

Zero-shot CT Super-Resolution using Diffusion-based 2D Projection Priors and Signed 3D Gaussians

Jeonghyun Noh*, Hyun-Jic Oh*, Won-Ki Jeong†

Department of Computer Science and Engineering, Korea University, Seoul, Korea
{wjdgus0967, hyunjic0127, wkjeong}@korea.ac.kr

Abstract. Computed tomography (CT) is important in clinical diagnosis, but acquiring high-resolution (HR) CT is constrained by radiation exposure risks. While deep learning-based super-resolution (SR) methods have shown promise for reconstructing HR CT from low-resolution (LR) inputs, supervised approaches require paired datasets that are often unavailable. Zero-shot methods address this limitation by operating on single LR inputs; however, they frequently fail to recover fine structural details due to limited LR information within individual volumes. To overcome these limitations, we propose a novel zero-shot 3D CT SR framework that integrates diffusion-based upsampled 2D projection priors into the 3D reconstruction process. Specifically, our framework consists of two stages: (1) LR CT projection SR, training a diffusion model on abundant X-ray data to upsample LR projections, thereby enhancing the scarce information inherent in the LR inputs. (2) 3D CT volume reconstruction, using 3D Gaussian splatting with our novel Negative Alpha Blending (NAB-GS), which models positive and negative Gaussian densities to learn signed residuals between diffusion-generated HR and upsampled LR projections. Our framework demonstrates superior quantitative and qualitative performance on two public datasets, and expert evaluations present the framework’s clinical potential at $4\times$. Github-code.

Keywords: Zero-shot CT Super-resolution · Diffusion Models · 3D Gaussian Splatting.

1 Introduction

High-quality computed tomography (CT) is an essential modality for accurate clinical diagnosis and treatment planning. However, acquiring high-resolution (HR) CT scans necessitates elevated radiation doses, risking potential DNA damage and radiation-induced malignancies [18]. While reducing radiation dose mitigates these concerns, it inevitably limits the data available for image reconstruction, thereby degrading spatial resolution [19]. This trade-off often obscures anatomical structures, potentially compromising diagnostic precision.

* Co-first authors.

† Corresponding author.

To address this, super-resolution (SR) techniques have emerged as a vital tool for recovering detailed structures from low-resolution (LR) inputs [11]. Deep learning has revolutionized a wide range of computer vision tasks [12,21,22]. In particular, SR has benefited from advances in convolutional neural networks [7,35], transformers [3,16], and diffusion models [14,17,24], these are primarily tailored to 2D images and often fail to exploit volumetric consistency in medical imaging. In the medical domain, supervised 3D SR approaches, ranging from spatially aware interpolation networks [23] to implicit neural representations (INR) [8,32], have been proposed. Despite such effectiveness, applicability is limited by scarce paired HR-LR training volumes. Consequently, zero-shot learning has gained traction as a data-efficient alternative [25,26,29], with recent innovations like CuNeRF [4] adapting neural radiance fields for continuous upsampling. Yet, relying solely on LR information, current zero-shot methods often struggle to recover structural details, yielding over-smoothed reconstructions.

To overcome these limitations, we propose a novel zero-shot 3D CT SR framework that reformulates volumetric SR as a reconstruction task using diffusion-based upsampled 2D projection priors. The proposed framework consists of two key stages: (1) LR CT projection SR, and (2) 3D reconstruction. First, given that acquiring 2D X-ray data is more feasible than collecting paired 3D HR-LR volumes, we adopt a diffusion model trained on large-scale 2D X-ray datasets. We then employ this diffusion prior within a Denoising Diffusion Null-space Model (DDNM) [30] to generate 2D HR CT projections from LR counterparts. This design choice intends LR projections to enforce data consistency, simultaneously enhancing the projection information through the diffusion model prior. Second, we introduce a Negative Alpha Blending Gaussian Splatting (NAB-GS) to reconstruct the HR 3D CT volume. NAB-GS learns the residual field between the diffusion-generated HR projections and the projections of the upsampled LR volume. To fully exploit this residual field, which inherently contains both positive and negative values, NAB-GS relaxes the non-negativity constraint of standard 3DGS [15], enabling precise encoding of signed residuals and substantially improving the recovery of structural details. Consequently, our framework achieves superior results on two public datasets, with expert evaluations suggesting its clinical potential at 4 $\times$. Our main contributions are summarized as follows:

