# LEARNING RESTORATION ACTIONS FROM DEGRADATION STATES FOR MULTI-WEATHER IMAGE RESTORATION

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Multi-weather image restoration aims to recover clean images from diverse weather degradations within a unified model. Existing degradation-aware methods mainly infer degradation representations and use them to condition restoration features. However, knowing what degradation is present does not fully determine how it should be restored, as images with similar degradations may require different restoration operations depending on degradation severity, spatial distribution, and image content. To bridge this gap, we propose D2AIR, an efficient framework that learns input-adaptive restoration actions from degradation states. Its core Degradation-to-Action (D2A) module consists of a Degradation State Interpreter (DSI) and a Restoration Action Executor (RAE). DSI learns degradation-discriminative channel responses to infer degradation states and adapts latent features through degradation-aware modulation. RAE further translates the inferred states into input-adaptive restoration actions through a shared operator bank. Specifically, degradation-conditioned routing provides a degradation-dependent prior over restoration operators, while instance-conditioned routing adapts their selection to the current input, enabling flexible operator composition across degradation instances. Designed to operate at the latent bottleneck, D2A can be readily integrated into diverse restoration backbones. Experiments on single-weather, composite-degradation, and real-world benchmarks demonstrate the effectiveness of the proposed method. The code will be publicly available.

## 1 INTRODUCTION

Restoring images corrupted by diverse weather conditions within a single model requires handling substantially different degradation characteristics and restoration demands (Su et al., 2023; Jiang et al., 2025; Zamir et al., 2021; Chen et al., 2023). Recent methods have increasingly introduced adaptive mechanisms (Zamir et al., 2021; Fan et al., 2019), including prompts (Potlapalli et al., 2023; Wu et al., 2025b;a), conditional representations (Guo et al., 2024; Luo et al., 2024), and dynamic routing (Zamfir et al., 2025; An et al., 2025), allowing restoration features or computational paths to vary with the input. Such flexibility improves the capacity of a unified model to handle different degradations. However, this adaptation is commonly derived from latent features without explicit degradation discrimination, and it may produce redundant responses (Wu et al., 2025a; Zhu et al., 2024). This has motivated degradation-aware restoration, where explicit degradation information is introduced to provide more targeted guidance.

Obtaining useful degradation information within a restoration model, however, remains non-trivial (Wu et al., 2025b; Li et al., 2020; Wu et al., 2025c). A dedicated degradation perception network introduces additional parameters and computation (Zhang et al., 2026; Ai et al., 2024), while auxiliary classification on reconstruction features imposes a discriminative objective on representations primarily optimized for image reconstruction (Hu et al., 2025). More importantly, degradation categories provide only a coarse description of complex degradations: for example, images affected by the same rain degradation may differ substantially in rain density, intensity, and spatial distribution. Meanwhile, directly assigning degradation categories to specific experts or restoration branches can also be restrictive (Zhu et al., 2023), since different degradations may share similar restoration requirements; for instance, low-light and haze both involve global illumination adjust-



Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

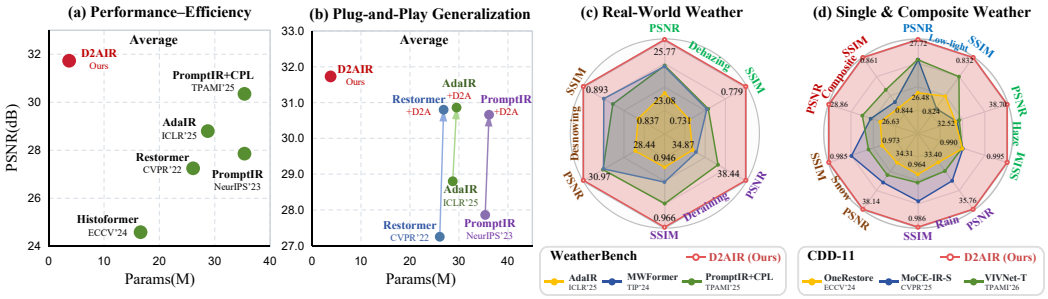



Figure 1: Overall performance comparison of D2AIR. (a) Performance–efficiency trade-off on WeatherBench (Guan et al., 2025). (b) Plug-and-play performance of D2A across different restoration backbones. (c) Real-world weather restoration on WeatherBench. (d) Single and composite degradation restoration on CDD-11 (Guo et al., 2024). Higher PSNR/SSIM indicates better performance.

ment, while rain and snow removal both require recovering locally occluded structures and textures (Zheng et al., 2019). These limitations suggest that effective adaptive restoration requires not only reliable degradation perception, but also a flexible mechanism that translates the inferred degradation state into suitable restoration responses.

We propose D2AIR, a unified multi-weather restoration framework that progressively transforms degradation perception into adaptive restoration. Its core Degradation-to-Action (D2A) module consists of a Degradation State Interpreter (DSI) and a Restoration Action Executor (RAE). DSI identifies degradation states from the channel-response patterns of latent features, avoiding direct classification constraints on spatial representations. A learnable degradation memory bank further enriches the inferred cues, while channel-spatial adaptation strengthens degradation-relevant responses to establish effective degradation awareness. RAE translates the inferred degradation state into adaptive restoration behaviors through a shared bank of lightweight operators. Rather than assigning a dedicated expert to each degradation category, it dynamically composes shared operators according to both degradation characteristics and image content, enabling cross-degradation reuse while adapting restoration to each individual image. In this way, D2AIR connects discriminative degradation perception with efficient, instance-adaptive restoration within a compact latent-space module that can be integrated into diverse restoration backbones. Extensive experiments on single-weather, composite-degradation, and real-world benchmarks (Guo et al., 2024; Valanarasu et al., 2022; Guan et al., 2025) demonstrate the effectiveness and generality of D2AIR, with consistent improvements across different restoration backbones at modest additional computational cost.

## 2 Related Work

### 2.1 All-in-One Image Restoration

All-in-one image restoration aims to handle multiple degradation types within a unified model. Recent methods have increasingly moved from shared restoration architectures toward input-adaptive modeling. Restormer (Zamir et al., 2022) establishes an efficient transformer architecture for high-resolution image restoration, which has been widely adopted as a backbone for unified restoration. PromptIR (Potlapalli et al., 2023) further introduces input-adaptive prompts to dynamically condition restoration features for different degradations. Subsequent methods improve such adaptation through dynamic feature transformation and specialized computation. AdaIR (Cui et al., 2025) employs input-adaptive modules for efficient all-in-one restoration, while MoCE-IR (Zamfir et al., 2025) introduces task-discriminative complexity experts to dynamically allocate restoration capacity. These developments demonstrate the importance of adapting representations and computation to heterogeneous restoration requirements.

### 2.2 Degradation-Aware Image Restoration

Degradation-aware methods explicitly model degradation characteristics to provide more targeted restoration guidance. AirNet (Li et al., 2022) learns degradation representations using contrastive

learning. DFPIR (Tian et al., 2025) uses degradation types to guide channel- and attention-wise feature perturbations, alleviating interference among heterogeneous restoration tasks. CPL (Wu et al., 2025a) addresses redundancy in adaptive degradation representations through sparse prompt modeling and contrastive regularization, improving their task discriminability. ExDA (Ren et al., 2026) also introduce explicit global degradation context into restoration features. Different from these approaches, D2AIR focuses on translating degradation perception into restoration execution: degradation states inferred from latent responses guide an instance-adaptive composition of shared restoration operators.

# 3 METHOD

## 3.1 OVERVIEW

As illustrated in Fig. 2, D2AIR adopts a generic encoder–decoder architecture and inserts a lightweight Degradation-to-Action (D2A) module at the latent bottleneck. Given a degraded image $I$, the encoder first extracts a latent representation $F$, which is then processed by the D2A module to infer the degradation characteristics and adapt the restoration behavior accordingly.

## 3.2 DEGRADATION STATE INTERPRETER

The Degradation State Interpreter (DSI) aims to identify degradation characteristics from latent features and enhance degradation-relevant responses for subsequent restoration. It performs two successive processes: degradation identification from channel responses and degradation enhancement with memory-guided feature adaptation.

**Degradation Identification.** Given the latent feature $F \in \mathbb{R}^{C\times H\times W}$, DSI first summarizes its channel responses through Global Average Pooling (GAP) and predicts a soft degradation distribution:

$$D_F = [d_1, d_2, \ldots, d_n] = \mathrm{Sigmoid}(\mathrm{MLP}(\mathrm{GAP}(F))), \tag{1}$$

where $D_F = [d_1, \ldots, d_n]$ describes the presence of $n$ degradation components. The sigmoid activation allows multiple components to be activated simultaneously, supporting both single and composite degradations. Performing degradation prediction on the globally aggregated channel responses reduces the direct classification constraint on spatially structured features used for reconstruction.

