

Aleth-NeRF: Illumination Adaptive NeRF with Concealing Field Assumption

Ziteng Cui^{1,2}, Lin Gu^{3,1}, Xiao Sun^{2*}, Xianzheng Ma⁴, Yu Qiao², Tatsuya Harada^{1,3}

¹The University of Tokyo

²Shanghai AI Laboratory

³RIKEN AIP

⁴University of Oxford

(cui, harada)@mi.t.u-tokyo.ac.jp, lin.gu@riken.jp, (sunxiao, maxianzheng)@pjlab.org.cn, yu.qiao@siat.ac.cn

Abstract

The standard Neural Radiance Fields (NeRF) paradigm employs a viewer-centered methodology, entangling the aspects of illumination and material reflectance into emission solely from 3D points. This simplified rendering approach presents challenges in accurately modeling images captured under adverse lighting conditions, such as low light or over-exposure. Motivated by the ancient Greek emission theory that posits visual perception as a result of rays emanating from the eyes, we slightly refine the conventional NeRF framework to train NeRF under challenging light conditions and generate normal-light condition novel views unsupervisedly. We introduce the concept of a ‘‘Concealing Field,’’ which assigns transmittance values to the surrounding air to account for illumination effects. In dark scenarios, we assume that object emissions maintain a standard lighting level but are attenuated as they traverse the air during the rendering process. Concealing Field thus compel NeRF to learn reasonable density and colour estimations for objects even in dimly lit situations. Similarly, the Concealing Field can mitigate over-exposed emissions during rendering stage. Furthermore, we present a comprehensive multi-view dataset captured under challenging illumination conditions for evaluation. Our code and proposed dataset are available at <https://github.com/cuiziteng/Aleth-NeRF>.

Introduction

Neural Radiance Field (NeRF) (Mildenhall et al. 2020) has been demonstrated to effectively understand 3D scenes from 2D images and generate novel views. However, the formulation of NeRF and its follow-up variants assume captured images are under normal light, often failing to work under low-light (Mildenhall et al. 2021) or over-exposure scenarios. This is because vanilla NeRF is *viewer-centered* which models the amount of light emission from a location to the viewer without disentangling illumination and material (Fig. 1(a)) (Lyu et al. 2022). As a result, the NeRF algorithm interprets a dark scene as insufficient radiation of the 3D object particles, violating the estimation of the object’s material and geometry. In practical applications, images are often taken under challenging lighting conditions.

Therefore, this paper aims to slightly modify vanilla NeRF for under & over-exposure scenes. As shown in Fig. 1(c, d), the proposed, Aleth-NeRF, renders normal-light novel views despite the severe input images.

The rendering process in NeRF (Fig. 1(b)) is similar to the *viewer-centered* emission theory held by ancient Greek. Emission theory ignores the incident light but postulates visual rays emitted from the eye travel in straight lines and interacts with objects to form the visual perception. Therefore, the darkness of an entity is solely caused by the particles between the object and the eye. In other words, all objects are visible by default unless concealed. Inspired by this worldview, we assume a simple but NeRF-friendly concept that it is the concealing fields (gray particle in Fig. 1(c)) in viewing direction that attenuates the emission and makes the viewer see a low-light scene. This is in contrast to the standard NeRF setting where the density of air (white particle in Fig. 1(a)) is usually zero. Introducing the Concealing Field, which assigns the air particles with transmittance value allows NeRF to accurately estimate the colour and density of objects (yellow particles in Fig. 1(c)) in low-light conditions, therefore when removing the concealing fields, or Aletheia ($\alpha\lambda\eta\theta\epsilon\iota\alpha$)¹, we are able to render novel views with normal-light. On the contrary, for the over-exposed scene, deliberately adding the concealing fields in rendering stage could correct the exposure.

Our proposed method Aleth-NeRF takes low-light & over-exposure images as inputs to train the model and learn the volumetric representation jointly with concealing fields. As shown in Fig. 1(b), we jointly train NeRF with concealing fields between the object and viewer. For the low light scenario, we remove the concealing fields during the rendering stage (Fig. 1(c)). When dealing with over-exposure images, Aleth-NeRF would add concealing fields to suppress overly bright (Fig. 1(d)). Our contributions are summarized as follows:

- We propose **Aleth-NeRF**, that trains under low-light & over-exposure conditions and generates novel views under normal-light. Inspired by ancient Greek philosophy, we naturally extend the transmittance function in vanilla NeRF by modelling concealing fields between the ob-

*Corresponding Author
Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹normally translated as ‘‘unconcealedness’’, ‘‘disclosure’’ or ‘‘revealing’’ (Heidegger, Stambaugh, and Schmidt 2010)

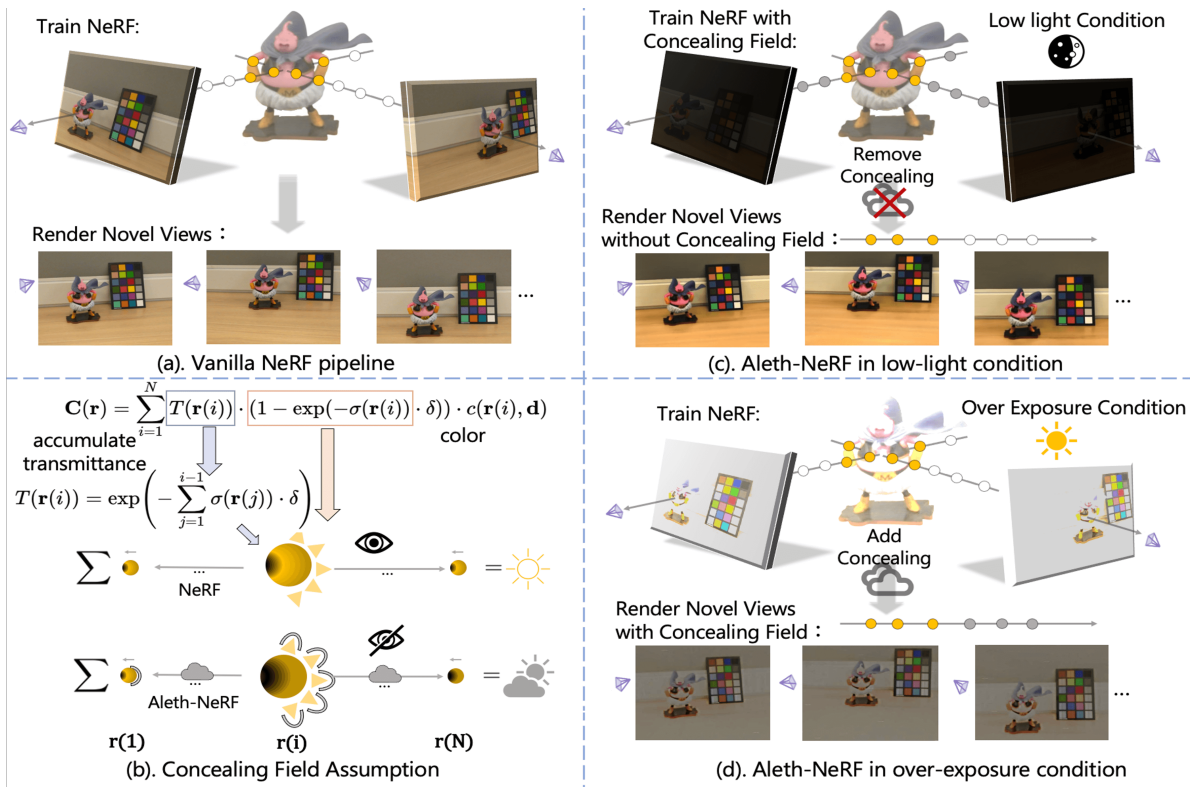


Figure 1: Utilizing the Concealing Field assumption, Aleth-NeRF is capable of processing both low-light & over-expose multi-view images as inputs and generating novel views with natural illumination.

jects and viewer to interpret lightness degradation.

