FoundIR-v2: Optimizing Pre-Training Data Mixtures for Image Restoration Foundation Model
- *Supplemental Material* -

Xiang Chen[1] Jinshan Pan[1*] Jiangxin Dong[1] Jian Yang[1] Jinhui Tang[2]
[1]Nanjing University of Science and Technology [2]Nanjing Forestry University

Overview

In this document, we first provide an explanation of data mixing laws in Section 1. In Section 2, we present a comparative summary of the training dataset configurations used by existing universal/all-in-one image restoration methods and our proposed FoundIR-v2. Next, we further discuss the comparison with image restoration agents in Section 3. Finally, we provide more experimental results in Section 4.

1. Data Mixing Laws

Data mixing laws, also referred to as data re-weighting, aim to refine the distribution of data characteristics within a curated dataset to improve the model's multi-capability performance. For training image restoration foundation models, when the model capacity is fixed, employing a fixed or imbalanced data mixing ratio p inevitably skews the gradient contributions towards specific tasks. This results in biased specialization during the early training stage, thereby diminishing the generality of early-stage representations. In this paper, we introduce a data equilibrium scheduling (DES) strategy that dynamically determines the sampling probability $p_k(t)$ for each degradation type k at iteration t. At the beginning of training, $p_k(t)$ is initialized close to a uniform distribution to ensure adequate gradient coverage across all degradation types. During the mid stage, $p_k(t)$ gradually shifts toward more challenging restoration tasks. In the final stage, $p_k(t)$ converges to a task-dependent steady-state distribution aligned with the intended target. The corresponding optimization procedure can be expressed in a time-integral form as follows:

$$\min_{\theta} \int_0^T \left[\sum_{k=1}^{K} p_k(t) \mathcal{L}_k(\theta(t)) \right] dt, \tag{1}$$

where $\mathcal{L}_k$ denotes the loss corresponding to degradation type k, $\theta(t)$ represents model parameters at iteration t. This mechanism is equivalent to maintaining a gradient equilibrium point throughout training, ensuring that the contribution of each degradation type remains controlled and relatively balanced, thereby preventing collapse and fostering more generalizable representations.

2. Training Data for FoundIR-v2

We summarize the statistics of the training datasets used by existing all-in-one/universal image restoration methods and our FoundIR-v2 (see Table 1). FoundIR [19] demonstrates the critical role of real-world training data in all-in-one image restoration by constructing a unified data collection pipeline that facilitates the acquisition of a million-scale high-quality paired dataset, thereby advancing image restoration foundation models. To further improve the performance of the image restoration foundation model, we augment the original FoundIR training data by incorporating more public high-quality datasets into the pre-training pool for FoundIR-v2. Compared with existing methods, the training data composition of FoundIR-v2 is significantly more diverse: beyond natural images, it also includes remote sensing images, underwater scenes, mural images, and other types. This broadened data distribution equips the foundation model with richer visual priors, leading to improved generalization, stronger robustness across heterogeneous degradations, and better adaptability to real-world scenarios.

Table 1. A summarization of existing all-in-one/universal image restoration methods capable of handling multiple degradations simultaneously, which could be regarded as potential foundation models for low-level vision tasks. "Self-Generated" and "Self-Collected" denotes images generated and collected by the author rather than common datasets.

Methods	Venue	Training Datasets	Degradation Data Types				Num. (k)
			Syn	Real	Isolated	Coupled	
Li et al. [20]	CVPR'20	Outdoor-Rain, Snow100K, Raindrop	✔	✔	✔	✗	110.1
IPT [3]	CVPR'21	Self-Generated	✔	✗	✔	✗	100.0
BSRGAN [49]	ICCV'21	DIV2K, Flick2K, WED, FFHQ	✔	✗	✔	✔	10.1
Real-ESRGAN [35]	ICCVW'21	DIV2K, Flickr2K, OutdoorSceneTraining	✔	✗	✔	✔	453.4
AirNet [18]	CVPR'22	BSD400, BSD68, WED, Urban100, Rain100L, RESIDE	✔	✗	✔	✗	77.4
Chen et al. [6]	CVPR'22	RESIDE, Rain1400, CSD	✔	✗	✔	✗	15.0
TransWeather [32]	CVPR'22	Snow100K, Raindrop, Outdoor-Rain	✔	✔	✔	✗	18.0
BIDeN [21]	ECCV'22	Self-Generated	✔	✗	✔	✔	10.1
TAPE [24]	ECCV'22	SIDD, Rain200L/H, Raindrop, Snow100K, ISTD, DIV2K, REDS	✔	✗	✔	✔	146.5
CPNet [33]	CVPR'23	Self-Generated	✔	✗	✔	✔	1000.0
GDP [12]	CVPR'23	ImageNet, LSUN, CelebA, USC, LOL, VE-LOL-V, LoLi-Phone, NTIRE	✔	✗	✔	✗	-
IDR [48]	CVPR'23	Rain200L, RESIDE, BSD400, WED, GoPro, LOL	✔	✗	✔	✗	109.1
Zhu et al. [53]	CVPR'23	SPA+, RealSnow, Outdoor-Rain, RainDrop, Snow100K, RESIDE, REVIDE	✔	✔	✔	✔	174.0
AWRCP [40]	ICCV'23	Snow100K, Outdoor-Rain, RainDrop	✔	✔	✔	✗	60.1
AMIRNet [47]	MM'23	RED4, SIDD, LOL, DPDD, WED	✔	✗	✔	✗	2.4
PromptIR [29]	NeurIPS'23	BSD400, WED, Rain100L, SOTS	✔	✗	✔	✗	77.4
Ozdenizci et al. [28]	TPAMI'23	Snow100K, Outdoor-Rain, RainDrop	✔	✗	✔	✗	59.8
DiffUIR [51]	CVPR'24	Rain13k, LOL, Snow100K, RESIDE, GoPro	✔	✗	✔	✗	138.4
SeeSR [38]	CVPR'24	LSDIR, FFHQ	✔	✗	✔	✔	184.9
SUPIR [41]	CVPR'24	Self-Generated	✔	✗	✔	✔	20000.0
PASD [39]	ECCV'24	DIV2K, Flickr2K, OST, FFHQ	✔	✗	✔	✔	13.4
AutoDIR [17]	ECCV'24	SIDD, DIV2K, Flickr2K, GoPro, LOL, RESIDE, Rain200L, Raindrop	✔	✗	✔	✗	114.7
GRIDS [2]	ECCV'24	Self-Generated	✔	✗	✔	✗	-
InstructIR [9]	ECCV'24	BSD400, WED, Rain200L, SOTS, GoPro, LOL, DSLR	✔	✗	✔	✔	10.7
OneRestore [14]	ECCV'24	Self-Generated, LOL, RESIDE, Rain1200, Snow100k	✔	✗	✔	✔	312.0
RAM [30]	ECCV'24	RESIDE, Rain13k, GoPro, LOL-V2, LSDIR	✔	✗	✔	✗	-
UniProcessor [11]	ECCV'24	DIV2K, Flicker2K, BSD500, WED	✔	✗	✔	✗	71.5
DA-CLIP [27]	ICLR'24	Self-Generated	✔	✗	✔	✗	52.8
StableSR [34]	IJCV'24	DIV2K, DIV8K, Flickr2K, OutdoorSceneTraining, FFHQ	✔	✗	✔	✔	458.4
DreamClear [1]	NeurIPS'24	Self-Generated, DIV2K, Flickr2K, LSDIR, DIV8K	✔	✗	✔	✔	1000.0
ResShift [42]	TPAMI'24	Self-Generated	✔	✗	✔	✔	-
Art [36]	MM'24	Self-Generated	✔	✗	✔	✔	-
GenLV [8]	MM'24	Self-Generated	✔	✗	✔	✔	-
DCPT [15]	ICLR'25	Rain200L/H, RESIDE, Snow100K, BSD400, WED, GoPro, DPDD, LOL	✔	✗	✔	✔	82.8
AdaIR [10]	ICLR'25	SOTS, Rain100L, BSD400, WED	✔	✗	✔	✗	77.4
FaithDiff [5]	CVPR'25	LSDIR, DIV2K, DIV8K, Flicker2K, FFHQ	✔	✗	✔	✔	99.9
GenDeg [31]	CVPR'25	Self-Generated	✔	✗	✔	✗	550.0
MoCE-IR [43]	CVPR'25	BSD400, WED, Rain100L, SOTS, GoPro, LOL, CDD11	✔	✗	✔	✔	79.5
VLU-Net [46]	CVPR'25	BSD400, WED, RESIDE, Rain100L, GoPro, LOL	✔	✗	✔	✔	79.6
UniRestore [4]	CVPR'25	DIV2K, Flickr2K, OST, Self-Generated	✔	✗	✔	✗	80.0
FoundIR [19]	ICCV'25	Self-Collected	✔	✔	✔	✔	1011.6
FoundIR-v2	-	FoundIR-TrainData, 4KRD, MC-Blur, Dense-HAZE, NH-HAZE, HQ-Rain, LSDIR, HQ-NightRain, UAV-Rain1k, WeatherBench, StateHaze1k, FFHQ, DIV2k, DIV8k, PolyU, UHD-LL, Flickr2k, Mural-Dunhuang, Mural-Shaanxi	✔	✔	✔	✔	1185.9

3. Discussion with Image Restoration Agents

Recently, image restoration agents have been emerging that automatically invoke different expert models to address a broader range of real-world degradation scenarios. However, we note that these agents perform chain-of-thought reasoning, sequentially carrying out image restoration steps based on the perceived types of degradation. In fact, as the complexity of real-world degradations increases, overly long reasoning chains may lead to snowball errors, where early-stage mistakes accumulate and ultimately degrade the quality of the final image reconstruction. As shown in Figure 1, compared with image restoration agents that rely on scheduling multiple experts, our foundation model achieves higher-quality outputs in a shorter time.

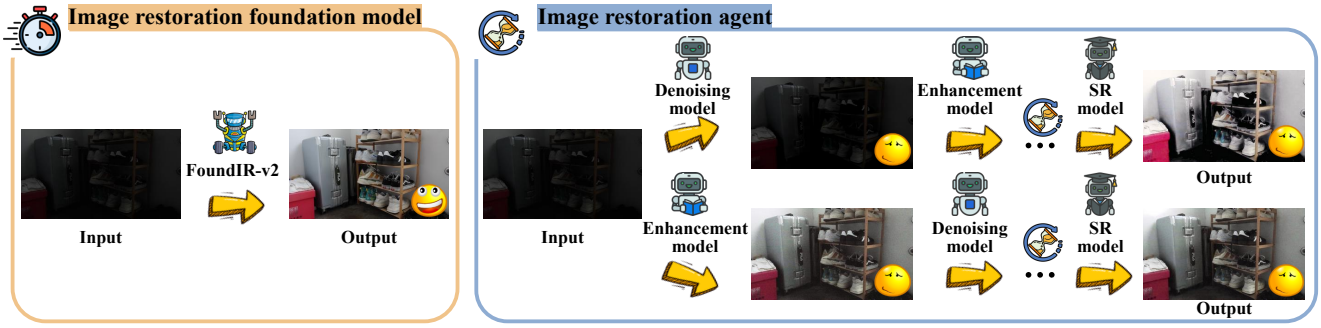

Figure 1. An example illustrating the restoration results produced by image restoration foundation model and image restoration agent.

4. More Experimental Results

Inference cost. We further compare our method with other diffusion-based models in terms of average inference time for 4K images on one H20 GPU. As shown in the Table 2, our method is more efficient. Our MoE-driven scheduler is built upon a lightweight attention-based gating mechanism and thus does not significantly introduces additional inference cost.

Table 2. Comparison of inference time for diffusion-based methods on high-resolution 4K image.

Methods	DreamClear [1]	SUPIR [41]	HYPIR [23]	Ours
Running Time (s)	578.3	467.1	91.8	68.8

Applications on downstream tasks. In fact, a well-trained image restoration foundation model serves as a versatile basis that can be fine-tuned for various downstream restoration tasks related to its pre-training objective. To demonstrate the applicability of image restoration foundation models to downstream tasks, we extend FoundIR [19] and FoundIR-v2 to laparoscopic surgery and biological microscopy image restoration. Since the training dataset in these tasks are typically limited and protected by privacy constraints, we fine-tune the pre-trained foundation models using only using a small amount of available data. The qualitative comparison results are provided in Figure 2. Compared with FoundIR, our FoundIR-v2 achieves better reconstruction results on these downstream tasks, highlighting the significance of developing effective image restoration foundation models.

More visual results. We show more experimental results to demonstrate the effectiveness of the proposed image restoration foundation model FoundIR-v2. Figures 3-51 show the visual comparison results on 50 restoration tasks across broader real-world scenarios. Table 3 summarizes the task categories addressed by our proposed FoundIR-v2. Compared with the 20 restoration tasks covered by FoundIR, our proposed FoundIR-v2 enables a broader range of application tasks. It can be observed that our method can handle various degradations and generate much clearer images with finer details and structures.

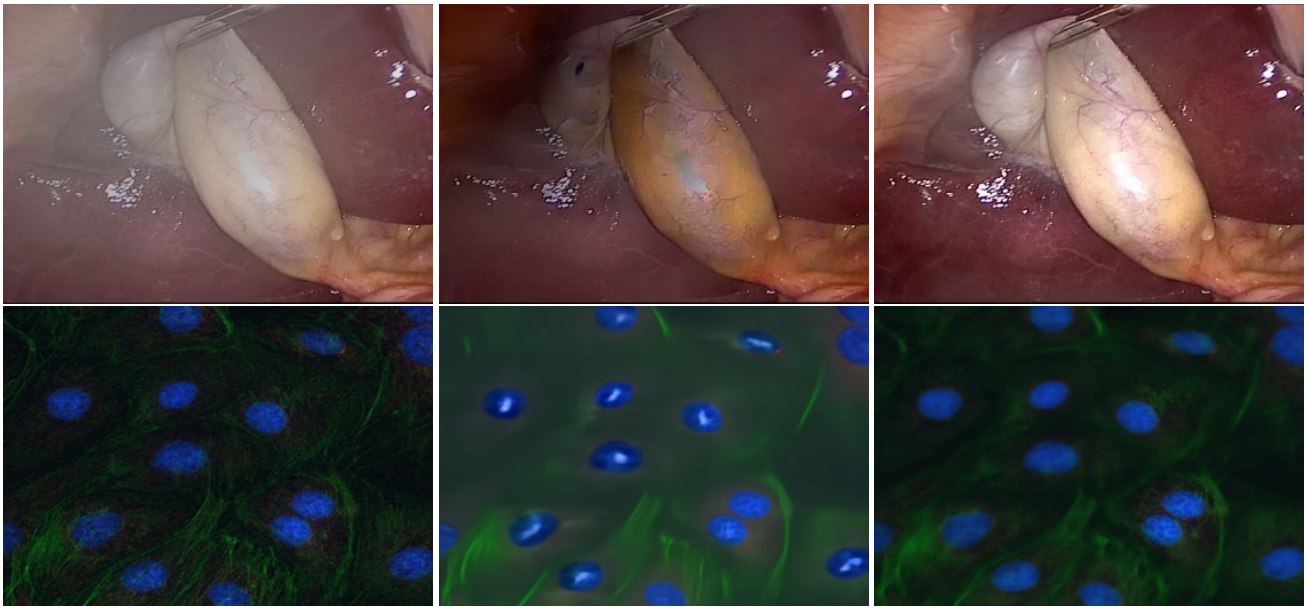

(a) LQ (b) FoundIR [19] (c) FoundIR-v2 (Ours)

Figure 2. Visual comparison of fine-tuning the image restoration foundation models on downstream restoration tasks (laparoscopic surgery and biological microscopy image restoration).

Table 3. Summary of image restoration task categories handled by FoundIR-v2.

Task Categories	Tasks		
Deblur and Denoise	1. Camera Motion Deblurring	2. Object Motion Deblurring	3. Defocus Deblurring
	4. UHD Deblurring	5. Lowlight Deblurring	6. Text Deblurring
	7. Denoising	8. Joint Deblurring and Denoising	
Weather Removal	9. Dehazing	10. Colored Dehazing	11. Non-Homogeneous Dehazing
	12. Lowlight Dehazing	13. Remote Sensing Image Dehazing	14. Remote Sensing Image Declouding
	15. Rain Streak Removal	16. Raindrop Removal	17. Rain Mist Removal
	18. Lowlight Deraining	19. Desnowing	20. Lowlight Desnowing
	21. All-in-One Weather Removal		
Lowlight Enhancement	22. Contrast Enhancement	23. Lowlight Enhancement	24. Nighttime Enhancement
	25. Extreme Dark Enhancement	26. UHD Enhancement	27. Joint Denoising and Enhancement
Super-Resolution	28. Super-Resolution ($\times 1$)	29. Super-Resolution ($\times 2$)	30. Super-Resolution ($\times 4$)
	31. Text Super-Resolution	32. Deblurring and SR	33. Denoising and SR
	34. DeJPEG and SR	35. Inpainting	36. JPEG Compression Removal
	37. Artifact Removal	38. Stain Removal	39. Scratch Removal
Real-world Restoration	40. Underwater Restoration	41. UAV Image Restoration	42. Old Photo Restoration
	43. Face Restoration	44. Classic Film Restoration	45. Social Media Enhancement
	46. Anime Image Enhancement	47. Cultural Relic Image Restoration	48. Crack Restoration
	49. Mural Restoration	50. Ancient Text Restoration	

