Learning Enriched Features for Real Image Restoration and Enhancement

Syed Waqas Zamir[1], Aditya Arora[1], Salman Khan[1,2], Munawar Hayat[1,2], Fahad Shahbaz Khan[1,2], Ming-Hsuan Yang[3,4], and Ling Shao[1,2]

[1] Inception Institute of Artificial Intelligence, UAE
[2] Mohamed bin Zayed University of Artificial Intelligence, UAE
[3] University of California, Merced, USA [4] Google Research

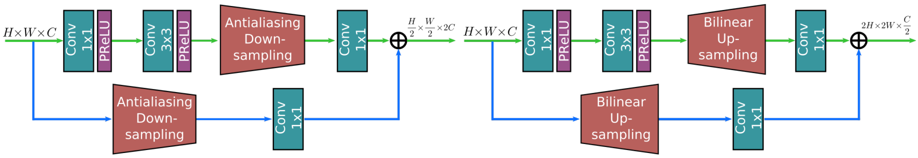

(a) Downsampling module (b) Upsampling module

Fig. 1: Residual resizing modules to perform downsampling and upsampling.

In this supplementary document, we provide:

- An ablation study on residual resizing modules.
- Several denoising comparisons on SIDD dataset [1].
- Additional denoising results on DND dataset [9].
- Super-resolved visual examples on RealSR dataset [4].
- Several image enhancement comparisons on MIT-Adobe FiveK dataset [3].

1. Additional Ablation Study

The proposed MIRNet employs a recursive residual design (with skip connections) to ease the flow of information during the learning process. In order to maintain the residual nature of our architecture, we introduce residual resizing modules to perform downsampling (Fig. 1a) and upsampling (Fig. 1b) operations.

In this section, we demonstrate the effectiveness of the *residual resizing modules*. This ablation experiment is performed for the super-resolution task with $\times 3$ scale factor. Table 1 shows that, when both the residual branch (blue in color) and the antialiasing operation are removed, the performance is relatively low (30.98 dB PSNR). After adding the antialiasing downsampling [13] to the main branch (green) of Fig. 1a, the PSNR score is increased from 30.98 dB to 31.05 dB. Finally, the combination of all the components in residual resizing modules yield significantly improved results (31.16 dB) than only using the main branch (30.98 dB).

Table 1: Ablation study on residual resizing modules.

Main branch (Green)	✓	✓	✓	✓
Residual branch (Blue)			✓	✓
Antialiasing [13]		✓		✓
PSNR (in dB)	30.98	31.05	31.11	**31.16**

2. Image Denoising

Here we provide additional results for image denoising on real image datasets.

SIDD dataset [1]: Figures 2 and 3 show results produced by our method and those of the state-of-the-art approaches (CBDNet [8], RIDNet [2], and VDN [11]). It can be seen that our method yields favorable results both visually and in terms of image quality metrics (PSNR and SSIM).

DND dataset [9]: Figures 4 and 5 demonstrate that our method is more effective in removing real noise than other competing algorithms.

3. Super-resolution

In Figure 6, we present the full-resolution versions of the super-resolved images provided in Fig. 8 of the main paper. Our method produces sharp and natural images. In contrast, the recent best method LP-KPN [4] has a tendency to over-enhance the contrast, and therefore yields images that are perceptually less faithful to the ground-truth, which is undesirable for several applications. For example, in Television industry, those restoration methods are preferred that preserve as much as possible the artistic intent (in terms of brightness, color and contrast) of the content creator.

4. Image Enhancement

We provide several visual comparisons of image enhancement on the MIT-Adobe FiveK [3] in Figures 7 and 8. Compared to other techniques, the proposed MIRNet makes better color and contrast adjustments and generates images that are vivid and natural in appearance.

