CURE: Controllable Unified Image Restoration for Complex Degradations

Boseong Kim and Donghyeon Cho*

Department of Computer Science, Hanyang University, Seoul, South Korea
{hosky789, doncho}@hanyang.ac.kr

Abstract. The presence of composite degradations poses a significant challenge, since the underlying corruption factors exhibit complex and interdependent interactions. Even when the degradation types are known, accurately restoring the image remains difficult due to the intertwined nature of their effects and the need for selective control during the recovery process. To address this, we introduce CURE, a unified framework that enables controllable restoration in complex degradation settings by learning disentangled and adjustable representations. CURE is driven by four complementary objectives. First, an identity embedding is incorporated, along with a reconstruction constraint, to ensure that the model can reproduce the input image when restoration is unnecessary. Second, the ratio control mechanism blends the identity embedding with degradation-specific embeddings using user-regulated mixing ratios, allowing continuous control over restoration intensity. Third, an intermediate loss is applied to supervise stepwise outputs, each encouraged to tackle the removal of only a single degradation factor within a composite mixture. Finally, a permutation-invariant loss ensures that the model achieves consistent restoration quality regardless of the order in which multiple degradations are addressed. Since CURE modifies only the training strategy and not the underlying network architecture, it can be seamlessly integrated into existing controllable restoration models. Experiments demonstrate that CURE delivers state-of-the-art performance on composite degradation benchmarks, while enabling both selective and jointly fused restoration through flexible modulation of embedding ratios.

Keywords: All-in-one Image Restoration · Controllable Image Restoration · Composite Degradation

1 Introduction

Image restoration is a core task in low-level vision that aims to recover high-quality images from inputs affected by various forms of degradation. This capability is important in real-world applications such as autonomous driving [17] and robotics [19], where obtaining clean and stable visual signals is often difficult. The problem remains inherently ill-posed because multiple plausible clean images

* Corresponding author.

can explain the same degraded content, which makes it difficult for models to generalize across diverse degradation patterns. Earlier approaches [12,21,34] rely on hand-crafted priors designed for specific degradation scenarios. With the introduction of transformer architectures [24], attention-based models such as Uformer [26] and Restormer [29] show strong performance improvements. However, these models are typically optimized for a single degradation type, requiring separate networks for different corruptions and limiting their flexibility when dealing with varied or mixed degradation conditions.

To address the limitations of single-task models, the all-in-one restoration paradigm aims to build a unified network capable of handling a broad spectrum of restoration tasks. Existing attempts include corruption-agnostic approaches, such as AirNet [15], and weather-oriented architectures, such as WGWS [40] and TKL [4]. Beyond these approaches, recent works aim to enhance task-discriminative capability through advanced learning objectives, such as mixture-of-experts and degradation classification [28,13]. From a different perspective, prompt-based methods, including PromptIR [20], AM-PromptIR [35], TextPromptIR [27], VLU-Net [33], and DFPIR [23] further extend this direction by using learned embeddings to guide the model's behavior in a task-adaptive manner. Despite these developments, most studies still treat individual degradation types independently and do not fully address scenarios where multiple corruptions coexist in complex combinations.

Recently, composite models such as OneRestore [10] have emerged to address simultaneous degradations. Nevertheless, they still struggle with clean disentanglement and controllability. While OneRestore represents meaningful progress, it still encounters difficulties in cleanly disentangling and removing mixed degradations. A key issue lies in its limited controllability. (1) The model does not provide an explicit mechanism to bypass restoration and preserve the input content when desired. While a "*clear*" prompt can be used to mimic an identity mapping, existing models do not allow users to intentionally preserve certain degradations or bypass the processing for specific components. (2) The model also lacks support for soft or partial removal. For instance, when a user prefers to retain a mild level of haze for stylistic purposes, the model offers no control over the degree of removal. These controllability limitations become even more problematic under composite degradations. (3) Selectively removing a single corruption within a mixture often introduces unintended artifacts. (4) Moreover, the restoration outcome is sensitive to the sequence in which degradations are processed; for example, removing snow before enhancing low-light yields noticeably different results from performing these steps in the opposite order. Such issues highlight the fundamental need for more explicit and reliable controllability when restoring images affected by multiple degradations.

To address these challenges and strengthen controllability in composite degradation settings, we propose Controllable Unified Image Restoration (CURE), a disentangled restoration framework that improves the adaptability and effectiveness of all-in-one models. CURE integrates several components that provide structured and user-controllable restoration: (1) Identity embedding and an identity loss, which constrain the model to preserve the input when no restoration is required. (2) A ratio-control loss, which enables proportional removal of degradations by mixing the identity embedding with degradation-specific embeddings at a controllable ratio. (3) An intermediate loss, which supervises partially restored outputs and guides the model to learn fine-grained restoration behaviors. (4) A permutation-invariant loss, which ensures consistent results even when degradations are addressed in different orders. Together, these components form a disentangled restoration pipeline that offers flexible, structured, and adaptive control across various degradation combinations. CURE can remove multiple degradations either step-by-step or in a single pass, while also allowing users to specify the desired removal intensity for each degradation. Extensive experiments on composite degradation benchmarks demonstrate that CURE achieves state-of-the-art restoration performance while offering significantly enhanced controllability and robustness compared to existing frameworks.

