REDU-IR: RECONSTRUCTIBLE FEATURE DECOMPOSITION AND STATE COMPLEMENTARY UPDATING FOR ALL-IN-ONE BLIND IMAGE RESTORATION

Anonymous authors
Paper under double-blind review

ABSTRACT

All-in-One Image Restoration (AiOIR) aims to recover high-quality images from inputs corrupted by heterogeneous degradations using a single model. In hierarchical restoration networks, downsampling may discard information needed for reconstruction. Standard skip connections typically pass encoder information directly to the corresponding decoder stage without updating it based on the restored main state. To address these issues, we propose ReDU-IR. It combines Reconstructible Feature Decomposition (RFD) with State Complementary Updating (SCU). RFD decomposes encoder features into main and complementary representations. The matched inverse transform reconstructs features at the corresponding decoder stage. RFD preserves the information needed for reconstruction in the paired main and complementary representations. SCU applies a shared mapping to the encoder and decoder main states. It uses the difference between the two mapping outputs to update the complementary representation without requiring degradation identification. Extensive experiments on multiple datasets demonstrate the effectiveness of our proposed method and its individual components.

1 INTRODUCTION

All-in-One Image Restoration (AiOIR) aims to recover high-quality images from inputs with diverse degradations using a single model (Li et al., 2022; Potlapalli et al., 2023). Noise, rain, haze, blur, and low light alter images through different processes. Therefore, a shared model must adapt its restoration behavior to the degradation present in each input image.

Recent methods mainly achieve such input-adaptive restoration through degradation representations, conditional modulation, and expert routing. Degradation representations and prompts modulate restoration features (Li et al., 2022; Zhang et al., 2023; Potlapalli et al., 2023). Degradation cues also modify attention computation or frequency processing (Ren et al., 2026b; Cui et al., 2025). Other approaches route each input to a different expert network (Zamfir et al., 2025; Lin et al., 2026). These methods mainly address how a shared model adapts its restoration behavior to each input. We further examine cross-scale feature reconstruction in hierarchical networks. We examine how to preserve reconstruction information during scale transformations and how to update skip-connection features using corresponding encoder and decoder features.

Prior work provides a foundation for addressing these issues. MWCNN uses wavelet analysis and synthesis for resolution changes, while invertible restoration networks organize information for subsequent reconstruction (Liu et al., 2018; 2021). WINNet further combines a reconstructible transform with detail-coefficient denoising (Huang & Dragotti, 2022). URWKV explores state-aware selective fusion, while SymUNet studies encoder feature reuse through symmetric scale alignment (Xu et al., 2025; Jiao et al., 2025). We build on these studies by combining reconstructible scale transformations with updates to skip-connection features. A shared mapping processes corresponding encoder and decoder features, and the difference between its outputs determines each update.

We formulate cross-scale feature reuse as state-coupled reconstruction, where the complementary representation is defined relative to the main restoration state. Based on this formulation, we propose ReDU-IR, which combines Reconstructible Feature Decomposition (RFD) and State Complementary Updating (SCU). RFD addresses information preservation by splitting encoder features into

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.

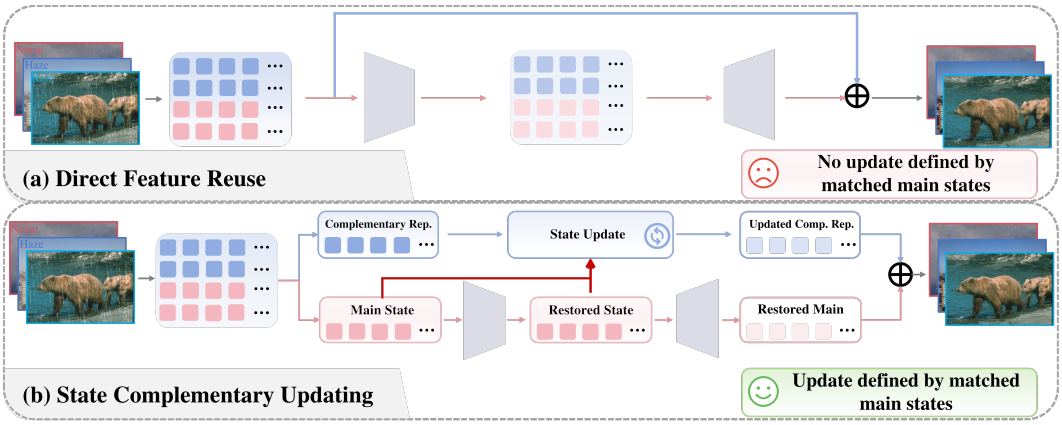

Figure 1: Conceptual comparison of feature reconstruction at matched encoder and decoder stages. (a) Direct reuse passes the retained encoder-side representation unchanged to the matched decoder stage. (b) State complementary updating uses matched encoder and decoder main states to update the complementary representation before feature reconstruction.

main and complementary representations. The main representation enters the deeper restoration pathway, while the complementary representation retains the remaining information needed for reconstruction. A matched inverse transform recombines the two representations at the corresponding decoder stage. SCU updates the complementary representation before this recombination. It applies the same mapping to the main representation before and after restoration, then adds the difference between the mapping outputs to the complementary representation. This update uses the change in the main representation without requiring degradation identification. Fig. 1 compares this process with direct reuse, which passes the complementary representation to the decoder unchanged.

We summarize our contributions as follows:

- We formulate cross-scale feature reconstruction in hierarchical networks as a pairing of main and complementary representations. The complementary representation retains reconstruction information not carried by the main pathway. Before decoder recombination, matched encoder- and decoder-side main states guide its update.
- We propose ReDU-IR, which consists of RFD and SCU. RFD performs a reconstructible main-complementary decomposition of encoder features. It completes feature reconstruction at the corresponding decoder stage through a matched inverse transform. SCU updates the complementary representation using the difference between shared-mapping outputs for the corresponding encoder and decoder main states.
- We evaluate ReDU-IR on three-degradation (3D), five-degradation (5D), CDD11, and WeatherBench restoration settings. Comparisons show gains in average PSNR, and ablations measure the contributions of RFD and SCU.

2 Related Work

All-in-One Image Restoration. All-in-One image restoration aims to handle different degradations with a single model. Early approaches mainly use degradation-aware representations or conditioning mechanisms. AirNet (Li et al., 2022) learns discriminative degradation representations through contrastive learning, IDR (Zhang et al., 2023) decomposes degradations into constituent ingredients, and PromptIR (Potlapalli et al., 2023) introduces learnable prompts for degradation-conditioned restoration. Subsequent methods explore neural degradation representations in NDR (Yao et al., 2024), quality perception in Perceive-IR (Zhang et al., 2026a), and frequency priors in AdaIR (Cui et al., 2025). Other approaches use expert routing in MoCE-IR (Zamfir et al., 2025), feature perturbation in DFPIR (Tian et al., 2025), and residual conditioning in DA-RCOT (Tang et al., 2025). More recent designs investigate efficient spatial–frequency degradation adaptation in AnyIR (Ren et al., 2026c), evolutionary frequency modulation in EvoIR (Ma et al., 2025), and deformable sliding-window attention in DSwinIR (Wu et al., 2026). BaryIR (Tang et al., 2026) instead constructs a Wasserstein barycenter space to improve generalization to unseen degra-

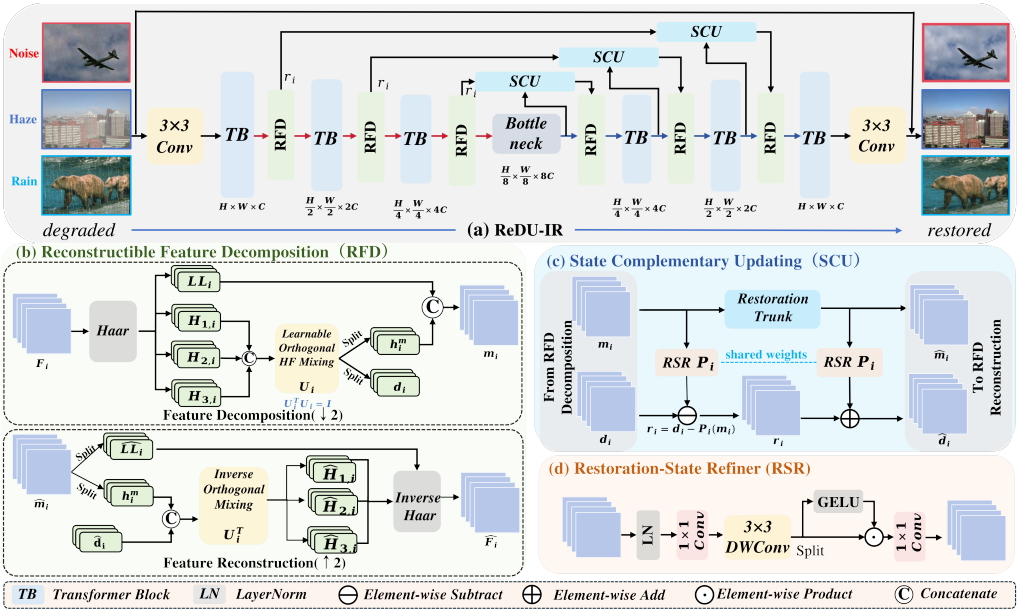

Figure 2: Overview of ReDU-IR. (a) Overall encoder–decoder framework. (b) Feature decomposition and reconstruction in RFD. (c) Complementary representation updating in SCU. (d) Structure of the Restoration-State Refiner (RSR).

dations. Recent works such as MIRAGE (Ren et al., 2026a), UniRestorer (Lin et al., 2026), and CWP-Net (Wang et al., 2026) further improve degradation-agnostic representation learning, adaptive expert selection, and degradation–semantic disentanglement, respectively. These methods primarily investigate how a shared model identifies degradation characteristics or adapts restoration strategies. We examine how to update the complementary representation after restoration changes the corresponding main representation in a hierarchical network.

