

NeRF in the Dark: High Dynamic Range View Synthesis from Noisy Raw Images

Ben Mildenhall Peter Hedman Ricardo Martin-Brualla Pratul P. Srinivasan Jonathan T. Barron
 Google Research

Abstract

Neural Radiance Fields (NeRF) is a technique for high quality novel view synthesis from a collection of posed input images. Like most view synthesis methods, NeRF uses tonemapped low dynamic range (LDR) as input; these images have been processed by a lossy camera pipeline that smooths detail, clips highlights, and distorts the simple noise distribution of raw sensor data. We modify NeRF to instead train directly on linear raw images, preserving the scene’s full dynamic range. By rendering raw output images from the resulting NeRF, we can perform novel high dynamic range (HDR) view synthesis tasks. In addition to changing the camera viewpoint, we can manipulate focus, exposure, and tonemapping after the fact. Although a single raw image appears significantly more noisy than a postprocessed one, we show that NeRF is highly robust to the zero-mean distribution of raw noise. When optimized over many noisy raw inputs (25-200), NeRF produces a scene representation so accurate that its rendered novel views outperform dedicated single and multi-image deep raw denoisers run on the same wide baseline input images. As a result, our method, which we call RawNeRF, can reconstruct scenes from extremely noisy images captured in near-darkness.

1. Introduction

View synthesis methods, such as neural radiance fields (NeRF) [39], typically use tonemapped low dynamic range (LDR) images as input and directly reconstruct and render new views of a scene in LDR space. This poses no issues for scenes that are well-lit and do not contain large brightness variations, since they can be captured with minimal noise using a single fixed camera exposure setting. However, this precludes many common capture scenarios: images taken at nighttime or in any but the brightest indoor spaces will have poor signal-to-noise ratios, and scenes with regions of both daylight and shadow have extreme contrast ratios that require high dynamic range (HDR) to represent accurately.

Our method, RawNeRF, modifies NeRF to reconstruct the scene in linear HDR color space by supervising directly on noisy raw input images. This bypasses the lossy postpro-

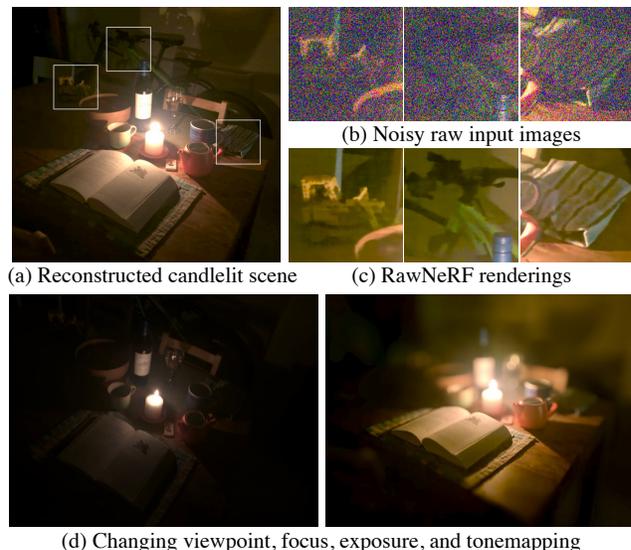


Figure 1. By jointly optimizing a single scene representation over many input images, NeRF is surprisingly robust to high levels of image noise. We exploit this fact to train RawNeRF directly on completely unprocessed HDR linear raw images. In this nighttime scene lit only by a single candle (a), RawNeRF can extract details from the noisy raw data that would have been destroyed by post-processing (b, c). RawNeRF recovers full HDR color information, enabling HDR view synthesis tasks such as changing focus and exposure for rendered novel views. The resulting renderings can be retouched like any raw photograph: here we show (d, left) a dark all-in-focus exposure with a simple global tonemap and (d, right) a brighter, synthetically refocused exposure postprocessed by HDRNet [18]. See our supplementary video for more results.

cessing that cameras apply to compress dynamic range and smooth out noise in order to produce visually palatable 8-bit JPEGs. By preserving the full dynamic range of the raw inputs, RawNeRF enables various novel HDR view synthesis tasks. We can modify the exposure level and tonemapping algorithm applied to rendered outputs and even create synthetically refocused images with accurately rendered bokeh effects around out-of-focus light sources.

Beyond these view synthesis applications, we show that training directly on raw data effectively turns RawNeRF into a multi-image denoiser capable of reconstructing

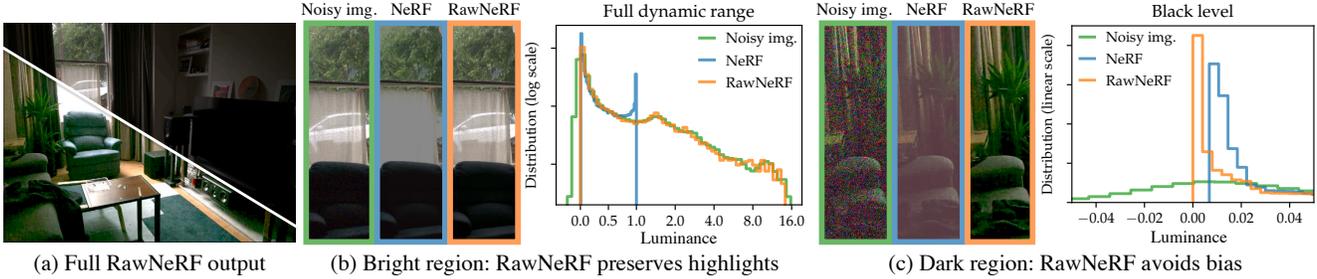


Figure 2. Failure modes of NeRF on a daytime indoor scene. (a) Here we show two exposures ($24\times$ apart) of a full RawNeRF output rendering, both passed through a global tonemapping curve. Training NeRF with postprocessed LDR images, as done in prior work, (b) prevents it from recovering bright highlights clipped above at 1, resulting in the missing car outside the window, and (c) corrupts the per-pixel noise distribution such that NeRF recovers incorrect colors due to the nonlinear tonemap and clipping below at 0, particularly in dark regions around the plant and sofa. In contrast, RawNeRF trains directly on HDR linear raw images and correctly recovers the radiance distribution in both extremely bright and extremely dark parts of the scene.

scenes captured in near-darkness (Figure 1). The standard camera postprocessing pipeline (*e.g.*, HDR+ [21]) corrupts the simple noise distribution of raw data, introducing significant bias in order to reduce variance and produce an acceptable output image. Feeding these images into NeRF thus produces a biased reconstruction with incorrect colors, particularly in the darkest regions of the scene (see Figure 2 for an example). We instead exploit NeRF’s ability to reduce variance by aggregating information across frames, demonstrating that it is possible for RawNeRF to produce a clean reconstruction from many noisy raw inputs.

