# GAURA: Generalizable Approach for Unified Restoration and Rendering of Arbitrary Views

Vinayak Gupta<sup>1\*</sup>, Rongali Simhachala Venkata Girish<sup>1\*</sup>, Mukund Varma T<sup>2\*</sup>, Ayush Tewari<sup>3</sup>, Kaushik Mitra<sup>1</sup>

<sup>1</sup>Indian Institute of Technology Madras <sup>2</sup>University of California San Diego <sup>3</sup>Massachusetts Institute of Technology



Fig. 1: Reconstructing 3D scenes from imperfect image captures is highly desirable, e.g., in low light, but is challenging. To address this, we introduce GAURA, a technique designed to render and restore novel views from degraded input views. Unlike previous attempts at this problem, we demonstrate generalization to different scenes and degradation types. In each example, we visualize the imperfect target capture (on the left) and its corresponding clear rendered view (on the right). Please note that the imperfect view from the target viewing angle is not used as input for our method, and we simply visualize the same for simplicity. GAURA faithfully recovers the underlying 3D scene with high geometric accuracy while still managing to generalize across several degradation types.

Abstract. Neural rendering methods can achieve near-photorealistic image synthesis of scenes from posed input images. However, when the images are imperfect, e.g., captured in very low-light conditions, stateof-the-art methods fail to reconstruct high-quality 3D scenes. Recent approaches have tried to address this limitation by modeling various degradation processes in the image formation model; however, this limits

<sup>\*</sup> Authors contributed equally to this work.

them to specific image degradations. In this paper, we propose a generalizable neural rendering method that can perform high-fidelity novel view synthesis under several degradations. Our method, *GAURA*, is learningbased and does not require any test-time scene-specific optimization. It is trained on a synthetic dataset that includes several degradation types. GAURA outperforms state-of-the-art methods on several benchmarks for low-light enhancement, dehazing, deraining, and on-par for motion deblurring. Further, our model can be efficiently fine-tuned to any new incoming degradation using minimal data. We thus demonstrate adaptation results on two unseen degradations, desnowing and removing defocus blur. Code and video results are available at vinayakvg.github.io/GAURA.

# 1 Introduction

Recent progress in Neural Radiance Fields (NeRF) [27] and other implicit representations [4, 8, 28, 31] has enabled the transformation of multi-view captures of real-world scenes into 3D models that produce photo-realistic and consistent renderings from novel viewpoints. However, methods like NeRF still require images captured in perfect conditions, which is often hard to achieve in the real world. Imperfections in the captured images can either arise from challenging scenes, e.g., lack of sufficient illumination, or from the capturing device, e.g., motion or defocus blur from commodity hand-held devices. These imperfections violate the image formation model of neural radiance fields, resulting in poor rendering quality.

Existing methods for novel view synthesis from degraded image captures, typically modify the image formation process by modeling the physical degradation. As examples, Seathru-NeRF [15] incorporates the underwater image formation model into their rendering process, and Deblur-NeRF [23] leverages a kernel at each spatial location to simulate blurring. However, explicitly modeling the physical degradation process results in more challenging inverse problems and limits these approaches to very specific degradations, limiting practical applications of these methods in the real world.

In this paper, we propose a novel method that *learns* to reconstruct scenes under any general *real-world* degradation, enabling the rendering of clean images from arbitrary viewpoints. We take inspiration from state-of-the-art 2D image restoration (IR) methods that leverage large networks trained on abundant paired data, i.e., degraded and their corresponding clean image pairs captured from the real world or generated synthetically [17, 18, 29].

Our learning-based reconstruction method builds on the progress in generalizable NeRF-based scene reconstruction [35, 37, 41]. Unlike per-scene optimization approaches [27], generalizable methods can reconstruct scenes from sparse input images using feed-forward network operations, which can be used to render images from novel viewpoints. We extend the idea of generalization across scenes to generalization across degradations as well. We call our technique GAURA: *Gen*-

 $\mathbf{2}$ 

eralizable Approach for Unified Restoration and Rendering of Arbitrary Views, a model capable of restoring many different degradations.

We base our method on the state-of-the-art generalizable model GNT [35], which leverages two transformers - view transformer for epipolar feature aggregation and ray transformer for learned ray-based rendering. We propose to condition the epipolar feature aggregation and learned rendering on degradationaware latent codes. These codes are learnable parameters that encode discriminative information about different image degradations, allowing the network to dynamically adapt its behavior to different degradations. GAURA casts a ray for each pixel in the target image plane, samples, and project points onto the degraded input views to fetch RGB features based on epipolar correspondences. The view transformer estimates coordinate-wise features from epipolar lines, which are additionally now conditioned on the input degradation type. For rendering a clean novel view, the ray transformer composes these point features to color while simultaneously incorporating the input degradation prior onto its learned rendering process. Our key insight is that our degradation-aware latent codes can interact with the input features and dynamically adjust their representations to adapt the restoration and rendering process for the relevant degradation. The entire pipeline is trained end-to-end on degraded and clean image pairs of several degradation types, synthetically generated from existing multi-view datasets [26, 27, 41].

Even though GAURA is trained on synthetic data, its performance transfers to real-world scenes, where it can simultaneously render and restore novel views from arbitrary viewpoints. Further, our method can be easily tuned to any unseen degradation with very little paired data (as little as eight scenes) by adapting the learned scene and rendering priors. On various rendering and restoration benchmarks like low-light enhancement, dehazing, and motion deblurring, the performance of our generalizable restoration technique is comparable and sometimes even better than methods that use per-scene optimization and have been designed for the specific degradation type. In other degradation types, like deraining and desnowing, we present the first 3D restoration technique to model these 2 corruptions effectively.

