

When are Diffusion Priors Helpful in Sparse Reconstruction? A Study with Sparse-view CT

Matt Y. Cheung^{1,3,†}, Sophia Zorek^{2,3,†}, Tucker J. Netherton³, Laurence E. Court³, Sadeer Al-Kindi⁴, Ashok Veeraraghavan¹, and Guha Balakrishnan¹

¹Department of Electrical & Computer Engineering, Rice University, Houston TX

²Department of Statistics, Rice University, Houston TX

³Department of Radiation Physics, The University of Texas MD Anderson Cancer Center, Houston TX

⁴Department of Cardiology, DeBaakey Heart and Vascular Center, Houston TX

February 12, 2025

Abstract

Diffusion models demonstrate state-of-the-art performance on image generation, and are gaining traction for sparse medical image reconstruction tasks. However, compared to classical reconstruction algorithms relying on simple analytical priors, diffusion models have the dangerous property of producing realistic looking results *even when incorrect*, particularly with few observations. We investigate the utility of diffusion models as priors for image reconstruction by varying the number of observations and comparing their performance to classical priors (sparse and Tikhonov regularization) using pixel-based, structural, and downstream metrics. We make comparisons on low-dose chest wall computed tomography (CT) for fat mass quantification. First, we find that classical priors are superior to diffusion priors when the number of projections is “sufficient”. Second, we find that diffusion priors can capture a large amount of detail with very few observations, significantly outperforming classical priors. However, they fall short of capturing all details, even with many observations. Finally, we find that the performance of diffusion priors plateau after extremely few (≈ 10 -15) projections. Ultimately, our work highlights potential issues with diffusion-based sparse reconstruction and underscores the importance of further investigation, particularly in high-stakes clinical settings.

1 Introduction

Medical image modalities, such as magnetic resonance imaging (MRI) and computed tomography (CT) scans, are typically reconstructed from a set of projections. For example, CT algorithms combine many x-ray projections to produce an image volume. These volumes may then be used on various downstream tasks, such as disease diagnosis and fat measurement. Image reconstruction accuracy and reliability can, therefore, directly impact the patient’s standard of care.

Generally, the more projections used in reconstruction, the better the image quality. However, in settings requiring limited resources [1] or radiation dose [2], the number of projections can be small (< 100 instead of 100s). In this scenario, image reconstruction problems are ill-posed - meaning the number of projections is insufficient to fully capture the underlying anatomy of the patient. Many past works have addressed this “ill-posedness” through different priors, which drive the algorithm towards a solution with certain desirable

†: Equal contribution. M.C. would like to acknowledge support from a fellowship from the Gulf Coast Consortia on the NLM Training Program in Biomedical Informatics and Data Science T15LM007093. S.Z. would like to acknowledge that this material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research under Award Number DE-SC0025528. T.N. would like to acknowledge the support of the NIH LRP award.

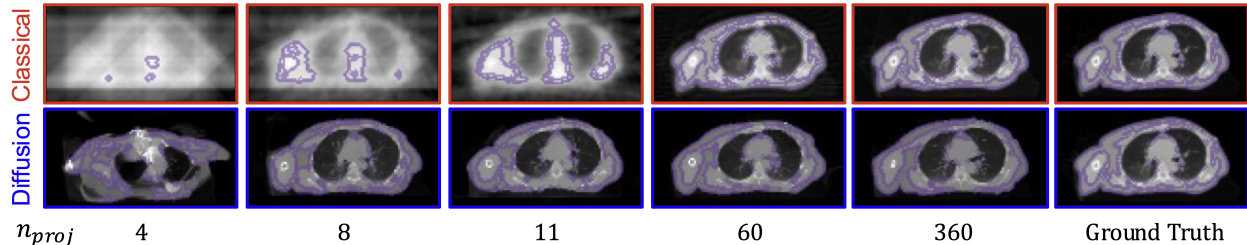


Figure 1: **Diffusion priors produce convincing results with extremely few projections.** We show reconstructions using classical (top) and diffusion (bottom) priors for varying number of projections n_{proj} . We show the fat segmentation in purple. Under limited n_{proj} , we observe the diffusion method perceptually outperforming the classical priors, but under sufficiently large n_{proj} , classical priors produce reconstructions whose fat segmentation’s adhere more closely to the ground truth.

characteristics [3]. For example, an L_1 prior drives the algorithm towards a solution that is sparse in some domain [4]. Another example is L_2 , or Tikhonov, prior that drives the algorithm towards a solution that is smooth [4].

Unlike classical priors based on simple analytical forms, an emerging alternative is to use deep generative models trained on large data distributions with neural networks. The most powerful current deep generative family are *diffusion models* [5], which demonstrate impressive results on a range of image generation tasks. Several recent studies show that diffusion models can produce convincing results in sparse image reconstruction with very limited projections [3] where classical priors struggle. However, diffusion models are not yet safe for deployment in practice because they can generate reconstructions that superficially look correct, but contain fabricated anatomical detail. This discrepancy highlights the need for a more nuanced evaluation framework.

