

A Large-Scale Outdoor Multi-modal Dataset and Benchmark for Novel View Synthesis and Implicit Scene Reconstruction

Chongshan Lu¹ Fukun Yin^{1,2} Xin Chen² Wen Liu² Tao Chen^{1*} Gang Yu² Jiayuan Fan³

¹School of Information Science and Technology, Fudan University, China

²Tencent PCG, China

³Academy for Engineering and Technology, Fudan University, China

Abstract

Neural Radiance Fields (NeRF) [24] has achieved impressive results in single object scene reconstruction and novel view synthesis, as demonstrated on many single modality and single object focused indoor scene datasets like DTU [14], BMVS [42], and NeRF Synthetic [24]. However, the study of NeRF on large-scale outdoor scene reconstruction is still limited, as there is no unified outdoor scene dataset for large-scale NeRF evaluation due to expensive data acquisition and calibration costs. In this work, we propose a large-scale outdoor **multi-modal** dataset, **OMMO dataset**, containing complex objects and scenes with calibrated images, point clouds and prompt annotations. A new benchmark for several outdoor NeRF-based tasks is established, such as novel view synthesis, diverse 3D representation, and multi-modal NeRF. To create the dataset, we capture and collect a large number of real fly-view videos and select high-quality and high-resolution clips from them. Then we design a quality review module to refine images, remove low-quality frames and fail-to-calibrate scenes through a learning-based automatic evaluation plus manual review. Finally, volunteers are employed to label and review the prompt annotation for each scene and keyframe. Compared with existing NeRF datasets, our dataset contains abundant real-world urban and natural scenes with various scales, camera trajectories, and lighting conditions. Experiments show that our dataset can benchmark most state-of-the-art NeRF methods on different tasks. The dataset can be found at the following link: <https://ommo.luchongshan.com/>.

1. Introduction


Recent advances in implicit neural representations have achieved remarkable results in photo-realistic novel view synthesis and high-fidelity surface reconstruction [44, 43]. Unfortunately, most of the existing methods focus on single

*Corresponding author.


Prompt annotation

Buildings at night in the background; two tall buildings in the middle; roads among buildings; cars are passing on the road; tall buildings decorated with blue and orange lights.

Calibrated image



Camera trajectory



More views



More scenes



Figure 1. A city scene example from our dataset captured with low illuminance and circle-shaped camera trajectory. We show multi-view calibrated images, the camera track, and text descriptions of the scene. Some details in colored boxes are zoomed in to indicate that our dataset can provide real-world high-fidelity texture details.

objects or indoor scenes [44, 43, 9, 15, 5], and their synthesis performance will decrease drastically if migrated to outdoor scenes. Although some very recent methods try to solve this problem and are well-designed for large scenes [36, 40], their performance is difficult to compare due to the lack of large-scale outdoor scene datasets and uniform benchmarks.

At present, the existing outdoor scene datasets are either

Table 1. Comparison with existing NeRF datasets, especially those outdoor datasets related to ours. The first group is single object datasets, the second group is large scenes datasets (or outdoor parts), and the last row is our dataset. For each dataset, we show the number of scenes and images, the diversity of scene types (D_T), camera trajectories (D_C), and lighting conditions (D_L), and whether they are real-world scenes (*Real*), and whether they have multi-modal data (*Mmodal*).

Datasets	#Scenes	#Images	D_T	D_C	D_L	Real	Mmodal
DTU [14]	124	4.2K	No	No	Yes	Yes	No
NeRF [24]	18	3551	No	Yes	No	Yes	No
Scannet [7]	1.5K	2.5M	No	Yes	No	Yes	No
T & T [18]	6	88k	No	No	No	Yes	No
BMVS [42]	28	5k	Yes	No	No	Part	No
Urban3D [22]	16	10.4K	Yes	No	No	Part	No
Quad 6k [6]	1	5.1K	No	Yes	No	Yes	No
Mill 19 [36]	2	3.6K	Yes	No	No	Yes	No
Block-NeRF [35]	an entire city	2.8M	Yes	No	Yes	Yes	No
Ours	33	14.7K	Yes	Yes	Yes	Yes	Yes

collected with a small geographical scale, or rendered from not-real virtual scenes. For example, Tanks and temples [18] provides a benchmark of realistic outdoor scenes captured by a high-precision industrial laser scanner, but its scene scale is still too small ($463m^2$ on average) and only focuses on a single outdoor object or building. The BlendedMVS [42] and UrbanScene3D [22] datasets contain scene images rendered from reconstructed or virtual scenes, which deviate from the real scene in both texture and appearance details. Collecting images from the Internet can theoretically build very effective datasets [13, 1], like ImageNet [8] and COCO [21], but these methods are not suitable for NeRF-based task evaluation due to the changes of objects and lighting conditions in the scene at different times. Our dataset acquisition method is similar to Mega-NeRF [36], which captures large real-world scenes by drones. But Mega-NeRF only provides two monotonic scenes, which hinders it from being a widely used baseline. Therefore, to our knowledge, no uniform and widely recognized large-scale scene dataset is built for NeRF benchmarking, causing large-scale NeRF research for outdoor far fall behind that for single objects or indoor scenes [14, 42, 24, 7].

To address the lack of large-scale real-world outdoor scene datasets, we introduce a well-selected fly-view multi-modal dataset. The dataset contains totally 33 scenes with prompt annotations, tags, and 14K calibrated images (*cf.* Fig. 4). Different from the existing methods mentioned above, the sources of our scenes are very extensive, including those collected on the Internet and captured by ourselves. Meanwhile, the collection indicators are also comprehensive and representative, including various scene types, scene scales, camera trajectories, lighting conditions, and multi-modal data that are not available in existing datasets (*cf.* Tab. 1). More importantly, we provide a generic pipeline to generate real-world NeRF-based data from drone videos on the Internet, which makes our dataset easily to be extensible by the community.

Further, to evaluate the applicability and performance

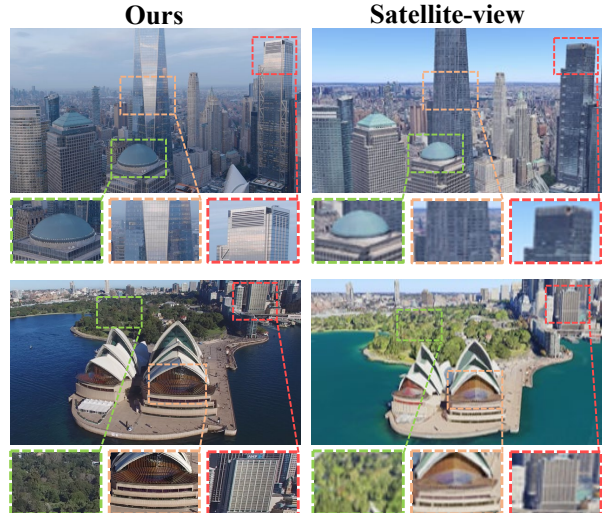


Figure 2. Visual comparison with existing large-scale satellite-view outdoor datasets [40] acquired from Google Earth Studio. The top row is from [40], and the bottom row is corresponding scenes from our fly-view dataset, which is more realistic with clear textures and rich details (zoom-in for the best of views).

of the built dataset for evaluating mainstream NeRF methods, we build all-around benchmarks including novel view synthesis, scene representations, and multi-modal synthesis based on the dataset. Moreover, we provide several detailed sub-benchmarks for each above task, according to different scene types, scene scales, camera trajectories and lighting conditions, to give a fine-grained evaluation of each method.

