

# Mambagym: Comparing Low-Dose CT Reconstruction Denoisers with Optimizing Priors

John Lin, Yukthi Wickramarachchi, Steven Yuan

**Abstract**—Computed Tomography (CT) is a diagnostic medical imaging technique using X-rays to produce internal images of the human body. However, the utility of CT scans are balanced by the adverse effects of using X-rays, which are a form of ionizing radiation that may induce cancer, due to the damage they cause to DNA. In order to reduce patient risk, low-dose CT is a newer technique that reduces ionizing radiation exposure. However, it comes at the cost of CT image quality, as low-dose CT produces less-detailed images compared to conventional CT, since lower radiation results in lower signal collected. At the same time, increasing the signal by increasing image gain results in a similar increase in background noise. Is it possible, therefore, to improve low-dose CT image quality through computational imaging techniques, while preserving the benefits of low-dose CT on patient radiation exposure? Here, we explore a novel denoising method based on state-space modeling, DenoMamba, using an Alternating Direction Method of Multipliers (ADMM) plug-and-play prior, on resolving low-dose CT images. We show that compared to state-of-the-art and status-quo denoisers, DenoMamba with ADMM performs better at resolving low-dose CT images, with similar subjective image resolution compared to conventional CT. This work highlights the role of advanced computational imaging techniques in addressing trade-offs in technical outcomes and patient risk in diagnostic medical imaging, and sets the foundation for further innovations in this field.

**Index Terms**—Computational Imaging, Low-Dose Computed Tomography, Denoising, Convolutional Neural Networks, State-Space Machines




Fig. 1. Anatomical planes as seen in magnetic resonance imaging (MRI) [1].

## 1 INTRODUCTION

CONVENTIONAL computed tomography (CT) techniques are the norm nowadays to resolve relatively high-resolution, high-quality images of the internal structures of the human body. Indeed, with today’s advances in medical imaging technology, it’s possible from a single scan to construct images in any of the three anatomical planes: axial, coronal, and sagittal *and* construct a 3D volume of the target organs (Fig. 1).

However, conventional CT is not without drawbacks. Primarily, the ionizing radiation that CT generates increases the risk of cancer, particularly in children [2]. Thus, in order to reduce radiation exposure, low-dose CT was introduced.

Nevertheless, simply reducing radiation exposure in CT is not without its own problems. Namely, with lower radiation dosages comes a proportional deterioration in image quality [3]. This can be somewhat mitigated by increasing image gain, but also results in excessive noise and artifacts upon reconstruction [4]. Thus, the goal of reducing clinical radiation exposure while still obtaining high-quality CT images is an area of active research.

The typical flow for CT imaging involves constructing X-ray projections into sinograms, which are then converted into the image domain (Fig. 2). Traditional algorithms used for this task include filtered back projection (FBP) and Feldkamm-Davis-Kress, which are notable for their computational efficiency and reconstruction speed, especially in conventional CT scans. However, since these algorithms rely on high quality projection data, low-dose CT scans are not suitable for these algorithms. Alternatively, iterative reconstruction is another technique that benefits from advances in computing power and robustness to diverse scanning conditions and noise levels. In other words, it is an improvement on FBP and Feldkamm-Davis-Kress especially in the realm of noise reduction. However, Chen et al. state that despite its benefits, iterative reconstruction suffers from slow speed due to multiple iterative optimization and intricate setup, which restrict its use in clinical practice [4].

The current state-of-the-art for low-dose CT imaging instead involves deep learning in various modalities. Chen et al. [4] categorize these modalities into four broad groups: projection domain preprocessing, image domain post-processing, hybrid projection and image domain pre- and post-processing and direct reconstruction using DL alone (Fig. 3). Some of these methods that we will explore include Residual Encoder-Decoder Convolutional Neural Network (RED-CNN) [5], which is an image domain DL




Fig. 2. CT projection to image flow




Fig. 3. Four domains of DL-based low-dose CT reconstruction.

model which has shown promising results in the past.

