

Physics-Guided Dual-Domain Plug-and-Play ADMM for Low-Dose CT Reconstruction

Sayantana Dutta¹ 

SAYANTAN.DUTTA1@GEHEALTHCARE.COM

Sudhanya Chatterjee¹

SUDHANYA.CHATTERJEE@GEHEALTHCARE.COM

Ashwini Galande¹

ASHWINI.GALANDE@GEHEALTHCARE.COM

K. S. Shriram¹

KS.SHRIRAM@GEHEALTHCARE.COM

Bipul Das¹

BIPUL.DAS@GEHEALTHCARE.COM

¹ Science and Technology Organization, GE HealthCare, Bangalore 560066, Karnataka, India.

Editors: Under Review for MIDL 2026

Abstract

Ultra-low-dose CT (ULDCT) imaging can greatly reduce patient radiation exposure, but the resulting scans suffer from severe structured and random noise that degrades image quality. To address this challenge, we propose a novel Plug-and-Play model-based iterative reconstruction framework (PnP-MBIR) that integrates a deep convolutional denoiser trained in a 2-stage self-supervised Noise-to-Noise (N2N) scheme. The method alternates between enforcing sinogram-domain data fidelity and applying the learned image-domain denoiser within an optimization, enabling artifact suppression while maintaining anatomical structure. The 2-stage protocol enables fully self-supervised training from noisy data, followed by high-dose fine-tuning, ensuring the denoiser’s robustness in the ultra-low-dose regime. Our method enables high-quality reconstructions at $\sim 70\text{--}80\%$ lower dose levels, while maintaining diagnostic fidelity comparable to standard full-dose scans. Quantitative evaluations using Gray-Level Co-occurrence Matrix (GLCM) features — including contrast, homogeneity, entropy, and correlation — confirm that the proposed method yields superior texture consistency and detail preservation over standalone deep learning and supervised PnP baselines. Qualitative and quantitative results on both simulated and clinical datasets demonstrate that our framework effectively reduces streaks and structured artifacts while preserving subtle tissue contrast, making it a promising tool for ULDCT reconstruction.

Keywords: Low-dose CT, Model-based deep learning, Plug-and-Play ADMM, Dual-domain reconstruction, Deep denoising, CT image reconstruction.

1. Introduction

X-ray computed tomography (CT) is essential for diagnosis but exposes patients to ionizing radiation, posing long-term risks, especially for children and those needing frequent scans (Smith-Bindman et al., 2010). Guided by the ALARA principle and initiatives like *Image Gently*, minimizing exposure while preserving diagnostic quality is critical (Brenner and Hall, 2007; Clement et al., 2025). Pediatric CT has been linked to higher risks of leukemia and brain tumors (Brenner and Hall, 2007), emphasizing the need for ultra-low-dose CT (ULDCT) protocols that improve safety without compromising utility.

However, reducing the tube current or voltage in CT comes at the cost of significantly degraded image quality. ULDCT images are affected by increased quantum noise (Poisson mottle) and severe streak artifacts from photon starvation, which obscure low-contrast features and anatomical details (Ma et al., 2025). The resulting low signal-to-noise ratio (SNR)

can mask subtle pathologies, making reliable diagnosis challenging. Various reconstruction strategies have been developed to address these issues, broadly classified into (1) analytical methods (Yazdanpanah et al., 2016), (2) model-based iterative reconstruction (MBIR) (Sidky and Pan, 2008; Zhang et al., 2014; Kim et al., 2015), and (3) modern learning-based techniques (Chen et al., 2017; Xu et al., 2012; Jin et al., 2017; Fu and De Man, 2022).

Filtered back projection (FBP) is computationally efficient but highly noise-sensitive. Sinogram-domain filtering (Balda et al., 2012; Manduca et al., 2009) partly alleviates this, though excessive smoothing often erodes fine structural details (Wang et al., 2006). MBIR enhances image quality through statistical noise modeling (Bouman and Sauer, 1996) and priors such as total variation (Zhang et al., 2016a), non-local similarity (Zhang et al., 2016b; Dutta et al., 2022a), dictionary learning (Xu et al., 2012), and low-rank constraints (Cai et al., 2014; Dutta et al., 2021b). Despite these advantages, MBIR remains computationally intensive and constrained by proprietary, vendor-specific implementations which hinders clinical adoption (Katsura et al., 2012).

Deep learning (DL) has become a powerful approach for low-dose CT (LDCT) denoising, offering fast, data-driven alternatives to traditional reconstruction. CNNs (Chen et al., 2017), U-Nets, and GANs (Jin et al., 2017; Ma et al., 2025) learn nonlinear mappings from low-dose to high-dose images, capturing priors that suppress noise while preserving anatomy. Hybrid methods (Fu and De Man, 2022) and plug-and-play (PnP) frameworks (Venkatakrishnan et al., 2013) embed pretrained denoisers (Floquet et al., 2024; Dutta et al., 2024a) within iterative schemes like ADMM (Dutta and Mamou, 2024, 2025), enforcing data fidelity without explicit regularization (Dutta et al., 2022b). Parameterized PnP (Ye et al., 2018; He et al., 2019) and deep unfolding models (Kong et al., 2022; Dutta et al., 2023) achieve state-of-the-art LDCT performance, with studies confirming PnP’s superiority in structure preservation and dose generalization (Xu et al., 2022). However, DL faces challenges: over-smoothing and loss of subtle textures or low-contrast lesions (Chen et al., 2017); limited generalization across scanners or dose levels (Ye et al., 2018; He et al., 2019); and reliance on large paired datasets, scarce in ULDCCT (Xu et al., 2022). Self-supervised approaches like Noise-to-Noise (N2N) (Chan et al., 2019; Lehtinen et al., 2018; Dutta et al., 2024b) mitigate these issues, but purely self-supervised CT models remain hard to train and risk texture loss without careful regularization (Li et al., 2020; Wu et al., 2021). Another concern is image-domain inconsistency with measured sinograms, causing CT number errors. Structured noise-aware denoising (Langoju et al., 2024) and dual-domain strategies (Fan et al., 2024; Wang et al., 2022) address this by enforcing measurement fidelity while preserving perceptual detail. Iterative refinement in both domains improves artifact suppression over single-domain methods (Wang et al., 2022), enhancing texture recovery, reducing structured noise, and preserving quantitative CT accuracy.

Building on these insights, we propose a DL-based MBIR framework (PnP-MBIR) that embeds a CNN denoiser as a proximal operator within an ADMM reconstruction loop for ULDCCT. The algorithm alternates between data-consistent projection updates and image-domain denoising, ensuring fidelity to the measured sinogram at each iteration. Unlike black-box networks, our modular design separates physics-based forward modeling from learned priors. The denoiser is trained using a N2N strategy on simulated or real low-dose data across varying dose levels, enabling suppression of both structured and unstructured noise/artifacts without clean ground-truth references—critical for ultra-low-dose

regime. This PnP architecture flexibly integrates realistic noise-trained priors, improving generalizability across dose levels and scan conditions. Validated on ULDCT datasets, our method achieves superior image quality: higher entropy in uniform regions (avoiding over-smoothing), improved soft-tissue contrast-to-noise ratio, and sharper anatomical edges, yielding better visibility of low-contrast lesions and crisper anatomical boundaries compared to FBP and standalone denoising. Key contributions include:

- *Dual-Domain PnP Recon*: We propose a novel MBIR framework that incorporates deep priors in the image domain while enforcing sinogram-domain data consistency. Unlike purely image-based methods, our approach performs corrective updates in projection space at each iteration, reducing artifacts and ensuring CT number accuracy.
- *Noise-to-Noise (N2N) Deep Prior*: We utilize a N2N-trained CNN denoiser as the pnp prior, trained entirely on noisy ULDCT data without requiring ground-truth high-dose images. This self-supervised setup captures real-world noise statistics (e.g., Poisson and electronic noise) of ULDCT, enhancing training practicality and robustness.
- *Dose-Level Generalization*: Our 2-stage N2N training spans a wide range of dose levels (ultra-low to standard), enabling the network to learn a clean image manifold from artifact-rich pairs. This broad training improves generalization across scan conditions. To prevent potential hallucinations from unspecific mappings, the sinogram-domain correction step in our PnP framework anchors reconstructions to measured data.
- *Improved Image Quality*: On ULDCT datasets (simulated and real), our method consistently outperforms FBP and baseline denoisers in both visual and quantitative evaluations. It yields higher entropy (indicating natural textures), improved soft-tissue contrast, and sharper anatomical boundaries, all while preserving attenuation accuracy and avoiding over-smoothing.

2. Methodology

2.1. Problem Formulation

In a fan-beam CT acquisition, the measured sinogram $\mathbf{y} \in \mathbb{R}^M$ is related to the object image $\mathbf{x} \in \mathbb{R}^N$ via a forward model involving the system matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$, representing the discrete fan-beam projection geometry. Under ultra-low-dose conditions, the measurement noise is no longer i.i.d. Gaussian; instead, it becomes *structured* and *correlated* due to factors such as extremely low photon counts (photon starvation) and unmodeled physics (beam hardening). These effects cause non-uniform noise variance and streak artifacts in the sinogram. In practice, noise in each ray often exhibits variance dependent on detected photon counts, and adjacent rays or views may share correlated components.

Reconstructing \mathbf{x} from such noisy projections is posed as a Maximum a posteriori (MAP) estimation problem (Boyd et al., 2011), which seeks the image that best fits the measurements while adhering to prior expectations of image appearance. Formally, we define a cost function comprising a data fidelity term and a regularization (*prior*) term:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{Ax} - \mathbf{y}\|_{\mathbf{W}}^2 + \lambda R(\mathbf{x}), \quad (1)$$

where $\|\mathbf{z}\|_{\mathbf{W}}^2 = \mathbf{z}^T \mathbf{W} \mathbf{z}$ denotes a weighted ℓ_2 norm. Here $\mathbf{W} = \Sigma^{-1}$ is a weighting (precision) matrix reflecting the noise covariance in \mathbf{y} . In the simplest case, \mathbf{W} is taken as diagonal,

$W_{ii} = 1/\sigma_i^2$, to weight each projection by its noise variance, though the framework can accommodate more general Σ capturing correlations. The term $R(\mathbf{x})$ represents the negative log-prior on \mathbf{x} (i.e., regularizer), which encourages images with desired properties (e.g., smoothness or sparsity in some domain (Dutta et al., 2021a)) and λ controls the penalty weight. The goal is to find $\hat{\mathbf{x}}$ that minimizes (1), balancing fidelity to the data with prior-based denoising. This optimization is challenging because N is large (high-dimensional images) and $R(\mathbf{x})$ is often nonconvex or implicitly defined by learned models. By selecting an appropriate regularization function, various proximal operator-based iterative schemes have been extensively studied to efficiently solve the MAP estimation problem (Chan et al., 2017). We address this using a model-based iterative reconstruction approach combined with a modern denoising prior, as described next.

2.2. Plug-and-Play (PnP) ADMM Framework

To solve (1), we adopt the PnP ADMM scheme, which enables the use of advanced denoisers as priors without requiring an explicit analytical form for $R(\mathbf{x})$. The PnP framework was first introduced by Venkatakrishnan et al. (Venkatakrishnan et al., 2013) and has since been successfully applied to CT reconstruction problems (He et al., 2019). It leverages the ADMM to split the problem into manageable substeps for data fidelity and denoising. We introduce an auxiliary variable \mathbf{v} to decouple the prior term from the data term:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}, \mathbf{v}} \frac{1}{2} \|\mathbf{Ax} - \mathbf{y}\|_{\mathbf{W}}^2 + \lambda R(\mathbf{v}) \quad \text{s.t.} \quad \mathbf{x} = \mathbf{v} . \quad (2)$$

This constrained formulation is equivalent to (1) with the appropriate $R(\cdot)$, but allows separate handling of data fidelity and prior terms. The augmented Lagrangian for (2) is:

$$\mathcal{L}(\mathbf{x}, \mathbf{v}, \mathbf{u}) = \frac{1}{2} \|\mathbf{Ax} - \mathbf{y}\|_{\mathbf{W}}^2 + \lambda R(\mathbf{v}) + \frac{\beta}{2} \|\mathbf{x} - \mathbf{v} + \mathbf{u}\|_2^2 , \quad (3)$$

where, $\beta > 0$ is the penalty parameter and \mathbf{u} is the scaled dual variable (Lagrange multiplier) for the constraint $\mathbf{x} = \mathbf{v}$. Under the ADMM framework, the augmented Lagrangian in eq. (3) decomposes into three distinct subproblems:

$$\mathbf{x}^{k+1} = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{Ax} - \mathbf{y}\|_{\mathbf{W}}^2 + \frac{\beta}{2} \|\mathbf{x} - \mathbf{v}^k + \mathbf{u}^k\|_2^2 , \quad (4a)$$

$$\mathbf{v}^{k+1} = \arg \min_{\mathbf{v}} \lambda R(\mathbf{v}) + \frac{\beta}{2} \|\mathbf{x}^{k+1} + \mathbf{u}^k - \mathbf{v}\|_2^2 , \quad (4b)$$

$$\mathbf{u}^{k+1} = \mathbf{u}^k + \mathbf{x}^{k+1} - \mathbf{v}^{k+1} . \quad (4c)$$

Eq. (4a) is a quadratic minimization that updates the image to enforce data fidelity. Setting the gradient to zero yields the normal-equation system:

$$(\mathbf{A}^T \mathbf{WA} + \beta I) \mathbf{x}^{k+1} = \mathbf{A}^T \mathbf{Wy} + \beta(\mathbf{v}^k - \mathbf{u}^k) , \quad (5)$$

is solved for \mathbf{x}^{k+1} using a Cholesky-based approximation for computational efficiency (Yeung et al., 2025). When \mathbf{W} includes ramp-filter weighting, $\mathbf{A}^T \mathbf{Wy}$ corresponds to an FBP image, which can serve as an initial guess or one-step approximation for (4a), accelerating convergence. This physics-guided integration ensures that each update remains consistent with the measured sinogram data.

Eq. (4b) is the prior enforcement step, which typically has no closed-form solution if $R(\mathbf{v})$ is a complex learned prior. In the PnP methodology, we replace this optimization with an equivalent denoising operation: instead of explicitly minimizing (4b), we set

$$\mathbf{v}^{k+1} = \mathcal{D}_\sigma(\mathbf{x}^{k+1} + \mathbf{u}^k) , \quad (6)$$

where $\mathcal{D}_\sigma(\cdot)$ is a learned image denoiser (with implicit parameter σ controlling its noise reduction strength) that acts as the proximal operator for R . In other words, \mathcal{D}_σ plays the role of $\text{prox } \lambda R / \beta$, producing a denoised image \mathbf{v}^{k+1} that balances staying close to $\mathbf{x}^{k+1} + \mathbf{u}^k$ with having desirable image properties encoded by the prior. The dual update (4c) then adjusts \mathbf{u} to enforce consistency between \mathbf{x} and \mathbf{v} .

