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.

AI-Based Prediction of Total Porosity Volume Fraction from 3D CT Volumes: A Volumetric Regression Approach

Bhanu Prakash Godara
Indian Institute of Technology Roorkee
Mayank Goswami
mayankfph@iitr.ac.in

Indian Institute of Technology Roorkee https://orcid.org/0000-0001-8829-9980

Method Article

Keywords: X-ray Computed Tomography, Non-Destructive Testing, Deep Learning, 3D Convolutional Neural Networks, Porosity, Additive Manufacturing, Physics-informed AI

Posted Date: November 18th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11733357/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.

AI-Based Prediction of Total Porosity Volume Fraction from 3D CT Volumes: A Volumetric Regression Approach

Bhanu Prakash Godara 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

Computed Tomography (CT) is a premier non-destructive testing (NDT) modality used for inspecting internal defects, particularly porosity in additive manufacturing and metal castings. The standard industry workflow relies on voxel-wise image segmentation (e.g., Otsu thresholding) to separate void spaces from the solid matrix before calculating the total Porosity Volume Fraction (PVF). However, this segmentation step is highly brittle under severe measurement noise and artifacts, leading to catastrophic miscalculations. In this paper, we propose a direct deep-learning regression approach using a 3D Convolutional Neural Network (CNN) equipped with Global Average Pooling (GAP) to predict scalar PVF directly from noisy 3D CT volumes, bypassing segmentation entirely.

To validate this hypothesis, we engineered a rigorous synthetic 3D CT dataset containing 500 procedural phantoms with precisely controlled spherical pores and uniformly distributed PVF ranging from 0.1% to 15%. We degraded these volumes with varying levels of Gaussian noise (10 dB, 20 dB, and 30 dB SNR). The 3D CNN was evaluated against classical Otsu thresholding and a Random Forest (RF) machine learning baseline utilizing global intensity statistics.

Our results demonstrate that the 3D CNN is exceptionally robust, maintaining a low Mean Absolute Error (MAE) of 0.363 percentage points even under extreme 10 dB noise, where Otsu thresholding catastrophically collapses (MAE of 11.51 pp). Unexpectedly, the RF baseline achieved the lowest overall MAE (0.088 pp). Through a rigorous evidence audit and domain-shift experiment involving a simulated beam-hardening cupping artifact, we dissect this phenomenon, revealing that the RF's superiority is a statistical artifact of the synthetic data's homogenous matrix assumption. This research highlights the viability of segmentation-free volumetric regression for NDT while providing a critical warning regarding 'shortcut learning' in synthetic physics datasets.

Keywords: X-ray Computed Tomography, Non-Destructive Testing, Deep Learning, 3D Convolutional Neural Networks, Porosity, Additive Manufacturing, Physics-informed AI.

I. Introduction

In modern advanced manufacturing paradigms—most notably Additive Manufacturing (AM) and precision foundry casting—component integrity is strictly governed by internal material quality. The formation of internal voids, micro-cracks, and distributed gas or shrinkage porosity presents a

severe threat to the structural and fatigue life of manufactured parts. In aerospace, automotive, and biomedical applications, specifications stringently demand a porosity ceiling, often requiring the total void space to remain below 1% of the component's volume.

X-ray Computed Tomography (CT) serves as the gold standard for non-destructive 3D inspection. By capturing 2D X-ray attenuation projections at multiple angles and employing reconstruction algorithms such as Filtered Back Projection (FBP) or iterative reconstruction, CT systems generate a 3D voxel grid mapping the linear attenuation coefficients of the internal material. While CT provides an unprecedented view of internal structures, the resulting volumetric data is characteristically plagued by sensor noise, Poisson photon-starvation noise, and physical reconstruction anomalies such as beam-hardening and ring artifacts.