In addition to these approaches, iterative optimization algorithms can be employed to solve the low-dose CT linear inverse problem. In this paper we will focus on this type of approach, exploring algorithms such as ADMM and Fixed-Point iteration in order to obtain higher quality reconstructions.

Finally, we intend to explore DenoMamba, a novel image domain denoiser that uses state-space modeling (SSM) to capture short-and long-range context in order to improve on the status quo of low-dose CT reconstruction algorithms [6]. By evaluating these cutting-edge techniques, we hope to push the boundaries of what is possible with low-dose CT, balancing image quality and radiation reduction while maintaining computational feasibility for real-world clinical use.

## 2 RELATED WORK

### 2.1 Inverse Image Problems

The primary objective of this work lies in the domain of inverse image problems, where the goal is to recover an unknown image  $\mathbf{x}$  from noisy measurements  $\mathbf{y}$ . In CT imaging, measurements are typically modeled as being corrupted by two primary types of noise: quantum noise and detector noise [7]. Quantum noise is commonly modeled as Poisson noise, while detector noise is typically modeled as Gaussian noise.

At lower radiation doses, the effect of the Poisson noise becomes more pronounced relative to detector noise. Consequently, the low-dose CT inverse problem is generally modeled as a Poisson inverse problem of the form:

$$\mathbf{y} = \mathbf{Ax} + \epsilon(\mathbf{Ax}), \quad (1)$$

where:

- $\mathbf{y}$  represents the noisy sinogram measurements,

- $\mathbf{x}$  is the true image to be reconstructed,
- $\mathbf{A}$  is the system matrix (in the CT case, a Radon transform) [8],
- $\epsilon(\mathbf{Ax})$  represents a data-dependent noise distribution modeled by:

$$\epsilon(\mathbf{Ax}) = -\mathbf{Ax} - \ln \left( \frac{\tilde{N}_1}{N_0} \right)$$

$$\text{where, } \tilde{N}_1 \sim \text{Poisson}(N_0 \exp(-\mathbf{Ax}))$$

However, numerical stability issues arose when explicitly solving for the poisson objective. As a result, we found that modeling the noise as additive Gaussian noise and subsequently optimizing a Gaussian-based objective function led to better results. Thus, in this paper, the inverse problem is modeled using the form:

$$\mathbf{y} = \mathbf{Ax} + \epsilon, \quad (2)$$

where  $\epsilon \sim N(0, \sigma^2)$  represents additive Gaussian noise. This formulation corresponds to what is commonly referred to as the **linear inverse imaging problem**.

Due to the presence of noise corrupting the measurements, direct inversion of the system matrix  $\mathbf{A}$  is not feasible. Therefore, alternative optimization techniques are required to address this ill-posed inverse problem.

### 2.2 Regularization by Denoising (RED)

Regularization by Denoising (RED) is a framework that utilizes denoising algorithms to regularize general inverse problems in image processing [9]. This method defines a regularization term as the cross-correlation between the image and its denoising residual, aiming to minimize this correlation. RED incorporates a denoising engine iteratively within the regularization term, making the overall optimization task clearer and better defined. The flexibility of RED allows for various optimization methods and is capable of achieving state-of-the-art results in image deblurring and super-resolution.

$$E(x) = \frac{1}{2\sigma^2} \|\mathbf{Hx} - \mathbf{y}\|_2^2 + \frac{\lambda}{2} \mathbf{x}^T (\mathbf{x} - f(\mathbf{x}))$$

### 2.3 Alternating Direction Method of Multipliers (ADMM)

The alternating direction method of multipliers (ADMM) is an algorithm that solves convex optimization problems by breaking them down into smaller, easier-to-handle subproblems [10], [11]. The algorithm achieves this by alternating between minimizing the objective function with respect to different subsets of variables, while also incorporating a dual variable update to enforce constraints. ADMM has a wide range of applications but is particularly useful for denoising since it allows us to mix and match various algorithms to assess their performance, especially in the Plug-and-Play Prior ( $P^3$ ) framework [12].

$$L_\rho(x, z, y) = f(x) + g(z) + y^T (Ax + Bz - c) + \frac{\rho}{2} \|Ax + Bz - c\|_2^2$$




Fig. 4. RED-CNN Architecture

$P^3$  with ADMM is a framework built on top of ADMM, with the core idea being that the explicit prior is replaced with a denoiser, which decomposes the inverse problem into simpler subproblems [9].

