DECOMPOSE TO RESTORE: FACTOR-SPECIFIC ADAPTATION FOR IMAGE RESTORATION UNDER COMPOSITE DEGRADATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

Image restoration under composite degradations requires a unified model to handle multiple degradation factors that often co-occur in real-world images. Existing all-in-one methods typically treat each degradation type as an independent restoration condition or use degradation cues only as implicit conditioning, which limits factor-specific restoration behavior under composite degradations. We observe that restoration models already encode degradation-factor information, but such information is not well organized for factor-wise adaptation. We propose DFLoRA, a unified restoration network that decomposes composite degradations into reusable factors with dedicated low-rank adaptations. DFLoRA explicitly organizes latent degradation information into reusable factor representations and composition coefficients, decomposes a composite degradation into its fundamental factors, and equips each factor with a dedicated low-rank restoration executor. The resulting factor-specific adaptations are then composed according to the inferred degradation composition, directly translating degradation composition into adaptation composition. On the composite degradation restoration benchmark, DFLoRA achieves a PSNR gain of 1.23 dB over the strongest competing method, and it also achieves state-of-the-art performance across multiple single-degradation tasks in the standard all-in-one restoration setting. Our code will be publicly available.

1 INTRODUCTION

Image restoration aims to recover a clean image from its degraded observation. Driven by large-scale synthetic data and increasingly powerful architectures, deep restoration models can now recover individual degradations such as rain, haze, snow, and low light with high fidelity (Zamir et al., 2021; Qiu et al., 2023; Chen et al., 2021; Cai et al., 2023; Chen et al., 2022). This progress has motivated *all-in-one* restoration, where a single network is trained to handle multiple degradation types jointly (Li et al., 2022; Potlapalli et al., 2023; Valanarasu et al., 2022). In the prevailing paradigm, however, "all-in-one" largely follows a *multi-single-degradation* setting, where training and evaluation are primarily organized around separate single-degradation tasks rather than explicitly modeling the composition of co-occurring degradation factors. Real-world images, in contrast, are rarely corrupted by a single factor. For example, night scenes may simultaneously suffer from low light and noise, while haze often co-occurs with rain or snow. Composite degradation cannot be fully characterized by independent single degradations because their effects can interact. For example, haze can obscure rain streaks, while low-light conditions can challenge the assumptions underlying conventional dehazing methods. As illustrated in Fig. 1(a), multi-single-degradation restoration can leave residual degradations in composite inputs when the underlying factors remain entangled.

Recent benchmarks such as CDD-11 (Guo et al., 2024) have introduced composite degradation settings for evaluating all-in-one restoration models. Existing approaches (Zhu et al., 2025; Mao et al., 2026; Lee et al., 2026) demonstrate that a shared model can handle different combinations of degradations, extending all-in-one restoration beyond isolated degradation types. Yet, these combinations are typically handled as joint degradation conditions, without explicitly modeling the underlying factors and their composition. For example, haze and rain may co-occur in one composite, while low light and haze may form another, suggesting that the same degradation factor can recur across

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.

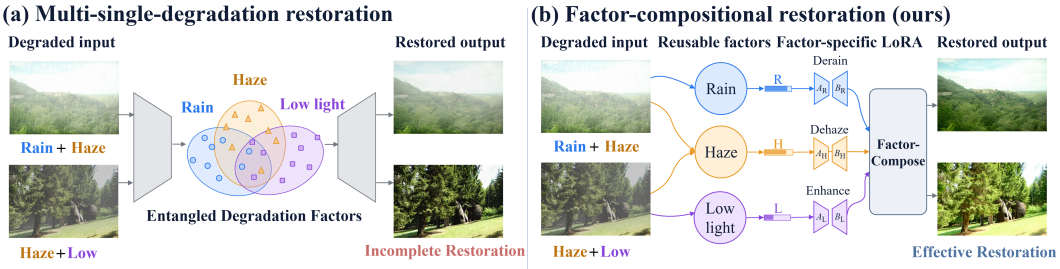

Figure 1: (a) Multi-single-degradation restoration with entangled degradation factors can leave residual degradations in composite inputs. (b) DFLoRA explicitly organizes reusable degradation factors and composes their corresponding low-rank adaptations for composite image restoration.

different composites. This motivates a more fundamental question: *can composite degradations be explicitly decomposed into reusable fundamental factors, and can restoration adaptation be organized according to their composition?*

Interestingly, we find that existing all-in-one restoration models already encode substantial degradation-factor information, yet this information is not well organized for factor-specific adaptation. A linear probe on a backbone trained under a composite degradation setting achieves a mean per-factor AUC of 0.997, indicating that degradation-factor information is already encoded and linearly recoverable (Appendix B). However, an 11-way nearest-centroid classifier achieves only 67.9% accuracy (Table 4), revealing that the representation remains poorly organized at the factor level. These observations highlight a gap between degradation awareness and factor-specific restoration, with factor information encoded but not explicitly decomposed into reusable factors.

Building on this insight, we propose DFLoRA, a factor-compositional low-rank adaptation framework that transforms implicit degradation knowledge into an explicit factor structure and provides each factor with a dedicated restoration adaptation. As illustrated in Fig. 1(b), we (i) explicitly organize latent degradation information using learnable factor prototypes and multi-label supervision, yielding factor-level representations and composition coefficients; (ii) decompose a composite degradation into its fundamental factors, e.g., low+haze+rain $\rightarrow$ low, haze, rain; (iii) equip each factor with a dedicated restoration executor implemented as a low-rank branch using the LoRA parameterization (Hu et al., 2021), establishing a direct correspondence between degradation factors and restoration adaptations; (iv) compose the factor-specific adaptations according to the inferred degradation composition, weighting each adaptation by its corresponding composition coefficient. Degradation composition is thereby directly translated into adaptation composition.

To summarize, our contributions are:

- We advance all-in-one restoration toward explicit modeling of composite degradations by identifying a gap between degradation awareness and factor-wise adaptation, motivating a compositional view based on reusable degradation factors.
- We propose DFLoRA, which explicitly organizes latent degradation information, assigns each fundamental factor an independent factor-specific LoRA, and composes these adaptations according to the degradation composition of the input.
- We comprehensively validate DFLoRA on CDD-11, achieving 32.01 dB PSNR and a 1.23 dB improvement over existing state-of-the-art methods, with consistent gains on composite degradations. We further achieve state-of-the-art performance under the multi-single-degradation all-in-one restoration setting, together with systematic analyses of factor representation, factor composition, and adaptation behavior.

2 Related Work

2.1 All-in-One Image Restoration

All-in-one image restoration aims to unify multiple restoration tasks within a single model. Early approaches focus on learning shared representations that accommodate different degradation types, while subsequent methods introduce degradation-aware feature modulation and conditional mechanisms to adapt a common restoration network to individual degradation conditions (Li et al., 2022; Potlapalli et al., 2023; Hu et al., 2025; Valanarasu et al., 2022). More recently, expert-based methods allocate specialized restoration capacity across different degradation conditions (Zamfir et al., 2025; Wang et al., 2025a; Lin et al., 2026), while diffusion-based approaches exploit pretrained generative priors for unified restoration (Cheng et al., 2026; Liu et al., 2026; Jiang et al., 2024). LoRA-based approaches equip models with degradation-specific restoration capabilities through low-rank adapters (Ai et al., 2024; Serrano-Lozano et al., 2026; Zhang et al., 2024). Despite these architectural advances, the prevailing formulation remains largely *multi-single-degradation*, where each degradation type is treated as an individual restoration condition with its corresponding restoration capability. This formulation has proven effective for handling diverse single degradations within one unified model, but it implicitly assumes a direct correspondence between a degradation condition and a restoration objective. When multiple degradation factors co-occur, this one-factor-to-one-restoration correspondence no longer directly holds, motivating explicit modeling of degradation composition.

