

SINGRAV: LEARNING A GENERATIVE RADIANCE VOLUME FROM A SINGLE NATURAL SCENE

Anonymous authors

Paper under double-blind review

ABSTRACT

We present a 3D generative model for general natural scenes. Lacking necessary volumes of 3D data characterizing the target scene, we propose to learn from a single scene. Our key insight is that a natural scene often contains multiple constituents whose geometry, texture, and spatial arrangements follow some clear patterns, but still exhibit rich variations over different regions within the same scene. This suggests localizing the learning of a generative model on substantial local regions. Hence, we exploit a multi-scale convolutional network, which possesses the spatial locality bias in nature, to learn from the statistics of local regions at multiple scales within a single scene. In contrast to existing methods, our learning setup bypasses the need to collect data from many homogeneous 3D scenes for learning common features. We coin our method SinGRAV, for learning a Generative *R*Adiance Volume from a *S*ingle natural scene. We demonstrate the ability of SinGRAV in generating plausible and diverse variations from a single scene, the merits of SinGRAV over state-of-the-art generative neural scene methods, as well as the versatility of SinGRAV by its use in a variety of applications, spanning 3D scene editing, composition, and animation. Code and data will be released to facilitate further research.

1 INTRODUCTION

Recently, 3D generative modeling has made great strides via gravitating towards neural scene representations, which boast unprecedentedly photo-realism. Generative neural scene models (Schwarz et al., 2020; Niemeyer & Geiger, 2021; Chan et al., 2021b; Gu et al., 2022; Chan et al., 2021a) can now draw class-specific realistic scenes (e.g., cars and portraits), offering a glimpse into the boundless universe in the virtual. Yet an obvious question is how we can go beyond class-specific scenes, and achieve similar success with *general natural* scenes, creating *at scale* diverse scenes of more sorts. This work presents, to our knowledge, the first endeavor towards answering this question.

While neural scene generation has boosted the field via learning from images with differentiable projections, there still exists strong dependence on datasets containing images of many homogeneous scenes. Collecting homogeneous data for each scene type ad hoc is cumbersome, and would become prohibitive when the scene type of interest varies dynamically. Fortunately, parallels can be drawn following nature, who implicitly maintains a rich variety of natural scenes. In nature, no scenes ever stayed absolutely unaltered — imagine, as the stars shift, a river changes its current and little pebbles get scattered and weathered variously. Analogically, in scene creation, new scenes can be created via "varying" an existing one, yet still resembling the original (see Figure 1).

Following this spirit, we present SinGRAV, for a *generative radiance volume* learned to synthesize variations from a *single* general natural scene. Building upon recent advances in 3D generative modeling, SinGRAV also learns from visual observations of the target scene via differentiable volume rendering, but the training supervision, scene representation, and network architecture all have to be altered fundamentally to meet unique challenges arising from learning with a single scene. Predominantly, the supervision comes from merely a single scene, which would manifestly elude state-of-the-art models that learn common features across homogeneous scenes.

Natural scenes often contain multiple constituents whose geometry, appearance, and spatial arrangements follow some clear patterns, but still exhibit rich variations over different regions within the same scene. This naturally inspires one to learn from the statistics of internal regions via *localizing*

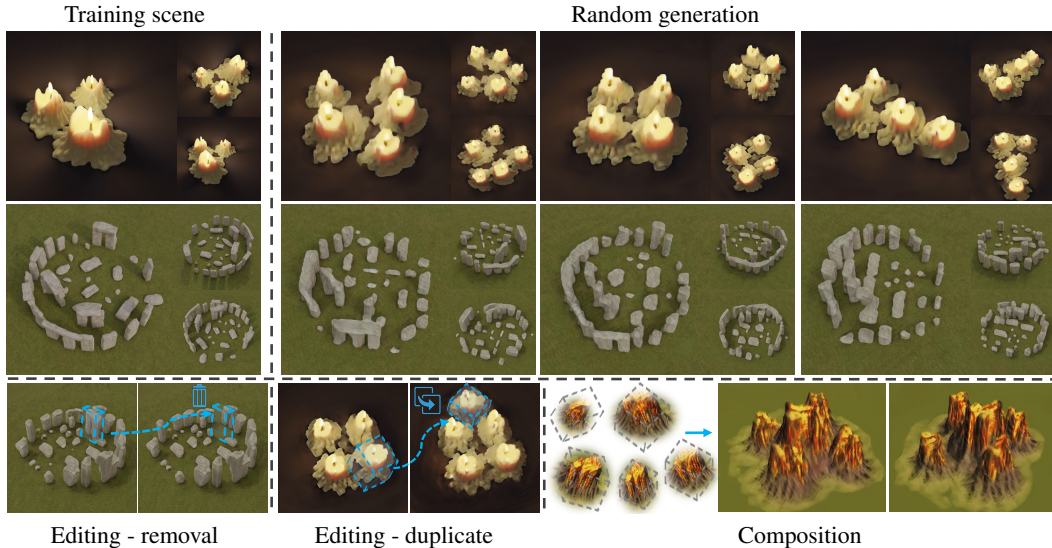


Figure 1: Top two rows: From observations of a single natural scene, we learn a generative model to synthesize highly plausible variations. Three views rendered from each of input/generated scenes are presented. Note how the global and object configurations vary in generated samples, yet still resemble the original. Bottom: Applications enabled by SinGRAV, including removal (left) and duplicate (middle) operation for editing a 3D scene sample, and scene composition (right) that combines 5 different generated samples to form a novel complex scene.

the training. As our supervision comes from 2D images, learning internal distributions consequently grounds SinGRAV on the assumption that multi-view observations share a consistent internal distribution for learning, which can be simply realized by uniformly placing the training cameras around the scene to obtain homogeneous images. On the other hand, MLP-based representations tend to synthesize holistically and perform better at modeling global patterns over local ones (Ding et al., 2022), thus favor modeling class-specific distributions. Hence, we resort to convolutional operations, which generate discrete radiance volumes from noise volumes with limited receptive fields, for learning local properties over a confined spatial extent, granting better out-of-distribution generation in terms of the global configuration. Moreover, we adopt a multi-scale architecture containing a pyramid of convolutional GANs to capture the internal distribution at various scales, alleviating the notorious mode-collapse issue. This framework is similar in spirit with Shaham et al. (2019); Shocher et al. (2019), however, important designs must be incorporated to *efficiently* and *effectively* improve the plausibility of the spatial arrangement and the appearance realism of the generated 3D scene.

We demonstrate SinGRAV enables us to easily generate plausible variations of the input scene in large quantities and varieties, with novel geometry, textures and configurations. Plausibility evaluations are conducted via perceptual studies, and comparisons are made to state-of-the-art generative models. The importance of various network design choices are validated. Finally, we show the versatility of SinGRAV by its use in a series of applications, spanning 3D scene editing, retargeting, and animation.