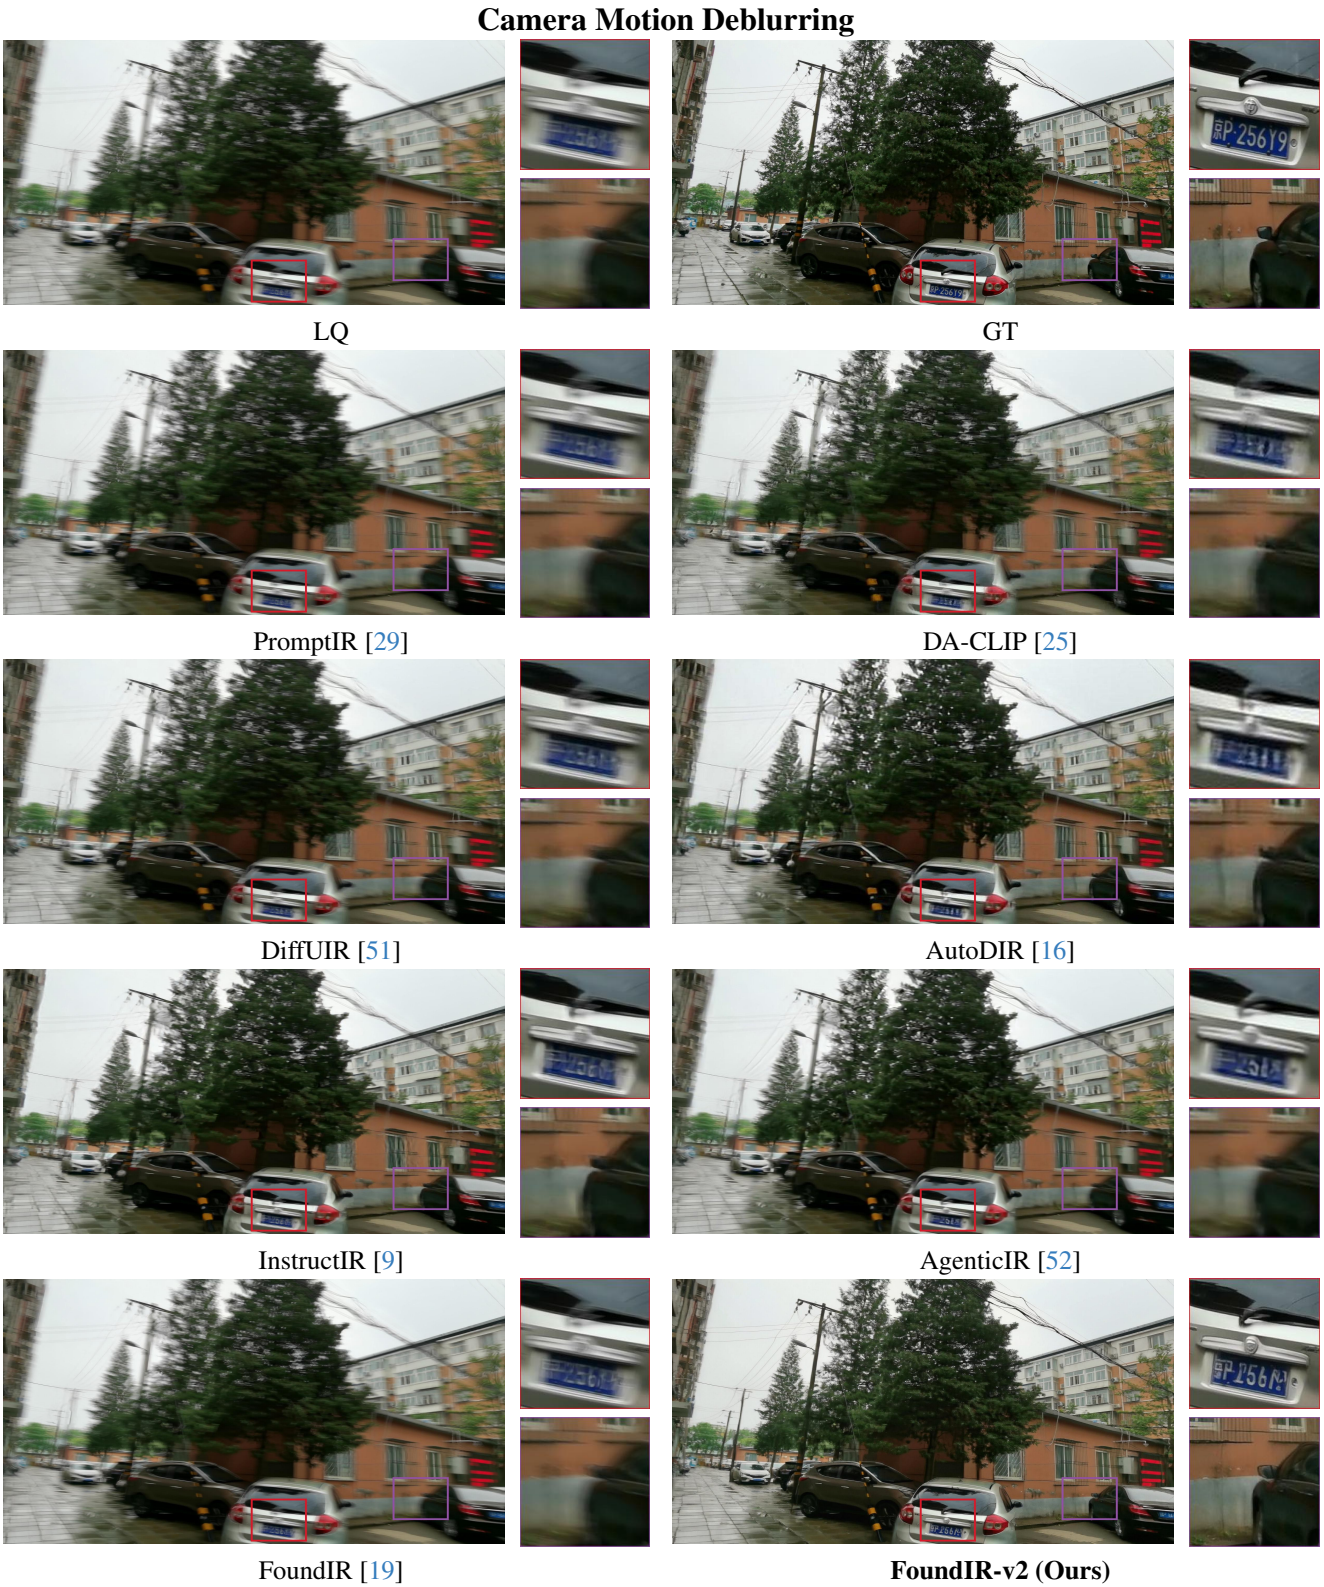

Figure 3. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

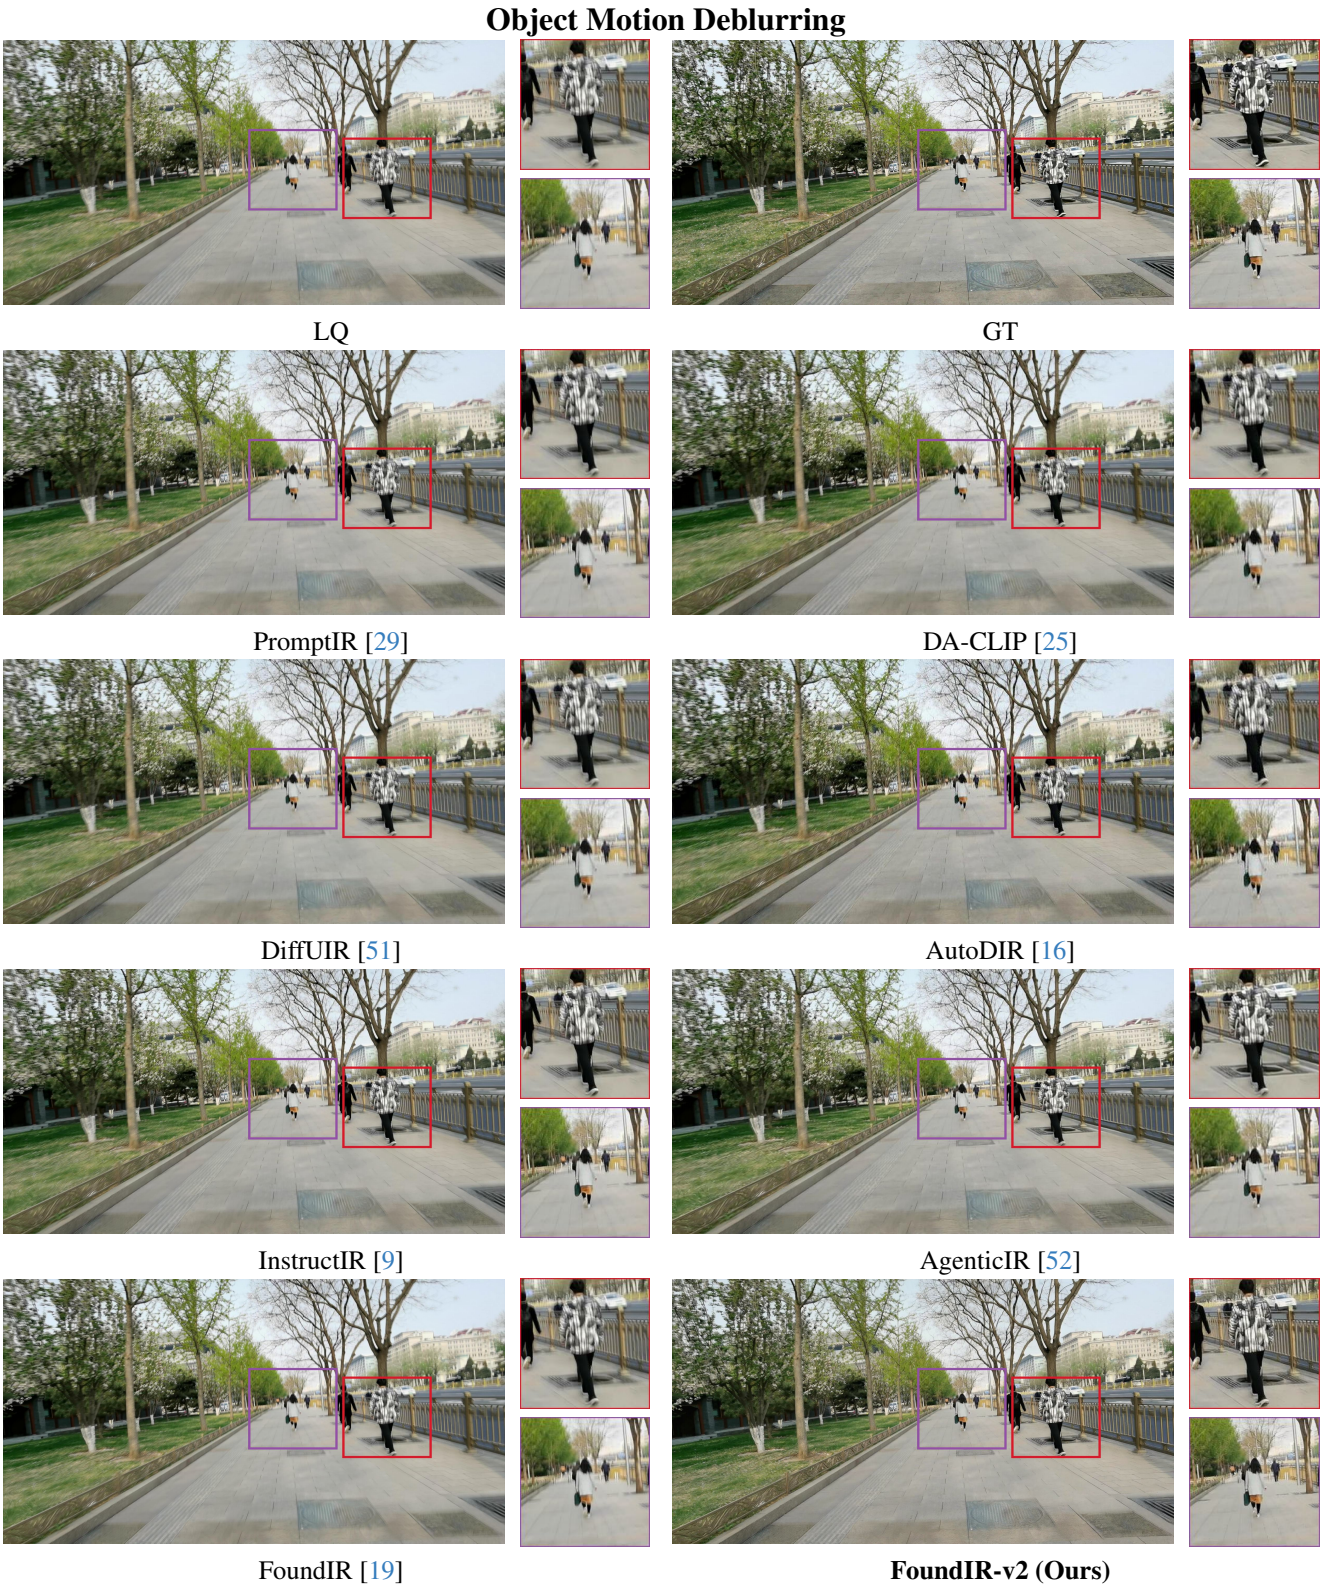

Figure 4. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

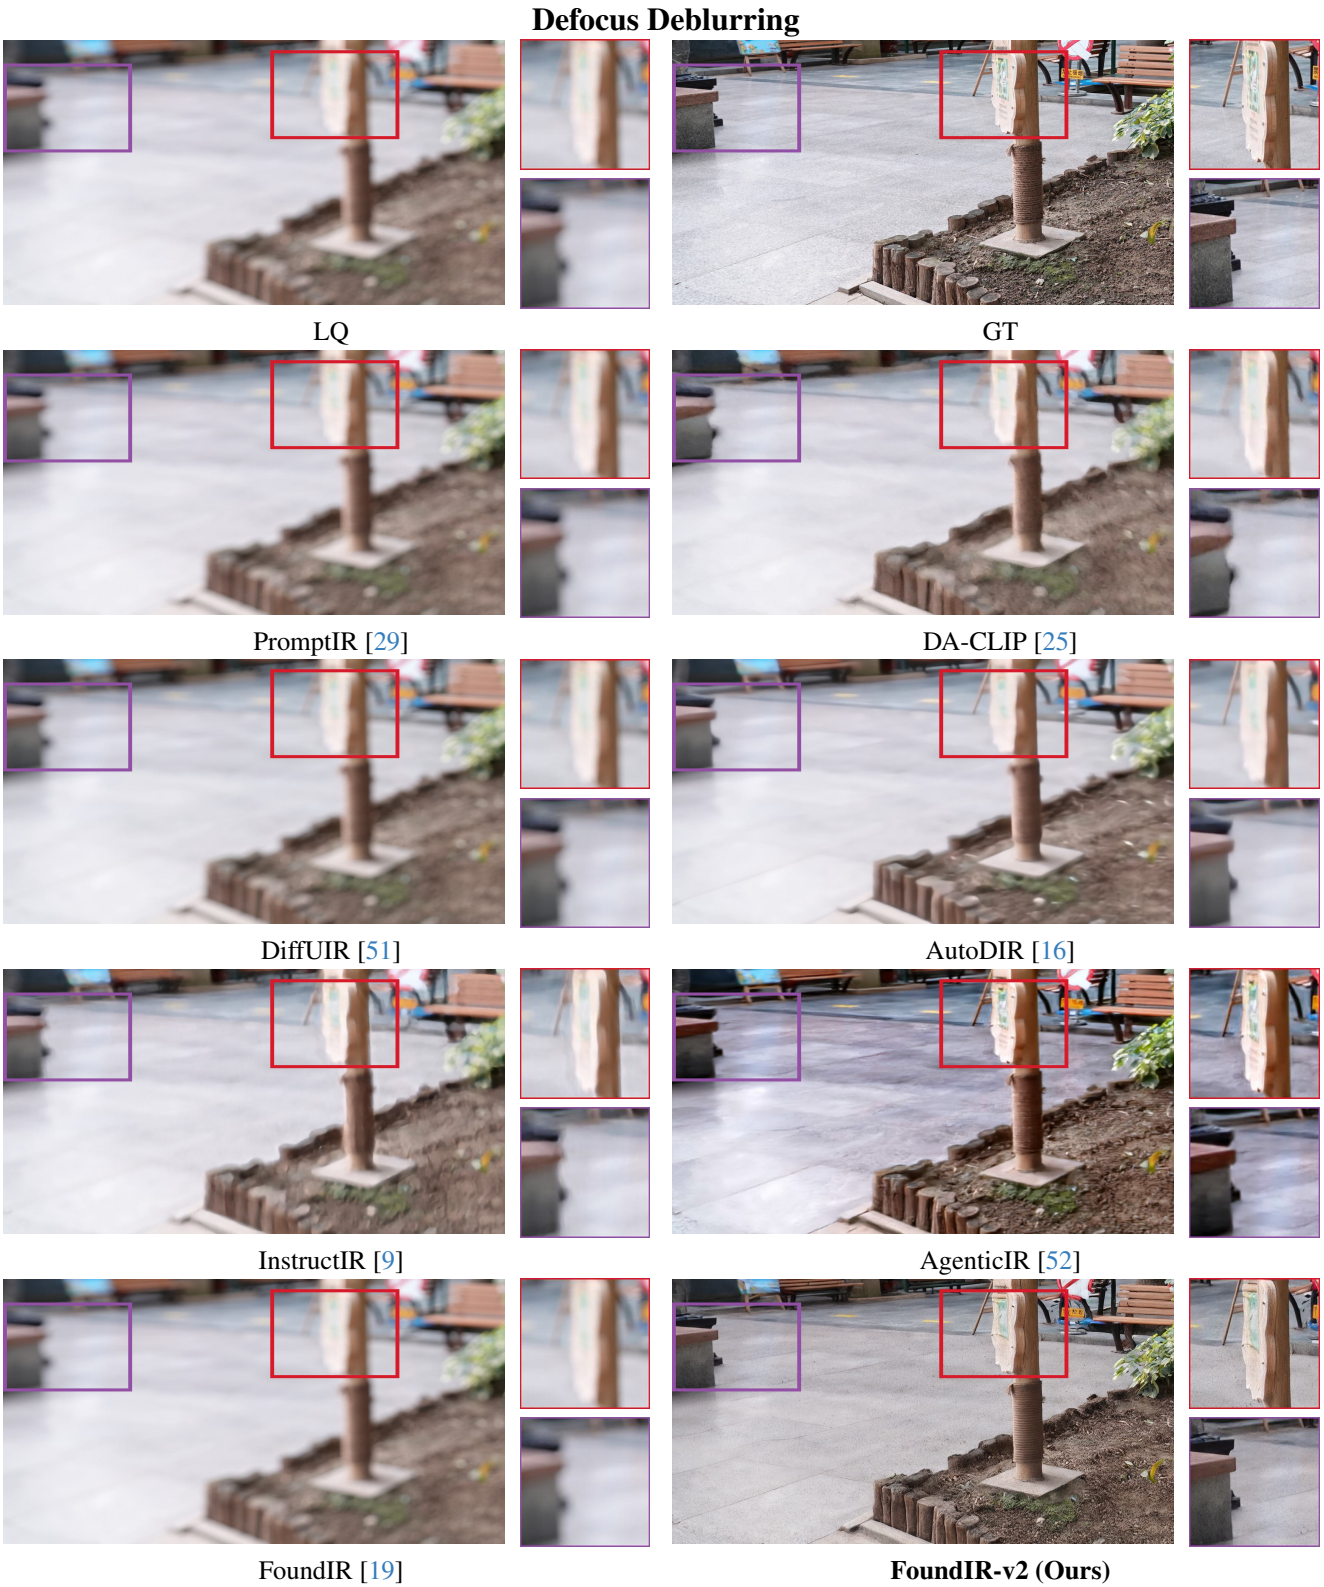

Figure 5. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 6. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 7. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 8. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

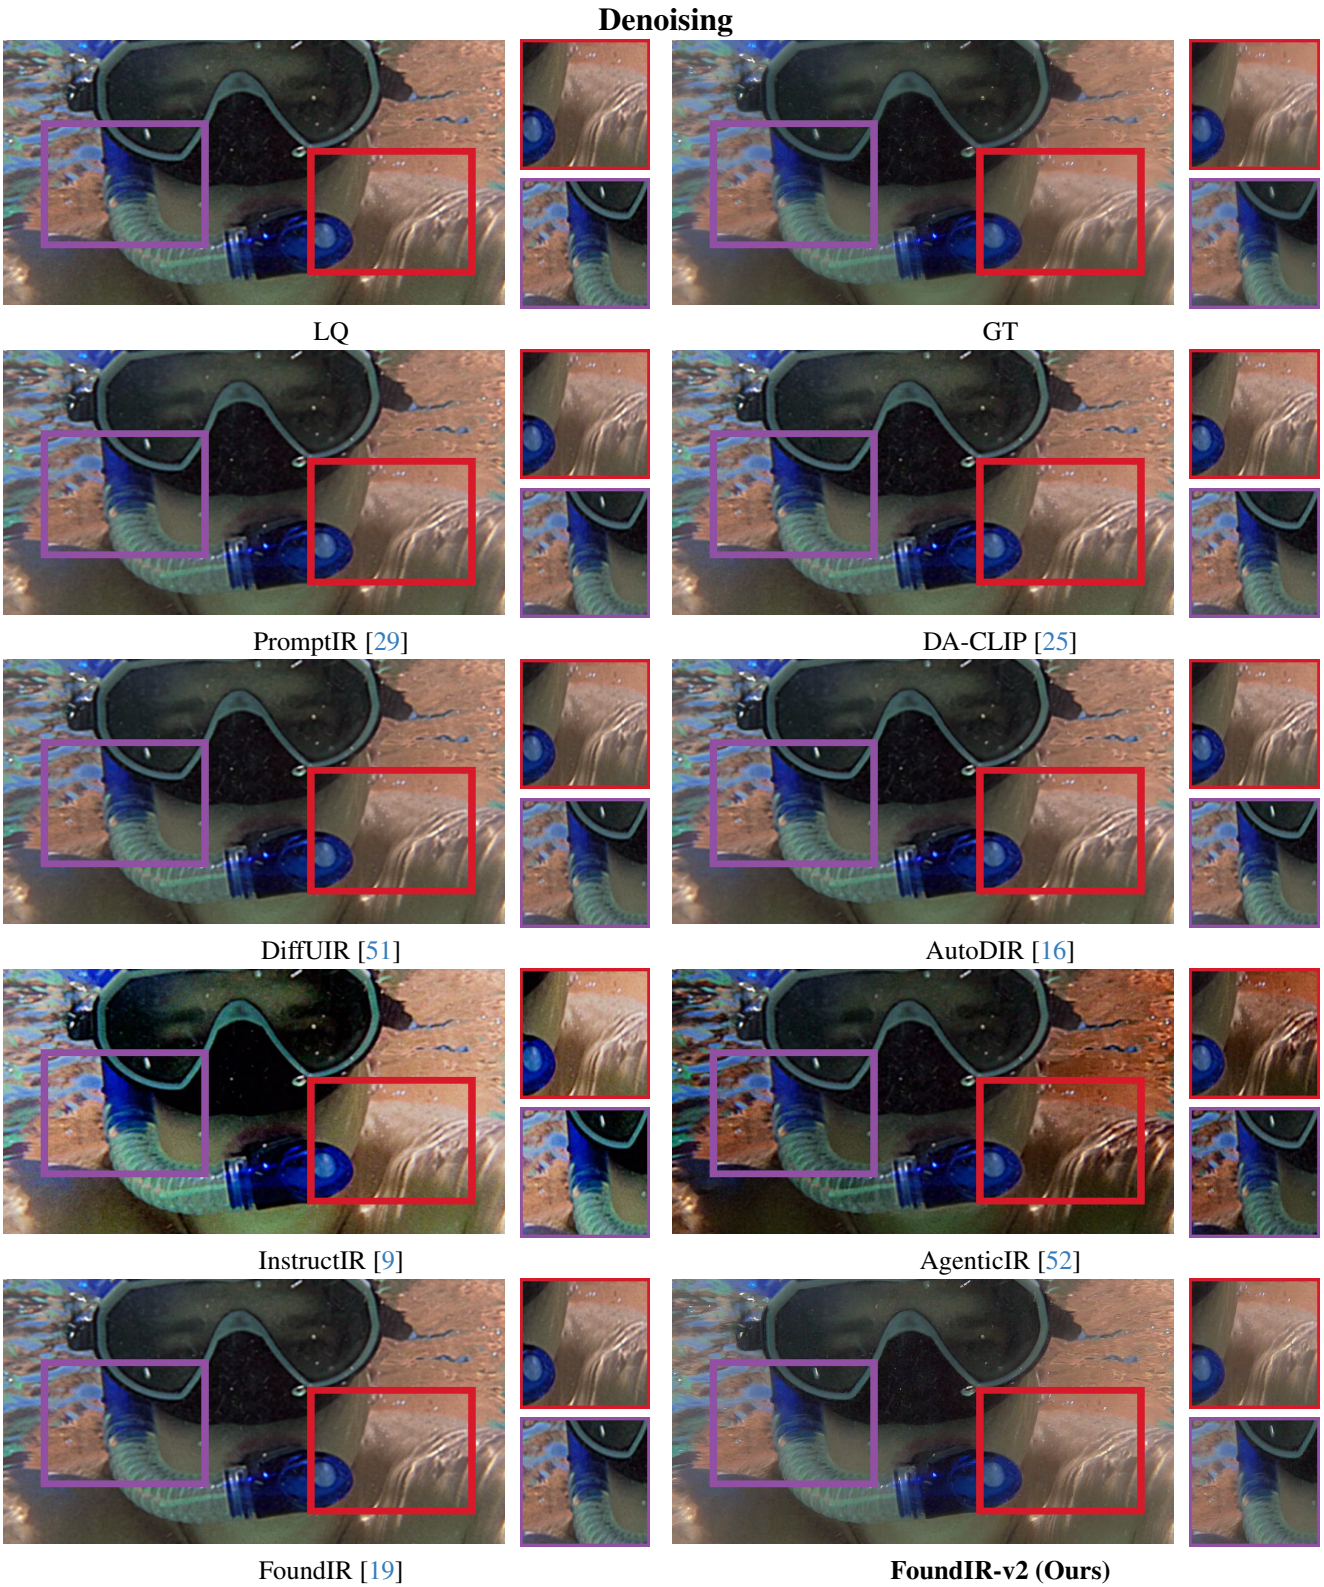

Figure 9. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

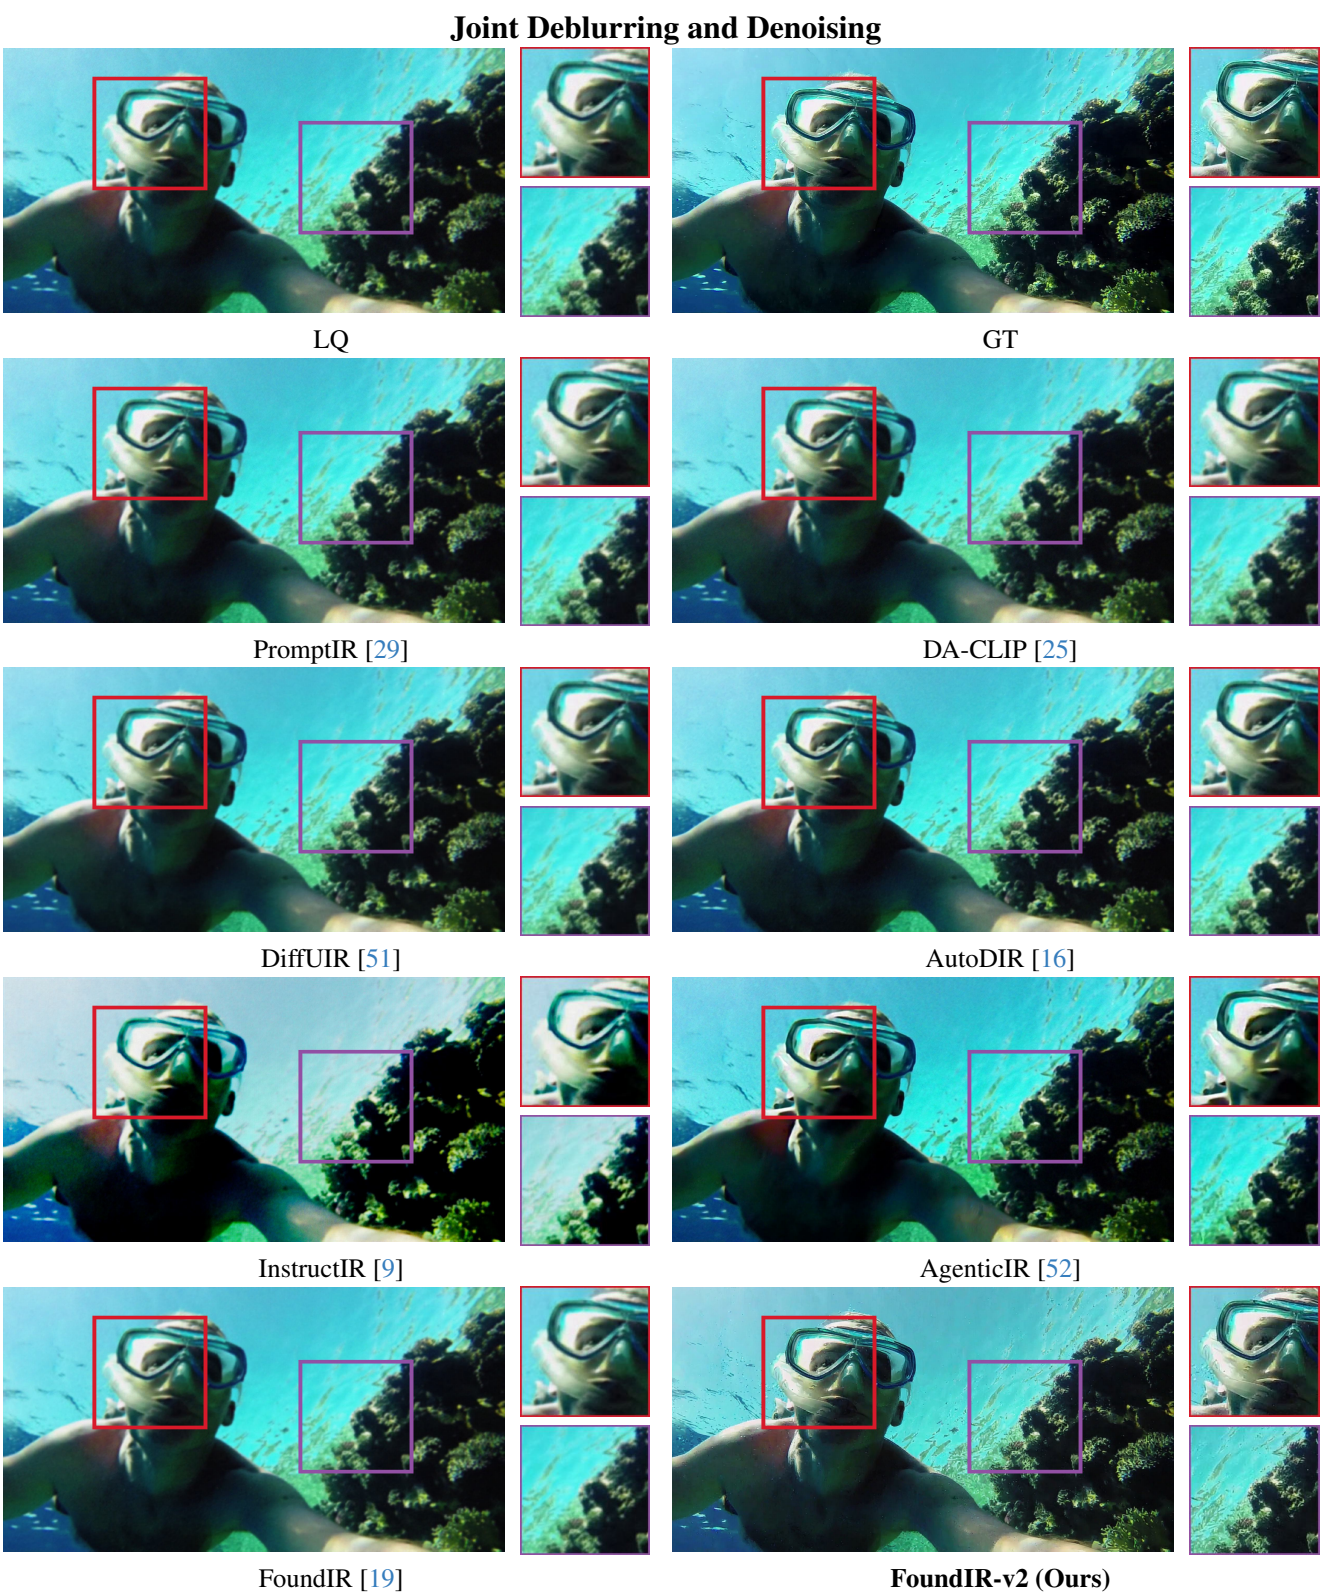

Figure 10. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

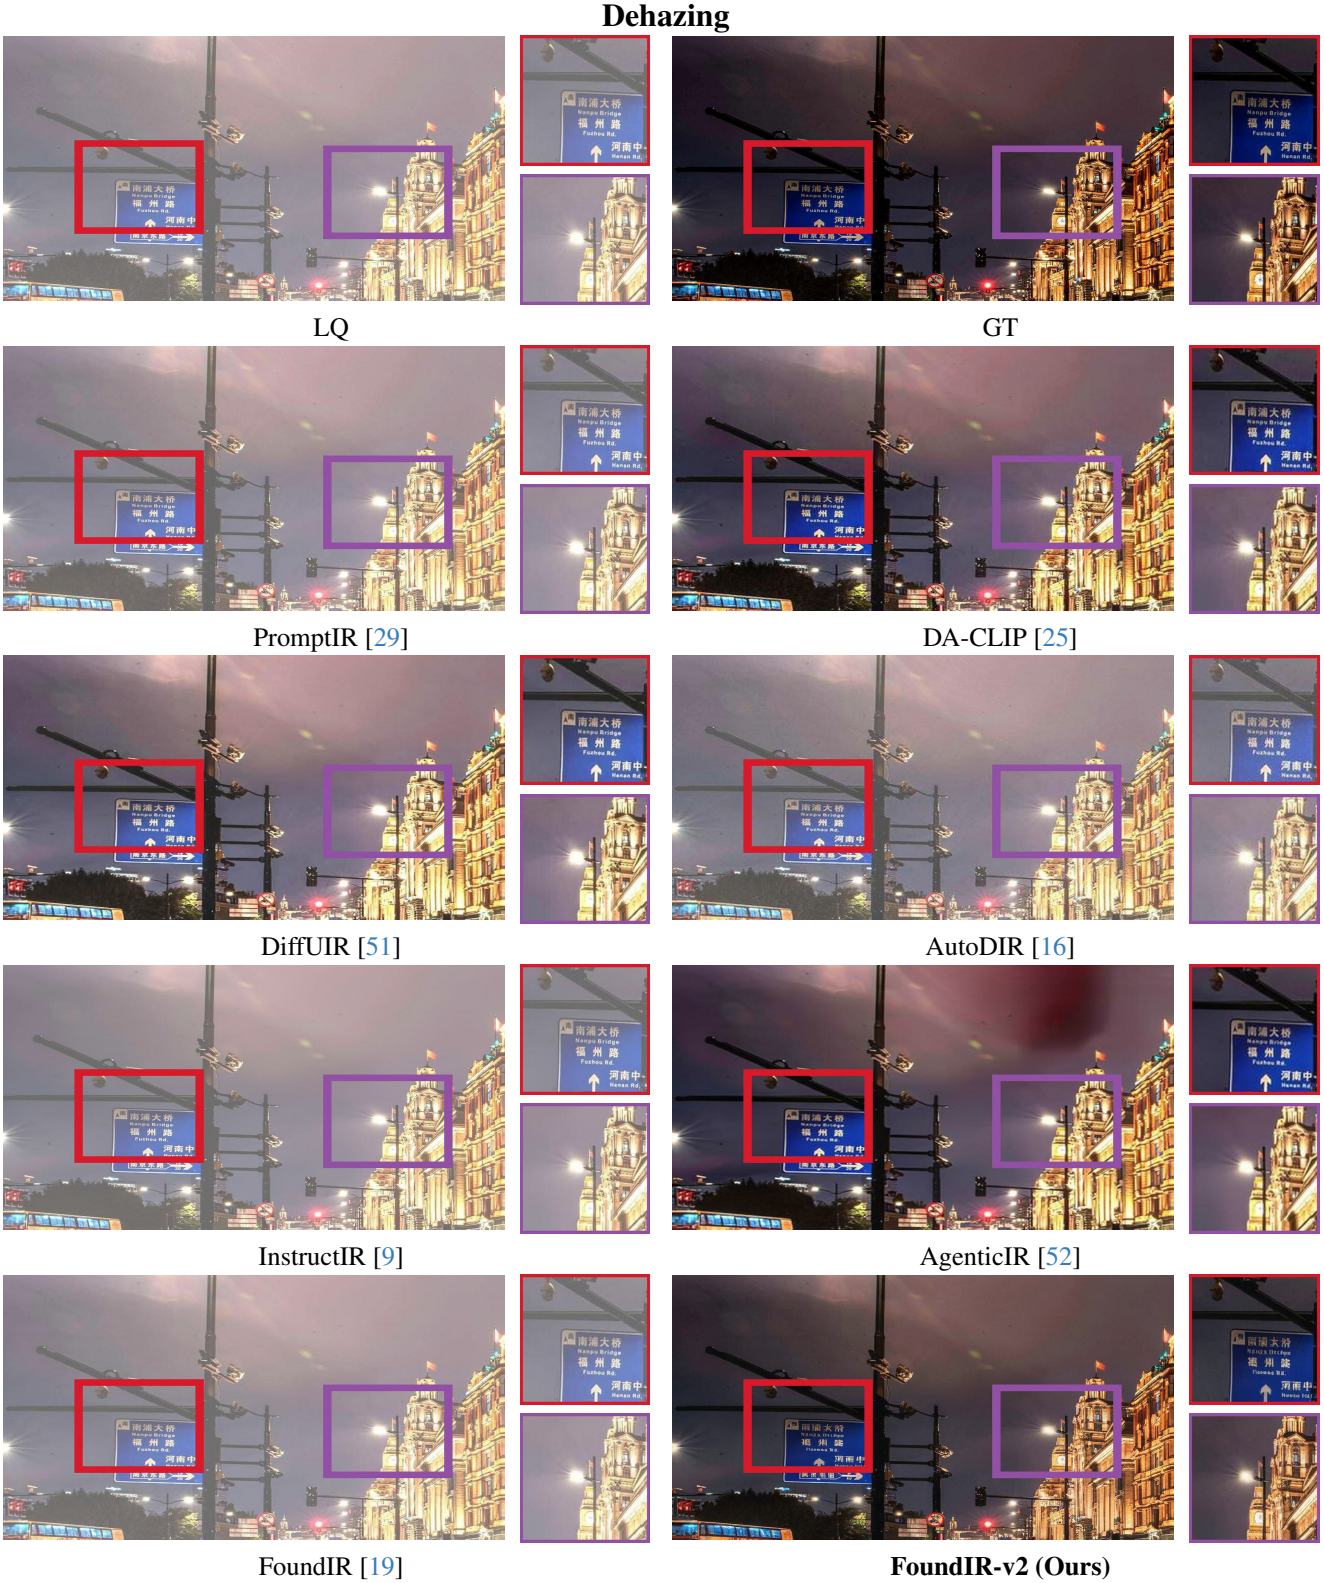

Figure 11. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 12. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 13. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

