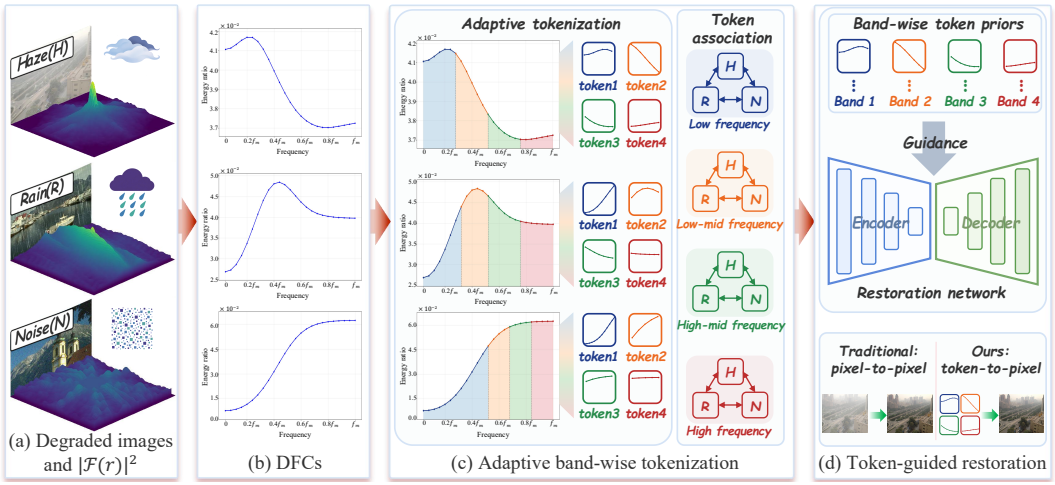
# DEGRADATION FREQUENCY CURVE: AN EXPLICIT FREQUENCY-QUANTIFIED REPRESENTATION FOR ALL-IN-ONE IMAGE RESTORATION

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




(a) Degraded images and $|\mathcal{F}(r)|^2$ (b) DFCs (c) Adaptive band-wise tokenization (d) Token-guided restoration

Figure 1: Motivation of the proposed DFC-guided restoration paradigm. (a) Degraded images and residual frequency-domain energy maps; (b) DFCs derived from residual-to-degraded energy ratios; (c) Adaptive DFC tokenization into band-wise spectral tokens, where same-band tokens across degradations form reusable spectral priors; (d) Token-guided degradation-aware restoration.

## ABSTRACT

A fundamental difficulty in all-in-one blind image restoration is that degradation is usually treated as an implicit factor hidden in degraded-to-clean mapping, rather than as an explicit object that can be measured and manipulated. This limitation becomes more pronounced under mixed, compound, or unseen degradation conditions, where degradation effects are hard to assign to predefined labels or task-specific parameters. We propose the Degradation Frequency Curve (DFC), a structured spectral representation that quantifies degradation responses by measuring band-wise residual-to-degraded energy ratios in the frequency domain. DFC maps visually entangled degradation responses into a measurable degradation coordinate space. Moreover, DFC can be adaptively decomposed into band-wise spectral tokens, allowing local degradation responses to be represented as reusable restoration priors. Based on this representation, we develop the DFC-guided Image Restorer (DFC-IR), a token-conditioned multi-scale framework that progressively estimates DFCs from intermediate restorations and uses the resulting spectral tokens to guide degradation-aware restoration in a coarse-to-fine manner. Extensive experiments on standard, composite, unseen, and real-world degradation benchmarks show that DFC provides an effective representation basis for all-in-one restoration, leading to state-of-the-art performance and improved generalization under complex degradation profiles.

## 1 INTRODUCTION

Real-world images are frequently degraded by complex, overlapping degradations such as noise (Lin et al., 2023; Zhang et al., 2024), blur (Cho et al., 2021; Kupyn et al., 2019; Wang et al., 2020),



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.

rain (Chen & Li, 2021; Jiang et al., 2020), and haze (Zeng & Dai, 2016), presenting a fundamental challenge for all-in-one image restoration. Most existing methods, whether based on feature-space unification or parameter specialization, still follow a pixel-to-pixel mapping paradigm. In this formulation, degradation cues are implicitly encoded in network features or parameters rather than explicitly represented. Consequently, degradation patterns themselves remain difficult to directly characterize and manipulate, particularly under compound or unseen degradations.

Current methods mainly follow two paradigms. Feature-space unification methods (Yang et al., 2024; Potlapalli et al., 2023; Conde et al., 2024) attempt to learn shared representations for diverse degradations through contrastive or prompt-based learning, but often suffer from feature entanglement and suboptimal performance. Parameter-specialization methods (Li et al., 2022; Park et al., 2023; Wang et al., 2023b) allocate dedicated sub-networks, experts, or parameter sets for different degradation types, but often incur substantial parameter overhead and remain prone to misclassification under ambiguous or compound scenarios. Crucially, neither stream explicitly formulates degradation itself as a quantifiable and structured representation, which would enable degradation patterns to be directly measured, compared, and used to guide degradation-aware restoration.

Our approach is motivated by a key empirical observation, as shown in Fig. 1: while different degradations are visually entangled in the pixel domain, they exhibit structured and distinguishable response patterns across frequency bands (Fig. 1(a)). To quantify these responses, we introduce the Degradation Frequency Curve (DFC), a frequency-quantified representation of degradation, as illustrated in Fig. 1(b). Specifically, DFC measures the band-wise energy ratio between residual and degraded images across discrete frequency bands. This relative formulation helps suppress content-dependent spectral variations and highlights degradation-induced responses. Importantly, unlike existing frequency-aware restoration methods that mainly use frequency as a processing operator for feature modulation (Cui et al., 2025) and routing (Cui et al., 2023b;a), DFC treats the frequency axis as a shared measurement coordinate for directly measuring and comparing degradation responses across samples. As quantified in Sec. 3.2, the average Wasserstein-1 distance between normalized DFC profiles of different degradation types is $6.07\times$ the average within-degradation distance, indicating distinct degradation patterns with reduced content sensitivity.

The explicit frequency coordinate provided by DFC naturally supports structured band-wise decomposition. As illustrated in Fig. 1(c), we introduce adaptive band-wise tokenization, which partitions each DFC along the frequency axis into a set of localized spectral tokens, as indicated by the color-coded band segments. Each token captures the frequency location, response intensity, and local curve pattern of a specific band. More importantly, tokens from the same frequency band across different degradations can be associated to capture shared spectral patterns and degradation-specific variations. In this way, the global DFC is converted into reusable band-wise degradation units, providing structured spectral priors for characterizing complex and unseen degradations.

To translate these band-wise spectral priors into restoration guidance, we develop a DFC-guided Image Restorer (DFC-IR), a multi-scale restoration framework, as illustrated in Fig. 1(d). Rather than relying solely on a direct pixel-to-pixel mapping, DFC-IR performs restoration in a token-to-pixel manner by using band-wise DFC tokens as degradation-aware priors. Specifically, at each scale of the coarse-to-fine decoder, a scale-specific DFC is obtained to describe the current degradation response, and the Frequency Curve-guided Mask Sampling (FCMS) module adaptively converts it into band-wise spectral tokens. These tokens are then fed into the Multi-band Modulated Decoding (MMD) module, which performs token-conditioned band-wise feature modulation and adaptive aggregation for restoration. Through this coarse-to-fine process, DFC-IR progressively translates explicit degradation responses into restoration guidance without requiring degradation-type labels.

The main contributions are summarized as follows:

- We propose DFC, an explicit, quantifiable, and interpretable spectral representation that uses frequency as a shared measurement coordinate to characterize degradation responses through band-wise residual-to-degraded energy ratios, enabling degradation patterns to be directly measured and compared.
- We develop DFC-IR, a multi-scale restoration framework built on DFC. It integrates FCMS for adaptive DFC tokenization and MMD for token-guided feature modulation and aggregation, enabling coarse-to-fine token-to-pixel restoration.

- Extensive experiments on standard all-in-one, composite, unseen-degradation, and real-world benchmarks demonstrate that the proposed framework achieves state-of-the-art restoration performance and strong generalization to complex and unseen degradations.

## 2 Related Work

**All-in-one Image Restoration.** All-in-one image restoration aims to handle diverse degradations using a unified model. Recent methods mainly follow two directions. The first introduces degradation-related conditioning into a shared backbone, such as learnable prompts in PromptIR (Potlapalli et al., 2023), human instructions in InstructIR (Conde et al., 2024), scene descriptors in OneRestore (Guo et al., 2024b), and vision-language guidance in VLU-Net (Zeng et al., 2025). The second enhances a unified model with adaptive specialization, for example through frequency-aware modulation in AdaIR (Cui et al., 2025), degradation-aware feature perturbation in DFPIR (Tian et al., 2025), and complexity experts in MoCE-IR (Zamfir et al., 2025). Together, these designs improve model flexibility by enriching degradation guidance or adapting processing pathways to different degradations. However, they mainly focus on incorporating degradation cues into restoration, rather than explicitly characterizing degradation itself in a quantifiable and structured form.

**Frequency-aware Image Restoration.** Frequency-domain processing has been widely explored in image restoration. Early studies focused on fixed spectral transformations, including local spectral estimation (Yoo et al., 2018), representation learning in Fourier space (Xu et al., 2020), and frequency-domain objectives such as Focal Frequency Loss (Jiang et al., 2021). Recent methods further introduce adaptive frequency manipulation, including selective frequency decomposition and routing (Cui et al., 2023b;a) and frequency-aware modulation in AdaIR (Cui et al., 2025). These studies demonstrate the effectiveness of frequency cues, from spectral reconstruction and optimization to adaptive feature modulation. However, they mainly exploit frequency information at the feature level, rather than explicitly characterizing degradation itself in the spectral domain. In contrast, DFC characterizes degradation as a compact band-wise spectral representation for restoration.

## 3 Degradation Frequency Curve Modeling

The core of our approach is the DFC, a frequency-domain representation that characterizes image degradation through band-wise spectral responses. By representing degradation as an explicit band-wise spectral curve, DFC provides an interpretable prior for all-in-one image restoration.

### 3.1 DFC Formulation

The DFC is defined as a $B$-dimensional vector $D(f) \in \mathbb{R}^B$ extracted from a degraded image $\boldsymbol{y} \in \mathbb{R}^{H\times W\times 3}$ and its clean counterpart $\boldsymbol{x} \in \mathbb{R}^{H\times W\times 3}$, where $B$ denotes the number of frequency bands and $f$ denotes the frequency. Each element of $D(f)$ quantifies the degradation response within a specific frequency band, and the whole curve provides a compact spectral description of the degradation. We first compute the residual image $\boldsymbol{r} = \boldsymbol{y} - \boldsymbol{x}$ to isolate the degradation-induced component, and then construct the DFC from $\boldsymbol{r}$ and $\boldsymbol{y}$ through two key steps:

**Spectral transformation and partitioning.** We perform two preparatory operations in the frequency domain: 1) obtaining the spectral energy maps of both $\boldsymbol{r}$ and $\boldsymbol{y}$ using the operator $|\mathcal{F}(\cdot)|^2$, where $\mathcal{F}(\cdot)$ denotes the Fourier transform; 2) partitioning the frequency domain into $B$ contiguous bands using a set of predefined ring-shaped Gaussian masks $\{\mathcal{M}_b\}_{b=1}^{B}$. The $b$-th mask is defined as:

$$\mathcal{M}_b(h,w) = \exp\left(-\frac{(d(h,w)-\mu_b)^2}{2\sigma_b^2}\right), \tag{1}$$

where $d(h,w)$ denotes the $\ell_2$ distance from a pixel $(h,w)$ to the DC component in the frequency domain, $\mu_b$ is the central frequency of the $b$-th band, and $\sigma_b$ controls its bandwidth.

**Band-wise energy ratio computation.** For the $b$-th frequency band, we obtain the band-wise energy maps of the residual and degraded images as

$$E_r^b = \mathcal{M}_b \odot |\mathcal{F}(\boldsymbol{r})|^2, \qquad E_y^b = \mathcal{M}_b \odot |\mathcal{F}(\boldsymbol{y})|^2, \tag{2}$$

where $\odot$ denotes element-wise multiplication. The degradation response of this band is then quantified by the residual-to-degraded energy ratio:

$$R_b = \frac{\sum_{h,w} E_r^b(h,w)}{\sum_{h,w} E_y^b(h,w)}. \tag{3}$$

Collecting the band-wise responses $\{R_b\}_{b=1}^{B}$ along the frequency axis forms the DFC $D(f)$.

