DSCE-Net: Dual-Scale Conditioned Evolving Network with Balanced Optimization for All-in-One Medical Image Restoration

Xulin Hu, Yi Chen(✉), Xianming Gu, Ying Cao, and Lihui Wang

Key Laboratory of Advanced Medical Imaging and Intelligent Computing of Guizhou Province, College of Computer Science and Technology, Guizhou University, Guiyang, China

Abstract. All-in-one medical image restoration (MedIR) recoveries diverse degraded multimodal medical images with a single model, improving efficiency and deployment. However, its performance is often hindered by cross-modality task interference and optimization imbalance across tasks. To address these challenges, We propose a Dual-Scale Conditioned Evolving Network (DSCE-Net) with balanced optimization for task-adaptive restoration across multiple modalities and degradations, which dynamically evolves its parameters and feature representations during training to accommodate the varying complexities of different tasks. Specifically, we introduce a frequency-prompt-conditioned dynamic convolution that leverages a learnable prototype kernel bank to generate sample- and spatiallly-adaptive kernels, enabling effective modeling of both globally consistent and locally heterogeneous degradations across tasks and therefore reducing task-specific representation interference. In addition, we develop a progressive reweighting strategy that jointly considers update strength and convergence pace to alleviate optimization imbalance. Extensive experiments show that DSCE-Net consistently outperforms mainstream all-in-one models on multi-modal, multi-task MedIR benchmarks, demonstrating its effectiveness and superiority. Code is available at XuLinHuu/DSCE-Net .

Keywords: All-In-One · Medical image restoration · Frequency prompt · Dynamic Convolution · Optimization balancing.

1 Introduction

All-in-one medical image restoration (MedIR) [30,29,27,1,28] intends to address multiple restoration tasks within a single model, enabling the simultaneous recovery of high-quality (HQ) medical images from their degraded low-quality (LQ) counterparts. Unlike conventional MedIR approaches that target individual tasks, such as MRI super-resolution [21,16,13,8], CT denoising [2,7,26,32,31], or PET image synthesis[17,10,19,14], all-in-one MedIR offers a versatile solution for complex clinical scenarios. It improves parameter efficiency and simplifies training and deployment. However, building a robust all-in-one MedIR model remains challenging due to task interference from modality/degradation discrepancies and task imbalance driven by varying optimization difficulty across tasks.

To mitigate task interference, existing all-in-one methods typically adopt two main strategies to alleviate negative transfer. The first focuses on explicit task conditioning, where task-related priors,such as degradation codes [33,11] and learned prompt vectors [18,25,23], are injected to modulate a shared backbone, enabling it to switch processing behaviors across heterogeneous tasks. The second emphasizes parameter decoupling with dynamic routing, which reserves partially task-specific capacity via mixture-of-experts [9] and lightweight adapter or gating mechanisms [27], thereby reducing conflicts caused by fully shared representations. Despite their effectiveness, these strategies still face practical limitations. Most of them rely on relatively coarse cues to refine features, which can be insufficient to capture both global consistent and local heterogeneous discrepancies in degradations and restoration behaviors. Without an explicit mechanism to disentangle and model global statistics and local variations, shared representations can still suffer interference, leading to performance trade-offs.

Regarding task imbalance, most existing work focuses on optimization-side techniques, broadly including loss reweighting and gradient-level control. Loss-weighting methods mitigate task dominance by assigning task-specific coefficients, with representative approaches such as uncertainty-based weighting [24,30] and training-progress reweighting [22]. Although easy to implement, they can be unstable in heterogeneous settings because loss magnitudes do not reliably reflect each task's contribution to parameter updates, limiting performance gains. Gradient-based methods address imbalance more directly. GradNorm [4], for instance, reweights task losses to match gradient norms and equalize their influence on shared parameters. However, similar gradient norms do not ensure synchronized convergence, near-converged tasks may be overly down-weighted while harder tasks still lag behind, making it difficult to balance update strength and learning progress, especially under disparate loss scales and noise.

