

Figure 2: The superiority of our approach in recovering luminance and chrominance under unknown degradations. (a) Input SDR image. (b) Results of LSNet (Guo et al. 2023), trained on a single degradation type, produces results that closely mimic the input SDR. (c) Results of ICTCPNet (Huang et al. 2023) with multi-degradation training, which still fail to recover realistic brightness and color. (d) Our proposed method (RealRep) achieves significantly better recovery of both luminance and chrominance, demonstrating strong generalization to unseen degradations. (e) Ground truth HDR image.

erality. As deep learning evolves (Li et al. 2025), recent deep learning approaches (Chen et al. 2021c; Shao et al. 2022; Xu et al. 2022a; Guo et al. 2023; He et al. 2022; Xu et al. 2022b, 2023, 2024a,b) treat iTM as a mapping problem and have improved visual fidelity. However, most are trained on SDRs synthesized by fixed tone-mapping (TM) curves, such as Reinhard or YouTube styles, and struggle to generalize to real-world SDRs with diverse degradations. As shown in Figure 1(b), different degradation types exhibit distinct statistical variations in luminance and chrominance. Prior works often overlook chroma variations and learn entangled representations, resulting in poor separation of inputs with different styles (Figure 1(c, middle)). In contrast, our method disentangles luminance and chrominance, learning attribute-aware representations that are clearly separated by degradation type (Figure 1(c, right)), enabling better generalization to realistic SDR scenarios.

To address these challenges, we propose RealRep, a novel SDR-to-HDR framework designed for real-world scenarios involving complex and diverse degradations. RealRep employs an attribute-disentangled representation learning strategy that separately encode luminance and chrominance, allowing the model to capture degradation-aware, generalizable priors for inverse tone mapping. To further enhance mapping robustness, we introduce DDACMNet, a degradation-domain-aware controlled mapping network that leverages these priors through a control-aware normalization mechanism. By learning from a wide spectrum of degradation types, RealRep constructs a robust representation space that models both large-scale brightness variations and fine-grained color structures. As shown in Figure 2, our method produces perceptually more faithful HDR reconstructions compared to state-of-the-art approaches, particularly under severe degradation and color distortion.

Our main contributions are summarized as follows:

- **Attribute-disentangled representation learning:** We propose RealRep, which disentangles luminance and chrominance to reduce style variation and builds a robust embedding space for accurate SDR-to-HDR conversion.
- **Degradation-controlled mapping:** We design DDACMNet to inject disentangled priors via a zero-

initialized controller and stabilize early training through a two-stage degradation-guided strategy.

- **Robust performance on diverse SDR styles:** Extensive experiments on datasets with varied and unseen degradations show that RealRep achieves superior visual quality and generalization over state-of-the-art methods.

## Related Work

### Multiple Styles SDR-to-HDR

SDR-to-HDR conversion has recently attracted growing attention due to its relevance in media enhancement and consumer displays. Kim et al. (Kim and Kim 2018) pioneered a joint framework for super-resolution and SDR-to-HDR conversion using CNNs. Deep SR-ITM (Kim, Oh, and Kim 2019a) introduced spatially adaptive contrast modulation, while JSI-GAN (Kim, Oh, and Kim 2019b) incorporated dynamic convolution and adversarial learning to enhance visual quality. To support learning-based methods, HDRTV1K (Chen et al. 2021c) and HDRTV4K (Guo et al. 2023) were proposed, along with HDRTVNet and LSNet. These models combine global color mapping, local enhancement, and luminance-aware strategies. FMNet (Xu et al. 2022a) further introduced frequency-aware modulation to reduce artifacts. Despite these advances, most approaches rely on synthetic SDRs generated by fixed tone-mapping curves (e.g., Reinhard, Hable), which limits generalization to real-world SDR types. Guo et al. (Guo et al. 2023) also emphasized this issue. Our method addresses this limitation by learning attribute-disentangled representations from diverse SDR types, enabling robust luminance and chrominance modeling without reliance on predefined tone-mapping operators.

### Image Enhancement for Multi-Degradations

Image enhancement tasks, such as super-resolution, denoising, dehazing, deraining, and deblurring, were traditionally addressed separately. Early CNN-based methods like SR-CNN (Dong et al. 2014) and its deeper extensions (Kim, Lee, and Lee 2016; Lim et al. 2017; Ahn, Kang, and Sohn 2018; Ma et al. 2023) achieved strong results in super-resolution, while similar task-specific networks were developed for other degradations (Tian et al. 2020; Li et al.

