

Correlation Matching Transformation Transformers for UHD Image Restoration

Cong Wang^{1*}, Jinshan Pan², Wei Wang³, Gang Fu¹,
Siyuan Liang⁴, Mengzhu Wang⁵, Xiao-Ming Wu¹, Jun Liu⁶

¹The Hong Kong Polytechnic University

²Nanjing University of Science and Technology

³Dalian University of Technology

⁴National University of Singapore

⁵Hebei University of Technology

⁶Singapore University of Technology and Design

Abstract

This paper proposes UHDformer, a general Transformer for Ultra-High-Definition (UHD) image restoration. UHDformer contains two learning spaces: (a) learning in high-resolution space and (b) learning in low-resolution space. The former learns multi-level high-resolution features and fuses low-high features and reconstructs the residual images, while the latter explores more representative features learning from the high-resolution ones to facilitate better restoration. To better improve feature representation in low-resolution space, we propose to build feature transformation from the high-resolution space to the low-resolution one. To that end, we propose two new modules: Dual-path Correlation Matching Transformation module (DualCMT) and Adaptive Channel Modulator (ACM). The DualCMT selects top C/r (r is greater or equal to 1 which controls the squeezing level) correlation channels from the max-pooling/mean-pooling high-resolution features to replace low-resolution ones in Transformers, which can effectively squeeze useless content to improve the feature representation in low-resolution space to facilitate better recovery. The ACM is exploited to adaptively modulate multi-level high-resolution features, enabling to provide more useful features to low-resolution space for better learning. Experimental results show that our UHDformer reduces about ninety-seven percent model sizes compared with most state-of-the-art methods while significantly improving performance under different training sets on 3 UHD image restoration tasks, including low-light image enhancement, image dehazing, and image deblurring. The source codes will be made available at <https://github.com/supersupercong/UHDformer>.

Introduction

In recent years, the rapid development of advanced imaging sensors and displays has greatly contributed to the progress of Ultra-High-Definition (UHD) imaging. However, UHD images captured under low-light, hazy, or high-speed movement conditions often suffer from undesirable degradation, resulting in visually low quality and hindering high-level vision tasks. This paper presents a unified framework to address the challenging UHD image restoration problem.

With the emergence of convolutional neural networks (CNNs) and Transformers, learning-based methods have

*supercong94@gmail.com.

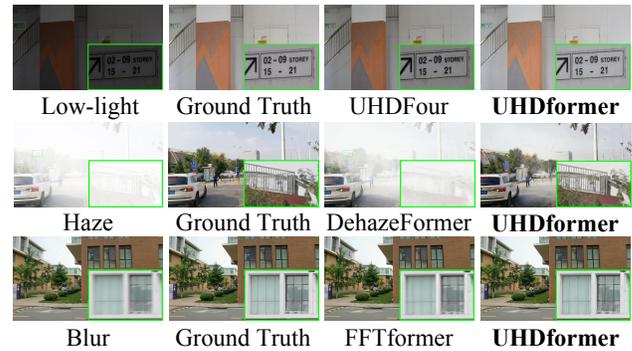


Figure 1: Challenging examples. Our UHDformer with only 0.3393M parameters outperforms SOTAs with about $50\times$ more parameters than ours (see Tabs. 2-4).

achieved impressive performance on general image restoration (Wang et al. 2020a,b; Zhu et al. 2020; Tsai et al. 2022; Jin et al. 2022, 2023; Jin, Yang, and Tan 2022; Cui et al. 2022; Wang et al. 2023). Unfortunately, these methods are typically based on general image restoration tasks, which are not capable of handling UHD image sizes or effectively recovering high-quality images. This limitation restricts the potential applications of UHD imaging systems.

Recently, with the demand for handling UHD degraded images, several approaches have been developed for UHD restoration (Zheng et al. 2021; Li et al. 2023). Zheng et al. (2021) propose a multi-guided bilateral upsampling model for UHD image dehazing. Different from the above methods which handle restoration in the spatial domain, Li et al. (2023) incorporate Fourier transform into low-light image enhancement by leveraging the amplitude and phase within a cascaded network. However, these methods usually fail to explore the valuable content for low-resolution space from the high-resolution one, which could contain useful information that significantly affects restoration quality. Furthermore, existing methods often rely on large-capacity models to achieve optimal performance. For instance, UHDFour (Li et al. 2023) and UHD (Zheng et al. 2021) methods have 17.7M and 34.5M parameters, respectively, making them unsuitable for deployment on small-capacity devices.

To solve the above problems, we propose the **UHD-**

former, a general Transformer for UHD image restoration. Our UHDformer contains two learning spaces: (a) learning in high-resolution space and (b) learning in low-resolution space. The former learns multi-level high-resolution features, fuses low-high features, and reconstructs the residual images, while the latter explores more representative features from the high-resolution space to facilitate better restoration. To learn more useful features from high-resolution space for the low-resolution one, we propose to build feature transformation from high-resolution space to low-resolution one in Transformers to improve low-resolution feature representations for better restoration. To that end, we propose two new modules: **Dual-path Correlation Matching Transformation module (DualCMT)** and **Adaptive Channel Modulator (ACM)**. Specifically, DualCMT selects the top C/r correlation channels (C denotes the number of channels; $r \geq 1$ controls the squeezing level) from the max-pooling/mean-pooling high-resolution features for low-resolution ones. These selected channels subsequently replace the low-resolution features within the Query vector of attention and forward networks within Transformers. This process effectively squeezes redundant features, improving feature representation within the low-resolution space for better recovery. To furnish the low-resolution space with more representative features, we propose the ACM that facilitates the adaptive modulation of multi-level high-resolution features by acting upon their channel dimensions. With the above designs, our UHDformer achieves superior performance (see Fig. 1) while significantly reducing the model sizes, achieving the best parameters-performance trade-off on low-light image enhancement, dehazing, and deblurring.

Our main contributions are summarized below:

- We propose **UHDformer**, the first general UHD image restoration Transformer to the best of our knowledge, by building feature transformation from a high-resolution space to a low-resolution one.
- We propose a dual-path correlation matching transformation module that can better transform the features from high- and low-resolution space to squeeze useless content, enabling it to improve feature representations in low-resolution space for better UHD image recovery.
- We propose an adaptive channel modulator to adaptively modulate multi-level high-resolution features to provide more representative content for low-resolution space.