2 RELATED WORK

Neural scene representation and rendering. In recent years, neural scene representations have been the de facto infrastructure in several tasks, including representing shapes (Chen & Zhang, 2019; Park et al., 2019; Michalkiewicz et al., 2019; Takikawa et al., 2021; Martel et al., 2021), novel view synthesis (Nguyen-Phuoc et al., 2019; Mildenhall et al., 2020; Wiles et al., 2020), and 3D generative modeling (Nguyen-Phuoc et al., 2020; Schwarz et al., 2020; Chan et al., 2021a; Niemeyer & Geiger, 2021; Gu et al., 2022; DeVries et al., 2021; Chan et al., 2021b). Paired with differentiable projection functions, the geometry and appearance of the underlying scene can be optimized based on the error derived from the downstream supervision signals. Chen & Zhang (2019); Park et al. (2019); Xu et al. (2019); Takikawa et al. (2021); Michalkiewicz et al. (2019) adopt neural implicit fields to represent

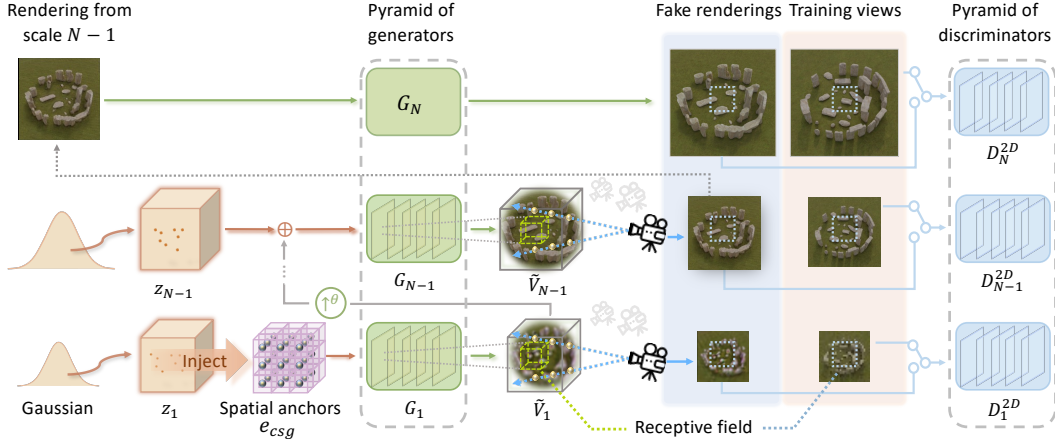


Figure 2: SinGRAV training setup. A series of convolutional generators are trained to generate a scene in a coarse-to-fine manner. At each scale, G_n learns to form a volume via generating realistic 3D overlapping patches, that collectively contribute to a volumetric-rendered imagery indistinguishable from the observation images of the input scene by the discriminator D_n . At the finest scale, the generator G_N operates purely on 2D domain to super-resolve the imagery produced from the previous scale (top), significantly reducing the computation overhead.

3D shapes and attain highly detailed geometries. On the other hand, Sitzmann et al. (2019); Thies et al. (2019); Wiles et al. (2020) work on discrete grids, UV maps, and point clouds, respectively, with attached learnable neural features that can produce pleasing novel view imagery. More recently, the Neural Radiance Field (NeRF) technique (Mildenhall et al., 2020) has revolutionized several research fields with a trained MLP-based radiance and opacity field, achieving unprecedented success in producing photo-realistic imagery. An explosion of NeRF techniques occurred in the research community since then that improves the NeRF in various aspects of the problem (Liu et al., 2020; Rebain et al., 2020; Zhang et al., 2020; Lindell et al., 2021; Wizarwongsa et al., 2021; Martin-Brualla et al., 2021; Lin et al., 2021; Wang et al., 2021; Lombardi et al., 2021). A research direction drawing increasing interest, which we discuss in the following, is to incorporate such neural representations to learn a 3D generative model possessing photo-realistic viewing effects.

Generative neural scene generation. With great success achieved in 2D image generation tasks, research have taken several routes to realize 3D-aware generative models. The heart of these methods is a 3D neural scene representation, paired with the volume rendering, makes the whole pipeline differentiable with respect to the supervision imposed in image domain. Schwarz et al. (2020); Chan et al. (2021a) integrate a neural radiance field into generative models, and directly produce the final images via volume rendering, yielding a 3D field modeled with great view-consistency under varying views. To overcome the low query efficiency and high memory cost issues, researchers proposed to adopt 2D neural renders to improve the inference speed, and achieved high-resolution outputs (Niemeyer & Geiger, 2021; Gu et al., 2022; Chan et al., 2021b). Niemeyer & Geiger (2021) utilizes multiple radiance sub-fields to represent a scene, and shows its potential to model scenes containing multiple objects on synthetic primitive datasets. More often than not, these methods are demonstrated on single-object scenes. DeVries et al. (2021) found that the capacity of a generative model conditioned on a global latent vector is limited, and instead propose to use a grid of locally conditioned radiance sub-fields to model more complicated scenes. All these work learn category-specific models, requiring training on sufficient volumes of images data collected from many homogeneous scenes. In this work, we target general natural scenes, which in general possess intricate and exclusive characteristic, suggesting difficulties in collecting necessary volumes of training data and rendering these data-consuming learning setups intractable. Moreover, as aforementioned, our task necessitates localizing the training over local regions, which is lacking in MLP-based representations, leading us to the use of voxel grids in this work.

3 METHOD

SinGRAV learns a powerful generative model for generating neural radiance volumes from multi-view observations $\mathcal{X} = \{x^1, \dots, x^m\}$ of a single scene. In contrast to learning class-specific priors, our specific emphasis is to learn the internal distribution of the input scene. To this end, we resort to convolutional networks, which inherently possess spatial locality bias, with limited receptive fields for learning over a variety of local regions within the input scene. The generative model is learned via adversarial generation and discrimination through 2D projections of the generated volumes. During training, the camera pose is randomly selected from the training set. We will omit the notion of the camera pose for brevity. Moreover, we use a multi-scale framework to learn properties at different scales, ranging from global configurations to local fine texture details. Figure 2 presents an overview.

3.1 NEURAL RADIANCE VOLUME AND RENDERING

The generated scene is represented by a discrete 3D voxel grid, and is to be produced by a 3D convolutional network. Each voxel center stores a 4-channel vector that contains a density scalar σ and a color vector \mathbf{c} . Trilinear interpolation is used to define a continuous radiance field in the volume. We use the same differentiable function as in NeRF for volume-rendering generated volumes. The expected color \hat{C} of a camera ray \mathbf{r} is approximated by integrating over M samples spreading along the ray: $\hat{C}(\mathbf{r}) = \sum_{i=1}^M T_i (1 - \exp(-\sigma_i \delta_i)) \mathbf{c}_i$, and $T_i = \exp\left(-\sum_{j=1}^{i-1} \sigma_j \delta_j\right)$, where the subscript denotes the sample index between the near t_n and far t_f bound, $\delta_i = t_{i+1} - t_i$ is the distance between two consecutive samples, and T_i is the accumulated transmittance at sample i , which is obtained via integrating over the preceding samples along the ray. This color rendering is reduced to traditional alpha compositing with alpha values $\alpha_i = 1 - \exp(-\sigma_i \delta_i)$. This function is differentiable and enables updating the volume based on the error derived from supervision signals.

3.2 HYBRID MULTI-SCALE ARCHITECTURE

We use the multi-scale framework widely followed in image generation. In addition, SinGRAV adopts a hybrid framework, which contains a series of 3D convolutional generators $\{G_n\}_{n=1}^{N-1}$ and a lightweight 2D convolutional generator G_N (see Figure 2). Specifically, 3D generators $\{G_n\}_{n=1}^{N-1}$ sequentially generate a volume with an increasing resolution at $N - 1$ coarser scales, with the volume resolution increased by a factor of θ between two consecutive scales and the rendering resolution increases by a factor of μ_r . At N th scale, to address the overly high computation cost, we use a lightweight 2D generator G_N to super-resolve the imagery from the preceding scale by a factor μ_s , achieving higher-resolution outputs. Essentially, each generator G_n learns to generate realistic outputs to fool an associate 2D discriminator D_n^{2D} , that is designed to distinguish the generated renderings from real ones. Importantly, the generators and discriminators are designed to be equipped with *limited receptive fields*, preventing them from over-fitting to the whole scene.