## 2.4 Residual Encoder-Decoder Convolutional Neural Network (RED-CNN)

The residual encoder-decoder convolutional neural network (RED-CNN) is a deep learning model designed to improve low-dose CT imaging [5]. Its architecture comprises ten layers, including five convolutional and five deconvolutional layers arranged symmetrically, with shortcuts connecting matching layers (Fig. 4). RED-CNN’s performance has been shown to be superior to other state-of-the-art methods in both simulated and clinical cases [5]. The model’s robustness allows it to handle different noise levels with a computational cost significantly lower than other methods. As such, it is a promising solution for denoising low-dose CT imaging and a good benchmark for comparing to our target framework.

## 2.5 Sequence Modeling for Images

A recently the use of sequence modeling architectures such as Vision Transformers [13] have been used for various computational imaging tasks. These use self-attention [14] operators in order to capture spatial contextual relationships within the image. These architectures process images as sequences of non-overlapping patches embedded into a higher dimensional space. There have been several successful uses of transformer architectures for LDCT denoising [15] [16], however due to the quadratic complexity of attention at training time, there are computational constraints to the use of this method.

## 2.6 Selective State-Space Models

Selective state space models (SSMs) are a class of neural architectures originally developed for modeling sequential data, such as time-series or text [17]. They employ a linear state-space system coupled with learnable gating mechanisms to selectively control the flow of information across input sequences.

For image processing, information flow is captured in a hidden state which is updated at each time step (in the case of images, each subsequent patch embedding). This hidden state captures the short and long range context

without the need for self-attention, significantly reducing the computational complexity.

The selective state-space mechanism is defined as follows:

$$\mathbf{h}_{t+1} = \mathbf{A}\mathbf{h}_t + \mathbf{B}\mathbf{x}_t, \quad \mathbf{y}_t = \mathbf{C}\mathbf{h}_t,$$

where:

$$\begin{aligned} \mathbf{h}_t &\in \mathbb{R}^d && \text{hidden state at time } t, \\ \mathbf{x}_t &\in \mathbb{R}^m && \text{input at time } t \text{ (e.g. an image patch)}, \\ \mathbf{y}_t &\in \mathbb{R}^n && \text{output at time } t, \\ \mathbf{A} &\in \mathbb{R}^{d \times d} && \text{state transition matrix,} \\ \mathbf{B} &\in \mathbb{R}^{d \times m} && \text{input mapping matrix,} \\ \mathbf{C} &\in \mathbb{R}^{n \times d} && \text{output mapping matrix.} \end{aligned}$$

Selective gating mechanism:

$$\mathbf{h}_{t+1} = g_t \odot (\mathbf{A}\mathbf{h}_t) + (1 - g_t) \odot (\mathbf{B}\mathbf{x}_t),$$

where  $g_t \in [0, 1]^d$  is a learned gating vector.

In the image context, the selective input gating allows for selection of both long range and short range spatial features to contribute to the state transition.




Fig. 5. Selective State-Space Mechanism.

## 2.7 DenoMamba

Öztürk et al. [6] proposed Denomamba, a state-space model based LDCT denoiser which leveraged a U-Net style encoder-decoder structure with each stage containing a sequence of “FuseSSM” blocks. Each FuseSSM block takes a sequence of image non-overlapping patches as an input and consists of 3 pathways which are pooled (or fused) together to obtain the an output. .




