This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Towards Multi-Scenario Generalization: Text-Guided Unified Framework for Low-Dose CT and Total-Body PET Reconstruction

Weitao Wang[1], Yanyan Huang[2], Shunjie Dong[3], Le Xue[4], Kuangyu Shi[5], and Yu Fu[1(✉)]

[1] Lanzhou University, Lanzhou, China
fuyu@lzu.edu.cn
[2] The University of Hong Kong, Hong Kong, China
[3] Shanghai Jiao Tong University, Shanghai, China
[4] Fudan University, Shanghai, China
[5] University of Bern, Bern, Switzerland

Abstract. Low-dose computed tomography (LDCT) and low-dose positron emission tomography (LDPET) imaging substantially reduce radiation exposure compared to their normal-dose counterparts, mitigating health risks such as elevated cancer incidence. However, the resulting LDCT and total-body LDPET images are often compromised by noise and artifacts stemming from photon starvation and electronic interference. While supervised reconstruction methods have tackled challenges like over-smoothing and training instability, their generalization is hindered by variations in imaging devices, dosage levels, and modality-specific characteristics. Recent advances in text-guided models have augmented traditional deep learning techniques, offering greater adaptability. Building on this, we propose a Text-guided Unified Framework (TUF) for high-precision reconstruction of LDCT and total-body LDPET images. Leveraging insights from cold diffusion paradigms, TUF introduces a novel mean-preserving degradation operator to model the physical process of image degradation. Additionally, we design a dual-domain fusion network that converts textual inputs into scaling and shifting factors, enabling seamless integration of text cues at each timestep. Extensive experiments across four publicly available datasets reveal that TUF surpasses state-of-the-art methods in both reconstruction quality and generalization across LDCT and total-body LDPET imaging scenarios. The code will be available at TUF-code.

Keywords: Low-dose CT · Low-dose PET · Image Reconstruction.

1 Introduction

Low-dose computed tomography (LDCT) and total-body low-dose positron emission tomography (LDPET) are vital for radiation-sensitive applications, including pediatric cancer screening and personalized dosimetry [23,8]. Compared to

their normal-dose CT (NDCT) and normal-dose PET (NDPET), these low-dose modalities significantly reduce radiation exposure while maintaining diagnostic capabilities. Yet, their clinical adoption is impeded by conventional reconstruction methods, which rely on isotropic Gaussian priors or single-domain processing and struggle to mitigate non-stationary quantum noise and aliasing artifacts prevalent in ultra-low photon counts or rapid-scan protocols [4,9,15]. This challenge is particularly pronounced in complex anatomical regions—such as the pulmonary hila or hepatic vasculature—where existing methods often blur critical microcalcifications (<1 mm) or submillimeter pulmonary nodules, key indicators of early malignancies [12,13]. Although recent deep learning approaches achieve high peak signal-to-noise ratios (PSNR), their lack of physics-based constraints and semantic grounding often results in clinically implausible reconstructions, especially across heterogeneous imaging systems (e.g., dual-source vs. photon-counting CT) or patient-specific anatomies [1,3]. These shortcomings highlight a pressing need for a reconstruction paradigm that transcends the limitations of task-specific models and integrates domain knowledge into a cohesive, adaptable framework.

To this end, we propose the Text-guided Unified Framework (TUF), a novel framework that achieves adaptive image reconstruction through text-conditioned cross-modal guidance, where clinical text semantics dynamically regulate reconstruction fidelity for anatomically accurate results. At its core, a cold diffusion paradigm employs a mean-preserving degradation operator to capture the first-order moment characteristics of LDCT and LDPET degradation, inverting traditional diffusion by treating the noisy low-dose image as the endpoint. This physics-informed strategy is paired with a dual-stream network that balances spatial and frequency-domain processing, preserving fine details while correcting global artifacts. Crucially, the Text Interaction Guidance Architecture (TIGA) implements CLIP-derived semantic conditioning with clinical linguistic priors, establishing an interpretable pathway for prompt-guided artifact correction. Extensive validation on clinical benchmarks establishes TUF as a new standard for unified, high-fidelity low-dose imaging reconstruction.

