

OneRestore: A Universal Restoration Framework for Composite Degradation

Yu Guo^{1,2,†} , Yuan Gao^{1,†} , Yuxu Lu³ , Huilin Zhu^{1,2} , Ryan Wen Liu^{1(✉)} , and Shengfeng He^{2(✉)} 

¹ Wuhan University of Technology, Wuhan, China
wenliu@whut.edu.cn

² Singapore Management University, Singapore
shengfenghe@smu.edu.sg

³ The Hong Kong Polytechnic University, Hong Kong, China
† Equal Contribution

<https://github.com/gy65896/OneRestore>

Abstract. In real-world scenarios, image impairments often manifest as composite degradations, presenting a complex interplay of elements such as low light, haze, rain, and snow. Despite this reality, existing restoration methods typically target isolated degradation types, thereby falling short in environments where multiple degrading factors coexist. To bridge this gap, our study proposes a versatile imaging model that consolidates four physical corruption paradigms to accurately represent complex, composite degradation scenarios. In this context, we propose OneRestore, a novel transformer-based framework designed for adaptive, controllable scene restoration. The proposed framework leverages a unique cross-attention mechanism, merging degraded scene descriptors with image features, allowing for nuanced restoration. Our model allows versatile input scene descriptors, ranging from manual text embeddings to automatic extractions based on visual attributes. Our methodology is further enhanced through a composite degradation restoration loss, using extra degraded images as negative samples to fortify model constraints. Comparative results on synthetic and real-world datasets demonstrate OneRestore as a superior solution, significantly advancing the state-of-the-art in addressing complex, composite degradations.

1 Introduction

Image restoration is an essential task that entails the recovery of high-quality visuals from compromised inputs. The fidelity of these restored images is vital for the subsequent downstream tasks, such as robot navigation and autonomous driving. Considerable progress has been made in addressing single degradation scenarios by a single model (*i.e.*, One-to-One models, see Fig. 2a), such as low-light conditions [16, 49, 59], haze [21, 41, 57], rain [11, 14, 31], and snow [4, 8, 38]. These methods have achieved remarkable success within their intended contexts. Nevertheless, real-world conditions often involve unpredictable and variable composite degradations that can severely impact image quality and clarity [28]. The



Fig. 1: Our OneRestore allows fully controllable image restoration using scene descriptors derived from both automatic visual attribute extraction (top) and manual text embeddings (bottom).

occurrence of mixed degrading conditions in one image can lead to intricate interactions, posing significant challenges to One-to-One solutions [22, 52, 54, 60]. The necessity for seamless switching and decision-making across specialized methods for each condition proves to be inefficient and ineffective in practical applications. Therefore, there is a growing imperative for robust, all-encompassing image restoration approaches that can universally address multiple forms of degradation through a singular, dynamic framework.

Recent research has been directed toward One-to-Many models, which address various individual degradation factors within a single framework. These models are categorized into: those with partial parameter-sharing (One-to-Many_p) [2, 19, 28, 45] and those with full parameter-sharing (One-to-Many_f) [6, 26, 35, 42, 43]. One-to-Many_p (left of Fig. 2b) restoration methods typically employ models with either multiple heads and tails or a single tail to address various types of image impairments. However, the complexity and size of these models tend to scale with the diversity of degradations they aim to correct, leading to increased training costs and parameter counts. Furthermore, when images suffer from a combination of degrading factors, these models may underperform. Some One-to-Many_p techniques also depend on user intervention to select appropriate restoration weights, which is impractical for automated or dynamic settings. On the other hand, One-to-Many_f (right of Fig. 2b) approaches train models using a direct learning strategy, where inputs are images with different individual degradations and the targets are their clear counterparts. However, this direct training approach is designed predominantly for single-category degradation issues, it does not adequately address the complex dynamics that arise when multiple types of impairments affect an image simultaneously. As a result, such methods can falter in practical applications where user control and precise directionality of restoration are crucial. For instance, in cases where users seek to improve low-light imagery, the model might inadvertently prioritize dehazing actions, deviating from user intent.

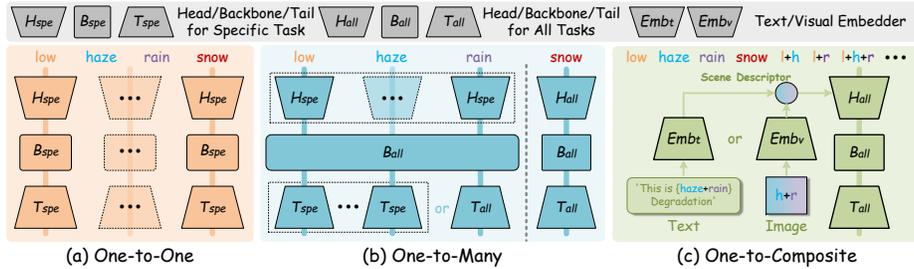


Fig. 2: Structural comparison between One-to-One, One-to-Many, and our One-to-Composite image restoration methods.

Drawing from the preceding discussion, we delineate the primary challenges in developing image restoration models for composite degradations in real-world conditions:

- How can the model be designed to adaptively discern the specific degradation afflicting the input image?
- What approaches enable a model to be trained efficiently across a spectrum of degradations while enhancing user control over the restoration process?
- How can the recovery process be optimized that not only approach the quality of clear images but also maintain distinction from various degraded forms?

In response to the challenges outlined, we present a unified model that builds upon established physical corruption models for degradations such as low light, haze, rain, and snow. This model effectively captures the complexity of real-world degradation conditions and forms the basis for constructing our Composite Degradation Dataset (CDD-11), which serves both training and evaluation purposes. To address the issue of composite degradation, our work introduces *OneRestore*, a transformer-based universal restoration architecture. This novel framework distinguishes itself from conventional methods by embodying a One-to-Composite strategy, as depicted in Fig. 2c, which allows for precise and controllable outcomes (see Fig. 1). It employs an innovative cross-attention block that fuses scene descriptors, which can be manual text embeddings or derived automatically from visual attributes. We further enhance the model’s robustness by integrating a composite degradation restoration loss, using additional degraded images as negative samples to establish stringent lower-bound constraints.