Related Works

In this section, we review Transformers for image restoration and Ultra-High-Definition restoration approaches.

Transformers for Image Restoration. CNN-based architectures (Liu et al. 2019; Cho et al. 2021; Zamir et al. 2021; Wang et al. 2022a) have been shown to outperform conventional restoration approaches (Pan et al. 2016) due to implicitly learning the priors from large-scale data. Recently, Transformer-based models (Liang et al. 2021; Wang et al. 2022b; Zamir et al. 2022; Kong et al. 2023) have dominated image restoration due to modeling long-range pixel dependencies which overcoming the shortage of performing

computation in local windows of CNN-based algorithms. Although these Transformer-based methods have achieved promising performance for general image restoration, they usually cannot handle UHD images, which limits further potential applications on UHD imaging devices.

Ultra-High-Definition Restoration. With the demand for processing UHD images on imaging systems, a few methods have been proposed to recover clear UHD images via various network designs, including bilateral learning for image dehazing (Zheng et al. 2021), multi-scale separable-patch integration networks for video deblurring (Deng et al. 2021a), Fourier embedding network for low-light image enhancement (Li et al. 2023). However, all methods do not explore useful content in high-resolution space for the contributions of low-resolution space. In this paper, we propose to build the feature transformation from high- to low-resolution space to provide more representative features to conduct effective Transformer computation in low-resolution space.

Methodology with UHDformer

Overall Pipeline

Fig. 2 shows the overall framework of our UHDformer, which contains two learning spaces: (a) learning in high-resolution space and (b) learning in low-resolution space. The former learns multi-level high-resolution features, fuses low-high features, and reconstructs the residual images, while the latter explores more representative low-resolution features by learning to match from the high-resolution space.

Learning in High-Resolution Space. Given a UHD input image $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$, we first apply a 3×3 convolution to obtain low-level embeddings $\mathbf{X}_0 \in \mathbb{R}^{H \times W \times C}$, where $H \times W$ denotes the spatial dimension and C is the number of channels. Next, the shallow features \mathbf{X}_0 are hierarchically encoded into multi-level features $\{\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3\} \in \mathbb{R}^{H \times W \times C}$ via 3 ConvNeXt blocks (Liu et al. 2022). Then, the multi-level features are modulated via an Adaptive Channel Modulator, and the modulated feature $\mathbf{X}_{\text{acm}} \in \mathbb{R}^{H \times W \times 3C}$ is sent to low-resolution space to participate in conducting DualCMT in Transformers. Finally, a low-high fusion and reconstruction layer containing a concatenation operation and two ConvNeXt blocks, and a 3×3 convolution receives \mathbf{X}_3 and shuffled-up features learned from low-resolution space and generates residual image $\mathbf{S} \in \mathbb{R}^{H \times W \times 3}$ to which UHD input image is added to obtain the restored image: $\hat{\mathbf{H}} = \mathbf{I} + \mathbf{S}$. Fig. 2(a) shows the learning in high-resolution space.

Learning in Low-Resolution Space. The low-resolution space, as shown in Fig. 2(b), first receives the shuffled-down features $\mathbf{X}_{\text{low}}^{\text{in}} \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{8} \times C}$ from \mathbf{X}_0 of the high-resolution space. Then, $\mathbf{X}_{\text{low}}^{\text{in}}$ is sent to several Correlation Matching Transformation Transformer Blocks (**CMT-TB**), shown in Fig. 2(c), to learn the low-resolution features. Meanwhile, each CMT-TB also matches the correlation between the modulated feature \mathbf{X}_{acm} in ACM and attention/forward networks in CMT-TBs to generate more representative features to facilitate better restoration. Finally, the learned feature $\mathbf{X}_{\text{low}}^{\text{out}} \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{8} \times C}$ is sent to high-resolution space to participate in reconstructing the final recovery images.

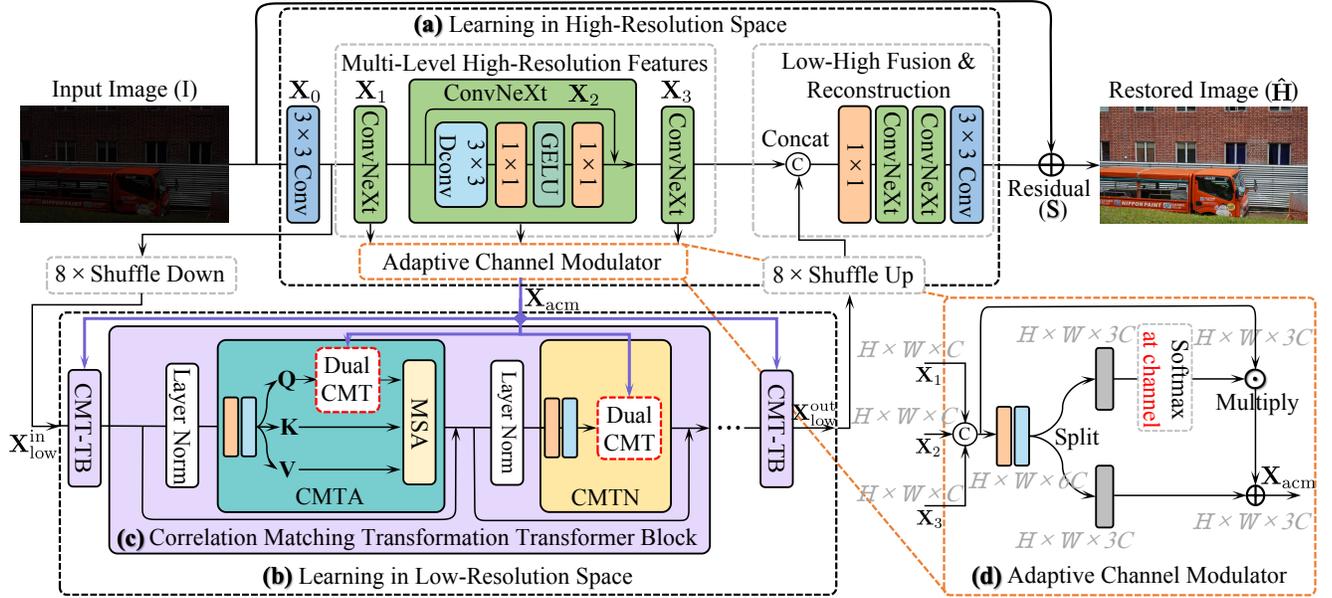


Figure 2: Overall framework of our UHDformer.