2 Methods

TUF is a radiation-aware framework designed for diagnostic-quality LDCT and LDPET reconstruction (Fig. 1). It integrates three core components: (1) a Cold Diffusion Sampling Process (CDSP) with learnable transition kernels, (2) a Dual-domain Reconstruction Network (DDRN) fusing Fourier and spatial features, and (3) a Text Interaction Guidance Architecture (TIGA) leveraging clinical linguistic priors through cross-modal alignment.

2.1 Cold Diffusion Sampling Process (CDSP)

The CDSP redefines diffusion modeling for low-dose imaging by introducing a statistically constrained degradation operator that preserves first-order moments, tailored to the structured noise in LDCT and LDPET. Unlike traditional

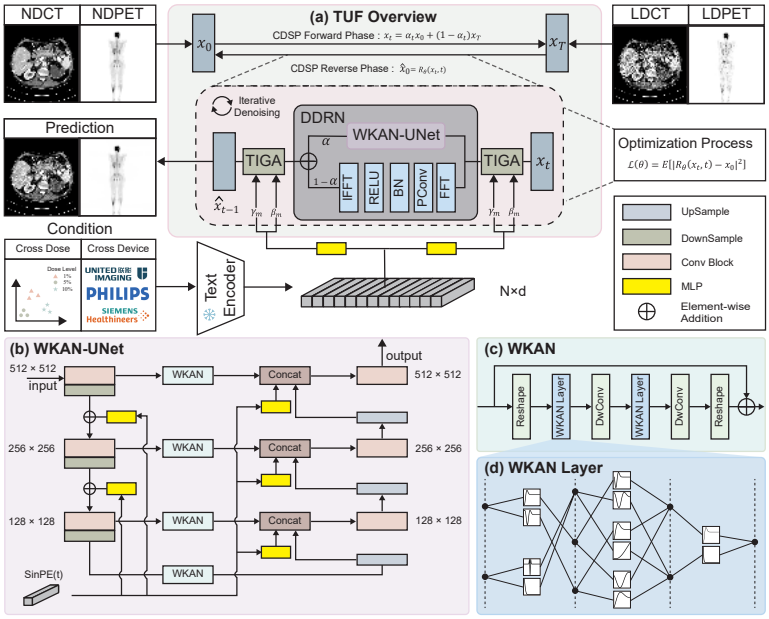

Fig. 1. The overall architecture of the proposed TUF network. (a) TUF overview; (b) WKAN-UNet framework; (c) WKAN framework.

diffusion models that add isotropic noise, CDSP simulates degradation from a full-dose image x_0 to a noisy low-dose image x_t. Formally, the forward process is:

$$x_t = \alpha_t x_0 + (1 - \alpha_t)x_T, \quad t = 1, \ldots, T, \tag{1}$$

where $\alpha_t \in (0, 1)$ decreases over time, blending x_0 into x_T. This inverted approach aligns with clinical imaging physics, enabling the reverse process to reconstruct high-quality images from noisy inputs. The reconstruction is performed by a parameterized operator R_θ, where the predicted full-dose image $\hat{x}_0$ can be formulated as:

$$\hat{x}_0 = R_\theta(x_t, t) \tag{2}$$

2.2 Dual-domain Reconstruction Network (DDRN)

The DDRN is engineered to reconstruct low-dose images with high fidelity by processing them in both frequency and spatial domains. This dual-stream approach leverages the complementary strengths of each domain: the frequency-domain stream excels at capturing global patterns and mitigating pervasive noise artifacts, while the spatial-domain stream preserves local details and textures critical for medical image interpretation. In the frequency-domain stream, we apply the Discrete Fourier Transform (DFT) to the input low-dose image x_t:

$$x_{\mathcal{R}}, x_{\mathcal{I}} = \mathcal{F}(TIGA(x_t)), \tag{3}$$

where $x_{\mathcal{R}}$ and $x_{\mathcal{I}}$ denote the real and imaginary components, respectively. These components are concatenated and processed through a phase-aware convolutional network: $\widehat{x_{\mathcal{R}}}, \widehat{x_{\mathcal{I}}} = \sigma \cdot \mathrm{BN}(\mathrm{PConv}(\mathrm{Concat}(x_{\mathcal{R}}, x_{\mathcal{I}})))$,where σ denotes the ReLU function, BN denotes batch normalization, and PConv denotes phase-aware convolution. Then, the reconstructed image $x_{\text{frequency}}$ output from the frequency-domain branch can be defined as:

$$x_{\text{frequency}} = \mathcal{F}^{-1}(\widehat{x_{\mathcal{R}}}, \widehat{x_{\mathcal{I}}}), \tag{4}$$

where $\mathcal{F}^{-1}$ is the inverse DFT

In the spatial domain, we propose WKAN-UNet, a novel architecture that integrates Wavelet-Kolmogorov-Arnold Networks (WKAN) into skip connection paths to enhance interpretability and model nonlinear dependencies[11]. The WKAN modules leverage wavelet-based operations to perform multi-scale nonlinear transformations, effectively refining both spatial details and semantic representations. As illustrated in Fig. 1 (c), each WKAN module comprises depthwise convolution (DWConv) followed by N stacked WKAN-Layers. The mathematical formulation of a WKAN with N layers can be expressed as:

$$WKAN(I) = \Phi_{N-1} \circ \Phi_{N-2} \cdots \Phi_1 \Phi_0, \tag{5}$$

where I is the input feature map, and Φ_i denotes the i-th WKAN layer. The transition between layers is defined as: $I_{i+1} = \mathrm{DWConv}(\Phi_i(I_i; w, \sigma, b))$,with the WKAN layer function given by:

$$\Phi(x; w, \sigma, b) = w \cdot \left((x-b)^2 - 1\right) e^{-\frac{(x-b)^2}{2\sigma^2}}. \tag{6}$$

Here, w, σ, and b are learnable parameters that adapt the transformation to the data. The final reconstructed image is obtained by fusing the outputs from the frequency-domain and spatial-domain branches, controlled by a learnable fusion parameter α, which balances the contributions from both streams:

$$\hat{x}_{\text{DDRN}} = \alpha \cdot \hat{x}_{\text{WKAN-UNet}} + (1-\alpha) \cdot x_{\text{frequency}}, \tag{7}$$

where $\hat{x}_{\text{WKAN-UNet}}$ is the reconstruction result from the spatial domain branch.

2.3 Text Interaction Guidance Architecture (TIGA)

We introduce the TIGA, a module designed to integrate clinical linguistic priors into image reconstruction through semantic-aware feature modulation. TIGA utilizes frozen CLIP embeddings $\mathbf{F}_{\text{text}} = \mathcal{F}_T(T_{\text{text}})$ from user-specified text prompts, exploiting the model's cross-modal alignment capabilities established during contrastive pretraining. Note that $\mathbf{F}_{\text{text}}$ denotes the text features generated by encoding the text prompt using the CLIP model. We define a nonlinear projection module to process these embeddings to derive dynamic modulation parameters:

$$\{\gamma_m, \beta_m\} = \Phi_m(\mathbf{F}_{\text{text}}), \tag{8}$$

where Φ_m consists of learnable multi-layer perceptrons (MLPs), producing scaling factors γ_m and bias terms β_m. These parameters adaptively recalibrate the fusion features as follows:

$$\hat{\mathbf{F}}_f = (1 + \gamma_m) \odot \hat{x} + \beta_m, \tag{9}$$

where $\hat{x}$ denotes the input features, $\odot$ represents element-wise multiplication, and $\hat{\mathbf{F}}_f$ denotes the modulated output features, enriched with textual semantics while preserving structural integrity.

3 Experiments

To demonstrate the effectiveness of the proposed TUF for the two different modalities, we perform the evaluation on four datasets: two open-access (OA) LDCT datasets and two OA LDPET datasets. Following, we conduct experiment under two different experimental settings: **(a) all-in-one experiments:** both LDCT and LDPET images, along with diverse dose level descriptions and reconstruction task descriptions provided by clinical scenarios.; **(b) single-task experiments:** single modality, single low-dose level, without incorporating diverse clinical scenario text descriptions.