- We contribute a challenging illumination multi-view dataset, with paired sRGB low-light & normal-light & over-exposure images, dataset would also be public.
- We compare with various image enhancement and exposure correction methods & previous NeRF-based method (Mildenhall et al. 2021). Extensive experiments show that our **Aleth-NeRF** achieves satisfactory enhancement quality and multi-view consistency.

Related Works

Novel View Synthesis with NeRF

NeRF (Mildenhall et al. 2020) is proposed for novel view synthesis from a collection of posed input images. The unique advantage of NeRF models exists in preserving the 3D geometry consistency thanks to its physical volume rendering scheme. In addition several methods have been proposed to speed up and improve NeRF training (Barron et al. 2021; Sara Fridovich-Keil and Alex Yu et al. 2022; Lindell, Martel, and Wetzstein 2021; Yu et al. 2021; Jain, Tancik, and Abbeel 2021; Deng et al. 2022; Müller et al. 2022).

Many of the latter works focus on improving NeRF’s performance under various degradation conditions, such as blurry (Ma et al. 2021), noisy (Pearl, Treibitz, and Korman 2022), reflection (Guo et al. 2022), glossy surfaces (Verbin et al. 2022), underwater (Levy et al. 2023), or use NeRF

to handle super-resolution (Wang et al. 2021a; Bahat et al. 2022) and HDR reconstruction (Xin et al. 2021; Jun-Seong et al. 2022) in 3D space. Another line of research extends NeRF for lightness editing in 3D space. Some work, like NeRF-W (Martin-Brualla et al. 2021), focuses on rendering NeRF with uncontrolled in-the-wild images, other relighting works (Srinivasan et al. 2021; Rudnev et al. 2022; Zhang et al. 2021b) rely on known illumination conditions and introduce additional physical elements (*i.e.* normal, light, albedo, etc.), along with complex parametric modeling of these elements. Meanwhile, these methods are not specifically designed for low-light & over-exposure conditions.

Among these, RAW-NeRF (Mildenhall et al. 2021) is more closer to our work, which proposes to render NeRF in HDR RAW domain and then post-process the rendered scene with image signal processor (ISP), RAW-NeRF has shown a preliminary ability to enhance the scene light but requires HDR RAW data for training, which make it hard to generalize on common used sRGB images. Instead our Aleth-NeRF could directly rendered on sRGB under & over exposure images and injection unsupervised enhancement into 3D space by an effective concealing fields manner.

Enhancement in Challenging Light Conditions

Challenging lightness can arise from multiple sources, encompassing natural lighting variances (such as low-light situations and overly bright scenes) as well as human-induced

factors (such as incorrect camera exposure settings). To tackle these challenge lighting conditions, numerous techniques for image enhancement and exposure correction have been developed and proposed.

Image Enhancement & Exposure Correction: Image enhancement methods aims to enhance images with poor illumination, traditional methods usually rely on RetiNex theory (Land 1986; Guo, Li, and Ling 2017) or Histogram Equalization (Gonzalez and Woods 2006), currently deep neural networks (DNNs) based methods become the mainstream solutions, series of CNN & Transformer-based methods have been developed (Wei et al. 2018; Moran et al. 2020; Wang et al. 2021b, 2022; Jiang et al. 2021; Guo et al. 2020; Jin, Yang, and Tan 2022; Ma et al. 2022; Yang et al. 2023). Meanwhile, several exposure correction methods have been proposed to consider both under & over exposure conditions (Afifi et al. 2021; Nsampi, Hu, and Wang 2021; Cui et al. 2022; Huang et al. 2023), which aims to correct under-exposure and its adverse overexposure images into normal-light condition. However, image enhancement & exposure correction methods almost build on 2D image space operations, which often fail to exploit the 3D geometry of the scene and could not deal with multi-view inputs.

Video Enhancement & Burst Enhancement: Beyond above techniques that focus on single image. Video enhancement methods have been proposed to optimize the temporal consistency between adjacent frames, ensuring stability when processing different frames. These methods employ various approaches such as 3D convolution (Lv et al. 2018), optical flow (Zhang et al. 2021a), and event guidance (Liu et al. 2023). Burst enhancement also plays a crucial role in modern computational photography area (Delbracio et al. 2021; Hasinoff et al. 2016), where multiple frames are captured during exposure and processed using an "align-merge-enhance" approach to produce a single output frame. In recent advancements, deep neural networks have been employed to replace traditional manual operation algorithms in these methods (Godard, Matzen, and Uyttendaele 2018; Mildenhall et al. 2018; Dudhane et al. 2022).

However, existing image & video & burst enhancement methods primarily focus on enhancing images in their original views, rather than generating coherent 3D scenes with novel views, for comparison we have to combine these enhancement methods with NeRF (see Table. 2 and Table. 3). In contrast, Aleth-NeRF is capable of directly synthesizing novel views under challenging light conditions while achieving state-of-the-art enhancement quality.

Methods

Neural Radiance Field Revisited

Radiance Field is defined as the density σ and RGB colour value c of a 3D location \mathbf{x} under a 2D viewing direction \mathbf{d} . The density σ , on the one hand, represents the radiation capacity of the particle itself at \mathbf{x} , and on the other hand, controls how much radiance is absorbed when other lights pass through \mathbf{x} .

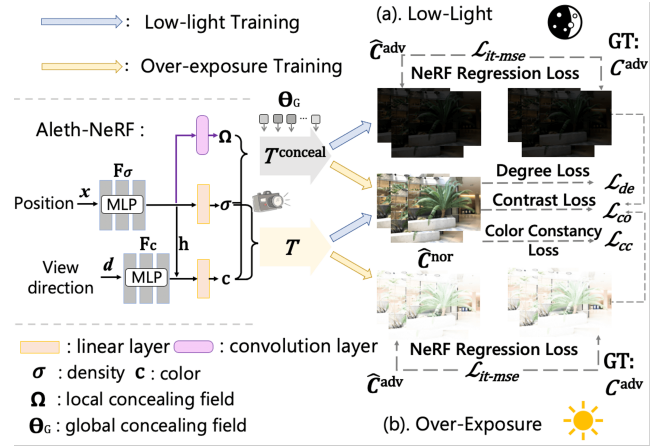


Figure 2: Train on adverse lighting condition images C^{adv} , Aleth-NeRF performs unsupervised lightness correction by (a). remove concealing fields in low-light conditions and (b). add concealing fields in over-exposure conditions.

