CACTUS Project: Novel methods for the monitoring of CASTOR casks using muons

AutorInnen: *Dr. Sarah Barnes, J.-M. Alameddine, M. Pérez Prada, Á. Bueno Rodríguez, F. Sattler,* German Aerospace Center (DLR) – Institute for the Protection of Maritime Infrastructures, Bremerhaven, Germany

Kontakt: sarah.barnes@dlr.de

Introduction

Muon tomography is an emerging non-destructive technique that exploits naturally occurring cosmic-ray muons to image the internal structure of large or shielded objects. Cosmic-ray muons are highly penetrating, charged particles produced in the Earth's atmosphere, with an ambient flux of approx. 1 muon per cm^2 per min at sea level. These properties make them particularly useful for applications where more conventional radiographic methods, such as X-rays or neutrons, are ineffective or impractical. For nuclear cask monitoring, muon scattering tomography (MST) is of particular interest, utilizing the multiple Coulomb scattering of muons as they traverse matter to infer material properties. The resulting angular deflection of the muon is dependent on the radiation length of the material, and hence on its atomic number and density. By tracking muon trajectories before and after passing through an object, three-dimensional density maps can be reconstructed revealing spent nuclear fuel or shielding components. This principle was first demonstrated for nuclear safeguards applications in the early 2000s and has since been extended to a range of security and industrial problems [BHM+03, SBB+07, MAB+08, BGG+23, BAA+25, MKF+25].

Despite its promise, applying MST to large, dense objects such as nuclear storage casks with the aim of achieving single fuel rod resolution (< 1 cm) presents significant challenges. The thickness and high density of cask materials leads to extensive scattering and significant absorption, which complicates image reconstruction, rendering many traditional approaches unusable for this application. Furthermore, the passive nature of the technique implies long scan times. This adds significant overhead in terms of simulation time and presents computational challenges when considering the implementation of more complex statistical reconstruction approaches that are essential to achieve high-resolution imaging for this domain. Current state-of-the-art tomographic analyses of spent nuclear fuel casks show the ability to safely reconstruct entire fuel assemblies, but struggle to detect single fuel rods within the assemblies [AAB+25, VSMD+25, UC25].

This work addresses the challenge of single fuel rod detection within a CASTOR® V/19 storage cask by presenting a robust reconstruction and analysis framework. First, a fast and efficient statistical reconstruction method based on gradient descent optimisation is developed to recover the underlying density distribution. Subsequently, a self-supervised AI model is employed to separate signal from noise, learning spatial correlations in data to enable reliable identification of anomalous or missing fuel rods without requiring clean reference measurements. The subsequent sections will give an overview of these methodologies and present initial results using 75 hours of simulated muon tomography data.

Reconstruction Methodology

Reconstruction in MST is defined as the creation of a voxelized density map from the positions and directions of the muons inside the detector planes. In this density map, each voxel is defined by a scattering density quantifying the density and composition of the underlying material. This task is an inverse problem, which is intrinsically hard to solve due to the stochastic nature of the involved muon interactions. For many applications, it is sufficient to rely on simple geometric reconstruction approaches. However, these algorithms often underly strong simplifications, such as the PoCA algorithm [ZZP+20], which assumes that the entire scattering process happens in a single, localized point. To obtain imaging results with a higher resolution, a more sophisticated approach is necessary.

A major advance beyond purely geometric reconstruction in MST was the introduction of a statistical framework based on the Expectation-Maximization (EM) algorithm [SBB+07]. Unlike earlier methods, the EM algorithm explicitly models muon scattering interactions using a statistical model, including angular deflection and displacement, and allows prior information on the examined geometry to be incorporated into the optimization. The EM algorithm performs a global probabilistic reconstruction by jointly considering all muon tracks and their interactions with the voxelized volume, enabling improved material property estimation. A refined EM approach demonstrated superior performance for high-Z material detection in shipping containers, further strengthening its usability [RADB+13]. Despite the clear advantages in imaging, the original EM formulation suffered from severe computational and memory scaling issues due to dense muon–voxel associations [SBB+07]. This was addressed introducing sparse hash-map representations, reducing memory usage but limiting efficient GPU parallelization [RADB+13]. More recently, a sparse and scalable EM implementation integrated with geometric methods in the PyTorch framework was proposed [PGM+19, SABR+25]. This refined implementation enables scalable reconstructions with millions of voxels and muon events, as well as compatibility with modern machine-learning-based approaches.

