Restoring Images in Adverse Weather Conditions via Histogram Transformer

Shangquan Sun^{1,2} , Wenqi Ren^{3,4†} , Xinwei Gao⁵, Rui Wang^{1,2}, and Xiaochun Cao³

¹ Institute of Information Engineering, Chinese Academy of Sciences, Beijing 100085, China

 $^2\,$ School of Cyber Security, University of Chinese Academy of Sciences, Beijing 100049, China

³ School of Cyber Science and Technology, Shenzhen Campus of Sun Yat-sen University, Shenzhen 518107, China

⁴ Guangdong Provincial Key Laboratory of Information Security Technology, Guangzhou 510006, China

⁵ Wechat Business Group, Tencent, Shenzhen, Guangdong, China, 518057 shangquansun@gmail.com, renwq3@mail.sysu.edu.cn,

Abstract. Transformer-based image restoration methods in adverse weather have achieved significant progress. Most of them use self-attention along the channel dimension or within spatially fixed-range blocks to reduce computational load. However, such a compromise results in limitations in capturing long-range spatial features. Inspired by the observation that the weather-induced degradation factors mainly cause similar occlusion and brightness, in this work, we propose an efficient Histogram Transformer (Histoformer) for restoring images affected by adverse weather. It is powered by a mechanism dubbed histogram selfattention, which sorts and segments spatial features into intensity-based bins. Self-attention is then applied across bins or within each bin to selectively focus on spatial features of dynamic range and process similar degraded pixels of the long range together. To boost histogram selfattention, we present a dynamic-range convolution enabling conventional convolution to conduct operation over similar pixels rather than neighbor pixels. We also observe that the common pixel-wise losses neglect linear association and correlation between output and ground-truth. Thus, we propose to leverage the Pearson correlation coefficient as a loss function to enforce the recovered pixels following the identical order as groundtruth. Extensive experiments demonstrate the efficacy and superiority of our proposed method. We have released the codes in Github.

Keywords: Image restoration · Adverse weather removal · Image Desnowing · Image deraining · Image dehazing · Raindrop removal

[†] Corresponding author.



Fig. 1: Given weather-degraded images in (a), traditional transformers perform selfattention either along the channel dimension or within a fixed-range block as shown in (b). In contrast, we observe that weather-induced degradation patterns tend to be similar but distinct from the background. So we categorize pixels affected by adverse weather and background pixels into distinct bins based on descending intensities (as depicted in (c)) and then conducts self-attention within and between bins.

1 Introduction

The field of computer vision witnessed growing interest in restoring images affected by adverse weather conditions like rain, fog, and snow. These weather conditions significantly degrade visual quality, impacting the performance of downstream tasks such as object detection [3,60], and depth estimation [16,18]. The restoration of images under adverse weather is thereby a vital problem for the sake of vision aesthetics and safety.

Early works leverage weather-related priors to model statistical characteristic of degradation and remove adverse weathers [1, 2, 22, 23, 35, 76, 83, 95, 99]. Subsequently, convolutional neural networks (CNNs) have emerged as powerful tools for addressing deraining [5, 17, 34, 40, 57, 73, 79, 80, 88, 89], dehazing [27, 29, 62–64, 77, 88, 90, 92] and desnowing [30, 44, 61, 94]. However, the need of separately training networks for each task and the complexity of switching among multiple models present challenges for real-world applications. Li *et al.* [33] thus introduced the challenge of adverse weather removal, which entails the restoration of images affected by various weather conditions using a single unified model.

Recently, Transformer-based approaches have also been investigated for the adverse weather removal task, surpassing the efficacy of CNNs [11, 19, 70, 72]. Nonetheless, these Transformer-based methods usually make concessions regarding efficient memory utilization by confining self-attention operations to a fixed spatial range or solely within the channel dimension, as depicted in Figure 1b. This compromise impedes the inherent potential of Transformers, which was originally designed for superior global feature modeling, and consequently, it leads to a deterioration in the performance of restoration.

To address these problems, based on the observation that weather-induced degradation often exhibits common patterns shown in Figure 1a, we develop an efficient **Histo**gram Transformer for unified adverse weather removal, named

Histoformer. Specifically, we introduce a Dynamic-range Histogram Self-Attention (DHSA) module, which endows self-attention with a dynamic-range spatial receptive field. We categorize pixel values proximate in intensity yet varied in spatial location into histogram bins. Self-attention is executed across the dimension of bin or frequency, whose process is illustrated in Figure 1c. To facilitate comprehensive feature extraction on both local and global scales, we devise two ways of reshaping for histogram self-attention: bin-wise histogram reshaping (BHR) and frequency-wise histogram reshaping (FHR). In BHR, the number of bins is configured to incorporate pixels spanning a more comprehensive intensity range, thereby facilitating global feature integration. In FHR, the number of frequencies is assigned such that each bin focuses on limited number of pixels, enhancing the utility of finer features. Consequently, the histogram self-attention attains the capability of modeling spatially dynamic ranges effectively.

To enable the convolution to extract dynamically-located weather-related dependencies, we develop a dynamic-range convolution layer, which involves sequential horizontal and vertical pixel sorting prior to the application of separable convolution. In order to capture multi-scale and multi-range information embedded within feature matrices, we introduce a Dual-scale Gated Feed-Forward (DGFF) module, enhancing its ability to model the visual characteristics effectively. Additionally, we note that conventional loss functions primarily focus on pixel-level closeness, overlooking the correlation at overall patch level. Consequently, we propose to leverage the Pearson correlation coefficient [12] to ensure the reconstruction of the linear relationship between restored and clean images.

Our contribution can be summarized in three folds:

- We propose a novel transformer targeted for unified adverse weather removal, equipped with a new histogram self-attention. It possesses dynamicrange spatial attention to weather-induced obstructions and thus can achieve degradation removal globally and efficiently.
- To capture multi-range information, we present a dual-scale feed-forward module. To enhance the comprehensive linear association between the recovered and ground-truth images, we develop a correlation loss.
- Our method attains state-of-the-art performance across various datasets. Additionally, we substantiate the efficacy of the proposed approach to restore real-world images and bolster the downstream application of detection.

2 Related Work

Extensive research has been dedicated to addressing adverse weather removal challenges in computer vision, including tasks like deraining [5, 17, 34, 40, 66, 73, 79, 80, 88, 89], dehazing [27, 29, 62–64, 67, 77, 88, 90, 92], desnowing [30, 44, 61, 94], raindrop removal [57, 59, 83, 93] and All-in-One weather removal [28, 33, 53, 70].

