Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Spatial Distribution of WO_3 Nanoparticles Governs the Gamma-Ray Shielding of PMMA Nanocomposites: Discrete-Particle Monte Carlo Modeling and Experimental Validation

Mohammad Mahdi Norouzi
Amirkabir University of Technology
Masoumeh Mohamadian
mohamadian@aut.ac.ir

Amirkabir University of Technology

Article

Keywords: Geant4, Monte Carlo simulation, radiation shielding, polymer nanocomposite, spatial distribution, linear attenuation coefficient, Bremsstrahlung, gamma-ray spectroscopy

Posted Date: November 16th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11393872/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Spatial Distribution of WO_3 Nanoparticles Governs the Gamma-Ray Shielding of PMMA Nanocomposites: Discrete-Particle Monte Carlo Modeling and Experimental Validation

Mohammad Mahdi Norouzi[1], Masoumeh Mohamadian[1*]
[1] Amirkabir University of Technology, Tehran, Iran
[*] mohamadian@aut.ac.ir

Abstract

Polymer nanocomposites with high-Z nanoparticles (e.g., WO_3) are leading lead-free candidates for gamma-ray shielding, yet their predictive design rests on the homogeneous-mixture assumption, which ignores the microstructural reality of nanoparticle dispersion. Here, chemical composition is decoupled from spatial geometry to quantify how the three-dimensional arrangement of WO_3 particles in a PMMA matrix dictates macroscopic attenuation. Ten spatial-distribution algorithms—idealized baselines, fabrication defects, bio-inspired phyllotaxis, and engineered multi-layer, phase-separated, and flake architectures—were implemented as discrete-particle Geant4 geometries at a fixed 10 vol% loading, with a saturating power-law extrapolation ($\mu(d) = \mu\infty - C \cdot d^q$, $R^2 = 0.996$) bridging the computational scale gap to true nanoscale fillers and predicting a ≈19% nanoscale enhancement, consistent with the 15–30% reported experimentally. At fixed composition, the effective linear attenuation coefficient varied by more than a factor of two at 1 MeV (0.033–0.074 cm^{-1}) solely by rearranging particles, widening to ×6.5 at 250 keV; fabrication defects obey quantitative laws (sedimentation penalty ≈4% of the uniform-random-distribution baseline per unit gradient exponent; dispersion quality ≈2.3× more potent than cluster fragmentation), and multi-layer architectures trade a small primary-attenuation deficit for secondary-Bremsstrahlung suppression. A hot-cell case study resolves two shielding regimes (room-scatter-limited at 140 keV and window-limited at 364–511 keV, where positron-emitter workloads exceed the occupational limit ×3.1), delimiting the applicability of single-material nanocomposite windows. Narrow-beam gamma spectroscopy (^{152}Eu, ^{137}Cs) validates the transport model within 6% in the photoelectric regime (122–245 keV). Spatial architecture, not composition alone, is thus a first-order design variable for next-generation radiation shields.

Keywords: Geant4; Monte Carlo simulation; radiation shielding; polymer nanocomposite; spatial distribution; linear attenuation coefficient; Bremsstrahlung; gamma-ray spectroscopy

1. Introduction

The expanding use of ionizing radiation in diagnostic and therapeutic nuclear medicine has intensified the demand for shielding materials that satisfy the ALARA (As Low As Reasonably Achievable) principle and the occupational dose limits recommended by the International Commission on Radiological Protection (20 mSv per year, averaged over five years) [1]. For more than a century, lead has served as the industry standard owing to its high atomic number (Z = 82) and density; yet its severe neurotoxicity, the generation of inhalable lead dust in occupational settings, its environmental burden, and its excessive weight, which is a documented cause of chronic musculoskeletal disorders among clinical personnel, have motivated a decisive search for alternatives [2-4].

Polymer-matrix composites (PMCs) embedding high-Z, low-toxicity fillers (principally tungsten (W) and bismuth (Bi) oxides) have emerged as the most promising replacement platform [2-7]. Within this family, polymethyl methacrylate (PMMA) is an attractive matrix: it is optically transparent, mechanically compliant, chemically inert, and already qualified for medical environments, while tungsten trioxide (WO_3) offers the electron density required for photon attenuation (W: Z = 74, K-edge at 69.5 keV) without the toxicological hazards of lead. Numerical and

experimental evaluations indicate that tungsten-based composites can achieve half-value layers (HVL) comparable to lead, particularly in the low-energy regime where the photoelectric effect dominates [5-7].

A further performance increment has been reported upon transitioning from micro- to nano-fillers, with mass attenuation coefficients 15–30% higher than those of their micro-structural counterparts [8,9]. Two structural mechanisms are consistently invoked: (i) the elevated surface-area-to-volume ratio at the nanoscale, and (ii) the superior dispersion homogeneity of nanofillers, which suppresses "tunneling"—the passage of photons through matrix-only channels in poorly mixed or crystalline macro-structures—and thereby shortens the effective mean free path of the incident radiation.

Despite this extensive materials characterization, a critical bottleneck persists in how shielding efficacy is predicted. Standard theoretical tools (e.g., NIST-XCOM [11]) and conventional Monte Carlo material definitions (G4Nist in Geant4, or mixture cards in MCNP) represent nanocomposites as perfectly homogeneous atomic mixtures keyed solely to weight fraction, as indeed implemented in recent Monte Carlo studies of WO_3- and Bi_2O_3-based polymer composites [7,12,13]. This assumption is computationally efficient and adequate for dilute systems, but it is blind to microstructure. In practical fabrication routes (solution casting, in-situ polymerization), absolute homogeneity is unattainable: van der Waals-driven agglomeration [10], gravitational sedimentation of WO_3 ($\rho = 7.16$ g/cm³) within the far lighter PMMA matrix ($\rho = 1.18$ g/cm³) prior to curing, and phase separation all generate localized heterogeneous zones. Moreover, explicit discrete-particle modeling of a true nanocomposite is computationally prohibitive: a 1 cm³ sample at 10 vol% with 20 nm radius fillers requires on the order of 10^{15} individual geometric volumes, far beyond the memory and initialization limits of any standard workstation.

Consequently, three questions remain open: (1) by how much do realistic microstructural deviations alter the macroscopic linear attenuation coefficient (LAC, μ) at fixed composition; (2) can deliberately engineered non-uniform architectures (e.g., layered designs that manage secondary Bremsstrahlung, or bio-inspired quasi-random packings) outperform homogeneous mixing; and (3) how can simulation access the nanoscale regime despite the scale gap?

This work addresses all three. Our specific contributions are:

- Decoupling of composition from geometry: ten spatial-distribution algorithms (spanning idealized baselines, real fabrication defects, bio-inspired topologies, and engineered architectures) are modeled as discrete particle ensembles in Geant4 at a fixed loading of 10 vol% WO_3/PMMA.

- A scale-bridging methodology: a saturating power-law model ($\mu(d) = \mu\infty - C \cdot d^q$) that predicts nanoscale shielding behavior from tractable meso-scale simulations, circumventing the 10^{15}-volume bottleneck.

- Experimental validation with a self-auditing protocol: narrow-beam gamma spectroscopy of solution-cast composites, confirming the transport model within 6% in the photoelectric regime and establishing the as-fabricated loading by XCOM inversion.

- Engineering translation: a hot-cell case study resolving room-scatter-limited versus window-limited shielding regimes and quantifying the compliance failure of single-material windows at positron-emission energies.

Collectively, the results argue for a paradigm shift in shield design, from "adding more filler" to "architecting the spatial distribution" of the filler.

2. Materials and Methods

To systematically investigate the structural dependency of shielding performance (beyond the limitations of standard theoretical databases) a dual-phase methodology was developed, integrating computational physics modeling to isolate geometric variables with empirical fabrication and gamma spectroscopy to validate the computational

predictions. The base material system consists of 10 vol% WO_3 nanoparticles dispersed in a PMMA matrix. Within this fixed chemical composition, ten distinct spatial-distribution functions (ranging from idealized baselines to realistic manufacturing defects) were mathematically defined and implemented to quantify exactly how particle arrangement dictates macroscopic attenuation.

2.1 Material System and Design Window

The shielding architecture was conceptualized as a two-phase system. PMMA (effective atomic number $Z_{eff} \approx 6.5$, $\rho = 1.18$ g/cm³) was selected as the low-Z matrix for its optical transparency, structural flexibility, and established viability in medical environments (e.g., hot-cell viewing windows). WO_3 ($\rho = 7.16$ g/cm³) was selected as the high-Z phase, exploiting the high electron density of tungsten ($Z = 74$; K-edge at 69.5 keV) while eliminating the toxicological and environmental liabilities of lead.

At a volume fraction $\varphi = 0.10$, the corresponding weight fraction is ≈ 40 wt% and the rule-of-mixtures composite density is $\rho_c \approx 1.78$ g/cm³. This loading was chosen as representative of practical, processable composites: high enough to yield measurable attenuation contrast, yet below the packing and percolation limits at which geometric artifacts begin to dominate (Section 3.1). The simulation energy window (250 keV–1 MeV) spans the diagnostically and therapeutically relevant photon range, while the experimental window (121.8–661 keV) covers the ^{152}Eu and ^{137}Cs photo-peaks. Fixing the composition across all simulations isolates geometry as the sole independent variable.

