

MDeRainNet: An Efficient Macro-pixel Image Rain Removal Network

TAO YAN, Jiangnan University, China

WEILONG HUANG, Jiangnan University, China

WEIJIANG HE, Jiangnan University, China

CHENGLONG WANG, Jiangnan University, China

CIHANG WEI, Jiangnan University, China

YIWEI LU, Jiangnan University, China

XIANGJIE ZHU, Jiangnan University, China

YINGHUI WANG, Jiangnan University, China

RYN SON W.H. LAU, City University of Hong Kong, China

Since raining weather always degrades image quality and poses significant challenges to most computer vision-based intelligent systems, image de-raining has been a hot research topic in computer vision community. Fortunately, in a rainy light field (LF) image, background obscured by rain streaks in one sub-view may be visible in the other sub-views, and implicit depth information and recorded 4D structural information may benefit rain streak detection and removal. However, existing LF image rain removal methods either do not fully exploit the global correlations of 4D LF data, or only utilize partial sub-views (i.e., under-utilization of the rich angular information), resulting in sub-optimal rain removal performance and no-equally good quality for all de-rained sub-views. In this paper, we propose an efficient neural network, called *MDeRainNet*, for rain streak removal from LF images. The proposed network adopts a multi-scale encoder-decoder architecture, which directly works on Macro-pixel images (MPIs) for improving the rain removal performance. To fully model the global correlation between the spatial information and the angular information, we propose an *Extended Spatial-Angular Interaction (ESAI)* module to merge the two types of information, in which a simple and effective Transformer-based *Spatial-Angular Interaction Attention (SALA)* block is also proposed for modeling long-range geometric correlations and making full use of the angular information. Furthermore, to improve the generalization performance of our network on real-world rainy scenes, we propose a novel semi-supervised learning framework for our *MDeRainNet*, which utilizes multi-level KL loss to bridge the domain gap between features of synthetic and that of real-world rain streaks and introduces colored-residue image guided contrastive regularization to reconstruct rain-free images. Extensive experiments conducted on both synthetic and real-world LFs demonstrate that our method outperforms the state-of-the-art methods both quantitatively and qualitatively.

CCS Concepts: • **Computing methodologies** → **Image processing**.

Additional Key Words and Phrases: Rain removal, light field images, macro-pixel image (MPI), deep learning, semi-supervised learning

1 Introduction

Unlike regular cameras, plenoptic cameras could record target scenes with 4D light field data, i.e., spatial and angular information. Comparing with regular images, LFs benefit many attractive applications, such as depth estimation [53], foreground de-occlusion [52] and saliency detection [24]. Raining as a common weather condition, may greatly affect the performance of downstream

Authors' Contact Information: Tao Yan, yantao.ustc@gmail.com, Jiangnan University, Wuxi, Jiangsu, China; Weilong Huang, 6233115012@stu.jiangnan.edu.cn, Jiangnan University, Wuxi, Jiangsu, China; Weijiang He, 6213113057@stu.jiangnan.edu.cn, Jiangnan University, Wuxi, Jiangsu, China; Chenglong Wang, 6223115020@stu.jiangnan.edu.cn, Jiangnan University, Wuxi, Jiangsu, China; Cihang Wei, 6243115022@stu.jiangnan.edu.cn, Jiangnan University, Wuxi, Jiangsu, China; Yiwei Lu, 6243112037@stu.jiangnan.edu.cn, Jiangnan University, Wuxi, Jiangsu, China; Xiangjie Zhu, 6213113162@stu.jiangnan.edu.cn, Jiangnan University, Wuxi, Jiangsu, China; Yinghui Wang, wangyh@jiangnan.edu.cn, Jiangnan University, Wuxi, Jiangsu, China; Rynson W.H. Lau, rynson.lau@cityu.edu.hk, City University of Hong Kong, Hong Kong, China.

computer vision systems, such as autonomous driving and video surveillance. Therefore, removing rainy artifacts (i.e., rain streaks and raindrops) from rainy images is an important pre-procedure for most outdoor vision systems and has received extensive research attention in recent years.

To remove rainy artifacts from images, early research mainly focused on single image rain removal (SIRR) [21, 42], and a few works focus on video rain removal [62, 63]. A simple solution for LFI rain streak removal (LFRSR) is to apply SIRR methods to each sub-view of the input LF image independently. However, despite the great progress of SIRR methods [5, 31, 40, 45, 50, 59, 68–70] based on CNNs and/or Transformers, their solutions are always sub-optimal due to one inherent limitation that these rain streak detection and removal are highly dependent on the appearance (e.g., textures and shapes) of rain streak observed in a single image. Therefore, directly utilizing these SIRR methods to remove rain streaks from LF images would ignore the intrinsic relationship between different sub-views of an LF image, which may lead to angular inconsistency in the de-rained LF image. To address this problem, some recently proposed methods [12, 23, 60, 61] have made great efforts to fully explore and exploit information from LF images for rain removal. An example of the performances of the single image de-raining method [68] and the LF image de-raining method [61] working for LF image de-raining are shown in Figure 1, [where our method can well preserve the angular consistency](#).

Despite these LF image rain removal methods have achieved great progress, there are still two limitations. First, the structure (e.g. high dimension) of 4D LF data leads to difficulty in feature learning. Thus, it is necessary to re-organize the input LFI reasonably to facilitate networks to learn the feature representation for removing rain streaks from all sub-views of an LF image. Second, the receptive fields of such CNN-based de-raining networks are always limited and their representation learning efficiencies are low. 2D or high-dimensional (3D or 4D) convolutions utilized in these methods may be inefficient for feature learning and lead to high computation costs. Recently, Transformer has been proven to be effective for modeling long-range dependency, and has been successfully applied to image restoration [32, 52, 56, 59, 68].

LF images can record depth information and [multiple](#) sub-views of [the target scene](#), but due to limited spatial resolutions of LF images, rain streaks appearing on real-world LF images may be always tiny and hard to accurately detect. Fortunately, we observed that rain streaks appearing in the MPI of an LF image are much larger and sharper than those in single sub-views of the LF image, as shown in Figure 2. [The Macro-pixel image \(MPI\) is constructed by partitioning each sub-aperture view of the light field into \$N \times N\$ pixel blocks and then re-arranging the pixels corresponding to the same spatial point across all angular views into a single macro-pixel. Within each macro-pixel, the relative pixel positions encode angular information, whereas the macro-pixel's absolute position in the image retains the spatial coordinates.](#) The motivation of our method is that to detect rain streaks on the MPI of a rainy LFI would be more effective and efficient than that on the sub-views or the EPIs of a rainy LF Image.

In this paper, we propose an effective LFI rain streak removal network, called *MDeRainNet*, for rain streak removal from MPIs of LFIs. Specifically, our network takes as input the MPI of a rainy LF image, and adopts an encoder-decoder architecture to extract multi-scale features. Our *MDeRainNet* can fully explore and exploit the angular and spatial information of an LF image. Moreover, it can better preserve the content coherence of subviews of an LF image than existing methods. Since the low-level stages can always extract rich detailed features and the size of the feature maps is large, a *Spatial Feature Extractor* (SFE) and a convolution-based *Modified Disentangling Block* (MDB) are used to extract low-level features. We also propose an *Extended Spatial-Angular Interaction* (ESAI) module for extracting high-level features in the later stage of our network, which could model the correlation between spatial information and angular information over long range to capture global patterns of rain streaks.

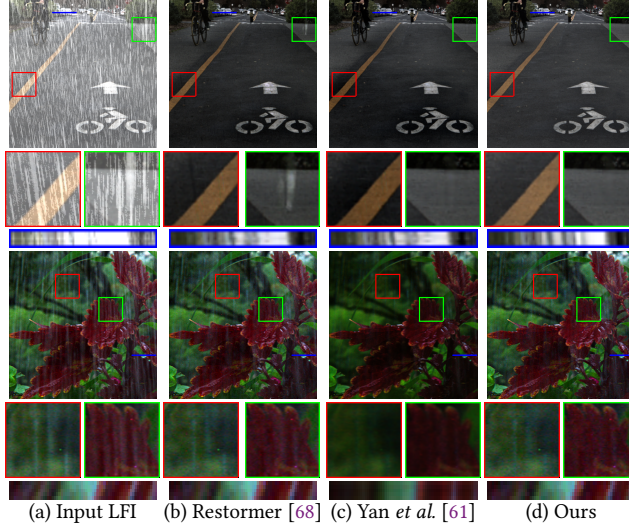


Fig. 1. Comparison of our method with two state-of-the-art de-raining methods [61, 68], tested on a synthetic rainy LFI (first row) and a real-world rainy LFI (second row) from RLMB [61]. The horizontal EPI (epipolar plane image) extracted along the blue line across the sub-views in the same row of the rainy LF image is also visualized to show the angular consistency (i.e., image content consistency across different de-rained sub-views).

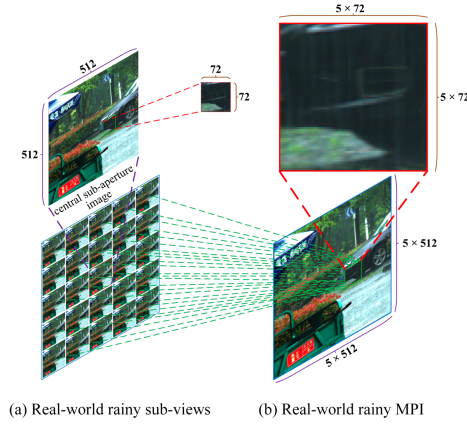


Fig. 2. An illustration of the relationship between the sub-views and the MPI (Macro-pixel image) representations of an LF image. The MPI consists of a large number of macro-pixels, and the number of pixels within a single macro-pixel is equal to the number of sub-views of the LF image. Each macro-pixel corresponds to a position sample of LF, and each pixel within a macro-pixel corresponds to a direction sample of LF at that position. It is worth noting that the rain streaks on MPI are always larger and sharper than the rain streaks in a single sub-view, as shown in the zoomed-in patch.