We aim to establish a general large-scale outdoor NeRF-based benchmark and promote NeRF-related algorithm research on novel view synthesis, surface reconstruction, and multi-modal. To summarize, our main contributions include:

- Aiming at advancing the large-scale NeRF research, we introduce an outdoor scene dataset captured from the real world with multi-modal data, which surpasses all existing relative outdoor datasets in both quantity and diversity, see Tab. 1 and Sec. 3.3.
- To form a uniform benchmarking standard for outdoor NeRF methods, we create multiple benchmark tasks for mainstream outdoor NeRF methods. Extensive experiments show that our dataset can well support common NeRF-based tasks and provide prompt annotations for future research, see Sec. 4.
- We provide a cost-effective pipeline for converting videos that can be flexibly accessed from the Internet to NeRF-purpose training data, which makes our dataset easily scalable, see Sec. 3.1 and Sec. 3.2.

2. Related Work

2.1. Neural Scene Representation and Rendering

Neural Radiance Fields (NeRF) [24] propose an effective implicit neural scene representation method to synthesize

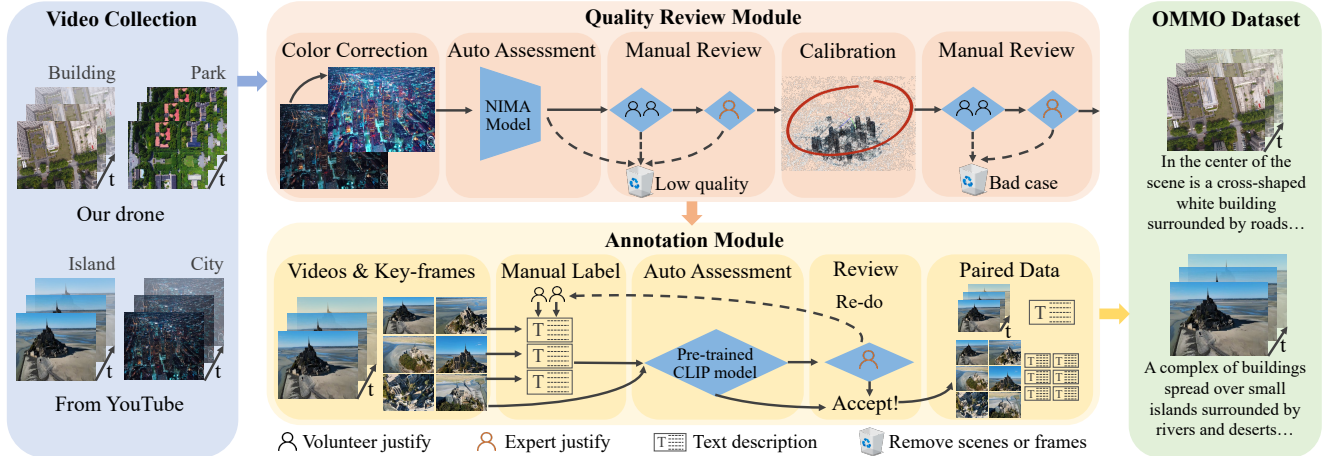


Figure 3. The pipeline for our dataset generation. The original videos are collected from both YouTube and captured by us, and then fed into the review and annotation module. The former mainly removes low-quality frames and failed scenes; the latter annotates text descriptions for scenes and keyframes. Currently, we have generated 33 scenes with 14K images and text descriptions.

novel views by the single scene optimization. Many subsequent coordinate-based methods are inspired by it, and we can classify them into neural surface fields [27, 43, 38] and neural radiance fields [26, 16, 45] based on the shape representation differences. However, both ways have poor performance on large scenes due to the limited representation capability of MLP-based networks adopted in NeRF. Even with prior information, it is difficult to directly apply small-scale scene-focused methods due to the increased scene complexity [44, 19].

Fortunately, some very recent works have started to study the neural representation for large-scale scenes. Mega-NeRF [36] divides the large fly-view scene into multiple small blocks to train specialized NeRFs in parallel. Block-NeRF [34] also adopts this simplified idea, dividing the neighborhood into blocks, and then novel views are sampled from overlapping blocks and combined according to inverse distance weights. CityNerf (BungeeNeRF) [40] introduces a progressive neural radiance field that starts from fitting distant views with a shallow base block, and appends new blocks to accommodate details. NeRF in the Wild (NeRF-W) [23] introduces a series of extensions to NeRF [24] to synthesize novel views of complex scenes, using only unstructured collections of in-the-wild photos. Recursive-NeRF [41] provides an efficient and adaptive rendering and training approach for NeRF, that forwards high uncertainties coordinates to a bigger network with more powerful representational capability.

However, the aforementioned large-scale NeRF methods [36, 34, 40, 23, 41] use different outdoor datasets with various capturing conditions and research focuses, causing difficulty to fairly compare these NeRF methods’ performance on common tasks based on a uniform benchmark.

2.2. NeRF-based Datasets and Benchmarks

There are widely used NeRF datasets and established benchmarks for single objects [14, 24], unbounded objects [42], human faces [29], and indoor scenes [7], see the first group in Tab. 1.

For large-scale outdoor scenes, some datasets provide high-fidelity models from accurate radar scans, but this expensive data acquisition makes the scale and size of these datasets still unsatisfactory [18]. Rendering images from optimized models will result in higher cost and unrealistic scenes [42, 22]. A low-cost and easy-to-expand way is to collect images of the same scene shot by different people and devices, at different times from the Internet [13, 1]. However, these methods do not meet the needs of common NeRF tasks due to the changes in weather, lighting, and objects in the scene. Mega-NeRF [36] builds two high-quality fly-view real scenes and calibrated images, but has not become a widely used benchmark due to its small size and single type. There are also large-scale NeRF datasets and benchmarks for specific problems, such as neighborhoods [34] or remote sensing [40], see the second group in Tab. 1.

In conclusion, for the reasons mentioned above, none of the above datasets have formed a widely used uniform benchmark. So a more comprehensive outdoor NeRF-based dataset is required, to facilitate the research and exploitation of larger-scale implicit scene representation. In contrast, the built dataset in this work provides 33 large-scale scenes, more than 14K fly-view images with camera poses, rich content, and text descriptions. Meanwhile, several new large-scale fly-view-based benchmarks for novel view synthesis, implicit scene representations, and multi-modal synthesis tasks are also proposed.



Figure 4. Details comparison between ours and other datasets, where the first row is from other datasets (BlendMVS, Mill-19, and Bungee NeRF), and the second row is from our dataset.

3. Dataset Generation

Our dataset acquisition, calibration, and annotation pipeline are shown in Fig. 3. We first decompose and enhance the original videos by time-sampling and color-correcting, and review the quality of the frames and scene calibrations by automated models followed by volunteers to remove low-quality frames and fail-to-calibrate scenes, see Sec. 3.1. Then volunteers provide prompt annotations for each scene and keyframe, and the CLIP [28] model is exploited to cooperate with human experts, to supervise the semantic consistency between labeled text and corresponding images, see Sec. 3.2. Finally, we introduce the dataset distribution from several aspects such as scene categories and collection cost, see Sec.3.3 and Sec. 3.4.

