Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

DA-LightNet: Real-Time Rain and Haze Removal on Edge Devices

Jianhong Li
Sichuan Vocational and Technical College
Yan Liang
Sichuan Vocational and Technical College
Miaoqin Yuan
Sichuan Vocational and Technical College
Fuyu Deng
Sichuan Vocational and Technical College
Ke Fang
Sichuan Vocational and Technical College
Guangqiang He
lijianhongsc@gmail.com

Sichuan Vocational and Technical College

Article

Keywords: image restoration, image deraining, image dehazing, lightweight network, real-time inference, edge devices

Posted Date: September 14th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10148896/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

DA-LightNet: Real-Time Rain and Haze Removal on Edge Devices

Jianhong Li · Yan Liang · Miaoqin Yuan · Fuyu Deng · Ke Fang · Guangqiang He

the date of receipt and acceptance should be inserted later

Abstract Images captured in real-world environments are often degraded by adverse weather conditions such as rain and haze, which reduce visual quality and negatively affect downstream vision tasks. Although recent deep learning based restoration methods have achieved promising results, many of them rely on heavy architectures and high computational cost, making them difficult to deploy on resource-constrained edge devices. To address this issue, this paper proposes DA-LightNet, a lightweight degradation-aware image restoration network for real-time rain and haze removal on edge devices. The proposed network adopts a unified restoration architecture for both deraining and dehazing, avoiding the need to deploy separate models for different weather degradations. A degradation-aware modulation module is designed to adaptively recalibrate intermediate features according to rain and haze characteristics, while an efficient restoration block based on depthwise and pointwise convolutions is introduced to reduce redundant computation. In addition, a lightweight reconstruction head is used to further reduce inference overhead. Experiments are conducted on the paired LHP-Rain and LMHaze datasets. The results show that DA-LightNet achieves competitive restoration quality with only 4.6M parameters and 21.5G FLOPs, while reaching 34.8 FPS on a Jetson edge platform. Compared with representative restoration methods, DA-LightNet provides a more favorable balance between restoration accuracy and real-time deployment efficiency, demonstrating its potential for practical edge-side image restoration applications.

Keywords image restoration · image deraining · image dehazing · lightweight network · real-time inference · edge devices

Jianhong Li
School of Electronics and the Internet of Things, Sichuan Vocational and Technical College, Suining, China
E-mail: 516331055@qq.com

Yan Liang
School of Electronics and the Internet of Things, Sichuan Vocational and Technical College, Suining, China
E-mail: 93447867@qq.com

Miaoqin Yuan
School of Electronics and the Internet of Things, Sichuan Vocational and Technical College, Suining, China
E-mail: 23985282@qq.com

Fuyu Deng
School of Electronics and the Internet of Things, Sichuan Vocational and Technical College, Suining, China
E-mail: 154895591@qq.com

Ke Fang
School of Electronics and the Internet of Things, Sichuan Vocational and Technical College, Suining, China
E-mail: 799335999@qq.com

Guangqiang He (**Corresponding author**)
School of Electronics and the Internet of Things, Sichuan Vocational and Technical College, Suining, China
E-mail: lijianhongsc@gmail.com

1 Introduction

Images captured in real-world outdoor environments are often degraded by adverse weather conditions such as rain and haze. Rain streaks introduce local high-frequency interference and may occlude important scene details, while haze reduces image contrast, weakens visibility, and causes color distortion [13, 18]. These degradations not only reduce visual quality, but also affect downstream vision tasks, including object detection, video surveillance, autonomous navigation, and intelligent transportation systems [21]. Therefore, effective

rain and haze removal is important for both image enhancement and practical visual perception systems.

In recent years, deep learning based image deraining and dehazing methods have achieved significant progress. CNN-based methods, progressive restoration networks, attention mechanisms, and transformer-based architectures have improved restoration quality on benchmark datasets [6, 16, 22–24]. For image dehazing, prior-based methods have gradually been replaced or enhanced by end-to-end CNN and transformer models, leading to stronger feature representation and better reconstruction capability [3, 9, 12, 15, 17]. However, many high-performance methods rely on large model capacity, complex attention operations, or repeated high-resolution feature interactions. As a result, they are difficult to deploy on resource-constrained edge devices where memory, computation, and power consumption are limited.

In practical applications, image restoration models are often expected to run on mobile phones, intelligent cameras, unmanned aerial vehicles, and embedded boards. In these scenarios, restoration quality alone is not sufficient. Model size, floating-point operations, memory usage, inference latency, and frames per second are also important factors [10, 20]. Moreover, rain and haze usually require different restoration behaviors. Rain removal needs to suppress local streak-like structures while preserving image details, whereas haze removal requires global contrast recovery and color correction. Deploying separate models for different degradations increases storage cost and complicates edge-side system design. Therefore, a unified and lightweight restoration model is needed for real-time adverse-weather image restoration.

To address these issues, this paper proposes DA-LightNet, a lightweight degradation-aware image restoration network for real-time rain and haze removal on edge devices. The proposed method adopts a unified restoration architecture for both deraining and dehazing, avoiding the need to deploy multiple task-specific models. A degradation-aware modulation module is introduced to adaptively recalibrate intermediate features according to different weather degradation characteristics. In addition, efficient restoration blocks based on depthwise and pointwise convolutions are designed to reduce redundant computation, and a lightweight reconstruction head is used to further lower inference overhead. The proposed method is evaluated on the paired real-world LHP-Rain and LMHaze datasets, and both restoration quality and edge-side inference efficiency are reported.

The main contributions of this paper are summarized as follows:

- We propose DA-LightNet, a lightweight unified restoration network for real-time rain and haze removal on resource-constrained edge devices.
- We design a degradation-aware modulation module to adaptively enhance restoration features for different weather degradations without introducing separate deraining and dehazing branches.
- We introduce efficient restoration blocks and a lightweight reconstruction head to reduce computational complexity while preserving image structures and local details.
- We conduct experiments on LHP-Rain and LMHaze, and evaluate the proposed method from both restoration quality and deployment efficiency perspectives.

The remainder of this paper is organized as follows. Section 2 reviews related work. Section **??** introduces the proposed DA-LightNet. Section 4 describes the datasets and evaluation protocol. Section 5 presents the training strategy. Section **??** reports the experimental results and analysis. Finally, Section 7 concludes this paper.

2 Related Work

2.1 Image Deraining and Dehazing

Single image deraining aims to recover a clean image from a rainy observation. Early CNN-based methods mainly learned rain-specific representations from paired data. DDN removed rain streaks by learning deep detail features [6]. JORDER introduced joint rain detection and removal to improve rain-aware restoration [22]. PReNet further adopted a progressive recurrent strategy and became an effective baseline for single image deraining [16]. Later, MPRNet improved restoration quality through a multi-stage progressive architecture [23]. Restormer introduced an efficient transformer structure for high-resolution image restoration [24]. Although these methods achieve strong restoration quality, many of them still rely on relatively heavy architectures or expensive feature interactions.

