Journal of Visual Computing

https://jvcjournal.org

Generative Image Priors for Artifact Suppression in Low-Dose CT Scans: Simulation Study

Anna Leppänen*

Department of Computer Science, Faculty of Science, University of Helsinki, Helsinki, Uusimaa, Finland

Corresponding author: anna.leppanen@helsinki.fi

Abstract

The widespread utilization of computed tomography in clinical diagnostics has brought immense benefits to patient care, yet the associated ionizing radiation poses non-negligible long-term health risks. Consequently, the transition toward low dose computed tomography has become an urgent clinical imperative, guided by the principle of keeping radiation doses as low as reasonably achievable. However, reducing the radiation dose intrinsically leads to severe degradation in image quality, primarily manifested as excessive quantum mottle, electronic noise, and complex streak artifacts resulting from photon starvation. Traditional analytical and iterative reconstruction methods often struggle to strike an optimal balance between aggressive noise reduction and the preservation of crucial anatomical details. In recent years, data-driven approaches, particularly those leveraging deep learning, have demonstrated remarkable potential in addressing these inverse problems. Among these, generative image priors derived from unsupervised generative models have emerged as highly promising tools, offering powerful regularization capabilities without relying on strictly paired training data. This paper presents a comprehensive simulation study investigating the efficacy of generative image priors in suppressing artifacts and restoring image fidelity in low dose computed tomography scans. By simulating a rigorous forward projection model incorporating realistic noise characteristics, we evaluate an iterative reconstruction framework guided by a score-based generative prior. Extensive quantitative and qualitative analyses demonstrate that the proposed method significantly outperforms conventional techniques in mitigating streak artifacts, preserving high-frequency structural edges, and maintaining natural image textures.

Keywords: Generative Image Priors; Artifact Suppression; Low-Dose CT; Image Reconstruction; Simulation Study

1. Introduction

1.1 Motivation and Clinical Context

The introduction of computed tomography into clinical practice marked a revolutionary milestone in medical imaging, providing unprecedented cross-sectional views of the human body and transforming diagnostic paradigms across multiple medical disciplines. Over the decades, continuous technological advancements have led to faster acquisition times, higher spatial resolutions, and broader clinical applications, ranging from routine anatomical screening to complex image-guided interventions. However, the diagnostic power of computed tomography is inextricably linked to the use of ionizing radiation. The cumulative exposure to such radiation has raised significant public health concerns, as extensive epidemiological studies have established a correlation between medical radiation exposure and an increased lifetime risk of radiation-induced malignancies [1]. In response to these concerns, the global medical community, alongside regulatory bodies, has strongly advocated for the optimization of radiation protocols, widely adopting the principle of maintaining radiation doses as low as reasonably achievable. This principle dictates that diagnostic images must be acquired using the lowest possible radiation exposure while ensuring that the resulting image quality remains sufficient for accurate clinical interpretation. Consequently, low dose computed tomography has emerged as a standard of care, particularly in scenarios requiring repeated examinations, pediatric imaging, and widespread screening programs such as those for lung cancer. Despite its clear clinical necessity, the physical reality of low dose computed tomography presents a formidable challenge: a reduction in the X-ray tube current or voltage inherently decreases the number of photons reaching the detector [2]. This reduction in photon flux leads to a fundamentally poorer signal-to-noise ratio, severely degrading the diagnostic utility of the reconstructed images. Addressing this challenge requires sophisticated computational strategies capable of recovering high-fidelity diagnostic information from highly degraded measurement data.