To tackle these challenges, we propose DSCE-Net, a Dual-Scale Conditioned Evolving Network with balanced optimization for all-in-one medical image restoration. It combines frequency prompts with prototype-bank–based dynamic convolution to model global and local degradations, and employs progressive optimization to jointly balance update strength and convergence speed across tasks. Our main contributions are:

(1) We propose a frequency-prompt-conditioned unified network that jointly captures global, modality/task-dependent degradation discrepancies and local texture heterogeneity, enabling task-adaptive restoration across multiple modalities and degradations within a single model.
(2) We introduce a frequency-prompt-conditioned dynamic convolution that combines a learnable prototype-kernel bank to generate sample- and channel-adaptive kernels, enabling expressive modeling of task-specific degradations.
(3) We develop a progressive reweighting strategy that simultaneously accounts update strength and convergence pace, thereby alleviating optimization imbalance and improving overall restoration performance.

(4) Extensive experiments demonstrate that our method consistently outperforms existing models on multi-task MedIR benchmarks, validating its effectiveness.

2 Method

The architecture of the proposed DSCE-Net is illustrated in Fig. 1(a). Given a low-quality input from different tasks, we first employ a Frequency Prompt Extraction (FPE) module to generate a conditioning vector that encodes modality-specific characteristics. This prompt is then fed into the subsequent Adaptive Global and Local Feature Aggregation (AGLFE) module, which integrates the prompt with prototype-bank-based dynamic convolution to jointly model degradation heterogeneity at both global and local levels, enabling a unified network for task-adaptive restoration. To tackle the optimization challenges in all-in-one MedIR, we further introduce a joint strategy that dynamically adjusts task weights based on gradient magnitudes and training progress, alleviating task imbalance and improving overall restoration performance. The FPE, AGLFE, and the optimization strategies are detailed in the following subsections.

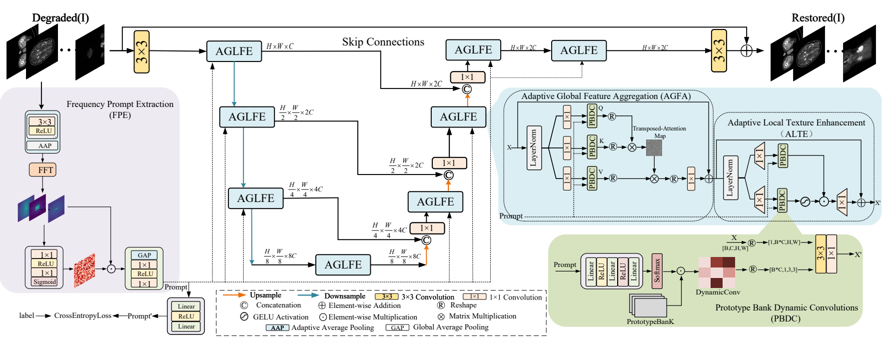

Fig. 1. Overview of the proposed all-in-one medical image restoration (DSCE-Net) network.

2.1 Frequency Prompt Extraction Module

Different modalities exhibit distinct frequency-domain statistics, therefore, we propose an FPE module to extract a modality-aware prompt from the spectrum, which conditions dynamic operators for unified multi-modality restoration. As illustrated in Fig. 1, given a low-quality input image $I^{LQ} \in \mathbb{R}^{B\times 1\times H\times W}$, FPE first employs a lightweight encoder $\psi(\cdot)$ (a few convolution layers with ReLU followed by adaptive average pooling) to extract a compact feature map. The extracted feature is then projected to a target channel dimension d_h via a 1×1

convolution and transformed into the frequency domain using a 2D FFT $\mathcal{F}(\cdot)$. Finally, we take the logarithm of the magnitude to obtain the spectral representation F_s, which explicitly captures the energy distribution across frequency bands:

$$F_s = \log\Big(1 + \big|\mathcal{F}\big(\mathrm{Conv}_{1\times1}\big(\psi(I^{LQ})\big)\big)\big|\Big) \in \mathbb{R}^{B\times d_h\times32\times32}. \tag{1}$$

