

All-In-One Medical Image Restoration via Task-Adaptive Routing

Zhiwen Yang¹, Haowei Chen¹, Ziniu Qian¹, Yang Yi¹, Hui Zhang², Dan Zhao³,
Bingzheng Wei⁴, and Yan Xu¹(✉,*)

¹ School of Biological Science and Medical Engineering, State Key Laboratory of Software Development Environment, Key Laboratory of Biomechanics and Mechanobiology of Ministry of Education, Beijing Advanced Innovation Center for Biomedical Engineering, Beihang University, Beijing 100191, China
xuyan04@gmail.com

² Department of Biomedical Engineering, Tsinghua University, Beijing 100084, China

³ Department of Gynecology Oncology, National Cancer Center/National Clinical Research Center for Cancer/Cancer Hospital, Chinese Academy of Medical Sciences and Peking Union Medical College, Beijing 100021, China

⁴ ByteDance Inc., Beijing 100098, China

Abstract. Although single-task medical image restoration (MedIR) has witnessed remarkable success, the limited generalizability of these methods poses a substantial obstacle to wider application. In this paper, we focus on the task of all-in-one medical image restoration, aiming to address multiple distinct MedIR tasks with a single universal model. Nonetheless, due to significant differences between different MedIR tasks, training a universal model often encounters task interference issues, where different tasks with shared parameters may conflict with each other in the gradient update direction. This task interference leads to deviation of the model update direction from the optimal path, thereby affecting the model’s performance. To tackle this issue, we propose a task-adaptive routing strategy, allowing conflicting tasks to select different network paths in spatial and channel dimensions, thereby mitigating task interference. Experimental results demonstrate that our proposed **All-in-one Medical Image Restoration (AMIR)** network achieves state-of-the-art performance in three MedIR tasks: MRI super-resolution, CT denoising, and PET synthesis, both in single-task and all-in-one settings. The code and data will be available at <https://github.com/Yaziwel/AMIR>.

Keywords: All-In-One · Medical Image Restoration · Task Routing.

1 Introduction

Medical image restoration (MedIR) refers to the process of transforming low-quality (LQ) medical images into high-quality (HQ) images. MedIR has achieved remarkable success in individual tasks such as MRI super-resolution [1,2], CT

* Corresponding author

denoising [3,4,5,6], and PET synthesis [7,8,9,10,11,12]. Within the scope of single-task MedIR, there are also studies addressing input degradation variations, such as varying noise levels [8], imaging protocols [12], and centers [12]. The well-defined settings of these MedIR tasks allow researchers to design specific models tailored to the unique characteristics of each individual task.

However, these single-task models designed for particular MedIR tasks often encounter significant performance drops when applied to other MedIR tasks. In complex application scenarios such as multimodal imaging (e.g., PET/CT and PET/MRI), multiple MedIR tasks coexist simultaneously. Single-task models struggle to address the differences between modalities and tasks, and utilizing separate models for each task may result in inefficiencies in both usage and maintenance. Apart from practical limitations, single-task models remain inherently constrained by their task-specific nature, hindering the evolution from specificity to a more generalized intelligence in the field of medical image restoration. Therefore, there is a pressing need to develop a single universal model capable of simultaneously handling multiple MedIR tasks.

Recently, all-in-one image restoration [13,14,15,16] (also known as multi-task image restoration) has gained prominence in natural images, attempting to address multiple different restoration tasks using a single universal model. It holds the potential to address multiple MedIR tasks with a single universal model. However, given the substantial disparities between medical and natural image restoration tasks, it is not advisable to directly apply all-in-one natural image restoration methods to MedIR tasks. In natural image restoration, task distinctions primarily arise from varied input image degradations, with their ground truths presumed to follow a uniform data distribution. However, in MedIR, alongside input degradation disparities, the ground truths of different MedIR tasks also showcase significantly varied data distributions due to modality differences (as shown in Fig. I (a) in the **supplement**). These significant differences between MedIR tasks can result in task interference, a common issue in multi-task learning [17,18], where the gradient update directions between tasks are inconsistent or even opposite. As shown in Fig. I (b) in the **supplement**, we quantify the task interference metric [18] among different MedIR tasks and observe significant task interference. This task interference issue will lead to an uncertain update direction that deviates from the optimal, resulting in suboptimal performance. While essential for handling multiple MedIR tasks, the task interference issue is rarely explicitly addressed by current all-in-one methods.

