Degradation-Consistent Test-Time Adaptation for All-in-One Image Restoration

Supplementary Material

This supplemental material mainly contains:

- Sec. A provides detailed descriptions of the Residual Denoising Diffusion Model (RDDM) used to construct the test-time degradation generator.
- Additional quantitative results across multiple image restoration tasks, including deraining, denoising, dehazing, deblurring, and low-light enhancement, are reported in Sec. B.1.
- More visual comparisons for five restoration tasks are presented in Sec. B.2.
- Sec. B.3 provides additional analyses, including ablation studies, loss sensitivity analysis, limited-data evaluation, and source-domain performance after adaptation.
- Limitations of the proposed approach are discussed in Sec. C.

A. Details of the Residual Denoising Diffusion Model

As introduced in the main paper, we utilize a Residual Denoising Diffusion Model (RDDM) [7] to construct our test-time degradation generator. RDDM is a conditional diffusion-based framework designed for image restoration tasks such as deraining, shadow removal, and low-light enhancement. Unlike conventional diffusion models that operate solely on additive Gaussian noise, RDDM explicitly models the residual between the degraded image and its clean counterpart as part of the forward diffusion process.

Forward diffusion process. Given a degraded image I_{in} and its corresponding clean image I_0, RDDM defines the residual as:

$$I_{\text{res}} = I_0 - I_{\text{in}}. \tag{1}$$

During the forward diffusion process, the model perturbs I_0 with both structured residuals and random noise. The diffusion at time step t is formulated as:

$$I_t = I_0 + \bar{\alpha}_t I_{\text{res}} + \bar{\beta}_t \epsilon, \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \tag{2}$$

where $\bar{\alpha}_t = \sum_{i=1}^{t} \alpha_i$ and $\bar{\beta}_t = \sqrt{\sum_{i=1}^{t} \beta_i^2}$. The parameters α_t and β_t control the magnitude of the residual and noise components injected at each timestep, respectively.

In this work, we only use the residual prediction branch of RDDM. Specifically, we train a neural network $I_{\text{res}}^{\theta}(I_t, t, I_{\text{in}})$ to estimate the residual component during the reverse denoising process. The noise prediction branch ϵ_θ is not used in our framework.

Reverse diffusion process. In the reverse diffusion process, the generative transition at each timestep is modeled as $p_\theta(I_{t-1} \mid I_t)$. The reconstruction from I_t to I_{t-1} is carried out using the predicted residual as follows:

$$I_{t-1} = I_t - (\bar{\alpha}_t - \bar{\alpha}_{t-1}) I_{\text{res}}^{\theta} + \sigma_t \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \tag{3}$$

where the variance term σ_t^2 is computed by:

$$\sigma_t^2 = \eta \cdot \frac{\beta_t^2 \bar{\beta}_{t-1}^2}{\bar{\beta}_t^2}. \tag{4}$$

This reverse step allows the model to generate a re-degraded version of the input image conditioned on I_{in}, effectively aligning the simulated degradation with the distribution encountered at test time.

B. Experiments

B.1. Quantitative Results

We provide additional quantitative evaluations to further verify the robustness of our method across several image restoration tasks. Table 1 reports results on denoising, dehazing, and deraining using nine metrics (PSNR, SSIM [10], LPIPS [15], DISTS [3], CLIPIQA [9], NIQE [14], MUSIQ [6], MANIQA [11], FIQ [5]). The best scores are in **bold**, and models marked with "*" use conventional test-time augmentations with inverse transformations. Our method achieves consistently strong performance across all metrics. Table 2 extends the comparison to five tasks: denoising, dehazing, deraining, deblurring, and low-light enhancement.