5. Joint Denoising and Super-resolution

To test our model for multiple degradations, we perform the following experiment. We take the paired data from RealSR dataset and synthesize realistic noise in LR images using [12], thus defining a joint denoising and SR task. We train RCAN [14] and MIRNet on this data, and evaluate on (noisy) test set of

Table 2: Joint denoising and super-resolution experiment. PSNR (dB) is reported.

SR	27.16	30.40	31.16
Denoising + SR	24.62	28.41	28.81

RealSR. Table 2 shows that our method achieves superior performance with a gain of 0.4 dB over RCAN [14].

References

1. Abdelhamed, A., Lin, S., Brown, M.S.: A high-quality denoising dataset for smartphone cameras. In: CVPR (2018) 1, 2, 4, 5
2. Anwar, S., Barnes, N.: Real image denoising with feature attention. ICCV (2019) 2, 4, 5, 6, 7
3. Bychkovsky, V., Paris, S., Chan, E., Durand, F.: Learning photographic global tonal adjustment with a database of input/output image pairs. In: CVPR (2011) 1, 2, 9, 10
4. Cai, J., Zeng, H., Yong, H., Cao, Z., Zhang, L.: Toward real-world single image super-resolution: A new benchmark and a new model. In: ICCV (2019) 1, 2, 8
5. Chen, Y.S., Wang, Y.C., Kao, M.H., Chuang, Y.Y.: Deep photo enhancer: Unpaired learning for image enhancement from photographs with gans. In: CVPR (2018) 9, 10
6. https://noise.visinf.tu-darmstadt.de/benchmark/ (2017), [Online; accessed 29-Feb-2020] 6, 7
7. Gharbi, M., Chen, J., Barron, J.T., Hasinoff, S.W., Durand, F.: Deep bilateral learning for real-time image enhancement. TOG (2017) 9, 10
8. Guo, S., Yan, Z., Zhang, K., Zuo, W., Zhang, L.: Toward convolutional blind denoising of real photographs. In: CVPR (2019) 2, 4, 5, 6, 7
9. Plotz, T., Roth, S.: Benchmarking denoising algorithms with real photographs. In: CVPR (2017) 1, 2, 6, 7
10. Wei, C., Wang, W., Yang, W., Liu, J.: Deep retinex decomposition for low-light enhancement. BMVC (2018) 9, 10
11. Yue, Z., Yong, H., Zhao, Q., Meng, D., Zhang, L.: Variational denoising network: Toward blind noise modeling and removal. In: NeurIPS (2019) 2, 4, 5, 6, 7
12. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H., Shao, L.: CycleISP: Real image restoration via improved data synthesis. In: CVPR (2020) 2
13. Zhang, R.: Making convolutional networks shift-invariant again. In: ICML (2019) 1, 2
14. Zhang, Y., Li, K., Li, K., Wang, L., Zhong, B., Fu, Y.: Image super-resolution using very deep residual channel attention networks. In: ECCV (2018) 2, 3