Fig. 6. FuseSSM Block




Fig. 7. DenoMamba Architecture

### 3 PROPOSED METHOD

#### 3.1 Motivation

In recent years, convolutional neural networks (CNNs) have primarily been the state-of-the-art in deep learning based LDCT denoising [18]. However, CNNs are inherently constrained by their limited receptive field, which may hinder their ability to capture long-range dependencies and global context in CT images. To address this limitation, we wanted to explore the use of an alternative architecture capable of capturing both short and long range contextual features. We hypothesized that this could better preserve complex tissue and structure details, resulting in better denoising and reconstruction.

#### 3.2 Our Method

We propose incorporating Denomamba into the RED framework as a de-noising prior for low-dose CT image reconstruction task. We experimented with using both ADMM and Fixed Point iteration algorithms to perform reconstructions. For our ADMM solution we chose to use a conjugate gradient solver to solve for the data fidelity term, and used a fixed point iterative solution to solve for the regularization term.

---

#### Algorithm 1 Denoising with ADMM

---

1: **Input:** Supply the following inputs:

- DenoMamba  $f_{DM}(\cdot)$
- Forward and adjoint system matrices  $\mathbf{H}$  and  $\mathbf{H}^T$
- Regularization parameter  $\lambda$
- Log-Likelihood parameter:  $\sigma$
- ADMM coefficient  $\rho$
- Number of outer and inner iterations:  $N, m$

2: **Initialization:** Set  $\tilde{\mathbf{x}}_0 = 0, \hat{\mathbf{v}}_0 = 0$ , and  $\hat{\mathbf{u}}_0 = 0$

3: **for**  $k = 1, 2, \dots, N$  **do**

4: **Part 1: Solve**  $\tilde{\mathbf{x}}_k = \arg \min_z \frac{1}{2} \|\mathbf{H}z - \mathbf{y}\|_2^2 + \frac{\rho}{2} \|z - \hat{\mathbf{v}}_{k-1} + \hat{\mathbf{u}}_{k-1}\|_2^2$   
**Solve via Conjugate Gradient Method**

5: **Part 2: Solve**  $\hat{\mathbf{v}}_k = \arg \min_z \lambda z^T (z - f_{\sigma_f}(z)) + \frac{\rho}{2} \|z - \tilde{\mathbf{x}}_k - \hat{\mathbf{u}}_{k-1}\|_2^2$

6: **Initialization:**  $\mathbf{z}_0 = \hat{\mathbf{v}}_{k-1}$ , and define  $\mathbf{z}^* = \tilde{\mathbf{x}}_k + \hat{\mathbf{u}}_{k-1}$

7: **for**  $j = 1, 2, \dots, m_2$  **do**

8: Apply DenoMamba  $\tilde{\mathbf{z}}_j = f_{DM}(\hat{\mathbf{z}}_{j-1})$

9: Compute the gradient  $\mathbf{z}_j = \frac{1}{\rho + \lambda} (\lambda \tilde{\mathbf{z}}_j + \rho \mathbf{z}^*)$

10: **end for**

11: Set  $\hat{\mathbf{v}}_k = \mathbf{z}_m$

12: **Part 3: Update**  $\hat{\mathbf{u}}_k = \hat{\mathbf{u}}_{k-1} + \tilde{\mathbf{x}}_k - \hat{\mathbf{v}}_k$

13: **end for**

---



---

#### Algorithm 2 Denoising with Fixed Point Iteration

---

1: **Input:** Supply the following ingredients and parameters:

- DenoMamba  $f_{DM}(\cdot)$
- Forward and adjoint system matrices  $\mathbf{H}$  and  $\mathbf{H}^T$
- Regularization parameter  $\lambda$
- Log-Likelihood parameter:  $\sigma$
- Number of outer and inner iterations:  $N$  and  $m$

2: **Initialization:**  $\tilde{\mathbf{x}}_0 = 0$

3: **for**  $k = 1, 2, \dots, N$  **do**

4: Run DenoMamba forward pass:  $\tilde{\mathbf{x}}_k = f_{DM}(\tilde{\mathbf{x}}_{k-1})$

