

SLER-IR: Spherical Layer-wise Expert Routing for All-in-One Image Restoration

Shurui Peng^{1*}, Xin Lin^{2*}, Shi Luo^{1*}, Jincen Ou¹,
Dizhe Zhang³, Lu Qi³, Truong Nguyen², Chao Ren^{1†}

¹Sichuan University, Chengdu, China

²University of California San Diego, La Jolla, CA, USA

³Insta360 Research, Shenzhen, China

*Equal contribution. †Corresponding author.

`chaoren@scu.edu.cn`

Abstract. Image restoration under diverse degradations remains challenging for unified all-in-one frameworks due to feature interference and insufficient expert specialization. We propose **SLER-IR**, a spherical layer-wise expert routing framework that dynamically activates specialized experts across network layers. To ensure reliable routing, we introduce a Spherical Uniform Degradation Embedding with contrastive learning, which maps degradation representations onto a hypersphere to eliminate geometry bias in linear embedding spaces. In addition, a Global-Local Granularity Fusion (GLGF) module integrates global semantics and local degradation cues to address spatially non-uniform degradations and the train-test granularity gap. Experiments on three-task and five-task benchmarks demonstrate that **SLER-IR** achieves consistent improvements over state-of-the-art methods in both PSNR and SSIM. Code and models will be available at <https://github.com/PSR666/SLER-IR>.

Keywords: All-in-one image restoration · Mixture of experts · Degradation representation learning

1 Introduction

Image restoration aims to reconstruct high-quality images from degraded inputs (e.g., noise, rain, blur), serving as a fundamental prerequisite for robust downstream vision tasks. Traditionally, it has been treated as a collection of task-specific problems, with models tailored exclusively for noise, haze, rain, or blur [3, 4, 7–9, 12, 19, 21–23, 28, 33–38, 43, 45, 47, 49, 50, 54, 57–60, 69–71]. While effective within their respective domains, these methods lack the flexibility to handle complex or unseen degradations. Extending them to new scenarios requires labor-intensive retraining, limiting their practicality in real-world applications.

To address this, all-in-one restoration frameworks [10, 11, 14–18, 27, 30, 40, 46, 51–53, 55, 67, 68, 72] have emerged to handle diverse degradations within a single model, offering superior generalization across real-world conditions. As illus-




Fig. 1: Illustration of representative all-in-one image restoration frameworks. (a) Prompt-based methods [10, 27, 40, 46, 52] adapt a unified parameter space via feature modulation. (b) Expert-based methods [1, 29, 39, 61–64] activate task-specific experts. (c) Our **SLER-IR** introduces spherical expert routing and global–local granularity fusion for robust restoration.




Fig. 2: Quantitative comparisons ($\text{PSNR}^{1-5}/\text{SSIM}^{1-5}$) of four models (AirNet [27], Restormer [65], MoCE-IR [64], and **SLER-IR** (ours)) across five image restoration tasks (Derain, Dehaze, Deblur, Low-light, Denoise).

trated in Figure 1, these methods generally fall into two paradigms: feature modulation (Figure 1(a)), which employs prompts or contrastive learning to adapt a unified parameter space (e.g., AirNet [27], PromptIR [46], InstructIR [10], DA-CLIP [40]), and architectural adaptation (Figure 1(b)), which utilizes Mixture-of-Experts (MoE) to activate task-specific components (e.g., WM-MoE [39], MoCE-IR [64]). However, both paradigms exhibit notable limitations. Firstly, feature modulation methods suffer from feature interference within the shared backbone. Conflicting objectives, exemplified by the trade-off between suppressing noise and restoring high-frequency details, hinder fine-grained, degradation-specific modeling. Secondly, expert-based architectures often restrict experts to local modules, failing to fully exploit degradation priors. Furthermore, existing routing mechanisms struggle to resolve adversarial relationships among tasks, leading to suboptimal specialization. Beyond these structural issues, most methods overlook locally non-continuous degradations (e.g., rain streaks) and neglect the granularity gap between training crops and full-resolution inference.

To address these limitations, we propose **SLER-IR**, an all-in-one image restoration framework equipped with a spherical layer-wise expert routing mechanism. Structurally, we transform the encoder–decoder by replacing each backbone block with multiple parameter-independent experts. To ensure precise expert selection, we employ a spherically uniform embedding strategy. This approach mitigates the class-distance bias often found in linear embeddings and, complemented by a triplet-based contrastive loss, significantly enhances degradation discriminability. Consequently, the network performs layer-wise dynamic activation, adaptively recalibrating its processing strategy as features evolve from global to local levels. As a result, **SLER-IR** enables a compositional inference paradigm, where expert specialization and path diversity emerge natu-

rally through layer-wise routing. Such exponentially diverse routing trajectories allow the model to flexibly adapt to complex and composite degradations without increasing inference overhead. Finally, a Global-Local Granularity Fusion module bridges the granularity gap between training patches and full-resolution inference images, ensuring consistent and fine-grained guidance. Extensive experiments show that our proposed method achieves comprehensive superiority across multiple types of degradation, as highlighted in the quantitative comparison in Figure 2.

The main contributions of this work are summarized as follows:

1. We propose **SLER-IR**, a unified all-in-one image restoration framework featuring Spherical Layer-wise Expert Routing. By expanding each backbone layer into parameter-independent experts and enabling layer-wise expert activation, the model supports compositional inference paths and achieves progressive specialization across diverse and composite degradations.
2. We introduce Spherical Uniform Degradation Embedding with Contrastive Learning for fine-grained degradation discrimination. By mapping degradation embeddings onto a unit hypersphere and optimizing them with a triplet-constrained contrastive objective, our method mitigates class-distance bias in linear spaces, thereby enhancing the separability and robustness of degradation representations for stable routing.
3. We design Global-Local Granularity Fusion (GLGF) to harmonize global image context (CLS tokens) with local degradation cues (patch tokens). GLGF bridges the granularity gap between patch-based training and full-resolution inference, leading to robust routing guidance under complex and spatially non-uniform degradations.