Unlike typical video or burst image denoising methods, RawNeRF assumes a static scene and expects camera poses as input. Provided with these extra constraints, RawNeRF is able to make use of 3D multiview consistency to average information across nearly *all* of the input frames at once. Since our captured scenes each contain 25-200 input images, this means RawNeRF can remove more noise than feed-forward single or multi-image denoising networks that only make use of 1-5 input images for each output.

In summary, we make the following contributions:

1. We propose a method for training RawNeRF directly on raw images that can handle high dynamic range scenes as well as noisy inputs captured in the dark.
2. We show that RawNeRF outperforms NeRF on noisy real and synthetic datasets and is a competitive multi-image denoiser for wide-baseline static scenes.
3. We showcase novel view synthesis applications made possible by our linear HDR scene representation (varying exposure, tonemapping, and focus).

2. Related Work

RawNeRF combines concepts from several areas of research. We build upon NeRF as a baseline for high quality view synthesis, bring in ideas from low level image processing to optimize NeRF directly on noisy raw data, and

take inspiration from uses of HDR in computer graphics and computational photography to showcase new applications made possible by an HDR scene reconstruction. We briefly cover relevant prior work across each of these areas.

2.1. Novel view synthesis

Novel view synthesis is the task of using a set of input images and their camera poses to reconstruct a scene representation capable of rendering novel views. When the input images are densely sampled, it is possible to use direct interpolation in pixel space for view synthesis [20, 33]. A more feasible capture scenario is to capture more widely spaced inputs and use a “proxy” geometry (*e.g.*, a reconstructed triangle mesh) to reproject and combine colors from the input images, using either a heuristic [6] or learned [22, 42, 43] blending function.

Recent work on applying deep learning to view synthesis has focused on volumetric rather than mesh-based scene representations [16, 36, 57]. NeRF [39] directly optimizes a *neural* volumetric scene representation to match all input images using gradient descent on a rendering loss. Various extensions have improved NeRF’s robustness to varying lighting conditions [37] or added supervision with depth [25, 51, 52], time-of-flight data [1], or semantic segmentation labels [56]. As of yet, no approach has extended NeRF to work with high dynamic range color data. Some previous view synthesis methods trained using LDR data jointly solve for per-image scaling factors to account for inconsistent lighting or miscalibration between cameras [29, 36]. ADOP [44] supervises with LDR images and solves for exposure through a differentiable tonemapping step to approximately recover HDR, but does not focus on robustness to noise or supervision with raw data.

2.2. Denoising

Early neural denoising approaches mostly focused on denoising sRGB images synthetically corrupted with additive

white Gaussian noise [54]. In 2017, Plötz and Roth [41] established a real raw image denoising benchmark, which showed that these deep denoisers failed to generalize beyond the synthetic data used during training and were outperformed by standard non-learned methods, such as BM3D [11]. Subsequent work on both single [5, 9] and multi-image [8, 19, 38, 53] denoising demonstrated the benefits of training networks to operate directly on noisy raw input data. Modern cellphone camera pipelines perform a robust averaging of multiple noisy input frames in the raw domain [21], though they typically cannot afford to employ deep networks due to speed and power limitations.

Another line of research investigated whether denoisers could be trained using *only* noisy data when no corresponding clean ground truth exists. Noise2Noise [32] demonstrated this was possible given a dataset of pairs of independent noisy observations of the same image, an insight Ehret *et al.* [14] applied to denoise videos by aligning consecutive noisy frames. Various followups to Noise2Noise proposed modified network architectures allowing supervision with a dataset of single noisy images [4, 30, 31]. Sheth *et al.* [46] showed that this paradigm could be applied to train a denoiser using a *single* noisy video, including an application to raw video data. Similarly, RawNeRF is optimized over a single set of images to both denoise and recover the 3D structure of the captured scene.

Concurrent work by Pearl *et al.* [40] modified the feed-forward IBRNet view synthesis method [49] to achieve competitive burst denoising results; this approach produces results more quickly for each new scene but can only process a limited number of input images (per output pixel) due to memory constraints.

2.3. Applications of raw and HDR image data

Computational photography The value of working directly with raw data has long been noted by digital photographers due to the fact that its preservation of dynamic range allows for maximum postprocessing flexibility, letting users modify exposure, white balance, and tonemapping after the fact. Many works have tried to automate this process by using heuristics or machine learning to map directly from raw data to postprocessed LDR images [7, 9, 18, 23].

Another line of work focuses on recovering HDR images from LDR inputs. This concept was pioneered by Debevec and Malik [12], who used a stack of aligned LDR images taken at different exposures to recover and invert the camera’s nonlinear response curve. Current approaches apply machine learning to produce HDR outputs from single [15] or multiple misaligned [24] LDR inputs, either recovering or hallucinating detail in clipped highlights.

Synthetic defocus Many modern cellphones include a postprocessing option to add synthetic defocus blur after

capture [48]. Though it is possible to accurately simulate defocus using a thin-lens model [10] or real multi-element camera lens [28] using ray tracing, most machine learning models use a much faster approximate rendering model, predicting a depth map and applying a depth-varying blur kernel to each discretized depth layer [2, 47]. Performing this blur in HDR space is critical to achieving the correct appearance of defocused bright highlights (known as “bokeh”), as demonstrated by Zhang *et al.* [55].

3. Noisy Raw Input Data

NeRF [39] takes postprocessed low dynamic range (LDR) sRGB color space images as input. This works well when using clean, noise-free images with minimal contrast. However, all real images contain some level of noise, and each step in the camera postprocessing pipeline corrupts this distribution in a certain way. Here we briefly outline the simplified pipeline stages relevant to our method (complete descriptions of the full camera pipeline can be found in recent work [13, 26]).