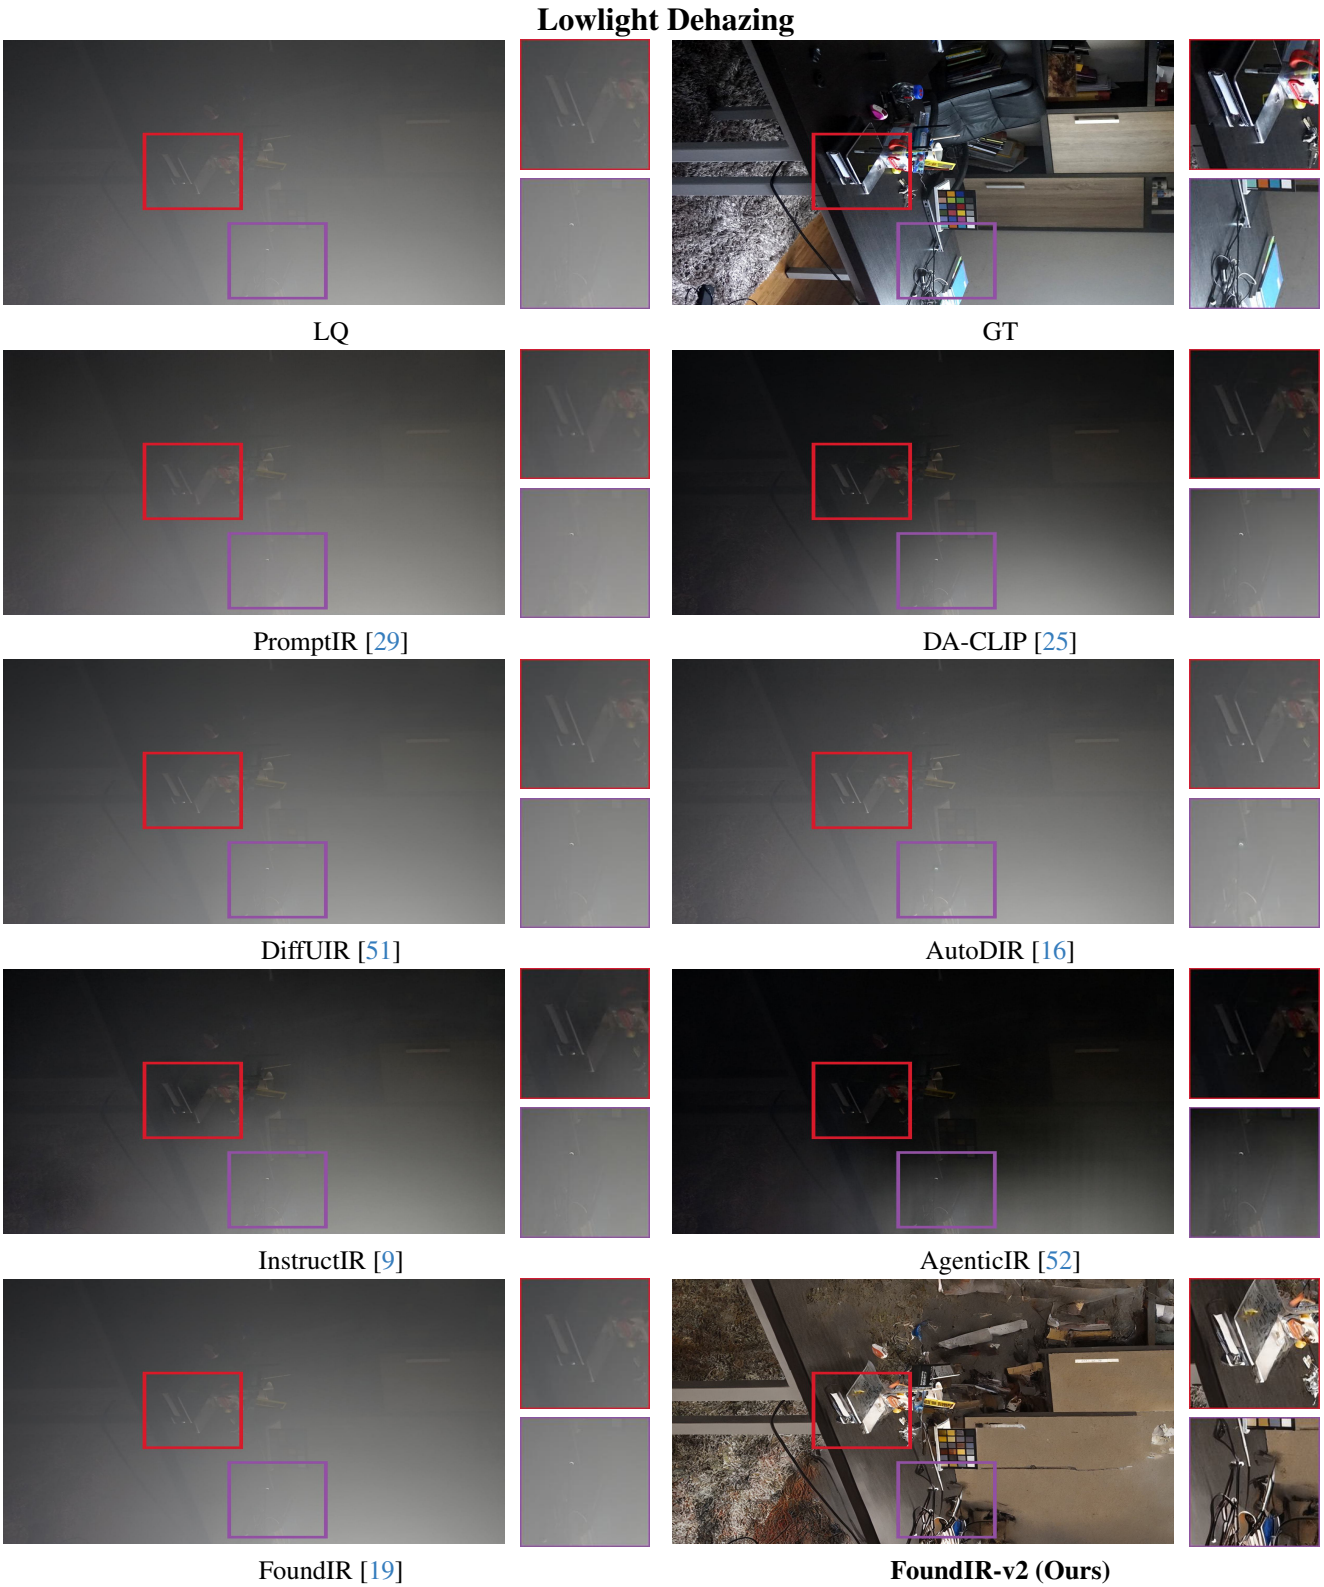

Figure 14. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

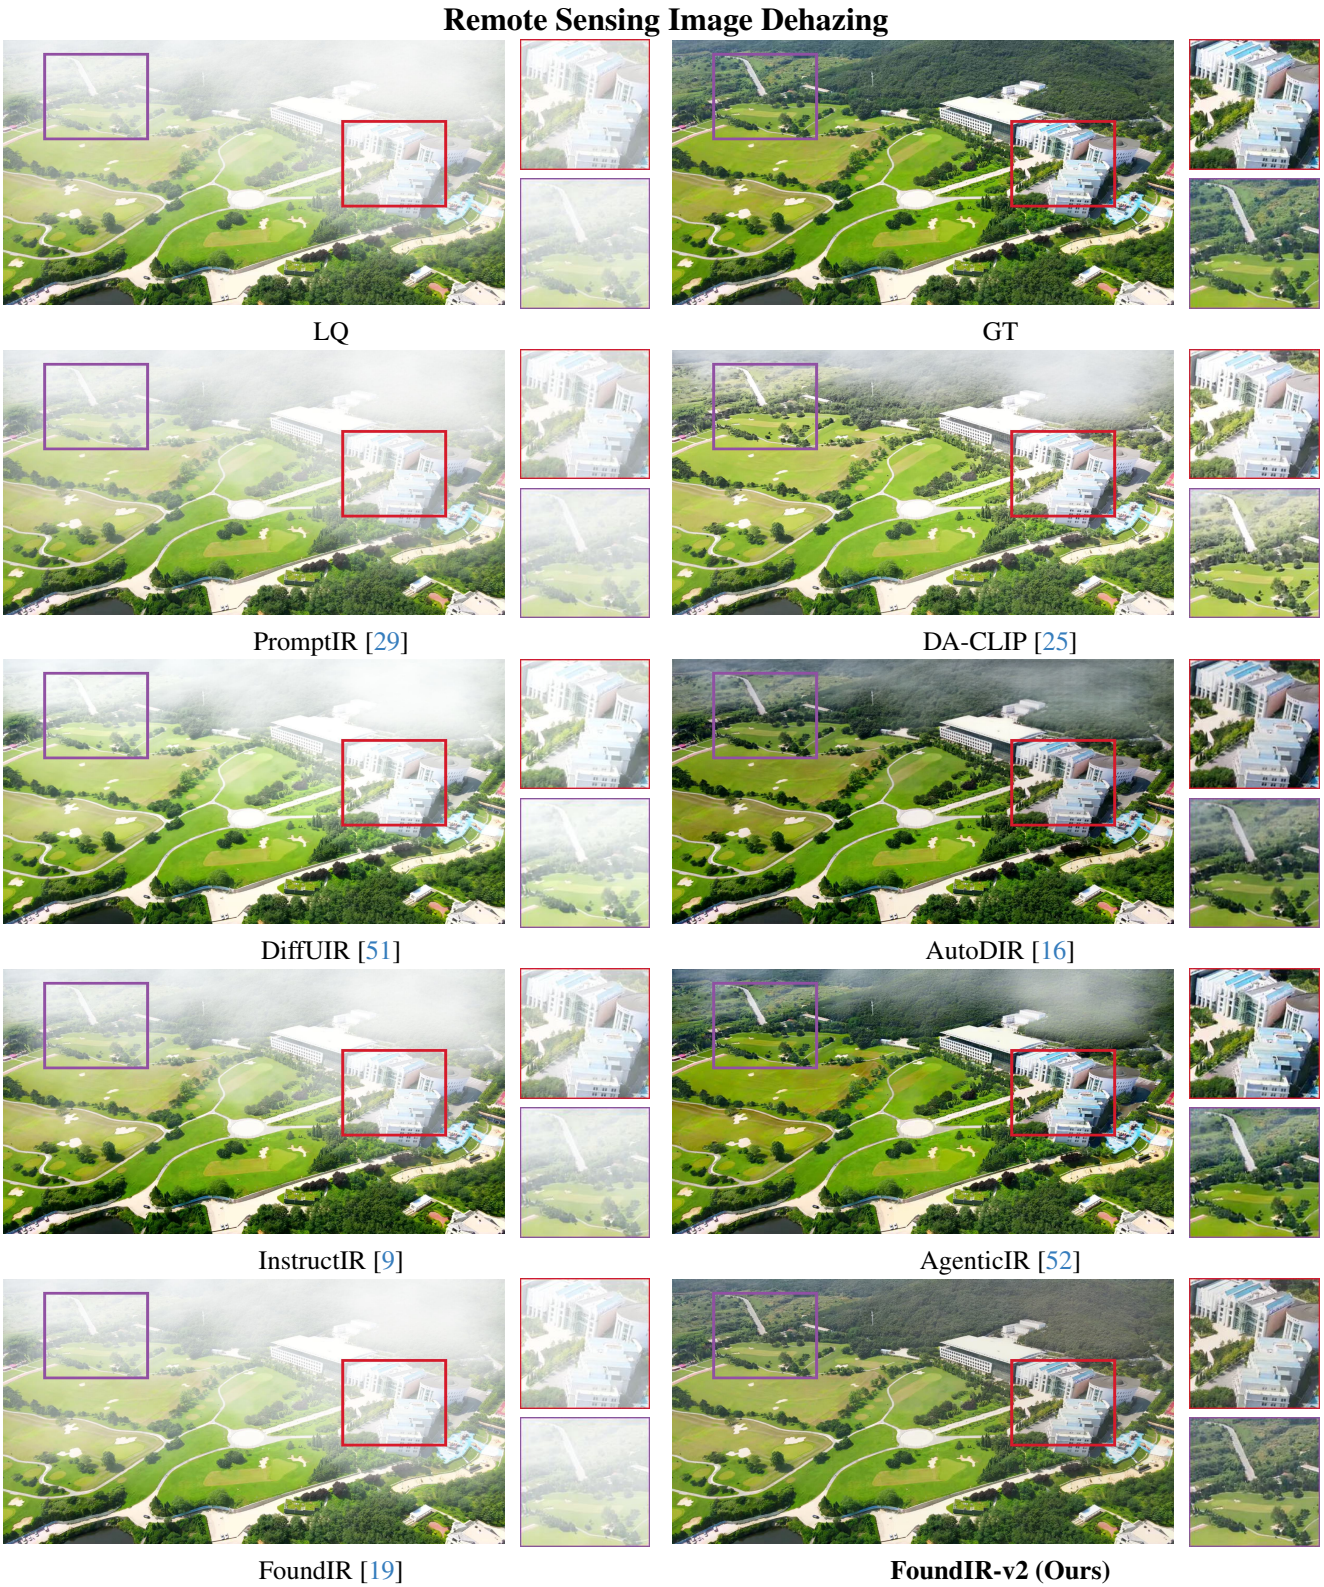

Figure 15. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

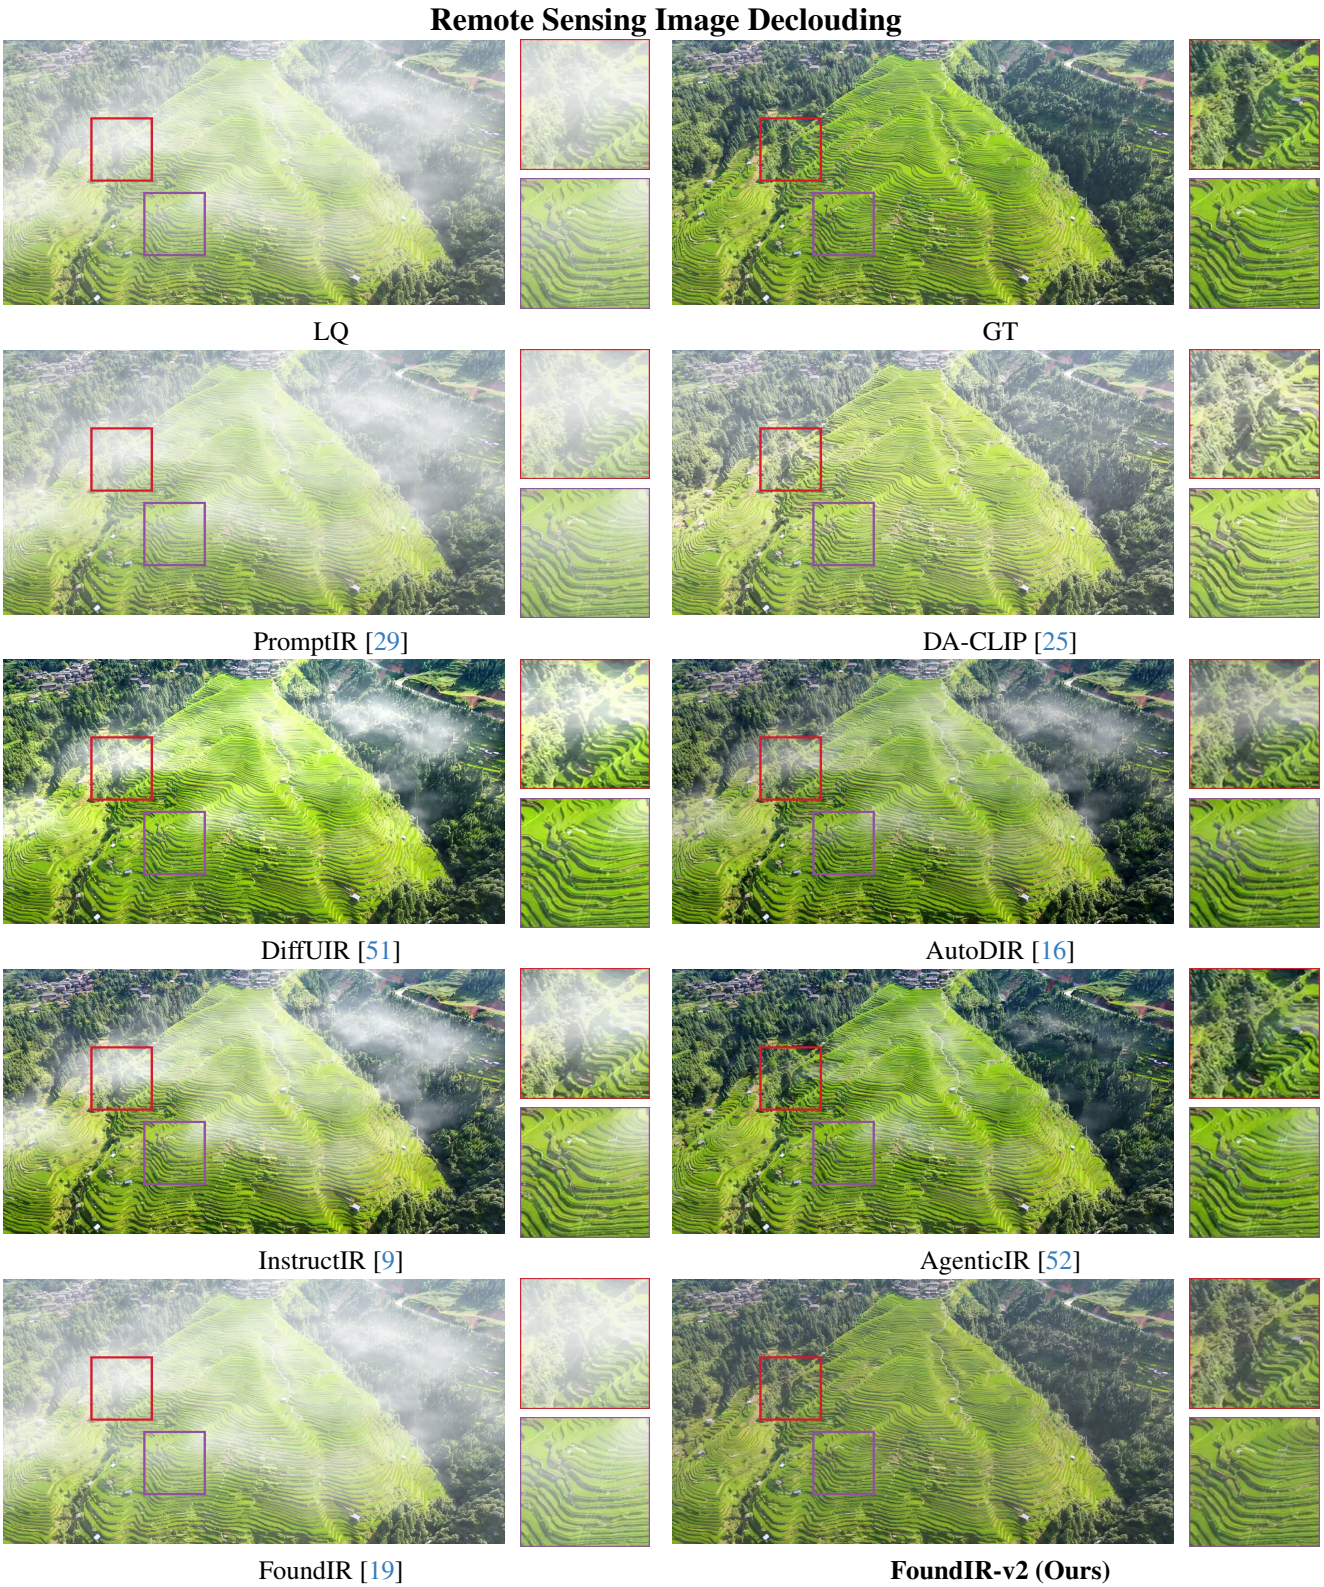

Figure 16. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 17. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

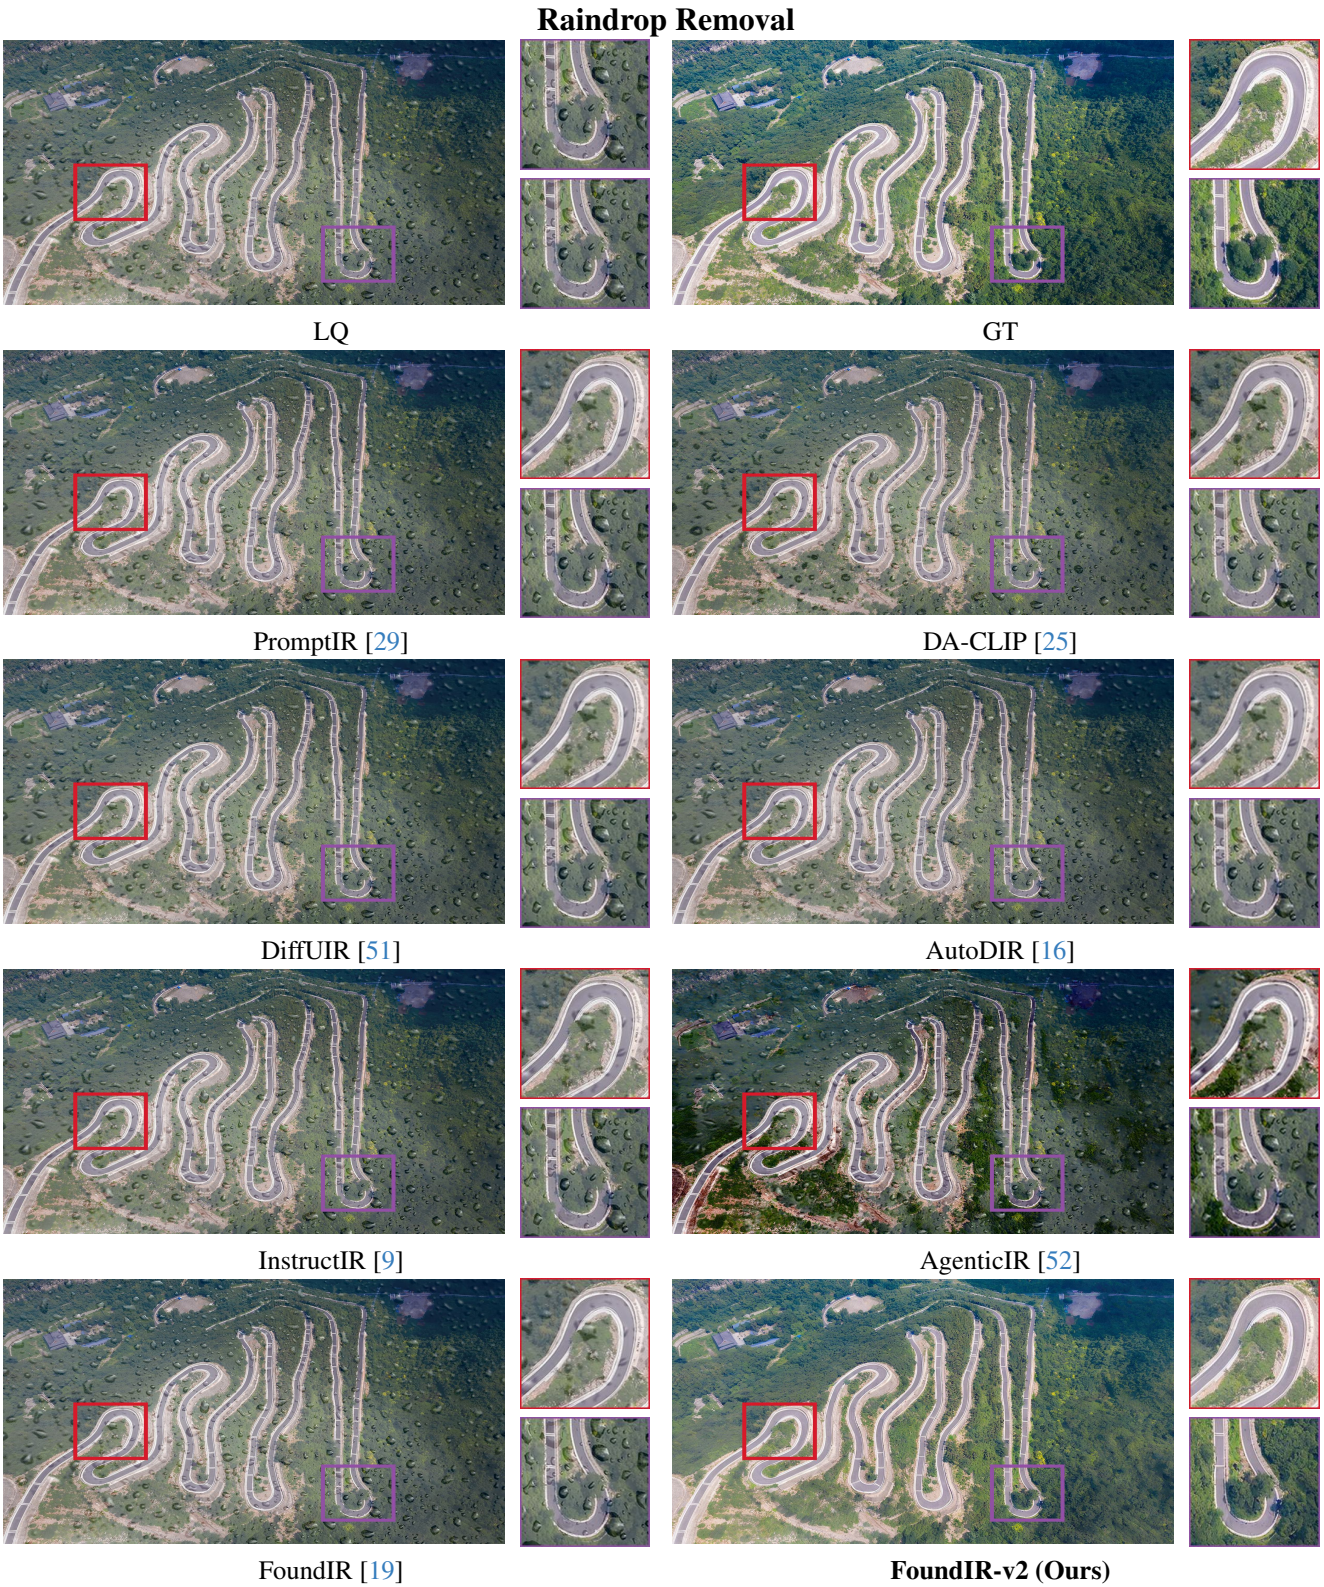

Figure 18. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 19. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