Single image dehazing has also been widely studied. Traditional methods usually rely on physical imaging models and handcrafted priors. The dark channel prior is a representative prior-based method for haze removal [9]. With the development of deep learning, DehazeNet learned haze-related features in an end-to-end manner [3]. AOD-Net simplified the dehazing pipeline by directly learning a reformulated image formation model [12]. FFA-Net used feature fusion attention to enhance haze-relevant features [15]. More recently, DehazeFormer investigated transformer designs for single

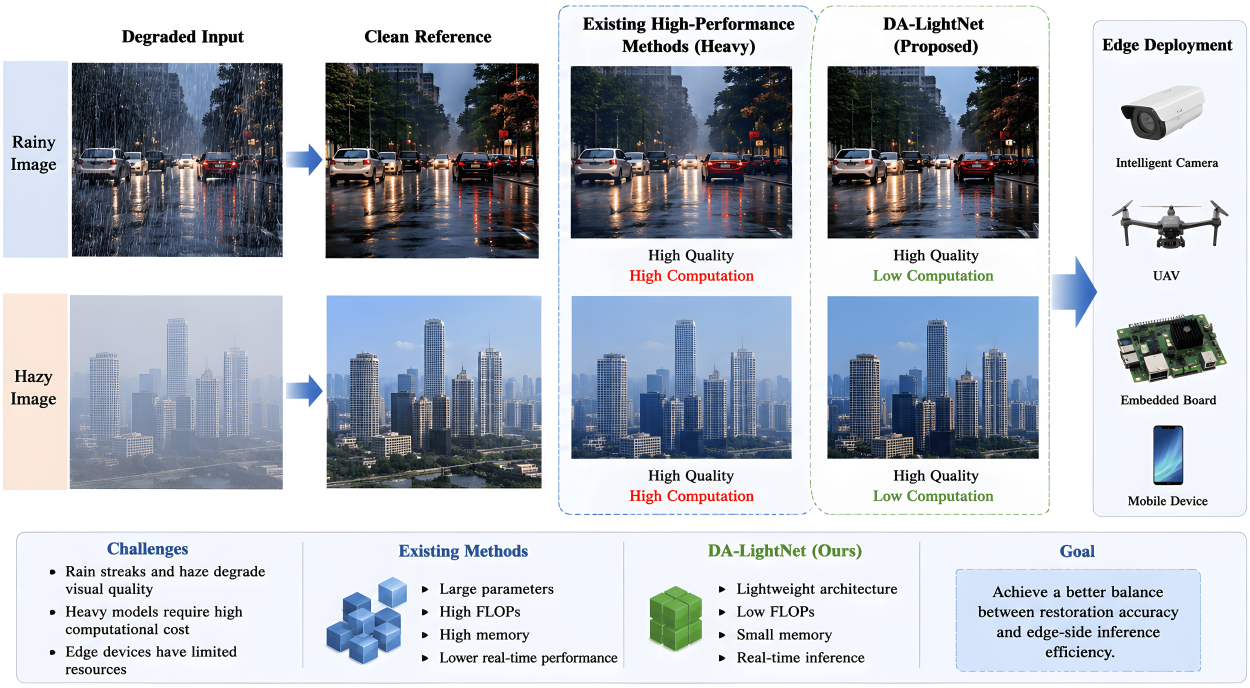


Fig. 1 Motivation of the proposed real-time rain and haze removal framework. Existing high-performance restoration models often obtain strong visual quality but require high computational cost. In contrast, DA-LightNet aims to achieve a better balance between restoration accuracy and edge-side inference efficiency.

image dehazing [17]. DEA-Net further improved dehazing performance through detail-enhanced convolution and content-guided attention [5]. However, these methods are mainly optimized for restoration accuracy, while their efficiency on low-computing edge devices is not always fully considered.

2.2 Real-World Weather Restoration Datasets

Datasets are important for training and evaluating weather image restoration models. For dehazing, RESIDE provides a systematic benchmark and has been widely used in single image dehazing research [13]. O-HAZE supports outdoor real-world dehazing evaluation with paired hazy and haze-free images [2]. I-HAZE provides paired indoor hazy and haze-free images for real-scene dehazing evaluation [1]. For deraining, SPA-Data provides high-quality real rain image pairs for practical rain removal evaluation [19]. More recently, LHP-Rain introduced a large-scale paired real rain benchmark for real-world deraining [8]. LMHaze provides paired real hazy and haze-free images with multiple haze intensities [25]. In this paper, LHP-Rain and LMHaze are used to evaluate DA-LightNet under realistic rain and haze degradation conditions.

2.3 Lightweight and Real-Time Image Restoration

Lightweight and real-time image restoration is essential for practical edge vision systems. Resource-constrained devices, such as mobile phones, smart cameras, UAVs, and embedded boards, require restoration models with low latency, small memory usage, and limited computational cost. Yang and Evans studied dehazing methods for resource-constrained platforms and emphasized deployment efficiency [20]. Jeong et al. proposed an end-to-end real-time single image dehazing method, showing the value of low-latency restoration in practical scenarios [10]. Zhou et al. further demonstrated that restoration algorithms for mobile devices should be evaluated from both visual quality and real deployment performance [26].

Unified restoration methods are also related to our work because rain and haze can be regarded as different adverse-weather degradations. AirNet handled unknown corruption within a unified image restoration framework [11]. PromptIR introduced prompt learning for all-in-one blind image restoration [14]. TransWeather focused on restoring images degraded by multiple adverse weather conditions [18]. However, many unified restoration methods still involve relatively heavy structures or are not primarily designed for real-time edge deployment. Different from these methods, DA-LightNet focuses on balancing restoration quality and

deployment efficiency through degradation-aware modulation and

3 Edge-Oriented DA-LightNet

3.1 Design Requirements

The design requirements of DA-LightNet are determined by the practical constraints of real-time image restoration on edge devices. Unlike cloud-based restoration systems, edge-side models must process degraded images locally under limited computation, memory, and power budgets. Therefore, the proposed network is designed by considering the following requirements.

Unified restoration architecture. Rain and haze are two representative weather degradations with different visual characteristics. Rain mainly introduces local streak-like interference and high-frequency occlusion, while haze causes global contrast degradation and color distortion. To avoid deploying separate models for different degradations, DA-LightNet adopts a unified restoration architecture for both deraining and dehazing. This design reduces storage cost and simplifies edge-side deployment.

Lightweight computation. Image restoration is a dense prediction task that requires pixel-level reconstruction, which is usually more computationally expensive than image classification or detection. For resource-constrained platforms, heavy backbones, large attention modules, and repeated high-resolution feature processing can lead to high latency and memory consumption. Therefore, the proposed model should maintain a compact structure with low parameters and FLOPs while preserving sufficient restoration ability.

Real-time inference. In practical applications such as intelligent cameras, UAV vision, mobile photography, and embedded surveillance, restoration models are expected to process continuous frames with low latency. Thus, FPS and inference latency are considered as key evaluation metrics, in addition to PSNR and SSIM. The model should achieve a favorable balance between restoration quality and real-time performance.

Real-world generalization. The model should not rely on handcrafted assumptions about rain streak direction, rain density, haze thickness, illumination condition, or scene category. Instead, it should learn degradation-aware representations from realistic paired datasets. In this work, LHP-Rain and LMHaze are used to evaluate whether the proposed method can handle real-world rain and haze degradations.

Deployment feasibility. The proposed model should be compatible with common edge inference toolchains. In addition to the original PyTorch implementation, the model should support deployment-oriented optimization, such as ONNX export, mixed-precision inference, or TensorRT acceleration. This requirement is important for reducing inference latency and improving hardware utilization on edge platforms.