2.2 Composite Degradation Restoration

Recent studies have increasingly investigated image restoration under composite degradations, where multiple degradation factors co-occur in a single image (Jiang et al., 2025; Wang et al., 2025b; Zhou et al., 2025). Existing approaches can be broadly grouped into three directions. Unified composite restoration methods typically retain a shared restoration backbone while introducing degradation-aware conditions or adaptive mechanisms to specialize the restoration process for different combinations of degradations. OneRestore (Guo et al., 2024) uses scene descriptors to condition the restoration network on the overall degradation scenario, including composite conditions. AllRestorer (Mao et al., 2026) models composite scene embeddings and further estimates adaptive weights for individual degradations to control their restoration intensity. CDIR (Cui et al., 2026) extracts degradation-related context priors from the input and iteratively updates them with intermediate restoration features to guide its dynamic operations. Prior-guided methods instead exploit additional degradation-related priors to complement the information available from the degraded input, with MMFE-IR (Lee et al., 2026) deriving such priors from a multimodal large language model. Agent-based methods (Zhu et al., 2025; Lu et al., 2026; Chen et al., 2026) formulate composite restoration as a sequence of perception, reasoning, and restoration operations. Despite these different strategies, existing approaches primarily use degradation information to condition restoration, provide external guidance, or organize restoration procedures, rather than establishing an explicit correspondence between reusable degradation factors and dedicated restoration adaptations. In contrast, we model composite degradation at the factor level and directly translate degradation composition into the composition of factor-specific restoration adaptations.

3 Method

3.1 Overview

Composite image restoration requires a unified model for multiple co-occurring degradation factors. Rather than treating each degradation combination as an independent restoration condition, we propose **D**egradation **F**actor-specific **LoRA** (DFLoRA), a unified restoration network that decomposes composite degradations into reusable factors and composes the corresponding factor-specific LoRA adaptations.

Given a degraded image $\mathbf{x} \in \mathbb{R}^{C_{\text{in}} \times H \times W}$, DFLoRA processes the input through an encoder–decoder hierarchy and produces an intermediate bottleneck feature $\mathbf{F}_{\text{mid}}$ together with multi-scale decoder features. We introduce two key components within this architecture. **Degradation Factor Organization (DFO)** decomposes the latent degradation information over K reusable degradation

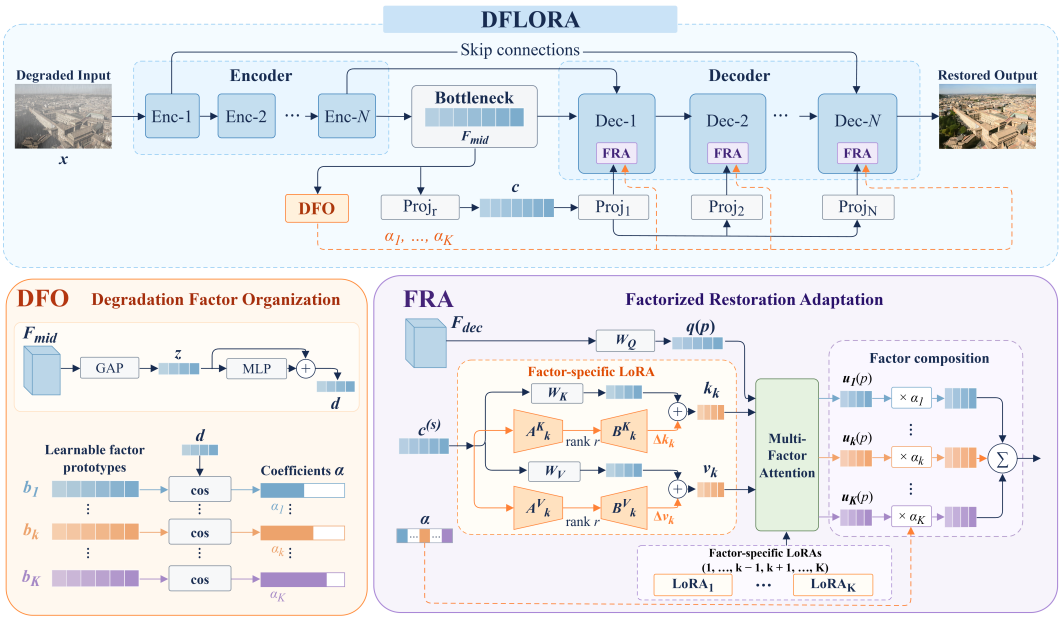

Figure 2: **Overview of DFLoRA.** DFLoRA integrates Degradation Factor Organization (DFO) and Factorized Restoration Adaptation (FRA). DFO organizes bottleneck features using learnable factor prototypes and predicts non-exclusive factor composition coefficients. FRA applies factor-specific LoRA to shared K/V projections of the stage-conditioned restoration context. Decoder queries spatially modulate the resulting factor-specific responses, which are composed according to the predicted coefficients to restore images under composite degradations.

factors and estimates non-exclusive decomposition coefficients $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_K]^\top$. **Factorized Restoration Adaptation (FRA)** translates this factor composition into restoration adaptations. It uses the bottleneck representation to provide stage-conditioned restoration context and introduces factor-specific LoRA residuals to specialize the shared K/V projections for individual degradation factors, and composes their restoration responses according to the estimated decomposition coefficients.

DFO determines the constituent factors of the input degradation, while FRA constructs and composes their restoration responses. Within FRA, decoder features spatially modulate the factor-specific responses, which are then combined according to the estimated decomposition coefficients. This design maps degradation composition directly to adaptation composition, allowing each factor-specific adaptation to be reused across different composite conditions without requiring a separate restoration branch for each degradation combination. Fig. 2 illustrates the overall architecture.

3.2 Degradation Factor Organization

We first organize the degradation information contained in the intermediate representation of DFLoRA into an explicit set of reusable constituent factors. Given the bottleneck feature $\mathbf{F}_{\text{mid}}$, we apply global average pooling (GAP) to obtain a compact descriptor

$$\mathbf{z} = \texttt{GAP}(\mathbf{F}_{\text{mid}}) \in \mathbb{R}^d, \tag{1}$$

where $\texttt{GAP}(\cdot)$ denotes global average pooling and d is the dimensionality of the bottleneck representation.

The descriptor $\mathbf{z}$ already contains degradation-related information, but this information is not explicitly separated according to constituent degradation factors. We therefore introduce DFO to construct a factor-oriented representation while keeping its supervision decoupled from the main restoration pathway. Specifically, we apply DFO to a detached copy of $\mathbf{z}$:

$$\mathbf{d} = \texttt{sg}(\mathbf{z}) + \texttt{MLP}(\texttt{LN}(\texttt{sg}(\mathbf{z}))), \tag{2}$$

where $\mathtt{sg}(\cdot)$ denotes the stop-gradient operation, $\mathtt{LN}(\cdot)$ denotes layer normalization, and $\mathtt{MLP}(\cdot)$ denotes a lightweight residual mapping whose output branch is initialized to zero. We isolate the gradients so that the restoration losses optimize only the restoration branch, while the factor supervision loss is dedicated to learning factor decomposition. The residual mapping then provides a compact feature space in which the latent degradation information can be reorganized according to constituent factors.

We maintain K learnable factor prototypes $\{\mathbf{b}_k\}_{k=1}^{K}$, where $\mathbf{b}_k \in \mathbb{R}^d$ represents the k-th degradation factor. For each factor, we measure the similarity between the organized degradation representation and its prototype using normalized cosine similarity:

$$\ell_k = \tau \frac{\mathbf{d}^\top \mathbf{b}_k}{\|\mathbf{d}\|_2 \|\mathbf{b}_k\|_2}, \qquad \alpha_k = \sigma(\ell_k), \tag{3}$$

where $\tau > 0$ is a learnable bounded logit scale and $\sigma(\cdot)$ denotes the sigmoid function.

We use independent sigmoid activations rather than a softmax classifier because multiple degradation factors may coexist in the same image. During training, we supervise the factor coefficients with a standard multi-label binary cross-entropy loss, denoted by $\mathcal{L}_{\mathrm{fac}}$, using the factor-presence labels of each training sample.

The resulting vector $\boldsymbol{\alpha} = [\alpha_1, \ldots, \alpha_K]^\top$ serves as the decomposition coefficients of the input degradation over the learned factor vocabulary. A composite degradation is therefore represented by the joint activation of its constituent factors rather than by assigning the entire input to a single degradation-combination class. We use these coefficients to compose the corresponding restoration adaptations in the decoder.

3.3 Factorized Restoration Adaptation

Given the degradation decomposition from DFO, we introduce FRA to associate each constituent factor with a dedicated low-rank restoration adaptation and compose their responses according to the inferred degradation composition. This allows the decoder to construct restoration behavior from the same factor structure used to describe the input degradation.