Our key contributions are summarized as follows:

- We present the first attempt to unify composite degradations, creating a versatile imaging model that simulates a range of degradation types. This innovation underpins the development of our comprehensive Composite Degradation Dataset.
- We propose a universal framework that introduces a novel level of controllability in image restoration, employing a cross-attention mechanism that integrates scene descriptors derived from either manual text embeddings or automated visual attribute extractions.

- We tailor a composite degradation restoration loss, enhancing the model’s ability to discern between various forms of image degradation.
- Our universal model not only sets a new state-of-the-art for composite dataset performance but also competes favorably with One-to-One models in single-degradation scenarios.

2 Related Work

Complex Degradation Image Restoration. Recent advancements in image restoration have introduced a variety of solutions for complexly degraded images. A notable direction is the partial parameter-sharing One-to-Many approach, employing architectures with multiple heads and tails, and a single backbone, to address various degradation factors. For instance, Li *et al.* [28] develop an all-weather network that restores scenes affected by haze, rain, and snow, using specialized encoders for each condition and a shared decoder. Chen *et al.* [2] introduce a transformer-based model for deraining, denoising, and super-resolving images, leveraging a multi-head and multi-tail design. Wang *et al.* [45] present a knowledge assignment strategy to differentiate and process rain and snow degradations adaptively. While these models represent strides in complex degradation image restoration, they stop short of being truly universal One-to-Many networks. With only partial weight sharing, they address different degradation types through distinct network pathways.

Conversely, universal image restoration has emerged with the One-to-Many methods sharing all network weights, thereby uniformly mitigating multiple degradation types. Examples include the unified two-stage knowledge learning framework by Chen *et al.* [6], which efficiently processes dehazing, deraining, and desnowing with a singular pretrained weight set. Li *et al.* [26] present a network capable of comprehensive image restoration without the need for pre-identifying degradation levels. Additionally, Özdenizci *et al.* [35] introduce a conditional denoising diffusion model tailored for patch-based image restoration under varied weather conditions. However, due to the significant differences and mutual interference between various degradation factors, this type of method can only solve specific types of degradation and cannot achieve image restoration in composite degradation scenarios. To this end, we propose a universal One-to-Composite image restoration method that allows the model to adapt to arbitrary combinations of composite degradations.

Visual Attribute Control. Visual attributes are pivotal in defining image features and enhancing the interpretability of models. They enrich feature representations and provide a hierarchical framework for image analysis. Early research highlighted visual attributes’ transformative impact on image generation, recognition, and retrieval [13, 24, 50]. More recent studies underscore their utility in improving object recognition and detection, particularly in data-scarce environments [1, 33, 48, 61]. Visual attributes have also been leveraged for zero-shot learning applications, advancing classification tasks beyond the constraints of labeled datasets [39, 44]. In our study, we harness visual attributes to delineate the intricate degradation patterns present in real-world imaging conditions,

enhancing the pertinence and efficacy of image restoration techniques. It is noteworthy that conventional discriminators, when tasked with identifying image degradation, are typically confined to image data, precluding autonomous feature control. Our approach, by contrast, accommodates both images and text descriptions as inputs, thereby facilitating the automatic or manual determination of degradation characteristics.

3 Composite Degradation Formulation

3.1 Imaging Formulation

In real-world scenarios, images are often affected by multiple degrading factors simultaneously, resulting in suboptimal performance from restoration methods designed for isolated degradations. To address the broad spectrum of degradation scenarios, we introduce a versatile imaging model that encapsulates all degradations, including low light, haze, rain, and snow. This model aims to narrow the disparity between simulated and authentic real-world data. Given a degraded image $I(x)$ and its corresponding clear state $J(x)$, the proposed imaging model is mathematically articulated as

$$I(x) = \mathcal{P}_h(\mathcal{P}_{rs}(\mathcal{P}_l(J(x)))) \tag{1}$$

where \mathcal{P}_l , \mathcal{P}_{rs} , and \mathcal{P}_h denote the low light, rain/snow, and haze degradations, respectively. Given that concurrent rain and snow conditions are infrequent, they are not addressed within the scope of this work. A comprehensive exposition of the proposed imaging model is detailed below.

Low-Light Conditions. Inspired by Retinex theory, low-light degraded image $I_l(x)$ can be generated by

$$I_l(x) = \mathcal{P}_l(J(x)) = \frac{J(x)}{L(x)}L(x)^\gamma + \varepsilon. \tag{2}$$

Here, we use a darkening coefficient $\gamma \in [2, 3]$ as the brightness controller to adjust the illumination map $L(x)$ generated via LIME [17]. To simulate poor light conditions realistically, we add the Gaussian noise ε with mean = 0 and variance $\in [0.03, 0.08]$.

Rain/Snow Streaks. Following the methodologies outlined in [5, 27], our process superimposes rain or snow streaks onto the image before synthesizing haze. To introduce rain streaks \mathcal{R} , we produce the rain-degraded image $I_{rs}(x)$ by overlaying \mathcal{P}_{rs} onto the low-light image $I_l(x)$, which is expressed as

$$I_{rs}(x) = \mathcal{P}_{rs}(I_l(x)) = I_l(x) + \mathcal{R}. \tag{3}$$

For snow streaks \mathcal{S} , the snow-degraded image is defined as

$$I_{rs}(x) = \mathcal{P}_{rs}(I_l(x)) = I_l(x)(1 - \mathcal{S}) + M(x)\mathcal{S}, \tag{4}$$

with $M(x)$ being the chromatic aberration map.