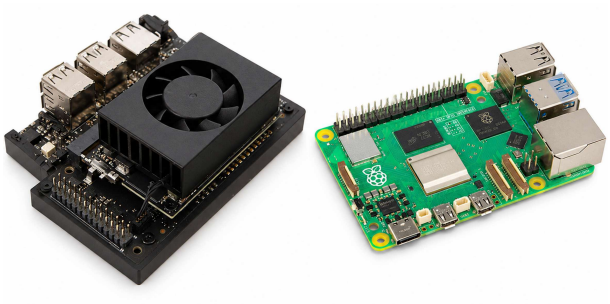


Fig. 2 Edge devices used for deployment evaluation. The NVIDIA Jetson platform is used as a GPU-accelerated edge AI device, while Raspberry Pi is used as a low-computing CPU-based embedded platform.

3.2 Device Selection

To evaluate the practical deployment capability of DA-LightNet, two representative edge devices are selected: an NVIDIA Jetson platform and a Raspberry Pi platform. These two devices represent different levels of edge computing capability and provide a comprehensive setting for analyzing the speed–accuracy trade-off of the proposed lightweight restoration model.

The NVIDIA Jetson platform is selected as the main edge AI device because it provides GPU acceleration and supports efficient deep learning inference optimization. Compared with desktop GPUs, Jetson devices have stricter constraints in terms of power consumption, memory capacity, and thermal design, but they are still capable of running neural networks in real time. Therefore, testing on this platform can directly reflect whether the proposed method is suitable for GPU-accelerated edge vision applications, such as smart cameras, UAVs, robots, and embedded monitoring systems.

The Raspberry Pi platform is selected as a more resource-constrained device. Unlike the Jetson platform, Raspberry Pi mainly relies on a low-power ARM CPU and has limited parallel acceleration capability. This makes it a challenging platform for image restoration, since deraining and dehazing require dense pixel-wise prediction. If a model can still maintain acceptable latency and memory usage on Raspberry Pi, it indicates stronger lightweight characteristics and better deployment potential for low-cost embedded systems.

The combination of these two platforms provides a more complete deployment evaluation. The Jetson plat-

form is used to verify real-time performance under edge GPU acceleration, while Raspberry Pi is used to examine the feasibility of the proposed model under stricter CPU-based computing constraints. In the experiments, we report FPS, inference latency, memory usage, and model size on these devices, together with restoration metrics such as PSNR and SSIM, to evaluate both visual quality and practical efficiency.

3.3 Overall Architecture

The overall architecture of the proposed DA-LightNet is shown in Fig. 3. The network is designed as a compact end-to-end restoration framework for both rain and haze removal. Given a degraded input image $I_{in} \in \mathbb{R}^{H\times W\times 3}$, DA-LightNet directly predicts the restored image $I_{out} \in \mathbb{R}^{H\times W\times 3}$ without using separate deraining and dehazing subnetworks. This unified design reduces model storage and avoids redundant feature extraction, which is important for real-time deployment on resource-constrained edge devices.

Specifically, the input image is first mapped into a low-dimensional feature space by a shallow feature extraction module. A 3×3 convolution is used to extract local texture and edge information while keeping the number of channels small. This process can be formulated as

$$F_0 = \phi_s(I_{in}), \tag{1}$$

where $\phi_s(\cdot)$ denotes the shallow feature extraction function and F_0 is the initial feature representation.

After shallow feature extraction, F_0 is fed into the degradation-aware restoration backbone. The backbone is composed of a degradation-aware feature modulation module and a series of efficient restoration blocks. The modulation module estimates degradation-related responses from the input features and adaptively adjusts the restoration process for different weather conditions. For rainy images, the network is encouraged to suppress local high-frequency rain streaks. For hazy images, the network focuses more on global contrast recovery and color correction. The restoration process can be written as

$$F_i = \mathcal{B}_i(F_{i-1}, M_d), \quad i = 1, 2, \ldots, N, \tag{2}$$

where $\mathcal{B}_i(\cdot)$ represents the i-th efficient restoration block, N is the number of restoration blocks, and M_d denotes the degradation-aware modulation information.

To improve information flow and avoid detail loss, a long skip connection is introduced between the shallow feature and the output of the restoration backbone. This design allows the network to learn residual degradation components instead of reconstructing the whole image from scratch, which improves convergence and reduces unnecessary computation. The fused feature is defined as

$$F_{out} = F_N + F_0. \tag{3}$$

Finally, a lightweight reconstruction head maps the restored feature back to the image space:

$$I_{out} = \phi_r(F_{out}), \tag{4}$$

where $\phi_r(\cdot)$ denotes the reconstruction head. In this work, the reconstruction head is implemented with lightweight convolutional layers to reduce inference latency.

Compared with heavy restoration networks that rely on large attention modules or repeated high-resolution feature interactions, DA-LightNet follows three lightweight principles. First, it uses a unified backbone for rain and haze removal, avoiding the need to deploy two independent models. Second, efficient restoration blocks are adopted to reduce spatial convolution cost. Third, degradation-aware modulation is introduced with limited additional parameters, allowing the network to adapt to different weather degradations without significantly increasing computational complexity. These designs make DA-LightNet suitable for real-time image restoration on edge devices.

3.3.1 Degradation-Aware Modulation

Rain and haze have different degradation characteristics. Rain usually appears as local high-frequency streaks and occlusions, while haze mainly causes global contrast reduction and color distortion. If the same feature transformation is used for both degradations, the network may fail to adaptively emphasize degradation-specific information. To address this problem, we introduce a degradation-aware modulation module, which estimates the degradation tendency from intermediate features and adaptively modulates the restoration backbone.

Given an intermediate feature map $F \in \mathbb{R}^{H\times W\times C}$, global average pooling is first applied to obtain a compact degradation descriptor:

$$z = \mathrm{GAP}(F), \tag{5}$$

where $z \in \mathbb{R}^C$ summarizes the global degradation distribution of the input feature. Then, a lightweight degradation estimator is used to predict degradation-aware responses:

$$s = \sigma(\mathrm{MLP}(z)), \tag{6}$$

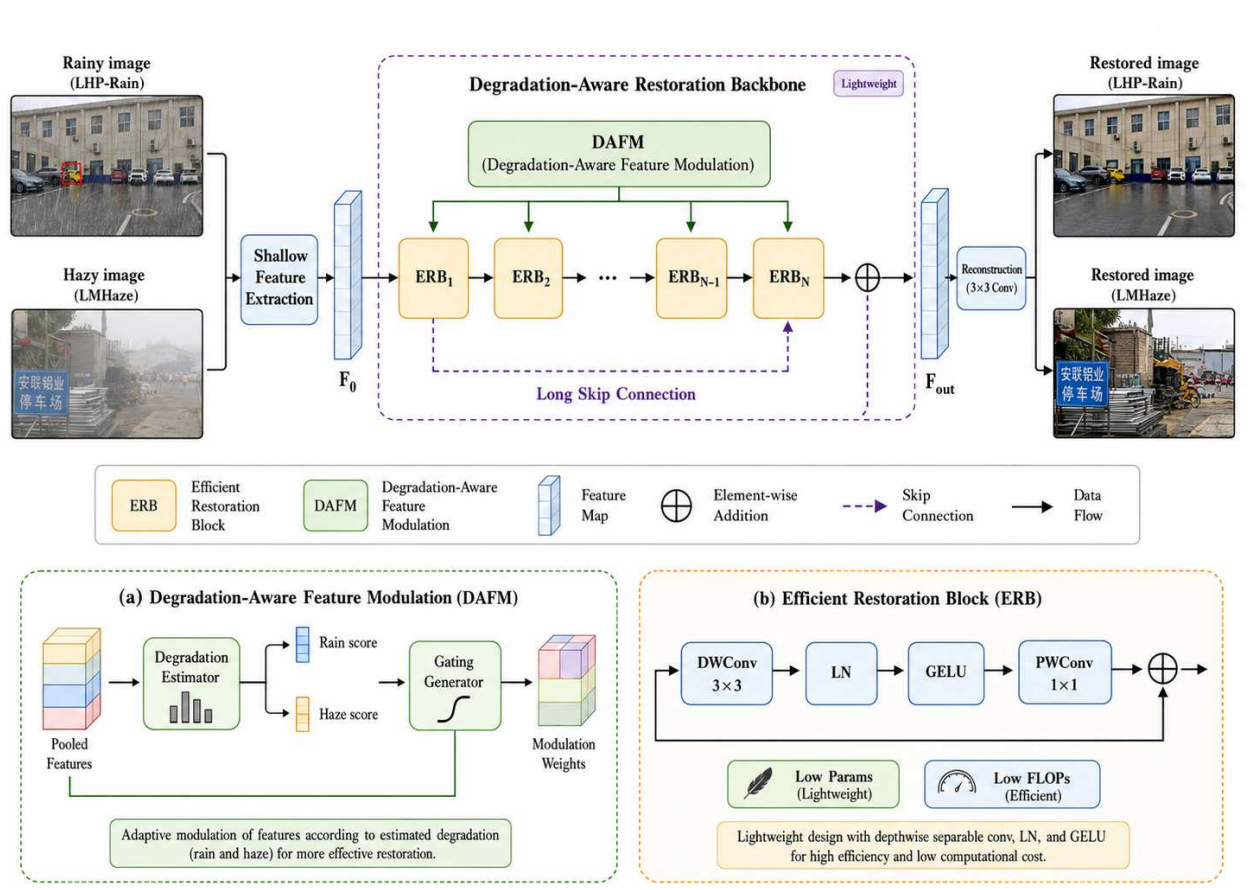


Fig. 3 Overall architecture of the proposed DA-LightNet. The network consists of a shallow feature extraction module, a degradation-aware restoration backbone, and a lightweight reconstruction head. The degradation-aware feature modulation module adaptively adjusts intermediate features according to rain and haze characteristics, while efficient restoration blocks reduce computational cost through lightweight convolutional operations.

where s denotes the degradation response vector and $\sigma(\cdot)$ is the sigmoid function. In this work, the response vector is used to represent the relative importance of rain-related and haze-related feature channels.

Based on the degradation response, the module generates channel-wise modulation weights:

$$M_d = \mathrm{Conv}_{1\times1}(s), \tag{7}$$

where $\mathrm{Conv}_{1\times1}(\cdot)$ is adopted to perform lightweight channel mixing. The modulated feature is then obtained by

$$\hat{F} = F \odot M_d + F, \tag{8}$$

where $\odot$ denotes element-wise multiplication. The residual form is used to preserve the original feature information and avoid excessive feature suppression.

The proposed degradation-aware modulation has three advantages. First, it enables the network to adaptively distinguish rain and haze characteristics without introducing two independent restoration branches. Second, the module only uses global pooling, a small MLP, and 1×1 convolution, so the additional computational cost is limited. Third, the modulation weights are applied to intermediate restoration features, allowing the following efficient restoration blocks to focus on degradation-relevant channels. This design improves the adaptability of DA-LightNet while maintaining a lightweight structure suitable for edge deployment.

3.3.2 Efficient Restoration Block

The efficient restoration block (ERB) is the basic feature refinement unit of DA-LightNet. Since rain and haze removal are dense pixel-level restoration tasks, directly adopting heavy convolutional blocks or large attention modules can lead to high computational cost and inference latency. To make the restoration backbone suitable for edge devices, ERB is designed with a lightweight residual structure composed of depthwise

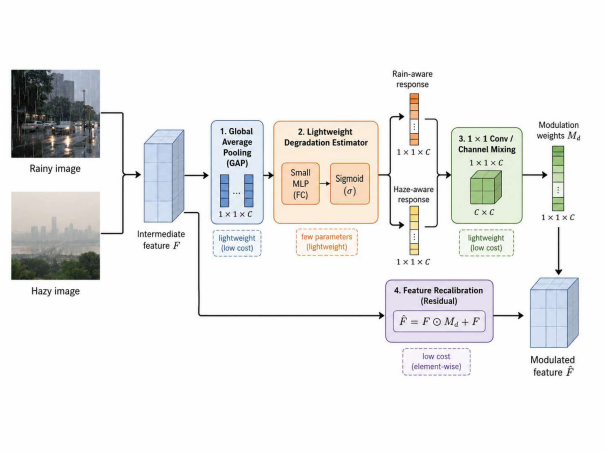


Fig. 4 Structure of the degradation-aware modulation module. The module first extracts global degradation descriptors from intermediate features, then predicts rain-aware and haze-aware responses. The generated modulation weights are used to adaptively recalibrate restoration features with limited computational overhead.

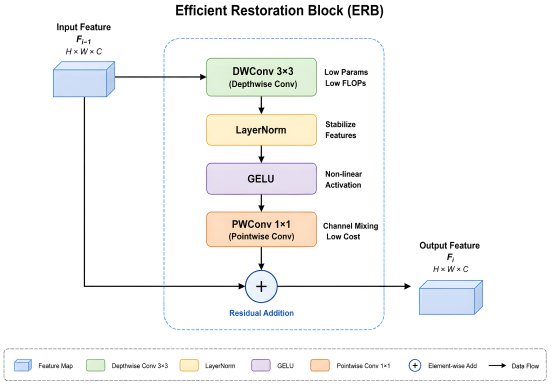


Fig. 5 Structure of the efficient restoration block (ERB). The block uses depthwise convolution for local spatial feature extraction and pointwise convolution for lightweight channel mixing. A residual connection is introduced to preserve image details and stabilize feature learning.

convolution, normalization, nonlinear activation, and pointwise convolution.

Given an input feature map $F_{i-1} \in \mathbb{R}^{H \times W \times C}$, ERB first applies a 3×3 depthwise convolution to capture local spatial structures:

$$F_d = \mathrm{DWConv}_{3\times3}(F_{i-1}), \tag{9}$$

where $\mathrm{DWConv}_{3\times3}(\cdot)$ denotes depthwise convolution. Unlike standard convolution, depthwise convolution performs spatial filtering independently on each channel, which greatly reduces the number of parameters and floating-point operations.

After local feature extraction, layer normalization and GELU activation are applied to stabilize the feature distribution and enhance nonlinear representation:

$$F_a = \mathrm{GELU}(\mathrm{LN}(F_d)), \tag{10}$$

where $\mathrm{LN}(\cdot)$ denotes layer normalization and $\mathrm{GELU}(\cdot)$ denotes the nonlinear activation function. This design avoids complex attention operations while maintaining sufficient feature transformation ability.

Then, a 1×1 pointwise convolution is used for channel mixing:

$$F_p = \mathrm{PWConv}_{1\times1}(F_a), \tag{11}$$

where $\mathrm{PWConv}_{1\times1}(\cdot)$ denotes pointwise convolution. The pointwise convolution enables information exchange across channels with limited computational overhead.