### 3.2 Properties of DFC

$D(f)$ exhibits three key properties for degradation representation: content robustness, degradation sensitivity, and severity-aware response. These properties show that DFC can provide a stable and discriminative description of degradation-related frequency responses.

**Content robustness.** Following the degradation synthesis strategy of AgenticIR (Zhu et al., 2025), we generate haze, rain, and noise degradations on the MiO100 (Kong et al., 2024) dataset and visualize the mean DFC profiles with standard deviation bands for each degradation. As shown in the first row of Fig. 2, the normalized DFC profiles of the same degradation remain within a relatively narrow range despite variations in image content. This demonstrates that DFC primarily captures degradation-related frequency responses rather than content-dependent spectral patterns.

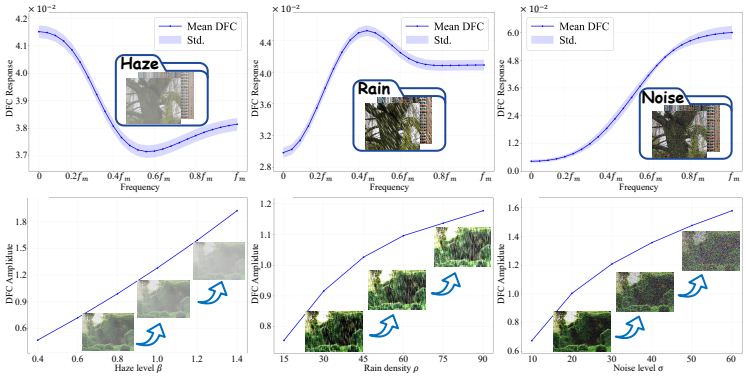



Figure 2: Statistical analysis of DFC properties. The first row shows mean normalized DFC profiles with standard deviation bands under different degradations, and the second row presents average DFC amplitudes under different severity levels.

**Degradation sensitivity.** Beyond content robustness, DFC also exhibits sensitivity to different degradation types. As shown in the first row of Fig. 2, different degradations exhibit distinct DFC profiles, indicating that they induce different frequency responses. To quantitatively evaluate this property, we measure the Wasserstein distance between normalized DFC profiles. The average distance between different degradation types is 6.07$\times$ the average distance within the same degradation type, confirming that DFC shapes are more discriminative across degradation types than across image contents. Detailed metric definitions and quantitative results are provided in Appendix B.

**Severity-aware response.** By adjusting the corresponding degradation parameters of haze, rain, and noise, we construct degraded samples with different severity levels and visualize the average DFC amplitudes under each severity level. As shown in the second row of Fig. 2, the DFC amplitudes exhibit consistent variation trends as the degradation severity increases. This indicates that DFC not only characterizes degradation types but also their severity levels.

### 3.3 Decomposing DFC into Spectral Tokens

Based on the above properties, we use DFC as an explicit degradation prior for frequency-aware restoration. However, using the entire DFC as a global condition may obscure its frequency-dependent local responses. Accordingly, as illustrated in the left side of Fig. 3, we decompose each DFC along the frequency axis into band-wise spectral tokens, where each token encodes the response intensity and spectral pattern within a local frequency segment. This tokenization transforms the global degradation curve into a set of localized spectral representations for more precise frequency-aware degradation modeling. The detailed tokenization strategy is introduced in Sec. 4.1.

Moreover, DFC tokenization provides reusable band-wise priors for unseen degradations. Although an unseen degradation may exhibit a distinct global DFC profile, its local frequency responses can

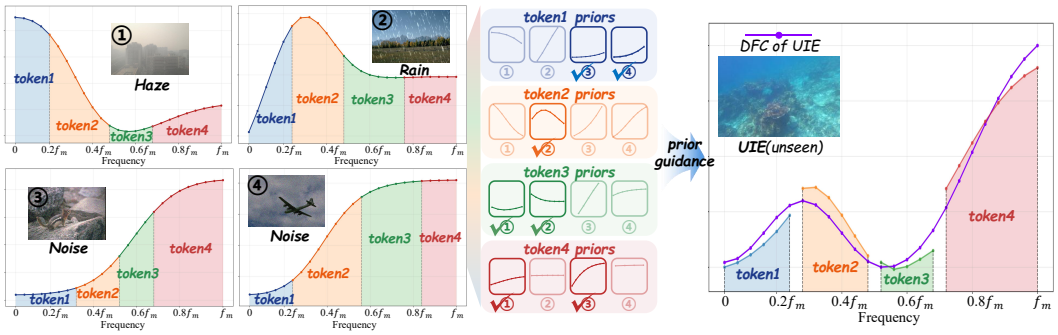



Figure 3: Illustration of DFC tokenization and reusable band-wise priors for unseen degradations. Seen degradation DFCs are decomposed into band-wise spectral tokens, and tokens within each frequency band form reusable spectral priors for characterizing unseen degradation responses.

share common patterns with seen degradations. As shown in the right side of Fig. 3, band-wise token priors are constructed from tokens extracted from seen degradations, providing reusable spectral patterns within each frequency band. These priors serve as references for characterizing the local frequency responses of unseen degradations, such as the underwater degradation shown in the figure. This structured band-wise representation allows local degradation patterns to collectively characterize complex and unseen degradations.

## 4 DFC-Guided Image Restoration

Building on the DFC representation, we develop DFC-IR to translate explicit degradation responses into restoration guidance. As shown in Fig. 4(a), DFC-IR adopts a multi-scale encoder-decoder architecture that takes a three-level degraded image pyramid $\boldsymbol{Y} = \{\boldsymbol{y}^i\}_{i=0}^2$ as input. The encoder extracts a latent feature $\boldsymbol{z}^3$ from $\boldsymbol{Y}$, while the decoder progressively restores the image from coarse to fine through three scales $(s = 3, 2, 1)$. The core of DFC-IR lies in its DFC-guided decoding process. At each scale, the Frequency Curve-guided Mask Sampling (FCMS) module first obtains the scale-specific DFC $D^s(f)$ to characterize the current degradation response and then adaptively converts $D^s(f)$ together with the latent feature $\boldsymbol{z}^s$ into $N$ band-wise spectral tokens:

$$\{\boldsymbol{t}_i^s\}_{i=1}^N = \mathcal{T}_{\text{FCMS}}\left(\boldsymbol{z}^s, D^s(f)\right). \tag{4}$$

The resulting tokens are then fed into the Multi-band Modulated Decoding (MMD) module, which performs token-conditioned feature modulation and produces the refined feature and intermediate restoration at the next finer scale:

$$\boldsymbol{z}^{s-1}, \hat{\boldsymbol{x}}^{s-1} = \mathcal{R}_{\text{MMD}}\left(\boldsymbol{z}^s, \{\boldsymbol{t}_i^s\}_{i=1}^N\right). \tag{5}$$

This process is repeated across scales, forming a coarse-to-fine restoration pipeline that progressively translates DFCs into band-wise restoration priors.

### 4.1 Adaptive DFC Tokenization

As shown in Fig. 4(b), FCMS obtains the scale-specific DFC $D^s(f)$ from the degraded image $\boldsymbol{y}^s$ and intermediate restoration $\hat{\boldsymbol{x}}^s$ as in Sec. 3.1 (uniform at the initial scale). It then adaptively partitions $D^s(f)$ into $N$ bands and extracts spectral tokens from the latent feature $\boldsymbol{z}^s$.

Specifically, we divide $D^s(f)$ into $N$ equal-area regions and initialize parameters $\hat{\boldsymbol{\theta}}_i^s = (\hat{\mu}_i^s, \hat{\sigma}_i^s)$ using the center and bandwidth of each region. This equal-area partition balances degradation-response mass across tokens. To further adapt the partitions, we sample $K$ candidate parameters $\{\boldsymbol{\theta}_{i,k}^s = (\mu_{i,k}^s, \sigma_{i,k}^s)\}_{k=1}^K$ around $\hat{\boldsymbol{\theta}}_i^s$. A lightweight network $\Phi(\cdot)$ predicts their weights:

$$\boldsymbol{\omega}_i^s = \text{softmax}\left(\Phi(\{\boldsymbol{\theta}_{i,k}^s\}_{k=1}^K)\right). \tag{6}$$

The refined Gaussian parameters $\boldsymbol{\theta}_{i*}^s = (\mu_{i*}^s, \sigma_{i*}^s)$ are then aggregated via weighted summation:

$$\mu_{i*}^s = \sum_{k=1}^K \omega_{i,k}^s \mu_{i,k}^s, \qquad \sigma_{i*}^s = \sum_{k=1}^K \omega_{i,k}^s \sigma_{i,k}^s. \tag{7}$$

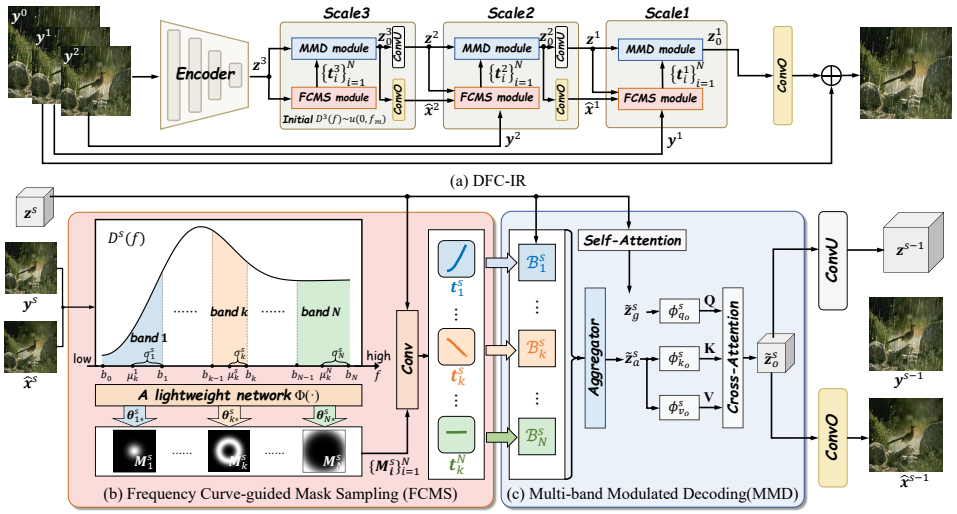



Figure 4: (a) The overall framework of DFC-IR. (b) FCMS for adaptive DFC tokenization and spectral token extraction. (c) MMD for token-guided feature modulation and aggregation.

Using these parameters, we construct the Gaussian masks $\{\boldsymbol{M}_i^s\}_{i=1}^N$ via Eq. 1 and extract tokens $\{\boldsymbol{t}_i^s\}_{i=1}^N$ through learnable convolution blocks. The token for the $i$-th band is formulated as:

$$\boldsymbol{t}_i^s = \mathrm{conv}_i\left(\mathcal{F}^{-1}\left(\mathcal{F}(\mathrm{conv}_s(\boldsymbol{z}^s)) \odot \boldsymbol{M}_i^s\right)\right), \tag{8}$$

where $\mathrm{conv}_s$ and $\mathrm{conv}_i$ are shared and band-specific convolutions. Each token couples local DFC responses with spectral components of $\boldsymbol{z}^s$, providing band-wise priors for subsequent decoding.

### 4.2 Band-wise Token Decoding

As shown in Fig. 4(c), MMD uses the DFC tokens $\{\boldsymbol{t}_i^s\}_{i=1}^N$ as band-wise priors to modulate the latent feature $\boldsymbol{z}^s$. Specifically, each token $\boldsymbol{t}_i^s$ modulates a band-wise branch via cross-attention:

$$\tilde{\boldsymbol{z}}_i^s = \mathrm{softmax}\left(\frac{\phi_{q_i}^s(\boldsymbol{z}^s)\left(\phi_{k_i}^s(\boldsymbol{t}_i^s)\right)^T}{\sqrt{d}}\right)\phi_{v_i}^s(\boldsymbol{t}_i^s), \tag{9}$$

where $\phi_{q_i}^s$, $\phi_{k_i}^s$, and $\phi_{v_i}^s$ denote linear projections. The resulting token-conditioned features $\{\tilde{\boldsymbol{z}}_i^s\}_{i=1}^N$ are then adaptively aggregated. A lightweight aggregator $\mathcal{G}(\cdot)$ predicts their weights:

$$\{\alpha_i^s\}_{i=1}^N = \mathrm{softmax}\left(\mathcal{G}(\{\boldsymbol{t}_i^s\}_{i=1}^N)\right). \tag{10}$$

The aggregated band-wise feature is computed as $\tilde{\boldsymbol{z}}_a^s = \sum_{i=1}^N \alpha_i^s \tilde{\boldsymbol{z}}_i^s$.

