WTNet: A Weather Transfer Network for Domain-Adaptive All-In-One Adverse Weather Image Restoration

Si-Yu Huang*[1]
stella900604.ii12@nycu.edu.tw
Fu-Jen Tsai*[2]
fjtsai@gapp.nthu.edu.tw
Chia-Wen Lin[2]
cwlin@ee.nthu.edu.tw
Yen-Yu Lin[1]
lin@cs.nycu.edu.tw

[1] National Yang Ming Chiao Tung University, Taiwan
[2] National Tsing Hua University, Taiwan

Abstract

All-in-one adverse weather image restoration has attracted increasing attention due to its potential to recover high-quality images with a single model. However, existing methods often exhibit significant performance drops due to the domain gap between training and testing weather conditions. Moreover, they typically achieve only average, rather than optimal, performance across different weather conditions, when compared to weather-specific approaches. To address these two issues, we propose a novel Weather Transfer Network (WTNet), which fine-tunes all-in-one models to enhance their performance during testing. Recognizing the unavailability of paired degraded-clean images at test time, WTNet transfers degradation patterns from the testing images in an unseen target domain to clean images in the source domain, thereby generating the fine-tuning sets for enabling domain adaptation. Additionally, by leveraging the fine-tuning sets, all-in-one models can be dynamically adapted to weather-specific or mixed weather models based on the transferred degradation patterns observed during testing. Experimental results demonstrate that WTNet can significantly enhance state-of-the-art all-in-one models across real-world image deraining, desnowing, and dehazing benchmarks. The source code is available at https://github.com/stellaahuang/WTNet.

1 Introduction

Adverse weather image restoration aims to remove undesirable artifacts caused by weather conditions like rain, haze, and snow from a degraded image. With the advancement of deep learning, significant progress has been made in addressing individual weather conditions, such as deraining [6, 11, 12, 19, 20, 21, 24, 38, 40], dehazing [8, 9, 14, 26, 31, 33, 35, 42, 47, 50], and desnowing [2, 3, 27, 41, 49]. However, the inherent unpredictability and

*equal contribution

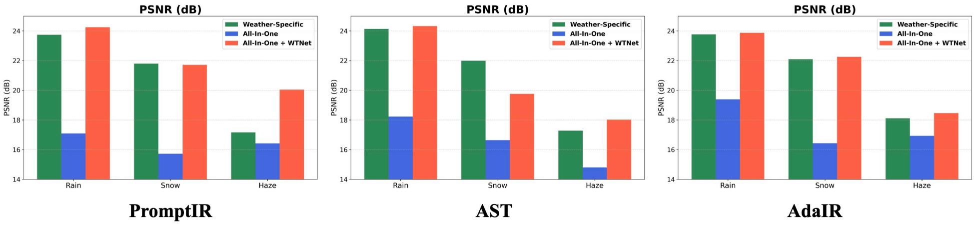

Figure 1: **Quantitative comparison of different training strategies.** Previous methods, including PromptIR [30], AST [51], and AdaIR [7], exhibit significant performance gaps between using weather-specific and all-in-one training strategies. In contrast, the proposed WTNet substantially improves these all-in-one models, even outperforming their weather-specific variants.

dynamic variations in weather pose a significant challenge, often limiting the generalization ability of models tailored for individual weather conditions in real-world applications. As a result, all-in-one image restoration approaches [4, 7, 23, 29, 30, 36], which address various types of degradation within a unified model, have gained increasing attention in recent years.

Recent all-in-one restoration methods fall into two categories: approaches addressing general degradations [7, 23, 29] and those tackling adverse weather conditions [4, 30, 36]. To handle diverse weather conditions, several studies have incorporated degradation-specific features, such as weather-type queries [36], degradation-specific prompts [30], or multi-teacher networks [4], into their unified models. While effective at addressing various degradations within a single network, these all-in-one methods tend to exhibit balanced but potentially lower performance across different degradations, compared to their weather-specific variants, i.e., blue vs. green bars in Figure 1. Additionally, existing all-in-one models often suffer from significant performance drops due to the domain gap between training and testing images. This issue is particularly pronounced in real-world scenarios where images frequently exhibit complex degradations resulting from mixed weather conditions, rather than a single type. These limitations hinder the potential of all-in-one models in diverse and uncontrollable weather conditions.

In this paper, we propose a novel Weather Transfer Network (WTNet), which fine-tunes all-in-one models to improve their performance during testing. Given that paired degraded-clean images are unavailable during testing, WTNet generates domain-adaptive fine-tuning sets, which encode the degradation patterns of the unseen target domain for effective domain adaptation. Unlike previous restoration methods that recover high-quality content from degraded inputs, WTNet takes a different approach: predicting degradation patterns. This task is generally simpler than content restoration because degradation patterns like rain, haze, and snow are typically visible, consistent, and less textured. After predicting degradation patterns, WTNet transfers them from unseen target domains to clean images in the source domain, thereby constructing fine-tuning sets, i.e., degradation-transferred and clean source-domain images, to adapt the restoration models during testing.