Reconstructible Representations and Feature Reuse. Hierarchical restoration networks commonly preserve fine details through encoder–decoder skip connections, as introduced by the symmetric encoder–decoder architecture (Mao et al., 2016) and widely adopted in restoration networks such as Restormer (Zamir et al., 2022). Another line of research aims to reduce information loss during hierarchical representation transformations. MWCNN (Liu et al., 2018) uses wavelet and inverse-wavelet transforms for resolution changes, while InvDN (Liu et al., 2021) uses invertible mappings to separate latent components for reconstruction. AdaptIR (Cui et al., 2026) further incorporates wavelet-domain processing for frequency-aware restoration and receptive-field expansion. Among invertible restoration approaches, WINNet (Huang & Dragotti, 2022) constructs a lifting-inspired reconstructible multi-scale transform and restores images by denoising transformed detail coefficients. These approaches primarily preserve or reconstruct information across scales. URWKV (Xu et al., 2025) revisits reuse of encoder representations through state-aware selective fusion, while SymUNet (Jiao et al., 2025) reports that symmetric scale alignment supports additive skip fusion. Together, prior work offers two perspectives on hierarchical restoration: information-preserving transformations and state-aware reuse of encoder representations. ReDU-IR connects these perspectives by formulating restoration as coordination between paired main and complementary representations.

3 Method

3.1 Overall Architecture

Fig. 2 shows the overall architecture of ReDU-IR, which adopts a symmetric multiscale encoder–decoder with Transformer blocks. At each encoder scale, RFD decomposes features into main and complementary representations, and the main representation enters the deeper restoration pathway. At the corresponding decoder scale, SCU updates the complementary representation using the paired encoder and decoder main representations. RFD then reconstructs features from the updated com-

Algorithm 1 Single-scale RFD decomposition, SCU update, and RFD reconstruction

Require: Encoder feature $\mathbf{F}_i \in \mathbb{R}^{B\times C_i\times H_i\times W_i}$ with even H_i, W_i
Require: Matched restored main state $\hat{\mathbf{m}}_i \in \mathbb{R}^{B\times 2C_i\times H_i/2\times W_i/2}$
Require: Orthogonal $\mathbf{U}_i \in \mathbb{R}^{3C_i\times 3C_i}$; shared mapping P_i
Ensure: Encoder states $(\mathbf{m}_i, \mathbf{r}_i)$ and $\hat{\mathbf{F}}_i \in \mathbb{R}^{B\times C_i\times H_i\times W_i}$

[Encoder] RFD decomposition and SCU encoding
$(\mathbf{LL}_i, \mathbf{H}_{1,i}, \mathbf{H}_{2,i}, \mathbf{H}_{3,i}) \leftarrow \mathcal{H}(\mathbf{F}_i)$
$[\mathbf{h}_i^m; \mathbf{d}_i] \leftarrow \mathrm{Split}_{C_i,2C_i}(\mathbf{U}_i[\mathbf{H}_{1,i}; \mathbf{H}_{2,i}; \mathbf{H}_{3,i}])$
$\mathbf{m}_i \leftarrow [\mathbf{LL}_i; \mathbf{h}_i^m]$
$\mathbf{r}_i \leftarrow \mathbf{d}_i - P_i(\mathbf{m}_i)$ ▷ Form the state-relative code
After the restoration pathway produces $\hat{\mathbf{m}}_i$*:*
[Decoder] SCU update and RFD reconstruction
$\hat{\mathbf{d}}_i \leftarrow \mathbf{r}_i + P_i(\hat{\mathbf{m}}_i)$ ▷ Reuse the same code
$[\hat{\mathbf{LL}}_i; \hat{\mathbf{h}}_i^m] \leftarrow \mathrm{Split}_{C_i,C_i}(\hat{\mathbf{m}}_i)$
$[\hat{\mathbf{H}}_{1,i}; \hat{\mathbf{H}}_{2,i}; \hat{\mathbf{H}}_{3,i}] \leftarrow \mathbf{U}_i^{\mathsf{T}}[\hat{\mathbf{h}}_i^m; \hat{\mathbf{d}}_i]$
$\hat{\mathbf{F}}_i \leftarrow \mathcal{H}^{-1}(\hat{\mathbf{LL}}_i, \hat{\mathbf{H}}_{1,i}, \hat{\mathbf{H}}_{2,i}, \hat{\mathbf{H}}_{3,i})$
return $(\mathbf{m}_i, \mathbf{r}_i, \hat{\mathbf{F}}_i)$

plementary representation and the restored main representation. The network produces the restored image through feature reconstruction across successive scales.

Algorithm 1 summarizes RFD decomposition, SCU updating, and RFD reconstruction at a single scale. Here, $\mathbf{U}_i$ is a learned orthogonal transform, and P_i is the shared mapping used by SCU. The following subsections define these operations, which ReDU-IR applies at each pair of corresponding encoder and decoder scales.

3.2 Reconstructible Feature Decomposition

RFD separates the information sent through the deeper restoration pathway from the complementary information needed for reconstruction. The decomposition must preserve both parts so that the decoder can recombine them after restoration. We therefore introduce Reconstructible Feature Decomposition (RFD), illustrated in Fig. 2(b), as a paired feature decomposition and reconstruction transform.

Feature decomposition. Let $\mathbf{F}_i \in \mathbb{R}^{B\times C_i\times H_i\times W_i}$ denote the feature entering the i-th RFD. A 2×2 Haar transform first decomposes $\mathbf{F}_i$ into a low-frequency component $\mathbf{LL}_i$ and three high-frequency components $\mathbf{H}_{1,i}$, $\mathbf{H}_{2,i}$, and $\mathbf{H}_{3,i}$. We concatenate the three high-frequency components and apply a learnable orthogonal transform $\mathbf{U}_i$. We then split the transformed coefficients into a main high-frequency component $\mathbf{h}_i^m$ and a complementary representation $\mathbf{d}_i$. Concatenating $\mathbf{h}_i^m$ with $\mathbf{LL}_i$ gives the main representation $\mathbf{m}_i$:

$$\begin{aligned}(\mathbf{LL}_i, \mathbf{H}_{1,i}, \mathbf{H}_{2,i}, \mathbf{H}_{3,i}) &= \mathcal{H}(\mathbf{F}_i),\\ [\mathbf{h}_i^m; \mathbf{d}_i] &= \mathrm{Split}(\mathbf{U}_i[\mathbf{H}_{1,i}; \mathbf{H}_{2,i}; \mathbf{H}_{3,i}]),\\ \mathbf{m}_i &= [\mathbf{LL}_i; \mathbf{h}_i^m],\end{aligned} \tag{1}$$

where $[\cdot;\cdot]$ denotes channel concatenation and $\mathbf{U}_i^{\mathsf{T}}\mathbf{U}_i = \mathbf{I}$. We use a Cayley parameterization to keep $\mathbf{U}_i$ orthogonal during training. Its transpose therefore serves as the corresponding inverse transform, recovering the learned orthogonal mixing. After decomposition, both $\mathbf{m}_i$ and $\mathbf{d}_i$ have spatial resolution $H_i/2\times W_i/2$ and $2C_i$ channels. The main representation $\mathbf{m}_i$ continues through the restoration pathway, while the complementary representation $\mathbf{d}_i$ retains reconstruction information not carried by the main pathway.

The orthogonal mixing lets the network learn which combinations of high-frequency coefficients enter each representation. Its transpose reverses the mixing, so adapting the high-frequency basis preserves paired reconstruction.

Feature reconstruction. At the corresponding decoder stage, RFD receives the restored main representation $\hat{\mathbf{m}}_i$ and the updated complementary representation $\hat{\mathbf{d}}_i$. We split $\hat{\mathbf{m}}_i$ back into its low-

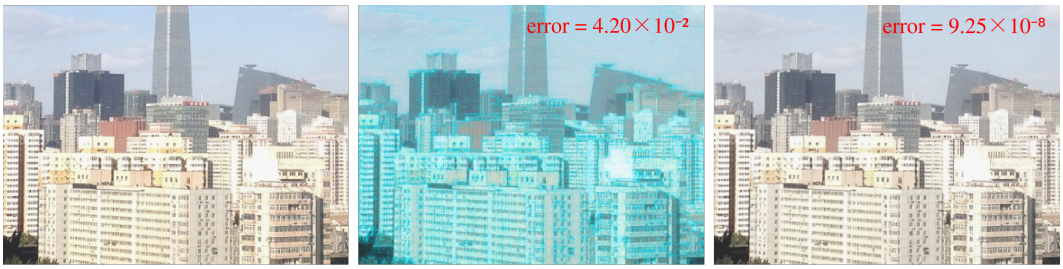

(a) Input (b) Rank-reducing down–up 4C→2C→4C (c) Joint RFD 4C→ $\underbrace{2C}_{m_i}$+$\underbrace{2C}_{d_i}$ →4C

Figure 3: Full-band reconstruction fidelity. (a) Input. (b) A rank-reducing down–up transition incurs visible distortion ($\epsilon_{\text{rel}} = 4.20 \times 10^{-2}$). (c) With both RFD branches unchanged, the inverse transform recovers the input up to 32-bit floating-point round-off ($\epsilon_{\text{rel}} = 9.25 \times 10^{-8}$). Cyan indicates the root-mean-square reconstruction error across channels at each pixel, using a shared color scale.

frequency and main high-frequency parts, recombine the latter with $\hat{\mathbf{d}}_i$, and apply the transpose of the same orthogonal transformation. This produces three updated high-frequency Haar components. We combine them with the updated low-frequency component through the inverse Haar transform:

$$\begin{aligned} [\hat{\mathbf{L}}\hat{\mathbf{L}}_i; \hat{\mathbf{h}}_i^m] &= \text{Split}(\hat{\mathbf{m}}_i), \\ [\hat{\mathbf{H}}_{1,i}; \hat{\mathbf{H}}_{2,i}; \hat{\mathbf{H}}_{3,i}] &= \mathbf{U}_i^{\mathsf{T}}[\hat{\mathbf{h}}_i^m; \hat{\mathbf{d}}_i], \\ \hat{\mathbf{F}}_i &= \mathcal{H}^{-1}(\hat{\mathbf{L}}\hat{\mathbf{L}}_i, \hat{\mathbf{H}}_{1,i}, \hat{\mathbf{H}}_{2,i}, \hat{\mathbf{H}}_{3,i}). \end{aligned} \tag{2}$$

Reconstruction fidelity. RFD retains the information needed for reconstruction in the paired main and complementary representations. To assess information preservation across scales, we reconstruct the input directly from these representations without intervening feature restoration or complementary updates. Fig. 3 compares the reconstruction errors over three successive scale transitions using a shared error scale. Retaining only $2C_i$ of the $4C_i$ coefficients associated with each 2×2 spatial block produces visible reconstruction errors. With both $\mathbf{m}_i$ and $\mathbf{d}_i$, the relative reconstruction error is 9.25×10^{-8} for the illustrated input. This comparison illustrates the role of the complementary representation in retaining information needed for reconstruction. In ReDU-IR, SCU updates this representation according to changes in the main representation before paired feature reconstruction.