**Memory-Guided Degradation Enhancement.** The predicted distribution identifies the degradation composition but does not explicitly encode how individual degradation components should be represented in the latent space. We therefore introduce a learnable degradation memory bank:

$$M = [m_1, m_2, \ldots, m_n], \quad m_k \in \mathbb{R}^{C_m}, \tag{2}$$

where each prototype $m_k$ stores a learnable representation associated with the $k$-th degradation component. Given $D_F$, the corresponding prototypes are softly retrieved and aggregated:

$$\mathbf{m} = \sum_{k=1}^{n} d_k \cdot m_k. \tag{3}$$

The retrieved prior $\mathbf{m}$ is used to enhance degradation-relevant channel responses. Specifically, degradation-conditioned affine parameters are generated by:

$$\boldsymbol{\gamma}, \boldsymbol{\beta} = Split(Linear(\mathbf{m})), \tag{4}$$

and applied to $F$ through:

$$F_c = F \odot (1 + \boldsymbol{\gamma}) + \boldsymbol{\beta}. \tag{5}$$

This channel-wise modulation selectively adjusts feature responses according to the inferred degradation composition. To further account for the spatially non-uniform nature of weather degradations, we estimate a spatial response map from the channel-adapted feature:

$$\sigma^{h,w} = Tanh(Conv_{1\times1}(Gelu(Conv_{3\times3}(F_c)))), \tag{6}$$

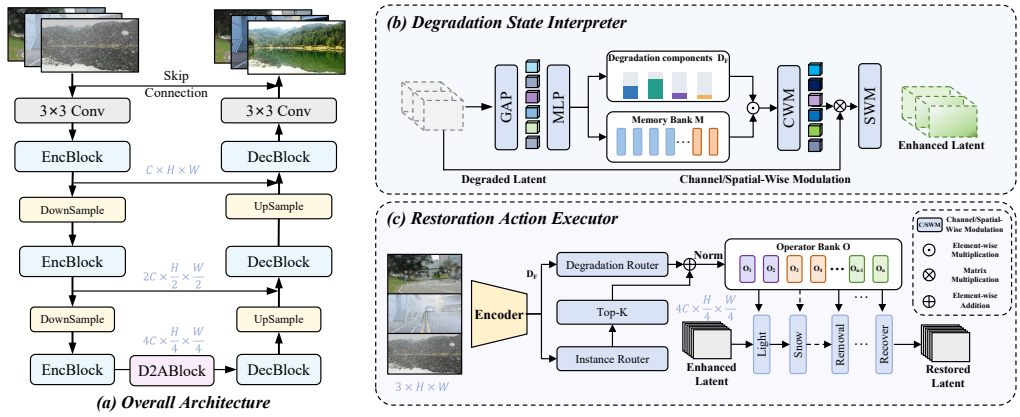



Figure 2: Overview of the proposed D2AIR framework. (a) D2A is inserted at the latent bottleneck of a generic encoder–decoder restoration network. (b) The Degradation State Interpreter (DSI) identifies degradation states from channel responses and enhances degradation-relevant features through memory-guided channel and spatial modulation. (c) The Restoration Action Executor (RAE) combines degradation-conditioned router and instance router to dynamically compose shared lightweight restoration operators.

and use it to obtain the degradation-aware representation:

$$F_d^{c,h,w} = F_c^{c,h,w} \left(1 + \sigma^{h,w}\right), \tag{7}$$

Thus, channel and spatial modulation respectively strengthen degradation-relevant responses and their spatial manifestations. DSI outputs $F_d$ for subsequent restoration and retains $D_F$ as explicit degradation guidance for operator routing in RAE.

### 3.3 Restoration Action Executor

The Restoration Action Executor (RAE) determines how the degradation-aware feature should be processed for restoration. While DSI identifies and enhances degradation characteristics, the required restoration operations can still vary across individual inputs. RAE therefore combines degradation-level guidance with the instance router to dynamically compose lightweight restoration operators. We first construct a lightweight restoration operator bank:

$$\mathcal{O} = [O_1, O_2, \ldots, O_N], \tag{8}$$

where each $O_k(\cdot)$ denotes a learnable restoration operator. The operators are jointly learned across all degradation conditions and serve as reusable restoration primitives rather than degradation-specific experts.

Given the degradation state $D_F$ estimated by DSI, we first predict a degradation-conditioned operator prior:

$$R_d = \text{Softmax}\left(\text{MLP}_d(D_F)\right), \tag{9}$$

where $R_d \in \mathbb{R}^N$ represents the degradation-driven preference over the operator bank.

The degradation state provides global guidance for operator selection, but inputs with similar degradation characteristics may still require different restoration responses due to variations in image content and local degradation patterns. We therefore introduce an instance router to refine the degradation-level routing by predicting input-specific operator preferences from the degradation-adapted latent feature. To keep this instance-level adaptation selective, only the top-$k$ ($k = N/n$) routing weights are retained, concentrating the input-dependent adjustment on a small subset of operators:

$$R_i = \text{top-k}(\text{Softmax}\left(\text{MLP}\left(\text{GAP}(F_d)\right)\right), k), \tag{10}$$

The final operation is obtained by fusing the two routing responses:

$$R = \text{Norm}\left(R_d + R_i\right), \tag{11}$$

Based on the fused routing weight, the executable operators are composed as:

$$F_o = \sum_{k=1}^{N} R_k O_k(F_d), \tag{12}$$

where $R_k$ is the weight of the $k$-th operator. The final restored feature is obtained through residual aggregation:

$$F_a = F_d + \lambda F_o, \tag{13}$$

where $\lambda$ is a learnable scaling factor.

RAE thus converts the degradation state into an input-adaptive composition of shared restoration operators, providing a flexible and efficient restoration mechanism for diverse degradation conditions. Together with DSI, this design establishes a structured transition from degradation estimation to input-adaptive restoration behavior.

### 3.4 Learning Objectives

D2AIR is trained with paired degraded-clean images and degradation labels $Y$, where the labels are used only during training to supervise degradation state learning. The overall objective consists of a pixel reconstruction loss, a frequency-domain loss, and a degradation classification loss:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{pix}} + \lambda_1 \mathcal{L}_{\text{fft}} + \lambda_2 \mathcal{L}_{\text{cls}}, \tag{14}$$

where

$$\mathcal{L}_{\text{pix}} = \|\hat{I} - I_{gt}\|_1, \qquad \mathcal{L}_{\text{cls}} = \text{BCE}(D_F, Y), \tag{15}$$

and $\mathcal{L}_{\text{fft}}$ measures the reconstruction discrepancy in the Fourier domain. We set $\lambda_1 = \lambda_2 = 0.1$ in all experiments.

## 4 Experiments

In this section, we conduct extensive experiments to evaluate the effectiveness of the proposed D2AIR framework for multi-weather image restoration. We first introduce the benchmark datasets, evaluation metrics, and implementation details used in our experiments. We then compare D2AIR with representative state-of-the-art image restoration methods under different degradation scenarios, including single-weather, composite-weather, and real-world adverse weather conditions. Finally, we carry out ablation experiments to analyze the effectiveness and rationality of each key component in our framework.

### 4.1 Experimental Setup

**Datasets and Metrics.** We evaluate D2AIR on three benchmarks covering different restoration scenarios: AllWeather (Valanarasu et al., 2022) for **Three-task All-in-one Restoration**, CDD-11 (Guo et al., 2024) for **Composite Degradation Removal**, and WeatherBench (Guan et al., 2025) for **Real-world Adverse Weather Restoration**. We follow the standard training and evaluation protocols of each benchmark. PSNR and SSIM are used as the main evaluation metrics, and LPIPS is additionally reported on WeatherBench for perceptual quality evaluation.

### 4.2 Implementation Details

We implement D2AIR using PyTorch 2.5.1 and conduct all experiments on four NVIDIA RTX 4090 GPUs. The proposed Degradation-to-Action module is inserted into the latent bottleneck of the restoration backbone (Cui et al., 2026a), and all components are optimized jointly in an end-to-end manner. During training, degradation category labels are used only for auxiliary degradation state supervision and are not required during inference.

For optimization, we use AdamW with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. All models are trained for 600K iterations with a batch size of 4 and $256 \times 256$ image patches. The learning rate is initialized to $1 \times 10^{-3}$ and decayed to $1 \times 10^{-7}$ following a cosine annealing schedule. Random horizontal and vertical flips are applied for data augmentation.

