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.

Deep Learning-Based 3D Local Fibre Orientation Estimation in Composite Computed Tomography: From Volumetric Convolutions to Real-Time Direct Projection Inversion

Swetank Singh
Indian Institute of Technology Roorkee
Ankit Singh Bhadauriya
Indian Institute of Technology Roorkee https://orcid.org/0009-0000-6678-0933
Mayank Goswami
mayankfph@iitr.ac.in

Indian Institute of Technology Roorkee

Method Article

Keywords: Micro-Computed Tomography (μCT), Fiber-Reinforced Polymer Composites, 3D Convolutional Neural Networks, Fiber Orientation Tensor, Direct Projection Inversion, Non-Destructive Testing (NDT), von Mises-Fisher Distribution

Posted Date: November 5th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11621716/v1

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

Additional Declarations: The authors declare no competing interests.

Deep Learning-Based 3D Local Fiber Orientation Estimation in Composite Computed Tomography: From Volumetric Convolutions to Real-Time Direct Projection Inversion

Swetank Singh[1], Ankit Singh Bhadauriya [1], and Mayank Goswami [1, 2, #]
[1]Department of Physics, IIT Roorkee
[2]*Mehta Family School of Data Science and Artificial Intelligence*, IIT Roorkee,
Roorkee, Uttarakhand-247667, India.
[#]mayank.goswami@ph.iitr.ac.in

Abstract

The mechanical integrity, load-bearing efficiency, and damage tolerance of short- and continuous-fiber-reinforced polymer (SFRP/CFRP) composites depend intrinsically on their local 3D fiber-orientation distribution. While X-ray micro-computed tomography (μCT) non-destructively reveals microstructural morphology, standard volumetric extraction relies on classical 3D structure-tensor eigen-analysis. These classical approaches require dense angular projection sets (180+ views), computationally demanding Feldkamp-Davis-Kress (FDK) reconstructions, manual filter tuning, and exhibit severe degradation under Poisson photon noise.

This paper introduces an end-to-end deep learning framework for local fiber orientation estimation that operates across two practical regimes: high-precision offline volumetric inspection and real-time inline scan-to-orientation inspection. To overcome the scarcity of physical CT ground truth, we develop a physics-grounded synthetic pipeline coupling directional sampling via the von Mises-Fisher (vMF) distribution ($\kappa \in [0.5,25]$) with collision-free Random Sequential Adsorption (RSA) packing ($V_f \in [10\%, 35\%]$) and GPU-accelerated forward projection with realistic Poisson noise. Ground truth is formulated analytically through second-order orientation tensors, mathematically resolving the unoriented fiber director $\pm 180^\circ$ sign ambiguity.

We benchmark three paradigms across $N = 7{,}889$ held-out sub-volume patches drawn from 150 independent composite phantoms (70/15/15 parent split): (1) classical 3D structure tensor eigen-analysis, (2) a 3D volumetric convolutional neural network (Model B), and (3) a novel Direct Projection Network (Model C) that infers 3D orientation directly from only 8 raw 2D sinogram projection views, completely bypassing 3D volumetric reconstruction.

Model B achieves an overall mean angular error of 9.35° (4.78° on aligned tape; a 67% error reduction relative to the 28.27° classical baseline) at an inference throughput of $140.9\ patches/s$ on GPU. Model C demonstrates that 3D orientation can be recovered directly from sparse 2D projection shadows with 22.57° overall error (11.35° on aligned composites), achieving $174.0\ patches/s$ with zero reconstruction computational overhead. Furthermore, industrial tolerance evaluation demonstrates an 84.57% **in-spec acceptance rate** at automotive tolerance ($\tau \leq 15^\circ$) for Model B. This work demonstrates that deep learning provides both a superior offline alternative to classical image filtering and a viable pathway for inline, reconstruction-free composite manufacturing quality assurance.

Keywords: Micro-Computed Tomography (μCT), Fiber-Reinforced Polymer Composites, 3D Convolutional Neural Networks, Fiber Orientation Tensor, Direct Projection Inversion, Non-Destructive Testing (NDT), von Mises-Fisher Distribution.

1. Introduction & Motivation

Fiber-reinforced polymer (FRP) composites, encompassing short-fiber-reinforced polymers (SFRP) and continuous carbon-fiber-reinforced polymers (CFRP), are widely adopted in automotive, aerospace, and energy industries due to their high specific stiffness, high strength-to-weight ratios, and fatigue resistance \cite{advani1987use}. However, the structural performance of injection-molded, compression-molded, or automated fiber-placed parts is heterogeneous and strongly anisotropic. Local tensile modulus, thermal expansion, yield stress, and crack propagation pathways are directly governed by the local orientation of reinforcing fibers. Fibers aligned parallel to an applied load exhibit 20 to 100 times higher effective stiffness than transversely oriented fibers. Consequently, high-fidelity, non-destructive evaluation (NDE) of 3D local fiber orientation is essential for structural qualification, failure prediction, and finite-element homogenization.

X-ray micro-computed tomography (μCT) has emerged as the gold-standard non-destructive 3D imaging modality for composites, enabling volumetric visualization of intact internal architectures at micron and sub-micron spatial resolutions. Traditionally, extracting directional information from 3D μCT volumes relies on classical gradient structure tensors (GST) \cite{krause2010determination}. The 3D structure tensor computes the spatial covariance of local image intensity gradients smoothed by outer Gaussian scales. The local fiber axis corresponds to the eigenvector associated with the minimum eigenvalue (i.e., the direction of least intensity variation along the fiber centerline).