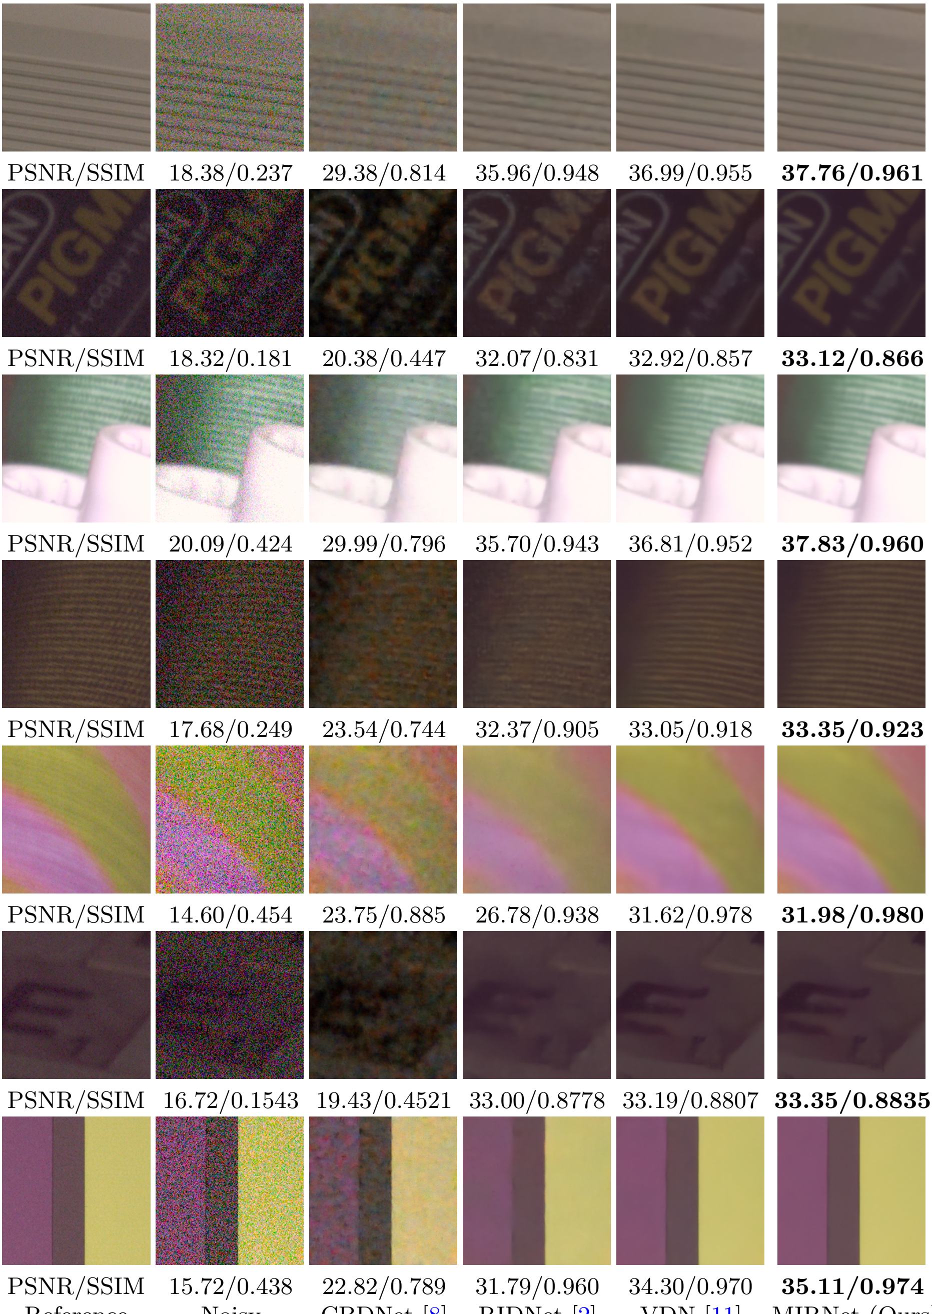

Fig. 2: Denoising examples from the SIDD benchmark dataset [1].