Starting from a coarsest volume, the generation sequentially passes through all generators and increases the resolution up to the finest, with noise injected at each scale. At the coarsest scale, the volume generation is purely generative, i.e. G_1 maps a Gaussian noise volume z_1 to a radiance volume. We observe that, compared to existing 3D category-level models, learning the internal distribution with spatial-invariant and receptive field-limited convolutional networks leads to more difficulties in producing plausible 3D structures. Following Xu et al. (2021), which alleviates a similar issue in image generation by introducing spatial inductive bias, we introduce spatial inductive bias into our pipeline by lifting the normalized Cartesian Spatial Grid (CSG) to 3D: $e_{csg}(x, y, z) = 2 \cdot [\frac{x}{W} - \frac{1}{2}, \frac{y}{H} - \frac{1}{2}, \frac{z}{U} - \frac{1}{2}]$, where W , H and U are size of the volume along the x -, y - and z -axis. As illustrated in Figure 3, the grid is equipped with distinct spatial anchors, empowering the model with better spatial localization. The spatial anchors provided by e_{csg} are injected into the noise volume z_1 at the coarsest level: $\tilde{V}_1 = G_1(z_1, e_{csg})$. Note we only inject the spatial inductive bias at the coarsest scale, as the positional-encoded information will be propagated through subsequent scales by convolution operations (Xu et al., 2021). The effective receptive field of the generator and discriminator at this




Figure 3: The spatial anchors provided by e_{csg} .

scale is around 40% of the volume of interest and the image size, thus G_1 learns to generate the overall layout and structure of the objects in the scene.

Subsequently, G_n at finer scales ($1 < n < N$) learns to add details missing from previous scales. Hence, each generator G_n takes as input a spatial noise volume z_n and an upsampled volume of $(\tilde{V}_{n-1})^\uparrow^\theta$ output from the preceding scale. Specially, prior to being fed into G_n , z_n is added to $(\tilde{V}_{n-1})^\uparrow^\theta$, and, akin to residual learning He et al. (2016), G_n only learns to generate missing details :

$$\tilde{V}_n = (\tilde{V}_{n-1})^\uparrow^\theta + G_n(z_n, (\tilde{V}_{n-1})^\uparrow^\theta), \quad 1 < n < N. \quad (1)$$

At the finest scale, a 2D convolutional generator G_N takes as input only the rendering \tilde{x}_{N-1} produced from the preceding scale, and outputs a super-resolved image \tilde{x}_N with enhanced details: $\tilde{x}_N = G_N(\tilde{x}_{N-1})$. G_N employs upsampling layers introduced in Gu et al. (2022), and produces the final image that twice the resolution of \tilde{x}_{N-1} . To better guarantee the view consistency, we adopt a joint discrimination strategy as in Chan et al. (2021b) for the discriminator D_N^{2D} . To achieve a progressive learning of the internal distributions at different scales, the receptive field of discriminators D_n^{2D} is limited so that the fineness of details in the generated scenes is gradually improved.

3.3 TRAINING LOSS

The multi-scale architecture is sequentially trained, from the coarsest to the finest scale. We construct a pyramid of resized input observations, $\mathcal{X}_n = \{x_n\}_{n=1}^N$, for providing supervisions. The GANs at coarser scales are frozen once trained. The training objective is as followed:

$$\min_{G_n} \max_{D_n^{2D}} \mathcal{L}_{adv}(G_n, D_n^{2D}) + \mathcal{L}_{rec}(G_n) + \mathbb{1}(n = N) \cdot \mathcal{L}_{sw}(G_n). \quad (2)$$

where \mathcal{L}_{adv} is an adversarial term, \mathcal{L}_{rec} is a reconstruction term as similar in Shaham et al. (2019), $\mathbb{1}(\cdot)$ is an indicator function to activate the associated term *iff* the condition is satisfied, \mathcal{L}_{sw} calculates the Sliced Wasserstein Distance (SWD) as in Heitz et al. (2021). SWD measures the distance between the textural distributions of two images, while neglecting the difference of the global layouts. Concretely, \mathcal{L}_{sw} is given by: $\mathcal{L}_{sw} = \mathcal{L}_{sw}(x_N, \tilde{x}_N) = \sum_{k=1}^K \mathcal{L}_{sw}(\tilde{f}^k, f^k)$, where \tilde{f}^k and f^k are features from layer k of a pre-trained VGG-19 network (Simonyan & Zisserman, 2015). Please refer to Heitz et al. (2021) for more details. In the following, we elaborate the designs of \mathcal{L}_{adv} and \mathcal{L}_{rec} .

Adversarial loss We use the WGAN-GP loss (Gulrajani et al., 2017) as \mathcal{L}_{adv} for stabilizing the training. The discrimination score is obtained by averaging over the patch discrimination map of D_n^{2D} . Similar to DeVries et al. (2021), we also condition the discriminators on the depth, of which the real samples d_n can be derived with multi-view geometry techniques trivially, to help improve the structural plausibility of the generation. During training, we render the depth from the generated volume, and concatenate the depth and color images for the input to $\{D_n^{2D}\}_{n=1}^{N-1}$. Note we do not use the depth discrimination for the 2D discriminator D_N^{2D} , and adopt the joint discriminator strategy as in Chan et al. (2021b) by concatenating the naively super-resolved $(\tilde{x}_{N-1})^\uparrow^{\mu_s}$ with the output from G_N . For preparing the real input to discriminator D_N^{2D} , we upsample the resized observation x_{N-1} via bilinear upsampling and concatenate the upsampled image with the ground truth observation x_N .

Reconstruction loss Inspired by Shaham et al. (2019), we introduce a specific set of input noise volumes to ensure that the noise within is able to reconstruct the underlying scene depicted in the observations \mathcal{X} . Specifically, a set of fixed noise volume is defined as $\{z_n^*\}_{n=1}^{N-1} = \{z_1^*, 0, \dots, 0\}$. The reconstructed radiance volumes and associated renderings are denoted as $\{V_n^*\}_{n=1}^{N-1}$ and $\{\tilde{x}_n^*, \tilde{d}_n^*\}_{n=1}^{N-1}$, respectively. Then the reconstruction loss \mathcal{L}_{rec} is defined as:

$$\mathcal{L}_{rec} = \lambda_c \|\tilde{x}_n^* - x_n\|_2^2 + \mathbb{1}(n < N) \cdot \lambda_d \|\tilde{d}_n^* - d_n\|_2^2, \quad (3)$$

where λ_c and λ_d are balance parameters and we set $\lambda_c = 10$ and $\lambda_d = 30$. As shown in Equation 3, we use supervisions on both color images and depth images for achieving higher quality. Note that the depth penalty term is removed at the last scale since G_N only works in the color image domain.

4 EXPERIMENTS

In the following, we explain experimental settings, and present evaluation results. More details and results can be found in the appendix A and the supplementary video.



Figure 4: Random scene generation. After training on multi-view observations of an input scene, SinGRAV learns to generate similar scenes with new objects and configurations. At each row, we show the input scene (left) and three diverse generated scenes (right) under the same two viewpoints.

Data. We collected observation images from a dozen of synthetic and real scenes, exhibiting ample variations over the global arrangements and constituents, for evaluating our method. For each scene, we assume the volume of interest is within a unit cube, and 200 observation images fully covering the scene are collected with cameras randomly distributed on a hemisphere. Random natural scene generation of our method is demonstrated upon all collected scenes, whereas more evaluations are conducted on a subset (*Stonehenge*, *grass and flowers*, and *island*).

Evaluation measures. We extend common metrics in single image generation to quantitatively assess $m = 50$ scenes generated from each input scene. For each generated scene, $k = 40$ images at random viewpoints are rendered for evaluation under a multi-view setting: a) *SIFID-MV* measures how well the model captures the internal statistics of the input by SIFID (Shaham et al., 2019) averaged over multi-view images of a generated scene; b) *Diversity-MV* measures the diversity of generated scenes by the averaged image diversity (Shaham et al., 2019) over multiple views.

