

Degradation-Aware Feature Perturbation for All-in-One Image Restoration

Supplementary Material

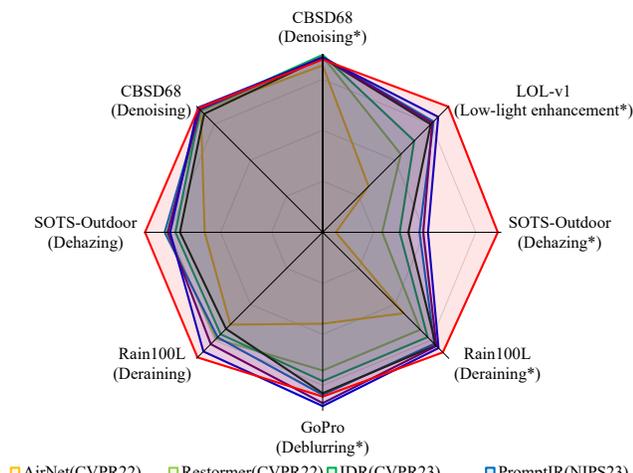


Figure 1. PSNR comparison with state-of-the-art all-in-one and task-specific methods. * denotes results obtained under All-in-One (five tasks) training setting, while unmarked results are from All-in-One (three tasks) training setting.

As shown in Fig. 1, our proposed DFPIR surpasses state-of-the-art general restoration methods in the All-in-One training setup, demonstrating superior performance. The three-task results for IDR [23] and the five-task results for PromptIR [13] are referenced from Perceive-IR [26].

1. Datasets

In our experiments, we closely follow prior works [9, 13] and evaluate our method using eight benchmark datasets: BSD400 [1], CBSD68 [11], WED [10], and Urban100 [6] for image denoising; Rain100L [20] for deraining; RESIDE [8] for dehazing; GoPro [12] for motion deblurring; and LOL [18] for low-light enhancement. Specifically, BSD400 consists of 400 clean natural images, while CBSD68 contains 68 images. WED includes 4,744 natural images, and Urban100 comprises 100 clean images. For image denoising, we combine BSD400 and WED for training, and CBSD68 and Urban100 for testing. Following [9, 13], noisy images are generated by adding white Gaussian noise to the clean images at three corruption levels: $\sigma = 15, 25, 50$. For image deraining, we use the Rain100L dataset, which contains 200 rainy-clean image pairs for training and 100 pairs for testing. In the case of image dehazing, we use the RESIDE dataset, which includes the Outdoor Training Set (OTS) with 72,135 hazy-clean pairs for training, and the Synthetic Objective Testing Set (SOTS) with 500 hazy-clean pairs for testing. We utilize the GoPro dataset for mo-

tion deblurring, which includes 2,103 images for training and 1,111 images for testing. Lastly, for low-light enhancement, we adopt the LOL-v1 dataset [18], using its official split of 485 training images and 15 testing images.

2. Degradation Type Prompt Generation

The role of the degradation type prompts is to guide the perturbation of the feature space. What we focus on is the degradation type, not the specific image details. For example, we aim to differentiate between rain degradation and hazy degradation, rather than the extent or location of rain or haze in the image. To simplify network training, we use the pre-trained language model CLIP [15] to encode text descriptions of the degradation types, thereby obtaining the degradation type prompts. We describe Gaussian noise with a standard deviation of 15 as “Gaussian noise with a standard deviation of 15”. For rain degradation, we use “Rain degradation with rain lines”. Hazy degradation is described as “Hazy degradation with normal haze”. Motion blur degradation is represented as “Blur degradation with motion blur”. Finally, low-light degradation is referred to simply as “Lowlight degradation”. This detailed description is sufficient to distinguish the degradation types of images in multitasking. The CLIP model can embed them into a common representation space: $\mathbf{P}_e = f_{CLIP}(\mathbf{d}_n)$, where \mathbf{d}_n represents degradation prompts. It is important to note that the degradation types mentioned above do not necessarily require detailed descriptions. Instead, they can be represented using simple degradation prompts, such as “Noise” for degradation levels with σ values of 15, 25, and 50, “Rain” for rainy conditions, and “Haze” for hazy conditions. We utilize detailed descriptions to fully leverage the flexibility of CLIP [15], enabling the model to handle complex or future mixed degradation scenarios for further research. In the supplementary ablation studies, we provide experimental results using these simplified degradation type descriptions (see Tab. 6). Experimental results indicate that there is only a minimal difference between using simple degradation descriptions and detailed ones.

3. Comparisons on Single-Task

Our primary goal is to design a powerful multi-task image restoration model, as single-task scenarios with specific degradation types fail to adequately demonstrate the effectiveness of our proposed feature perturbation method in mitigating the effects of multiple degradations. Nevertheless, we compare DFPIR with single-task methods. Since InstructIR [3] does not provide single-task results, we pri-

Table 1. Dehazing results in the single-task setting on the SOTS benchmark dataset [8]. Compared to PromptIR [13], our method generates a 0.69 dB PSNR improvement. PSNR (dB, \uparrow) and SSIM (\uparrow) metrics are reported on the full RGB images.