Though the semi-supervised de-raining method recently proposed by Yan *et al.* [61] achieves good performance on real-world rainy LFIs, its de-rained results may be over-smoothed and cannot maintain angular consistency well. Given that the high-level (semantic) features of rain streaks

extracted from synthetic and real-world rainy images should be similar, we enforce the consistency of the distributions of high-level features of rain streaks from synthetic rainy images and that from real-world images by minimizing Kullback–Leibler Divergence loss (KL-Loss) at multi-level scales [9, 57]. Furthermore, we design a contrastive regularization guided by colored rain-free residue images [28]. There are two de-raining methods[27, 66] that utilize rain-free residue images to decompose and maintain high-frequency background details. We propose to utilize rain-free residue images as positive samples in contrastive regularization for semi-supervised learning-based rain removal in our method. A real-world de-rained image should be closer to its corresponding rain-free residue image but away from the original rainy image in feature space. Therefore, by combining the proposed multi-level KL divergence and colored-residue image guided contrastive regularization, both the rain streaks domain and clean background domain can be fully utilized for supervision in our method. Thus, our *MDeRainNet* is able to well solve the domain-shift problem and improve generalization on real-world rainy LF images.

The main contributions of our work can be summarized as follows:

- We propose a novel and efficient rain removal network, called *MDeRainNet*, for rain streak removal from LFIs, which takes the MPI of an LFI as input.
- We propose an Extended Spatial-Angular Interaction (ESAI) module for fusing spatial and angular features, which can interact with spatial and angular features over long-range.
- We propose a novel semi-supervised learning framework for improving the generalization of our network on real-world rainy images, which utilizes multi-level KL loss to bridge the domain gap between features of synthetic and that of real-world rain streaks and introduces colored-residue image guided contrastive regularization to reconstruct the rain-free image.
- Extensive experiments conducted on synthetic and real-world rainy datasets demonstrate that our network outperforms other state-of-the-art methods both quantitatively and qualitatively, with low computation cost and less inference time.

2 Related Works

In this section, we briefly review the existing LF image rain removal methods, single image rain removal methods and video rain removal methods. In addition, some recently proposed multiple weather removal methods are also surveyed.

2.1 LFI Rain Removal

The research on LFI rain removal has not yet flourished, but several landmark works have emerged recently. Ding *et al.* [12] proposed a semi-supervised rain streak removal method based on GAN and Gaussian Mixture Model. The disadvantage is that it can only remove rain streaks by using sub-views in the same row/column of the rainy LFI at a time. Most recently, Yan *et al.* [61] proposed a semi-supervised rain streak removal network, which uses 4D convolutional layers to make full use of all sub-views of the input LF image for rain streak detection and removal. First, a dense sub-network based on multi-scale Gaussian processes is proposed to extract rain streaks from all input LF sub-views, and then the depth map and the fog map are estimated from the rough de-rained LF image. Finally, an adversarial and recurrent neural sub-network is constructed to progressively recover the rain-free LFI. Jing *et al.* [23] proposed a 4D re-sampling network for LFI raindrop removal, which exploits the complementary pixel information of the raindrop-free regions of the input LFI to produce the raindrop-free LFI. All these de-raining methods rely on 3D/4D convolution to exploit sub-views of an LFI, which may not fully utilize the abundant information embedding in rainy LFIs and are non-efficient due to high computation cost.

Wang *et al.* [52] proposed a novel LFI de-occlusion network, which combines Swin Transformer with CNN to extract both global and local features for restoring occlusion-free images. The network first uses CNNs to extract local features at shallow layers, and then utilizes Transformers to capture global features of large size occlusions at deep layers. Wang *et al.* [54] proposed a general disentangling mechanism to extract the spatial feature, the angular feature and the EPI feature from the input MPI, for LFI super-resolution and depth estimation.

2.2 Single-Image Rain Removal

Unlike LF-based methods that utilize complementary information among different sub-view images, single-image de-raining methods remove rain from the background by analyzing the visual information of a single image. Traditional de-raining methods [11, 25, 30, 39, 72] usually separate rain layer from the background layers by exploring the physical properties or prior knowledge of rain streaks (raindrops). In recent years, deep learning-based methods have dominated the single-image de-raining research topic. Fu *et al.* [14, 15] proposed the earliest deep learning-based methods, which adopt CNN to remove rain streaks from the high-frequency detail layer of a rainy image. Later, a lot of deep learning-based networks have been proposed to improve the single image de-raining performance, including multi-scale deep networks [21, 64], multi-stage neural networks [42, 69], conditional generative adversarial networks [71] and attention-based encoder-decoder networks [18, 49, 50]. For the effectiveness and efficiency of image deraining, Jiang *et al.* [22] firstly devised a novel coupled representation module named CRM to investigate the blending correlations between rain streaks and rain-free details. Most recently, several semi-supervised and unsupervised de-raining methods [7, 20, 65, 67] have also been proposed for de-raining real-world rainy images. Most recently, Fu *et al.* [16] designed a patchwise hypergraph convolutional module and a biological brain-inspired continual learning algorithm to achieve better generalizability and adaptability in real-world scenes.

These single image de-raining methods could not be directly utilized for LFI de-raining due to two main reasons. First, they can not make full use of abundant information embedding in LFIs for high-performance de-raining. Second, they could not ensure image content consistency between all sub-views of a de-rained LFI.

2.3 Multiple Weather Removal

Recently, some unified image restoration networks, such as IPT [3], SANet [47], RAHC [48], Uformer [56], TransWeather [45] and Restormer [68] have been proposed. IPT [3] is the first network, which uses standard Transformer blocks to form a multi-task learning framework for image super-resolution, de-raining and de-noising. SANet [47] addresses the joint task of rain streak removal and low-light illumination enhancement in a single forward pass guided by a degradation-localization mask and contrastive learning, which ensures faithful texture and lighting recovery. RAHC [48] employs multi-head aggregation to learn diverse degradation subspaces and leverages output-space discrimination for unified handling of multiple adverse weather conditions. Though these methods achieve impressive performance for multi-task single image restoration, they can not be directly utilized for LFI de-raining.

3 Our Proposed Network

In this section, we describe our *MDeRainNet* and its semi-supervised learning framework in detail.

3.1 Background and Motivation

A 4D LFI can be expressed as $L \in \mathbb{R}^{U \times V \times H \times W \times C}$, where (H, W) , (U, V) and C denotes the spatial resolution, the angular resolution and the RGB channels, respectively. As shown in Figure 2, an

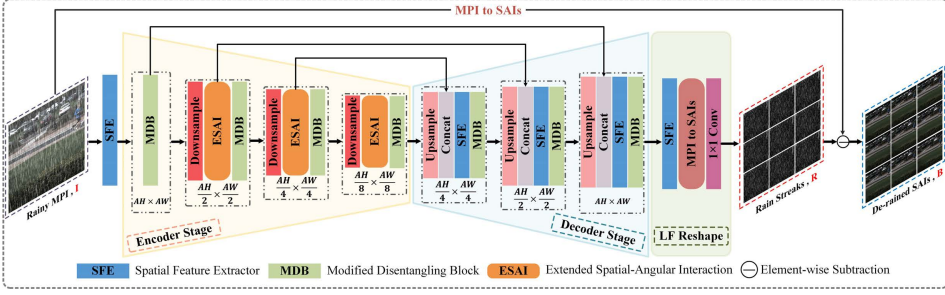


Fig. 3. Overall architecture of our *MDeRainNet*. The input of our network is rainy MPI, and the network learns LF rain streak residuals as a whole. We feed the MPI into an encoder-decoder network to extract hierarchical features. Firstly, the combination of SFE and MDB is used to extract low-level features from the input MPI in the encoder. Subsequently, we combine efficient ESAI and MDB to complement the receptive fields of these two modules on multiple scales of the encoder, so as to obtain locality while capturing long-range dependencies. Finally, at the end of the encoder, the rain residuals estimated on MPI are reshaped into a sub-aperture images array to get the final de-rained SAs.

array of sub-views with different angular coordinates (u, v) is shown, each of which is called a sub-aperture image (SAI) or sub-view, termed as $L(u, v) \in \mathbb{R}^{H \times W \times C}$. Recently, most LFI processing methods take as input a subset of sub-views of an input LFI. However, such form of input is not convenient for extracting spatial and angular information, since angular information is implicitly embedded in different sub-views. On the other hand, some LFI processing methods take EPIs (Epipolar Plane Images) as input. Since an EPI is only a 2D horizontal or vertical slice of a 4D LF image, it is difficult to exploit both spatial and angular information from EPIs. Therefore, in order to make full use of both the spatial and angular information for rain streak removal, following previous super-resolution methods[54, 55], our network directly takes the MPI of the input LFI as input. As far as we know, our *MDeRainNet* is the first method, which directly takes as input MPIS for LF image rain streak removal.

In Figure 4, we provide an illustration of the angular and spatial feature extractors. By taking the MPI of an LFI as input, the spatial information and angular information embedding in the input LFI can be well extracted by leveraging domain-specific convolutions, i.e., spatial feature extractor (SFE) and angular feature extractor (AFE) [54, 55].

Motivated by our observations that rain streaks appearing in MPI not only have larger sizes than that of rain streaks appearing in each single sub-view of the input LFI, but also have sharper edges (as shown in the zoomed-in area of Figure 2), it may be much effective to remove rain streaks from MPIS. In another word, although both the MPI and sub-views of an LFI are 2D images, the rain streaks in MPIS contain richer structure information than that in single sub-views, which may help our method to extract rain streak features from MPIS. Furthermore, since the MPI contains both spatial and angular information of an LFI in a 2D image format, extracting rain streaks from it will be more convenient than that from sub-views of the input rainy LFI.