B.2. Qualitative Results

Fig. 1 provides a visual comparison of five image restoration tasks: denoising, dehazing, deraining, deblurring, and low-light image enhancement. The datasets utilized include Kodak24 [4], O-HAZY [1], DenseHaze [2], Rain100H [12], RealBlur [8], and LOLv2-real [13]. It is evident from the figure that the incorporation of our proposed DCTTA method significantly enhances the performance of both AirNet and PromptIR. For instance, in the deraining task, the base model AirNet initially struggled to completely remove rain streaks, but with DCTTA, a noticeable reduction in rain artifacts is observed. In the deblurring task, the original PromptIR method introduced erroneous information in the top-left corner during restoration, which is effectively mitigated by the addition of DCTTA, resulting in a restored image that closely matches the Ground Truth (GT). These results demonstrate that DCTTA, as an effective image restoration enhancement technology, can markedly improve the performance of existing methods and achieve superior image restoration outcomes.

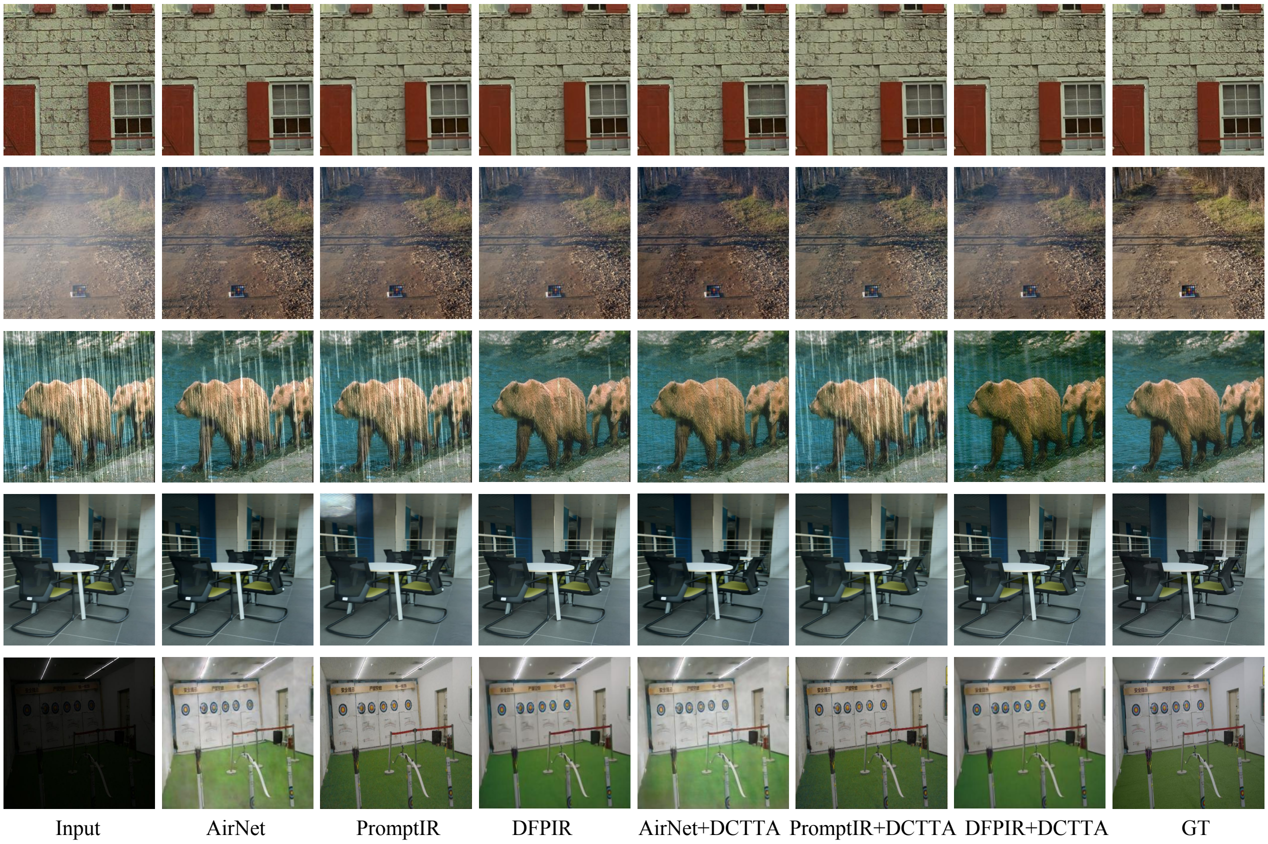

Figure 1. Visual comparison of five image restoration tasks: denoising, dehazing, deraining, deblurring and low-light image enhancement.