WTNet employs a physics-based model to disentangle and reassemble key weather components, including snow masks, rain streaks, haze density, and atmospheric light. Thus, WTNet can exploit the inherent inductive biases of weather formation, leading to more accurate weather pattern prediction and transfer. WTNet generates fine-tuning sets for test-time adaptation and offers two primary advantages. First, it effectively reduces the domain gap between training and testing sets by transferring degradation patterns from the target domain to the source domain. Second, it dynamically adapts all-in-one models to either weather-specific or mixed-weather scenarios, guided by the degradation patterns observed during testing, thereby improving restoration performance, i.e., blue vs. orange bars in Figure 1.

Our contributions are summarized as follows: First, we propose WTNet, a novel framework that transfers degradation patterns from unseen target domains to source-domain clean images. This process creates domain-adaptive fine-tuning sets that enhance the performance of all-in-one models on unseen domains during testing. Second, WTNet enables dynamic adaptation of all-in-one models to specific or mixed weather conditions during testing, thereby unlocking their full potential and improving performance under uncontrollable weather scenarios. Third, extensive experimental results demonstrate that WTNet significantly boosts the performance of all-in-one restoration models across benchmark real-world image deraining, desnowing, and dehazing datasets.

2 Related Works

2.1 All-in-One Image Restoration

All-in-one image restoration aims to tackle various types of degradation using a single, unified model. To this end, several studies [4, 7, 23, 29, 30, 36, 39] have explored incorporating degradation-specific features into unified architectures to handle diverse degradations adaptively. For instance, Li *et al.* [23] use contrastive learning to extract degradation-specific features, which are subsequently used to guide their all-in-one restoration model. Valanarasu *et al.* [36] utilize learnable weather-type queries to encode weather-specific information within a Transformer framework. Potlapalli *et al.* [30] employ learnable prompts to inject degradation-specific features into a unified restoration model. Cui *et al.* [7] adaptively restore degraded images by harnessing frequency-aware cues tailored to specific types of degradation. Although these methods handle diverse degradations using a single model, they often achieve a compromise in performance, resulting in balanced rather than optimal results across different degradation types. Moreover, they frequently suffer from significant performance drops due to the domain gap between training and testing images. This issue is even exacerbated in real-world scenarios, where images often suffer from complex degradations caused by mixed weather conditions.

2.2 Domain Adaptation for Restoration

Domain adaptation seeks to reduce the discrepancy between source and target domains. Since paired degraded-clean images are typically unavailable at test time, many studies employ generative models [18, 37] to synthesize pseudo training pairs for domain adaptation. For example, Chen *et al.* [5] and Shao *et al.* [32] leverage generative adversarial networks (GANs) [37] to synthesize paired training data from unpaired images for deraining and dehazing tasks, respectively. However, GAN-based methods may suffer from unstable optimization [1, 13, 28] and mode collapse [34, 45]. As an alternative, some recent studies have explored diffusion models to achieve domain adaptation. For instance, He *et al.* [15] propose a domain adaptation approach for video deblurring based on a diffusion-based blurring model [43]. Although their method can generate domain-adaptive training pairs for deblurring, the reliance on video motion flow limits its applicability to image restoration.

Despite the effectiveness, the aforementioned methods [5, 15, 32] focus on mitigating domain gaps for specific degradations, such as deraining, dehazing, or deblurring, while overlooking the practicality and generalizability of all-in-one models for image restoration. In

