

Perceptual Quality Assessment of NeRF and Neural View Synthesis Methods for Front-Facing Views

H. Liang¹  T. Wu¹  P. Hanji¹  F. Banterle²  H. Gao¹  R. Mantiuk¹  and C. Öztireli¹ 

¹University of Cambridge, UK

²ISTI-CNR, Italy

Abstract

Neural view synthesis (NVS) is one of the most successful techniques for synthesizing free viewpoint videos, capable of achieving high fidelity from only a sparse set of captured images. This success has led to many variants of the techniques, each evaluated on a set of test views typically using image quality metrics such as PSNR, SSIM, or LPIPS. There has been a lack of research on how NVS methods perform with respect to perceived video quality. We present the first study on perceptual evaluation of NVS and NeRF variants. For this study, we collected two datasets of scenes captured in a controlled lab environment as well as in-the-wild. In contrast to existing datasets, these scenes come with reference video sequences, allowing us to test for temporal artifacts and subtle distortions that are easily overlooked when viewing only static images. We measured the quality of videos synthesized by several NVS methods in a well-controlled perceptual quality assessment experiment as well as with many existing state-of-the-art image/video quality metrics. We present a detailed analysis of the results and recommendations for dataset and metric selection for NVS evaluation.

CCS Concepts

• **Computing methodologies** → *Image-based rendering; Image and video acquisition; Perception;*

1. Introduction

Synthesizing photorealistic novel views of a complex scene from a sparse set of RGB images is a fundamental challenge in image-based rendering. Various representations and methods have been developed to accurately model the image formation process and handle complex geometry, materials, and lighting conditions [CW93, SGHS98, SK00, YKM*20, VSJ22, WBF*17, KLR*22, KKLD23]. More recently, Neural View Synthesis (NVS) via implicit representations has shown promising results. In particular, methods such as Neural Radiance Field (NeRF) [MST*20] and its successors [BMV*22a, BMV*22b, FYT*22, SSC22, WWG*21, WPYS21] have attracted great interest due to their outstanding fidelity and robustness. However, assessing the performance of contemporary NVS methods is not straightforward as it is closely tied to the final applications. Specifically, NVS methods are increasingly developed in immersive and realistic AR/VR applications, it is thus crucial for the methods to synthesize high-quality free-viewpoint videos with unnoticeable artifacts to human users.

The current protocol for comparing NVS methods involves computing image quality metrics, such as PSNR, SSIM [WBSS04b] and LPIPS [ZIE*18], on a subset of hold-out views for a few scenes. Even dedicated benchmarks [DLBD*23, WWL*23] follow the same evaluation protocol. Since the main objective of NVS methods is to offer interactive exploration of novel views,

we argue that those methods should be evaluated on *video sequences* rather than individual sparse views, ideally in a subjective quality evaluation experiment. Thus, we identify two key limitations in existing evaluation protocols. ❶ They rely exclusively on image quality metrics, which can be problematic because these metrics may not correlate well with subjective judgments, especially when used for a task they are not designed for [CHM*12, HME*22, PJJ*15]. Since most of the image quality metrics have not been calibrated or validated on the distortions specific to novel view synthesis, their predictions could be too noisy to quantify perceived quality. ❷ The evaluation protocol lacks assessment on video sequences, which can reveal temporal artifacts and subtle distortions, such as flickering or floating ghost images, that are easily noticeable in video but difficult to spot in static images [CAD19, DM20, LAK*16, MDC*21a]. This issue is compounded by the limited nature of commonly used NVS datasets, which do not have reference videos for evaluating NVS methods.

To address these problems, we first collect two new datasets with front-facing views: a Lab dataset captured using 2D gantry in well-controlled laboratory conditions, and a Fieldwork dataset, captured in-the-wild with the help of either a gimbal or a camera slider (Section 3). We mainly focus on the front-facing setup because this is applicable to more NVS methods (e.g., NeX [WPYS21]) and has significant applications in free-viewpoint video capture for AR/VR applications. Each captured scene contains several sparse training

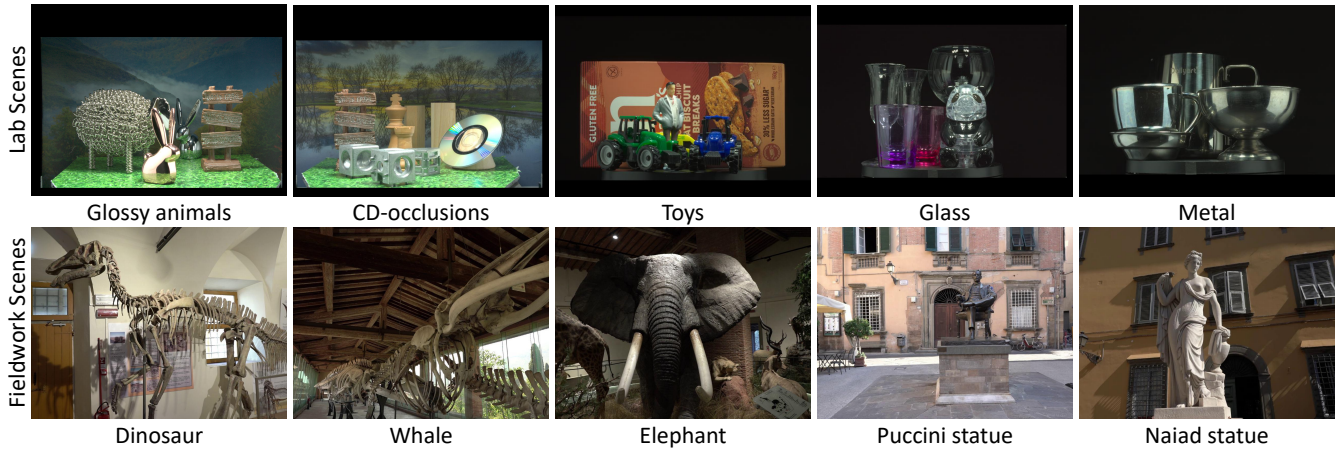


Figure 1: A subset of scenes from our collected Lab (first row) and Fieldwork (second row) datasets. Our datasets include both controlled laboratory and in-the-wild scenes, each with reference video sequences. We selected a diverse range of objects composed of various materials such as wood, marble, metal, and glass etc.

views and a reference test video intended to evaluate NVS methods. We use the two new datasets together with the popular LLFF dataset [MSOC*19] to reconstruct the video sequences by 8 NVS methods and 2 variants of generalizable NVS methods. The output videos of these methods are then evaluated by human participants in a subjective quality assessment experiment (Section 4). The results of that experiment serve as ground-truth scores for testing how well the existing image and video quality metrics can predict the perceptual performance of NVS methods (Section 6).

In summary, the main contributions of this work are:

- Two new datasets with front-facing views and video references for full-reference evaluation of synthesized videos,
- A subjective quality assessment of videos synthesized by 8 NVS methods (and two generalizable NeRF variants) measured via a perceptual quality assessment experiment,
- An objective evaluation of existing image/video quality metrics on synthesized videos to assess how well these metrics correlate with subjective quality.
- A thorough analysis of metrics that elucidates the limitation in the current NVS evaluation protocol and reveals the crucial need for video assessment and video metrics. In light of this, we provide practical recommendations to enhance the efficacy of evaluation processes.

Experimental results and the dataset can be found at the project web page[†].

2. Related Work

2.1. Quality assessment of NVS methods

Most works on NVS methods and NVS benchmarks [MST*20, DLBD*23, MSOC*19, WPYS21, FBD*19, LGZL*20, AAB23]

[†] Project web page: <https://www.cl.cam.ac.uk/research/rainbow/projects/perceptualnerf/>

evaluate on sparse hold-out views using image quality metrics [FTS*23, ANAM*20, MDC*21b, SB13, LAK*16]. An exception is the Light Field Benchmark [YKGB*20], where light field interpolation methods were evaluated on video sequences. On the contrary, our focus is on assessing the perceptual quality of NVS methods and evaluating how well current objective metrics can predict subjective quality. Such subjective benchmarks have previously motivated and advanced other areas such as tone mapping [LCTS05, EWMU13], image compression [AMR*16], and single-image HDR [HME*22]. To the best of our knowledge, we present the first study on perceptual assessment of NVS methods and hope that our study will similarly inspire improvements that better meet the needs of human users.