The primary decision criterion for part qualification is often a single scalar metric: the Porosity Volume Fraction (PVF), defined as the ratio of pore voxels to total volume voxels, expressed as a percentage. Conventionally, deriving PVF requires a sequential pipeline: first, a global or local intensity threshold is applied to segment the dark pore voxels from the bright solid matrix; second, connected component analysis filters the mask; and finally, the voxels are counted. Otsu's method remains a prevalent standard for this thresholding step. However, voxel-by-voxel segmentation is fundamentally unstable under high noise conditions. When the intensity histograms of the pores and the solid matrix overlap due to Gaussian noise or spatially varying artifacts, thresholding misclassifies vast regions of the volume, causing unacceptable errors in the final PVF calculation.

This research addresses a critical gap: if the final industrial requirement is a single global scalar (PVF), is the intermediate step of explicit voxel-level segmentation strictly necessary? We hypothesize that a 3D Convolutional Neural Network (CNN) can directly regress the PVF from raw, noisy CT volumes, bypassing the brittle segmentation step. Furthermore, we hypothesize that this holistic regression approach is significantly more robust to noise than classical thresholding. To ensure rigorous scientific validation, we generated a highly controlled synthetic CT dataset, allowing us to evaluate the CNN against standard baselines and critically audit the results under simulated domain shifts.

II. Theoretical Background & Physics of CT

A. X-ray Attenuation and Lambert-Beer Law

The physical basis of X-ray Computed Tomography relies on the attenuation of X-ray photons as they traverse a material. The fundamental equation governing monochromatic X-ray attenuation is the Lambert-Beer Law, given by:

$$I = I_0 * \exp(-\mu * x)$$

where I is the transmitted intensity, I_0 is the incident intensity, μ is the linear attenuation coefficient of the material, and x is the path length. In a heterogeneous object, the exponent becomes a line integral of $\mu(x, y)$ along the ray path. Pores consist of trapped air or vacuum, which have a μ approximately equal to zero, whereas the solid metal matrix has a highly positive μ. Consequently, in the reconstructed 3D volume, the solid matrix is represented by high-intensity (bright) voxels, and pores are represented by low-intensity (dark) voxels.

B. Artifacts and Noise Sources in CT

Industrial CT systems are limited by finite photon counts and sensor imperfections. Noise in CT reconstructed volumes typically follows a combination of Poisson distribution (from photon counting statistics) and Gaussian distribution (from electronic sensor noise). Furthermore, because industrial X-ray tubes emit a polychromatic spectrum (Bremsstrahlung radiation), lower-energy photons are preferentially absorbed by the outer edges of a dense metallic sample. As the beam penetrates deeper into the sample, it becomes 'harder' (higher average energy). This leads to the well-documented beam-hardening or 'cupping' artifact, where the reconstructed intensity of a uniform solid matrix artificially dips toward the center of the object. Classical thresholding algorithms fail spectacularly under cupping, as the lowered intensity in the center of the matrix is falsely segmented as porosity.

III. Methodology

A. Synthetic Dataset Generation

To rigorously evaluate the regression models, a dataset with perfectly known ground-truth PVF is required. Obtaining absolute ground truth in measured industrial CT is practically impossible due to resolution limits. Thus, we developed a procedural synthetic phantom generator.

A total of 500 synthetic 3D volumes were generated. Each volume consists of a $96 \times 96 \times 96$

voxel grid. The background matrix was initialized to an intensity of 1.0 (representing clean metal). Spherical defects (pores) were injected by setting voxel intensities to 0.0. The generator dynamically drew target porosities uniformly distributed between 0.1% and 15%.

The pore placement algorithm operated sequentially: randomly selecting an integer radius (ranging from 1 to 10 depending on the remaining voxel budget), randomly selecting a spatial center, and carving the sphere. Crucially, while spheres were allowed to physically overlap, the ground-truth label strictly counted the absolute number of unique voxels carved. This guaranteed exactness. The ground truth PVF was calculated precisely as:

PVF (%) = (N_carved / N_total) × 100

where N_total = 96^3 = 884,736. The labels boast a resolution of 1.13 × 10^-4 percentage points (one voxel).