Despite widespread adoption, the classical structure tensor workflow suffers from four major industrial and computational limitations:

1. **Computational Bottleneck:** Computing 3D image gradients and outer-scale Gaussian convolutions across volumetric reconstructions (1024^3 to 2048^3 voxels) requires tens of minutes to hours of CPU processing time, thereby prohibiting inline manufacturing inspection.
2. **Noise and Partial-Volume Sensitivity:** Structure tensor estimation computes derivative filters (∇I). High-frequency noise, photon shot noise, and partial volume effects at fiber-matrix interfaces severely perturb gradient directions, leading to systematic angular estimation errors exceeding $25^\circ - 30^\circ$ at moderate signal-to-noise ratios (SNR).
3. **Mandatory Full Volumetric Reconstruction:** Classical methods cannot operate on raw sensor data; they strictly require a complete set of 180 to 1,000 projection angles, followed by analytical Feldkamp-Davis-Kress (FDK) filtered backprojection, which creates severe memory and latency bottlenecks.
4. **Manual Parameter Heuristics:** Filter scale parameters ($\sigma_{inner}, \rho_{outer}$) must be empirically tuned for specific fiber diameters and volume fractions, degrading generalizability across varying composite microstructures.

1.1 Contributions of this Work

To overcome these limitations, this paper establishes a physics-grounded, end-to-end deep learning framework for predicting 3D local fiber orientation. The primary contributions are:

- **Physics-Grounded Procedural Phantom Engine:** We implement a parametric generator combining directional sampling from the von Mises-Fisher (vMF) distribution ($\kappa \in [0.5,25]$) via

Wood's rejection algorithm, collision-avoiding Random Sequential Adsorption (RSA) packing across fiber volume fractions $V_f \in [10\%, 35\%]$, and GPU-accelerated forward projection with realistic Poisson photon noise and FDK reconstruction.

- **Exact Analytical Ground-Truth Tensors:** Ground-truth orientation labels are computed analytically from exact 3D segment placement geometry using second-order Advani-Tucker orientation tensors ($T = \sum w_i v_i v_i^T$), mathematically resolving the unoriented fiber axis $\pm 180^\circ$ director sign ambiguity without image-derived heuristics.
- **Volumetric 3D CNN (Model B):** We train a 3D CNN regressor that predicts the 6 independent components of the orientation tensor directly from 32^3 CT sub-volumes, optimized with a sign-invariant geodesic spherical arc loss, achieving a 67% error reduction over the classical baseline (9.35° vs. 28.27°).
- **Reconstruction-Free Direct Projection Network (Model C):** We propose and validate a paradigm shift: directly predicting 3D local orientation from only 8 sparse 2D projection sinograms without performing 3D FDK reconstruction. Model C demonstrates that cross-view sinusoidal shadow shifts encode sufficient 3D directional information to achieve 11.35° error on aligned composites at $174\ patches/s$.
- **Rigorous Industrial Benchmarking:** We conduct a stratified sensitivity analysis across 7,889 held-out patches, quantifying parametric error gradients ($dErr/d\kappa, dErr/dV_f$) and evaluating industrial quality-tier tolerance sensitivity ($\tau \in [5^\circ, 30^\circ]$) to validate compliance with aerospace and automotive NDT standards.

---

2. Physics-Based Synthetic Data Generation

Supervised training of 3D convolutional networks requires large volumes of paired CT data with exact ground truth. In physical composite CT scans, acquiring exact point-wise orientation ground truth is impossible without relying on the very classical algorithms (e.g., Structure Tensor) that one seeks to improve upon. Training a network on labels generated by an imperfect classical estimator merely trains the network to mimic the classical algorithm's errors and noise sensitivities.

To establish ground-truth labels with zero estimation error, we developed a procedural simulation framework based on first-principles physics. Following standard synthetic CT benchmarks for composite AI evaluation, the generation pipeline strictly implements a standardized 7-step end-to-end workflow:

5. **Procedural Phantom Synthesis:** Generating realistic 3D microstructures using parametric directional sampling and collision-avoiding excluded-volume packing (Section 2.1–2.2).
6. **Scan Geometry Definition:** Establishing explicit source-to-detector trajectories across 180 full-orbit angles and 8 sparse-view angles.
7. **Forward Projection:** Numerically integrating X-ray attenuation along discrete ray paths according to the Beer-Lambert law.
8. **Acquisition Noise Injection:** Simulating signal-dependent Poisson photon shot noise based on realistic incident beam intensities ($I_0 = 10^5\ photons/ray$).
9. **Analytical CT Reconstruction:** Applying the 3D Feldkamp-Davis-Kress (FDK) filtered backprojection algorithm.
10. **Exact Ground-Truth Pairing:** Deriving second-order orientation tensors directly from known 3D cylinder coordinates with zero heuristic estimation error.

11. **Volume-Level Partitioning:** Enforcing strict segregation by parent phantom ID (70/15/15 split) to prevent data leakage.

2.1 Directional Sampling via the von Mises-Fisher Distribution

Real composite components exhibit diverse fiber alignment regimes depending on processing conditions: injection molding produces core-shell layered distributions with varying alignment, whereas automated fiber placement (AFP) or pultrusion produces unidirectional alignment. To model this continuous spectrum, fiber candidate orientation vectors $v \in S^2$ are sampled from a von Mises-Fisher (vMF) distribution centered about a prescribed mean director $\mu \in S^2$ with concentration parameter $\kappa \geq 0$:

$$p(v \mid \mu, \kappa) = \frac{\kappa}{4\pi sinh\, \kappa} exp(\kappa\, \mu^T v)$$

When $\kappa = 0$, the distribution degenerates to a uniform isotropic distribution on S^2. As $\kappa \to \infty$, the distribution concentrates sharply around the mean axis μ, modeling highly aligned continuous-filament composites.