1.2 Challenges in Low Dose Computed Tomography

The degradation of image quality in low dose computed tomography is a multifaceted problem rooted in the physics of X-ray attenuation and the limitations of detector technology. When the radiation dose is substantially reduced, the statistical fluctuations in the number of X-ray photons reaching the detector become highly prominent. This phenomenon, commonly referred to as quantum noise, follows a Poisson distribution and represents the primary source of image degradation in low dose scenarios. Furthermore, the inherent electronic noise of the data acquisition system, which can be modeled as Gaussian noise, begins to dominate the signal in regions where photon counts are critically low. The combination of Poisson and Gaussian noise sources leads to a complex, non-stationary noise profile in the acquired projection data [3]. Beyond simple random noise, low dose computed tomography is uniquely susceptible to severe structured artifacts, most notably streak artifacts. These artifacts occur due to photon starvation, a condition where highly attenuating structures within the patient absorb nearly all incident X-ray photons along specific projection paths. The resulting near-zero photon counts at the detector produce extreme outliers in the logarithmic projection data, which, upon reconstruction via standard algorithms, manifest as intense dark and bright streaks radiating across the image. These streak artifacts not only obscure adjacent soft tissues but can also mimic or mask pathological lesions, severely compromising diagnostic accuracy. Traditional reconstruction techniques, such as filtered back projection, amplify high-frequency noise due to their reliance on a high-pass ramp filter, rendering them highly inadequate for low dose data [4]. While post-processing filters can smooth out noise, they inevitably blur critical structural boundaries and compromise the spatial resolution necessary for detecting small abnormalities. Therefore, overcoming these specific physical and computational challenges remains a central focus of modern medical imaging research.

1.3 Objectives and Contributions of the Study

The primary objective of this simulation study is to systematically investigate and validate the integration of advanced generative image priors into the iterative reconstruction pipeline for low dose computed tomography. By framing the image reconstruction process as an ill-posed inverse problem, we aim to demonstrate how prior knowledge captured by state-of-the-art generative models can effectively constrain the solution space, yielding reconstructed images that are both visually natural and diagnostically accurate. The contributions of this work are manifold. First, we design a comprehensive and highly realistic simulation framework that accurately models the forward projection process, incorporating complex non-stationary noise characteristics that mimic real-world low dose clinical environments [5]. Second, we introduce an unsupervised reconstruction methodology that leverages score-based generative models to encapsulate the underlying distribution of high-quality, routine-dose computed tomography images. By functioning as a powerful regularization term, this generative prior guides the optimization process toward artifact-free solutions. Third, we present a detailed comparative analysis against conventional artifact suppression techniques, utilizing a robust suite of quantitative image quality metrics and qualitative visual assessments. This study specifically emphasizes the algorithm's capability to suppress photon starvation streak artifacts while simultaneously preserving fine anatomical textures, a persistent bottleneck in current deep learning-based image reconstruction methods. Through this comprehensive simulation study, we aim to provide robust theoretical and empirical evidence supporting the clinical translation of generative image priors in low-dose medical imaging applications.

2. Background and Literature Review

2.1 Evolution of Artifact Suppression Techniques

The pursuit of high-quality image reconstruction from noisy projection data has driven decades of innovation in the field of computed tomography. Historically, the filtered back projection algorithm has been the undisputed standard for image reconstruction due to its computational efficiency and robust mathematical foundation in the Radon transform. However, the fundamental assumption of filtered back projection requires projection data with a high signal-to-noise ratio and dense angular sampling [6]. In low dose scenarios, the violation of these assumptions leads to the catastrophic amplification of quantum noise and the generation of severe streak artifacts. To mitigate these limitations, the research community shifted its focus toward statistical iterative reconstruction methods. These methods mathematically model the statistical properties of the measurement noise and the physical characteristics of the imaging system. By formulating an objective function that balances data fidelity with a regularization penalty, iterative reconstruction algorithms iteratively update the image estimate to minimize the discrepancy between the forward-projected estimate and the actual measured data. Early regularization techniques relied on simple mathematical assumptions about the image, such as total variation, which assumes that medical images consist of piecewise constant regions separated by sharp boundaries. While total variation is highly effective at smoothing uniform regions and preserving prominent edges, it notoriously suffers from the staircasing effect, imparting an artificial, blocky, and plasticky appearance to the reconstructed images that radiologists find diagnostically unappealing [7]. To overcome this, more sophisticated patch-based priors, such as dictionary learning, were developed. Dictionary learning algorithms adaptively learn a set of overcomplete basis functions directly from the image data, allowing for a more flexible and robust representation of complex anatomical textures. Despite their improved performance over total variation, traditional iterative reconstruction and dictionary learning methods demand immense computational resources and require careful, often empirical, tuning of hyperparameters to balance noise suppression and texture preservation.