For decoder stage $s \in \{1, \ldots, S\}$, we map the undetached bottleneck descriptor $\mathbf{z}$ to a stage-conditioned restoration context

$$\mathbf{c}^{(s)} = \mathtt{LN}(\mathtt{Proj}_s(\mathtt{Proj}_{\mathrm{r}}(\mathbf{z}))), \tag{4}$$

where $\mathtt{Proj}_{\mathrm{r}}(\cdot)$ denotes the restoration projection shared across decoder stages and $\mathtt{Proj}_s(\cdot)$ denotes the projection associated with decoder stage s. We share the restoration context across all degradation factors and introduce factor-specific specialization only through lightweight LoRA residuals.

For a decoder block, let $\mathbf{W}_K$ and $\mathbf{W}_V$ denote the key and value projection matrices shared across all degradation factors. Omitting the stage index for clarity, we define the key and value representations for factor k as

$$\mathbf{k}_k = \mathbf{W}_K \mathbf{c} + \frac{\gamma}{r} \mathbf{B}_k^K \mathbf{A}_k^K \mathbf{c}, \qquad \mathbf{v}_k = \mathbf{W}_V \mathbf{c} + \frac{\gamma}{r} \mathbf{B}_k^V \mathbf{A}_k^V \mathbf{c}, \tag{5}$$

where $\mathbf{A}_k^K$ and $\mathbf{A}_k^V$ project the restoration context into a rank-r latent space, while $\mathbf{B}_k^K$ and $\mathbf{B}_k^V$ project it back to the decoder feature dimension. The factor γ/r controls the scale of the low-rank residual. The shared projections model restoration information common to different degradation factors, whereas the LoRA residuals introduce factor-specific deviations. We therefore obtain factor-specific restoration behavior with only lightweight low-rank parameters rather than separate restoration branches.

We further make each factor-specific adaptation responsive to the current decoder feature. Let $\mathbf{q}_h(p)$ denote the query vector at spatial position p and attention head h, obtained from the decoder feature through the shared query projection $\mathbf{W}_Q$. For factor k, we compute its spatial response as

$$g_{h,k}(p) = 2\,\sigma\left(\frac{\mathbf{q}_h(p)^\top \mathbf{k}_{h,k}}{\sqrt{d_h}}\right), \qquad \mathbf{u}_{h,k}(p) = g_{h,k}(p)\mathbf{v}_{h,k}, \tag{6}$$

where d_h denotes the attention-head dimension. The image-level decomposition coefficients specify which degradation factors should contribute to restoration, while $g_{h,k}(p)$ provides a spatially varying response for each factor within the decoder feature.

Let a_k denote the factor coefficient used by the decoder. At inference time, $a_k = \alpha_k$, while training may use supervised or mixed coefficients as described in the experimental settings. We compose the factor-specific responses as

$$Z_a = \max\left(1, \sum_{k=1}^{K} a_k\right), \qquad \mathbf{o}_h(p) = Z_a^{-1} \sum_{k=1}^{K} a_k \, \mathbf{u}_{h,k}(p). \tag{7}$$

We project the composed response back to the decoder feature space and add it as a residual update. The normalization keeps the magnitude of the factor-specific update from increasing simply because more degradation factors are active. When the total factor activation is below one, the update remains proportionally attenuated rather than being renormalized to full strength. If no factor is active, the factor-specific branch contributes no additional update.

With this formulation, DFO provides the degradation composition and the decoder uses the same coefficients to combine the corresponding factor-specific LoRA responses. A factor-specific adaptation can therefore be reused across different composite conditions whenever the associated degradation factor is present.

4 EXPERIMENTS

We evaluate DFLoRA under two complementary restoration settings.

Composite Degradation Restoration. Our main experiments are conducted on CDD-11 (Guo et al., 2024), which contains 11 degradation conditions constructed from four constituent degradation factors: rain, haze, snow, and low light. The benchmark includes single-, double-, and triple-degradation cases, providing a controlled setting for evaluating degradation decomposition and composite restoration. We use the CDD-11 training set for optimization and evaluate on its corresponding test sets.

Multi-Single-Degradation Restoration. We further evaluate DFLoRA across five individual restoration tasks with distinct degradation types and datasets. Specifically, BSD400 (Arbelaez et al., 2010) and WED (Ma et al., 2016) are used for denoising training, with BSD68 (Martin et al., 2001) under Gaussian noise of $\sigma = 25$ for evaluation; Rain100L (Yang et al., 2019) is used for deraining; SOTS (Li et al., 2018) for dehazing; GoPro (Nah et al., 2017) for deblurring; and LOL-v1 (Wei et al., 2018) for low-light enhancement. Unlike CDD-11, each image in this setting corresponds to a single degradation type. This setting evaluates whether DFLoRA can maintain strong task-specific restoration capability across heterogeneous degradation types within a unified network. Detailed dataset configurations for both settings are provided in Appendix A.1.

Implementation Details. We train DFLoRA with three complementary objectives, including pixel-space reconstruction for image fidelity, frequency-domain reconstruction for structural and high-frequency consistency, and degradation-factor supervision for learning the factor coefficients in Sec. 3.2. Training progressively transitions from ground-truth factor compositions to DFO-predicted coefficients, while inference relies entirely on the predicted coefficients. Detailed loss formulations and training configurations are provided in Appendix A.2.

4.1 COMPOSITE DEGRADATION RESTORATION

We compare DFLoRA with representative all-in-one and composite image restoration methods in Table 1. DFLoRA achieves the best overall performance, reaching an average PSNR/SSIM of 32.01 dB/0.900 and outperforming CDIR-S (Cui et al., 2026), the strongest competing method in average PSNR, by 1.23 dB. The improvement remains consistent as degradation complexity increases, with DFLoRA achieving 36.15, 30.70, and 27.01 dB on single-, double-, and triple-degradation cases, respectively. On the composite subset, it further improves PSNR from 28.65 to 29.65 dB and SSIM from 0.860 to 0.870. These results support the effectiveness of explicitly decomposing composite degradations into constituent factors and translating the inferred degradation composition

Table 1: Quantitative comparison under the composite degradation setting. PSNR (dB) / SSIM are reported. Higher values indicate better restoration quality. Bold and underlined values denote the **best** and second-best results.

Method	Single PSNR/SSIM	Double PSNR/SSIM	Triple PSNR/SSIM	Composite PSNR/SSIM	Avg. PSNR/SSIM
AirNet (Li et al., 2022)	25.59/.885	22.97/.796	22.02/.716	22.70/.773	23.75/.814
PromptIR (Potlapalli et al., 2023)	28.88/.920	24.46/.834	23.53/.750	24.19/.810	25.90/.850
WGWSNet (Zhu et al., 2023)	29.97/.923	25.76/.850	23.93/.772	25.24/.828	26.96/.863
OneRestore (Guo et al., 2024)	31.81/.939	27.65/.871	25.23/.796	26.96/.850	28.72/.882
MoCE-IR-S (Zamfir et al., 2025)	32.53/.941	27.73/.869	25.40/.790	27.06/.846	29.05/.881
MIRAGE (Ren et al., 2026)	32.79/.944	28.00/.875	25.73/.800	27.35/.854	29.33/.887
DCPT-NAFNet (Hu et al., 2025)	34.50/.947	29.33/.881	26.31/.806	28.47/.859	30.66/.891
CDIR-S (Cui et al., 2026)	34.50/.946	29.56/.881	26.39/.808	28.65/.860	30.78/.892
DFLoRA (Ours)	**36.15/.952**	**30.70/.891**	**27.01/.820**	**29.65/.870**	**32.01/.900**

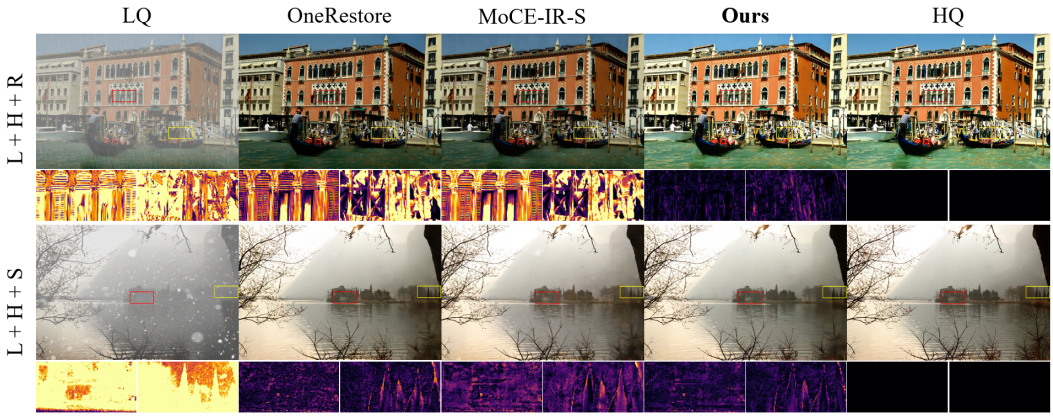

Figure 3: Visual comparison on CDD-11 under triple degradations (low-haze-rain / low-haze-snow). Insets show error maps relative to HQ for the marked 128×64 patches, and brighter colors indicate larger errors.

into factor-specific restoration behavior. DFLoRA achieves this performance with 29.7M parameters and 64.77 GFLOPs at a 224×224 input resolution. As shown in Fig. 3, it also more effectively suppresses co-occurring degradation effects while preserving structural details and local contrast.