3.2 Overview of Our *MDeRainNet*

Our *MDeRainNet* adopts an encoder-decoder structure, as shown in Figure 3. The complete architecture of our *MDeRainNet* with the semi-supervised learning framework is shown in Figure 9. To effectively detect rain streaks from MPIS of LFIs, we propose an *Extended Spatial-Angular Interaction* (ESAI) module, which is able to model the correlation between spatial information and angular information over long-range to capture global patterns of rain streaks. We also introduced a Modified




Fig. 4. The working principle of spatial and angular feature extractors. We show an example of an LF with an angular resolution of 3×3 (i.e., $A = 3$) and a spatial resolution of 4×4 (i.e., $H = 4, W = 4$). For better visualization, different macro-pixels are filled with different background colors, and pixels from different sub-views are represented by different labels (e.g., moon or pentagram). SFE extracts spatial features within the same view, while AFE extracts angular features within the same macro-pixel, thereby decoupling spatial and angular information.

Disentangling Block (MDB) for extracting and integrating spatial and angular features. The ESAI and MDB are used to form the encoder of our *MDeRainNet*, which combines the advantages of Transformer and CNN to ensure that both **global** features and local features can be retained. Then, at each scale of the decoder of our *MDeRainNet*, the SFE and the MDB are adopted to progressively estimate rain streaks from the input MPI. Skip-connections are also adopted to concatenate the encoded feature extracted at the same scale with the decoded feature to avoid losing meaningful image details and make our network converge fast. Pixel-unshuffling and pixel-shuffling operations [43, 54] are applied for feature downsampling and upsampling, respectively. Finally, at the LF Reshape stage, an SFE is used to further refine the input feature, and then the enhanced MPI feature is rearranged into the format of an array of sub-views. A 1×1 convolution is also applied to squeeze the number of feature channels to 3 in order to produce the rain streak (RGB) sub-views, termed as R . The de-rained sub-views, **denoted as B , are** obtained by converting the input rainy MPI I into an array of sub-views, **from which the rain component R is then subtracted.**

To improve the generalization performance of our *MDeRainNet* for real-world scenes, **we propose a semi-supervised learning framework (detailed in Section 3.5). This framework utilizes two key components: Multi-level KL divergence to enforce feature distribution consistency between rain streaks from synthetic images and those from real-world rainy images, and colored rain-free residue image-guided contrastive regularization to aid in the reconstruction the rain-free image.**

3.3 Modified Disentangling Block (MDB)

We introduced a modified disentangling block (MDB), as shown in Figure 5, for disentangling features of the input MPI. The typical disentangling block [54] proposed for disentangling features of an LF consists of an SFE, an AFE and an epipolar feature extractor (EFE) for extracting the spatial feature, the angular feature and the EPI feature from the MPI of an LFI, respectively. **Given that AFE can extract angular information from all sub-views while EFE can only extract angular information from EPIS, we believe that the angular information extracted by AFE would be more comprehensive. Thus, our MDB omits the EFE while extending both the SFE and the AFE.**

Specifically, as shown in Figure 5, given the MPI feature $F_{in} \in \mathbb{R}^{AH \times AW \times C'}$, a spatial feature extraction branch and an angular feature extraction branch of the MDB are used to disentangle the spatial and angular features, respectively. In the angular feature extraction branch, an AFE is first applied to generate an intermediate angular feature $F_{inter}^a \in \mathbb{R}^{H \times W \times \frac{C'}{2}}$. Then, the output




Fig. 5. The structure of our Modified Disentangling Block (MDB).

feature $F_{inter,u}^a \in \mathbb{R}^{AH \times AW \times \frac{C'}{2}}$ can be obtained by upsampling this intermediate feature F_{inter}^a by a factor of A , with a 1×1 convolution layer and an upsample layer. Note that compared to the typical disentangling block in [54], the channel number of the angular feature $F_{inter,u}^a$ is doubled in our MDB. Next, the feature $F_{inter,u}^a$ is concatenated with the intermediate spatial feature $F_{inter}^s \in \mathbb{R}^{AH \times AW \times C'}$ extracted by two SFEs of the spatial branch. The concatenated features are then passed through a 1×1 convolution to reduce the number of channels. A spatial residual block, comprising two SFEs and a skip connection, is introduced to fuse $F_{inter,u}^a$ and F_{inter}^s adaptively. Finally, the fusion feature, F_{fused} , is added to the input feature F_{in} , to obtain the enhanced feature F_{out} .

In our network, MDB is first employed to extract shallow features at the beginning of the encoder. It is also used after each ESAI module to enhance features produced by ESAIs.

3.4 Extended Spatial-Angular Interaction (ESAI) Module

We propose an ESAI module for interacting with and fusing spatial and angular feature of LFIs, as shown in Figure 6. Our ESAI contains two key blocks: a *Spatial-Angular Interaction Attention* (SAIA) block and a *Deep Fusion* (DF) block.




Fig. 6. The structure of our proposed ESAI module.

The input feature F_m is first separated into its spatial component F_m^s and angular component F_m^a using an SFE and an AFE, respectively. The angular feature F_m^a is upsampled with a factor of A to obtain the upscaled angular feature $F_{m,u}^a$. Then, F_m^s and $F_{m,u}^a$ are fed into the SAIA block, passing through each attention unit in SAIA to make cross-feature attention process and feature reinforcement attention process. The spatial and angular features generated by the i -th attention unit of SAIA (as shown in Figure 7), denoted as U_i^s and U_i^a , are then concatenated with F_m^s and $F_{m,u}^a$, respectively. Finally, the concatenated spatial feature F_c^s and angular feature F_c^a are fused by a DF block to obtain the fused feature F_d , which is later added to the input feature F_m by a skip connection to obtain the final fused feature F_d' .

SAIA Block. The SAIA block, as shown in Figure 7, contains L attention units, each consisting of two main components: cross-feature attention and feature reinforcement attention. Specifically, first, the cross-feature attention based on the cross-attention mechanism [2, 4, 8, 19, 29, 36, 38] is

used to make interaction between the spatial feature and the angular feature. Second, the feature reinforcement attention based on the self-attention mechanism [3, 13, 46] is executed to further enhance the spatial and angular features output by the cross-feature attention.

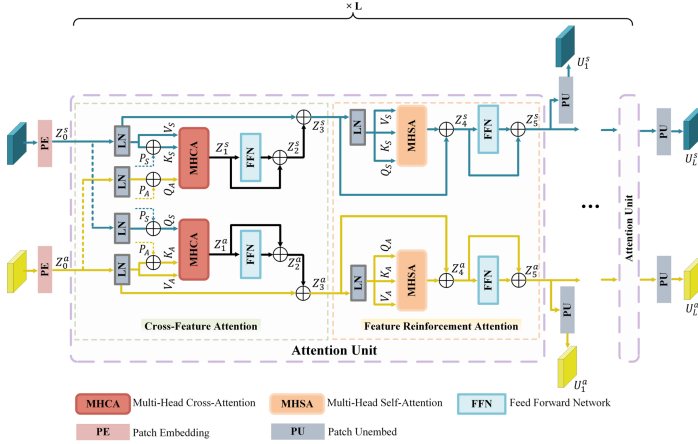


Fig. 7. The structure of SAIA (Spatial-Angular Interaction Attention) block.

Following [51], the Patch Embedding (PE) of our SAIA block first splits the input spatial and angular features, $F_m^s \in \mathbb{R}^{H' \times W' \times C'}$ and $F_{m,u}^a \in \mathbb{R}^{H' \times W' \times C'}$, into patches with the spatial size of $P \times P$. Then, it reshapes each patch into a flattened 1-D feature, and a total of $N = \frac{H'W'}{P^2}$ flattened features can be obtained. Finally, the array of flattened features, termed as F_p , is mapped to the latent D -dimensional embedding space via a trainable linear projection, which plays a role in feature transformation and avoids the flattened features being too large, as follows:

$$Z_0^s = [F_{p,1}^s E, F_{p,2}^s E, \dots, F_{p,N}^s E], \quad (1)$$

$$Z_0^a = [F_{p,1}^a E, F_{p,2}^a E, \dots, F_{p,N}^a E], \quad (2)$$

where $E \in \mathbb{R}^{(P^2 C') \times D}$ denotes the linear projection used for patch embedding. After completing the patch embedding, the features Z_0^s and Z_0^a are then fed into the cross-feature attention.

The cross-feature attention, as shown in Figure 7, contains two branches, i.e., one spatial branch and one angular branch, each of which consists of a multi-head cross-attention (MHCA) and a feed-forward network (FFN). In the spatial branch, the features Z_0^a and Z_0^s pass through a layer normalization (LN) [1] and are then directly added with the learnable angular position encoding P_A and spatial position encoding P_S , respectively, to generate query Q_A and key K_S . Here, the angular and spatial position encodings can make the global correlation of features calculated cross-feature attention stronger. V_S is directly assigned as Z_0^s after processed by the normalization layer without adding the spatial position encoding P_S . Afterwards, the produced Q_A , K_S , and V_S are fed into MHCA to learn the relationship between different spatial and angular tokens, as follows:

$$Q_A = \text{LN}(Z_0^a) + P_A, \quad (3)$$

$$K_S = \text{LN}(Z_0^s) + P_S, \quad (4)$$

$$V_S = \text{LN}(Z_0^s), \quad (5)$$

$$\text{MHCA}(Q_A, K_S, V_S) = \text{Concat}(h_1 \dots h_n) W_O, \quad (6)$$

where n is the number of heads, and the embedding dimension of K_S , V_S and Q_A is split into n groups. h_i refers to the attention matrix calculated by the i -th head, and similar to [32, 56], it can be defined as:

$$h_i = \text{Softmax}\left(\frac{Q_{A,i}W_{Q,i}(K_{S,i}W_{K,i})^T}{\sqrt{d}}\right)V_{S,i}W_{V,i}, \quad (7)$$

where $W_{Q,i}$, $W_{K,i}$ and $W_{V,i} \in \mathbb{R}^{d \times d}$ denote the linear projection matrices of the queries, keys and values for the i -th head of MHCA, respectively. $d = D/n$ is the feature dimension for each head. $W_O \in \mathbb{R}^{D \times D}$ denotes the output projection matrix.