- We propose a novel zero-shot 3D CT SR framework that reformulates volumetric SR as a 3D reconstruction task driven by diffusion-based upsampled 2D projection priors. Leveraging a diffusion model trained on large-scale X-ray datasets, we mitigate the need for paired HR-LR volumes and enhance the limited information of LR inputs.
- We introduce NAB-GS, a Gaussian splatting that relaxes the non-negativity constraint of standard 3DGS to learn signed residual fields between diffusion-generated HR projections and upsampled LR projections, enabling precise residual encoding and improved structural details.
- We validate our framework on the UHRCT [6] and MELA[†] public datasets, achieving superior quantitative and qualitative performance compared to

[†] <https://mela.grand-challenge.org/>




Fig. 1. Overview of our framework. (a) LR projection SR using diffusion model: A pre-trained diffusion model with 2D X-ray data is employed within the DDNM to generate HR 2D CT projection images from LR counterparts. (b) 3D CT reconstruction via NAB-GS: Using both positive and negative density Gaussians, we model a signed residual field between diffusion-generated HR projections and LR counterparts. For HR volume generation, the learned residual field is added onto the upsampled LR volume.

state-of-the-art (SOTA) zero-shot methods, alongside expert evaluations supporting its CT volume quality and clinical potential.

2 Zero-Shot 3D CT SR Framework

2.1 LR Projection SR using Diffusion Model

The first stage focuses on generating high-fidelity 2D CT projections to serve as reliable guidance for 3D reconstruction, as shown in Fig. 1(a). We train a diffusion model on large-scale 2D X-ray datasets specifically to model a generative projection prior. Subsequently, this prior is adapted to our CT projection SR task using DDNM [30], enhancing the limited information in LR CT projections.

SR using DDNM. Let $y = Ax$ denote an LR CT projection, where x is the unknown HR projection and A is the downsampling operator. Our goal is to estimate an HR projection $\hat{x}$ that both matches the measured LR projection and lies on the manifold modeled by the diffusion prior.

Following the null-space based diffusion formulation in [30], we combine the learned prior with a data-consistency constraint defined by y . At each reverse diffusion step, the estimate is decomposed into range and null-space components:

$$\hat{x}_{0|t} = A^\dagger y + (I - A^\dagger A)x_{0|t}, \quad (1)$$

where $\mathbf{A}^\dagger$ denotes the pseudo-inverse of $\mathbf{A}$ and $\mathbf{x}_{0|t}$ is the denoised estimate predicted by the diffusion model at step t . The first term enforces data consistency in the range space of $\mathbf{A}$, while the second term allows the diffusion prior to refine details in the null space, injecting high-frequency structures.

To account for measurement noise or slight inconsistencies in LR projections, we further adopt the DDNM+ [30]:

$$\hat{\mathbf{x}}_{0|t} = \mathbf{x}_{0|t} - \Sigma_t \mathbf{A}^\dagger (\mathbf{A} \mathbf{x}_{0|t} - \mathbf{y}), \quad (2)$$

where Σ_t controls the strength of the correction at each diffusion step, balancing data fidelity and perceptual quality. This null-space diffusion process allows us to generate HR CT projections that inherit realistic projection appearance from the trained diffusion prior while leveraging LR projection information to enforce data consistency. In this way, an X-ray trained diffusion prior coupled with DDNM offers a practical design to enhance the limited information in LR CT inputs under a zero-shot SR setting.

2.2 3D CT Reconstruction via NAB-GS

The second stage performs 3D CT reconstruction via the proposed NAB-GS, as shown in Fig. 1(b). We obtain an initial upsampled LR volume by cubic upsampling of the LR volume and initialize a radiative Gaussian field from this upsampled LR volume. To reconstruct the HR volume, we learn a residual field between the diffusion-generated HR projections and the upsampled LR baseline. The Gaussians are optimized via rasterizer and voxelizer (following R²-GS [33]) under reconstruction and total variation (TV) losses to refine structural details.