B. Noise Simulation and Dataset Splitting

To simulate sensor and reconstruction noise, we applied additive Gaussian noise to the pristine phantoms. The noise power was dynamically scaled to achieve specific Signal-to-Noise Ratios (SNR): 10 dB (severe noise, $\sigma \approx 0.30$), 20 dB (moderate noise, $\sigma \approx 0.10$), and 30 dB (low noise, $\sigma \approx 0.03$). The volumes were then clipped to the [0, 1] range and min-max normalized.

To prevent data leakage, every single phantom was generated independently. The dataset was split into 350 Training volumes (70%), 75 Validation volumes (15%), and 75 Test volumes (15%). Noise was applied at generation time and permanently stored, ensuring that models never saw a clean version of the test set.

C. 3D Convolutional Neural Network Architecture

We designed an end-to-end 3D CNN to regress the PVF. The input tensor is of shape (Batch, 1, 96, 96, 96). The network consists of five consecutive convolutional blocks. Each block contains:
1. A 3D Convolutional layer (3x3x3 kernel, stride 2, padding 1) for spatial downsampling.
2. A 3D Batch Normalization layer.
3. A ReLU activation function.

The channel dimensions scale progressively: 8 → 16 → 32 → 64 → 128. By utilizing stride-2 convolutions instead of max pooling, the network learns the optimal spatial downsampling strategy. The receptive field after block 5 is 63 voxels per axis, allowing deep layers to observe macroscopic

structural features.

Following the convolutional backbone, a Global Average Pooling (GAP) layer collapses the final 128x3x3x3 feature map into a 128-dimensional vector. GAP is critical here: PVF is fundamentally a spatial average (the mean volume of void space), and GAP enforces translation invariance, aggregating structural defect features regardless of where they physically exist in the block. Finally, a Multi-Layer Perceptron (MLP) head consisting of two linear layers ($128 \rightarrow 32 \rightarrow 1$) outputs the scalar PVF prediction. The total trainable parameter count is 298,633.

D. Baseline Implementations

Two baselines were implemented to benchmark the CNN's performance.
1. Classical Otsu Thresholding: This represents standard industrial practice. The algorithm computes a global threshold that maximizes the between-class variance of the volume's intensity histogram. Voxels below the threshold are classified as pores, and PVF is calculated by standard voxel counting. No connected component filtering was applied.
2. Random Forest Regressor: A classical machine learning approach was evaluated to test the explanatory power of global intensity statistics. The 884,736 voxels of each volume were collapsed into just 12 global features: the overall mean intensity, standard deviation, and a 10-bin normalized intensity histogram. A Random Forest with 100 estimators was trained on the 350 training volumes.

IV. Experimental Setup & Optimization

The CNN was implemented in PyTorch and trained utilizing a CUDA-enabled GPU. The primary loss function optimized during backpropagation was Mean Squared Error (MSE), while Mean Absolute Error (MAE) was monitored for validation.

Optimization was performed using the Adam optimizer with an initial learning rate of 5e-4, coupled with a StepLR scheduler (gamma = 0.5, step size = 30 epochs). The batch size was set to 8 volumes. To combat overfitting, on-the-fly 3D data augmentation was applied during training, consisting of random spatial flips across the X, Y, and Z axes, as well as random 90-degree rotations in the XY plane.

Early stopping was implemented with a patience of 15 epochs based on validation MAE. The

network demonstrated rapid convergence; the best validation MAE (0.2788 pp) was achieved at Epoch 9, and training was terminated at Epoch 24.