2 Related work

Single Task Image Restoration. Image restoration encompasses various low-level vision tasks, including dehazing [11,36], deraining [9,5], desnowing [1,22], and low-light enhancement [14,38]. While the aforementioned methods perform well on individual or at most two degradations, they require separate models for different tasks. This paper explores a unified model that handles multiple degradations simultaneously while allowing selective removal of specific ones.

All-in-One Image Restoration. Recent developments in image restoration have explored various techniques to handle images affected by multiple types of degradation. One prominent approach is the partial parameter-sharing One-to-Many strategy, which utilizes a shared backbone with multiple input and output pathways to address different degradation factors [16,2,25]. Beyond this, recent efforts have shifted toward fully shared-weight architectures for universal restoration, enabling more flexible solutions. For instance, TKL [4] proposed a unified model for adverse weather removal using a two-stage knowledge-learning framework, while AirNet [15] introduced a model that handles unknown corruptions without prior knowledge. To improve performance, MoCE-IR [28] employs mixture-of-experts to discriminate tasks based on complexity, while DCPT [13] utilizes degradation classification to explicitly guide the restoration process. Meanwhile, UniRes [39] handles complex degradations by aggregating diffusion latents during inference, and WeatherDiff [41] leverages diffusion priors to restore patch-wise degradations across diverse conditions.

Prompt-based Approaches. Since existing universal methods struggle with interference between different degradation types in complex scenarios, prompt-based approaches are emerging to adaptively guide restoration across multiple degradations. PromptIR [20] is a prompt-based learning approach for all-in-one image restoration that encodes degradation-specific information to dynamically guide restoration. To further advance this concept, text-guided restoration approaches [6,35,27,33] take advantage of text embeddings and language models to

provide more precise guidance. Specifically, VLU-Net [33] addresses the manual tuning constraints of deep unfolding networks using text-driven gradients, while preserving interpretability. On the other hand, AM-PromptIR [35], TextPromptIR [27], InstructIR [6] and DFPIR [23] leverage the capabilities of language models to significantly boost restoration performance. Nevertheless, since controllable restoration is achieved implicitly rather than being treated as a core objective, these methods lack the level of precision required for effective control.
Composite Image Restoration. Composite image restoration deals with images degraded by multiple factors simultaneously. This task poses a unique challenge as the entanglement of different degradations leads to severe interference during the restoration process. To tackle this, OneRestore [10] successfully introduces a text-embedding strategy utilizing versatile scene descriptors to enable adaptive restoration for such diverse composite degradations. Despite its effectiveness in a unified context, similar to previous controllable frameworks, it lacks the capability for precise selective degradation removal and suffers from order dependency in restoration sequences. This deficiency primarily stems from an insufficient consideration of the complex entanglement between degradations, where the blended characteristics of multiple corruptions make it difficult to isolate individual factors. Addressing these limitations, we propose CURE, a controllable restoration framework that leverages text-guided embeddings to achieve fine-grained, disentangled removal of complex degradations.

3 Method

To decouple intertwined degradations and achieve fine-grained control, we propose a unified framework leveraging disentangled learning. Specifically, we construct degradation-aware embeddings via a pre-trained text embedder and ratio-control mechanism, as illustrated in Fig. 1(a). The identity loss shown in Fig. 1(b) preserves input content, essential for disentangled control. Finally, these elements are unified into a restoration pipeline, as shown in Fig. 1(c), allowing the model to adaptively handle composite degradations under precise user guidance.

3.1 Baseline Framework

Our approach operates within a text-guided restoration framework based on an encoder-decoder architecture. The framework receives an image I and a text prompt as inputs. The prompt is encoded into a text embedding, which is then used to modulate the restoration process through cross-attention or feature-injection mechanisms. Following existing text-guided restoration models, the baseline is optimized using a combination of a reconstruction objective and a task-specific auxiliary loss. We formulate this general baseline loss as:

$$\mathcal{L}_{\text{base}} = \mathcal{L}_{\text{rec}}(J, \hat{J}) + \lambda \mathcal{L}_{\text{aux}}, \tag{1}$$

where J and $\hat{J}$ denote the ground truth (GT) and the model output, respectively. In addition, $\mathcal{L}_{\text{rec}}$ represents the pixel-wise reconstruction loss, $\mathcal{L}_{\text{aux}}$ denotes

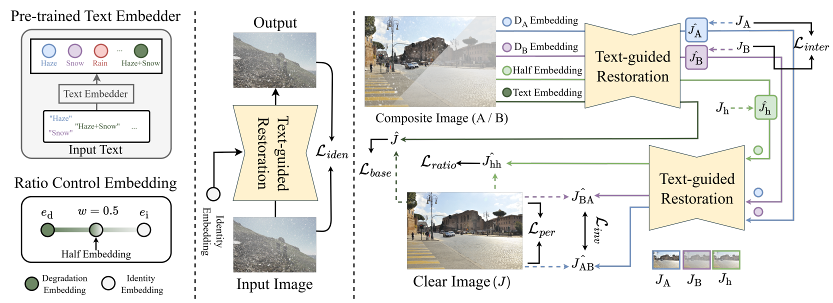

(a) Embedding Generation. (b) Identitiy Loss. (c) Overall Pipeline.