When rendering an image, a camera ray $\mathbf{r}(t) = \mathbf{o} + t \cdot \mathbf{d}$ ($\mathbf{r} \in \mathbf{R}$) cast from the given camera position \mathbf{o} towards direction \mathbf{d} . All the radiance is accumulated along the ray to render its corresponding pixel value $\mathbf{C}(\mathbf{r})$. Formally,

$$\mathbf{C}(\mathbf{r}) = \int_{t_n}^{t_f} T(\mathbf{r}(t))\sigma(\mathbf{r}(t))c(\mathbf{r}(t), \mathbf{d})dt, \quad (1)$$

where

$$T(\mathbf{r}(t)) = \exp\left(-\int_{t_n}^t \sigma(\mathbf{r}(s))ds\right), \quad (2)$$

is known as the *accumulated transmittance* that denotes the radiance decay rate of the particle at $\mathbf{r}(t)$ when it is occluded by particles closer to the camera (at $\mathbf{r}(s)$, $s < t$). The integrals are computed by a discrete approximation over sampled 3D points along the ray \mathbf{r} (see Fig. 1(b)). The discrete form of Eq. 1 and Eq. 2 are represented as follows:

$$\mathbf{C}(\mathbf{r}) = \sum_{i=1}^N T(\mathbf{r}(i))(1 - \exp(-\sigma(\mathbf{r}(i)) \cdot \delta)) \cdot c(\mathbf{r}(i), \mathbf{d}), \quad (3)$$

$$T(\mathbf{r}(i)) = \exp\left(-\sum_{j=1}^{i-1} \sigma(\mathbf{r}(j)) \cdot \delta\right). \quad (4)$$

Value of $\sigma(\mathbf{r}(i))$ reflects the object occupancy in $\mathbf{r}(i)$. δ is a constant distance value between adjacent sample points under uniform sampling.

For the network structure, NeRF learns two multilayer perceptron (MLP) networks: density MLP F_σ and colour MLP F_c , to map the 3D location $\mathbf{r}(i)$ and 2D viewing direction \mathbf{d} to its density σ and colour c , specifically:

$$F_\sigma(\mathbf{r}(i)) \rightarrow \sigma(\mathbf{r}(i)), \mathbf{h} \quad (5)$$

$$F_c(\mathbf{h}, \mathbf{d}) \rightarrow c(\mathbf{r}(i), \mathbf{d}), \quad (6)$$

where \mathbf{h} is a hidden feature vector that is sent to colour MLP F_c , color $c(\mathbf{r}(i), \mathbf{d})$ and density $\sigma(\mathbf{r}(i))$ are further activated

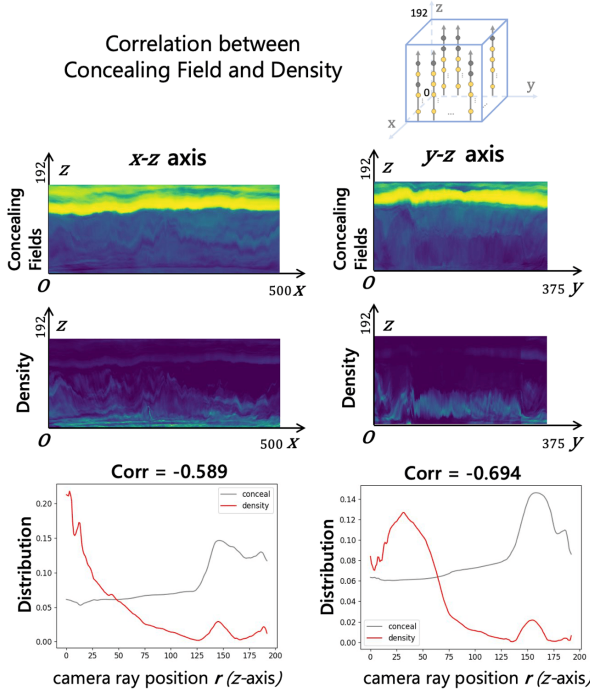


Figure 3: Along camera ray r (z axis), concealing fields and density σ exhibit a negative correlation, (x, y) denotes training images’ width and height.

by Sigmoid and ReLU functions to regularise their ranges into $[0, 1)$ and $[0, \infty)$ respectively. Given ground truth images \mathbf{C} , NeRF is optimised by minimising the MSE loss between predicted images $\hat{\mathbf{C}}$ and ground truth images \mathbf{C} :

$$\mathcal{L}_{mse} = \sum_{\mathbf{r}} \|\hat{\mathbf{C}}(\mathbf{r}) - \mathbf{C}(\mathbf{r})\|^2. \quad (7)$$

We refer more details such as positional encoding and hierarchical volume sampling to NeRF paper (Mildenhall et al. 2020).

Aleth-NeRF with Concealing Field

Given adverse lighting condition images $\mathbf{C}^{adv}(\mathbf{r})$ taken in low-light & over-exposure conditions, our goal is to generate novel views $\mathbf{C}^{nor}(\mathbf{r})$ in normal light condition. The key idea is that Aleth-NeRF assumes \mathbf{C}^{adv} and \mathbf{C}^{nor} are rendered with the same underlying density field σ (yellow particle in Fig. 1) along each camera ray \mathbf{r} , but with or without the proposed concealing fields (gray particle in Fig. 1).

We model two types of concealing fields to reduce light transport in the volume rendering stage: local Concealing Field denoted as Ω at the voxel level, and global Concealing Field denoted as Θ_G at the scene level. Multiplying two fields, Ω and Θ_G gives the final concealing field value which reduce the *accumulated transmittance*. Fig. 2 shows an overview of the Aleth-NeRF training strategy. Fig. 3 shows concealing fields $\Omega \cdot \Theta_G$ ’s distribution.

Local Concealing Field denoted as $\Omega(\mathbf{r}(i))$, defines an extra light concealing capacity of a particle at 3D location $\mathbf{r}(i)$. As depicted in Fig. 2, Ω is individually learned for each 3D position and is generated using the density MLP F_σ . To create the local concealing field Ω , we introduce an additional large kernel convolution layer (with a size of 7) built upon F_σ . This convolution process establishes spatial relationships between pixels and enriches the Concealing Field with predominantly light-related information rather than structural information (Zamir et al. 2020). This larger kernel convolution also effectively suppresses noise and contributes to smoother rendering outcomes.

$$conv_{size:7}(F_\sigma(\mathbf{r}(i))) \rightarrow \Omega(\mathbf{r}(i)) \quad (8)$$

Global Concealing Field denoted as $\Theta_G(i)$, is defined as a set of learnable parameters corresponding to the camera distance i for all camera rays in \mathbf{R} . Θ_G remains constant within a scene and is independent of voxels, as we posit that a specific degree of lighting influence remains consistent across the same scene. In our experiments, we initialize $\Theta_G(i)$ with a value of 0.5 for each camera distance i . With the global and local concealing fields, the *accumulated transmittance* T in Eq. 4 is blighted by Ω and Θ_G to mimic the process of light suppression (see Fig. 1(b)), as follow:

$$T^{conceal}(\mathbf{r}(i)) = \exp \left(- \sum_{j=1}^{i-1} \sigma(\mathbf{r}(j)) \cdot \delta \right) \cdot \prod_{j=1}^{i-1} \Omega(\mathbf{r}(j)) \Theta_G(j). \quad (9)$$

Train and Test Schemes of Aleth-NeRF for both low-light and over-exposure scenes are illustrated in Fig. 2. For low-light conditions (blue arrow in Fig. 2), Aleth-NeRF incorporates $T^{conceal}$ during the volume rendering stage to generate low-light images $\hat{\mathbf{C}}^{adv}$, while removing the concealing fields to produce normal-light images $\hat{\mathbf{C}}^{nor}$ with the original T in Eq. 4. Conversely, for over-exposure scenes (yellow arrow in Fig. 2), Aleth-NeRF employs T to render very-bright images $\hat{\mathbf{C}}^{adv}$ and incorporates the concealing fields along with $T^{conceal}$ to render normal-light images $\hat{\mathbf{C}}^{nor}$. The NeRF regression loss is applied to compare predicted adverse-light images $\hat{\mathbf{C}}^{adv}$ with the corresponding ground truth \mathbf{C}^{adv} , while unsupervised lightness correction losses ensure the generation of predicted normal-light images $\hat{\mathbf{C}}^{nor}$, more details please refer to next section.