In this paper, we introduce an innovative All-in-one Medical Image Restoration (AMIR) network capable of handling multiple MedIR tasks with a single universal model. The key idea behind AMIR is the incorporation of a task-adaptive routing strategy, which dynamically directs inputs from conflicting tasks to different network paths, explicitly mitigating interference between tasks. Specifically, the proposed task-adaptive routing involves routing instruction learning, spatial routing, and channel routing. Routing instruction learning aims to adaptively learn task-relevant instructions based on input images, while spatial routing and channel routing utilize these learned instructions to guide the routing




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

of network features at spatial and channel levels, respectively, thereby alleviating potential interference. Extensive experiments demonstrate that our proposed AMIR achieves state-of-the-art performance in three MedIR tasks: MRI super-resolution, CT denoising, and PET synthesis, both in single-task and all-in-one settings. Our contribution can be three-fold:

- We propose a novel All-in-one Medical Image Restoration (AMIR) network, which allows handling multiple different MedIR tasks with a single unified model. To the best of our knowledge, AMIR could be one of the first methods to handle multiple MedIR tasks in an all-in-one fashion.
- We propose a novel task-adaptive routing strategy to mitigate interference between different tasks. It is achieved by assigning conflicting tasks to different network paths.
- Extensive experiments show that our proposed AMIR achieves state-of-the-art performance in both single-task MedIR and all-in-one MedIR tasks.

2 Method

2.1 Network Architecture

The architecture of the proposed All-In-One Medical Image Restoration (AMIR) network is shown in Fig. 1 (a). AMIR adopts Restormer [2], a Unet-style network, as the baseline model for medical image restoration. The key difference lies in AMIR’s inclusion of a Routing Instruction Network (RIN) alongside Spatial Routing Modules (SRMs; as shown in Fig. 1 (b)) and Channel Routing Modules

(CRMs; as shown in Fig. 1 (c)) for network path routing. Specifically, given an input LQ medical image $I^{LQ} \in \mathbb{R}^{H \times W \times 1}$, AMIR first extracts shallow features $I^{SF} \in \mathbb{R}^{H \times W \times C}$ by applying a 3×3 convolution, where H , W , and C denote the height, width, and channel, respectively. Subsequently, I^{SF} undergoes a hierarchical encoder-bottleneck-decoder structure to be transformed into deep features I^{DF} , with multiple Restormer Transformer blocks [2] utilized at each level for feature extraction. To mitigate task interference, AMIR employs a task-adaptive routing strategy, incorporating SRMs before each encoder level and CRMs before the bottleneck and each decoder level. SRMs and CRMs adaptively select propagation paths for different task features based on the task-relevant instructions from RIN, thus reducing potential interference. Finally, a 3×3 convolution layer is applied to deep features I^{DF} to generate residual image $I^R \in \mathbb{R}^{H \times W \times 1}$, which is added to the input LQ image to obtain the restored image $\hat{I}^{HQ} = I^{LQ} + I^R$. Next, we will describe the task-adaptive routing strategy and its key components.

2.2 Task-Adaptive Routing

To address potential interference caused by multiple tasks sharing the same parameters but with different optimization directions, we propose a task-adaptive routing strategy. This strategy enables different tasks to select distinct paths for customized processing, thereby avoiding interference between tasks. We will introduce the task-adaptive routing strategy from three aspects: routing instruction learning, spatial routing, and channel routing.

