

Log NeRF: Comparing Spaces for Learning Radiance Fields

Sihe Chen
chen.sihe1@northeastern.edu

Northeastern University
Boston, MA, USA

Luv Verma
verma.lu@northeastern.edu

Bruce A. Maxwell
b.maxwell@northeastern.edu

Abstract

Neural Radiance Fields (NeRF) have achieved remarkable results in novel view synthesis, typically using sRGB images for supervision. However, little attention has been paid to the color space in which the network is learning the radiance field representation. Inspired by the Bi-Illuminant Dichromatic Reflection (BIDR) model, which suggests that a logarithmic transformation simplifies the separation of illumination and reflectance, we hypothesize that log-RGB space enables NeRF to learn a more compact and effective representation of scene appearance.

To test this, we captured approximately 30 videos using a GoPro camera, ensuring linear data recovery through inverse encoding. We trained NeRF models under various color space interpretations—linear, sRGB, GPLog, and log RGB—by converting each network output to a common color space before rendering and loss computation, enforcing representation learning in different color spaces. Quantitative and qualitative evaluations demonstrate that using a log RGB color space consistently improves rendering quality, exhibits greater robustness across scenes, and performs particularly well in low-light conditions while using the same bit-depth input images. Further analysis across different network sizes and NeRF variants confirms the generalization and stability of the log-space advantage.

1 Introduction

Neural Radiance Fields (NeRF) are a popular technique for novel view synthesis and 3D scene representation [VR21]. NeRF methods use a multi-layer perceptron (MLP) to learn a continuous representation of a scene’s volumetric structure and appearance from a set of training images. For a single query, the input to the NeRF network is a 3D coordinate (x, y, z) and a 2D viewing direction (θ, ϕ) , typically combined with a positional encoding. The MLP then outputs a single color value (R, G, B) and a scalar volume density σ for that specific point and direction. To render a novel high-resolution view, a volumetric renderer will sample multiple points along each camera ray, querying the network at each point to gather a set of colors and densities. It then integrates the values along the ray to compute a final color for each pixel.

There are many variations on NeRF models that have explored a wide variety of issues that arise when trying to model the light field of a scene. With a few exceptions that have used RAW image data (e.g. [22]), NeRF work to date has focused on either synthetic data or data captured with consumer devices and processed for human viewing. The implied assumption in most prior work is that the network representing the light field is learning it in the input color space, also assumed to be the output color space, which is generally sRGB.

One reason that most NeRF work has used videos processed for human viewing is that obtaining linear RGB videos has been challenging without high-end video cameras. While capturing single images in RAW mode is possible with most consumer devices, capturing RAW video has not been as simple (though it is becoming easier [11]). Recently, however, GoPro released consumer cameras that can capture video in their GProLog encoding format, which is a 10-bit transformation of the linear RGB data. The transformation is invertible into linear RGB without significant loss, which enables transformations into other spaces like log RGB and sRGB. This makes it possible to easily capture linearizable videos appropriate for NeRF processing. While more time-consuming than capturing a scene with video, it is also possible to linearize a large set of photos of a scene taken with any camera that can capture RAW data. However, pose estimation can be more difficult compared to video sequences.

In this work, we explore whether the representation space in which the network learns the radiance field impacts the training, compactness, robustness, dynamic range, and quality of the NeRF process. For each of these characteristics, we execute a series of experiments and compare the GProLog encoding, sRGB, linear RGB, and log RGB. In order to achieve fair comparisons, the training data in all cases is sRGB derived from the GProLog encoded video, and losses are computed in the sRGB color space.

To make our work replicable, we use the BiLaRF NeRF model as the basis for all of our experiments [6]. We control the radiance field color space implicitly using two small modifications. First, we use a hard-coded transformation from the implicit color space to linear space between the radiance field MLP and the differentiable renderer. We then execute an sRGB transform on the linear renderer output in order to compute losses in sRGB.

We explore compactness by evaluating the relationship between the learned space, network size, network training time, and Peak Signal-to Noise Ratio (PSNR). We explore robustness by executing multiple training runs using each color space and computing statistics on performance. We compare dynamic range by evaluating each space on a video with controlled variation in intensity. Finally, we compare quality through PSNR for the output images and a qualitative comparison of the depth maps learned by the network on a range of real videos taken indoors and outdoors with both simple and complex lighting situations.