This ADMM-based splitting enables integration of advanced denoisers within a rigorous iterative framework. PnP-ADMM offers convergence guarantees under mild conditions (Chan et al., 2017), and our implementation exhibits stable behavior similar to prior LDCT studies (He et al., 2019). Penalty β and denoiser strength σ are tuned to ensure diminishing residuals across iterations. Overall, the framework combines physics-based data consistency with learned priors for robust ULDCCT reconstruction.

2.3. Denoiser Training Protocol

The performance of PnP-MBIR depends on the denoiser prior \mathcal{D}_σ , built on a residual dense network (RDN) backbone (Zhang et al., 2018) with attention modules (Woo et al., 2018) to emphasize informative features like edges and soft-tissue textures (Ma et al., 2021). It comprises residual dense groups with interleaved attention layers and skip connections for learning residual noise maps. We adopt an image-domain residual learning strategy where the network predicts noise rather than mapping low-dose to high-dose directly, focusing on structured artifacts with lower dynamic range and distinct patterns (Wagner et al., 2023). Following the RED-CNN paradigm (Chen et al., 2017), the predicted residual is added to the input, enabling suppression of stochastic and structured noise while preserving diagnostic details. However, training \mathcal{D}_σ is challenging due to the absence of clean ground-truth CT data for low-dose scans. To overcome this, we employ a 2-stage N2N strategy inspired by Lehtinen et al. (Lehtinen et al., 2018), tailored for ultra-low-dose conditions.

Stage 1 – Criss-cross N2N pre-training: The denoiser is first trained in a self-supervised manner using pairs of independently simulated noisy data. Two low-dose reconstructions or FBP images of the same object are used—one as input, the other as target—minimizing their discrepancy drives the network toward the underlying noise-free signal. To increase diversity, input and target roles are randomly swapped (criss-cross strategy) (Calvarons, 2021), doubling the dataset and exposing the network to both noise-to-noise directions. This approach expands training data and removes the need for ground-truth high-dose images. The networks in this stage are trained to minimize a combined mean-squared error (MSE) and structural similarity (SSIM) loss between the network output and the target noisy measurement, $\mathcal{L}_{\text{stage1}} = \mathcal{L}_{\text{MSE}} + \lambda_{\text{SSIM}} \mathcal{L}_{\text{SSIM}}$, with λ_{SSIM} as tunable weight. Since the true signal is the only consistent component across noise realizations, this stage promotes meaningful noise suppression. By leveraging diverse noise patterns in a criss-cross manner, the denoiser becomes robust to mild and severe noise and learns realistic LDCT artifacts (structured, textured, and correlated), unlike synthetic additive noise.

Stage 2 – Supervised fine-tuning: After pre-training, networks are fine-tuned on paired low-dose and high-dose data, predicting residual maps from low-dose inputs to high-dose outputs. Initialized with Stage 1 weights, training uses a hybrid loss combining \mathcal{L}_{MSE} for general denoising with terms that improve residual accuracy. This step calibrates denoisers to dose-specific noise (correlated and structured artifacts) and corrects Stage 1 bias, mitigating over-smoothing while preserving texture learned earlier and driving outputs closer to high-dose quality in ultra-low-dose conditions.

The denoiser is optimized using a composite loss $\mathcal{L}_{\text{stage2}}$ balancing fidelity, texture realism, and soft-tissue contrast (Langoju et al., 2024). To handle structured streaks from photon starvation—often misinterpreted as true features and varying with dose—additional loss terms are introduced alongside \mathcal{L}_{MSE} .

Entropy-based texture regularization: This term encourages the retention of anatomical texture by maximizing output entropy, preventing over-smoothing. It is defined as the negative Shannon entropy of local patches: $\mathcal{L}_{\text{tex}} = -\sum_{i \in \Omega} H(\mathcal{D}_{\sigma}(v_{\text{low}})_i)$, where $H(\cdot)$ denotes entropy in the neighborhood of pixel i (Abid et al., 2025) and $\mathcal{D}_{\sigma}(v_{\text{low}})$ is the denoised low-dose image v_{low} . Minimizing $-H$ forces the network to maximize the output entropy, thus retaining a level of randomness/texture similar to that of real high-dose images instead of producing overly smooth results. This leverages the fact that entropy is a quantitative measure of image complexity and texture richness (Abid et al., 2025). By preserving entropy, the model is discouraged from eliminating fine stochastic details that characterize soft tissues, thereby maintaining realistic textures.

Soft-tissue contrast-aware loss: To protect low-contrast anatomical boundaries (e.g., liver-kidney or gray-white matter), we employ a contrast-aware term that emphasizes soft-tissue intensity using a weighted SSIM loss: $\mathcal{L}_{\text{ca}} = \text{SSIM}(w_j \mathcal{D}_{\sigma}(v_{\text{low}})_j, w_j v_j^{\text{HD}})$, where w_j increases for HU values typical of soft tissues and v^{HD} is the high-dose target, placing more importance on subtle differences (Liu et al., 2020). This helps preserve diagnostically relevant details vulnerable to suppression under pure MSE training in soft tissue – effectively acting like a perceptual loss focusing on low-contrast detail. Thus, the total image-domain loss becomes: $\mathcal{L}_{\text{stage2}} = \mathcal{L}_{\text{MSE}} + \lambda_{\text{tex}} \mathcal{L}_{\text{tex}} + \lambda_{\text{ca}} \mathcal{L}_{\text{ca}}$, with λ_{tex} and λ_{ca} balancing the trade-off between fidelity and perceptual quality. This combined objective helps the network suppress noise at different dose-levels while retaining texture and contrast in soft-tissue structures.

3. Experimental Setup and Results

3.1. Data Generation

We generated a diverse LDCT dataset using the CatSim X-ray CT simulation platform (Zhang et al., 2024). Multiple anatomically realistic anthropomorphic digital 3D phantoms (e.g., thorax, abdomen, chest, head) were scanned under varied acquisition conditions. The simulated scanner employed a fan-beam geometry with axial scanning and a detector array comprising 900 columns and 32 rows. Complete coverage of each phantom was achieved by varying the starting z-position during axial scans. To emulate dose reduction, tube current and tube voltage were systematically varied (e.g., 20–200 mA, 80–120 kVp), where reduced X-ray flux resulted in increased image noise. The acquired projection data (sinograms) were reconstructed using standard FBP with a slice thickness of 0.625 mm to produce LDCT images, which served as inputs for subsequent experiments. By varying mA and

kVp, the simulated dataset spans a broad range of noise levels and spectral characteristics, including correlated and structured artifacts. A high-dose scan at 800 mA was designated as the ground-truth reference. Furthermore, a scan at 800mA serves as the ground-truth reference. This range captures clinical low-dose conditions while providing a high-dose baseline for supervision. This diversity addresses the data-scarcity problem noted in DL reconstruction (He et al., 2019) and ensures that the PnP framework remains robust when deployed on real patient data. Additionally, Appendix. A and Appendix. B detail the training setup and integration into the PnP-MBIR reconstruction framework, respectively.