Table 1: Quantitative comparison on the AllWeather (Valanarasu et al., 2022) benchmark. Best and second-best results are highlighted in **bold** and underline, respectively.

| Method | Venue | Params. | Snow100K-S | | Snow100K-L | | Outdoor-Rain | | RainDrop | | Average | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | PSNR | SSIM | PSNR | SSIM | PSNR | SSIM | PSNR | SSIM | PSNR | SSIM |
| TransWeather (Valanarasu et al., 2022) | CVPR'22 | 38M | 32.51 | 0.934 | 29.31 | 0.888 | 28.83 | 0.900 | 30.17 | 0.916 | 30.20 | 0.909 |
| TKL (Chen et al., 2022) | CVPR'22 | 29M | 34.42 | 0.947 | 30.22 | 0.907 | 31.81 | 0.931 | 29.27 | 0.915 | 31.43 | 0.925 |
| WeatherDiff$_{64}$ (Özdenizci & Legenstein, 2023) | TPAMI'23 | 83M | 35.83 | 0.957 | 30.09 | 0.904 | 29.64 | 0.931 | 30.71 | 0.931 | 31.57 | 0.931 |
| GridFormer (Wang et al., 2024) | IJCV'24 | 30M | 37.46 | 0.964 | 31.71 | 0.923 | 31.87 | 0.934 | 32.39 | 0.936 | 33.36 | 0.939 |
| Histoformer (Sun et al., 2024) | ECCV'24 | 17M | 37.41 | 0.966 | 32.16 | 0.926 | 32.08 | 0.939 | **33.06** | **0.944** | 33.68 | 0.944 |
| AdaIR (Cui et al., 2025) | ICLR'25 | 29M | 36.12 | 0.960 | 30.92 | 0.916 | 30.84 | 0.920 | 31.45 | 0.931 | 32.33 | 0.932 |
| PromptIR+CPL (Wu et al., 2025a) | TPAMI'25 | 35M | 37.82 | 0.967 | 32.27 | 0.928 | 32.16 | 0.942 | 32.73 | 0.943 | 33.74 | 0.944 |
| ClearAIR (Zhang et al., 2026) | AAAI'26 | 31M | 37.79 | 0.967 | 32.53 | **0.932** | 32.45 | 0.941 | 32.82 | 0.942 | 33.90 | 0.946 |
| ExDA (Ren et al., 2026) | ICLR'26 | 11M | 37.97 | 0.968 | 32.49 | 0.930 | 32.70 | **0.945** | 32.66 | **0.944** | 33.92 | **0.947** |
| **D2AIR** | **Ours** | **4M** | **38.54** | **0.970** | **32.63** | **0.932** | **33.08** | 0.942 | 32.64 | **0.944** | **34.22** | **0.947** |

Table 2: Quantitative comparison on the CDD-11 (Guo et al., 2024) benchmark.

| Method | Venue | *CDD11-Single* | | | | | | | | *CDD11-Double* | | | | | | | | | | *CDD11-Triple* | | | | Avg. | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | Low (L) | | Haze (H) | | Rain (R) | | Snow (S) | | L+H | | L+R | | L+S | | H+R | | H+S | | L+H+R | | L+H+S | | | |
| AirNet | CVPR'22 | 24.83 | .778 | 24.21 | .951 | 26.55 | .891 | 26.79 | .919 | 23.23 | .779 | 22.82 | .710 | 23.29 | .723 | 22.21 | .868 | 23.29 | .901 | 21.80 | .708 | 22.24 | .725 | 23.75 | .814 |
| PromptIR | NeurIPS'23 | 26.32 | .805 | 26.10 | .969 | 31.56 | .946 | 31.53 | .960 | 24.49 | .789 | 25.05 | .771 | 24.51 | .761 | 24.54 | .924 | 23.70 | .925 | 23.74 | .752 | 23.33 | .747 | 25.90 | .850 |
| WGWSNet | CVPR'23 | 24.39 | .774 | 27.90 | .982 | 33.15 | .964 | 34.43 | .973 | 24.27 | .800 | 25.06 | .772 | 24.60 | .765 | 27.23 | .955 | 27.65 | .960 | 23.90 | .772 | 23.97 | .771 | 26.96 | .863 |
| WeatherDiff | TPAMI'23 | 23.58 | .763 | 21.99 | .904 | 24.85 | .885 | 24.80 | .888 | 21.83 | .756 | 22.69 | .730 | 22.12 | .707 | 21.25 | .868 | 21.99 | .868 | 21.23 | .716 | 21.04 | .698 | 22.49 | .799 |
| OneRestore | ECCV'24 | 26.48 | .826 | 32.52 | .990 | 33.40 | .964 | 34.31 | .973 | 25.79 | .822 | 25.58 | .799 | 25.19 | .789 | 29.99 | .957 | 30.21 | .964 | 24.78 | .788 | 24.90 | .791 | 28.47 | .878 |
| MoCE-IR-S | CVPR'25 | 27.26 | .824 | 32.66 | .990 | 34.31 | .970 | 35.91 | .980 | 26.24 | .817 | 26.25 | .800 | 26.04 | .793 | 29.93 | .964 | 30.19 | .970 | 25.41 | .789 | 25.39 | .790 | 29.05 | .881 |
| VIVNet-T | TPAMI'26 | 27.25 | .829 | 33.00 | .990 | 33.86 | .966 | 35.36 | .976 | 26.45 | .824 | 26.46 | .806 | 26.22 | .797 | 30.62 | .963 | 30.70 | .969 | 25.55 | .795 | 25.74 | .796 | 29.20 | .883 |
| ExDA | ICLR'26 | 27.60 | **.837** | 34.44 | .991 | 35.33 | .973 | 36.85 | .981 | 26.95 | .834 | 26.80 | .817 | 26.73 | .809 | 31.17 | .968 | 31.40 | .972 | 25.59 | .801 | 25.86 | .799 | 29.97 | .891 |
| Focal-CD | IJCV'26 | 27.28 | .836 | 36.90 | .994 | 35.21 | .973 | 36.86 | .980 | 26.98 | **.836** | 26.68 | .816 | 26.53 | .808 | 32.51 | .972 | 33.51 | .976 | 26.27 | **.809** | 26.25 | **.809** | 30.45 | .892 |
| **D2AIR** | **Ours** | **27.72** | .832 | **38.70** | **.995** | **35.76** | **.976** | **38.14** | **.985** | **27.24** | .831 | **26.98** | .813 | **27.06** | **.810** | **33.42** | **.975** | **34.44** | **.981** | **26.45** | .807 | **26.46** | .807 | **31.13** | **.893** |

## 4.3 Experimental Results

### 4.3.1 Three-Task All-in-One Restoration

Table 1 compares D2AIR with task-specific and all-in-one restoration methods on the AllWeather benchmark. D2AIR achieves the highest average PSNR of 34.22 dB and a joint-best average SSIM of 0.947, outperforming PromptIR+CPL (Wu et al., 2025a) and ClearAIR (Zhang et al., 2026) by 0.48 dB and 0.32 dB, respectively. It obtains the best PSNR on Snow100K-S, Snow100K-L (Liu et al., 2018), and Outdoor-Rain (Li et al., 2019), reaching 38.54 dB, 32.63 dB, and 33.08 dB, while remaining competitive on RainDrop (Qian et al., 2018). Notably, these results are achieved with approximately 4M parameters, substantially fewer than the 35M and 31M parameters of PromptIR+CPL and ClearAIR. These results demonstrate a favorable balance between restoration performance and model size across diverse weather degradations.

### 4.3.2 Composite Degradation Removal

Table 2 presents the results on CDD-11, which covers both single and composite degradations. D2AIR achieves the best average PSNR of 31.13 dB, outperforming Focal-CD (Cui et al., 2026b) by 0.68 dB. Under single degradations, D2AIR reaches 38.70 dB on Haze and 38.14 dB on Snow. The improvements extend to all evaluated composite conditions, with gains of 0.91 dB on Haze+Rain and 0.93 dB on Haze+Snow over Focal-CD. These results support the effectiveness of D2AIR in handling co-occurring degradations, where jointly modeling degradation components and adaptively combining shared operators helps accommodate mixed restoration requirements.