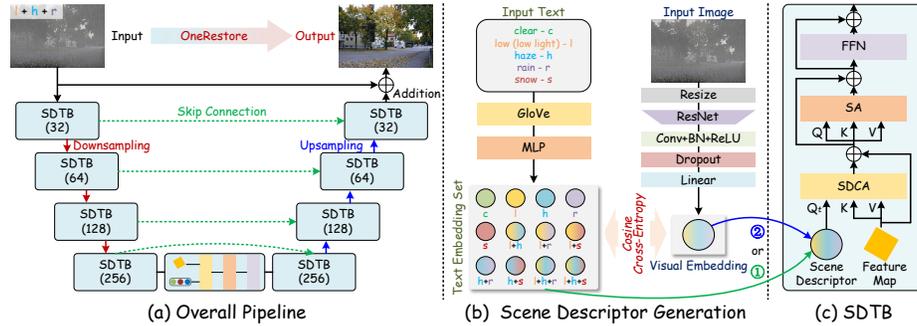


Fig. 3: Architecture of our OneRestore. (a) Overall pipeline, where 32, 64, 128, and 256 represent the number of channels. (b) Scene descriptor generation, where scene descriptors are fed into each (c) Scene Descriptor-guided Transformer Block (SDTB) by manual text embeddings or automatic extraction based on visual attributes.

Haze Degradations. We incorporate haze degradation by employing the atmospheric scattering model as

$$I(x) = \mathcal{P}_h(I_{rs}(x)) = I_{rs}(x)t + A(1 - t), \quad (5)$$

where t denotes the transmission map and A represents the atmospheric light. The transmission map t is defined by the exponential decay of light, $t = e^{-\beta d(x)}$, where β represents the haze density coefficient and $d(x)$ is the scene depth (estimated by MegaDepth [29]). The haze density coefficient β is varied within the range of [1.0, 2.0] and the atmospheric light A is constrained to [0.6, 0.9] for realistic simulation.

3.2 Composite Degradation Dataset

Drawing on Eq. (1), we have developed the Composite Degradation Dataset (CDD-11), encompassing 11 categories of image degradations and their clear counterparts. These degraded samples include *low* (low-light), *haze*, *rain*, *snow*, *low+haze*, *low+rain*, *low+snow*, *haze+rain*, *haze+snow*, *low+haze+rain*, and *low+haze+snow*. From the RAISE database [10], we selected 1,383 high-resolution clear images for producing 11 composite degradations. They are resized to a uniform resolution of 1080×720 . The overall dataset is split into 13,013 image pairs for training and 2,200 for testing.

4 OneRestore

The encoder-decoder network, validated as an effective architectural paradigm for image restoration [7, 12, 18], serves as the overall structure for our model, as depicted in Fig. 3a. Building upon the encoder-decoder structure, we devise

an innovative Scene Descriptor-guided Transformer Block (SDTB) as the basic unit for feature extraction (Fig. 3c). Our approach diverges from existing transformer-based methods [4, 7, 41] by concurrently integrating image features and scene descriptors into the SDTB basic block for supporting fine-grained image restoration. In particular, the generation of scene descriptors ranges from manual text embeddings to automatic extraction based on visual attributes, as illustrated in Fig. 3b. Further details on scene descriptor generation and SDTB are elaborated in Sections 4.1 and 4.2. For capturing multi-scale representations, the encoder employs three downsampling operations, mirrored by the decoder’s upsampling operations. To be specific, we utilize "Max-Pooling + Conv" and "Bilinear Interpolation + Conv" for downsampling and upsampling, respectively. Simultaneously, three skip connections and a global residual enhance model convergence by bridging shallow and deep features. Additional details regarding the configuration of our OneRestore can be found in **supplementary materials**.

Ultimately, the overall loss \mathcal{L} for model training is defined as follows

$$\mathcal{L} = \alpha_1 \mathcal{L}_1^s(J, \hat{J}) + \alpha_2 \mathcal{L}_M(J, \hat{J}) + \alpha_3 \mathcal{L}_c(J, \hat{J}, I, \{I_o\}), \quad (6)$$

where \mathcal{L}_1^s , \mathcal{L}_M , and \mathcal{L}_c represent the smooth l_1 loss [15], Multi-Scale Structural Similarity (MS-SSIM) loss, and the proposed composite degradation restoration loss (see Sec. 4.3), respectively, α_{1-3} denote the penalty coefficients, J , \hat{J} , I , and $\{I_o\}$ are the positive, model output anchor, model input negative, and other negatives, respectively.

4.1 Scene Descriptor Generation

To augment the controllability of our OneRestore, we introduce the scene description embedding as an additional input to the model, alongside the degraded image. As delineated in Fig. 3b, the generation of the scene descriptor offers two selectable modes: the manual mode ① and the automatic mode ②. In the manual mode, user involvement is necessitated to supply scene description text, generating the corresponding text embedding employed as an input. Conversely, the automatic mode involves the extraction of visual attributes from the image to generate visual embedding, estimating the most proximate text embedding. To quantify the difference between visual and text embeddings and train optimal text and visual embedders, we introduce cosine cross-entropy loss for model weight optimization.

Text Embedder. For the text embedding generation task, we employ a set of 5 scene description texts as input to generate 12 text embeddings. Initially, GloVe [36] is leveraged to generate initial text embeddings for 1 clear scene and 4 scenes with single degradation types. Subsequently, 7 additional composite degradation text embeddings (i.e., *low+haze*, *low+rain*, *low+snow*, *haze+rain*, *haze+snow*, *low+haze+rain*, and *low+haze+snow*) are generated by averaging the element values in the corresponding text embeddings. A Multi-Layer Perceptron (MLP) is then employed to refine the 12 text embeddings.

Visual Embedder. Taking inspiration from the work of Saini et al. [39], we first resize the input image to 224×224 and extract image features using a

ResNet-18 [20] network without average pooling. The initial weights of ResNet-18 are pre-trained on the ImageNet dataset. Subsequently, a single convolutional layer, followed by a dropout layer and a linear layer, is introduced to generate the final visual embedding. By calculating the cosine similarity between visual embeddings and all text embeddings, and selecting the one with the highest similarity, we can automatically generate scene descriptors to achieve image restoration based on visual attribute control.