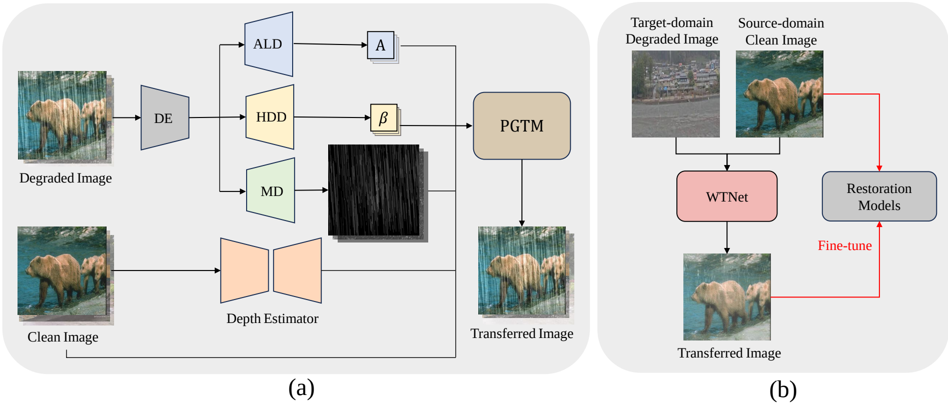

Figure 2: (a) WTNet Architecture: WTNet employs a Degradation Encoder (DE), Atmospheric Light Decoder (ALD), Haze Density Decoder (HDD), and Mask Decoder (MD) to parameterize weather-related features. These parameters are then transferred to the clean image via the Physics-Guided Transfer Module (PGTM). (b) Test-Time Domain Adaptation: WTNet transfers degradation patterns from target-domain images to source-domain clean images, generating domain-adaptive fine-tuning sets for restoration model adaptation.

contrast, we propose WTNet, a weather-transfer-based adaptation method that fine-tunes all-in-one adverse weather image restoration models during testing, enabling dynamic adaptation to a wider range of adverse weather conditions, including deraining, dehazing, desnowing, and even mixed-weather scenarios.

3 Proposed Method

This section introduces the proposed Weather Transfer Network (WTNet), a novel framework designed to transfer degradation patterns from weather-degraded images in unseen target domains to clean images from the source domain. The generated clean-degraded image pairs, embedded with target-domain degradation characteristics, are then used as domain-adaptive fine-tuning data to adapt all-in-one restoration models for improved performance on previously unseen target domains during testing. The remainder of this section provides an overview of the proposed approach, details each module of WTNet, and outlines the associated loss functions and fine-tuning strategy.

3.1 Overview

As illustrated in Figure 2, WTNet begins by employing a Degradation Encoder (DE) to extract degradation features from a target-domain weather-degraded image. These features are then projected into a parametric space using multiple specialized decoders. Specifically, the Atmospheric Light Decoder (ALD) and Haze Density Decoder (HDD) extract haze-related parameters corresponding to atmospheric light and haze density, while the Mask Decoder (MD) estimates occlusion masks caused by rain or snow. To transfer these weather-degradation patterns to source-domain clean images, we introduce the Physics-Guided Transfer Module (PGTM), inspired by the atmospheric scattering models [10, 16]. By leveraging the inductive bias inherent in weather formation processes, PGTM enables effective and unified transfer of diverse weather patterns. During testing, WTNet facilitates adaptation of a restoration model by generating paired training samples—source-domain clean images and their degraded counterparts with target-domain degradation characteristics—thus improving generalization to unseen weather conditions.

3.2 Degradation Encoder and Parametric Decoders

Given a degraded image $I^d \in \mathbb{R}^{H\times W\times 3}$, WTNet employs the Degradation Encoder (DE) to extract degradation features $F^d \in \mathbb{R}^{(H/32)\times(W/32)\times 256}$. These features are subsequently processed by three specialized parametric decoders: the Atmospheric Light Decoder (ALD), Haze Density Decoder (HDD), and Mask Decoder (MD), which collectively parameterize F^d as follows:

$$A = \mathrm{ALD}(F^d), \quad \beta = \mathrm{HDD}(F^d), \quad \text{and} \quad M = \mathrm{MD}(F^d), \tag{1}$$

where $A \in \mathbb{R}$ and $\beta \in \mathbb{R}$ denote the atmospheric light and haze density, respectively, and $M \in \mathbb{R}^{H\times W\times 1}$ with values in $[0,1]$ represents the occlusion mask induced by rain or snow.

In our implementation, DE consists of six convolutional blocks, each comprising two residual blocks followed by a downsampling convolutional layer. Each of the three decoders, ALD, HDD, and MD, is composed of five convolutional blocks, where each block includes a bilinear upsampling layer followed by two residual blocks. To produce compact representations, a global average pooling layer is applied to the outputs of both ALD and HDD. To guide the disentanglement of degradation features, we incorporate PGTM, which maps the predicted parameters back to the degraded image space. This ensures that the parameterization produced by ALD, HDD, and MD remains physically meaningful and interpretable.

3.3 Physics-Guided Transfer Module (PGTM)

After retrieving weather parameters, A, β, and M, WTNet employs PGTM to transfer these parameters onto the clean image.. As the first step, the occlusion mask M is applied to generate an initial degradation-transferred image $O^{\mathrm{ini}} \in \mathbb{R}^{H\times W\times 3}$ as follows:

$$O^{\mathrm{ini}}(x) = I^c(x)(1-M(x)) + S\cdot M(x), \tag{2}$$

where $I^c \in \mathbb{R}^{H\times W\times 3}$ denotes the clean image. x represents the pixel index, and S is a scalar randomly sampled from the range $[1.0, 2.61]$ to control the intensity of the occlusion effect, following the setting in [24].