### 4.3.3 Real-world Adverse Weather Removal

We further evaluate D2AIR on the real-world WeatherBench (Guan et al., 2025) benchmark. As shown in Table 3, D2AIR achieves the best average PSNR and SSIM of 31.73 dB and 0.879, respectively, while maintaining competitive perceptual quality with an LPIPS of 0.221. Compared with PromptIR+CPL, D2AIR improves the average PSNR by 1.38 dB, with consistent gains across dehazing, deraining, and desnowing. Real-world weather degradations exhibit larger intra-class

Table 3: Quantitative comparison on the WeatherBench (Guan et al., 2025) benchmark.

| Method | Venue | Dehaze | | | Derain | | | Desnow | | | Average | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | PSNR | SSIM | LPIPS | PSNR | SSIM | LPIPS | PSNR | SSIM | LPIPS | PSNR | SSIM | LPIPS |
| ***General Restoration Methods*** | | | | | | | | | | | | | |
| DCMPNet | CVPR'24 | 21.18 | 0.506 | 0.491 | 32.04 | 0.876 | 0.282 | 24.81 | 0.614 | 0.546 | 26.01 | 0.665 | 0.440 |
| DRSformer | CVPR'23 | 19.95 | 0.694 | 0.404 | 33.98 | 0.943 | 0.209 | 28.00 | 0.836 | 0.197 | 27.31 | 0.824 | 0.270 |
| NeRD-Rain | CVPR'24 | 21.52 | 0.718 | 0.386 | 35.74 | 0.950 | 0.182 | 28.87 | 0.851 | 0.182 | 28.71 | 0.840 | 0.250 |
| MPRNet | CVPR'21 | 23.27 | 0.739 | 0.355 | 36.14 | 0.954 | 0.171 | 29.18 | 0.860 | 0.177 | 29.53 | 0.851 | 0.234 |
| DehazeFormer | TIP'23 | 24.12 | 0.745 | 0.345 | 36.05 | 0.954 | 0.181 | 28.88 | 0.849 | 0.178 | 29.68 | 0.849 | 0.235 |
| Restormer | CVPR'22 | 19.30 | 0.687 | 0.412 | 34.49 | 0.945 | 0.197 | 27.95 | 0.836 | 0.197 | 27.25 | 0.823 | 0.269 |
| ***All-in-One Restoration Methods*** | | | | | | | | | | | | | |
| AirNet | CVPR'22 | 20.94 | 0.705 | 0.383 | 33.59 | 0.942 | 0.224 | 22.06 | 0.780 | 0.291 | 25.53 | 0.809 | 0.299 |
| TransWeather | CVPR'22 | 19.79 | 0.680 | 0.397 | 29.34 | 0.903 | 0.294 | 24.96 | 0.796 | 0.231 | 24.70 | 0.793 | 0.307 |
| WGWS-Net | CVPR'23 | 13.79 | 0.603 | 0.535 | 37.08 | 0.961 | **0.117** | 20.81 | 0.780 | 0.248 | 23.89 | 0.781 | 0.300 |
| PromptIR | NeurIPS'23 | 21.11 | 0.713 | 0.375 | 34.54 | 0.944 | 0.198 | 27.93 | 0.836 | 0.195 | 27.86 | 0.831 | 0.256 |
| DiffUIR | CVPR'24 | 22.74 | 0.744 | 0.355 | 35.93 | 0.955 | 0.172 | 29.50 | 0.870 | 0.162 | 29.39 | 0.856 | 0.230 |
| MWFormer | TIP'24 | 24.42 | 0.746 | **0.284** | 35.15 | 0.951 | 0.153 | 29.98 | 0.872 | **0.133** | 29.85 | 0.856 | **0.190** |
| Histoformer | ECCV'24 | 17.69 | 0.669 | 0.437 | 30.70 | 0.916 | 0.279 | 25.39 | 0.808 | 0.225 | 24.59 | 0.798 | 0.314 |
| AdaIR | ICLR'25 | 23.08 | 0.731 | 0.351 | 34.87 | 0.946 | 0.192 | 28.44 | 0.837 | 0.179 | 28.80 | 0.838 | 0.240 |
| PromptIR+CPL | TPAMI'25 | 24.48 | 0.747 | 0.339 | 36.57 | 0.958 | 0.168 | 30.02 | 0.863 | 0.172 | 30.35 | 0.856 | 0.226 |
| **D2AIR** | **Ours** | **25.77** | **0.779** | 0.335 | **38.44** | **0.966** | 0.166 | **30.97** | **0.893** | 0.161 | **31.73** | **0.879** | 0.221 |

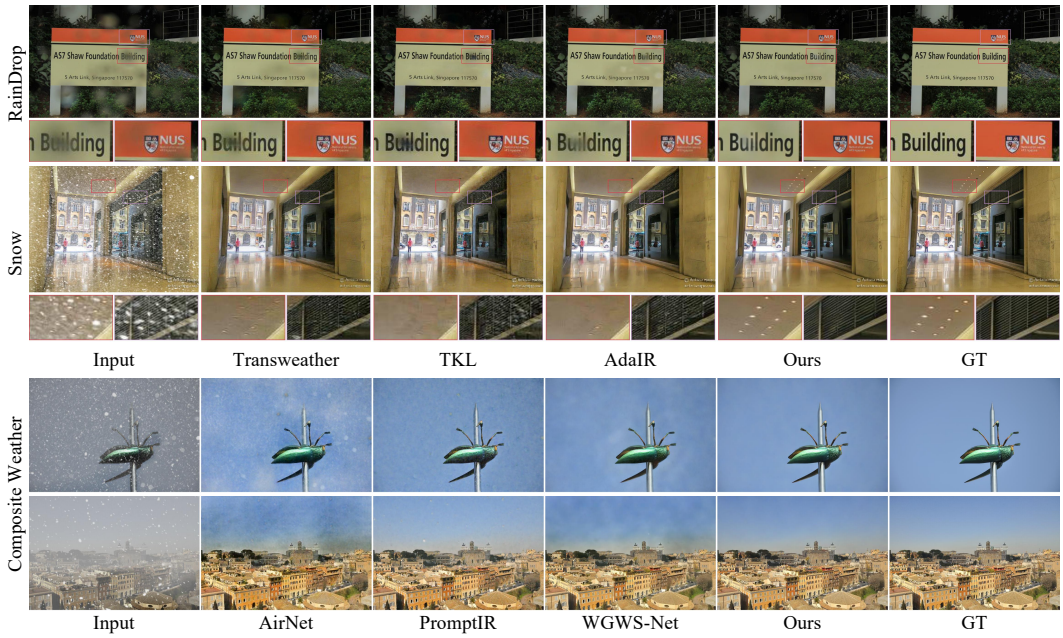



Figure 3: Visual comparisons on single and composite weather restoration. D2AIR effectively removes diverse weather degradations while preserving scene structures and fine details.

variations in severity and spatial distribution than synthetic ones. By adapting restoration actions to individual degradation patterns, D2AIR effectively accommodates such variations in real-world scenarios.

#### 4.3.4 Visual Results

Fig. 3 presents qualitative comparisons under diverse weather conditions. Existing methods may leave visible degradation residuals or introduce local over-/under-restoration, particularly under heavy and composite degradations. In contrast, D2AIR more effectively removes rain, snow, and mixed weather artifacts while preserving scene structures and fine details. The consistent visual quality across different degradation conditions further supports the ability of D2AIR to adapt restoration behavior to individual inputs.

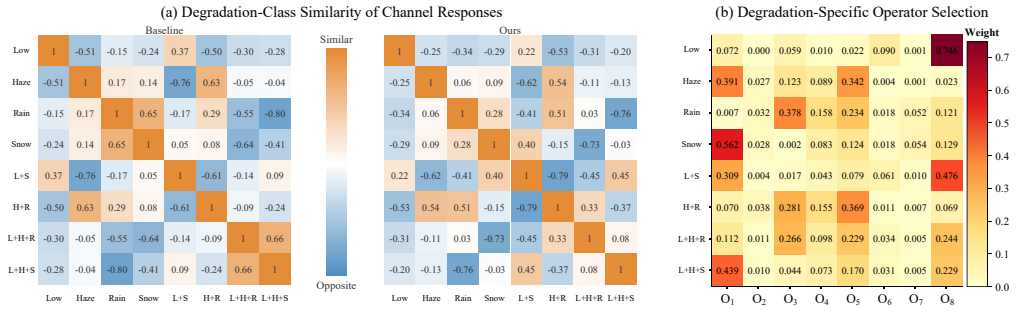



Figure 4: (a) Comparison of degradation-class channel-response similarities between the baseline (Cui et al., 2026a) and our D2AIR. (b) Degradation-specific operator selection, showing distinct and compositional operator preferences across degradations.