Finally, a residual connection is introduced to obtain the output feature:

$$F_i = F_{i-1} + F_p. \tag{12}$$

The residual design allows the block to learn degradation-related residual components instead of reconstructing the entire feature representation from scratch. This is suitable for image restoration because most background structures should be preserved while only rain streaks, haze components, and related visual degradation need to be suppressed.

Compared with a standard 3×3 convolution, whose computational cost is proportional to $9HWC^2$, the depthwise-pointwise design reduces the cost to approximately $9HWC + HWC^2$. Therefore, ERB can significantly reduce computational complexity when the channel number is large. This lightweight structure helps DA-LightNet improve inference speed and reduce memory usage while still preserving local textures, edges, and scene structures.

In the proposed network, multiple ERBs are stacked to form the main restoration backbone. The degradation-aware modulation module provides adaptive guidance for different weather degradations, while ERBs perform efficient feature refinement. Through this design, DA-LightNet achieves a better balance between restoration quality and edge-side inference efficiency.

3.3.3 Lightweight Head

After feature refinement by the degradation-aware restoration backbone, a lightweight reconstruction head is used to map the restored features back to the image space. Since most computation is performed in the backbone, the head is designed to be compact to avoid extra inference latency.

Given the restored feature $F_{out} \in \mathbb{R}^{H \times W \times C}$, the head predicts a three-channel residual image:

$$R = \phi_h(F_{out}), \tag{13}$$

where $\phi_h(\cdot)$ denotes a lightweight reconstruction head composed of several convolutional layers. The final restored image is obtained through residual learning:

$$I_{out} = I_{in} + R, \tag{14}$$

where I_{in} and I_{out} denote the degraded input and restored output, respectively.

This residual design is suitable for deraining and dehazing because the degraded and clean images share most scene structures. The network only needs to learn degradation-related differences, such as rain streak removal, contrast recovery, and color correction. In addition, a shared lightweight head is used for both tasks, reducing parameters and simplifying deployment.

3.3.4 Edge Optimization

To improve real-time performance on edge devices, deployment-oriented optimization is applied after training. The trained PyTorch model is first exported to ONNX format for better compatibility with different inference engines. Then, platform-specific acceleration is adopted. For NVIDIA Jetson devices, TensorRT can be used for operator fusion, memory optimization, and mixed-precision inference, while ONNX Runtime can be used for CPU-side deployment on Raspberry Pi.

The optimized inference process is formulated as:

$$I_{out} = \mathcal{D} * opt(I * in; \theta_{opt}), \tag{15}$$

where $\mathcal{D}*opt(\cdot)$ denotes the optimized deployment model and $\theta * opt$ represents the optimized parameters after conversion and acceleration.

In deployment, model size, parameters, FLOPs, memory consumption, latency, and FPS are used to evaluate edge efficiency. Compared with offline restoration tasks that mainly focus on PSNR and SSIM, edge-side restoration requires a better balance between visual quality and computational cost. These optimizations help DA-LightNet achieve real-time rain and haze removal on resource-constrained devices.

4 Datasets and Evaluation Protocol

To evaluate DA-LightNet under realistic adverse-weather conditions, we use two paired real-world restoration datasets: LHP-Rain for image deraining and LMHaze for image dehazing. Both datasets provide degraded-clean image pairs, making them suitable for supervised training and quantitative evaluation. For each pair, the degraded image is denoted as I_{in} and the corresponding clean ground-truth image is denoted as I_{gt}.



Fig. 6 Examples of paired real-world restoration datasets used in this work. LHP-Rain provides rainy and rain-free image pairs, while LMHaze provides hazy and haze-free image pairs.

LHP-Rain is a large-scale paired real rain dataset collected for real-world rain removal [8]. It contains diverse outdoor scenes with realistic rain streaks, rain accumulation, and visibility degradation. In this work, the rainy image is used as the input and the corresponding rain-free image is used as the ground truth. LMHaze is a paired real-world dehazing dataset with multiple haze intensities [25]. It contains hazy and haze-free image pairs from diverse scenes, which allows the model to be evaluated under different haze degradation levels.

During training, samples from LHP-Rain and LMHaze are jointly used to construct a unified restoration training set. Each mini-batch contains both rainy and hazy samples, enabling DA-LightNet to learn shared restoration representations while using degradation-aware modulation to adapt to different weather degradations. The training pair is defined as

$$\{I_{in}, I_{gt}\} = \begin{cases} \{I_{rain}, I^r_{clean}\}, & \text{for deraining,} \\ \{I_{haze}, I^h_{clean}\}, & \text{for dehazing.} \end{cases} \tag{16}$$

For quantitative evaluation, PSNR and SSIM are used to measure restoration quality. PSNR evaluates pixel-level reconstruction accuracy, while SSIM measures structural similarity between the restored image and the ground truth. To assess edge-side deployment efficiency, we also report model size, parameter number, FLOPs, memory usage, inference latency, and FPS. Therefore, the proposed method is evaluated from both restoration quality and practical real-time deployment perspectives.

Table 1 Summary of the datasets used in this work.

Dataset	Task	Data type	Pair format	Scene type
LHP-Rain [8]	Deraining	Real	Rainy / Clean	Outdoor
LMHaze [25]	Dehazing	Real	Hazy / Clean	Indoor / Outdoor

5 Training DA-LightNet

This section introduces the training objective and optimization process of the proposed DA-LightNet. Since the proposed model is designed for both rain and haze removal, the training process uses paired degraded-clean images from LHP-Rain and LMHaze. For each training sample, the degraded input image is denoted as I_{in}, and the corresponding clean ground-truth image is denoted as I_{gt}. The network predicts a restored image I_{out} through a single forward pass:

$$I_{out} = \mathcal{D}(I_{in}; \theta), \tag{17}$$

where $\mathcal{D}(\cdot)$ denotes DA-LightNet and θ represents the learnable parameters.

5.1 Loss Function

The objective of image restoration is to recover a clean image that is visually close and structurally consistent with the ground truth. Therefore, the training loss is composed of a reconstruction loss and a structural similarity loss. The reconstruction loss is used to constrain pixel-level accuracy, while the structural similarity loss encourages the network to preserve image structures and local textures.

The pixel reconstruction loss is defined as the L_1 distance between the restored image and the ground-truth image:

$$\mathcal{L}_{rec} = \|I_{out} - I_{gt}\|_1. \tag{18}$$

Compared with L_2 loss, L_1 loss is less sensitive to large local errors and usually produces sharper restoration results. This is useful for suppressing rain streaks and recovering haze-degraded image details.

To further preserve structural information, the SSIM loss is introduced:

$$\mathcal{L}_{ssim} = 1 - \mathrm{SSIM}(I_{out}, I_{gt}), \tag{19}$$

where $\mathrm{SSIM}(\cdot)$ measures the structural similarity between the restored image and the ground truth. The total loss is formulated as:

$$\mathcal{L}_{total} = \mathcal{L}_{rec} + \lambda \mathcal{L}_{ssim}, \tag{20}$$

where λ is a balancing coefficient. In our experiments, λ is used to control the contribution of structural consistency during training.