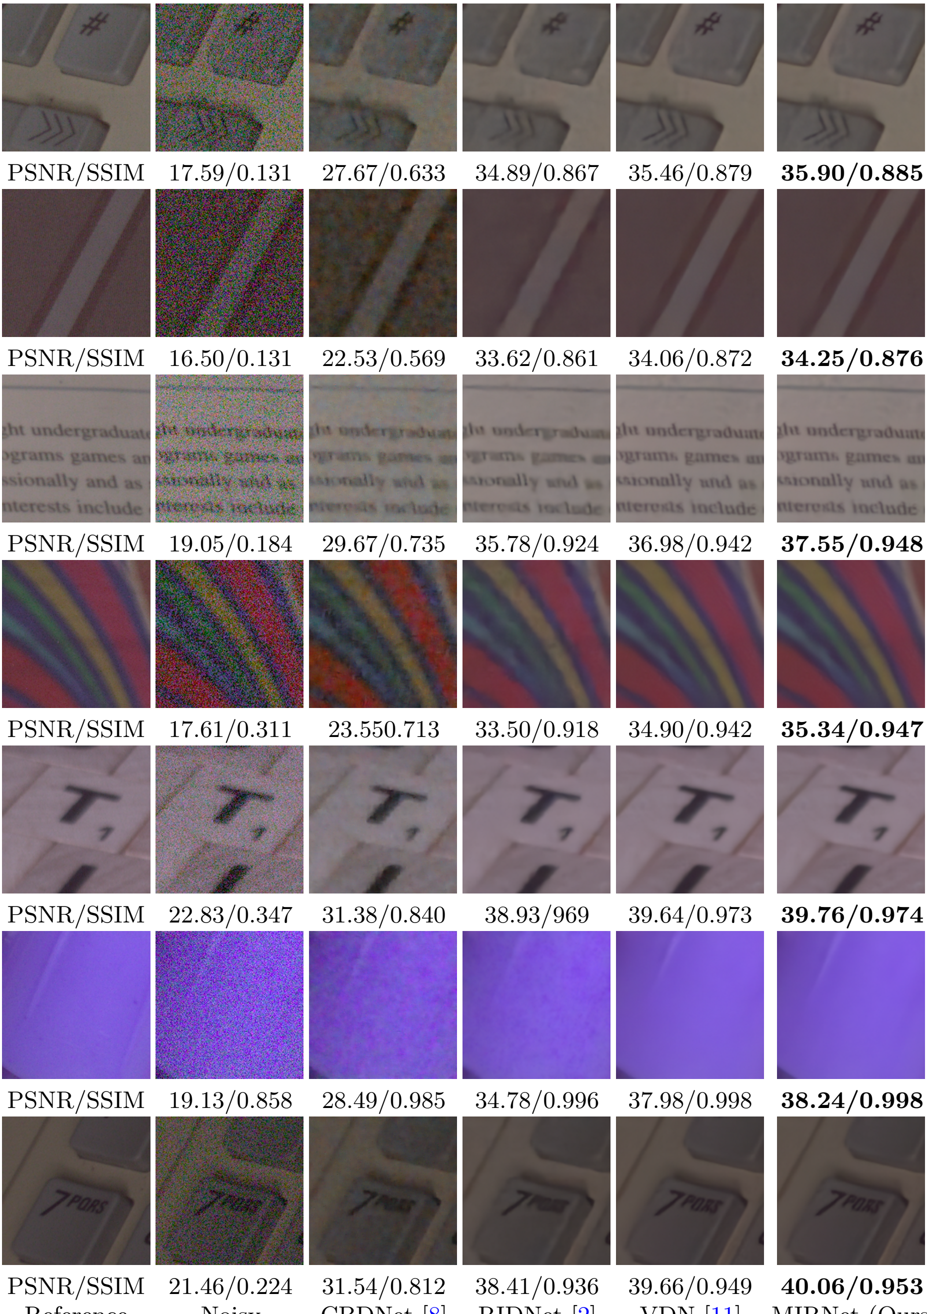

Fig. 3: Denoising examples from the SIDD benchmark dataset [1].