2.2 Emergence of Deep Learning in Image Reconstruction

The advent of deep learning, particularly the proliferation of convolutional neural networks, has instigated a paradigm shift in the domain of computed tomography artifact suppression and image reconstruction. The extraordinary capacity of deep neural networks to learn complex, non-linear mappings from vast amounts of data has been extensively harnessed to address the limitations of conventional iterative methods. The initial wave of deep learning applications in this field focused primarily on image domain post-processing. In this approach, convolutional neural networks were trained in a supervised manner to map low dose, noisy reconstructed images directly to their high-quality, routine-dose counterparts. While these post-processing networks achieved remarkable reductions in standard noise metrics and provided exceptionally fast inference times, they frequently exhibited a critical flaw: the tendency to produce overly smoothed images [8]. Because these networks typically optimized a mean squared error loss function, they invariably learned to output the statistical average of possible clean images, thereby blurring out subtle, high-frequency diagnostic details and fine textures. Furthermore, image-domain only methods lack access to the original projection data, limiting their ability to recover information lost during the initial, flawed analytical reconstruction. To address this, subsequent research explored dual-domain networks and unrolled iterative networks. Dual-domain approaches apply convolutional neural networks in both the projection domain, to repair corrupted sinogram bins before reconstruction, and the image domain, for final refinement [9]. Unrolled networks, on the other hand, integrate convolutional neural networks directly into the iterative reconstruction loop, replacing traditional mathematical regularizers with learned, data-driven regularizers. While these advanced architectures have significantly improved edge preservation and artifact reduction, they predominantly rely on supervised learning, necessitating massive datasets of perfectly paired low dose and routine dose scans. In clinical reality, acquiring perfectly aligned pairs of scans from the same patient under identical physiological conditions is practically impossible, thus restricting the scalability and generalizability of supervised deep learning models in computed tomography reconstruction.

2.3 Generative Image Priors and Unsupervised Learning

To circumvent the strict requirement for paired training data inherent in supervised learning, researchers have increasingly turned toward unsupervised learning frameworks and generative models. Generative image priors represent a conceptually elegant approach to the inverse problem of image reconstruction. Instead of learning a direct mapping from noisy to clean images, generative models are trained exclusively on a large, unpaired dataset of high-quality images to learn the underlying probability distribution of natural anatomical structures. Once trained, the generative model acts as a powerful mathematical prior that penalizes reconstructions deviating from the learned distribution of realistic diagnostic images. Early explorations in this domain utilized generative adversarial networks, where a generator network learned to synthesize realistic images while a discriminator network evaluated their authenticity. In the context of artifact suppression, the discriminator could be employed as an adversarial loss term to enforce textural realism [10]. However, generative adversarial networks are notoriously unstable during training and are prone to mode collapse, where the generator produces a limited variety of outputs, potentially hallucinating anatomical structures that do not exist in the patient data. More recently, diffusion models, particularly score-based generative models, have revolutionized the field of unsupervised image generation. Score-based models learn the gradient of the log probability density of the data distribution, a quantity known as the score function [11]. By learning to denoise images perturbed with varying levels of Gaussian noise, these models capture highly detailed and structurally complex image priors without the instability of adversarial training. During the reconstruction of a low dose computed tomography scan, the pre-trained score network can be utilized to iteratively push the noisy, artifact-ridden image estimate toward the manifold of high-quality images. This score-matching approach provides a robust, principled, and unsupervised mechanism to inject highly realistic anatomical priors into the iterative reconstruction process, ensuring superior artifact suppression while strictly maintaining fidelity to the measured projection data.

3. Methodology and Simulation Framework

3.1 Simulation Design and Data Acquisition

The foundation of this study relies on a meticulously designed simulation framework intended to replicate the complex physical processes and data acquisition mechanisms of a modern medical computed tomography scanner. Utilizing synthetic forward projection, we generate highly realistic raw projection data, commonly referred to as sinograms, from high-resolution digital anatomical phantoms. To ensure clinical relevance, we employ a diverse set of computational phantoms that accurately represent the intricate anatomical variations, tissue densities, and complex structural boundaries found in human thoracic and abdominal regions. The forward projection process is mathematically modeled based on the fundamental principles of the X-ray transform, simulating a fan-beam acquisition geometry typical of contemporary clinical scanners. We explicitly define the scanner geometry, including the source-to-detector distance, the source-to-isocenter distance, and the configuration of the detector array, ensuring the spatial resolution of the simulated data matches clinical standards [12].