Rain Streak Removal. The evolution of approaches is notable in rain streak removal techniques in computer vision. Kang et al. [23] pioneered a single image

4 S. Sun et al.

deraining method using bilateral filters to decompose images into low and highfrequency components. However, recent advancements have seen a dominance of deep neural networks. An early deep CNN was introduced by Fu *et al.* [17] for extracting features from the high-frequency rain component, while Yang *et al.* [79] utilized recurrent networks to decompose rain layers and remove various streak types. Li *et al.* [32] proposed a method that addresses rain streaks and veiling effects in heavy rain scenes by integrating physics-based rain models and adversarial learning. A conditional generative adversarial network was also employed to solve rain streak removal [89]. Yasarla *et al.* [81] explored Gaussian processes for transfer learning from synthetic to real-world rain data. Quan *et al.* [58] used a cascaded network to remove both rain streaks and raindrops. Recently an image deraining Transformer [78] featuring a dual Transformer architecture was intricately formulated, incorporating both window-based and spatial-based mechanisms, thereby attaining exemplary outcomes. A sparse deraining Transformer is also proposed to enhance feature aggregation [11].

Raindrop Removal. Raindrop removal from single images has been addressed through various methods, with some relying on traditional hand-crafted features. An early work incorporated temporal information to address video-based raindrop removal [83]. Eigen *et al.* [15] employed a shallow CNN trained with image pairs containing raindrop-degraded and raindrop-free versions, though the results often exhibited blurriness. Qian *et al.* [57] introduced an attention GAN and a new dataset. Their method was later improved by Quan *et al.* [59] via generating attention maps based on mathematical raindrop descriptions and combining them with detected raindrop edges.

Snow Removal. Desnow-Net [44] was among the pioneering CNN-based approaches for snow removal, followed by Li *et al.*'s stacked dense network [31] and Chen *et al.*'s JSTASR [8], which introduced a size and transparency aware method. More recently, DDMSNet [94] introduced a dense multi-scale network that leverages semantic and geometric priors to enhance snow removal. A hierarchical decomposition paradigm involving the dual-tree wavelet transform for snow removal is also proposed [9]. Chen *et al.* [7] designed SnowFormer, a framework that used cross-attentions to establish local-global context interaction.

Fog Removal. Li et al. [26] presented a CNN that takes into account both atmospheric luminosity and transmission maps to conduct dehazing. Ren et al. [63] advocated a pre-processing approach for hazy image manipulation, thereby engendering multiple input modalities and, in the process, inducting chromatic aberrations as part of their dehazing procedure. A hierarchical density-aware network is also introduced, specializing in the domain of image dehazing [92]. Zheng et al. [97] formulated a curriculum-based contrastive regularization dehazing method aimed at fostering agreement within a contrastive space.

All-in-One Weather Removal. Some recent works attempted to address various weather-induced degradations by a singular network. Li et al. [33] proposed an

All-in-One network, containing a generator comprising multiple task-specific encoders and a shared decoder. Valanarasu *et al.* [70] presented TransWeather, a transformer-based model featuring a solitary encoder-decoder structure, capable of rejuvenating images afflicted by various atmospheric conditions. A pipeline for the automatic selection of weather-degraded data was also proposed to enhance existing models [91]. Zhu *et al.* developed WGWS-Net [100] capable of learning weather-general and weather-specific in two separate stages. Some other recent works also trial addressing adverse weather removal by adopting probabilistic denoising diffusion model [53], knowledge distillation [10], large-scale Pre-trained model [68], mixture of experts [49], few-shot learning [24], codebooks [41,71,82], adaptive filters [54], knowledge assignment [74] and domain translation [56].

Transformer-based Image Restoration. Since the inception of the Vision Transformer (ViT) [14] for visual recognition, transformers have gained substantial traction across a spectrum of computer vision tasks [25, 37, 45, 46, 50, 51, 86]. Particularly within the realm of low-level vision, the Image Processing Transformer [4] exemplifies how pre-training a transformer on extensive datasets can significantly enhance performance for low-level applications. U-former [75], on the other hand, introduced a transformer architecture based on the U-Net design for restoration tasks. Swin-IR [38] employed the Swin Transformer [45] for image restoration. Some latest Transformer-based methods were proposed for deraining [11, 39], desnowing [7], dehazing [19, 43, 65] and All-in-One weather removal [70, 72].

Unlike the existing Transfomer-based approaches whose self-attention is applied within either fixed spatial ranges or merely channel dimension, our method enables dynamic-range spatial attention to adaptively focus on weather-induced degradation with similar patterns.

3 Method

3.1 Overall Architecture

The overall architectural framework of our Histoformer is illustrated in Figure 2. Suppose the input is a low-quality image $I^{lq} \in \mathbb{R}^{3 \times H \times W}$, we pass the input through a 3×3 convolution to conduct the overlapping image patch-embedding. Within both the encoder and decoder of the network backbone, we arrange Histogram Transformer Blocks (HTBs) to extract intricate features and capture dynamically distributed degradation factors. Within the same stage, encoders and decoders are interlinked through skip-connections, thereby establishing connections between consecutive intermediate features to enhance the stability of the training process. Between each stage, we apply pixel-unshuffle and pixel-shuffle operations for the purpose of feature down-sampling and up-sampling.

Within each HTB, we introduce Dynamic-range Histogram Self-Attention (DHSA) to extract spatially dynamic weather degradation and enhance both local and global feature aggregation. Moreover, a Dual-scale Gated Feed-Forward



Fig. 2: The overall architecture of our Histoformer for weather removal. The main component is the Histogram Transformer block, and it comprises the Dynamic-range Histogram Self-Attention (DHSA) module and the Dual-scale Gated Feed-Forward (DGFF) module. Within DHSA, we present two types of reshaping mechanism, *i.e.*, Bin-wise Histogram Reshaping and Frequency-wise Histogram Reshaping.

(DGFF) module is integrated into the HTB to enrich the representation of multirange features, contributing to the process of image restoration. During each stage of encoding phases, our model is equipped with a crude skip-connection for supplementing original features from input, comprised of a sequence of operations, including average pooling, pixel-wise convolution, and depth-wise convolution. We start the crude skip-connection after the first stage, and this setup enables the encoders to focus more effectively on learning the weather-induced residuals. Through this hybrid formulation, Histoformer is empowered to exploit both the adaptive contents of weather-irrelevant background and the inherent characteristics of weather-degraded patterns, facilitating the separation of undesired degradation from the latent clear background.

3.2 Histogram Transformer Block

As the key component of our Histoformer, HTB contains two pivotal modules, *i.e.*, DHSA and DGFF. These two components are arranged to interact with layer normalization and can be formulated as follows:

$$F_l = F_{l-1} + \text{DHSA}\left(\text{LN}\left(F_{l-1}\right)\right),\tag{1}$$

$$F_l = F_l + \text{DGFF}\left(\text{LN}\left(F_l\right)\right),\tag{2}$$

where LN denotes layer normalization and F_l represents the feature at *l*-th stage. The details of DHSA and DGFF are presented in Section 3.2 and 3.2 respectively.