Table 4: Ablation studies on the key components of D2AIR. Mem. denotes the memory-guided channel calibration, Spa. denotes the spatial modulation, Ope. denotes the restoration operator bank, DeR denotes the degradation router, and InR denotes the instance router.

| Method | DSI | | RAE | | | Low-light (L) | | Haze (H) | | Rain (R) | | Snow (S) | | Average | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Mem. | Spa. | Ope. | DeR | InR | PSNR | SSIM | PSNR | SSIM | PSNR | SSIM | PSNR | SSIM | PSNR | SSIM |
| Baseline | × | × | × | × | × | 27.08 | 0.815 | 35.35 | 0.992 | 33.82 | 0.962 | 35.16 | 0.976 | 32.85 | 0.936 |
| w/o Mem. Calibration | × | ✓ | ✓ | ✓ | ✓ | 27.60 | 0.832 | 38.16 | 0.994 | 35.50 | 0.975 | 37.64 | 0.983 | 34.73 | 0.946 |
| w/o Spatial Modulation | ✓ | × | ✓ | ✓ | ✓ | 27.57 | 0.831 | 36.69 | 0.993 | 35.27 | 0.975 | 37.50 | 0.983 | 34.26 | 0.946 |
| w/o RAE | ✓ | ✓ | × | × | × | 27.44 | 0.829 | 36.44 | 0.994 | 34.79 | 0.972 | 36.88 | 0.982 | 33.89 | 0.944 |
| w/o Degradation Router | ✓ | ✓ | ✓ | × | ✓ | 27.32 | 0.829 | 36.17 | 0.994 | 34.99 | 0.973 | 37.14 | 0.983 | 33.90 | 0.944 |
| w/o Instance Router | ✓ | ✓ | ✓ | ✓ | × | 27.65 | 0.831 | 38.54 | 0.994 | 35.29 | 0.975 | 37.73 | 0.983 | 34.80 | 0.946 |
| D2AIR | ✓ | ✓ | ✓ | ✓ | ✓ | **27.72** | **0.832** | **38.70** | **0.995** | **35.76** | **0.976** | **38.14** | **0.985** | **35.08** | **0.947** |

## 4.4 Ablation Studies

### 4.4.1 Impact of Core Components

Table 4 evaluates the contribution of the major components in D2AIR. The full model improves the average PSNR from 32.85 dB to 35.08 dB over the baseline, demonstrating the effectiveness of explicit degradation interpretation and adaptive restoration execution. Removing memory-guided channel calibration or spatial modulation reduces the average PSNR to 34.73 dB and 34.26 dB, respectively, showing that both degradation-conditioned channel modulation and spatial adaptation contribute to learning degradation-aware representations. The larger performance drop without spatial modulation further highlights the importance of adapting features to spatially varying weather corruption. The RAE ablations examine the benefit of translating degradation awareness into adaptive restoration operations. Retaining the operator bank and instance router without degradation-conditioned routing yields only 33.90 dB, nearly identical to removing RAE entirely (33.89 dB). This suggests that input-adaptive operator selection alone provides limited benefit without explicit degradation guidance. In contrast, using degradation-conditioned routing improves the average PSNR to 34.80 dB, highlighting the importance of translating the inferred degradation state into restoration actions. Further incorporating the instance router improves the performance to 35.08 dB, showing that instance-specific adaptation complements the degradation-level prior by refining restoration actions for individual inputs.

### 4.4.2 Analysis of Degradation-Aware Representations and Restoration Actions

As shown in Fig. 4(a), baseline latent features already exhibit degradation-dependent channel responses, indicating that degradation information is implicitly encoded in restoration features. D2AIR further reorganizes these channel-response relationships, reducing strong correlations between easily confused degradations, such as rain and snow, while retaining selective relationships among related degradation states. Such structured degradation responses provide more informative

Table 5: Plug-and-play analysis of the proposed D2A module on the WeatherBench (Guan et al., 2025) dataset. We additionally report parameters, MACs, and inference time on $256 \times 256$ inputs.

| Method | Venue | Dehaze | | Derain | | Desnow | | Average | | Params | MACs | Time |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | PSNR | SSIM | PSNR | SSIM | PSNR | SSIM | PSNR | SSIM | (M) | (G) | (ms) |
| Restormer | CVPR'22 | 19.30 | 0.687 | 34.49 | 0.945 | 27.95 | 0.836 | 27.25 | 0.823 | 26.09 | 140.99 | 45.80 |
| Restormer+D2A | | **24.38** | **0.763** | **37.45** | **0.961** | **30.58** | **0.885** | **30.80** | **0.870** | 26.88 | 143.64 | 50.35 |
| PromptIR | NeurIPS'23 | 21.11 | 0.713 | 34.54 | 0.944 | 27.93 | 0.836 | 27.86 | 0.831 | 35.37 | 158.14 | 51.05 |
| PromptIR+D2A | | **24.19** | **0.761** | **37.35** | **0.961** | **30.45** | **0.882** | **30.66** | **0.868** | 36.16 | 160.79 | 56.83 |
| AdaIR | ICLR'25 | 23.08 | 0.731 | 34.87 | 0.946 | 28.44 | 0.837 | 28.80 | 0.838 | 28.74 | 147.18 | 59.22 |
| AdaIR+D2A | | **24.44** | **0.763** | **37.58** | **0.962** | **30.55** | **0.884** | **30.86** | **0.870** | 29.52 | 149.83 | 63.90 |
| D2AIR | Ours | 25.77 | 0.779 | 38.44 | 0.966 | 30.97 | 0.893 | 31.73 | 0.879 | 3.79 | 32.98 | 28.98 |

Table 6: Ablation study on the kernel size of executable restoration operators.

| Kernel Size | Haze (H) | | Rain (R) | | Snow (S) | | Average | |
|---|---|---|---|---|---|---|---|---|
| | PSNR | SSIM | PSNR | SSIM | PSNR | SSIM | PSNR | SSIM |
| $1 \times 1$ | 37.85 | 0.993 | 35.58 | 0.976 | 37.99 | 0.984 | 37.14 | 0.984 |
| $3 \times 3$ | **38.70** | **0.995** | **35.76** | **0.976** | **38.14** | **0.985** | **37.53** | **0.985** |
| $5 \times 5$ | 38.28 | 0.994 | 35.47 | 0.975 | 37.83 | 0.983 | 37.19 | 0.984 |
| $7 \times 7$ | 38.00 | 0.994 | 35.68 | **0.976** | 37.99 | **0.985** | 37.22 | **0.985** |

Table 7: Ablation study on the number of executable restoration operators.

| Operator Count | Haze (H) | | Rain (R) | | Snow (S) | | Average | |
|---|---|---|---|---|---|---|---|---|
| | PSNR | SSIM | PSNR | SSIM | PSNR | SSIM | PSNR | SSIM |
| $N = 4$ | 37.77 | **0.996** | 35.60 | 0.976 | 38.05 | 0.985 | 37.14 | **0.986** |
| $N = 8$ | **38.70** | 0.995 | **35.76** | 0.976 | **38.14** | 0.985 | **37.53** | 0.985 |
| $N = 12$ | 38.00 | 0.995 | 35.71 | 0.976 | 38.09 | 0.985 | 37.27 | 0.985 |

guidance for restoration action selection. As shown in Fig. 4(b), different degradations exhibit distinct operator preferences, while composite degradations flexibly reuse operators across different restoration requirements.

#### 4.4.3 Generalizability Across Diverse Architectures

We further integrate D2A into three representative restoration backbones to evaluate its architectural generality. As shown in Table 5, D2A consistently improves Restormer (Zamir et al., 2022), PromptIR (Potlapalli et al., 2023), and AdaIR (Cui et al., 2025), increasing their average PSNR by 3.55 dB, 2.80 dB, and 2.06 dB, respectively. These improvements are achieved with only about 0.8M additional parameters and 2.65G MACs for each backbone. The consistent gains across restoration architectures suggest that D2A does not depend on a particular backbone design and can serve as a lightweight latent-space adaptation module.

#### 4.4.4 Impact of Restoration Operators

Tables 6 and 7 study the configuration of the restoration operator bank. A $3 \times 3$ kernel achieves the highest average PSNR of 37.53 dB. Compared with $1 \times 1$, it enables effective local spatial modeling, while larger kernels may introduce excessive spatial aggregation that is less favorable for preserving fine-grained details. Regarding bank size, increasing the number of operators from 4 to 8 improves the average PSNR from 37.14 dB to 37.53 dB, whereas further increasing it to 12 reduces performance to 37.27 dB, suggesting that an oversized bank may introduce redundant operators with overlapping restoration behaviors. We therefore adopt eight $3 \times 3$ operators by default.

## 5 Conclusion

In this work, we presented D2AIR, a lightweight framework that moves beyond degradation-aware representation by explicitly connecting degradation interpretation with restoration execution. D2AIR interprets degradation states from latent feature responses and translates them into adaptive compositions of shared restoration operators, enabling restoration behavior to adapt to both degradation characteristics and individual inputs. Extensive experiments on single, composite, and real-world weather degradations demonstrate strong restoration performance with substantially fewer parameters, while consistent improvements across different backbones further validate the generality of the proposed design. These results demonstrate the effectiveness of translating degradation understanding into adaptive restoration actions for all-in-one image restoration.

### AI USE STATEMENT

We used ChatGPT to assist with language polishing and to check the manuscript for grammatical errors, unclear wording, and inconsistencies in notation and presentation. All AI-assisted revisions were reviewed by the authors, who take responsibility for the final content of the paper.

### ETHICS STATEMENT

This work studies multi-weather image restoration using the existing AllWeather, CDD-11, and WeatherBench datasets. We do not collect new data or conduct studies involving human participants.

### REPRODUCIBILITY STATEMENT

The main text and appendix describe the D2AIR architecture, training settings, datasets, and evaluation protocols. We also provide the configuration of the D2A module and describe its integration into different restoration backbones. Upon acceptance, we will release the source code, pretrained models, training configurations, and evaluation scripts to support reproduction of the reported results.