Raw camera measurements When capturing an image, the number of photons hitting a pixel on the camera sensor is converted to an electrical charge, which is recorded as a high bit-depth digital signal (typically 10 to 14 bits). These values are offset by a “black level” to allow for negative measurements due to noise. After black level subtraction, the signal is a noisy measurement y_i of a quantity x_i proportional to the expected number of photons arriving while the shutter is open. This noise results from both the physical fact that photon arrivals are a Poisson process (“shot” noise) and noise in the readout circuitry that converts the analog electrical signal to a digital value (“read” noise). The combined shot and read noise distribution can be well modeled as a Gaussian whose variance is an affine function of its mean [17]; importantly, this implies that the distribution of the error $y_i - x_i$ is zero mean.

Color filter demosaicking Color cameras contain a Bayer color filter array in front of the image sensor such that each pixel’s spectral response curve measures either red, green or blue light. The pixel color values are typically arranged in 2×2 squares containing two green pixels, one red, and one blue pixel (known as a Bayer pattern), resulting in “mosaicked” data. To generate a full-resolution color image, the missing color channels are interpolated using a demosaicking algorithm [34]. This interpolation correlates noise spatially, and the checkerboard pattern of the mosaic leads to different noise levels in alternating pixels.

Color correction and white balance The spectral response curves for each color filter element vary between different cameras, and a color correction matrix is used to

convert the image from this camera-specific color space to a standardized color space. Additionally, because human perception is robust to the color tint imparted by different light sources, cameras attempt to account for this tint (*i.e.*, make white surfaces appear RGB-neutral white) by scaling each color channel by an estimated white balance coefficient. These two steps are typically combined into a single linear 3×3 matrix transform, which further correlates the noise between color channels.

Gamma compression and tonemapping Humans are able to discern smaller relative differences in dark regions compared to bright regions of an image. This fact is exploited by sRGB gamma compression, which optimizes the final image encoding by clipping values outside $[0, 1]$ and applying a nonlinear curve to the signal that dedicates more bits to dark regions at the cost of compressing bright highlights. In addition to gamma compression, tonemapping algorithms can be used to better preserve contrast in high dynamic range scenes (where the bright regions are several orders of magnitude brighter than the darkest) when the image is quantized to 8 bits [12, 21].

In a slight abuse of terminology, we will refer both of these steps jointly as “tonemapping” in the rest of the paper, indicating the process by which linear HDR values are mapped to nonlinear LDR space for visualization. We will refer to signals before tonemapping as high dynamic range (HDR) and signals after as low dynamic range (LDR). Of all postprocessing operations, tonemapping has the most drastic effect on the noise distribution: clipping completely discards information in the brightest and darkest regions, and after the non-linear tonemapping curve the noise is no longer guaranteed to be Gaussian or even zero mean.

4. RawNeRF

A neural radiance field (NeRF) [39] is a neural network based scene representation that is optimized to reproduce the appearance of a set of input images with known camera poses. The resulting reconstruction can then be used to render novel views from previously unobserved poses. NeRF’s multilayer perceptron (MLP) network takes 3D position and 2D viewing direction as input and outputs volume density and color. To render each pixel in an output image, NeRF uses volume rendering to combine the colors and densities from many points sampled along the corresponding 3D ray.

Standard NeRF takes clean, low dynamic range (LDR) sRGB color space images with values in the range $[0, 1]$ as input. Converting raw HDR images to LDR images (*e.g.*, using the pipeline described in Section 3) has two significant consequences:

1. Detail in bright areas is lost when values are clipped from above at one, and detail across the image is

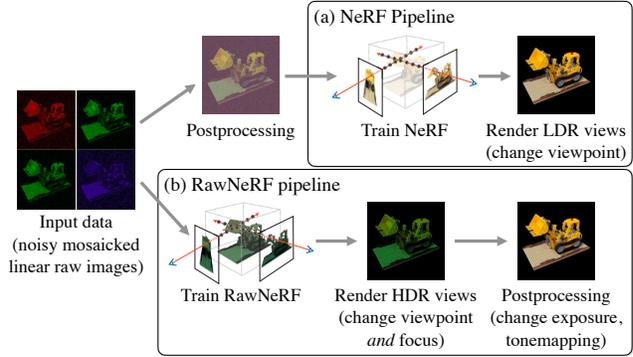


Figure 3. The standard NeRF training pipeline (a) takes in LDR images that have been sent through a camera processing pipeline, reconstructing the scene and rendering new views in LDR color space. As such, its renderings are effectively already postprocessed and cannot be significantly retouched. In contrast, our method RawNeRF (b) modifies NeRF to train directly on linear raw HDR input data. The resulting scene representation produces novel views that can be edited like any raw photograph.

compressed by the tonemapping curve and subsequent quantization to 8 bits.

2. The per-pixel noise distribution becomes biased (no longer zero-mean) after passing through a nonlinear tonemapping curve and being clipped from below at zero.

The goal of RawNeRF is to make use of this information rather than discarding it, optimizing NeRF directly on linear raw input data in HDR color space (Figure 3). In Section 5, we will show that reconstructing NeRF in raw space makes it much more robust to noisy inputs and allows for novel HDR view synthesis applications. First, we detail the changes required to make NeRF work with raw data.

4.1. Loss function

Since the color distribution in an HDR image can span many orders of magnitude, a standard L2 loss applied in HDR space will be completely dominated by error in bright areas and produce an image that has muddy dark regions with low contrast when tonemapped (see Figure 4). Instead, we apply a loss that more strongly penalizes errors in dark regions to align with how human perception compresses dynamic range. One way to achieve this is by passing both the rendered estimate \hat{y} and noisy observed intensity y through a tonemapping curve ψ before the loss is applied:

$$L_\psi(\hat{y}, y) = \sum_i (\psi(\hat{y}_i) - \psi(y_i))^2. \quad (1)$$

However, in low-light raw images the observed signal y is heavily corrupted by zero-mean noise, and a nonlinear tonemap will introduce bias that changes the noisy signal’s