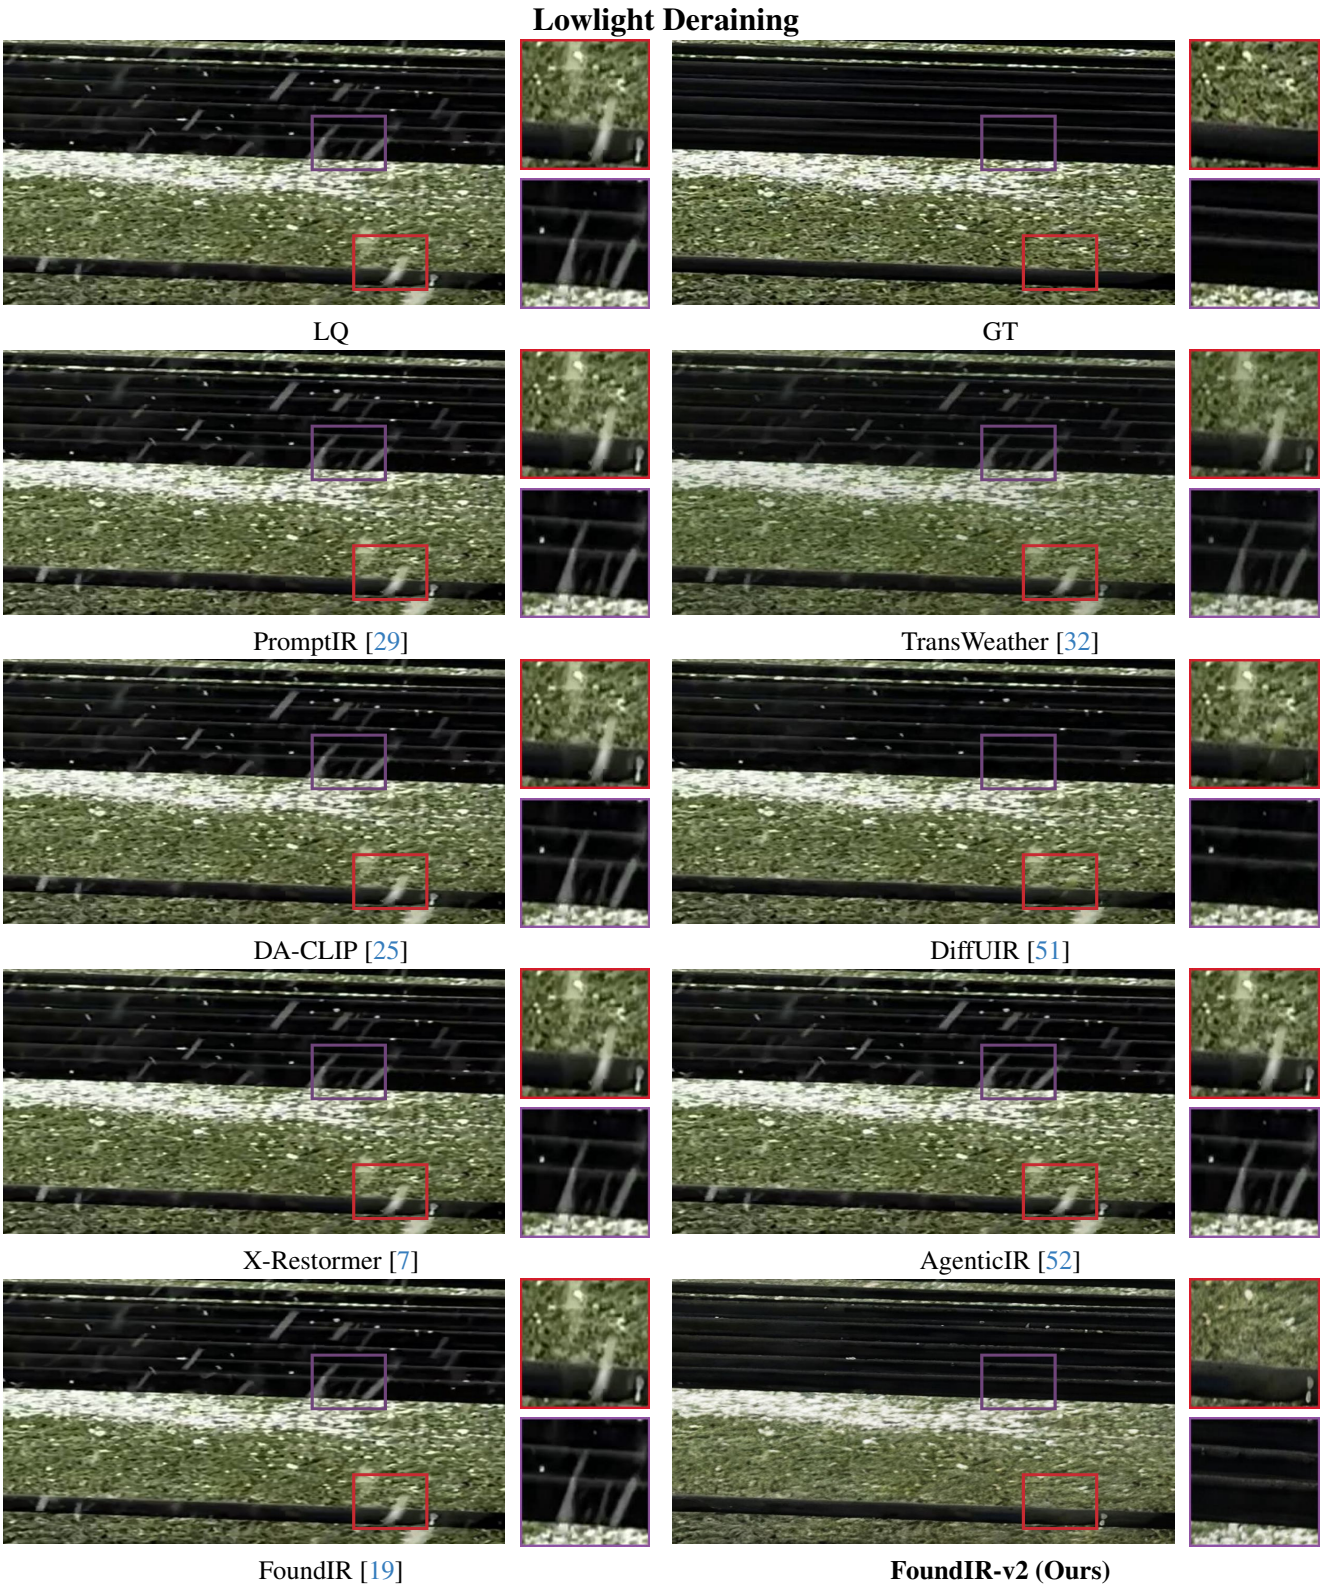

Figure 20. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

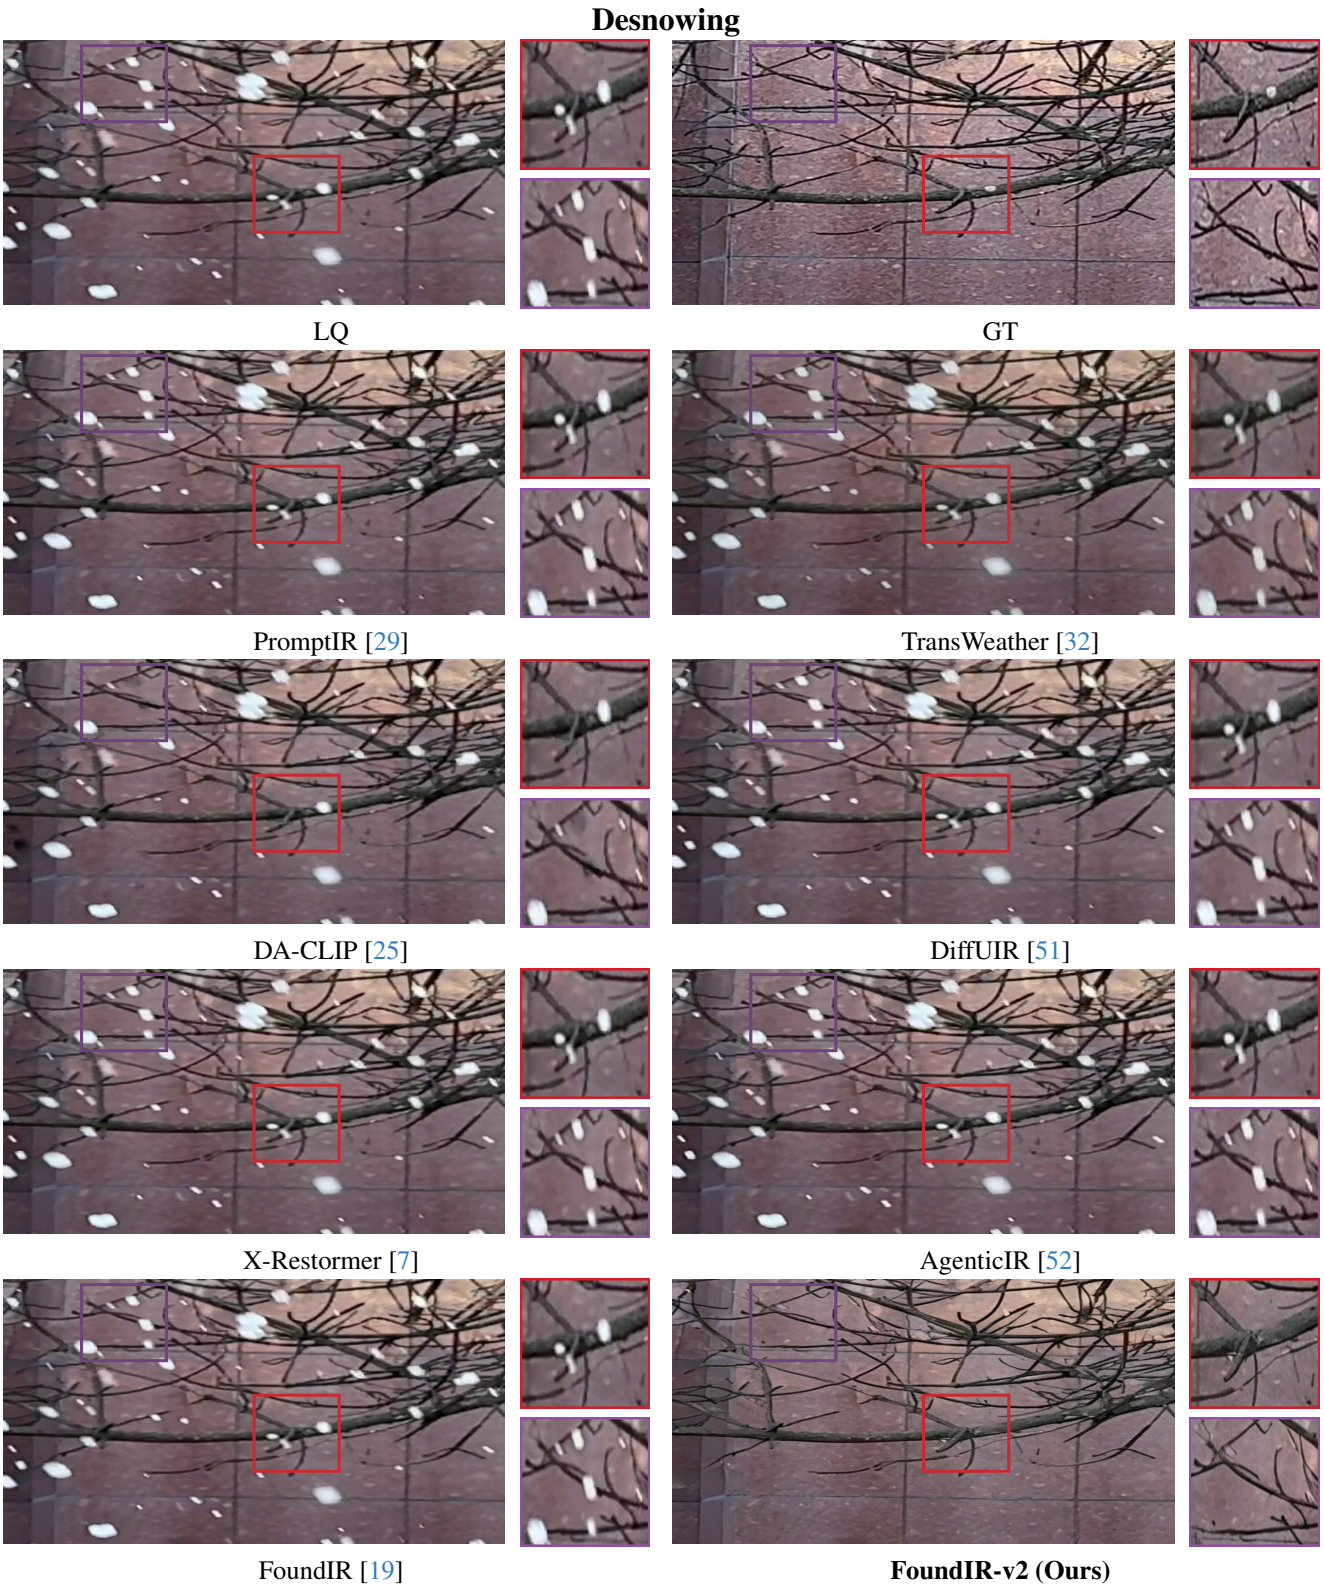

Figure 21. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

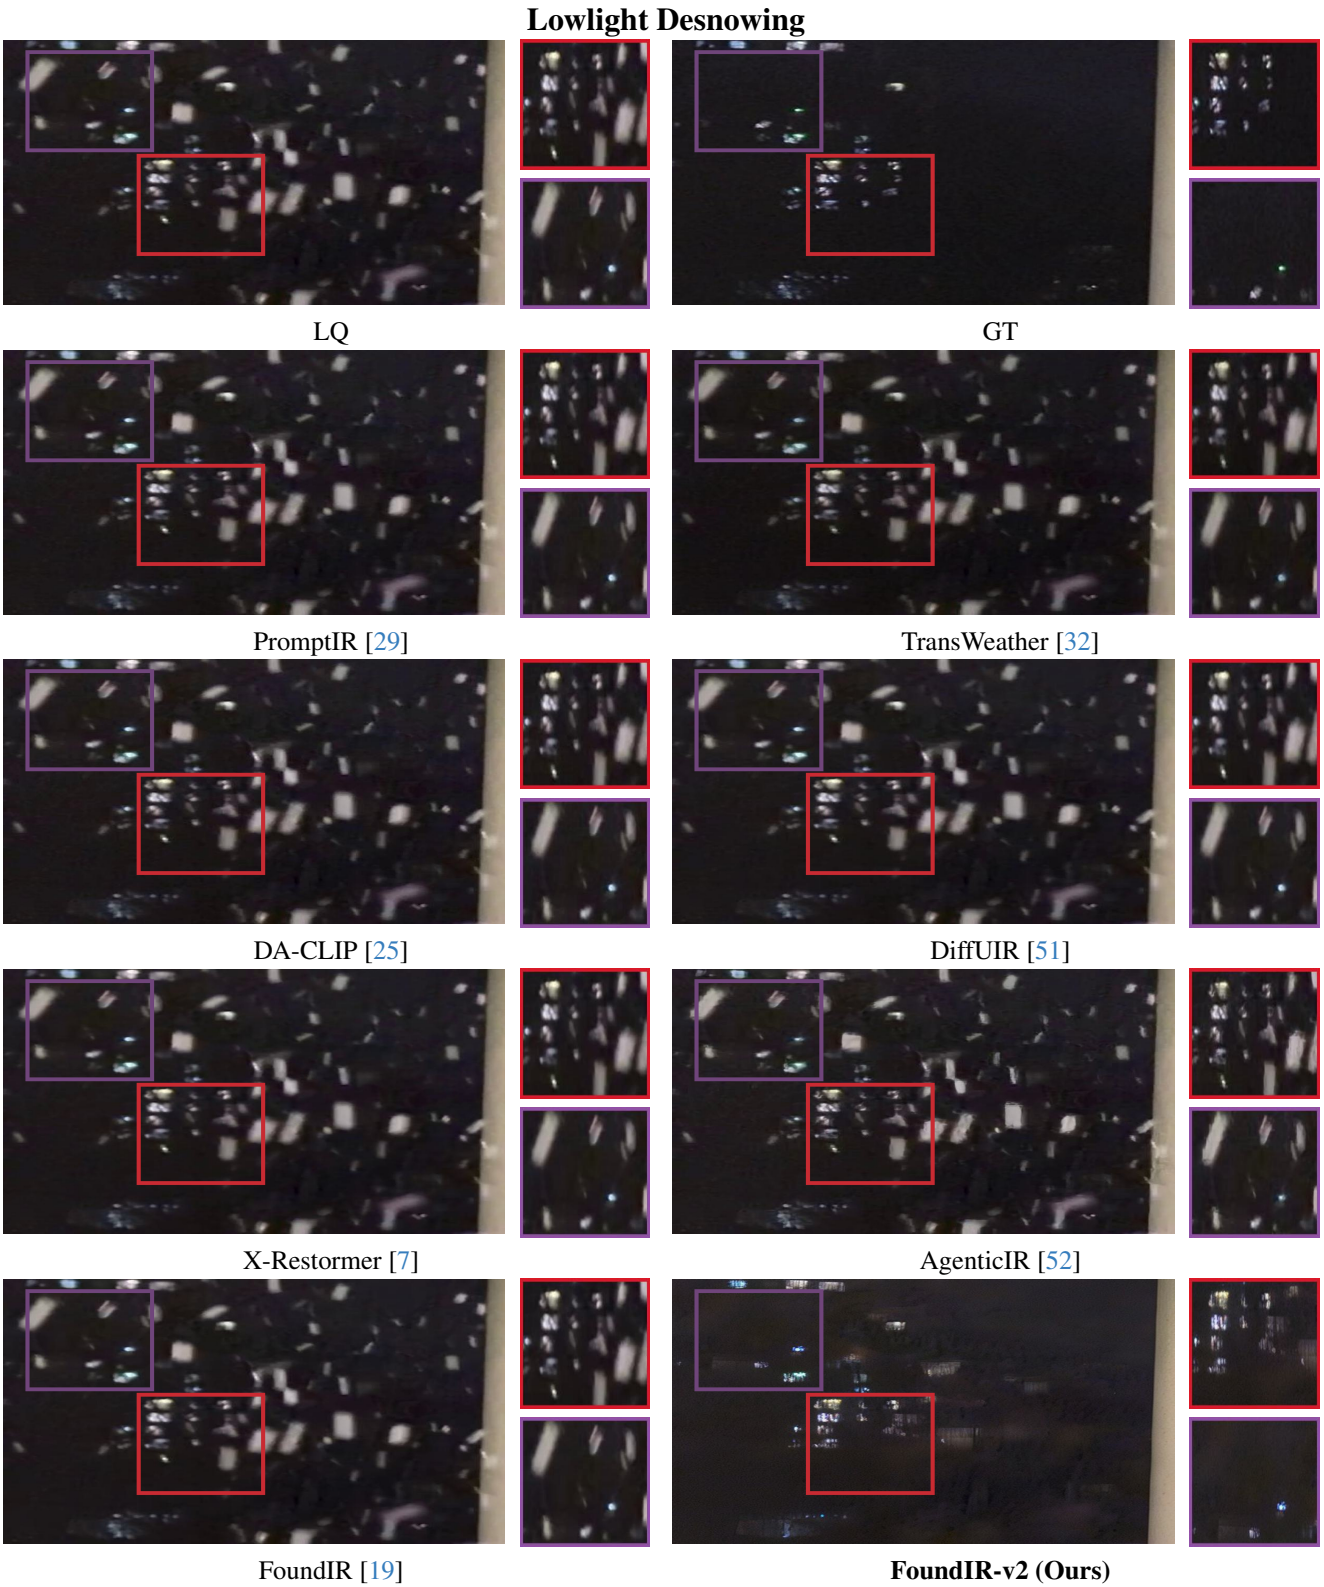

Figure 22. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 23. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 24. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 25. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 26. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

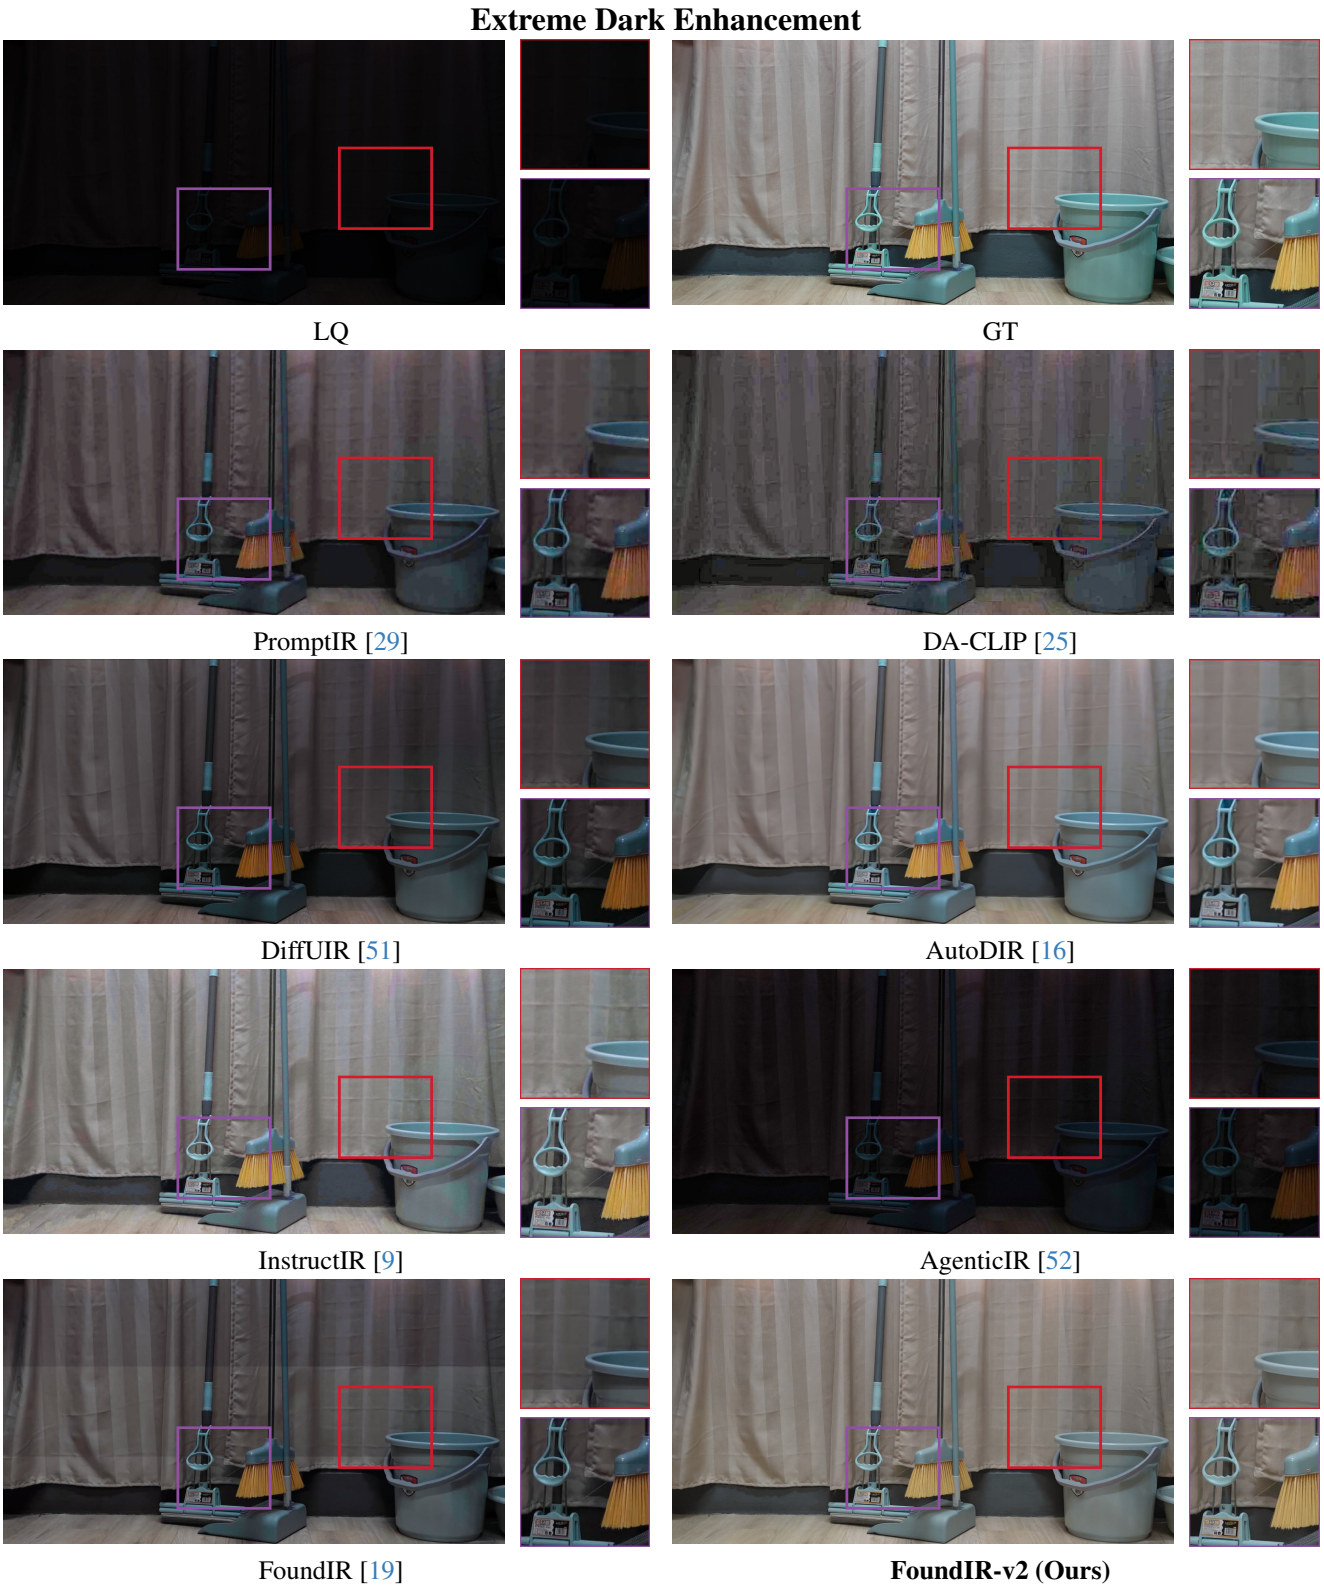

Figure 27. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 28. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 29. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