Method	DehazeNet [2]	MSCNN [16]	EPDN [14]	FDGAN [4]	AirNet [9]	Restormer [22]	PromptIR [13]	DFPIR(Ours)
PSNR	22.46	22.06	22.57	23.15	23.18	30.87	31.31	32.00
SSIM	0.851	0.908	0.863	0.921	0.900	0.969	0.973	0.981

Table 2. Deraining results in the single-task setting on Rain100L [20]. Our DFPIR obtains a significant performance boost of 2.04 dB PSNR over PromptIR [13]. PSNR (dB, \uparrow) and SSIM (\uparrow) metrics are reported on the full RGB images.

Method	UMR [21]	SIRR [19]	MSPFN [7]	LPNet [5]	AirNet [9]	Restormer [22]	PromptIR [13]	DFPIR(Ours)
PSNR	32.39	32.37	33.50	33.61	34.90	36.74	37.04	39.08
SSIM	0.921	0.926	0.948	0.958	0.977	0.978	0.979	0.984

marily focus on comparisons with PromptIR [13]. This is to show that under the single-task setting, DFPIR remains effective for images with varying degrees of degradation within the same degradation type. We compare DFPIR against various general image restoration methods (DehazeNet [2], MSCNN [16], EPDN [14], FDGAN [4], AirNet [9], Restormer [22] and PromptIR [13]) for dehazing. Tab. 1 reports dehazing results. Compared to the previous approaches PromptIR [13] and AirNet [9], our method obtains PSNR gains of 0.69 dB and 1.13 dB, respectively. Similarly, on the deraining task, our DFPIR surpasses the state-of-the-art [13] by 2.04 dB, as indicated in Tab. 2. The main methods compared are UMR [21], SIRR [19], MSPFN [7], LPNet [5], AirNet [9], Restormer [22] and PromptIR [13] for deraining. A comparable performance trend is evident in the image quality scores presented in Tab. 3 for denoising. The denoising methods compared include IRCNN [24], FFDNet [25], BRDNet [17], AirNet [9], PromptIR [13], etc. It should be noted that the metrics on the CBSD68 dataset [11] are on par with PromptIR [13], but on the Urban100 dataset [6], the metrics exceed PromptIR [13] by 0.14 dB for noise level $\sigma = 50$.

4. Ablation Studies

We conduct several ablation experiments to demonstrate the effectiveness of our proposed degradation-guided perturbation block. We report the results of training an all-in-one model on combined datasets from three restoration tasks.

Impact of key components. As illustrated in the Tab. 4, using channel attention (Method (a)) directly improves by 0.36 dB compared to the baseline [22], but it is 0.15 dB lower than channel shuffle (Method (c)). This also validates the effectiveness of the channel shuffle strategy we proposed. Channel shuffle preserves inherent image features with degradation info but offers limited reduction in cross-degradation interference. By applying attention-wise perturbation, restoration quality is significantly enhanced (DGCPM+CAAPM). The average PSNR increases

from 32.49 to 32.88, reflecting an improvement of 0.39 dB. However, method (b) (CA+CAAPM) results in a lower performance than DFPIR, indicating that the perturbations in channel and attention dimensions produce a synergistic enhancement effect. Furthermore, as observed, the method using channel attention (CA) (Method (b)) shows limited ability to mitigate cross-degradation interference, resulting in an imbalance in multi-task performance. For instance, while it achieves competitive dehazing results (close to DFPIR), its denoising and deraining performance is significantly lower than that of DFPIR. This further demonstrates that our proposed channel shuffle strategy is more effective in reducing the impact of multiple degradations. We also replaced CAAPM with spatial attention (method (d)), and the results indicate a significant performance gap in restoration quality compared to DFPIR. Additionally, it struggled to achieve a balanced performance across multiple tasks (e.g., achieving higher dehazing metrics but significantly lower results in other tasks). This is due to its limitations in effectively fusing original features and mitigating the impact of various degradations. We further validate the experimental results of retaining only CAAPM (w/o DGCPM, Method(e)). The results indicate that the absence of degradation-aware channel perturbation leads to increased interference among multiple tasks, resulting in significantly lower performance compared to DFPIR. This further demonstrates the effectiveness of channel-wise perturbation in enhancing multi-task restoration performance.

Impact of parameter γ . We conducted ablation experiments on the perturbation factor γ as well, as illustrated in the Tab. 5. If the perturbation in the attention dimension is too high ($\gamma = 0.5$) or absent ($\gamma = 1.0$), the performance is not optimal. This is because excessive perturbation, while reducing the interference between images with different degradations, increases information loss, leading to suboptimal performance. Similarly, if the perturbation is too small, the interference between tasks becomes more significant, resulting in suboptimal performance as well. After removing degradation prompts (w/o DGCPM), the absence

Table 3. Denoising comparisons in the single-task setting on CBSD68 [11] and Urban100 datasets [6]. On Urban100 [6] for the noise level 50, DFPIR yields a 0.14 dB gain over PromptIR [13]. PSNR (dB, \uparrow) and SSIM (\uparrow) metrics are reported on the full RGB images.

