

# Patch-based 3D Natural Scene Generation from a Single Example

Weiyu Li<sup>1\*</sup> Xuelin Chen<sup>2\*†</sup> Jue Wang<sup>2</sup> Baoquan Chen<sup>3</sup>  
<sup>1</sup> Shandong University    <sup>2</sup> Tencent AI Lab    <sup>3</sup> Peking University

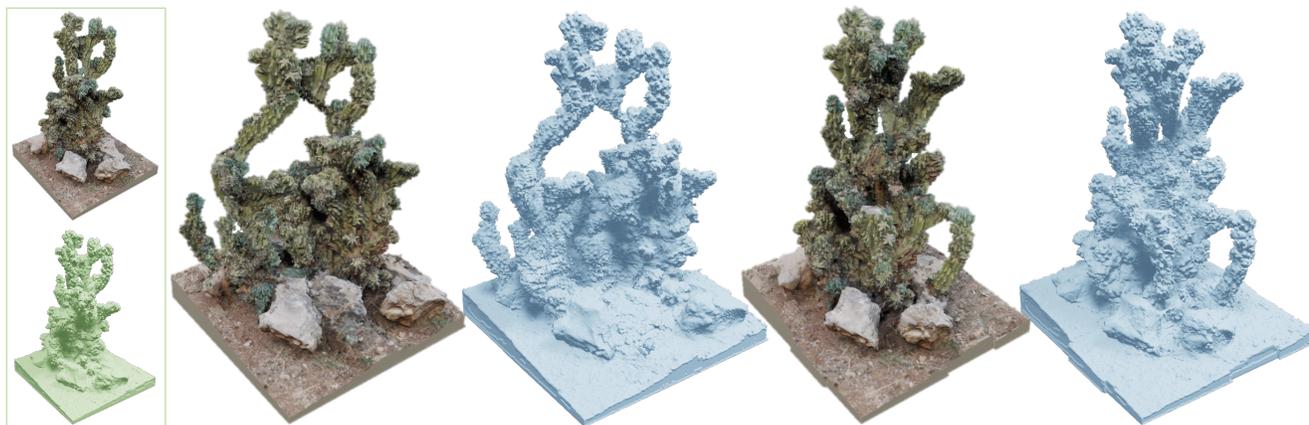


Figure 1. From a single example (inset), our method can generate diverse samples with highly realistic geometric structure and visual appearance. None of existing 3D generative models can achieve these results. More generated samples that are intricate and intriguing, ranging from single-object scenes (e.g., plants) to mid-scale ones, and large terrains, can be found in Section 4 and the supplementary.

## Abstract

We target a 3D generative model for general natural scenes that are typically unique and intricate. Lacking the necessary volumes of training data, along with the difficulties of having ad hoc designs in presence of varying scene characteristics, renders existing setups intractable. Inspired by classical patch-based image models, we advocate for synthesizing 3D scenes at the patch level, given a single example. At the core of this work lies important algorithmic designs w.r.t the scene representation and generative patch nearest-neighbor module, that address unique challenges arising from lifting classical 2D patch-based framework to 3D generation. These design choices, on a collective level, contribute to a robust, effective, and efficient model that can generate high-quality general natural scenes with both realistic geometric structure and visual appearance, in large quantities and varieties, as demonstrated upon a variety of exemplar scenes. Data and code can be found at <http://wyysf-98.github.io/Sin3DGen>.

## 1. Introduction

3D scene generation generally carries the generation of both realistic geometric structure and visual appearance. A wide assortment of scenes on earth, or digital ones across the internet, exhibiting artistic characteristics and ample variations over geometry and appearance, can be easily listed. Being able to populate these intriguing scenes in the virtual universe has been a long pursuit in the community.

Research has taken several routes, among which a prevalent one is learning to extract common patterns of the geometry *or* appearance from homogeneous scene samples, such as indoor scenes [14, 25, 34, 37, 59, 63, 71, 72, 75], terrains [15, 19, 21, 26], urban scenes [12, 30, 44], etc. Another line learns to generate single objects [6, 7, 16, 25, 33, 35, 45, 74]. A dominant trend in recent has emerged that learns 3D generative models to jointly synthesize 3D structures and appearances via differentiable rendering [4, 5, 8, 18, 43, 50]. Nevertheless, all these learning setups are limited in their ability to generalize in terms of varied scene types. While a more promising direction is the exemplar-based one, where one or a few exemplars featuring the scene of interest are provided, algorithm designs tailored for certain scene types in existing methods [38–40, 73] again draw clear boundaries

\*Joint first authors

†Corresponding author

‡Work done during an internship at Tencent AI Lab

of scene characteristics they can handle.

This work seeks to generate *general natural* scenes, wherein the geometry and appearance of constituents are often tightly entangled and contribute jointly to unique features. This uniqueness hinders one from collecting sufficient homogeneous samples for learning common features, directing us to the exemplar-based paradigm. On the other hand, varying characteristics across different exemplar scenes restrain us from having ad hoc designs for a certain scene type (e.g., terrains). Hence, we resort to classical patch-based algorithms, which date long before the deep learning era and prevail in several image generation tasks even today [13, 17, 20]. Specifically, given an input 3D scene, we synthesize novel scenes at the patch level and particularly adopt the multi-scale generative patch-based framework introduced in [17], where the core is a *Generative Patch Nearest-Neighbor* module that maximizes the bidirectional visual summary [54] between the input and output. Nevertheless, key design questions yet remain in the 3D generation: What representation to work with? And how to synthesize *effectively* and *efficiently*?

In this work, we exploit a grid-based radiance field – Plenoxels [69], which boasts great visual effects, for representing the input scene. While its simplicity and regular structure benefit patch-based algorithms, important designs must be adopted. Specifically, we construct the exemplar pyramid via coarse-to-fine training Plenoxels on images of the input scene, instead of trivially downsampling a pretrained high-resolution one. Furthermore, we transform the high-dimensional, unbounded, and noisy features of the Plenoxels-based exemplar at each scale into more well-defined and compact geometric and appearance features, improving the robustness and efficiency in the subsequent patch matching.

On the other end, we employ heterogeneous representations for the synthesis inside the generative nearest neighbor module. Specifically, the patch matching and blending operate in tandem at each scale to gradually synthesize an intermediate *value*-based scene, which will be eventually converted to a *coordinate*-based counterpart at the end. The benefits are several-fold: a) the transition between consecutive generation scales, where the value range of exemplar features may fluctuate, is more stable; b) the transformed features in the synthesized output is inverted to the original “renderable” Plenoxels features; so that c) the visual realism in the Plenoxels-based exemplar is preserved intactly. Last, working on voxels with patch-based algorithms necessarily leads to high computation issues. So we use an *exact-to-approximate* patch nearest-neighbor module in the pyramid, which keeps the search space under a manageable range while introducing negligible compromise on the visual summary optimality. These designs, on a collective level, essentially lay a solid foundation for an effective and

efficient 3D generative model.

To our knowledge, our method is the *first* 3D generative model that can generate 3D general natural scenes from a *single* example, with *both* realistic geometry and visual appearance, in large quantities and varieties. We validate the efficacy of our method on random scene generation with an array of exemplars featuring a variety of general natural scenes, and show the superiority by comparing to baseline methods. The importance of each design choice is also validated. Extensive experiments also demonstrates the versatility of our method in several 3D modeling applications.

## 2. Related Work

**3D Generative Models.** The goal of 3D generative models is to synthesize 3D contents with realistic geometric structures and visual appearances. While procedural models are capable of mass-producing particular 3D models, they take expertise and time to obtain rules and elementary assets. Hence, automating this process has been an active area of research, resulting in a vast body of work.

A prevalent route is the learning-based one, assuming having access to sufficient homogeneous samples for training. Some learn to generate realistic 3D geometric structures, such as indoor scenes [14, 34, 37, 59, 63, 71, 72], terrains [15, 19], urban scenes [12, 44], etc. Others focus on the visual appearance, attempting to automatically texturize or assign materials for geometric scaffolds [21, 25, 26, 30, 75]. Another line has been directed at generating single objects with realistic structures or/and textures [6, 7, 16, 33, 35, 45, 74], showing the potential in enriching the elementary asset library. A dominant trend in recent has also emerged [4, 5, 8, 18, 43, 50], where deep generative models are trained on large volumes of images collected from scenes of a specific category, to allow joint synthesis of realistic 3D structure and appearance with neural radiance fields. Nevertheless, all these learning setups require large volumes of training data, and are limited in their ability to generalize, especially in terms of varied scene types.

A more relevant direction is the exemplar-based one, where one or a few exemplars featuring the scene of interest are provided. However, existing methods with algorithm designs tailored for certain scene types again draw clear boundaries of scene characteristics they can handle. [73] extract height field patches from exemplars to synthesize terrains, but the synthesis is guided with particular emphasis on dominant visual elements in terrains. [38–40] use structured units specified in the input exemplar to facilitate architecture model synthesis. Extending texture image synthesis, [32] synthesizes signed distance fields from an input geometry, the method can not generalize to complex general natural scenes, and the result is inadequate for displaying due to the lack of appearance properties.

In this paper, we aim for 3D general natural scenes, with

an emphasis on generating both realistic geometry and appearance. Lacking the necessary volume of data characterizing the target scene, along with the difficulties of having ad hoc designs in presence of varying scene characteristics, we advocate for synthesizing novel 3D scenes at the patch level, given a single exemplar scene.

**Generative Image Models.** Generative image models have made great strides in the past years. State-of-the-art methods can now learn generative models from large volumes of homogeneous image samples, achieving unprecedented success in producing realistic images [9, 28, 29, 31, 57, 58]. On the other end, there has also been a surge of developments to learn a generative model from a single training image [24, 51, 53, 66]. But, these learning-based single image models typically require a long training time. Differing from these learning-based paradigms, a classical patch-based approach, that dates back long before the deep learning era, is revived in [11, 17, 20], showing amazing performance. The core of these models is to maximize the bidirectional patch similarity between the input and synthesized output in a coarse-to-fine manner, and have demonstrated their capability to generate diverse outputs from a single exemplar image, with orders of magnitude faster than learning-based ones. Our work is particularly inspired by this line of work but must address challenges arising from lifting the multi-scale generative patch-based framework to *effective* and *efficient* 3D scene generation.

**3D Scene Representations.** While it is common to represent an image as a distributed amplitude of colors over a 2D grid, more often than not, the 3D representation varies. Polygon meshes and points offer a compact representation, with precedents in patch-based synthesis [22, 52], but the irregularity makes them intractable for high-quality 3D generation. The same holds for point clouds. Recently, the community has indeed witnessed a revolution started by an emerging representation, i.e. neural radiance field [41], which approximates the 5D plenoptic function [3] of the underlying scene with neural networks and shows unprecedentedly photo-realistic visual results. An explosion of techniques occurred since then that improves the representation in various aspects [23, 42, 47, 48, 55, 67, 68, 70]. We refer readers to [56, 64] for more in-depth summaries. Among these variants, we opt for a simple yet expressive voxel-based representation – Plenoxels [69], which has shown great competence on novel view synthesis. Its simplicity and regular structure benefit patch-based algorithms, however, important designs must be taken to fit it into our framework for high-quality generation of general natural scenes.

**Concurrent Work.** Concurrent works [27, 60] propose to learn a 3D generative model from images of an input scene, producing variations that can be rendered with realistic imagery. [62] focus on generating diverse geometric struc-

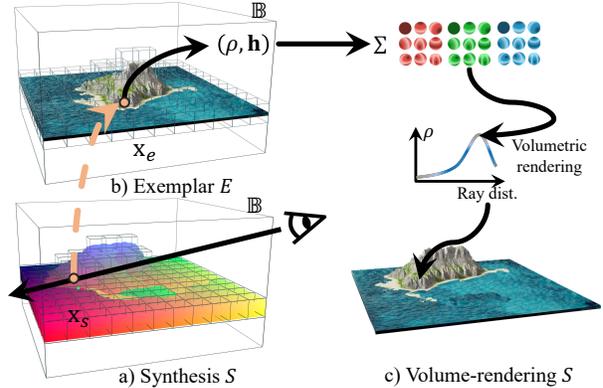


Figure 2. a) The synthesized scene  $S$  is represented as a field mapping a coordinate in  $S$  to one in  $E$ . b) The Plenoxels-based exemplar  $E$  uses a sparse grid, where each occupied voxel stores a scalar opacity  $\rho$  and spherical harmonic coefficients  $\mathbf{h}$ . c) Appealing imagery of  $S$  can be produced via the volumetric rendering function. Empty voxels are omitted for simplicity.

tures from an input shape. Their core idea is to extend 2D SinGAN [51] for learning the internal distribution in the 3D exemplar, differing significantly from our technical route. While these methods require a long training time (typically days), our method can generate high-quality samples in minutes, without offline training. Last, [46] can generate arbitrary 3D models represented by NeRF, with pretrained powerful image diffusion models as priors.

### 3. Method

The input 3D scene to our method can be a real-world or digital scene, as we first train Plenoxels on the images of the input scene to obtain a Plenoxels-parameterized exemplar. Then, our method synthesizes novel variations at the patch level, with a multi-scale generative patch-based framework. In the following, we describe important designs w.r.t. the scene representation (Section 3.1 & 3.2) and the generative patch nearest-neighbor field module (Section 3.3), that, integrated into the multi-scale patch-based framework (Section 3.2), contribute collectively to our success.