In all of these metrics we found that log RGB outperformed the other color spaces, in many cases by a significant margin using the same base architecture and training procedure. In order to determine the extent to which log RGB is special, we also tested similar, but not exactly log functions. Only true log RGB demonstrated improvement over the other spaces.

Overall, our contributions are as follows: (1) we explore the impact of the internal representation space for radiance fields, (2) we demonstrate how to control the implicit representation space of a NeRF network, (3) we show that log RGB outperforms other color spaces on multiple metrics, (4) we show that log RGB is special; similar transforms do not show the same benefits, and (5) we contribute a new set of NeRF videos in GProLog encoding format along with a process for linearizing them.

The idea of learning radiance fields in a different color space seems simple. But the impacts of this small change are profound across multiple metrics and methods of evaluation.

2 Related and Prior Work

Novel View Synthesis (NeRFs) Novel view synthesis aims to construct scene representations capable of rendering new viewpoints from a set of images and camera poses. Early methods relied on dense sampling and direct pixel interpolation [9] [10]. Recent deep learning approaches have shifted towards volumetric scene representations [2, 13, 52]. Neural Radiance Fields (NeRF) [21] pioneered the use of neural volumetric scene representations optimized through gradient descent on a rendering loss. Subsequent extensions enhanced NeRF’s capabilities, addressing varying lighting conditions [16], incorporating depth [6, 10, 52], utilizing time-of-flight data [2], and integrating semantic segmentation [53].

A recent innovation in NeRF is the use of a 3D bilateral grid to model per-frame camera processing such as gain or color balance changes representable using a 3x4 affine transformation of the color space [51]. The BiLaRF implementation is publicly available and we use it as the basis for our analysis of color spaces. However, we actually turn off the 3-D bilateral grid used during training to avoid clipping, simplifying the model.

RawNeRF [22] has demonstrated the potential of incorporating raw HDR data into NeRF frameworks. By training on raw linear data instead of tone-mapped sRGB images, RawNeRF can handle high dynamic range scenes and produce HDR novel views. This approach has shown remarkable robustness to high noise levels, effectively functioning as a multi-image denoiser for wide-baseline static scenes. We show that using log RGB doesn’t require higher bit depths to achieve high quality results on darker sequences.

Logarithmic RGB (log RGB) representations have been explored in various computer vision applications due to their illumination-invariant properties. In particular, differences in log space are ratios in linear space and are not sensitive to overall intensity. Log RGB has been utilized for illumination-invariant feature extraction [30], material prior computation [25], and the creation of illumination-invariant color spaces [15] [8] [17]. These techniques have proven valuable in tasks such as skin lesion detection [22] [24], color constancy [29], intrinsic image decomposition [12], and shadow removal [18]. The log transformation aligns with the human eye’s logarithmic response to light [26], suggesting potential benefits for computational models.

Recent advancements have explored the use of log RGB in deep neural networks for classical vision tasks. Equivariant networks have been proposed to achieve invariance to input data offsets [4], while other studies have demonstrated the effectiveness of log RGB in providing robustness to intensity and color variations without modifying network architectures [8] [19] [20]. These findings suggest that log RGB representations can enhance the performance and generalization capabilities of deep learning models. However, to date, this has not been evaluated for learning radiance fields.

3 Methods and Experimental Design

3.1 Data Capture and Pre-processing

The videos were captured using a GoPro Hero 13 Black. Each video consists of a 360-degree rotation around a specific object, recorded under varying lighting conditions. In total, approximately 30 videos were collected for training and testing our method. Each video has a duration of 20–30 seconds. To reverse-engineer the GoPro-log color representation, we

referred to the official documentation on GPLog 10-bit encoding and validated our interpretation using a color checker.

We used COLMAP to extract camera poses from video frames. However, for some very dark videos, COLMAP reported a low percentage of successfully registered frames. To address this, we converted the videos from GPLog to TrueLog color space, effectively brightening the images. Among the 13 videos with initially low pose recovery rates, 3 showed no improvement due to scenes being nearly completely dark, another 3 improved but still failed to reach the 100% registration required by the network, and the remaining 7 videos successfully reached full registration after applying COLMAP in the log color space.

3.2 Color Space Analysis

Much of the light from a scene captured by a camera is reflected light that is a result of the interaction of illumination and materials. The Dichromatic Reflection Model, suitable for many natural and human-made materials, divides this into two types of reflection: body reflection and surface reflection [28]. The Bi-Illuminant Dichromatic Reflection Model (BIDR) extends the Dichromatic model and separates the ambient light A from the direct illuminant D , allowing them to be different colors, as occurs in most natural scenes (e.g. blue skylight and yellow sunlight) [17].