Method	Denoising on CBSD68			Denoising on Urban100		
	$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	$\sigma = 15$	$\sigma = 25$	$\sigma = 50$
IRCNN [24]	33.87 / 0.929	31.18 / 0.882	27.88 / 0.790	27.59 / 0.833	31.20 / 0.909	27.70 / 0.840
FFDNet [25]	33.87 / 0.929	31.21 / 0.882	27.96 / 0.789	33.83 / 0.942	31.40 / 0.912	28.05 / 0.848
BRDNet [17]	34.10 / 0.929	31.43 / 0.885	28.16 / 0.794	34.42 / 0.946	31.99 / 0.919	28.56 / 0.858
AirNet [9]	34.14 / 0.936	31.48 / 0.893	28.23 / 0.806	34.40 / 0.949	32.10 / 0.924	28.88 / 0.871
PromptIR [13]	34.34 / 0.938	31.71 / 0.897	28.49 / 0.813	34.77 / 0.952	32.49 / 0.929	29.39 / 0.881
DFPIR(Ours)	34.32 / 0.934	31.71 / 0.897	28.49 / 0.814	34.79 / 0.952	32.57 / 0.930	29.53 / 0.883

Table 4. Ablation experiments on the impact of key components in the multi-task restoration setting across three tasks. PSNR (dB, \uparrow) and SSIM (\uparrow) metrics are reported on the full RGB images. CA stands for Channel Attention, and SA stands for Spatial Attention.

Method	CA	DGCPM	CAAPM	SA	Dehazing on SOTS	Deraining on Rain100L	Denoising on CBSD68 dataset			Average
							$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	
Baseline	×	×	×	×	30.43 / 0.976	36.55 / 0.975	33.84 / 0.931	31.18 / 0.886	27.90 / 0.790	31.98 / 0.911
(a)	✓	×	×	×	31.01 / 0.978	37.47 / 0.979	33.94 / 0.933	31.28 / 0.888	27.99 / 0.793	32.34 / 0.914
(b)	✓	×	✓	×	31.86 / 0.979	37.86 / 0.980	34.01 / 0.934	31.37 / 0.890	28.14 / 0.801	32.65 / 0.917
(c)	×	✓	×	×	31.34 / 0.977	38.42 / 0.983	33.52 / 0.926	31.24 / 0.883	27.95 / 0.783	32.49 / 0.910
(d)	×	✓	×	✓	31.94 / 0.980	37.58 / 0.979	34.11 / 0.935	31.44 / 0.890	28.23 / 0.805	32.66 / 0.918
(e)	×	×	✓	×	31.17 / 0.979	38.64 / 0.983	33.74 / 0.925	31.31 / 0.887	28.07 / 0.796	32.59 / 0.914
DFPIR(Ours)	×	✓	✓	×	31.87 / 0.980	38.65 / 0.982	34.14 / 0.935	31.47 / 0.893	28.25 / 0.806	32.88 / 0.919

of preceding channel-wise perturbation intensifies inter-task interference. Therefore, enhancing perturbation in the attention dimension is necessary, with $\gamma = 0.7$ achieving the optimal average PSNR, effectively mitigating interference across multiple tasks. Furthermore, it can be observed that under different perturbation factor γ , the performance metrics for various tasks either improve or decline simultaneously. This further validates that our proposed perturbation method effectively mitigates cross-degradation interference, achieving balanced performance across multiple tasks.

Model Parameters and Computational Complexity.

In our proposed method, the channel shuffle only involves simple top-K channel rearrangement and 1×1 convolution for channel dimension transformation, resulting in a minimal increase in the number of parameters. The primary increase in parameters prompts in the attention dimension perturbation, specifically through cross-attention. From the perspective of the overall model’s parameter count and computational complexity, the additional parameters in our method, compared to Restormer, are less than 20%, and the computational cost is less than 8%. Both the parameter count and computational complexity are smaller than those of our main reference, PromptIR [13]. For detailed information, please refer to the Tab. 7. The use of a pre-trained CLIP text encoder to encode degradation descriptions and provide degradation prompts has a certain impact on the model’s inference time. However, the overall runtime remains relatively fast. Compared to the Baseline (Restormer)

and PromptIR, the inference time only increased by 11.6% and 5.0%, respectively, for the most time-consuming dehazing task. Detailed test results are presented on the Tab. 8.

Simple and detailed degradation prompts. We conducted ablation experiments on the prompt content, as shown in Tab. 6. We replaced the detailed degradation prompts (Method (dp)) with simple degradation prompts (Method (sp)), i.e., Noise, Haze, and Rain. The experimental results show that the overall performance remains nearly the same, with a slight improvement in dehazing performance and a minor decline in denoising performance across multiple noise levels. Although the performance slightly decreased across multiple noise levels, this also demonstrates some degree of generalization. The slight improvement in dehazing performance may be due to the reduced number of degradation types for the entire network, resulting in a minor upward trend in single-task performance. Overall, the results suggest that the content of the prompts does not need to be precisely designed; it is sufficient to provide the network with a degradation-type prompt. By using pre-trained CLIP [15] text prompts, we cleverly leverage the flexibility of text-based prompts, which can help address the restoration of complex multi-degradation or mixed-degradation images in future research.

5. Additional Visual Results

We provide additional qualitative results for the multi-task setting, including deraining, dehazing, and denoising (three-task setting), as well as deblurring and low-light en-

Table 5. Ablation experiments on the perturbation factor γ in the multi-task restoration setting across three tasks. PSNR (dB, \uparrow) and SSIM (\uparrow) metrics are reported on the full RGB images. * denotes w/o DGCPM, retaining only CAAPM.

