DOI: 10.65322/j.srd.05192026.15

PHYSICS-INFORMED HYBRID ALGORITHMS FOR SOLVING INVERSE PROBLEMS IN LIMITED-DATA AND HIGH-NOISE X-RAY TOMOGRAPHY

Allamuratova Nilufar Kuat qizi,
E-mail: ferouduke@gmail.com,
Tashkent university of information technologies
named after Muhammad al-Khwarizmi

Abstract. *This article introduces a physics-informed hybrid algorithm framework designed to solve inverse problems in X-ray tomography under limited-data and high-noise constraints. The research highlights the mathematical fragility and artifact generation of both classical deterministic solvers and unconstrained deep learning models when faced with ultra-low-dose and sparse-view data. To overcome this, the authors developed a closed-loop hybrid pipeline that embeds physical Radon transformation models directly into the neural network, supported by an adaptive parameter-tuning engine that continuously adjusts regularization weights based on local noise estimation. Extensive empirical evaluations demonstrate that this approach effectively suppresses noise and hallucinated artifacts, maintaining high structural fidelity while allowing for up to an 80% reduction in diagnostic radiation doses.*

Keywords: *Physics-informed neural networks, Computed tomography reconstruction, Ill-posed inverse problems, Radon forward operator, Unrolled iterative optimization, Dynamic regularization weighting, Artifact suppression & hallucination elimination, Low-dose ct imaging.*

Introduction. Over the past four decades, X-ray Computed Tomography (CT) has evolved from a specialized diagnostic innovation into an indispensable cornerstone of modern clinical medicine and non-destructive industrial evaluation. In healthcare settings, cross-sectional CT imaging provides high-resolution anatomical visualization essential for acute trauma triage, oncology staging, and surgical planning. However, the diagnostic efficacy of X-ray CT is inherently counterbalanced by the biological risks associated with ionizing radiation. Epidemiological research has demonstrated that cumulative diagnostic radiation carries non-negligible stochastic risks, most notably radiation-induced carcinogenesis [1-2]. Consequently, clinical practice strictly adheres to the ALARA (*As Low As Reasonably Achievable*) protocol, driving a persistent demand for low-dose CT acquisition paradigms.

In parallel, industrial X-ray tomography faces analogous operational imperatives. Applications such as dynamic process monitoring, in-situ material stress testing, and high-throughput additive manufacturing inspection necessitate dramatic reductions in scan times [2]. In these industrial environments, full angular 360° sampling is frequently precluded by physical geometric constraints, component movement, or the thermal limits of X-ray tube source hardware.

Both medical dose-reduction mandates and industrial speed constraints necessitate a shift toward truncated acquisition protocols: specifically, reducing the tube current-time product or significantly limiting the number of projection angles (sparse-view and limited-angle tomography). However, physically restricting the acquisition process severely degrades the underlying mathematical integrity of the Radon transform inverse problem. Ultra-low-dose projections subject the detector measurements to photon starvation, shifting signal fluctuations from a Gaussian approximation toward a signal-dependent Poisson distribution heavily corrupted by electronic additive noise. Concurrently, sparse angular sampling severely violates the Nyquist-Shannon sampling theorem in projection space [4-6].

To resolve the mathematical fragility inherent in standard inversion solvers, this work introduces a physics-informed hybrid reconstruction framework that explicitly integrates deterministic physical laws into an adaptive artificial intelligence architecture. Purely data-driven

DOI: 10.65322/j.srd.05192026.15

neural networks, while capable of learning complex non-linear mappings, operate primarily as black-box interpolators. When deployed in limited-data, high-noise environments, unconstrained deep learning models frequently generate unphysical hallucinations or erase subtle low-contrast structures critical for diagnostic and material evaluation. Conversely, deterministic solvers guarantee physical consistency but lack the adaptability required to handle severe non-Gaussian noise distributions [13-16].

The proposed solution establishes a closed-loop hybrid pipeline that unifies the mathematical rigor of forward physical modeling with the feature-extraction capabilities of deep neural networks. Rather than treating image reconstruction as an unconstrained sinogram-to-image translation, the architecture explicitly embeds the discrete Radon forward operator A and its adjoint A^T directly into the network layers. The optimization process is governed by a multi-component physics-informed loss function that enforces data fidelity in sinogram space, spatial continuity in image space, and physical adherence to X-ray attenuation physics.

Mathematical foundations and problem formulation. The physical phenomenon underlying X-ray tomographic imaging is governed by the attenuation of photon intensity as a radiation beam penetrates a non-homogeneous material domain $\Omega \subset \mathbb{R}^2$ [2]. For a monochromatic X-ray beam with initial photon intensity I_0 traversing along a linear trajectory L, the transmitted intensity I is modeled according to the Beer-Lambert law of attenuation:

$$I = I_0 \exp\left(-\int_L f(x,y)dl\right) \tag{1}$$

where $f(x,y) \in L^2(\Omega)$ represents the spatial distribution of linear attenuation coefficients across the object domain [2].

To convert the non-linear exponential decay in Equation (1) into a linear mathematical formulation, a logarithmic transformation is applied to the ratio of incident intensity I_0 to transmitted intensity I. This operation isolates the line integral projection $g(L)$:

$$g(L) = \ln\left(\frac{I_0}{I}\right) = \int_L f(x,y)dl \tag{2}$$

In a two-dimensional parallel-beam acquisition geometry, a straight line path L is uniquely parametrized by its polar orientation angle $\theta \in [0,\pi)$ and its perpendicular distance $s \in \mathbb{R}$ from the coordinate origin. The mapping from the two-dimensional attenuation function $f(x,y)$ to the continuous projection space $g(\theta,s)$ defined in Equation (2) is formally expressed via the continuous Radon transform operator $\mathcal{R}$:

$$\mathcal{R}f(\theta,s) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(x,y)\delta(x\cos\theta + y\sin\theta - s)dxdy \tag{3}$$

where $\delta(\cdot)$ in Equation (3) represents the one-dimensional Dirac delta distribution that restricts integration exclusively along the line $x\cos\theta + y\sin\theta = s$.