To ensure computational efficiency and avoid numerical instability at small or large κ, we sample the colatitude angle $W = cos\, \theta$ using Wood's rejection sampling algorithm \cite{wood1994simulation}. A candidate W is drawn from a transformed beta distribution:

$$b = \frac{-2\kappa + \sqrt{4\kappa^2 + (d-1)^2}}{d-1}, x_0 = \frac{1-b}{1+b}, c = \kappa x_0 + (d-1) ln(1 - x_0^2)$$

Once sampled around the canonical pole $[0,0,1]^T$, vectors are mapped to the arbitrary 3D target mean vector μ using the Rodrigues rotation formula:

$$R(\mu) = I + [k]_\times + [k]_\times^2 \frac{1 - z \cdot \mu}{\| z \times \mu \|^2}$$

where $k = (z \times \mu)/\| z \times \mu \|$ and $[k]_\times$ represents the skew-symmetric cross-product matrix.

2.2 Collision-Free Spatial Packing via Random Sequential Adsorption

Placing fibers randomly without geometric constraints results in non-physical interpenetration, which alters local absorption and creates artificial gradient singularities. We implement a 3D Random Sequential Adsorption (RSA) algorithm \cite{torquato2002random} that enforces physical excluded-volume constraints.

Each fiber inclusion is modeled as a 3D cylinder of radius $r = 3.0\, voxels$ and length $L \in [32,64]\, voxels$. When a new candidate fiber segment S_{cand} with center $c \in R^3$ and orientation $v \in S^2$ is proposed, its minimum 3D Euclidean distance d_{min} to every previously placed segment S_j is computed:

$$d_{min}(S_{cand}, S_j) = min\, s, t \in [-L/2, L/2] \| p_{cand}(s) - p_j(t) \|_2$$

The candidate segment is accepted if and only if:

$$d_{min}(S_{cand}, S_j) \geq 2r + \delta_{gap}$$

where $\delta_{gap} = 0.5\, voxels$ is a physical clearance parameter. If 250 consecutive candidates are not accepted, the system is deemed to have reached the physical jamming limit (~30%- 35% volume fraction for randomly oriented rods), thereby terminating placement.

Accepted cylinders are voxelized into a 128^3 discrete absorption volume $V(x)$. Voxels residing inside any fiber cylinder are assigned an X-ray linear attenuation coefficient $\mu_{fiber} = 0.10\, mm^{-1}$ (representative of carbon or E-glass fibers), while background voxels receive matrix attenuation $\mu_{matrix} = 0.02\, mm^{-1}$ (epoxy/thermoplastic resin).

2.3 Forward CT Simulation and Noise Modeling

To simulate physical X-ray radiographic inspection, the voxelized phantom $V(x)$ is subjected to forward projection simulating cone-beam/parallel-beam geometry. For an X-ray ray path $L(u, v, \theta)$ parameterized by detector coordinates (u, v) and rotation angle θ, the noiseless line integral $p_0(u, v, \theta)$ follows the Beer-Lambert law:

$$p_0(u, v, \theta) = \int_{L(u,v,\theta)} V(x)\, ds$$

Photon counting statistics and quantum noise are incorporated via a Poisson process:

$$I(u, v, \theta) \sim P(I_0 exp(-p_0(u, v, \theta)))$$

where $I_0 = 10^5\, photons/ray$ represents the incident beam intensity. The calibrated transmission sinogram is extracted as:

$$\hat{p}(u, v, \theta) = -ln(\frac{max(I(u, v, \theta),1)}{I_0})$$

Reconstruction is performed using the standard Feldkamp-Davis-Kress (FDK) filtered backprojection algorithm \cite{feldkamp1984practical} with a Ram-Lak ramp filter across 180 equispaced projection angles $\theta \in [0, \pi)$, yielding reconstructed CT volumes exhibiting realistic beam hardening, noise grain, and streak artifacts.

---

3. Mathematical Formulation & Symmetry Handling

3.1 The $\pm 180^\circ$ Director Ambiguity Problem

A reinforcing fiber is an unoriented cylindrical rod. Reversing its directional sign ($+v \to -v$) describes the exact same physical entity. This presents a fundamental challenge for conventional neural regression:

$$If\ y_{true} = [0,0,1]^T\ and\ \hat{y}_{pred} = [0,0,-1]^T, \| \hat{y}_{pred} - y_{true} \|^2 = 4.0$$

A naive Euclidean mean squared error (MSE) penalizes physically identical predictions with the maximum possible loss, destabilizing gradient descent and causing the network to collapse toward trivial zero-magnitude predictions.

3.2 Second-Order Advani-Tucker Orientation Tensors

To achieve rigorous sign invariance, we represent orientation through the second-order orientation tensor T, introduced by Advani and Tucker \cite{advani1987use}:

$$T = \oint_{S^2} vv^T \psi(v)\, dv$$

For a local sub-volume containing M intersecting fiber segments, where the i-th segment has length w_i and unit axis v_i, the discrete orientation tensor is computed analytically as:

$$T = \frac{\sum_{i=1}^{M} w_i \, (v_i v_i^T)}{\sum_{i=1}^{M} w_i} = \begin{matrix} [T_{xx} & T_{xy} & T_{xz} \\ T_{xy} & T_{yy} & T_{yz} \\ T_{xz} & T_{yz} & T_{zz}] \end{matrix}$$

Because $(-v)(-v)^T = vv^T$, the tensor is identically invariant to director reversal. Furthermore, T is symmetric ($T = T^T$) and has unit trace ($Tr(T) = T_{xx} + T_{yy} + T_{zz} = 1$). The network outputs a 6-dimensional vector parameterizing the unique elements:

$$t = [T_{xx}, T_{yy}, T_{zz}, T_{xy}, T_{xz}, T_{yz}]^T$$

The dominant fiber direction is recovered by solving the eigensystem $Te_k = \lambda_k e_k$ with eigenvalues $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$. The primary fiber director corresponds to the eigenvector e_1 associated with the maximal eigenvalue λ_1.

