

Geant4-Based Gamma-Radiography Detectability of Defects in Aerospace CFRP Panels: Analytical Benchmarking and Electromagnetic-Model Robustness

Yusuf Havvat

Department of Physics, Cukurova University, Adana 01330, Türkiye
havvatyusuf249@gmail.com

<https://doi.org/10.62753/ctp.2026.05.3.3>

Abstract

Carbon fibre reinforced polymer (CFRP) structures may contain internal delaminations, foreign inclusions, local thickness or density variations, and bond-line discontinuities that are difficult to distinguish upon external inspection. A controlled Geant4/ROOT benchmark was developed to quantify how five representative defect classes appear in single-projection gamma radiography under a common exposure budget. A 300 keV parallel gamma field rastered a $100 \times 100 \times 10 \text{ mm}^3$ homogeneous-equivalent CFRP panel onto a pixelated silicon transmission plane. Seven independent four-million-history simulations comprised two matched references and five defect conditions. Defect contrast was evaluated both in a pre-specified region of interest (ROI) and by a two-sided blind scan over 141,094 position-scale candidates, with a Bonferroni global 5σ decision criterion. At four million histories, 2 mm delamination and 2 mm surface thinning produced ROI contrasts of 2.579 ± 0.147 and 2.441 ± 0.147 percentage points, with blind global significances of 16.88σ and 15.87σ , respectively. A 10% density deficit reached 7.99σ , while a 2 mm aluminium inclusion generated a sign-reversed contrast and only narrowly passed the blind threshold at 5.12σ . A 1 mm air gap inside a 1.2 mm epoxy bond line remained below the threshold (1.51σ global), despite a physically consistent positive ROI contrast. As an internal verification, the seven modelled configurations were repeated with three electromagnetic physics constructors, giving 21 independent runs with 4×10^6 incident histories each. The uncollided-primary transmission agreed with the Beer–Lambert predictions within 0.36 percentage points, while the defect-contrast signs and broad hierarchy remained stable, with a maximum inter-model spread of 0.24 percentage points. The study provides a reproducible numerical ranking of idealized defect visibility and clearly separates known-location sensitivity from blind detectability; it is a simulation benchmark rather than an experimentally validated probability-of-detection qualification.

Keywords: carbon-fibre-reinforced polymer; gamma radiography; non-destructive testing; Geant4; blind detection; Monte Carlo simulation

1 Introduction

Carbon-fibre-reinforced polymer structures combine low mass with high specific stiffness and strength, making them attractive for aerospace panels, fairings, pressure-containing structures, and launch-vehicle components. Their manufacturing and joining processes can nevertheless introduce voids, delaminations, foreign material, resin or density variations, local thickness loss, and adhesive disbonds. Such discontinuities may alter structural performance while remaining hidden beneath an intact external surface. Consequently, non-destructive testing (NDT) is required throughout manufacture and service. No single NDT modality is universally optimal because detectability depends on the defect composition, geometry, depth, accessibility, and the relevant contrast mechanism [1–3].

Radiographic methods directly probe projected attenuation and provide a non-contact view of the internal areal-density distribution. X-ray computed tomography (CT) has been used to characterize porosity, adhesion defects, and damage in carbon-based composites [4–6]; projection radiography has also shown that CFRP damage often appears as a low-contrast signal requiring careful image processing [7]. A single projection is faster and simpler than CT but collapses depth information. It can therefore detect an attenuation anomaly without necessarily identifying whether the missing material is internal or located at a surface.

Monte Carlo transport provides a controlled route for separating these physical and statistical effects before expensive experimental trials. Geant4 supports detailed geometry, material construction, electromagnetic transport, and event-level scoring [8–10], and recent work has demonstrated its use as the transport core of digital radiography simulation platforms [11]. The present work does not attempt to reproduce a particular commercial scanner or a proprietary CFRP lay-up. Instead, it asks a narrower benchmark question: under identical idealized 300 keV projection conditions, which representative defect classes generate a statistically reliable localized signal, and how does the answer change when the defect position and size are unknown?