In summary, by incorporating learnable degradation latent codes in both the epipolar feature aggregation and rendering steps, GAURA can universally render and restore novel views from degraded input views and demonstrates versatility to several degradation types. Such a data-driven pipeline enables zeroshot inference on any new scene containing real-world degradations from any degradation types seen during training without any extra optimization. Further, the useful priors learned by the trained model can help adapt GAURA to any unseen degradation using a fast and efficient fine-tuning strategy. To the best of our knowledge, our method is the first that does not need to deal with each degradation type separately in a 3D setting.

# 2 Related Works

**2D Image Restoration.** Image Restoration (IR), a long-standing problem in computer vision, attempts to reconstruct a high-quality, clean image from a

degraded observation. The 2D image restoration domain boasts of several benchmark datasets [30,44,46,50], and performant methods for individual corruption types, e.g., for deblurring [6,19], low-light enhancement [24,43], dehazing [16,47] and many others [1,3,34]. In contrast to these degradation-specific techniques, a recent line of work [17,21,29] tries to build a single unified model that can generalize to several input imperfections and has showcased promising results. Naively applying these techniques to 3D data, i.e., multi-view images, yields inconsistent predictions as they are originally only intended for single-image restoration.

Novel View Synthesis. A pioneering work, Neural Radiance Field or NeRF [27] synthesizes highly realistic and consistent novel views by fitting each scene as a continuous 5D radiance field parameterized by an MLP. Subsequent works have improved NeRF's rendering quality [2, 36, 39, 45]. Recently, there have been efforts to generalize NeRFs to arbitrary scenes by learning priors from large-scale multi-view image datasets [9, 32, 41]. However, all these methods are rarely applied to input images captured in imperfect conditions (e.g., lack of sufficient illumination, motion blur, etc), thereby violating the image formation model and resulting in poor novel view synthesis quality. In this paper, we attempt to reconstruct a 3D scene under any real-world degradation and render clean images from arbitrary viewpoints without any additional optimization.

**3D** Restoration. Some recent efforts have attempted to reconstruct a radiance field using imperfect image captures, e.g., from blurry [23,40], and low-light corruptions [7,25,38]. These methods explicitly simulate the degradation process by modifying the rendering process in NeRF. This requires significant creativity, is challenging for more complex imperfections, and limits approaches to a single degradation type, affecting the practical applications of these techniques in the real world. In contrast, our method GAURA implicitly encodes the degradation process into learned modules, enabling generalization - to both scenes and corruption types.

# 3 Preliminary: Generalizable Novel View Synthesis

We formulate our technique around Generalizable Novel View Synthesis (GNVS) to render clean images from imperfect input views and introduce the background below. Given N calibrated input (or source) views with known pose information  $\{I_i, P_i\}_{i=1}^N$ , the goal of GNVS is to synthesize a target novel view  $I_T$  based on these source views, even for scenes not observed during training, thereby achieving generalizability. These methods first extract deep convolutional features  $F_i = \mathcal{F}_{conv}(I_i)$  from each input view. To render a target view, several rays are cast into the scene, and K points  $\{r(t_i)\}_{i=1}^K$  are sampled along each ray  $r(t) = \mathbf{o} + t\mathbf{d}$ , where  $\mathbf{o}$  is the camera center and  $\mathbf{d}$  is the ray's unit direction vector. Each point is then projected onto source view  $I_i$  using projection function  $\Pi_i$ , and the nearest feature in the image plane is queried. These multi-view features are aggregated into point feature f as follows:

$$\boldsymbol{f}(t) = \mathcal{F}_{view}(\{\boldsymbol{F}_{\Pi_i(\boldsymbol{r}(t))}\}_{i=1}^N)$$
(1)

where  $\mathcal{F}_{view}$  is a permutation-invariant aggregation mapping that learns to be occlusion-aware to combine epipolar features. Finally, these point features are accumulated into color c using aggregation function  $\mathcal{F}_{point}$ :

$$\boldsymbol{c}(r) = \mathcal{F}_{point}(\{\boldsymbol{f}(t_i)\}_{i=1}^K)$$
(2)

Early works like IBRNet [41] adopt simple MLPs and an explicit volume rendering as aggregation functions  $\mathcal{F}_{view}$  and  $\mathcal{F}_{point}$  respectively, while more recent techniques [12,32,35] use complex transformer mechanisms due to their stronger representation capacity. In this work, we adapt GNT [35] that specifically employs a view transformer to aggregate epipolar features and a ray transformer for "attention-based" rendering.



Fig. 2: Overview of GAURA: 1) Given multi-view images of a scene captured in imperfect conditions, we first extract deep convolutional features for each input view, 2) Using the source view features, we estimate the target clean rendered view via an extended epipolar-based rendering pipeline conditioned on the input imperfection type, 3) By supervising this pipeline end-to-end on paired synthetic data, our degradation-specific latent codes encode discriminative information about each imperfection type and dynamically adapts the rendering modules based on the input corruption type. Our method can directly generalize to any new scene containing real-world degradations during inference.

