Learning Hazing to Dehazing: Towards Realistic Haze Generation for Real-World Image Dehazing

Supplementary Material

A. Additional Experimental Results

This supplementary material provides additional experimental results, including further visual comparisons using the RTTS dataset [5], the ablation study of HazeGen, the analysis of the effectiveness of HazeGen, and the influence of hyperparameters in AccSamp.

A.1. Additional Visual Comparisons on RTTS

Further qualitative comparisons on the RTTS dataset [5] are presented in Figure 11, showcasing our method against recent state-of-the-art approaches, including D4 [7], RIDCP [6], and PTTD [3]. D4, RIDCP, and PTTD leave substantial residual haze and exhibit poor visual quality, including unpleasant color shifts. In contrast, our method achieves superior haze removal performance, delivering results with natural color restoration, improved visual coherence, and overall enhanced aesthetic appeal.

A.2. Ablation Study of HazeGen

Beyond quantitative results, we present visual comparisons of various ablation configurations against the complete HazeGen framework in Figure 12. Specifically, columns (b) and (c) depict results from HazeGen without the hybrid training or the blended sampling approach, respectively. Column (d) illustrates outcomes when both strategies are omitted, and column (e) provides results without HazeGen, where DiffDehaze is directly trained on the synthetic data provided by RIDCP [6]. These visual results clearly demonstrate that the full HazeGen framework consistently achieves the highest performance, providing superior dehazing quality and enhanced detail rendering, as exemplified by clearer tree textures in the second row. Additionally, the absence of the entire HazeGen framework results in notably inferior performance.

A.3. Effectiveness of HazeGen

To further illustrate the advantage of training data generated by HazeGen, we evaluate the performance of three models—MSBDN [4], NAFNet [1], and DiffDehaze—trained separately on synthetic data from RIDCP [6] and realistic data generated by HazeGen. As shown in Figure 13, models trained on HazeGen data consistently exhibit superior haze removal and enhanced visual quality compared to models trained on synthetic data.

A.4. Effect of Hyperparameters (ω and s) in AccSamp

This section analyzes two crucial parameters of AccSamp: ω and s. Parameter ω determines the number of refinement sampling steps, while s controls the strength of the fidelity guidance applied.

Effect of ω. Visual sampling results obtained using various ω values (1000, 800, 600, 400, and 200) are illustrated in Figure 14. Adjusting ω enables a trade-off between *detail quality* and *sampling speed*. Since inference time is approximately proportional to the number of sampling steps, reducing ω improves both sampling speed and image fidelity. The high-quality training data allows even small ω values to yield strong dehazing performance. Nevertheless, larger values of ω produce richer details due to increased sampling iterations. Empirically, we find setting ω around 600 provides the best balance between speed and detail quality.

Effect of s. Figure 15 shows sampling results with different values of fidelity guidance strength s (0.0, 0.1, 0.2, 0.5, and 1.0). Adjusting s provides control over the balance between *quality* and *fidelity*. Larger values of s enhance fidelity to the input image but can lead to less detailed or softer results compared to smaller values of s.

References

[1] Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. In *ECCV*, 2022. 1, 3

[2] Zeyuan Chen, Yangchao Wang, Yang Yang, and Dong Liu. Psd: Principled synthetic-to-real dehazing guided by physical priors. In *CVPR*, 2021. 2

[3] Zixuan Chen, Zewei He, Ziqian Lu, Xuecheng Sun, and Zhe-Ming Lu. Prompt-based test-time real image dehazing: A novel pipeline. In *ECCV*, 2024. 1, 2

[4] Hang Dong, Jinshan Pan, Lei Xiang, Zhe Hu, Xinyi Zhang, Fei Wang, and Ming-Hsuan Yang. Multi-scale boosted dehazing network with dense feature fusion. In *CVPR*, 2020. 1, 3

[5] Boyi Li, Wenqi Ren, Dengpan Fu, Dacheng Tao, Dan Feng, Wenjun Zeng, and Zhangyang Wang. Benchmarking single-image dehazing and beyond. *TIP*, pages 492–505, 2019. 1, 2

[6] Rui-Qi Wu, Zheng-Peng Duan, Chun-Le Guo, Zhi Chai, and Chongyi Li. Ridcp: Revitalizing real image dehazing via high-quality codebook priors. In *CVPR*, 2023. 1, 2, 3

[7] Yang Yang, Chaoyue Wang, Risheng Liu, Lin Zhang, Xiaojie Guo, and Dacheng Tao. Self-augmented unpaired image dehazing via density and depth decomposition. In *CVPR*, 2022. 1, 2

(a) Input (c) PSD [2] (b) D4 [7] (c) RIDCP [6] (d) PTTD [3] (e) Ours

Figure 11. More visual comparisons on the RTTS dataset [5]. Zoomed-in for details.