4.2 Multi-Single-Degradation Restoration

We further evaluate DFLoRA across five single-degradation restoration tasks and compare it with representative single-task and all-in-one restoration methods in Table 2. DFLoRA achieves the best average performance at 31.12 dB/0.922, improving the average PSNR of FIT (He et al., 2026) by 0.40 dB while also obtaining the highest average SSIM. It achieves the best dehazing result at 32.97 dB/0.982 and the highest SSIM on denoising (0.894) and low-light enhancement (0.861), while remaining competitive on deraining and deblurring. These results show that DFLoRA is not restricted to composite-degradation inputs and can maintain effective task-specific restoration across heterogeneous degradation types within a unified network. The qualitative examples in Fig. 4 further show sharper structures and cleaner local details without obvious restoration artifacts.

4.3 Ablation Studies

4.3.1 Factor-Specific LoRA

We first investigate the contribution of factor-specific low-rank adaptation within DFLoRA through a progressive ablation. We compare three variants under the same training and evaluation setting: (1) **Baseline**, which removes DFO, restoration context, and factor-specific LoRA; (2) **+ DFO &**

Table 2: Quantitative comparison under the multi-single-degradation setting. PSNR (dB) / SSIM are reported. Higher values indicate better restoration quality. Bold and underlined values denote the **best** and second-best results.

Method	Denoising PSNR/SSIM	Deraining PSNR/SSIM	Dehazing PSNR/SSIM	Deblurring PSNR/SSIM	Low-light PSNR/SSIM	Avg. PSNR/SSIM
AirNet (Li et al., 2022)	30.91/.882	32.98/.951	21.04/.884	24.35/.781	18.18/.735	25.49/.847
Restormer (Zamir et al., 2022)	31.49/.884	34.81/.962	24.09/.927	27.22/.829	20.41/.806	27.60/.881
PromptIR (Potlapalli et al., 2023)	31.47/.886	36.37/.970	26.54/.949	28.71/.881	22.68/.832	29.15/.904
InstructIR (Conde et al., 2024)	31.40/.887	36.84/.973	27.10/.956	29.40/.886	23.00/.836	29.55/.907
AdaIR (Cui et al., 2025)	31.35/.888	38.02/.981	30.53/.978	28.12/.858	23.00/.845	30.20/.910
MoCE-IR (Zamfir et al., 2025)	31.34/.887	38.04/.982	30.48/.974	**30.05/.899**	23.00/.852	30.58/.919
DFPIR (Tian et al., 2025)	31.29/.889	37.62/.978	31.64/.979	28.82/.873	**23.82**/.843	30.64/.913
FIT (He et al., 2026)	**31.57**/.893	**38.56/.983**	30.81/.979	29.52/.890	23.13/.848	30.72/.919
DFLoRA (Ours)	31.53/**.894**	38.07/.982	**32.97/.982**	29.79/.892	23.25/**.861**	**31.12/.922**

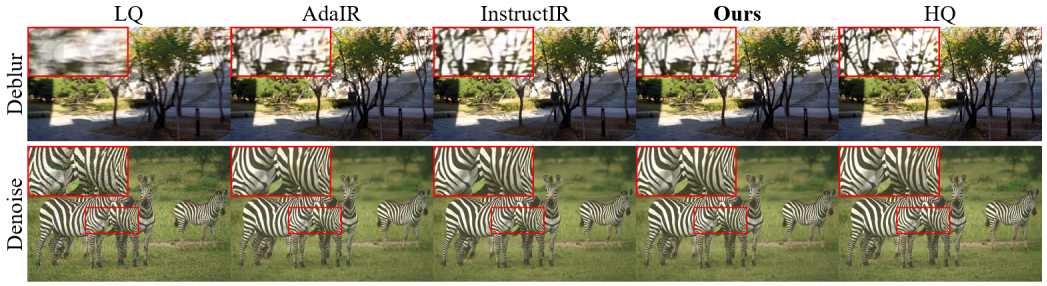

Figure 4: Visual comparison on deblurring (top) and denoising (bottom) in the multi-single-degradation setting. Insets show 128×64 patches from the red-boxed regions.

restoration context, which adds DFO and stage-conditioned restoration context to the baseline, using shared K/V projections without factor-specific LoRA residuals; and (3) **+ Factor-specific LoRA**, which further adds dedicated low-rank K/V adaptations for individual degradation factors to the second variant, yielding the complete DFLoRA.

The results are reported in Table 3. Introducing DFO and restoration context improves the average PSNR from 31.65 to 31.87 dB, confirming that the intermediate restoration representation contains useful degradation information. Adding factor-specific LoRA further increases the average PSNR to 32.01 dB, with only 0.52M additional parameters over the preceding variant. The same trend is observed on the composite subset, where PSNR improves from 29.54 to 29.65 dB.

These results demonstrate that DFLoRA makes more effective use of explicitly decomposed degradation factors through factor-specific LoRA, associating each factor with a dedicated low-rank adaptation instead of using the factors only as a common conditioning signal.

4.3.2 Degradation Factor Organization and Composition

Factor Organization. We next examine whether DFO transforms the implicit degradation information encoded in the intermediate representation of DFLoRA into a more structured factor representation. We compare the original bottleneck descriptor $\mathbf{z}$ with the organized representation $\mathbf{d}$ produced by DFO. Factor organization is evaluated using the maximum cross-factor cosine similarity, 11-way nearest-centroid classification accuracy, and exact-match accuracy of the predicted factor composition. Detailed evaluation protocols are provided in Appendix B.

As reported in Table 4, the maximum cross-factor cosine similarity decreases substantially from 0.712 to 0.090, indicating stronger separation among degradation factors. Meanwhile, nearest-centroid accuracy increases from 67.9% to 98.8%, and the predicted factor composition reaches 99.00% exact-match accuracy. Fig. 5 provides a complementary visualization, showing that the representation after DFO exhibits substantially more organized degradation-dependent clusters than the original bottleneck representation. These results show that DFO restructures the degradation

Table 3: Progressive ablation of factor-specific LoRA in terms of PSNR (dB). Each successive row builds on the preceding variant, and ΔP denotes the parameter increase relative to that variant (M).

Setting	Comp.	Avg.	ΔP
Baseline	29.33	31.65	–
+ DFO & restoration context	29.54	31.87	+6.93
+ Factor-specific LoRA	**29.65**	**32.01**	**+0.52**

Table 4: Analysis of DFO. Max cross-factor cos. measures inter-factor similarity; 11-way centroid measures degradation classification accuracy; composition exact-match measures complete factor-set prediction accuracy.

Metric	w/o DFO	DFO
Max cross-factor cos.	0.712	0.090
11-way centroid (%)	67.9	98.8
Composition exact-match (%)	–	99.00

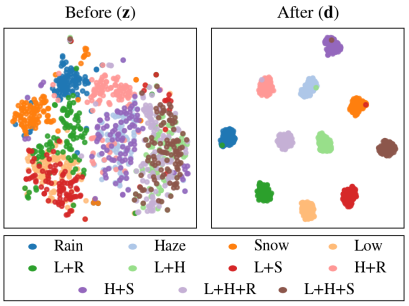

Figure 5: t-SNE visualization of the bottleneck features colored by the eleven degradation conditions, before (z) and after (d) DFO.

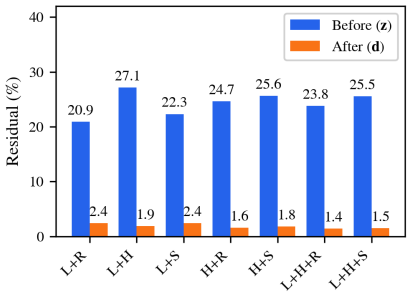

Figure 6: Reconstruction residual of composite features on the span of their constituent factors, before (z) and after (d) DFO.

information already present in the intermediate restoration representation into an explicit factor representation suitable for subsequent factor-specific adaptation.

