# City-on-Web: Real-time Neural Rendering of Large-scale Scenes on the Web

Kaiwen Song<sup>1</sup><sup>(0)</sup>, Xiaoyi Zeng<sup>1</sup><sup>(0)</sup>, Chenqu Ren<sup>2</sup><sup>(0)</sup>, and Juyong Zhang<sup>1</sup><sup>\*</sup><sup>(0)</sup>

<sup>1</sup> University of Science and Technology of China <sup>2</sup> East China Normal University

## 1 Derivation of Block-based Volume Rendering

Traditional volume rendering methods, like those used in Block-NeRF [9] and Mega-NeRF [11], necessitate accessing radiance and opacity for all sampling points, even if the points belong to different blocks. This requirement means that when the rendering resources of multiple shaders cannot communicate with each other, their volume rendering algorithm cannot work. On the contrary, our block-based volume rendering algorithm is tailored for such resource-independent settings, allowing for the independent rendering of each block before blending their outputs. In this section, we will prove the equivalence between this resourceindependent volume rendering and the traditional volume rendering under Lambertian conditions.

In traditional volume rendering, consider a ray  $\boldsymbol{r}$  with M sampling points. For each sampling point  $p_i$ , the diffuse color, feature, and opacity are represented as  $\boldsymbol{c}_i$ ,  $\boldsymbol{f}_i$ , and  $\alpha_i$  respectively. The final diffuse color  $\boldsymbol{C}_d$  of the ray  $\boldsymbol{r}$  can be obtained by

$$\hat{\boldsymbol{C}}_{d}(\boldsymbol{r}) = \sum_{i=1}^{M} \prod_{j=1}^{i-1} (1 - \alpha_{j}) \alpha_{i} \boldsymbol{c}_{i}$$
(1)

In block-based volume rendering, let's assume the ray, with M sampling points, traverses through K blocks. Within block  $\mathcal{B}_k$ , there are  $N_k$  sampling points. For sampling point  $p_i^k$  in block  $\mathcal{B}_k$ , the diffuse color, feature, and opacity output by the shader or submodel of block  $\mathcal{B}_k$  are represented as  $c_i^k$ ,  $f_i^k$ , and  $\alpha_i^k$ , respectively.

$$c_i^k = c_{N_1 + \dots + N_{k-1} + i}$$

$$f_i^k = f_{N_1 + \dots + N_{k-1} + i}$$

$$\alpha_i^k = \alpha_{N_1 + \dots + N_{k-1} + i}$$
(2)

<sup>\*</sup> Corresponding Author

# 2 K. Song et al.

Then, for this ray, diffuse color  $C_d^k$ , specular feature  $F^k$  and opacity  $\alpha^k$  of block  $\mathcal{B}_k$  are calculated as follows:

$$\begin{aligned} \boldsymbol{C}_{d}^{k} &= \sum_{i=1}^{N_{k}} \prod_{j=1}^{i-1} (1 - \alpha_{j}^{k}) \cdot \alpha_{i}^{k} \boldsymbol{c}_{i}^{k} \\ \boldsymbol{F}^{k} &= \sum_{i=1}^{N_{k}} \prod_{j=1}^{i-1} (1 - \alpha_{j}^{k}) \cdot \alpha_{i}^{k} \boldsymbol{f}_{i}^{k} \\ \boldsymbol{\alpha}^{k} &= \sum_{i=1}^{N_{k}} \prod_{j=1}^{i-1} (1 - \alpha_{j}^{k}) \cdot \alpha_{i}^{k} \end{aligned}$$
(3)

By blending the rendering results of each block using Eq. (8), the final diffuse color  $C_d$  and specular feature F of the ray r can be obtained.

$$C_d(\mathbf{r}) = \sum_{k=1}^K \prod_{j=1}^{k-1} (1 - \alpha^j) \cdot C_d^k$$

$$F(\mathbf{r}) = \sum_{k=1}^K \prod_{j=1}^{k-1} (1 - \alpha^j) \cdot F^k$$
(4)