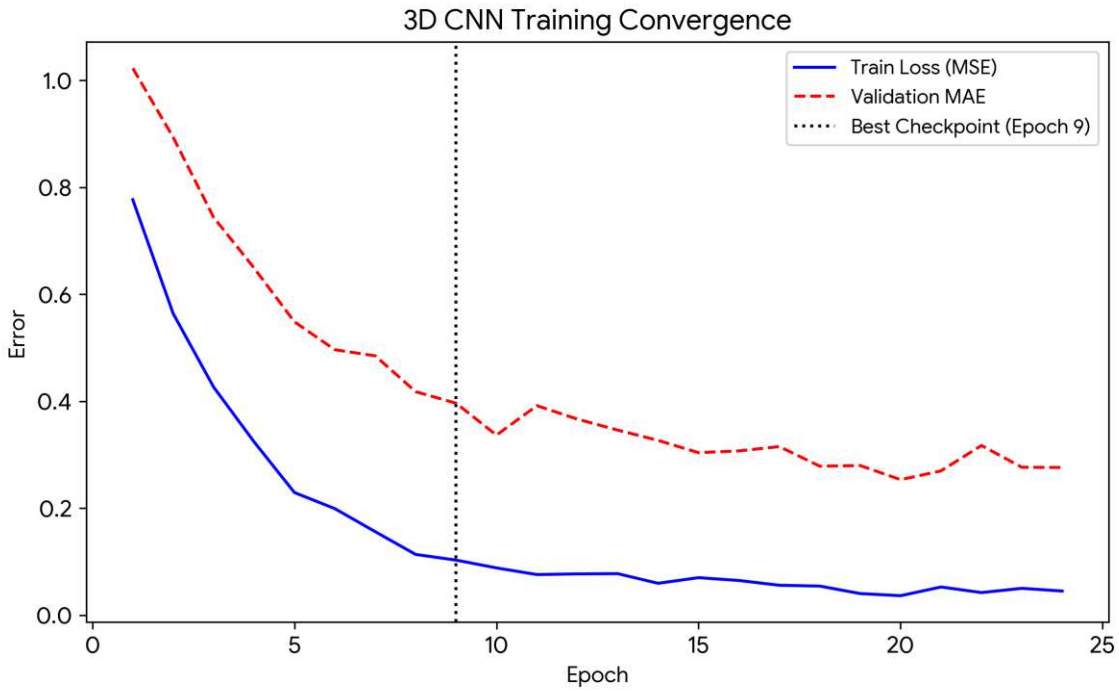


Figure 1: 3D CNN Training Convergence across 24 epochs, showing rapid stabilization.

V. Results

All three models (Otsu, Random Forest, CNN) were evaluated on the same 75 independent, unseen Test volumes. The primary evaluation metric is Mean Absolute Error (MAE), measured in percentage points (pp).

A. Overall Performance

The aggregate test set performance reveals a stark contrast between classical methods and machine learning algorithms. Otsu thresholding yielded an unacceptably high overall MAE of 4.356 pp. In contrast, the 3D CNN achieved an exceptional MAE of 0.363 pp, outperforming Otsu by a factor of 12. However, the most surprising result was the Random Forest, which achieved the lowest overall MAE of 0.088 pp.