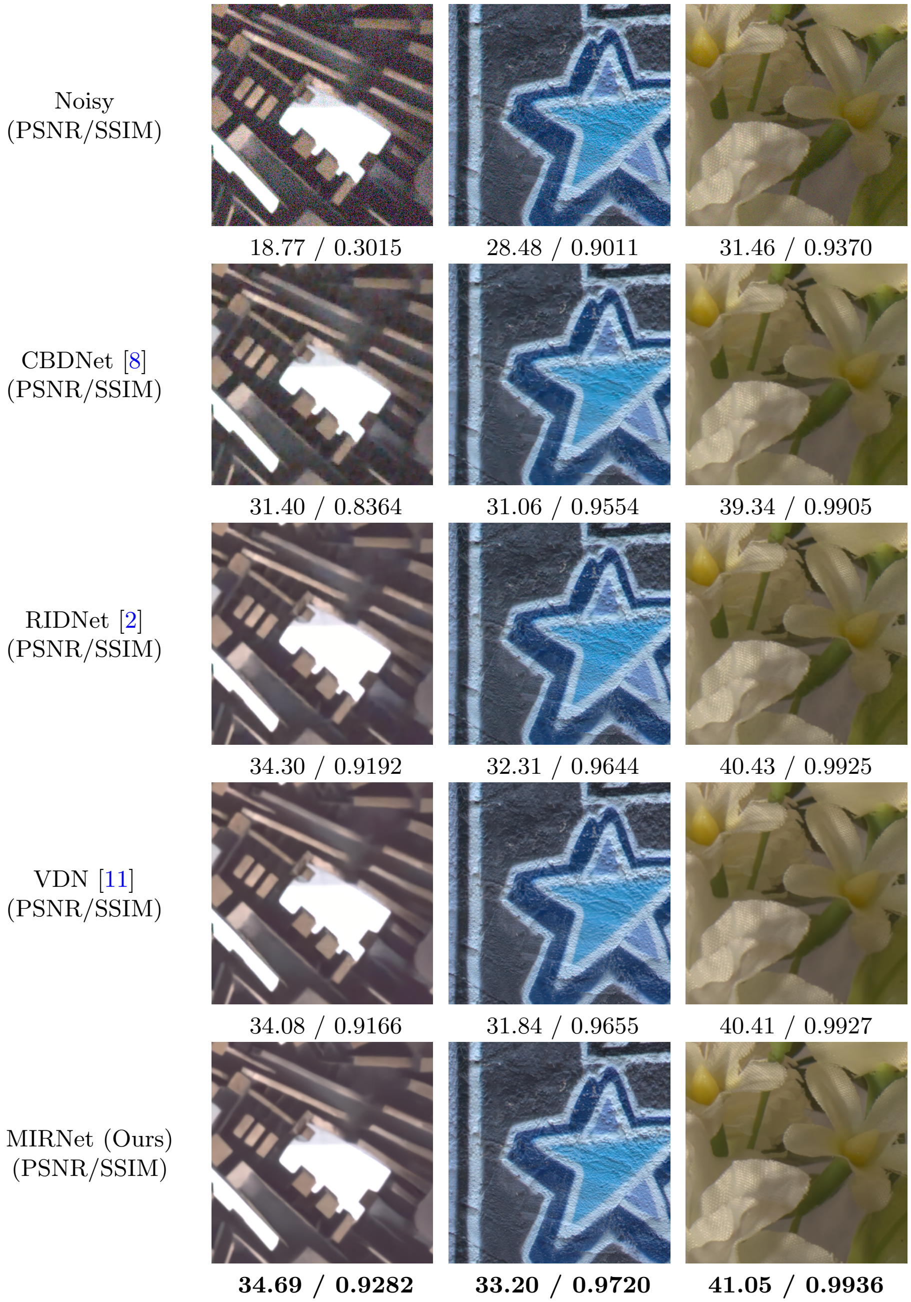

Fig. 4: Denoising examples from the DND benchmark dataset [9]. PSNR and SSIM scores for all competing methods are obtained from the website of the DND evaluation server [6].