Negative alpha blending. Conventional 3DGS [15] enforces non-negativity of density ρ by softplus activation ($\rho = \ln(1 + e^z)$) and alpha blending, where z is the raw density. While physically plausible, this restricts the representation to strictly positive densities. However, our residual formulation targets the discrepancy between the diffusion-generated HR projection and the upsampled LR baseline. Since the upsampled LR can locally over- or underestimate true intensities, the residual inherently contains both positive and negative values.

To faithfully encode signed residuals, first, we replace the softplus with Parametric ReLU (PReLU), allowing for negative densities:

$$\phi(z) = \begin{cases} z, & \text{if } z \geq 0 \\ \gamma z, & \text{otherwise} \end{cases}, \quad \phi^{-1}(\rho) = \begin{cases} \rho, & \text{if } \rho \geq 0 \\ \frac{\rho}{\gamma}, & \text{otherwise} \end{cases}, \quad (3)$$

where $\phi(\cdot)$ and $\phi^{-1}(\cdot)$ denote the PReLU and its inverse, and γ is a learnable negative-slope parameter that controls the gradient magnitude. This regulation prevents divergence, stabilizing the optimization process across the residual field.

While z can now take negative values, standard alpha blending still constrains the accumulation to be non-negative for physical plausibility. Thus, we modify

the rendering formulation to permit negative contributions within the blending process. Following R²-GS, we model the linear integral of X-ray projection without the standard transmittance term:

$$C = \sum_{i=1}^N \alpha_i, \quad \text{subject to } (\alpha_i \geq \epsilon), \quad (4)$$

where α_i and ϵ denote the i -th Gaussian’s contribution and a heuristic threshold, respectively. To actively accommodate negative values, we specifically remove the strict non-negativity filter ($\alpha_i \geq \epsilon$), conventionally imposed to ensure physical plausibility, from Eq. (4). Despite this change, since our formulation is a linear accumulation, the gradient remains constant ($\frac{\partial C}{\partial \alpha_i} = 1$) regardless of whether α_i is positive or negative, ensuring stable optimization. Accordingly, we adapt our pruning process, discarding Gaussians only if $|z| < 10^{-5}$. Finally, the output HR volume is reconstructed by adding the learned residual field to the upsampled LR volume, followed by clipping ($\max(0, \cdot)$) to ensure physical plausibility. This mechanism, NAB-GS, can selectively amplify or suppress local intensities of the upsampled LR volume, successfully recovering fine structure details.

Loss function. To optimize 3D CT reconstruction, the total loss function is:

$$\mathcal{L}_{total} = \mathcal{L}_{recon} + \lambda_1 \mathcal{L}_{tv}, \quad \mathcal{L}_{recon} = \mathcal{L}_1(y, x) + \mathcal{L}_{res}(\hat{y}, \hat{x}) + \lambda_2 \mathcal{L}_{ssim}(y, x), \quad (5)$$

where y and x denote the diffusion-generated HR projection and our predicted projection, while $\hat{y}$ and $\hat{x}$ represent residual projection (*i.e.*, y – upsampled projection) and residual projection. The terms $\mathcal{L}_1(\cdot, \cdot)$, $\mathcal{L}_{res}(\cdot, \cdot)$, $\mathcal{L}_{ssim}(\cdot, \cdot)$, and $\mathcal{L}_{tv}(\cdot)$ correspond to the L1, L1-based residual, structural similarity index measure (SSIM), and TV loss function, respectively. The hyperparameters are empirically set to $\lambda_1 = 0.05$ and $\lambda_2 = 0.5$. Following R²-GS, we compute $\mathcal{L}_{tv}(\cdot)$ efficiently using randomly cropped 32^3 subvolumes during training.

3 Experiments

3.1 Setup

Dataset. [2D] We train our unconditional diffusion model on two large-scale X-ray datasets resized to 512×512 . ChestX-ray14[†] (112,120 frontal images) and CheXpert[†] (80,845 frontal and lateral images). **[3D]** Evaluation is performed on UHRCT [6] (20/10 train/test split) and MELA (60/20/20 train/val/test split). Volumes are resized to 512^3 , clipped to $[-512, 3071]$, and normalized to $[0, 1]$. LR volumes are generated via Gaussian smoothing followed by sinc interpolation. Note that our zero-shot learning uses only test sets.