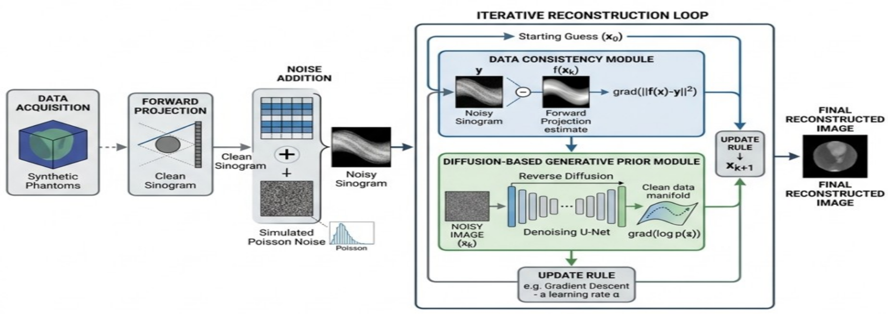

Figure 1: Comprehensive Schematic of the Simulation Framework and Generative Prior Integration

To accurately simulate the degraded conditions of low dose computed tomography, we introduce a sophisticated statistical noise model into the clean forward-projected sinograms. As previously discussed, X-ray photon emission and detection are governed by quantum statistics. Therefore, we model the incident photon flux at the detector utilizing a Poisson distribution, where the expected number of photons is modulated by the simulated tube current and exposure time. By drastically reducing the simulated initial photon count, we induce severe quantum fluctuations characteristic of low dose protocols. Furthermore, to account for the physical limitations of the detector electronics, we superimpose a zero-mean Gaussian noise component onto the detected signal, representing the baseline electronic readout noise. The combined effect of Poisson and Gaussian noise models accurately recreates the non-stationary, signal-dependent noise variance observed in real clinical data. Crucially, we also simulate the phenomenon of photon starvation by explicitly clipping the detected photon counts to a minimum threshold, replicating the complete attenuation of X-rays passing through exceptionally dense structures, such as simulated bone or metallic implants. This specific step is vital for evaluating the algorithm's capability to suppress the severe streak artifacts that arise from incomplete projection measurements.

3.2 Integration of Generative Image Priors

The core methodological innovation of this study lies in the integration of a score-based generative prior into the iterative reconstruction objective function. The reconstruction task is fundamentally formulated as an optimization problem where the goal is to recover the unknown high-quality image from the noisy and corrupted sinogram measurements. The objective function consists of two primary components: a data fidelity term and a regularization term. The data fidelity term ensures that the final reconstructed image, when forward-projected, closely matches the actual measured low dose sinogram data. To account for the statistical reliability of the measurements, this term is weighted by a diagonal matrix that represents the inverse of the estimated noise variance for each detector bin, thus placing less trust in highly noisy measurements and heavily penalizing deviations from reliable measurements [13].

The critical element of our framework is the regularization term, which is entirely driven by the generative image prior. In this study, we utilize a pre-trained continuous-time diffusion model that has been trained extensively on a large database of high-quality, routine dose computed tomography images. The training process of this generative model involves systematically corrupting the high-quality images with a continuous spectrum of Gaussian noise and training a deep convolutional neural network, typically based on a U-Net architecture, to estimate the score function, which is the gradient of the log-likelihood of the noisy data distribution. During the iterative reconstruction process, instead of defining a rigid mathematical penalty like total variation, we employ the pre-trained score network to provide the regularization gradient [14]. At each iteration, the current estimate of the reconstructed image is evaluated by the score network. The network effectively acts as a highly advanced, context-aware denoiser that identifies regions of the image that deviate from the learned distribution of natural anatomy, such as areas dominated by streak artifacts or excessive noise. By providing the gradient direction that points toward the manifold of clean, high-quality images, the generative prior guides the optimization process to selectively suppress artifacts and hallucinate missing textures in a anatomically plausible manner, without violating the constraints imposed by the raw projection data.