3.3 State Complementary Updating

RFD produces a main representation $\mathbf{m}_i$ and a complementary representation $\mathbf{d}_i$ at each encoder scale. The main representation passes through the deeper restoration pathway to yield the restored main representation $\hat{\mathbf{m}}_i$. Under direct reuse, $\mathbf{d}_i$ remains unchanged and participates in feature reconstruction together with $\hat{\mathbf{m}}_i$ at the corresponding decoder scale. SCU uses the change from $\mathbf{m}_i$ to $\hat{\mathbf{m}}_i$ to adapt the complementary representation before reconstruction. As shown in Fig. 2(c), it applies a shared mapping to the corresponding encoder and decoder main states. SCU uses the difference between the two mapping responses to update $\mathbf{d}_i$, yielding $\hat{\mathbf{d}}_i$ for paired reconstruction with $\hat{\mathbf{m}}_i$.

Restoration-State Refiner. A lightweight Restoration-State Refiner (RSR) implements the shared mapping $P_i(\cdot)$ from a main state to the complementary feature space. As illustrated in Fig. 2(d), RSR applies LayerNorm, a 1×1 projection, and a 3×3 depth-wise convolution. It splits the resulting feature into gate and value branches, applies the Gaussian error linear unit (GELU) to the gate branch, multiplies the two branches element-wise, and uses a final 1×1 projection:

$$\begin{aligned} [\mathbf{g}_i, \mathbf{v}_i] &= \text{Split}\left(\text{DWConv}_{3\times3}\left(\text{Conv}_{1\times1}^{\text{in}}\left(\text{LN}(\mathbf{m}_i)\right)\right)\right), \\ P_i(\mathbf{m}_i) &= \text{Conv}_{1\times1}^{\text{out}}\left(\text{GELU}(\mathbf{g}_i) \odot \mathbf{v}_i\right), \end{aligned} \tag{3}$$

where $\odot$ denotes element-wise multiplication. RSR takes only the main representation as input. The encoder and decoder evaluations at the same scale share the parameters of P_i. The difference between their outputs therefore reflects a change in the input state under the same mapping.

State Complementary Updating. We express a residual update conditioned on the corresponding main states as

$$\hat{\mathbf{d}}_i = \mathbf{d}_i + G_i(\mathbf{m}_i, \hat{\mathbf{m}}_i),$$

where G_i predicts a complementary correction from the encoder- and decoder-side main states. Direct reuse is the special case $G_i = 0$.

SCU constrains this residual updater using a single shared mapping. Instead of transmitting the raw complementary representation $\mathbf{d}_i$, the encoder side subtracts the mapping response at the current main state and stores the result as a state-relative complementary code. At the matched decoder stage, the same RSR maps the restored main representation $\hat{\mathbf{m}}_i$, and SCU adds the resulting response to the same code. Eliminating the stored code gives the equivalent response-difference form:

$$\begin{aligned} \mathbf{r}_i &= \mathbf{d}_i - P_i(\mathbf{m}_i), \\ \hat{\mathbf{d}}_i &= \mathbf{r}_i + P_i(\hat{\mathbf{m}}_i), \\ \hat{\mathbf{d}}_i - \mathbf{d}_i &= P_i(\hat{\mathbf{m}}_i) - P_i(\mathbf{m}_i). \end{aligned} \tag{4}$$

Structural properties of SCU. Equation 4 parameterizes the complementary update as the shared-mapping response difference $G_i^{\mathrm{SCU}}(\mathbf{a}, \mathbf{b}) = P_i(\mathbf{b}) - P_i(\mathbf{a})$. For arbitrary states $\mathbf{a}, \mathbf{b}, \mathbf{c}$, the above difference formulation has zero self-difference and satisfies antisymmetry and path additivity:

$$\begin{aligned} &G_i^{\mathrm{SCU}}(\mathbf{a}, \mathbf{a}) = 0, \\ &G_i^{\mathrm{SCU}}(\mathbf{a}, \mathbf{b}) = -G_i^{\mathrm{SCU}}(\mathbf{b}, \mathbf{a}), \\ &G_i^{\mathrm{SCU}}(\mathbf{a}, \mathbf{c}) = G_i^{\mathrm{SCU}}(\mathbf{a}, \mathbf{b}) + G_i^{\mathrm{SCU}}(\mathbf{b}, \mathbf{c}). \end{aligned} \tag{5}$$

The third identity follows because intermediate responses cancel across successive state transitions. If P_i is L_i-Lipschitz, the update further satisfies $\|G_i^{\mathrm{SCU}}(\mathbf{a}, \mathbf{b})\| \leq L_i \|\mathbf{b} - \mathbf{a}\|$. Under this condition, the norm of the complementary update is at most L_i times the norm of the main-state change. These properties describe the constraints imposed by SCU; they do not establish improved restoration performance. Section 4 evaluates whether this update improves restoration quality.

RFD and SCU form a paired decomposition–updating–reconstruction pathway. RFD decomposes encoder features into main and complementary representations. SCU uses the shared-mapping response difference between the matched main states to update the complementary representation. RFD then reconstructs the higher-resolution feature from the restored main representation $\hat{\mathbf{m}}_i$ and the updated complementary representation $\hat{\mathbf{d}}_i$. The decoder Transformer blocks continue restoration.

4 Experiments

We evaluate ReDU-IR under four all-in-one restoration settings: 3D, 5D, CDD11, and WeatherBench. The 3D setting covers denoising, dehazing, and deraining; 5D adds deblurring and low-light enhancement. CDD11 tests single and composite degradations, while WeatherBench tests real-world haze, rain, and snow. We report peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) (Wang et al., 2004). The appendix provides datasets, implementation details, evaluation protocols, model configurations, and computational complexity comparisons.

4.1 Comparison with State of the Art

Three-degradation restoration. The results in Tab. 1 show that ReDU-IR-B ranks first in PSNR under all five test conditions and achieves the best average result of 33.25 dB/0.922. At a comparable model size (23M vs. 22M), it outperforms ExDA (Ren et al., 2026b) by 0.29 dB/0.001. ReDU-IR-S achieves 33.07 dB/0.921 with 10M parameters, exceeding ExDA by 0.11 dB while matching its average SSIM.

Five-degradation restoration. Tab. 2 shows that ReDU-IR-B achieves the highest average PSNR of 31.03 dB within the main comparison group. It exceeds ReDU-IR-S and ExDA (Ren et al., 2026b), which tie at 30.83 dB, by 0.20 dB. ReDU-IR-S reaches this result with 10M parameters,

Table 1: Comparison under the 3D protocol. Unless noted otherwise, all tables report RGB PSNR (dB) and SSIM; Avg. is the arithmetic mean, and "–" denotes unavailable results. **Bold**: best; underlined: second best. ReDU-IR-S/B denote the 10M/23M configurations.

Method	Venue	Params.	Dehazing		Deraining		Denoising on CBSD68						Average	
			SOTS-Outdoor		Rain100L		$\sigma = 15$		$\sigma = 25$		$\sigma = 50$		PSNR	SSIM
BRDNet (2020)	NN'20	–	23.23	.895	27.42	.895	32.26	.898	29.76	.836	26.34	.693	27.80	.843
LPNet (2020)	TNNLS'20	–	20.84	.828	24.88	.784	26.47	.778	24.77	.748	21.26	.552	23.64	.738
FDGAN (2020)	AAAI'20	–	24.71	.929	29.89	.933	30.25	.910	28.81	.868	26.43	.776	28.02	.883
DL (2021)	TPAMI'19	2M	26.92	.931	32.62	.931	33.05	.914	30.41	.861	26.90	.740	29.98	.876
MPRNet (2021)	CVPR'21	16M	25.28	.955	33.57	.954	33.54	.927	30.89	.880	27.56	.779	30.17	.899
Restormer (2022)	CVPR'22	26M	27.78	.958	33.78	.958	33.72	.930	30.67	.865	27.63	.792	30.72	.901
NAFNet (2022)	ECCV'22	17M	24.11	.928	33.64	.956	33.03	.918	30.47	.865	27.12	.754	29.67	.884
AirNet (2022)	CVPR'22	9M	27.94	.962	34.90	.968	33.92	.933	31.26	.888	28.00	.797	31.20	.910
NDR (2024)	TIP'24	28M	25.01	.860	28.62	.848	28.72	.826	27.88	.798	26.18	.720	27.28	.810
PromptIR (2023)	NeurIPS'23	36M	30.58	.974	36.37	.972	33.98	.933	31.31	.888	28.06	.799	32.06	.913
MoCE-IR-S (2025)	CVPR'25	11M	30.98	.979	38.22	.983	34.08	.933	31.42	.888	28.16	.798	32.57	.916
AdaIR (2025)	ICLR'25	29M	31.06	.980	38.64	.983	34.12	.935	31.45	.892	28.19	.802	32.69	.918
MoCE-IR (2025)	CVPR'25	25M	31.34	.979	38.57	.984	34.11	.932	31.45	.888	28.18	.800	32.73	.917
DSwinIR (2026)	TPAMI'26	24M	31.86	.980	37.73	.983	34.12	.933	31.59	.890	28.31	.803	32.72	.917
AnyIR-S (2026c)	TMLR'26	9M	31.85	**.982**	38.56	.983	34.15	.936	31.49	.893	28.24	.806	32.86	.920
ClearAIR (2026b)	AAAI'26	31M	31.08	.981	38.61	.984	34.18	.935	31.50	.891	28.31	.804	32.74	.919
ExDA (2026b)	ICLR'26	22M	31.58	**.982**	39.13	.985	34.22	**.937**	31.56	.895	28.32	.808	32.96	.921
MIRAGE-S (2026a)	ICLR'26	10M	31.86	.981	38.94	.985	34.12	.935	31.46	.891	28.19	.803	32.91	.919
ReDU-IR-S (Ours)	–	10M	32.19	**.982**	39.02	.985	34.23	**.937**	31.59	.895	28.34	.809	33.07	.921
ReDU-IR-B (Ours)	–	23M	**32.47**	**.982**	**39.47**	**.986**	**34.28**	**.937**	**31.64**	**.896**	**28.40**	**.810**	**33.25**	**.922**
Methods with the assistance of vision language, multi-task learning, natural language prompts, or multi-modal control														
InstructIR-3D (2024)	ECCV'24	16M	30.22	.959	37.98	.978	34.15	.933	31.52	.890	28.30	.804	32.43	.913
VLU-Net (2025)	CVPR'25	35M	30.71	.980	38.93	.984	34.13	.935	31.48	.892	28.23	.804	32.70	.919
Perceive-IR (2026a)	TIP'26	42M	30.87	.975	38.29	.980	34.13	.934	31.53	.890	28.31	.804	32.63	.917
DFPIR (2025)	CVPR'25	31M	31.87	.980	38.65	.982	34.14	.935	31.47	.893	28.25	.806	32.88	.919

compared with 22M for ExDA. ReDU-IR-B ranks first in dehazing PSNR and denoising SSIM and second in deraining PSNR and average SSIM; neither ReDU-IR variant leads on deblurring or low-light enhancement.