The contribution is threefold. First, five physically distinct defects are embedded in matched CFRP or bonded-joint reference families and evaluated with a common geometry and exposure scale. Second, local contrast is defined as defect–reference difference-in-differences, suppressing small global normalization shifts. Third, a pre-specified ROI result is reported separately from a two-sided, multi-scale blind search with an explicit trials correction. This avoids interpreting a high significance obtained at a known defect location as automatic blind detectability. In addition to the detectability analysis, analytical Beer–Lambert predictions and three alternative electromagnetic physics constructors were used to examine the internal physical consistency and model dependence of the reported results. These checks are treated as analytical and numerical verification rather than as a substitute for experimental validation.

2 Materials and methods

2.1 Transport model and imaging geometry

The simulation was implemented with Geant4 11.2.2 and analysed with ROOT 6.34. [12]. The reference physics list was FTFP_BERT; its default electromagnetic constructor was replaced by G4EmLivermorePhysics to model low-energy photon transport. The beam was monoenergetic at 300 keV, normally incident, and directed along +z. The source plane at $z = -200$ mm and the detector plane at $z = +100$ mm each covered 100×100 mm². The nominal panel was centred at $z = 0$ and was 10 mm thick (Fig. 1).

The 300 keV energy was selected as a representative controlled benchmark rather than as an optimized inspection energy. For the modelled 10 mm homogeneous-equivalent CFRP panel, the corresponding analytical transmission is approximately 0.84, retaining substantial transmitted counts while preserving measurable attenuation differences between the prescribed material configurations. The beam energy, panel thickness, and the 1 mm detector-pixel pitch were held fixed to isolate the effects of the defect class and photon-history count. A systematic energy–thickness–pixel-size sensitivity study would be required to optimize a particular inspection system and lies outside the scope of this benchmark.




Figure 1. Normal-incidence single-projection geometry. Balanced planar raster emits equal histories into each 1×1 mm² source pixel. The highlighted region shows the common 15.6×15.6 mm² final-defect footprint; its through-thickness material depends on the defect class.

The source field was divided into 100×100 pixels. Consecutive blocks of 10,000 events illuminated every pixel exactly once, with a uniform random sub-pixel offset. This construction gave exactly 50, 100, 200, and 400 incident histories per pixel for the nested 0.5, 1, 2, and 4 million-history exposures. A $100 \times 100 \times 0.5$ mm³ silicon

slab represented the transmission plane. A history was counted as detected when at least one gamma crossed the detector entrance boundary; no deposited-energy threshold, detector blur, electronic noise, or energy-resolution response was imposed.

2.2 Equivalent CFRP and defect models

The nominal panel was a homogeneous-equivalent CFRP. Carbon fibre and a reproducible DGEBA-based epoxy surrogate were assigned densities of $1.80 \text{ g}\cdot\text{cm}^{-3}$ and $1.20 \text{ g}\cdot\text{cm}^{-3}$, respectively. Fibre and epoxy volume fractions of 0.60 and 0.40 yield a composite density of $1.56 \text{ g}\cdot\text{cm}^{-3}$ and mass fractions of 0.6923 and 0.3077. The epoxy surrogate composition was $\text{C}_{21}\text{H}_{24}\text{O}_4$; it is not intended to represent a specific cured commercial resin.

Five final defect conditions were fixed before production of the final random seeds (Table 1). All square defects had a 15.6 mm side and centre $(x, y) = (+10.9, +5.1)$ mm. Air-filled delamination and an aluminium inclusion replaced the same CFRP volume. Surface thinning was produced by Boolean removal from the $+z$ face, while the density anomaly retained a nominal composition and external dimensions but used 90% of the CFRP density. The bonded-joint family contained two 4.40 mm CFRP adherends and a 1.20 mm central epoxy layer. Its defect was a 1.00 mm local air gap inside that layer, leaving 0.20 mm of epoxy in the ROI. The bonded defect was compared only with the intact bonded reference, never with the homogeneous panel.

Table 1. Locked final defect definitions.