Dynamic-range Histogram Self-Attention To better capture dynamically distributed weather-induced degradation, we develop a Dynamic-range Histogram

Self-Attention (DHSA) module. This module consists of a process involving dynamic-range convolution, which reorders the spatial distribution of fractional features, and a dual-path histogram self-attention mechanism that combines global and local dynamic feature aggregation. Prior to the final output projection of a 1×1 point-wise convolution, the reordered features are sorted back into their original locations to maintain spatial consistency.

Dynamic-range Convolution. Traditional convolution operations employ fixed kernel sizes, resulting in a limited receptive field range and consequently perform local and small-range computations. This restricted scope of convolution, which primarily focuses on local information, does not naturally align with the self-attention mechanism's capacity to model long-range dependencies. To address this limitation, we devise a dynamic-range convolution technique by meticulously reordering the input features prior to the traditional convolution operation. Given an input feature $F \in \mathbb{R}^{C \times H \times W}$, we divide it into two branches along the channel dimension, namely F_1 and F_2 . For the first branch of features, we perform sorting operations both horizontally and vertically, subsequently concatenating the sorted features with the second branch of feature. The resulting recombined features are then passed through the subsequent separable convolution. The entire process is articulated as follows:

$$F_1, F_2 = \text{Split}(F), \ F_1 = \text{Sort}_v(\text{Sort}_h(F_1)),$$

$$F = \text{Conv}_{3\times3}^d(\text{Conv}_{1\times1}(\text{Concat}(F_1, F_2))),$$
(3)

where $\operatorname{Conv}_{1\times 1}$ is 1×1 point-wise convolution, $\operatorname{Conv}_{3\times 3}^d$ represents 3×3 depthwise convolution, Concat is the concatenation operation along channel, Split denotes the operation of splitting features along channel dimension, and $\operatorname{Sort}_{i\in\{h,v\}}$ represents the horizontal or vertical sorting operation. This approach organizes pixels of high and low intensities into regular patterns at the diagonal corners of the matrices, thereby allowing convolution to perform computations across dynamic ranges. Given that weather-induced degradation typically exhibits closely related patterns, degraded pixels tend to concentrate in neighboring locations, separated with those clean pixels. As a result, this arrangement enables convolution kernels to partially focus on preserving clean information and separately recovering degraded features.

Histogram Self-Attention. Existing vision Transformers [11,75,78,78,86,96] typically leverage fixed range of attention or merely the attention along channel dimension due to the compromise of computation and memory efficiency. However, the fixed setting restricts the self-attention to span adaptively long range to associate desired features. We notice that weather-induced degradation causes similar patterns and that those pixels containing either background feature or weather degradation of different intensities had better be assigned with various extents of attention. We thus propose a histogram self-attention mechanism to categorize spatial elements into bins and allocate varying attention within and across bins. For the sake of parallel computing, we set each bin contains identical number of pixels during implement.

8 S. Sun et al.

Given the output of dynamic-range convolution, we can separate them into Value feature $V \in \mathbb{R}^{C \times H \times W}$ and two pairs of Query-key $F_{QK,1}, F_{QK,2} \in \mathbb{R}^{2C \times H \times W}$ subsequently passing to two branches. We firstly sort the sequence of V and based on its index arrange the Query-Key pairs accordingly, expressed as follows:

$$V, d = \text{Sort} \left(\mathbf{R}_{C \times H \times W}^{C \times H W}(V) \right),$$

$$Q_1, K_1 = \text{Split} \left(\text{Gather} \left(\mathbf{R}_{C \times H \times W}^{C \times H W}(F_{QK,1}), d \right) \right),$$

$$Q_2, K_2 = \text{Split} \left(\text{Gather} \left(\mathbf{R}_{C \times H \times W}^{C \times H W}(F_{QK,2}), d \right) \right),$$
(4)

where $\mathbf{R}_{C\times,H\times W}^{C\times HW}$ represents the operation of reshaping features from $\mathbb{R}^{C\times H\times W}$ to $\mathbb{R}^{C\times HW}$, d is the index of sorted Value, and Gather denotes the operation of retrieving elements of tensor based on a given index.

Then given the number of bins B, we reshape the sorted features from $C \times HW$ into $C \times B \times HW/B$. To extract both global and local information, we define two types of reshaping, i.e., bin-wise histogram reshaping (BHR) and frequency-wise histogram reshaping (FHR). The first is to assign the number of bins equal to B and each bin contains HW/B elements, while the second is to set the frequency of each bin equal to B and the number of bins is HW/B. By this way, we can extract large-scale information by BHR where each bin contains large number of dynamically located pixels and fine-grained information by FHR where each bins contains modicum pixels neighboring in terms of intensity. The two pairs of Query-Key features are passed through two types of reshaping and subsequent self-attention process respectively, and their outputs are element-wisely multiplied to yield the final output. The process can be formulated as the following expressions:

$$A_{\rm B} = \operatorname{softmax} \left(\frac{\mathbf{R}_{\rm B}(Q_1)\mathbf{R}_{\rm B}(K_1)^{\top}}{\sqrt{k}} \right) \mathbf{R}_{\rm B}(V),$$

$$A_{\rm F} = \operatorname{softmax} \left(\frac{\mathbf{R}_{\rm F}(Q_2)\mathbf{R}_{\rm F}(K_2)^{\top}}{\sqrt{k}} \right) \mathbf{R}_{\rm F}(V),$$

$$A = A_{\rm B} \odot A_{\rm F},$$
(5)

where k is the number of heads, $\mathbf{R}_{i \in \{B,F\}}$ denotes the reshaping operation of either BHR or FHR, and $A_{i \in \{B,F\}}$ represents the obtained attention map.

Dual-scale Gated Feed-Forward Previous studies [11, 75, 78, 86] typically leverage single-range or single-scale convolution in the standard feed-forward network to bolster local context. Nonetheless, these methods often disregard the correlations among dynamically distributed weather-induced degradation. In practice, multi-scale information can be extracted by not only enlarging the kernel size but also leveraging the dilation mechanism [36, 84, 85]. As a result, we conceive a Dual-scale Gated Feed-Forward (DGFF) module, which integrates two distinct multi-range and multi-scale depth-wise convolution pathways within the transmission process.

Table 1: Quantitative comparisons on three weather removal tasks in terms of PSNR and SSIM, where higher values indicate better performance. The top halves of tables display the results of task-specific methods, while the bottom halves present evaluations of the unified multi-weather models. The best and the second best results are in **bold** and <u>underlined</u>. Those with * indicate the methods whose source codes are unavailable.