2 Related Works

Single-Task Image Restoration. Single-task image restoration methods target specific degradation types and typically assume that the degradation pattern is known and consistent with the training data. While these approaches achieve strong performance within their respective domains, they often lack the flexibility to generalize to unseen or mixed degradations. Existing methods can be categorized by degradation type. For denoising, representative works include DnCNN [69], FFDNet [71], CasaPuNet [21], and Neighbor2neighbor [22]. For dehazing, typical methods include DehazeNet [3], FFA-Net [48], GridDehazeNet [37], and PMHLD [7]. For deraining, representative models include PReNet [49] and ERL-Net [54]. DesnowNet [38] addresses image desnowing, while DeblurGAN [24] and DeblurGANv2 [25] focus on deblurring. For low-light enhancement, representative works include LIME [20], DeepLPF [43], UretinexNet [58], and RetinexFormer [4]. In super-resolution, typical methods include SRCNN [12] and EDSR [32].

More recently, several general-purpose restoration frameworks, such as MPRNet [66], SwinIR [31], NAFNet [6], and Restormer [65], have been proposed to




Fig. 3: Overview of the proposed **SLER-IR** framework. Given an input LQ image, the degradation extractor produces a raw routing vector, which is projected onto the unit hypersphere for cosine-similarity gating, enabling layer-wise expert selection under Stage 1 (Probabilistic Routing) and Stage 2 (Deterministic Routing). In parallel, a Global-Local Map Construction (GLMC) module derives a content semantic patch (CSP) map and a degradation severity patch (DSP) map, which are fused through Content-Guided Degradation Fusion (CGDF) to guide restoration.

improve model generality across multiple restoration tasks. Despite their strong performance, these models still require retraining or task-specific adaptation when encountering unseen or composite degradations. To overcome this limitation and better address complex real-world scenarios, all-in-one image restoration frameworks have recently emerged.

All-in-One Image Restoration. All-in-one restoration aims to train a unified model capable of handling multiple degradations (e.g., noise, blur, haze, and low-light) without explicit task identification or separate fine-tuning.

One research direction employs intrinsic or extrinsic degradation prompts to modulate features within a unified parameter space. AirNet [27] learns degradation representations via a contrastive encoder, while PromptIR [46] and Instruc-tIR [10] use prompt-based or language-guided embeddings to steer restoration. DA-CLIP [40] decouples degradation knowledge from content, and UniProces-sor [13] further achieves explicit prompt-driven modulation. Another direction focuses on architectural designs tailored for heterogeneous degradations. Early works, such as [29], adopt multi-encoder-single-decoder structures, while IPT [5] employs a shared backbone with multi-head and multi-tail branches. MoE-based methods—including WM-MoE [39], LoRA-IR [1], DaAIR [63], AMIR [61], MEASNet [62], and MoCE-IR [64]—activate task-specific experts via routing or gating, with the latter introducing complexity-aware routing. Despite their

progress, these approaches still suffer from limited use of degradation priors and incomplete parameter sharing. Experts often differ only in isolated modules, restricting their adaptability across diverse degradations, and existing balanced-loss strategies inadequately address directional conflicts, hindering effective routing and expert specialization.

3 Method

As illustrated in Fig. 3, **SLER-IR** comprises four integral components: Spherical Layer-wise Expert Routing (SLER) (Sec. 3.1), which serves as the core routing backbone and constructs a unified multi-expert architecture with compositional inference paths; Router Design and Similarity-based Expert Gating (Sec. 3.2), which performs layer-wise expert selection via similarity-based gating; Hyperspherical Degradation Representation Learning (Sec. 3.3), which learns geometry-consistent degradation embeddings to mitigate class-distance bias and stabilize routing decisions; and Global-Local Granularity Fusion (GLGF) (Sec. 3.4), which integrates global CLS tokens and local patch cues to bridge the train-test granularity gap and provide robust routing guidance under varying degradations.

3.1 Spherical Layer-wise Expert Routing

To enable degradation-specific processing within a unified model, we transform the encoder-decoder by expanding each backbone layer into multiple parameter-independent expert nodes. Unlike conventional static networks, each layer contains three candidate experts with distinct latent feature representations. This multi-expert backbone, governed by Spherical Layer-wise Expert Routing (SLER), forms the core routing architecture of **SLER-IR**.




Fig. 4: Visualization of dynamic routing trajectories across different degradations. Stage 1 learns degradation-aware routing via probabilistic expert selection, while Stage 2 forms specialized expert paths through deterministic routing across layers for different degradations.

As shown in Fig. 4, during inference the network does not activate all experts; instead, it dynamically selects a specific expert path for each layer based

on the input’s degradation characteristics. Since routing decisions are made independently across the eight backbone layers, the network can compose up to $3^8 = 6561$ unique inference paths. This exponential path diversity allows the model to flexibly specialize to complex and composite degradations without increasing the inference computational burden. Importantly, the routing trajectories are not fixed per task but adapt progressively across layers, allowing composite degradations to activate hybrid expert combinations.

While the proposed multi-expert backbone provides exponential path diversity and compositional inference capability, its effectiveness critically depends on how experts are selected. Without a reliable degradation-aware routing signal, the large routing space may lead to unstable specialization or suboptimal path selection. Therefore, we next introduce the routing mechanism that enables stable and progressive expert activation.

3.2 Router Design and Similarity-based Expert Gating

Router Inputs and Routing Vector. Given an input degraded image x , we extract semantic priors using a Content CLIP encoder and degradation cues using a Degraded CLIP encoder. The degradation-related features are then fed into a Router to produce a layer-wise routing representation. Concretely, we take the CLS token from Degraded CLIP and pass it through an MLP followed by L2 normalization to obtain the raw branch selection vector $f \in \mathbb{R}^{K \times d}$:

$$f = \text{L2}(\text{MLP}(\text{CLIP}(x))), \quad (1)$$

where K is the number of backbone layers (e.g., $K=8$) and d is the routing embedding dimension. Importantly, the L2 normalization maps routing features onto the unit hypersphere, forming our *spherical uniform embedding* and enabling stable cosine-similarity-based expert selection.