Composite-degradation restoration. Tab. 3 shows that ReDU-IR-S ranks first in both metrics under all eleven CDD11 conditions and reaches 31.55 dB/0.897 on average. With 10M parameters, it outperforms ExDA (Ren et al., 2026b) by 1.66 dB/0.008; the largest PSNR margin, 4.94 dB, occurs under haze.

Real-world adverse-weather restoration. For a consistent comparison, we recompute the DSwinIR (Wu et al., 2026) scores from its released outputs using the same protocol. As shown in Tab. 4, ReDU-IR-S leads on both metrics under all three conditions, improving over DSwinIR by 0.5618 dB/0.0153 on average. The largest improvement occurs on desnowing, with gains of 1.2359 dB/0.0254.

Appendix C reports results for models trained separately on each restoration task.

Visual comparisons. Appendix I shows examples from the four all-in-one settings and three single-task settings. These examples illustrate degradation removal and the preservation of fine structures and boundaries.

4.2 Ablation and Mechanism Analysis

RFD Design. Tab. 5(a) evaluates decomposition designs without complementary updates. RFD improves PSNR over PixelUnshuffle by 0.22 dB while using 0.34M fewer parameters, indicating that it improves restoration quality without increasing model size. Fixed Haar preserves information for reconstruction, but its detail representation is determined by predefined subbands. Within the same main–complementary structure, RFD learns this allocation, allowing the detail representation entering the main pathway to adapt to the restoration objective while maintaining paired reconstruction. Compared with fixed Haar, RFD raises PSNR from 33.06 to 33.17 dB, achieving the highest PSNR among the tested decomposition designs.

Table 2: Comparison under the 5D protocol. * denotes general image restorers. Ranks refer to the main comparison group; a separate group lists methods using natural-language prompts or multi-task learning.

Method	Venue	Params.	Dehazing	Deraining	Denoising	Deblurring	Low-Light	Average
			SOTS-Outdoor	Rain100L	CBSD68 ($\sigma = 25$)	GoPro	LOL-v1	PSNR/SSIM
NAFNet* (2022)	ECCV'22	22M	25.23/.939	35.56/.967	31.02/.883	26.53/.808	20.49/.809	27.77/.881
DGUNet* (2022)	CVPR'22	17M	24.78/.940	36.62/.971	31.10/.883	27.25/.837	21.87/.823	28.32/.891
SwinIR* (2021)	ICCVW'21	1M	21.50/.891	30.78/.923	30.59/.868	24.52/.773	17.81/.723	25.04/.835
Restormer* (2022)	CVPR'22	26M	24.09/.927	34.81/.962	31.49/.884	27.22/.829	20.41/.806	27.60/.882
MambaIR* (2024a)	ECCV'24	27M	25.81/.944	36.55/.971	31.41/.884	28.61/.875	22.49/.832	28.97/.901
FSNet* (2024a)	TPAMI'23	13M	25.53/.943	36.07/.968	31.33/.883	28.32/.869	22.29/.829	28.71/.898
DL (2021)	TPAMI'19	2M	20.54/.826	21.96/.762	23.09/.745	19.86/.672	19.83/.712	21.05/.743
TransWeather (2022)	CVPR'22	38M	21.32/.885	29.43/.905	29.00/.841	25.12/.757	21.21/.792	25.22/.836
TAPE (2022)	ECCV'22	1M	22.16/.861	29.67/.904	30.18/.855	24.47/.763	18.97/.621	25.09/.801
AirNet (2022)	CVPR'22	9M	21.04/.884	32.98/.951	30.91/.882	24.35/.781	18.18/.735	25.49/.847
IDR (2023)	CVPR'23	15M	25.24/.943	35.63/.965	**31.60**/.887	27.87/.846	21.34/.826	28.34/.893
PromptIR (2023)	NeurIPS'23	36M	26.54/.949	36.37/.970	31.47/.886	28.71/.881	22.68/.832	29.15/.904
MoCE-IR-S (2025)	CVPR'25	11M	31.33/.978	37.21/.978	31.25/.884	28.90/.877	21.68/.851	30.08/.913
AdaIR (2025)	ICLR'25	29M	30.53/.978	38.02/.981	31.35/.889	28.12/.858	23.00/.845	30.20/.910
MoCE-IR (2025)	CVPR'25	25M	30.48/.974	38.04/.982	31.34/.887	**30.05/.899**	23.00/.852	30.58/**.919**
DSwinIR (2026)	TPAMI'26	24M	30.09/.975	37.77/.982	31.34/.885	29.17/.879	22.64/.843	30.20/.913
ClearAIR (2026b)	AAAI'26	31M	30.12/.978	38.20/.982	31.53/.888	29.67/.887	22.83/.846	30.47/.916
AnyIR-S (2026c)	TMLR'26	9M	31.77/**.982**	39.00/.983	31.44/.892	28.52/.867	23.03/.857	30.75/.916
ExDA (2026b)	ICLR'26	22M	31.14/.981	**39.23/.985**	31.52/.893	29.06/.878	23.19/**.859**	30.83/**.919**
MIRAGE-S (2026a)	ICLR'26	10M	31.45/.980	38.92/**.985**	31.41/.892	28.10/.858	**23.59**/.858	30.69/.915
ReDU-IR-S (Ours)	–	10M	32.13/**.982**	38.76/.984	31.50/.893	28.47/.864	23.27/.856	30.83/.916
ReDU-IR-B (Ours)	–	23M	**32.29/.982**	39.10/**.985**	31.56/**.894**	28.75/.871	23.47/.857	**31.03**/.918
Methods with the assistance of natural-language prompts or multi-task learning								
InstructIR-5D (2024)	ECCV'24	16M	27.10/.956	36.84/.973	31.40/.887	29.40/.886	23.00/.836	29.55/.908
Perceive-IR (2026a)	TIP'26	42M	28.19/.964	37.25/.977	31.44/.887	29.46/.886	22.88/.833	29.84/.909
DFPIR (2025)	CVPR'25	31M	31.64/.979	37.62/.978	31.29/.889	28.82/.873	23.82/.843	30.64/.912
DA-RCOT (2025)	TPAMI'25	41M	30.96/.975	37.87/.980	31.23/.888	28.68/.872	23.25/.836	30.40/.910

Table 3: *Comparison on CDD11.* Entries show PSNR (dB ↑)/SSIM (↑). Avg. denotes the arithmetic mean over all eleven conditions.

Method	Params.	Single				Double					Triple		Avg.
		Low (L)	Haze (H)	Rain (R)	Snow (S)	$L+H$	$L+R$	$L+S$	$H+R$	$H+S$	$L+H+R$	$L+H+S$	
AirNet (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 (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 (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 (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/.798
OKNet (2024b)	5M	25.64/.813	28.27/.983	30.98/.947	32.35/.963	24.08/.808	24.54/.773	24.20/.766	25.83/.934	27.14/.949	23.13/.766	23.48/.764	26.33/.861
OneRestore (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 (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
MIRAGE-S (2026a)	10M	27.41/.833	33.12/.992	34.66/.971	35.98/.981	26.55/.828	26.53/.810	26.33/.803	30.32/.965	30.27/.969	25.59/.801	25.86/.799	29.33/.887
ExDA (2026b)	22M	27.60/.837	34.44/.991	35.33/.973	36.85/.981	26.95/.834	26.80/.817	26.73/.809	31.17/.968	31.40/.972	25.59/.801	25.86/.799	29.88/.889
ReDU-IR-S (Ours)	10M	**27.90/.841**	**39.38/.995**	**36.36/.978**	**38.41/.985**	**27.70/.839**	**27.18/.823**	**27.15/.816**	**34.22/.977**	**35.04/.980**	**26.77/.816**	**26.90/.816**	**31.55/.897**

Table 4: Comparison on WeatherBench. We compute metrics on the luminance (Y) channel after cropping a four-pixel border. † ExDA scores and averages follow the original paper at its reported precision.

Method	Venue	Dehaze		Derain		Desnow		Average	
		PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑
AirNet (2022)	CVPR'22	20.9399	0.7054	33.5855	0.9418	22.0594	0.7799	25.5283	0.8090
TransWeather (2022)	CVPR'22	19.7904	0.6800	29.3372	0.9034	24.9631	0.7959	24.6969	0.7931
PromptIR (2023)	NeurIPS'23	21.1132	0.7128	34.5374	0.9443	27.9270	0.8363	27.8592	0.8311
WGWSNet (2023)	CVPR'23	13.7943	0.6027	37.0770	0.9608	20.8051	0.7799	23.8921	0.7811
DiffUIR (2024)	CVPR'24	22.7374	0.7439	35.9333	0.9547	29.5040	0.8696	29.3916	0.8561
MWFormer (2024)	TIP'24	24.4223	0.7462	35.1464	0.9507	29.9783	0.8721	29.8490	0.8563
Histoformer (2024)	ECCV'24	17.6861	0.6691	30.6981	0.9160	25.3939	0.8076	24.5927	0.7976
AdaIR (2025)	ICLR'25	23.0837	0.7311	34.8685	0.9458	28.4436	0.8372	28.7986	0.8380
ExDA (2026b)†	ICLR'26	23.74	0.739	35.86	0.946	29.42	0.868	29.68	0.851
DSwinIR (2026)	TPAMI'26	25.0723	0.7662	38.4960	0.9651	30.1578	0.8729	31.2420	0.8681
ReDU-IR-S (Ours)	–	**25.3381**	**0.7861**	**38.6797**	**0.9657**	**31.3937**	**0.8983**	**31.8038**	**0.8834**

Table 5: Ablation of RFD and SCU in the ReDU-IR-B configuration, with SymUNet (Jiao et al., 2025) as the baseline. All variants are evaluated under the 3D setting using the same training protocol. (a) Comparison of feature decomposition designs without complementary updates. (b) Comparison of complementary update designs with the learned RFD fixed. HF denotes high frequency.