Factor Composition. We further investigate whether the organized representation exhibits the compositional structure assumed by DFLoRA. For each composite degradation, we construct a constituent subspace by combining the principal-component bases of its constituent single-degradation tasks. We project the centered composite features onto this subspace and measure the normalized squared reconstruction residual, averaged over samples. A lower residual indicates that the constituent subspace captures more of the composite feature variation.

As shown in Fig. 6, DFO reduces the reconstruction residual from 20.9%–27.1% for the original representation $\mathbf{z}$ to only 1.4%–2.4% for the organized representation $\mathbf{d}$. This substantial reduction indicates that composite feature variation is more closely captured by the constituent single-degradation subspaces after DFO. Together with the restoration ablations above, this analysis supports the decomposition-to-adaptation formulation of DFLoRA, in which the inferred degradation composition directly determines the composition of factor-specific restoration adaptations.

5 Conclusion

In this work, we revisit all-in-one image restoration from the perspective of degradation composition. We observe that restoration backbones already encode rich degradation-factor information, but this information is not explicitly organized into reusable factors. Based on this observation, we propose DFLoRA, which explicitly organizes the latent degradation representation into fundamental factors and their composition coefficients, assigns each factor a dedicated low-rank restoration adaptation, and composes these adaptations according to the inferred degradation composition. This establishes a direct correspondence between degradation factors and restoration capabilities, allowing the same factor-specific adaptation to be reused across different composite conditions. Experiments on CDD-11 demonstrate improvements across diverse degradation combinations, while the multi-single-degradation evaluation demonstrates general restoration capability. Overall, DFLoRA provides a factor-level perspective on image restoration under composite degradations by explicitly relating the composition of degradation factors to the composition of restoration adaptations.

AI USE STATEMENT

In this work, we used generative AI tools to assist with translation and to aid and polish writing. All AI-assisted translations and edits were carefully reviewed by the authors for accuracy and consistency with the intended technical meaning. We take responsibility for the final content of this work, including text, claims, or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

We describe the datasets, preprocessing procedures, network architecture, and training configurations in Appendix A. The protocols for our representation analyses are provided in Appendix B. We will publicly release the training and evaluation code, configuration files, and pretrained model weights upon acceptance to facilitate reproduction of our results.

REFERENCES

Yuang Ai, Huaibo Huang, and Ran He. Lora-ir: Taming low-rank experts for efficient all-in-one image restoration. *arXiv preprint arXiv:2410.15385*, 2024.

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.

Yuanhao Cai, Hao Bian, Jing Lin, Haoqian Wang, Radu Timofte, and Yulun Zhang. Retinexformer: One-stage retinex-based transformer for low-light image enhancement. In *2023 IEEE/CVF international conference on computer vision (ICCV)*, pp. 12470–12479. IEEE, 2023.

I-Hsiang Chen, Isma Hadji, Enrique Sanchez, Adrian Bulat, Sy-Yen Kuo, Radu Timofte, Georgios Tzimiropoulos, and Brais Martinez. Restore, assess, repeat: A unified framework for iterative image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15506–15515, June 2026.

Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. In *European conference on computer vision*, pp. 17–33. Springer, 2022.

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.

Zihan Cheng, Liangtai Zhou, Dian Chen, Ni Tang, Xiaotong Luo, Yuan Xie, and Yanyun Qu. Unildiff: Unlocking the power of diffusion priors for all-in-one image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 37465–37475, 2026.

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, Syed Waqas Zamir, Salman Khan, Alois Knoll, Mubarak Shah, and Fahad Khan. Adair: Adaptive all-in-one image restoration via frequency mining and modulation. In *International Conference on Learning Representations*, volume 2025, pp. 101306–101327, 2025.

Yuning Cui, Wenqi Ren, Boxin Shi, Jianhou Gan, and Alois Knoll. Cdir: Lora-inspired attention for efficient composite degradation image restoration. *IEEE Transactions on Image Processing*, 2026.

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, 2024.

Zihao He, Yunfeng Wu, Xinchao Wang, and Songhua Liu. Bend the basics: Degradation-aware deformable tokenization for all-in-one image restoration. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=ZFOA8DZqhl.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.

JiaKui Hu, Lujia Jin, Zhengjian Yao, and Yanye Lu. Universal image restoration pre-training via degradation classification. In *International Conference on Learning Representations*, volume 2025, pp. 64647–64668, 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 Transactions on Pattern Analysis and Machine Intelligence*, 2025.

Yitong Jiang, Zhaoyang Zhang, Tianfan Xue, and Jinwei Gu. Autodir: Automatic all-in-one image restoration with latent diffusion. In *European Conference on Computer Vision*, pp. 340–359. Springer, 2024.

Eunho Lee, Rei Kawakami, and Youngbae Hwang. Leveraging multimodal large language models for all-in-one image restoration via a mixture of frequency experts. *arXiv preprint arXiv:2605.11444*, 2026.

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 *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 17431–17441. IEEE, 2022.

Jingbo Lin, Zhilu Zhang, Wenbo Li, Renjing Pei, Hang Xu, Hongzhi Zhang, and Wangmeng Zuo. Unirestorer: Universal image restoration via adaptively estimating image degradation at proper granularity. In *International Conference on Learning Representations*, volume 2026, pp. 2369–2409, 2026.

Jingren Liu, Shuning Xu, Qirui Yang, Yun Wang, Xiangyu Chen, and Zhong Ji. Fape-ir: frequency-aware planning and execution framework for all-in-one image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15561–15573, 2026.

Jianglin Lu, Yuanwei Wu, Ziyi Zhao, Hongcheng Wang, Felix Jimenez, Abrar Majeedi, and Yun Fu. Restore-r1: Efficient image restoration agents via reinforcement learning with multimodal llm perceptual feedback. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8629–8639, 2026.

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.

Jiawei Mao, Yu Yang, Xuesong Yin, Ling Shao, and Hao Tang. All-in-one transformer for image restoration under adverse weather degradations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 48(6):6628–6641, 2026. doi: 10.1109/TPAMI.2026.3658598.

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.

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.

Vaishnav Potlapalli, Syed Waqas Zamir, Salman Khan, and Fahad Khan. Promptir: Prompting for all-in-one image restoration. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.

Yuwei Qiu, Kaihao Zhang, Chenxi Wang, Wenhan Luo, Hongdong Li, and Zhi Jin. Mb-taylorformer: Multi-branch efficient transformer expanded by taylor formula for image dehazing. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 12756–12767. IEEE, 2023.

Bin Ren, Yawei Li, Xu Zheng, Yuqian Fu, Danda Pani Paudel, Hong Liu, Ming-Hsuan Yang, Luc Van Gool, and Nicu Sebe. Efficient degradation-agnostic image restoration via channel-wise functional decomposition and manifold regularization. In *International Conference on Learning Representations*, volume 2026, pp. 126404–126430, 2026.

David Serrano-Lozano, Luis Herranz, Shaolin Su, and Javier Vazquez-Corral. Adaptive blind all-in-one image restoration. *Computer Vision and Image Understanding*, pp. 104795, 2026.

Xiangpeng Tian, Xiangyu Liao, Xiao Liu, Meng Li, and Chao Ren. Degradation-aware feature perturbation for all-in-one image restoration. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 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 *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2343–2353. IEEE, 2022.

Tao Wang, Peiwen Xia, Bo Li, Peng-Tao Jiang, Zhe Kong, Kaihao Zhang, Tong Lu, and Wenhan Luo. Moerl: When mixture-of-experts meet reinforcement learning for adverse weather image restoration. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 13673–13683. IEEE, 2025a.

Zhongze Wang, Haitao Zhao, Lujian Yao, Jingchao Peng, and Kaijie Zhao. Dual-level prototype learning for composite degraded image restoration. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 14006–14061. IEEE, 2025b.

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

Wenhan Yang, Robby T Tan, Jiashi Feng, Zongming Guo, Shuicheng Yan, and Jiaying Liu. Joint rain detection and removal from a single image with contextualized deep networks. *IEEE transactions on pattern analysis and machine intelligence*, 42(6):1377–1393, 2019.

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 *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 12753–12763. IEEE, 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 *2021 IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pp. 14816–14826. IEEE, 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.