(a) Image Desnowing			(b) Deraining & Dehazing			(c) Raindrop Removal				
	Snow10 PSNR	00K-S [44 SSIM	ISnow10 PSNR	0K-L [44] SSIM		$\mathcal{O}utdoo$ PSNR	r-Rain [32] SSIM		$_{\mathrm{PSNR}}^{\mathrm{RainD}}$	rop [57] SSIM
SPANet [73] JSTASR [8] RESCAN [34] DesnowNet [44] DDMSNet [94] NAFNet [6] Restormer [86]	$\begin{array}{c} 29.92 \\ 31.40 \\ 31.51 \\ 32.33 \\ 34.34 \\ 34.79 \\ 36.01 \end{array}$	$\begin{array}{c} 0.8260\\ 0.9012\\ 0.9032\\ 0.9500\\ 0.9445\\ 0.9497\\ 0.9579\end{array}$	23.70 25.32 26.08 27.17 28.85 30.06 30.36	$\begin{array}{c} 0.7930 \\ 0.8076 \\ 0.8108 \\ 0.8983 \\ 0.8772 \\ 0.9017 \\ 0.9068 \end{array}$	CycleGAN [98] pix2pix [20] HRGAN [32] PCNet [21] MPRNet [87] NAFNet [6] Restormer [86]	$\begin{array}{c} 17.62 \\ 19.09 \\ 21.56 \\ 26.19 \\ 28.03 \\ 29.59 \\ 30.03 \end{array}$	0.6560 0.7100 0.8550 0.9015 0.9192 0.9027 0.9215	pix2pix [20] DuRN [42] RaindropAttn [59] AttentiveGAN [57] IDT [78] MAXIM [69] Restormer [86]	28.02 31.24 31.44 31.59 31.87 31.87 32.18	$\begin{array}{c} 0.8547 \\ 0.9259 \\ 0.9263 \\ 0.9170 \\ 0.9313 \\ 0.9352 \\ \underline{0.9408} \end{array}$
All-in-One [33]* TransWeather [70] Chen et al. [10] WGWSNet [100] WeatherDiff ₆₄ [53] WeatherDiff ₁₂₈ [53 AWRCP [82]* Histoformer (Ours)	32.51 34.42 34.31 35.83 35.02 <u>36.92</u> 37.41	0.9341 0.9469 0.9460 0.9566 0.9516 <u>0.9652</u> 0.9656	28.33 29.31 30.22 30.16 30.09 29.58 <u>31.92</u> 32.16	0.8820 0.8879 0.9071 0.9007 0.9041 0.8941 0.9341 <u>0.9261</u>	All-in-One [33]* TransWeather [70] Chen <i>et al.</i> [10] WGWSNet [100] WeatherDiff ₆₄ [53] WeatherDiff ₁₂₈ [53 AWRCP [82]* Histoformer (Ours)	24.71 28.83 29.27 29.32 29.64] 29.72 <u>31.39</u> 32.08	0.8980 0.9000 0.9147 0.9207 0.9312 0.9216 <u>0.9329</u> 0.9389	All-in-One [33]* TransWeather [70] Chen et al. [10] WGWSNet [100] WeatherDiff ₆₄ [53] WeatherDiff ₁₂₈ [53 AWRCP [82]* Histoformer (Ours)	31.12 30.17 31.81 <u>32.38</u> 30.71] 29.66 31.93 33.06	0.9268 0.9157 0.9309 0.9378 0.9312 0.9225 0.9314 0.9441

Given an input tensor $F_l \in \mathbb{R}^{C \times H \times W}$, we initially employ a point-wise convolution operation to augment the channel dimension by a factor of r. Following this augmentation, the expanded tensor is directed into two parallel branches. Throughout the feature transformation process, 5×5 and dilated 3×3 depthwise convolutions are employed to enhance the extraction of multi-range and multi-scale information. Following the gating mechanism [13], the output of the second branch after passing through an activation act as a gating map for the other branch. Thus, the complete feature fusion process within the DGFF module is formulated as follows:

$$F_{l,1}, F_{l,2} = \text{Split} \left(\text{Shuffle}(\text{Conv}_{1 \times 1}(F_l)) \right),$$

$$F_{l,1} = \text{Conv}_{5 \times 5}^{d}(F_{l,1}), \ F_{l,2} = \text{Conv}_{3 \times 3}^{d,\text{dilated}}(F_{l,2}),$$

$$F_{l+1} = \text{Conv}_{1 \times 1} \left(\text{Unshuffle} \left(\text{Mish}(F_{l,2}) \odot F_{l,1} \right) \right),$$
(6)

where $\operatorname{Conv}_{5\times 5}^{d}$ represents 5×5 depth-wise convolution, $\operatorname{Conv}_{3\times 3}^{d,\operatorname{dilated}}$ is 3×3 dilated depth-wise convolution, Shuffle and Unshuffle represent respectively the operations of pixel-shuffling and unshuffling, Mish denotes the Mish activation [52], and F_{l+1} is the output of current stage passing to l + 1-th stage.

3.3 Reconstruction Loss and Correlation Loss

We use the L_1 norm of the pixel-wise difference between the restored high-quality image I^{hq} and ground-truth I^{gt} as the reconstruction loss, i.e.,

$$\mathcal{L}_{rec} = \left\| I^{hq} - I^{gt} \right\|_{1}. \tag{7}$$

Furthermore, we notice that the \mathcal{L}_{rec} only regulates the pixel-level similarity between the restored image and the ground-truth, while neglecting the patchlevel linear correlations. The innate relationships of intensity within the image are disrupted by the consistent patterns of weather-induced degradation. By



Fig. 3: Visual comparison for desnowing on Snow100K [44]. The samples from (b) to (e) are Restormer [86], TransWeather [70], WGWSNet [100], WeatherDiff [53].



Fig. 4: Visual comparison for deraining and dehazing on Outdoor-Rain [32]. The samples from (b) to (e) are Restormer [86], TransWeather [70], WGWSNet [100], WeatherDiff [53].

emulating the intensity relationships within the ground-truth, we compel the degraded pixels to occupy their original positions according to the original intensity ranking. Consequently, we introduce the Pearson correlation [12] between images as a means to regulate the linear relationship, expressed as follows:

$$\rho\left(I^{hq}, I^{gt}\right) = \frac{\sum_{i=1}^{3HW} \left(I_i^{hq} - \overline{I}^{hq}\right) \left(I_i^{gt} - \overline{I}^{gt}\right)}{3HW\sigma\left(I^{hq}\right)\sigma\left(I^{gt}\right)},\tag{8}$$

where $I_i^{\{\cdot\}}$ represents the *i*-th pixel of image, $\overline{I}^{\{\cdot\}}$ and $\sigma(I^{\{\cdot\}})$ denotes respectively the mean and the standard deviation of image sequence. Its value falls within the [-1, 1] range. When two images exhibit perfect correlation, the value of function ρ attains a value of 1, while in the case of negative correlation, its value reaches -1. Hence, we formulate the correlation loss as follows:

$$\mathcal{L}_{cor} = \frac{1}{2} \left(1 - \rho \left(I^{hq}, I^{gt} \right) \right), \tag{9}$$

such that $\mathcal{L}_{cor} = 0$ when the recovered image perfectly aligns with the ground-truth. The overall loss function is thus defined as:

$$\mathcal{L} = \mathcal{L}_{rec} + \alpha \mathcal{L}_{cor},\tag{10}$$

where α is the weight of correlation loss.