Variant	Module configuration	Params. (M)	PSNR (dB) ↑
(a) RFD ablation			
w/o RFD	SymUNet baseline (PixelUnshuffle)	22.29	32.95
w/o learnable HF mixing	Fixed Haar decomposition	21.51	33.06
RFD	Learnable orthogonal HF mixing	21.95	**33.17**
(b) SCU ablation			
w/o SCU	Direct complementary feature reuse	21.95	33.17
w/o encoder-side RSR	Decoder-only RSR	22.54	33.20
w/o SCU (fusion replacement)	Concatenation-based fusion	22.34	33.18
w/o RSR weight sharing	Separate encoder and decoder RSRs	23.14	33.23
ReDU-IR (full model)	SCU with shared RSR	22.54	**33.25**

SCU Update Design. Tab. 5(b) compares how different update designs use the encoder and decoder main states. Direct reuse leaves the complementary representation unchanged as the main representation undergoes restoration. Decoder-only RSR predicts the correction from the decoder main state, while concatenation-based fusion also uses the complementary representation. SCU instead applies a shared RSR to the encoder and decoder main states and uses their response difference to update the complementary representation. At the same 22.54M parameter count, SCU improves PSNR from 33.20 to 33.25 dB over decoder-only RSR, supporting the use of the encoder main state as a reference for the update. Sharing also saves 0.60M parameters relative to separate RSRs, with PSNR of 33.25 versus 33.23 dB. These results support computing the complementary update from both main states with a shared RSR.

SCU updates improve restoration. We compare disabled and default SCU updates during inference while keeping the trained model parameters fixed (Tab. 6).

Table 6: SCU updates with fixed trained parameters at inference. $\lambda = 0$ disables updates, and $\lambda = 1$ uses the default updates. Values average five test conditions for 3D and eleven for CDD11. Δ denotes the gain from enabling updates.

Setting	Model	$\text{PSNR}_{\lambda=0}$ (dB)	$\text{PSNR}_{\lambda=1}$ (dB)	ΔPSNR (dB)	ΔSSIM
3D	ReDU-IR-B	33.10	33.25	+0.15	+0.0046
CDD11	ReDU-IR-S	31.20	31.55	+0.35	+0.0064

Enabling updates improves average PSNR by 0.15 dB on 3D and 0.35 dB on CDD11. On CDD11, both PSNR and SSIM improve across all eleven single and composite degradation conditions. These results support updating the complementary representation during reconstruction across both evaluation settings. Appendix F provides the full update-strength sweep, per-image analysis, and per-condition CDD11 results.

5 Conclusion

We propose ReDU-IR for all-in-one blind image restoration. RFD decomposes encoder features into main and complementary representations and reconstructs features through a matched inverse transform, preserving reconstruction information during scale transformations. SCU updates the complementary representation using the response difference of a shared mapping applied to the corresponding encoder and decoder main states. This update adapts the complementary representation to the restored main representation before reconstruction without requiring degradation identification. Across the four all-in-one settings, ReDU-IR improves average PSNR over the compared baselines. Ablations support both components, while tests with fixed parameters show that the effect of SCU varies across stages and degradation types.

Ethics Statement

This work does not involve human-participant studies or the collection of personal or sensitive information by the authors. All experiments use existing image-restoration benchmark datasets cited in the paper, and we do not collect or create a new dataset. We follow the applicable licenses and usage terms of the datasets used in our experiments. Model training incurs computational and energy costs; we therefore report the training hardware and model-complexity measurements to support transparent assessment of computational requirements.

Reproducibility Statement

We provide the model architecture and the single-scale RFD–SCU procedure in Fig. 2, Algorithm 1, and Section 3. Section 4 summarizes the evaluation settings and primary comparisons. The appendix documents the datasets, implementation details, evaluation procedures, and model configurations. The appendix also provides single-task, qualitative, and sensitivity-analysis results. Upon acceptance, we will release the training and evaluation code, configuration files, and pretrained checkpoints used to produce the reported results.

USE OF LARGE LANGUAGE MODELS (LLMS)

We used large language models to assist with translating, and polishing parts of the manuscript, as well as LaTeX and figure formatting. We also used them to obtain feedback on experimental design and to assist in interpreting experimental results. The authors take full responsibility for the research methodology, scientific claims, experimental results, and all content of the final manuscript, including material prepared with AI assistance.

References

Pablo Arbeláez, 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, 2011.

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.

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, 2024a. doi: 10.1109/TPAMI.2023.3330416.

Yuning Cui, Wenqi Ren, and Alois Knoll. Omni-kernel network for image restoration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 1426–1434, 2024b.

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 *The Thirteenth International Conference on Learning Representations*, 2025.

Yuning Cui, Wenqi Ren, and Alois Knoll. Efficient all-in-one image restoration with adaptive frequency enhancement. *IEEE Transactions on Image Processing*, 35:8080–8092, 2026. doi: 10.1109/TIP.2026.3715328.

Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Color image denoising via sparse 3D collaborative filtering with grouping constraint in luminance-chrominance space. In *2007 IEEE International Conference on Image Processing*, pp. I–313–I–316, 2007. doi: 10.1109/ICIP.2007.4378954.

Yu Dong, Yihao Liu, He Zhang, Shifeng Chen, and Yu Qiao. FD-GAN: Generative adversarial networks with fusion-discriminator for single image dehazing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 10729–10736, 2020.

Qingnan Fan, Dongdong Chen, Lu Yuan, Gang Hua, Nenghai Yu, and Baoquan Chen. A general decoupled learning framework for parameterized image operators. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(1):33–47, 2021. doi: 10.1109/TPAMI.2019.2925793.

Xueyang Fu, Borong Liang, Yue Huang, Xinghao Ding, and John Paisley. Lightweight pyramid networks for image deraining. *IEEE Transactions on Neural Networks and Learning Systems*, 31 (6):1794–1807, 2020. doi: 10.1109/TNNLS.2019.2926481.

Qiyuan Guan, Qianfeng Yang, Xiang Chen, Tianyu Song, Guiyue Jin, and Jiyu Jin. WeatherBench: A real-world benchmark dataset for all-in-one adverse weather image restoration. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pp. 12607–12613, 2025. doi: 10.1145/3746027.3758196.

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. doi: 10.1007/978-3-031-72649-1_13.

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

Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5197–5206. IEEE, 2015.

Jun-Jie Huang and Pier Luigi Dragotti. WINNet: Wavelet-inspired invertible network for image denoising. *IEEE Transactions on Image Processing*, 31:4377–4392, 2022.

Wenlong Jiao, Heyang Lee, Ping Wang, Pengfei Zhu, Qinghua Hu, and Dongwei Ren. Unleashing degradation-carrying features in symmetric U-Net: Simpler and stronger baselines for all-in-one image restoration. *arXiv preprint arXiv:2512.10581*, 2025.

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, 2019.

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.

Jingyun Liang, Jiezhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. SwinIR: Image restoration using Swin Transformer. In *2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, pp. 1833–1844. IEEE, 2021.

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 *The Fourteenth International Conference on Learning Representations*, 2026.

Lin Liu, Lingxi Xie, Xiaopeng Zhang, Shanxin Yuan, Xiangyu Chen, Wengang Zhou, Houqiang Li, and Qi Tian. TAPE: Task-agnostic prior embedding for image restoration. In *European Conference on Computer Vision*, pp. 447–464. Springer, 2022. doi: 10.1007/978-3-031-19797-0_26.

Pengju Liu, Hongzhi Zhang, Kai Zhang, Liang Lin, and Wangmeng Zuo. Multi-level wavelet-CNN for image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pp. 773–782, 2018.

Yang Liu, Zhenyue Qin, Saeed Anwar, Pan Ji, Dongwoo Kim, Sabrina Caldwell, and Tom Gedeon. Invertible denoising network: A light solution for real noise removal. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 13360–13369. IEEE, 2021.

Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations*, 2017. URL `https://openreview.net/pdf?id=Skq89Scxx`.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/pdf?id=Bkg6RiCqY7.

Jiaqi Ma, Shengkai Hu, Xu Zhang, Jun Wan, Jiaxing Huang, Lefei Zhang, and Salman Khan. EvoIR: towards all-in-one image restoration via evolutionary frequency modulation. *arXiv preprint arXiv:2512.05104*, 2025.

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, 2017.

Xiaojiao Mao, Chunhua Shen, and Yu-Bin Yang. Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. *Advances in Neural Information Processing Systems*, 29, 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 of the 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.

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.

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 *The Fourteenth International Conference on Learning Representations*, 2026a.

Bin Ren, Runyi Yang, Qi Ma, Xu Zheng, Mengyuan Liu, Danda Pani Paudel, Luc Van Gool, Rita Cucchiara, and Nicu Sebe. Rethinking expressivity and degradation-awareness in attention for All-in-One blind image restoration. In *International Conference on Learning Representations*, 2026b.

Bin Ren, Eduard Zamfir, Zongwei Wu, Yawei Li, Yidi Li, Danda Pani Paudel, Radu Timofte, Ming-Hsuan Yang, Luc Van Gool, and Nicu Sebe. Any image restoration via efficient spatial-frequency degradation adaptation. *Transactions on Machine Learning Research*, 2026c. ISSN 2835-8856.

Shangquan Sun, Wenqi Ren, Xinwei Gao, Rui Wang, and Xiaochun Cao. Restoring images in adverse weather conditions via histogram transformer. In *European Conference on Computer Vision*, pp. 111–129. Springer, 2024.

Xiaole Tang, Xiang Gu, Xiaoyi He, Xin Hu, and Jian Sun. Degradation-aware residual-conditioned optimal transport for unified image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(8):6764–6779, 2025.

Xiaole Tang, Xiaoyi He, Jiayi Xu, Xiang Gu, and Jian Sun. Learning continuous Wasserstein barycenter space for generalized all-in-one image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 48(7):8395–8410, 2026. doi: 10.1109/TPAMI.2026.3669121.

Chunwei Tian, Yong Xu, and Wangmeng Zuo. Image denoising using deep CNN with batch renormalization. *Neural Networks*, 121:461–473, 2020. doi: 10.1016/j.neunet.2019.08.022.

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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2353–2363, 2022.

Bingnan Wang, Bin Qin, Jiangmeng Li, Fanjiang Xu, Fuchun Sun, and Hui Xiong. All-in-one image restoration via causal-deconfounding wavelet-disentangled prompt network. *IEEE Transactions on Image Processing*, 35:3421–3436, 2026.