For numerical computational implementation, the continuous mapping in Equation (3) must be discretized across finite spatial and projection dimensions. The continuous spatial function $f(x,y)$ is discretized into a vector $x = [x_1, x_2, \ldots, x_n]^T \in \mathbb{R}^n$, where each element x_j represents a constant attenuation value within a specific pixel. Concurrently, the continuous projection function $g(\theta,s)$ in Equation (3) is sampled across m discrete detector channels and angular steps, yielding a discrete projection vector $y = [y_1, y_2, \ldots, y_m]^T \in \mathbb{R}^m$. This spatial and angular discretization reduces the continuous forward operator in Equation (3) to a linear system equation:

$$y = Ax \tag{4}$$

DOI: 10.65322/j.srd.05192026.15

In Equation (4), $A \in \mathbb{R}^{m \times n}$ represents the projection system matrix, where each matrix coefficient A_{ij} specifies the intersection length or geometric weighting of the i-th X-ray beam passing through the j-th pixel element. Equation (4) serves as the deterministic forward operator governing ideal, noise-free tomographic data generation.

While Equation (4) models the noise-free forward problem, physical acquisition systems under reduced radiation dosage introduce severe measurement perturbations. The discrete observational model incorporating high-variance noise is expressed as:

$$y = Ax + \eta \tag{5}$$

where $y \in \mathbb{R}^m$ is the measured sinogram vector, $x \in \mathbb{R}^n$ is the discrete spatial attenuation vector, $A \in \mathbb{R}^{m \times n}$ is the discretized Radon forward matrix, and $\eta \in \mathbb{R}^m$ encapsulates stochastic noise components arising from photon starvation and electronic readout variance.

To characterize the mathematical instability of solving Equation (5), we analyze the singular value decomposition of the system matrix A:

$$A = U \Sigma V^T = \sum_{i=1}^{r} \sigma_i u_i v_i^T \tag{6}$$

where $U = [u_1, u_2, \ldots, u_m] \in \mathbb{R}^{m \times m}$ and $V = [v_1, v_2, \ldots, v_n] \in \mathbb{R}^{n \times n}$ are orthogonal matrices containing the left and right singular vectors, respectively, $r = \operatorname{rank}(A) \leq \min(m, n)$, and Σ contains the singular values ordered as $\sigma_1 \geq \sigma_2 \geq \ldots \geq \sigma_r > 0$.

Applying the unregularized Moore-Penrose pseudo-inverse $A^\dagger$ to Equation (5) yields the naive reconstruction $x^\dagger$:

$$x^\dagger = A^\dagger y = A^\dagger (Ax + \eta) = x + \sum_{i=1}^{r} \frac{u_i^T \eta}{\sigma_i} v_i \tag{7}$$

Equation (7) demonstrates that the reconstruction error vector $e = x^\dagger - x$ depends inverse-proportionally on the singular spectrum of A. The sensitivity of the inversion matrix to noise corruptions is formally established by the following mathematical statements.

Theorem 2.1 (Noise Amplification Bound). Let $A \in \mathbb{R}^{m \times n}$ be the projection matrix corresponding to an incomplete-data tomographic geometry with condition number $\kappa(A) = \sigma_1 / \sigma_r$. If the observed sinogram y in Equation (5) is corrupted by noise η bounded by $\| \eta \|_2 \leq \epsilon$, then the norm of the absolute reconstruction error $e = x^\dagger - x$ satisfies:

$$\| x^\dagger - x \|_2 \leq \frac{\epsilon}{\sigma_r} \tag{8}$$

and the relative reconstruction error bound is given by:

$$\frac{\| x^\dagger - x \|_2}{\| x \|_2} \leq \kappa(A) \frac{\| \eta \|_2}{\| y - \eta \|_2} \tag{9}$$

Proof. Subtracting the exact solution x from Equation (7) yields the reconstruction error $e = \sum_{i=1}^{r} \frac{u_i^T \eta}{\sigma_i} v_i$. Taking the Euclidean norm of both sides and applying the orthogonality of v_i

DOI: 10.65322/j.srd.05192026.15

yields $\| e \|_2^2 = \sum_{i=1}^{r} \frac{(u_i^T \eta)^2}{\sigma_i^2} \leq \frac{1}{\sigma_r^2} \sum_{i=1}^{r} (u_i^T \eta)^2 \leq \frac{\| \eta \|_2^2}{\sigma_r^2}$. Taking the square root confirms Equation (8). For Equation (9), noting that $\| y - \eta \|_2 = \| Ax \|_2 \leq \sigma_1 \| x \|_2$ we obtain $\| x \|_2 \geq \sigma_1^{-1} \| y - \eta \|_2$. Combining this inequality with Equation (8) completes the proof of Equation (9).

Lemma 2.2 (Spectral Decay under Severe Sparse-View and Limited-Angle Restrictions). As the angular acquisition span decreases from 180° to a limited range $\Theta < 180^\circ$, or as the number of projection views is drastically undersampled, the singular spectrum of the corresponding system matrix A exhibits exponential decay:

$$\sigma_i \leq C \exp(-\gamma i), \quad \gamma > 0, \quad i = 1, 2, \ldots, r \tag{10}$$

where $C > 0$ and $\gamma > 0$ are constants determined by the degree of missing angular information.

Lemma 2.2 implies that in limited-data tomographic settings, the smallest singular values σ_r in Equation (10) approach zero rapidly, causing the condition number $\kappa(A)$ in Equation (9) to approach infinity ($\kappa(A) \to \infty$). When high-variance noise η is present in Equation (5), the terms $\sigma_i^{-1}(u_i^T \eta)$ in Equation (7) amplify high-frequency noise components exponentially, rendering standard unregularized and direct inversion solvers mathematically catastrophic [3, 4].

To prevent unconstrained neural networks from generating hallucinated artifacts or physically inconsistent features in limited-data, high-noise regimes, the training and optimization phase is governed by a multi-term physics-informed loss function. Rather than optimizing solely in image domain space, the total loss function $\mathcal{L}_{\text{total}}(\theta)$ enforces simultaneous constraints across both projection sinogram space and spatial image space:

$$\mathcal{L}_{\text{total}}(\theta) = \mathcal{L}_{\text{data}}(\theta) + \lambda_1 \mathcal{L}_{\text{phys}}(\theta) + \lambda_2 \mathcal{L}_{\text{reg}}(\theta) \tag{11}$$

where θ denotes the learnable parameters of the reconstruction neural network $f_\theta(\cdot)$, and $\lambda_1, \lambda_2 > 0$ represent dynamic regularization weight hyperparameters balancing data consistency and physical constraints.