From Eq. (3), we get the following equation:

$$1 - \alpha^{k} = 1 - \sum_{i=1}^{N_{k}} \prod_{j=1}^{i-1} (1 - \alpha_{j}^{k}) \cdot \alpha_{i}^{k}$$
  

$$= 1 - \alpha_{1}^{k} - (1 - \alpha_{1}^{k})\alpha_{2}^{k} - (1 - \alpha_{1}^{k})(1 - \alpha_{2}^{k})\alpha_{3}^{k} - \cdots$$
  

$$= (1 - \alpha_{1}^{k})(1 - \alpha_{2}^{k} - (1 - \alpha_{2}^{k})\alpha_{3}^{k} - \cdots)$$
  

$$= (1 - \alpha_{1}^{k})(1 - \alpha_{2}^{k})(1 - \alpha_{3}^{k} - \cdots)$$
  

$$= \cdots$$
  

$$= \prod_{i=1}^{N_{k}} (1 - \alpha_{i}^{k})$$
(5)

Consequently, we can obtain the following by substituting Eq. (5) into Eq. (4).

$$C_{d}(\mathbf{r}) = \sum_{k=1}^{K} \prod_{j=1}^{k-1} (1 - \alpha^{j}) \cdot C_{d}^{k}$$
  
=  $\sum_{k=1}^{K} \prod_{j=1}^{k-1} \prod_{i=1}^{N_{j}} (1 - \alpha_{i}^{j}) \cdot C_{d}^{k}$   
=  $\sum_{k=1}^{K} \prod_{i=1}^{N_{1} + \dots + N_{k-1}} (1 - \alpha_{i}) \cdot C_{d}^{k}$  (6)

By substituting  $C_d^k$  from Eq. (3) into Eq. (6), we demonstrate the equivalence between Eq. (1) and Eq. (4).