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. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2018. URL `https://bmva-archive.org.uk/bmvc/2018/contents/papers/0451.pdf`.

Gang Wu, Junjun Jiang, Kui Jiang, Xianming Liu, and Liqiang Nie. DSwinIR: Rethinking window-based attention for image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 48(4):4350–4366, 2026. doi: 10.1109/TPAMI.2025.3646016.

Rui Xu, Yuzhen Niu, Yuezhou Li, Huangbiao Xu, Wenxi Liu, and Yuzhong Chen. URWKV: Unified RWKV model with multi-state perspective for low-light image restoration. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 21267–21276. IEEE, 2025.

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 *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1685–1694. IEEE, 2017.

Mingde Yao, Ruikang Xu, Yuanshen Guan, Jie Huang, and Zhiwei Xiong. Neural degradation representation learning for all-in-one image restoration. *IEEE Transactions on Image Processing*, 33:5408–5423, 2024. doi: 10.1109/TIP.2024.3456583.

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 *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.

Jinghao Zhang, Jie Huang, Mingde Yao, Zizheng Yang, Hu Yu, Man Zhou, and Feng Zhao. Ingredient-oriented multi-degradation learning for image restoration. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5825–5835. IEEE, 2023.

Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017a. doi: 10.1109/TIP.2017.2662206.

Kai Zhang, Wangmeng Zuo, Shuhang Gu, and Lei Zhang. Learning deep CNN denoiser prior for image restoration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017b.

Kai Zhang, Wangmeng Zuo, and Lei Zhang. FFDNet: Toward a fast and flexible solution for CNN-based image denoising. *IEEE Transactions on Image Processing*, 27(9):4608–4622, 2018. doi: 10.1109/TIP.2018.2839891.

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*, 35:2018–2033, 2026a. doi: 10.1109/TIP.2025.3566300.

Xu Zhang, Huan Zhang, Guoli Wang, Qian Zhang, and Lefei Zhang. ClearAIR: A human-visual-perception-inspired all-in-one image restoration. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(15):12861–12869, 2026b. doi: 10.1609/aaai.v40i15.38284.

Dian Zheng, Xiao-Ming Wu, Shuzhou Yang, Jian Zhang, Jian-Fang Hu, and Wei-Shi Zheng. Selective hourglass mapping for universal image restoration based on diffusion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 25445–25455, 2024.

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

Yurui Zhu, Tianyu Wang, Xueyang Fu, Xuanyu Yang, Xin Guo, Jifeng Dai, Yu Qiao, and Xiaowei Hu. Learning weather-general and weather-specific features for image restoration under multiple adverse weather conditions. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 21747–21758. IEEE, 2023.

A Datasets and Evaluation Protocols

Tab. 7 summarizes the training and test datasets for all-in-one and single-task restoration.

For the **3D setting**, one model jointly handles Gaussian denoising, dehazing, and deraining. BSD400 (Arbeláez et al., 2011) and the Waterloo Exploration Database (WED) (Ma et al., 2017) provide clean denoising images, with Gaussian noise synthesized at $\sigma \in \{15, 25, 50\}$. We evaluate denoising on CBSD68 (Martin et al., 2001), dehazing on SOTS-Outdoor from RESIDE-β (Li et al., 2019), and deraining on Rain100L (Yang et al., 2017).

The **5D setting** extends 3D with motion deblurring and low-light enhancement. We evaluate a single model on SOTS-Outdoor for dehazing, Rain100L for deraining, CBSD68 at $\sigma = 25$ for denoising, GoPro (Nah et al., 2017) for deblurring, and LOL-v1 (Wei et al., 2018) for low-light enhancement.

The **CDD11 setting** follows the benchmark introduced by OneRestore (Guo et al., 2024b). It contains four single degradations (L, H, R, and S), five double degradations ($L + H$, $L + R$, $L + S$, $H + R$, and $H + S$), and two triple degradations ($L + H + R$ and $L + H + S$). We use the original training and test split.

For the **WeatherBench setting**, we use the real-world paired benchmark of haze, rain, and snow (Guan et al., 2025). WeatherBench contains 41,402 training pairs and 600 test pairs, with 200 test pairs for each weather condition.

For **single-task restoration**, we train separate models for dehazing, deraining, and denoising using the datasets in Tab. 7. We evaluate denoising on Urban100 (Huang et al., 2015) and CBSD68 at $\sigma \in \{15, 25, 50\}$.

Evaluation protocols. We report PSNR and SSIM (Wang et al., 2004) for all-in-one restoration and single-task dehazing and deraining. For single-task denoising, we report PSNR.

Table 7: Training and test datasets for all-in-one and single-task restoration.

Protocol	Degradation	Training dataset (Number)	Test dataset (Number)
All-in-one restoration			
3D	Noise	BSD400 (400) + WED (4,744)	CBSD68 (68)
	Haze	RESIDE-β OTS (72,135)	SOTS-Outdoor (500)
	Rain	Rain100L (200)	Rain100L (100)
5D	Noise	BSD400 (400) + WED (4,744)	CBSD68 (68)
	Haze	RESIDE-β OTS (72,135)	SOTS-Outdoor (500)
	Rain	Rain100L (200)	Rain100L (100)
	Blur	GoPro (2,103)	GoPro (1,111)
	Low light	LOL-v1 (485)	LOL-v1 (15)
CDD11	Single and composite (11 types)	CDD11 (13,013; 1,183 per type)	CDD11 (2,200; 200 per type)
WeatherBench	Haze, rain, snow	WeatherBench (41,402)	WeatherBench (600; 200 per condition)
Single-task restoration			
Single-task	Haze	RESIDE-β OTS (72,135)	SOTS-Outdoor (500)
	Rain	Rain100L (200)	Rain100L (100)
	Noise	BSD400 (400) + WED (4,744)	Urban100 (100), CBSD68 (68)

B Implementation Details

We implement ReDU-IR in PyTorch and use distributed data-parallel training on two NVIDIA GeForce RTX 4090 GPUs. We augment training images with random 128×128 crops, horizontal and vertical flips, and rotations. We optimize the model with AdamW Loshchilov & Hutter (2019), using $\beta_1 = 0.9$, $\beta_2 = 0.999$, and a weight decay of 1×10^{-3}. A cosine annealing schedule Loshchilov & Hutter (2017) decreases the learning rate from 1×10^{-3} to 1×10^{-7} without warm-up. We clip gradients to a maximum ℓ_2 norm of 0.01 to stabilize optimization. The objective combines an ℓ_1 reconstruction loss with a fast Fourier transform (FFT) domain reconstruction loss, with a weight of 0.1 for the FFT loss.

We train the 3D model for 300K iterations and the 5D, CDD11, and WeatherBench models for 400K iterations. Each all-in-one setting uses a total batch size of 32, with 16 samples per GPU. We train each single-task model for 400K iterations with a total batch size of 8, with four samples per GPU. Unless otherwise specified, the all-in-one and single-task experiments use the same remaining optimization and augmentation settings.

C Single-Task Restoration Results

Single-task image dehazing. Tab. 8 shows that ReDU-IR-S achieves 33.66 dB/0.984, outperforming the runner-up MIRAGE (Ren et al., 2026a) by 2.13 dB/0.004.

Table 8: Single-task dehazing on SOTS-Outdoor (Li et al., 2019).

Metric	DehazeNet	MSCNN	AODNet	EPDN	FDGAN	AirNet	Restormer	PromptIR	MIRAGE	**ReDU-IR-S**
PSNR	22.46	22.06	20.29	22.57	23.15	23.18	30.87	31.31	<u>31.53</u>	**33.66**
SSIM	.851	.908	.877	.863	.921	.900	.969	.973	<u>.980</u>	**.984**

Single-task image deraining. Tab. 9 shows that ReDU-IR-S achieves 39.06 dB/0.985, outperforming the runner-up MIRAGE (Ren et al., 2026a) by 1.05 dB/0.003.

Table 9: Single-task deraining on Rain100L (Yang et al., 2017).

Metric	DID-MDN	UMR	SIRR	MSPFN	LPNet	AirNet	Restormer	PromptIR	MIRAGE	**ReDU-IR-S**
PSNR	23.79	32.39	32.37	33.50	33.61	34.90	36.74	37.04	<u>38.01</u>	**39.06**
SSIM	.773	.921	.926	.948	.958	.977	.978	.979	<u>.982</u>	**.985**

Gaussian image denoising. As shown in Table 10, ReDU-IR-S achieves the best performance in five of the six test conditions, with an average PSNR of 32.12 dB. Compared with AdaIR, it improves PSNR by 0.16–0.40 dB on Urban100 while performing comparably on CBSD68. It also consistently outperforms MIRAGE-S across all six conditions. These results demonstrate stronger denoising performance on Urban100 while maintaining competitive performance on CBSD68.

Table 10: Single-task Gaussian color image denoising on Urban100 and CBSD68 at $\sigma \in \{15, 25, 50\}$. † denotes MIRAGE-S retrained by us for denoising only using its official code.

	Urban100			CBSD68			
Method	$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	Average
CBM3D (Dabov et al., 2007)	33.93	31.36	27.93	33.50	30.69	27.36	30.80
DnCNN (Zhang et al., 2017a)	32.98	30.81	27.59	33.89	31.23	27.92	30.74
IRCNN (Zhang et al., 2017b)	33.78	31.20	27.70	33.86	31.16	27.86	30.93
FFDNet (Zhang et al., 2018)	33.83	31.40	28.05	33.87	31.21	27.96	31.05
BRDNet (Tian et al., 2020)	34.42	31.99	28.56	34.10	31.43	28.16	31.44
AirNet (Li et al., 2022)	34.40	32.10	28.88	34.14	31.48	28.23	31.54
PromptIR (Potlapalli et al., 2023)	34.77	32.49	29.39	34.34	31.71	<u>28.49</u>	31.87
AdaIR (Cui et al., 2025)	<u>34.96</u>	<u>32.74</u>	<u>29.70</u>	**34.36**	<u>31.72</u>	<u>28.49</u>	<u>32.00</u>
MIRAGE-S† (Ren et al., 2026a)	33.90	32.02	28.97	34.13	31.55	28.30	31.48
ReDU-IR-S (Ours)	**35.36**	**32.93**	**29.86**	<u>34.35</u>	**31.73**	**28.51**	**32.12**

D Model Configurations

Table 11: Architecture specifications of ReDU-IR-S and ReDU-IR-B. The decoder mirrors the encoder with $[6, 6, 4]$ Transformer blocks. We count floating-point operations (FLOPs) on 256×256 rainy images. FFN denotes the feed-forward network.