2.2 Monte Carlo Computational Framework

2.2.1 Justification for Geant4 and Rejection of Voxel-Based Methods

Conventional transport codes such as MCNP define composite materials through cell-based or lattice-based (voxel) mixture definitions. While computationally efficient, this voxelization homogenizes the sub-voxel structure and destroys the discrete physical boundaries of individual nanoparticles, rendering "2D geometric cross-section" effects and photon tunneling inaccessible to study. We therefore adopted the Geant4 toolkit [14, 15]. Its object-oriented C++ architecture permits each nanoparticle to be instantiated as an independent solid volume (G4Sphere) algorithmically placed inside a continuous matrix volume (G4Box), with exact boundary tracking. This enables the explicit resolution of inter-particle voids, particle shadows, and secondary-electron transport across WO_3/PMMA interfaces, all first-order effects in the present study.

2.2.2 Physics List and Electromagnetic Processes

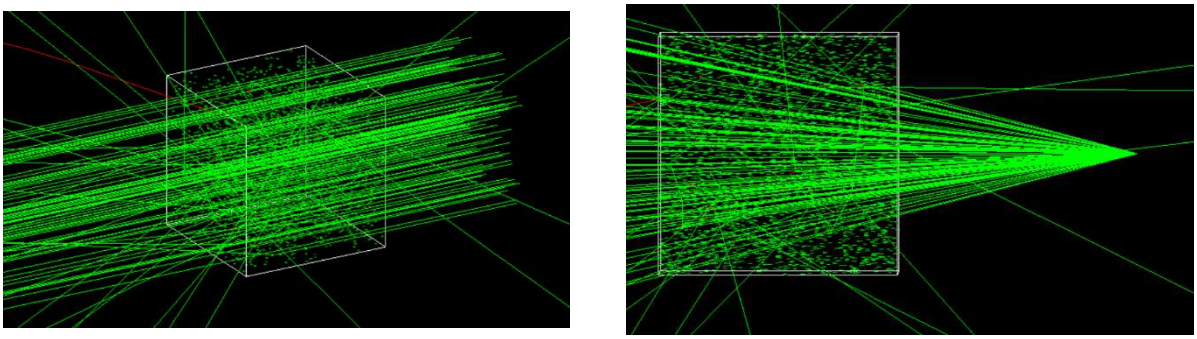
Photon and electron transport were handled with the G4EmLivermorePolarized electromagnetic constructor, which employs evaluated data libraries (EPDL/EEDL-based) and is optimized for low-energy electromagnetic processes. The modeled processes include shell-resolved photoelectric absorption with fluorescence (critical for accurately treating tungsten near its K-edge at 69.5 keV), Compton scattering with atomic binding corrections, Rayleigh scattering, pair production, and full secondary-electron transport with bremsstrahlung and ionization losses. Electron/photon production cuts were reduced below the characteristic particle dimensions so that secondary electrons generated inside WO_3 particles are correctly propagated across material interfaces, a strict requirement for the Bremsstrahlung-suppression analysis of Section 3.5.

2.2.3 Geometrical Configuration and Source Terms

Two distinct irradiation geometries were implemented (Fig. 1):

1. Narrow-beam planar source (Fig. 1a): a monoenergetic, collimated beam directed perpendicular to the shield face, sized to the shield cross-section, satisfying "good-geometry" irradiation conditions. The associated scoring convention and its build-up implications are detailed in Section 2.2.4.

2. Isotropic point source (Fig. 1b): used for the hot-cell dose assessment (Section 2.6), where extended-field irradiation makes the build-up factor significant. Because variance-reduction treatment of thick shields in Geant4 is impractical for this configuration, the dose behind the shield was analytically corrected using the Taylor approximation for the build-up factor [16].



(a) Narrow-beam geometry (b) Point-source geometry

Figure 1. Irradiation geometries implemented in Geant4. (a) Planar source covering the shield cross-section, used for the attenuation campaign (scoring convention in Section 2.2.4). (b) Point-source configuration for extended-field and dose calculations (Sections 3.6 and 3.8), with emission restricted to the solid angle subtended by the shield and build-up corrected analytically (Taylor approximation). Solid and dashed arrows denote primary and scattered photons. Not to scale.

For each configuration, the matrix thickness was varied parametrically, and monoenergetic simulations were performed across 250 keV–1 MeV.

2.2.4 Tallies, Statistical Treatment and Scoring Convention

In the structural campaign (Sections 3.1–3.5), the transmitted tally registered every photon exiting the rear face of the shield, irrespective of scattering history or energy. The extracted quantities are therefore effective, transmission-inferred linear attenuation coefficients (μ_{eff}), obtained from the ratio of transmitted to incident counts through identical scoring planes behind the shield:

$$\mu = -\frac{1}{x}\ln\left(\frac{I}{I_0}\right) \quad (1)$$

where x is the shield thickness (5.0 cm) and at least 10^7 primary histories were generated per configuration, keeping relative counting uncertainties below 0.2%. Because this convention is applied uniformly to all ten distributions, all loadings, and all energies, every structural comparison in this work is internally consistent. Absolute magnitudes, however, are not directly comparable with narrow-beam databases: for the 5-cm URD (Uniform Random Distribution) baseline, the ratio of NIST-XCOM narrow-beam mixture values to μ_{eff} corresponds to an effective build-up factor B ≈ 1.35–1.67 over 250–1000 keV, consistent with the shield's optical depth ($\mu x \approx 0.6$–2) [11]. Where narrow-beam references are shown for context (Figs. 3 and 7), this offset is annotated. The half-value layer (HVL = ln2/μ) and mean free path (λ = 1/μ) are derived from the same coefficients (Figs. S5 and S6).

2.2.5 The Computational Bottleneck and Scale-Bridging Extrapolation Model

The number of discrete volumes required to represent a particulate composite grows inversely with the cube of the particle radius:

$$N(r) = \frac{3\,\phi_f V_m}{4\pi r^3} \quad (2)$$

For a 1 cm^3 matrix at 10 vol%, reducing the radius from 100 μm to 1 μm already demands $N > 10^{10}$ discrete volumes, and a true nanoscale filler (r = 20 nm) would require $N \approx 10^{15}$, far beyond the memory and initialization limits of standard workstations.

To overcome this barrier, we simulated the filler-size dependence over the tractable meso-scale range (d = 79 μm–10 mm; fixed 10 vol%, E = 1 MeV, URD) and fitted it with a saturating power law:

$$\mu(d) = \mu\infty - C \cdot d^q \quad (3)$$

Here $\mu\infty$ is the nanoscale asymptote—the coefficient of the limitingly fine dispersion, and $C \cdot d^q$ quantifies the heterogeneity (channeling) penalty, which vanishes as the characteristic heterogeneity scale falls far below the photon mean free path and the discrete structure becomes optically indistinguishable from a homogeneous medium. The regression ($R^2 = 0.996$; $\mu\infty = 0.0545\ cm^{-1}$, $C = 1.65\times10^{-6}\ cm^{-1}\cdot\mu m^q$, $q \approx 0.93$) extrapolates to the true nanoscale ($d \rightarrow 20$ nm) without generating nanoscale geometry; the extrapolation is restricted to $d \geq 10$ nm, below which continuum material definitions cease to be meaningful. A log-linear form ($\mu = A \ln r + B$), by contrast, fails to capture the observed saturation ($R^2 = 0.73$) and would diverge unphysically in the nanoscale limit.

2.2.6 Overlap Control and Error Quantification

Particles were placed by random sequential addition (RSA) with rejection of overlapping placements. At high volume fractions, packing frustration causes rejection sampling to skew the achieved volume fraction, while permitting overlaps artificially inflates local density. The magnitude of this geometric artifact was quantified as a function of loading (Section 3.1); it remains negligible below ~20 vol% but contributes materially to model deviations above ~25 vol%, and is explicitly flagged wherever relevant.

2.3 Spatial Distribution Functions

Standard calculations (XCOM) and conventional Monte Carlo definitions treat the nanocomposite as a perfectly homogeneous mixture whose μ follows solely from the weight fraction. Physically, however, a photon beam experiences the shield as a stochastic sequence of encounters with discrete high-Z particles: the macroscopic μ is implicitly a function of the effective 2D geometric cross-section (σ_2D) presented to the beam. Ten distribution algorithms were implemented via C++ placement loops, grouped into four physical paradigms (Table 1); two-dimensional cross-sections of all ten realizations are shown in Fig. 2.

Table 1. The ten spatial-distribution algorithms investigated at fixed composition (PMMA/10 vol% WO_3).