Case	Local replacement/removal	Matched reference
Delamination	2.0 mm CFRP replaced by air	Homogeneous CFRP
Al inclusion	2.0 mm CFRP replaced by Al	Homogeneous CFRP
Thinning	2.0 mm removed at exit face	Homogeneous CFRP
Density deficit	9.98 mm at $0.90\rho_0$	Homogeneous CFRP
Bond-line gap	1.0 mm air in 1.2 mm epoxy	Intact bonded panel

Seven independent four-million-history simulations were produced: one homogeneous reference, four homogeneous-panel defects, one intact bonded reference, and one bonded defect. Their random seeds were independent of the exploratory geometry, size, exposure, and blind-search studies used to fix the analysis. The final files stored cumulative incident and detected count maps rather than event-level rows. The data validation required exactly four million histories per case, equal per-pixel exposure, integer counts, and $0 \leq N_{\text{det}} \leq N_{\text{inc}}$ in every bin.

2.3 Contrast, uncertainty, and blind search

For a set of incident and detected counts, the transmission was $T = N_{\text{det}}/N_{\text{inc}}$, with binomial standard uncertainty $\sigma_T = [T(1 - T)/N_{\text{inc}}]^{1/2}$. The contrast for a candidate region R , relative to outside region O , was defined as

$$\Delta T = (T_R - T_O)_{\text{defect}} - (T_R - T_O)_{\text{reference}}. \quad (1)$$

The four binomial variances were summed for $\sigma_{\Delta T}$. The known-location contrast-to-noise ratio was $|\Delta T|/\sigma_{\Delta T}$, using the true $15.6 \times 15.6 \text{ mm}^2$ footprint. Because this test employs prior knowledge of the defect, it carries no search-trials penalty and is not called a blind detection.

For the blind analysis, every integer square side from 6 to 24 mm was translated in 1 mm increments over all the positions that remained inside the image. This yielded 141,094 position-scale candidates. Contrast sign, position, and size were not supplied to the search. The candidate with the largest $|z_{\text{local}}| = |\Delta T|/\sigma_{\Delta T}$ was selected. Its two-sided local probability was $p_{\text{local}} = \text{erfc}(|z_{\text{local}}|/\sqrt{2})$. Conservative family-wise probability and equivalent two-sided significance were then calculated as

$$p_{\text{global}} = \min(1, 141094 p_{\text{local}}) \quad (2)$$

$$Z_{\text{global}} = \sqrt{2} \text{erfc}^{-1}(p_{\text{global}}) \quad (3)$$

A blind detection required $Z_{\text{global}} \geq 5$. The reported first stable threshold was the first tested exposure at which the threshold was met and remained met at all higher tested prefixes. Since the exposure points were nested prefixes of the same run, they show information accumulation and are not independent repeat measurements or a formal probability-of-detection (PoD) curve.

Before the final runs, the unchanged blind algorithm was also tested on a mixed set of ten delamination cases and ten defect-free controls at one million histories per case. This small test was included solely as an illustrative end-to-end check of the pre-declared search procedure; it was not used to support the final defect ranking or to estimate the population-level inspection performance. Five of the ten defects and all the ten controls were classified correctly, corresponding to 50% sensitivity (exact 95% interval 18.7–81.3%) and a 0% observed false-positive rate (95% interval 0–30.9%). Among the true positives, the mean localization error was 0.96 mm and the mean intersection-over-union was 0.855. With only ten cases in each class and consequently wide confidence intervals, these values are statistically under-powered and must not be interpreted as a PoD estimate or as evidence of a low population false-positive rate.

2.4 Analytical and model cross-checks

Two complementary internal checks were performed to examine the physical consistency of the transport calculation and its dependence on the selected electromagnetic constructor. The homogeneous reference, 2.0-mm delamination, 2.0-mm Al inclusion, 2.0-mm surface thinning, 10% density deficit, intact bond line, and 1.0-mm bond-line gap configurations were simulated using G4EmLivermorePhysics, G4EmPenelopePhysics, and G4EmStandardPhysics_option4. Each of the resulting 21 model–geometry combinations was evaluated with 4×10^6 incident photon histories. The source definition, material geometry, scoring plane, image pixelization, regions of interest, and event-selection criteria were unchanged between the three constructors.