Probabilistic Routing (Stage I: Router Optimization). In the first stage, we adopt probabilistic routing to optimize the Router and learn stable degradation-aware representations. We compute cosine similarities between each layer-wise routing vector $f_i \in \mathbb{R}^d$ and a set of expert centers $\mathbf{C} \in \mathbb{R}^{C \times d}$ (one center per candidate expert), and apply a *row-wise* Softmax to obtain selection probabilities $\mathbf{p} \in \mathbb{R}^{K \times C}$:

$$\text{Sim}_{i,j} = \frac{f_i \cdot \mathbf{C}_j^\top}{\|f_i\|_2 \|\mathbf{C}_j\|_2}, \quad \mathbf{p}_{i,:} = \text{Softmax}(\text{Sim}_{i,:}), \quad (2)$$

where $\mathbf{C}_j \in \mathbb{R}^d$ denotes the j -th center and $\mathbf{p}_{i,:} \in \mathbb{R}^C$ sums to 1. The expert centers $\mathbf{C}$ are updated by Monte Carlo-regularized uniform clustering: we perform Monte Carlo sampling over hyperspherical routing embeddings and impose a uniformity constraint, encouraging a more balanced assignment of samples across centers and yielding well-separated, uniformly distributed expert prototypes on the unit hypersphere. During this stage, the restoration objective (e.g., L_1 loss)

is jointly optimized with a triplet-constrained contrastive loss to enhance the discriminability of degradation representations on the hypersphere. The probabilistic formulation allows all experts to receive gradients, preventing premature specialization and stabilizing routing learning.

Deterministic Routing (Stage II: Restoration Refinement). In the second stage, we switch to deterministic routing and freeze the Router to obtain fixed expert assignments. Specifically, the expert index with the maximum probability is selected at each layer:

$$y_i = \arg \max_{j \in \{1, 2, \dots, C\}} \mathbf{p}_{i,j}, \quad (3)$$

where $\mathbf{y} \in \mathbb{N}^{K \times 1}$ denotes the hard selection indices across layers. At each backbone layer l , only the expert corresponding to y_l is activated, while the Router parameters remain frozen. Under fixed routing trajectories, the restoration network is further optimized to refine expert specialization and improve reconstruction quality. Because routing decisions are made independently across depths, the network dynamically recalibrates its processing strategy as feature representations evolve from low-level textures to high-level structures, yielding progressive specialization without activating all experts.

However, the effectiveness of deterministic routing ultimately depends on the quality of the learned degradation representations. This raises a fundamental question: what representation structure is required to ensure stable and reliable expert selection? We address this in Sec. 3.3.

3.3 Hyperspherical Degradation Representation Learning

Motivation. Our router performs similarity-based, degradation-aware expert gating (Sec. 3.2). Hence, routing quality hinges on whether the degradation representation space provides a geometry-consistent similarity structure: expert selection should reflect true degradation affinity rather than artifacts induced by embedding geometry. We observe that commonly used linear embedding strategies can introduce biased inter-degradation geometry, which directly destabilizes similarity-based gating.

Limitations of Existing Routing Methods. We identify two key limitations:

(i) Rigid mapping and coarse supervision. As shown in Fig. 5(a), many schemes enforce rigid degradation-to-expert assignments (e.g., manually categorized degradations or task-ID gating). When the number of experts is smaller than the number of degradation types, distinct degradations must be merged into the same expert, ignoring fine-grained similarities and differences. This limitation becomes more severe for composite degradations, where the optimal processing path should be progressively composed across layers rather than fixed to a single expert.

(ii) Geometry bias in unconstrained linear embedding spaces. As illustrated in Fig. 5(b), representations learned in an unconstrained linear space often exhibit class-distance bias: inter-class distances/similarities are not explicitly




Fig. 5: Comparison of existing degradation routing methods and our hyperspherical degradation representation-based branch selection strategy. (a) Traditional expert-based categorization and assignment neglect inter-degradation similarities. (b) Linear-space embeddings may introduce class-distance bias. (c) Our method maps degradations onto a unit hypersphere, enabling geometrically balanced distributions and robust branch selection via cosine similarity.

regularized and can become highly uneven. Since our gating is similarity-based, such bias contaminates routing scores and may trigger unstable or incorrect expert selection.

Counterexample aligned with Fig. 5(b): distance bias misleads similarity gating. Consider three degradation types A (Denoise), B (Derain), and C (Dehaze) embedded in a linear space. A biased yet plausible configuration is

$$d(A, C) \gg d(A, B), d(B, C), \quad (4)$$

meaning Denoise and Dehaze are artificially pushed far apart while Derain lies in between. Under a similarity-based gating rule

$$j^* = \arg \max_{j \in \{1, \dots, C\}} \text{sim}(\mathbf{f}, \mathbf{c}_j), \quad (5)$$

the embedding bias in Eq. (4) exaggerates dissimilarity between A and C regardless of their actual restoration affinity. As a result, the router may over-separate degradations that share processing components, or over-merge those that should be separated, leading to branch-selection flips and cross-degradation interference.

Spherical Uniform Embedding. To eliminate geometry-induced bias, we introduce Spherical Uniform Embedding, mapping degradation representations onto a unit hypersphere and performing angle-based comparison. Given a degradation feature vector $\mathbf{f}$ extracted by the degradation encoder, we project it onto

the hypersphere via L_2 normalization:

$$\hat{\mathbf{f}} = \frac{\mathbf{f}}{\|\mathbf{f}\|_2}. \quad (6)$$

Routing is then performed by cosine similarity to expert centers:

$$s_j = \cos(\hat{\mathbf{f}}, \hat{\mathbf{c}}_j) = \hat{\mathbf{f}}^\top \hat{\mathbf{c}}_j, \quad j^* = \arg \max_j s_j. \quad (7)$$

Importantly, we optimize expert centers to be mutually dissimilar (low pairwise cosine similarity), encouraging uniformly distributed expert directions on the hypersphere (spherical-code-like), which increases the minimum inter-center angle and stabilizes expert gating.