To maintain global structural coherence, MMD derives a global context feature $\tilde{\boldsymbol{z}}_g^s$ from $\boldsymbol{z}^s$ through self-attention and fuses it with $\tilde{\boldsymbol{z}}_a^s$ via cross-attention. The fused feature $\boldsymbol{z}_o^s$ is then processed by ConvU and ConvO to produce $\boldsymbol{z}^{s-1}$ and the intermediate restoration $\hat{\boldsymbol{x}}^{s-1}$, respectively.

### 4.3 Training Objective

DFC-IR is optimized end-to-end with multi-scale supervision. To preserve spatial-domain restoration accuracy, we apply a spatial reconstruction loss to the DFC-guided outputs at all scales:

$$\mathcal{L}_s = \sum_{s\in\{1,2,3\}} \left\|\mathcal{R}_{\mathrm{MMD}}^x\left(\boldsymbol{z}^s, \mathcal{T}_{\mathrm{FCMS}}\left(\boldsymbol{z}^s, D^s(f)\right)\right) - \boldsymbol{x}^{s-1}\right\|_1, \tag{11}$$

where $\mathcal{R}_{\mathrm{MMD}}^x$ denotes the image-output branch of MMD. We further introduce a frequency-domain consistency loss to reduce spectral discrepancies between the restored output and clean target:

$$\mathcal{L}_f = \sum_{s\in\{1,2,3\}} \left\|\mathcal{F}\left(\mathcal{R}_{\mathrm{MMD}}^x\left(\boldsymbol{z}^s, \mathcal{T}_{\mathrm{FCMS}}\left(\boldsymbol{z}^s, D^s(f)\right)\right)\right) - \mathcal{F}\left(\boldsymbol{x}^{s-1}\right)\right\|_1. \tag{12}$$

This term complements spatial supervision by enforcing frequency-domain consistency with the clean target. The total objective is $\mathcal{L} = \lambda_s\mathcal{L}_s + \lambda_f\mathcal{L}_f$, where $\lambda_s$ and $\lambda_f$ balance the two losses.

Table 1: Quantitative comparison (PSNR/SSIM) of all-in-one image restoration on the three-degradation setting. The best and second-best results are **highlighted** and underlined, respectively; missing values are denoted by "–".

| Type | Method | Reference | Params. | Dehazing | | Deraining | | Denoising | | | | | | Average | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | SOTS | | Rain100L | | $BSD68_{\sigma=15}$ | | $BSD68_{\sigma=25}$ | | $BSD68_{\sigma=50}$ | | PSNR | SSIM |
| General | MPRNet (Zamir et al., 2021) | CVPR'21 | 16.3M | 28.00 | .958 | 33.86 | .958 | 33.27 | .920 | 30.76 | .871 | 27.29 | .761 | 30.63 | .894 |
| | Restormer (Zamir et al., 2022) | CVPR'22 | 26.1M | 27.78 | .958 | 33.78 | .958 | 33.72 | .930 | 30.67 | .865 | 27.63 | .792 | 30.75 | .901 |
| | MambaIR (Guo et al., 2024a) | ECCV'24 | 27.2M | 29.57 | .970 | 35.42 | .969 | 33.88 | .931 | 30.95 | .874 | 27.74 | .793 | 31.51 | .907 |
| All-in-one | InstructIR (Conde et al., 2024) | ECCV'24 | 15.8M | 30.22 | .959 | 37.98 | .978 | 34.15 | .933 | 31.52 | .890 | 28.30 | .804 | 32.43 | .913 |
| | AdaIR (Cui et al., 2025) | ICLR'25 | 28.8M | 31.06 | .980 | 38.64 | .983 | 34.12 | .935 | 31.45 | .892 | 28.19 | .802 | 32.69 | .918 |
| | VLU-Net (Zeng et al., 2025) | CVPR'25 | 35.4M | 30.71 | .980 | **38.93** | .984 | 34.13 | .935 | 31.48 | .892 | 28.23 | .804 | 32.70 | .919 |
| | MoCE-IR (Zamfir et al., 2025) | CVPR'25 | 25.4M | 31.34 | .979 | 38.57 | .984 | 34.11 | .932 | 31.45 | .888 | 28.18 | .800 | 32.73 | .917 |
| | DFPIR (Tian et al., 2025) | CVPR'25 | 31.1M | 31.87 | .980 | 38.65 | .982 | 34.14 | .935 | 31.47 | .893 | 28.25 | **.806** | 32.88 | .919 |
| | R2R (Wang et al., 2026) | CVPR'26 | 19.7M | 31.40 | .977 | 37.46 | .980 | 34.10 | **.936** | 31.45 | **.895** | 28.22 | **.806** | 32.53 | .918 |
| | FIT (He et al., 2026) | ICML'26 | – | 31.11 | .981 | 38.92 | .984 | **34.23** | **.936** | **31.55** | .894 | **28.32** | **.806** | 32.83 | **.920** |
| | DFC-IR | Ours | 33.1M | **32.28** | **.982** | 38.74 | **.985** | 34.19 | .933 | 31.53 | .890 | 28.28 | .804 | **33.00** | .919 |

Table 2: Quantitative comparison (PSNR/SSIM) for all-in-one restoration on five degradation tasks.

| Type | Method | Reference | Params. | Dehazing | | Deraining | | Denoising | | Deblurring | | Low-light | | Average | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | SOTS | | Rain100L | | $BSD68_{\sigma=25}$ | | GoPro | | LOLv1 | | PSNR | SSIM |
| General | DGUNet (Mou et al., 2022) | CVPR'22 | 17.3M | 24.78 | .940 | 36.62 | .971 | 31.10 | .883 | 27.25 | .837 | 21.87 | .823 | 28.32 | .891 |
| | Restormer (Zamir et al., 2022) | CVPR'22 | 26.1M | 24.09 | .927 | 34.81 | .962 | 31.49 | .884 | 27.22 | .829 | 20.41 | .806 | 27.60 | .881 |
| | MambaIR (Guo et al., 2024a) | ECCV'24 | 27.2M | 25.81 | .944 | 36.55 | .971 | 31.41 | .884 | 28.61 | .875 | 22.49 | .832 | 28.97 | .901 |
| All-in-one | InstructIR (Conde et al., 2024) | ECCV'24 | 15.8M | 27.10 | .956 | 36.84 | .973 | 31.40 | .887 | 29.40 | .886 | 23.00 | .836 | 29.55 | .907 |
| | VLU-Net (Zeng et al., 2025) | CVPR'25 | 35.4M | 30.84 | .980 | 38.54 | .982 | 31.43 | .891 | 27.46 | .840 | 22.29 | .833 | 30.11 | .905 |
| | AdaIR (Cui et al., 2025) | ICLR'25 | 28.8M | 30.53 | .978 | 38.02 | .981 | 31.35 | .889 | 28.12 | .858 | 23.00 | .845 | 30.20 | .910 |
| | MoCE-IR (Zamfir et al., 2025) | CVPR'25 | 25.4M | 30.48 | .974 | 38.04 | .982 | 31.34 | .887 | 30.05 | .899 | 23.00 | .852 | 30.58 | .919 |
| | DFPIR (Tian et al., 2025) | CVPR'25 | 31.1M | 31.64 | .979 | 37.62 | .978 | 31.29 | .889 | 28.82 | .873 | **23.82** | .843 | 30.64 | .913 |
| | R2R (Wang et al., 2026) | CVPR'26 | 19.7M | 30.64 | .974 | 36.61 | .975 | 31.35 | .891 | **30.93** | **.911** | 22.88 | .856 | 30.48 | .921 |
| | FIT (He et al., 2026) | ICML'26 | – | 30.81 | .979 | **38.56** | **.983** | **31.57** | **.893** | 29.52 | .890 | 23.13 | .848 | 30.72 | .919 |
| | DFC-IR | Ours | 33.1M | **32.27** | **.982** | 37.91 | **.983** | 31.55 | .891 | 29.62 | .892 | **23.82** | **.863** | **31.03** | **.922** |

Table 3: Quantitative comparison on the CDD11 dataset under complex composite degradations.

| Method | Params. | CDD11-Single | | | | | | | | CDD11-Double | | | | | | | | CDD11-Triple | | | | Average | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | Low (L) | | Haze (H) | | Rain (R) | | Snow (S) | | L+H | | L+R | | L+S | | H+R | | H+S | | L+H+R | | L+H+S | | PSNR | SSIM |
| AirNet (Li et al., 2022) | 9M | 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 (Potlapalli et al., 2023) | 36M | 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 (Zhu et al., 2023) | 26M | 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 (Özdenizci & Legenstein, 2023) | 83M | 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 (Guo et al., 2024b) | 6M | 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 (Zamfir et al., 2025) | 11M | 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 |
| R2R (Wang et al., 2026) | 19.7M | 27.29 | **.836** | 32.88 | .992 | 35.12 | .971 | 36.54 | .981 | 26.21 | **.833** | 26.70 | **.816** | 26.33 | **.809** | 30.41 | .966 | 30.37 | .971 | 25.47 | **.804** | 25.21 | **.803** | 29.32 | .889 |
| DFC-IR | 33.1M | **27.71** | .831 | **34.94** | **.993** | **35.41** | **.976** | **37.22** | **.983** | **26.87** | .828 | **26.87** | .812 | **26.85** | .806 | **31.94** | **.973** | **32.15** | **.977** | **26.27** | **.804** | **26.25** | **.803** | **30.22** | **.890** |

## 5 EXPERIMENTS

We conduct extensive experiments following the protocols of prior all-in-one image restoration works (Li et al., 2022; Potlapalli et al., 2023; Zhang et al., 2025; 2023). Our evaluation covers both standard restoration settings, including three-degradation, five-degradation, and composite-degradation benchmarks, and challenging generalization settings involving unseen degradations and real-world compound degradations. Detailed dataset descriptions, implementation settings, evaluation metrics, and additional results are provided in the Appendix.

### 5.1 COMPARISON TO STATE-OF-THE-ART METHODS

**Three degradations.** We first evaluate DFC-IR under the three-degradation setting, including dehazing, deraining, and denoising. As shown in Table 1, DFC-IR achieves the highest average PSNR of 33.00 dB among all compared methods, outperforming recent DFPIR (Tian et al., 2025) and FIT (He et al., 2026) by 0.12 dB and 0.17 dB, respectively, with competitive SSIM performance.

**Five degradations.** We further extend the three-degradation setting with deblurring and low-light enhancement following prior works (Li et al., 2022; Zhang et al., 2023). As shown in Table 2, DFC-IR achieves the highest average PSNR of 31.03 dB and the best SSIM of 0.922, outperforming recent FIT (He et al., 2026) and DFPIR (Tian et al., 2025) by 0.31 dB and 0.39 dB in PSNR, respectively.