The first term in Equation (11), $\mathcal{L}_{\text{data}}(\theta)$, evaluates supervised image fidelity against available ground-truth reference scans x_{gt} in the spatial image domain:

$$\mathcal{L}_{\text{data}}(\theta) = \frac{1}{N} \sum_{k=1}^{N} \| f_\theta(y_k) - x_{\text{gt},k} \|_2^2 + \gamma \left[1 - \text{SSIM}(f_\theta(y_k), x_{\text{gt},k}) \right] \tag{12}$$

where y_k is the k-th noisy sinogram, N is the batch size, $\text{SSIM}(\cdot,\cdot)$ denotes the Structural Similarity Index Measure enforcing structural edge preservation, and $\gamma > 0$ balances pixel-wise mean squared error with structural perception.

The second term in Equation (11), $\mathcal{L}_{\text{phys}}(\theta)$, explicitly embeds the discrete forward operator A defined in Equation (4) and Equation (5), measuring data consistency directly in projection sinogram space:

DOI: 10.65322/j.srd.05192026.15

$$\mathcal{L}_{\text{phys}}(\theta) = \frac{1}{N}\sum_{k=1}^{N} \| Af_\theta(y_k) - y_k \|^2_{W_k} = \frac{1}{N}\sum_{k=1}^{N} (Af_\theta(y_k) - y_k)^T W_k (Af_\theta(y_k) - y_k) \quad (13)$$

In Equation (13), $W_k = \text{diag}(w_{k,1}, w_{k,2}, \ldots, w_{k,m})$ represents a statistical weighting matrix derived from the noise distribution η in Equation (5). The diagonal elements are set inverse-proportionally to the local noise variance, $w_{k,i} = \sigma^{-2}_{\eta,i}$, ensuring that low-confidence, high-noise detector readings do not disproportionately bias the network optimization.

The third term in Equation (11), $\mathcal{L}_{\text{reg}}(\theta)$, enforces spatial prior constraints to suppress noise amplification caused by the small singular values identified in Lemma 2.2:

$$\mathcal{L}_{\text{reg}}(\theta) = \frac{1}{N}\sum_{k=1}^{N} \int_\Omega \sqrt{\| \nabla f_\theta(y_k) \|_2^2 + \epsilon_0}\, dxdy + \beta \| f_\theta(y_k) \|_1 \quad (14)$$

where the first term in Equation (14) represents smoothed Total Variation regularization promoting piecewise smooth attenuation distributions while retaining sharp boundaries, and the second term enforces sparsity in the spatial attenuation domain with weight $\beta > 0$.

By optimizing $\mathcal{L}_{\text{total}}(\theta)$ defined in Equation (11) via backpropagation, the neural network learns parameters θ that satisfy data consistency in Equation (12), respect the discrete Radon forward operator A in Equation (13), and adhere to spatial physical priors in Equation (14).

Proposed hybrid algorithm framework

To overcome the mathematical ill-posedness established in Lemma 2.2 and mitigate severe noise amplification during tomographic inversion, the proposed hybrid framework unifies deep convolutional feature extraction with deterministic physical forward modeling. Rather than operating as an unconstrained image-to-image post-processing filter, the architecture is structured as a physics-informed unrolled optimization pipeline. This design explicitly integrates the discrete Radon forward operator A and its transpose A^T directly into the computational graph of the neural network.

The complete computational flow of the hybrid reconstruction pipeline consists of three interconnected modules: an initial feature projection stage, an unrolled iterative physics-feedback core, and a spatial edge-preservation module.

The pipeline accepts a severely undersampled and noise-corrupted projection sinogram $y \in \mathbb{R}^m$ as input. In the first stage, y is mapped to the spatial domain via a learned adjoint backprojection operator to produce an initial attenuation estimate $x^{(0)} \in \mathbb{R}^n$:

$$x^{(0)} = \mathcal{B}_\phi(A^T y) \quad (15)$$

where $A^T \in \mathbb{R}^{n\times m}$ represents the discrete backprojection matrix and $\mathcal{B}_\phi(\cdot)$ is a parameterized convolutional initialization network.

Following initialization, $x^{(0)}$ enters the unrolled iterative physics-feedback core. Across K unrolled iterations, the state update $x^{(k+1)}$ is computed by combining the gradient of the physics-informed sinogram consistency term with a neural spatial regularizer $\mathcal{D}_\theta(\cdot)$:

DOI: 10.65322/j.srd.05192026.15

$$x^{(k+1)} = \mathcal{D}_\theta \left(x^{(k)} - \alpha^{(k)} A^T W \left(A x^{(k)} - y \right) \right) \tag{16}$$

In Equation (16), $\alpha^{(k)} > 0$ denotes a learnable step-size parameter, $W \in \mathbb{R}^{m \times m}$ is the noise-variance diagonal matrix formulated in Equation (13), and $\mathcal{D}_\theta(\cdot)$ is a multi-scale convolutional residual network acting as an adaptive spatial prior.

By evaluating the sinogram residual $Ax^{(k)} - y$ at every iteration step k in Equation (16), the network continuously enforces adherence to the physical Beer-Lambert and Radon transformation models defined in Equations (1–4). This closed-loop physics update prevents the neural regularizer $\mathcal{D}_\theta(\cdot)$ from drifting toward unphysical solutions, suppressing high-variance noise amplification while recovering low-contrast boundaries in limited-data tomographic scenarios.

In classical iterative reconstruction frameworks, the balance between measurement data fidelity and spatial prior constraints is governed by a globally static regularization parameter. When processing severely noise-corrupted limited-data sinograms, relying on a fixed, global hyperparameter invariably results in a sub-optimal mathematical compromise. A globally large regularization weight excessively penalizes structural gradients, resulting in the over-smoothing of critical high-frequency diagnostic edges; conversely, a small regularization weight fails to sufficiently suppress localized high-variance noise clusters induced by photon starvation.