Configuration	ReDU-IR-S	ReDU-IR-B
Transformer blocks per encoder scale	$[4, 6, 6]$	$[4, 6, 6]$
Input embedding dimension	32	48
FFN expansion factor	2.66	2.66
Transformer blocks at the bottleneck	8	8
Params. (M) (↓)	10.1	22.5
FLOPs (G) (↓)	36.3	79.5

E Computational Complexity

Tab. 12 shows that ReDU-IR-S achieves the lowest latency among the compared methods at 0.033 s. ReDU-IR-B requires 79.5 G FLOPs and 0.039 s, compared with 282.0 G and 0.050 s for the similarly sized Restormer.

Table 12: Computational complexity comparison. We count FLOPs on a 256×256 rainy image, with two floating-point operations per multiply-accumulate, and measure inference time using an NVIDIA GeForce RTX 4090 GPU.

Method	Params. (M)	FLOPs (G)	Time (s)
Restormer (Zamir et al., 2022)	26.1	282.0	0.050
PromptIR (Potlapalli et al., 2023)	35.6	316.3	0.054
MoCE-IR (Zamfir et al., 2025)	25.4	86.4	0.106
AnyIR-S (Ren et al., 2026c)	**8.5**	49.8	0.048
MIRAGE-S (Ren et al., 2026a)	<u>9.7</u>	**33.7**	<u>0.035</u>
ReDU-IR-S (Ours)	10.1	<u>36.3</u>	**0.033**
ReDU-IR-B (Ours)	22.5	79.5	0.039

F SENSITIVITY TO SCU UPDATE STRENGTH WITH FIXED PARAMETERS

Table 13: Inference-time sensitivity to the SCU update scale λ under the 3D protocol.

λ	SOTS-Outdoor	Rain100L	CBSD68			Avg.
			$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	
0.00	32.40 / .981	39.19 / .984	34.15 / .934	31.52 / .891	28.25 / .799	33.101 / .9177
0.25	32.45 / .981	39.33 / .985	34.21 / .936	31.57 / .893	28.32 / .804	33.175 / .9199
0.50	32.47 / .981	39.41 / .986	34.25 / .937	31.61 / .895	28.36 / .808	33.222 / .9213
0.75	32.48 / .982	39.46 / .986	34.27 / .937	31.63 / .896	28.39 / .810	33.247 / .9221
1.00	32.47 / .982	39.47 / .986	34.28 / .937	31.64 / .896	28.40 / .810	**33.251 / .9223**
1.25	32.44 / .981	39.45 / .986	34.27 / .937	31.64 / .896	28.39 / .810	33.237 / .9220

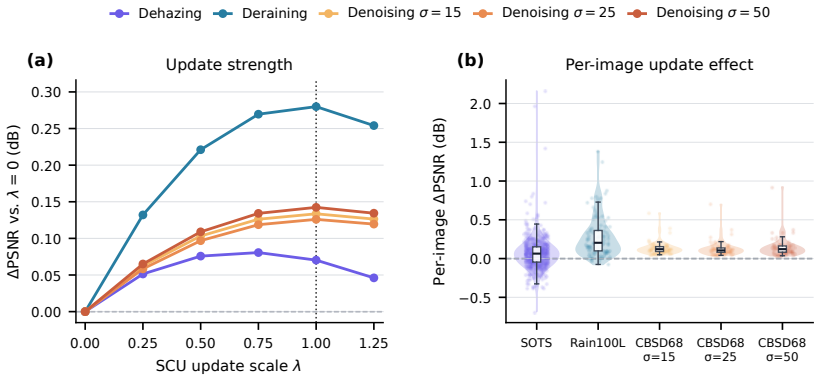

Figure 4: Sensitivity to SCU updates with fixed parameters. (a) PSNR changes relative to $\lambda = 0$; the dotted line marks $\lambda = 1$. (b) Per-image PSNR differences between $\lambda = 1$ and $\lambda = 0$. Violins show distributions; boxes show medians and interquartile ranges, with whiskers extending to the most extreme values within 1.5 times the interquartile range. Dots denote individual images.

Under the 3D protocol, we fix the parameters of ReDU-IR-B and scale its SCU updates at all stages by a common coefficient λ during inference. Setting $\lambda = 0$ retains the original complementary representation, whereas $\lambda = 1$ recovers the default update. This experiment examines the trained model's sensitivity to the update strength.

Tab. 13 shows that $\lambda = 1$ improves the five-condition average PSNR and SSIM by 0.150 dB and 0.0046, respectively, over $\lambda = 0$. These gains indicate that enabling SCU updates improves restoration in the model with fixed parameters. Fig. 4(a) further shows that the response to update strength varies across tasks. Dehazing PSNR peaks at $\lambda = 0.75$, while deraining and denoising at all three noise levels reach their highest PSNR at $\lambda = 1$. The average PSNR differs by only 0.004 dB between $\lambda = 0.75$ and $\lambda = 1$, and average performance decreases slightly at $\lambda = 1.25$. Performance remains similar near the default strength; increasing the tested strength to $\lambda = 1.25$ gives no additional average benefit.

Fig. 4(b) shows the per-image effects of the default update. Compared with $\lambda = 0$, $\lambda = 1$ improves PSNR for 96% of Rain100L images and all CBSD68 images at each noise level. The corresponding proportion on SOTS-Outdoor is 63.4%, indicating that improvements cover most test images but occur less consistently in dehazing than in deraining and denoising. Overall, SCU updates improve restoration near the default strength, with the gain magnitude and consistency of per-image improvements varying across degradation types.

Table 14: SCU intervention with fixed parameters on CDD11. Each condition contains 200 test images. Scores are PSNR (dB)/SSIM.

Condition	$\lambda = 0$	$\lambda = 1$	Δ
L	27.85 / .8356	27.90 / .8409	+0.05 / +.0054
H	38.41 / .9937	39.38 / .9950	+0.97 / +.0013
R	35.23 / .9625	36.36 / .9783	+1.14 / +.0158
S	38.15 / .9829	38.41 / .9846	+0.26 / +.0017
$L + H$	27.62 / .8362	27.70 / .8391	+0.08 / +.0030
$L + R$	27.03 / .8071	27.18 / .8226	+0.15 / +.0156
$L + S$	27.07 / .8097	27.15 / .8164	+0.08 / +.0068
$H + R$	33.75 / .9694	34.22 / .9772	+0.46 / +.0079
$H + S$	34.63 / .9778	35.04 / .9803	+0.41 / +.0025
$L + H + R$	26.63 / .8088	26.77 / .8158	+0.14 / +.0070
$L + H + S$	26.82 / .8130	26.90 / .8162	+0.08 / +.0032
Average	31.20 / .8906	**31.55 / .8970**	**+0.35 / +.0064**

L, H, R, and S denote low light, haze, rain, and snow. Average gives the arithmetic mean over the eleven conditions. Differences give scores at $\lambda = 1$ minus those at $\lambda = 0$ and use unrounded scores.

SCU updates improve restoration on CDD11. Tab. 14 compares ReDU-IR-S with and without SCU updates while keeping its parameters fixed. Enabling SCU improves PSNR and SSIM across all eleven conditions, with average gains of 0.35 dB and 0.0064, respectively. Rain and haze show the largest PSNR gains of 1.14 and 0.97 dB, while the seven composite conditions improve by 0.08–0.46 dB. These results support updating the complementary representation during reconstruction under both single and composite degradations.

G Stage-specific SCU Sensitivity with Fixed Parameters

We disable one SCU update at a time in the same trained ReDU-IR-B model used in Appendix F, keeping all parameters fixed. Fig. 5 shows that disabling shallow, middle, and deep updates changes the five-condition mean PSNR by -0.123, -0.011, and $+0.006$ dB, respectively. The model is most sensitive to the shallow update. Disabling the deep update improves SOTS-Outdoor by 0.033 dB, indicating variation across stages and test conditions. The effect of an SCU update therefore depends on both the stage and the test condition.

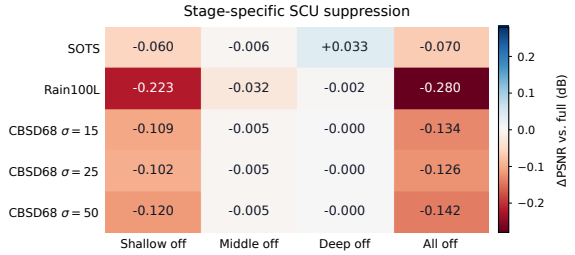

Figure 5: Stage-specific SCU sensitivity with fixed parameters under the 3D protocol. Columns disable the indicated update or all updates; other stages retain their default updates. Shallow, middle, and deep follow the encoder scale order. Values are PSNR differences from the default model. Positive values indicate an improvement after disabling updates.

H Limitations and Future Work

We have evaluated ReDU-IR only in hierarchical encoder–decoder networks. The tests with fixed parameters show that the effect of SCU varies across stages and test conditions. Future work will examine update strengths that adapt to scale and degradation characteristics, applicability to other architectures, and generalization to unseen degradations.

I QUALITATIVE RESULTS



Figure 6: Qualitative comparison under the 3D protocol. Rows show groups of haze, rain, and Gaussian noise, with two examples per degradation. Columns show Input, Restormer, PromptIR, AdaIR, MIRAGE, ReDU-IR, and ground truth from left to right. Outlined boxes mark corresponding regions.