γ	Dehazing on SOTS	Deraining on Rain100L	Denoising on CBSD68 dataset			Average
			$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	
0.5	31.65 / 0.979	38.50 / 0.982	33.96 / 0.932	31.39 / 0.890	27.88 / 0.781	32.67 / 0.913
0.7	31.75 / 0.980	38.52 / 0.982	34.11 / 0.935	31.46 / 0.892	28.21 / 0.803	32.81 / 0.919
0.8	31.87 / 0.980	38.53 / 0.982	34.11 / 0.935	31.46 / 0.893	28.23 / 0.804	32.84 / 0.919
0.9	31.87 / 0.980	38.65 / 0.982	34.14 / 0.935	31.47 / 0.893	28.25 / 0.806	32.88 / 0.919
1.0	31.78 / 0.980	38.57 / 0.982	34.11 / 0.935	31.46 / 0.893	28.23 / 0.805	32.83 / 0.919
0.5*	31.17 / 0.979	38.58 / 0.983	33.52 / 0.919	31.16 / 0.880	27.90 / 0.785	32.47 / 0.909
0.7*	31.17 / 0.979	38.64 / 0.983	33.74 / 0.925	31.31 / 0.887	28.07 / 0.796	32.59 / 0.914
0.8*	31.01 / 0.978	38.60 / 0.983	33.70 / 0.924	31.20 / 0.882	28.12 / 0.800	32.53 / 0.913
0.9*	31.11 / 0.979	38.63 / 0.983	33.62 / 0.923	31.13 / 0.879	28.11 / 0.798	32.52 / 0.912
1.0*	30.87 / 0.978	38.60 / 0.983	33.67 / 0.923	31.17 / 0.880	28.13 / 0.799	32.49 / 0.913

Table 6. Ablation experiments in the multi-task restoration setting using simple degradation prompts, i.e., Noise, Rain, and Haze, across three tasks. PSNR (dB, \uparrow) and SSIM (\uparrow) metrics are reported on the full RGB images. Method(sp) indicates using simple degradation prompts, while method(dp) indicates using detailed degradation prompts.

Method	Dehazing on SOTS	Deraining on Rain100L	Denoising on CBSD68 dataset			Average
			$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	
(sp)	31.93 / 0.981	38.63 / 0.982	34.12 / 0.935	31.45 / 0.891	28.18 / 0.801	32.86 / 0.918
(dp)	31.87 / 0.980	38.65 / 0.982	34.14 / 0.935	31.47 / 0.893	28.25 / 0.806	32.88 / 0.919

Table 7. The model parameters and computational complexity are evaluated. GMACS are computed on a 256×256 input image using an NVIDIA RTX 3090 GPU.

Method	Params.	GMACS
Restormer(Baseline)	26.1M	141.2G
PromptIR	35.6M	158.4G
DFPIR(Ours)	31.1M	151.4G

Table 8. The inference time is measured on a single NVIDIA RTX 3090 GPU. The inference time refers to the total runtime on the entire dataset.

Method	Dehazing on SOTS	Deraining on Rain100L	Denoising on CBSD68 dataset		
			$\sigma = 15$	$\sigma = 25$	$\sigma = 50$
Restormer	155.78s	19.19s	13.07s	12.99s	13.02s
PromptIR	165.52s	20.60s	13.93s	13.97s	13.98s
DFPIR(ours)	173.80s	21.71s	14.84s	14.66s	14.74s

hancement (five-task setting). Visual examples are presented in Fig. 2 for dehazing, Fig. 3 for deraining, Fig. 4 for denoising, Fig. 5 for deblurring and Fig. 6 for lowlight enhancement. These examples demonstrate the effectiveness of our DFPIR method in removing degradations and producing images that are visually closer to the ground truth compared to other approaches [3, 13]. Notably, the restored images from our method exhibit superior structural fidelity

and preserve fine textures more effectively.

We also visualize the features before and after perturbation for multiple degradations in Fig. 7. As shown in the figure, the features before perturbation clearly contain severe degradation artifacts, such as noise, rain streaks, and blur. After the perturbation in the channel dimension (DGCPM), the degradation features are significantly suppressed, while the inherent features of the image, such as textures, are enhanced. For example, regarding the features of blur degradation, the features after the channel-wise perturbation (DGCPM) almost achieve a deblurring effect. This means that the channel-wise perturbation we proposed can significantly reduce the impact of various degradation features and robustly extract the inherent features of the image. Although the perturbation in the channel dimension enhances the inherent features of the image, it can be observed that some detail information is lost. However, after the perturbation in the attention dimension (DGCPM+CAAPM), the features not only retain the inherent characteristics of the image and further suppress the degradation features, but also enhance the feature details. These results contribute to the improvement of image restoration quality in the decoding stage. Overall, our feature visualization results support the claims in the ‘‘Analysis of Our Feature Perturbation Strategy’’ section, showing that the proposed perturbation-based method preserves the inherent features of the image, reduces degradation feature interference, and improves image restoration performance in multi-degradation scenarios.