# 4 Method

**Overview.** We introduce *GAURA*, a novel framework to perform generalizable novel view synthesis from imperfect input captures, as illustrated in Fig. 2. Formally, given N calibrated multi-view images of a scene with known camera poses  $\{I_i, P_i\}_{i=1}^N$  captured in an imperfect environment containing a given degradation type  $D^1$ , our goal is to synthesize clear novel views from arbitrary angles and simultaneously generalize to different scenes and corruption types. Our method builds upon existing GNVS technique [35] but conditions the scene representation and rendering processes upon some learned degradation-aware

<sup>&</sup>lt;sup>1</sup> we assume a single degradation type for simplicity, while in practice it could be a combination of multiple (see Sec. 5.5)

latent codes (Sec. 4.1, 4.2). These latent codes encode discriminative information about each degradation type and dynamically adapt the remaining parts of the network based on the input imperfection (Sec. 4.3). Such a pipeline can be trained end-to-end on manually altered scenes and directly transferred to any real scene without any additional optimization (Sec. 4.4). Our degradation-aware latents are fully disentangled from the remaining parts of the network, implying that they can be easily expanded for any new imperfection type and therefore be fine-tuned with minimal training (Sec. 4.4). This further underscores the practical applicability of our technique, i.e., GAURA acts as an *all-in-one unified restoration and rendering pipeline*.

## 4.1 Degradation-aware Latent Module

Existing work for novel view synthesis and restoration often explicitly models the physical degradation process into the rendering equation [7, 23]. Although the image formation process in imperfect environments is well-studied [13, 14, 33], it still requires significant creativity to tailor the rendering process for each corruption type and forces no generalization to different degradation types. In this work, we take an alternate route and attempt to implicitly model the image formation process under several corruption types (assuming M different corruption types) into learnable parameters. A naive strategy would be to create M different network clones and individually allow each to model the rendering process under different degradation types, but this is inefficient. Moreover, the image formation processes under different imperfect environments sometimes share several similarities [5]. Motivated by this observation, we encode degradation-specific information into M latent codes  $\{L_i\}_{i=1}^M$  that interact with the remaining parts of the network whose parameters are shared irrespective of the input imperfection type. Drawing inspiration from HyperNetworks [10], these latent codes are mapped to weight W of an MLP that performs a degradation-specific transformation on the input data X. This module, dubbed Degradation-aware Latent Module (or DLM) can be formally described as:

$$DLM(\boldsymbol{X}, \boldsymbol{D}) = W \cdot \boldsymbol{X}, \text{ where } W = \mathcal{F}_{latent}(\boldsymbol{L}_{\boldsymbol{D}}|\{\boldsymbol{L}_i\}_{i=1}^M)$$
 (3)

where  $\mathcal{F}_{\text{latent}}$  denotes a fully-connected layer, and  $L_D$  represents the latent code for the degradation type D. Such a technique is efficient, as the latent codes contribute to a small portion of extra parameters while still ensuring sufficient capacity to represent individual imperfection types.

### 4.2 Adaptive Residual Module

Although we use a set of degradation-aware latents to encode imperfectionspecific information, they are independent of the actual input captures. We argue that several variations exist within each corruption type, and the previous static approach might yield rather sub-optimal results. Therefore, we propose to add a residual feature S obtained from the input view closest to the target view to be rendered and reformulate Eq. 3 to:

$$DLM_{w/residue}(\boldsymbol{X}, \boldsymbol{D}) = DLM(\boldsymbol{X}, \boldsymbol{D}) + \boldsymbol{S}, \text{ where } \boldsymbol{S} = pool(\mathcal{F}_{residue}(\boldsymbol{I}_{nearest}))$$
  
(4)

where pool(.) denotes the global average pooling operation,  $\mathcal{F}_{\text{residue}}$  denotes a tiny convolutional network that encodes the input view closest to the target viewing angle  $I_{\text{nearest}}$ . In the remainder of this paper, we refer to the DLM module with *Adaptive Residual Module* (ARM) as simply DLM for the sake of simplicity unless stated otherwise.

#### 4.3 Degradation-aware Generalizable Novel View Synthesis

We formulate our method around the recent SOTA GNVS technique - GNT [35], that follows a three-stage pipeline as highlighted in Sec. 3. More specifically, GNT leverages a UNet-based feature extractor ( $\mathcal{F}_{conv}$ ) to obtain geometry and appearance features from input images that are aggregated into point features using a view transformer ( $\mathcal{F}_{view}$ ) and then accumulated along the ray to estimate ray color using the ray transformer ( $\mathcal{F}_{point}$ ). This section details the different strategies employed to augment these components with degradation-specific priors that constitute *GAURA*.

The first key challenge is extracting meaningful features from imperfect input views that can later be used to embed a scene representation and render a clean novel view. Therefore, we incorporate the DLM blocks in the feature UNet, specifically at each decoder level, before increasing the spatial resolution of the intermediate feature maps. At each level, the DLM implicitly enriches the input features with information about the input degradation for guided recovery. To obtain point feature F(t) at the t-th point sampled along the ray, we aggregate features along the epipolar line using the view transformer. However, this is not trivial since the features on the image plane are not necessarily multi-view consistent due to imperfections in the input captures and need to be rectified. Following our recurrent theme of implicitly capturing such imperfection-specific aggregation processes, we replace the vanilla MLP query, key, and value feature mapping with DLM modules and, by doing so, incorporate a degradation prior onto the scene representation phase. Finally, the ray transformer accumulates these point features along the ray and then decodes them into color. By viewing the attention scores as blending weights and values as color features, such a learned aggregation represents a "generalized" volume rendering process. Therefore, we only replace the value feature mapping with DLM modules in the ray transformer since the underlying scene geometry represented as attention scores obtained from the query and key features are independent of input imperfection type.