WTNet then transfers the haze-related parameters A and β onto O^{ini} to generate the final weather-transferred image $O \in \mathbb{R}^{H\times W\times 3}$ based on the atmospheric scattering model:

$$O(x) = O^{\mathrm{ini}}(x)T(x) + A(1-T(x)), \tag{3}$$

where $T(x) = e^{-\beta d(x)}$ denotes the transmission map, which quantifies the proportion of scene radiance that reaches the camera, as defined by the atmospheric scattering model [10, 16], and $d(x) = \theta^{\mathrm{dep}}(I^c(x))$ denotes the depth estimated from a pre-trained depth estimation network θ^{dep} [17].

By incorporating this physically grounded formulation, PGTM enables WTNet to simulate the formation process of adverse weather conditions, thereby facilitating the unified transfer of diverse degradation types in a physically interpretable manner.

Method		Rain PSNR ↑	Rain SSIM ↑	Snow PSNR ↑	Snow SSIM ↑	Haze PSNR ↑	Haze SSIM ↑	Average PSNR ↑	Average SSIM ↑
PromptIR [30]	Baseline	17.09	0.524	15.73	0.465	16.42	0.383	16.41	0.457
	+WTNet	**24.12 (+7.03)**	**0.766 (+0.242)**	**22.10 (+6.37)**	**0.730 (+0.265)**	**19.81 (+3.39)**	**0.635 (+0.252)**	**22.01 (+5.60)**	**0.710 (+0.253)**
AST [51]	Baseline	18.23	0.592	16.64	0.542	14.80	0.405	16.56	0.513
	+WTNet	**24.27 (+6.04)**	**0.762 (+0.170)**	**21.29 (+4.65)**	**0.698 (+0.156)**	**17.68 (+2.88)**	**0.575 (+0.170)**	**21.08 (+4.52)**	**0.678 (+0.165)**
AdaIR [8]	Baseline	19.39	0.614	16.43	0.546	16.93	0.398	17.58	0.519
	+WTNet	**23.21 (+3.82)**	**0.752 (+0.138)**	**22.11 (+5.68)**	**0.724 (+0.178)**	**19.49 (+2.56)**	**0.650 (+0.252)**	**21.60 (+4.02)**	**0.709 (+0.190)**
Average Gain		**+5.63**	**+0.183**	**+5.57**	**+0.200**	**+2.94**	**+0.225**	**+4.71**	**+0.203**

Table 1: Quantitative comparison of image restoration performance under three real-world weather types: Rain, Snow, and Haze. Results are reported in PSNR (dB) and SSIM. WTNet consistently improves performance across all backbone models and weather conditions.

3.4 Loss Function

WTNet is trained on synthetic datasets of degraded-clean pairs $\{I_i^d, I_i^c\}_{i=1}^N$. We supervise the network in both the parametric and image spaces with the objective function:

$$\mathcal{L} = ||A_i - A_i^{\mathrm{GT}}||_1 + ||\beta_i - \beta_i^{\mathrm{GT}}||_1 + ||M_i - M_i^{\mathrm{GT}}||_1 + ||O_i - I_i^d||_1, \tag{4}$$

where (A_i, A_i^{GT}), (β_i, β_i^{GT}), (M_i, M_i^{GT}), and (O_i, I_i^d) denote the estimated results in (1) and (3) and their corresponding ground-truth.

3.5 Domain-Adaptive Fine-Tuning Strategy

After training WTNet, we use it to transfer degradation patterns from target-domain degraded images $\{\hat{I}_i^d\}_{i=1}^N$ to source-domain clean images $\{I_i^c\}_{i=1}^N$, where we select N clean images randomly from the source domain. This process yields a domain-adaptive fine-tuning set $\{\hat{O}_i, I_i^c\}_{i=1}^N$, where $\hat{O}_i = \mathrm{WTNet}(\hat{I}_i^d, I_i^c)$ denotes the weather-transferred image with degradation patterns from $\hat{I}_i^d$ and scene content from I_i^c. We then fine-tune the all-in-one restoration model using this synthesized dataset, as illustrated in Figure 2(b). To ensure adaptation efficiency in test time, we fine-tune each restoration model for only a single epoch.

4 Experiments

4.1 Experiment Settings

Datasets. We train WTNet and restoration models using synthetic datasets: Rain100H [46] for deraining, Snow100K [27] for desnowing, and RESIDE [22] for dehazing. Specifically, Rain100H and Snow100K provide paired degraded-clean images along with their corresponding rain and snow masks, which serve as ground-truth occlusion masks for supervising WTNet's occlusion modeling. For dehazing, we synthesize hazy images from the clean images in RESIDE following the procedure described in [44], allowing us to generate hazy-clean image pairs with known atmospheric light and haze density. These datasets provide the necessary supervision signals—occlusion masks, atmospheric light, and haze density—for training WTNet to disentangle and transfer weather-related degradation parameters.