This work employs a statistical reconstruction approach for the imaging of the CASTOR® V/19 cask, which is based on and further improves the PyTorch implementation [ASSB25, Koi10]. The core of this approach is the formulation of a multi-dimensional likelihood, similar to the formalism utilized in the original EM algorithm [SBB+07], which denotes the probability that the current density map is compatible with the detected muon tracks. The likelihood is constructed on the basis that the muon scattering within each voxel is described by a Gaussian distribution with zero-mean and a variance given by the scattering density λi. In combination with the assumed muon track from the detector hits, the likelihood is extended to describe the influence of each passed voxel and its scattering density on the accumulated deflection of the entire muon path. This way, the entire muon track information is processed based on a well-motivated statistical foundation. As all muon events are independent of each other, the likelihoods from each event are combined to obtain a single likelihood that includes the information for the entire cask measurement. Since this likelihood describes the agreement of the density map with the detected muon hits, the scattering density values λi are treated as free parameters that need to be optimized. The optimal reconstruction result is obtained by maximizing the likelihood iteratively using automatic differentiation. This technique allows for a quick, uncertainty-free evaluation of the gradients, and therefore enables the optimization of λi even for millions of voxel parameters. With that, a powerful reconstruction method is obtained, which efficiently uses the available statistics to create an image of the scanned cask. Notably, the iterative character of this approach enables processing the muon data in chunks. Firstly, this allows the approach to be scalable, as the chunk size can be adapted to comply with the available computational resources. Secondly, for real-life measurements, the reconstruction can be performed in real-time, since not all data batches need to be available at the beginning of the reconstruction process.

Simulation Setup and Results

The interactions between particles and matter are simulated with the GEANT4 Monte Carlo simulation toolkit [AAA+03, Koi10]. The muons from cosmic-ray air showers are created with the EcoMug generator [PBD+21] to reproduce the natural muon flux at sea level. In order to obtain the realistic muon flux through the spent nuclear fuel cask from all directions, 21.48 billion muons are generated from a hemisphere-shaped surface with a diameter of 7.5 m, which is equivalent to an exposure time of around 75 hours. The cask is integrated into the simulation with the B2G4 toolkit [18] using a simplified three-dimensional model based on the dimensions and materials of the CASTOR® V/19 cask [GNS25, NS25]. Two different cask configurations are simulated: a cask fully loaded with all possible fuel rods inside the basket and a cask randomly loaded with 90% of all possible fuel rods inside the basket. The fuel rod configurations for these two scenarios are shown in Fig. 1.

The CASTOR® is placed inside a detector configuration covering the entirety of the cask. The detectors are based on three infinitesimally thin detector layers with a 10 cm spacing forming a detector set. Four of these detector sets with an area of 3x7 m^2 are placed vertically on the sides of the cask, while one detector set with an area of 3x3 m^2 is placed horizontally above and one set below the cask. For simplicity, each layer is assumed to hold perfect spatial resolution and 100% muon detection efficiency, and the muon information and kinematics are directly retrieved from GEANT4. Muons failing to intersect all three detector layers for a given detector set are excluded from the analysis and events with a muon passing through only one detector set are vetoed. The incoming and outgoing positions and trajectories of the muon into the volume of interest created by the detectors is estimated with a least-squares-fit based on the three positions across the detection layers in a detector set.

For each simulated scene, the gradient descent reconstruction method described in the previous section is applied to reconstruct the voxel density maps. To allow for the processing to be computationally feasible, the dataset is split into 2500 randomly selected batches, which are processed sequentially. The momenta of the muons are assumed to be known precisely, although reconstruction is also possible with less accurate muon momentum information. The reconstruction results are shown in Fig. 2, showing a clear picture of the cask and its features. Differences between the fully and partially loaded casks are visible by eye.

Single Rod Analysis

In the context of muon tomography, self-supervised learning is essential, as one fundamental limitation is the lack of a reliable ground truth to support conventional supervised approaches. Supervised methods tend to suffer from two significant drawbacks: They impose assumptions that do not accurately reflect noise statistics, and they risk introducing simulation-driven biases into the reconstructed structure.

Self-supervised frameworks, by contrast, learn directly from noisy data, exploiting the intrinsic statistical properties such as the spatial coherence of the physical signal and the incoherence of the noise. To this end, a self-supervised deep learning model is proposed to learn directly from raw data without requiring labelled reference data. In the context of single fuel rod detection for nuclear fuel casks, this is paramount, as the reconstructions provide a tomographic image of the cask's current state, for which a ground truth is not readily available.