$$\begin{split} C_{d}(r) &= \sum_{k=1}^{K} \prod_{l=1}^{N_{1}+\dots+N_{k-1}} (1-\alpha_{l}) \sum_{i=1}^{N_{k}} \prod_{j=1}^{i-1} (1-\alpha_{j}^{k}) \cdot \alpha_{i}^{k} c_{i}^{k} \\ &= \underbrace{\alpha_{1}^{1} c_{1}^{1} + \dots + \prod_{i=1}^{N_{1}-1} (1-\alpha_{i}^{1}) \alpha_{N_{1}}^{1} c_{N_{1}}^{1}}_{block1} \\ &+ \underbrace{\prod_{j=1}^{N_{1}} (1-\alpha_{j}^{1})}_{l} \left( \alpha_{1}^{2} c_{1}^{2} + \dots + \prod_{i=1}^{N_{2}-1} (1-\alpha_{i}^{2}) \alpha_{N_{2}}^{2} c_{N_{2}}^{2} \right) + \dots \\ &= \underbrace{\prod_{j=1}^{N_{1}} (1-\alpha_{j}^{1}) \cdots \prod_{j=1}^{N_{K-1}} (1-\alpha_{j}^{K-1})}_{block2} \left( \alpha_{1}^{K} c_{1}^{K} + \dots + \prod_{i=1}^{N_{K}-1} (1-\alpha_{i}^{K}) \alpha_{N_{K}}^{K} c_{N_{K}}^{K} \right) \\ &= \underbrace{\alpha_{1}^{1} c_{1}^{1} + \dots + \prod_{i=1}^{N_{1}-1} (1-\alpha_{i}^{1}) \alpha_{N_{1}}^{1} c_{N_{1}}^{1}}_{block1} \\ &+ \underbrace{\prod_{j=1}^{N_{1}} (1-\alpha_{j}^{1}) \alpha_{1}^{2} c_{1}^{2} + \dots + \prod_{j=1}^{N_{1}} (1-\alpha_{j}^{1}) \prod_{i=1}^{N_{2}-1} (1-\alpha_{i}^{2}) \alpha_{N_{2}}^{2} c_{N_{2}}^{2} + \dots \\ &= \underbrace{\alpha_{1}^{1} c_{1}^{1} + \dots + \prod_{i=1}^{N_{K-1}} (1-\alpha_{i}^{1}) \alpha_{N_{1}}^{1} c_{N_{1}}^{1}}_{block2} \\ &+ \underbrace{\prod_{j=1}^{N_{1}} (1-\alpha_{j}^{1}) \cdots \prod_{i=1}^{N_{K-1}} (1-\alpha_{i}^{K-1}) \alpha_{1}^{K} c_{1}^{K} + \dots \\ &= \underbrace{\alpha_{1} c_{1} + \dots + \prod_{i=1}^{N_{K-1}} (1-\alpha_{i}^{K-1}) \alpha_{N_{K}}^{K-1} (1-\alpha_{i}^{K}) \alpha_{N_{K}}^{K} c_{N_{K}}^{K})}_{blockK} \\ &= \underbrace{\alpha_{1} c_{1} + \dots + \prod_{i=1}^{N_{1}-1} (1-\alpha_{i}) \alpha_{N_{1}} c_{N_{1}}}_{block1} \\ &+ \underbrace{\prod_{j=1}^{N_{1}} (1-\alpha_{j}) \alpha_{N_{1}+1} c_{N_{1}+1} + \dots + \prod_{i=1}^{N_{1}+N_{2}-1} (1-\alpha_{i}) \alpha_{N_{1}+N_{2}} c_{N_{1}+N_{2}} + \dots \\ &= \underbrace{\alpha_{1} c_{1} + \dots + \prod_{i=1}^{N_{1}-1} (1-\alpha_{i}) \alpha_{N_{1}} c_{N_{1}}}_{block1} \\ &+ \underbrace{\prod_{i=1}^{N_{1}} (1-\alpha_{i}) \alpha_{N_{1}+1} c_{N_{1}+1} + \dots + \prod_{i=1}^{N_{1}+N_{2}-1} (1-\alpha_{i}) \alpha_{N_{1}+N_{2}} c_{N_{1}+N_{2}} + \dots \\ &= \underbrace{\sum_{i=1}^{N_{1}} (1-\alpha_{i}) \alpha_{N_{1}+1} c_{N_{1}+1} + \dots + \prod_{i=1}^{N_{1}+N_{2}-1} (1-\alpha_{i}) \alpha_{N_{1}+N_{2}} c_{N_{1}+N_{2}} + \dots \\ &= \underbrace{\sum_{i=1}^{N_{1}} (1-\alpha_{i}) \alpha_{N_{1}+1} c_{N_{1}+1} + \dots + \underbrace{\sum_{i=1}^{N_{1}+N_{2}-1} (1-\alpha_{i}) \alpha_{N_{1}+N_{2}} c_{N_{1}+N_{2}} + \dots \\ &= \underbrace{\sum_{i=1}^{N_{1}} (1-\alpha_{i}) \alpha_{N_{1}+1} c_{N_{1}+1} + \dots + \underbrace{\sum_{i=1}^{N_{1}+N_{2}-1} (1-\alpha_{i}) \alpha_{N_{1}+N_{2}} c_{N_{1}+N_{2}} + \dots \\ &= \underbrace{\sum_{i=1}^{N_{1}} (1-\alpha_{i}) \alpha_{N_{1}+1} c_{N_{1}+1} + \dots + \underbrace{\sum_{i=1}^{N_{$$

4 K. Song et al.

$$+ \underbrace{\prod_{j=1}^{N_{1}+\dots+N_{K-1}} (1-\alpha_{j})\alpha_{N_{1}+\dots+N_{K-1}+1}\boldsymbol{c}_{N_{1}+\dots+N_{K-1}+1} + \cdots}_{blockK} \\ + \underbrace{\prod_{i=1}^{N_{1}+\dots+N_{K}-1} (1-\alpha_{i})\alpha_{N_{1}+\dots+N_{K}}\boldsymbol{c}_{N_{1}+\dots+N_{K}}}_{blockK} \\ = \underbrace{\sum_{i=1}^{N_{1}+\dots+N_{K}} \prod_{j=1}^{i-1} (1-\alpha_{j})\alpha_{i}\boldsymbol{c}_{i}}_{j=1} \\ = \sum_{i=1}^{M} \prod_{j=1}^{i-1} (1-\alpha_{j})\alpha_{i}\boldsymbol{c}_{i} = \hat{\boldsymbol{C}}_{d}(\boldsymbol{r})$$

It can be observed that block-based volume rendering approach in the resourceindependent environment, such as in our case where multiple shaders are used to render the entire scene, is equivalent to the MERF's volume rendering method [7] without deferred rendering.

(8)

# 2 More Details on Experiments

In this section, we provide more details on our custom *Campus* dataset, implemention details and settings of comparison methods.