Implementation details. [Diffusion] is conducted in PyTorch 1.11.0. We trained an unconditional DDPM [13] using four NVIDIA A6000 GPUs. For inference, we used DDIM [27] sampling with 50 steps, employing DDNM+ [30] for

[†] <https://nihcc.app.box.com/v/ChestXray-NIHCC>

[†] <https://stanfordmlgroup.github.io/competitions/chexpert/>

Table 1. Quantitative comparison of 3D CT SR with previous methods.

Methods	UHRCT 4×		UHRCT 8×		MELA 4×		MELA 8×	
	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑
Trilinear	24.56	<u>0.8877</u>	21.14	<u>0.8125</u>	33.64	<u>0.9472</u>	30.27	<u>0.9063</u>
Cubic	24.61	0.8783	21.46	0.8013	33.55	0.9431	<u>30.49</u>	0.9005
NeRF [20]	20.86	0.6745	18.92	0.5797	29.76	0.8088	28.73	0.7873
CuNeRF [4]	<u>25.25</u>	0.8459	21.04	0.7572	<u>33.76</u>	0.9096	30.11	0.8535
Ours	25.42	0.8957	21.96	0.8172	34.17	0.9525	30.81	0.9115
Supervised (ArSSR [32])	22.78	0.8936	21.10	0.8336	33.00	0.9658	30.70	0.9343

both 4× and 8× SR with a noise level of 0.0015. We upsampled 100 X-ray projections, uniformly spaced between 0° to 180°, from each LR volume using LR projection by TIGRE [1]. **[NAB-GS]** is implemented in PyTorch and trained for 5k iterations on a single RTX 3090 GPU. We adopt the default training parameter of R²-GS [33], except at 8×, where the learning rates for scaling (from 5×10^{-4} to 5×10^{-5}) and rotation (from 1×10^{-4} to 1×10^{-5}). Our framework takes around 15 minutes per volume (Diffusion 10 min, NAB-GS 5 min). Performance is evaluated using PSNR and SSIM [31].

3.2 Results and Evaluation

Quantitative results. Table 1 presents the comparison against conventional upsampling (Trilinear, Cubic), zero-shot (NeRF [20], CuNeRF [4]), and supervised (ArSSR [32]) methods, where CuNeRF is the only publicly available zero-shot baseline for 3D CT SR. ArSSR was trained on the train set of each dataset; NeRF, implemented as in CuNeRF, used standard ray sampling for SR, unlike CuNeRF’s cube sampling. Overall, our method achieves superior PSNR and SSIM across both the UHRCT and MELA datasets. Notably, when compared to CuNeRF, which relies solely on internal volume information, our approach yields substantially higher SSIM (+0.04-0.06). Compared to ArSSR, our zero-shot framework demonstrates competitive performance. Furthermore, our framework is computationally efficient, completing in 15 minutes, whereas CuNeRF takes about 1 hour and ArSSR takes about 25 minutes per volume.

Qualitative results. Figure 2 visualizes the slice of reconstructed HR volumes and their corresponding L2 error maps. Cubic (b) suffers from severe over-smoothing, and CuNeRF (d) introduces noticeable high-frequency artifacts, particularly at 8×. In contrast, our method (e) yields more faithful structural details. Notably, at 4× upsampling (rows 1 and 3), ours more precisely restores fine structures such as bone boundaries. These quantitative and qualitative results clearly demonstrate the effectiveness of the proposed method.

Experts’ evaluation. A user study with two domain experts was conducted to assess the clinical potential of the proposed method, comparing MELA (4×, 8×) volumes produced by three methods (Cubic, CuNeRF, and Ours). The experts commented that (1) our method at 4× shows potential for real-world clinical use, while further improvements are required at 8× and (2) ours yields sharper




Fig. 2. Visual comparisons of 3D CT reconstruction results. (a) Ground truth (GT), (b) Cubic interpolation, (c) ArSSR [32], (d) CuNeRF [4], and (e) Ours. The green box and zoom-in highlight regions where our method excels at reconstruction. The error map is computed by the L2 norm between the prediction and the ground truth.

results than Cubic, but the current visual gain does not substantially contribute to overall utility. They further noted that (3) it clearly outperforms CuNeRF by more accurately preserving fine structural details and (4) enhancing inter-slice consistency is an important direction for further improving clinical utility.