#### 3.1. Scene Representations

**Exemplar Scene Representation.** We assume the exemplar scene  $E$  lies within an axis-aligned box  $\mathbb{B}$  centered at the origin, around which we can distribute cameras to capture images for training Plenoxels. As per Plenoxels,  $E$  is represented by a sparse voxel grid, where each occupied voxel center stores features including a scalar opacity  $\rho$  and a vector of spherical harmonic (SH) coefficients  $\mathbf{h}$  for each color channel:  $E : \mathbf{x} \rightarrow (\rho, \mathbf{h})$ , where  $\mathbf{x}$  indicates a voxel center within  $\mathbb{B}$ . These features can be further trilinearly interpolated to model the full plenoptic function continuously

in space. Notably, the appearance feature uses 2-degree harmonics, which requires 9 coefficients per color channel for a total of 27 harmonic coefficients per voxel.

**Exemplar Transformation.** While Plenoxels features can be used to render pleasing imagery, naively using them for the patch distance is unsuitable. Density values are not well-bounded, contain outliers, and can not accurately describe the geometric structure within a patch. On the other hand, high-dimensional SH coefficients are excessively consumptive for patch-based frameworks. Hence, we transform the exemplar features for the input to the generative patch nearest neighbor module. First, the density field is converted to a signed distance field (SDF). Specifically, the signed distance at each voxel is computed against the surface mesh extracted from the density field by Marching Cubes [36]. Note that Plenoxels prunes unnecessary voxels during training, which creates holes and irregular structures in invisible regions. So we flood-fill these regions with high-density values, prior to the mesh extraction. Last, we rescale and truncate the signed distance to ignore distance values far away from the surface. Formally, the geometry transformation is as follow:  $G(\mathbf{x}) = \max(-1, \min(1, SDF(\mathbf{x})/t))$ , where the truncated scale  $t$  is set to 3 times of the voxel size at each generation scale. Moreover, we normalize SH coefficient vectors and use the principal component analysis (PCA) to reduce the dimensionality (from 27 to 3 by default), significantly reducing the computation overhead. Finally, the transformed exemplar  $\hat{E}$  is now given as:

$$\hat{E} : \mathbf{x} \rightarrow (G(\mathbf{x}), P(\mathbf{h})), \quad (1)$$

where  $G(\cdot)$  denotes transforming of the geometric feature, and  $P(\cdot)$  transforming the appearance feature.

**Synthesized Scene Representation.** In the multi-scale generation, the output scene  $S$  at each scale is represented by a coordinate-based mapping field, instead of a value-based one that stores features. Specifically,  $S$  is represented as a field that maps a 3D voxel center in the synthesis grid to one in the exemplar  $E$ ,  $S : \mathbf{x}_s \rightarrow \mathbf{x}_e$ , with which the original Plenoxels features  $E(S(\mathbf{x}_s))$  can be queried for  $S$ .

Note, in addition to discrete grid samplings, dense samplings  $\mathbf{x}_s$  in  $S$  can also be mapped to the continuous exemplar space, by simply considering the local offset  $\delta$  to the nearest voxel center, i.e.,  $S(\mathbf{x}) = S(\mathbf{N}(\mathbf{x}) + \delta)$ , where  $\mathbf{N}(\cdot)$  returns the nearest voxel center of  $\mathbf{x}$ . This is particularly useful, as it enables upscaling  $S$  to finer grids in the multi-scale framework, and sufficient sampling for rendering the final generation result with high-quality imagery.

**Viewing Synthesized Results.** The synthesized scene can be projected onto 2D through the volume rendering equation as in NeRF [41], yielding highly photo-realistic imagery under varying views. We refer readers to [69] for

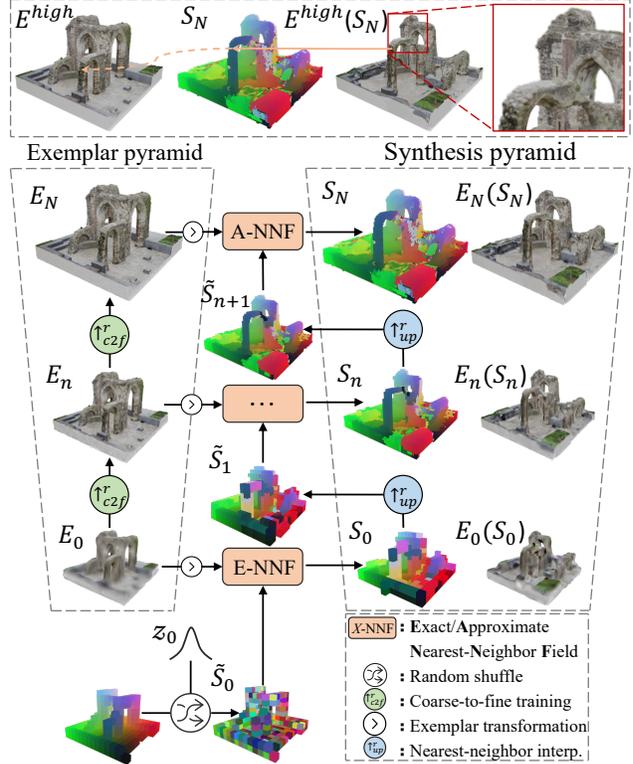


Figure 3. Multi-scale generation. At each scale  $n$ , the NNF module updates the generation based on the transformed exemplar  $\hat{E}_n$  and an initial guess  $\tilde{S}_n$  upscaled from the previous. The coarsest scale takes a shuffled identity mapping field as input. Note that our coordinate-based representation  $S_N$  can map to patches in a higher-resolution exemplar for higher quality (top row).

more details. Figure 2 illustrates how a synthesized result, paired with the exemplar, can display appealing imagery.

### 3.2. Multi-scale Generation

We use the same multi-scale framework as in previous works [10, 17, 51], which generally employs a coarse-to-fine process, so we have the opportunity to synthesize a more detailed scene based on an initial guess upscaled from the previous scale. In this pyramidal pipeline, different information is captured and reproduced at varying scales, spanning from global layouts at coarser scales to fine geometric and appearance details at finer scales (See Figure 3).

**Exemplar Pyramid Construction.** Given the input scene, we build a pyramid  $(E_0, \dots, E_N)$ , where  $E_{n-1}$  is a downsampled version of  $E_n$  by a factor  $r^{-1}$  ( $r = 4/3$ ). By default, we use  $N = 7$  (8 scales in total) for balancing quality and efficiency. Specific resolutions in the pyramid are listed in the supplementary. When working with an exemplar pyramid obtained by recursively downsampling a pretrained high-resolution exemplar, we observed lots of

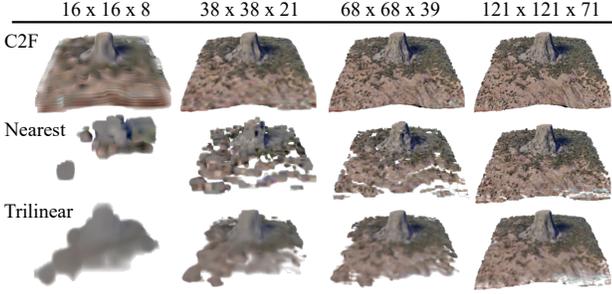


Figure 4. Exemplar pyramid. Coarse-to-fine training (top) shows more consistency between consecutive exemplars, whereas common downsampling algorithms (mid and bottom) result in missing geometry (e.g., the ground) and blurry appearance.

artifacts due to missing features at coarser exemplars, and severe feature inconsistency between exemplars at consecutive scales. Hence, we build the exemplar pyramid by coarse-to-fine training Plenoxels, at increasing resolutions synchronized with the multi-scale framework. Such exemplar pyramid prevents losing thin structures at coarser scales, and offers rather smoother transition and consistent features between consecutive exemplars, leading to stable transition in the multi-scale generation (See Figure 4).

**Coarse-to-fine Generation.** At each scale  $n$ , an initial guess  $\tilde{S}_n$  is produced by upsampling the output in the previous scale:  $\tilde{S}_n = S_{n-1} \uparrow^r$ , with the same factor  $r$  to match with the exemplar. Then, the mapping field in  $\tilde{S}_n$  is updated by the generative nearest neighbor module with matched coordinates in the exemplar. The patch size at all scales is  $p^3$  ( $p = 5$  by default), which captures around 1/3 of the content in the coarsest exemplar. Unlike adding noise to raw exemplar values as in the image synthesis, our initial guess  $\tilde{S}_0$  at the coarsest scale is an identity mapping field shuffled with Gaussian noise  $z_0 = \mathcal{N}(0, \sigma^2)$ ,  $\sigma = 0.5$  by default, scaled by the extents of the bounding box  $\mathbb{B}$ , which is natural for our coordinate-based synthesis.

### 3.3. Generative Patch Nearest-Neighbor Field

Usually, two stages, namely the patch *matching* and *blending*, operate in tandem in the nearest neighbor field (NNF) search of patch-based algorithms. Specifically, the matching finds the most suitable patch from the exemplar for each in  $S$ , and then the latter blends of the *values* of overlapping patches together. This is vital to a robust EM-like optimization in patch-based image synthesis [1, 2], leading to converging synthesis results in several iterations.

**Value-/Coordinate-based Synthesis.** In this work, we use heterogeneous representations for the synthesis in NNF. Specifically, at each scale  $n$ , the patch matching and blend-

ing first operate in tandem for  $T - 1$  iterations, to gradually synthesize an intermediate *value*-based scene with averaged values over overlapping patches. Then, when the synthesis is stable at the last iteration, the final output of NNF uses coordinate-based representation, which stores only the center location of the nearest patch in  $\hat{E}_n$ . As aforementioned, this design offers stable transition between consecutive generation scales, where the value range of exemplar features may fluctuate, and, importantly, helps us trace back to the original Plenoxels features that can be rendered into photo-realistic imagery, via simply mapping to the original exemplar, even to a higher-resolution version for the final generated scene (See top of Figure 3). Specifically, each iteration in NNF at each scale proceeds as follows:

(1) *Extract Patches:* Patches in  $\hat{E}_n(\tilde{S}_n)$  are extracted to form a query patch set  $Q$ , and ones in  $\hat{E}_n$  form a key set  $K$ .

(2) *Match Nearest Neighbors:* We first compute distance between each query patch  $Q_i$  and each key patch  $K_j$  as the weighted sum of the appearance and geometric features using  $L2$  distance:

$$D_{i,j} = w_a \|Q_{i,j}^a - K_{i,j}^a\|^2 + (1 - w_a) \|Q_{i,j}^g - K_{i,j}^g\|^2, \quad (2)$$

where  $w_a$  (0.5 by default) is the trade-off parameter. To control the visual completeness in the synthesis by the bidirectional similarity [54], the final patch similarity scores normalize the distance with a per-key factor:

$$C_{i,j} = \frac{D_{i,j}}{(\alpha + \min_l(D_{l,j}))}, \quad (3)$$

where  $\alpha$  (0.01 by default) controls the degree of completeness, and smaller  $\alpha$  encourages completeness.

(3) *Update  $\tilde{S}_n$ :* For each query patch  $Q_i$  in  $\hat{E}_n(\tilde{S}_n)$ , we find its nearest patch in  $K_l$ , then update  $\tilde{S}_n$  with averaged values over overlapping patches for the first  $T - 1$  iterations, and with the nearest patch center for the last iteration.

**Exact-to-Approximate NNF.** Although the computation above can be in parallel performed on GPUs, brutally enumerating all pairs of patches would apparently lead to surprisingly huge distance matrices as the resolution increases, preventing us from obtaining high-resolution synthesis even with modern powerful GPUs. Hence, to avoid searching in tremendous space, we propose to perform the NNF in an *exact-to-approximate* manner. Specifically, at first 5 coarser scales, *exact nearest-neighbor field* (E-NNF) search is performed with  $T_e = 10$  times to stabilize global layout synthesis when the memory consumption is low. At rest 3 finer scales, an *approximate nearest-neighbor field* (A-NNF) search – PatchMatch [1] – with jump flood [49] is used for  $T_a = 2$  times to reduce memory footprint from  $O(M^2)$  to  $O(M)$  ( $M$  is the number patches), which is equivalent to only considering visual coherence.



Figure 5. Random generation. Our method generalizes to all these scenes with highly varied structures and appearances, producing highly diverse and realistic scenes. The supplementary presents exemplars and more artistic imagery rendered with these 3D scenes.

## 4. Experiments

We collected a rich variety of 3D scene models to examine the performance of our method on random scene generation, ranging from rocks to plants, sculptures, landscapes, terrains, artistic scenes, etc. Some are digitalized *real-world* scenes, e.g., the *Devil's Tower*. These scenes possess varying degrees of complexity in terms of geometry and appearance. In the following, we present experiments conducted to evaluate various aspects of the proposed solution. Unless specified, we use the default parameters described above, 512 for the resolution along the max dimension of the  $E^{high}$ , and  $512 \times 512$  image resolution for rendering. Full visualization of all exemplars, more technical details and experimental results can be found in the supplementary.

**Random Generation.** Figure 5 presents results obtained by our method on exemplar-based random scene generation. These results show our method can generalize to scenes of highly varied features, yielding high-quality and diverse scenes similar to the exemplar. A particular feature of our method is the photo-realism and view-dependent effects of the exemplar are inherited in the results, as evidenced by Figure 5 and 7. Each sample is generated in 1~3 minutes on a V100 GPU depending on the scenes, and viewing the results can be executed at an interactive rate (15 fps).