Cheng Zhang, Dong Gong, Jiumei He, Yu Zhu, Jinqiu Sun, and Yanning Zhang. Uir-lora: Achieving universal image restoration through multiple low-rank adaptation. *arXiv preprint arXiv:2409.20197*, 2024.

Mo Zhou, Keren Ye, Mauricio Delbracio, Peyman Milanfar, Vishal M Patel, and Hossein Talebi. Unires: Universal image restoration for complex degradations. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 13237–13247. IEEE, 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 *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 21747–21758. IEEE, 2023.

A Experimental Details

A.1 Datasets

Composite degradation setting. For the composite degradation setting, we train and evaluate our model on the CDD-11 benchmark (Guo et al., 2024), which covers eleven degradation tasks constructed from four elementary factors: rain, haze, snow, and low-light. The training set contains 1,183 image pairs per task (13,013 pairs in total), where each degraded image is paired with its clean counterpart. During training, we randomly crop paired patches of 256×256 from the degraded and clean images with spatial alignment, and apply random horizontal/vertical flipping and 90° rotation. The test set contains 200 images per task (2,200 in total), and we evaluate on the full-resolution images without cropping.

Multi-single-degradation setting. For the multi-single-degradation setting, we train on the all-in-one training set used by prior methods, which pools 5,144 denoising pairs from BSD400(Arbelaez et al., 2010) and WED (Ma et al., 2016), 200 deraining pairs from Rain100L (Yang et al., 2019), 72,135 hazy images from SOTS (Li et al., 2018), 2,103 blurry pairs from GoPro (Nah et al., 2017), and 485 low-light pairs from LOL-v1 (Wei et al., 2018). Since the five tasks differ substantially in dataset size, we balance their contributions by repeating each task's pool a fixed number of times per epoch. We use repetition factors of 3, 120, 1, 1, and 200 for denoising, deraining, dehazing, deblurring, and low-light enhancement, respectively, yielding 210,670 samples per epoch. We apply the same data augmentation as in CDD-11 training, including 256×256 paired random cropping, flipping, and rotation. For evaluation, we use BSD68 for denoising, Rain100L (Yang et al., 2019) for deraining, SOTS (Li et al., 2018) for dehazing, GoPro (Nah et al., 2017) for deblurring, and LOL-v1 (Wei et al., 2018) for low-light enhancement. Gaussian noise with $\sigma = 25$ is synthesized from the clean BSD68 images during testing. All test images are processed at their original resolution.

A.2 Implementation Details and Training Protocol

A.2.1 Implementation Details

Composite degradation setting. The full model adopts a U-Net architecture with a channel width of 48 and transformer blocks as the basic building units. The encoder, bottleneck, and decoder contain $[4, 6, 6]$, 8, and $[6, 6, 4]$ blocks, respectively, with $[1, 2, 4]$ attention heads in the encoder stages and 8 heads in the bottleneck. The decoder integrates transformer blocks with the proposed factor-specific restoration adaptation modules and feed-forward networks for degradation-aware restoration. Each factor-specific restoration adaptation employs a low-rank branch with rank 8. The LoRA residual scale in Eq. (5) is set to $\gamma = 8$. The logit scale in Eq. (3) is computed as $\tau = \min(\exp(t_\tau), 100)$, where t_τ is a learnable scalar initialized to zero, giving an initial scale of $\tau = 1$. The feed-forward expansion factor is 2.66. The degradation factor organizer produces a 384-dimensional factor representation for four elementary degradation factors. The factor matching module estimates four factor composition coefficients that jointly represent the eleven degradation conditions in CDD-11. The full model contains 29.7M parameters in total.

Multi-single-degradation setting. We use the same backbone architecture with five degradation factors corresponding to denoising, deraining, dehazing, deblurring, and low-light enhancement. Each input has a single ground-truth degradation label. The matching head predicts a softmax distribution over the five factors, whose probabilities serve as the predicted composition coefficients for

the decoder. The matching-guided low-rank adaptation mechanism otherwise remains the same as in the composite setting.

A.2.2 Training Protocol

Composite degradation setting. We train the model with the AdamW optimizer (β_1=0.9, β_2=0.999, weight decay 10^{-3}) for 1.2M iterations with a batch size of 6 on 2 GPUs. The learning rate starts from 10^{-3} with 1,000 warm-up iterations and follows a cosine annealing schedule down to 10^{-7}, with gradient clipping at norm 0.5. The total loss is defined as

$$\mathcal{L} = \mathcal{L}_{\text{rec}} + \lambda_1 \mathcal{L}_{\text{fft}} + \lambda_2 \mathcal{L}_{\text{fac}}, \tag{8}$$

where $\mathcal{L}_{\text{rec}}$ denotes the ℓ_1 reconstruction loss, $\mathcal{L}_{\text{fft}}$ the frequency-domain loss, and $\mathcal{L}_{\text{fac}}$ the factor-matching supervision. For the composite setting, λ_1=0.125, λ_2=1, and $\mathcal{L}_{\text{fac}}$ is a multi-label binary cross-entropy loss over the factor coefficients,

$$\mathcal{L}_{\text{fac}} = -\frac{1}{K}\sum_{k=1}^{K}\Big[c_k \log \sigma(\ell_k) + (1-c_k)\log\big(1-\sigma(\ell_k)\big)\Big], \tag{9}$$

where K=4, c_k denotes the ground-truth coefficient of the k-th factor, and ℓ_k the predicted logit. The two supervision paths are deliberately separated. $\mathcal{L}_{\text{fac}}$ is applied only to a detached degradation-readout branch, so that it is dedicated to learning factor decomposition without disturbing the restoration features, while $\mathcal{L}_{\text{rec}}$ and $\mathcal{L}_{\text{fft}}$ are dedicated to restoration quality.

We adopt a matching curriculum. Throughout training, $\mathcal{L}_{\text{fac}}$ remains supervised by ground-truth factor labels. For the first 320K iterations, the decoder uses ground-truth factor coefficients. These coefficients are then gradually mixed with predicted coefficients until iteration 800K, after which the decoder uses only predicted coefficients. The initial ground-truth guidance stabilizes restoration training by preventing unreliable early factor predictions from misleading the composition of factor-specific adaptations.

Multi-single-degradation setting. We use the same optimizer and learning-rate schedule and train for 600K iterations with a batch size of 6 on 2 GPUs. The objective combines the same ℓ_1 and frequency-domain losses with a cross-entropy task-classification loss with weight 0.1 for the five-way factor matching. The matching curriculum uses ground-truth labels for the first 75K iterations and mixes in predictions until iteration 225K. In this setting, the matching head is a five-way classifier. Its softmax outputs define the predicted composition coefficients for the five factors, replacing the independent sigmoid outputs used in the composite setting. The decoder uses these predictions according to the curriculum described above and exclusively at inference. The classification loss remains supervised by ground-truth task labels throughout training.

B Representation Analysis Protocol

All analyses are conducted on features extracted on the CDD-11 test set (200 images per task, 2,200 in total). For each image, we record the global spatial average of the bottleneck features of the backbone trained under the composite degradation setting, denoted $\boldsymbol{z}$, and the factor representation $\boldsymbol{d}$ produced by DFO on top of this bottleneck. Both $\boldsymbol{z}$ and $\boldsymbol{d}$ are ℓ_2-normalized before all measurements, and every classifier or probe is fit on a random 70% split (seed 42) and evaluated on the remaining 30%.

AUC Measurement. For each of the four degradation factors, we fit a least-squares linear probe with a bias term to predict its binary presence label. The eleven degradation-condition labels are converted into four-dimensional multi-hot labels according to their constituent factors. We compute the rank-based area under the ROC curve (AUC) on the held-out split and report the mean across the four factors. The bottleneck representation $\boldsymbol{z}$ achieves a mean per-factor AUC of 0.997, indicating that degradation-factor information is already readily accessible before explicit organization by DFO.

Factor Distinguishability Measurement. We compute the mean normalized feature for each of the four single-degradation conditions and measure the pairwise cosine similarities between these means. We report the maximum off-diagonal cosine similarity in Table 4 to quantify the strongest alignment between distinct factor directions. Lower values indicate reduced cross-factor alignment.

11-Way Degradation Classification. We compute the class centroid of each of the eleven degradation conditions as the mean of the normalized training-split features, and classify each held-out sample by its nearest centroid under cosine similarity. We report the top-1 classification accuracy over the eleven conditions in Table 4.