4 Experiments

4.1 Experimental settings

Datasets. We train our model on the same datasets as the previous works [33,53, 70] to ensure a fair comparison. The training set encompasses 9,000 images drawn



Fig. 5: Visual comparison for raindrop removal on RainDrop [57]. The samples from (b) to (e) are Chen *et al.* [10], TransWeather [70], WGWSNet [100], WeatherDiff [53].



Fig. 6: A qualitative comparison for real-world adverse weather removal on Snow100K [44]. The samples from (b) to (e) are Chen *et al.* [10], Restormer [86], TransWeather [70], WGWSNet [100], WeatherDiff [53].

from Snow100K [44], 1,069 images sourced from Raindrop [57], and 9,000 images from Outdoor-Rain [32]. Snow100K contains synthetic images deteriorated by snow, while Raindrop comprises real raindrop-affected images. Outdoor-Rain features synthetic images afflicted by both fog and rain streaks. For evaluation, we employ the Test1 dataset [32, 33], the RainDrop test dataset [57], and the Snow100K-L and -S test sets [44]. Snow100K also provides a real-world test set containing 1,329 images affected by adverse weather.

Comparison Baselines. We assess the performance of our method against stateof-the-art approaches designed specifically for distinct weather removal tasks: raindrop removal, snow removal, and rain&fog removal. Specifically, for snow removal, we benchmark against SPANet [73], JSTASR [8], RESCAN [34], Desnow-Net [44], and DDMSNet [94]. In the case of rain&fog removal, we compare with CycleGAN [98], pix2pix [20], HRGAN [32], MPRNet [87] and Restormer [86]. For raindrop removal, we evaluate against the methods such as pix2pix [20], DuRN [42], RaindropAttn [59], AttentiveGAN [57]. Additionally, we include some recent transformer or multi-degradation restoration networks, IDT [78], NAFNet [6], MAXIM [69], and Restormer [86], in our comparative analysis. It is worth noting that all these methods are single-task networks fine-tuned for specific datasets.

Furthermore, we conduct a performance comparison with the All-in-One network [33], Chen *et al.* [10], TransWeather [70], WGWS-Net [100], WeatherDiff [53] and AWRCP [82], which are trained to handle all the aforementioned tasks using a unified model. Note that our approach is also trained to tackle all these tasks using a single model.

Training details. Our implementation is realized by PyTorch [55] and on NVIDIA Tesla V100 GPU. The network is trained for a total of 300,000 iterations, with an initial batch size of 8 and an initial patch size of 128 akin to the progressive learning pipeline [86]. We employ the AdamW optimizer [48] with an initial learning rate of $3e^{-4}$ for the first 92,000 iterations, which is gradually reduced to $1e^{-6}$ using cosine annealing schedule [47] during the remaining 208,000 iterations. The number of blocks at each stage $L_{i \in \{1,2,3,4\}}$ is set to $\{4,4,6,8\}$ and the channel size C is 36. The channel expansion factor r in DGFF is set to 2.667. The numbers of heads in self-attention at different stages are set to $\{1,2,4,8\}$ respectively. We randomly apply horizontal and vertical flips as the technique of data augmentation.

4.2 Comparisons with the state-of-the-arts

Quantitative Evaluation. In our study, we provide a comprehensive comparative analysis of metrics applied to both synthetic and real datasets, as summarized in Table 1. For a fair and well-founded comparison, we utilize recent multiple degradation removal methods such as MPRNet [87], MAXIM [69], and Restormer [86], treating them as weather-specific approaches for each benchmark. Additionally, we retrain the all-in-one adverse weather removal methods including Chen *et al.* [10] and WGWS-Net [100] using the all-weather training dataset [33, 53, 70]. This exhaustive comparison reveals that our proposed method exhibits a significant performance advantage over existing approaches across three different types of degradation.

Qualitative Evaluation. Furthermore, we conduct a visual comparison on three tasks, and the outcomes are showcased in Figure 3, 4 and 5 respectively. Figure 6 shows a case of real-world weather removal. These results highlight that our method excels in comprehensively eliminating snow degradation, including fine and large snow spots. In contrast, the recent WeatherDiff [53] method still exhibits some residual snow degradation, and its capability to restore details is not optimal. When it comes to the restoration of challenging weather conditions, our method excels in removing complex haze and rain marks, yielding visually appealing results in comparison to prior approaches.

4.3 Ablation studies

To substantiate the effectiveness of each component within Histoformer, we conduct a sequence of ablation studies on Outdoor-Rain [32]. In particular, we examine the impact of the dynamic-range convolution, the DHSA module, the number of bins in DHSA, the DGFF module, and the correlation loss.

Dynamic-range Convolution. We experiment on two settings of dynamic-range convolution, namely, sorting horizontally first and then vertically before convolution, and the reverse order. Additionally, we compared them with vanilla convolution, and the results are displayed in Table 2. The operations of regular sorting led to a performance improvement of 0.14 dB, and the order of sorting operations does not significantly affect the outcome.

Verti.+Horiz. Vanilla Conv Horiz.+Verti. PSNR 31.9432.03 32.08 SSIM 0.93770.93850.9389

Table 2: Ablation studies on the dynamic-range convolution.

Table 3: Ablation studies on the design of self-attention

SA type	PSNR	SSIM
MDTA [86] TKSA [11] w/o BHR w/o EHB	30.94 31.12 31.05 31.79	0.9278 0.9295 0.9301 0.9364
DHSA	32.08	0.9389

 Table 4: Ablation studies on the number
 of Bins in DHSA

#Bins in DHSA	PSNR	SSIM
12	30.43	0.9283
20	31.56	0.9312
28	31.94	0.9379
36	32.08	0.9389
44	44 Out of memo	

Table 5: Ablation studies on the choice Table 6: Ablation studies on the setting of feed-forward module

Feed-Forward	PSNR	SSIM
FFN [38]	31.32	0.9313
GDFN [86]	31.42	0.9347
DANB [96]	31.56	0.9351
MSFN [11]	31.78	0.9367
DGFF	32.08	0.9389

of correlation loss

	α	PSNR	SSIM
w/o \mathcal{L}_{cor}	0	31.77	0.9358
w/ \mathcal{L}_{cor}	$0.1 \\ 1 \\ 5 \\ 10$	32.01 32.08 32.03 31.96	0.9369 0.9389 0.9392 0.9375

DHSA. To evaluate the effectiveness of the proposed DHSA module, we conduct a comparison with two baselines, i.e., a multi-Dconv head transposed attention (MDTA) [86] and a top-k sparse attention (TKSA) [11]. Additionally, we explore two additional settings of DHSA by excluding either the BHR branch or the FHR branch. The quantitative analysis results are presented in Table 3.