Correlation Matching Transformation Transformer

To better improve low-resolution feature representation, we propose constructing the feature transformation from the high-resolution space to the low-resolution one since we observe that multi-level high-resolution features implicitly contain better content. This transformation aims to substitute low-resolution features with more representative high-resolution counterparts by filtering out redundant high-resolution features. To achieve this objective, we introduce the Correlation Matching Transformation Transformer Block (CMT-TB), depicted in Fig. 2(c). The CMT-TB is to endow the low-resolution space with more representative features, acquired from multi-level high-resolution features. These features are intended to replace the existing low-resolution features via Dual-path Correlation Matching Transformation (DualCMT), illustrated in Fig. 3(a). Each CMT-TB contains Correlation Matching Transformation Attention (CMTA) and Correlation Matching Transformation Forward Network (CMTN) to respectively explore the correlation matching transformation in both attention and forward networks:

$$\begin{aligned} \mathbf{X}' &= \text{CMTA}\left(\text{LN}(\mathbf{X}_{\text{low}}^l), \mathbf{X}_{\text{acm}}\right) + \mathbf{X}_{\text{low}}^l, \\ \mathbf{X}_{\text{low}}^{l+1} &= \text{CMTN}\left(\text{LN}(\mathbf{X}'), \mathbf{X}_{\text{acm}}\right) + \mathbf{X}', \end{aligned} \quad (1)$$

where $\mathbf{X}_{\text{low}}^l$ means the output features of l^{th} ($l = 1, 2, \dots, L$) CMT-TB; Especially, $\mathbf{X}_{\text{low}}^0$ is the $\mathbf{X}_{\text{low}}^{\text{in}}$ and $\mathbf{X}_{\text{low}}^L$ is the $\mathbf{X}_{\text{low}}^{\text{out}}$; $\text{CMTA}(\cdot, \cdot)$ and $\text{CMTN}(\cdot, \cdot)$ respectively denote the operations of CMTA and CMTN, which are respectively defined in Eq. (2) and Eq. (3); $\text{LN}(\cdot)$ means the operation of layer normalization (Ba, Kiros, and Hinton 2016).

Correlation Matching Transformation Attention. The query typically describes its relationship with others in the attention (Ding et al. 2021). Hence, one more powerful query may significantly influence the results. In this paper, we empower the query with more representative content by replacing it with adjusted multi-level high-resolution features \mathbf{X}_{acm} . To that end, we propose the CMTA to improve the query representation to better conduct attention. Specifically, the CMTA first generates the *query* (\mathbf{Q}), *key* (\mathbf{K}), and *value* (\mathbf{V}) projections from the normalized low-resolution features $\mathbf{X}_{\text{low}}^l$ via 1×1 convolution W_p and 3×3 depth-wise convolution W_d , and then conducts the DualCMT between \mathbf{Q} and \mathbf{X}_{acm} . Next, the CMTA conducts attention (Zamir et al. 2022):

$$\begin{aligned} \text{CMTA}(\mathbf{T}_1, \mathbf{T}_2) &= \mathcal{A}\left(\text{DualCMT}(\mathbf{Q}, \mathbf{T}_2), \mathbf{K}, \mathbf{V}\right), \\ \text{where } \mathbf{Q}, \mathbf{K}, \mathbf{V} &= \text{Split}\left(W_d W_p(\mathbf{T}_1)\right), \end{aligned} \quad (2)$$

where \mathbf{T}_1 is the $\mathbf{X}_{\text{low}}^l$ in Eq. (1) and \mathbf{T}_2 is the \mathbf{X}_{acm} in Eq. (1); $\text{DualCMT}(\cdot, \cdot)$ means the DualCMT; $\mathcal{A}\left(\hat{\mathbf{Q}}, \hat{\mathbf{K}}, \hat{\mathbf{V}}\right) = \hat{\mathbf{V}} \cdot \text{Softmax}\left(\hat{\mathbf{K}} \cdot \hat{\mathbf{Q}} / \alpha\right)$; Here, α is a learnable scaling parameter to control the magnitude of the dot product of $\hat{\mathbf{K}}$ and $\hat{\mathbf{Q}}$; $\text{Split}(\cdot)$ denotes the split operation;

Correlation Matching Transformation Forward Network. Similarly to CMTA, we advocate improving feature representation in forward networks would help better restore images. To this end, we introduce CMTN, which leverages DualCMT to learn more representative features from high-resolution space, thereby improving recovery quality. Initially, CMTN generates features from the output of CMTA through layer normalization, 1×1 convolution, and 3×3

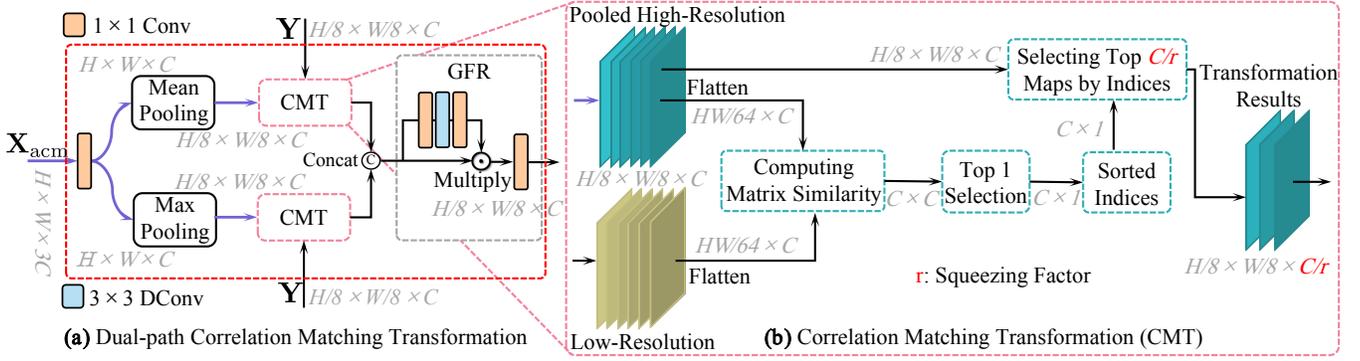


Figure 3: (a) Dual-path Correlation Matching Transformation and (b) Correlation Matching Transformation.

depth-wise convolution. It subsequently employs DualCMT between the generated features and ACM features to yield more representative forward features:

$$CMTN(\mathbf{T}_1, \mathbf{T}_2) = DualCMT(W_d W_p(\mathbf{T}_1), \mathbf{T}_2), \quad (3)$$

where \mathbf{T}_1 is the \mathbf{X}' in Eq. (1) and \mathbf{T}_2 is the \mathbf{X}_{acm} in Eq. (1).