**Comparisons.** We particularly compare to GRAF and StyleNeRF, which are representative GAN-based 3D gen-

erative models. We cast them into exemplar-based models via training separately on images of each exemplar. In addition, we also compare to GPNN-3D, which trivially extends [17] for our task. We investigate the advantages of exemplar-based scene generation using our method against these alternatives, on various exemplars listed in Table 1. Figure 8 presents part of their visual results. Generally, GAN-base baselines suffer from notorious mode collapse, producing almost identical results due to lacking diverse training scenes. The visuals also tend to be more blurry and noisy, compared to our sharp imagery. GPNN-3D can not synthesize high-resolution results due to computational efficiency issues, and quickly fails at coarse scales, producing meaningless content. For quantitative comparisons, we produce 50 generated scenes from each exemplar with each method, render multi-view images and extract 3D surface points of the exemplar and of each generated scene, and then rate the *Visual Quality* (V-Qua.) *Visual Diversity* (V-Div.), *Geometry Quality* (G-Qua.), and *Geometry Diversity* (G-Div.) using common metrics employed in both 2D [51] and 3D [61] generation. The supplementary contains more details. Table 1 presents quantitative results, where, by rating with the combination of these established metrics, ours outperforms baselines by large margins, suggesting high quality and diversity from both 2D and 3D perspective.

**Ablation.** We compare to several variants derived from our full method: 1) *Ours (w/o TSDF)* uses an occupancy

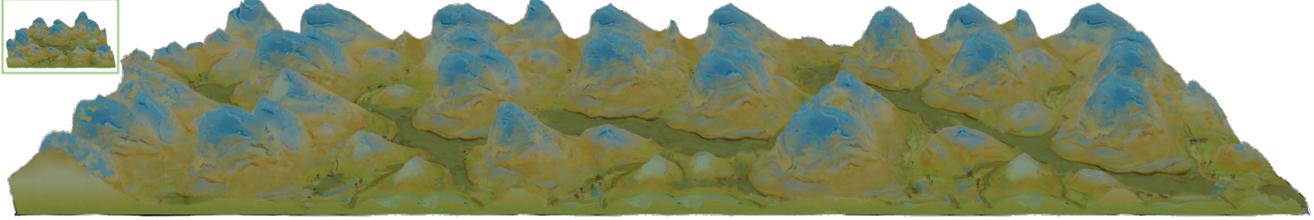


Figure 6. A novel "A Thousand Li of Rivers and Mountains" [65] is rendered from a generated 3D sample, that is of a different size, resolution and aspect ratio to the *Vast Land* exemplar (inset). Specification:  $E_N$  -  $288 \times 288 \times 112$ ,  $E^{high}$  -  $512 \times 512 \times 200$ ,  $S_N$  -  $747 \times 288 \times 112$ ,  $E^{high}(S_N)$  -  $1328 \times 512 \times 200$ , final rendering resolution -  $4096 \times 1024$ .

Table 1. Quantitative comparisons. Ours outperforms baselines by large margins, with high quality and diversity scores in terms of both visual and geometric content. We highlight top two in bold and underline the top one. Note GPNN-3D’s high diversity scores can be explained by noisy contents shown in the visual results.

	GRAF				StyleNeRF				GPNN-3D				Ours			
	V-Qua.↓	V-Div.↑	G-Qua.↓	G-Div.↑	V-Qua.↓	V-Div.↑	G-Qua.↓	G-Div.↑	V-Qua.↓	V-Div.↑	G-Qua.↓	G-Div.↑	V-Qua.↓	V-Div.↑	G-Qua.↓	G-Div.↑
St Alphege	<b>0.078</b>	<b>0.046</b>	<b>0.473</b>	0.040	0.206	0.032	0.769	0.012	3.929	0.041	134.997	0.059	<b>0.022</b>	<b>0.312</b>	<b>0.612</b>	<b>0.473</b>
Devil’s Tower	<b>0.233</b>	0.084	<b>0.480</b>	0.021	0.470	0.021	1.000	0.011	2.495	<b>0.694</b>	6.545	<b>1.974</b>	<b>0.032</b>	<b>0.203</b>	<b>0.304</b>	<b>0.207</b>
Desert Lowpoly	<b>0.057</b>	0.048	0.721	0.043	0.255	0.026	0.813	0.018	1.312	<b>0.405</b>	<b>0.344</b>	<b>1.048</b>	<b>0.020</b>	<b>0.312</b>	<b>0.568</b>	<b>0.454</b>
Green Island	<b>0.294</b>	0.047	<b>0.277</b>	0.015	0.606	0.015	0.669	0.014	0.254	<b>1.136</b>	18.228	<b>17.673</b>	<b>0.044</b>	<b>0.172</b>	<b>0.097</b>	<b>0.081</b>
Stone Arch	0.101	0.055	0.603	0.029	<b>0.060</b>	0.011	<b>0.339</b>	0.005	1.608	<b>0.504</b>	53.943	<b>29.448</b>	<b>0.003</b>	<b>0.146</b>	<b>0.126</b>	<b>0.100</b>
Mountain	0.410	0.060	<b>0.498</b>	0.022	<b>0.222</b>	0.037	<b>0.757</b>	0.010	2.602	<b>0.787</b>	5.947	<b>1.674</b>	<b>0.105</b>	<b>0.391</b>	0.935	<b>0.467</b>
Vast Land	<b>0.072</b>	0.058	0.994	0.229	0.219	0.023	1.047	0.017	0.907	<b>0.348</b>	<b>0.690</b>	<b>1.840</b>	<b>0.014</b>	<b>0.124</b>	<b>0.177</b>	<b>0.105</b>

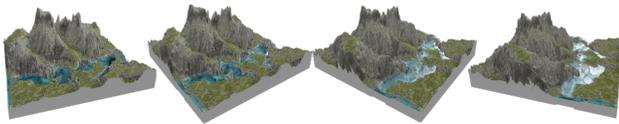


Figure 7. View-dependent effects in our synthesized results. See the reflection on the river changing under spinning cameras.

field, instead of TSDF, converted from the exemplar density field for geometric features; 2) *Ours (w/o c2f)* drops the deep coarse-to-fine exemplar training, and instead recursively trilinearly interpolates a high-resolution exemplar; 3) *Ours (value-only)* uses only value-based synthesis in NNF, and does not use TSDF and PCA as we can not trace back to original Plenxels features, and the maximum resolution is limited to 68; 4) *Ours (coord.-only)* uses only coordinate-based synthesis in NNF. Figure 9 and Table 2 present the qualitative and quantitative comparison results,

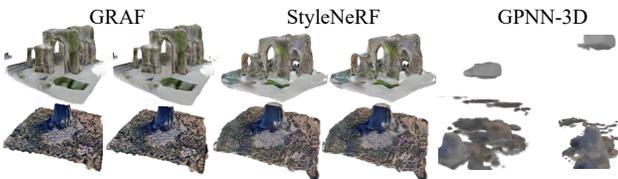


Figure 8. Visual comparisons. GAN-based baselines suffer from severe mode collapse, producing samples (two shown) almost identical to the input. GPNN-3D fails on the task.

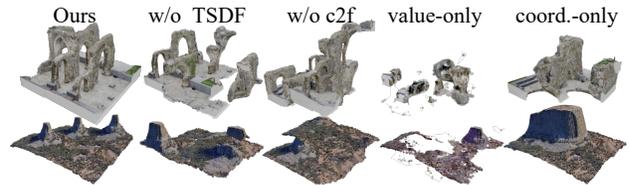


Figure 9. Ablation study. Ours (w/o TSDF) and (w/o c2f) can not well preserve the geometric structures. Ours (value-only) fails and produces with noisy content, while Ours (coord.-only) is unstable, easily leading to bulky structures or holes.

respectively, showing the importance of each algorithmic design.

**Higher-resolution Generation.** 1) In Figure 6, we show that our method supports generating a result scene of different size to the exemplar, and particularly of a much higher resolution and different aspect ratio. See specifications in the caption. 2) In addition, we also stress test with a very high-resolution setting, where  $E_N$  has 288 voxels along the max dimension, and our method can still synthesize a highly plausible sample in  $\sim 10$  minutes. We observed slightly improved visual quality over the default setting, as the default is sufficient for most complicated scenes. Results and details can be found in the supplementary.

**Applications.** In Figure 10, we demonstrate the versatility of our method in several 3D modeling applications with our unified generation framework (more details in the supplementary): 1) *Retargeting*: The goal is to resize a 3D scene

Table 2. Quantitative ablation results. While some variants produce higher diversity scores with meaningless noisy contents, ours consistently produce *diverse* results with *highest* quality scores.

		V-Qua.↓	V-Div.↑	G-Qua.↓	G-Div.↑
St Alpage	Ours	<b>0.022</b>	0.312	<b>0.612</b>	0.473
	w/o TSDF	0.114	<b>0.568</b>	1.176	<b>1.105</b>
	w/o c2f	<b>0.024</b>	<b>0.353</b>	<b>0.847</b>	0.639
	value-only	3.779	0.054	56.304	<b>3.823</b>
	coord.-only	0.044	0.336	1.003	0.719
Devil's Tower	Ours	<b>0.032</b>	0.203	<b>0.304</b>	0.207
	w/o TSDF	0.047	0.263	0.422	0.350
	w/o c2f	0.082	<b>0.547</b>	2.101	<b>3.500</b>
	value-only	1.795	<b>0.344</b>	14.122	<b>7.650</b>
	coord.-only	<b>0.041</b>	0.201	<b>0.256</b>	0.492
Desert Lowpoly	Ours	<b>0.020</b>	0.312	<b>0.568</b>	0.454
	w/o TSDF	<b>0.041</b>	<b>0.462</b>	<b>1.347</b>	1.007
	w/o c2f	0.049	0.457	1.745	1.097
	value-only	0.763	0.419	29.047	<b>6.674</b>
	coord.-only	0.100	<b>0.487</b>	2.754	<b>1.526</b>

to a target size (typically of a different aspect ratio), while maintaining the local patches in the exemplar. We simply change the size of the identity mapping field and use it as the initial guess  $\tilde{S}_0$  without shuffling. 2) *Editing*: Users can manipulate on a 3D proxy, which can be the underlining mapping field or mesh, for editing an exemplar or generated scene, such as removal, duplication, and modification. The manually manipulated proxy is then converted and fed as the initial guess at the coarsest scale for synthesizing the final scene. 3) *Structural analogies*: Given two scenes A and B, we create a scene with the patch distribution of A, but which is structurally aligned with B. This is realized by using the exemplar pyramid of A, and an identity mapping as the initial guess, but by replacing  $\hat{E}_0(\tilde{S}_0)$  with the transformed features in B, and vice versa. 4) *Re-decoration*: With the coordinate-based representation, we can re-decorate the generated ones with ease, via simply remapping to exemplars of different appearance.

## 5. Discussion, Limitations and Future Work

This work makes a first attempt towards a generic generative model for synthesizing highly realistic general natural scenes from only one exemplar. Building upon Plenoxels, our method can efficiently synthesize diverse and high-quality scenes. The generated samples particularly inherit photo-realism and view-dependent effects from the exemplar. Despite success demonstrated, we note a few shortcomings. We can not handle scenes eluding Plenoxels (e.g., transparent fluids, strong reflection), which is the actual input to our framework. Particularly, the Plenoxels-based representation is not suitable for large and unbounded scenes,

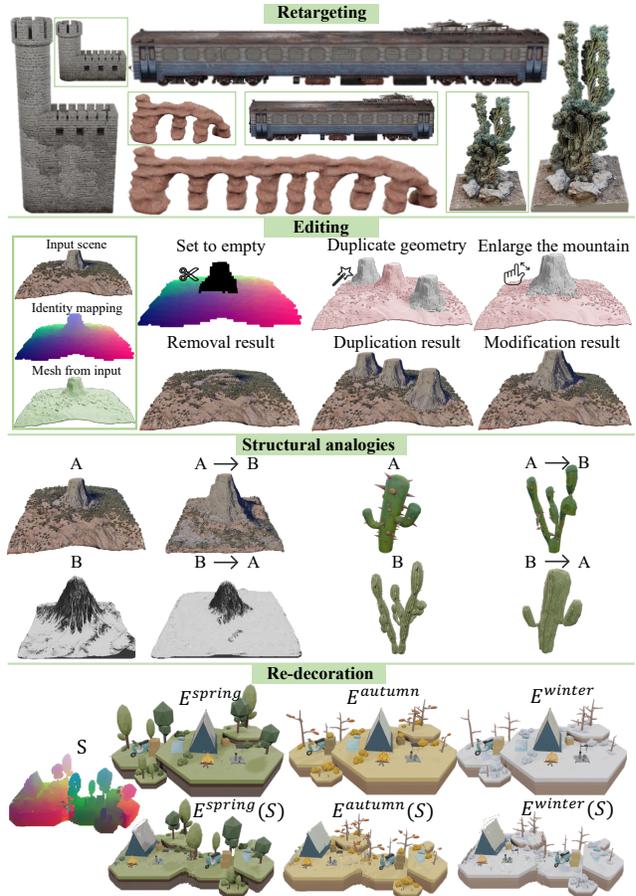


Figure 10. Applications. 1<sup>st</sup>: Retargeting 3D scenes (marked in boxes). 2<sup>nd</sup>: Editing a 3D scene (removal, duplication and modification). 3<sup>rd</sup>: Structural analogies. A  $\rightarrow$  B = Visual content of A + Structure of B, and vice versa. 4<sup>th</sup>: Re-decoration is realized by simply re-mapping to exemplars of different appearance.

leading to artifacts in the results (more discussion in supplementary). With voxelized volumetric representations, we can not perfectly synthesize scenes with tiny thin structures, and ones with highly semantic or structural information, e.g., human body and modern buildings. Moreover, in contrast to *continuous* distributions learned in neural-based methods, we work on *discrete* patch distributions and thus lack the capability of generating novel patches/pixels. A future direction is to learn a continuous distribution from a large number of homogeneous samples produced by our method, with GANs, VQ-VAEs, or diffusion models. Last, the view-dependent effects of the results are inherited from the input Plenoxels, although SH features have already *implicitly* considered the view-dependent lighting, consistent global illumination can not be guaranteed in our results, leading to another future direction.