Table 1. Comparisons under all-in-one reconstruction setting for CT-Dataset.

Method	**20%(Mayo-2020)**	**25%(Mayo-2016)**
CLIP [2]	30.36/0.725	41.14/0.959
Restore-RWKV [18]	30.29/0.718	41.58/0.960
Restormer [20]	30.81/0.729	40.91/0.950
PromptIR [16]	31.93/0.766	42.38/0.963
TUF (Ours)	**33.04/0.795**	**43.08/0.965**

Table 2. Comparisons under all-in-one reconstruction setting for Bern and UI-Datasets.

Method	**Bern-Dataset**		**UI-Dataset**	
	5%	**10%**	**5%**	**10%**
CLIP [2]	38.13/0.887	39.36/0.943	35.29/0.846	37.01/0.855
Restore-RWKV [18]	37.11/0.885	39.77/0.938	34.30/0.819	36.88/0.852
Restormer [20]	36.88/0.852	39.12/0.929	33.83/0.804	38.10/0.897
PromptIR [16]	38.12/0.885	39.77/0.935	34.97/0.835	37.11/0.876
TUF (Ours)	**39.76/0.943**	**41.09/0.949**	**36.49/0.866**	**38.43/0.889**

3.1 Dataset and Training Details

Mayo2016 and Mayo2020 Datasets. The experimental design employs the '2016 NIH-AAPM-Mayo Clinic Low-Dose CT Grand Challenge' dataset containing 5,936 normal-dose CT slices and corresponding quarter-dose counterparts (1mm slice thickness) from 10 subjects. Testing data consists of LDCT slices from patient L506, while the remaining 9 patients' data are partitioned into 80% training and 20% validation subsets. The Mayo2020 dataset expands this scope with 299 patient cases featuring dual-dose acquisitions [14]. Our selection strategy includes 36 20% dose samples (19 thoracic and 17 abdominal) for model training, complemented by a dedicated test set containing 1,562 abdominal slices reconstructed from 20% dose data.

MICCAI2022 Challenge Datasets. The Bern Dataset, acquired from the Department of Nuclear Medicine, University of Bern, Switzerland, using the Siemens Biograph Vision Quadra, and the UI Dataset, obtained from the Department of Nuclear Medicine, Ruijin Hospital, Shanghai Jiao Tong University School of Medicine, using the United Imaging uEXPLORER, both provide raw LDPET projections across three reduced-dose levels (5% and 10% of standard dose). Further details on these datasets can be found at: `https://ultra-low-dose-pet.grand-challenge.org`.

Training Details. The TUF training utilized a learning rate of 2×10^{-4}, a batch size of 7, and a GPU memory consumption of 21.33 GB, with the Adam optimizer ensuring efficient convergence. All datasets were divided into 80% for training, 10% for validation, and 10% for testing. Reconstruction quality of each image was assessed using peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM).

3.2 All-in-one experiment results

In comprehensive evaluations of multi-task medical image reconstruction (i.e., all-in-one experiment), the TUF demonstrates state-of-the-art performance across diverse datasets and degradation scenarios. On the Mayo-2020 CT dataset under 20% dose conditions (see Table 1), our approach achieves PSNR of 33.04 dB and SSIM of 0.7949, significantly surpassing PromptIR (31.93/0.766) and other baselines by significant margins. At 25% dose level, we further set a new benchmark with PSNR of 43.08 dB and SSIM of 0.9652, outperforming all all-in-one competitors.

For PET reconstruction tasks on the Bern and UI datasets, TUF exhibits exceptional robustness under extreme low-dose conditions (5% dose level), attaining 39.76/0.943 (Bern) and 36.49/0.866 (UI) (see Table 2), which exceed the second-best method (CLIP) by 1.63 dB and 1.2 dB, respectively. Our method advances medical imaging by effectively restoring anatomical details and suppressing noise artifacts, as evidenced in Fig. 2. Quantitative metrics and visualizations confirm that our framework not only generalizes across modalities and dose levels but also addresses the high-precision demands of clinical scenarios, establishing a new technical benchmark for all-in-one low-dose medical image reconstruction.