**Composite degradations.** We additionally evaluate DFC-IR on CDD11 (Guo et al., 2024b), which contains eleven single and composite degradation settings. As reported in Table 3, DFC-IR achieves the highest average PSNR of 30.22 dB, outperforming R2R (Wang et al., 2026) and MoCE-IR-S

Table 4: Quantitative comparison on unseen tasks: UIE (UIEB) and desnowing (CSD).

| Method | UIE (UIEB) | | | | Desnowing (CSD) | | | |
|---|---|---|---|---|---|---|---|---|
| | PSNR↑ | SSIM↑ | LPIPS↓ | FID↓ | PSNR↑ | SSIM↑ | LPIPS↓ | FID↓ |
| Restormer | 17.28 | 0.768 | 0.304 | 49.56 | 18.89 | 0.746 | 0.359 | 130.1 |
| PromptIR | 17.51 | 0.774 | 0.289 | 45.77 | 19.12 | 0.750 | 0.356 | 126.2 |
| DFPIR | 17.74 | 0.771 | 0.296 | 45.01 | 19.78 | 0.759 | 0.337 | 114.7 |
| AdaIR | 18.44 | 0.800 | 0.251 | 32.76 | 20.44 | 0.783 | 0.322 | 101.6 |
| VLU-Net | 18.51 | 0.802 | 0.248 | 30.01 | 20.52 | 0.771 | 0.328 | 94.87 |
| MoCE-IR | 18.65 | 0.804 | 0.256 | 33.74 | 20.60 | 0.784 | 0.324 | 92.16 |
| R2R | 18.60 | 0.806 | 0.255 | 30.99 | 20.54 | 0.779 | 0.326 | 96.11 |
| DFC-IR (Ours) | **19.24** | **0.817** | **0.231** | **27.76** | **21.05** | **0.802** | **0.291** | **70.68** |

Table 5: Quantitative comparison on real-world composite degradations.

| Method | Real haze and rain | | | Real blur and noise | | |
|---|---|---|---|---|---|---|
| | MANIQA↑ | CLIPIQA↑ | MUSIQ↑ | MANIQA↑ | CLIPIQA↑ | MUSIQ↑ |
| Restormer | 0.262 | 0.381 | 50.16 | 0.150 | 0.201 | 26.74 |
| PromptIR | 0.274 | 0.399 | 52.37 | 0.139 | 0.194 | 25.99 |
| DFPIR | 0.268 | 0.395 | 54.27 | 0.147 | 0.198 | 26.91 |
| AdaIR | 0.304 | 0.412 | 55.74 | 0.164 | 0.219 | 31.53 |
| VLU-Net | 0.316 | 0.457 | 57.74 | 0.168 | 0.225 | 32.27 |
| MoCE-IR | 0.324 | 0.448 | 58.02 | 0.177 | 0.236 | 33.38 |
| R2R | 0.318 | 0.433 | 56.93 | 0.171 | 0.230 | 32.94 |
| DFC-IR (Ours) | **0.341** | **0.474** | **60.23** | **0.189** | **0.251** | **35.22** |




Figure 5: Visual comparisons on unseen degradation types with desnowing on CSD.

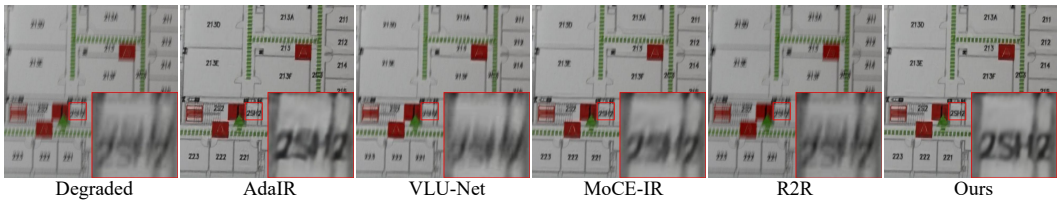



Figure 6: Visual comparisons on real-world compound degradations with blur-and-noise cases.

(Zamfir et al., 2025) by 0.90 dB and 1.17 dB, respectively. These results demonstrate the effectiveness of DFC-guided spectral tokenization in handling complex mixed degradation scenarios.

### 5.2 Generalization under Challenging Scenarios

**Unseen degradation types.** We evaluate the five-degradation model on two unseen tasks without retraining: underwater image enhancement (UIE) on UIEB (Li et al., 2019) and desnowing on CSD (Chen et al., 2021). As shown in Table 4, DFC-IR outperforms existing methods across all metrics. Specifically, DFC-IR improves the PSNR over MoCE-IR (Zamfir et al., 2025) by 0.59 dB on UIEB and 0.45 dB on CSD, demonstrating strong generalization beyond seen degradation types. Fig. 5 illustrates effective desnowing with better detail preservation; see Appendix for more results.

**Real-world compound degradations.** We further evaluate the same five-degradation model on real-world compound degradations, including haze-and-rain images from SPANet (Wang et al., 2019) and blur-and-noise images from Lai (Lai et al., 2016). We evaluate with three no-reference metrics: MANIQA, CLIPIQA, and MUSIQ. As shown in Table 5, DFC-IR achieves the highest scores across all scenarios, improving MUSIQ over MoCE-IR (Zamfir et al., 2025) by 2.21 and 1.84 on haze-and-rain and blur-and-noise images, respectively. Fig. 6 demonstrates effective restoration of real-world compound degradations with fewer artifacts; see Appendix for more results.

### 5.3 Ablation Studies

**Component ablation.** Table 6 evaluates the contributions of FCMS and MMD. As shown in variant (b), FCMS improves the PSNR by 0.21 dB over the baseline (a) through DFC-guided frequency partitioning and adaptive band-wise tokenization. Variant (c) shows that employing MMD alone yields a 0.78 dB improvement, demonstrating the effectiveness of band-specific feature modulation and adaptive aggregation for degradation-aware restoration. Finally, the full model combines FCMS and MMD and achieves the best performance with a 1.02 dB gain over the baseline. These results confirm the complementary roles of DFC-guided tokenization and token-conditioned decoding.

**Token granularity.** As shown in Table 7, partitioning the DFC into 4 band-wise tokens achieves the best average PSNR of 33.00 dB, striking a good balance between spectral granularity and model efficiency. Specifically, it provides sufficient resolution to distinguish local DFC responses across frequency regions, such as high-frequency responses dominated by noise and middle-to-high-

Table 6: Ablation of the proposed components.

| Method | FCMS | MMD | Params. | FLOPs | PSNR↑ | SSIM↑ |
|---|---|---|---|---|---|---|
| (a) | ✗ | ✗ | 26.5M | 120.9G | 31.98 | 0.911 |
| (b) | ✓ | ✗ | 29.3M | 129.1G | 32.19 | 0.912 |
| (c) | ✗ | ✓ | 31.6M | 155.5G | 32.76 | 0.916 |
| Ours | ✓ | ✓ | 33.1M | 163.7G | **33.00** | **0.919** |

Table 7: Ablation on token granularity.

| DFC Tokens | Params. | SOTS | | Rain100L | | $BSD68_{15\text{-}50}$ | | Average | |
|---|---|---|---|---|---|---|---|---|---|
| | | PSNR↑ | SSIM↑ | PSNR↑ | SSIM↑ | PSNR↑ | SSIM↑ | PSNR↑ | SSIM↑ |
| 2 | 30.6M | 31.21 | 0.977 | 38.34 | 0.980 | 31.05 | 0.869 | 32.54 | 0.913 |
| 3 | 31.8M | 31.69 | 0.978 | 38.47 | 0.981 | 31.18 | 0.873 | 32.74 | 0.916 |
| 4 (ours) | 33.1M | **32.28** | **0.982** | **38.74** | **0.985** | **31.33** | **0.876** | **33.00** | **0.919** |
| 5 | 34.4M | 32.25 | 0.981 | 38.71 | 0.984 | 31.32 | 0.874 | 32.98 | 0.917 |

Table 8: Ablation on degradation representations for unseen degradations.

| Representation | UIE (UIEB) | | | | Desnowing (CSD) | | | |
|---|---|---|---|---|---|---|---|---|
| | PSNR↑ | SSIM↑ | LPIPS↓ | FID↓ | PSNR↑ | SSIM↑ | LPIPS↓ | FID↓ |
| Learned Prompt | 18.61 | 0.802 | 0.251 | 31.08 | 20.45 | 0.781 | 0.328 | 93.01 |
| Uniform Tokens | 18.76 | 0.809 | 0.251 | 30.99 | 20.51 | 0.774 | 0.328 | 90.41 |
| DFC-guided Tokens | **19.24** | **0.817** | **0.231** | **27.76** | **21.05** | **0.802** | **0.291** | **70.68** |

Table 9: Ablation on DFC estimation reliability.

| Variant | MS | $\mathcal{L}_f$ | NMAE↓ | $W_1$ ↓ | PSNR↑ |
|---|---|---|---|---|---|
| (a) | ✗ | ✗ | 0.3562 | 0.1066 | 32.22 |
| (b) | ✗ | ✓ | 0.2776 | 0.0835 | 32.38 |
| (c) | ✓ | ✗ | 0.0543 | 0.0196 | 32.81 |
| (d) | ✓ | ✓ | **0.0163** | **0.0047** | **33.00** |

frequency responses caused by rain streaks, while avoiding redundant tokens and unnecessary modulation branches. In comparison, using 2 or 3 tokens results in overly broad frequency segments that merge heterogeneous degradation responses and reduce token discriminability. Increasing the number to 5 raises the parameter count to 34.4M but slightly degrades performance, indicating that over-partitioning increases parameter redundancy and complicates token aggregation.

**Effect of degradation representation.** We further investigate the influence of different degradation representations on generalization to unseen degradation types. Specifically, the model trained on five degradation types is directly evaluated on UIEB and CSD without adaptation. Table 8 compares three strategies: Learned Prompt, Uniform Tokens, and DFC-guided Tokens. Both Learned Prompt and Uniform Tokens consistently underperform DFC-guided Tokens on the two unseen datasets, indicating that implicit learned degradation cues or fixed frequency partitioning alone provide limited adaptability to degradation patterns beyond the training distribution. In contrast, DFC-guided Tokens achieve the best results across all metrics, improving PSNR over Uniform Tokens by 0.48 dB on UIEB and 0.54 dB on CSD, together with notable improvements in LPIPS and FID. The visual comparison in Fig. 7 further shows that DFC-guided Tokens more effectively correct unseen underwater color shifts and recover more faithful details. These results demonstrate that adapting the degradation representation according to input-specific DFC responses provides more effective degradation guidance and improves generalization to unseen degradations.




Figure 7: Restoration results with different degradation representations on UIE.

**Reliability of DFC estimation.** Since DFC-IR relies on DFCs estimated from intermediate restorations to guide subsequent restoration stages, we further evaluate the accuracy of these DFCs and their effect on restoration performance. We use normalized mean absolute error (NMAE) and Wasserstein-1 distance $W_1$ to measure the discrepancy between the estimated and reference DFCs. As shown in Table 9, without multi-scale supervision, the intermediate restorations are not explicitly constrained, resulting in relatively large DFC estimation errors. Introducing multi-scale supervision substantially reduces NMAE from 0.3562 to 0.0543 and $W_1$ from 0.1066 to 0.0196, while the frequency loss further improves the estimation accuracy. With both objectives, the estimated DFC achieves an NMAE of 0.0163 and a $W_1$ of 0.0047, demonstrating that DFC can be reliably estimated from intermediate restorations. Meanwhile, increasingly accurate DFC estimation is consistently accompanied by improved restoration performance, with PSNR increasing from 32.22 to 33.00 dB. This trend supports the effectiveness of reliable DFC guidance.

## 6 Conclusion

In this paper, we introduced DFC, an explicit, quantifiable, and interpretable spectral representation for characterizing degradation responses in the frequency domain. To translate DFC into restoration guidance, we further developed DFC-IR, which decomposes scale-specific DFCs into band-wise spectral tokens and leverages these tokens as degradation-aware priors for coarse-to-fine restoration. Extensive experiments on standard all-in-one benchmarks, composite degradations, unseen degradations, and real-world scenarios demonstrate that the proposed framework achieves state-of-the-art restoration performance and strong generalization under complex degradation conditions.

## AI USE STATEMENT

In this work, we used generative AI tools for language polishing and proofreading during manuscript preparation. We have not used generative AI tools for developing the methodology, conducting experiments, analyzing experimental results, or generating scientific claims. Other uses of generative AI tools beyond writing assistance are not applicable to this work. All AI-assisted suggestions were carefully reviewed and revised by the authors before inclusion in the manuscript.

## ETHICS STATEMENT

This work studies fundamental methods for all-in-one image restoration, aiming to improve the robustness and generalization ability of restoration models under diverse degradation conditions. Our research focuses on algorithmic improvements for degradation representation and restoration guidance, and does not involve the collection or annotation of new data. All datasets used in our experiments are publicly available benchmark datasets and are used according to their corresponding licenses. The proposed method is designed to improve image restoration performance and does not raise foreseeable ethical concerns beyond the general considerations of responsible AI research.

## REPRODUCIBILITY STATEMENT

We aim to ensure reproducibility and transparency of our results. DFC-IR is implemented in PyTorch, and detailed descriptions of the proposed representation, model architecture, training settings, datasets, evaluation metrics, and experimental protocols are provided in the main paper and supplementary material. Additional analyses and experimental results are included in the appendix to facilitate understanding and reproduction of our method. Upon acceptance, we will release the source code, pretrained models, and instructions for reproducing the reported results.