For a layered material sequence, the analytical uncollided transmission was calculated from the Beer–Lambert relation,

$$T_{\text{BL}} = \exp \left[- \sum_i \left(\frac{\mu}{\rho} \right)_i \rho_i t_i \right], \quad (4)$$

where $(\mu/\rho)_i$, ρ_i , and t_i denote the mass attenuation coefficient, density, and photon path length of layer i , respectively. The calculation utilised the modelled layer thicknesses, densities, material compositions, and 300-keV attenuation coefficients. Agreement was quantified by means of the signed transmission residual

$$\Delta T = 100(T_{\text{G4}}^{\text{uncollided}} - T_{\text{BL}}), \quad (5)$$

expressed in percentage points. The Beer–Lambert calculation describes uncollided-primary attenuation. By contrast, the detected-history contrast used in the radiographic assessment includes accepted transmitted histories according to the detector-plane scoring definition and can therefore contain contributions from interacted photons. Exact agreement between these two contrast definitions is consequently not expected. The Beer–Lambert comparison tests the consistency of the modelled attenuation and primary-photon accounting, whereas the three-constructor comparison evaluates the robustness of the detected-history results against the electromagnetic-model choice. Neither check constitutes experimental validation or a certified probability-of-detection analysis.

3 Results

3.1 Radiographs and localized contrast

All seven final data files passed the count and exposure checks. The whole-image transmissions were close because the defect footprint occupied only about 2.4% of the panel area: 0.861659 for the homogeneous CFRP and 0.864506 for the intact bonded panel, with the defective cases spanning 0.861421–0.865262. Consequently, global transmission alone did not provide a meaningful defect decision. Source-normalized I/I_0 radiographs and defect–reference maps localized the changes (Figs. 2 and 3). For demonstration only, the counts were summed in 5×5 original-pixel blocks before division, giving 10,000 incident histories per displayed cell at the four-million-history exposure.




Figure 2. Four-million-history source-normalized transmission radiographs for two references and five locked defect cases. Demonstration rebinning preserves incident and detected count totals.




Figure 3. Defect-minus-matched-reference transmission maps. Delamination, thinning, density deficit, and bond-line gap increase transmission; aluminium produces the opposite sign. Common matched-reference rule is used within each physical panel family.




Figure 4. Local standardized residual maps after 5×5 pixel count-preserving display aggregation. These cell-level values visualize where contrast accumulates; they are not the 141,094-trial-corrected blind significances used for the formal decisions.

The sign and magnitude ordering were physically consistent (Table 2). Replacing CFRP or epoxy with air and reducing the local areal density increased the transmission. The embedded aluminium replacement increased attenuation and therefore generated a negative contrast. At four million histories, the delamination and thinning had the largest absolute ROI contrasts (2.579 and 2.441 percentage points), followed by the density deficit (1.390 percentage points), aluminium (1.082 percentage points), and the bond-line gap (0.719 percentage points). The standardized residual maps in Fig. 4 show the same spatial ordering, but no formal image-wide claim was made from the largest single display cell.

Table 2. Final four-million-history results and first stable tested 5σ exposure. Contrast is defect minus matched reference according to Eq. (1). “—” means the criterion was not reached up to four million histories.

Defect	ΔT (percentage points)	ROI CNR	Blind local $ Z $	Blind global Z	Known ROI	Blind global
2 mm delamination	2.579 ± 0.147	17.57	17.57	16.88	0.5M	0.5M
2 mm surface thinning	2.441 ± 0.147	16.59	16.59	15.87	0.5M	1M
10% density deficit	1.390 ± 0.150	9.28	9.34	7.99	2M	2M
2 mm Al inclusion	-1.082 ± 0.156	6.96	7.02	5.12	4M	4M
1 mm bond-line gap	0.719 ± 0.149	4.84	4.91	1.51	—	—

3.2 Exposure dependence and blind decisions

The known-ROI and blind thresholds gave different practical conclusions (Fig. 5). Delamination was robustly detectable from the lowest tested exposure. Surface thinning passed the known-ROI criterion at 0.5 million histories but required one million for a stable blind detection. The 10% density deficit fulfilled both criteria at two million histories. Aluminium was the weakest successful blind case: its global significance was 5.12σ at four million histories and did not meet the criterion at a lower tested prefix. The bond-line gap remained below even the known-ROI 5σ level at four million histories, and its blind global significance fell to 1.51σ after accounting for the search volume.