Hyperspherical Contrastive Optimization. To sharpen angular decision boundaries on the hypersphere, we adopt a triplet-constrained contrastive objective. For each layer l , similarity between samples i and j is measured by the dot product of normalized vectors (equivalent to cosine similarity on the unit hypersphere):

$$S_{i,j,l} = \hat{\mathbf{f}}_{i,l}^\top \hat{\mathbf{f}}_{j,l}. \quad (8)$$

We define the hyperspherical contrastive loss as

$$L_{\text{HC},l} = \max(\mathbb{E}_N[S_{i,j,l}] - \mathbb{E}_P[S_{i,j,l}] + \lambda, 0), \quad (9)$$

where λ is a margin, and $\mathbb{E}_P/\mathbb{E}_N$ denote the average similarity over positive/negative pairs. The overall objective combines restoration and contrastive terms:

$$L = L_1 + \alpha \cdot L_{\text{HC}}. \quad (10)$$

This joint optimization yields embeddings that are compact within class and well-separated across classes in angular space, producing more reliable routing scores and reducing cross-degradation interference. To the best of our knowledge, such hyperspherical (spherical-uniform) degradation embedding explicitly designed for expert routing remains underexplored in the image restoration literature, where representations are typically learned in unconstrained linear spaces.

3.4 Global–Local Granularity Fusion

While global representations provide scene-level consistency, restoration under spatially non-uniform degradations requires guidance that is both *content-aware* and *spatially localized*. In real scenarios, corruptions such as rain streaks, localized haze, or mixed degradations often appear only in certain regions. Moreover, patch-based augmentation during training and full-image inference at test time may introduce a granularity mismatch, where local degradation patterns observed in crops are not consistently reflected when processing the entire image.

To address this, we propose Global–Local Granularity Fusion (GLGF), which constructs a *content-guided, degradation-conditioned* restoration prior by aligning global content context with localized degradation evidence. As shown in Fig. 6, GLGF includes two components: (i) GLMC to construct two aligned patch-level maps (CSP/DSP), and (ii) CGDF to fuse them into a restoration prior map.




Fig. 6: Architecture of GLGF. GLGF consists of GLMC (Global-Local Map Construction) that derives a CSP map (content semantic patch map) and a DSP map (degradation severity patch map), and CGDF (Content-Guided Degradation Fusion) that fuses them into a *restoration prior map* for robust restoration under spatially varying degradations.

Global-Local Map Construction (GLMC). Given an input image x , a Content Extractor processes the full image and outputs a content CLS token $\mathbf{c}$ and content patch tokens $\mathbf{T}_c \in \mathbb{R}^{H \times W \times d}$. The global token $\mathbf{c}$ encodes scene semantics and is used to calibrate content patch tokens via a lightweight FiLM modulation, yielding the CSP map:

$$\gamma, \beta = \text{MLP}_c(\mathbf{c}), \quad \hat{\mathbf{T}}_c = \text{LN}(\gamma \odot \mathbf{T}_c + \beta). \quad (11)$$

This calibration stabilizes the content-side query representation across diverse scenes and texture statistics.

In parallel, we partition x into a fixed 4×4 grid of crops and feed each crop into a Degradation Extractor to obtain 16 crop-wise degradation CLS tokens $\{\mathbf{d}_i\}_{i=1}^{16}$. We spatially rearrange them into a coarse degradation map $\mathbf{E} \in \mathbb{R}^{4 \times 4 \times d}$ and upsample it to match the patch resolution, producing the DSP map $\mathbf{D} \in \mathbb{R}^{H \times W \times d}$. Importantly, $\mathbf{E}$ captures the *relative* regional distribution of degradation evidence; upsampling resamples this coarse distribution onto the patch grid to serve as a spatial prior for restoration rather than a pixel-accurate degradation estimation.

Content-Guided Degradation Fusion (CGDF). With the CSP/DSP maps constructed, we fuse calibrated content patches with the degradation prior tokens using cross-attention:

$$\mathbf{F}' = \hat{\mathbf{T}}_c + \text{Attn}(Q = \hat{\mathbf{T}}_c, K = \mathbf{D}, V = \mathbf{D}), \quad (12)$$

followed by a standard LN+FFN block. The resulting $\mathbf{F}'$ forms a *restoration prior map*, providing patch-level guidance signals for subsequent restoration. Finally, we inject this prior into the restoration backbone via a guided cross-attention mechanism. At layer l , we first align the prior $\mathbf{F}'$ to the dimension




Fig. 7: Visual comparison of **SLER-IR** with state-of-the-art methods. Zoom in for better view.

of the intermediate restoration features $\mathbf{F}_l$, and then use the aligned prior as queries with $\mathbf{F}_l$ as keys/values to produce a guided update, which is merged back to $\mathbf{F}_l$ through a lightweight output projection and a residual connection. This injection enables region-adaptive restoration by letting the prior selectively aggregate backbone responses under spatially varying degradations.

4 Experiment

This section presents the overall experimental analysis, including the experimental setup, quantitative comparisons with state-of-the-art methods, ablation studies, and qualitative visualizations to comprehensively validate the effectiveness of the proposed framework.

Table 1: Quantitative comparison of three-task all-in-one image restoration on SOTS [26] (dehazing), Rain100L [59] (deraining), and CBSD68 [42] (denoising). The best and second-best results are highlighted in red and blue, respectively.

Method	Reference	Dehazing on SOTS	Deraining on Rain100L	Denoising on CBSD68			Average (PSNR/SSIM)
		(PSNR/SSIM)	(PSNR/SSIM)	$\sigma=15$	$\sigma=25$	$\sigma=50$	
MPRNet	CVPR'21	28.00 / 0.958	33.86 / 0.958	33.27 / 0.920	30.76 / 0.871	27.29 / 0.761	30.63 / 0.894
Restormer	CVPR'22	27.78 / 0.958	33.78 / 0.958	33.72 / 0.865	30.67 / 0.865	27.63 / 0.792	30.75 / 0.901
Airnet	CVPR'22	27.94 / 0.962	34.90 / 0.967	33.92 / 0.933	31.26 / 0.888	28.00 / 0.797	31.20 / 0.910
IDR	CVPR'23	29.87 / 0.970	36.03 / 0.971	33.89 / 0.931	31.32 / 0.884	28.04 / 0.798	31.83 / 0.911
PromptIR	NeurIPS'23	30.58 / 0.974	36.37 / 0.972	33.98 / 0.933	31.31 / 0.888	28.06 / 0.799	32.06 / 0.913
InstructIR	ECCV'24	30.22 / 0.959	37.98 / 0.978	34.15 / 0.933	31.52 / 0.890	28.30 / 0.804	32.43 / 0.913
VLU-Net	CVPR'25	30.71 / 0.980	38.93 / 0.984	34.13 / 0.935	31.48 / 0.892	28.23 / 0.804	32.70 / 0.919
MoCEIR	CVPR'25	31.34 / 0.979	38.57 / 0.984	34.11 / 0.932	31.45 / 0.888	28.18 / 0.800	32.73 / 0.917
SLER-IR	Ours	33.00 / 0.982	38.34 / 0.983	34.25 / 0.936	31.68 / 0.897	28.45 / 0.813	33.14 / 0.922