3.1. Acquisition and Calibration Method

Original Videos. Our outdoor fly-view dataset mainly comes from two sources: captured by ourselves and collected on YouTube. Among them, YouTube videos are from worldwide, including urban and natural scenes of various scales, camera moving trajectories and lighting conditions. However, these videos often have limited resolution due to the compression when uploading, while our captured videos are all 4K HD and can meet more high-fidelity NeRF needs. Benefiting from the diversity of videos, the proposed scenes in our dataset include various elements, such as buildings, roads, trees, islands, mountains, rivers, etc.

Compared to synthetic data, our scenes’ space layout, dynamic range, moving objects, surface reflection, and lighting conditions are completely real, which supports NeRF methods’ learning about real-world outdoor scenes. In addition, to have different lighting conditions in the same scene, we use drones to capture videos of the same building at different times and weather conditions. To explore the impact of different camera trajectories, we scan the same scene with different flight strategies. Overall, our original videos come from real scans from around the world, covering a variety of scene types, camera trajectories, and lighting conditions.

Color Correction. We set different sampling intervals according to the video length and frames per second(FPS, ranges from 24 to 30 in our dataset), so that the remaining

frame number is between 800-1000 to meet the needs of calibration. Since NeRF is subject to low dynamic range, we first employ an image enhancement model [17] to enhance and correct ill-illumination, overexposure, distortion, rain and fog weather frames. Image enhancement models are effective in reducing the dynamic range of images to a more general level, especially for low-light or overexposed images. Human review and scene calibration can remove images that deviate from the color range of the scene. By doing so, we ensure that the dataset is consistent and suitable for training and evaluation purposes. For HDR NeRF research, raw images and physical attributes will also be available.

Auto Assessment. Then the NIMA [33] model is employed to evaluate image quality and remove blur, ghosting, and low-quality images.

Manual Quality Review. After the quality auto-evaluation, we employ three volunteers to manually discard frames that do not meet requirements i.e. blurring, artifacts, focusing outside the scene, or out of scene dynamic range. Specifically, the three volunteers consist of 2 professional data labelers and a domain expert. For each frame, 2 professional data labelers judge whether to keep it or not, and the decision is made if both labelers agree, otherwise it is up to the domain expert for the decision.

Calibration. Large-scale scene calibration based on highly dynamic images has always been a difficult problem, especially for outdoor scenes with moving objects and no obvious local detail differences [4, 20, 12]. As a common solution, we use COLMAP [31, 32] to achieve multi-view 3D reconstruction, image calibration and depth image rendering. It is foreseeable that the reconstruction of some scenes with insufficient overlap and textures, or forwardly moving camera motion will fail. These fail-to-calibrate scenes cannot meet the requirements of NeRF-based methods, which need to be removed manually.

Manual Scene Review. We invite the above three volunteers to review the calibration quality based on the completeness of scene point clouds, following the same decision-making process. Finally, 33 real-world NeRF-friendly scans with nearly 14K calibrated images constitute our outdoor dataset, which surpasses existing large-scale datasets in both quantity and diversity. Viewed comprehensively, the evaluation of image quality primarily depends on the quality assessment model, which offers a relatively objective approach and reduces the requirement for manual labor. In parallel, multi-level manual reviews are solely aimed at addressing potential oversights of the model, ensuring a holistic quality standard.

3.2. Prompt Annotation Method

To provide prompt annotations for multi-modal NeRF, we annotate text descriptions to each scene and keyframes.

Manual Label. Generating descriptions is a more sub-

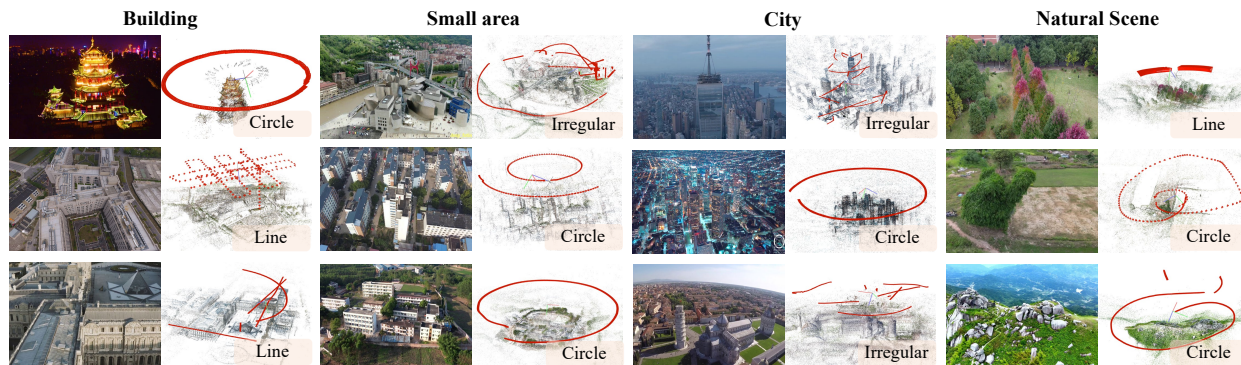


Figure 5. Examples of different types from our dataset. We visualize some scenes and camera trajectories from our dataset, which contain both urban and natural scenes with various scales, camera trajectories, and lighting conditions.

Table 2. Distribution of our dataset. We divide our dataset into subsets based on scene type, camera trajectory, and lighting condition and count the number of each subset.

Scene Type	# Scenes	Camera Trajectory	# Scenes	Lighting Condition	# Scenes
Building	8	Circle	15	Day	30
Small area	9	Line	10	Night	3
City	8	Irregular	8		
Natural scene	8				

jective and time-consuming task, so we need to employ more volunteers. Thus six trained volunteers participate in this progress, who non-repeatedly label and record the corresponding descriptive texts from scenes and keyframes, respectively. These volunteers include senior Ph.D. students and professors in computer vision and natural language processing fields who can handle this job well.

Auto Assessment. Each frame and annotated text is fed into a CLIP [28] model finetuned on large scene-text datasets [30, 11] to compute their similarity scores. If the similarity score is above the threshold, we accept this annotation. Otherwise, we leave it to experts to double-check.

Expert Review. For annotations that the CLIP model cannot judge, experts will evaluate whether the text can describe the image comprehensively and objectively. If possible, we accept the label, otherwise, we hand it over to another volunteer to label again. Fortunately, the frames and scenes in our dataset are only re-annotated at most 2 times.

We have annotated scenes and keyframes with corresponding descriptions and tags, which can well meet the training needs of multi-modal NeRF. We are still annotating the remaining frames for more potentially complex needs.

3.3. Distribution

According to different division methods, the distribution of our dataset is shown in Tab. 2. However, some scene types are relatively ambiguous. For example, when a building is surrounded by plenty of trees, warehouses, etc., it is difficult to say whether such an image belongs to a building type or

not. To resolve the ambiguity, we design a questionnaire and invite 50 voters to determine the attributes of the scenes. Invited volunteers range from 19 to 53 years old, and we recommend a very typical reference for each attribute.