Neural scene variations. Figure 1 and 4 presents qualitative results of SinGRAV, where the generated scenes depict reasonably new global layouts, and objects with various shapes. Note the observation images of the *beach castle* at the top row are taken from a real-world scene, and novel beach castles with various layouts are generated by SinGRAV, and realistic looking. These results suggest the efficacy of SinGRAV on modeling the internal patch distribution within the input scene. On *grass*




Figure 5: Qualitative results of baselines. All baselines encounter severe mode-collapse issues.

and flowers, which exhibits uniform yet complicated geometries and textures over an open field, SinGRAV produces high-quality random generation results. Moreover, SinGRAV is able to capture the global illumination to some extent, as evidenced by the shadows around stones and islands, along with the illumination changes under spinning cameras on samples of *mushroom*. The supplementary video presents more visual results under spinning cameras.

Comparisons. We compare to three state-of-the-art neural scene generative models, namely, GRAF Schwarz et al. (2020), pi-GAN Chan et al. (2021a), and GIRAFFE Niemeyer & Geiger (2021). We use their official codes for training these baselines. Training details of each baseline can be found in the supplementary. For fair comparisons, similar to SinGRAV, we use ground truth cameras when training baselines, instead of using random cameras from a predefined camera distribution. The quantitative and qualitative results are presented in Table 1 and Figure 5, respectively. Table 1 shows that pi-GAN produces the best *SIFID-MV* score, but the value of *Diversity-MV* drastically degrades. GRAF and GIRAFFE also exhibit a significantly degraded diversity score. The qualitative results in Figure 5 show that these baselines suffer from severe mode collapse, due to the lack of diverse scene samples for learning category-level priors.

Table 1: Quantitative comparisons. Top two on each metric are highlighted. All baselines suffer from severe mode collapse, producing low scores on Diversity-MV.

Method	Img. res.	SIFID-MV ↓	Diversity-MV ↑
GRAF	320 ²	0.4447	0.1337
pi-GAN	128 ²	0.0133	0.1157
GIRAFFE	320 ²	0.4710	0.3198
SinGRAV	320 ²	0.1113	0.7769

Table 2: Numerical results for variants of SinGRAV.

Variant	Img. res.	SIFID-MV ↓	Diversity-MV ↑
(wo. SWD)	320 ²	0.2713	0.7730
(wo. CSG)	320 ²	0.1307	0.8434
(wo. depth sup.)	320 ²	0.1157	0.7910
(NeRF-depth)	320 ²	0.1290	0.8046
(self-depth)	320 ²	0.1120	0.7862
(MLP)	160 ²	0.1843	0.5235
(MLP-LLG)	108 ²	0.2013	0.4621
SinGRAV	320 ²	0.1113	0.7769

Validation of design choices. We conduct experiments to evaluate several key design choices:

SWD loss. If \mathcal{L}_{swd} is eliminated, the internal distribution of generated scenes would differ greatly from that of the input, leading to significantly increased *SIFID-MV* at the first row in Table 2. Correspondingly, the visual results in Figure 6 also show that SinGRAV (wo. \mathcal{L}_{swd}) produces blurry and less realistic textures, suggesting the efficacy of \mathcal{L}_{swd} in improving the visual quality.

Spatial inductive bias. We build a variant – SinGRAV (wo.CSG), which is trained without the spatial anchor volume e_{csg} . As shown in Table 2, compared to SinGRAV, SinGRAV (wo.CSG) produces a worse *SIFID-MV* score (increases by 17%) and an increased *Diversity-MV* score. While the latter suggests increased diversity, we observe from the visual results that the spatial arrangements of objects deviate significantly from that of the input, producing floating stones, as shown in the middle of Figure 7.

Depth supervision. The third row in 2 presents numerical results of a variant SinGRAV (wo. depth sup.), which is trained without depth supervision. SinGRAV (wo. depth sup.) achieves comparable numerical performance to SinGRAV. But, the extracted point clouds in Figure 7 reveal implausible geometric structures in the generated scenes of SinGRAV (wo. depth sup.), showing the importance of depth supervision. On the other hand, we also show that our framework can work with depth data

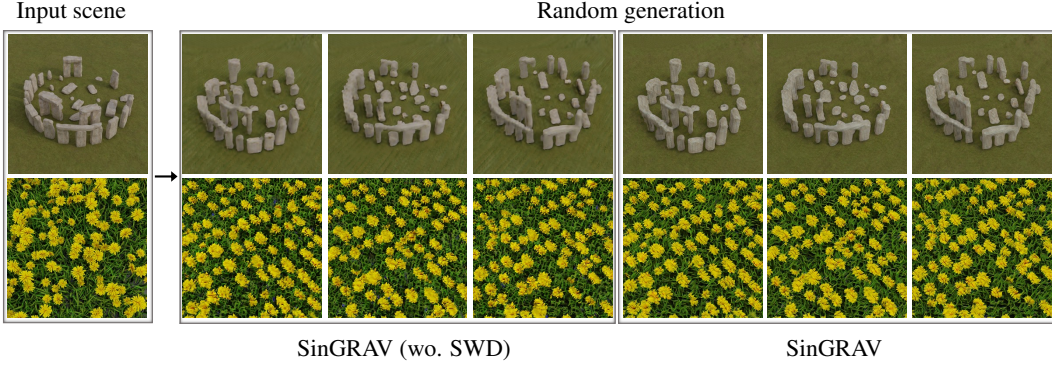


Figure 6: SinGRAV vs. SinGRAV (wo. SWD). Left: Without the SWD loss at the finest scale, SinGRAV (wo. SWD) produces blurry textures and undesired artifacts. Right: Training with SWD loss significantly improves the texture quality.

obtained from various sources, in addition to rather clean depth map in our default setting. Specifically, we designed SinGRAV (NeRF-depth), which uses the depth obtained from a NeRF Mildenhall et al. (2020) reconstruction of the input scene, and SinGRAV (self-depth), which uses the depth obtained from the reconstructed volume with z^* and $\lambda_d = 0$ in our framework. Table 2 shows that SinGRAV (NeRF-depth) and SinGRAV (self-depth) are able to achieve comparable performance to SinGRAV, implying that SinGRAV is robust to the quality of the depth data. That said, all these depth data of different quality provide sufficient supervision for learning global geometric structures.

MLP vs. voxel. The numerical results of SinGRAV (w. MLP), which uses a conditional MLP-based radiance field as in GRAF, are reported in Table 2. The highest image resolution is 160×160 due to overly high memory consumption. SinGRAV (w. MLP) degrades in both the quality and diversity, which is also reflected by the visuals in Figure 8. Note that, although the same patch discrimination strategy is used, SinGRAV (w. MLP) suffers from severe mode collapse. We believe this is because the generative output of coordinate-based MLPs is very likely to be dominated by the input coordinates. The poor generation is also possibly a product of the conflict between the lack of locality in the fully connected layers Ding et al. (2022) and the limited receptive field in the adversarial training. In addition, we train a variant SinGRAV (w. MLP-LLG), which incorporates a local latent grid proposed in DeVries et al. (2021) to increase the capacity of the MLP-based representation. Figure 8 shows that the generated layouts are improved with the increased locality, however, severe mode collapse still exists. On the other hand, we believe more effort can be made in the future to adapt MLP-based representations for our task.

Influence of varying pyramid depth. We train variants with various numbers of scales. Specifically, for training a variant with t ($t < N$) scales, we use generators $\{G_n\}_{n=N-t+1}^N$ and discriminators $\{D_n^{2D}\}_{n=N-t+1}^N$ to preserve the final image resolution. As shown in Figure 9, with less scales, the




Figure 7: Influence of spatial inductive bias and depth supervision. One generated sample with the corresponding point cloud is shown for each variant. Without the inductive bias (middle) or depth supervision (right), the generated scenes exhibit implausible geometric structures, while SinGRAV (left) preserves the spatial arrangements well.