$$\boldsymbol{F}_{i} = \mathcal{F}_{dec}(\mathcal{F}_{enc}(\boldsymbol{I}_{i}), DLM(\boldsymbol{I}_{i}, D)), \text{ where } \mathcal{F}_{dec} \circ \mathcal{F}_{enc} := \mathcal{F}_{conv}$$
$$\boldsymbol{f}(t) = \mathcal{F}_{view}(\{DLM(\boldsymbol{G})_{\boldsymbol{q}}, DLM(\boldsymbol{G})_{\boldsymbol{k}}, DLM(\boldsymbol{G})_{\boldsymbol{v}}\}), \text{ where } \boldsymbol{G} = \{\boldsymbol{F}_{\Pi_{i}(\boldsymbol{r}(t))}\}_{i=1}^{N}$$
$$\boldsymbol{c}(r) = \mathcal{F}_{point}(\{\boldsymbol{H}_{\boldsymbol{q}}, \boldsymbol{H}_{\boldsymbol{k}}, DLM(\boldsymbol{H})_{\boldsymbol{v}}\}), \text{ where } \boldsymbol{H} = \{\boldsymbol{f}(t_{i})\}_{i=1}^{K}$$
(5)

Since  $\mathcal{F}_{\text{view}}$ ,  $\mathcal{F}_{\text{point}}$  both represent transformer mechanisms, the operands within each explicitly indicate the query (q), key (k), value (v) features with their respective subscripts. We defer more implementation details to the supplementary.

### 4.4 Training and Inference

The entire pipeline can be end-to-end supervised across multiple scenes with multi-view images. At each training iteration, we sample from the M degradation types, apply the same to the input images (with reasonable intra-corruption variability), and supervise the rendered view with its corresponding clean ground truth image.

$$\mathcal{L} = \mathcal{L}_{\text{MSE}}(I_{target}, I_{target}) + \mathcal{L}_{\text{LPIPS}}(I_{target}, I_{target})$$
(6)

where  $\mathcal{L}_{\text{MSE}}$ ,  $\mathcal{L}_{\text{LPIPS}}$  denote the mean-squared error and LPIPS feature distance [49] criterion respectively, and  $I_{target}$  represents the rendered target view and its corresponding ground truth view  $\hat{I}_{target}$ . Given sufficient training data, the degradation-aware latent codes must encode specific information regarding each imperfection type and guide the recovery process. We observe that by simply training on scenes containing synthetic realizations of several corruption types, GAURA learns to generalize to scenes containing real-world imperfections. Similar results have been seen in 2D image restoration [20, 22, 42, 48] and confirms the hypothesis in the 3D domain. More significantly, our DLM blocks are plugand-play and can be easily expanded to accommodate new imperfection types. Being disentangled from other parts of the network, the new latent code can be efficiently fine-tuned with limited training (as little as 6-8 synthetic scenes) by re-purposing the priors learned during pretraining.

## 5 Results

To demonstrate the effectiveness of GAURA, we evaluate several image restoration tasks, including low-light enhancement, motion blur removal, image dehazing, and deraining. The useful priors learned by the trained model can help adapt GAURA to any unseen degradation using an efficient finetuning strategy (discussed in Sec. 5.4), and we verify the same in two representative image restoration tasks: defocus blur removal and image desnowing. In some of these restoration tasks, like image desnowing and deraining, we present the first 3D extension with no comparable previous 3D method. Our method is never finetuned on each scene for all the experiments listed below, and we only perform inferences using a trained model.

## 5.1 Implementation Details

We train our entire pipeline end-to-end on datasets of multi-view posed images. Specifically, we use the training data from IBRNet [41] and LLFF [26] and synthetically alter the original clean scene to contain imperfections. At every training iteration, we sample from one of 4 different corruption types (i.e., lowlight, motion blur, haze, or rain) and apply it to the input images. To simulate real-world data, we incorporate variations within each corruption type by randomly sampling the strength of the corruption. This helps the network generalize better (more implementation details in the supplementary materials). We follow the same input view sampling strategy from IBRNet [41], and select between 8-12 source views during training and a fixed number of 10 views during inference on any new scene. Rather than training from scratch, we reuse the pretrained checkpoint from GNT [35] and add our proposed degradation-aware modules to the necessary parts of the network. Our method is optimized to render clean images using the Adam optimizer with an initial learning rate of  $5 \times 10^{-4}$ , decayed over the course of 400,000 training steps.

#### 5.2 Comparisons

Since no previous works deal with generalizable restoration and view synthesis, we devise several reasonable baselines for comparison.

**NeRF**-Restore. refers to a Vanilla NeRF [27] trained on images restored using a single-image restoration technique, specialized for the given degradation type. The restoration technique in the study is different for different input imperfections. This baseline does not generalize to scenes or corruption types.

**GNT**-**Restore.** uses a similar strategy as Restore-NeRF but replaces the Vanilla NeRF with a generalizable novel view synthesis technique GNT [35]. Such a base-line can generalize across scenes but is specific to a given restoration type.

GNT-(All-in-One) Restore. Recent single-image restoration methods are All-in-One i.e., can handle multiple degradation types. We leverage state-of-the-art all-in-one methods [11, 17, 21, 29] followed by GNT for novel view synthesis. Such a baseline is neither scene nor corruption-specific and is, therefore, a more suitable comparison to our method.

**3D** Restore. We also compare against specialized NeRF techniques for the given degradation type, wherever available. Similar to NeRF-Restore, this base-line is both scene and degradation-specific.

The best-performing method choices for each baseline across individual degradation types are discussed in the supplementary materials.