2.2. NVS Datasets

NVS methods are typically evaluated using synthetic and real-world datasets with sparse views [MSOC*19, WPYS21, FBD*19, LGZL*20, YLL*20, JDV*14, MRAQ23]. The NeRF synthetic dataset [MST*20] consists of 8 inwards facing scenes rendered with blender [Com18], each containing 200 test images rendered at viewpoints located spirally at the upper hemisphere around the object. The LLFF dataset [MSOC*19] is a forward-facing dataset of real scenes, but with very sparse test views. The DTU [JDV*14] Stereo dataset is also widely used to evaluate novel view synthesis performance, but its captured views are too sparse to create a continuous video. Recently, De Luigi et al. [DLBD*23] set up a resource-efficient system to capture 360-degree dense views of various objects, but only for simple objects in a controlled lab environment and without video references. RealEstate10K [ZTF*18] only contains a fly-through style video per scene, where objects are only visible in a few frames and from limited angles, making it inaccessible to split into distinct train and test frames. Tanks and Temples [KPZK17] is designed for large-scale scenes and can introduce unfairness in comparison. The presence of moving people/objects in these datasets also makes them unsuitable for many NVS methods. In contrast, our dataset is the first forward-facing dataset that

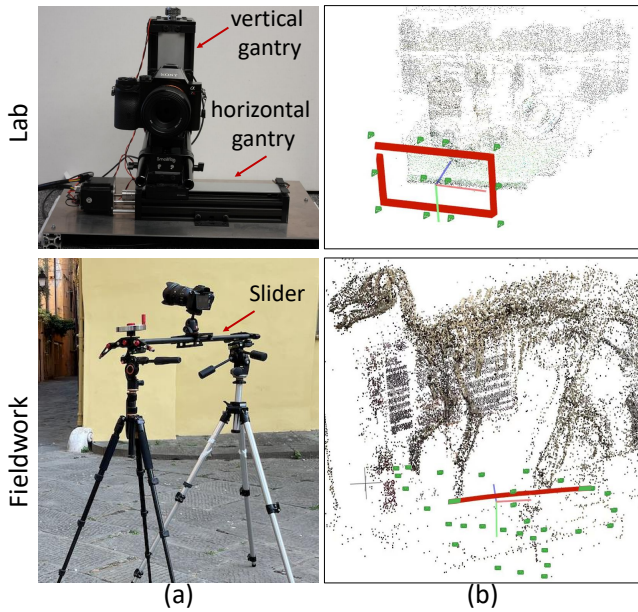


Figure 2: The camera rig (a) and camera poses (b) that are used to capture scenes for the Lab (top) and Fieldwork (bottom) datasets. On the right, the green dots represent training camera poses and the red dots represent the poses of the reference video frames used for testing.

captures scenes with reference videos in both laboratory and fieldwork environments, with accurately calibrated poses for each video frame.

3. Forward-Facing Video Dataset

To evaluate NVS methods on video rather than individual views, we collected two new datasets: Lab, captured using a 2D gantry in a laboratory with controlled lighting and background; and Fieldwork, captured in-the-wild, consisting of both indoor and outdoor scenes. Both datasets were captured with Sony A7RIII. Images of selected scenes from both datasets are shown in Figure 1.

3.1. Lab Dataset

Capture Setup The Lab dataset was captured in our laboratory using a 2D gantry (upper-left of Figure 2), which allowed horizontal and vertical movement of a camera. To minimize the amount of noise and avoid saturated pixels, we captured each view with a RAW image stack consisting of 2 exposures at constant ISO. The RAW image stacks were merged into an HDR image using an estimator that accounts for the photon noise [HZM20]. All images were color-corrected using a reference white point and cropped to 4032×3024 px. To map linear images to display-referred units, gamma was applied ($\gamma = 2.2$).

A sparse set of training views were taken on a uniform grid, shown as green dots in upper-right of Figure 2. The training views cover a horizontal range of 100 mm and a vertical range of 80 mm. The reference videos consisted of 300 to 500 frames, captured in a

rectangular camera motion, as shown by red dots in the figure. The camera traveled about 0.6 mm between each frame. Since we only consider a view interpolation task (no extrapolation), the reference frames were positioned within the range of the training views.

Scenes The lab scenes were placed inside a box of $30\text{cm} \times 41.5\text{cm} \times 38\text{cm}$ for capturing. As illustrated in the first row of Figure 1, they were designed to cover a wide range of objects with various materials, including glass, metal, wood, ceramic, and plastics. The layout of the objects was selected to introduce occlusions and to offer a good range of depth, which would fit within the depth-of-field of the camera. The dataset contains challenging view-dependent effects, such as diffraction on the surface of a CD-ROM, specular reflections from metallic and ceramic surfaces, and transparency of the glass. Six scenes were captured in this dataset.

Pose Estimation For accurate pose estimation, 4 sets of 4 AprilTag markers were placed in each corner of the scene. The camera positions were selected to ensure that all markers were visible in each view and the images were later cropped to remove the markers. By detecting the position of AprilTags [Ols11], we obtained camera poses with standard camera calibration methods [Zha00]. According to our pose estimation results, we got an acceptable mean re-projection error of 0.2174 px across all scenes.

3.2. Fieldwork Dataset

Our in-the-wild Fieldwork dataset was captured in both outdoor city areas and indoor rooms of a public museum[‡]. Typically, such scenes are challenging due to complex backgrounds, occlusions, and uncontrolled illuminations.

Capture Setup Different from the Lab dataset, we captured video sequences instead of individual images for the Fieldwork scenes. The video sequences were captured with resolution 1920×1080 px and framerate 30 fps. To reduce camera shake, we used either a DJI RS3 gimbal or a 90 cm manual slider, which was fixed on two tripods, see lower-left of Figure 2. For each scene, we captured several video sequences with different trajectories. One of these sequences, whose trajectory is well within the bounds of the scene, is selected as the test sequence. The bottom-right of Figure 2 shows the test sequence of one scene from the Fieldwork dataset (red dots). Images for training are sampled from the remaining videos (green dots). We also moved the first and last 15 frames from the test video sequence to the training set to ensure that the test views can always be interpolated from training views. In total, we have around 120 frames reserved as test views.

Scenes The second row of Figure 1 shows selected examples of the captured scenes, which cover both indoor and outdoor scenarios with a high variability of materials including wood, marble, window glasses, metals, etc. and complex geometries such as a whale skeleton, posing challenging scenarios for NVS methods. Nine scenes were captured in this dataset.

Pose Estimation We employed COLMAP [SF16] to perform joint calibration of camera poses for both the training and testing frames, so that all the calibrated poses share the same scale with

[‡] Natural History Museum of the University of Pisa

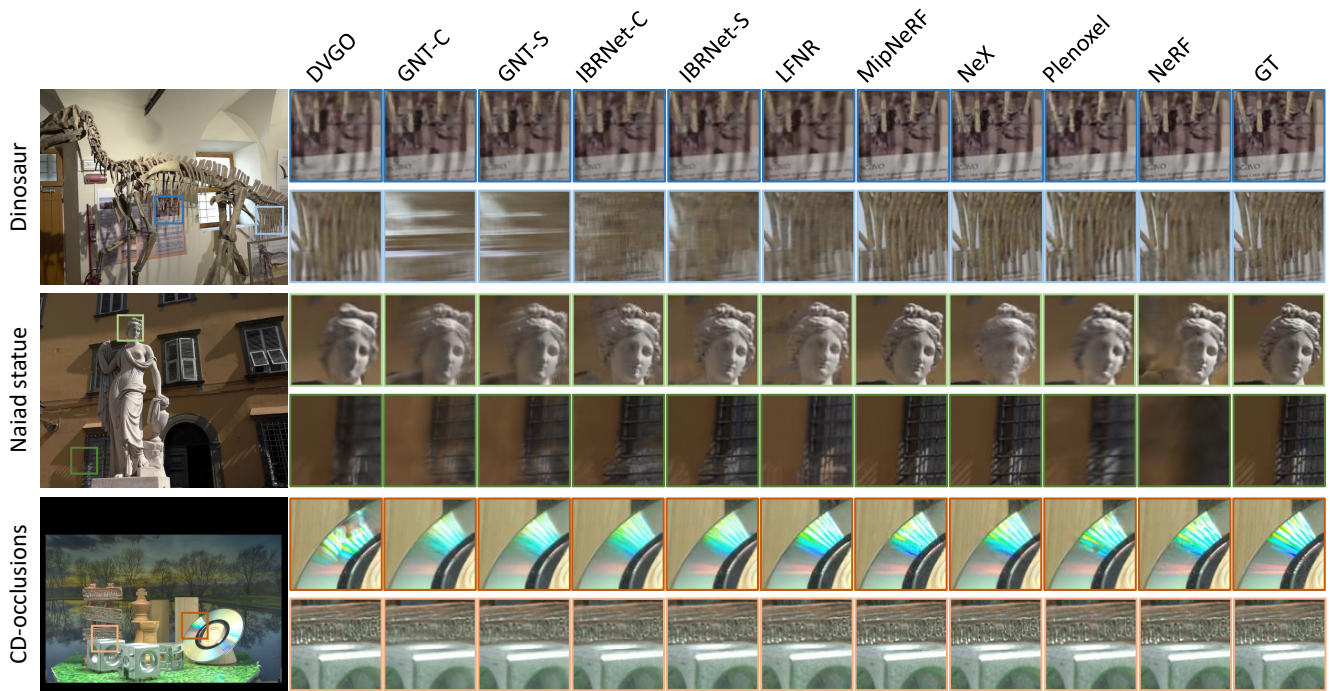


Figure 3: Examples of reconstructions by various NVS methods on selected scenes from Fieldwork dataset (first two rows) and Lab dataset (third row). We only show three scenes due to limited space, please refer to the supplementary for more visual results.