Cosine Cross-Entropy. Similar to [39], the score of visual embedding and each text embedding is obtained by ‘‘Cosine Similarity + Softmax’’. Given e_v and e_t being the visual and text embeddings, the similarity score $S(e_v, e_t)$ can be given by

$$\begin{aligned} \cos(e_v, e_t) &= \delta \cdot \frac{e_v \cdot e_t^\top}{\|e_v\| \|e_t\|}, \\ S(e_v, e_t) &= \frac{e^{\cos(e_v, e_t)}}{\sum_{t_i=1}^{N_t} e^{\cos(e_v, e_{t_i})}}, \end{aligned} \quad (7)$$

with δ and N_t being the temperature factor and the number of text embeddings, respectively. Based on the calculated $S(e_v, e_t)$, we use the cross-entropy loss for model training.

4.2 Scene Descriptor-guided Transformer Block

To further harness the capabilities of scene description embeddings in image restoration tasks and facilitate a more adaptive and nuanced restoration process, we introduce the innovative Scene Descriptor-guided Transformer Block (SDTB), illustrated in Fig. 3c. The SDTB consists of three primary components, namely Scene Descriptor-guided Cross-Attention (SDCA), Self-Attention (SA), and Feed-Forward Network (FFN). Diverging from traditional self-attention mechanisms, the proposed SDCA module utilizes a query generated from scene descriptors instead of queries generated from image features. The mathematical representation of our SDCA can be written as

$$\text{SDCA}(\mathbf{Q}_t, \mathbf{K}, \mathbf{V}) = \text{Softmax}\left(\frac{\mathbf{Q}_t \cdot \mathbf{K}^\top}{\lambda}\right)\mathbf{V}, \quad (8)$$

where λ is a temperature factor, \mathbf{Q}_t represents the scene descriptor-generated query matrix, \mathbf{K} and \mathbf{V} denote the image feature-generated key and value matrices, respectively. \mathbf{Q}_t is obtained from the input scene description embedding e_t via a linear layer. During the process of matrix multiplication between \mathbf{Q}_t and \mathbf{K} . To ensure consistency in the number of tokens between \mathbf{K} and \mathbf{Q}_t , we resize the original image before performing the reshaping operation to generate \mathbf{K} . By incorporating the SDCA module, each transformer block in OneRestore can effectively integrate scene descriptors and image features for targeted image restoration. As for the SA and FFN modules, we adopt structures akin to those in Restormer [52]. More details of our SDTB are introduced in the **supplement**.

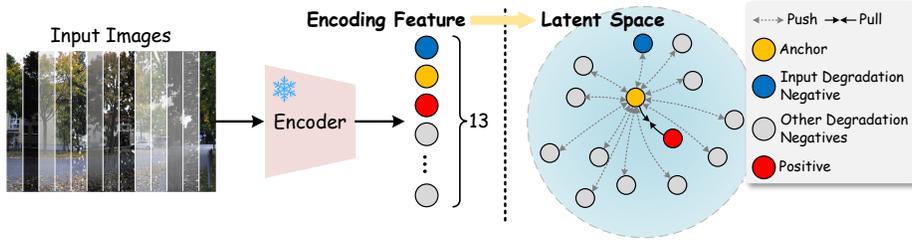


Fig. 4: Illustration of the composite degradation restoration loss.

4.3 Composite Degradation Restoration Loss

Using contrastive loss has been demonstrated as an effective way of improving image restoration performance [47, 57]. Essentially, contrastive loss views the degraded input image as a negative sample, the clear image as a positive sample, and the restored image as an anchor. This loss function works to minimize the distance between the anchor and the positive sample while maximizing its distance from the negative sample. In the composite degradation image restoration task, the traditional contrastive loss, however, may inadvertently reduce the distance between the anchor and other degraded images. To make the anchor closer to the positive sample while maintaining a clear separation from all degraded negatives, we design a composite degradation restoration loss, illustrated in Fig. 4. Considering I , I_o , J , and \hat{J} as the input negative, other negatives, positive, and anchor, the proposed loss can be written as

$$\mathcal{L}_c(J, \hat{J}, I, \{I_o\}) = \sum_{k=1}^K \xi_k \frac{\mathcal{L}_1(V_k(J), V_k(\hat{J}))}{\xi_c \mathcal{L}_1(V_k(\hat{J}), V_k(I)) + \sum_{o=1}^O \xi_o \mathcal{L}_1(V_k(I_o), V_k(\hat{J}))}, \quad (9)$$

where $V_k(\cdot)$ denotes the extraction operation of k -th hidden feature from the VGG-16 pre-trained on the ImageNet, \mathcal{L}_1 represents the l_1 loss, K and O are the numbers of used VGG-16 feature layers and other degraded negatives, ξ_k , ξ_c , and ξ_o denote the hyper-parameters, respectively. In this work, the number of other degraded negatives O is 10. We use the output features of 3-rd, 8-th, and 15-th layers from VGG-16 and set $K = 3$. Furthermore, we consider that the features of each feature layer of each degraded negative have equal importance. Therefore, we set $\xi_c = \xi_o = \frac{1}{11}$ and $\xi_k = \frac{1}{3}$. According to our ablation study, the proposed loss confers a significant advantage over using a single input degradation as a negative sample, which can easily capture positive sample features.

Table 1: Settings of model training. “lr” denotes the learning rate.

Terms	Text/Visual Embedder	OneRestore
Epoch	200	120
Initial lr	0.0001	0.0002
lr Decay	0.5/50 epochs	0.5/20 epochs
Batch Size	256	4

Table 2: Comparison of quantitative results on CDD-11 dataset. OneRestore[†] means using the corresponding scene description text as additional input, and WGWSNet requires the scene type as model input. **Red**, **green**, and **blue** indicate the best, second-best, and third-best results, respectively.