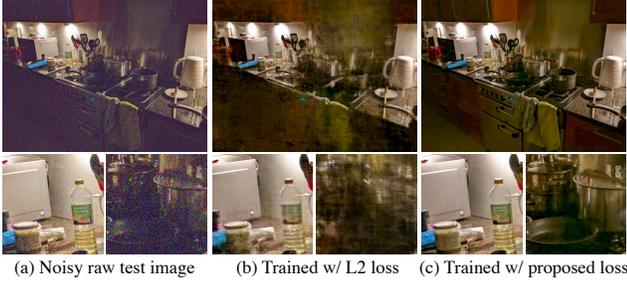


Figure 4. This challenging scene (a) has a 7000 \times ratio between its 90th and 10th raw color percentiles. (b) When faced with such high-contrast inputs, the standard L2 loss from NeRF manages to recover the bright parts of the scene but produces poor results in darker regions, which becomes particularly apparent after LDR tonemapping. (c) Our proposed loss (4), reweighted according to the gradient of a log tonemap curve, successfully reconstructs all parts of the scene. Both rendered images are tonemapped using HDR+ [21] for visualization.

expected value ($E[\psi(y)] \neq \psi(E[y])$). In order for the network to converge to an unbiased result [32], we instead use a weighted L2 loss of the form

$$L(\hat{y}, y) = \sum_i w_i (\hat{y}_i - y_i)^2. \quad (2)$$

We can approximate the tonemapped loss (1) in this form by using a linearization of the tone curve ψ around each \hat{y}_i :

$$\tilde{L}_\psi(\hat{y}, y) = \sum_i [\psi'(\text{sg}(\hat{y}_i))(\hat{y}_i - y_i)]^2, \quad (3)$$

where $\text{sg}(\cdot)$ indicates a stop-gradient that treats its argument as a constant with zero derivative, preventing it from influencing the loss gradient during backpropagation. We find that a “gradient supervision” tone curve $\psi(z) = \log(y + \epsilon)$ with $\epsilon = 10^{-3}$ produces perceptually high quality results with minimal artifacts, implying a loss weighting term of $\psi'(\text{sg}(\hat{y}_i)) = (\text{sg}(\hat{y}_i) + \epsilon)^{-1}$ and final loss

$$\tilde{L}_\psi(\hat{y}, y) = \sum_i \left(\frac{\hat{y}_i - y_i}{\text{sg}(\hat{y}_i) + \epsilon} \right)^2. \quad (4)$$

This corresponds exactly to the relative MSE loss used to achieve unbiased results when training on noisy HDR path-tracing data in Noise2Noise [32]. The curve ψ is proportional to the μ -law function used for range compression in audio processing, and has previously been applied as a tonemapping function when supervising a network to map from a burst of LDR images to an HDR output [24].

4.2. Variable exposure training

In scenes with very high dynamic range, even a 10-14 bit raw image may not be sufficient for capturing both bright

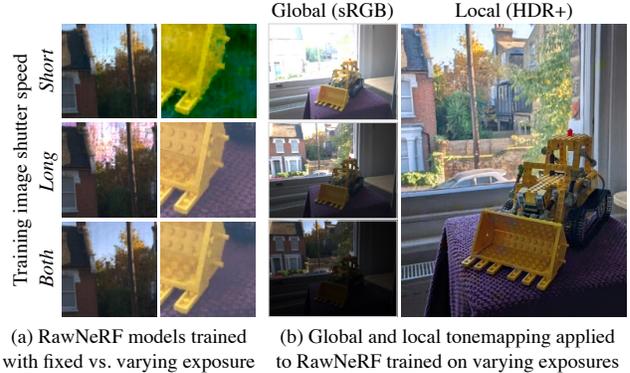


Figure 5. A fixed shutter speed is not sufficient for capturing the full dynamic range in scenes with extreme brightness variation. (a) For example, this scene requires variable exposure capture to avoid either poor quality in dark indoor regions or blown-out sky highlights. Only a RawNeRF model optimized using both short and long exposures recovers the full dynamic range. (b) This brightness variation is too high to visualize in a single image using a simple global sRGB gamma curve, requiring a more sophisticated local tonemapping algorithm (e.g., HDR+ postprocessing [21]).

Figure 5. A fixed shutter speed is not sufficient for capturing the full dynamic range in scenes with extreme brightness variation. (a) For example, this scene requires variable exposure capture to avoid either poor quality in dark indoor regions or blown-out sky highlights. Only a RawNeRF model optimized using both short and long exposures recovers the full dynamic range. (b) This brightness variation is too high to visualize in a single image using a simple global sRGB gamma curve, requiring a more sophisticated local tonemapping algorithm (e.g., HDR+ postprocessing [21]).

and dark regions in a single exposure. This is addressed by the “bracketing” mode included in many digital cameras, where multiple images with varying shutter speeds are captured in a burst, then merged to take advantage of the bright highlights preserved in the shorter exposures and the darker regions captured with more detail in the faster exposures.

We can similarly take advantage of variable exposures in RawNeRF (Figure 5). Given a sequence of images I_i with exposure times t_i (and all other capture parameters held constant), we can “expose” RawNeRF’s linear space color output to match the brightness in image I_i by scaling it by the recorded shutter speed t_i . In practice, we find that varying exposures cannot be precisely aligned using shutter speed alone due to sensor miscalibration (see supplement). To correct for this, we add a learned per-color-channel scaling factor for each unique shutter speed present in the set of captured images, which we jointly optimize along with the NeRF network. The final RawNeRF “exposure” given a output color \hat{y}_i from the network is then $\min(\hat{y}_i^c \cdot t_i \cdot \alpha_{t_i}^c, 1)$, where c indexes color channels, and $\alpha_{t_i}^c$ is the learned scaling factor for shutter speed t_i and channel c (we constrain $\alpha_{t_{\max}}^c = 1$ for the longest exposure). We clip from above at 1 to account for the fact that pixels saturate in overexposed regions. This scaled and clipped value is passed to the previously described loss (Equation 4).