**Acknowledgements.** This work was supported in part by National Key R&D Program of China 2022ZD0160801.

## References

- [1] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics (TOG)*, 28(3):24, 2009. 5
- [2] Connelly Barnes and Fang-Lue Zhang. A survey of the state-of-the-art in patch-based synthesis. *Computational Visual Media*, 3(1):3–20, 2017. 5
- [3] James R Bergen and Edward H Adelson. The plenoptic function and the elements of early vision. *Computational models of visual processing*, 1:8, 1991. 3
- [4] Eric Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2
- [5] Eric R Chan, Connor Z Lin, Matthew A Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, et al. Efficient geometry-aware 3d generative adversarial networks. *arXiv preprint arXiv:2112.07945*, 2021. 1, 2
- [6] Siddhartha Chaudhuri, Evangelos Kalogerakis, Leonidas Guibas, and Vladlen Koltun. Probabilistic reasoning for assembly-based 3d modeling. In *SIGGRAPH*, pages 1–10, 2011. 1, 2
- [7] Xuelin Chen, Daniel Cohen-Or, Baoquan Chen, and Niloy J Mitra. Towards a neural graphics pipeline for controllable image generation. In *Computer Graphics Forum (Eurographics)*, volume 40, pages 127–140. Wiley Online Library, 2021. 1, 2
- [8] Terrance DeVries, Miguel Angel Bautista, Nitish Srivastava, Graham W Taylor, and Joshua M Susskind. Unconstrained scene generation with locally conditioned radiance fields. In *International Conference on Computer Vision (ICCV)*, pages 14304–14313, 2021. 1, 2
- [9] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. 3
- [10] Alexei A Efros and William T Freeman. Image quilting for texture synthesis and transfer. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, pages 341–346, 2001. 4
- [11] Ariel Elnekave and Yair Weiss. Generating natural images with direct patch distributions matching. In *European Conference on Computer Vision (ECCV)*, 2022. 3
- [12] Tian Feng, Lap-Fai Yu, Sai-Kit Yeung, KangKang Yin, and Kun Zhou. Crowd-driven mid-scale layout design. *ACM Transactions on Graphics (TOG)*, 35(4):132–1, 2016. 1, 2
- [13] Noa Fish, Lilach Perry, Amit Bermano, and Daniel Cohen-Or. Sketchpatch: Sketch stylization via seamless patch-level synthesis. *ACM Transactions on Graphics (TOG)*, 39(6):1–14, 2020. 2
- [14] Qiang Fu, Xiaowu Chen, Xiaotian Wang, Sijia Wen, Bin Zhou, and Hongbo Fu. Adaptive synthesis of indoor scenes via activity-associated object relation graphs. *ACM Transactions on Graphics (TOG)*, 36(6):1–13, 2017. 1, 2
- [15] Eric Galin, Eric Guérin, Adrien Peytavie, Guillaume Cor-donnier, Marie-Paule Cani, Bedrich Benes, and James Gain. A review of digital terrain modeling. In *Computer Graphics Forum*, volume 38, pages 553–577. Wiley Online Library, 2019. 1, 2
- [16] Lin Gao, Tong Wu, Yu-Jie Yuan, Ming-Xian Lin, Yu-Kun Lai, and Hao Zhang. Tm-net: Deep generative networks for textured meshes. *ACM Transactions on Graphics (TOG)*, 40(6):1–15, 2021. 1, 2
- [17] Niv Granot, Ben Feinstein, Assaf Shocher, Shai Bagon, and Michal Irani. Drop the gan: In defense of patches nearest neighbors as single image generative models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 3, 4, 6
- [18] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. Stylenerf: A style-based 3d aware generator for high-resolution image synthesis. In *International Conference on Learning Representations (ICLR)*, 2022. 1, 2
- [19] Éric Guérin, Julie Digne, Éric Galin, Adrien Peytavie, Christian Wolf, Bedrich Benes, and Benoît Martinez. Interactive example-based terrain authoring with conditional generative adversarial networks. *ACM Transactions on Graphics (TOG)*, 36(6):228:1–228:13, Nov. 2017. 1, 2
- [20] Niv Haim, Ben Feinstein, Niv Granot, Assaf Shocher, Shai Bagon, Tali Dekel, and Michal Irani. Diverse generation from a single video made possible. *arXiv preprint arXiv:2109.08591*, 2021. 2, 3
- [21] Zekun Hao, Arun Mallya, Serge Belongie, and Ming-Yu Liu. Gancraft: Unsupervised 3d neural rendering of minecraft worlds. In *International Conference on Computer Vision (ICCV)*, pages 14072–14082, 2021. 1, 2
- [22] Gur Harary, Ayellet Tal, and Eitan Grinspun. Context-based coherent surface completion. *ACM Transactions on Graphics (TOG)*, 33(1):1–12, 2014. 3
- [23] Peter Hedman, Pratul P. Srinivasan, Ben Mildenhall, Jonathan T. Barron, and Paul Debevec. Baking neural radiance fields for real-time view synthesis. *International Conference on Computer Vision (ICCV)*, 2021. 3
- [24] Tobias Hinz, Matthew Fisher, Oliver Wang, and Stefan Wermter. Improved techniques for training single-image gans. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1300–1309, 2021. 3
- [25] Arjun Jain, Thorsten Thormählen, Tobias Ritschel, and Hans-Peter Seidel. Material memex: Automatic material suggestions for 3d objects. *ACM Transactions on Graphics (TOG)*, 31(6):1–8, 2012. 1, 2
- [26] Konrad Kapp, James Gain, Eric Guérin, Eric Galin, and Adrien Peytavie. Data-driven authoring of large-scale ecosystems. *ACM Transactions on Graphics (TOG)*, 39(6):1–14, 2020. 1, 2
- [27] Animesh Karnewar, Tobias Ritschel, Oliver Wang, and Niloy Mitra. 3inGAN: Learning a 3D generative model from images of a self-similar scene. In *Proc. 3D Vision (3DV)*, 2022. 3
- [28] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4401–4410, 2019. 3

- [29] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8110–8119, 2020. 3
- [30] Tom Kelly, Paul Guerrero, Anthony Steed, Peter Wonka, and Niloy J Mitra. Frankengan: guided detail synthesis for building mass-models using style-synchronized gans. *ACM Transactions on Graphics (TOG)*, 2018. 1, 2
- [31] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018. 3
- [32] Sung-Ho Lee, Taejung Park, Jong-Hyun Kim, and Chang-Hun Kim. Adaptive synthesis of distance fields. *IEEE Trans. Vis. & Comp. Graphics*, 18(7):1135–1145, 2011. 2
- [33] Jun Li, Kai Xu, Siddhartha Chaudhuri, Ersin Yumer, Hao Zhang, and Leonidas Guibas. Grass: Generative recursive autoencoders for shape structures. *ACM Transactions on Graphics (TOG)*, 36(4):1–14, 2017. 1, 2
- [34] Manyi Li, Akshay Gadi Patil, Kai Xu, Siddhartha Chaudhuri, Owais Khan, Ariel Shamir, Changhe Tu, Baoquan Chen, Daniel Cohen-Or, and Hao Zhang. Grains: Generative recursive autoencoders for indoor scenes. *ACM Transactions on Graphics (TOG)*, 38(2):1–16, 2019. 1, 2
- [35] Ruihui Li, Xianzhi Li, Ka-Hei Hui, and Chi-Wing Fu. Srgan: Sphere-guided 3d shape generation and manipulation. *ACM Transactions on Graphics (TOG)*, 40(4):1–12, 2021. 1, 2
- [36] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM Transactions on Graphics (SIGGRAPH)*, 21(4):163–169, 1987. 4
- [37] Rui Ma, Akshay Gadi Patil, Matthew Fisher, Manyi Li, Sören Pirk, Binh-Son Hua, Sai-Kit Yeung, Xin Tong, Leonidas Guibas, and Hao Zhang. Language-driven synthesis of 3d scenes from scene databases. *ACM Transactions on Graphics (TOG)*, 37(6):1–16, 2018. 1, 2
- [38] Paul Merrell. Example-based model synthesis. In *Proceedings of the 2007 symposium on Interactive 3D graphics and games*, pages 105–112, 2007. 1, 2
- [39] Paul Merrell and Dinesh Manocha. Continuous model synthesis. In *SIGGRAPH Asia*, pages 1–7. 2008. 1, 2
- [40] Paul Merrell and Dinesh Manocha. Model synthesis: A general procedural modeling algorithm. *IEEE Trans. Vis. & Comp. Graphics*, 17(6):715–728, 2010. 1, 2
- [41] 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 *European Conference on Computer Vision (ECCV)*, 2020. 3, 4
- [42] 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. 3
- [43] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2
- [44] Gen Nishida, Ignacio Garcia-Dorado, Daniel G Aliaga, Bedrich Benes, and Adrien Bousseau. Interactive sketching of urban procedural models. *ACM Transactions on Graphics (TOG)*, 35(4):1–11, 2016. 1, 2
- [45] Keunhong Park, Konstantinos Rematas, Ali Farhadi, and Steven M Seitz. Photoshape: Photorealistic materials for large-scale shape collections. *arXiv preprint arXiv:1809.09761*, 2018. 1, 2
- [46] Ben Poole, Ajay Jain, Jonathan T. Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. *arXiv*, 2022. 3
- [47] Gernot Riegler and Vladlen Koltun. Free view synthesis. In *European Conference on Computer Vision (ECCV)*, 2020. 3
- [48] Gernot Riegler and Vladlen Koltun. Stable view synthesis. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [49] Guodong Rong and Tiow-Seng Tan. Jump flooding in gpu with applications to voronoi diagram and distance transform. In *Proceedings of the 2006 Symposium on Interactive 3D Graphics and Games*, I3D '06, page 109–116, New York, NY, USA, 2006. Association for Computing Machinery. 5
- [50] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:20154–20166, 2020. 1, 2
- [51] Tamar Rott Shaham, Tali Dekel, and Tomer Michaeli. Singan: Learning a generative model from a single natural image. In *International Conference on Computer Vision (ICCV)*, 2019. 3, 4, 6
- [52] Andrei Sharf, Marc Alexa, and Daniel Cohen-Or. Context-based surface completion. In *ACM Transactions on Graphics (TOG)*, pages 878–887. 2004. 3
- [53] Assaf Shocher, Shai Bagon, Phillip Isola, and Michal Irani. Ingan: Capturing and retargeting the “dna” of a natural image. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4492–4501, 2019. 3
- [54] Denis Simakov, Yaron Caspi, Eli Shechtman, and Michal Irani. Summarizing visual data using bidirectional similarity. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8. IEEE, 2008. 2, 5
- [55] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5459–5469, 2022. 3
- [56] Ayush Tewari, Ohad Fried, Justus Thies, Vincent Sitzmann, Stephen Lombardi, Kalyan Sunkavalli, Ricardo Martin-Brualla, Tomas Simon, Jason Saragih, Matthias Nießner, et al. State of the art on neural rendering. In *Computer Graphics Forum*, volume 39, pages 701–727. Wiley Online Library, 2020. 3
- [57] Arash Vahdat and Jan Kautz. NVAE: A deep hierarchical variational autoencoder. In *Neural Information Processing Systems (NeurIPS)*, 2020. 3
- [58] Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. In *Neural Information Processing Systems (NeurIPS)*, 2021. 3