Table 1. Quantitative comparison across multiple restoration tasks: deraining, denoising and dehazing. Using Nine Evaluation Metrics (PSNR↑, SSIM↑, LPIPS↓, DISTS↓, CLIPIQA↑, NIQE↓, MUSIQ↑, MANIQA↑, FIQ↓). The best results are highlighted in **bold**. Models marked with "*" apply conventional test-time augmentations (e.g., rotation and flipping) followed by inverse transformations for the final outputs.

Method	Rain100H									Kodak24								
	PSNR↑	SSIM↑	LPIPS↓	DISTS↓	CLIPIQA↑	NIQE↓	MUSIQ↑	MANIQA↑	FIQ↓	PSNR↑	SSIM↑	LPIPS↓	DISTS↓	CLIPIQA↑	NIQE↓	MUSIQ↑	MANIQA↑	FIQ↓
AdaIR	15.79	0.492	0.489	0.287	0.730	3.150	68.54	0.678	196.59	25.43	0.868	0.185	0.166	0.509	8.430	65.13	0.670	132.95
VLU-Net	15.55	0.483	0.495	0.290	0.502	7.773	63.51	0.593	200.33	25.22	0.864	0.185	0.166	0.507	8.826	65.33	0.662	136.54
DACLIP	15.37	0.493	0.486	0.286	0.481	6.744	62.55	0.575	202.26	24.65	0.839	0.195	0.169	0.507	8.356	64.84	0.658	130.16
AirNet	14.62	0.490	0.504	0.298	0.485	6.306	61.24	0.555	226.02	25.57	0.862	0.170	0.143	0.520	7.385	63.57	0.650	100.99
AirNet*	15.14	0.515	**0.478**	**0.285**	0.481	**5.924**	61.83	**0.560**	**211.05**	25.71	0.863	0.168	0.145	**0.530**	7.469	**63.72**	**0.652**	100.45
AirNet + DCTTA	**17.34**	**0.531**	0.496	0.308	**0.485**	6.652	**62.15**	0.500	225.48	**25.79**	**0.863**	**0.166**	**0.132**	0.523	**6.516**	60.94	0.641	**91.76**
PromptIR	15.64	0.498	0.483	0.288	0.492	6.966	62.22	0.577	204.28	25.41	0.859	0.190	0.164	0.492	7.894	63.90	0.645	133.63
PromptIR*	15.72	0.504	0.480	0.286	0.484	7.160	**62.64**	**0.581**	199.10	25.48	0.860	0.188	0.167	0.487	8.113	**64.24**	**0.647**	133.69
PromptIR + DCTTA	**20.21**	**0.663**	**0.357**	**0.227**	**0.538**	**5.710**	62.54	0.541	**156.89**	**25.78**	**0.862**	**0.169**	**0.135**	**0.525**	**7.073**	62.04	0.643	**101.20**
DFPIR	15.94	0.503	0.472	0.285	0.500	7.012	63.00	0.579	197.93	24.57	0.861	0.195	0.174	0.553	8.920	65.38	**0.673**	138.48
DFPIR*	16.35	0.524	0.459	0.274	0.510	7.022	**63.58**	**0.593**	188.07	24.62	0.862	0.195	0.176	0.550	8.995	**65.69**	0.673	137.71
DFPIR + DCTTA	**20.27**	**0.630**	**0.359**	**0.223**	**0.545**	**5.377**	63.31	0.574	**151.51**	**24.66**	**0.875**	**0.148**	**0.124**	**0.614**	**6.712**	62.79	0.657	**82.03**