## 5.3 Generalization to Multiple Degradations

Low-light Enhancement. Given multi-view images of a scene captured in low light, we aim to render images from arbitrary viewpoints as if they were viewed with sufficient illumination. We evaluate our method and the baselines discussed in Sec. 5.2 on the dataset from AlethNeRF [7], containing in total 5 scenes and discuss results in Table 1. In the case of scene-specific techniques, we optimize a radiance field for each scene (i.e., 2D Restore-NeRF and 3D-Restore), while the generalizable methods, including ours, are directly evaluated

**Table 1:** Quantitative results on scenes containing *real-world* degradations - specifically low-light enhancement, motion blur removal, and dehazing. The **best** scores and second best scores are highlighted with their respective colors only amongst the generalizable methods.

Models	Generalize to		Low-Light			Motion Blur				Haze		
	Scene	Corr.	$\mathrm{PSNR}\uparrow$	$\mathrm{SSIM}\uparrow$	$\mathrm{LPIPS}{\downarrow}$	$ PSNR\uparrow$	$\mathrm{SSIM}\uparrow$	$\mathrm{LPIPS}{\downarrow}$	$ PSNR\uparrow$	$\mathrm{SSIM}\uparrow$	LPIPS↓	
NeRF-Restore	X	X	15.42	0.702	0.393	23.27	0.717	0.331	13.87	0.690	0.333	
3D Restore	X	X	17.64	0.736	0.415	25.65	0.758	0.181	-	-	-	
GNT–Restore	1	X	16.36	0.730	0.344	21.97	0.705	0.392	14.16	0.678	0.279	
GNT-(All-in-one) Restore	1	1	17.90	0.573	0.354	20.88	0.632	0.41	16.68	0.729	0.300	
Ours	1	1	19.91	0.738	0.352	22.12	0.712	0.346	16.82	0.759	0.288	

 Table 2: Quantitative results on the LLFF-Corrupted benchmark, containing synthetically altered scenes of several imperfection types.

Models	Low-Light			Motion Blur			Haze			Rain		
hiodolo	$\mathrm{PSNR}\uparrow$	$\mathrm{SSIM}\uparrow$	$\mathrm{LPIPS}{\downarrow}$	$ PSNR\uparrow$	$\mathrm{SSIM}\uparrow$	$\mathrm{LPIPS}{\downarrow}$	$ PSNR\uparrow$	$\mathrm{SSIM}\uparrow$	$\mathrm{LPIPS}{\downarrow}$	$ PSNR\uparrow$	$\mathrm{SSIM}\uparrow$	LPIPS↓
GNT–Airnet [17]	18.2	0.590	0.405	21.08	0.604	0.410	14.55	0.501	0.400	20.71	0.612	0.319
GNT-PromptIR [29]	17.67	0.571	0.389	21.01	0.604	0.378	15.81	0.514	0.361	20.73	0.612	0.276
GNT–DA-CLIP [21]	12.458	0.527	0.473	21.96	0.634	0.412	8.36	0.443	0.563	20.24	0.608	0.434
GNT-AutoDIR [11]	10.48	0.453	0.552	21.85	0.628	0.427	12.59	0.511	0.538	19.25	0.558	0.501
Ours	21.98	0.598	0.464	22.61	0.680	0.349	18.95	0.642	0.383	22.61	0.710	0.342

on the benchmark dataset, not observed during training. GAURA, although neither scene nor corruption-specific, manages to outperform other comparison methods significantly, even 3D restoration methods tuned specifically for lowlight enhancement [7] (3D Restore in Table 1) This indicates the effectiveness of the learned rendering process compared to explicitly modeling the degradation in the image formation process. In Fig. 3 (row 1), we observe that GAURA can faithfully recover the scene much closer to the ground truth from poorly lit scenes.

**Deblurring.** We evaluate our method on the task of blur-removal, particularly motion blur on the scenes from the DeblurNeRF dataset [23], and average results across 10 different scenes in Table 1. We observe that our method performs on par with methods that require any optimization and significantly better than a fully generalizable baseline [21] indicated as All-in-one-GNT in Table 1 (~10%  $\uparrow$  PSNR). We wish to reiterate that although the 3D Restore technique [23] performs better than GAURA, it requires per-scene optimization and is specific to the task of deblurring, while our method is generalizable.

Fig. 3 (row 2) visualizes the rendered novel views, and our method can recover fine details in the scene (e.g., woven patterns in the ball) with superior quality compared to other baselines.

**Haze Removal.** In this task, we look at novel view synthesis from input images captured in a hazy environment. We compare GAURA against different baselines on the dataset from [50] containing 5 scenes, initially not intended for novel view synthesis. We run COLMAP on the video frames given in the dataset



Fig. 3: Qualitative results on three restoration tasks: low-light enhancement, motion blur removal, and dehazing on *real-world* datasets. In the first row, our method successfully recovers the underlying 3D scene from poorly lit images much closer to the ground truth than the baselines. In the second row, our method can reconstruct fine details (woven patterns on the ball) with higher fidelity while in the third row, GAURA successfully removes haze from the synthesized views and can accurately match the

GNT-(All-in-One) Restore

Ours

Ground Truth

Haze

Corrupted View

GNT-Restore

colors on the palette more closely to the ground truth color.

to extract camera poses and test on every 8th image in the dataset. Unlike the previous two degradations, there are no existing 3D restoration techniques, perhaps indicating the complexity of explicitly modeling the degradation process and further supporting the practical applicability of our method. From Table 1, we see that our method showcases better quantitative results when compared to other baselines. From Fig. 3 (row 3), we observe that *GAURA* successfully removes haze from the synthesized views and accurately matches the color palette more closely to the ground truth color.

