CDTNet: Compositional Degradation Transfer Network for Test-Time All-in-One De-weathering Adaptation

Kuan-Hung Lin[1] Fu-Jen Tsai[2] Jin-Ting He[1] Yen-Yu Lin[1]
[1]National Yang Ming Chiao Tung University [2]MediaTek Inc.
guan40214@gmail.com fu-jen.tsai@mediatek.com
jinting.cs12@nycu.edu.tw lin@cs.nycu.edu.tw

Abstract

All-in-one adverse weather image restoration aims to handle diverse weather degradations, such as rain, haze, and snow, within a unified model. However, existing methods often struggle with composite weather conditions and domain gaps between training and testing data. To address these limitations, we propose the Compositional Degradation Transfer Network (CDTNet), a domain adaptation framework based on compositional degradation transfer. Specifically, CDTNet decomposes target-domain degraded patterns into weather-specific residuals corresponding to rain, haze, and snow, which are then composed and transferred onto source-domain clean images to generate paired data for model adaptation. CDTNet employs weather-specific Degradation Residual Extractors (DREs) with specialized adapters, degradation tokens, and Spatial-Channel Coupled Attention to explicitly model individual degradation components. Moreover, we propose a Degradation Decomposition Loss (DDL) that encourages each DRE to capture its corresponding degradation while suppressing irrelevant residuals. Experimental results demonstrate that CDTNet consistently improves multiple restoration backbones across benchmark datasets and outperforms existing domain adaptation methods, with substantial gains under composite weather conditions. The source code is available at https://github.com/homer40214/CDTNet.

1. Introduction

Images captured under adverse weather conditions, such as rain, haze, and snow, often suffer from degraded visual quality and reduced scene visibility, adversely affecting downstream vision systems. Previous studies have developed specialized methods for deraining [4, 8, 9, 16, 18, 19, 22, 37, 39], dehazing [6, 7, 12, 15, 25, 31, 33, 42, 45, 47], and desnowing [1, 2, 26, 40, 46]. However, deploying a separate model for each degradation is inefficient and requires prior knowledge of the degradation type.

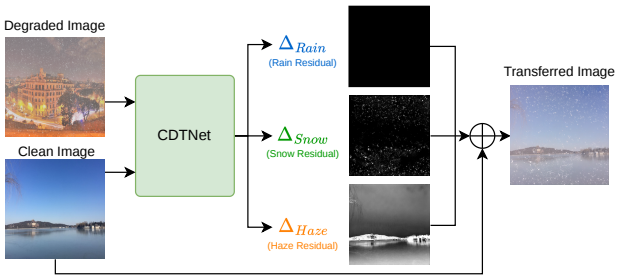


Figure 1. **Compositional Degradation Transfer.** CDTNet decomposes target-domain degradations into rain, haze, and snow residuals and composes them with source-domain clean images to synthesize individual or composite weather degradations.

To address these limitations, all-in-one image restoration methods handle multiple degradations within a unified model. Existing approaches commonly employ degradation-aware representations, prompts, or adaptive feature modulation to distinguish and handle different degradation types [3, 5, 21, 30, 35, 36]. Despite these advances, existing all-in-one methods still face two challenges. First, multiple weather degradations may coexist, resulting in composite weather conditions. Second, real-world degradations often differ from training data, leading to a domain gap that limits model generalization.

Recent studies have explored domain adaptation to improve the generalization of all-in-one adverse weather image restoration [17, 24]. Rather than recovering clean content from target-domain degraded images, these methods transfer their degradation patterns onto source-domain clean images to generate paired data for domain adaptation. However, existing methods transfer target-domain degradations without explicitly modeling individual degradation components. This motivates us to explore compositional degradation transfer, where different degradations are separately modeled and combined to better handle composite weather conditions.

In this paper, we propose the **Compositional Degradation Transfer Network (CDTNet)**, a domain adaptation

framework based on compositional degradation transfer. As shown in Figure 1, CDTNet takes a target-domain degraded image and a source-domain clean image as unpaired inputs. It decomposes the target-domain degradation into weather-specific residuals corresponding to rain, haze, and snow. These residuals are then composed according to the target weather conditions and transferred onto the source-domain clean image. By transferring only the degradation patterns, CDTNet retains the source image as the clean reference without requiring target-domain clean images. This process generates paired training samples for adapting restoration models during testing.

To achieve compositional degradation transfer, CDTNet employs weather-specific Degradation Residual Extractors (DREs) to capture individual degradation patterns. Each DRE integrates a weather-specific adapter and weather-specific degradation tokens with Spatial-Channel Coupled Attention (SCCA) to capture degradation features across spatial and channel dimensions. A Shared Residual Decoder (SRD) then decodes the extracted features into their corresponding degradation residuals, which are composed to synthesize individual or composite degradations. Moreover, we introduce a Degradation Decomposition Loss (DDL) to ensure that each DRE focuses on its corresponding degradation while suppressing irrelevant residuals. This enables rain, haze, and snow residuals to be separately modeled and composed according to the target weather conditions. Experimental results demonstrate that CDTNet consistently improves multiple restoration models across benchmark datasets and outperforms existing domain adaptation methods, particularly under composite weather conditions. Our contributions are summarized as follows:

- We propose CDTNet, which decomposes and transfers target-domain weather degradations to generate domain-adaptive training pairs for fine-tuning restoration models during testing.
- We propose weather-specific Degradation Residual Extractors (DREs) and a Degradation Decomposition Loss (DDL) to separately extract rain, haze, and snow residuals for compositional degradation transfer.
- Extensive experiments demonstrate that CDTNet consistently improves multiple restoration models and outperforms existing adaptation methods, particularly under composite weather conditions.

2. Related Work

2.1. All-in-One Image Restoration

All-in-one image restoration handles multiple degradations within a unified model without prior knowledge of the degradation type. Existing methods typically incorporate degradation-aware representations, prompts, or adaptive feature modulation to accommodate diverse degradations [3, 5, 13, 21, 27, 28, 30, 35, 36]. For example, PromptIR [30] conditions the restoration process on degradation-specific prompts, while AdaIR [5] employs frequency mining and adaptive modulation to capture degradation-dependent characteristics. DFPIR [35] modulates restoration features based on estimated degradation information, enabling adaptive processing across different degradation types. However, these methods rely heavily on source-domain training data and may generalize poorly to unseen target-domain degradations. Therefore, recent studies explore domain adaptation to adapt pretrained restoration models to target domains.