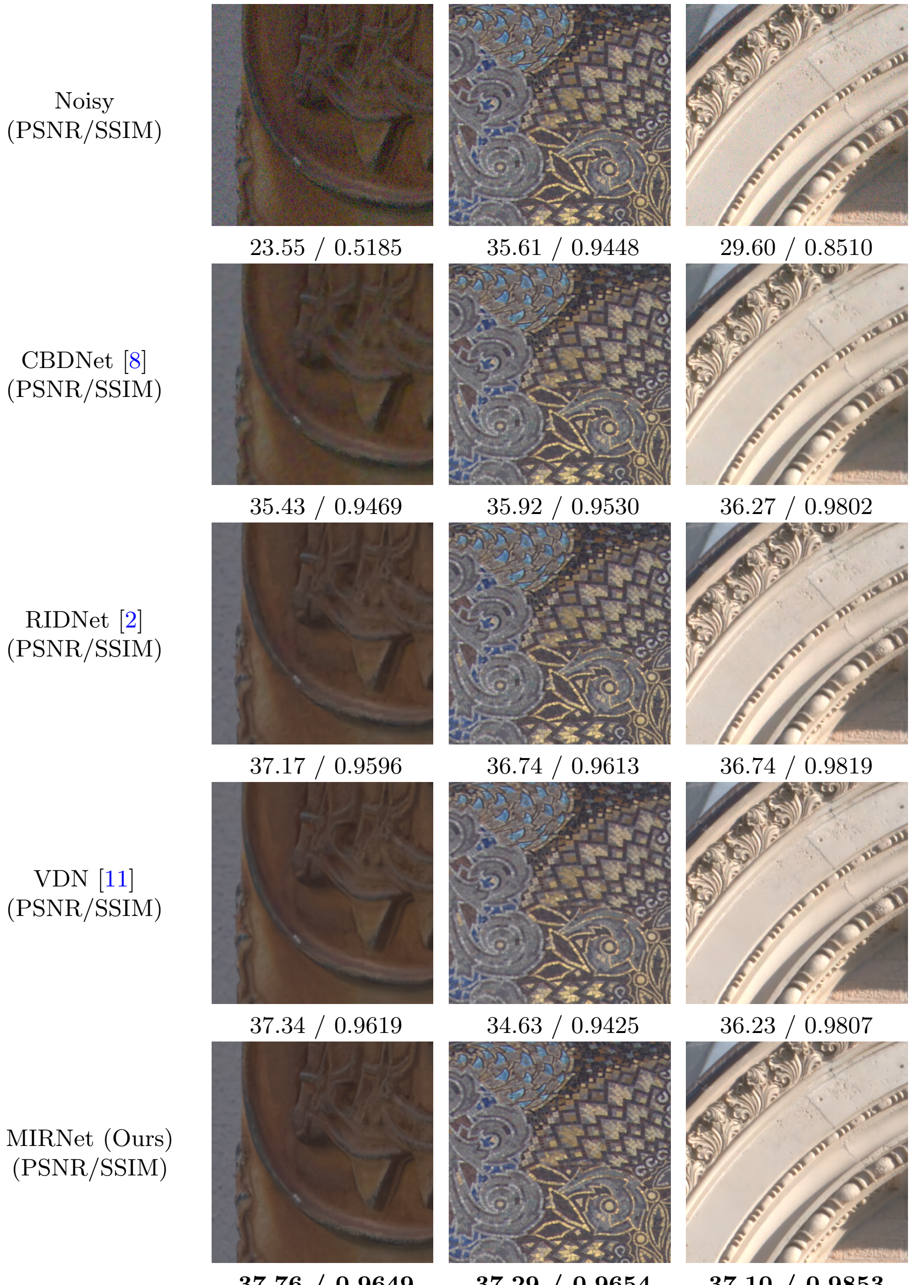

Fig. 5: Denoising examples from the DND benchmark dataset [9]. PSNR and SSIM scores for all competing methods are obtained from the website of the DND evaluation server [6].