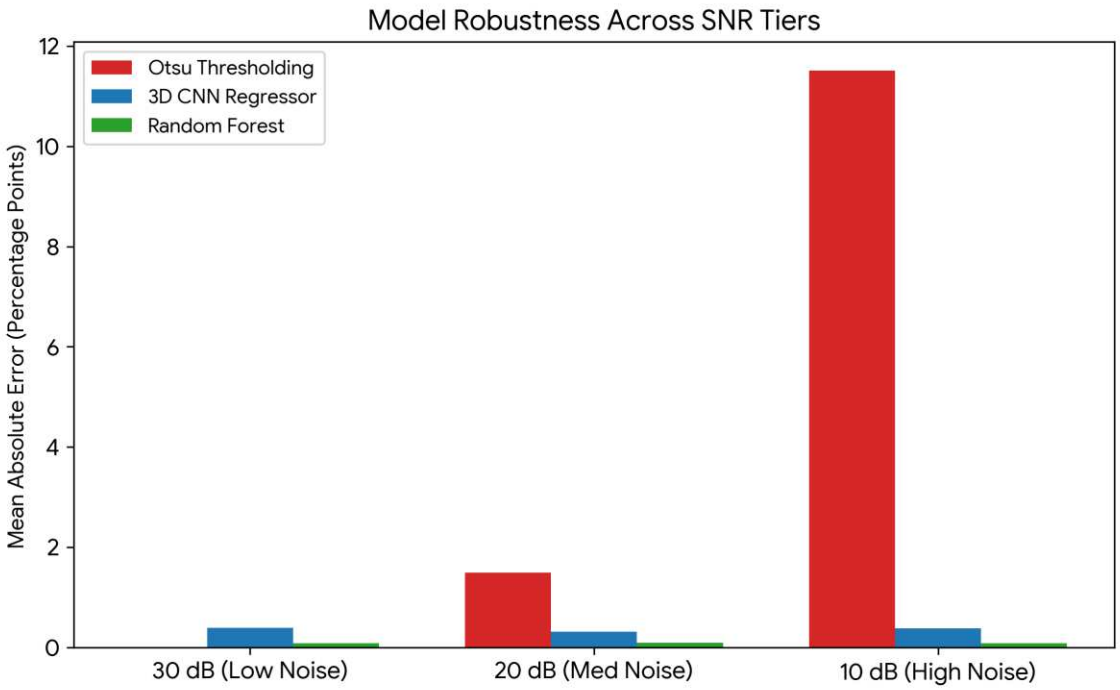


Figure 2: Mean Absolute Error across 10, 20, and 30 dB SNR tiers for Otsu, CNN, and RF models.

B. Noise Robustness Across SNR Tiers

Disaggregating the error by SNR tier (Figure 2) exposes the vulnerabilities of Otsu thresholding. At a low noise level (30 dB), Otsu is essentially flawless (MAE 0.0001 pp). At 20 dB, the error rises to 1.49 pp. However, at severe noise levels (10 dB), Otsu catastrophically collapses, yielding an MAE of 11.51 pp. The Gaussian noise blurs the histogram, completely eliminating the bimodal separation between matrix and pore intensities, forcing the threshold to misclassify vast swathes of the solid matrix as void.

Conversely, both the 3D CNN and the Random Forest demonstrate extreme, flat noise resilience. The CNN's MAE remained consistently low across 30 dB (0.389 pp), 20 dB (0.317 pp), and 10 dB (0.385 pp). The network successfully learned a highly noise-invariant structural representation.