2.2. Domain Adaptation for Restoration

Domain adaptation aims to reduce domain gap between source and target domains. Existing restoration methods employ image translation or feature alignment to transfer knowledge across domains, often focusing on a single degradation [29, 32, 41]. Additionally, recent approaches utilize unlabeled target-domain images to adapt restoration models through self-supervised objectives [10, 14, 23, 34, 38]. Gou et al. [10] adapts a degradation adapter to each test image during diffusion denoising and uses it to guide the restoration process. However, these methods lack explicit paired supervision from the target domain, which may limit their adaptation effectiveness.

Alternatively, recent methods synthesize pseudo degraded-clean pairs by transferring degradation patterns from unlabeled target-domain images onto source-domain clean images, enabling supervised adaptation with reliable clean references. WTNet [17] introduces the concept of weather degradation transfer and employs physics-based modeling to transfer degradation patterns from test images onto source-domain clean images, constructing domain-adaptive training pairs for fine-tuning all-in-one restoration models. DDTNet [24] introduces a learning-based degradation disentanglement mechanism to separate degradation information from scene content before transferring it onto source-domain clean images. While DDTNet enables degradation transfer for domain adaptation, it does not explicitly model the composition of different weather degradations, limiting its ability to represent composite weather conditions. In contrast, CDTNet decomposes target-domain degradations into weather-specific rain, haze, and snow residuals, which are then composed to synthesize individual or composite degradations. This compositional formulation enables more flexible degradation transfer across diverse target weather conditions.

3. Proposed Method

We propose the Compositional Degradation Transfer Network (CDTNet), a domain adaptation framework that generates degraded-clean pairs aligned with target-domain

degradation patterns through compositional degradation transfer. As shown in Fig. 2, CDTNet consists of three main components. First, weather-specific Degradation Residual Extractors (DREs) integrate weather-specific adapters, degradation tokens, and Spatial-Channel Coupled Attention (SCCA) to capture degradation features across spatial and channel dimensions. Second, a Shared Residual Decoder (SRD) decodes these features into weather-specific degradation residuals, which are composed to generate degraded images with individual or composite weather conditions. Third, a Degradation Decomposition Loss (DDL) encourages each DRE to capture its corresponding degradation while suppressing irrelevant residuals.

Specifically, given a target-domain degraded image $I^d \in \mathbb{R}^{H\times W\times 3}$ and a source-domain clean image $I^c \in \mathbb{R}^{H\times W\times 3}$, CDTNet extracts degradation patterns from I^d and transfers them onto I^c. This degradation transfer process can be formulated as

$$I^o = \mathrm{CDTNet}(I^c, I^d), \tag{1}$$

where $I^o \in \mathbb{R}^{H\times W\times 3}$ preserves the scene content of I^c while incorporating the target-domain degradation patterns from I^d. The generated image I^o and its clean source I^c form paired data (I^o, I^c), which is used to fine-tune the all-in-one restoration model with its original supervised objective. This enables the restoration model to adapt to unseen target-domain degradations without modifying its architecture. The components of CDTNet are detailed in the following subsections.

3.1. Weather-Specific Degradation Residual Extractors

As shown in Fig. 2, CDTNet first employs a Degradation Encoder (DE) and a Content Encoder (CE) to extract multi-scale degradation features $F_i^d \in \mathbb{R}^{H_i\times W_i\times C_i}$ from I^d and content features $F_i^c \in \mathbb{R}^{H_i\times W_i\times C_i}$ from I^c as

$$\{F_i^d\}_{i=0}^2 = \mathrm{DE}(I^d), \quad \{F_i^c\}_{i=0}^2 = \mathrm{CE}(I^c), \tag{2}$$

where, for each scale $i \in \{0,1,2\}$, $H_i = H/2^i$, $W_i = W/2^i$, and $C_i = 32\cdot 2^i$. The degradation features F_i^d encode the weather degradation of the target-domain degraded image, whereas the content features F_i^c represent the scene content of the source-domain clean image.

Next, the degradation features F_i^d are fed into three weather-specific DREs, each dedicated to rain, haze, or snow. Let $\mathcal{G} = \{\mathrm{rain}, \mathrm{haze}, \mathrm{snow}\}$ denote the set of weather types. Within each DRE, a weather-specific adapter $A_{i,g}$, where $g \in \mathcal{G}$, generates adapted features $\bar{F}_{i,g}^d \in \mathbb{R}^{H_i\times W_i\times C_i}$ as

$$\bar{F}_{i,g}^d = A_{i,g}(F_i^d) = F_i^d + \mathrm{Conv}_{1\times 1}\big(\mathrm{GELU}\big(\mathrm{Conv}_{1\times 1}(F_i^d)\big)\big), \tag{3}$$

Spatial-Channel Coupled Attention. Next, each adapted feature $\bar{F}_{i,g}^d$ is fed into a Spatial-Channel Coupled Attention (SCCA) module. SCCA contains a spatial branch and a channel branch that capture complementary structural and appearance characteristics, respectively. First, the spatial branch reshapes $\bar{F}_{i,g}^d$ into a sequence of spatial tokens to generate spatial tokens $\bar{F}_{i,g}^{d,s} \in \mathbb{R}^{H_iW_i\times C_i}$. These spatial tokens facilitate the modeling of spatially varying degradation patterns, such as rain streaks, snow particles, and nonuniform haze. Additionally, the channel branch pools $\bar{F}_{i,g}^d$ to smaller resolutions, then flattens and transposes the pooled features to generate channel tokens $\bar{F}_{i,g}^{d,c} \in \mathbb{R}^{C_i\times P_i^2}$, where $P_i \in \{32, 24, 16\}$ for each i as

$$\bar{F}_{i,g}^{d,c} = \mathrm{transpose}\big(\mathrm{reshape}\big(\mathrm{Pool}_{P_i}(\bar{F}_{i,g}^d)\big)\big). \tag{4}$$

These channel tokens facilitate the modeling of channel-wise degradation characteristics.

For both the spatial and channel branches, weather-specific Degradation Information Tokens (DITs) interact with the corresponding features through attention to capture degradation-specific information. For the spatial branch, SCCA performs

$$\widetilde{T}_{i,g}^s = T_{i,g}^s + \mathrm{Attn}\Big(T_{i,g}^s, \bar{F}_{i,g}^{d,s}\Big), \tag{5}$$

$$\widetilde{F}_{i,g}^{d,s} = \bar{F}_{i,g}^{d,s} + \mathrm{Attn}\Big(\bar{F}_{i,g}^{d,s}, \widetilde{T}_{i,g}^s\Big), \tag{6}$$

where $T_{i,g}^s$ denotes the DITs for the spatial branch, and Attn denotes the cross-attention. The first attention operation aggregates degradation information from the spatial features into the DITs, while the second uses the updated DITs to modulate the spatial features. The channel branch follows the same attention formulation using the corresponding channel features $\bar{F}_{i,g}^{d,c}$ and DITs $T_{i,g}^c$.