**Results on LLFF-Corrupted.** In addition to the results discussed above on existing benchmarks, we manually alter all eight scenes from LLFF [26] into several corrupted versions, containing one low-light, motion blur, haze, or rain imperfections. We follow the same train-test split from the original LLFF dataset and evaluate GAURA's generalization capabilities on this benchmark (dubbed LLFF-Corrupted). We compare against a similar generalizable baseline as ours, an All-in-One restoration technique [11,17,21,29] followed by generalizable rendering [35], and retrain both restoration and rendering components on the same synthetic training dataset as GAURA to ensure a fair comparison. Table 2 shows that GAURA significantly outperforms all baselines across all imperfection types. This indicates the capacity of our model to implicitly encode different degradation processes into neural networks, in this case, transformers, and that such a learned renderer is performing well. We visualize qualitative results in the supplementary materials.



Fig. 4: Qualitative results of GAURA on additional restoration tasks. Our pretrained model can remove imperfections caused due to rain particles (row 1). With limited finetuning, GAURA can learn to successfully restore any new degradation type (unseen during training). For representation purposes, we visualize results on snow (row 2) and defocus blur (row 3), and our method outperforms degradation-specific restoration baselines.

## 5.4 Fine-tuning to New Degradations

As discussed in Sec. 4.4, our method allows for efficient fine-tuning to any new degradation type and we provide more details below. We expand the Degradation Aware Latent Block for an incoming corruption type to contain an additional embedding for the given input imperfection,

**Table 3:** Quantitative results upon fine-tuning pretrained GAURA on two unseen restoration tasks - snow removal and defocus deblurring.

Models		Snow		Defocus				
	$\mathrm{PSNR}\uparrow$	$\mathrm{SSIM}\uparrow$	$\mathrm{LPIPS}{\downarrow}$	$ PSNR\uparrow$	$\mathrm{SSIM}\uparrow$	LPIPS.		
Vanilla-GNT [35]	21.96	0.688	0.295	20.95	0.754	0.283		
GNT–Restore	20.24	0.608	0.434	21.13	0.762	0.257		
Ours	22.61	0.710	0.342	21.34	0.768	0.251		

effectively adding less than 5% parameters to the GAURA model. By freezing the remaining parameters, one can efficiently optimize GAURA on the new degradation with a limited number of scenes and training iterations. Simply put, our architecture design only requires a latent code to be learned that encodes degradation-specific information and can modulate the remaining parts of the network based on the input imperfection. Note that in this case, we do not utilize paired real-world data and rely only on synthetically corrupted data for supervision. To verify the same, we select two representative restoration tasks: defocus blur removal and desnowing and finetune on 8 scenes from LLFF's training dataset [26], manually altered to contain these imperfections. For evaluation, we benchmark on the scenes from DeblurNeRF [23] (for defocus blur removal), synthetically altered scenes from LLFF containing snow imper-



Fig. 5: Qualitative results on scenes containing multiple imperfections in the input image captures. Upon interpolating the latent codes corresponding to the individual imperfection types, GAURA can "combine" multiple image formation processes and recover the original clean scene effectively.

fections (for desnowing), and report average scores in Table 3. In addition to the baseline GNT–Restore (see Sec. 5.2), we compare against a Vanilla GNT model trained from scratch on the given degradation type. We observe that, indeed, the pre-trained priors learned by GAURA can be repurposed to learn any new degradation type with limited fine-tuning (in practice, as little as 10,000 steps). Due to the absence of benchmarks with ground truth clean views, we simply visualize our rendered outputs on real scenes containing snow and blur imperfections in Fig. 4 (rows 2 and 3) and our method faithfully reconstructs the clean novel view with high quality when compared to baselines. Notably, since the additional latent codes are only optimized during fine-tuning, GAURA can retain its original performance on the previous degradation types.



Fig. 6: Qualitative comparisons for ablation study.

# 5.5 Extension to Multiple Degradations

In a practical scenario, there are often cases where multiple imperfections are present in the input image captures. Explicitly modeling such a complex image formation process is challenging, making it difficult to naively extend existing NeRF restoration techniques to such a scenario. However, GAURA simply embeds degradation-specific priors onto individual latent codes that can be "interpolated" to combine their individual image formation processes. Formally, given two degradation types  $D_1$ , and  $D_2$  present in the input images, we query the

DLM blocks with a new combined latent code  $\mathbf{L} = \alpha \times \mathbf{L}_{D_1} + (1 - \alpha) \times \mathbf{L}_{D_2}$ , where  $\alpha$  is a scaling term. We visualize qualitative results in Fig. 5, where the input images of a scene are manually altered to contain more than one imperfection type. By manually tweaking  $\alpha$ , our pretrained GAURA can effectively remove multiple imperfections and synthesize clean novel views.

## 5.6 Ablation Studies

We conduct the following ablations to validate our architectural design and provide quantitative results in Table 4. In each case, we train the model following the details mentioned in Sec. 5.1, and then evaluate performance on a single representative restoration task, i.e. low-light enhancement on the AlethNeRF dataset [7]. First, we train the original GNT [35], as-is without any degradationaware modules to automatically render and restore novel views (indicated as Vanilla GNT). Without explicitly encoding degradation-specific information, GNT fails to restore novel views reasonably. By inserting Degradation-aware Latent Blocks (from Sec. 4.1) into several parts of the network ( $\mathcal{F}_{conv}$ ,  $\mathcal{F}_{view}$ ,  $\mathcal{F}_{point}$ ), GAURA can now learn to implicitly embed degradation-specific information into its learned rendering process. Finally, our Adaptive Residual block captures variations within each degradation type that can further improve performance (note the color differences in the case w/o ARM in Fig. 5.4).