Routing Instruction Learning. Instructions are representations relevant to the task, crucial for helping the network better understand the current task and adjust the direction of restoration. Previous all-in-one natural image restoration methods often utilize representations from contrastive learning [19, 13] as instructions. However, the task of contrastive learning exhibits significant differences from image restoration tasks, making it challenging to balance their relationship and potentially introducing undesirable task interference. To address this, we propose a Routing Instruction Network (RIN; as shown in Fig. 1 (a)), which can adaptively generate task-relevant instructions from the input image without the need for additional supervision. The mechanism of the RIN can be formulated as follows:

$$I^{IR} = \sum_{i=1}^N \alpha_i D_i, \quad \alpha_i = \text{Softmax}(\text{GAP}(\text{E}(I))), \quad (1)$$

where $\text{E}(\cdot)$ is a five-layer CNN encoder, while $\text{GAP}(\cdot)$ denotes the global average pooling layer. $D = [D_1, D_2, \dots, D_N]$ constitutes an instruction dictionary, where the N instructions $D_i \in \mathbb{R}^{256}$ in the dictionary are learnable parameters. During this process, RIN dynamically predicts weights α_i from the input image I and applies them to the instruction dictionary D to generate input-conditioned instruction I^{IR} . This process does not require any supervision, yet we demonstrate in Fig. III (a) in the **supplement** that the learned instructions I^{IR} are task-relevant.

Spatial Routing. Task interference arises when different tasks share network parameters but have different update directions. A fundamental solution is to assign conflicting tasks to separate parameters. Mixture-of-Experts (MoE) [20] provides a potential solution, which learns to dynamically route inputs to different expert networks. However, vanilla MoE selects the experts solely relying on input token representation, neglecting crucial global and task-relevant information. Hence, we propose to enhance the routing strategy of the vanilla MoE by incorporating the learned global task-relevant instruction I^{IR} . Given a local spatial token $x_i \in \mathbb{R}^{C'}$ of the input feature $X \in \mathbb{R}^{H' \times W' \times C'}$, along with the global instruction I^{IR} and an expert bank $E = [E_1, E_2, \dots, E_M]$ comprising M expert networks, our spatial routing module (SRM; as shown in Fig. 1 (b)) can be formulated as:

$$G(x_i, I^{IR}) = \text{Top-K}(\text{Softmax}(\text{FC}([x_i, \text{FC}(I^{IR})])), \quad (2)$$

$$x'_i = \sum_{e=1}^M G(x_i, I^{IR})_e E_e(x_i), \quad (3)$$

where $\text{FC}(\cdot)$ indicates the fully connected layer. $\text{Top-K}(\cdot)$ operator sets all values to be zero except the largest K values. $G(\cdot)$ denotes the routing function that produces a sparse weight for different experts. E_e refers to the e -th expert in the expert bank, with each being a multi-layer perceptron (MLP). $G(x_i, I^{IR})_e$ determines how much the e -th expert contributes to the output. In this process, SRM routes each spatial token from the input feature X to the corresponding top- K selected experts for separate processing with the guidance of the global instruction I^{IR} , and then combines each expert's knowledge through weighted summation. Notably, very dissimilar tokens will be routed to distinct experts, thereby avoiding potential interference. We integrate SRM before each encoder level in the UNet to mitigate task interference during the encoding process.

Channel Routing. Although SRM can effectively mitigate task interference and obtain strong interpretability, it suffers from two drawbacks. Firstly, routing different spatial tokens to different experts disrupts the feature spatial continuity. Secondly, SRM introduces multiple expert networks, leading to a linear increase in parameters with the growth of expert networks. To address these issues, we propose a more efficient Channel Routing Module (CRM; as shown in Fig. 1 (c)). With the guidance of task instruction I^{IR} , CRM dynamically routes data from different tasks to different channels without introducing excessive additional parameters, while also preserving feature spatial continuity. Given an input feature $X \in \mathbb{R}^{H' \times W' \times C'}$, the mechanism of CRM can be described as:

$$X' = (X \odot m), \quad m = \text{Sigmoid}(\text{FC}(I^{IR})), \quad (4)$$

where $m \in \mathbb{R}^{C'}$ is a channel-wise (soft-)binary mask estimated from the instruction I^{IR} . In this process, CRM conducts channel-wise task routing through feature masking depending on the task instruction. We incorporate the CRM before the bottleneck and each decoder level of the UNet to address task interference in the decoding process.