Analysis of Concealing Fields is shown in Fig. 3, we analyse the distribution of concealing fields and density along the camera ray $\mathbf{r}(i)$ (z axis), where x, y are the training images’ width and height. In Fig. 3, brighter region denotes the larger possibility there exists concealing fields (gray line) or density (red line). We find that concealing fields exist in location $\mathbf{r}(i)$ with lower density value $\sigma(\mathbf{r}(i))$. And the Pearson correlation coefficient **Corr** between concealing fields

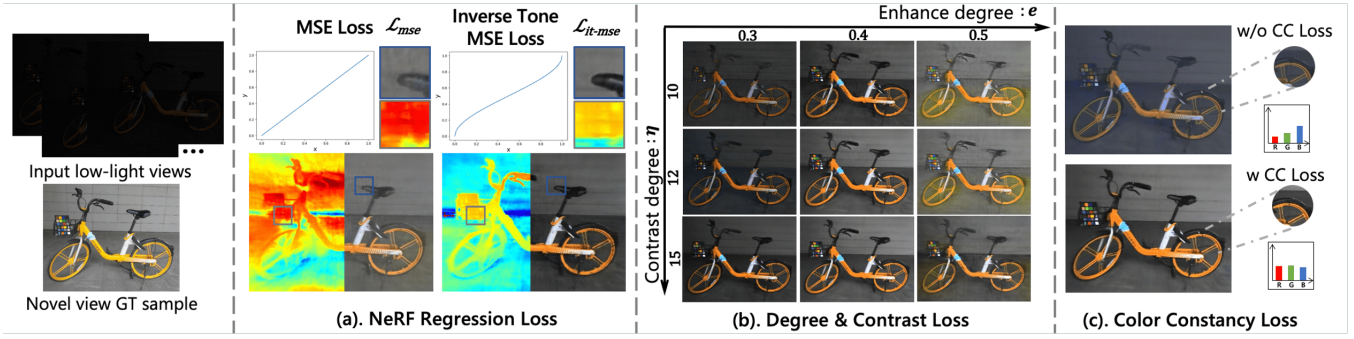


Figure 4: A low-light “bike” scene for example, we show the ablation analyze of different loss functions’ effectiveness.

and density σ are also negative along both $x - z$ axis (-0.589) and $y - z$ axis (-0.694). This validates that concealing fields are separated from density, thus rarely participating in scene rendering. Concealing fields exists more in locations $\mathbf{r}(i)$ with sparse density, *i.e.* air outside the objects.

Effective Losses for Unsupervised Training

In this section, we present the loss functions that guide the unsupervised lightness correction of Aleth-NeRF, facilitating improve novel views synthesis and enhancement in low-light & over-exposure conditions. An overview of the unsupervised training strategy is shown in Fig. 2.

NeRF Regression Loss: Direct apply NeRF’s MSE loss function \mathcal{L}_{mse} in low-light and over-bright scenes would result in an imbalance of pixel weights. Dark pixels with small values would contribute minimally, while bright pixels with larger values would dominate the contribution. To rectify this imbalance and ensure effective NeRF training under non-standard lighting conditions, we propose the incorporation of an inverse tone curve (Brooks et al. 2019; Cui et al. 2021) denoted as Φ . This curve serves to re-balance pixel weights, additionally we introduce a small value $\epsilon = 1e^{-3}$ within the regression loss to facilitate more accurate estimations. As a result, the formulated inverse tone MSE loss \mathcal{L}_{it-mse} for NeRF training is presented as follows:

$$\Phi(x) = \frac{1}{2} - \sin\left(\frac{\sin^{-1}(1 - 2x)}{3}\right),$$

$$\mathcal{L}_{it-mse} = \sum_{\mathbf{r}} \|\Phi(\hat{\mathbf{C}}(\mathbf{r}) + \epsilon) - \Phi(\mathbf{C}(\mathbf{r}) + \epsilon)\|^2, \quad (10)$$

the comparison of \mathcal{L}_{mse} and \mathcal{L}_{it-mse} is shown in Fig. 4(a), inverse tone MSE loss \mathcal{L}_{it-mse} can enhance scene clarity while also suppress noise to a certain extent.

Degree & Contrast Loss: In order to regulate the level of enhancement, we incorporate a degree loss denoted as \mathcal{L}_{de} on the predict normal-light images $\hat{\mathbf{C}}^{nor}$. \mathcal{L}_{de} involves a hyperparameter e , which signifies the enhancement degree inherent to Aleth-NeRF. To compute \mathcal{L}_{de} , a global average pooling operation is applied to the $\hat{\mathbf{C}}^{nor}$ of each training batch. Subsequently, we minimize the degree loss \mathcal{L}_{de} by

comparing the pooled $\hat{\mathbf{C}}^{nor}$ with the specified degree value e . We set e to 0.45 in all experiments, ablation analyse of e is in Fig. 4(b).

$$\mathcal{L}_{de} = \|\text{avgpool}(\hat{\mathbf{C}}^{nor}(\mathbf{r})) - e\|^2. \quad (11)$$

Then we additional design a contrast degree loss \mathcal{L}_{co} between the adjacent sampling pixels of predict normal-light images $\hat{\mathbf{C}}^{nor}$ and GT adverse-light images \mathbf{C}^{adv} :

$$\mathcal{L}_{co} = \sum_{k \in \{+1, -1\}} (\hat{\mathbf{C}}^{nor}(\mathbf{r}) - \hat{\mathbf{C}}^{nor}(\mathbf{r} + k)) - e \cdot \eta \cdot (\mathbf{C}^{adv}(\mathbf{r}) - \mathbf{C}^{adv}(\mathbf{r} + k)), \quad (12)$$

the contrast degree loss would maintain the structure similarity between $\hat{\mathbf{C}}^{nor}$ and \mathbf{C}^{adv} , meanwhile control the contrast degree of $\hat{\mathbf{C}}^{nor}$ by parameter η , ablation analyze of η in low-light scene is shown in Fig. 4, when enhance degree e is fixed, higher η represents higher contrast in $\hat{\mathbf{C}}^{nor}$. While in over-exposure scene, $e \cdot \eta$ would fixed to 0.5.

Color Constancy Loss: Photos taken in low-light & over exposure conditions would easily lose their color information, directly remove & add the concealing fields may cause color imbalance. To regularize the color of the predict normal-light images $\hat{\mathbf{C}}^{nor}$, we introduce color constancy loss \mathcal{L}_{cc} on $\hat{\mathbf{C}}^{nor}$. Here we assume that $\hat{\mathbf{C}}^{nor}$ obey the gray-world assumption (Buchsbaum 1980; Guo et al. 2020; Jin, Yang, and Tan 2022), as follows:

$$\mathcal{L}_{cc} = \sum_{p,q} (\hat{\mathbf{C}}^{nor}(\mathbf{r})^p - \hat{\mathbf{C}}^{nor}(\mathbf{r})^q)^2, \quad (13)$$

where $(p, q) \in \{(R, G), (G, B), (B, R)\}$ represents any pair of color channels, an ablation of color constancy loss \mathcal{L}_{cc} can be found in Fig. 4(c).