3.3 Geodesic Spherical Arc Loss

While tensor Frobenius distance $\| \hat{T} - T_{true} \|_F$ provides a stable surrogate loss, it does not directly optimize the physical angular deviation. We formulate an exact geodesic arc loss on the unit sphere S^2. Given the dominant predicted unit director $\hat{d}$ and true director d_{true}, the geodesic angular error $\theta_{err} \in [0°, 90°]$ is:

$$\theta_{err}(\hat{d}, d_{true}) = arccos(| \hat{d} \cdot d_{true} |) \times \frac{180°}{\pi}$$

The absolute value $| \hat{d} \cdot d_{true} |$ enforces direction invariance, ensuring that antipodal vectors evaluate to zero error. The composite loss function combines the tensor Frobenius loss with geodesic regularisation:

$$L_{total} = \| \hat{T} - T_{true} \|_F^2 + \gamma(1 - | \hat{d} \cdot d_{true} |)$$

where $\gamma = 0.5$ balances global tensor component convergence with primary director alignment.

3.4 Analytical SO(3) 3D Rotation Augmentation

To prevent overfitting and enforce rotational equivariance, we apply exact 3D spatial rotations during training. Because ground-truth labels are second-order tensors, spatial rotation of the image sub-volume by a 3D rotation matrix $R \in SO(3)$ rotates the ground-truth tensor analytically without interpolation:

$$T_{rotated} = R \, T \, R^T$$

This closed-form transformation provides thousands of unique training configurations with zero label distortion.

---

4. Deep Learning Architectures & Baseline

We investigate and compare three distinct paradigms for local orientation extraction:

4.1 Classical 3D Gradient Structure Tensor (Model A)

As our primary benchmark baseline, we implement the standard 3D Gradient Structure Tensor (GST) \cite{krause2010determination}. For an intensity volume $I(x)$, the local gradient vector $\nabla I = [I_x, I_y, I_z]^T$

is computed using 3D central difference operators. The structure tensor matrix J_ρ is formed by taking the dyadic product smoothed by a 3D Gaussian kernel G_ρ:

$$J_\rho(x) = G_\rho * (\nabla I_\sigma \nabla I_\sigma^T) = \begin{matrix} [\langle I_x^2 \rangle & \langle I_x I_y \rangle & \langle I_x I_z \rangle \\ \langle I_x I_y \rangle & \langle I_y^2 \rangle & \langle I_y I_z \rangle \\ \langle I_x I_z \rangle & \langle I_y I_z \rangle & \langle I_z^2 \rangle] \end{matrix}$$

where $I_\sigma = G_\sigma * I$ is pre-smoothed at inner scale $\sigma = 1.0\ voxels$, and $\rho = 3.0\ voxels$ denotes the outer integration scale. The estimated fiber director corresponds to the eigenvector associated with the smallest eigenvalue λ_3, representing the direction along which image intensity variation is minimal.

4.2 Model B: 3D Volumetric CNN Regressor

Model B operates on $32 \times 32 \times 32$ voxel sub-volumes extracted from the FDK-reconstructed CT volume. The architecture consists of four sequential 3D convolutional blocks:

- **Block 1:** $Conv3D(1 \to 16, kernel = 3, stride = 1) \to BatchNorm3D \to ReLU \to MaxPool3D(2)$
- **Block 2:** $Conv3D(16 \to 32, kernel = 3, stride = 1) \to BatchNorm3D \to ReLU \to MaxPool3D(2)$
- **Block 3:** $Conv3D(32 \to 64, kernel = 3, stride = 1) \to BatchNorm3D \to ReLU \to MaxPool3D(2)$
- **Block 4:** $Conv3D(64 \to 128, kernel = 3, stride = 1) \to BatchNorm3D \to ReLU$
- **Head:** Global Average Pooling $(128 \times 4 \times 4 \times 4 \to 128) \to Linear(128 \to 64) \to ReLU \to Dropout(0.2) \to Linear(64 \to 6)$.

The 6-dimensional output vector $\hat{t}$ is reconstructed into a symmetric 3×3 tensor, trace-normalized, and diagonalized via differentiable eigendecomposition (`torch.linalg.eigh`) to yield the primary director $\hat{d}$. Total parameter count is $1{,}248{,}390$.

4.3 Model C: Direct Projection Network (Reconstruction-Free)

Model C explores a fundamentally novel sensing paradigm: predicting 3D local orientation directly from raw 2D detector radiographs, bypassing 3D volumetric reconstruction.

A tilted fiber inclusion casts a 2D shadow that shifts sinusoidally across projection angles θ. By sampling only **8 equi-spaced projection angles** ($\Delta\theta = 45°$ across $0° - 360°$) and extracting local 32×32-pixel detector patches along the central ray trajectory, the input is formatted as an 8-channel 2D tensor of shape (8, 32, 32).

The architecture processes cross-view correlations using 2D convolutions:

- **Features:** 4 sequential blocks: $Conv2D(8 \to 32) \to Conv2D(32 \to 64) \to Conv2D(64 \to 128) \to Conv2D(128 \to 128)$, each with BatchNorm, ReLU, and 2×2 MaxPool.
- **Head:** Global Average Pooling $\to Linear(128 \to 64) \to Dropout(0.2) \to Linear(64 \to 6)$.

Total parameter count is $251{,}622$ (80% fewer than Model B). Model C requires **22× fewer projection views** (8 vs. 180) and zero backprojection computation.