(a) Input (b) w/o hybrid (c) w/o blended (d) w/o both (e) w/o HazeGen (f) Full

Figure 12. Visual comparisons of ablation settings of HazeGen against the full version. Zoomed-in for details.

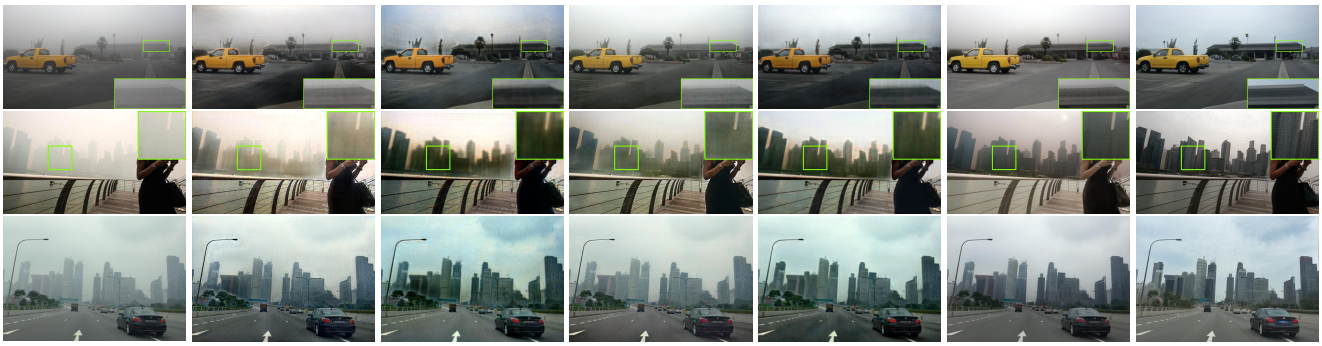

(a) Input Hazy Image | (b) MSBDN [4] (RIDCP [6]) | (c) MSBDN [4] (HazeGen) | (d) NAFNet [1] (RIDCP [6]) | (e) NAFNet [1] (HazeGen) | (f) Ours (RIDCP [6]) | (g) Ours (HazeGen)

Figure 13. Visual comparisons of dehazing models trained on data produced by RIDCP [6] and HazeGen. Zoomed-in for details.

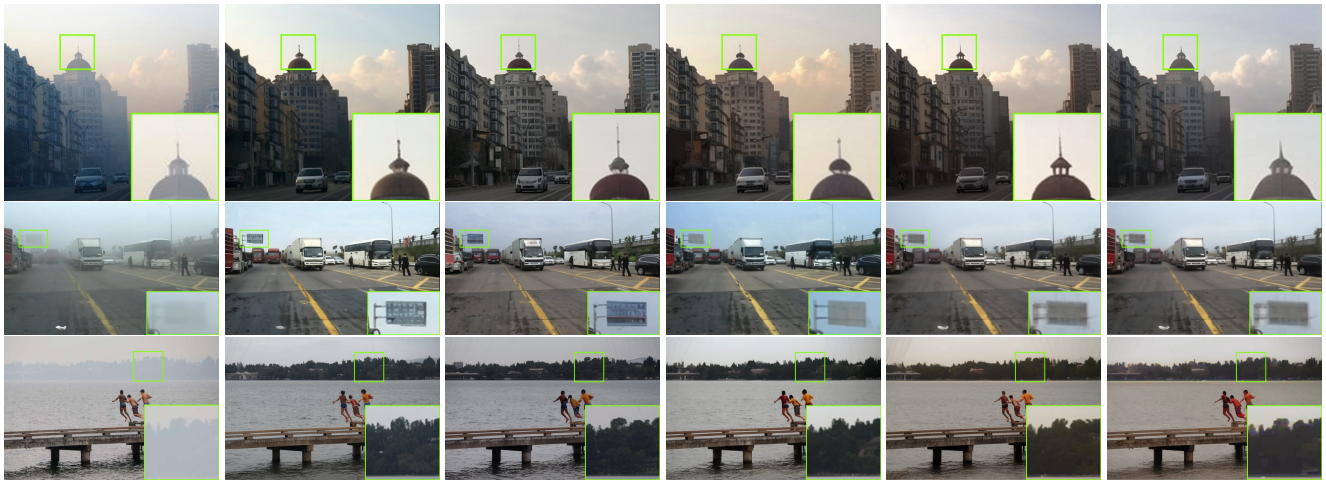

(a) Input | (b) $\omega = 1000$ | (c) $\omega = 800$ | (d) $\omega = 600$ | (e) $\omega = 400$ | (f) $\omega = 200$

Figure 14. Dehazing results with varying timestep ω in AccSamp. Zoomed-in for details.

(a) Input | (b) $s = 0.0$ | (c) $s = 0.1$ | (d) $s = 0.2$ | (e) $s = 0.5$ | (f) $s = 1.0$

Figure 15. Dehazing results with varying strength s of the fidelity guidance. Zoomed-in for details.