5: Solve  $\mathbf{A}\mathbf{z} = \mathbf{b}$  for  $\mathbf{A} = \frac{1}{\sigma^2} \mathbf{H}^T \mathbf{H} + \lambda \mathbf{I}$  and  $\mathbf{b} = \frac{1}{\sigma^2} \mathbf{H}^T \mathbf{y} + \lambda \tilde{\mathbf{x}}_k$

6: **Initialization:**  $\mathbf{z}_0 = \tilde{\mathbf{x}}_k$

7: **for**  $j = 1, 2, \dots, m$  **do**

8: Compute the residual:  $\mathbf{r}_j = \mathbf{A}\mathbf{z}_{j-1} - \mathbf{b}$

9: Compute the vector:  $\mathbf{e}_j = \mathbf{A}\mathbf{r}_j$

10: Compute the optimal step size:  $\mu = \frac{\mathbf{r}_j^T \mathbf{e}_j}{\mathbf{e}_j^T \mathbf{e}_j}$

11: Update the solution:  $\mathbf{z}_j = \mathbf{z}_{j-1} + \mu \cdot \mathbf{r}_j$

12: **end for**

13: Set  $\tilde{\mathbf{x}}_k = \mathbf{z}_m$

14: **end for**

---

### 4 EXPERIMENTS

To evaluate Denomamba as a denoising prior within the RED framework, we compared it to another deep denoiser, RED-CNN. In addition to this, we evaluated the performance of DenoMamba in standard Plug-and-Play ADMM, as well as some classical CT reconstruction techniques for a baseline comparison.

#### 4.1 Evaluation Metrics

We evaluated performance on commonly used measures such as Peak Signal-to-Noise Ratio (PSNR) and structural similarity index measure (SSIM).

PSNR is the most commonly used metric for evaluating the quality of reconstructions. It is computed as:

$$\text{PSNR} = 10 \log_{10} \left( \frac{\text{MAX}_I^2}{\text{MSE}} \right),$$

where:

$$\text{MSE} = \frac{1}{N} \sum_{i=1}^N (\mathbf{I}_{\text{ref}}(i) - \mathbf{I}_{\text{test}}(i))^2,$$

where  $\text{MAX}_I$  denotes the maximum possible pixel value of the image. PSNR is expressed in decibels (dB) and a higher PSNR indicates better reconstruction quality.

SSIM is a perceptual metric that evaluates the similarity between two images by considering their structural information, luminance, and contrast [19].

#### 4.2 Datasets

We evaluated the performance of each method using the LoDoPaB-CT reconstruction dataset [7]. This dataset consists of parallel-beam projection CT images generated from thoracic slices of 800 patients at full dose. To simulate low-dose CT conditions, the projections were corrupted with Poisson noise.

The dataset includes projections captured at 1,000 equidistant angles with 513 equidistant detector bins, resulting in 1000x513 noisy projection data. The corresponding "ground truth" images have a resolution of 362x362.

For training of DenoMamba and RED-CNN we used the 2016 Low-Dose CT Grand Challenge dataset [20]. This

dataset consisted of 512x512 paired image examples taken from 30 patient scans. Each image pair contained a full dose "ground truth" image and a quarter dose image simulated by adding poisson noise.

### 4.3 Training details

We trained DenoMamba for 60 epochs with a learning rate of  $5e^{-5}$  on an Nvidia RTX-A4500 GPU on the UofT GPU Clusters. We trained RED-CNN for 100 epochs on an Nvidia RTX-4090 with a learning rate of  $1e^{-5}$ .

### 4.4 Experiment Details

For the ADMM Plug-and-Play and RED+ADMM experiments with DenoMamba, we ran for 50 ADMM iterations and 5 fixed-point iterations at each z-update step with parameters  $\lambda = 0.1$  and  $\rho = 0.1$ .