## REFERENCES

Codruta O Ancuti, Cosmin Ancuti, and Radu Timofte. Nh-haze: An image dehazing benchmark with non-homogeneous hazy and haze-free images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pp. 444–445, 2020.

Pablo Arbelaez, Michael Maire, Charless Fowlkes, and Jitendra Malik. Contour detection and hierarchical image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 33(5):898–916, 2010.

Chenghao Chen and Hao Li. Robust representation learning with feedback for single image deraining. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 7742–7751, 2021.

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 *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4196–4205, 2021.

Sung-Jin Cho, Seo-Won Ji, Jun-Pyo Hong, Seung-Won Jung, and Sung-Jea Ko. Rethinking coarse-to-fine approach in single image deblurring. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4641–4650, 2021.

Marcos V Conde, Gregor Geigle, and Radu Timofte. Instructir: High-quality image restoration following human instructions. In *European Conference on Computer Vision*, pp. 1–21. Springer, 2024.

Yuning Cui, Wenqi Ren, Xiaochun Cao, and Alois Knoll. Image restoration via frequency selection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(2):1093–1108, 2023a.

Yuning Cui, Yi Tao, Zhenshan Bing, Wenqi Ren, Xinwei Gao, Xiaochun Cao, Kai Huang, and Alois Knoll. Selective frequency network for image restoration. In *The eleventh international conference on learning representations*, 2023b.

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 *13th International Conference on Learning Representations, ICLR 2025*, pp. 57335–57356. International Conference on Learning Representations, ICLR, 2025.

Hang Guo, Jinmin Li, Tao Dai, Zhihao Ouyang, Xudong Ren, and Shu-Tao Xia. Mambair: A simple baseline for image restoration with state-space model. In *European conference on computer vision*, pp. 222–241. Springer, 2024a.

Yu Guo, Yuan Gao, Yuxu Lu, Huilin Zhu, Ryan Wen Liu, and Shengfeng He. Onerestore: A universal restoration framework for composite degradation. In *European conference on computer vision*, pp. 255–272. Springer, 2024b.

Zihao He, Yunfeng Wu, Xinchao Wang, and Songhua Liu. Bend the basics: Degradation-aware deformable tokenization for all-in-one image restoration. *arXiv preprint arXiv:2608.06832*, 2026.

Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.

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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8346–8355, 2020.

Liming Jiang, Bo Dai, Wayne Wu, and Chen Change Loy. Focal frequency loss for image reconstruction and synthesis. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 13919–13929, 2021.

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

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Xiangtao Kong, Chao Dong, and Lei Zhang. Towards effective multiple-in-one image restoration: A sequential and prompt learning strategy. *arXiv preprint arXiv:2401.03379*, 2024.

Orest Kupyn, Tetiana Martyniuk, Junru Wu, and Zhangyang Wang. Deblurgan-v2: Deblurring (orders-of-magnitude) faster and better. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 8878–8887, 2019.

Wei-Sheng Lai, Jia-Bin Huang, Zhe Hu, Narendra Ahuja, and Ming-Hsuan Yang. A comparative study for single image blind deblurring. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1701–1709, 2016.

Boyi Li, Wenqi Ren, Dengpan Fu, Dacheng Tao, Dan Feng, Wenjun Zeng, and Zhangyang Wang. Benchmarking single-image dehazing and beyond. *IEEE transactions on image processing*, 28 (1):492–505, 2018.

Boyun Li, Xiao Liu, Peng Hu, Zhongqin Wu, Jiancheng Lv, and Xi Peng. All-in-one image restoration for unknown corruption. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 17452–17462, 2022.

Chongyi Li, Chunle Guo, Wenqi Ren, Runmin Cong, Junhui Hou, Sam Kwong, and Dacheng Tao. An underwater image enhancement benchmark dataset and beyond. *IEEE transactions on image processing*, 29:4376–4389, 2019.

Xin Lin, Chao Ren, Xiao Liu, Jie Huang, and Yinjie Lei. Unsupervised image denoising in real-world scenarios via self-collaboration parallel generative adversarial branches. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 12642–12652, 2023.

Kede Ma, Zhengfang Duanmu, Qingbo Wu, Zhou Wang, Hongwei Yong, Hongliang Li, and Lei Zhang. Waterloo exploration database: New challenges for image quality assessment models. *IEEE Transactions on Image Processing*, 26(2):1004–1016, 2016.

David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings eighth IEEE international conference on computer vision. ICCV 2001*, volume 2, pp. 416–423. IEEE, 2001.

Chong Mou, Qian Wang, and Jian Zhang. Deep generalized unfolding networks for image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 17399–17410, 2022.

Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3883–3891, 2017.

Ozan Özdenizci and Robert Legenstein. Restoring vision in adverse weather conditions with patch-based denoising diffusion models. *IEEE transactions on pattern analysis and machine intelligence*, 45(8):10346–10357, 2023.

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 *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5815–5824. IEEE, 2023.

Vaishnav Potlapalli, Syed Waqas Zamir, Salman H Khan, and Fahad Shahbaz Khan. Promptir: Prompting for all-in-one image restoration. *Advances in Neural Information Processing Systems*, 36:71275–71293, 2023.

Xiangpeng Tian, Xiangyu Liao, Xiao Liu, Meng Li, and Chao Ren. Degradation-aware feature perturbation for all-in-one image restoration. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 28165–28175, 2025.

Chenxu Wang, Kai Zhang, and Jian Yang. Retrieve-to-restore: Efficient all-in-one image restoration with a retrieval-based degradation bank. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1277–1287, 2026.

Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring clip for assessing the look and feel of images. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pp. 2555–2563, 2023a.

Simin Wang, Xiaoguang Li, Hui Zhang, and Li Zhuo. A novel blind motion blur restoration algorithm for text images. *Chinese Journal of Electronics*, 29(1):161–167, 2020.

Tianyu Wang, Xin Yang, Ke Xu, Shaozhe Chen, Qiang Zhang, and Rynson WH Lau. Spatial attentive single-image deraining with a high quality real rain dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 12270–12279, 2019.

Yinglong Wang, Chao Ma, and Jianzhuang Liu. Smartassign: Learning a smart knowledge assignment strategy for deraining and desnowing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3677–3686, 2023b.

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

Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. *arXiv preprint arXiv:1808.04560*, 2018.

Kai Xu, Minghai Qin, Fei Sun, Yuhao Wang, Yen-Kuang Chen, and Fengbo Ren. Learning in the frequency domain. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 1740–1749, 2020.

Sidi Yang, Tianhe Wu, Shuwei Shi, Shanshan Lao, Yuan Gong, Mingdeng Cao, Jiahao Wang, and Yujiu Yang. Maniqa: Multi-dimension attention network for no-reference image quality assessment. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 1190–1199. IEEE, 2022.

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

Wenhan Yang, Wenjing Wang, Haofeng Huang, Shiqi Wang, and Jiaying Liu. Sparse gradient regularized deep retinex network for robust low-light image enhancement. *IEEE Transactions on Image Processing*, 30:2072–2086, 2021.

Zhiwen Yang, Haowei Chen, Ziniu Qian, Yang Yi, Hui Zhang, Dan Zhao, Bingzheng Wei, and Yan Xu. All-in-one medical image restoration via task-adaptive routing. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 67–77. Springer, 2024.

Jaeyoung Yoo, Sang-ho Lee, and Nojun Kwak. Image restoration by estimating frequency distribution of local patches. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 6684–6692, 2018.

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 *Proceedings of the Computer Vision and Pattern Recognition Conference*, 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 14821–14831, 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 5728–5739, 2022.

Haijin Zeng, Xiangming Wang, Yongyong Chen, Jingyong Su, and Jie Liu. Vision-language gradient descent-driven all-in-one deep unfolding networks. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 7524–7533, 2025.

Liang Zeng and Yongzhen Dai. Single image dehazing based on combining dark channel prior and scene radiance constraint. *Chinese Journal of Electronics*, 25(6):1114–1120, 2016.

Jinghao Zhang, Jie Huang, Mingde Yao, Zizheng Yang, Hu Yu, Man Zhou, and Feng Zhao. Ingredient-oriented multi-degradation learning for image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 5825–5835, 2023.

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

Tao Zhang, Ying Fu, and Jun Zhang. Deep guided attention network for joint denoising and demosaicing in real image. *Chinese Journal of Electronics*, 33(1):303–312, 2024.

Xu Zhang, Jiaqi Ma, Guoli Wang, Qian Zhang, Huan Zhang, and Lefei Zhang. Perceive-ir: Learning to perceive degradation better for all-in-one image restoration. *IEEE Transactions on Image Processing*, 2025.

Kaiwen Zhu, Jinjin Gu, Zhiyuan You, Yu Qiao, and Chao Dong. An intelligent agentic system for complex image restoration problems. In *International Conference on Learning Representations*, volume 2025, pp. 57985–58013, 2025.

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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 21747–21758, 2023.

# APPENDIX

This appendix provides additional analyses and experimental results to further validate the effectiveness of DFC-IR. Sec. A presents additional details of DFC, including its construction process and Gaussian mask design. Sec. B provides quantitative analyses of DFC properties and compares DFC with alternative spectral representations in terms of content–degradation separability. Sec. C describes the complete experimental protocols, including dataset preparation, implementation details, and evaluation metrics. Sec. D reports additional quantitative results under challenging scenarios, including unseen degradation levels and real-world individual degradations. Sec. E presents the computational efficiency analysis, and Sec. F provides extensive visual comparisons under various restoration settings.

## A ADDITIONAL DETAILS OF DFC

### A.1 DFC CONSTRUCTION

We provide a detailed illustration of the DFC construction process in Fig. 8. Given a degraded image $\boldsymbol{y}$ and its clean counterpart $\boldsymbol{x}$, we first obtain the residual image $\boldsymbol{r} = \boldsymbol{y} - \boldsymbol{x}$ to isolate the degradation-induced component. The residual and degraded images are then transformed into the frequency domain to obtain their spectral energy maps $|\mathcal{F}(\boldsymbol{r})|^2$ and $|\mathcal{F}(\boldsymbol{y})|^2$. To characterize degradation responses at different frequencies, the spectral domain is decomposed using a set of $B$ ring-shaped Gaussian masks $\{\mathcal{M}_b\}_{b=1}^{B}$. Each mask extracts the corresponding band-wise energy components from both spectral energy maps. The degradation response within each band is subsequently quantified by the residual-to-degraded energy ratio defined in Eq. 3 of the main paper. Collecting these responses along the frequency axis yields the final DFC $D(f)$. This process converts spatially entangled degradation components into a compact one-dimensional representation that explicitly describes their frequency-dependent responses.

### A.2 GAUSSIAN MASK DESIGN AND VISUALIZATION

As described in the main paper, the $b$-th ring-shaped Gaussian mask $\mathcal{M}_b$ is parameterized by its central frequency $\mu_b$ and bandwidth $\sigma_b$, where $\mu_b$ determines the radial frequency location of the maximum response and $\sigma_b$ controls the bandwidth. To make the frequency partition independent of image resolution, we normalize the radial frequency coordinate by the maximum distance from the frequency center to the image boundary. We use $B = 25$ predefined Gaussian masks for DFC construction in this normalized frequency domain. The central frequencies $\{\mu_b\}_{b=1}^{B}$ are linearly distributed in $[0, 1]$, while the bandwidth parameters $\{\sigma_b\}_{b=1}^{B}$ are linearly spaced in $[0.05, 0.3]$. All Gaussian masks are constructed in this normalized frequency coordinate system.

Unlike rigid frequency partitioning, Gaussian masks provide smooth transitions between adjacent bands through soft spectral assignment, reducing sensitivity to manually defined frequency boundaries. The progressively increasing bandwidth further accounts for the non-uniform spectral distribution of degraded images. Specifically, narrower masks in the energy-dominant low-frequency regions enable finer discrimination of local degradation responses, whereas broader masks toward higher frequencies provide more stable energy estimation and avoid excessive spectral fragmentation. Together, these masks provide smooth and comprehensive coverage of the frequency domain while retaining sufficient resolution for band-wise degradation quantification. The complete set of 25 Gaussian masks is visualized in Fig. 9.

## B QUANTITATIVE ANALYSIS OF DFC PROPERTIES

To clarify the different forms of DFC used in the following analysis, we derive two representations from the original DFC $D(f)$.