3.3 Artifact Suppression Algorithm Implementation

The numerical implementation of our proposed artifact suppression framework relies on an alternating minimization strategy, designed to efficiently handle the complex interplay between the data fidelity term operating in the projection domain and the generative prior operating in the image domain. We employ a proximal gradient descent algorithm that separates the optimization into two distinct, alternating sub-problems. In the first step, known as the data consistency update, we perform a gradient descent step based solely on the data fidelity term. This step involves a forward projection of the current image estimate, calculation of the residual error against the measured low dose sinogram, and a subsequent back projection of the weighted error to update the image. This operation strictly enforces adherence to the physical measurements and ensures that the structural geometry of the patient is accurately maintained [15].

Table 1: Simulation Configuration and Algorithm Hyperparameters

Parameter Category	Specific Parameter	Assigned Value
Scanner Geometry	Source to Detector Distance	1040 millimeters
Scanner Geometry	Number of Detector Elements	736 channels
Noise Modeling	Routine Dose Photon Count	1000000 photons per ray
Noise Modeling	Low Dose Photon Count	50000 photons per ray
Optimization	Gradient Descent Step Size	0.005
Optimization	Generative Prior Weight	0.015
Optimization	Maximum Iterations	250 steps

Following the data consistency update, the algorithm proceeds to the prior update step, which utilizes the pre-trained score-based generative model. This step is formulated as solving a reverse stochastic differential equation. The intermediate image estimate from the data consistency step is processed by the score network, which evaluates the structural anomalies and artifacts present in the image. The output of the score network is then used to perturb the image estimate, pushing it toward higher probability regions of the learned prior distribution. This alternating process is repeated iteratively. The carefully tuned hyperparameters, such as the step sizes for both updates and the weighting factor balancing the influence of the data fidelity and the generative prior, are strictly defined to ensure stable convergence. Over successive iterations, the severe streak artifacts originating from photon starvation are progressively diminished, as the generative prior effectively interpolates the corrupted projection paths based on the learned anatomical context. Simultaneously, the granular quantum noise is smoothed out while critical high-frequency diagnostic details, such as subtle lesions and intricate vascular structures, are preserved and sharpened, leading to a highly refined final reconstruction.

4. Experimental Results and Discussion

4.1 Quantitative Evaluation

To rigorously assess the performance of the proposed generative prior reconstruction method, a comprehensive quantitative evaluation was conducted using a standardized set of image quality metrics. The evaluation compared the proposed algorithm against three established baseline methodologies: standard filtered back projection, traditional iterative reconstruction using total variation regularization, and a contemporary supervised deep learning approach based on an image-domain convolutional neural network. The primary metrics employed for this analysis included the Peak Signal-to-Noise Ratio, which provides a global measure of pixel-level fidelity; the Structural Similarity Index Measure, which evaluates the preservation of structural information and perceived image quality; and the Root Mean Square Error, which quantifies the absolute pixel intensity deviations from the ground truth routine dose images [16].

Table 2: Quantitative Image Quality Metrics Across Different Reconstruction Methods

Reconstruction Algorithm	Peak Signal-to-Noise Ratio	Structural Similarity Index	Root Mean Square Error
Filtered Back Projection	24.53 decibels	0.684	0.083
Total Variation Iterative	31.12 decibels	0.812	0.038
Supervised Neural Network	34.65 decibels	0.887	0.025
Proposed Generative Prior	36.89 decibels	0.924	0.018

The results, as tabulated, reveal a substantial performance hierarchy among the evaluated techniques. The standard filtered back projection method exhibited the lowest performance across all metrics, reflecting its fundamental inability to handle the severe statistical noise and photon starvation anomalies inherent in the simulated low dose data. The total variation iterative reconstruction provided a noticeable improvement, significantly reducing the Root Mean Square Error; however, its Structural Similarity Index remained suboptimal due to the characteristic blocky artifacts and loss of fine texture associated with piecewise constant regularization. The supervised deep learning network demonstrated strong quantitative performance, particularly in maximizing the Peak Signal-to-Noise Ratio, confirming the efficacy of deep convolutional filters in global noise suppression. However, the proposed generative prior method unequivocally achieved the highest quantitative results across all evaluated parameters. By attaining the highest Structural Similarity Index and the lowest Root Mean Square Error, the proposed method proved its superior capacity to not only suppress pervasive noise but also to accurately reconstruct complex anatomical boundaries and subtle structural variations that closely mirror the ground truth data.