4.3. Implementation details

Our implementation is based on the mip-NeRF [3] codebase, which improves upon the positional encoding used in the original NeRF method. Please see that paper for further details on the MLP scene representation and volumetric

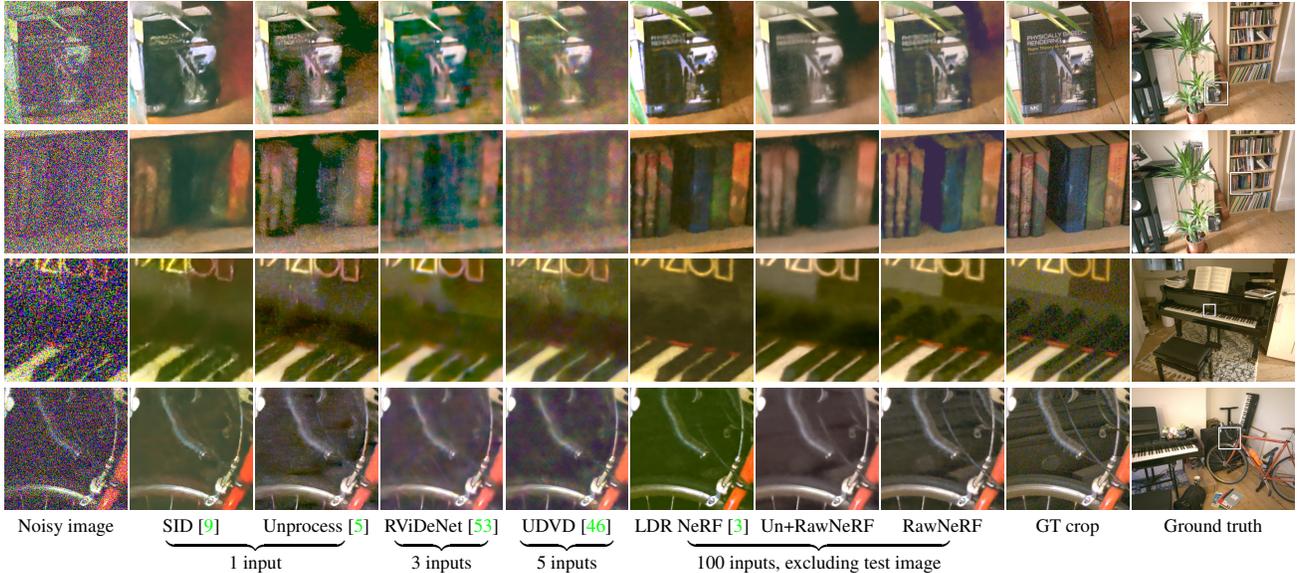


Figure 6. Example postprocessed and color-aligned patches from our real denoising dataset. RawNeRF produces the most detailed output in each case. All deep denoising methods (columns 2-5) receive the noisy test image as input, whereas NeRF variants (columns 6-8) perform both novel view synthesis and denoising.

Method	Num. inputs	Raw PSNR \uparrow	Affine-aligned sRGB		
			PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Noisy input	-	54.38	10.24	0.035	0.733
SID [9]	1	-	21.62	0.525	0.547
Unprocess [5]	1	70.80	23.02	0.491	0.489
RViDeNet [53]	3	68.29	22.20	0.516	0.545
UDVD [46]	5	70.68	22.75	0.514	0.507
LDR NeRF [3]	$N - 1$	-	19.43	0.518	0.544
Un+RawNeRF	$N - 1$	67.99	23.35	0.531	0.507
RawNeRF	$N - 1$	67.20	23.53	0.536	0.501

Table 1. We compare RawNeRF’s denoising performance to various single and multi-image denoisers and NeRF ablations. Despite only being optimized on a single scene and never having seen even a noisy version of the test view, RawNeRF achieves results competitive with deep denoising methods trained on large image datasets. RawNeRF also outperforms NeRF trained on LDR sRGB images (LDR NeRF) and an ablation where RawNeRF’s inputs have been denoised using “Unprocess” (Un+RawNeRF).

rendering algorithm. Our only network architecture change is to modify the activation function for the MLP’s output color from a sigmoid to an exponential function to better parameterize linear radiance values. We use the Adam optimizer [27] with batches of 16k random rays sampled across all training images and a learning rate decaying from 10^{-3} to 10^{-5} over 500k steps of optimization.

We find that extremely noisy scenes benefit from a regularization loss on volume density to prevent partially transparent “floater” artifacts. We apply a loss on the variance of the weight distribution used to accumulate color values along the ray during volume rendering; please see the sup-

plement for details.

As our raw input data is mosaicked, it only contains one color value per pixel. We only apply the loss to the active color channel for each pixel, such that optimizing NeRF effectively demosaics the input images. Since any resampling steps will effect the raw noise distribution, we do not undistort or downsample the inputs, and instead train using the full resolution mosaicked images (usually 12MP for our scenes). To achieve this, we use camera intrinsics to account for radial distortion when generating rays. We use full resolution postprocessed JPEG images to calculate camera poses as COLMAP [45] does not support raw images.

5. Results

We present results exploring two consequences of supervising NeRF with raw HDR data. First, we show that RawNeRF is surprisingly robust to high levels of noise, to the extent that it can act as a competitive multi-image denoiser when applied to wide-baseline images of a static scene. Second, we demonstrate the HDR view synthesis applications enabled by recovering a scene representation that preserves high dynamic range color values.

5.1. Denoising

Recent years have seen an increasing focus on developing deep learning methods for denoising images directly in the raw linear domain [5, 9]. This effort has expanded to include multi-image denoisers that can be applied to burst images or video frames [8, 46, 53]. These multi-image denoisers typically assume that there is a relatively small

amount of motion between frames, but that there may be large amounts of object motion within the scene. When nearby frames can be well aligned, these methods merge information from similar image patches (typically across 2-8 neighboring images) to outperform single image denoisers.

By comparison, NeRF (and by extension, RawNeRF) optimizes for a single scene reconstruction that is consistent with *all* input images. By specializing to wide-baseline static scenes and taking advantage of 3D multiview information, RawNeRF can aggregate observations from much more widely spaced input images than a typical multi-image denoising method.

Real dataset We collect a real world denoising dataset with 3 different scenes, each consisting of 101 noisy images and a clean reference image merged from stabilized long exposures. The first 100 images are taken handheld across a wide baseline (a standard forward-facing NeRF capture), using a fast shutter speed to accentuate noise. We then capture a stabilized burst of 50-100 longer exposures on a tripod and robustly merge them using HDR+ [21] to create a clean ground truth frame. One additional tripod image taken at the original fast shutter speed serves as a noisy input “base frame” for the deep denoising methods. All images are taken with an iPhone X (2017) or iPhone SE (2020) at 12MP resolution using the wide-angle lens and saved as 12-bit raw DNG files.