To adaptively identify modality-/degradation-relevant frequency bands, FPE further introduces a gating branch $\mathcal{G}(\cdot)$ that produces a attention map using 1×1 convolutions and a Sigmoid function. The attention map is applied to the spectrum via element-wise multiplication, suppressing redundant components while emphasizing informative ones:

$$\hat{F}_s = \mathcal{G}(F_s) \odot F_s, \tag{2}$$

Finally, the re-weighted spectral feature $\hat{F}_s$ is mapped to a conditioning vector $P \in \mathbb{R}^{B\times d_p}$ through a 1×1 convolution with ReLU followed by global average pooling, where $d_p = 128$. The vector P serves as the frequency-domain prompt for subsequent conditional modulation of dynamic operators and feature aggregation.

2.2 Adaptive Global and Local Feature Extraction Module

AGLFE leverages the frequency prompt P to condition dynamic operators and jointly model cross-task degradation discrepancies at both global and local levels. As shown in Fig. 1, it is inserted into each encoder/decoder stage and consists of two cascaded blocks: adaptive global feature aggregation (AGFA) and adaptive local texture enhancement (ALTE). AGFA captures spatially consistent yet modality/task-dependent degradation cues, while ALTE further refines spatially varying textures and fine structures.

AGFA. Given an input feature $X \in \mathbb{R}^{B\times C\times H\times W}$, AGFA first applies layer normalization and then forms $Q/K/V$ using 1×1 convolutions followed by a prompt-conditioned prototype bank dynamic convolution (PBDC):

$$Q, K, V = \mathrm{PBDC}(\mathrm{Conv}_{1\times1}(\mathrm{LN}(X)); P)\,. \tag{3}$$

PBDC constructs dynamic group-wise convolution kernels by mixing a learnable bank of N prototypes. Specifically, the prompt P is mapped by a lightweight MLP $\phi(\cdot)$ to channel-wise mixture coefficients, which are used to linearly combine the prototype kernels $\mathbf{K} \in \mathbb{R}^{N\times3\times3}$ to obtain a sample- and channel-specific dynamic kernel $\mathbf{W} \in \mathbb{R}^{B\times C\times3\times3}$:

$$\mathbf{a} = \mathrm{Softmax}(\phi(P)) \in \mathbb{R}^{B\times C\times N}, \qquad W_{b,c,:,:} = \sum_{n=1}^{N} a_{b,c,n}\, K_{n,:,:}. \tag{4}$$

Here, $W_{b,c,:,:}$ denotes the dynamic 3×3 kernel for the c-th channel of the b-th sample. With $\mathbf{W}$, PBDC performs efficient per-sample dynamic filtering via grouped convolution:

$$\mathrm{PBDC}(X; P) = \mathrm{GroupConv}_{3\times3}(X, \mathbf{W}), \tag{5}$$

where $\mathbf{W}$ is computed from P by Eq. (4).

We then compute transposed attention in the channel domain. After reshaping $Q, K, V \in \mathbb{R}^{B\times C\times HW}$, the prompt-adaptive global aggregation is given by

$$\hat{X} = X + \mathrm{Conv}_{1\times 1}\Big(\mathrm{Reshape}\big(\mathrm{Softmax}(\tfrac{QK^\top}{\sqrt{HW}})\,V\big)\Big). \tag{6}$$

By conditioning $Q/K/V$ with PBDC, AGFA yields a prompt-adaptive global channel aggregation that aligns feature mixing with modality-specific spectral statistics and degradation characteristics.

ALTE. ALTE further enhances spatially varying textures and fine structures with a prompt-adaptive gated formulation. Given $\hat{X}$, we apply LayerNorm and use two 1×1 projections (expansion ratio γ) to form a feature branch and a gating branch, both conditioned by PBDC:

$$T = \mathrm{PBDC}(\mathrm{Conv}^{f}_{1\times 1}(\mathrm{LN}(\hat{X})); P) \odot \sigma\Big(\mathrm{PBDC}(\mathrm{Conv}^{g}_{1\times 1}(\mathrm{LN}(\hat{X})); P)\Big), \tag{7}$$

where $\sigma(\cdot)$ denotes Sigmoid. The output is projected back to C channels and added residually:

$$X' = \mathrm{ALTE}(\hat{X}; P) = \hat{X} + \mathrm{Conv}_{1\times 1}(T). \tag{8}$$

In this way, AGFA provides prompt-adaptive global aggregation, while ALTE selectively enhances local textures and boundaries, together enabling robust multi-modality restoration under heterogeneous degradations.