	Distribution	Algorithm / parameters	Physical analogue
1	Simple-cubic (SC) lattice	Grid spacing $a = (V/N)^{1/3}$; beam || [100]	Crystalline regularity; channeling bound
2	URD	MT19937 sampling; RSA overlap rejection	Ideal homogeneous dispersion (baseline)
3	Sedimentation	$n(z) \propto z^{-s}$; s = 1-5 (reference s = 2)	Gravity-induced phase separation before curing
4	Centralized Gaussian	Single 3D normal distribution at matrix center	Poor mixing; central agglomeration
5	Clustered	N_c Gaussian foci (10 → 10^4; reference 150) at fixed loading; intra-cluster dispersion σ_c (reference 1/12)	Van der Waals agglomeration
6	Golden-ratio (phyllotaxis)	$\theta_i = 2\pi i/\varphi$; $\rho_i \propto \sqrt{i}$	Bio-inspired quasi-random packing
7	Multi-layer	Gaussian density × periodic step function along beam axis	Engineered impedance layering

8	Phase-separated layered ⊥	WO_3 confined to dense planes; beam ⊥ planes	Best-case orientation
9	Phase-separated layered ||	Same structure; beam || planes	Worst-case orientation
10	Flake-morphology URD	Platelets, broad faces ⊥ beam	High-σ_2D filler morphology

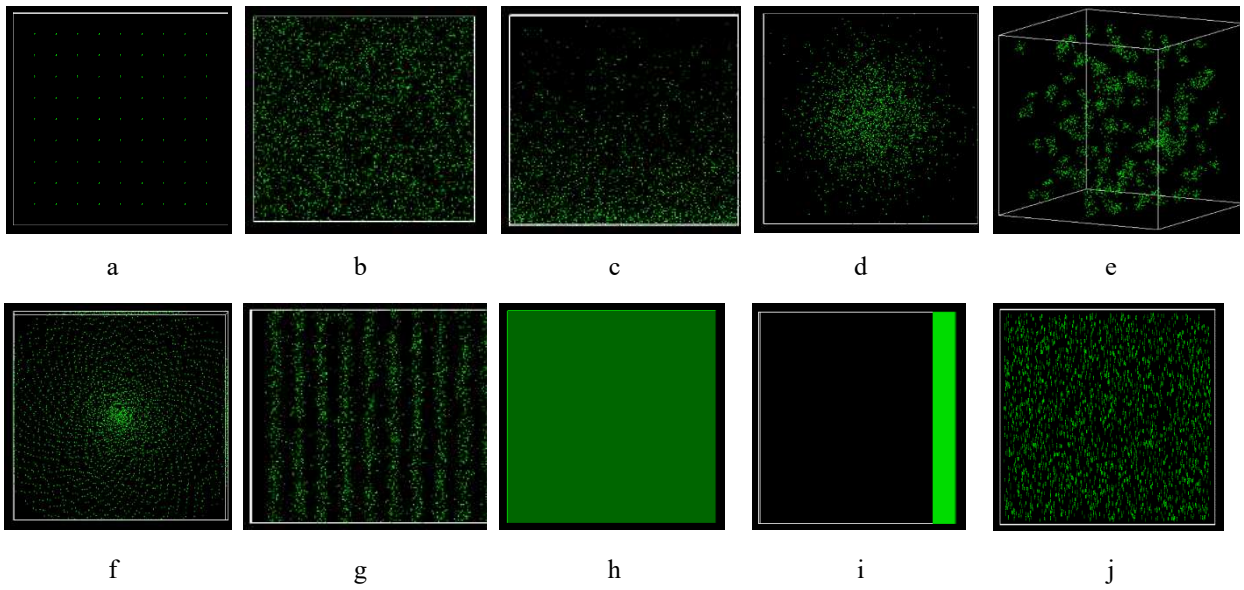


Figure 2. Cross-sectional profiles of the ten WO_3 spatial-distribution algorithms at fixed composition (PMMA/10 vol% WO_3). (a–e) Idealized and defective structures: simple-cubic lattice, URD, sedimentation (s = 2), centralized Gaussian, clustered (N_c = 150, σ_c = 1/12). (f–j) Engineered architectures: golden-ratio (phyllotaxis), multi-layer (side view), phase-separated layered with the beam perpendicular and parallel to the layers, flake-morphology URD. Arrows indicate the incident beam direction. Three-dimensional renderings are provided in Fig. S1.

2.3.1 Idealized and Baseline Distributions

URD serves as the primary control: particles are positioned with the Mersenne Twister pseudo-random generator, assuming ideal fabrication with no agglomeration, and providing an isotropic (3D) σ_2D. Regular Cubic Lattice establishes the lower structural bound: particles are arranged on a simple-cubic grid. Although volumetrically homogeneous over large differential volumes, this arrangement creates directional channels; a beam parallel to the grid axes can traverse the matrix via tunneling, drastically reducing μ at constant volume fraction.

2.3.2 Real-World Fabrication Defect Models

Sedimentation models gravity-induced phase separation during solution casting, implemented as a density gradient along the vertical axis:

$$n(z) \propto z^{-s} \quad (4)$$

with sedimentation coefficient s (s = 2 as the reference case, varied over s=1-5). The result is a composite with a WO_3-rich base and a depleted, low-attenuation upper region (a realization at the reference case s = 2 is shown in Fig. S3). Centralized Gaussian placement concentrates the entire filler inventory at the matrix center, mimicking severe mixing

failure. Clustered distributions simulate agglomeration: N_c focal points are defined, each seeding a local 3D Gaussian cloud of particles (reference dispersion $\sigma_c = 1/12 \approx 0.083$), creating localized high-density zones separated by macroscopic matrix voids. The cluster population (N_c=10 → 10^4) and the dispersion coefficient (σ_c = 1/48–1/2) were varied independently (Section 3.4).

2.3.3 Bio-Inspired Quasi-Random Topologies

Golden-Ratio (Phyllotaxis) Distribution explores directional optimization. Particle positions follow the phyllotaxis (Vogel) spiral [17]:

$$\theta_i = \frac{2\pi i}{\phi}, \quad \rho_i = \rho_{max}\sqrt{\frac{i}{N}} \qquad 5$$

where $\varphi \approx 1.618$ is the Golden Ratio (golden angle ≈ 137.5°). Because φ is irrational, the angular sequence never repeats, preventing periodic void alignment; the $\sqrt{i}$ radial law optimally minimizes void fractions in the projected 2D cross-section along the beam axis, maximizing first-interaction probability without increasing the filler loading. The model inherently sacrifices 3D homogeneity, producing cylindrical voids at the structural periphery.

2.3.4 Engineered Architectures

Multi-layer distributions (Section 2.4) alternate WO_3-rich planes with pure-PMMA spacers along the beam axis (their orientation dependence is illustrated in Fig. S2). Phase-separated layered composites represent the fully heterogeneous limit (dense particle planes versus continuous matrix corridors) evaluated at two orthogonal beam orientations (⊥ and ∥) to bound the anisotropy. Flake-morphology URD replaces spheres with platelets whose broad faces are oriented perpendicular to the beam, maximizing the projected cross-section at identical volume fraction.

2.4 Physics-Driven Multi-Layer Design: Secondary-Electron Management

In conventional homogeneous nanocomposites, a secondary hazard emerges in the diagnostic energy range (≈50–300 keV): secondary electrons ejected by photoelectric or Compton events in WO_3 decelerate in the vicinity of other high-Z atoms, converting a significant fraction of their kinetic energy into Bremsstrahlung. Since the radiative stopping power scales approximately with Z^2, a poorly structured composite can partially behave as a secondary X-ray source, degrading its own net shielding efficiency.

The multi-layer paradigm addresses this by spatially separating the generation and thermalization of secondary electrons. Alternating WO_3-rich planes (primary photon attenuation) with pure-PMMA spacers forces ejected electrons to cross the interface and decelerate in the low-Z matrix, where losses occur predominantly through ionization and excitation rather than radiation, trapping the energy as heat. In Geant4, this architecture was implemented as a spatially modulated probability density rather than as solid macroscopic slabs (which would suppress nano-scale interface effects): the WO_3 placement distribution was multiplied by a periodic step function along the beam axis,

$$p(x) = G(x) \cdot \Pi(x) \backslash tag6 \qquad 6$$

creating well-defined particle-rich layers separated by zones of zero placement probability, with spacer thicknesses commensurate with the ranges of the secondary electrons generated at the energies of interest.

2.5 Experimental Synthesis and Gamma Spectroscopy

2.5.1 Composite Fabrication and Microstructural State

We fabricated the PMMA/10 vol% WO_3 nanocomposite by solution casting. PMMA powder was first dissolved in acetone; WO_3 nanoparticles were then added gradually while the suspension underwent continuous magnetic stirring and ultrasonication to break down agglomerates. After initial solvent evaporation at elevated temperature, the dried composite was repeatedly re-dissolved in acetone and re-cast through several melt/re-dissolution cycles, iterative processing that was essential to fragment micro-scale agglomerates and approach a uniform dispersion. A pure PMMA reference sample was fabricated with an identical protocol for baseline measurements (the complete fabrication sequence is documented in Fig. S8).

Owing to the absence of an industrial vacuum-curing furnace, the final samples exhibited minor optical scattering (incomplete transparency); the dimensional and structural integrity required for accurate attenuation measurements was fully preserved. Based on the physical behavior of the viscous suspension during casting and evaporation, the as-fabricated microstructure is expected to align more closely with the mildly sedimented or mildly clustered simulation models than with the ideal URD, an expectation tested quantitatively in Section 3.7.

The effective filler loading of the as-fabricated samples was established post hoc by XCOM inversion of the measured attenuation coefficients (Section 3.7) at ≈15 wt% WO_3 (≈4 vol%). All validation comparisons in Section 3.7 are therefore performed like-for-like at this loading, while the structural campaign (Sections 3.1–3.5) is conducted at the design loading of 10 vol.%

2.5.2 Narrow-Beam Gamma Spectroscopy

We measured experimental LAC values with a narrow-beam geometry (Fig. S9) to suppress scattered-photon contributions. The detection chain comprised a NaI(Tl) scintillator coupled to a photomultiplier tube, with signals processed through a pre-amplifier and linear amplifier and recorded by a multi-channel analyzer (MCA). Lead collimators defined the beam aperture between source, shield, and detector.

Two calibrated sources provided energy-dependent data: ^{152}Eu (photo-peaks at 121.78, 244.69, 344.27, and 443.96 keV) and ^{137}Cs (661 keV). Rather than raw count rates, the net area under Gaussian-fitted photo-peaks was integrated with and without the shield under identical geometrical conditions and acquisition times. The attenuation factor then yields μ via Eq. (1), with the sample thickness taken as the mean of multiple caliper measurements. Background subtraction and repeated acquisitions were applied to bound counting statistics; dead-time losses in the high-rate unshielded acquisitions were identified post hoc as the dominant residual artifact and are treated explicitly in Section 3.7.