4.1 Experimental Setup

Datasets and Evaluation. Following previous work [27, 46], we prepare datasets for five image degradation tasks, including denoising (BSD400 [2] and WED [41]), dehazing (SOTS [26]), deraining (Rain100L [59]), deblurring (GoPro [44]), and low-light enhancement (LOL [57]). Following the experimental protocols established in prior works [27, 46], we conduct evaluations under two distinct configurations: (a) the three-task all-in-one and (b) five-task all-in-one setting. For fair comparison, we adopt standard image quality metrics PSNR and SSIM [56] to evaluate restoration performance.

Implementation Details. Following the configuration of PromptIR [46], our **SLER-IR** architecture features a 4-level encoder-decoder structure, comprising varying numbers of Transformer blocks at each level, specifically [4, 6, 6, 8] from level-1 to level-4. We conduct our experiments using PyTorch on H20 GPUs. The training is divided into two stages. In Stage 1, we train for 15 epochs with a batch size of 10 under uniform data sampling, using softmax-based probabilistic routing where all expert branches receive probability-weighted gradients to jointly learn degradation-aware branch selection. A combined L1 and contrastive loss is used. The Adam optimizer ($\beta_1 = 0.9$, $\beta_2 = 0.999$) is employed with an initial learning rate of 2×10^{-4} and a cosine decay schedule. In Stage 2, we freeze the router and train for 80 epochs with a batch size of 20, adopting deterministic top-1 (hard) selection where only the selected branch is updated while the gradients of non-selected branches are masked, using only the L1 loss. Data augmentation is applied throughout training.

4.2 Qualitative and Quantitative Evaluation.

Three-tasks. As presented in Table 1, our method consistently outperforms both general restoration models (MPRNet [66] and Restormer [65]) and recent all-in-one approaches such as AirNet [27], PromptIR [46], and MoCE-IR [64]. Specifically, **SLER-IR** achieves the highest average performance of 33.14/0.922 in PSNR/SSIM [56], performs better than MoCE-IR by 0.41 dB and 0.005,

Table 2: Quantitative comparison on the five-task all-in-one image restoration setting, including dehazing (SOTS [26]), deraining (Rain100L [59]), denoising (CBSD68 [42]), deblurring (GoPro [44]), and low-light enhancement (LOL [57]). The best and second-best results are highlighted in red and blue, respectively.

Method	Reference	Dehazing on SOTS	Deraining on Rain100L	Denoising on CBSD68	Deblurring on Gopro	Low-light Enh. on LOL	Average
NAFNet	ECCV'22	25.23 / 0.939	35.56 / 0.967	31.02 / 0.883	26.32 / 0.808	20.49 / 0.809	27.60 / 0.881
FSNet	TPAMI'23	25.53 / 0.943	36.77 / 0.968	31.33 / 0.883	28.32 / 0.869	22.29 / 0.829	28.71 / 0.898
Restormer	ECCV'22	24.09 / 0.927	34.81 / 0.960	31.49 / 0.884	27.22 / 0.829	20.41 / 0.806	27.60 / 0.881
Transweather	CVPR'22	21.32 / 0.885	29.43 / 0.905	29.00 / 0.841	25.12 / 0.757	21.21 / 0.792	25.22 / 0.836
AirNet	CVPR'22	21.04 / 0.884	32.98 / 0.951	30.91 / 0.882	24.35 / 0.781	18.18 / 0.735	25.49 / 0.846
IDR	CVPR'23	25.24 / 0.943	35.63 / 0.965	31.60 / 0.887	27.87 / 0.846	21.34 / 0.826	28.34 / 0.893
PromptIR	NeurIPS'23	26.54 / 0.949	36.37 / 0.970	31.47 / 0.886	28.71 / 0.881	22.68 / 0.832	29.15 / 0.904
InstructIR	ECCV'24	26.14 / 0.973	37.10 / 0.966	31.44 / 0.887	29.40 / 0.886	23.00 / 0.836	29.55 / 0.907
VLU-Net	CVPR'25	30.84 / 0.980	38.54 / 0.982	31.43 / 0.891	27.46 / 0.840	22.29 / 0.833	30.11 / 0.905
MoCEIR	CVPR'25	30.48 / 0.974	38.04 / 0.982	31.34 / 0.887	30.05 / 0.899	23.00 / 0.852	30.58 / 0.919
SLER-IR	Ours	33.43 / 0.985	38.47 / 0.983	31.50 / 0.894	31.27 / 0.918	23.96 / 0.861	31.73 / 0.928

respectively. On the SOTS test set for dehazing, our approach improves PSNR by 1.66 dB over the baseline, and for denoising (CBSD68) it achieves consistent gains across all noise levels ($\sigma = 15, 25, 50$).

Five-tasks. As shown in Table 2, **SLER-IR** still achieves the best performance, which attains the highest overall average of 31.73/0.928 in PSNR/SSIM [56], surpassing the previous approach by a margin of 1.15 dB and 0.009, respectively. Notably, the proposed method achieves remarkable gains of +2.59 dB on the SOTS dataset for dehazing and +1.22 dB on GoPro for deblurring. Even on the challenging low-light enhancement task (LOL), it maintains a clear advantage with an improvement of +0.96 dB. These results demonstrate the effectiveness and robustness of the proposed framework across diverse restoration tasks.