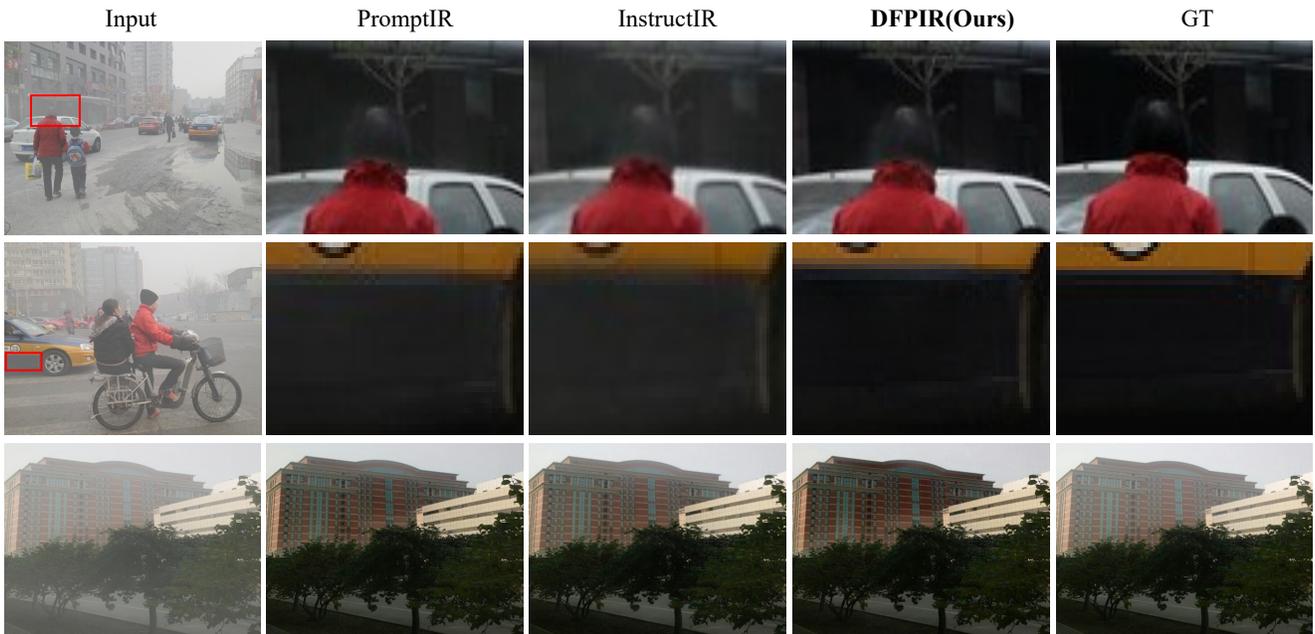


Figure 2. Image dehazing comparisons on SOTS [8] in the three-degradation setting. The image quality of the results produced by our DFPIR is visually better than the previous state-of-the-art approach InstructIR [3].

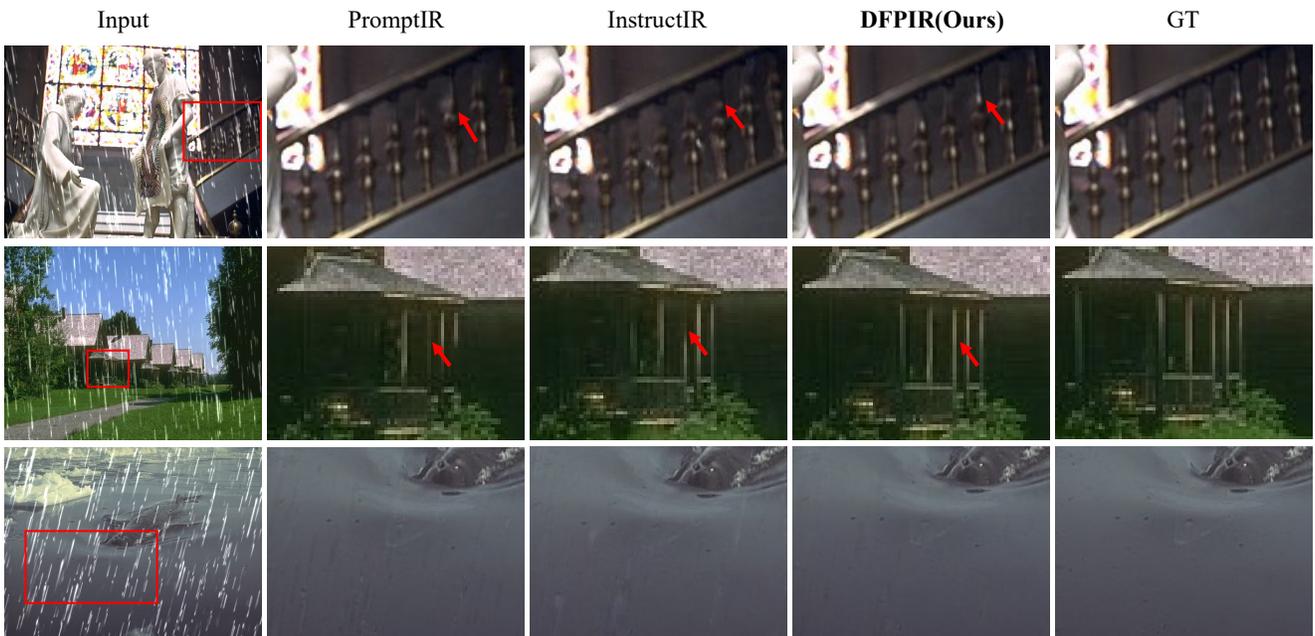


Figure 3. Image deraining comparisons on Rain100L [20] in the three-degradation setting. Our method effectively removes rain streaks to generate rain-free images.

