facility," Nucl. Sci. Tech. **30**, 1–9 (2019).

[19] S. Riggi, V. Antonuccio-Delogu, M. Bandieramonte, U. Becciani, A. Costa, P. L. Rocca, P. Massimino, C. Petta, C. Pistagna, F. Riggi *et al.*, "Muon tomography imaging algorithms for nuclear threat detection inside large volume containers with the muon portal detector," Nucl. Instrum. Methods Phys. Res. Sect. A: Accel. Spectrom. Detect. Assoc. Equip. **728**, 59–68 (2013).

[20] C. Thomay, J. Velthuis, P. Baesso, D. Cussans, P. Morris, C. Steer, J. Burns, S. Quillin, and M. Stapleton, "A binned clustering algorithm to detect high-Z material using cosmic muons," J. Instrum. **8**, P10013 (2013).

[21] G. Wang, L. Schultz, and J. Qi, "Statistical image reconstruction for muon tomography using a Gaussian scale mixture model," IEEE Trans. Nucl. Sci. **56**, 2480–2486 (2009).

[22] L. Schultz, G. Blanpied, K. Borozdin, A. Fraser, N. Hengartner, A. Klimenko, C. Morris, C. Orum, and M. Sossong, "Statistical reconstruction for cosmic ray muon tomography," IEEE Trans. Image Process. **16**, 1985–1993 (2007).

[23] G. C. Strong, M. Lagrange, A. Orio, A. Bordignon, F. Bury, T. Dorigo, A. Giammanco, M. Heikal, J. Kieseler, M. Lamparth *et al.*, "Tomopt: Differential optimisation for task-and constraint-aware design of particle detectors in the context of muon tomography," Mach. Learn.: Sci. Technol. **5**, 035002 (2024).

[24] C. Poole, I. Cornelius, J. Trapp, and C. Langton, "A CAD interface for GEANT4," Australas. Phys. Eng. Sci. Med. **35**, 329–334 (2012).

[25] R. Chytracek, 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).
[26] M. Pinto and P. Gonçalves, "GUIMesh: A tool to import STEP geometries into Geant4 via GDML," Comput. Phys. Commun. **239**, 150–156 (2019).
[27] M. Constantin, D. Constantin, P. Keall, A. Narula, M. Svatos, and J. Perl, "Linking computer-aided design (CAD) to Geant4-based Monte Carlo simulations for precise implementation of complex treatment head geometries," Phys. Med. Biol. **55**, N211 (2010).
[28] Y. Wang, A. Fathi, A. Kundu, D. Ross, C. Pantofaru, T. Funkhouser, and J. Solomon, "Pillar-based object detection for autonomous driving," in *Computer Vision—ECCV 2020*, edited by A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm (Springer International Publishing, Cham, 2020), pp. 18–34.
[29] C. Hagmann, D. Lange, and D. Wright, "Cosmic-ray shower generator (CRY) for Monte Carlo transport codes," in *2007 IEEE Nuclear Science Symposium Conference Record* (IEEE, 2007), Vol. 2, pp. 1143–1146.
[30] A. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "Pointpillars: Fast encoders for object detection from point clouds," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (IEEE, 2019), pp. 12697–12705.
[31] M. Nixon and A. Aguado, *Feature Extraction and Image Processing for Computer Vision* (Academic Press, 2019).