## REFERENCES

Yuang Ai, Huaibo Huang, Xiaoqiang Zhou, Jiexiang Wang, and Ran He. Multimodal prompt perceiver: Empower adaptiveness generalizability and fidelity for all-in-one image restoration. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 25432–25444, 2024.

Tao An, Hongbo Gao, Ruqi Liu, Kun Dai, Tao Xie, Ruifeng Li, Ke Wang, and Lijun Zhao. An moe-driven unified image restoration framework for adverse weather conditions. *IEEE Trans. Circuit Syst. Video Technol.*, 2025.

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 *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022.

Xiang Chen, Hao Li, Mingqiang Li, and Jinshan Pan. Learning A sparse transformer network for effective image deraining. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 5896–5905, 2023.

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 *Int. Conf. Learn. Represent.*, 2025.

Yuning Cui, Wenqi Ren, and Alois Knoll. Bio-inspired image restoration. *NIPS*, 38:80452–80481, 2026a.

Yuning Cui, Wenqi Ren, and Alois Knoll. Focal modulation for image restoration: Y. cui et al. *Int. J. Comput. Vis.*, 134(1):6, 2026b.

Qingnan Fan, Dongdong Chen, Lu Yuan, Gang Hua, Nenghai Yu, and Baoquan Chen. A general decoupled learning framework for parameterized image operators. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019.

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 Int. Conf. Multimedia*, pp. 12607–12613, 2025.

Yu Guo, Yuan Gao, Yuxu Lu, Ryan Wen Liu, and Shengfeng He. OneRestore: A universal restoration framework for composite degradation. In *Eur. Conf. Comput. Vis.*, pp. 255–272, 2024.

JiaKui Hu, Lujia Jin, Zhengjian Yao, and Yanye Lu. Universal image restoration pre-training via degradation classification. In *Int. Conf. Learn. Represent.*, 2025.

Junjun Jiang, Zengyuan Zuo, Gang Wu, Kui Jiang, and Xianming Liu. A survey on all-in-one image restoration: Taxonomy, evaluation and future trends. *IEEE Trans. Pattern Anal. Mach. Intell.*, 47 (12):11892–11911, 2025.

Boyun Li, Xiao Liu, Peng Hu, Zhongqin Wu, Jiancheng Lv, and Xi Peng. All-in-one image restoration for unknown corruption. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 17431–17441, 2022.

Ruoteng Li, Loong-Fah Cheong, and Robby T Tan. Heavy rain image restoration: Integrating physics model and conditional adversarial learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 1633–1642, 2019.

Ruoteng Li, Robby T Tan, and Loong-Fah Cheong. All in one bad weather removal using architectural search. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 3175–3185, 2020.

Yun-Fu Liu, Da-Wei Jaw, Shih-Chia Huang, and Jenq-Neng Hwang. Desnownet: Context-aware deep network for snow removal. *IEEE Trans. Image Process.*, 27(6):3064–3073, 2018.

Ziwei Luo, Fredrik K. Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B. Schön. Controlling vision-language models for multi-task image restoration. In *Int. Conf. Learn. Represent.*, 2024.

Ozan Özdenizci and Robert Legenstein. Restoring vision in adverse weather conditions with patch-based denoising diffusion models. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(8):10346–10357, 2023.

Vaishnav Potlapalli, Syed Waqas Zamir, Salman H. Khan, and Fahad Shahbaz Khan. PromptIR: Prompting for all-in-one image restoration. In *Adv. Neural Inform. Process. Syst.*, 2023.

Rui Qian, Robby T. Tan, Wenhan Yang, Jiajun Su, and Jiaying Liu. Attentive generative adversarial network for raindrop removal from a single image. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 2482–2491, 2018.

Bin Ren, Runyi Yang, Qi Ma, Xu Zheng, Mengyuan Liu, Danda Pani Paudel, Luc Van Gool, Rita Cucchiara, and Nicu Sebe. Rethinking expressivity and degradation-awareness in attention for all-in-one blind image restoration. In *Int. Conf. Learn. Represent.*, volume 2026, pp. 35910–35935, 2026.

Zhipeng Su, Yixiong Zhang, Jianghong Shi, and Xiao-Ping Zhang. A survey of single image rain removal based on deep learning. *ACM Computing Surveys*, 56(4):1–35, 2023.

Shangquan Sun, Wenqi Ren, Xinwei Gao, Rui Wang, and Xiaochun Cao. Restoring images in adverse weather conditions via histogram transformer. In *Eur. Conf. Comput. Vis.*, pp. 111–129, 2024.

Xiangpeng Tian, Xiangyu Liao, Xiao Liu, Meng Li, and Chao Ren. Degradation-aware feature perturbation for all-in-one image restoration. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 28165–28175. IEEE, 2025.

Jeya Maria Jose Valanarasu, Rajeev Yasarla, and Vishal M Patel. Transweather: Transformer-based restoration of images degraded by adverse weather conditions. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 2353–2363, 2022.

Tao Wang, Kaihao Zhang, Ziqian Shao, Wenhan Luo, Björn Stenger, Tong Lu, Tae-Kyun Kim, Wei Liu, and Hongdong Li. GridFormer: Residual dense transformer with grid structure for image restoration in adverse weather conditions. *Int. J. Comput. Vis.*, 132(10):4541–4563, 2024.

Gang Wu, Junjun Jiang, Kui Jiang, Xianming Liu, and Liqiang Nie. Beyond degradation redundancy: Contrastive prompt learning for all-in-one image restoration. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2025a.

Gang Wu, Junjun Jiang, Kui Jiang, Xianming Liu, and Liqiang Nie. Learning dynamic prompts for all-in-one image restoration. *IEEE Trans. Image Process.*, 2025b.

Gang Wu, Junjun Jiang, Yijun Wang, Kui Jiang, and Xianming Liu. Debiased all-in-one image restoration with task uncertainty regularization. In *AAAI*, pp. 8386–8394, 2025c.

Eduard Zamfir, Zongwei Wu, Nancy Mehta, Yuedong Tan, Danda Pani Paudel, Yulun Zhang, and Radu Timofte. Complexity experts are task-discriminative learners for any image restoration. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 12753–12763, 2025.

Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021.

Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 5718–5729, 2022.

Xu Zhang, Huan Zhang, Guoli Wang, Qian Zhang, and Lefei Zhang. Clearair: A human-visual-perception-inspired all-in-one image restoration. *AAAI*, 2026.

Qian Zheng, Boxin Shi, Xudong Jiang, Ling-Yu Duan, and Alex C. Kot. Denoising adversarial networks for rain removal and reflection removal. In *IEEE Int. Conf. Image Process.*, pp. 2766–2770, 2019.

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 *IEEE Conf. Comput. Vis. Pattern Recog.*, pp. 2952–2963, 2024.

Ruoxi Zhu, Zhengzhong Tu, Jiaming Liu, Alan C. Bovik, and Yibo Fan. Mwformer: Multi-weather image restoration using degradation-aware transformers. *IEEE Trans. Image Process.*, 33:6790–6805, 2024.

Yurui Zhu, Tianyu Wang, Xueyang Fu, Xuanyu Yang, Xin Guo, Jifeng Dai, Yu Qiao, and Xiaowei Hu. Learning weather-general and weather-specific features for image restoration under multiple adverse weather conditions. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2023.

# A Experimental Settings

## A.1 Datasets

**Three-task All-in-one Restoration.** AllWeather (Valanarasu et al., 2022) is a widely used benchmark for all-in-one adverse weather image restoration, constructed by combining images from existing weather restoration datasets, including Snow100K (Liu et al., 2018), Outdoor-Rain (Li et al., 2019), and RainDrop (Qian et al., 2018). These subsets cover heterogeneous weather degradations, including snow, rain/haze, and raindrops. Following the standard protocol, we jointly train a single model on the mixed training set and evaluate it on the corresponding weather-specific test sets.

**Composite Degradation Removal.** CDD-11 (Guo et al., 2024) is a composite degradation image restoration benchmark designed to evaluate restoration under both individual and coexisting degradations. It contains 11 degradation configurations constructed from low-light, haze, rain, and snow, including four individual degradations and seven compound combinations: Low, Haze, Rain, Snow, Low+Haze, Low+Rain, Low+Snow, Haze+Rain, Haze+Snow, Low+Haze+Rain, and Low+Haze+Snow. Each degraded image is paired with its corresponding clean ground truth. Compared with benchmarks dominated by individual weather degradations, CDD-11 provides a more challenging setting for evaluating whether a unified model can interpret and handle multiple coexisting degradation factors.