The BIDR model represents the body reflection as $I = R_B(A + \gamma D)$, where I is the captured color, R_B is the body reflectance, A is the ambient illumination, D is the direct illuminant, and $\gamma \in [0, 1]$ is the fraction of the direct illuminant incident on the surface point. In linear RGB space, the line or cylinder representing each material under varying illumination starts at a unique point $R_B A$, points in a unique direction defined by $R_B D$, and is of a unique length, also defined by $R_B D$. Therefore, each material requires two colors to define it: $R_B A$ and $R_B D$, and these colors conflate material and illumination.

In log RGB space, however, the reflectance and illumination terms separate.

$$\log I = \log R_B(A + \gamma D) = \log R_B + \log(A + \gamma D) \quad (1)$$

The cylinder representing each material’s body reflection under varying illumination still starts in a unique location. However, the illumination term remains the same across materials lit by the same ambient/direct illumination pair. All of the body cylinders in that set point in the same direction and are the same length. From a representational point of view, in log RGB each material is defined by a single point R_B , all materials in shadow share the offset $\log(A)$ and the body cylinders of all materials move from shadow to lit along the vector defined by $\log(A + \gamma D) - \log(A)$. Therefore, it is possible to represent all body reflection with one material color per pixel, one ambient color for the scene, and one illumination ratio vector for the scene, almost half the information required in linear space.

In sRGB space, which is approximated by a gamma-correction with a factor of 2.2 applied to linear RGB, the body reflection across varying illumination forms a curve, and each material requires learning a different curve of different length. As defining a curve is more complex than defining a line, we hypothesize that sRGB space requires more information to represent a single material under varying illumination than linear RGB.

The GoPro-log encoding is an intermediate space that is closer to linear than to sRGB or log RGB. As such, it likely requires at least as much information as linear RGB space in order to represent the appearance of a material under varying illumination.

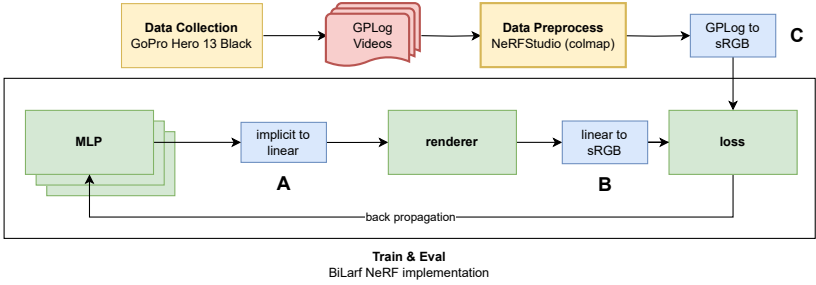


Figure 1: Modified BiLaRF NeRF architecture

3.3 Network Modifications

We use the BiLaRF code base from [51] as the base architecture for all experiments. They all train using a batch size of 16,384, a Charbonnier loss function with significant Total Variation regularization, and a learning rate scheduled to decay from 0.01 to 0.001 over 25,000 epochs. As shown in Figure 1, we added three components: (A) the representation space transform, (B) the sRGB transform, and (C) an input data transform. The representation space transform converts from the implicit representation space learned by the NeRF network into linear RGB space, which is the most appropriate space for the differentiable renderer. The sRGB transform converts the linear RGB signal into sRGB. The input data transform modifies the input data format (e.g. GLog or Linear) into sRGB using the standard sRGB transformation.

We adopt four different color spaces as implicit representations within the network: GoPro-log (GLog), linear RGB, sRGB, and true log RGB (TrueLog). All conversions between color spaces are performed in two steps: the source space is first transformed into linear RGB, followed by a mapping from linear RGB to the target space. Below, we define the forward transformations from linear RGB to each of the three non-linear spaces. The inverse transformations are simply the analytical inverses of these mappings.