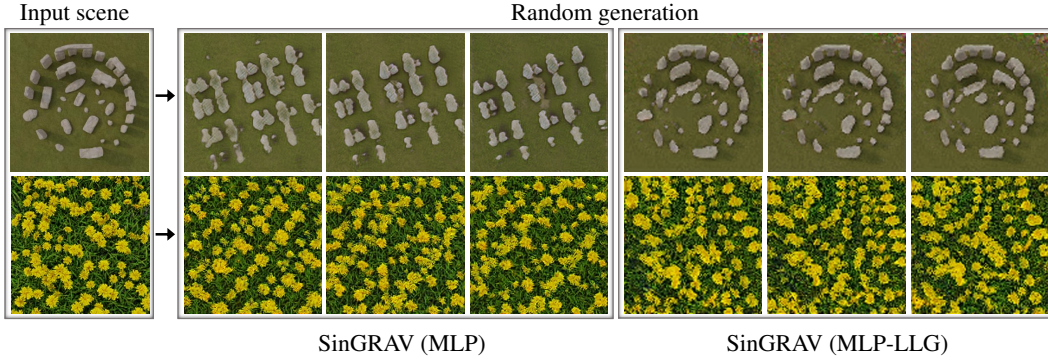


Figure 8: Qualitative results of using MLP-based representations. SinGRAV (MLP) with globally conditioned MLPs produces grid patterns, which are alleviated by the use of local latent grid in SinGRAV (MLP-LLG). Nevertheless, both variants suffer from severe mode collapse, generating almost identical scene samples.

effective receptive field at the coarsest scale is rather small, resulting in a model that only captures local properties, whereas, more scales allows modeling plausible global arrangements.

Applications. We show the application potential of SinGRAV in common 3D modeling process. Specifically, we derive three applications with SinGRAV, including 3D scene editing with removal and duplicate operations, composition, and animation. Figure 1 presents the results and more can be found in the supplementary video.



Figure 9: Influence of different number of scales. Training with more scales is beneficial to the modeling of global arrangements, while a model with less scales tends to capture only local textures.

5 CONCLUSION

In this work, we make a first attempt to learn a deep generative neural scene model from visual observations of a single scene. Once trained on a single scene, the model can generate novel scenes with plausible geometries and arrangements, that can be rendered with pleasing viewing effects. The importance of key design choices are validated. Despite successful demonstrations, our technique has a few limitations of its current designs.

While our model learns from a single scene, bypassing the need for collecting data from many homogeneous 3D samples, a rather complete 3D scene is yet required to obtain the multi-view data for training. Moreover, albeit validated, the use of voxel grids inherently limits the network ability in modeling fine details, consequently hindering the model from achieving high-resolution imagery. Our remedy is to incorporate a 2D neural renderer that operates on 2D domain to super-resolve the imagery, which inevitably introduces the multi-view inconsistency. There are view inconsistencies on complex textural areas, such as the thin structures in the Grass and Flowers scene. A future direction would be to overturn this design, with more endeavor on exploiting MLP-based representations to model continuous volumes. Besides, the proposed method, in its current form, sometimes produces artifacts in the background, as it does not incorporate special designs for modeling the background. Hence, it would also be worth addressing such an issue, especially for scenes with complicated background, potentially with designs inspired from other generative neural scene models.

REFERENCES

- 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)*, 2021a.
- Eric R. Chan, Connor Z. Lin, Matthew A. Chan, Koki Nagano, Boxiao Pan, Shalini De Mello, Orazio Gallo, Leonidas Guibas, Jonathan Tremblay, Sameh Khamis, Tero Karras, and Gordon Wetzstein. Efficient geometry-aware 3D generative adversarial networks. *arXiv:2112.07945*, 2021b.
- Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *CVPR*, 2019.
- 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)*, pp. 14304–14313, 2021.
- Xiaohan Ding, Honghao Chen, Xiangyu Zhang, Jungong Han, and Guiguang Ding. Repmlpnet: Hierarchical vision mlp with re-parameterized locality. In *CVPR*, 2022.
- Jiatao Gu, Lingjie Liu, Peng Wang, and Christian Theobalt. Stylenerf: A style-based 3d-aware generator for high-resolution image synthesis. In *ICLR*, 2022.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans. 2017.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
- Eric Heitz, Kenneth Vanhoey, Thomas Chambon, and Laurent Belcour. A sliced wasserstein loss for neural texture synthesis. In *CPVR*, 2021.
- Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural radiance fields. *arXiv preprint arXiv:2104.06405*, 2021.
- David B. Lindell, Julien N.P. Martel, and Gordon Wetzstein. Autoint: Automatic integration for fast neural volume rendering. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Stephen Lombardi, Tomas Simon, Gabriel Schwartz, Michael Zollhoefer, Yaser Sheikh, and Jason Saragih. Mixture of volumetric primitives for efficient neural rendering. *arXiv preprint arXiv:2103.01954*, 2021.
- Julien N. P. Martel, David B. Lindell, Connor Z. Lin, Eric R. Chan, Marco Monteiro, and Gordon Wetzstein. Acorn: Adaptive coordinate networks for neural scene representation. *TOG*, 40(4), 2021.
- Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections. In *CVPR*, 2021.
- Mateusz Michalkiewicz, Jhony K Pontes, Dominic Jack, Mahsa Baktashmotlagh, and Anders Eriksson. Deep level sets: Implicit surface representations for 3d shape inference. In *ICCV*, 2019.
- 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.
- Thu Nguyen-Phuoc, Chuan Li, Lucas Theis, Christian Richardt, and Yong-Liang Yang. Hologan: Unsupervised learning of 3d representations from natural images. In *ICCV*, 2019.

- Thu Nguyen-Phuoc, Christian Richardt, Long Mai, Yong-Liang Yang, and Niloy Mitra. Blockgan: Learning 3d object-aware scene representations from unlabelled images. In *NIPS*, 2020.
- Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019.
- Daniel Rebain, Wei Jiang, Soroosh Yazdani, Ke Li, Kwang Moo Yi, and Andrea Tagliasacchi. Derf: Decomposed radiance fields. *arXiv preprint arXiv:2011.12490*, 2020.
- 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.
- 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.
- Assaf Shocher, Shai Bagon, Phillip Isola, and Michal Irani. Ingan: Capturing and retargeting the "dna" of a natural image. In *International Conference on Computer Vision (ICCV)*, pp. 4492–4501, 2019.
- K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- Vincent Sitzmann, Justus Thies, Felix Heide, Matthias Nießner, Gordon Wetzstein, and Michael Zollhöfer. Deepvoxels: Learning persistent 3d feature embeddings. In *CVPR*, 2019.
- Towaki Takikawa, Joey Litalien, Kangxue Yin, Karsten Kreis, Charles Loop, Derek Nowrouzezahrai, Alec Jacobson, Morgan McGuire, and Sanja Fidler. Neural geometric level of detail: Real-time rendering with implicit 3D shapes. 2021.
- Justus Thies, Michael Zollhöfer, and Matthias Nießner. Deferred neural rendering: Image synthesis using neural textures. *TOG*, 2019.
- Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. Nerf-: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064*, 2021.
- Olivia Wiles, Georgia Gkioxari, Richard Szeliski, and Justin Johnson. SynSin: End-to-end view synthesis from a single image. In *CVPR*, 2020.
- Suttisak Wizadwongsa, Pakkapon Phongthawee, Jiraphon Yenphraphai, and Supasorn Suwajanakorn. Nex: Real-time view synthesis with neural basis expansion. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. volume 32, 2019.
- 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)*, pp. 13569–13578, 2021.
- Kai Zhang, Gernot Riegler, Noah Snaveley, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. *arXiv:2010.07492*, 2020.

A APPENDICES

In this section, we provide additional implementation details in Section A.1, additional experimental results in Section A.2 and list outlines in Section A.3 for video demonstrations provided in the supplementary videos.

A.1 IMPLEMENTATION DETAILS

We describe more implementation details including detailed network structures for SinGRAV in Section A.1.1, details for data creation in Section A.1.2, details for evaluation measures in Section A.1.3 and details for demonstrated applications in A.1.4. In addition, we provide the detailed training settings for SinGRAV in Section A.1.5 and training details for three baselines in Section A.1.6.