For evaluating content robustness and degradation sensitivity, we use the normalized DFC, defined as:

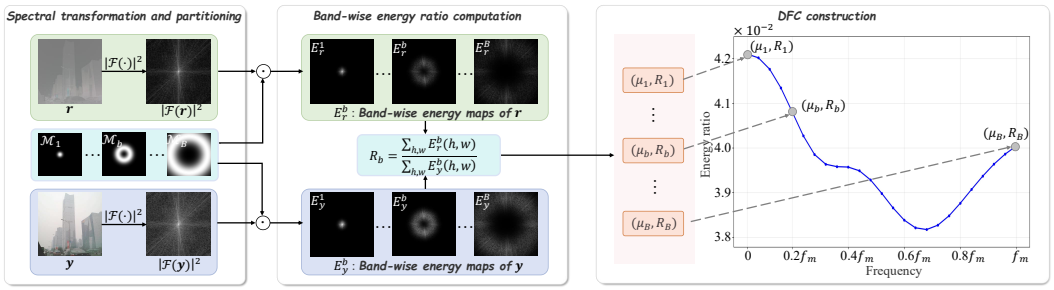



Figure 8: Overview of DFC construction. The DFC is obtained by transforming the residual and degraded images into the frequency domain and computing band-wise energy ratios with Gaussian masks to form the final curve.

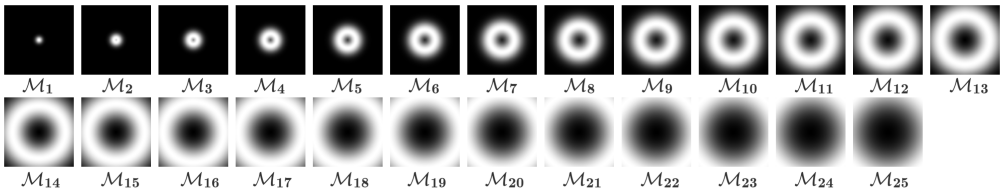


Figure 9: Visualization of the 25 predefined ring-shaped Gaussian masks used for DFC construction, arranged from low to high frequencies.

$$\tilde{D}(f_b) = \frac{R_b}{\sum_{j=1}^{B} R_j}, \tag{13}$$

where $\tilde{D}(f)$ denotes the normalized DFC profile. The normalization removes the overall magnitude variation while preserving the relative spectral distribution across frequency bands.

In addition, to measure the overall degradation response strength, we use the average magnitude of DFC, defined as:

$$A(D) = \frac{1}{B}\sum_{b=1}^{B} R_b, \tag{14}$$

where $A(D)$ denotes the average response magnitude of $D(f)$.

### B.1 Content Robustness

We first quantitatively evaluate the content robustness of DFC. According to the definition above, the normalized DFC profiles $\tilde{D}$ are used for this analysis.

To measure the discrepancy between two normalized DFC profiles $\tilde{D}_i$ and $\tilde{D}_j$, we employ the Wasserstein-1 distance:

$$W_1(\tilde{D}_i, \tilde{D}_j) = \sum_{b=1}^{B-1} \left| \sum_{k=1}^{b} \left( \tilde{D}_{i,k} - \tilde{D}_{j,k} \right) \right| (\mu_{b+1} - \mu_b), \tag{15}$$

where $\mu_b$ is the central frequency of the $b$-th Gaussian frequency band. A smaller $W_1$ indicates more similar DFC shapes, whereas a larger value indicates a greater difference between their spectral response distributions.

Table 10: Within-degradation Wasserstein-1 distances for evaluating DFC content robustness.

| Metric | Haze | Rain | Noise | Average |
|---|---|---|---|---|
| $W_1 \downarrow$ | 0.0192 | 0.0411 | 0.0259 | **0.0287** |

Table 11: Between-degradation Wasserstein-1 distances for evaluating DFC degradation sensitivity.

| Metric | Haze–Rain | Haze–Noise | Rain–Noise | Average |
|---|---|---|---|---|
| $W_1 \uparrow$ | 0.1596 | 0.1816 | 0.1821 | **0.1744** |

To quantify the influence of image content, we compare normalized DFC profiles from different image contents while keeping the degradation type fixed. For degradation type $c$, the within-degradation distance is defined as the mean pairwise Wasserstein-1 distance among all $N$ images:

$$d^c_{\text{within}} = \frac{2}{N(N-1)} \sum_{i<j} W_1(\tilde{D}^c_i, \tilde{D}^c_j). \tag{16}$$

The overall within-degradation distance is obtained by averaging the results over haze, rain, and noise:

$$d_{\text{within}} = \frac{1}{3} \left( d^H_{\text{within}} + d^R_{\text{within}} + d^N_{\text{within}} \right). \tag{17}$$

As shown in Table 10, the within-degradation distances are $0.0192$, $0.0411$, and $0.0259$ for haze, rain, and noise, respectively, with an average of only $0.0287$. These consistently small distances indicate that normalized DFC profiles remain stable across different image contents under the same degradation, quantitatively supporting the content robustness observed in the main paper.

### B.2 Degradation Sensitivity

Using the same Wasserstein-1 distance, we further evaluate the sensitivity of DFC to different degradation types while controlling image content. Specifically, we compare normalized DFC profiles generated from the same clean image under different degradations. For two degradation types $c$ and $c'$, the between-degradation distance is defined as:

$$d^{c,c'}_{\text{between}} = \frac{1}{N} \sum_{i=1}^{N} W_1(\tilde{D}^c_i, \tilde{D}^{c'}_i). \tag{18}$$

The overall between-degradation distance is obtained by averaging the three degradation pairs:

$$d_{\text{between}} = \frac{1}{3} \left( d^{H,R}_{\text{between}} + d^{H,N}_{\text{between}} + d^{R,N}_{\text{between}} \right). \tag{19}$$

As reported in Table 11, the between-degradation distances are $0.1596$, $0.1816$, and $0.1821$ for haze–rain, haze–noise, and rain–noise, respectively, yielding an average of $0.1744$. To jointly characterize degradation discrimination and content robustness, we further define the between-to-within separation ratio as

$$\mathcal{S}_{\text{sep}} = \frac{d_{\text{between}}}{d_{\text{within}}}. \tag{20}$$

A larger $\mathcal{S}_{\text{sep}}$ indicates that degradation-related variation dominates content-induced variation. For DFC, the average between-degradation distance is substantially larger than the average within-degradation distance of $0.0287$, resulting in $\mathcal{S}_{\text{sep}} = 6.07$. These results show that DFC maintains clear degradation-specific response patterns while remaining relatively stable across image contents.

### B.3 Severity-Aware Response

We further quantitatively evaluate whether the average magnitude of DFC $A(D)$ consistently reflects degradation severity. For each image, we generate degraded samples at six ordered severity levels

Table 12: Spearman rank correlation between degradation severity and average DFC amplitude.

| Degradation | SRC↑ |
|---|---|
| Haze | 1.0000 |
| Rain | 0.8480 |
| Noise | 1.0000 |
| Average | **0.9493** |

Table 13: Comparison of content–degradation separability across different spectral representations. Lower within-degradation distance and higher between-degradation distance indicate stronger content robustness and degradation sensitivity, respectively. The separation ratio measures degradation separability relative to content variation.

| Representation | Within Dist. ↓ | Between Dist. ↑ | Separation Ratio ↑ |
|---|---|---|---|
| Degraded Spectrum | 0.0380 | 0.0659 | 1.73 |
| Residual Spectrum | 0.0309 | 0.1083 | 3.50 |
| DFC (Ours) | **0.0287** | **0.1744** | **6.07** |

and measure the monotonic relationship between degradation severity and the corresponding DFC magnitudes using Spearman rank correlation (SRC).

Let $\boldsymbol{s} = [1, \ldots, 6]$ denote the ordered severity levels and $\boldsymbol{A}_i = [A_i^1, \ldots, A_i^6]$ denote the corresponding average DFC magnitudes of the $i$-th image, where $A_i^k$ represents the value of $A(D)$ at the $k$-th severity level. The Spearman rank correlation for each image is computed as

$$\rho_i = \mathrm{corr}\left(\mathrm{rank}(\boldsymbol{s}), \mathrm{rank}(\boldsymbol{A}_i)\right), \tag{21}$$

and the reported SRC is averaged over all $N$ test images:

$$\mathrm{SRC} = \frac{1}{N} \sum_{i=1}^{N} \rho_i. \tag{22}$$

A larger SRC indicates that DFC magnitude follows the progression of degradation severity more consistently, with $\mathrm{SRC} = 1$ corresponding to a perfectly monotonic relationship.

As shown in Table 12, haze and noise both achieve an SRC of $1.0000$, indicating that their average DFC magnitudes vary perfectly monotonically with degradation severity across the evaluated levels. Rain achieves an SRC of $0.8480$, showing a strong monotonic relationship despite larger variation. Averaged over the three degradation types, DFC obtains an SRC of $0.9493$. These results quantitatively demonstrate that the average magnitude of DFC consistently reflects changes in degradation severity.

### B.4 Comparison with Alternative Degradation Representations

The above analyses establish the intrinsic properties of DFC. We further investigate whether its content robustness and degradation sensitivity specifically benefit from the residual-to-degraded ratio formulation, rather than merely from using frequency-domain information. To this end, we compare DFC with two direct spectral representations: the degraded-image spectrum and the residual spectrum.

Following the controlled comparison protocol in Secs. B.1 and B.2, we measure content-induced variation using the within-degradation distance $d_{\mathrm{within}}$ and degradation discrimination using the between-degradation distance $d_{\mathrm{between}}$. For all three representations, the corresponding band-wise profiles are normalized to unit sum and evaluated using the Wasserstein-1 distance defined above. We further report the separation ratio $\mathcal{S}_{\mathrm{sep}}$ defined in Sec. B.2, where a larger value indicates stronger degradation separability relative to content variation.

As shown in Table 13, directly using the degraded-image spectrum yields a relatively large within-degradation distance of $0.0380$ and a separation ratio of only $1.73$, indicating substantial sensitivity

to image content. Using the residual spectrum reduces the within-degradation distance to $0.0309$ and increases the between-degradation distance to $0.1083$, suggesting that isolating the residual already suppresses part of the content-dependent spectral variation. Nevertheless, DFC further reduces the within-degradation distance to $0.0287$ while substantially increasing the between-degradation distance to $0.1744$, resulting in a separation ratio of $6.07$. This is markedly higher than the corresponding ratios of the degraded spectrum ($1.73$) and residual spectrum ($3.50$).

These results show that the advantage of DFC does not arise merely from using frequency-domain information or isolating the residual component. Instead, the residual-to-degraded ratio formulation further suppresses content-dependent spectral variations while enhancing degradation-specific differences, providing a more stable and discriminative degradation representation.

# C Experimental Protocols

## C.1 Dataset Preparation

**Three-degradation and five-degradation settings.** Following established all-in-one image restoration protocols (Li et al., 2022; Potlapalli et al., 2023), we construct unified restoration models for three-degradation and five-degradation settings. For image denoising, we combine BSD400 (Arbelaez et al., 2010) and WED (Ma et al., 2016) for training and synthesize noisy images with Gaussian noise levels $\sigma \in \{15, 25, 50\}$, while BSD68 (Martin et al., 2001) is used for evaluation. For deraining and dehazing, we adopt Rain100L (Yang et al., 2017) and SOTS (Li et al., 2018), respectively. The three-degradation setting jointly handles denoising, deraining, and dehazing with a single model. The five-degradation setting further incorporates deblurring and low-light enhancement using GoPro (Nah et al., 2017) and LOL-v1 (Wei et al., 2018), respectively. Following the standard five-degradation protocol, denoising performance in this setting is evaluated at $\sigma = 25$.

**Composite-degradation setting.** For composite degradation restoration, we conduct experiments on CDD11 (Guo et al., 2024b), which contains eleven degradation conditions with different levels of compositional complexity. Specifically, it includes four single degradations, i.e., low-light (L), haze (H), rain (R), and snow (S); five double degradations, i.e., L+H, L+R, L+S, H+R, and H+S; and two triple degradations, i.e., L+H+R and L+H+S. A unified model is trained to handle all eleven degradation conditions.