The workflow utilizes a two-dimensional residual U-Net high-capacity deep learning architecture [RPB15, AYH+19], which is trained independently for each cask configuration to reconstruct masked pixels from their local context. The model is built with a three-level depth, residual connections, and 64 channels, optimized for stable convergence. An attention mechanism is incorporated for dynamic recalibration

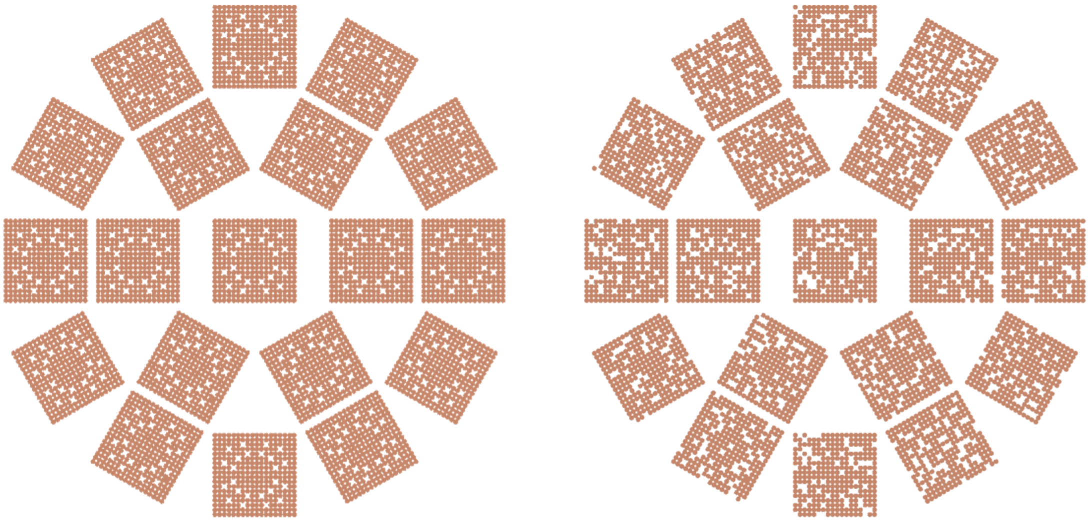

Fig. 1: Fuel rod configuration of the fully (a) and the partially (b) loaded CASTOR® cask. Each orange dot represents one fuel rod. The fuel assembly layout is based on the configuration of the CASTOR® V/19 cask [19,20]

of feature maps within the neural network [WZL21, JZ22]. The training of this neural network follows canonical self-supervised approaches, yet is adapted for the denoising task at hand. First, a set of muon tomography projections are derived from the gradient descent reconstruction results, using 64x64 pixel patches with a 32-pixel overlap. Then, for each of these projections, stochastic masking is applied to 20% of the pixels within an 11x11 pixel sized region of interest, excluding the central pixel, to create a blind-spot configuration. The loss is computed based on these masked pixels, ensuring that the network learns to predict their values solely from the surrounding spatial context. By doing this, the network is forced to infer the underlying signal from its surrounding spatial context.

The effectiveness of this workflow relies on the assumption that the detector noise is statistically independent across pixels, whereas the rod geometry exhibits strong spatial autocorrelation. The proposed architecture exploits this assumption to discriminate between the signal, representing fuel rods, and stochastic noise. Critically, the attention mechanism and the network depth allow the circular structure of the rods to be modeled as a strong topological prior. This ensures that, in the physical absence of a rod, the model does not interpret the resulting void as noise to be suppressed or induce structural hallucinations. Fig. 3 and Fig. 4 show the denoising results from the reconstruction of the fully and the partially loaded cask. The results presented in Fig. 4 show clearly that the reconstruction and analysis methods allow for a successful discrimination between single rods and voids for the partially loaded cask scenario at various fuel rod locations.