Comparisons In Table 1 and Figure 6, we compare RawNeRF’s joint view synthesis and denoising performance to several recent deep single and multi-image denoising methods. Note that all denoisers require the noisy version of the test image as input, whereas RawNeRF and its ablations only require its camera pose.

We focus our comparison on methods explicitly designed to handle raw input images. Chen *et al.* [9] (SID) present a single image denoiser that maps from raw inputs to postprocessed LDR images and is trained on a large dataset of noisy raw and clean postprocessed image pairs collected by the authors. Brooks *et al.* [5] (Unprocess) is a method for training a raw single image denoiser on simulated raw data created from internet image datasets that transfers well to real raw images. RViDeNet [53] trains a raw video denoiser on a combination of Unprocessing-style synthetic data and a new real raw video dataset. Sheth *et al.* [46] (UDVD) present a “self-supervised” method for training a video denoiser only using noisy data, building on ideas from Noise2Noise [32] and blind-spot networks [31]. UDVD provides network weights specifically trained on the raw video dataset from RViDeNet. For all methods, we use publicly available code and pretrained model weights.

We also compare to two ablations of our method. LDR NeRF represents mip-NeRF [3] trained (as usual) in LDR

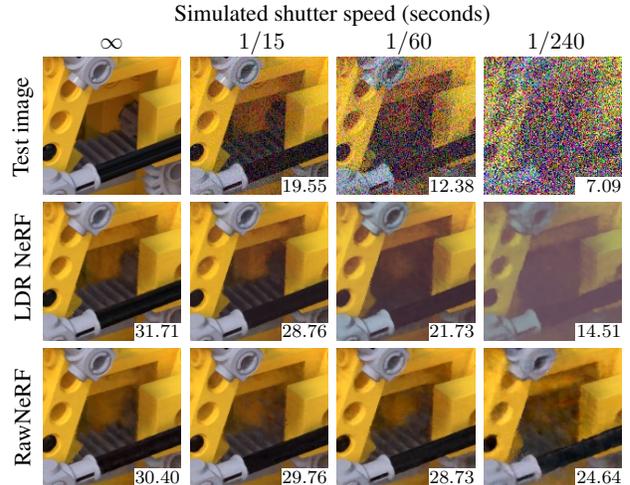


Figure 7. Example patches from the synthetic scene used in Table 2, annotated with sRGB PSNR for each inset. With perfectly clean inputs, training on LDR images is superior, but with any nonzero amount of noise, it is more beneficial to optimize NeRF in raw space, where the noise distribution remains unbiased.

Method	Simulated shutter speed (seconds)						
	∞	1/7	1/15	1/30	1/60	1/120	1/240
Noisy input	-	23.33	19.65	16.03	12.51	9.40	7.18
LDR NeRF	33.16	31.25	29.14	26.10	22.31	18.27	14.87
RawNeRF	32.15	32.11	31.94	31.59	30.94	29.69	27.73

Table 2. We perform an ablation study on a synthetically rendered raw dataset with 120 training images, simulating shot and read noise for 8 different shutter speeds. Here we report PSNR values in LDR sRGB space.

sRGB space on images postprocessed by a minimal sRGB tonemapping pipeline. “Un+RawNeRF” preprocesses the training images using the single image raw denoiser from Brooks *et al.* [5] (“Unprocess”) before training RawNeRF.

All compared methods take mosaicked raw images as input. Every deep denoiser [5, 9, 46, 53] uses the noisy “base frame” as input, and the two multi-image denoising networks [46, 53] also receive the nearest images from the wide-baseline capture (based on camera position). We convert the 12-bit raw input to floating point by normalizing with the white and black levels. Since each method was trained on raw data from a different source, they impart different color tints to the output. So this does not affect metrics, we calculate a per-color-channel affine transform that best matches each method’s raw output to the ground truth raw image. (The exceptions are SID and LDR NeRF, whose sRGB output we match to the postprocessed sRGB ground truth.) Our basic postprocessing pipeline for visualization and computing sRGB metrics is to apply a bilinear demosaic (when necessary), perform white balance/color correction, rescale white level, clip to [0, 1], and apply the sRGB

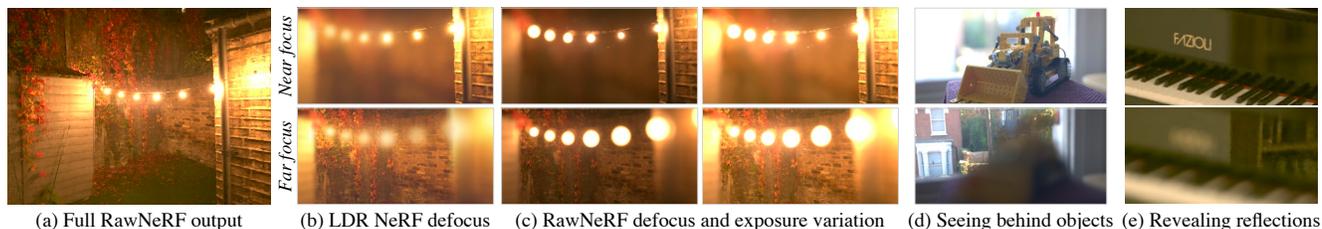


Figure 8. Synthetic defocus examples. In this nighttime garden scene (a), LDR NeRF cannot accurately render defocused bright highlights since it is trained on images that have already been tonemapped and clipped (b). RawNeRF recovers the linear intensity of the light sources such that applying defocus blur produces correctly oversaturated “bokeh balls” (c). Since RawNeRF is optimized for view synthesis from wide-baseline inputs, it can achieve 3D defocus effects not possible with a single image and depth map, such as revealing occluded parts of the background by focusing behind the foreground bulldozer (d) or focusing on the bookshelves reflected above the piano keys (e).

gamma curve. Please see the supplement for details.

Analysis Despite simultaneously performing denoising and novel view synthesis, our method is competitive with all compared deep denoisers (Table 1, Figure 6). We suspect that the multi-image denoisers struggle to make use of the additional frames provided from the wide-baseline capture, as the camera movement is larger than in a typical sub-second burst or video clip. By comparison, RawNeRF, despite lacking any explicitly learned image priors, clean training data, or even a “base frame” input image, produces high quality outputs by combining information from across all input images in its reconstruction. Despite the fact that LDR NeRF is directly trained to minimize mean-squared error in sRGB space, RawNeRF achieves significantly better sRGB metrics. We also find that applying a single image denoiser to the inputs before training RawNeRF results in oversmoothed renderings (Un+RawNeRF).