Finally, we fuse the spatial and channel branch outputs, $\widetilde{F}_{i,g}^{d,s}$ and $\widetilde{F}_{i,g}^{d,c}$ through concatenation followed by a convolutional layer to generate the degradation features $\widetilde{F}_{i,g}^d$ as

$$\widetilde{F}_{i,g}^d = \mathrm{Conv}_{1\times 1}\Big(\mathrm{Concat}\Big[\widetilde{F}_{i,g}^{d,s}, \widetilde{F}_{i,g}^{d,c}\Big]\Big). \tag{7}$$

The fused features capture degradation information across both spatial and channel dimensions. While SCCA is shared across weather types, the weather-specific adapters and DITs specialize each DRE toward its corresponding degradation.

3.2. Degradation Residual Decoding and Composition

After extracting degradation features $\widetilde{F}_{i,g}^d$ for each weather type $g \in \mathcal{G} = \{\mathrm{rain}, \mathrm{haze}, \mathrm{snow}\}$, a Shared Residual Decoder (SRD) is used to decode degradation and content features, which progressively upsamples the fused degradation

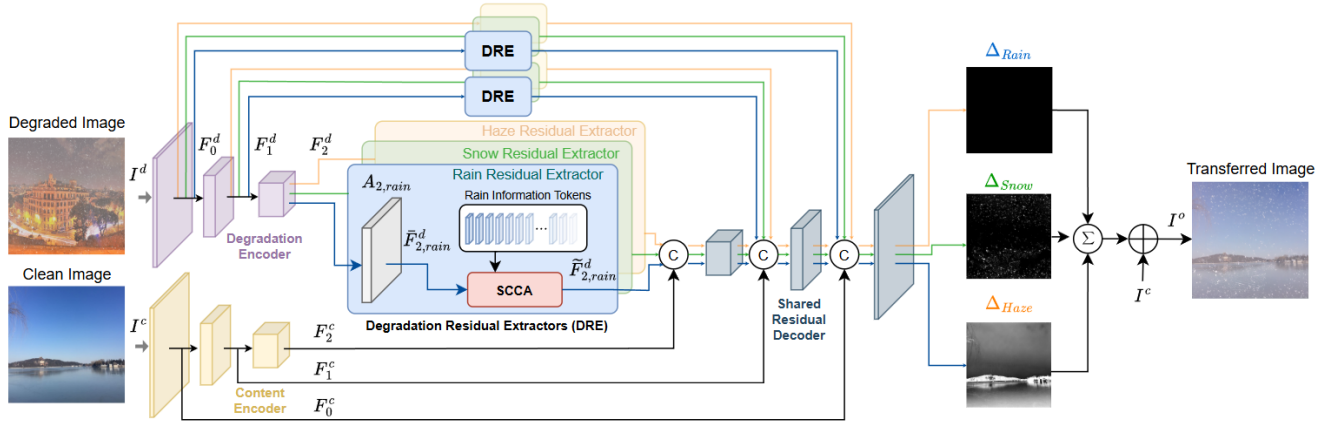


Figure 2. **Architecture of CDTNet.** The Degradation Encoder (DE) and Content Encoder (CE) extract features from the target-domain degraded image I^d, and source-domain clean image I^c. At each scale, weather-specific Degradation Residual Extractors (DREs) model rain, haze, and snow features, which are decoded by the Shared Residual Decoder (SRD) into degradation residuals and composed with I^c to generate I^o.

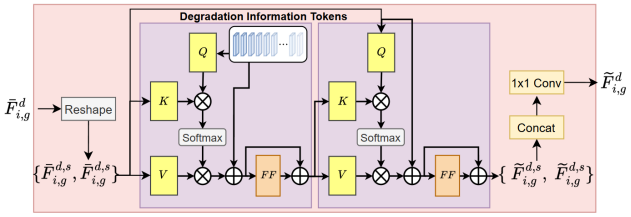


Figure 3. **Architecture of Spatial-Channel Coupled Attention (SCCA).** Weather-specific Degradation Information Tokens interact with image features through spatial and channel attention branches to capture complementary degradation patterns across both dimensions.

and content features as

$$F_{i,g}^{\text{dec}} = \begin{cases} \phi_i\left(\widetilde{F}_{i,g}^{d}, F_i^c\right), & i = 2, \\ \phi_i\left(\widetilde{F}_{i,g}^{d}, F_i^c, \uparrow F_{i+1,g}^{\text{dec}}\right), & i = 1, 0, \end{cases} \tag{8}$$

where ϕ_i denotes the decoding block at scale i, and $\uparrow$ denotes upsampling. The final degradation residual $\Delta_g \in \mathbb{R}^{H\times W\times 3}$ is predicted from the finest-scale decoder feature $F_{0,g}^{\text{dec}}$ as

$$\Delta_g = \text{Conv}_{1\times 1}\left(F_{0,g}^{\text{dec}}\right), \tag{9}$$

Finally, the weather-specific degradation residuals are composed with the source-domain clean image to generate the degradation-transferred image I^o as

$$I^o = I^c + \sum_{g\in\mathcal{G}} \Delta_g. \tag{10}$$

This formulation allows multiple weather-specific residuals to be combined when composite degradations are present.

3.3. Degradation Decomposition Loss

We introduce the Degradation Decomposition Loss (DDL) to encourage each DRE to capture its corresponding weather degradation while suppressing irrelevant residuals. During training, we use only individual-weather samples, where each degraded image contains a single weather degradation (rain, haze, or snow). Based on this setting, DDL consists of three objectives: degradation reconstruction loss, degradation suppression loss, and identity preservation loss. The degradation reconstruction loss supervises the overall degradation-transferred output, while the degradation suppression loss prevents each DRE from capturing degradations of other weather types. The identity preservation loss further prevents degradation generation when the reference image is clean.

To train CDTNet, we construct training triplets (I_n^d, I_n^c, I_n^m), where I_n^d provides degradation reference, I_n^c is the clean image, and I_n^m is the ground-truth degradation-transferred image, where n indexes a sample in the current mini-batch. In particular, I_n^d and I_n^m share the same weather degradation, while I_n^c and I_n^m share the same scene content. Next, we describe the three objectives of DDL in detail.