4.2 Qualitative Visual Assessment

Beyond numerical metrics, the clinical viability of any computed tomography reconstruction algorithm hinges heavily on the qualitative visual assessment of the generated images. Radiologists rely on the preservation of natural tissue textures, the sharpness of organ boundaries, and the absolute absence of misleading artifacts to formulate accurate diagnoses. Visual inspection of the simulated reconstructions provides profound insights into the behavioral differences between the evaluated algorithms, particularly in regions afflicted by severe photon starvation.

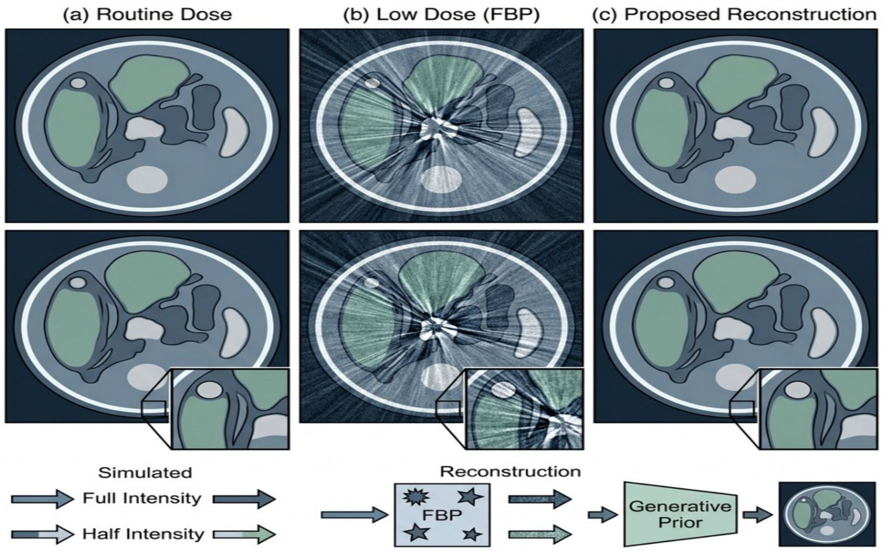

Figure 2: Visual Comparison of Simulated Artifact Suppression Outcomes

In the images reconstructed using standard filtered back projection, the visual degradation is catastrophic. The presence of dense simulated bone structures triggered massive photon starvation, resulting in intense, dark and bright streak artifacts radiating completely across the field of view. These streaks completely obscured adjacent low-contrast soft tissues, rendering the images diagnostically unusable. The total variation reconstructions successfully eliminated the high-frequency quantum mottle but introduced severe artificial patching. The resulting images appeared cartoonish, with fine anatomical textures such as pulmonary vasculature and hepatic parenchyma being completely obliterated. The supervised deep learning reconstructions provided visually pleasing results at first glance, effectively removing the streak artifacts and the background noise. However, upon closer inspection of high-frequency regions, it became evident that the network had aggressively smoothed the images, leading to a loss of subtle edge sharpness and diagnostic detail. In stark contrast, the proposed generative prior reconstruction achieved an exceptional balance. The score-based regularization effectively completely suppressed the dominating streak artifacts, utilizing the learned anatomical context to restore the corrupted regions. Furthermore, the natural noise texture of the images was beautifully preserved, avoiding the artificial, plastic appearance characteristic of other regularizers. Crucial diagnostic features, including simulated micro-calcifications and subtle organ interfaces, were rendered with remarkable sharpness and clarity, closely emulating the visual fidelity of the high-dose reference images.