Types	Methods	Venue & Year	PSNR \uparrow	SSIM \uparrow	#Params
	Input		16.00	0.6008	-
One-to-One	MIRNet [53]	ECCV2020	25.97	0.8474	31.79M
	MPRNet [54]	CVPR2021	25.47	0.8555	15.74M
	MIRNetv2 [55]	TPAMI2022	25.37	0.8335	5.86M
	Restormer [52]	CVPR2022	26.99	0.8646	26.13M
	DGUNet [34]	CVPR2022	26.92	0.8559	17.33M
	NAFNet [3]	ECCV2022	24.13	0.7964	17.11M
	SRUDC [40]	ICCV2023	27.64	0.8600	6.80M
	Fourmer [58]	ICML2023	23.44	0.7885	0.55M
	OKNet [9]	AAAI2024	26.33	0.8605	4.72M
One-to-Many	AirNet [26]	CVPR2022	23.75	0.8140	8.93M
	TransWeather [42]	CVPR2022	23.13	0.7810	21.90M
	WeatherDiff [35]	TPAMI2023	22.49	0.7985	82.96M
	PromptIR [37]	NIPS2023	25.90	0.8499	38.45M
	WGWSNet [62]	CVPR2023	26.96	0.8626	25.76M
One-to-Composite	OneRestore		28.47	0.8784	5.98M
	OneRestore [†]		28.72	0.8821	5.98M

5 Experiments

5.1 Experiment Settings

Implementation Details. Our OneRestore is implemented by PyTorch 1.12.0 and trained on a PC with 2 AMD EPYC 7543 32-Core Processors and 8 NVIDIA L40 GPUs. We use the Adam with exponential decay rates being $\beta_1 = 0.9$ and $\beta_2 = 0.999$ for optimization. The training parameter settings for text/visual embedder and OneRestore are shown in Table 1. In particular, all images used for training are cropped into 256×256 image patches with a sampling stride of 200 and randomly flipped by 0, 90, 180, and 270 degrees to generate 312k patch pairs. More details about training and inference can be found in the **supplement**.

Evaluation Datasets. To facilitate a thorough comparison, we execute image restoration experiments across various datasets. These include the Composite Degradation Dataset (CDD-11) crafted by us, established One-to-One benchmarks, and authentic real-world datasets. The standard One-to-One benchmarks are designated for four principal tasks: the LOw-Light dataset (LOL) [46], the REalistic Single Image DEhazing Outdoor Training Set (RESIDE-OTS) [25], the Rain1200 dataset [56], and the Snow100k dataset [30]. Owing to the dearth of specialized datasets for real-world composite degradation, we have curated a selection from the LIME low-light enhancement dataset [17], the RESIDE Real-world Task-driven Testing Set (RESIDE-RTTS) [25], Yang’s rainy dataset [51], and the realistically snowy Snow100k dataset [25].

Competitors and Evaluation Metrics. We compare our OneRestore with 9 One-to-One image restoration methods (MIRNet [53], MPRNet [54], MIRNetv2 [55], Restormer [52], DGUNet [34], NAFNet [3], SRUDC [40], Fourmer [58], and OKNet [9]) and 5 One-to-Many image restoration methods (AirNet [26],

TransWeather [42], WeatherDiff [35], PromptIR [37], and WGWSNet [62]). In particular, all compared state-of-the-art methods are retrained on our CDD-11 train set. For PromptIR and WGWSNet, we adjust the number of prompts and the weather-specific parameters of each layer for corresponding to our task, respectively. Furthermore, we adopt the Peak Signal-to-Noise Ratio (PSNR) and Structure Similarity Index Measure (SSIM) as evaluation metrics.

5.2 Comparison with SOTAs

The quantitative evaluation results on our CDD-11 test dataset are displayed in Table 2. OneRestore[†] with text embedder achieves the highest performance, followed by OneRestore with visual embedder as the second-best. Both OneRestore[†] and OneRestore demonstrate substantial leads over the third-ranked methods in terms of both PSNR and SSIM. Further detailed restoration performance comparison of our method and the One-to-Many methods in each degradation scenario is illustrated in Fig. 5 using radar charts. Notably, image restoration for composite degradation scenarios proves challenging. By incorporating scene description embeddings into image restoration models and utilizing the proposed composite degradation restoration loss, our method robustly reconstructs clear images across various situations. For compelling evidence, Fig. 6 provides the visual comparison of image restoration in two composite degradation samples. It is evident that SRUDC, OKNet, AirNet, and WGWSNet are unable to suppress all degenerative interferences, while Restormer and PromptIR are prone to color distortion and noise residual. In contrast, the proposed OneRestore can produce more natural results and fully preserve image textures.

Moreover, Fig. 7 displays the image restoration results in two real-world scenarios. Compared with existing One-to-Many image restoration methods [26, 35, 37], the proposed OneRestore can simultaneously focus on multiple degradation factors and robustly reconstruct scene details. More results on CDD-11, real-world scenarios, and other One-to-One benchmarks of low-light, hazy, rainy, and snowy scenes are included in **supplementary materials**.

5.3 Ablation Study

Effectiveness of Network Modules. Table 3 reports the quantitative evaluation of different module configurations. While the Feed-Forward Network (FFN) with convolution as the basic unit serves as a crucial module in the transformer network, it alone fails to generate satisfactory reconstruction results. The ordinary Self-Attention module (SA) significantly strengthens network performance,

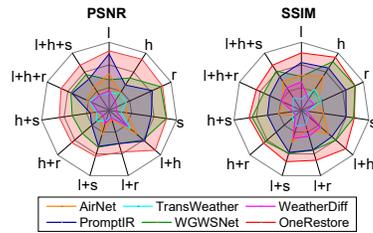


Fig. 5: Comparison of quantitative results for different degradation scenarios on CDD-11 dataset.