2.6 Hot-Cell Case Study: From Micro-Defects to Macro-Safety

To translate the spatial framework into engineering practice, we assessed a nuclear-medicine dispensing hot-cell room model [18], evaluated at the three workhorse photon energies of nuclear medicine: 140 keV (^{99m}Tc), 364 keV (^{131}I), and 511 keV (positron emitters, with dedicated shielding requirements per AAPM TG-108 [19]). Operator dose rates per unit activity were scored with the isotropic point-source configuration (Section 2.2.3), build-up-corrected via the Taylor approximation, both without shielding and behind a transparent PMMA/WO_3 window in the cell wall. Annual occupational doses were derived for a representative workload of $W \cdot T = 3.7 \times 10^6$ GBq·min per year (≈1.7×10^3 Ci·h, e.g., a 1 Ci-class ^{99}Mo/^{99m}Tc generator over a full working year) and assessed against the ICRP occupational limit of 20 mSv per year [1]. The URD model provided the ideal-manufacturing scenario and the sedimentation model ($s = 2$) the defective-manufacturing scenario; their difference quantifies the safety margin consumed by microstructural

neglect. Composite attenuation coefficients at the three energies were taken from the campaign spectra, and the implied window thickness was back-calculated from the modeled attenuation factors as an internal consistency check.

3 Results and Discussion

3.1 Benchmarking the Discrete Engine: Loading Dependence and Limits of Homogeneous Definitions

To establish the reliability of the discrete-particle geometry engine, the URD model and a hand-built homogeneous material definition were evaluated as a function of WO_3 volume fraction at 1 MeV (Fig. 3), alongside NIST-XCOM narrow-beam mixture values.

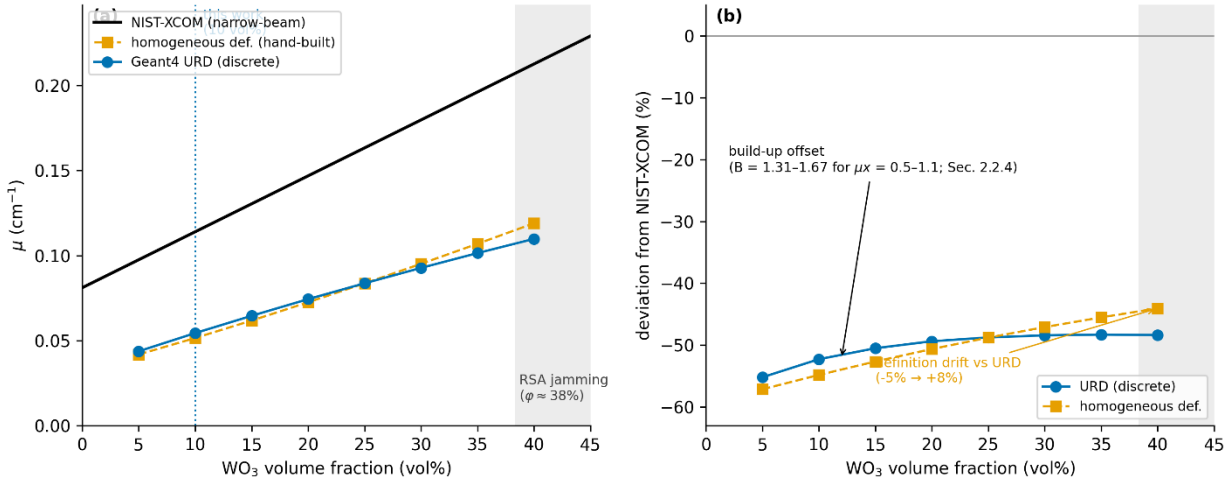


Figure 3. Loading dependence of the effective linear attenuation coefficient at 1 MeV. (a) Discrete URD Geant4 model (blue circles), hand-built homogeneous material definition (orange squares), and NIST-XCOM narrow-beam mixture values (black line); the dotted vertical line marks the working loading (10 vol%), and the shaded region beyond φ ≈ 38 vol% marks the random-sequential-addition jamming limit. (b) Relative deviation of the two simulated columns from the NIST-XCOM reference: the systematic offset (−48 to −55%) is the build-up signature of the transmission-inclusive scoring convention (implied B = 1.35–1.67 over μx = 0.5–1.1; Section 2.2.4), while the progressive crossing of the two curves (homogeneous definition: −5% to +8% relative to the URD) reveals a material-definition inconsistency.

At the working loading (10 vol%), the discrete model yields an effective coefficient of 0.0544 cm^{-1} against a narrow-beam XCOM value of 0.114 cm^{-1}, an offset corresponding to a build-up factor B = 1.35 at a narrow-beam optical depth $\mu_x = 0.57$. This offset is not constant but scales with optical depth as expected: the implied build-up factor $B = \exp[(\mu_{nb} - \mu_{eff})x]$ rises monotonically from 1.35 at 10 vol% to 1.67 at 40 vol% (Table S6; $\mu_x = 0.57 \rightarrow 1.06$), with the effective coefficients running −55% (5 vol%) to −48% (40 vol%) below the narrow-beam reference, precisely the Taylor-type behavior of the transmission-inclusive scoring convention (Section 2.2.4), and an independent confirmation that the discrete engine is internally consistent with database physics under the campaign convention.

Two deviations require explicit qualification. First, random sequential addition jams at φ ≈ 0.383 for monodisperse spheres: the two highest loadings in Fig. 3 necessarily contain overlapping (locally densified) configurations or reduced effective loadings, consistent with the overlap-error analysis of Fig. S7, and discrete results beyond ~25 vol% should be interpreted with this geometric caveat. Second, the hand-built homogeneous definition drifts monotonically from −5% at 10 vol% to +8% at 40 vol% relative to the discrete URD, which shares its scoring convention and must coincide with it, a smooth 13-point swing diagnostic of a density or fraction-basis inconsistency; this column was therefore excluded from quantitative use. The methodological lesson (homogeneous definitions must be audited against database references before use) is itself a finding of practical importance for simulation studies of nanocomposites [7].

3.2 Particle Size Effect and the Scale-Bridging Extrapolation

At fixed 10 vol% loading (E = 1 MeV, URD), decreasing the filler size monotonically increased the effective LAC over the simulated meso-scale range (d = 79 μm–10 mm), but the dependence saturates rather than continuing log-linearly (Fig. 4a). The data are captured by Eq. (3) with $R^2 = 0.996$ ($\mu_\infty = 0.0545$ cm^{-1}, $C = 1.65\times10^{-6}$ cm$^{-1}\cdot\mu m^q$, $q \approx 0.93$); a log-linear fit reproduces only 73% of the variance and would diverge unphysically in the nanoscale limit. (One tabulated value at d = 79.4 μm is inconsistent with its own transmitted counts at the 21σ level; the count-derived value was adopted; Table S5.)

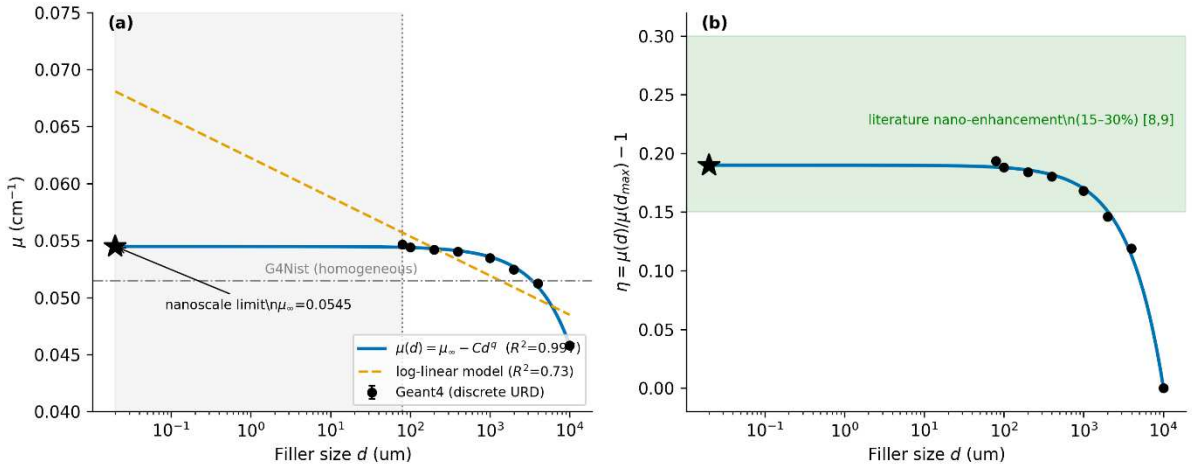


Figure 4. Filler-size dependence of the effective LAC (URD, PMMA/10 vol% WO_3, E = 1 MeV) and scale-bridging extrapolation. (a) Discrete-particle results over the meso-scale range (circles, ±1σ), the saturating power-law fit of Eq. (3) (shaded band: 95% CI), its extrapolation to the nanoscale (star; $\mu_\infty = 0.0545$ cm^{-1}), and the homogeneous-definition reference (dotted). (b) Size-induced enhancement $\eta = \mu(d)/\mu(d_{max}) - 1$; the green band marks the 15–30% nano-enhancement range reported in the literature [8, 9].