Fig. 1: **Overview of our CURE.** (a) Embedding Generation constructs degradation-aware embeddings via a pre-trained text embedder and our ratio-control mechanism. This enables continuous intensity control by interpolating between identity and degradation embeddings. (b) The identity loss is a loss function designed specifically for the identity input, which preserves the input without any modifications. It serves to suppress changes in the text-guided restoration model's output when the model receives an identity embedding as input. (c) Overall Pipeline of our CURE. The text-guided restoration model handles composite images with multiple degradations (D_A and D_B simultaneously, denoted as D_{AB}) using embeddings selected from Pre-trained Text Embedder sets (a). The Half Embedding is formed by weighted combination of Identity and D_{AB} embeddings, and each restoration flow is processed following the guidance of its corresponding same-colored pipeline

the auxiliary regularization term tailored to specific model designs, and λ is a balancing hyperparameter. Specifically, the implementation of Eq. (1) varies across different baseline methods. For instance, TextPromptIR [27] primarily relies on the L_1 distance for $\mathcal{L}_{\text{rec}}$ without an auxiliary term (i.e., $\lambda = 0$). In the case of OneRestore [10], the reconstruction term $\mathcal{L}_{\text{rec}}$ combines the Smooth L_1 and MS-SSIM losses, while the auxiliary term $\mathcal{L}_{\text{aux}}$ is formulated as a contrastive loss utilizing negative degradation samples. In this work, we retain the inherent $\mathcal{L}_{\text{base}}$ of the selected baseline to preserve its foundational restoration capabilities, while introducing our proposed objectives to enable disentangled control.

3.2 Identity Embedding

The identity operation refers to the case where the model outputs the input image without any modifications. While baseline text embedders are designed for restoration tasks, they do not account for the identity operation. To support this, we assign an embedding vector that explicitly represents the identity operation. This embedding vector is predefined as a constant vector of ones and requires no additional training. We refer to this constant vector as the identity embedding.

The baseline restoration models are trained to preserve the input image unchanged when the identity embedding is fused into their text-guided frameworks.

Ratio-Control Embedding. Existing frameworks are designed to fully restore a clean image regardless of the severity of degradation, as long as a corresponding degradation embedding is provided. This observation suggests the potential to control the degree of restoration by adjusting the degradation embedding vector. To enable this, we leverage the identity embedding that preserves the input without any modification. Specifically, we linearly interpolate between the identity embedding $\mathbf{e}_{\mathrm{i}}$ and the degradation embedding $\mathbf{e}_{\mathrm{d}}$ according to the desired restoration intensity ratio, as follows:

$$\mathbf{e}_{\mathrm{r}} = (1 - w) \cdot \mathbf{e}_{\mathrm{i}} + w \cdot \mathbf{e}_{\mathrm{d}}, \tag{2}$$

where w is an intensity ratio control parameter ranging from identity operation ($w = 0$) to full restoration ($w = 1$), and $\mathbf{e}_{\mathrm{r}}$ denotes the corresponding ratio-control embedding. Adjusting w enables the model to continuously control the restoration intensity for each degradation type, allowing for proportional rather than complete removal of degradations. Note that the ratio-control embedding, like other degradation text embeddings, is integrated into the text-guided modules of the baseline models. This allows it to explicitly guide the network to remove degradations proportionally to the specified ratio.

3.3 Loss Functions

Based on the identity embedding and ratio-control embedding, CURE is trained through disentangled learning guided by the following four loss functions.

Identity Loss. The identity loss is designed to ensure that the model preserves the original content when no degradation needs to be removed. In other words, it enforces the model to maintain the input unchanged when the identity embedding is used as a prompt in the restoration model, as follows:

$$\mathcal{L}_{\mathrm{iden}} = \mathcal{L}_{\mathrm{rec}}(I, \hat{I}), \tag{3}$$

where I is the input image, $\hat{I}$ represents the model output with identity embedding. Through this loss term, the model learns to behave as an identity function when necessary, which can be utilized for the intermediate and permutation-invariant losses in cases involving a single degradation type.

Ratio-Control Loss. The goal of the ratio-control loss is to ensure that, when the ratio-control embedding is used as a prompt, the model removes degradations proportionally to the specified ratio, rather than performing full restoration. Ideally, this would require GT images corresponding to each value of w in Eq. (2), where only the w proportion of degradation is removed. However, in practice, constructing GT images for every possible w is infeasible. Therefore, during training, we only generate and use GT images corresponding to $w = 0.5$. In the experiments, we demonstrate that training with only the $w = 0.5$ case generalizes well, and the model performs seamlessly for other w values during inference.