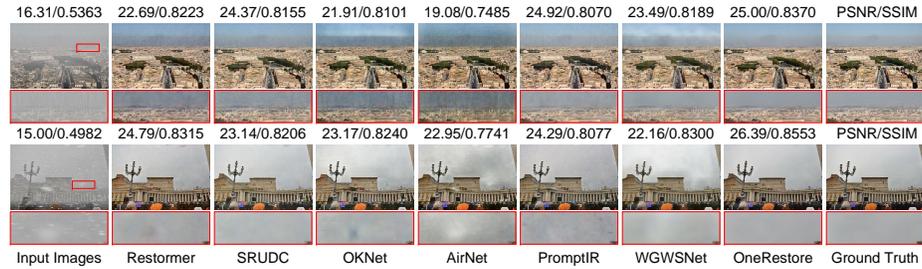


Fig. 6: Comparison of image restoration on *low+haze+rain* (top) and *low+haze+snow* (bottom) synthetic samples.

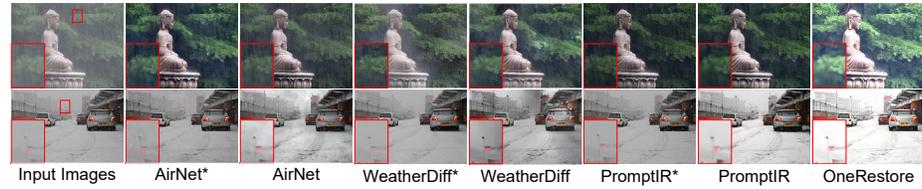


Fig. 7: Comparison of image restoration on *low+haze+rain* (top) and *low+haze+snow* (bottom) samples in real-world scenarios. * represents the utilization of original weights published in the author’s code.

but this method regrets that the model fails to be flexibly controlled. Compared to depending only on SA, relying exclusively on Scene Description-guided Cross-Attention (SDCA) utilizing scene description embedding yields better outcomes. This further confirms that addressing complex composite degradations can be done more effectively and flexibly by integrating degradation scene descriptors into the model. Ultimately, the best performance can be achieved by using a combination of SDCA, SA, and FFN.

Selection of Description Embedding Strategies. In this section, we compare the image restoration performance by using different strategies to generate scene descriptor, including Classifier (utilizing the embedding before the last linear layer of the pretrained AlexNet [23] as the scene descriptor), Visual Embedder (employing visual embedding for model training), Visual Embedder[‡] (deriving text embedding via visual embedder), and Text Embedder (employing text embedding for model training). The quantitative evaluation results are presented in Table 4. It is observed that employing Classifier and Visual Embedder to generate scene descriptors

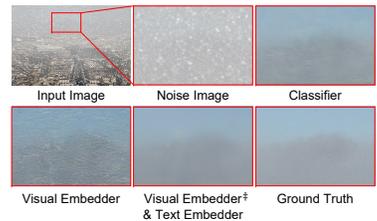


Fig. 8: Comparison of image restoration on different description embedding strategies.

Table 3: Ablation study for different model configurations. SDCA, SA, and FFN denote the scene description-guided cross-attention, self-attention, and feed-forward network, respectively.

SDCA	SA	FFN	PSNR \uparrow	SSIM \uparrow	Controllability
		✓	24.81	0.8607	
	✓	✓	27.19	0.8697	
✓		✓	27.93	0.8767	✓
✓	✓	✓	28.72	0.8821	✓



Fig. 9: Comparison of image restoration on *low+haze+rain* (top) and *low+haze+snow* (bottom) synthetic samples by using different texts.

yields comparable outcomes. However, these approaches have a drawback in that each image has a distinct embedding as input, making it challenging for the model to recognize the cause of degradation. Fig. 8 displays a classic case comparison using different embedding strategies, where the abnormal artifacts of the results produced by Classifier and Visual Embedder are caused by the model’s fuzzy identification of degradation factors. Moreover, the restoration model lacks controllability. The discrepancy between Visual Embedder[‡] and Text Embedder stems from errors in scene descriptor estimation. Visual Embedder[‡] has an accuracy of 97.55% on the test dataset, with misestimation occurring when some degradation types are insignificant. For instance, a *haze+snow* scene may be evaluated as a *haze* or *snow* scene. In contrast, the text embedding-based approach is superior in describing degradation scenes and recovering images, using a fixed number of scene descriptors. This allows the model to employ degraded scene descriptors as switches to achieve controlled restoration.

Effectiveness of Loss Functions. Table 5 presents a performance comparison of different loss functions. Notably, the integration of the MS-SSIM loss is instrumental in elevating SSIM scores at the cost of reduced PSNR performance. In contrast, the CL loss demonstrates minimal impact, providing only marginal improvements in SSIM and PSNR. Upon replacing CL with CDRL, a substantial enhancement is observed in both PSNR and SSIM. The proposed CDRL loss yields an output that closely resembles the clear image while avoiding other forms of degradation.

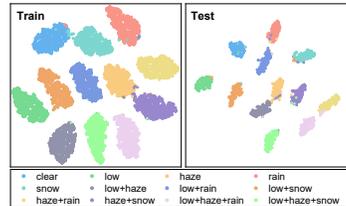


Fig. 10: Visualization of visual embeddings based on t-SNE.

Table 4: Ablation study for description embedding strategies.

Models	PSNR \uparrow	SSIM \uparrow	Controllability
Classifier	28.19	0.8783	
Visual Embedder	28.24	0.8781	
Visual Embedder [‡]	28.47	0.8784	
Text Embedder	28.72	0.8821	✓

Table 5: Ablation study for different loss functions. CL and CDRL denote the ordinary contrastive loss and the proposed composite degradation restoration loss, respectively.