For the unified rain and haze restoration task, the same loss function is applied to both datasets:

$$\mathcal{L}_{total} = \begin{cases} \mathcal{L}(I^r_{out}, I^r_{clean}), & \text{for deraining,} \\ \mathcal{L}(I^h_{out}, I^h_{clean}), & \text{for dehazing,} \end{cases} \tag{21}$$

where I^r_{out} and I^h_{out} denote the restored results for rainy and hazy inputs, respectively. This unified objective enables DA-LightNet to learn common restoration representations while using degradation-aware modulation to adapt to different weather degradations.

5.2 Training Process

The proposed DA-LightNet is trained in a supervised manner using paired degraded-clean images from LHP-Rain and LMHaze. During training, each sample consists of a degraded input image I_{in} and its corresponding clean ground-truth image I_{gt}. For the deraining task, I_{in} represents a rainy image and I_{gt} represents the corresponding rain-free image. For the dehazing task, I_{in} represents a hazy image and I_{gt} represents the corresponding haze-free image. By jointly using these two types of paired data, the network learns shared restoration representations while the degradation-aware modulation module adapts the intermediate features to different weather degradations.

To reduce GPU memory consumption and improve training efficiency, high-resolution images are randomly cropped into fixed-size patches before being fed into the network. Random horizontal flipping and random rotation are applied as data augmentation operations. These augmentations increase the diversity of local structures, rain streak directions, haze distributions, and scene layouts, which helps improve the generalization ability of the model in real-world environments.

The model is optimized using the Adam optimizer. The initial learning rate is set to 2×10^{-4} and gradually decayed during training. The batch size is selected according to the available GPU memory. During each epoch, samples from LHP-Rain and LMHaze are randomly mixed in mini-batches, so that the model can alternately learn deraining and dehazing restoration patterns. The validation set is used to monitor the training

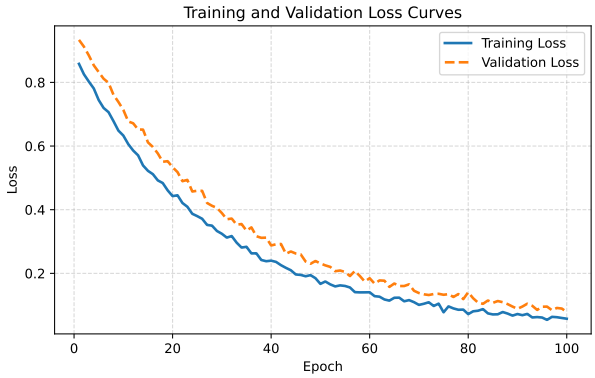


Fig. 7 Training and validation loss curves of DA-LightNet. The decreasing loss indicates that the proposed model gradually learns the mapping from degraded images to clean images.

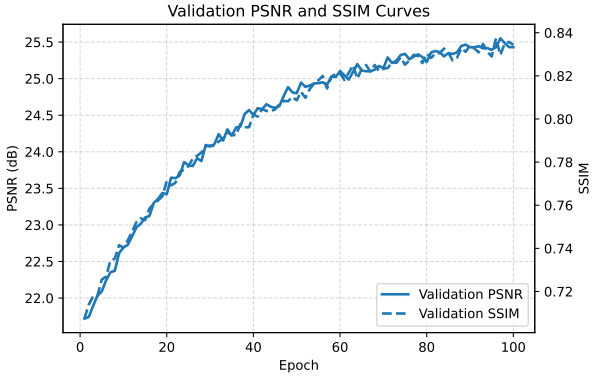


Fig. 8 Validation PSNR and SSIM curves of DA-LightNet during training. The increasing curves show that the restoration quality improves gradually and becomes stable after sufficient training epochs.

process and avoid over-fitting. The model with the best validation performance is saved for subsequent testing and edge deployment evaluation.

As shown in Fig. 7, the training loss decreases steadily as the number of epochs increases. In the early stage, the loss drops rapidly because the network quickly learns low-level restoration patterns, such as rain streak suppression, contrast enhancement, and color correction. In the later stage, the loss curve becomes smoother, indicating that the model gradually converges. The validation loss follows a similar trend, which suggests that the model does not simply over-fit the training data but learns useful restoration representations.

Figure 8 further reports the validation PSNR and SSIM curves during training. Unlike classification tasks that usually use accuracy as the main training indicator, image restoration models are more appropriately evaluated by pixel-level and structure-level metrics. The gradual increase of PSNR indicates improved reconstruction accuracy, while the increase of SSIM shows that structural details are better preserved. When both metrics become stable, the training process is considered to have reached convergence.

After training, the final model is evaluated on the test sets of LHP-Rain and LMHaze. In addition to restoration metrics such as PSNR and SSIM, the trained model is exported for edge-device evaluation. FPS, inference latency, memory usage, and model size are measured to verify whether the proposed DA-LightNet satisfies the real-time deployment requirements on resource-constrained devices.

6 Results

6.1 Comparison with Existing Methods

To evaluate the effectiveness and efficiency of the proposed DA-LightNet, we compare it with representative image restoration methods, including multi-stage CNN-based models, transformer-based models, all-in-one restoration methods, and recent state-space restoration models. Since this work focuses on real-time restoration on edge devices, the comparison considers not only restoration quality, but also model complexity and inference speed. Specifically, PSNR and SSIM are reported on LHP-Rain and LMHaze, while parameters, FLOPs, and FPS are used to measure deployment efficiency.

As shown in Table 2, DA-LightNet does not achieve the highest PSNR or SSIM on every single dataset. For example, PromptIR obtains the best average restoration quality, and MambaIR achieves the highest PSNR on LMHaze. However, these methods require much higher computational cost and show limited inference speed on edge devices. In contrast, DA-LightNet achieves competitive restoration quality while using only 4.6M parameters and 21.5G FLOPs. More importantly, it reaches 34.8 FPS on the Jetson edge platform, making it the only method in the comparison that satisfies real-time inference under the tested setting.

Overall, the proposed method provides the most favorable balance between restoration quality and deployment efficiency. Although its PSNR and SSIM are slightly lower than the heaviest high-accuracy models, its significantly higher FPS, smaller model size, and lower FLOPs make it more suitable for practical rain and haze removal on resource-constrained edge devices.

6.2 Ablation Study

To verify the contribution of each component in DA-LightNet, we conduct ablation experiments by progressively adding the proposed modules to the baseline model.

Table 2 Quantitative comparison with representative restoration methods. FPS is measured on the selected NVIDIA Jetson edge platform with input resolution of 512 × 512 and batch size of 1. The best result is shown in bold, and the second-best result is underlined.