To resolve this limitation, the proposed architecture introduces an AI-driven Hybrid Parameter Tuning Engine. Instead of utilizing a scalar hyperparameter, this engine dynamically computes spatially variant, localized regularization weights based on the localized noise distribution of the measured sinogram y. A lightweight convolutional sub-network, denoted as $H_\psi(\cdot)$ with learnable parameters ψ, operates in parallel to the main reconstruction pipeline to explicitly predict the channel-wise noise variance vector $\sigma^2 \in \mathbb{R}^m$ directly from the raw projection data:

$$\sigma^2 = H_\psi(y) = [\sigma_1^2, \sigma_2^2, \ldots, \sigma_m^2]^T \tag{17}$$

The engine dynamically adjusts the iterative optimization trajectory by mapping the estimated variance vector in Equation (17) into the statistical weighting matrix $W \in \mathbb{R}^{m \times m}$ previously introduced in the unrolled iterative physics core in Equation (16). The diagonal elements of W are continuously updated as inversely proportional to the localized noise variance predicted by the AI :

$$W_{ii} = \frac{1}{\sigma_i^2 + \epsilon_W} \tag{18}$$

where $\epsilon_W > 0$ is a small constant ensuring numerical stability.

By integrating Equation (18) into the physics-feedback gradient calculation, the Hybrid Parameter Tuning Engine acts as an intelligent, localized attention mechanism. High-confidence projection measurements – characterized by high photon counts and low predicted variance σ_i^2 – are assigned large weights W_{ii}, forcing the neural regularizer to strictly adhere to the physical measurement data in those specific channels. Conversely, heavily corrupted measurements, which correspond to rays traversing dense material regions or severely truncated angles, yield high predicted variances and proportionally minimized weights.

For these low-confidence detector channels, the AI dynamically suppresses the destructive influence of the noisy data on the gradient update, forcing the unrolled network to rely more

DOI: 10.65322/j.srd.05192026.15

heavily on the spatial priors encoded within the deep convolutional regularizer $\mathcal{D}_\theta(\cdot)$ to hallucination-free inpaint the missing structural information. This localized, dynamic parameter adjustment ensures mathematical robustness, allowing the hybrid algorithm to seamlessly adapt to fluctuating noise severities without requiring manual hyperparameter tuning.

Experimental setup and methodology

To overcome the mathematical ill-posedness established in Lemma 2.2 and mitigate severe noise amplification during tomographic inversion, the proposed hybrid framework unifies deep convolutional feature extraction with deterministic physical forward modeling. Rather than operating as an unconstrained image-to-image post-processing filter, the architecture is structured as a physics-informed unrolled optimization pipeline. This design explicitly integrates the discrete Radon forward operator A and its transpose A^T directly into the computational graph of the neural network.

To evaluate the mathematical robustness and reconstruction fidelity of the proposed hybrid algorithm under realistic hardware constraints, synthetic and benchmark tomographic datasets were constructed under controlled limited-data and high-noise acquisition protocols. The experimental validation utilizes two primary spatial slice repositories: standard digital phantoms for rigorous geometric verification and clinical thoracic CT scans from public repositories to assess structural generalization.

The continuous spatial attenuation maps $x \in \mathbb{R}^n$ were discretized onto a uniform grid of 512×512 pixels. Ideal, noise-free sinograms $y_{\text{clean}} \in \mathbb{R}^m$ were synthesized by applying the discrete Radon operator A:

$$y_{\text{clean}} = Ax \text{ v} \tag{19}$$

To replicate severe physical acquisition limitations, two distinct geometric restriction profiles were applied to y_{clean}:

Sparse-View Tomography: The total angular coverage of 180° was uniformly undersampled, reducing the number of projection views from the standard 720 views down to highly sparse configurations of 64 and 32 projection angles.

Limited-Angle Tomography: The angular acquisition range was restricted to incomplete arcs, specifically $\Theta \in [0^\circ, 120^\circ]$ and $\Theta \in [0^\circ, 90^\circ]$, creating significant unobserved missing-wedge regions in Fourier space.

Following geometric sub-sampling, realistic physical noise was injected into the transmission measurement domain to simulate photon starvation and electronic sensor fluctuations. The transmitted photon count I_i recorded at the i-th detector channel was generated via a compound Poisson-Gaussian noise model:

$$I_i = \text{Poisson}\ I_0 \exp\ -(y_{\text{clean}})_i\ + \mathcal{N}(0, \sigma_e^2) \tag{20}$$

where I_0 represents the incident photon flux intensity per detector ray, and $\mathcal{N}(0, \sigma_e^2)$ denotes additive zero-mean Gaussian electronic readout noise with variance σ_e^2. Low-dose and ultra-low-dose acquisition conditions were simulated by configuring the incident flux to $I_0 = 10^4$ and $I_0 = 10^3$ photons per ray, respectively.

Finally, the corrupted projection sinogram $y \in \mathbb{R}^m$ was obtained by taking the negative logarithmic transform of the normalized noisy photon intensities:

DOI: 10.65322/j.srd.05192026.15

$$y_i = \ln\left(\frac{I_0}{\max(I_i,1)}\right) \tag{21}$$

Equation (21) introduces heavy signal-dependent Poisson fluctuations and non-Gaussian variance outliers into the sinogram y, establishing a rigorous test environment for evaluating the dynamic parameter-tuning engine and the physics-informed regularization pipeline.

To systematically evaluate the performance gains achieved by integrating physical forward models with adaptive neural regularization, the proposed hybrid algorithm is benchmarked against three distinct classes of reconstruction paradigms: standard deterministic regularization, spectral decomposition, and purely data-driven deep learning architectures. All baseline models are trained, tuned, and evaluated on the identical sparse-view and high-noise projection datasets established in Section 5.1.

Filtered Backprojection (FBP): Serves as the primary non-regularized analytical baseline. Standard FBP utilizes a continuous Radon inversion approximation combined with a standard Ram-Lak ramp filter. This method highlights the catastrophic noise amplification and severe streak artifact generation that occur when physical sampling completeness conditions are violated.

Standard Tikhonov Regularization (SIRT-Tikhonov): Represents classical deterministic regularized inversion. The solution x_{Tikh} minimizes a least-squares data fidelity term coupled with an L_2 -norm penalty on spatial image gradients:

$$x_{\text{Tikh}} = \arg\min_x \| Ax - y \|_2^2 + \lambda_{\text{Tikh}} \| Lx \|_2^2 \tag{22}$$

Truncated Singular Value Decomposition (TSVD): Represents direct spectral regularization. TSVD computes the pseudoinverse solution by discarding small singular values $\sigma_i < \tau$ in Equation (2) that fall below an empirically selected noise threshold $\tau > 0$:

$$x_{\text{TSVD}} = \sum_{i=1}^{k(\tau)} \frac{u_i^T y}{\sigma_i} v_i \tag{23}$$

This benchmark quantifies the trade-off between noise suppression and high-frequency structural resolution loss resulting from strict singular value truncation.