Smooth l_1	MS-SSIM	CL	CDRL	PSNR \uparrow	SSIM \uparrow
✓				28.16	0.8633
✓	✓			27.54	0.8708
✓	✓	✓		27.61	0.8723
✓	✓		✓	28.72	0.8821

5.4 Model Control via Scene Descriptor

To demonstrate the controllability of the scene descriptor in our OneRestore, we conduct experiments on two intricate synthetic degradation scenarios, as depicted in Fig. 9. By manually adjusting specific text to generate corresponding degraded scene description embeddings, our model can selectively focus on different degradation factors, achieving targeted restorations. Our approach excels in controllable restoration without the need for multiple sets of weight screening, making it more concise compared to other One-to-Many methods that rely on partial parameter sharing. Meanwhile, our method can support adaptive restoration through the estimation of text embeddings based on visual attribute extraction. Fig. 10 employs the t-SNE [32] to visualize the embeddings generated by the visual embedder on the CDD-11 train and test sets. Obviously, the constructed visual embedder can clearly distinguish between different scenarios.

6 Conclusion

Our proposed OneRestore framework represents a significant step forward in image restoration, synthesizing a range of degradation patterns to accurately simulate complex scenarios. The innovative use of a scene descriptor-guided cross-attention block within a transformer-based model facilitates an adaptable and fine-grained restoration process, yielding results that surpass prior methods in both synthetic and real-world datasets. However, despite its advancements, OneRestore exhibits limitations in processing extremely high-density corruption scenarios and high-complex corruption scenarios containing unconsidered degradations, where the model’s predictive capabilities can be challenged. Future work will aim to enhance the robustness of OneRestore against such extreme conditions and to reduce computational overhead, further extending its applicability and efficiency. Our findings underscore the potential of universal scene restoration, with the caveat that continued refinement is essential for tackling the full spectrum of real-world image degradation.

Acknowledgement. This project is supported by the Guangdong Natural Science Funds for Distinguished Young Scholar under Grant 2023B1515020097 and National Research Foundation Singapore under the AI Singapore Programme under Grant AISG3-GV-2023-011.

References

1. Bertasius, G., Torresani, L.: Cobe: Contextualized object embeddings from narrated instructional video. In: NeurIPS. pp. 15133–15145 (2020)
2. 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. pp. 12299–12310 (2021)
3. Chen, L., Chu, X., Zhang, X., Sun, J.: Simple baselines for image restoration. In: ECCV. pp. 17–33 (2022)
4. Chen, S., Ye, T., Liu, Y., Liao, T., Jiang, J., Chen, E., Chen, P.: Msp-former: Multi-scale projection transformer for single image desnowing. In: ICASSP. pp. 1–5 (2023)
5. 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: CVPR. pp. 4196–4205 (2021)
6. 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 multi-contrastive regularization: Toward a unified model. In: CVPR. pp. 17653–17662 (2022)
7. Chen, X., Li, H., Li, M., Pan, J.: Learning a sparse transformer network for effective image deraining. In: CVPR. pp. 5896–5905 (2023)
8. Chen, Z., Sun, Y., Bi, X., Yue, J.: Lightweight image de-snowing: A better trade-off between network capacity and performance. *Neural Networks* **165**, 896–908 (2023)
9. Cui, Y., Ren, W., Knoll, A.: Omni-kernel network for image restoration. In: AAAI. vol. 38, pp. 1426–1434 (2024)
10. Dang-Nguyen, D.T., Pasquini, C., Conotter, V., Boato, G.: Raise: A raw images dataset for digital image forensics. In: MMSys. pp. 219–224 (2015)
11. Du, Y., Deng, J., Zheng, Y., Dong, J., He, S.: Dsdnet: Toward single image deraining with self-paced curricular dual stimulations. *CVIU* **230**, 103657 (2023)
12. Du, Y., Han, G., Tan, Y., Xiao, C., He, S.: Blind image denoising via dynamic dual learning. *IEEE TMM* **23**, 2139–2152 (2020)
13. Ferrari, V., Zisserman, A.: Learning visual attributes. In: NeurIPS. pp. 433–440 (2007)
14. Fu, X., Xiao, J., Zhu, Y., Liu, A., Wu, F., Zha, Z.J.: Continual image deraining with hypergraph convolutional networks. *IEEE TPAMI* **45**(8), 9534–9551 (2023)
15. Girshick, R.: Fast r-cnn. In: ICCV. pp. 1440–1448 (2015)
16. Guo, X., Hu, Q.: Low-light image enhancement via breaking down the darkness. *IJCV* **131**(1), 48–66 (2023)
17. Guo, X., Li, Y., Ling, H.: Lime: Low-light image enhancement via illumination map estimation. *IEEE TIP* **26**(2), 982–993 (2016)
18. Guo, Y., Gao, Y., Liu, W., Lu, Y., Qu, J., He, S., Ren, W.: Scanet: Self-paced semi-curricular attention network for non-homogeneous image dehazing. In: CVPRW. pp. 1884–1893 (2023)