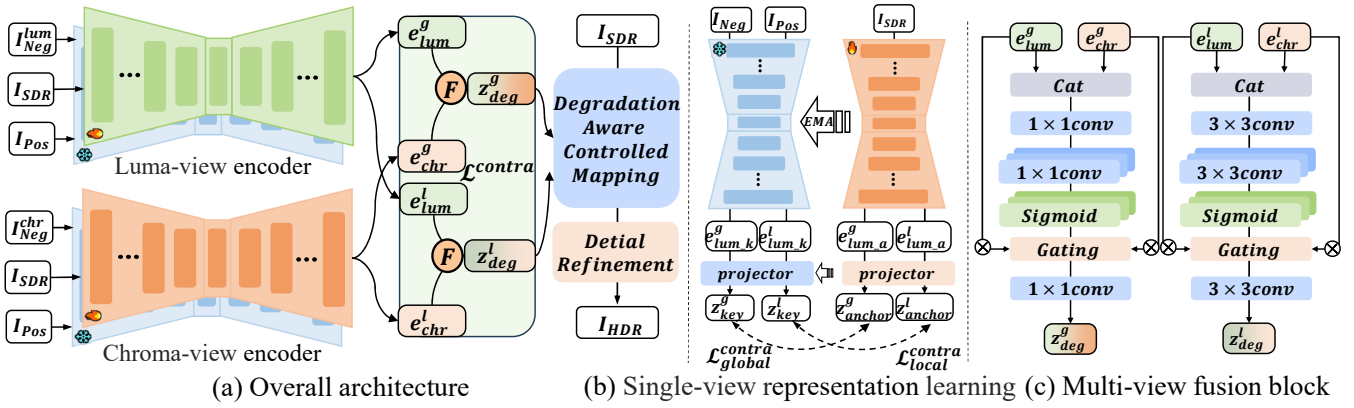


Figure 3: Overview of the proposed multi-degradation SDR-to-HDR framework. Our architecture consists of a pair of multi-view degradation encoders that separately extract global and local representations of luminance and chrominance, a multi-view fusion module that facilitates spatial and channel-wise interaction, and the Degradation-Domain Aware Controlled Mapping Network (DDACMNet), which performs adaptive mapping and detail refinement conditioned on degradation-aware features. The encoders are jointly trained with a disentanglement objective to ensure cross-domain consistency. Our framework enables dynamic luminance expansion and perceptually faithful color gamut reconstruction across a wide range of real-world SDR degradations.

2017; Yang et al. 2019; Eboli, Sun, and Ponce 2020). To improve generalization, recent works explore unified models. Transformer-based frameworks like IPT (Chen et al. 2021a), AirNet (Li et al. 2022), and PromptIR (Potlapalli et al. 2023) handle multiple degradations within a single network. While unified models improve scalability under unknown degradations, they offer limited control in HDR contexts. Our method targets HDR enhancement by disentangling luminance and chrominance, enabling more robust and interpretable reconstruction.

## Degradation Representation

Degradation representation is critical for guiding effective enhancement, especially under diverse or unknown degradation conditions. Recent approaches learn implicit representations that encode degradation characteristics without explicit labels. AirNet (Li et al. 2022) leverages contrastive learning, while PromptIR (Potlapalli et al. 2023) conditions restoration with learned prompts. We adopt a similar strategy for multi-style SDR-to-HDR enhancement, learning a degradation-aware representation that captures style-specific information to guide adaptive HDR reconstruction.

## Methodology

### Overview

Figure 3 shows an overview of our RealRep framework. It features: (1) a contrastive multi-view encoder that disentangles luminance and chrominance across global and local views, and (2) a degradation-aware mapping network (DDACMNet) guided by hierarchical priors.

Given an SDR input  $x$ , the encoder extracts luminance and chrominance features from global and local views. These are projected and normalized into a multi-view representation  $z$ , and supervised via contrastive loss to encourage

degradation-invariant and attribute-disentangled learning. A fusion block aggregates  $z$  into a unified latent code.

DDACMNet performs progressive SDR-to-HDR mapping, modulated by global and local degradation priors derived from the encoder. This enables both global consistency and spatial adaptivity to diverse SDR degradations.

The overall process is defined as:

$$\hat{y} = f_{\text{DDACM}} \left( f_{\text{fusion}} \left( e_{\text{lum}}^{g,l}, e_{\text{chr}}^{g,l} \right), x \right). \quad (1)$$

### Attribute-Disentangled Representation Learning

**Attribute Feature Extraction.** We adopt a UNet-based encoder (Ronneberger, Fischer, and Brox 2015) to extract hierarchical features from SDR inputs. The encoder-decoder structure, enhanced with skip connections, is well-suited for SDR-to-HDR tasks that require both global tone mapping and local detail preservation (Chen et al. 2021c).

Given an input SDR image  $x$ , we extract two modality-specific features: the luminance feature  $e_{\text{lum}} \in \mathbb{R}^{H \times W \times C}$  and the chrominance feature  $e_{\text{chr}} \in \mathbb{R}^{H \times W \times C}$ . Each feature is further decomposed into global and local views. Specifically, the global features  $e_c^g$  are extracted from the UNet bottleneck via a  $1 \times 1$  convolution followed by global average pooling, while the local features  $e_c^l$  are obtained from the decoder and fused with encoder outputs through skip connections, where  $c \in \{\text{lum}, \text{chr}\}$ . Each view-specific feature is then passed through a projection head  $\phi_c^s$  and normalized using  $\ell_2$ -normalization, resulting in the following embeddings:

$$z_c^s = \frac{\phi_c^s(e_c^s)}{\|\phi_c^s(e_c^s)\|_2}, \quad \text{for } c \in \{\text{lum}, \text{chr}\}, s \in \{g, l\}. \quad (2)$$

These four embeddings  $\{z_{\text{lum}}^g, z_{\text{lum}}^l, z_{\text{chr}}^g, z_{\text{chr}}^l\}$  form the basis of our multi-view representation, and are later supervised using contrastive learning to enforce attribute-level disentanglement.