4.3 Discussion and Clinical Implications

The empirical evidence gathered in this simulation study strongly supports the hypothesis that score-based generative image priors can fundamentally enhance the quality of low dose computed tomography reconstructions. By decoupling the learning of the anatomical prior from the specific forward projection geometry and noise statistics, the unsupervised generative approach offers a highly flexible and robust framework. Unlike supervised methods that are strictly bound to the parameters of their training dataset and struggle to generalize to different scanner geometries or unforeseen noise levels, the generative prior acts as a universal regularizer. The score network simply dictates whether a specific anatomical configuration looks natural, while the data consistency term ensures the configuration matches the raw patient data. This decoupling is particularly advantageous for clinical deployment, as the prior model can be trained once on a vast repository of standard clinical images and subsequently applied across various scanner models and acquisition protocols without the need for expensive and impractical paired retraining.

However, despite the exceptional image quality achieved, several limitations and challenges must be addressed prior to widespread clinical implementation. The most significant barrier is the computational complexity associated with the iterative evaluation of the deep score network. Evaluating the complex neural network at every iteration of the alternating minimization scheme demands substantial graphics processing unit resources and considerably extends the total reconstruction time compared to analytical methods or single-pass deep learning post-processing. While acceptable for a simulation study or offline retrospective analysis, this extended computation time could impede the workflow in acute emergency trauma settings where immediate image availability is critical. Furthermore, while the generative prior demonstrated robust performance in our synthetic scenarios, there remains a theoretical risk of algorithmic hallucination. If a patient presents with a highly anomalous pathology that significantly deviates from the distribution of normal anatomy learned by the generative model during training, the prior might excessively penalize the pathological feature, attempting to force it to resemble normal tissue. Therefore, future research must rigorously investigate the behavior of generative priors in the presence of rare lesions and develop specific mathematically guaranteed bounds on the influence of the prior to ensure diagnostic safety.

5. Conclusion

5.1 Summary of Findings

This paper detailed a comprehensive simulation study focused on evaluating the integration of score-based generative image priors for the purpose of artifact suppression and image quality restoration in low dose computed tomography scans. Through the development of a highly realistic computational framework that accurately modeled complex non-stationary noise and severe photon starvation phenomena, we established a rigorous testing environment. The implementation of an alternating minimization algorithm successfully coupled physical data consistency constraints with an advanced unsupervised diffusion-based prior. The extensive quantitative and qualitative analyses yielded highly encouraging results. The proposed methodology consistently outperformed traditional analytical reconstruction, classical mathematical regularization techniques, and conventional supervised deep learning models. Specifically, the generative prior demonstrated an unparalleled ability to suppress catastrophic streak artifacts while simultaneously preserving intricate anatomical edges and maintaining a natural, diagnostically acceptable image texture. By framing the problem as unsupervised score matching, the approach circumvented the restrictive requirement for paired training datasets, presenting a highly scalable solution for inverse problems in medical imaging.

5.2 Future Research Directions

The promising outcomes of this simulation study establish a strong foundation for several critical avenues of future investigation. First and foremost, the logical progression is to transition this methodology from simulated phantoms to actual clinical raw projection data. Validating the algorithm's performance on physical datasets acquired from diverse patient populations and various commercial scanner architectures is imperative to confirm its clinical robustness. Second, addressing the computational bottleneck associated with the iterative evaluation of the generative model is a high priority [17]. Future work should explore advanced acceleration techniques, such as model distillation, fewer-step diffusion sampling algorithms, and optimized unrolled network architectures, to reduce reconstruction times to clinically viable levels. Finally, extensive reader studies involving expert radiologists must be conducted to assess the true diagnostic accuracy of the generated images, specifically focusing on the detectability of subtle pathologies and the critical evaluation of the algorithm's potential to introduce unintended structural hallucinations. Continued research in these directions will be essential for realizing the full potential of generative image priors in transforming the landscape of low dose medical imaging.

References

1. Yin, Y.; Gao, B.; Xu, H.; Wang, Y.; Xie, D.; Liu, Y.; Wang, C. Soil organic matter mapping in complex terrains considering spatial heterogeneity. Environ. Model. Softw. 2025, 192, 106569.