2.3 Loss Function

The overall loss can be summarized as follows:

$$L = L_1 + \gamma L_{Balance}, \quad (5)$$

where L_1 depicts the difference between the restored $\hat{I}^{HQ}$ and the HQ image I^{HQ} . $L_{Balance}$ [20] is a regularization term in MoE that prevents static routing where the same few experts are always selected. γ is a weighting parameter.

3 Experiments and Results

3.1 Dataset

Our all-in-one medical image restoration experiment encompasses three tasks: MRI super-resolution, CT denoising, and PET Synthesis. In the subsequent section, we introduce the corresponding datasets for each task.

MRI Super-Resolution. We use the publicly available IXI MRI dataset ¹, which comprises 578 HQ T2 weighted MRI images. Each 3D MRI volume is sized at $256 \times 256 \times n$, from which we extract the central 100 2D slices sized 256×256 to exclude side slices. The LQ image is obtained by cropping the k -space with a downsampling factor of $4 \times$ (retaining the central 6.25% data points). The dataset is divided into 405 for training, 59 for validation, and 114 for testing.

CT Denoising. We employ the dataset from the 2016 NIH AAPM-Mayo Clinic Low-Dose CT Grand Challenge [21], which comprises paired standard-dose HQ CT images and quarter-dose LQ CT images, each with an image size of 512×512 . These images originate from 10 patients, with 8 allocated for training, 1 for validation, and 1 for testing purposes.

PET Synthesis. We acquire 159 HQ PET images using the PolarStar m660 PET/CT system in list mode, with an injection dose of 293MBq ^{18}F -FDG. LQ PET images are generated through random list mode decimation with a dose reduction factor of 12. Both HQ and LQ PET images are reconstructed using the standard OSEM method [22]. Each PET image has 3D shapes of $192 \times 192 \times 400$, with a voxel size of $3.15\text{mm} \times 3.15\text{mm} \times 1.87\text{mm}$, and is divided into 192 2D slices sized 192×400 . Slices containing only air are excluded. Patient data are partitioned into 120 for training, 10 for validation, and 29 for testing.

3.2 Implementation

In our AMIR architecture, as shown in Fig. 1, the number of Transformer blocks is configured as follows: $L_1 = 5$, $L_2 = L_3 = 7$, $L_4 = 9$, and $L_{refinement} = 4$. The input channel number is specified as $C = 42$, and the length of the instruction dictionary is set to $N = 16$. Within the SRM component, we designate the expert number as $M = 4$ and the selected number of experts as $K = 2$. During

¹ <https://brain-development.org/ixi-dataset/>

Table 1. All-in-one medical image restoration results. The best results are **bolded**, and the second-best results are underlined.

Method	MRI Super-Resolution			CT Denoising			PET Synthesis			Average		
	PSNR $\uparrow$	SSIM $\uparrow$	RMSE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	RMSE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	RMSE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	RMSE $\downarrow$
Restormer [2]	31.7177	0.9362	30.0549	33.6142	0.9177	8.5329	37.1368	0.9473	0.0872	34.1562	0.9337	12.8917
Eformer [5]	29.1922	0.8728	39.0983	32.4438	0.9078	9.7565	35.1096	0.9091	0.1085	32.2485	0.8966	16.3211
Spach Transformer [10]	31.1799	0.9290	31.8342	33.4740	0.9155	8.6677	37.0547	0.9445	0.0874	33.9029	0.9297	13.5298
DRMC [12]	29.5466	0.9032	38.1691	33.2770	0.9153	8.8674	36.1909	0.9376	0.0960	33.0048	0.9187	15.7108
AirNet [13]	31.3921	0.9316	31.1141	33.6222	0.9176	8.5226	37.1721	0.9451	0.0864	34.0621	0.9314	13.2410
AMIR	32.0262	0.9396	29.0988	33.7011	0.9182	8.4520	37.1193	0.9475	0.0876	34.2822	0.9351	12.5461

training, we utilize patches sized at 128×128 with a batch size of 8. Our model undergoes training via the Adam optimizer for 2×10^5 iterations, starting with an initial learning rate of 2×10^{-4} , gradually decreasing to 1×10^{-6} using cosine annealing. All experiments are conducted in PyTorch, utilizing NVIDIA A100 with 40GB memory.