a consistent coordinate system. We used the “OPENCV” camera model, which supports separate x and y focal lengths as well as radial and tangential distortions. We also used COLMAP to undistort the captured images after pose estimation. Our reconstructed camera parameters have a mean reprojection error of 0.5327 px across all scenes.

3.3. Evaluated NVS Methods

We tested ten representative NVS methods (including two generalizable NeRF variants) that encompass a diverse range of models, which feature both explicit and implicit geometric representations, distinct rendering modelings, as well as generalizable and per-scene optimization strategies. NeRF [MST*20] is a neural volumetric representation that excels in image-based scene reconstruction and novel view synthesis. Mip-NeRF [BMV*22a] builds upon NeRF and provides a multiscale representation for anti-aliased view synthesis. DVGO [SSC22] and Plenoxels [FYT*22] use hybrid representations to achieve fast training and rendering. NeX [WPYS21] utilizes multi-plane images and trainable basis functions, which is intended for rendering view-dependent effects in forward-facing scenes. LFNR [SESM22] operates on a light field representation and uses an epipolar constraint to guide the rendering process. IBRNet [WWG*21] and GNT [WCC*22] are both generalizable NeRF models. IBRNet aggregates nearby source views to estimate radiance and density and GNT extends this idea by proposing a unified transformer-based architecture that replaces both multi-view feature aggregation and volume rendering. For

IBRNet and GNT, we tested both their published cross-scene models (labeled as GNT-C and IBRNet-C) and also models fine-tuned on each scene (labeled as GNT-S and IBRNet-S).

We use these methods to reconstruct videos of scenes from both our collected datasets, Lab and Fieldwork, as well as from the popular forward-facing LLFF [MSOC*19] dataset. For a fair comparison between methods, we downsampled images from the Lab dataset by a factor of 4 and cropped images from the Fieldwork dataset, so that they all have the same training image resolution of 1008×756 px. as that in LLFF scene. In this way, we were able to adopt the same training setup (network architecture, training iterations, optimizer, etc.) on LLFF scenes proposed by the respective authors. Please refer to the supplementary materials for more details about the training setup.

4. Subjective Evaluation

To attain precise subjective quality scores of the videos synthesized by the aforementioned NVS methods, we conducted a controlled quality assessment experiment with human participants. We relied on a pairwise comparison experiment, as it has been shown to be more accurate and robust than direct rating methods [POMZ*20].

4.1. Experimental Procedure

We employed pairwise comparison experiments for subjective evaluation. Particularly, in each trial of our experiment, a participant was shown a pair of videos side-by-side on the same display and

was instructed to pick the video of higher quality — “better resembles a natural scene and contains fewer distortions” (exact wording on the briefing form). To reduce the number of comparisons and maximize the information gained from each trial, we used ASAP [MWPO*21], an active sampling method. Participants could press the space bar to view the reference video of the displayed scene (except for LLFF dataset as reference videos were not available). The reference videos were included as one of the compared conditions. Please refer to supplementary for more details about pairwise comparison and ASAP sampling.

4.2. Display and Videos

Our videos were displayed on a 27" Eizo ColorEdge CS2740 4K monitor, which was colorimetrically calibrated to reproduce BT.709 color space with a peak luminance of 200 cd/m². The average viewing distance was 70 cm, restricted by a table in front of the display.

We used 14 scenes from our two datasets and 8 scenes from the LLFF dataset. For the scenes in our dataset, videos were synthesized on the same views as in ground-truth video frames. As LLFF dataset does not have reference videos, we combined 120 frames rendered in a spiral trajectory around the mean pose (similar to other NVS methods). All the video frames were cropped to a resolution of 960×756 px. and then up-scaled to 1920×1512 px. (bilinear filter) so that two videos could be shown side by side on our 4K monitor. The upscaling was necessary to obtain a more realistic and effective resolution of 40 pixels per degree with respect to the original video resolution. Please note that when we evaluate image/video metrics in Section 6, all the computation is also done on up-scaled images/videos to ensure fairness. Each video was between 3 to 15 seconds long, with a framerate of 30 fps. In total, each participant assessed the quality of 22 scenes reconstructed by 10 NVS methods as well as 14 reference videos.

4.3. Participants

We invited 39 volunteers (20 males and 19 females) with normal color vision (confirmed by running the *Ishihara Test*). Each participant completed 4–5 full batches of comparisons scheduled by ASAP [MWPO*21]. The experiment was authorized by an external institutional review board and the participants were rewarded for their participation.

4.4. Subjective Score Scaling and Calculation

We scaled the results of the pairwise comparison and expressed the subjective evaluation score in Just-Objectable-Difference (JOD) units using the Thurstone Case V observer model [POM17]. A difference of 1 JOD unit means that 75% participants preferred one method over another. The model assumes that participants made their selections by assigning a single quality value to each video and approximates this quality by a normally distributed random variable with the same inter- and intra-observer variance.

5. Perceptual Benchmark Results

Figure 4 shows the perceptual preference for different methods averaged across our collected Lab and Fieldwork datasets, as well

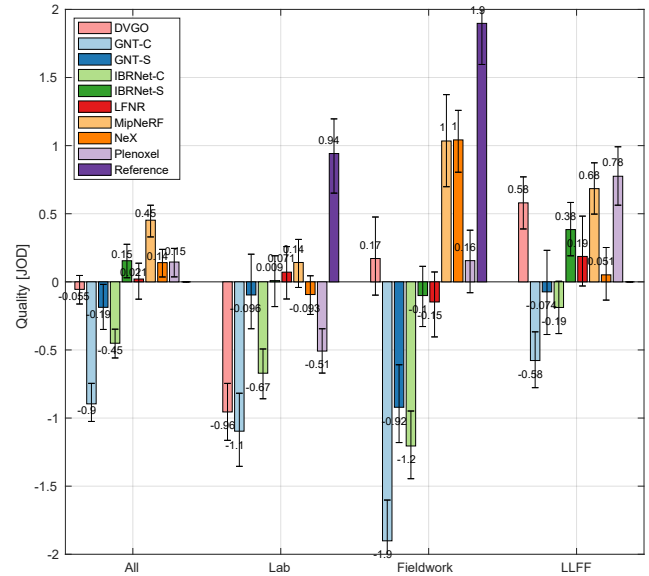


Figure 4: Perceptual preferences for NVS methods. The bars indicate preference in JOD units, relative to the original NeRF method [MST*20], which is at 0 JOD. Negative values indicate that, on average, the method produced less preferable results than NeRF. The error bars show 95% confidence intervals.

as the LLFF dataset. We report both per-dataset performance and the overall performance across all three datasets. To view results on the individual scenes, we refer to Figures 2–4 in the supplementary. The baseline (0 JOD line) in Figure 4 is the original NeRF model [MST*20], so positive JOD values indicate improvement and negative values indicate degradation in quality (on average) with respect to NeRF.

The results on both of our datasets show that despite the impressive performance of NVS methods, their results can still be easily distinguished from the reference videos. There is about 0.85 JOD difference between the best neural rendering methods and the reference; 0.94 vs. 0.14 for Lab, 1.9 vs. 1 for Fieldwork. This indicates that the reference will be selected as better in 70% of the cases across the population. On average, only five out of nine methods produced better results than the original NeRF. It is evident that existing generalizable models require further refinement, as an additional per-scene optimization step is needed to achieve desirable outcomes. It is noteworthy that the discrepancies among the methods are more noticeable in the more challenging Fieldwork dataset, which implies that a challenging dataset is essential to distinguish between methods.