Subsequently, the features passing through MHCA are further fed to an FFN for further feature transformation, which consists of an LN layer and a multi-layer perception (MLP) layer. In summary, the calculation formula of the spatial branch in the cross-feature attention can be written as:

$$Z_1^s = \text{MHCA}(Q_A, K_S, V_S), \quad (8)$$

$$Z_2^s = \text{MLP}(\text{LN}(Z_1^s)) + Z_1^s, \quad (9)$$

$$Z_3^s = Z_2^s + \text{LN}(Z_0^s). \quad (10)$$

The calculation of the angular branch is similar to that of the spatial branch, except that in the calculation of multi-head cross-attention, the key (K_A) and value (V_A) are calculated from the angular features, while the query (Q_S) is calculated from the spatial features. Through the above calculation process in the spatial and angular branches, the long-range dependencies between spatial and angular features can be modeled to obtain the cross-feature attention.

The calculation of feature reinforcement attention in the spatial branch is the same as that in the angular branch. We take the spatial branch as an example to describe it. The structure of our feature reinforcement attention is the same as a single Transformer layer, which consists of a multi-head self-attention (MHSA) and an FFN. First, the output of the cross-feature attention is utilized to generate query Q_S , key K_S and value V_S via a LN, i.e., $Q_S = K_S = V_S = \text{LN}(Z_3^s)$. The process of the Feature Reinforcement Attention with an MHSA and an FFN is as follows:

$$Z_4^s = \text{MHSA}(Q_S, K_S, V_S) + Z_3^s, \quad (11)$$

$$Z_5^s = \text{MLP}(\text{LN}(Z_4^s)) + Z_4^s. \quad (12)$$

Similarly, the enhanced feature Z_5^a can be obtained by the angular branch. Finally, Z_5^s and Z_5^a are passed to their respective Patch Unembedding (PU) operations to produce the output features U_i^s and U_i^a .

Deep Fusion. The output of the SAIA block is fed into the final deep fusion block to fully integrate the interacted features, as shown in Figure 8. Specifically, first, the spatial and angular features output by all attention units of the SAIA are concatenated with the input spatial and angular features (F_m^s and $F_{m,u}^a$), respectively, to form the inputs of DF as follows:

$$F_c^s = [F_m^s, U_1^s, \dots, U_L^s] \in \mathbb{R}^{H' \times W' \times (L+1)C'}, \quad (13)$$

$$F_c^a = [F_{m,u}^a, U_1^a, \dots, U_L^a] \in \mathbb{R}^{H' \times W' \times (L+1)C'}. \quad (14)$$

Later, the concatenated angular feature F_c^a is fed into a 1×1 convolution followed by a ReLU layer to reduce the channels and generate the intermediate feature $F_t^a \in \mathbb{R}^{H' \times W' \times C'}$. Finally, F_t^a is concatenated with the spatial features F_c^s , and then the combined feature is further processed by an SFE and a ReLU layer to obtain the final fused feature F_d .




Fig. 8. The structure of our Deep Fusion (DF) block. As the bottleneck of ESAI, DF is used to fully merge the spatial and angular features.

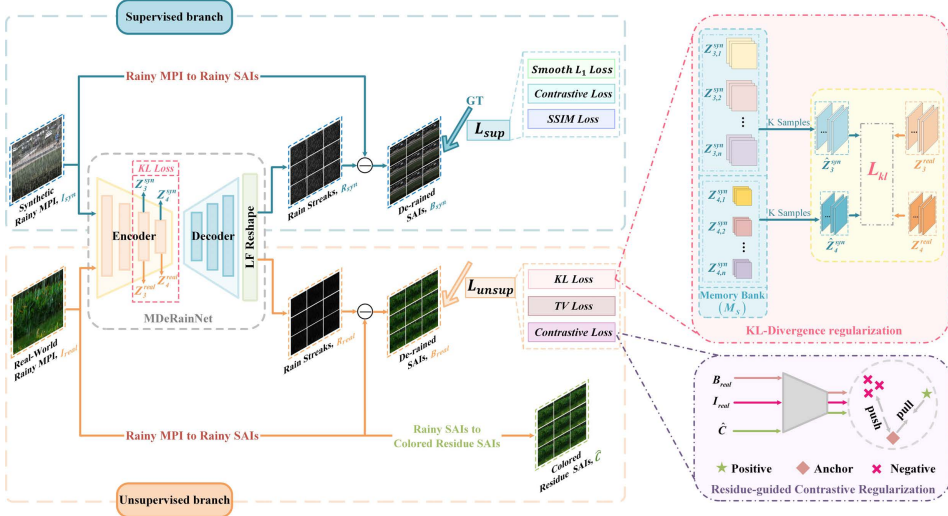


Fig. 9. An overview of our semi-supervised learning framework for our *MDerainNet* for improving its generalization on real-world rainy LF images.

3.5 Semi-supervised Learning

We propose a novel semi-supervised learning framework for our network to improve its generalization on real-world rainy images without ground-truths. As shown in Figure 9, the semi-supervised framework consists of a supervised branch and an unsupervised branch, both of which share the same rain removal network shown in Figure 3. Specifically, our network is trained with two stages. In the first stage, supervised learning is conducted on paired synthetic rainy LF images to obtain the best de-raining model on the synthetic datasets. In the second stage, the de-raining model obtained in the previous stage is further trained by utilizing both synthetic and real-world rainy LF images simultaneously. For the second phase of training, we utilize a multi-level KL divergence regularization to enhance the feature distribution consistency between synthetic and real-world rain streaks, and introduce a contrastive regularization guided by colored-residue images to achieve high-quality background restoration.

Multi-level KL. In our task, the synthetic and real-world rain streak domains can be regarded as two different probability distributions. By modeling the closing process on the distribution probability of deep features extracted from synthetic LF images and that from real-world images, the interdomain distance can be reduced. To achieve this goal, we introduce a multi-level KL-divergence regularization to align the deep features of rain streaks extracted from synthetic rainy

LF images with that of real-world rainy LF images by the shared encoder at multiple scales to achieve distribution consistency.

Specifically, the shared encoder first extracts multi-scale rain streak features from synthetic and real-world rainy LFIs, and the extracted synthetic rain streak features Z_3^{syn} and Z_4^{syn} (at 3rd and 4th scales) are stored in M_S . Here, M_S is a memory bank that stores the latent space vectors at each level of the synthetic rain streaks to generate pseudo-GT for the real-world rain streaks, which is then used for KL divergence regularization. The creation of the pseudo-GT is based on the assumption that while projecting high-level features of real-world rain streaks to the latent space, they can be approximately represented as a weighted combination of high-level feature of synthetic rain streaks [65]. In M_S , there are two scales of rain streak features, termed $Z_{3,i}^{syn} \in \mathbb{R}^{\frac{AH}{4} \times \frac{AW}{4} \times C'}$ and $Z_{4,i}^{syn} \in \mathbb{R}^{\frac{AH}{8} \times \frac{AW}{8} \times C'}$, where i represents i -th batch. Before each epoch of training in the second stage, the encoder of our network will extract latent features of rain streaks from the synthetic rainy LFIs to totally replace the features of synthetic rain streaks stored in the memory bank.

Following the semi-supervised learning-based deraining method [65], we use the top K (set to 16 by default) nearest synthetic rain streak features to obtain a weighted combination, termed as \hat{Z}^{syn} , for the real-world rain streak feature Z^{real} . Finally, \hat{Z}^{syn} and Z^{real} are used to calculate the KL loss. The above process can be described as:

$$\hat{Z}_j^{syn} = \sum_{i=1}^K \alpha_i Z_{j,i}^{syn}, \quad Z_{j,i}^{syn} \in \text{nearest}(Z_j^{real}, M_S, k), \quad (15)$$

$$\mathcal{L}_{kl} = \sum_{j=3}^4 \text{KL}(\text{Softmax}(\hat{Z}_j^{syn}), \text{Softmax}(Z_j^{real})), \quad (16)$$

where i is used to index the i -th nearest feature for Z^{real} . j refers to the j -th scale of the encoder. α_i is the weighting coefficients, and $\text{nearest}(P, Q, N)$ is a function for finding the top N nearest neighbors of P in Q .

Residue Image. After narrowing the gap of the rain streak domain in the latent space, a contrastive regularization guided by rain-free residue image is introduced into the unsupervised branch to further bridge the gap of the clean background domain between synthetic and the real-world rainy images. The residue image was first proposed by Li *et al.* [28]. It is a single-channel image, which is the residual result of the maximum channel value and the minimum channel value of the input rainy image. Given a rainy LF image I in the form of a SAI array, the residue image of I can be defined as:

$$I'_k(x) = \max_{i \in \{R, G, B\}} I_k^i(x) - \min_{j \in \{R, G, B\}} I_k^j(x), \quad (17)$$

where I_k^i, I_k^j denote the color channels of the k -th sub-view of I , $k \in \{1, 2, \dots, A^2\}$. x denotes pixel coordinates of the sub-view. A refers to the angular dimension, and I' represents the residue image. The appearance of rain streaks in residue images is colorless, i.e., I' is free of rain streaks and contains only a transformed version of the background details. In practice, we use colored-residue images [28] to achieve better de-raining performance in real-world rainy scenarios. Here, we take the k -th rainy sub-view as an example to introduce the acquisition process of its colored-residue image.

For the k -th original rainy sub-view I_k , we first convert it to the YCbCr domain. Considering the fact that the impact of rain streaks on the image is mainly concentrated in the luminance channel (Y channel) while the chrominance channel (CbCr channels) is not affected much, we replace the Y channel with the single-channel residue image I'_k defined in Eq. 17. After obtaining the intensity

values for the residual image in the I^r CbCr domain, we convert them back to the RGB domain to obtain the colored-residue image C_k^r . Though most of the background details are well preserved in C_k^r , there is still some color deviation compared to the original rainy sub-view. The reason is that when some background areas are achromatic (i.e., white, gray or black), they are cancelled out along with the rain streaks, as defined in Equation 17, resulting in some background areas being darkened. To alleviate this problem, we utilize a rain-free residue image-guided filter [27] working on the original rainy sub-view I_k , guided by I_k^r , to obtain another colored rain-free image F_k^r , which can be defined as:

$$F_k^r = \text{GF}(I_k, I_k^r), \quad (18)$$

where $\text{GF}(\cdot)$ is the guided filtering function [27]. The loss of details in F_k^r is severe, but its color fidelity is good. Compared to C_k^r , F_k^r can maintain better color consistency between the de-rained sub-views and the input rainy sub-views, but the loss of background details is more severe due to the filtering operation.