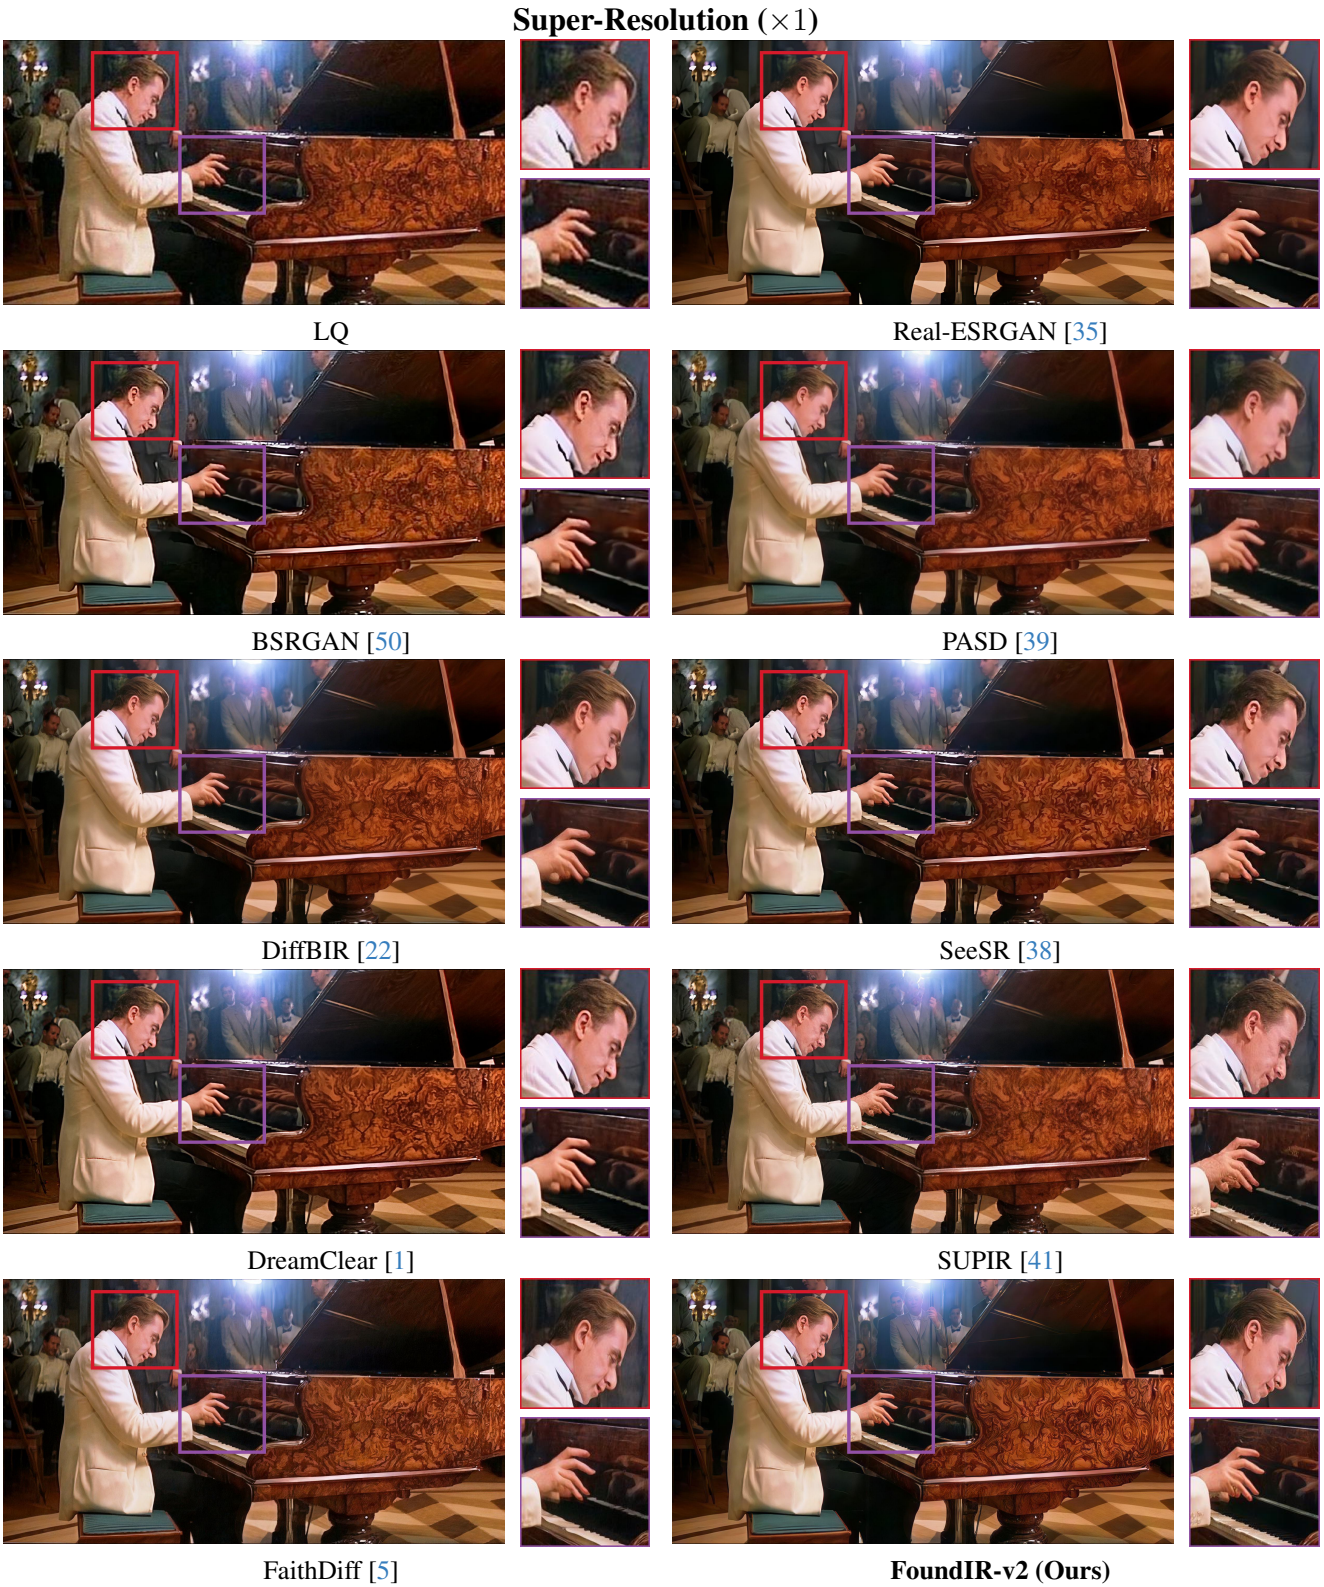

Figure 30. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

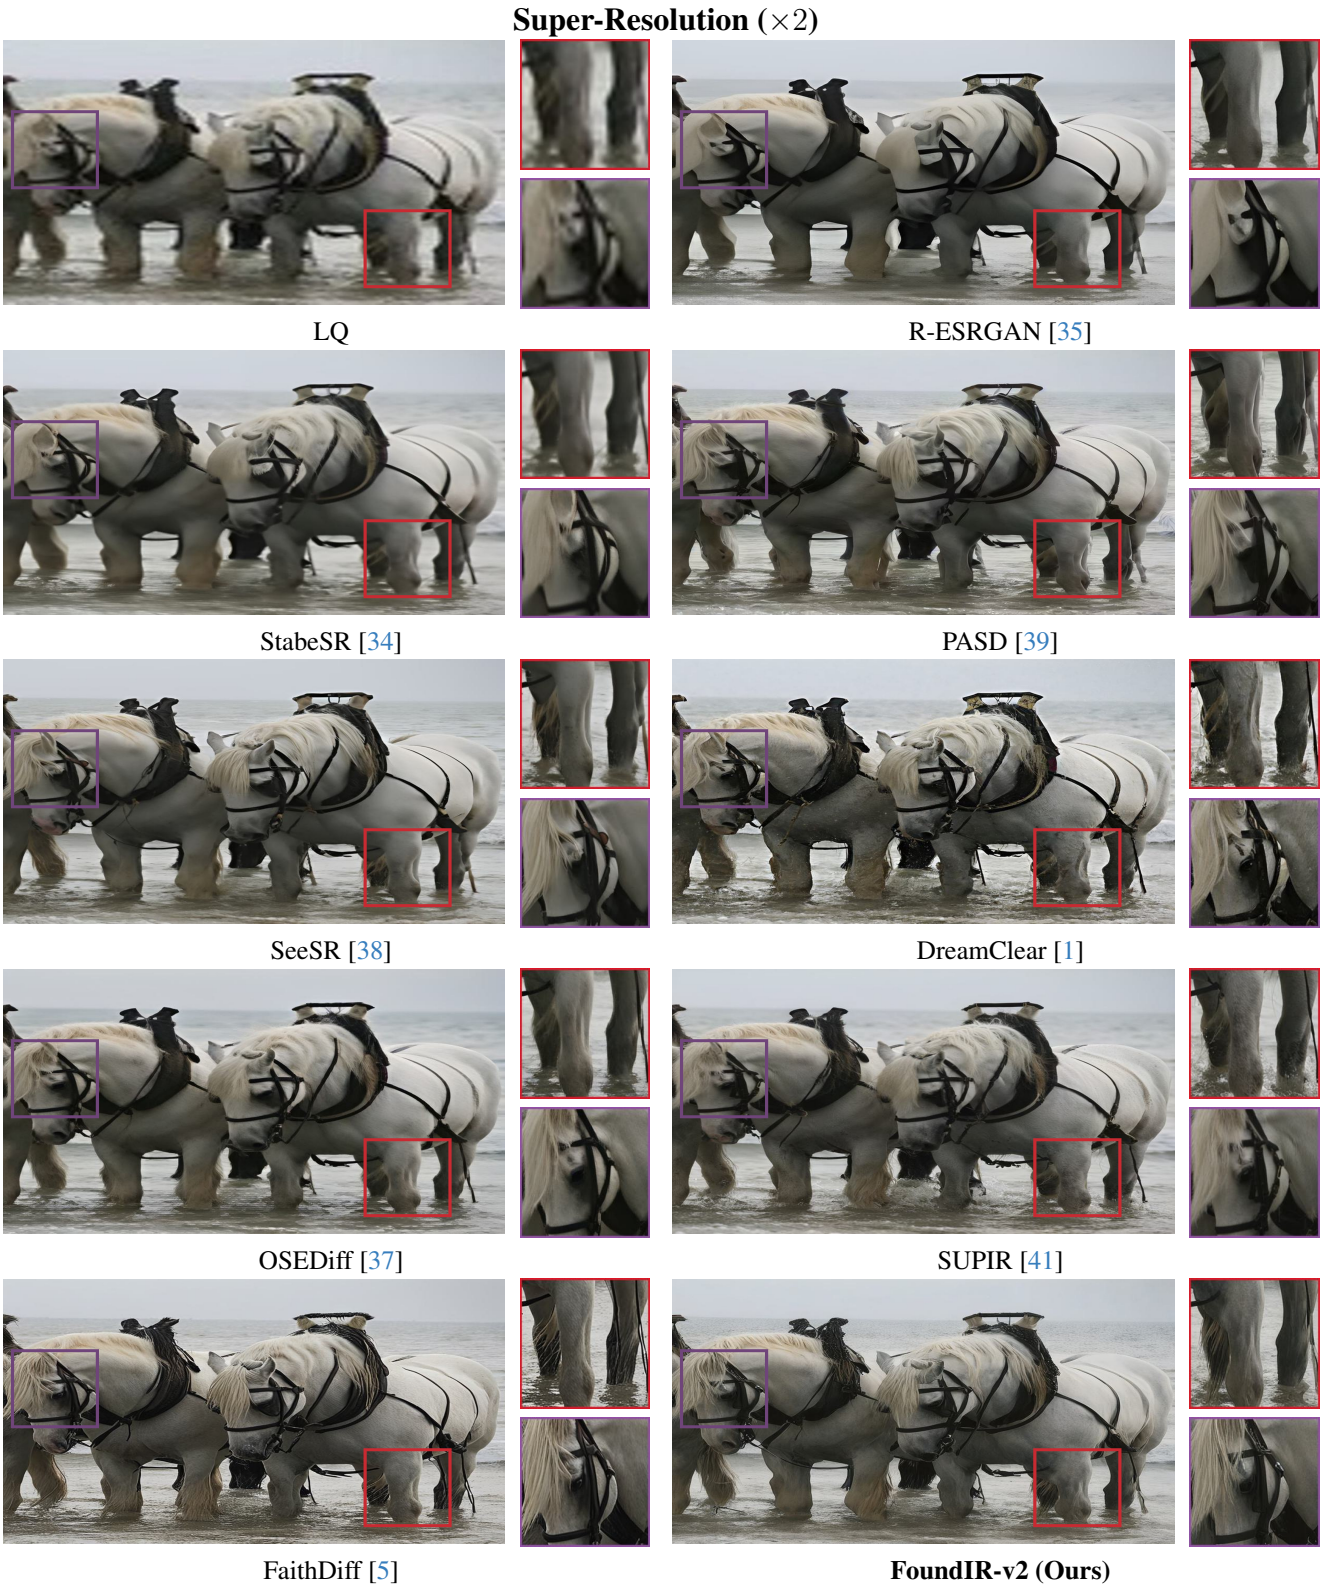

Figure 31. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 32. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