While PSNR and SSIM provide important quantitative evaluation for a machine learning practitioner, they do not link image reconstruction with downstream clinical applications and are, therefore, not meaningful and practical [6]. Since prior work does not fully cover the entire range of projections or assess downstream metrics comprehensively, key questions emerge: 1) Do diffusion priors outperform classical priors across a wide range of projections (e.g., hundreds of X-rays for CT scans)? and 2) Do downstream metrics align with trends observed in pixel- and structure-based metrics? We tackle these questions by comparing classical (L_1 and L_2) and diffusion priors over a wide range of projections and find interesting results.

We explore the real-life, practical, and clinically relevant problem of quantifying thoracic fat mass and its distribution using low-dose CT. Quantifying thoracic fat mass and its distribution is crucial to understanding its role in cardiometabolic risk and other health outcomes [7, 8]. Traditional methods, such as full-projection computed tomography (CT) scans, offer high accuracy but are associated with significant levels of ionizing radiation. This concern is particularly important for individuals with obesity, who often require higher radiation doses due to increased tissue density [9]. Reducing radiation exposure while maintaining diagnostic accuracy is, therefore, critical for population-level screening and individual-level longitudinal assessments. Using a low number of projections (low-dose CT) offers a promising solution to balance radiation safety and diagnostic effectiveness.

Our key findings are as follows: 1) diffusion priors excel in extremely sparse settings often succeeding only with several projections, 2) classical priors are superior to diffusion priors when the number of projections is “sufficient”, and 3) diffusion prior performance plateaus after few ($\approx 10-15$) projections. Ultimately, our work paves the way for more meaningful, practical, and clinically relevant comparisons between medical image reconstruction methods in sparse regimes.

2 Experimental Setup

Dataset. We use a de-identified CT dataset of 935 patients from The University of Texas MD Anderson Cancer Center. All CT images were of patients who had received surgical mastectomy to the right side of the body, and radiotherapy to the post-mastectomy chest wall and/or axillary lymph nodes. We resample

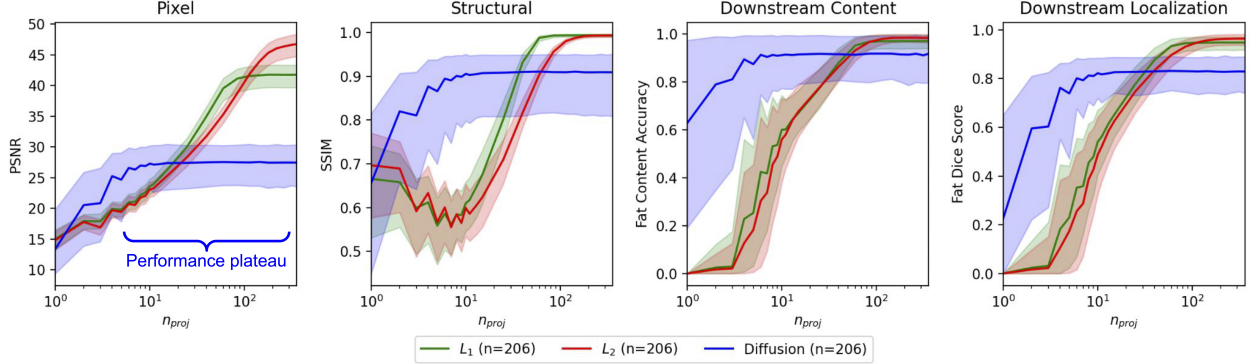


Figure 2: **Diffusion priors prevail with extremely few projections but may have wrong content and structure.** We plot evaluation metric versus number of projections for classical and diffusion priors. For 206 patients (3 slices/patient), we plot the metric and interquartile range (0.05 and 0.95 quantile, $IQR_{5,95}$) for PSNR, SSIM, fat content accuracy, and fat dice score.

each scan to 1 mm^3 , remove table artifacts and resize each slice to 128^2 pixels. We split our dataset into 729 patients consisting of over 300k slices for training and 206 patients consisting of over 618 slices testing. We evaluate our models on test data only.

Unconditional Diffusion Model Training. We train a classic 2D U-Net diffusion model to reconstruct CT slices unconditionally. Our U-Net consists of four downsampling and upsampling modules with attention heads in the third and fourth layers. We use 1000 training timesteps and a squared cosine noise scheduler as outlined in [10]. Model choice was determined based on loss convergence.