Method	O-HAZY									Dense_Haze								
	PSNR↑	SSIM↑	LPIPS↓	DISTS↓	CLIPIQA↑	NIQE↓	MUSIQ↑	MANIQA↑	FIQ↓	PSNR↑	SSIM↑	LPIPS↓	DISTS↓	CLIPIQA↑	NIQE↓	MUSIQ↑	MANIQA↑	FIQ↓
AdaIR	16.54	0.709	0.363	0.275	0.416	4.093	54.11	0.671	154.81	11.65	0.388	0.756	0.490	0.303	9.475	42.07	0.572	258.09
VLU-Net	16.61	0.714	0.357	0.273	0.421	4.032	55.55	0.677	154.31	11.63	0.386	0.753	0.486	0.314	9.186	42.39	0.580	254.47
DACLIP	16.61	0.714	0.366	0.264	0.446	3.972	52.69	0.671	151.96	11.11	0.348	0.770	0.482	0.325	8.895	39.98	0.585	255.14
AirNet	16.70	0.713	0.366	0.273	0.441	4.065	55.12	0.673	160.07	10.90	0.360	0.770	0.496	**0.331**	8.545	39.94	0.491	288.09
AirNet*	16.71	0.713	0.363	0.273	0.433	4.138	**55.33**	**0.678**	**157.57**	10.91	0.361	**0.761**	**0.495**	0.330	9.297	**40.93**	**0.546**	**269.95**
AirNet + DCTTA	**16.95**	**0.723**	**0.350**	**0.273**	**0.460**	**4.049**	54.26	0.658	165.21	**12.49**	**0.377**	0.798	0.503	0.303	**8.498**	40.40	0.452	310.94
PromptIR	16.70	0.715	0.358	0.271	**0.433**	4.058	**55.39**	0.678	**155.27**	11.48	0.376	**0.755**	0.486	0.313	9.426	41.41	0.568	260.52
PromptIR*	16.71	0.715	**0.357**	0.271	0.432	4.044	55.39	0.679	155.35	11.47	**0.377**	0.755	0.488	0.308	9.479	**41.51**	0.569	**257.46**
PromptIR + DCTTA	**16.78**	**0.715**	0.359	**0.271**	0.431	**3.968**	55.13	**0.680**	156.81	**11.65**	0.377	0.766	**0.485**	**0.320**	**9.405**	41.21	**0.571**	258.93
DFPIR	16.72	0.717	**0.354**	0.271	0.427	**3.989**	54.84	0.675	**153.58**	11.61	**0.381**	**0.754**	**0.489**	**0.318**	**9.231**	42.24	**0.574**	259.79
DFPIR*	16.81	**0.718**	0.356	0.271	0.427	3.998	55.21	0.677	154.18	11.52	0.380	0.756	0.493	0.313	9.565	**42.42**	0.572	**257.60**
DFPIR + DCTTA	**16.84**	0.718	0.356	**0.271**	**0.435**	4.007	**55.33**	**0.678**	155.16	**11.73**	0.368	0.789	0.537	0.295	11.200	40.32	0.537	277.04

Table 2. Quantitative comparison across five restoration tasks: deraining, denoising, dehazing, low-light enhancement and deblurring.

Method	Rain100H								
	PSNR↑	SSIM↑	LPIPS↓	DISTS↓	CLIPIQA↑	NIQE↓	MUSIQ↑	MANIQA↑	FIQ↓
AdaIR	15.57	0.496	0.485	0.285	0.503	7.286	63.15	0.590	200.72
VLU-Net	16.31	0.517	0.466	0.274	0.527	6.889	63.75	0.595	190.32
DiffUIR	14.26	0.453	0.530	0.311	0.449	8.076	62.93	0.577	222.49
AirNet	14.49	0.494	0.506	**0.305**	**0.450**	**5.710**	60.05	**0.523**	**232.68**
AirNet+DCTTA	**18.96**	**0.555**	**0.504**	0.328	0.372	7.974	**62.28**	0.414	244.13
PromptIR	13.29	0.404	0.586	0.363	0.375	9.390	**60.32**	**0.513**	273.43
PromptIR+DCTTA	**14.22**	**0.442**	**0.549**	**0.352**	**0.386**	**7.956**	58.95	0.465	**261.00**
DFPIR	16.27	0.523	0.459	0.275	**0.511**	6.618	**62.80**	**0.577**	193.73
DFPIR+DCTTA	**20.13**	**0.649**	**0.352**	**0.226**	0.474	**5.208**	61.79	0.546	**152.90**