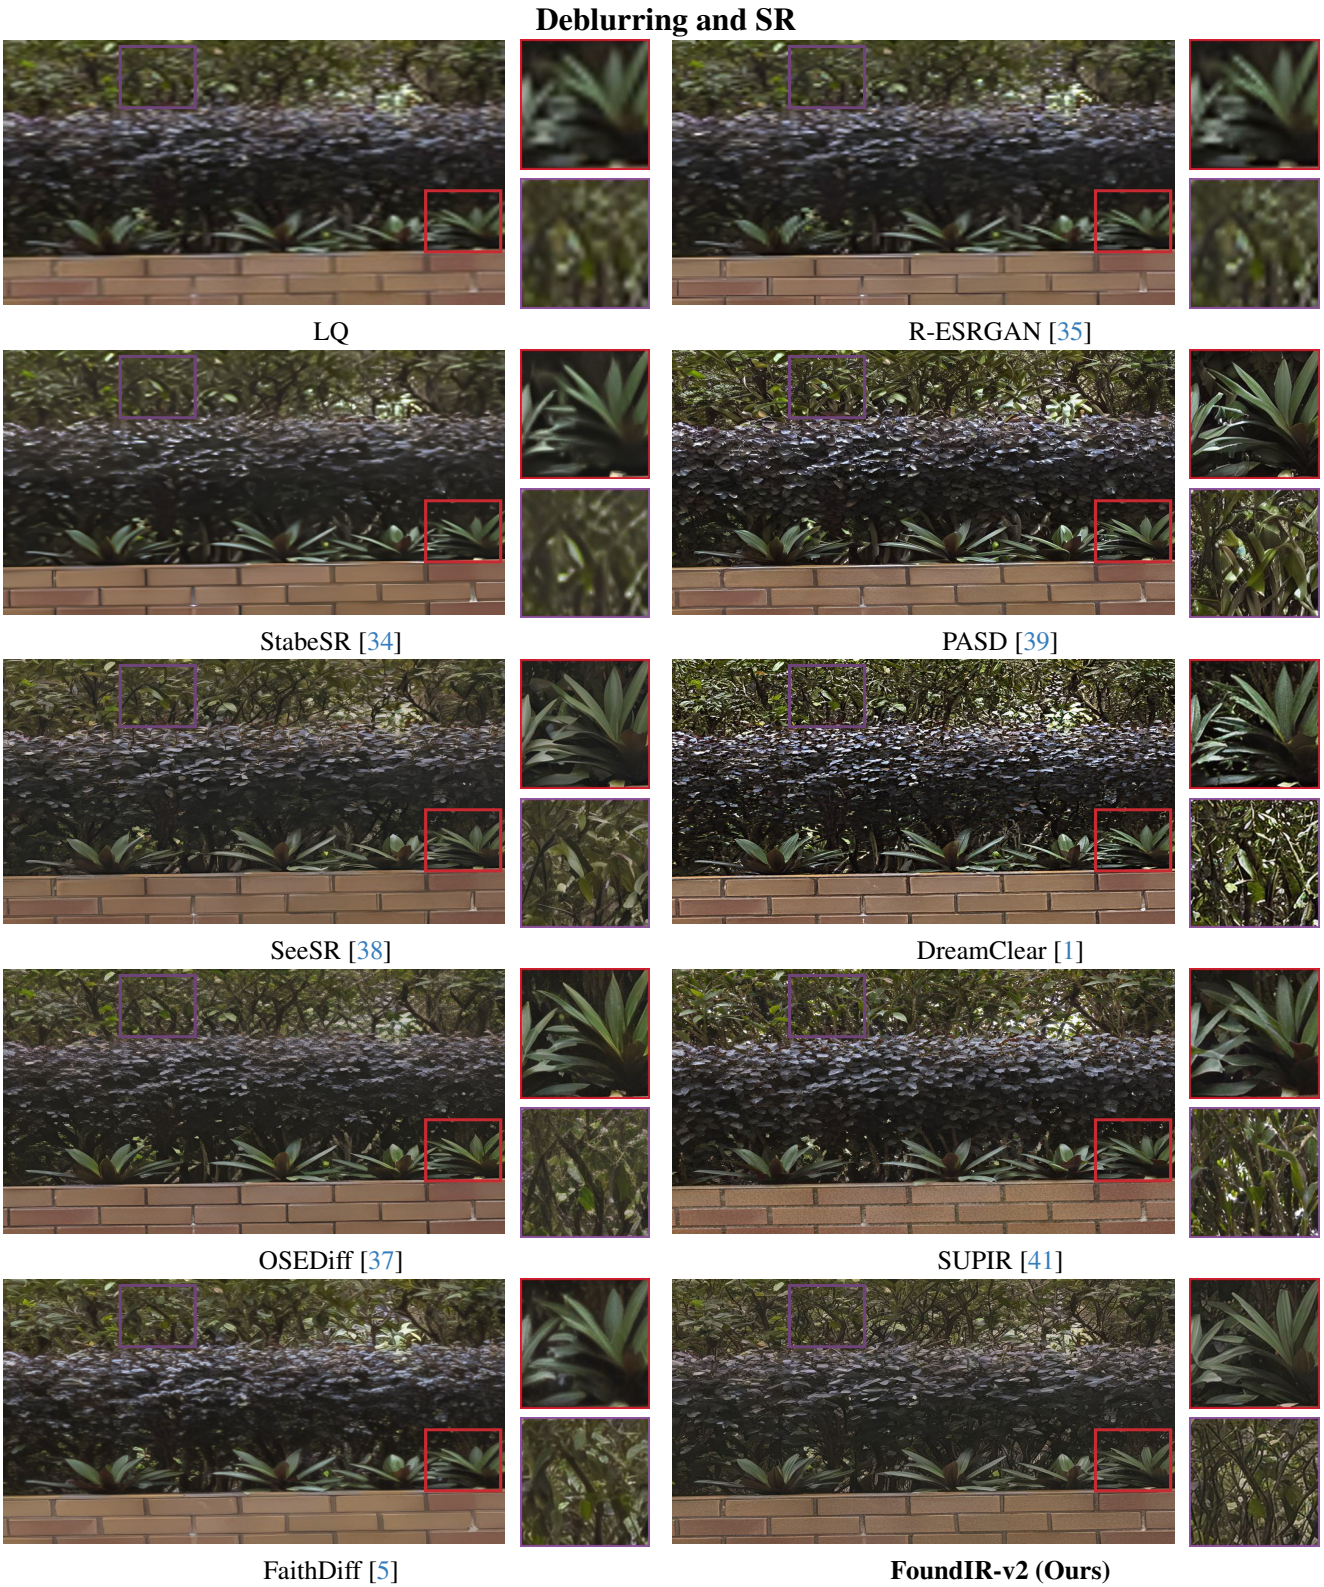

Figure 33. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

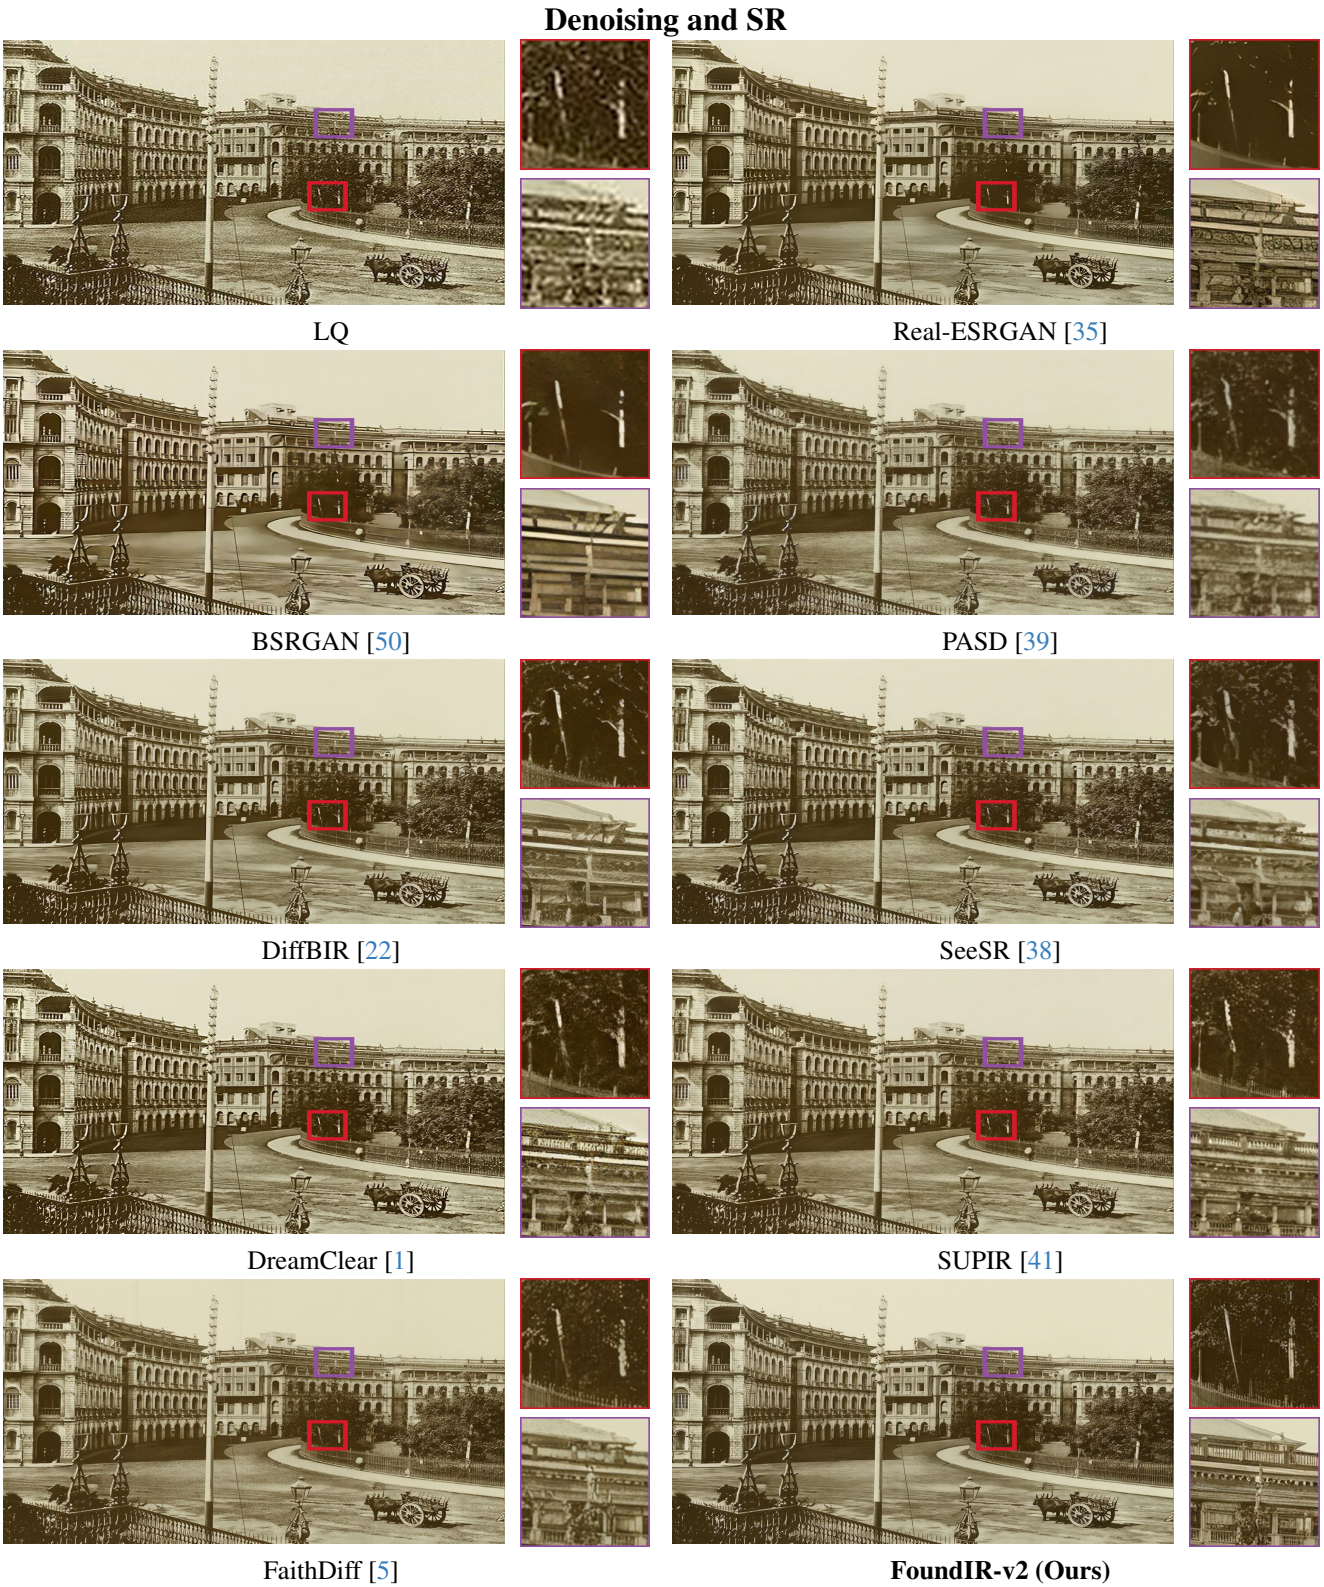

Figure 34. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

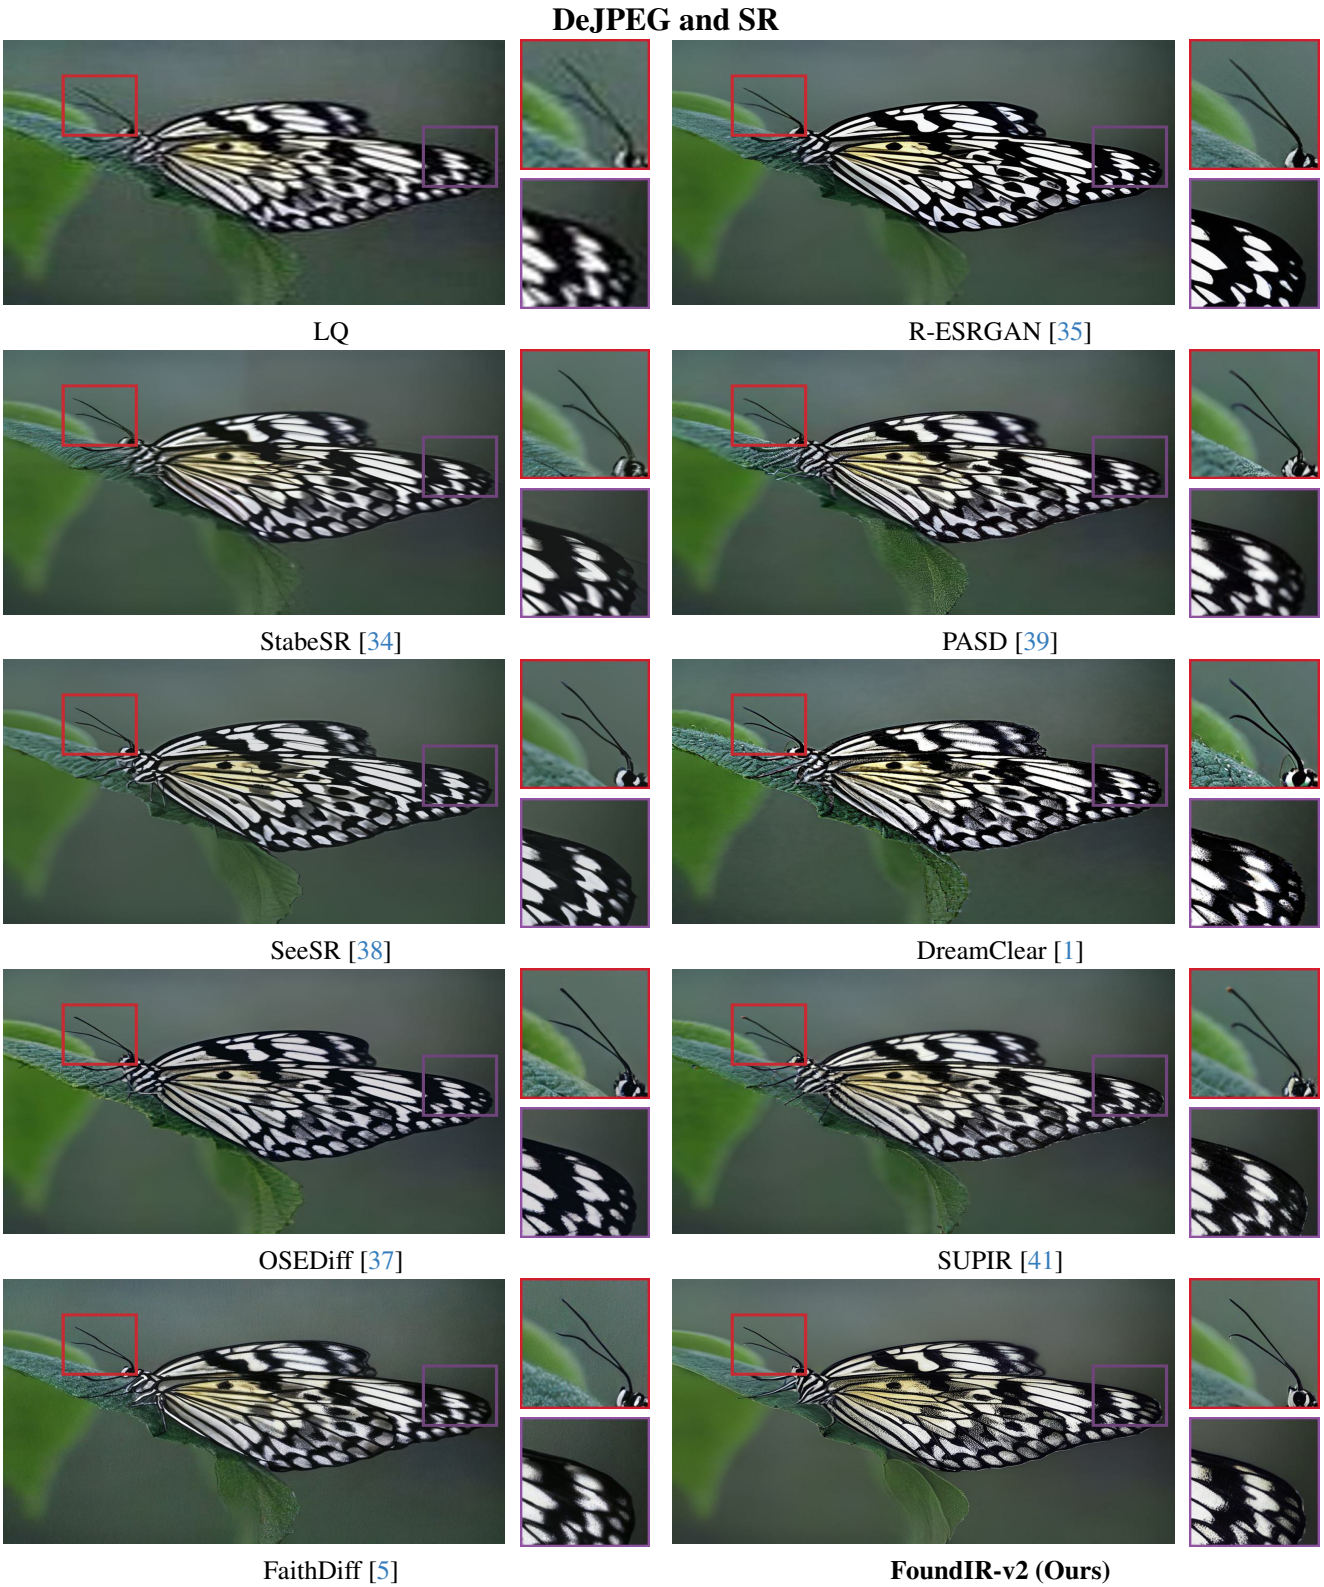

Figure 35. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

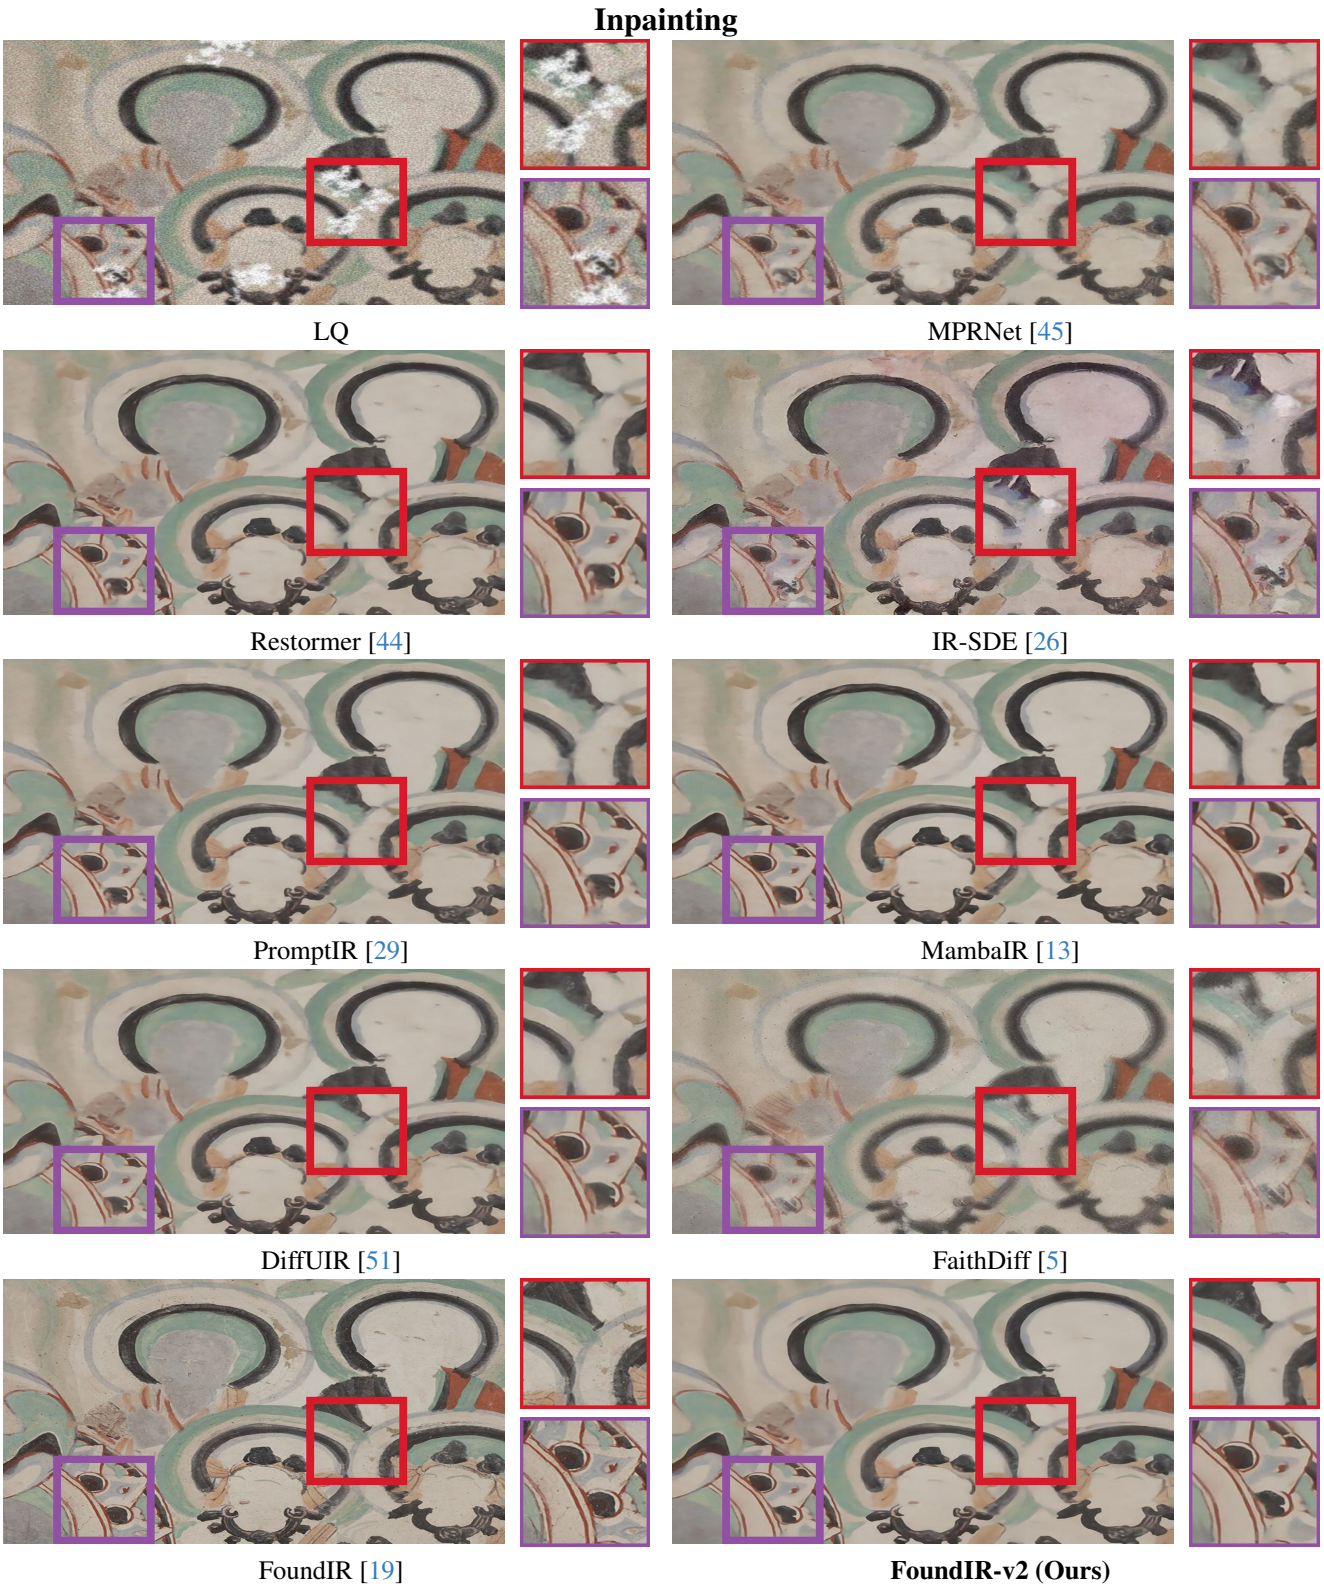

Figure 36. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

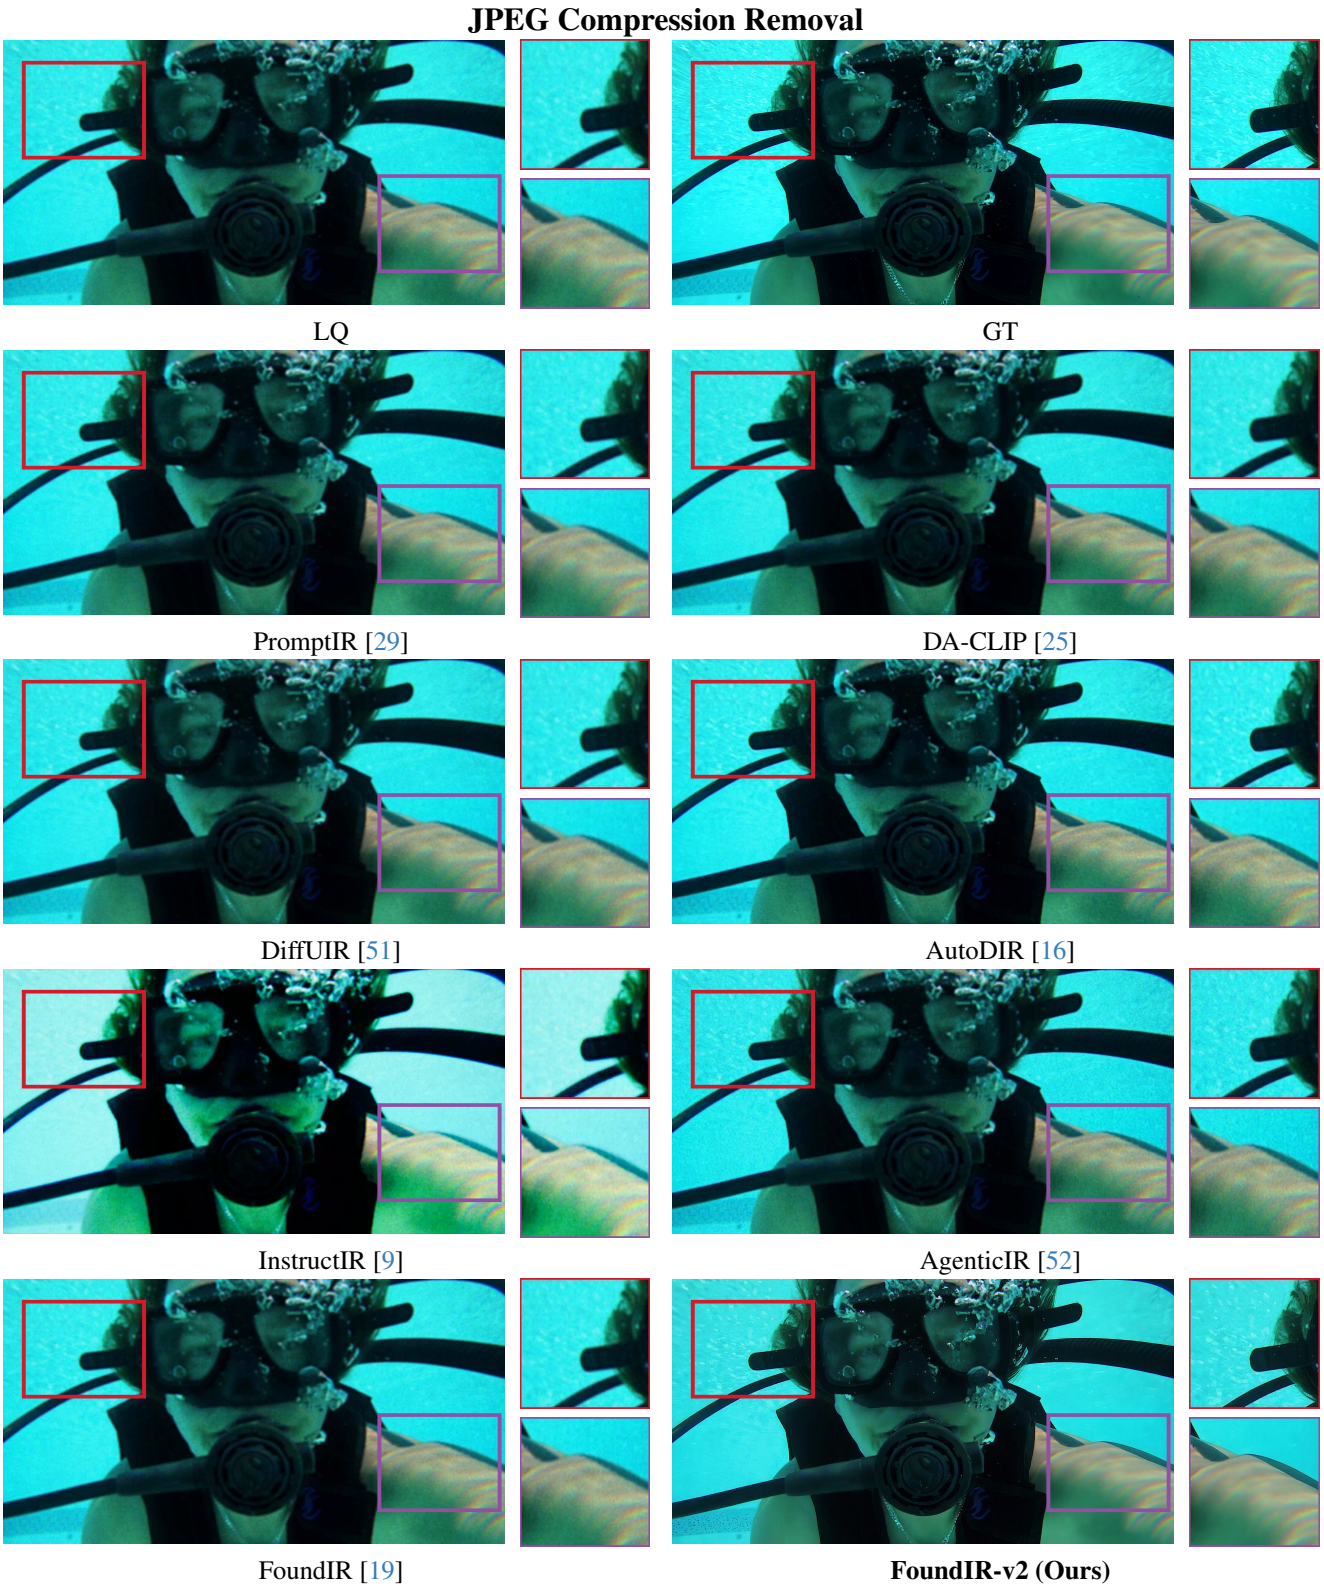

Figure 37. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

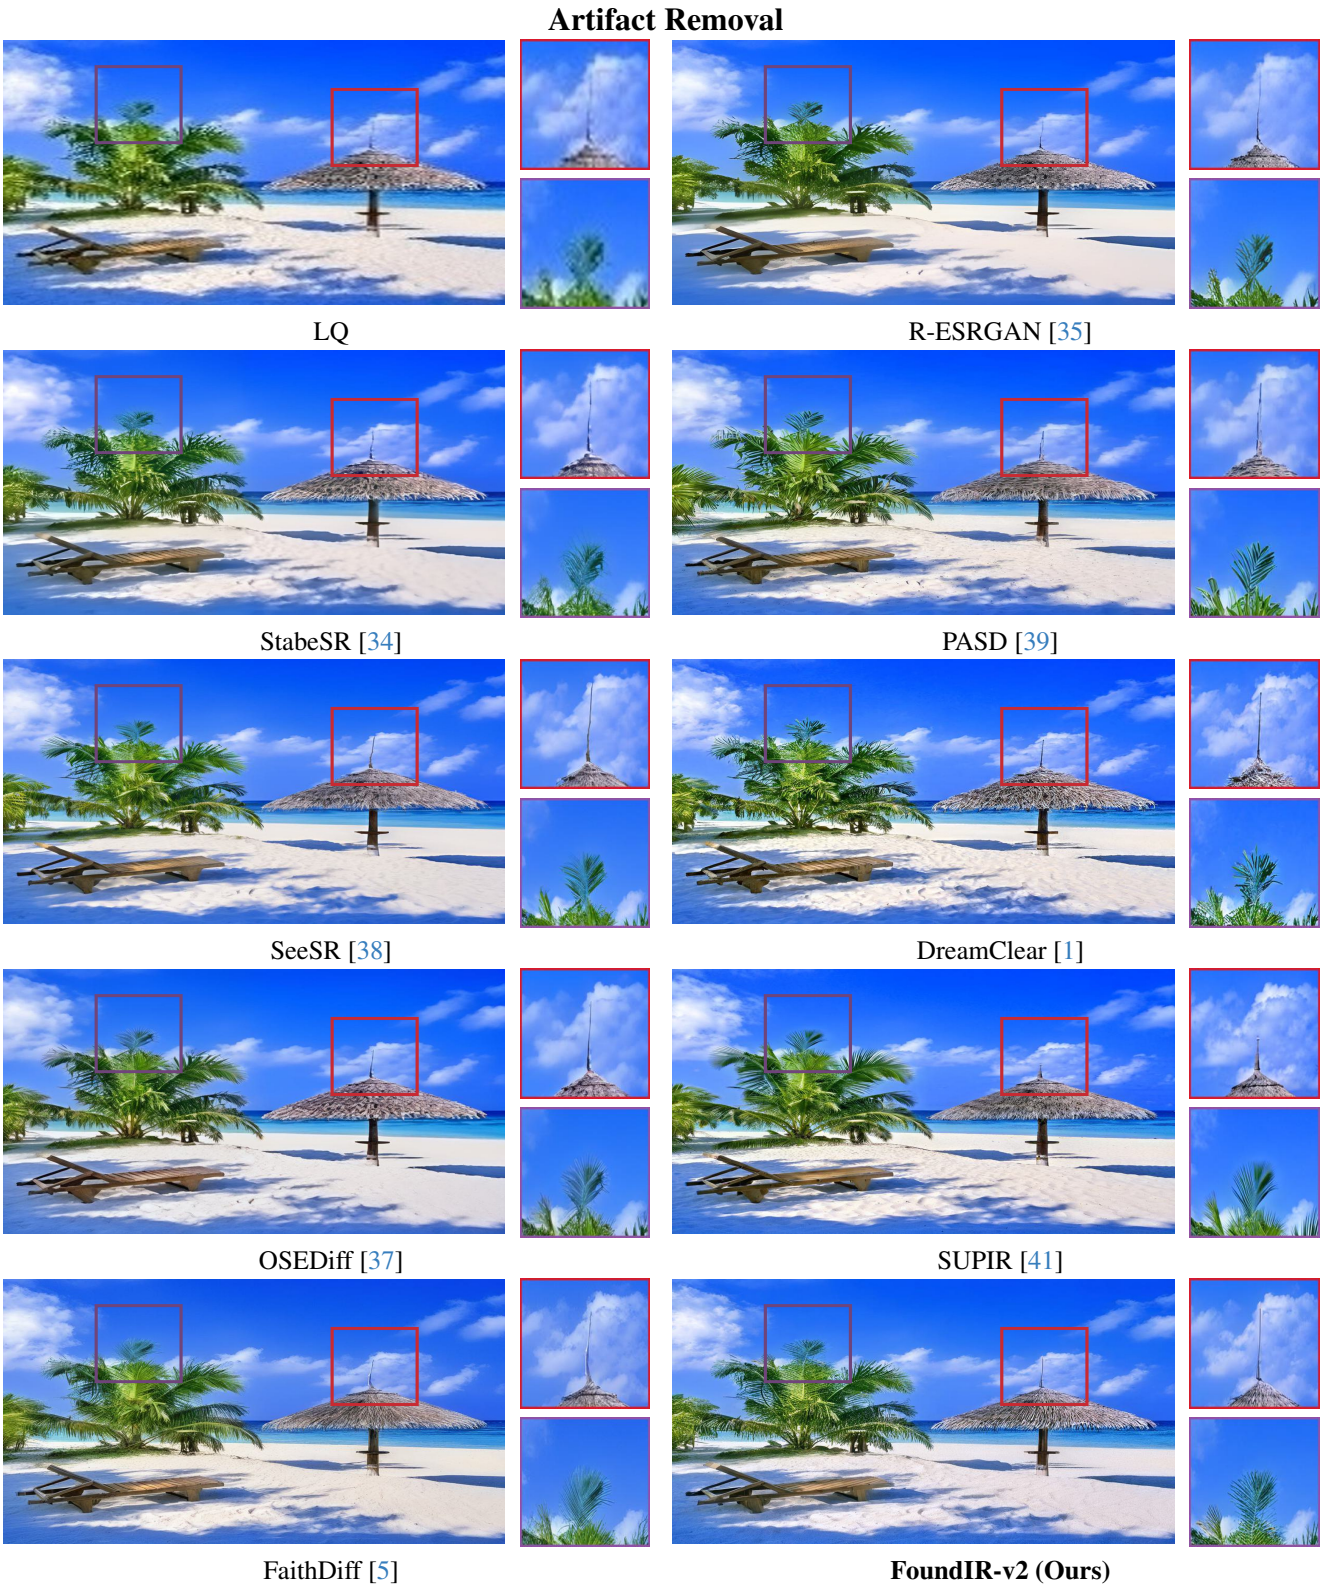

Figure 38. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

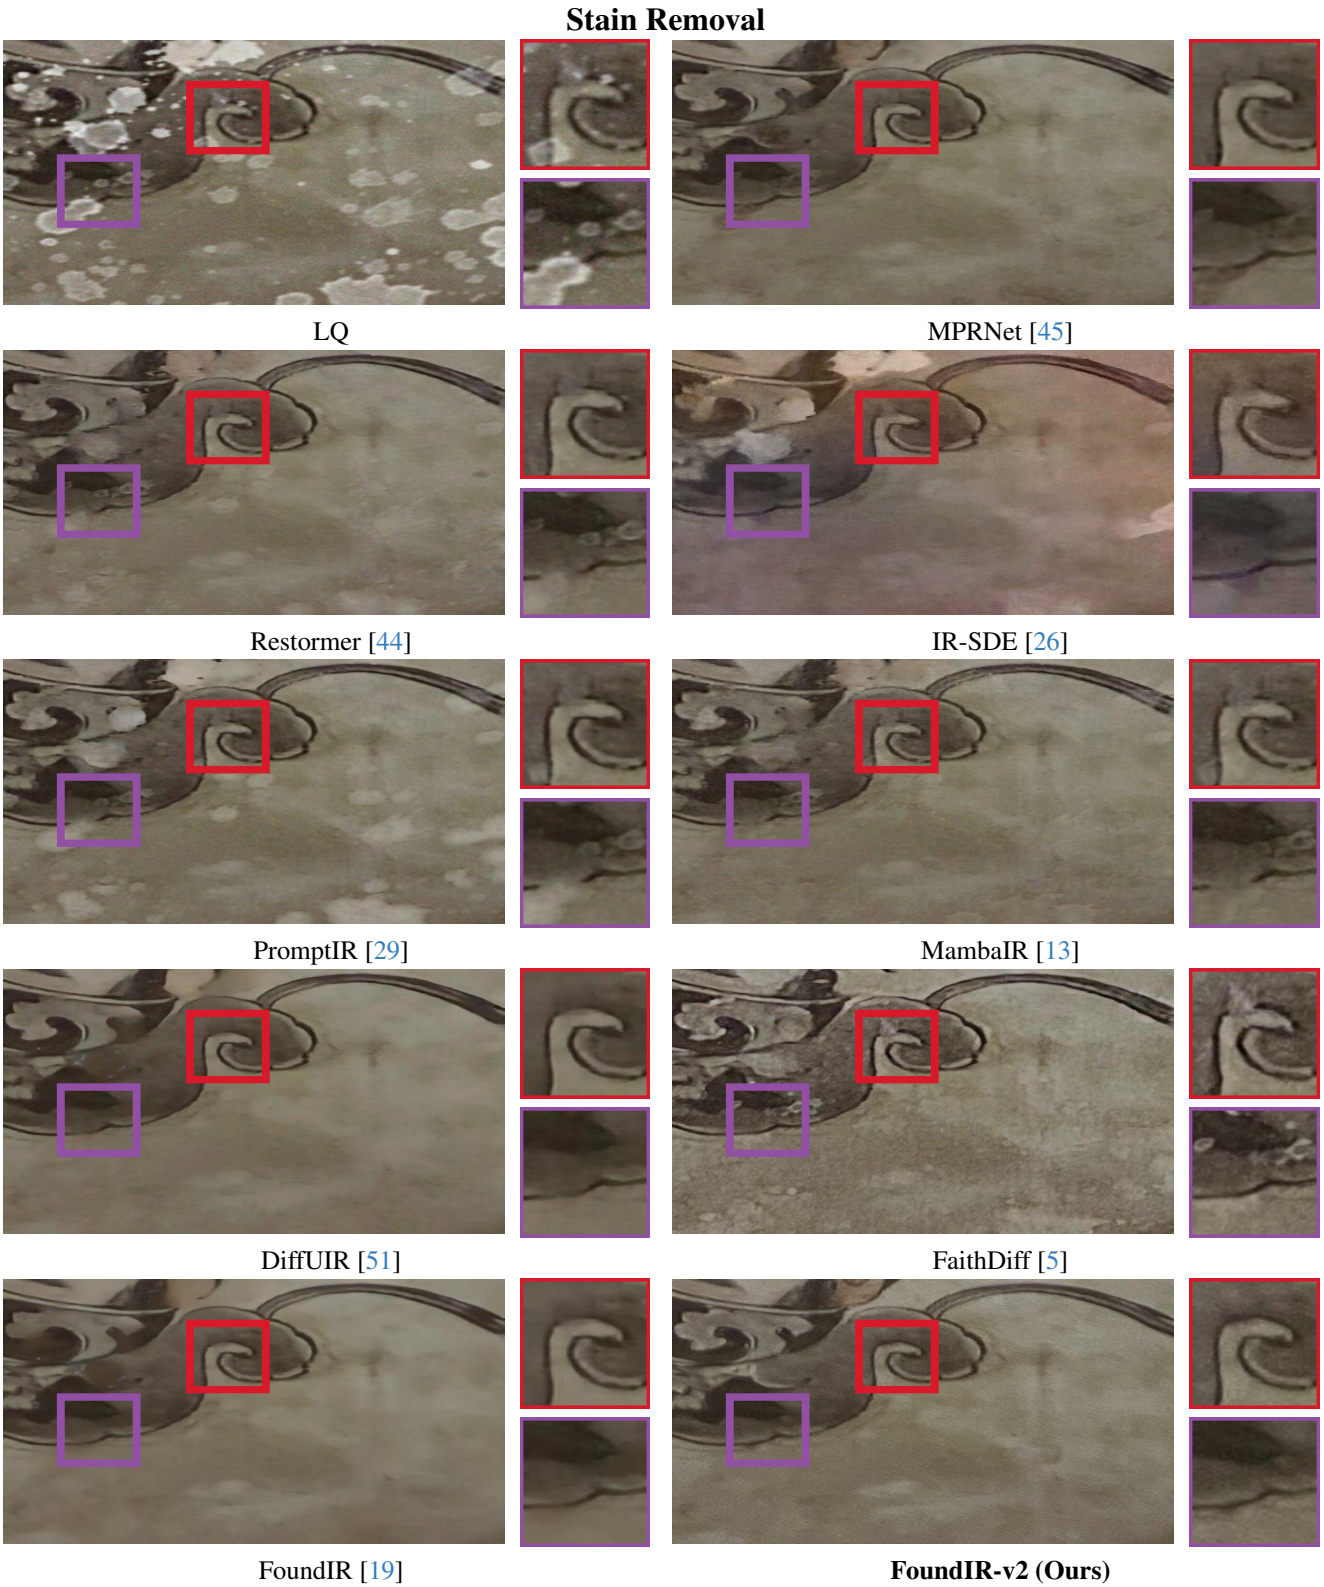