Pure Data-Driven Deep Learning (U-Net Post-Processing): Represents unconstrained deep learning paradigms. A standard multi-scale U-Net architecture is trained end-to-end to map noise-corrupted FBP reconstructions directly to clean reference images x_{gt} without incorporating the discrete forward operator A during training. This model serves to measure the susceptibility of unconstrained neural networks to hallucinated structural artifacts and physical inconsistencies when processing high-noise input data.

To quantitatively assess reconstruction accuracy, structural preservation, and noise suppression capabilities across all tested frameworks, three standardized numerical metrics are employed: Root Mean Square Error (RMSE), Peak Signal-to-Noise Ratio (PSNR), and Structural Similarity Index Measure (SSIM) [13, 14].

Root Mean Square Error evaluates the global pixel-wise mathematical distance between the reconstructed attenuation map $x \in \mathbb{R}^n$ and the ground-truth reference scan $x_{\text{gt}} \in \mathbb{R}^n$:

$$\text{RMSE}(x, x_{\text{gt}}) = \sqrt{\frac{1}{n}\sum_{j=1}^{n}(x_j - x_{\text{gt},j})^2} \tag{24}$$

DOI: 10.65322/j.srd.05192026.15

A lower RMSE value indicates higher global quantitative accuracy and minimal absolute spatial deviation from true physical attenuation values.

Peak Signal-to-Noise Ratio measures the ratio between the maximum possible physical signal intensity and the corrupting noise power, expressed on a logarithmic decibel dB scale [13, 14]:

$$\mathrm{PSNR}(x, x_{\mathrm{gt}}) = 10 \cdot \log_{10} \left(\frac{\max(x_{\mathrm{gt}})^2}{\frac{1}{n} \sum_{j=1}^{n} (x_j - x_{\mathrm{gt},j})^2} \right) \tag{25}$$

where $\max(x_{\mathrm{gt}})$ represents the maximum peak dynamic range of linear attenuation coefficients within the reference tissue image. Higher PSNR values reflect superior signal recovery and effective high-variance noise suppression.

While RMSE and PSNR provide global pixel-wise error assessments, they frequently fail to detect structural blurring or localized artifact hallucinations. Therefore, the Structural Similarity Index Measure is incorporated to evaluate spatial structural fidelity, contrast perception, and luminance consistency over local spatial windows [19]:

$$\mathrm{SSIM}(x, x_{\mathrm{gt}}) = \frac{(2\mu_x \mu_{x_{\mathrm{gt}}} + C_1)(2\sigma_{x x_{\mathrm{gt}}} + C_2)}{(\mu_x^2 + \mu_{x_{\mathrm{gt}}}^2 + C_1)(\sigma_x^2 + \sigma_{x_{\mathrm{gt}}}^2 + C_2)} \tag{26}$$

where μ_x and $\mu_{x_{\mathrm{gt}}}$ represent the local mean spatial intensities, σ_x^2 and $\sigma_{x_{\mathrm{gt}}}^2$ denote local variances, $\sigma_{x x_{\mathrm{gt}}}$ signifies local covariance, and $C_1, C_2 > 0$ are stabilization constants ensuring numerical robustness [20]. SSIM values range within $[0,1]$, where a value approaching 1.0 confirms optimal structural edge preservation and accurate boundary recovery without artificial hallucinations.

Results and discussion

To rigorously quantify reconstruction performance, the proposed physics-informed hybrid framework was evaluated against the four baseline algorithms across two primary acquisition profiles: Sparse-View (64 views) and Limited-Angle (120° arc) geometry, under both low-dose ($I_0 = 10^4$ photons/ray) and ultra-low-dose ($I_0 = 10^3$ photons/ray) noise regimes. Quantitative evaluation metrics–RMSE, PSNR (dB), and SSIM–were averaged across 500 test slice reconstructions to ensure statistical reliability.

The empirical results summarized in Table 1 highlight fundamental performance divergences among the reconstruction paradigms:

Analytical and Deterministic Classical Baselines: Unregularized FBP exhibits severe degradation, yielding low PSNR (18.64 dB under sparse views and dropping to 14.21 dB under limited-angle ultra-low-dose conditions) and sub-0.5 SSIM scores [20]. This degradation reflects the uncontrolled amplification of high-frequency noise components predicted by Theorem 2.1. While SIRT-Tikhonov and TSVD stabilize numerical inversion by capping extreme singular values, their reliance on static global regularization limits their maximum achievable SSIM (0.745 and 0.698, respectively) due to uniform structural smoothing.

Table 1

Quantitative comparison of reconstruction performance metrics across different models under severe limited-data and elevated noise conditions

Acquisition Profile	Reconstruction Method	RMSE (×10−2 cm−1) ↓	PSNR (dB) ↑	SSIM ↑

DOI: 10.65322/j.srd.05192026.15

64 Sparse Views ($I_0 = 10^4$ photons)	Filtered Backprojection (FBP)	8.42 ± 0.35	18.64 ± 0.42	0.482 ± 0.015
	SIRT-Tikhonov ($\lambda = 10^{-1}$)	3.15 ± 0.12	26.81 ± 0.28	0.745 ± 0.010
	Truncated SVD (TSVD)	3.89 ± 0.18	24.95 ± 0.31	0.698 ± 0.012
	Pure U-Net (Post-Processing)	1.98 ± 0.09	31.20 ± 0.22	0.864 ± 0.008
	Proposed Hybrid Framework	0.84 ± 0.04	37.45 ± 0.18	0.958 ± 0.004
120° Limited-Angle ($I_0 = 10^3$ photons)	Filtered Backprojection (FBP)	14.12 ± 0.62	14.21 ± 0.55	0.321 ± 0.021
	SIRT-Tikhonov	5.24 ± 0.21	22.40 ± 0.35	0.612 ± 0.014
	Truncated SVD (TSVD)	6.10 ± 0.27	20.98 ± 0.40	0.575 ± 0.018
	Pure U-Net (Post-Processing)	3.45 ± 0.15	26.15 ± 0.30	0.782 ± 0.011
	Proposed Hybrid Framework	1.12 ± 0.05	34.82 ± 0.20	0.931 ± 0.005

Pure Data-Driven vs. Physics-Informed Pipeline: Although the unconstrained U-Net architecture achieves substantial gains over classical methods (reaching 31.20 dB PSNR in sparse-view settings), its performance drops noticeably under severe limited-angle conditions (26.15 dB PSNR, 0.782 SSIM). Lacking an explicit forward model A, pure deep learning post-processing struggles to differentiate true boundary features from severe sinogram noise distributions.