**Real-world Adverse Weather Restoration.** WeatherBench (Guan et al., 2025) is a large-scale real-world benchmark for all-in-one adverse weather image restoration. Unlike benchmarks assembled primarily from synthetic or independently collected single-weather datasets, WeatherBench provides precisely aligned degraded/clean image pairs captured under real-world rain, snow, and haze conditions, covering diverse outdoor scenes and illumination conditions. The dataset contains 42,002 paired images at $512 \times 512$ resolution, including 41,402 pairs for training and 600

pairs for testing, with 200 test pairs for each weather condition. Its real-world acquisition and diverse weather conditions provide a challenging benchmark for evaluating restoration performance in practical adverse-weather scenarios.

Table A1: Detailed architecture configuration of D2AIR. DSI and RAE denote the Degradation State Interpreter and Restoration Action Executor, respectively.

| | Item | Configuration |
|---|---|---|
| Backbone | Architecture | BioIR (Cui et al., 2026a) |
| | Feature dimensions | [32, 64, 128] |
| | Encoder / decoder blocks | [3, 3, 6] / [3, 3, 6] |
| | Refinement blocks | 3 |
| | FFN expansion factor | 3 |
| DSI | Placement | Bottleneck latent space |
| | Degradation memory bank | $4 \times 224$ |
| | State predictor | MLP: $128 \rightarrow 160 \rightarrow 4$; sigmoid |
| | Channel modulation | Linear: $224 \rightarrow 256$; 128-D scale ($\gamma$) and bias ($\beta$) |
| | Spatial modulation | Conv $3\times3$ ($128 \rightarrow 32$) $\rightarrow$ GELU $\rightarrow$ Conv $1\times1$ ($32 \rightarrow 1$) $\rightarrow$ tanh |
| RAE | Restoration operators | 8 |
| | Operator architecture | DWConv $3\times3 \rightarrow$ PWConv $\rightarrow$ GELU $\rightarrow$ PWConv |
| | Instance router | MLP: $128 \rightarrow 64 \rightarrow 8$; Top-2 selection |
| | Degradation router | Linear: $4 \rightarrow 8$; softmax |

Feature dimensions and block numbers are listed from fine to coarse resolution. DWConv and PWConv denote depthwise and $1\times1$ pointwise convolutions, respectively.

### A.2 Evaluation Metrics

Following previous evaluation protocols (Li et al., 2020; Wang et al., 2024; Guo et al., 2024; Zhou et al., 2024), we use PSNR and SSIM as the main quantitative metrics. For CDD-11 (Guo et al., 2024), PSNR and SSIM are computed in the RGB color space, while for the other benchmarks they are computed on the Y channel to ensure fair comparison with prior methods. In addition, we report LPIPS on WeatherBench (Guan et al., 2025) to evaluate perceptual quality under real-world adverse weather conditions. Higher PSNR and SSIM indicate better reconstruction fidelity, while lower LPIPS indicates better perceptual similarity.

### A.3 Implementation Details

Table A1 summarizes the detailed architecture configuration of D2AIR. For the generalizability experiments, D2A is consistently inserted into the deepest latent feature space of Restormer (Zamir et al., 2022), PromptIR (Potlapalli et al., 2023), and AdaIR (Cui et al., 2025), while the remaining backbone architectures are kept unchanged from their corresponding baselines.

## B Additional Visual Analysis

We provide additional qualitative comparisons on AllWeather, CDD-11, and WeatherBench to further evaluate D2AIR under diverse adverse weather conditions. The visual results cover both single and composite degradations, as well as real-world weather conditions. Overall, D2AIR consistently removes weather-induced artifacts while preserving image structures and fine details, demonstrating its effectiveness across different degradation scenarios.

### B.1 Analysis of Learned Restoration Operators

To better understand the behavior of the shared operator bank, we visualize the spatial responses of individual restoration operators under different degradation conditions. As shown in Fig. A1, dif-

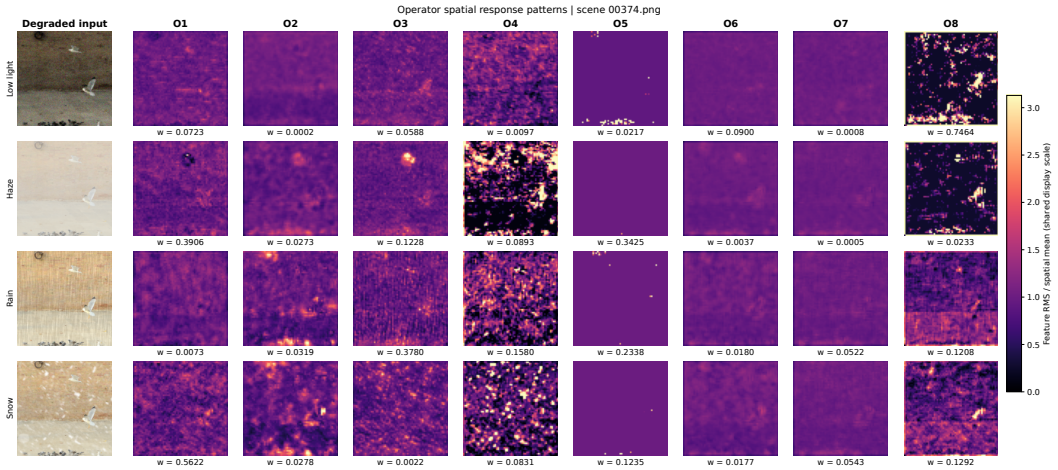



Figure A1: Spatial responses of learned restoration operators under different degradation conditions. Each column corresponds to one restoration operator, and each row represents a degradation type. Different operators exhibit distinct spatial response patterns, while the same operator can respond to multiple degradation types with degradation-dependent distributions, indicating complementary and reusable restoration behaviors.

ferent operators exhibit clearly different response patterns for the same degraded input, suggesting that the operator bank learns different transformations rather than redundant branches. For example, some operators produce relatively smooth and globally distributed responses, whereas others exhibit stronger localized or high-frequency responses. Meanwhile, the same operator can respond to multiple degradation types, but its spatial response varies with the degradation presented in the input. This indicates that the learned operators are not exclusively tied to individual degradation categories. Instead, they act as reusable restoration primitives whose effects are modulated by the current degradation characteristics and image content. Such behavior is consistent with the design of RAE, which dynamically composes a shared set of operators rather than assigning a dedicated expert to each degradation category.

### B.2 Results on AllWeather

Fig. A2 presents additional visual comparisons on the AllWeather benchmark. We include representative examples from different weather conditions to evaluate the restoration capability of different methods. D2AIR effectively removes weather artifacts while better preserving structural and textural details, producing visually cleaner and more faithful restoration results.

### B.3 Results on CDD-11

Fig. A3 presents comparisons under challenging composite degradation conditions involving low-light, haze, rain, and snow. Existing methods can alleviate the dominant degradation but may leave residual weather artifacts or exhibit imperfect brightness and contrast recovery when multiple degradations coexist. In comparison, D2AIR more effectively handles the coupled restoration requirements, simultaneously improving visibility and suppressing weather artifacts while preserving scene structures. The results on different degradation compositions further demonstrate the effectiveness of degradation-aware restoration actions for composite degradation removal.

### B.4 Results on WeatherBench

Fig. A4 shows qualitative comparisons on WeatherBench under real-world haze, rain, and snow conditions. These cases contain spatially non-uniform weather effects and diverse scene content, making restoration more challenging than controlled synthetic degradation. Under haze, D2AIR improves scene visibility while maintaining natural structures and appearance. For real rain and

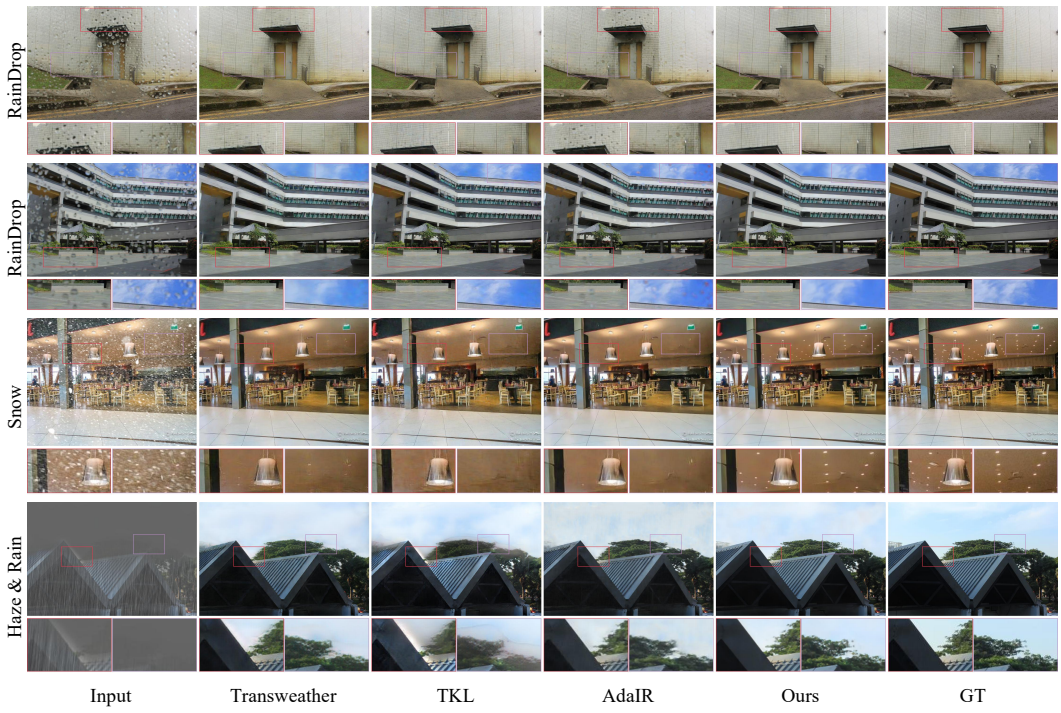



Figure A2: Visual comparisons on the AllWeather dataset. Our method achieves prominent visual gains on raindrop removal, snow removal and composite-degradation tasks. It successfully removes snow artifacts while retaining lamp structures visually similar to snowflakes, which are incorrectly erased by other approaches, verifying our model's adaptiveness to image content.