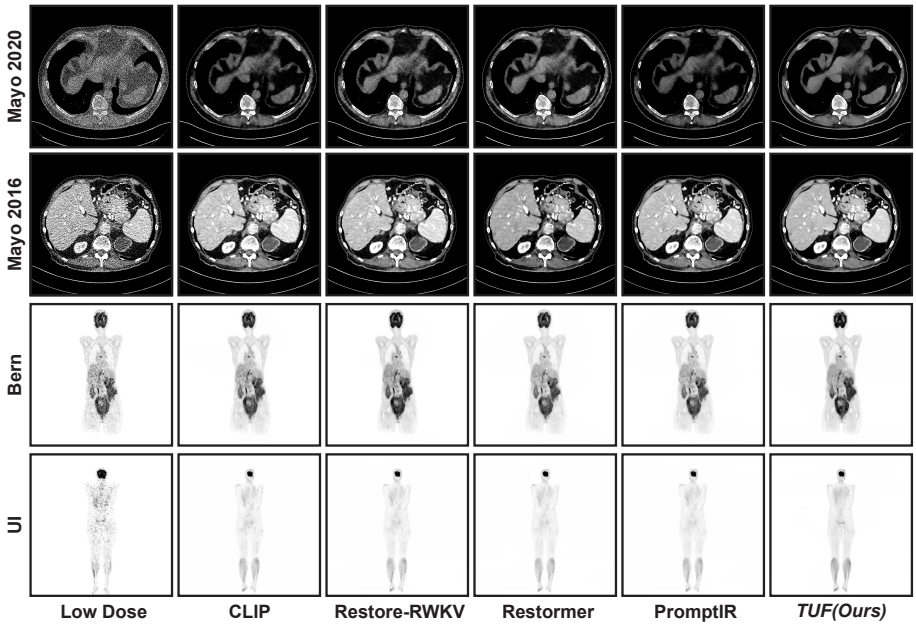

Fig. 2. Visual and quantitative comparisons of TUF against 4 competing models on the LDCT and total-body LDPET datasets.

3.3 Single-task experiment results

To assess the effectiveness of TUF in specialized reconstruction scenarios, we further evaluate it in single-task settings, comparing its performance against recent state-of-the-art methods developed for single-task reconstruction. On the Mayo-2016 CT dataset, the proposed TUF achieves PSNR of 43.09 dB and SSIM of 0.967 (see Table 3), surpassing CoreDiff (41.27/0.958) and UNAD (41.28/0.958) by 1.82 dB and 1.81 dB, respectively. For the Bern dataset under extreme low-dose conditions (5% dose), the TUF attains 40.56 dB PSNR and 0.959 SSIM, outperforming 3D DDPM (37.61/0.898) by 2.95 dB and WF-Diff (38.04/0.912) by 2.52 dB (see Table 3). These results demonstrate consistent superiority across diverse modalities (CT and PET) and dose levels, highlighting its adaptability to specialized clinical reconstruction tasks.

3.4 Ablation Studies

We conduct systematic ablations to validate the contributions of key components in the TUF framework. As shown in Table 4, removing the TIGA module degrades the performance of TUF by up to 0.96 dB PSNR on the UI dataset (5% dose), highlighting its critical role in low-dose feature enhancement. Omitting the WKAN-UNet backbone reduces reconstruction quality across all scenarios, particularly under 20% CT dose (PSNR drops from 33.04 to 32.71), confirming its

Table 3. Reconstruction results in the single-task setting on Mayo2016 and Bern 5% dose datasets.

Mayo2016 Dataset					
Method	LDMANet [21]	CoreDiff [5]	UNAD [7]	RDDM [10]	**TUF** (Ours)
PSNR	40.89	41.27	41.28	41.12	**43.09**
SSIM	0.956	0.958	0.958	0.949	**0.967**
Bern 5% Dose Dataset					
Method	DDPM [6]	WF-Diff [22]	AMIR [17]	3D DDPM [19]	**TUF**(Ours)
PSNR	38.20	38.04	37.88	37.61	**40.56**
SSIM	0.914	0.912	0.901	0.898	**0.959**

Table 4. Ablation experiments of TUF on mayo and UI datasets.