Ablation on UHRCT. (1) 2D projection SR. We evaluate the impact of 2D SR methods on 3D SR performance in Tab. 2. We compare our method against non-diffusion baselines, including bilinear interpolation, RDN [35], and Omni [28], as well as diffusion-based optimization methods, including GDP [9] and BIRD [5]. RDN and Omni were trained on projections generated from the


Table 2. Ablation on 2D projection SR methods on 3D CT SR performance.

Methods	UHRCT 4×		UHRCT 8×	
	PSNR↑	SSIM↑	PSNR↑	SSIM↑
Bilinear	24.57	0.8244	21.34	0.8015
RDN [35]	23.78	0.6410	21.41	0.6946
Omni [28]	22.74	0.5621	20.56	0.5992
BIRD [5]	20.18	0.8111	19.62	0.7845
GDP [9]	22.84	0.8485	21.00	0.8075
Ours	25.42	0.8957	21.96	0.8172

Table 3. Ablation on 3D reconstruction methods and activation functions.

Methods	UHRCT 4×		UHRCT 8×	
	PSNR↑	SSIM↑	PSNR↑	SSIM↑
FDK [10]	21.88	0.5375	19.39	0.4337
NAF [34]	24.23	0.8559	21.66	0.7982
SAX-NeRF [2]	24.81	0.8719	21.40	0.7826
R ² -GS [33]	24.89	0.8839	21.25	0.8056
Softplus	24.71	0.8790	21.61	0.8022
ReLU	24.62	0.8784	21.57	0.8021
Sine	25.19	0.8475	21.26	0.7208
Ours	25.42	0.8957	21.96	0.8172

UHRCT training set via TIGRE. RDN and Omni substantially degrade SSIM due to the scarcity of training data. While optimization-based methods (BIRD and GDP) with the same diffusion prior degrade both PSNR and SSIM, our method achieves clear gains in 3D SR performance, justifying our design choice.

**Fig. 3.** Visual comparisons of 3D CT reconstruction across activation functions. (a) Ground truth, (b) Softplus, (c) ReLU, (d) Sine, and (e) Ours (PReLU). Ours effectively enhances structural details while suppressing grainy noise.

(2) Effectiveness of NAB-GS. Table 3 validates NAB-GS. First, we compare to existing reconstruction methods (FDK [10], NAF [34], SAX-NeRF [2], and R²-GS [33]). Specifically, ours outperforms R²-GS in both PSNR (4×: +0.53, 8×: +0.71) and SSIM (4×: +0.0118, 8×: +0.0116), demonstrating the effectiveness of the residual learning-based strategy. Second, we compare alternative activations (softplus, ReLU, and sine) with PReLU. Although the raw density z can take negative values, softplus and ReLU strictly enforce non-negativity, failing to capture boundaries in overestimated regions (Fig. 3(b) and (c)). Although sine allows non-negative outputs, it introduces grainy artifacts as shown in Fig. 3(d). In contrast, PReLU leverages a learnable negative slope to accommodate negative values and achieves the best performance as in Fig. 3(e).

4 Conclusion

In this paper, we have proposed a novel zero-shot 3D CT SR framework that reformulates volumetric SR as a 3D reconstruction task driven by diffusion-

upsampled 2D projection priors. Furthermore, we introduced NAB-GS, which represents the negative densities to learn signed residual fields that enhance structural details. Our framework outperforms previous zero-shot methods in both quantitative and qualitative results on two public datasets. Moreover, an experts' evaluation highlights the clinical potential of the proposed method, especially at $4\times$. For future work, we will enhance inter-slice continuity for practical clinical use and conduct evaluations on real-world clinical data.