Qualitative Results. Figure 7 shows visual comparisons on low-light enhancement, deblurring, dehazing, and deraining tasks. In low-light scenes, competing methods often produce insufficient brightness or lose fine structural details, whereas SLER-IR achieves balanced enhancement with well-preserved structures and natural illumination. For deblurring, existing approaches often leave noticeable residual blur, while our method restores clearer textures and sharper edges. In dehazing, other methods tend to retain haze artifacts and color distortions, whereas **SLER-IR** delivers cleaner results with more faithful color reconstruction. Under rainy conditions, the proposed model removes rain streaks more thoroughly while preserving underlying scene details. These visual results, together with the quantitative comparisons, further validate the effectiveness and robustness of our approach.

4.3 Ablation Studies

To validate the effectiveness of the proposed components, we conduct ablation experiments under the five-task setting. Additional results and analyses are provided in the supplementary material.

Table 3: Ablation study on loss functions. **Table 4:** Ablation study on different model components.

Loss	Branch Sel.	PSNR	SSIM	Model	PSNR	SSIM
L_1	Classifier	28.30	0.879	w/o GLGF	30.64	0.917
$L_1 + L_{\text{cls}}$	Classifier	28.76	0.884	SLER-IR	31.73	0.928
$L_1 + L_{\text{HC}}$	Ours	31.73	0.928			

Table 5: Ablation study on the number of experts per layer. **Table 6:** Ablation study on hyperspherical contrastive loss weight (α).

C	PSNR	SSIM	α	PSNR	SSIM
2	31.22	0.924	0.5	31.03	0.923
3	31.73	0.928	1	31.73	0.928
4	31.32	0.926	2	31.15	0.924

Effectiveness of the proposed losses. Table 3 evaluates the impact of different training objectives. Starting from the baseline model trained with the restoration loss (L_1), introducing the auxiliary classifier loss improves performance by providing additional degradation-aware supervision. Further incorporating the proposed hyperspherical contrastive loss yields the best PSNR and SSIM results. This demonstrates that contrastive degradation representation learning helps enhance the discriminability of degradation embeddings, leading to more reliable expert routing and improved restoration performance.

Impact of Global–Local Granularity Fusion (GLGF). Table 4 evaluates the effectiveness of the GLGF module. As shown in results, integrating GLGF consistently improves restoration performance compared with the baseline model without this module. This improvement indicates that jointly modeling global semantic cues (CLS tokens) and local degradation evidence (patch-level cues) provides more informative restoration guidance via the learned restoration prior.

Ablation of Expert Number per Layer. Table 5 examines the effect of experts per layer ($C \in 2, 3, 4$). Under identical settings, $C=3$ achieves the best trade-off between specialization and efficiency. Fewer experts limit discrimination, while more experts increase computation without clear gains.

Ablation of Hyperspherical Contrastive Loss Weight. Table 6 analyzes the influence of the hyperspherical contrastive loss weight α . When the weight is too small ($\alpha=0.5$), the degradation embedding receives insufficient supervision, resulting in less discriminative routing representations. Conversely, an excessively large weight ($\alpha=2$) overemphasizes contrastive optimization and weakens the restoration objective. Empirically, $\alpha=1$ achieves the best balance between degradation representation learning and restoration performance, and is therefore used as the default setting.

5 Conclusion

In this paper, we propose **SLER-IR**, a spherical layer-wise expert routing framework for unified all-in-one image restoration. By leveraging hyperspherical degradation embedding for geometry-consistent routing and a global–local granularity fusion module for spatially robust guidance, our model enables progressive expert specialization and improves adaptability to complex degradation scenarios. Moreover, the routing mechanism allows flexible compositional inference paths without increasing inference overhead. Extensive experiments on multiple benchmarks show that our approach consistently outperforms state-of-the-art methods in both quantitative metrics and visual quality.

Acknowledgements. This work was supported in part by the National Natural Science Foundation of China under Grant 62171304, the Local Scientific and Technological Development Funding Projects for Central guidance in Sichuan under Grant 2024ZYD0004, the Sichuan Province Science and Technology Support Program under Grant 2025JDRC0011, the Science and Technology Project of Xizang Autonomous Region under Grant XZ202601ZY0092, and the Young Faculty Technology Innovation Capacity Enhancement Program of Sichuan University under Grant 2024SCUQJTX025.

References

1. Ai, Y., Huang, H., He, R.: Lora-ir: taming low-rank experts for efficient all-in-one image restoration. arXiv preprint arXiv:2410.15385 (2024)
2. Arbelaez, P., Maire, M., Fowlkes, C., Malik, J.: Contour detection and hierarchical image segmentation. *IEEE transactions on pattern analysis and machine intelligence* **33**(5), 898–916 (2010)
3. Cai, B., Xu, X., Jia, K., Qing, C., Tao, D.: Dehazenet: An end-to-end system for single image haze removal. *IEEE transactions on image processing* **25**(11), 5187–5198 (2016)
4. Cai, Y., Bian, H., Lin, J., Wang, H., Timofte, R., Zhang, Y.: Retinexformer: One-stage retinex-based transformer for low-light image enhancement. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 12504–12513 (2023)
5. Chen, H., Wang, Y., Guo, T., Xu, C., Deng, Y., Liu, Z., Ma, S., Xu, C., Xu, C., Gao, W.: Pre-trained image processing transformer. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 12299–12310 (2021)
6. Chen, L., Chu, X., Zhang, X., Sun, J.: Simple baselines for image restoration. In: *European conference on computer vision*. pp. 17–33. Springer (2022)
7. Chen, W.T., Fang, H.Y., Ding, J.J., Kuo, S.Y.: Pmhld: Patch map-based hybrid learning dehazenet for single image haze removal. *IEEE Transactions on Image Processing* **29**, 6773–6788 (2020)
8. Chen, W.T., Fang, H.Y., Hsieh, C.L., Tsai, C.C., Chen, I., Ding, J.J., Kuo, S.Y., 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)