Figure 5. (a) Signed known-ROI contrast. (b) Bonferroni-corrected blind significance. Exposure dependence was obtained from nested prefixes of locked simulation runs. Points at different exposure levels are therefore not independent repetitions. The dashed line in panel (b) indicates the global 5σ detection threshold. First stable tested crossing is marked directly on each successful curve: 0.5, 1, 2, and 4 million histories for delamination, surface thinning, density deficit, and aluminium inclusion, respectively; bond-line gap did not reach the threshold.

3.3 Analytical and model robustness

Figure 6 compares the uncollided-primary transmission obtained with Geant4 against the corresponding Beer–Lambert predictions. Across the seven geometries and three electromagnetic constructors, all 21 calculations followed the analytical attenuation trend. The maximum absolute transmission residual was 0.360 percentage points. The largest residual occurred for the 2.0-mm Al-inclusion configuration and corresponded to approximately 2.66 Monte Carlo standard uncertainties. This level of agreement supports the consistency of the modelled material path lengths, areal densities, and uncollided-primary scoring.




Figure 6. Comparison of Geant4 uncollided-primary transmission with the Beer-Lambert prediction for seven modelled configurations and three electromagnetic physics constructors. Upper panel presents the identity comparison, whereas the lower panel shows the signed Geant4-minus-analytical residual in percentage points. Monte Carlo error bars represent binomial standard uncertainties from runs containing 4×10^6 incident histories each.

Figure 7 presents the detected-history ROI contrasts obtained with the three electromagnetic constructors. The sign of every defect contrast was preserved: delamination, surface thinning, density deficit, and the bond-line gap produced positive contrasts relative to their corresponding references, whereas the Al inclusion produced a negative contrast. The broad response hierarchy was also stable, with the largest absolute contrasts associated with delamination and surface thinning. The maximum model-to-model spread was 0.242 percentage points, occurring for the delamination case.

The difference between the Monte Carlo detected-history contrasts and the Beer-Lambert curve is physically expected because the analytical prediction represents only uncollided attenuation, whereas the detector-plane tally also accepts transmitted histories according to the established scoring definition. The close responses of delamination and surface thinning should therefore be interpreted only within the present matched-footprint geometry and similar projected areal-density reduction; they do not establish that these two defect classes are generally indistinguishable.




Figure 7. Detected-history ROI contrasts for five defect classes obtained with Livermore, Penelope, and Standard EM Option 4 constructors, together with corresponding uncollided Beer–Lambert predictions. Monte Carlo error bars show propagated binomial standard uncertainties. Horizontal marker offsets are used only for visual separation, and dashed analytical curve is a guide to the eye. Exact agreement with the analytical curve is not expected because the two quantities employ different transmitted-history definitions.

4 Discussion

The common-condition ranking was approximately delamination > thinning > density deficit > aluminium inclusion > bond-line gap for the selected defects. This is a detectability ranking for the stated geometry, not a universal ranking of NDT difficulty. The contrast primarily followed the change in projected attenuation. The delamination and surface thinning both removed 2 mm of CFRP along the beam path and produced similar positive contrasts. The similarity is itself an important limitation of a normal, single projection: when lateral footprints and lost areal densities are matched, the image contains little depth information with which to distinguish an internal planar air replacement from exit-surface material loss. This near-equivalence is specific to the matched projected material loss examined here, rather than a general equivalence of the two defect classes. The projections acquired at additional angles could introduce parallax and angular changes in the apparent footprint, providing constraints on the defect depth and orientation. Limited-angle tomography or full CT could reconstruct the through-thickness attenuation distribution more directly and thereby help separate an internal planar delamination from external surface thinning. The present method should therefore be interpreted as projected-contrast detection, not depth-resolved defect classification.