2.3 Progressive Optimization Balancing Strategy

Our training objective combines a progressive reconstruction loss with an auxiliary modality classification loss:

$$\mathcal{L}_{\mathrm{total}} = \mathcal{L}_{\mathrm{recon}} + \lambda_{\mathrm{cls}}\mathcal{L}_{\mathrm{cls}}, \tag{9}$$

where $\lambda_{\mathrm{cls}} = 0.01$. The classification loss encourages the frequency prompt to be modality-discriminative, facilitating modality-adaptive dynamic operators:

$$\mathcal{L}_{\mathrm{cls}} = \mathrm{CE}(\hat{y}_{\mathrm{mod}}, y_{\mathrm{mod}}), \quad y_{\mathrm{mod}} \in \{0, 1, 2\}, \tag{10}$$

where $\hat{y}_{\mathrm{mod}}$ is predicted from the prompt via a lightweight classification head. Let $\mathcal{L}^{(m)}_{\mathrm{recon}}(t)$ denote the progressive reconstruction loss of task $m \in \{0, \ldots, M\}$ at iteration t. To mitigate task imbalance, we dynamically reweight tasks by jointly considering their parameter update strength depicted with the gradient norm $G_m(t)$ and convergence pace reflected with the relative loss ratio $r_m(t)$. Every $T = 20$ iterations, we compute

$$G_m(t) = \left\|\nabla_\theta \mathcal{L}^{(m)}_{\mathrm{recon}}(t)\right\|_2, \qquad r_m(t) = \left(\frac{\mathcal{L}^{(m)}_{\mathrm{recon}}(t)}{\mathcal{L}^{(m)}_{\mathrm{recon}}(0)}\right)^{\alpha},\ \alpha = 0.6, \tag{11}$$

and set the target gradient $\tilde{G}_m(t) = r_m(t)\,\bar{G}(t)$ with $\bar{G}(t) = \frac{1}{M+1}\sum_{m=0}^{M} G_m(t)$. The weight is updated by

$$w_m(t) \leftarrow w_m(t-T) \cdot \frac{\tilde{G}_m(t)}{G_m(t)+\epsilon}, \qquad \epsilon = 10^{-8}. \tag{12}$$

The weighted reconstruction loss is

$$\mathcal{L}_{\text{recon}}(t) = \sum_{m=0}^{M} w_m(t)\,\mathcal{L}_{\text{recon}}^{(m)}(t). \tag{13}$$

3 Experiments

3.1 Dataset

MRI Super-Resolution. Public IXI dataset(available at IXI Dataset Website)with 578 HQ T2-weighted MRI volumes are used. LQ images are generated by 4× k-space downsampling. The data are split into 405/59/114 for training/validation/testing. For each volume, we extract the central 100 slices (256×256) for experiments. **CT Denoising.** 50 abdomen scans from the LDCT dataset[15] are selected, which provides paired standard-dose and quarter-dose CT scans. The dataset is divided into 40/5/5 cases for training/validation/testing, and each slice has a size of 512×512. **PET Synthesis.** PET data provided in the paper[27] is used in this work. LQ images are simulated by reducing the acquisition dose by a factor of 12, each image has a size of 192×400. The dataset is split into 120/10/29 for training/validation/testing, respectively.

3.2 Implementations

DSCE-Net uses a four-level encoder-decoder with AGLFE blocks $[4, 6, 6, 8]$ from level-1 to level-4, an input channel width of 48, and a PBDC prototype bank size of 16. We train with 128×128 patches (batch size=8) for 2×10^5 iterations using Adam, with the learning rate cosine-annealed from 2×10^{-4} to 1×10^{-6}. We compare DSCE-Net with six all-in-one MedIR methods: CSNet[5], ConvIR[6], DFPIR[20], Restore-RWKV[29],AMIR[27], and TAT[30], they are trained using their default settings. All experiments are implemented in PyTorch and run on an NVIDIA L20 GPU.