Physically, $C\cdot d^q$ quantifies the heterogeneity penalty: once the characteristic heterogeneity scale falls far below the photon mean free path, the discrete structure becomes optically indistinguishable from a homogeneous medium and μ converges to the nanoscale asymptote μ_∞. Relative to the largest simulated filler (10 mm), the asymptote represents a +19% enhancement (within the 15–30% band reported for nano- versus macro-particulate composites [8, 9]) while the residual benefit over ~1-mm fillers is only ≈2% (Fig. 4b). The frequently reported "nano advantage" in primary attenuation is therefore largely a heterogeneity-scale effect, essentially exhausted once the filler size falls below ~100 μm at this energy and loading. The framework bridges the ~10^5-order-of-magnitude volume gap between tractable simulation and physical nanocomposites (≈10^{15} volumes), and μ_∞ provides an upper-bound design estimate per distribution.

3.3 Decoupling Composition from Geometry: The Dominance of σ_2D

The central result of this study is presented in Fig. 5 and Table 2: at a fixed composition of 10 vol% WO_3, the effective 2D geometric cross-section (σ_2D) presented to the incident beam, rather than the composition itself, governs the effective LAC, which spans ×2.2 at 1 MeV (0.033–0.074 cm^{-1}). The spread widens monotonically as energy decreases (to ×6.5 at 250 keV) tracking the growth of the WO_3/PMMA attenuation contrast (≈6 → ≈19), which amplifies both shadowing gains and channeling penalties; structural design is therefore most consequential precisely in the diagnostic energy window (the complete energy-resolved dataset is tabulated in Table S1).

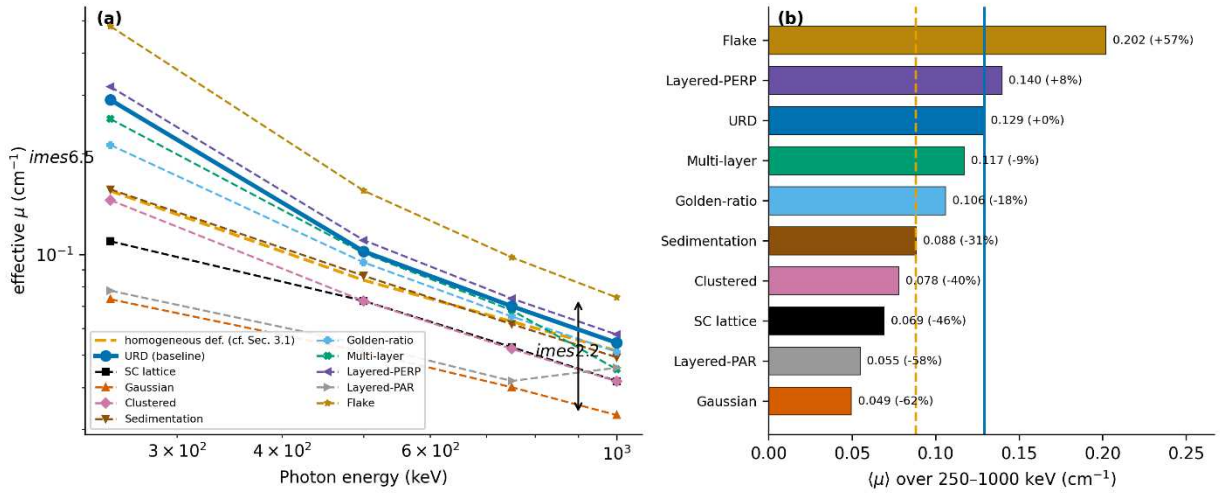


Figure 5. Structural control of the effective LAC at fixed composition (PMMA/10 vol% WO_3). (a) Energy-resolved effective LAC for the ten spatial distributions (log–log axes); the orange dashed curve is the hand-built homogeneous definition (Section 3.1). Double-headed arrows mark the max/min spread (×2.2 at 1 MeV, ×6.5 at 250 keV). (b) Energy-averaged effective LAC over 250–1000 keV, ranked in descending order; the blue line marks the URD baseline and labels give deviations from it. All values follow the campaign scoring convention (Section 2.2.4).

Table 2. Simulated effective LAC at 1 MeV for PMMA/10 vol% WO_3 as a function of spatial distribution.

Distribution	**μ (cm⁻¹)**	**Δ vs. URD**	**Physical regime**
Flake-morphology URD	0.074	+37%	Upper bound (max. projected area)
Phase-separated layered ⊥	0.058	+6%	Guaranteed traversal of dense planes
URD	0.054	—	Isotropic benchmark
Homogeneous material definition	0.0515	−5%	Mixture-rule reference
Golden-ratio (phyllotaxis)	0.051	−6%	Optimized 2D projection
Sedimentation (s = 2)	0.049	−9%	Systematic under-protected zone
Phase-separated layered ∥	0.046	−16%	Continuous matrix corridors
Multi-layer	0.045	−17%	Bremsstrahlung-management architecture
Clustered (N_c = 150, σ_c = 1/12)	0.042	−23%	Distributed agglomerates
Simple-cubic lattice	0.042	−23%	Aligned tunneling channel
Centralized Gaussian	0.033	−39%	Lower bound (central void)

The upper bound. Flake morphology maximizes the projected area per unit filler mass (μ = 0.074 cm^{-1}, +37% over URD). The perpendicular layered configuration (0.058 cm^{-1}) benefits from staggered particle shadows across consecutive dense planes: a photon that evades one plane is intercepted by the next, eliminating the through-going matrix channels that arise by chance in 3D random dispersions.

The lower bound. The centralized Gaussian distribution (0.033 cm^{-1}, −39%) is the worst case: density peaks at the center and decays outward, opening a macroscopic annular void around the beam axis through which photons cross essentially pure PMMA. We note that the clustered distribution (0.042 cm^{-1}) still outperforms the single centralized Gaussian: many small voids are statistically less damaging than one large void, because distributed clusters maintain a nonzero interception probability everywhere.

Anisotropy and tunneling. The simple-cubic lattice (0.0418 cm^{-1}) performs as poorly as severely clustered material despite being mathematically homogeneous over large differential volumes: beam-parallel alignment of lattice

channels enables tunneling, proving that volumetric homogeneity does not guarantee shielding efficacy. Likewise, rotating the phase-separated layered structure by 90° collapses performance from 0.058 to 0.046 cm^{-1}, as continuous matrix corridors open parallel to the layers. This orientational dependence is consistent with the anisotropic attenuation recently reported for layered WO_3/PP and Bi_2O_3/PP composites [20].

Bio-inspired quasi-randomness. The golden-ratio distribution (0.051 cm^{-1}) essentially matches the homogeneous material definition (0.0515 cm^{-1}) despite its 3D cylindrical inhomogeneity. Its optimized 2D projection fully compensates for peripheral voids, demonstrating that engineering the geometric projection of particles can substitute for chemical homogeneity.

Omnidirectional reliability. Flake and perpendicular-layered architectures deliver their maxima only for a known beam direction. In realistic environments (scattered, multi-directional fields), the URD maintains a high, angle-independent μ and thus remains the most reliable structural paradigm; directional architectures are preferable only where the source geometry is fixed and known, an explicit design trade-off introduced by this framework.

Energy-resolved ratios sharpen these regimes (Fig. 5a): all channeling-prone structures fall monotonically relative to the URD as energy decreases (Gaussian/URD: 0.61 → 0.25 from 1 MeV to 250 keV; layered-∥: 0.84 → 0.27; simple-cubic: 0.77 → 0.38; sedimentation: 0.91 → 0.54), whereas the σ_2D-maximizing architectures improve (flake: 1.37 → 1.66; layered-⊥: 1.06 → 1.10), and the orientational anisotropy of the layered structure grows from ×1.3 to ×4.1. The structural penalty is thus governed by the photon mean free path relative to the characteristic void dimension, rather than being a fixed property of the geometry.

3.4 Parametric Analysis of Fabrication Defects

The severity of the two clustering defects was mapped parametrically at 1 MeV (Fig. 6). Fragmenting the agglomerate population at fixed dispersion (σ_c = 1/12) monotonically restores attenuation (from −37% relative to the URD at N_c = 10 to −0.7% at $N_c = 10^4$) with the ten-cluster limit approaching the centralized-Gaussian penalty (−39%) to within 2.4 percentage points (Fig. S4): a small number of clusters is statistically equivalent to one large centralized void, identifying the two defect families as endpoints of a single heterogeneity continuum. Widening the intra-cluster dispersion at fixed N_c = 150 acts analogously, from −39% at σ_c = 1/48 to −4% at σ_c = 1/2, with a mild turnover at σ_c= 1/2 consistent with boundary truncation of overlapping cluster clouds.