Conclusions and Next Steps

In this work, the feasibility to identify individual fuel rods in a CASTOR® V/19 cask using MST has been demonstrated for the first time. A simplified, but realistic 3D model of the CASTOR V/19 cask has been used to generate two MST datasets with 75 h scan time each: One simulation with a fully loaded fuel basket and one with a 90% partially loaded cask. To reconstruct individual fuel rods from these simulations, a workflow consisting of two key parts has been developed: The first part is an efficient implementation of a gradient-descent maximum likelihood optimization for a powerful and scalable statistical recon-

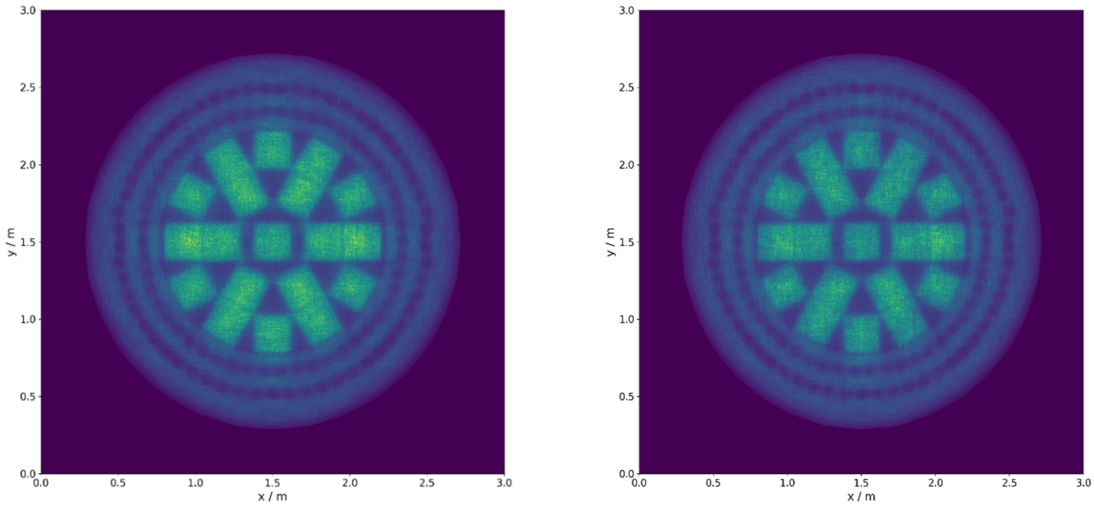

(a) **(b)**

Fig. 2: Two-dimensional projection of the reconstruction results for the fully (a) and partially (b) loaded CASTOR® cask , using a voxel size of 15 mm and based on simulation data corresponding to 75 h of scan time. Brighter colours correspond to higher density scattering values, i.e. a higher material density. Features such as the detailed assembly configuration and the moderator rods in the cask are clearly visible.

struction. This allowed to obtain a high-resolution density map, which was passed to a self-supervised neural network capable of distinguishing between fuel rods and background solely on the features of the available data.

Using the B2G4 toolkit, large-scale datasets of realistic and relevant scenarios of spent nuclear fuel casks can be generated. Based on these datasets, further general capabilities of the workflow can be evaluated and, if necessary, improved. The developed statistical reconstruction method currently works without any prior information. However, since an approximate ground truth of the cask is known, incorporating this knowledge into the reconstruction and machine learning analysis methods will be crucial to further improve the achievable image quality. In order to verify the performance of the workflow under real-world conditions, conducting experimental measurement with real casks will be required.

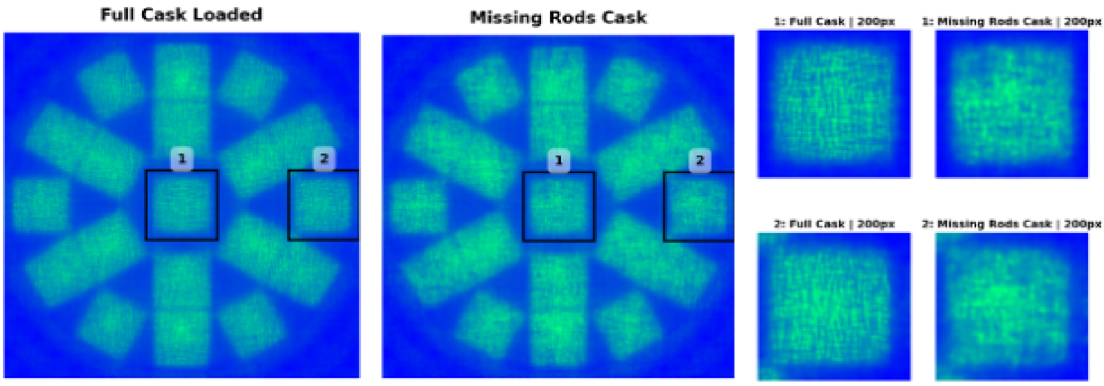

(a) (b) (c)

Fig. 3: The denoised reconstruction of the fully (a) and partially (b) loaded cask , as well as two zoomed-in regions (c). The gaps with lower scattering density corresponding to missing fuel rods in the partially loaded cask compared to the fully loaded scenario are visible.