During training, we select 1,800 paired samples from each dataset, resulting in a total of 5,400 paired degraded-clean images for training WTNet and the restoration models. To demonstrate the generalizability of WTNet, we adopt the real-world dataset WeatherStream [48] during testing, which contains 3,000 rainy, 4,500 hazy, and 3,960 snowy images, along with their corresponding clean images captured by fixed webcams. Notably, WeatherStream presents additional challenges due to the presence of mixed-weather conditions, such as haze co-occurring with rain or snow, making it a suitable benchmark for assessing the robustness of domain adaptation and weather-specific restoration performance.

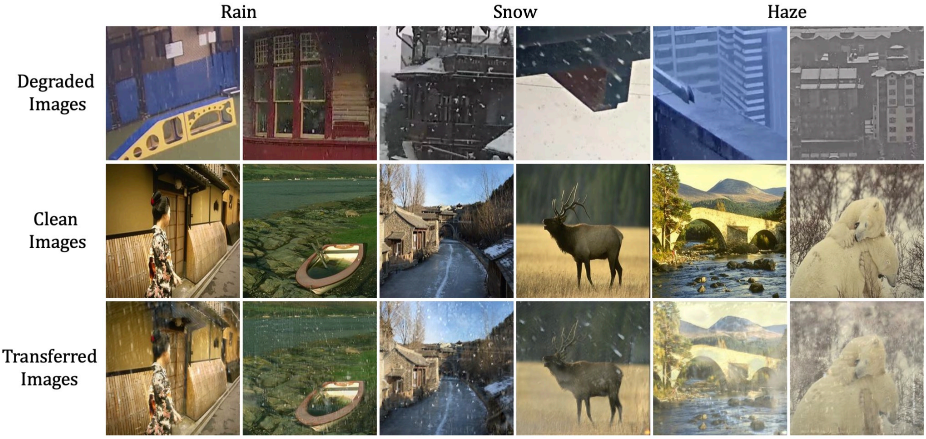

Figure 3: Qualitative results of weather-transferred images. We transfer degradation patterns presented in the target-domain dataset WeatherStream [48] to the clean images in source domain datasets: Rain100H [46], Snow100K [27], and RESIDE [22].

Implementation details. We optimize WTNet for 300 epochs with a batch size of 8 using the AdamW optimizer and a learning rate of 5×10^{-4}. During training and testing, all input images to WTNet are resized to 256×256. WTNet contains 16 million parameters and has an inference time of 30 milliseconds (ms) to generate a weather-transferred image on an NVIDIA 2080Ti GPU.

Restoration Models. To assess the effectiveness of WTNet for test-time adaptation, we conduct experiments using three state-of-the-art restoration models: PromptIR [30], AST [51], and AdaIR [7], which serve as the evaluation backbones for domain adaptation performance. PromptIR and AdaIR are all-in-one frameworks capable of handling multiple degradation types within a single unified network, while AST is a degradation-specific model applied individually to each degradation type. All models are trained in an all-in-one fashion using a mixed set of 5,400 synthetic image pairs, following their default training configurations. During testing, each model is fine-tuned for one epoch using the domain-adaptive fine-tuning sets generated by WTNet, enabling adaptation to unseen target-domain degradations.

4.2 Experimental Results

Quantitative Comparison. We compare the restoration performances of three baseline models and their WTNet-enhanced counterparts in Table 1, where "Baseline" refers to models trained without WTNet, and "+WTNet" denotes models fine-tuned using the proposed WTNet framework. As shown in Table 1, WTNet consistently and significantly improves the performance of all three SoTA restoration models: PromptIR [30], AST [51], and AdaIR [7]. In particular, WTNet boosts the average PSNR of PromptIR, AST, and AdaIR by 5.59 dB, 4.15 dB, and 3.95 dB, respectively, on WeatherStram. Task-wise, WTNet also yields substantial performance gains, achieving average improvements of 5.91 dB for deraining, 4.97 dB for desnowing, and 2.79 dB for dehazing. In Figure 1, we further compare restoration performance across three training strategies: (i) weather-specific training, (ii) all-in-one training, and (iii) all-in-one training with WTNet. Models enhanced with WTNet not only outperform conventional all-in-one models but also approach, or in some cases surpass, the performance of weather-specific models. These results demonstrate WTNet's effective-