A.1.1 MORE DETAILS FOR NETWORK DESIGN

Overview All 3D generators $\{G_n\}_{n=1}^{N-1}$ share the same structure, with the same number of 3D convolutional layers (kernel size 3 and stride 1). The number of layers is set to 5 (small receptive field for small objects in the scene), or 7 (large receptive field for large objects), according to different scene constituents. All 2D discriminators $\{D_n^{2D}\}_{n=1}^N$ have the same number (5 or 7) of 2D convolutional layers (kernel size 3 and stride 1). The resolution of the volume and the image start from $40 \times 40 \times 40$ and $32 \times 32 \times 32$, and increase by a factor of $\theta = 4/3$ and $\mu_r = 1.5$, respectively. We set $N = 6$ by default, and at scale $N - 1$, the volume and image resolution reach 126^3 and 160^2 separately. The above setting is used on all data, except the *grass and flows*, and *island* that possess rather fine-grained texture details.

Detailed structure for $\{G_n\}_{n=1}^{N-1}$ For generators ranging from the coarsest scale to scale $N - 1$, we adopt the same structure. Specifically, for generators $\{G_n\}_{n=1}^{N-1}$, we adopt an L -layer structure, which is designed with 3D convolutional layers, Batch Normalization and LeakyReLU activation. The detailed structure of $\{G_n\}_{n=1}^{N-1}$ is summarized in Table 3. As shown in Table 3, the number of input channels, IC_n^V , is adjusted according to the difference of the input volume for different scales. At the coarsest level, to make the volume adapt to the spatial anchor volume e_{csg} , we set $IC_1^V = 3$. For the remaining 3D generators, we set $IC_n^V = 4$. In addition, for generators whose scale is in range $[2, N - 1]$, a residual connection is set up between the output volume and the upsampled \tilde{V}_{n-1} , *i.e.*, the generators $\{G_n\}_{n=2}^{N-1}$ learn to add more details to the synthesized scene.

Table 3: Detailed structure of the generators $\{G_n\}_{n=1}^{N-1}$. Specifically, “conv3d”, “batchNorm3d” represent the 3D convolution layer and batch normalization layer respectively. IC_n^V denotes the number of channels within the input volume. We set $IC_n^V = 3$ for $n = 1$ and $IC_n^V = 4$ for $n > 1$. L is set to 5 or 7.

Index	Name	Kernel	Stride	Pad	Input ch.	Output ch.
1	Conv3d + BatchNorm3d + LeakyReLU (0.2)	$3 \times 3 \times 3$	1	1	IC_n^V	32
2	Conv3d + BatchNorm3d + LeakyReLU (0.2)	$3 \times 3 \times 3$	1	1	32	32
$[3, L - 1]$	Conv3d + BatchNorm3d + LeakyReLU (0.2)	$3 \times 3 \times 3$	1	1	32	32
L	Conv3d	$3 \times 3 \times 3$	1	1	32	4

Structure of G_N At the finest level, we adopt a 2D neural generator G_N to directly lift the rendered images to the output resolution 320×320 . We adopt an upsampling layer proposed in Gu et al. (2022) at the first layer, after which we design 4 layers with 2D convolutional layers and Instance Normalization. The detailed structure is given in Table 4. Besides, we utilize a residual connection between the final output and the resized input image, *i.e.*, $\tilde{x}_{N-1} \uparrow$.

Detailed structures for and $\{D_n^{2D}\}_{n=1}^N$ We design $\{D_n^{2D}\}_{n=1}^N$ in a similar manner to the design of $\{G_n\}_{n=1}^{N-1}$, *i.e.*, all of $\{D_n^{2D}\}_{n=1}^N$ are designed with the same structure implemented with 2D convolutional layers and LeakyReLU() activation. In Table 5, we provide the detailed information of the structure. The discriminators have the same number of layers as the generators $\{G_n\}_{n=1}^{N-1}$. For the first layer of the generators $\{D_n^{2D}\}_{n=1}^{N-1}$, we set $IC_1^I = 4$ since their input is formed by concatenating the rendered color image \tilde{x}_n and depth map \tilde{d}_n . For the discriminator D_N^{2D} at the finest scale, we set $IC_N^I = 6$ as its input is the concatenation of \tilde{x}_N and $\tilde{x}_{N-1} \uparrow$, which is the output from G_N and the super-resolved image from the preceding scale respectively.

Table 4: Detailed structure of the generators G_N . Specifically, "conv2d" represent the 2D convolution layer. "InsNorm2d" denotes Instance Normalization.

Index	Name	Kernel	Stride	Pad	I/O	Input res	Output res
1	Upsampling layer	-	-	-	3/3	160×160	320×320
2	Conv2d + InsNorm2d + LeakyReLU (0.2)	3×3	1	1	32/16	320×320	320×320
3	Conv2d + InsNorm2d + LeakyReLU (0.2)	3×3	1	1	16/8	320×320	320×320
4	Conv2d + InsNorm2d + LeakyReLU (0.2)	3×3	1	1	8/8	320×320	320×320
5	Conv2d	3×3	1	1	8/3	320×320	320×320

Table 5: Detailed structure of the generators $\{G_n\}_{n=1}^{N-1}$. Specifically, "conv2d" represent the 2D convolution layer. IC_n^I denotes the number of channels within the input for D_n^{2D} . We set $IC_n^I = 4$ for $n < N$ and $IC_n^V = 6$ for $n = N$. L is set to 5 or 7.

Index	Name	Kernel	Stride	Pad	Input ch.	Output ch.
1	Conv2d + LeakyReLU (0.2)	3×3	1	0	IC_n^I	32
2	Conv2d + LeakyReLU (0.2)	3×3	1	0	32	32
[3, $L - 1$]	Conv2d + LeakyReLU (0.2)	3×3	1	0	32	32
L	Conv2d	3×3	1	0	32	1

A.1.2 ADDITIONAL DETAILS FOR DATA CREATION

The 3D models are collected from this website¹, under TurboSquid 3D Model License². During the process of data creation, we scale the scenes so that the volume of interest of them stay within a cube with side width=2 (within the range $[-1, 1]$). To obtain the data for training our model, on synthetic input scenes, we utilize path tracing renderer in Blender to get the multi-view RGB-D observation. As for real-world scenes, the visual images are collected from the internet (Google Earth and Sketchfab), we then utilize COLMAP to reconstruct high-quality depth and camera poses, which is a common practice in many neural scene representation works. Specifically, for *Stonehenge, grass and flowers, island, mushroom and water lily*, we use a camera with FOV=33.40° (focal length=76mm). For *crystal, candies and volcano*, we use a camera with FOV=49.13° (focal length=50mm). The viewpoints are sampled randomly on a hemisphere with a radius of 3.5.

A.1.3 MORE DETAILS FOR EVALUATION METRICS

In this section, we provide more details for calculating the scores for the evaluation metrics including *SIFID-MV* and *Diversity-MV*. Specifically, for each dataset, we randomly sample $M = 40$ pairs of color image x_m and camera pose ps_m , and m denotes the index running over the sampled viewpoints. Then, we generate $J = 50$ scenes, which are rendered under the sampled viewpoints to get 40 rendered images \tilde{x}_m^j per scene. Note that the notions for scales are omit here for brevity since we only use the rendered image at the finest scale for evaluation. Herein, we calculate the score for *SIFID-MV* according to

$$\text{SIFID-MV} = \frac{1}{M} \frac{1}{J} \sum_{m=1}^M \sum_{j=1}^J \text{SIFID}(\tilde{x}_m^j, x_m), \quad (4)$$

where $\text{SIFID}(\cdot, \cdot)$ denotes the SIFID metric proposed in Shaham et al. (2019), which is designed for measuring the distance between the internal distributions of two single images.