Compared with other models, MipNeRF performs quite well in most scenes, particularly those with high-frequency geometric details (Figure 3, statue’s face and hair in Naiad statue, bones and background poster in Dinosaur, fence in CD-occlusions etc.). In comparison, techniques that lack explicit volume rendering (e.g., LFNR, and GNT) and those with coarse geometric modeling (e.g., NeX) exhibit suboptimal performance in these situations. Nonetheless, LFNR and NeX do provide more natural outcomes

Table 1: The list of evaluated objective metrics.

Metric	Reference required	Video metric	Details
PSNR	✓	✗	Widely used ratio to measure noise relative to the signal in log units
PSNR-L	✓	✗	PSNR computed on image luma values
SSIM [WBSS04a]	✓	✗	Popular quality measure that perceives structural similarity
MS-SSIM [WSB03]	✓	✗	Multi-scale version of SSIM
VIF [SB06]	✓	✗	Natural Scene Statistics (NSS) models on information-theoretic setting
FSIM [ZZMZ11]	✓	✗	Low-level image feature similarity based on the human visual system
LPIPS-VGG [ZIE*18]	✓	✗	Perceptual similarity metric based on deep network of VGG model
LPIPS-ALEX [ZIE*18]	✓	✗	Perceptual similarity metric based on deep network of AlexNet model
DISTS [DMWS20]	✓	✗	Unify texture and structure similarity with deep network
HDR-VDP-3 [MKRH11]	✓	✗	Low-level vision model on HDR images
FLIP [ANAM*20]	✓	✗	Metric that considers HVS, viewing distance and monitor conditions
FovVideoVDP [MDC*21b]	✓	✓	Spatial-temporal metric that accounts for foveation effect
STRRED [SB13]	✓	✓	Hybrid metric measures temporal motion and spatial difference
VMAF [LAK*16]	✓	✓	Support Vector Machine combination of multiple image and video metrics
HDR-VQM [NPDSL15]	✓	✓	Spatial-temporal metric that considers human eye fixation behavior
BRISQUE [MMB12]	✗	✗	Support vector regression trained on IQA dataset
NIQE [MSB13]	✗	✗	Distance between NSS-based features to those from a database
PIQE [VDB*15]	✗	✗	Averaged block-wise distortion estimation

for scenes with complex lighting and specular reflections such as Metal and CD in CD-occlusions; see Figure 3. For certain Fieldwork scenes, techniques founded on multi-view epipolar geometry constraints, such as IBRNet and GNT, tend to fail and exhibit conspicuous artifacts (Dinosaur in Figure 3, Vespa and Giraffe shown in supplementary). This is due to the considerable distance between the test and source views, which renders the epipolar features inaccurate. For a more in-depth examination, we encourage the reader to review the quality results of individual scenes provided in the supplementary.

6. Assessing Quality Metrics for Neural View Synthesis

Our collected datasets with video references, together with perceptual quality results of reconstructed videos, allow us to test how well the existing image/video quality metrics can measure the perceived quality. We test a range of existing objective metrics, full-reference and non-reference, image and video metrics, as listed in Table 1. We look into the widely used image similarity metrics such as PSNR and SSIM [WBSS04b], and also deep-learning related metrics, such as LPIPS [ZIE*18] and DISTS [DMWS20]. We test LPIPS using two different backbone models, VGG and AlexNet, as we noted they differ in their predictions. PSNR-L converts image RGB values into luma values before computing PSNR similarity. We include several video quality metrics, includ-

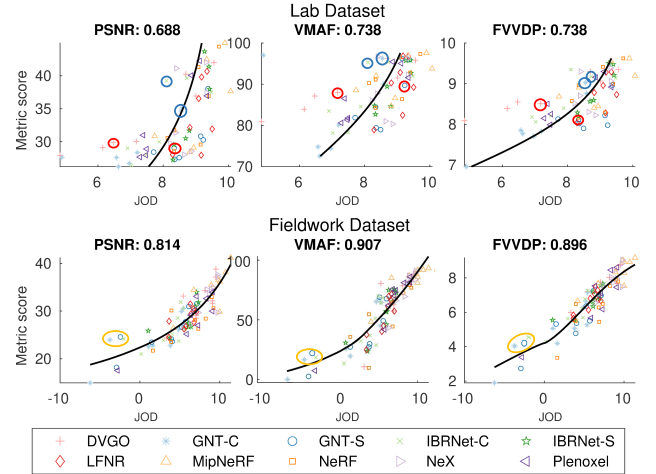


Figure 5: Selected metric correlations for our Lab (top row) and Fieldwork (bottom row) datasets. The black lines are obtained through fitting a logistic function which helps to detect outliers that affect correlations. Glossy animals produced by DVGO and GNT-S are denoted in red circles in the first row. Metal produced by IBRNet-C and GNT-C are denoted in blue circles. Vespa produced by IBRNet-C, GNT-S, and GNT-S are denoted in orange circles in the second row.

ing FovVideoVDP (v1.2, labeled FVVD) [MDC*21b], VMAF (v0.6.1) [LLWK14,LAK*16,LBJN*18], STRRED and HDRVQM. Given that FVVD and VMAF are applicable to both videos and images, we assess their performance in two distinct contexts: one is directly evaluated on videos (denoted with a "video" suffix), while the other is evaluated on each video frame and average among them (denoted with an "image" suffix). The absence of a suffix implies the evaluation is conducted on video content. In contrast to its image-based counterpart, a direct evaluation of videos entails the consideration of temporal distortions between successive video frames. For instance, VMAF-video incorporates an additional measurement of temporal differences within the luminance component when compared to VMAF-image. We also evaluated on several blind or non-reference image metrics (BRISQUE [MMB12], PIQE [VDB*15] and NIQE [MSB13]), that directly compute scores without comparing to the reference images. For metrics that require display parameters (e.g., HDRVDP-3, and FovVideoVDP), we matched the effective resolution of the videos in our experiment (40 ppd, see Section 4.2).

For our collected Lab and Fieldwork datasets, which include reference videos, we computed the quality scores on the captured test video sequences. Since LLFF dataset lacks reference videos, we computed the quality scores on the sparse test image set as done in NVS papers. Thus, it is inherently unfeasible to consider the temporal distortions within this dataset, which is likely to be inferior to testing on videos.

To test the reliability of popular metrics, we followed the standard protocol used to evaluate quality metrics [HME*22,PJ*15], and computed rank-order (Spearman) correlations between metric predictions and perceptual JOD values. Figure 5 shows scatter plots