Composition Exact-Match Accuracy. The DFO-predicted coefficients are binarized using $\hat{y}_k = \mathbb{I}[\alpha_k \geq 0.5]$, where α_k is the predicted coefficient for the k-th degradation factor, $\mathbb{I}[\cdot]$ is the indicator function that returns 1 when the condition holds and 0 otherwise, and $\hat{y}_k$ is the resulting binary factor-presence prediction. These predictions are compared against the ground-truth factor-presence labels. An image is counted as correct only when all four factor predictions match their corresponding labels. We compute this accuracy over all 2,200 CDD-11 test images, without fitting an additional classifier. The results are reported in Table 4.

Subspace Residual Measurement. For each degradation condition, we center its normalized features and compute the top $k = 10$ principal components as its subspace basis. For a composite condition, we concatenate the bases of its constituent single-degradation conditions and apply QR orthogonalization. We compute the ratio of the squared projection residual to the squared norm of each centered composite feature and average it over samples within that condition. The results are presented in Fig. 6, where lower residuals indicate that the constituent subspaces capture more of the composite feature variation.

C Additional Visual Results

C.1 Visual Results on Composite Degradation Restoration

We present additional qualitative results on CDD-11 under seven combinations of degradations, including haze and rain in Fig. 7, haze and snow in Fig. 8, low-light and haze in Fig. 9, low-light and rain in Fig. 12, low-light and snow in Fig. 13, low-light with haze and rain in Fig. 10, and low-light with haze and snow in Fig. 11. These results illustrate the restoration performance of our method across diverse scenes and degradation combinations.

C.2 Visual Results on Multi-Single-Degradation Restoration

We present additional qualitative results in the multi-single-degradation setting, including deblurring in Fig. 14, dehazing in Fig. 15, denoising in Fig. 16, deraining in Fig. 17, and low-light enhancement in Fig. 18. These results illustrate the restoration performance of our method across diverse scenes and degradation types.

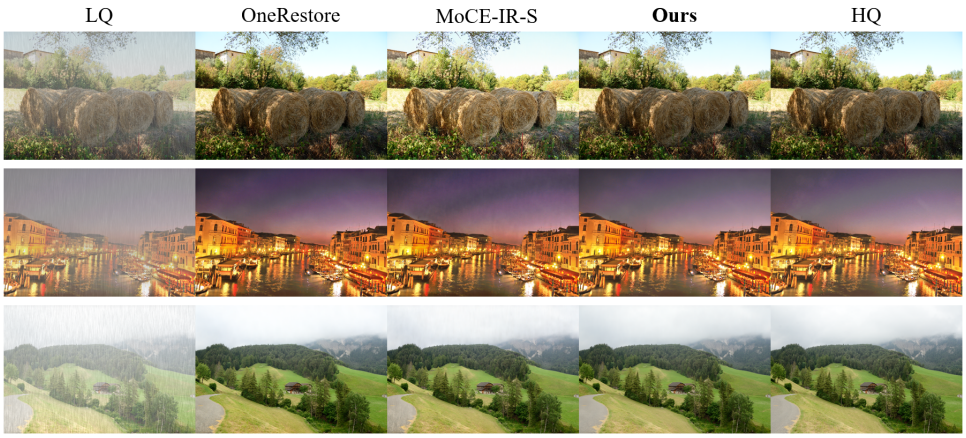

Figure 7: Qualitative comparisons under haze and rain degradation on CDD-11.

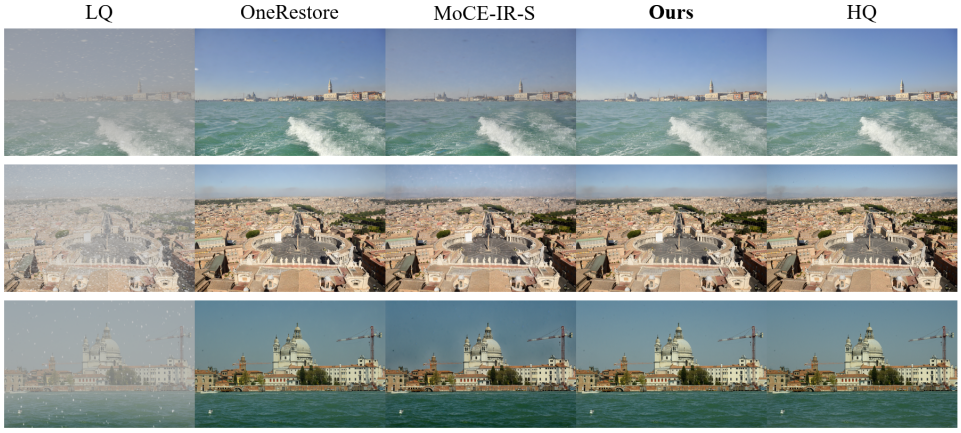

Figure 8: Qualitative comparisons under haze and snow degradation on CDD-11.

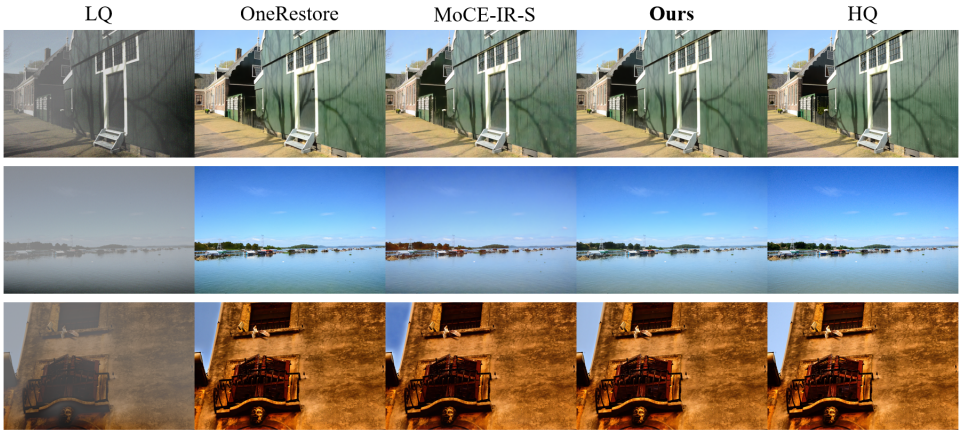

Figure 9: Qualitative comparisons under low-light and haze degradation on CDD-11.



Figure 10: Qualitative comparisons under low-light, haze, and rain degradation on CDD-11.

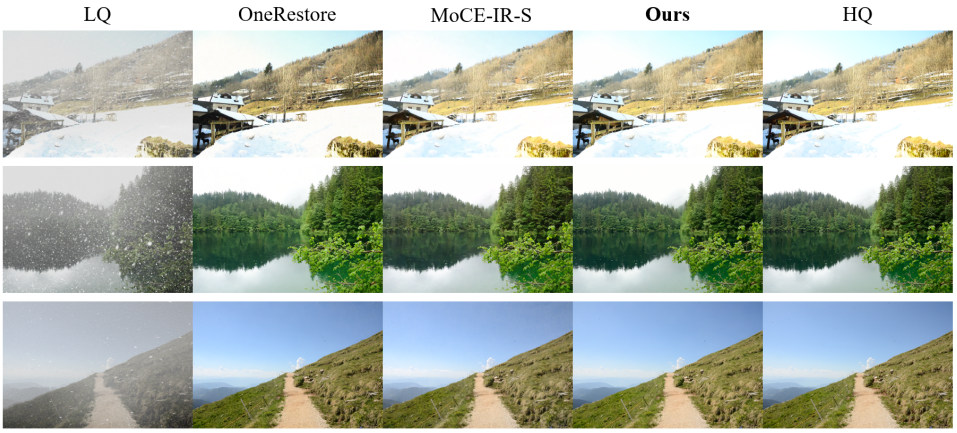

Figure 11: Qualitative comparisons under low-light, haze, and snow degradation on CDD-11.