---

**Algorithm 1: Luma-/Chroma-aware Negative Exemplars Generation**


---

**Input:** Multi-degradation dataset  $D$ , parameters  $k_l$  (luminance negatives),  $k_c$  (chrominance negatives)

**Output:** Luminance negatives  $N_l$ , Chrominance negatives  $N_c$

- 1: Initialize  $N_l \leftarrow \emptyset, N_c \leftarrow \emptyset$
  - 2: **for** each image  $I \in D$  **do**
  - 3:   Convert to YCbCr: Extract brightness ( $L_{\text{ref}}$ ) and color components ( $C_{\text{ref}}$ )
  - 4:   Initialize lists  $List_l \leftarrow \emptyset, List_c \leftarrow \emptyset$
  - 5:   **for** each degradation  $x_d \in I.\text{degradations}$  **do**
  - 6:     Extract brightness ( $L_d$ ) and color components ( $C_d$ ) of  $x_d$
  - 7:     Compute differences:  $D_l \leftarrow \|L_d - L_{\text{ref}}\|_1, D_c \leftarrow \|C_d - C_{\text{ref}}\|_1$
  - 8:     Append ( $D_l, x_d$ ) to  $List_l$ , ( $D_c, x_d$ ) to  $List_c$
  - 9:   **end for**
  - 10:   Select top  $k_l$  luminance-based and top  $k_c$  chrominance-based samples
  - 11:    $I_l \leftarrow \text{TopK}(List_l, k_l), I_c \leftarrow \text{TopK}(List_c, k_c)$
  - 12:   **for** each  $(x_l, x_c)$  in  $I_l$  and  $I_c$  **do**
  - 13:     Add ( $L_l, C_{\text{ref}}$ ) to  $N_l$
  - 14:     Add ( $L_{\text{ref}}, C_c$ ) to  $N_c$
  - 15:   **end for**
  - 16: **end for**
  - 17: Convert  $N_l, N_c$  back to RGB
  - 18: **Return**  $N_l, N_c$
- 

**Disentangled Representation Learning.** To learn degradation-invariant and attribute-disentangled representations for SDR-to-HDR mapping, we introduce a contrastive representation separation paradigm that explicitly leverages the statistical inconsistency between luminance and chrominance under different SDR degradations.

While SDR images derived from the same HDR source should share consistent scene semantics, they often exhibit significant variations in local brightness and color distributions due to diverse tone mapping processes. We exploit this observation by constructing contrastive samples that isolate such variations. Specifically, positive samples are generated by applying spatial augmentations (e.g., flips) to the same SDR input, which preserve both luminance and chrominance identity while introducing geometric diversity.

To construct hard negatives, we develop a Luma-/Chroma-aware mining strategy based on the assumption that SDR degradations primarily distort either brightness or color in a content-independent way. For each anchor, we compute the L1 distance over the luminance or chrominance channels with respect to a candidate pool, and select the top- $k$  dissimilar samples as negatives. To further disentangle the contrastive supervision, we synthesize attribute-specific hard negatives by replacing either the luminance or chrominance channel of the anchor with that of a distinct SDR sample. This ensures that the contrastive signal focuses on a single modality at a time.

Formally, the contrastive loss is defined over all attribute-




Figure 4: Illustration of (a) Dense Controlled Mapping (DCM) and (b) Sparse Controlled Mapping (SCM).

view combinations:

$$\mathcal{L}_{\text{contra}} = \sum_{c \in \{\text{lum}, \text{chr}\}} \sum_{s \in \{\text{g}, \text{l}\}} -\log \left( \frac{e^{z_c^s \cdot z_c^{s+}}}{e^{z_c^s \cdot z_c^{s+}} + \sum_{j=1}^{N_c} e^{z_c^s \cdot z_{c,j}^s}} \right) \quad (3)$$

where  $c \in \{\text{lum}, \text{chr}\}$  denotes the feature type, i.e., luminance (lum) or chrominance (chr);  $s \in \{\text{g}, \text{l}\}$  denotes the feature scope, i.e., global (g) or local (l);  $z_c^s$  is the feature vector of the anchor sample for feature type  $c$  and scope  $s$ ;  $z_c^{s+}$  is the positive sample, which is generated by applying data augmentation to the anchor sample  $z_c^s$ ;  $z_{c,j}^s$  are negative samples, where  $j = 1, 2, \dots, N_c$ , and  $N_c$  is the total number of negative samples for feature type  $c$ .

This objective encourages the model to align samples that share semantic content while explicitly penalizing variations in luminance and chrominance caused by SDR degradation. As a result, the learned embedding space becomes both semantically consistent and attribute-disentangled, providing a robust foundation for high-fidelity HDR reconstruction.

**Multi-View Attribute Fusion.** After contrastive training, we fuse the view-specific embeddings into a unified latent representation. The fusion module leverages multi-head convolutions and learnable gating mechanisms to dynamically aggregate global-local and luminance-chrominance cues. The output embedding  $z$  integrates attribute-aware priors and is passed to the degradation-adaptive HDR mapping network described in Section 3(c).

### Degradation-Domain Aware Controlled Mapping Network

**Degradation-Aware Feature Modulation** To handle spatially-varying SDR degradations, we design a hierarchical degradation-aware feature modulation mechanism composed of two sequential stages: a stack of  $N$  Dense Controlled Mapping (DCM) modules followed by  $N$  Sparse Controlled Mapping (SCM) modules. These modules inject degradation priors into the network via affine transformations over layer-normalized features, enabling both global and local adaptation to degradation patterns. Inspired by ControlNet (Zhang, Rao, and Agrawala 2023), we insert zero-initialized control layers before each modulation to gradually introduce prior information in a stable manner.