Accordingly, the ratio-control loss is defined using only the half-intensity case during training, as follows:

$$\mathcal{L}_{\text{ratio}} = \mathcal{L}_{\text{rec}}(J_{\text{h}}, \hat{J}_{\text{h}}) + \mathcal{L}_{\text{rec}}(J, \hat{J_{\text{hh}}}), \tag{4}$$

where J_{h} denotes the half-degraded GT image, $\hat{J}_{\text{h}}$ is the corresponding half-restored output, J is the fully clean GT image, and $\hat{J}_{\text{hh}}$ is the fully restored output obtained by reapplying the restoration model to $\hat{J}_{\text{h}}$. Note that both $\hat{J}_{\text{h}}$ and $\hat{J}_{\text{hh}}$ are generated using the ratio-control embedding for $w = 0.5$. The first term supervises the half-restored output $\hat{J}_{\text{h}}$ to match the corresponding half-degraded GT J_{h}. This explicitly trains the model to remove approximately half of the degradation, establishing fine-grained controllability as intended by the ratio-control embedding. The second term applies the same half-intensity embedding again, guiding the final output $\hat{J}_{\text{hh}}$ to progressively approach the fully clean image J. This enforces a self-consistency constraint that reinforces the linearity of the control space.

Intermediate Loss. To improve controllability and prevent error propagation in sequential restoration, we propose an intermediate loss function that provides precise and explicit supervision for removing each specific degradation at its respective stage in the restoration process. Let D_A and D_B denote two independent degradations present in a composite image. If the image is restored sequentially by first removing D_A and then D_B, or in reverse order, any errors introduced in the first step are likely to propagate into the second. To mitigate this issue, the intermediate loss is designed to help the model handle each degradation independently during the restoration process, as follows:

$$\mathcal{L}_{\text{inter}} = \mathcal{L}_{\text{rec}}(J_{\text{A}}, \hat{J}_{\text{A}}) + \mathcal{L}_{\text{rec}}(J_{\text{B}}, \hat{J}_{\text{B}}), \tag{5}$$

where J_{A} denotes the GT image with degradation D_B removed (but not completely clean), and $\hat{J}_{\text{A}}$ is the model output when removing only D_B. Similarly, J_{B} is the GT image with degradation D_A removed, and $\hat{J}_{\text{B}}$ is the corresponding output when the model restores only D_A. Note that Eq. (5) can be easily extended to cases involving more than two degradations. Also, for single-degradation scenarios, one of D_A or D_B can be treated as "clear," and the identity embedding can be applied accordingly.

Permutation-Invariant Loss. To ensure consistency in sequential restoration of composite degradations, we propose a permutation-invariant loss that enforces the final output to remain unaffected by the order in which degradations are removed. To ensure consistent results regardless of the restoration order, our permutation-invariant loss consists of two terms as follows:

$$\mathcal{L}_{\text{per}} = \mathcal{L}_{\text{rec}}(J, \hat{J_{\text{AB}}}) + \mathcal{L}_{\text{rec}}(J, \hat{J_{\text{BA}}}), \quad \mathcal{L}_{\text{inv}} = \mathcal{L}_{L1}(\hat{J_{\text{AB}}}, \hat{J_{\text{BA}}}), \tag{6}$$

where $\hat{J}_{\text{AB}}$ and $\hat{J}_{\text{BA}}$ represent the model outputs obtained by restoring the image in the order of D_A followed by D_B, and in the reversed order. The $\mathcal{L}_{\text{per}}$ ensures that both restoration sequences produce outputs close to the J. The invariance loss $\mathcal{L}_{\text{inv}}$, calculated via L_1 or Smooth L_1 distance depending on the

Table 1: Quantitative comparisons on CCDD-11 dataset[1].

Types	Methods	PSNR ↑	SSIM ↑	#Params
	Input	16.04	0.61	-
One-to-One	MIRNet [30]	26.17	0.87	31.79M
	MPRNet [31]	27.07	0.86	15.74M
	MIRNetV2 [32]	24.90	0.83	15.74M
	Restormer [29]	25.43	0.84	26.13M
	DGUNet [18]	27.17	0.86	17.33M
	NAFNet [3]	26.78	0.80	17.11M
	Fourmer [37]	24.49	0.82	0.55M
	OKNet [7]	27.54	0.87	4.72M
One-to-Many	AirNet [15]	26.33	0.84	8.93M
	WGWSNet-WG [40]	25.40	0.85	25.76M
	WGWSNet-WS [40]	20.20	0.74	25.76M
	PromptIR [20]	28.07	0.87	38.45M
	AdaIR [8]	27.25	0.86	28.77M
	DFPIR [23]	27.51	0.85	31.10M
Text-guided	AM-PromptIR [35]	26.11	0.84	64.54M
	TextPromptIR [27]	27.55	0.86	93.55M
	OneRestore (Visual) [10]	27.38	0.86	5.98M
	OneRestore (Text) [10]	27.74	0.87	5.98M
Fine Control	AM-PromptIR+ CURE	26.79	0.84	64.54M
	TextPromptIR + CURE	28.15	0.87	93.55M
	OneRestore + CURE (Visual)	27.89	0.86	5.98M
	OneRestore + CURE (Text)	**28.28**	0.87	5.98M

baseline, enforces consistency between the two restoration outputs, $\hat{J}_{\mathrm{AB}}$ and $\hat{J}_{\mathrm{BA}}$. Finally, the restoration model is trained using the baseline loss Eq. (1), along with the four proposed loss functions: identity Eq. (3), ratio-control Eq. (4), intermediate Eq. (5), and permutation-invariant loss functions Eq. (6).

3.4 Experiment Settings

We conduct the training on 4 NVIDIA A100 GPUs, requiring approximately two days for our method. We initialize the process using the checkpoints of each baseline model where performance has reached saturation. During the training of CURE, we set the learning rate to a value moderately lower than the baseline's initial setting. This ensures that the rate is sufficient for learning new tasks yet conservative enough to maintain stability.