3.3 Comparative Experiment

To validate our proposed AMIR, we conduct evaluations on three tasks: MRI super-resolution, CT denoising, and PET synthesis. Baselines for MRI super-resolution include SRCNN [23], VDSR [24], SwinIR [25], and Restormer [2]; for CT denoising, CNN [3], REDCNN [4], Eformer [5], and CTformer [6]; and for PET synthesis, Xiang’s method [7], DCNN [8], ARGAN [9], and Spach Transformer [10]. We conduct experiments under two settings: single-task and all-in-one. In the single-task setting, we train different models for each MedIR task and compare them with their respective baselines. In the all-in-one setting, we train a universal model to address all tasks simultaneously. We retrain the best-performing baseline models from each task into the all-in-one setting for comparison. Additionally, we utilize two universal models for comparison: DRMC [12], initially developed for multi-center PET synthesis, and AirNet [13], originally designed for all-in-one natural image restoration. PSNR, SSIM, and RMSE scores are calculated to assess the restoration performance.

Single-Task Medical Image Restoration. The results of single-task MedIR are showcased in Table. I in the **supplement**, revealing that our proposed AMIR surpasses the best-performing baseline models — Restormer [2], Eformer [5], and Spach Transformer [10] — in MRI super-resolution, CT denoising, and PET synthesis tasks, respectively. Despite not being specifically designed for single-task MedIR, AMIR’s outstanding performance suggests that the proposed routing strategy is also helpful for handling the sample differences within a single task.

All-In-One Medical Image Restoration. The results of all-in-one MedIR are presented in Table. 1, where AMIR demonstrates state-of-the-art performance when averaged across the three tasks. Specifically, in MRI super-resolution and CT denoising, AMIR outperforms all comparative methods in terms of PSNR, SSIM, and RMSE metrics. In PET synthesis, although AMIR’s PSNR and RMSE metrics are lower than AirNet [13], it achieves the best result in SSIM. Visual comparisons in Fig. II in the **supplement** demonstrate that AMIR better restores image structure and details across all three tasks. The superiority of

Table 2. Ablation study on the combinations of training tasks. "✓" denotes AMIR training with the task and "—" indicates unavailable results.

Training Task				Testing Task								
MRI Super-Resolution	CT Denoising	PET	Synthesis	MRI Super-Resolution			CT Denoising			PET Synthesis		
				PSNR↑	SSIM↑	RMSE↓	PSNR↑	SSIM↑	RMSE↓	PSNR↑	SSIM↑	RMSE↓
✓				31.9923	0.9393	29.2095	—	—	—	—	—	—
	✓			—	—	—	33.6738	0.9183	8.4773	—	—	—
		✓		—	—	—	—	—	—	37.2121	0.9473	0.0863
✓			✓	32.0683	0.9404	29.0296	33.6934	0.9183	8.4592	—	—	—
✓			✓	<u>32.0404</u>	<u>0.9399</u>	<u>29.0864</u>	—	—	—	37.1054	0.9473	0.0875
	✓		✓	—	—	—	33.7537	0.9189	8.4027	37.0738	0.9473	0.0880
✓			✓	32.0262	0.9396	29.0988	33.7011	0.9182	8.4520	37.1193	0.9475	0.0876

Table 3. Ablation study on the task-adaptive routing strategy.

Method		Params	MRI Super-Resolution			CT Denoising			PET Synthesis		
			PSNR↑	SSIM↑	RMSE↓	PSNR↑	SSIM↑	RMSE↓	PSNR↑	SSIM↑	RMSE↓
Instruction	Restormer [2] (Baseline)	26.12 M	31.7177	0.9362	30.0549	33.6142	0.9177	8.5329	37.1368	0.9473	0.0872
	w/o D	23.53 M	31.9088	0.9382	29.4575	33.6800	0.9181	8.4712	37.1233	0.9474	0.0876
	w/o D and w/ Contrastive Learning [19]	23.53 M	31.9545	0.9388	29.3359	33.5987	0.9161	8.5495	37.1388	0.9465	0.0872
Routing Module		22.74 M	31.8678	0.9379	29.5851	33.6556	0.9171	8.4942	37.1639	0.9470	0.0871
	w/o SRM	23.37 M	31.9802	0.9391	29.2251	33.6814	0.9181	8.4712	37.1443	0.9476	0.0874
	w/o CRM	23.54 M	32.0262	0.9396	29.0988	33.7011	0.9182	8.4520	37.1193	0.9475	0.0876
AMIR											

AMIR over other methods lies in its ability to mitigate task interference more effectively, thereby preserving the specificity of each task.