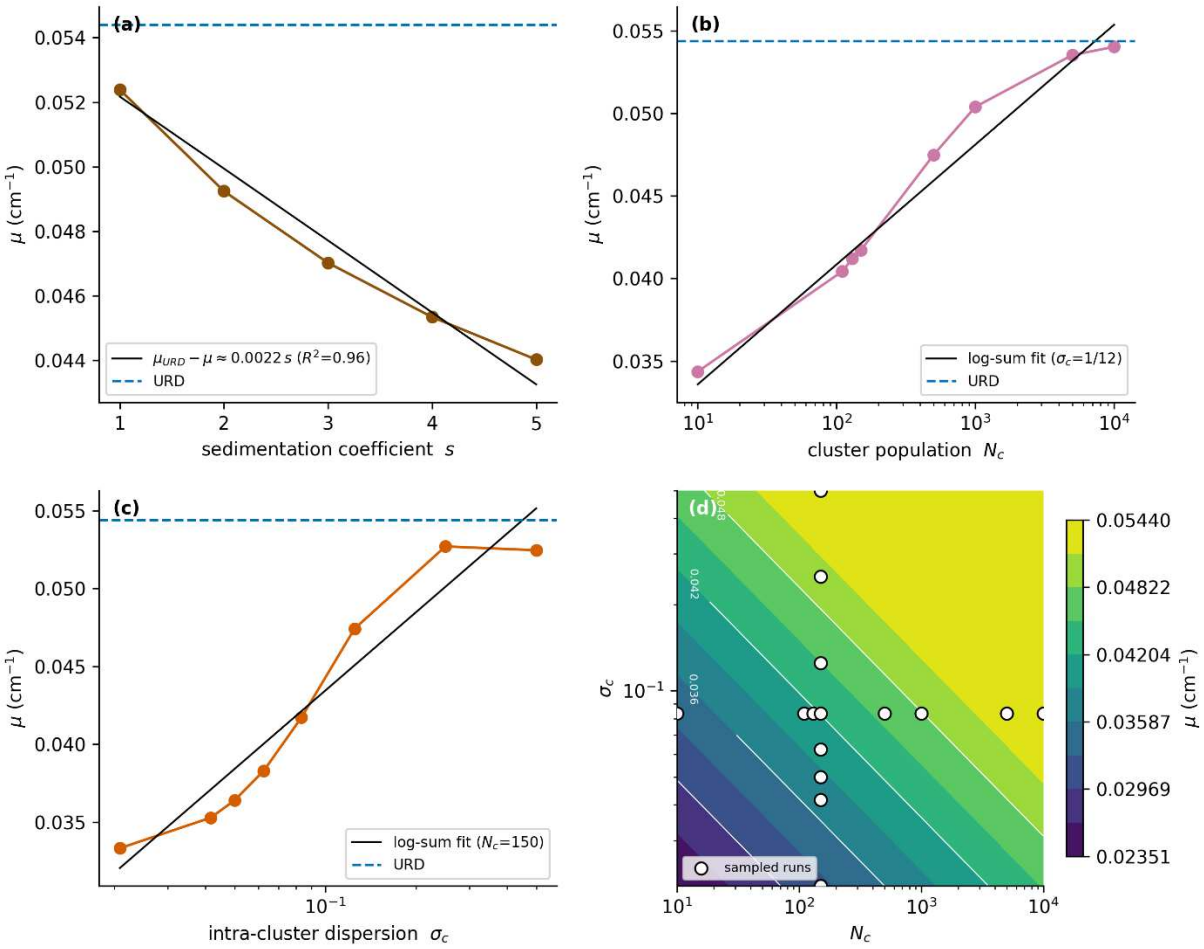


Figure 6. Parametric analysis of fabrication defects (E = 1 MeV; URD baseline μ = 0.0544 cm⁻¹). (a) Sedimentation sweep (s = 1–5). (b) Cluster-population sweep at $\sigma_c = 1/12$. (c) Dispersion sweep at $N_c = 150$. (d) Separable regression surface $\mu = A \ln N_c + B \ln \sigma_c + C$ ($R^2 = 0.95$) over the sampled domain (white markers); the dashed contour marks the URD level. Per decade, the dispersion lever is ≈2.3× more potent than cluster fragmentation.

Both sweeps collapse onto a compact separable regression, $\mu = (3.2\times10^{-3}) \ln N_c + (7.3\times10^{-3}) \ln \sigma_c + 4.4\times10^{-2}$ cm⁻¹ ($R^2 = 0.95$), whose coefficients quantify the relative potency of the two processing levers: per decade, dispersion recovers $\Delta\mu = B \ln 10 \approx 0.017$ cm⁻¹ (≈31% of the URD level), whereas fragmentation recovers (Table S4) only ≈0.007 cm⁻¹ (≈13%), a factor $B/A \approx 2.3$ in favor of dispersion quality over cluster count (Table S3). The logarithmic dependence implies diminishing returns, and the fitted surface overshoots the URD asymptote in the far corner where the clustered geometry converges to the URD; the regression is therefore strictly an interpolation over the sampled domain (residuals $\lesssim 3\times10^{-3}$ cm⁻¹ reflect model-form error, two orders of magnitude above the counting uncertainty). Within the sampled domain, recovering 95% of the URD attenuation requires either $N_c \gtrsim 2\times10^3$ (at $\sigma_c = 1/12$) or $\sigma_c \gtrsim 0.2$ (at $N_c = 150$), a direct defect-budget criterion for process qualification.

The sedimentation family obeys an even simpler law (the complete sweep is tabulated in Table S2): over s = 1–5, the penalty grows essentially linearly, $\mu_{URD} - \mu(s) \approx 0.0022 \cdot s$ cm⁻¹ ($R^2 = 0.96$, anchored at s = 0 ≡ URD), i.e., each unit of gradient exponent costs ≈4% of the URD attenuation; the reference case s = 2 corresponds to −9.5% at 1 MeV, deepening to −46% at 250 keV as the attenuation contrast grows. We thus find that defect severity is governed by the characteristic void dimension relative to the photon mean free path—a single criterion spanning both defect families.

3.5 Multi-Layer Architectures: Secondary-Radiation Management at a Primary-Attenuation Cost

The multi-layer architecture was evaluated as a secondary-radiation management scheme rather than a primary-attenuation maximizer. In a URD nanocomposite, photoelectrons generated inside a WO_3 particle almost immediately enter the high-Z field of neighboring particles and radiate Bremsstrahlung ($\sigma \propto Z^2$), partially converting the shield into a secondary X-ray source. The multi-layer design decouples generation from thermalization: electrons ejected in WO_3-rich planes cross into pure-PMMA spacers, where the Z^2-scaled radiative loss is replaced by collisional (ionization/excitation) losses, trapping the energy as heat, a strategy aligned with recent layered-structure engineering of γ-ray shields [20].

This management carries a quantified primary-attenuation cost: relative to the URD at identical loading, the multi-layer configuration exhibits −17% at 1 MeV, −1 to −3% at 500–750 keV, and −12% at 250 keV (Fig. 5a). The energy structure of this deficit is itself informative: at 250 keV, photoelectric absorption generates low-energy electrons with short ranges that are reliably contained within the spacers (maximizing the Bremsstrahlung suppression per unit filler), while at ~1 MeV Compton electrons traverse both spacers and subsequent WO_3 planes, bypassing the spatial separation. The net benefit of the trade-off is intrinsically a dose-based quantity that the LAC alone cannot capture, and is addressed through the secondary-particle scoring of the hot-cell analysis (Section 3.8). The architecture is therefore specifically motivated for the diagnostic energy window, precisely where ^{99m}Tc (140 keV) and most ^{131}I emissions lie.

3.6 Energy Dependence and Source-Geometry Effects

Across 250–1000 keV, the ranking of the ten distributions is largely preserved (Spearman $\rho = 0.99$ at 500–750 keV; 0.85 between 250 keV and 1 MeV), with two systematic trends (Fig. 5a): channeling-prone structures fall monotonically relative to the URD as energy decreases, σ_2D-maximizing architectures improve, and the multi-layer climbs from 7th to 4th place—confirming that the σ_2D mechanism is structural rather than spectroscopic in origin and scales with the photon mean free path relative to the characteristic void dimension. The energy-averaged LAC (Fig. 5b) reproduces the same ordering, and the extracted mean free paths ($\lambda = 1/\mu$, Fig. S6) directly visualize the void-channeling penalty of the Gaussian, clustered, and lattice models. Source geometry also matters: the two irradiation configurations of Fig. 1 probe different shielding responses: the collimated planar beam isolates the intrinsic structural attenuation, whereas the isotropic point-source configuration of the hot-cell analysis (Section 3.8) additionally weights oblique incidence paths and scattered photons. Intrinsic coefficients must therefore be extracted under good-geometry conditions, while design calculations must adopt realistic source terms.

3.7 Experimental Validation

We characterized the fabricated PMMA/WO_3 nanocomposite by narrow-beam gamma spectroscopy (^{152}Eu and ^{137}Cs), which yielded linear attenuation coefficients of 0.599, 0.303, 0.201, 0.140, and 0.0386 cm^{-1} at 121.8, 244.7, 344.3, 444.0, and 661 keV, respectively (Fig. 7; complete dataset: Table S7; representative spectra in Fig. S10).