Evaluation Dataset and Metrics. We create Controllable Composite Degradation Dataset (CCDD-11), which follows the pipeline introduced in CDD-11 [10] (see Supp. Material for details). Specifically, our CCDD-11 adopts the same

[1] The performance gap compared to the originally reported OneRestore [10] results is primarily due to modifications made to all rain-related datasets. Unlike the original setup, which uses a limited variety of rain masks in the CDD-11 dataset, our version incorporates a broader and more diverse set of rain mask variations.

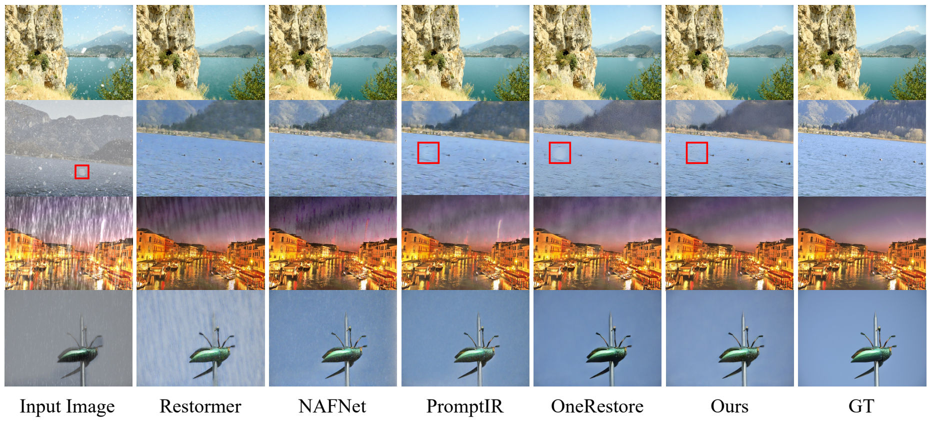

Fig. 2: Qualitative comparisons on the CCDD-11 dataset. Ours refers to the results of OneRestore integrated with our CURE. From top to bottom, each row corresponds to a case in which the input images are degraded by $snow$, $low + haze + snow$, $rain$, and $low + haze + rain$, respectively.

Table 2: Ablation study on different loss functions.

Identity	Ratio control	Intermediate	Permutation-invariant	PSNR
✓				27.85
✓	✓			27.89
✓		✓		28.13
✓			✓	28.20
✓	✓	✓	✓	**28.28**

pipeline as CDD-11, while generating all pairs corresponding to individual degradation removal from composite degradation images. However, due to the limited variety of rain masks in the original CDD-11, we incorporate a more diverse range of rain masks into our CCDD-11. Restoration performance is evaluated using PSNR and SSIM two standard metrics.

Compared Methods. We compare our CURE with seven one-to-one models: MIRNet [30], MPRNet [31], MIRNetV2 [32], Restormer [29], DGUNet [18], NAFNet [3], and Fourmer [37]; five one-to-many models: AirNet [15], AdaIR [8], PromptIR [20], WGWS [40], and DFPIR [23]; and three Text-guided models: TextPromptIR [27], AM-PromptIR [35], and OneRestore [10]. All models are trained and evaluated on the CCDD-11 training and test sets. During CURE training, we exclude images with triple degradations due to combinatorial complexity; however, these are included in the test set to assess generalization. For PromptIR, as the correspondence between the number of learnable prompts and degradation types is unspecified, we evaluated both five and eleven prompts. The version with five prompts achieves better performance and is thus reported.

Table 3: PSNR with identity embedding (*clear* prompt for OneRestore).

Method	Low	Haze	Rain	Snow	Low+Haze	Low+Rain
OneRestore [10]	21.31	23.87	25.04	24.94	24.45	23.89
Ours	55.27	57.31	57.35	57.44	57.58	57.24
Method	**Low+Snow**	**Haze+Rain**	**Haze+Snow**	**Low+Haze+Rain**	**Low+Haze+Snow**	**Average**
OneRestore [10]	23.34	23.60	23.80	23.72	23.67	23.78
Ours	57.00	57.32	57.59	57.66	57.73	**57.05**

Table 4: Intensity ratio and Embedder classification accuracy.

Intensity Ratio w	1.0	0.9	0.8	0.7	0.6	0.5	0.4	0.3	0.2	0.1	0.0
Accuracy	0.3 %	0.4 %	0.8 %	2.5 %	16.4 %	43.8 %	56.2 %	77.9 %	94.7 %	95.0 %	94.9 %

For WGWS, weather-specific parameters are adjusted for CCDD-11, where WG and WS denote the first and second stages, respectively. For AM-PromptIR and TextPromptIR, we align their input prompts with the degradation types in CCDD-11. We construct these text prompts following the protocols specified in their respective papers.