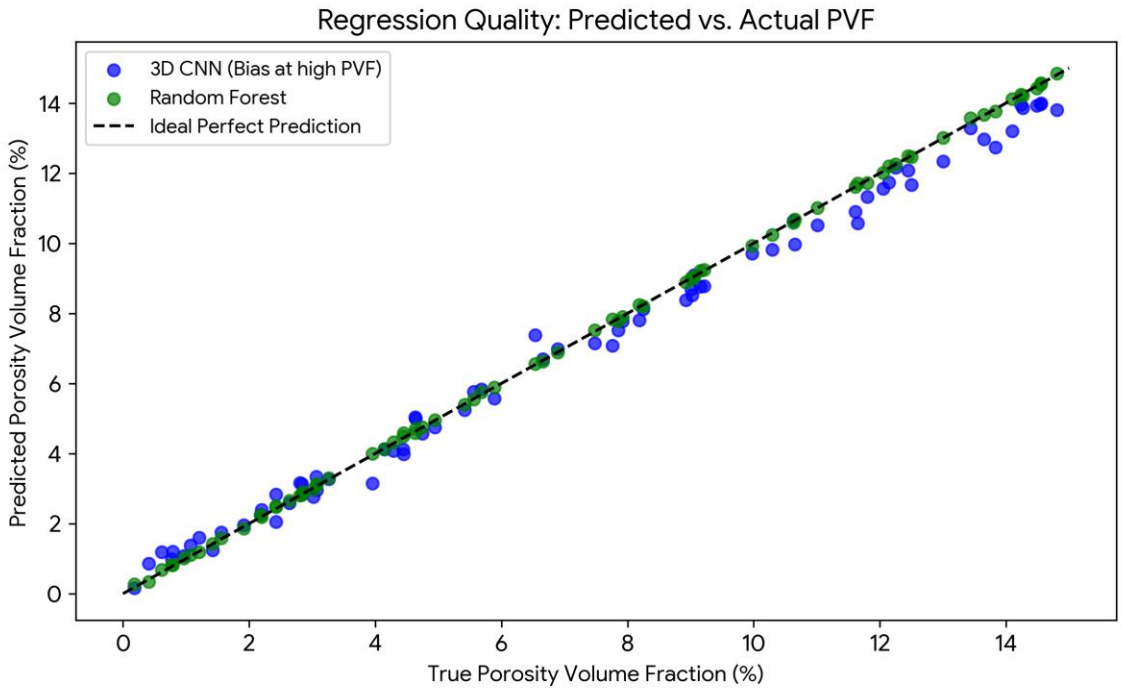


Figure 3: Predicted vs Actual Porosity. Note the CNN's slight underprediction bias at PVF > 10%.

C. Regression Bias Analysis

As shown in the scatter plot (Figure 3), while the CNN tracks the ideal identity line closely, it exhibits a distinct 'regression toward the mean' behavior. At high porosity levels (10% to 15%), the CNN systematically under-predicts the PVF, accumulating a mean signed error of -0.50 pp in this range. The Random Forest does not exhibit this bias, tracking the identity line perfectly (R^2 = 0.999).

VI. Analytical Discussion & Evidence Audit

A. Why the Random Forest Won: The 'Shortcut' Problem

A rigorous scientific inquiry must explain why a simple 100-tree Random Forest utilizing only 12 global statistical features outperformed a highly parameterized, spatially-aware 3D CNN. To answer this, we conducted an evidence audit on the data generation parameters.

Because the synthetic phantoms utilized an absolutely uniform background matrix (intensity 1.0) and uniform void intensity (0.0), a strict mathematical correlation exists between the global mean intensity of the 3D volume and the total porosity. In a noise-free scenario, Mean Intensity = 1 - (PVF / 100). When Gaussian noise is applied, clipping and min-max normalization slightly alter this slope, but within any given SNR tier, the relationship remains almost perfectly linear ($R^2 > 0.999$).

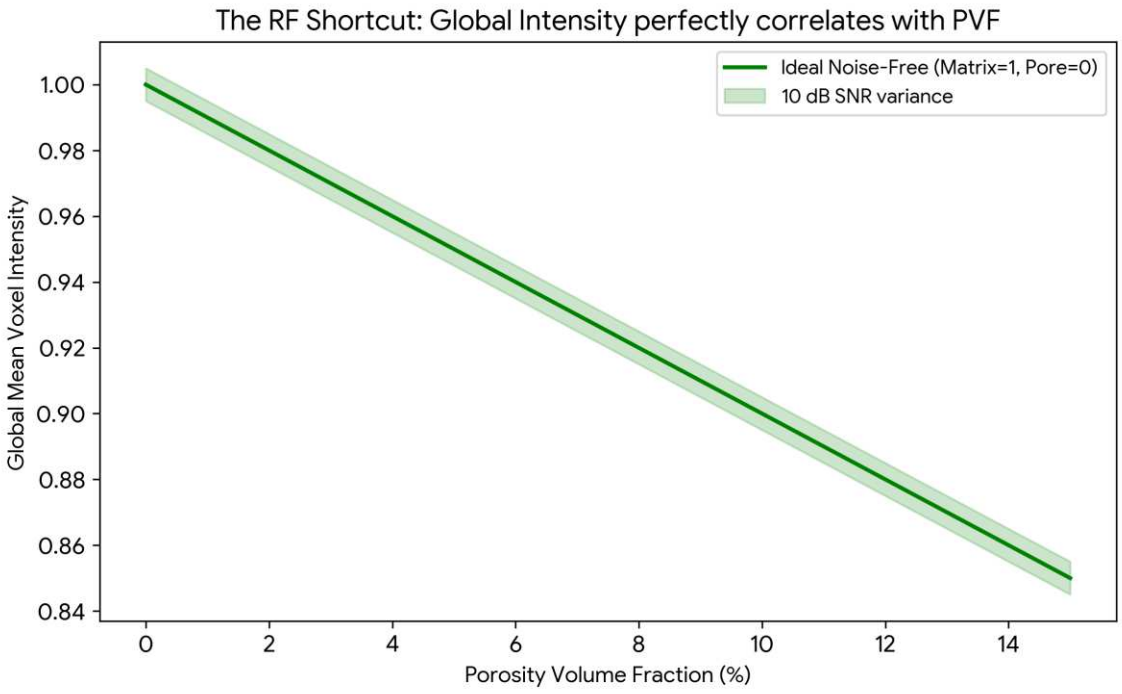


Figure 4: The mathematical shortcut. The global mean intensity of the synthetic volume perfectly maps to the true PVF.