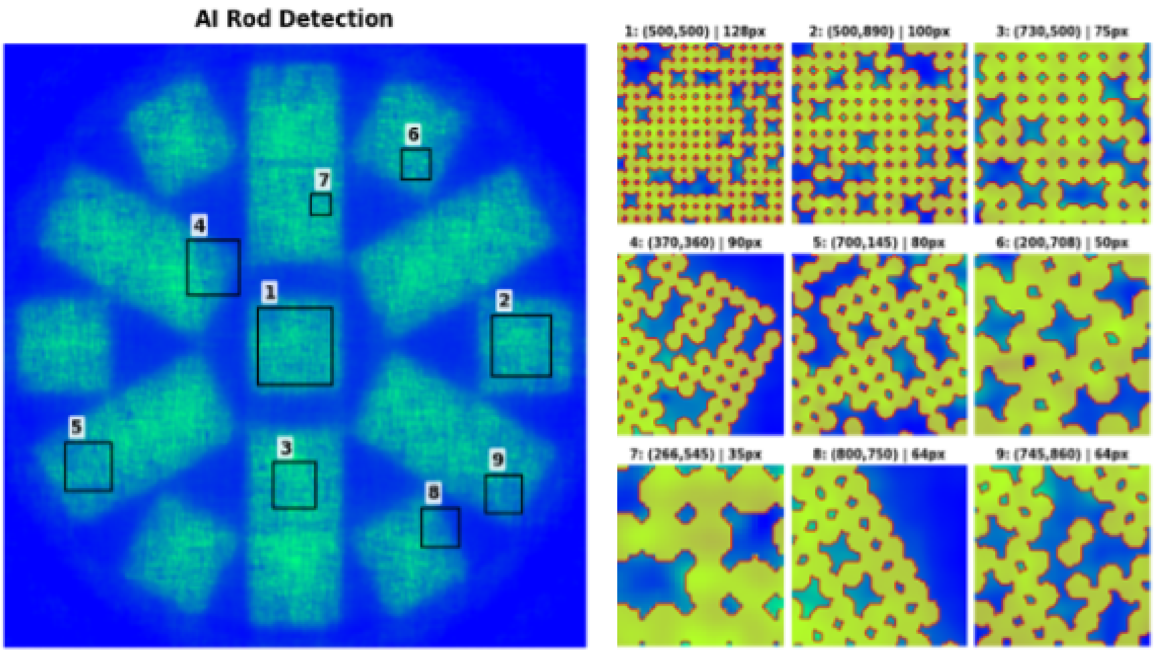

(a) (b)

Fig. 4: The denoised reconstruction of the partially (a) loaded CASTOR® cask and (b) nine zoomed-in regions with an overlay of the corresponding fuel rod configuration. The gaps with lower scattering density corresponding to missing fuel rods are clearly visible.

Bibliography

[AAA+03] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrand, F. Behner, L. Bellagamba, J. Boudreau, L. Broglia, A. Brunengo, H. Burkhardt, S. Chauvie, J. Chuma, R. Chytracek, G. Cooperman, G. Cosmo, P. Degtyarenko, A. Dell'Acqua, G. Depaola, D. Dietrich, R. Enami, A. Feliciello, C. Ferguson, H. Fesefeldt, G. Folger, F. Foppiano, A. Forti, S. Garelli, S. Giani, R. Giannitrapani, D. Gibin, J.J. Gómez Cadenas, I. González, G. Gracia Abril, G. Greeniaus, W. Greiner, V. Grichine, A. Grossheim, S. Guatelli, P. Gumplinger, R. Hamatsu, K. Hashimoto, H. Hasui, A. Heikkinen, A. Howard, V. Ivanchenko, A. Johnson, F.W. Jones, J. Kallenbach, N. Kanaya, M. Kawabata, Y. Kawabata, M. Kawaguti, S. Kelner, P. Kent, A. Kimura, T. Kodama, R. Kokoulin, M. Kossov, H. Kurashige, E. Lamanna, T. Lampén, V. Lara, V. Lefebure, F. Lei, M. Liendl, W. Lockman, F. Longo, S. Magni, M. Maire, E. Medernach, K. Minamimoto, P. Mora de Freitas, Y. Morita, K. Murakami, M. Nagamatu, R. Nartallo, P. Nieminen, T. Nishimura, K. Ohtsubo, M. Okamura, S. O'Neale, Y. Oohata, K. Paech, J. Perl, A. Pfeiffer, M.G. Pia, F. Ranjard, A. Rybin, S. Sadilov, E. Di Salvo, G. Santin, T. Sasaki, N. Savvas, Y. Sawada, S. Scherer, S. Sei, V. Sirotenko, D. Smith, N. Starkov, H. Stoecker, J. Sulkimo, M. Takahata, S. Tanaka, E. Tcherniaev, E. Safai Tehrani, M. Tropeano, P. Truscott, H. Uno, L. Urban, P. Urban, M. Verderi, A. Walkden, W. Wander, H. Weber, J.P. Wellisch, T. Wenaus, D.C. Williams, D. Wright, T. Yamada, H. Yoshida, and D. Zschiesche. Geant4—a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003.