Our dataset contains both natural and urban scenes, which are further divided into buildings, small areas, and whole cities. The performance on different subsets can verify the most suitable scene type and scale for different NeRF method. Some methods only target scenes with camera trajectories moving in a ring or matrix. For a fair baseline, we divide the scene into circles, lines, and irregular, according to the camera trajectories. In addition, few recent methods focus on low-illuminance NeRF research, so we also provide scenes with different lighting conditions. In particular, we collect some scans of the same scene under different lighting conditions, which can evaluate the ability of the method against poor lighting.

We provide benchmarks for novel view synthesis, scene representations, and multi-modal synthesis tasks. At the same time, according to the above scene types, our dataset can also produce corresponding sub-benchmarks to evaluate methods under different conditions and settings, see Sec. 4.

3.4. Cost

Since our dataset requires volunteers to review scenes' quality and annotate prompts, it is unavoidably time-consuming and labor-intensive.

However, data collection which involves drone purchasing and shooting, and computing costs introduced by pre-training scene-text CLIP model, account for a large part of the total cost. But once the drone and CLIP model is ready, we can just use them to prepare more new scenes and expand the dataset, without involving additional cost. Besides, our proven pipeline can directly process videos from the Internet, which is more general and reduces collection costs. Therefore, although our method may be expensive in the early stage, it has good potential to be automated and save future costs that may be introduced when expanding the dataset.

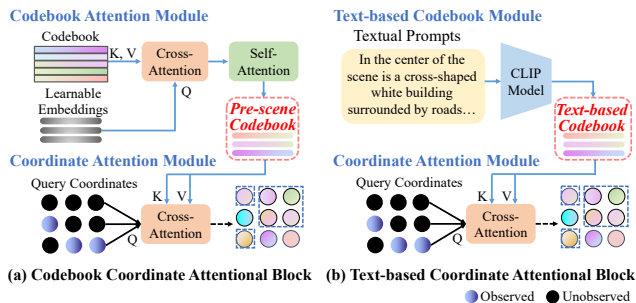


Figure 6. Structure of coordinate attentional blocks. The left sub-figure (a) is the CoCo-INR’s [44] codebook coordinate attentional block. The right sub-figure (b) is our text-based coordinate attentional block, which obtains the scene-related text-based codebook by encoding textual prompts of each scene.

Furthermore, in light of the swift advancement observed in large language models (LLMs) such as ChatGPT, we are also exploring to use these models to generate rich and representative prompts for more calibrated scenes, which will be released for non-multimodal NeRF research

4. Experiments

4.1. Implementation Details

To verify the applicability and performance of the built dataset for NeRF methods, and provide a baseline for NeRF-based tasks, we train and evaluate recent NeRF [24], NeRF++ [45], Mip-NeRF [2], Mip-NeRF 360 [3], Mega-NeRF [36] and Ref-NeRF [37] on our datasets.

NeRF [24] presents the first continuous MLP-based neural network to represent the scene, that is able to synthesize semantic-consistent novel views by volume rendering.

NeRF++ [45] separately models the foreground and background neural representations to address the challenge of modeling large-scale unbounded scenes.

Mip-NeRF [2] reduces aliasing artifacts and improves NeRF’s [24] ability to represent fine details, by rendering anti-aliased conical frustums instead of rays.

Mip-NeRF 360 [3] uses a non-linear scene parameterization, online distillation, and a distortion-based regularizer, to model and produce realistic synthesized views for unbounded real-world scenes.

Mega-NeRF [36] proposes a framework for training large-scale 3D scenes by introducing a sparse structure and geometric clustering algorithm, to partition training pixels into different parallel NeRF submodules.

Ref-NeRF [37] improves the quality of appearance and normal in synthesized views of the scene, by a new parameterization and structuring of view-dependent outgoing radiance, as well as a regularizer on normal vectors.

Implementations. Since there is no official PyTorch implementation of NeRF [24], we use the widely recognized

third-party implementation [10]. But for other methods, we use the official implementation from GitHub.

In our dataset, reviewed posed images are numbered from 0 to $\#image - 1$ under timing sequence, and test views are evenly sampled according to the view ID. That is we sample one from every eight for testing, and the rest are used as training views (i.e., for testing: 0, 8, 16, 24, ...).

All training hyper-parameters follow the original paper’s settings in our experiments. Each scene is trained on a single Nvidia V100 GPU device for around 6-33 hours, depending on the time complexity of each method, and about 32 V100 GPU devices are used in parallel.

Multi-modal NeRF Synthesis. Since there is no available NeRF-based method for text-assisted fidelity novel view synthesis, inspired by CoCo-INR [44], we replace its image-based codebook with text-based codebook and apply it in NeRF [24] and CoCo-INR [44] as our benchmark. Specifically, we apply a text-based attentional coordinate module in front of the last MLP layer of volume rendering network, where the text-based codebook is encoded by textual prompts through the pre-trained CLIP [28] model, (*cf.* Fig. 6).

Evaluation Metrics. To evaluate the performance, we use three common metrics: Peak Signal-to-Noise Ratio (PSNR), Structural Similarity (SSIM) [39], and LPIPS [46] on novel view synthesis. Higher PSNR and SSIM mean better performance, while a lower LPIPS means better.

4.2. Novel View Synthesis

Benchmark. To establish a benchmark for the large-scale outdoor novel view synthesis, we comprehensively evaluate and report quantitative performances of the above six state-of-the-art methods in our dataset (*cf.* Tab. 3).

It can be seen that except for the failure of NeRF [24] in 4 scenes (7, 26, 27, and 33), other results show that NeRF can synthesize reasonable novel views, which means that OMMO dataset can support various NeRF-based methods. NeRF++ [45], Mip-NeRF 360 [3], Mega-NeRF [36], and Ref-NeRF [37] perform well on our dataset with an average PSNR of beyond 20, and can maintain the view consistency of each scene (*cf.* Fig. 7). Among them, Mip-NeRF 360 [3] can synthesize more realistic detailed texture features for large-scale scenes and its quantitative evaluation is more than 6 points higher than other methods on PSNR, SSIM, LPIPS. So, we not only provide a challenging large-scale outdoor scene dataset, but also further evaluate the performance and robustness of existing methods, which are not very discriminative on previous datasets especially for recent methods. Our benchmarks are open to all NeRF-based methods, and we are also ready to evaluate newer large-scale scene NeRF methods once they are proposed.