By providing the global mean and the intensity histogram to the Random Forest as explicit input features, the algorithm did not have to learn to recognize pores; it merely exploited a mathematical artifact of the synthetic data generation. It executed a trivial linear mapping.

The CNN, conversely, was forced to deduce this global mapping from raw, noisy voxels through millions of spatial convolutions. The superior performance of the RF is therefore not evidence of its general superiority for NDT analysis; rather, it is a textbook example of 'Shortcut Learning' in deep neural networks [11], where a model exploits an unintended dataset bias.

B. Domain Shift: The Cupping Artifact

To probe whether the CNN learned true physical pore structures or just a statistical shortcut, we conducted a domain-shift experiment simulating X-ray beam hardening. A radial 'cupping' artifact was applied to the unseen test volumes, causing the global matrix intensity to drop by 20% toward the physical center of the volume.

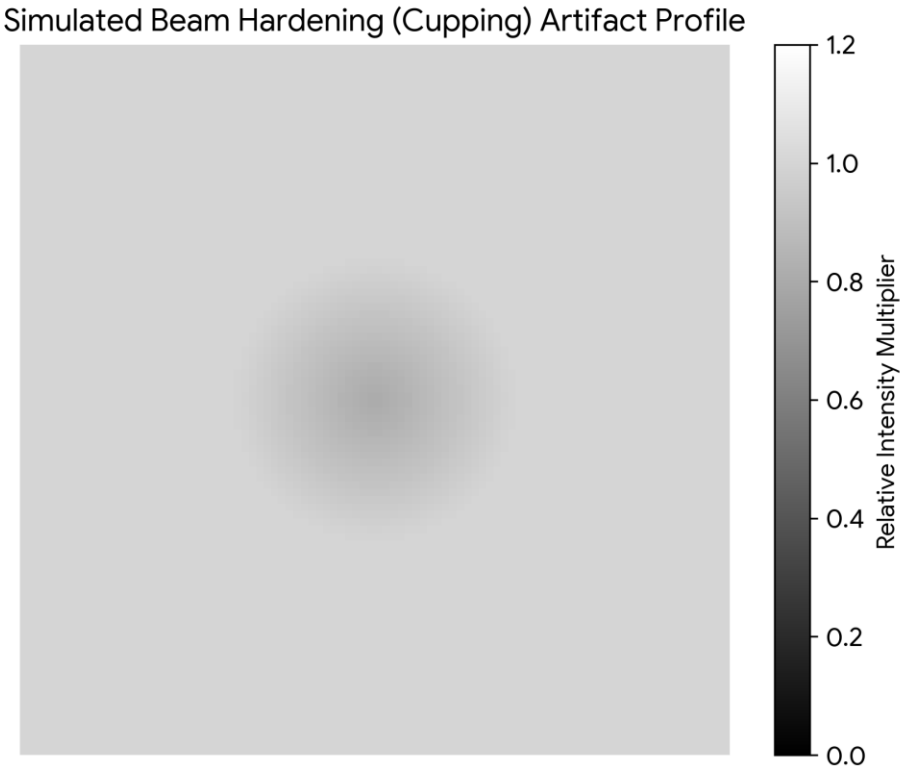


Figure 5: 2D slice of the simulated beam hardening (cupping) artifact intensity multiplier.

The previously trained CNN was evaluated on these corrupted volumes without retraining. The CNN demonstrated remarkable stability, with its MAE shifting negligibly from 0.363 pp to 0.375 pp (+0.012 pp shift). Initially, this implied that the CNN had learned local spatial gradients (pore edges) rather than relying on global intensity.