Both MDTA and TKSA integrate rich information across channels, which may result in a loss of the exploitation of long-range information across spatial dimensions. While our histogram self-attention can capture spatially long-range information, the use of either a single BHR or a single FHR branch neglects the inter-bin or inner-bin relationships, leading to inferior results. By incorporating dynamic-range convolution and dual-branch histogram self-attention, capable of extracting long-range spatial features, our DHSA enhances performance, resulting in a PSNR improvement of 0.96 dB compared to TKSA.

Number of Bins. To assess the influence of the number of bins in DHSA, we conduct experiments with five different values: 12, 20, 28, 36, and 44. The results are presented in Table 4. It is observed that increasing the number of bins consistently improves performance. However, when the number of bins exceeds 44, it results in an out-of-memory error.

DGFF. To assess the effectiveness of the proposed DGFF module, we conduct a comparison with four baselines: (i) the vanilla feed-forward network (FN) [38], (ii) a gated-Dconv feed-forward network (GDFN) [86], (iii) a dual adaptive neural block (DANB) [96], and (iv) a mixed-scale feed-forward network (MSFN) [11].



Fig. 7: Real-world deweathering on two snowy images [44] and their downstream detection results on Google API.

The quantitative analysis results are presented in Table 5. While MSFN integrates mixed-scale information, it may still miss out on the exploitation of multirange spatial knowledge. Through the inclusion of pixel-shuffling and feature aggregation across different ranges, our DGFF further enhances performance, resulting in a PSNR gain of 0.3 dB over MSFN.

Correlation Loss. Table 6 shows the effectiveness of the correlation loss \mathcal{L}_{cor} and the influence of its weight. It is evident that \mathcal{L}_{cor} consistently improves the performance, while the specific loss weight does not have a substantial impact on the final results. We therefore keep the loss weight as 1 by default.

4.4 Real-world Application

To further demonstrate the practical applicability of our method for real-world adverse weather removal and its potential to improve downstream detection task, we provide two samples in Figure 7. As depicted, our Histoformer effectively eliminates snowflakes from the scene and assists the detector in recognizing omitted door and building.

5 Conclusion

In this research, we introduce a novel mechanism called histogram self-attention and devise a new histogram transformer named Histoformer to tackle the challenge of all-in-one weather removal. Our histogram self-attention involves segmenting spatial features into multiple bins, and allocating varying attention along the bin or frequency dimension, allowing it to selectively focus on weatherrelated features with a dynamic range. To facilitate learning both multi-range and multi-scale information, we present DGFF module and a correlation loss. Through extensive experimentation, we demonstrate the effectiveness and superiority of our approach.

15

Acknowledgement

This work has been supported in part by National Natural Science Foundation of China (No. 62322216, 62172409, 62025604, 62306308, 62311530686), in part by Shenzhen Science and Technology Program (Grant No. JCYJ20220818102012025, KQTD20221101093559018), and in part by Guangdong Provincial Key Laboratory of Information Security Technology (No. 2023B1212060026).

References

- 1. Ancuti, C.O., Ancuti, C.: Single image dehazing by multi-scale fusion. IEEE TIP (2013)
- 2. Berman, D., Avidan, S., et al.: Non-local image dehazing. In: CVPR (2016)
- 3. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV (2020)
- Chen, H., Wang, Y., Guo, T., Xu, C., Deng, Y., Liu, Z., Ma, S., Xu, C., Xu, C., Gao, W.: Pre-trained image processing transformer. In: CVPR (2021)
- 5. Chen, J., Tan, C.H., Hou, J., Chau, L.P., Li, H.: Robust video content alignment and compensation for rain removal in a cnn framework. In: CVPR (2018)
- Chen, L., Chu, X., Zhang, X., Sun, J.: Simple baselines for image restoration. In: ECCV (2022)
- Chen, S., Ye, T., Liu, Y., Chen, E., Shi, J., Zhou, J.: Snowformer: Scale-aware transformer via context interaction for single image desnowing. arXiv preprint arXiv:2208.09703 (2022)
- Chen, W.T., Fang, H.Y., Ding, J.J., Tsai, C.C., Kuo, S.Y.: Jstasr: Joint size and transparency-aware snow removal algorithm based on modified partial convolution and veiling effect removal. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.M. (eds.) ECCV (2020)
- Chen, W.T., Fang, H.Y., Hsieh, C.L., Tsai, C.C., Chen, I., Ding, J.J., Kuo, S.Y., et al.: All snow removed: Single image desnowing algorithm using hierarchical dual-tree complex wavelet representation and contradict channel loss. In: ICCV (2021)
- Chen, W.T., Huang, Z.K., Tsai, C.C., Yang, H.H., Ding, J.J., Kuo, S.Y.: Learning multiple adverse weather removal via two-stage knowledge learning and multicontrastive regularization: Toward a unified model. In: CVPR (2022)
- 11. Chen, X., Li, H., Li, M., Pan, J.: Learning a sparse transformer network for effective image deraining. In: CVPR (2023)
- Cohen, I., Huang, Y., Chen, J., Benesty, J., Benesty, J., Chen, J., Huang, Y., Cohen, I.: Pearson correlation coefficient. Noise reduction in speech processing (2009)
- Dauphin, Y.N., Fan, A., Auli, M., Grangier, D.: Language modeling with gated convolutional networks. In: ICML (2017)
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
- 15. Eigen, D., Krishnan, D., Fergus, R.: Restoring an image taken through a window covered with dirt or rain. In: ICCV (2013)