Dual-path Correlation Matching Transformation

The DualCMT shown in Fig. 3(a) aims to transform the features from pooled high-resolution features to low-resolution ones to provide the low-resolution space with more representative features via a correlation matching scheme for better restoration. Given the learned features $\mathbf{X}_{acm} \in \mathbb{R}^{H \times W \times 3C}$ from ACM in high-resolution space, we first exploit 1×1 convolution to generate channel-reduced features, and then utilize dual-branch pooling including max-pooling and mean-pooling to reduce the spatial dimension to \mathbf{Y}_{max} and $\mathbf{Y}_{mean} \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{8} \times C}$. Then, the \mathbf{Y}_{max} and \mathbf{Y}_{mean} are matched with the low-resolution features $\mathbf{Y} \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{8} \times C}$ to squeeze useless features to select more representative features to replace the low-resolution features via CMT ($\mathcal{M}(\cdot, \cdot)$) which is computed as the process in Eq. (5):

$$\begin{aligned} \hat{\mathbf{Y}} &= W_p(\mathbf{X}_{acm}), \\ \mathbf{Y}_{max} &= Max-Pool(\hat{\mathbf{Y}}); \mathbf{Y}_{mean} = Mean-Pool(\hat{\mathbf{Y}}), \\ \mathbf{Y}_{max}^{selected} &= \mathcal{M}(\mathbf{Y}_{max}, \mathbf{Y}); \mathbf{Y}_{mean}^{selected} = \mathcal{M}(\mathbf{Y}_{mean}, \mathbf{Y}). \end{aligned} \quad (4)$$

The CMT shown in Fig. 3(b) first computes the matrix similarity ($MatSim(\cdot, \cdot)$) at the channel dimension between two tensors $\mathbf{R}_1 \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{8} \times C}$ and $\mathbf{R}_2 \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{8} \times C}$ after flattening to $\tilde{\mathbf{R}}_1 \in \mathbb{R}^{\frac{HW}{64} \times C}$ and $\tilde{\mathbf{R}}_2 \in \mathbb{R}^{\frac{HW}{64} \times C}$ to generate similarity matrix $\mathbf{M} \in \mathbb{R}^{C \times C}$. Then, we select the Top-1 ($Top_1(\cdot)$) vector $\mathbf{D} \in \mathbb{R}^{C \times 1}$ for each row tensor in \mathbf{M} and sort the values of \mathbf{D} to produce the sorted indices $\mathbf{S} \in \mathbb{R}^{C \times 1}$. Finally, we select ($Select_{C/r}(\cdot|\cdot)$) top C/r ($r \geq 1$ means the squeezing factor which controls the squeezing level) features $\mathbf{Y}^{selected} \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{8} \times \frac{C}{r}}$ from \mathbf{R}_1 by the sorted indices

S:

$$\begin{aligned} \mathbf{M} &= MatSim(\tilde{\mathbf{R}}_1, \tilde{\mathbf{R}}_2), \# \text{ matrix similarity} \\ \mathbf{D} &= Top_1(\mathbf{M}); \mathbf{S} = Sort(\mathbf{D}), \# \text{ sorted indices} \\ \mathbf{Y}^{selected} &= Select_{C/r}(\mathbf{R}_1|\mathbf{S}). \# \text{ select Top } C/r \text{ maps} \end{aligned} \quad (5)$$

Here \mathbf{R}_1 and \mathbf{R}_2 can be respectively regarded as the pooled high-resolution and low-resolution features in Fig. 3(b).

Finally, with the selected features from dual-path matching results $\mathbf{Y}_{max}^{selected} \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{8} \times \frac{C}{r}}$ and $\mathbf{Y}_{mean}^{selected} \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{8} \times \frac{C}{r}}$, we concatenate ($Concat[\cdot]$) them and then use Gated Feature Refinement ($GFR(\cdot)$) to refine the features:

$$\begin{aligned} \mathbf{Y}_{concat}^{selected} &= Concat[\mathbf{Y}_{max}^{selected}, \mathbf{Y}_{mean}^{selected}], \\ GFR(\mathbf{Y}_{concat}^{selected}) &= W_p \left(W_p W_d W_p(\mathbf{Y}_{concat}^{selected}) \odot \mathbf{Y}_{concat}^{selected} \right). \end{aligned} \quad (6)$$

Adaptive Channel Modulator

The ACM, shown in Fig. 2(d), aims to adaptively modulate the high-resolution features to better balance the importance of channel-wise features, enabling to provide more representative features for low-resolution space for better restoration. Given the multi-level high-resolution tensors $\{\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3\} \in \mathbb{R}^{H \times W \times C}$, we first concatenate them to generate a wider tensor $\mathbf{X}_{concat} \in \mathbb{R}^{H \times W \times 3C}$. Then, we use 1×1 convolution and 3×3 depth-wise convolution to expand the concatenated features \mathbf{X}_{concat} to wider features and split the features into two tensors: $\mathbf{Z}_1 \in \mathbb{R}^{H \times W \times 3C}$ and $\mathbf{Z}_2 \in \mathbb{R}^{H \times W \times 3C}$. Next, we conduct the softmax $\mathcal{S}_{channel}$ at channel dimension on \mathbf{Z}_1 to obtain channel-wise weights and finally conduct element-wise addition and element-wise multiplication among \mathbf{X}_{concat} , \mathbf{Z}_1 , and \mathbf{Z}_2 :

$$\begin{aligned} ACM(\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3) &= \mathbf{X}_{concat} \odot \mathcal{S}_{channel}(\mathbf{Z}_1) + \mathbf{Z}_2, \\ \text{where } \mathbf{X}_{concat} &= Concat[\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3], \\ \mathbf{Z}_1, \mathbf{Z}_2 &= Split(W_d W_p(\mathbf{X}_{concat})). \end{aligned} \quad (7)$$

Experiments

We evaluate **UHDformer** on benchmarks for 3 UHD image restoration tasks: (a) low-light image enhancement, (b) image dehazing, and (c) image deblurring.