Inference. We performed a grid search to determine the best optimization, regularization, and guidance parameters for L_1 ($\lambda = 10$), L_2 ($\lambda = 10$) and diffusion ($\lambda = 0.1$). These corresponded to the lowest mean squared error at 360 projections. We used the `tomosipo` package [11] to compute the projection operator. For L_1 , we ran the total-variation regularized least squares for reconstruction for 1000 iterations. At each time step of our guided diffusion inference we take the projection of the predicted reconstruction using our projection operator and compute its mean squared error with the ground truth projections, varying the number of projections used. We then compute the gradients of this loss with respect to the predicted reconstruction, and guide the diffusion process by adding the gradients back to the prediction [5, 12]. We iterate through this process over 50 timesteps.

Metrics We investigate pixel (peak-signal-to-noise ratio or PSNR), structural (structural similarity index or SSIM), and downstream metrics (content and localization). We use the percentage of correct pixels classified as fat (fat content accuracy) and the fat segmentation dice score (fat localization accuracy). Fat has Hounsfield units between -150 and -50.

3 Results

Diffusion priors work well with extremely few projections. We find that diffusion priors excel in extremely sparse settings to estimate pixel-based, structural-based, and downstream metrics, often already succeeding in capturing many details only with several projections (Fig. 1). However, they fall short of capturing all details accurately, like vasculature, even with many observations. On the other hand, we find that classical priors are superior to diffusion priors when the number of projections is “sufficient” (i.e., exceeds extremely few projections). We leverage results from Fig. 2 to identify the regions where diffusion priors are *guaranteed* to be better than classical priors. We can define the number of projections (n_{proj}) as the interval $I = [LB, UB]$ where diffusion priors perform statistically better than classical priors. The lower and upper bounds are denoted as $LB, UB \in \mathbb{R} \times \mathbb{R}$. For a monotonically increasing (can be positive or zero slope) metric (M) as number of projections (n_{proj}) increases, the lower bound is given by n at the minimum acceptable metric (τ_M), and the upper bound is given by the maximum n_{proj} when the Mann-Whitney U-test has a p-value of less than a user-defined threshold (τ_p ; for example, 0.05): $I = [\min\{n_{proj} \in$

Metric	LB	UB (p-value)	
		L_1 Prior	L_2 Prior
MSE	7	15 (3.7e-40)	15 (2.4e-42)
SSIM	7	25 (4.24e-120)	60 (7.4e-40)
Fat Content	9	40 (9.1e-43)	40 (1.25e-72)
Fat Localization	7	25 (1.2e-69)	25 (5.39e-120)

Table 1: **Diffusion priors perform statistically better than classical priors for extremely low number of projections.** Upper (UB) and Lower (LB) bounds for number of projections where diffusion performs better than classical priors. The lower bound is from the “corner” of the performance curve and upper bound is based on the Mann-Whitney U-test.

$\mathbb{R} : M \geq \tau_M\}, \max\{n_{proj} \in \mathbb{R} : p \leq \tau_p\}$. The upper bound of the interval gives the highest number of projections where the probability that a randomly chosen metric value generated from one method is larger than a randomly chosen metric value generated from another $P(M_1 > M_2)$. For a monotonically decreasing M , maximum and minimum are used for the lower and upper bounds. We choose the lower bound to be the “corner” of the plots where performance starts to plateau. We show these bounds in Table. 3. We find that diffusion priors perform better than classical priors between extremely few observations (5 to 10) and 10s of projections. Beyond 10s of projections, classical priors perform better overall.

Diffusion prior performance plateaus after extremely few projections. We observe that as the number of projections gets larger, the performance curves plateau and do not reach perfect reconstruction (Fig. 2). In contrast, the performance of classical priors continually increases and supersedes diffusion priors beyond extremely few projections. Furthermore, we show the inter-quantile range for quantiles 0.05 to 0.95 ($IQR_{5,95}$) in Fig. 2. While the diffusion prior curves are heteroscedastic for extremely few projections, they get more homoscedastic when the performance plateaus. These observations show that the diffusion priors do not leverage the extra information well for reconstruction beyond extremely few projections. Thus, diffusion priors get high-level structures (i.e., the shape of the body) correct but not low-level structures (i.e., vasculature) (Fig. 1).

4 Conclusion and Discussion

In this paper, we answer the question “when are diffusion priors helpful for sparse reconstruction?” by comparing classical priors used in practice with state-of-the-art diffusion priors for different number of projections and metrics. First, we find that classical priors are superior to diffusion priors when the number of projections is “sufficient”. Second, we find that diffusion priors can capture a large amount of detail with very few observations, significantly outperforming classical priors. However, they fall short of capturing all details, even with many observations. Finally, we find that the performance of downstream metrics plateau after few ($\approx 10-15$) projections for diffusion priors.