GLog: We approximate the logarithmic encoding used by GoPro cameras with (2)

$$GLog(x) = \ln(x \cdot (e - 1) + 1) \quad (2)$$

sRGB: We approximate the sRGB transformation using a gamma correction of 2.22.

$$sRGB(x) = x^{\frac{1}{2.22}} \quad (3)$$

TrueLog: The TrueLog transformation is defined in (4).

$$f(x) = \ln \left(\frac{e^{\ln(\max(255x, 1)) - 1} \cdot 255}{e - 1} \right) \quad (4)$$

The TrueLog transform converts a linear signal in the range $[0, 1]$ to a log signal in the range $[0, 1]$. To constrain the output of this function to the range $[0, 1]$ and avoid mathematical instability, we apply the further normalization step:

$$TrueLog(x) = \frac{f(x) - f(x_{\min})}{f(x_{\max}) - f(x_{\min})} \quad (5)$$

Video ID	$PSNR_{GPLog}$	$PSNR_{linear}$	$PSNR_{sRGB}$	$PSNR_{TrueLog}$	$\Delta PSNR$
GX010032	27.47	25.4	25.35	34.43	9.08
GX010055	34.55	34.56	34.57	39.54	4.97
GX010059	33.85	33.89	33.87	37.19	3.32
GX010108	26.11	24.11	24.83	27.55	2.72
GX010650	26.43	26.42	26.43	28.31	1.88
GX010101	29.72	29.73	29.74	30.9	1.16
GX010025	41.54	41.54	41.54	41.54	1.04
GX010090	31.08	31.06	31.08	32.09	1.01
GX010097	27.23	27.43	27.44	28.29	0.85
GX010103	28.19	28.18	28.2	28.84	0.64

Table 1: Overall PSNR results of captured videos (sampled)

where $x_{\min} = \frac{1}{255}$ and $x_{\max} = 1$ correspond to the lower and upper bounds of the input domain.

For all of our experiments we turned off the `bilarf_grid` flag and set the padding to zero to ensure the intermediate network results were in the range $[0, 1]$ to facilitate the representation space transformations. The bilateral grid can cause clipping to occur, losing information and corrupting the implicit log representation of the network.

4 Experiments and Results

Overall Quality: We assess the quality of our network using PSNR, color renderings, and depth map visualizations. To balance reconstruction quality and training efficiency, we empirically found that 5000 training iterations provides consistently good results. To ensure robustness, we ran each experiment at least five times and report the best PSNR.

The results in Table 1 are uniformly sampled from a total of 28 processed videos and ordered by the difference $PSNR_{TrueLog} - PSNR_{sRGB}$ from highest to lowest. They demonstrate consistent improvements across diverse scenes. The most significant gain is observed in one of the darkest recordings. As shown in the color renderings in Figure 2 and depth map visualizations Figure 3, the TrueLog representation preserves scene details more effectively than other color spaces under low-light conditions. Note, in particular, the reconstruction of the objects on the table in video GX010032.

Even in cases with minimal PSNR improvement, such as video GX010103 in Figure 2, the TrueLog renderings reveal more visual detail, particularly in the foliage, compared to the blurrier outputs from other color spaces. The corresponding depth maps in Figure 3 column GX010103 appear largely similar across all color spaces, indicating that the perceptual advantage of TrueLog in this case is primarily in appearance rather than geometry.

Illumination Sensitivity and Dynamic Range: To study the effect of illumination, we recorded videos of the same scene at different times of day to simulate varying lighting conditions (video IDs: GX010090 to GX010109). Each video was processed using all four color spaces. If an exception occurred during training, the experiment was rerun to ensure valid results. To quantify brightness, we computed average luminance from a randomly sampled 25% of the frames, using the Y weighting: $I = 0.299 \cdot R + 0.587 \cdot G + 0.114 \cdot B$

Figure 4 plots the PSNR improvement of TrueLog over sRGB as a function of scene brightness. We observe that TrueLog provides significant gains in extremely low or high illumination conditions, while improvements are more modest under well-balanced lighting. Notably, in one of the darkest scenes (GX010109), the TrueLog representation reveals much finer details in both the RGB rendering in Figure 2 and the depth map in Figure 3 compared to other color spaces, similar to the types of results reported by RAWNeRF.



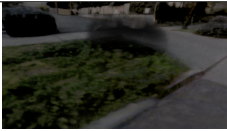












Color Space	GX010032	GX010103	GX010109
GPLog			
linear			
sRGB			
TrueLog			
Ground Truth			

Figure 2: RGB renderings and ground truth of three videos

Robustness: We conducted experiments on nine videos using four different color spaces. For each video–color space pair, we ran the training process 10 times and report the mean, minimum, maximum, and standard deviation of the PSNR values. Representative results of four of the videos are summarized in Table 2, with the remainder in the supplemental.