Models	CT-dataset		UI-dataset	
	20%	25%	5%	10%
TUF	33.04/0.795	43.08/0.965	36.49/0.866	38.43/0.895
Ablation for Modules				
TUF w/o TIGA	32.61/0.789	42.60/0.963	35.53/0.852	37.80/0.890
TUF w/o WKAN-UNet	32.71/0.791	42.77/0.965	35.59/0.854	37.94/0.884
TUF w/o frequency stream	32.83/0.797	42.86/0.966	35.73/0.851	38.07/0.888
Feedforward Type				
TUF w/ MLP	32.44/0.787	42.35/0.966	35.87/0.852	37.61/0.879
TUF w/ KAN	32.33/0.782	42.52/0.966	35.94/0.849	37.79/0.889
WKAN layer number				
TUF w/ 1	32.94/0.798	43.03/0.966	36.22/0.869	38.39/0.894
TUF w/ 3	32.99/0.795	43.05/0.966	36.38/0.867	38.44/0.896

effectiveness in multi-scale modeling. The frequency stream further contributes to artifact suppression, as evidenced by a 0.76 dB improvement on the Bern dataset. Comparisons of feedforward types demonstrate that our WKAN layers outperform MLP and KAN by 0.6–0.8 dB, while experiments on layer depth reveal that two WKAN layers achieve an optimal balance between complexity and performance. These results collectively validate the necessity of each proposed component in TUF.

4 Conclusion

This study presents the TUF to address the generalization challenges of existing LDCT and LDPET reconstruction models, particularly across varying dose levels, device specifications, and anatomical variations. By integrating a novel cold diffusion sampling process with a dual-domain fusion network, TUF

dynamically adapts its reconstruction strategies through input-driven feature analysis and textual guidance, eliminating the need for prior noise characterization. Extensive evaluations across multiple datasets demonstrate the framework's effectiveness, achieving significant performance improvements in both LDCT and LDPET tasks. TUF not only surpasses current state-of-the-art methods in terms of PSNR and SSIM, but also outperforms both all-in-one and single-task approaches by substantial margins. These results underscore TUF's potential for broad clinical deployment, offering a powerful solution for low-dose medical imaging that enhances diagnostic accuracy while reducing radiation exposure.

Acknowledgments This work was supported in part by the Gansu Provincial Key Research and Development Program (Grant No. 25YFWA005), and in part by the Open Research Fund of the State Key Laboratory of Brain-Machine Intelligence, Zhejiang University (Grant No. BMI2400002).

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

References

1. Cai, Q., Triphuridet, N., Zhu, Y., You, N., Yip, R., Yankelevitz, D.F., Henschke, C.I.: Bronchiectasis in low-dose ct screening for lung cancer. Radiology **304**(2), 437–447 (2022)
2. Cheng, J., Liang, D., Tan, S.: Transfer clip for generalizable image denoising. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 25974–25984 (2024)
3. Fu, Y., Dong, S., Huang, Y., Niu, M., Ni, C., Yu, L., Shi, K., Yao, Z., Zhuo, C.: Mpgan: multi pareto generative adversarial network for the denoising and quantitative analysis of low-dose pet images of human brain. Medical Image Analysis **98**, 103306 (2024)
4. Fu, Y., Dong, S., Niu, M., Xue, L., Guo, H., Huang, Y., Xu, Y., Yu, T., Shi, K., Yang, Q., et al.: Aigan: Attention–encoding integrated generative adversarial network for the reconstruction of low-dose ct and low-dose pet images. Medical Image Analysis **86**, 102787 (2023)
5. Gao, Q., Li, Z., Zhang, J., Zhang, Y., Shan, H.: Corediff: Contextual error-modulated generalized diffusion model for low-dose ct denoising and generalization. IEEE Transactions on Medical Imaging (2023)
6. Gong, K., Johnson, K., El Fakhri, G., Li, Q., Pan, T.: Pet image denoising based on denoising diffusion probabilistic model. European Journal of Nuclear Medicine and Molecular Imaging **51**(2), 358–368 (2024)
7. Gu, L., Deng, W., Wang, G.: Unad: Universal anatomy-initialized noise distribution learning framework towards low-dose ct denoising. In: ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 1671–1675. IEEE (2024)
8. Hricak, H., Mayerhoefer, M.E., Herrmann, K., Lewis, J.S., Pomper, M.G., Hess, C.P., Riklund, K., Scott, A.M., Weissleder, R.: Advances and challenges in precision imaging. The Lancet Oncology **26**(1), e34–e45 (2025)