3.5 Comparisons

As reported in Tab. 1, the integration of our CURE framework consistently enhances restoration performance across all three models. Beyond achieving superior quantitative metrics, our framework empowers existing architectures with fine-grained controllability. Interestingly, although not originally designed as a One-to-Composite model, PromptIR outperforms OneRestore on the CCDD-11 dataset. However, it is worth noting that PromptIR employs a much larger model and, by design, is not easily adaptable for controllability. As shown in Fig. 2, our CURE more effectively removes degradations compared to existing approaches in terms of visual quality. Especially, in the second row of Fig. 2, the red box highlights a challenging case where both PromptIR and OneRestore leave noticeable snow artifacts, whereas our CURE achieves more effective removal under composite degradation ($low + haze + snow$). We believe our disentangled learning approach provides more effective guidance, offering not only enhanced controllability but also superior performance in composite degradation restoration.

3.6 Ablation Studies

In this subsection, we conduct ablation studies to evaluate the effectiveness of the proposed loss functions, identity embedding, and ratio-control embedding. Detailed settings for each ablation experiment and additional qualitative results are provided in the Supplementary Material.

Effect of Losses. To verify the effectiveness of the proposed loss functions, we conduct an ablation study by excluding each component, as shown in Tab. 2. While identity loss alone leads to improved performance compared to OneRestore [10],

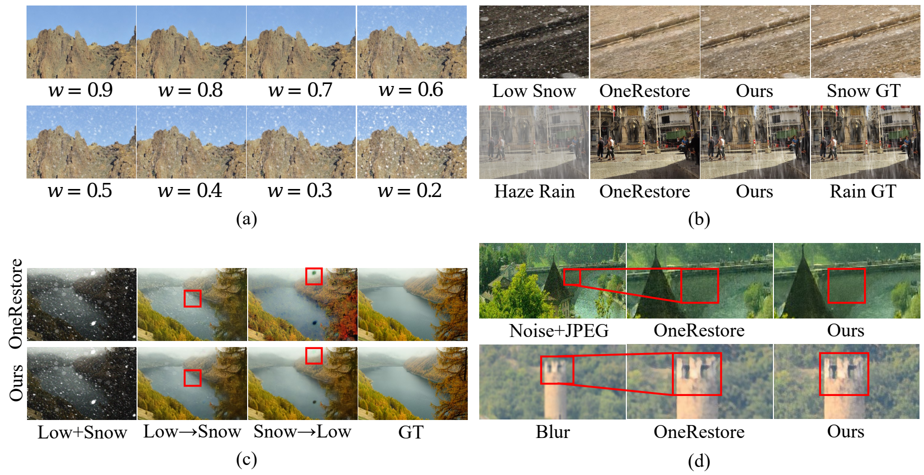

Fig. 3: (a) Ratio-control embedding. (b) Selective restoration. (c) Restoration order dependency. (d) Blur-noise-JPEG composite degradation.

Table 5: PSNR for selective restoration on two combined degradations

Degradation	Task	OneRestore [10]	+CURE	Task	OneRestore	+CURE
Low+Haze	Delow	20.89	30.61	Dehaze	22.56	35.50
Low+Rain	Delow	24.67	26.74	Derain	28.76	36.37
Low+Snow	Delow	23.42	26.68	Desnow	27.53	34.30
Haze+Rain	Dehaze	27.73	33.01	Derain	34.69	40.30
Haze+Snow	Dehaze	25.56	33.13	Desnow	31.61	37.81
Average	-	24.45	**30.03**	-	29.03	**36.86**

incorporating additional loss functions results in further performance gains. The best performance is achieved when all loss components are used together, indicating that the proposed loss functions effectively enable the text-guided restoration model to disentangle mixed degradations.

Effect of Identity Embedding. The identity embedding is designed to enforce the preservation of the input image, even in the presence of degradation. To verify its effectiveness, we measure the PSNR between the input and the output generated using the identity embedding. For comparison, we evaluate OneRestore [10] with the "*clear*" prompt, which most closely matches the identity condition in its framework. As shown in Tab. 3, our CURE yields outputs with minimal deviation from the inputs, demonstrating superior identity preservation compared to OneRestore and validating our approach.

Effect of Ratio-Control Embedding. The ratio-control embedding is designed to guide the restoration intensity by removing only a specified proportion of the degradation. To achieve this, we construct a prompt by linearly combining the identity and target degradation embeddings according to Eq. (2). To quantitatively evaluate this controllability, we measure classification accuracy across varying values of w in Eq. (2). Note that $w = 1$ leads to full degradation removal and

Table 6: Classification performance after selective restoration. The top-row degradation types represent the remaining degradations after selective restoration.

Method	Low	Haze	Rain	Snow	Low+Haze	Low+Rain	Low+Snow	Haze+Rain	Haze+Snow	Average
OneRestore [10]	76.7%	61.7%	73.8%	48.6%	76.7%	42.0%	24.5%	60.5%	15.5%	53.33%
+ CURE	98.6%	85.5%	93.7%	94.3%	88.7%	91.6%	98.4%	95.6%	94.6%	**93.44%**

Table 7: PSNR evaluation with varying restoration orders.