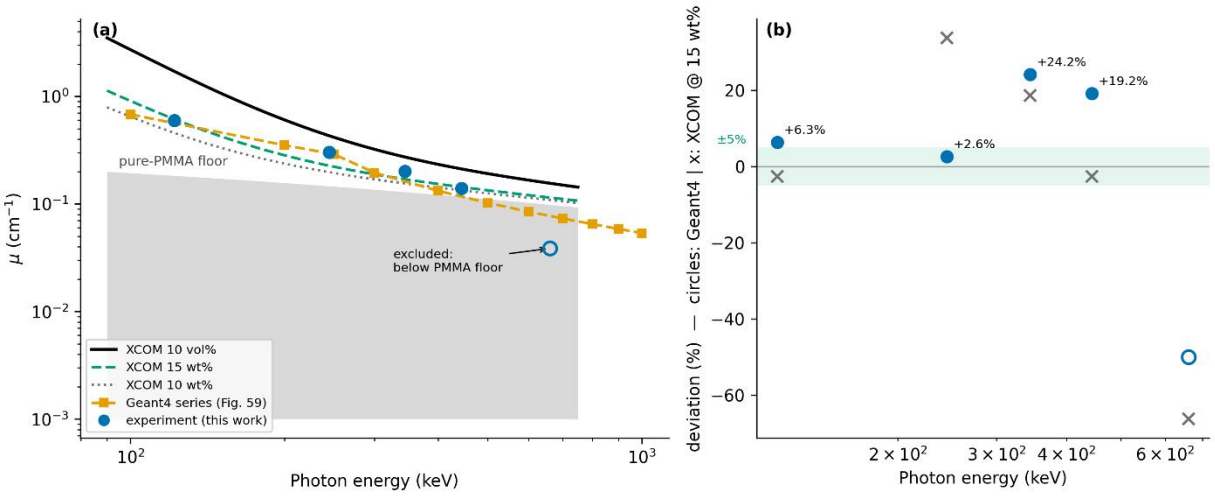


Figure 7. Experimental validation by narrow-beam gamma spectroscopy. (a) Measured LACs from ^{152}Eu and ^{137}Cs photo-peaks (circles) versus the Geant4 series (squares, log–log interpolation) and NIST-XCOM references (10 vol%, 15 wt%); the shaded region is the pure-PMMA physical bound, and the open symbol at 661 keV is excluded (below the bound; dead-time artifact). (b) Relative deviation from the Geant4 series (±5% band); crosses show the deviation from XCOM at 15 wt%. See Table S7.

Composition forensics. Before quantitative comparison, the effective filler loading of both the sample and the simulation series was established independently by XCOM inversion (Table S7): the loading implied by the artifact-free photo-peaks (121.8 and 444.0 keV) converges on ≈15 wt% WO_3, and the two lowest-energy points of the Geant4 series (0.679 and 0.351 cm^{-1} at 100 and 200 keV) back-propagate to the same loading. This mutual consistency confirms that the comparison is like-for-like at the fabricated composition.

Validation in the photoelectric regime. In the 122–245 keV range (where tungsten's photoelectric cross-section dominates and the composite derives its shielding advantage) the measured LACs agree with the Geant4 series (log–log interpolated) within +6.3% and +2.6% (Fig. 7b; complete comparison: Table S7). This agreement validates the low-energy photon transport, the material definition, and the attenuation-extraction procedure at the fabricated loading.

Mid-energy deviations and spectral artifacts. The +24.2% and +19.3% deviations at 344.3 and 444.0 keV follow a rate-dependent pattern consistent with documented spectrometric artifacts rather than transport-model error: the intense 121.78-keV line of ^{152}Eu produces random summing at 2 × 121.78 = 243.6 keV, directly beneath the 244.7-keV photopeak, and cascade sums contaminate the 344.3-keV region; these mechanisms inflate the unshielded reference peak areas and bias the extracted LACs upward. The pattern of deviations (maximal at the summing-exposed peaks, minimal at the isolated low-energy peak) is itself diagnostic of the artifact origin. We additionally note the mixed provenance of the simulation series (its 250- and 500-keV points coincide with the 10 vol% URD campaign values; Section 2.2.4), which bounds the precision of the mid-energy comparison (Table S7).

Excluded point. The 661-keV value lies below the pure-PMMA physical bound ($\mu \geq 0.093$ cm^{-1} at 661 keV for any PMMA-based composite) and is excluded from quantitative comparison, attributed to uncorrected dead-time losses in the high-rate unshielded ^{137}Cs acquisition (Table S7). The detection of this internal inconsistency by the physical-bound test demonstrates the value of the multi-reference validation protocol adopted here.

Scope. The present validation certifies the framework in the photoelectric regime (within 6%); the mid-energy agreement is bounded by spectrometric artifacts whose correction requires reduced-rate re-acquisition, and the structural campaign is benchmarked internally (Section 3.1). Two refinements (reduced-rate ^{152}Eu acquisition and

explicit dead-time correction of the ^{137}Cs pair) would tighten the chain without altering these conclusions. These artifacts notwithstanding, the photoelectric-regime agreement is consistent with the broader experimental record for WO_3-reinforced polymer shielding composites [13,21–23].

3.8 Case Study: Hot-Cell Shielding — Two Regimes and a Compliance Boundary

The engineering relevance of the framework was assessed with the dispensing hot-cell room model (Fig. 8a). Operator dose rates per unit activity were scored at 140, 364, and 511 keV, unshielded (0.87, 2.39, and 3.33 µSv (GBq·min)$^{-1}$) and behind the transparent PMMA/WO_3 window (5.4×10^{-7}, 8.6×10^{-4}, and 1.7×10^{-2} µSv (GBq·min)$^{-1}$; Fig. 8b). The ^{131}I/^{99m}Tc unshielded ratio (2.75) reproduces the point-source air-kerma prediction for the single 364-keV line (2.64) within 4%, validating the source term and room transport.

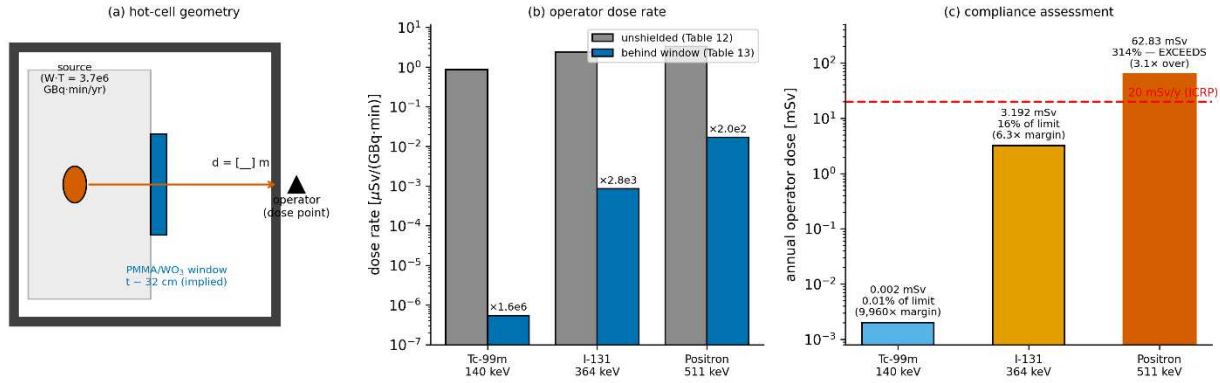


Figure 8. Hot-cell case study (workload 3.7×10^6 GBq·min per year; ICRP limit 20 mSv/y). (a) Simulated dispensing geometry. (b) Operator dose rates per unit activity, unshielded and behind the PMMA/WO_3 window (attenuation $\times1.6\times10^6$, $\times2.8\times10^3$, ×197 at 140, 364, and 511 keV). (c) Annual doses: ^{99m}Tc 0.002 mSv (0.01% of the limit), ^{131}I 3.19 mSv (16%), 511-keV positron emitters 62.8 mSv (314%, exceeding the limit ×3.1).

The shielded data reveal two distinct regimes. At 364 and 511 keV, the attenuation factors ($\times2.8\times10^3$ and ×197) back-propagate to a mutually consistent window thickness of ≈32 cm (narrow-beam), comparable to commercial hot-cell viewing windows; these energies are *window-limited*, and the attenuation follows the composite's energy spectrum as expected. At 140 keV, by contrast, the primary transmission through such a window is below 10^{-14}, while the observed attenuation is 1.6×10^6: the residual dose lies in the 10^{-5}–10^{-7} band characteristic of room scatter and streaming; the 140-keV design point is *room-limited*. Beyond this regime, additional window thickness (and equally, window microstructure (sedimentation, clustering; Section 3.4)) has no effect on the 140-keV operator dose; the design levers become room geometry, collimation, and the window's lateral extent.

Annual doses at the assumed workload ($W\cdot T = 3.7\times10^6$ GBq·min per year) expose a sharp compliance divide (Fig. 8c): ^{99m}Tc 0.002 mSv/y (0.01% of the 20 mSv/y limit; a 10^4-fold margin, room-scatter limited), ^{131}I 3.19 mSv/y (16%; 6.3-fold margin, window-limited), and 511-keV positron emitters 62.8 mSv/y (314% of the limit) consistent with the dedicated PET/CT shielding analyses of AAPM Task Group 108 [19]. The failure is structural: above the tungsten photoelectric regime the composite LAC collapses from ≈1.0 cm^{-1} (140 keV) to ≈0.19 cm^{-1} (511 keV), and PET compliance demands thicker sections, higher-Z or higher-loading formulations, or hybrid window designs.

For thin, material-limited shields (e.g., syringe shields and aprons), where the composite itself is the barrier, the defect analysis of Section 3.4 applies with full force: the s = 2 sedimentation penalty deepens from −9.5% at 1 MeV to −46% at 250 keV, i.e., a factor-≈1.9 thickness penalty at SPECT-relevant energies. The hot-cell window and the thin shield thus bracket the two design regimes in which microstructural architecture is, respectively, irrelevant and decisive; ignoring this distinction leads either to occupational overexposure or to systematic material overuse.