Synthetic noise ablation In Table 2 and Figure 7, we demonstrate the impact of noise level on RawNeRF image quality. For training, we render 120 linear HDR images using the *Lego* scene from NeRF [39], borrowing color correction, white balance, and noise parameters from our iPhone captures’ EXIF metadata to “unprocess” this data into raw space [5]. Since the renderings have a large amount of empty space, we report sRGB PSNR on the object only, by using the provided alpha masks (otherwise error from the background pixels heavily penalizes LDR NeRF). Even in this synthetic setting free from camera miscalibration issues, we can clearly observe the color bias and loss of detail caused by training LDR NeRF on postprocessed noisy data.

5.2. HDR view synthesis applications

Modifying exposure and tonemapping Figures 1, 2, 4, 5, and 8 include examples of varying the exposure level and tonemapping algorithm for images output by RawNeRF, which exist in linear HDR space and can thus be postprocessed like a raw photo from a digital camera. Please see

our supplement and video for many more examples.

Synthetic defocus Given a full 3D model of a scene, physically-based renderers accurately simulate camera lens defocus effects by tracing rays refracted through each lens element [28], but this process is extremely computationally expensive. A reasonably convincing and much cheaper solution is to apply a varying blur kernel to different depth layers of the scene and composite them together [2, 48]. In Figure 8, we apply this synthetic defocus rendering model to sets of RGBA depth layers precomputed from trained RawNeRF models (similar to a multiplane image [57]). As shown by Zhang *et al.* [55], recovering linear HDR color is critical for achieving the characteristic oversaturated “bokeh balls” around defocused bright light sources.

6. Discussion

We have demonstrated the benefits of training NeRF directly on linear raw camera images. However, this modification is not without tradeoffs. Most digital cameras can only save raw images at full resolution with minimal compression, resulting in huge storage requirements when capturing tens or hundreds of images per scene. Our method is also dependent on COLMAP’s [45] robustness for computing camera poses, preventing us from capturing scenes below a certain light level. This could potentially be addressed by jointly optimizing RawNeRF and the input camera poses [35, 50]. Finally, despite its robustness to noise, RawNeRF cannot be considered a general purpose denoiser as it cannot handle scene motion and requires orders of magnitude more computation than a feed-forward network.

Despite these shortcomings, we believe that RawNeRF represents a step toward robust, high quality capture of real world environments. Training on raw images with variable exposure allows us to capture scenes with a much wider dynamic range, and robustness to noise makes reconstructing dark nighttime captures possible. Lifting these constraints greatly increases the fraction of the world that can be reconstructed and explored with photorealistic view synthesis.

References

- [1] Benjamin Attal, Eliot Laidlaw, Aaron Gokaslan, Changil Kim, Christian Richardt, James Tompkin, and Matthew O’Toole. Törf: Time-of-flight radiance fields for dynamic scene view synthesis. *NeurIPS*, 34, 2021. 2
- [2] Jonathan T. Barron, Andrew Adams, YiChang Shih, and Carlos Hernández. Fast bilateral-space stereo for synthetic defocus. *CVPR*, 2015. 3, 8
- [3] Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. Mip-NeRF: A multiscale representation for anti-aliasing neural radiance fields. *ICCV*, 2021. 5, 6, 7
- [4] Joshua Batson and Loic Royer. Noise2self: Blind denoising by self-supervision. *ICML*, 2019. 3
- [5] Tim Brooks, Ben Mildenhall, Tianfan Xue, Jiawen Chen, Dillon Sharlet, and Jonathan T. Barron. Unprocessing images for learned raw denoising. *CVPR*, 2019. 3, 6, 7, 8
- [6] Chris Buehler, Michael Bosse, Leonard McMillan, Steven Gortler, and Michael Cohen. Unstructured lumigraph rendering. *SIGGRAPH*, 2001. 2
- [7] Vladimir Bychkovsky, Sylvain Paris, Eric Chan, and Frédo Durand. Learning photographic global tonal adjustment with a database of input / output image pairs. *CVPR*, 2011. 3
- [8] Chen Chen, Qifeng Chen, Minh Do, and Vladlen Koltun. Seeing motion in the dark. *ICCV*, 2019. 3, 6
- [9] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. *CVPR*, 2018. 3, 6, 7
- [10] Robert L. Cook, Thomas Porter, and Loren Carpenter. Distributed ray tracing. *SIGGRAPH*, 1984. 3
- [11] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *TIP*, 2007. 3
- [12] Paul E. Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. *SIGGRAPH*, 1997. 3, 4
- [13] Mauricio Delbracio, Damien Kelly, Michael S. Brown, and Peyman Milanfar. Mobile computational photography: A tour. *Annual Review of Vision Science*, 2021. 3
- [14] Thibaud Ehret, Axel Davy, Jean-Michel Morel, Gabriele Facciolo, and Pablo Arias. Model-blind video denoising via frame-to-frame training. *CVPR*, 2019. 3
- [15] Gabriel Eilertsen, Joel Kronander, Gyorgy Denes, Rafał Mantiuk, and Jonas Unger. Hdr image reconstruction from a single exposure using deep cnns. *SIGGRAPH*, 2017. 3
- [16] John Flynn, Ivan Neulander, James Philbin, and Noah Snavely. Deepstereo: Learning to predict new views from the world’s imagery. *CVPR*, 2016. 2
- [17] Alessandro Foi, Mejdí Trimeche, Vladimir Katkovnik, and Karen O. Egiazarian. Practical poissonian-gaussian noise modeling and fitting for single-image raw-data. *IEEE Transactions on Image Processing*, 2008. 3
- [18] Michaël Gharbi, Jiawen Chen, Jonathan T Barron, Samuel W Hasinoff, and Frédo Durand. Deep bilateral learning for real-time image enhancement. *SIGGRAPH*, 2017. 1, 3
- [19] Clement Godard, Kevin Matzen, and Matt Uyttendaele. Deep burst denoising. *ECCV*, 2018. 3
- [20] Steven J Gortler, Radek Grzeszczuk, Richard Szeliski, and Michael F Cohen. The lumigraph. *SIGGRAPH*, pages 43–54, 1996. 2
- [21] Samuel W. Hasinoff, Dillon Sharlet, Ryan Geiss, Andrew Adams, Jonathan T. Barron, Florian Kainz, Jiawen Chen, and Marc Levoy. Burst photography for high dynamic range and low-light imaging on mobile cameras. *SIGGRAPH Asia*, 2016. 2, 3, 4, 5, 7
- [22] Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow. Deep blending for free-viewpoint image-based rendering. *SIGGRAPH Asia*, 2018. 2
- [23] Yuanming Hu, Hao He, Chenxi Xu, Baoyuan Wang, and Stephen Lin. Exposure: A white-box photo post-processing framework. *SIGGRAPH*, 2018. 3
- [24] Nima Khademi Kalantari and Ravi Ramamoorthi. Deep high dynamic range imaging of dynamic scenes. *SIGGRAPH*, 36(4), 2017. 3, 5
- [25] Jun-Yan Zhu Kangle Deng, Andrew Liu and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. *arXiv:2107.02791*, 2021. 2
- [26] Hakki Can Karaimer and Michael Brown. A software platform for manipulating the camera imaging pipeline. volume 9905, pages 429–444, 10 2016. 3
- [27] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015. 6
- [28] Craig Kolb, Don Mitchell, and Pat Hanrahan. A realistic camera model for computer graphics. *SIGGRAPH*, 1995. 3, 8
- [29] Georgios Kopanas, Julien Philip, Thomas Leimkühler, and George Drettakis. Point-based neural rendering with per-view optimization. *Computer Graphics Forum*, 2021. 2
- [30] Alexander Krull, Tim-Oliver Buchholz, and Florian Jug. Noise2void-learning denoising from single noisy images. *CVPR*, 2019. 3
- [31] Samuli Laine, Tero Karras, Jaakko Lehtinen, and Timo Aila. High-quality self-supervised deep image denoising. *NeurIPS*, 2019. 3, 7
- [32] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2noise: Learning image restoration without clean data. *ICML*, 2018. 3, 5, 7
- [33] Marc Levoy and Pat Hanrahan. Light field rendering. *SIGGRAPH*, 1996. 2
- [34] Xin Li, Bahadır Gunturk, and Lei Zhang. Image demosaicing: A systematic survey. *Visual Communications and Image Processing*, 2008. 3
- [35] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. BARF: Bundle-adjusting neural radiance fields. *ICCV*, 2021. 8
- [36] Stephen Lombardi, Tomas Simon, Jason Saragih, Gabriel Schwartz, Andreas Lehrmann, and Yaser Sheikh. Neural volumes: Learning dynamic renderable volumes from images. *SIGGRAPH*, 2019. 2
- [37] Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections. *CVPR*, 2021. 2