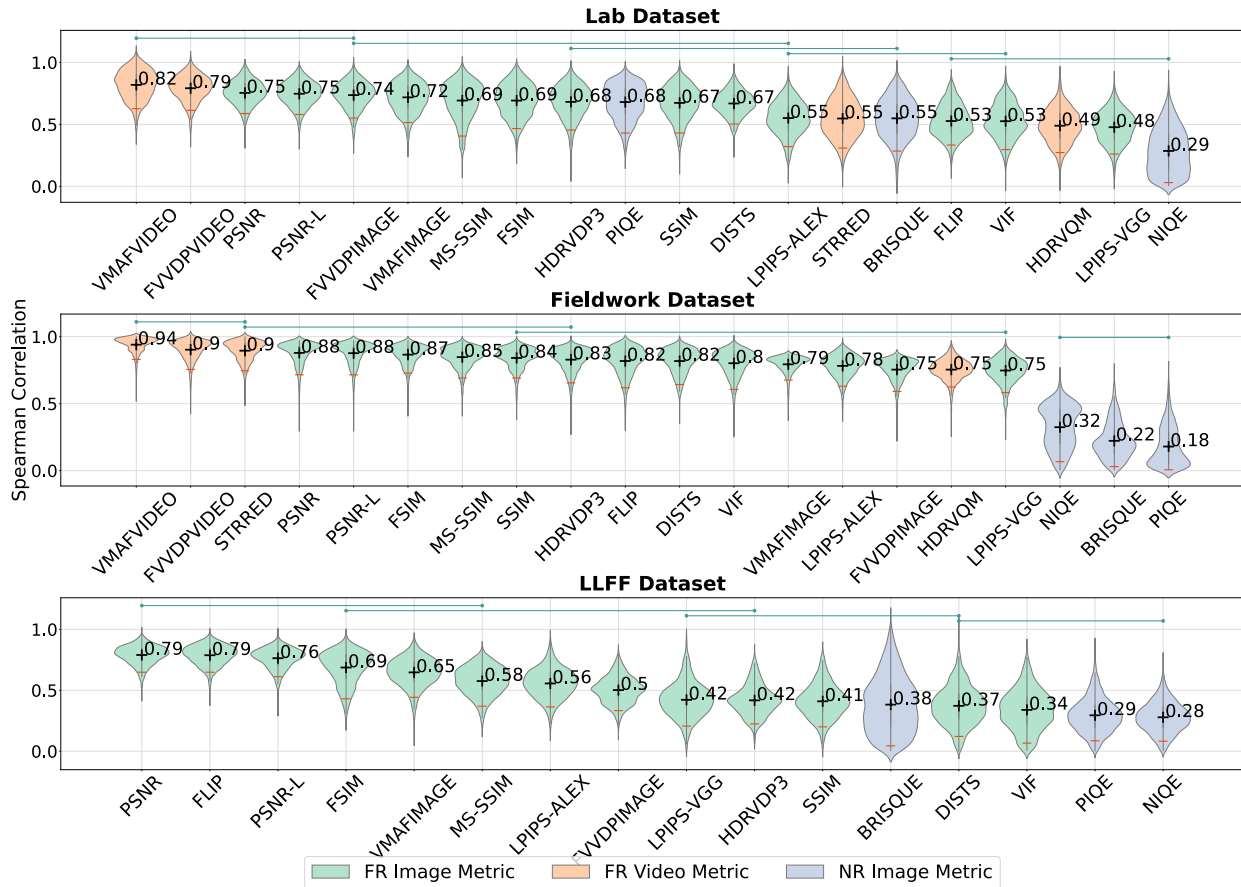


Figure 6: Bootstrapped distributions of correlation coefficients for all metrics computed on (a) Lab, (b) Fieldwork, and (c) LLFF. The “+” in black denotes mean correlation, and “-” in red denotes the 5th percentile (an estimate of the bad-case performance). Full-reference video metrics are missing for LLFF because this dataset has no ground-truth videos. The lines connect metrics where differences cannot be deemed statistically significant in a non-parametric test, with a p -value of 0.05.

and correlations of popular quality metrics w.r.t. subjective scores for the Lab (top row) and Fieldwork (bottom row) datasets. Please refer to the supplementary for similar plots for other metrics. We use least square optimization to fit a logistic function between metric score and subjective JOD score, which helps us find the outliers that affect the correlations. However, these point estimates of correlations conceal measurement noise due to: (a) the selection of scenes, (b) the subjective experiment results. Thus we cannot draw conclusions solely based on these correlations.

6.1. Averaged Bootstrapped Correlations

For each dataset and each NVS method, we averaged subjective scores and quality metric predictions across all scenes and then computed a single correlation per dataset per metric. This serves two purposes: (a) it mitigates the effects of measurement noise and improves the predictions of quality metrics as shown in previous works [HME*22], and (b) NVS methods are typically compared on averaged scores across scenes that reduce per-scene bias, making it more relevant for us.

When comparing quality metrics, it is essential to account for the variance in our data (subjective score variance and scene selection). We estimate the distribution of correlation values using bootstrapping [MR93]: we generated 2000 bootstrap samples for each estimated correlation by randomizing (sampling with replacement) both the participants and the selection of scenes. Within each bootstrap sample, we independently scaled the JOD values (following Section 4.4). In this way, our bootstrapping simulates 2000 outcomes of the experiment to capture the variance we can expect due to measurement noise. To determine whether the differences between the metrics are statistically significant, we performed a non-parametric test at the $\alpha = 0.05$ level by directly computing the distribution of the difference of bootstrapped samples. Please refer to supplementary for more details about non-parametric test. The results of that test are visualized as horizontal green lines in Figure 6.

6.2. Quality Metrics Performance

The correlations between metrics scores and subjective scores are shown in Figure 6. Below, we discuss the main observations that can be made based on that data.

PSNR is more accurate than SSIM and LPIPS NVS methods are typically evaluated using image quality metrics such as PSNR, SSIM, and LPIPS. The results in Figure 6 show that the simplest metric, PSNR, performed significantly better than more complex SSIM and LPIPS. NVS evaluation clearly does not benefit from the statistics extracted by SSIM or deep features extracted by LPIPS. Poor performance of SSIM has been noted before [PJI*15, LHS19], but it is still a popular metric because of its simplicity. The poor performance of LPIPS could be attributed to its training data consisting of small image patches with specific distortion types (noise, blur, compression-related, etc.) that are unlike NVS artifacts. We did not observe a statistically significant performance difference between PSNR-L (computed on luma) and PSNR (computed on RGB).

Importance of video reference dataset. In Figure 6, we observe that per-metric correlations are the lowest for the LLFF dataset and the highest for Fieldwork dataset. The low correlations for LLFF could be partially explained by the fact that while the subjective experiment measured video quality, the metrics could only be run on individual test views due to the lack of reference videos in LLFF.

To further investigate the importance of video reference, we experimented with sequences in which the number of available frames varied. We recomputed metric scores on progressively denser subsets of the test video frames (ranging from 10% to 100% of frames). Figure 7 illustrates the effect of increasing frames for representative image metrics. Note that the reported values are bootstrapped correlations (see Section 6.1) between metric predictions and subjective scores. The figure shows a gradual increase in correlation as we use more video frames. This observation highlights the limitations of using sparse image sets for assessing perceived quality, which degrades the predictions of image metrics.

All the above results indicate that the current objective evaluation protocol using a sparse image set is inadequate for assessing the perceptual quality of NVS methods applied to video generation. This underscores the rationale behind the development of our new datasets, which incorporate reference videos for testing.

Video metrics outperform image metrics. For both datasets with video reference (Lab and Fieldwork), video quality metrics VMAF-video and FVVDP-video demonstrate the strongest correlations with human perceptual assessments. As shown in Figure 7, the predictions of video metrics (shown as star markers) were more accurate than those of the same metrics run on video frames (shown as square markers). This compelling evidence suggests that the subjective assessment of NVS-generated video is significantly influenced by temporal distortions and highlights again the importance of video assessment for NVS evaluation. This also emphasizes the significance of using video metrics that take temporal artifacts into account.

Importance of challenging datasets. Similar to point estimates of correlations (Figure 5), the bootstrapped correlations are the highest for the Fieldwork dataset (Figure 6). The simple explanation for this result is that the Fieldwork dataset was more challenging for NVS methods and resulted in larger, more objectionable artifacts, as can also be seen in the subjective results (Figure 4). Such large differences make it much easier for the quality metrics to differentiate between the methods. In fact, most full-reference qual-

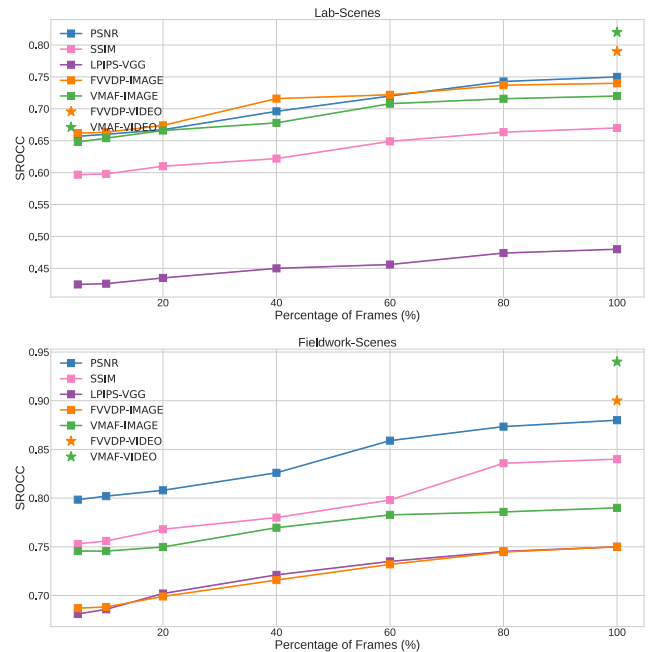


Figure 7: The performance of selected image quality metrics (in terms of SROCC) as the function of the number of frames used for the quality prediction. The star symbols on the right denote the performance of video metrics run on all frames. The performance of image quality metrics improves as we supply more images. Still, the highest performance is achieved by video metrics (FVVDP-video, VMAF-video), which can detect temporal distortions.

ity metrics performed well on this dataset. The Lab dataset, with its highly specular materials, was designed to pose a challenge to NVS methods. However, as it has a denser and more regular set of training views, most NVS methods performed well on those scenes, making it harder for the metrics to differentiate between the methods.