In order to preserve both color information and background textural details, we perform a weighted summation of C_k^r and F_k^r to obtain our final colored-residue image. Thus, the colored-residue image of the k -th rainy sub-view can be expressed as:

$$\hat{C}_k = \mu \cdot C_k^r + (1 - \mu) \cdot F_k^r, \quad (19)$$

where μ is a hyper-parameter used to balance C^r and F^r for the formation of the final colored-residue image \hat{C} . For more details on the formation and visualization of colored-residue images used in our method, please refer to the Supplementary Material.

3.6 Loss Function

Based on our proposed semi-supervised learning framework, we use supervised and unsupervised losses to separately train the corresponding branches. Next, we will introduce these two types of loss functions respectively.

Supervised Loss. In the supervised branch, we combine smooth L_1 loss [17], SSIM loss, and contrastive regularization [58] to constrain our network, and each one is designed for a specific purpose. We first use a smooth L_1 loss between the de-rained image (B) and the ground truth (G), which is a good way to avoid the shortcomings of MSE and MAE loss, such as generating over-smoothed results and failing to distinguish the rain streaks and the background objects. Smooth L_1 loss is defined as follows:

$$\mathcal{L}_{smoothL_1} = \begin{cases} 0.5(B - G)^2 & \text{if } |B - G| < 1 \\ |B - G| - 0.5 & \text{otherwise.} \end{cases} \quad (20)$$

Second, to make the network generate high-quality results, we also adopt SSIM loss to constrain the structural similarities. The SSIM loss can be expressed as:

$$\mathcal{L}_{ssim} = 1 - \text{SSIM}(B, G). \quad (21)$$

Third, we employ contrastive regularization (CR) [58] to further improve the de-raining performance of our method. By adding a contrastive regularization term, our method is able to generate more natural de-rained images, and improve the visual quality of the restored images. CR is based on contrastive learning, which ensures that the predicted de-rained result is pulled closer to the clear image and pushed farther away from the rainy image in the representation space. Therefore, CR needs to construct ‘positive’ and ‘negative’ pairs and find the potential feature space of these pairs for contrast. In the LF image rain removal task, the positive pair is generated by the group of clean LFI (G) and de-rained LFI (B), and the negative pair is generated by the group of B and a

rainy LFI (I). The same fixed pre-trained model $M(\cdot)$ (e.g., VGG-19 [44]) is used as a latent feature space for common intermediate feature extraction. Therefore, CR can be formulated as:

$$\mathcal{L}_{cr} = \sum_{i=1}^n w_i \cdot \frac{D(M_i(G), M_i(B))}{D(M_i(I), M_i(B))}, \quad (22)$$

where M_i , $i = 1, 2, \dots, n$ denotes extracting the i -th hidden features from the pre-trained model. w_i is the weight coefficient, and $D(\cdot)$ denotes the L1 distance.

Finally, the supervised loss is obtained by adding the above three functions:

$$\mathcal{L}_{sup} = \mathcal{L}_{smoothL_1} + \lambda_{ssim} \mathcal{L}_{ssim} + \lambda_{cr} \mathcal{L}_{cr}, \quad (23)$$

where λ_{ssim} and λ_{cr} are two hyper-parameters used to control the contributions (weights) of SSIM loss and CR.

Unsupervised Loss. For the unsupervised branch, we introduce multi-level KL loss, colored rain-free residue image-guided contrastive regularization (UCR) and total variation loss (TV loss) to constrain the real-world de-rained image generation process. In our UCR, the colored rain-free residue image (\hat{C}) is regarded as a “positive sample”. During the contrastive learning, the de-rained image is close to \hat{C} but away from the input rain image in the feature spaces. Therefore, UCR can be defined as:

$$\mathcal{L}_{ucr} = \sum_{i=1}^n w_i \cdot \frac{D(M_i(\hat{C}), M_i(B))}{D(M_i(I), M_i(B))}. \quad (24)$$

To further remove residual rain streaks and preserve structures and details of real-world rainy images, we adopt a TV loss. The TV loss is defined as follows:

$$\mathcal{L}_{tv} = \|\nabla_h B\|_1 + \|\nabla_v B\|_1, \quad (25)$$

where ∇_h and ∇_v denote the horizontal and vertical differential operators, respectively.

Overall, the loss of unsupervised branch shown in Figure 9 can be expressed as:

$$\mathcal{L}_{unsup} = \alpha \mathcal{L}_{kl} + \beta \mathcal{L}_{ucr} + \lambda_{tv} \mathcal{L}_{tv}, \quad (26)$$

where α , β , and λ_{tv} are three hyper-parameters.

Total Loss. To sum up, we combine supervised and unsupervised losses as our total loss, which can be described as:

$$\mathcal{L}_{total} = \mathcal{L}_{sup} + \mathcal{L}_{unsup}. \quad (27)$$

4 Experiments and Analysis

We compare our method with the state-of-the-art single image de-raining methods [6, 10, 40–42, 59, 65, 68, 69] and LF image de-raining methods [12, 34, 35, 37, 54, 61]. For comparing with the single image de-raining methods, every sub-view of an LFI is taken as a training sample. Since there are few methods for LFI de-raining, four LFI spatial super-resolution methods, DistgSSR [54], SASRNet [34], DSMNet [37] and PDistgNet [35], which take the MPI of an LF image as input, are slightly modified by removing the upsampling operations and learned to recover rainy LF images. All methods were evaluated under a unified protocol, including angular resolution and patch cropping.

Table 1. De-raining results of different methods on the datasets RLFDB [12] and RLMB [61] were evaluated using PSNR and SSIM metrics. The best and second-best results are marked in **red** and **cyan**, respectively. \uparrow means the higher the better, and our method achieves the best performance on both datasets. GFLOPs and inference time were measured on a 320×320 image patch.

Input Image	Methods	Training Type	RLFDB		RLMB		GFLOPs \downarrow	Time(ms) \downarrow
			PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow		
Single Image	PreNet [42]	Supervised	30.13	0.942	29.69	0.954	103.81	23.66
	Syn2Real [65]	Semi-supervised	28.94	0.923	26.74	0.883	33.24	10.87
	RecDerain [41]	Supervised	30.18	0.946	30.24	0.957	306.94	28.45
	CCN [40]	Supervised	30.22	0.936	29.08	0.959	290.84	40.29
	MPRNet [69]	Supervised	31.12	0.910	30.30	0.944	857.27	35.57
	IDT [59]	Supervised	30.02	0.931	31.84	0.960	90.50	177.75
	Restormer [68]	Supervised	30.33	0.914	30.74	0.951	255.04	222.65
	NeRD-Rain [6]	Supervised	30.93	0.945	31.06	0.957	231.22	181.68
	PoolNet [10]	Supervised	28.29	0.923	28.47	0.931	103.36	23.48
LF Image	Ding <i>et al.</i> [12]	Semi-supervised	31.81	0.936	28.46	0.940	259.25	218.94
	DistgSSR [54]	Supervised	28.18	0.897	29.85	0.948	549.93	71.24
	SASRNet [34]	Supervised	26.98	0.873	28.19	0.922	3004.14	526.54
	DSMNet [37]	Supervised	27.28	0.884	28.43	0.931	1,228.92	2,569.28
	Yan <i>et al.</i> [61]	Semi-supervised	30.78	0.931	29.89	0.959	341.54	101.46
	PDistgNet [35]	supervised	28.56	0.933	29.67	0.947	62.76	15.03
	Ours	Semi-supervised	32.19	0.950	32.27	0.962	189.75	23.39

4.1 Datasets

To verify the effectiveness of our *MDeRainNet*, we conduct experiment on two recently proposed light field rainy datasets [12, 61]. Ding *et al.* [12] proposed the first LF rainy dataset, termed RLFDB. RLFDB contains 80 synthetic and real-world rainy LF images, but there is no paired ground-truth for real-world scenes. In their synthetic rainy LF image generation approach, the particle system of Blender is adopted to simulate synthetic rain streaks with varying shapes, directions and densities, which are then superimposed onto real-world LFI to produce realistic rain streak LFIs. Later, Yan *et al.* [61] further considered the motion blur of rain streaks and global atmospheric light, so they proposed a new and large rainy LF image dataset, called RLMB, including 400 synthetic rainy LFI and 200 real-world LF images captured by a Lytro ILLUM camera without ground-truth LF images.

4.2 Implementation Details

Our proposed network with the semi-supervised learning framework, called *MDeRainNet*, is implemented with Pytorch, and all experiments are conducted on a PC with an NVIDIA GeForce RTX 3090 GPU card. Our network is trained using the Adam optimizer [26] for a total of 200 epochs with the batch size set to 6. The initial learning rate is $4.0e-5$, and the learning rate is reduced by a factor of 0.5 every 45 epochs. In Equations 19, 23 and 26, the hyper-parameters μ , λ_{ssim} , λ_{cr} , λ_{tv} are empirically set to 0.5, 1, 0.1, and $1.0e-6$, respectively. Inspired by [33], we do not set α and β in Equation 26 to fixed values, but change them dynamically (please refer to our Supplementary Material to see more details). In the training phase, we keep the angular resolution at 5×5 and crop each sub-view into patches of 64×64 with a stride of 32. Then, an array of patches is reorganized into MPI pattern (the resolution of the MPI is 320×320) to form the input MPI of our network. In experiments, the peak signal-to-noise ratio (PSNR) and structural similarity (SSIM) metrics are used to evaluate the performance of all de-raining methods.