- 16 S. Sun et al.
- Fu, H., Gong, M., Wang, C., Batmanghelich, K., Tao, D.: Deep ordinal regression network for monocular depth estimation. In: CVPR (2018)
- Fu, X., Huang, J., Zeng, D., Huang, Y., Ding, X., Paisley, J.: Removing rain from single images via a deep detail network. In: CVPR (2017)
- Godard, C., Mac Aodha, O., Brostow, G.J.: Unsupervised monocular depth estimation with left-right consistency. In: CVPR (2017)
- 19. Guo, C., Yan, Q., Anwar, S., Cong, R., Ren, W., Li, C.: Image dehazing transformer with transmission-aware 3d position embedding. In: CVPR (2022)
- Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR (2017)
- Jiang, K., Wang, Z., Yi, P., Chen, C., Wang, Z., Wang, X., Jiang, J., Lin, C.W.: Rain-free and residue hand-in-hand: A progressive coupled network for real-time image deraining. IEEE TIP (2021)
- Jiang, T.X., Huang, T.Z., Zhao, X.L., Deng, L.J., Wang, Y.: A novel tensor-based video rain streaks removal approach via utilizing discriminatively intrinsic priors. In: CVPR (2017)
- Kang, L.W., Lin, C.W., Fu, Y.H.: Automatic single-image-based rain streaks removal via image decomposition. IEEE TIP (2011)
- Kim, Y., Cho, Y., Nguyen, T.T., Lee, D.: Metaweather: Few-shot weatherdegraded image restoration via degradation pattern matching. arXiv preprint arXiv:2308.14334 (2023)
- Lai, Z., Wu, J., Chen, S., Zhou, Y., Hovakimyan, N.: Residual-based language models are free boosters for biomedical imaging tasks. In: CVPRW (2024)
- Li, B., Peng, X., Wang, Z., Xu, J., Feng, D.: Aod-net: All-in-one dehazing network. In: ICCV (2017)
- 27. Li, B., Ren, W., Fu, D., Tao, D., Feng, D., Zeng, W., Wang, Z.: Benchmarking single-image dehazing and beyond. IEEE TIP (2018)
- Li, B., Liu, X., Hu, P., Wu, Z., Lv, J., Peng, X.: All-in-one image restoration for unknown corruption. In: CVPR (2022)
- Li, L., Dong, Y., Ren, W., Pan, J., Gao, C., Sang, N., Yang, M.H.: Semi-supervised image dehazing. IEEE TIP (2019)
- Li, M., Cao, X., Zhao, Q., Zhang, L., Meng, D.: Online rain/snow removal from surveillance videos. IEEE TIP (2021)
- Li, P., Yun, M., Tian, J., Tang, Y., Wang, G., Wu, C.: Stacked dense networks for single-image snow removal. Neurocomputing (2019)
- 32. Li, R., Cheong, L.F., Tan, R.T.: Heavy rain image restoration: Integrating physics model and conditional adversarial learning. In: CVPR (2019)
- 33. Li, R., Tan, R.T., Cheong, L.F.: All in one bad weather removal using architectural search. In: CVPR (2020)
- 34. Li, X., Wu, J., Lin, Z., Liu, H., Zha, H.: Recurrent squeeze-and-excitation context aggregation net for single image deraining. In: ECCV (2018)
- Li, Y., Tan, R.T., Guo, X., Lu, J., Brown, M.S.: Rain streak removal using layer priors. In: CVPR (2016)
- Li, Y., Lu, J., Chen, H., Wu, X., Chen, X.: Dilated convolutional transformer for high-quality image deraining. In: CVPRW (June 2023)
- 37. Li, Z., Guan, B., Wei, Y., Zhou, Y., Zhang, J., Xu, J.: Mapping new realities: Ground truth image creation with pix2pix image-to-image translation. arXiv preprint arXiv:2404.19265 (2024)
- Liang, J., Cao, J., Sun, G., Zhang, K., Van Gool, L., Timofte, R.: Swinir: Image restoration using swin transformer. In: ICCV (2021)

- 39. Liang, Y., Anwar, S., Liu, Y.: Drt: A lightweight single image deraining recursive transformer. In: CVPR (2022)
- 40. Liu, J., Yang, W., Yang, S., Guo, Z.: Erase or fill? deep joint recurrent rain removal and reconstruction in videos. In: CVPR (2018)
- 41. Liu, K., Jiang, Y., Choi, I., Gu, J.: Learning image-adaptive codebooks for classagnostic image restoration. arXiv preprint arXiv:2306.06513 (2023)
- 42. Liu, X., Suganuma, M., Sun, Z., Okatani, T.: Dual residual networks leveraging the potential of paired operations for image restoration. In: CVPR (2019)
- Liu, Y., Liu, H., Li, L., Wu, Z., Chen, J.: A data-centric solution to nonhomogeneous dehazing via vision transformer. In: CVPR (2023)
- 44. Liu, Y.F., Jaw, D.W., Huang, S.C., Hwang, J.N.: Desnownet: Context-aware deep network for snow removal. IEEE TIP (2018)
- 45. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: ICCV (2021)
- Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., Hu, H.: Video swin transformer. In: CVPR (2022)
- Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. arXiv preprint arXiv:1608.03983 (2016)
- Loshchilov, I., Hutter, F.: Fixing weight decay regularization in adam. arXiv preprint arXiv:1711.05101 (2017)
- 49. Luo, Y., Zhao, R., Wei, X., Chen, J., Lu, Y., Xie, S., Wang, T., Xiong, R., Lu, M., Zhang, S.: Mowe: Mixture of weather experts for multiple adverse weather removal. arXiv preprint arXiv:2303.13739 (2023)
- 50. Lyu, W., Zheng, S., Ling, H., Chen, C.: Backdoor attacks against transformers with attention enhancement. In: ICLR Workshop (2023)
- Ma, H., Zeng, D., Liu, Y.: Learning individualized treatment rules with many treatments: A supervised clustering approach using adaptive fusion. NeurIPS (2022)
- 52. Misra, D.: Mish: A self regularized non-monotonic activation function. arXiv preprint arXiv:1908.08681 (2019)
- 53. Özdenizci, O., Legenstein, R.: Restoring vision in adverse weather conditions with patch-based denoising diffusion models. IEEE TPAMI (2023)
- Park, D., Lee, B.H., Chun, S.Y.: All-in-one image restoration for unknown degradations using adaptive discriminative filters for specific degradations. In: CVPR (2023)
- 55. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, highperformance deep learning library. NeurIPS (2019)
- 56. Patil, P.W., Gupta, S., Rana, S., Venkatesh, S., Murala, S.: Multi-weather image restoration via domain translation. In: ICCV (2023)
- 57. Qian, R., Tan, R.T., Yang, W., Su, J., Liu, J.: Attentive generative adversarial network for raindrop removal from a single image. In: CVPR (2018)
- Quan, R., Yu, X., Liang, Y., Yang, Y.: Removing raindrops and rain streaks in one go. In: CVPR (2021)
- 59. Quan, Y., Deng, S., Chen, Y., Ji, H.: Deep learning for seeing through window with raindrops. In: ICCV (2019)
- Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. NeurIPS (2015)
- Ren, W., Tian, J., Han, Z., Chan, A., Tang, Y.: Video desnowing and deraining based on matrix decomposition. In: CVPR (2017)