[AAB+25] P. Andreetto, K. Aymanns, M. Balling, M. Benettoni, N. Bez, G. Bonomi, L. Castellani, P. Checchia, E. Conti, J. Dackner, F. Gonella, A. Jussofie, A. Lorenzon, F. Montecassiano, M. Mosconi, M. Murtezi, J. Niedermeier, I. Niemeyer, J. Pekkarinen, A. Rigoni Garola, D. Scarpa, M. Stuke, M. Turcato, and G. Zumerle. Results on re-verification tests of spent fuel casks with muon tomography: Mutomca project. *Journal of Applied Physics*, 138(12):124903, 09 2025.

[ASSB25] Jean-Marco Alameddine, Felix Sattler, Maurice Stephan, and Sarah Barnes. Gradient-descent-based reconstruction for muon tomography based on automatic differentiation in PyTorch. In *5th MODE Workshop on Differentiable Programming for Experiment Design*, 8–13 June 2025.

[AYH+19] Md Zahangir Alom, Chris Yakopcic, Mahmudul Hasan, Tarek M. Taha, and Vijayan K. Asari. Recurrent residual U-Net for medical image segmentation. *Journal of Medical Imaging*, 6(1):014006, 2019.

[BAA+25] D. Borselli, F. Ambrosino, P. Andreetto, L. Bonechi, G. Bonomi, S. Bottai, T. Buhles, I. Calliari, P. Checchia, U. Chiarotti, C. Cialdai, R. Ciaranfi, L. Cimmino, V. Ciulli, R. D'Alessandro, P. G. De Seta Cosentino, E. Faraci, F. Finke, A. Franzen, C. Frosin, C. Gennari, B. Glaser, S. Gonzi, A. Lorenzon, V. Masone, V. Moroli, O. Nechyporuk, A. Paccagnella, L. Pezzato, B. V. Rangavittal, D. Ressegotti, G. Saracino, R. P. Santos Ferreira, J. Sauerwarld, O. Starodubtsev, H. Upadhyay, F. Volzone, and M. Vynnycky. Muon imaging of a blast furnace: The european project BLEMAB. *Journal of Applied Physics*, 138(8):084902, 08 2025.

[BGG+23] Sarah Barnes, Anzori Georgadze, Andrea Giammanco, Madis Kiisk, Vitaly A. Kudryavtsev, Maxime Lagrange, and Olin Lyod Pinto. Cosmic-ray tomography for border security. *Instruments*, 7(1), 2023.

[BHM+03] Konstantin N. Borozdin, Gary E. Hogan, Christopher Morris, William C. Priedhorsky, Alexander Saunders, Larry J. Schultz, and Margaret E. Teasdale. Radiographic imaging with cosmic-ray muons. *Nature*, 422(6929):277–277, 2003.

[GNS25] GNS. Product info CASTOR V/19. https://www.gns.de/behaelter-equipment/brennelemente-haw/lwr/castor-v/castor-v19/, 2025. aufgerufen am 20.11.2025.

[JZ22] David John and Ce Zhang. An attention-based u-net for detecting deforestation within satellite sensor imagery. *International Journal of Applied Earth Observation and Geoinformation*, 107:102685, 2022.