One possible explanation for the performance saturation of diffusion priors is guided diffusion with naive gradient descent may not converge to the optimal solution. While the unconditional model steers the gradient toward the data space, the gradient of the external loss (i.e., the mean square error between predicted projections and ground truth projections) may point away from the data space [13]. The non-optimal convergence may explain why, even with more projections, $IQR_{5,95}$ remains constant because the guidance process eventually becomes “stuck” balancing the steering of the diffusion prior and minimizing the external loss. Therefore, when the number of projections is high, they still fail to capture the low-level features like vasculature. A possible solution could be adaptive fine-tuning [13].

Clinical Implications. While diffusion priors may not capture low-level details well, we demonstrate the potential of diffusion priors for accurate assessment of thoracic fat quantity. Using extremely few projections significantly reduces the radiation exposure required for imaging, lowering the dose by up to 97% compared with full projection and by up to 75% compared to classical priors. The reduction in radiation is particularly advantageous for individuals who require repeated imaging or those with obesity, where standard imaging protocols typically involve higher radiation doses. The enhanced performance of low-dose CT scalability makes it a viable solution for large-scale studies and clinical applications where minimizing radiation is

critical. Fixed planar X-ray systems could be employed in lieu of expensive CT machines, making diffusion priors especially valuable in low-resource settings where access to advanced imaging technologies is limited.

5 Compliance with Ethical Standards

This research was conducted retrospectively using human data under an approved IRB protocol.

References

- [1] Laurence Court, Ajay Aggarwal, Hester Burger, Carlos Cardenas, Christine Chung, Raphael Douglas, Monique du Toit, David Jaffray, Anuja Jhingran, Michael Mejia, et al., “Addressing the global expertise gap in radiation oncology: the radiation planning assistant,” *JCO Global Oncology*, vol. 9, pp. e2200431, 2023.
- [2] Cristiano Rampinelli, Daniela Origgi, and Massimo Bellomi, “Low-dose ct: technique, reading methods and image interpretation,” *Cancer imaging*, vol. 12, no. 3, pp. 548, 2012.
- [3] Yiran Sun, Hana Baroudi, Tucker Netherton, Laurence Court, Osama Mawlawi, Ashok Veeraraghavan, and Guha Balakrishnan, “Difr3ct: Latent diffusion for probabilistic 3d ct reconstruction from few planar x-rays,” *arXiv preprint arXiv:2408.15118*, 2024.
- [4] Hao Zhang, Jing Wang, Dong Zeng, Xi Tao, and Jianhua Ma, “Regularization strategies in statistical image reconstruction of low-dose x-ray ct: A review,” *Medical physics*, vol. 45, no. 10, pp. e886–e907, 2018.
- [5] Prafulla Dhariwal and Alexander Nichol, “Diffusion models beat gans on image synthesis,” *Advances in neural information processing systems*, vol. 34, pp. 8780–8794, 2021.
- [6] Matt Y Cheung, Tucker J Netherton, Laurence E Court, Ashok Veeraraghavan, and Guha Balakrishnan, “Metric-guided image reconstruction bounds via conformal prediction,” *arXiv preprint arXiv:2404.15274*, 2024.
- [7] H Wang, YE Chen, and Daniel T Eitzman, “Imaging body fat: techniques and cardiometabolic implications,” *Arteriosclerosis, thrombosis, and vascular biology*, vol. 34, no. 10, pp. 2217–2223, 2014.
- [8] Damini Dey, Ryo Nakazato, Debiao Li, and Daniel S Berman, “Epicardial and thoracic fat-noninvasive measurement and clinical implications,” *Cardiovascular diagnosis and therapy*, vol. 2, no. 2, pp. 85, 2012.
- [9] Jacquelyn C Yanch, Richard H Behrman, Michael J Hendricks, and John H McCall, “Increased radiation dose to overweight and obese patients from radiographic examinations,” *Radiology*, vol. 252, no. 1, pp. 128–139, 2009.
- [10] Jonathan Ho, Ajay Jain, and Pieter Abbeel, “Denoising diffusion probabilistic models,” *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [11] Allard A Hendriksen, Dirk Schut, Willem Jan Palenstijn, Nicola Viganó, Jisoo Kim, Daniël M Pelt, Tristan Van Leeuwen, and K Joost Batenburg, “Tomosipo: fast, flexible, and convenient 3d tomography for complex scanning geometries in python,” *Optics Express*, vol. 29, no. 24, pp. 40494–40513, 2021.
- [12] Alexandros Graikos, Nikolay Malkin, Nebojsa Jojic, and Dimitris Samaras, “Diffusion models as plug-and-play priors,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 14715–14728, 2022.
- [13] Yingqing Guo, Hui Yuan, Yukang Yang, Minshuo Chen, and Mengdi Wang, “Gradient guidance for diffusion models: An optimization perspective,” *arXiv preprint arXiv:2404.14743*, 2024.