3.4 Ablation Study

We conduct ablation studies on different training task combinations to analyze their impact on AMIR outcomes. Also, we study the task-adaptive routing strategy through ablation studies to understand the specific roles of its components.

Ablation Study on Task Combinations. We train AMIR on various task combinations and list the results in Table. 2. It is surprising that, despite the increase in the number of tasks, AMIR maintains its performance without significant degradation. This can be attributed to the proposed task-adaptive routing strategy, which enables multiple tasks to share a universal model with minimal interference. Furthermore, Table. 2 reveals that certain task combinations yield better results than single-task models. This is because universal models benefit from more training data and task synergy compared to single-task models.

Ablation Study on the Task-Adaptive Routing Strategy. We conduct ablations on instruction learning and routing modules to analyze the components of the task-adaptive routing strategy. For instruction learning, we remove the dictionary D and adopt contrastive learning [19], similar to AirNet [13], for guidance. Table. 3 demonstrates the superior effectiveness of adaptive instruction learning. The instruction representation I^{IR} is visualized in Fig. III (a) in the **supplement** using t-SNE, revealing discriminations between tasks, thus demonstrating its relevance to tasks. Regarding the routing modules SRM and CRM, we assess their effectiveness by removing them individually. Table. 2 indicates that both of them effectively enhance the model’s performance. Additionally, it is shown in Fig. III (b) in the **supplement** that different tasks select distinct expert networks within SRM, confirming its interpretability. Thanks to

the well-designed task-adaptive routing strategy, our proposed AMIR achieves better results than the baseline model Restormer [2] even with fewer parameters (although introducing additional parameters, AMIR utilizes fewer channels).

4 Conclusion

In this paper, we propose an all-in-one medical image restoration (AMIR) network capable of handling multiple MedIR tasks with a single universal model. To mitigate task interference, we introduce a task-adaptive routing strategy that dynamically routes different tasks to distinct network paths. Experiments demonstrate that the proposed AMIR achieves state-of-the-art performance in both single-task MedIR and all-in-one MedIR tasks. In the future, we will explore the effectiveness of the proposed AMIR as more MedIR tasks are involved.

Acknowledgments. This work is supported by the National Natural Science Foundation in China under Grant 62371016, U23B2063, 62022010, and 62176267, the Beijing Natural Science Foundation Haidian District Joint Fund in China under Grant L222032, the Beijing hope run special fund of cancer foundation of China under Grant LC2018L02, the Fundamental Research Funds for the Central University of China from the State Key Laboratory of Software Development Environment in Beihang University in China, the 111 Project in China under Grant B13003, the SinoUnion Healthcare Inc. under the eHealth program, the high performance computing (HPC) resources at Beihang University.

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

References

1. Chen, Y., Shi, F., Christodoulou, A.G., Xie, Y., Zhou, Z., Li, D.: Efficient and accurate mri super-resolution using a generative adversarial network and 3d multi-level densely connected network. In: International conference on medical image computing and computer-assisted intervention. pp. 91–99. Springer (2018)
2. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: Restormer: Efficient transformer for high-resolution image restoration. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5728–5739 (2022)
3. 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)
4. Chen, H., Zhang, Y., Kalra, M.K., Lin, F., Chen, Y., Liao, P., Zhou, J., Wang, G.: Low-dose ct with a residual encoder-decoder convolutional neural network. IEEE transactions on medical imaging **36**(12), 2524–2535 (2017)
5. Luthra, A., Sulakhe, H., Mittal, T., Iyer, A., Yadav, S.: Eformer: Edge enhancement based transformer for medical image denoising. arXiv preprint arXiv:2109.08044 (2021)
6. Wang, D., Fan, F., Wu, Z., Liu, R., Wang, F., Yu, H.: Ctformer: convolution-free token2token dilated vision transformer for low-dose ct denoising. Physics in Medicine & Biology **68**(6), 065012 (2023)