While some color spaces exhibit instability, the TrueLog representation consistently demonstrates equal or superior robustness across diverse videos and scenes. Notably, in darker videos such as GX010098, GX010099, and GX010109, TrueLog yields lower standard deviations, indicating more stable performance compared to the other color spaces.

Effect of Training Iterations: We analyzed how training duration affects model performance across color spaces by training two representative videos (GX010106 and GX010405) using seven iteration settings: 500, 1000, 2000, 5000, 10000, 20000, and 25000. For one video, we trained to 30000 iterations but observed negligible change (~ 0.01 dB PSNR).

Figure 5 demonstrates distinct convergence patterns across color spaces. TrueLog outperforms others throughout the training process, with stable improvement across iterations. GPLog maintains relatively stable performance, but the linear and sRGB color spaces exhibit instability in the second dataset (GX010405), with degradation after initial peaks.

Our findings suggest that TrueLog provides the optimal balance between computational efficiency and reconstruction quality, with particularly strong performance emerging at approximately 5000 iterations. This represents an important efficiency finding, as it indicates that high-quality results can be achieved with fewer iterations than previously assumed.

Specialness of Log: We evaluated alternative scaled logarithmic conversions across four videos to examine the uniqueness of TrueLog. Specifically, we evaluated (6) using $k =$

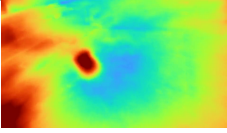
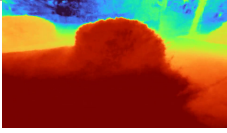
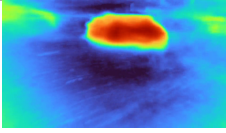
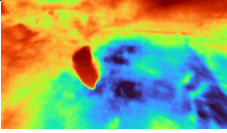
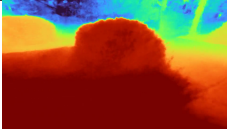
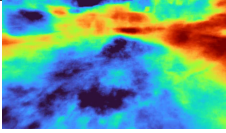
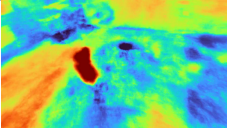
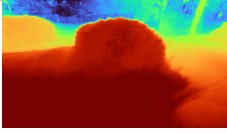
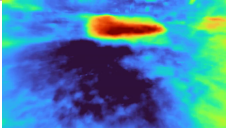
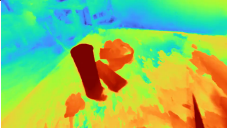
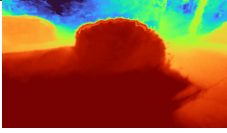
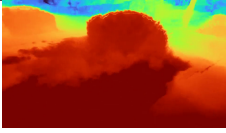
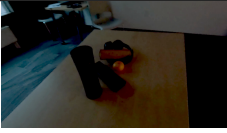


Color Space	GX010032	GX010103	GX010109
GLog			
linear			
sRGB			
TrueLog			
Ground Truth			

Figure 3: Depth map renderings of three videos

255×100 (Log100) and $k = 255 \times 0.1$ (Log01).

$$\text{ScaledLog}(x) = \frac{\log(kx + 1)}{\log(k + 1)} \quad (6)$$

As shown in Table 3, TrueLog consistently outperforms similar logarithmic transformations with different scaling coefficients. The performance gap demonstrates that log space is special for this task, and it does not generalize to other similar transformations.

Network Compactness Analysis: We evaluated compactness using the BiLaRF architecture by reducing MLP widths from 256 to 16 neurons (Figure 6 A). TrueLog consistently outperformed other color spaces by 2-3 dB PSNR, demonstrating stronger scaling with increased width. Grid encoder size reduction experiments from 8192 to 256) revealed that




Figure 4: PSNR improvement of TrueLog over sRGB as a function of scene brightness