Acknowledgments. This work was supported in part by the National Research Foundation of Korea under Grant RS-2024-00349697 and Grant RS-2021-NR060143; in part by the Institute for Information and Communications Technology Planning and Evaluation under Grant IITP-2026-RS-2020-II201819; in part by the Technology Development Program funded by the Ministry of SMEs and Startups (MSS), South Korea, under Grant RS-2024-00437796; in part by the National Research Council of Science and Technology (NST) grant funded by Korean Government [Ministry of Science and Information and Communications Technology (MSIT)] under Grant GTL24031-000; and in part by Korea University Grant.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Biguri, A., Dosanjh, M., Hancock, S., Soleimani, M.: TIGRE: a MATLAB-GPU toolbox for CBCT image reconstruction. *Biomed. Phys. Eng. Express (BPEE)* **2**(5), 055010 (2016)
2. Cai, Y., Wang, J., Yuille, A., Zhou, Z., Wang, A.: Structure-aware sparse-view x-ray 3d reconstruction. In: *Proc. Comput. Vis. Pattern Recognit. (CVPR)*. pp. 11174–11183 (2024)
3. Chen, X., Wang, X., Zhou, J., Qiao, Y., Dong, C.: Activating more pixels in image super-resolution transformer. In: *Proc. Comput. Vis. Pattern Recognit. (CVPR)*. pp. 22367–22377 (2023)
4. Chen, Z., Yang, L., Lai, J.H., Xie, X.: CuNeRF: Cube-based neural radiance field for zero-shot medical image arbitrary-scale super resolution. In: *Proc. Comput. Vis. Pattern Recognit. (CVPR)*. pp. 21185–21195 (2023)
5. Chihaoui, H., Lemkhenter, A., Favaro, P.: Blind image restoration via fast diffusion inversion. *Adv. Neural Inf. Process. Syst. (NeurIPS)* **37**, 34513–34532 (2024)
6. Chu, Y., Zhou, L., Luo, G., Qiu, Z., Gao, X.: Topology-preserving computed tomography super-resolution based on dual-stream diffusion model. In: *Proc. Int'l Conf. Med. Image Comput. Comput. Assist. Interv. (MICCAI)*. pp. 260–270 (2023)
7. Dong, C., Loy, C.C., He, K., Tang, X.: Image super-resolution using deep convolutional networks. *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)* **38**(2), 295–307 (2015)
8. Fang, W., Tang, Y., Guo, H., Yuan, M., Mok, T.C., Yan, K., Yao, J., Chen, X., Liu, Z., Lu, L., et al.: CycleINR: Cycle implicit neural representation for arbitrary-scale volumetric super-resolution of medical data. In: *Proc. Comput. Vis. Pattern Recognit. (CVPR)*. pp. 11631–11641 (2024)