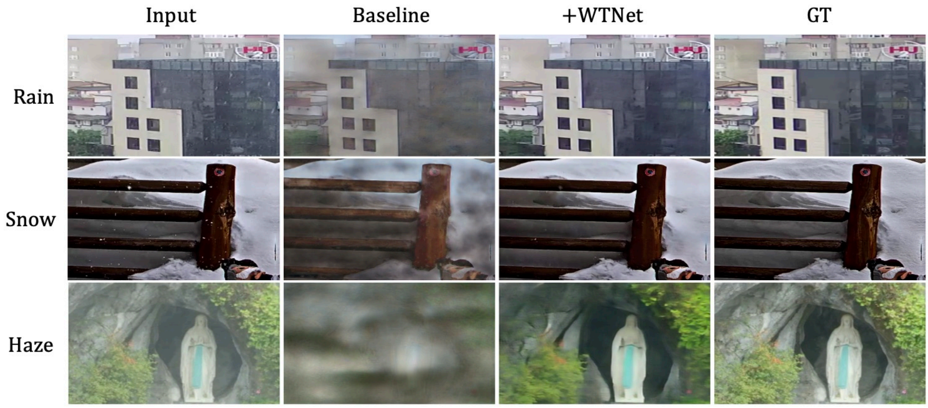

Figure 4: Qualitative results of PromptIR [30] on WeatherStream [48].

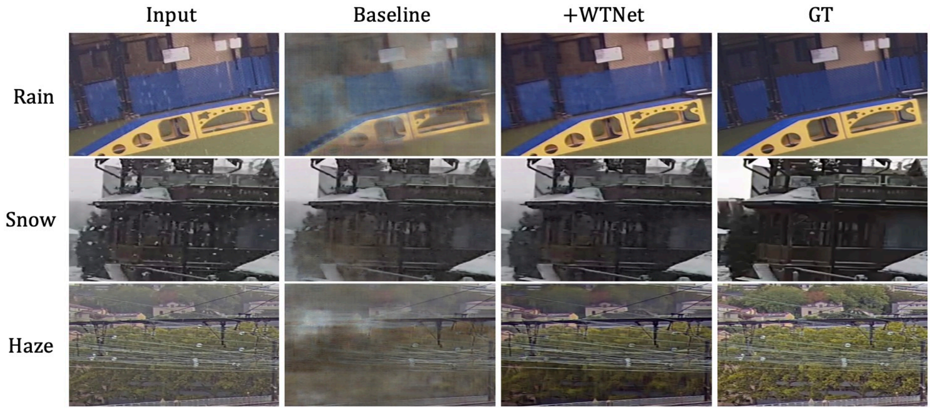

Figure 5: Qualitative results of AST [51] on the WeatherStream [48] dataset.

ness in narrowing the performance gap between general-purpose and task-specific restoration strategies.

Qualitative Comparison. We present qualitative results of the weather-transferred images in Figure 3. As shown, degradation patterns from WeatherStream [48] are faithfully transferred onto clean images from Rain100H [46], Snow100K [27], and RESIDE [22]. WTNet effectively disentangles weather-related degradations from scene content, enabling clean degradation transfer without introducing semantic distortion. Notably, real-world degraded images often exhibit mixed-weather conditions, such as haze co-occurring with rain or snow. WTNet successfully captures and transfers these compound degradations to clean images, thereby generating domain-adaptive fine-tuning sets that accurately reflect the degradation distributions present in the target domain.

We present qualitative de-weathering results on WeatherStream for three baseline models and their corresponding WTNet-enhanced versions. The baseline results are denoted as "Baseline," while the WTNet-enhanced outputs are denoted as "+WTNet." The comparisons are shown for PromptIR in Figure 4, AST in Figure 5, and AdaIR 6. Across these visualizations, the baseline models reveal several common limitations. In rain and snow scenarios, they often fail to fully remove degradations, resulting in residual artifacts and mild color distortions. In hazy scenes, they tend to struggle with recovering background structures, frequently producing overly smoothed or faded outputs. By contrast, WTNet-enhanced models

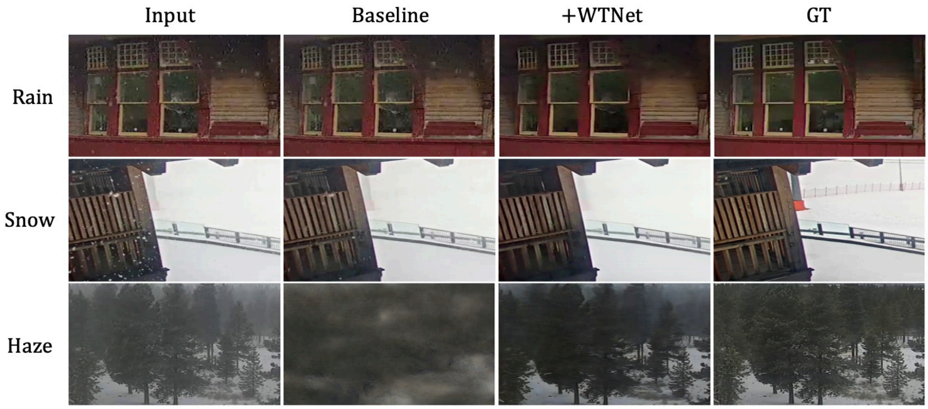

Figure 6: Qualitative results of AdaIR [7] on the WeatherStream [48] dataset.