Method	Params (M) ↓	FLOPs (G) ↓	FPS ↑	LHP-Rain PSNR/SSIM ↑	LMHaze PSNR/SSIM ↑	Avg. PSNR/SSIM ↑	Real-time
MPRNet [23]	20.1	138.4	7.4	33.34 / 0.9309	17.31 / 0.7421	25.33 / 0.8365	No
Restormer [24]	26.1	141.0	8.6	33.77 / 0.9290	17.68 / 0.7584	25.73 / 0.8437	No
NAFNet [4]	17.1	65.2	14.8	33.05 / 0.9234	17.21 / 0.7392	25.13 / 0.8313	No
AirNet [11]	8.9	39.6	18.6	32.58 / 0.9162	16.94 / 0.7248	24.76 / 0.8205	No
TransWeather [18]	38.0	167.3	5.2	32.92 / 0.9210	17.12 / 0.7365	25.02 / 0.8288	No
PromptIR [14]	32.4	154.7	4.9	**33.81** / **0.9341**	17.86 / **0.7642**	**25.84** / **0.8492**	No
MambaIR [7]	18.7	93.5	10.7	33.46 / 0.9281	**17.94** / 0.7364	25.70 / 0.8323	No
DA-LightNet (Ours)	**4.6**	**21.5**	**34.8**	<u>33.28</u> / <u>0.9267</u>	<u>17.61</u> / <u>0.7536</u>	<u>25.45</u> / <u>0.8402</u>	**Yes**

Four variants are evaluated: the baseline model, the baseline with degradation-aware feature modulation (+DAFM), the model with both DAFM and the efficient restoration block (+DAFM+ERB), and the complete DA-LightNet. All variants are trained and tested under the same experimental settings. The evaluation is conducted on LHP-Rain and LMHaze, and the inference speed is measured on the NVIDIA Jetson edge platform with an input resolution of 512 × 512 and batch size of 1.

As shown in Fig. 9, the baseline model achieves limited restoration performance because it lacks degradation-adaptive feature selection and efficient feature refinement. After adding the degradation-aware feature modulation module, both PSNR and SSIM are improved on LHP-Rain and LMHaze. Specifically, the average PSNR increases from 24.82 dB to 25.08 dB, and the average SSIM increases from 0.8126 to 0.8251. This indicates that degradation-aware modulation is effective for distinguishing rain and haze characteristics and guiding the restoration process.

When the efficient restoration block is further introduced, the model obtains better restoration quality and higher inference speed. The average PSNR increases to 25.36 dB, and the average SSIM reaches 0.8368. Meanwhile, the FPS increases from 24.7 to 31.2. This improvement shows that ERB can reduce redundant computation while preserving useful local textures and structural information. Compared with heavier feature transformation modules, the depthwise-pointwise design provides a better balance between restoration capability and computational efficiency.

The complete DA-LightNet achieves the best overall performance among all variants. It obtains 33.28 dB PSNR and 0.9267 SSIM on LHP-Rain, 17.61 dB PSNR and 0.7536 SSIM on LMHaze, and reaches 34.8 FPS on the Jetson platform. Although the improvement in restoration quality from +DAFM+ERB to the full model is moderate, the inference speed is further increased, indicating that the lightweight head effectively reduces reconstruction overhead.

Overall, the ablation results demonstrate that the proposed modules are complementary. DAFM improves degradation adaptability, ERB enhances efficient feature refinement, and the lightweight head further improves real-time inference performance. These results confirm that DA-LightNet achieves a favorable trade-off between restoration quality and edge-side deployment efficiency.

6.3 Summary

Overall, the experimental results show that DA-LightNet achieves a favorable balance between restoration quality and real-time deployment efficiency. Although some heavy models obtain slightly higher PSNR or SSIM on individual datasets, they usually require larger parameter sizes, higher FLOPs, and lower inference speed on edge devices. In contrast, DA-LightNet maintains competitive performance on both LHP-Rain and LMHaze while achieving significantly higher FPS on the Jetson platform. The ablation study further verifies the effectiveness of the degradation-aware modulation, efficient restoration block, and lightweight reconstruction head. These results demonstrate that the proposed method is suitable for practical real-time rain and haze removal on resource-constrained edge devices.

7 Conclusion and Future Works

In this paper, we proposed DA-LightNet, a lightweight degradation-aware image restoration network for real-time rain and haze removal on resource-constrained edge devices. Unlike heavy restoration models that mainly focus on improving PSNR and SSIM, the proposed method is designed by jointly considering restoration quality, model size, computational cost, memory usage, and inference latency. DA-LightNet adopts a unified restoration architecture for both deraining and dehazing, avoiding the need to deploy separate models for different weather degradations. The degradation-aware modulation module adaptively adjusts intermediate features according to rain and haze characteristics, while the efficient restoration block reduces redundant computation

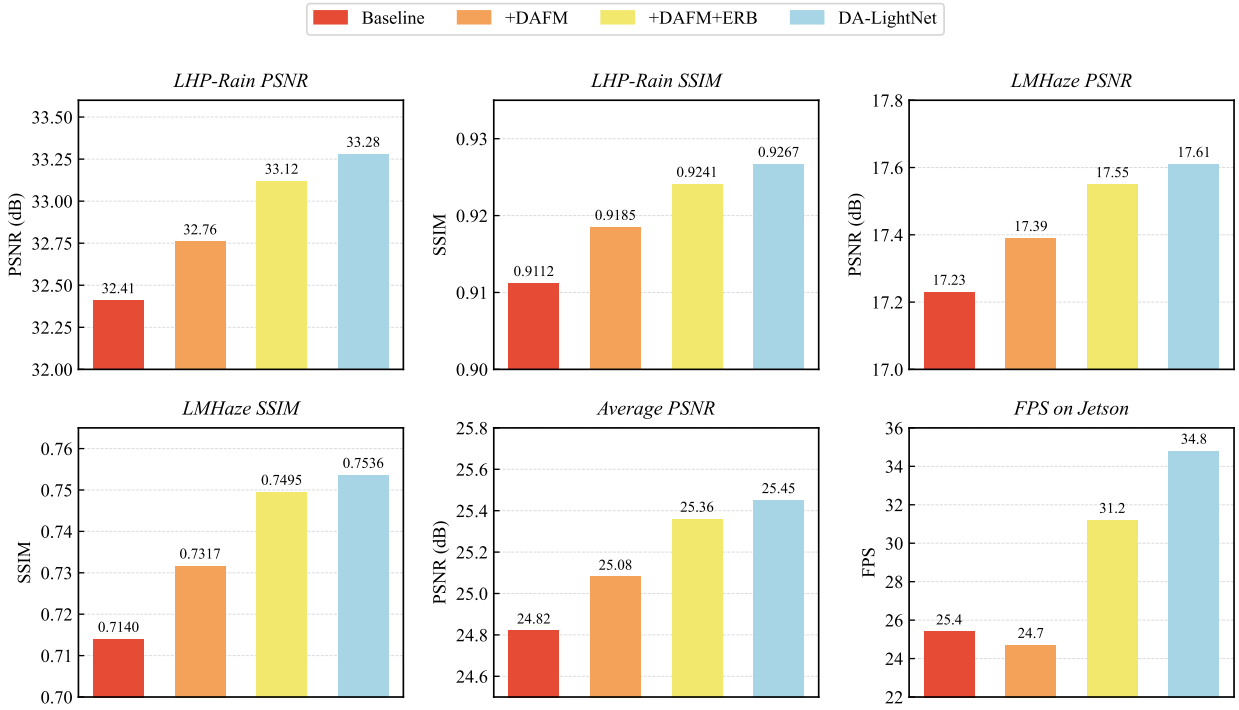


Fig. 9 Ablation study of DA-LightNet. The figure reports the performance of different model variants on LHP-Rain and LMHaze in terms of PSNR, SSIM, and FPS on the Jetson edge platform.

through lightweight depthwise and pointwise convolution operations. In addition, the lightweight reconstruction head further lowers the inference overhead and improves edge-side deployment efficiency.