And we calculate the score for *Diversity-MV* according to

$$\text{Diversity-MV} = \frac{1}{M} \sum_{m=1}^M \frac{\text{std}(\{\tilde{x}_m^j\}_{j=1}^J)}{\text{std}(x_m)}, \quad (5)$$

¹<https://www.turbosquid.com>

²<https://blog.turbosquid.com/turbosquid-3d-model-license/#3d-model-license>

where $\text{std}(\cdot)$ denotes calculating the standard deviation of each pixel over all of the rendered images under the same viewpoints. As shown in Equation (5), before calculating the averages over all of the viewpoints, we normalize the score under single viewpoint with the standard deviation over all of the pixels within the corresponding input observation.

A.1.4 IMPLEMENTATION DETAILS OF APPLICATIONS

In this section, we present the implementation details of the demonstrated applications, including *Scene Animation* and *Scene Editing*. Since the generator G_N at the finest scale only operates in 2D rendering domain and functions mainly as a faithful super-resolving module, we omit it for brevity in the descriptions below.

Scene Animation. To plausibly animate a generated 3D scene for T consecutive time steps, we need to construct a sequence of radiance volumes $\{\tilde{V}_{N-1}^{(t)}\}_{t=1}^T$ with similar layouts but slightly changing content. In our implementation, we achieve this by utilizing the generators to produce $\{\tilde{V}_{N-1}^{(t)}\}_{t=1}^T$ from the modified versions of the pyramid of noise volumes $\{z_n\}_{n=1}^{n=N-1}$, which is used to generate the original scene \tilde{V}_{N-1} . Specifically, we inject random changes smoothly for T time steps $\{t\}_{t=1}^{T-1}$, and restrict the resulting noise volumes to stay close to the original ones $\{z_i\}_{i=1}^{n=N-1}$. Formally, the way we construct the random walk is given by:

$$\begin{aligned} z_n^{(t+1)} &= \alpha z_n^{(1)} + (1 - \alpha)(z_n^{(t)} + \delta_n^{(t+1)}), \\ \delta_n^{(t+1)} &= \xi(z_n^{(t)} - z_n^{(t-1)}) + (1 - \xi)\mu_n^{(t+1)}, \end{aligned} \quad (6)$$

where n denotes the index running over the pyramid of the noise volumes. $z_n^{(1)}$ is the noise volume at the first step and is set as the original noise volume at scale n , i.e., $z_n^{(1)} = z_n$. $\delta_n^{(t+1)}$ is the injected change at time step $t + 1$, and $\mu_n^{(t)}$ is a noise volume that has the same distribution as z_n . α and ξ are two parameters that can be tuned to achieve different animation effects. Concretely, α controls the degree of preserving the original scene, i.e., the similarity between the original scene \tilde{V}_{N-1} and the animated scene at time step $t + 1$. ξ controls the smoothness between the adjacent frames and the rate of introduced changes in the generated frame. Note that the noise volumes at the coarser scales are able to control the global layout of the generated scenes, on which slight modifications might drastically alter the scene. Thus we choose to inject the random changes in finer scales while keeping the noise volumes at the coarser scales unchanged for each time step. For obtaining the demonstrated animation effects for *candles* and *crystal*, we let the random change injection start from the third scale and set $\alpha = 0.58$ and $\xi = 0.45$.

Scene Editing. For editing a scene, we first need to localize the target area, then simply manipulate the volume \tilde{V}_{N-1} , and finally downscale the edited volume and let the volume go through the pyramid of generators to harmonize the edited scene. We provide detailed description of these steps below.

Step1: object area localization. For finding an area of interest, we first extract the mesh from a scene volume \tilde{V}_{N-1} , and interactively locate the bounding box of the area of interest in a 3D mesh visualizer, e.g., Blender in our prototype implementation. Then we construct a 3D mask \mathcal{M} , where the area within the bounding box is filled with 1 and the rest is set to 0. Besides, we also sample a point within the area representing 'air' and use its color and density as \mathbf{c}_{empty} and σ_{empty} .

Step2: volume manipulation.

For the application *Removal*, we manipulate the volume via

$$\tilde{V}_{N-1}(\mathcal{M} = 1) = (\mathbf{c}_{empty}, \sigma_{empty}). \quad (7)$$

For the application *Duplicate*, we select an area with empty voxels within the volume, which is masked by \mathcal{M}_{empty} . Then we set the empty area to the values extracted from the object area covered by \mathcal{M} . That is, we manipulate the volume via

$$\tilde{V}_{N-1}(\mathcal{M}_{empty} = 1) = \tilde{V}_{N-1}(\mathcal{M} = 1). \quad (8)$$

For the application *Move*, we achieve the effects by first conducting the *Duplicate* operation and then performing the *Removal* operation to \tilde{V}_{N-1} .

For the application *Composition*, we select the object of interest in k generated scenes and construct a group of masks $\{\mathcal{M}_i\}_{i=1}^k$. Then we set k target bounding boxes within an extra scene, which are destinations where we want to inject the new objects in. Consequently, we get k target masks $\{\mathcal{M}_i^{dest}\}_{i=1}^k$. Then we combine the objects into a single scene by putting the extracted sub-volume covered by \mathcal{M}_i for each object in corresponding area \mathcal{M}_i^{dest} .

Step3: neural editing. After each kind of volume manipulation mentioned above, we downscale it to the size of \tilde{V}_3 and then let $\{G_n\}_{n=3}^{N-1}$ to generate the final volume \tilde{V}_{N-1} . This step will smooth the discontinuities caused by the naive volume manipulations. Then we can generate the renderings from the newly generated scene.

A.1.5 TRAINING DETAILS FOR SINGRAV

Weight initialization. For the networks at the coarsest level, $\{G_1, D_1\}$, we adopt a random initialization. For generator G_n and discriminator D_n ($n > 1$), we initialize the weights using the weights from trained networks of the preceding scale, *i.e.*, weights of G_{n-1} and D_{n-1} . We empirically find that using the weights from the pre-trained model of the preceding scale is helpful to stabilize the training.

Parameter settings. The weight of gradient penalty term in the adversarial loss is set to 0.1. We adopt the Adam optimizer for optimizing both the generator and discriminator, for both of which the learning rate is set to 0.0005 and the momentum parameters are set as $\beta_1 = 0.9, \beta_2 = 0.999$.

Training. We sequentially train the networks at different scales in a coarse-to-fine manner and the networks at coarser levels are fixed after being trained. Fixing the networks at the coarser scales is helpful to stabilize the training process of the network at finer scales, since the results at the beginning of the training at each scale is messy. At each scale s , we train the networks for 80 epochs. Specifically, we first train G_n using the reconstruction loss defined in Equation (3) for 20 epochs and then train $\{G_n, D_n^{2D}\}$ together using the loss defined in Equation (2). In each iteration, we alternate between 3 gradient steps for the discriminator D_n^{2D} and 3 steps for the generator G_n .

Time and platform. We train SinGRAV with 4 NVIDIA Tesla V100 GPUs (32GB memory) for around 2 days. For generating a new scene, the inference time is 0.32 seconds and the memory consumption is about 5 GB.

Parameter settings for different data In Table 6, we list the detailed parameter settings used for training SinGRAV on the collected datasets. For most of the data, the settings shown in the first row of Table 6 are applicable. For training on *island*, we slightly adjust the number of layers in $\{G_n\}_{n=1}^{N-1}$ and $\{D_n^{2D}\}_{n=1}^N$ to make them have a smaller receptive field, which is helpful to learn the fine textures existed in *island* dataset. For training on *water lily*, we slightly increase the resolution of \tilde{V}_1 considering the complicated geometries of the water lily. Besides, with the gradually increased resolution of volumes and rendered images, we adjust the batch size at each scale in order to accommodate the training to the memory budget of the hardwares. Note that the parameter tweaking for *island* and *water lily* does not bring significant gains, instead, it just slightly improve the results.