snow, it effectively suppresses weather artifacts while retaining object boundaries and fine textures. The consistent visual performance across the three real-world weather conditions demonstrates the applicability of D2AIR to practical adverse weather restoration.

## C Limitations and Broader Impact

D2AIR currently learns degradation states from a predefined degradation space. Although it can model both individual and composite degradations, its behavior on completely unseen degradation types remains to be explored. In addition, the degradation-aware modulation and restoration operator bank introduce additional computation despite their relatively small parameter overhead. Future work may extend D2AIR toward open-set degradation modeling and more efficient action execution.

Multi-weather restoration can improve image quality for vision systems operating under adverse weather, with potential applications in autonomous driving, robotics, and outdoor monitoring. However, restoration models may modify ambiguous visual details under severe degradation. Restored images should therefore be used cautiously in safety-critical or forensic scenarios where faithful preservation of the original observation is essential.

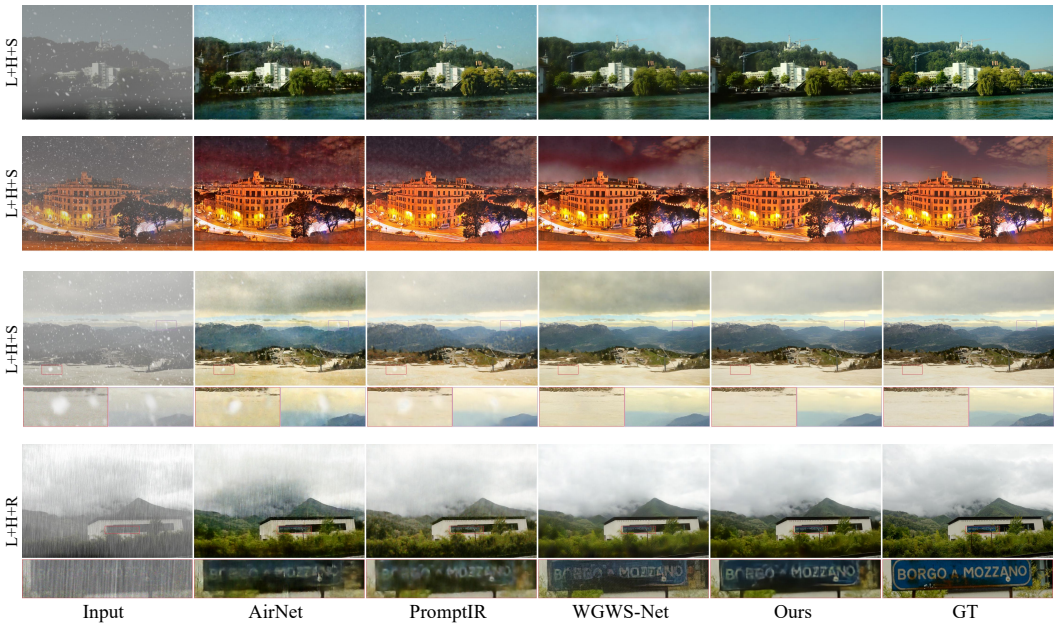



Figure A3: Visual comparisons on composite weather restoration.

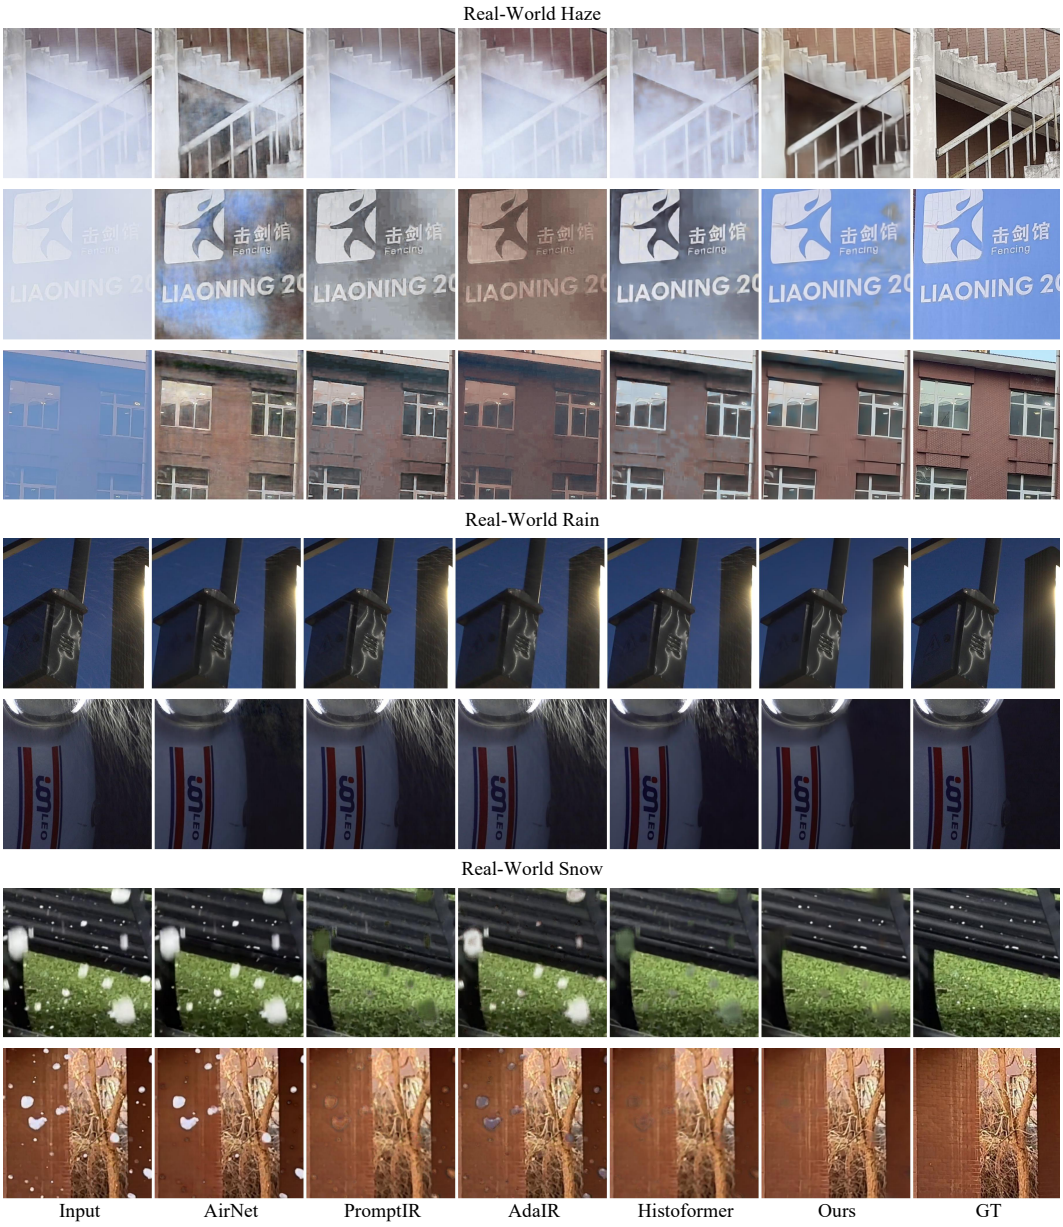



Figure A4: Visual comparisons on real-world weather restoration. Our method is evaluated on complex real-world dehazing, deraining and desnowing scenarios, which are more challenging than synthetic counterparts. It obtains noticeable gains in color restoration for heavy haze and rain-streak elimination. For highly similar patterns including snow spots and paint-peeling textures on black chairs, our model effectively differentiates artifacts to be removed from real scene content to be retained.