Implementation Details. The number of CMT-TBs, i.e., L , is 15, where we use residual learning (He et al. 2016) to connect every 3 CMT-TB. The number of attention heads is 8, and the number of channels C is 16. The matching factor, i.e., r , is set as 4. We train models using AdamW optimizer with the initial learning rate $5e^{-4}$ gradually reduced to $1e^{-7}$ with the cosine annealing (Loshchilov and Hutter 2017). The patch size is set as 512×512 . To constrain the training of UHDformer, we use the same loss function (Kong et al. 2023) with default parameters. All experiments are conducted on two NVIDIA 3090 GPUs.

Datasets. We use UHD-LL (Li et al. 2023) to conduct UHD low-light image enhancement. For image dehazing and deblurring, we respectively re-collect the samples from the datasets of (Zheng et al. 2021) and (Deng et al. 2021b) to form new benchmarks, named as UHD-Haze and UHD-Blur. The statistics of these 3 datasets are summarised in Tab. 1. Besides using UHD images to train the models, we also use commonly-used general image restoration datasets to train and then test on UHD images. Here, we respectively use well-known LOL (Wei et al. 2018), SOTS-ITS (Li et al. 2019), and GoPro (Nah, Hyun Kim, and Mu Lee 2017) as general low-light image enhancement, dehazing, and deblurring datasets, where their training samples are used to train deep models.

Evaluation. Following (Li et al. 2023), we adopt commonly-used IQA PyTorch Toolbox¹ to compute the PSNR (Huynh-Thu and Ghanbari 2008) and SSIM (Wang et al. 2004) scores of all compared methods and also report the trainable parameters (**Param**). Since some methods cannot directly process full-resolution UHD images, we have to adopt an additional manner to conduct the experiments. According to UHDFour (Li et al. 2023), resizing (**RS**) the input to the largest size that the model can handle produces better results than splitting the input into four patches and then stitching the result. Hence, we adopt the resizing strategy for these methods and report whether models need to resize or not.

Main Results

Low-Light Image Enhancement Results. We evaluate UHD low-light image enhancement results on UHD-LL with two training dataset sets, including LOL and UHD-LL. Tab. 2 shows that our UHDformer advances state-of-the-art approaches in both these two training sets. Compared with recent state-of-the-art UHDFour (Li et al. 2023), UHDformer saves at least 98% training parameters while consistently advancing it under different training sets. We note that UHDFour trained on LOL (Wei et al. 2018) performs inferior results on UHD images. In contrast, our UHDformer still keeps excellent enhancement performance on this setting, indicating the strong robustness of our UHDformer. Moreover, UHDformer outperforms the methods which can directly handle full-resolution UHD images, e.g., Zhao et al. and URetinex (Wu et al. 2022), which further demonstrates the superiority of our UHDformer. Fig. 4 presents visual

comparisons on UHD-LL, where UHDformer is able to generate results with more natural colors.

Image Dehazing Results. Tab. 3 summarises the quantitative dehazing results on UHD-Haze with SOTS-ITS and UHD-Haze training sets. Compared to recent work DehazeFormer (Song et al. 2023) which cannot directly handle the UHD images, our UHDformer reduces 86% model sizes while achieving 0.774dB and 7.241dB PSNR gains on both SOTS-ITS and UHD-Haze training sets, respectively. We notice that although GridNet (Liu et al. 2019), UHD (Zheng et al. 2021), and MSBDN (Dong et al. 2020) do not need to resize the UHD input images, they are less effective to handle the UHD images on the SOTS-ITS training set, while our UHDformer still keeps competitive performance. Fig. 5 shows that UHDformer is capable of producing clearer results, while other methods always hand down extensive haze.

Image Deblurring Results. We evaluate UHD image deblurring with GoPro (Nah, Hyun Kim, and Mu Lee 2017) and UHD-Blur training sets. Tab. 4 summarises the results, where UHDformer significantly advances current state-of-the-art approaches on different training sets. Compared with the recent state-of-the-art deblurring approach FFTformer (Kong et al. 2023), UHDformer respectively obtains 2.811dB and 3.412dB PSNR gains on GoPro and UHD-Blur training sets. It is worth noticing that although DMPHN (Zhang et al. 2019), MIMO-Unet++ (Cho et al. 2021), and MPRNet (Zamir et al. 2021) can handle the full-resolution UHD images, they consume at least 97.8% more training parameters compared UHDformer on the GoPro training set while dropping at least 0.946dB PSNR. Fig. 6 provides several visual UHD deblurring examples, where our UHDformer is able to produce sharper results while existing state-of-the-art approaches cannot handle the UHD images well.

Ablation Study

We use the UHD-LL dataset to conduct the ablation study on the main designs of UHDformer.

Effect on Dual-path Correlation Matching Transformation. Since one of the core designs of our UHDformer is the DualCMT, it is of great interest to analyze its effect on restoration. To understand the impact of DualCMT, we disable it in the CMTA or CMTN in Transformers or disable the Max-Pooling or Mean-Pooling operations in DualCMT to compare with the full model. Tab. 5 shows that our full model with DualCMT (Tab. 5(g)) gets the 1.682dB PSNR gains compared with the model without transformation via DualCMT (Tab. 5(a)). Note that the DualCMT in both CMTA (Tab. 5(b)) and CMTN (Tab. 5(c)) plays a positive effect on image restoration, while Max-Pooling and

Dataset	Training samples	Testing samples	Resolution
UHD-LL	2,000	150	$3,840 \times 2,160$
UHD-Haze	2,290	231	$3,840 \times 2,160$
UHD-Blur	1,964	300	$3,840 \times 2,160$

Table 1: Datasets Statistics.

¹<https://github.com/chaofengc/IQA-PyTorch>

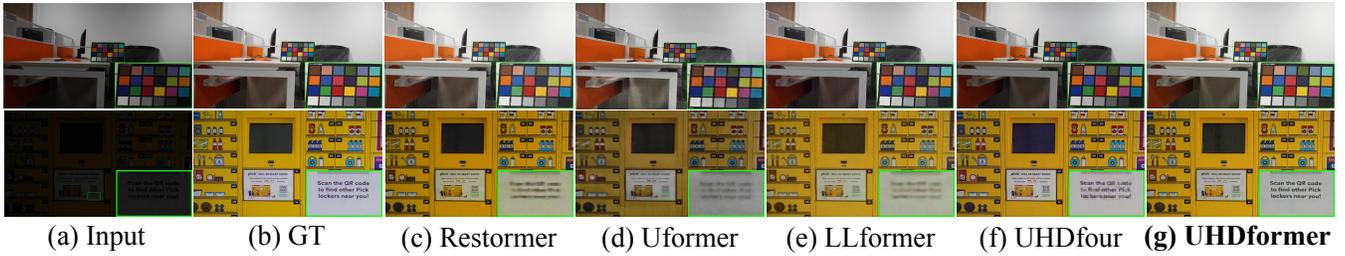


Figure 4: Low-light image enhancement on UHD-LL. UHDformer is able to generate cleaner results with finer details.