Above all, Aleth-NeRF’s loss function \mathcal{L} comprises four components: NeRF rendering loss \mathcal{L}_{it-mse} , along with three unsupervised lightness correction losses (\mathcal{L}_{de} , \mathcal{L}_{co} , and \mathcal{L}_{cc}). The overall training loss is then represented as:

$$\mathcal{L} = \mathcal{L}_{it-mse} + \lambda_1 \cdot \mathcal{L}_{de} + \lambda_2 \cdot \mathcal{L}_{co} + \lambda_3 \cdot \mathcal{L}_{cc}, \quad (14)$$

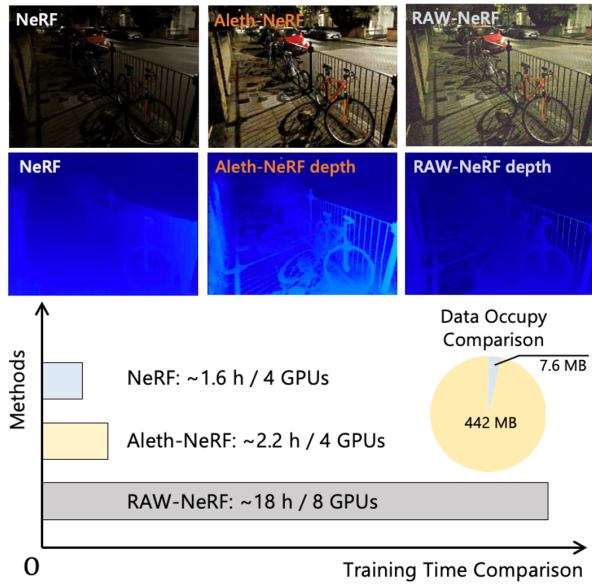


Figure 5: An comparison of enhancement results and model efficiency with RAW-NeRF (Mildenhall et al. 2021), note that RAW-NeRF take 16-bits HDR image as inputs while NeRF & Aleth-NeRF take 8-bits LDR image as inputs.

where λ_1 , λ_2 and λ_3 are three non-negative parameters to balance total loss weights, which we set to $1e^{-3}$, $1e^{-3}$ and $1e^{-8}$ respectively.

Experiments

In this section, we first introduce our collected dataset with various exposure paired multi-view images, then we present the novel view synthesis results of Aleth-NeRF under both low-light and under-exposure settings. For practical implementation, our framework builds on the open-source PyTorch toolbox `NeRF-Factory`². We utilize the Adam optimizer with an initial learning rate of $5e^{-4}$ and employ a cosine learning rate decay strategy every 2500 iterations. The training batch size is set at 4096 for a total of 62500 iterations. To ensure fairness, all comparative experiments within the proposed dataset use the same training strategies and parameter configurations.

Challenging Illumination Multi-view Dataset

In this section, we introduce our collected paired low-light & normal-light & over-exposure multi-view dataset, names LOM dataset. Previous work (Mildenhall et al. 2021) also includes some low-light scenes, but their dataset concentrates on RAW denoising rather than sRGB low-light enhancement and does not include normal-light sRGB ground truth, making it hard to evaluate enhancement performance in novel view synthesis.

In our proposed LOM dataset, we collected 5 scenes (“*buu*”, “*chair*”, “*sofa*”, “*bike*”, “*shrub*”) in real-world.

²<https://github.com/kakaobrain/nerf-factory>

scene	“ <i>buu</i> ”	“ <i>chair</i> ”	“ <i>sofa</i> ”	“ <i>bike</i> ”	“ <i>shrub</i> ”
collected views	25	48	33	40	35
training views	22	43	29	36	30
evaluation views	3	5	4	4	5

Table 1: Details of the dataset split for LOM.

Each scene includes 25 ~ 48 images, for each view’s image, we generate low-light & normal-light & over-exposure pairs by adjusting exposure time and ISO while other configurations of the camera are fixed. We capture multi-view images by moving and rotating the camera tripod. Images are collected with resolution 3000×4000 . We down-sample the original resolution with ratio 8 to 375×500 for convenience, and generate the ground truth view and angle information by adopting COLMAP (Schönberger and Frahm 2016; Schönberger et al. 2016) on the normal-light part images. For dataset split, in each scene, we choose 3 ~ 5 images as the testing set, 1 image as the validation set, and other images to be the training set, details of training and evaluation views split is shown in Table. 1. More details of our dataset please refer to supplementary part.

Novel View Synthesis in Low-light Condition

In this section, we show novel view synthesis results in low-light scenes, and make comparison between predicted normal-light views with ground truth normal-light views (see Table. 2). We design multiple comparison experiments in proposed dataset as follow:

We start by training vanilla NeRF (Mildenhall et al. 2020) under low-light conditions as a baseline (first row in Table 2). ① Subsequently, we apply various state-of-the-art enhancement methods to post-process the rendered NeRF results (denote as “NeRF + ” in Table 2), includes RetinexNet (Wei et al. 2018), Zero-DCE (Guo et al. 2020), SCI (Ma et al. 2022) and IAT (Cui et al. 2022). ② Then we pre-process low-light images with enhancement methods and then train NeRF on these enhanced images (denote as “* + NeRF” in Table. 2). Notably, we observe that the “NeRF + *” methods tend to underperform compared to the “* + NeRF” methods in low-light setting, especially in perceptual similarity (LPIPS \downarrow). This potentially arise from NeRF’s production of poor-quality images in dark scenes, causing enhancement methods to struggle when applied to the generated dark scene by NeRF. ③ At last, we apply two video enhancement methods MBLLN (Lv et al. 2018) and LLVE (Zhang et al. 2021a) to pre-process low-light images and then trained with NeRF. ④ Above all, our Aleth-NeRF could gain the best or second best performance on most scenes (see Table. 2), at the same time Aleth-NeRF does not require any training data as prior knowledge, and trained in an end-to-end manner without any pre-processing nor post-processing.

Beyond enhancement methods, RAW-NeRF (Mildenhall et al. 2021) also conducted experiments under low-light conditions, due to their model is trained on RAW data and

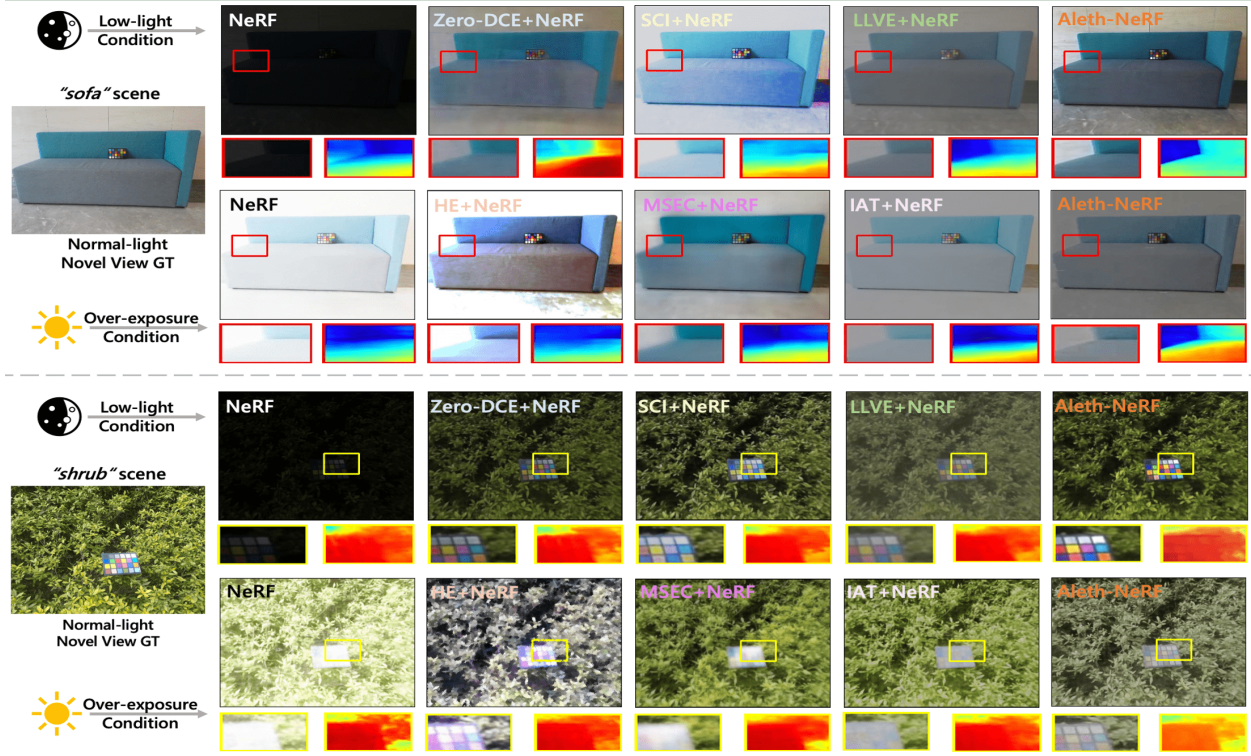


Figure 6: Novel view synthesis comparison in low-light and over-exposure conditions.