Each DCM module performs global, channel-wise modulation conditioned on a compact degradation embedding  $z_{\text{deg}}^g \in \mathbb{R}^{C \times 1 \times 1}$ . A pair of zero-initialized  $1 \times 1$  convolutions predict affine parameters  $\gamma$  and  $\beta$ , which are applied to normalized features as  $\hat{x} = \gamma \cdot x + \beta$ . To enhance modulation

Methods	2390eetf_gm	2446a	2446c_gm	davinci	hc_gm	ocio2	reinhard	youtube	Average
	PSNR↑ / SSIM↑	PSNR↑ / SSIM↑	PSNR↑ / SSIM↑	PSNR↑ / SSIM↑	PSNR↑ / SSIM↑	PSNR↑ / SSIM↑	PSNR↑ / SSIM↑	PSNR↑ / SSIM↑	PSNR↑ / SSIM↑
HDRUNet	29.91 / 0.9219	25.55 / 0.8861	24.17 / 0.8509	23.17 / 0.8790	24.67 / 0.8548	25.12 / 0.8676	23.66 / 0.8887	22.76 / 0.8758	24.88 / 0.8781
HDRTVNet	28.97 / 0.9059	23.73 / 0.8502	26.05 / 0.8597	25.50 / 0.8815	26.91 / 0.8621	26.46 / 0.8699	23.23 / 0.8722	25.31 / 0.8709	25.77 / 0.8716
FMNet	27.96 / 0.9050	23.84 / 0.8753	22.29 / 0.8327	24.48 / 0.8835	23.16 / 0.8390	24.88 / 0.8603	22.26 / 0.8603	24.31 / 0.8755	24.15 / 0.8665
ICTCPNet	29.58 / 0.8548	27.02 / 0.8548	27.76 / 0.8594	25.67 / 0.8705	<u>28.27</u> / 0.8612	<u>29.97</u> / 0.8709	25.63 / 0.8694	25.71 / 0.8639	27.45 / 0.8631
LSNet	<u>33.23</u> / <u>0.9408</u>	<u>31.18</u> / <u>0.9080</u>	28.05 / <u>0.8731</u>	25.63 / 0.9000	27.34 / 0.8712	<u>29.56</u> / <u>0.8913</u>	<u>25.67</u> / <u>0.8970</u>	27.01 / 0.9019	28.46 / <u>0.8979</u>
AirNet	23.00 / 0.8594	21.72 / 0.8519	21.02 / 0.8044	22.22 / 0.8672	21.59 / 0.8107	22.18 / 0.8407	18.72 / 0.8379	22.20 / 0.8512	21.58 / 0.8404
PromptIR	30.06 / 0.9148	30.72 / 0.9073	<u>28.23</u> / 0.8691	27.83 / 0.9028	28.00 / 0.8740	29.30 / 0.8838	24.97 / 0.8931	27.62 / 0.9068	28.34 / 0.8940
RAM-PromptIR	30.25 / 0.9188	30.83 / 0.9103	28.20 / 0.8711	<u>28.07</u> / <u>0.9088</u>	28.09 / <u>0.8770</u>	29.31 / 0.8828	25.29 / 0.9001	<u>27.77</u> / <u>0.9108</u>	<u>28.48</u> / <u>0.8975</u>
Ours	<b>34.13</b> / <b>0.9507</b>	<b>34.41</b> / <b>0.9333</b>	<b>30.38</b> / <b>0.8936</b>	<b>30.05</b> / <b>0.9266</b>	<b>30.47</b> / <b>0.8976</b>	<b>31.36</b> / <b>0.9083</b>	<b>27.59</b> / <b>0.9292</b>	<b>30.03</b> / <b>0.9358</b>	<b>31.05</b> / <b>0.9219</b>

Table 1: Comparison of quantitative results on HDRTV4K (known degradation benchmark). The first group contains CNN-based methods, while the second group includes transformer-based methods. All methods were trained using all degraded SDR versions from the HDRTV4K dataset. **Bold** and underline indicate the best and second-best results, respectively.

Method	HDR-VDP3↑	EHL↓	FHLP↓	EWG↓	FWGP↓
HDRTVNet	8.69	0.0079	0.0630	0.0024	0.0287
FMNet	8.75	0.0074	0.0621	0.0022	0.0273
ICTCPNet	8.90	0.0067	0.0580	0.0019	0.0250
LSNet	9.19	0.0060	0.0567	0.0017	<u>0.0241</u>
AirNet	8.48	0.0086	0.0662	0.0027	0.0303
PromptIR	9.11	0.0078	0.0735	<b>0.0012</b>	0.0345
RAM-PromptIR	9.17	0.0070	0.0716	<u>0.0013</u>	0.0329
<b>Ours</b>	<b>9.35</b>	<b>0.0032</b>	<b>0.0420</b>	<u>0.0013</u>	<b>0.0087</b>

Table 2: Quantitative comparison using HDR-specific perceptual quality metrics. ↑ indicates higher is better, ↓ indicates lower is better. **Bold** and underline indicate the best and second-best results, respectively.

flexibility, we extend this to degradation-aware LayerNorm, where the affine parameters are dynamically predicted from  $z_{\text{deg}}^g$ . We stack  $N$  such DCM modules to progressively inject global degradation priors.

After global modulation, we apply  $N$  SCM modules to further refine features under spatially-varying degradations. Each SCM module utilizes a local degradation map  $z_{\text{deg}}^l \in \mathbb{R}^{C \times H \times W}$  to generate pixel-wise affine parameters. These are obtained via two parallel branches, each consisting of a shared  $1 \times 1$  convolution, a ReLU activation, and a zero-initialized  $3 \times 3$  convolution. The resulting parameters modulate the features using  $\hat{x} = \gamma \odot x + \beta$ , enabling spatially adaptive transformation. Outputs from both DCM and SCM stages are fused via residual connections, forming a unified degradation-aware modulation pipeline.