4 Conclusions

This study quantitatively demonstrates that the macroscopic shielding efficacy of polymer nanocomposites is dictated as much by the effective 2D geometric cross-section presented to the incident radiation as by chemical composition. The principal findings are:

1. Internal consistency and definitional pitfalls. At the working loading (10 vol%), the discrete-particle engine and the hand-built homogeneous definition differ by ≈5% (a definitional offset that widens to a 13-point swing (−5% to +8%) at 40 vol%) while the ratio of narrow-beam database values to the campaign's transmission-inferred coefficients corresponds to build-up factors of 1.35–1.67 ($\mu x = 0.57–1.06$), consistent with the shield's optical depth. At high loadings, RSA jamming ($\varphi \approx 0.383$) further limits overlap-free discrete placement, and discrete modeling becomes necessary beyond ~25 vol%.
2. Geometry dominates at fixed composition. At constant 10 vol% WO_3, the effective LAC at 1 MeV spans 0.033–0.074 cm^{-1} (×2.2) solely by rearranging particles, widening to ×6.5 at 250 keV as the filler/matrix attenuation contrast grows (≈19≈⟶6).
3. Anisotropy is a design variable. Lattice channeling and layered structures are strongly directional (layered ⊥/∥ anisotropy growing from ×1.3 to ×4.1 from 1 MeV to 250 keV); the golden-ratio distribution matches the homogeneous benchmark through 2D-projection optimization alone; the URD remains the most reliable omnidirectional paradigm.
4. Engineered architectures trade primary attenuation for secondary-radiation management. Multi-layer designs carry a −1% to −17% primary-attenuation deficit relative to the URD across 250–1000 keV but suppress secondary Bremsstrahlung by thermalizing photoelectrons in low-Z spacers; their net benefit is dose-based and accrues in the diagnostic energy window.
5. Micro-defects obey quantitative design laws. The sedimentation penalty grows linearly (≈4% of the URD level per unit gradient exponent, deepening to −46% at 250 keV), and the clustered regime follows $\mu = A \ln N_c + B \ln \sigma_c + C$ ($R^2 = 0.95$), with dispersion quality ≈2.3× more potent per decade than cluster fragmentation; defect severity is governed by the characteristic void dimension relative to the photon mean free path.
6. A scale-bridging methodology. The saturating power-law model ($\mu = \mu\infty - C \cdot d^q$, $R^2 = 0.996$) bridges the ~10^{15}-volume computational gap, predicting a ≈19% nanoscale enhancement consistent with experimental literature [8,9] and showing that the size benefit is essentially exhausted below ~100 μm.
7. Experimental validation with a self-auditing protocol. Narrow-beam gamma spectroscopy confirmed the transport model within 6% in the photoelectric regime (122–245 keV) at the as-fabricated loading (≈15 wt%, established by XCOM inversion); mid-energy deviations were attributed to rate-dependent spectral artifacts, and one point violating the physical lower bound was identified and excluded by the validation protocol itself.
8. Two shielding regimes in practice. In the hot-cell case study, the nanocomposite window is room-scatter-limited at 140 keV (0.01% of the occupational limit) but window-limited at 364–511 keV, where positron-emitter workloads exceed the limit ×3.1 [19], delimiting the applicability of single-material nanocomposite windows and motivating hybrid designs for PET.

Future work should focus on processing-controlled architecture (surfactant-assisted dispersion, vacuum-assisted curing, in-situ polymerization, additive manufacturing of graded and multi-layer shields); on importing real 3D morphologies (e.g., micro-CT) directly into discrete Monte Carlo models; on primary-only re-simulation and reduced-rate spectroscopy to tighten the validation chain; and on extending the framework to mixed neutron–gamma fields. The design philosophy that follows from this work is unambiguous: the next generation of shielding materials will be architected, not merely composed.

Data availability

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

Funding Declaration

The authors declare that no funds, grants, or other support were received during the preparation of this manuscript.

References

[1] ICRP (2007) The 2007 recommendations of the International Commission on Radiological Protection. ICRP Publication 103. Ann ICRP 37(2–4)

[2] More CV, Alsayed Z, Badawi MS, et al (2021) Polymeric composite materials for radiation shielding: a review. Environ Chem Lett 19:2057–2090. https://doi.org/10.1007/s10311-021-01189-9

[3] Jayakumar S, Saravanan T, Philip J (2023) A review on polymer nanocomposites as lead-free materials for diagnostic X-ray shielding: recent advances, challenges and future perspectives. Hybrid Adv 4:100100. https://doi.org/10.1016/j.hybadv.2023.100100

[4] Bawazeer O, Makkawi K, Aga ZB, et al (2023) A review on using nanocomposites as shielding materials against ionizing radiation. J Umm Al-Qura Univ Appl Sci 9:325–340. https://doi.org/10.1007/s43994-023-00042-9

[5] Wang B, Qiu T, Yuan L, Fang Q, Wang X, Guo X, et al (2023) A comparative study between pure bismuth/tungsten and the bismuth tungsten oxide for flexible shielding of gamma/X rays. Radiat Phys Chem 208:110906. https://doi.org/10.1016/j.radphyschem.2023.110906

[6] Kim DW, Jeon H, Joo JH, Ki Y, Yoon SH, Kim YS, Kang SW, Lee T, Seol HY (2026) Tungsten-bismuth shielding sheets for radiation safety in fluoroscopy-guided bronchoscopy: a phantom study. J Thorac Dis 18(2):130. https://doi.org/10.21037/jtd-2025-1763

[7] Marshall SK, Kwandee P, Songphum N, Chuaymuang J (2025) Comprehensive simulation-based evaluation of gamma radiation shielding performance of bismuth oxide- and tungsten oxide-reinforced polymer composites for nuclear medicine occupational safety. Polymers 17(11):1491. https://doi.org/10.3390/polym17111491

[8] Cheewasukhanont W, Limkitjaroenporn P, Kothan S, Kedkaew C, Kaewkhao J (2020) The effect of particle size on radiation shielding properties for bismuth borosilicate glass. Radiat Phys Chem 172:108791. https://doi.org/10.1016/j.radphyschem.2020.108791

[9] Li R, Gu Y, Wang Y, Yang Z, Li M, Zhang Z (2017) Effect of particle size on gamma radiation shielding property of gadolinium oxide dispersed epoxy resin matrix composite. Mater Res Express 4:035035. https://doi.org/10.1088/2053-1591/aa5d09

[10] Zare Y (2016) Study of nanoparticles aggregation/agglomeration in polymer particulate nanocomposites by mechanical properties. Compos Part A Appl Sci Manuf 84:158–164. https://doi.org/10.1016/j.compositesa.2016.07.012

[11] Berger MJ, Hubbell JH, Seltzer SM, Chang J, Coursey JS, Sukumar R, Zucker DS, Olsen K (2010) XCOM: photon cross sections database (version 1.5). NIST Standard Reference Database 8 (XGAM). National Institute of Standards and Technology, Gaithersburg. https://doi.org/10.18434/T48G6X

[12] Kazemi F, Malekie S, Hosseini SMA (2019) A Monte Carlo study on the shielding properties of a novel polyvinyl alcohol (PVA)/WO_3 composite against gamma rays, using the MCNPX code. J Biomed Phys Eng 9:465–472. https://doi.org/10.31661/jbpe.v0i0.1114

[13] Mokhtari Dorostkar M, Abdi Saray A (2025) Development of PMMA-based polymer composite incorporating WO_3 for gamma radiation shielding using synthesis and Monte Carlo simulation. Sci Rep 15:27346. https://doi.org/10.1038/s41598-025-11155-y

[14] Agostinelli S, Allison J, Amako K, et al (2003) Geant4 — a simulation toolkit. Nucl Instrum Methods Phys Res A 506:250–303. https://doi.org/10.1016/S0168-9002(03)01368-3

[15] Allison J, Amako K, Apostolakis J, et al (2016) Recent developments in Geant4. Nucl Instrum Methods Phys Res A 835:186–224. https://doi.org/10.1016/j.nima.2016.06.125

[16] Trubey LK, Eisenhauer C, Foderaro A, Gopinath D, Harima Y, Hubbell JH, Shure K, Su S (1991) Gamma-ray attenuation coefficients and buildup factors for engineering materials. ANSI/ANS-6.4.3-1991. American Nuclear Society, La Grange Park, IL

[17] Vogel H (1979) A better way to construct the sunflower head. Math Biosci 44:179–189. https://doi.org/10.1016/0025-5564(79)90075-4

[18] ICRP (2007) Radiological protection in medicine. ICRP Publication 105. Ann ICRP 37(6)

[19] Madsen MT, Anderson JA, Halama JR, Kleck J, Simpkin DJ, Votaw JR, Wendt RE III, Williams LE, Yester MV (2006) AAPM Task Group 108: PET and PET/CT shielding requirements. Med Phys 33:4–15. https://doi.org/10.1118/1.2135911

[20] Duan X, et al (2025) Layered structural engineering of Bi_2O_3/PP and WO_3/PP composites for γ-ray shielding in high energy range: anisotropic attenuation mechanisms via Monte Carlo simulation and experiments. J Mater Res Technol 36:10324–10336. https://doi.org/10.1016/j.jmrt.2025.05.223

[21] Isazadeh F, Saray AA (2025) Experimental and simulation assessment of gamma-radiation shielding properties of the polyester polymer reinforced with WO_3. Sci Rep 15. https://doi.org/10.1038/s41598-025-21003-8

[22] Al-Ghamdi H, Hemily HM, Saleh I, Ghataas ZF, Abdel-Halim A, Sayyed M, Yasmin S, Almuqrin A, Elsafi M (2022) Impact of WO_3-nanoparticles on silicone rubber for radiation protection efficiency. Materials 15:5706. https://doi.org/10.3390/ma15165706

[23] Khudhair I, Ahmed RH, Aziz AA (2026) The effect of WO_3 nanoparticles on the structural and morphological properties of epoxy composites for gamma shielding application. Phys Scr 101. https://doi.org/10.1088/1402-4896/ae32cf

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- Supplementary.docx