Method	“ <i>buu</i> ”	“ <i>chair</i> ”	“ <i>sofa</i> ”	“ <i>bike</i> ”	“ <i>shrub</i> ”	<i>mean</i>
	P ↑/ S ↑/ L ↓	P ↑/ S ↑/ L ↓	P ↑/ S ↑/ L ↓	P ↑/ S ↑/ L ↓	P ↑/ S ↑/ L ↓	P ↑/ S ↑/ L ↓
NeRF	7.51/ 0.291/ 0.448	6.04/ 0.147/ 0.594	6.28/ 0.210/ 0.568	6.35/ 0.072/ 0.623	8.03/ 0.031/ 0.680	6.84/ 0.150/ 0.582
NeRF + Image Enhancement Methods						
NeRF+ RetinexNet	11.64/ 0.646/ 0.395	11.04/ 0.631/ 0.561	12.85/ 0.738/ 0.527	17.79/ 0.653/ 0.550	11.85/ 0.211/ 0.583	13.03/ 0.576/ 0.523
NeRF+ Zero-DCE	17.81/ 0.833/ 0.357	12.44/ 0.684/ 0.547	14.43/ 0.787/ 0.539	10.16/ 0.468/ 0.557	12.58/ 0.282/ 0.540	13.48/ 0.610/ 0.488
NeRF+ SCI	7.84/ 0.660/ 0.562	12.07/ 0.699/ 0.584	10.25/ 0.737/ 0.626	<u>18.84</u> / 0.637/ 0.565	12.38/ 0.358/ 0.587	12.27/ 0.618/ 0.585
NeRF+ IAT	14.03/ 0.656/ 0.453	19.08/ 0.800/ 0.565	10.49/ 0.528/ 0.678	17.16/ 0.657/ 0.534	16.15/ 0.344/ 0.573	15.38/ 0.597/ 0.561
Image Enhancement Methods + NeRF						
RetinexNet+ NeRF	16.19/ 0.780/ 0.396	16.89/ 0.756/ 0.543	16.98/ 0.807/ 0.577	18.00/ <u>0.707</u> / 0.482	14.86/ 0.284/ 0.518	16.58/ 0.667/ 0.503
Zero-DCE+ NeRF	17.90/ 0.858/ 0.376	12.58/ 0.721/ 0.460	14.45/ 0.831/ <u>0.419</u>	10.39/ 0.518/ <u>0.464</u>	12.32/ 0.308/ 0.481	13.53/ 0.649/ <u>0.432</u>
SCI+ NeRF	7.76/ 0.692/ 0.525	19.77/ <u>0.802</u> / 0.674	10.08/ 0.772/ 0.520	13.44/ 0.658/ 0.435	<u>18.16</u> / <u>0.503</u> / <u>0.475</u>	13.84/ 0.689/ 0.510
IAT+ NeRF	14.46/ 0.705/ 0.386	18.70/ 0.780/ 0.665	17.88/ 0.829/ 0.547	13.65/ 0.616/ 0.528	13.87/ 0.317/ 0.536	15.71/ 0.649/ 0.532
Video Enhancement Methods + NeRF						
MBLLEN+ NeRF	22.39 / 0.877 / <u>0.353</u>	23.59 / 0.788/ 0.559	19.99 / 0.836/ 0.542	14.09/ 0.636/ 0.525	13.17/ 0.501/ 0.555	<u>18.65</u> / <u>0.728</u> / 0.507
LLVE+ NeRF	19.97/ 0.848/ 0.393	15.17/ 0.764/ 0.610	18.17/ <u>0.855</u> / 0.465	13.84/ 0.638/ 0.492	15.35/ 0.287/ 0.577	16.50/ 0.678/ 0.507
Our Proposed End-to-end Method						
Aleth-NeRF	<u>20.22</u> / 0.859/ 0.315	<u>20.93</u> / 0.818 / <u>0.468</u>	<u>19.52</u> / 0.857 / 0.354	20.46 / 0.727 / 0.499	18.24 / 0.511 / 0.448	19.87 / 0.754 / 0.417

Table 2: We assess novel view synthesis performance in low-light settings by comparing generated images with ground truth normal-light views. Our evaluation metrics include PSNR (P) ↑, SSIM (S) ↑ and LPIPS (L) ↓. Bold denotes the best result, Underline denotes the second best result.

Method	“ <i>buu</i> ”	“ <i>chair</i> ”	“ <i>sofa</i> ”	“ <i>bike</i> ”	“ <i>shrub</i> ”	<i>mean</i>
	P ↑/S ↑/L ↓	P ↑/S ↑/L ↓	P ↑/S ↑/L ↓	P ↑/S ↑/L ↓	P ↑/S ↑/L ↓	P ↑/S ↑/L ↓
NeRF	7.12/0.674/0.499	11.05/0.741/ 0.418	10.22/0.783/0.475	9.65/0.698/0.416	9.96/0.405/ <u>0.480</u>	9.60/0.660/ <u>0.457</u>
NeRF + Exposure Correction Methods						
NeRF+ HE	14.34/0.613/0.673	15.37/0.661/0.590	16.69/0.733/0.558	15.32/0.627/0.458	11.97/0.468/0.556	14.74/0.620/0.567
NeRF+ IAT	14.11/0.780/ <u>0.433</u>	<u>19.24</u> / <u>0.810</u> /0.491	16.60/0.837/ <u>0.459</u>	17.73/0.760/0.394	14.05/0.381/0.499	16.35/0.714/ 0.455
NeRF+ MSEC	16.13/0.800/ 0.427	15.60/0.786/ <u>0.472</u>	16.56/0.807/0.495	12.60/0.716/0.465	13.66/0.332/0.509	14.91/0.688/0.473
Exposure Correction Methods + NeRF						
HE+ NeRF	14.65/0.743/0.519	15.55/0.736/0.497	17.11/0.781/0.477	15.77/0.692/ 0.367	12.61/ 0.506 /0.622	15.13/0.692/0.496
IAT+ NeRF	<u>16.22</u> / <u>0.815</u> /0.486	18.98/0.799/0.503	<u>18.45</u> / <u>0.849</u> /0.478	<u>19.63</u> / 0.776 /0.408	<u>15.63</u> /0.434/ 0.477	<u>17.78</u> / <u>0.734</u> /0.470
MSEC+ NeRF	15.53/ 0.817 /0.499	16.95/0.758/0.580	19.60 /0.817/0.498	18.90/0.725/0.483	15.48/0.400/0.499	17.29/0.703/0.512
Our Proposed End-to-end Method						
Aleth-NeRF	16.78 /0.805/0.611	20.08 / 0.820 /0.499	17.85/ 0.852 / 0.458	19.85 / <u>0.773</u> / <u>0.392</u>	15.91 / <u>0.477</u> /0.483	18.09 / 0.745 /0.488

Table 3: We assess novel view synthesis performance in over-exposure settings by comparing generated images with ground truth normal-light views. Our evaluation metrics include PSNR ↑, SSIM ↑ and LPIPS ↓. Bold denotes the best result, Underline denotes the second best result.

hard to implement on our dataset, a comparison with RAW-NeRF on their proposed dataset scene is shown in Fig. 5, while NeRF and Aleth-NeRF are trained with 8-bits LDR inputs and RAW-NeRF is trained with 16-bits HDR inputs, our Aleth-NeRF could also gain delight enhancement results while save much data occupy and training time.