- [38] Ben Mildenhall, Jonathan T. Barron, Jiawen Chen, Dillon Sharlet, Ren Ng, and Robert Carroll. Burst denoising with kernel prediction networks. *CVPR*, 2018. 3
- [39] 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. *ECCV*, 2020. 1, 2, 3, 4, 8
- [40] Naama Pearl, Tali Treibitz, and Simon Korman. Nan: Noise-aware nerfs for burst-denoising. In *CVPR*, 2022. 3
- [41] Tobias Plötz and Stefan Roth. Benchmarking denoising algorithms with real photographs. *CVPR*, 2017. 3
- [42] Gernot Riegler and Vladlen Koltun. Free view synthesis. *ECCV*, 2020. 2
- [43] Gernot Riegler and Vladlen Koltun. Stable view synthesis. *CVPR*, 2021. 2
- [44] Darius Rückert, Linus Franke, and Marc Stamminger. Adop: Approximate differentiable one-pixel point rendering. *arXiv preprint arXiv:2110.06635*, 2021. 2
- [45] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. *CVPR*, 2016. 6, 8
- [46] Dev Yashpal Sheth, Sreyas Mohan, Joshua Vincent, Ramon Manzorro, Peter A. Crozier, Mitesh M. Khapra, Eero P. Simoncelli, and Carlos Fernandez-Granda. Unsupervised deep video denoising. *ICCV*, 2021. 3, 6, 7
- [47] Pratul P. Srinivasan, Rahul Garg, Neal Wadhwa, Ren Ng, and Jonathan T. Barron. Aperture supervision for monocular depth estimation. *CVPR*, 2018. 3
- [48] Neal Wadhwa, Rahul Garg, David E. Jacobs, Bryan E. Feldman, Nori Kanazawa, Robert Carroll, Yair Movshovitz-Attias, Jonathan T. Barron, Yael Pritch, and Marc Levoy. Synthetic depth-of-field with a single-camera mobile phone. *SIGGRAPH*, 2018. 3, 8
- [49] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul Srinivasan, Howard Zhou, Jonathan T. Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *CVPR*, 2021. 3
- [50] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. NeRF—: Neural radiance fields without known camera parameters. *arXiv:2102.07064*, 2021. 8
- [51] Yi Wei, Shaohui Liu, Yongming Rao, Wang Zhao, Jiwen Lu, and Jie Zhou. Nerfingmvs: Guided optimization of neural radiance fields for indoor multi-view stereo. *ICCV*, 2021. 2
- [52] Christophehr Choy Animashree Anandkumar Minsu Cho Yoonwoo Jeong, Seokjun Ahn and Jaesik Park. Self-calibrating neural radiance fields. *ICCV*, 2021. 2
- [53] Huanjing Yue, Cong Cao, Lei Liao, Ronghe Chu, and Jingyu Yang. Supervised raw video denoising with a benchmark dataset on dynamic scenes. *CVPR*, 2020. 3, 6, 7
- [54] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Transactions on Image Processing*, 2017. 3
- [55] Xuaner Zhang, Kevin Matzen, Vivien Nguyen, Dillon Yao, You Zhang, and Ren Ng. Synthetic defocus and look-ahead autofocus for casual videography. *SIGGRAPH*, 2019. 3, 8
- [56] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew Davison. In-place scene labelling and understanding with implicit scene representation. *ICCV*, 2021. 2
- [57] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *SIGGRAPH*, 2018. 2, 8