For the RED+ADMM experiment using RED-CNN we ran for the same number of iterations but used parameters  $\lambda = 0.1$  and  $\rho = 1.0$ .

For the RED-Fixed-Point experiment with Denomamba we ran for 50 outer iterations and 25 inner iterations per sample with parameters  $\sigma^2 = 0.3$  and  $\lambda = 0.5$ .

For the RED-Fixed-Point experiment with RED-CNN we ran for the same number of iterations with parameters  $\sigma^2 = 1.0$  and  $\lambda = 0.1$ .

For the ADMM with a Total-Variation prior experiment we ran for 75 ADMM iterations with parameters  $\lambda = 0.1$  and  $\rho = 1.0$ .

For our experiment with the SART reconstruction algorithm we ran for 100 iterations per sample with  $\omega = 1$ .

## 5 RESULTS AND DISCUSSION

We evaluated eight reconstruction algorithms on 1,000 test examples using three metrics: PSNR, SSIM, and RMSE. Our method, Fixed-Point Iteration with DenoMamba Prior, achieved the best results across all metrics (Table. 1). With a PSNR of 35.62 dB, our method indicates superior reconstruction quality and noise suppression, as PSNR is inversely related to reconstruction error. The SSIM of 0.77 demonstrates that our method can retain image details better than other methods that we have tested, displaying the highest structural fidelity to the reference images. The low RMSE further substantiates our method's precision - the lowest error was produced in pixel intensity reconstruction.

### 5.1 Qualitative Analysis of Reconstruction Performance

In Fig. 8, ADMM + DenoMamba and Fixed-point + DenoMamba produce clearer and more detailed structures in the lung region compared to the other methods. Both of these methods use a learned denoising model that is trained to handle the kind of noise and artifacts that appear in low-dose CT scans. We observe that these algorithms are

TABLE 1  
Evaluation of Reconstruction Algorithms on 1000 Test Examples Using PSNR, SSIM, and RMSE Metrics

Reconstruction Algorithm	PSNR (dB)	SSIM	RMSE
FBP	27.70	0.40	0.22
ADMM TV	33.51	0.64	0.11
SART	33.59	0.71	0.11
ADMM + RED-CNN	31.20	0.69	0.15
Fixed-Point + RED-CNN	29.28	0.60	0.18
$P^3$ + DenoMamba	35.04	0.74	0.10
ADMM + DenoMamba	35.20	0.74	0.09
Fixed-Point + DenoMamba	<b>35.62</b>	<b>0.77</b>	<b>0.08</b>

better able to remove noise while keeping subtle anatomical details intact. The DenoMamba model likely learns patterns specific to lung structures, such as fine vessel details, thin boundaries, and soft texture variations that traditional models or generic priors fail to preserve. This leads to a higher PSNR and SSIM, as seen in the provided metrics, because the final reconstructed image more closely matches the underlying true anatomy without introducing smooth or blocky artifacts.

However, In Fig. 9, we observe that RED-CNN has better performance. The smoother boundaries and higher SSIM value indicate that RED-CNN preserved structural details, such as the curved and high-contrast regions, better than other algorithms. This aligns with the strength of RED-CNN, which is great at capturing hierarchical features and retaining spatial coherence in noise suppression [3]. For homogeneous areas in the image, such as regions with lower intensity variation, ADMM + RED-CNN avoided excessive smoothing while preserving global consistency, whereas the fixed-point strategy introduced more artifacts among these regions.

### 5.2 Analysis of RED-CNN

The RED-CNN method showed surprisingly lower performance, with a PSNR of 31.47 dB and SSIM of 0.71 in the ADMM configuration. This discrepancy can be attributed to multiple factors. The RED-CNN implementation we used was an off-the-shelf implementation last updated in 2020, and not optimized for the current environment [21]. According to the original RED-CNN paper by Chen et. al. [5], the model was initially implemented in MATLAB, which involves dependencies and methodologies that may no longer be compatible or optimal in modern frameworks such as PyTorch. The reliance on older computational environments potentially led to discrepancies in performance when porting the code to newer tools. A complete rewrite would be required to faithfully replicate the original RED-CNN performance, which was outside the scope of this project due to time constraints.