**Detail Refinement** To enhance textures and remove residual artifacts, we append a detail refinement module consisting of stacked ResBlocks:

$$I_{\text{HDR}} = F_{\text{refine}}(I_{\text{SCM}}) = \text{Conv}_{3 \times 3} \circ [\text{ResBlock}]^{N_r} \circ \text{Conv}_{3 \times 3}(I_{\text{SCM}}) \quad (4)$$

where  $I_{\text{SCM}}$  is the output after the final SCM module, and  $N_r$  is the number of residual blocks.

**Training Strategy** We adopt a two-stage training scheme. In Stage 1, contrastive learning is disabled and control layers are initialized as standard convolutions, allowing stable optimization of the encoder and modulation modules. In Stage 2, contrastive supervision is introduced via an EMA-updated negative encoder (He et al. 2019), and all control layers are

Metrics	HDRUNet	HDRTVNet	FMNet	ICTCPNet
PSNR	20.83	<u>23.78</u>	21.61	21.71
SSIM	0.8667	<u>0.8824</u>	0.8609	0.7961
Metrics	LSNet	PromptIR	RAM-PromptIR	Ours
PSNR	20.36	21.24	21.39	<b>29.12</b>
SSIM	0.8603	0.7946	0.8089	<b>0.9412</b>

Table 3: Comparison of quantitative results on HDRTV1K (unknown degradation benchmark). All methods were trained using all degraded SDR versions from the HDRTV4K dataset. **Bold** and underline indicate the best and second-best results, respectively.

re-initialized as zero-convs to softly inject degradation priors.

The total loss combines pixel-level supervision and contrastive regularization:

$$\mathcal{L}_{\text{total}} = \|I_{\text{HDR}}^{\text{pred}} - I_{\text{HDR}}^{\text{G.T.}}\|_1 + \lambda_l \|I_{\text{DCM}} - I_{\text{HDR}}^{\text{G.T.}}\|_1 + \lambda_{\text{contra}} \mathcal{L}_{\text{contra}} \quad (5)$$

where  $I_{\text{DCM}}$  denotes the output of the final DCM module. We empirically set  $\lambda_l = 0.7$ ,  $\lambda_{\text{contra}} = 0.2$ .

## Experiments

### Experimental Setup

**Dataset.** We evaluate the effectiveness of our proposed framework on the HDRTV4K dataset (Guo et al. 2023), which consists of 3878 4K HDR images encoded in BT.2020/PQ1000. All degradation types included in HDRTV4K are utilized for training and testing. The training set contains 3478 images from HDRTV4K, while the testing set comprises 400 images from HDRTV4K. Additionally, we include 117 images from the HDRTV1K dataset (BT.2020/PQ1000) (Chen et al. 2021c) in the testing set to evaluate the performance on unknown degradations, resulting in a total of 517 test images. The SDR content is configured with BT.709 gamut, gamma 2.2, and a peak brightness of 100 nits.

**Implementation Details.** Each multi-view encoder outputs a global vector and a local feature map with 64 and 16 channels, respectively. The mapping network consists of convolutional layers with 32 channels. We implement our framework using the PyTorch library and adopt the Adam

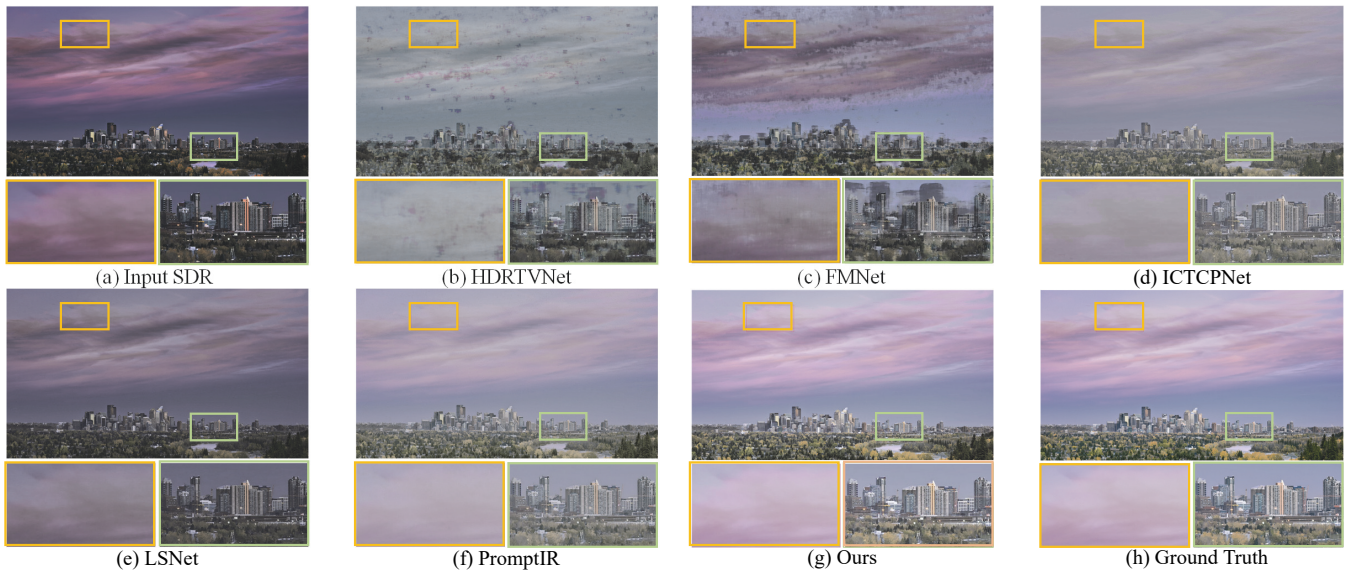


Figure 5: Visual comparison with the state-of-the-art methods. Following `hdrtvdm` (Guo et al. 2023), the HDR images are shown in their original encoding without tone mapping to preserve highlights. Our method delivers visually superior results, with clearer highlight recovery, more vivid colors, and better structure preservation compared to prior methods. Visualization is best viewed on an HDR screen.

optimizer with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.99$ , and a learning rate initialized to  $2 \times 10^{-4}$ . The learning rate is scheduled using a MultiStepLR strategy, with decay milestones at 200K and 400K iterations and a decay factor of 0.5. We employ exponential moving average (EMA) with a decay factor of 0.999 to stabilize training. The model is trained for 400K iterations in total, with the first 40K iterations designated as Stage 1. Mixed-precision training with `bfloat16` is enabled to improve computational efficiency. We use a mini-batch size of 16, and the training is conducted on an NVIDIA RTX 4090 GPU, taking approximately 3 days to complete.