Degradation Reconstruction Loss. Given a training triplet (I_n^d, I_n^c, I_n^m), the predicted weather-specific residuals are added to the clean image I_n^c to reconstruct the ground-truth degraded image I_n^m. We define the degradation reconstruction loss as

$$\mathcal{L}_{\text{rec}} = \frac{1}{N}\sum_{n=1}^{N}\left\| I_n^c + \sum_{g\in\mathcal{G}} \Delta_{n,g} - I_n^m \right\|_1, \tag{11}$$

where N denotes the number of training triplets in the mini-batch.

Degradation Suppression Loss. Since each training sample contains only one weather degradation, the residuals of the other weather types should be suppressed. We define the degradation suppression loss as

$$\mathcal{L}_{\text{sup}} = \frac{1}{N}\sum_{n=1}^{N}\sum_{g\neq g_n}\|\Delta_{n,g}\|_1\,, \tag{12}$$

where g_n denotes the weather type of the n-th training sample. For example, for a rainy sample, the haze and snow residuals are encouraged to approach zero.

Identity Preservation Loss. In addition to degraded reference images, we intentionally introduce clean images as degradation references during training. Since a clean reference does not provide degradation information for transfer, CDTNet should produce near-zero degradation residuals for all weather types. We therefore define the identity preservation loss as

$$\mathcal{L}_{\text{id}} = \frac{1}{|\mathcal{B}_c||\mathcal{G}|}\sum_{n\in\mathcal{B}_c}\sum_{g\in\mathcal{G}}\|\Delta_{n,g}\|_1\,, \tag{13}$$

where $\mathcal{B}_c$ denotes the set of training samples in which I_n^d is a clean reference image. This objective prevents CDTNet from generating irrelevant degradation residuals when no degradation information is provided by the reference image.

Total Loss. The overall training loss of CDTNet combines the degradation reconstruction, degradation suppression, and identity preservation losses as

$$\mathcal{L}_{\text{DDL}} = \lambda_{\text{rec}}\mathcal{L}_{\text{rec}} + \lambda_{\text{sup}}\mathcal{L}_{\text{sup}} + \lambda_{\text{id}}\mathcal{L}_{\text{id}}, \tag{14}$$

where $\lambda_{\text{rec}} = 1$, $\lambda_{\text{sup}} = 0.1$, and $\lambda_{\text{id}} = 0.1$ control the contributions of the three losses, respectively. The resulting losses enable CDTNet to reconstruct the target degradation while learning separable weather-specific residuals for compositional degradation transfer.

4. Experiments

4.1. Experimental Protocol

Training Datasets. Following DDTNet [24], we construct the training dataset for CDTNet and restoration models using a degraded–clean–mixed triplet synthesis procedure. Each triplet consists of a degraded image I^d and a clean image I^c, together with a ground-truth degradation-transferred image I^m synthesized by transferring the weather degradation of I^d onto I^c. Specifically, we use the rain and snow masks provided by Rain100H [44], Rain100L [44], and Snow100K [26] to synthesize the corresponding rainy and snowy images. For haze synthesis, we use the haze-density and atmospheric-light parameters provided by RESIDE [20]. We generate 5,000 triplets for each degradation type, resulting in a total of 15,000 weather-transfer triplets.

Testing Dataset. For evaluation, we use CDD [13], Weather30K [43], and WeatherBench [11]. CDD and Weather30K are synthetic composite-degradation benchmarks containing images degraded by a single weather type, as well as images with co-occurring rain and haze or snow and haze. WeatherBench serves as real-world benchmark and comprises rainy, snowy, and hazy images. These benchmarks evaluate adaptation to both individual and composite weather degradations. During adaptation, only target-domain degraded images are used as weather references, without access to their corresponding clean ground truths.

CDTNet Configuration. CDTNet is trained on our synthesized degraded–clean–mixed triplets for 150 epochs using the Adam optimizer with a learning rate of 1×10^{-4} and a batch size of 8. During training, all input images are resized to 256×256 pixels. The resulting model comprises 16.8 million parameters and requires 101 ms to perform a single degradation transfer on an NVIDIA A5000 GPU.

Domain Adaptation Protocol. To evaluate the effectiveness of CDTNet for domain adaptation, we employ three state-of-the-art all-in-one image restoration backbones: PromptIR [30], AdaIR [5], and DFPIR [35]. Each backbone is trained in an all-in-one manner on the same 15,000 synthesized training pairs, consistent with the training data used for CDTNet, following its respective default training configuration. In addition, we compare CDTNet against two existing weather-transfer methods, WTNet [17] and DDTNet [24]. For each method, the generated degradation-transferred images are paired with their corresponding clean source images to construct target-adaptive training pairs for fine-tuning the restoration backbones. For a fair comparison, each backbone is fine-tuned for one epoch using the training pairs generated by each weather-transfer method under the same adaptation setting.

4.2. Quantitative Comparisons

Table 1 presents quantitative results on the CDD and Weather30K datasets. CDTNet achieves the best average performance across all weather conditions and restoration backbones. The "Baseline" denotes the restoration models trained without any adaptation process. Compared with the baseline models, CDTNet improves the average PSNR of PromptIR, AdaIR, and DFPIR by 1.17, 0.85, and 1.85 dB, respectively. Compared with DDTNet, CDTNet further improves PSNR by 0.74, 0.59, and 0.81 dB on PromptIR, AdaIR, and DFPIR, respectively. These results demonstrate the effectiveness of CDTNet in adapting all-in-one restoration models to both individual and composite weather degradations.

Table 1. Quantitative evaluation on the CDD and Weather30K datasets. Results are reported in terms of PSNR/SSIM. *Baseline* denotes the results without adaptation; *Single* denotes the average performance across the rain, snow, and haze subsets; *Composite* denotes the average across the rain+haze and snow+haze subsets; and *Average* denotes the overall average performance across single and composite subsets.