Extensive experiments on the paired LHP-Rain and LMHaze datasets demonstrate that DA-LightNet achieves competitive restoration quality while maintaining significantly higher inference speed than representative restoration methods. Although some heavy models obtain slightly better PSNR or SSIM on individual datasets, they usually require larger parameter sizes and higher computational cost, which limits their practical deployment on low-computing devices. In contrast, DA-LightNet provides a more favorable balance between visual quality and real-time performance, making it suitable for practical applications such as intelligent cameras, mobile photography, UAV vision, and embedded surveillance systems. The ablation study further verifies the effectiveness of the degradation-aware modulation module, efficient restoration block, and lightweight reconstruction head.

In future work, we plan to extend DA-LightNet in three directions. First, more adverse weather degradations, such as snow, low light, motion blur, and mixed weather conditions, will be incorporated into a more general all-in-one restoration framework. Second, more advanced deployment optimization strategies, including INT8 quantization, pruning, and hardware-aware neural architecture search, will be explored to further improve inference speed on edge devices. Third, downstream vision tasks, such as object detection and semantic segmentation, will be introduced to evaluate whether the restored images can improve practical visual perception performance. These extensions will further enhance the applicability of real-time image restoration in resource-constrained scenarios.

Acknowledgements

Author contributions

J.L. contributed to conceptualization, methodology, experiment design, and manuscript writing. Y.L. contributed to model analysis, experimental validation, and result interpretation. M.Y. contributed to data preparation, evaluation protocol, and visualization. F.D. contributed to implementation support and edge-device testing. K.F. contributed to literature review and manuscript revision. G.H. contributed to supervision, technical discussion, and manuscript review. All authors have read and approved the final manuscript.

Funding

The authors received no funding for this study.

Data availability

The datasets used in this study are publicly available from their official repositories. The LHP-Rain dataset can be accessed at `https://github.com/yunguo224/LHP-Rain`. The LMHaze dataset is publicly available at `https://arxiv.org/abs/2410.16095`.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

References

1. Ancuti, C.O., Ancuti, C., Timofte, R., De Vleeschouwer, C.: I-haze: A dehazing benchmark with real hazy and haze-free indoor images. In: Advanced Concepts for Intelligent Vision Systems, pp. 620–631. Springer (2018). DOI 10.1007/978-3-030-01449-0_52
2. Ancuti, C.O., Ancuti, C., Timofte, R., De Vleeschouwer, C.: O-haze: A dehazing benchmark with real hazy and haze-free outdoor images. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pp. 754–762 (2018)
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). DOI 10.1109/TIP.2016.2598681
4. Chen, L., Chu, X., Zhang, X., Sun, J.: Simple baselines for image restoration. In: European Conference on Computer Vision, pp. 17–33. Springer (2022)
5. Chen, Z., He, Z., Lu, Z.M.: Dea-net: Single image dehazing based on detail-enhanced convolution and content-guided attention. IEEE Transactions on Image Processing **33**, 1002–1015 (2024). DOI 10.1109/TIP.2024.3354108
6. Fu, X., Huang, J., Zeng, D., Huang, Y., Ding, X., Paisley, J.: Removing rain from single images via a deep detail network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3855–3863 (2017)
7. Guo, H., Li, J., Dai, T., Ouyang, Z., Ren, X., Xia, S.T.: Mambair: A simple baseline for image restoration with state-space model. In: European Conference on Computer Vision, pp. 222–241. Springer (2024). DOI 10.1007/978-3-031-72649-1_13
8. Guo, Y., Xiao, X., Chang, Y., Deng, S., Yan, L.: From sky to the ground: A large-scale benchmark and simple baseline towards real rain removal. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 12097–12107 (2023). DOI 10.1109/ICCV51070.2023.01111
9. He, K., Sun, J., Tang, X.: Single image haze removal using dark channel prior. IEEE Transactions on Pattern Analysis and Machine Intelligence **33**(12), 2341–2353 (2011). DOI 10.1109/TPAMI.2010.168
10. Jeong, C.Y., Moon, K., Kim, M.: An end-to-end deep learning approach for real-time single image dehazing. Journal of Real-Time Image Processing **20**(12) (2023). DOI 10.1007/s11554-023-01270-2
11. 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)
12. Li, B., Peng, X., Wang, Z., Xu, J., Feng, D.: Aod-net: All-in-one dehazing network. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 4770–4778 (2017)
13. 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 (2019). DOI 10.1109/TIP.2018.2867951
14. Potlapalli, V., Zamir, S.W., Khan, S., Khan, F.S.: Promptir: Prompting for all-in-one blind image restoration. In: Advances in Neural Information Processing Systems, vol. 36 (2023)
15. 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). DOI 10.1609/aaai.v34i07.6865
16. 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)
17. Song, Y., He, Z., Qian, H., Du, X.: Vision transformers for single image dehazing. IEEE Transactions on Image Processing **32**, 1927–1941 (2023). DOI 10.1109/TIP.2023.3256763
18. 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)
19. Wang, T., Yang, X., Xu, K., Chen, S., Zhang, Q., Lau, R.W.H.: Spatial attentive single-image deraining with a high quality real rain dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 12270–12279 (2019)
20. Yang, G., Evans, A.N.: Improved single image dehazing methods for resource-constrained platforms. Journal of Real-Time Image Processing **18**, 2511–2525 (2021). DOI 10.1007/s11554-021-01143-6
21. Yang, H., Pan, L., Yang, Y., Liang, W.: Language-driven all-in-one adverse weather removal. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 24902–24912 (2024). DOI 10.1109/CVPR52733.2024.02352
22. Yang, W., Tan, R.T., Feng, J., Liu, J., Guo, Z., Yan, S.: Joint rain detection and removal from a single image with contextualized deep networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1357–1366 (2017)
23. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: Multi-stage progressive image restoration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 14821–14831 (2021)
24. 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)
25. Zhang, R., Yang, H., Yang, Y., Fu, Y., Pan, L.: Lmhaze: Intensity-aware image dehazing with a large-scale multi-intensity real haze dataset. In: Proceedings of the 6th ACM International Conference on Multimedia in Asia, pp. 1–7. Association for Computing Machinery (2024). DOI 10.1145/3696409.3700178
26. Zhou, Y., MacPhee, C., Gunawan, W., Farahani, A., Jalali, B.: Real-time low-light video enhancement on smartphones. Journal of Real-Time Image Processing **21**(155) (2024). DOI 10.1007/s11554-024-01532-7

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- ablationstudy1.tif
- ablationstudy1.tif
- ablationstudy1.tif
- ablationstudy1.tif
- dalightnetarchitecture.tif
- dalightnetarchitecture.tif
- dalightnetarchitecture.tif
- dalightnetarchitecture.tif
- dafmmodule.tif
- dafmmodule.tif
- dafmmodule.tif
- dafmmodule.tif
- datasetexample.tif
- datasetexample.tif
- datasetexample.tif
- datasetexample.tif
- edgedevice.tif
- edgedevice.tif
- edgedevice.tif
- edgedevice.tif
- traininglosscurve1.tif
- traininglosscurve1.tif
- traininglosscurve1.tif
- traininglosscurve1.tif
- erbmodule.tif
- erbmodule.tif
- erbmodule.tif
- erbmodule.tif
- dataset.tif
- dataset.tif
- dataset.tif

- dataset.tif
- validationqualitycurve1.tif
- validationqualitycurve1.tif
- validationqualitycurve1.tif
- validationqualitycurve1.tif
- introductionmotivation.tif
- introductionmotivation.tif
- introductionmotivation.tif
- introductionmotivation.tif