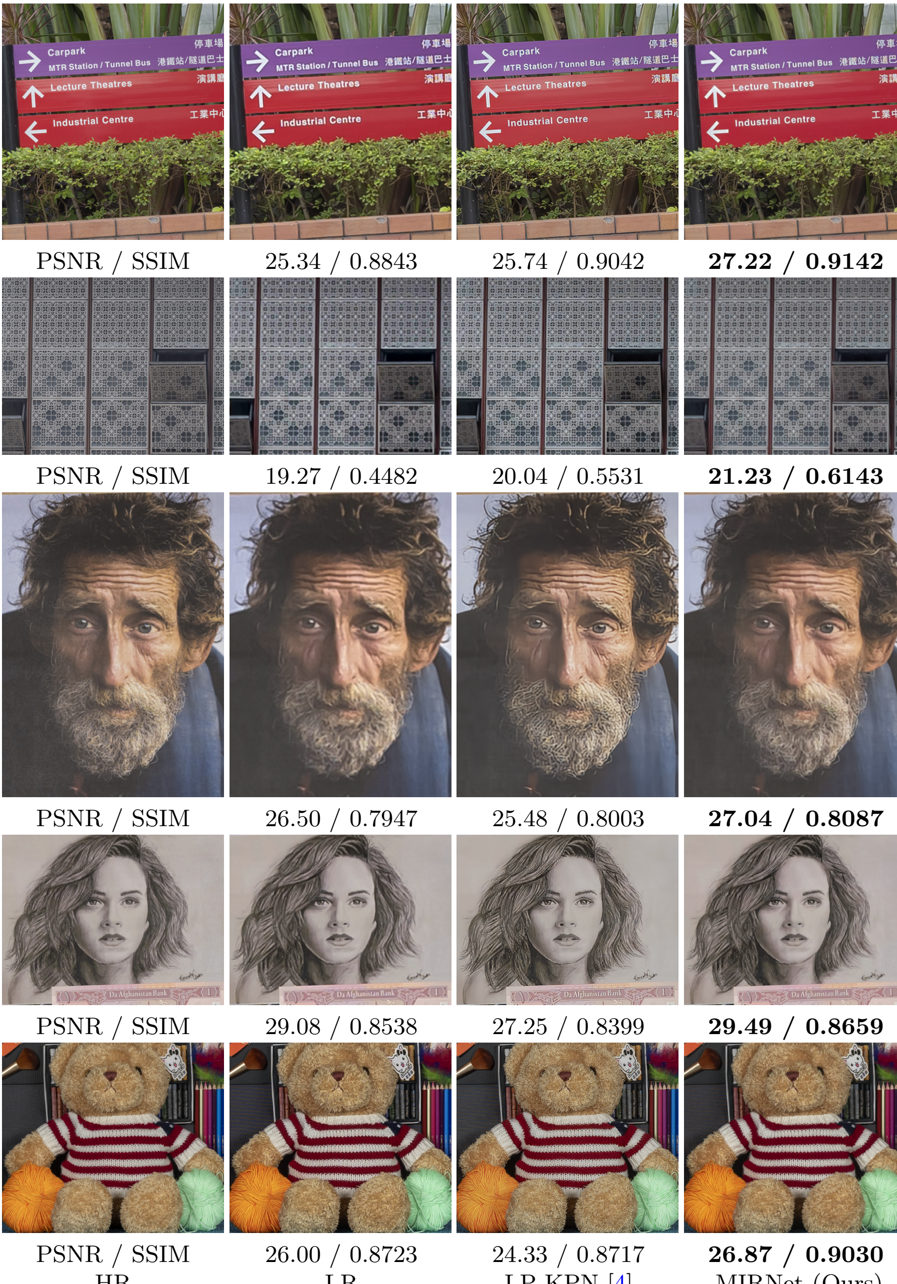

Fig. 6: Super-resolution ($\times 4$). The full-resolution versions of examples provided in Fig. 8 of the main paper. Zoom-in for better visualization.