19. Han, J., Li, W., Fang, P., Sun, C., Hong, J., Armin, M.A., Petersson, L., Li, H.: Blind image decomposition. In: ECCV. pp. 218–237 (2022)
20. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
21. Hoang, T., Zhang, H., Yazdani, A., Monga, V.: Transer: Hybrid model and ensemble-based sequential learning for non-homogenous dehazing. In: CVPR. pp. 1670–1679 (2023)
22. Jiao, J., Tu, W.C., Liu, D., He, S., Lau, R.W.H., Huang, T.S.: Formnet: Formatted learning for image restoration. *IEEE TIP* **29**, 6302–6314 (2020)
23. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. *CACM* **60**(6), 84–90 (2017)
24. Kumar, N., Berg, A., Belhumeur, P.N., Nayar, S.: Describable visual attributes for face verification and image search. *IEEE TPAMI* **33**(10), 1962–1977 (2011)
25. Li, B., Ren, W., Fu, D., Tao, D., Feng, D., Zeng, W., Wang, Z.: Benchmarking single-image dehazing and beyond. *IEEE TIP* **28**(1), 492–505 (2018)
26. Li, B., Liu, X., Hu, P., Wu, Z., Lv, J., Peng, X.: All-in-one image restoration for unknown corruption. In: CVPR. pp. 17452–17462 (2022)
27. Li, R., Cheong, L.F., Tan, R.T.: Heavy rain image restoration: Integrating physics model and conditional adversarial learning. In: CVPR. pp. 1633–1642 (2019)
28. Li, R., Tan, R.T., Cheong, L.F.: All in one bad weather removal using architectural search. In: CVPR. pp. 3175–3185 (2020)
29. Li, Z., Snavely, N.: Megadepth: Learning single-view depth prediction from internet photos. In: CVPR. pp. 2041–2050 (2018)
30. Liu, Y.F., Jaw, D.W., Huang, S.C., Hwang, J.N.: Desnownet: Context-aware deep network for snow removal. *IEEE TIP* **27**(6), 3064–3073 (2018)
31. Luo, Y., Huang, Q., Ling, J., Lin, K., Zhou, T.: Local and global knowledge distillation with direction-enhanced contrastive learning for single-image deraining. *Knowledge-Based Syst.* **268**, 110480 (2023)
32. Van der Maaten, L., Hinton, G.: Visualizing data using t-sne. *JMLR* **9**(11), 2579–2605 (2008)
33. Meng, Q., Shin’ichi, S.: Adinet: Attribute driven incremental network for retinal image classification. In: CVPR. pp. 4033–4042 (2020)
34. Mou, C., Wang, Q., Zhang, J.: Deep generalized unfolding networks for image restoration. In: CVPR. pp. 17399–17410 (2022)
35. Özdenizci, O., Legenstein, R.: Restoring vision in adverse weather conditions with patch-based denoising diffusion models. *IEEE TPAMI* **45**(8), 10346–10357 (2023)
36. Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. In: EMNLP. pp. 1532–1543 (2014)
37. Potlapalli, V., Zamir, S.W., Khan, S.H., Shahbaz Khan, F.: PromptIR: Prompting for all-in-one image restoration. *NeurIPS* **36** (2024)
38. Quan, Y., Tan, X., Huang, Y., Xu, Y., Ji, H.: Image desnowing via deep invertible separation. *IEEE TCSVT* **33**(7), 3133–3144 (2023)
39. Saini, N., Pham, K., Shrivastava, A.: Disentangling visual embeddings for attributes and objects. In: CVPR. pp. 13658–13667 (2022)
40. Song, B., Chen, X., Xu, S., Zhou, J.: Under-display camera image restoration with scattering effect. In: ICCV. pp. 12580–12589 (2023)
41. Song, Y., He, Z., Qian, H., Du, X.: Vision transformers for single image dehazing. *IEEE TIP* **32**, 1927–1941 (2023)
42. Valanarasu, J.M.J., Yasarla, R., Patel, V.M.: Transweather: Transformer-based restoration of images degraded by adverse weather conditions. In: CVPR. pp. 2353–2363 (2022)

43. Wang, C., Zheng, Z., Quan, R., Sun, Y., Yang, Y.: Context-aware pretraining for efficient blind image decomposition. In: CVPR. pp. 18186–18195 (2023)
44. Wang, Q., Liu, L., Jing, C., Chen, H., Liang, G., Wang, P., Shen, C.: Learning conditional attributes for compositional zero-shot learning. In: CVPR. pp. 11197–11206 (2023)
45. Wang, Y., Ma, C., Liu, J.: Smartassign: Learning a smart knowledge assignment strategy for deraining and desnowing. In: CVPR. pp. 3677–3686 (2023)
46. Wei, C., Wang, W., Yang, W., Liu, J.: Deep retinex decomposition for low-light enhancement. In: BMVC. pp. 1–12 (2018)
47. 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. pp. 10551–10560 (2021)
48. Xu, W., Xian, Y., Wang, J., Schiele, B., Akata, Z.: Attribute prototype network for zero-shot learning. In: NeurIPS. pp. 21969–21980 (2020)
49. Xu, X., Wang, R., Lu, J.: Low-light image enhancement via structure modeling and guidance. In: CVPR. pp. 9893–9903 (2023)
50. Yan, X., Yang, J., Sohn, K., Lee, H.: Attribute2image: Conditional image generation from visual attributes. In: ECCV. pp. 776–791 (2016)
51. 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. pp. 1357–1366 (2017)
52. 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. pp. 5728–5739 (2022)
53. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H., Shao, L.: Learning enriched features for real image restoration and enhancement. In: ECCV. pp. 492–511 (2020)
54. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H., Shao, L.: Multi-stage progressive image restoration. In: CVPR. pp. 14821–14831 (2021)
55. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H., Shao, L.: Learning enriched features for fast image restoration and enhancement. *IEEE TPAMI* **45**(2), 1934–1948 (2022)
56. Zhang, H., Patel, V.M.: Density-aware single image de-raining using a multi-stream dense network. In: CVPR. pp. 695–704 (2018)
57. Zheng, Y., Zhan, J., He, S., Dong, J., Du, Y.: Curricular contrastive regularization for physics-aware single image dehazing. In: CVPR. pp. 5785–5794 (2023)
58. Zhou, M., Huang, J., Guo, C.L., Li, C.: Fourmer: An efficient global modeling paradigm for image restoration. In: ICML. pp. 42589–42601 (2023)
59. Zhou, M., Wu, X., Wei, X., Xiang, T., Fang, B., Kwong, S.: Low-light enhancement method based on a retinex model for structure preservation. *IEEE TMM* pp. 1–13 (2023)
60. Zhou, Y., Xu, Y., Du, Y., Wen, Q., He, S.: Pro-pulse: Learning progressive encoders of latent semantics in gans for photo upsampling. *IEEE TIP* **31**, 1230–1242 (2022)
61. Zhu, Y., Min, W., Jiang, S.: Attribute-guided feature learning for few-shot image recognition. *IEEE TMM* **23**, 1200–1209 (2020)
62. 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. pp. 21747–21758 (2023)