		CDD				Weather30K			
Backbone	Method	Single	Composite	Average	Gain	Single	Composite	Average	Gain
PromptIR	Baseline	24.99/.9079	19.72/.8398	22.88/.8807	–	19.61/.8004	16.22/.6864	18.25/.7548	–
	+WTNet	**26.22/.9364**	20.75/.8686	24.03/.9093	+1.15	19.65/.8083	18.04/.7198	19.01/.7729	+0.76
	+DDTNet	26.09/.9270	20.48/.8538	23.84/.8977	+0.96	19.35/.8104	16.35/.7062	18.15/.7687	-0.10
	+CDTNet	26.10/.9275	**21.49/.8655**	**24.25/.9027**	**+1.37**	**19.91/.8218**	**18.19/.7285**	**19.22/.7845**	**+0.97**
AdaIR	Baseline	24.72/.9072	19.81/.8394	22.76/.8801	–	19.54/.8067	16.56/.6851	18.35/.7580	–
	+WTNet	24.96/.9097	19.97/.8407	22.96/.8821	+0.20	18.53/.7792	17.38/.6798	18.07/.7394	-0.28
	+DDTNet	25.58/.9166	20.13/.8508	23.40/.8903	+0.64	19.46/.8101	16.35/.6978	18.22/.7652	-0.13
	+CDTNet	**25.78/.9320**	**20.82/.8430**	**23.79/.8964**	**+1.03**	**19.61/.8141**	**18.12/.6983**	**19.01/.7678**	**+0.66**
DFPIR	Baseline	23.10/.8554	17.21/.7726	20.74/.8223	–	19.16/.7315	15.52/.6373	17.70/.6938	–
	+WTNet	23.60/.8747	19.85/.8349	22.10/.8588	+1.36	18.61/.7277	17.81/.6754	18.29/.7068	+0.59
	+DDTNet	23.78/.8739	19.79/.8315	22.18/.8569	+1.44	19.03/.7563	17.31/.6785	18.34/.7252	+0.64
	+CDTNet	**24.61/.9003**	**20.71/.8459**	**23.05/.8785**	**+2.31**	**19.59/.7675**	**18.35/.6985**	**19.09/.7399**	**+1.39**

Table 2. Quantitative evaluation on the WeatherBench dataset. Results are reported in terms of PSNR/SSIM. *Baseline* denotes the results without adaptation.

Backbone	Method	Rain	Snow	Haze	Average
PromptIR	Baseline	27.49/.8327	22.52/.7975	13.58/.5783	21.20/.7362
	+DDTNet	28.29/.8592	23.83/.8020	13.78/.5858	21.97/.7490
	+CDTNet	**28.62/.8613**	**23.84/.8084**	**14.24/.5961**	**22.23/.7553** (+1.03)
AdaIR	Baseline	27.71/.8433	22.53/.7981	13.93/.5885	21.39/.7433
	+DDTNet	27.84/.8466	22.73/.7997	13.96/.5955	21.51/.7473
	+CDTNet	**28.39/.8650**	**24.10/.8036**	**14.87/.6121**	**22.45/.7602** (+1.06)
DFPIR	Baseline	25.20/.7232	22.18/.7785	12.24/.5146	19.87/.6721
	+DDTNet	27.14/.7958	22.21/.7800	13.40/.5431	20.92/.7063
	+CDTNet	**27.42/.8052**	**22.82/.7837**	**13.85/.5576**	**21.36/.7155** (+1.49)

Table 2 presents the quantitative results on the real-world WeatherBench dataset. CDTNet consistently achieves the best performance across all weather conditions and restoration backbones. Compared with the baseline models, CDTNet improves the average PSNR of PromptIR, AdaIR, and DFPIR by 1.03, 1.06, and 1.49 dB, respectively. CDTNet further outperforms DDTNet across all backbones, improving the average PSNR by 0.26, 0.94, and 0.44 dB for PromptIR, AdaIR, and DFPIR, respectively. These results show that CDTNet consistently improves different restoration backbones when adapting to real-world weather conditions. Overall, these results demonstrate that CDTNet consistently improves different restoration backbones under both composite and real-world weather conditions.

4.3. Qualitative Analysis

Estimated Residuals and Transferred Results. Figure 4 visualizes the rain, snow, and haze residuals estimated by CDTNet, together with the transferred images produced by CDTNet, WTNet, and DDTNet. For degraded images containing rain and haze, CDTNet simultaneously extracts the corresponding rain and haze residuals. Similarly, for degraded images containing snow and haze, CDTNet simultaneously extracts the corresponding snow and haze residuals. The estimated residuals exhibit degradation patterns consistent with those observed in the corresponding degraded images, showing that CDTNet can separate composite weather degradations into meaningful weather-specific components. Additionally, the transfer results further demonstrate the advantage of CDTNet under composite weather conditions. WTNet and DDTNet cannot effectively reproduce multiple degradation components simultaneously. In contrast, CDTNet jointly transfers the decomposed weather components to the clean source image while preserving its underlying scene content and structure.

Restoration Results. Figure 5 presents qualitative restoration results on Weather30K. The baseline, WTNet, and DDTNet generally remove only one type of weather degradation, leaving the remaining degradations visible in the restored images. In contrast, models adapted using the composite degradation images generated by CDTNet effectively remove multiple weather degradations and produce cleaner restoration results. These results highlight the advantage of CDTNet for adaptation to composite weather conditions. Figure 6 presents qualitative restoration results on WeatherBench. Compared with the baseline models and those adapted using WTNet or DDTNet, the models adapted using CDTNet produce clearer restored images with fewer remaining weather degradations. In particular, CDTNet better removes rain, snow, and haze while preserving the underlying scene content and image details. These improvements are consistently observed