- [59] Kai Wang, Manolis Savva, Angel X Chang, and Daniel Ritchie. Deep convolutional priors for indoor scene synthesis. *ACM Transactions on Graphics (TOG)*, 37(4):1–14, 2018. 1, 2
- [60] Yujie Wang, Xuelin Chen, and Baoquan Chen. Singrav: Learning a generative radiance volume from a single natural scene. *arXiv preprint arXiv:2210.01202*, 2022. 3
- [61] Rundi Wu, Xuelin Chen, Yixin Zhuang, and Baoquan Chen. Multimodal shape completion via conditional generative adversarial networks. In *European Conference on Computer Vision (ECCV)*, pages 281–296. Springer, 2020. 6
- [62] Rundi Wu and Changxi Zheng. Learning to generate 3d shapes from a single example. *ACM Transactions on Graphics (TOG)*, 41(6), 2022. 3
- [63] Wenming Wu, Xiao-Ming Fu, Rui Tang, Yuhan Wang, Yu-Hao Qi, and Ligang Liu. Data-driven interior plan generation for residential buildings. *ACM Transactions on Graphics (TOG)*, 38(6):1–12, 2019. 1, 2
- [64] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Or Litany, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond. *arXiv preprint arXiv:2111.11426*, 2021. 3
- [65] Wang Ximeng. A thousand li of rivers and mountains. [https://en.wikipedia.org/wiki/Wang\\_Ximeng](https://en.wikipedia.org/wiki/Wang_Ximeng), 1113. 7
- [66] Rui Xu, Xintao Wang, Kai Chen, Bolei Zhou, and Chen Change Loy. Positional encoding as spatial inductive bias in gans. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13569–13578, 2021. 3
- [67] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 3
- [68] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems (NeurIPS)*, 33, 2020. 3
- [69] Alex Yu, Sara Fridovich-Keil, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2, 3, 4
- [70] Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *International Conference on Computer Vision (ICCV)*, 2021. 3
- [71] Lap Fai Yu, Sai Kit Yeung, Chi Keung Tang, Demetri Terzopoulos, Tony F Chan, and Stanley J Osher. Make it home: automatic optimization of furniture arrangement. *ACM Transactions on Graphics (TOG)*, 30(4), 2011. 1, 2
- [72] Zaiwei Zhang, Zhenpei Yang, Chongyang Ma, Linjie Luo, Alexander Huth, Etienne Vouga, and Qixing Huang. Deep generative modeling for scene synthesis via hybrid representations. *ACM Transactions on Graphics (TOG)*, 39(2):1–21, 2020. 1, 2
- [73] Howard Zhou, Jie Sun, Greg Turk, and James M Rehg. Terrain synthesis from digital elevation models. *IEEE Trans. Vis. & Comp. Graphics*, 13(4):834–848, 2007. 1, 2
- [74] Chenyang Zhu, Kai Xu, Siddhartha Chaudhuri, Renjiao Yi, and Hao Zhang. Scores: Shape composition with recursive substructure priors. *ACM Transactions on Graphics (TOG)*, 37(6):1–14, 2018. 1, 2
- [75] Jie Zhu, Yanwen Guo, and Han Ma. A data-driven approach for furniture and indoor scene colorization. *IEEE Trans. Vis. & Comp. Graphics*, 24(9):2473–2486, 2017. 1, 2

## Supplementary Material: Patch-based 3D Natural Scene Generation from a Single Example

Weiyu Li<sup>1\*</sup>† Xuelin Chen<sup>2\*</sup>† Jue Wang<sup>2</sup> Baoquan Chen<sup>3</sup>  
<sup>1</sup> Shandong University    <sup>2</sup> Tencent AI Lab    <sup>3</sup> Peking University



Figure 1. Artistic imagery created with various 3D scenes generated by our method (background sky post-added). The original exemplar scenes for generating these results are: (from left to right and top to bottom): The Vast Land [35], Heal Mountain [17], Devil’s Tower ©2022 Google, Callanish [24], Meteora ©2022 Google. Green Island [18],

## A. More Visual Results

More artistic pieces, that are created with high-quality and diverse general natural scenes generated by our method, are shown in Figure 1. Moreover, in Figure 2, we also visualize the underlying high-quality and diverse geometry of more generated samples. Figure 15, 16 presents more samples generated with our method.

In addition to more dynamic 3D viewing of a large collection of generated samples presented in the supplementary video, please also see the anonymous project website <http://wyysf-98.github.io/Sin3DGen> for a more immersive view into our 3D results.

## B. Implementation Details

All synthesis results presented in this paper share the following default setting, unless specified. *We will release the code for reproducing the results presented in this paper, upon the publication of this work.*

**Random Generation.** By default, the synthesized scene shares the same bounding box with in the random synthesis task. Each scene is located inside a cuboid, of which the aspect ratio varies according to different exemplars. In Table 1, we list the final resolution of  $N$  in the pyramidal generation framework for each exemplar scene. At the  $N$ -th scale in the multi-scale framework, the resolution along the maximum dimension of  $N$  is set to 121, considering the trade-off between the quality and computational efficiency of the generation. The resolution along the maximum dimension of the higher-resolution *high* is 512. The scaling factor between consecutive scales in the pyramid is  $r = 4/3$ , and the coarsest resolution is 16. We use  $N = 7$  in the pyramid, which results in 8 scales in total. The patch size at all scales is set to  $p = 5$ . We set the number of PCA components to 3, the truncate scale of SDF  $t = 3 \times w$ , where  $w$  is the voxel size. The weight of the appearance feature is  $w_a = 0.5$ , the completeness trade-off weight  $\alpha = 0.01$ , and the initial noise  $\sigma = 0.5$ . At coarser scales ( $n < 5$ ), exact NNF is applied  $T_e = 10$  times, which means the value-based NNF is performed with  $T_e - 1 = 9$  times and followed by one mapping-based NNF search; At the finest scale ( $n \geq 5$ ), approximate NNF via PatchMatch is performed  $T_a = 2$  times. The “jump flood” radius is 8, and the random search radius is fit to the max resolution of the current exemplar.

**Applications.** In contrast to the random synthesis task, the  $\sigma$  for noise used in all applications is set to 0. More specifically:

- 1) **Retargeting:** The goal is to resize a 3D scene to a target size (typically of a different aspect ratio), while

maintaining the local patches in the exemplar. We simply set the resolution of  $N$  to the target size, and the resolution of  $0, \dots, N-1$  is adapted accordingly with the default scaling factor  $r$ .

- 2) **Editing:** Users can manipulate on a 3D proxy, which can be the underlining mapping field or mesh, for editing an exemplar or generated scene, such as removal, duplication, and modification. The manually manipulated proxy is then converted and fed as the initial guess at the coarsest scale for synthesizing the final scene. As editing the 3D scene requires more meticulous 3D interaction, we set the resolution at the coarsest scale to a higher value (resolution along the max dimension is 28), and use 6 scales in total. We perform the exact NNF at the first 3 scales, followed by 3 finer scales with the approximate NNF.
- 3) **Structural analogies:** Given two scenes A and B, we create a scene with the patch distribution of A, but which is structurally aligned with B. This is realized by using the exemplar pyramid of A, and an identity mapping as the initial guess, but replacing  $o(0)$  with the transformed features in B, and vice versa. As the content of the generated scene at the coarsest scale is already specified by  $o(0)$ , the pyramidal generation starts with a higher-resolution scale (51 voxels along the max dimension), finishes the generation with 4 scales in total, and performs exact NNF in the first scale.
- 4) **Re-decoration:** Trivially, we do not need to re-synthesize the scene in the re-decoration application. Given an already generated scene, a novel scene can be obtained by simply remapping the coordinate-based synthesis result to an exemplar of different appearance.

## C. Datasets

We collected a rich variety of 3D scene models to examine the performance of our method on random scene generation, ranging from rocks to plants, sculptures, landscapes, terrains, artistic scenes, etc. For each 3D scene model, we render 200 images at the resolution  $1024 \times 1024$ , with cameras distributed on a sphere in Blender [10]. Then the Plenoxels-based exemplar pyramid is obtained via coarse-to-fine training on these images. Notably, in Figure 11, we also demonstrate our method on real-images collected from a real-world scenic site. To this end, we collect 300 images with the resolution  $1280 \times 720$  from Google Earth Studio [1], where we can manually specify cameras for simulating a drone programmed to fly over a scenic spot, for training the Plenoxels-based exemplar. Specifically, we move the camera in a spiral motion and gradually elevate the camera from a high altitude to a low altitude. Then, we

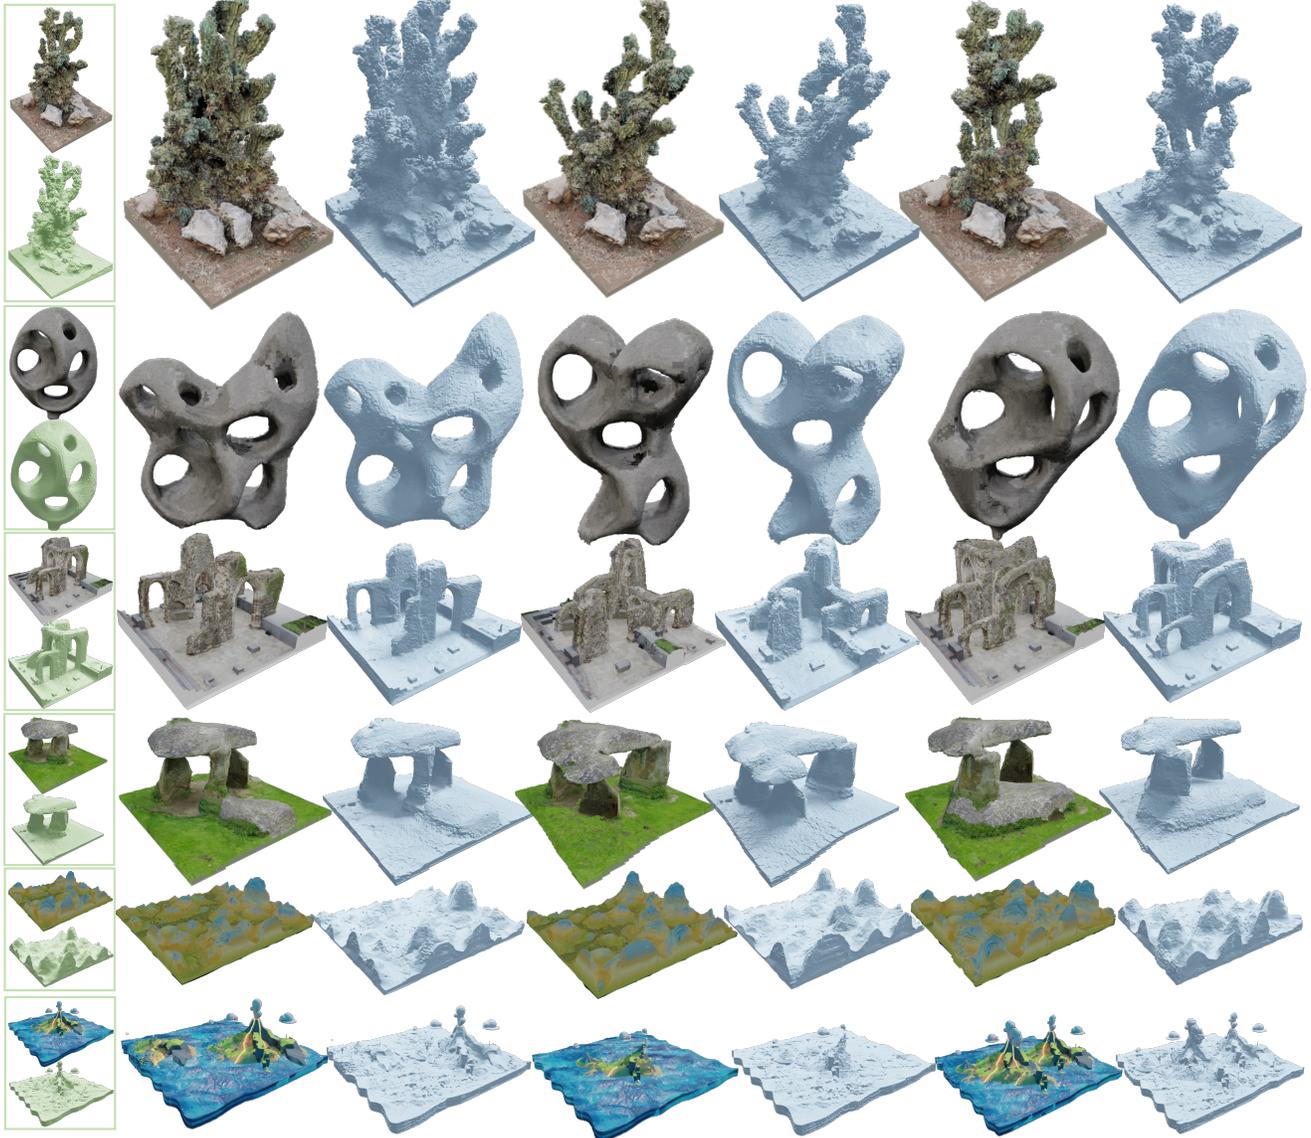


Figure 2. Visualization of the high-quality geometries generated by our methods. The original exemplar scenes for generating these results are: (from top to bottom) Cactus Cereus [25], Stone Sculpture [3], St Alphege [6], Spinsters Rock [4], The Vast Land [35], Volcano Island Lowpoly [5].

use COLMAP [31, 32] to estimate the camera parameters. More details can be found in the video. Figure 13 presents the visuals of all exemplars used in this paper.

#### D. More Analysis

In general, our method is robust to varying hyperparameters to some extent. We shall show the effects of using different parameters in the following.

**Effects of Different Noise  $z_0$ .** In Figure 3, we show the results obtained by different initial guesses, i.e., the identity

mapping shuffled with different noises, at the coarsest level.

**Effects of Different  $w_a$ .** Empirically, the trade-off parameter  $w_a$  for balancing the appearance and geometry feature is set to 0.5 by default. While we have shown this setting yields robust and high-quality generation, we shall also demonstrate the effects of varying weights. Generally, in Figure 4, we can see that our method is robust to varying  $w_a$  to some extent. However, setting  $w_a$  to an extremely small value, which pays much attention to the geometry feature,

Table 1. Resolution configuration for figures in the **main paper**.