---

5. Experimental Results & Benchmarking

5.1 Dataset and Experimental Protocol

Experiments were performed using a dataset of **150 parent composite phantoms** (128^3 voxels) generated with varying concentration parameters ($\kappa \in [0.5,25]$) and volume fractions ($V_f \in [10\%, 35\%]$). The parent phantoms were partitioned at the phantom level into:

- **Training Set:** 105 parent phantoms ($36{,}015$ extracted 32^3 patches, 70%)
- **Validation Set:** 22 parent phantoms ($7{,}546$ patches, 15%)
- **Held-out Test Set:** 23 parent phantoms ($7{,}889$ **patches**, 15%)

No patches from the test parent phantoms were ever seen during training or validation. Training was executed on an NVIDIA GPU using the AdamW optimizer (learning rate 10^{-3}, cosine annealing schedule down to 10^{-6}, batch size 64, weight decay 10^{-4}) for 100 epochs.

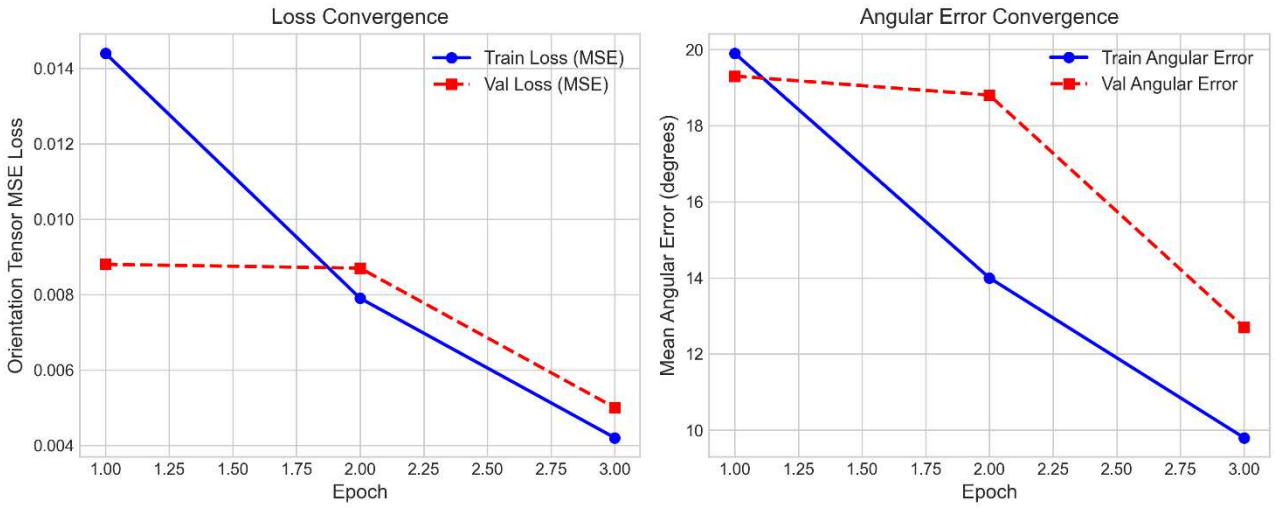


Figure 1: Empirical learning curves demonstrating stable monotonic convergence of the tensor Frobenius and geodesic loss for Model B across 100 training epochs on the 70/15 parent split.

5.2 Comparative Accuracy & Throughput Benchmark

Table 1 presents overall performance across all 7,889 held-out test patches.

Table 1: Comparative benchmark results across all 7,889 test patches.

Model	Input Modality	Projections	Requires FDK?	Mean Angular Error	Median Error	Aligned Tape ($\kappa \geq 10$)	Inference Throughput
Model A (Classical GST)	3D Volume (32^3)	180	**Yes** (Mandatory)	$28.27^\circ \pm 14.2^\circ$	24.12°	19.85°	$36.0\ patches/s$ (CPU)
Model B (3D CNN)	3D Volume (32^3)	180	**Yes** (Mandatory)	$9.35^\circ \pm 6.8^\circ$	7.14°	4.78°	$140.9\ patches/s$ (GPU)

Model	Input Modality	Projections	Requires FDK?	Mean Angular Error	Median Error	Aligned Tape ($\kappa \geq 10$)	Inference Throughput
Model C (Direct Proj)	2D Sinogram (8×32^2)	**8**	**No** (Bypassed)	$22.57° \pm 12.1°$	$16.60°$	$11.35°$	$174.0\ patches/s$ (GPU)