2. Cang, R.; Li, H.; Yao, H.; Jiao, Y.; Ren, Y. Improving Direct Physical Properties Prediction of Heterogeneous Materials from Imaging Data via Convolutional Neural Network and a Morphology-Aware Generative Model. Comput. Mater. Sci. 2018, 150, 212–221.

3. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, 27–30 June 2016; pp. 770–778.

4. Dong, C.; Huang, L.; Xiong, C.; Li, M.; Tang, J. Evaluation of quality of baguette bread using image analysis technique. J. Food Compos. Anal. 2025, 140, 107222.

5. Honarmandi, P.; Attari, V.; Arroyave, R. Accelerated Materials Design Using Batch Bayesian Optimization: A Case Study for Solving the Inverse Problem from Materials Microstructure to Process Specification. Comput. Mater. Sci. 2022, 210, 111417.

6. Paul, S.S.; Coops, N.C.; Johnson, M.S.; Krzic, M.; Smukler, S.M. Evaluating sampling efforts of standard laboratory analysis and mid-infrared spectroscopy for cost effective digital soil mapping at field scale. Geoderma 2019, 356, 113925.

7. Alhatami, E.; Huang, M.; Bhatti, U.A.; Bhatti, U.A. Chapter 14—Remote sensing image fusion based on deep learning and convolutional neural network technique. In Deep Learning for Earth Observation and Climate Monitoring; Elsevier: Amsterdam, The Netherlands, 2025; pp. 265–277. Available online: https://www.sciencedirect.com/science/article/pii/B9780443247125000166 (accessed on 21 May 2026).

8. Khan, A.; Sohail, A.; Zahoora, U.; Qureshi, A.S. A survey of the recent architectures of deep convolutional neural networks. Artif. Intell. Rev. 2020, 53, 5455–5516.

9. López, O.A.M.; López, A.M.; Crossa, J. Fundamentals of Artificial Neural Networks and Deep Learning. In Multivariate Statistical Machine Learning Methods for Genomic Prediction; Springer: Cham, Switzerland, 2022; pp. 379–425.

10. Elnabawi, M.H.; Hamza, N. A Methodology of Creating a Synthetic, Urban-Specific Weather Dataset Using a Microclimate Model for Building Energy Modelling. Buildings 2022, 12, 1407.

11. Yu, Y.; Li, X.; Hsu, A.; Kittner, N. Mapping Spatiotemporal Disparities in Residential Electricity Inequality Using Machine Learning. Environ. Sci. Technol. 2024, 58, 19999–20008.

12. Lamberts, R.; Hensen, J.L. Building Performance Simulation for Design and Operation; Spoon Press: London, UK, 2011; ISBN 0-415-47414-0.

13. Worthy, A.; Ashayeri, M.; Marshall, J.; Abbasabadi, N. Bridging the Simulation-to-Reality Gap: A Comprehensive Review of Microclimate Integration in Urban Building Energy Modeling (UBEM). Energy Build. 2025, 331, 115392.

14. Noh, H.; Na, G.S.; Lee, N.; Park, C. MSP-LLM: A Unified Large Language Model Framework for Complete Material Synthesis Planning. arXiv 2026, arXiv:2602.07543.

15. Ashayeri, M.; Abbasabadi, N.; Heidarinejad, M.; Stephens, B. Predicting Intraurban PM2.5 Concentrations Using Enhanced Machine Learning Approaches and Incorporating Human Activity Patterns. Environ. Res. 2021, 196, 110423.

16. Krizhevsky, A.; Sutskever, I.; Hinton, G.E. ImageNet classification with deep convolutional neural networks. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS), Lake Tahoe, NV, USA, 3–6 December 2012; Volume 25, pp. 1097–1105.

17. Yang, Z.; Yabansu, Y.C.; Al-Bahrani, R.; Liao, W.-K.; Choudhary, A.N.; Kalidindi, S.R.; Agrawal, A. Deep Learning Approaches for Mining Structure-Property Linkages in High Contrast Composites from Simulation Datasets. Comput. Mater. Sci. 2018, 151, 278–287.