Figure	Data	Resolution of $S_N$
Fig. 1	Cactus Cereus [25]	$92 \times 108 \times 121$
Fig. 2 & Fig. 5	Green Island [18]	$121 \times 121 \times 47$
Fig. 3 & Fig. 5	St Alphege [6]	$121 \times 121 \times 92$
Fig. 5	Calda House [7]	$121 \times 121 \times 71$
	Callanish [24]	$121 \times 121 \times 47$
	Stone Arch [36]	$121 \times 51 \times 71$
	Desert Lowpoly [12]	$121 \times 121 \times 92$
	Meteora ©2022 Google	$121 \times 121 \times 47$
	Spinsters Rock [4]	$121 \times 121 \times 84$
	Stone Sculpture [3]	$108 \times 84 \times 121$
Fig. 7	Volcano Island Lowpoly [5]	$121 \times 121 \times 71$
	The Vast Land [35]	$121 \times 121 \times 47$
Fig. 10	Mountain with Lakes [40]	$121 \times 121 \times 72$
Fig. 10	Autumn Camping [21]	$121 \times 99 \times 72$
	Winter Camping [23]	$121 \times 99 \times 72$
Fig. 5 & Fig. 10	Devil's Tower ©2022 Google	$121 \times 121 \times 63$
	Cactus Saguaro [13]	$71 \times 72 \times 121$
	Camping Lowpoly [22]	$121 \times 99 \times 72$
	Mountain [29]	$121 \times 121 \times 84$
	Stylized Cactus [30]	$121 \times 121 \times 121$

Table 2. Resolution configuration for *high-resolution synthesis* and the *retargeting application* in the **main paper**.

Figure	Exemplar	Resolution of $S_N$
Fig. 6	The Vast Land [35]	$288 \times 288 \times 112$
Fig. 10	Cactus Cereus [25]	$92 \times 108 \times 152$
	Stone Arch [36]	$242 \times 51 \times 71$
	Tiny Castle [9]	$121 \times 63 \times 237$
	Train Wagon [2]	$47 \times 182 \times 47$

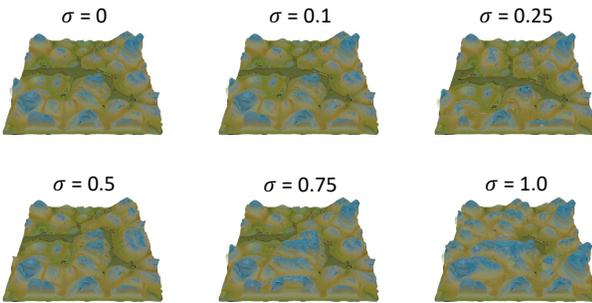


Figure 3. Effects of different initial noise. Intuitively, noise with smaller  $\sigma$  values leads to more similar scenes to the exemplar, while larger values result in more diverse ones.

will lead to inconsistent appearance and artifacts. On the other end, if we put all emphasis on the geometry feature, the synthesis fails and produces empty scenes.

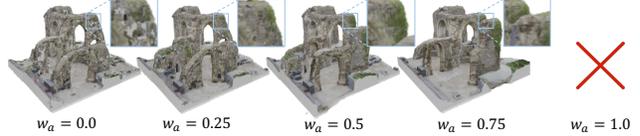


Figure 4. Effects of varying  $w_a$ . While  $w_a$  around 0.5 produces relatively stable results, extreme values introduce visual inconsistency and artifacts (see left  $w_a$ ), or even fail (see right  $w_a = 1.0$ ).

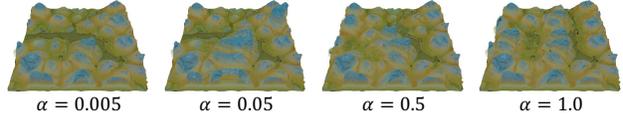


Figure 5. Effects of varying  $\alpha$ .  $\alpha$  serves as a coarse control knob for visual completeness. The river becomes shorter, suggesting lower visual completeness, as  $\alpha$  increases.

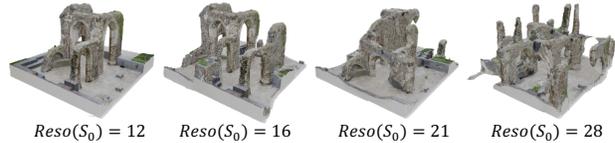


Figure 6. Effects of different resolutions of  $o$ . Lower-resolution  $o$  (larger receptive field at coarse scale) results in less structural diversity, producing almost identical to the exemplar. On the contrary, with smaller receptive fields at the coarsest scale, the global arrangement can not be well preserved (see messy structures on the right).

**Effects of Different  $\alpha$ .** Figure 5 presents the effects of varying  $\alpha$  for different levels of visual completeness. Nevertheless, we also found the degrees of such control may not be always perfectly explicit, which we also observed with [14].

**Effects of Different Resolution for  $o$  at the Coarsest Scale.** Given a fixed patch size, which is  $p = 5$  in our work, a larger resolution at the coarsest scale suggests a smaller effective receptive field (the same concept as in convolutional neural networks) and less-considered global layouts at the coarser scales, and vice versa. In our work, we by default use the setting where the patch at the coarsest scale captures 1/3 of the content in the exemplar, balancing the local diversity and global layout. In Figure 6, we also show the impact of varying resolutions at the coarsest scale, that capture contents of different sizes in the generation.

**Effects of Different Resolution of  $N$  at the Finest Scale.** The synthesis at finer scales only considers visual coherence and adds local details. We have shown that synthesizing

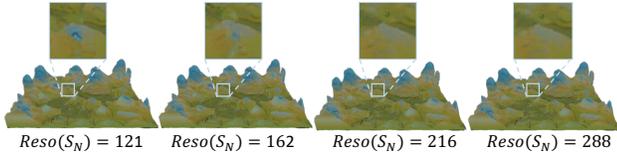


Figure 7. Effects of different resolutions of  $N$ . The overall visual synthesis is stable at a resolution of 121. Some minor aliasing can be found in lower resolution (left), with a larger scale for synthesis, more details can be obtained.

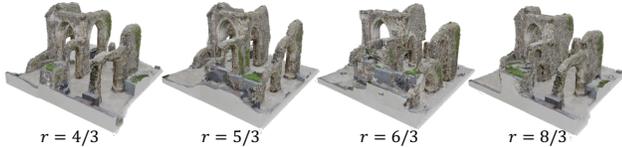


Figure 8. Effects of different downscale ratios  $r$ . As  $r$  increasing, fine structures, e.g., the arch doors, gradually disappears, producing bulky structures and less diversity in each instance.

with a maximum resolution 121 by default in the pyramid is sufficient in most cases for the trade-off between quality and efficiency. Moreover, in Figure 7, we show that using higher resolutions for  $N$  only leads to negligible visual gains. Besides, we also observed that as we use approximate NNF at finer scales, the inaccurate NNF search may introduce some wrong patches and lead to performance degradation in the generation.

**Effects of Different Downscale Ratio  $r$ .** The downscale ratio  $r$  used for building the pyramid affects the transition between scales. As the ratio increases, the transition of the generation between scales becomes more inconsistent and unstable due to large gaps between consecutive scales, leading to the loss of fine structures and less diversity as shown in Figure 8.

**Effects of the Truncated Scale  $t$ .** The truncated scale  $t$  controls the range of geometric features we keep for patching matching. Smaller truncate scales only consider information near to surface and degrade to the occupancy field, which may produce many tiny pieces and incomplete instances, while larger values lead to blurry results. Figure 9 presents the visual results.

**Only Exact or Approximate NNF-3D.** The mix use of exact NNF and approximate NNF in our framework has shown the efficacy and efficiency in 3D generation. Using only exact NNF would quickly lead to prohibitive computational cost and prevent us from synthesizing high-resolution results. See Table 3 for the detailed computation overhead. On the other end, only using approximate NNF all the time

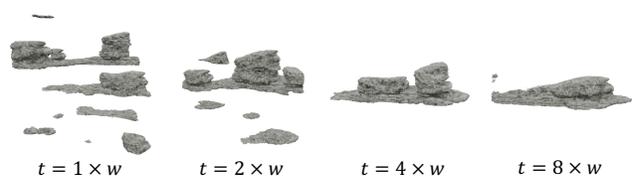


Figure 9. Effects of different truncate scales  $t$ , where  $w$  is the current voxel size at each scale. A small  $t$  will lead to separated fragments due to the loss of geometry information. On the contrary, a very larger truncate scale will blur the geometric feature and only synthesize instances with a coarse shape.

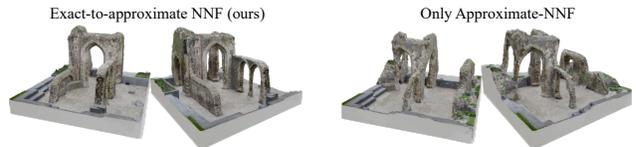


Figure 10. Approximate NNF-only generation. See the distorted arches occurred in the results on the right.

will harm the generation, producing distorted results, as the approximate NNF is inaccurate. In Figure 9, we show the visual results when only using approximate NNF-3D.

## E. More Experiments

**Working with Unbounded Scenes.** Benefiting from using Plenoxels, which trains on 2D images, for representing the input scene, our method can also work on images collected from a *real-world unbounded* scene. To this end, we use COLMAP [31, 32] to estimate the camera parameters, and model the background using an implicit neural network, similar to NeRF++ [39]. Figure 11 presents the results, more visual details can be found in the video. Note that, existing NeRF-based models often struggle in handling "unbounded" real-world scenes, and disentangling the foreground and background. Nevertheless, some works [8, 20] attempt to tackle these problems, showing promising results. We believe these methods can help boost the performance of our method on more real-world scenes, which, however, is not in the scope of this work and stimulates future research.

**Computational Overhead.** In Table 3, we reported the detailed time and memory usage for the exact-only NNF and approximate-only NNF. As aforementioned, using either exact-only or approximate-only NNF would not be satisfying, and our exact-to-approximate scheme is the key to enable synthesizing high-quality results with limited computational resources.



Figure 11. Samples generated with images collected from a real-world scenic site – Bryce canyon ©2022 Google. Notably, we only synthesize the region of interest (i.e., the odd rocks) and the background is disentangled out by modeling with an independent implicit neural network.

Table 3. Computational overhead. We report the time and memory consumption in each NNF iteration for exact-only and approximate-only NNF.

Resolution	Only Exact NNF		Only Approximate NNF	
	Time (s)	Memory (GB)	Time (s)	Memory (GB)
16 × 16 × 5	0.01	0.07	0.55	0.07
21 × 21 × 7	0.01	0.12	0.59	0.14
28 × 28 × 10	0.02	0.68	0.69	0.21
38 × 38 × 14	0.08	1.48	0.98	0.43
51 × 51 × 19	0.57	3.66	1.69	0.92
68 × 68 × 26	4.06	9.74	3.62	2.02
91 × 91 × 35	27.47	25.17	8.47	4.77
121 × 121 × 47	N/A	N/A	20.07	11.28
162 × 162 × 63	N/A	N/A	48.91	17.05
216 × 216 × 84	N/A	N/A	116.09	21.95
288 × 288 × 112	N/A	N/A	278.53	23.27
384 × 384 × 150	N/A	N/A	662.23	26.60

**Failure Cases.** As mentioned in the paper, our method favors scenes with complex structures and textures for matching the internal distribution, lacking sufficient diverse patch candidates will lead to mode collapse issues. Besides, with voxelized volumetric representations, our method can not perfectly synthesize scenes with tiny thin structures. Moreover, our method operates on the patch level, so we can not guarantee that highly semantic or structural features in the exemplar can be preserved intactly in the synthesized results. Figure 12 shows some failure cases when working with our method.

## F. Evaluation

### F.1. Baselines

**GRAF [33]** We use the official implementation, and replace the camera poses with ones in our work. We follow the default setting for training, one model for each exemplar scene is trained with renderings of resolution  $512^2$  for  $7200k$  samples, which takes about 3 days in a single V100 GPU. The final visuals are rendered at the resolution  $512^2$ .

**StyleNeRF [15]** We use the official release of StyleNeRF. Same as GRAF, we replace the camera extrinsic and intrinsic

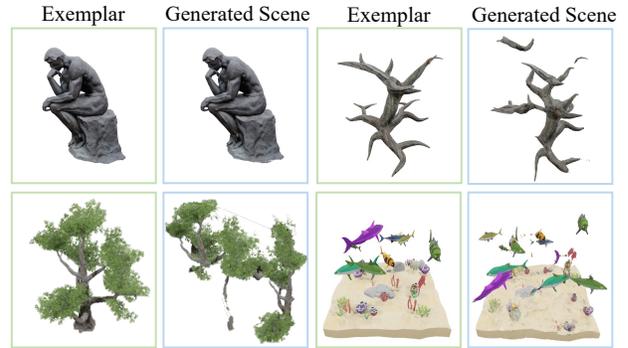


Figure 12. Failure cases. An exemplar scene, that does not have sufficient diverse patch candidates, would result in identical generation results to the input (See top left sculpture [27]). Our method failed on scenes with tiny thin structures, such as branches [16] and trees [26]. Scenes with highly semantic or structural information can not be correctly handled (See broken fishes on the bottom right [11,28]).

parameters with the real distribution and set the background to white. All models are trained in the resolution of  $512^2$  following the default setting by going through  $6000k$  images for about 3 days using 4 V100 GPUs.

**GPNN-3D** We naively extend the GPNN [14] for working on Plenoxels-based exemplar scenes. The density value and SH values are normalized to fit  $[-1, 1]$ , and we follow all parameters as described in the original paper. The maximum resolution reached by GPNN-3D is only 38 due to computational efficiency issues.

### F.2. Camera Pose Sampling

To quantitatively evaluate the synthesized scenes from 2D projections, we uniformly sample  $K = 50$  camera poses on the upper hemisphere with radius  $R = 2.5$ , and use elevation angles range from  $0^\circ$  to  $90^\circ$ . The focal length of the camera is set to 512 times the pixel size, equivalently  $\text{FoV} \approx 39.6^\circ$ , for all exemplar scenes.