3.2. Results and Evaluation

Fig. 1 qualitatively compares the reconstructions from FBP, supervised denoising, 2-stage N2N denoising, PnP-MBIR with a supervised prior, and our proposed PnP-MBIR with 2-stage (N2N) denoiser for dose-levels 100 mA, 70 mA and 40 mA respectively in the top,




Figure 1: Qualitative comparison of different reconstruction approaches. Each row shows a different anatomical image: the top three rows correspond to simulated datasets acquired at 100 mA, 70 mA, and 40 mA respectively with 120 kVp, and the bottom row shows an *in-vivo* acquisition at 55 mA with 100 kVp. Columns (from left to right) show reconstructions from: FBP, image-domain supervised denoising, 2-stage N2N denoising, PnP-MBIR with supervised prior, and our proposed PnP-MBIR using the 2-stage (N2N) denoiser. The final column shows the ultra-high-dose ground truth image for reference (not available for the *in-vivo* case). The proposed method clearly suppresses structured artifacts and preserves fine anatomical details. All images are displayed with [WL 30, WW 300].

second and third rows. Finally a *in-vivo* data, acquired with 55 mA and 100 kVp, are presented in the bottom row. As expected, the low-dose FBP images (left column) are severely degraded by noise and streak artifacts. The standalone supervised denoiser (column 2) removes most of the random noise but leaves substantial structured streak artifacts and creates distortion and anomalies in the underlying anatomical structure (see blue arrows and blue boxes), since conventional supervised training framework cannot effectively learn to remove correlated streaks and noise. The standalone 2-stage N2N denoiser (column 3) further suppresses structured artifacts, but it introduces noticeable blurring and loss of fine details (see blue arrows and blue boxes). PnP reconstruction using the supervised denoiser as a prior (column 4) improves on image-domain denoising – some low-contrast features become clearer – but still leaves visible correlated artifacts in complex regions (e.g. liver, inter-tissue edges), resulting distorted reconstructions (magenta box) and loss of fine details (blue box). In contrast, the proposed PnP-MBIR reconstruction with the 2-stage N2N denoiser (column 5) yields the best visual quality: streaks and structured noise are substantially removed while high-frequency anatomical details (highlighted by blue boxes and arrows) are preserved. This output closely matches the ultra-high-dose reference (column 6) and, in practice, even exceeds it by eliminating subtle artifacts present in the ground truth reference (red arrows) due to photon starvation and attenuation in dense bone regions. A similar pattern holds *in-vivo* (bottom row): our method retains sharp bone edges and soft-tissue contrast (blue box and arrow) far better than other low-dose methods. In addition to noise and artifact reduction, our method maintains the subtle contrast between soft and hard bone tissues, ensuring better visualization of intra-bone structures (green boxes). Additional zoomed-in sections of the reconstructed images are provided in Appendix C for a comprehensive visual analysis.

Performance is quantified using texture and histogram-based metrics. We compute gray-level co-occurrence matrix (GLCM) features—contrast, entropy, energy, homogeneity, correlation, dissimilarity, and angular second moment (ASM)—to capture textural content and smoothness. Additionally, Earth Mover’s Distance (EMD) is calculated between reconstructed and ground-truth intensity histograms. All evaluations are performed on simulated cases with known ground truth under consistent hyperparameters: GLCM with distance = 1, angles = $\{0, \frac{\pi}{4}, \frac{\pi}{2}, \frac{3\pi}{4}\}$, and 256 gray levels; EMD using 256-bin normalized histograms.

The quantitative GLCM texture analysis confirms these observations. The FBP baseline has the highest contrast (1314.19) and lowest texture uniformity (energy 0.0162, homogeneity 0.0429), reflecting severe noise and artifacts. The relative changes in the GLCM measures against the FBP are summarized in Fig. 2. Compared to FBP, all methods significantly reduce contrast and increase uniformity. The ground-truth high-dose image already reduces contrast by 69.2% (to 405.08) and boosts energy by 46.7%. Our PnP-MBIR with 2-stage N2N denoiser achieves the largest gains: 93.1% contrast reduction (to 356.08), 113.1% increase in energy (to 0.0248), and 277.7% increase in homogeneity (to 0.0736). Dissimilarity (texture variation) is similarly minimized (14.76, 75.5% improvement), and the ASM (texture order) is maximized (0.0006, 357.2% improvement). All reconstruction methods yield high GLCM correlation (≥ 0.75), indicating good structural fidelity, with our method achieving 0.7534 (234.1% improvement over FBP). Entropy (information content) is reduced in all cases (by 10–13%), with our proposed PnP-MBIR with N2N reaching the largest drop, suggesting effective noise removal without loss of detail. Finally, the EMD from the FBP

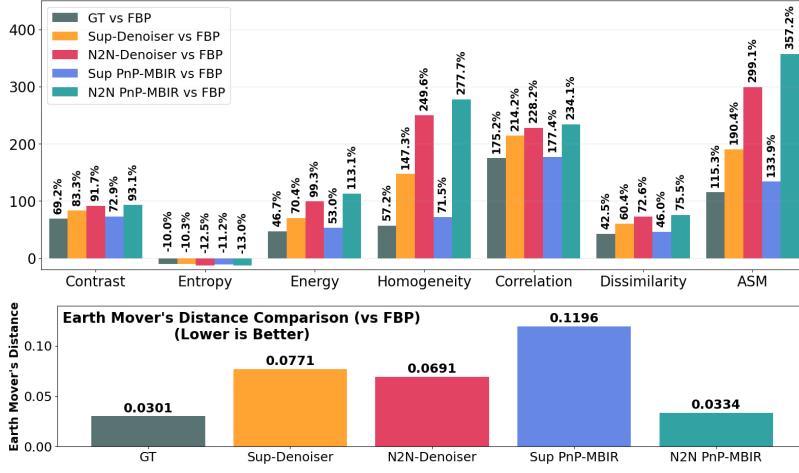


Figure 2: Quantitative texture analysis using GLCM and Earth Mover’s Distance (EMD). The top plot shows relative improvement (percentage change) in key GLCM features—contrast, energy, homogeneity, correlation, dissimilarity, angular second moment (ASM), and entropy—compared to FBP. The bottom plot reports EMD values indicating distributional similarity between each reconstruction and the ground truth. Our proposed PnP-MBIR method with 2-stage N2N denoiser consistently yields the best or near-best performance across all metrics.

reference is lowest for our PnP-MBIR (0.0334), indicating its intensity distribution most closely matches the reference. In summary, our 2-stage N2N PnP-MBIR reconstruction attains the most uniform textures, minimal noise/artifacts, and highest structural similarity among the tested methods.

4. Discussion and Conclusion

The GLCM texture analysis and qualitative comparisons together demonstrate the clear superiority of the proposed 2-stage PnP-MBIR framework. Integrating a physics-based iterative reconstruction with our learned denoiser yields consistently better image quality than either approach alone. On average, the proposed PnP-MBIR N2N method achieved the largest improvements across all texture metrics (an average of $\sim 162.5\%$ improvement over FBP baseline, Fig. 2). In particular, it achieved 93.1% noise (contrast) reduction, 113.1% energy increase, and a remarkable 357.2% increase in ASM, far surpassing the gains of standalone denoisers. These results underscore the synergy of model-based reconstruction and learned priors: the iterative use of the learned denoiser exploits knowledge of the CT physics to remove noise beyond what image-domain denoising can achieve. This is consistent with prior findings that PnP-MBIR can significantly reduce noise and artifacts compared to standard MBIR, and better preserve texture than deep denoising alone (Xu et al., 2022).