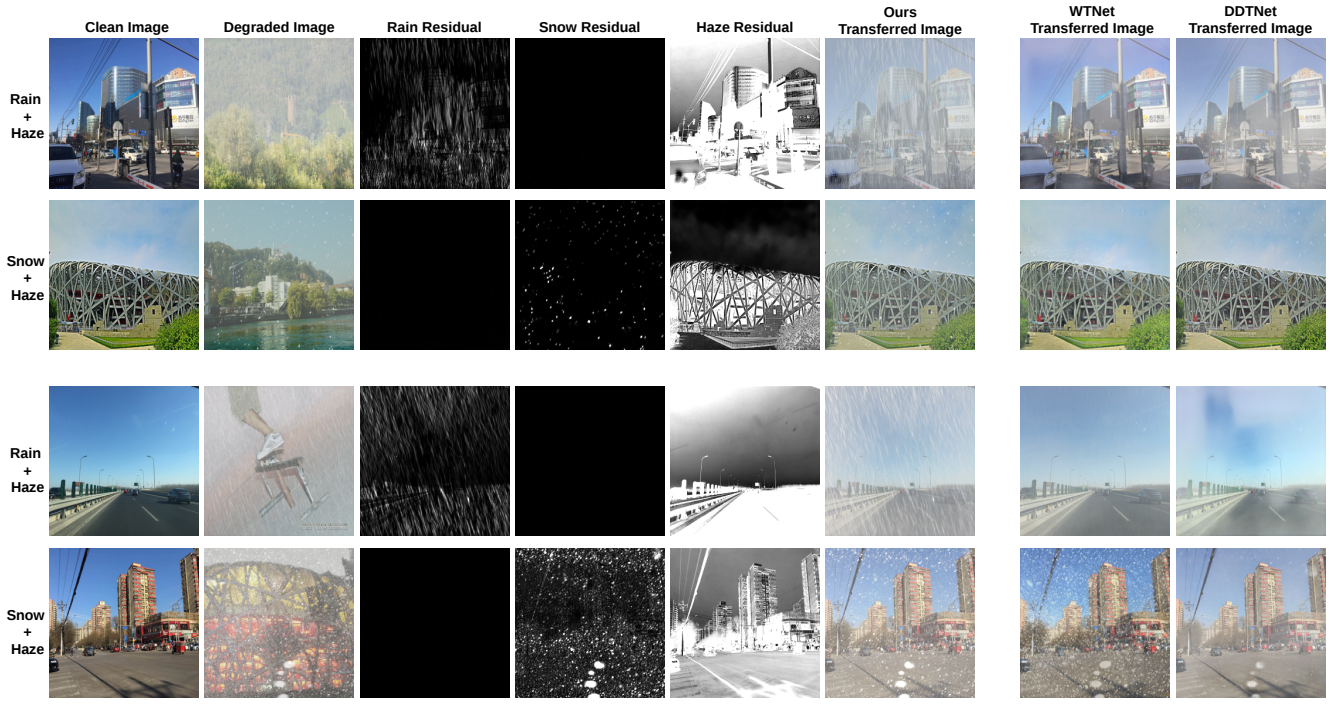


Figure 4. Qualitative comparison of degradation-transferred results on CDD [13] (top) and Weather30K [43] (bottom). From left to right, each example shows the clean images, degraded images, rain, snow, and haze residuals estimated by CDTNet, and the degradation-transferred images generated by CDTNet, WTNet, and DDTNet.

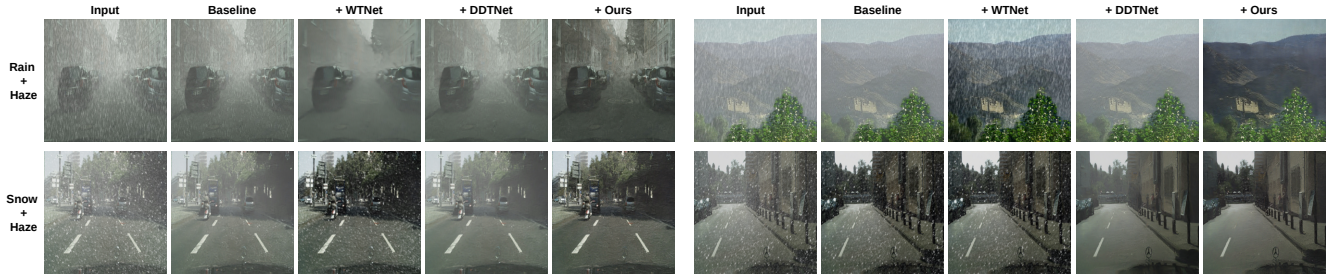


Figure 5. Qualitative restoration comparison on Weather30K using PromptIR (left) and AdaIR (right). Each panel shows the degraded input, followed by the restoration results from the baseline model and the models adapted using WTNet, DDTNet, and CDTNet.

across different real-world weather conditions, further supporting the effectiveness of CDTNet for real-world adaptation.

4.4. Ablation Studies

To evaluate the effectiveness of each component in CDTNet, we follow the evaluation protocol of DDTNet [24] and construct a test set using Rain100H, Rain100L, Snow100K, and RESIDE, containing 900 triplets (I^d, I^c, I^m). We additionally construct composite weather samples by combining rain with haze and snow with haze, enabling the evaluation of compositional degradation transfer. Transfer quality is measured by the PSNR between the transferred images and their corresponding ground-truth images. Next, we conduct ablation studies on the key components of CDTNet, including the Degradation Residual Extractors (DREs), weather-specific adapters, Spatial-Channel Coupled Attention (SCCA), and weather-specific Degradation Information Tokens (DITs).

Effect of DRE and Weather-Specific Adapters. Table 3 evaluates the effects of the Degradation Residual Extractor (DRE) and weather-specific adapters. Introducing DREs improves the baseline by 0.57 dB, showing the benefit of separately extracting different weather degradations. Incorporating weather-specific adapters further improves the performance by 0.63 dB, enabling each DRE to better capture the characteristics of its corresponding weather type. These improvements demonstrate the importance of weather-specific degradation modeling for compositional degradation transfer.

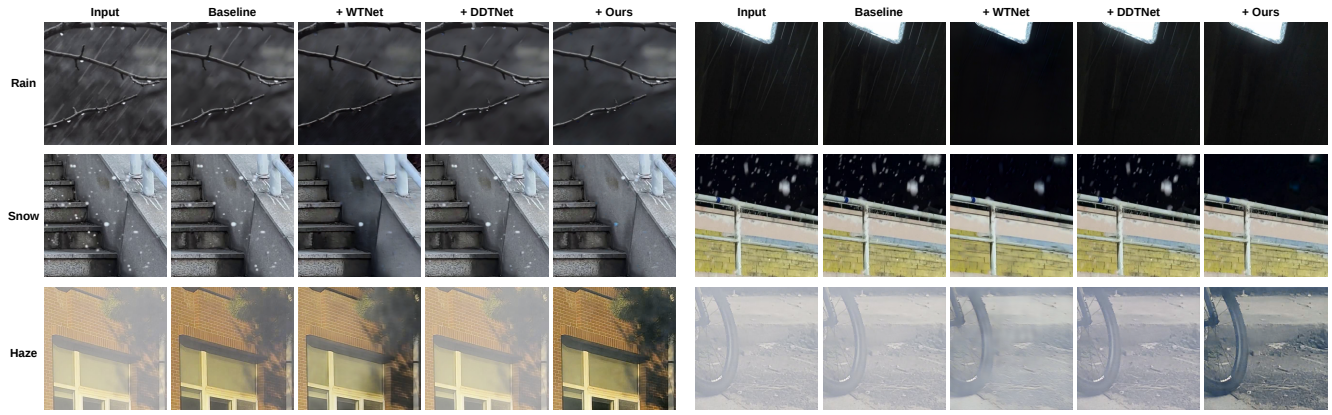


Figure 6. Qualitative restoration comparison on WeatherBench using PromptIR (left) and DFPIR (right). Each panel shows the degraded inputs, followed by restoration results from the baseline model and the models adapted using WTNet, DDTNet, and CDTNet.