Table 6: Training parameters for SinGRAV on different datasets. L denotes the number of layers in G_n or D_n^{2D} .

Datasets	L	Res. of $\tilde{V}_1 / \tilde{x}_1$	Res. of $\tilde{V}_{N-1} / \tilde{x}_{N-1}$	Final image Res.	batch size for \tilde{x}_n at 6 scales	batch size for \tilde{x}_n^* at 6 scales
<i>Stonehenge, grass and flowers, mushroom, crystal, candies, volcano</i>	7	40/32	126/160	320	[6, 6, 6, 5, 2, 2]	[2, 2, 2, 1, 1, 1]
<i>Island</i>	5	40/25	126/160	320	[6, 6, 6, 6, 3, 2]	[2, 2, 2, 1, 1, 1]
<i>Water lily</i>	7	53/32	156/160	320	[6, 6, 4, 2, 1, 1]	[2, 2, 2, 1, 1, 1]

Sampling strategy for volume rendering. In the rendering process, we adopt a uniform sampling strategy, *i.e.*, we uniformly sample a number of points within the range of the *near* point and *far* point along each ray cast from each pixel. Considering the increased fineness of textures within the generated scene at different scales, we progressively enlarge the number of sampling. Specifically, the number of sampling points start from 64 at scale 1 and reaches 128 at the finest scale.

A.1.6 DETAILED TRAINING SETTINGS FOR BASELINES

For training the baseline methods, we keep most of the parameters the same as those provided in the exemplar configurations. The parameters we modified according to our settings are listed below:

GRAF Schwarz et al. (2020): On each dataset, we train more than 4500 epochs until convergence. The patch size used during training is set to 32. Due to the patch based training process, it supports to be trained at 320×320 resolution and convergences slower.

pi-GAN Chan et al. (2021a): We train 685 epochs on each dataset and adopt the progressive training strategy as in Chan et al. (2021a). Specifically, we train 150 epochs with batch size = 30 at image resolution 32×32 , 270 epochs with batch size = 12 at image resolution 64×64 and 265 epochs with batch size = 8 at image resolution 128×128 .

GIRAFFE Niemeyer & Geiger (2021): We set the parameters according to the instructions provided in the official implementation³. Specifically, we set the number of object field to 1 and use a background field. For training on our datasets, we separately set the parameters for the foreground field in the following manner:

- *Stonehenge*: The scaling range is set to $[0.45, 0.55]$ and the translation range is set to $[0.0, 0.0]$ for x-axis, y-axis and z-axis.
- *Grass and flowers*: The scaling range is set to $[0.7, 0.8]$ and the translation range is set to $[0.0, 0.0]$ for x-axis, y-axis and z-axis.
- *island*: The scaling range is set to $[0.45, 0.55]$ for x-axis, y-axis and z-axis. And the translation range for x-/y- axis is set to $[-0.15, 0.15]$, while the translation range for z-axis is set to $[0.0, 0.0]$.

Besides, for training with GIRAFFE Niemeyer & Geiger (2021) method, we let the low-res image output from the volume rendering have a resolution of 20×20 , so that the resolution of the super-resolved image reaches 320×320 . We train the framework GIRAFFE for more than 1.3 million iterations.

A.2 ADDITIONAL EXPERIMENTAL RESULTS

In this section, we present the results on more complicated real-world scenes in A.2.1 and investigate the robustness of the framework to the underlying camera pose distribution for the multi-view observations used for training in A.2.1.

A.2.1 RESULTS ON REAL-WORLD OUTDOOR SCENES

In this section, we provide the results for SinGRAV on two more complicated real scenes, which are collected from Google Earth⁴, namely *Devil Tower* and *Bagana Volcano*, respectively. After training on the collected data, SinGRAV achieves quantitative results shown in Table 7.

Table 7: Numerical results on real-world data.

Variant	Data	SIFID-MV ↓	Diversity-MV ↑
SinGRAV	Devil Tower	0.1940	0.5801
SinGRAV	Bagana Volcano	0.1480	0.7359
SinGRAV (wo. depth sup.)	Bagana Volcano	0.1710	0.8391

³<https://github.com/autonomousvision/giraffe>

⁴<https://earth.google.com>




Figure 10: Random scene samples from SinGRAV trained on real-world data. At each row, we show two views of the input scene on the left, followed by three diverse scenes generated by SinGRAV on the right. Each scene is rendered with the same two cameras.

The first row and the second row in Table 7 show that the numerical results on real-world scenes approaches those demonstrated in Table 2. This suggests that SinGRAV also produces reasonable results on more complex real-world scenes. On the qualitative perspective, we notice that, while our method produces plausible generation visual results, the fine textures of the real-world scene indeed challenge the proposed voxel-based framework as also discussed in the main paper. Note these real-world results are obtained using the default hyper-parameter setting, without ad hoc fine-tuning.

In addition, we also investigate the influence of the depth supervision on one of the above scenes – *Bagana Volcano*, and report the numerical results in the third row of Table 7. It can be seen that the numerical results drop slightly, which is similar to the trend shown in Table 2. Again, despite the slight drop in numerical results, totally eliminating depth supervisions on this real-world data also drastically affects the plausibility of the spatial arrangements in the results, which can be observed from the qualitative results shown in Figure 11. Nevertheless, the depth data obtained with COLMAP in our default setting is sufficient for producing reasonable geometric structures in the results, as also stated in the main paper.




Figure 11: Influence of the depth supervisions on the real-world scene. Two generated scenes are provided for SinGRAV and three results are provided for SinGRAV (wo. depth sup.). Red boxes highlight the implausible geometry – holes in the results of SinGRAV (wo. depth sup.), that are caused by the lack of depth supervision.

A.2.2 ROBUSTNESS TO CAMERA DISTRIBUTION

In real-world scenarios, it is likely that the cameras would deviate from the planned trajectories to some extent, i.e., the hemisphere camera distribution may not be guaranteed under a less controlled setting. Nevertheless, we shall show that SinGRAV does NOT necessitate an absolutely strict distribution for the training cameras, and can work properly if two criteria can be met: a) *high observation completeness*, which suggests the scene of interest needs to be covered as much as possible, and is a reasonable assumption with the prevalent "render-and-discriminate" framework; b) *consistent multi-scale internal distribution*, which means the captured images need to share a similar internal distribution at each scale for learning the generative model.

To account for such factors, we investigate the robustness of SinGRAV under such less controlled environments by simulating with perturbed camera poses. Specifically, we randomly add noise drawn from a Gaussian distribution with different standard deviations to the camera poses during




Figure 12: Performance for SinGRAV trained on data obtained with perturbed camera poses. Within a reasonable perturbation range (± 20 cm), the performance is stable, denoting that the SinGRAV is robust to the reasonable perturbations for camera pose distribution during the data acquisition.

the multi-view images acquisition. We conduct the study on *Stonehenge*, of which the results are demonstrated in Figure 12. Figure 12 shows that the scores for SIFID-MV and Diversity-MV stay similar when the scale of the Gaussian noise changes within a range of $[-20, +20]$ cm. Experimental results show that SinGRAV is robust to adequate perturbations of the camera poses, which means that SinGRAV just needs the camera poses to come from a roughly satisfying distribution rather than strictly demands that the camera poses come from exactly controlled settings.

A.3 VIDEO DEMONSTRATIONS

Here we give outlines for video results we provide in the supplementary videos. In Supplementary Video, we organize the results in the following manner:

- Qualitative results on various datasets from SinGRAV.
- Qualitative results from SinGRAV-derived variants.
- Qualitative results from baselines including GRAF Schwarz et al. (2020), pi-GAN Chan et al. (2021a) and GIRAFFE Niemeyer & Geiger (2021).
- Qualitative results for applications.