4 Results and Analysis

4.1 Comparison Results

As shown in Table 1, DSCE-Net achieves the best performance across all three tasks and delivers the highest PSNR/SSIM and the lowest RMSE. Compared with prior all-in-one models (e.g., AMIR and TAT), DSCE-Net shows more evident gains on CT denoising and PET synthesis, indicating stronger capability in

Table 1. Quantitative Comparison of Different All-in-One Image Restoration Methods.

Method	CT Denoising			MRI Super-resolution			PET Synthesis			Average		
	PSNR↑	SSIM↑	RMSE↓	PSNR↑	SSIM↑	RMSE↓	PSNR↑	SSIM↑	RMSE↓	PSNR↑	SSIM↑	RMSE↓
CSNet[5]	33.1532	0.9072	9.1880	31.4894	0.9312	30.7943	37.1487	0.9430	0.0863	33.9304	0.9271	13.3562
ConvIR[6]	33.3117	0.9090	9.0361	31.6561	0.9333	30.1966	37.1138	0.9435	0.0870	34.0272	0.9286	13.1066
DFPIR[20]	33.6993	0.9105	8.6439	31.9245	0.9382	29.4066	37.1755	0.9469	0.0867	34.2664	0.9319	12.7124
AMIR[27]	33.7113	0.9103	8.6139	31.9936	0.9394	29.2125	37.1735	0.9476	0.0868	34.2928	0.9324	12.6377
Restore-RWKV[29]	33.6812	0.9097	8.6450	31.9709	0.9390	29.2749	37.2372	0.9478	0.0862	34.2964	0.9322	12.6687
TAT[30]	33.8523	0.9119	8.4898	32.0566	0.9401	29.0020	37.1808	0.9476	0.0867	34.3632	0.9332	12.5262
DSCE-Net	**33.9257**	**0.9133**	**8.3994**	**32.0662**	**0.9401**	**28.9918**	**37.3411**	**0.9489**	**0.0852**	**34.4443**	**0.9341**	**12.4921**

modeling cross-modality degradation and noise discrepancies. The visual comparisons in Fig. 2 further confirm that DSCE-Net better preserves fine structures and edges while suppressing residual noise/artifacts, producing results closer to the HQ references.

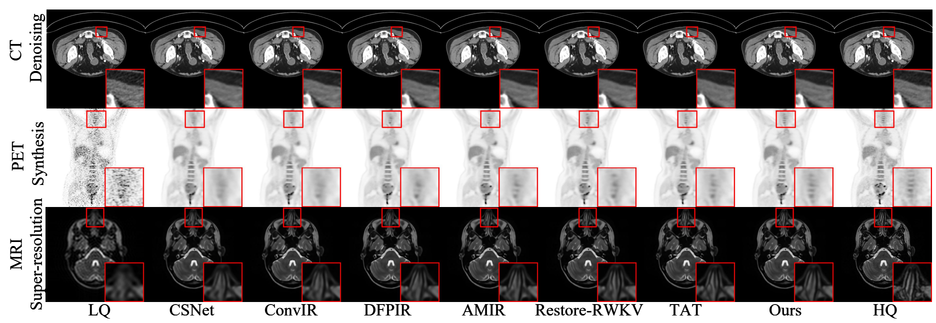

Fig. 2. Visual comparison between different methods across three MedIR tasks.

4.2 Ablation Results and Analysis

Table 2 reports the performance of the variants w/o AGFA and w/o ALTE, and further examines how different prompts, convolution operators, and optimization strategies affect MedIR performance. In the module ablation study, Baseline is constructed by replacing the PBDC used in the AGFA and ALTE modules with conventional group convolutions and removing the prompt. Re-

Table 2. Quantitative Ablation Results Averaged Over All Tasks.