4.3 Comparison to the State-of-the-Arts

A total of nine regular image de-raining methods [6, 10, 40–42, 59, 65, 68, 69] and four LFI de-raining methods [12, 34, 35, 37, 54, 61] are chosen to compare with our proposed *MDeRainNet* by evaluating their performance on two public LF rainy datasets, RLFDB [12] and RLMB [61]. The number of rainy LFIs in RLFDB [12] is relative small and rain streaks are usually thin and sparse. In contrast, RLMB [61] contains a large number of rainy LFIs, and the rain streaks are always large and dense.

4.3.1 Quantitative Evaluation. Table 1 reports the PSNR and SSIM values of different methods evaluated on two datasets. It can be seen that our method obviously outperforms all the other competing methods. The performances of the competing single image de-raining methods [6, 10, 40–42, 59, 65, 68, 69] do not perform well. The shortcoming of such single image de-raining methods is exposed, while learning on a small sub-set. In another word, these single image de-raining methods rely more on the visual information of a single input image to recover the regions contaminated by rain streaks, which may lead to unnatural or unreliable de-rained results. In contrast, better rain removal results could be obtained by the LFI rain removal methods, since the LFI rain removal methods can fully exploit the complementary information of the input LFI rather than relying solely on the visual information of single images. While testing on the RLMB dataset and measuring the results by PSNR metric, our method outperforms the single image rain removal methods IDT [59] and Restormer [68] by 0.43dB and 1.53dB, respectively, and also outperforms the LFI rain removal methods [12, 61] by 2.38dB and 3.81dB, respectively.

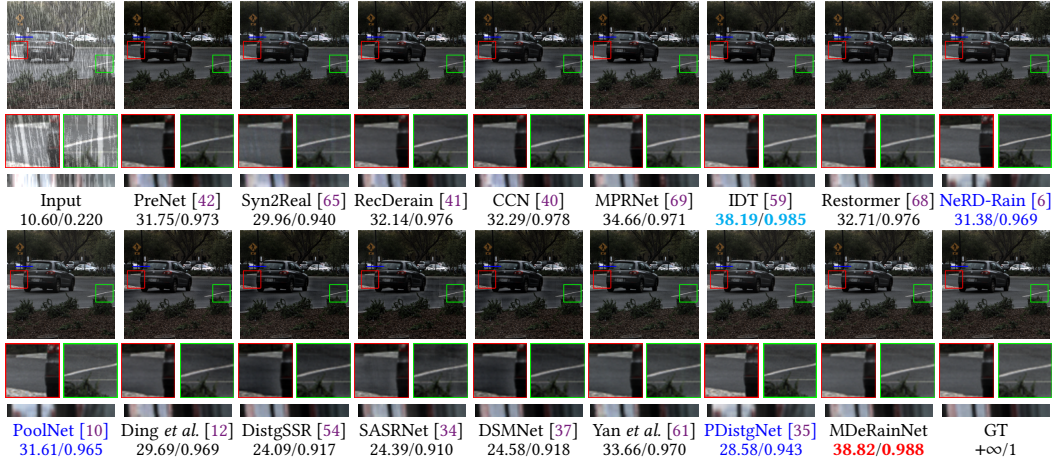


Fig. 10. De-raining results of different de-raining methods on synthetic LFIs from RLMB [61]. The de-rained central sub-view as well as the zoomed-in patches and an EPI produced by each method are shown. The horizontal EPIs extracted along the blue lines across the sub-views in the same row are also visualized to show the angular (image content) consistency between different sub-views. The PSNR and SSIM scores obtained by different methods on the presented scenes are reported under the zoomed-in area. The best value is colored in red, and the second best value is colored in cyan.

4.3.2 Qualitative Evaluation. We visualize the partial prediction results of different methods on RLMB [61], as shown in Figure 10. Since the rainy image contains large and/or dense rain streaks, some competing methods [6, 10, 12, 35, 59, 61, 65, 68, 69] do not perform well. Concretely, some large rain streaks still left in the de-rained results obtained by the competing methods [6, 10, 34, 35,




Fig. 11. Visual results obtained by different methods tested in the real-world rainy LF images coming from RLMB [61]. For each scene, the de-rained central sub-view as well as the zoomed-in patches and an EPI obtained by each method, are shown. For more results, please refer to our supplemental material.

37, 42, 54, 61, 65, 68, 69], especially the zoom-in patches marked by red and green color boxes and the details in the background can not be well recovered. The exhibited EPIs obtained by our method demonstrate that our method can well preserve the consistency of the de-rained sub-views of the input rainy LFI. Concretely, our method outperforms the competing methods [6, 10, 12, 34, 35, 37, 40–42, 59, 65, 68, 69], and is comparable with the competing methods [54, 61].

We also compare our method to the competing methods by evaluating them on several real-world rainy LF images from the test set of RLMB [61], as shown in Figure 11. Our *MDeRainNet* outperforms all supervised methods [54, 68] and the semi-supervised learning de-raining method [12, 61, 65]. There are still some rain streaks left in the de-rained images produced by other methods [12, 54, 65, 68]. Moreover, although the method [61] also removes most of the rain streaks, texture details of the de-rained results seem to be over-smoothed and are not as sharp as ours.

We also visualize the horizontal EPI under the zoomed-in patches of results produced by different de-raining methods. Compared with the competing methods [12, 42, 61, 65, 69], our method can generate straighter and clearer line patterns with richer textures. Though the state-of-the-art method [61] removes the most of rain streaks in real-world LFIs, its EPI is a little blurry and some distortions are introduced into line patterns, as shown in Figure 11.

4.3.3 Efficiency and Equilibrium. We compare our method with several competing methods in terms of run-times and PSNR/SSIM, as shown in Table 1. It can be seen that our method obtains the highest average PSNR and SSIM scores while taking less running time to process each rainy LF image.

The consistency between different sub-views is also investigated. An example is shown in Figure 12, where a synthetic rainy LF image (the scene of Figure 10). The PSNR values for every de-rained sub-view are exhibited in the colored map. The average PSNR value and the standard deviation (Std) of PSNR for all sub-views of the de-rained LF image are also reported. It can be seen that our proposed method obtains the best average PSNR value and Std of PSNR value.




Fig. 12. Visualization of PSNR values calculated from different sub-views of the de-rained LF image. The scene comes from Figure 10. For each method, we also exhibit the average (Avg) and standard deviation (Std) of PSNR values distributed among different sub-views, and the standard deviation can represent the consistency of different de-rained sub-views. Our MDeRainNet achieves the best de-raining results with a relatively balanced distribution among sub-views.

Table 2. The PSNR and SSIM values obtained by several variants of our proposed MDeRainNet tested on the RLMB dataset [61]. “MDB→Distg-Block” refers to replacing our MDB with Distg-Block in [54]. “w/o ESAI” means that the encoder of our network does not use ESAI module. “w/o CCR” means that Colored-Residue Contrastive Regularization is removed while all other settings and the network remain unchanged. “sub-views input” means only the input format is changed to sub-views while keeping the network identical. Baseline means the network composed of Swin Transformer. The best performances are marked in **bold**.

Method	FLOPs↓	PSNR↑	SSIM↑
MDB→Distg-Block [54]	186.22G	31.83	0.958
w/o ESAI	56.93G	27.24	0.896
w/o CCR	189.75G	29.80	0.942
sub-views input	189.75G	25.84	0.831
Baseline	45.96G	28.54	0.916
Ours	189.75G	32.27	0.962

Since the recently proposed method [61] and DistgSSR [54] fully exploit angular information of the input rainy LFI, they also obtain better Std of PSNR values, i.e., relatively balanced PSNR distributions among different sub-views. In contrast, since the single image de-raining methods, MPRNet [69], Restormer [68] and IDT [59], ignore the relationship among different sub-views (angular information) of an LF image, the PSNR values calculated from the different sub-views of the de-rained LF image result in an imbalanced PSNR distribution. Thus, though IDT [59] obtains the second-best Std of PSNR value, it produces a poor PSNR value. More results can be found in our Supplementary Material.

4.4 Ablation Study

In this subsection, we compare our MDeRainNet with its different variants to investigate the benefits of the modules of our network. In addition, we investigate the influence of different angular resolutions of input LF images on the performance of our MDeRainNet. The quantitative results are exhibited in Tables 2, 3 and 4. It should be noted that a baseline composed of Swin Transformer was introduced as a fundamental reference in Table 2. To further demonstrate the effectiveness of our ESAI and MDB, we visualize the derained results in Figure 13.

4.4.1 Influence of input format. To quantify the sole contribution of MPI, we performed an input-format ablation while keeping the full network intact. Table 2 shows that replacing MPI with sub-aperture views drops PSNR by 6.43 dB and SSIM by 0.131, owing to the disruption of the fixed spatial-angular correspondence that ESAI and MDB exploit. Experimental results quantitatively

Table 3. The PSNR and SSIM values obtained by several variants of the attention unit in SAIA tested on the RLMB dataset [61]. * indicates that the execution order of cross-feature attention and feature reinforcement attention of the attention unit in SAIA is exchanged. The best performances are marked in **bold**.

	Cross-Feature Attn	Feature Reinforcement Attn	AngPos	SpaPos	FLOPs↓	PSNR↑	SSIM↑
1		√	√	√	146.06G	30.93	0.950
2	√		√	√	146.09G	31.32	0.954
3	√	√			189.75G	30.84	0.948
4	√*	√*	√	√	189.75G	31.56	0.957
5	√	√	√	√	189.75G	32.27	0.962

demonstrate that, within the proposed architecture, MPI input synergistically interacts with the specifically designed ESAI and MDB modules to significantly boost performance.