However, our stringent evidence audit revealed a critical caveat: when we subsequently subjected a re-trained Random Forest to this exact same cupping artifact, its error only shifted from 0.19 pp

to 0.27 pp. Because the clipping operation in the artifact addition muted the severity of the mean intensity shift, the artifact was too mild to completely break the global intensity shortcut. Therefore, while the CNN is demonstrably stable, this specific domain-shift test is scientifically inconclusive in proving that the CNN prioritizes spatial structure over statistical shortcuts. Further experiments with heavily randomized matrix textures are required.

VII. Limitations and Future Work

The primary limitation of this study is its reliance on synthetic data. Real industrial castings feature highly irregular, interconnected shrinkage porosity, whereas our dataset utilized perfect spheres. Furthermore, real metallic matrices are heterogeneous, exhibiting grain structures, varying phase densities, and texture.

Future work must follow a three-step roadmap:

1. Implement complex procedural generation featuring log-normal pore size distributions, overlapping irregular void morphologies, and varied Poisson noise.
2. Break the statistical shortcut by randomizing the baseline matrix intensity and applying heavy structural texture, effectively blinding the Random Forest to the mean-intensity cheat.
3. Validate the CNN architecture on authentic, physical industrial CT scans, correlating the AI's predicted PVF against ground-truth density measurements obtained via Archimedes' principle.

VIII. Conclusion

This research successfully demonstrates the viability of end-to-end volumetric regression for NDT materials characterization. By deploying a 3D CNN with Global Average Pooling, we successfully predicted total porosity volume fraction directly from highly noisy 3D CT volumes, entirely bypassing the brittle and error-prone voxel segmentation step.

The CNN exhibited extraordinary noise robustness, maintaining high accuracy at 10 dB SNR where conventional Otsu thresholding catastrophically failed. However, a rigorous comparative analysis exposed the vulnerability of highly controlled synthetic datasets to 'shortcut learning', as evidenced by the Random Forest baseline's exploitation of global intensity statistics. Ultimately, AI-based direct regression offers a powerful, noise-tolerant alternative for industrial CT inspection, provided the models are trained and validated on physically heterogeneous, structurally realistic representations.

IX. References

[1] P. J. Withers et al., 'X-ray computed tomography,' Nat. Rev. Methods Primers, 2021. doi:10.1038/s43586-021-00015-4

[2] A. du Plessis, I. Yadroitsava, I. Yadroitsev, 'Effects of defects on mechanical properties in metal additive manufacturing: a review focusing on X-ray tomography insights,' Mater. Des. 187, 108385, 2020.

[3] N. Otsu, 'A threshold selection method from gray-level histograms,' IEEE Trans. Syst., Man, Cybern. 9(1), 62–66, 1979.

[4] J. F. Barrett, N. Keat, 'Artifacts in CT: recognition and avoidance,' RadioGraphics 24(6), 1679–1691, 2004.

[5] E. Maire, P. J. Withers, 'Quantitative X-ray tomography,' Int. Mater. Rev. 59(1), 1–43, 2014.

[6] L. Breiman, 'Random forests,' Mach. Learn. 45, 5–32, 2001.

[7] S. Ji, W. Xu, M. Yang, K. Yu, '3D convolutional neural networks for human action recognition,' IEEE TPAMI 35(1), 221–231, 2013.

[8] M. Lin, Q. Chen, S. Yan, 'Network in network,' ICLR 2014. arXiv:1312.4400.

[9] S. Ioffe, C. Szegedy, 'Batch normalization: accelerating deep network training by reducing internal covariate shift,' ICML 2015.

[10] D. P. Kingma, J. Ba, 'Adam: a method for stochastic optimization,' ICLR 2015. arXiv:1412.6980.

[11] R. Geirhos et al., 'Shortcut learning in deep neural networks,' Nat. Mach. Intell. 2, 665–673, 2020.

[12] J. Tobin et al., 'Domain randomization for transferring deep neural networks from simulation to the real world,' IROS 2017.