7. Xiang, L., Qiao, Y., Nie, D., An, L., Lin, W., Wang, Q., Shen, D.: Deep auto-context convolutional neural networks for standard-dose pet image estimation from low-dose pet/mri. *Neurocomputing* **267**, 406–416 (2017)
8. Chan, C., Zhou, J., Yang, L., Qi, W., Kolthammer, J., Asma, E.: Noise adaptive deep convolutional neural network for whole-body pet denoising. In: 2018 IEEE Nuclear Science Symposium and Medical Imaging Conference Proceedings (NSS/MIC). pp. 1–4. IEEE (2018)
9. 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)
10. Jang, S.I., Pan, T., Li, Y., Heidari, P., Chen, J., Li, Q., Gong, K.: Spach transformer: spatial and channel-wise transformer based on local and global self-attentions for pet image denoising. *IEEE transactions on medical imaging* (2023)
11. Zhou, Y., Yang, Z., Zhang, H., Eric, I., Chang, C., Fan, Y., Xu, Y.: 3d segmentation guided style-based generative adversarial networks for pet synthesis. *IEEE Transactions on Medical Imaging* **41**(8), 2092–2104 (2022)
12. Yang, Z., Zhou, Y., Zhang, H., Wei, B., Fan, Y., Xu, Y.: Drmc: A generalist model with dynamic routing for multi-center pet image synthesis. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 36–46. Springer (2023)
13. Li, B., Liu, X., Hu, P., Wu, Z., Lv, J., Peng, X.: All-in-one image restoration for unknown corruption. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17452–17462 (2022)
14. Potlapalli, V., Zamir, S.W., Khan, S., Khan, F.S.: Promptir: Prompting for all-in-one blind image restoration. *arXiv preprint arXiv:2306.13090* (2023)
15. Park, D., Lee, B.H., Chun, S.Y.: All-in-one image restoration for unknown degradations using adaptive discriminative filters for specific degradations. In: 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5815–5824. IEEE (2023)
16. Kong, X., Dong, C., Zhang, L.: Towards effective multiple-in-one image restoration: A sequential and prompt learning strategy. *arXiv preprint arXiv:2401.03379* (2024)
17. Yu, T., Kumar, S., Gupta, A., Levine, S., Hausman, K., Finn, C.: Gradient surgery for multi-task learning. *Advances in Neural Information Processing Systems* **33**, 5824–5836 (2020)
18. Zhu, J., Zhu, X., Wang, W., Wang, X., Li, H., Wang, X., Dai, J.: Uni-perceiver-moe: Learning sparse generalist models with conditional moes. *Advances in Neural Information Processing Systems* **35**, 2664–2678 (2022)
19. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9729–9738 (2020)
20. Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q., Hinton, G., Dean, J.: Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538* (2017)
21. McCollough, C.H., Bartley, A.C., Carter, R.E., Chen, B., Drees, T.A., Edwards, P., Holmes III, D.R., Huang, A.E., Khan, F., Leng, S., et al.: Low-dose ct for the detection and classification of metastatic liver lesions: results of the 2016 low dose ct grand challenge. *Medical physics* **44**(10), e339–e352 (2017)
22. Hudson, H.M., Larkin, R.S.: Accelerated image reconstruction using ordered subsets of projection data. *IEEE transactions on medical imaging* **13**(4), 601–609 (1994)

23. Dong, C., Loy, C.C., He, K., Tang, X.: Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence* **38**(2), 295–307 (2015)
24. Kim, J., Lee, J.K., Lee, K.M.: Accurate image super-resolution using very deep convolutional networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1646–1654 (2016)
25. Liang, J., Cao, J., Sun, G., Zhang, K., Van Gool, L., Timofte, R.: Swinir: Image restoration using swin transformer. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 1833–1844 (2021)