## 6 CONCLUSION

In this work, we investigated the potential of DenoMamba as a novel approach to low-dose CT denoising and evaluated its performance compared to the current SOTA meth-




Fig. 8. Visual comparison of denoised low-dose CT scans by 8 reconstruction algorithms (sample 500) compared to ground truth.

ods. The results indicate that DenoMamba achieves competitive performance, outperforming several established denoisers in key metrics such as PSNR and SSIM. However, assessing its true impact also requires more comprehensive comparisons to other DL-based denoisers, some of which are not CNNs but rather are transformers or generative adversarial networks as well [15], [16], [22].

From a computational performance perspective, the computational feasibility of DenoMamba plays a pivotal role in its potential for real-world application. While its training demands are quite high, its inference speed suggests it could integrate into clinical workflows with appropriate hardware support. Still, further optimization may be necessary to reduce resource costs without compromising

performance. Although we did not set out to evaluate training or inference speed, our implementation DenoMamba took roughly 18 hours to train and over 24 hours to evaluate on 1000 samples. By comparison, RED-CNN took roughly 40 minutes to train and approximately 2 hours to evaluate the same samples.

Importantly, the impact on patient outcomes by reducing radiation exposure cannot be overstated. DenoMamba and other low-dose CT reconstruction algorithms are necessary for continuing to improve safety in medical imaging, particularly in vulnerable populations.

Future work exploring additional denoisers and hybrid models may also yield valuable insights. By integrating advanced optimization strategies such as ADMM and lever-




Fig. 9. Visual comparison of denoised low-dose CT scans by 8 reconstruction algorithms (sample 0) compared to ground truth.

aging emerging architectures like state-space models and transformers, we may continue to advance the medical imaging field by enhancing both the quality and efficiency of low-dose CT reconstruction such that it becomes the new standard-of-care for CT imaging, ultimately enabling safer, more accurate diagnostic imaging for patients worldwide.

## REFERENCES

- [1] M. I. Iacono, E. Neufeld, E. Akinnagbe, K. Bower, J. Wolf, I. Vogiatzis Oikonomidis, D. Sharma, B. Lloyd, B. J. Wilm, M. Wyss, K. P. Pruessmann, A. Jakab, N. Makris, E. D. Cohen, N. Kuster, W. Kainz, and L. M. Angelone, "Mida: A multimodal imaging-based detailed anatomical model of the human head and neck," *PLOS ONE*, vol. 10, no. 4, pp. 1–35, 04 2015. [Online]. Available: <https://doi.org/10.1371/journal.pone.0124126>
- [2] D. J. Brenner and E. J. Hall, "Computed tomography — an increasing source of radiation exposure," *New England Journal of Medicine*, vol. 357, no. 22, p. 2277–2284, Nov. 2007. [Online]. Available: <http://dx.doi.org/10.1056/NEJMra072149>
- [3] H. Chen, Y. Zhang, W. Zhang, P. Liao, K. Li, J. Zhou, and G. Wang, "Low-dose ct denoising with convolutional neural network," in *2017 IEEE 14th International Symposium on Biomedical Imaging (ISBI 2017)*, 2017, pp. 143–146.
- [4] H. Chen, Q. Li, L. Zhou, and F. Li, "Deep learning-based algorithms for low-dose ct imaging: A review," *European Journal of Radiology*, vol. 172, p. 111355, Mar. 2024. [Online]. Available: <http://dx.doi.org/10.1016/j.ejrad.2024.111355>
- [5] H. Chen, Y. Zhang, M. K. Kalra, F. Lin, Y. Chen, P. Liao, J. Zhou, and G. Wang, "Low-dose ct with a residual encoder-decoder convolutional neural network," *IEEE Transactions on Medical Imaging*, vol. 36, no. 12, p. 2524–2535, Dec. 2017. [Online]. Available: <http://dx.doi.org/10.1109/TMI.2017.2715284>
- [6] Şaban Öztürk, O. C. Duran, and T. Çukur, "Denomamba: A fused