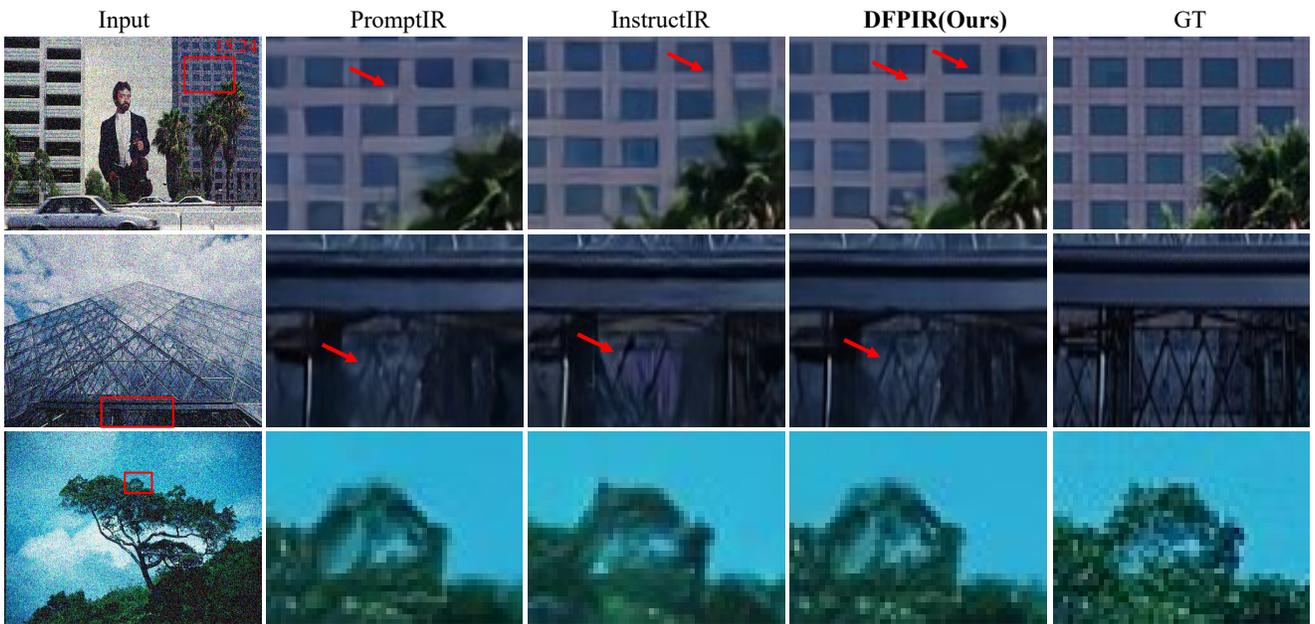


Figure 4. Image denoising comparisons on CBSD68 [11] in the three-degradation setting. The image reproduction quality of our DFPIR is more visually faithful to the ground truth.

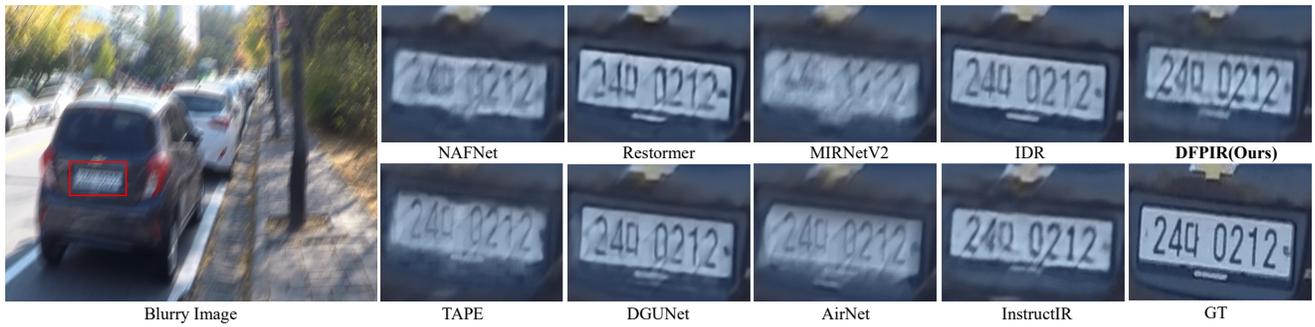


Figure 5. Image Deblurring Results. Comparison with other methods on the GoPro [12] dataset (GOPR0854-11-00-000001.png).