## Experimental Results

**Quantitative Results.** We validate RealRep on the HDRTV4K (known degradation benchmark) and HDRTV1K (unknown degradation benchmark) datasets, comparing it to state-of-the-art SDR-to-HDR methods, including HDRUNet (Chen et al. 2021b), HDRTVNet (Chen et al. 2021c), FMNet (Xu et al. 2022a), ICTCPNet (Huang et al. 2023), LSNet (Guo et al. 2023), PromptIR (Potlapalli et al. 2023), and RAM-PromptIR (Qin et al. 2024). To comprehensively evaluate visual quality, we report both traditional metrics such as PSNR and SSIM, as well as HDR-specific perceptual quality metrics, including HDR-VDP3 and the error-aware HDR metrics EHL, FHLP, EWG, and FWGP proposed by (Guo et al. 2023). Additionally, we evaluate perceptual color difference using  $\Delta E_{ITP}$ , which is reported in the supplementary material.

As shown in Table 1, RealRep outperforms all competing methods on HDRTV4K, demonstrating its ability to handle diverse degradations and produce high-quality HDR reconstructions. This stems from its disentangled multi-view representation learning, which separates luminance and chrominance distributions, enabling robust adaptation to

varying degradation characteristics. On HDRTV1K (Table 3), RealRep excels in generalizing to unknown degradations, outperforming methods like HDRUNet, HDRTVNet, and FMNet, which suffer significant performance drops. PromptIR, while effective in simpler tasks like noise removal, struggles with complex luminance and chrominance transformations. RealRep consistently delivers superior HDR outputs, leveraging its disentangled encoder and domain-adaptive transfer framework.

**Qualitative Results.** Figure 5 compares HDR/WCG images generated by state-of-the-art methods and our proposed RealRep framework. Existing methods, such as HDRTVNet, FMNet, and ICTCPNet, struggle with diverse degradations, causing artifacts like banding, color distortion, and uneven luminance due to their reliance on fixed tone mapping strategies, which limits adaptability and leads to inconsistent performance. For instance, HDRTVNet produces overexposure and banding, FMNet shows luminance and chrominance mismatches, and ICTCPNet renders desaturated colors. In contrast, RealRep combines disentangled multi-view degradation representation learning with a domain-adaptive hierarchical transfer framework to dynamically adapt to varying degradations, producing HDR images that are natural, consistent, and artifact-free. These qualities are evident in Figure 5, where RealRep achieves smoother luminance transitions, better highlight recovery, and accurate color distributions, surpassing other methods in handling unknown degradations and producing results closer to the ground truth.

**Joint Luma-/Chroma-aware Representation Visualization.** To demonstrate the effectiveness of RealRep in constructing the inverse tone mapping (iTM) embedding space and its generalization to unknown degradations, we performed t-SNE visualization of the feature space. As shown

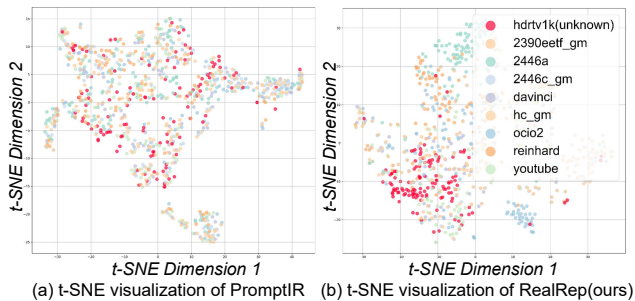


Figure 6: Comparison of t-SNE visualizations for PromptIR (left) and our method (right). Each point represents a frame-level feature projected via PCA (16 components) and t-SNE (perplexity=50). Red points indicate unknown degradations. Our method yields a more compact and well-separated feature space, clearly distinguishing known and unknown degradations.

Configuration	2390eetf_gm	2446a	2446c_gm	davinci	Avg.
Baseline	27.37	27.72	25.96	25.67	25.74
RealRep w/o fusion	32.27	32.76	29.68	27.71	29.83
RealRep w/o $z_{lum}$	32.96	27.88	27.53	29.37	28.96
RealRep w/o $z_{chr}$	32.92	29.48	28.40	29.42	29.58
RealRep w/o $\mathcal{L}_{contra}$	33.73	29.71	28.23	29.77	30.04
RealRep	<b>34.29</b>	<b>34.83</b>	<b>31.10</b>	<b>30.17</b>	<b>31.42</b>

Table 4: Ablation on multi-view representations and multi-view fusion, evaluated by PSNR (dB). The final column reports the average across all eight degradation types. Due to space limitations, only four representative degradations are shown.

in Figure 6, the left panel depicts the prompt representations extracted by PromptIR, while the right panel illustrates the degradation features extracted by RealRep. In PromptIR’s feature space, unknown degradations (e.g., HDRTV1K) are scattered and overlap with other clusters, indicating poor generalization to unseen scenarios. In contrast, RealRep produces compact and well-separated clusters for HDRTV1K, showcasing superior generalization to unknown degradations and a more structured feature space. These findings highlight the strength of RealRep in addressing diverse real-world degradations.