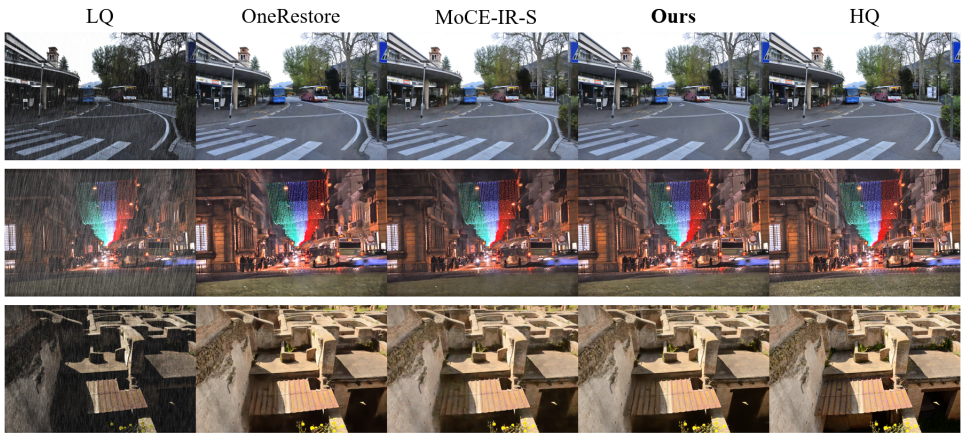

Figure 12: Qualitative comparisons under low-light and rain degradation on CDD-11.

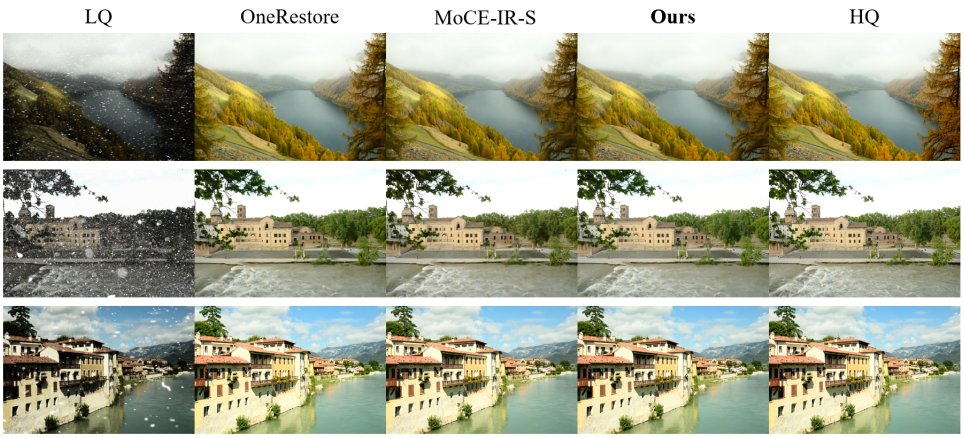

Figure 13: Qualitative comparisons under low-light and snow degradation on CDD-11.

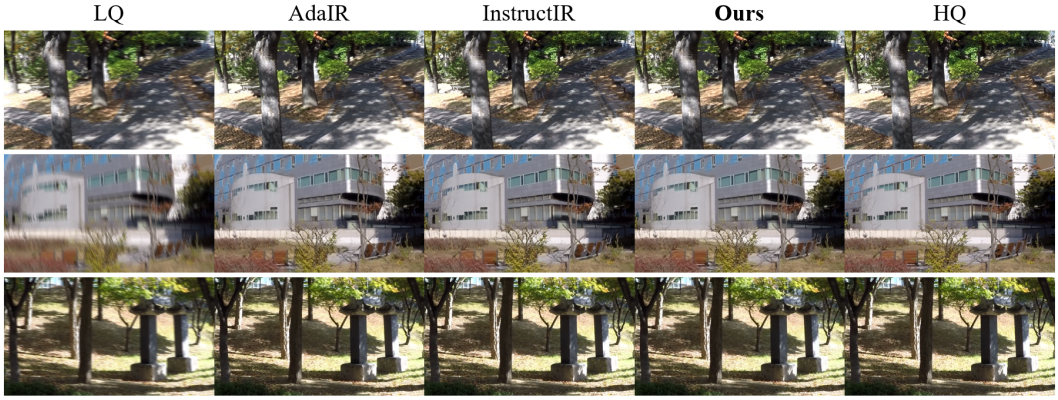

Figure 14: Qualitative comparisons on image deblurring.



Figure 15: Qualitative comparisons on image dehazing.

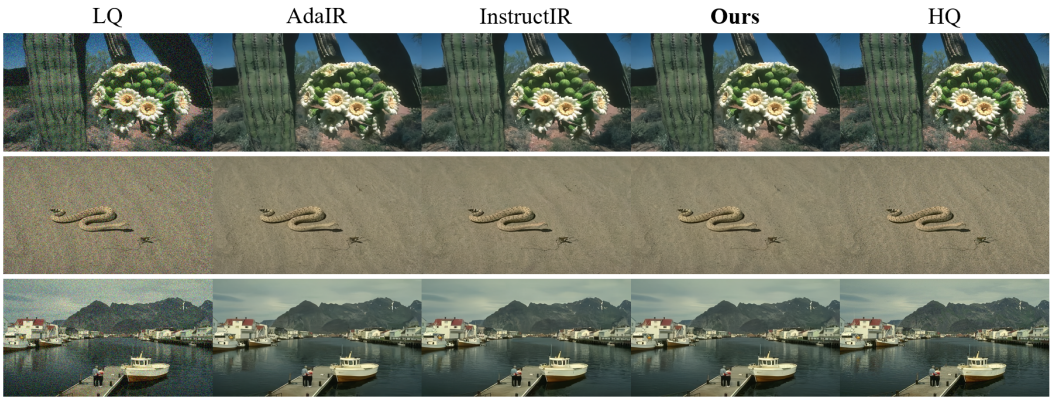

Figure 16: Qualitative comparisons on image denoising.

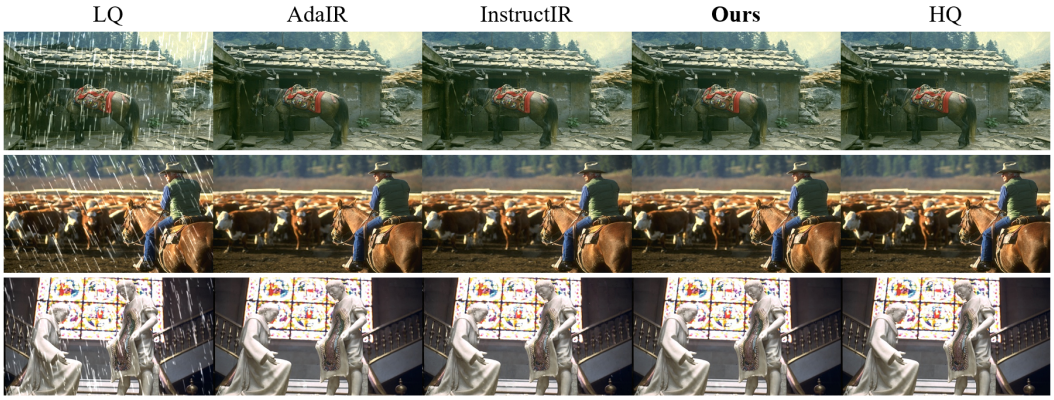

Figure 17: Qualitative comparisons on image deraining.

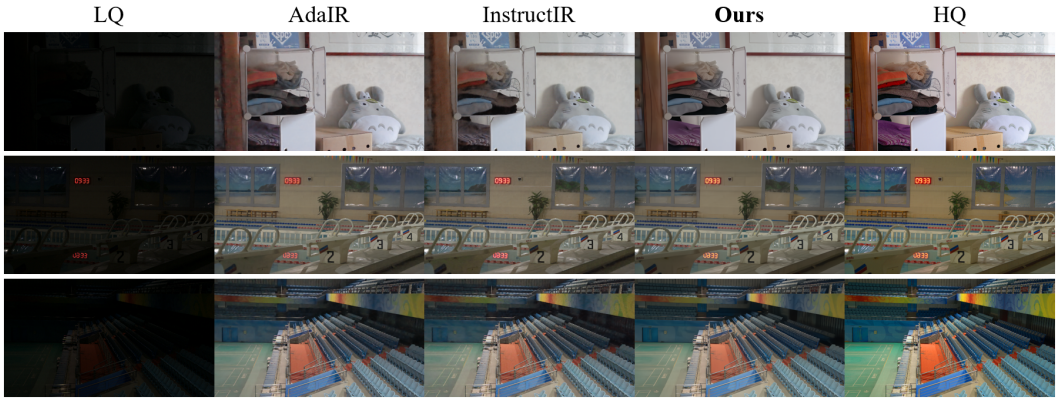

Figure 18: Qualitative comparisons on low-light image enhancement.