In particular, we notice that most of the scenes where NeRF fails are based on irregular camera trajectories, which suggests that NeRF may be more suitable for scenes captured

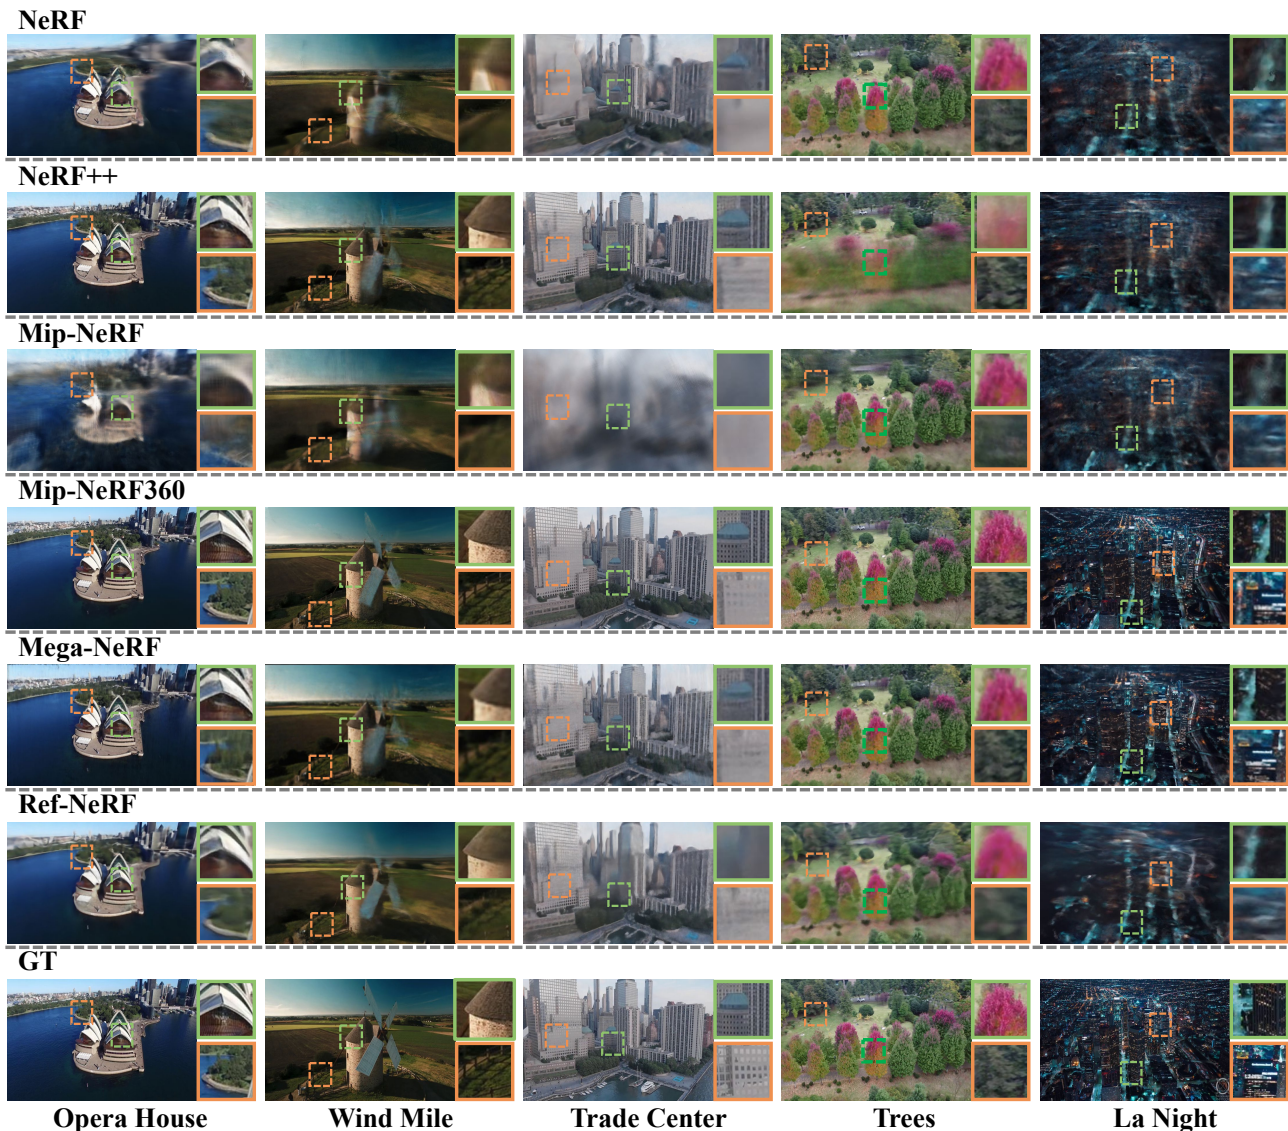


Figure 7. Qualitative visualization results for novel view synthesis (zoom-in for the best of views) on our OMMO dataset.

with stronger trajectory consistency constraints and more overlap (such as equidistant circular acquisitions). So we divide OMMO dataset into subsets according to the data types, and provide sub-benchmarks to study the most suitable setting for each method.

Sub-benchmark split by scene types. According to different scales of urban and natural scenes, we propose 4 sub-benchmarks for buildings, small areas, cities and natural scenes. It can be seen from Tab. 5 that all methods perform worse in cities than in smaller-scale subsets, i.e. buildings and small areas. These performance differences show that the large-scale scene implicit representation is still not as well resolved as for single objects or small scenes.

Sub-benchmark split by camera tracks. Circular camera trajectories tend to present better experiment perfor-

mance than other types, especially irregular ones. This is because 360-degree views contain richer texture features from different angles, and the focus of views is overlapped to maintain the view consistency.

Sub-benchmark split by lighting conditions. Intuitively, daytime scenes are richer in texture and easier to learn their representation than dark ones. However, we find that almost every method performs better on the low-light subset than on the normal-light subset. We visualize two different scans of the same scene generated by Mega-NeRF [36] during day and night (*cf.* Fig. 8). It is not difficult to see that in low light settings, the implicit network uses black areas to erase details when generating the RGB images, which reduces the synthesis difficulty and tricks the evaluation metrics, while synthesised poor depth map illustrates the network’s incapac-

Table 3. Benchmark for novel view synthesis. We present the performance of six state-of-the-art and representative methods on our dataset. \uparrow means the higher, the better.