The aluminium case illustrates why a two-sided analysis is required. A one-sided search designed only for transmission increases would miss dense foreign material, even though such an inclusion can be structurally relevant. Although aluminium has a higher effective atomic number and attenuation than the carbon/epoxy matrix at 300 keV, it is not conventionally classified with high-Z radiographic materials such as tungsten or lead. Its relatively small contrast in this benchmark also shows that higher attenuation does not guarantee effortless detection when only a small CFRP sub-volume is replaced and a stringent image-wide trials penalty is applied. The result passed by only 0.12σ at the highest tested exposure, hence it should be described as detected under the simulated condition but close to the decision boundary.

The bond-line result is the clearest example of the distinction between signal presence and blind detectability. The positive ROI contrast agrees with the replacement of 1 mm of epoxy by air, and the local score approached 5σ . Nevertheless, searching through 141,094 possible positions and scales reduced the global significance to 1.51σ . The correct conclusion is therefore not that the geometry has no radiographic effect, but that the effect is insufficient for reliable blind detection within the four-million-history budget and the conservative decision rule.

Several limitations define the scope of the benchmark. The CFRP was treated as a homogeneous mixture, thus the fibre orientation, ply interfaces, tow-scale structure, resin-rich layers, surface curvature, and real roughness were not modelled. The source was perfectly parallel and monoenergetic; the detector had perfect geometric counting without an energy threshold, efficiency loss, a point-spread response, dead pixels, or electronic noise. Scatter rejection, collimation, the finite focal spot, source spectrum, environmental background, and mechanical misalignment were also absent. In addition, the beam energy, panel thickness, and detector sampling were not varied; the selected values define a controlled reference configuration rather than an optimized inspection system. Most importantly, no physical CFRP coupon or calibrated imaging system was used for validation. Thus, numerical significances cannot be translated directly into inspection time, dose, hardware specification, or certified PoD.

These limitations do not invalidate the controlled internal comparisons: every defect was transported with the same code path, exposure definition, and analysis, while the bonded and homogeneous panel families retained their own matched references. The agreement of the uncollided-primary transmission with the Beer–Lambert calculation, with a maximum residual of 0.360 percentage points, supports the consistency of the material definitions, photon path lengths, and primary-transmission tally. The preservation of the contrast signs and broad response hierarchy across three electromagnetic constructors, with a maximum inter-model spread of 0.242 percentage points, also indicates that the main conclusions are not an artefact of a single electromagnetic model.

These analytical and numerical checks are nevertheless not experimental validation. They do not calibrate the source output, detector response, inspection time, absorbed dose, or probability of detection. A future validation programme should therefore measure CFRP coupons with known inserts and machined thinning, calibrate the source spectrum and detector response, as well as repeat the blind decision protocol across independent exposures. Such measurements would support a genuine probability-of-detection analysis without changing the pre-declared search rule.

5 Conclusions

A Geant4/ROOT gamma-radiography benchmark was constructed for five representative CFRP defect classes using matched references, locked final seeds, and an explicit separation between known-ROI sensitivity and blind image-wide detection. At 300 keV and four million histories, 2 mm delamination, 2 mm surface thinning, a 10% density deficit, and a 2 mm aluminium inclusion passed the conservative blind global 5σ criterion. The aluminium case passed only at the highest tested exposure. A 1 mm air gap within a 1.2 mm epoxy bond line produced a correctly signed local contrast but did not meet the blind criterion.

The additional internal checks strengthened the physical basis of the benchmark. Across all the 21 model–geometry combinations, the uncollided-primary transmission differed from the Beer–Lambert prediction by no more than 0.360 percentage points. The maximum inter-model spread in the detected-history contrast was 0.242 percentage points, while the contrast signs and broad response hierarchy were preserved across the three electromagnetic constructors.

The results demonstrate that a physically consistent local signal is not equivalent to reliable blind detection. Under the present matched-footprint geometry, the internal material loss and surface thinning produced similar single-projection contrasts because they generated comparable projected areal-density reductions. This is a limitation specific to the present configuration rather than a general claim that the two defect classes are indistinguishable. The framework is suitable for controlled feasibility and exposure-scaling studies, but experimental detector-response and CFRP-coupon validation remain necessary before the numerical values can be used for aerospace inspection qualification.