	TM	Depth	HDD	ALD	MD	Rain	Snow	Haze	Avg.
Net1						17.09	15.73	16.42	16.41
Net2	✓					18.90	17.43	16.00	17.44
Net3		✓	✓			21.71	16.21	17.53	18.48
Net4		✓	✓	✓		20.29	18.16	16.77	18.41
Net5		✓	✓		✓	22.25	16.23	16.52	18.33
Net6		✓	✓	✓	✓	**24.24**	**21.71**	**20.04**	**22.00**

Table 2: Component analysis of WTNet for restoration performance of PromptIR [30] in PSNR (dB) on WeatherStream [48].

generate significantly improved results with sharper edges, clearer textures, and better structural preservation. These qualitative results highlight WTNet's superior generalization and restoration capability across a variety of challenging real-world weather conditions.

4.3 Ablation Studies

Component Analysis. In Table 2, we conduct an ablation study using PromptIR as the backbone model to evaluate each component in WTNet. Net1 represents the baseline model trained without WTNet. We first compare two strategies for generating transmission maps (TM): direct estimation (Net2) versus reconstruction using the Haze Density Decoder (HDD) and a pre-trained depth estimator (Depth) (Net3). The latter (Net3) outperforms direct estimation (Net2) and improves upon the baseline (Net1) by 2.07 dB on average.

Next, we assess the individual impact of the Atmospheric Light Decoder (ALD) and the Mask Decoder (MD) by incorporating them into Net3, forming Net4 and Net5, respectively. While neither component alone yields significant improvements over Net3, their combination in Net6 leads to a substantial average performance gain of 3.52 dB. Given that real-world adverse weather often involves compound degradations (e.g., haze with rain or snow), Net6, the final version of WTNet, integrates all weather-related parameters in a unified framework. This comprehensive design enables WTNet to robustly handle diverse degradation types, resulting in the best all-in-one performance across adverse weather restoration tasks.

Representation Analysis of A and β. To assess how different representations of atmospheric light A and haze density β affect restoration, we compare their scalar and spatial map forms in Table 3. Net1 and Net2 use mixed forms (scalar A with map β, and vice versa) and yield similar performance, suggesting limited individual impact. While using

	A: Map	A: Scalar	β: Map	β: Scalar	Rain	Snow	Haze	Avg.
Net1		✓	✓		22.15	21.34	18.65	20.71
Net2	✓			✓	22.62	21.24	18.49	20.78
Net3	✓		✓		23.22	21.58	18.64	21.15
Net4		✓		✓	**24.12**	**22.10**	**19.81**	**22.01**

Table 3: Representation analysis of atmospheric light A and haze density β, comparing spatial map and scalar forms, for restoration performance of PromptIR [30] in PSNR (dB) on WeatherStream [48].

maps for both parameters (Net3) improves results, scalar forms for both (Net4) achieve the best PSNR. This suggests that although spatial maps offer flexibility, scalar representations may generalize better on WeatherStream [48] due to their simplicity and robustness.

Method	Rain	Snow	Haze	Average
Baseline	19.31	17.31	12.37	16.33
+Noise-DA	19.63 (+0.32)	18.15 (+0.84)	13.48 (+1.11)	17.09 (+0.76)
+WTNet	**25.32 (+6.01)**	**23.04 (+5.73)**	**19.01 (+6.64)**	**22.46 (+6.13)**

Table 4: Quantitative comparison of domain adaptation methods, WTNet (ours) and Noise-DA [25], in terms of PSNR (dB) on WeatherStream [48], using the restoration backbone from Noise-DA as the baseline.

Comparison with other domain adaptation methods. We compare WTNet with Noise-DA [25], a training-time domain adaptation method, using the restoration backbone adopted in the original Noise-DA. As shown in Table 4, Noise-DA improves the baseline by only 0.76 dB, showing limited adaptation under complex weather conditions. In contrast, WTNet conducts test-time adaptation and enhances PSNR by 6.13 dB on average by transferring degradation patterns and constructing domain-adaptive fine-tuning sets.

Limitations. WTNet is designed to address adverse weather conditions by leveraging the inductive biases of weather formation. However, it may not directly apply to other types of degradation, such as blur, low-light, or noise, which remain open for future research.