Texture and structure preservation are also significantly improved. All compared methods yield high correlation (>0.75 , Fig. 2) with the reference, indicating that anatomical structures are largely preserved after denoising. Notably, the PnP framework particularly enhances texture fidelity: as observed by Xu et al. (Xu et al., 2022), PnP reconstructions

preserve native image texture better than standalone CNN denoising. In our results, the PnP-MBIR with 2-stage N2N output maintains subtle soft-tissue contrast and sharp bone boundaries that are blurred in other outputs (shown with blue and green boxes in Fig. 1), which is reflected in its high homogeneity and low dissimilarity. The consistent entropy reduction (around 10–13%, Fig. 2) across methods implies that the denoisers remove random noise without discarding diagnostically relevant information. Moreover, the proposed PnP-MBIR N2N method exhibits the lowest EMD to the high-dose reference, indicating that its intensity histogram and texture patterns most closely match the “ideal” image. Together, the quantitative metrics and visual evidence show that our approach yields statistically cleaner, more uniform images while retaining clinically important features.

These improvements have direct clinical implications. Enhanced contrast and uniformity in ultra-low-dose reconstructions can improve the detectability of subtle lesions and anatomical structures, potentially enabling further dose reduction in screening protocols. As a result, the proposed PnP-MBIR method substantially reduces noise and streak artifacts, enabling high-quality reconstructions at ~ 70 – 80% lower radiation dose that are comparable to standard full-dose images. Moreover, our method’s ability to preserve soft-tissue contrast and bone detail (e.g. subtle intra-bone structures) should increase radiologist confidence in interpretation. Interestingly, our reconstruction even outperform the nominal “ground truth” high-dose scan in texture metrics (shown with red arrows in Fig. 1); this suggests that learned priors can partly overcome limitations of the reference data (e.g. photon starvation artifacts) and achieve more diagnostically valuable images.

Despite its strong performance, the proposed 2-stage PnP-MBIR framework has some limitations. While it significantly enhances image quality over conventional FBP and standalone denoising, the spatial resolution of the reconstructed images remains bounded by the fidelity of the forward model. Future work could explore integrating resolution-enhancing mechanisms directly within the MBIR forward operator to further improve fine structural recovery. Additionally, the current approach leverages deep priors only in the image domain. Extending the framework to incorporate learned priors in the sinogram domain—trained using the same two-stage N2N strategy—could enable more powerful regularization by jointly exploiting complementary information from both domains. A modified dual-domain PnP-MBIR architecture, wherein denoisers operate in parallel on both the sinogram and image spaces, holds promise for further reducing artifacts and enhancing texture fidelity.

In conclusion, the proposed dual-domain PnP-MBIR framework offers a robust and generalized framework for ULDCT reconstructions. By embedding a self-supervised N2N deep denoiser as a prior within an ADMM-based iterative reconstruction, our method achieves substantial suppression of both stochastic noise and structured noise and artifacts while preserving soft-tissue contrast and fine anatomical detail. Experimental results on both simulated phantoms and *in-vivo* patient data demonstrate that the reconstructed images exhibit markedly improved texture quality across GLCM-based metrics, significantly outperforming conventional FBP and image-domain DL approaches in preserving anatomical detail and reducing structured noise. Qualitatively, the approach effectively eliminates streaks and mottling noise without blurring edges or subtle features, closely matching the appearance of full-dose CT. These findings demonstrate that the 2-stage N2N-trained PnP prior can reliably enable ULDCT by restoring diagnostic image quality without sacrificing clinically relevant structures.

References

- Muqaddas Abid, Muhammad Suzuri Hitam, Rozniza Ali, and Muhammad Hammad. Advancements in texture analysis and classification: A bibliometric review of entropy-based approaches. *International Journal of Advanced Computer Science and Applications*, 16(9), 2025. doi: 10.14569/IJACSA.2025.0160957. URL <http://dx.doi.org/10.14569/IJACSA.2025.0160957>.
- Michael Balda, Joachim Hornegger, and Bjoern Heismann. Ray contribution masks for structure adaptive sinogram filtering. *IEEE Transactions on Medical Imaging*, 31(6):1228–1239, 2012. doi: 10.1109/TMI.2012.2187213.
- C.A. Bouman and K. Sauer. A unified approach to statistical tomography using coordinate descent optimization. *IEEE Transactions on Image Processing*, 5(3):480–492, 1996. doi: 10.1109/83.491321.
- Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011. ISSN 1935-8237. doi: 10.1561/22000000016. URL <http://dx.doi.org/10.1561/22000000016>.
- David J. Brenner and Eric J. Hall. Computed tomography — an increasing source of radiation exposure. *New England Journal of Medicine*, 357(22):2277–2284, 2007. doi: 10.1056/NEJMra072149.
- Jian-Feng Cai, Xun Jia, Hao Gao, Steve B. Jiang, Zuowei Shen, and Hongkai Zhao. Cine cone beam CT reconstruction using low-rank matrix factorization: Algorithm and a proof-of-principle study. *IEEE Transactions on Medical Imaging*, 33(8):1581–1591, 2014. doi: 10.1109/TMI.2014.2319055.
- Adria Font Calvarons. Improved Noise2Noise denoising with limited data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 796–805, 2021.
- Chung Chan, Jian Zhou, Li Yang, Wenyuan Qi, and Evren Asma. Noise to noise ensemble learning for pet image denoising. In *2019 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, pages 1–3, 2019. doi: 10.1109/NSS/MIC42101.2019.9059779.
- Stanley H. Chan, Xiran Wang, and Omar A. Elgendy. Plug-and-play admm for image restoration: Fixed-point convergence and applications. *IEEE Transactions on Computational Imaging*, 3(1):84–98, 2017. doi: 10.1109/TCI.2016.2629286.
- Hu Chen, Yi Zhang, Mannudeep K. Kalra, Feng Lin, Yang Chen, Peixi Liao, Jiliu Zhou, and Ge Wang. Low-dose CT with a residual encoder-decoder convolutional neural network. *IEEE Transactions on Medical Imaging*, 36(12):2524–2535, 2017. doi: 10.1109/TMI.2017.2715284.