Method	Kodak24								
	PSNR↑	SSIM↑	LPIPS↓	DISTS↓	CLIPIQA↑	NIQE↓	MUSIQ↑	MANIQA↑	FIQ↓
AdaIR	25.52	0.872	0.182	0.160	0.508	8.247	64.84	0.667	125.78
VLU-Net	25.04	0.872	0.182	0.162	0.510	8.568	64.76	0.667	125.08
DiffUIR	23.92	0.800	0.204	0.163	0.603	7.758	59.75	0.624	113.63
AirNet	26.68	0.875	0.138	0.125	0.575	6.522	**63.83**	**0.656**	**76.85**
AirNet+DCTTA	**26.86**	**0.877**	**0.131**	**0.121**	**0.592**	**6.121**	63.61	0.643	77.84
PromptIR	25.09	0.822	0.232	0.184	0.480	**8.194**	59.54	**0.616**	154.54
PromptIR+DCTTA	**25.21**	**0.822**	**0.231**	**0.183**	**0.481**	8.207	**59.75**	0.613	**149.76**
DFPIR	24.82	0.851	0.196	0.172	0.558	8.655	**64.57**	**0.664**	139.11
DFPIR+DCTTA	**24.95**	**0.868**	**0.164**	**0.139**	**0.610**	**7.235**	62.66	0.657	**96.58**

Method	O-HAZY								
	PSNR↑	SSIM↑	LPIPS↓	DISTS↓	CLIPIQA↑	NIQE↓	MUSIQ↑	MANIQA↑	FIQ↓
AdaIR	16.42	0.704	0.365	0.277	0.418	4.091	54.31	0.671	154.54
VLU-Net	16.56	0.710	0.358	0.274	0.423	4.013	55.55	0.676	156.01
DiffUIR	16.91	0.726	0.356	0.267	0.439	3.978	55.58	0.680	154.84
AirNet	15.69	0.688	0.381	0.285	**0.464**	4.261	**54.58**	**0.661**	**164.99**
AirNet+DCTTA	**15.69**	**0.689**	**0.380**	**0.285**	0.460	**4.237**	54.50	0.660	165.58
PromptIR	16.80	0.715	0.361	0.272	**0.429**	**4.035**	54.75	0.676	**155.46**
PromptIR+DCTTA	**16.81**	**0.715**	**0.361**	**0.272**	0.427	4.036	**54.75**	**0.676**	155.62
DFPIR	16.78	**0.716**	0.359	0.272	**0.431**	4.049	**55.01**	0.674	**154.08**
DFPIR+DCTTA	**16.93**	0.715	**0.357**	**0.271**	0.427	**4.038**	54.75	**0.676**	154.91

Method	LOL-v2-real								
	PSNR↑	SSIM↑	LPIPS↓	DISTS↓	CLIPIQA↑	NIQE↓	MUSIQ↑	MANIQA↑	FIQ↓
AdaIR	28.76	0.931	0.072	0.107	0.356	7.162	55.94	0.674	65.74
VLU-Net	25.09	0.907	0.090	0.119	0.365	6.344	51.14	0.644	63.53
DiffUIR	13.77	0.582	0.348	0.222	0.545	4.904	66.54	0.639	100.77
AirNet	14.36	0.754	0.236	**0.198**	**0.342**	**3.628**	**59.97**	0.525	**120.38**
AirNet+DCTTA	**17.97**	**0.798**	**0.221**	0.200	0.313	6.025	48.61	**0.572**	133.97
PromptIR	20.47	0.715	0.217	0.182	0.327	5.806	42.00	0.589	104.38
PromptIR+DCTTA	**20.67**	**0.723**	**0.211**	**0.177**	**0.335**	**5.747**	**42.69**	**0.593**	**102.51**
DFPIR	28.30	0.930	0.082	0.115	0.361	7.137	56.64	0.673	68.90
DFPIR+DCTTA	**28.46**	**0.930**	**0.079**	**0.112**	**0.372**	**6.982**	**57.18**	**0.678**	**67.51**