9. Fei, B., Lyu, Z., Pan, L., Zhang, J., Yang, W., Luo, T., Zhang, B., Dai, B.: Generative diffusion prior for unified image restoration and enhancement. In: Proc. Comput. Vis. Pattern Recognit. (CVPR). pp. 9935–9946 (2023)
10. Feldkamp, L.A., Davis, L.C., Kress, J.W.: Practical cone-beam algorithm. *Journal of the Optical Society of America A* **1**(6), 612–619 (1984)
11. Frazer, L.L., Louis, N., Zbijewski, W., Vaishnav, J., Clark, K., Nicolella, D.P.: Super-resolution of clinical CT: Revealing microarchitecture in whole bone clinical CT image data. *Bone* **185**, 117115 (2024)
12. Gharib, O.A., Remedios, S.W., Dewey, B.E., Prince, J.L., Carass, A.: Exploring the feasibility of zero-shot super-resolution in preclinical imaging. In: Proc. Int’l Conf. Med. Image Comput. Comput. Assist. Interv. (MICCAI). pp. 186–196 (2025)
13. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Adv. Neural Inf. Process. Syst. (NeurIPS)* **33**, 6840–6851 (2020)
14. Kavar, B., Elad, M., Ermon, S., Song, J.: Denoising diffusion restoration models. *Adv. Neural Inf. Process. Syst. (NeurIPS)* **35**, 23593–23606 (2022)
15. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. on Graph. (ToG)* **42**(4), 139–1 (2023)
16. Liang, J., Cao, J., Sun, G., Zhang, K., Van Gool, L., Timofte, R.: SwinIR: Image restoration using swin transformer. In: Proc. of Int’l Conf. on Comput. Vis. (ICCV). pp. 1833–1844 (2021)
17. Liu, G., Sun, H., Li, J., Yin, F., Yang, Y.: Accelerating diffusion models for inverse problems through shortcut sampling. *arXiv preprint arXiv:2305.16965* (2023)
18. Martin, D.R., Semelka, R.C.: Health effects of ionising radiation from diagnostic CT. *The Lancet* **367**(9524), 1712–1714 (2006)
19. McCollough, C.H., Yu, L., Kofler, J.M., Leng, S., Zhang, Y., Li, Z., Carter, R.E.: Degradation of ct low-contrast spatial resolution due to the use of iterative reconstruction and reduced dose levels. *Radiology* **276**(2), 499–506 (2015)
20. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
21. Noh, J., Jeon, W., Park, J.: Dual interaction network with cross-image attention for medical image segmentation. *Pattern Recognition Letters (PRL)* (2025)
22. Oh, H.J., Jeong, W.K.: Diffmix: Diffusion model-based data synthesis for nuclei segmentation and classification in imbalanced pathology image datasets. In: Proc. Int’l Conf. Med. Image Comput. Comput. Assist. Interv. (MICCAI). pp. 337–345 (2023)
23. Peng, C., Lin, W.A., Liao, H., Chellappa, R., Zhou, S.K.: Saint: Spatially aware interpolation network for medical slice synthesis. In: Proc. Comput. Vis. Pattern Recognit. (CVPR). pp. 7750–7759 (2020)
24. Saharia, C., Ho, J., Chan, W., Salimans, T., Fleet, D.J., Norouzi, M.: Image super-resolution via iterative refinement. *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)* **45**(4), 4713–4726 (2022)
25. Shocher, A., Cohen, N., Irani, M.: Zero-shot super-resolution using deep internal learning. In: Proc. Comput. Vis. Pattern Recognit. (CVPR). pp. 3118–3126 (2018)
26. Soh, J.W., Cho, S., Cho, N.I.: Meta-transfer learning for zero-shot super-resolution. In: Proc. Comput. Vis. Pattern Recognit. (CVPR). pp. 3516–3525 (2020)
27. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502* (2020)

28. Wang, H., Chen, X., Ni, B., Liu, Y., Liu, J.: Omni aggregation networks for lightweight image super-resolution. In: Proc. Comput. Vis. Pattern Recognit. (CVPR). pp. 22378–22387 (2023)
29. Wang, J., Wang, R., Tao, R., Zheng, G.: UASSR: Unsupervised arbitrary scale super-resolution reconstruction of single anisotropic 3d images via disentangled representation learning. In: Proc. Int’l Conf. Med. Image Comput. Comput. Assist. Interv. (MICCAI). pp. 453–462 (2022)
30. Wang, Y., Yu, J., Zhang, J.: Zero-shot image restoration using denoising diffusion null-space model. arXiv preprint arXiv:2212.00490 (2022)
31. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE Trans. Image Process. (TIP) **13**(4), 600–612 (2004)
32. Wu, Q., Li, Y., Sun, Y., Zhou, Y., Wei, H., Yu, J., Zhang, Y.: An arbitrary scale super-resolution approach for 3d MR images via implicit neural representation. IEEE J. Biomed. Health Inform. (JBHI) **27**(2), 1004–1015 (2022)
33. Zha, R., Lin, T.J., Cai, Y., Cao, J., Zhang, Y., Li, H.: R²-Gaussian: Rectifying radiative gaussian splatting for tomographic reconstruction. Adv. Neural Inf. Process. Syst. (NeurIPS) (2024)
34. Zha, R., Zhang, Y., Li, H.: NAF: Neural attenuation fields for sparse-view CBCT reconstruction. In: Proc. Int’l Conf. Med. Image Comput. Comput. Assist. Interv. (MICCAI). pp. 442–452 (2022)
35. Zhang, Y., Tian, Y., Kong, Y., Zhong, B., Fu, Y.: Residual dense network for image super-resolution. In: Proc. Comput. Vis. Pattern Recognit. (CVPR). pp. 2472–2481 (2018)