Declarations

Author contributions. The sole author performed the conceptualization, methodology, software development, validation, formal analysis, visualization, interpretation, and manuscript preparation.

Data and code availability. The Geant4/ROOT source code, locked run definitions, analysis macros, and derived tables required to reproduce the reported calculations are available from the corresponding author on reasonable request.

Declaration of generative-AI-assisted work. OpenAI ChatGPT was used for assistance with code review, language editing, and manuscript organization. The author checked the simulation definitions, numerical outputs, references, interpretation, as well as the final text and takes full responsibility for the work. No AI system is listed as an author.

References

- [1] Wang, B., Zhong, S., Lee, T.-L., Fancey, K. S., and Mi, J., “Non-destructive testing and evaluation of composite materials/structures: A state-of-the-art review,” *Adv. Mech. Eng.*, 12 (4), 1–28, 2020. <https://doi.org/10.1177/1687814020913761>.
- [2] Chen, J., Yu, Z., and Jin, H., “Nondestructive testing and evaluation techniques of defects in fiber-reinforced polymer composites: A review,” *Front. Mater.*, 9 986645, 2022. <https://doi.org/10.3389/fmats.2022.986645>.
- [3] Gupta, R., Mitchell, D., Blanche, J., Harper, S., Tang, W., Pancholi, K., Baines, L., Bucknall, D. G., and Flynn, D., “A review of sensing technologies for non-destructive evaluation of structural composite materials,” *J. Compos. Sci.*, 5 319, 2021. <https://doi.org/10.3390/jcs5120319>.
- [4] More, N., Basse-Cathalinat, B., Baquey, C., Lacroix, F., and Ducassou, D., “Application of novel techniques of medical imaging to the non-destructive analysis of carbon-carbon composite materials,” *Nucl. Instrum. Methods Phys. Res.*, 214 (2–3), 531–536, 1983. [https://doi.org/10.1016/0167-5087\(83\)90628-2](https://doi.org/10.1016/0167-5087(83)90628-2).
- [5] Liu, X. and Chen, F., “Defects characterization in CFRP using x-ray computed tomography,” *Polym. Polym. Compos.*, 24 (2), 149–154, 2016. <https://doi.org/10.1177/096739111602400210>.
- [6] Petrò, S., Reina, C., and Moroni, G., “X-ray CT-based defect evaluation of continuous CFRP additive manufacturing,” *J. Nondestruct. Eval.*, 40 7, 2021. <https://doi.org/10.1007/s10921-020-00737-7>.
- [7] Kosmachev, P. V., Stepanov, D. Yu., Tyazhev, A. V., Vinnik, A. E., Eremin, A. V., Tolbanov, O. P., and Panin, S. V., “Radiographic inspection of carbon fiber-reinforced polymer composites (laminates) with epoxy and PEEK binders after impact and subsequent compression loading,” *Polymers*, 16 3262, 2024. <https://doi.org/10.3390/polym16233262>.
- [8] Agostinelli, S. and others, “Geant4—a simulation toolkit,” *Nucl. Instrum. Methods Phys. Res. A*, 506 250–303, 2003. [https://doi.org/10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8).
- [9] Allison, J. and others, “Geant4 developments and applications,” *IEEE Trans. Nucl. Sci.*, 53 270–278, 2006. <https://doi.org/10.1109/TNS.2006.869826>.
- [10] Allison, J. and others, “Recent developments in Geant4,” *Nucl. Instrum. Methods Phys. Res. A*, 835 186–225, 2016. <https://doi.org/10.1016/j.nima.2016.06.125>.
- [11] Li, O., Gao, J., and Zhang, Q., “A Geant4-based monte carlo x-ray imaging simulation platform,” *Appl. Radiat. Isot.*, 210 111376, 2024. <https://doi.org/10.1016/j.apradiso.2024.111376>.
- [12] Brun, R. and Rademakers, F., “ROOT—an object oriented data analysis framework,” *Nucl. Instrum. Methods Phys. Res. A*, 389 81–86, 1997. [https://doi.org/10.1016/S0168-9002\(97\)00048-X](https://doi.org/10.1016/S0168-9002(97)00048-X).