[Koi10] T. Koi. A geant4 physics list for shielding calculation. In *Proceedings of the Shielding Aspects of Accelerators, Targets and Irradiation Facilities (SATIF)*. NEA, 2-4 June 2010.

[MAB+08] C. L. MORRIS, C. C. ALEXANDER, J. D. BACON, K. N. BOROZDIN, D. J. CLARK, R. CHARTRAND, C. J. ESPINOZA, A. M. FRASER, M. C. GALASSI, J. A. GREEN, J. S. GONZALES, J. J. GOMEZ, N. W. HENGARTNER, G. E. HOGAN, A. V. KLIMENKO, M. F. MAKELA, P. McGAUGHEY, J. J. MEDINA, F. E. PAZUCHANICS, W. C. PRIEDHORSKY, J. C. RAMSEY, A. SAUNDERS, R. C. SCHIRATO, L. J. SCHULTZ, M. J. SOSSONG, and G. S. BLANPIED. Tomographic imaging with cosmic ray muons. *Science & Global Security*, 16(1-2):37–53, 2008.

[MKF+25] Haruo Miyadera, Naoto Kume, Takuro Fujimaki, Yuki Nakai, Kyohei Noguchi, Souichi Ueno, Tsukasa Sugita, Kenji Kurihara, Takahiro Tanaka, Ryunosuke Mizuguchi, and Masaki Yoda. Cosmic-ray muon scattering method to evaluate nuclear materials in fuel debris. *Journal of Applied Physics*, 138(14):144905, 10 2025.

[NS25] J. Niedermeier and M. Stuke. Multiple-scattering muography of a thick-walled spent-fuel cask: A high statistics geant4 simulation. *Journal of Applied Physics*, Special Topic: Muography: Discoveries, Innovations, and Applications, 2025.

[PBD+21] D. Pagano, G. Bonomi, A. Donzella, A. Zenoni, G. Zumerle, and N. Zurlo. EcoMug: An efficient cosmic muon generator for cosmic-ray muon applications. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1014:165732, 2021.

[PGM+19] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. *PyTorch: an imperative style, high-performance deep learning library*. Curran Associates Inc., Red Hook, NY, USA, 2019.

[RADB+13] S. Riggi, V. Antonuccio-Delogu, M. Bandieramonte, U. Becciani, A. Costa, P. La Rocca, P. Massimino, C. Petta, C. Pistagna, F. Riggi, E. Sciacca, and F. Vitello. Muon tomography imaging algorithms for nuclear threat detection inside large volume containers with the muon portal detector. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 728:59–68, 2013.

[RPB15] O. Ronneberger, P.Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, volume 9351 of *LNCS*, pages 234–241. Springer, 2015. (available on arXiv:1505.04597 [cs.CV]).

[SABR+25] F. A. Sattler, J. M. Alameddine, Á. Bueno Rodríguez, M. Stephan, and S. Barnes. A comprehensive framework toward the seamless integration of muon reconstruction algorithms with machine learning. *Journal of Applied Physics*, 138(14):144904, 10 2025.

[SBB+07] Larry J. Schultz, Gary S. Blanpied, Konstantin N. Borozdin, Andrew M. Fraser, Nicolas W. Hengartner, Alexei V. Klimenko, Christopher L. Morris, Chris Orum, and Michael J. Sossong. Statistical reconstruction for cosmic ray muon tomography. *IEEE Transactions on Image Processing*, 16(8):1985–1993, 2007.

[UC25] Reshma Ughade and Stylianos Chatzidakis. µtrec: A muon trajectory reconstruction algorithm for enhanced scattering tomography. *Journal of Applied Physics*, 138(6):064909, 08 2025.

[VSMD+25] Jesus J. Valencia, Jacob W. Sperow, J. Matthew Durham, Andrew G. Osborne, Chris L. Morris, Daniel Poulson, and Adam A. Hecht. Simulations of muon imaging with the lanl gmt detector for spent nuclear fuel cask content verification. *Journal of Applied Physics*, 138(18):184502, 11 2025.

[WZL21] Zekun Wang, Yanni Zou, and Peter X. Liu. Hybrid dilation and attention residual u-net for medical image segmentation. *Computers in Biology and Medicine*, 134:104449, 2021.

[ZZP+20] W. Zeng, M. Zeng, X. Pan, Z. Zeng, H. Ma, and J. Cheng. Principle study of image reconstruction algorithms in muon tomography. *Journal of Instrumentation*, 15(02):T02005, feb 2020.