### F.3. Metrics

For each method, we produce 50 generated scenes on each of the evaluated exemplars, render 50 multi-view images and extract the 3D geometric surface points of the exemplar and each in the generated, and then rate the performance using a combination of several common metrics in both 2D and 3D generation:

**Visual Quality** measures how well the model captures the internal statistics of the input exemplar from the 2D perspective, by simply computing the averaged SIFID [34] over multiple views of a generated scene. Concretely, for

each image rendered from a generated scene, we compute the single image SIFID of this image against the image rendered at the associate viewpoint in the exemplar scene. Then the SIFID-MV for a generated scene is the average over the multiple views. We finally report the mean SIFID-MV averaged over multiple generated scenes.

**Visual Diversity** of the set of generated scenes is measured via extending the image diversity score as in [34] to multi-view images of a scene. First, under each view, we calculated the standard deviation (std) of the intensity values of each pixel over 50 images rendered from 50 generated scenes, averaged it over all pixels, and normalized by the std of the intensity values of the image rendered from the exemplar. Then, we report the average of std values obtained at 50 views as the Visual Diversity of a set of generated scenes.

**Geometry Quality** of a generated scene is measured as the Minimal Matching Distance [37] (multiplied by  $10^2$ ) between the set of generated patches and exemplar patches (represented as point clouds sampled on the surface). As mentioned in the paper, Plenoxels often produce invisible noise, so we only pick point cloud patches on the surface. Specifically, for a scene represented by a discrete volume of resolution  $256^3$ , we extract mesh using Marching Cubes [19], and evenly sample 102400 points from the mesh surface. To extract patches, we randomly pick 1000 points center, then combined them with the nearest 1024 points via k-NN search. Then the geometry quality of a generated scene is calculated as the MMD between the set of patches in a generated scene and the set of patches in the exemplar scene. We then report the averaged geometry quality score over 50 generated scenes.

**Geometry Diversity** of generated scenes is measured by summing up all the differences among the 50 generated scenes, i.e., Total Mutual Difference as in [37]. Specifically, we evenly sample 10240 points on the surface of each generated mesh to obtain a point cloud, forming a set of 50 scene point clouds. Empty scenes are deprecated to calculate the geometry diversity. Then the Geometry Diversity of 50 generated scenes is reported as the TMD calculated on this set of point clouds.

## References

- [1] Google earth studio. <https://earth.google.com/studio/>. 2
- [2] 3ddominator. Train wagon. <https://sketchfab.com/3d-models/train-wagon-42875c098c33456b84bcfcdc4c7f1c58>, 2019. License: **CC Attribution**. 4
- [3] 3dhdsan. Stone sculpture i. <https://sketchfab.com/3d-models/stone-sculpture-i-fb6e253b75db40c3b38e3977ee60bc5c>, 2018. License: **CC Attribution**. 3, 4
- [4] Alec. Spinsters rock. <https://sketchfab.com/3d-models/spinsters-rock-14f654eb8c4d42d991c61f907379644a>, 2018. License: **CC Attribution**. 3, 4
- [5] Animateria. Volcano island lowpoly. <https://sketchfab.com/3d-models/volcano-island-lowpoly-4a6591dc9fee40d8bfda8350683af9af>, 2020. License: **CC Attribution**. 3, 4
- [6] artfletch. <https://sketchfab.com/3d-models/st-alphage-church-ruin-6979d7daf8ad4bd887694ef65cleeeae>, 2018. License: **CC Attribution**. 3, 4
- [7] artfletch. Calda house. <https://sketchfab.com/3d-models/calda-house-439588b47dfb4c0984b20555163db514>, 2020. License: **CC Attribution**. 4
- [8] Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. *CVPR*, 2022. 5
- [9] bilgehan.korkmaz. Tiny castle. <https://sketchfab.com/3d-models/tiny-castle-2836b219403d4fd487d0434f88b65363>, 2018. License: **CC Attribution**. 4
- [10] Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Blender Institute, Amsterdam, 2022. 2
- [11] EdwinRC. L-p shark cage diving scene. <https://sketchfab.com/3d-models/l-p-shark-cage-diving-scene-5bf4fcbbc6ea40b69c734f54180fa38c>, 2018. License: **CC Attribution**. 6
- [12] EdwinRC. Low poly desert. <https://sketchfab.com/3d-models/low-poly-desert-c25524a942704f5c833ac53b645cda82>, 2019. License: **CC Attribution**. 4
- [13] evolveduk. Cactus 2 (saguaro). <https://sketchfab.com/3d-models/cactus-2-saguaro-e7ba6d448e2c44ce89c716fa71cbb716>, 2022. License: **CC Attribution**. 4
- [14] Niv Granot, Ben Feinstein, Assaf Shocher, Shai Bagon, and Michal Irani. Drop the gan: In defense of patches nearest neighbors as single image generative models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 4, 6
- [15] Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. Stylenerf: A style-based 3d aware generator for high-resolution image synthesis. In *International Conference on Learning Representations (ICLR)*, 2022. 6
- [16] Harry\_L. Old dead tree. <https://sketchfab.com/3d-models/old-dead-tree-8c680561463a4e909fd24b4cc8bc7e48>, 2018. License: **CC Attribution**. 6

- [17] Hirnlaich. Heal mountain. <https://sketchfab.com/3d-models/stacked-heal-mountain-4ea3d857952347578dbb57f4e642ce9f>, 2022. License: **CC Attribution. 1**
- [18] Bilgehan Korkmaz. Green island. <https://sketchfab.com/3d-models/green-island-79445f7196e24f4e833247f82aeaea0b>, 2019. License: **CC Attribution. 1, 4**
- [19] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM Transactions on Graphics (SIGGRAPH)*, 21(4):163–169, 1987. **7**
- [20] Inc. Luma AI. Luma ai. <https://lumalabs.ai/>, 2022. **5**
- [21] marzipanne. Autumn camping. <https://sketchfab.com/3d-models/autumn-camping-e41d700dcc0342cb92b8e0af59c1dabc>, 2020. License: **CC Attribution. 4**
- [22] marzipanne. Camping low poly. <https://sketchfab.com/3d-models/camping-low-poly-d05f47384a6046ecaf39f857bd06c5a9>, 2020. License: **CC Attribution. 4**
- [23] marzipanne. Winter camping. <https://sketchfab.com/3d-models/winter-camping-01c176fd6b204446b495f60056a7a064>, 2020. License: **CC Attribution. 4**
- [24] Megalitharchive. Callanish. <https://sketchfab.com/3d-models/callanish-bf4e9b9a0e764b39b196537377c082ea>, 2021. License: **CC Attribution. 1, 4**
- [25] netgis. Cactus cereus. <https://sketchfab.com/3d-models/cactus-cereus-1a51f9da8f62466797b271648e652a3d>, 2021. License: **CC Attribution. 3, 4**
- [26] Nicholas-3D. More realistic trees free! <https://sketchfab.com/3d-models/more-realistic-trees-free-b5b506fc4f5d4af9b546283bdf0c6a15>, 2022. License: **CC Attribution. 6**
- [27] noe 3d.at. Der denker. <https://sketchfab.com/3d-models/der-denker-434542838ef140d4a706dc42e775eb12>, 2020. License: **CC Attribution. 6**
- [28] Oneeeee. Mini low poly fish pack. <https://sketchfab.com/3d-models/mini-low-poly-fish-pack-8a9b0be8fa6c4e3e9b6cd9f1f4e0967a>, 2022. License: **CC Attribution. 6**
- [29] Resistance Studios. Black mountain. <https://sketchfab.com/3d-models/mountain-vr-test-aad376d6ad634d0bb5964a863f1e9a51>, 2018. License: **CC Attribution. 4**
- [30] rubenve. Stylized cactus. <https://sketchfab.com/3d-models/stylized-cactus-d515ed6d19184b80830a6feba04863a8>, 2019. License: **CC Attribution. 4**
- [31] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. **3, 5**
- [32] Johannes Lutz Schönberger, Enliang Zheng, Marc Pollefeys, and Jan-Michael Frahm. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision (ECCV)*, 2016. **3, 5**
- [33] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3d-aware image synthesis. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:20154–20166, 2020. **6**
- [34] Tamar Rott Shaham, Tali Dekel, and Tomer Michaeli. Singan: Learning a generative model from a single natural image. In *International Conference on Computer Vision (ICCV)*, 2019. **6, 7**
- [35] Shahriar Shahrabi. The vast land. <https://sketchfab.com/3d-models/the-vast-land-733b802f5a4743ef99ad574279d49920>, 2021. License: **CC Attribution. 1, 3, 4, 10**
- [36] Werniech. Large stone arch 3d. <https://www.turbosquid.com/3d-models/large-stone-arch-3d-1908101>, 2022. License: **3D Model License: Standard. 4**
- [37] Rundi Wu, Xuelin Chen, Yixin Zhuang, and Baoquan Chen. Multimodal shape completion via conditional generative adversarial networks. In *European Conference on Computer Vision (ECCV)*, pages 281–296. Springer, 2020. **7**
- [38] Wang Ximeng. A thousand li of rivers and mountains. [https://en.wikipedia.org/wiki/Wang\\_Ximeng](https://en.wikipedia.org/wiki/Wang_Ximeng), 1113. **10**
- [39] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv:2010.07492*, 2020. **5**
- [40] Šimon Ustal. Grassy mountains with lakes. <https://sketchfab.com/3d-models/grassy-mountains-with-lakes-408364942d794699bbf99aaced81881d>, 2021. License: **CC Attribution. 4**

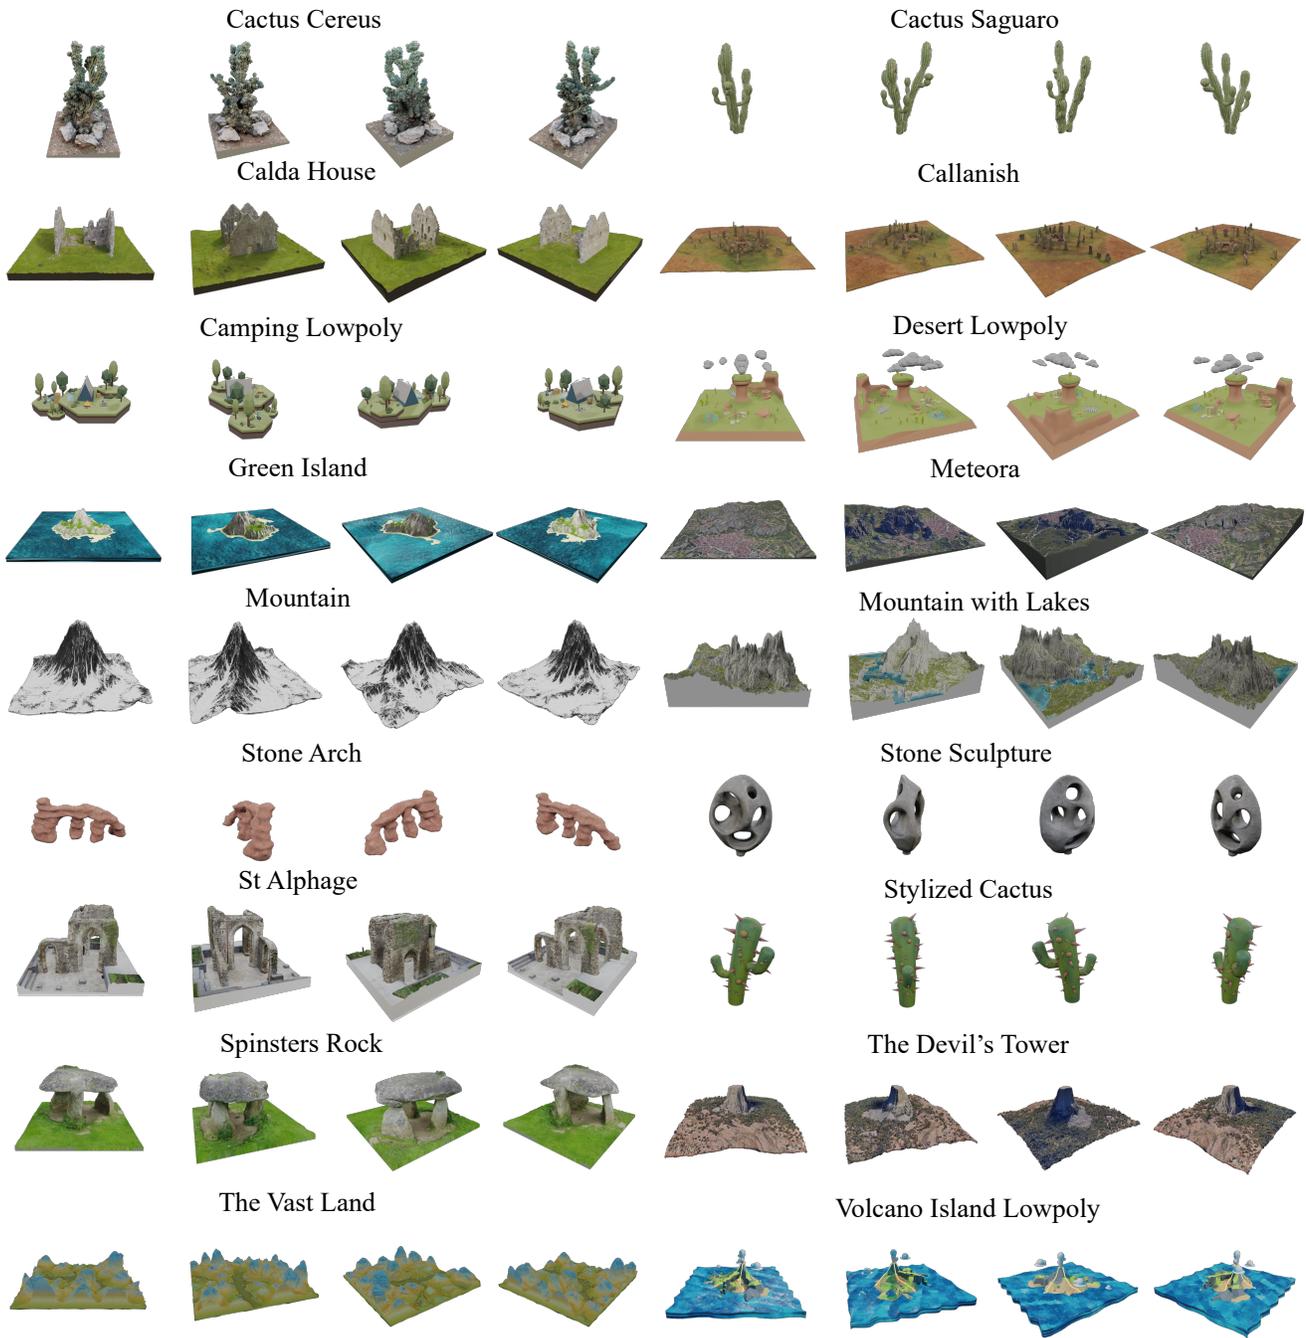


Figure 13. Visualization of the exemplars used in the **main paper**. More scenes can be found in the project page and video.