Module				**Prompt**				**Dynamic Conv**				**Loss**			
Variant	PSNR↑	SSIM↑	RMSE↓	Variant	PSNR↑	SSIM↑	RMSE↓	Variant	PSNR↑	SSIM↑	RMSE↓	Variant	PSNR↑	SSIM↑	RMSE↓
Baseline	34.2856	0.9322	12.6870	w/o FPE	34.3568	0.9329	12.5577	w/o PBDC	34.3067	0.9327	12.6607	w/o Ours loss	34.4101	0.9337	12.4926
w/o AGFA	34.3849	0.9331	12.6062	AMIR's RIN[27]	34.4020	0.9334	12.5051	DYConv[3]	34.3555	0.9329	12.6173	TUR[24]	34.3958	0.9336	12.5808
w/o ALTE	33.7530	0.9261	13.9092	TAT's TREN[30]	34.3771	0.9330	12.8340	ODConv[12]	34.1568	0.9297	13.0410	TAT's loss[30]	34.3933	0.9335	12.4773
Ours	34.4443	0.9341	12.4921	Ours	34.4443	0.9341	12.4921	Ours	34.4443	0.9341	12.4921	Ours	34.4443	0.9341	12.4921

moving local degradation adaptation (w/o ALTE) causes a marked performance drop across all tasks; for MRI super-resolution in particular, fine details become severely blurred (Fig.3). Without the global degradation adaptation module (w/o AGFA), the average performance decreases slightly, accompanied by reduced contrast (PET synthesis in Fig. 3) and more visible artifacts (MRI super-resolution in Fig. 3), confirming the importance of modeling both spatially variant and globally consistent degradations. Moreover, compared with the Baseline, our approach highlights the benefit of prompt-guided dynamic convolution with PBDC. This is further supported by the w/o FPE and w/o PBDC variants: removing either module blurs the outputs and introduces spurious details (Fig. 3). Our prompt strategy also outperforms AMIR's task-routing RIN [27] and TAT's stop-gradient TREN [30], achieving the best PSNR (34.4443). The t-SNE results (Fig. 4) show clear modality-wise clusters, indicating stronger feature decoupling. Compared with DYConv [3] and ODConv [12], PBDC still delivers superior performance. For optimization, removing our proposed weighting scheme (w/o Ours loss) causes a dramatic drop. In addition, our loss outperforms TUR's task-uncertainty regularization [24] and TAT's adversarial perceptual objective [30], yielding better reconstruction fidelity and structural consistency on validation metrics.

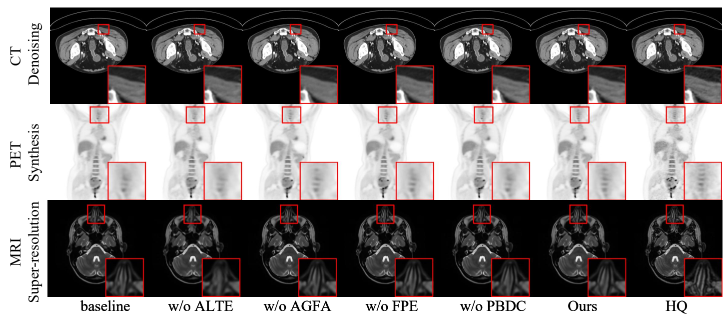

Fig. 3. Visual comparison between different methods across three MedIR tasks.

5 Conclusion

We propose DSCE-Net, a dual-scale conditioned evolving network for all-in-one medical image restoration across modalities and degradations. A frequency prompt-conditioned dynamic convolution with a learnable prototype kernel bank produces sample- and spatial-adaptive kernels, capturing global and local degradation patterns and reducing cross-modality interference. A progressive reweighting strategy that accounts for update strength and convergence pace further balances optimization and stabilizes multi-task training. Experiments on multi-modal, multi-task MedIR benchmarks show consistent improvements over main-

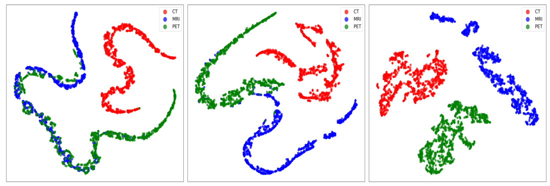

Fig. 4. t-SNE visualization of the proposed FPE vs. AMIR's RIN and TAT's TREN on different task inputs.

stream all-in-one restoration methods. However, it is worth noting that the incorporation of the learnable prototype kernel bank and dynamic convolution mechanisms inevitably increases the model complexity, resulting in a relatively large number of parameters. Future work will study finer-grained condition modeling and broader clinical scenarios.