9. Li, M., Wang, J., Chen, Y., Tang, Y., Wu, Z., Qi, Y., Jiang, H., Zheng, J., Tsui, B.M.: Low-dose ct image synthesis for domain adaptation imaging using a generative adversarial network with noise encoding transfer learning. IEEE transactions on medical imaging **42**(9), 2616–2630 (2023)
10. Liu, J., Wang, Q., Fan, H., Wang, Y., Tang, Y., Qu, L.: Residual denoising diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2773–2783 (June 2024)
11. Liu, Z., Wang, Y., Vaidya, S., Ruehle, F., Halverson, J., Soljačić, M., Hou, T.Y., Tegmark, M.: Kan: Kolmogorov-arnold networks. arXiv preprint arXiv:2404.19756 (2024)
12. Mascalchi, M., Zompatori, M.: Mediastinal lymphadenopathy in lung cancer screening: a red flag (2022)
13. Mazzone, P.J., Lam, L.: Evaluating the patient with a pulmonary nodule: a review. Jama **327**(3), 264–273 (2022)
14. McCollough, C., Chen, B., Holmes III, D., Duan, X., Yu, Z., Yu, L., Leng, S., Fletcher, J.: Low dose ct image and projection data (ldct-and-projection-data). The Cancer Imaging Archive pp. 1–10 (2020)
15. Meng, M., Wang, Y., Zhu, M., Tao, X., Mao, Z., Liao, J., Bian, Z., Zeng, D., Ma, J.: Ddt-net: Dose-agnostic dual-task transfer network for simultaneous low-dose ct denoising and simulation. IEEE Journal of Biomedical and Health Informatics (2024)
16. Potlapalli, V., Zamir, S.W., Khan, S., Khan, F.: Promptir: Prompting for all-in-one image restoration. In: Thirty-seventh Conference on Neural Information Processing Systems (2023)
17. Yang, Z., Chen, H., Qian, Z., Yi, Y., Zhang, H., Zhao, D., Wei, B., Xu, Y.: All-in-one medical image restoration via task-adaptive routing. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 67–77. Springer (2024)
18. Yang, Z., Zhang, H., Zhao, D., Wei, B., Xu, Y.: Restore-rwkv: Efficient and effective medical image restoration with rwkv. arXiv preprint arXiv:2407.11087 (2024)
19. Yu, B., Ozdemir, S., Dong, Y., Shao, W., Shi, K., Gong, K.: Pet image denoising based on 3d denoising diffusion probabilistic model: Evaluations on total-body datasets. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 541–550. Springer (2024)
20. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: Restormer: Efficient transformer for high-resolution image restoration. In: CVPR (2022)
21. Zhang, J., Ye, L., Gong, W., Chen, M., Liu, G., Cheng, Y.: A novel network for low-dose ct denoising based on dual-branch structure and multi-scale residual attention. Journal of Imaging Informatics in Medicine pp. 1–20 (2024)
22. Zhao, C., Cai, W., Dong, C., Hu, C.: Wavelet-based fourier information interaction with frequency diffusion adjustment for underwater image restoration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8281–8291 (2024)
23. Zhou, X., Fu, Y., Dong, S., Li, L., Xue, S., Chen, R., Huang, G., Liu, J., Shi, K.: Intelligent ultrafast total-body pet for sedation-free pediatric [18f] fdg imaging. European journal of nuclear medicine and molecular imaging **51**(8), 2353–2366 (2024)