In summary, we recommend testing NVS methods on video sequences and using video quality metrics, such as VMAF and FVVDP. We still recommend using PSNR because of its simplicity, relatively good performance, and because it allows comparison with existing studies. We further recommend conducting testing of NVS methods on challenging datasets equipped with reference videos, such as our Fieldwork dataset.

6.3. Failure Cases of PSNR

Although PSNR is an effective image metric for evaluating NVS methods with respect to the perceived subjective quality, it is still beneficial to investigate when PSNR can fail to reflect subjective preferences. To do so, we compute per-scene correlations between PSNR scores and bootstrapped perceptual JOD values, and find scenes for which the metric results in poor correlations.

On the Lab dataset, we find that PSNR fails to accurately assess perceived quality on Glossy animals and Metal scenes. Particularly, for Glossy animals, we observe that the NVS gen-

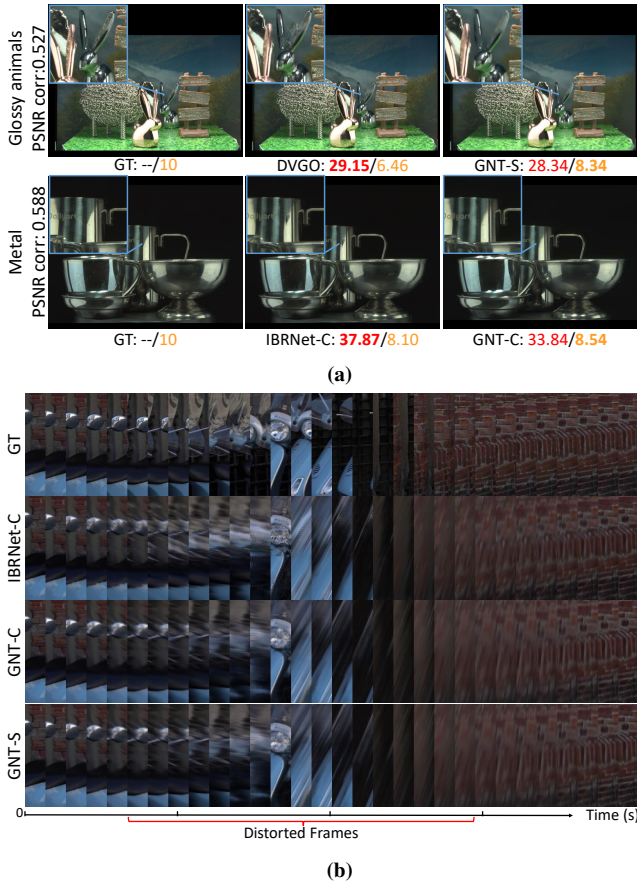


Figure 8: Representative scenes where PSNR fails to predict quality. (a) Scenes from Lab dataset: The values under the plot show PSNR / subjective-JOD values. The images in the middle column have higher PSNR values than those in the right column, but they contain obvious artifacts and are less preferred (compare JOD values) — see the blurry artifact in *Glossy animals* and unnatural local shading in *Metal*. We show the Spearman correlation between PSNR and subjective score on the left of the GT image (i.e. 0.527 on *Glossy animals* and 0.588 on *Metal*). (b) *Vespa* from Fieldwork dataset produced by GNT-C, GNT-S, and IBRNet-C show a degradation in image quality from 1s to 4s with severe temporal distortions between consecutive frames. These distortions are markedly disfavored by study participants. However, PSNR fails to capture the distortions adequately and yields relatively favorable metric scores. Please note that all these artifacts are more noticeable in videos.

erated by DVGO exhibits a higher PSNR score compared to that produced by GNT-S, as depicted in Figure 5, however, the latter is more preferred by participants. This discrepancy is exemplified in Figure 8a, which presents a representative reconstruction of this scene. Notably, PSNR appears to lack sensitivity to the blurry artifacts introduced by NVS methods, as observed in the top row’s middle column, specifically between the rabbit’s ears. However, this artifact is easily noticeable (especially when shown in video)

and significantly impacts human preferences, with a clear preference for images without such distortions (top row, right column in Figure 8a). The inadequacy of PSNR is also underscored when examining the *Metal* scene, notably in the NVS results produced by IBRNet-C and GNT-C, as illustrated in the second row of Figure 8a. In this scene, participants are highly sensitive to localized distortions, such as the unnatural local shading, as depicted in the bottom row’s middle column in Figure 8a. These subtle yet crucial artifacts may not be effectively captured by PSNR due to its reliance on averaging pixel-level information across the entire image.

On the Fieldwork dataset, although PSNR effectively evaluates the perceived quality across most scenes, it exhibits limitations in adequately assessing the scene *Vespa*, as can be regarded as an outlier illustrated in the orange circles in Figure 5. Specifically, the NVS results generated by GNT-C, GNT-S, and IBRNet-C, as depicted in Figure 8b, show a progressive degradation in image quality from 1s to 4s, accompanied by severe temporal distortions between consecutive frames. These observed distortions are markedly disfavored by study participants with low perceptual JOD value around -5. However, PSNR fails to capture the temporal distortions adequately, continuing to yield relatively favorable metric scores around 25. In contrast, both VMAF and FVVDP metrics prove to be effective in detecting and quantifying the presence of temporal artifacts (*Vespa* produced by IBRNet-C/GNT-C/GNT-S has VMAF scores of 15.09/16.10/18.35, and FVVDP scores of 4.73/4.23/4.38), offering a better correlation with subjective scores.

7. Conclusions

The primary application of NVS methods is the interactive exploration of 3D scenes. Yet, those methods are typically tested on isolated views instead of videos, which could mimic such 3D exploration. In this work, we collected two new datasets with reference videos and used them to evaluate 8 representative NVS methods (and two variants) in a subjective quality assessment experiment. The results helped us to identify the strengths and weaknesses of tested NVS methods, but also to evaluate 18 image/video quality metrics. We found that (a) existing quality metrics struggle to differentiate between the NVS methods when they are tested on datasets with a dense set of training views; and (b) SSIM and LPIPS, which are two commonly used quality metrics, perform worse than PSNR when evaluating NVS methods; (c) our analysis elucidates the limitation in the current NVS evaluation protocol and reveals the crucial need for video assessment and video metric. Our recommendation is to evaluate NVS methods on challenging datasets with sparsely sampled views and to use both PSNR and video metrics, such as VMAF and FovVideoVDP. Our work mainly focuses on evaluating front-facing scenes because this is applicable to more NVS methods (e.g., NeX). Moreover, this setup has significant applications in free-viewpoint video capture, particularly for AR/VR applications. While capturing 360° scenes is equally important, we will investigate such a setup in the future.

Acknowledgements

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie

Sklodowska-Curie grant agreement N^o 765911 (RealVision) and from the Royal Society, award IES202141. This project is also supported by a UKRI Future Leaders Fellowship [grant number G104084]. We greatly thank Prof. Damiano Marchi and Dr. Simone Farina for granting access to collect NeRFs inside the Natural History Museum of the University of Pisa.