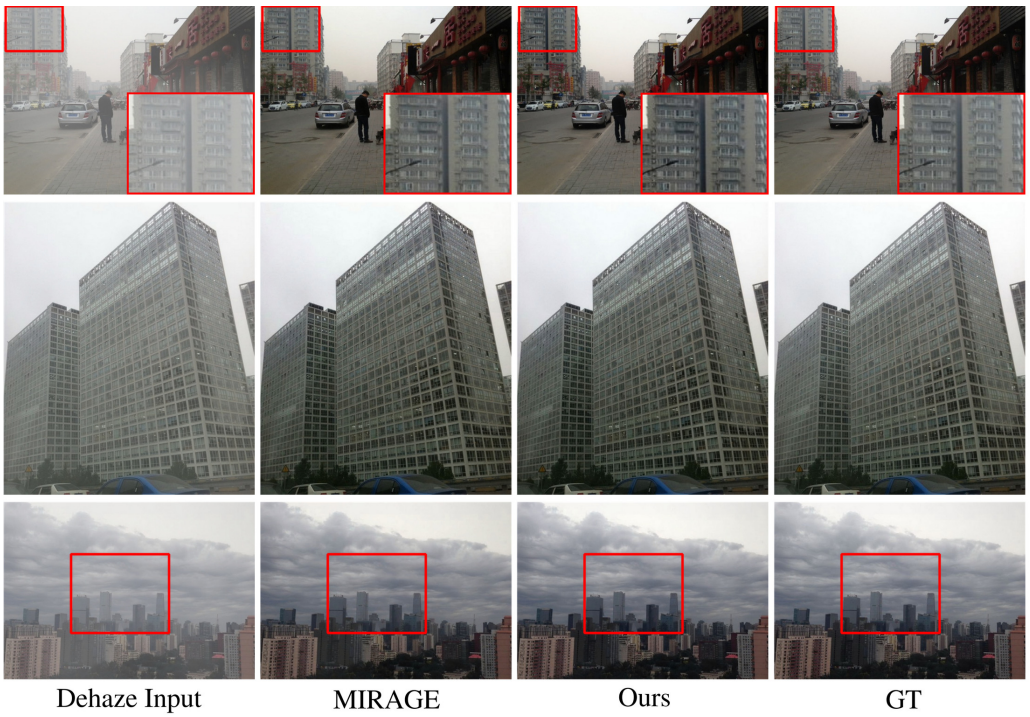

Figure 7: Qualitative comparison for single-task dehazing on three SOTS-Outdoor examples. Columns show Input, MIRAGE, ReDU-IR, and ground truth from left to right. Outlined boxes mark corresponding regions and enlarged details.



Figure 8: Qualitative comparison under the 5D protocol. Rows show groups of haze, rain, Gaussian noise, blur, and low light, with two examples per degradation. Columns show Input, Restormer, PromptIR, AdaIR, MIRAGE, ReDU-IR, and ground truth from left to right. Outlined boxes mark corresponding regions for detail comparison.

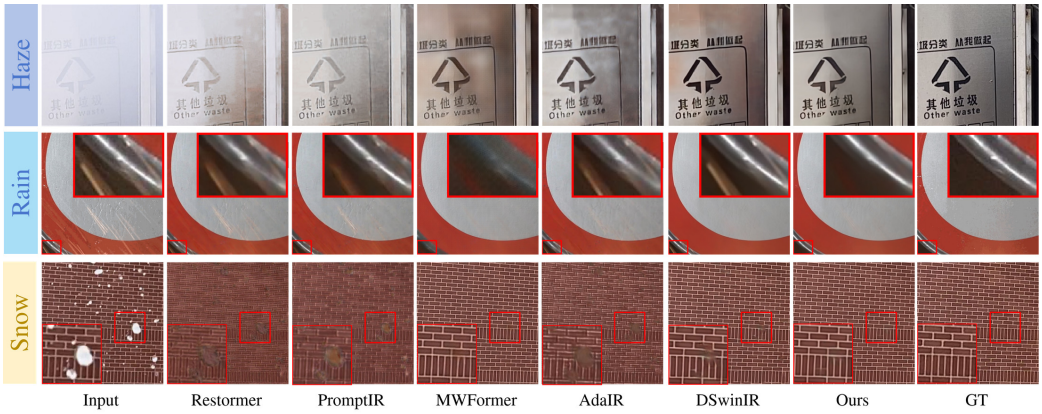

Figure 9: Qualitative comparison on WeatherBench. Rows show real paired examples with haze, rain, and snow, respectively. Columns show Input, Restormer, PromptIR, MWFormer, AdaIR, DSwinIR, ReDU-IR, and ground truth from left to right. Outlined boxes mark corresponding regions and enlarged details.

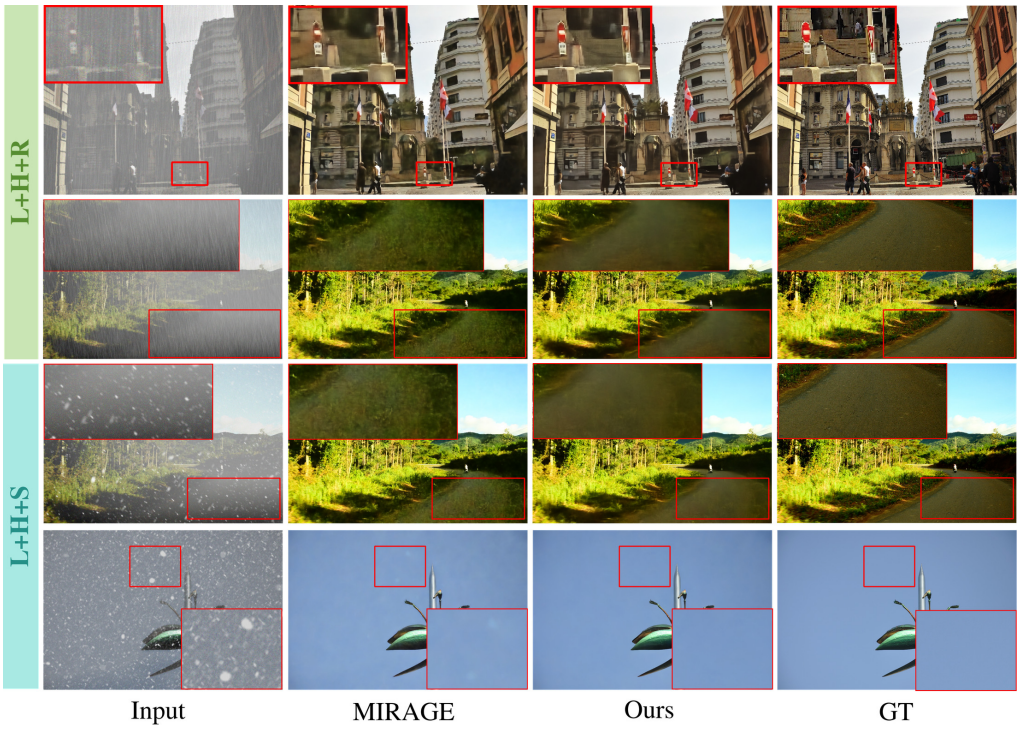

Figure 10: Qualitative comparison on CDD11. The top two rows contain low-light, haze, and rain ($L+H+R$); the bottom two contain low-light, haze, and snow ($L+H+S$). Columns show Input, MIRAGE, ReDU-IR, and ground truth from left to right. Outlined boxes mark corresponding regions and enlarged details.



Figure 11: Qualitative comparison for single-task deraining on three Rain100L examples. Columns show Input, MIRAGE, ReDU-IR, and ground truth from left to right. Outlined boxes mark corresponding regions and enlarged details.

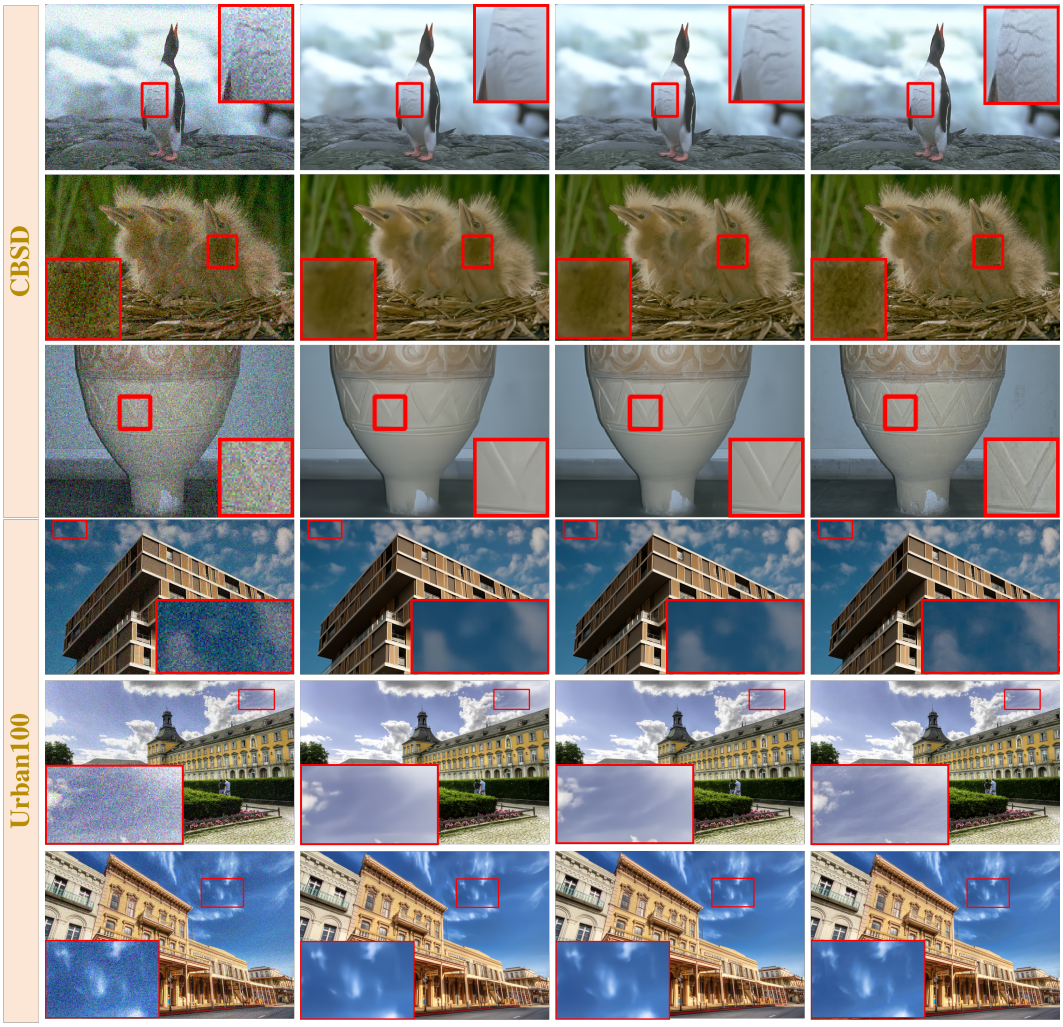

Figure 12: Qualitative comparison for single-task Gaussian denoising. Rows group examples by dataset, as labeled on the left, with three examples per dataset. Columns show Input, MIRAGE, ReDU-IR, and ground truth from left to right. Outlined boxes mark corresponding regions and enlarged details.