Figure 39. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

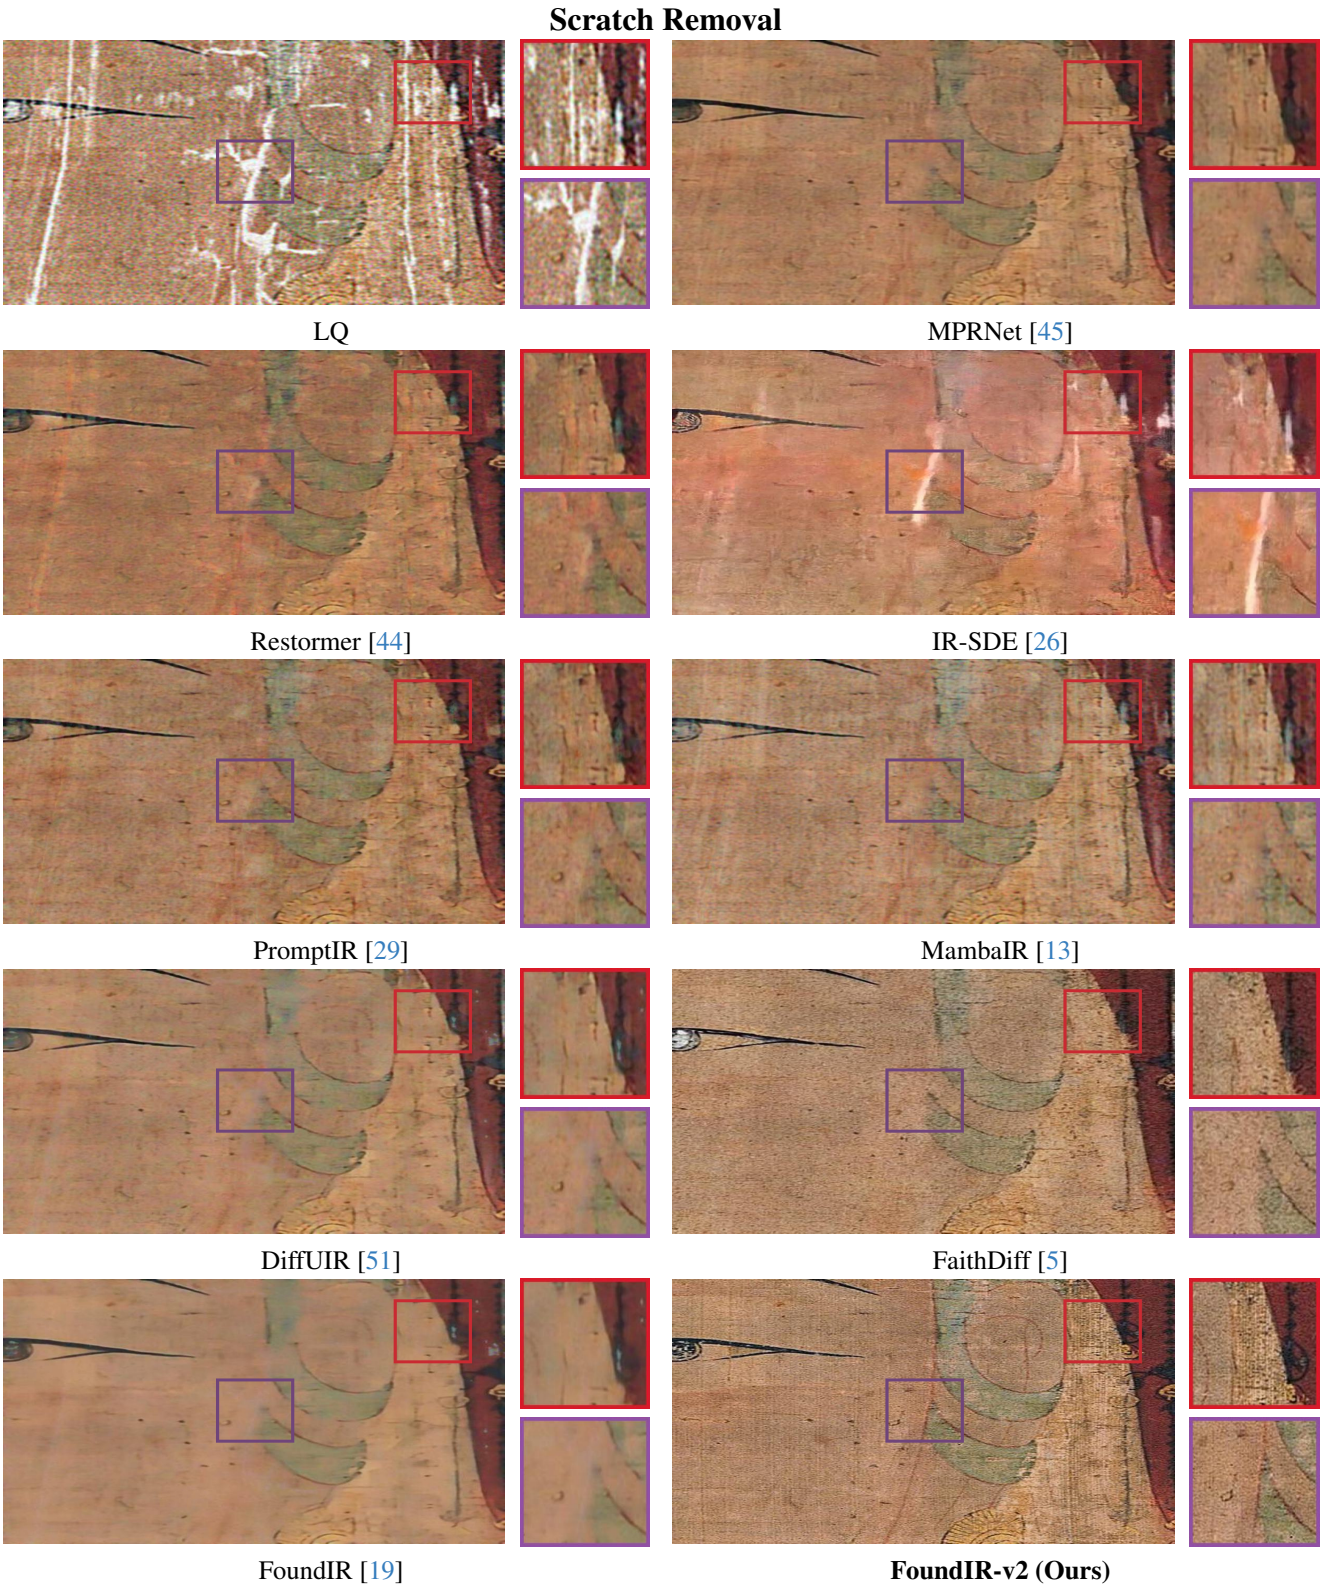

Figure 40. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

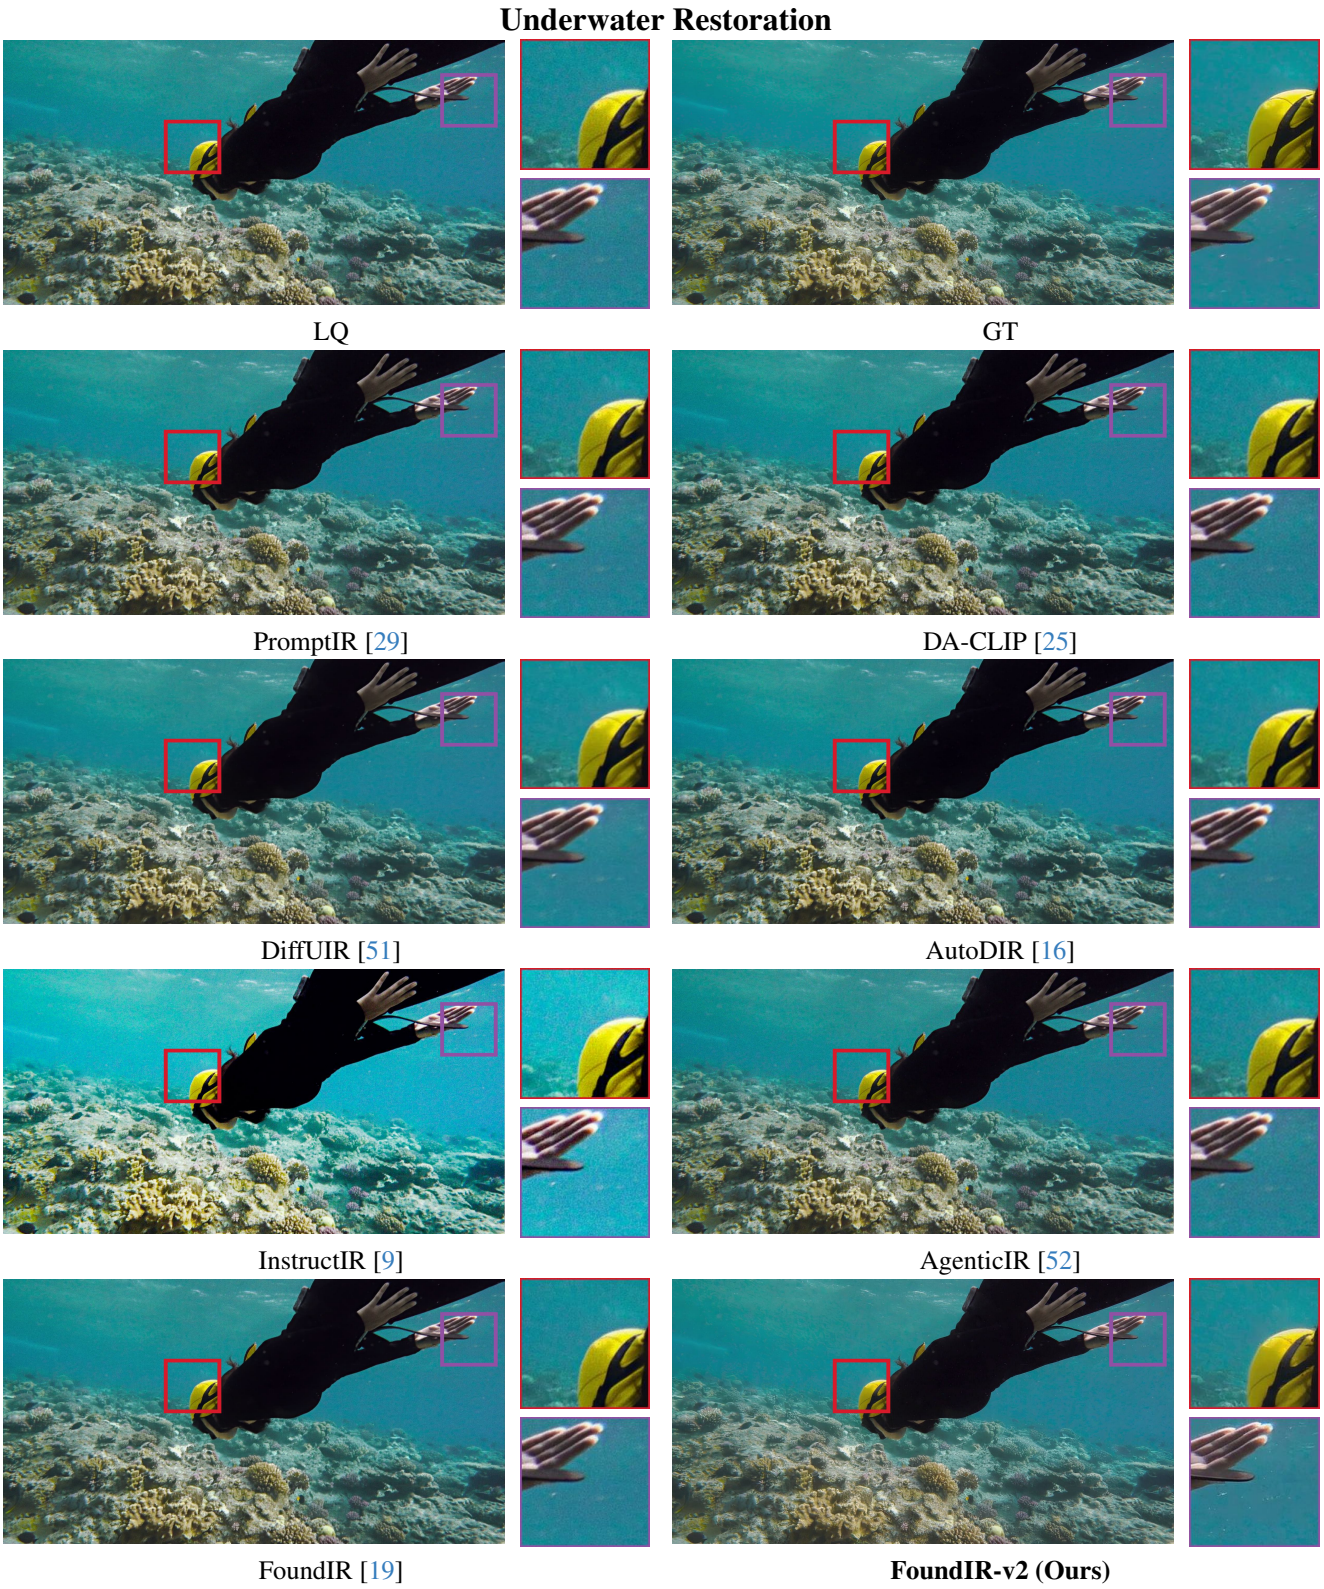

Figure 41. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

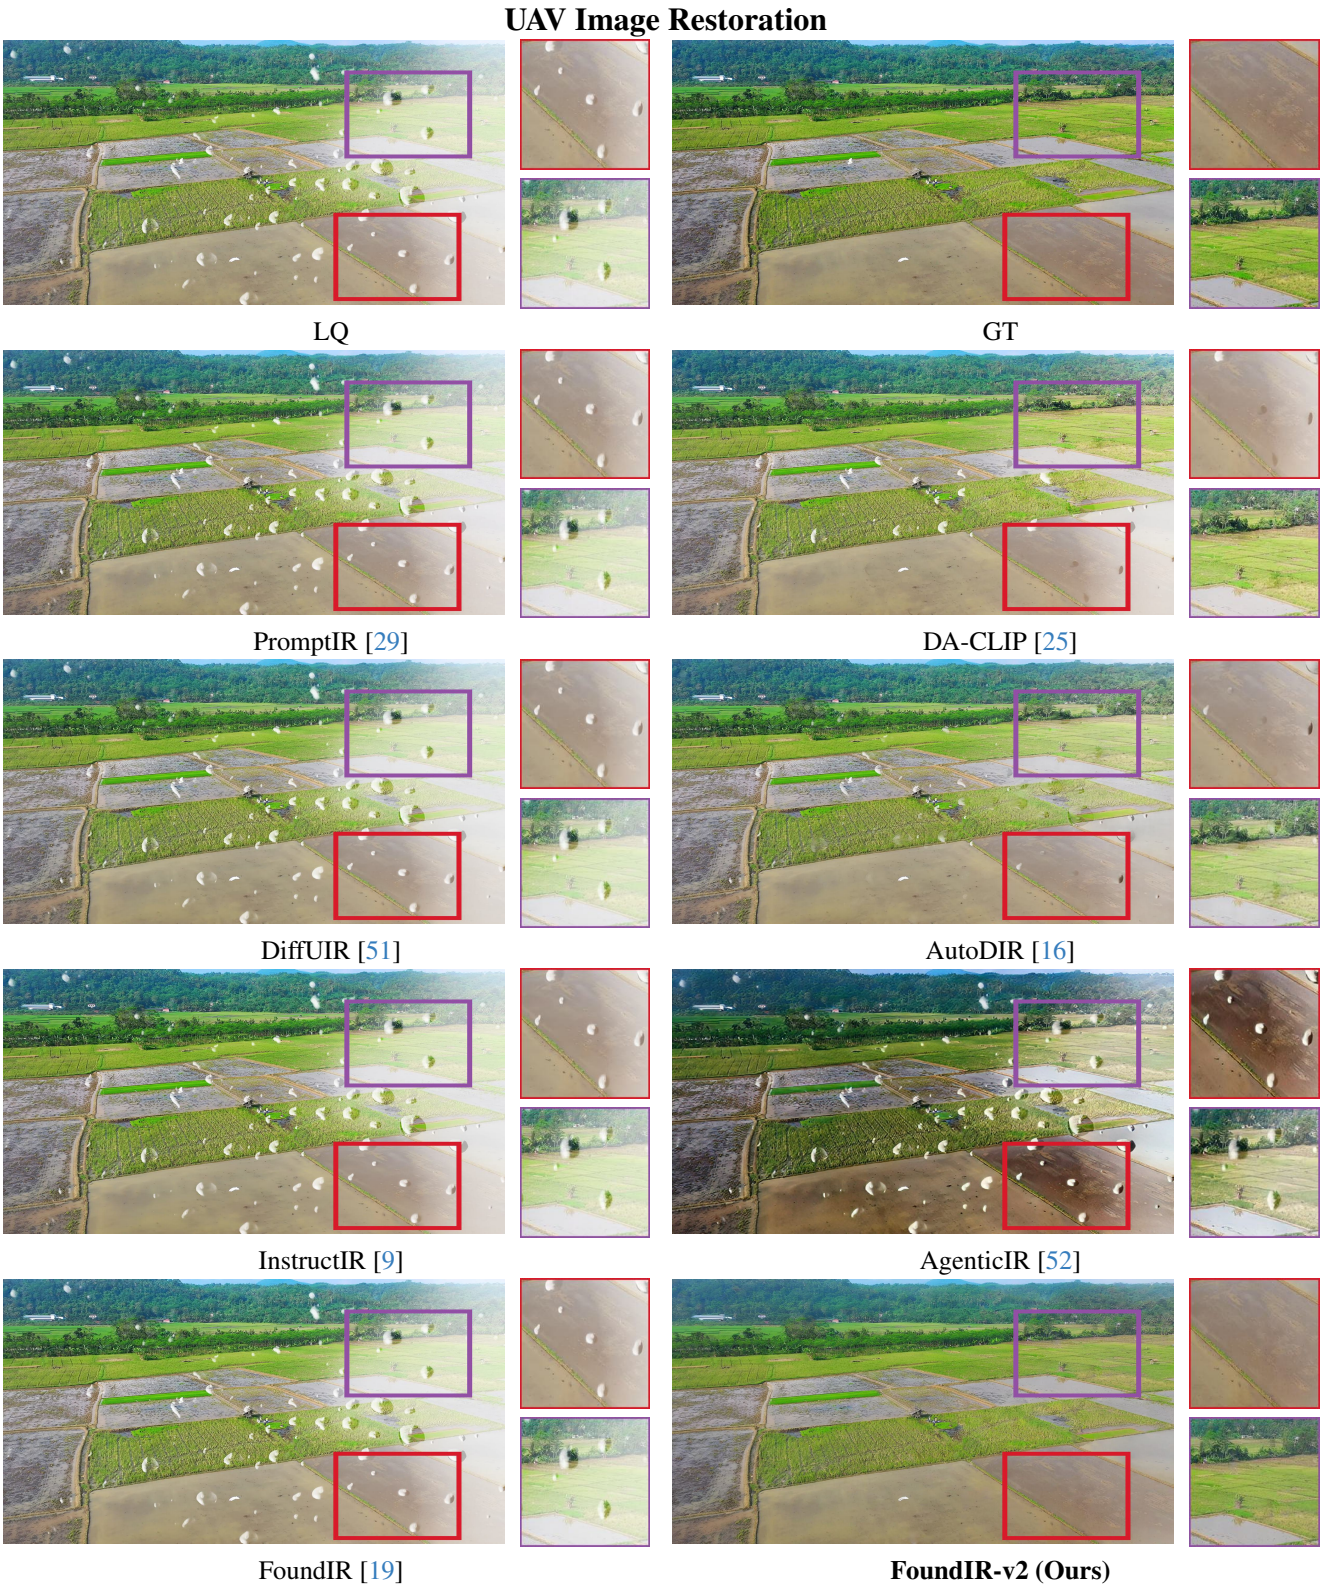

Figure 42. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

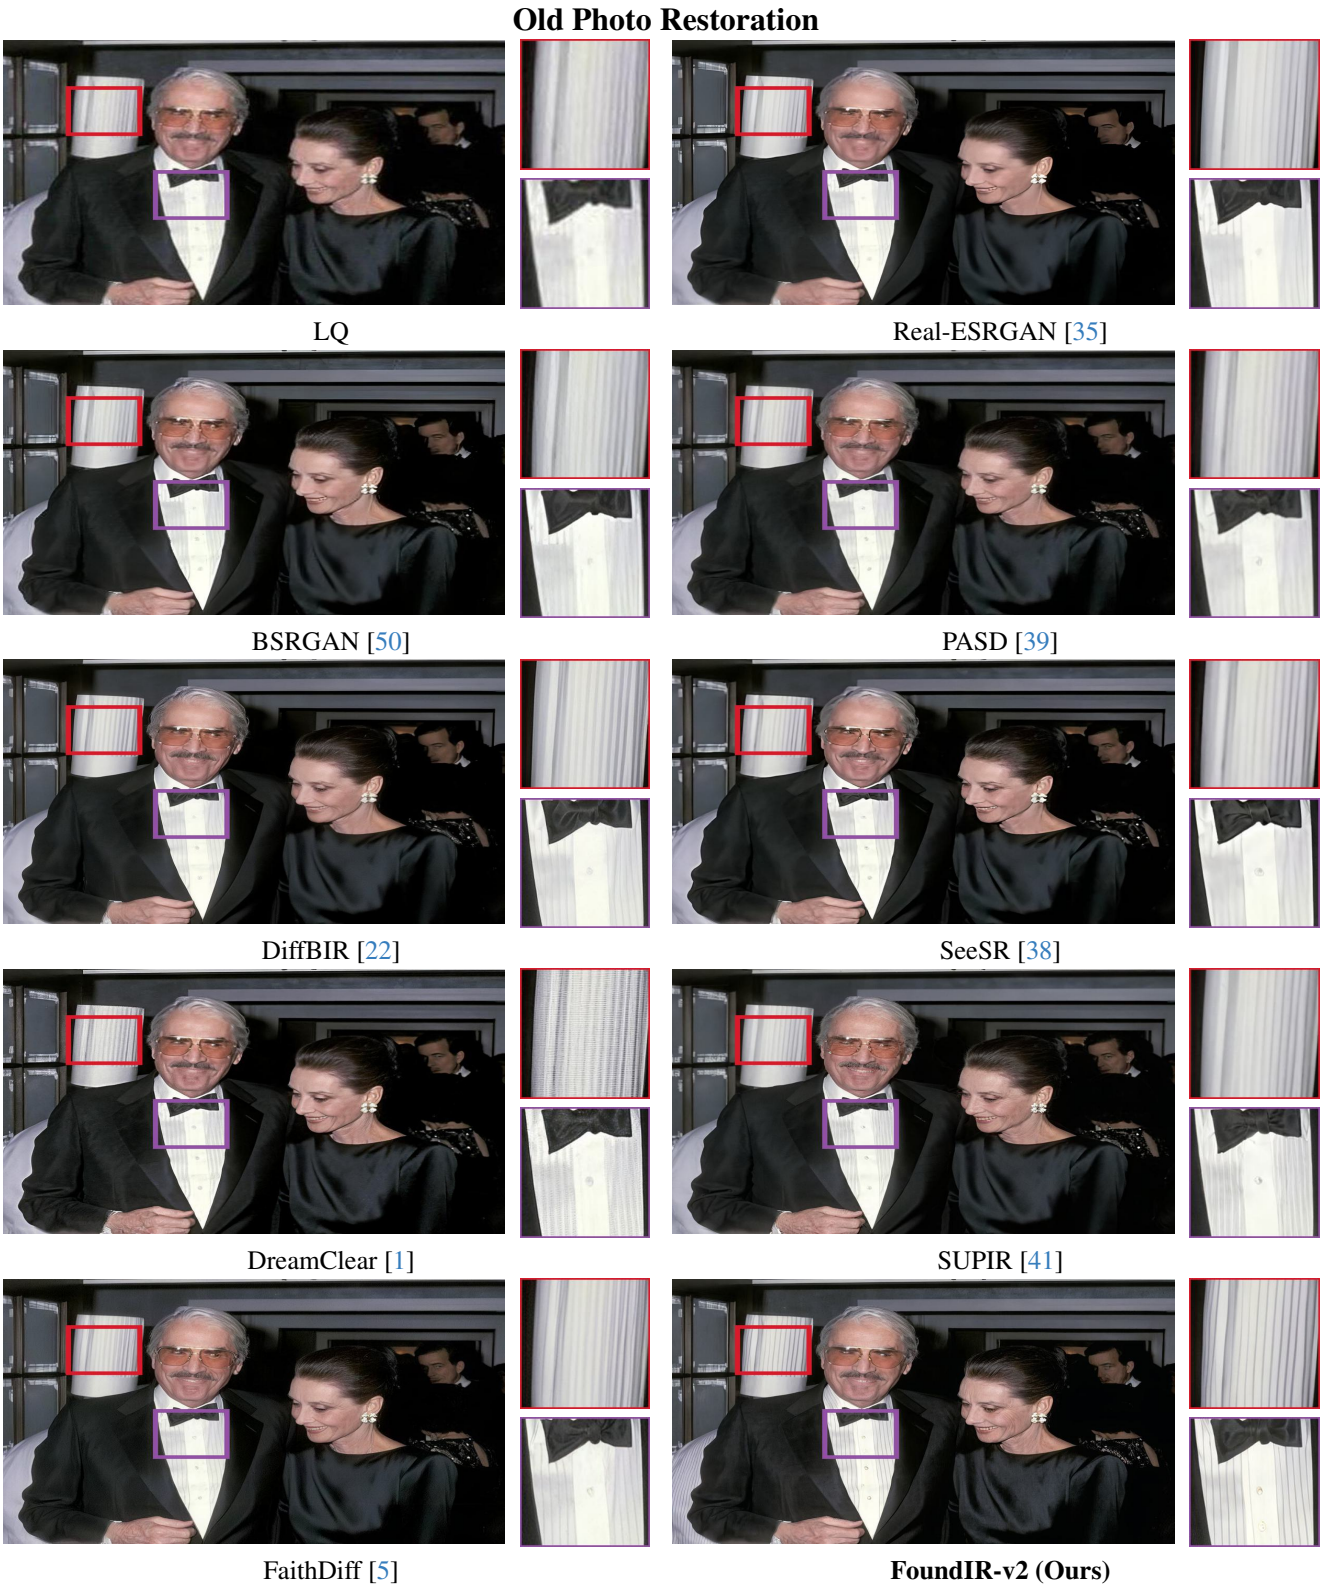

Figure 43. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

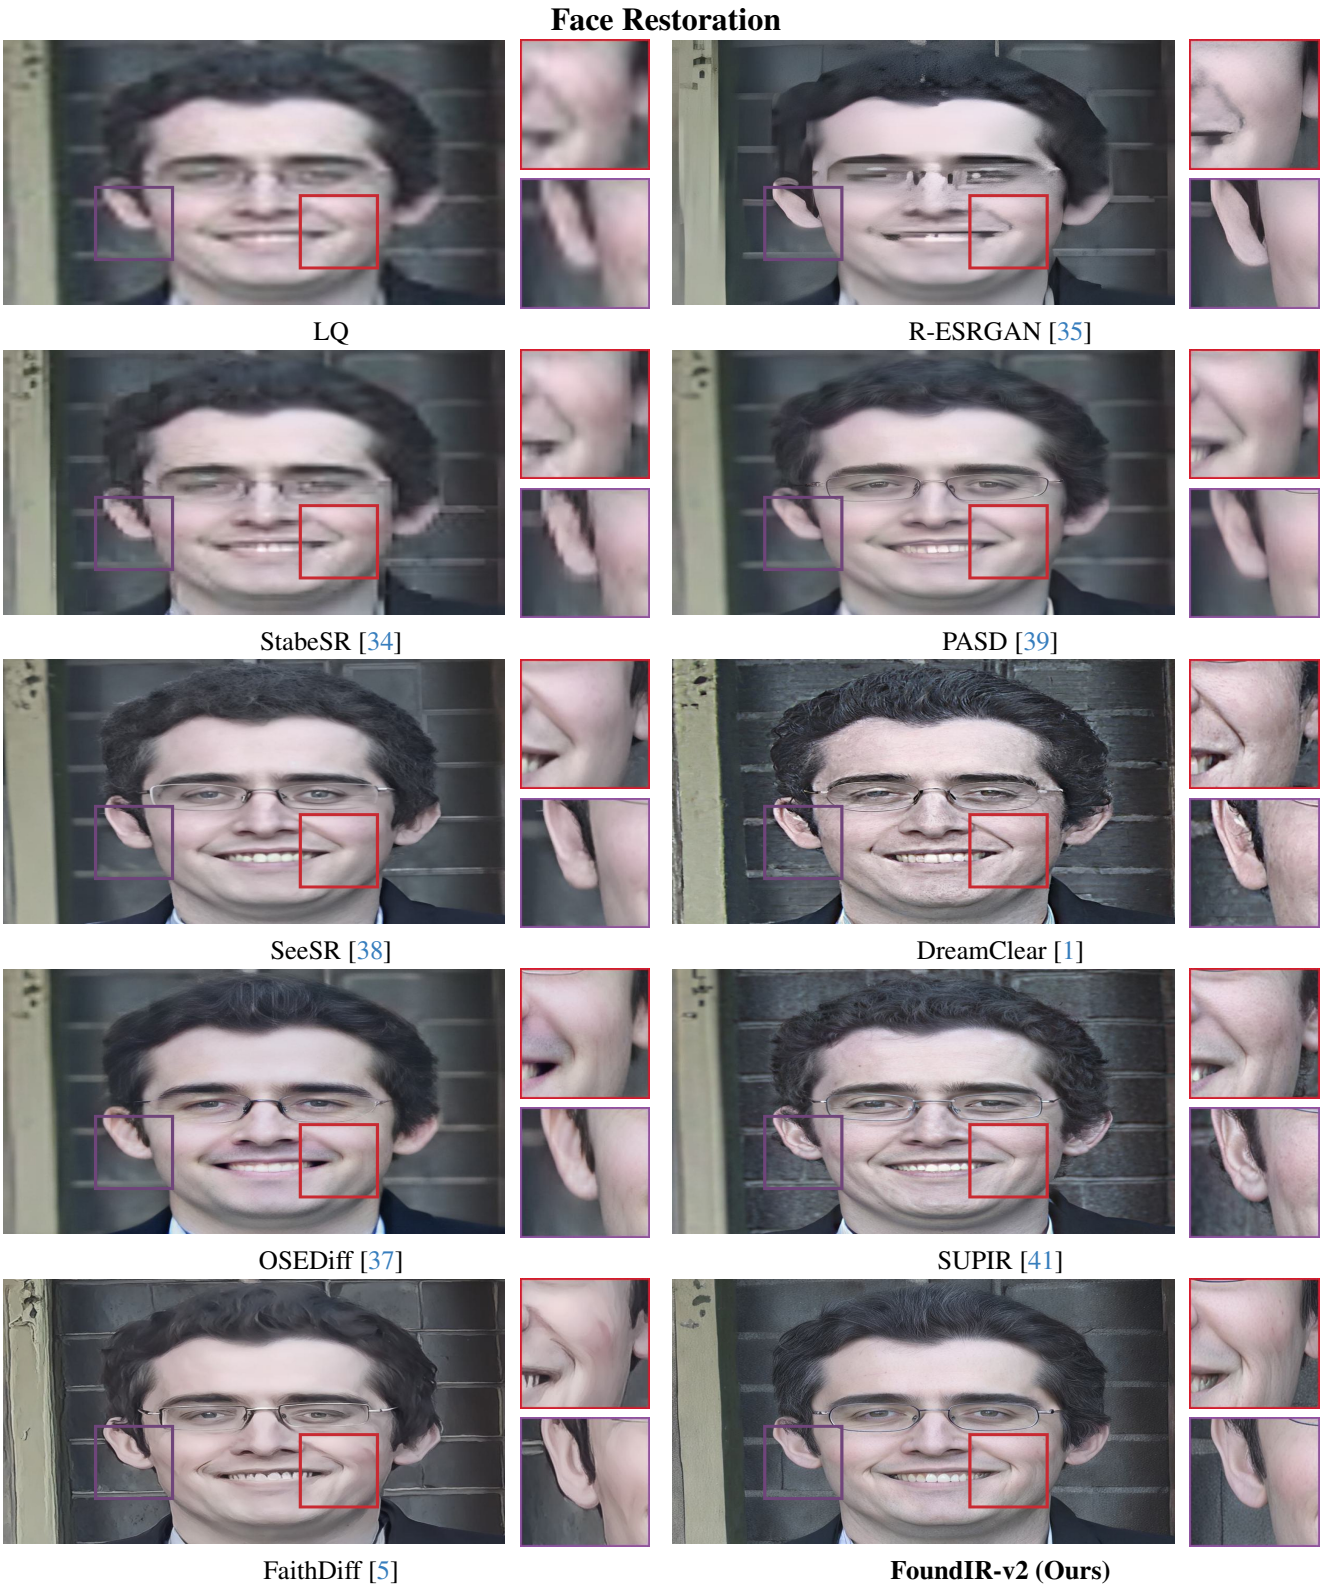

Figure 44. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 45. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