Fig. 6 shows qualitative visualization results in LOM dataset’s “*sofa*” and “*shrub*” scenes, with the comparison of current SOTA image enhancement methods Zero-DCE (Guo et al. 2020), SCI (Ma et al. 2022) and video enhancement method LLVE (Zhang et al. 2021a), we can see that our Aleth-NeRF could gain more vivid enhancement results and closer to ground truth, meanwhile also gain much better in multi-view consistency (see depth maps). More visualization results please refer to supplementary part.

Novel View Synthesis in Over-Exposure Condition

Then we show novel view synthesis results in over-exposure scenes, similar to low-light scenes, we make comparison with 3 mainstream exposure correction methods, including traditional histogram equalization (HE) method (Gonzalez and Woods 2006) and two current SOTA deep network methods MSEC (Afifi et al. 2021) and IAT (Cui et al. 2022).

As it shown in Table. 3, we use vanilla NeRF trained on over-exposure images to serve as a baseline. ① Then we adopt exposure correction methods to post-process the generated results of NeRF (denote as “NeRF + *”). ② We also adopt exposure correction methods to pre-process over-exposure images and then train NeRF on the processed images (denote as “* + NeRF”). ③ Our proposed Aleth-NeRF method. could gain best results in PSNR and SSIM, but the performance in LPIPS is not very good. Some visualization results are shown in Fig. 6, our Aleth-NeRF could suppress over-bright meanwhile maintain multi-view consistency.

One major reason for the less pronounced LPIPS results

on over-exposed images maybe that Aleth-NeRF’s training directly based on degraded overexposed images. Therefore, unsupervised learning faces challenges in compensating for these degraded pieces of information. Aleth-NeRF’s assumption about darkness is grounded in the Concealing Field present in the air, we believe that a more effective modeling of darkness might offer improved solutions for addressing and handling overexposure issues.

Conclusion

In this paper, we propose Aleth-NeRF model. Inspired by wisdom of ancient Greeks, we introduce a concept: Concealing Field. After incorporating the concealing fields into NeRF, we can effectively address novel view synthesis tasks in scenes with low-light & over-exposure conditions. Additionally, we proposed several unsupervised loss functions to constrain the generation of the concealing fields, make Aleth-NeRF could achieve SOTA results on novel view synthesis under low-light & over-exposure conditions.

One limitation is that Aleth-NeRF should be specifically trained for each scene, which is the same as vanilla NeRF. Besides, Aleth-NeRF may fail in scenes with non-uniform lighting conditions (Wang, Xu, and Lau 2022) or shadow conditions (Qu et al. 2017), we believe that this is also a valuable research topic for future exploration.

Acknowledgements

This work was partially supported by JST Moonshot R&D Grant Number JPMJPS2011, CREST Grant Number JPMJCR2015 and Basic Research Grant (Super AI) of Institute for AI and Beyond of the University of Tokyo. And this work is supported in part by the National Key R&D Program of China (NO. 2022ZD0160100).

References

- Affi, M.; Derpanis, K. G.; Ommer, B.; and Brown, M. S. 2021. Learning Multi-Scale Photo Exposure Correction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9157–9167.
- Bahat, Y.; Zhang, Y.; Sommerhoff, H.; Kolb, A.; and Heide, F. 2022. Neural Volume Super-Resolution.
- Barron, J. T.; Mildenhall, B.; Tancik, M.; Hedman, P.; Martin-Brualla, R.; and Srinivasan, P. P. 2021. Mip-NeRF: A Multiscale Representation for Anti-Aliasing Neural Radiance Fields. arXiv:2103.13415.
- Brooks, T.; Mildenhall, B.; Xue, T.; Chen, J.; Sharlet, D.; and Barron, J. T. 2019. Unprocessing Images for Learned Raw Denoising. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Buchsbaum, G. 1980. A spatial processor model for object colour perception. *Journal of the Franklin Institute*, 310(1): 1–26.
- Cui, Z.; Li, K.; Gu, L.; Su, S.; Gao, P.; Jiang, Z.; Qiao, Y.; and Harada, T. 2022. You Only Need 90K Parameters to Adapt Light: a Light Weight Transformer for Image Enhancement and Exposure Correction. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022*. BMVA Press.
- Cui, Z.; Qi, G.-J.; Gu, L.; You, S.; Zhang, Z.; and Harada, T. 2021. Multitask AET With Orthogonal Tangent Regularity for Dark Object Detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2553–2562.
- Delbracio, M.; Kelly, D.; Brown, M. S.; and Milanfar, P. 2021. Mobile Computational Photography: A Tour. *Annual Review of Vision Science*, 7(1): 571–604. PMID: 34524880.
- Deng, K.; Liu, A.; Zhu, J.-Y.; and Ramanan, D. 2022. Depth-supervised NeRF: Fewer Views and Faster Training for Free. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Dudhane, A.; Zamir, S. W.; Khan, S.; Khan, F. S.; and Yang, M.-H. 2022. Burst image restoration and enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5759–5768.
- Godard, C.; Matzen, K.; and Uyttendaele, M. 2018. Deep Burst Denoising. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Gonzalez, R. C.; and Woods, R. E. 2006. *Digital Image Processing (3rd Edition)*. USA: Prentice-Hall, Inc. ISBN 013168728X.
- Guo, C. G.; Li, C.; Guo, J.; Loy, C. C.; Hou, J.; Kwong, S.; and Cong, R. 2020. Zero-reference deep curve estimation for low-light image enhancement. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 1780–1789.
- Guo, X.; Li, Y.; and Ling, H. 2017. LIME: Low-Light Image Enhancement via Illumination Map Estimation. *IEEE Transactions on Image Processing*, 26(2): 982–993.
- Guo, Y.-C.; Kang, D.; Bao, L.; He, Y.; and Zhang, S.-H. 2022. NeRFReN: Neural Radiance Fields With Reflections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 18409–18418.
- Hasinoff, S. W.; Sharlet, D.; Geiss, R.; Adams, A.; Barron, J. T.; Kainz, F.; Chen, J.; and Levoy, M. 2016. Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Transactions on Graphics (ToG)*, 35(6): 1–12.
- Heidegger, M.; Stambaugh, J.; and Schmidt, D. 2010. *Being and Time*. SUNY Series in Contemporary Co. State University of New York Press.
- Huang, J.; Zhao, F.; Zhou, M.; Xiao, J.; Zheng, N.; Zheng, K.; and Xiong, Z. 2023. Learning Sample Relationship for Exposure Correction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9904–9913.
- Jain, A.; Tancik, M.; and Abbeel, P. 2021. Putting NeRF on a Diet: Semantically Consistent Few-Shot View Synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 5885–5894.
- Jiang, Y.; Gong, X.; Liu, D.; Cheng, Y.; Fang, C.; Shen, X.; Yang, J.; Zhou, P.; and Wang, Z. 2021. EnlightenGAN: Deep light enhancement without paired supervision. *IEEE Transactions on Image Processing*, 30: 2340–2349.
- Jin, Y.; Yang, W.; and Tan, R. T. 2022. Unsupervised Night Image Enhancement: When Layer Decomposition Meets Light-Effects Suppression. *arXiv preprint arXiv:2207.10564*.
- Jun-Seong, K.; Yu-Ji, K.; Ye-Bin, M.; and Oh, T.-H. 2022. HDR-Plenoxels: Self-Calibrating High Dynamic Range Radiance Fields. In *ECCV*.
- Land, E. H. 1986. An Alternative Technique for the Computation of the Designator in the Retinex Theory of Color Vision. *Proceedings of the National Academy of Sciences of the United States of America*.
- Levy, D.; Peleg, A.; Pearl, N.; Rosenbaum, D.; Akkaynak, D.; Korman, S.; and Treibitz, T. 2023. SeaThru-NeRF: Neural Radiance Fields in Scattering Media. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 56–65.
- Lindell, D. B.; Martel, J. N. P.; and Wetzstein, G. 2021. AutoInt: Automatic Integration for Fast Neural Volume Rendering. In *Proc. CVPR*.
- Liu, L.; An, J.; Liu, J.; Yuan, S.; Chen, X.; Zhou, W.; Li, H.; Wang, Y. F.; and Tian, Q. 2023. Low-Light Video Enhancement with Synthetic Event Guidance. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(2): 1692–1700.
- Lv, F.; Lu, F.; Wu, J.; and Lim, C. 2018. MBLLEN: Low-light Image/Video Enhancement Using CNNs. In *British Machine Vision Conference*.
- Lyu, L.; Tewari, A.; Leimkuehler, T.; Habermann, M.; and Theobalt, C. 2022. Neural Radiance Transfer Fields for Relightable Novel-view Synthesis with Global Illumination. In *ECCV*.