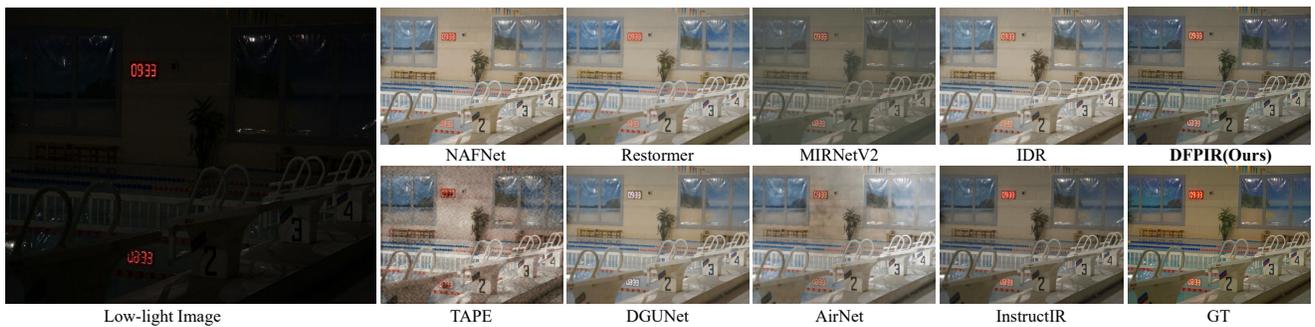
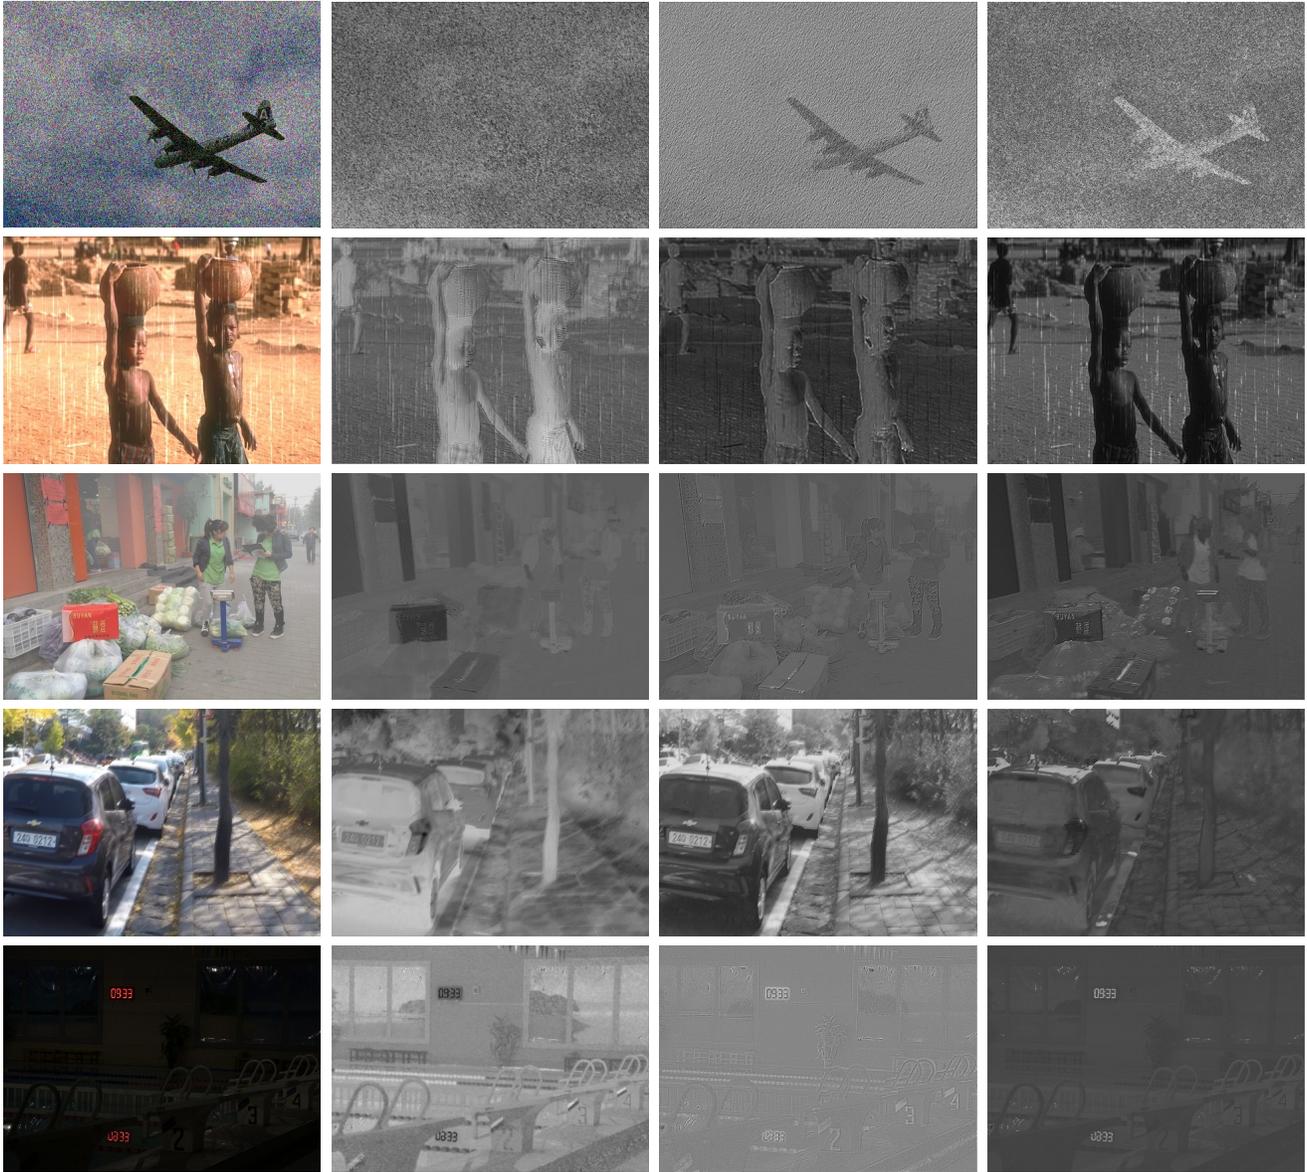


Figure 6. Low-light Image Enhancement Results. We compare with other methods on LOL [18].