## Ablation Studies

In this section, we perform ablation studies to evaluate the impact of multi-view representations and the degradation-aware controlled mapping module.

### Ablation of Disentangled Multi-View Representations

We evaluate the impact of disentangled multi-view representations through ablations on the luminance representation  $z_{lum}$ , chrominance representation  $z_{chr}$ , their fusion, the contrastive loss  $\mathcal{L}_{contra}$ , and the fusion module. As shown in Table 4, the baseline without multi-view components yields the lowest PSNR (25.74dB), confirming their necessity. Removing  $z_{lum}$  or  $z_{chr}$  reduces PSNR to 28.96dB and 29.58dB,

Configuration	2390eetf_gm	2446a	2446c_gm	davinci	hdrv1k	Avg.
RealRep w/o $z_{global}$	24.56	23.43	28.41	22.41	25.03	25.00
RealRep w/o $z_{local}$	32.95	34.00	30.90	27.33	28.42	30.28
RealRep w/o control	33.24	33.55	30.35	29.49	28.37	30.47
RealRep	34.29	34.83	31.10	30.17	29.12	31.16

Table 5: Ablation on the degradation-aware controlled mapping module, evaluated by PSNR (dB).

respectively, indicating their individual contributions. Without the fusion module, PSNR drops to 29.83dB, showing its role in integrating luminance and chrominance features. Excluding the contrastive loss  $\mathcal{L}_{contra}$  leads to a notable drop from 31.42dB to 30.04dB, highlighting the benefit of contrastive supervision. The full RealRep model achieves the highest PSNR (31.42dB), validating the joint effectiveness of all components in enhancing SDR-to-HDR performance.

### Ablation of Degradation-Domain Aware Controlled Mapping Network

We conduct ablation studies on the proposed degradation-aware controlled module, focusing on its global/local mapping components and the control mechanism. As shown in Table 5, removing either the global or local mapping leads to notable performance drops (25.00dB and 30.28dB average PSNR, respectively), confirming their complementary roles in modeling hierarchical domain knowledge. The control mechanism, composed of a zero-initialized convolution and a two-stage training strategy, also proves essential. Its removal reduces average PSNR from 31.16dB to 30.78dB, and performance on HDRTV1K from 29.12dB to 28.87dB, indicating weakened generalization. Overall, the complete RealRep model achieves the best results, demonstrating the effectiveness of the full degradation-aware controlled mapping design.

### Effectiveness on Unknown TM-Style SDR Inputs

Our method is designed to generalize better to unknown tone-mapped SDR inputs. As shown in Table 6, our method significantly outperforms a classification-based baseline (ResNet+LSNet) on HDRTV1K, achieving higher PSNR and HDR-VDP3 scores.

Method	PSNR	HDR-VDP3
ResNet+LSNet	21.57	9.24
Ours	<b>29.12</b>	<b>9.60</b>

Table 6: Generalization to unknown TM-style SDR inputs on HDRTV1K.

## Concluding Remarks

This paper presents RealRep, a novel framework for converting SDR to HDR, designed to address real-world variations in style and degradation. By using style disentangled representation learning and a degradation-aware mapping, RealRep effectively expands luminance and color ranges. Extensive evaluations show that RealRep outperforms existing methods in visual quality and generalization.

## Acknowledgements

This work was supported in part by China Postdoctoral Science Foundation under Grant 2025M773501; in part by the Fundamental Research Funds for the Central Universities under Grant ZYTS25270.