References

- [AAB23] AZZARELLI A., ANANTRASIRICHAI N., BULL D. R.: Towards a robust framework for nerf evaluation. *arXiv preprint arXiv:2305.18079* (2023). 2
- [AMR*16] ARTUSI A., MANTIUK R. K., RICHTER T., KORSHUNOV P., HANHART P., EBRAHIMI T., AGOSTINELLI M.: JPEG XT: A compression standard for HDR and WCG images [standards in a nutshell]. *IEEE Signal Process. Mag.* 33, 2 (2016), 118–124. URL: <https://doi.org/10.1109/MSP.2015.2506199>, doi:10.1109/MSP.2015.2506199. 2
- [ANAM*20] ANDERSSON P., NILSSON J., AKENINE-MÖLLER T., OSKARSSON M., ÅSTRÖM K., FAIRCHILD M. D.: Flip: A difference evaluator for alternating images. *Proc. ACM Comput. Graph. Interact. Tech.* 3, 2 (aug 2020). URL: <https://doi.org/10.1145/3406183>, doi:10.1145/3406183. 2, 6
- [BMV*22a] BARRON J. T., MILDENHALL B., VERBIN D., SRINIVASAN P. P., HEDMAN P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022* (2022), IEEE, pp. 5460–5469. URL: <https://doi.org/10.1109/CVPR52688.2022.00539>, doi:10.1109/CVPR52688.2022.00539. 1, 4
- [BMV*22b] BARRON J. T., MILDENHALL B., VERBIN D., SRINIVASAN P. P., HEDMAN P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 5470–5479. 1
- [CAD19] CHAPIRO A., ATKINS R., DALY S.: A luminance-aware model of judder perception. *ACM Trans. Graph.* 38, 5 (jul 2019). URL: <https://doi.org/10.1145/3338696>, doi:10.1145/3338696. 1
- [ČHM*12] ČADÍK M., HERZOG R., MANTIUK R. K., MYŠKOWSKI K., SEIDEL H.-P.: New Measurements Reveal Weaknesses of Image Quality Metrics in Evaluating Graphics Artifacts. *ACM Transactions on Graphics (proc. of SIGGRAPH Asia)* 31, 6 (2012), 147. doi:10.1145/2366145.2366166. 1
- [Com18] COMMUNITY B. O.: *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018. URL: <http://www.blender.org>. 2
- [CW93] CHEN S. E., WILLIAMS L.: View interpolation for image synthesis. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques* (1993), pp. 279–288. 1
- [DLBD*23] DE LUIGI L., BOLOGNINI D., DOMENICONI F., GREGORIO D. D., POGGI M., STEFANO L. D.: Scannerf: a scalable benchmark for neural radiance fields. In *2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)* (2023), pp. 816–825. doi:10.1109/WACV56688.2023.00088. 1, 2
- [DM20] DENES G., MANTIUK R. K.: Predicting visible flicker in temporally changing images. *Electronic Imaging 2020*, 11 (2020), 233–1–233–8. URL: <https://www.ingentaconnect.com/content/ist/ei/2020/00002020/00000011/art00012>, doi:doi:10.2352/ISSN.2470-1173.2020.11.HVEI-233. 1
- [DMWS20] DING K., MA K., WANG S., SIMONCELLI E. P.: Image quality assessment: Unifying structure and texture similarity. *IEEE transactions on pattern analysis and machine intelligence* 44, 5 (2020), 2567–2581. 6
- [EWMU13] EILERTSEN G., WANAT R., MANTIUK R. K., UNGER J.: Evaluation of tone mapping operators for hdr-video. *Comput. Graph. Forum* 32, 7 (2013), 275–284. URL: <https://doi.org/10.1111/cg.f.12235>, doi:10.1111/cg.f.12235. 2
- [FBD*19] FLYNN J., BROXTON M., DEBEVEC P., DUVALL M., FYFFE G., OVERBECK R., SNAVELY N., TUCKER R.: Deepview: View synthesis with learned gradient descent. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), pp. 2367–2376. 2
- [FTS*23] FU S., TAMIR N., SUNDARAM S., CHAI L., ZHANG R., DEKEL T., ISOLA P.: Dreamsim: Learning new dimensions of human visual similarity using synthetic data. *arXiv preprint arXiv:2306.09344* (2023). 2
- [FYT*22] FRIDOVICH-KEIL S., YU A., TANCİK M., CHEN Q., RECHT B., KANAZAWA A.: Plenoxels: Radiance fields without neural networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022* (2022), IEEE, pp. 5491–5500. URL: <https://doi.org/10.1109/CVPR52688.2022.00542>, doi:10.1109/CVPR52688.2022.00542. 1, 4
- [HME*22] HANJI P., MANTIUK R., EILERTSEN G., HAJISHARIF S., UNGER J.: Comparison of single image hdr reconstruction methods — the caveats of quality assessment. In *ACM SIGGRAPH 2022 Conference Proceedings* (2022), pp. 1–8. 1, 2, 6, 7
- [HZM20] HANJI P., ZHONG F., MANTIUK R. K.: Noise-aware merging of high dynamic range image stacks without camera calibration. In *European Conference on Computer Vision* (2020), Springer, pp. 376–391. 3
- [JDV*14] JENSEN R., DAHL A., VOGIATZIS G., TOLA E., AANÆS H.: Large scale multi-view stereopsis evaluation. In *2014 IEEE Conference on Computer Vision and Pattern Recognition* (2014), IEEE, pp. 406–413. 2
- [KKLD23] KERBL B., KOPANAS G., LEIMKÜHLER T., DRETTAKIS G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* 42, 4 (2023). 1
- [KLR*22] KOPANAS G., LEIMKÜHLER T., RAINER G., JAMBON C., DRETTAKIS G.: Neural point catacaustics for novel-view synthesis of reflections. *ACM Transactions on Graphics (TOG)* 41, 6 (2022), 1–15. 1
- [KPZK17] KNAPITSCH A., PARK J., ZHOU Q.-Y., KOLTUN V.: Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)* 36, 4 (2017), 1–13. 2
- [LAK*16] LI Z., AARON A., KATSAVOUNIDIS I., MOORTHY A., MANOHARA M.: Toward a practical perceptual video quality metric, June 2016. URL: <https://netflixtechblog.com/toward-a-practical-perceptual-video-quality-metric-653f20126>. 1, 2, 6
- [LBJN*18] LI Z., BAMPIS C., JULIE NOVAK A. A., SWANSON K., MOORTHY A., COCK J. D.: Vmaf: The journey continues, October 2018. URL: <https://netflixtechblog.com/vmaf-the-journey-continues-44b51ee9ed12>. 6
- [LCTS05] LEDDA P., CHALMERS A., TROSCIANKO T., SEETZEN H.: Evaluation of tone mapping operators using a high dynamic range display. *ACM Trans. Graph.* 24, 3 (2005), 640–648. URL: <https://doi.org/10.1145/1073204.1073242>, doi:10.1145/1073204.1073242. 2
- [LGZL*20] LIU L., GU J., ZAW LIN K., CHUA T.-S., THEOBALT C.: Neural sparse voxel fields. *Advances in Neural Information Processing Systems* 33 (2020), 15651–15663. 2
- [LHS19] LIN H., HOSU V., SAUPE D.: KADID-10k: A Large-scale Artificially Distorted IQA Database. In *2019 Eleventh International Conference on Quality of Multimedia Experience (QoMEX)* (jun 2019), vol. 161, IEEE, pp. 1–3. URL: <https://ieeexplore.ieee.org/document/8743252/>, doi:10.1109/QoMEX.2019.8743252. 8