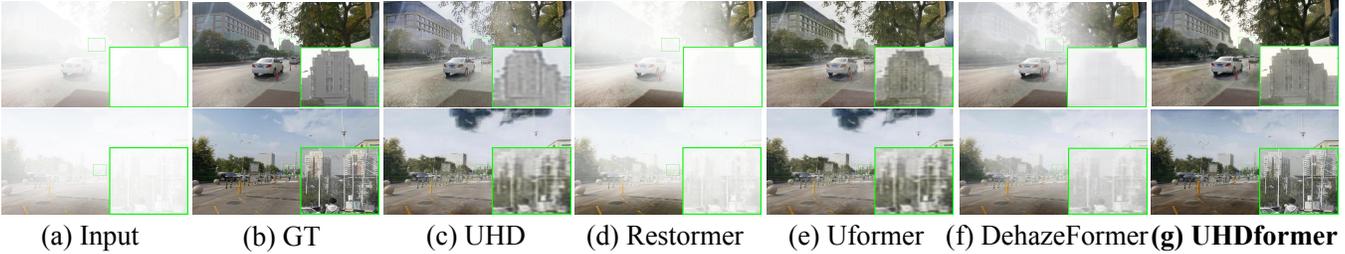


Figure 5: Image dehazing on UHD-Haze. UHDformer is able to generate much clearer dehazing results with finer structures.

Method	Venue	PSNR \uparrow	SSIM \uparrow	Param \downarrow	RS
Training Set on LOL					
SwinIR	ICCVW'21	17.900	0.7379	11.5M (-97%)	✓
Zhao et al.	ICCV'21	18.604	0.6940	11.6M (-97%)	✗
Restormer	CVPR'22	19.728	0.7703	26.1M (-99%)	✓
Uformer	CVPR'22	18.168	0.7201	20.6M (-98%)	✓
LLFlow	AAAI'22	19.596	0.7333	17.4M (-98%)	✗
LLformer	AAAI'23	21.440	0.7763	13.2M (-97%)	✓
UHDfour	ICLR'23	14.771	0.3760	17.5M (-98%)	✗
UHDformer	-	22.615	0.7754	0.3393M	✗
Training Set on UHD-LL					
Restormer	CVPR'22	21.536	0.8437	26.1M (-99%)	✓
Uformer	CVPR'22	21.303	0.8233	20.6M (-98%)	✓
LLformer	AAAI'23	24.065	0.8580	13.2M (-97%)	✓
UHDfour	ICLR'23	26.226	0.9000	17.5M (-98%)	✗
UHDformer	-	27.113	0.9271	0.3393M	✗

Table 2: Low-light image enhancement. UHDformer with at least 97% fewer parameters achieves the SOTA.

Method	Venue	PSNR \uparrow	SSIM \uparrow	Param \downarrow	RS
Training Set on SOTS-ITS					
GridNet	ICCV'19	14.783	0.8466	0.96M (-64%)	✗
MSBDN	CVPR'20	15.043	0.8570	31.4M (-99%)	✗
UHD	ICCV'21	11.708	0.6569	34.5M (-99%)	✗
Restormer	CVPR'22	13.875	0.6405	26.1M (-99%)	✓
Uformer	CVPR'22	15.264	0.6724	20.6M (-98%)	✓
D4	CVPR'22	13.656	0.8290	22.9M (-99%)	✗
DehazeFormer	TIP'23	14.551	0.6710	2.5M (-86%)	✓
UHDformer	-	15.325	0.8560	0.3393M	✗
Training Set on UHD-Haze					
UHD	ICCV'21	18.043	0.8113	34.5M (-99%)	✗
Restormer	CVPR'22	12.718	0.6930	26.1M (-99%)	✓
Uformer	CVPR'22	19.828	0.7374	20.6M (-98%)	✓
DehazeFormer	TIP'23	15.372	0.7245	2.5M (-86%)	✓
UHDformer	-	22.586	0.9427	0.3393M	✗

Table 3: Image dehazing. UHDformer with the fewest parameters significantly advances state-of-the-art methods.

Mean-Pooling are also useful for further improving recovery quality. Moreover, we also compare the model with direct max-pooling and mean-pooling transformation features from high-resolution space to low-resolution one without DualCMT (Tab. 5(f)), where our DualCMT achieves 0.91dB PSNR gains, further demonstrating the effectiveness of our DualCMT. Fig. 7 shows that the DualCMT is able to help better enhance the visual quality towards more natural colors.

The DualCMT involves exploiting the squeezing factor r

to control the matching number of features, one may wonder to know the effect of r . Tab. 6 shows that the PSNR reaches the best when r is 4, revealing that squeezing to fewer features which may reduce useless content to keep more representative features for post-learning is better than more.

Effect on Adaptive Channel Modulator. The ACM adaptively modulates the multi-level high-resolution features from the channel-wise perspective. Hence, it is necessary to analyze the impact of ACM on restoration quality by disabling the component or replacing it with other exist-

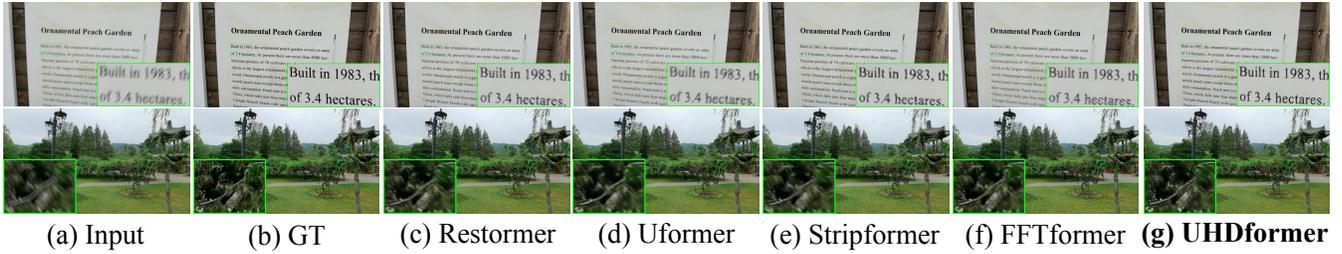


Figure 6: Image deblurring on UHD-Blur. UHDformer is able to generate deblurring results with sharper structures.