Table 3. Effect of DREs and weather-specific adapters.

	DRE	Adapter	PSNR
Net1			18.58
Net2	✓		19.15
Ours	✓	✓	**19.78**

Table 4. Effect of the Spatial-Channel Coupled Attention.

	Spatial	Channel	PSNR
Net1			18.58
Net3	✓		18.90
Net4		✓	19.61
Ours	✓	✓	**19.78**

Effect of Spatial-Channel Coupled Attention. Table 4 demonstrates the effectiveness of the spatial and channel branches of SCCA. Using spatial or channel alone improves the baseline by 0.32 dB and 1.03 dB, respectively, while combining both branches yields the best performance of 19.78 dB. These results indicate that the spatial and channel branches capture complementary degradation information, leading to improved performance when combined.

Effect of Weather-Specific DITs. Table 5 evaluates the effects of weather-specific Degradation Information Tokens (DITs) in CDTNet. "Shared DITs" uses the same DITs across rain, haze, and snow, whereas "Specific DITs" uses separate DITs for each weather type. Compared with "Shared DITs", "Specific DITs" outperform "Shared DITs" DITs by 0.99 dB. This improvement highlights the importance of weather-specific DITs for composite degradation transfer, where multiple weather degradations need to be separately modeled and jointly transferred.

4.5. Limitations.

CDTNet currently focuses on three representative weather degradations, including rain, haze, and snow, as well as their compositions. These weather types exhibit diverse degradation characteristics and provide a practical setting for investigating both individual and compositional degradation transfer. Nevertheless, real-world adverse weather conditions may involve a broader range of degradation types and combinations. Extending CDTNet to additional weather types and more diverse compositions therefore represents a promising direction for future work.

Table 5. Effect of weather-shared and weather-specific DITs.

	Shared DITs	Specific DITs	PSNR
Net5	✓		18.79
Ours		✓	**19.78**

5. Conclusion

We presented CDTNet, a domain adaptation framework for addressing domain gaps and composite weather degradations in all-in-one image de-weathering. CDTNet decomposes target-domain weather degradations into weather-specific residuals and transfers them to source-domain clean images, enabling the construction of domain-adaptive training pairs without requiring clean target images. Through weather-specific degradation modeling and compositional residual transfer, CDTNet enables multiple weather degradations to be separately captured and jointly transferred for effective adaptation under composite weather conditions. Extensive experiments on CDD, Weather30K, and WeatherBench show that CDTNet consistently improves different restoration backbones and outperforms existing weather-transfer methods under both individual and composite weather conditions. These results demonstrate the potential of compositional degradation transfer for adapting all-in-one restoration models to diverse target-domain weather degradations.

Acknowledgments

This work was supported in part by the National Science and Technology Council (NSTC) under grants 114-2221-E-A49-038-MY3, 115-2221-E-A49-119-MY3, and 115-2634-F-A49-011, and by MediaTek.

References

[1] Wei-Ting Chen, Hao-Yu Fang, Jian-Jiun Ding, Cheng-Che Tsai, and Sy-Yen Kuo. JSTASR: Joint size and transparency-aware snow removal algorithm based on modified partial convolution and veiling effect removal. In *ECCV*, 2020. 1

[2] Wei-Ting Chen, Hao-Yu Fang, Cheng-Lin Hsieh, Cheng-Che Tsai, I-Hsiang Chen, Jian-Jiun Ding, and Sy-Yen Kuo. All snow removed: Single image desnowing algorithm using hierarchical dual-tree complex wavelet representation and contradict channel loss. In *ICCV*, 2021. 1

[3] 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. 1, 2

[4] Xiang Chen, Jinshan Pan, and Jiangxin Dong. Bidirectional multi-scale implicit neural representations for image deraining. In *CVPR*, 2024. 1

[5] 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. 1, 2, 5

[6] Qili Deng, Ziling Huang, Chung-Chi Tsai, and Chia-Wen Lin. HardGAN: A haze-aware representation distillation GAN for single image dehazing. In *ECCV*, 2020. 1

[7] Wenxuan Fang, JunKai Fan, Yu Zheng, Jiangwei Weng, Ying Tai, and Jun Li. Guided real image dehazing using ycbcr color space. In *AAAI*, 2025. 1

[8] Xueyang Fu, Jiabin Huang, Delu Zeng, Yue Huang, Xinghao Ding, and John Paisley. Removing rain from single images via a deep detail network. In *CVPR*, 2017. 1

[9] Ning Gao, Xingyu Jiang, Xiuhui Zhang, and Yue Deng. Efficient frequency-domain image deraining with contrastive regularization. In *ECCV*, 2024. 1

[10] Yuanbiao Gou, Haiyu Zhao, Boyun Li, Xinyan Xiao, and Xi Peng. Test-time degradation adaptation for open-set image restoration. In *ICML*, 2024. 2

[11] Qiyuan Guan, Qianfeng Yang, Xiang Chen, Tianyu Song, Guiyue Jin, and Jiyu Jin. Weatherbench: A real-world benchmark dataset for all-in-one adverse weather image restoration. In *ACM MM*, 2025. 5

[12] Chun-Le Guo, Qixin Yan, Saeed Anwar, Runmin Cong, Wenqi Ren, and Chongyi Li. Image dehazing transformer with transmission-aware 3d position embedding. In *CVPR*, 2022. 1

[13] Yu Guo, Yuan Gao, Yuxu Lu, Huilin Zhu, Ryan Wen Liu, and Shengfeng He. Onerestore: A universal restoration framework for composite degradation. In *ECCV*, pages 255–272, 2024. 2, 5, 7

[14] Jin-Ting He, Fu-Jen Tsai, Jia-Hao Wu, Yan-Tsung Peng, Chung-Chi Tsai, Chia-Wen Lin, and Yen-Yu Lin. Domain-adaptive video deblurring via test-time blurring. In *ECCV*, 2024. 2

[15] Jin-Ting He, Fu-Jen Tsai, Yan-Tsung Peng, Min-Hung Chen, Chia-Wen Lin, and Yen-Yu Lin. Hndiff: Haze-noise diffusion for image dehazing. In *ECCV*, 2026. 1

[16] Xiaowei Hu, Chi-Wing Fu, Lei Zhu, and Pheng-Ann Heng. Depth-attentional features for single-image rain removal. In *CVPR*, 2019. 1

[17] Si-Yu Huang, Fu-Jen Tsai, Chia-Wen Lin, and Yen-Yu Lin. Wtnet: A weather transfer network for domain-adaptive all-in-one adverse weather image restoration. In *BMVC*, 2025. 1, 2, 5