- Ma, L.; Li, X.; Liao, J.; Zhang, Q.; Wang, X.; Wang, J.; and Sander, P. V. 2021. Deblur-NeRF: Neural Radiance Fields from Blurry Images. *arXiv preprint arXiv:2111.14292*.
- Ma, L.; Ma, T.; Liu, R.; Fan, X.; and Luo, Z. 2022. Toward Fast, Flexible, and Robust Low-Light Image Enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5637–5646.
- Martin-Brualla, R.; Radwan, N.; Sajjadi, M. S. M.; Barron, J. T.; Dosovitskiy, A.; and Duckworth, D. 2021. NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections. In *CVPR*.
- Mildenhall, B.; Barron, J. T.; Chen, J.; Sharlet, D.; Ng, R.; and Carroll, R. 2018. Burst denoising with kernel prediction networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2502–2510.
- Mildenhall, B.; Hedman, P.; Martin-Brualla, R.; Srinivasan, P. P.; and Barron, J. T. 2021. NeRF in the Dark: High Dynamic Range View Synthesis from Noisy Raw Images. *arXiv*.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *ECCV*.
- Moran, S.; Marza, P.; McDonagh, S.; Parisot, S.; and Slabaugh, G. 2020. DeepLPF: Deep Local Parametric Filters for Image Enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Müller, T.; Evans, A.; Schied, C.; and Keller, A. 2022. Instant Neural Graphics Primitives with a Multiresolution Hash Encoding. *ACM Trans. Graph.*, 41(4): 102:1–102:15.
- Nsampi, N. E.; Hu, Z.; and Wang, Q. 2021. Learning exposure correction via consistency modeling. In *Proc. Brit. Mach. Vision Conf.*
- Pearl, N.; Treibitz, T.; and Korman, S. 2022. NAN: Noise-Aware NeRFs for Burst-Denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 12672–12681.
- Qu, L.; Tian, J.; He, S.; Tang, Y.; and Lau, R. W. H. 2017. DshadowNet: A Multi-context Embedding Deep Network for Shadow Removal. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2308–2316.
- Rudnev, V.; Elgharib, M.; Smith, W.; Liu, L.; Golyanik, V.; and Theobalt, C. 2022. NeRF for Outdoor Scene Relighting. In *European Conference on Computer Vision (ECCV)*.
- Sara Fridovich-Keil and Alex Yu; Tancik, M.; Chen, Q.; Recht, B.; and Kanazawa, A. 2022. Plenoxels: Radiance Fields without Neural Networks. In *CVPR*.
- Schönberger, J. L.; and Frahm, J.-M. 2016. Structure-from-Motion Revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Schönberger, J. L.; Zheng, E.; Pollefeys, M.; and Frahm, J.-M. 2016. Pixelwise View Selection for Unstructured Multi-View Stereo. In *European Conference on Computer Vision (ECCV)*.
- Srinivasan, P. P.; Deng, B.; Zhang, X.; Tancik, M.; Mildenhall, B.; and Barron, J. T. 2021. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7495–7504.
- Verbin, D.; Hedman, P.; Mildenhall, B.; Zickler, T.; Barron, J. T.; and Srinivasan, P. P. 2022. Ref-NeRF: Structured View-Dependent Appearance for Neural Radiance Fields. *CVPR*.
- Wang, C.; Wu, X.; Guo, Y.-C.; Zhang, S.-H.; Tai, Y.-W.; and Hu, S.-M. 2021a. NeRF-SR: High-Quality Neural Radiance Fields using Supersampling. *arXiv*.
- Wang, H.; Xu, K.; and Lau, R. W. 2022. Local Color Distributions Prior for Image Enhancement. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Wang, T.; Zhang, K.; Shen, T.; Luo, W.; Stenger, B.; and Lu, T. 2022. Ultra-High-Definition Low-Light Image Enhancement: A Benchmark and Transformer-Based Method. *arXiv preprint arXiv:2212.11548*.
- Wang, Y.; Wan, R.; Yang, W.; Li, H.; Chau, L.-P.; and Kot, A. C. 2021b. Low-Light Image Enhancement with Normalizing Flow. *arXiv preprint arXiv:2109.05923*.
- Wei, C.; Wang, W.; Yang, W.; and Liu, J. 2018. Deep Retinex Decomposition for Low-Light Enhancement. In *British Machine Vision Conference*.
- Xin, H.; Qi, Z.; Ying, F.; Hongdong, L.; Xuan, W.; and Qing, W. 2021. HDR-NeRF: High Dynamic Range Neural Radiance Fields. *arXiv preprint arXiv:2111.14451*.
- Yang, S.; Ding, M.; Wu, Y.; Li, Z.; and Zhang, J. 2023. Implicit Neural Representation for Cooperative Low-light Image Enhancement. *arXiv:2303.11722*.
- Yu, A.; Li, R.; Tancik, M.; Li, H.; Ng, R.; and Kanazawa, A. 2021. PlenOctrees for Real-time Rendering of Neural Radiance Fields. In *ICCV*.
- Zamir, S. W.; Arora, A.; Khan, S.; Hayat, M.; Khan, F. S.; Yang, M.-H.; and Shao, L. 2020. CycleISP: Real Image Restoration via Improved Data Synthesis. In *CVPR*.
- Zhang, F.; Li, Y.; You, S.; and Fu, Y. 2021a. Learning Temporal Consistency for Low Light Video Enhancement From Single Images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4967–4976.
- Zhang, X.; Srinivasan, P. P.; Deng, B.; Debevec, P.; Freeman, W. T.; and Barron, J. T. 2021b. NeRFactor: Neural Factorization of Shape and Reflectance under an Unknown Illumination. *ACM Trans. Graph.*, 40(6).