# 6 Discussion

Limitations and Future Work. While GAURA can successfully perform novel view synthesis using input images captured in degraded environments, prior information about the imperfection type is still required. While this is not difficult to obtain, an ideal pipeline should be able to automatically restore any corruption type observed in the input images (formally called blind image restoration). Our method relies on epipolar-based rendering techniques, and therefore, its potential is capped by its flaws; for example, our method may not handle sparse 360-degree scenes or objects with complex light transport where epipolar geometry no longer holds. The recent advancements in 3D reconstruction use an explicit representation, e.g., 3D Gaussians, that can encode high-frequency details and enable faster rendering. Therefore, incorporating our degradation-aware modules into such a pipeline could be an interesting direction. Lastly, (and more importantly) the need for benchmarks to evaluate methods on different real-world restoration tasks is glaring, and future work must focus on this direction.

**Conclusion.** We propose GAURA, a pipeline that combines degradationaware modules with epipolar-based rendering to perform novel view synthesis from imperfect input images, showcasing generalization to any scene and corruption type. Our method demonstrates that several of the inductive biases that are necessary for rendering and restoration (e.g. explicit modeling of physical degradation process, hard-coded rendering equation) can be implicitly encoded into neural networks. We believe this insight could open up new possibilities for a universal 3D restoration and rendering model.

15

# References

- Aharon, S., Ben-Artzi, G.: Hypernetwork-based adaptive image restoration. In: ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 1–5. IEEE (2023) 4
- Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mipnerf 360: Unbounded anti-aliased neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5470– 5479 (2022) 4
- Chen, C., Li, H.: Robust representation learning with feedback for single image deraining. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7742–7751 (2021) 4
- Chen, Z., Zhang, H.: Learning implicit fields for generative shape modeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5939–5948 (2019) 2
- Cheng, Y., Shao, M., Wan, Y., Wang, C., Zuo, W.: Drm-ir: Task-adaptive deep unfolding network for all-in-one image restoration. arXiv preprint arXiv:2307.07688 (2023) 6
- Cho, S.J., Ji, S.W., Hong, J.P., Jung, S.W., Ko, S.J.: Rethinking coarse-to-fine approach in single image deblurring. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 4641–4650 (2021) 4
- Cui, Z., Gu, L., Sun, X., Ma, X., Qiao, Y., Harada, T.: Aleth-nerf: Illumination adaptive nerf with concealing field assumption. In: Proceedings of the AAAI Conference on Artificial Intelligence (2024) 4, 6, 9, 10, 14
- Genova, K., Cole, F., Sud, A., Sarna, A., Funkhouser, T.: Local deep implicit functions for 3d shape. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4857–4866 (2020) 2
- Gupta, V., Goel, R., Dhawal, S., Narayanan, P.: Gsn: Generalisable segmentation in neural radiance field. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 2013–2021 (2024) 4
- 10. Ha, D., Dai, A., Le, Q.V.: Hypernetworks. arXiv preprint arXiv:1609.09106 (2016) $_{6}$
- 11. Jiang, Y., Zhang, Z., Xue, T., Gu, J.: Autodir: Automatic all-in-one image restoration with latent diffusion. arXiv preprint arXiv:2310.10123 (2023) 9, 10, 11
- Johari, M.M., Lepoittevin, Y., Fleuret, F.: Geonerf: Generalizing nerf with geometry priors. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18365–18375 (2022) 5
- Koschmieder, H.: Theorie der horizontalen sichtweite. Beitrage zur Physik der freien Atmosphare pp. 33–53 (1924)
- Land, E.H.: The retinex theory of color vision. Scientific american 237(6), 108–129 (1977) 6
- Levy, D., Peleg, A., Pearl, N., Rosenbaum, D., Akkaynak, D., Korman, S., Treibitz, T.: Seathru-nerf: Neural radiance fields in scattering media. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 56–65 (2023) 2
- Li, B., Gou, Y., Gu, S., Liu, J.Z., Zhou, J.T., Peng, X.: You only look yourself: Unsupervised and untrained single image dehazing neural network. International Journal of Computer Vision 129, 1754–1767 (2021) 4
- Li, B., Liu, X., Hu, P., Wu, Z., Lv, J., Peng, X.: All-in-one image restoration for unknown corruption. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17452–17462 (2022) 2, 4, 9, 10, 11