Method	RealBlur_J								
	PSNR↑	SSIM↑	LPIPS↓	DISTS↓	CLIPIQA↑	NIQE↓	MUSIQ↑	MANIQA↑	FIQ↓
AdaIR	16.94	0.619	0.300	0.177	0.213	5.576	37.06	0.435	49.84
VLU-Net	14.55	0.546	0.340	0.194	0.198	5.214	36.68	0.412	57.06
DiffUIR	25.86	0.796	0.200	0.144	0.217	5.727	41.39	0.473	33.70
AirNet	17.54	0.653	0.326	0.201	0.200	**5.654**	38.09	0.400	**64.97**
AirNet+DCTTA	**20.33**	**0.736**	**0.282**	**0.174**	**0.218**	5.740	**38.30**	**0.427**	93.20
PromptIR	21.81	0.695	0.272	0.161	0.205	**5.379**	38.05	0.440	79.79
PromptIR+DCTTA	**21.89**	**0.698**	**0.270**	**0.160**	**0.205**	5.427	**38.14**	**0.444**	**78.92**
DFPIR	26.86	0.829	0.175	0.126	0.222	**5.682**	44.27	0.509	**28.36**
DFPIR+DCTTA	**27.60**	**0.847**	**0.153**	**0.112**	**0.246**	5.757	**47.10**	**0.544**	46.62

B.3. Ablation Study

Effectiveness of DA, EMA and TIPS Modules. Table 3 provides a supplementary analysis for the main-text section "Effectiveness of the DA, EMA, and TIPS Modules." Consistent trends are observed across different restoration tasks, further demonstrating the effectiveness of the proposed components.

Table 3. Ablation study on key components of our method across different restoration tasks.

Task	**Metric**	PromptIR	w/o TIPS	w/o DA	w/o EMA	**Ours**
Deraining	PSNR↑	15.64	14.24	19.57	19.87	**20.21**
	SSIM↑	0.498	0.487	0.606	0.609	**0.663**
Denoising	PSNR↑	25.41	25.27	25.67	25.61	**25.78**
	SSIM↑	0.859	0.855	0.857	0.859	**0.862**
Dehazing	PSNR↑	16.70	16.66	16.70	16.73	**16.78**
	SSIM↑	0.715	0.702	0.711	0.715	**0.715**

Loss balancing. In TTA, GT supervision is unavailable, making it necessary to rely on carefully designed self-supervised objectives. Although multiple loss terms are involved, they can be grouped into two objectives: the adaptation loss L_s, and the stabilization loss L_a. Other weights λ are used only to align loss scales rather than for fine-grained tuning. The key hyperparameter is α, as defined in Eq. (10) of the main paper. We further provide a sensitivity analysis of α using the PromptIR backbone with DCTTA enabled, as shown in Table 4. The best trade-off is achieved at $\alpha = 1$.

Table 4. Sensitivity analysis of the loss balancing parameter α on the PromptIR backbone with DCTTA.

α	0.1	0.5	1	2
Rain100H	16.76 / 0.638	18.87 / 0.657	20.21 / 0.663	19.85 / 0.661

Flexibility. The degradation generator does not require pretraining and instead performs fast degradation learning at test time using target-domain images. With a few diffusion steps, it generates re-degraded samples to adapt the pretrained restoration model. To evaluate performance under limited data, we test on Rain100H using a single image, 10%, and 20% of the test set, as shown in Table 5. PromptIR+DCTTA achieves PSNR gains of 0.02 dB, 0.23 dB, and 0.70 dB over PromptIR, respectively, showing that the method remains effective even with very few test samples.