- state-space model for low-dose ct denoising," 2024. [Online]. Available: <https://arxiv.org/abs/2409.13094>
- [7] J. Leuschner, M. Schmidt, D. Otero Baguer, and P. Maass, "Lodopab-ct, a benchmark dataset for low-dose computed tomography reconstruction," *Scientific Data*, vol. 8, p. 109, 04 2021.
- [8] C. L. Epstein, 6: *The Radon Transform*. SIAM, 2008.
- [9] Y. Romano, M. Elad, and P. Milanfar, "The little engine that could: Regularization by denoising (red)," *SIAM Journal on Imaging Sciences*, vol. 10, no. 4, pp. 1804–1844, 2017. [Online]. Available: <https://doi.org/10.1137/16M1102884>
- [10] S. Boyd, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, p. 1–122, 2010. [Online]. Available: <http://dx.doi.org/10.1561/22000000016>
- [11] J. Wang, F. Yu, X. Chen, and L. Zhao, "Admm for efficient deep learning with global convergence," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '19. ACM, Jul. 2019. [Online]. Available: <http://dx.doi.org/10.1145/3292500.3330936>
- [12] S. V. Venkatakrisnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *2013 IEEE Global Conference on Signal and Information Processing*. IEEE, Dec. 2013. [Online]. Available: <http://dx.doi.org/10.1109/GlobalSIP.2013.6737048>
- [13] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," 2021. [Online]. Available: <https://arxiv.org/abs/2010.11929>
- [14] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," 2023. [Online]. Available: <https://arxiv.org/abs/1706.03762>
- [15] D. Wang, F. Fan, Z. Wu, R. Liu, F. Wang, and H. Yu, "Ctformer: convolution-free token2token dilated vision transformer for low-dose ct denoising," *Physics in Medicine and Biology*, vol. 68, no. 6, p. 065012, Mar. 2023. [Online]. Available: <http://dx.doi.org/10.1088/1361-6560/acc000>
- [16] H. Li, X. Yang, S. Yang, D. Wang, and G. Jeon, "Transformer with double enhancement for low-dose ct denoising," *IEEE Journal of Biomedical and Health Informatics*, vol. 27, no. 10, pp. 4660–4671, 2023.
- [17] A. Gu and T. Dao, "Mamba: Linear-time sequence modeling with selective state spaces," 2024. [Online]. Available: <https://arxiv.org/abs/2312.00752>
- [18] J. Zhang, W. Gong, L. Ye, F. Wang, Z. Shangguan, and Y. Cheng, "A review of deep learning methods for denoising of medical low-dose ct images," *Computers in Biology and Medicine*, vol. 171, p. 108112, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0010482524001963>
- [19] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *Image Processing, IEEE Transactions on*, vol. 13, pp. 600 – 612, 05 2004.
- [20] C. H. McCollough, A. C. Bartley, R. E. Carter, B. Chen, T. A. Drees, P. Edwards, D. R. Holmes, 3rd, A. E. Huang, F. Khan, S. Leng, K. L. McMillan, G. J. Michalak, K. M. Nunez, L. Yu, and J. G. Fletcher, "Low-dose CT for the detection and classification of metastatic liver lesions: Results of the 2016 low dose CT grand challenge," *Med. Phys.*, vol. 44, no. 10, pp. e339–e352, Oct. 2017.
- [21] SSinyu, "Red-cnn," 2020. [Online]. Available: <https://github.com/SSinyu/RED-CNN>
- [22] C. Tang, W. Zhang, L. Wang, A. Cai, N. Liang, L. Li, and B. Yan, "Generative adversarial network-based sinogram super-resolution for computed tomography imaging," *Physics in Medicine Biology*, vol. 65, 11 2020.