#### 2.1 Dataset

The Campus dataset is captured at an altitude of about 180 meters, covering an area of approximately 960,000  $m^2$ . We adopt a circular data capture method for areial photography, as shown in Fig. 1. We find that this method often results in a higher overlap rate, allowing for a more accurate estimation of camera poses. Our dataset was captured over 8 hours on a cloudy day, with a fixed exposure setting to ensure almost identical appearance of photos taken at different time. We used Colmap [8] to estimate camera poses. Feature matching was done using a vocabulary tree, followed by a hierarchical mapper followed by a few iterations of triangulation and bundle adjustment to estimate camera poses.

#### 2.2 Implementation Details

For blocks at the boundaries of the entire scene, an unbounded scene representation is required to represent areas outside the block boundaries. We follow the same approach as MERF [7] to compute ray-AABB intersections trivially. To be specific, we employ the scene contraction function to project the scene external to the unit sphere into a cube, which has a radius of 2. The definition of the j - th coordinate for a contracted point is as follows:

$$\operatorname{contract}(\mathbf{x})_{j} = \begin{cases} x_{j} & \text{if } \|\mathbf{x}\|_{\infty} \leq 1\\ \frac{x_{j}}{\|\mathbf{x}\|_{\infty}} & \text{if } x_{j} \neq \|\mathbf{x}\|_{\infty} > 1\\ \left(2 - \frac{1}{|x_{j}|}\right) \frac{x_{j}}{|x_{j}|} & \text{if } x_{j} = \|\mathbf{x}\|_{\infty} > 1 \end{cases}$$
(9)

Our method takes posed multi-view images captured using a fly-through camera as input. The training code is built on the nerfstudio framework [10] with tiny-cuda-nn extension. Our real-time viewer is a JavaScript web application whose rendering is implemented through GLSL. We set the  $512^3$  resolution for the voxel and  $2048^2$  resolution for the triplane within each block. We use a 4layer MLP with 64 hidden dimensions as an encoder after multi-resolution hash encoding to output density, color, and specular feature. Moreover, a 3-layer MLP with 16 hidden dimensions tiny deferred MLP is developed to predict residual view-dependent color. We sample 16384 rays per batch and use Adam optimizer with an initial learning rate of  $1 \times 10^{-2}$  decaying exponentially to  $1 \times 10^{-3}$ . The global deferred MLP is a 6-layer MLP with 128 hidden dimensions. Our model is trained with 50k iterations on one NVIDIA A100 GPU. We split the scene into 24 blocks for *Campus* scene, and split other scenes into four blocks to reconstruct these scenes. We perform qualitative comparisons between our method and existing SOTA methods for large-scale reconstruction. The Campus dataset is partitioned into six parts. NeRFacto, Instant-NGP, and Grid-NeRF were applied to reconstruct one of these parts. NeRFacto and Instant-NGP are utilized with the highest hash encoding resolution of  $8192^3$ . Mega-NeRF divides the *Campus* scene into 24 blocks and splits another dataset into four blocks to evaluate its performance.

We use an A100 GPU for training. In Sections 5.1 and 5.2, we perform training for 50,000 iterations with a batch size of 16384 pixels for the *Campus* dataset, taking approximately 48 hours, while the training for other datasets takes around 24 hours. The experiments described in Sections 5.2 and 5.4 are subjected to training for 30,000 iterations with the same batch size, which is completed in approximately 12 hours. Training losses are initially balanced with  $\lambda_1 = 1.0, \lambda_2 = 1.0, \lambda_3 = 0.01, \lambda_4 = 0.05, \lambda_5 = 0.001$  and sample 2<sup>14</sup> samples for computing sparsity loss. During LOD genreation phase, we utilize the same loss function and hyperparameters as those employed during the training stage. We freeze the training of submodels for 5,000 iterations and then refine them and shared global deferred mlp jointly for an additional 10,000 iterations.

#### 2.3 Comparative Method Settings

In Sec. 5.1, we adopt the official implementations of Mega-NeRF [11], NeR-Facto [10], and Grid-NeRF [12]. Additionally, we use an unofficial implementation of Instant-NGP<sup>3</sup>. Specifically, NeRFacto is trained with a batch size of

<sup>&</sup>lt;sup>3</sup> https://github.com/ashawkey/torch-ngp



Fig. 1: Visulization of Campus Dataset.