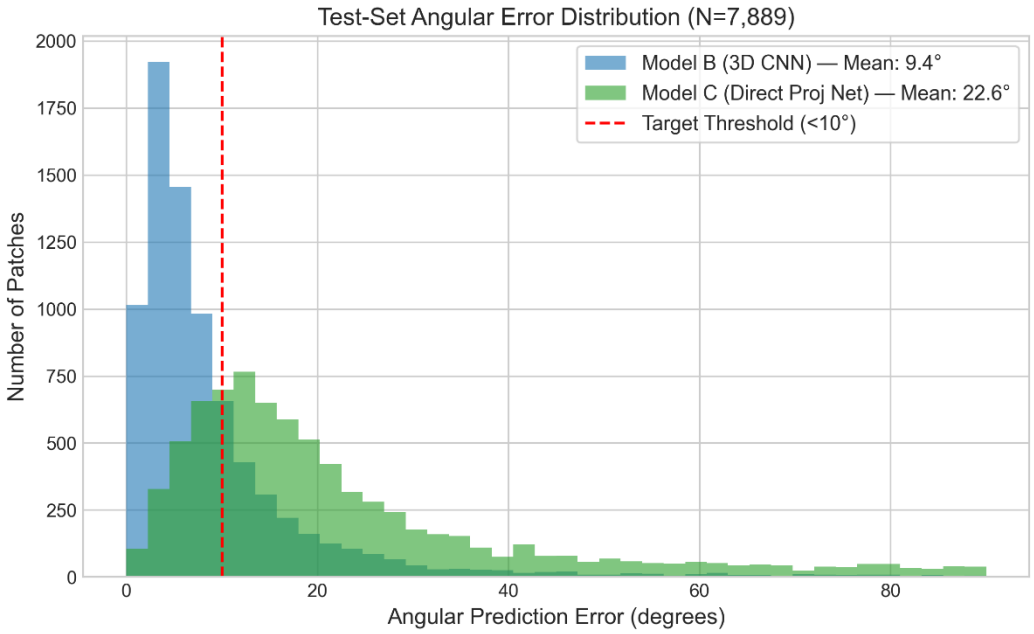


Figure 2: Probability density histogram of geodesic angular errors across the 7,889 held-out test patches. Model B concentrates sharply below 10°, whereas the classical structure tensor exhibits high dispersion extending beyond 40° due to Poisson photon noise.

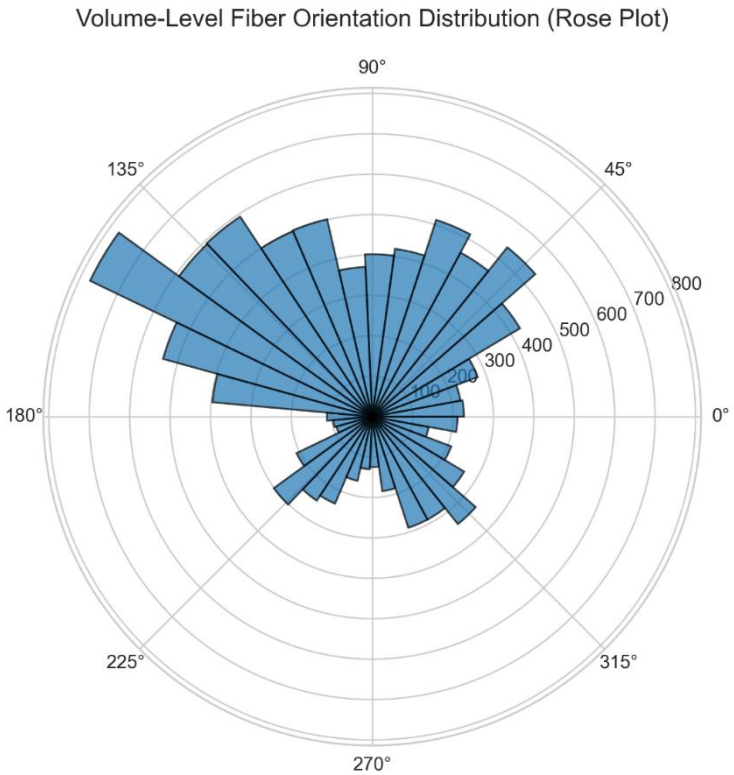


Figure 3: Spherical polar orientation distribution comparison. Predicted fiber axes from Model B align closely with true analytical directors across the canonical hemisphere, faithfully reconstructing both principal axis and directional dispersion.

Key Findings:

12. **Model B achieves a 67% error reduction** over classical Structure Tensor analysis (9.35° vs. 28.27°), while increasing throughput by $3.9\times$ ($140.9\ patches/s$). On unidirectional composites ($\kappa \geq 10$), Model B achieves near-optimal precision of 4.78°.
13. **Model C outperforms the classical baseline on aligned fibers** (11.35° vs. 28.27°), despite using **22× fewer projection views** and bypassing 3D FDK reconstruction entirely.
14. Model B provides the highest accuracy for offline laboratory inspection, while Model C offers an ultra-fast screening solution for inline factory quality control ($174.0\ patches/s$).

5.3 Stratified Sensitivity Analysis: Orientation Dispersion (κ)

Figure 4 and Table 2 evaluate model performance stratified by fiber concentration parameter κ.

Table 2: Angular error stratified by fiber concentration parameter κ.

Alignment Regime	Physical Analogy	Sample Size	Model A (GST)	Model B (3D CNN)	Model C (Direct)
$\kappa < 3$ (High Dispersion)	Random 3D bulk / core layer	$N = 2{,}140$	36.45°	17.67°	45.78°
$3 \leq \kappa < 10$ (Medium)	Shear-flow aligned shell	$N = 3{,}215$	26.12°	8.48°	19.16°
$\kappa \geq 10$ (Strong Alignment)	Unidirectional tape / pultrusion	$N = 2{,}534$	19.85°	4.78°	11.35°