- Aanuoluwapo Clement, David B. Olawade, Laura Vanderbloemen, Oluwayomi B. Rotifa, Sandra Chinaza Fidelis, Eghosasere Egbon, Akwaowo Owoidighe Akpan, Sola Adeleke, Aruni Ghose, and Stergios Boussios. Ai-driven advances in low-dose imaging and enhancement—a review. *Diagnostics*, 15(6), 2025. ISSN 2075-4418. doi: 10.3390/diagnostics15060689. URL <https://www.mdpi.com/2075-4418/15/6/689>.
- S. Dutta, A. Basarab, B. Georgeot, and D Kouamé. Plug-and-play quantum adaptive denoiser for deconvolving poisson noisy images. *IEEE Access*, 9:139771–139791, 2021a. doi: 10.1109/ACCESS.2021.3118608.
- S. Dutta, A. Basarab, B. Georgeot, and D. Kouamé. Quantum mechanics-based signal and image representation: Application to denoising. *IEEE Open J. of Signal Process.*, 2: 190–206, 2021b. doi: 10.1109/OJSP.2021.3067507.
- S. Dutta, A. Basarab, B. Georgeot, and D. Kouamé. A novel image denoising algorithm using concepts of quantum many-body theory. *Signal Processing*, 201:108690, 2022a. ISSN 0165-1684. doi: <https://doi.org/10.1016/j.sigpro.2022.108690>.
- S. Dutta, K.T. Nwigbo, J. Michetti, B. Georgeot, D.-H. Pham, A. Basarab, and D. Kouamé. Quantum denoising-based super-resolution algorithm applied to dental tomography images. In *Proc. Int. Symp. Biomed. Imaging (ISBI)*, pages 1–4, 2022b. doi: 10.1109/ISBI52829.2022.9761623.
- S. Dutta, A. Basarab, B. Georgeot, and D. Kouamé. DIVA: Deep unfolded network from quantum interactive patches for image restoration. *Pattern Recognition*, 155:110676, 2024a. doi: <https://doi.org/10.1016/j.patcog.2024.110676>.
- Sayantana Dutta and Jonathan Mamou. A quantum denoising-based resolution enhancement framework for 250-MHz and 500-MHz quantitative acoustic microscopy. *IEEE Transactions on Computational Imaging*, 10:1489–1504, 2024. doi: 10.1109/TCI.2024.3473312.
- Sayantana Dutta and Jonathan Mamou. Enhancing 3-D radio frequency data in quantitative acoustic microscopy using quantum-driven prior at 250 MHz and 500 MHz. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 72(8):1119–1133, 2025. doi: 10.1109/TUFFC.2025.3576239.
- Sayantana Dutta, Kenule Tuador Nwigbo, Jérôme Michetti, Bertrand Georgeot, Duong Hung Pham, Denis Kouamé, and Adrian Basarab. Computed tomography image restoration using a quantum-based deep unrolled denoiser and a plug-and-play framework. In *2023 31st European Signal Processing Conference (EUSIPCO)*, pages 845–849, 2023. doi: 10.23919/EUSIPCO58844.2023.10289892.
- Sayantana Dutta, Bertrand Georgeot, Jérôme Michetti, Adrian Basarab, and Denis Kouamé. Unsupervised physics-inspired deep learning network with application to dental computed tomography image restoration. In *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 1–5, 2024b. doi: 10.1109/ISBI56570.2024.10635786.

- Xiaohong Fan, Ke Chen, Huaming Yi, Yin Yang, and Jianping Zhang. MVMS-RCN: A dual-domain unified CT reconstruction with multi-sparse-view and multi-scale refinement-correction. *IEEE Transactions on Computational Imaging*, 10:1749–1762, 2024. doi: 10.1109/TCI.2024.3507645.
- A. Floquet, S. Dutta, E. Soubies, D.-H. Pham, and D. Kouamé. Automatic tuning of denoising algorithms parameters without ground truth. *IEEE Signal Process. Lett.*, 31: 381–385, 2024. doi: 10.1109/LSP.2024.3354554.
- Lin Fu and Bruno De Man. Deep learning tomographic reconstruction through hierarchical decomposition of domain transforms. *Visual Computing for Industry, Biomedicine, and Art*, 5(1):30, 2022. doi: 10.1186/s42492-022-00127-y.
- Ji He, Yan Yang, Yongbo Wang, Dong Zeng, Zhaoying Bian, Hao Zhang, Jian Sun, Zongben Xu, and Jianhua Ma. Optimizing a parameterized plug-and-play admm for iterative low-dose CT reconstruction. *IEEE Transactions on Medical Imaging*, 38(2):371–382, 2019. doi: 10.1109/TMI.2018.2865202.
- Kyong Hwan Jin, Michael T. McCann, Emmanuel Froustey, and Michael Unser. Deep convolutional neural network for inverse problems in imaging. *IEEE Transactions on Image Processing*, 26(9):4509–4522, 2017. doi: 10.1109/TIP.2017.2713099.
- Masaki Katsura, Izuru Matsuda, Masaaki Akahane, Jiro Sato, Hiroyuki Akai, Koichiro Yasaka, Akira Kunimatsu, and Kuni Ohtomo. Model-based iterative reconstruction technique for radiation dose reduction in chest CT: comparison with the adaptive statistical iterative reconstruction technique. *European Radiology*, 22(8):1613–1623, 2012. doi: 10.1007/s00330-012-2452-z.
- Kyungsang Kim, Jong Chul Ye, William Worstell, Jinsong Ouyang, Yothin Rakvongthai, Georges El Fakhri, and Quanzheng Li. Sparse-view spectral CT reconstruction using spectral patch-based low-rank penalty. *IEEE Transactions on Medical Imaging*, 34(3): 748–760, 2015. doi: 10.1109/TMI.2014.2380993.
- S. Kong, W. Wang, X. Feng, and X. Jia. Deep red unfolding network for image restoration. *IEEE Trans. Image Process.*, 31:852–867, 2022. doi: 10.1109/TIP.2021.3136623.
- Rajesh Langoju, Utkarsh Agrawal, and Bipul Das. Structured noise reduction with texture preservation in low-dose x-ray computed tomography. In *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 1–4, 2024. doi: 10.1109/ISBI56570.2024.10635554.
- Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2Noise: Learning image restoration without clean data. In *International Conference on Machine Learning (ICML)*, pages 2965–2974, 2018.
- Meng Li, William Hsu, Xiaodong Xie, Jason Cong, and Wen Gao. SACNN: Self-attention convolutional neural network for low-dose CT denoising with self-supervised perceptual loss network. *IEEE Transactions on Medical Imaging*, 39(7):2289–2301, 2020. doi: 10.1109/TMI.2020.2968472.

- Qiuyue Liu, Zhen Zhou, Feng Liu, Xiangming Fang, Yizhou Yu, and Yizhou Wang. Multi-stream progressive up-sampling network for dense CT image reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 518–528. Springer, 2020.
- Yin-Jin Ma, Yong Ren, Peng Feng, Peng He, Xiao-Dong Guo, and Biao Wei. Sinogram denoising via attention residual dense convolutional neural network for low-dose computed tomography. *Nuclear Science and Techniques*, 32(4):41, 2021. doi: 10.1007/s41365-021-00874-2.
- Yinjin Ma, Yajuan Zhang, Lin Chen, Qiang Jiang, Fengjuan Shi, and Biao Wei. Deep plug-and-play denoising prior with total variation regularization for low-dose ct. *Frontiers in Physics*, 13, 2025. doi: 10.3389/fphy.2025.1563756. URL <https://www.frontiersin.org/journals/physics/articles/10.3389/fphy.2025.1563756>.
- Armando Manduca, Lifeng Yu, Joshua D. Trzasko, Natalia Khaylova, James M. Kofler, Cynthia M. McCollough, and Joel G. Fletcher. Projection space denoising with bilateral filtering and CT noise modeling for dose reduction in CT. *Medical Physics*, 36(11):4911–4919, 2009. doi: <https://doi.org/10.1118/1.3232004>. URL <https://aapm.onlinelibrary.wiley.com/doi/abs/10.1118/1.3232004>.
- Emil Y. Sidky and Xiaochuan Pan. Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Physics in Medicine & Biology*, 53(17):4777, aug 2008. doi: 10.1088/0031-9155/53/17/021. URL <https://doi.org/10.1088/0031-9155/53/17/021>.
- R. Smith-Bindman, J. Lipson, and R. Marcus. Radiation dose associated with common computed tomography examinations and the associated lifetime attributable risk of cancer. *Journal of Vascular Surgery*, 51(3):783, 2010. doi: 10.1016/j.jvs.2010.01.042.
- S.V. Venkatakrishnan, C.A. Bouman, and B. Wohlberg. Plug-and-play priors for model based reconstruction. In *Proc. IEEE Global Conf. Signal Inf. Process.*, pages 945–948, 2013. doi: 10.1109/GlobalSIP.2013.6737048.
- Fabian Wagner, Mareike Thies, Laura Pfaff, Oliver Aust, Sabrina Pechmann, Daniela Weidner, Noah Maul, Maximilian Rohleder, Mingxuan Gu, Jonas Utz, Felix Denzinger, and Andreas Maier. On the benefit of dual-domain denoising in a self-supervised low-dose CT setting. In *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5, 2023. doi: 10.1109/ISBI53787.2023.10230511.
- Jing Wang, Tianfang Li, Hongbing Lu, and Zhengrong Liang. Penalized weighted least-squares approach to sinogram noise reduction and image reconstruction for low-dose x-ray computed tomography. *IEEE Transactions on Medical Imaging*, 25(10):1272–1283, 2006. doi: 10.1109/TMI.2006.882141.
- Tao Wang, Zexin Lu, Ziyuan Yang, Wenjun Xia, Mingzheng Hou, Huaiqiang Sun, Yan Liu, Hu Chen, Jiliu Zhou, and Yi Zhang. IDOL-Net: An interactive dual-domain parallel network for CT metal artifact reduction. *IEEE Transactions on Radiation and Plasma Medical Sciences*, 6(8):874–885, 2022. doi: 10.1109/TRPMS.2022.3171440.

- Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. CBAM: Convolutional block attention module. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- Dufan Wu, Kyungsang Kim, and Quanzheng Li. Low-dose CT reconstruction with Noise2Noise network and testing-time fine-tuning. *Medical Physics*, 48(12):7657–7672, 2021. doi: <https://doi.org/10.1002/mp.15101>.
- Qifan Xu, Qihui Lyu, Dan Ruan, and Ke Sheng. To what extent can plug-and-play methods outperform neural networks alone in low-dose CT reconstruction. In *2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI)*, pages 1–5, 2022. doi: 10.1109/ISBI52829.2022.9761612.
- Qiong Xu, Hengyong Yu, Xuanqin Mou, Lei Zhang, Jiang Hsieh, and Ge Wang. Low-dose x-ray CT reconstruction via dictionary learning. *IEEE Transactions on Medical Imaging*, 31(9):1682–1697, 2012. doi: 10.1109/TMI.2012.2195669.
- A. Pour Yazdanpanah, Emma E. Regentova, and George Bebis. Algebraic iterative reconstruction-reprojection (airr) method for high performance sparse-view CT reconstruction. *Appl. Math*, 10(6):1–8, 2016.
- Dong Hye Ye, Somesh Srivastava, Jean-Baptiste Thibault, Ken Sauer, and Charles Bouman. Deep residual learning for model-based iterative CT reconstruction using plug-and-play framework. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6668–6672, 2018. doi: 10.1109/ICASSP.2018.8461408.
- Kylie Yeung, Christine Tobler, Rolf F Schulte, Benjamin White, Anthony McIntyre, Sébastien Serres, Peter Morris, Dorothee Auer, Fergus V Gleeson, Damian J Tyler, et al. Algebraic methods and computational strategies for pseudoinverse-based mr image reconstruction (pinv-recon). *Scientific Reports*, 15(1):37997, 2025. doi: 10.1038/s41598-025-21929-z.
- Hua Zhang, Jing Huang, Jianhua Ma, Zhaoying Bian, Qianjin Feng, Hongbing Lu, Zhenrong Liang, and Wufan Chen. Iterative reconstruction for x-ray computed tomography using prior-image induced nonlocal regularization. *IEEE Transactions on Biomedical Engineering*, 61(9):2367–2378, 2014. doi: 10.1109/TBME.2013.2287244.
- Jiayong Zhang, Mingye Wu, Paul FitzGerald, Stephen Araujo, and Bruno De Man. Development and tuning of models for accurate simulation of ct spatial resolution using catsim. *Physics in Medicine & Biology*, 69(4):045014, feb 2024. doi: 10.1088/1361-6560/ad2122. URL <https://doi.org/10.1088/1361-6560/ad2122>.
- Yi Zhang, Yan Wang, Weihua Zhang, Feng Lin, Yifei Pu, and Jiliu Zhou. Statistical iterative reconstruction using adaptive fractional order regularization. *Biomed. Opt. Express*, 7(3):1015–1029, Mar 2016a. doi: 10.1364/BOE.7.001015. URL <https://opg.optica.org/boe/abstract.cfm?URI=boe-7-3-1015>.
- Yi Zhang, Yan Xi, Qingsong Yang, Wenxiang Cong, Jiliu Zhou, and Ge Wang. Spectral CT reconstruction with image sparsity and spectral mean. *IEEE Transactions on Computational Imaging*, 2(4):510–523, 2016b. doi: 10.1109/TCI.2016.2609414.

Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pages 2472–2481, June 2018.

Appendix A. Training Setup and Implementation

Our image-domain denoiser is trained with 512×512 full resolution images. Training uses Adam optimization (initial learning rate 1×10^{-4} with a decay factor 0.99), and typical batch sizes of 8 for 500 epochs. The full-resolution data with five random augmentation followed by the 2-stage N2N training strategy make the denoisers more robust for the ultra-low-dose regime, particularly to learn the structured artifacts. The network is implemented in PyTorch and trained on an NVIDIA RTX 600 Ada generation cluster. Furthermore, for comprehensive analysis of the proposed self-supervised training protocol (2-stage N2N, as described in Section 2.3), we trained the image-based denoisers supervised way (trained on low-dose and high-dose data) under the same dataset and training setup.

Appendix B. Integration into PnP-MBIR Reconstruction

The trained denoiser \mathcal{D}_σ is integrated as a PnP prior within the ADMM framework. At each iteration of the proposed PnP-MBIR scheme, the projection data enforces the measurement fidelity and suppress noise in the Radon domain. This is followed by an image update step where \mathcal{D}_σ refines the reconstruction by imposing natural image priors. This al-




Figure 3: Qualitative comparison of reconstruction methods. The left column shows a simulated FBP image acquired at 40 mA and 120 kVp. Zoomed-in regions (highlighted in red) from the reconstructed outputs are presented for: FBP, image-domain supervised denoising, 2-stage N2N denoising, PnP-MBIR with supervised prior, our proposed PnP-MBIR using the 2-stage (N2N) denoiser, and the ultra-high-dose ground truth image for reference. The proposed approach effectively suppresses structured artifacts while preserving fine anatomical details and soft tissue contrast. All images are displayed with [WL 30, WW 300].