Degradation	Task	OneRestore [10]	+CURE	Task	OneRestore	+CURE	Task	OneRestore	+CURE
Low+Haze	Delow → Dehaze	20.95	25.80	Dehaze → Delow	21.20	24.76	Dehaze + Delow	25.07	25.28
Low+Rain	Delow → Derain	22.66	25.77	Derain → Delow	24.45	25.54	Derain + Delow	25.27	25.41
Low+Snow	Delow → Desnow	21.63	25.26	Desnow → Delow	22.63	24.87	Desnow + Delow	24.64	24.95
Haze+Rain	Dehaze → Derain	27.74	30.12	Derain → Dehaze	28.92	30.75	Derain + Dehaze	29.62	30.52
Haze+Snow	Dehaze → Desnow	26.26	28.66	Desnow → Dehaze	27.13	29.59	Desnow + Dehaze	28.65	29.42
Average	-	23.85	**27.12**	-	24.87	**27.10**	-	26.65	**27.12**

low classifier accuracy, while $w = 0$ preserves the degradation and yields high accuracy. As reported in Tab. 4, the classification performance varies consistently with w, indicating that CURE effectively enables fine-grained intensity control. This controllability is visually confirmed in Fig. 3(a), which demonstrates the precise adjustment of restoration intensity.

Selective Restoration. By enforcing explicit constraints on degradation independence through our disentangled learning objectives, our framework empowers models to precisely isolate and remove a target degradation while preserving other components. This capability is crucial for addressing composite degradations. In this section, we compare our method against OneRestore [10], which serves as a strong baseline due to its capability to handle composite degradations. For instance, given an input with composite degradations of *low* and *haze*, if the task is *delow*, the goal is to remove only the *low*-light effect while retaining the *haze*. In Tab. 5, we report the quantitative results on selective restoration tasks involving two degradations. Here, the GT for evaluation is synthesized by applying only the remaining degradation types to the clean image. While the baseline supports single-degradation embeddings, it lacks explicit disentanglement supervision. Consequently, as shown in Tab. 5, our model achieves superior performance across all selective tasks. We also validate our method on triple-composite scenarios. Detailed results provided in the Supplementary Material confirm that our model maintains its performance advantage over the baseline even in these composie settings. To further validate the disentanglement capability of CURE, we conduct a classification-based analysis. Specifically, we perform selective restoration on composite images to remove a target degradation and then classify the remaining degradation types in the output using a pre-trained classifier. This protocol objectively verifies whether the non-target degradations are successfully preserved without distortion. As shown in Tab. 6, our method consistently outperforms the baseline in correctly identifying the remaining degradations, regardless of the type being removed. Moreover, qualitative comparisons in Fig. 3(b) confirm that the selective restoration closely matches the corresponding GT. These results demonstrate that CURE effectively decomposes complex composite degradations into their individual components for fine-grained control.

Restoration-Order Dependency. For composite degradations, restoration can be performed by sequentially removing each degradation type. As shown in Tab. 7, our method maintains consistent and high performance regardless of the restoration sequence, whereas OneRestore [10] exhibits significant performance variance depending on the order. We attribute this order-invariant robustness to the proposed permutation-invariant loss defined in Eq. (6), which encourages the model to treat each degradation independently. Notably, even when performing simultaneous restoration in a single step (using composite embeddings), our CURE outperforms the baseline, further validating its superior controllability. This can also be visually confirmed in Fig. 3(c).

Blur-Noise-JPEG Degradation. To further validate the generalizability of the proposed CURE pipeline, we conduct experiments across not only diverse weather removal scenarios but also digital degradation scenarios involving $blur$, $noise$, and $JPEG$. It is worth noting that CURE consistently outperforms the baselines in this setting as well. As illustrated in Fig. 3(d), our method achieves successful restoration by effectively removing complex digital artifacts while preserving image details. We also perform the same evaluations as CCDD-11 including intensity analysis, classification accuracy, selective restoration, and restoration order to ensure consistency. Detailed quantitative results and real-world evaluations are provided in the Supplementary Material.

4 Conclusion and Limitations

In this paper, we have proposed CURE, a disentangled learning framework for controllable all-in-one image restoration under composite degradation scenarios. CURE is built upon four specialized restoration objectives: (1) identity loss for preserving the input when no degradation is present, (2) ratio-control loss for adjusting the intensity level of restoration through prompt-based interpolation, (3) intermediate loss for guiding selective removal of individual degradations, and (4) permutation-invariant loss to ensure consistent results regardless of the degradation removal order. These components collectively enable the model to effectively disentangle and selectively restore target degradations without affecting unrelated image content, offering a significant leap in restoration flexibility. Our method supports fine-grained intensity control and achieves competitive performance in comparison on composite degradation benchmarks. Our findings underscore the importance of disentangled representations and targeted guidance in building flexible and controllable restoration systems tailored for complex scenarios. However, our approach has several limitations. First, the model may struggle to generalize to unknown degradations not encountered during training, indicating a need to extend the framework toward handling unseen or open-set degradation types. Second, while effective for up to three concurrent degradations, the performance may decrease when a larger number of types are combined, particularly in the presence of complex spatially varying degradations. Addressing these challenges would further enhance the robustness and scalability of controllable restoration systems in real-world applications.