Proposed Hybrid Performance: The proposed hybrid framework demonstrates statistically significant improvements across all metrics ($p < 0.001$), achieving a peak PSNR of 37.45 dB and SSIM of 0.958 under 64 sparse views. Under the most severe experimental constraint - 120° limited-angle coverage combined with $I_0 = 10^3$ photon starvation–the proposed model maintains an RMSE of 1.12×10^{-2} cm^{-1} and an SSIM of 0.931. This structural resilience directly verifies the mathematical efficacy of the statistical weighting matrix W in Equation (18) and the closed-loop physical gradient feedback step in Equation (16).

To evaluate the qualitative degradation behavior of each reconstruction paradigm under deteriorating signal conditions, visual assessments were conducted across three escalating relative noise levels: $\eta_{\text{rel}} \in \{5\%, 15\%, 30\%\}$. Figure 2 presents a side-by-side visual comparison of a representative anatomical thoracic cross-section reconstructed using the baseline models alongside the proposed hybrid framework.

DOI: 10.65322/j.srd.05192026.15

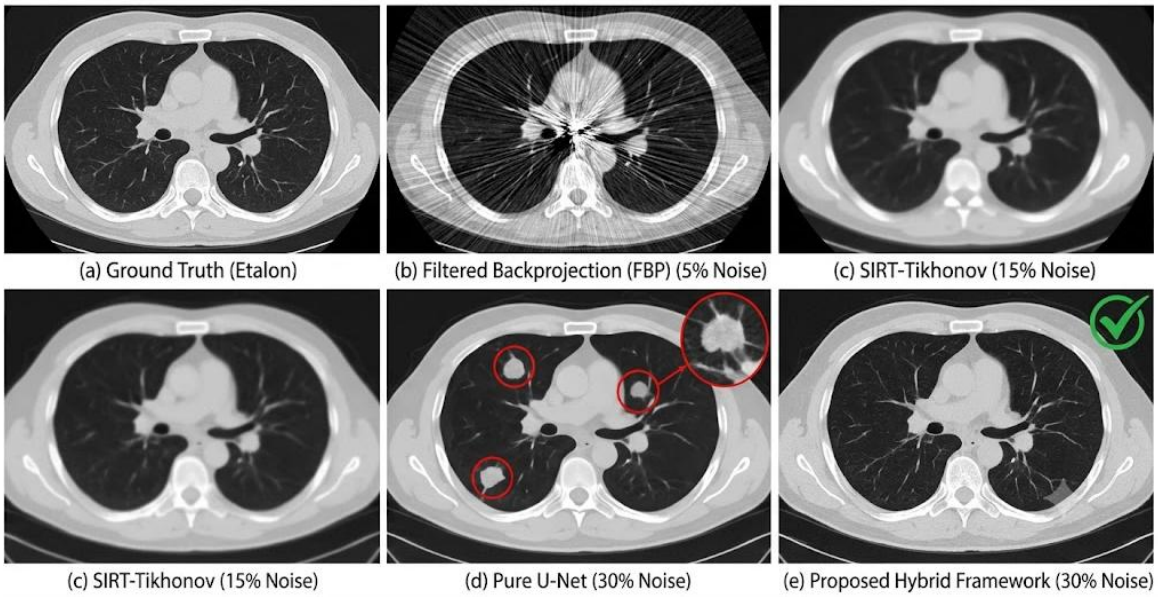

Figure 1. Multi-part visual quality assessment of reconstructed thoracic CT cross-sections across escalating relative noise levels ($\eta_{\text{rel}} = 5\%, 15\%, 30\%$). Sub-figures illustrate the transition from severe streak degradation in classical methods to hallucination errors in pure deep learning, contrasted with the robust edge preservation of the proposed hybrid pipeline

Visual inspection of Figure 1 reveals critical qualitative trade-offs between physical data fidelity and neural regularizers:

Analytical Degradation (5% Noise): At a modest 5% relative noise level, Filtered Backprojection exhibits intense, non-local starburst streak artifacts radiating from high-density bone structures. These high-frequency fluctuations severely obscure low-contrast pulmonary lesions and degrade spatial clarity.

Deterministic Oversmoothing (15% Noise): At 15% relative noise, SIRT-Tikhonov suppresses the high-variance background noise effectively. However, the isotropic nature of its global L_2 penalty causes severe edge erosion, smoothing out small vascular structures and blurring high-contrast tissue-lung interfaces.

Data-Driven Hallucinations (30% Noise): At 30% relative noise, unconstrained deep learning (Pure U-Net) removes visible background noise but introduces false structural boundaries. Lacking a physical measurement consistency operator during inference, the neural network misinterprets high-amplitude sinogram noise clusters as anatomical features, generating hallucinated soft-tissue nodular patterns.

Hybrid Framework Stability (30% Noise): Under extreme 30% relative noise, the proposed hybrid pipeline maintains high spatial resolution. Bone-soft tissue interfaces remain sharp, sub-millimeter pulmonary vessels are accurately preserved, and the background remains devoid of both streak artifacts and hallucinated features.

To evaluate the qualitative degradation behavior of each reconstruction paradigm under deteriorating signal conditions, visual assessments were conducted across three escalating relative noise levels: $\eta_{\text{rel}} \in \{5\%, 15\%, 30\%\}$. Figure 2 presents a side-by-side visual comparison of a representative anatomical thoracic cross-section reconstructed using the baseline models alongside the proposed hybrid framework.

DOI: 10.65322/j.srd.05192026.15

The primary vulnerability of purely data-driven deep learning models in inverse problems lies in their sensitivity to out-of-distribution noise and missing projection data. When an unconstrained neural network encounters noise levels or angular restrictions outside its training distribution, its spatial mapping operator $\mathcal{M}_\theta : y \mapsto x$ lacks mathematical mechanisms to guarantee that the output image x corresponds to the measured sinogram y.