Method	Venue	PSNR \uparrow	SSIM \uparrow	Param \downarrow	RS
Training Set on GoPro					
DMPHN	CVPR'19	26.490	0.7985	21.7M (-98%)	\times
MIMO-Unet++	ICCV'21	24.290	0.7354	16.1M (-98%)	\times
MPRNet	CVPR'21	24.571	0.7426	20.1M (-98%)	\times
Restormer	CVPR'22	24.872	0.7484	26.1M (-99%)	\checkmark
Uformer	CVPR'22	24.382	0.7209	20.6M (-98%)	\checkmark
Stripformer	ECCV'22	24.915	0.7463	19.7M (-98%)	\checkmark
FFTformer	CVPR'23	24.625	0.7396	16.6M (-98%)	\checkmark
UHDformer	-	27.436	0.8231	0.3393M	\times
Training Set on UHD-Blur					
MIMO-Unet++	ICCV'21	25.025	0.7517	16.1M (-98%)	\times
Restormer	CVPR'22	25.210	0.7522	26.1M (-99%)	\checkmark
Uformer	CVPR'22	25.267	0.7515	20.6M (-98%)	\checkmark
Stripformer	ECCV'22	25.052	0.7501	19.7M (-98%)	\checkmark
FFTformer	CVPR'23	25.409	0.7571	16.6M (-98%)	\checkmark
UHDformer	-	28.821	0.8440	0.3393M	\times

Table 4: Image deblurring. UHDformer with at least 98% fewer parameters significantly advances state-of-the-arts.

Experiment	PSNR \uparrow	SSIM \uparrow
(a) w/o DualCMT in CMTA & CMTN	25.431	0.9203
(b) w/o DualCMT in CMTA	26.740	0.9264
(c) w/o DualCMT in CMTN	26.616	0.9262
(d) w/o Max-Pooling in DualCMT	25.642	0.9217
(e) w/o Mean-Pooling in DualCMT	26.291	0.9263
(f) Max- & Mean-Pooling Transformation	26.203	0.9252
(g) Full Model (Ours)	27.113	0.9271

Table 5: Ablation study on DualCMT.

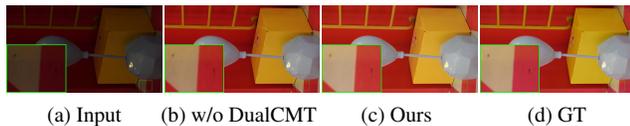


Figure 7: Visual effect on DualCMT.

ing channel attention modules. Tab. 7 shows that our ACM is more effective than well-known ECA (Qilong Wang and Hu 2020) and SE (Hu, Shen, and Sun 2018) modules. Espe-

r	1	2	3	4	5
PSNR \uparrow	26.634	26.742	26.826	27.113	26.874
SSIM \uparrow	0.9269	0.9249	0.9252	0.9271	0.9275

Table 6: Effect on squeezing factor r in DualCMT.

Experiment	PSNR \uparrow	SSIM \uparrow
(a) w/o ACM	26.485	0.9252
(b) w/ ECA (Qilong Wang and Hu 2020)	26.727	0.9240
(c) w/ SE (Hu, Shen, and Sun 2018)	26.116	0.9255
(d) w/ ACM (Ours)	27.113	0.9271

Table 7: Ablation study on ACM.

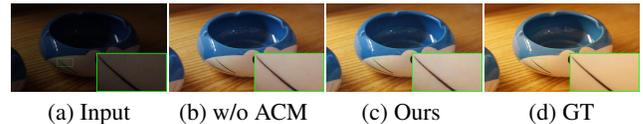


Figure 8: Visual effect on ACM.

cially, our ACM produces about 1dB PSNR gains compared with the SE channel attention, which adequately demonstrates the effectiveness of our ACM. Fig. 8 shows that our model with ACM generates more natural results (Fig. 8(c)), while the model without ACM tends to underestimate the results (Fig. 8(b)).

Concluding Remarks

We have proposed a general UHDformer to solve UHD image restoration problems. To generate more useful and representative features for low-resolution space to facilitate better restoration, we have proposed to build the feature transformation from the high-resolution space to the low-resolution one by the proposed DualCMT and ACM. Extensive experiments have demonstrated that our UHDformer significantly reduces model sizes while favoring against state-of-the-art approaches under different training sets on 3 UHD image restoration tasks, including low-light image enhancement, image dehazing, and image deblurring.