References

1. Bossu, J., Hautière, N., Tarel, J.P.: Rain or snow detection in image sequences through use of a histogram of orientation of streaks. IJCV **93**(3), 348–367 (2011)
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. Springer (2022)
4. 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)
5. Chen, X., Pan, J., Dong, J.: Bidirectional multi-scale implicit neural representations for image deraining. In: CVPR. pp. 25627–25636 (2024)
6. Conde, M.V., Geigle, G., Timofte, R.: Instructir: High-quality image restoration following human instructions. In: ECCV. pp. 1–21. Springer (2024)
7. Cui, Y., Ren, W., Knoll, A.: Omni-kernel network for image restoration. In: AAAI. vol. 38, pp. 1426–1434 (2024)
8. Cui, Y., Zamir, S.W., Khan, S., Knoll, A., Shah, M., Khan, F.S.: Adair: Adaptive all-in-one image restoration via frequency mining and modulation. In: ICLR. pp. 57335–57356. International Conference on Learning Representations, ICLR (2025)
9. 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)
10. Guo, Y., Gao, Y., Lu, Y., Zhu, H., Liu, R.W., He, S.: Onerestore: A universal restoration framework for composite degradation. In: ECCV. pp. 255–272. Springer (2024)
11. He, K., Sun, J., Tang, X.: Single image haze removal using dark channel prior. In: CVPR. pp. 1956–1963 (2009)
12. He, K., Sun, J., Tang, X.: Single image haze removal using dark channel prior. IEEE TPAMI **33**(12), 2341–2353 (2010)
13. Hu, J., Jin, L., Yao, Z., Lu, Y.: Universal image restoration pre-training via degradation classification. arXiv preprint arXiv:2501.15510 (2025)
14. Land, E.H.: The retinex theory of color vision. Scientific American **237**(6), 108–129 (1977)
15. 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)
16. Li, R., Tan, R.T., Cheong, L.F.: All in one bad weather removal using architectural search. In: CVPR. pp. 3175–3185 (2020)
17. Lin, Y., Lin, Z., Chen, H., Pan, P., Li, C., Chen, S., Wen, K., Jin, Y., Li, W., Ding, X.: Jarvisir: Elevating autonomous driving perception with intelligent image restoration. In: CVPR. pp. 22369–22380 (2025)
18. Mou, C., Wang, Q., Zhang, J.: Deep generalized unfolding networks for image restoration. In: CVPR. pp. 17399–17410 (2022)
19. Porav, H., Bruls, T., Newman, P.: I can see clearly now: Image restoration via de-raining. In: ICRA. pp. 7087–7093. IEEE (2019)
20. Potlapalli, V., Zamir, S.W., Khan, S., Khan, F.S.: Promptir: Prompting for all-in-one blind image restoration. In: NeurIPS. pp. 71275–71293 (2023)
21. Qu, L., Tian, J., He, S., Tang, Y., Lau, R.W.: Deshadownet: A multi-context embedding deep network for shadow removal. In: CVPR. pp. 4067–4075 (2017)

22. Quan, Y., Tan, X., Huang, Y., Xu, Y., Ji, H.: Image desnowing via deep invertible separation. IEEE TCSVT **33**(7), 3133–3144 (2023)
23. Tian, X., Liao, X., Liu, X., Li, M., Ren, C.: Degradation-aware feature perturbation for all-in-one image restoration. In: CVPR. pp. 28165–28175 (2025)
24. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. NeurIPS **30** (2017)
25. Wang, Y., Ma, C., Liu, J.: Smartassign: Learning a smart knowledge assignment strategy for deraining and desnowing. In: CVPR. pp. 3677–3686 (2023)
26. Wang, Z., Cun, X., Bao, J., Zhou, W., Liu, J., Li, H.: Uformer: A general u-shaped transformer for image restoration. In: CVPR. pp. 17683–17693 (2022)
27. Yan, Q., Jiang, A., Chen, K., Peng, L., Yi, Q., Zhang, C.: Textual prompt guided image restoration. EAAI **155**, 110981 (2025)
28. Zamfir, E., Wu, Z., Mehta, N., Tan, Y., Paudel, D.P., Zhang, Y., Timofte, R.: Complexity experts are task-discriminative learners for any image restoration. In: CVPR. pp. 12753–12763 (2025)
29. 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)
30. 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. Springer (2020)
31. 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)
32. 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)
33. Zeng, H., Wang, X., Chen, Y., Su, J., Liu, J.: Vision-language gradient descent-driven all-in-one deep unfolding networks. In: CVPR. pp. 7524–7533 (2025)
34. Zhang, H., Patel, V.M.: Density-aware single image de-raining using a multi-stream dense network. In: CVPR. pp. 695–704 (2018)
35. Zhang, S., Wang, J., He, N., Xu, S., Liu, S., Yu, P., Ma, X.: Adaptive multi-modal prompting for universal image restoration amidst diverse degradations: S. zhang et al. Vis. Comput. pp. 1–13 (2025)
36. 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)
37. Zhou, M., Huang, J., Guo, C.L., Li, C.: Fourmer: An efficient global modeling paradigm for image restoration. In: ICML. pp. 42589–42601. PMLR (2023)
38. 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 **26**, 650–662 (2023)
39. Zhou, M., Ye, K., Delbracio, M., Milanfar, P., Patel, V.M., Talebi, H.: Unires: Universal image restoration for complex degradations. arXiv preprint arXiv:2506.05599 (2025)
40. 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)
41. Özdenizci, O., Legenstein, R.: Restoring vision in adverse weather conditions with patch-based denoising diffusion models. IEEE TPAMI **45**(8), 10346–10357 (2023)