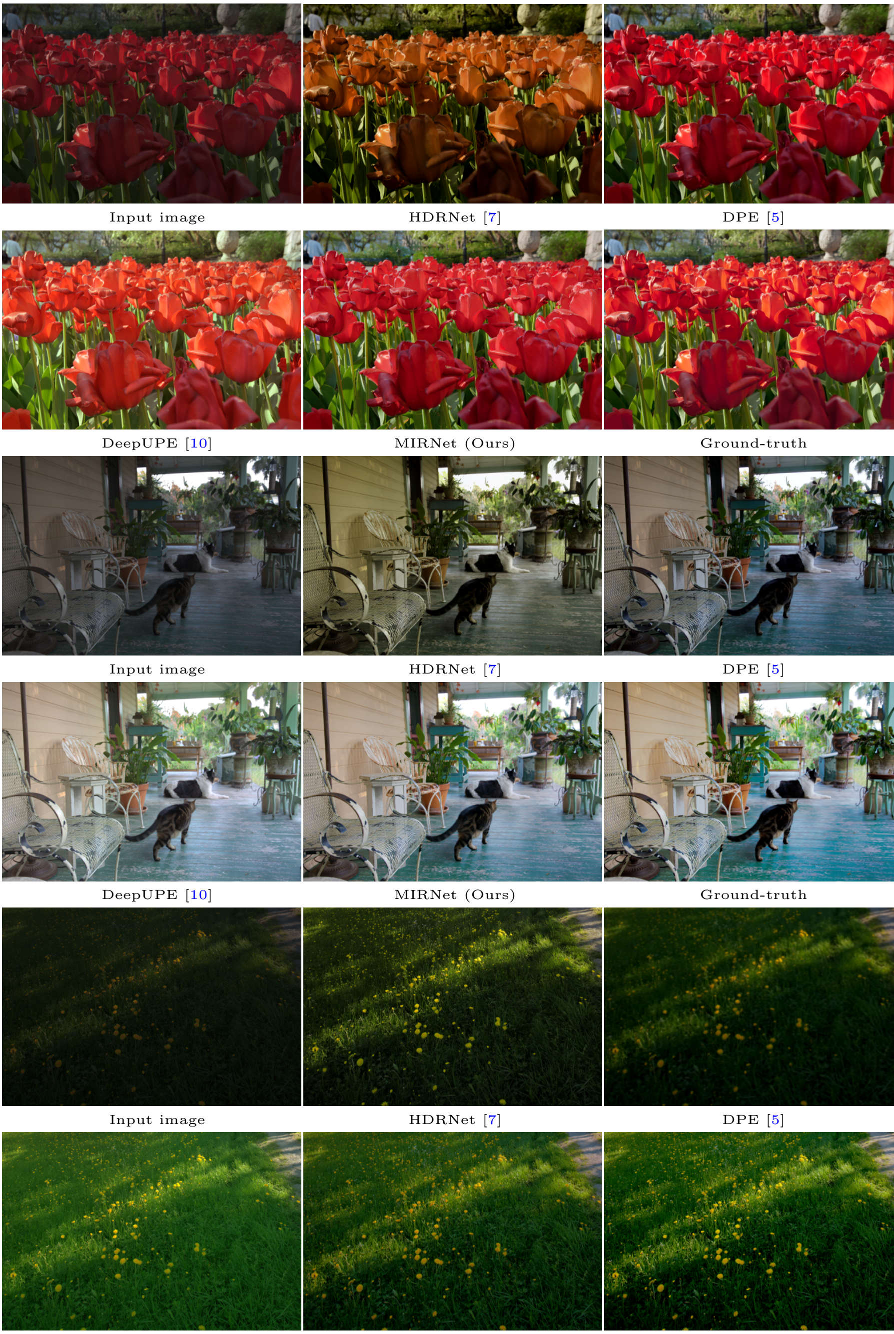

Fig. 7: Visual results of image enhancement on the MIT-Adobe FiveK [3] dataset. Compared to the state-of-the-art, our MIRNet makes better brightness, color and contrast adjustments, while staying more faithful to the ground-truth.

Fig. 8: Visual results of image enhancement on the MIT-Adobe FiveK [3] dataset. Compared to the state-of-the-art, our MIRNet makes better brightness, color and contrast adjustments, while staying more faithful to the ground-truth.