Scene ID	Scene Types	Camera Tracks	Lighting Conditions	NeRF [24]			NeRF++ [45]			Mip-NeRF [2]			Mip-NeRF 360 [3]			Mega-NeRF [36]			Ref-NeRF [40]		
				PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
1	Buildings	Irregular	Day	16.93	0.369	0.744	16.86	0.359	0.780	16.84	0.369	0.793	13.91	0.311	0.771	16.12	0.341	0.782	15.10	0.344	0.755
2	Small area	Circles	Day	15.31	0.442	0.694	14.89	0.471	0.653	15.16	0.396	0.731	15.06	0.438	0.646	15.64	0.467	0.679	15.90	0.490	0.632
3	Cities	Lines	Day	14.38	0.278	0.556	14.64	0.294	0.547	14.56	0.288	0.533	14.25	0.309	0.526	15.21	0.325	0.517	15.44	0.371	0.526
4	Buildings	Circle	Night	25.39	0.859	0.431	27.47	0.898	0.380	21.78	0.758	0.469	27.68	0.943	0.292	23.36	0.855	0.419	27.86	0.905	0.404
5	Small area	Circles	Day	22.26	0.670	0.531	24.32	0.729	0.450	14.98	0.544	0.633	25.76	0.801	0.317	25.78	0.763	0.436	23.54	0.706	0.491
6	Natural scenes	Circles	Day	24.09	0.679	0.504	25.59	0.749	0.396	23.18	0.658	0.529	28.86	0.896	0.211	24.92	0.772	0.393	26.07	0.716	0.459
7	Buildings	Lines	Day	5.36	0.166	0.747	21.93	0.707	0.542	15.57	0.643	0.624	23.05	0.734	0.523	22.33	0.691	0.552	25.79	0.731	0.511
8	Cities	Circle	Day	21.14	0.496	0.594	22.91	0.568	0.509	19.82	0.462	0.638	25.07	0.714	0.354	16.65	0.478	0.431	21.21	0.489	0.606
9	Cities	Lines	Day	14.92	0.344	0.744	14.57	0.341	0.732	14.58	0.338	0.746	15.40	0.303	0.706	17.32	0.491	0.673	20.34	0.432	0.649
10	Cities	Irregular	Day	22.26	0.550	0.626	24.37	0.599	0.578	19.80	0.528	0.643	26.68	0.719	0.420	21.78	0.615	0.558	24.23	0.578	0.597
11	Buildings	Circles	Night	22.36	0.816	0.420	24.61	0.852	0.342	22.81	0.822	0.423	27.06	0.931	0.217	24.37	0.844	0.392	23.81	0.843	0.355
12	Small area	Circles	Day	22.41	0.594	0.533	24.29	0.675	0.447	22.13	0.601	0.526	28.12	0.825	0.274	21.60	0.619	0.493	23.06	0.604	0.524
13	Buildings	Lines	Day	22.27	0.592	0.608	23.52	0.623	0.581	18.90	0.537	0.673	26.63	0.771	0.403	25.50	0.722	0.517	23.29	0.605	0.594
14	Small area	Lines	Day	19.85	0.554	0.569	23.89	0.737	0.417	17.06	0.481	0.655	28.06	0.894	0.224	24.42	0.746	0.411	21.76	0.625	0.508
15	Small area	Circles	Day	20.35	0.527	0.552	21.71	0.612	0.490	19.44	0.489	0.594	28.63	0.888	0.179	22.69	0.665	0.445	20.33	0.497	0.576
16	Natural scenes	Circles	Day	17.86	0.397	0.631	18.75	0.405	0.597	18.49	0.399	0.610	10.01	0.344	0.850	20.26	0.532	0.509	19.64	0.428	0.572
17	Natural scenes	Circles	Day	22.02	0.571	0.610	24.20	0.671	0.461	17.01	0.526	0.696	29.53	0.833	0.247	17.23	0.574	0.529	23.17	0.589	0.529
18	Small area	Lines	Day	26.06	0.754	0.428	25.57	0.730	0.461	24.61	0.732	0.469	28.55	0.855	0.265	24.76	0.733	0.448	22.79	0.674	0.569
19	Small area	Circles	Day	14.20	0.399	0.726	13.86	0.373	0.703	13.84	0.394	0.738	14.72	0.367	0.676	23.81	0.682	0.465	14.34	0.386	0.691
20	Cities	Circles	Day	22.84	0.613	0.499	23.28	0.642	0.475	22.41	0.603	0.519	28.33	0.862	0.228	21.11	0.633	0.490	21.54	0.553	0.574
21	Natural scenes	Circles	Day	22.59	0.514	0.532	21.84	0.473	0.593	22.31	0.513	0.537	25.64	0.747	0.344	21.92	0.506	0.478	21.07	0.436	0.672
22	Buildings	Lines	Day	16.53	0.466	0.733	20.66	0.558	0.575	13.37	0.420	0.776	24.79	0.766	0.362	20.84	0.597	0.527	20.31	0.530	0.615
23	Natural scenes	Lines	Day	18.99	0.405	0.669	19.51	0.417	0.597	18.09	0.389	0.671	21.25	0.514	0.539	20.13	0.438	0.585	19.94	0.409	0.622
24	Natural scenes	Lines	Day	19.32	0.386	0.696	23.14	0.522	0.535	16.89	0.374	0.715	25.86	0.707	0.373	23.87	0.563	0.518	22.17	0.452	0.616
25	Natural scenes	Lines	Day	24.72	0.550	0.528	22.42	0.509	0.613	24.24	0.541	0.542	28.91	0.789	0.306	25.98	0.629	0.457	23.62	0.502	0.598
26	Buildings	Irregular	Day	8.56	0.242	0.564	19.94	0.586	0.513	13.43	0.353	0.688	14.59	0.459	0.626	19.23	0.669	0.467	21.00	0.615	0.489
27	Cities	Irregular	Day	4.54	0.006	0.705	21.25	0.548	0.546	14.82	0.453	0.674	21.26	0.599	0.235	20.59	0.606	0.543	20.82	0.519	0.590
28	Small area	Circles	Day	24.48	0.660	0.479	23.28	0.642	0.475	24.76	0.659	0.406	29.62	0.874	0.240	25.87	0.723	0.442	22.17	0.452	0.616
29	Buildings	Circle	Day	22.98	0.608	0.540	23.17	0.617	0.529	23.01	0.609	0.539	25.51	0.740	0.400	21.57	0.611	0.557	21.11	0.543	0.631
30	Natural scenes	Irregular	Day	20.23	0.522	0.605	23.27	0.639	0.476	18.63	0.461	0.675	26.54	0.837	0.296	24.04	0.686	0.459	21.62	0.535	0.586
31	Cities	Circles	Night	18.97	0.365	0.645	19.05	0.371	0.643	18.91	0.358	0.659	13.08	0.234	0.708	20.93	0.596	0.545	19.18	0.372	0.645
32	Cities	Irregular	Day	17.99	0.582	0.621	18.99	0.605	0.540	11.28	0.424	0.687	17.16	0.566	0.601	21.29	0.702	0.475	18.98	0.595	0.565
33	Cities	Irregular	Day	5.79	0.007	0.745	20.19	0.497	0.597	14.31	0.42	0.755	22.76	0.629	0.457	22.89	0.635	0.478	21.23	0.522	0.578
Mean	-	-	-	18.72	0.484	0.600	21.45	0.576	0.538	18.39	0.501	0.623	23.10	0.672	0.418	21.63	0.621	0.508	21.28	0.546	0.574

Table 4. More sub-benchmarks for novel view synthesis. We divide our dataset into subsets based on different scene types, camera trajectories, and lighting conditions, and provide sub-benchmarks under different settings. \uparrow means the higher, the better.

Sub-benchmark	NeRF [24]			NeRF++ [45]			Mip-NeRF [2]			Mip-NeRF 360 [3]			Mega-NeRF [36]			Ref-NeRF [40]		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Buildings	17.32	0.501	0.605	22.24	0.644	0.528	17.82	0.546	0.636	22.85	0.704	0.444	21.05	0.650	0.511	22.30	0.633	0.541
Small areas	20.88	0.579	0.561	21.66	0.621	0.514	19.44	0.545	0.588	24.89	0.742	0.358	22.90	0.668	0.486	20.56	0.553	0.582
Cities	15.87	0.360	0.637	19.92	0.496	0.574	16.72	0.430	0.650	20.44	0.548	0.471	19.75	0.565	0.523	20.33	0.492	0.592
Natural scenes	21.23	0.503	0.597	22.34	0.548	0.534	19.86	0.483	0.622	24.58	0.708	0.396	22.29	0.588	0.504	22.16	0.508	0.582
Circles	21.08	0.573	0.559	22.00	0.609	0.508	19.80	0.545	0.581	23.81	0.713	0.386	21.74	0.647	0.483	21.53	0.564	0.556
Lines	18.24	0.450	0.628	20.99	0.544	0.560	17.79	0.474	0.640	23.68	0.664	0.423	22.04	0.594	0.521	21.55	0.533	0.581
Irregular	14.91	0.361	0.644	21.01	0.556	0.570	16.52	0.452	0.682	21.05	0.608	0.476	20.94	0.608	0.540	20.51	0.531	0.599
Day	18.37	0.465	0.610	21.23	0.563	0.547	18.12	0.487	0.634	23.15	0.670	0.420	21.51	0.607	0.514	21.05	0.531	0.585
Night	22.24	0.680	0.499	23.71	0.707	0.455	21.17	0.646	0.517	22.61	0.703	0.406	22.89	0.765	0.452	23.62	0.707	0.468

Table 5. Performance of novel view synthesis on BlendMVS.