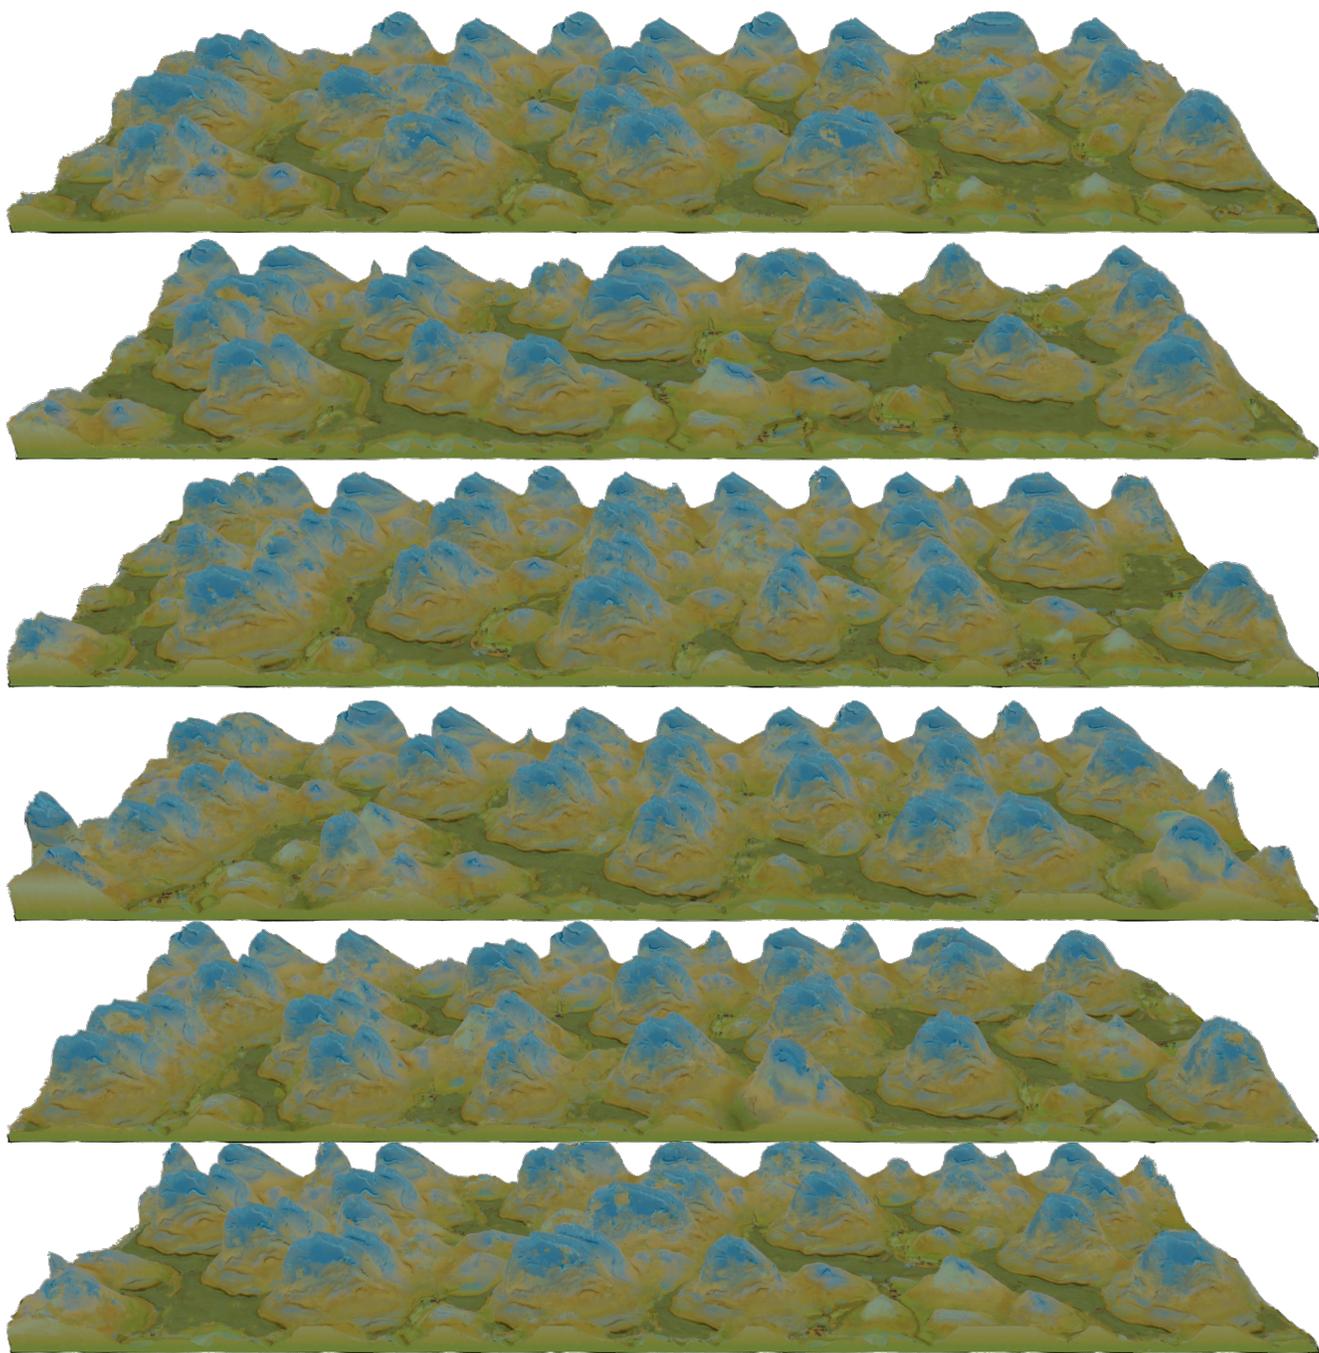


Figure 14. Diverse "A Thousand Li of Rivers and Mountains" [38] generated from The Vast Land [35] by our method. Specification:  $N$  -  $288 \times 288 \times 112$ ,  $^{high}$  -  $512 \times 512 \times 200$ ,  $N$  -  $747 \times 288 \times 112$ ,  $^{high}(N)$  -  $1328 \times 512 \times 200$ , final rendering resolution -  $4096 \times 1024$ .

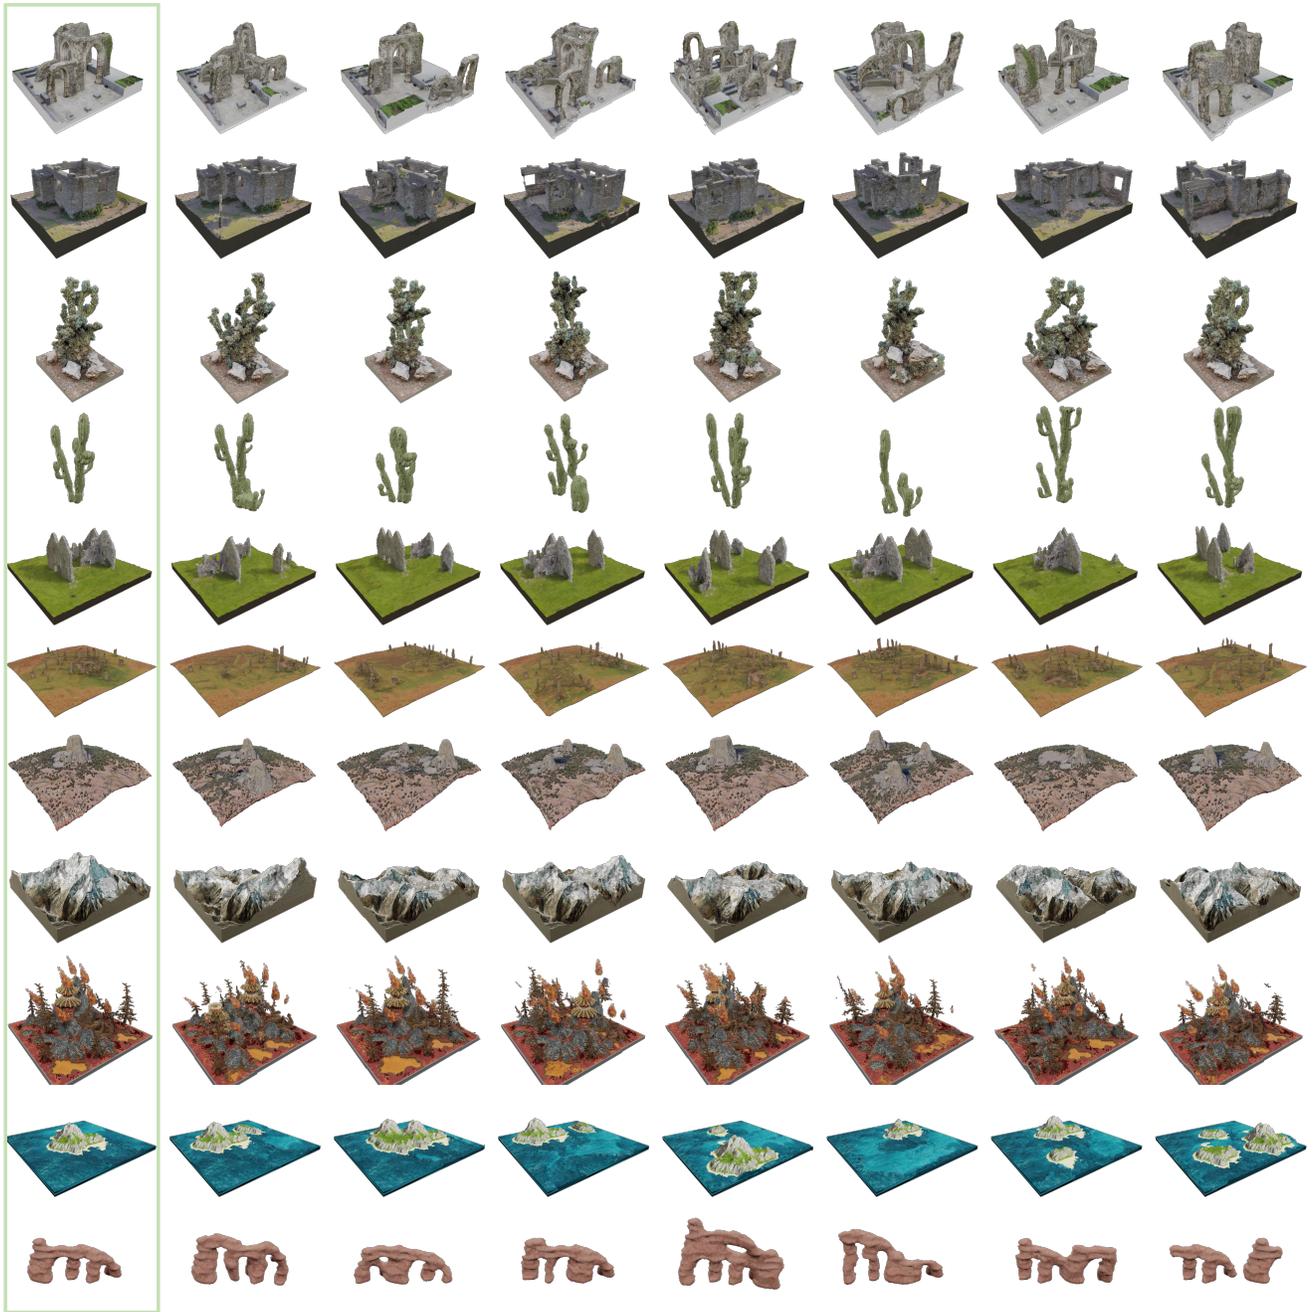


Figure 15. Diverse samples generated by our method. The input is shown in the green box on the left, followed by 7 generated novel scenes.