Video ID	Color	StdDev	Avg	Min	Max
GX010090	GPLog	0.01	31.08	31.08	31.09
	linear	0.01	31.06	31.05	31.07
	sRGB	0.01	31.07	31.05	31.09
	TrueLog	0.02	32.09	32.06	32.12
GX010098	GPLog	4.15	24.98	13.95	26.78
	linear	0.03	27.52	27.48	27.56
	sRGB	0.11	27.42	27.19	27.5
	TrueLog	0.04	28.31	28.25	28.38
GX010099	GPLog	0.07	25.56	25.45	25.67
	linear	2.55	22.95	17.63	25.42
	sRGB	4.09	22.60	11.36	24.75
	TrueLog	0.04	27.87	27.83	27.94
GX010109	GPLog	0.26	26.45	25.85	26.68
	linear	2.01	23.42	20.11	25.41
	sRGB	2.83	21.83	16.55	25.01
	TrueLog	0.03	29.07	29.03	29.13

Table 2: Consistency of training results across representation spaces




Figure 5: PSNR values across training iterations for videos (A) GX010106 & (B) GX010405.

while GPLog, linear, and sRGB maintained around 28 dB performance across all configurations, TrueLog scaled more effectively, reaching optimal performance (around 31 dB) at size 2048 (Figure 6 B). Despite longer training times and less refined outputs compared to BiLaRF, TrueLog consistently produced sharper, less blurry results in both Mip-NeRF [9] and Robust-NeRF [27] (Figures 7) [23]. This result shows the potential generalization of TrueLog representation to other NeRF implementations.

5 Discussion and Summary

Our experiments show that using a log representation space during NeRF training yields consistent improvement over other color spaces, including linear, sRGB, and GPLog. Quantitatively, TrueLog improves PSNR by an average of 10% across various scenes for the same data and network. Qualitatively, it produces sharper edges, clearer visualizations, and more

Video	TrueLog	Log100	Log01
GX010032	38.06	26.41	26.67
GX010099	27.83	25.27	12.23
GX010109	29.05	24.93	27.03
GX010405	31.07	27.75	27.93

Table 3: PSNR results comparing TrueLog against two scaled log conversions





Figure 6: Network complexity versus PSNR for video GX010405. (A) shows the impact of changes in network width, (B) shows the impact of changes in the grid encoder size.



Mip-NeRF

Robust NeRF

Figure 7: Video GX010405 trained with different NeRF approaches

accurate depth maps, particularly in challenging lighting conditions. While the BiLaRF framework provides a strong foundation for efficient training and fast convergence, incorporating log-space representations further enhances the network’s performance and stability.

For future work, we plan to extend our investigation of log-space generalization to a broader range of NeRF architectures and tasks. Although preliminary tests on mip-NeRF and robust-NeRF suggest that TrueLog maintains its advantages, exploring other approaches could provide additional insights. Additionally, being able to directly visualize the network’s output in its native color space before sRGB conversion would be valuable for analyzing how different representations influence learning dynamics.

References

- [1] August 2025. URL <https://www.motioncamapp.com>.
- [2] Benjamin Attal, Eliot Laidlaw, Aaron Gokaslan, Changil Kim, Christian Richardt, James Tompkin, and Matthew O’Toole. Torf: Time-of-flight radiance fields for dynamic scene view synthesis. *arXiv preprint arXiv:2109.15271*, 2021.
- [3] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. *CVPR*, 2022.
- [4] Marco Cotogni and Claudio Cusano. Offset equivariant networks and their applications. *Neurocomputing*, 502:110–119, 2022.

- [5] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. *arXiv preprint arXiv:2107.02791*, 2021.
- [6] G. D. Finlayson, M. S. Drew, and L. Cheng. Intrinsic images by entropy minimization. In T. Pajdla and J. Matas, editors, *Proc. of European Conf. on Computer Vision*, LNCS 3023, pages 582–595, 2004.
- [7] John Flynn, Ivan Neulander, James Philbin, and Noah Snavely. Deepstereo: Learning to predict new views from the world’s imagery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5515–5524, 2016.
- [8] Brian Funt and Ligeng Zhu. Laplacian of logarithm as illumination-invariant input space. In *Proc. of 30th Color and Imaging Conference*, 2022.
- [9] Steven J Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F Cohen. The lumigraph. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 43–54. ACM, 1996.
- [10] Yoonwoo Jeong, Seokjun Ahn, Christopher Choy, Anima Anandkumar, Minsu Cho, and Jaesik Park. Self-calibrating neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5846–5855, 2021.
- [11] Marc Levoy and Pat Hanrahan. Light field rendering. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 31–42. ACM, 1996.
- [12] Yuanliu Liu, Ang Li, Zejian Yuan, Badong Chen, and Nanning Zheng. Consistency-aware shading orders selective fusion for intrinsic image decomposition. *ArXiv*, abs/1810.09706, 2018.
- [13] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *ACM Transactions on Graphics (TOG)*, 38(4):1–14, 2019.
- [14] Ali Madooei, Mark S. Drew, Maryam Sadeghi, and M. Stella Atkins. Intrinsic melanin and hemoglobin colour components for skin lesion malignancy detection. *Med Image Comput Assist Interv*, 15, 2012.
- [15] John A. Marchant and Christine M. Onyango. Shadow-invariant classification for scenes illuminated by daylight. *J. of the Optical Society of America A*, 17(11), November 2000.
- [16] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7210–7219, 2021.
- [17] Bruce A Maxwell, Richard M Friedhoff, and Casey A Smith. A bi-illuminant dichromatic reflection model for understanding images. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.