4.4.2 Modified Disentangling Block (MDB). To validate the benefits of the MDB in our MDeRainNet, we replace the MDB in our network with the Distg-Block [54]. As shown in Table 2, the performance degrades by 0.44dB and 0.004 in terms of PSNR and SSIM, respectively, which indicates that MDB is beneficial to the performance of our MDeRainNet. In fact, MDB is a variant of the Distg-Block [54]. MDB abandons the EPI Feature Extractor (EFE) of the Distg-Block, but increases the number of channels of the angular features extracted by the Angular Feature Extractor (AFE). The angular features extracted by AFE are more comprehensive than that extracted by EFE. AFE can efficiently integrate angular information from all sub-views, while EFE can only incorporate angular information from EPIs. In addition, to make the complementary angular information better guide spatial information extraction, a spatial residual block consisting of two SFEs and a residual connection is added to our MDB.

4.4.3 Extended Spatial-Angular Interaction (ESAI). As a core module of our network, the ESAI module can fully model the long-range correlation between spatial and angular information for integrating them. To evaluate the effectiveness of our ESAI, we remove the ESAI module from our network while keeping all other blocks unchanged, as shown in Table 2. It can be seen that removing the ESAI module significantly reduces the de-raining performance of our MDeRainNet, with PSNR and SSIM degraded by 5.03dB and 0.066, respectively.

Attention Unit of SAIA. As shown in Table 3, we have also verified the effectiveness of the attention unit in our SAIA. While using cross-feature attention alone or feature reinforcement attention alone, the performance of the variant of our network is inferior to that of the complete Attention Unit of SAIA. On the other hand, switching the execution order of *Cross-Feature Attention* and *Feature Reinforcement Attention* also degrades the performance of our MDeRainNet.

Spatial and Angular Positional Encoding. In Transformers, each flattened patch is regarded as token, and its position information is crucial for networks to capture the image semantics. To model the position correlation of spatial feature and angular feature of the same patch, two learnable positional encodings [3, 13] for spatial and angular features are added to the corresponding feature, respectively.) As shown in Table 3, compared with the variant of our MDeRainNet with position encoding, it without position encoding significantly degrades the PSNR value (> 1.4 dB).

4.4.4 Angular Resolution. In Table 4, we analyze the performance of our MDeRainNet while taking different angular resolutions of the input LFI as input. Specifically, the central $A \times A$ ($A = 3, 5, 7, 9$) sub-views of each LFI from the RLMB [61] is taken as input for learning and testing our network. In general, the PSNR value increases along with the increasing of the angular resolution of input LFIs. This is because the newly added sub-views would provide more angular information, which

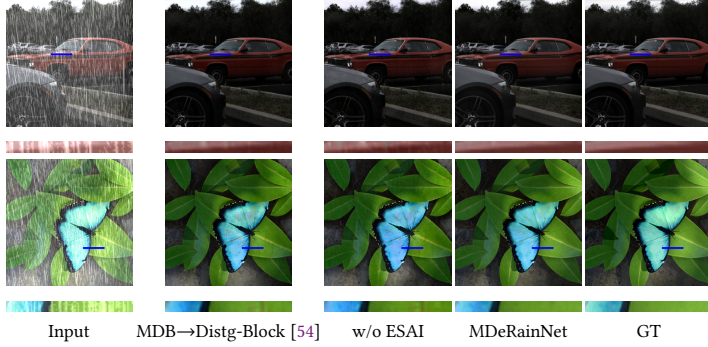


Fig. 13. Derained visualizations of several ablation studies.

Table 4. The PSNR and SSIM results obtained by our MDeRainNet on the RLMB dataset [61] with different angular resolutions. The best performances are marked in **bold**.

	3×3	5×5	7×7	9×9
PSNR \uparrow	31.48	32.27	32.55	32.68
SSIM \uparrow	0.956	0.962	0.965	0.968

facilitates the detection of rain streaks and background image recovery. It is worth noting that while the angular resolution increases from 5×5 to 7×7 and from 7×7 to 9×9 , the performance improvement gradually saturates (measured by the PSNR, increased by only 0.28dB and 0.13dB, respectively). Thus, utilizing much more sub-views (e.g., 9×9) may be not the best choice, since it requires greater computational cost and memory but with slight gain. To balance the computational complexity and performance of our network, only the central 5×5 sub-views of a rainy LFI are chosen as the input of our network.

4.5 Limitation

Although our proposed MDeRainNet demonstrates superior performance across most synthetic and real-world test scenes, some limitations persist when processing certain real-world rainy images with complex background textures. Our method cannot clearly remove the rain streaks and recover the obscured background, as the enlarged region highlighted with a red rectangle in Figure 14. Residual rain streaks can still be observed in our results. By comparison, the method [61] leaves




Fig. 14. A failure case in real-world rainy LFI from RLMB [61].

less pronounced residual rain streaks, yielding a slightly better result than our approach. The incomplete removal of rain streaks can likely be attributed to severe occlusion across all sub-views, which prevents the Macro-pixel Image (MPI) from retaining sufficient background information for the model to effectively restore the scene.

5 Conclusion

In this paper, we have proposed a novel light field rain streak removal network, called *MDeRainNet*, which takes the MPI of an rainy LFI as input. Our network can disentangle and incorporate the spatial feature and the angular feature from the MPI by introducing a Modified Disentangling Block (MDB). Moreover, we have also proposed a Transformer-based feature interaction (ESAI) mechanism to fully model long-range correlations between spatial and angular information. Our ESAI is not only able to effectively incorporate spatial and angular information, but also solves the limitation of receptive field of CNNs. Furthermore, a semi-supervised learning paradigm for rain removal from light field images is proposed for improving the generalization of our method on real-world rainy LF images. Extensive experiments have demonstrated the effectiveness and superiority of our proposed method compared with competing methods. In the future, we will investigate how to reduce the GPU memory usage (i.e., number of parameters) of our network in order to deploy it in downstream computer vision systems.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant No. 61902151 and 62172190) and the Natural Science Foundation of Jiangsu Province, China (Grant No. BK20170197).

References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv:1607.06450* (2016).
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-end object detection with transformers. In *Proc. ECCV*.
- [3] Hanting Chen, Yunhe Wang, Tianyu Guo, Chang Xu, Yiping Deng, Zhenhua Liu, Siwei Ma, Chunjing Xu, Chao Xu, and Wen Gao. 2021. Pre-trained image processing transformer. In *Proc. CVPR*.
- [4] Weidong Chen, Xiaofeng Xing, Xiangmin Xu, and Jichen Yang. 2021. Key-sparse transformer with cascaded cross-attention block for multimodal speech emotion recognition. *arXiv: 2106.11532* (2021).
- [5] Xiang Chen, Hao Li, Mingqiang Li, and Jinshan Pan. 2023. Learning A Sparse Transformer Network for Effective Image Deraining. In *Proc. CVPR*.
- [6] Xiang Chen, Jinshan Pan, and Jiangxin Dong. 2024. Bidirectional Multi-Scale Implicit Neural Representations for Image Deraining. In *Proc. CVPR*. 25627–25636.
- [7] Xiang Chen, Jinshan Pan, Kui Jiang, Yufeng Li, Yufeng Huang, Caihua Kong, Longgang Dai, and Zhentao Fan. 2022. Unpaired deep image deraining using dual contrastive learning. In *Proc. CVPR*.
- [8] Junyan Cheng, Iordanis Fostiropoulos, Barry Boehm, and Mohammad Soleymani. 2021. Multimodal phased transformer for sentiment analysis. In *Proc. EMNLP*.
- [9] Xin Cui, Cong Wang, Dongwei Ren, Yunjin Chen, and Pengfei Zhu. 2022. Semi-supervised image deraining using knowledge distillation. *IEEE TCSVT* 32, 12 (2022), 8327–8341.
- [10] Yuning Cui, Wenqi Ren, and Alois Knoll. 2025. Exploring the Potential of Pooling Techniques for Universal Image Restoration. *IEEE Transactions on Image Processing* 34 (2025), 3403–3416.
- [11] Liang-Jian Deng, Ting-Zhu Huang, Xi-Le Zhao, and Tai-Xiang Jiang. 2018. A directional global sparse model for single image rain removal. *Applied Mathematical Modelling* 59 (2018), 662–679.
- [12] Yuyang Ding, Mingyue Li, Tao Yan, Fan Zhang, Yuan Liu, and Rynson WH Lau. 2021. Rain streak removal from light field images. *IEEE TCSVT* 32, 2 (2021), 467–482.
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv:2010.11929* (2020).
- [14] Xueyang Fu, Jiabin Huang, Xinghao Ding, Yinghao Liao, and John Paisley. 2017. Clearing the skies: A deep network architecture for single-image rain removal. *IEEE TIP* 26, 6 (2017), 2944–2956.