[18] 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 *CVPR*, 2020. 1

[19] Kui Jiang, Zhongyuan Wang, Chen Chen, Zheng Wang, Laizhong Cui, and Chia-Wen Lin. Magic ELF: Image deraining meets association learning and transformer. In *ACM MM*, 2022. 1

[20] Boyi Li, Wenqi Ren, Dengpan Fu, Dacheng Tao, Dan Feng, Wenjun Zeng, and Zhangyang Wang. Benchmarking single-image dehazing and beyond. *IEEE TIP*, 2019. 5

[21] Boyun Li, Xiao Liu, Peng Hu, Zhongqin Wu, Jiancheng Lv, and Xi Peng. All-In-One Image Restoration for Unknown Corruption. In *CVPR*, 2022. 1, 2

[22] Ruoteng Li, Loong-Fah Cheong, and Robby T. Tan. Heavy rain image restoration: Integrating physics model and conditional adversarial learning. In *CVPR*, 2019. 1

[23] Kang Liao, Zongsheng Yue, Zhouxia Wang, and Chen Change Loy. Denoising as adaptation: Noise-space domain adaptation for image restoration. In *ICLR*, 2025. 2

[24] Kuan-Hung Lin, Fu-Jen Tsai, Yan-Tsung Peng, Min-Hung Chen, Chia-Wen Lin, and Yen-Yu Lin. Ddtnet: Degradation disentanglement and transfer network for test-time all-in-one de-weathering adaptation. *arXiv preprint arXiv:2606.16298*, 2026. 1, 2, 5, 7

[25] Xiaohong Liu, Yongrui Ma, Zhihao Shi, and Jun Chen. Griddehazenet: Attention-based multi-scale network for image dehazing. In *ICCV*, 2019. 1

[26] Yun-Fu Liu, Da-Wei Jaw, Shih-Chia Huang, and Jenq-Neng Hwang. DesnowNet: Context-aware deep network for snow removal. *IEEE TIP*, 2018. 1, 5

[27] Wenyang Luo, Haina Qin, Zewen Chen, Libin Wang, Dandan Zheng, Yuming Li, Yufan Liu, Bing Li, and Weiming Hu. Visual-instructed degradation diffusion for all-in-one image restoration. In *CVPR*, 2025. 2

[28] Dongwon Park, Byung Hyun Lee, and Se Young Chun. All-in-one image restoration for unknown degradations using adaptive discriminative filters for specific degradations. In *CVPR*, 2023. 2

[29] Prashant W. Patil, Sunil Gupta, Santu Rana, Svetha Venkatesh, and Subrahmanyam Murala. Multi-weather image restoration via domain translation. In *ICCV*, 2023. 2

[30] Vaishnav Potlapalli, Syed Waqas Zamir, Salman Khan, and Fahad Khan. PromptIR: Prompting for all-in-one image restoration. In *NeurIPS*, 2023. 1, 2, 5

[31] Xu Qin, Zhilin Wang, Yuanchao Bai, Xiaodong Xie, and Huizhu Jia. FFA-Net: Feature fusion attention network for single image dehazing. In *AAAI*, 2020. 1

[32] Yuanjie Shao, Lerenhan Li, Wenqi Ren, Changxin Gao, and Nong Sang. Domain adaptation for image dehazing. In *CVPR*, 2020. 2

[33] Yuda Song, Zhuqing He, Hui Qian, and Xin Du. Vision transformers for single image dehazing. *IEEE TIP*, 2023. 1

[34] Ni Tang, Shenghao Nie, Xiaotong Luo, Yuan Xie, and Yanyun Qu. Degradation-consistent test-time adaptation for all-in-one image restoration. In *CVPR*, pages 15476–15485, 2026. 2

[35] Xiangpeng Tian, Xiangyu Liao, Xiao Liu, Meng Li, and Chao Ren. Degradation-aware feature perturbation for all-in-one image restoration. In *CVPR*, 2025. 1, 2, 5

[36] Jeya Maria Jose Valanarasu, Rajeev Yasarla, and Vishal M. Patel. Transweather: Transformer-based restoration of images degraded by adverse weather conditions. In *CVPR*, 2022. 1, 2

[37] Hong Wang, Qi Xie, Qian Zhao, and Deyu Meng. A model-driven deep neural network for single image rain removal. In *CVPR*, 2020. 1

[38] Qin Wang, Olga Fink, Luc Van Gool, and Dengxin Dai. Continual test-time domain adaptation. In *CVPR*, pages 7201–7211, 2022. 2

[39] Tianyu Wang, Xin Yang, Ke Xu, Shaozhe Chen, Qiang Zhang, and Rynson W.H. Lau. Spatial attentive single-image deraining with a high quality real rain dataset. In *CVPR*, 2019. 1

[40] Yinglong Wang, Shuaicheng Liu, Chen Chen, and Bing Zeng. A hierarchical approach for rain or snow removing in a single color image. *IEEE TIP*, 2017. 1

[41] Yanyan Wei, Zhao Zhang, Yang Wang, Mingliang Xu, Yi Yang, Shuicheng Yan, and Meng Wang. Deraincyclegan: Rain attentive cyclegan for single image deraining and rainmaking. *IEEE TIP*, 2021. 2

[42] Haiyan Wu, Yanyun Qu, Shaohui Lin, Jian Zhou, Ruizhi Qiao, Zhizhong Zhang, Yuan Xie, and Lizhuang Ma. Contrastive learning for compact single image dehazing. In *CVPR*, 2021. 1

[43] Jiaqi Xu, Xiaowei Hu, Lei Zhu, and Pheng-Ann Heng. Unifying physically-informed weather priors in a single model for image restoration across multiple adverse weather conditions. *IEEE TCSVT*, pages 9575–9591, 2025. 5, 7

[44] 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 *CVPR*, 2017. 5

[45] Hu Yu, Naishan Zheng, Man Zhou, Jie Huang, Zeyu Xiao, and Feng Zhao. Frequency and spatial dual guidance for image dehazing. In *ECCV*, 2022. 1

[46] Kaihao Zhang, Rongqing Li, Yanjiang Yu, Wenhan Luo, and Changsheng Li. Deep dense multi-scale network for snow removal using semantic and geometric priors. *IEEE TIP*, 2021. 1

[47] Yafei Zhang, Shen Zhou, and Huafeng Li. Depth information assisted collaborative mutual promotion network for single image dehazing. In *CVPR*, 2024. 1