Input

Pre-perturbation

DGCPM

DGCPM+CAAPM

Figure 7. More feature visualization. Our degradation-aware strategy is capable of extracting the inherent features of images from various degraded inputs, suppressing degradation-specific features, thereby enhancing restoration performance.

References

- [1] Pablo Arbelaez, Michael Maire, Charless Fowlkes, and Jitendra Malik. Contour detection and hierarchical image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 33(5):898–916, 2010. 1
- [2] Bolun Cai, Xiangmin Xu, Kui Jia, Chunmei Qing, and Dacheng Tao. Dehazenet: An end-to-end system for single image haze removal. *IEEE transactions on image processing*, 25(11):5187–5198, 2016. 2
- [3] Marcos V Conde, Gregor Geigle, and Radu Timofte. Instructir: High-quality image restoration following human instructions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2024. 1, 4, 5
- [4] Yu Dong, Yihao Liu, He Zhang, Shifeng Chen, and Yu Qiao. Fd-gan: Generative adversarial networks with fusion-discriminator for single image dehazing. In *Proceedings of the AAAI conference on artificial intelligence*, pages 10729–10736, 2020. 2
- [5] Hongyun Gao, Xin Tao, Xiaoyong Shen, and Jiaya Jia. Dynamic scene deblurring with parameter selective sharing and nested skip connections. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3848–3856, 2019. 2
- [6] Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5197–5206, 2015. 1, 2, 3
- [7] Kui Jiang, Zhongyuan Wang, Peng Yi, Chen Chen, Baojin Huang, Yimin Luo, Jiayi Ma, and Junjun Jiang. Multi-scale progressive fusion network for single image deraining. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8346–8355, 2020. 2
- [8] Boyi Li, Wenqi Ren, Dengpan Fu, Dacheng Tao, Dan Feng, Wenjun Zeng, and Zhenyang Wang. Benchmarking single-image dehazing and beyond. *IEEE Transactions on Image Processing*, 28(1):492–505, 2018. 1, 2, 5
- [9] Boyun Li, Xiao Liu, Peng Hu, Zhongqin Wu, Jiancheng Lv, and Xi Peng. All-in-one image restoration for unknown corruption. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17452–17462, 2022. 1, 2, 3
- [10] Kede Ma, Zhengfang Duanmu, Qingbo Wu, Zhou Wang, Hongwei Yong, Hongliang Li, and Lei Zhang. Waterloo exploration database: New challenges for image quality assessment models. *IEEE Transactions on Image Processing*, 26(2):1004–1016, 2016. 1
- [11] David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, pages 416–423. IEEE, 2001. 1, 2, 3, 6
- [12] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3883–3891, 2017. 1, 6
- [13] Vaishnav Potlapalli, Syed Waqas Zamir, Salman Khan, and Fahad Khan. Promptir: Prompting for all-in-one image restoration. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. 1, 2, 3, 4
- [14] Yanyun Qu, Yizi Chen, Jingying Huang, and Yuan Xie. Enhanced pix2pix dehazing network. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8160–8168, 2019. 2
- [15] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 1, 3
- [16] Wenqi Ren, Si Liu, Hua Zhang, Jinshan Pan, Xiaochun Cao, and Ming-Hsuan Yang. Single image dehazing via multi-scale convolutional neural networks. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, pages 154–169. Springer, 2016. 2
- [17] Chunwei Tian, Yong Xu, and Wangmeng Zuo. Image denoising using deep cnn with batch renormalization. *Neural Networks*, 121:461–473, 2020. 2, 3
- [18] Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. *arXiv preprint arXiv:1808.04560*, 2018. 1, 6
- [19] Wei Wei, Deyu Meng, Qian Zhao, Zongben Xu, and Ying Wu. Semi-supervised transfer learning for image rain removal. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3877–3886, 2019. 2
- [20] Fuzhi Yang, Huan Yang, Jianlong Fu, Hongtao Lu, and Bain-ing Guo. Learning texture transformer network for image super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5791–5800, 2020. 1, 2, 5
- [21] Rajeev Yasarla and Vishal M Patel. Uncertainty guided multi-scale residual learning-using a cycle spinning cnn for single image de-raining. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8405–8414, 2019. 2
- [22] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5728–5739, 2022. 2
- [23] Jinghao Zhang, Jie Huang, Mingde Yao, Zizheng Yang, Hu Yu, Man Zhou, and Feng Zhao. Ingredient-oriented multi-degradation learning for image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5825–5835, 2023. 1
- [24] Kai Zhang, Wangmeng Zuo, Shuhang Gu, and Lei Zhang. Learning deep cnn denoiser prior for image restoration. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3929–3938, 2017. 2, 3
- [25] Kai Zhang, Wangmeng Zuo, and Lei Zhang. Ffdnet: Toward a fast and flexible solution for cnn-based image denoising.

IEEE Transactions on Image Processing, 27(9):4608–4622, 2018. [2](#), [3](#)

- [26] Xu Zhang, Jiaqi Ma, Guoli Wang, Qian Zhang, Huan Zhang, and Lefei Zhang. Perceive-ir: Learning to perceive degradation better for all-in-one image restoration. *arXiv preprint arXiv:2408.15994*, 2024. [1](#)