- [15] Xueyang Fu, Jiabin Huang, Delu Zeng, Yue Huang, Xinghao Ding, and John Paisley. 2017. Removing rain from single images via a deep detail network. In *Proc. CVPR*.
- [16] Xueyang Fu, Jie Xiao, Yurui Zhu, Aiping Liu, Feng Wu, and Zheng-Jun Zha. 2023. Continual image deraining with hypergraph convolutional networks. *IEEE TPAMI* 45, 8 (2023), 9534–9551.
- [17] Ross Girshick. 2015. Fast r-cnn. In *Proc. ICCV*.
- [18] Yun Guo, Xueyao Xiao, Yi Chang, Shumin Deng, and Luxin Yan. 2023. From sky to the ground: A large-scale benchmark and simple baseline towards real rain removal. In *Proc. ICCV*.
- [19] Jiaxuan He, Sijie Mai, and Haifeng Hu. 2021. A unimodal reinforced transformer with time squeeze fusion for multimodal sentiment analysis. *IEEE Signal Processing Letters* 28 (2021), 992–996.
- [20] Huaibo Huang, Aijing Yu, and Ran He. 2021. Memory oriented transfer learning for semi-supervised image deraining. In *Proc. CVPR*.
- [21] Kui Jiang, Zhongyuan Wang, Peng Yi, Chen Chen, Baojin Huang, Yimin Luo, Jiayi Ma, and Junjun Jiang. 2020. Multi-scale progressive fusion network for single image deraining. In *Proc. CVPR*.
- [22] Kui Jiang, Zhongyuan Wang, Peng Yi, Chen Chen, Zheng Wang, Xiao Wang, Junjun Jiang, and Chia-Wen Lin. 2021. Rain-free and residue hand-in-hand: A progressive coupled network for real-time image deraining. *IEEE TIP* 30 (2021), 7404–7418.
- [23] Dong Jing, Shuo Zhang, Song Chang, and Youfang Lin. 2022. Light Field Raindrop Removal via 4D Re-sampling. *arXiv:2205.13165* (2022).
- [24] Dong Jing, Shuo Zhang, Runmin Cong, and Youfang Lin. 2021. Occlusion-aware bi-directional guided network for light field salient object detection. In *Proc. ACM Multimedia*.
- [25] Li-Wei Kang, Chia-Wen Lin, and Yu-Hsiang Fu. 2011. Automatic single-image-based rain streaks removal via image decomposition. *IEEE TIP* 21, 4 (2011), 1742–1755.
- [26] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv:1412.6980* (2014).
- [27] Ruoteng Li, Loong-Fah Cheong, and Robby T Tan. 2019. Heavy rain image restoration: Integrating physics model and conditional adversarial learning. In *Proc. CVPR*.
- [28] Ruoteng Li, Robby T Tan, and Loong-Fah Cheong. 2018. Robust optical flow in rainy scenes. In *Proc. ECCV*.
- [29] Xuejian Li, Shiqiang Ma, Jijun Tang, and Fei Guo. 2022. TranSiam: Fusing multimodal visual features using transformer for medical image segmentation. *arXiv:2204.12185* (2022).
- [30] Yu Li, Robby T Tan, Xiaojie Guo, Jiangbo Lu, and Michael S Brown. 2016. Rain streak removal using layer priors. In *Proc. CVPR*.
- [31] Yuanchu Liang, Saeed Anwar, and Yang Liu. 2022. DRT: A Lightweight Single Image Deraining Recursive Transformer. In *Proc. CVPR*.
- [32] Zhengyu Liang, Yingqian Wang, Longguang Wang, Jungang Yang, and Shilin Zhou. 2022. Light field image super-resolution with transformers. *IEEE Signal Processing Letters* 29 (2022), 563–567.
- [33] Huangxing Lin, Yanlong Li, Xueyang Fu, Xinghao Ding, Yue Huang, and John Paisley. 2020. Rain o’er me: Synthesizing real rain to derain with data distillation. *IEEE TIP* 29 (2020), 7668–7680.
- [34] Gaosheng Liu, Huanjing Yue, Kun Li, and Jingyu Yang. 2024. Adaptive pixel aggregation for joint spatial and angular super-resolution of light field images. *Information Fusion* 104 (2024), 102183.
- [35] Gaosheng Liu, Huanjing Yue, and Jingyu Yang. 2024. Efficient Light Field Image Super-Resolution via Progressive Disentangling. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. 6277–6286.
- [36] Nian Liu, Ni Zhang, Kaiyuan Wan, Ling Shao, and Junwei Han. 2021. Visual saliency transformer. In *Proc. ICCV*.
- [37] Yutong Liu, Zhen Cheng, Zeyu Xiao, and Zhiwei Xiong. 2024. Light Field Super-Resolution Using Decoupled Selective Matching. *IEEE TCSVT* 34, 5 (2024), 3313–3326.
- [38] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Proc. NeurIPS*.
- [39] Yu Luo, Yong Xu, and Hui Ji. 2015. Removing rain from a single image via discriminative sparse coding. In *Proc. ICCV*.
- [40] Ruijie Quan, Xin Yu, Yuanzhi Liang, and Yi Yang. 2021. Removing raindrops and rain streaks in one go. In *Proc. CVPR*.
- [41] Dongwei Ren, Wei Shang, Pengfei Zhu, Qinghua Hu, Deyu Meng, and Wangmeng Zuo. 2020. Single image deraining using bilateral recurrent network. *IEEE TIP* 29 (2020), 6852–6863.
- [42] Dongwei Ren, Wangmeng Zuo, Qinghua Hu, Pengfei Zhu, and Deyu Meng. 2019. Progressive image deraining networks: A better and simpler baseline. In *Proc. CVPR*.
- [43] Wenzhe Shi, Jose Caballero, Ferenc Huszar, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. 2016. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proc. CVPR*.
- [44] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556* (2014).

- [45] Jeya Maria Jose Valanarasu, Rajeev Yasarla, and Vishal M Patel. 2022. Transweather: Transformer-based restoration of images degraded by adverse weather conditions. In *Proc. CVPR*.
- [46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proc. NeurIPS*.
- [47] Yecong Wan, Yuanshuo Cheng, Mingwen Shao, and Jordi González. 2022. Image rain removal and illumination enhancement done in one go. *Knowledge-Based Systems* 252 (2022), 109244.
- [48] Yecong Wan, Mingwen Shao, Yuanshuo Cheng, Yuexian Liu, and Zhiyuan Bao. 2025. Learning to Restore Arbitrary Hybrid adverse weather Conditions in one go. *Pattern Recognition* 164 (2025), 111504.
- [49] Guoqing Wang, Changming Sun, and Arcot Sowmya. 2020. Cascaded attention guidance network for single rainy image restoration. *IEEE TIP* 29 (2020), 9190–9203.
- [50] Guoqing Wang, Changming Sun, and Arcot Sowmya. 2021. Attentive feature refinement network for single rainy image restoration. *IEEE TIP* 30 (2021), 3734–3747.
- [51] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. 2021. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *Proc. ICCV*.
- [52] Xingzheng Wang, Jiehao Liu, Songwei Chen, and Guoyao Wei. 2022. Effective Light Field De-Occlusion Network based on Swin Transformer. *IEEE TCSVT* 33, 6 (2022), 2590–2599.
- [53] Yingqian Wang, Longguang Wang, Zhengyu Liang, Jungang Yang, Wei An, and Yulan Guo. 2022. Occlusion-Aware Cost Constructor for Light Field Depth Estimation. In *Proc. CVPR*.
- [54] Yingqian Wang, Longguang Wang, Gaochang Wu, Jungang Yang, Wei An, Jingyi Yu, and Yulan Guo. 2022. Disentangling light fields for super-resolution and disparity estimation. *IEEE TPAMI* 45, 1 (2022), 425–443.
- [55] Yingqian Wang, Longguang Wang, Jungang Yang, Wei An, Jingyi Yu, and Yulan Guo. 2020. Spatial-angular interaction for light field image super-resolution. In *Proc. ECCV*.
- [56] Zhendong Wang, Xiaodong Cun, Jianmin Bao, Wengang Zhou, Jianzhuang Liu, and Houqiang Li. 2022. Uformer: A general u-shaped transformer for image restoration. In *Proc. CVPR*.
- [57] Wei Wei, Deyu Meng, Qian Zhao, Zongben Xu, and Ying Wu. 2019. Semi-supervised transfer learning for image rain removal. In *Proc. CVPR*.
- [58] Haiyan Wu, Yanyun Qu, Shaohui Lin, Jian Zhou, Ruizhi Qiao, Zhizhong Zhang, Yuan Xie, and Lizhuang Ma. 2021. Contrastive learning for compact single image dehazing. In *Proc. CVPR*.
- [59] Jie Xiao, Xueyang Fu, Aiping Liu, Feng Wu, and Zheng-Jun Zha. 2022. Image de-raining transformer. *IEEE TPAMI* 45, 11 (2022), 12978–12995.
- [60] Tao Yan, Yuyang Ding, Fan Zhang, Ningyu Xie, Wenxi Liu, Zhengtian Wu, and Yuan Liu. 2019. Snow Removal From Light Field Images. *IEEE Access* 7 (2019), 164203–164215.
- [61] Tao Yan, Mingyue Li, Bin Li, Yang Yang, and Rynson WH Lau. 2023. Rain Removal from Light Field Images with 4D Convolution and Multi-scale Gaussian Process. *IEEE TIP* 32 (2023), 921–936.
- [62] Wending Yan, Robby T Tan, Wenhan Yang, and Dengxin Dai. 2021. Self-aligned video deraining with transmission-depth consistency. In *Proc. CVPR*.
- [63] Wenhan Yang, Robby T Tan, Shiqi Wang, and Jiaying Liu. 2020. Self-learning video rain streak removal: When cyclic consistency meets temporal correspondence. In *Proc. CVPR*.
- [64] Rajeev Yasarla and Vishal M Patel. 2019. Uncertainty guided multi-scale residual learning-using a cycle spinning cnn for single image de-raining. In *Proc. CVPR*.
- [65] Rajeev Yasarla, Vishwanath A Sindagi, and Vishal M Patel. 2020. Syn2real transfer learning for image deraining using gaussian processes. In *Proc. CVPR*.
- [66] Qiaosi Yi, Juncheng Li, Qinyan Dai, Faming Fang, Guixu Zhang, and Tiejong Zeng. 2021. Structure-preserving deraining with residue channel prior guidance. In *Proc. ICCV*.
- [67] Changfeng Yu, Yi Chang, Yi Li, Xile Zhao, and Luxin Yan. 2021. Unsupervised image deraining: Optimization model driven deep cnn. In *Proc. ACM Multimedia*.
- [68] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. 2022. Restormer: Efficient transformer for high-resolution image restoration. In *Proc. CVPR*.
- [69] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. 2021. Multi-stage progressive image restoration. In *Proc. CVPR*.
- [70] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. 2022. Learning enriched features for fast image restoration and enhancement. *IEEE TPAMI* 45, 2 (2022), 1934–1948.
- [71] He Zhang, Vishwanath Sindagi, and Vishal M Patel. 2019. Image de-raining using a conditional generative adversarial network. *IEEE TCSVT* 30, 11 (2019), 3943–3956.
- [72] Lei Zhu, Chi-Wing Fu, Dani Lischinski, and Pheng-Ann Heng. 2017. Joint bi-layer optimization for single-image rain streak removal. In *Proc. ICCV*.