65,536 for 30,000 iterations. Instant-NGP [6] is trained for 500,000 iterations with a batch size of 4096. Grid-NeRF [12] is trained for 500,000 iterations with a batch size of 16384. Mega-NeRF [11] is trained for 500,000 iterations for each block, using a batch size of 1024. In Sec. 5.2, For MobileNeRF [1], We initialize a  $192^3$  grid to generate polygonal meshes while adhering to default parameters for other setups. We use the open-source version<sup>4</sup> for BakedSDF [13], conducting training in two phases: 20,000 and 50,000 iterations, respectively with a batch size of 16,384 and use  $1024^3$  to extract meshes by marching cubes [5]. We employ the official implementation of MERF [7] and 3D Gaussian Splatting [3]. During the MERF [7] experiments, the data is divided into four parts for a fair comparison with our method, and each block is trained with 32,768 batch size for 30,000 iterations, using default parameter settings. For 3D Gaussian splatting [3], we demonstrate the results after training for 30,000 and 100,000 iterations using the default settings.

Residence Sci-art Building PSNR↑ SSIM↑ LPIPS↓ PSNR↑ SSIM↑ LPIPS↓ PSNR↑ SSIM↑ LPIPS↓ NeRFacto 0.3120.50221.920.6420.47524.650.78117.700.44222.0820.93Mega-NeRF 0.6280.48925.600.7700.3900.5470.50422.540.680 0.386 25.710.8020.24120.130.397 Ours 0.578

Table 1: Comparison on Residence, Building and Sci-art scenes.

# 3 More Results

We present more qualitative comparisons among our method, MERF [7] (4 blocks), and 3D Gaussian Splatting [3] (100k iters) on the *Campus* Dataset

<sup>&</sup>lt;sup>4</sup> https://github.com/hugoycj/torch-bakedsdf



Fig. 2: Qualitative comparison of real-time rendering methods.

in Fig. 2. Additionally, we showcase rendering results from different LODs to demonstrate the appearance consistency across various LODs in Fig. 3. We conduct more experiments on the *Residence*, *Building*, and *Sci-art* scenes datasets. The hyperparameters used in these additional experiments are the same as described in the main paper. The results of these extended experiments are presented in Tab. 1. Our method outperforms Mega-NeRF and NeRFacto in all the scenes.

Ablations on shared heads. In contrast to Grid-NeRF [12], we don't share heads. This strategy reduces the model's representational capacity, consequently decreasing reconstruction quality as shown in the first two rows of Tab. 2. Additionally, we tested the speed of generating LOD under the shared heads condition in Tab. 2. It can be observed that sharing heads does not exert an influence on the quality of lod generation.

# 4 Real-Time Viewer

Our real-time viewer platform is based on the MERF volume renderer [7], leveraging an OpenGL fragment shader to execute ray marching and deferred rendering. It accesses feature and density information through texture look-ups. K. Song et al.

-	Table 2	2: Ab	olation	$\mathbf{study}$	on	sharing	heads.	
---	---------	-------	---------	------------------	----	---------	--------	--

	$\mathrm{PSNR}\uparrow$	$\mathrm{SSIM}\uparrow$	$\mathrm{LPIPS}{\downarrow}$	Training Time $\downarrow$
training with shared heads (4 blocks) training with non-shared heads (4 blocks)	$\begin{array}{c} 24.36\\ 24.82 \end{array}$	$0.734 \\ 0.741$	$0.201 \\ 0.190$	-
generate lod with shared heads (2 hours) generate lod with shared heads (4 hours) generate lod with non-shared heads	24.19 24.22 24.20	$0.718 \\ 0.729 \\ 0.724$	$0.213 \\ 0.204 \\ 0.204$	2 hours 4 hours 4 hours

However, we have implemented several improvements tailored for large scenes to enhance performance.

**Dynamic Loading.** We employ dynamic loading strategy to determine the level-of-detail and blocks to be loaded, thereby reducing VRAM usage. We use the center points obtained in Section 4.1 for each block on the xy plane as the xy-components, and assign the height of the highest part within the block as the z-component. This defines the 3D center point for each block during the rendering stage. Additionally, we use the rectangle formed by the block's four xy corner points and the previously determined z value to define the block's region. By projecting this region onto the camera plane, we can determine whether the block is visible to the camera.