**Unseen-degradation settings.** We evaluate generalization from both unseen degradation types and unseen degradation levels. For unseen degradation types, the five-degradation model is directly applied without retraining to underwater image enhancement on UIEB (Li et al., 2019) and desnowing on CSD (Chen et al., 2021), neither of which is included during training. For unseen degradation levels, we further evaluate the robustness under degradation severities beyond the training distribution. Specifically, the five-degradation model, whose deraining task is trained on Rain100L, is directly evaluated on Rain100H (Yang et al., 2017). We also evaluate the three-degradation model trained with $\sigma \in \{15, 25, 50\}$ on BSD68 corrupted by an unseen noise level of $\sigma = 75$.

**Real-world settings.** For real-world individual degradations, we directly evaluate the five-degradation model on NH-HAZE (Ancuti et al., 2020) for dehazing, SPANet (Wang et al., 2019) for deraining, and LOL-v2-real (Yang et al., 2021) for low-light enhancement. We further evaluate the same model on real-world compound degradations using 45 mixed-degradation images collected from SPANet and Lai (Lai et al., 2016), covering haze-and-rain and blur-and-noise cases, respectively.

## C.2 Implementation Details

DFC-IR is implemented as an end-to-end trainable multi-scale encoder-decoder framework with four encoder levels and three decoding stages. The encoder progressively stacks $[4, 6, 6, 8]$ transformer blocks (Zamir et al., 2022; Potlapalli et al., 2023) from high to low resolution, while the three decoder stages contain $[2, 4, 4]$ FCMS and MMD modules from shallow to deep stages. The lightweight network $\Phi(\cdot)$ in FCMS consists of three MLP layers, and the number of sampled candidate parameters is set to $K = 5$. We use $B = 25$ predefined Gaussian masks for DFC construction and $N = 4$ band-wise DFC tokens in FCMS, with corresponding band-wise modulation branches in MMD. Here, $B$ controls the frequency resolution of the DFC representation, while $N$ specifies

the number of adaptive spectral tokens generated from the DFC by FCMS for subsequent token-conditioned restoration.

All models are trained for 100 epochs with a batch size of 32 using Adam (Kingma & Ba, 2014) with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. The initial learning rate is set to $2 \times 10^{-4}$ and gradually decayed using a cosine schedule. During training, we randomly crop $128 \times 128$ patches and apply random horizontal and vertical flips for data augmentation. Following the training objective defined in the main paper, the model is optimized using the multi-scale spatial reconstruction loss $\mathcal{L}_s$ and frequency-domain consistency loss $\mathcal{L}_f$, with $\lambda_s = 1$ and $\lambda_f = 0.1$. All experiments are conducted on 4 NVIDIA A800 GPUs.

### C.3 Evaluation Metrics

We adopt different evaluation metrics according to the characteristics of each experimental setting. For the standard three-degradation, five-degradation, and composite-degradation benchmarks, we report PSNR and SSIM (Wang et al., 2004) to evaluate restoration fidelity. For unseen degradation types, unseen degradation levels, and real-world individual degradations with reference images, we additionally report LPIPS (Zhang et al., 2018) and FID (Heusel et al., 2017) to assess perceptual quality. For real-world compound degradations without reference images, we employ three no-reference image quality metrics: MANIQA (Yang et al., 2022), CLIPIQA (Wang et al., 2023a), and MUSIQ (Ke et al., 2021).

## D Additional Quantitative Results under Challenging Scenarios

### D.1 Unseen Degradation Levels

To further evaluate the generalization capability of DFC-IR beyond the degradation conditions observed during training, we consider two unseen degradation level settings: heavy rain and severe noise. Specifically, the five-degradation model is directly evaluated on Rain100H (Yang et al., 2017), where the deraining branch is trained only on Rain100L. We further evaluate the three-degradation model, trained with Gaussian noise levels $\sigma \in \{15, 25, 50\}$, on BSD68 with an unseen noise level of $\sigma = 75$.

As shown in Table 14, DFC-IR consistently achieves superior performance under both unseen degradation levels. On Rain100H, DFC-IR obtains a PSNR improvement of 0.45 dB over the strongest competing method, while under severe noise with $\sigma = 75$, it further improves PSNR by 1.06 dB compared with the best baseline. These results demonstrate that DFC provides reliable degradation representation for generalizing to unseen degradation severities.

### D.2 Real-world Individual Degradations

We further evaluate the transferability of DFC-IR on real-world individual degradations using the five-degradation model without additional adaptation. Specifically, we conduct experiments on NH-HAZE (Ancuti et al., 2020) for dehazing, SPANet (Wang et al., 2019) for deraining, and LOL-v2-real (Yang et al., 2021) for low-light enhancement. PSNR, SSIM, LPIPS, and FID are adopted for quantitative evaluation.

As reported in Table 15, DFC-IR consistently achieves competitive performance across all three real-world scenarios. Compared with existing all-in-one restoration methods, DFC-IR obtains the best overall results on NH-HAZE, SPANet, and LOL-v2-real, with consistent improvements in both fidelity and perceptual metrics. These results indicate that DFC provides effective degradation guidance under real-world degradation conditions.

## E Efficiency Analysis

We further evaluate the computational efficiency of DFC-IR in terms of parameters, FLOPs, GPU memory consumption, and inference latency. As reported in Table 16, DFC-IR achieves the highest

Table 14: Quantitative comparison on unseen degradation levels, including heavy rain and severe noise.

| Method | Rain100H | | | | $BSD68_{\sigma=75}$ | | | |
|---|---|---|---|---|---|---|---|---|
| | PSNR↑ | SSIM↑ | LPIPS↓ | FID↓ | PSNR↑ | SSIM↑ | LPIPS↓ | FID↓ |
| Restormer | 14.57 | 0.469 | 0.481 | 247.6 | 13.79 | 0.359 | 0.472 | 201.6 |
| PromptIR | 14.69 | 0.478 | 0.471 | 243.4 | 18.46 | 0.401 | 0.406 | 170.1 |
| DFPIR | 15.56 | 0.494 | 0.403 | 197.5 | 18.67 | 0.416 | 0.392 | 164.1 |
| AdaIR | 15.88 | 0.515 | 0.386 | 179.8 | 21.57 | 0.498 | 0.341 | 129.5 |
| VLU-Net | 15.97 | 0.524 | 0.374 | 164.2 | 21.81 | 0.506 | 0.338 | 120.6 |
| MoCE-IR | 16.04 | 0.531 | 0.397 | 184.6 | 20.18 | 0.467 | 0.376 | 156.1 |
| R2R | 15.95 | 0.522 | 0.393 | 177.7 | 21.14 | 0.486 | 0.349 | 137.2 |
| DFC-IR (Ours) | **16.49** | **0.541** | **0.352** | **141.7** | **22.87** | **0.516** | **0.315** | **112.9** |

Table 15: Quantitative comparison of five-degradation models on real-world individual degradations.

| Method | NH-HAZE | | | | SPANet | | | | LOL-v2-real | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | PSNR↑ | SSIM↑ | LPIPS↓ | FID↓ | PSNR↑ | SSIM↑ | LPIPS↓ | FID↓ | PSNR↑ | SSIM↑ | LPIPS↓ | FID↓ |
| Restormer | 16.04 | 0.501 | 0.399 | 218.8 | 34.35 | 0.912 | 0.032 | 43.51 | 27.11 | 0.877 | 0.112 | 85.67 |
| PromptIR | 16.45 | 0.507 | 0.394 | 214.1 | 35.32 | 0.936 | 0.027 | 33.20 | 27.61 | 0.868 | 0.109 | 81.43 |
| DFPIR | 17.04 | 0.514 | 0.384 | 224.5 | 34.96 | 0.911 | 0.037 | 49.26 | 27.74 | 0.871 | 0.104 | 76.29 |
| AdaIR | 17.96 | 0.541 | 0.356 | 186.5 | 36.99 | 0.951 | 0.022 | 23.02 | 28.09 | 0.899 | 0.095 | 70.11 |
| VLU-Net | 18.09 | 0.536 | 0.361 | 191.2 | 37.31 | 0.962 | **0.017** | 19.96 | 27.25 | 0.881 | 0.110 | 88.49 |
| MoCE-IR | 18.17 | 0.545 | 0.350 | 180.1 | 37.54 | **0.971** | **0.017** | 24.18 | 28.18 | 0.902 | 0.106 | 84.57 |
| R2R | 17.99 | 0.548 | 0.470 | 177.2 | 37.09 | 0.958 | **0.017** | 22.31 | 27.99 | 0.894 | 0.097 | 77.54 |
| DFC-IR (Ours) | **18.44** | **0.551** | **0.343** | **168.6** | **38.01** | 0.969 | **0.017** | **19.12** | **29.27** | **0.929** | **0.062** | **38.95** |

Table 16: Computational efficiency under the five-degradation setting. Evaluated on one NVIDIA A800 GPU with an input size of $256 \times 256$.

| Method | Params. | FLOPs | GPU Mem. | Latency | PSNR↑ | SSIM↑ |
|---|---|---|---|---|---|---|
| Restormer | 26.1M | 141.0G | 667.5MiB | 85.77ms | 27.60 | 0.881 |
| PromptIR | 35.6M | 158.1G | 711.5MiB | 81.52ms | 29.15 | 0.904 |
| VLU-Net | 35.4M | 168.1G | 762.8MiB | 127.7ms | 30.11 | 0.905 |
| AdaIR | 28.8M | 147.2G | 678.3MiB | 92.25ms | 30.20 | 0.910 |
| MoCE-IR | 25.4M | 86.25G | 314.9MiB | 67.21ms | 30.58 | 0.919 |
| DFPIR | 31.1M | 151.1G | 709.7MiB | 139.8ms | 30.64 | 0.913 |
| DFC-IR (Ours) | 33.1M | 163.7G | 413.1MiB | 122.6ms | **31.03** | **0.922** |

restoration performance while maintaining competitive computational efficiency. Compared with DFPIR, DFC-IR improves PSNR by 0.39 dB, while reducing GPU memory consumption and inference latency by 41.8% and 12.3%, respectively, with only moderate increases in parameters and FLOPs. Compared with VLU-Net, DFC-IR achieves a 0.92 dB PSNR improvement while using fewer parameters and FLOPs, together with substantially lower GPU memory consumption and inference latency.

# F Additional Visual Results

## F.1 Challenging Scenarios

**Unseen degradation types.** We first provide additional visual comparisons on unseen degradation types in Fig. 10. The first and second rows correspond to underwater image enhancement and desnowing, respectively. When directly generalized to these unseen degradations, existing methods often exhibit inaccurate color correction in underwater scenes or leave noticeable snow artifacts in

restored images. In comparison, DFC-IR produces more faithful color restoration with improved visual clarity for underwater images, and more effectively removes snow streaks and particles while preserving scene structures and fine details. The restored results contain fewer residual artifacts and maintain more natural local textures. These observations further demonstrate the generalization capability of DFC-guided restoration under degradation types beyond the training distribution.

**Real-world individual degradations.** Figure 11 presents visual comparisons on real-world individual degradations. Compared with synthetic degradations, real-world cases usually contain more complex and spatially varying degradation patterns, which can lead to residual artifacts, inaccurate color reconstruction, or over-smoothed details for existing methods. In contrast, DFC-IR more thoroughly suppresses the degradation effects while maintaining natural image appearance. In particular, the restored images exhibit cleaner structures, more faithful colors, and sharper local details, with fewer visible degradation residues. These results further demonstrate the effectiveness of DFC-guided degradation representation under real-world scenarios.

**Real-world composite degradations.** Additional visual results on real-world composite degradations are shown in Fig. 12. Such scenarios are particularly challenging because multiple degradation effects coexist in the same image and may interfere with each other during restoration. Existing approaches may incompletely remove one degradation or introduce additional artifacts when simultaneously handling multiple degradation components. DFC-IR more effectively suppresses the mixed degradations while preserving important scene structures, local textures, and visually natural appearance. The restored results exhibit fewer residual artifacts and clearer scene details, demonstrating the robustness of DFC-guided restoration under complex real-world degradation combinations.