Acknowledgments. This work is supported by National Natural Science Foundation of China (Grant No. 62571150)

Disclosure of Interests. We have no conflicts of interest to disclose.

References

1. Chen, H., Yang, Z., Hou, H., Zhang, H., Wei, B., Zhou, G., Xu, Y.: All-in-one medical image restoration with latent diffusion-enhanced vector-quantized codebook prior. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 67–77. Springer (2025)
2. Chen, H., Zhang, Y., Zhang, W., Liao, P., Li, K., Zhou, J., Wang, G.: Low-dose ct denoising with convolutional neural network. In: 2017 IEEE 14th international symposium on biomedical imaging (ISBI 2017). pp. 143–146. IEEE (2017)
3. Chen, Y., Dai, X., Liu, M., Chen, D., Yuan, L., Liu, Z.: Dynamic convolution: Attention over convolution kernels. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11030–11039 (2020)
4. Chen, Z., Badrinarayanan, V., Lee, C.Y., Rabinovich, A.: Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In: International conference on machine learning. pp. 794–803. PMLR (2018)
5. Cui, Y., Liu, M., Ren, W., Knoll, A.: Hybrid frequency modulation network for image restoration. In: IJCAI. pp. 722–730 (2024)
6. Cui, Y., Ren, W., Cao, X., Knoll, A.: Revitalizing convolutional network for image restoration. IEEE Transactions on Pattern Analysis and Machine Intelligence **46**(12), 9423–9438 (2024)
7. Diwakar, M., Kumar, M.: A review on ct image noise and its denoising. Biomedical Signal Processing and Control **42**, 73–88 (2018)
8. Greenspan, H.: Super-resolution in medical imaging. The computer journal **52**(1), 43–63 (2009)