- 16 V. Gupta, R. Girish, M. Varma T et al.
- Li, R., Tan, R.T., Cheong, L.F.: All in one bad weather removal using architectural search. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3175–3185 (2020) 2
- Liu, R., He, Y., Cheng, S., Fan, X., Luo, Z.: Learning collaborative generation correction modules for blind image deblurring and beyond. In: Proceedings of the 26th ACM international conference on Multimedia. pp. 1921–1929 (2018) 4
- Lou, Z., Xu, H., Mu, F., Liu, Y., Zhang, X., Shang, L., Li, J., Guan, B., Li, Y., Hu, Y.H.: Simhaze: game engine simulated data for real-world dehazing. arXiv preprint arXiv:2305.16481 (2023) 8
- Luo, Z., Gustafsson, F.K., Zhao, Z., Sjölund, J., Schön, T.B.: Controlling visionlanguage models for universal image restoration. arXiv preprint arXiv:2310.01018 (2023) 4, 9, 10, 11
- Lv, F., Li, Y., Lu, F.: Attention guided low-light image enhancement with a large scale low-light simulation dataset. International Journal of Computer Vision 129(7), 2175–2193 (2021) 8
- Ma, L., Li, X., Liao, J., Zhang, Q., Wang, X., Wang, J., Sander, P.V.: Deblurnerf: Neural radiance fields from blurry images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12861–12870 (2022) 2, 4, 6, 10, 12
- Ma, L., Ma, T., Liu, R., Fan, X., Luo, Z.: Toward fast, flexible, and robust low-light image enhancement. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5637–5646 (2022) 4
- Mildenhall, B., Hedman, P., Martin-Brualla, R., Srinivasan, P.P., Barron, J.T.: Nerf in the dark: High dynamic range view synthesis from noisy raw images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16190–16199 (June 2022) 4
- Mildenhall, B., Srinivasan, P.P., Ortiz-Cayon, R., Kalantari, N.K., Ramamoorthi, R., Ng, R., Kar, A.: Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. ACM Transactions on Graphics (TOG) (2019) 3, 8, 11, 12
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I. pp. 405–421 (2020) 2, 3, 4, 9
- Niemeyer, M., Mescheder, L., Oechsle, M., Geiger, A.: Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3504–3515 (2020) 2
- Potlapalli, V., Zamir, S.W., Khan, S., Khan, F.S.: Promptir: Prompting for all-inone blind image restoration. arXiv preprint arXiv:2306.13090 (2023) 2, 4, 9, 10, 11
- Rim, J., Lee, H., Won, J., Cho, S.: Real-world blur dataset for learning and benchmarking deblurring algorithms. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020) 4
- Sitzmann, V., Zollhöfer, M., Wetzstein, G.: Scene representation networks: Continuous 3d-structure-aware neural scene representations. Advances in Neural Information Processing Systems 32 (2019) 2
- Suhail, M., Esteves, C., Sigal, L., Makadia, A.: Generalizable patch-based neural rendering. In: European Conference on Computer Vision. pp. 156–174. Springer (2022) 4, 5

- Sun, L., Cho, S., Wang, J., Hays, J.: Edge-based blur kernel estimation using patch priors. In: IEEE international conference on computational photography (ICCP). pp. 1–8. IEEE (2013) 6
- Tian, C., Xu, Y., Zuo, W.: Image denoising using deep cnn with batch renormalization. Neural Networks 121, 461–473 (2020) 4
- 35. Varma, M., Wang, P., Chen, X., Chen, T., Venugopalan, S., Wang, Z.: Is attention all that nerf needs? In: The Eleventh International Conference on Learning Representations (2023) 2, 3, 5, 7, 9, 11, 12, 14
- Verbin, D., Hedman, P., Mildenhall, B., Zickler, T., Barron, J.T., Srinivasan, P.P.: Ref-nerf: Structured view-dependent appearance for neural radiance fields. arXiv preprint arXiv:2112.03907 (2021) 4
- 37. Wang, D., Cui, X., Salcudean, S., Wang, Z.J.: Generalizable neural radiance fields for novel view synthesis with transformer. arXiv preprint arXiv:2206.05375 (2022) 2
- Wang, H., Xu, X., Xu, K., Lau, R.W.: Lighting up nerf via unsupervised decomposition and enhancement. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 12632–12641 (2023) 4
- Wang, P., Liu, L., Liu, Y., Theobalt, C., Komura, T., Wang, W.: Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. arXiv preprint arXiv:2106.10689 (2021) 4
- Wang, P., Zhao, L., Ma, R., Liu, P.: BAD-NeRF: Bundle Adjusted Deblur Neural Radiance Fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4170–4179 (June 2023) 4
- 41. Wang, Q., Wang, Z., Genova, K., Srinivasan, P.P., Zhou, H., Barron, J.T., Martin-Brualla, R., Snavely, N., Funkhouser, T.: Ibrnet: Learning multi-view image-based rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4690–4699 (2021) 2, 3, 4, 5, 8, 9
- Wang, Y., Ma, C., Liu, J.: Removing rain streaks via task transfer learning. arXiv preprint arXiv:2208.13133 (2022) 8
- 43. Wang, Y., Wan, R., Yang, W., Li, H., Chau, L.P., Kot, A.: Low-light image enhancement with normalizing flow. In: Proceedings of the AAAI conference on artificial intelligence. vol. 36, pp. 2604–2612 (2022) 4
- 44. Wei, Chen, W., Wenjing, Y., Wenhan, L., Jiaying: Deep retinex decomposition for low-light enhancement. In: British Machine Vision Conference (2018) 4
- 45. Wizadwongsa, S., Phongthawee, P., Yenphraphai, J., Suwajanakorn, S.: Nex: Realtime view synthesis with neural basis expansion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8534–8543 (2021) 4
- 46. Yang, W., Tan, R.T., Feng, J., Liu, J., Guo, Z., Yan, S.: Deep joint rain detection and removal from a single image. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1357–1366 (2017) 4
- Zhang, H., Patel, V.M.: Densely connected pyramid dehazing network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3194–3203 (2018) 4
- Zhang, K., Luo, W., Zhong, Y., Ma, L., Stenger, B., Liu, W., Li, H.: Deblurring by realistic blurring. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2737–2746 (2020) 8
- 49. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018) 8
- 50. Zhang, X., Dong, H., Pan, J., Zhu, C., Tai, Y., Wang, C., Li, J., Huang, F., Wang, F.: Learning to restore hazy video: A new real-world dataset and a new method.

In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9239–9248 (2021)  $\,4,\,10$