- [18] Bruce A. Maxwell, Casey A. Smith, Maan Qraitem, Ross Messing, Spencer Whitt, Nicolas Thien, and Richard M. Friedhoff. Real-time physics-based removal of shadows and shading from road surfaces. In *CVPR Workshop on Autonomous Driving*. IEEE / CVF, 2019.
- [19] Bruce A. Maxwell, Sumegha Singhanian, Heather Fryling, and Haonan Sun. Log rgb images provide invariance to intensity and color balance variation for convolutional networks. In *British Machine Vision Conference*, 2023.
- [20] Bruce A. Maxwell, Sumegha Singhanian, Avnish Patel, Rahul Kumar, Heather Fryling, Sihan Li, Haonan Sun, Ping He, and Zewen Li. Logarithmic Lenses: Exploring Log RGB Data for Image Classification. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17470–17479, Los Alamitos, CA, USA, June 2024. IEEE Computer Society. doi: 10.1109/CVPR52733.2024.01654. URL <https://doi.ieeecomputersociety.org/10.1109/CVPR52733.2024.01654>.
- [21] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. URL <http://arxiv.org/abs/2003.08934v2>.
- [22] Ben Mildenhall, Peter Hedman, Ricardo Martin-Brualla, Pratul P Srinivasan, and Jonathan T Barron. 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*, pages 16190–16199, 2022.
- [23] Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, Peter Hedman, Ricardo Martin-Brualla, and Jonathan T. Barron. MultiNeRF: A Code Release for Mip-NeRF 360, Ref-NeRF, and RawNeRF, 2022. URL <https://github.com/google-research/multinerf>.
- [24] Luisa F. Polanía, Raja Bala, Ankur Purwar, Paul Matts, and Martin Maltz. Skin chromophore estimation from mobile selfie images using constrained independent component analysis. *Electronic Imaging*, 14, 2020.
- [25] Jerome Put, Nick Michiels, and Philippe Bekaert. Material-specific chromaticity priors. In *British Machine Vision Conference*, 2016.
- [26] Albert Rose. The sensitivity performance of the human eye on an absolute scale. *Journal of the Optical Society of America*, 28(2):196–208, 1948.
- [27] Sara Sabour, Suhani Vora, Daniel Duckworth, Ivan Krasin, David J. Fleet, and Andrea Tagliasacchi. Robustnerf: Ignoring distractors with robust losses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20626–20636, June 2023.
- [28] Steven A. Shafer. Using color to separate reflection components. *Color Research Applications*, 10:210–218, 1985.
- [29] Wu Shi, Chen Change Loy, and Xiaoou Tang. Deep specialized network for illuminant estimation. In *European Conf. on Computer Vision*, 2016.

- [30] X. Wang, G. Doretto, T. Sebastian, J. Rittscher, and P. Tu. Shape and appearance context modeling. In *International Conference on Computer Vision*, 2007.
- [31] Yuehao Wang, Chaoyi Wang, Bingchen Gong, and Tianfan Xue. Bilateral guided radiance field processing. *ACM Transactions on Graphics (TOG)*, 43(4):1–13, 2024.
- [32] Yi Wei, Shaohui Liu, Yongming Rao, Wang Zhao, Jiwen Lu, and Jie Zhou. Nerfin-gmvs: Guided optimization of neural radiance fields for indoor multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5610–5619, 2021.
- [33] Shuai Feng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew J Davison. In-place scene labelling and understanding with implicit scene representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15838–15847, 2021.
- [34] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *ACM Transactions on Graphics (TOG)*, 37(4):1–12, 2018.