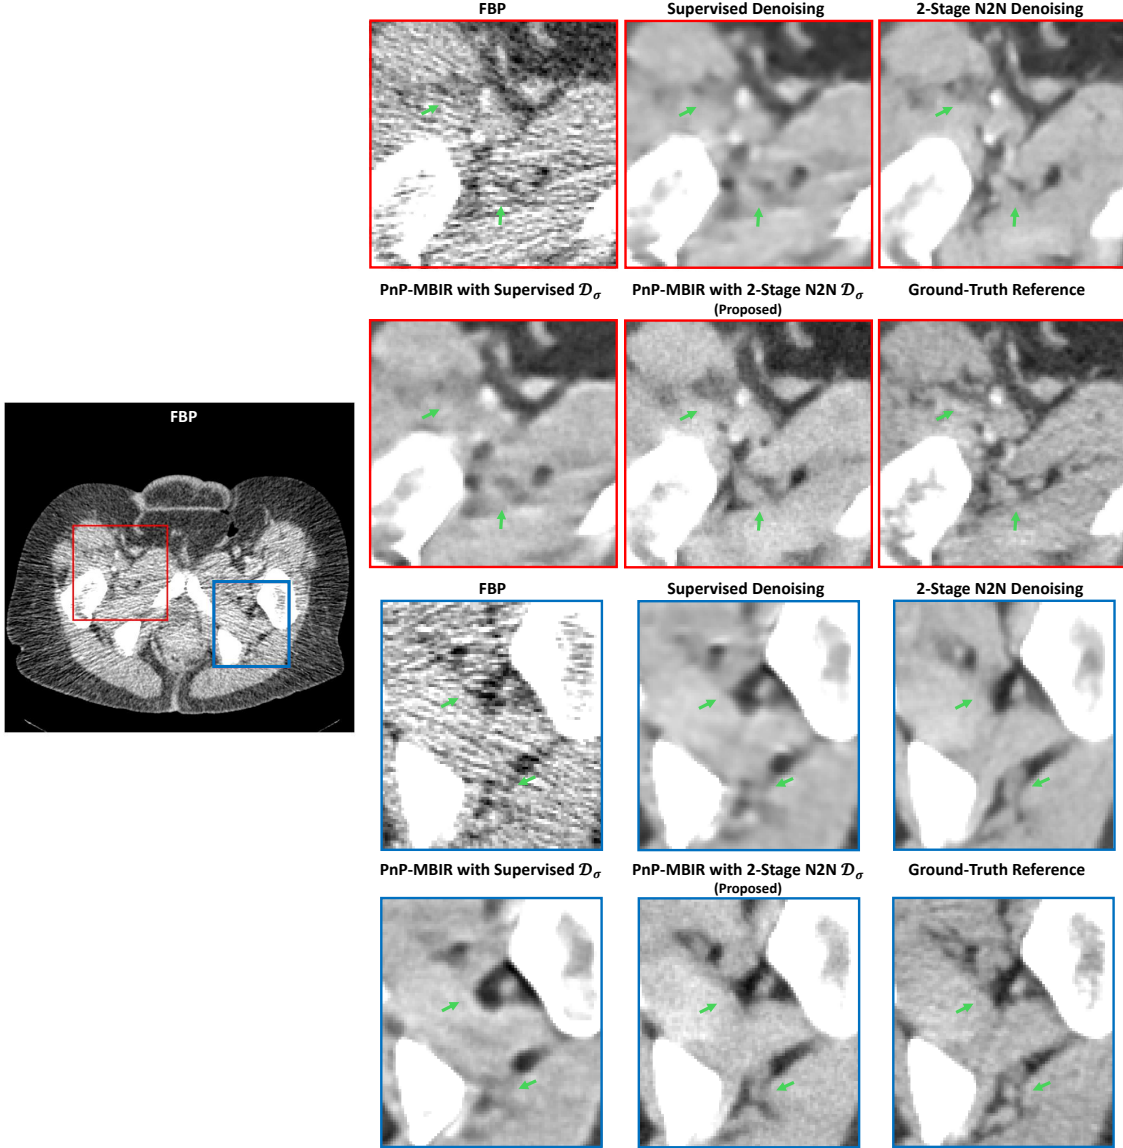


Figure 4: Qualitative comparison of reconstruction methods. The left column shows a simulated FBP image acquired at 100,mA and 120 kVp. Zoomed-in regions (highlighted in red and blue) from the reconstructed outputs are presented for: FBP, image-domain supervised denoising, 2-stage N2N denoising, PnP-MBIR with supervised prior, our proposed PnP-MBIR using the 2-stage (N2N) denoiser, and the ultra-high-dose ground truth image for reference. The proposed approach effectively suppresses structured artifacts while preserving fine anatomical details and soft tissue contrast. All images are displayed with [WL 30, WW 300].

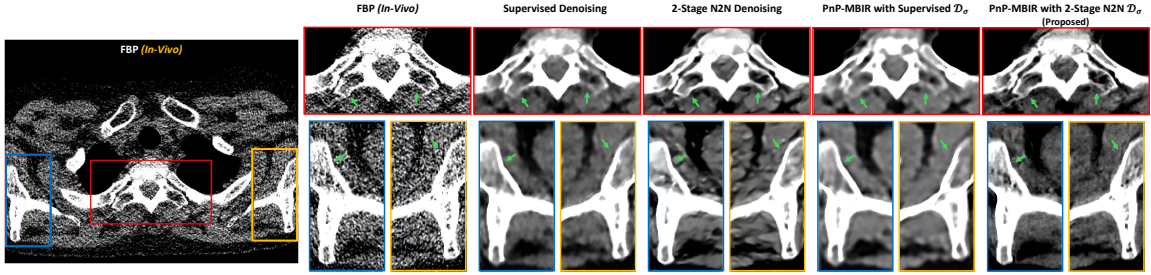


Figure 5: Qualitative comparison of reconstruction methods on an *in-vivo* acquisition. The left column shows an *in-vivo* scan acquired at 55 mA and 100 kVp. Zoomed-in regions (highlighted in red, blue, and yellow) from the reconstructed outputs are presented for: FBP, image-domain supervised denoising, 2-stage N2N denoising, PnP-MBIR with supervised prior, and our proposed PnP-MBIR using the 2-stage (N2N) denoiser. The proposed approach effectively suppresses structured artifacts while maintaining subtle contrast between soft and hard bone tissues, enabling improved visualization of intra-bone structures. All images are displayed with [WL 30, WW 300].

ternating update process effectively regularizes both domains: eq. (4a) ensures consistency with measured projections while removing projection noise, and \mathcal{D}_σ (eq. (4b)) suppresses residual image-domain artifacts. By leveraging the complementary strengths of both steps, the dual-domain PnP-MBIR algorithm achieves high-quality CT reconstructions with significantly reduced noise and preserved anatomical structures, even under ultra-low-dose conditions. The iterative process continues until convergence is reached, as defined by the condition: $\frac{1}{3N}(\|\mathbf{x}^{k+1} - \mathbf{x}^k\|_2^2 + \|\mathbf{v}^{k+1} - \mathbf{v}^k\|_2^2 + \|\mathbf{u}^{k+1} - \mathbf{u}^k\|_2^2) \leq 10^{-4}$ (Chan et al., 2017). Empirically, we found that the proposed model typically converges within five iterations when using $\beta = 1.2$.

Appendix C. More Experimental Results

Figs. 3–5 present zoomed-in sections from the reconstructed images obtained using different methods for qualitative evaluation. From the regions highlighted in Figs. 3–4, it is evident that our proposed PnP-MBIR reconstruction with the 2-stage N2N denoiser delivers superior visual quality. Specifically, streak artifacts and structured noise are significantly suppressed, while high-frequency anatomical details—indicated by green arrows—are well preserved. The resulting image closely approximates the ultra-high-dose reference, demonstrating the effectiveness of our approach. A similar trend is observed in the *in-vivo* results shown in Fig. 5 (zoomed-in red, blue, and yellow regions). Our method maintains sharp bone edges and soft-tissue contrast (blue and yellow sections) far better than competing low-dose reconstruction techniques. Beyond noise and artifact reduction, the proposed approach preserves subtle contrast between soft and hard bone tissues, enabling improved visualization of intra-bone structures (highlighted by green arrows).