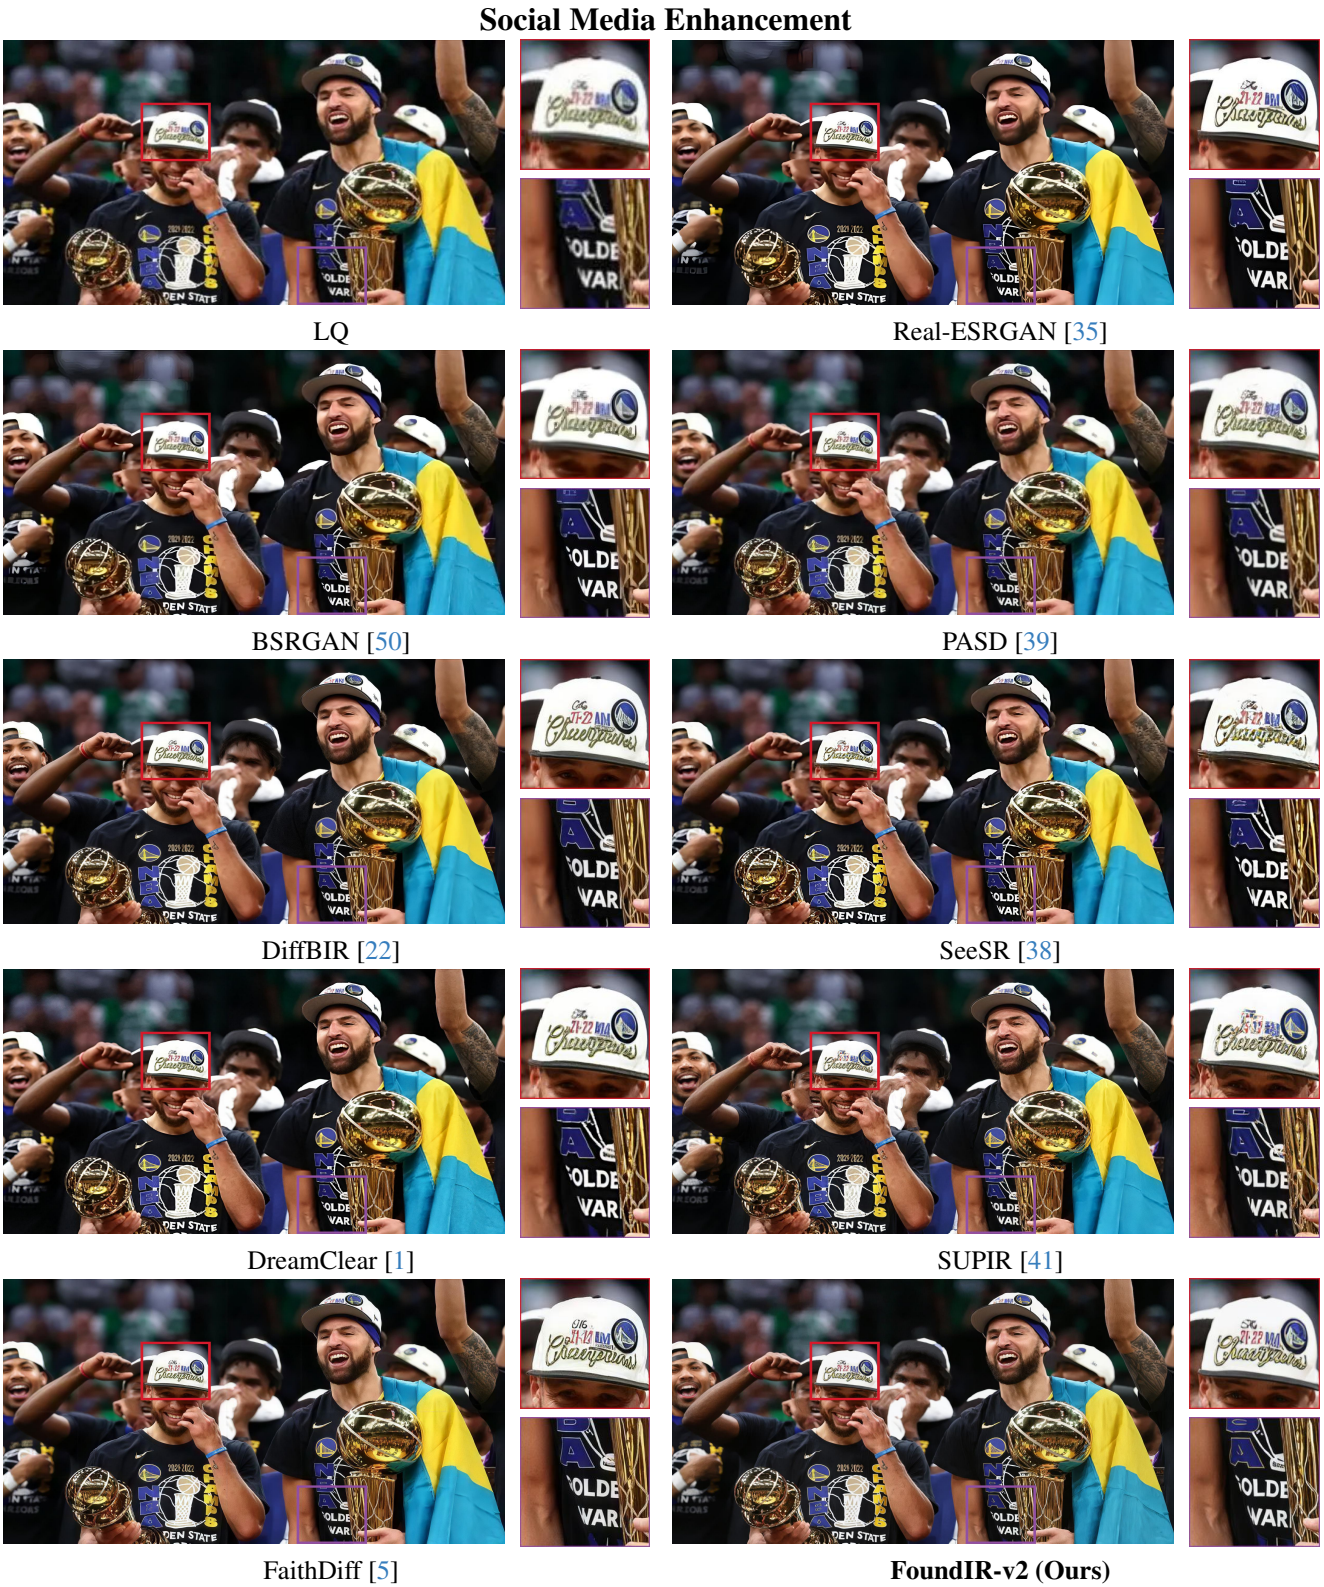

Figure 46. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 47. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

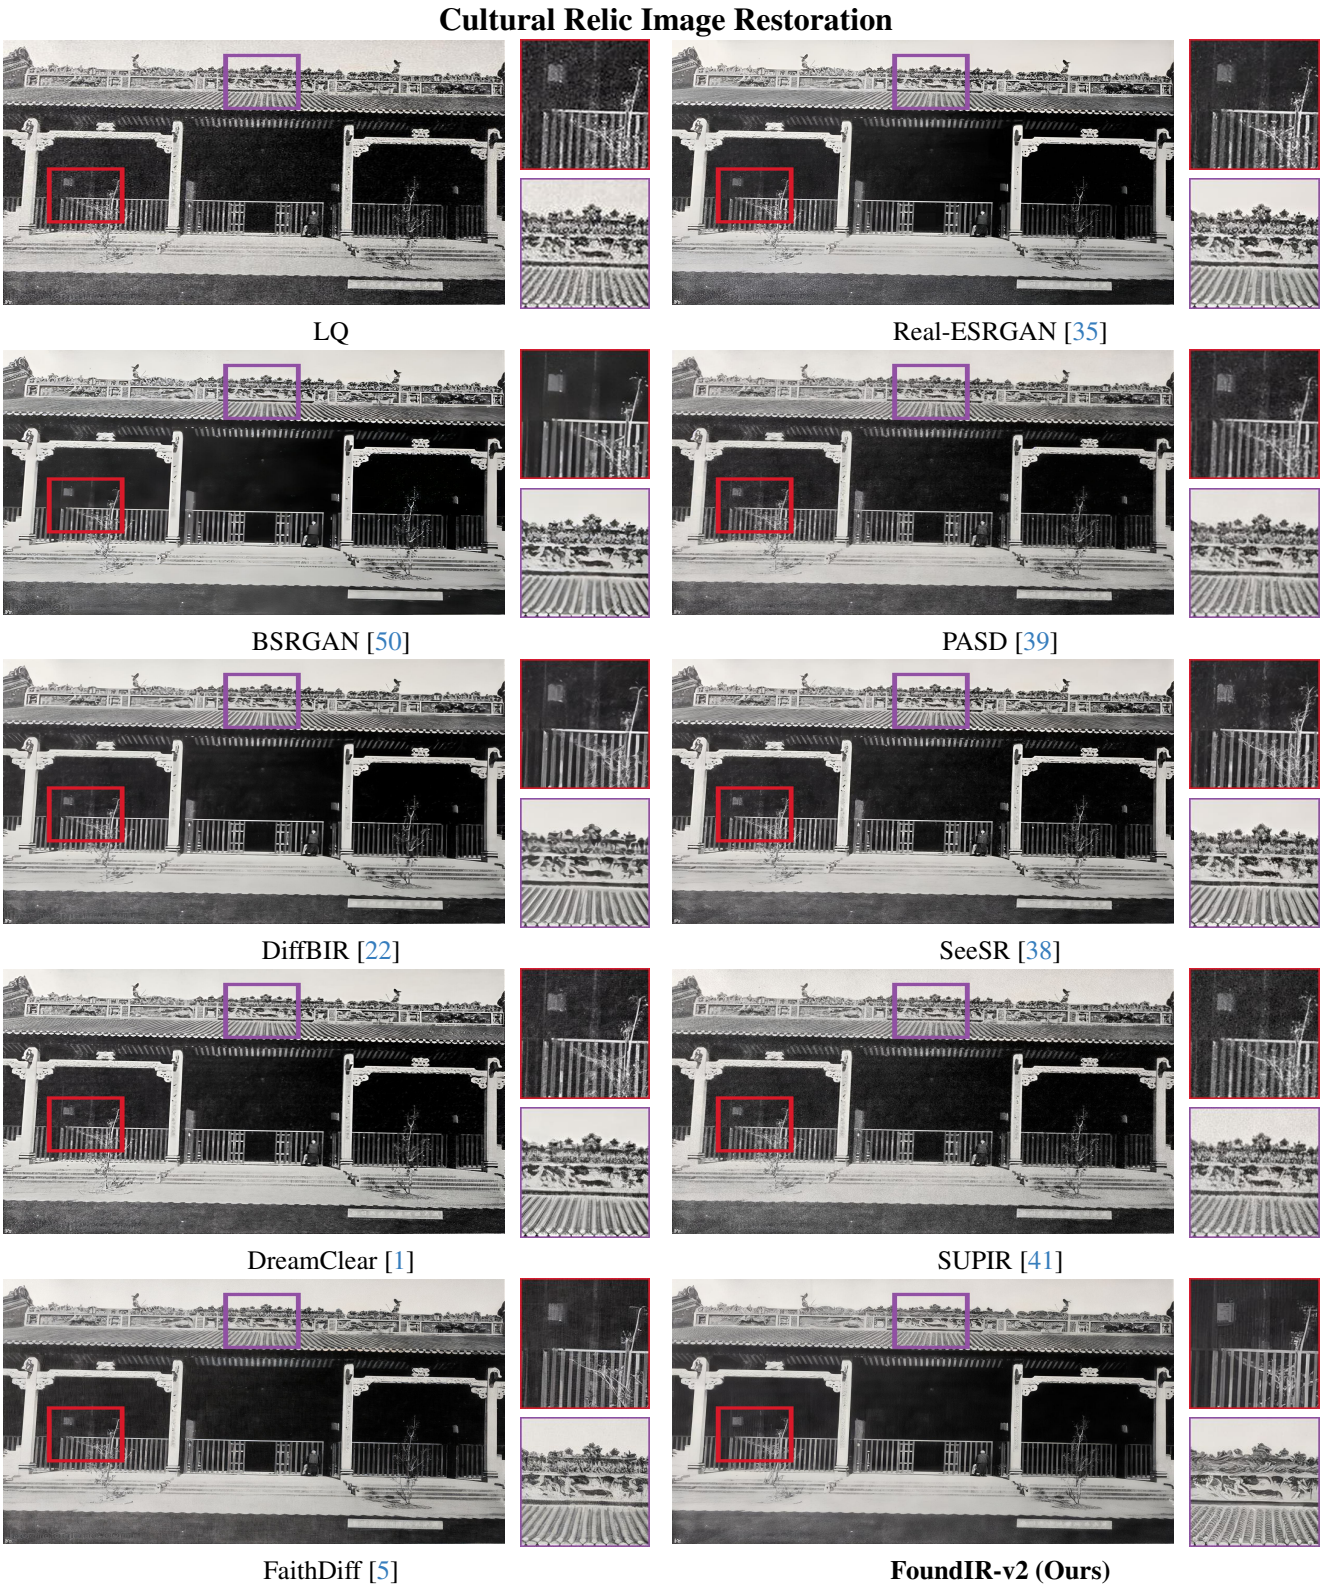

Figure 48. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 49. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

Figure 50. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

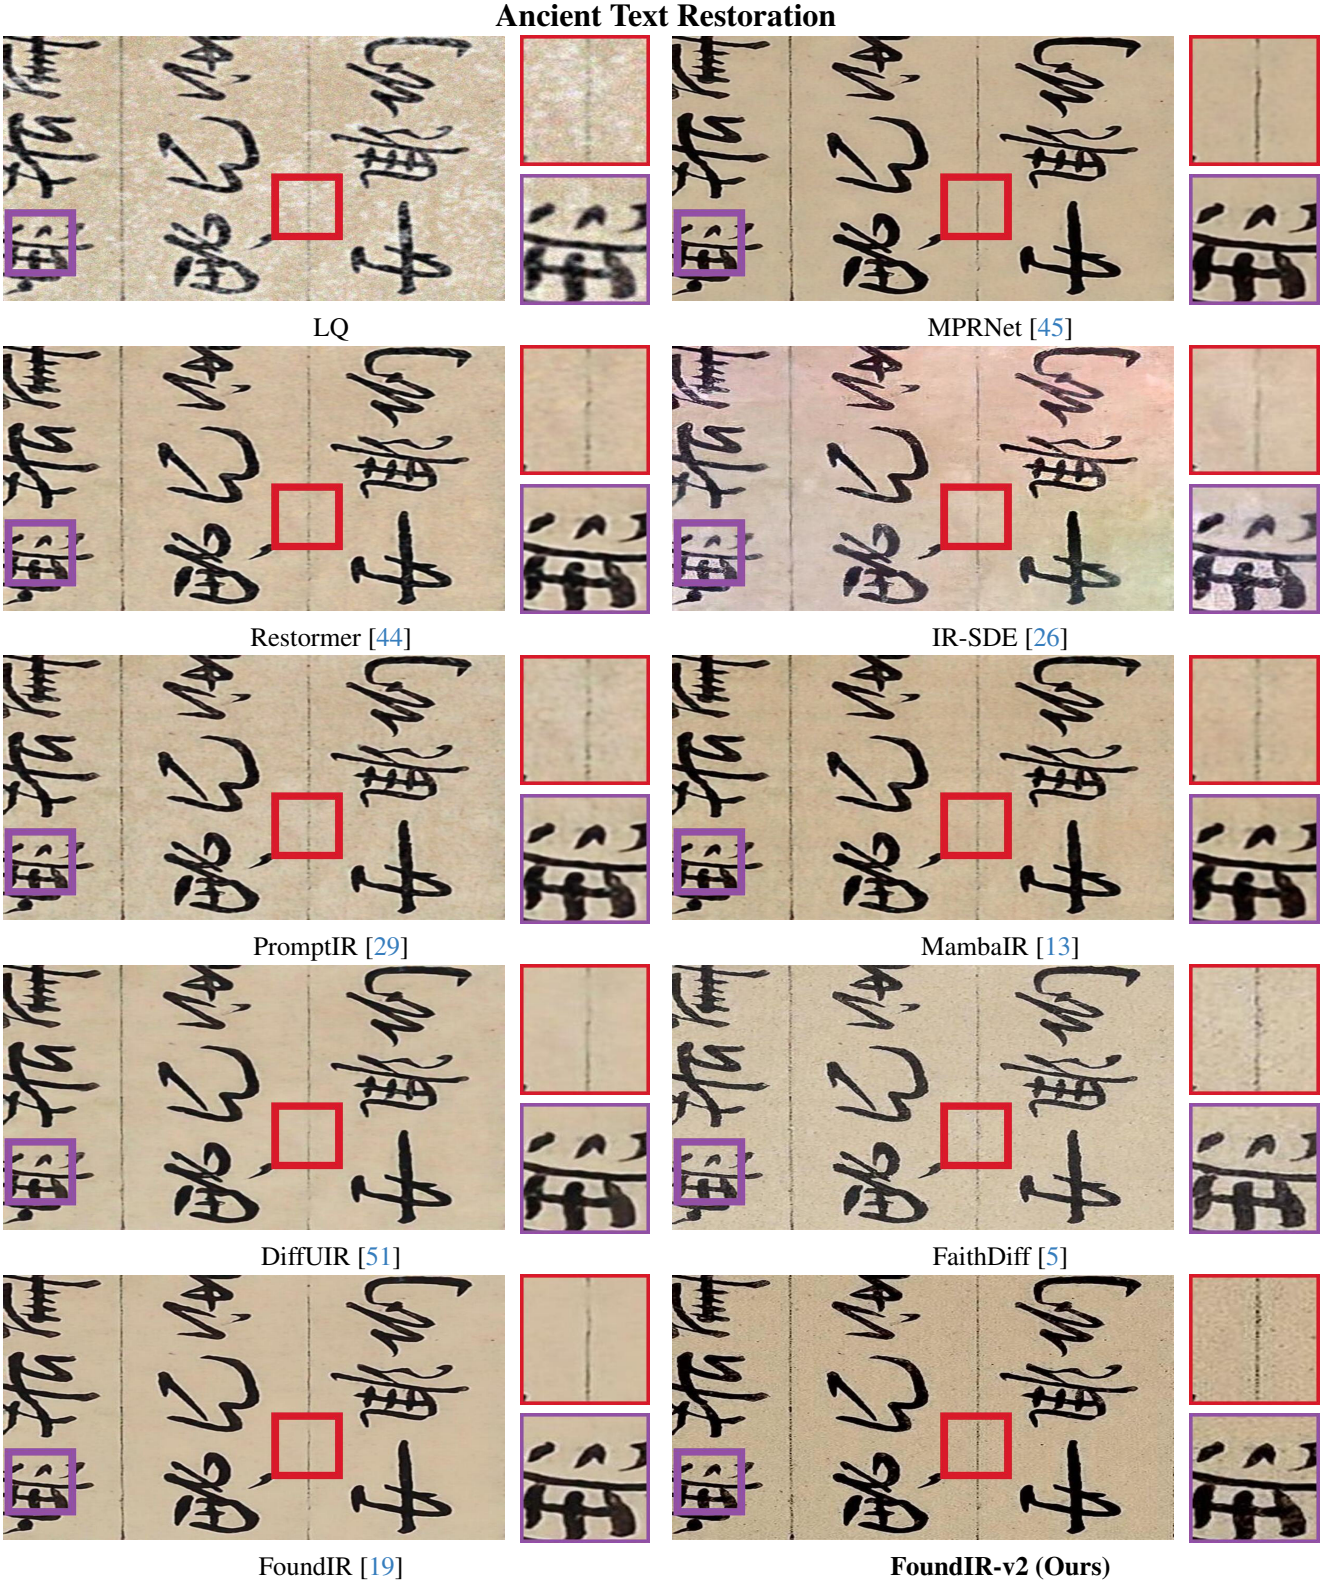

Figure 51. Qualitative comparison results. Compared to the results restored by other methods, our FoundIR-v2 generates a clearer image.

References

[1] Yuang Ai, Xiaoqiang Zhou, Huaibo Huang, Xiaotian Han, Zhengyu Chen, Quanzeng You, and Hongxia Yang. Dreamclear: High-capacity real-world image restoration with privacy-safe dataset curation. *NeurIPS*, 2024. 2, 3, 32, 33, 34, 35, 36, 37, 40, 45, 46, 47, 48, 49, 50

[2] Shuo Cao, Yihao Liu, Wenlong Zhang, Yu Qiao, and Chao Dong. Grids: Grouped multiple-degradation restoration with image degradation similarity. In *ECCV*, 2024. 2

[3] Hanting Chen, Yunhe Wang, Tianyu Guo, Chang Xu, Yiping Deng, Zhenhua Liu, Siwei Ma, Chunjing Xu, Chao Xu, and Wen Gao. Pre-trained image processing transformer. In *CVPR*, 2021. 2

[4] I Chen, Wei-Ting Chen, Yu-Wei Liu, Yuan-Chun Chiang, Sy-Yen Kuo, Ming-Hsuan Yang, et al. Unirestore: Unified perceptual and task-oriented image restoration model using diffusion prior. In *CVPR*, 2025. 2

[5] Junyang Chen, Jinshan Pan, and Jiangxin Dong. Faithdiff: Unleashing diffusion priors for faithful image super-resolution. In *CVPR*, 2025. 2, 32, 33, 34, 35, 36, 37, 38, 40, 41, 42, 45, 46, 47, 48, 49, 50, 51, 52, 53

[6] Wei-Ting Chen, Zhi-Kai Huang, Cheng-Che Tsai, Hao-Hsiang Yang, Jian-Jiun Ding, and Sy-Yen Kuo. Learning multiple adverse weather removal via two-stage knowledge learning and multi-contrastive regularization: Toward a unified model. In *CVPR*, 2022. 2

[7] Xiangyu Chen, Zheyuan Li, Yuandong Pu, Yihao Liu, Jiantao Zhou, Yu Qiao, and Chao Dong. A comparative study of image restoration networks for general backbone network design. In *ECCV*, 2024. 22, 23, 24

[8] Xiangyu Chen, Yihao Liu, Yuandong Pu, Wenlong Zhang, Jiantao Zhou, Yu Qiao, and Chao Dong. Learning a low-level vision generalist via visual task prompt. In *ACM MM*, 2024. 2

[9] Marcos V Conde, Gregor Geigle, and Radu Timofte. Instructir: High-quality image restoration following human instructions. In *ECCV*, 2024. 2, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 25, 26, 27, 28, 29, 30, 31, 39, 43, 44

[10] Yuning Cui, Syed Waqas Zamir, Salman Khan, Alois Knoll, Mubarak Shah, and Fahad Shahbaz Khan. Adair: Adaptive all-in-one image restoration via frequency mining and modulation. In *ICLR*, 2025. 2

[11] Huiyu Duan, Xiongkuo Min, Sijing Wu, Wei Shen, and Guangtao Zhai. Uniprocessor: a text-induced unified low-level image processor. In *ECCV*, 2024. 2

[12] Ben Fei, Zhaoyang Lyu, Liang Pan, Junzhe Zhang, Weidong Yang, Tianyue Luo, Bo Zhang, and Bo Dai. Generative diffusion prior for unified image restoration and enhancement. In *CVPR*, 2023. 2

[13] Hang Guo, Jinmin Li, Tao Dai, Zhihao Ouyang, Xudong Ren, and Shu-Tao Xia. Mambair: A simple baseline for image restoration with state-space model. In *ECCV*, 2024. 38, 41, 42, 51, 52, 53

[14] Yu Guo, Yuan Gao, Yuxu Lu, Huilin Zhu, Ryan Wen Liu, and Shengfeng He. Onerestore: A universal restoration framework for composite degradation. In *ECCV*, 2024. 2

[15] JiaKui Hu, Lujia Jin, Zhengjian Yao, and Yanye Lu. Universal image restoration pre-training via degradation classification. *arXiv preprint arXiv:2501.15510*, 2025. 2

[16] Yitong Jiang, Zhaoyang Zhang, Tianfan Xue, and Jinwei Gu. Autodir: Automatic all-in-one image restoration with latent diffusion. *arXiv preprint arXiv:2310.10123*, 2023. 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 25, 26, 27, 28, 29, 30, 31, 39, 43, 44

[17] Yitong Jiang, Zhaoyang Zhang, Tianfan Xue, and Jinwei Gu. Autodir: Automatic all-in-one image restoration with latent diffusion. In *ECCV*, 2024. 2

[18] Boyun Li, Xiao Liu, Peng Hu, Zhongqin Wu, Jiancheng Lv, and Xi Peng. All-in-one image restoration for unknown corruption. In *CVPR*, 2022. 2

[19] Hao Li, Xiang Chen, Jiangxin Dong, Jinhui Tang, and Jinshan Pan. Foundir: Unleashing million-scale training data to advance foundation models for image restoration. In *ICCV*, 2025. 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 38, 39, 41, 42, 43, 44, 51, 52, 53

[20] Ruoteng Li, Robby T Tan, and Loong-Fah Cheong. All in one bad weather removal using architectural search. In *CVPR*, 2020. 2

[21] Wenbo Li, Xin Lu, Shengju Qian, Jiangbo Lu, Xiangyu Zhang, and Jiaya Jia. On efficient transformer-based image pre-training for low-level vision. *arXiv preprint arXiv:2112.10175*, 2021. 2

[22] Xinqi Lin, Jingwen He, Ziyan Chen, Zhaoyang Lyu, Bo Dai, Fanghua Yu, Yu Qiao, Wanli Ouyang, and Chao Dong. Diffbir: Toward blind image restoration with generative diffusion prior. In *ECCV*, 2024. 32, 34, 36, 45, 47, 48, 49, 50

[23] Xinqi Lin, Fanghua Yu, Jinfan Hu, Zhiyuan You, Wu Shi, Jimmy S Ren, Jinjin Gu, and Chao Dong. Harnessing diffusion-yielded score priors for image restoration. *arXiv preprint arXiv:2507.20590*, 2025. 3

[24] Lin Liu, Lingxi Xie, Xiaopeng Zhang, Shanxin Yuan, Xiangyu Chen, Wengang Zhou, Houqiang Li, and Qi Tian. Tape: Task-agnostic prior embedding for image restoration. In *ECCV*, 2022. 2

[25] Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Controlling vision-language models

for universal image restoration. *arXiv preprint arXiv:2310.01018*, 2023. 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 39, 43, 44

[26] Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Image restoration with mean-reverting stochastic differential equations. *arXiv preprint arXiv:2301.11699*, 2023. 38, 41, 42, 51, 52, 53

[27] Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Controlling vision-language models for universal image restoration. In *ICLR*, 2025. 2

[28] Ozan Özdenizci and Robert Legenstein. Restoring vision in adverse weather conditions with patch-based denoising diffusion models. *IEEE TPAMI*, 2023. 2

[29] Vaishnav Potlapalli, Syed Waqas Zamir, Salman Khan, and Fahad Khan. Promptir: Prompting for all-in-one image restoration. In *NeurIPS*, 2023. 2, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 38, 39, 41, 42, 43, 44, 51, 52, 53

[30] Chu-Jie Qin, Rui-Qi Wu, Zikun Liu, Xin Lin, Chun-Le Guo, Hyun Hee Park, and Chongyi Li. Restore anything with masks: Leveraging mask image modeling for blind all-in-one image restoration. In *ECCV*, 2024. 2

[31] Sudarshan Rajagopalan, Nithin Gopalakrishnan Nair, Jay N Paranjape, and Vishal M Patel. Gendeg: Diffusion-based degradation synthesis for generalizable all-in-one image restoration. In *CVPR*, 2025. 2

[32] Jeya Maria Jose Valanarasu, Rajeev Yasarla, and Vishal M Patel. Transweather: Transformer-based restoration of images degraded by adverse weather conditions. In *CVPR*, 2022. 2, 22, 23, 24

[33] Chao Wang, Zhedong Zheng, Ruijie Quan, Yifan Sun, and Yi Yang. Context-aware pretraining for efficient blind image decomposition. In *CVPR*, 2023. 2

[34] Jianyi Wang, Zongsheng Yue, Shangchen Zhou, Kelvin CK Chan, and Chen Change Loy. Exploiting diffusion prior for real-world image super-resolution. *IJCV*, 2024. 2, 33, 35, 37, 40, 46

[35] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *CVPR*, 2021. 2, 32, 33, 34, 35, 36, 37, 40, 45, 46, 47, 48, 49, 50

[36] Gang Wu, Junjun Jiang, Kui Jiang, and Xianming Liu. Harmony in diversity: Improving all-in-one image restoration via multi-task collaboration. In *ACM MM*, 2024. 2

[37] Rongyuan Wu, Lingchen Sun, Zhiyuan Ma, and Lei Zhang. One-step effective diffusion network for real-world image super-resolution. In *NeurIPS*, 2024. 33, 35, 37, 40, 46

[38] Rongyuan Wu, Tao Yang, Lingchen Sun, Zhengqiang Zhang, Shuai Li, and Lei Zhang. Seesr: Towards semantics-aware real-world image super-resolution. In *CVPR*, 2024. 2, 32, 33, 34, 35, 36, 37, 40, 45, 46, 47, 48, 49, 50

[39] Tao Yang, Rongyuan Wu, Peiran Ren, Xuansong Xie, and Lei Zhang. Pixel-aware stable diffusion for realistic image super-resolution and personalized stylization. In *ECCV*, 2024. 2, 32, 33, 34, 35, 36, 37, 40, 45, 46, 47, 48, 49, 50

[40] Tian Ye, Sixiang Chen, Jinbin Bai, Jun Shi, Chenghao Xue, Jingxia Jiang, Junjie Yin, Erkang Chen, and Yun Liu. Adverse weather removal with codebook priors. In *ICCV*, 2023. 2

[41] Fanghua Yu, Jinjin Gu, Zheyuan Li, Jinfan Hu, Xiangtao Kong, Xintao Wang, Jingwen He, Yu Qiao, and Chao Dong. Scaling up to excellence: Practicing model scaling for photo-realistic image restoration in the wild. In *CVPR*, 2024. 2, 3, 32, 33, 34, 35, 36, 37, 40, 45, 46, 47, 48, 49, 50

[42] Zongsheng Yue, Jianyi Wang, and Chen Change Loy. Efficient diffusion model for image restoration by residual shifting. *IEEE TPAMI*, 2024. 2

[43] Eduard Zamfir, Zongwei Wu, Nancy Mehta, Yuedong Tan, Danda Pani Paudel, Yulun Zhang, and Radu Timofte. Complexity experts are task-discriminative learners for any image restoration. In *CVPR*, 2025. 2

[44] 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 *CVPR*, 2022. 38, 41, 42, 51, 52, 53

[45] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *CVPR*, 2021. 38, 41, 42, 51, 52, 53

[46] Haijin Zeng, Xiangming Wang, Yongyong Chen, Jingyong Su, and Jie Liu. Vision-language gradient descent-driven all-in-one deep unfolding networks. In *CVPR*, 2025. 2

[47] Cheng Zhang, Yu Zhu, Qingsen Yan, Jinqiu Sun, and Yanning Zhang. All-in-one multi-degradation image restoration network via hierarchical degradation representation. In *ACM MM*, 2023. 2

[48] Jinghao Zhang, Jie Huang, Mingde Yao, Zizheng Yang, Hu Yu, Man Zhou, and Feng Zhao. Ingredient-oriented multi-degradation learning for image restoration. In *CVPR*, 2023. 2

[49] Kai Zhang, Jingyun Liang, Luc Van Gool, and Radu Timofte. Designing a practical degradation model for deep blind image super-resolution. In *ICCV*, 2021. 2

[50] Kai Zhang, Jingyun Liang, Luc Van Gool, and Radu Timofte. Designing a practical degradation model for deep blind image super-resolution. In *ICCV*, 2021. 32, 34, 36, 45, 47, 48, 49, 50

[51] Dian Zheng, Xiao-Ming Wu, Shuzhou Yang, Jian Zhang, Jian-Fang Hu, and Wei-shi Zheng. Selective hourglass map-

ping for universal image restoration based on diffusion model. In *CVPR*, 2024. 2, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 38, 39, 41, 42, 43, 44, 51, 52, 53

[52] Kaiwen Zhu, Jinjin Gu, Zhiyuan You, Yu Qiao, and Chao Dong. An intelligent agentic system for complex image restoration problems. *arXiv preprint arXiv:2410.17809*, 2024. 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 39, 43, 44

[53] Yurui Zhu, Tianyu Wang, Xueyang Fu, Xuanyu Yang, Xin Guo, Jifeng Dai, Yu Qiao, and Xiaowei Hu. Learning weather-general and weather-specific features for image restoration under multiple adverse weather conditions. In *CVPR*, 2023. 2