9. Cho, S.J., Ji, S.W., Hong, J.P., Jung, S.W., Ko, S.J.: Rethinking coarse-to-fine approach in single image deblurring. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 4641–4650 (2021)
10. Conde, M.V., Geigle, G., Timofte, R.: Instructir: High-quality image restoration following human instructions. In: European Conference on Computer Vision. pp. 1–21. Springer (2024)
11. Cui, Y., Ren, W., Cao, X., Knoll, A.: Image restoration via frequency selection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **46**(2), 1093–1108 (2023)
12. Dong, C., Loy, C.C., He, K., Tang, X.: Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence* **38**(2), 295–307 (2015)
13. Duan, H., Min, X., Wu, S., Shen, W., Zhai, G.: Uniprocessor: a text-induced unified low-level image processor. In: European Conference on Computer Vision. pp. 180–199. Springer (2024)
14. Gu, Y., Meng, Y., Chen, S., Ji, J., Sun, X., Ruan, W., Ji, R.: SFIR: Optimizing spatial and frequency domains for image restoration. *Pattern Recognition* **171**, 112188 (2026). <https://doi.org/10.1016/j.patcog.2025.112188>
15. Gu, Y., Meng, Y., Ji, J., Sun, X.: ACL: Activating capability of linear attention for image restoration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17913–17923 (2025)
16. Gu, Y., Meng, Y., Sun, X., Ji, J., Ji, R.: Boosting filter optimization and prompt-guided decoding for mixed degradation image restoration. *International Journal of Computer Vision* **134**(6), 280 (2026). <https://doi.org/10.1007/s11263-026-02832-0>
17. Gu, Y., Meng, Y., Sun, X., Ji, J., Ruan, W., Ji, R.: Mixed degradation image restoration via local dynamic optimization and conditional embedding. *arXiv preprint arXiv:2411.16217* (2024)
18. Gu, Y., Meng, Y., Zheng, K., Sun, X., Ji, J., Ruan, W., Cao, L., Ji, R.: An efficient and mixed heterogeneous model for image restoration. *arXiv preprint arXiv:2504.10967* (2025)
19. Guo, C., Li, C., Guo, J., Loy, C.C., Hou, J., Kwong, S., Cong, R.: Zero-reference deep curve estimation for low-light image enhancement. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1780–1789 (2020)
20. Guo, X., Li, Y., Ling, H.: Lime: Low-light image enhancement via illumination map estimation. *IEEE Transactions on image processing* **26**(2), 982–993 (2016)
21. Huang, J., Liu, X., Pan, Y., He, X., Ren, C.: Casapynet: channel affine self-attention-based progressively updated network for real image denoising. *IEEE Transactions on Industrial Informatics* **19**(8), 9145–9156 (2022)
22. Huang, T., Li, S., Jia, X., Lu, H., Liu, J.: Neighbor2neighbor: Self-supervised denoising from single noisy images. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 14781–14790 (2021)
23. Jiang, K., Wang, Z., Yi, P., Chen, C., Huang, B., Luo, Y., Ma, J., Jiang, J.: Multi-scale progressive fusion network for single image deraining. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 8346–8355 (2020)
24. Kupyn, O., Budzan, V., Mykhailych, M., Mishkin, D., Matas, J.: Deblurgan: Blind motion deblurring using conditional adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8183–8192 (2018)

25. Kupyn, O., Martyniuk, T., Wu, J., Wang, Z.: Deblurgan-v2: Deblurring (orders-of-magnitude) faster and better. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 8878–8887 (2019)
26. Li, B., Ren, W., Fu, D., Tao, D., Feng, D., Zeng, W., Wang, Z.: Benchmarking single-image dehazing and beyond. *IEEE transactions on image processing* **28**(1), 492–505 (2018)
27. Li, B., Liu, X., Hu, P., Wu, Z., Lv, J., Peng, X.: 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)
28. Li, R., Cheong, L.F., Tan, R.T.: Heavy rain image restoration: Integrating physics model and conditional adversarial learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1633–1642 (2019)
29. Li, R., Tan, R.T., Cheong, L.F.: All in one bad weather removal using architectural search. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3175–3185 (2020)
30. Li, Z., Lei, Y., Ma, C., Zhang, J., Shan, H.: Prompt-in-prompt learning for universal image restoration. *arXiv preprint arXiv:2312.05038* (2023)
31. Liang, J., Cao, J., Sun, G., Zhang, K., Van Gool, L., Timofte, R.: Swinir: Image restoration using swin transformer. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 1833–1844 (2021)
32. Lim, B., Son, S., Kim, H., Nah, S., Mu Lee, K.: Enhanced deep residual networks for single image super-resolution. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 136–144 (2017)
33. Lin, X., Ren, C., Liu, X., Huang, J., Lei, Y.: Unsupervised image denoising in real-world scenarios via self-collaboration parallel generative adversarial branches. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12642–12652 (2023)
34. Lin, X., Yue, J., Chan, K.C., Qi, L., Ren, C., Pan, J., Yang, M.H.: Multi-task image restoration guided by robust dino features. *arXiv preprint arXiv:2312.01677* (2023)
35. Lin, X., Yue, J., Ding, S., Ren, C., Qi, L., Yang, M.H.: Dual degradation representation for joint deraining and low-light enhancement in the dark. *IEEE transactions on circuits and systems for video technology* (2024)
36. Lin, X., Zhou, Y., Yue, J., Ren, C., Chan, K.C., Qi, L., Yang, M.H.: Re-boosting self-collaboration parallel prompt gan for unsupervised image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2025)
37. Liu, X., Ma, Y., Shi, Z., Chen, J.: Griddehazenet: Attention-based multi-scale network for image dehazing. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 7314–7323 (2019)
38. Liu, Y.F., Jaw, D.W., Huang, S.C., Hwang, J.N.: Desnownet: Context-aware deep network for snow removal. *IEEE Transactions on Image Processing* **27**(6), 3064–3073 (2018)
39. Luo, Y., Zhao, R., Wei, X., Chen, J., Lu, Y., Xie, S., Wang, T., Xiong, R., Lu, M., Zhang, S.: Wm-moe: Weather-aware multi-scale mixture-of-experts for blind adverse weather removal. *arXiv preprint arXiv:2303.13739* (2023)
40. Luo, Z., Gustafsson, F.K., Zhao, Z., Sjölund, J., Schön, T.B.: Controlling vision-language models for multi-task image restoration. *arXiv preprint arXiv:2310.01018* (2023)
41. Ma, K., Duanmu, Z., Wu, Q., Wang, Z., Yong, H., Li, H., Zhang, L.: Waterloo exploration database: New challenges for image quality assessment models. *IEEE Transactions on Image Processing* **26**(2), 1004–1016 (2016)