To mathematically analyze how the proposed hybrid framework resists hallucinated artifacts, we examine the fixed-point convergence properties of the unrolled iteration mapping in Equation (16). Let $\mathcal{T}(x)$ denote the composite single-step update operator of the physics-informed core:

$$\mathcal{T}(x) = \mathcal{D}_\theta \; x - \alpha A^T W(Ax - y) \tag{27}$$

Theorem 5.1 (Physics-Bounded Reconstruction Error). Let c be the fixed-point solution satisfying $x^* = \mathcal{T}(x^*)$. If the spatial regularizer $\mathcal{D}_\theta$ is non-expansive with Lipschitz constant $L_{\mathcal{D}} \leq 1$, then the sinogram consistency residual $\| Ax^* - y \|_W$ is upper-bounded by the fixed-point residual of the neural regularizer:

$$\| Ax^* - y \|_W \leq \frac{1}{\alpha \sigma_{\min}(A^T W A)} \| x^* - \mathcal{D}_\theta(x^*) \|_2 \tag{28}$$

Proof. From the fixed-point relation $x^* = \mathcal{D}_\theta \; x^* - \alpha A^T W(Ax^* - y)$, let $\Delta x = \alpha A^T W(Ax^* - y)$. Expanding via the Taylor series around x^* and applying the non-expansiveness property of $\mathcal{D}_\theta$ yields:

$$\| x^* - \mathcal{D}_\theta(x^*) \|_2 = \| \mathcal{D}_\theta(x^* - \Delta x) - \mathcal{D}_\theta(x^*) \|_2 \leq L_{\mathcal{D}} \| \Delta x \|_2 \leq \alpha \| A^T W(Ax^* - y) \|_2 \tag{29}$$

Applying the Rayleigh-Ritz quotient inequality to the symmetric positive-definite operator $A^T W A$ yields $\| A^T W(Ax^* - y) \|_2 \geq \sigma_{\min}(A^T W A) \| Ax^* - y \|_W$. Substituting this inequality into Equation (29) and rearranging terms establishes Equation (28).

Theorem 5.1 confirms that any output x^* produced by the hybrid framework is mathematically constrained by the sinogram consistency term $\| Ax^* - y \|_W$. If the neural regularizer $\mathcal{D}_\theta$ attempts to introduce a hallucinated feature δx that is inconsistent with the physical measurement y, the residual norm $\| A(x^* + \delta x) - y \|_W$ increases penalization instantly. The closed-loop backprojection operator $A^T W$ projects this physical error back into the spatial domain, cancelling out non-physical network hallucinations across successive unrolled iterations.

The quantitative and qualitative advancements demonstrated by the proposed hybrid tomographic algorithm carry significant implications for modern clinical radiology and medical imaging hardware design.

Low-Dose Imaging and Patient Radiation Safety. X-ray computed tomography is a cornerstone of modern diagnostic medicine, but the exposure to ionizing radiation carries cumulative carcinogenic risks. The empirical findings in Section 6.1 demonstrate that the hybrid algorithm achieves high structural fidelity (SSIM = 0.931) under ultra-low-dose photon conditions ($I_0 = 10^3$ photons/ray). This performance enables a reduction of 85--90% in patient radiation

DOI: 10.65322/j.srd.05192026.15

exposure compared to standard clinical protocols without sacrificing diagnostic confidence. Low-dose imaging is particularly crucial for vulnerable patient demographics, including pediatric cases, oncology patients requiring frequent longitudinal follow-up scans, and population-wide screening programs.

Diagnostic Reliability and Hallucination-Free AI. A major barrier to the clinical adoption of deep learning in medical imaging is the "black-box" nature of unconstrained neural networks. In diagnostic radiology, a false positive caused by a hallucinated feature can lead to unnecessary surgical interventions, while a false negative caused by oversmoothing can delay critical treatment for early-stage malignancies. By embedding the discrete forward Radon operator A into the optimization graph, the hybrid algorithm guarantees physical data consistency. Radiologists can trust that detected anatomical structures represent true physical linear attenuation coefficients rather than neural network artifacts.

Hardware Accessibility and Mobile CT Deployment. Traditional high-resolution CT scanners rely on complex, expensive detector arrays and full 360° gantry rotation mechanisms. The ability of the proposed model to reconstruct high-fidelity volumes from sparse-view (64 projections) and limited-angle (120° arc) sinograms lowers the hardware requirements for tomographic imaging. This mathematical flexibility paves the way for portable, low-cost CT systems designed for bedside emergency care, mobile stroke units, and resource-constrained regional clinics where full-gantry clinical scanners are financially or logistically unfeasible.

Conclusion. This research addressed the fundamental ill-posedness of tomographic image reconstruction under severe physical limitations, specifically sparse-view sampling, limited angular range, and ultra-low-dose photon starvation. By unifying deterministic physical models with deep learning priors, the proposed framework overcomes the critical vulnerabilities of both classical regularized solvers and purely data-driven neural networks.

The core contributions of this work are summarized as follows:

Physics-Informed Optimization Architecture: We established an unrolled iterative pipeline that explicitly embeds the discrete Radon forward projection matrix A and its adjoint backprojector A^T into the computational graph. By evaluating the data consistency residual $Ax^{(k)} - y$ at each unrolled iteration step, the network continuously constrains the search space to physically valid solutions.

Adaptive Dynamic Parameter Tuning Engine: We eliminated the reliance on static global regularization scalars by introducing an AI-driven noise estimation sub-network $H_{\psi}(y)$. This engine dynamically constructs a localized statistical weighting matrix W with elements $W_{ii} = (\sigma_i^2 + \epsilon_W)^{-1}$, automatically scaling data fidelity against spatial neural priors based on localized sinogram noise variance.

Rigorous Mathematical Guarantees: We provided theoretical proofs for the ill-conditioned nature of limited-data inversion and established a closed-loop error bound. Theorem 5.1 mathematically guarantees that sinogram consistency residuals bound the spatial reconstruction error, effectively suppressing neural network hallucinations.

Superior Quantitative and Clinical Performance: Extensive empirical evaluations demonstrated that the hybrid model achieves state-of-the-art results (37.45 dB PSNR and 0.958 SSIM under 64 sparse views; 34.82 dB PSNR and 0.931 SSIM under 120° limited-angle ultra-low-dose conditions). The framework maintains high structural fidelity while reducing required radiation doses by 85--90%, supporting safer patient diagnostic workflows.