5 Conclusion

This paper introduces Weather Transfer Network (WTNet), a domain adaptation framework that enhances all-in-one image restoration under adverse weather conditions at test time. WTNet transfers degradation patterns from target-domain images to source-domain clean images to construct domain-adaptive fine-tuning sets for test-time adaptation. It explicitly disentangles and reassembles key weather components—including snow masks, rain streaks, haze density, and atmospheric light—leveraging the inductive biases of weather formation for accurate degradation transfer. By utilizing these adaptive fine-tuning sets, WTNet dynamically adapts restoration models to both weather-specific and mixed-weather scenarios, improving generalization and performance. Experiments on real-world deraining, desnowing, and dehazing benchmarks show that WTNet consistently enhances restoration quality and outperforms state-of-the-art methods, validating its effectiveness and practicality for real-world adverse weather image restoration.

Acknowledgements

This work was supported in part by the National Science and Technology Council (NSTC) under grants 112-2221-E-A49-090-MY3, 114-2221-E-A49-038-MY3, and 113-2634-F-006-002. This work was funded in part by Qualcomm Technologies, Inc..

References

[1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *ICML*, 2017.

[2] Wei-Ting Chen, Hao-Yu Fang, Jian-Jiun Ding, Chen-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.

[3] Wei-Ting Chen, Hao-Yu Fang, Cheng-Lin Hsieh, Cheng-Che Tsai, I Chen, Jian-Jiun Ding, Sy-Yen Kuo, et al. ALL Snow Removed: Single image desnowing algorithm using hierarchical dual-tree complex wavelet representation and contradict channel loss. In *ICCV*, 2021.

[4] 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.

[5] Xiang Chen, Jinshan Pan, Kui Jiang, Yufeng Li, Yufeng Huang, Caihua Kong, Longgang Dai, and Zhentao Fan. Unpaired deep image deraining using dual contrastive learning. In *CVPR*, 2022.

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

[7] 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.

[8] 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.

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

[10] Raanan Fattal. Single image dehazing. *ACM TOG*, 2008.

[11] 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.

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

[13] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *NeurIPS*, 2017.

[14] 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.

[15] 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.

[16] Kaiming He, Jian Sun, and Xiaoou Tang. Single image haze removal using dark channel prior. In *ICCV*, 2009.

[17] Mu He, Le Hui, Yikai Bian, Jian Ren, Jin Xie, and Jian Yang. Ra-depth: Resolution adaptive self-supervised monocular depth estimation. In *ECCV*, 2022.

[18] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020.

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

[20] 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.

[21] 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.

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

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

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

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

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

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

[28] Xudong Mao, Qing Li, Haoran Xie, Raymond Y.K. Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *ICCV*, 2017.

[29] 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.

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

[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.

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

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

[34] Hoang Thanh-Tung, Truyen Tran, and Svetha Venkatesh. Improving generalization and stability of generative adversarial networks. In *ICLR*, 2019.

[35] Fu-Jen Tsai, Yan-Tsung Peng, Yen-Yu Lin, and Chia-Wen Lin. PHATNet: A physics-guided haze transfer network for domain-adaptive real-world image dehazing. In *ICCV*, 2025.

[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.

[37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.

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

[39] Hsing-Hua Wang, Fu-Jen Tsai, Yen-Yu Lin, and Chia-Wen Lin. TANet: Triplet attention network for all-in-one adverse weather image restoration. In *ACCV*, 2024.

[40] 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.

[41] 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.

[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.

[43] Jia-Hao Wu, Fu-Jen Tsai, Yan-Tsung Peng, Chung-Chi Tsai, Chia-Wen Lin, and Yen-Yu Lin. ID-Blau: Image deblurring by implicit diffusion-based reblurring augmentation. In *CVPR*, 2024.

[44] Ruiqi Wu, Zhengpeng Duan, Chunle Guo, Zhi Chai, and Chongyi Li. RIDCP: Revitalizing real image dehazing via high-quality codebook priors. In *CVPR*, 2023.

[45] Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with denoising diffusion GANs. In *ICLR*, 2022.

[46] 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.

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

[48] Howard Zhang, Yunhao Ba, Ethan Yang, Varan Mehra, Blake Gella, Akira Suzuki, Arnold Pfahnl, Chethan Chinder Chandrappa, Alex Wong, and Achuta Kadambi. WeatherStream: Light transport automation of single image deweathering. In *CVPR*, 2023.

[49] 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.

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

[51] Shihao Zhou, Duosheng Chen, Jinshan Pan, Jinglei Shi, and Jufeng Yang. Adapt or Perish: Adaptive sparse transformer with attentive feature refinement for image restoration. In *CVPR*, 2024.