42. Martin, D., Fowlkes, C., Tal, D., Malik, J.: 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. vol. 2, pp. 416–423. IEEE (2001)
43. Moran, S., Marza, P., McDonagh, S., Parisot, S., Slabaugh, G.: Deeplpf: Deep local parametric filters for image enhancement. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12826–12835 (2020)
44. Nah, S., Hyun Kim, T., Mu Lee, K.: 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)
45. Park, D., Kang, D.U., Kim, J., Chun, S.Y.: Multi-temporal recurrent neural networks for progressive non-uniform single image deblurring with incremental temporal training. In: European conference on computer vision. pp. 327–343. Springer (2020)
46. Potlapalli, V., Zamir, S.W., Khan, S.H., Shahbaz Khan, F.: Promptir: Prompting for all-in-one image restoration. *Advances in Neural Information Processing Systems* **36**, 71275–71293 (2023)
47. Qin, C.J., Wu, R.Q., Liu, Z., Lin, X., Guo, C.L., Park, H.H., Li, C.: Restore anything with masks: Leveraging mask image modeling for blind all-in-one image restoration. In: European Conference on Computer Vision. pp. 364–380. Springer (2024)
48. Qin, X., Wang, Z., Bai, Y., Xie, X., Jia, H.: Ffa-net: Feature fusion attention network for single image dehazing. In: Proceedings of the AAAI conference on artificial intelligence. vol. 34, pp. 11908–11915 (2020)
49. Ren, D., Zuo, W., Hu, Q., Zhu, P., Meng, D.: Progressive image deraining networks: A better and simpler baseline. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3937–3946 (2019)
50. Song, Y., He, Z., Qian, H., Du, X.: Vision transformers for single image dehazing. *IEEE Transactions on Image Processing* **32**, 1927–1941 (2023)
51. Tang, N., Luo, X., Cheng, Z., Zhou, L., Zhang, D., Qu, Y.: Diffusion once and done: Degradation-aware lora for all-in-one image restoration. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 40, pp. 9448–9456 (2026). <https://doi.org/10.1609/aaai.v40i11.37905>
52. Tian, Y., Han, J., Chen, H., Xi, Y., Ding, N., Hu, J., Xu, C., Wang, Y.: Instruct-ipt: All-in-one image processing transformer via weight modulation. *arXiv preprint arXiv:2407.00676* (2024)
53. Valanarasu, J.M.J., Yasarla, R., Patel, V.M.: 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)
54. Wang, G., Sun, C., Sowmya, A.: Erl-net: Entangled representation learning for single image de-raining. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5644–5652 (2019)
55. Wang, S., Wang, Y., Hu, B., Zhang, L., Zhang, X.P., Wei, M.: M2IR: Proactive all-in-one image restoration via mamba-style modulation and mixture-of-experts. *arXiv preprint arXiv:2603.14816* (2026)
56. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**(4), 600–612 (2004)
57. Wei, C., Wang, W., Yang, W., Liu, J.: Deep retinex decomposition for low-light enhancement. *arXiv preprint arXiv:1808.04560* (2018)

58. Wu, W., Weng, J., Zhang, P., Wang, X., Yang, W., Jiang, J.: Uretinex-net: Retinex-based deep unfolding network for low-light image enhancement. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5901–5910 (2022)
59. Yang, W., Tan, R.T., Feng, J., Liu, J., Guo, Z., Yan, S.: Deep joint rain detection and removal from a single image. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1357–1366 (2017)
60. Yang, Y., Wang, X., Lin, X., Chen, H.: Demnet: A degradation difference enabled multi-stage network for multiple degradation image restoration. *Knowledge-Based Systems* **318**, 113426 (2025)
61. Yang, Z., Chen, H., Qian, Z., Yi, Y., Zhang, H., Zhao, D., Wei, B., Xu, Y.: All-in-one medical image restoration via task-adaptive routing. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 67–77. Springer (2024)
62. Yu, X., Zhou, S., Li, H., Zhu, L.: Multi-expert adaptive selection: Task-balancing for all-in-one image restoration. *IEEE Transactions on Circuits and Systems for Video Technology* (2024)
63. Zamfir, E., Wu, Z., Mehta, N., Paudel, D.P., Zhang, Y., Timofte, R.: Efficient degradation-aware any image restoration. arXiv preprint arXiv:2405.15475 (2024)
64. Zamfir, E., Wu, Z., Mehta, N., Tan, Y., Paudel, D.P., Zhang, Y., Timofte, R.: Complexity experts are task-discriminative learners for any image restoration. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 12753–12763 (2025)
65. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: 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)
66. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H., Shao, L.: Multi-stage progressive image restoration. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 14821–14831 (2021)
67. Zeng, H., Wang, X., Chen, Y., Su, J., Liu, J.: 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)
68. Zhang, J., Huang, J., Yao, M., Yang, Z., Yu, H., Zhou, M., Zhao, F.: Ingredient-oriented multi-degradation learning for image restoration. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5825–5835 (2023)
69. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing* **26**(7), 3142–3155 (2017)
70. Zhang, K., Zuo, W., Gu, S., Zhang, L.: Learning deep cnn denoiser prior for image restoration. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3929–3938 (2017)
71. Zhang, K., Zuo, W., Zhang, L.: Ffdnet: Toward a fast and flexible solution for cnn-based image denoising. *IEEE Transactions on Image Processing* **27**(9), 4608–4622 (2018)
72. Zhang, X., Ma, J., Wang, G., Zhang, Q., Zhang, H., Zhang, L.: Perceive-ir: Learning to perceive degradation better for all-in-one image restoration. *IEEE Transactions on Image Processing* (2025)