Dataset	Metrics	NeRF	NeRF++	Mip-NeRF	Mip-NeRF 360	Mega-NeRF	Ref-NeRF
		BlendMVS	PSNR \uparrow 22.11	22.59	22.07	24.05	23.64
	SSIM \uparrow 0.446	0.469	0.444	0.594	0.577	0.339	
	LPIPS \downarrow 0.588	0.572	0.589	0.389	0.471	0.726	
OMMO	PSNR \uparrow 18.72	21.45	18.39	23.10	21.63	21.28	
	SSIM \uparrow 0.484	0.576	0.501	0.672	0.621	0.546	
	LPIPS \downarrow 0.600	0.538	0.623	0.418	0.508	0.574	

ity to understand and represent the scene. So efficient low-light NeRF methods are urged to be proposed. Results across diverse scene types, camera trajectories, and lighting conditions have demonstrated the robustness of each method as shown in Tab. 3 and 4. It is noteworthy that the calibration of nocturnal scenes frequently encounters challenges stemming from limited overlap and textural cues. To undertake an initial assessment of nocturnal performance, a subset of three 4K scenes was meticulously chosen from an ensemble of eleven calibrated nighttime scenes. Furthermore, the forthcoming release of five additional 2K scenes and three 1080P scenes, as an adjunctive nocturnal sub-baseline, is contingent upon the completion of manual prompt annotations. These resources are poised to fortify future investigations in the




Figure 8. Different scans of the same scene during day and night. Both RGB and depth images are synthesized by Mega-NeRF [36].

realm of Dark-NeRF research.

Performance on Blend-MVS. We evaluate baselines on the first three scenes from the large scene part of the BlendMVS dataset. In alignment with our findings, Mip-NeRF 360 demonstrates optimal performance. However, as this dataset is synthetic and the scene scale is relatively




Figure 9. Examples of various scene representations from our dataset through different methods.

small, each method presents similar performance with PSNR centered from 22 to 24 (*cf.* Tab. 5). This imposes a growing expectation for a more discerning, real-world, large-scale outdoor scene dataset.

4.3. Scene Representation

To verify that our dataset can be used by a series of extensions of NeRF or traditional method to generate richer explicit 3D representations, such as mesh, normal surface, and dense point cloud, we conducted experiments on Mega-NeRF [36], Instant-NGP [25], which are based on implicit representations and COLMAP [31, 32] respectively. It can be seen from Fig. 9 that neither the implicit network nor the feature matching reconstruction method can reconstruct the large scene finely. Theoretically, the advantage of implicit scene representation is that, the scene can be reconstructed with high resolution benefiting from the continuous representation. However, due to the limitation of the viewpoint, the 3D information of the farthest edges of large-scale scenes is difficult to converge and has low accuracy. Therefore, it is challenging to quantitatively evaluate the explicit 3D representation of the entire scene. The benchmark for this aspect is still under construction, but our dataset provides diverse 3D representation data that can be used to compare the reconstruction quality of the central region of the scene after processing. We are now able to generate dense point clouds for large-scale scenes, with 9 million to 17 million points per scene, by using sub-region reconstruction, multi-region fusion, and human de-noising. We employ the Iterative Closest Point (ICP) as the evaluation metric, and BakedSDF achieves 0.210. More methods are being evaluated to establish additional 3D reconstruction benchmarks for further research. We will make this part of the data publicly available for further research. Please refer to the supplementary materials for more visual 3D representation data.

4.4. Multi-modal NeRF Synthesis

We design the text-based coordinate attentional block shown in Fig. 6 to test the quality of prompt annotations and verify whether favorable prompt annotations can provide the

Table 6. Performance comparison of with or without textual prompts for novel view synthesis on our OMMO dataset. We report the results of each method and the performance improvement after injecting textual prompts. \uparrow means the higher, the better.

Mthod	Without / With Textual Prompts			Improvement (%) \uparrow		
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	PSNR	SSIM	LPIPS
NeRF [24]	18.72 / 19.01	0.484 / 0.500	0.600 / 0.591	1.5	3.2	1.5
CoCo-INR [44]	16.80 / 16.97	0.489 / 0.490	0.681 / 0.678	1.0	0.2	0.4

missing global semantic information in the point-wise NeRF training process and improve performance.

It can be seen from Tab. 6, the performance of both NeRF [24] and CoCo-INR [44] methods have improved even for not-well-designed injecting methods. Since the textual prompts contain more global features about rich geometry or appearance information, which are shared by different views in the scene to guarantee the network to generate view-consistency results. We hope to inspire more image-text multi-modal NeRF methods to synthesize photo-realistic rendering results and decent geometry by exploring effective ways to make full use of textual prompts. The benchmark on each scene and the sub-benchmarks on different scene types are shown in supplementary.

5. Discussion

Conclusion. We introduce a well-selected large-scale outdoor multi-modal fly-view dataset, OMMO, to address the problem of no widely-used benchmark for outdoor NeRF-based methods. The built OMMO surpasses the previous datasets in several key indicators such as quantity, quality and variety, by providing 33 real-world scenes with more than 14K posed images and text description. With the help of our cost-effective data collection pipeline, it is easy to expand our dataset by continuously converting new internet videos into NeRF-purpose training data. We provide benchmarks on multiple tasks such as novel view synthesis, implicit scene representations, and multi-modal synthesis by evaluating existing methods. Experiments show that our dataset can well support mainstream NeRF-based tasks.

Limitation. Scenes with low-illumination, rain and fog are still few in the current dataset due to the limited calibration and reconstruction ability of COLMAP [31, 32]. We will try more reconstruction and calibration methods to solve this problem. Meanwhile, we are continuing to expand our dataset thanks to our cost-effective pipeline. Besides, it is a promising direction to study the large-scale scene NeRF generalization and generate/edit large-scale scenes from text based on our paired scene images and textual prompts.

Acknowledgements

This work is supported by National Natural Science Foundation of China (No.62101137 and 62071127), Shanghai Natural Science Foundation(No. 23ZR1402900), Zhejiang Lab Project (No. 2021KH0AB05).