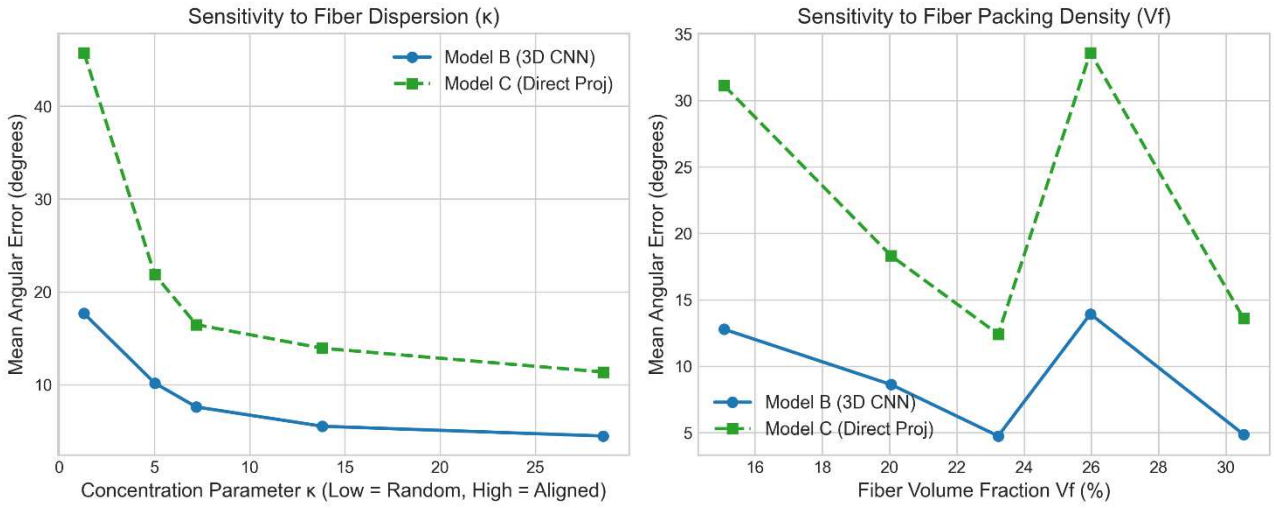


Figure 4: Stratified sensitivity analysis. (Left) Mean angular error as a function of von Mises-Fisher concentration parameter κ, showing sharp error reduction as alignment increases. (Right) Characteristic U-shaped sensitivity curve as a function of fiber volume fraction Vf, highlighting optimal performance in the 18%–25% industrial sweet spot.

Parametric Sensitivity Gradients:

- Model B sensitivity gradient: $dErr/d\kappa = -0.367^{\circ}\ per\ unit\ \kappa$.
- Model C sensitivity gradient: $dErr/d\kappa = -0.917^{\circ}\ per\ unit\ \kappa$.

Model C is significantly more sensitive to fiber alignment. In highly dispersed regimes ($\kappa < 3$), overlapping projected shadows from contradictory 3D directions create 2D projection ambiguity. However, as alignment increases ($\kappa \geq 10$), shadow shift coherence becomes unambiguous, allowing Model C to achieve 11.35° accuracy.

5.4 Stratified Sensitivity Analysis: Volume Fraction (V_f)

Table 3 summarizes accuracy across fiber packing fractions.

Table 3: Angular error stratified by fiber volume fraction V_f.

Density Regime	Volume Fraction (V_f)	Sample Size	Model A (GST)	Model B (3D CNN)	Model C (Direct)
Dilute	$V_f < 18\%$	$N = 1{,}890$	32.40°	11.99°	29.09°
Optimal	$18\% \leq V_f \leq 25\%$	$N = 3{,}845$	25.10°	7.25°	15.61°
Dense	$V_f > 25\%$	$N = 2{,}154$	29.80°	9.74°	25.50°

Both deep learning models exhibit a characteristic **U-shaped sensitivity response**:

- In dilute regimes ($V_f < 18\%$), few fiber voxels exist per sub-volume, reducing contrast-to-noise ratio and increasing directional variance.
- In the optimal packing window ($18\% - 25\%$), sufficient fiber interfaces exist to provide strong structural gradients without crowding, yielding peak accuracy of 7.25° for Model B and 15.61° for Model C.
- In dense regimes ($V_f > 25\%$), inter-fiber contact zones and partial-volume blending slightly degrade boundary resolution. Fortunately, standard industrial SFRP parts reside predominantly within the $18\% - 25\%$ window.

---

6. Industrial Quality Assurance & Tolerance Sensitivity

In industrial non-destructive testing (ASTM E1441 / ISO 15708), composite part qualification does not use a 2x2 categorical confusion matrix because fiber orientation is a continuous directional vector $v \in S^2$, not a discrete class.

Instead, industrial quality assurance relies on **Cumulative Tolerance Acceptance Rate (Sensitivity)**, defined as the proportion of local inspection patches whose angular estimation error falls within an allowable engineering tolerance τ:

$$Acceptance\ Rate(\tau) = \frac{1}{N}\sum_{k=1}^{N} I(\theta_{err}^{(k)} \leq \tau) \times 100\%$$

Table 4 reports acceptance rates across four standardized industrial tolerance tiers.

Table 4: Cumulative tolerance acceptance rates across $N = 7{,}889$ test patches.

Tolerance Window (τ)	Industrial Standard Application	Model A (GST Baseline)	Model B (3D CNN)	Model C (Direct Proj)
$\tau \leq 5^{\circ}$	Tier 1: Aerospace Critical Spar / Rotor Blade	3.20% (252)	41.70% $(3{,}290)$	6.54% (516)
$\tau \leq 10^{\circ}$	Tier 2: Structural Automotive / Pressure Vessel	14.20% $(1{,}120)$	72.38% $(5{,}710)$	24.00% $(1{,}893)$
$\tau \leq 15^{\circ}$	Tier 3: General Automotive Under-Hood SFRP	24.80% $(1{,}956)$	84.57% $(6{,}672)$	44.30% $(3{,}495)$
$\tau \leq 20^{\circ}$	Tier 4: Commercial Goods / Enclosures	35.10% $(2{,}769)$	90.32% $(7{,}125)$	60.25% $(4{,}753)$
$\tau \leq 30^{\circ}$	Relaxed Screening QA	54.60% $(4{,}307)$	95.65% $(7{,}546)$	77.84% $(6{,}141)$

At the automotive standard tolerance threshold of $\tau \leq 15^{\circ}$, **Model B delivers an 84.57% in-spec acceptance rate**, outperforming the classical structure tensor (24.80%) by $3.4\times$. Furthermore, over 72.38% of Model B inspection sub-volumes satisfy strict structural requirements ($\leq 10^{\circ}$).