References

- Ba, J. L.; Kiros, J. R.; and Hinton, G. E. 2016. Layer normalization. *arXiv:1607.06450*.
- Cho, S.-J.; Ji, S.-W.; Hong, J.-P.; Jung, S.-W.; and Ko, S.-J. 2021. Rethinking Coarse-to-Fine Approach in Single Image Deblurring. In *ICCV*.
- Cui, X.; Wang, C.; Ren, D.; Chen, Y.; and Zhu, P. 2022. Semi-Supervised Image Deraining Using Knowledge Distillation. *IEEE TCSVT*, 32(12): 8327–8341.
- Deng, S.; Ren, W.; Yan, Y.; Wang, T.; Song, F.; and Cao, X. 2021a. Multi-Scale Separable Network for Ultra-High-Definition Video Deblurring. In *ICCV*, 14010–14019.
- Deng, S.; Ren, W.; Yan, Y.; Wang, T.; Song, F.; and Cao, X. 2021b. Multi-Scale Separable Network for Ultra-High-Definition Video Deblurring. In *ICCV*, 14030–14039.
- Ding, H.; Liu, C.; Wang, S.; and Jiang, X. 2021. Vision-Language Transformer and Query Generation for Referring Segmentation. In *ICCV*, 16301–16310.
- Dong, H.; Pan, J.; Xiang, L.; Hu, Z.; Zhang, X.; Wang, F.; and Yang, M. 2020. Multi-Scale Boosted Dehazing Network With Dense Feature Fusion. In *CVPR*, 2154–2164.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *CVPR*, 770–778.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-Excitation Networks. In *CVPR*, 7132–7141.
- Huynh-Thu, Q.; and Ghanbari, M. 2008. Scope of validity of PSNR in image/video quality assessment. *Electronics Letters*, 44(13): 800–801.
- Jin, Y.; Lin, B.; Yan, W.; Yuan, Y.; Ye, W.; and Tan, R. T. 2023. Enhancing visibility in nighttime haze images using guided apsf and gradient adaptive convolution. In *ACM MM*, 2446–2457.
- Jin, Y.; Yan, W.; Yang, W.; and Tan, R. T. 2022. Structure Representation Network and Uncertainty Feedback Learning for Dense Non-Uniform Fog Removal. In *ACCV*, 2041–2058.
- Jin, Y.; Yang, W.; and Tan, R. T. 2022. Unsupervised night image enhancement: When layer decomposition meets light-effects suppression. In *ECCV*, 404–421.
- Kong, L.; Dong, J.; Ge, J.; Li, M.; and Pan, J. 2023. Efficient Frequency Domain-Based Transformers for High-Quality Image Deblurring. In *CVPR*, 5886–5895.
- Li, B.; Ren, W.; Fu, D.; Tao, D.; Feng, D.; Zeng, W.; and Wang, Z. 2019. Benchmarking Single-Image Dehazing and Beyond. *TIP*, 28(1): 492–505.
- Li, C.; Guo, C.-L.; Zhou, M.; Liang, Z.; Zhou, S.; Feng, R.; and Loy, C. C. 2023. Embedding Fourier for Ultra-High-Definition Low-Light Image Enhancement. In *ICLR*.
- Liang, J.; Cao, J.; Sun, G.; Zhang, K.; Van Gool, L.; and Timofte, R. 2021. SwinIR: Image restoration using swin transformer. In *ICCV Workshops*.
- Liu, X.; Ma, Y.; Shi, Z.; and Chen, J. 2019. GridDehazeNet: Attention-Based Multi-Scale Network for Image Dehazing. In *ICCV*, 7313–7322.
- Liu, Z.; Mao, H.; Wu, C.-Y.; Feichtenhofer, C.; Darrell, T.; and Xie, S. 2022. A ConvNet for the 2020s. *CVPR*.
- Loshchilov, I.; and Hutter, F. 2017. SGDR: Stochastic gradient descent with warm restarts. In *ICLR*.
- Nah, S.; Hyun Kim, T.; and Mu Lee, K. 2017. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *CVPR*.
- Pan, J.; Sun, D.; Pfister, H.; and Yang, M.-H. 2016. Blind image deblurring using dark channel prior. In *CVPR*.
- Qilong Wang, P. Z. P. L. W. Z., Banggu Wu; and Hu, Q. 2020. ECA-Net: Efficient Channel Attention for Deep Convolutional Neural Networks. In *CVPR*.
- Song, Y.; He, Z.; Qian, H.; and Du, X. 2023. Vision Transformers for Single Image Dehazing. *IEEE TIP*, 32: 1927–1941.
- Tsai, F.-J.; Peng, Y.-T.; Lin, Y.-Y.; Tsai, C.-C.; and Lin, C.-W. 2022. Stripformer: Strip Transformer for Fast Image Deblurring. In *ECCV*.
- Wang, C.; Pan, J.; Wang, W.; Dong, J.; Wang, M.; Ju, Y.; and Chen, J. 2023. PromptRestorer: A Prompting Image Restoration Method with Degradation Perception. In *NeurIPS*.
- Wang, C.; Wu, Y.; Su, Z.; and Chen, J. 2020a. Joint Self-Attention and Scale-Aggregation for Self-Calibrated Deraining Network. In Chen, C. W.; Cucchiara, R.; Hua, X.; Qi, G.; Ricci, E.; Zhang, Z.; and Zimmermann, R., eds., *ACM MM*, 2517–2525. ACM.
- Wang, C.; Xing, X.; Wu, Y.; Su, Z.; and Chen, J. 2020b. DCSFN: Deep Cross-scale Fusion Network for Single Image Rain Removal. In *ACM MM*, 1643–1651. ACM.
- Wang, Y.; Wan, R.; Yang, W.; Li, H.; Chau, L.-P.; and Kot, A. C. 2022a. Low-Light Image Enhancement with Normalizing Flow. In *AAAI*.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; and Simoncelli, E. P. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 13(4): 600–612.
- Wang, Z.; Cun, X.; Bao, J.; and Liu, J. 2022b. Uformer: A General U-Shaped Transformer for Image Restoration. In *CVPR*.
- Wei, C.; Wang, W.; Yang, W.; and Liu, J. 2018. Deep Retinex Decomposition for Low-Light Enhancement. In *BMVC*, 155.
- Wu, W.; Weng, J.; Zhang, P.; Wang, X.; Yang, W.; and Jiang, J. 2022. URetinex-Net: Retinex-Based Deep Unfolding Network for Low-Light Image Enhancement. In *CVPR*, 5901–5910.
- Zamir, S. W.; Arora, A.; Khan, S.; Hayat, M.; Khan, F. S.; and Yang, M.-H. 2022. Restormer: Efficient Transformer for High-Resolution Image Restoration. In *CVPR*, 5718–5729.
- Zamir, S. W.; Arora, A.; Khan, S.; Hayat, M.; Khan, F. S.; Yang, M.-H.; and Shao, L. 2021. Multi-Stage Progressive Image Restoration. In *CVPR*.
- Zhang, H.; Dai, Y.; Li, H.; and Koniusz, P. 2019. Deep Stacked Hierarchical Multi-Patch Network for Image Deblurring. In *CVPR*.

Zhao, L.; Lu, S.; Chen, T.; Yang, Z.; and Shamir, A. 2021. Deep Symmetric Network for Underexposed Image Enhancement with Recurrent Attentional Learning. In *ICCV*, 12055–12064.

Zheng, Z.; Ren, W.; Cao, X.; Hu, X.; Wang, T.; Song, F.; and Jia, X. 2021. Ultra-High-Definition Image Dehazing via Multi-Guided Bilateral Learning. In *CVPR*, 16185–16194.

Zhu, H.; Wang, C.; Zhang, Y.; Su, Z.; and Zhao, G. 2020. Physical Model Guided Deep Image Deraining. In *ICME*.