Source-domain performance. We report the performance of the adapted PromptIR on the source-domain benchmarks, as shown in Table 6. A certain performance drop is observed, which is expected since TTA performs domain-specific adjustments toward the target degradation. Once the model is adapted to a new degradation domain, its per-

Table 5. Effectiveness of DCTTA under limited test samples on Rain100H.

Method	Single Image	10%	20%
PromptIR	17.65/0.445	16.53/0.512	15.49/0.455
PromptIR+DCTTA	17.67/0.446	16.76/0.523	16.19/0.482

Table 6. Source-domain performance of PromptIR after test-time adaptation.

Method	$BSD68_{15}$	$BSD68_{25}$	$BSD68_{50}$	Rain100L	SOTS
PromptIR	33.98/0.933	31.31/0.888	28.06/0.799	36.37/0.972	30.58/0.974
PromptIR+DCTTA	33.75/0.928	31.13/0.883	27.93/0.796	31.11/0.949	30.26/0.977

formance on the original distribution may be affected. Our work focuses on TTA rather than continual learning, and thus does not enforce cross-domain performance preservation.

C. Limitation

From Table 5, it can be observed that although the proposed method still improves performance under limited data, the performance gain becomes noticeably smaller compared to using the full dataset. This indicates that the effectiveness of the proposed adaptation strategy may depend on the availability of sufficient target-domain samples. In future work, we plan to further investigate more data-efficient adaptation strategies that can achieve strong performance with only a small number of target images.

References

[1] Codruta O Ancuti, Cosmin Ancuti, Radu Timofte, and Christophe De Vleeschouwer. O-haze: a dehazing benchmark with real hazy and haze-free outdoor images. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 754–762, 2018. 1

[2] Codruta O Ancuti, Cosmin Ancuti, Mateu Sbert, and Radu Timofte. Dense-haze: A benchmark for image dehazing with dense-haze and haze-free images. In *2019 IEEE international conference on image processing (ICIP)*, pages 1014–1018. IEEE, 2019. 1

[3] Keyan Ding, Kede Ma, Shiqi Wang, and Eero P Simoncelli. Image quality assessment: Unifying structure and texture similarity. *IEEE transactions on pattern analysis and machine intelligence*, 44(5):2567–2581, 2020. 1

[4] R. Franzen. Kodak lossless true color image suite. http://r0k.us/graphics/kodak/, 1999. Accessed: 2025-07-18. 1

[5] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, 2017. 1

[6] Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5148–5157, 2021. 1

[7] Jiawei Liu, Qiang Wang, Huijie Fan, Yinong Wang, Yandong Tang, and Liangqiong Qu. Residual denoising diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2773–2783, 2024. 1

[8] Jaesung Rim, Haeyun Lee, Jucheol Won, and Sunghyun Cho. Real-world blur dataset for learning and benchmarking deblurring algorithms. In *European conference on computer vision*, pages 184–201. Springer, 2020. 1

[9] Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring clip for assessing the look and feel of images. In *Proceedings of the AAAI conference on artificial intelligence*, number 2, pages 2555–2563, 2023. 1

[10] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 1

[11] Sidi Yang, Tianhe Wu, Shuwei Shi, Shanshan Lao, Yuan Gong, Mingdeng Cao, Jiahao Wang, and Yujiu Yang. Maniqa: Multi-dimension attention network for no-reference image quality assessment. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1191–1200, 2022. 1

[12] Wenhan Yang, Robby T Tan, Jiashi Feng, Jiaying Liu, Zongming Guo, and Shuicheng Yan. Deep joint rain detection and removal from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1357–1366, 2017. 1

[13] Wenhan Yang, Wenjing Wang, Haofeng Huang, Shiqi Wang, and Jiaying Liu. Sparse gradient regularized deep retinex network for robust low-light image enhancement. *TIP*, 2021. 1

[14] Lin Zhang, Lei Zhang, and Alan C Bovik. A feature-enriched completely blind image quality evaluator. *IEEE Transactions on Image Processing*, 24(8):2579–2591, 2015. 1

[15] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 1