During the rendering phase, we first eliminate blocks that are not visible for cameras. Starting with the coarsest LOD, we check if the distance from the camera to all visible blocks within the finer LOD exceeds a certain threshold. For blocks beyond this threshold, rendering is performed using the resources loaded at the current LOD. Otherwise, the evaluation progresses to the subsequent, finer LOD. This stepwise process continues until we reach the finest LOD or complete the rendering of the entire image.

**Depth sorting.** We compute the distance from each block's center point to the camera using only their xy-components, leveraging this calculated distance for depth sorting. We adopt this approach as our scene segmentation occurs solely in the xy plane, without division along the z-axis.

### 5 Discussion

UE4-NeRF [2] and NeuRas [4], both mesh-based rendering approaches, offer fast rendering speed but face challenges when representing large, detail-rich scenes like those including dense foliage. These mesh representations can consume extensive memory, often amounting to several gigabytes. Additionally, 3D Gaussian Splatting [3], typically requires millions to tens of millions of points. The inherent properties of Gaussian splatting, which involve numerous Gauss-related attributes, further increase the VRAM needed for rendering. The substantial VRAM consumption characteristic of these methods poses significant challenges for implementing rendering on web platforms and resource-constrained devices, as they struggle to accommodate the high memory demands.

8



Fig. 3: Qualitative comparison of different levels-of-detail.

# References

- Chen, Z., Funkhouser, T., Hedman, P., Tagliasacchi, A.: Mobilenerf: Exploiting the polygon rasterization pipeline for efficient neural field rendering on mobile architectures. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16569–16578 (2023)
- Gu, J., Jiang, M., Li, H., Lu, X., Zhu, G., Shah, S.A.A., Zhang, L., Bennamoun, M.: Ue4-nerf: Neural radiance field for real-time rendering of large-scale scene. Advances in Neural Information Processing Systems 36 (2024)
- Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics 42(4) (July 2023), https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/
- 4. Liu, J.Y., Chen, Y., Yang, Z., Wang, J., Manivasagam, S., Urtasun, R.: Neural scene rasterization for large scene rendering in real time. In: The IEEE International Conference on Computer Vision (ICCV) (2023)
- Lorensen, W.E., Cline, H.E.: Marching cubes: A high resolution 3d surface construction algorithm. In: Seminal graphics: pioneering efforts that shaped the field, pp. 347–353 (1998)
- Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. ACM Transactions on Graphics (ToG) 41(4), 1– 15 (2022)
- Reiser, C., Szeliski, R., Verbin, D., Srinivasan, P., Mildenhall, B., Geiger, A., Barron, J., Hedman, P.: Merf: Memory-efficient radiance fields for real-time view synthesis in unbounded scenes. ACM Transactions on Graphics (TOG) 42(4), 1–12 (2023)
- Schönberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
- Tancik, M., Casser, V., Yan, X., Pradhan, S., Mildenhall, B., Srinivasan, P.P., Barron, J.T., Kretzschmar, H.: Block-nerf: Scalable large scene neural view synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8248–8258 (2022)
- Tancik, M., Weber, E., Ng, E., Li, R., Yi, B., Wang, T., Kristoffersen, A., Austin, J., Salahi, K., Ahuja, A., et al.: Nerfstudio: A modular framework for neural radiance field development. In: ACM SIGGRAPH 2023 Conference Proceedings. pp. 1–12 (2023)

- 10 K. Song et al.
- Turki, H., Ramanan, D., Satyanarayanan, M.: 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. pp. 12922–12931 (2022)
- Xu, L., Xiangli, Y., Peng, S., Pan, X., Zhao, N., Theobalt, C., Dai, B., Lin, D.: Grid-guided neural radiance fields for large urban scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8296– 8306 (2023)
- Yariv, L., Hedman, P., Reiser, C., Verbin, D., Srinivasan, P.P., Szeliski, R., Barron, J.T., Mildenhall, B.: Bakedsdf: Meshing neural sdfs for real-time view synthesis. In: ACM SIGGRAPH 2023 Conference Proceedings, SIGGRAPH 2023, Los Angeles, CA, USA, August 6-10, 2023. pp. 46:1–46:9. ACM (2023)