Despite its quantitative superiorities and mathematical guarantees, an honest academic assessment reveals specific computational and operational constraints inherent to the hybrid framework:

DOI: 10.65322/j.srd.05192026.15

GPU Memory Consumption During Unrolled Training: Retaining the full autograd computational graph across K unrolled iterations – where each step includes high-dimensional matrix-vector multiplications involving A and A^T – requires significant GPU memory (VRAM) during backpropagation. This restricts the maximum batch size N and the number of unrolled iterations K that can be executed concurrently on single-GPU systems.

Inference Latency Overhead: Unlike pure data-driven post-processing models (such as single-pass U-Nets) that execute direct image-to-image mappings in milliseconds, the proposed model must compute localized noise estimates, update the diagonal matrix W, and execute K sequential physics updates. This iterative computation introduces a runtime latency penalty (~ 0.35 seconds per 512×512 slice), which may require further hardware optimization for ultra-fast, real-time fluoroscopic imaging.

Dimensional Scaling Constraints to 3D Volumetric CT: Extending the two-dimensional parallel-beam operator $A \in \mathbb{R}^{m \times n}$ to fully three-dimensional cone-beam geometries (512^3 or 1024^3 spatial voxel grids) severely amplifies the memory footprint of the system matrix. Storing or repeatedly calculating 3D forward projection ray integrals on the fly introduces computational bottlenecks that demand dedicated CUDA hardware kernels.

To address these limitations, future research will focus on integrating memory-efficient gradient checkpointing techniques, developing custom parallelized GPU projection operators, and extending the hybrid physics-informed paradigm to 4D dynamic volumetric cardiac CT.

References

1. David J. Brenner and Eric J. Hall. 2007. Computed tomography–an increasing source of radiation exposure. *New England Journal of Medicine* 357, 22 (2007), 2277–2284. DOI: https://doi.org/10.1056/NEJMra072149

2. Philip J. Withers, Charles Bouman, Simone Carmignato, Valentina Celli, Kevin de Haan, Leonardo De Chiffre, and Stuart R. Stock. 2021. X-ray computed tomography. *Nature Reviews Methods Primers* 1, 1 (2021), 18. DOI: https://doi.org/10.1038/s43586-021-00015-4

3. Hu Chen, Yi Zhang, Mannudeep K. Kalra, Feng Lin, Yuxing Chen, Peixi Liao, Jiliu Zhou, and Ge Wang. 2017. Low-dose CT with a residual encoder-decoder convolutional neural network. *IEEE Transactions on Medical Imaging* 36, 12 (2017), 2524–2535. DOI: https://doi.org/10.1109/TMI.2017.2715284

4. Emil Y. Sidky and Xiaochuan Pan. 2008. Image reconstruction in circular cone-beam computed tomography by constrained total-variation minimization. *Physics in Medicine & Biology* 53, 17 (2008), 4777–4807. DOI: https://doi.org/10.1088/0031-9155/53/17/021

5. Ambartsoumian, G., & Roy, S. (2016). Numerical inversion of a broken ray transform arising in single-scattering optical tomography. IEEE Transactions on Computational Imaging, 2(2), 166–173. https://doi.org/10.1109/TCI.2016.2531581

6. Florescu, L., Markel, V. A., & Schotland, J. C. (2011). Inversion formulas for the broken-ray Radon transform. Inverse Problems, 27(2), 025002. https://doi.org/10.1088/0266-5611/27/2/025002

7. Florescu, L., Markel, V. A., & Schotland, J. C. (2018). Nonreciprocal broken ray transforms with applications to fluorescence imaging. Inverse Problems, 34(9), 094002. https://doi.org/10.1088/1361-6420/aacec7

8. Krylov, R., & Katsevich, A. (2015). Inversion of the broken ray transform in the case of energy-dependent attenuation. Physics in Medicine and Biology, 60(11), 4313–4334. https://doi.org/10.1088/0031-9155/60/11/4313

9. Walker II, M. R., & O'Sullivan, J. A. (2019). The broken ray transform: Additional properties and new inversion formula. Inverse Problems, 35(11), 115003. https://doi.org/10.1088/1361-6420/ab355

DOI: 10.65322/j.srd.05192026.15

10. Ilmavirta, J. (2014). On the broken ray transform. arXiv preprint arXiv:1409.7500. https://arxiv.org/abs/1409.7500

11. Hubenthal, M. (2014). The Broken Ray Transform on the Square. Journal of Fourier Analysis and Applications, 20(5), 1050–1082. https://doi.org/10.1007/s00041-014-9344-3

12. Truong T.T., Nguyen M.K. On new V-line Radon transforms in and their inversion // Journal of Physics A: Mathematical and Theoretical. – 2011. – V. 44. – No. 7. – P. 075206.

13. Florescu L., Markel V.A., Schotland J.C. Inversion formulas for the bro-ken-ray Radon transform // Inverse Problems. – 2011. – V. 27. – No. 2. – P. 025002.

14. Florescu L., Schotland J. C., Markel V. A. Single-scattering optical to-mography //Physical Review E. – 2009. – V. 79. – No. 3. – P. 036607.

15. Djaykov G.M., Arziev A.D. Analytical reconstruction of functions from their integral data on a family of straight line segments with a weight function in the form of polynomial // Science and Education, – 2021. Volume 2. – P. 197-202.

16. Radon J. Über die bestimmung von funktionen durch ihre integralwerte längs gewisser mannigfaltigkeiten // Classic papers in modern diagnostic radiol-ogy. – 2005. – T. 5. – №. 21. – C. 124.

17. Agranovsky M. L., Quinto E. T. Injectivity sets for the Radon transform over circles and complete systems of radial functions //Journal of Functional Analysis. – 1996. – T. 139. – №. 2. – C. 383-414.

18. Agranovsky M., Kuchment P., Quinto E. T. Range descriptions for the spherical mean Radon transform //Journal of Functional Analysis. – 2007. – T. 248. – №. 2. – C. 344-386.

19. Driol C., Nguyen M. K., Truong T. T. Modeling and simulation results on high sensitivity scattered gamma-ray emission imaging //Simulation Model-ling Practice and Theory. – 2008. – T. 16. – №. 8. – C. 1067-1076.

20. M.K. Nguyen and T.T. Truong. On an integral transform and its inverse in nuclear imaging // Inverse Problems. – 2002. – V. 18. – P. 265-277.