## References

- Ahn, N.; Kang, B.; and Sohn, K.-A. 2018. *Fast, Accurate, and Lightweight Super-Resolution with Cascading Residual Network*, 256–272.
- Chen, H.; Wang, Y.; Guo, T.; Xu, C.; Deng, Y.; Liu, Z.; Ma, S.; Xu, C.; Xu, C.; and Gao, W. 2021a. Pre-Trained Image Processing Transformer. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Chen, X.; Liu, Y.; Zhang, Z.; Qiao, Y.; and Dong, C. 2021b. HDRUNet: Single Image HDR Reconstruction with Denoising and Dequantization. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 354–363.
- Chen, X.; Zhang, Z.; Ren, J. S. J.; Tian, L.; Qiao, Y.; and Dong, C. 2021c. A New Journey from SDRTV to HDRTV. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 4480–4489.
- Dong, C.; Loy, C. C.; He, K.; and Tang, X. 2014. *Learning a Deep Convolutional Network for Image Super-Resolution*, 184–199.
- Eboli, T.; Sun, J.; and Ponce, J. 2020. End-to-end Interpretable Learning of Non-blind Image Deblurring. *Le Centre pour la Communication Scientifique Directe - HAL - Diderot, Le Centre pour la Communication Scientifique Directe - HAL - Diderot*.
- Guo, C.; Fan, L.; Xue, Z.; and Jiang, X. 2023. Learning a Practical SDR-to-HDRTV Up-conversion using New Dataset and Degradation Models. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 22231–22241.
- He, G.; Xu, K.; Xu, L.; Wu, C.; Sun, M.; Wen, X.; and Tai, Y.-W. 2022. SDRTV-to-HDRTV via hierarchical dynamic context feature mapping. In *Proceedings of the 30th ACM international conference on multimedia*, 2890–2898.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. B. 2019. Momentum Contrast for Unsupervised Visual Representation Learning. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9726–9735.
- Huang, P.; Cao, G.; Zhou, F.; and Qiu, G. 2023. Video Inverse Tone Mapping Network with Luma and Chroma Mapping. *Proceedings of the 31st ACM International Conference on Multimedia*.
- (ITU), I. T. U. 2015. Parameter Values for Ultra-High Definition Television Systems for Production and International Programme Exchange. Technical Report ITU-R BT.2020-2, International Telecommunication Union, Geneva, Switzerland.
- (ITU), I. T. U. 2018. Image Parameter Values for High Dynamic Range Television for Use in Production and International Programme Exchange. Technical Report ITU-R BT.2100-2, International Telecommunication Union, Geneva, Switzerland.
- Kim, J.; Lee, J. K.; and Lee, K. M. 2016. Accurate Image Super-Resolution Using Very Deep Convolutional Networks. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kim, S. Y.; and Kim, M. 2018. A Multi-purpose Convolutional Neural Network for Simultaneous Super-Resolution and High Dynamic Range Image Reconstruction. In *Asian Conference on Computer Vision*.
- Kim, S. Y.; Oh, J.; and Kim, M. 2019a. Deep SR-ITM: Joint Learning of Super-Resolution and Inverse Tone-Mapping for 4K UHD HDR Applications. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 3116–3125.
- Kim, S. Y.; Oh, J.; and Kim, M. 2019b. JSI-GAN: GAN-Based Joint Super-Resolution and Inverse Tone-Mapping with Pixel-Wise Task-Specific Filters for UHD HDR Video. In *AAAI Conference on Artificial Intelligence*.
- Li, B.; Liu, X.; Hu, P.; Wu, Z.; Lv, J.; and Peng, X. 2022. All-In-One Image Restoration for Unknown Corruption. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 17431–17441.
- Li, B.; Peng, X.; Wang, Z.; Xu, J.; and Feng, D. 2017. An All-in-One Network for Dehazing and Beyond. *Cornell University - arXiv, Cornell University - arXiv*.
- Li, B.; Zhao, H.; Wang, W.; Hu, P.; Gou, Y.; and Peng, X. 2025. MaLR: A Locality- and Continuity-Preserving Mamba for Image Restoration. In *IEEE Conference on Computer Vision and Pattern Recognition*. Nashville, TN.
- Lim, B.; Son, S.; Kim, H.; Nah, S.; and Lee, K. M. 2017. Enhanced Deep Residual Networks for Single Image Super-Resolution. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*.
- Ma, S.; Gao, J.; Wang, R.; et al. 2023. Overview of intelligent video coding: from model-based to learning-based approaches. *Visual Intelligence*, 1.
- Potlapalli, V.; Zamir, S. W.; Khan, S.; and Khan, F. 2023. PromptIR: Prompting for All-in-One Image Restoration. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Qin, C.; Wu, R.; Liu, Z.; Lin, X.; Guo, C.-L.; Park, H. H.; and Li, C. 2024. Restore Anything with Masks: Leveraging Mask Image Modeling for Blind All-in-One Image Restoration. In *European Conference on Computer Vision*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, 234–241. Cham: Springer International Publishing. ISBN 978-3-319-24574-4.
- Series, B. 2012. Parameter values for ultra-high definition television systems for production and international programme exchange. Technical Report BT.2020, International Telecommunication Union - Telecommunication Standardization Sector (ITU-T).
- Shao, T.; Zhai, D.; Jiang, J.; and Liu, X. 2022. Hybrid Conditional Deep Inverse Tone Mapping. *Proceedings of the 30th ACM International Conference on Multimedia*.

Standard, S. 2014. High dynamic range electro-optical transfer function of mastering reference displays. Technical Report ST 2084, Society of Motion Picture and Television Engineers (SMPTE).

Tian, C.; Xu, Y.; Li, Z.; Zuo, W.; Fei, L.; and Liu, H. 2020. Attention-guided CNN for image denoising. *Neural Networks*, 117–129.

Xu, G.; Hou, Q.; Zhang, L.; and Cheng, M.-M. 2022a. FM-Net: Frequency-Aware Modulation Network for SDR-to-HDR Translation. *Proceedings of the 30th ACM International Conference on Multimedia*.

Xu, K.; Xu, L.; He, G.; Wu, C.; Ma, Z.; Sun, M.; and Tai, Y.-W. 2022b. Sdrtv-to-hdrtv conversion via spatial-temporal feature fusion. *arXiv preprint arXiv:2211.02297*.

Xu, K.; Xu, L.; He, G.; Yu, W.; and Li, Y. 2023. Towards robust sdrtv-to-hdrtv via dual inverse degradation network. *arXiv e-prints*, arXiv–2307.

Xu, K.; Xu, L.; He, G.; Yu, W.; and Li, Y. 2024a. Beyond Alignment: Blind Video Face Restoration via Parsing-Guided Temporal-Coherent Transformer. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, {IJCAI-24}*, 1489–1497.

Xu, K.; Xu, L.; He, G.; Zhang, Z.; Yu, W.; Wang, S.; Zhou, D.; and Li, Y. 2024b. Beyond Feature Mapping GAP: Integrating Real HDRTV Priors for Superior SDRTV-to-HDRTV Conversion. *arXiv preprint arXiv:2411.10775*.

Yang, W.; Liu, J.; Yang, S.; and Guo, Z. 2019. Scale-Free Single Image Deraining Via Visibility-Enhanced Recurrent Wavelet Learning. *IEEE Transactions on Image Processing*, 2948–2961.

Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding Conditional Control to Text-to-Image Diffusion Models. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 3813–3824.