- 18 S. Sun et al.
- Ren, W., Liu, S., Zhang, H., Pan, J., Cao, X., Yang, M.H.: Single image dehazing via multi-scale convolutional neural networks. In: ECCV (2016)
- Ren, W., Ma, L., Zhang, J., Pan, J., Cao, X., Liu, W., Yang, M.H.: Gated fusion network for single image dehazing. In: CVPR (2018)
- 64. Shao, Y., Li, L., Ren, W., Gao, C., Sang, N.: Domain adaptation for image dehazing. In: CVPR (2020)
- Song, Y., He, Z., Qian, H., Du, X.: Vision transformers for single image dehazing. IEEE TIP (2023)
- 66. Sun, S., Ren, W., Li, J., Zhang, K., Liang, M., Cao, X.: Event-aware video deraining via multi-patch progressive learning. IEEE TIP (2023)
- 67. Sun, S., Ren, W., Wang, T., Cao, X.: Rethinking image restoration for object detection. NeurIPS (2022)
- Tan, Z., Wu, Y., Liu, Q., Chu, Q., Lu, L., Ye, J., Yu, N.: Exploring the application of large-scale pre-trained models on adverse weather removal. arXiv preprint arXiv:2306.09008 (2023)
- Tu, Z., Talebi, H., Zhang, H., Yang, F., Milanfar, P., Bovik, A., Li, Y.: Maxim: Multi-axis mlp for image processing. In: CVPR (2022)
- 70. Valanarasu, J.M.J., Yasarla, R., Patel, V.M.: Transweather: Transformer-based restoration of images degraded by adverse weather conditions. In: CVPR (2022)
- Wan, Y.C., Shao, M.W., Cheng, Y.S., Liu, Y.X., Bao, Z.Y., Meng, D.Y.: Restoring images captured in arbitrary hybrid adverse weather conditions in one go. arXiv preprint arXiv:2305.09996 (2023)
- 72. Wang, T., Zhang, K., Shao, Z., Luo, W., Stenger, B., Lu, T., Kim, T.K., Liu, W., Li, H.: Gridformer: Residual dense transformer with grid structure for image restoration in adverse weather conditions. arXiv preprint arXiv:2305.17863 (2023)
- 73. Wang, T., Yang, X., Xu, K., Chen, S., Zhang, Q., Lau, R.W.: Spatial attentive single-image deraining with a high quality real rain dataset. In: CVPR (2019)
- Wang, Y., Ma, C., Liu, J.: Smartassign: Learning a smart knowledge assignment strategy for deraining and desnowing. In: CVPR (2023)
- 75. Wang, Z., Cun, X., Bao, J., Zhou, W., Liu, J., Li, H.: Uformer: A general u-shaped transformer for image restoration. In: CVPR (2022)
- 76. Wei, W., Yi, L., Xie, Q., Zhao, Q., Meng, D., Xu, Z.: Should we encode rain streaks in video as deterministic or stochastic? In: ICCV (2017)
- 77. Wu, H., Qu, Y., Lin, S., Zhou, J., Qiao, R., Zhang, Z., Xie, Y., Ma, L.: Contrastive learning for compact single image dehazing. In: CVPR (2021)
- Xiao, J., Fu, X., Liu, A., Wu, F., Zha, Z.J.: Image de-raining transformer. IEEE TPAMI (2022)
- 79. Yang, W., Tan, R.T., Feng, J., Liu, J., Guo, Z., Yan, S.: Deep joint rain detection and removal from a single image. In: CVPR (2017)
- Yasarla, R., Patel, V.M.: Uncertainty guided multi-scale residual learning-using a cycle spinning cnn for single image de-raining. In: CVPR (2019)
- Yasarla, R., Sindagi, V.A., Patel, V.M.: Syn2real transfer learning for image deraining using gaussian processes. In: CVPR (2020)
- 82. Ye, T., Chen, S., Bai, J., Shi, J., Xue, C., Jiang, J., Yin, J., Chen, E., Liu, Y.: Adverse weather removal with codebook priors. In: ICCV (2023)
- You, S., Tan, R.T., Kawakami, R., Mukaigawa, Y., Ikeuchi, K.: Adherent raindrop modeling, detection and removal in video. IEEE TPAMI (2015)
- Yu, F., Koltun, V.: Multi-scale context aggregation by dilated convolutions. arXiv preprint arXiv:1511.07122 (2016)
- 85. Yu, F., Koltun, V., Funkhouser, T.: Dilated residual networks. In: CVPR (2017)

Restoring Images in Adverse Weather Conditions via Histogram Transformer

- 86. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: Restormer: Efficient transformer for high-resolution image restoration. In: CVPR (2022)
- Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H., Shao, L.: Multi-stage progressive image restoration. In: CVPR (2021)
- Zhang, H., Patel, V.M.: Density-aware single image de-raining using a multistream dense network. In: CVPR (2018)
- Zhang, H., Sindagi, V., Patel, V.M.: Image de-raining using a conditional generative adversarial network. IEEE TCSVT (2019)
- Zhang, H., Sindagi, V., Patel, V.M.: Joint transmission map estimation and dehazing using deep networks. IEEE TCSVT (2019)
- Zhang, H., Ba, Y., Yang, E., Mehra, V., Gella, B., Suzuki, A., Pfahnl, A., Chandrappa, C.C., Wong, A., Kadambi, A.: Weatherstream: Light transport automation of single image deweathering. In: CVPR (2023)
- 92. Zhang, J., Ren, W., Zhang, S., Zhang, H., Nie, Y., Xue, Z., Cao, X.: Hierarchical density-aware dehazing network. IEEE Transactions on Cybernetics (2021)
- 93. Zhang, K., Li, D., Luo, W., Ren, W.: Dual attention-in-attention model for joint rain streak and raindrop removal. In: IEEE TIP (2021)
- 94. Zhang, K., Li, R., Yu, Y., Luo, W., Li, C.: Deep dense multi-scale network for snow removal using semantic and depth priors. IEEE TIP (2021)
- Zhang, X., Li, H., Qi, Y., Leow, W.K., Ng, T.K.: Rain removal in video by combining temporal and chromatic properties. In: ICME (2006)
- 96. Zhao, H., Gou, Y., Li, B., Peng, D., Lv, J., Peng, X.: Comprehensive and delicate: An efficient transformer for image restoration. In: CVPR (2023)
- Zheng, Y., Zhan, J., He, S., Dong, J., Du, Y.: Curricular contrastive regularization for physics-aware single image dehazing. In: CVPR (2023)
- 98. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICCV (2017)
- Zhu, L., Fu, C.W., Lischinski, D., Heng, P.A.: Joint bi-layer optimization for single-image rain streak removal. In: ICCV (2017)
- 100. Zhu, Y., Wang, T., Fu, X., Yang, X., Guo, X., Dai, J., Qiao, Y., Hu, X.: Learning weather-general and weather-specific features for image restoration under multiple adverse weather conditions. In: CVPR (2023)