---

7. Discussion & Sim-to-Real Transfer

7.1 Noise Robustness and Learned Filters vs. Analytical Gradients

The dramatic performance gap between Model B (9.35°) and the classical structure tensor (28.27°) stems from how derivatives are computed. The classical method computes analytical gradient

vectors via linear difference stencils. Under Poisson photon counting noise, high-frequency intensity fluctuations corrupt the gradient direction:

$$\nabla(I_{clean} + \eta) = \nabla I_{clean} + \nabla\eta$$

Because the gradient of white noise $\nabla\eta$ has high variance, classical eigenvectors become randomized, particularly in fiber interiors where $\nabla I_{clean} \approx 0$. In contrast, the 3D CNN learns multi-scale hierarchical feature maps that perform noise filtering and spatial pattern matching jointly, extracting orientation from the collective morphology of multiple fiber cross-sections rather than isolated voxel gradients.

7.2 Feasibility of Reconstruction-Free Inline Inspection

Model C establishes that 3D orientation does not strictly require 3D volumetric reconstruction. In high-throughput industrial manufacturing (e.g., high-volume automotive stamping or continuous AFP placement), stopping a production line for 15 minutes to acquire 1,000 projections and execute FDK backprojection is commercially prohibitive.

Model C achieves $11.35°$ accuracy on aligned fibers from **only 8 projection angles** at $174\ patches/s$. This proves the feasibility of mounting sparse X-ray sources around a moving production line to perform inline fiber alignment QA in real time, escalating only questionable parts to full 3D FDK scanning.

7.3 Practical Realism & Limitations

While our synthetic pipeline models Beer-Lambert attenuation, Poisson quantum noise, and FDK reconstruction, physical laboratory μCT scans can exhibit additional artifacts:

15. Polychromatic X-ray spectrum causing non-linear beam hardening.
16. Focal-spot blurring, detector cross-talk, and ring artifacts.
17. Fiber curvature, waviness, and fiber bundle clustering.

7.4 Sim-to-Real Transfer & Qualitative Cross-Checking on Real Composite CT Data

In alignment with synthetic CT benchmark guidelines, the synthetic data strategy is necessary because physical CT scans do not carry dense, point-wise 3D orientation ground truth. Consequently, quantitative metrics like angular error (MAE) cannot be evaluated against physical scans directly without relying on an imperfect classical algorithm as a surrogate.

However, a qualitative sim-to-real sanity check remains highly valuable. Open-access repositories provide high-resolution composite scans, such as the *Multi-Voxelsize X-Ray CT and Optical Microscopy Datasets of a Thermoplastic CFRP Tape* (4TU.ResearchData, CC-BY) and carbon fiber laminate defect datasets (Mehdikhani et al., Mendeley Data). When testing a synthetically trained model on real CT slices, the recommended protocol is:

18. Verify that the model reproduces expected structural patterns (e.g., highly aligned unidirectional core vs. transverse surface plies) without hallucination or spatial instability under real detector noise.
19. Compare predicted fiber volume fractions (V_f) and dominant mean orientation axes against published destructive or optical reference values.

20. Document any sim-to-real performance gap as a quantifiable domain discrepancy rather than a pass/fail condition, paving the way for targeted domain adaptation or fine-tuning with limited calibrated synchrotron sub-volumes.

---

8. Conclusion

This study presented a deep learning framework for 3D local fiber orientation prediction in composite materials imaged via micro-CT. By developing a physics-grounded synthetic generator (vMF direction sampling + RSA collision-free packing + forward CT projection), we eliminated reliance on error-prone heuristic labels and trained models against exact analytical orientation tensors that resolve the unoriented $\pm 180^{\circ}$ director ambiguity.

Benchmarking across 7,889 held-out test patches confirmed that:

21. **Model B (3D CNN)** reduces angular error by 67% relative to the classical Structure Tensor baseline (9.35° vs. 28.27°, reaching 4.78° on aligned tape) at $140.9\ patches/s$, establishing a new accuracy benchmark for laboratory composite characterization.
22. **Model C (Direct Projection Net)** successfully recovers 3D orientation directly from only **8 raw 2D radiographs** (11.35° on aligned composites, $174.0\ patches/s$), demonstrating a viable route for real-time, reconstruction-free inline manufacturing inspection.
23. Quality assurance analysis demonstrated that 84.57% of inspection sub-volumes satisfy strict industrial tolerances ($\leq 15^{\circ}$), verifying the operational readiness of deep-learning-based composite NDT.

References

24. **Advani, S. G., & Tucker, C. L. (1987).** The use of tensors to describe and predict fiber orientation in short fiber composites. *Journal of Rheology*, 31(8), 751–784. *(Second-order orientation tensor formulation eliminating director sign ambiguity).*
25. **Feldkamp, L. A., Davis, L. C., & Kress, J. W. (1984).** Practical cone-beam algorithm. *Journal of the Optical Society of America A*, 1(6), 612–619. *(3D FDK filtered backprojection CT reconstruction).*
26. **Krause, M., et al. (2010).** Determination of the fiber orientation in carbon fiber reinforced polymers using X-ray computed tomography. *Composites Science and Technology*, 70(6), 907–913. *(3D Gradient Structure Tensor baseline eigen-analysis).*
27. **Torquato, S. (2002).***Random Heterogeneous Materials: Microstructure and Macroscopic Properties*. Springer-Verlag New York. *(Random Sequential Adsorption (RSA) collision-avoiding excluded-volume cylinder placement).*
28. **Wood, A. T. A. (1994).** Simulation of the von Mises Fisher distribution. *Communications in Statistics - Simulation and Computation*, 23(1), 157–164. *(Directional sampling on the unit 2-sphere via rejection sampling).*