9. Guo, H., Dai, T., Bai, Y., Chen, B., Ren, X., Zhu, Z., Xia, S.T.: Parameter efficient adaptation for image restoration with heterogeneous mixture-of-experts. Advances in Neural Information Processing Systems **37**, 13522–13547 (2024)
10. Hu, S., Shen, Y., Wang, S., Lei, B.: Brain mr to pet synthesis via bidirectional generative adversarial network. In: International conference on medical image computing and computer-assisted intervention. pp. 698–707. Springer (2020)
11. Lang, H., Zhou, Y., Yu, Y., Su, Z., Zhuge, H., Wang, W., Fang, D., Qin, J., Wei, M., Lin, R., et al.: Multi-modal low-dose medical imaging through instruction-guided unified ai. Frontiers in Medicine **13**, 1691143 (2026)
12. Li, C., Zhou, A., Yao, A.: Omni-dimensional dynamic convolution. arXiv preprint arXiv:2209.07947 (2022)
13. Li, G., Rao, C., Mo, J., Zhang, Z., Xing, W., Zhao, L.: Rethinking diffusion model for multi-contrast mri super-resolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11365–11374 (2024)
14. Luo, Y., Zhou, L., Zhan, B., Fei, Y., Zhou, J., Wang, Y., Shen, D.: Adaptive rectification based adversarial network with spectrum constraint for high-quality pet image synthesis. Medical image analysis **77**, 102335 (2022)
15. Moen, T.R., Chen, B., Holmes III, D.R., Duan, X., Yu, Z., Yu, L., Leng, S., Fletcher, J.G., McCollough, C.H.: Low-dose ct image and projection dataset. Medical physics **48**(2), 902–911 (2021)
16. Muhammad, A., Aramvith, S., Duangchaemkarn, K., Sun, M.T.: Brain mri image super-resolution reconstruction: A systematic review. IEEE Access (2024)
17. Pan, S., Abouei, E., Peng, J., Qian, J., Wynne, J.F., Wang, T., Chang, C.W., Roper, J., Nye, J.A., Mao, H., et al.: Full-dose whole-body pet synthesis from low-dose pet using high-efficiency denoising diffusion probabilistic model: Pet consistency model. Medical Physics **51**(8), 5468–5478 (2024)
18. 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), `https://openreview.net/forum?id=KAlSIL4tXU`
19. Sun, H., Jiang, Y., Yuan, J., Wang, H., Liang, D., Fan, W., Hu, Z., Zhang, N.: High-quality pet image synthesis from ultra-low-dose pet/mri using bi-task deep learning. Quantitative Imaging in Medicine and Surgery **12**(12), 5326 (2022)
20. Tian, X., Liao, X., Liu, X., Li, M., Ren, C.: Degradation-aware feature perturbation for all-in-one image restoration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 28165–28175 (2025)
21. Van Reeth, E., Tham, I.W., Tan, C.H., Poh, C.L.: Super-resolution in magnetic resonance imaging: a review. Concepts in Magnetic Resonance Part A **40**(6), 306–325 (2012)
22. Wu, G., Jiang, J., Jiang, K., Liu, X.: Harmony in diversity: Improving all-in-one image restoration via multi-task collaboration. In: Proceedings of the 32nd ACM international conference on multimedia. pp. 6015–6023 (2024)
23. Wu, G., Jiang, J., Jiang, K., Liu, X., Nie, L.: Learning dynamic prompts for all-in-one image restoration. IEEE Transactions on Image Processing (2025)
24. Wu, G., Jiang, J., Wang, Y., Jiang, K., Liu, X.: Debiased all-in-one image restoration with task uncertainty regularization. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 39, pp. 8386–8394 (2025)
25. Wu, Z., Liu, W., Wang, J., Li, J., Huang, D.: Freprompter: Frequency self-prompt for all-in-one image restoration. Pattern Recognition **161**, 111223 (2025). `https://doi.org/https://doi.org/10.1016/j.patcog.2024.111223`, `https://www.sciencedirect.com/science/article/pii/S0031320324009749`

26. Yang, L., Li, Z., Ge, R., Zhao, J., Si, H., Zhang, D.: Low-dose ct denoising via sinogram inner-structure transformer. IEEE transactions on medical imaging **42**(4), 910–921 (2022)
27. 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 (2024)
28. Yang, Z., Chen, H., Qian, Z., Zhou, Y., Zhang, H., Zhao, D., Wei, B., Xu, Y.: Region attention transformer for medical image restoration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 603–613. Springer (2024)
29. Yang, Z., Li, J., Zhang, H., Zhao, D., Wei, B., Xu, Y.: Restore-rwkv: Efficient and effective medical image restoration with rwkv. IEEE Journal of Biomedical and Health Informatics (2025)
30. Yang, Z., Zhang, J., Yi, Y., Liang, J., Wei, B., Xu, Y.: Tat: Task-adaptive transformer for all-in-one medical image restoration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 565–575. Springer (2025)
31. Yang, Z., Zhou, Y., Chen, H., Zhang, H., Zhao, D., Wei, B., Xu, Y.: Unipet: a universal network for high-quality pet image denoising across varied dose reduction factors. Medical Image Analysis p. 104059 (2026)
32. You, C., Yang, Q., Shan, H., Gjesteby, L., Li, G., Ju, S., Zhang, Z., Zhao, Z., Zhang, Y., Cong, W., et al.: Structurally-sensitive multi-scale deep neural network for low-dose ct denoising. IEEE access **6**, 41839–41855 (2018)
33. Zhang, C., Zhu, Y., Yan, Q., Sun, J., Zhang, Y.: All-in-one multi-degradation image restoration network via hierarchical degradation representation. In: Proceedings of the 31st ACM international conference on multimedia. pp. 2285–2293 (2023)