- [LLWK14] LIN J. Y., LIU T., WU E. C., KUO C. J.: A fusion-based video quality assessment (fvqa) index. In *Asia-Pacific Signal and Information Processing Association Annual Summit and Conference, APSIPA 2014, Chiang Mai, Thailand, December 9-12, 2014* (2014), IEEE, pp. 1–5. URL: <https://doi.org/10.1109/APSIPA.2014.7041705>, doi:10.1109/APSIPA.2014.7041705. 6
- [MDC*21a] MANTIUK R. K., DENES G., CHAPIRO A., KAPLANYAN A., RUFO G., BACHY R., LIAN T., PATNEY A.: Fovvideovp: A visible difference predictor for wide field-of-view video. *ACM Trans. Graph.* 40, 4 (jul 2021). URL: <https://doi.org/10.1145/3450626.3459831>, doi:10.1145/3450626.3459831. 1
- [MDC*21b] MANTIUK R. K., DENES G., CHAPIRO A., KAPLANYAN A., RUFO G., BACHY R., LIAN T., PATNEY A.: Fovvideovp: A visible difference predictor for wide field-of-view video. *ACM Trans. Graph.* 40, 4 (jul 2021). URL: <https://doi.org/10.1145/3450626.3459831>, doi:10.1145/3450626.3459831. 2, 6
- [MKRH11] MANTIUK R., KIM K. J., REMPEL A. G., HEIDRICH W.: HDR-VDP-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM Transactions on graphics (TOG)* 30, 4 (2011), 1–14. 6
- [MMB12] MITTAL A., MOORTHY A. K., BOVIK A. C.: No-reference image quality assessment in the spatial domain. *IEEE Transactions on Image Processing* 21, 12 (2012), 4695–4708. doi:10.1109/TIP.2012.2214050. 6
- [MR93] MOONEY C. Z., ROBERT D. DUVAL: *Bootstrapping: A Non-parametric Approach to Statistical Inference*. Sage, 1993. 7
- [MRAQ23] MARTIN P., RODRIGUES A., ASCENSO J., QUELUZ M. P.: Nerf-qa: Neural radiance fields quality assessment database. *arXiv preprint arXiv:2305.03176* (2023). 2
- [MSB13] MITTAL A., SOUNDARARAJAN R., BOVIK A. C.: Making a “completely blind” image quality analyzer. *IEEE Signal Processing Letters* 20, 3 (2013), 209–212. doi:10.1109/LSP.2012.2227726. 6
- [MSOC*19] MILDENHALL B., SRINIVASAN P. P., ORTIZ-CAYON R., KALANTARI N. K., RAMAMOORTHY R., NG R., KAR A.: Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–14. 2, 4
- [MST*20] MILDENHALL B., SRINIVASAN P. P., TANCIK M., BARRON J. T., RAMAMOORTHY R., NG R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV* (2020). 1, 2, 4, 5
- [MWPO*21] MIKHAILIUK A., WILMOT C., PEREZ-ORTIZ M., YUE D., MANTIUK R.: Active sampling for pairwise comparisons via approximate message passing and information gain maximization. In *2020 IEEE International Conference on Pattern Recognition (ICPR)* (Jan 2021). 5
- [NPDSL15] NARWARIA M., PERREIRA DA SILVA M., LE CALLET P.: Hdr-vqm: An objective quality measure for high dynamic range video. *Signal Processing: Image Communication* 35 (05 2015). doi:10.1016/j.image.2015.04.009. 6
- [Ols11] OLSON E.: Apriltag: A robust and flexible visual fiducial system. In *2011 IEEE international conference on robotics and automation* (2011), IEEE, pp. 3400–3407. 3
- [PJI*15] PONOMARENKO N., JIN L., IEREMEIEV O., LUKIN V., EGAZARIAN K., ASTOLA J., VOZEL B., CHEHDI K., CARLI M., BATTISTI F., JAY KUO C.-C.: Image database TID2013: Peculiarities, results and perspectives. *Signal Processing: Image Communication* 30 (jan 2015), 57–77. URL: <http://linkinghub.elsevier.com/retrieve/pii/S0923596514001490>, doi:10.1016/j.image.2014.10.009. 1, 6, 8
- [POM17] PEREZ-ORTIZ M., MANTIUK R. K.: A practical guide and software for analysing pairwise comparison experiments. *arXiv preprint arXiv:1712.03686* (2017). 5
- [POMZ*20] PÉREZ-ORTIZ M., MIKHAILIUK A., ZERMAN E., HULUSIC V., VALENZISE G., MANTIUK R. K.: From pairwise comparisons and rating to a unified quality scale. *IEEE Transactions on Image Processing* 29 (2020), 1139–1151. doi:10.1109/TIP.2019.2936103. 4
- [SB06] SHEIKH H., BOVIK A.: Image information and visual quality. *IEEE Transactions on Image Processing* 15, 2 (2006), 430–444. doi:10.1109/TIP.2005.859378. 6
- [SB13] SOUNDARARAJAN R., BOVIK A.: Video quality assessment by reduced reference spatio-temporal entropic differencing. *Circuits and Systems for Video Technology, IEEE Transactions on* 23 (04 2013), 684–694. doi:10.1109/TCSVT.2012.2214933. 2, 6
- [SESM22] SUHAIL M., ESTEVES C., SIGAL L., MAKADIA A.: Light field neural rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 8269–8279. 4
- [SF16] SCHÖNBERGER J. L., FRAHM J.-M.: Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)* (2016). 3
- [SGHS98] SHADE J., GORTLER S., HE L.-W., SZELISKI R.: Layered depth images. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques* (1998), pp. 231–242. 1
- [SK00] SHUM H., KANG S. B.: Review of image-based rendering techniques. In *Visual Communications and Image Processing 2000* (2000), vol. 4067, SPIE, pp. 2–13. 1
- [SSC22] SUN C., SUN M., CHEN H.: Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022* (2022), IEEE, pp. 5449–5459. URL: <https://doi.org/10.1109/CVPR52688.2022.00538>, doi:10.1109/CVPR52688.2022.00538. 1, 4
- [VDB*15] VENKATANATH N., D P., BH M. C., CHANNAPPAYYA S. S., MEDASANI S. S.: Blind image quality evaluation using perception based features. 1–6. doi:10.1109/NCC.2015.7084843. 6
- [VSJ22] VICINI D., SPEIERER S., JAKOB W.: Differentiable signed distance function rendering. *Transactions on Graphics (Proceedings of SIGGRAPH)* 41, 4 (July 2022), 125:1–125:18. doi:10.1145/3528223.3530139. 1
- [WBF*17] WAECHTER M., BELJAN M., FUHRMANN S., MOEHRLE N., KOPF J., GOESELE M.: Virtual rephotography: Novel view prediction error for 3d reconstruction. *ACM Transactions on Graphics (TOG)* 36, 1 (2017), 1–11. 1
- [WBSS04a] WANG Z., BOVIK A., SHEIKH H., SIMONCELLI E.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* 13, 4 (2004), 600–612. doi:10.1109/TIP.2003.819861. 6
- [WBSS04b] WANG Z., BOVIK A. C., SHEIKH H. R., SIMONCELLI E. P.: Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.* 13, 4 (2004), 600–612. URL: <https://doi.org/10.1109/TIP.2003.819861>, doi:10.1109/TIP.2003.819861. 1, 6
- [WCC*22] WANG P., CHEN X., CHEN T., VENUGOPALAN S., WANG Z., ET AL.: Is attention all nerf needs? *arXiv preprint arXiv:2207.13298* (2022). 4
- [WPYS21] WIZADWONGSA S., PHONGTHAWEE P., YENPHRAPHAI J., SUWAJANAKORN S.: Nex: Real-time view synthesis with neural basis expansion. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021* (2021), Computer Vision Foundation / IEEE, pp. 8534–8543. doi:10.1109/CVPR46437.2021.00843. 1, 2, 4
- [WSB03] WANG Z., SIMONCELLI E., BOVIK A.: Multiscale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems Computers, 2003* (2003), vol. 2, pp. 1398–1402 Vol.2. doi:10.1109/ACSSC.2003.1292216. 6

- [WWG*21] WANG Q., WANG Z., GENOVA K., SRINIVASAN P. P., ZHOU H., BARRON J. T., MARTIN-BRUALLA R., SNAVELY N., FUNKHOUSER T. A.: Ibrnet: Learning multi-view image-based rendering. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021* (2021), Computer Vision Foundation / IEEE, pp. 4690–4699. URL: https://openaccess.thecvf.com/content/CVPR2021/html/Wang_IBRNet_Learning_Multi-View_Image-Based_Rendering_CVPR_2021_paper.html, doi:10.1109/CVPR46437.2021.00466. 1,4
- [WWL*23] WANG C., WANG A., LI J., YUILLE A., XIE C.: Benchmarking robustness in neural radiance fields, 2023. URL: <https://arxiv.org/abs/2301.04075>, doi:10.48550/ARXIV.2301.04075. 1
- [YKGB*20] YUE D., KHAN GUL M. S., BÄTZ M., KEINERT J., MANTIUK R.: A benchmark of light field view interpolation methods. In *2020 IEEE International Conference on Multimedia & Expo Workshops (ICMEW) (2020)*, pp. 1–6. doi:10.1109/ICMEW46912.2020.9106041. 2
- [YKM*20] YARIV L., KASTEN Y., MORAN D., GALUN M., ATZMON M., RONEN B., LIPMAN Y.: Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems 33* (2020). 1
- [YLL*20] YAO Y., LUO Z., LI S., ZHANG J., REN Y., ZHOU L., FANG T., QUAN L.: Blendedmvs: A large-scale dataset for generalized multi-view stereo networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), pp. 1790–1799. 2
- [Zha00] ZHANG Z.: A flexible new technique for camera calibration. *IEEE Transactions on pattern analysis and machine intelligence* 22, 11 (2000), 1330–1334. 3
- [ZIE*18] ZHANG R., ISOLA P., EFROS A. A., SHECHTMAN E., WANG O.: The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR* (2018). 1, 6
- [ZTF*18] ZHOU T., TUCKER R., FLYNN J., FYFFE G., SNAVELY N.: Stereo magnification: Learning view synthesis using multiplane images. *arXiv preprint arXiv:1805.09817* (2018). 2
- [ZZMZ11] ZHANG L., ZHANG L., MOU X., ZHANG D.: Fsim: A feature similarity index for image quality assessment. *IEEE Transactions on Image Processing* 20, 8 (2011), 2378–2386. doi:10.1109/TIP.2011.2109730. 6