References

- [1] Sameer Agarwal, Yasutaka Furukawa, Noah Snavely, Ian Simon, Brian Curless, Steven M Seitz, and Richard Szeliski. Building rome in a day. *Communications of the ACM*, 54(10):105–112, 2011. 2, 3
- [2] 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. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. 6, 8
- [3] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022. 6, 8
- [4] Aljaz Bozic, Pablo Palafox, Justus Thies, Angela Dai, and Matthias Nießner. Transformerfusion: Monocular rgb scene reconstruction using transformers. *Advances in Neural Information Processing Systems*, 34:1403–1414, 2021. 4
- [5] Di Chen, Yu Liu, Lianghua Huang, Bin Wang, and Pan Pan. Geoaug: Data augmentation for few-shot nerf with geometry constraints. In *European Conference on Computer Vision*, pages 322–337. Springer, 2022. 1
- [6] David Crandall, Andrew Owens, Noah Snavely, and Dan Huttenlocher. Discrete-continuous optimization for large-scale structure from motion. In *CVPR 2011*, pages 3001–3008. IEEE, 2011. 2
- [7] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. 2, 3
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 2
- [9] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12882–12891, 2022. 1
- [10] Github. nerf-pytorch. <https://github.com/yenchenlin/nerf-pytorch>. 6
- [11] Michael Grubinger, Paul Clough, Henning Müller, and Thomas Deselaers. The iapr tc-12 benchmark: A new evaluation resource for visual information systems. In *International workshop ontoImage*, volume 2, 2006. 5
- [12] Haoyu Guo, Sida Peng, Haotong Lin, Qianqian Wang, Guofeng Zhang, Hujun Bao, and Xiaowei Zhou. Neural 3d scene reconstruction with the manhattan-world assumption. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5511–5520, 2022. 4
- [13] Jared Heinly, Johannes L Schonberger, Enrique Dunn, and Jan-Michael Frahm. Reconstructing the world* in six days*(as captured by the yahoo 100 million image dataset). In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3287–3295, 2015. 2, 3
- [14] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 406–413, 2014. 1, 2, 3
- [15] Mohammad Mahdi Johari, Yann Lepoittevin, and François Fleuret. Geonerf: Generalizing nerf with geometry priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18365–18375, 2022. 1
- [16] Kacper Kania, Kwang Moo Yi, Marek Kowalski, Tomasz Trzcíński, and Andrea Tagliasacchi. Conerf: Controllable neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18623–18632, 2022. 3
- [17] Hanul Kim, Su-Min Choi, Chang-Su Kim, and Yeong Jun Koh. Representative color transform for image enhancement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4459–4468, 2021. 4
- [18] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4):1–13, 2017. 2, 3
- [19] Jonáš Kulháněk, Erik Derner, Torsten Sattler, and Robert Babuška. Viewformer: Nerf-free neural rendering from few images using transformers. *arXiv preprint arXiv:2203.10157*, 2022. 3
- [20] Quentin Legros, Julián Tachella, Rachael Tobin, Aongus McCarthy, Sylvain Meignen, Gerald S Buller, Yoann Altmann, Stephen McLaughlin, and Michael E Davies. Robust 3d reconstruction of dynamic scenes from single-photon lidar using beta-divergences. *IEEE Transactions on Image Processing*, 30:1716–1727, 2020. 4
- [21] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 2
- [22] Yilin Liu, Fuyou Xue, and Hui Huang. Urbanscene3d: A large scale urban scene dataset and simulator. *arXiv preprint arXiv:2107.04286*, 2021. 2, 3
- [23] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7210–7219, 2021. 3
- [24] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 1, 2, 3, 6, 8, 9
- [25] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, July 2022. 9
- [26] Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi SM Sajjadi, Andreas Geiger, and Noha Radwan. Reg-nerf: Regularizing neural radiance fields for view synthesis

- from sparse inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5480–5490, 2022. 3
- [27] Michael Oechsle, Songyou Peng, and Andreas Geiger. Unisurf: Unifying neural implicit surfaces and radiance fields for multi-view reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5589–5599, 2021. 3
- [28] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 4, 5, 6
- [29] Eduard Ramon, Gil Triginer, Janna Escur, Albert Pumarola, Jaime Garcia, Xavier Giro-i Nieto, and Francesc Moreno-Noguer. H3d-net: Few-shot high-fidelity 3d head reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5620–5629, 2021. 3
- [30] Ramon Sanabria, Ozan Caglayan, Shruti Palaskar, Desmond Elliott, Loïc Barrault, Lucia Specia, and Florian Metze. How2: a large-scale dataset for multimodal language understanding. In *Proceedings of the Workshop on Visually Grounded Interaction and Language (ViGIL)*. NeurIPS, 2018. 5
- [31] Johannes L Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4104–4113, 2016. 4, 9
- [32] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *European conference on computer vision*, pages 501–518. Springer, 2016. 4, 9
- [33] Hossein Talebi and Peyman Milanfar. Nima: Neural image assessment. *IEEE transactions on image processing*, 27(8):3998–4011, 2018. 4
- [34] Matthew Tancik, Vincent Casser, Xinchun Yan, Sabeek Pradhan, Ben Mildenhall, Pratul P Srinivasan, Jonathan T Barron, and Henrik Kretschmar. Block-nerf: Scalable large scene neural view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8248–8258, 2022. 3
- [35] Matthew Tancik, Vincent Casser, Xinchun Yan, Sabeek Pradhan, Ben P. Mildenhall, Pratul P. Srinivasan, Jonathan T. Barron, and Henrik Kretschmar. Block-nerf: Scalable large scene neural view synthesis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 8238–8248. IEEE, 2022. 2
- [36] Haithem Turki, Deva Ramanan, and Mahadev Satyanarayanan. Mega-nerf: Scalable construction of large-scale nerfs for virtual fly-throughs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12922–12931, 2022. 1, 2, 3, 6, 7, 8, 9
- [37] Dor Verbin, Peter Hedman, Ben Mildenhall, Todd Zickler, Jonathan T Barron, and Pratul P Srinivasan. Ref-nerf: Structured view-dependent appearance for neural radiance fields. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5481–5490. IEEE, 2022. 6
- [38] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *Advances in Neural Information Processing Systems*, 34:27171–27183, 2021. 3
- [39] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6
- [40] Yuanbo Xiangli, Linning Xu, Xingang Pan, Nanxuan Zhao, Anyi Rao, Christian Theobalt, Bo Dai, and Dahua Lin. Bungeenerf: Progressive neural radiance field for extreme multi-scale scene rendering. In *The European Conference on Computer Vision (ECCV)*, volume 2, 2022. 1, 2, 3, 8
- [41] Guo-Wei Yang, Wen-Yang Zhou, Hao-Yang Peng, Dun Liang, Tai-Jiang Mu, and Shi-Min Hu. Recursive-nerf: An efficient and dynamically growing nerf. *IEEE Transactions on Visualization and Computer Graphics*, 2022. 3
- [42] Yao Yao, Zixin Luo, Shiwei Li, Jingyang Zhang, Yufan Ren, Lei Zhou, Tian Fang, and Long Quan. Blendedmvs: A large-scale dataset for generalized multi-view stereo networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1790–1799, 2020. 1, 2, 3
- [43] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021. 1, 3
- [44] Fukun Yin, Wen Liu, Zilong Huang, Pei Cheng, Tao Chen, and Gang YU. Coordinates are not lonely—codebook prior helps implicit neural 3d representations. *arXiv preprint arXiv:2210.11170*, 2022. 1, 3, 6, 9
- [45] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv preprint arXiv:2010.07492*, 2020. 3, 6, 8
- [46] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. 6