### F.2 Three-Degradation Restoration

We provide additional visual comparisons under the three-degradation setting in Fig. 13, Fig. 14, and Fig. 15. DFC-IR consistently produces high-quality restoration results across denoising, deraining, and dehazing tasks. Compared with existing approaches, our method more effectively removes degradation artifacts while preserving fine textures, structural details, and natural color appearance. In challenging regions with dense degradations or complex local structures, competing methods may leave visible residual artifacts or over-smooth informative image details, whereas DFC-IR maintains a better balance between degradation removal and detail preservation. In denoising scenarios, DFC-IR suppresses noise while maintaining sharp edges and local details, avoiding excessive smoothing in textured regions. For deraining, our method removes rain streaks more thoroughly while preserving background structures with fewer visible artifacts. In dehazing cases, DFC-IR better recovers scene contrast and color fidelity while reducing residual haze. Moreover, the corresponding error maps further show that DFC-IR produces fewer restoration errors compared with other approaches, particularly around structural boundaries and degradation-heavy regions. These visual results further demonstrate the effectiveness of DFC-guided degradation representation for handling diverse degradation types within a unified restoration framework.

### F.3 Five-Degradation Restoration

Figure 16 and Fig. 17 present additional visual comparisons under the five-degradation setting, which further includes deblurring and low-light enhancement. DFC-IR achieves visually superior restoration results across these degradations, producing clearer structures and fewer residual artifacts than existing methods. In the original denoising, deraining, and dehazing tasks, DFC-IR continues to provide effective degradation removal while preserving local textures and scene structures. For deblurring, our method better recovers sharp edges and fine textures from blurred inputs, with reduced ringing or over-smoothing around object boundaries. For low-light enhancement, DFC-IR effectively restores illumination and improves visibility while preserving image details and maintaining natural appearance. Existing methods may produce insufficient enhancement, residual blur, or noticeable artifact patterns under these degradations, whereas DFC-IR yields more stable restoration quality across tasks. The corresponding error maps further reveal that DFC-IR leaves fewer restoration errors across different degradation conditions, indicating more accurate recovery of clean image content. Together with the results on the three-degradation setting, these visual comparisons verify that DFC-guided degradation representation provides effective guidance for all-in-one image restoration under diverse and complex degradations.

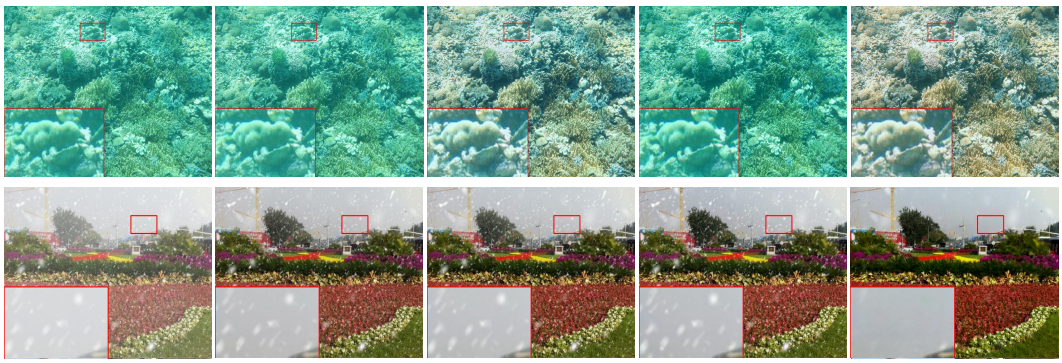



Figure 10: Visual comparisons on unseen restoration tasks, with underwater image enhancement on the UIEB dataset shown in the first row and desnowing on the CSD dataset shown in the second row. Existing methods often leave color distortion or residual snow artifacts, while DFC-IR achieves more natural color recovery and cleaner artifact removal with better detail preservation. Please zoom in for the best view.




Figure 11: Visual comparisons of generalization evaluation with five-degradation models on real-world NH-HAZE and SPANet. Existing methods often retain visible degradation residues or produce inaccurate colors, whereas DFC-IR yields cleaner restoration with more natural appearance and sharper details. Please zoom in for the best view.

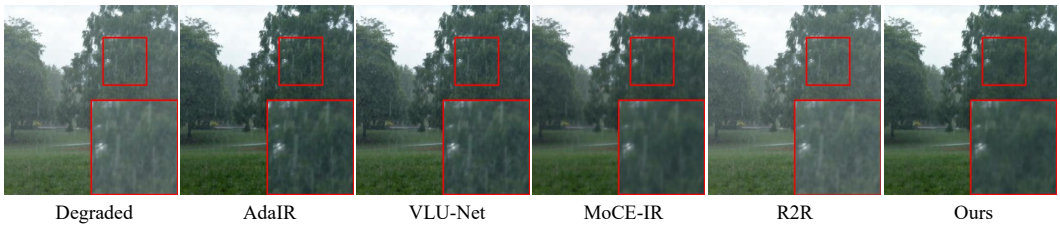



Figure 12: Visual comparisons on real-world haze-and-rain composite degradation scenarios. Existing methods may leave residual haze or rain artifacts, while DFC-IR more effectively removes both degradations and preserves clearer structures and finer textures. Please zoom in for the best view.

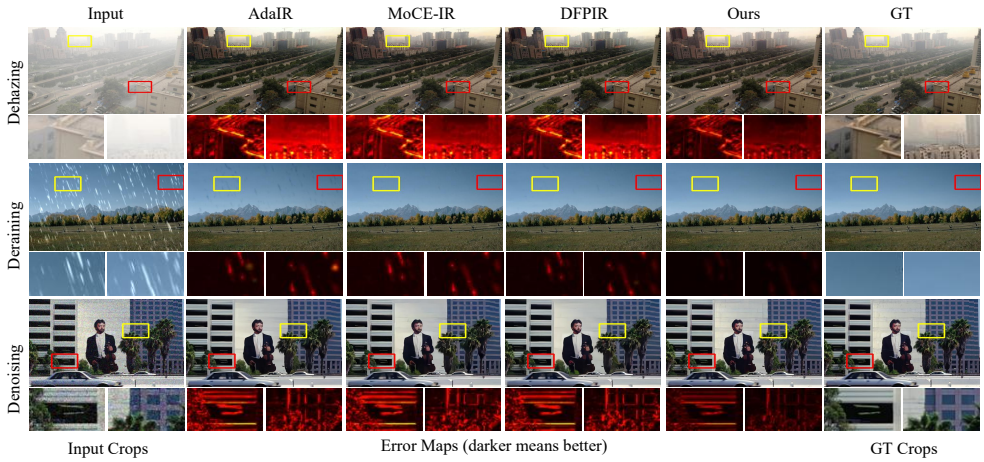



Figure 13: Visual comparison on the three-degradation all-in-one restoration setting. Please zoom in for the best view.

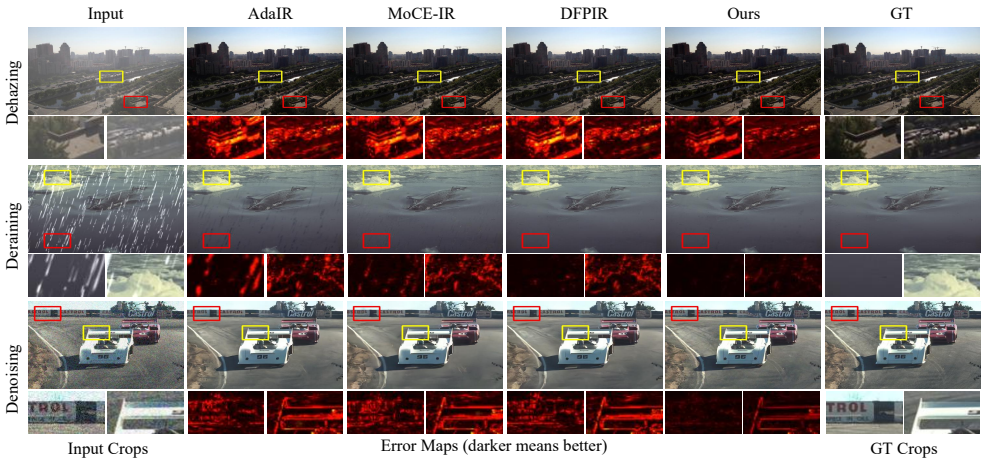



Figure 14: Additional visual comparison on the three-degradation all-in-one restoration setting. Please zoom in for the best view.

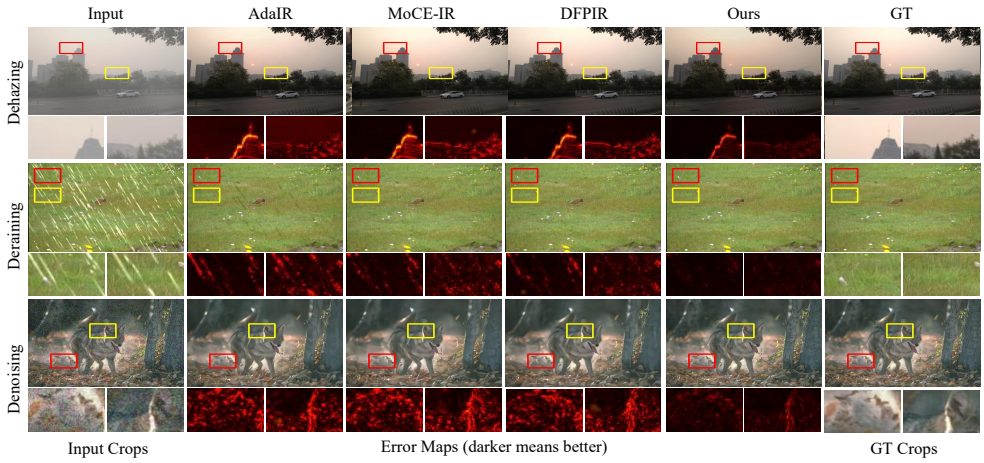



Figure 15: Additional visual comparison on the three-degradation all-in-one restoration setting. Please zoom in for the best view.

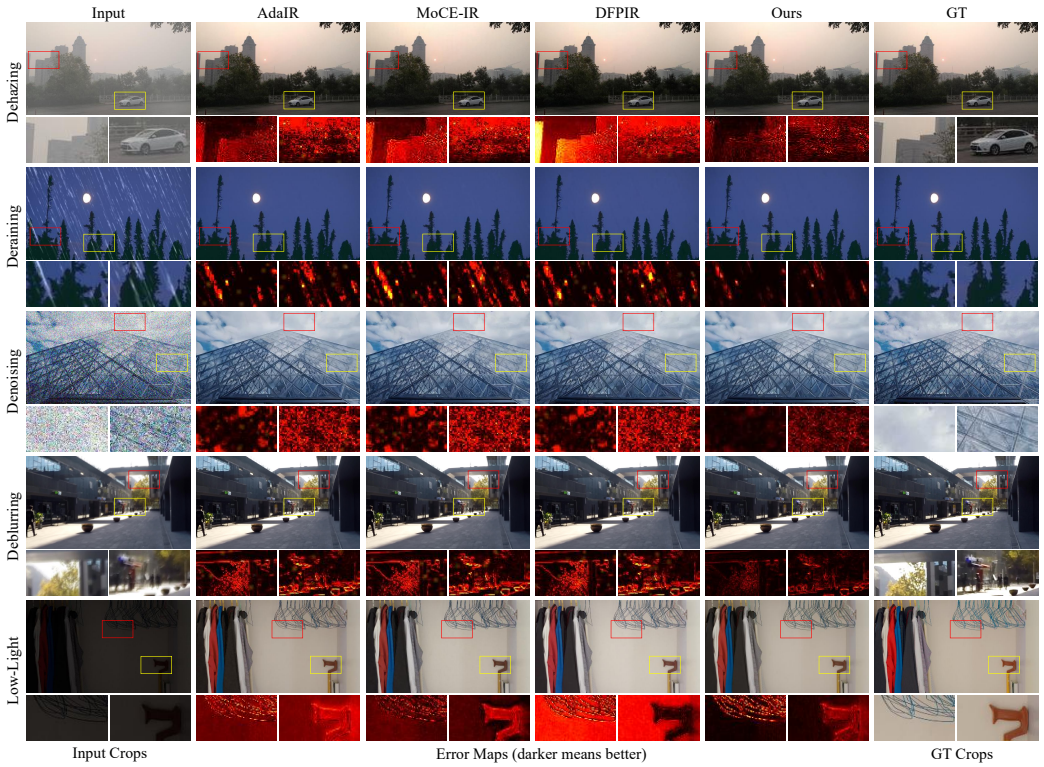



Figure 16: Visual comparison on the five-degradation all-in-one restoration setting. Please zoom in for the best view.




Figure 17: Additional visual comparison on the five-degradation all-in-one restoration setting. Please zoom in for the best view.